

Electronic Thesis and Dissertation Repository

---

10-25-2021 10:00 AM

## Evolutionary Design of Search and Triage Interfaces for Large Document Sets

Jonathan A. Demelo, *The University of Western Ontario*

Supervisor: Sedig, Kamran, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Computer Science

© Jonathan A. Demelo 2021

Follow this and additional works at: <https://ir.lib.uwo.ca/etd>



Part of the [Graphics and Human Computer Interfaces Commons](#)

---

### Recommended Citation

Demelo, Jonathan A., "Evolutionary Design of Search and Triage Interfaces for Large Document Sets" (2021). *Electronic Thesis and Dissertation Repository*. 8214.  
<https://ir.lib.uwo.ca/etd/8214>

This Dissertation/Thesis is brought to you for free and open access by Scholarship@Western. It has been accepted for inclusion in Electronic Thesis and Dissertation Repository by an authorized administrator of Scholarship@Western. For more information, please contact [wlsadmin@uwo.ca](mailto:wlsadmin@uwo.ca).

## Abstract

This dissertation is concerned with the design of visual interfaces for searching and triaging large document sets. Data proliferation has generated new and challenging information-based tasks across various domains. Yet, as the document sets of these tasks grow, it has become increasingly difficult for users to remain active participants in the information-seeking process, such as when searching and triaging large document sets. During information search, users seek to understand their document set, align domain knowledge, formulate effective queries, and use those queries to develop document set mappings which help generate encounters with valued documents. During information triage, users encounter the documents mapped by information search to judge relevance to information-seeking objectives. Yet, information search and triage can be challenging for users. Studies have found that when using traditional design strategies in tool interfaces for search and triage, users routinely struggle to understand the domain being searched, apply their expertise, communicate their objectives during query building, and assess the relevance of search results during information triage. Users must understand and apply domain-specific vocabulary when communicating information-seeking objectives. Yet, task vocabularies typically do not align with those of users, especially in tasks of complex domains. Ontologies can be valuable mediating resources for bridging between the vocabularies of users and tasks. They are created by domain experts to provide a standardized mapping of knowledge that can be leveraged both by computational- as well as human-facing systems. We believe that the activation of ontologies within user-facing interfaces has a potential to help users when searching and triaging large document sets, however more research is required.

This dissertation is structured in the form of an integrated article, encompassing a chaptered set of research over five materials either published or prepared for publishing, along with support chapters. The first of these materials is concerned with exploring the design of visual interfaces for information search and triage, and the integration of ontology files during query building. The second of these materials is concerned with examining how cognitive map formation helps form knowledge of complex ontological space, and how visual interfaces can support users encounter, understand, and explore their ontologies. The third of these materials is concerned with framing the high-level requirements for generalized search interfaces, which is then use in the generation of a prototype interface. The fourth of these materials is concerned with understanding the multi-staged information seeking process, distilling its design requirements for search and triage interfaces. Additionally, the design of a novel visual interface integrating progressive disclosure and ontology mediation is described. The fifth and final of these materials is concerned with describing a three-staged evolutionary design of a VAT interface for searching and triaging large document sets. Specifically, this material describes the formulization, realization, and validation of three interface stages generated with guidance from in-depth topic analysis, design criteria, and formative assessment. Additional support materials within the dissertation include introduction, background, as well as a chapter presenting summary, contributions, and future research directions.

## **Lay Summary**

From the ancient practice of food foraging to the information-seeking objectives of the modern world, search and triage has always been a cornerstone of the human experience. The slow evolutionary march has held humans largely to the same form, yet through tools, humanity has been able to break free of physiological binds, both physical and cognitive. Examples such as shovels, cars, and pencils reflect the power of physical tool augmentation, just as language, mathematics, and computers reflect the power of cognitive tool augmentation. User wellbeing is tied to the quality of their tools. When mindfully designed, tools can help users complete their tasks – perhaps at a quicker pace, more effectively, for a longer time, or with less effort. Thus, if a task can be better understood, then designers can be cognizant of constraining features, formalize requirements which work around those constraints, then use those requirements to assess the quality of existing tools and in turn form new strategies that better serve users.

The 1950s initiated the use of computers as search tools, proliferating over the decades. Notably, the rise of reasonably accessible global networking in the 80s and 90s culminated in the first exposure to the power of generalized search engines, and in particular, for exploring the interconnected web. Yet even as tools continue to improve – with faster computations, more intelligence, less power, and over larger document sets – users must still be the primary consideration when designing effective search and triage solutions. That is, technological innovation does not matter if the user does not know how to use their tools, understand its results, or activate that understanding in a manner that is meaningful to completion of their task.

It has never been more important for researchers to provide guidance for designers to help realize more powerful and effective search and triage tools. Designers must concentrate first on addressing the user-facing considerations of search and triage interface design. For this effort, we center this dissertation on investigating how the designs of visual interfaces can help users improve their searching and triaging of large document sets.

## **Keywords**

Information Search, Information Triage, Ontologies, Human-Information Interaction, Human-Computer Interaction, Visual Representation Design, Interaction Design, Visual Analytics Tools, Large Document Sets; Machine Learning

## **Co-Authorship Statement**

**Chapter 1** and **Chapter 2** are original materials which introduce and provide background for the dissertation. **Chapter 3** was a collaborative effort between supervisor Dr. Kamran Sedig, Dr. Paul Parsons, post-doctorate, and myself. This research was crafted to explore future directions of investigation. This research involved efforts in the design of a prototype visual interface using existing software materials created by Dr. Paul Parsons. Figures were conceptualized and generated in a collaborative manner. **Chapter 4, Chapter 5, Chapter 6, and Chapter 7** were a collaborative effort with my supervisor, Dr. Kamran Sedig. These research materials were conceptualized and directed in a collaborative manner, framing, and generating the materials for research objectives. Portions of software material for Chapter 4 was generated in collaboration with Insight Lab member Joe Ling, M.Sc. **Chapter 8** is original research which provides conclusions on research objectives within the dissertation.

## **Acknowledgements**

I would first like to express my deepest appreciation to my supervisor, Dr. Kamran Sedig. Together, we have faced an incredible number of challenges, many of which we could never have anticipated. Yet with your leadership and teachings, you have supported me to become more vigilant, mindful, and resolute. For this and much more, I am eternally grateful.

Secondly, I would like to express gratitude to my examiners: Dr. Louis Capretz, Dr. Hai-Ning Liang, Dr. Dan Lizotte, and Dr. Mostafa Milani. Completing a doctorate program is never a simple endeavor. Yet it would be hard to find a time more challenging than our present circumstances. Thank you all for your time and efforts.

Finally, I thank my family for their support during my academic journey.

“Once men turned their thinking over to machines in the hope that this would set them free. But that only permitted other men with machines to enslave them.”

Frank Herbert, 'Dune'.

# Table of Contents

|  |     |
|--|-----|
| Abstract.....  | i   |
| Lay Summary .....  | ii  |
| Keywords .....   | ii  |
| Co-Authorship Statement .....  | iii |
| Acknowledgements .....   | iv  |
| Table of Contents.....   | v   |
| List of Tables .....   | ix  |
| List of Figures.....   | xi  |
| Chapter 1    Introduction.....   | 1   |
| 1.1 Motivation .....   | 1   |
| 1.2 Structure of the Dissertation.....   | 2   |
| 1.3 References .....   | 3   |
| Chapter 2    Background and Terminology .....  | 4   |
| 2.1 Conceptual and Theoretical Background .....  | 4   |
| 2.2 Terminology .....  | 4   |
| 2.3 Acronyms .....   | 5   |
| Chapter 3    OVERT-MED: An Ontology-Driven Visual Search and Triage Interface for MEDLINE..... | 6   |
| 3.1 Introduction .....   | 6   |
| 3.1.1 Ontologies.....  | 8   |
| 3.1.2 Document Triage .....  | 9   |
| 3.1.3 Search Result Visualization .....  | 9   |
| 3.1.4 Related Work .....   | 9   |
| 3.2 Methods.....   | 10  |
| 3.2.1 Indexing of MEDLINE and HPO .....  | 10  |
| 3.2.2 Development and Architecture .....   | 11  |
| 3.3 Results .....  | 11  |
| 3.3.1 Ontology Term Suggestion .....   | 11  |
| 3.3.2 Sensitivity Encoding for Query Refinement.....   | 12  |
| 3.3.3 Interactive Triageing Support to Mitigate Information Overload .....                     | 14  |
| 3.4 Discussion .....   | 19  |
| 3.4.1 Validation.....  | 20  |
| 3.4.2 Limitations .....  | 21  |
| 3.4.3 Future Work.....   | 21  |
| 3.4.4 Conclusions.....   | 21  |
| 3.5 References .....   | 22  |

|           |  |     |
|-----------|--|-----|
| Chapter 4 | Forming Cognitive Maps of Ontologies Using Interactive Visualizations .....            | 26  |
| 4.1       | Introduction .....   | 26  |
| 4.2       | Background .....   | 28  |
| 4.2.1     | Cognitive Map Formation .....  | 28  |
| 4.2.2     | Ontologies .....   | 29  |
| 4.2.3     | Interactive Visualization Tools .....  | 30  |
| 4.3       | Methods .....  | 31  |
| 4.3.1     | Related Work .....   | 31  |
| 4.3.2     | Task Analysis .....  | 34  |
| 4.3.3     | Existing Tool Review .....   | 34  |
| 4.4       | Materials .....  | 44  |
| 4.4.1     | PRONTOVISE Technologies .....  | 44  |
| 4.4.2     | PRONTOVISE Workflow and Design .....   | 44  |
| 4.5       | Usage Scenario .....   | 57  |
| 4.6       | Discussion and Conclusions .....   | 68  |
| 4.5       | References .....   | 70  |
| Chapter 5 | Design of Generalized Search Interfaces for Health Informatics .....                   | 75  |
| 5.1       | Introduction .....   | 75  |
| 5.2       | Background .....   | 76  |
| 5.2.1     | Health Informatics .....   | 76  |
| 5.2.2     | Machine Learning .....   | 77  |
| 5.2.3     | Ontologies .....   | 78  |
| 5.3       | Methods .....  | 79  |
| 5.3.1     | Task Review .....  | 79  |
| 5.3.2     | High-Level Criteria .....  | 83  |
| 5.3.3     | Analysis of Traditional Interface Strategies for Health Informatics Search Tasks ..... | 83  |
| 5.4       | Results .....  | 89  |
| 5.4.1     | Design Scope .....   | 89  |
| 5.4.2     | Technical Scope .....  | 89  |
| 5.4.3     | Functional Workflow of ONTSI .....   | 90  |
| 5.4.4     | Front-End Subviews .....   | 91  |
| 5.4.5     | Back-End Systems .....   | 94  |
| 5.5       | Usage Scenario .....   | 97  |
| 5.5.1     | User Profile .....   | 97  |
| 5.5.2     | Ontology File and Document Set .....   | 97  |
| 5.5.3     | Usage Scenario .....   | 98  |
| 5.6       | Discussion and Conclusions .....   | 103 |

|   |     |
|---|-----|
| 5.6.1 Evaluation of ONTSI .....   | 103 |
| 5.6.2 Limitations .....   | 105 |
| 5.6.3 Conclusions.....  | 105 |
| 5.7 References .....  | 106 |
| Chapter 6 Searching and Triage of Large Document Sets: An Ontology-Supported Visual Analytics Approach        | 112 |
| 6.1 Introduction .....  | 112 |
| 6.2 Background .....  | 114 |
| 6.2.1 Information Search and Triage.....  | 114 |
| 6.2.2 Machine Learning .....  | 115 |
| 6.2.3 Ontologies.....   | 116 |
| 6.3 Methods.....  | 117 |
| 6.3.1 Task Analysis.....  | 117 |
| 6.3.2. Design Criteria.....   | 124 |
| 6.4 Materials.....  | 125 |
| 6.4.1 Technical Scope .....   | 126 |
| 6.4.2 VisualQUEST Functional Workflow .....   | 126 |
| 6.4.3 Back-End Systems .....  | 127 |
| 6.4.4 Front-End Subviews.....   | 129 |
| 6.5 Discussion and Summary .....  | 136 |
| 6.5.1 Formative Evaluations of VisualQUEST .....  | 136 |
| 6.5.2 Limitations .....   | 139 |
| 6.5.3 Future Work.....  | 140 |
| 6.5.4 Summary.....  | 140 |
| 6.6 References .....  | 141 |
| Chapter 7 Evolutionary Design of Search and Triage Interfaces for Large Document Sets: A Formative Assessment | 147 |
| 7.1 Introduction.....   | 147 |
| 7.2 Background .....  | 148 |
| 7.2.1 Evolutionary Design .....   | 148 |
| 7.2.2 Information Search and Triage in ML-supported VAT interfaces .....                                      | 149 |
| 7.3 Materials and Methods .....   | 151 |
| 7.3.1 Formative Assessment .....  | 151 |
| 7.3.2 Shared Back-End Components .....  | 153 |
| 7.4 Results .....   | 155 |
| 7.4.1 Stages of the Evolutionary Design.....  | 155 |
| 7.4.2 Stage 1.....  | 155 |
| 7.4.3 Stage 2.....  | 161 |

|   |     |
|---|-----|
| 7.4.3 Stage 3.....  | 171 |
| 7.5 Discussion and Conclusions .....  | 178 |
| 7.5.1 Overall .....   | 178 |
| 7.5.2 Progressive Disclosure and the Stages of Information Search and Triage .....        | 179 |
| 7.5.3 Ontology Integration .....  | 180 |
| 7.5.4 Limitations .....   | 181 |
| 7.5.5 Conclusions.....  | 182 |
| 7.6 References .....  | 182 |
| Chapter 8 Summary, Contributions, and Future Research .....                               | 187 |
| 8.1 Chapter Summaries .....   | 187 |
| 8.2 General Contributions .....   | 188 |
| 8.2.1 Research Contributions .....  | 188 |
| 8.2.2 Practical Contributions.....  | 189 |
| 8.3 Future Research.....  | 189 |
| 8.3.1 Formal User Studies and Expanded Investigations.....                                | 190 |
| 8.3.2 Prescriptive Frameworks.....  | 190 |
| 8.3.3 Novel Modes of Representation and Interaction for Search and Triage Interfaces..... | 191 |
| Appendices .....  | 192 |
| Appendix A Formal Study Letter of Information and Consent .....                           | 192 |
| Appendix B Task Questions.....  | 197 |
| Appendix C Post-Task Questions.....   | 202 |
| Appendix D List of Documents for Task 7 .....   | 205 |
| References.....   | 205 |
| Appendix E Interview Questions .....  | 207 |
| Curriculum Vitae .....  | 208 |
| Education .....   | 208 |
| Doctor of Philosophy, Computer Science (Candidate) .....                                  | 208 |
| Bachelor of Science, Honours, Computer Science and Visual Arts: 2013 .....                | 208 |
| Ontario Secondary School Diploma: 2009 .....  | 208 |
| Publications .....  | 208 |
| Talks and Posters.....  | 208 |
| Teaching.....   | 209 |
| Course Design Support & Management .....  | 209 |
| TA Positions .....  | 209 |
| Awards and Grants .....   | 209 |
| Professional and Organizational Activities .....  | 209 |

## List of Tables

|  |     |
|--|-----|
| Table 4-1 Summary of the cognitive activities performed within spaces. Included is the name and description, the underlying processes of creative thinking which relate to the cognitive activity, and the types of spatial knowledge which must be developed within a cognitive map of a space before the activity can be performed within that space. ....   | 33  |
| Table 4-2 The high-level criteria for designing interactive visualization tools of ontology datasets which promote the stages of cognitive map formation for learning tasks. ....  | 34  |
| Table 4-3 The types of subviews within an ontology dataset visualization interface. ....   | 35  |
| Table 4-4 A summary of the high-level criteria within PRONTOVISE, which can be used for designing interactive visualization tools which support complex learning and the stages of cognitive map formation. The satisfaction of these criteria at the implementation level are discussed in detail later within the workflow. ....   | 46  |
| Table 4-5 A summary of the subviews of PRONTOVISE, describing their subview type, supported cognitive activities, and their relationship to the stages of cognitive map formation. ....  | 48  |
| Table 5-1 The criteria for guiding the design of interfaces for health informatics search tasks involving large document sets. For abbreviation purposes, design criteria will be referenced in the text as DC#, where # is its assigned number. ....  | 83  |
| Table 5-2 A summary matrix of alignment between the criteria and interface strategies. Full descriptions are found within their respective sections. “Strong” is assigned if a characteristic of the interface strategy promotes alignment with the requirements of the design criterion. “Weak” is assigned if a characteristic of the strategy does not promote alignment with the requirements of the design criterion. “Variable” is assigned if the interface strategy has the potential to align with the criterion; however, such an alignment is not innate and must be actively pursued. .... | 83  |
| Table 5-3 The role of each criterion within the design of ONTSI. The incorporation of these criteria in ONTSI’s implementation is discussed within the workflow and usage scenario. ....   | 89  |
| Table 6-1 The stages of the information-seeking and triage process: their associated task, alignment with existing models, and functional descriptions. ....   | 118 |
| Table 6-2 The set of design criteria for creating VAT interfaces for searching and triaging large document sets. For reference and clarity, a numerical value is assigned to each design criteria. Each DC# describes the design criteria, provides an integration classification, and suggests value in ontology integration for that criteria. ....  | 124 |
| Table 6-3 General description of tasks performed during formative evaluations. ....  | 137 |
| Table 7-1 General description of tasks. ....   | 153 |
| Table 7-3 The criteria for guiding the design of generalized search interfaces for tasks involving large document sets. ....   | 156 |
| Table 7-4 The role of each criterion within the design of Stage 1. ....  | 156 |
| Table 7-5 How users completed the tasks using Stage 1. ....  | 159 |
| Table 7-6 Design criteria for creating VAT interfaces for searching and triaging large document sets. Each DC# describes the design criteria and provides an integration classification. ....  | 162 |
| Table 7-7 The role of each criterion within the design of Stage 2. ....  | 163 |

|   |     |
|---|-----|
| Table 7-8 How users completed the tasks using Stage 2 .....   | 169 |
| Table 7-9 The configuration of Stage 3, and its alignment to unsatisfied criterion of prior formulization efforts. .... | 172 |
| Table 7-10 How users completed the tasks using Stage 3. ....  | 175 |

## List of Figures

|  |    |
|--|----|
| Figure 3-1 A screenshot of PubMed showing results from searching for “liver”. In exploratory contexts, when users have open-ended information needs and imprecise vocabulary, PubMed’s interface style is not very helpful. Users have difficulty articulating information needs and must deal with long lists of results that span many pages. ....   | 7  |
| Figure 3-2 Client-server architecture of the OVERT-MED system. The webserver fetches MEDLINE updates from the NLM ftp site and indexes them in a Lucene index. HPO is also indexed. The API handles requests and connects to the Solr search server. The client displays visualizations and handles user interactions. ....  | 11 |
| Figure 3-3 The Ontology Term Suggester, showing results from typing “liver”. Many of the resulting phenotypes do not have “liver” in their names, but are in some way associated with liver, and are displayed to help users improve vocabulary to better articulate information needs. ....   | 12 |
| Figure 3-4 A set of Sensitivity Encoded Query Selectors for ‘progressive external ophthalmoplegia’. The number of matching articles for each combination is provided numerically and encoded visually using the length of the bar next to each combination. ....   | 13 |
| Figure 3-5 The result of adding a second phenotype via the Ontology Term Suggester, which leads to more Sensitivity Encoded Query Selectors. ....  | 14 |
| Figure 3-6 The Query Result Heatmap: Six documents are represented by six rows, where each column represents a term (progressive external ophthalmoplegia). ....   | 14 |
| Figure 3-7 State of the interface after a user has selected “progressive+ophthalmoplegia”. ....  | 15 |
| Figure 3-8 The Term Distribution Matrix for two different documents within the same result set. ....   | 16 |
| Figure 3-9 The result of a user adding documents containing “congenital” and “fibrosis” for comparison. It is immediately clear through quick visual scanning that not many documents contain both “congenital fibrosis” and “ophthalmoplegia”. ....   | 16 |
| Figure 3-10 Close-up view of the Query Result Heatmap. When users pause on a document for five or more seconds a small mark is placed beside the document as a reminder. ....  | 17 |
| Figure 3-11 Twenty-nine documents have been selected to examine in closer detail. They are encoded in the Selected Pile Heatmap, and the Term-Encoded Abstract is displayed for the first document. It is easy to see that ‘renin’ and ‘hypertension’ both appear frequently, indicating that they are important. ....   | 18 |
| Figure 3-12 Full screen capture showing all components of OVERT-MED where a user has traversed all stages of a search and triaging activity. ....  | 18 |
| Figure 3-13 The interface after a user has examined a document in detail in the final stage, discovered a link, and has returned to the initial stage with a new information need. The user discovers a phenotype named “elevated mean arterial pressure”, selects it, then adds documents satisfying “arterial+pressure” to the Query Result Heatmap. ....  | 19 |
| Figure 3-14 Overall search and triage process supported by OVERT-MED. Users take seven main steps—the first two within a vocabulary building phase, and the next five within a triaging phase. The triaging phase can be broken down into the three key stages. Each step is supported by a technique in the interface. At any point users can return to any step, which supports the iterative and cyclical nature of search and triaging tasks. .... | 20 |

|  |    |
|--|----|
| Figure 4-1 Protégé Entity Browser: List and details subviews loaded with the Human Phenotype Ontology. Source: Image with permission courtesy of Center for Biomedical Informatics Research, Stanford University School of Medicine, Protégé Team, <a href="https://protege.stanford.edu">https://protege.stanford.edu</a> .....   | 37 |
| Figure 4-2 Protégé OntoGraf: List and context subviews loaded with the Human Phenotype Ontology. Source: Image with permission courtesy of Center for Biomedical Informatics Research, Stanford University School of Medicine, Protégé Team, <a href="https://protege.stanford.edu">https://protege.stanford.edu</a> .....   | 38 |
| Figure 4-3 WebVOWL: Context and details subviews loaded with supplied Friend of a Friend Ontology. Source: Image generated from WebVOWL: Web-based Visualization of Ontologies resource <a href="http://vowl.visualdataweb.org/webvowl.html">http://vowl.visualdataweb.org/webvowl.html</a> .....  | 39 |
| Figure 4-4 Ontodia (now contained within the Metaphactory software suite): List, context, and details subviews. Source: Image with permission courtesy of Metaphacts, <a href="http://metaphacts.com">metaphacts.com</a> .....   | 41 |
| Figure 4-5 WebProtégé Entity Graph: List and context subviews. Source: Image with permission courtesy of Center for Biomedical Informatics Research, Stanford University School of Medicine, Protégé Team, <a href="https://protege.stanford.edu">https://protege.stanford.edu</a> .....   | 42 |
| Figure 4-6 OntoViewer: List, overview, and relations context subviews. Source: Image with permission courtesy of “Visualization and analysis of schema and instances of ontologies for improving user tasks and knowledge discovery”, School of Informatics, UniRitter Laureate International Universities. ....   | 44 |
| Figure 4-7 Depiction of the workflow of PRONTOVISE (PRogressive ONTOlogy VISualization Explorer). Yellow boxes represent the processes performed within the back-end computation system. Blue boxes represent the object types which are persisted within browser storage. Orange boxes represent the various subviews within the front-end visualization system. The green box represents the types of low-level interactions which can be made to the system. ....   | 45 |
| Figure 4-8 An overall view of the PRONTOVISE ontology visualization system, which has seven subviews: Search and Pinning Panel subview (a); Ontology Sections Panel subview (b); Section Levels Panel subview (c); Level Landmark Entities Panel subview (d); Entity Network Panel subview (e); Path Explorer Panel subview (f); and Entity Details Panel subview (g). ....  | 47 |
| Figure 4-9 A depiction of Ontology Entity Search, showing a search activity which has resulted in three ‘Pin’ interactions. ....   | 49 |
| Figure 4-10 A depiction of Ontology Entity Pinning, where pinned entities are represented in the same fashion as they were found in Ontology Entity Search, with a name label, an Unpin button, and a unique color. We have included a button located at the topmost position of Ontology Entity Pinning labeled ‘Remove All Pinned Entities’. When clicked, this button removes all pinned landmarks from the system. When an ontology entity is removed, its annotated representations will be removed from all subviews. .... | 49 |
| Figure 4-11 A depiction of the distortion technique within the Ontology Sections Panel. This technique can be adjusted through interaction. This is achieved by holding the Shift key while directing the mouse over a section. Releasing the Shift key will end the interaction event and lock in the sizing adjustments. If adjustments have been  |    |

made, yet the user would like to return to the original distortion scaling, we have provided a Reset button at the top left corner of the Ontology Sections Panel. ....50

Figure 4-12 A depiction of the magic lens within the Ontology Sections Panel. We have designed an interaction to occur when users drag the bottom portion of the magic lens horizontally across the ontology sections, which will both refresh the ontology section’s distortion technique and expose information for that section. ....51

Figure 4-13 The initial list of levels within the Ontology Section Levels Panel. ....52

Figure 4-14 A level within the Ontology Section Levels Panel. The level has three main interactions. First, when we move our cursor over a circle, a label is generated which displays the name of the entity and annotates the location of the ontology entity within the Level Landmark Entities Panel. Second, if we click on a circle, we perform an interaction which selects that entity as the initial position within the Entity Network Panel, the Path Explorer Panel, and the Entity Details Panel. Additionally, when a level has many ontology entities, the available visual space may become very crowded. To address this, we designed an interaction which allows us to distort the space by holding Shift and activating our mouse scroll wheel, which expands and contracts the horizontal scale of the plot graphs. We then can drag the plot graphs left or right using a single mouse click and drag action, allowing us the ability to closely inspect the full set of ontology entities and relations within the level. ....52

Figure 4-15 The magic lens within the Ontology Section Levels Panel. We have designed the magic lens with an interaction which generates lines to represent the set of ontology relations within the level as we drag horizontally across the visual space. ....53

Figure 4-16 The Level Landmark Entities Panel depicting an ontology level with significant connectivity within its matrix representation. We can see, using the magic lens of the Section Levels Panel, that the level has countless numbers of ontology entities and relations, which have been analyzed and presented in a usable manner with the Level Landmark Entities Panel. Within the matrix representation, each ontology entity is represented as a circle accompanied by a text label. Each circle maintains a red to white fill encoding reflecting its importance calculation, red being the most inherited, and white being the least. When we move our cursor over a circle, the text label grows and boldens for rapid association between the matrix position and its row and column labels. A color spectrum and text label is provided at the intersection points of the matrix representing ontological distance. This distance calculation is determined by the number of inheritances performed when defining the ontologies up to their nearest common parentage. For example, when a matrix position reflects the intersection between two ontology entities which inherit from the same immediate parent, their distance will be calculated and displayed as 2. ....54

Figure 4-17 The Entity Network Panel depicting a low-level graph-like abstraction of specific ontology entities within the ontology network: parents (super classes), children (sub classes), and shared inheritances. Ontology entities are represented by a blue filled circle and text label. In these regions, relations are depicted with lines which link ontology entities to reflect relationship and localized structure. When two ontology entities are selected, their shared network is depicted with an additional representation maintaining the relations between each, their shared inherited parent, and distance. Additional information is exposed by moving the mouse over. ....55

|  |    |
|--|----|
| Figure 4-18 The Path Explorer Panel depicting all the levels within the selected path of the ontology, where each level maintains a set of rectangles representing all sibling entities and a text label reflecting its depth within the levels of the ontology.....   | 56 |
| Figure 4-19 An overall view of the PRONTOVISE ontology visualization system in its initial state after the Human Phenotype Ontology has been uploaded. ....  | 58 |
| Figure 4-20 The initial stage of the Ontology Sections Panel subview. ....   | 58 |
| Figure 4-21 Adjusting the scaling of the Ontology Sections Panel subview to enlarge the “Abnormality of the Musculature” section, which is normally represented as a much smaller portion of the visual space.....   | 59 |
| Figure 4-22 The result of a magic lens levels’ activation on the “Abnormality of Limbs” ontology section.....  | 59 |
| Figure 4-23 The result of a magic lens on the “Abnormality of Limbs” ontology section.....   | 60 |
| Figure 4-24 The selection of the “Abnormality of the Skeletal System” section. ....  | 60 |
| Figure 4-25 An overview of the Section Levels Panel subview depicting the “Skeletal System” ontology section levels after navigating from the Ontology Sections Panel subview.....   | 61 |
| Figure 4-26 Inspecting the 13th level of the “Skeletal System” ontology section.....   | 61 |
| Figure 4-27 Using the magic lens to inspect a set of ontology entities and their inheritance relations. ....   | 62 |
| Figure 4-28 An overview of the Level Landmark Entity Subview, representing the 3rd level of the “Abnormality of Skeletal System” ontology section. ....  | 63 |
| Figure 4-29 Route between of “Abnormal Joint Morphology” and “Abnormal Epiphyseal Stippling”.....  | 63 |
| Figure 4-30 The initial state of the Entity Network Panel subview when an ontology entity is chosen as the initial position. In this case, “Abnormal Appendicular Skeleton Morphology” has been selected. ....   | 64 |
| Figure 4-31 The initial state of the Entity Network Panel subview when two ontology entities are chosen as the initial positions. In this case, “Abnormal Appendicular Skeleton Morphology” and “Abnormal Joint Morphology” have been selected.....  | 65 |
| Figure 4-32 The Path Explorer Panel subview after selecting “Abnormal Appendicular Skeleton Morphology” as its current position.....   | 66 |
| Figure 4-33 The Ontology Landmark Search results from typing “skeleton”. ....  | 67 |
| Figure 4-34 “Broad Forearm Bones” has been pinned, designating it as an important ontology entity, which is to be highlighted with its assigned color whenever it appears in a PRONTOVISE subview. ....  | 67 |
| Figure 4-35 When “Broad Forearm Bones” has been selected within any subview, the Entity Details Panel subview depicts all information for that ontology entity as provided by HPO. ....  | 68 |
| Figure 5-1 PubMed Advanced Search Builder: an example of a structured interface strategy for a search task. In this use case, a query item was generated for a MeSH term for heart abnormalities, a completion date after August 8, 2015, and in the English language, with the publisher of Oxford University Press soon to be added. Source: Image generated on January 18, 2021, using the public web portal provided by the National Center for Biotechnical Information, <a href="https://pubmed.ncbi.nlm.nih.gov/advanced/">https://pubmed.ncbi.nlm.nih.gov/advanced/</a> (accessed on January 18, 2021). .... | 85 |
| Figure 5-2 WebMD Search Interface: an example of an unstructured interface strategy for a search task. In this use case, the free-text query “heart condition” was generated. Source: Image generated on January 18, 2021, using the   |    |

public web portal provided by WebMD, [https://www.webmd.com/search/search\\_results/default.aspx?query=heart%20condition](https://www.webmd.com/search/search_results/default.aspx?query=heart%20condition) (accessed on January 18, 2021). ..... 86

Figure 5-3 WebMD’s Symptom Checker: an example of a structured–like query expansion interface strategy for a search task. In this use case, users are guided along a series of query building opportunities, allowing them to enter various symptoms and personal health criteria while aligning their personal vocabulary with the information resource vocabulary. Source: Image generated on January 18, 2021, using the public web portal provided by WebMD, <https://symptoms.webmd.com/> (accessed on January 18, 2021). ..... 88

Figure 5-4 Depiction of the functional workflow of ONTSI. Labelled arrows reflect the steps of the interaction loop. Brown boxes represent the processes performed within the back-end computation systems. Blue boxes represent the object types that persist within the browser and external index database storage. Pink boxes represent the back-end computational systems of ONTSI. Yellow boxes represent the various subviews within the front end of ONTSI. The green box represents the types of interactions that can be conducted with the system. .... 91

Figure 5-5 The overall view of ONTSI in full use and outlined coverage of its five subviews: Upload subview, partial (a); Search subview (b); Result List subview (c); Result Item subview (d); and Saved List subview, partial (e). ..... 92

Figure 5-6 Pseudocode of clustering spanning the workflow of ONTSI (front end), ONTSI server, and Solr server. 96

Figure 5-7 The initial state of ONTSI. .... 98

Figure 5-8 Selecting a document set and ontology file for upload. .... 98

Figure 5-9 The results of a search task using the query, “How does chromosomal instability drive tumor progression?” ..... 99

Figure 5-10 ONTSI after opening the documents “Cancer morphology, carcinogenesis and genetic instability: a background” and “Kaposi’s sarcoma–associated herpesvirus–encoded latency-associated nuclear antigen induces chromosomal instability through inhibition of p53 function.” ..... 100

Figure 5-11 ONTSI while the user is presented with mediating opportunities from the expert–defined Human Phenotype Ontology. .... 101

Figure 5-12 ONTSI after running a new search after the user took advantage of mediation opportunities presented in the generation of the query items. .... 102

Figure 5-13 ONTSI after the user has selected new document entries for deeper inspection. .... 103

Figure 6-1 WebMD’s Symptom Checker: An example of a progressive disclosure implementation of a task. In this case, users are guided along a series of query building opportunities, input-ting symptoms and personal health criteria. Source: Image generated on 2021/01/18 using the public web portal provided by WebMD, <https://symptoms.webmd.com/> (accessed on January 18, 2021) ..... 119

Figure 6-2 OVERT-MED, a sensitivity encoded ontology-driven search interface. Depicted is the result of users comparing the queries of “Progressive + ophthalmoplegia” and “congenital + fibrosis”. Source: Image generated on 2021/07/05 with permission courtesy of Insight Lab, Western University, London Ontario Canada <http://insight.uwo.ca/> (accessed on July 5, 2021) ..... 121

Figure 6-3 PubMed: Searching the MEDLINE document set for “heart” provides a linear triaging interface of over 1.5 million results spread over 100,000 pages within a ten per page system. Source: Image generated on 2021/07/05

using the public web portal provided by the National Center for Bio-technical Information, <https://pubmed.ncbi.nlm.nih.gov/?term=heart> (accessed on July 5, 2021) ..... 122

Figure 6-4 ChartingHockey Player Production Rates: ChartingHockey supports non-linear triaging of statistical information in player evaluation tasks. Source: Image generated on 2021/07/05 using the public web portal provided by ChartingHockey, <https://www.chartinghockey.ca/daily-skater-charts/> (accessed on July 5, 2021) ..... 123

Figure 6-5 Depiction of the dynamic functional workflow of VisualQUEST. Labelled arrows reflect transitional actions of systems and users. Users begin by uploading their ontology files and document sets (bottom center). These activate their respective Ontology and Document Process Systems (brown boxes), which prepare their content for back-end storage (blue boxes) and activation (pink boxes). Once achieved, the front-end interface activates its various subview functionalities (yellow boxes). The user can act upon the interface (green box). The system processes those actions, formulates adjustments to its display, and returns a visual response to be perceived by users. Some-times, these adjustments must connect with back-end systems, ask computations to be performed, and the results of those computations sent back to the display level..... 127

Figure 6-6 Pseudocode of clustering functionality spanning the workflow of VisualQUEST (front-end), Analytics Server, and Document Server. .... 128

Figure 6-7 An overview of Query Building subview within VisualQUEST. Here, an ontology and document set have been uploaded, from which a set of query items have been generated, with a subset of those added in Search. .... 129

Figure 6-8 Expanding a query item to assess related ontology elements..... 130

Figure 6-9 An overview of the Search subview within VisualQUEST. Prior query building has produced a set of query items, which have been inserted into the search formulation for preview. .... 131

Figure 6-10 A preview of a search formulation which uses all query items..... 131

Figure 6-11 An overview of the High-Level Triage subview within VisualQUEST. Previous searching has produced groupings of the document set. A subset of those groupings has been selected, producing further listing of their contained documents. Some of these documents have been inspected and added into Low-Level Triage. .... 132

Figure 6-12 High-Level Triage after groupings have been inspected and opened for further examination. .... 133

Figure 6-13 A close look at High-Level Triage, showing a listing of documents contained within a selected grouping of the document set..... 133

Figure 6-14 An overview of the Low-Level Triage subview within VisualQUEST. A set of saved documents have been produced in prior triaging. Each of these documents have been provided a timeline-like summary reflecting words or phrases aligning with the search formulation within its content. When selected, the Document Content viewer depicts the document itself, either in the summarized or full document mode. The summarized mode is shown. .... 134

Figure 6-15 A closer view of the Selected Documents listing, where the user can see timeline-like abstraction of documents and their alignment with the search formulation. .... 135

Figure 6-16 An overview of Low-Level Triage after documents have been added and opened for viewing. .... 136

Figure 7-1 Pseudocode of clustering functionality for Stage 1, Stage 2, and Stage 3, Analytics Server, and Document Server. Step 5 is not performed for Stage 1 and Stage 2, which do not use ontologies during query expansion. .... 154

|   |     |
|---|-----|
| Figure 7-2 The overall view of Stage 1 components: Upload, partial (a); Query Building (b); Result List (c); Result Item (d); Saved List, partial (e).....  | 158 |
| Figure 7-3 An overview of Query Building within Stage 2. Users can see that a document set has been uploaded, that a set of query items have been generated for inspection, one having been added in Search. ....   | 165 |
| Figure 7-4 An overview of Search within Stage 2. Users can see that Query Building has generated a set of query items, which have been injected into the preview search environment. ....   | 166 |
| Figure 7-5 An overview of High-Level Triage within Stage 2. Users can see that Search has generated a set of document groupings reflecting the document set. A subset of those groupings has been selected for deeper investigation, producing further listings of documents from the set. These documents have been inspected and have been determined to be valuable enough to be added into Low-Level Triage. .... | 167 |
| Figure 7-6 An overview of Low-Level Triage within Stage 2. Users can see High-Level Triage has generated a subset of the document set for deeper inspection. Each of these documents have been provided a timeline-like summary reflecting the presence of query items within its content. If selected, a document-level viewer depicts all available document content, either in summarized or full form.....        | 168 |
| Figure 7-7 An overview of Query Building within Stage 3. Users can see that an ontology and document set has been uploaded, that a set of query items have been generated for inspection, and a subset of those have been added in Search. ....   | 174 |

# Chapter 1 Introduction

This dissertation is presented in the form of an integrated article, reflecting a series of individual materials which together synthesize into a singular research direction. These materials establish research objectives encompassing a variety of diverse and interdisciplinary domains, such as human-computer interaction, human-information interaction, visual analytics, cognitive sciences, alongside computational considerations such data management, computational linguistics, and machine learning. Each material is presented in its original, self-contained form. That is, each maintains its original set of sections (e.g., introductory, background, materials, results, and conclusion) presented in a fashion that can support both individual and progressive reading. We now provide a brief description of the general motivation of the dissertation, as well as an outline of subsequent chapters.

## 1.1 Motivation

From the ancient practice of food foraging to the information-seeking objectives of the modern world, search and triage has always been and continues to be a cornerstone of the human experience. The slow evolutionary march has held humans largely to the same form, yet through tools, humanity has been able to break free of physiological binds, both physical and cognitive. Examples such as shovels, cars, and pencils reflect the power of physical tool augmentation, just as language, mathematics, and computers reflect the power of cognitive tool augmentation. User wellbeing is tied to the quality of their tools. When mindfully designed, tools help users complete their tasks – perhaps at a quicker pace, more effectively, for a longer time, or with less effort. Thus, if a task can be better understood, then designers can be cognizant of constraining features, formalize requirements which work around those constraints, then use those requirements to assess the quality of existing tools and in turn form new strategies that better serve users.

The 1950s initiated the use of computers as search tools, with research in academia and industry investigating strategies for indexing, weighting, and modeling the information of document collections (Sanderson & Croft, 2012). Technological and algorithmic solutions continued to mature and proliferate over the decades. Notably, the rise of reasonably accessible global networking in the 80s and 90s culminated in the first exposure to the power of generalized search engines, and in particular, for exploring the interconnected web (Haigh, 2011). Yet even as the computational technologies used by tools continue to improve – with faster computations, more intelligence, less power, and over larger document sets – users must still be the primary consideration when designing effective search and triage solutions. That is, technological innovation does not matter if the user does not know how to use those technologies for their task, understand their results, or activate that understanding in a manner that is meaningful to completion of their task.

It has never been more important for researchers to provide structured analysis on information-seeking tasks, from which guidance can be generated for designers to help realize more powerful and effective search and triage tools. Designers must concentrate first on addressing the user-facing considerations of search and triage interface design. For this effort, we center this dissertation on investigating how the designs of visual interfaces can

help users improve their searching and triaging of large document sets. In particular, we structure our research direction using the evolutionary design model. Models for evolutionary design, inspired by biological origins, encapsulate a four-staged, iterative process of formulization, realization, validation, and refinement (Baldominos, Saez, & Isasi, 2020). First, with guidance from prior interfaces, research, and framing devices, designers distill the necessary requirements of the task, decide how they are best addressed within the interface and its components, and formulize decisions within a design. From this design, a working prototype is realized, which is then validated through user-based methods of formative assessment. The findings of this verification are then propagated back to the designer for further formulization, realization, and verification, until requirements are satisfied (Guerrero-García, 2014). Using evolutionary design, VAT interface designers can provide more opportunities for novel thinking, de-couple design from prototyping, and promote user-centered design through formative assessment. Through this lens, we hope to contribute new insight which can be used within the designs of the next generation of information search and triage tools.

## 1.2 Structure of the Dissertation

This dissertation continues with the following chapters:

**Chapter 2** is concerned with directing background reading for the topics covered in this dissertation. Specifically, readers are provided an index to locations where appropriate background material is provided, and describes common terms and acronyms found within the material.

**Chapter 3** is concerned with exploring the design of generalized visual interfaces for information search and triage, and the activation of ontology file formats during query building. Included in this chapter is a presentation of OVERT-MED, an ontology-driven visual interface for searching and triaging MEDLINE. The creation of OVERT-MED examines the process of designing visual interfaces for information search and triage, with particular concentration on the use of ontologies as a mediating resource for aligning vocabularies during query building, and progressive disclosure in interface design. The research in this chapter is used as an exploratory proof-of-concept which inspired future work in later chapters.

**Chapter 4** is concerned with investigating how humans understand ontologies through cognitive map formation, and how visual interfaces can support users to encounter, form knowledge of, and explore complex ontological space. Included in this chapter is a presentation of PRONTOVISE, a progressively disclosed and scaffolded generalized visualization environment for exploring the ontological space of user-provided ontology files. This research is important to the overall research direction of this dissertation, as it helps establish if users can reasonably learn the complex ontological space described by ontology files. The research in this chapter inspires further investigations towards ontology activation within visual interfaces for information search and triaging tasks.

**Chapter 5** is concerned with framing the high-level requirements for generalized interfaces for information search and triage. Included in this chapter is an in-depth topic analysis which explores existing materials on data sources, information characteristics, types of search tasks, and considerations for generalized interfaces in the health domain. From this analysis, we distill a set of high-level requirements, which is then demonstrated by structuring the design of ONTSI, a generalized search interface which integrates ontologies for mediating between common and domain vocabularies. ONTSI is inspired by the first stage of the evolutionary design described in a later chapter.

**Chapter 6** is concerned with expanded investigations on the multi-staged information-seeking process and its design requirements. These investigations explore novel techniques which reduce the need for tedious search and triage. Specifically, it establishes how users struggle within traditional design strategies, align configurations with the requirements of each stage of the information-seeking process, and learns of strategies to promote the activation of domain expertise when searching and triaging. Furthermore, the use of ontologies to support these requirements is examined. Included in this chapter is a presentation of VisualQUEST, a novel visual interface which generalizes the support of information search and triage using the progressive disclosure and ontology mediation over user provided ontology and document datasets. VisualQUEST is built from the evolutionary design described in a later chapter.

**Chapter 7** is concerned with describing a three-staged evolutionary design of a Visual Analytics Tool (VAT) interface for searching and triaging large document sets. Specifically, the chapter outlines the formulation, realization, and validation of three interface stages. The formation of these stages is guided by in-depth topic analysis, design criteria, and formative assessment. Confirmatory evidence from formative assessment is presented.

**Chapter 8** concludes with chapter summaries, contributions, and future research directions stemming from the research of this dissertation.

### 1.3 References

- Baldominos, A., Saez, Y., & Isasi, P. (2020). On the automated, evolutionary design of neural networks: past, present, and future. *Neural Computing and Applications*, 32(2), 519–545. <https://doi.org/10.1007/s00521-019-04160-6>
- Guerrero-García, J. (2014). Evolutionary design of user interfaces for workflow information systems. *Science of Computer Programming*, 86, 89–102. <https://doi.org/10.1016/j.scico.2013.07.003>
- Haigh, T. (2011). The history of information technology. *Annual Review of Information Science and Technology*, 45, 431–487. <https://doi.org/10.1002/aris.2011.1440450116>
- Sanderson, M., & Croft, W. B. (2012). The history of information retrieval research. *Proceedings of the IEEE*, 100(SPL CONTENT), 1444–1451. <https://doi.org/10.1109/JPROC.2012.2189916>

## Chapter 2      **Background and Terminology**

### **2.1 Conceptual and Theoretical Background**

This dissertation follows an integrated article format, and therefore present material in an original, self-contained form. That is, each chapter maintains its original set of introductory, background, materials, results, and conclusion sections, presented in a fashion that can support both individual and progressive reading. For readers who seek background on the topics discussed in this dissertation, we ask that they be directed to the associated background and topic analysis written to contextualize that specific material in its original form. Yet, as many of the topics discussed within these materials share similar conceptual and theoretical background, we present an effective reading order for which can guide the reader through the primary thematic tones within this dissertation:

- 4.4.1 Cognitive Map Formation
- 5.2.1 Health Informatics
- 5.2.3 Ontologies
- 5.3.1.1 Health Data, Information Management, and Information-Centric Interfaces
- 5.3.1.2 Search Tasks and Structuring the Design of Interfaces for Health Informatics
- 5.3.1.3 Aligning Vocabularies for Health Informatics Search Tasks
- 6.2.1 Information Search and Triage
- 6.2.2 Machine Learning
- 6.3.1.1 Information-Seeking Process Models and Progressive Disclosure
- 6.3.1.2 Stages of Query Building and Search
- 6.3.1.3 Stages of High-Level and Low-Level Triage
- 7.2.1 Evolutionary Design

### **2.2 Terminology**

**User, human:** Within this dissertation, care is put to understanding how an interface is being acting upon, and how it can react in a manner that promotes the performance of whoever is acting upon it. Minor distinctions may arise in the reference of the ‘who’ in this relationship. The term ‘user’ and its plural ‘users’ may be found in cases when the generic actor is being discussed for their ‘using of’ an interface, whereas the usage of the term ‘human’ may prevail in discussions of a more general understanding of cognition, nature, and experience. This being said, within the contents of this dissertation, the terms are used interchangeably and should be understood uniformly.

**Cognitive tool, visual analytics tool, interactive visualization tool, health informatics tool:** Discussions within the various chapters of this dissertation may prefix tool with terms such as cognitive, visual analytics, interactive visualization, and health informatics. These terms maintain unique considerations which can differentiate their appropriate use. Yet, the research objective of this dissertation is the design of visual interfaces for information

search and triage in a generalized sense, regardless of the context or domain in question. Some chapters may maintain material that caters to the specific vocabulary and thematic tones of its published location. That is, while an individual chapter may apply a specific prefix, these specificities are used to align with the materials of that chapter and its published location. In general, these terms can and should be understood uniformly within the generalized theme of the dissertation.

## **2.3 Acronyms**

- HPO: Human Phenotype Ontology
- IVT: Interactive Visualization Tool
- MEDLINE: Medical Literature Analysis and Retrieval System Online
- MeSH: Medical Subject Header
- ML: Machine Learning
- NLM: National Library of Medicine
- ONTSI: ONTology-driven Search Interface
- OVERT-MED: Ontology-driven Visual sEaRch and Triage interface for MEDLINE
- PRONTOVISE: PRogressive ONTology VISualization Explorer
- PubMed: Publisher/Public MEDLINE
- VisualQUEST: Visual interface for QUERy, Search, Triage
- VAT: Visual Analytics Tool
- VR: Visual Representation

## Chapter 3      **OVERT-MED: An Ontology-Driven Visual Search and Triage Interface for MEDLINE**

This chapter has been published as: **Demelo, J.**, Parsons, P., & Sedig, K. (2017). Ontology-driven search and triage: Design of a web-based visual interface for MEDLINE. *JMIR medical informatics*, 5(1), e4.

We have made minor adjustments to the original material of this chapter to provide cohesion with the overall integrated article structure of this dissertation. Specifically, to distinguish between chapters, figures and tables have been provided an additional prepend reflecting the chapter number. Readers should be aware that chapter text will maintain original numbering references. For instance, “Figure 3-1” is equivalent to “Figure 1” in the chapter text.


### **3.1 Introduction**

Seeking information within the published medical literature is important in many domains and contexts (Islamaj Dogan, Murray, Neveol, & Lu, 2009; Krupski, Dahm, Fesperman, & Schardt, 2008). Diverse users need to search the literature including physicians (Kritz, Gschwandtner, Stefanov, Hanbury, & Samwald, 2013), medical students (William R. Hersh et al., 2002), cytogeneticists (Parsons et al., 2015), and patients and their relatives (Palotti, Hanbury, Müller, & Kahn, 2016). Searches can be roughly categorized into two types: *lookup* and *exploratory* (Marchionini, 2006). Lookup searches are closed-ended, having precise results and little need for examining and comparing result sets. Exploratory searches, however, are open-ended, having imprecise results and often requiring significant time and effort to work with result sets in order to satisfy the original information need. Examples of exploratory searches with open-ended goals include making evidence-based decisions and updating knowledge to stay abreast of current research findings (W R Hersh & Hickam, 1998; Islamaj Dogan et al., 2009). While significant progress has been made in supporting lookup searches, exploratory searches are still not well supported, and open-ended search goals are often quite difficult to achieve (Cui, Carter, & Zhang, 2014; Islamaj Dogan et al., 2009; Pang, Chang, Verspoor, & Pearce, 2016). Common barriers to finding relevant medical information include the time it takes to perform searches (Ely, Osheroff, Chambliss, Ebell, & Rosenbaum, 2005; Kritz et al., 2013), the increasing scope of topical coverage (Islamaj Dogan et al., 2009), and the information overload that arises from dealing with large result sets (K. Davies, 2007; Dietze et al., 2009; Ely et al., 2005; Islamaj Dogan et al., 2009; Kritz et al., 2013).

One of the most popular collections of published medical literature is MEDLINE, which comprises more than 25 million documents and is growing every year. The most common means of searching MEDLINE is PubMed, a free search engine and web interface (“PubMed,” n.d.). Although the search capabilities in PubMed have improved in recent years, there can still be a considerable burden on users when seeking information in the context of exploratory search, due to at least two major problems: 1) the difficulty in articulating information needs using accurate vocabulary; and 2) the large number of documents that can be returned from searches. Many users do not have the proper vocabulary to construct effective queries (Belkin, 2000; Furnas, Landauer, Gomez, & Dumais, 1987), which is especially true in medical and health contexts (Patrick, Monga, Sievert, Hall, & Longo, 2001; Plovnick &

Zeng, 2004; Sievert, Patrick, & Reid, 2001; Zeng & Tse, 2006a). When uncontrolled vocabularies are used, there is no guarantee that concepts are expressed with the same terms in different contexts (Dietze et al., 2009; Lowe & Barnett, 1994). For instance, if an article contains the term *eye hamartoma*, and a user searches for the vaguer term *eye growth*, there may not be a close match. Thus, without proper terminological knowledge, effective searching can be quite difficult. Adding to the difficulty of searching effectively is the large number of documents that can be returned, which leads to an information overload problem (Cui et al., 2014; Lu, 2011a; Malhotra et al., 2015). Dogan et al. (Islamaj Dogan et al., 2009) note that at least one third of PubMed searches return 100 or more documents. In our own testing, searches for common terms (e.g., “breast cancer” or “brain tumor”) return many thousands of documents.

Interfaces to most search engines, including PubMed, use simple text boxes into which users enter query terms. This interface style does not assist users in articulating their information needs (Hoerber & Khazaei, 2015), and works well only for lookup search tasks (M. Hearst et al., 2002; Hoerber, 2014). For example, if a user is interested in finding information about “liver” but is not sure what terms are relevant in articulating a query, she must simply enter “liver” into the search box. Because the query is vague, a very large set of documents is returned—almost one million documents spanning over 4900 pages when using PubMed (see Figure 1).



**Figure 3-1** A screenshot of PubMed showing results from searching for “liver”. In exploratory contexts, when users have open-ended information needs and imprecise vocabulary, PubMed’s interface style is not very helpful. Users have difficulty articulating information needs and must deal with long lists of results that span many pages.

Multiple strategies have been employed to help support query formation in exploratory search contexts by replacing the standard text box, including faceted search (Yee, Swearingen, Li, & Hearst, 2003), visualization widgets (Dork, Williamson, & Carpendale, 2009), query previews (Diriye, Tombros, & Blandford, 2012), and hierarchical presentation of expansion terms (Joho, Coverson, Sanderson, & Beaulieu, 2002). The common theme among these

strategies is that meaningful information is extracted from the document collection, and then represented in a manner that can help the searcher recognize terms that will more accurately describe the information they are seeking. Such strategies promote recognition over recall, not relying on users having to know and retrieve correct vocabulary from memory (Hoeber & Khazaei, 2015).

We present OVERT-MED (Ontology-driven Visual sEaRch and Triage interface for MEDLINE), a web-based visualization tool that addresses two major difficulties in searching large document collections: 1) the difficulty in articulating information needs with useful vocabulary, and 2) the difficulty in dealing with large search result sets. To address the first difficulty, we propose the idea of using a formal ontology to help users build domain-specific knowledge and vocabulary. To test this, we have implemented a searchable index of the Human Phenotype Ontology (HPO) that provides users with suggestion terms that are related to their information needs. To address the second difficulty, OVERT-MED supports multi-stage interactive triaging of search results using interactive visualization techniques. We use a custom-built index of MEDLINE, which comprises approximately 25 million documents, as our searchable collection of medical literature. Although OVERT-MED has been initially developed for use with a particular ontology and document collection, we expect that our design ideas will transfer to other contexts. The following subsections provide background information and discuss related work.

### 3.1.1 Ontologies

One way to meaningfully extract and model information from a domain is to construct an ontology (Chandrasekaran, Josephson, & Benjamins, 1999; Gruber, 1991). An ontology represents concepts and their relationships using a standard vocabulary (Chandrasekaran et al., 1999). Ontologies serve many practical functions, including clarifying the structure of knowledge within a domain, providing a common vocabulary, enabling computational analysis, and supporting knowledge sharing (Chandrasekaran et al., 1999; Gruber, 1991; Guarino, Oberle, & Staab, 2009). Ontologies often capture concepts within a domain at multiple levels of abstraction. For instance, an anatomy ontology may have a concept *body*, a sub-concept *face*, a further sub-concept *nose*, and so on. The concepts in an ontology can be represented using many different structures, including trees and different types of graphs.

The ontology we are using, HPO, has been curated by domain experts in an attempt to capture all phenotypic abnormalities that are commonly encountered in human monogenic disease (Robinson et al., 2008). In our previous work with genomics researchers, we learned of the importance of HPO in their workflow, including in activities involving literature search (Xxx & Xxx, n.d.). HPO is widely used in the biomedical field, is regularly updated, and has a high level of quality control. It is also available for download in the popular OBO (Open Biomedical Ontologies) and OWL (Web Ontology Language) formats. For these reasons, we believe HPO is ideal for testing our proposal of using ontologies to address the vocabulary problem. It should be noted that we are not suggesting HPO is better than other ontologies, or that it should be used in all contexts. HPO is only one of many ontologies that could be used to support exploratory search, and search systems should make use of whichever ontologies are most appropriate for given contexts.

### **3.1.2 Document Triage**

Triaging is an activity that involves determining the relevance of documents to an information need (Aekaterini Mavri, Fernando Loizides, Thomas Fotiades, 2014). Triaging activities are often time-constrained and require quick assessment of relevance with incomplete knowledge. For example, a search may return hundreds or thousands of potentially relevant documents. As it is not feasible to read each one in detail, users must sort through the documents and quickly assess their relevance based on incomplete knowledge of their contents. Research suggests that triaging takes place in three successive stages: 1) the “multiple document” stage, where initial relevance judgments are made to select documents from a set without careful examination; 2) the “individual document” stage, where individual documents are examined in more detail and categorized (e.g., kept or rejected); and 3) the “further reading” stage, where a small set of documents are read in-depth to extract relevant information and satisfy the original information need (Loizides & Buchanan, 2013). Additionally, research shows that triaging often occurs in a cyclical and iterative fashion, where the above stages are revisited multiple times (Loizides & Buchanan, 2009a).

### **3.1.3 Search Result Visualization**

Most search interfaces present results in a traditional list-based manner, where documents are ranked and textually represented using title and various metadata. While not a problem for simple lookup search tasks, traditional list-based representations are not effective in supporting exploratory search tasks, which are typically open-ended and involve complex information needs (Khazaei & Hoerber, 2016). Although lists are familiar and simple, studies show that users rarely examine lists fully or carefully (Spink, Wolfram, Jansen, & Saracevic, 2001) and seldom venture past the first few pages of results (Silverstein, Marais, Henzinger, & Moricz, 1999). Scanning through long lists can be tedious and cognitively demanding. Visualizations of search results can overcome some of the problems associated with textual list-based representations by shifting cognitive burden onto the perceptual system. For instance, while visualizations can be scanned freely by the eyes, text must be scanned sequentially, requiring more time and cognitive effort to detect patterns and relationships (Larkin & Simon, 1987; Scaife & Rogers, 1996). Additionally, visualizations can encode a significant amount of information within a small space, removing the need to navigate multiple pages to view search results. Previous work has demonstrated the utility of visualizations in document search, exploration, and analysis (Görg, Liu, & Stasko, 2013; M. A. Hearst, 1995).

### **3.1.4 Related Work**

Some researchers have recognized the value of using ontologies to better support search activities (e.g., (Dietze et al., 2009; Thomas, Alexopoulou, Dietze, & Schroeder, 2009)). The central focus of this research is term extraction and mapping, which is done using text mining and natural language processing techniques. In this body of work, ontologies are used to improve search performance computationally without involving users. The fundamental difference compared to our work is that we use ontologies to help users develop knowledge and domain-specific vocabulary—i.e., the focus is on the user rather than on algorithms and other computational processes. Our approach is important in contexts where users have valuable knowledge and context-specific goals that cannot be replaced by computation—in other words, users need to be kept “in the loop”.

Other researchers have focused on developing interfaces to MEDLINE as alternatives to PubMed. For example, Wei and colleagues have developed PubTator, a PubMed replacement interface that uses multiple text mining algorithms to improve search results (C. H. Wei, Kao, & Lu, 2013). PubTator also offers some support for document triaging. While PubTator appears interesting and useful, it relies on queries being input into the standard text box, and it presents results in a typical list-based fashion. Thus, it is not aimed at addressing either of the two problems we are attempting to address with OVERT-MED—i.e., the vocabulary problem and the information overload problem. Other alternative interfaces that offer interesting features but do not address either of the two problems include SLIM (Muin, Fontelo, Liu, & Ackerman, 2005) and HubMed (Eaton, 2006). An alternative interface that potentially provides support in addressing the first problem is iPubMed (J. Wang et al., 2011), which provides fuzzy matches to search results. An alternative interface that may provide support in addressing the second problem is refMED (Yu, Kim, Oh, Ko, & Kim, 2009), which provides minimal triaging support through relevance ranking. A for-profit private tool, Quertle, appears to use visualizations to mitigate the information overload problem, although very few details are publicly available. Lu (Lu, 2011b) provides a detailed survey that includes many other alternative interfaces to MEDLINE, although none are aimed at solving either of the two problems that we are addressing here.

In summary, no extant research explores the combination of (a) ontologies to help build domain-specific knowledge and vocabulary when users need to be kept “in the loop”; and (b) triaging support using interactive visualizations to help mitigate the information overload problem. The following sections provide details about our approach to addressing these issues.

## **3.2 Methods**

We developed OVERT-MED to test our proposed solutions to the two problems described previously. To anchor our research in a specific context, we chose MEDLINE as our document collection. MEDLINE offers an interesting testbed because of its popularity and size. We developed a custom index of MEDLINE so that it can be queried from the front end of OVERT-MED. We have also indexed HPO to help users build knowledge and domain-specific vocabulary.

### **3.2.1 Indexing of MEDLINE and HPO**

We downloaded the entire MEDLINE database, which is made freely available by the NLM for research purposes. The MEDLINE database consists of article “citations”, which are essentially article metadata, including authors, journal title, Medical Subject Heading (MeSH) keywords, publication date, and other fields. Also included in each citation is the abstract text. We developed a custom index using the open-source Apache Solr/Lucene project. Lucene supports full-text indexing and search functionality, and Solr is a search platform that runs on the Lucene index. To rank documents, Lucene uses the well-known term frequency-inverse document frequency (tf-idf) scheme (Gerard Salton & Buckley, 1988). Lucene also ranks results based on an internal similarity measure that generates a vector space model (VSM) score (G. Salton, Wong, & Yang, 1975), using index terms as dimensions and tf-idf values as weights. We have described our indexing strategy in more detail previously (Xxx & Xxx, n.d.).

HPO is a formal ontology of human phenotypic abnormalities found in human disease (Robinson et al., 2008). Each entry in HPO describes a phenotypic abnormality, such as melanoma or hepatoblastoma. HPO is under active development and currently contains over 11,000 terms. We have also indexed HPO in our Lucene index. HPO contains multiple fields for each phenotype in the ontology, including name, definition, id, synonyms, and commentary from domain experts. We index all fields to provide robust vocabulary suggestions—when a user enters a term, all fields in the index are examined, which provides much more useful information than would result from looking for only exact matches on the phenotype name. This will be described using an example in more detail below.

### 3.2.2 Development and Architecture

We developed OVERT-MED as a web-based tool that runs in any modern browser. It connects to a web server that stores our indices and handles search requests (via our Solr search server). We have developed a series of scripts to retrieve MEDLINE updates from the NLM public ftp site and to construct the indices for MEDLINE and HPO in our Lucene index. We have also developed an API that handles requests for searches and other basic functions. The front-end has been developed using HTML5, CSS, and JavaScript. The visualizations have been developed using D3.js (Michael Bostock, Ogievetsky, & Heer, 2011), a popular JavaScript visualization library. Figure 2 provides a diagrammatic overview of the architecture of the OVERT-MED system.





Figure 3-2 Client-server architecture of the OVERT-MED system. The webservice fetches MEDLINE updates from the NLM ftp site and indexes them in a Lucene index. HPO is also indexed. The API handles requests and connects to the Solr search server. The client displays visualizations and handles user interactions.

## 3.3 Results

### 3.3.1 Ontology Term Suggestion

OVERT-MED uses HPO to help users better articulate their search needs through a technique we call *Ontology Term Suggester*. Users enter terms into a text box, and a set of suggestions (phenotypes) are provided. The suggestions are updated in real-time as a user types each character. Additionally, to provide better terminological support, we look

for matches on both the phenotype names as well as descriptions and expert commentary on the phenotypes (these are not shown to users but are indexed on our server). For example, a user may be interested in finding articles related to the term “liver” but may not have sufficient vocabulary to articulate a useful query involving relevant terms. Figure 3 shows the Ontology Term Suggester after typing “liver” into the search box. Phenotypes related to the liver are displayed. Results such as “Growth hormone deficiency” and “Ascites” are displayed because they have a connection to the liver—the effects of growth hormone are mediated by insulin-like growth factor, which is produced primarily in the liver; and ascites is commonly associated with liver disease. Many of the returned phenotypes do not have the term liver in their name but are related to the liver. In a traditional search interface, there is no way for a user to get from “liver” to “ascites” or “growth hormone deficiency”. Finally, because users may not understand a particular phenotype (e.g., congenital diaphragmatic hernia), selecting the ‘?’ button will open a new tab and load the official entry in the HPO online browser. From there users can find more details, including associated genes and diseases. This search strategy can help users build knowledge of the domain and vocabulary that can be used to enhance cognitive performance and exploration.




**Figure 3-3 The Ontology Term Suggester, showing results from typing “liver”. Many of the resulting phenotypes do not have “liver” in their names, but are in some way associated with liver, and are displayed to help users improve vocabulary to better articulate information needs.**

### 3.3.2 Sensitivity Encoding for Query Refinement

A well-known problem in open-ended search tasks is that potentially relevant results may not be displayed if they do not meet the specified search criteria. For example, when searching for a house to buy, users often have ill-formed criteria, such as price range, number of bedrooms and bathrooms, yard size, location, and so on. Although certain search criteria may be specified (e.g., 4 bedrooms, under \$200,00), results that do not meet the criteria may also be relevant, such as a house that has only three bedrooms but is a great price. When using visualizations to support such search tasks, certain criteria can be relaxed and results that do not meet certain criteria can be visually encoded in

different ways. For instance, results that do not meet number of bedrooms can be encoded with one color, results that do not meet yard size can be encoded with another, and so on. Visually encoding this type of information can provide cues to users to adjust their search criteria so that potentially relevant results are included. This visualization strategy, known as sensitivity encoding, has been shown to be beneficial in a number of contexts (Spence, 2002; Spence & Tweedie, 1998).

Although OVERT-MED supports the selection of precise phenotype names, the exact combination of words in a name may be too restrictive and may not provide the most relevant results. For example, a user may select the phenotype *progressive external ophthalmoplegia*. Our index shows 811 articles associated with this specific phenotype. However, users may be interested in articles associated with different variations of the words—e.g., *progressive ophthalmoplegia* or *external ophthalmoplegia*. We employ a set of *Sensitivity Encoded Query Selectors* in OVERT-MED to handle this issue. When a phenotype is selected, we perform searches on our index using all possible combinations of the words, then visually encode the size of the result set. Figure 4 shows the result of a user selecting “progressive external ophthalmoplegia”. The number of matching articles for each combination is provided numerically and encoded visually using the length of the bar next to each combination. From Figure 4 we can see that if the user relaxes the term to only “progressive ophthalmoplegia” an additional 104 articles are available; with “external ophthalmoplegia” and additional 418 articles are available. Without such a sensitivity encoding strategy, many of these potentially relevant results would not be made available. Because users are often interested in more than one phenotype, multiple phenotypes can be selected, each of which is subjected to the same sensitivity encoding process. Figure 5 shows a second phenotype, congenital fibrosis of extraocular muscles, being added.



**Figure 3-4** A set of Sensitivity Encoded Query Selectors for ‘progressive external ophthalmoplegia’. The number of matching articles for each combination is provided numerically and encoded visually using the length of the bar next to each combination.




Figure 3-5 The result of adding a second phenotype via the Ontology Term Suggester, which leads to more Sensitivity Encoded Query Selectors.

### 3.3.3 Interactive Triaging Support to Mitigate Information Overload

OVERT-MED provides multi-stage triaging support to mitigate the information overload problem. Multiple design strategies support the first stage of triaging—the “multiple document” stage. First, when a specific set of terms is chosen, the metadata from up to 250 documents is visualized. Each document is encoded using a small bar, and the presence of each term is encoded using a section of the bar. Figure 6 shows how six documents are represented where there are three terms (progressive external ophthalmoplegia). Within the visualization, each row represents one document, and each column represents one of the phenotype words. The words are color coded—in this case green for progressive, teal for external, and red for ophthalmoplegia. A white cell indicates no occurrence of the word. The visualization functions as a type of heatmap (Wilkinson & Friendly, 2009), where the color saturation encodes the frequency of a term within a document. We call this technique the *Query Result Heatmap*. In Figure 6, a darker red means higher occurrence of the word ophthalmoplegia. This type of encoding can aid in rapid visual scanning and identification of potentially relevant documents (M. A. Hearst, 1995; Hoerber & Yang, 2009).



Figure 3-6 The Query Result Heatmap: Six documents are represented by six rows, where each column represents a term (progressive external ophthalmoplegia).

To further support the triaging activity, OVERT-MED allows users to interactively explore metadata associated with the matching documents. Figure 7 shows the state of the interface after a user has selected “progressive+ophthalmoplegia”. The first 250 documents (ranked by our indexing algorithm) are encoded in the Query Result Heatmap. Each row functions as an individual document heatmap, showing the occurrence of the seven phenotype terms within the document. Because the user has selected “progressive” and “ophthalmoplegia”, all documents indicate occurrences of both terms. It is readily apparent that most of the documents also contain the term “external”. Approximately twenty also contain “muscles”, four contain “extraocular”, one contains “fibrosis”, and one contains “congenital”.




Figure 3-7 State of the interface after a user has selected “progressive+ophthalmoplegia”.

OVERT-MED also provides a *Term Distribution Matrix* to help users quickly determine document relevance while browsing the Query Result Heatmap. Within the Term Distribution Matrix, users can see the occurrence of terms in four places within the document metadata: 1) title, 2) journal name, 3) MeSH terms, and 4) abstract text. The document title, journal, year, and MeSH terms are also displayed. This representation helps users make decisions about relevance via quick visual scanning. For example, if a term appears only in the journal name it may not be very relevant, but if a term appears five times in the abstract text it is more likely to be relevant. Users can perceive this type of information quickly due to the categorical color encodings. Figure 8 shows the Term Distribution Matrix for two different documents within the same result set. Through rapid visual scanning, even without reading the text, it is apparent that the terms are quite important in the document on the right.




Figure 3-8 The Term Distribution Matrix for two different documents within the same result set.

To support rapid exploration—a fundamental goal of triaging—the keyboard arrow keys can be used to move quickly through the documents while the metadata is dynamically updated. If a relevant document is detected, users can hit the ‘enter’ key or click the button to add the document to a pile for subsequent investigation (this stage will be explained in more detail below). This stage of triaging also allows for quick comparison of co-occurring phenotypes within documents. For example, Figure 9 shows the result of a user adding documents containing “congenital” and “fibrosis”. It is immediately clear through quick visual scanning that not many documents contain both “congenital fibrosis” and “ophthalmoplegia”.





Figure 3-9 The result of a user adding documents containing “congenital” and “fibrosis” for comparison. It is immediately clear through quick visual scanning that not many documents contain both “congenital fibrosis” and “ophthalmoplegia”.

While browsing the Query Result Heatmap, it may be difficult to remember which documents have been visited previously. This is especially true in the context of iterative triaging, where users may return to the heatmap after being away for some time. In OVERT-MED, when users pause on a document for five or more seconds, a small mark is placed beside the document to serve as a visual reminder (see Figure 10). When revisiting the heatmap, users

can quickly recognize which documents they have previously examined. We assume that five seconds is a reasonable threshold for determining when a user has examined the Term Distribution Matrix.



**Figure 3-10 Close-up view of the Query Result Heatmap. When users pause on a document for five or more seconds a small mark is placed beside the document as a reminder.**

The next stage in the triaging activity—the “individual document” stage—involves examining individual abstracts of previously chosen articles. At this stage, users are likely to have narrowed down the number of documents significantly. Documents are encoded via a *Selected Pile Heatmap* in the same manner as in the Query Result Heatmap, and each can be selected to view its abstract. In this *Term-Encoded Abstract*, matching terms are color-coded to facilitate quick identification, especially within the abstract text. Figure 11 shows an example in which the user has selected 29 documents, which are encoded in the Selected Pile Heatmap and the Term-Encoded Abstract is displayed for the first document. Even before reading the text in detail, it is easy to see that ‘renin’ and ‘hypertension’ both appear frequently, indicating that they are important. Thus, users can scan the text quickly to get a sense of the appearance of the query terms, without having to necessarily read the text sequentially. An important aspect of this stage of triaging is the ability to quickly categorize documents. In OVERT-MED, users can quickly reject a paper by selecting the orange ‘x’ button or can quickly add a paper to the next stage by selecting the green button or pressing the ‘enter’ key.




Figure 3-11 Twenty-nine documents have been selected to examine in closer detail. They are encoded in the Selected Pile Heatmap, and the Term-Encoded Abstract is displayed for the first document. It is easy to see that ‘renin’ and ‘hypertension’ both appear frequently, indicating that they are important.

The final stage of triaging is the “further reading” stage, where a small set of documents are read in-depth to extract relevant information and satisfy the original information need. While this stage could be supported in various ways, we support this stage in OVERT-MED by presenting a PubMed entry for a selected document in an embedded frame directly within the interface of OVERT-MED. This allows for quick inspection of any PubMed details that are important to the user, such as full text links, citation details, and PubMed Commons links, and allows users to login to their NCBI account to save the article to a collection, compare to other saved articles, and so on. There is also a button to open the PubMed link in a new browser tab if a user needs more space. Figure 12 shows a full screen capture of OVERT-MED in which a user has traversed all stages of a search and triaging activity.




Figure 3-12 Full screen capture showing all components of OVERT-MED where a user has traversed all stages of a search and triaging activity.

As research shows that triaging activities are cyclical and iterative, we have designed OVERT-MED to be flexible in this regard. At any point during an activity, users may adjust their query or document selections, and each component of the interface will dynamically reflect any changes. For example, a user may reach the final stage of

triaging and find a term within a document that seems relevant to the original information need. The user can return to the initial stage of entering the term and selecting phenotypes. In doing so, the rest of the interface will remain stable, and the user can proceed through any of the triaging stages. Figure 13 shows the interface after a user has examined a document in detail in the final stage, discovered a link between renin level (the original phenotype of interest) and arterial pressure, and has returned to the initial stage to find a phenotype related to arterial pressure. The user discovers a phenotype named “elevated mean arterial pressure” and selects it. At this stage, the user is not particularly interested in whether the arterial pressure is elevated or not, and simply wants to explore the relationship between renin level and arterial pressure. Due to our sensitivity encoding strategy, the user can select “arterial+pressure” to add documents with those two terms. From this point, the user can continue through the triaging stages or return to the initial stage again.





Figure 3-13 The interface after a user has examined a document in detail in the final stage, discovered a link, and has returned to the initial stage with a new information need. The user discovers a phenotype named “elevated mean arterial pressure”, selects it, then adds documents satisfying “arterial+pressure” to the Query Result Heatmap.

### 3.4 Discussion

OVERT-MED was developed to address two major problems that are known to exist in complex, exploratory search activities: 1) the difficulty in articulating information needs due to insufficient knowledge and domain-specific vocabulary, and 2) the difficulty in dealing with information overload due to the large number of results returned. To address the first difficulty, we proposed the idea of using a formal ontology to help users build domain-specific terminology and knowledge for constructing search queries. To assist in this process, we indexed HPO and provided a search feature that provides robust results to terms that are entered. To address the problem of search criteria being too restrictive in open-ended contexts, we employed a visual sensitivity encoding strategy to help users see possibilities with different combinations of terms.

There are seven main steps that users take when performing search and triaging tasks with OVERT-MED—the first two within a vocabulary building phase, and the next five within a triaging phase. The triaging phase can be broken down into the three key stages. Figure 14 provides an overview of this process and shows the techniques we employ to help users at each step. To help users build vocabulary and generate queries, we use an *Ontology Term*

*Suggester* and *Sensitivity Encoded Query Selectors*. After selecting a query, users move to the triaging phase, where they traverse through three stages. During the first stage—the multi-document stage—users are presented with a *Query Result Heatmap* that encodes the appearance and frequency of query terms within the document result set. A keyboard interaction technique enables rapid navigation through the documents. To facilitate assessment at this stage, a *Term Distribution Matrix* provides more information about each document within the heatmap. Together these techniques allow for rapid scanning to assess relevance and select documents for the next stage. During the second triaging stage—the individual document stage—users are presented with a *Selected Pile Heatmap* that encodes only the selected documents from the previous stage. As users browse the heatmap, they can inspect a *Term-Encoded Abstract* of each individual document. The term-encoding supports quick detection of the appearance of query terms within the document abstract. After assessing the relevance of individual documents, users select documents to move to the next stage. During the third triaging stage—the further reading stage—users focus on a single document by viewing details in depth. Here the PubMed entry for a document can be retrieved directly within OVERT-MED or within a new browser tab. At any point in the overall activity, users can return to any step and continue from there, which supports the iterative and cyclical nature of search and triaging tasks.



**Figure 3-14 Overall search and triage process supported by OVERT-MED. Users take seven main steps—the first two within a vocabulary building phase, and the next five within a triaging phase. The triaging phase can be broken down into the three key stages. Each step is supported by a technique in the interface. At any point users can return to any step, which supports the iterative and cyclical nature of search and triaging tasks.**

### 3.4.1 Validation

Ongoing formative evaluation suggests that the design features in OVERT-MED can mitigate the two problems mentioned above. We tested OVERT-MED with a small group of users who are not domain-experts, and our proposal to use a formal ontology to help users articulate their information needs does seem to be useful. As mentioned previously, different types of users are known to search the scientific literature, many of which are not domain experts. For example, pediatricians will often try to identify abnormal phenotypes in patients before referring them to a clinical geneticist. However, because they are not domain experts, pediatricians may not have very extensive knowledge and vocabulary of phenotypes. Even if they search the literature to identify phenotype names (e.g., via

PubMed), they may still not find phenotypes that are related to one another. As another example, patients are known to search the literature to learn more about their own conditions. Because they are not domain experts, patients could also benefit from having access to an ontology such as HPO to help them build domain-specific knowledge and vocabulary. Thus, testing with users who are not domain experts can give an indication of the usefulness of our design strategies.

In our testing, we noticed that although an ontology can help users develop more appropriate vocabulary, users do not necessarily develop a good understanding of the ontology itself. Because a robust mental model of the ontology may lead to even better search performance (e.g., by knowing which entities are highly connected to others, knowing relationships among entities at multiple levels of abstraction, and so on), we have decided to pursue a solution to this as future work (see Future Work section below). Additionally, our multi-stage triaging shows promise in mitigating the information overload problem. Users were able to go back and forth through the triaging stages to satisfy information needs without being overwhelmed by long lists of documents.

### **3.4.2 Limitations**

There is one current limitation of OVERT-MED that should be noted: the MEDLINE data is limited to metadata and abstract text only and does not include full texts. This is simply because the NLM does not release full-texts due to copyright issues. There is little we can do to address this issue. Empirical evidence does suggest, however, that the document title and abstract are among the most important features of a document in determining its relevance (Loizides & Buchanan, 2009a), so perhaps it is not a critical limitation.

### **3.4.3 Future Work**

We envision at least three lines of valuable future research: First, developing interactive visualization techniques to support ontology sensemaking. The intention behind the current version of OVERT-MED is to help address the common problem of lack of adequate vocabulary. Although OVERT-MED appears to support users in improving their search terms and potentially developing some domain knowledge, it does not necessarily support users in making sense of the ontology itself — i.e., understanding its size, organization, types of relationships, significant and insignificant entities, and so on. Interactive visualizations of ontologies may enhance search and triaging activities. Second, testing OVERT-MED with different ontologies in different contexts. This will help to assess the transferability of the design features of OVERT-MED. Third, conducting formal testing of OVERT-MED. Although our informal testing has been useful, more formal testing will provide validation of the design strategies.

### **3.4.4 Conclusions**

We have developed a web-based interactive visualization tool, OVERT-MED, to address two common problems in exploratory search — namely, the lack of adequate vocabulary to construct useful queries, and the difficulty of dealing with very large result sets. The novelty of our approach is in the combination of (a) using an ontology to help build domain-specific knowledge and vocabulary when users need to be kept “in the loop”; and (b) providing multi-

stage triaging support using interactive visualizations to help mitigate the information overload problem. We anticipate these ideas can be applied successfully in other contexts where either of these issues exist.

### 3.5 References

- Aekaterini Mavri, Fernando Loizides, Thomas Fotiades, P. Z. (2014). We Have the Content...Now What? The role of Structure and Interactivity in Academic Document Triage Interfaces. *Information Design Journal*, 20(3), 247–265. <https://doi.org/10.1075/idj.20.3.05mav>
- Belkin, N. J. (2000). Helping people find what they don't know. *Communications of the ACM*, 43(8), 58–61. <https://doi.org/10.1145/345124.345143>
- Bostock, M., Ogievetsky, V., & Heer, J. (2011). D3 data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), 2301–2309. <https://doi.org/10.1109/TVCG.2011.185>
- Chandrasekaran, B., Josephson, J. R., & Benjamins, V. R. (1999). What are ontologies, and why do we need them? *Intelligent Systems and Their Applications, IEEE*, 14(1), 20–26. <https://doi.org/10.1109/5254.747902>
- Cui, L., Carter, R., & Zhang, G. Q. (2014). Evaluation of a novel conjunctive exploratory navigation interface for consumer health information: A crowdsourced comparative study. *Journal of Medical Internet Research*, 16(2), 1–15. <https://doi.org/10.2196/jmir.3111>
- Davies, K. (2007). The information-seeking behaviour of doctors: A review of the evidence. *Health Information and Libraries Journal*, 24(2), 78–94. <https://doi.org/10.1111/j.1471-1842.2007.00713.x>
- Dietze, H., Alexopoulou, D., Alvers, M. R., Barrio-Alvers, L., Andreopoulos, B., Doms, A., ... Schroeder, M. (2009). GoPubMed: Exploring PubMed with Ontological Background Knowledge. In S. Krawetz (Ed.), *Bioinformatics for Systems Biology* (pp. 385–399). [https://doi.org/10.1007/978-1-59745-440-7\\_20](https://doi.org/10.1007/978-1-59745-440-7_20)
- Diriye, A., Tombros, A., & Blandford, A. (2012). A Little Interaction Can Go a Long Way: Enriching the Query Formulation Process. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics): Vol. 7224 LNCS* (pp. 531–534). [https://doi.org/10.1007/978-3-642-28997-2\\_57](https://doi.org/10.1007/978-3-642-28997-2_57)
- Dork, M., Williamson, C., & Carpendale, S. (2009). Towards Visual Web Search : Interactive Query Formulation and Search Result Visualization. *WSSP 2009: WWW Workshop on Web Search Result Summarization and Presentation*, 2–5.
- Eaton, A. D. (2006). PubMed: A web-based biomedical literature search interface. *Nucleic Acids Research*, 34(WEB. SERV. ISS.). <https://doi.org/10.1093/nar/gkl037>
- Ely, J. W., Osheroff, J. A., Chambliss, M. L., Ebell, M. H., & Rosenbaum, M. E. (2005). Answering physicians' clinical questions: Obstacles and potential solutions. *Journal of the American Medical Informatics Association*, 12(2), 217–224. <https://doi.org/10.1197/jamia.M1608>
- Furnas, G. W., Landauer, T. K., Gomez, L. M., & Dumais, S. T. (1987). The vocabulary problem in human-system communication. *Communications of the ACM*, 30(11), 964–971. <https://doi.org/10.1145/32206.32212>
- Görg, C., Liu, Z., & Stasko, J. T. (2013). Reflections on the evolution of the Jigsaw visual analytics system. *Information Visualization*. <https://doi.org/10.1177/1473871613495674>

- Gruber, T. (1991). The role of common ontology in achieving sharable, reusable knowledge bases. *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, 601–602.
- Guarino, N., Oberle, D., & Staab, S. (2009). What is an ontology? In S. Staab & R. Studer (Eds.), *Handbook on Ontologies* (pp. 1–17). <https://doi.org/10.1007/978-3-540-92673-3>
- Hearst, M. A. (1995). TileBars: Visualization of Term Distribution Information in Full Text Information Access. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '95*, 59–66. <https://doi.org/10.1145/223904.223912>
- Hearst, M., Elliott, A., English, J., Sinha, R., Swearingen, K., & Yee, K.-P. (2002). Finding the flow in web site search. *Communications of the ACM*, 45(9). <https://doi.org/10.1145/567498.567525>
- Hersh, W R, & Hickam, D. H. (1998). How well do physicians use electronic information retrieval systems? A framework for investigation and systematic review. *JAMA : The Journal of the American Medical Association*, 280(15), 1347–1352. <https://doi.org/10.1001/jama.280.15.1347>
- Hersh, William R., Katherine Crabtree, M., Hickam, D. H., Sacherek, L., Friedman, C. P., Tidmarsh, P., ... Kraemer, D. (2002). Factors associated with success in searching MEDLINE and applying evidence to answer clinical questions. *Journal of the American Medical Informatics Association*, 9(3), 283–293. <https://doi.org/10.1197/jamia.M0996>
- Hoeber, O. (2014). Visual Search Analytics: Combining Machine Learning and Interactive Visualization to Support Human-Centred Search. *Beyond Single-Shot Text Queries: Bridging the Gap(s) Between Research Communities*, 37–43.
- Hoeber, O., & Khazaei, T. (2015). Evaluating citation visualization and exploration methods for supporting academic search tasks. *Online Information Review*, 39(2), 229–254. <https://doi.org/10.1108/OIR-10-2014-0259>
- Hoeber, O., & Yang, X. D. (2009). HotMap: Supporting visual exploration of Web search results. *Journal of the American Society for Information Science and Technology*, 60(1), 90–110. <https://doi.org/10.1002/asi>
- Islamaj Dogan, R., Murray, G. C., Neveol, A., & Lu, Z. (2009). Understanding PubMed user search behavior through log analysis. *Database, 2009*, bap018–bap018. <https://doi.org/10.1093/database/bap018>
- Joho, H., Coverson, C., Sanderson, M., & Beaulieu, M. (2002). Hierarchical presentation of expansion terms. *Proceedings of the 2002 ACM Symposium on Applied Computing - SAC '02*, 645. <https://doi.org/10.1145/508791.508916>
- Khazaei, T., & Hoeber, O. (2016). Supporting academic search tasks through citation visualization and exploration. *International Journal on Digital Libraries*, 1–14. <https://doi.org/10.1007/s00799-016-0170-x>
- Kritz, M., Gschwandtner, M., Stefanov, V., Hanbury, A., & Samwald, M. (2013). Utilization and perceived problems of online medical resources and search tools among different groups of european physicians. *Journal of Medical Internet Research*, 15(6), e122. <https://doi.org/10.2196/jmir.2436>
- Krupski, T. L., Dahm, P., Fesperman, S. F., & Schardt, C. M. (2008). Users ' Guide to the Urological Literature How to Perform a Literature Search. *Clinical Evidence*, 179(April), 1264–1270. <https://doi.org/10.1016/j.juro.2007.11.087>
- Larkin, J., & Simon, H. (1987). Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive Science*,

11(1), 65–100.

- Loizides, F., & Buchanan, G. (2009). An empirical study of user navigation during document triage. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 5714 LNCS, 138–149. [https://doi.org/10.1007/978-3-642-04346-8\\_15](https://doi.org/10.1007/978-3-642-04346-8_15)
- Loizides, F., & Buchanan, G. (2013). Towards a framework for human (manual) information retrieval. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8201 LNCS, 87–98. [https://doi.org/10.1007/978-3-642-41057-4\\_10](https://doi.org/10.1007/978-3-642-41057-4_10)
- Lowe, H. J., & Barnett, G. O. (1994). Understanding and Using the Medical Subject Headings ( MeSH ) Vocabulary to Perform Literature Searches. *Journal of the American Medical Association*, 271(14), 1103–1108.
- Lu, Z. (2011a). PubMed and beyond: a survey of web tools for searching biomedical literature. *Database: The Journal of Biological Databases and Curation*, 2011(0), baq036. <https://doi.org/10.1093/database/baq036>
- Lu, Z. (2011b). PubMed and beyond: A survey of web tools for searching biomedical literature. *Database*, 2011, baq036. <https://doi.org/10.1093/database/baq036>
- Malhotra, A., Gündel, M., Rajput, A. M., Mevissen, H. T., Saiz, A., Pastor, X., ... Villoslada, P. (2015). Knowledge retrieval from pubmed abstracts and electronic medical records with the multiple sclerosis ontology. *PLoS ONE*, 10(2), e0116718. <https://doi.org/10.1371/journal.pone.0116718>
- Marchionini, G. (2006). Exploratory search: from finding to understanding. *Communications of the ACM*, 49(4), 46.
- Muin, M., Fontelo, P., Liu, F., & Ackerman, M. (2005). SLIM: an alternative Web interface for MEDLINE/PubMed searches - a preliminary study. *BMC Medical Informatics and Decision Making*, 5, 37. <https://doi.org/10.1186/1472-6947-5-37>
- Palotti, J., Hanbury, A., Müller, H., & Kahn, C. E. (2016). How users search and what they search for in the medical domain: Understanding laypeople and experts through query logs. *Information Retrieval*, 19(1–2), 189–224. <https://doi.org/10.1007/s10791-015-9269-8>
- Pang, P. C. I., Chang, S., Verspoor, K., & Pearce, J. (2016). Designing health websites based on users' web-based information-seeking behaviors: A mixed-method observational study. *Journal of Medical Internet Research*, 18(6). <https://doi.org/10.2196/jmir.5661>
- Parsons, P., Sedig, K., Mercer, R. E., Khordad, M., Knoll, J., & Rogan, P. (2015). Visual analytics for supporting evidence-based interpretation of molecular cytogenomic findings. *Proceedings of the 2015 Workshop on Visual Analytics in Healthcare - VAHC '15*, 1–8. <https://doi.org/10.1145/2836034.2836036>
- Patrick, T. B., Monga, H. K., Sievert, M. C., Hall, J. H., & Longo, D. R. (2001). Evaluation of controlled vocabulary resources for development of a consumer entry vocabulary for diabetes. *Journal of Medical Internet Research*, 3(3), 5–26. <https://doi.org/10.2196/jmir.3.3.e24>
- Plovnick, R. M., & Zeng, Q. T. (2004). Reformulation of consumer health queries with professional terminology: A pilot study. *Journal of Medical Internet Research*, 6(3). <https://doi.org/10.2196/jmir.6.3.e27>
- PubMed. (n.d.). Retrieved from <https://pubmed.ncbi.nlm.nih.gov/>
- Robinson, P. N., Köhler, S., Bauer, S., Seelow, D., Horn, D., & Mundlos, S. (2008). The Human Phenotype Ontology: A Tool for Annotating and Analyzing Human Hereditary Disease. *American Journal of Human Genetics*, 83,

- 610–615. <https://doi.org/10.1016/j.ajhg.2008.09.017>
- Salton, G., Wong, A., & Yang, C. S. (1975). A vector space model for automatic indexing. *Communications of the ACM*, 18(11), 613–620. <https://doi.org/10.1145/361219.361220>
- Salton, Gerard, & Buckley, C. (1988). Term-weighting approaches in automatic text retrieval. *Information Processing & Management*, 24(5), 513–523. [https://doi.org/10.1016/0306-4573\(88\)90021-0](https://doi.org/10.1016/0306-4573(88)90021-0)
- Scaife, M., & Rogers, Y. (1996). External cognition: how do graphical representations work? *International Journal of Human-Computer Studies*, 45(2), 185–213. <https://doi.org/10.1006/ijhc.1996.0048>
- Sievert, M. C., Patrick, T. B., & Reid, J. C. (2001). Need a bloody nose be a nosebleed? or, lexical variants cause surprising results. *Bulletin of the Medical Library Association*, 89(1), 68–71.
- Silverstein, C., Marais, H., Henzinger, M., & Moricz, M. (1999). Analysis of a very large web search engine query log. *ACM SIGIR Forum*, 33(1), 6–12. <https://doi.org/10.1145/331403.331405>
- Spence, R. (2002). Sensitivity encoding to support information space navigation: a design guideline. *Information Visualization*, 1(2), 120–129. <https://doi.org/10.1057/palgrave.ivs.9500019>
- Spence, R., & Tweedie, L. (1998). The Attribute Explorer: information synthesis via exploration. *Interacting with Computers*, 11(2), 137–146. [https://doi.org/10.1016/S0953-5438\(98\)00022-8](https://doi.org/10.1016/S0953-5438(98)00022-8)
- Spink, A., Wolfram, D., Jansen, M. B. J., & Saracevic, T. (2001). Searching the Web: The Public and Their Queries. *Journal of the American Society for Information Science and Technology*, 52(3), 226–234. [https://doi.org/10.1002/1097-4571\(2000\)9999:9999<::AID-ASII591>3.0.CO;2-R](https://doi.org/10.1002/1097-4571(2000)9999:9999<::AID-ASII591>3.0.CO;2-R)
- Thomas, W., Alexopoulou, D., Dietze, H., & Schroeder, M. (2009). Searching Biomedical Literature with Anatomy Ontologies. *Anatomy Ontologies for Bioinformatics: Principles and Practice*.
- Wang, J., Cetindil, I., Ji, S., Li, C., Xie, X., Li, G., ... Bateman, A. (2011). Interactive and fuzzy search: A dynamic way to explore MEDLINE. *Bioinformatics*, 27(13), 2321–2327. <https://doi.org/10.1093/bioinformatics/btq414>
- Wei, C. H., Kao, H. Y., & Lu, Z. (2013). PubTator: a web-based text mining tool for assisting biocuration. *Nucleic Acids Research*, 41(Web Server issue), 518–522. <https://doi.org/10.1093/nar/gkt441>
- Wilkinson, L., & Friendly, M. (2009). The History of the Cluster Heat Map. *The American Statistician*, 63(2), 179–184. <https://doi.org/10.1198/tas.2009.0033>
- Xxx, X., & Xxx, X. (n.d.). *hidden for blinded review*.
- Yee, K.-P., Swearingen, K., Li, K., & Hearst, M. A. (2003). Faceted metadata for image search and browsing. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*, 401–408. <https://doi.org/10.1145/642679.642681>
- Yu, H., Kim, T., Oh, J., Ko, I., & Kim, S. (2009). RefMed: relevance feedback retrieval system fo PubMed. *Proceeding of the 18th ACM Conference on Information and Knowledge Management - CIKM '09*, 2099–2100. <https://doi.org/10.1145/1645953.1646322>
- Zeng, Q. T., & Tse, T. (2006). Exploring and developing consumer health vocabularies. *Journal of the American Medical Informatics Association*, 13(1), 24–29. <https://doi.org/10.1197/jamia.M1761>

## Chapter 4      **Forming Cognitive Maps of Ontologies Using Interactive Visualizations**

This chapter has been published as: **Demelo, J., & Sedig, K. (2021).** Forming Cognitive Maps of Ontologies Using Interactive Visualizations. *Multimodal Technologies and Interaction*, 5(1), 2.

We have made minor adjustments to the original material of this chapter to provide cohesion with the overall integrated article structure of this dissertation. Specifically, to distinguish between chapters, figures and tables have been provided an additional prepend reflecting the chapter number. Readers should be aware that chapter text will maintain original numbering references. For instance, “Figure 4-1” is equivalent to “Figure 1” in the chapter text.

### **4.1 Introduction**

Ontologies are external representations of domain knowledge created by experts through a collaborative examination process (Rector, Schulz, Rodrigues, Chute, & Solbrig, 2019). When creating ontologies, experts define an explicit and standardized common vocabulary which they use to transcribe their knowledge into a set of mappings which reflect the entities, relations, and structures of the domain. Ontology datasets are collections of software files which encode the complex objects of ontologies for use in digital environments (Priya & Kumar, 2015). Ontology datasets are increasingly being used to help the performance of challenging knowledge-based tasks. For instance, ontology datasets are being applied towards both system-facing computation tasks like information extraction on unstructured text and behavior modeling of intellectual agents, as well as an increasing number of human-facing visualization tasks like decision support systems within critical care environments (Jusoh, Awajan, & Obeid, 2020; Lytvyn, Dosyn, Vysotska, & Hryhorovych, 2020; Román-Villarán et al., 2019). Yet for domains of high complexity, such as biomedical research, environmental sciences, and medical triage, the ontology datasets can be challenging to understand. This is because their complex objects can combine to reflect countless ontology entities, relations, and any number of additional domain-specific concepts (Dessimoz & Škunca, 2017).

For ontology datasets to be used effectively, we need to understand them. When we encounter unfamiliar complex objects, we use our perception, intuition, and reasoning to form a mental model of their parts, relationships, and behaviors (Sedig, Parsons, Liang, & Morey, 2016). When encounters present us complex objects that describe a space, like distance, position, or orientation, our cognitive processes form a specific type of mental model, the cognitive map. Through theoretical and experimental work, researchers have explored how our cognitive processes organize our knowledge of spaces and our performances of cognitive activities like sensemaking, navigation, and exploration within spaces. The cognitive map framework describes formation as a set of stages. We first develop landmark knowledge of a space through the internalizing of the complex objects which describe location. Next, we develop route knowledge by building associations of the relationships which connect locations. Finally, we develop our understanding of the overall structure and layout of the space, referred to within the framework as survey knowledge (Behrens et al., 2018; Craig, Dewar, Harris, Della Sala, & Wolbers, 2016; Epstein, Patai, Julian, & Spiers,

2017). We form cognitive maps during encounters with physical spaces, like when encountering the unfamiliar districts, streets, and buildings of a new city. Yet cognitive maps can also form for spaces that we perceive as spatial, yet does not directly exist within the physical dimension, like a website and its webpages (Epstein et al., 2017; Weisberg & Newcombe, 2016). Critically, encounters with ontology datasets and the knowledge encoded within their complex objects reflect spatial qualities like location, relation, and structure and thus encapsulate the conditions for cognitive map formation.

Ontology dataset visualizations are increasingly finding use during the performance of challenging knowledge-based tasks. Yet, up until recently, the design of ontology dataset visualizations have typically only considered tasks which presume a level of expertise and experience of the ontology dataset and the ontology it describes, such as ontology management or clinical treatment interfaces (Borland, Christopherson, & Schmitt, 2019; Schlegel & Elkin, 2016). As a result, leading considerations towards the design of ontology dataset visualization have not targeted the specific problem space of learning tasks which help us in building understanding of the ontology dataset itself. Therefore, our motivation for this paper is to consider how the design of interactive visualizations of ontology datasets can promote conditions for cognitive map formation so that we can be helped in developing our understanding of the ontological space described within ontology datasets.

This paper begins with an introduction of the topics of cognitive maps, ontologies, and interactive visualization tools. We find that a wide range of theoretical and experimental disciplines have directed their efforts towards understanding the functionality of our cognitive processes and their effect on the performance of our cognitive activities. Next, we introduce the theoretical framework of the cognitive map and its application towards understanding how our brains organized knowledge of complex spaces. Then, we explore the use, creation, and limitations of ontologies, an expert-defined standardized common vocabulary describing the knowledge of a domain. The introductory content concludes with an examination of the fields of information visualization and visual analytics, discussing how designers can create visualization tools using visual representation and interaction design to support our performance of our tasks and their underlying cognitive activities.

Next, we examine existing work on cognitive load and the use of interactive visualizations to support learning tasks. Here, we find that recent studies show that there is no one specific level of cognitive load that is proper for supporting learning tasks. Instead, cognitive load is a set of extraneous, intrinsic, and germane loads which must adjust for the specific conditions of the tool and task context. Through this, we discover that interactive visualization tools are a valuable resource for learning tasks. Studies have found that if designed correctly, interactive visualization tools can be an effective environment for engaging learners. The examination of existing work concludes with an exploration of insight towards the design of visual representations and interactions to support cognitive mapping of spatial knowledge, alongside a summary of the cognitive activities performed within spaces. From these findings, we formalize a set of high-level design criteria for designing interactive visualization tools to support learning tasks through alignment of the cognitive map framework and its formation process. We then perform a review of existing tools which visualize ontology datasets. This review categorizes each tool based on their generalized subview combinations, and for each, we supply analysis of their strengths and weaknesses towards promoting cognitive map formation.

Following this, we present PRONTOVISE (PRogressive ONTOlogy VISualization Explorer), an interactive visualization tool which applies the criteria in its design to support us in understanding unfamiliar ontologies. In this, we explain the technological features of the PRONTOVISE, and describe its workflow and design within the context of our high-level design criteria. We describe PRONTOVISE, an interactive visualization tool that represents ontology datasets using a combination ‘List+Overview+Context+Details’ design. The presentation continues with a detailed description of the considerations made when designing the novel ontology dataset visual representations and interactions within each subview of PRONTOVISE. Through a usage scenario, we describe a set of encounters with the Human Phenotype Ontology (HPO) and show how the underlying design of PRONTOVISE can support the requirements for cognitive map formation. We conclude with a discussion of the implications of the generalized criteria and assess the strengths and limitations of the design of PRONTOVISE.

## **4.2 Background**

This section supplies background on the concepts and terminology used when building our design criteria and its use for the design of PRONTOVISE. First, we summarize the current understanding of the cognitive map framework as shown by theoretical and experimental work. Second, we describe the value of ontologies and their constitute parts. Finally, we explore how visual representation and interaction design can be used to create interactive visualization tools.

### **4.2.1 Cognitive Map Formation**

A long-standing yet nebulous problem space across a wide range of theoretical and experimental disciplines is that of understanding the cognitive processes which form our knowledge of spaces and help us perform our activities within them. These disciplines, like neuroscience, experimental psychology, and human-information interaction, have examined how we internalize our encounters with complex spaces and their constituent parts, and use that knowledge to perform our activities within those spaces [8,10,11]. From these examinations, understanding has been built towards how our brain states process our experiences within unfamiliar spaces, and how our memory of those spaces is encoded internally within our cognitive systems. It has been found that internal representations mapping spatial relationships do, in fact, form when navigating unfamiliar environments and that the quality of that formation is directly affected by external conditions (Craig et al., 2016). Studies have also been made on specific parts of our brain, like the hippocampus, to improve our understanding of our cognitive processes which involve space and time. From these studies, it has been found that when processing experiences, our brains leverage externally represented information which describes spatial and temporary knowledge to distill and organize that knowledge within our internal cognitive systems (Ekstrom & Ranganath, 2017).

Integrating leading experimental evidence, current understanding towards how our brains organized knowledge of complex spaces aligns with that of the theoretical work for cognitive map formation and its general coding mechanisms (Behrens et al., 2018). The cognitive map formation is a staged process which occurs over repeated encounters with external representations of a space. During these encounters, our sensory and cognitive systems process our experiences into internal representations, which promotes the formation mechanisms associated

with a cognitive map (Waller & Nadel, 2013). Under these mechanisms, cognitive map formation occurs in stages of increased fidelity, depends on the complexity of the space and the level of granularity which we want to understand it. For a space which is unfamiliar, the formation of our cognitive map begins by forming our initial understanding of high-level objects of the space. Our cognitive processes use existing mental models to begin to distinguish distinct locations of the space, and from them, process locations of importance as landmark knowledge. Landmark knowledge is used for activities involving static information of specific locations and objects with a space, like comparison and sensemaking (Fast, 2010; Kallioniemi et al., 2013). Once forming our initial level of landmark knowledge, our cognitive processes begin to form associations towards the links between locations, contextualizing their pathing and relationships, which the framework refers to as route knowledge. Route knowledge is used for transitional activities involving movement between locations and objects of a space like wayfinding and navigation (Fast, 2010; Weisberg & Newcombe, 2016). As landmark and route knowledge grows, we start to form extended associations which map the locations and objects across their relationships and paths. These associations form our survey knowledge of the overall structure and layout of a space. Survey knowledge is involved with generalizing our understanding of a space, and allows us to perform activities which require a refined level of landmark and route knowledge, such as orientation, exploration, and comparison of spaces (Fast, 2010; Sedig, Rowhani, & Liang, 2005).

#### **4.2.2 Ontologies**

Ontologies are an expert-defined standardized common vocabulary describing the knowledge of a domain. Ontologies are increasingly being used to help the performance of challenging knowledge-based tasks. This is because they provide the flexibility, extensibility, generality, and expressiveness necessary to bridge the gap between the requirements for mapping domain knowledge into forms which are generalized for effective computer-facing and human-facing use (Saleemi, Rodríguez, Lilius, & Porres, 2011). After defining an ontology, experts record the complex objects of the ontology into data files which supply standardized ontology specifications. Once generated, these data files can be packaged in a dataset, shared amongst domain stakeholders, and integrated into computation and human-facing resources to support performances of challenging domain tasks. For instance, they are being used within towards both system-facing computation tasks like information extraction on unstructured text, behavior modeling of intellectual agents, as well as an increasing number of human-facing visualization tasks like decision support systems within critical care environments (Jusoh et al., 2020; Lytvyn et al., 2020; Román-Villarán et al., 2019).

The common vocabulary of an ontology is composed of a network of complex objects produced by a systematic review of domain content (Jakus, Milutinovic, Omerović, & Tomazic, 2013; Rector et al., 2019). Experts construct this network using two types of complex objects: the ontology entity and the ontology relation, which together yield various ontology structures. Ontology entities reflect the distinct concepts within the domain, like a phenotype in a medical triage ontology, a processor in a computer architecture ontology, or a precedent in a legal ontology (Tobergte & Curtis, 2013). Ontology entities will typically encode information about their role in the vocabulary, definitions, descriptions, and contexts, as well as metadata that can be used to inform the performance of future ontology engineering tasks. Ontology relations are the links between ontology entities which express the

quality of interaction between them and towards the domain as a whole (Katifori, Torou, Vassilakis, Lepouras, & Halatsis, 2008). One of the most common types of ontology relations is that of inheritance. In this relation, the characteristics of one ontology entity act as a template to define another. For instance, an ontology entity in an animal ontology standing for the concept of a ‘dog’ may inherit from an ontology entity reflecting the concept of a ‘domesticated animal’. Typically, ontology relation types are domain-dependent and emerge out of unique interoperability between ontology entities. For instance, an animal ontology may also have an ontology entity reflecting the concept of a ‘human’, which may have the ontology relations ‘domesticates/is domesticated by’ between it and the ‘dog’ ontology entity.

When the size and complexity of a domain rises, so too does the complexity of its ontology. As a result, ontological datasets can become very large and complex, supporting countless complex objects describing ontology entities and relations. When interacting with highly complex spaces like ontologies, the limitations of human cognition can create a bottleneck in human-facing analytic workflows (Zhao, Ward, Rundensteiner, & Higgins, 2017). Therefore, a leading challenge for those who look to use ontologies is maintaining an ontology dataset which accurately describes its domain while still being useful for both computation and human-facing tasks.

### **4.2.3 Interactive Visualization Tools**

Our daily lives are permeated by encounters with external representations that connect to us through our visual perception, auditory, and other sensory systems. Designers encode their knowledge as information within their external representations, in the hopes that this knowledge can be transferred to the sensing observer. Information visualization is an area of research which concentrates on investigating the use of visual representation as an interface to our cognitive processes, and the mental representation space which they manage (Parsons & Sedig, 2014). Through theoretical and experimental work, researchers investigate strategies for designing visual representations to support of our cognitive processes (Sedig et al., 2016; Shneiderman, 1996).

We use tools to improve our ability to complete challenging tasks, both physical and cognitive. Rollerblades, hammers, and pencils are examples of tools which augment the physicality of the human body to perform difficult physical activities (i.e., dexterity, strength, speed, precision, etc.). Similarly, we can use tools like language, books, and computational devices to support the performance of activities which are cognitive in nature. We achieve cognitive augmentation through the activation of distributed cognition, where through interaction, complex cognition is offloaded from the internal processes of our mental representation space and into the external representation and computation space of our tools (J. Davies & Michaelian, 2016). By offloading complex cognition to tools designed to support complex cognitive activities, this distribution of cognitive responsibility allows us to direct our mental concentration towards other activities which are more aligned with our natural cognitive abilities (Pereira Rocha, de Paula, & Sirihal Duarte, 2016).

These days, designers take advantage of readily available technologies like high resolution monitors, standardized operating systems, and internet services to produce powerful cognitive tools. Research spaces like visual analytics, which concentrate on the using of visual representation to support analytic reasoning, are using these technologies to design visualization tools that support the performance of our complex cognitive activities. For

instance, visualization tools have been used for sensemaking activities towards misinformation within the medical domain, search activities on large document sets, and decision-making activities using health data (Demelo, Parsons, & Sedig, 2017; Ninkov & Sedig, 2019; Ola & Sedig, 2016). Providing the opportunity for interaction with visual representations allows us to become an active participant in our encounters with encoded information. That is, by integrating interactive components within a visualization tool, designers can formulate a dynamic and evolving dialectic between us and the encoded information. Interactive visualization tools allow us to perform actions onto the interface based on our perception of encoded information. Based on an action event, a tool can ingest that action into its internal logic, move into the computation space, formulate potential responses, and then adjust its interface in a way in which we can perceive. These three stages: perception, action, and tool reaction, form an interaction loop, which can be explored by researchers to establish generalized patterns, frameworks, and methodologies which better support our needs as we perform complex cognitive activities (Parsons & Didandeh, 2015; Sedig & Parsons, 2013).

## **4.3 Methods**

In this section, we describe the methods used for formulating a set of high-level design criteria for designing interactive visualizations of ontology datasets which support the performance of activities which promoting cognitive map formation. We begin with related work about cognitive load during complex learning and the use of interactive visualizations to support complex learning. Based on these findings, we outline a set of criteria for designing interactive visualization tools for complex learning by supporting the stages of cognitive map formation. We review existing ontology dataset visualization tools and analyze how they align or mis-align with the conditions that support the cognitive activities performed during complex learning and their promotion of cognitive map formation for unfamiliar ontologies.

### **4.3.1 Related Work**

Cognitive load theory, a framework for understanding the functional interplay between working and long-term memory, describes that our working memory can be understood as a cognitive load put onto us that forms out of the complexity of a learning task (Mayer, 2014). Within the framework, cognitive load is explained as a combination of three loads: intrinsic load as the mental effort associated with task performance, germane load as the mental effort required for processing an encounter for conversion into long-term memory, and extraneous load as the task-irrelevant activities resulting from poor encounter design (Seufert, 2018).

Recent work has targeted the challenging dynamics of cognitive load and its impact on learning. A study by Wang et al. explored the impact of cognitive load and affordance design on the performance of learning tasks using collaborative tools. Within their three-cohort study, three unique interfaces were prepared and assigned to a cohort to support the performance of the same learning task. For the three interfaces, one was noninteractive video-based, one noninteractive text-based, and one providing an interactive interface integrating various multimedia. It was found that the noninteractive video-based interface cohort expressed significant overloaded working memory and performed poorly in their post-task scoring assessment. The text-based interface cohort expressed that they experienced low cognitive loads for working memory and performed adequately in their assessment scores. Yet, they

found that the interactive multimedia group produced the highest assessment scores of any cohort, even though they expressed a moderate level of cognitive load (C. Wang, Fang, & Gu, 2020). In addition, recent research efforts by Seufert explored the problem space, targeting the performance of self-regulated learning tasks, and arrived at similar conclusions (Seufert, 2018). As such, neither high nor low cognitive load is definitively correlated with the conditions for effective learning. Instead, leading guidance prescribes that proper cognitive load can vary task to task. Therefore, when creating learning environments, designers should: Take care to minimize extraneous load which is unrelated to the learning task, direct intrinsic load towards supporting the specific cognitive activities of the learning task, and unify affordances to best align the information, learning process, and the learner towards maximizing germane load for converting working to long-term memory (Mutlu-Bayraktar, Cosgun, & Altan, 2019). Yet, this care is not often observed within the design of interactive visualization tools which support learning tasks involving ontology datasets. This will be examined in depth within our review of existing tools.

Interactive visualization tools can be a valuable resource for learning tasks. We gain a deeper level of understanding when performing learning tasks when we engage mixtures of deeply textured information formats within a flexible learning environment (C. Wang et al., 2020). When we learn, we seek to move beyond our prior knowledge and into the unfamiliar through cognitive engagement (Mayer, 2014). For this, it is critical to consider creative thinking and the underlying processes of divergent thinking, which is the generation of ideas, and convergent thinking, the evaluation of ideas. A two-cohort study was performed by Sun et al. which asked each cohort to perform the same learning task involving divergent thinking. Specifically, one cohort was provided an online system without any assistive support, yet the other was provided an interactive visualization tool to support cognitive mapping during task performance. The results from the study directly exhibited that members of the cognitive mapping resource cohort had an improved task performance over their corresponding non-resource cohort members. It was concluded that members of the cohort were able to manage their working memory through a moderation of cognitive load during cognitive mapping (Sun, Wang, & Wegerif, 2019). As such, tasks which involve creative thinking can be improved through the use of interactive visualization tools. This can be achieved by aligning with the requirements for cognitive mapping and its underlying cognitive activities like association, decomposition, combination, and adjustment during divergent thinking, and selection and evaluation for convergent thinking. This is especially important in self-regulated learning environments with interactive visualization tools, where we must guide our own learning tasks through the setting of goals, the planning of our learning process, enacting our process by using our resources to interact with new information, and evaluating our learning achievements (Seufert, 2018). We, however, find that the requirements for supporting cognitive mapping are not accounted for within the design of interactive visualization tools which support learning tasks involving ontology datasets. This will be examined in depth within our review of existing tools.

Visualization can improve our capacity to encounter new information, yet poorly designed visual representation and interaction can also harm learning and the performance of its necessary cognitive activities (Yalçın, Elmqvist, & Bederson, 2016). This is still true for interactive visualizations of ontology datasets and their ontological spaces. Studies have shown that the inclusion of supplementation information describing a space in visual representations not only affects how our cognitive processes handle new information, such as with memorization and

decision-making, but also provides meaningful improvements towards the performance of cognitive activities during cognitive mapping within learning tasks (Ragan, Bowman, & Huber, 2012; Sun et al., 2019). It is important that designers account for the way novel information is processed by learners when designing their visualizations, and be cognizant towards how specific design strategies can facilitate conditions for effective learning (Mutlu-Bayraktar et al., 2019). We summarize, in Table 1, the cognitive activities performed within spaces, expounding their relationship to divergent and convergent thinking, and the types of spatial knowledge required for their performance within the framework of cognitive map formation.

**Table 4-1 Summary of the cognitive activities performed within spaces. Included is the name and description, the underlying processes of creative thinking which relate to the cognitive activity, and the types of spatial knowledge which must be developed within a cognitive map of a space before the activity can be performed within that space.**

| <b>Name</b> | <b>Description</b>   | <b>Related Thinking Processes</b> | <b>Required Spatial Knowledge</b> |
|-------------|--|-----------------------------------|-----------------------------------|
| Sensemaking | Reasoning and the mental manipulation of representations to develop, build upon, and refine mental models (Sedig et al., 2016).      | Convergent                        | None                              |
| Navigation  | Observing, orientating, and decision-making for directed movement towards a known objective [4,11,31].                               | Convergent                        | Landmark, Route                   |
| Exploration | Observing, orientating, and decision-making for undirected movement without an objective (Lytvyn et al., 2020; Yalçın et al., 2016). | Divergent, Convergent             | None                              |
| Search      | Observing, orientating, and decision-making for directed movement towards an unknown objective (Sedig & Parsons, 2013).              | Divergent, Convergent             | Landmark, Route, Survey           |
| Wayfinding  | Constructing and memorizing movement sequences for future objective-oriented activities [16,39,40].                                  | Divergent, Convergent             | Landmark, Route, Survey           |

The visualization of ontology datasets is an active problem with an expansive set of research themes. New publications are consistently taking the creation, activation, and visualization of ontology datasets into novel and varied directions. Specifically, a literature review performed by Pesquita et al. highlights the range of discussion towards semantic web research. They describe the two leading challenges for supporting semantic web tasks. The first is the challenging of support users of varying levels of expertise. The second is the challenging of generalizing findings across different task contexts, such as different types of information within datasets and what the task wants to do with that information. Additionally, they note that there is a shortfall of research directed towards understanding the performance of open-ended tasks using semantic web visualizations, when considering the users, information, and task context (Pesquita, Ivanova, Lohmann, & Lambrix, 2018).

### 4.3.2 Task Analysis

We find five high-level criteria for designing interactive visualization tools of ontology datasets that promote the stages of cognitive map formation for learning tasks. They are provided in Table 2.

**Table 4-2 The high-level criteria for designing interactive visualization tools of ontology datasets which promote the stages of cognitive map formation for learning tasks.**

| Criteria  | Description  |
|---|--|
| Provide generalized support for ontology datasets                                 | Designs should provide a generalized environment which facilitate the loading of ontology datasets of any size under the guidance of existing ontology file specifications. This is so that we may build our understanding of ontology datasets which are relevant to our challenging knowledge-based tasks.   |
| Tune cognitive load to specific needs   | Designs should provide a cognitive load which is aligned with the conditions for an effective learning environment for ontology datasets. Specifically, extraneous load which is unrelated to the learning task should be minimized, intrinsic load should be tuned to support the specific cognitive activities of the learning task, and germane load should provide affordances which unify the needs of the learner, space, and chosen process for learning. |
| Afford the spatial knowledge within ontological space                             | Designs should supply encounters which afford to us an authentic internal encoding of the entities, relations, and structures of the ontology dataset to support our development of spatial knowledge for the formation of our cognitive maps.   |
| Facilitate the performance of the cognitive activities necessary to learn a space | Designs should provide encounters which allow us to perform the cognitive activities necessary to build understanding of a space. This is because not supporting any one of sensemaking, navigation, exploration, wayfinding, and search would lessen our ability to leverage our various cognitive processes and hamper the stages of cognitive map formation.  |
| Support self-regulated learning   | Designs should provide encounters which allow us to guide our own learning tasks: through setting goals, planning our learning process, enacting our process by using our resources to interact with new information, and evaluating our learning achievements.  |

### 4.3.3 Existing Tool Review

We consider prior survey work by Katifori et al., since updated by Dudáš et al., which provides a high-level collection of design strategies for visualizing ontologies and assist in the record keeping of active tools (Dudáš, Lohmann, Svátek, & Pavlov, 2018; Katifori, Halatsis, Lepouras, Vassilakis, & Giannopoulou, 2007). Additionally, we consider recent work by Po et al. which provides a thorough investigation of linked data visualization with dedicated portions directed towards ontology visualization tools (Po, Bikakis, Desimoni, & Papastefanatos, 2020). These resources aid in our determination towards our coverage of existing tools within in our examination, based on three conditions: The tool is currently accessible, is still in a working state, and both loads and represents ontologies of any size. We require the tool to be accessible, as they must still be available for our examination. This eliminates tools like GrOWL and OntoTriX, which are no longer accessible. We require the tool to be in a working state, as it would be unfair to

assess a tool that can no longer fulfill its functional purpose in the manner it was intended. This condition eliminates tools like OntoViz and OntoSphere, which are still accessible, but are no longer supported in their original Protégé suite environment. The final condition specifies that the tool must load and represent ontology of any size, as our scope is of a generalized design for visualizing ontology datasets of all sizes. This condition removes a tool like SOVA, which, while accessible and working, cannot load large ontology datasets. Using these criteria, we filter from the full set of ontology dataset visualization tools constructed by Dudáš et al., to produce the following list of ontology visualization tools: Protégé Entity Browser, Protégé OntoGraf, Ontodia (now maintained under the name Metaphactory), OntoStudio, WebVOWL, and TopBraid Explorer (Falconer, 2010; Lohmann, Negru, Haag, & Ertl, 2016; Mouromtsev et al., 2015; Musen & Team, 2015; Semafora Systems, 2020; TopQuadrant, 2020). Additionally, we add consideration towards OntoViewer, a demonstrative tool from a recent publication by Silva et al. (Silva, Santucci, & Freitas, 2019). Furthermore, we consider WebProtégé Entity Graph, built within the latest edition of the Protégé software suite (Tudorache, Nyulas, Noy, & Musen, 2013). This review is a targeted review of existing tools and their underlying designs towards supporting complex learning and their promotion of the stages of cognitive map formation. Within this review, we categorize the tools based on their included subview types. Table 3 provides a description of each subview type.


**Table 4-3 The types of subviews within an ontology dataset visualization interface.**

| <b>Type</b> | <b>Description</b>  | <b>Typical Implementation Strategy</b>   | <b>Cognitive Activities</b>                              | <b>Use in Review Tools</b>   |
|-------------|---|--|--|--|
| List        | A subview that depicts components of the ontology datasets like entities and relations within a list. | A text-based visual representation strategy with interactions for selection and management.      | Sensemaking, Navigation, Exploration, Search, Wayfinding | Protégé Entity Browser, Protégé OntoGraf, Ontodia OntoStudio, TopBraid Explorer, WebProtégé Entity Graph, OntoViewer |
| Overview    | A subview that depicts the full contents of an ontology dataset.                                      | A pictorial-based visual representation strategy with interactions for selection and filtering.  | Sensemaking, Navigation, Exploration, Search, Wayfinding | WebVOWL, Ontodia, OntoViewer   |
| Context     | A subview that depicts a subset of the ontology dataset contents determined through interaction.      | A pictorial-based visual representation strategy with interactions for selection and comparison. | Sensemaking, Exploration, Wayfinding                     | Protégé OntoGraf, OntoStudio, TopBraid Explorer, WebProtégé Entity Graph, OntoViewer                                 |
| Details     | A subview that depicts the information of a specific object within the ontology dataset.              | A text-based visual representation strategy with minimal opportunities for interaction.          | Sensemaking  | WebVOWL, Ontodia OntoStudio, TopBraid Explorer, WebProtégé Entity Graph, OntoViewer                                  |

### ***4.3.3.1 List+Details Designs***

Protégé Entity Browser is an interactive visualization tool which uses a legacy version of Protégé software suite. It represents ontology datasets using a combination ‘List+Details’ design, as depicted in Figure 1 (Musen & Team, 2015). The system has two subviews, a list and a details subview. The visual space of the list subview maintains a tree-like list of either entities or relations with standard expand-collapse interactions. When an interaction is made on an entity label, the details subview to the right of the list is shown. When this occurs, the information associated with the selected ontology entity is represented in the details subview, accompanied by buttons which allow for various creation, edit, and removal interactions. If any of the text-based labels refers to an alternative entity within the ontology, selecting it will change the details view to show the information of that entity.

An advantage of the design of Protégé Entity Browser is that it supplies encounters which do not depict any novel visual representations or interactions. Little to no training is needed, as we can apply intuition from mental models of standard text-based interfaces. A disadvantage of Protégé Entity Browser is that its list subview does not scale well to ontologies with high numbers of complex objects, as only a limited number can be represented before going ‘off the screen’. To address this, collapsing interactions are provided; however, this reduces the opportunity for encounters with large sections of the ontological space that may be relevant to activities performed during learning like navigation, exploration, and search. Furthermore, the design represents only one of the entities or relations at a time. This reduces our ability to perform sensemaking on entities, forming landmark knowledge, and relations, which form route knowledge. Additionally, this design consideration harms our ability to build strong associations between entities using their shared relations which form the structure and layout for activities which require survey knowledge.




**Figure 4-1 Protégé Entity Browser: List and details subviews loaded with the Human Phenotype Ontology.** Source: Image with permission courtesy of Center for Biomedical Informatics Research, Stanford University School of Medicine, Protégé Team, <https://protege.stanford.edu>.

#### 4.3.3.2 List+Context Designs

Protégé OntoGraf is an interactive visualization tool that also uses a legacy version of the Protégé software suite to visualize ontology datasets using a combination 'List+Context' design, as depicted in Figure 2 (Falconer, 2010). The system supplies two subviews, a list and a context subview. The list subview of Protégé OntoGraf uses the same representation and interaction design as the earlier Protégé Entity Browser. While Protégé Entity Browser provided a text-based details subview, Protégé OntoGraf instead provides a subview which supports visualizations with interactions that depict representations of ontology entities and relations. As a representation of an ontology entity is interacted with in the list subview, it appears in the context subview, encoded with interactions which adjust it in various ways. A double-clicking interaction will request the tool to center the selected entity and its relations. A right-clicking interaction allows for the use of some entity-specific actions, like generating its full network of ontology relations. Additionally, a hold-and-drag interaction allows for entities to move within the representation space, and a zoom interaction can make visual representations larger in their display. Protégé OntoGraf also includes a text-based search, though it does not supply any autocomplete or suggestive capabilities.

The advantages and disadvantages of the list subview is shared with the earlier Protégé Entity Browser. However, unlike the Protégé Entity Browser, Protégé OntoGraf does not allow for the list to represent ontology relations. As such, it can only provide encounters with interactions directed towards ontology entities, reducing opportunities for sensemaking activities, and, in turn, the development of non-landmark knowledge within the list

subview. A strong advantage for Protégé OntoGraf is its ability to support initial landmark and route knowledge development within its context subview. When an ontology entity is provided in the context subview, we can directly interact with it as an object in space. This allows us to encounter inheritance relationships backwards, as well as open multiple entities at the same time, allowing for navigation, exploration, and even more challenging activities like wayfinding. However, a disadvantage of Protégé OntoGraf's design is that it is very hard to establish a detailed understanding of any one specific entity, as there is no way to encounter information of ontology entities beyond their text label, as was available in the Protégé Entity Browser. Additionally, the size of a box generated to represent an ontology entity is based on the length of its text-based label. At best, this supplies little to no value during sensemaking activities, but, at worst, may mislead us while we try to understand which ontology entities are important locations within the space, promoting poor landmark knowledge. Furthermore, it is not possible within the context subview to navigate to ontology entities which inherit from a target entity. Finally, the context subview does not scale well to large ontologies. If the generation of any sizable number of ontology entities and relations is requested, the display must be zoomed out to such an extreme point that all individual clarity is lost, reducing the effectiveness in activities like search.



**Figure 4-2 Protégé OntoGraf: List and context subviews loaded with the Human Phenotype Ontology. Source: Image with permission courtesy of Center for Biomedical Informatics Research, Stanford University School of Medicine, Protégé Team, <https://protege.stanford.edu>.**

#### 4.3.3.3 Overview+Details Designs

WebVOWL is an interactive visualization tool that visualizes ontology datasets using a combination 'Overview+Details' design, as depicted in Figure 3 (Falconer, 2010). The system supplies two subviews, an overview and a details subview. The details subview of WebVOWL uses a similar representation and interaction design as the

earlier Protégé Entity Browser, where information is supplied for ontology entities and relations selected within alternate subviews. Unlike previously discussed tools, WebVOWL maintains an overview subview instead of a list subview. That is, when WebVOWL loads, it applies a similar representation and interaction design for its ontology entities and relations to that of OntoGraf’s context subview, except that the ontology is visualized in full. This adjustment removes the need for users to specifically target ontology entities or relations from a list subview before they are represented in alternate subviews, immediately allowing users to encounter ontology structure. As an entity or relation is interacted with in the overview subview, its information is presented within the details subview. A hold-and-drag interaction allows for entities to move within the representation space, and a zoom interaction can make visual representations larger in their display.

The advantages and disadvantages of the details subview of the design is shared with the earlier Protégé Entity Browser. A strong advantage for WebVOWL is its ability to fully support the performance of many cognitive activities like search, navigation, and exploration within its overview subview, thereby supporting the development of all stages of cognitive map formation. When an ontology dataset loads, we can directly interact with it as an object in space. However, a disadvantage of the design of WebVOWL is that it is very hard to target cognitive activities on specific entities and relations, as the tool forces an overview representation with a very low quality of interactivity towards divergent thinking. Finally, the overview subview does not scale well to large ontologies, because if the generation of any sizable number of ontology entities and relations is requested, the display must be zoomed out to such an extreme point that all individual clarity is lost.





Figure 4-3 WebVOWL: Context and details subviews loaded with supplied Friend of a Friend Ontology. Source: Image generated from WebVOWL: Web-based Visualization of Ontologies resource <http://vowl.visualdataweb.org/webvowl.html>.

#### ***4.3.3.4 List+Overview+Details Design***

Ontodia is an interactive visualization tool which represents ontology datasets using a combination ‘List+Overview+Details’ design, as depicted in Figure 4 (Mouromtsev et al., 2015). Like previously discussed systems, Ontodia maintains a list subview using the same expand-collapse representation and interaction encodings. Ontodia improves on this by including a text-based search with auto-complete and suggestive features. Ontodia provides a context subview populate based on interactions with the list subview. Ontodia follows the model of Protégé OntoGraf yet differentiates by annotating added high-level information. That is, it supplies a summarized label of inherited ontology entities, an assigned color, and imagery of the entity, if included, and support for all ontology relation types, each relation including a text-based label. Ontodia does supply new interactions for ontology entities, like access to a details subview, a removal interaction for ontology entities, and filters the list subview based on a specific ontology entity. More so, Ontodia represents a high-level overview of the contents of the context subview, which simplifies the active ontology entities based on position in space, their current size, and assigned color. Attached to this is a basic panning interaction. Finally, Ontodia supplies a details subview, which, upon request, will show in text the information content of the selected entity or relation, just as was provided in Protégé Entity Explorer.

The advantages of the list subview is shared with the earlier systems. Yet, Ontodia goes above and beyond with the inclusion of a panel which keeps track of active ontology entities and relations. This addition allows users to reference in text the current state of the view and quickly associate listed entities and relations, supporting sensemaking and orientation activities. The tool aligns with the formation of route knowledge through its explicit labeling of ontology relations within the context subview. Finally, Ontodia includes access to an ontology details subview that was not provided in the prior Protégé OntoGraf. This allows users to have encounters with information during the performance of the more complex activities associated with the later stages of cognitive map formation which receive help from the use of more familiar visual representations. On the side of disadvantages, Ontodia struggles to provide visual representations of large ontologies, as just as Protégé OntoGraf, the only solution for providing a wide view of the overall ontology is to zoom out the context subview at the expense of clarity. Ontodia attempts to provide some solution to this issue with the high-level overview of the contents of the context subview. However, this inclusion does not represent any new information which would support the expansion of survey knowledge, nor include interactions that provide support for any added activities within the space that can support our learning tasks.




**Figure 4-4 Ontodia (now contained within the Metaphactory software suite): List, context, and details subviews.**  
 Source: Image with permission courtesy of Metaphacts, metaphacts.com.

#### 4.3.3.5. List+Context+Details Designs

OntoStudio, TopBraid Explorer, and WebProtégé Entity Graph are three interactive visualization tools which represent ontology datasets using a combination ‘List+Context+Details’ design (Semafora Systems, 2020; TopQuadrant, 2020; Tudorache et al., 2013) While each of these tools include distinctive qualities, in general, they all have a similar high-level design, as seen with WebProtégé Entity Graph in Figure 5. Like previously discussed systems, each of the three tools support a list subview using the same expand-collapse representation and interaction encodings. WebProtégé improve on this by including a text-based search with auto-complete and suggestive features. These three tools provide a context subview populate based on interactions with the list subview. Each system has slight variations in their representation strategy, but largely follow the model of Protégé OntoGraf. All three build upon Protégé OntoGraf by providing support for all ontology relation types, and with each relation, including a text-based label. However, like Protégé OntoGraf, WebProtégé Entity Graph, OntoStudio and TopBraid Explorer do not supply interactions which support any expansion of inheritance, centering, or sorting. Finally, all three systems supply a details subview, which, upon request, will show in text the information content of the selected entity or relation, just as was provided in Protégé Entity Explorer.

The advantages of the list subview within these systems is shared with the earlier systems. All three tools better align with the formation of route knowledge through their explicit labeling of ontology relations. Additionally, the tools include access to an ontology details subview that were not provided in the prior Protégé OntoGraf. This

allows users to have encounters with information during the performance of the more complex activities associated with the later stages of cognitive map formation which receive help from the use of more familiar visual representations. On the side of disadvantages, all three systems struggle to supply visual representations of large ontologies. Like Protégé OntoGraf, the only solution for supplying a wide view of the overall ontology is to zoom out the context subview at the expense of clarity. This harms the performance of cognitive activities within spaces like search, wayfinding, and exploration.



**Figure 4-5 WebProtégé Entity Graph: List and context subviews.** Source: Image with permission courtesy of Center for Biomedical Informatics Research, Stanford University School of Medicine, Protégé Team, <https://protege.stanford.edu>.

#### 4.3.3.6 List+Overview+Context+Details Design

OntoViewer is an interactive visualization tool that visualizes ontology datasets using a combination ‘List+Overview+Context+Details’ design, as partially depicted in Figure 6 (Silva et al., 2019). As with earlier systems, OntoViewer maintains a list subview using similar expand-collapse representation and interaction. Like Ontodia, this list includes a text-based search which supplies filter interactions over the full ontology. However, unlike Ontodia, the list subview does not support both a list of the full ontology and a list of all active ontology entities at the same time. Instead, only a single list is supplied which filters down from the full set to a filtered set. OntoViewer improves on its list subview by encoding more information about the ontology relations associated with each ontology entity. OntoViewer includes a dedicated overview subview which represents the full network of ontology entities and relations, and changes based on the current selections within alternate subviews. The overview represents a node-link radial tree which maps the network of the ontology entities and relations out from the root ontology entity. Yet, this overview is limited to two ‘steps’ of ontology relations out from the root ontology, nor does it maintain any interaction on it. OntoViewer maintains a context subview which supplies novel visual representations

and interactions for ontology entities, relations, and structure. OntoViewer supplies a stacked visual space which allows users to select one of three dedicated context subviews to highlight the qualities of ontology entities, relations, structure, and instances. For relations, OntoViewer attempts to improve on the node-link graph representation, as seen in the earlier systems, into what they referred to as 2.5-dimensional space. That is, it represents its network of ontology entities and relations in a radial distribution, just as was done in the overview, except at a perspective which mimics a three-dimensional plane within the two-dimensional display. Concerns with encoding overlap are addressed with interactions which shift the perceived perspective of the representation in various directions and orientations. For ontology entities, OntoViewer provides a dedicated context subview with an icicle tree diagram representing a selected ontology entity, the ontology entities it inherited from, and the ontology entities which inherited from it. Finally, for structure, OntoViewer supplies a dedicated context subview which presents a bar chart describing the spatial calculations for a set of ontology entities like distance between entities, number of relations, and their distance from the root ontology entities.

OntoViewer provides numerous advantages for supporting the stages of cognitive map formation. The advantages and disadvantages of the list subview largely aligns with the advantages of previous systems. However, by improving by including encodings within its list that signal when ontology relations are assigned to an entity ontology, encounters with entities can provide opportunities for initial comparison and sensemaking and can guide navigation activities which develop route knowledge. OntoViewer adds an overview which shows a high-level abstraction of the ontological space. This abstraction can aid us during encounters to orient our activities within the space, act as a wayfinding resource towards further encounters, and suggest structural patterns which can help develop survey knowledge. However, the representation strategy used within the overview subview does not represent the full ontology mapped within the dataset, but instead limits its representation of the network ontology entities and relations at a certain distance away from the root ontology entity. By providing a representation which does not depict a complete mapping of ontology structure, nor provide encodings which clearly afford the existence of obscured ontological features, the potential for bad encounters with ontological space rises. These encounters can lead to misunderstandings towards the information which describes the space and may lead us to misalign our development of spatial knowledge. In its context subviews, OntoViewer advances past other systems in its support of cognitive map formation by appointing three dedicated subviews for ontology entity, relation, and structure. By splitting concerns, each subview can supply encounters that best demonstrate the unique qualities of each form of spatial knowledge. A disadvantage with the specific implementation of the context subviews of OntoViewer is that they share a stacked visual space, where subviews are occluded when not selected. This breaks with the value of distributed presentation, as it is reducing our ability to receive feedback when performing activities which should afford a coupling between multiple types of spatial knowledge.




Figure 4-6 OntoViewer: List, overview, and relations context subviews. Source: Image with permission courtesy of “Visualization and analysis of schema and instances of ontologies for improving user tasks and knowledge discovery”, School of Informatics, UniRitter Laureate International Universities.

## 4.4 Materials


In this section, we describe the materials of our generalized design of PRONTOVISE and its implementation. We begin with an outline of the technologies used within the PRONTOVISE implementation. Next, we present a high-level look at the workflow of PRONTOVISE and supply a general overview of its design details.

### 4.4.1 PRONTOVISE Technologies

We developed PRONTOVISE as a generalized web-based tool which allows for the uploading of correctly formatted OWL RDF ontology data resources, either individually, or within a .zip compression file. The tool processes the uploaded files and indexes its contents for use. We have created the front end of PRONTOVISE using the latest HTML5, CSS, and JavaScript technologies, allowing for cross-browser (Firefox, Chrome, Opera) and cross-platform support. Its back-end technology is also developed with JavaScript. We use the Lunr.js JavaScript service as our ontology entity indexer and search engine (Nightingale, 2017). We used the D3.js JavaScript visualization library to create the visualization and interaction experiences found throughout PRONTOVISE (Mike Bostock, 2016).

### 4.4.2 PRONTOVISE Workflow and Design

PRONTOVISE maintains several systems within its workflow. We will now briefly describe each of their designs in the context of their workflow, as depicted in Figure 7, and highlight their satisfaction of the criteria for designing interactive visualization tools which support complex learning and the stages of cognitive map formation.



**Figure 4-7 Depiction of the workflow of PRONTOVISE (PRogressive ONTOLOGY VISualization Explorer). Yellow boxes represent the processes performed within the back-end computation system. Blue boxes represent the object types which are persisted within browser storage. Orange boxes represent the various subviews within the front-end visualization system. The green box represents the types of low-level interactions which can be made to the system.**

The workflow of PRONTOVISE begins first by one loading the PRONTOVISE web application using their computer and browser of choice. PRONTOVISE presents a starting page which asks to upload a valid OWL RDF ontology file. When an upload interaction is performed, the ‘back-end’ ontology processing system uses several technologies to read, validate, and initiate the processing of the data encoded in the uploaded file. Using third-part resources like the Lunr.js library, the data of the ontology file is analyzed, indexed, and stored within browser memory. This temporary storage allows our visualization software to access ontology content and structure, as well as provide an

index for text-based search functionality. As we do not consider our work in the ‘back-end’ of PRONTOVISE a novel pursuit, no further description will be directed towards this system.

Once the ontology processing subsystem has completed its work, the ‘front-end’ system which performs the visualization of the ontology dataset starts. The system accesses the stored ontology data, analyzes it, and directs that data to each of the subviews which are shown across the available visual space.

PRONTOVISE is an interactive visualization tool that represents ontology datasets using a combination ‘List+Overview+Context+Details’ design. This combination aligns with the OntoViewer interactive visualization tool, described in the Methods sections within our review of existing tools. Within the review, two concerns were presented towards the implementation of the ‘List+Overview+Context+Details’ design within OntoViewer. To recall, the examination of the overview subview within OntoViewer was that it did not depict a complete mapping of ontology structure, nor supply encodings which clearly afford the existence of obscured ontological features. Thereby, it afforded in such a way that could result in bad encounters with ontological space. Additionally, our examination of its context subview highlighted a concern with its choice to restrict each context subview within a shared visual space, where subviews are occluded when not selected. We stated that this breaks with the value of distributed presentation because it reduces opportunity for feedback when performing activities which should afford a coupling between multiple types of spatial knowledge. PRONTOVISE differentiates from OntoViewer by facilitating improvements on the visual representation and representation designs for each of ‘List+Overview+Context+Details’. Concentration is directed towards improving the quality of affordances within the overview and addressing the concerns which arise due to stacked visual spaces using a distributed series of context subviews. We summarize, in Table 4, the high-level criteria within PRONTOVISE, which can be used for designing interactive visualization tools which support complex learning and the stages of cognitive map formation.

**Table 4-4 A summary of the high-level criteria within PRONTOVISE, which can be used for designing interactive visualization tools which support complex learning and the stages of cognitive map formation. The satisfaction of these criteria at the implementation level are discussed in detail later within the workflow.**

| <b>Criteria</b>                                   | <b>PRONTOVISE</b>  | <b>Related Systems/Views</b>                       |
|---|--|--|
| Provide generalized support for ontology datasets | PRONTOVISE provides a generalized environment which supports the loading of ontology datasets of any size and from any domain when they fulfill the requirements of OWL RDF, the leading ontology dataset format. Additionally, its visual representation and interaction designs are built to scale for any number of encoded complex objects.  | Ontology processing system; all front-end subviews |
| Tune cognitive load to specific needs             | Cognitive load is actively considered within the design of PRONTOVISE. PRONTOVISE is designed to be a complex learning environment, so design features which produce extraneous load unrelated to learning tasks are minimized. PRONTOVISE provides a level intrinsic load which targets a promotion of the stages of cognitive map formation. PRONTOVISE accounts for germane load by specifically being designed to provide a learning environment for those who are unfamiliar with an ontology dataset. This is achieved through | All front-end subviews                             |

visualizations which address the specific spatial knowledge of the various complex objects within ontology datasets.

Afford the spatial knowledge within ontological space

PRONTOVISE includes numerous subviews which provide encounters that afford perspectives of authentic internal encodings of the entities, relations, and structures of the ontology dataset.

Various front-end subviews

Facilitate the performance of the cognitive activities necessary to learn a spacethinking processes and the stages of cognitive map formation.

PRONTOVISE facilitates the performance of sensemaking, navigation, exploration, wayfinding, and search within ontological space over numerous subviews to support our necessary to learn a spacethinking processes and the stages of cognitive map formation.

Various front-end subviews

Support self-regulated learning

The design of PRONTOVISE includes a modular set of subviews which support nonlinear interaction loops, which together provide the freedom to set, plan, enact, and evaluate any set of learning tasks for ontological space, all while following the requirements for cognitive map formation.

Ontology processing system; all front-end subviews

There are seven subviews within PRONTOVISE: Search and Pinning Panel, Ontology Sections Panel, Section Levels Panel, Level Landmark Entities Panel, Entity Network Panel, Path Explorer Panel, and Entity Details Panel. They are presented together in Figure 8. The full set of subviews remain context aware of their neighboring subviews and manage their internal logic to align with the user as they move between each. We summarize each subview in relation to our task analysis in Table 5, followed by discussion for each subview.




Figure 4-8 An overall view of the PRONTOVISE ontology visualization system, which has seven subviews: Search and Pinning Panel subview (a); Ontology Sections Panel subview (b); Section Levels Panel subview (c); Level Landmark Entities Panel subview (d); Entity Network Panel subview (e); Path Explorer Panel subview (f); and Entity Details Panel subview (g).

**Table 4-5 A summary of the subviews of PRONTOVISE, describing their subview type, supported cognitive activities, and their relationship to the stages of cognitive map formation.**

| <b>Subview</b>                | <b>Type of Subview</b> | <b>Cognitive Activities</b>                                 | <b>Spatial Knowledge</b>   |
|-------------------------------|------------------------|---|----------------------------|
| Search and Pinning Panel      | List                   | Sensemaking, Navigation,<br>Search, Wayfinding              | Landmark                   |
| Ontology Sections Panel       | Overview               | Sensemaking, Navigation,<br>Exploration, Search, Wayfinding | Landmark, Survey           |
| Section Levels Panel          | Context                | Sensemaking, Exploration, Search,<br>Wayfinding             | Landmark, Route,<br>Survey |
| Level Landmark Entities Panel | Context                | Sensemaking, Navigation, Exploration,<br>Wayfinding         | Landmark, Route            |
| Entity Network Panel          | Context                | Sensemaking, Navigation, Exploration,<br>Wayfinding         | Landmark, Route            |
| Path Explorer Panel           | Overview               | Sensemaking, Navigation, Exploration,<br>Wayfinding         | Route, Survey              |
| Entity Details Panel          | Details                | Sensemaking   | Landmark                   |

#### ***4.4.2.1 Search and Pinning Panel***

Search and Pinning Panel, found to the furthest left of PRONTOVISE, maintains a visual space which stacks two ‘list’ subviews called Search and Pinned. These two subviews will now be discussed in more detail.

#### **Ontology Entity Search**

We have designed Ontology Entity Search to support text-based search of ontology entities within PRONTOVISE. This interaction is critical to cognitive map formation as it allows us to direct our encounters for self-regulated learning by using our existing understanding of the ontology dataset. We are presented with a search input field where we can type to perform search activities. During search activities, we are also provided with a type-ahead system that suggests possible ontology entities related to our current input and an interaction to Pin the suggestion. After a search is performed, the ontology entities contained within the result list are placed into the Ontology Sections Panel subview. This helps us perform sensemaking activities within the space, orient the position of the entity within the ontology, and to begin activities like wayfinding, navigation, and exploration. When selected, the Pin button found within each result item adds the chosen entity into PRONTOVISE’s pinning system. We reflect this by changing the Pin button into an Unpin button, as well as by assigning a unique color to that entity wherever it is found in PRONTOVISE to support wayfinding activities using that ontology entity. These considerations are depicted in Figure 9.




Figure 4-9 A depiction of Ontology Entity Search, showing a search activity which has resulted in three ‘Pin’ interactions.

### Ontology Entity Pinning

We have designed the Ontology Entity Pinning to support the management of ontology entities within PRONTOVISE. This feature is critical to cognitive map formation as it allows us a dedicated interface to manage the selection interactions we have made to initiate or continue our various cognitive activities during self-regulated learning. We initialize Ontology Entity Pinning with an empty pinned list which fills as users add entities into the pinning system through Ontology Entity Search. These color assignments are used as cross-subview encodings which support the performance of cognitive activities across PRONTOVISE, as seen in Figure 10.





Figure 4-10 A depiction of Ontology Entity Pinning, where pinned entities are represented in the same fashion as they were found in Ontology Entity Search, with a name label, an Unpin button, and a unique color. We have included a button located at the topmost position of Ontology Entity Pinning labeled ‘Remove All Pinned Entities’. When clicked, this button removes all pinned landmarks from the system. When an ontology entity is removed, its annotated representations will be removed from all subviews.

#### 4.4.2.2 Ontology Sections Panel


The Ontology Sections Panel, which is found at the top center position of PRONTOVISE, presents us with an ‘overview’ subview which fully affords ontology structure and promotes highly connected ontology entities as potential landmarks. We have designed the Ontology Sections Panel to represent a series of ontology sections, headed by its high-level ontology entity, determined from the set of entities associated with direct routes from the root entity

of the ontology. We supply information regarding the depth and comparative size of each ontology section to help our sensemaking activities towards the distribution of entities within the ontology, as well as preview the ontology relations between groupings of entities. During encounters with the ontology sections and their structure, the concept of distance from the root ‘super classes’ becomes an important assessment metric. For each ontology section, we are supplied a series of vertically distributed blue bars that are sized proportionately to the number of ontology entities found at that distance from the root superclass. By default, we distort the width of each ontology section by the percentage of entities within it compared to the total number of entities of the ontology. As some ontology sections have a significantly smaller percentage of entities, they can sometimes be adversely affected by this distortion technique. To address this, we have included an interaction which allows us to adjust the scaling from its default state into a fish-eye distortion concentrated on our selected section, as seen in Figure 11.



**Figure 4-11** A depiction of the distortion technique within the Ontology Sections Panel. This technique can be adjusted through interaction. This is achieved by holding the Shift key while directing the mouse over a section. Releasing the Shift key will end the interaction event and lock in the sizing adjustments. If adjustments have been made, yet the user would like to return to the original distortion scaling, we have provided a Reset button at the top left corner of the Ontology Sections Panel.

Additionally, we have supplied a magic lens tool that appears when we click the checkbox labeled ‘lens’ in the subview (Ukrop, Číková, & Kapec, 2013). With this overlaid magic lens, we can scan the magic lens over specific ontology sections to reveal more information of the ontology structure. There are two types of information available, as seen in Figure 12. First, by selecting the levels radio button, users will be presented with exact level depth, including the total number of entities within each level. The second available magic lens choice is landmarks, which annotates the important ontology entities from the sections, as determined by the entities with the most relations assigned to them within the ontology.




**Figure 4-12** A depiction of the magic lens within the Ontology Sections Panel. We have designed an interaction to occur when users drag the bottom portion of the magic lens horizontally across the ontology sections, which will both refresh the ontology section’s distortion technique and expose information for that section.

When we select a section header, a specific level bar, or an ontology entity annotated within a section, a red border is placed around it to support wayfinding activities. Then, all relevant subviews within the tool adjust to match the selected position within the ontology.

#### **4.4.2.3 Section Levels Panel**


The Section Levels Panel, which is found directly to the right of the Search and Pinning Panel, and below the Ontology Sections Panel, provides a ‘context’ subview depicting the levels of a selected ontology section produced by ontology relations, and the entities contained within them. We have designed this to provide us with the ability to inspect ontology entities and their shared relations within the scope of a section level to promote activities which use and develop landmark and route knowledge.

The Section Levels Panel depicts a list of levels ordered by their depth from the root entity, where an ontology level is the set of entities which share the same distance from the root ontology entity of a section. Each level has a line plot representing a summary distribution of the ontology entities within that level, as well as a red circle plotting the entity which is calculated to be the most linked ontology entity in that level. This metric of importance is calculated as the entity which maintains the highest total number of entities which are descendants through inheritance relations. These levels can be seen in Figure 13.



**Figure 4-13** The initial list of levels within the Ontology Section Levels Panel.

When we select a level, it expands to show a connected entities chart, allowing us to inspect ontology entities and their shared ontology relations. We are also provided with a count of the number of entities contained within the level. These ontology entities are positioned vertically based on the number of entities that inherit from it and distributed horizontally with others within the level which share the same immediate inheritance. Plotted in the graph as circles, these supply a series of useful interactions, as depicted and described in Figure 14.




**Figure 4-14** A level within the Ontology Section Levels Panel. The level has three main interactions. First, when we move our cursor over a circle, a label is generated which displays the name of the entity and annotates the location of the ontology entity within the Level Landmark Entities Panel. Second, if we click on a circle, we perform an interaction which selects that entity as the initial position within the Entity Network Panel, the Path Explorer Panel, and the Entity

**Details Panel.** Additionally, when a level has many ontology entities, the available visual space may become very crowded. To address this, we designed an interaction which allows us to distort the space by holding Shift and activating our mouse scroll wheel, which expands and contracts the horizontal scale of the plot graphs. We then can drag the plot graphs left or right using a single mouse click and drag action, allowing us the ability to closely inspect the full set of ontology entities and relations within the level.

An expanded level presents a plot graph above the entity chart with rectangles representing the ontology entities that share relations to the ontology entities within the current level, also seen in Figure 14. A line axis representing specific level's depth intersects the approximate middle of vertical range of the rectangular plot graph. Furthermore, to represent the effect of similar inheritance within the level, the width of each rectangle is scaled to reflect the number of entities which inherited from it, larger widths representing more inheritances. These representations and interactions provide many encounters with information describing the ontological space and can help the performance of the cognitive activities which promote spatial knowledge for cognitive map formation. For instance, we can use this to determine the impact an entity has on the current level by moving our cursor over a rectangle, which will annotate a blue border around the ontology entities which share an ontology relation, as well as text labels.

Finally, we have developed a magic lens tool within the subview (Ukrop et al., 2013). When we click the checkbox labeled 'lens' in the subview header, we can expose more information by scanning the magic lens over a specific range of the level, as depicted in Figure 15. This magic lens also keeps its scope when using zooming and panning interactions.



**Figure 4-15** The magic lens within the Ontology Section Levels Panel. We have designed the magic lens with an interaction which generates lines to represent the set of ontology relations within the level as we drag horizontally across the visual space.

#### **4.4.2.4 Level Landmark Entities Panel**

The Level Landmark Entities Panel, which is directly to the right of the Section Levels Panel and below the Ontology Sections Panel, provides a 'context' subview which allows us to inspect the connectivity between the ontology entities at a specific level of the ontology. The Level Landmark Entities Panel maintains representations and

interactions which are particularly useful for ontology levels of significant connectivity, as it can be challenging to navigate through levels which possess large numbers of ontology entities and relations. By supplying an ordered perspective into the connectivity of a level, we can more effectively direct our cognitive activities as we move through our learning task.

The Level Landmark Entities Panel will generate a triangular matrix collecting the 13 most important entities within the level, where the metric of importance is calculated as the total number of entities which are descendants through inheritance relations. When we pin ontology entities, they will be included within the matrix. The Level Landmark Entities Panel also includes a representation maintaining a node-link graph and two text areas which helps us build associations between ontology entities. When we interact with a matrix position, the node-link graph and text areas update to visually represent the inheritance tree between the ontology entities of that matrix position up to their nearest common parent. These representations are depicted and described in Figure 16.





Figure 4-16 The Level Landmark Entities Panel depicting an ontology level with significant connectivity within its matrix representation. We can see, using the magic lens of the Section Levels Panel, that the level has countless numbers of ontology entities and relations, which have been analyzed and presented in a usable manner with the Level Landmark Entities Panel. Within the matrix representation, each ontology entity is represented as a circle accompanied by a text label. Each circle maintains a red to white fill encoding reflecting its importance calculation, red being the most inherited, and white being the least. When we move our cursor over a circle, the text label grows and bolds for rapid association between the matrix position and its row and column labels. A color spectrum and text label is provided at the intersection points of the matrix representing ontological distance. This distance calculation is determined by the number of inheritances performed when defining the ontologies up to their nearest common parentage. For example, when a matrix position reflects the intersection between two ontology entities which inherit from the same immediate parent, their distance will be calculated and displayed as 2.

When we select a specific ontology entity or set of entities within the Level Landmark Entities Panel, this will tell the tool to use them as the initial position or positions in the Entity Network Panel, the Path Explorer Panel, and the Entity Details Panel.

#### 4.4.2.5 Entity Network Panel

The Entity Network Panel is found directly to the right of the Ontology Sections Panel and above the Path Explorer Panel. This panel provides a ‘context’ subview that allows us to interact with a representation of a network of ontology entities and their relations as we perform our cognitive activities within the space.

The Entity Network Panel maintains three regions: the selected ontologies at the center position, their parents above, and their children below. When the Entity Network Panel is initialized from the Section Levels Panel, a single selected ontology entity is represented as a circle. When the Entity Network Panel is initialized from the Level Landmark Entities Panel, the two ontology entities from that selection are depicted. In either case, all ontology entities are represented by a circle with text label. In these regions, links are depicted between ontology entities reflecting inheritance relations. When two ontology entities are represented, it also represents their lowest common parent. These regions can be seen in Figure 17.




**Figure 4-17** The Entity Network Panel depicting a low-level graph-like abstraction of specific ontology entities within the ontology network: parents (super classes), children (sub classes), and shared inheritances. Ontology entities are represented by a blue filled circle and text label. In these regions, relations are depicted with lines which link ontology entities to reflect relationship and localized structure. When two ontology entities are selected, their shared network is depicted with an additional representation maintaining the relations between each, their shared inherited parent, and distance. Additional information is exposed by moving the mouse over.

The Entity Network Panel supplies a movement interaction which, when an ontology entity is selected, that ontology entity becomes the new position. This interaction will also adjust positions of the Path Explorer Panel and Entity Details Panel.

#### 4.4.2.6 Path Explorer Panel

The Path Explorer Panel, which is directly below the Entity Network Panel, provides a ‘context’ subview that allows us to examine the full set of inheritance relations from the root of the ontology section and down to the selected position. These sets of representations and interactions fully expose the low-level structure of the ontological space, and can promote the final, more granular, stages of cognitive map formation.

The Path Explorer Panel represents the set of ontology levels traversed when navigating the inheritance path of the current entity, and in each, the full set of sibling ontology entities. Ontology entities which share a relation to the current position are represented in the bottom as rectangles, where their width and color represents the number of ontology entities that are their children. This color fill ranges from white to red, where red is the highest number of inheritances within the level. When an ontology entity is never inherited, it is a leaf of the ontology structure. For these, their height is slightly increased to promote visibility and are given a purple fill. Ontology entities which are within inheritance lineage of the selected ontology entity have a slightly increased height to improve the visibility of the inheritance path. Links are provided between ontology entities to reflect their connectivity throughout the inheritance path. These design considerations can be seen in Figure 18.



**Figure 4-18** The Path Explorer Panel depicting all the levels within the selected path of the ontology, where each level maintains a set of rectangles representing all sibling entities and a text label reflecting its depth within the levels of the ontology.

The Path Explorer Panel has an interaction which, when an ontology entity is selected, that ontology entity becomes the new position. This interaction also adjusts the positioning of the Entity Network Panel and Entity Details Panel.

#### **4.4.2.7. Entity Details Panel**

The Entity Details Panel, which is to the far right of the system, presents a ‘details’ subview which depicts the information of an ontology entity in the form of a standard listing. This listing is based on the specification of the ontology, although will typically reflect information like Index, Name, Definition, Synonym, Superclass, Subclass, and External Link.

Each of the Synonym, Superclass, Subclass, and Link listings within the Entity Details Panel provide two interactions. The first interaction, provided by Synonym, Superclass, and Subclass, allows us to select an ontology entity represented in text, which will show it as the new position in the Entity Details Panel, Entity Network Panel, and the Path Explorer Panel. The second interaction, if supported by the ontology specification, allows users to leave PRONTOVISE and inspect the ontology creator’s official documentation on the web.


### **4.5 Usage Scenario**

In this section, we will describe a usage scenario which demonstrates how PRONTOVISE can support the stages of cognitive map formation of an ontology. For an expanded demonstration of PRONTOVISE, we also provide a demonstration video “Visual Demonstration of PRONTOVISE” in the supplementary materials.

Domain expertise can be assessed as a spectrum of knowledge, ranging from a member of the general public with no expertise, up to domain expert such as a geneticist, doctor, or medical researcher. We will collapse this range into two general user types—the ‘non-expert’ and the ‘expert’. A scenario will be presented to demonstrate their ability to begin, or in the case of the ‘expert’, build upon, their cognitive map of an ontology.

For purposes of demonstration, the Human Phenotype Ontology (HPO) will act as the ontology dataset of choice within this usage scenario and shared demonstration materials. HPO has been selected because of its high complexity resulting from its exhaustive and expert-defined domain coverage. HPO is a controlled and standardized vocabulary reflecting the human disease and phenotypic abnormality domain, and includes associated annotations in the domains of bioinformatics, biochemistry, and human genetics. HPO is an active ontology, consisting of over 11,000 ontology entities, as well as over 110,000 disease annotations (Kohler et al., 2014). For instance, HPO maintains an ontology entity for Blindness, which possesses a superclass of Visual Impairment, a subclass of Congenital Blindness, and is annotated to be associated with a variety of diseases, such as a variant of colorblindness defined as Achromatopsia 2 (Köhler & Robinson, 2016). Each HPO ontology entity and relation is accompanied by attributes such as names, definitions, ontology indexing, synonyms, class relationships, logical definitions, and domain expert commentary, to name a few. For additional details on the Human Phenotype Ontology, see (Köhler et al., 2018; “The Human Phenotype Ontology,” 2020).


PRONTOVISE allows users to upload valid RDF OWL file types, such as the ones produced by Stanford University’s Protégé Editor and Cognitum’s Fluent Editor. For this usage scenario, we will be using the Human Phenotype Ontology (HPO) as our selected ontology dataset describing an unfamiliar ontology. We will begin these usage scenarios with the tool in its initial state, as seen in Figure 19.



**Figure 4-19** An overall view of the PRONTOVISE ontology visualization system in its initial state after the Human Phenotype Ontology has been uploaded.

In our usage scenario, we take on the role of a user who has no prior experience with HPO. This means that we have not developed any level of understanding towards the ontology. Our initial interactions will require the tool to support our cognitive map formation through encounters with the ontological entities and relations to promote the early stages of landmark, route, and survey knowledge.


In the Ontology Sections Panel subview, we see an overview of the ontology structure. From Figure 20, we see that each of the sections of HPO are headed by a root ontology entity. To begin, we would like to examine the various top ontology Section entities, so that we can become aware of the entities of HPO which could act as initial landmarks.



**Figure 4-20** The initial stage of the Ontology Sections Panel subview.


During encounters within the Ontology Sections Panel, we can begin to develop initial survey knowledge towards the structure of HPO. For example, we see that there are quite a few sections in HPO. We also see that the sections represented by ontology entities like “Abnormality of the Skeletal System” and “Abnormality of Limbs” consume significantly larger portions of HPO. Notably, there are a few sections with very small visual spaces. Additionally, we see the general shapes of each section, suggesting the potential ontology entities and relations which form its structure. Some sections like “Skeletal System” have very extended paths, going up to 14 levels away from the top

entity acting as our landmark for the section, while other sections expand out only five steps away. We further inspect the smaller sections like the section headed by the “Abnormality of the Musculature” by holding the Shift key while directing our mouse over a section, as seen in Figure 21. Releasing the Shift key will end the interaction event and lock in the sizing adjustments.




**Figure 4-21 Adjusting the scaling of the Ontology Sections Panel subview to enlarge the “Abnormality of the Musculature” section, which is normally represented as a much smaller portion of the visual space.**

We inspect the contents of each section. PRONTOVISE provides us with information regarding the number of ontology entities and ontological distance of each level relative to the entity acting as a landmark for that section. Additionally, when we select the checkbox labeled ‘lens’ in the subview header, a magic lens is added to our mouse, allowing us to rapidly scan each section. This reveals structural information of the ontology and previews the potential information in that section which could support the building of route knowledge. Figure 22 shows a magic lens activation on the “Abnormality of Limbs” ontology section, depicting the depths and the number of individual ontology entities within each level.




**Figure 4-22 The result of a magic lens levels’ activation on the “Abnormality of Limbs” ontology section.**

We also select a second available magic lens choice labeled “landmarks”, which changes the functionality of the scanning lens to preview the potential landmark knowledge available in that section. Figure 23 depicts our interaction where the magic lens is placed over the “Abnormality of Limbs” ontology section, encoding the most prominent ontology entities which we can use as our landmarks for each level.




**Figure 4-23** The result of a magic lens on the “Abnormality of Limbs” ontology section.

As we select a section header, a specific level bar, or an annotated entity within a section, the subview updates to represent the confirmation of this interaction, as well as signal other subviews to change their positions to match. As seen in Figure 24, we further our exploration of HPO within the “Abnormality of the Skeletal System” section by selecting for deeper inspection.




**Figure 4-24** The selection of the “Abnormality of the Skeletal System” section.

We now move to the Sections Levels Panel subview, directly below the Ontology Sections Panel subview, which magnifies the levels of the selected ontology section and the ontology entities contained within it. We begin our inspection of the various levels of this ontology section, seeing the number of ontology entities in each level. This provides us the potential to build route knowledge of the ontology entities we have encountered and used as landmarks. Figure 25 shows the initial state of the Sections Level Panel subview. We see that there are 14 levels within the section and that there are 9 ontology entities in the second level of the section. We also see a more detailed representation of the contents of each level, where levels like depth 11 (d11) carry a significant number of ontology entities, while others carry less.




**Figure 4-25** An overview of the Section Levels Panel subview depicting the “Skeletal System” ontology section levels after navigating from the Ontology Sections Panel subview.

We are interested in d13, so we select that level for inspection. In response, PRONTOVISE has expanded the subview to display the ontology entity and relations of the 13th level of the section. Here, the connected parentage and landmark chart supplies many opportunities for insight. Using mouse over and click interactions, we interact with the different ontology entities in the level. This allows us to form associations between ontology entities, preview an overview of the ontology relations between its ontology entities, as well as signal other subviews to preview our target ontology entity. As seen in Figure 26, then we select the “Duplication of Phalanx of 3rd Finger” found in the 9th level of the section to see what ontology entities in this level have inherited from it. We also use the magic lens that shows ontology relations between ontology entities to promote route and survey knowledge, shown in Figure 27.



**Figure 4-26** Inspecting the 13th level of the “Skeletal System” ontology section.



**Figure 4-27 Using the magic lens to inspect a set of ontology entities and their inheritance relations.**

We select the 3rd level of the ontology section for deeper inspection. We direct our attention to the Level Landmark Entities Panel subview, found to the right of the Section Levels Panel subview. The Level Landmark Entities Panel subview allows us to inspect the ontology relations between the major ontology entities contained within the chosen level. As depicted in Figure 28, we begin to encounter the ontological space in increasingly lower levels of visual abstraction. We see individual ontology entities with numerical values describing the ontological distance between each of the ontology entities of the level. We select the meeting point between “Abnormal Appendicular Skeleton Morphology” and “Abnormal Bone Structure”, highlighting that those ontology entities have a two-step separation. Additionally, we see that this choice has provided a description of the matching ontology entities, as well as a line demonstrating the full network of ontology relations. We see that their distance separation of 2 is because they each inherit from a shared parent, that of “Abnormality of Skeletal Morphology”. We also investigate a different pair of ontology entities, “Abnormal Joint Morphology” and “Epiphyseal Stippling” and see those ontology entities sharing a more distant relationship, as depicted in Figure 29.




Figure 4-28 An overview of the Level Landmark Entity Subview, representing the 3rd level of the “Abnormality of Skeletal System” ontology section.





Figure 4-29 Route between of “Abnormal Joint Morphology” and “Abnormal Epiphyseal Stippling”.

So far, PRONTOVISE has presented us encounters with ontology entities, providing us opportunities to build our landmark knowledge, make associations between our landmarks helping the development of route knowledge, and combining for initial survey knowledge of the structure of HPO. We now want to begin comparing specific ontology entities and relations in the context of the full HPO system. Therefore, we direct our interest towards the Entity Network Panel subview, positioned to the right of the previously encountered subviews. The Entity Network Panel subview allows us to inspect the low-level abstractions of specific ontology entities within the system, gathering insight towards exact entity positioning for parent and child ontology entities and the shared ontology relations which


reflect ontology structure. We select “Abnormal Appendicular Skeleton Morphology” in a prior subview, generating that entity as the target of the Entity Network Panel subview, as seen in Figure 30.



**Figure 4-30** The initial state of the Entity Network Panel subview when an ontology entity is chosen as the initial position. In this case, “Abnormal Appendicular Skeleton Morphology” has been selected.

We see that the chosen “Abnormal Appendicular Skeleton Morphology” entity directly inherits from a single ontology entity and is inherited from four other ontology entities on various ontology section levels. We then select two entities for the entity network within the prior Level Landmark Entities Panel subview, and the two ontology entities are used as positions side-by-side, as depicted in Figure 31.






**Figure 4-32 The Path Explorer Panel subview after selecting “Abnormal Appendicular Skeleton Morphology” as its current position.**


PRONTOVISE provides us with three support subviews that can extend our ability to generate encounters. The first two of these exist within the Search and Pinning Panel subview, Search and Pinning, respectively, and the third is the Entity Details Panel subview.

After using PRONTOVISE, we have developed some level of understanding towards HPO. The search functionality found with the Search and Pinning Panel subview allows us to use a text-based search bar to specifically target ontology entities with the assistance of suggestions from a type-ahead. Based on our experiences so far, we are interested to see if there are any other “skeleton”-related entities existing in HPO outside of the section we have already encountered. We type in the search bar, as seen in Figure 33. We see that many relate to “skeleton”, the Ontology Sections Panel, which has updated in response to the search query, suggesting that there are indeed some outside of the “Skeletal System” section that relate to “skeleton”.



**Figure 4-33 The Ontology Landmark Search results from typing “skeleton”.**

We record our interest in a new ontology entity, “Broad Forearm Bones”, by pinning it within PRONTOVISE. Clicking the “Pin” button, the ontology entity is added to the Pinning Panel, which we access by selecting the tab in the subview, as seen in Figure 34. The ontology entity has received a permanent point of reference within the tool and has been assigned a unique color which will be used whenever we encounter it within the tool.



**Figure 4-34 “Broad Forearm Bones” has been pinned, designating it as an important ontology entity, which is to be highlighted with its assigned color whenever it appears in a PRONTOVISE subview.**

After we select the “Broad Forearm Bones” ontology entity during our interaction with PRONTOVISE, a third and final support subview, the Entity Details Panel subview, becomes active. At this subview, we are presented with the full set of HPO information for our ontology entity. In Figure 35, we see that the subview supplies us the following ontology details for a specific entity: HPO Index Number, Name, Definition, Synonym, Superclass, Subclass, and HPO Link.




Figure 4-35 When “Broad Forearm Bones” has been selected within any subview, the Entity Details Panel subview depicts all information for that ontology entity as provided by HPO.

## 4.6 Discussion and Conclusions

In this paper, we began with an introduction of the topics of cognitive maps, ontologies, and interactive visualization tools. It was found that various research at varying levels of granularity across a wide range of theoretical and experimental disciplines has been directed towards understanding the functionality of our cognitive processes and the effect they have on the performance of our cognitive activities. The theoretical framework of the cognitive map and its formation process was introduced, which leads current understanding towards how our brains organized knowledge of complex spaces. Then, an introduction to the use, creation, and limitations of ontologies was presented. This section described that ontologies are an expert-defined standardized common vocabulary describing the knowledge of a domain. The introductory content concluded with an examination of the fields of information visualization and visual analytics, discussing how designers can create visualization tools using visual representation and interaction design to support our performance of cognitive activities.

Next, we examined existing work on cognitive load and the use of interactive visualizations to support learning tasks. Here, it was found that recent studies have established that for supporting learning tasks, there is no one specific level of cognitive load that is appropriate. Instead, cognitive load should be understood as a set of extraneous, intrinsic, and germane loads which adjust for the specific conditions of the tool and task context. It was found that interactive visualization tools are a valuable resource for complex learning, as studies have found that if designed correctly, they can provide an effective environment for engaging learners with the types of information encodings which best align with the needs for learning tasks. The examination of existing work concluded with an exploration of leading insight towards the design of visual representations and interactions to support cognitive

mapping of spatial knowledge, alongside a summary of the cognitive activities performed within spaces. From these findings, we formalized a set of high-level design criteria for designing interactive visualization tools which support learning tasks through alignment with the cognitive map framework and its formation process. A review was performed on existing tools which visualize ontology datasets. This review categorized each tool based on generalized subview components and, for each, analyzed their strengths and weaknesses towards supporting the conditions for cognitive map formation.

Following this, we presented PRONTOVISE (PRogressive ONTOlogy VISualization Explorer), an interactive visualization tool which applied the criteria in its design to support us in understanding unfamiliar ontologies. In this, we explained the technological features of the PRONTOVISE, and described its workflow and design within the context of our high-level design criteria. PRONTOVISE was described as an interactive visualization tool that represents ontology datasets using a combination ‘List+Overview+Context+Details’ design. The presentation continued with a detailed description of the considerations made when designing each of the subviews of PRONTOVISE to satisfy the established high-level criteria. Through a usage scenario which describes an initial set of encounters with the Human Phenotype Ontology (HPO), we demonstrated how the design of PRONTOVISE uses novel ontology dataset visual representations and interactions to provide us valuable encounters which support the requirements for cognitive map formation.

From our investigation of related work and existing tools, as well as through our description and usage scenario of the PRONTOVISE design, several implications arise. We find that there is value in design criteria which generalize a set of high-level requirements yet refrain from specifying the exceedingly granular patterns and processes which are often associated with low-level design frameworks. Through this higher level, we were able to see that many fashions of interactive visualization tools were possible, each providing us encounters with visual representations and interactions at various levels of novelty within their design. We believe a strength within the design of PRONTOVISE fully encompasses the strengths presented within the spread of existing designs, while having the opportunity to appropriately assess and address the weaknesses present in preceding work. We also believe that using the design criteria, PRONTOVISE provides the support for any ontology dataset learning task, as it provides a stable, iterative, and scalable design which supports use for any appropriately encoded domain, and for use by any level of expertise.

Yet, from our work on the design and PRONTOVISE, we also find limitations. First, we acknowledge the limited scope of evaluation within the paper. It is our intention to explore expanded evaluations of the design within future work. Next, we found that a number of subviews were needed to fulfill the requirements of the design criteria. As a result, we find that PRONTOVISE requires a significant amount of display space, such that it would not be able to facilitate the same level of quality towards the performance of learning tasks within a reduced display space. This means that the current design of PRONTOVISE is not practical for small screens like notepad laptops or mobile devices. A target for future work may then be to investigate this problem space for small screens. Another limitation with PRONTOVISE is that it takes full advantage visualization technologies to produce its many novel visual representations and interactions within its subviews and thus demands an attunement period before it can be used optimally. This aspect of design, that of attunement, may also be a valuable topic of interest for future research.

Finally, PRONTOVISE currently supports the OWL RDF ontology dataset format, however, there are many formats available. PRONTOVISE would be improved by expanding its support to all formats which are used to digitally encode ontologies.

In conclusion, we hope that insight gathered from this paper inspires innovative research and provides valuable guidance in the design of future work which visualizes ontologies for learning tasks. We hope to continue exploring this problem space, including but not limited to a deeper inspection of PRONTOVISE through an expanded usability study and further investigation towards the design of interactive visualization tools that support the performance of other challenging knowledge-based tasks.

## 4.5 References

- Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*, *100*(2), 490–509. <https://doi.org/10.1016/j.neuron.2018.10.002>
- Borland, D., Christopherson, L., & Schmitt, C. (2019). Ontology-Based Interactive Visualization of Patient-Generated Research Questions. *Applied Clinical Informatics*, *10*(03), 377–386.
- Bostock, M. (2016). D3.js Data-Driven Documents. Retrieved from <https://d3js.org/>
- Craig, M., Dewar, M., Harris, M. A., Della Sala, S., & Wolbers, T. (2016). Wakeful rest promotes the integration of spatial memories into accurate cognitive maps. *Hippocampus*, *26*(2), 185–193. <https://doi.org/10.1002/hipo.22502>
- Davies, J., & Michaelian, K. (2016). Identifying and individuating cognitive systems: a task-based distributed cognition alternative to agent-based extended cognition. *Cognitive Processing*, *17*(3), 307–319. <https://doi.org/10.1007/s10339-016-0759-4>
- Demelo, J., Parsons, P., & Sedig, K. (2017). Ontology-Driven Search and Triage: Design of a Web-Based Visual Interface for MEDLINE. *JMIR Medical Informatics*, *5*(1), e4. <https://doi.org/10.2196/medinform.6918>
- Dessimoz, C., & Škunca, N. (2017). *The Gene Ontology Handbook* (Vol. 1446). <https://doi.org/10.1007/978-1-4939-3743-1>
- Dudáš, M., Lohmann, S., Svátek, V., & Pavlov, D. (2018). Ontology visualization methods and tools: a survey of the state of the art. In *The Knowledge Engineering Review* (Vol. 33). <https://doi.org/10.1017/s0269888918000073>
- Ekstrom, A. D., & Ranganath, C. (2017). Space, Time and Episodic Memory: the Hippocampus is All Over the Cognitive Map CA , USA 2 . Department of Psychology , University of California , Davis , CA , USA 3 . Neuroscience Graduate Group , University of California , Davis , CA , USA \* Each autho. *Hippocampus*, *093052*, 1–16.
- Epstein, R. A., Patai, E. Z., Julian, J. B., & Spiers, H. J. (2017). The cognitive map in humans: Spatial navigation and beyond. *Nature Neuroscience*, *20*(11), 1504–1513. <https://doi.org/10.1038/nn.4656>
- Falconer, S. (2010). OntoGraf. Retrieved May 21, 2020, from <https://protegewiki.stanford.edu/wiki/OntoGraf>
- Fast, K. V. (2010). Human-Information Interaction for Digital Libraries : From Document Repositories to Knowledge Environments. In *Knowledge Creation Diffusion Utilization*.

- Jakus, G., Milutinovic, V., Omerović, S., & Tomazic, S. (2013). *Concepts, Ontologies, and Knowledge Representation*. <https://doi.org/10.1007/978-1-4614-7822-5>
- Jusoh, S., Awajan, A., & Obeid, N. (2020). The Use of Ontology in Clinical Information Extraction. *Journal of Physics: Conference Series*, 1529(5). <https://doi.org/10.1088/1742-6596/1529/5/052083>
- Kallioniemi, P., Hakulinen, J., Keskinen, T., Turunen, M., Heimonen, T., Pihkala-Posti, L., ... Lulea University of, T. (2013). Evaluating landmark attraction model in collaborative wayfinding in virtual learning environments. *PUZZLE: Revista Hispana de La Inteligencia Competitiva*, 4(17), 15–18. <https://doi.org/10.1145/2541831.2541849>
- Katifori, A., Halatsis, C., Lepouras, G., Vassilakis, C., & Giannopoulou, E. (2007). Ontology visualization methods--a survey. *ACM Computing Surveys*, 39(4), 10-es. <https://doi.org/10.1145/1287620.1287621>
- Katifori, A., Torou, E., Vassilakis, C., Lepouras, G., & Halatsis, C. (2008). Selected results of a comparative study of four ontology visualization methods for information retrieval tasks. *2008 Second International Conference on Research Challenges in Information Science*, (September 2016), 133–140. <https://doi.org/10.1109/RCIS.2008.4632101>
- Köhler, S., Carmody, L., Vasilevsky, N., Jacobsen, J. O. B., Danis, D., Gourdine, J.-P., ... Robinson, P. N. (2018). Expansion of the Human Phenotype Ontology (HPO) knowledge base and resources. *Nucleic Acids Research*, 47(D1), D1018–D1027. <https://doi.org/10.1093/nar/gky1105>
- Kohler, S., Doelken, S. C., Mungall, C. J., Bauer, S., Firth, H. V., Bailleul-Forestier, I., ... Robinson, P. N. (2014). The Human Phenotype Ontology project: Linking molecular biology and disease through phenotype data. *Nucleic Acids Research*, 42(D1), 966–974. <https://doi.org/10.1093/nar/gkt1026>
- Köhler, S., & Robinson, P. (2016). HPO Web Browser: Blindness Infopage. Retrieved from Human Phenotype Ontology website: <http://compbio.charite.de/hpweb/showterm?id=HP:0000618>
- Lohmann, S., Negru, S., Haag, F., & Ertl, T. (2016). Visualizing Ontologies with VOWL. *Semantic-Web-Journal.Net*, (August 2016), 21. <https://doi.org/10.3233/SW-150200>
- Lytvyn, V., Dosyn, D., Vysotska, V., & Hryhorovych, A. (2020). *Method of Ontology Use in OODA*. 409–413. <https://doi.org/10.1109/dsmp47368.2020.9204107>
- Mayer, R. E. (2014). *The Cambridge Handbook of Multimedia Learning*. Cambridge University Press.
- Mouromtsev, D., Pavlov, D., Emelyanov, Y., Morozov, A., Razdyakonov, D., & Galkin, M. (2015). The Simple Web-based Tool for Visualization and Sharing of Semantic Data and Ontologies. *International Semantic Web Conference (Posters & Demos)*. , Bethlehem, PA, USA.
- Musen, M. A., & Team, P. (2015). The Protégé Project: A Look Back and a Look Forward. *AI Matters*, 1(4), 4–12. <https://doi.org/10.1145/2757001.2757003>
- Mutlu-Bayraktar, D., Cosgun, V., & Altan, T. (2019). Cognitive load in multimedia learning environments: A systematic review. *Computers and Education*, 141(July), 103618. <https://doi.org/10.1016/j.compedu.2019.103618>
- Nightingale, O. (2017). Lunr.js. Retrieved August 10, 2020, from <http://lunrjs.com/>
- Ninkov, A., & Sedig, K. (2019). VINCENT: A visual analytics system for investigating the online vaccine debate.

- Online Journal of Public Health Informatics*, 11(2). <https://doi.org/10.5210/ojphi.v11i2.10114>
- Ola, O., & Sedig, K. (2016). Beyond simple charts: Design of visualizations for big health data. *Online Journal of Public Health Informatics*, 8(3). <https://doi.org/10.5210/ojphi.v8i3.7100>
- Parsons, P., & Didandeh, A. (2015). *Interactivity in Visual Analytics : Use of Conceptual Frameworks to Support Human-Centered Design of a Decision-Support Tool*. (Hicss 48), 1138–1147.
- Parsons, P., & Sedig, K. (2014). *Handbook of Human Centric Visualization*. 671–691. <https://doi.org/10.1007/978-1-4614-7485-2>
- Pereira Rocha, J. A., de Paula, C. P., & Sirihal Duarte, A. B. (2016). Distributed Cognition As a Theoretical Framework for Information Behavior Studies. *Informacao & Sociedade-Estudos*, 26(2), 91–105.
- Pesquita, C., Ivanova, V., Lohmann, S., & Lambrix, P. (2018). A framework to conduct and report on empirical user studies in semantic web contexts. *European Knowledge Acquisition Workshop, 11313*, 567–583. [https://doi.org/10.1007/978-3-030-03667-6\\_36](https://doi.org/10.1007/978-3-030-03667-6_36)
- Po, L., Bikakis, N., Desimoni, F., & Papastefanatos, G. (2020). Linked Data Visualization: Techniques, Tools, and Big Data. *Synthesis Lectures on the Semantic Web: Theory and Technology*, 10(1), 1–157. <https://doi.org/10.2200/s00967ed1v01y201911wbe019>
- Priya, M., & Kumar, C. A. (2015). A survey of state of the art of ontology construction and merging using Formal Concept Analysis. *Indian Journal of Science and Technology*, 8(24). <https://doi.org/10.17485/ijst/2015/v8i24/82808>
- Ragan, E. D., Bowman, D. A., & Huber, K. J. (2012). Supporting cognitive processing with spatial information presentations in virtual environments. *Virtual Reality*, 16(4), 301–314. <https://doi.org/10.1007/s10055-012-0211-8>
- Rector, A., Schulz, S., Rodrigues, J. M., Chute, C. G., & Solbrig, H. (2019). On beyond Gruber: “Ontologies” in today’s biomedical information systems and the limits of OWL. *Journal of Biomedical Informatics: X*, 2, 100002. <https://doi.org/https://doi.org/10.1016/j.yjbinox.2019.100002>
- Román-Villarán, E., Pérez-Leon, F. P., Escobar-Rodríguez, G. A., Martínez-García, A., Álvarez-Romero, C., & Parra-Calderón, C. L. (2019). An ontology-based personalized decision support system for use in the complex chronically ill patient. *Studies in Health Technology and Informatics*, 264, 758–762. <https://doi.org/10.3233/SHTI190325>
- Saleemi, M. M., Rodríguez, N. D., Lilius, J., & Porres, I. (2011). A framework for context-aware applications for smart spaces. In *Smart Spaces and Next Generation Wired/Wireless Networking: Vol. 6869 LNCS* (pp. 14–25). [https://doi.org/10.1007/978-3-642-22875-9\\_2](https://doi.org/10.1007/978-3-642-22875-9_2)
- Schlegel, D. R., & Elkin, P. L. (2016). Ontology Visualization Tools to Assist in Creating and Maintaining Textual Term Definitions. Retrieved November 28, 2020, from <https://pdfs.semanticscholar.org/3ffb/60e33e8009761b4465dee57c7ef86c4eda22.pdf>
- Sedig, K., & Parsons, P. (2013). Interaction Design for Complex Cognitive Activities with Visual Representations: A Pattern-Based Approach. *AIS Transactions on Human-Computer Interaction*, 5(2), 85–133. <https://doi.org/10.5121/ijfcst.2014.4403>

- Sedig, K., Parsons, P., Liang, H.-N., & Morey, J. (2016). Supporting Sensemaking of Complex Objects with Visualizations: Visibility and Complementarity of Interactions. *Informatics*, 3(4), 20. <https://doi.org/10.3390/informatics3040020>
- Sedig, K., Rowhani, S., & Liang, H. (2005). *Designing interfaces that support formation of cognitive maps of transitional processes : an empirical study*. 17, 419–452. <https://doi.org/10.1016/j.intcom.2005.02.002>
- Semafora Systems. (2020). OntoStudio. Retrieved May 21, 2020, from <https://www.semafora-systems.com/ontobroker-and-ontostudio-x>
- Seufert, T. (2018). The interplay between self-regulation in learning and cognitive load. *Educational Research Review*, 24(August 2017), 116–129. <https://doi.org/10.1016/j.edurev.2018.03.004>
- Shneiderman, B. (1996). The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. *Proceedings of the IEEE Symposium on Visual Languages*, 336–343.
- Silva, I. C. S., Santucci, G., & Freitas, C. M. D. S. (2019). Visualization and analysis of schema and instances of ontologies for improving user tasks and knowledge discovery. *Journal of Computer Languages*, 51(February 2018), 28–47. <https://doi.org/10.1016/j.cola.2019.01.004>
- Sun, M., Wang, M., & Wegerif, R. (2019). Using computer-based cognitive mapping to improve students' divergent thinking for creativity development. *British Journal of Educational Technology*, 50(5), 2217–2233. <https://doi.org/10.1111/bjet.12825>
- The Human Phenotype Ontology. (2020). Retrieved May 10, 2020, from The Jackson Laboratory website: <https://hpo.jax.org/app/>
- Tobergte, D. R., & Curtis, S. (2013). A Survey on Ontologies for Human Behavior Recognition. *Journal of Chemical Information and Modeling*, 53(9), 1689–1699. <https://doi.org/10.1017/CBO9781107415324.004>
- TopQuadrant. (2020). TopBraid Explorer. Retrieved May 21, 2020, from <https://www.topquadrant.com/products/topbraid-explorer/>
- Tudorache, T., Nyulas, C., Noy, N. F., & Musen, M. A. (2013). WebProtégé: A collaborative ontology editor and knowledge acquisition tool for the web. *Semantic Web*, 4(1), 89–99.
- Ukrop, J., Číková, Z., & Kapec, P. (2013). Visual Access to Graph Content Using Magic Lenses and Filtering. *Proceedings of the 28th Spring Conference on Computer Graphics*, 1(212), 23–30. <https://doi.org/10.1145/2448531.2448534>
- Waller, D. E., & Nadel, L. E. (2013). *Handbook of spatial cognition*. Washington DC, United States of America: American Psychological Association.
- Wang, C., Fang, T., & Gu, Y. (2020). Learning performance and behavioral patterns of online collaborative learning: Impact of cognitive load and affordances of different multimedia. *Computers and Education*, 143(5), 103683. <https://doi.org/10.1016/j.compedu.2019.103683>
- Weisberg, S. M., & Newcombe, N. S. (2016). How do (some) people make a cognitive map? Routes, places, and working memory. *Journal of Experimental Psychology: Learning Memory and Cognition*, 42(5), 768–785. <https://doi.org/10.1037/xlm0000200>
- Yalçın, M. A., Elmqvist, N., & Bederson, B. B. (2016). Cognitive stages in visual data exploration. *ACM*

*International Conference Proceeding Series, 24-October(i)*, 86–95. <https://doi.org/10.1145/2993901.2993902>

Zhao, K., Ward, M., Rundensteiner, E., & Higgins, H. (2017). MaVis: Machine Learning Aided Multi-Model Framework for Time Series Visual Analytics. *Electronic Imaging*, 2016(1), 1–10. <https://doi.org/10.2352/issn.2470-1173.2016.1.vda-493>

# Chapter 5      Design of Generalized Search Interfaces for Health Informatics

This chapter has been published as: **Demelo, J., & Sedig, K.** (2021). Design of Generalized Search Interfaces for Health Informatics. *Information*, 12(8):317. Featured Article of Special Issue: The Digital Health New Era: Where We Stand and the Challenges.

We have made minor adjustments to the original material of this chapter to provide cohesion with the overall integrated article structure of this dissertation. Specifically, to distinguish between chapters, figures and tables have been provided an additional prepend reflecting the chapter number. Readers should be aware that chapter text will maintain original numbering references. For instance, “Figure 5-1” is equivalent to “Figure 1” in the chapter text.

## 5.1 Introduction

Health informatics is concerned with emergent technological systems that improve the quality and availability of care, promote the sharing of knowledge, and support the performance of proactive health and wellness tasks by motivated individuals (Wickramasinghe, 2019). Subareas of health informatics may include medical informatics, nursing informatics, consumer informatics, cancer informatics, and pharmacy informatics, to name a few. Simply put, health informatics is concerned with finding new ways to help stakeholders work with health information to be able to perform health-related tasks more effectively. Users in the health domain are increasingly taking advantage of computer resources in their tasks. For instance, a 2017 Canadian survey found that 32% of respondents within their last month had used at least one mobile application for health-related tasks. Even more, those under the age of 35 are twice as likely to do so (“The Future of Technology in Health and Health Care: A Primer,” 2018). Furthermore, studies have calculated that over 58% of Americans have used tools like Google and other domain-specific tools to support their health informatics search tasks—with search being one of the most important and central tasks in most health informatics activities (Demiris, 2016; Zuccon & Koopman, 2014).

Yet, search can be challenging, particularly for health informatics tasks that utilize large and complex document sets. For such tasks, health informatics tools may require the use of domain-specific vocabulary. Aligning with this vocabulary can be a significant challenge within health tasks, as they can involve a lexicon of intricate nomenclature, deeply layered relations, and lengthy descriptions that are misaligned with common vocabulary. For instance, one highly cited medical research paper defines the term “chromosomal instability” as “an elevated rate of chromosome mis-segregation and breakage, results in diverse chromosomal aberrations in tumor cell populations.” In this example, those unfamiliar with the defined term could find parsing its definition just as significant a challenge as the term itself (Gao, 2017). Thus, when communicating across vocabularies, users may struggle to describe the requirements of their search task in a way that is understandable by health informatics tools (Mehta & Pandit, 2018; Thiébaud & Cossin, 2019). To deal with this challenge, ontologies can be a valuable mediating resource in the design of user-facing interfaces of health informatics tools (Saleemi, Rodríguez, Lilius, & Porres, 2011). That is, ontologies

can bridge the vocabularies of users with the vocabulary of their task and its tools. Yet, the use of ontologies in user-facing interface design is not well established. Furthermore, health informatics tools that present a generalized interface, one that can support search tasks across any number of domain vocabularies and document sets, can allow users to transfer their experience between tasks, presenting users with information-centric perspectives during their performances rather than technology-centered perspectives (Fang, Pouyanfar, Yang, Chen, & Iyengar, 2016; Gibson, Dixon, & Abrams, 2015). For this, there is a need to distill criteria that can guide designers during the creation of ontology-supported interfaces for health informatics search tasks involving large document sets.

The goal of this paper is to investigate the following research questions:

- What are the criteria for the structure and design of generalized ontology-supported interfaces for health informatics search tasks involving large document sets?
- If such criteria can be distilled, can they then be used to help create such interfaces?

In this paper, we examine health informatics, machine learning, and ontologies. We then review leading research on health informatics search tasks. From this analysis, we formulate criteria for the design of ontology-supported interfaces for health informatics search tasks involving large document sets. We then use these criteria to contrast the traditional design strategies for search interfaces. To demonstrate the utility of the criteria in design, we will use them to structure the design of a tool, ONTSI (ONTology-supported Search Interface). ONTSI allows users to plug-and-play their document sets and expert-defined ontology files to perform health informatics search tasks. We describe ONTSI through a functional workflow and an illustrative usage scenario. We conclude with a summary of ongoing evaluation efforts, future research, and our limitations (Köhler et al., 2018).

## **5.2 Background**

In this section, we describe the concepts and terminology used when discussing ontology-supported interfaces for health informatics search tasks involving large document sets. We begin with background on health informatics. Next, we examine machine learning. We conclude with coverage of ontologies and their utility as a mediating resource for both human- and computer-facing use.

### **5.2.1 Health Informatics**

Health informatics is broadly concerned with emergent technological systems for improving the quality and availability of care, promoting the sharing of knowledge, and supporting the performance of proactive health and wellness practices by motivated individuals (Wickramasinghe, 2019). Initially, the need for expanded health and wellness services stemmed from rising population levels combined with the growing complexity of medical sciences. These issues made it challenging to maintain quality care within increasingly stressed medical systems (Carayon & Hoonakker, 2019). Thus, a central objective for health informatics is the development of strategies to tackle large-scale problems that harm trained medical professionals' ability to perform their tasks in a timely and effective manner.

For instance, tele-health services allowed doctors to practice remote medicine, providing care to those without local medical services. Another early innovation was standardized health care records, where patient records were given standardized encodings to provide an increased ability to track, compare, manage, and share personal health information (Demiris, 2016). Some examples of current research directions are the push for stronger patient privacy, personalized medicine, and the expansion of healthcare into under-served regions and communities (Brewer et al., 2020; Demiris, 2016; Gamache, Kharrazi, & Weiner, 2018; “The Future of Technology in Health and Health Care: A Primer,” 2018; Wickramasinghe, 2019).

The rising production and availability of health-related data has resulted in a growing number of data-intensive tasks within health. Both private and public entities like health industry companies, government bodies, and everyday citizens are turning to health informatics tools as they manage and activate their health data (“The Future of Technology in Health and Health Care: A Primer,” 2018). A growing number of health-related tasks involve searching document sets. During these tasks, the aim of the user is to use the information described within their document set to increase their understanding of a topic or concept. For example, a search task could be a practitioner searching the electronic health records of their patients, a member of the general public using public materials for their general health concerns, or a researcher performing a literature review [12,15,16]. In general, a search task involves the generation of a query based on an information-seeking objective. The computation systems of these tools then use this query within their computation systems to map and extract relevant documents out from the document set (Wu, Meder, Filimon, & Nelson, 2017). Powerful technologies like machine learning are increasingly being integrated within tools to help perform rapid and automated computation on document sets (Zuccon & Koopman, 2014). Yet, when taking advantage of these technologies, designers must be mindful of human factors when generating the user-facing interfaces of their tools, as a task cannot be performed effectively without direction from an empowered user (Carayon & Hoonakker, 2019).

### **5.2.2 Machine Learning**

Machine learning techniques are increasingly being utilized to tackle analytic problems once considered too complex to solve in an effective and timely manner (Talbot, Lee, Kapoor, & Tan, 2009). Yet, recent analysis (Endert et al., 2017; Hohman, Kahng, Pienta, & Chau, 2018; Yuan et al., 2020) on the human factors in machine learning environments have found that the current design strategies continually limit users’ ability to take part in the analytic process. More so, it has produced a generation of machine learning-integrated tools that are failing to provide users a complete understanding on how computational systems of their tools arrive at their results. This has significantly reduced users’ control and lowered the ability to achieve task objectives. In response, there is a growing desire to promote the “human-in-the-loop,” bringing the benefits of human reasoning back to the forefront of the design process (Jusoh, Awajan, & Obeid, 2020; Lytvyn, Dosyn, Vysotska, & Hryhorovych, 2020; Román-Villarán et al., 2019).

When considering the interaction loop of a machine learning-integrated tool, Sacha et al. (Sacha et al., 2016) present a five-stage conceptual framework: producing and accessing data, preparing data for tool use, selecting a machine learning model, visualizing computation in the tool interface, and users applying analytic reasoning to

validate and direct further use. Assessing this framework, a machine learning-integrated search tool must provide users with a functional workflow where:

1. Users communicate their task requirements as a query.
2. Users ask their tool to apply that query as input within its computational system.
3. The tool performs its computation, mapping the features against the document set.
4. The tool represents the results of the computation in its interface.
5. Users assess whether they are or are not satisfied with the results.
6. Users restart the interaction loop with adjustments or conclude their use of the tool.

Thus, a primary responsibility for users within machine learning environments is the need to assess how well the results of machine learning have aligned with their task objectives. A systematic review by Amershi et al. (Amershi, Cakmak, Knox, & Kulesza, 2014) suggests six considerations for the user's role in arbitrating machine learning performance:

1. Users are people, not oracles (should not be expected to repeatedly answer whether a model is right or wrong).
2. People tend to give more positive than negative feedback.
3. People need a demonstration of how machine learning should behave.
4. People naturally want to provide more than just data labels.
5. People value transparency.
6. Transparency can help people provide better labels.

### **5.2.3 Ontologies**

In search tasks involving large document sets, many challenges can arise that reduce performance quality, harm user satisfaction, and increase the time for task completion (Endert et al., 2017; Hohman et al., 2018; Yuan et al., 2020). Often, these challenges result from misalignment between the vocabularies used by the document sets, storage maintainers, interface designers, and users. For instance, Qing et al. (Zeng & Tse, 2006) outline the difficulties faced when translating between common and domain vocabularies in health tasks. They describe a study that found that up to 50% of health expressions by consumers were not represented by public health vocabularies (Zeng & Tse, 2006). Within the pipeline of a search task, both the human and computational system can only perform optimally if communication is strong (Arp, Smith, Spear, & American Journal of Sociology, 2015). Ontologies are representational artifacts that reflect the entities, relations, and structures of its domain. Ontologies are of three types: a philosophical ontology for describing and structuring reality, a domain ontology for structuring the entities and relations of a knowledge base, and a top-level ontology for interfacing between different domain ontologies (Arp et al., 2015). Ontologies provide the flexibility, extensibility, generality, and expressiveness necessary to bridge the gap when mapping domain knowledge for effective computer-facing and human-facing use (Saleemi et al., 2011). For this purpose, ontologies are increasingly being used within tools to help users perform their challenging search tasks (Bikakis & Sellis, 2016; Carpendale et al., 2014; Dou, Wang, & Liu, 2015; Ebner, 2015; Livingston, Bada, Baumgartner, & Hunter, 2015).

When creating an ontology, experts construct a network of entities and relations, which together yield various structures (Jakus, Milutinovic, Omerović, & Tomazic, 2013; Rector, Schulz, Rodrigues, Chute, & Solbrig, 2019). Ontology entities reflect the objects of the domain, like a phenotype in a medical abnormality ontology, a processor in a computer architecture ontology, or a precedent in a legal ontology (Tobergte & Curtis, 2013). In some ontologies, like the top-level ontology, Basic Formal Ontology, designers go as far as denoting qualities such as materiality, object composition, and spatial qualities in reality (Arp et al., 2015). Ontology entities are encoded with information about their role in the vocabulary, definitions, descriptions, and contexts, as well as metadata that can inform the performance of future ontology engineering tasks. Ontology relations are the links between ontology entities that express the quality of interaction between them and the domain as a whole (Katifori, Torou, Vassilakis, Lepouras, & Halatsis, 2008). When assessing ontology relations, *Arp et al.* (Arp et al., 2015) distinguish relations under the categories of universal–universal (dog “is\_a” animal), particular–universal (this dog “instance\_of” dog), and particular–particular (this dog “continuant\_parts” of this dog grouping). Domain ontology relations are realized through unique interoperability between ontology entities. For instance, an animal ontology may have an ontology entity reflecting the concept of a “human,” which may have the ontology relations “domesticates/is domesticated by” between it and the “dog” ontology entity.

After defining the entities, relations, and other features of an ontology, experts record their work in ontology files of standardized data formats like RDF, OWL, and OBO. These ontology files are then distributed amongst users. They can then be integrated into the computational and human-facing systems of tools for use during tasks. Some examples of current ontology use are information extraction on unstructured text, behavior modeling of intellectual agents, and an increasing number of human-facing visualization tasks such as decision support systems within critical care environments (Jusoh et al., 2020; Lytvyn et al., 2020; Román-Villarán et al., 2019).

## **5.3 Methods**

In this section, we describe the methods used for criteria formulation. We begin with a review of literature for health informatics search tasks. Based on the insights gained from this review, we distill a set of criteria. We then use these criteria to contrast traditional design strategies for interfaces of search tasks.

### **5.3.1 Task Review**

Here, we review some research on interfaces for health informatics search tasks. We used Google Scholar, IEEE Xplore, and PubMed to conduct an exhaustive search of articles and reviews published between 2015 and 2021. We have divided our findings into three sections. First, we explore research on health data, information management, and information-centric interfaces. This is followed by research discussing the types of search tasks and their use in structuring the design of interfaces for health informatics. Finally, we investigate the requirements for aligning vocabularies for health informatics search tasks.

### ***5.3.1.1 Health Data, Information Management, and Information-Centric Interfaces***

Health data is constantly generated, highlighted by reports that within just a year the US healthcare system created 150 new exabytes of data (Raghupathi & Raghupathi, 2014). Yet, the information that is expressed by this data, such as personal medical records, research publications, and consumer health media, is not useful unless it can be effectively understood and utilized by users. As such, it is critical to examine the challenges facing users when performing their tasks, and through this establish novel strategies for supporting the activation of health data.

Fang et al. (Fang et al., 2016) explore the pressing challenges for accessing health data under the four categories: volume, variety, velocity, and veracity. They find that the volume of health care data creates challenges in the management of data sources and stores. They describe that existing strategies are struggling, and that novel designs should be established for scaling data services. They explain that a variety of challenges come with the management of data characteristics, ranging from unstructured datapoints generated from sources like sensors, to structured data entities like research papers and medical documents. For this, they state that designers should concentrate on aligning with the characteristics of the information being encountered. Next, they explore the challenges of velocity, which involves the rate at which users require their data to move from source to activation within their task. They highlight novel research in the networking and data management space. Finally, they explore veracity challenges, such as the assessment and validation of data quality and the quality of information that the data may produce.

Gibson et al. (Gibson et al., 2015) provide a review of the evolving fields of health information management and informatics. They review the topics of data capture, digital e-record systems, aggregate health management, healthcare funding models, data-oriented evidence-based medicine, consumer health applications, health governance, personal health access, and genomic personalized medicine. Similar to Fang et al. (Fang et al., 2016), they note that the predominant work for health informatics should be concerned with presenting users with information-centric perspectives during their performances rather than technology-centered perspectives. More so, they describe that users in healthcare “must often navigate and understand complex clinical workflows to effectively ... capture, store, or exchange information.” In other words, task workflows are already complex; therefore, effective interfaces should promote information encounters that help users perform better, rather than engage in unrelated technical details.

From the above research, we distill the criterion: Designs should maintain an information-centric interface that is flexible with respect to the dynamic requirements of search tasks like veracity of data sources, variety of data types, and evolving needs of users for health informatics.

### ***5.3.1.2 Search Tasks and Structuring the Design of Interfaces for Health Informatics***

Russell-Rose et al. (Russell-Rose, Chamberlain, & Azzopardi, 2018) describe professional health workplace tasks. They find that the most prevalent types of search tasks are literature reviews for overviewing a topic, scoping reviews for rapidly inspecting the possible relevance of an information source, rapid evidence reviews for appraising the overall quality of a scoping review, and, finally, systematic reviews for exploring a topic in a robust manner.

During search tasks, users often lack the ability to perceive how their query decisions impact, relate, and interact with the document set. This is an important consideration for users who might want to adjust a query to better align with their information-seeking objectives. A further study by Russell-Rose et al. (Russell-Rose & Chamberlain,

2017) analyzes search strategies performed by healthcare professionals. They find that a large majority of participants have a general desire to utilize advanced search functionalities when available. This suggests that users are not hesitant to take advantage of resources that they believe help optimize their task performance. Hurdeman (Hurdeman, 2017) outlines that for this, a good course of action is to leverage query corrections, autocomplete, and suggestions. Yet, they find that such additions can be harmful if those features do not provide appropriate domain context. That is, resources must allow users to be contextually aware of how their query aligns with the contents of document sets, as well as the conditions of computational technologies used by interfaces.

In the same research, Hurdeman (Hurdeman, 2017) investigates complex tasks involving information search and information-seeking models when using multistage search systems. In this research, they explore requirements that designers must account for when supporting users. Challenging search tasks require users to learn about the searched domain, understand how their objectives align, and formulate their objectives into a way that can be used by their tool. In other words, query building requires users to be domain cognizant, as they must communicate information-seeking objectives in a way that is understood by the tool, yet also aligns with the information found within the document sets. Thus, a health informatics tool that supports search tasks should provide the opportunity for understanding the domain of the document set being explored.

Zahabi et al. (Zahabi, Kaber, & Swangnetr, 2015) describe a set of nine requirements for designers when considering how to design usable interfaces for health informatics search tasks, summarized as:

- Naturalness: The workflow of the system must present a natural task progression.
- Consistency: The parts of the system should present similar functional language.
- Prevent errors: Be proactive in the prevention of potential errors.
- Minimizing cognitive load: Align cognitive load to the requirements of the task.
- Efficient interaction: Be efficient in the number of steps to complete a task.
- Forgiveness and feedback: Supply proper and prompt feedback opportunities.
- Effective use of language: Promote clear and understandable communication.
- Effective information presentation: Align with information characteristics.
- Customizability/flexibility: The system should remain flexible to the task requirements.

Additional research by Dudley et al. (Dudley & Kristensson, 2018) reviews user interface design for machine learning environments. They provide a set of principles that can be used by designers:

- Make task goals and constraints explicit.
- Support user understanding of model uncertainty and confidence.
- Capture intent rather than input.
- Provide effective data representations.
- Exploit interactivity and promote rich interactions.
- Engage the user.

From this research, we can distill two criteria: First, designs should provide interaction loops that promote prompt and effective feedback opportunities for the user. Second, designs should provide representations that are natural and consistent to the requirements of the information source, the user, and the task.

### ***5.3.1.3 Aligning Vocabularies for Health Informatics Search Tasks***

When considering interfaces for health informatics search tasks, a major challenge for users is the need to overcome problem formulation deficiencies when encountering unfamiliar domains. This is because, according to Harvey et al. (Harvey, Hauff, & Elswailer, 2015), users have been found to consistently suffer from four major issues during the performance of search tasks:

- Difficulty understanding the domain being searched.
- An inability to apply their domain expertise.
- Lacking the capacity to formulate an effective search query within the interface that accurately reflects their information-seeking objective.
- Deficient understanding of how to assess results produced by search, to decide whether the search has or has not satisfied their objective.

Harvey (Harvey et al., 2015) shows that in domains with complex vocabularies, such as health and medicine, the disparity of potential users' prior knowledge is extreme. They find that non-expert users routinely do not possess enough domain knowledge to address their information-seeking needs. This can cause significant issues during query formulation. As a result, non-expert users must first step away from their tool to learn specialized vocabulary before they can begin query building. Both Soldaini and Anderson (Anderson & Wischgoll, 2020; Soldaini, Yates, Yom-Tov, Frieder, & Goharian, 2016) describe that this issue can still affect even experts. This is because experts often must make assumptions when attuning to their tool.

There is growing research targeting the generation and application of mediation resources to help reduce the communication gap while using health informatics tools. Zeng et al. (Zeng & Tse, 2006) investigate the development of consumer health vocabularies for reducing the discourse gap between lay people and medical information document sets. Furthermore, Soldaini et al. (Soldaini, Cohan, Yates, Goharian, & Frieder, 2015) explore the use of novel query computation strategies to improve the quality of medical literature retrieval during search tasks. In their quantitative study, they contrast models generated using combinations of algorithms, vocabularies, and feature weights, assessing the computational performance of different query reformulation techniques. The results of their study suggest "greatly improved retrieval performance" when utilizing combined machine learning and bridged vocabularies. More so, they provide insight regarding the quality of options that can support computational systems for health informatics search tasks.

From the above research, we can distill two criteria. First, designs should provide interactions that allow users to efficiently prepare, perform, assess, and adjust their machine learning to align with information-seeking objectives of search tasks. Second, designs should provide mediation opportunities that assist users in communicating information-seeking objectives into the domain-specific vocabulary of the document set.

### 5.3.2 High-Level Criteria

In Table 1, we provide five criteria based on the above review.

**Table 5-1 The criteria for guiding the design of interfaces for health informatics search tasks involving large document sets. For abbreviation purposes, design criteria will be referenced in the text as DC#, where # is its assigned number.**

| DC# | Design Criteria  |
|-----|--|
| DC1 | Provide an information-centric interface that shows flexibility towards the evolving needs of users and the dynamic requirements of search tasks like the veracity of data sources and variety of information types. |
| DC2 | Provide interaction loops that supply prompt and effective feedback for users during the performance of search tasks.  |
| DC3 | Provide natural and consistent representations that allow users to understand the constraints, processes, and results provided by the interface.   |
| DC4 | Provide interactions that allow users to efficiently prepare, perform, assess, and adjust their machine learning to align with the information-seeking objectives of search tasks.                                   |
| DC5 | Provide mediation opportunities that assist users in communicating and bridge their information-seeking objectives into the vocabulary of the document set.  |

### 5.3.3 Analysis of Traditional Interface Strategies for Health Informatics Search Tasks

We now assess the traditional design strategies for interfaces for health informatics search tasks. Wilson’s comprehensive Search User Interface Design (Wilson, 2011) provides a complete survey of the history and current state of search interfaces. Based on their survey, and in particular their discussion of input and control features within the modern search user interfaces, two base strategies and one extension strategy for search interfaces are realized: “structured” interfaces, “unstructured” interfaces, and, in extension, “query expansion” interfaces. Table 2 provides a summary of how the above criteria align with each interface strategy.

**Table 5-2 A summary matrix of alignment between the criteria and interface strategies. Full descriptions are found within their respective sections. “Strong” is assigned if a characteristic of the interface strategy promotes alignment with the requirements of the design criterion. “Weak” is assigned if a characteristic of the strategy does not promote alignment with the requirements of the design criterion. “Variable” is assigned if the interface strategy has the potential to align with the criterion; however, such an alignment is not innate and must be actively pursued.**


| DC# | Structured | Unstructured | Query Expansion |
|-----|------------|--------------|-----------------|
| DC1 | Weak       | Variable     | Strong          |

|     |          |          |          |
|-----|----------|----------|----------|
| DC2 | Strong   | Strong   | Variable |
| DC3 | Variable | Variable | Variable |
| DC4 | Variable | Strong   | Strong   |
| DC5 | Weak     | Weak     | Strong   |

### 5.3.3.1 Structured Interface Strategy

The structured interface strategy creates designs that regulate input during query building. This is achieved by maintaining heavily restricted input control profiles. Designers who implement the structured interface strategy into their interfaces presuppose a search task with specific expectations for input, bounding queries to a limited input profile. One common bounding technique is to constrain query lengths and limiting query content to a controlled set of terms (Zielstorff, 2003). This restricted scope is considered the sole acceptable input profile, and thus allows designers to generate interfaces that limit the possible range of inputs and restrict all inputs that fall outside of that range. Designers typically achieve this by using interface elements like dropdowns, checkboxes, and radio buttons instead of elements like text boxes with free typing. For example, Figure 1 depicts the PubMed Advance Search Builder, which implements the structured interface strategy in its design. This interface requires users to select specific query term types from a restricted list, which then guides user input (“PubMed,” n.d.).

Since input control is restricted, a strength of the structured interface strategy is that designers can use information characteristics to prescribe the full range of query formulations. This allows for the use of representational and computational designs that optimize for the expected characteristics of the restricted input profiles, per DC2. This strategy provides a designer-friendly environment that is hardened against unwanted queries, which, if effectively communicated in the design of result representations, could allow for alignment with DC3 and DC4. Yet, it can be challenging to designers to use structured interface strategies in a generalized setting. This is because when a document set is swapped, hardened approaches may not align with the information characteristics of the new document set. This negatively affects the flexibility of the interface, and in turn alignment with DC1. A potential weakness of the structured interface strategy is that it requires users to possess expertise on both the controlled vocabulary of the interface as well as the vocabulary of the document set being searched. If this is not known, user experience can suffer, drastically affecting alignment with DC5. Within the context of health informatics, such weaknesses reduce the users’ ability to effectively perform search tasks. This is because the controlled vocabularies within the health and medical domains demand significant expertise and result in numerous points of failure during the query formation process (Keselman, Browne, & Kaufman, 2008; McCray & Tse, 2003).




**Figure 5-1 PubMed Advanced Search Builder: an example of a structured interface strategy for a search task. In this use case, a query item was generated for a MeSH term for heart abnormalities, a completion date after August 8, 2015, and in the English language, with the publisher of Oxford University Press soon to be added. Source: Image generated on January 18, 2021, using the public web portal provided by the National Center for Biotechnical Information, <https://pubmed.ncbi.nlm.nih.gov/advanced/> (accessed on January 18, 2021).**

### 5.3.3.2 Unstructured Interface Strategy

The unstructured interface strategy creates designs that provide limited input regulation. Unlike the structured interface strategy, it provides an open input control that accepts most input profiles during query building. Designers who implement the unstructured interface strategy do so without presupposing particular input, only accounting for general user error. That is, this input can originate from anywhere, such as common vocabulary, rather than from a pre-determined set of terms provided by the designer. Often, this input is directed to a single interface element. Implementations of the unstructured interface strategy typically present a text box that allows users to freely type their own text into the interface. These implementations will perform some input processing prior to use; however, the presentation of this processing to users is usually limited to correcting typographical errors rather than semantic ones. For example, the interface of Google aligns with the unstructured interface strategy, presenting users with an open, text-box input control without domain-specific assumptions or requirements. Of course, Google's computational systems use extensive processing between receiving input from users and presenting the results of computation back to users (Luo, Wu, Gopukumar, & Zhao, 2016). Yet, users themselves are not informed of how their results came to be, even after changing to Google Instant (Qvarfordt, Golovchinsky, Dunnigan, & Agapie, 2013). Another example of an implementation of the unstructured interface strategy is WebMD's search interface. This interface processes a free-text input with basic sanitization techniques before generating features for its search engine system, as depicted in Figure 2.

A strength of the unstructured interface strategy is that it supports the use of any vocabulary during query building, allowing for the natural activation of common vocabulary during task performance, in alignment with DC4. Additionally, this removes the requirement for users to possess input expertise and control profiles that typically come with a structured interface strategy, per DC1 and DC4. Designers can still implement prompt and effective feedback during task performance, thereby supporting DC2. If the constraints, processes, and results of their task performance are effectively communicated in result representations, DC3 and DC4 can be well supported. However, by allowing for the direct use of common vocabulary in lieu of a presupposed controlled vocabulary, the unstructured interface strategy suffers where the structured interface strategy excels. That is, poor implementations of the unstructured interface strategy can produce interfaces that do not provide mediation for users to translate their common vocabulary into the domain-specific vocabulary. In doing so, users are not being helped in understanding how their query building has impacted their search performance. For example, these poorly implemented interfaces may take input literally and bring users directly to a result page without providing context as to how the results were found, negatively affecting DC1 and DC5. This potential for promoting weak alignment between user and information source can lead to a significant drop in the quality of search performance. This can be an especially important requirement to address for health informatics interfaces, as it has been found that users routinely struggle to craft effective query terms during their health-related search tasks (Jimmy, Zuccon, & Koopman, 2018).




**Figure 5-2 WebMD Search Interface: an example of an unstructured interface strategy for a search task. In this use case, the free-text query “heart condition” was generated. Source: Image generated on January 18, 2021, using the public web portal provided by WebMD, [https://www.webmd.com/search/search\\_results/default.aspx?query=heart%20condition](https://www.webmd.com/search/search_results/default.aspx?query=heart%20condition) (accessed on January 18, 2021).**

### ***5.3.3.3 Query Expansion Interface Strategy***

The query expansion interface strategy is an extension of both the structured and unstructured interface strategies. That is, this strategy expands by adding mediation opportunities to bridge the vocabulary of the user with the vocabulary of the document set both within the representational as well as the computational systems (Azad & Deepak, 2019). These mediation opportunities are typically implemented within two parts of the interaction loop. The first is during input, where mediating opportunities present during query building. Often, these mediation opportunities come as cues that suggest to users how their common vocabulary could align with the vocabulary of the domain, and visa-versa. An example of an implementation of the structured-like query expansion interface strategy is WebMD's Symptom Checker, shown in Figure 3. This example interface goes through a series of controlled stages of query building that are structured by numerous opportunities for mediation. The second is during the processing prior to document set mapping. Like other strategies, a system can apply natural language processing techniques to the input, where the text string provided as input is tokenized into its parts. From this, the system sanitizes token parts to remove trivial tokens like the stop words "the," "a," and "an," and any remaining tokens are then inserted as features in search engine systems. In more complex systems, additional sanitization techniques can be used (Jimmy, Zuccon, Palotti, Goeuriot, & Kelly, 2018). Yet instead of immediately inserting the remaining tokens as features into the computational systems, the query expansion interface strategy builds upon the input profile by injecting insight provided by mediating resources, such as related terms, synonyms, and other expansion opportunities (Capuano, Longhi, Salerno, & Toti, 2015). In other words, these systems utilize mediating resources to computationally expand the query. Some examples of mediating resources are knowledge bases like WordNet and Wikipedia, and ontologies like The Human Phenotype Ontology (Azad & Deepak, 2019; Köhler et al., 2018; "The Human Phenotype Ontology," 2020).

A strength of the combined approach of the query expansion interface strategy is its strong efforts to eliminate the weaknesses associated with the structured and unstructured interface strategies while still maintaining their strengths. That is, by allowing the continued use of common vocabulary during the process of query building, users can have higher confidence about what the interface is asking of them, and what they are telling the interface to do, helping with DC4 and DC5. Furthermore, by integrating the use of mediating resources like ontologies, designers can demonstrate to users the quality of their query building and how their vocabulary decisions affect the performance of their search tasks, supporting DC2 and DC3 (Lüke, Schaer, & Mayr, 2012). Yet, with the added complexities of query expansion, computational systems may be required to perform more work before arriving at a final set of search results. Therefore, designers of systems taking advantage of query expansion should consider the impact on performance and responsiveness and counteract them to maintain alignment with DC2. For the query expansion interface strategy to be successful, designers must clearly communicate to users how exactly their query building has affected their search. If this communication is not provided, it can leave users confused regarding how their decisions have affected their search and can make it challenging for them to assess task performance, negatively affecting DC2. Such limitations may not provide optimal alignment in communication between the system, the user, and the information resource (Jimmy, Zuccon, & Koopman, 2018). That is, if a selected mediating resource does not provide an effective mapping between vocabularies, then query expansion can weaken the quality of search tasks. To address this challenge, designers can utilize user-supplied ontologies, as per DC1. This provides users the freedom to select mediating

resources that they believe can best support their task performance, rather than being restricted to a tool-provided mediating resource. A user study by Jay et al. (Jay et al., 2016) compares users as they perform the same task set using two interfaces, one with a structured multiple variable input profile, the other with an unstructured single variable input profile. In this study, they find that users felt their needs and expectations were better fulfilled using the single-input profile, performing their tasks quicker, with more ease of use and learnability, and with a higher appraisal of results. Designers must carefully select how they activate query expansion such that it addresses the needs of the task, the information, and the user.



**Figure 5-3** WebMD’s Symptom Checker: an example of a structured-like query expansion interface strategy for a search task. In this use case, users are guided along a series of query building opportunities, allowing them to enter various symptoms and personal health criteria while aligning their personal vocabulary with the information resource vocabulary. Source: Image generated on January 18, 2021, using the public web portal provided by WebMD, <https://symptoms.webmd.com/> (accessed on January 18, 2021).

## 5.4 Results

In this section we describe ONTSI, a generalized ontology-supported interface for health informatics search tasks involving large document sets created using the above-discussed criteria. We outline how the criteria were used to structure ONTSI’s design. We then discuss the technical scope of ONTSI, concluding with ONTSI’s functional workflow.

### 5.4.1 Design Scope

Table 3 highlights the role of each criterion in the design of ONTSI.

**Table 5-3 The role of each criterion within the design of ONTSI. The incorporation of these criteria in ONTSI’s implementation is discussed within the workflow and usage scenario.**

| DC# | ONTSI   |
|-----|---|
| DC1 | ONTSI leverages powerful third-party computational technology. Specifically, pre-built machine learning packages like SciKit-Learn are integrated within ONTSI, and highly optimized indexing is provided by The Apache Software Foundation’s Solr product (“Solr Cloud,” 2020). Additionally, ONTSI’s interface provides users with clear text-based alerts, which reflect their current performance status.   |
| DC2 | ONTSI supports an iterative interaction loop to allow users perform repeated sets of search tasks. That is, within iterative interactions, users can save the results they regard relevant in a persistent location within the tool, while still allowing further performances to occur.  |
| DC3 | ONTSI provides visual representations to help analyze and judge the relevance of search results.  |
| DC4 | ONTSI utilizes modern visualization and computational technologies like D3.js to provide powerful interaction opportunities.  |
| DC5 | ONTSI supports the use of a common vocabulary during query building using the query expansion strategy. Specifically, when using ONTSI, users upload both a document set and an ontology file, which are then integrated into the workflow of the computational systems of ONTSI. Users can interact with a search textbox that allows for unstructured text input. ONTSI provides domain-specific vocabulary suggestions that can assist users in guiding their performance and promote alignment between their vocabulary and domain-specific vocabulary. |

### 5.4.2 Technical Scope

ONTSI is developed as a web-based tool that provides a generalized, plug-and-play support of user-supplied ontology files and document sets. That is, ONTSI allows for the uploading of ontology files, either individually or within a .zip compressed file, as well as any compressed document set in the ZIP format. ONTSI then processes and


indexes their contents for use within the interface. ONTSI's front end uses the latest HTML5, CSS, and JavaScript technologies, allowing for cross-browser (i.e., Firefox, Chrome, Opera) and cross-platform support. The D3.js JavaScript library is used to create the visualization and interaction experiences found throughout the front end of ONTSI (Bostock, 2016). ONTSI's back-end technology is developed using a custom Python-based computational server that maintains data transfer and machine learning APIs, and with the use of Apache's Solr system as the search indexer and engine ("Solr Cloud," 2020). The current ONTSI system maintains support for the live uploading of well-formed ontologies in the Ontology Web Language (OWL) format.

### **5.4.3 Functional Workflow of ONTSI**

ONTSI encompasses several subsystems and subviews within its workflow. Recalling the workflow description of a machine learning-integrated search tool, ONTSI allows:

1. Users to communicate their task requirements as a query within its Upload and Search subview.
2. Users to ask their tool to apply that query as input within its computational system within its Search Subview.
3. The tool to perform its computation, mapping the features against the document set within its ONTSI server and Solr server.
4. The tool to represent the results of the computation in its interface within its Result List and Result Item subviews.
5. Users to assess whether they are or are not satisfied with the results within its Result List, Result Item, and Saved List subview.
6. Users to restart the interaction loop with adjustments within the Upload and Search subviews or conclude their use of the tool.

We will now describe the overall functional workflow of ONTSI and its parts, as depicted in Figure 4.



**Figure 5-4** Depiction of the functional workflow of ONTSI. Labeled arrows reflect the steps of the interaction loop. Brown boxes represent the processes performed within the back-end computation systems. Blue boxes represent the object types that persist within the browser and external index database storage. Pink boxes represent the back-end computational systems of ONTSI. Yellow boxes represent the various subviews within the front end of ONTSI. The green box represents the types of interactions that can be conducted with the system.

### 5.4.4 Front-End Subviews

ONTSI consists of a series of interconnected subviews, shown in Figures 4 and 5. We will now describe the functional workflow of each subview.

The screenshot displays the ONTSI interface with the following subviews:

- (a) Upload:** A button with an upload icon.
- (b) Search:** A search bar containing the text "difference between fatigue and sleepiness x" and a "Run" button.
- (c) Result List:** A list of search results, each with a "44.00% Relevant" badge and a document snippet. Navigation buttons "Front", "Previous", "Page 1 of 500 (Documents 1-20)", "Next", and "End" are visible.
- (d) Result Item:** A detailed view of a document titled "Distinguishing sleepiness and fatigue: focus on definition and measurement." It includes a summary and a snippet of the document content.
- (e) Saved List:** A list of saved documents, each with a "33.00% Relevant" badge and a document snippet. Navigation buttons "Front", "Previous", "Page 1 of 500 (Documents 1-20)", "Next", and "End" are visible.

Figure 5-5 The overall view of ONTSI in full use and outlined coverage of its five subviews: Upload subview, partial (a); Search subview (b); Result List subview (c); Result Item subview (d); and Saved List subview, partial (e).

#### 5.4.4.1 Upload Subview

The Upload subview supports the plug-and-play of user-supplied ontology files and document sets. This subview can be found at the top left of ONTSI, Figure 5a. When clicked, the upload button opens a file selection window. The window limits uploading to valid ontology files under the OWL ontology format and the .zip compression format. When a compressed file is uploaded, it is inspected for OWL files. This allows the upload system to not only take in individual OWL ontology files, but also sets of OWL files that are combined in a compressed format. Ontology file contents are put through a custom OWL to JSON processor, and then indexed into a local storage

system within the browser memory. If it is a document set, it is transferred to the back end ONTSI server. Once at least one ontology file and one document set are uploaded, the Search subview and the system become active.

#### ***5.4.4.2 Search Subview***

The Search subview facilitates query building using an ontology-supported unstructured-like query expansion strategy. Three points of interaction are maintained: Query Input, the Run button, and the Clear button. The Search subview is located to the right of the Upload subview at the top center, Figure 5b, and becomes available for interaction after the requirements of the Upload system are fulfilled.

Query Input is a text input box. As text is typed, ONTSI cross-references that text against the uploaded ontological content for mediation opportunities. If found, those mediations are provided within an expanding dropdown. When a user values a suggested mediation, it can be selected and locked in as a query term. If none are desired, they can be ignored. When a user is satisfied with their own typed text, it can also be added. Each query term is depicted with the text of the term and a removal interaction, represented by a trailing “x” button. If multiple terms require removal, this can be done either with individual removal actions, repeated backspacing actions from the keyboard, or the red “trash can” button, which clears all query terms.

When at least one query term has been entered, the green “Run” button becomes active. This initiates the performance of computation on the uploaded document set using the ontology file for query expansion. Query terms are collected and sent to the back end ONTSI server system. The Result List subview updates when the computation is complete.

#### ***5.4.4.3 Result List Subview***

The Result List subview provides a paged listing the search results. The Result List subview is found directly under the Search, Upload, and Saved List subviews, Figure 5c.

Once a search is performed, the Result List subview changes from an informational alert to the results of a search. The list itself is bounded above and below by buttons and text that describe and support paging interactions. Specifically, the buttons and text describe information about the current page position, the number of pages used to divide the document set, and the number of documents in the current page, and allow for various navigation interactions on the pages.

The search results are sorted by their relevance calculation generated during clustering, such that the results assigned to document clusters that have the highest predicted relevance rating are prioritized. Then, the list is paged. Instantaneous navigation between pages is provided. Color-coded relevance ratings accompany each document, ranging from best to worst within a green–red color spectrum. Each result represents the document title with annotations highlighting terms or phrases that are believed to align with the provided query terms. A button is also provided that allows the user to access additional document content and open the document for deeper inspection. Finally, each result has a “pin” button, which allows for the saving of documents for future use within the Saved List subview.

#### ***5.4.4.4 Result Item Subview***

The Result Item subview provides document-level information. This allows users to rapidly assess the content of individual documents during their search task. When a user selects a document within the Result List subview, ONTSI will request the full document content of that result from the Solr server using its HTTP-based API. Query terms are then used by annotation services within the Solr server to wrap HTML-based annotation tags into the document content, which is then returned to the Result Item subview. When a document is selected for inspection, the Result Item subview expands that document in place within the Result List subview, pushing down trailing items, Figure 5d.

The content of the selected document is represented in the following order: the file name of the document within the uploaded document set, the full document title, and a summarized version of the document content. The summarized version of the document content restricts the document to the passages of content that surround or have associations with the query terms provided during query building. Terms are highlighted through capitalization and with bolded font. In addition, the Result Item subview provides a dropdown at the top right, which collects all web links found within the document content for quick access. Any number of documents can be opened within the Result Item subview for comparison.

#### ***5.4.4.5 Saved List Subview***

Each result within the Result Item subview includes a green “pin” button, which saves documents for future reference. ONTSI collects these saved documents within the Saved List subview. The Saved List subview can be accessed at the top right of ONTSI’s overall view, directly to the right of the Search subview. There, a green “pin” Saved List can be found that allows us to request ONTSI to open the Saved List modal, Figure 5e. Upon request, the Saved List modal displays saved documents. Here, documents can be recalled, removed from the list, or copied for external use.

### **5.4.5 Back-End Systems**

ONTSI consists of two back-end systems that support the various front-end subviews and their controlling logic: the ONTSI server and the Solr server. Through their use, heavy computation is moved away from the browser and into dedicated computational systems. This allows for a reduction in computational overhead within the browser to improve response times and allows ONTSI to access computational technology that is not readily supported in the browser.

#### ***5.4.5.1 ONTSI Server***

The ONTSI server is created using the Python-based Flask framework. It exposes an API supporting communication between the various systems of ONTSI. The API satisfies two major roles: preparing the uploaded document set for indexing within the Solr server and handling machine learning requests for search tasks.

When a document set has been signaled for upload within the Upload subview, it is packaged and sent through the API of the ONTSI server. Incoming document sets are assessed and provided a suitable decompression

algorithm. Next, for each document within the document set, the ONTSI server assesses the encoding of that file (e.g., UTF-8, UTF-16, PDF, etc.). Based on this assessment, a suitable transcription algorithm is applied to that document. The indexing process for the Solr server is a pull interaction, so documents are stored in a static location from which they can be pulled. Therefore, the documents are sanitized, packaged, and then inserted into a temporary PostgreSQL database. The ONTSI server then requests the Solr server to begin indexing the new document set.

When a search task is initiated, the request is sent to the API of the ONTSI server. There, requests are read for settings like the clustering algorithm, the document set being searched, and query specifications. The ONTSI server then prepares the machine learning environment. Next, ONTSI performs query expansion. This involves a set of natural language preprocessing steps on the query and its individual query items, such as tokenization and the application of stop word limiters. Then, each query item is examined against the provided ontology file for mediating opportunities alongside a complete synonym ring analysis on each query item using WordNet. The original query terms and their associated ontology and synonym terms are then packaged together. These packages are then applied during the performance of unsupervised K-means clustering computation from SciKit-Learn, a third-party machine learning suite. The computed weighting characteristics of clusters are then propagated back as a package of clusters and their associated documents for the ONTSI front end for use within its various subviews. We include a pseudocode representation of these steps in Figure 6.

---

**Algorithm 1: CLUSTERING** pseudocode between ONTSI, ONTSI Server, and Solr Server

---

```
Input: A set Q of user inputted queries
Output: Signal to update interface with cluster assignments
1 targets ← chain(Q).unique().difference(getStopWords())
2 documents ← getDocuments()
   /* Prepare bag of words using target, related entities, and
   their generated WordNet synsets */
3 for i = 0 to targets.length do
4   target = targets[i]
5   targetCoverage ← target + target.getDirectlyRelatedEntities()
6   targetSpread[target] ←
     targetCoverage + targetCoverage.getWordNetSynsets()
7   targetSpread[target] ←
     targetSpread[target].unique().difference(getStopWords())
   /* Gather counts from pre-indexed documents, then fit and
   predict clusters using Scikit.Learn KMeans clustering */
8 documentCounts ←
   getIndexesFromSolrAPI(targetSpread, documents).scaleRange(0, 1)

9 reducedPCA ← SciKitLearn.PCA(nComponents =
   2).fit_transform(documentCounts)
10 kmeansPCA ← SciKitLearn.KMeans(init = 'k-
   means ++', nClusters = 7, nInit = 10)
11 clusterAssignments ← kmeansPCA.fit_predict(reducedPCA)
12 for i = 0 to targets.length do
13   target = targets[i]
14   for j = 0 to clusterAssignments.length do
15     cluster = clusterAssignments[j]
16     yPred = cluster.yPred[target]
     /* Generate weighting scale using x5 multiplier */
17     clusterAssignments[j].weighting[target] ←
       generateClusterWeighting(yPred)
18 return updateInterface(clusterAssignments)
```

---

Figure 5-6 Pseudocode of clustering spanning the workflow of ONTSI (front end), ONTSI server, and Solr server.

### 5.4.5.2 Solr Server

ONTSI uses Solr, a third-party document indexing software developed by The Apache Software Foundation. Solr is a scalable indexing system that provides a valuable array of features like a REST-like API supporting many HTTP-based communication interfaces. Solr also provides and a wide range of customizable settings and schemas that supports any number of storing, searching, filtering, analysis, optimization, and monitoring tasks. For a more information regarding Solr and its various permutations, seek out their official website and documentation (“Solr Cloud,” 2020).

A cloud-based permutation of the Solr server is used to handle the indexing and serving of uploaded document sets. Indexing occurs when a request is made to the Solr server from the ONTSI server. The Solr server schema will seek out the location of the temporary PostgreSQL database hosted by the ONTSI server, extract all new documents not already indexed, and apply a processing schema on those documents for indexing. Then, signals are sent out to the relevant ONTSI systems. Solr also handles serving requests when ONTSI requires document content, either at the metadata level when loading the Result List subview, or full, annotated document

content in the Result Item subview. Requests are communicated to the Solr server through its HTTP-based API. The Solr server then handles the request, packages the results under the conditions specified in the request, and returns its response.

## 5.5 Usage Scenario

In this section, we provide a health informatics search task scenario using ONTSI. We begin with a description of the user profile, as well as the ontology file and document set in the usage scenario. We then present the usage scenario.

### 5.5.1 User Profile

The user profile we select here is that of a health stakeholder, a researcher within a professional workplace setting performing a scoping review as an information-seeking objective. A scoping review is concerned with establishing an initial idea of the amount of information on a topic within a document set (Russell-Rose et al., 2018). The user has a general level of knowledge, typical of other health stakeholders. For instance, the user understands and can communicate phenotypic abnormalities like a broken leg, light-headedness, or loss of vision. The user understands how to perform typical actions on the interface like clicking, typing, and saving, but does not possess knowledge of the technical concerns typical of backend computational technologies.

The objective of the user is to learn whether there are any documents within a document set that are relevant to a research question. Let us assume the user's research question is, "How does chromosomal instability drive tumor progression?" We selected this question from recently published materials on topical examples within the health domain, using "The 150 most important questions in cancer research and clinical oncology series," published in 2017 by the *Chinese Journal of Cancer* (Gao, 2017).

### 5.5.2 Ontology File and Document Set

ONTSI requires the user to upload an ontology file and a document set. We used the Human Phenotype Ontology (HPO) in the usage scenario. We selected HPO because of its high complexity resulting from its exhaustive and expert-defined domain coverage of terms and their relationships. HPO is a controlled and standardized vocabulary encoding human disease and phenotypic abnormalities. It also includes annotations in bioinformatics, biochemistry, and human genetics. HPO is an active ontology, consisting not only of over 11,000 terms, but also over 110,000 disease annotations (Köhler et al., 2014). An example of an HPO term is "blindness," which possesses a superclass of "visual impairment," a subclass of "congenital blindness," and is annotated to be associated with a variety of diseases, such as a variant of colorblindness termed Achromatopsia 2 (Köhler & Robinson, 2016). Each HPO term describes attributes such as names, conceptual definitions, ontology indexing, term synonyms, class relationships, logical definitions, and expert commentary, to name a few. For additional details on the Human Phenotype Ontology, see (Köhler et al., 2018; "The Human Phenotype Ontology," 2020).

The National Library of Medicine's PubMed is selected as the document set within the usage scenario. PubMed is chosen because of its prominence within the health domain, maintaining more than 30 million citations

used within a wide scope of literature and active research endeavors. Data availability limits this usage scenario to a subset of PubMed representing 10,000 document entries. These entries maintain the document title, abstract, and various metadata like authors, published date, and keywords (“PubMed,” n.d.).

### 5.5.3 Usage Scenario

The user loads ONTSI, finding it in its initial state, as seen in Figure 7.




Figure 5-7 The initial state of ONTSI.

The user uploads their document set and ontology file by clicking the Upload button, activating the Upload subview, as seen in Figure 8. After confirming a selection, the upload process begins.




Figure 5-8 Selecting a document set and ontology file for upload.

After the upload process is complete, the user begins typing the research question “How does chromosomal instability drive tumor progression?” into the textbox, finishing with a click of the “Run” button. In response, ONTSI provides

the results of its computation, as seen in Figure 9. It presents the first page of 500 pages of documents, which totals 20 document entries. There are percentages to the left of each entry that use text and color to annotate the relevance of each document. This scalar is based on the cluster weightings within the dimensional space of the document set, where a 100% would be produced by documents within a cluster that aligns with every input feature. The scalar maintains a color scale between red and green, where red is at the zero point and green at 100%. For instance, at the top of the first page there are five documents that present an orange 45.25% relevance rating. Looking at these documents, the user scans the titles of the documents, where some have terms within their titles that relate to the research question.

The screenshot shows a search interface with a search bar containing the query "how does chromosomal instability drive tumor progression?". Below the search bar are navigation buttons: "Front", "Previous", "Page 1 of 500 (Documents 1-20)", "Next", and "End". The main area displays a list of 20 search results, each with a relevance percentage and a title snippet. The relevance percentages are color-coded: orange for 45.25%, red for 32.25%, and red for 15.25%. Each result has a dropdown arrow and a green checkmark icon.

| Relevance       | Title Snippet  |
|-----------------|--|
| 45.25% Relevant | Cancer morphology, carcinogenesis and genetic <b>instability</b> : a background.                       |
| 45.25% Relevant | Kaposi's sarcoma-associated herpesvirus-encoded latency-associated nuclear antigen induc...            |
| 45.25% Relevant | Complex formation of Plk1 and INCENP required for metaphase-anaphase transition.                       |
| 45.25% Relevant | STK15 gene overexpression, centrosomal amplification, and <b>chromosomal instability</b> in ...        |
| 45.25% Relevant | Incidence of chromosome abnormalities in the Sultanate of Oman.  |
| 32.25% Relevant | Monte Carlo charged-particle tracking and energy deposition on a Lagrangian mesh.                      |
| 32.25% Relevant | Impulse absorption by tapered horizontal alignments of elastic spheres.                                |
| 32.25% Relevant | Variable practice in learning the forehand <b>drive</b> in tennis.                                     |
| 32.25% Relevant | Electromechanical properties of high coupling single crystals under large electric <b>drive</b> and... |
| 32.25% Relevant | How does the sound pressure generated by circumaural, supra-aural, and insert earphones...             |
| 32.25% Relevant | Superscaling in charged current neutrino quasielastic scattering in the relativistic impulse a...      |
| 32.25% Relevant | Current- <b>drive</b> efficiency in a degenerate plasma.   |
| 32.25% Relevant | Three-junction SQUID rocking ratchet.  |
| 32.25% Relevant | Nonlinear amplifier and frequency shifter using a tunable periodic <b>drive</b> .                      |
| 32.25% Relevant | Rapid shear viscosity calculation by momentum impulse relaxation molecular dynamics.                   |
| 15.25% Relevant | Cycles of chromosome <b>instability</b> are associated with a fragile site and are increased by d...   |
| 15.25% Relevant | Dynamical <b>instability</b> and domain formation in a spin-1 Bose-Einstein condensate.                |
| 15.25% Relevant | Ventricular dilation as an <b>instability</b> of intracranial dynamics.                                |
| 15.25% Relevant | Alterations in DNA repair gene expression under hypoxia: elucidating the mechanisms of h...            |
| 15.25% Relevant | Arthroscopic dorsal radiocarpal ligament repair.   |

At the bottom of the results list, there are navigation buttons: "Front", "Previous", "Page 1 of 500 (Documents 1-20)", "Next", and "End".

Figure 5-9 The results of a search task using the query, “How does chromosomal instability drive tumor progression?”

Some documents at the top of the results could align with the user’s research question. To explore further, the user selects a few of the top documents, generating additional document information for inspection, as seen in Figure 10.

Doing so, the user encounters a summarized version of their selected documents, which provides metadata and abstracts annotated with words and phrases related to the research question.


The screenshot displays the ONTSI search interface. At the top, a search bar contains the query "how does chromosomal instability drive tumor progression?". Below the search bar, navigation buttons for "Front", "Previous", "Page 1 of 500 (Documents 1-20)", "Next", and "End" are visible. The main content area shows three search results, each with a "45.25% Relevant" tag. The first result is titled "Cancer morphology, carcinogenesis and genetic instability: a background." and includes a summary of the document content. The second result is titled "Kaposi's sarcoma-associated herpesvirus-encoded latency-associated nuclear antigen induces chromosomal instability through inhibition of p53 function." and also includes a summary. The third result is titled "Complex formation of Plk1 and INCENP required for metaphase-anaphase transition." Each result has a "pin" icon on the right side, which the user is likely to click to save the reference.

Figure 5-10 ONTSI after opening the documents “Cancer morphology, carcinogenesis and genetic instability: a background” and “Kaposi’s sarcoma-associated herpesvirus-encoded latency-associated nuclear antigen induces chromosomal instability through inhibition of p53 function.”

The user estimates that these top documents may align with their research question. Therefore, they click on the green “pin” button found at the rightmost point of each document to save their reference for future retrieval from the

document set. These references are accessible by clicking the green “pin” button found at the top right of ONTSI to open the Saved List subview.

Although the user has now encountered some documents relevant to their research question, they choose to continue searching. This time, the user decides to take advantage of mediation opportunities when building their query. After closing the Saved List subview, the user begins a new search. After assessing the important words in their research question, the user types in the term “chromosomal.” At the point that they have typed “chromo,” they are presented with mediation opportunities, as seen in Figure 11. They inspect these mediation opportunities and add phenotypic terms that align with their research question.



**Figure 5-11 ONTSI while the user is presented with mediating opportunities from the expert-defined Human Phenotype Ontology.**

With the aid of mediation, the user builds a three-item query consisting of “abnormal chromosome morphology,” “chromosomal instability,” and “tumor progression.” After asking ONTSI to run with this query, the user encounters a set of results different from the one produced by their earlier search, as seen in Figure 12. Notably, an increased set

of 10 documents at a 48.75% rating is encountered. Looking at these documents, the user notices some that are familiar, such as the saved “Genetic instability in human tumors.”

The screenshot shows the ONTSI search interface. At the top, there is a search bar with the terms "abnormal chromosome morphology", "chromosomal instability", and "tumor progression". Below the search bar, there are navigation buttons: "Front", "Previous", "Page 1 of 500", "Next", and "End". The main content area displays a list of search results, each with a relevance percentage and a document title. The results are as follows:

| Relevance       | Document Title   |
|-----------------|--|
| 48.75% Relevant | Microsatellite <b>instability</b> in multiple nonfamilial malignancies.  |
| 48.75% Relevant | Rayleigh-Taylor <b>instability</b> in elastic solids.  |
| 48.75% Relevant | Cancer <b>morphology</b> , carcinogenesis and genetic <b>instability</b> : a background.                             |
| 48.75% Relevant | The influence of vertebral <b>instability</b> on peridural circulation and concomitant peridural fibrosis formation. |
| 48.75% Relevant | Kaposi's sarcoma-associated herpesvirus-encoded latency-associated nuclear antigen induces <b>chromoso...</b>        |
| 48.75% Relevant | STK15 gene overexpression, centrosomal amplification, and <b>chromosomal instability</b> in the absence of...        |
| 48.75% Relevant | Arthroscopic dorsal radiocarpal ligament repair.   |
| 48.75% Relevant | Cycles of <b>chromosome instability</b> are associated with a fragile site and are increased by defects in DN...     |
| 48.75% Relevant | Genetic <b>instability</b> in human <b>tumors</b> .  |
| 48.75% Relevant | Orthostatic <b>instability</b> in a population-based study of chronic fatigue syndrome.                              |
| 41.75% Relevant | [Studies of treatment strategy and prognosis on acute myeloid leukemia with <b>chromosome</b> 8 and 21 tr...         |
| 41.75% Relevant | Multicolor fluorescence in situ hybridization (SKY) in mycosis fungoides and Sézary syndrome: search for...          |
| 41.75% Relevant | Shwachman syndrome as mutator phenotype responsible for myeloid dysplasia/neoplasia through karyot...                |
| 41.75% Relevant | Robust estimation of experimentwise P values applied to a genome scan of multiple asthma traits identifi...          |
| 41.75% Relevant | [Research progress in human artificial <b>chromosomes</b> (HACs) and the potentials in application].                 |
| 41.75% Relevant | The role of foetal red blood cells in protecting cultured lymphocytes against diepoxybutane-induced <b>chro...</b>   |
| 41.75% Relevant | Y- <b>chromosome</b> diversity is inversely associated with language affiliation in paired Austronesian- and Pa...   |
| 41.75% Relevant | Use of plant genotoxicity bioassay for the evaluation of efficiency of algal biofilters in bioremediation of t...    |
| 41.75% Relevant | A scan of <b>chromosome</b> 10 identifies a novel locus showing strong association with late-onset Alzheimer ...     |
| 41.75% Relevant | [Clinical and experimental study of 7 cases of acute lymphoblastic leukemia with dic(7;9) (p11;p11)].                |

Figure 5-12 ONTSI after running a new search after the user took advantage of mediation opportunities presented in the generation of the query items.

From this listing, the user selects two new documents for deeper inspection, as seen in Figure 13. They notice that terms such as “morphology” and “neoplasm” are now being highlighted within the document annotations. The adjusted query based on mediation opportunities has helped promote documents that align with their research question. In this case, the user finds value in the two documents, so they are saved.

Before concluding the search task, the user can upload a different ontology file to investigate how alternate vocabularies may bridge them to their document set. Their encounters could have also allowed them to make the assessment that the document set may not be best to help with their research question. If that is the case, they may

upload a different document set, performing another scoping review. In any case, ONTSI provides a search task interface that has been generalized to support plug-and-play capabilities for user-provided ontology files and document sets, allowing users to customize the interface to match their search task objectives.




Figure 5-13 ONTSI after the user has selected new document entries for deeper inspection.

## 5.6 Discussion and Conclusions

### 5.6.1 Evaluation of ONTSI

We have conducted ongoing, formative, task-driven user evaluations of ONTSI. These evaluations were informally conducted with a few people associated with our research lab; they have provided initial insights into how ontology-supported interfaces for health informatics can support users to perform elaborate search tasks involving large document sets. In these evaluations, we asked the users to perform a targeted set of tasks, such as researching questions outlined in the presented usage scenario. Initial sessions provided general insight into how users search and how ontologies can help mediate such tasks. From these sessions, we have learned a few things, which are itemized below:

- Users are able to quickly transfer their experiences with previous interfaces to use ONTSI (e.g., A, B).

- Users are capable of utilizing ontology files to align their vocabulary with the vocabulary of the domain, even if they are not initially familiar with the ontology's domain or its structure and content (e.g., C, D).
- Users are capable of understanding the requirements of the information-seeking process, expressing their valuation of the support they are provided by the interface as they performed their search tasks (e.g., E, F, G).
- Users felt that mediating ontologies make search tasks more manageable and easier, and not having them would negatively affect their task performance (e.g., E, F, G).

The following are some informal excerpts of some of the comments of those who have used ONTSI:

(A) *"I think with ONTSI, I can immediately it matches my mental models of how I use search interfaces. I type things in, I click run, I go through pages of results."*

(B) *"Once I understood what it was showing me, it helped me. Usually with new tools I tend to read through the documentation or watch videos. And then it still takes me like a while to pick up on them. Like, just running through them and using them a few times. Once you get the hang of it, usually you find success in whatever it's providing you."*

(C) *"But ... you can get lost in the information too, right? So, if you have like so much so many things related in that ontology, it's like, well, it can be useful. But it could also be a distraction for something that you know. There's this flip side, but I think that's on the searcher to know what they're using and why they're using it. So ... for me to complete these tasks, if I hadn't had the ontologies listed, then I would have had a much more difficult time. It essentially provided guidance ... and a structure to something I was unfamiliar with in this case."*

(D) *"I wasn't necessarily intimidated, but I was just like -- I don't know what this is. But the background information for the context helped a little bit. A lot of big words, but they did help me when I was looking at the documents that I had to search for to find out which ones I felt best. So even though I did not have full understanding of the words, having them there in that background provided me a kind of help towards finding myself in the space of the question."*

(E) *"I was thinking ... where (the ontology) would have been helpful. So ... it would have possibly brought up some of those other terms just from searching a few words and they would be able to make some connections between the text that was provided and some of my search terms (to see) ... how relevant they were. So, if I was shooting in the dark and hoping for the best, which is what I was kind of doing (without the ontology), at the very least, it would have given you confidence of your actions. Yeah, I think so. A little bit more confidence."*

(F) *"I thought it would be like pretty easy because I (am used to) answering ... open questions like ... find the things most relevant. So, this research question is for me ... just an easier thing to do because I have background in doing that kind of stuff. ONTSI kind of functions like ... a library tool that is available. This kind of tool felt very familiar to me. I wouldn't say that I'm an expert when it comes to medical knowledge, but ... I understand*

*... basic terminology. So... what the terms meant or what they refer to wasn't really ... an issue. It wasn't really alienating. I have like some general level of confidence just using the terms and trusting the tool as you went along."*

*(G) "Yeah, this (ontology) would have helped because I can find the things that ... share in common, and that can make it probably much easier to find the relevant documents. Yeah, being able to see the things that certain phrases ... or words share in common. You can find that common link ... that can find you the relevant documents."*

In the future, we plan to perform formal, empirical evaluations with users comparing ONTSI to other systems. Such evaluations will help generate new insights into features of interface designs and their qualitative and quantitative measures of how search task performances are affected. Beyond that, such evaluation studies may provide prescriptive guidelines for the design of optimal and effective interfaces for health informatics search tasks. In addition, we intend to further investigate how ontologies and machine learning should be integrated into elaborate and challenging search tasks that need domain-specific knowledge for optimal performance.

### **5.6.2 Limitations**

The first limitation of ONTSI is the scaling of computational resources. ONTSI in its current state provides a plug-and-play experience that can handle the uploading and processing of both document sets and ontology files of large sizes. For instance, ONTSI easily handles HPO and its more than 11,000 ontology terms, alongside an extracted subset of PubMed of more than 10,000 documents. Yet, under the load of large-volume document sets and connected suites of ontology files, ONTSI's computational systems may provide reduced responsiveness. To deal with such scenarios, further work is needed to solve overhead limitations—strategies such as pre-hosting common ontology files, establishing API connections to access externally hosted document sets, as well as simply expanding the computational power of our systems.

The second limitation of ONTSI is the support of ontology file formats. ONTSI in its current state can process the core encoded elements within the OWL format, a leading format for encoding ontologies. Yet, the format is quite verbose in its specification, requiring developments beyond the scope of our immediate research objectives. In addition, there are other formats used to encode ontologies that would be valuable to support ontology-supported interfaces for health informatics search tasks.

### **5.6.3 Conclusions**

In summary, in this paper we began with an examination of the background on the topics of health informatics, machine learning, and ontologies. We then reviewed recent research on health informatics search tasks. Based on this review, we formalized a set of criteria for guiding designers when creating ontology-supported interfaces for health informatics search tasks involving large document sets. We then used these criteria to contrast traditional design strategies for interfaces of search tasks.

To demonstrate the utility of the criteria in the design process, we applied them to structure the creation of ONTSI (ONTology-supported Search Interface), an ontology-supported interface for health informatics search tasks involving large document sets. ONTSI combines five front-end subviews and two back-end computational systems. With these systems, ONTSI supplies a generalized interface that supports users' ability to plug-and-play their provided document sets and an ontology file as a mediating resource within the interface when performing their health informatics search tasks.

The workflow of ONTSI was described and illustrated in a usage scenario. For our scenario, we used the Human Phenotype Ontology to mediate a search task on a subset of the PubMed document set. This usage scenario presented a narrative of a health professional performing a scoping review. Within the scenario, we found that ONTSI allows the user to utilize their ontology resource in a manner that aligns with both the unstructured and structured-like query expansion interface strategy. In the former, the user entered a research question without participating in mediation opportunities. In that case, ONTSI used HPO and WordNet as mediating resources to extend the user's query within an expansion model to generate the results of a search task. In the latter case, the user took advantage of mediation opportunities during their query building. Although this usage scenario provides a single health informatics narrative, we believe value can be generated from both the criteria and ONTSI for health informatics in a broad sense. In this sense, we envision that our efforts can be further expanded to encompass tasks in informatics such as consumer informatics, nursing informatics, and ontology-supported domains beyond health and medicine, to name but a few. In conclusion, in this paper we generated and proposed a set of criteria that can provide guidance to designers in creating ontology-supported interfaces for health informatics search tasks involving large document sets. We illustrated the utility of these criteria in the context of the creation and demonstration of ONTSI. We provided general insight from ongoing, formative, task-driven user evaluations of ONTSI. We hope to continue this research to promote the design of generalized ontology-supported interfaces for health informatics search tasks involving large document sets.

## 5.7 References

- Amershi, S., Cakmak, M., Knox, W. B., & Kulesza, T. (2014). Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine*, 35(4), 105–120. <https://doi.org/10.1609/aimag.v35i4.2513>
- Anderson, J. D., & Wischgoll, T. (2020). Visualization of Search Results of Large Document Sets. *IS and T International Symposium on Electronic Imaging Science and Technology*, 2020(1), 1–7. <https://doi.org/10.2352/ISSN.2470-1173.2020.1.VDA-388>
- Arp, R., Smith, B., Spear, A. D., & American Journal of Sociology. (2015). Building Ontologies with Basic Formal Ontology. In *Journal of Chemical Information and Modeling* (Vol. 53).
- Azad, H. K., & Deepak, A. (2019). A New Approach for Query Expansion using Wikipedia and WordNet. *Proceedings of the Working Notes of CLEF 2018—Conference and Labs of the Evaluation Forum, Avignon, France, 10–14 September 2018*, 492, 147–163. <https://doi.org/10.1016/j.ins.2019.04.019>
- Bikakis, N., & Sellis, T. (2016). Exploration and Visualization in the Web of Big Linked Data: A Survey of the State of the Art. *CEUR Workshop Proceedings*, 1558(Proceedings of the EDBT/ICDT 2016 Joint Conference,

Bordeaux, France, 15 March 2016).

- Bostock, M. (2016). D3.js Data-Driven Documents. Retrieved from <https://d3js.org/>
- Brewer, L. P. C., Fortuna, K. L., Jones, C., Walker, R., Hayes, S. N., Patten, C. A., & Cooper, L. A. (2020). Back to the Future: Achieving Health Equity Through Health Informatics and Digital Health. *JMIR MHealth and UHealth*, 8(1), 1–16. <https://doi.org/10.2196/14512>
- Capuano, N., Longhi, A., Salerno, S., & Toti, D. (2015). Ontology-driven Generation of Training Paths in the Legal Domain. *International Journal of Emerging Technologies in Learning*, 10(7), 14–22. <https://doi.org/10.3991/ijet.v10i7.4609>
- Carayon, P., & Hoonakker, P. (2019). Human Factors and Usability for Health Information Technology: Old and New Challenges. *Yearbook of Medical Informatics*, 28(1), 71–77. <https://doi.org/10.1055/s-0039-1677907>
- Carpendale, S., Chen, M., Evanko, D., Gehlenborg, N., Görg, C., Hunter, L., ... Strobel, H. (2014). Ontologies in Biological Data Visualization. *IEEE Computer Graphics and Applications*, 34(2), 8–15. <https://doi.org/10.1109/MCG.2014.33>
- Demiris, G. (2016). Consumer Health Informatics: Past, Present, and Future of a Rapidly Evolving Domain. *Yearbook of Medical Informatics*, S42–S47. <https://doi.org/10.15265/IYS-2016-s005>
- Dou, D., Wang, H., & Liu, H. (2015). Semantic Data Mining: A Survey of Ontology-based Approaches. *Proceedings of the 2015 IEEE 9th International Conference on Semantic Computing, Anaheim, CA, USA, 7–9 February 2015*, 244–251. <https://doi.org/10.1109/ICOSC.2015.7050814>
- Dudley, J., & Kristensson, P. (2018). A Review of User Interface Design for Interactive Machine Learning. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 8.2, 1(1), 1–37. Retrieved from <https://www.repository.cam.ac.uk/handle/1810/274032>
- Ebner, M. (2015). Linked Data Driven Visual Analytics for Tracking Learners in a PLE. *Proceedings of the 13th E-Learning Conference, Munich, Germany, 1 September 2015*, 329–331.
- Endert, A., Ribarsky, W., Turkay, C., Wong, B. L. W., Nabney, I., Blanco, I. D., & Rossi, F. (2017). The State of the Art in Integrating Machine Learning into Visual Analytics. *Computer Graphics Forum*, 36(8), 458–486. <https://doi.org/10.1111/cgf.13092>
- Fang, R., Pouyanfar, S., Yang, Y., Chen, S. C., & Iyengar, S. S. (2016). Computational Health Informatics in the Big Data Age: A Survey. *ACM Computing Surveys*, 49(1). <https://doi.org/10.1145/2932707>
- Gamache, R., Kharrazi, H., & Weiner, J. P. (2018). Public and Population Health Informatics: The Bridging of Big Data to Benefit Communities. *Yearbook of Medical Informatics*, 27(1), 199–206. <https://doi.org/10.1055/s-0038-1667081>
- Gao, C.-F. (2017). The 150 Most Important Questions in Cancer Research and Clinical Oncology Series: Questions 67–75. *Chinese Journal of Cancer*, 36(1), 86. <https://doi.org/10.1186/s40880-017-0254-z>
- Gibson, C. J., Dixon, B. E., & Abrams, K. (2015). Convergent Evolution of Health Information Management and Health Informatics. *Applied Clinical Informatics*, 06(01), 163–184. <https://doi.org/10.4338/aci-2014-09-ra-0077>
- Harvey, M., Hauff, C., & Elswiler, D. (2015). Learning by Example: Training Users with High-quality Query

- Suggestions. *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, New York, NY, USA, 9–13 August 2015*, 133–142.
- Hohman, F., Kahng, M., Pienta, R., & Chau, D. H. (2018). Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers. *IEEE Transactions on Visualization and Computer Graphics*, 25(8), 2674–2693.
- Huurdean, H. C. (2017). Dynamic Compositions: Recombining Search User Interface Features for Supporting Complex Work Tasks. *CEUR Workshop Proceedings*, 1798, 22–25.
- Jakus, G., Milutinovic, V., Omerović, S., & Tomazic, S. (2013). *Concepts, Ontologies, and Knowledge Representation*. <https://doi.org/10.1007/978-1-4614-7822-5>
- Jay, C., Harper, S., Dunlop, I., Smith, S., Sufi, S., Goble, C., & Buchan, I. (2016). Natural Language Search Interfaces: Health Data Needs Single-Field Variable Search. *Journal of Medical Internet Research*, 18(1), e13. <https://doi.org/10.2196/jmir.4912>
- Jimmy, Zuccon, G., & Koopman, B. (2018). Choices in Knowledge-base Retrieval for Consumer Health Search. *Proceedings of the European Conference on Information Retrieval, Grenoble, France, 25–29 March 2018*, 10772 LNCS, 72–85. [https://doi.org/10.1007/978-3-319-76941-7\\_6](https://doi.org/10.1007/978-3-319-76941-7_6)
- Jimmy, Zuccon, G., Palotti, J., Goeuriot, L., & Kelly, L. (2018). Overview of the CLEF 2018 Consumer Health Search Task. *CEUR Workshop Proceedings*, 2125.
- Jusoh, S., Awajan, A., & Obeid, N. (2020). The Use of Ontology in Clinical Information Extraction. *Journal of Physics: Conference Series*, 1529(5). <https://doi.org/10.1088/1742-6596/1529/5/052083>
- Katifori, A., Torou, E., Vassilakis, C., Lepouras, G., & Halatsis, C. (2008). Selected Results of a Comparative Study of Four Ontology Visualization Methods for Information Retrieval Tasks. *Proceedings of the Second International Conference on Research Challenges in Information Science, Marrakech, Morocco, 3–6 June 2008*, (September 2016), 133–140. <https://doi.org/10.1109/RCIS.2008.4632101>
- Keselman, A., Browne, A. C., & Kaufman, D. R. (2008). Consumer Health Information Seeking as Hypothesis Testing. *Journal of the American Medical Informatics Association*, 15(4), 484–495. <https://doi.org/10.1197/jamia.M2449>
- Köhler, S., Carmody, L., Vasilevsky, N., Jacobsen, J. O. B., Danis, D., Gourdine, J.-P., ... Robinson, P. N. (2018). Expansion of the Human Phenotype Ontology (HPO) Knowledge Base and Resources. *Nucleic Acids Research*, 47(D1), D1018–D1027. <https://doi.org/10.1093/nar/gky1105>
- Kohler, S., Doelken, S. C., Mungall, C. J., Bauer, S., Firth, H. V., Bailleul-Forestier, I., ... Robinson, P. N. (2014). The Human Phenotype Ontology Project: Linking Molecular Biology and Disease Through Phenotype Data. *Nucleic Acids Research*, 42(D1), 966–974. <https://doi.org/10.1093/nar/gkt1026>
- Köhler, S., & Robinson, P. (2016). HPO Web Browser: Blindness Infopage. Retrieved from Human Phenotype Ontology website: <http://compbio.charite.de/hpoweb/showterm?id=HP:0000618>
- Livingston, K. M., Bada, M., Baumgartner, W. A., & Hunter, L. E. (2015). KaBOB: Ontology-based Semantic Integration of Biomedical Databases. *BMC Bioinformatics*, 16(1), 1–21. <https://doi.org/10.1186/s12859-015-0559-3>
- Lüke, T., Schaer, P., & Mayr, P. (2012). Improving Retrieval Results with Discipline-Specific Query Expansion.

- International Conference on Theory and Practice of Digital Libraries*, 408–413.
- Luo, J., Wu, M., Gopukumar, D., & Zhao, Y. (2016). Big Data Application in Biomedical Research and Health Care: A Literature Review. *Biomedical Informatics Insights*, 8, BII.S31559. <https://doi.org/10.4137/bii.s31559>
- Lytvyn, V., Dosyn, D., Vysotska, V., & Hryhorovych, A. (2020). Method of Ontology Use in OODA. *Proceedings of the IEEE Third International Conference on Data Stream Mining & Processing, Lviv, Ukraine, 21–25 August 2020*, 409–413. <https://doi.org/10.1109/dsmp47368.2020.9204107>
- McCray, A. T., & Tse, T. (2003). Understanding Search Failures in Consumer Health Information Systems. *AMIA ... Annual Symposium Proceedings / AMIA Symposium. AMIA Symposium*, 430–434.
- Mehta, N., & Pandit, A. (2018). Concurrence of Big Data Analytics and Healthcare: A Systematic Review. *International Journal of Medical Informatics*, 114(March), 57–65. <https://doi.org/10.1016/j.ijmedinf.2018.03.013>
- PubMed. (n.d.). Retrieved from <https://pubmed.ncbi.nlm.nih.gov/>
- Qvarfordt, P., Golovchinsky, G., Dunnigan, T., & Agapie, E. (2013). Looking Ahead: Query Preview in Exploratory Search. *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval, New York, NY, USA, 28 July–1 August 2013*, 243. <https://doi.org/10.1145/2484028.2484084>
- Raghupathi, W., & Raghupathi, V. (2014). Big Data Analytics in Healthcare: Promise and Potential. *Health Information Science and Systems*, 2(1), 3.
- Rector, A., Schulz, S., Rodrigues, J. M., Chute, C. G., & Solbrig, H. (2019). On beyond Gruber: “Ontologies” in Today’s Biomedical Information Systems and the Limits of OWL. *Journal of Biomedical Informatics: X*, 2, 100002. <https://doi.org/https://doi.org/10.1016/j.yjbinx.2019.100002>
- Román-Villarán, E., Pérez-Leon, F. P., Escobar-Rodríguez, G. A., Martínez-García, A., Álvarez-Romero, C., & Parra-Calderón, C. L. (2019). An Ontology-based Personalized Decision Support System for Use in the Complex Chronically Ill Patient. *Studies in Health Technology and Informatics*, 264, 758–762. <https://doi.org/10.3233/SHTI190325>
- Russell-Rose, T., & Chamberlain, J. (2017). Expert Search Strategies: The Information Retrieval Practices of Healthcare Information Professionals. *JMIR Medical Informatics*, 5(4), e33. <https://doi.org/10.2196/medinform.7680>
- Russell-Rose, T., Chamberlain, J., & Azzopardi, L. (2018). Information Retrieval in the Workplace: A Comparison of Professional Search Practices. *Information Processing and Management*, 54(6), 1042–1057. <https://doi.org/10.1016/j.ipm.2018.07.003>
- Sacha, D., Sedlmair, M., Zhang, L., Lee, J. A., Weiskopf, D., North, S., & Keim, D. (2016). Human-centered Machine Learning Through Interactive Visualization: Review and Open Challenges. *ESANN 2016: 24th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning Bruges, Belgium April 27-28-29, 2016 Proceedings*, (27–29 April 2016), 641–646. Retrieved from <http://eprints.mdx.ac.uk/20763/>
- Saleemi, M. M., Rodríguez, N. D., Lilius, J., & Porres, I. (2011). A Framework for Context-aware Applications for

- Smart Spaces. In *Smart Spaces and Next Generation Wired/Wireless Networking: Vol. 6869 LNCS* (pp. 14–25). [https://doi.org/10.1007/978-3-642-22875-9\\_2](https://doi.org/10.1007/978-3-642-22875-9_2)
- Soldaini, L., Cohan, A., Yates, A., Goharian, N., & Frieder, O. (2015). Retrieving Medical Literature for Clinical Decision Support BT - Advances in Social Computing. *Advances in Social Computing*, 9022(Chapter 59), 538–549. Retrieved from [http://link.springer.com/10.1007/978-3-319-16354-3\\_59%5Cnpapers3://publication/doi/10.1007/978-3-319-16354-3\\_59](http://link.springer.com/10.1007/978-3-319-16354-3_59%5Cnpapers3://publication/doi/10.1007/978-3-319-16354-3_59)
- Soldaini, L., Yates, A., Yom-Tov, E., Frieder, O., & Goharian, N. (2016). Enhancing Web Search in the Medical Domain via Query Clarification. *Information Retrieval*, 19(1–2), 149–173. <https://doi.org/10.1007/s10791-015-9258-y>
- Solr Cloud. (2020). Retrieved from <https://lucene.apache.org/solr/>
- Talbot, J., Lee, B., Kapoor, A., & Tan, D. S. (2009). EnsembleMatrix: Interactive Visualization to Support Machine Learning with Multiple Classifiers. *Proceedings of the 27th International Conference on Human Factors in Computing Systems - CHI 09*, 1283. <https://doi.org/10.1145/1518701.1518895>
- The Future of Technology in Health and Health Care: A Primer. (2018). *Canadian Medical Association*. Retrieved from <https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&cad=rja&uact=8&ved=2ahUKEwjj8K7khsDxAhVBWs0KHamXDCcQFjACegQIAxAD&url=https%3A%2F%2Fwww.cma.ca%2Fsites%2Fdefault%2Ffiles%2Fpdf%2Fhealth-advocacy%2Factivity%2F2018-08-15-future-technology-health>
- The Human Phenotype Ontology. (2020). Retrieved May 10, 2020, from The Jackson Laboratory website: <https://hpo.jax.org/app/>
- Thiébaud, R., & Cossin, S. (2019). Artificial Intelligence for Surveillance in Public Health. *Yearbook of Medical Informatics*, 28(1), 232–234. <https://doi.org/10.1055/s-0039-1677939>
- Tobergte, D. R., & Curtis, S. (2013). A Survey on Ontologies for Human Behavior Recognition. *Journal of Chemical Information and Modeling*, 53(9), 1689–1699. <https://doi.org/10.1017/CBO9781107415324.004>
- Wickramasinghe, N. (2019). Essential Considerations for Successful Consumer Health Informatics Solutions. *Yearbook of Medical Informatics*, 28(1), 158–164. <https://doi.org/10.1055/s-0039-1677909>
- Wilson, M. L. (2011). Search User Interface Design. In *Synthesis Lectures on Information Concepts, Retrieval, and Services* (Vol. 3). <https://doi.org/10.2200/s00371ed1v01y201111icr020>
- Wu, C. M., Meder, B., Filimon, F., & Nelson, J. D. (2017). Asking Better Questions: How Presentation Formats Influence Information Search. *Journal of Experimental Psychology: Learning Memory and Cognition*, 43(8), 1274–1297. <https://doi.org/10.1037/xlm0000374>
- Yuan, J., Chen, C., Yang, W., Liu, M., Xia, J., & Liu, S. (2020). A Survey of Visual Analytics Techniques for Machine Learning. *ArXiv*, 1, 1–34.
- Zahabi, M., Kaber, D. B., & Swangnetr, M. (2015). Usability and Safety in Electronic Medical Records Interface Design: A Review of Recent Literature and Guideline Formulation. *Human Factors*, 57(5), 805–834. <https://doi.org/10.1177/0018720815576827>
- Zeng, Q. T., & Tse, T. (2006). Exploring and Developing Consumer Health Vocabularies. *Journal of the American*

*Medical Informatics Association*, 13(1), 24–29. <https://doi.org/10.1197/jamia.M1761>

Zielstorff, R. D. (2003). Controlled Vocabularies for Consumer Health. *Journal of Biomedical Informatics*, 36(4–5), 326–333. <https://doi.org/10.1016/j.jbi.2003.09.015>

Zuccon, G., & Koopman, B. (2014). Integrating Understandability in the Evaluation of Consumer Health Search Engines. *CEUR Workshop Proceedings*, 1276, 32–35.

## Chapter 6      Searching and Triaging Large Document Sets: An Ontology-Supported Visual Analytics Approach

This chapter has been prepared for publishing as: **Demelo, J., & Sedig, K. (2022).** Searching and Triaging Large Document Sets: An Ontology-Supported Visual Analytics Approach.

We have made minor adjustments to the original material of this chapter to provide cohesion with the overall integrated article structure of this dissertation. Specifically, to distinguish between chapters, figures and tables have been provided an additional prepend reflecting the chapter number. Readers should be aware that chapter text will maintain original numbering references. For instance, “Figure 6-1” is equivalent to “Figure 1” in the chapter text.

### 6.1 Introduction

Visual analytics combines the strengths of machine learning (ML) techniques, visualizations, and interaction to help users explore information interactively and achieve their knowledge building goals (Parsons et al., 2015). This joint human-machine coupling is more complicated than an internal automated analysis augmented with an external visualization of results seen by users. It is both data-driven and user-driven and requires re-computation when users manipulate the information through the visual interface (Sedig & Ola, 2014; Sedig & Parsons, 2016). Visual analytics tools (VATs) help users form valuable connections with their information and be more active participants in the analysis process (Golitsyna, Maksimov, & Monankov, 2018; Ramanujan, Chandrasegaran, & Ramani, 2017). They are used in a wide variety of domain tasks, such as supporting the sensemaking of misinformation, searching large document sets, and decision-making using health data, to name a few (Boschee et al., 2019; Demelo, Parsons, & Sedig, 2017; Ninkov & Sedig, 2019; Ola & Sedig, 2016). More than ever, researchers are investigating strategies to combat the rising computational needs of analytic tasks (Talbot, Lee, Kapoor, & Tan, 2009). ML technologies can be helpful in increasing the computational power of VATs; however, their utilization can often come at the cost of clarity and usability for users. Recent studies (A. Endert et al., 2017; Hohman, Kahng, Pienta, & Chau, 2018; Yuan et al., 2020) have found that traditional interface designs which integrate ML technologies limit user participation in the information analysis process and can lead to reduced user satisfaction. That is, users may struggle to understand and control ML when performing their tasks. In response, there is a growing desire to strengthen “human-in-the-loop” when designing interfaces (Alex Endert et al., 2014; Hohman et al., 2018; Wall, Blaha, Franklin, & Endert, 2018).

Users perform information search by communicating their information-seeking needs into the tool interface, which then generates a document set mapping used to direct document encounters which further knowledge building. This generation of document set mappings arise from the results of computational search, where investigations are concerned with the design of algorithms and processes which improve computational power. Yet for information search, investigations center on how decisions within the design process can allow users to better participate in the direction of computation search through the interface of their tools. While there is great value in novel computational

investigations, this paper directs investigation exclusively to the design process of visual interfaces, and how the active use of human reasoning affordances strengthens the “human-in-the-loop” during task performance.

Furthermore, when information search results are too numerous to be immediately useful, information triage must be performed. That is, document sets must sometimes be further triaged into more manageable sizes for the overall task — just as doctors must perform an initial, rapid triage on their intake before their concentration can shift to the details of individual cases. During information triage, users inspect, contextualize, and make timely relevance decisions on documents to produce a reduced, task-relevant set. However, information search and triage can be challenging for users, particularly in analytic tasks involving large document sets. Studies (Harvey, Hauff, & Elswailer, 2015) have found that when using traditional interfaces for search and triage, such as those with multiple input profiles, paged sets of documents, and linear inspection flows, users struggle to complete their domain tasks. Specifically, users routinely struggle to understand the domain being searched, apply their expertise, communicate their objectives during query building, and understand how to assess the relevance of search results during information triage.

For these concerns, we believe users can benefit from VAT interfaces which promote a novel combination of two considerations: 1) the use of progressively disclosure for the multi-staged information-seeking process, and 2) the use of ontologies to bridge user and task vocabularies. For the former, research (Huurdean, 2017) suggests that information seeking should be understood as a multi-staged process with distinct functional roles and human-centered requirements. Progressive disclosure is an organizational design technique which manages the visual space of an interface by occluding unnecessary elements of past and future stages, allowing users to concentrate on the task at hand. For multi-staged tasks like those involved in the information-seeking process, progressive disclosure has been found to effectively support users to perceive and plan their task performances, and thus can benefit users when searching and triaging (Chuang, Ramage, Manning, & Heer, 2012). For the latter consideration, users must understand and apply domain-specific vocabulary when communicating their information-seeking objectives. Yet, task vocabularies typically do not align with user vocabularies, particularly in tasks within complex domains such as health (Zeng & Tse, 2006). Ontologies are created by domain experts to provide a standardized mapping of knowledge that can be leveraged both by computational- as well as human-facing systems. Thus, ontologies can be valuable mediating resources to assist users in bridging their own vocabulary to the task vocabulary (Khan et al., 2016; Saleemi, Rodríguez, Lilius, & Porres, 2011). Yet, if these considerations are to be activated, there is first a need to establish high-level criteria for guiding the design of VAT interfaces for searching and triaging large document sets.

Therefore, we propose the following research questions:

- What are the criteria for the design of VAT interfaces that support the process of searching and triaging large document sets?
- If such a novel collection of design criteria can be distilled, can they be used to help guide the design of a progressively disclosed and ontology-supported interface?

We investigate the design of ontology-supported, progressively disclosed visual analytics interfaces for searching and triaging large document sets, distill a novel collection of design criteria, and generate a demonstrative visual interface with its guidance. We begin with background on information search, information triage, ML, and ontologies. We review leading research on the multi-staged information-seeking process to distill criteria. To illustrate the utility of the design criteria, we apply them to the design of a demonstrative prototype: VisualQUEST (Visual interface for QUery, Search, and Triage). VisualQUEST enables users to build queries, search, and triage document sets. Users can plug-and-play document sets and expert-defined ontology files within a domain-independent, progressively disclosed environment for multi-staged information search and triage tasks. We describe VisualQUEST through a functional workflow and culminate with a discussion of on-going formative evaluations, limitations, future work, and summary.

## **6.2 Background**

This section provides conceptual and terminological background. We begin with a discussion of information search and triage, explore the ML pipeline, then conclude with an examination of ontologies.

### **6.2.1 Information Search and Triage**

Numerous models exist which describe operational, temporal, and sequential frameworks of the information-seeking process (Hurdeman, Wilson, & Kamps, 2016). We mention two here. For example, Kuhlthau's six-part model describes the stages of initiation, selection, exploration, formulation, collection, and presentation. This model was then refined by Vakkari into their three-part model of pre-focus (initiation, selection, exploration), focus formulation (formulation), and post-formulation (collection and presentation) (Hurdeman, 2017).

When utilizing a VAT interface to search a document set, users' primary objective is to encounter documents that are most relevant to their task. For this, users must first communicate their information-seeking needs via the VAT interface. Computational components of the VAT can generate a mapping between users' input and the qualified and relevant documents in the document set; then presented to users at the interface level of the tool (Wu, Meder, Filimon, & Nelson, 2017). Existing research (Harvey et al., 2015) describes user requirements, and in turn, design considerations of interfaces that mediate information search on document sets. Namely, users must first establish an understanding of the document set and how it relates to their existing domain knowledge. Next, users must learn how to effectively communicate their objectives in a way that can be understood by the tool. Finally, users must comprehend how the VAT applied their input in its computational component so that they can effectively assess and guide their analytics process.

However, as document sets increase in size within analytic reasoning tasks, it has become more challenging for users to arrive at a final set of relevant documents without additional intervention. That is, even after computational components have reduced the document set down to a subset of documents, these subsets are still too large to be of value to users. For this issue, information triage may be required to further reduce the number of documents into a usable collection of task-relevant documents. During information triage, a user's primary objective is to inspect, contextualize, and make timely relevance decisions on search results (Herceg, Allison, Belvin, & Tzoukermann, 2018). For this to occur, existing research (Badi et al., 2006; Bae et al., 2010; Buchanan & Owen, 2008) describes that

tools must allow users to encounter and perform rapid triage on large sets of documents in a non-linear fashion, while still being able to assess document relevance to information-seeking objectives. Notably, supporting information triage within tools can also help users assess the quality of their searching and triaging, and inform them on how to improve further information seeking (Loizides, Buchanan, & Mavri, 2016).

## 6.2.2 Machine Learning

ML technologies are increasingly being applied to challenging analytic problems once considered too complex to solve in an effective and timely manner (Talbot et al., 2009). For instance, ML is utilized in developing pathways for drug discovery, for rapid design and analysis within materials science, and to improve the performance of search tasks on large document sets (Tang et al., 2019; Vamathevan et al., 2019; Wang et al., 2020; Wei et al., 2019).

ML processes are traditionally described as a three-staged pipeline covering the preparation, utilization, and assessment of models (Yuan et al., 2020). First, the primary objective when preparing a model is to analyze and sanitize incoming data. During this stage, responsibilities include data reformatting, minimizing signal noise, organizing common feature labels, and removing features which misalign with the task domain (Holzinger, 2014).

The next stage in the process is the utilization of a selected ML algorithm. There are many ML algorithms, typically categorized under either supervised or unsupervised learning (A. Endert et al., 2017). With supervised ML, labelled data is typically ingested and fit to train the model to optimally arrive at a gold standard output. The objective is that, given the same task and a new dataset with similar labels, a trained model can then repeat its performance (Fiebrink, Cook, & Trueman, 2011). Supervised models are best used in ML pipelines that require repeated predictions of label classification or the regression of numerical data points. Examples of supervised learning algorithms are Support Vector Machines, Naïve Bayes, and K-Nearest Neighbor (Dey, 2016). On the other hand, unsupervised ML relies on probabilistic adjustment techniques rather than training to a specific gold standard. The most common application of unsupervised learning is when an algorithm can analyze data points to learn of their shared associations within the structure of the input space. Thus, unsupervised models are best applied in ML pipelines whose goal is to make sense of data-point clusters and densities within the input space, such as mapping inputs to document sets to generate document groupings during information search (A. Endert et al., 2017). Examples of unsupervised learning models are hierarchical clustering algorithms which calculate distances between data points, centroid-based clustering (*e.g.*, K-Means) which converge data points to centralized nodes, and density-based clustering (*e.g.*, MeanShift) which re-weight data points based on proximity to densities within the input space (Celebi & Aydin, 2016; Dey, 2016).

In the final stage in the ML pipeline, users assess the effectiveness of the selected model so that they can conclude their task or provide feedback to their tool for future computations. For this stage, a research area receiving attention is how to design interactive ML interfaces such that a balance is struck between the computational power of the machine and the perceptual and decision-making power of humans. That is, by supporting the “human-in-the-loop” aspects of the design appropriately such that analysis can be provided by the user. Through the lens of visual analytics, such ML pipelines can further expand to a five-staged pipeline. That of, data collecting, cleaning, storage, analysis, followed by the final step of visualizing the data in both macro and micro forms depending on the needs of the visual analytics task (Sun et al., 2020). A generalized and human-centered interaction loop for interactive ML involves a set of stages where (Sacha et al., 2016):

1. Users specify their needs as a set of terms understood by the tool.
2. Users ask the tool to apply them as input features within its computational components.
3. The tool performs computation which maps the features against the document set.
4. The tool displays the results of the computation to users, and how it arrived at them.
5. Users assess if they are satisfied with the results, or if they would like to adjust their set of terms to generate an alternate mapping.
6. Users either restart the interaction loop or complete the task.

Despite knowledge of the above-mentioned stages, there are still many challenges in supporting effective engagement with ML components of tools, such as those within VATs. One challenge is the design of interfaces that can help users engage with ML processes effectively (Sacha et al., 2016). That is, if users cannot understand ML characteristics and requirements of their tool, they cannot perform their tasks effectively or maximally. Another significant challenge is supporting users in communicating their information-seeking objectives to a VAT's ML components. This is especially a concern in visual analytic tasks which involve direct interaction with ML processes (Hoerber, 2014; Holzinger, 2016; Mehta & Pandit, 2018; Tresp et al., 2016).

### **6.2.3 Ontologies**

When using VATs for searching and triaging large document sets, both the human and computational component can only perform optimally if their communication is strong (Arp, Smith, Spear, & American Journal of Sociology, 2015). Simply put, users can only perform well if they understand what their tool's interface is presenting to them. Furthermore, a tool can only optimize its computational components if the vocabulary and instruction applied by users to express their knowledge truly align with their intended information-seeking objectives. Tools are not typically designed to adapt to changing vocabularies. That is, when searching and triaging a specific document set, tools are traditionally designed to fit a singular vocabulary and task. It is often up to the users to understand the tool's domain-specific vocabulary and use that understanding to communicate their information-seeking objectives. Yet, learning the often unfamiliar vocabulary of the tool can be a significant challenge for users, particularly in tasks of complex domains (*e.g.*, health) which encapsulate terminology, relationships, axioms, and knowledge structures which diverge from common vocabulary typical of general users (Zeng & Tse, 2006).

To address this challenge, expert-defined ontologies are increasingly being used as mediating resources within human- and system-facing interfaces (Xing et al., 2019). Ontologies are created by domain experts to provide a standardized mapping of knowledge that can be leveraged both by computational- and human-facing resources (Khan et al., 2016). Ontologies are being used in a variety of applications, such as information extraction of unstructured text, behavior modeling of intellectual agents, decision support systems within critical care environments, as well as an increasing number of human-facing visualization tasks (Jusoh, Awajan, & Obeid, 2020; Lytvyn, Dosyn, Vysotska, & Hryhorovych, 2020; Román-Villarán et al., 2019). They can fall under one of three types: traditional ontologies to describe the structure of reality, domain ontologies created by experts to describe the entities, relations, and structures of a given domain, and top-level ontologies which interface domain ontologies (Arp et al., 2015).

Ontologies can provide flexibility, extensibility, generality, and expressiveness necessary when trying to effectively bridge domain knowledge between computational tools and humans (Saleemi et al., 2011).

When creating a domain ontology, experts must navigate the terms and characteristics of their domain to conceptualize a generalized and universal mapping of their knowledge (Arp et al., 2015). During the creation process, experts construct a structured network formed largely of ontology entities and relations (Jakus, Milutinovic, Omerović, & Tomazic, 2013; Rector, Schulz, Rodrigues, Chute, & Solbrig, 2019). Ontology entities reflect the conceptual objects of a domain, and will typically encode information about their role in the vocabulary, definitions, descriptions, contexts, as well as metadata that can inform the performance of future ontology engineering tasks (Tobergte & Curtis, 2013). Ontology relations express the type and quality of interaction between entities and unfold numerous unique interoperability of axioms within a domain (Katifori, Torou, Vassilakis, Lepouras, & Halatsis, 2008). Arp *et al.* (Arp et al., 2015) distinguish relations as universal-universal (*e.g.*, a rabbit “is a” animal), particular-universal (*e.g.*, this rabbit is an “instance of” a rabbit), and particular-particular (*e.g.*, this rabbit is a “continuant parts” of this grouping of rabbits).

After transcribing an ontology, designers prepare its entities, relations, and other descriptive attributes into standardized data file formats (*e.g.*, OWL: the W3C Web Ontology Language). These data files can be then shared amongst knowledge users and utilized in domain task tools.

## **6.3 Methods**

This section provides an analysis of leading research within the task space. We use this analysis to distill design criteria to guide the creation of VAT interfaces for searching and triaging large document sets.

### **6.3.1 Task Analysis**

To conduct a task-analysis review, we used Google Scholar, IEEE Xplore, and MEDLINE to search exhaustively for articles published between 2015 and 2021. We divide our analysis into three topic sections. First, we summarize leading research on the models of information-seeking process and the importance of progressive disclosure for supporting multi-staged tasks. Next, we examine the substages of information search: query building and search. This is followed by the substages of information triage: high-level and low-level triage.

#### ***6.3.1.1 Information-Seeking Process Models and Progressive Disclosure***

Huurdeman (Huurdeman, 2017) suggests that information seeking should be understood as a multi-staged process with distinct functional requirements. This research explores existing models of the information-seeking process, beginning with a summary of six-stage model of: initiation, selection, exploration, formulation, collection, and presentation. This is further refined and summarized as a three-stage model: pre-focus (initiation, selection, and exploration), focus formulation (formulation), and post-formulation (collection and presentation). Huurdeman concludes with an analysis of each of these stages, stating that “information sought for evolves during different stages”. Specifically, Huurdeman describes how users in the pre-focus stage concentrate on conceptualizing their topic using search tactics like browsing, querying, and deciding on search models. Next, the focus formulation stage investigates

the broad concepts that are being searched. Finally, during the post-focus stage, users are concerned with searching for specific information, increasing from low specificity (high-level assessments of relevance) to high specificity (low-level assessments of relevance). For our purposes, we can use the above-mentioned analysis to break down the information search and triage task into four stages: *query building*, *search*, *high-level triage*, and *low-level triage* (see Table 1).

**Table 6-1 The stages of the information-seeking and triage process: their associated task, alignment with existing models, and functional descriptions.**


| Stage             | Associated Task    | Alignment with Existing Models                                  | Functional Descriptions   |
|-------------------|--------------------|---|---|
| Query Building    | Information Search | Pre-focus (Initiation, Selection, Exploration)                  | Users communicate their information-seeking objectives via the tool's interface.  |
| Search            | Information Search | Focus Formulation (Formulation)                                 | Users specify the formulation of their search, and when satisfied, initiate the performance of computational search.  |
| High-Level Triage | Information Triage | Low-specificity Post-formulation (Collection and Presentation)  | Users encounter sets of similar information entities generated from computational search, make initial high-level assessments of general alignment with information-seeking objectives, and direct further triaging encounters. |
| Low-Level Triage  | Information Triage | High-specificity Post-formulation (Collection and Presentation) | Users encounter individual information entities, previously encountered in high-level triage, to perform final, low-level assessments of relevance to information-seeking objectives.   |

Progressive disclosure is a technique for organizing and managing the visual space of an interface. When implementing the technique, designers abstract and sequence the stages of a complex task. Afterwards, views can be generated with information encodings and controls to promote the performance of each individual stage. Views can be placed in sequence in the visual space, with controls to direct their activation. This enables unnecessary interface elements, reflecting past and future stages, to be minimized, while still maintaining transparency of other relevant stages (Springer & Whittaker, 2018). The goal of progressive disclosure is to minimize users' distraction, allowing them to concentrate on the critical decisions of the task at hand (William Lidwell, Kritina Holden, 2010). When designed well, it can effectively support users to perceive, plan, and navigate complex, multi-staged tasks (Chuang et al., 2012).

Early research on progressive disclosure insisted on full occlusion of future stages. The belief was that in cases when future stages do not directly support a current, active stage, they should be fully concealed and only be accessible by request (William Lidwell, Kritina Holden, 2010). However, recent research suggests otherwise. Specifically, user studies (Springer & Whittaker, 2018) have compared full occlusion strategies with strategies that do not fully occlude future stages. These studies find that the former strategies distract users, lack information

transparency, and reduce user opportunity for feedback. This suggests that effective progressive disclosure strategies should maintain a balanced transparency between performances of previous stages and impact of current stage on future stages.

An example of a progressively disclosed interface is WebMD's Symptom Checker. This interface used by the general public supports the task of health diagnosis – a complex, multi-staged task. In particular, WebMD's Symptom Checker's implementation provides hints for how current task decisions may impact the performance of future stages within the diagnosis sequence (see Figure 1).



**Figure 6-1 WebMD's Symptom Checker: An example of a progressive disclosure implementation of a task. In this case, users are guided along a series of query building opportunities, input-ting symptoms and personal health criteria. Source: Image generated on 2021/01/18 using the public web portal provided by WebMD, <https://symptoms.webmd.com/> (accessed on January 18, 2021)**

Next, we review research on the individual stages of information search and triage.

### **6.3.1.2 Stages of Query Building and Search**

When searching for information in the document sets of unfamiliar domains, users generally possess some level of knowledge deficiency. This makes it difficult to formulate and communicate their problem, a challenge that users must work to overcome. According to Harvey *et al.* (Harvey et al., 2015), this is because users consistently suffer from four major issues during search:

- Difficulty understanding the domain being searched.
- Inability applying domain expertise.

- Inability to accurately formulate queries matching information-seeking objectives.
- Deficiency assessing and determining if search results satisfy objectives, or if adjustments are required.

We now explore these in more depth.

While searching information, users must learn about the searched domain, understand how their information-seeking objectives align, and formulate how to communicate their knowledge in a way that can be understood by their tools. Thus, designers must provide users with the opportunity to understanding the document set being explored. Huurdeman (Huurdeman, 2017) highlights that search interface designers must consider both novices and experts. Yet, studies by Harvey *et al.* (Harvey et al., 2015) show that, in domains with complex vocabularies (*e.g.*, health and medicine), the disparity of prior domain knowledge in potential users is extreme – that is, users routinely do not possess enough domain knowledge to satisfy their information-seeking needs. This can cause significant problems during query formulation for both domain experts and non-expert users. As a result, non-expert users must first step away from their VAT to learn to express their knowledge using specialized domain vocabulary before they can begin query building. Both Soldaini and Anderson (Anderson & Wischgoll, 2020; Soldaini, Yates, Yom-Tov, Frieder, & Goharian, 2016) note that this issue can still affect even domain experts. This is because they must often make assumptions regarding the appropriateness and specificity of their information-seeking communications and/or inputs to the VAT.

Another commonly cited challenge for users is their inability to perceive how their query decisions impact, relate, and interact with the document set being searched. This is a particularly important consideration for users who want and need to adjust their previously communicated query to better align with their information-seeking objectives. Huurdeman (Huurdeman, 2017) also describes potential strategies to address these concerns, prescribing the use of query corrections, autocomplete, and suggestions. Yet these strategies can be ineffective if they do not allow users to be cognizant of how their query-making decisions achieve the results that they seek.

Seha *et al.* (Saha et al., 2016) examine considerations for designing ontology-supported information retrieval systems. They suggest that natural-language interfaces to information sources provide novice users a pathway to avoid complex tool-dependent query languages. They highlight the benefit of shifting the vocabulary of query building away from the tool and towards the semantics of the domain being searched. Munir *et al.* (Munir & Sheraz Anjum, 2018) summarize the benefits of computational strategies which integrating ontologies into information retrieval systems for tasks involving both information search and triage. They state that ontology-supported information retrieval has become an advantageous strategy for supporting domain-specific over tool-specific vocabulary due to their improvement to the effectiveness of human-computer communication. Using the medical domain as a framing device, Soldaini *et al.* [48] investigate the use of novel, ontology-supported query computation strategies in improving the quality of literature retrieval during search tasks. They apply combinations of algorithms, vocabularies, and feature weights to assess the computational performance of different query-reformulation techniques. Their findings suggest that the utilization of bridged vocabularies within ML components improve retrieval performance.

In a systematic review of search interfaces that have ML components, Amershi *et al.* (Amershi, Cakmak, Knox, & Kulesza, 2014) compile a few considerations:

- Users are people, not oracles.
- Users should not be expected to repeatedly answer if ML results are right or wrong without an opportunity to explore and understand the results.
- Users tend to give more positive than negative feedback to interactive ML.
- Users need demonstration of the behavior of ML components.
- Users value transparency in ML components of tools, as transparency helps users provide better labels to ML components.

The application of sensitivity encoding within human-facing interface design can also be of benefit to users when searching and triaging (Cortez & Embrechts, 2013). Sensitivity encoding is a design strategy to provide a visual preview of possible results if particular actions are taken (Spence, 2014). Within query building interfaces with sensitivity encoding, a visual interface can provide the number of current query search results against ones with minor adjustments. In such cases, users can relax one query item, thereby changing the size of the result set. By providing such meaningful context cues, sensitivity encoding can guide query building and search formulation and enhance the perceived value of the search results, particularly when used in combination with techniques like progressive disclosure (Spence, 2002, 2004). An example of sensitivity encoding in a search and triage tool interface is OVERT-MED (Demelo, Sedig, & Parsons, 2017). This tool uses sensitivity encoding to provide alignment cues within its query generator between the expressed vocabulary and the search space (see Figure 2).




Figure 6-2 OVERT-MED, a sensitivity encoded ontology-driven search interface. Depicted is the result of users comparing the queries of “Progressive + ophthalmoplegia” and “congenital + fibrosis”. Source: Image generated on 2021/07/05 with permission courtesy of Insight Lab, Western University, London Ontario Canada <http://insight.uwo.ca/> (accessed on July 5, 2021)

### 6.3.1.3 Stages of High-Level and Low-Level Triage

When performing information search on large document sets, the result set can remain large and overwhelming, even after significant refinement efforts by users during query building. To make the result set smaller and more manageable, information triage is often a necessary extra step in the information-seeking process. Research by Azzopardi *et al.* (Azzopardi & Zuccon, 2016) indicates that the prevailing design language for triage interfaces provides users with linear inspection flows maintaining interactions with long lists of documents. That is, traditional triage interfaces present document search results as ordered sets of individual documents. However, this linear inspection flow has been found to significantly hamper both efficiency and effectiveness of the users making relevant decisions using the result sets, and does not scale for tasks with large document sets (Bae et al., 2010). Poorly designed tools typically attempt to hide their scaling weaknesses by paging away large percentages of their results, implementing smooth scrolling interactions, or worse, by showing only the first result in an attempt to avoid triaging altogether. Yet, tools cannot ignore information triage, as it is critical in helping users assess the quality of search results (Loizides et al., 2016). An example of a tool which forces a linear inspection of document results within a paged system is PubMed – with Figure 3 showing the default triaging interface of this tool.

The screenshot displays the PubMed search results for the query "heart". The search bar at the top contains the term "heart" and a "Search" button. Below the search bar, there are options for "Advanced", "Create alert", "Create RSS", and "User Guide". The results are sorted by "Best match" and there are buttons for "Save", "Email", "Send to", and "Display options".

On the left side, there are several filter sections:


- MY NCBI FILTERS**: A link to filter settings.
- RESULTS BY YEAR**: A graph showing the number of results over time, with a peak around 2021. The x-axis ranges from 1785 to 2021.
- TEXT AVAILABILITY**: Checkboxes for "Abstract", "Free full text", and "Full text".
- ARTICLE ATTRIBUTE**: A checkbox for "Associated data".
- ARTICLE TYPE**: Checkboxes for "Books and Documents" and "Clinical Trial".

The main results list shows three items:

- 1 The heart and development.**  
Anderson PA.  
Cite: *Semin Perinatol.* 1996 Dec;20(6):482-509. doi: 10.1016/s0146-0005(96)80064-4.  
PMID: 9090776 Review.  
Share: The perspective from which the developing **heart** is viewed can lead to differing conclusions about the effects of development on cardiac function. The **hearts** of the embryo, fetus and adult, viewed from a global perspective, sustain the circulation through the same ba ...
- 2 The heart's 'little brain' controlling cardiac function in the rabbit.**  
Brack KE.  
Cite: *Exp Physiol.* 2015 Apr 1;100(4):348-53. doi: 10.1113/expphysiol.2014.080168. Epub 2014 Oct 29.  
PMID: 25833107 **Free PMC article.** Review.  
Share: To investigate this, an isolated, perfused, innervated rabbit Langendorff **heart** preparation was considered the best option. Although ganglionic plexuses have been well described for several species, there was no full description of the anatomy and histochemistry of rabbit ...
- 3 Ontogenetic shifts of heart position in snakes.**  
Lillywhite HB, Lillywhite SM.  
Cite: *J Morphol.* 2017 Aug;278(8):1105-1113. doi: 10.1002/jmor.20697. Epub 2017 May 3.  
PMID: 28470824 **Free PMC article.**  
Share: **Heart** position relative to total body length (TL) varies among snakes, with anterior **hearts** in

Figure 6-3 PubMed: Searching the MEDLINE document set for “heart” provides a linear triaging interface of over 1.5 million results spread over 100,000 pages within a ten per page system. Source: Image generated on 2021/07/05 using the public web portal provided by the National Center for Bio-technical Information, <https://pubmed.ncbi.nlm.nih.gov/?term=heart> (accessed on July 5, 2021)

When designing information triage interfaces that provide non-linear inspection flows, the two primary concerns for designers are how best to represent documents to users, and how users can interact with those documents to further their information-seeking objectives. Critically, designers must re-approach triaging as a part of the multi-staged information-seeking process. In this regard, Chandrasegaran *et al.* (Chandrasegaran, Badam, Kisselburgh, Ramani, & Elmqvist, 2017) describe the value of high-level triage of full document sets prior to individual document inspection. Specifically, they suggest that designers should avoid having users open individual documents in full initially, and instead create visual abstractions which provide an overview display of all documents for high-level relevance assessment. Anderson (Anderson & Wischgoll, 2020) aligns with this consideration, stating that ideal high-level triage strategies should abstract out shared characteristics to structure groupings of comparable documents. Users can be provided interactions which allow them to simultaneously traverse, preview, contrast, and judge relevance for groupings of documents at a time, rather than individually. Figure 4 shows the ChartingHockey Production Rates interface to support visual search for non-linear triaging of player performance within the NHL, as opposed to a traditional linear approach which would have listed each player in an ordered list of statistical information. This can be regarded as an example of the high-level triage stage within a visual interface supporting an encompassing information seeking task.



**Figure 6-4 ChartingHockey Player Production Rates:** ChartingHockey supports non-linear triaging of statistical information in player evaluation tasks. Source: Image generated on 2021/07/05 using the public web portal provided by ChartingHockey, <https://www.chartinghockey.ca/daily-skater-charts/> (accessed on July 5, 2021)

After users perform high-level triaging on groupings of comparable documents, they should be able to perform low-level triaging – that is, assessing the contents of groupings at the individual document level. For low-level triaging, Huurdeman [45] suggests that users should be aided by displays which allow them to save, annotate, and/or provide other personalized interactions. Other useful strategies include displaying document metadata, titles, and short snippets of result sets (Huurdeman, 2017). Oftentimes, traditional interfaces do not provide scaffolding supports that help users make distinctions between the high-level and low-level triage stages. Specifically, Huurdeman (Huurdeman, 2017) states that interfaces often use one of three strategies when displaying search results:

- Underload documents: Little to no content of each document is displayed, making it difficult to compare documents.
- Overload documents: Too much of each document is displayed, making it difficult to rapidly understand each document.
- Distort documents: A summarization, weighting, or filtering strategy is used to either demote or promote certain document attributes, providing different tradeoffs: making some attributes easier to perceive, creating poor decontextualized generalizations, hiding away value, and sometimes promoting harmful attributes.

The design of low-level triage interfaces can be challenging. Namely, designers must consider the characteristics of the information to be represented, the knowledge domain of users, the task to be performed, as well as the interactions that can effectively support the former (Bae et al., 2010; Kekäläinen, 2014; Loizides et al., 2016). When designing document displays, it is important that any summarization, weighting, and filtering techniques highlight attributes which best assist in the application of domain expertise for relevance judgement. For this, both Loizides and Mavri (Loizides & Buchanan, 2009; Mavri, Loizides, Photiadis, & Zaphiris, 2013) suggest some best-practice design strategies. They note that factors such as section types, content positioning, font weight, and font size (among other factors) influence final document relevance decision-making. Furthermore, document titles, captions, abstracts, section snippets, conclusions, and a decreasing valuation for document pages are the most important factors that affect decision-making regarding relevance of documents (Loizides, 2012; Mavri et al., 2013). Simply put, a tool which supports effective document triaging should maximize and highlight important content and minimize elements that do not support rapid decision-making.

### 6.3.2. Design Criteria

Using the task-analysis review in the previous section, we distill the following set of design criteria, presented in Table 2.

**Table 6-2 The set of design criteria for creating VAT interfaces for searching and triaging large document sets. For reference and clarity, a numerical value is assigned to each design criteria. Each DC# describes the design criteria, provides an integration classification, and suggests value in ontology integration for that criteria.**

| # | Design Criteria   | Integration | Value in Ontology Integration   |
|---|---|-------------|---|
| 1 | Use progressive disclosure when sequencing the stages of the information-seeking process. | All Stages  | Ontology entities and relations can be consistent and transparent guideposts between stages, particularly for non-active stages which must be pruned of unnecessary elements. |

---

|    |   |                   |   |
|----|---|-------------------|---|
| 2  | Attune users to the characteristics and domain of the document set before beginning search formulation.                   | Query Building    | Ontology entities and relations can promote the characteristics and domain of document sets.  |
| 3  | Be cognizant of users' domain expertise.  | Query Building    | Ontology entities and relations can provide a bridge between task vocabularies and the common vocabularies of non-expert users, as well as previously formed domain vocabularies of expert users.                                     |
| 4  | Create search formulation and refinement environments supplemented by query building.                                     | Search            | Ontology entities and relations can be useful within interface elements that suggest expansions and refinements to their search formulation.  |
| 5  | Leverage sensitivity encoding when previewing the document set mappings of search formulations.                           | Search            | Ontology entities and relations can be useful sensitivity encoded displays which can suggest refinement opportunities for re-aligning their search formulation to the document set being searched and information-seeking objectives. |
| 6  | Present overview displays which arrange and compare document groupings using shared characteristics.                      | High-level Triage | Ontology entities and relations can be useful in abstraction, such as locating shared document characteristics and when forming document groupings.   |
| 7  | Utilize non-linear inspection flows which support actions for traversing, previewing, contrasting, and judging relevance. | High-level Triage | Ontology entities and relations can help users connect to and assess the general characteristics and contents of a document grouping, allowing them to inspect, assess, and judge relevance on multiple documents at a time.          |
| 8  | Offer document-level displays which allow users to apply domain expertise during relevance decision making.               | Low-level Triage  | Ontology entities and relations can be useful for directing summation and annotation actions, as well as provide familiar cues for interactions like sorting and relevance judgement.   |
| 9  | Persist relevance decision making results externally to allow for repeat information-seeking sequences.                   | Low-level Triage  | Ontology entities and relations can be useful for indexing document selections as well as for recordkeeping users' prior search and triage sequences.   |
| 10 | Allow users to encounter search results without a demand for immediate appraisal.   | All Triage        | Ontology entities and relations can help direct search formulation previews and the results of a full mapping of the document set, allowing users to quickly associate their predictions against search results.                      |
| 11 | Promote positive feedback over negative feedback.   | All Stages        | Ontology entities and relations can provide familiar cues to direct positive feedback interactions within information-seeking sequences.  |

---

## 6.4 Materials

In this section, we describe VisualQUEST, an ontology-supported and progressively disclosed VAT created to demonstrate the utility of the design criteria for searching and triaging large document sets. We describe the technical scope and functional workflow of VisualQUEST.


### **6.4.1 Technical Scope**

VisualQUEST is a web-based generalized plug-and-play interface with user-provided ontology files and document sets. VisualQUEST provides cross-browser (Firefox, Chrome, Opera) and cross-platform support. VisualQUEST's front-end views use HTML5, CSS, and JavaScript. D3.js JavaScript visualization library is used in VisualQUEST's interactive displays (Bostock, 2016). A custom Python Flask-based server is used for data storage and ML computations. VisualQUEST also uses Apache's Solr system as its indexer and search engine ("Solr Cloud," 2020).

### **6.4.2 VisualQUEST Functional Workflow**

VisualQUEST's workflow encompasses several system and view components. The front-end interface of VisualQUEST maintains an accordion-like design which sequences view stages using the progressive disclosure technique (DC1). Only one VisualQUEST subview is active at a time, assigning it the majority of the visual interface. Following progressive disclosure best practices, inactive stages are not occluded. Instead, they are assigned reduced, yet still present display space which highlights any task-relevant value generated within the stage.

We now describe the overall workflow of VisualQUEST and its parts (see Figure 5).



**Figure 6-5** Depiction of the dynamic functional workflow of VisualQUEST. Labeled arrows reflect transitional actions of systems and users. Users begin by uploading their ontology files and document sets (bottom center). These activate their respective Ontology and Document Process Systems (brown boxes), which prepare their content for back-end storage (blue boxes) and activation (pink boxes). Once achieved, the front-end interface activates its various subview functionalities (yellow boxes). The user can act upon the interface (green box). The system processes those actions, formulates adjustments to its display, and returns a visual response to be perceived by users. Some-times, these adjustments must connect with back-end systems, ask computations to be performed, and the results of those computations sent back to the display level.

### 6.4.3 Back-End Systems

VisualQUEST is supported by two servers which move heavy computation away from the browser: Analytics Server and Document Server.

#### 6.4.3.1 Analytics Server

Analytics Server is built using the Python-based Flask framework. It is accessed through an API which offers two functionalities: 1) uploading user-provided document sets and 2) performing ML computations. Document sets are

uploaded from users' computer file system to Analytics Server. This server validates the type, format, size, and encoding of uploaded documents, and then stores them into a temporary PostgreSQL database. This database is accessed by Document Server (a Solr server) during indexing procedures. Analytics Server can request Document Server to index all new documents. When users request ML computations during search, Analytics Server assesses the selected algorithm, the document set, and the search formulation generated by users during query building and search. It then performs sanitization and query expansion. During this process, query items within the search formulation are expanded using user-provided ontology files and WordNet for synonym ring analysis. The search formulation in both its original and expanded form is packaged and applied within ML computations. The resulting clusters are propagated back to VisualQUEST. These ML computation services are facilitated by the Scikit-Learn library, which we do not express to be a part of the novel contributions in this material (Pedregosa et al., 2012). That is, we connect with existing ML tool sets to enable the required ML computation in support of our investigations on design processes for visual interfaces. We include pseudocode describing this process (Figure 6).

```

Algorithm 1: CLUSTERING pseudocode between QUEST, Analytics Server, and Document Server


---


Input: A set Q of user inputted queries
Output: Signal to update interface with cluster assignments
1 targets ← chain(Q).unique().difference(getStopWords())
2 documents ← getDocuments()
   /* Prepare bag of words using target, related entities, and
   their generated WordNet synsets */
3 for i = 0 to targets.length do
4   target = targets[i]
5   targetCoverage ← target + target.getDirectlyRelatedEntities()
6   targetSpread[target] ←
     targetCoverage + targetCoverage.getWordNetSynsets()
7   targetSpread[target] ←
     targetSpread[target].unique().difference(getStopWords())
   /* Gather counts from pre-indexed documents, then fit and
   predict clusters using Scikit.Learn KMeans clustering */
8 documentCounts ←
   getIndexesFromSolrAPI(targetSpread, documents).scaleRange(0, 1)

9 reducedPCA ← SciKitLearn.PCA(nComponents =
   2).fit_transform(documentCounts)
10 kmeansPCA ← SciKitLearn.KMeans(init = 'k-
   means ++', nClusters = 7, nInit = 10)
11 clusterAssignments ← kmeansPCA.fit_predict(reducedPCA)
12 for i = 0 to targets.length do
13   target = targets[i]
14   for j = 0 to clusterAssignments.length do
15     cluster = clusterAssignments[j]
16     yPred = cluster.yPred[target]
     /* Generate weighting scale using x5 multiplier */
17     clusterAssignments[j].weighting[target] ←
       generateClusterWeighting(yPred)
18 return signalInterfaceUpdate(clusterAssignments)

```

Figure 6-6 Pseudocode of clustering functionality spanning the workflow of VisualQUEST (front-end), Analytics Server, and Document Server.

### 6.4.3.2 Document Server


VisualQUEST’s Document Server is a cloud-based Solr server for indexing, storing, and serving documents from user-provided document sets. Solr is a prepackaged, scalable indexing solution developed by The Apache Software Foundation. It provides a valuable array of features like a REST-like API that support numerous HTTP-based communication interfaces. Solr also supports a wide range of customizable settings and schemas for storing, searching, filtering, analyzing, optimizing, and monitoring tasks (“Solr Cloud,” 2020). During indexing procedures, Document Server uses a prepared schema to extract new documents from a temporary PostgreSQL database hosted by Solr. These documents are then treated and stored within an index. Document Server also handles document serving requests. When VisualQUEST is displaying documents, Document Server provides metadata such as titles, word counts, as well as content for document-level displays (See Apache’s official website and document for more information on Solr (“Solr Cloud,” 2020)).

## 6.4.4 Front-End Subviews

This section describes VisualQUEST’s subviews with reference to relevant design criteria (DC#) discussed before.

### 6.4.4.1 Query Building Subview

Query Building is the first subview within VisualQUEST (Figure 7, DC1).



**Figure 6-7** An overview of Query Building subview within VisualQUEST. Here, an ontology and document set have been uploaded, from which a set of query items have been generated, with a subset of those added in Search.

In this subview, two functions are performed: uploading user-provided files and query building. Upon clicking the upload button, users can select ontology files and document sets which are then inserted into a file management listing, accompanied by file name, type, and any available descriptions. Once a document set has been uploaded, users can begin query building by inputting text into a search bar. This leads to the generation of query items from all combinations of the inputted words (DC2). For instance, if a two-word input is provided, a query item is generated for each, as well as two-word query items in both possible orders (e.g., A, B, A B, B A). Each query item is

accompanied by a count of its *verbatim* presence within the document set (DC10). Furthermore, VisualQUEST analyzes the alignment of query items with the entities, relations, and descriptions of user-provided ontology files. If a feature of the ontology is found to align with a query item, it is placed within a drop-down menu attached to the listing (DC3). In this menu, users are presented with ontology terms which are conceptually similar to that query item (Figure 8).





Figure 6-8 Expanding a query item to assess related ontology elements.

Encountering these terms, users can learn more about the domain vocabulary, appraise how their research problem may or may not align with their document set and adjust their query item selections (DC2). Both direct-input and ontology-mediated query items can be saved and are assigned unique colors that are used throughout all subviews (DC11).


#### 6.4.4.2 Search Subview

Search is the second subview within VisualQUEST (Figure 9, DC1).



**Figure 6-9** An overview of the Search subview within VisualQUEST. Prior query building has produced a set of query items, which have been inserted into the search formulation for preview.

In this subview, users can control the formulation of search queries, encounter sensitivity-encoded previews of the formulation, and initialize search on the full document set (DC4). In the Search subview, a list allows users to manage query items, including insertion into the search formulation (DC11). After at least one query item has been selected, a preview of the current search formulation is activated (Figure 10).




**Figure 6-10** A preview of a search formulation which uses all query items.

This preview is a matrix-like display describing a cluster analysis of document groupings within the document set (DC5, DC10). By adding and removing query items from the search formulation, users can investigate: 1) how individual query items align with the document set, 2) how differing query item combinations change document grouping arrangements, as well as 3) estimate how many documents may be found in a full search. If not satisfied, users can refine their formulation using existing query items or generate new query items within Query Building (DC4, DC10). Finally, users can initialize a full search, with the results of ML computations sent to the triage stages (DC11).

### 6.4.4.3 High-Level Triage Subview

High-Level Triage is the third subview within VisualQUEST (Figure 11, DC1).



**Figure 6-11** An overview of the High-Level Triage subview within VisualQUEST. Previous searching has produced groupings of the document set. A subset of those groupings has been selected, producing further listing of their contained documents. Some of these documents have been inspected and added into Low-Level Triage.

In this subview, users triage the results of ML computations at the grouping level. A full document mapping is displayed within Query Result Heatmap, providing users with a high-level abstraction of the document set (Figure, 12, DC6, DC10) (Demelo, Sedig, et al., 2017).




Figure 6-12 High-Level Triage after groupings have been inspected and opened for further examination.

The abstraction is divided into horizontal slices representing document groupings. For each document grouping, a set of color cues highlight query item presence. Users can inspect each document grouping to assess its size and alignment with the search formulation. Listings can be re-ordered to prioritize specific query items. A cursor marks the current position, trailing dots mark previously viewed listings, and a green mark for those selected for further triaging. Document groupings can be opened within an additional Query Result Heatmap, which provides high-level abstractions of individual documents from selected groupings (Figure 13, DC7).




Figure 6-13 A close look at High-Level Triage, showing a listing of documents contained within a selected grouping of the document set.

Users can use this additional collection of documents to individually assess alignment with the search formulation, as well as inspect metadata such as titles and document-specific counts (DC6). Documents can then be saved to Low-Level Triage (DC7, DC11).


#### 6.4.4.4 Low-Level Triage Subview

Low-Level Triage is the fourth subview within VisualQUEST (Figure 14, DC1).




**Figure 6-14** An overview of the Low-Level Triage subview within VisualQUEST. A set of saved documents have been produced in prior triaging. Each of these documents have been provided a timeline-like summary reflecting words or phrases aligning with the search formulation within its content. When selected, the Document Content viewer depicts the document itself, either in the summarized or full document mode. The summarized mode is shown.

In this subview, users triage the documents produced in High-Level Triage. Users are provided a timeline-like visual abstraction of individual documents (Figure 15).



**Figure 6-15** A closer view of the Selected Documents listing, where the user can see timeline-like abstraction of documents and their alignment with the search formulation.

In this visual abstraction, colored marks are placed along a timeline to reflect the position of words or phrases which align with the query items of the search formulation. Documents with strong alignment will produce numerous markings, resulting in color-heavy and densely annotated timelines. Utilizing the timelines, users can perform rapid analysis of the thematic themes of individual documents. Namely, users can assess the presence of query items used within the search formulation, where in the document they are, and the density of their usage (DC8, DC10). The Document Content viewer is activated after selecting a document for deeper inspection (Figure 16, DC10, DC11).



**Figure 6-16 An overview of Low-Level Triage after documents have been added and opened for viewing.**

This viewer will present all available document content, such as document text, title, authors, file name, URLs, and published date. Users can toggle between a full document and a summarized mode. The full document mode displays all available information, annotated to reflect alignment with the search formulation. The summarized mode condenses documents to just content in proximity to aligning words or phrases (DC8, DC10). Users may open documents to inspect, compare, and make final relevance decisions on its content. Relevant documents can be added to a persistently saved list, allowing users to continue searching and triaging without the risk of losing progress (DC9).

## 6.5 Discussion and Summary

This section provides discussion of on-going formative evaluations and its initial conclusions, as well as limitations, future work, and summary.

### 6.5.1 Formative Evaluations of VisualQUEST

We have had formative user evaluations of VisualQUEST—that is, ongoing, task-driven assessments. These evaluations were informal involving volunteers associated with our research lab. They have provided initial insights into how ontology-supported and progressively disclosed VAT interfaces generated by the novel collection of design criteria can help users perform complex, multi-staged information-seeking tasks on large document sets. In these evaluations, users are asked to perform tasks aligning with the stages of the information-seeking process, including the fulfillment of a research, question-driven scoping review. The restrictions of the on-going COVID-19 pandemic

placed significant limitations on our available scope of evaluation. This is described in detail with the Limitations section. Our research objectives for expanded evaluations beyond the COVID-19 pandemic are described within the Future Work section.

### **6.5.1.1 Experimental Settings**

Users were asked to complete seven tasks using their assigned interface with the same document set and ontology. Users were directed through an automated task set. Users were tasked with exploring how their interface could help them relate their knowledge into a domain vocabulary and apply their information seeking objectives when query building. Next, users were tasked with assessing how their queries align with a document set, and how such alignments may impact search quality. Third, users were asked to perform high-level triage on search results, followed by opportunities to perform low-level triage. The task set culminated with the completion of a task involving all stages of information search and triage, guided by a research question. We provide a general description of each task within the task set (Table 3).

**Table 6-3 General description of tasks performed during formative evaluations.**

| <b>T#</b> | <b>Target Stage</b> | <b>Task Description</b>   |
|-----------|---------------------|---|
| T1        | Query Building      | The first task asked users to consider two terms and contrast their rate of occurrence within the document set.   |
| T2        | Query Building      | The second task asked users to consider a term and determine its alignment with a set of provided definitions.  |
| T3        | Search              | The third task asked users to consider how provided set of terms aligned with the document set, both individually and in combinations.  |
| T4        | High-Level Triage   | The fourth task brought users to a specific document, and without allowing them to open the full document, were asked to predict its alignment to a provided set of terms.  |
| T5        | High-Level Triage   | The fifth task brought users to a specific pair of documents, and without opening them, were asked to compare and then predict which of them would contain a higher rate of occurrence of a specific term.  |
| T6        | Low-Level Triage    | The sixth task brought users to a specific document and were asked to count and order the rate of occurrences of a provided set of terms within that document.  |
| T7        | Multi-Staged        | The seventh task gave users a domain research question and asked them to produce five relevant documents from the document set. This task required users to progress through each stage of the information-seeking process, requiring the use of all available functionalities of their interface. A topic background was provided for optional domain context. |

### **6.5.1.2 Formative Evaluation Metrics and Conclusions**

Through initial sessions, formative evaluation metrics were generated, bounded the limitations of COVID-19 restrictions. We describe an expanded set of evaluation metrics within the Future Work section. During on-going

formative evaluations, task completion metrics were informally accounted, noting that all users completed each assigned task. Furthermore, users were asked about their general experiences after using their assigned tool, akin to qualitative usability metrics typical of user studies. Users described their experiences with progressive disclosure, the use of novel visual abstractions, view sequences, and ontology mediations. We itemize initial conclusions below, followed by some informal quotes which informed those conclusions (A ... J):

- 1) Users are able to translate their experience with traditional interfaces to the use of unfamiliar interfaces (*e.g.*, A, B).
- 2) Users are able to interpret complex visual abstractions to heighten their participation in the analytic process and form a stronger connection to their ML tool in contrast to traditional “black box” approaches (*e.g.*, C, D).
- 3) Users are able to differentiate between the individual stages of their information-seeking process and utilize VisualQUEST’s domain-independent, progressively disclosed interface to search and triage large document sets (*e.g.*, E, F, G).
- 4) Users are able to utilize ontology files when aligning their vocabulary with the vocabulary of the domain, even if they are not initially familiar with the ontology’s domain or its structure and content (*e.g.*, H, I, J).
- 5) Users felt that mediating ontologies make search tasks more manageable and easier, and not having them would negatively affect their task performance (*e.g.*, H, I, J).

The following are excerpts from these informal sessions:

A) *“Once I understood what it was showing me, it helped me. Usually with new tools I tend to read through the documentation or watch videos. And then it still takes me like a while to pick up on them. Like, just running through them and using them a few times. Once you get the hang of it, usually you find success in whatever it's providing you.”*

B) *“It's a tool that I'm not used to and I'm kind of going back and forth and for me when there's kind of a lot of little moving parts. I mean it, it feels that way right now because I'm not familiar with the tool and I'm kind of taking in this information and figure out how the way different pieces of information needs to go together.”*

C) *“If I was using (VisualQUEST) against other search (interfaces), I would just use (VisualQUEST) constantly. Being able to really filter down exactly what I need ... like that's really on point. Especially with like the different colorings of the words. It's like telling you like what each document (grouping) is like. It gives me the ability to better align with the documents and gives me more confidence when I'm creating my queries. I'm making the correct query decision even before even running the search.”*

D) *“A lot of other tools you use, they kind of do predictive searches for you. They build a filter. So, for example, like Google doing a predictive search. It's predicting based on top results from previous searches, so with that, it can be finagled with. Where you know you could have a bot farm or whatever finagling those search results and making them be what they want them to be. Whereas with (VisualQUEST) you get to parse those results and make your own educated decisions versus (the search engine) doing it for you. So, I feel more control in the experience then you would typically. I feel more certain in the end goal.”*

E) *“(With VisualQUEST), going through low-level triage and seeing and reading the abstract and the actual documents... I wanted to check and have a comparison between the documents that I chose... (because) usually my style of choosing (is that) I usually choose more than what I have to choose and then I remove that and the extra ones.”*

F) *“At first, I just copy and pasted the entire keywords ... and my first thought I wasn't really seeing the results of the all the keywords mixed in together the same way. So, (with VisualQUEST) I was able to go back and see if the words individually not together had brought in any difference in the search.”*

G) *“I think (VisualQUEST) is a benefit. In my previous experiences with searching queries, if you just type it in and it blurts out the answer it prioritizes in whatever way that it wanted. I like this because it is a little bit more specific, and you are able to choose more of the options that you want to use. I think you have more control of how to exactly to find the answer and what exactly you are looking for, instead of starting from just a general basis of all the answers that are possible. So, I think it's better to be able to narrow down exactly what you're looking for and find more appropriate answer towards your question.”*

H) *“I was thinking ... where (the ontology) would have been helpful. So ... it would have possibly brought up some of those other terms just from searching a few words and they would be able to make some connections between the text that was provided and some of my search terms (to see) ... how relevant they were. So, if I was shooting in the dark and hoping for the best, which is what I was kind of doing (without the ontology), at the very least, it would have given you confidence of your actions. Yeah, I think so. A little bit more confidence.”*

I) *“Yeah, actually on second thought, yeah, this (ontology) would have helped because I ... can find the things that ... share in common, and that can make it probably much easier to find the relevant documents. Yeah, being able to see the things that certain phrases ... or words share in common. You can find that common link ... that can find you the ... the relevant documents.”*

J) *“I do really like how (VisualQUEST) gives you the ability to find different vocabulary or other words that may not have been the first thing you thought of when you were building the query. “*

## **6.5.2 Limitations**

The first and most primary limitation of this material is the scope of evaluation. The active COVID-19 pandemic has restricted access to study resources, user pools, and study environments spaces. Therefore, at this time, we do not have the ability to perform expanded formal user studies which could highlight further qualitative and quantitative evaluation metrics. Once pandemic conditions end, we intend to address this limitation by conducting expanded formal user studies of VisualQUEST. We describe this in greater detail within the Future Work section.

We also highlight technical limitations. First, the Analytics server of VisualQUEST can handle uploading, processing, and serving document sets and ontology files of large sizes. For example, our usage scenario demonstrates VisualQUEST handling of HPO and its 11000 ontology terms and a large document set reflecting a subset of MEDLINE. However, if document sets and ontology files were to be increased to an extreme scale, overhead limits within the local browser could produce a notable wait before users could begin to search and triage. Therefore, the first limitation of our research is VisualQUEST's current ability to handle document sets and ontology files of extremely large sizes. Additional technical efforts would be needed to address this limitation. For instance, additional

efforts could be made to shift computation away from the local browser, improvements to general computational efficiency, and seeking centralized solutions to eliminate document set and ontology file indexing prior to task performance. VisualQUEST's second technical limitation is its current level of support for ontology file formats. As previously described, VisualQUEST in its current state can process the core elements of OWL, a leading format for encoding ontologies within the digital space. Yet, the OWL specification is overly verbose, particularly in regard to its extensive base of axiom relations. Therefore, we believe VisualQUEST's ontology processing system can be improved to supply even more value for mediation and query expansion opportunities. In addition, there are other RDF-based ontology formats which would be valuable to support.

### **6.5.3 Future Work**

Once COVID-19 pandemic restrictions are relaxed, we plan to continue our research efforts with formal, empirical evaluations. These evaluations will implement task-driven formal evaluations which compare VisualQUEST to other interfaces that facilitate searching and triaging large document sets. We hope to generate qualitative and quantitative results within a formal evaluation setting to provide insight into the impact interface design can have on users when performing information-seeking tasks. Specifically, we seek to expand our evaluation metrics to also include quantitative results which track how user performance changes when presented with alternative interfaces for a task set. During user study sessions, we intend to log metrics such as task performance scores, task completion timings, expand qualitative metrics through user-reported ease, satisfaction, and assessment logs, as well as more in-depth interview sessions. From these efforts, we believe criteria can be bolstered and in turn expanded upon to generate prescriptive guidelines and frameworks for the design of visual interfaces.

Beyond an immediate scope, we assess several potential research directions. First, we believe there is value in deeper investigations of lower-level design considerations and their impact on the performance of challenging information-seeking tasks on large documents. Second, the criteria and the demonstrative interface VisualQUEST provide an initial exploration into how ontology mediation can be presented to information seekers within the visual interface of their tool. Yet we believe more can be done to establish novel designs which further ontology mediation opportunities. Future research may explore additional points of ontology integration such as for the stages of high and low-level triage within the multi-staged information-seeking process. Third, information-seeking tasks can sometimes require refined levels of domain knowledge for effective performance. Future research could investigate how domain-specific considerations affect the performance of information search and triage, and what design approaches can be used to benefit those requirements.

### **6.5.4 Summary**

We investigated the design of ontology-supported, progressively disclosed visual analytics interfaces for searching and triaging large document sets, distilled design criteria, and with its guidance generated a demonstrative visual interface. That is, we proposed the following research questions:

- What are the criteria for the design of VAT interfaces that support the process of searching and triaging large document sets?

- If such criteria can be distilled, can they be used to help guide the design of a progressively disclosed and ontology-supported interface?

We began with a background discussion of information search, information triage, machine learning, and ontologies. We reviewed leading research on the multi-staged information-seeking process to distill high-level design criteria. In this review, investigations of existing models of the information-seeking process distilled the requirements of a four-stage model encompassing query building, search, high-level triage, and low-level triage. The organizational design technique of progressive disclosure was assessed for its value for complex multi-staged tasks, such as those involving the information-seeking process. Next, leading research related to the individual stages of the information-seeking process and their requirements were reviewed, generating best practices in interface design. These findings were then distilled into 11-part criteria for interface designers, each accompanied by a description of ontology integration value.

To illustrate the utility of the criteria, we applied them to the design of a demonstrative prototype: VisualQUEST (Visual interface for QUery, Search, and Triage). VisualQUEST enables users to build queries, search, and triage document sets both at a high as well as low levels. Users can plug-and-play document sets and expert-defined ontology files within a domain-independent, progressively disclosed environment for multi-staged information search and triage tasks. We described VisualQUEST through a functional workflow.

We culminated with discussion of on-going formative evaluations, limitations, future work, and summary. Initial evaluations have found that users have responded positively to the design criteria applied within VisualQUEST. Namely, their experiences with progressive disclosure and the use of novel visual abstractions. As well as multi-stage sequences, and ontology mediation when using VisualQUEST for searching and triaging large document sets. We hope to continue to promote novel thinking for VAT interface design within our future research, such as those for complex multi-staged information-seeking tasks on large document sets.

## 6.6 References

- Amershi, S., Cakmak, M., Knox, W. B., & Kulesza, T. (2014). Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine*, 35(4), 105–120. <https://doi.org/10.1609/aimag.v35i4.2513>
- Anderson, J. D., & Wischgoll, T. (2020). Visualization of Search Results of Large Document Sets. *IS and T International Symposium on Electronic Imaging Science and Technology*, 2020(1), 1–7. <https://doi.org/10.2352/ISSN.2470-1173.2020.1.VDA-388>
- Arp, R., Smith, B., Spear, A. D., & American Journal of Sociology. (2015). Building Ontologies with Basic Formal Ontology. In *Journal of Chemical Information and Modeling* (Vol. 53).
- Azzopardi, L., & Zuccon, G. (2016). An Analysis of the Cost and Benefit of Search Interactions. *ICTIR 2016 - Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval*, 59–68. <https://doi.org/10.1145/2970398.2970412>
- Badi, R., Bae, S., Moore, J. M., Meintanis, K., Zacchi, A., Hsieh, H., ... Marshall, C. C. (2006). Recognizing User Interest and Document Value from Reading and Organizing Activities in Document Triage. *Proceedings of the 11th International Conference on Intelligent User Interfaces - IUI '06*, (1), 218.

<https://doi.org/10.1145/1111449.1111496>

- Bae, S., Kim, D., Meintanis, K., Moore, J. M., Zacchi, A., Shipman, F., ... Marshall, C. C. (2010). Supporting Document Triage via Annotation-based Multi-application Visualizations. *Proceedings of the 10th ACM/IEEE Joint Conference on Digital Libraries*, 177–186. <https://doi.org/10.1145/1816123.1816150>
- Boschee, E., Barry, J., Billa, J., Freedman, M., Gowda, T., Lignos, C., ... Miller, S. (2019). Saral: A Low-resource cross-lingual domain-focused information retrieval system for effective rapid document triage. *ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics, Proceedings of System Demonstrations*, 19–24. <https://doi.org/10.18653/v1/p19-3004>
- Bostock, M. (2016). D3.js Data-Driven Documents. Retrieved May 10, 2020, from <https://d3js.org/>
- Buchanan, G., & Owen, T. (2008). Improving Skim Reading for Document Triage. *Proceedings of the Second International Symposium on Information Interaction in Context - IiX '08*, 83. <https://doi.org/10.1145/1414694.1414714>
- Celebi, M. E., & Aydin, K. (2016). *Unsupervised learning algorithms*. Berlin: Springer International Publishing.
- Chandrasegaran, S., Badam, S. K., Kisselburgh, L., Ramani, K., & Elmquist, N. (2017). Integrating Visual Analytics Support for Grounded Theory Practice in Qualitative Text Analysis. *Computer Graphics Forum*, 36(3), 201–212. <https://doi.org/10.1111/cgf.13180>
- Chuang, J., Ramage, D., Manning, C., & Heer, J. (2012). Interpretation and Trust: Designing Model-driven Visualizations for Text Analysis. *Proceedings of the SIGCHI ...*, 443–452. Retrieved from <http://dl.acm.org/citation.cfm?id=2207738>
- Cortez, P., & Embrechts, M. J. (2013). Using Sensitivity Analysis and Visualization Techniques to Open Black Box Data Mining Models. *Information Sciences*, 225, 1–17. <https://doi.org/10.1016/j.ins.2012.10.039>
- Demelo, J., Parsons, P., & Sedig, K. (2017). Ontology-Driven Search and Triage: Design of a Web-Based Visual Interface for MEDLINE. *JMIR Medical Informatics*, 5(1), e4. <https://doi.org/10.2196/medinform.6918>
- Demelo, J., Sedig, K., & Parsons, P. (2017). OVERT-MED : An Ontology-Driven Visual Search and Triage Interface for MEDLINE. *Journal of Medical Internet Research Medical Informatics*, 5(1). <https://doi.org/10.2196/medinform.6918>
- Dey, A. (2016). Machine Learning Algorithms: A Review. *International Journal of Computer Science and Information Technologies*, 7(3), 1174–1179. Retrieved from [www.ijcsit.com](http://www.ijcsit.com)
- Endert, A., Ribarsky, W., Turkay, C., Wong, B. L. W., Nabney, I., Blanco, I. D., & Rossi, F. (2017). The State of the Art in Integrating Machine Learning into Visual Analytics. *Computer Graphics Forum*, 36(8), 458–486. <https://doi.org/10.1111/cgf.13092>
- Endert, Alex, Hossain, M. S., Ramakrishnan, N., North, C., Fiaux, P., & Andrews, C. (2014). The Human is the Loop: New Directions for Visual Analytics. *Journal of Intelligent Information Systems*, 43(3), 411–435. <https://doi.org/10.1007/s10844-014-0304-9>
- Fiebrink, R., Cook, P. R., & Trueman, D. (2011). Human Model Evaluation in Interactive Supervised Learning. *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems - CHI '11*, 147. <https://doi.org/10.1145/1978942.1978965>

- Golitsyna, O. L., Maksimov, N. V., & Monankov, K. V. (2018). Focused on Cognitive Tasks Interactive Search Interface. *Procedia Computer Science*, 145, 319–325. <https://doi.org/10.1016/j.procs.2018.11.078>
- Harvey, M., Hauff, C., & Elsweiler, D. (2015). Learning by Example: Training Users with High-quality Query Suggestions. *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, New York, NY, USA, 9–13 August 2015*, 133–142.
- Herceg, P. M., Allison, T. B., Belvin, R. S., & Tzoukermann, E. (2018). Collaborative Exploratory Search for Information Filtering and Large-scale Information Triage. *Journal of the Association for Information Science and Technology*, 69(3), 395–409. <https://doi.org/10.1002/asi.23961>
- Hoeber, O. (2014). Visual Search Analytics: Combining Machine Learning and Interactive Visualization to Support Human-Centred Search. *Beyond Single-Shot Text Queries: Bridging the Gap(s) Between Research Communities*, 37–43.
- Hohman, F., Kahng, M., Pienta, R., & Chau, D. H. (2018). Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers. *IEEE Transactions on Visualization and Computer Graphics*, 25(8), 2674–2693.
- Holzinger, A. (2014). *Knowledge Discovery and Data Mining in Biomedical Informatics: The Future Is in Integrative, Interactive Machine Learning Solutions*. 8401(March 2016). <https://doi.org/10.1007/978-3-662-43968-5>
- Holzinger, A. (2016). Interactive Machine Learning for Health Informatics: When Do We Need the Human-in-the-loop? *Brain Informatics*, 3(2), 119–131. <https://doi.org/10.1007/s40708-016-0042-6>
- Hurdeman, H. C. (2017). Dynamic Compositions: Recombining Search User Interface Features for Supporting Complex Work Tasks. *CEUR Workshop Proceedings*, 1798, 22–25.
- Hurdeman, H. C., Wilson, M. L., & Kamps, J. (2016). Active and passive utility of search interface features in different information seeking task stages. *CHIIR 2016 - Proceedings of the 2016 ACM Conference on Human Information Interaction and Retrieval*, 3–12. <https://doi.org/10.1145/2854946.2854957>
- Jakus, G., Milutinovic, V., Omerović, S., & Tomazic, S. (2013). *Concepts, Ontologies, and Knowledge Representation*. <https://doi.org/10.1007/978-1-4614-7822-5>
- Jusoh, S., Awajan, A., & Obeid, N. (2020). The Use of Ontology in Clinical Information Extraction. *Journal of Physics: Conference Series*, 1529(5). <https://doi.org/10.1088/1742-6596/1529/5/052083>
- Katifori, A., Torou, E., Vassilakis, C., Lepouras, G., & Halatsis, C. (2008). Selected Results of a Comparative Study of Four Ontology Visualization Methods for Information Retrieval Tasks. *Proceedings of the Second International Conference on Research Challenges in Information Science, Marrakech, Morocco, 3–6 June 2008*, (September 2016), 133–140. <https://doi.org/10.1109/RCIS.2008.4632101>
- Kekäläinen, J. (2014). Browsing Patterns in Retrieved Documents. *Proceedings of the 5th Information Interaction in Context Symposium*, 299–302.
- Khan, S., Kanturska, U., Waters, T., Eaton, J., Bañares-Alcántara, R., & Chen, M. (2016). Ontology-assisted Provenance Visualization for Supporting Enterprise Search of Engineering and Business Files. *Advanced Engineering Informatics*, 30(2), 244–257. <https://doi.org/10.1016/j.aei.2016.04.003>
- Loizides, F. (2012). Information Seekers' Visual Focus During Time Constraint Document Triage. *Lecture Notes in*

- Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 7489 LNCS, 25–31. [https://doi.org/10.1007/978-3-642-33290-6\\_3](https://doi.org/10.1007/978-3-642-33290-6_3)
- Loizides, F., & Buchanan, G. (2009). *What Patrons Want : Supporting Interaction for Novice Information Seeking Scholars*. 60558.
- Loizides, F., Buchanan, G., & Mavri, K. (2016). Theory and Practice in Visual Interfaces for Semi-structured Document Discovery and Selection. *Information Services & Use*, 35(4), 259–271. <https://doi.org/10.3233/ISU-150783>
- Lytvyn, V., Dosyn, D., Vysotska, V., & Hryhorovych, A. (2020). Method of Ontology Use in OODA. *Proceedings of the IEEE Third International Conference on Data Stream Mining & Processing, Lviv, Ukraine, 21–25 August 2020*, 409–413. <https://doi.org/10.1109/dsmp47368.2020.9204107>
- Mavri, A., Loizides, F., Photiadis, T., & Zaphiris, P. (2013). *We Have the Content ... Now What?* 20(3), 247–265.
- Mehta, N., & Pandit, A. (2018). Concurrence of Big Data Analytics and Healthcare: A Systematic Review. *International Journal of Medical Informatics*, 114(March), 57–65. <https://doi.org/10.1016/j.ijmedinf.2018.03.013>
- Munir, K., & Sheraz Anjum, M. (2018). The Use of Ontologies for Effective Knowledge Modelling and Information Retrieval. *Applied Computing and Informatics*, 14(2), 116–126. <https://doi.org/10.1016/j.aci.2017.07.003>
- Ninkov, A., & Sedig, K. (2019). VINCENT: A Visual Analytics System for Investigating the Online Vaccine Debate. *Online Journal of Public Health Informatics*, 11(2). <https://doi.org/10.5210/ojphi.v11i2.10114>
- Ola, O., & Sedig, K. (2016). Beyond Simple Charts: Design of Visualizations for Big Health Data. *Online Journal of Public Health Informatics*, 8(3). <https://doi.org/10.5210/ojphi.v8i3.7100>
- Parsons, P., Sedig, K., Mercer, R. E., Khordad, M., Knoll, J., & Rogan, P. (2015). Visual analytics for supporting evidence-based interpretation of molecular cytogenomic findings. *Proceedings of the 2015 Workshop on Visual Analytics in Healthcare - VAHC '15*, 1–8. <https://doi.org/10.1145/2836034.2836036>
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Duchesnay, É. (2012). *Scikit-learn: Machine Learning in Python*. 12, 2825–2830. <https://doi.org/10.1007/s13398-014-0173-7.2>
- Ramanujan, D., Chandrasegaran, S. K., & Ramani, K. (2017). Visual Analytics Tools for Sustainable Lifecycle Design: Current Status, Challenges, and Future Opportunities. *Journal of Mechanical Design, Transactions of the ASME*, 139(11). <https://doi.org/10.1115/1.4037479>
- Rector, A., Schulz, S., Rodrigues, J. M., Chute, C. G., & Solbrig, H. (2019). On beyond Gruber: “Ontologies” in Today’s Biomedical Information Systems and the Limits of OWL. *Journal of Biomedical Informatics: X*, 2, 100002. <https://doi.org/https://doi.org/10.1016/j.yjbix.2019.100002>
- Román-Villarán, E., Pérez-Leon, F. P., Escobar-Rodriguez, G. A., Martínez-García, A., Álvarez-Romero, C., & Parra-Calderón, C. L. (2019). An Ontology-based Personalized Decision Support System for Use in the Complex Chronically Ill Patient. *Studies in Health Technology and Informatics*, 264, 758–762. <https://doi.org/10.3233/SHTI190325>
- Sacha, D., Sedlmair, M., Zhang, L., Lee, J. A., Weiskopf, D., North, S., & Keim, D. (2016). Human-centered Machine Learning Through Interactive Visualization: Review and Open Challenges. *ESANN 2016: 24th European*

- Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning Bruges, Belgium April 27-28-29, 2016 Proceedings*, (27–29 April 2016), 641–646. Retrieved from <http://eprints.mdx.ac.uk/20763/>
- Saha, D., Floratou, A., Sankaranarayanan, K., Minhas, U. F., Mittal, A. R., Özcan, F., & Deacy, S. (2016). ATHENA: An Ontology-Driven System for Natural Language Querying over Relational Data Stores. *Proceedings of the VLDB Endowment*, 9(12), 1209–1220. <https://doi.org/10.4324/9780203932148>
- Saleemi, M. M., Rodríguez, N. D., Lilius, J., & Porres, I. (2011). A Framework for Context-aware Applications for Smart Spaces. In *Smart Spaces and Next Generation Wired/Wireless Networking: Vol. 6869 LNCS* (pp. 14–25). [https://doi.org/10.1007/978-3-642-22875-9\\_2](https://doi.org/10.1007/978-3-642-22875-9_2)
- Sedig, K., & Ola, O. (2014). The Challenge of Big Data in Public Health: An Opportunity for Visual Analytics. *Online Journal of Public Health Informatics*, 5(3), 1–21. <https://doi.org/10.5210/ojphi.v5i3.4933>
- Sedig, K., & Parsons, P. (2016). Design of Visualizations for Human-Information Interaction: A Pattern-Based Framework. In *Synthesis Lectures on Visualization* (Vol. 4). <https://doi.org/10.2200/S00685ED1V01Y201512VIS005>
- Soldaini, L., Yates, A., Yom-Tov, E., Frieder, O., & Goharian, N. (2016). Enhancing Web Search in the Medical Domain via Query Clarification. *Information Retrieval*, 19(1–2), 149–173. <https://doi.org/10.1007/s10791-015-9258-y>
- Solr Cloud. (2020). Retrieved May 10, 2020, from <https://lucene.apache.org/solr/>
- Spence, R. (2002). Sensitivity Encoding to Support Information Space Navigation: A Design Guideline. *Information Visualization*, 1(2), 120–129. <https://doi.org/10.1057/palgrave.ivs.9500019>
- Spence, R. (2004). Sensitivity, Residue and Scent. *Information Design Journal + Document*, 12(3). <https://doi.org/10.1075/idjdd.12.3.01spe>
- Spence, R. (2014). Interaction. In *Information Visualization* (pp. 173–223). Springer.
- Springer, A., & Whittaker, S. (2018). *Progressive Disclosure: Designing for Effective Transparency*. Retrieved from <http://arxiv.org/abs/1811.02164>
- Sun, G., Lv, H., Wang, D., Fan, X., Zuo, Y., Xiao, Y., ... Guo, Z. (2020). Visualization analysis for business performance of chinese listed companies based on gephi. *Computers, Materials and Continua*, 63(2), 959–977. <https://doi.org/10.32604/cmc.2020.08619>
- Talbot, J., Lee, B., Kapoor, A., & Tan, D. S. (2009). EnsembleMatrix: Interactive Visualization to Support Machine Learning with Multiple Classifiers. *Proceedings of the 27th International Conference on Human Factors in Computing Systems - CHI 09*, 1283. <https://doi.org/10.1145/1518701.1518895>
- Tang, Q., Wylie, S., Hoover, J., Ghamsari, A., Price, M., & Dagley, G. (2019, November 26). *Techniques to Leverage Machine Learning for Search Engine Optimization*. Google Patents.
- Tobergte, D. R., & Curtis, S. (2013). A Survey on Ontologies for Human Behavior Recognition. *Journal of Chemical Information and Modeling*, 53(9), 1689–1699. <https://doi.org/10.1017/CBO9781107415324.004>
- Tresp, V., Marc Overhage, J., Bundschuh, M., Rabizadeh, S., Fasching, P. A., & Yu, S. (2016). Going Digital: A Survey on Digitalization and Large-Scale Data Analytics in Healthcare. *Proceedings of the IEEE*, 104(11),

2180–2206. <https://doi.org/10.1109/JPROC.2016.2615052>

- Vamathevan, J., Clark, D., Czodrowski, P., Dunham, I., Ferran, E., Lee, G., ... Spitzer, M. (2019). Applications of Machine Learning in Drug Discovery and Development. *Nature Reviews Drug Discovery*, 18(6), 463–477.
- Wall, E., Blaha, L. M., Franklin, L., & Endert, A. (2018). Warning, Bias May Occur: A Proposed Approach to Detecting Cognitive Bias in Interactive Visual Analytics. *Proceedings of the IEEE Conference on Visual Analytics Science and Technology, Phoenix, AZ, USA, 3–6 October, 2017*, 104–115. <https://doi.org/10.1109/VAST.2017.8585669>
- Wang, J., Li, M., Diao, Q., Lin, H., Yang, Z., & Zhang, Y. J. (2020). Biomedical document triage using a hierarchical attention-based capsule network. *BMC Bioinformatics*, 21(Suppl 13), 1–20. <https://doi.org/10.1186/s12859-020-03673-5>
- Wei, J., Chu, X., Sun, X., Xu, K., Deng, H., Chen, J., ... Lei, M. (2019). Machine Learning in Materials Science. *InfoMat*, 1(3), 338–358.
- William Lidwell, Kritina Holden, J. B. (2010). Progressive Disclosure. In *Universal Principles of Design* (pp. 188–189). Rockport Publishers.
- Wu, C. M., Meder, B., Filimon, F., & Nelson, J. D. (2017). Asking Better Questions: How Presentation Formats Influence Information Search. *Journal of Experimental Psychology: Learning Memory and Cognition*, 43(8), 1274–1297. <https://doi.org/10.1037/xlm0000374>
- Xing, Z., Wang, L., Xing, W., Ren, Y., Li, T., & Xia, J. (2019). Application of Ontology in the Web Information Retrieval. *Journal on Big Data*, 1(2), 79–88. <https://doi.org/10.32604/jbd.2019.05806>
- Yuan, J., Chen, C., Yang, W., Liu, M., Xia, J., & Liu, S. (2020). A Survey of Visual Analytics Techniques for Machine Learning. *ArXiv*, 1, 1–34.
- Zeng, Q. T., & Tse, T. (2006). Exploring and Developing Consumer Health Vocabularies. *Journal of the American Medical Informatics Association*, 13(1), 24–29. <https://doi.org/10.1197/jamia.M1761>

# Chapter 7      Evolutionary Design of Search and Triage

## Interfaces for Large Document Sets: A Formative Assessment

This chapter has been prepared for submission as: **Demelo, J., & Sedig, K. (2022).** Evolutionary Design of Search and Triage Interfaces for Large Document Sets: A Formative Assessment.

We have made minor adjustments to the original material of this chapter to provide cohesion with the overall integrated article structure of this dissertation. Specifically, to distinguish between chapters, figures and tables have been provided an additional prepend reflecting the chapter number. Readers should be aware that chapter text will maintain original numbering references. For instance, “Figure 7-1” is equivalent to “Figure 1” in the chapter text.

### 7.1 Introduction

Visual analytics combines interactive visualization with powerful computational technologies to help users synthesize meaning from information, explore new ways to apply analytical reasoning, and become more active participants in the analytics process (Golitsyna, Maksimov, & Monankov, 2018; Ramanujan, Chandrasegaran, & Ramani, 2017). It is both user- and data-driven, interlinking a complex set of computational, data management, interaction, visualization, and human factor requirements. Designers create visual analytics tools (VATs) to help users connect to valuable information and promote decision making. These tools maintain interactive visual interfaces where information can be viewed, manipulated, re-computed using computational technologies, and re-displayed back to users. Machine learning (ML) is increasingly being utilized to address the growing computational needs of analytics tasks, such as when searching and triaging large document sets (Talbot, Lee, Kapoor, & Tan, 2009). These technologies can be invaluable for increasing the computational power of VATs yet if not appropriately communicated, can be equally damaging to users’ analytic reasoning. In particular, recent studies (A. Endert et al., 2017; Hohman, Kahng, Pienta, & Chau, 2018; Yuan et al., 2020) have found that the application of traditional interface design strategies for ML-supported VATs result in user experiences which weaken participation in the analytic process, as well as limit the capacity to understand, control, and be satisfied with task performances. For these concerns, current research (Alex Endert et al., 2014; Hohman et al., 2018; Wall, Blaha, Franklin, & Endert, 2018) promotes the benefits of user-centered design, and in particular, the re-activation of the “human-in-the-loop” for VAT interfaces (Sedig & Ola, 2014; Sedig & Parsons, 2016).

When searching and triaging large document sets, users desire to encounter, judge, and extract documents from a larger set in a rapid, yet productive manner. When using traditionally designed interfaces, studies (Harvey, Hauff, & Elswiler, 2015) have shown that users routinely struggle to understand their searched domain, use their expertise to communicate their information-seeking objectives, and in turn assess the relevance of search results during information triage. If users cannot effectively communicate information-seeking objectives, their tool will fail to perform best analysis on the document set being searched, reducing the quality of search results. Furthermore, if users cannot understand the domain, nor understand how the tool produced its results, users will struggle to effectively make

rapid relevance decisions during information triage (Herceg, Allison, Belvin, & Tzoukermann, 2018; Wu, Meder, Filimon, & Nelson, 2017).

In this paper, we present the formulization, realization, and validation efforts of a three-staged evolutionary design, producing a VAT interface for searching and triage large document sets:

- Stage 1: Traditionally designed search and triage interface.
- Stage 2: Progressively disclosed search and triage interface.
- Stage 3: Progressively disclosed, ontology-supported search and triage interface.

These interfaces allow users to upload a document set, and if supported, an ontology file for vocabulary mediation. Then users may begin to query build and finalize a search formulation. Upon directing systems to perform ML computations for computational search, a document set mapping is returned to users for triaging.

We begin with background on information search and triage within machine learning (ML)-supported VATs, where we examine challenges facing users and potential solutions. We outline the evolutionary design process and specify a task-driven formative assessment. In particular, we specify the document set and ontology files used, that of a static subset of MEDLINE digital library along with the Human Phenotype Ontology (HPO) for ontology-supported interfaces. We then describe the results of our evolutionary design, spanning the formulization, realization, and validation of three VAT interfaces: Stage 1, Stage 2, and Stage 3. We provide a general discussion of user responses to the evolutionary design, the value of ontology-supported interfaces for information search, and the promotion of progressive disclosure in interfaces for multi-staged information search and triage on large document sets.

The structure of this paper continues: Section 2 provides topic background. Section 3 describes materials and methods encompassing the evolutionary design process and a task-driven formative assessment. Section 4 provides the results of our evolutionary design, spanning the formulization, realization, and validation of three VAT interfaces: Stage 1, Stage 2, and Stage 3. Section 5 provides discussion of general implications, limitations, and a concluding summary.

## **7.2 Background**

This section first provides background on evolutionary design. This is followed by background on information search and triage within ML supported VATs, examining the responsibilities of users, what challenges present, and potential solutions for addressing those challenges.

### **7.2.1 Evolutionary Design**

Designers must align their approach to the requirements of the problem space, affecting how potential solutions can be formulized, realized, and validated. This is a particular challenge for VAT designers, who must account for requirements arising from a multitude of user, data, task, and tool constraints (Zhang et al., 2021). Many models exist which attempt to standardize design processes, such as breath-first strategies which perform a wide exploration of the design space prior to deeper prototyping efforts, and depth-first strategies which begin with comprehensive yet narrow prototyping on controlled positions of the design space (Stouffs & Rafiq, 2015). Yet as problem spaces continue to grow in scope and complexity, it has become increasingly costly, both in terms of time and effort, to address them in

a non-iterative manner. That is, the number of complex, interlinked requirements of ill-defined problem spaces generally do not allow for optimized, predictable design solutions prior to first verification (Guerrero-García, 2014). Instead, designers must plan for numerous configurations prior to the generation of a solution which satisfies requirements.

For this, evolutionary design models are increasingly being used to address the ill-defined design problems and the challenges which rise. For example, evolutionary models are being activated for the design of cloud systems, neural networks, and interface workflows (Baldominos, Saez, & Isasi, 2020; Guerrero-García, 2014; Zhang et al., 2021). Models for evolutionary design, inspired by biological origins, encapsulate a four-staged, iterative process of formulization, realization, validation, and refinement (Baldominos et al., 2020). First, with guidance from prior interfaces, research, and framing devices, designers distill the necessary requirements of the task, decide how they are best addressed within the interface and its components, and formulize decisions within a design. From this design, a working prototype is realized, which is then validated through user-based methods of formative assessment. The findings of this verification are then propagated back to the designer for further formulization, realization, and verification, until requirements are satisfied (Guerrero-García, 2014). Through evolutionary design, VAT interface designers can provide more opportunities for novel thinking, de-couple design from prototyping, and promote user-centered design through formative assessment. In this paper, we describe the structure and findings of an evolutionary design.

### **7.2.2 Information Search and Triage in ML-supported VAT interfaces**

When searching a document set using VAT interfaces, users' primary desire is to encounter documents most relevant to their task. For this, users are required to describe their information-seeking needs to the VAT through its interface. Once communicated, the computational components can then generate a mapping between users' input and the qualified and relevant documents in the document set; presented back to users at the interface level (Wu et al., 2017).

Current research (Harvey et al., 2015) describes user requirements when searching and triaging large document sets, which can be used to formulize the characteristics of an interface design which best supports users as they perform information search on large document sets. To begin, users must establish an understanding of the document set being searched and how it relates to their existing domain knowledge. Then, users must learn how to effectively communicate their information-seeking objectives in a way that can be understood by the tool. Optimal search (and triage) can only be achieved when communication between both human and computational components of VATs are strong (Arp, Smith, Spear, & American Journal of Sociology, 2015). That is, user performance is at its best only when what the interface is displaying is understandable, just as a tool can only optimize its computations, visualizations, and interactions if user-supplied instructions truly align with their analytic intentions. VATs are typically designed to support tasks maintaining static data sources with consistent vocabulary and are not typically adaptable to changing data sources and variable task vocabularies. In these settings, it is users' responsibility to learn the tool's domain-specific vocabulary and then apply that required understanding to communicate information-seeking objectives. Yet, learning unfamiliar vocabulary can be a significant challenge for users, particularly in tasks of complex domains (e.g., health) which present both domain experts and non-experts a significant barrier of entry when

trying to understand a lexicon and its unique linguistic considerations (Zeng & Tse, 2006). For these issues, expert-defined ontologies are increasingly being targeted as mediating resources within visual analytics and ML computation components (Xing et al., 2019). Ontologies are representational artifacts which can be leveraged both by the computational and human-facing systems to describe a standardized mapping of the entities, relations, and structures of a domain (Jakus, Milutinovic, Omerović, & Tomazic, 2013; Khan et al., 2016; Rector, Schulz, Rodrigues, Chute, & Solbrig, 2019). Some examples of ontology uses are information extraction, behavior modeling, and decision support systems within critical care environments (Jusoh, Awajan, & Obeid, 2020; Lytvyn, Dosyn, Vysotska, & Hryhorovych, 2020; Román-Villarán et al., 2019). Experts create ontologies by navigating the terms, relations, and contextualizing characteristics of their domain, and through those efforts, formulate a generalized knowledge map of complex ontological space (Arp et al., 2015). Ontology entities encode information about their role in the domain, often with metadata such as its vocabulary, definition, and description. Ontology relations are the links between entities that express interactions between them and within context of the overall domain (Katifori, Torou, Vassilakis, Lepouras, & Halatsis, 2008). Arp et al. (Arp et al., 2015) summarizes the types of relations as universal-universal (ex. this cat is an animal), particular-universal (ex. this cat is an instance of a cat), and particular-particular (ex. this cat is a continuant part of this cat grouping). After definition, ontologies are typically encoded within ontology creation software, exported into standardized formats such as RDF, OWL and OBO, and used within the tools of domain-specific tasks. Leading research (Saleemi, Rodríguez, Lilius, & Porres, 2011) describes that ontologies provide the flexibility, extensibility, generality, and expressiveness required for mapping domain knowledge into forms effective for computer-facing and human-facing use.

Next, users must comprehend how the VAT applied their input in its computational component so that they can effectively assess and guide their analytics process. That is, interface designs must help users engage with ML processes, as if users cannot understand ML characteristics and requirements of their tool, they cannot perform to the best of their abilities (Sacha et al., 2016). Additionally, users must be able to communicate their information-seeking objectives easily and accurately to a VAT's ML components. This is especially a concern in visual analytic tasks which involve direct interaction with ML processes (Hoerber, 2014; Holzinger, 2016; Mehta & Pandit, 2018; Tresp et al., 2016). Current research (Sacha et al., 2016) suggests that a balance must be struck between the computational technologies of a tool and users' perceptual and decision-making needs during analytic reasoning – that is, by supporting the “human-in-the-loop” aspects of the design appropriately. Therefore, a generalized and human-centered analytics process within a VAT involves a set of stages where:

1. Users specify their needs as a set of terms understood by the tool.
2. Users ask the tool to apply them into a search formulation.
3. The tool performs computations using the search formulation to produce a document set mapping.
4. The tool displays the document set mapping to users.
5. Users assess if they are satisfied, or if they would like to adjust their set of terms to generate an alternate mapping.
6. Users either restart their analytics process or complete the task.

Finally, as document sets within analytic reasoning tasks increase in size, it has become challenging for users to arrive at a final set of relevant documents without additional intervention. That is, even after computational components have produced a mapping which reduced the document set down to a subset of documents, these subsets are still too large to be of value to users. For this issue, information triage may be required to further reduce the number of documents into a usable size. During information triage, users' primary objective is to inspect, contextualize, and make timely relevance decisions on search results (Herceg et al., 2018). For this to occur, current research (Badi et al., 2006; Bae et al., 2010; Buchanan & Owen, 2008) describes that tools must allow users to encounter and perform rapid triage on large sets of documents in a non-linear fashion, while still being able to assess document relevance to information-seeking objectives. Notably, supporting information triage within tools can also help users assess the quality of their searching and triaging, and inform them on how to improve further information seeking (Loizides, Buchanan, & Mavri, 2016).

## **7.3 Materials and Methods**

### **7.3.1 Formative Assessment**

Users were asked to perform a controlled information search and triage task set for formative assessment. This set encompassed seven distinct tasks, combining to reflect the parts of a typical multi-staged information-seeking process. Users were assigned an interface to complete their tasks to their best of their abilities. During their performances, users were provided the same document set to be searched and triaged, and for ontology-supported interfaces, the same ontology file. No user performed a task set more than once, and all users completed all seven tasks. Users were asked about their thought processes as they received and performed their tasks and how they felt their interface helped or hindered their performances.

#### ***7.3.1.1 Users***

Within the evolutionary design model, users are central to the formative assessment of realized prototypes, where they can direct future formulations, realizations, and verifications within the design process. Periods of formative assessment were informally conducted using people associated with our research group, yet unaffiliated with the research. No user repeated a performance or used another interface at a later stage.

Users had no prior knowledge of the document set, and only possessed a general level of knowledge to the medical domain it describes. For instance, users understood and could communicate phenotypic abnormalities like a broken leg, light-headedness, or loss of vision, yet were not experts such that they could naturally use domain vocabulary without assistance. Users understood how to perform typical actions on an interface like clicking, typing, and saving, but were not provided information regarding the technical aspects of their interfaces.

#### ***7.3.1.2 Document Set and Ontology File***

A controlled information search and triage task set was used to guide users through their information search and triage. These tasks used the same document set, in addition to an ontology file for the ontology-support Stage 3 interface.

The health domain was selected as an equally unfamiliarity topic space to users, yet one which information search and triage is a common though challenging endeavor.

For all three of Stage 1, Stage 2, and Stage 3, a 10000-document subset of The National Library of Medicine's MEDLINE was used. MEDLINE was chosen because of its use within a wide scope of literature and active research. Each document within the chosen document set includes the document title, abstract, as well various metadata elements like authors, published date, and keywords ("PubMed," n.d.).

For Stage 3, the Human Phenotype Ontology (HPO) was selected for use as the ontology file. HPO maintains an exhaustive and expert defined domain coverage of over 11,000 phenotype terms, their relationships, as well as over 110,000 disease annotations (Kohler et al., 2014). HPO is a controlled and standardized vocabulary representing the human disease and phenotypic abnormality domain, and includes associated annotations in the domains of bioinformatics, biochemistry, and human genetics. For example, Blindness is an ontology entities within HPO which possesses a superclass of Visual Impairment, a subclass of Congenital Blindness, alongside various domain descriptions and annotations (Köhler & Robinson, 2016). Each HPO term is accompanied by attributes such as names, conceptual definitions, ontology indexing, term synonyms, class relationships, logical definitions, and domain expert commentary, to name a few (Köhler et al., 2018; "The Human Phenotype Ontology," 2020).

### ***7.3.1.3 Task Set***

Users were asked to complete seven tasks using their assigned interface. These tasks directed users to answer a multiple-choice question within an automated questionnaire. Users were tasked with exploring how their interface could help them relate their information-seeking objectives into a domain vocabulary and apply their information seeking objectives when query building. Next, users were tasked with assessing how their queries align with a document set, and how such alignments may impact search quality. Third, users were asked to perform high-level triage on search results, followed by opportunities to perform low-level triage. The task period culminated with the completion of a task involving all stages of information search and triage, guided by a research question. We provide a general description of each task within the task set (Table 1).

**Table 7-1 General description of tasks.**

| <b>T#</b> | <b>Target Stage</b> | <b>Task Description</b>   |
|-----------|---------------------|---|
| T1        | Query Building      | The first task asked users to consider two terms and contrast their rate of occurrence within the document set.   |
| T2        | Query Building      | The second task asked users to consider a term and determine its alignment with a set of provided definitions.  |
| T3        | Search              | The third task asked users to consider how provided set of terms aligned with the document set, both individually and in combinations.  |
| T4        | High-Level Triage   | The fourth task brought users to a specific document, and without allowing them to open the full document, were asked to predict its alignment to a provided set of terms.  |
| T5        | High-Level Triage   | The fifth task brought users to a specific pair of documents, and without opening them, were asked to compare and then predict which of them would contain a higher rate of occurrence of a specific term.  |
| T6        | Low-Level Triage    | The sixth task brought users to a specific document and were asked to count and order the rate of occurrences of a provided set of terms within that document.  |
| T7        | Multi-Staged        | The seventh task gave users a domain research question and asked them to produce five relevant documents from the document set. This task required users to progress through each stage of the information-seeking process, requiring the use of all available functionalities of their VAT interface. A topic background was provided for optional domain context. |

### **7.3.2 Shared Back-End Components**

All three interfaces of the evolutionary design share a ‘back-end’ encompassing two components: Analytics Server and Document Server.

Analytics Server is built using the Python-based Flask framework. It is accessed through an API which offers two functionalities: 1) uploading user-provided document sets and 2) performing ML computations. Document sets are uploaded from users’ computer file system to Analytics Server. This server validates the type, format, size, and encoding of uploaded documents, and then stores them into a temporary PostgreSQL database. This database is accessed by Document Server (a Solr server) during indexing procedures. Analytics Server can request Document Server to index all new documents. When users request ML computations during search, Analytics Server assesses the selected algorithm, the document set, and the search formulation generated by users during query building and search. It then performs sanitization and query expansion. During this process, query items within the search formulation are expanded using WordNet for synonym ring analysis, as well as with any ontology files, if provided. The search formulation in both its original and expanded form is packaged and applied within ML computations. The

resulting clusters are propagated back to the interface-level. We include pseudocode describing this process, using Stage 2/+ as an example (Figure 1).

```

Algorithm 1: CLUSTERING pseudocode between QUEST, Analytics Server, and Document Server


---


Input: A set Q of user inputted queries
Output: Signal to update interface with cluster assignments
1 targets ← chain(Q).unique().difference(getStopWords())
2 documents ← getDocuments()
   /* Prepare bag of words using target, related entities, and
   their generated WordNet synsets */
3 for i = 0 to targets.length do
4   target = targets[i]
5   targetCoverage ← target + target.getDirectlyRelatedEntities()
6   targetSpread[target] ←
     targetCoverage + targetCoverage.getWordNetSynsets()
7   targetSpread[target] ←
     targetSpread[target].unique().difference(getStopWords())
   /* Gather counts from pre-indexed documents, then fit and
   predict clusters using Scikit.Learn KMeans clustering */
8 documentCounts ←
   getIndexesFromSolrAPI(targetSpread, documents).scaleRange(0, 1)

9 reducedPCA ← SciKitLearn.PCA(nComponents =
   2).fit.transform(documentCounts)
10 kmeansPCA ← SciKitLearn.KMeans(init = 'k-
   means ++', nClusters = 7, nInit = 10)
11 clusterAssignments ← kmeansPCA.fit.predict(reducedPCA)
12 for i = 0 to targets.length do
13   target = targets[i]
14   for j = 0 to clusterAssignments.length do
15     cluster = clusterAssignments[j]
16     yPred = cluster.yPred[target]
     /* Generate weighting scale using x5 multiplier */
17     clusterAssignments[j].weighting[target] ←
       generateClusterWeighting(yPred)
18 return signalInterfaceUpdate(clusterAssignments)

```

**Figure 7-1 Pseudocode of clustering functionality for Stage 1, Stage 2, and Stage 3, Analytics Server, and Document Server. Step 5 is not performed for Stage 1 and Stage 2, which do not use ontologies during query expansion.**

Stage 1, Stage 2, and Stage 3’s Document Server is a cloud-based Solr server for indexing, storing, and serving documents from user-provided document sets. Solr is a prepackaged, scalable indexing solution developed by The Apache Software Foundation. It provides a valuable array of features like a REST-like API that support numerous HTTP-based communication interfaces. Solr also supports a wide range of customizable settings and schemas for storing, searching, filtering, analyzing, optimizing, and monitoring tasks (“Solr Cloud,” 2020). During indexing procedures, Document Server uses a prepared schema to extract new documents from a temporary PostgreSQL database hosted by Solr. These documents are then treated and stored within an index. Document Server also handles document serving requests. Document Server provides metadata such as titles, word counts, as well as content for document-level displays (See Apache’s official website and document for more information on Solr (“Solr Cloud,” 2020)).

## 7.4 Results

### 7.4.1 Stages of the Evolutionary Design

We present the results of the evolutionary design, spanning the formulization, realization, and validation of three VAT interfaces. For each interface, we summarize a topic analysis and if available, prior formative assessment, and its role in formulization. We then provide a functional description of the realized working prototype. Finally, we outline the findings and confirmatory evidence generated by formative assessment.

To assist recall, realized working prototypes were assigned stage identifiers:

- Stage 1: Traditionally designed search and triage interface.
- Stage 2: Progressively disclosed search and triage interface.
- Stage 3: Progressively disclosed, ontology-supported search and triage interface.

Each interface allows users to upload a document set, and if supported, an ontology file for vocabulary mediation. Then users may begin to query build and finalize a search formulation. Upon directing systems to perform search computations, a document set mapping is returned to users for further triaging. Stage 1, Stage 2, and Stage 3 are three web-based VATs which provide a generalized environment for search and triage through a plug-and-play support of user-supplied document sets, and for Stage 3, ontology files. Each interface is developed using HTML5, CSS, and JavaScript technologies, allowing for cross-platform, cross-browser support (ex., Firefox, Chrome, Opera). The D3.js JavaScript library is used to create visualization and interaction experiences (Bostock, 2016).

### 7.4.2 Stage 1

#### 7.4.2.1 Formulization

To initiate the evolutionary design process, we began by investigating the general design requirements of traditional search interfaces. A formative assessment of a traditionally designed interface would allow for a preliminary understanding of the issues facing users when searching and triaging large document sets, and thus was selected as the initial starting point in the evolutionary design. For this, we performed an in-depth topic analysis (Demelo & Sedig, 2021) on the types of search tasks and the traditional design strategies of generalized search interfaces. Collecting the findings of this analysis, we distilled a set of high-level criteria which could be used to guide the design of interfaces for search tasks involving large document sets (Table 2).

**Table 7-2 The criteria for guiding the design of generalized search interfaces for tasks involving large document sets.**

| <b>DC#</b> | <b>Design Criteria</b>   |
|------------|--|
| DC1        | Provide an information-centric interface that shows flexibility towards the evolving needs of users and the dynamic requirements of search tasks like the veracity of data sources and variety of information types. |
| DC2        | Provide interaction loops that supply prompt and effective feedback for users during the performance of search tasks.  |
| DC3        | Provide natural and consistent representations that allow users to understand the constraints, processes, and results provided by the interface.   |
| DC4        | Provide interactions that allow users to efficiently prepare, perform, assess, and adjust their machine learning to align with the information-seeking objectives of search tasks.                                   |
| DC5        | Provide mediation opportunities that assist users in communicating and bridge their information-seeking objectives into the vocabulary of the document set.  |

#### **7.4.2.2 Realization**

We describe the role of each criterion within the design of Stage 1 (Table 3).

**Table 7-3 The role of each criterion within the design of Stage 1.**

| <b>DC#</b> | <b>Stage 1</b>  |
|------------|---|
| DC1        | Stage 1 leverages powerful third-party computational technologies. Specifically, pre-built machine learning packages like SciKit-Learn and highly optimized indexing is provided by The Apache Software Foundation’s Solr product (“Solr Cloud,” 2020). Additionally, Stage 1’s interface provides users with clear text-based alerts, which reflect their current performance status.  |
| DC2        | Stage 1 supports an iterative analytics process for the performance repeated sets of search and triage. That is, within iterative interactions, users can save the results they regard as relevant in a persistent location within the tool, while still allowing further performances to occur.  |
| DC3        | Stage 1 supplies visualizations to help analyze and judge the relevance of search results.  |
| DC4        | Stage 1 utilizes modern visualization and computational technologies like D3.js to provide powerful interaction opportunities.  |
| DC5        | Stage 1 supports the use of a common vocabulary during query building within an unstructured input control. Since Stage 1 allows users to save relevant documents in a persistent list, users can repeat search and triage actions at a desired pace, without risking the lose of valued documents. This allows users to rapidly assess their vocabulary, assess the use of their vocabulary against search results, and adjust the search formulation as required. |

We now provide a functional description of Stage 1.

Upload is a component that supports the plug-and-play of user-supplied document sets (Figure 2a). When clicked, the upload button opens a file selection window. The window limits uploading to valid.zip compression format. When a compressed document set is uploaded, it is transferred to the back-end Analytics Server. Once at least one document set is uploaded, Search becomes active.

Search is a component that allows users to query build within a traditional search bar using unstructured input control (Figure 2b). Three interaction points are maintained: Query Input, Run button, and Clear button. Query Input is a search bar which accepts unstructured input from users. Users can remove previously typed input using typical text interactions like repeated backspacing actions from the keyboard, or by selecting the Clear button. After users provide input, the green “Run” button becomes active. If selected, this button will activate ML computations on the uploaded document set using the provided query input. Query terms are collected and sent to the Analytics Server. Result List component updates when the computations are completed.

Stage 1 supplies components to support information triage on the results of information search, also following traditional design practices. Result List is a component that provides users with a linear, multi-paged workflow of search results. Stage 1 uses the query created with Search within its Analytics Server to perform unsupervised K-Means clustering computations. Specifically, these computations move query input through various steps of sanitization, natural language processing, and ML to generate a mapping of the document set. Stage 1 then uses this mapping to calculate relevance weightings, sorts documents based on this weighting, broken up into 20 document pages. These documents are then retrieved from the Document Server, then displayed within a paged order accompanied by buttons to navigate between pages (Figure 2c). For each page, documents are placed within a listing that describe their assigned relevance weighting and document title. Words and phrases in the title which align with the query are highlighted. If users desire to examine a document in full before judging its relevance, they can select a document from the list.

Result Item is a component that presents to users an expanded display of a selected document, providing to them annotated document content (Figure 2d). Result Item allows users to rapidly assess the content of individual documents. When a user selects a document within Result List, Stage 1 will request the full document content of that result from Document Server. Query terms are then used by annotation services within Document Server to wrap HTML-based annotation tags into the document content, which is then returned to Result Item. The content of the selected document is represented in the following order: the file name of the document within the uploaded document set, the full document title, and a summarized version of the document content. The summarized version restricts the document to the passages of content that surround or have associations with the query terms provided in Search. Terms are highlighted through capitalization and with bolded font. In addition, Result Item collects all web links for quick access within a dropdown menu. For comparison, any number of documents can be opened within Result Item.

Each result within Result Item includes a green “pin” button, which saves documents for future reference. At any time, users may store a document in a global save list and to export for external use (Figure 2e). Stage 1 collects these saved documents within Saved List. Saved List can be accessed at the top right of Stage 1, directly to the right

of Search. Upon request, Saved List displays saved documents, which can be recalled, removed from the list, or copied for external use (not pictured).




Figure 7-2 The overall view of Stage 1 components: Upload, partial (a); Query Building (b); Result List (c); Result Item (d); Saved List, partial (e).

### 7.4.2.3 Validation

We asked users to perform a controlled information search and triage task set using Stage 1 (Table 1), then describe their experiences to inform future formulization efforts.

We provide a description of how users completed their task set using Stage 1 (Table 4).

**Table 7-4 How users completed the tasks using Stage 1.**

| <b>T#</b> | <b>Target Stage</b> | <b>Task Performance</b>  |
|-----------|---------------------|--|
| T1        | Query Building      | Users could use the provided terms within a query to generate a set of search results. From this set, they could scan the pages to access rates of occurrence. They could then be contrasted to generate an answer to which term had a higher rate of occurrence within the document set.  |
| T2        | Query Building      | Users could use the term, as well as the important terms within the definitions, to create search formulations. From these, they could inspect top documents to infer alignment between terms and their definitions.   |
| T3        | Search              | Users could use the provided terms within a query to generate search results. They could then scan pages to access rates of occurrence.  |
| T4        | High-Level Triage   | Users could take advantage of a relevance prediction and the annotated document title.   |
| T5        | High-Level Triage   | Users could take advantage of a relevance prediction and the annotated document within the paged set of search results.  |
| T6        | Low-Level Triage    | Users could open the document content, presented in summary form, to count terms and then use those counts to infer ordering.  |
| T7        | Multi-Staged        | Users could use the research question and any value they found in the background description to produce search formulations. They should then direct their interface to use those formulations on the document set, at which point they would be able to triage search results to encounter the most relevant documents to their task. From these results, users could triage the result pages and their documents, judging relevance on individual documents, and finalize their set of five documents. |

Formative assessment provided insight into how users approached the use of Stage 1 for their tasks. We itemize a summary of assessment findings its confirmatory evidence:

1. Users articulated a reduced easiness in their ability to perform their tasks, where they were required to establish indirect strategies involving manual and imprecise estimations of terms and documents, and tedious experiences within a linear inspection flow (e.g., A, B).
2. Users struggled during the performance their tasks, yet did not believe they struggled, accepting that was a necessary to search and triage. Yet after time to contemplate, they realized that they did not notice because they were familiar experiences from prior encounters with traditional interface designs. In particular, users observed that their implicit trust to how document results are rated, ordered, and presented within the overall structure of traditional interfaces, finding that they can unwillingly, or in some cases, willingly, avoid the critical step of confirming if that trust has been earned by their interface (e.g., C, D, E, F).

3. Users expressed they were restricted in their task performances by particular configuration details of the interface, such as the limitations of triaging large document sets within a paged layout, as well as a lacking use of color its effect on their ability to connect and process documents (e.g., F, G).
4. Users felt that the any opportunity to apply or take advantage of domain expertise would be helpful, yet they did not believe the interface provided enough to promote its use, nor help bridge between vocabularies (e.g., H, I).

(A) *"I realized looking at leukemia (in Stage 1), after I got like the third or fourth page, only like the first two (pages) even had 30% relevance. And then after like a few pages it went down to .25% and just a lot of these documents just (didn't) have information (when) I opened them up. I think a lot of things change when (you're looking at) such a large amount of the pool of documents, and since this tool really seems to focus on just finding: is this word in this document? It's going to be very challenging. I think if you have a much larger portion of (documents), it's going to be harder to find ... what you're trying to find. They're going to be buried inside the all the documents. "*

(B) *"I could look at some of those documents and try to determine whether or not they actually met that definition, but I found it arduous to both do that correlation and then try to delve into the document to understand the meaning."*

(C) *"I usually only go by the first page and having the relevant relevancy percentage beside the document. (These relevance rankings) kind of helped give me an idea of whether or not going any further would be a good idea. (Stage 1) highlighted the words, which I don't usually see that when I use other search engines. As somebody who is a visual person, having those bolded words help me (know) exactly what I'm looking for, and also the percentages beside (them)."*

(D) *"I thought it would be like pretty easy because I (am used to) answering ... open questions like ... find the things most relevant. So, this research question is like for me like that's just an easier thing to do because I have background in doing that kind of stuff. Uh, and (Stage 1) kind of functions like ... a library tool that is available. This kind of tool felt very familiar to me. I wouldn't say that I'm an expert when it comes to medical knowledge, but like I understand like basic terminology. So, like it's not something like the like what the terms meant or like what they refer to wasn't really like an issue like it wasn't really alienating. I have like some general level of confidence just using the terms and trusting the tool as you went along."*

(E) *"Looking at the Stage 1, I do not see any of the keywords that I typed in except for the very end, but they do not highlight any of the keywords that were searched, so I wouldn't even know if that that article is actually related without the relying on the number that's provided right beside the article that says that it's relevant. You would just have to trust that they're being honest with their assessment. I actually have no empirical proof of it myself."*

(F) *"I didn't consider it, because I assumed the relevancy (rankings) would show you the most relevant. There's no point in like going (down from the first ones) to the other ones. Because then it would be like almost impossible (see everything) ... that's not the front page. But then that would just take forever, right? If there were even more documents, that would be like I guess more difficult, because like you just have so much to choose from and ... finding the best articles that fit your research question like that would be more work, for sure. But like at the same time, it's like if these all are saying like there was like 1000 articles that had the same percentage of relevancy. So, like ... there's no point in assuming that ... the ones of the front are more relevant than the ones like that are two or three pages back."*

(G) *“Imagining the bolded words were colored instead, that would have greatly, greatly helped. Seeing color within the black and white text. I actually highlighted a lot of the tests and homework that I ever had to do. And anything that I had highlighted or drawn something beside in a colorful whatever color, I would remember a lot easier. And I’ve noticed (it’s) a lot easier than just black and white.”*

(H) *“At first my first mindset was just go by what was bolded, but that didn’t necessarily say everything. I found ... as I was looking at it more, I had to look for other words inside of the documents. That’s what was going through my head anyways. In the title (of) one, (a document) had something about ‘pediatric children’ ... but didn’t have children in it. Pediatric is kind of related to youth and children, right? I was able to take advantage of some external knowledge that I know in life that pediatric equaling children to maybe get a better answer.”*

(I) *“I think the knowledge and use that I got from (Stage 1) before. I knew that relevance would help with finding (documents) ... I can find relevance between certain terms. The first thing I searched for was ‘chromosomal instability’ and ‘tumor progression’. Just those four terms because they were going to be the major terms that I wanted to find. But then I also went through the background because that’s where there’s a lot of like terms and stuff that would be, you know, very significant in the knowledge, such as things like genome and like, karyotype. I don’t completely know what the words mean, but searching up and adding these terms I could find documents that were frequently showing up... I tried to figure out which documents are constantly showing up when I’m adding terms... The background really provides me guidance because I was unfamiliar with the research space... I think it was very efficient at finding documents that you know had a lot of these terms.”*

## **7.4.3 Stage 2**

### **7.4.3.1 Formulization**

Guided by Stage 1 formative assessment, the evolutionary design process continued with expanded investigations on the multi-staged information-seeking process and its design requirements. In particular, we sought to explore novel techniques which reduce the need for tedious search and triage, establish a stronger awareness of how users struggle within traditional design strategies, align configurations with the requirements of each stage of the information-seeking process, and learn of ways to promote the use of domain expertise. With these objectives, we performed an in-depth topic analysis (Chapter 6) on models of the information-seeking process, novel applications of progressive disclosure within complex multi-stages task interfaces, and the general requirements for supporting query building, search, high-level triage, and low-level triage within VAT interface design.

Portions of the analysis described the information-seeking process as a sequenced, multi-staged model with functional roles and human-centered requirements distinct to each stage (Hurdeman, 2017). Progressive disclosure is used to manage the visual space of an interface by occluding unnecessary elements of past and future stages, allowing users to satisfy the immediate requirements their task. For multi-staged tasks like those involved in the information-seeking process, progressive disclosure has been found to effectively support users to perceive and plan their task performances (Chuang, Ramage, Manning, & Heer, 2012). Users also benefit from the application of sensitivity encoding within human-facing interface design (Cortez & Embrechts, 2013). Sensitivity encoding provides

a visual preview of available actions for the user, and how those actions may affect future stages of multi-staged processes (Spence, 2014). If designed mindfully, sensitivity encoding can promote effective and timely query building and search formulation (Spence, 2002, 2004).

Collecting the findings of this analysis, we distilled an expanded set of high-level criteria which could be used to guide for designing VAT interfaces for searching and triaging large document sets (Table 5).

**Table 7-5 Design criteria for creating VAT interfaces for searching and triaging large document sets. Each DC# describes the design criteria and provides an integration classification.**

| <b>DC#</b> | <b>Design Criteria</b>  | <b>Integration</b> |
|------------|---|--------------------|
| DC1        | Use progressive disclosure when sequencing the stages of the information-seeking process.                                 | All Stages         |
| DC2        | Attune users to the characteristics and domain of the document set before beginning search formulation.                   | Query Building     |
| DC3        | Be cognizant of users' domain expertise.  | Query Building     |
| DC4        | Create search formulation and refinement environments supplemented by query building.                                     | Search             |
| DC5        | Leverage sensitivity encoding when previewing the document set mappings of search formulations.                           | Search             |
| DC6        | Present overview displays which arrange and compare document groupings using shared characteristics.                      | High-level Triage  |
| DC7        | Utilize non-linear inspection flows which support actions for traversing, previewing, contrasting, and judging relevance. | High-level Triage  |
| DC8        | Offer document-level displays which allow users to apply domain expertise during relevance decision making.               | Low-level Triage   |
| DC9        | Persist relevance decision making results externally to allow for repeat information-seeking sequences.                   | Low-level Triage   |
| DC10       | Allow users to encounter search results without a demand for immediate appraisal.   | All Triage         |
| DC11       | Promote positive feedback over negative feedback.   | All Stages         |

#### **7.4.2.2 Realization**

We describe the role of each criterion within the design of Stage 2 (Table 6).

**Table 7-6 The role of each criterion within the design of Stage 2.**

---

| <b>DC#</b> | <b>Stage 2</b>  |
|------------|---|
| DC1        | Stage 2 uses an accordion-like design which uses progressive disclosure best practices to sequence the multi-staged information-seeking process in the Query Building, Search, High-Level Triage, and Low-Level Triage components.  |
| DC2        | Stage 2's Query Building generates query items from all combinations of the inputted words, allowing users to attune to the characteristics and domain of the document set before beginning search formulation. This allows users to encounter terms and appraise how their information-seeking objectives may or may not align with the document set.  |
| DC3        | Stage 2's Query Building maintains an unstructured input control which allows users to describe their information-seeking objectives using their personal vocabulary, and any associated domain expertise that they might possess, rather than being restricted to a pre-set querying language.   |
| DC4        | Stage 2's Search allows users to control the search formulation, encounter previews of its use on a subset of the document set, then initialize ML computations on the full document set. From this preview, users can assess if they are satisfied with an existing search formulation, or if adjustments are required.  |
| DC5        | Stage 2's Search provides users a sensitivity-encoded matrix-like display which previews a cluster analysis of document groupings within the document set. By adding and removing query items from the search formulation, users can investigate how individual query items align with the document set, how differing query item combinations change document grouping arrangements, as well as estimate how many documents may be found if a full search were to be performed.  |
| DC6        | Stage 2's High-Level Triage displays a full document mapping within Query Result Heatmap, providing users with a high-level abstraction of the document set which arrange and compare document groupings using shared characteristics.  |
| DC7        | Stage 2's High-Level Triage empowers users to inspect, re-order, and open each document grouping to assess its size and alignment with the search formulation. When opened, an additional Query Result Heatmap provides high-level abstractions of individual documents. Together, these abstractions allow for non-linear inspection flows which support actions for traversing, previewing, contrasting, and judging relevance.   |
| DC8        | Stage 2's Low-Level Triage supplies users a timeline-like visual abstraction of individual documents which reflect the position of words or phrases which align with the query items of the search formulation. With this, users can use their vocabulary and domain expertise to assess the presence of query items used within the search formulation, where in the document they are, and the density of their usage. Furthermore, users can access full document-level displays which present all available document content. |
| DC9        | Stage 2's Low-Level Triage permits users to add relevant documents to a persistently saved list, allowing users to continue searching and triaging without the risk of losing progress.   |
| DC10       | Stage 2 implements a multi-modal interface of separate visual spaces for the stages of information-seeking process. Each visual space provides users the freedom to manipulate, re-compute, and display   |

---

---


the contents of its stages without affecting the progress of alternate stages. Furthermore, there is no externalized demand for immediate appraisal, decision making, or other requirements for action unless user-directed.

- DC11 Stage 2 provides users a progressive series of positive feedback opportunities. Specifically, these points of positive feedback span each component reflecting the complete information-seeking process. In them, the results of current stage components are promoted as the initial configurations of upcoming stage components, starting with initial query building, and resulting in a final collection of relevant documents from the document set.
- 

We now provide a functional description of Stage 2.


Stage 2 maintains a progressively disclosed accordion-like design which presents a sequenced set of modal visual components. These components allow users to transition between the stages of their information-seeking process. Namely, the stages of query building and search during information search, and the stages of high-level and low-level triage during information triage. Only one component of Stage 2 is active at a time. When a component is active, it becomes the primary respondent to interaction and is assigned a majority portion of the visual space. Since Stage 2 takes advantage of progressive disclosure within its design, the components of non-active stages are not hidden, but instead are provided a portion of the remaining visual space scaled by its proximity to the active component. For example, when the first Query Building component is active, it is provided the majority of the visual space. Yet the next component in the sequence, Search, still reflects information which would be beneficial to actions in Query Building.

Query Building is the first component within Stage 2 (Figure 3). In this component, two functions are performed: uploading user-provided document sets and query building. Upon clicking the upload button, users can select a document set which is then inserted into a file management listing, accompanied by file name, type, and any available descriptions. Once a document set has been uploaded, users can begin query building by inputting text into a search bar. This leads to the generation of query items from all combinations of the inputted words. Each query item is accompanied by a predicted presence within the document set. Encountering these terms, users can learn more about the vocabulary used within the document set, appraise how their research problem may or may not align, and adjust their query item selections. Query items can be saved and are assigned unique colors that are used throughout all components.




**Figure 7-3 An overview of Query Building within Stage 2. Users can see that a document set has been uploaded, that a set of query items have been generated for inspection, one having been added in Search.**

Search is the second component within Stage 2 (Figure 4). In this component, users can control the formulation of search queries, encounter sensitivity-encoded previews of the formulation, and initialize search on the full document set. In Search, a list allows users to manage query items, including insertion into the search formulation. After at least one query item has been selected, a preview of the current search formulation is activated. This preview is a matrix-like display describing a cluster analysis of document groupings within the document set. By adding and removing query items from the search formulation, users can investigate: 1) how individual query items align with the document set, 2) how differing query item combinations change document grouping arrangements, as well as 3) estimate how many documents may be found in a full search. If not satisfied, users can refine their formulation using existing query items or generate new query items within Query Building. Finally, users can initialize a full search, with the results of ML computations sent to triage components.



**Figure 7-4** An overview of Search within Stage 2. Users can see that Query Building has generated a set of query items, which have been injected into the preview search environment.

High-Level Triage is the third component within Stage 2 (Figure 5). In this component, users triage the results of ML computations at the grouping level. A full document mapping is displayed within Query Result Heatmap, providing users with a high-level abstraction of the document set (Demelo, Sedig, & Parsons, 2017). The abstraction is divided into horizontal slices representing document groupings. For each document grouping, a set of color cues highlight query item presence. Users can inspect each document grouping to assess its size and alignment with the search formulation. Listings can be re-ordered to prioritize specific query items. A cursor marks the current position, trailing dots mark previously viewed listings, and a green mark for those selected for further triaging. Document groupings can be opened within an additional Query Result Heatmap, which provides high-level abstractions of individual documents from selected groupings. Users can use this additional collection of documents to individually assess alignment with the search formulation, as well as inspect metadata such as titles and document-specific counts. Documents can then be saved to Low-Level Triage component.



**Figure 7-5** An overview of High-Level Triage within Stage 2. Users can see that Search has generated a set of document groupings reflecting the document set. A subset of those groupings has been selected for deeper investigation, producing further listings of documents from the set. These documents have been inspected and have been determined to be valuable enough to be added into Low-Level Triage.

Low-Level Triage is the fourth component within Stage 2 (Figure 6). In this component, users triage the documents produced in High-Level Triage. Users are provided a timeline-like visual abstraction of individual documents. In this visual abstraction, colored marks are placed along a timeline to reflect the position of words or phrases which align with the query items of the search formulation. Documents with strong alignment will produce numerous markings, resulting in color-heavy and densely annotated timelines. Utilizing the timelines, users can perform rapid analysis of the thematic themes of individual documents. Namely, users can assess the presence of query items used within the search formulation, where in the document they are, and the density of their usage. The Document Content viewer is activated after selecting a document for deeper inspection. This viewer will present all available document content, such as document text, title, authors, file name, URLs, and published date. Users can toggle between a full document and a summarized mode. The full document mode displays all available information, annotated to reflect alignment with the search formulation. The summarized mode condenses documents to just content in proximity to aligning words or phrases. Users may open documents to inspect, compare, and make final relevance decisions on its content. Relevant documents can be added to a persistently saved list, allowing users to continue searching and triaging without the risk of losing progress.

**Low-Level Triage**

4a. Selected Documents

Remove All

4b. Document Content

Add to Saved List View Full Document

Summarized Document annotated with query terms 0 Document Links

File Name: PUBMED2005DEMO.zip/16375848

Title: Acute effect of fatty acids on metabolism and mitochondrial coupling in skeletal muscle.

Acute effect of fatty acids on metabolism and mitochondrial coupling in **SKELETAL MUSCLE**...

... Acute effects of free fatty acids (FFA) were investigated on: (1) glucose oxidation, and UCP-2 and -3 mRNA and protein levels in 1 h incubated rat soleus and extensor digitorum longus (EDL) muscles, (2) mitochondrial membrane potential in cultured **SKELETAL MUSCLE** cells, (3) respiratory activity and transmembrane electrical potential in mitochondria isolated from rat **SKELETAL MUSCLE**, and (4) oxygen consumption by anesthetized rats. Long-chain FFA increased both basal and insulin-stimulated glucose oxidation in incubated rat soleus and EDL muscles and reduced mitochondrial membrane potential in C2C12 myotubes and rat **SKELETAL MUSCLE** cells. Caprylic, palmitic, oleic, and linoleic acid increased O<sub>2</sub> consumption and decreased electrical membrane potential in isolated mitochondria from rat **SKELETAL** muscles. FFA did not alter UCP-2 and -3 mRNA and protein levels in rat soleus and EDL muscles. Palmitic acid increased oxygen consumption by anesthetized rats. These results suggest that long-chain FFA acutely lead to mitochondrial uncoupling in **SKELETAL MUSCLE**...

... Mitochondria, **MUSCLE**...

... **MUSCLE** Cells...

... **MUSCLE** Fibers, **SKELETAL**...

... **MUSCLE**, **SKELETAL**

**Figure 7-6** An overview of Low-Level Triage within Stage 2. Users can see High-Level Triage has generated a subset of the document set for deeper inspection. Each of these documents have been provided a timeline-like summary reflecting the presence of query items within its content. If selected, a document-level viewer depicts all available document content, either in summarized or full form.

### 7.4.2.3 Validation

We asked users to perform a controlled information search and triage task set using Stage 2 (Table 1), then describe their experiences to inform future formulation efforts.

We provide a description of how users completed their task set using Stage 2 (Table 7).

**Table 7-7 How users completed the tasks using Stage 2.**

| <b>T#</b> | <b>Target Stage</b> | <b>Task Performance</b>   |
|-----------|---------------------|---|
| T1        | Query Building      | Users could use the provided terms to create potential query items, which would then generate counts of that term within the document set. If they required further confirmation, users could use those query items to generate a search prediction matrix, or even perform a full search.  |
| T2        | Query Building      | Users could use the term, as well as the important terms within the definitions, to create search formulations. From these, they could inspect top documents to infer alignment between terms and their definitions.  |
| T3        | Search              | Users could use the terms to create query items which could be used to generate a search prediction matrix. Using this prediction matrix, users could compare and contrast various term combinations, or they could perform a full search which could allow them to inspect document groupings.   |
| T4        | High-Level Triage   | Users could take advantage of the annotated document title, a document-level count matrix, and a document timeline.   |
| T5        | High-Level Triage   | Users could take advantage of the annotated document title, a document-level count matrix, and a document timeline, within the saved document collection.   |
| T6        | Low-Level Triage    | Users could take advantage of the document-level count matrix within High-Level Triage or the document timeline within Low-Level Triage. They could also open the full document content to manually count term occurrences.   |
| T7        | Multi-Staged        | Users could use the research question and background context to generate potential query items. With their available resources, users could then compare and contrast the relevance of query items to the document set. Then, they could insert query items into a search formulation, supported by prediction matrixes. If satisfied with their search formulation, users could direct ML computations on the full document set to generate document groupings. Next, users could perform high-level triage on these document groupings, extracting from them individual documents for deeper inspection. These documents could then be inspected, allowing users to finalize their set of five documents. |

Formative assessment provided insight into how users approached the use of Stage 2 for their tasks. We itemize a summary of assessment findings its confirmatory evidence:

1. Users expressed initial lack of confidence when using the novel elements within the interface, yet voicing appreciation for them after appropriate time to attune (e.g., A, B)
2. Users felt that they were able to mold their interface to match their personal needs and expectations when searching and triaging, in contrast to their typical limiting experiences with traditional interfaces (e.g., C, D)
3. Users believed that Stage 2 allowed for a heightened level of interaction and transparency with the “search engine”, which allowed them to build a more trusting relationship with their tool, in contrast over the more “black box” approach commonly used when combining ML with traditional interfaces (e.g., E, F).

4. Users indicated they their task performances were promoted by the configuration details of the interface, such as the automated term counts, predictive search formulation displays, compartmentation of document groupings, and document preview displays (e.g., G, H, I, J, K, L, M).
5. Users felt that the any opportunity to apply or take advantage of domain expertise would be helpful, yet they did not believe the interface provided enough to promote its use, nor help bridge between vocabularies (e.g., N, O).

(A) *“When starting the initial questions, I felt like I wasn't totally confident with my usage of the tool, especially during question two when I was facing so much difficulty. I was trying to really understand, I guess what the tool was allowing me to do and the how to interface with it. But by the time I got to task seven I was pretty confident that I understood how to use the tool.”*

(B) *“Once I understood what it was showing me, it helped me. Usually with new tools I tend to read through the documentation or watch videos. And then it still takes me like a while to pick up on them. Like, just running through them and using them a few times. Once you get the hang of it, usually you find success in whatever it's providing you.”*

(C) *“I think it in a strange way though, that (Stage 2), what I was using, actually feels like it would be more helpful (than traditional interfaces). It would slow things down. You know, I might not always get the answer I'm going to want, but I think in some ways it's going to be faster than you having to tab through the pages, previewing the articles and asking: is that is that relevant or not?”*

(D) *“One thing that I always do when I'm searching for different things, even using you know regular search engines like Google or whatever, is that I try to be as comprehensive as possible. I'm striving to avoid failure as much as possible. I tried to be as comprehensive as possible in both doing Task 7 and other tasks.”*

(E) *“A lot of other tools you use, they kind of do predictive searches for you. They build a filter. So, for example, like Google doing a predictive search. It's predicting based on top results from previous searches, so with that, it can be finagled with. Where you know you could have a bot farm or whatever finagling those search results and making them be what they want them to be. Whereas with (Stage 2) you get to parse those results and make your own educated decisions versus (the search engine) doing it for you. So, I feel more control in the experience then you would typically. I feel more certain in the end goal.”*

(F) *“I think (Stage 2) is a benefit. In my previous experiences with searching queries, if you just type it in and it blurts out the answer it prioritizes in whatever way that it wanted. I like this because it is a little bit more specific, and you are able to choose more of the options that you want to use. I think you have more control of how to exactly to find the answer and what exactly you are looking for, instead of starting from just a general basis of all the answers that are possible. So, I think it's better to be able to narrow down exactly what you're looking for and find more appropriate answer towards your question.”*

(G) *“Being able to see (in Stage 2) how many hits there were for each word grouped together... like seeing like the four words together and seeing that there was like two or three hits for a document and being able to sort that way -- it was pretty helpful.”*

(H) *“Only having the title would have probably hammered me. I would have been without those counts; the counts were crucial in this task.”*

(I) *“In my (Stage 2), I was able to type in cancer I saw a number in the seven hundreds, and when I typed in leukemia it was a number in the four hundreds. So, I was able to see a clear numerically different value between the two, whereas in Tool A there is no such value indicated.”*

(J) *“First, I did a query based on some keyword, and based on the query I found I went through and I removed some keywords that I found misleading for the query. Then I'll (went) to the search, (and found) the query result and the prediction. So, I found (it and thought): OK, there is a good possibility of these terms together.”*

(K) *“It was also helpful that you could see all those document (groupings) at the same time ... colored so that it would help you instead of having to manually count.”*

(L) *“I did find it very helpful to be able to go through and see the boxes. Whenever there's a dark one -- that that seems interesting. And then the ones that have all of them, ... you can see that it's already a fairly small cluster, so that that made sense.”*

(M) *“It really helped... that you were able to compartmentalize the different document clusters. Not have to deal with all of them at the same time... I think that was the part that really helps. To focus on the diagram, the clusters that contain all of the words together.”*

(N) *“For something like this where this was more involved towards the knowledge of medicine, I had no idea between the different options of my answers for the question. So, I was looking for some kind of relevance to be there in the 1st place without diving further into the subjects within the documents. I was trying to look for a combination of the two searches between one of the options of answers and the original question to see if there's actual relevance in any kind of way between the terms. I believe I was just lucky at with this question is specific because when I typed in the possible answers that were given only one of them had any connection options. The other one did not come up with being related to one another in a further search. So, in essence I was being guided by the document set.”*

(O) *“If the tool could provide me with a translator or with a dictionary of some sort, I could look up these words that I don't know, if they're related to the query or not, and by looking those up and you know, getting the information for the meaning of those, I could manually check to see if these can relate. I would have searched for the query or one of the terms and then go all the way to the end. Then I would see the actual documents and read some of the summaries, titles, and any portion of the text to see if the if the two concepts or if the term and the query are related to each other, or not. Though, given that I have to go through the clusters, and I have to you know review all of them and all of these -- it would be laborious task.”*

### **7.4.3 Stage 3**

#### **7.4.3.1 Formulation**

Continuing investigations concentrated on addressing the remaining challenges highlighted by prior formative assessment, and in particular, how existing design criteria (Table 2 DC5, Table 5: DC2, DC3) could be better promoted by ontology-supported VAT interfaces. Namely, we explored challenges of misaligned vocabularies during search tasks, both between the common vocabularies of users and the domain vocabularies of tasks, as well as ways to take advantage of domain expertise when searching and triaging.

Therefore, we performed an in-depth topic analysis (Demelo & Sedig, 2021) to establish greater understanding on the user-facing activation of ontologies for searching and triaging large document sets, and their benefit to helping address design requirements within the multi-staged information-seeking process. When communicating across vocabularies, users may struggle to describe the requirements of their search task in a way that is understandable by their VATs [6,7]. To deal with this challenge, ontologies can be a valuable mediating resource in the design of user-facing interfaces [8]. That is, ontologies can bridge the vocabularies of users with the vocabulary of their task and its tools. For this, query expansion strategies were explored, both for computational- and user-facing activation. Query expansion interface strategies provide users mediation opportunities which bridge their vocabulary to the vocabulary of the document set [54]. These mediation opportunities come during query building and search within the information-seeking process, where users can be presented mediating opportunities which suggest to users how their common vocabulary could align with the vocabulary of the domain. Query expansion interface strategies benefit users by allowing them to continue to use common vocabulary during the process of query building. This gives users higher confidence about what the interface is asking of them, and what they are telling the interface to do. Furthermore, by integrating the use of mediating resources like ontologies, designers can demonstrate to users the quality of their query building and how their vocabulary decisions affect the performance of their search tasks [58]. For the query expansion interface strategy to be successful, designers must clearly communicate to users how exactly their query building has affected their search. If this communication is not provided, it can leave users confused regarding how their decisions have affected their search and can make it challenging for them to assess task performance.

### 7.4.3.2 Realization

We describe how the configuration of Stage 2 adjusted into that of Stage 3, and its alignment to unsatisfied criterion (Table 2 DC5, Table 5: DC2, DC3) of prior formulization efforts (Table 8).

**Table 7-8 The configuration of Stage 3, and its alignment to unsatisfied criterion of prior formulization efforts.**

| <b>Table</b> | <b>DC#</b> | <b>Stage 3</b>  |
|--------------|------------|---|
| 2            | DC5        | Stage 3 expands on Stage 2 by allowing the uploading of user-provided ontology files which are processed and used to provide mediation opportunities within Query Building that assist users in communicating and bridge their information-seeking objectives into the vocabulary of the document set.  |
| 5            | DC2        | Stage 3 expands on Stage 2 by not just generating query items from all combinations of the inputted words, but also attaching mediation opportunities to each query item aligning to the content of user-provided ontologies files. This allows users to attune to the characteristics and domain of the document set before beginning search formulation, both from their own vocabulary as well as from guidance from expert-defined ontologies. This allows users to encounter terms from their input and from the ontology to help them appraise how their information-seeking objectives may or may not align with the document set. |


- 5 DC3 Stage 3 expands on Stage 2 by not only allowing users to describe their information-seeking objectives using their vocabulary, but also through the activation of a provided vocabulary described within their ontology files.
- 

We now provide a functional description of points of adjustment from Stage 2 within the configuration of Stage 3. Functionality remains in line with Stage 2 where no description is provided.

Query Building is the first component within Stage 3 (Figure 7). In this component, three functions are performed: uploading user-provided document sets, uploading user-provided ontology files, and query building. Upon clicking the upload button, users can select a document set or ontology file, which are then inserted into a file management listing, accompanied by file name, type, and any available descriptions. Users can then begin query building by inputting text into a search bar, using the same design considerations as Stage 2.

If an ontology file has been uploaded, then new mediation opportunities become available. Stage 3 analyzes the alignment of query items with the entities, relations, and descriptions of user-provided ontology files. If a feature of the ontology is found to align with a query item, it is placed within a drop-down menu attached to the listing (DC3). In this menu, users are presented with ontology terms which are conceptually similar to that query item. Encountering these terms, users can learn more about the domain vocabulary, appraise how their research problem may or may not align with their document set, and adjust their query item selections (DC2).

Query items both from user input as well as from ontology mediation can be saved and are assigned unique colors that are used throughout all components.



**Figure 7-7** An overview of Query Building within Stage 3. Users can see that an ontology and document set has been uploaded, that a set of query items have been generated for inspection, and a subset of those have been added in Search.

Search is the second component within Stage 3, sharing the same design considerations as Stage 2. Within Stage 3, ontology terms generated by mediation in Query Building will be displayed alongside user-defined query items in Search. Similarly, they can be selected for use within search formulation, and used within a final formulation for a full search of the document set. In addition, Stage 3 also adds ontology-provided query expansion opportunities alongside those provided by WordNet when making ML computation requests to “back-end” components.

High-Level Triage is the third component within Stage 3, sharing the same design considerations as Stage 2. Within Stage 3, ontology terms can be used as query items in the search formulation, and thus can in turn be used within document grouping computations and visualizations. Furthermore, functions which support term counts within document groups and individual documents use ontologies to expand their tabulations, just as the “back end” components used ontologies in query expansion.

Low-Level Triage is the fourth component within Stage 3, also sharing the same design considerations as Stage 2. Similar to High-Level Triage, ontology files are used to expand the tabulations of term counts for individual documents, including those that produce timeline-like visual abstractions of documents. Furthermore, this expansion is also used within the summation and annotation services of the Document Content viewer. When users toggle a summation of a document, words and phrases associated with ontology content are also reflected, helping users rapidly assess the relevance of a document to their information-seeking objectives.

### 7.4.3.3 Validation

We asked users to perform a controlled information search and triage task set using Stage 3 (Table 1), then describe their experiences.

We provide a description of how users completed their task set using Stage 3 (Table 9).

**Table 7-9 How users completed the tasks using Stage 3.**

| <b>T#</b> | <b>Target Stage</b> | <b>Task Performance</b>  |
|-----------|---------------------|--|
| T1        | Query Building      | Users could use the provided terms to create potential query items, which would then generate counts of that term within the document set. If they required further confirmation, users could use those query items to generate a search prediction matrix, or even perform a full search. If desired, users could take advantage of ontology mediation to bridge their vocabulary with that of the task space.  |
| T2        | Query Building      | Users could use the term, as well as the important terms within the definitions, to create search formulations. From these, they could inspect top documents to infer alignment between terms and their definitions. If desired, users could take advantage of ontology mediation to bridge their vocabulary with that of the task space.  |
| T3        | Search              | Users could use the terms to create query items which could be used to generate a search prediction matrix. Using this prediction matrix, users could compare and contrast various term combinations, or they could perform a full search which could allow them to inspect document groupings. If desired, users could take advantage of ontology mediation to bridge their vocabulary with that of the task space.   |
| T4        | High-Level Triage   | Users could take advantage of the annotated document title, a document-level count matrix, and a document timeline.  |
| T5        | High-Level Triage   | Users could take advantage of the annotated document title, a document-level count matrix, and a document timeline, within the saved document collection.  |
| T6        | Low-Level Triage    | Users could take advantage of the document-level count matrix within High-Level Triage or the document timeline within Low-Level Triage. They could also open the full document content to manually count term occurrences.  |
| T7        | Multi-Staged        | Users could use the research question and background context to generate potential query items. Additionally, ontology mediation was available to users. With their available resources, users could then compare and contrast the relevance of query items to the document set. Then, they could insert query items into a search formulation, supported by prediction matrixes. If satisfied with their search formulation, users could direct ML computations on the full document set to generate document groupings. Next, users could perform high-level triage on these document groupings, extracting from them individual documents for deeper inspection. These documents could then be inspected, allowing users to finalize their set of five documents. |

Formative assessments provided insight into how users approached the use of Stage 3 for their tasks. We itemize a summary of assessment findings its confirmatory evidence:

1. Users continued to express initial lack of confidence when using the novel elements within the interface, yet voicing appreciation for them after appropriate time to attune (e.g., A, B, C)
2. Users continued to feel that they were able to mold their interface to match their personal needs and expectations when searching and triaging, in contrast to their typical limiting experiences with traditional interfaces (e.g., D, E, F, G)
3. Users are capable of utilizing ontology files to align their vocabulary with the vocabulary of the domain, even if they are not initially familiar with the ontology's domain or its structure and content (e.g., H, I, J).
4. Users conveyed that the use of ontologies and the mediation opportunities they provided help provide guidance during their information-seeking process, making tasks manageable and easier, and not having them would negatively affect their task performance (e.g., K, L, M, N, O).

(A) *"It's a tool that I'm not used to and I'm kind of going back and forth and for me when there's kind of a lot of little moving parts. I mean it, it feels that way right now because I'm not familiar with the tool and I'm kind of taking in this information and figure out how the way different pieces of information needs to go together."*

(B) *"(A traditional tool) would be an easier to jump in with, but potentially could be more complicated to actually arrive at the end of the task with the correct understanding. Yeah, so if I just did these tasks (with a traditional tool) I can immediately it matches my mental models of how I use search interfaces. I type things in, I click run, I go through pages of results. But I get less control for all the important things that I was using during completing the tasks like selecting phrases of words and whatnot. It's a mixture. For simplicity sake you get things quicker, so (a traditional tool) is good because I can just go ahead and use it. But I'm not going to be able to get as much from it than (Stage 3); the one that I used."*

(C) *"I think at first I was a little stressed because I wasn't sure about the background that I have in this field so I thought maybe I can find the documents that are relevant to the question. But as I went through (Stage 3) and worked with the searching bar and I chose (the query) and I saw that there are only few clusters that contain the words. Then, I was confident that I can choose some clusters that contain the word(s), and then as I went through the task, it was much easier for me and I felt more confident. I can find the answer. I can choose a set of relevant documents."*

(D) *"I think how (a traditional tool is designed), this would have definitely taken longer... I would look at the first page and that's what I'm going to use to determine my results. The other way that (Stage 3 has) it, with the clustering and the progressive triaging, it killed the first page bias. So, I think that's a real advantage."*

(E) *"I found that Task 6 was (easy enough) ... (but) I think if I (had) taken the time to add families into the query and gotten the exact document back, that would have made that task much easier... I was going through and counting by hand and the highlight helps. Definitely once you have it in the query, it's very much easier to kind of get a sense of the proportion. So ... I was counting 'families' by hand, but ... really, had I added family to the query, I think that (any remaining challenge) would have been mitigated."*

(F) *"If I was using (Stage 3) against other search engines, I would just use this constantly. Being able to really filter down exactly what I need... like that's really on point. Especially with like the different colorings of the words. It's like telling you like what each*

document (cluster) is like. It gives me the ability to better align with the documents and gives me more confidence when I'm creating my queries. I'm making the correct query decision even before even running the search.”

(G) “I still really liked the being able to see where words were being used in the document. I think that was an important part and I think as somebody who actually does ... (search) for documents in databases all the time. That's something that is (valued). Sometimes (I use) control F, searching the document for phrases to see if it's relevant to me or not, and this was giving me that really quickly and effectively. So I think ... for the type of deep exploration of a data set like this, I think the effort you have to put in to figure out what you're doing and what it's telling you is useful in the end, because you're sifting through it a lot easier.”

(H) “I knew that when I searched the term ‘nevus’, the option that we can look for the terms and their definitions here was something that helped me for finding the term. I could easily search a term and look for the for its definition here. In the process I found that very helpful for something like this, where you did to find information fast because I don't have any background knowledge on this science. I think it's really necessary for someone who doesn't know about one term to have this information. It was really beneficial for me.”

(I) “For Task 2, since these terms were not so much familiar to me, I started off using the synonyms of the query that I was given. So, I was given an abnormality of the liver. Obviously, I wouldn't search for “an abnormality of the liver”. I started searching for the term “liver abnormality” and then I came across in the results different terms. I couldn't find the terms that were provided in the answers, so I started ... to see if there is any entry with one of these options in the ontology, and if so, (if) I could understand the meaning of this term and then see if it is related to the abnormality of the liver or not.”

(J) “First I did a query based on some keyword. And based on the query I found I went through (the) ontology and I removed some keywords that I found misleading, for the query and then I'll go to the search, found the query result and (used) the prediction... (where I thought), OK, there is a good possibility of these terms together. And then I went to triage level and based on two to two or three cluster I was able to find the document that that I was looking for.”

(K) “I think (the ontology) was really helpful ... because I didn't know what a lot of those terms even meant, so I was like pretty ‘fresh’. So, when I was looking at these then I clicked in and saw the ontologies. It was really useful because ... it gave me a sense of what we were talking about, right? This is of course useful to have those extra terms.”

(L) “Oh, oh, that one task with the birthmark, I would not have gotten anywhere. I probably would've put “I can't answer” as my answers came from the ontology, expanding ‘nevus’. So that's one thing right there that that task was difficult but doable with domain knowledge. If I didn't have the ontology, I would have said I can't do this. I probably would have tried to pull up the document at the end and had to review them to see what is the actual what documents have this in the context and then (find): what is nevus? So, if you had documents, eventually could review them to answer. But... I would've spent longer doing it without the ontology, for sure.”

(M) “If I didn't have the ontology, I think I would look for the documents that have the most occurrence of that term and then I think I could do is was to going through the documents and read the abstracts and see what I learn about I mean what other words that I see in the other documents. Then, I could come back to this query building query generator and then add those terms to my query and then see if there are more documents that contain both those words. Then I could go again and see it look for more terms that are relevant to these terms and add them to find some relevant documents.”

(N) *“When approaching Task 7, my thoughts went: What’s my domain knowledge, what word we’re in the query and then yeah, if I can’t remember a word like let’s see if it shows up in the suggested ones and see if I can jog my memory.”*

(O) *“I was thinking towards (the end) where (the ontology) would have been helpful. So ... it would have possibly brought up some of those other terms just from searching a few words and they would be able to make some connections between the text that was provided and some of my search terms (to see) ... how relevant they were. So, if I was shooting in the dark and hoping for the best, which is what I was kind of doing (without the ontology), at the very least, it would have given you confidence of your actions. Yeah, I think so. A little bit more confidence.”*

## **7.5 Discussion and Conclusions**

In this section, we provide discussion of user responses to the evolutionary design. We begin with Section 5.1., which provides a summation of the overall findings. This is followed by Section 5.2., where we present a discussion of the separate stages of information search and triage, and the benefits of interfaces which distribute their individual needs within a progressively disclosed transitional structure. In Section 5.3., we discuss ontology integration, and how users can find benefit from ontology mediation when working in unfamiliar task domains. We finish with Sections 5.4. and 5.5. where we describe the limitations of this research, as well as present considerations for future interfaces and objectives that this research can inspire.

### **7.5.1 Overall**

- Overall, users found benefit from the use of the novel characteristics of the evolved interfaces. They described improved overall easiness, highlighted how they benefited from interacting with their tools, and recounted where their mindsets were as they were tasked to search and triage (e.g., A, B, C, D).
- Users noted that searching and triaging experiences can vary depending on factors like their familiarity with the document set, how the size of the document set impacts the type of interface they would most value, and how there is a benefit to interfaces which can dynamically adjust based on your personal needs when searching and triaging (e.g., E, F, G).

(A) *“I was ... using (Stage 2) to try to give myself the best chance -- using probability and searching for the terms, finding what got the most hits and giving it a read over and seeing if it looked like it contains something that could match up with the correct answer... using the tool to try to give myself the best chance. So, I’d search for the terms and then find ones that had a match for all of them. And then go through the document and make sure that the highlighted ones look like they all kind of work together and made sense as far as I could understand.”*

(B) *“I do really like how (Stage 3) gives you the ability to find different vocabulary or other words that may not have been the first thing you thought of when you were building the query. “*

(C) *“I would say that ... I used (the) high-level triage section mostly. To find the documents I used all of the parts, but the main focus of my task I think was on the high-level triage... I chose the relevant words and then I had to go through the documents that have those words, and I think that's important.”*

(D) *“... I'm coming at this from a perspective of, let's say, using Google to do most of my searches. I don't know the machinery that Google uses. I just know as a user, if I were to type in: Damage to the optic nerve? ... I'm pretty sure it's going to give me the answer I'm looking for... I just do something that's fairly intuitive and most of the time I get the answer I'm looking for. So, when I'm thinking about (an) interface, I'm just thinking: what it is doing based on the intuitive input I'm giving it?”*

(E) *“I would find more success from a smaller document set in a more traditional tool. But if there were more documents, probably it seems like (Stage 2) would do better with lots of them.”*

(F) *“(With Stage 3), going through low-level triage and seeing and reading the abstract and the actual documents... I wanted to check and have a comparison between the documents that I chose... (because) usually my style of choosing (is that) I usually choose more than what I have to choose and then I remove that and the extra ones.”*

(G) *“At first I just copy and pasted the entire keywords like the chromosomal instability drive tumor progression, and my first thought I wasn't really seeing the results of the all the keywords mixed in together the same way. So, (with Stage 2) I was able to go back and see if the words individually not together had brought in any difference in the search.”*

### **7.5.2 Progressive Disclosure and the Stages of Information Search and Triage**

- Users quickly understood and articulated how well their needs as information seekers were supported by their interface while performing their tasks. Users connected strongly with the progressively-disclosed design of the evolved interfaces. This allowed them to concentrate on their current stage of information search and triage, instead of having to manage the features for all stages at once, as would be the case in traditional interfaces. Furthermore, users found great benefit in features like the search formulation “sandbox” in Stage 2/3’s Query Building, the search prediction matrix, and High-Level Triaging’s timeline-like document abstractions. Together, these considerations allowed users to be much more effective and confident in the quality of their performances as they transitioned between their tasks (e.g., A, B, C).
- Users of Stage 1 found that they were routinely weakened by the characteristics of its traditionally designed interface. Three common observations emerged: First, they found that elements of traditional design — like basic search bars and long results lists displaying minor amounts of information — occluded information which could have assisted them in completing their tasks. Second, the traditionally-paged system made it challenging to navigate and perform their tasks in preferred manners, particularly when comparing between documents during triage. Finally, they felt a lack of trust of the traditional interface, resulting from a ‘black box’ style relationship to ML computation components. These interface elements, such as undefined relevance metrics, non-interactive document sorting, and static summarization strategies, do not align with the objectives of keeping the ‘human-in-the-loop’. As a result, users felt a severe reduction in transparency which created uncertainty in their decision making and suboptimal performance (e.g., D, E, F, G).

(A) *“I value a tool that allows you to jump between different parts of the process and being able to adjust on the fly without having to like add a term and then search again.”*

(B) *“Then when I opened a document, like a summary of what's in it -- I really appreciate that because sometimes I can ... click it and be like: oh this ... has what I need, but it's not exactly what I need. So, having the summary that I can just open and shut really kind of helps. (It allows me to) be like -- OK, this is exactly what I need, or this is close to what I need, or this isn't what I need. I like it on request when it's appropriate.”*

(C) *“Once I was triaging my final documents, the line chart is definitely interesting to get a going because there's a chance for feedback, which I really liked. You can go in and see it's like OK, ‘there's a lot of colors’, and from there it's ... and you kind of compare them. I this is where I went back to the research question and I was going OK. ‘What is it exactly that I wanted?’ And I used that to drive it. So, I ... used the low-level (triage) part of this to try and distinguish between genetic instability or problems in gene regulation over chromosome, and that's where having the summary I think was helpful because it let me kind of dig into questions like: Is this about chromosome instability? ... Is it even cancer related?”*

(D) *“For me, (with Stage 1), it's actually more distracting having all of the information (of the full document) all at once, because I could be like, oh, this is actually like this in my mind is connected to the next thing, even though it isn't.”*

(E) *“Articles quality (in Stage 1) doesn't really necessarily mean its higher quality because of how many keywords are present. I would say that I have seeing the full ... document is better than just seeing the number. Uh, like just I was just like looking up what the focus was for articles and ... I just don't want to overthink the articles selected... But having things like counts and relevance ratings it made it a lot quicker. (I) didn't ... know if the articles I selected were higher quality or ... more beneficial to someone researching that topic, but it does make it makes the search faster and it allows you to sift through (and) give you some confidence.”*

(F) *“Well, (Stage 1 is) sorted by relevance, so starting at the top indicated that most of these entries had to do with the concept of chromosomal instability. And just skimming through the first 20, it's clear, at least on this first page that it's really the first half or so that are relevant. So I think I might have poked ahead to the next page or two just to kind of get a sense of whether something had slipped through, but it really did seem like the first ten of these were the most relevant... So, I just started looking at looking at those one at a time. Just doing a quick skim of the abstracts or summaries.”*

(G) *“Yeah, (with Stage 1), that was something that I really wanted to just ... open another window and start another search because there was times where I wanted to compare certain documents with certain search terms... to really see what kind of impact are these words really having on each other. The only thing I could really do (in Stage 1) was save the documents and then go up to another search. Yeah, which is not a very efficient way of going about it.”*

### **7.5.3 Ontology Integration**

- Users found the ontology integration within Stage 3 to be incredibly valuable during the performance of many tasks. Users was able to effectively apply available ontology mediation into the performance of tasks which benefited from additional domain knowledge, such as Task 2, Task 6, and Task 7. Users described

their experience of searching and triaging a document set from a domain that they did not have immediate familiarity, and how domain context can be a significant difference maker in their successes (e.g., A, B).

- While none of the users were domain experts, some were able to take advantage of prior experience from searching and triaging large document sets, as well as applying limited domain knowledge into their task performances. In tasks where domain context was beneficial, these users found that that access to the ontology allowed them to translate their existing expertise into new task contexts and was critical to the successes of their performances. One user described how they used their domain experience to align with ontology mediation during the final task (e.g., C).

(A) *“But ... you can get loss in the information too, right? So, if you have like so much so many things related in that ontology, it's like, well, it can be useful. But it could also be a distraction for something that you know. There's this flip side, but I think that's on the searcher to know what they're using and why they're using it. So ... for me to complete these tasks, if I hadn't had the ontologies listed, then I would have had a much more difficult time. It essentially provided guidance ... and a structure to something I was unfamiliar with in this case.”*

(B) *“I wasn't necessarily intimidated, but I was just like -- I don't know what this is. But the background in information for the context helped a little bit. A lot of big words, but they did help me when I was looking at the documents that I had to search for to find out which ones I felt best. So even though I did not have full understanding of the words, having them there in that background provided me a kind of help towards finding myself in the space of the question.”*

(C) *“... we know what we want: Cancer. that's part of the research question. Chromosome, that's a key one... Damage comes with instability: That's (me using) domain knowledge -- instability leads to genetic damage in some way... (For) cell division -- I mean, all cancers involve abnormalities in cell division. If they don't have that they're not a cancer. They are just dead cells. An abnormality of chromosome stability: that seemed really relevant when it showed up when ... browsing the ontology list. And that's kind of how the query came together. It's what terms kind of capture the same ideas.”*

#### **7.5.4 Limitations**

The primary limitation for the efforts described in this evolutionary design were the restrictions placed upon research due to the COVID-19 pandemic. Specifically, these restrictions limited our ability to facilitate formal user studies with in-person quantitative evaluations using an expanded set of users. To combat the issues presented by this limitation, we propose two solutions: First, we simply must wait until a return to normalcy. Second, the use of low-latency, high-resolution remote access technologies to distribute a centralized study environment and maintained strict minimum technology requirements in the participant approval criteria. With either of these solutions, we would be able to facilitate an expanded, formal user study to the standards we envisioned.

A secondary effect of this limitation of this study is the reduced availability level of domain expertise for participants, which restricted the types of tasks we could ask of users. These restrictions limited the investigation of deeper, domain-dependent aspects of information search and triage. That is, the investigations of this paper target information search and triage tasks in a general sense. If in the future we are able to gain access to users with specialized domain expertise, we will be able to investigate domain-specific considerations for interface design.

### 7.5.5 Conclusions

In this paper, we presented a three-staged evolutionary design, producing a VAT interface for searching and triage large document sets. We began with background on information search and triage within machine learning (ML) - supported VATs, where we examined challenges facing users and potential solutions. We outlined the evolutionary design process and specify a task-driven formative assessment. We then described the results of our evolutionary design, spanning the formulization, realization, and validation of three interfaces: Stage 1, Stage 2, and Stage 3. We provided a general discussion of user responses to the evolutionary design, the value of ontology-supported interfaces for information search, and the promotion of progressive disclosure in interfaces for multi-staged information search and triage on large document sets. Confirmatory evidence from formative assessment described: Users were able to understand and benefit from novel interface designs. Users connected strongly with the progressively disclosed design of the evolved interfaces, found great benefit in novel configurations of evolved interfaces, and were more effective and confident in their task performances. Furthermore, users found the use of ontologies to be incredibly valuable during their performances, which they were able to effectively into tasks which benefited from additional domain context. We ended with limitations and conclusions.

Based on these results, we assess several potential research directions. First, further research can be done to investigate how more specific, lower-level design considerations can impact users as they perform challenging information search and triaging tasks on large document sets. Second, this research provided an initial exploration of ontology mediation for information seekers. However, in the interfaces described in this evolutionary design, some stages of information search and triage were not provided opportunities user-facing ontology mediation. Therefore, future research may find insight in exploring additional points of ontology integration; particularly for the high and low-level triage stages of the information-seeking process. Third, information-seeking tasks can require users to possess extremely refined levels of domain knowledge. Thus, future research could investigate how domain-specific task requirements may affect the performance of information search and triage, and how users can benefit from designs which adjust to variable levels of domain expertise. Fourth and finally, as new visualization technologies like AR, VR, and non-traditional interaction technologies like touch and voice become commonplace, there will be a need to understand how such technologies can be useful in interfaces for searching and triaging large document sets.

### 7.6 References

- Arp, R., Smith, B., Spear, A. D., & American Journal of Sociology. (2015). Building Ontologies with Basic Formal Ontology. In *Journal of Chemical Information and Modeling* (Vol. 53).
- Badi, R., Bae, S., Moore, J. M., Meintanis, K., Zacchi, A., Hsieh, H., ... Marshall, C. C. (2006). Recognizing User Interest and Document Value from Reading and Organizing Activities in Document Triage. *Proceedings of the 11th International Conference on Intelligent User Interfaces - IUI '06*, (1), 218. <https://doi.org/10.1145/1111449.1111496>
- Bae, S., Kim, D., Meintanis, K., Moore, J. M., Zacchi, A., Shipman, F., ... Marshall, C. C. (2010). Supporting

- Document Triage via Annotation-based Multi-application Visualizations. *Proceedings of the 10th ACM/IEEE Joint Conference on Digital Libraries*, 177–186. <https://doi.org/10.1145/1816123.1816150>
- Baldominos, A., Saez, Y., & Isasi, P. (2020). On the automated, evolutionary design of neural networks: past, present, and future. *Neural Computing and Applications*, 32(2), 519–545. <https://doi.org/10.1007/s00521-019-04160-6>
- Bostock, M. (2016). D3.js Data-Driven Documents. Retrieved May 10, 2020, from <https://d3js.org/>
- Buchanan, G., & Owen, T. (2008). Improving Skim Reading for Document Triage. *Proceedings of the Second International Symposium on Information Interaction in Context - IIX '08*, 83. <https://doi.org/10.1145/1414694.1414714>
- Chuang, J., Ramage, D., Manning, C., & Heer, J. (2012). Interpretation and Trust: Designing Model-driven Visualizations for Text Analysis. *Proceedings of the SIGCHI ...*, 443–452. Retrieved from <http://dl.acm.org/citation.cfm?id=2207738>
- Cortez, P., & Embrechts, M. J. (2013). Using Sensitivity Analysis and Visualization Techniques to Open Black Box Data Mining Models. *Information Sciences*, 225, 1–17. <https://doi.org/10.1016/j.ins.2012.10.039>
- Demelo, J., & Sedig, K. (2021). Design of Generalized Search Interfaces for Health Informatics. *Multidisciplinary Digital Publishing Institute Information*, 12(8), 1–30. <https://doi.org/10.3390/info12080317>
- Demelo, J., Sedig, K., & Parsons, P. (2017). OVERT-MED : An Ontology-Driven Visual Search and Triage Interface for MEDLINE. *Journal of Medical Internet Research Medical Informatics*, 5(1). <https://doi.org/10.2196/medinform.6918>
- Endert, A., Ribarsky, W., Turkay, C., Wong, B. L. W., Nabney, I., Blanco, I. D., & Rossi, F. (2017). The State of the Art in Integrating Machine Learning into Visual Analytics. *Computer Graphics Forum*, 36(8), 458–486. <https://doi.org/10.1111/cgf.13092>
- Endert, Alex, Hossain, M. S., Ramakrishnan, N., North, C., Fiaux, P., & Andrews, C. (2014). The Human is the Loop: New Directions for Visual Analytics. *Journal of Intelligent Information Systems*, 43(3), 411–435. <https://doi.org/10.1007/s10844-014-0304-9>
- Golitsyna, O. L., Maksimov, N. V., & Monankov, K. V. (2018). Focused on Cognitive Tasks Interactive Search Interface. *Procedia Computer Science*, 145, 319–325. <https://doi.org/10.1016/j.procs.2018.11.078>
- Guerrero-García, J. (2014). Evolutionary design of user interfaces for workflow information systems. *Science of Computer Programming*, 86, 89–102. <https://doi.org/10.1016/j.scico.2013.07.003>
- Harvey, M., Hauff, C., & Elswailer, D. (2015). Learning by Example: Training Users with High-quality Query Suggestions. *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, New York, NY, USA, 9–13 August 2015*, 133–142.
- Herceg, P. M., Allison, T. B., Belvin, R. S., & Tzoukermann, E. (2018). Collaborative Exploratory Search for Information Filtering and Large-scale Information Triage. *Journal of the Association for Information Science and Technology*, 69(3), 395–409. <https://doi.org/10.1002/asi.23961>
- Hoeber, O. (2014). Visual Search Analytics: Combining Machine Learning and Interactive Visualization to Support Human-Centred Search. *Beyond Single-Shot Text Queries: Bridging the Gap(s) Between Research Communities*, 37–43.

- Hohman, F., Kahng, M., Pienta, R., & Chau, D. H. (2018). Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers. *IEEE Transactions on Visualization and Computer Graphics*, 25(8), 2674–2693.
- Holzinger, A. (2016). Interactive Machine Learning for Health Informatics: When Do We Need the Human-in-the-loop? *Brain Informatics*, 3(2), 119–131. <https://doi.org/10.1007/s40708-016-0042-6>
- Huurdeman, H. C. (2017). Dynamic Compositions: Recombining Search User Interface Features for Supporting Complex Work Tasks. *CEUR Workshop Proceedings*, 1798, 22–25.
- Jakus, G., Milutinovic, V., Omerović, S., & Tomazic, S. (2013). *Concepts, Ontologies, and Knowledge Representation*. <https://doi.org/10.1007/978-1-4614-7822-5>
- Jusoh, S., Awajan, A., & Obeid, N. (2020). The Use of Ontology in Clinical Information Extraction. *Journal of Physics: Conference Series*, 1529(5). <https://doi.org/10.1088/1742-6596/1529/5/052083>
- Katifori, A., Torou, E., Vassilakis, C., Lepouras, G., & Halatsis, C. (2008). Selected Results of a Comparative Study of Four Ontology Visualization Methods for Information Retrieval Tasks. *Proceedings of the Second International Conference on Research Challenges in Information Science, Marrakech, Morocco, 3–6 June 2008*, (September 2016), 133–140. <https://doi.org/10.1109/RCIS.2008.4632101>
- Khan, S., Kanturska, U., Waters, T., Eaton, J., Bañares-Alcántara, R., & Chen, M. (2016). Ontology-assisted Provenance Visualization for Supporting Enterprise Search of Engineering and Business Files. *Advanced Engineering Informatics*, 30(2), 244–257. <https://doi.org/10.1016/j.aei.2016.04.003>
- Köhler, S., Carmody, L., Vasilevsky, N., Jacobsen, J. O. B., Danis, D., Gourdine, J.-P., ... Robinson, P. N. (2018). Expansion of the Human Phenotype Ontology (HPO) Knowledge Base and Resources. *Nucleic Acids Research*, 47(D1), D1018–D1027. <https://doi.org/10.1093/nar/gky1105>
- Kohler, S., Doelken, S. C., Mungall, C. J., Bauer, S., Firth, H. V., Bailleul-Forestier, I., ... Robinson, P. N. (2014). The Human Phenotype Ontology Project: Linking Molecular Biology and Disease Through Phenotype Data. *Nucleic Acids Research*, 42(D1), 966–974. <https://doi.org/10.1093/nar/gkt1026>
- Köhler, S., & Robinson, P. (2016). HPO Web Browser: Blindness Infopage. Retrieved from Human Phenotype Ontology website: <http://compbio.charite.de/hpweb/showterm?id=HP:0000618>
- Loizides, F., Buchanan, G., & Mavri, K. (2016). Theory and Practice in Visual Interfaces for Semi-structured Document Discovery and Selection. *Information Services & Use*, 35(4), 259–271. <https://doi.org/10.3233/ISU-150783>
- Lytvyn, V., Dosyn, D., Vysotska, V., & Hryhorovych, A. (2020). Method of Ontology Use in OODA. *Proceedings of the IEEE Third International Conference on Data Stream Mining & Processing, Lviv, Ukraine, 21–25 August 2020*, 409–413. <https://doi.org/10.1109/dsmp47368.2020.9204107>
- Mehta, N., & Pandit, A. (2018). Concurrence of Big Data Analytics and Healthcare: A Systematic Review. *International Journal of Medical Informatics*, 114(March), 57–65. <https://doi.org/10.1016/j.ijmedinf.2018.03.013>
- PubMed. (n.d.). Retrieved May 10, 2020, from <https://pubmed.ncbi.nlm.nih.gov/>
- Ramanujan, D., Chandrasegaran, S. K., & Ramani, K. (2017). Visual Analytics Tools for Sustainable Lifecycle Design: Current Status, Challenges, and Future Opportunities. *Journal of Mechanical Design, Transactions of*

- the ASME*, 139(11). <https://doi.org/10.1115/1.4037479>
- Rector, A., Schulz, S., Rodrigues, J. M., Chute, C. G., & Solbrig, H. (2019). On beyond Gruber: “Ontologies” in Today’s Biomedical Information Systems and the Limits of OWL. *Journal of Biomedical Informatics: X*, 2, 100002. <https://doi.org/https://doi.org/10.1016/j.yjbinx.2019.100002>
- Román-Villarán, E., Pérez-Leon, F. P., Escobar-Rodríguez, G. A., Martínez-García, A., Álvarez-Romero, C., & Parra-Calderón, C. L. (2019). An Ontology-based Personalized Decision Support System for Use in the Complex Chronically Ill Patient. *Studies in Health Technology and Informatics*, 264, 758–762. <https://doi.org/10.3233/SHTI190325>
- Sacha, D., Sedlmair, M., Zhang, L., Lee, J. A., Weiskopf, D., North, S., & Keim, D. (2016). Human-centered Machine Learning Through Interactive Visualization: Review and Open Challenges. *ESANN 2016: 24th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning Bruges, Belgium April 27-28-29, 2016 Proceedings*, (27–29 April 2016), 641–646. Retrieved from <http://eprints.mdx.ac.uk/20763/>
- Saleemi, M. M., Rodríguez, N. D., Lilius, J., & Porres, I. (2011). A Framework for Context-aware Applications for Smart Spaces. In *Smart Spaces and Next Generation Wired/Wireless Networking: Vol. 6869 LNCS* (pp. 14–25). [https://doi.org/10.1007/978-3-642-22875-9\\_2](https://doi.org/10.1007/978-3-642-22875-9_2)
- Sedig, K., & Ola, O. (2014). The Challenge of Big Data in Public Helth: An Opportunity for Visual Analytics. *Online Journal of Public Health Informatics*, 5(3), 1–21. <https://doi.org/10.5210/ojphi.v5i3.4933>
- Sedig, K., & Parsons, P. (2016). Design of Visualizations for Human-Information Interaction: A Pattern-Based Framework. In *Synthesis Lectures on Visualization* (Vol. 4). <https://doi.org/10.2200/S00685ED1V01Y201512VIS005>
- Solr Cloud. (2020). Retrieved May 10, 2020, from <https://lucene.apache.org/solr/>
- Spence, R. (2002). Sensitivity Encoding to Support Information Space Navigation: A Design Guideline. *Information Visualization*, 1(2), 120–129. <https://doi.org/10.1057/palgrave.ivs.9500019>
- Spence, R. (2004). Sensitivity, Residue and Scent. *Information Design Journal + Document*, 12(3). <https://doi.org/10.1075/idjdd.12.3.01spe>
- Spence, R. (2014). Interaction. In *Information Visualization* (pp. 173–223). Springer.
- Stouffs, R., & Rafiq, Y. (2015). Generative and evolutionary design exploration. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing: AIEDAM*, 29(4), 329–331. <https://doi.org/10.1017/S0890060415000360>
- Talbot, J., Lee, B., Kapoor, A., & Tan, D. S. (2009). EnsembleMatrix: Interactive Visualization to Support Machine Learning with Multiple Classifiers. *Proceedings of the 27th International Conference on Human Factors in Computing Systems - CHI 09*, 1283. <https://doi.org/10.1145/1518701.1518895>
- The Human Phenotype Ontology. (2020). Retrieved May 10, 2020, from The Jackson Labratory website: <https://hpo.jax.org/app/>
- Tresp, V., Marc Overhage, J., Bundschus, M., Rabizadeh, S., Fasching, P. A., & Yu, S. (2016). Going Digital: A Survey on Digitalization and Large-Scale Data Analytics in Healthcare. *Proceedings of the IEEE*, 104(11),

2180–2206. <https://doi.org/10.1109/JPROC.2016.2615052>

- Wall, E., Blaha, L. M., Franklin, L., & Endert, A. (2018). Warning, Bias May Occur: A Proposed Approach to Detecting Cognitive Bias in Interactive Visual Analytics. *Proceedings of the IEEE Conference on Visual Analytics Science and Technology, Phoenix, AZ, USA, 3–6 October, 2017*, 104–115. <https://doi.org/10.1109/VAST.2017.8585669>
- Wu, C. M., Meder, B., Filimon, F., & Nelson, J. D. (2017). Asking Better Questions: How Presentation Formats Influence Information Search. *Journal of Experimental Psychology: Learning Memory and Cognition*, 43(8), 1274–1297. <https://doi.org/10.1037/xlm0000374>
- Xing, Z., Wang, L., Xing, W., Ren, Y., Li, T., & Xia, J. (2019). Application of Ontology in the Web Information Retrieval. *Journal on Big Data*, 1(2), 79–88. <https://doi.org/10.32604/jbd.2019.05806>
- Yuan, J., Chen, C., Yang, W., Liu, M., Xia, J., & Liu, S. (2020). A Survey of Visual Analytics Techniques for Machine Learning. *ArXiv*, 1, 1–34.
- Zeng, Q. T., & Tse, T. (2006). Exploring and Developing Consumer Health Vocabularies. *Journal of the American Medical Informatics Association*, 13(1), 24–29. <https://doi.org/10.1197/jamia.M1761>
- Zhang, Y., He, H., Legunsen, O., Li, S., Dong, W., & Xu, T. (2021). *An Evolutionary Study of Configuration Design and Implementation in Cloud Systems*. 188–200. <https://doi.org/10.1109/icse43902.2021.00029>

## Chapter 8      Summary, Contributions, and Future Research

In this chapter, we provide summaries of the five integrated article chapters, the general contributions of this dissertation, and thoughts of future research beyond the scope of this dissertation.

### 8.1 Chapter Summaries

The following provides high-level summary of the five integrated article chapters.

In **Chapter 3**, we explored the design of generalized visual interfaces for information search and triage, and the activation of ontologies during query building. Included in this chapter was a presentation of OVERT-MED, an ontology-driven visual interface for searching and triaging MEDLINE. OVERT-MED examined the process of designing visual interfaces for information search and triage. A particular concentration in its design was the use of ontologies as mediating resources during query building, and progressive disclosure in interface design. The research of this chapter was used as an exploratory piece which inspired future research in later chapters.

In **Chapter 4**, we investigated how humans understand ontologies through cognitive map formation, and how visual interfaces can support users to encounter, form knowledge of, and explore complex ontological space. Included in this chapter was a presentation of PRONTOVISE, a progressively disclosed and scaffolded generalized visualization environment for exploring the ontological space of user-provided ontology files. This research was important to the overall research direction of this dissertation, as it helped establish if users can reasonably learn the complex ontological space described by ontology files. From its findings, we were able to perform further investigations of their activation within visual interfaces for information search and triaging tasks.

In **Chapter 5**, we framed the high-level requirements for generalized search interfaces. Included in this chapter was a topic analysis which explored existing materials on data sources, information characteristics, types of search tasks, and considerations for generalized interfaces in the health domain. From this analysis, we distilled a set of high-level requirements, which we then used to structure the design of ONTSI, a demonstrative generalized search interface which integrates ontologies for mediating between common and domain vocabularies. ONTSI was inspired by the first stage of the evolutionary design described in a later chapter.

In **Chapter 6**, we investigated the multi-staged information-seeking process and its design requirements. These investigations explored novel techniques which reduce the need for tedious search and triage. Specifically, it established how users struggle within traditional design strategies, align configurations with the requirements of each stage of the information-seeking process, and learned of strategies to promote the activation of domain expertise when searching and triaging. Furthermore, the use of ontologies to support these requirements was

examined. Included in this chapter was a presentation of VisualQUEST, a novel visual interface which generalizes the support of information search and triage using the progressive disclosure and ontology mediation over user provided ontology and document datasets. VisualQUEST was built from the evolutionary design described in a later chapter.

In **Chapter 7**, we described a three-staged evolutionary design of a Visual Analytics Tool (VAT) interface for searching and triaging large document sets. Specifically, the chapter outlined the formulization, realization, and validation of three stages of an interface generated with guidance from in-depth topic analysis, design criteria, and formative assessment. Confirmatory evidence from formative assessment described: Users were able to understand and benefit from novel interface designs. Users connected strongly with the progressively disclosed design of the evolved interfaces, found great benefit in novel configurations of evolved interfaces, and were more effective and confident in their task performances. Furthermore, users found the use of ontologies to be incredibly valuable during their performances, which they were able to effectively into tasks which benefited from additional domain context.

## **8.2 General Contributions**

Each chapter describes contributions to the research objective. The following is a high-level summary of the general contributions of this dissertation. The discussion is divided into two categories of contribution: research and practical.

### **8.2.1 Research Contributions**

The overarching research objective of this dissertation is to improve our understanding of how the interface designs of information search and triage tools can impact users as they search and triage large document sets. Over the course of the five materials included within this integrated article, we have made efforts to pursue this research objective: (1) We investigated existing designs of information search and triage interfaces to establish an initial assessment of high-level design requirements and proposed exploratory proof-of-concepts for the activation of ontologies during query building visual interfaces; (2) We explored how knowledge of complex ontological space is formed and established novel designs for supporting users to encounter, form knowledge of, and explore complex ontological space; (3) We performed in-depth topic analysis to explore existing materials on data sources, information characteristics, types of search tasks, and considerations for interfaces in the health domain. From that analysis, we distilled criteria to guide the design of generalized search interfaces; (4) We provided in-depth topic analysis on the multi-staged information-seeking process and its design requirements, where we sought to explore novel techniques which reduce the need for tedious search and triage, establish a stronger awareness of how users struggle within traditional design strategies, align configurations with the requirements of each stage of the information-seeking process, and learn of ways to promote the use of domain expertise; (5) We described a three-staged evolutionary design of a Visual Analytics Tool (VAT) interface for searching and triaging large document sets. Specifically, the formulization, realization, and validation of three stages of an interface generated with guidance from in-depth topic analysis, design criteria, and formative assessment. Confirmatory evidence from formative assessment described: Users were able to understand and benefit from novel interface designs. Users

connected strongly with the progressively disclosed design of the evolved interfaces, found great benefit in novel configurations of evolved interfaces, and were more effective and confident in their task performances. Furthermore, users found the use of ontologies to be incredibly valuable during their performances, which they were able to effectively into tasks which benefited from additional domain context. We hope that designers will find inspiration and guidance from these research contributions within future investigations of information search and triage.

## **8.2.2 Practical Contributions**

Within this dissertation, practical contributions were produced which help users pursue information search and triage on large document sets. When investigating the potential for novel designs which support users in their information-seeking process, we generated OVERT-MED, an ontology-driven visual interface for searching and triaging MEDLINE. We facilitated an initial exploration into the design process of an information search and triage interface. In doing so, we identified functional requirements for integrating and activating the OWL ontology file format and the MEDLINE document structure within visual interfaces. With those efforts, we gathered understanding which could assist our ability to provide plug-and-play capabilities of user-provided ontology files and document sets within future practical contributions. Next, our investigations of cognitive map formation of complex ontological space allowed us to generate PRONTOVISE, a progressively disclosed and scaffolded generalized visualization tool for learning the complex ontological space of user-provided ontology files. We demonstrated the use of dynamic and deeply layered visual representation and interaction techniques such as progressive disclosure and scaffolding to support users in the performance of complex learning using visual interfaces. Based on these efforts, we were able to initialize an evolutionary design which navigated the formulization, realization, and validation of three stages of an interface generated with guidance from in-depth topic analysis, design criteria, and formative assessment. Through this design process, we formalized three distinct interface designs which were then realized as working prototypes. These prototypes were applied within task-driven formative assessment, generating confirmatory evidence to support future formulization and realization efforts. The product of this evolutionary design resulted in the progressively disclosed, ontology-supported VisualQUEST, which allows for the activation of user-supplied ontology files and document sets within a generalized search and triage VAT interface. Additional prototypes diverging from the first stage of the overall evolutionary design were also made to explore the integration of ontology-support within generalized search interfaces, in the form of the traditionally designed, yet ontology-supported ONTSI. We hope that designers will find inspiration and guidance from these practical contributions when creating the information search and triage interfaces of the future.

## **8.3 Future Research**

The following is an exploration of future research inspired by the content of this dissertation. The discussion is divided into three categories: Future research involving expanded empirical studies on existing topics, micro-level

frameworks for guiding interface design, and added investigations of novel modes of representation and interaction in interfaces for information search and triage.

### **8.3.1 Formal User Studies and Expanded Investigations**

Chapter 7 described a three-staged evolutionary design of a VAT interface for searching and triaging large document sets. Following evolutionary design practices, repeated formulization, realization, and validation efforts were performed with guidance from in-depth topic analysis, design criteria, and formative assessment. These assessments produced confirmatory evidence on a variety of topics, as described in their original published materials and summarizing material. Yet, there were limitations placed upon our efforts. Namely, the onset and active continuation of the COVID-19 pandemic dramatically impacted our ability to pursue research objectives in its originally envisioned form. Before the pandemic, we had planned to generate qualitative and quantitative metrics within formal user studies which could fortify our understanding of the information-seeking process and the role visual interface design can play when searching and triaging large document sets. Study procedures were described (Appendix A). Test environments were prepared (Appendices B, D, E). Recording equipment and software was developed which could facilitate the collection of qualitative and quantitative metrics (Appendix C, F). However, as the severity of the situation became clear to all, faculties began to close, ethics review boards were halted, and physical distancing restrictions were activated. This eliminated all possibility to progress under the same direction. Therefore, we adjusted research directions in a manner which respected the limitations placed upon us. Specifically, we reorganized our efforts under the perspective of an evolutionary design, performed in-depth topic analyses of relevant published research and their formal user studies, integrating findings with available informal accounts generated from formative assessment periods. We believe that under limiting circumstances, we have put forth the greatest effort possible to maximize research value in the spirit of our original research objectives.

Still, the results of our adjusted direction brought to light additional investigation opportunities not originally conceptualized prior to our research efforts. Expanded investigations could explore questions such as: Are there aspects of information search and triage that are domain-specific? If so, how can the characteristics of the document set and its information be presented during information search and triage to match the domain-specific requirements? What effects do alternate mediation sources have on the performance of an information search and triage task? How significantly are users affected by their existing level of domain expertise within the information-seeking process? How would visual interface design help or hinder users of different domain expertise, and if so, can those requirements and their solutions be described in a generalized, prescriptive form to assist designers?

We believe the research of this dissertation is merely a first step into a larger effort to understanding information search and triage, and we hope these materials can be useful for researchers and designers as they explore novel applications in visual interface design for searching and triaging large document sets.

### **8.3.2 Prescriptive Frameworks**

Alongside investigations of the high-level requirements for user-centered information search and triage visual interfaces, deeper investigation could be made into how specific, micro-level design elements impact the performance

of information search and triage tasks. Within the confirmatory evidence produced during formative assessment, as described in Chapter 7, we were able to learn of some benefits provided by micro-level design considerations. For instance, users benefited from and in turn expressed their appreciation for encounters with novel interactive visual representations like the ontology listing in Query Building, the predictive search matrix in the Search, the Query Result Heatmap in High-Level Triage, and the various document summations strategies in Low-Level Triage. Designers can benefit from prescriptive resources like frameworks when assessing the characteristics of their tasks, users, and data sources. Frameworks can guide the design process and help organize relevant concepts in the design space, supporting a generative role in design thinking, aiding in reflection, helping interpret requirements, and expand upon existing designs. These resources can give direction to how the requirements of those characteristics translate into concrete design decisions. For instance, researchers have previously performed studies on document summation techniques, where specific parts of documents have been studied to determine which parts of document content should be kept or removed when summarizing. If similar investigations were to be made across the stages of the information-seeking process, levels of domain expertise, vocabularies, information characteristics, and datasets, then overarching frameworks can be formulated which could provide guidance during the design process.

### **8.3.3 Novel Modes of Representation and Interaction for Search and Triage Interfaces**

The research of this dissertation concentrated on investigations of information search and triage interfaces for computer systems maintaining a computer monitor, mouse, keyboard, network connectivity, and other standard items of the typical peripheral suite. However, it would be unwise to assume that the computer as we currently know it, will remain the computer of the future. Already, we are beginning to see paradigm shifts to novel displays, new types of interaction, and new modes of interconnectivity. Like all major technological innovations of their time, these novel modes of representation and interaction will certainly be leveraged within tools in whatever form best satisfied stakeholder requirements. Thus, further research must be conducted to examine how novel modes of representation, such as small displays, large displays, virtual reality, and augmented reality, and modes of interaction such as touch, voice command, and cognitive implants could be used to interface users with their computational technologies which help them achieve their information-seeking objectives.

# Appendices

## Appendix A Formal Study Letter of Information and Consent

### Project Title

Investigating Visual Analytics Tool Interface Design for Searching and Triaging Large Document Sets

### Document Title

Letter of Information and Consent – Exploration Session

### Principal Investigator

Dr. Kamran Sedig, Professor, Computer Science, Faculty of Information and Media Studies  
Western University

### 1. Sponsor/Funder Information

1.1. The study is self-funded.

### 2. Conflict of Interest

2.1. There are no conflicts of interest for any of the investigators, study staff, or member(s) of their immediate family.

### 3. Invitation to Participate

3.1 You are being invited to participate in this research study about the design of visual analytics tool interfaces for searching and triaging large document sets.

### 4. Why is this study being done?

4.1 There is a growing desire for novel visual analytics tools (VATs) which help us to complete our increasingly challenging information search and triage tasks. Yet, as the document sets of our tasks grow ever larger and the computational processes of our VATs rise in complexity, there is a need to examine how novel designs of human-facing visual interfaces impact the performances of information search and triage.

This study is being conducted to examine and compare the effectiveness of a set of three interfaces. We are looking to explore the use of these unique interfaces, which possess varying design considerations either based on “current best practices” or novel techniques which we believe can help users perform their tasks. We hope that we may gather insight towards their potential use in establishing new paradigms for the design of interfaces which support information search and triage. We hope that your participation in our study will help us achieve these objectives.

### 5. How long will you be in this study?

5.1. It is expected that you will be in the study for a single day, there will be 1 study visit during your participation in this study, and the visit will take approximately 1 hour. If you are randomly selected for a post-performance interview, the visit duration may increase in time approximately 15 minutes, for a total of approximately 75-85 minutes.

### 6. What will happen during this study?

**6.1.** If you decide to participate then you will be “randomized” into one of three groups, as described in Section 7. Randomization means that you are put into a group by chance (like flipping a coin). There is no way to predict which group you will be assigned to. You will have a 33%, or 1/3 chance of being placed in either/any group. Neither you nor the researchers can choose what group you will be in. During the study, the investigator(s) will know which group you are in.

**6.2.** Members of the development team for the study tool are excluded from participating in the study.

**6.3.** We anticipate the study to include 12-45 participants.

## **7. What are the study procedures?**

**7.1.** If you agree to participate you will be asked to: Attend one study session of approximately 1 hour for the randomly determined non-interview variant, or 75-85 minutes for a post-study interview variant. Potential participants will be provided the required study disclosure materials, asked to review those materials, and acknowledge their willingness to participate in the study session in the form of their written signature. Once a signature is provided, the investigator will randomly assign the participant to one of three interface variants, each of which possesses a unique interface for performing the task set. At this time, the investigator will also randomly generate if the participant will be asked to perform a post-study interview; and knowledge of this assignment will not be known to the participant until after the completion of their assigned tasks. The participant will then be asked to fill out a pre-study questionnaire which will collect basic personal information from the participant, as well as information regarding the participants experience with the technologies and domains covered in the study. The participant will then be provided a general overview description of the study and will be allowed a short period of time to explore their assigned interface. The participant will then be provided the specifics of their task set. The participant will then be asked to begin their tasks. The task portion of the session will end when a student has completed their objectives. If the participant was randomly determined to not require a post-study interview, then the participant will be told that they will not require a post-study interview, and their session will conclude. If the participant was randomly assigned to a post-study interview, they will be informed of this assignment at this time. The participant may choose not to do the post-study audio recorded interview and will be presented that option. If they accept the request for the post-study interview, the investigator will prepare and perform that interview, which will require the participant to provide their perspective on their study experience. This recording will be in the form of an audio recording. The session will conclude after the completion of the interview.

**7.2.** If you are selected to participate in a post-study interview, the question-driven discussion will be recorded in audio form, to later be transcribed as text.

**7.3.** The study takes place in a computer lab environment. This may include accessing computer hardware, audio recording hardware, written content, physical documentation materials, etc. The participant will not be required to bring any equipment into the study.

**7.4.** Other than basic personal information (name, age, education, prior experience with involved technologies, etc.), no sensitive information will be required from participants.

## **8. What are the risks and harms of participating in this study?**

**8.1.** Feelings and emotions which may arise from a) encountering and/or performing within a test-like environment pertaining to information which you might consider unfamiliar (akin to the feeling of taking a test without prior study) and/or b) interacting with a software tool interface for the first time. Otherwise, there are no known or anticipated risk or discomfort associated with participating in this study. If discomfort associated with your participation in this study does arise, support resources can be found for Western students at Western Health and Wellness (<https://www.uwo.ca/health/index.html>).

## **9. What are the benefits?**

**9.1.** You may not directly benefit from participating in this study, but insight gathered may provide benefits towards the design of visual analytics tools for challenging information-based search and triage tasks on large document sets.

#### **10. Can participants choose to leave the study?**

**10.1.** If you decide to withdraw from the study, the information that was collected prior to you leaving the study will still be used as the researchers will be unable to identify an individual participant's responses. No new information will be collected without your permission.

#### **11. How will participants' information be kept confidential?**

**11.1.** Representatives of Western University's Non-Medical Research Ethics Board may require access to your study-related records to monitor the conduct of the research. Otherwise, no people/groups/organizations outside the study team will have access to information collection. Your basic, non-identifiable information (e.g., age, department of study, quotes, etc.) will be collected during your participation of the study, alongside the observation, questioned, and/or analytic results produced by your actions while you participate.

**11.2.** a) Basic, non-identifiable information (e.g., age, department of study, quotes, etc.) will be collected during your participation of the study. b) The study dissemination will combine the collected non-identifiable information alongside any observation, questioned, and/or analytic results produced by your participation. c) Consent will be acquired to disclose non-identifiable information for the dissemination of the study in an anonymous and collective fashion.

**11.3.** No identifiable information will be shared with others outside the study team.

**11.4.** Identifiable information (e.g., name, email address) will be kept for the length of the overall study, to facilitate the requirements of running the study (e.g., ensure no duplicate study performances, pursue follow-up email inquiries if required, etc.).

**11.5.** If the results of the study are published, your name will not be used.

**11.6.** Participation in this study will be performed individually.

**11.7.** Data collection in this study will be performed internally by the study team.

**11.8.** If randomly selected for a follow-up interview, quotes may be included within the dissemination in an anonymous fashion.

**11.9.** All data will be collected anonymously and neither the researchers nor anyone else will be able to identify you as a research participant. The data will be stored on a secure server at Western University and will be retained for a minimum of 7 years. Your data may be retained indefinitely and could be used for future research purposes (e.g., to answer a new research question). By consenting to participate in this study, you are agreeing that your data can be used beyond the purposes of this present study by either the current or other researchers. The study team will retain the anonymous, non-identifiable information collected during your participation for future use.

**11.10.** All identifiable information will be deleted from the dataset collected so that individual participant's anonymity will be protected. The de-identified data will be accessible by the study investigators as well as the broader scientific community. More specifically, the data available to other researchers upon publication so that data may be inspected and analyzed by other researchers. The data that will be shared will not contain any information that can identify you.

#### **12. Are participants compensated to be in this study?**

This study does not involve active compensation.

### **13. What are the Rights of Participants?**

**13.1.** Your participation in this study is voluntary. You may decide not to be in this study. Even if you consent to participate you have the right to withdraw from the study at any time. If you choose not to participate or to leave the study at any time it will have no effect on your employment status and/or academic standing.

You do not waive any legal right by consenting to this study.

### **14. Commercialization**

**14.1.** There is no immediate claim for commercialization from the results of this study.

### **15. Whom do participants contact for questions?**

**15.1.** If you have questions about this research study, please contact:

**Principal Investigator**

Dr. Kamran Sedig, Professor, Computer Science, Faculty of Information and Media Studies  
Western University

If you have any questions about your rights as a research participant or the conduct of this study, you may contact The Office of Human Research Ethics (519) 661-3036, 1-844-720-9816, email: [ethics@uwo.ca](mailto:ethics@uwo.ca). This office oversees the ethical conduct of research studies and is not part of the study team. Everything that you discuss will be kept confidential.

**This letter is yours to keep for future reference.**

**16. Consent**

**Written Consent**

**1. Project Title**

Investigating Visual Analytics Tool Interface Design for Searching and Triaging Large Document Sets

**2. Document Title**

Letter of Information and Consent

**3. Principal Investigator**

Dr. Kamran Sedig, Professor, Computer Science, Faculty of Information and Media Studies  
Western University

I have read the Letter of Information, have had the nature of the study explained to me and I agree to participate. All questions have been answered to my satisfaction.

I agree to be audio-recorded in this research.

YES       NO

I consent to the use of unidentified quotes obtained during the study in the dissemination of this research.

YES       NO

I consent to the use of my data for future research purposes.

YES       NO

\_\_\_\_\_  
Print Name of Participant      Signature      Date (DD-MMM-YYYY)

My signature means that I have explained the study to the participant named above. I have answered all questions.

\_\_\_\_\_  
Print Name of Person      Signature      Date (DD-MMM-YYYY)  
Obtaining Consent

## Appendix B Task Questions

Q1 Task 1.1 Using the tool for assistance, which term between "cancer" and "leukemia" has the highest rate of document occurrence within the document set?

Cancer.

Leukemia.

I cannot answer.

Q2 Task 1.2 Using the tool for assistance, which term between "cancer" and "treatment" has the highest rate of document occurrence within the document set?

Cancer.

Treatment.

I cannot answer.

Q3 Task 1.3 Using the tool for assistance, which term between "leukemia" and "treatment" has the highest rate of document occurrence within the document set?

Leukemia.

Treatment.

I cannot answer.

Q5 Task 2.1 Using the tool for assistance, which of the following definitions best aligns with the term "nevus"?

A mole or birthmark on the body.

Abnormal umbilical morphology.

The position of nerve damage in a cerebral cortex.

I cannot answer.

Q6 Task 2.2 Using the tool for assistance, which of the following terms best aligns with the definition "damage to the optic nerve head"?

Psoriasis.

Glaucoma.

Sarcoma.

I cannot answer.

Q7 Task 2.3 Using the tool for assistance, which of the following terms is not "an abnormality of the liver"?

Cirrhosis.

Portal Fibrosis.

Ankylosis.

I cannot answer.

Q10 Task 3.1 After performing a search using the terms "leukemia", "fever", and "symptoms" (but with the limitation that you cannot open any individual document to view its content), approximately what is the percentage of documents from the 10000 documents aligns with "one or more" of the three provided terms?

0% - 0.01% (0-1 documents)

0.02% - 1% (2-100 documents)

- 1.01% - 10% (101-1000 documents)
- 10.01% - 25% (1001-2500 documents)
- 25.01% - 40% (2501-4000 documents)
- 40.01% - 70% (4001-7000 documents)
- 70.01% - 100% (7001-10000 documents)

I cannot answer.

Q11 Task 3.2 After performing a search using the terms "leukemia", "fever", and "symptoms" (but with the limitation that you cannot open any individual document to view its content), approximately what is the percentage of documents from the 10000 documents aligns with the term "leukemia"?

- 0% - 0.01% (0-1 documents)
- 0.02% - 1% (2-100 documents)
- 1.01% - 10% (101-1000 documents)
- 10.01% - 25% (1001-2500 documents)
- 25.01% - 40% (2501-4000 documents)
- 40.01% - 70% (4001-7000 documents)
- 70.01% - 100% (7001-10000 documents)

I cannot answer.

Q12 Task 3.3 After performing a search using the terms "leukemia", "fever", and "symptoms" (but with the limitation that you cannot open any individual document to view its content), approximately what is the percentage of documents from the 10000 documents aligns with "the combination of all three" of the provided terms?

- 0% - 0.01% (0-1 documents)
- 0.02% - 1% (2-100 documents)
- 1.01% - 10% (101-1000 documents)
- 10.01% - 25% (1001-2500 documents)
- 25.01% - 40% (2501-4000 documents)
- 40.01% - 70% (4001-7000 documents)
- 70.01% - 100% (7001-10000 documents)

I cannot answer.

Q14 Task 4.1 After performing a search using the terms "Overweight", "Children", and "Prevention", and then locating, but not opening, the document titled "Overweight children reduce their activity levels earlier in life than healthy weight children." select from the following the combination of terms which you predict will align with the contents of the document:

- Overweight, Children.
- Children, Prevention.
- Overweight, Prevention.
- Overweight, Children, Prevention.
- None of these term combinations.

I cannot answer.

Q15 Task 4.2 After performing a search using the terms "Overweight", "Children", and "Prevention", and then locating, but not opening, the document titled "First lessons from the Kiel Obesity Prevention Study (KOPS)." select from the following the combination of terms which you predict will align most with the contents of the document:

Overweight, Children.

Children, Prevention.

Overweight, Prevention.

Overweight, Children, Prevention.

None of these term combinations.

I cannot answer.

Q16 Task 4.3 After performing a search using the terms "Overweight", "Children", and "Prevention", and then locating, but not opening, the document titled "Food behaviors and other strategies to prevent and treat pediatric overweight." select from the following the combination of terms which you predict aligns most with the contents of the document:

Overweight, Children.

Children, Prevention.

Overweight, Prevention.

Overweight, Children, Prevention.

None of these term combinations.

I cannot answer.

Q18 Task 5.1 After performing a search using the terms "Overweight", "Children", and "Prevention", and then locating, but not opening, the document titled "Overweight children reduce their activity levels earlier in life than healthy weight children." (referred in short as D1) and the document titled "First lessons from the Kiel Obesity Prevention Study (KOPS)." (referred in short as D2), predict the rate of occurrence for the term "Overweight" between the two documents.

D1 will have a higher rate of occurrence of the term Overweight than D2.

D2 will have a higher rate of occurrence of the term Overweight than D1.

Both D1 and D2 will have no occurrences of the term Overweight.

Both D1 and D2 will have some level of occurrence to the term Overweight, and it appears to be approximately equal.

I cannot answer.

Q19 Task 5.2 After performing a search using the terms "Overweight", "Children", and "Prevention", and then locating, but not opening, the document titled "First lessons from the Kiel Obesity Prevention Study (KOPS)." (referred in short as D1) and the document titled "Food behaviors and other strategies to prevent and treat pediatric overweight." (referred in short as D2), predict the rate of occurrences for the term "Children" between the two documents.

D1 will have a higher rate of occurrence of the term Children than D2.

D2 will have a higher rate of occurrence of the term Children than D1.

Both D1 and D2 will have no occurrences of the term Children.

Both D1 and D2 will have some level of occurrence to the term Children, and it appears to be approximately equal.

I cannot answer.

Q21 Task 6.1 After performing a search using the terms "Overweight", "Children", and "Prevention", locate and open the document titled "Food behaviors and other strategies to prevent and treat pediatric overweight.", which of the following correctly orders the rate of occurrence of each term within the document content (not including the title), where the order is highest rate of occurrence as the first and leftmost term, down to the lowest rate of occurrence as the last and rightmost term.

Children, Overweight, Prevention.

Overweight, Children, Prevention.

Prevention, Children, Overweight.

Overweight, Prevention, Children.

None of these are the correct order.

I cannot answer.

Q22 Task 6.2 After performing a search using the terms "Overweight", "Children", and "Prevention", locate and open the document titled "First lessons from the Kiel Obesity Prevention Study (KOPS)." which of the following correctly orders the rate of occurrence of each term plus a new term – "family"/"families" (consider both to be in the same count) within the document content (not including the title), where the order is highest rate of occurrence as the first and leftmost term, down to the lowest rate of occurrence as the last and rightmost term.

Children, Overweight, Prevention, Family/Families.

Overweight, Family/Families, Children, Prevention.

Prevention, Children, Family/Families, Overweight.

Prevention, Children, Overweight, Family/Families.

None of these are the correct order.

I cannot answer.

Task 7 We ask that you use the full tool to rapidly produce a set of 5 documents which you deemed most relevant to the research question. There is a 20-minute time limit, yet we do not intend for you to spend a significant duration of time performing this task.

Research Question: How does chromosomal instability drive tumor progression?

Background: "Chromosomal instability (CIN), defined by an elevated rate of chromosome mis-segregation and breakage, results in diverse chromosomal aberrations in tumor cell populations. Accumulating cytogenetic analyses of over 60,000 cases of human cancer have indicated that most solid tumors contain chromosomal aberrations, with each tumor displaying a distinct abnormal karyotype. In typical human cancers, one-quarter of the genome is affected by arm-level copy number aberrations. Cancer genome sequencing has revealed

dynamic chromosomal content changes during clonal evolution of the tumor cell population. However, how chromosomal loss or gain drives tumor progression to metastasis remains unknown. It is technically difficult to determine the biological function of a specific chromosomal content change, which may influence the expression of hundreds to thousands of genes. Recently, advanced genome-editing techniques have been used to delete large chromosomal region, even whole chromosomal arm. With the application of new methodology, the findings on chromosomal content changes in continuously isolated phenotypic variants of tumor cells might shed some light on the role of CIN in driving tumor metastatic phenotypic switching. These studies have proved the concept that CIN is playing an important role in cancer progression. Since CIN is one of the most common features of cancer cells, it is believed that CIN could be a potential therapeutic target.”

Q24 Task 7.1 What was the file name of your 1st document?

Q25 Task 7.2 What was the file name of your 2nd document?

Q26 Task 7.3 What was the file name of your 3rd document?

Q27 Task 7.4 What was the file name of your 4th document?

Q28 Task 7.5 What was the file name of your 5th document?

## Appendix C Post-Task Questions

Q4 Post-Task 1.4 I found Tasks 1.1, 1.2, and 1.3 easy to complete.

Strongly Agree.

Agree.

Somewhat Agree.

Neither Agree nor Disagree.

Somewhat Disagree.

Disagree.

Strongly Disagree.

Q8 Post-Task 2.4 Did you use the tool to assist your selections in Tasks 2.1 through 2.3?

Yes, and it significantly helped me.

Yes, and it somewhat helped me.

Yes, although it did not help me.

No, I knew it would not help me based on prior use of the tool interface.

No, for other reasons I elected to not use the assistance of the tool interface at all.

Q9 Post-Task 2.5 I found Tasks 2.1 through 2.3 easy to complete.

Strongly Agree.

Agree.

Somewhat Agree.

Neither Agree nor Disagree.

Somewhat Disagree.

Disagree.

Strongly Disagree.

Q13 Post-Task 3.4 I found Tasks 3.1 through 3.3 easy to complete.

Strongly Agree.

Agree.

Somewhat Agree.

Neither Agree nor Disagree.

Somewhat Disagree.

Disagree.

Strongly Disagree.

Q17 Post-Task 4.3 I found Tasks 4.1 through 4.3 easy to complete.

Strongly Agree.

Agree.

Somewhat Agree.

Neither Agree nor Disagree.

Somewhat Disagree.

Disagree.

Strongly Disagree.

Q20 Post-Task 5.3 I found Task 5.1 and 5.2 easy to complete.

Strongly Agree.

Agree.

Somewhat Agree.

Neither Agree nor Disagree.

Somewhat Disagree.

Disagree.

Strongly Disagree.

Q23 Post-Task 6.3 I found Task 6.1 and 6.2 easy to complete.

Strongly Agree.

Agree.

Somewhat Agree.

Neither Agree nor Disagree.

Somewhat Disagree.

Disagree.

Strongly Disagree.

Q29 Post-Task 7.1 I found Task 7 easy to complete.

Strongly Agree.

Agree.

Somewhat Agree.

Neither Agree nor Disagree.

Somewhat Disagree.

Disagree.

Strongly Disagree.

Q30 Post-Task 7.2 I am confident that the documents I selected for the final set of 5 align with the provided research question.

Strongly Agree.

Agree.

Somewhat Agree.

Neither Agree nor Disagree.

Somewhat Disagree.

Disagree.

Strongly Disagree.

Q31 Post-Task 7.3 While selecting the documents, I believe was able to assess the full document set.

Strongly Agree.

Agree.

Somewhat Agree.

Neither Agree nor Disagree.

Somewhat Disagree.

Disagree.

Strongly Disagree.

Q32 Post-Task 7.4 There are no documents that were more relevant to the research question than the five I selected.

Strongly Agree.

Agree.

Somewhat Agree.

Neither Agree nor Disagree.

Somewhat Disagree.

Disagree.

Strongly Disagree.

## Appendix D List of Documents for Task 7

- Determinants and clinical implications of chromosomal instability in cancer (Sansregret, Vanhaesebroeck, & Swanton, 2018)
- Chromosomal instability: A common feature and a therapeutic target of cancer (Tanaka & Hirota, 2016)
- The role of chromosomal instability in tumor initiation (Nowak et al., 2002)
- Chromosomal instability (CIN): what it is and why it is crucial to cancer evolution (Heng et al., 2013)
- Cancer morphology, carcinogenesis and genetic instability: a background (Bignold, Coghlan, & Jersmann, 2006)
- Chromosomal instability and transcriptome dynamics in cancer (Stevens, Horne, Abdallah, Ye, & Heng, 2013)
- Autophagy suppresses tumor progression by limiting chromosomal instability (Mathew et al., 2007)
- A double-edged sword: how oncogenes and tumor suppressor genes can contribute to chromosomal instability (Orr & Compton, 2013)
- Defining 'chromosomal instability' (Geigl, Obenauf, Schwarzbraun, & Speicher, 2008)
- Role of chromosomal instability in cancer progression (McClelland, 2017)

## References

- Bignold, L. P., Coghlan, B. L. D., & Jersmann, H. P. A. (2006). Cancer Morphology, Carcinogenesis and Genetic Instability: A Background. *Cancer: Cell Structures, Carcinogens and Genomic Instability*, 1–24.
- Geigl, J. B., Obenauf, A. C., Schwarzbraun, T., & Speicher, M. R. (2008). Defining “Chromosomal Instability.” *Trends in Genetics*, 24(2), 64–69. <https://doi.org/10.1016/j.tig.2007.11.006>
- Heng, H. H., Bremer, S. W., Stevens, J. B., Horne, S. D., Liu, G., Abdallah, B. Y., ... Ye, C. J. (2013). Chromosomal instability (CIN): What It Is and Why It Is Crucial to Cancer Evolution. *Cancer and Metastasis Reviews*, 32(3–4), 325–340. <https://doi.org/10.1007/s10555-013-9427-7>
- Mathew, R., Kongara, S., Beaudoin, B., Karp, C. M., Bray, K., Degenhardt, K., ... White, E. (2007). Autophagy Suppresses Tumor Progression by Limiting Chromosomal Instability. *Genes and Development*, 21(13), 1701. <https://doi.org/10.1101/gad.1545107.in>
- McClelland, S. E. (2017). Role of Chromosomal Instability in Cancer Progression. *Endocrine-Related Cancer*, 24(9), T23–T31. <https://doi.org/10.1530/ERC-17-0187>
- Nowak, M. A., Komarova, N. L., Sengupta, A., Jallepalli, P. V., Shih, L. M., Vogelstein, B., & Lengauer, C. (2002). The Role of Chromosomal Instability in Tumor Initiation. *Proceedings of the National Academy of Sciences of the United States of America*, 99(25), 16226–16231. <https://doi.org/10.1073/pnas.202617399>
- Orr, B., & Compton, D. A. (2013). A Double-Edged Sword: How Oncogenes and Tumor Suppressor Genes Can Contribute to Chromosomal Instability. *Frontiers in Oncology*, 3 JUN(June), 1–14. <https://doi.org/10.3389/fonc.2013.00164>
- Sansregret, L., Vanhaesebroeck, B., & Swanton, C. (2018). Determinants and Clinical Implications of Chromosomal Instability in Cancer. *Nature Reviews Clinical Oncology*, 15(3), 139–150. <https://doi.org/10.1038/nrclinonc.2017.198>
- Stevens, J. B., Horne, S. D., Abdallah, B. Y., Ye, C. J., & Heng, H. H. (2013). Chromosomal Instability and

Transcriptome Dynamics in Cancer. *Cancer and Metastasis Reviews*, 32(3–4), 391–402.  
<https://doi.org/10.1007/s10555-013-9428-6>

Tanaka, K., & Hirota, T. (2016). Chromosomal Instability: A Common Feature and a Therapeutic Target of Cancer. *Biochimica et Biophysica Acta - Reviews on Cancer*, 1866(1), 64–75.  
<https://doi.org/10.1016/j.bbcan.2016.06.002>

## Appendix E Interview Questions

*[Depending on where the discussion heads, seek out interesting points of contrast between tasks.]*

1. Can you describe to me your initial thoughts as you received <Task #>, and perhaps walk me through how you thought about and enacted your plan to complete this task using your provided tool.
2. Were there any parts of this task that you felt were challenging because of how the tool did or did not support you?
3. Were there any parts of this task that you felt were easy because of how the tool did or did not support you.
4. What task did you feel was the hardest to perform? Easiest?

*[Introduce one interface which is alternate to their session interface.]*

1. What are your initial thoughts on this new interface?
2. How do you think this interface compares to the interface you were given for the tasks?
3. Are there any tasks that you've done today that this new interface would've improved your performance over when you used your original?
4. How about worse?

*[Adjust discussions to match any interesting points of contrast generated in initial interview questions, but now relate to alternate interface. If deemed valuable, continue by showing second alternate interface, repeating the question set. Continue once discussion has completed for alternate interfaces.]*

1. Do you have any final thoughts on how your experience went today?

# Curriculum Vitae

## Education

### Doctor of Philosophy, Computer Science (Candidate)

Dissertation: Evolutionary Design of Search and Triage Interfaces for Large Document Sets

Supervisor: Dr. Kamran Sedig

Western University, London, Ontario, Canada

### Bachelor of Science, Honours, Computer Science and Visual Arts: 2013

Western University, London, Ontario, Canada

### Ontario Secondary School Diploma: 2009

Mother Teresa Catholic High School, London, Ontario, Canada

## Publications

1. **Demelo, J.,** Parsons, P., & Sedig, K. (2017). Ontology-driven search and triage: Design of a web-based visual interface for MEDLINE. *JMIR medical informatics*, 5(1), e4.
2. **Demelo, J.,** & Sedig, K. (2021). Forming Cognitive Maps of Ontologies Using Interactive Visualizations. *Multimodal Technologies and Interaction*, 5(1), 2.
3. **Demelo, J.,** & Sedig, K. (2021). Design of Generalized Search Interfaces for Health Informatics. *Information*, 12(8):317. Featured Article of Special Issue: The Digital Health New Era: Where We Stand and the Challenges.
4. (Prepared for Submission to Publisher) **Demelo, J.,** & Sedig, K. (2022). Searching and Triaging Large Document Sets: An Ontology-Supported Visual Analytics Approach.
5. (Prepared for Submission to Publisher) **Demelo, J.,** & Sedig, K. (2022). Evolutionary Design of Search and Triage Interfaces for Large Document Sets: A Formative Assessment.

## Talks and Posters

1. UWORCS2021: From Theory to Design: Using Interactive Visualization Tools to Support Complex Learning of Ontological Space
2. Topic Survey Proposal: Searching and Triaging Large Document Sets: Investigating the Design of Ontology-supported Interfaces
3. Visual Analytics in Biomedical Applications – 1<sup>st</sup> Place Poster 2017 – Western Science Department
4. Extracurricular Talk: On Lexical Semantics and Computational Lexical Semantics
5. Extracurricular Talk: On Assessing the Predictive Capabilities of Machine Learning Classification for NHL Game Results

## **Teaching**

### **Course Design Support & Management**

Object-Oriented Design and Analysis CS3307A – 2014/2015, 2015/2016, 2016/2017.

Human-Computer Interaction CS4474(A/B)/CS9552(A/B) – 2016/2017, 2017/2018, 2018/2019, 2019/2020, 2020/2021.

### **TA Positions**

Object-Oriented Design and Analysis CS3307A with Dr. Nazim Madhavji – 2014/2015, 2015/2016, 2016/2017.

Object-Oriented Design and Analysis CS3307A with Dr. Mike Katchabaw – 2017/2018, 2018/2019, 2019/2020.

Introduction to Medical Computing CS2124B/CS2125G with Dr. Stephen Watt – 2014/2015.

Introduction to Medical Computing CS2124B/CS2125G with Dr. Charles Ling – 2015/2016.

Approachable Apps: A Gentle Introduction to Programming CS1046B with Dr. Laura Reid – 2016/2017.

Human-Computer Interaction CS4474(A/B)/CS9552(A/B) with Dr. Kamran Sedig – 2016/2017, 2017/2018, 2018/2019, 2019/2020, 2020/2021.

## **Awards and Grants**

UWO Graduate Student Teaching Award Nomination – 2018

Google Summer of Code Award – 2013

Jenny Donald Memorial Award in Visual Arts – 2011

The Western Scholarship of Distinction – 2009

Aiming for the Top Tuition Scholarship – 2009

## **Professional and Organizational Activities**

Member of Insight Lab – 2013-2021

Member of Society of Graduate Students – 2013-2021

Mozilla Thunderbird Developer of Google Summer of Code – 2013

Review Board Developer of University Capstone Open Source Projects – 2013

Member of Computer Science Undergraduate Society – 2009-2013

Department Representative of UWOSC Board of First Years – 2009