5-23-2017 12:00 AM

Contributions of Continuous Max-Flow Theory to Medical Image Processing

John SH Baxter, *The University of Western Ontario*

Supervisor: Dr. Terry Peters, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Biomedical Engineering

Abstract

Discrete graph cuts and continuous max-flow theory have created a paradigm shift in many areas of medical image processing. As previous methods limited themselves to analytically solvable optimization problems or guaranteed only local optimizability to increasingly complex and non-convex functionals, current methods based now rely on describing an optimization problem in a series of general yet simple functionals with a global, but non-analytic, solution algorithms. This has been increasingly spurred on by the availability of these general-purpose algorithms in an open-source context. Thus, graph-cuts and max-flow have changed every aspect of medical image processing from reconstruction to enhancement to segmentation and registration.

To wax philosophical, continuous max-flow theory in particular has the potential to bring a high degree of mathematical elegance to the field, bridging the conceptual gap between the discrete and continuous domains in which we describe different imaging problems, properties and processes. In Chapter 1, we use the notion of *infinitely dense and infinitely densely connected* graphs to transfer between the discrete and continuous domains, which has a certain sense of mathematical pedantry to it, but the resulting variational energy equations have a sense of elegance and charm. As any application of the principle of duality, the variational equations have an enigmatic side that can only be decoded with time and patience.

The goal of this thesis is to show the contributions of max-flow theory through image enhancement and segmentation, increasing incorporation of topological considerations and increasing the role played by user knowledge and interactivity. These methods will be rigorously grounded in calculus of variations, guaranteeing fuzzy optimality and providing multiple solution approaches to addressing each individual problem.

Keywords: optimization-based segmentation, image enhancement, variational optimization, convex optimization

In effect, description is to the object what the proposition is to the representation it expresses: its arrangement in a series, elements succeeding elements.

MICHEL FOUCAULT
THE ORDER OF THINGS

Acknowledgments

Thank you to Drs. Terry Peters, Elvis Chen, Roy Eagleson, and Sandrine de Ribaupierre. Terry, your lab has offered me countless opportunities to explore my ideas, medical image processing, and image-guided interventions, and for that, I am eternally grateful. It has been a pleasure to spend these years at Robarts and in the VASST lab. To Elvis, Roy, and Sandrine, your advice has been invaluable and your patience inexhaustible.

Thank you to Jonathan McLeod, Golafsoun Ameri, Zahra Hosseini and Drs. Kamyar Abhari, Eli Gibson, Jiro Inoue, Martin Rajchl, and Ali Khan. It's been awesome working with you guys and I wish you all the best in the future. Thank you to John Moore for all the philosophical discussions and life advice, and to Chris Wedlake for the technical support in the little time between jokes.

Lastly, I must thank my family, especially my mother and father, for their support and to Nicholas-Conrad Rheault, for always being there for me.

Contents

Abstract ii

Acknowledgments iv

List of Figures x

List of Algorithms xiv

List of Tables xv

List of Appendices xvi

List of Abbreviations, Symbols, and Nomenclature xvii

Preface xviii

1 Introduction 1
- 1.1 Images and Image Labelling Problems . . . 1
  - 1.1.1 Partitioning and Segmentation Problems . . . 2
  - 1.1.2 Image Enhancement, Restoration and Filtering Problems . . . 3
  - 1.1.3 Image Registration Problems . . . 4
- 1.2 Markov Random Fields and Gibbs Distributions . . . 4
  - 1.2.1 Unary and Binary Energies . . . 6
- 1.3 Discrete Graph-Cuts . . . 7
  - 1.3.1 Discrete Ising Model . . . 7
  - 1.3.2 Applications of Graph-Cuts in Medical Image Processing . . . 8
- 1.4 Continuous Max-Flow Theory . . . 9
  - 1.4.1 From Discrete Graph-Cuts to Continuous Max-Flow . . . 10
  - 1.4.2 Duality and Convex Variational Optimization . . . 14
  - 1.4.3 Early Approaches to Max-Flow Optimization . . . 15
  - 1.4.4 Augmented Lagrangian Multipliers . . . 16
  - 1.4.5 Proximal Bregman Projections . . . 17
  - 1.4.6 Applications of Continuous Max-Flow in Medical Image Processing . . . 19
- 1.5 Contrasting Graph-Cuts and Continuous Max-Flow . . . 19
- 1.6 Thesis Outline . . . 20

2 Cyclic Continuous Max-Flow Image Enhancement 22

2.1 Introduction . . . 22
2.2 Previous Graph-Cuts and Max-Flow Image Enhancement Approaches . . . 24
2.2.1 Convex Max-Flow Image Restoration . . . 24
2.2.2 Discrete Potts Model . . . 25
2.2.3 Continuous Potts Model . . . 26
2.2.4 Discrete Ishikawa Model . . . 28
2.2.5 Continuous Ishikawa Model . . . 29
2.3 Susceptibility and MRI Phase Processing . . . 32
2.3.1 Homodyne Filtering Paradigm . . . 33
2.3.2 Phase Unwrapping Paradigm . . . 34
2.4 Cyclic Continuous Max-Flow Formulation . . . 36
2.5 Cyclic Continuous Max-Flow Algorithm . . . 37
2.6 Cyclic Continuous Max-Flow Synthetic Validation . . . 38
2.6.1 Images . . . 39
2.6.2 Methods . . . 40
2.6.3 Results . . . 40
2.7 Cyclic Continuous Max-Flow in MRI Phase Processing . . . 42
2.7.1 Images . . . 42
2.7.2 Methods . . . 43
2.7.3 Single Channel Qualitative Results . . . 44
2.7.4 Channel Combined Qualitative Results . . . 44
2.7.5 Channel Combined Quantitative Results . . . 48
2.8 Discussion . . . 49
2.8.1 Future Work . . . 49

3 Hierarchical Continuous Max-Flow Image Segmentation 51
3.1 Introduction . . . 51
3.2 Label Orderings and Hierarchical Topologies . . . 53
3.2.1 Label Ordering Operators . . . 55
3.3 Previous Approaches to Hierarchical Topologies . . . 56
3.3.1 Graph-Cuts and the *h*-Fusion Algorithm . . . 56
3.3.2 Hard-Coded Hierarchies in Continuous Max-Flow . . . 57
3.3.3 Gestalt Computer Vision . . . 57
3.4 Hierarchical Continuous Max-Flow Formulation . . . 58
3.5 Hierarchical Continuous Max-Flow Solution Algorithms . . . 59
3.6 Hierarchical Continuous Max-Flow in Brain Tissue Segmentation . . . 63
3.6.1 MICCAI 2012 OASIS Images . . . 63
3.6.2 MICCAI 2012 OASIS Methods . . . 64
3.6.3 MICCAI 2012 OASIS Results . . . 67
3.6.4 MICCAI 2013 MRBrainS Images . . . 71
3.6.5 MICCAI 2013 MRBrainS Methods . . . 71
3.6.6 MICCAI 2013 MRBrainS Results . . . 73
3.7 Discussion . . . 73
3.7.1 Future Work . . . 75

4 Optimization Based Interactive Segmentation with Anatomical Knowledge 76
4.1 Introduction . . . 76
4.2 Philosophy of Segmentation . . . 77
4.2.1 Semiotics in Segmentation . . . 78
4.2.2 Input Signs . . . 79
4.2.3 Output Signs . . . 82
4.2.4 Heterogeneity and Sign Graphs . . . 83
4.2.5 Philosophical Call to Action . . . 87
4.3 Interactive Segmentation Interface . . . 88
4.3.1 Hierarchical max-flow segmentation . . . 88
4.3.2 Definition of Cost Terms . . . 89
4.3.3 Plane Selection . . . 90
4.3.4 Interface Description . . . 90
4.4 Example Applications of Interactive Segmentation . . . 92
4.4.1 Cardiac Segmentation . . . 92
4.4.2 Neonatal Cranial MRI Segmentation . . . 94
4.5 Automatic Hierarchy Refinement . . . 96
4.6 Discussion . . . 97
4.6.1 Clinical Integration . . . 98
4.7 Future Work . . . 99

5 Directed Acyclic Graph Continuous Max-Flow Image Segmentation 101
5.1 Introduction . . . 101
5.2 Previous Graph-Cuts and Max-Flow Approaches with More Complex Topologies 103
5.2.1 Submodular Graph Construction . . . 103
5.2.2 Generalized Ordering Constraints in Continuous Min-Cut . . . 104
5.3 Directed Acyclic Graph Max-Flow Formulation . . . 105
5.3.1 Arbitrary Region Regularization . . . 106
5.4 Directed Acyclic Graph Max-Flow Algorithm . . . 107
5.5 Validation . . . 111
5.5.1 Synthetic Image: Venn Diagram . . . 111
5.5.2 Medical Images - Brain Tissue Segmentation . . . 112
5.5.3 Natural Images: Scene Decomposition . . . 116
5.5.4 Natural + Synthetic Images: Hue Reconstruction . . . 119
5.6 Discussion . . . 122
5.6.1 Future Work . . . 122

6 Shape Complexes in Max-Flow Image Segmentation 124
6.1 Introduction . . . 124
6.2 Prior Approaches to Shape Information in Segmentation . . . 125
6.3 Prior Work on Shape Information in Max-Flow . . . 126
6.3.1 Discrete Domain . . . 126
6.3.2 Continuous Domain . . . 127
6.4 Shape Complexes . . . 129
6.5 Shape Complexes Implementation . . . 129

6.6 Validation . . . 130
6.6.1 Synthetic Image Segmentation . . . 131
6.6.2 Ultrasound Vessel Segmentation . . . 135
6.6.3 Cardiac Valve Segmentation from Ultrasound . . . 137
6.6.4 Atrial Wall Segmentation from Cardiac CT . . . 140
6.7 Discussion . . . 143
6.7.1 Future Work . . . 144

7 Conclusions **146**
7.1 Recurrent Themes . . . 147
7.1.1 The Principle of Topology . . . 147
7.1.2 The Principle of Interactivity . . . 148
7.2 Unaddressed Aspects of Continuous Max-Flow . . . 148
7.2.1 Max-Flow and Graph-Cuts Propagated Level-Sets . . . 149
7.2.2 Multi-Resolution Deformable Registration . . . 149
7.2.3 Other Regularization Functions . . . 150
7.2.4 Existential Priors . . . 151
7.3 The Future of Continuous Max-Flow in Medical Image Processing . . . 151

A Use of the Principle of Duality and Derivations of Algorithms **153**
A.1 CCMF Algorithm Derivation . . . 153
A.1.1 Primal and Primal-Dual Models . . . 153
A.1.2 Equivalence to Dual Model . . . 154
A.1.3 Augmented Lagrangian . . . 155
A.1.4 Proximal Bregman . . . 156
A.2 HMF Algorithm Derivation . . . 158
A.2.1 Primal and Primal-Dual Models . . . 158
A.2.2 Equivalence to Dual Model . . . 158
A.2.3 Augmented Lagrangian . . . 160
A.2.4 Proximal Bregman . . . 163
A.3 DAGMF Algorithm Derivation . . . 166
A.3.1 Primal and Primal-Dual Models . . . 166
A.3.2 Equivalence to Dual Model . . . 167
A.3.3 Augmented Lagrangian . . . 168
A.3.4 Proximal Bregman . . . 172
A.4 Discretization and Memory Consumption . . . 175
A.5 p-Norm Regularization Terms . . . 175
A.5.1 Directional Regularization . . . 177
A.5.2 Geodesic Star Convexity . . . 178

B Combinatorial and Complexity Analysis of Label Ordering Structures **180**
B.1 Combinatorics of Unconstrained Hierarchies . . . 180
B.2 Grouping Graph . . . 181
B.3 NP-Hardness of Maximum Hierarchy Selection . . . 181

C Ethics Approvals for Patient Images **182**
C.1 Magnetic Resonance Imaging of Multiple Sclerosis at 7 Tesla 183
C.2 new technologies in the management of post-haemorrhagic hydrocephalus in preterm infants 184
C.3 Image-Guidance in Cardiac Interventions 185
C.4 Anatomical measurements of the heart for radiofrequency cathetar ablation . . 186

D Copyright Transfers and Reprint Permissions **187**

Bibliography **202**

Curriculum Vitae **219**

List of Figures

1.1 Example of the difference between segmentation and partitioning. Segmentation problems may, in general, involve regions that overlap or may not cover the image. Partitioning problems require non-overlapping regions which altogether cover the entire image. 2
1.2 Example of a 4-connected graph construction and cut representing a binary image partition problem. 8
1.3 MRF as the spacing between nodes, Δr, approaches 0. 10
1.4 Unit Spheres Under Various Norms 12
1.5 MRF as the number of elements in any given $\mathcal{N}(x)$, approaches infinity. 12
1.6 Example of metrification artifacts demonstrated. The graph cut segmentation, which used a 4-connected neighbourhood, is unnecessarily blocky whereas the continuous max-flow solution is more natural in appearance. 19

2.1 MRI phase image, background phase and local phase shift map. 23
2.2 Illustrative example of non-linear behaviour in cyclic range topologies. 23
2.3 Example graph used in the Potts model with labels $\mathbb{L} = \{A, B, C\}$ 27
2.4 Example graph used in the Ishikawa model with labels $\{L_0, L_1, L_2, L_3\}$. 30
2.5 Example of an non-unwrappable phase image (left) with corresponding demonstration of path dependence (right 35
2.6 Topology with which CCMF indicator functions are equipped 37
2.7 Comparison of phase smoothing using the Potts, Ishikawa, and CCMF models. The Potts model is excessively blocky and the Ishikawa model is error-prone surrounding the phase wraps. 40 phase bins were used in each model. 38
2.8 Phantom experiment gold standard and noisy images. Low-pass filtered results on the noisy image are shown in Fig. 2.9. 39
2.9 Example low-pass filtering results using the Augmented Lagrangian solver (a to d) and the Proximal Bregman solver (e to h). Each pair includes the low-pass filtered image and a difference image between the result and the noise-free phantom image, Fig. 2.8b. 41
2.10 Error reduction with varying regularization weight 42
2.11 Single channel cranial MR image including magnitude (a) and phase (b) components. 43
2.12 Local phase shift maps on the single channel MR image shown in Figure 2.11. . 45
2.13 Local Phase Shift Map computed via channel combination of single channel local phase shift maps such as those presented in Figure 2.12. 46

2.14 Residual phase wrapping artifacts present in channel combined images created using phase unwrapping paradigm. 47
2.15 CNR performance over echo time 48
2.16 Confidence intervals on mean CNR difference between methods 48

3.1 Example of anatomical knowledge in the form of *part/whole relationships* expressing the decomposition of the cerebral cortex into hemispheres and then into individual lobes. 52
3.2 Example label ordering from Figure 3.1 shown in diagrammatic form. 53
3.3 Potts and Ishikawa models in terms of label ordering 55
3.4 Example of the different label ordering operators 56
3.5 Label orderings used in cardiac segmentation by Rajchl *et al.* [149, 150] 57
3.6 Example of transforming a hierarchical label ordering into a flow network. . . . 60
3.7 Pipeline used in multi-atlas brain tissue segmentation. 65
3.8 Hierarchical label ordering used in segmentation of the OASIS database 66
3.9 Best and worst case results for the OASIS database (top row: best case T1w image, gold standard, JLF+HMF, worst case T1w image, gold standard, JLF+IM+HMF; bottom row: enlarged ROIs.) 69
3.10 Hierarchical label ordering used in segmentation of the OASIS database 72
3.11 Best case results for the MRBrainS database (top row: best case T1w image, T1IR image, T2FLAIR image, gold standard, JLF+IM+HMF. bottom row: enlarged ROI.) 74
3.12 Worst case results for the MRBrainS database (top row: best case T1w image, T1IR image, T2FLAIR image, gold standard, JLF+IM+HMF. bottom row: enlarged ROI.) 74

4.1 Segmentation Interactivity Spectrum 77
4.2 Example Sign Graphs in Segmentation Interfaces 84
4.3 Segmentation interface with user seeds before segmentation (a) and after segmentation (b). The hierarchy definition widget (bottom left corner of (a) and (b)) is shown enlarged in (c). 91
4.4 Cardiac segmentation with underlying (a) CT, (d) MRA, and (g) TEE. Manual segmentations are in (b), (e), and (h) respectively, and interactive segmentation results in (c), (f), and (i). 93
4.5 Neonatal Ventricle Segmentation with (a) the MR, (b) the manual segmentation, and (c) interactive segmentation results. (d) shows surface renderings of both the fully manual (left) and interactive (right) segmentation results. 94
4.6 Pathological Neonatal Ventricle Segmentation with (a) the MR, (b) the manual segmentation, and (c) interactive segmentation results. (d) shows surface renderings of both the fully manual (left) and interactive (right) segmentation results. 95
4.7 Hierarchies used in (a) healthy and (b) pathological neonatal ventricle segmentation. 95

4.8 Example of automatic hierarchy segmentation. (a) is the original hierarchy reproduced from [150] and (b) the optimized version, (c) an LGE-MRI with (d) manual segmentation and (e) interactive segmentations results. 96

5.1 Schematic vessel bifurcation 102
5.2 Two possible hierarchical label orderings for the vessel bifurcation segmentation example shown in Figure 5.1. 102
5.3 DAG for segmentation into labels $\mathbb{L} = \{A, B, C, D, E\}$ in which label groups $\mathbb{G} = \{AB, BC, CD\}$ are regularized. Note that this would be impossible in a hierarchical model since the regularization groups conflict with each other. Figure 5.3a shows the intermediate multi-edged, unweighted DAG. Figure 5.3b shows this DAG with weights explicitly recorded rather than through multiplicity which is used by the solution algorithms. 106
5.4 Synthetic image (a) polluted with noise (b) and reconstructed using a Potts model (c), DAGMF (d) and HMF models with either the red square (e) or green circle (f) regularized.Weighted DSC is given for each segmentation. . . . 111
5.5 Directed acyclic graph and weights used for DAGMF segmentation shown in Figure 5.4. The nodes *circle* and *square* denote the labels associated with the union of green with yellow and red with yellow respectively. 112
5.6 DAG representing the brain tissue segmentation problem in Fig. 5.8. 112
5.7 Bayesian data terms used in Fig. 5.8. 113
5.8 Brain tissue segmentation using DAGMF using data terms in Fig. 5.7 and constant smoothness terms. Note the improvement in the pink subcortical gray matter region. 114
5.9 Segmentation uncertainty (entropy) from Figure 5.8. The Potts model has much higher uncertainty in the background segmentation around the frontal lobe. . . . 115
5.10 Segmentation structure used in scene decomposition into F-front, T-top, B-bottom, L-left, R-right. The color code corresponds to that used in Figure 5.11. 116
5.11 Example outdoor scene segmentation. Accuracy rate is given for each segmentation. The color code for the segmented images are shown in Figure 5.10. . . . 117
5.12 Example scene segmentation from the Stanford indoor dataset [42]. DTO refers to the "data term only" method. Accuracy rates are given for each segmentation. Label orderings used in the first and second HMF segmentation are shown in (f) and (i) respectively. The color code for the segmented images are shown in Figure 5.10 and in (f) and (i). 118
5.13 Example DAG for hue reconstruction with $N = 9$ discrete hues and a truncated linear model of width $M = 3$. Although not shown, the weight of each edge on the top level is 1, and $1/M$ on the bottom layer. 120
5.14 Hue reconstruction on synthetic image with corresponding normalized hue error. 120
5.15 Example hue reconstruction on natural images with DAGMF model ($N = 36, M = 16$). 121

6.1 Simple star convex objects with vantage points indicated with an '$\mathbf{X}$'. 126

6.2 Synthetic image segmentation problem using DAGMF (2e) and DAGMF augmented with shape complexes (2f) according to the label ordering in (2d) with α referring to the level of regularization. (* a simple star convexity constraint is applied to this label.) Any overlap between segmentations can cause false colors, e.g. green occurs when the result is 50% exterior (cyan) and 50% interior (yellow). The 'X' marks the vantage point for the simple star convexity constraint. 134
6.3 Quantitative segmentation results for each region based on regularization strength. The Dice similarity coefficients are shown on a logarithmic scale approaching 100% DSC. 135
6.4 Venn diagram segmentation with and without shape complexes. The label ordering is given in Figure 4c (* a simple star convexity constraint is applied to this label) with the vantage point for the shape complex was the centroid of the region. Similar to Figure 6.2, any overlap between segmentations can cause false colors. 136
6.5 Vessel segmentation in ultrasound with and without shape complexes. The label ordering is given in Figure 5c (* a simple star convexity constraint is applied to this label) with the vantage point for the shape complex is marked with an '**X**'. Similar to Figure 6.2, any overlap between segmentations can cause false colors. 137
6.6 Quantitative results for the segmentation problem shown in Fig. 6.5 varying regularization and bias parameters. Blue indicates DSC $\approx$ 0 and yellow indicates DSC $\approx$ 1 as shown in Fig. 6g. 138
6.7 Synthetic valve annulus segmentation. The label K indicates the background (in cyan), TW and BW indicates the top and bottom walls respectively (in magenta), TB and BB indicate the top and bottom blood pools respectively (in yellow), and V indicates the valve annulus (in green). 139
6.8 Mitral valve labelling using trans-esophageal ultrasound images. The model (Figure 6.7a - previous figure) includes label K indicates the background (in cyan), TW and BW indicates the top and bottom walls respectively (in magenta), TB and BB indicate the top and bottom blood pools respectively (in yellow), and V indicates the valve annulus. 139
6.9 Atrial wall segmentation DAGMF augmented with shape complexes with α representing the regularization strength. (* a simple star convexity constraint is applied to this label.) 141
6.10 Best and worst case atrial wall segmentation results. The atrial blood pool is shown in magenta and the atrial wall in cyan. The black regions are user-provided seed points for the atrial blood label. 142

List of Algorithms

1.1 Split-Merge solution algorithm for binary max-flow proposed by Pock and Chambolle [32, 137] . . . 16
1.2 Augmented Lagrangian solution algorithm for binary max-flow proposed by Yuan *et al.* [192] . . . 17
1.3 Proximal Bregman solution algorithm for binary max-flow proposed by Bae *et al.* [13] . . . 18

2.1 Split-Merge solution algorithm proposed by Pock and Chambolle [32, 137] for convex image restoration problems . . . 25
2.2 Augmented Lagrangian solution algorithm proposed by Yuan *et al.* [193] for the continuous Potts model . . . 27
2.3 Proximal Bregman solution algorithm proposed by Baxter *et al.* [18] for the continuous Potts model . . . 27
2.4 Augmented Lagrangian solution algorithm proposed by Bae *et al.* [14] for the continuous Ishikawa model . . . 31
2.5 Proximal Bregman solution algorithm proposed by Baxter *et al.* [18] for the continuous Ishikawa model . . . 31
2.6 Augmented Lagrangian solution algorithm for the CCMF functional in terms of indicator functions. . . . 37
2.7 Proximal Bregman solution algorithm for the CCMF functional in terms of indicator functions. . . . 38

3.1 Augmented Lagrangian solution algorithm for the HMF functional. . . . 61
3.2 Proximal Bregman solution algorithm for the HMF functional. . . . 62

5.1 Augmented Lagrangian solution algorithm for the DAGMF functional. . . . 108
5.2 **InitializeSolution()** subroutine in Algorithm 5.1. . . . 108
5.3 **UpdateFlows()** subroutine in Algorithm 5.1. . . . 109
5.4 Proximal Bregman solution algorithm for the DAGMF functional. . . . 110

6.1 Augmented Lagrangian solution algorithm for Eq. (6.3). . . . 130
6.2 **InitializeSolution()** subroutine in Algorithm 6.1. . . . 131
6.3 **UpdateFlows()** subroutine in Algorithm 6.1. . . . 132
6.4 Proximal Bregman algorithm for Eq. (6.3). . . . 133

List of Tables

2.1 Computation times for 2D image slices at varying echo time 44
2.2 Three-way ANOVA Table of factors affecting CNR 48

3.1 Segmentation Results - OASIS: *significantly better metrics ($p \leq 0.05$ after Bonferroni correction) between HMF/Potts pairs are shown in bold and significantly better metrics ($p \leq 0.05$) with/without the intensity model are denoted with an asterix for JLF.* 69
3.2 Comparison Results - OASIS *DSC*: *significant difference to JLF+IM+HMF ($p \leq 0.05$ after Bonferroni correction) is shown in bold* 70
3.3 Comparison Results - OASIS *AVD*: *significant difference to JLF+IM+HMF ($p \leq 0.05$ after Bonferroni correction) is shown in bold* 70
3.4 Comparison Results - OASIS *MHD*: *significant difference to JLF+IM+HMF ($p \leq 0.05$ after Bonferroni correction) is shown in bold* 70
3.5 Segmentation Results - MRBrainS: *significant difference to JLF+IM+HMF ($p \leq 0.05$ after Bonferroni correction) is shown in bold* 73

4.1 Classification with examples of different input sign types 80
4.2 Examples of different output sign types 82
4.3 Cardiac Segmentation Numerical Results. Results shown in **bold** indicate metrics that are common across SEGUE, manual segmentations, and [149] which can be used as a reference for comparison 92
4.4 Scar Tissue Segmentation Results. Results shown in **bold** indicate metrics that are common across SEGUE, manual segmentations, and [150] which can be used as a reference for comparison 97

5.1 Dice coefficient for segmentations in Fig. 5.8. The results for the subcortical gray matter are shown in **bold** which reflect the quantitative improvement from using a more nuanced regularization model with DAGMF. 114
5.2 Value of the constant regularization terms used in the various max-flow models. 117
5.3 Accuracy rates for segmentations in the Stanford indoor dataset such as that shown in Fig. 5.12. DTO refers to the "data term only" method. The results shown in **bold** represent those statistically significantly different from the DTO method under a paired t-test with Bonferroni correction. 118

6.1 Mean distance error results for the blood pool and whole atrium labels. These are reflective of the errors seen on the inner and outer boundary of the atrial wall label. 141

List of Appendices

Appendix A Use of the Principle of Duality and Derivations of Algorithms 153
Appendix B Combinatorial and Complexity Analysis of Label Ordering Structures . . . 180
Appendix C Ethics Approvals for Patient Images 182
Appendix D Copyright Transfers and Reprint Permissions 187

List of Abbreviations, Symbols, and Nomenclature

Mathematical and Computer Graphics/Vision Terminology

$A^{-\top}$	Inverse transpose of a square matrix (i.e. $A^{-\top} = (A^{-1})^{\top} = (A^{\top})^{-1}$)
2^S	The power set of S, or set of all subsets of S
2D	Two dimensional
3D	Three dimensional
Ω	Spatial domain
ANR	Artefact-to-noise ratio
CNR	Contrast-to-noise ratio
DAG	Directed Acyclic Graph
GPGPU	General-Purpose Graphics Processing Unit
GPU	Graphics Processing Unit
MAP	Maximum *A Posteriori* Probability
MRF	Markov Random Field
*N*D	Arbitrary N dimensional
NP	Non-Deterministic Polynomial
NPC	Non-Deterministic Polynomial Complete
PGM	Probabilistic Graphical Model

Medical Imaging Terminology

CT	Computed Tomography
CTA	Computed Tomographic Angiography
MR	Magnetic Resonance
MRI	Magnetic Resonance Imaging
QSM	Quantitative Susceptibility Mapping
ROI	Region of Interest
STI	Susceptibility Tensor Imaging
SWI	Susceptibility Weighted Imaging
T1	Longitudinal Relaxation Time
T1IR	T1 Inversion Recovery
T2	Transverse Relaxation Time
T2FLAIR	T2 Fluid Attenuated Inversion Recovery
US	Ultrasound

Terms Introduced by this Thesis

$\mathbb{I}$	Image Intensity Range
$\mathbb{L}$	Set of labels forming a partition in a segmentation problem
$\mathbb{G}$	Set of superlabels requiring separate regularization
CCMF	Cyclic Continuous Max-Flow
DAGMF	Directed Acyclic Graphical Continuous Max-Flow
HMF	Hierarchical Continuous Max-Flow

Preface

To the reader, this thesis focuses on a relatively recent technical and theoretical development in image processing: continuous max-flow. Continuous max-flow theory is a continuous analogue to graph-cut theory which has received wide attention from computer scientists since the 1950s and computer vision scientists since the 2000s. The continuous max-flow community however grew out of the convergence of many different mathematical theories, notably graph-cuts, variational optimization and level set optimization, all of which tended towards conflicting notation.

I have attempted to reconcile the notation used by the many authors in the continuous max-flow community, to express discrete graph-cuts in an analogous notation, and to enforce some notion of stylistic integrity. If I do not use the reader's preferred notation in any particular section, I apologize. My goal is only to emphasize the growth of the field and the many similarities contained within.

On a related note, the literature contains many a phrase like *"Optimize f(x) which can be done analytically"* and leaves it at that. Wherever possible, I have replaced this with the analytic solution. I present pseudocode for many algorithms in a style that I hope is readily understandable and implementable. Once again, the goal is clarity and completeness, not obfuscation. To make this possible, one must assume a certain degree of background knowledge of the reader, such as the standard notation in formal logic and set theory.

I hope you find this thesis to be informative and hopefully entertaining.

Regards,

John S.H. Baxter

Chapter 1

Introduction

1.1 Images and Image Labelling Problems

What is an image? To begin on a formal note, an image is an *intensity function*, $I(x) : \Omega \to \mathbb{I}$, which maps a spatial location $x \in \Omega$, to an intensity $i \in \mathbb{I}$, often a real vector of pre-defined dimensionality (i.e. $\mathbb{I} = \mathbb{R}^n$). How one understands or models this intensity function and its underlying spatial domain Ω has a large effect on what form image analysis problems can take.

Changing how Ω is interpreted changes what mathematical formalisms can be expressed. For example, one could consider a photograph or *natural image* to be a functional of a finite 2D lattice of pixels (that is, Ω is a finite set). Because the domain Ω is a discrete lattice, the traditional concept of derivatives and curvature do not exist, requiring infinitesimally small but non-zero distances. (Approximations or analogues such as *calculus of finite differences* exist and somewhat bridge the gap in formalisms, but still illustrate that differences in the topology of the spatial domain have an effect on how the labelling problem is formulated.) If instead the natural image is interpreted as having a 2D continuum domain Ω, rather than a lattice, these traditional notions based on infinitesimally small distances are restored, encouraging us to use a different collection of theoretical tools to address labelling problems using this image as input.

This thesis will concentrate on low-level labelling (or just *labelling*) problems. These are image analysis problems that can be expressed in the form of a *labelling function*, $u(x)$, which maps spatial locations on the image, $x \in \Omega$, to one of:

- an element of a predefined finite set of values (i.e. image partitioning),
- a subset of a predefined finite set of values (i.e. image segmentation),
- a continuous interval of values (i.e. image enhancement, restoration, filtering, etc...), or
- a location in another spatial domain (i.e. image registration).

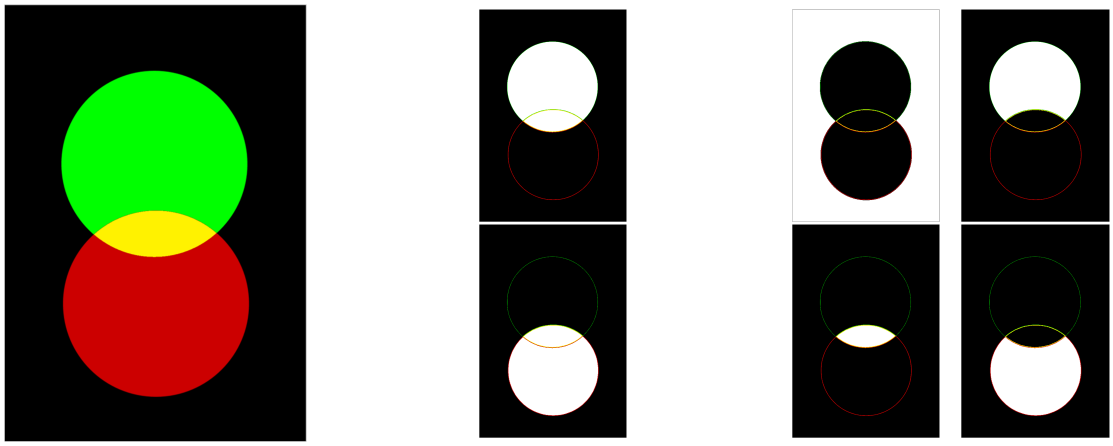

(a) Image featuring two overlapping circles.

(b) Segmentation with segmented objects shown in white.

(c) Partition with partition regions shown in white.

Figure 1.1: Example of the difference between segmentation and partitioning. Segmentation problems may, in general, involve regions that overlap or may not cover the image. Partitioning problems require non-overlapping regions which altogether cover the entire image.

1.1.1 Partitioning and Segmentation Problems

The reader may confused over why there is a difference between a partitioning problem and a segmentation problem. In fact, a partitioning problem can be expressed as a subset of segmentation problems in which the subsets mapped to always contain only one element. In some segmentation problems, this may not be the case, especially if objects overlap or if one object is described as spatially containing another such as a vessel containing a lumen. Although this distinction may appear pedantic, it will become important later on.

Figure 1.1 shows an example of the difference between segmentation and partitioning problems. The key difference is that segmentation problems are more general; they do not require the regions being segmented to be disjoint or to cover the entire image. Partitions on the other hand, must cover the entire image and cannot be overlapping.

Often, the labelling function will be described in terms of an indicator function. In those cases, the indicator function corresponding to a particular label, L, will be written as $u_L(x)$. In a segmentation or partitioning problem, these indicator functions can be defined as:

$$u_L(x) = \begin{cases} 1, & \text{if } x \in L \\ 0, & \textit{else} \end{cases} \tag{1.1}$$

which have the property that if a subset of L's, $\mathbb{L}$, forms a partition:

$$\sum_{\forall L \in \mathbb{L}} u_L(x) = 1. \tag{1.2}$$

Fuzzy indicator functions are a less constrained type of indicator function that take on a value in the interval $[0, 1]$ with higher values being indicative of membership in the corresponding label. Fuzzy indicator functions have the same partitioning property described in Equation 1.2. The subscript notation on the indicator functions was intended to be similar to that commonly used for indexing vector valued functions, although in this case, the indices are not integers but labels.

Partitioning problems are obviously a constrained class of segmentation problem. However, any segmentation problem can be converted into a partitioning problem by considering the power set 2^L. This is an important technique as many of the elements of 2^L may not be feasible and their indicator function is constrained to equal 0.

1.1.2 Image Enhancement, Restoration and Filtering Problems

Image enhancement, restoration and filtering problems involve assigning locations with a new intensity function. Often, image restoration and enhancement problems use the same intensity function range as the original image. A useful representation of these problems involves approximating them with a partitioning or segmentation problem once. That is, the labelling function (i.e. the processed image), $u(x)$, can piecewise-constant approximated as:

$$u(x) \approx \sum_{i=1}^{N} \tilde{u}_i \, u_i(x) \tag{1.3}$$

where $\tilde{u}_i$ is an intensity and $u_i(x)$ is an indicator function for the partitioning problem expressing that location x has a processed intensity of approximately u_i. That is, $u_i(x) = 1$ if and only if $u(x) \approx u_i$. (If fuzzy indicator functions are used, this becomes a linear approximation.) This approximation can become arbitrarily good given a larger and larger number of indicator functions with closer and closer values of u_i. An important result of this is that a function of a location and its label value can also be approximated:

$$f(u(x), x) \approx \sum_{i=1}^{N} f(u_i, x) \, u_i(x) \tag{1.4}$$

noting that the term $f(u_i, x)$ no longer depends on $u(x)$. This is akin to a partitioning problem in which the labels represent regions with an intensity proximal to a pre-defined value. Thus, any image enhancement problem can be approximated through an image partitioning problem and this representation is crucial for using partitioning and segmentation approaches in enhancement, restoration and filtering.

1.1.3 Image Registration Problems

Image registration involves associating every spatial location in an image to a spatial location in another image or atlas. In the case of *rigid* and *affine* registration problems, this mapping has a very particular form and thus can be represented more readily as a constant transform. In that case, it is rare to see it represented as an image labelling problem. *Deformable registration*, on the other hand, is often represented as a labelling through the use of a *deformation field*. That is, the spatial location being mapped to is equal to the original spatial location with some translational offset, or $y = x + d(x)$ where $y, x \in \Omega$ [179].

Maintz and Viergever [116, 179] provided a popular taxonomy for medical image registration techniques composed of the interrelated criteria:

1. Dimensionality of the image(s) involved,
2. Modalities of the images(s) involved,
3. Subjects involved (intra-subject, inter-subject, atlas),
4. Objects-of-interest involved,
5. Degree of interaction (interactive, semi-automatic, automatic),
6. Nature of the registration basis (extrinsic fiducials, intrinsic landmarks, segmentation-based, image-based, calibrated co-ordinate system based),
7. Nature of the transformation (rigid, affine, projective, spline-interpolated, fully deformable),
8. Domain of the transformation (local or global), and
9. Optimization procedure involved in determining said transformation.

The first four criteria roughly align to clinical context of the registration problem, the next five to what constraints are placed on the labelling, and the last pertains to how such a labelling is determined. A comprehensive review of medical image registration approaches is beyond the scope of this thesis, but the interested reader can consult the recent review article by Viergever *et al.* [179] or the review article specifically concerning deformable registration by Sotiras *et al.* [164].

1.2 Markov Random Fields and Gibbs Distributions

One invaluable theoretical tool for labelling problems in general is the *Markov Random Field* (MRF). An MRF is an undirected probabilistic graphical model, that is, a probability distribution concerning the values of a collection of variables represented as nodes in a graph, G. The edges in a graph define the three *Markov properties* [101]:

1. *Pairwise Markov Property:* The values of any two non-adjacent variables ($v, u \in G$ such that $w \notin \mathcal{N}_G(V)$) are conditionally independent given all other variables.

2. *Local Markov Property:* The value of any variable ($v \in G$) is conditionally independent of all other variables given its neighbours.
3. *Global Markov Property:* The values of any two disjoint sets of variables ($V, W \in 2^G$) are conditionally independent given a third separating set.

These properties increase in strength, that is, the global Markov property implies the local, which in turn implies the pairwise. In image labelling problems, the goal of using an MRF is to express more desirable labelling configurations as being more probable.

The probability distribution expressed by a MRF must be a *Gibbs distribution*. Specifically, the probability distribution can be expressed in the form:

$$P(u) = \frac{e^{-E(u)}}{Z} \tag{1.5}$$

where $E(u)$ is an energy function and Z is a normalization factor. The energy function maps the state configuration to a real number, with higher numbers representing less probable configurations. Representing an MRF via a Gibbs distribution is important because many MRF's of interest can be specified using a particular constrained set of energy functions, specifically those which can be decomposed as:

$$E(u) = \sum_{\forall V \in \mathrm{cl}(G)} E(u_V) \tag{1.6}$$

where the *clique* operator, $\mathrm{cl}(\cdot)$, is the set of all sets of variables in G such that every variable is adjacent to every other variable. More formally:

$$\mathrm{cl}(G) = \left\{ V \in 2^G \mid \forall (v, w) \in V, w = v \text{ or } w \in \mathcal{N}_G(v) \right\} \tag{1.7}$$

This is called the *clique factorization* of the MRF. Terms in which V has one or two variables are called *unary* and *binary* energies respectively. One particularly important result along these lines is the *Hammersly-Clifford theorem* [68] which states that the MRF has a strictly positive probability for every configuration if and only if it can be expressed via a clique factorization.

So, why are MRF's so important for medical image analysis? The importance of these constructs stem from their ability to encode contextual constraints and regularization in image labelling problems. Clique factorization often provides a comprehensible system of penalizing particular configurations based on solely local properties of the labelling and of the image. Some of these MRFs have fast maximum *a posteriori* probability (MAP) solution procedures given that their clique factorization adheres to particular constraints.

1.2.1 Unary and Binary Energies

Unary energies (often called *data terms*) are important in labelling problems because they directly relate the labelling value at a particular location to local properties such as the spatial location or image intensity. One particularly useful form of data term used in segmentation is the *Bayesian* data term, which can be written as:

$$D_L(x) = -\ln\left(I(x)|x \in L\right) + \ln P(x \in L). \tag{1.8}$$

This type of data term uses the probability distribution encoded in the MRF directly, that is, it defines the probability of a particular region taking on a particular label in the absence of neighbours to affect it. In image enhancement problems, a common data term is the difference between the intensity and the labelling taken to some power:

$$D(u(x), x) = |u(x) - I(x)|^p. \tag{1.9}$$

Being based on a single voxel, data terms can be sensitive to noise in the image.

Binary energies (often called first-order terms or *regularization terms*) control how much adjacent variables effect each other. The reason these are called regularization terms is that they are used to smooth away any overfitting caused by the data terms. A common regularization term in partitioning problems is the uniform term:

$$R(x, y) = \begin{cases} 0, & \text{if } u(x) = u(y) \\ \alpha, & \text{else} \end{cases} \tag{1.10}$$

where α is a positive constant. Other common regularization terms replace the constant α with a positive monotonically decreasing function of the difference in intensity between the two locations:

$$R(x, y) = \begin{cases} 0, & \text{if } u(x) = u(y) \\ f(|I(x) - I(y)|), & \text{else} \end{cases}. \tag{1.11}$$

These terms penalizes variation in the labelling less if it is associated with variation in the image intensity, thus encouraging edges in the labelling to align to edges in the image.

Many complex segmentation, registration and image enhancement problems can be specified as MRF's in which there are only data and regularization terms. That is, all terms associated with *higher order cliques* of size three or greater, have zero value.

1.3 Discrete Graph-Cuts

For MRFs with zero higher order clique energies and non-negative regularization terms taking the form specified in the previous section, graph-cut approaches have been favored due to their speed and optimality guarantees. This is largely because the MRF can be expressed as an edge-weighted graph in which a binary labelling ($u(x) \in \{0, 1\}$) is related to a cut through this graph where a cut is a minimal set of edges that, if removed, disconnect the source and sink nodes. The MAP optimum (i.e. energy minimum) is thus related to the minimum weighted cut.

1.3.1 Discrete Ising Model

The Ising model [79] is arguably the simplest MRF, in which each variable can only take on a single value, either a 0 or a 1, and was designed to model polarization in ferromagnetic materials. This model has great implications to image segmentation in that it can be seen as it could be seen as a representation of a binary segmentation problem, one with a single object-of-interest and the background. The equation for this model is:

$$E(u) = \sum_{x\in\Omega} D(x)u(x) + \sum_{x\in\Omega}\sum_{y\in\mathcal{N}(x)} \frac{R(x,y)}{2}|u(x) - u(y)| \tag{1.12}$$

which contains a data term, $\sum_{x\in\Omega} D(x)u(x)$, and regularization term, $\sum_{x\in\Omega}\sum_{y\in\mathcal{N}(x)} \frac{R(x,y)}{2}|u(x) - u(y)|$. (See Section 1.2.1.)

Greig *et al.* [64] were the first to use max-flow to address an MRF in image processing, specifically the Ising model. The Ising model can be easily recrafted in a graphical representation given the smoothness term is non-negative and symmetric. In this graph, there is a 'spatial' vertex for each voxel in the image which are adjacent to a source vertex, s, and a sink vertex, t. The weight of the edges between spatial vertices would be the smoothness term, the weight between spatial vertices and s being $\max\{0, D(x)\}$, and the weight between spatial vertices and t being $\max\{0, -D(x)\}$. In this form, the MAP labelling can be expressed as the graph-cut problem:

$$\underset{u(x)}{\operatorname{argmin}} \left(\sum_{x\in\Omega} \max\{0, D(x)\}u(x) + \sum_{x\in\Omega} \max\{0, -D(x)\}(1 - u(x)) + \sum_{x\in\Omega}\sum_{y\in\mathcal{N}(x)} \frac{R(x,y)}{2}|u(x) - u(y)| \right) \tag{1.13}$$

Such a minimization problem was found to be equivalent to an $s - t$ flow maximization problem through the same graph in which each edge has a maximum capacity of its weight [57]. The edges where the capacity of the edge is the limiting factor on the flow through said

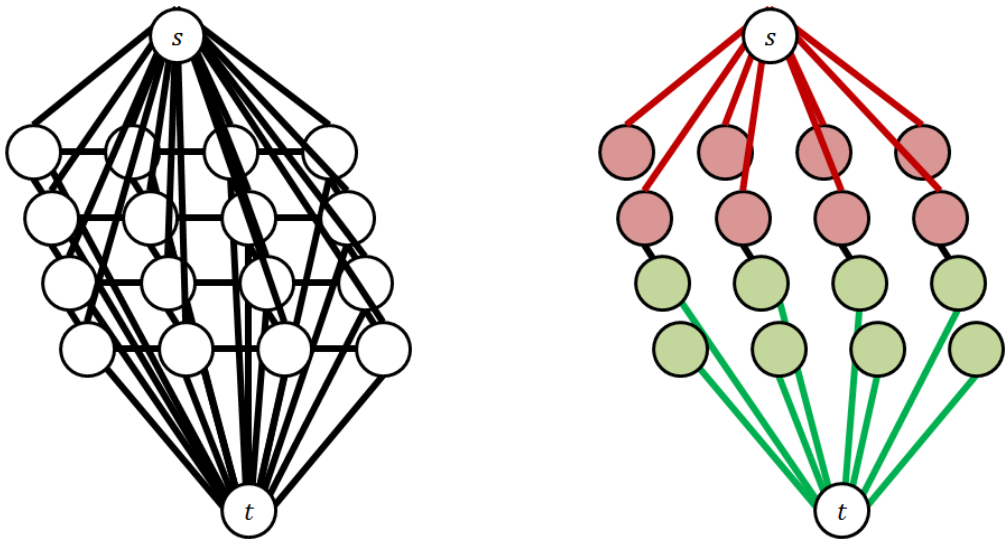

(a) Graph construction for binary image partition

(b) Example cut on said graph with colours illustrating resultant labelling

Figure 1.2: Example of a 4-connected graph construction and cut representing a binary image partition problem.

edge, that edge is part of the minimum cut, that is, the edges that, when removed, partition the graph into two sub-graphs, one with s and the other with t, such that the weight of the edges removed is minimized. Techniques for handling this max-flow/min-cut problem include:

- Augmenting paths algorithms based on the original algorithm by Ford and Fulkerson [57], where paths with excess capacity are identified and saturated; and
- Push-relabel algorithms [61], which order vertices by 'height' expressing their capacity to accept additional incoming flow.

In binary ND image segmentation problems, the most common neighbourhood configuration used is the rectilinear lattice. In this lattice, each variable (i.e. pixel/voxel) is connected to $2 \times N$ others, one pointing upwards and another downwards in each of the N orthogonal directions [24, 26]. An example of said graph is given in Figure 1.2a. A cut through said graph must associate each node with either the source or sink vertex, thus partitioning the image as shown in Figure 1.2b. The remaining edges in the cut represent the boundary of the object.

1.3.2 Applications of Graph-Cuts in Medical Image Processing

Because of this simple graph construction and the existence of many max-flow/min-cut solvers, graph-cuts has seen a wide array of medical applications. These applications include, but are certainly not limited to:

- segmentation of a single object of interest [23, 27] including the prostate [114, 200], liver [166], lungs [4, 29], kidneys [5, 154], cardiac ventricles [11, 84, 104, 115], the whole brain [112, 155] and various neurological structures such as the hippocampus [176],
- segmentation of multiple objects of interest [43, 44] including the abdominal cavity [126],
- MRI phase unwrapping [22, 188],
- fat-water separation in MRI [71],
- medical image fusion for visualization [121], and
- deformable image registration [110, 163, 170].

1.4 Continuous Max-Flow Theory

This section explores the max-flow equation:

$$\begin{aligned} E(u) &= \int_\Omega D(x)u(x)dx + \int_\Omega R(x)|\nabla u(x)|dx \\ &\text{s.t. } u(x) \in \{0,1\} \end{aligned} \tag{1.14}$$

its relationship to the MRFs examined in the previous section, and previous work in minimizing this equation. This equation is the analogue to the Ising model in equation 1.12 in that contains contains a data term, $\int_\Omega D(x)u(x)dx$, and regularization term, $\int_\Omega R(x)|\nabla u(x)|dx$ which fulfil the same practical roles as they do in the discrete model. (See Section 1.2.1.) The labelling function, similarly, takes on a value of either 0 or 1.

Before exploring continuous max-flow, the notions of *convexity* and *convex relaxation* should be introduced. Convexity is a property of both sets and of real-valued functions. A set is convex if and only if, for any two elements in the set, any point on the line segment connecting those elements is also in the set. Formally, a set, C is convex if and only if $\forall(c_1, c_2) \in C, \forall\lambda \in [0,1](\lambda c_1 + (1-\lambda)c_2 \in C)$. A function is convex if and only if the epigraph (the set $\{(x,y)|y \geq f(x)\}$) is a convex set. This is equivalent to the definition that:

$$\begin{aligned} &f(x) \text{ is convex if and only if} \\ &\qquad \forall(x_1, x_2) \in \Omega, \forall\lambda \in [0,1]\,(\lambda f(x_1) + (1-\lambda)f(x_2) \geq f(\lambda x_1 + (1-\lambda)x_2))\,. \end{aligned} \tag{1.15}$$

A convex relaxation is taking a non-convex set, such as the set $\{0,1\}$ and replacing it with the smallest convex set containing it, in this case, the interval $[0,1]$. This smallest convex set is called the *convex hull* of the original set. The first step in exploring a continuous model is

to first apply convex relaxation to the indicator function constraint $u(x) \in \{0, 1\}$ which yields the constraints that $u(x) \in [0, 1]$ and that $u(x)$ must have bounded variation, that is, it must be approximately smooth. In the context of segmentation, convex relaxation transforms a standard segmentation problem into a fuzzy segmentation problem that is easier to solve.

1.4.1 From Discrete Graph-Cuts to Continuous Max-Flow

As a thought-experiment, consider taking a space and sampling it with a rectilinear lattice as described by Boykov *et al.* [24, 26]. The spacing between points on the lattice will be denoted Δr. A continuum model can be developed by letting $\Delta r \to 0$, that is, by making a denser and denser lattice, as shown in Figure 1.3. Before doing this, one must specify how the data terms and regularization terms change as this lattice gets denser. In particular, one must develop *effective* data and smoothness terms $\tilde{D}(x)$ and $\tilde{R}(x, y)$ with desirable limiting case properties, including:

- Both $\tilde{D}(x)$ and $\tilde{R}(x, y)$ grow proportional to the volume element, for the rectilinear sampling case of a D-dimensional space with an isotropic inter-sampling distance of Δr, $\Delta V = (\Delta r)^D$. This intuitively means that as the graph is sampled more, each individual node and edge has a lower and lower importance, there being more of them.

- $\tilde{R}(x, y)$ is proportional to an underlying smoothness field $R\left(\frac{x+y}{2}\right)$ which is differentiable, i.e. $\nabla R(x)$ is finite. $\tilde{D}(x)$ should also be proportional to some underlying function $D(x)$.

- $\tilde{R}(x, y)$ grows inversely proportional to the spacing in between neighbours, that is $|x - y|$. That is, as the spacing gets smaller, knowing the label of an adjacent variable has more value.

Fulfilling these properties yield the terms $\tilde{D}(x) = D(x)(\Delta r)^D$ and $\tilde{R}(x, y) = \frac{R(\frac{x+y}{2})(\Delta r)^D}{|x-y|}$. For now, assume that $\mathcal{N}(x)$ contains only the variables immediately adjacent to x in each of the D

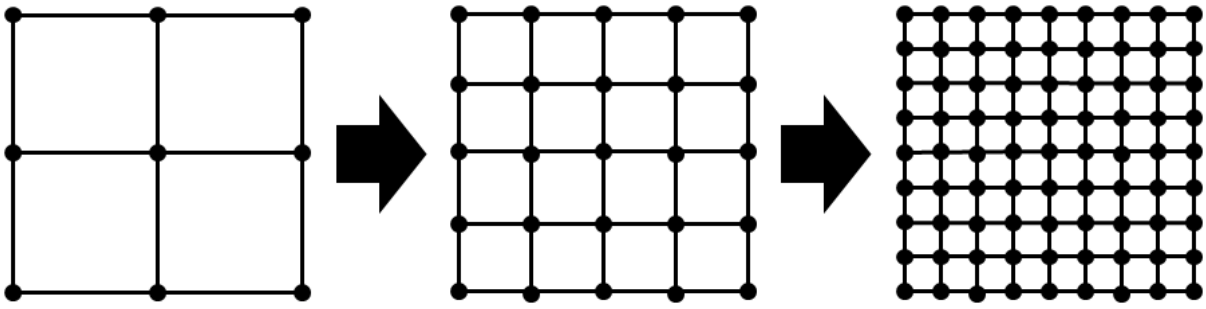

Figure 1.3: MRF as the spacing between nodes, Δr, approaches 0.

directions. Using these assumptions, one can take the limiting case of the energy equation:

$$
\begin{aligned}
E(u) &= \lim_{\Delta r \to 0} \left(\sum_{x \in \Omega} \tilde{D}(x) u(x) + \sum_{x \in \Omega} \sum_{y \in \mathcal{N}(x)} \frac{\tilde{R}(x,y)}{2} |u(x) - u(y)| \right) \\
&= \lim_{\Delta r \to 0} \left(\sum_{x \in \Omega} D(x) (\Delta r)^D u(x) + \sum_{x \in \Omega} \sum_{y \in \mathcal{N}(x)} \frac{R(\frac{x+y}{2})(\Delta r)^D}{2|x-y|} |u(x) - u(y)| \right) \\
&= \left(\int_\Omega D(x) u(x) dx + \int_\Omega \sum_{y \in \mathcal{N}(x)} \frac{R(x)}{2} \left| \nabla u(x) \cdot \frac{x-y}{|x-y|} \right| dx \right) \\
&= \left(\int_\Omega D(x) u(x) dx + \int_\Omega \frac{R(x)}{2} \sum_{i=1}^{D} |\nabla u(x) \cdot e_i| \, dx + \int_\Omega \frac{R(x)}{2} \sum_{i=1}^{D} |\nabla u(x) \cdot -e_i| \, dx \right) \\
E(u) &= \int_\Omega D(x) u(x) dx + \int_\Omega R(x) \sum_{i=1}^{D} \left| \frac{\delta u(x)}{\delta x_i} \right| dx
\end{aligned}
\tag{1.16}
$$

where e_i is the unit vector along the i^{th} axis. This formula implies that the regularization is applied to the L1 norm of the gradient magnitude. This L1 norm arises through the rectilinear nature of the lattice. This may not be appropriate in that it is not rotation invariant, like the L2 norm, and can lead to metrication artefacts.

From a geometric sense, the neighbourhood function, $\mathcal{N}(x)$, has a quite distinct effect, especially in terms of the concept of norms. On a global level, the immediate adjacent neighbourhood, that is, the 4-connected (in 2D) or 6-connected (in 3D) is associated with the L1 norm. As the neighbourhood gets more connected, the associated norm can be thought of as a polygonal norm, that is, a norm in which the unit sphere appears to be a regular polygon. These norms, a selection of which are shown in Figure 1.4, approach the L2 norm as the neighbourhood size increases.

Thus, an additional limit is needed. Not only should $\Delta r \to 0$, but also that the number of neighbours should approach infinity as shown in Figure 1.5, constructing *infinitely dense, infinitely densely connected* lattices. The last assumption is therefore:

- $\tilde{R}(x,y)$ grows proportional to a surface area density (Δa) of the neighbourhood around x. Intuitively, this means that as a neighbourhood is more densely connected, knowing the labelling value of only one neighbour has less of an effect. The surface area density is assumed to converge to the L2 unit sphere, specifically having the property that, for any vector $\theta \in \mathcal{B}^D$, as the size of the neighbourhood gets arbitrarily large, $\sum_{y \in \mathcal{N}(x)} \frac{\Delta a |(x-y) \cdot \theta|}{|x-y|} \to 2$. For some cases, such as an unbounded rectilinear 2D lattice, a formula for this can be

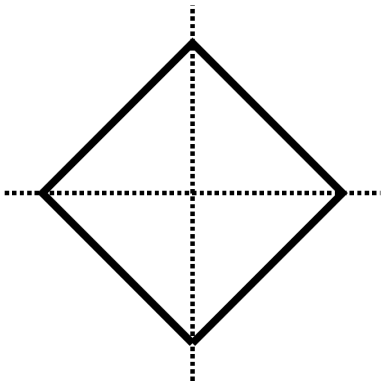

(a) L1 Unit Sphere (4-connected in 2D)

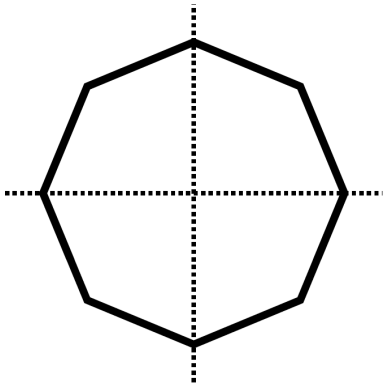

(b) Octagonal Unit Sphere (8-connected in 2D)

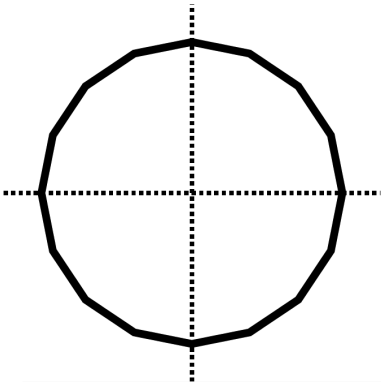

(c) Hexadecagonal (16 sided) Unit Sphere (24-connected in 2D)

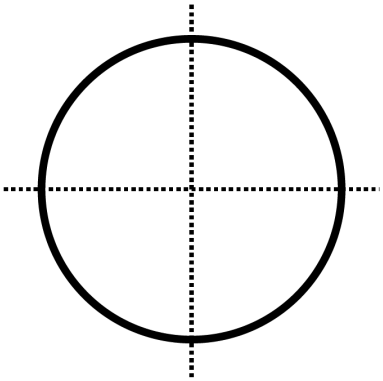

(d) L2 Unit Sphere

Figure 1.4: Unit Spheres Under Various Norms

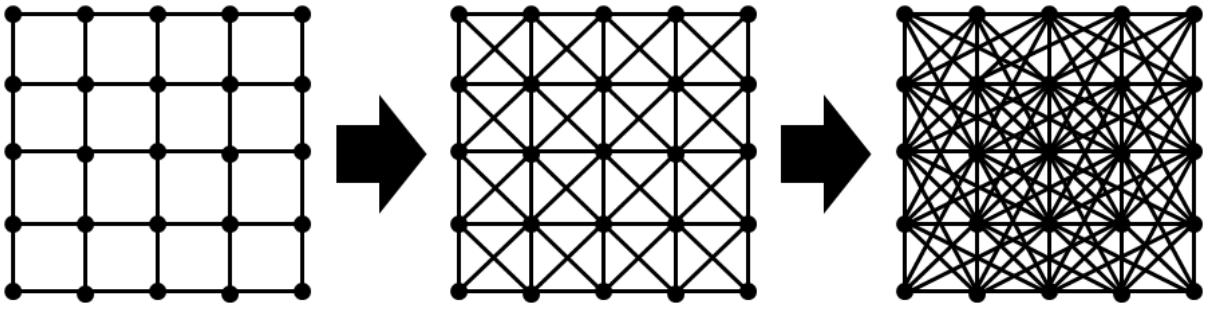

Figure 1.5: MRF as the number of elements in any given $\mathcal{N}(x)$, approaches infinity.

found analytically using an inscribed regular polygon, specifically:

$$\Delta a = \begin{cases} 2\sin(\frac{\pi}{n}), & \text{where } x \text{ has a clear path to } y \\ 0, & \text{else} \end{cases} \tag{1.17}$$

where n is the number of said paths in the first case where x has a clear path to y without intersecting any other element in the neighbourhood. (Note that as $\mathcal{N}(x)$ gets an arbitrarily large number of elements, n also approaches infinity.) This approximation converges quickly, in that for as small as $n = 8$ (as in Figure 1.4b), the best possible in a square neighbourhood of radius 1, $\sum_{y\in\mathcal{N}(x)} \frac{\Delta a|(x-y)\cdot\theta|}{|x-y|}$ is in the interval $[1.84, 2]$. Extending the neighbourhood to have radius 2, that is $n = 16$ (as in Figure 1.4c), $\sum_{y\in\mathcal{N}(x)} \frac{\Delta a|(x-y)\cdot\theta|}{|x-y|}$ is in the interval $[1.96, 2]$. In 3D and higher dimensions, determining this surface area element can be more difficult as the neighbourhood size grows. However, for finite dimensionality, it is always possible.

This constraint yields the terms $\tilde{D}(x) = D(x)(\Delta r)^D$ and $\tilde{R}(x, y) = \frac{R(\frac{x+y}{2})(\Delta r)^D \Delta a}{|x-y|}$. Using these four assumptions, one can take the limit of the energy equation:

$$\begin{aligned} E(u) &= \lim_{\Delta r\to 0}\lim_{n\to\infty}\left(\sum_{x\in\Omega}\tilde{D}(x)u(x) + \sum_{x\in\Omega}\sum_{y\in\mathcal{N}(x)}\frac{\tilde{R}(x,y)}{2}|u(x)-u(y)|\right) \\ &= \lim_{\Delta r\to 0}\lim_{n\to\infty}\left(\sum_{x\in\Omega}D(x)(\Delta r)^D u(x) + \sum_{x\in\Omega}\sum_{y\in\mathcal{N}(x)}\frac{R(\frac{x+y}{2})(\Delta r)^D\Delta a}{2|x-y|}|u(x)-u(y)|\right) \\ &= \lim_{n\to\infty}\left(\int_\Omega D(x)u(x)dx + \int_\Omega\sum_{y\in\mathcal{N}(x)}\frac{R(x)\Delta a}{2}\left|\nabla u(x)\cdot\frac{x-y}{|x-y|}\right|dx\right) \\ &= \lim_{n\to\infty}\left(\int_\Omega D(x)u(x)dx + \int_\Omega R(x)|\nabla u(x)|\sum_{y\in\mathcal{N}(x)}\frac{\Delta a}{2}\left|\frac{\nabla u(x)}{|\nabla u(x)|}\cdot\frac{x-y}{|x-y|}\right|dx\right) \\ E(u) &= \int_\Omega D(x)u(x)dx + \int_\Omega R(x)|\nabla u(x)|dx \end{aligned} \tag{1.18}$$

This demonstrates that continuous max-flow is a limiting case of graph-cuts as the density and connectivity of the lattice involved both approach infinity. Unfortunately, this does not imply that graph-cut methods on extremely large MRFs will be the most effective at addressing continuous max-flow problems, especially considering the memory required to store the extremely large number of edges.

1.4.2 Duality and Convex Variational Optimization

Efficiently minimizing the max-flow energy, Eq (1.14), involves a branch of mathematics known of *calculus of variations*. This is a branch of mathematical optimization in which one tries to optimize a functional with an infinite number of degrees of freedom, rather than a function with a finite number. Because the energy, $E(u)$ takes a function $u(x)$ as its argument, minimizing $E(u)$ involves an infinite number of degrees of freedom, as there are an infinite number of spatial locations, x for which $u(x)$ could take on a different value [51]. An important feature of $E(u)$ is its convexity; it satisfies the definition given in Eq (1.15).

The concept of *duality* is an especially powerful tool in the analysis and optimization of convex functionals. In optimization theory, duality implies that any optimization problem can be viewed from two perspectives, a maximization and a minimization perspective, both of which form a bound on the value of the other. The *duality gap* refers to the difference in the maximum value of the maximization problem and the minimum value of the minimization problem, which in the case of convex optimization problems is guaranteed to be 0. Throughout this thesis, three versions of the max-flow problem will be explored:

1. the *primal* problem of maximizing flow,

2. the *dual* problem of minimizing energy, and

3. the *primal-dual* problem of doing both simultaneously.

These problems are formulated through the use of Lagrangian multipliers. That is, given a primal problem:

$$\begin{aligned} &\max_x f(x) \\ &\text{such that } g_i(x) = 0,\ i = 1, 2, ...n \end{aligned} \tag{1.19}$$

can be transformed into a primal-dual problem:

$$\min_u \max_x f(x) + \sum_{i=1}^{n} u_i g_i(x) \tag{1.20}$$

where u_i are the *Lagrangian multipliers*. From this, one can theoretically construct the *dual function* of $f(x)$ which is a function $f'(u)$ with the definition:

$$f'(u) = \max_x \left(f(x) + \sum_{i=1}^{n} u_i g_i(x) \right). \tag{1.21}$$

which yields the third equivalent problem, the dual problem:

$$\min_{u} f'(u). \tag{1.22}$$

The reader may notice that I used the same variable name u to refer both to the dual variables and to the labelling function. This is because ***they are the same thing***. Dual optimization is such a powerful tool for addressing functionals like Eq 1.14 because it allows for a flow maximization problem through a particular network of continua to be used as the basic optimization problem, and the minimum energy labelling ***is a result of*** the computational of the Lagrangian multipliers on various constraints in the network. That $f'(u)$ is the same functional as $E(u)$, with the exception that it takes on an infinite value whenever $u(x)$ is not a feasible labelling.

1.4.3 Early Approaches to Max-Flow Optimization

The first person to study max-flow optimization in the continuous domain was Gilbert Strang [167] who formalized the analogy between the discrete and continuous cases, providing the intuitive geometrical interpretation of the duality between maximizing $\int_\Omega \operatorname{div} q(x)dx$ under the constraint $|q(x)| \leq 1$ and minimizing $\int_\Omega |\nabla u(x)|dx$. However, it wasn't until Antonin Chambolle [31] who developed the Chambolle iteration (now an essential component of current max-flow solution algorithms) that a truly primal-dual approach was developed. (Previous approaches involved estimating the solution of partial differential equations [7, 8, 33].) Chambolle's approach looked at a particular image restoration problem:

$$E(u) = \int_\Omega \frac{|I(x) - u(x)|^2}{2\lambda} dx + \int_\Omega |\nabla u(x)| dx. \tag{1.23}$$

Chambolle used the primal function $\int_\Omega \operatorname{div} q(x)dx$ under the constraint $|q(x)| \leq 1$ to find the dual function $\int_\Omega |\nabla u(x)|dx$. From this, a projected gradient descent operator can be derived:

$$q(x) \leftarrow \operatorname{Proj}_{|q(x)|\leq 1}\left(q(x) + \tau\nabla\left(\operatorname{div} q(x) - I(x)/\lambda\right)\right) \tag{1.24}$$

which is repeated until $q(x)$ converges, guaranteed if $\tau \leq 1/8$. The final step in this process is to calculate $u(x)$ simply by $u(x) = I(x) - \lambda \operatorname{div} q(x)$

The success of this image restoration algorithm led immediately to Pock and Chambolle developing a general max-flow solver using a *split-merge* approach [32, 137]. In this approach, the max-flow functional is estimated by another functional of the form:

$$E(u, v) = \int_\Omega D(x)u(x)dx + \int_\Omega \left|\frac{u(x) - v(x)}{2c}\right|^2 dx + \int_\Omega R(x)|\nabla v(x)|dx \tag{1.25}$$

in which two similar solutions $u(x)$ and $v(x)$ are iterated using gradient descent. The additional quadratic term in the middle encourages $u(x)$ and $v(x)$ to approximate each other (or becoming equivalent as $c \to 0$). If $u(x)$ is fixed, the optimization problem becomes equivalent to the previous image restoration problem, and a series of Chambolle iterations can be used. If $v(x)$ is fixed, the optimal value of $u(x)$ can be found analytically. The split-merge approach, shown in Algorithm 1.1, is to switch between fixing these two variables, both eventually converging.

Algorithm 1.1: Split-Merge solution algorithm for binary max-flow proposed by Pock and Chambolle [32, 137]

while *not converged* **do**
  while *not converged* **do**
    $q(x) \leftarrow \text{Proj}_{|q(x)|\leq 1}\left(q(x) + \tau\nabla\left(\text{div}\, q(x) - {}^{u(x)}/_{c}\right)\right)$;
  end
  $v(x) \leftarrow u(x) - c\,\text{div}\, q(x)$;
  $u(x) \leftarrow \max\{\,0,\ \min\{\,1,\ v(x) + cD(x)\,\}\,\}$;
end

1.4.4 Augmented Lagrangian Multipliers

Yuan *et al.* [192] took a similar approach, but provided a more computationally efficient ***fully*** primal-dual framework. In this framework, the primal model is taken to be the same max-flow problem suggested by Strang [167], that is:

$$\max_{p_S, p_T, q} \int_\Omega p_S(x)dx \tag{1.26}$$

subject to the *flow conservation constraint*:

$$G(x) = \text{div}\, q(x) + p_T(x) - p_S(x) = 0 \tag{1.27}$$

and the *capacity constraints*:

$$\begin{aligned} p_T(x) &\leq \max\{0, D(x)\} \\ p_S(x) &\leq \max\{0, -D(x)\} \\ |q(x)| &\leq R(x). \end{aligned} \tag{1.28}$$

Using the flow conservation constraint, the primal-dual model can be written as:

$$\begin{aligned} &\min_{u} \max_{p_S, p_T, q} \int_\Omega p_S(x)dx + \int_\Omega u(x)G(x)dx \\ &\text{s.t. } p_T(x) \leq \max 0, D(x) \\ &\quad p_S(x) \leq \max 0, -D(x) \\ &\quad |q(x)| \leq R(x). \end{aligned} \tag{1.29}$$

which can be shown to be equivalent to the dual problem of minimizing Eq (1.14).

Yuan *et al.* [192] perform the optimization on the primal-dual model with one small change; if an additional penalty term $-\int_\Omega \frac{c}{2}G^2(x)dx$ with positive constant c is applied to the equation, the maximization of $p_S(x)$ and $p_T(x)$ given all other variables can be determined analytically. This additional penalty is called an *augmentation* and the entire formula is called the *augmented Lagrangian*. Because it is independent of $u(x)$, it does not effect the minimization component of the primal-dual model, and because it takes on its maximum value (of 0) only at $G(x) = 0$, it does not affect the optimal solution space of the maximization component of the primal-dual model. The only effect the augmentation has is the improved convergence rate [20]. The result is the optimization algorithm shown in Algorithm 1.2.

From this primal-dual view of the entire max-flow problem, Yuan *et al.* [192] were able to show that the global optimum to the max-flow functional under integrality constraints $u(x) \in \{0, 1\}$ can be found by rounding the solution given by Algorithm 1.2 at a predefined threshold. This theorem is important for formal reasons, especially in partitioning problems, because it implies that global optimality can be guaranteed for both strict and fuzzy partitions.

Algorithm 1.2: Augmented Lagrangian solution algorithm for binary max-flow proposed by Yuan *et al.* [192]

while *not converged* **do**
  $q(x) \leftarrow \text{Proj}_{|q(x)| \leq R(x)} \left(q(x) + \tau\nabla \left(\text{div}\, q(x) + p_T(x) - p_S(x) - u(x)/c\right)\right)$;
  $p_T(x) \leftarrow \min\{\max\{0, D(x)\},\ p_S(x) - \text{div}\, q(x) + u(x)/c\}$;
  $p_S(x) \leftarrow \min\{\max\{0, -D(x)\},\ p_T(x) + \text{div}\, q(x) - u(x)/c\}$;
  $u(x) \leftarrow u(x) - c(\text{div}\, q - p_S(x) + p_T(x))$;
end

1.4.5 Proximal Bregman Projections

Another crucial approach to the max-flow functional was developed by Bae *et al.* [13] taking inspiration from message-passing, an alternative approach for MAP estimation in MRFs. In

terms of variational optimization, the precise technique used is the *proximal Bregman projection* [28]. These projections are developed in a specific way. In order to optimize a function $f(x)$, one can take a suboptimal solution x' and improve it by finding a near-by solution with better value. Formally, each proximal Bregman projection can be written as:

$$x \leftarrow \underset{x}{\operatorname{argmin}} \Big(f(x) + c d_g(x, x') \Big) \tag{1.30}$$

where c is a positive constant and $d_g(\cdot, \cdot)$ is a *Bregman distance* constructed out of a convex function $g(\cdot)$. Specifically, this Bregman distance must have the form $d_g(x, x') = g(x) - g(x') - \delta g \cdot (x - y)$. [28] Bae *et al.* uses the entropy function:

$$g(u(x)) = \int_\Omega (u(x) \ln u(x) + (1 - u(x)) \ln(1 - u(x)))\, dx \tag{1.31}$$

which yields the Bregman distance

$$d_g(u(x), v(x)) = \int_\Omega \Big(u(x) \ln({}^{u(x)}/_{v(x)}) + (1 - u(x)) \ln({}^{1-u(x)}/_{1-v(x)}) \Big) dx. \tag{1.32}$$

The benefit of this Bregman distance is that it provides an infinite penalty on any solution $u(x)$ which violates the constraint $u(x) \in (0, 1)$. Similar to the algorithm proposed by Yuan *et al.* [192], the function being optimized is derived from the primal-dual model (which happens to have an analytic solution) and a Chambolle iteration is used to maximize the spatial flows. Unlike Yuan *et al.*'s approach however, Bae *et al.*'s formulation allows for the source and sink flow terms $p_S(x)$ and $p_T(x)$ respectively to cancel out of the optimization. The resulting algorithm is given in Algorithm 1.3.

Because the source and sink flows are implicitly represented rather than explicitly optimized, Bae *et al.* dub this the *pseudo-flow* approach. In addition, the cancellation of the source and sink flows implies that less memory is required to use this algorithm.

Algorithm 1.3: Proximal Bregman solution algorithm for binary max-flow proposed by Bae *et al.* [13]

while *not converged* **do**

$q(x) \leftarrow \operatorname{Proj}_{|q(x)| \le R(x)} \Big(q(x) + \tau \nabla \exp\Big(-\frac{D(x) + \operatorname{div} q(x)}{c} \Big) \Big)$;

$u(x) \leftarrow \frac{u(x) \exp\left(-\frac{D(x) + \operatorname{div} q}{c} \right)}{1 - u(x) + u(x) \exp\left(-\frac{D(x) + \operatorname{div} q}{c} \right)}$;

end

1.4.6 Applications of Continuous Max-Flow in Medical Image Processing

Similar to graph-cuts, max-flow has received a significant amount of attention for a variety of medical image processing tasks. Often, these tasks are similar to those mentioned in Section 1.3.2. Recent applications of continuous max-flow include but are not limited to:

- Segmentation of a single object of interest including the prostate [143, 196], liver [135], cardiac ventricles [124, 148], cerebral ventricles [141], vasculature [175], and spine [12],
- Segmentation of multiple objects of interest including the prostate [142], brain [145] and lungs [66],
- Medical image fusion for visualization [195, 197], and
- Deformable registration, particularly for prostate [169] and cranial [146] MRI.

1.5 Contrasting Graph-Cuts and Continuous Max-Flow

The goal of this introduction was to show the smooth evolution of graph-cut, specifically the Ising model, into the corresponding continuous max-flow analogues. Thus, it has focused primarily on the conceptual equivalences and practical similarities between the two approaches. The goal of this section is to do the opposite, to focus on where the two methods diverge.

As illustrated in Section 1.4.1, in order to use graph cuts to approximate continuous max-flow, one must use not only an infinitely dense, but also an infinitely densely connected graph. This is infeasible in practice, resulting in a phenomenon known as a *metrification artifact*. These artifacts, an example of which is given in Figure 1.6, manifest as unwanted 'blockiness.'

(a) Original Image (b) Graph Cuts Segmentation (c) Max-Flow Segmentation

Figure 1.6: Example of metrification artifacts demonstrated. The graph cut segmentation, which used a 4-connected neighbourhood, is unnecessarily blocky whereas the continuous max-flow solution is more natural in appearance.

These metrification artifacts can also manifest as a preference for creating segmentation edges in a defined set of orientations, rather than at potentially arbitrary angles, as shown by Yuan *et al* [193]. These orientations are dependant on the connectivity of the neighbourhood used. For example, a 4-connected neighbourhood in 2D image segmentation results in a preference for axis-aligned edges.

Another key difference between graph-cuts and max-flow is the computational complexity. The continuous max-flow algorithmic paradigms illustrated earlier (e.g. augmented Lagrangian and proximal Bregman projections) are iterative and numerical with computation time dominated by the convergence rate. Graph-cuts are know to have polynomial time algorithms for some configurations, while other configurations (such as the Potts model in Section 2.2.2) are known to be NP-hard meaning they can take a prohibitive amount of time to solve exactly. For these NP-hard problems, approximate solvers exist which also vary in their runtimes. Thus the computational complexity is much more varied across models for graph-cuts.

From a practical standpoint, the computation time difference between graph-cuts and max-flow depends heavily on the definition of the problem under investigation. For the 2D binary segmentation in Figure 1.6 which contains 256x256 pixels, the computation time for continuous max-flow with GPU acceleration was 0.4 seconds whereas graph-cuts using the Edmond-Karp algorithm [46] took 0.7 seconds, both implemented in MATLAB. Whether or not a particular graph-cut or max-flow method is usable in practice heavily depends on the solution model being used. The additional benefit of the max-flow algorithms presented in this thesis are that they are all trivially parallelizable, meaning that they can readily be implemented using GPGPU programming, unlike graph-cuts solvers based on the original Ford-Fulkerson algorithm [57].

1.6 Thesis Outline

This thesis will rely heavily upon the information presented in this introduction, especially that of Section 1.4 which is the theoretical and technical jumping-off point for the majority of the work presented.

The chapters are as follows:

- ***Chapter 2*** develops a cyclic continuous max-flow image enhancement model and apply it to processing phase information in MRI.

- ***Chapter 3*** develops a continuous max-flow segmentation algorithm which can address label orderings that allow for objects to be hierarchically decomposed into simpler objects.

- ***Chapter 4*** motivates and develops an interactive segmentation interface that allows users to encode their own anatomical knowledge in an abstract form.
- ***Chapter 5*** develops a continuous max-flow segmentation algorithm which can address all possible label orderings.
- ***Chapter 6*** develops a framework for encoding shape information into the previous continuous max-flow segmentation algorithms, allowing for more complicated shapes to be specified.

Each chapter is prefaced with an additional introduction illustrating the theory and recent work similar to the topic at hand.

Chapter 2

Cyclic Continuous Max-Flow Image Enhancement

This chapter is largely based on:

- John S.H. Baxter, Zahra Hosseini, Junmin Liu, Maria Drangova and Terry M. Peters. "Cyclic Continuous Max-Flow: Phase Processing Using the Inherent Topology of Phase." Proceedings of ISMRM (2016).
- John S.H. Baxter, Zahra Hosseini, Terry M. Peters and Maria Drangova. "Cyclic Continuous Max-Flow Phase Processing: A Third Paradigm in Generating Local Phase Shift Maps in MRI." in revision for IEEE Transactions in Medical Imaging.

with additional material from:

- John SH Baxter, A. Jonathan McLeod, and Terry M. Peters. "A continuous max-flow approach to cyclic field reconstruction." arXiv preprint arXiv:1511.03629 (2015).

2.1 Introduction

Image enhancement and restoration has long been a key problem in image processing literature. However, one aspect of image enhancement has fallen by the way-side, specifically the *range topology*, the topology equipped to the range ($\mathbb{I}$) of the labelling function rather than its domain (Ω) . Most images, such as photographs, have a Euclidean range topology and this topology is often implicitly assumed. However, some medical images, such as MR phase images, have a fundamentally different range topology.

Susceptibility-weighted imaging (SWI) and Quantitative Susceptibility Mapping (QSM) are types of MRI sensitive to tissue magnetic susceptibility which is encoded in the phase information encapsulated in the raw MRI data. Because slight changes in tissue magnetic

susceptibility result in small variations in the local magnetic field, small deviations in phase can appear. These small deviations are masked by large variations caused by larger susceptibility differences at distant tissue-air interfaces. Image processing is necessary to isolate the clinical relevant contrast, the *local phase shift maps*, from the smoothly-varying *background phase*. However, phase images do not have a Euclidean range topology, but a cyclic one.

Linear image processing methods do not function in a cyclic topology because many linear concepts, such as linear spatial frequency, no longer have meaning. To illustrate this, consider Figure 2.2 displaying a sawtooth wave. In Figure 2.2a, a Euclidean signal intensity is doubled and the resultant signal has the same frequency and double the amplitude. In Figure 2.2b, the same signal is equipped with a cyclic topology. When this signal is 'doubled' (i.e. taking the Euclidean doubled signal and wrapping it back to the cyclic range $[-\pi, \pi]$) the resultant signal appears to have double the frequency and unchanged amplitude!

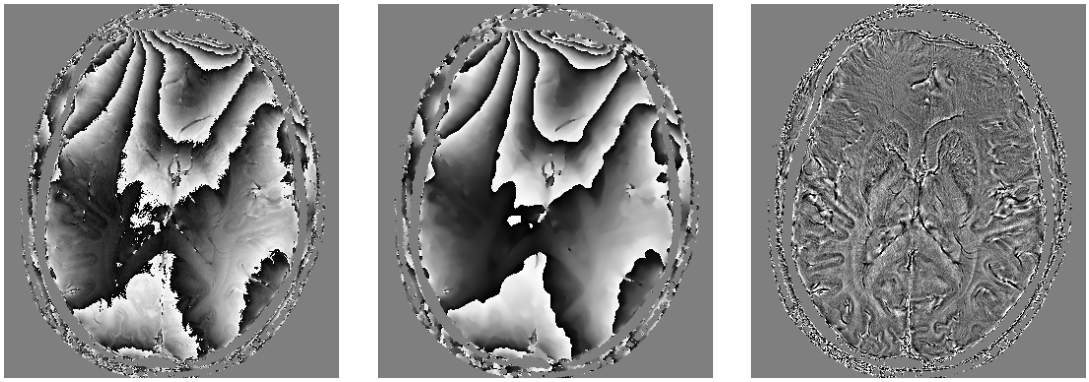

(a) Raw phase data (b) Background phase image (c) Local phase shift map

Figure 2.1: MRI phase image, background phase and local phase shift map.

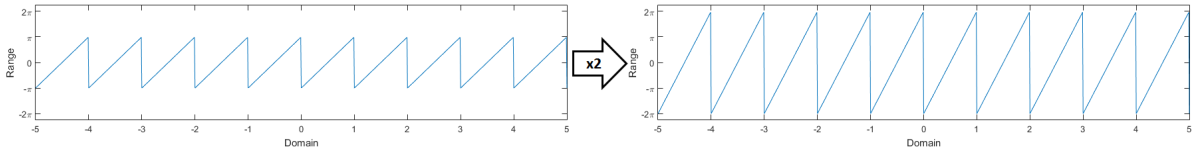

(a) Sawtooth signal equipped with Euclidean range topology. In this topology, when the signal is doubled, it always results in the same frequency spectrum, but with twice the amplitude.

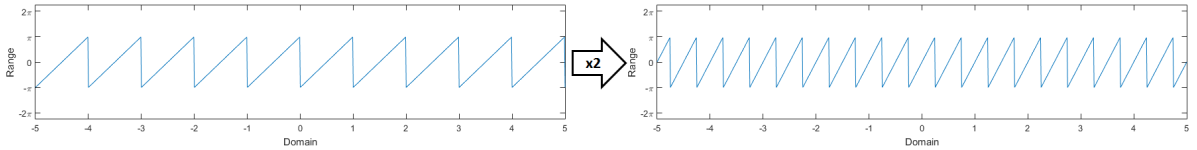

(b) Sawtooth signal equipped with cyclic range topology where π and $-\pi$ wrap to the same value. When this signal is 'doubled', it can result in frequency changes. In this case, the frequency has doubled.

Figure 2.2: Illustrative example of non-linear behaviour in cyclic range topologies.

Phenomena such as that presented in Figure 2.2b are troubling in that they necessitate using a different paradigm for phase images. Currently, there are two paradigms for phase processing. *Homodyne filtering* couples the magnitude and phase together into a complex image with a Euclidean (albeit 2D) topology. *Phase unwrapping* approaches perform a pre-processing step to remove the jumps by adding integer multiplies of 2π to parts of the image, ideally restoring a 1D Euclidean topology. The main contribution of this chapter is the proposal of a third distinct paradigm in which the range topology remains cyclic.

This chapter begins with an overview of how graph-cuts and max-flow techniques have been used in image enhancement, including the Potts and Ishikawa models. This is followed by an overview of MRI phase processing focusing on the two major paradigms: homodyne filtering and phase unwrapping. Cyclic continuous max-flow (CCMF) phase processing is presented as a third, distinct paradigm arising from the use of max-flow techniques on the appropriate range topology, which is validated on both synthetic and 7T MR phase images. This chapter ends with a discussion of max-flow based phase processing and future directions for the approach.

2.2 Previous Graph-Cuts and Max-Flow Image Enhancement Approaches

When considering graph-cut and max-flow image enhancement approaches, there are two particular facets to keep in mind. The first is the range topology, which motivates this chapter. The second is *extendibility*, that is, whether or not such a solution algorithm allows for an arbitrary number of labels to be used (in problems with a discrete range) or achieve an arbitrarily fine intensity-resolution (in problems with a continuous range). Historically, there have been two important models for graph-cut and max-flow image enhancement: the Potts model and the Ishikawa model, although this discussion with a prior continuous max-flow model of certain historical import.

2.2.1 Convex Max-Flow Image Restoration

As stated in Chapter 1 Section 1.4.3, image enhancement and restoration was one of the first uses of continuous max-flow techniques. Specifically, Chambolle's [31] image restoration algorithm and Pock and Chambolle's [32, 137] split-merge algorithm. The latter can be specified

in a way that is more fully generalized, addressing the optimization problem:

$$\min_u \int_\Omega f(u(x), x)dx + \int_\Omega R(x)|\nabla u(x)|dx \tag{2.1}$$

where $f(u(x), x)$ is a convex with respect to $u(x)$ which has bounds of $u(x) \in [u_{\min}, u_{\max}]$. This algorithm now does not have a analytic solution to the coupling of the $u(x)$ and $v(x)$ labellings, but uses a gradient descent procedure to update $u(x)$. The general algorithm is given in Algorithm 2.1.

Algorithm 2.1: Split-Merge solution algorithm proposed by Pock and Chambolle [32, 137] for convex image restoration problems

while *not converged* **do**
  while *not converged* **do**
    $q(x) \leftarrow \text{Proj}_{|q(x)| \le R(x)} \left(q(x) + \tau_1 \nabla \left(\text{div}\, q(x) - {}^{u(x)}/_{c} \right) \right)$;
  end
  $v(x) \leftarrow u(x) - c\, \text{div}\, q(x)$;
  while *not converged* **do**
    $u(x) \leftarrow \max\{ u_{\min}, \min\{ u_{\max}, v(x) + \tau_2 \frac{\delta f(u(x),x)}{\delta u(x)} \} \}$;
  end
end

2.2.2 Discrete Potts Model

The first extendable model for labelling was the Potts model [139] as an extension to the Ising model in which there is an arbitrary finite number of configurations (rather than solely a binary configuration) forming a partition, $\mathbb{L}$. The regularization takes on a particular structure in which every variable encourages its adjacent variables to have the same configuration. Using indicator functions $u_L(x)$ the energy equation can be written as:

$$\begin{aligned} &\min_u \sum_{x\in\Omega}\sum_{L\in\mathbb{L}} D_L(x)u_L(x) + \sum_{x\in\Omega}\sum_{y\in\mathcal{N}(x)}\sum_{L\in\mathbb{L}} \frac{R_L(x,y)}{2}|u_L(x) - u_L(y)| \\ &\text{s.t.} \sum_{L\in\mathbb{L}} u_L(x) = 1 \text{ and } u_L(x) \in \{0, 1\} \end{aligned} \tag{2.2}$$

where $R_L(x, y)$ is positive. In the original model, as well as many extensions, $R_L(x, y)$ does not strictly depend on L, taking the same value regardless of label. If so, this problem can be interpreted as a *multi-way cut* through a graph separating $|\mathbb{L}|$ nodes, each connected to every $x \in \Omega$ with weight $D_L(x)$ for each $L \in \mathbb{L}$. [26]

Boykov *et al.* [26] proposed two algorithms for addressing problems of this form: the α-expansion algorithm and the $\alpha\beta$-swap algorithm. Both algorithms are local optimizers which use binary graph-cuts as a subroutine to select the optimal neighbour iteratively. The neighbourhood of each optimizer is however quite large. In the α-expansion, a neighbour of solution $u(x)$ with respect to a label α is any solution $u'(x)$ with the property:

$$\forall x \in \Omega,\ (u(x) \neq u'(x)) \implies (u'_\alpha(x) = 1) \tag{2.3}$$

and in the $\alpha\beta$-swap algorithm, a neighbour of solution $u(x)$ with respect to two labels α, β is any solution $u'(x)$ with the property:

$$\forall x \in \Omega,\ (u(x) \neq u'(x)) \implies \Big((u'_\alpha(x) = 1 \wedge u_\beta(x) = 1) \vee (u'_\beta(x) = 1 \wedge u_\alpha(x) = 1)\Big) \tag{2.4}$$

Both allowing large numbers of variables to change value in a single move.

One particular result regarding the Potts model is that, although the α-expansion and $\alpha\beta$-swap algorithms perform well in practice, they cannot guarantee global optimality. Solving the Potts model happens to be NP-hard [26, 89] although α-expansion is guaranteed to find a solution to equation (2.4) with less than $2c$ times the minimum energy, where:

$$c = \max_{x \in \Omega, y \in \mathcal{N}(x)} \left\{ \frac{\max_{\alpha \neq \beta \in \mathbb{L}} R_\alpha(x, y) + R_\beta(x, y)}{\min_{\alpha \neq \beta \in \mathbb{L}} R_\alpha(x, y) + R_\beta(x, y)} \right\} \tag{2.5}$$

which has a value of $c = 1$ if $R_L(x, y)$ is independent of L. [26]

2.2.3 Continuous Potts Model

The continuous Potts model [137, 193] is analogous to the discrete Potts model in which the spatial domain Ω is taken to be continuous. Thus, it follows a very similar energy equation:

$$\begin{aligned} &\min_u \sum_{L \in \mathbb{L}} \int_\Omega D_L(x) u_L(x) dx + \sum_{L \in \mathbb{L}} \int_\Omega R_L(x) |\nabla u_L(x)| dx \\ &\text{s.t} \sum_{L \in \mathbb{L}} u_L(x) = 1 \text{ and } u_L(x) \in [0, 1] \end{aligned} \tag{2.6}$$

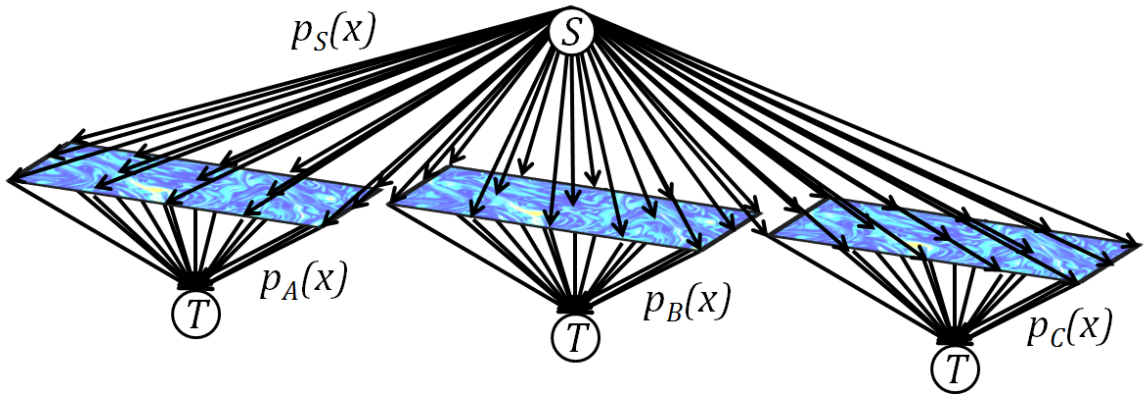

Figure 2.3: Example graph used in the Potts model with labels $\mathbb{L} = \{A, B, C\}$

Algorithm 2.2: Augmented Lagrangian solution algorithm proposed by Yuan *et al.* [193] for the continuous Potts model

$\forall L,\ u_L(x) \leftarrow 0$;
while *not converged* **do**
  $\forall L,\ q_L(x) \leftarrow \text{Proj}_{|q_L(x)| \leq R_L(x)} \left(q_L(x) + \tau \nabla \left(\text{div}\, q_L(x) + p_L(x) - p_S(x) - u_L(x)/c\right)\right)$;
  $\forall L,\ p_L(x) \leftarrow \min\{\, D_L(x),\ p_S(x) - \text{div}\, q_L(x) + u_L(x)/c \,\}$;
  $p_S(x) = \frac{1}{|\mathbb{L}|}\left(\frac{1}{c} + \sum_{L \in \mathbb{L}} (p_L(x) + \text{div}\, q_L(x) - u_L(x)/c)\right)$;
  $\forall L,\ u_L(x) \leftarrow u_L(x) - c(\text{div}\, q_L(x) - p_S(x) + p_L(x))$;
end

Algorithm 2.3: Proximal Bregman solution algorithm proposed by Baxter *et al.* [18] for the continuous Potts model

$\forall L,\ u_L(x) \leftarrow 1/|\mathbb{L}|$;
while *not converged* **do**
  $\forall L,\ u_L(x) \leftarrow u_L(x) \exp\left(-\frac{D_L(x) + \text{div}\, q_L(x)}{c}\right)$;
  $\forall L,\ q_L(x) \leftarrow \text{Proj}_{|q_L(x)| \leq R_L(x)} \left(q_L(x) - c\tau \nabla u_L(x)\right)$;
  $a(x) \leftarrow \sum_L u_L(x)$;
  $\forall L,\ u_L(x) \leftarrow u_L(x)/a(x)$;
end

which can be interpreted as a multi-flow problem with $|\mathbb{L}|$ continuous spaces receiving the same unconstrained ingoing source flow $p_S(x)$ but can have separate spatial flows within them $q_L(x)$ and separate constrained outgoing sink flows $p_L(x)$ as shown in Figure 2.3. This interpretation was used by Yuan *et al.* [193] to develop the Augmented Lagrangian algorithm shown in Algorithm 2.2. Baxter *et al.* [18] later developed a pseudoflow approach shown in Algorithm 2.3. (Pock and Chambolle [137] used a variant of the split-merge algorithm, optimizing the indicator functions $u_L(x)$.)

Because the Potts model is extendable, being capable of incorporating an arbitrarily high number of labels, it can also be used for image enhancement. Consider the image enhancement problem:

$$\min_u \int_\Omega f(u(x), x)dx + \int_\Omega R(u(x), x)|\nabla u(x)|dx \tag{2.7}$$

with any smooth bounded function $f(u(x), x)$. The Potts model can be immediately used for this problem using the following indicator function scheme:

$$\begin{aligned} u(x) \approx u_L \Leftrightarrow & u_L(x) \approx 1 \\ D_L(x) = f(u_L, x) \implies & f(u(x), x) \approx \sum_{L\in\mathbb{L}} D_L(x)u_L(x) \end{aligned} \tag{2.8}$$

which allows for an arbitrarily close approximation of $f(u(x), x)$. However, the regularization component, $R(u(x), x)|\nabla u(x)|$, is not approximated because of the notion of *range topology* mentioned earlier. The *range topology* of the filtering problem assumed by the Potts model is the *discrete topology*. That is, given two intensity values i and j, the metric relating their proximity is:

$$d(i, j) = \begin{cases} c & \text{if } i \neq j \\ 0 & \text{if } i = j \end{cases}. \tag{2.9}$$

where c is a positive constant. In this topology, there is no such thing as a gradient as any two distinct values of i and j are equally far apart, implying that Potts model approximations can appear blocky, especially at higher regularization values.

2.2.4 Discrete Ishikawa Model

Ishikawa *et al.* [78] were the first to examine linear label orderings in segmentation, a similar approach later been taken by Schlesinger and Flach [160]. Specifically, these are problems in which there is a series of labels L_i where $i \in \{0..N\}$, and membership in L_i implies membership

in all labels L_j where $j < i$. In terms of indicator functions, this can be written more simply as $u_{L_i}(x) \geq u_{L_{i+1}}(x)$. Formally, the model can be expressed as:

$$\begin{aligned} &\min_u \sum_{i=1}^{N} \sum_{x\in\Omega} D_{L_i}(x)u_{L_i}(x)dx + \sum_{i=1}^{N} \sum_{x\in\Omega} \sum_{y\in\mathcal{N}(x)} \frac{R_{L_i}(x,y)}{2}|u_{L_i}(x) - u_{L_i}(y)|dx \\ &\text{s.t. } u_{L_i}(x) \geq u_{L_{i+1}}(x) \\ &\quad u_{L_0}(x) = 1 \text{ and } u_{L_{N+1}}(x) = 0 \\ &\quad u_{L_i} \in \{0, 1\} \end{aligned} \tag{2.10}$$

These linear orderings are important in terms of image enhancement because, in the limiting case of $N \to \infty$, they can be used to define a one dimensional Euclidean range topology for image enhancement problems. Ishikawa determined that these problems could be addressed by a single graph-cut through a graph with a particular layered structure shown in Figure 2.4a in which the directed edges could only be cut in a particular way that ensured the linear label ordering property, $u_{L_i}(x) \geq u_{L_{i+1}}(x)$, is upheld.

2.2.5 Continuous Ishikawa Model

The continuous Ishikawa model [14] arose in the context of extending max-flow techniques from Section 2.2.1 in cases where the non-negative function $f(u(x), x)$ is not convex with respect to the first argument, and to do so while maintaining some form of optimality guarantee. Bae *et al.* [14] created a max-flow analogue to the discrete Ishikawa model [78]. The first approach maintained the finite number of label functions, addressing the optimization problem:

$$\begin{aligned} &\min_u \sum_{i=0}^{N} \int_\Omega D_{L_i}(x)u_{L_i}(x)dx + \sum_{i=1}^{N} \int_\Omega R_{L_i}(x)|\nabla u_{L_i}(x)|dx \\ &\text{s.t. } u_{L_0}(x) = 1 \text{ and } u_{L_i}(x) \geq u_{L_{i+1}}(x) \text{ and } u_{L_{N+1}}(x) = 0 \end{aligned} \tag{2.11}$$

which is an approximation of $u(x)$ given some increasing sequence $u_i \in [u_{\min}, u_{\max}]$:

$$\begin{aligned} u(x) &\approx u_0 + \sum_{i=1}^{N} (u_{i-1} - u_i)u_{L_i}(x) \text{ and} \\ f(u(x), x) &\approx \sum_{i=0}^{N} f(u_i, x)(u_{L_i}(x) - u_{L_{i+1}}(x)) = \sum_{i=1}^{N} D_{L_i}(x)u_{L_i}(x). \end{aligned} \tag{2.12}$$

Although this looks similar to the Potts model approximation, the key difference is that the Ishikawa model also approximates the regularization component:

$$R(x)|\nabla u(x)| \approx R(x)\left|\nabla\left(u_0 + \sum_{i=1}^{N}(u_{i-1} - u_i)u_{L_i}(x)\right)\right| = \sum_{i=1}^{N} R_{L_i}(x)|\nabla u_{L_i}(x)| \tag{2.13}$$

which the Potts model approximation cannot claim. Because of the linear range topology, the continuous Ishikawa model has been used for many medical image enhancement problems.

Similar to its discrete counter-part, a labelling in the continuous Ishikawa model can be constructed using a cut through a multi-layered graph using the L_0 node as the flow source as shown in Figure 2.4b. Bae *et al.* [14] extended the augmented Lagrangian approach to the binary max-flow problem to develop the solution algorithm shown in Algorithm 2.4. This algorithm has been used in medical image enhancement applications such as the fusion of spinal MRI and CT volumes [195]. From this, Bae *et al.* were also able to extend the thresholding

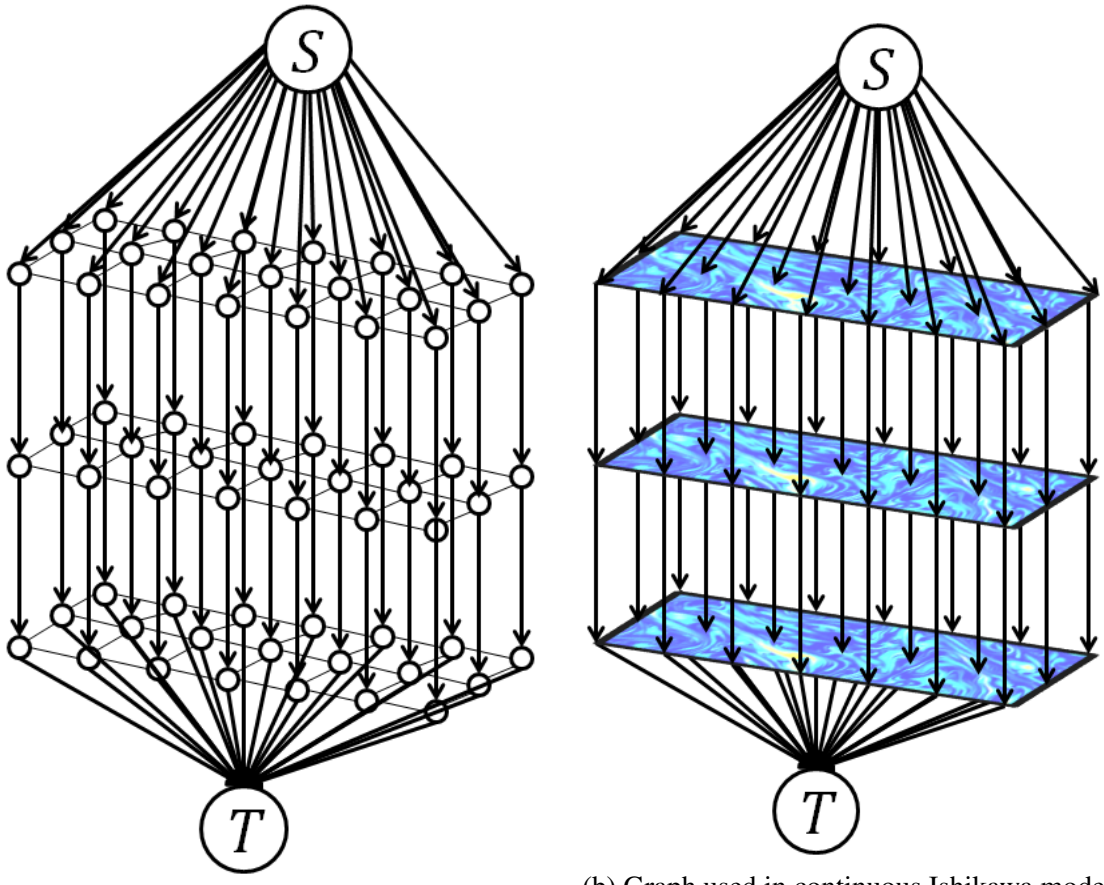

(a) Graph used in discrete Ishikawa model

(b) Graph used in continuous Ishikawa model

Figure 2.4: Example graph used in the Ishikawa model with labels $\{L_0, L_1, L_2, L_3\}$.

Algorithm 2.4: Augmented Lagrangian solution algorithm proposed by Bae *et al.* [14] for the continuous Ishikawa model

$\forall L_i, i \in 0..N,\ u_{L_i}(x) \leftarrow 1/(N+1)$;
$\forall L_i, i \in 0..N,\ F_{L_i}(x) = \sum_{j=0}^{i} D_{L_j}(x)$;
while *not converged* **do**
  $p_{L_0}(x) \leftarrow \min\left\{F_{L_0}(x),\ p_{L_1} + \operatorname{div} q_{L_1}(x) + \frac{1-u_{L_i}(x)}{c}\right\}$;
  for *i from* 1 *to N* **do**
    $q_{L_i}(x) \leftarrow \operatorname{Proj}_{|q_{L_i}(x)| \le R_{L_i}(x)}\left(q_{L_i}(x) - c\tau\nabla\left(\operatorname{div} q_{L_i}(x) - p_{L_{i-1}}(x) + p_{L_i}\right)\right)$;
    if $i \neq N$ **then**
      $p_{L_i}(x) \leftarrow$
      $\min\left\{F_{L_i}(x),\ \frac{1}{2}\left(p_{L_{i+1}} + p_{L_{i-1}}(x) + \operatorname{div} q_{L_{i+1}}(x) - \operatorname{div} q_{L_i}(x) + \frac{u_{L_i}(x) - u_{L_{i+1}}(x)}{c}\right)\right\}$;
    else
      $p_{L_N}(x) \leftarrow \min\left\{F_{L_i}(x),\ p_{L_{N-1}}(x) - \operatorname{div} q_{L_N}(x) + \frac{u_{L_N}(x)}{c}\right\}$;
      $u_{L_i}(x) \leftarrow u_{L_i}(x) - c(\operatorname{div} q_{L_i}(x) - p_{L_{i-1}}(x) + p_{L_i})(x)$;
    end
  end
end

Algorithm 2.5: Proximal Bregman solution algorithm proposed by Baxter *et al.* [18] for the continuous Ishikawa model

$\forall L_i, i \in 0..N,\ u_{L_i}(x) \leftarrow 1/(N+1)$;
$\forall L_i, i \in 0..N,\ F_{L_i}(x) = \sum_{j=0}^{i} D_{L_j}(x)$;
while *not converged* **do**
  $d_{L_0}(x) \leftarrow 0$;
  for *i from* 1 *to N* **do**
    $d_{L_i}(x) \leftarrow d_{L_{i-1}}(x) + \operatorname{div} q_{L_i}(x)$;
  end
  $\forall L_i,\ d_{L_i}(x) \leftarrow u_{L_i}(x) \exp\left(-\frac{d_{L_i}(x) + F_{L_i}}{c}\right)$;
  $a(x) \leftarrow \sum_{L_i} d_{L_i}(x)$;
  $\forall L_i,\ u_{L_i}(x) \leftarrow d_{L_i}(x)/a(x)$;
  for *i from* $N-1$ *to* 1 **do**
    $d_{L_i}(x) \leftarrow d_{L_i}(x) + d_{L_{i+1}}(x)$;
  end
  $\forall L_i, i \in 1..N,\ q_{L_i}(x) \leftarrow \operatorname{Proj}_{|q_{L_i}(x)| \le R_{L_i}(x)}\left(q_{L_i}(x) - c\tau\nabla d_{L_i}(x)\right)$;
end

result of Yuan *et al.* [192] to this model, thus providing a globally optimal solution framework for the continuous Ishikawa model in both the integer and fuzzy cases. A pseudo-flow approach was later derived by Baxter *et al.* [18] using entropic Bregman distances, shown in Algorithm 2.5.

2.3 Susceptibility and MRI Phase Processing

Now that one can use segmentation algorithms to handle non-convex image enhancement problems with either ***discrete*** or ***Euclidean*** range topologies, the question remains, is that good enough for the purposes of medical image processing. The answer turns out to be ***no***, as there remains at least one more range topology, the ***cyclic*** topology, of interest in medical imaging. This cyclic topology appears in MRI phase images, such as the one shown in Figure 2.1a, which are strongly affected by small changes and heterogeneities in the underlying magnetic field. This is beneficial from a clinical standpoint because it implies that these local phase shift maps (LPSM) contain inherent natural contrast dependent on the tissue's *magnetic susceptibility*, or how the tissue creates a parallel or anti-parallel local magnetic field in response to being exposed to another magnetic field.

Measuring this susceptibility, either qualitatively or quantitatively has become crucial to the diagnosis and staging of many neurological disorders associated with differential brain iron levels and myelination [103]. It has seen widespread use in investigating Parkinson's disease [16] and Alzheimer's disease [93]. The differing susceptibility between deoxygenated blood and surrounding tissue has also made susceptibility-based imaging techniques useful in imaging cerebral vasculature specifically for detecting micro-bleeds [109].

Early approaches to quantifying susceptibility with MRI were largely constrained by assuming a simple geometry and a single source of differential susceptibility [21]. The geometry constraints were removed as the image processing techniques became more complex [99] as well as the removal of the simplistic assumptions about the susceptibility distribution [100]. Although working with computer simulations and susceptibility phantoms, these methods laid the conceptual groundwork for the development of quantitative susceptibility mapping (QSM). The conceptual paradigm for QSM was developed at this early stage, consisting of the steps [183]:

1. Acquisition of a series of (often gradient echo (GRE) sequences) which produce phase maps sensitive to differential susceptibility,
2. Separation of the background phase image and the LPSM, and
3. Solving the inverse convolution problem which relates the susceptibility distribution to the LPSM.

Susceptibility-weighted imaging (SWI) follows a similar paradigm although the third quantitative step is replaced with a more readily computable qualitative visualization step, merging the LPSM and magnitude images. Susceptibility tensor imaging (STI) extends QSM assuming that the susceptibility of tissue at any given point is anisotropic and thus changes with respect to the relative orientation of the magnetic field [106].

The type of phase processing concerning this chapter is the second step, the separation of the background phase image and the LPSM in the original phase image. This can be interpreted as a variant of the low-pass and high pass filtering problem although complicated by the non-Euclidean range topology of the phase image. To address this topological issue, there are two traditional paradigms:

1. Homodyne filtering used in SWI, and
2. Phase unwrapping used in both SWI and QSM.

For the sake of notation, the original MRI magnitude, phase, and complex image will be referred to as $M_O(x)$, $\theta_O(x)$ and $C_O(x)$ respectively, the background phase as $\theta_B(x)$ and the LPSM as $\theta_\Delta(x)$.

2.3.1 Homodyne Filtering Paradigm

SWI as proposed by Haacke *et al.* [67] is designed to be a fast, immediately clinically-applicable susceptibility imaging technique. In order to be feasible, it relies on *homodyne filtering*, a very computationally efficient phase filtering technique, to find the background phase image (BPI). This extended the previous work by Wang *et al.* on cerebral venous system visualization [182].

Homodyne filtering in SWI calculates the LPSM using the formula: [67, 182]

$$\begin{aligned}\theta_B(x) =& \angle\left(C_O(x) * w(x)\right)\\ \theta_\Delta(x) =& \angle\left(\frac{C_O(x)}{C_O(x) * w(x)}\right) = \theta_O(x) - \theta_B(x)\end{aligned} \tag{2.14}$$

where $w(x)$ is a spatial low-pass filter kernel. The most common filter used is the k-space Hanning window:

$$\mathcal{F}\{w(x)\} = W(k) = \begin{cases} 0.5\left(1 - \cos\left(\frac{2\pi k}{f k_{\max}}\right)\right), & |k| \le f k_{\max} \\ 0, & \text{else} \end{cases} \tag{2.15}$$

where the parameter $f \in [0, 1]$ indicates the width of the window in terms of the fraction of k-space falling within the window's support. This value is normally set in the range 0.1 for

high contrast to 0.3 for regular SWI.

Aside from being computationally simple, homodyne filtering can also be interpreted using a theoretical basis in wrapped Gaussian distributions. Specifically, $\theta_B(x)$ as computed can be interpreted as the expected value of an unknown wrapped Gaussian distribution where $\theta_O(y)$ in the surrounding area $y \in \mathcal{N}(x)$ are weighted samples thereof with weights $M_O(y)w(|y-x|)$. (This is not the only theoretical interpretation which yields homodyne filtering as an cyclic expected value operator.) With that interpretation, an uncertainty measurement of $\theta_B(x)$ can be found through the biased standard deviation estimator of said wrapped Gaussian distribution:

$$\begin{aligned} &\text{SD} \approx -\ln(R^2) \text{ where} \\ &R^2 = \frac{\left(\left(\sum_{y\in\mathcal{N}(x)} M_O(y)w(|y-x|)\cos\theta_O(y)\right)^2 + \left(\sum_{y\in\mathcal{N}(x)} M_O(y)w(|y-x|)\sin\theta_O(y)\right)^2\right)}{\left(\sum_{y\in\mathcal{N}(x)} M_O(y)w(|y-x|)\right)^2} \end{aligned} \tag{2.16}$$

Unfortunately, homodyne filtering suffers from artifacts in scenarios in which the background phase varies quickly, such as the area surrounding the sinus or the cerebral cortex. In these scenarios, the standard deviation of the wrapped Gaussian distribution is large, implying a poorer quality estimate of the mean, i.e. that of $\theta_B(x)$. Additionally, homodyne filtering techniques do not currently separate the background phase image and LPSM in a physically justifiable manner that lends itself to quantification.

2.3.2 Phase Unwrapping Paradigm

The goal of phase unwrapping is to represent the original phase image not in terms of a cyclic topology, but an equivalent image with a Euclidean topology. That is, it attempts to find an unwrapped phase image $\theta_U(x)$ with the property that:

$$\exists k(x) \in \mathbb{Z},\ \theta_U(x) = \theta_O(x) + 2\pi k(x) \tag{2.17}$$

where $\theta_U(x)$ is smooth and equipped with a standard Euclidean range topology. Several approaches have been used to address this problem including graph-cut based techniques [22], region-growing [199, 185] and minimum spanning trees [6]. Theoretically, the phase unwrapping paradigm has many desirable features. By expressing the phase information with a 1D Euclidean topology, traditional linear image processing techniques can be used and many are well-developed specifically for MRI phase processing used in QSM. Laplacian based phase unwrapped techniques [102] take advantage of the relationship between the Laplacians of $\theta_O(x)$

and $\theta_U(x)$:

$$\begin{aligned}\nabla^2\theta_U(x) &= \cos\theta_U(x)\,\nabla^2\sin\theta_U(x) - \sin\theta_U(x)\,\nabla^2\cos\theta_U(x)\\ &= \cos\theta_O(x)\,\nabla^2\sin\theta_O(x) - \sin\theta_O(x)\,\nabla^2\cos\theta_O(x)\end{aligned} \tag{2.18}$$

and the relationship between the phase Laplacian and the underlying isotropic susceptibility distribution, $\chi(x)$ [157]:

$$\nabla^2\theta_U(x) = \gamma \times \mathrm{TE} \times \mu_0 \times H_0 \times \left(\nabla^2\chi(x)/3 - \delta^2\chi(x)/\delta z^2\right) \tag{2.19}$$

where z is the direction parallel to the main magnetic field.

However, the question arises: is this process always possible? With the previous homodyne filtering technique, the process is trivially possible, but with phase unwrapping, it is unclear if any given phase image is *unwrappable*. One condition for unwrappability is the second fundamental theorem of calculus, that is:

$$\int_C \nabla(\theta_O(x)) \cdot ds = \theta_O(x_2) - \theta_O(x_1) \tag{2.20}$$

where C is any curve from x_1 to x_2. This theorem holds if and only if both the domain and range of the function $\theta_O(\cdot)$ can both equipped with a Euclidean topology. This isn't always the case as shown in Figure 2.5 due to the presence of noise or absence of appreciable magnitude. *Branch-cut* methods developed in the radar interferometry community [62] were the

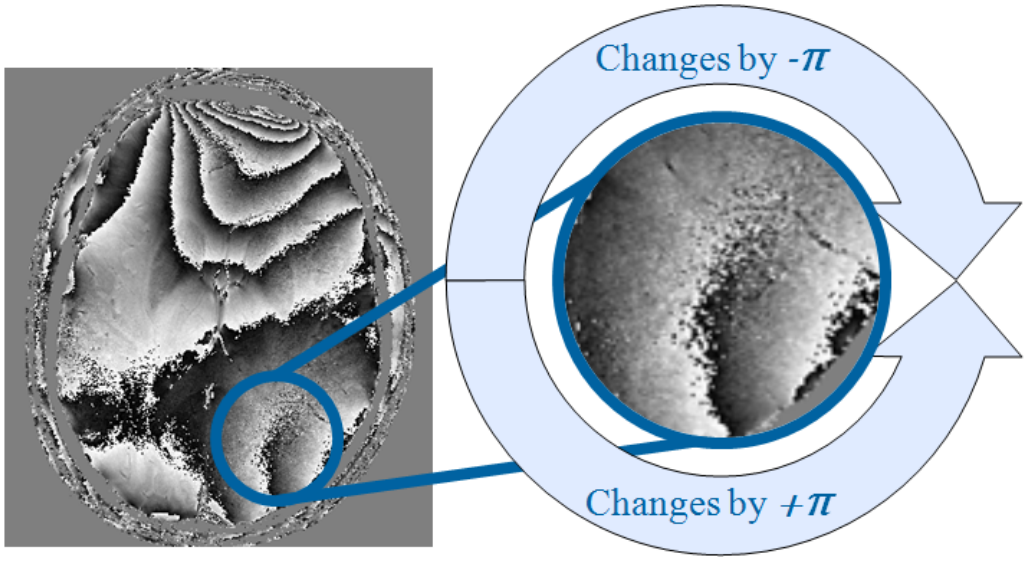

Figure 2.5: Example of an non-unwrappable phase image (left) with corresponding demonstration of path dependence (right

first approaches to handle this problem and are still common in MRI research [36, 41, 156] at least as a comparative approach [55]. In these approaches, the poles at the end of open-ended fringe lines (a sure-fire sign of non-Euclideanity) are joined together, creating a 'cut' which is removed from the domain of the image, rendering the domain non-Euclidean, but preserving an unwrappable Euclidean range. In 3D images, these cuts must form a contiguous surface and are often optimized through the notion of minimal surfaces [156]. Visually, these approaches appear to create a discontinuity in the region of the cut, even if such a discontinuity is not apparent locally in the original phase image. The motivation behind this chapter, instead of removing the Euclidean domain, is to use a cyclic, rather than Euclidean, range.

2.4 Cyclic Continuous Max-Flow Formulation

The goal of *cyclic continuous max-flow* (CCMF) is to address image enhancement functionals of the form:

$$\min_u \int_\Omega f(u(x), x)dx + \int_\Omega |\mathbf{R}(u(x), x)\nabla u(x)|_p dx \tag{2.21}$$

where the gradient operator $\nabla u(x)$ is cognizant that the range topology of $u(x)$ is cyclic and is therefore only locally Euclidean. These topological considerations would make CCMF particularly suited for MRI phase smoothing, in which such a cyclic range topology holds.

Similar to the Potts and Ishikawa models, CCMF approximates this function in terms of a series of indicator functions, $u_\theta(x)$. As the number of indicator functions grows arbitrarily large, this energy can be re-expressed as:

$$\begin{aligned} u(x) &= \int_\Theta \theta u_\theta(x)d\theta \\ \int_\Omega f(u(x), x)dx &= \int_\Theta \int_\Omega D_\theta(x)u_\theta(x)dxd\theta \\ \int_\Omega |\mathbf{R}(u(x), x)\nabla u(x)|_p dx &\approx \int_\Theta \int_\Omega |\mathbf{R}_\theta(x)\nabla u_\theta(x)|_p dxd\theta \end{aligned} \tag{2.22}$$

where the indicator functions $u_\theta(x)$ exist on a cylindrical manifold as shown in Figure 2.6 and the gradient operator ∇ is applied across both the spatial, x, and range, θ, dimensions. This leads to the energy functional:

$$\begin{aligned} &\min_u \int_\Theta \int_\Omega D_\theta(x)u_\theta(x)dxd\theta + \int_\Theta \int_\Omega |\mathbf{R}_\theta(x)\nabla u_\theta(x)|_p dxd\theta \\ &\text{s.t.} \int_\Theta u_\theta(x)d\theta = 1 \text{ and } u_\theta(x) \geq 0 \end{aligned} \tag{2.23}$$

2.5 Cyclic Continuous Max-Flow Algorithm

As with the Potts and Ishikawa models, they can be addressed through primal-dual optimization with either augmented Lagrangians or proximal Bregman projections. Additionally, the Chambolle iteration [31] appears in both which minimizes the (dual) regularization term through the constrained maximization of its primal spatial flow variables.

The augmented Lagrangian solution algorithm to the cyclic continuous max-flow formulation in the form of indicator functions is given in Algorithm 2.6. As with the previous augmented Lagrangian algorithms, it involves two positive parameters: τ, the gradient descent step size set to approximately $1/8$; and c, the quadratic penalty constant set to a default value of $1/4$.

Algorithm 2.6: Augmented Lagrangian solution algorithm for the CCMF functional in terms of indicator functions.

$\forall x, \theta,\ u_\theta(x) = 1/2\pi$;
while *not converged* **do**
  $\forall x, \theta,\ q_\theta(x) \leftarrow \text{Proj}_{|\mathbf{R}_\theta^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} \left(q_\theta + \tau \nabla \left(\text{div}\, q_\theta(x) + p_\theta(x) - p_S(x) - u_\theta(x)/c\right)\right)$;
  $\forall x, \theta,\ p_\theta(x) \leftarrow \min\{D_\theta(x), p_S(x) - \text{div}\, q_\theta(x) + u_\theta(x)/c\}$;
  $\forall x,\quad p_S(x) \leftarrow \frac{1}{2\pi}(1/c + \int_\Theta (p_\theta(x) + \text{div}\, q_\theta(x) - u_\theta(x)/c) d\theta$;
  $\forall x, \theta,\ u_\theta(x) \leftarrow u_\theta(x) - c\left(\text{div}\, q_\theta(x) - p_S(x) + p_\theta(x)\right)$;
end

The proximal Bregman solution algorithm to the cyclic continuous max-flow formulation in the form of indicator functions is given in Algorithm 2.7. Similar to the augmented Lagrangian algorithm, τ is the gradient descent step size which is set to approximately $1/8$. However, c is

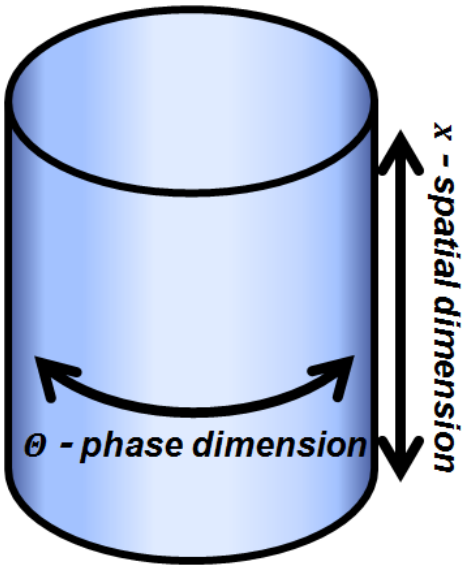

Figure 2.6: Topology with which CCMF indicator functions are equipped

now the weight of the proximity term in each projection step which is set to be on the order of the difference in data terms between adjacent values of θ.

Algorithm 2.7: Proximal Bregman solution algorithm for the CCMF functional in terms of indicator functions.

$\forall x, \theta,\ u_\theta(x) = {}^1/_{2\pi}$;
while *not converged* **do**
  $\forall x, \theta,\ u_\theta(x) \leftarrow u_\theta(x) \exp\left({}^{-(D_\theta(x) + \operatorname{div} q_\theta(x))}/_{c}\right)$;
  $\forall x, \theta, q_\theta(x) \leftarrow \operatorname{Proj}_{|\mathbf{R}_\theta^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1}\left(q_\theta - c\tau \nabla u_\theta(x)\right)$;
  $\forall x,\quad a(x) \leftarrow \int_\Theta u_\theta(x) d\theta$;
  $\forall x, \theta,\ u_\theta(x) \leftarrow u_\theta(x)/a(x)$;
end

For a full theoretical proof of correctness of these algorithms and any additional technical information, see Appendix Section A.1.

2.6 Cyclic Continuous Max-Flow Synthetic Validation

Continuing the motivation of phase processing using a CCMF, Figure 2.7 shows the results of applying prior continuous max-flow image enhancement approaches (the Potts model and the Ishikawa model). Due to the large number of phase bins the thickness of the spatial region corresponding to each phase bin is minimal. This causes excessively blockiness in the Potts model due to the shrinking bias. The Ishikawa model is mostly smooth but experiences errors about the phase wraps because of the assumption of a linear topology. That is, the Ishikawa model only assumes a smooth transition in intensities between $-pi$ and pi exclusive, viewing the phase wrap as a discontinuity. CCMF alleviates both these issues, as intended.

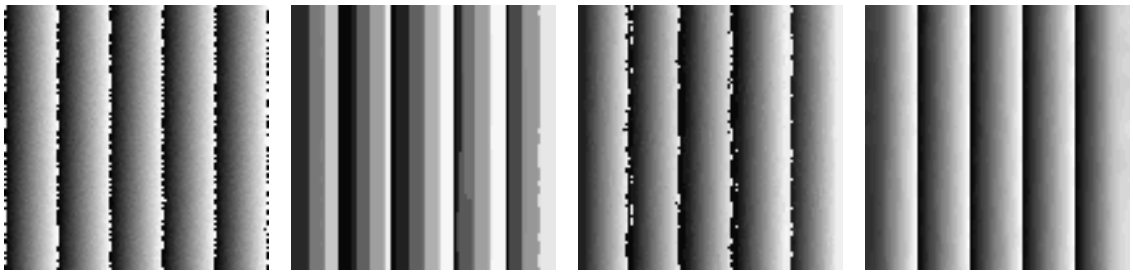

(a) Noisy Phase Image (b) Potts Model (c) Ishikawa Model (d) CCMF Model

Figure 2.7: Comparison of phase smoothing using the Potts, Ishikawa, and CCMF models. The Potts model is excessively blocky and the Ishikawa model is error-prone surrounding the phase wraps. 40 phase bins were used in each model.

Each method takes $O(NM)$ space where N is the size of the image and M is the number of phase bins used for interpolation. The Potts and CCMF methods both required 2 seconds for processing, whereas the Ishikawa model required 3 seconds, although all methods displayed linear convergence time and linear iteration time.

2.6.1 Images

As CCMF reflects a variety of techniques both in terms of optimization functionals and solvers, it is important to gauge its efficacy using highly controlled digital phantom experiments visually similar to MR images. The phantoms (one of which is shown in Figure 2.8) consist of a phase model with linearly increasing frequency in addition to heavily smoothed Gaussian noise and randomly placed phase poles outside of a circular mask. This constitutes the smoothly varying background phase signal. Independent and identically distributed complex Gaussian noise is then added to each pixel with a CNR of 300%.

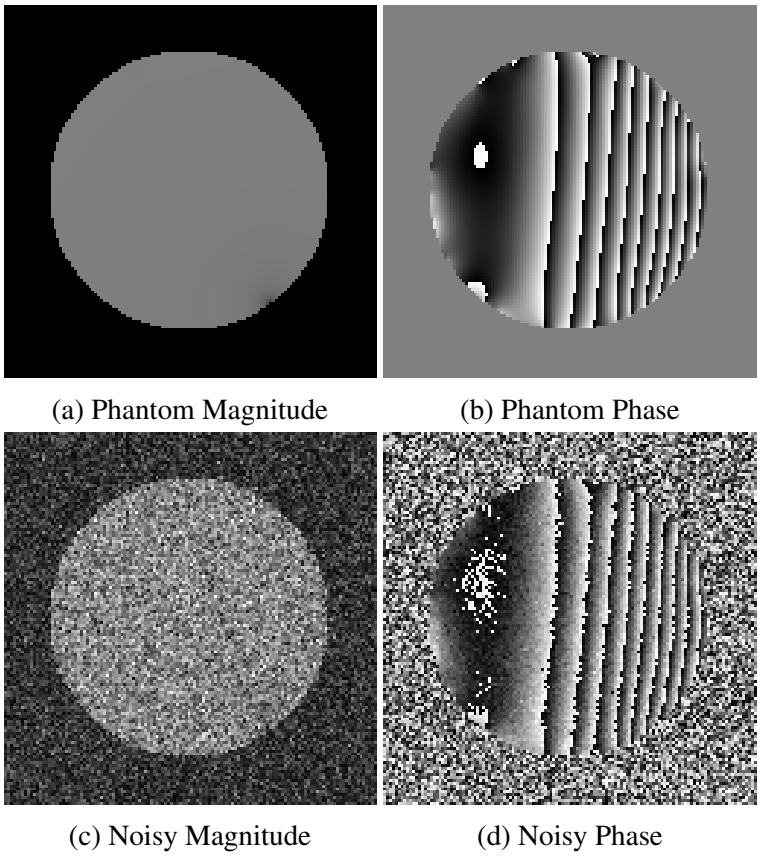

(a) Phantom Magnitude (b) Phantom Phase

(c) Noisy Magnitude (d) Noisy Phase

Figure 2.8: Phantom experiment gold standard and noisy images. Low-pass filtered results on the noisy image are shown in Fig. 2.9.

2.6.2 Methods

Two different data terms were used, each based either on a complex Gaussian noise model (similar to that assumed by homodyne filtering):

$$\begin{aligned} D_\theta(x) =& -\ln P(\angle(C(x)+\eta)=\theta) \text{ where } \eta \sim \mathcal{N}(0,\sigma) \\ =& \frac{m^2}{2} - \ln\left(1+\sqrt{\pi}\left(qe^{q^2}(1+\operatorname{erf}(q))\right)\right) \text{ where } m=\frac{|C(x)|}{\sigma}, q=\frac{m\cos(\theta-\angle C(x))}{\sqrt{2}}, \end{aligned} \tag{2.24}$$

or a model similar to that used by phase unwrapping:

$$D_\theta(x) = |\angle C(x) - \theta| \tag{2.25}$$

which is independent of the signal magnitude.

Additionally, the regularization terms may be isotropic, penalizing background phase image variation uniformly, or anisotropic in which the variation is penalized differently based on its orientation. (For mathematical details regarding anisotropic regularization, see Appendix A.5.1.) The preferred (i.e. low penalization) direction was extracted as the gradient direction of the homodyne-smoothed image. (In this case, phase wraps have a negligible effect on the direction of the gradient excepting 180° flips.)

Ultimately, these two considerations (along with the two algorithms) provide eight alternative CCMF methods. Each method was run multiple times on each randomly generated phantom varying only the regularization weight. The smoothed images are compared to the gold standard via *normalized error* with the unprocessed image reflecting 100% error and the gold standard reflecting 0% error. Thus, it is possible to receive greater than 100% error if the image is over-smoothed to a degree that it is further away from the gold standard than the original noisy image.

2.6.3 Results

All eight methods were run a total of $n = 25$ randomly generated phantoms. For qualitative results, a single phantom was generated and each method was applied to it using the optimal parameterization in Figure 2.10. These qualitative results are shown in Figure 2.9.

The augmented Lagrangian method with anisotropic smoothing and a Gaussian data term performed the best. The selection of data term and regularization term is intuitive in that the phantom had Gaussian noise applied to it (as assumed in the Gaussian data term formulation) and had a sufficiently large frequency in the right side that disallowed higher regularization weights with isotropic smoothing. The proximity of the Uniform data term's performance to

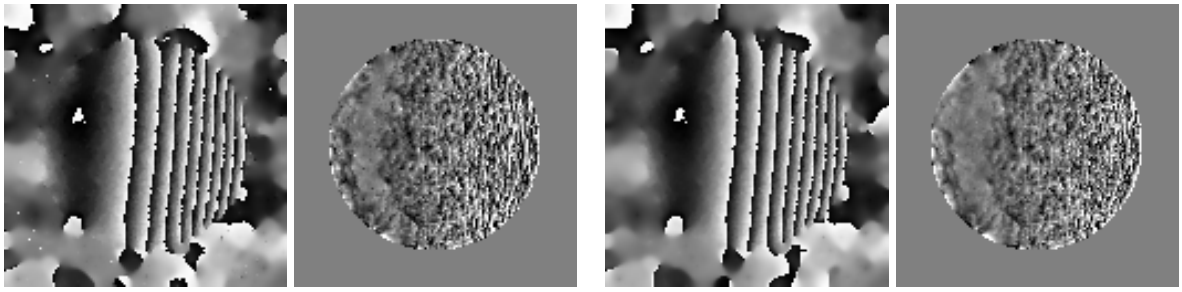

(a) Gaussian data and isotropic regularization

(b) Gaussian data and anisotropic regularization

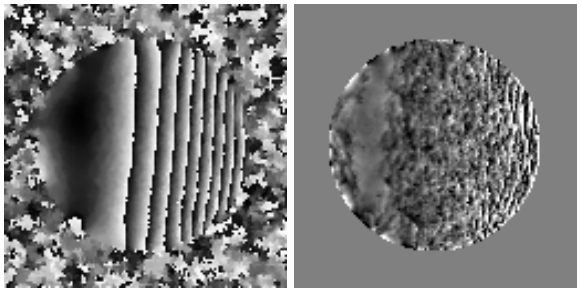

(c) Uniform data and isotropic regularization

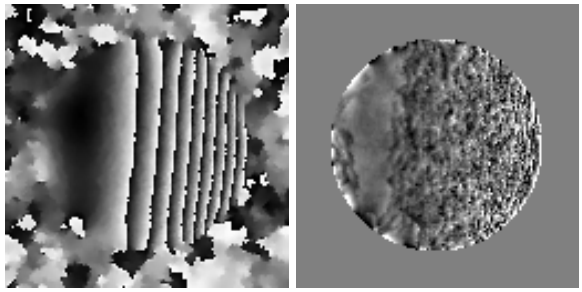

(d) Uniform data and anisotropic regularization

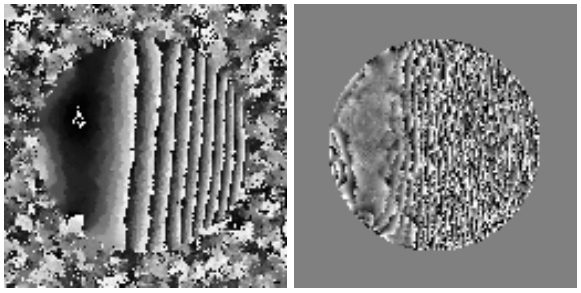

(e) Gaussian data and isotropic regularization

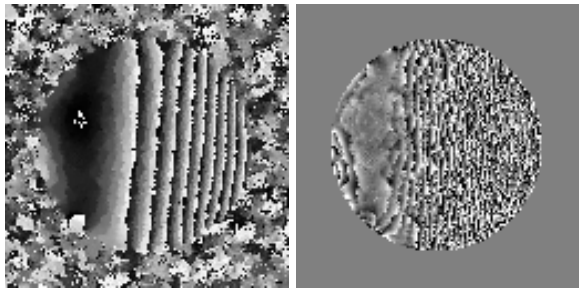

(f) Gaussian data and anisotropic regularization

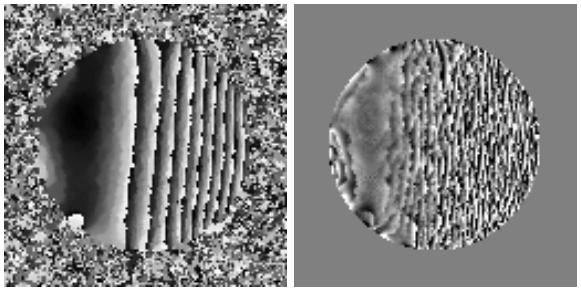

(g) Uniform data and isotropic regularization

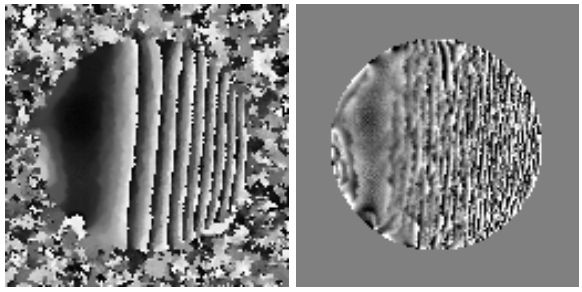

(h) Uniform data and anisotropic regularization

Figure 2.9: Example low-pass filtering results using the Augmented Lagrangian solver (a to d) and the Proximal Bregman solver (e to h). Each pair includes the low-pass filtered image and a difference image between the result and the noise-free phantom image, Fig. 2.8b.

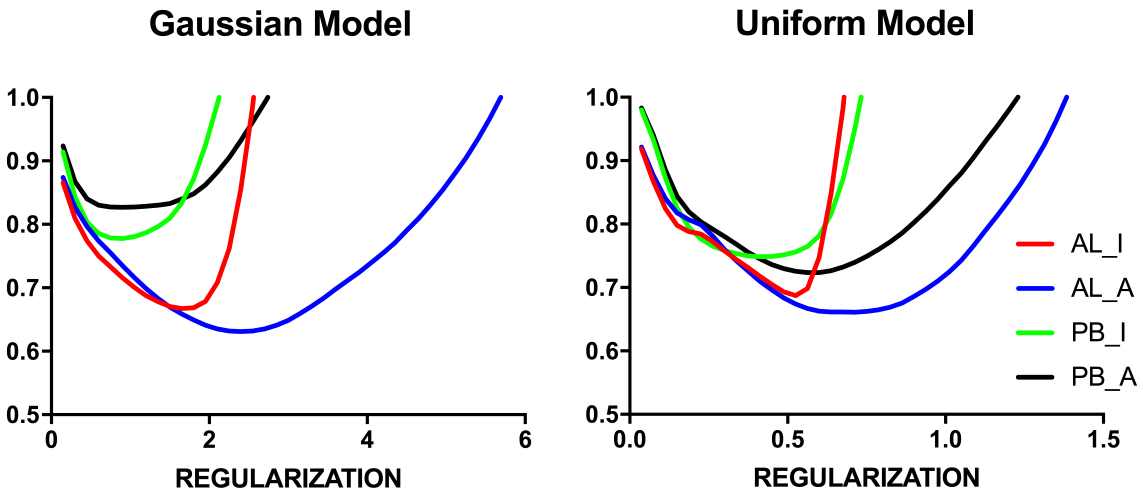

Figure 2.10: Error reduction with varying regularization weight

that of the Gaussian data term, combined with its simplicity and magnitude independence, made it a better candidate for later experiments.

2.7 Cyclic Continuous Max-Flow in MRI Phase Processing

2.7.1 Images

All data acquisition was performed in accordance to the requirements of the research ethics board at Western University and informed consent was obtained. A healthy volunteer was scanned at 7 T (7-T neurospecialized MR imaging system, Siemens, Erlangen, Germany) using a 16-channel transmit/receive head coil. A three-dimensional flow compensated six-echo gradient echo imaging sequence was used with the following parameters: TR/TE1/ESP: 40/3.7/4.1 (ms), Flip angle 13, voxel size 0.5x0.5x1.25 mm3. Image acquisition at both field strengths were performed with an acceleration factor of 2. In line with recent work highlighting the advantages of channel-by-channel processing of MR phase data [32], the reconstructed channel images (using GRAPPA) were saved and each pipeline was set up to process the individual channel data prior to application of channel combination. (These images were collected as part of a larger study investigating MR imaging for multiple sclerosis at 7T. The ethics approval for this data is given in Appendix C Section C.1.)

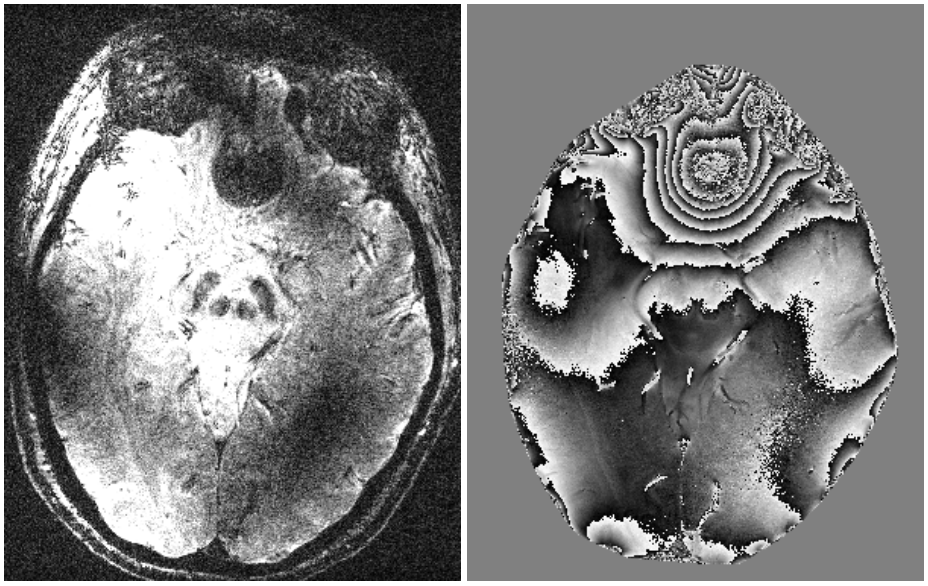

(a) Single Channel Magnitude Image (b) Single Channel Phase Image

Figure 2.11: Single channel cranial MR image including magnitude (a) and phase (b) components.

2.7.2 Methods

Four methods were applied to single channel data. CCMF was given a uniform model data term and anisotropic regularization and solved via augmented Lagrangian multipliers (Algorithm 2.6) with results shown in Figure 2.12. The three comparative methods included homodyne filtering [67, 182], Gaussian high pass filtering of the unwrapped phase with *phase unwrapping using recursive orthogonal referring* (PUROR) [107], and Gaussian high pass filtering of the unwrapped phase with *phase unwrapping max-flow* (PUMA) [22].

The quantitative comparison was performed using contrast to noise ratio calculated for vessels in different regions of the brain in a manner similar to that described previously [77]. Briefly, six representative imaging slices were selected from regions of the brain near the sinuses, in the areas of the ventricles, and slices superior to the ventricles. For each of the image slices the outputs of each of the processing paradigm were stacked, ensuring the images are co-registered. For each image slice, five veins of ≤ 3 mm diameter were selected and three line segments were drawn perpendicular to the vein, between 1 and 3 mm apart. For each segment the absolute contrast between the vein and the surrounding tissue was calculated as follows: The minimum intensity on the profile corresponded to the vessel signal intensity and the mean of the pixels coinciding with the tails of the profile (defined as two pixels away from the minimum intensity to the ends of the profile) defined the signal intensity of the surrounding tissue. The absolute difference between the two defined the absolute contrast. Contrast-to-noise ratio (CNR) was also evaluated, with noise defined as the standard deviation of the tails.

The quantitative CNR data were compared between the different paradigms using a three-way ANOVA test. Factors in this test included method (each paradigm), the echo time, and the vein being measured. Confidence intervals on the mean CNR difference between each method were also computed.

Both phase unwrapping approaches and homodyne filtering were given optimized parameters with respect to the quantitative experiment [77]. The parameters for the CCMF model were chosen as the optimal values in the phantom experiment described in Section 2.6.

2.7.3 Single Channel Qualitative Results

Qualitative results for single channel images are shown in Figure 2.12. The homodyne and PUROR based techniques both experienced artifacts in the area surrounding the sinus likely due to a combination of the depressed magnitude and more quickly varying background phase variation. Both phase unwrapping approaches, PUROR and PUMA, suffered from an artifactual phase wrap in the right occipital lobe corresponding to an open-ended fringe line in the original single channel phase image. In addition, the CCMF results appear to capture high contrast between the white and gray matter throughout the image similar to the phase unwrapping methods. (Note that homodyne, PUMA, and PUROR have been previously optimized for this exact problem ensuring their performance.) This demonstrates that CCMF can be considered comparable to the clinical and state-of-the-art phase processing methods for single channel data. It is worth noting that CCMF outperformed both phase unwrapping approaches based on computation time.

Algorithm	CCMF (s)	PUROR (s)	PUMA (s)
Echo 1	2.04 ± 0.38	2.88 ± 0.14	6.48 ± 2.34
Echo 2	2.01 ± 0.27	2.91 ± 0.09	6.63 ± 2.48
Echo 3	2.01 ± 0.28	2.92 ± 0.08	6.11 ± 1.57
Echo 4	2.02 ± 0.28	2.93 ± 0.09	6.20 ± 1.68
Echo 5	2.05 ± 0.37	2.98 ± 0.12	6.29 ± 1.70
Echo 6	2.06 ± 0.28	3.01 ± 0.10	6.52 ± 1.63
All Echos	2.03 ± 0.31	2.94 ± 0.11	6.37 ± 1.94

Table 2.1: Computation times for 2D image slices at varying echo time

2.7.4 Channel Combined Qualitative Results

Each channel was processed and the channels combined afterwards using *inter-echo weighted variance* (IEV) [108] to produce multi-channel phase maps. (The corresponding images, i.e. the channel combined qualitative results, are shown in Figure 2.13). Both the higher contrast in

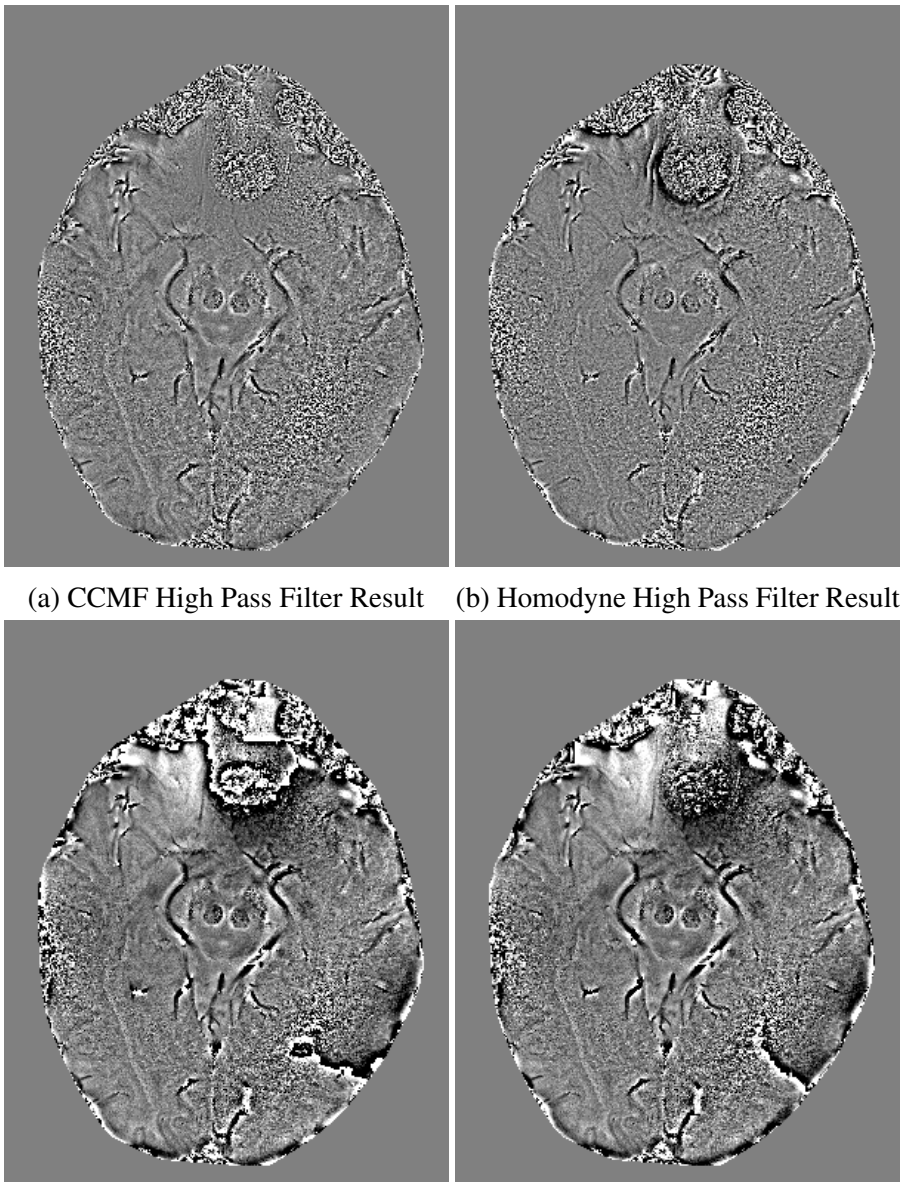

(a) CCMF High Pass Filter Result (b) Homodyne High Pass Filter Result

(c) PUROR High Pass Filter Result (d) PUMA High Pass Filter Result

Figure 2.12: Local phase shift maps on the single channel MR image shown in Figure 2.11.

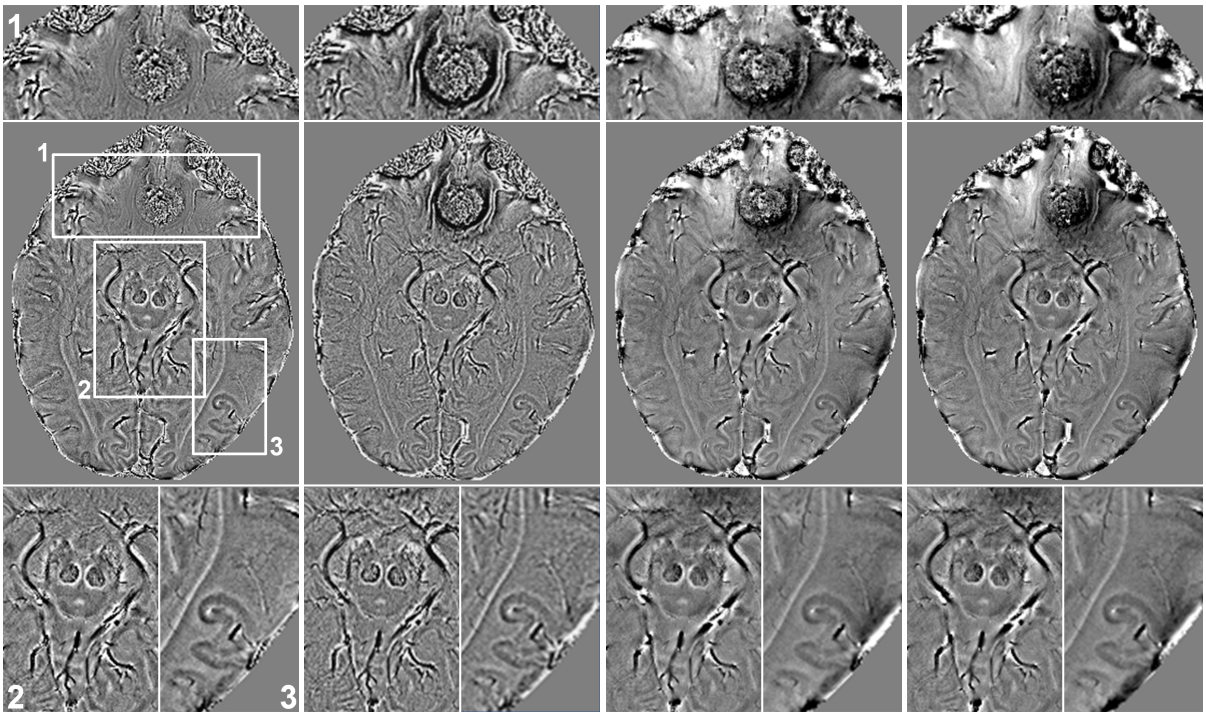

(a) CCMF High Pass Filter Result (b) Homodyne High Pass Filter Result (c) PUROR High Pass Filter Result (d) PUMA High Pass Filter Result

Figure 2.13: Local Phase Shift Map computed via channel combination of single channel local phase shift maps such as those presented in Figure 2.12.

the CCMF single channel images and the artifacts around the sinuses in the homodyne images are propagated to the channel combined images, being consistent across channels. Smaller cortical veins (seen in Figure 2.13 enlarged region 3) appear much more salient in CCMF images that in the comparative approaches. This is beneficial because it indicates that CCMF may be better suited for the investigation of the venous system outside of the major vessels.

The phase unwrapping artifacts present in both PUROR and PUMA do not visibly appear in the channel combined image, indicating that they are not placed in the same spatial location across channels. The channel combination, through averaging, thus decreases their signal to below that of the noise floor. (If one considers each pixel being an approximate average of those in N channels, the artifact signal decreases with $O(1/N)$ and the noise with $O(1/\sqrt{N})$ indicating that the artifact-to-noise ration would decrease with $O(1/\sqrt{N})$ and thus quickly approach the noise floor.) The extent to which these artifacts remain can be seen in Figure 2.14 in which the channel combined PUROR and PUMA images are subtracted thus cancelling out the signal and the vast majority of noise. In this image, the residual phase wrapping artifacts from both techniques can be readily seen.

Although these residual artifacts in phase unwrapping techniques are subtle, they may

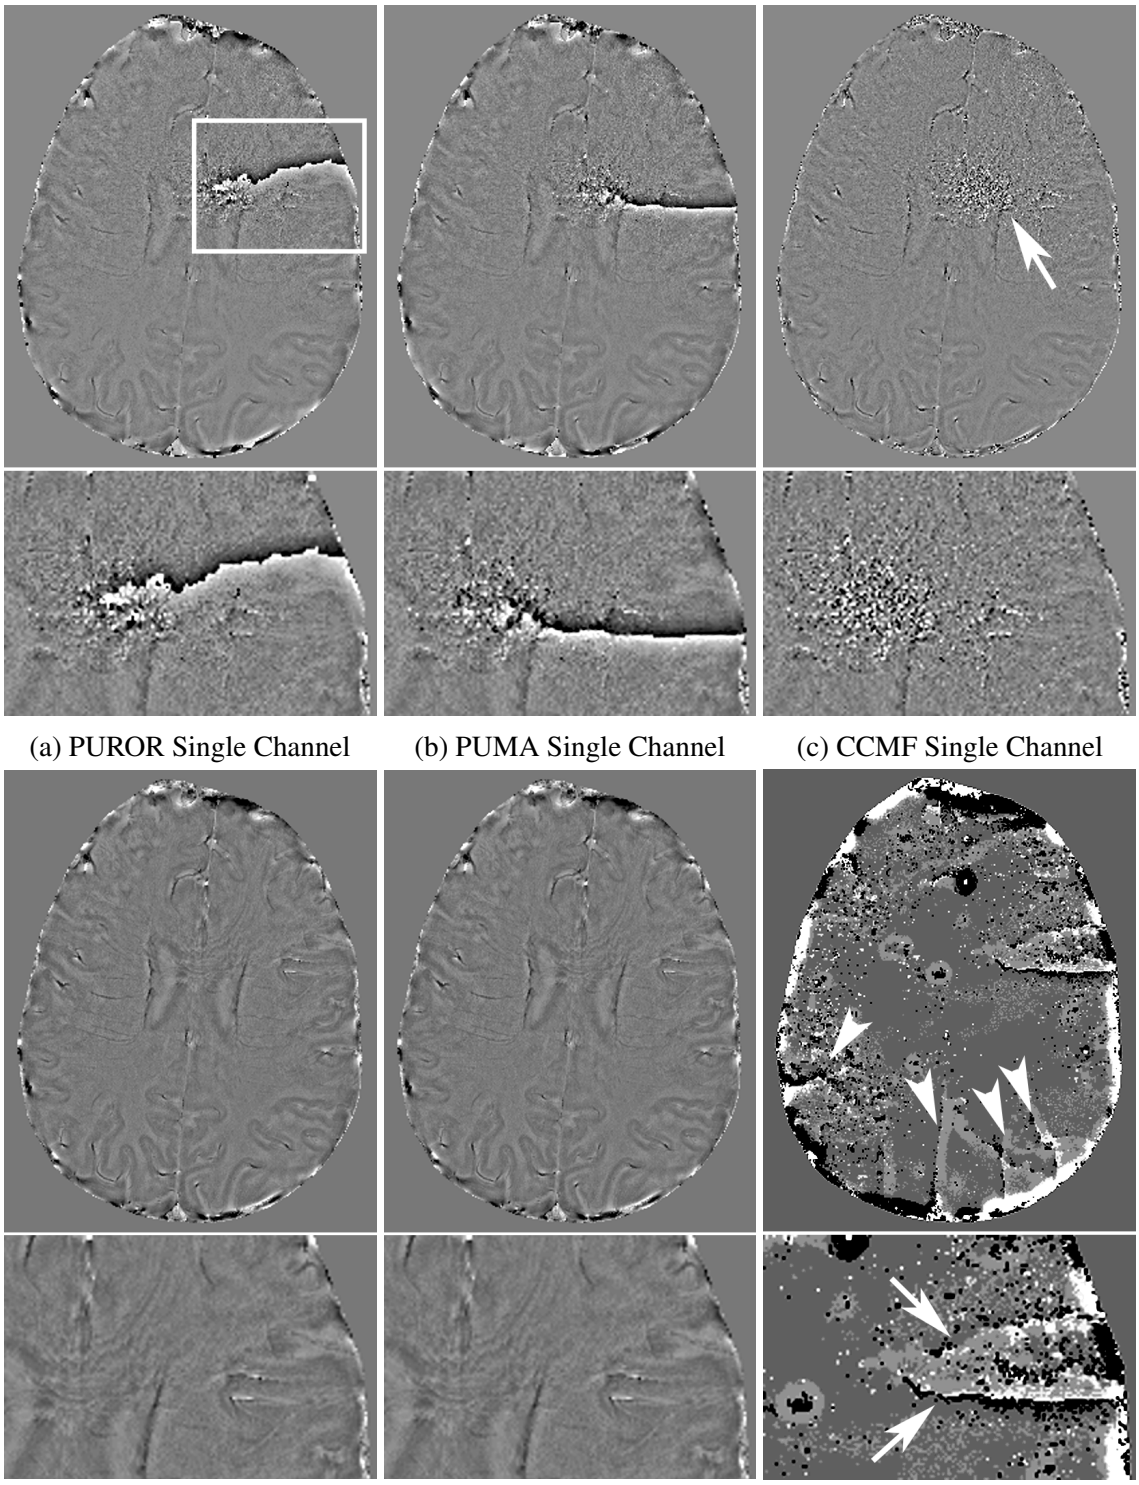

(a) PUROR Single Channel (b) PUMA Single Channel (c) CCMF Single Channel

(d) PUROR Channel Combined (e) PUMA Channel Combined (f) Difference of (d) and (e)

Figure 2.14: Residual phase wrapping artifacts present in channel combined images created using phase unwrapping paradigm.

create issues for other techniques, such as deconvolution, further down the phase processing pipeline as they represent highly spatially correlated error invalidating any assumption of strong independence.

2.7.5 Channel Combined Quantitative Results

Quantitative results for each of the six echos is shown in Figure 2.15. In order to compare each method accurately taking into account variations in the vein contrast and echo, a three-way ANOVA test was performed on the CNR data, the results of which are shown in Table 2.2. The factors in said test were the method (CCMF, homodyne, high pass filtered PUROR, and high

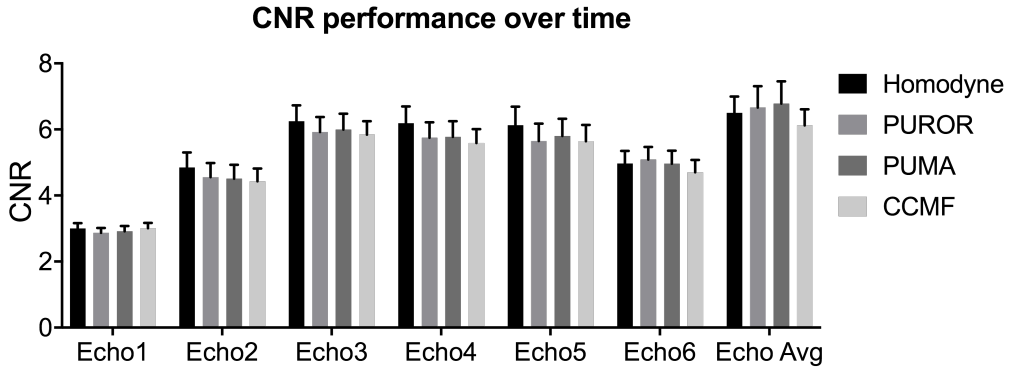

Figure 2.15: CNR performance over echo time

Source	F	p
Method	1.87	0.13
Echo Time	68.58	< 0.001
Vein	38.62	< 0.001

Table 2.2: Three-way ANOVA Table of factors affecting CNR

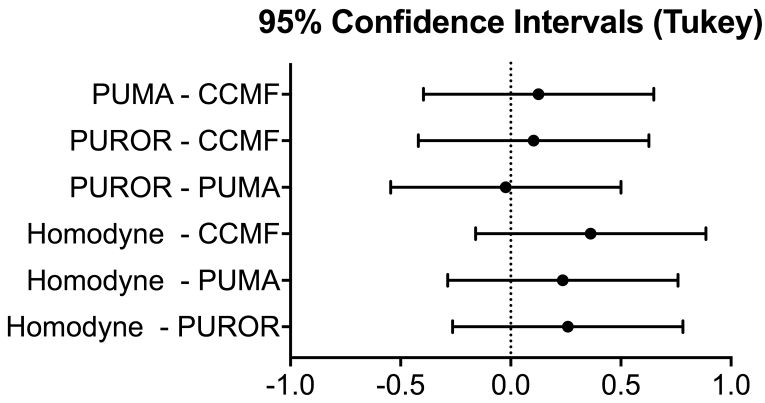

Figure 2.16: Confidence intervals on mean CNR difference between methods

pass filtered PUMA), the echo time, and the vein being measured.

As expected, the vein and echo time had a significant effect on the CNR, validating their inclusion as a factor. Differences between methods were not statistically significant in this analysis, indicating that all four methods compared comparably with no method showing dramatic quantitative improvement. Confidence intervals on the mean CNR differences between methods are shown in Figure 2.16 with a Tukey correction for multiple comparisons. As evidenced by the insignificance of the method as a factor in the three-way ANOVA test, each confidence interval overlaps the origin.

The vessels selected for quantitative evaluation likely had an impact on the results underestimating the performance of CCMF. A number of cortical veins are significantly more salient in the channel combined CCMF images and were almost imperceptible in the comparative approaches. These veins were smaller than 3 mm in diameter and thus did not meet inclusion criteria.

2.8 Discussion

The results have demonstrated the competitiveness of CCMF phase processing to that currently used in clinic and the state-of-the-art used in research settings. This is of particular note because CCMF does not comfortably fit into either existing paradigm for phase processing. That is, it is based neither on processing the complex image nor on processing a phase-only image which is equipped with a Euclidean topology. This inherent topological difference (CCMF using a cyclic, rather than Euclidean, topology) implies that CCMF represents a viable new *third paradigm* in phase processing.

CCMF is novel in another sense in that it is a max-flow based image enhancement algorithm for images with cyclic range topologies, complimenting the Potts and Ishikawa models for discrete and Euclidean topologies respectively. This may allow it to form the basis of models of hue enhancement and reconstruction.

2.8.1 Future Work

In terms of future work, incorporating a model of the relationship between channels can be envisioned using CCMF and may lead to higher contrast filtering. These may allow for filtering and channel combination to be done in tandem using a single optimization-based approach. The goal in this case is to combine information from channels which agree to create a common background phase image applicable to each channel.

Additionally, one especially important area of development would be to create physics-

informed objective functions. One important consideration is *Laplace's equation* (i.e. $\nabla^2 B = 0$) which states that the magnetic field sufficiently distant from any source of magnetic susceptibility is a harmonic function. For an unwrapped, Euclidean phase, this term can be optimized, having the dual form:

$$\int_\Omega R(x)|\nabla^2 u(x)|dx = \max_{|q(x)|\leq R(x)} \int_\Omega u(x)\nabla^2 q(x)dx \tag{2.26}$$

Given that the observed magnetic field superposition is a superposition of those generated from susceptibility sources inside and outside the brain, the background field should satisfy this constraint. An important area of future work would be to formulation this constraint in the context of a cylindrical manifold, incorporating it into CCMF. Such advances will be necessary if CCMF is to be used in the context of QSM as opposed to qualitative SWI.

Chapter 3

Hierarchical Continuous Max-Flow Image Segmentation

This chapter is largely based on:

- Rajchl, Martin, John SH Baxter, A. Jonathan McLeod, Jing Yuan, Wu Qiu, Terry M. Peters, and Ali R. Khan. "Hierarchical max-flow segmentation framework for multi-atlas segmentation with Kohonen self-organizing map based Gaussian mixture modelling." Medical Image Analysis 27 (2016): 45-56.

with additional material from:

- John SH Baxter, Martin Rajchl, Jing Yuan, and Terry M. Peters. "A continuous max-flow approach to general hierarchical multi-labelling problems." arXiv preprint arXiv:1404.0336 (2014).
- John SH Baxter, Martin Rajchl, Jing Yuan, and Terry M. Peters. "A proximal bregman projection approach to continuous max-flow problems using entropic distances." arXiv preprint arXiv:1501.07844 (2015).

3.1 Introduction

The previous chapter concerned itself with solving an image enhancement problem on MRI phase images, equipped with a *cyclic range topology*. The notion of a topology plays a vital role in more than just image enhancement problems. In segmentation problems, one particularly important topological notion is that of a *label ordering* expresses a segmentation problem using the mathematical notion of a *partially ordered set*.

In terms medical image segmentation, the philosophic notion of an *ontology* is starting to take hold [153]. In a radiographic ontology, RadLex [94] in particular, anatomical location is

referred to via segmentation in which the image is decomposed into multiple objects of interest which themselves can be recursively decomposed. That is, the ontology defines what objects are being segmented and a label ordering relating those objects. Label orderings are a powerful way of expressing *part/whole relationships* in anatomical knowledge; they encode knowledge that describes how larger anatomical components are composed of simpler ones. Statements such as "The cerebral cortex has two hemispheres, each having four lobes: the frontal, parietal, temporal, and occipital" shown in Figure 3.1 are an example of this type of knowledge. Generally speaking, part/whole relationships are a beneficial in that they can rely solely on the clinician's abstract knowledge, rather than requiring a pre-annotated or pre-labeled atlas to operate from. This abstract also allows them to have modality independent properties and encourage general use. Because of this general accessibility and level of abstraction, they can form the foundation for a *common language* between clinicians and computers.

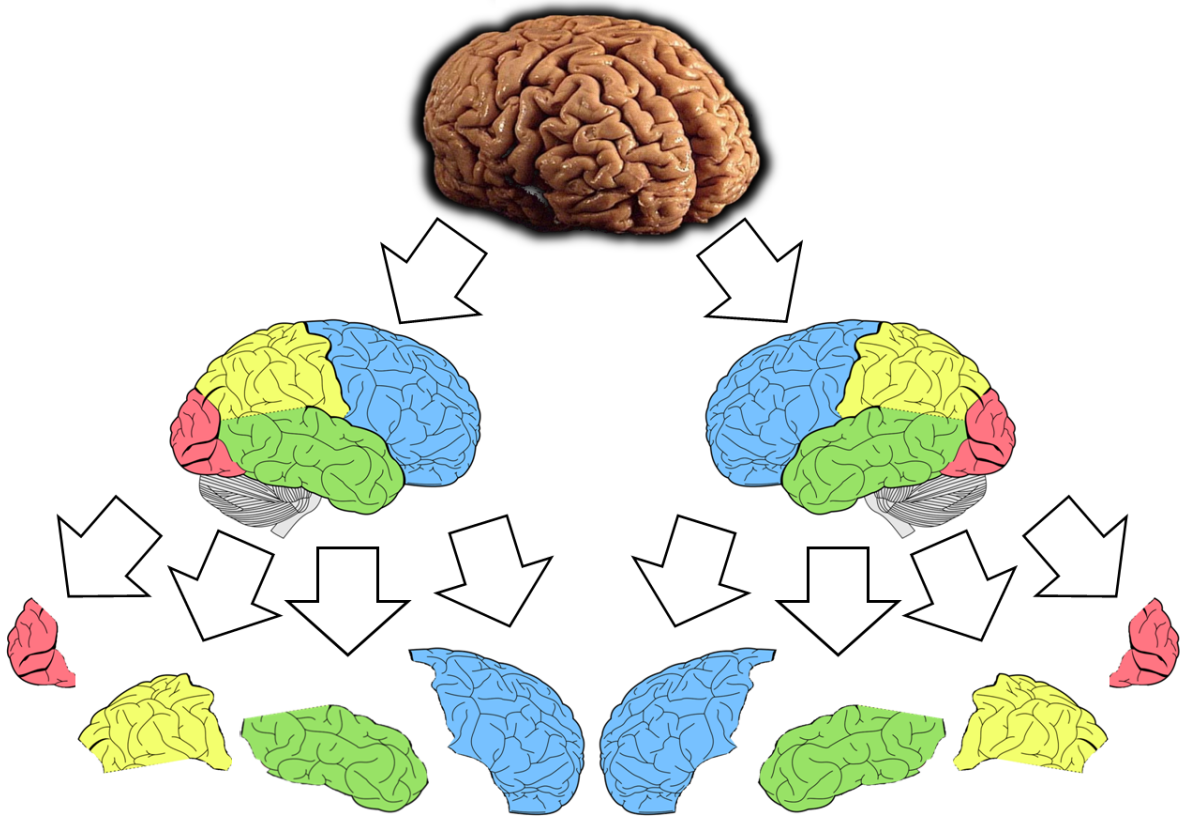

Figure 3.1: Example of anatomical knowledge in the form of *part/whole relationships* expressing the decomposition of the cerebral cortex into hemispheres and then into individual lobes.

That being said, these topological forms can be quite different than the set topologies investigated in the previous chapter. The techniques discussed there, the continuous Potts, Ishikawa, and cyclic models, all express a particular constrained family of label orderings. These constraints mean that these models are not suitable for capturing the spectrum of part/whole rela-

tionships present in multi-region medical segmentation problems. Just as using the incorrect range topology can cause errors and artefacts in image enhancement problems, using an insufficient model can lower the accuracy of a max-flow medical segmentation algorithm.

This chapter begins with an overview of label orderings and hierarchical topologies, their relationship to anatomic knowledge, and how graph-cuts and max-flow techniques have previously incorporated hierarchical topologies. Hierarchical continuous max-flow (HMF) segmentation is presented as an extendable framework for incorporating hierarchical topologies. HMF is then applied to the problem of multi-atlas brain tissue segmentation using two open brain MRI databases.

3.2 Label Orderings and Hierarchical Topologies

Formally, a label ordering is a binary relation defined over the set of labels in a segmentation making that set a *partially ordered set*. I will use the binary relation $\subseteq$ to denote the partial order operator. In order to be a partial ordering, that relation must have the properties of:

- *Reflexivity:* that for all labels, $L \subseteq L$,
- *Antisymmetry:* that if $A \subseteq B$ and $B \subseteq A$ then $A = B$, and
- *Transitivity:* that if $A \subseteq B$ and $B \subseteq C$ then $A \subseteq C$.

A label ordering in particular has the definition that $A \subseteq B$ if and only if $\forall x \in \Omega$, $u_A(x) \leq u_B(x)$. So if a spatial location is assigned the label A and $A \subseteq B$, it must also be assigned the label B as well. Often, one can express a label ordering through diagrammatic notation. Figure 3.2 expresses the label ordering of Figure 3.1 in a more abstract manner.

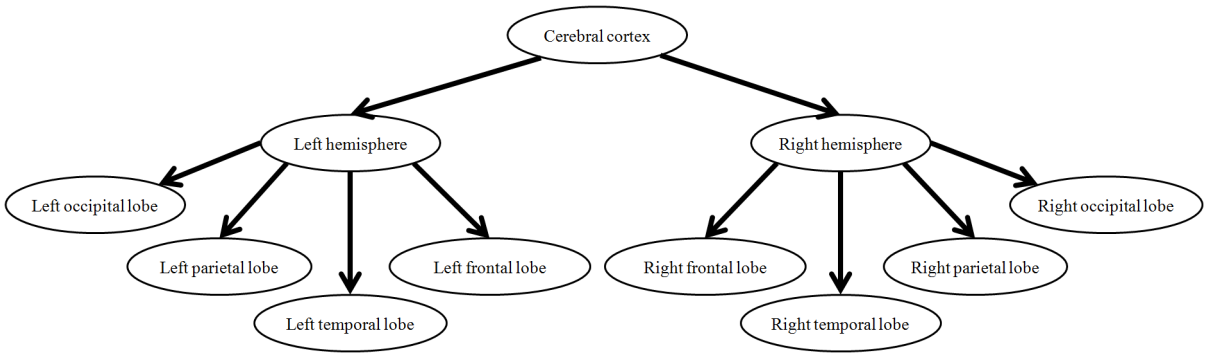

Figure 3.2: Example label ordering from Figure 3.1 shown in diagrammatic form.

In segmentation, to create a valid *hierarchical ordering* or *hierarchy*, these diagrams must be constrained in three ways:

- There must be a root node (with no incoming edges) that represents the entire image, the corresponding label is denoted as S.
- All other nodes must have one and only one incoming edge. The node at the other side of said edge will be called the *parent* of the label.
- Each node will either have no outgoing edges or the outgoing edges must form a partition of said label. The node at the other side of each outgoing edges will collectively be the called the *children* of the label.

One result of these three constraints is that the set of all labels with no children (called *end-labels* or *leaves*) must form a partition of the entire image. In that sense, there is a clear way of expressing this segmentation problem into a partitioning problem.

From an optimization point-of-view, these types of relationships are relatively easy to work with because their can be expressed as a convex constraint on the space of feasible segmentations. Because of this convexity, this constraint can be appended to any convex optimization problem while still maintaining its solvability. That is, if the original space of feasible segmentations is convex, say in fuzzy segmentation problems, the resulting space with these constraints applied is also convex.

The Potts and Ishikawa models introduced in the previous chapter (see Sections 2.2.3 and 2.2.5 respectively) can both be expressed in terms of these hierarchical label ordering diagrams which can elucidate their underlying topologies. In addition, this shows that they are both special cases of the general hierarchical model. In a Potts model, as shown in Figure 3.3a, the N segmentation labels, L_1 to L_N, are all disjoint and arranged as a single partition of the image. In an Ishikawa model, as shown in Figure 3.3b, the N segmentation labels, L_1 to L_N are nested in a particular structure, where each level bifurcates into two labels, the subsequent level and a label which expressed the set difference between the two levels, such as the labels $L_0 - L_1$ and $L_1 - L_2$, etc... This constrained class of label orderings makes the Ishikawa model useful when segmenting linearly nested structures.

However, I would like a *general* solution algorithm that can address any type of hierarchical model, given its hierarchy. The goal of such an endeavour would be a single unified solution framework to a wide array of medical image segmentation problem agnostic to the particular problem being addressed.

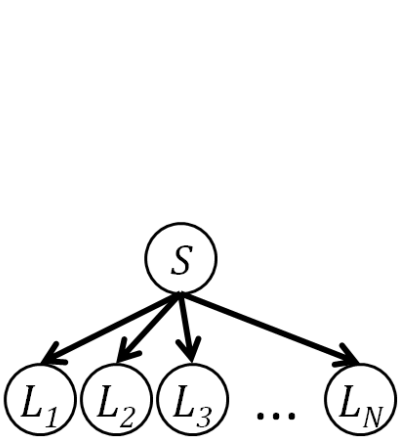

(a) Potts model label ordering

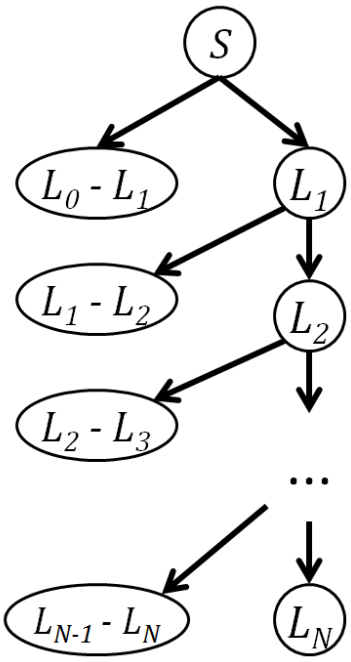

(b) Ishikawa model label ordering

Figure 3.3: Potts and Ishikawa models in terms of label ordering

3.2.1 Label Ordering Operators

Before investigating general hierarchical label ordering solution algorithms, there are five particularly useful operators in using label orderings:

- $L.C$ refers to the children of label L,
- $L.C^*$ refers to the descendants (children, children of children, etc...) of label L, which can be recursively defined as $L.C^* = L.C \cup \{L'' \in L'.C^* | L' \in L.C\} \cup \ldots\}$
- $\mathbb{L}$ is the set of all labels with no children, that is, $L \in \mathbb{L} \implies L.C = \emptyset$,
- $L.P$ refers to the parent of label L, and
- $L.P^*$ refers to the set of ancestors (parents, parents of parents, etc...) of label L, which can be recursively defined as $L.P^* = L.P \cup \{L'' \in L'.P^* | L' \in L.P\} \cup \ldots\}$.

As stated earlier, in a hierarchical ordering, the only constraint is that $L.P$ contains a single element for all labels, excepting the source label S. A concrete example of these operators is given in Figure 3.4.

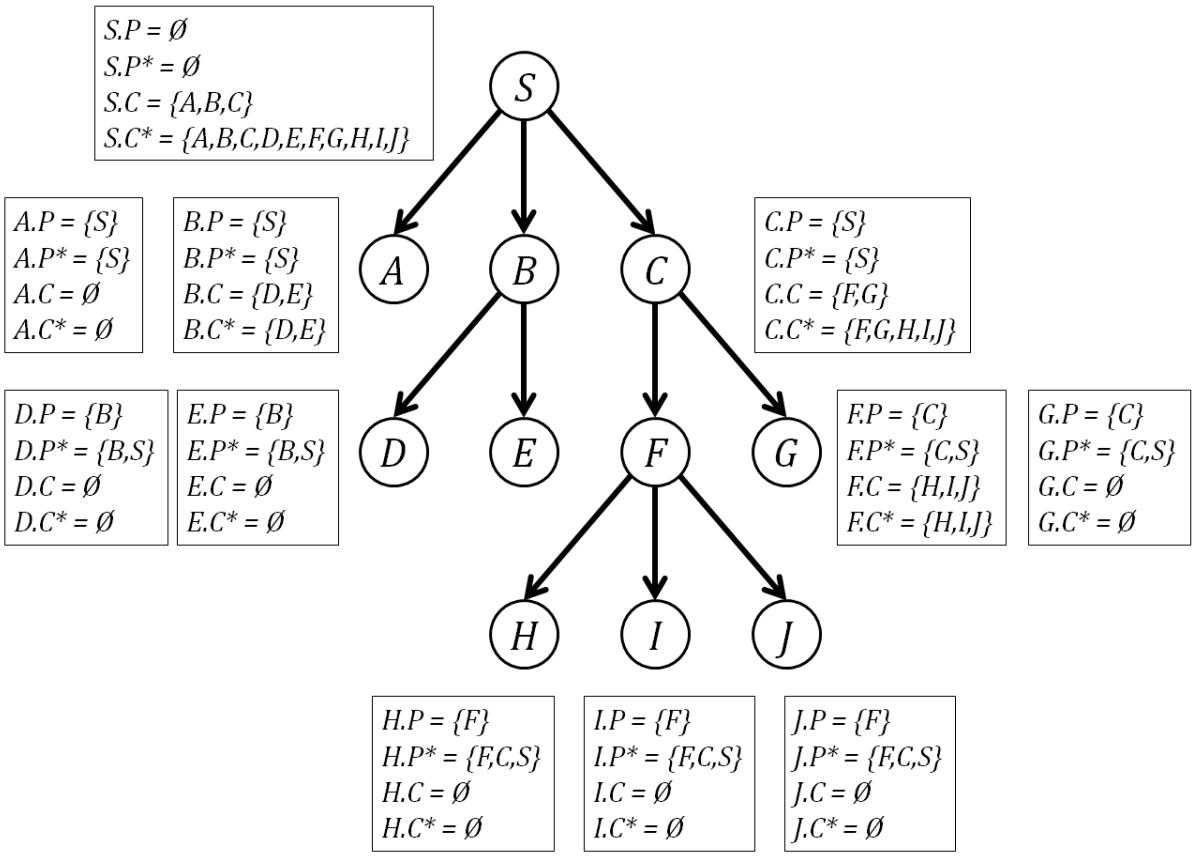

$$\mathbb{L} = \{A, D, E, G, H, I, J\}$$

Figure 3.4: Example of the different label ordering operators

3.3 Previous Approaches to Hierarchical Topologies

3.3.1 Graph-Cuts and the h-Fusion Algorithm

Although the α-expansion algorithm is applicable to discrete hierarchical models in theory (by creating a partitioning problem across the set of end-labels $\mathbb{L}$ with a complex regularization term), the guarantee of approximate optimality becomes weaker and weaker. Thus, an alternative solver is required for these hierarchical problems which takes advantage of the hierarchical structure of the label ordering.

Delong *et al.* [44] developed the h-fusion algorithm which uses the α-expansion algorithm as a subroutine. In the h-fusion algorithm, a sequence of multi-label subproblems are defined at each iteration which are then merged together. These subproblems are designed to maintain the hierarchical structure of the label ordering. In particular, each subproblem is created from partitioning each label into its children. The approximation guarantee is greatly improved for

hierarchical costs over directly using α-expansion over the labels in $\mathbb{L}$.

3.3.2 Hard-Coded Hierarchies in Continuous Max-Flow

The concept of hierarchical label orderings have been used in max-flow segmentation before under the moniker of the *partially ordered Potts model* for cardiac MRI segmentation. Rajchl *et al.* [150] used a hierarchical label ordering as shown in Figure 3.5a as an extension to their earlier Ishikawa-based approach [149] shown in Figure 3.5b.

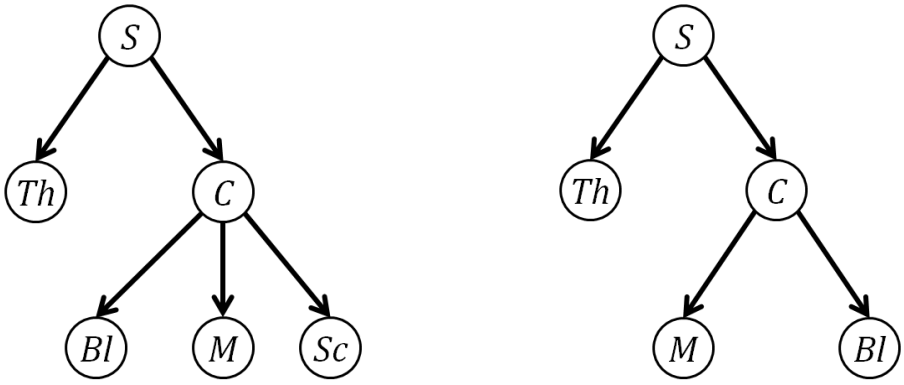

(a) Hierarchical label ordering used by Rajchl *et al.* [150]

(b) Ishikawa label ordering used by Rajchl *et al.* [149]

Figure 3.5: Label orderings used in cardiac segmentation by Rajchl *et al.* [149, 150]

The issue with this approach was not the label ordering but that the solution algorithm proposed was tailored solely to that particular label ordering and thus lacked extendibility and generality outside of its particular problem domain.

3.3.3 Gestalt Computer Vision

A fundamentally different approach to hierarchical label orderings has been inspired by Gestalt psychology of vision. The *Gestalt principles* or *grouping laws* are a structure for combining together smaller units into larger wholes. In addition, these larger wholes are interpretted via a *figure-ground* divide, that each whole can be decomposed into foreground components against a common background. These principles include:

1. *Connectedness:* units that are visually connected to each other are more likely to be a part of a whole than disconnected regions.
2. *Proximity:* units that are close together are more likely to be part of a larger whole than those which are further apart,
3. *Similarity:* units that share visual properties such as brightness, colour, and texture are likely to be part of a larger whole than those which do not share these properties,

4. *Closure:* units should be joined together in such a way that their contours enclose a region, that 'imaginary' edges can be created in order to ensure a complete boundary.
5. *Good continuance:* contours or lines tend to be smooth and continuous, thus 'X' shapes are more likely to be two lines crossing as opposed to two sharp objects touching at a point.
6. *Common Region:* units that occupy a common region or ground are more likely to be part of a larger whole than those in different regions, and
7. *Symmetry:* contours that are symmetric are more likely to be a part of a common object than contours which are not.

These principles can interact and contradict each other, but have led to the rigorous development of computer vision algorithms in which structure is extracted from the image such as super-pixel hierarchies and agglomerative clustering in computer vision [158].

The Gestalt approach is complementary to but distinct from the other previous approaches and that of this chapter to incorporating hierarchical topological information in that it is fundamentally more *bottom-up* as opposed to *top-down*. That is, the Gestalt grouping principles largely concern how to determine structure given the image, whereas in this chapter and prior work such as h-fusion [44], the hierarchical structure is already given and is imposed upon the image. The hierarchical structure in that sense is largely semantic and is a result of more abstract anatomical knowledge.

3.4 Hierarchical Continuous Max-Flow Formulation

In *hierarchical continuous max-flow* (HMF), the energy functional can be expressed as the minimization problem:

$$\begin{aligned} &\min_{u} \sum_{L\in\mathbb{L}} \int_{\Omega} D_L(x) u_L(x) dx + \sum_{L} \int_{\Omega} |\mathbf{R}_L(x) \nabla u_L(x)|_p dx \\ &\text{s.t. } \forall L \in \mathbb{L},\ u_L(x) \geq 0 \\ &\qquad \forall L \notin \mathbb{L},\ u_L(x) = \sum_{L'\in L.C} u_{L'}(x) \\ &\qquad u_S(x) = 1 \end{aligned} \tag{3.1}$$

which may be subject to the integrality constraint $u_L(x) \in \{0, 1\}$, however, such a constraint renders the energy minimization problem NP-hard as it is a strict generalization of the (NP-hard) Potts model. Each label is associated with a regularization term, but only end-labels have both a data and regularization term. This is because any data term associated with a non-end-label can be pushed down to its children, such a process can be repeated iteratively until

reaching the end-labels and thus the problems are equivalent. That is:

$$\begin{aligned} &\text{if } L.C \neq \emptyset \text{ then } \tilde{D}_L(x)u_L(x) = \sum_{L' \in L.C} \tilde{D}_L(x)u_{L'}(x) \\ &\text{therefore } \sum_L \int_\Omega \tilde{D}_L(x)u_L(x)dx = \sum_{L \in \mathbb{L}} \int_\Omega D_L(x)u_L(x)dx \\ &\text{where } D_L(x) = \tilde{D}_L(x) + \sum_{L' \in L.P^*} \tilde{D}_{L'}(x) \end{aligned} \tag{3.2}$$

Because the hierarchy has been divorced from the source code for the solution algorithm, HMF segmentation methods can be evaluated with a much higher throughput, and allows for the method to be extremely general-purpose. For example, in a segmentation research environment, one would have to create an entirely new solver for every hierarchical method investigated, posing a significant barrier to the evaluation of multiple methods. Because of the general-purpose and hierarchy-agnostic code, the process for determining the hierarchy is separate from the solution module, which would decrease the amount of complex and error-prone programming required. This additional modularity would allow for multiple hierarchy modules feeding into a single solver, greatly reducing the development complexity.

There are a few other considerations to take into account when designing a label ordering for any particular segmentation problem. These considerations, as well as some basic combinatorial and theoretical considerations when using hierarchical label orderings are presented in Appendix B. One important result in this appendix is that determining the optimal hierarchical label ordering given *a priori* knowledge of what groups of end-labels may be form a super-label is NP-hard (See Section B.3), implying that user knowledge in defining hierarchies is unavoidable for computational reasons.

3.5 Hierarchical Continuous Max-Flow Solution Algorithms

The HMF energy equation can be minimized using primal-dual optimization on a particular flow network which is derived from the label ordering diagram by replacing each label $L \neq S$ with a continuum, connecting those $L \in \mathbb{L}$ to a sink node. Those sink node connections are then constrained by the data terms. An example of this process is shown in Figure 3.6. The algorithms presented in this section are based on the maximization of flow through these particular flow networks. The main contribution of these *generalized* solution algorithms in particular is that the label ordering diagram, in the form of a tree data structure, can be given to them in run time rather than having been previously hard-coded as in Section 3.3.2.

The augmented Lagrangian solution algorithm to the HMF formulation is given in Algo-

rithm 3.1. As with the previous augmented Lagrangian algorithms, it involves two positive parameters: τ, the gradient descent step size set to approximately $1/8$; and c, the quadratic penalty constant which has a default value of $1/4$.

The proximal Bregman solution algorithm to the HMF formulation Algorithm 3.2. Similar to the augmented Lagrangian algorithm, τ is the gradient descent step size used in the Chambolle iteration step which is set to approximately $1/8$. However, c is now the weight of the proximity term in each projection step which is set to be on the order of the difference in data terms.

For a full theoretical proof of correctness of these algorithms and any additional technical information, see Appendix A.2.

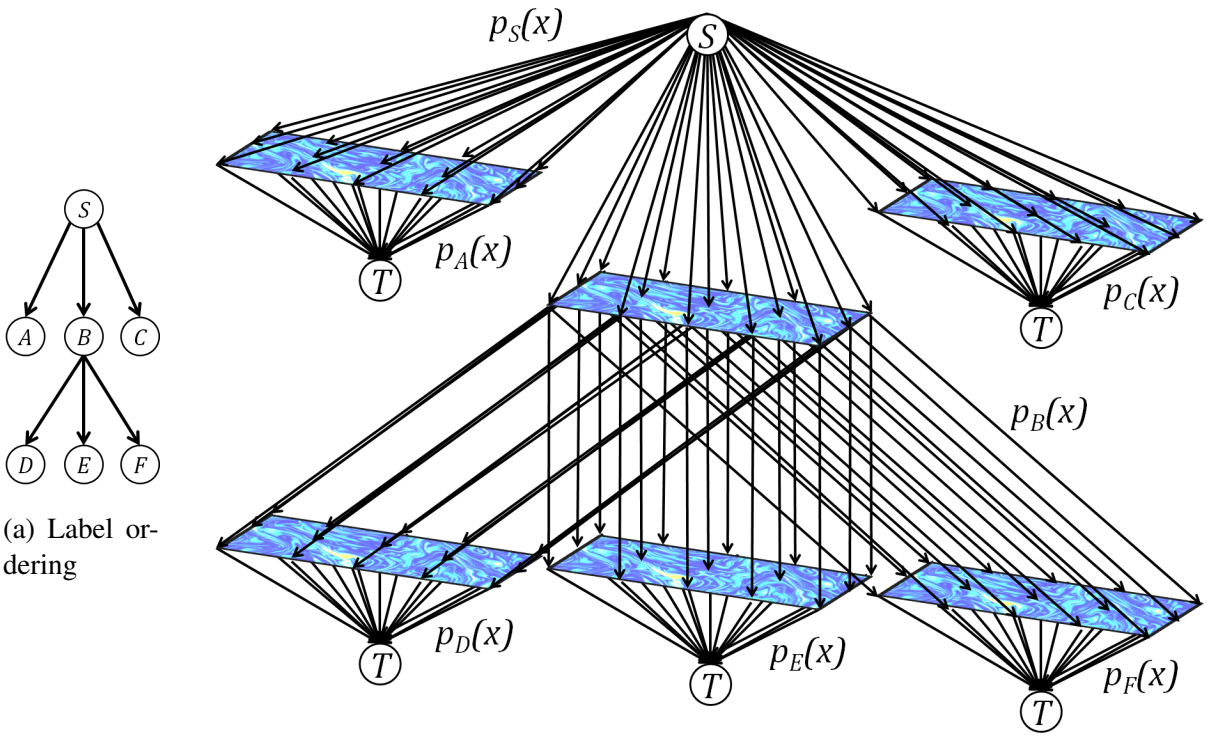

(a) Label ordering

(b) Corresponding flow network

Figure 3.6: Example of transforming a hierarchical label ordering into a flow network.

Algorithm 3.1: Augmented Lagrangian solution algorithm for the HMF functional.

InitializeSolution(S) ;
while *not converged* **do**
  for $\forall L \neq S$ **do**
    $\forall x, q_L(x) \leftarrow \text{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} \left(q_L + \tau \nabla \left(\text{div}\, q_L(x) + p_L(x) - p_{L.P}(x) - u_L(x)/c \right) \right)$;
  end
  UpdateFlows(S) ;
  for $\forall L \neq S$ **do**
    $\forall x, u_L(x) \leftarrow u_L(x) - c \left(\text{div}\, q_L(x) - p_{L.P}(x) + p_L(x) \right)$;
  end
end

InitializeSolution(L)
for $\forall L' \in L.C$ **do**
  InitializeSolution(L');
end
$\forall x,\ p_L(x) \leftarrow 0$;
if $L \neq S$ **then**
  $\forall x,\ u_L(x) \leftarrow 0$;
  $\forall x,\ q_L(x) \leftarrow 0$;
end

UpdateFlows(L)
for $\forall L' \in L.C$ **do**
  UpdateSinkFlows(L') ;
end
if $L \in \mathbb{L}$ **then**
  $\forall x, p_L(x) \leftarrow \min\{D_L(x), p_{L.P}(x) - \text{div}\, q_L(x) + u_L(x)/c\}$;
else if $L = S$ **then**
  $\forall x, p_S(x) \leftarrow 1/c$;
  for $\forall L' \in S.C$ **do**
    $\forall x, p_S(x) \leftarrow p_S(x) + p_{L'}(x) + \text{div}\, q_{L'}(x) - u_{L'}(x)/c$;
  end
  $\forall x, p_S(x) \leftarrow \frac{1}{|S.C|} p_S(x)$;
else
  $\forall x, p_L(x) \leftarrow p_{L.P}(x) - \text{div}\, q_L(x) + u_L(x)/c$;
  for $\forall L' \in L.C$ **do**
    $\forall x, p_L(x) \leftarrow p_L(x) + p_{L'}(x) + \text{div}\, q_{L'}(x) - u_{L'}(x)/c$;
  end
  $\forall x, p_L(x) \leftarrow \frac{1}{|L.C|+1} p_L(x)$;
end

Algorithm 3.2: Proximal Bregman solution algorithm for the HMF functional.

$\forall L \in \mathbb{L},\ u_L(x) \leftarrow 1/|\mathbb{L}|;$
$\forall x,\ d_S(x) = 0;$
while *not converged* **do**
  PushDownCosts(S);
  $\forall x, L \in \mathbb{L},\ d_L(x) \leftarrow u_L(x) \exp\left(-\frac{d_L(x)}{c}\right);$
  $\forall x, L \in \mathbb{L},\ u_L(x) \leftarrow d_L(x) / \sum_{L' \in \mathbb{L}} d_{L'}(x);$
  $\forall x, L \notin \mathbb{L},\ d_L(x) \leftarrow 0;$
  PushUpCapacities(S);
end

PushDownCosts(L)
if $L \in L$ **then**
  $\forall x,\ d_L(x) \leftarrow d_{L.P}(x) + \operatorname{div} q_L(x) + D_L(x);$
else if $L \neq S$ **then**
  $\forall x,\ d_L(x) \leftarrow d_{L.P}(x) + \operatorname{div} q_L(x);$
end
for $\forall L' \in L.C$ **do**
  PushDownCosts(L');
end

PushUpCapacities(L)
for $\forall L' \in L.C$ **do**
  PushUpCapacities(L');
  if $L \neq S$ **then**
    $\forall x,\ q_L(x) \leftarrow \operatorname{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} \left(q_L(x) - c\tau \nabla d_L(x)\right);$
    $\forall x,\ d_{L.P}(x) \leftarrow d_{L.P} + d_L(x);$
  end
end

3.6 Hierarchical Continuous Max-Flow in Brain Tissue Segmentation

The segmentation of the brain with the assistance of an atlas or multi-atlas has been the mainstay of neuro-imaging research, especially in cross-patient studies using techniques such as volumetric morphometry [76, 97], voxel-based morphometry [76, 119] and relaxometry [86]. In these population studies, an atlas is used to provide corresponding anatomical regions between different patients, allowing for population analysis to occur specific to particular regions of interest. For these studies to be effective, they must contain multi-atlas-based registration and segmentation protocols that are both efficient and robust to pathology.

Brain tissue segmentation is one of the simplest delineations of regions of interest, but an important one as it reflects the major anatomical delineations in the brain which can be used to constrain more fine-grained parcellations. These segmentation problems contain sufficient inherent complexity with readily available anatomical knowledge in the form of part-whole relationships that label orderings can come into play, encoding this knowledge in the segmentation problem. Thus a multi-atlas-based brain tissue segmentation framework has been implemented for both the OASIS and MRBrainS databases associated with two consecutive MICCAI segmentation challenges.

3.6.1 MICCAI 2012 OASIS Images

The Open Access Series of Imaging Studies (OASIS) database [118] contains a cross-sectional collection of over 400 T1w MR images acquired at 1.5T using a magnetization prepared rapid gradient-echo (MP-RAGE) sequence. A subset of this database was used for the MICCAI 2012 Grand Challenge on Multi-Atlas Labelling which provided 15 manually segmented T1w images for method development and training, with 20 testing T1w images. Although originally containing a large list of anatomical labels, comparative criteria were also produced for the subproblems of tissue segmentation (between gray matter, white matter, and cerebro-spinal fluid) and basic structure segmentation (cortical gray matter, subcortical gray matter, and brain stem). The list of labels used in brain tissue segmentation on this database were:

- K representing the background,
- cGM representing cortical gray matter,
- sGM representing subcortical gray matter,
- V representing the ventricles,
- WM representing white matter, and
- BS representing the cerebellum and brain stem.

3.6.2 MICCAI 2012 OASIS Methods

The segmentation framework developed for the OASIS database (summarized in Figure 3.7) contains multiple components, but revolves around creating a continuous max-flow segmentation model with three distinct components:

- An *intensity term* which encourages segmented objects to have the same intensity distribution as equivalent objects in the *a priori* segmented atlas images,
- A *spatial term* which encourages segmented objects to be in the same general location as equivalent objects in the *a priori* segmented atlas images, and
- A *smoothness term* which encourages smooth segmentation boundaries aligned with visible edges in the image.

In terms of the max-flow model applied, either the Potts model or the HMF model shown in Figure 3.8 was used.

The intensity term was created through Gaussian Mixture Modeling (GMM) using dense Kohonen self-organizing mixture networks (KSOMNs). [186] These KSOMNs have inherent intensity-domain regularization which makes them better suited for modelling more irregularly shaped distributions. The GMM was learned on a *synthetic* multichannel image with ten scalars: the T1w images from the atlas, three smoothed versions of the T1w image (Gaussian kernels of $\sigma = 1, 2, 3\,\text{mm}^3$, the gradient magnitude of said smoothed images, and the Laplacian of said smoothed images. Each image was normalized to unit standard deviation within the brain mask to ensure consistency.

The spatial term was created through the registration of the image to each of the images in the atlas first via an affine transformation achieved using NiftyReg [131] and then through the deformable registration framework RANCOR [147]. To merge the resulting registration-based labellings, two label fusion methods were investigated. *Mean label fusion* (MLF) simply averages the indicator functions; for example, if three of five atlases label a particular location as grey matter, MLF marks that location as being 60% grey matter. *Joint label fusion* (JLF) [181] is more complex, weighing each atlas by the joint probability of multiple atlases making incorrect labellings simultaneously. This probability is estimated from the difference in intensities in a local neighbourhood between the atlases and the target image, corresponding to a local intensity model.

The intensity and spatial terms can be combined together into a single data term for the model using the Bayesian formula:

$$D_L(x) = -\ln P(I(x)|x \in L) - \beta_L \ln P(x \in L) \tag{3.3}$$

where the first term is the negative log-likelihood of the probability determined by the GMM, and the second is that determined by the label fusion algorithm, β_L weighting their relative importance. The smoothness terms were derived from the gradient magnitude of the image normalized by that of the smoothed image:

$$R_L(x) = \alpha_L \exp\left(\frac{-\lambda|\nabla I(x)|}{1 + |\nabla(G * I(x))|}\right) \tag{3.4}$$

where G is a Gaussian kernel of standard deviation 1 mm^3.

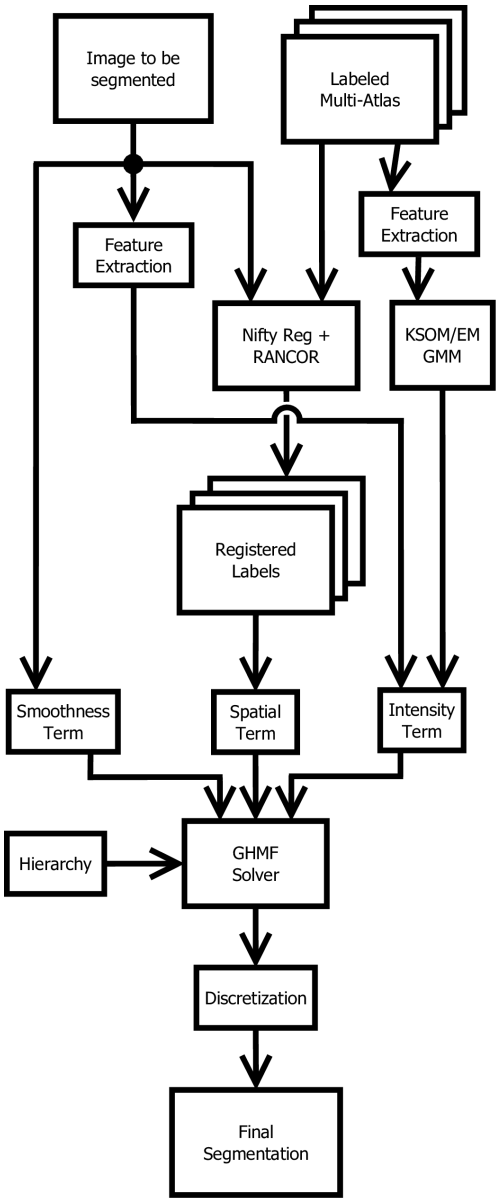

Figure 3.7: Pipeline used in multi-atlas brain tissue segmentation.

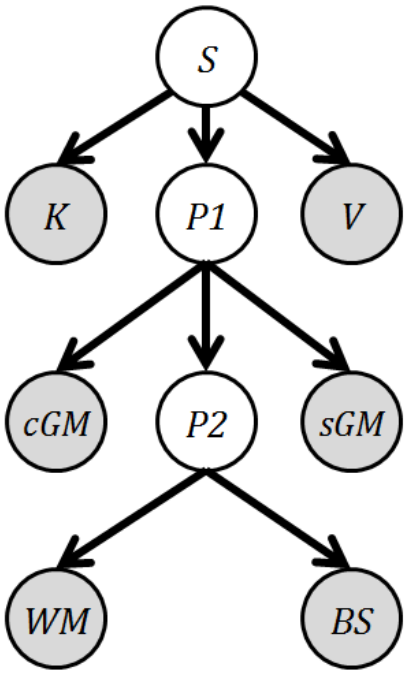

Figure 3.8: Hierarchical label ordering used in segmentation of the OASIS database

In order to quantify segmentation accuracy, three complimentary metrics were used:

- *Dice similarity coefficient* (DSC) which measures the normalized amount of volumetric overlap between segmentations, ranging between 0% indicating no spatial agreement to 100% indicating the equivalence of the segmentations,
- *Absolute volume difference* (AVD) which measures the normalized difference in the volume of the segmentations ranging between 0% indicating that the segmentations have the same volume to 100% indicating that the segmentations have completely incompatible volume estimates,
- *Modified Hausdorff distance* (MHD) which measures the proximity of the segmented surfaces, 0 mm indicating the equivalence of the segmentations.

Competitive methods were selected from the official MICCAI 2012 OASIS segmentation challenge [92]. The top seven methods include:

- *CIS JHU* presented by Tang *et al.* uses a large deformation diffeomorphic metric mapping (LDDMM) framework that iteratively optimizes atlas selection using an expectation-maximization algorithm.
- *MALP EM* presented by Ledig *et al.* extends their previous multi-atlas label propagation (MALP) [96] framework using expectation-maximization to refine their local intensity model.

- *MAPER* presented by Heckemann *et al.* iteratively re-registers atlases based on their agreement in terms of label fusion.
- *NonLocal STAPLE* presented by Asman and Landman [10] extends the STAPLE algorithm [184] to include non-local means-based intensity information.
- *PICSL BC* and *PICSL Joint* both proposed by Wang *et al.* [180] use joint label fusion and corrective learning to account for label fusion errors.
- *STEPS* proposed by Cardoso *et al.* [30] uses locally normalized cross-correlation to extend the STAPLE algorithm [184] to introduce local intensity information.

These competitive methods provide context for what is reasonable in terms of state-of-the-art multi-atlas brain tissue segmentation algorithms. Evaluation was performed on the basis of a set of paired t-tests for each evaluation metric. Paired t-tests allow for differences between methods to be extracted even when there is a large amount of variability in the datasets. By using a paired approach, performance variations across datasets are controlled for, isolating the variability resulting from differing methods.

3.6.3 MICCAI 2012 OASIS Results

The segmentation framework was run on the remaining images in the OASIS database to obtain segmentation results and accuracies. Six variants of the algorithm were applied:

1. *MLF+IM+Potts*: Mean label fusion with the additional intensity model and Potts regularization,
2. *MLF+IM+HMF*: Mean label fusion with the additional intensity model and HMF regularization,
3. *JLF+Potts*: Joint label fusion and Potts regularization and no additional intensity model,
4. *JLF+HMF*: Joint label fusion and HMF regularization and no additional intensity model,
5. *JLF+IM+Potts*: Joint label fusion with the additional intensity model and Potts regularization, and
6. *JLF+IM+HMF*: Joint label fusion with the additional intensity model and HMF regularization.

Each experiment has a paired Potts vs. HMF component, elucidating under what conditions more expressive regularization leads to improved segmentation performance. In addition, there is a pair to determine if JLF truly outperforms MLF given an auxilliary intensity model, and whether or not that intensity model improves JLF. Quantitative segmentation results are shown in Figure 3.1. From these results, two methods were selected as having superior performance, *JLF+HMF* and *JLF+IM+HMF*. The latter improved segmentation results marginally, but not

significantly. This may be because the images are univariate so the intensity model did not contain much more information than the local similarity model used by JLF.

Visual results for the best and worst case images using *JLF+IM+HMF* are shown in Figure 3.9. In the best case, the segmentation results adhere very closely to the manual segmentation results with some possible slight over-smoothing of the subcortical gray matter. In the worst case images, the presence of white matter lesions greatly reduced the performance of the segmentation algorithm, even without an *a priori* intensity model assuming their non-existence. This indicates that the lesions were present with sufficient size and hypointensity to cause large registration errors that could not be compensated for by the segmentation model. In addition, the feathering artefacts in the manual segmentation of the subcortical greatly decreases apparent segmentation performance due to manual segmentation variability. The increased smoothing applied to limit these artefacts has caused some erroneous combination of otherwise disconnected regions of the subcortical gray matter.

In terms of comparing with methods found in the literature, *JLF+IM+HMF* did quite well. Comparative segmentation accuracies on the seven highest ranked segmentation algorithms on the OASIS database challenge are shown in Tables 3.2, 3.3, and 3.4. Our method performed particularly well using the MHD metric on the ventricle component. This is largely due to the way in which parameters were manually selected for this problem. As the ventricle tended to have high edge contrast, it required less manual effort to improve performance via parameter optimization than it would for more convoluted or blurrred regions. These results indicate that our method would have likely been placed in the top five methods in the challenge in terms of brain tissue and basic structure segmentation, all of which are extremely comparable. Given that each image in the dataset had a single manual segmentation as a gold standard, it is impossible to quantify manual segmentation variability for this dataset. However, viewing the manual segmentations provided as training, one can readily see feathering artifacts on what should be smooth regions which tend to be indicative of higher manual segmentation variability, as they indicate that only 2D segmentation tools were provided and that no standard operating procedure was created to mitigate for out-of-plane errors like feathering.

OASIS	cGM	sGM	WM	V	BS	mean
DSC (%)						
MLF+IM+Potts	86.0±4.7	85.5±4.3	89.7±4.0	87.4±3.6	89.1±2.2	87.5±3.8
MLF+IM+HMF	**86.4±4.4**	**86.1±3.8**	**89.9±3.5**	**87.5±3.6**	89.1±2.3	**87.8±3.5**
AVD (%)						
MLF+IM+Potts	7.0±5.6	8.9±8.4	11.6±11.8	14.3±10.6	8.2±5.8	10.0±8.5
MLF+IM+HMF	**5.7±4.6**	**7.2±7.3**	**10.2±10.9**	14.7±10.5	**6.9±5.9**	**8.9±7.9**
MHD (mm)						
MLF+IM+Potts	2.2±0.3	2.0±0.5	1.6±0.4	**1.8±0.6**	2.2±0.5	1.9±0.5
MLF+IM+HMF	2.2±0.2	2.0±0.5	1.6±0.4	1.9±0.6	2.2±0.4	2.0±0.4
DSC (%)						
JLF+Potts	87.8±3.4	89.5±2.1	91.9±3.2	89.9±3.7	91.1±0.7	90.0±1.7
JLF+HMF	**88.5±3.5**	**90.2**±1.9	91.2±3.9	90.7±3.3	**93.0±0.8**	90.7±1.4
JLF+IM+Potts	88.4±3.4	*91.1±1.8	92.5±2.7	*91.1±3.1	*93.0±0.9	*91.2±1.5
JLF+IM+HMF	**88.9±3.3**	***91.1±1.8**	**92.6±2.7**	91.2±3.0	93.0±0.9	**91.3±1.5**
AVD (%)						
JLF+Potts	3.7±3.3	6.0±2.4	7.8±8.2	6.9±5.1	3.5±2.0	5.6±2.5
JLF+HMF	3.6±3.6	**3.5±2.3**	7.3±9.2	6.7±6.3	2.1±2.0	4.7±2.3
JLF+IM+Potts	4.2±2.8	*3.5±2.0	7.7±7.1	7.0±5.3	2.7±2.0	5.0±2.3
JLF+IM+HMF	3.8±2.5	3.2±1.9	7.5±7.0	6.8±5.6	2.7±2.1	4.8±2.2
MHD (mm)						
JLF+Potts	2.4±0.4	1.7±0.4	1.4±0.6	1.4±0.3	2.0±0.2	1.8±0.3
JLF+HMF	**2.2±0.4**	**1.5±0.3**	1.7±1.0	1.4±0.3	**1.6±0.3**	1.7±0.3
JLF+IM+Potts	*2.2±0.4	*1.3±0.3	1.2±0.3	*1.2±0.2	*1.6±0.3	*1.5±0.2
JLF+IM+HMF	***2.0±0.3**	*1.3±0.2	1.2±0.4	*1.2±0.2	1.5±0.3	***1.4±0.2**

Table 3.1: Segmentation Results - OASIS: *significantly better metrics ($p \leq 0.05$ after Bonferroni correction) between HMF/Potts pairs are shown in bold and significantly better metrics ($p \leq 0.05$) with/without the intensity model are denoted with an asterix for JLF.*

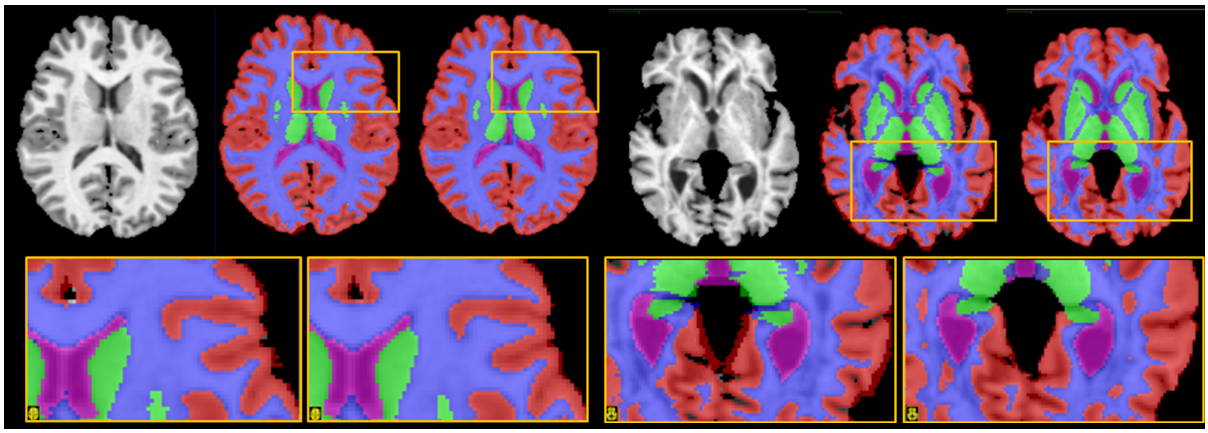

Figure 3.9: Best and worst case results for the OASIS database (top row: best case T1w image, gold standard, JLF+HMF, worst case T1w image, gold standard, JLF+IM+HMF; bottom row: enlarged ROIs.)

OASIS *DSC (%)*	cGM	sGM	WM	V	BS	mean
JLF+IM+HMF	88.9±3.3	91.1±1.8	92.6±2.7	91.2±3.0	93.0±0.9	91.4±1.5
CIS JHU	**92.5±1.5**	**89.5±2.3**	92.8±1.2	**88.1±3.3**	92.6±0.7	91.1±1.0
MALP EM	**92.8±2.2**	90.6±2.8	93.9±2.0	90.2±3.4	93.3±0.7	92.2±1.4
MAPER	91.0±1.8	90.3±3.0	92.6±1.5	**89.8±3.3**	93.4±0.9	91.4±1.2
NonLocalSTAPLE	91.6±3.4	91.3±2.3	93.6±3.5	90.6±2.8	**94.2±0.8**	92.2±2.0
PICSL BC	**92.6±2.1**	91.3±2.3	93.4±1.8	90.9±3.1	**94.6±0.5**	**92.6±1.3**
PICSL Joint	91.2±2.3	91.2±2.3	92.5±1.8	**89.7±3.5**	**94.5±0.5**	91.8±1.4
STEPS	90.2±1.7	90.4±2.2	92.0±1.1	**88.9±3.6**	93.8±0.7	91.1±1.0

Table 3.2: Comparison Results - OASIS *DSC*: *significant difference to JLF+IM+HMF ($p \leq 0.05$ after Bonferroni correction) is shown in bold*

OASIS *AVD (%)*	cGM	sGM	WM	V	BS	mean
JLF+IM+HMF	3.8±2.5	3.2±1.9	7.5±7.0	6.8±5.6	2.7±2.1	4.8±2.2
CIS JHU	3.0±2.7	2.5±2.1	4.0±2.9	9.0±5.7	2.6±1.8	4.2±2.2
MALP EM	4.8±2.8	2.3±2.4	5.2±2.8	7.7±6.4	2.1±1.4	4.4±1.8
MAPER	5.3±2.9	2.4±2.4	9.8±3.2	7.9±6.2	3.1±1.8	5.7±1.6
NonLocalSTAPLE	4.5±5.3	3.3±1.9	7.6±6.9	7.7±6.1	2.0±1.4	5.0±2.5
PICSL BC	7.2±4.6	2.5±2.6	9.2±3.3	5.7±5.2	2.1±1.5	5.3±2.1
PICSL Joint	5.8±3.7	2.3±1.5	9.2±3.8	8.4±5.1	2.1±1.5	5.5±2.2
STEPS	3.1±2.2	5.4±3.0	4.1±2.8	6.0±4.4	2.1±1.4	4.1±1.4

Table 3.3: Comparison Results - OASIS *AVD*: *significant difference to JLF+IM+HMF ($p \leq 0.05$ after Bonferroni correction) is shown in bold*

OASIS *MHD (mm)*	cGM	sGM	WM	V	BS	mean
JLF+IM+HMF	2.0±0.3	1.3±0.2	1.2±0.4	1.2±0.2	1.5±0.3	1.4±0.2
CIS JHU	**1.4±0.2**	**1.7±0.5**	1.2±0.2	**2.7±1.1**	1.6±0.3	1.7±0.3
MALP EM	**1.2±0.3**	1.4±0.6	1.1±0.2	**2.5±1.0**	1.6±0.3	1.6±0.3
MAPER	**1.4±0.2**	1.5±0.6	1.1±0.2	**2.6±1.0**	1.6±0.3	1.6±0.3
NonLocalSTAPLE	**1.5±0.2**	1.4±0.4	1.1±0.4	**2.5±1.0**	1.4±0.3	1.6±0.3
PICSL BC	**1.3±0.2**	1.2±0.3	1.1±0.2	**2.4±1.0**	**1.2±0.1**	1.5±0.2
PICSL Joint	**1.4±0.2**	1.3±0.3	1.1±0.2	**2.5±1.1**	**1.3±0.2**	1.5±0.3
STEPS	**1.6±0.2**	1.4±0.3	1.2±0.2	**2.7±1.1**	**1.3±0.2**	1.6±0.3

Table 3.4: Comparison Results - OASIS *MHD*: *significant difference to JLF+IM+HMF ($p \leq 0.05$ after Bonferroni correction) is shown in bold*

3.6.4 MICCAI 2013 MRBrainS Images

The MRBrainS database [120] is smaller but, unlike the OASIS database, is multi-variate. Each subject had three co-registered images: a T1w image, a T1 inversion recovery (T1IR) image, and a T2 fluid attenuated inversion recovery (T2FLAIR) image. All images were acquired at 3T with voxel-sizes of 0.96x0.96x3.00 mm. These images were segmented with the labels:

- K representing the background,
- cGM representing cortical gray matter,
- sGM representing subcortical gray matter,
- $eCSF$ representing cerebro-spinal fluid external to the cortex,
- V representing the ventricles,
- WM representing white matter, and
- WML representing the white matter lesions.

Because the white matter lesions were manually segmented in the database, they can be actively accounted for in training the intensity model. This additional training information ameliorates some of the fundamental issues experienced in the worst case results on the OASIS database.

3.6.5 MICCAI 2013 MRBrainS Methods

The segmentation approach taken was the same as that of the OASIS database outlined in Section 3.6.2 with some modifications listed below:

- The intensity model was trained using all three images as features with no smoothed versions.
- The affine and deformable registration steps along with the label fusion were performed on the T1w images without information from the T1IR and T2FLAIR images.
- The highest performing method on the OASIS database (*JLF+IM+HMF*) was the only one evaluated on the MRBrainS database, due to the limited number of entries allowed.

As the label fusion is performed using information solely from the T1w image, it no longer contains a full picture of the local intensity, missing information from the T1IR and T2FLAIR images. Thus, the intensity model should have a much higher impact on the multi-variate MRBrainS segmentation problem compared to the univariate OASIS segmentation problem. As the MRBrainS segmentation problem contains a different number of configuration of labels, a different label ordering must be used. The label ordering used in the MRBrainS segmentation experiment is shown in Figure 3.10.

The segmentation results were evaluated for accuracy using the same three metrics as in the OASIS database: the Dice Similarity Coefficient (DSC), the Absolute Volume Difference

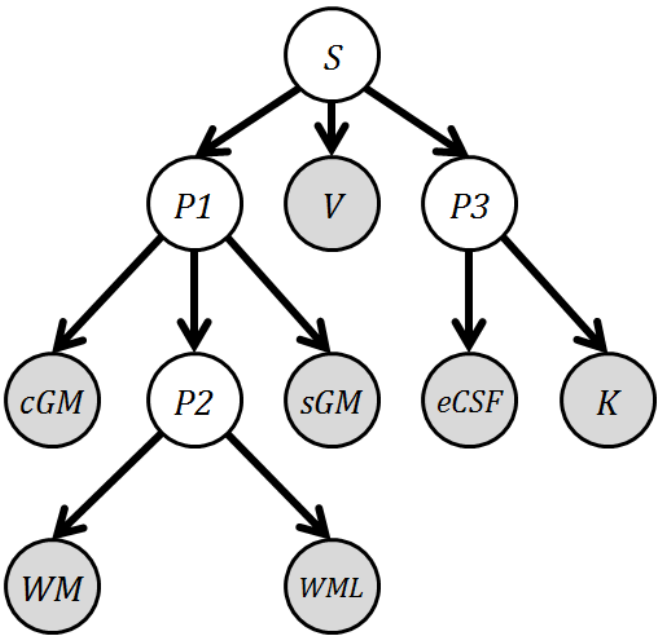

Figure 3.10: Hierarchical label ordering used in segmentation of the OASIS database

(AVD), and the Modified Hausdorff Distance (MHD). Comparative methods include the other four members of the top five best performing prior submitted methods in the MICCAI 2013 challenge:

- *ISI Neo* presented by Moeskops *et al.* [122] uses a three-stage segmentation framework originally proposed for neonatal brain tissue segmentation.
- *MNAB2* presented by Pereira *et al.* [136] uses a random decision forest using local intensity information in all three images.
- *BIGR2* presented by van Opbroek *et al.* [177] uses a support vector machine using Gaussian scale-space descriptors of all three images.
- *U of L* presented by Alansary *et al.* [3] uses a MAP-based approach in which the intensity distribution is modeled as a Linear Combination of Discrete Gaussians with an additional homogeneity descriptor.

Evaluation was performed on the basis of a set of paired t-tests for each evaluation metric. Paired t-tests allow for differences between methods to be extracted even when there is a large amount of variability in the datasets. By using a paired approach, performance variations across datasets are controlled for, isolating the variability resulting from differing methods.

3.6.6 MICCAI 2013 MRBrainS Results

JLF+IM+HMF performed favourably compared to the other top-ranked methods. In terms of the quantitative results (presented in Table 3.5), 48 comparisons were performed, 6 of which showed *JLF+IM+HMF* was superior to another method and 2 of which showed the opposite. (Bonferroni correction was performed on all tests.) That being said, the difference in results between the top five methods are all well within the range of expected manual segmentation variability inherent in the database.

MRBrainS	JLF+IM+HMF	ISI Neo	MNAB2	BIGR2	UofL
DSC CSF	82.10±4.42	81.08±3.25	82.18±3.47	**78.31±4.98**	**78.86±4.23**
GM	84.13±1.46	**85.77±1.89**	84.28±2.12	84.65±1.31	82.96±1.53
WM	87.96±1.11	88.66±1.59	88.18±1.22	88.42±1.21	87.88±1.96
mean	84.73±3.65	**85.17±3.91**	84.88±3.47	**83.79±5.17**	**83.23±4.64**
AVD CSF	12.78±11.57	9.77±10.18	8.81±8.71	**22.76±18.04**	9.69±10.32
GM	5.44±3.81	6.62±3.61	6.97±4.09	6.14±3.39	8.61±5.55
WM	6.59±4.81	6.96±5.92	7.87±6.79	6.02±5.02	8.74±6.59
mean	8.27±8.08	7.79±7.10	7.88±6.69	11.64±13.36	9.01±7.60
MHD CSF	2.71±0.72	2.65±0.52	2.46±0.57	**3.19±0.77**	2.66±0.50
GM	1.92±0.00	1.62±0.45	1.94±0.40	1.89±0.37	1.70±0.28
WM	2.49±0.46	2.07±0.44	2.22±0.51	2.36±0.46	2.24±0.61
mean	2.37±0.59	2.12±0.63	2.20±0.53	2.48±0.77	2.20±0.62

Table 3.5: Segmentation Results - MRBrainS: *significant difference to JLF+IM+HMF ($p \leq 0.05$ after Bonferroni correction) is shown in bold*

Best and worst case segmentation results are presented in Figures 3.11 and 3.12 respectively. In both the best and worst cases, there is slight oversmoothing of the cortical gray matter. Although better than in the OASIS dataset, white matter lesions still have a severe impact on the worst case results where they, coupled with partial voluming, cause erroneous gray-matter like signal in both the T1w and T2FLAIR images, the latter being the most indicative of white matter lesions.

3.7 Discussion

HMF provides a single algorithm for addressing segmentation problems which are equipped with a hierarchical label ordering. This is a fairly general class of segmentation problems which fundamentally includes both the Potts and Ishikawa models as sub-classes. The benefit of having such a general solver is important for medical image processing research, as it allows for multiple labels orderings to be attempted and compared without requiring a large degree

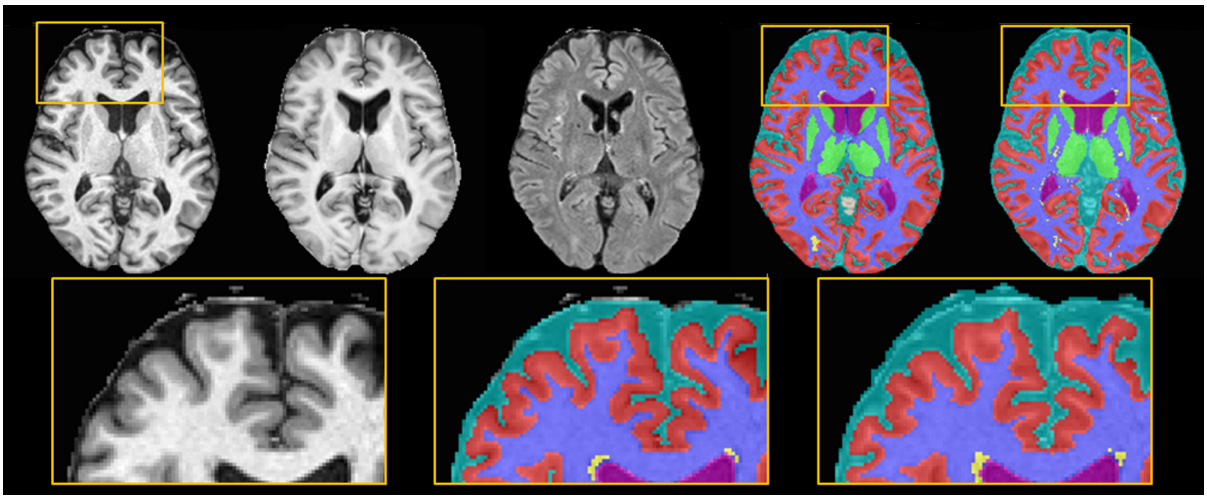

Figure 3.11: Best case results for the MRBrainS database (top row: best case T1w image, T1IR image, T2FLAIR image, gold standard, JLF+IM+HMF. bottom row: enlarged ROI.)

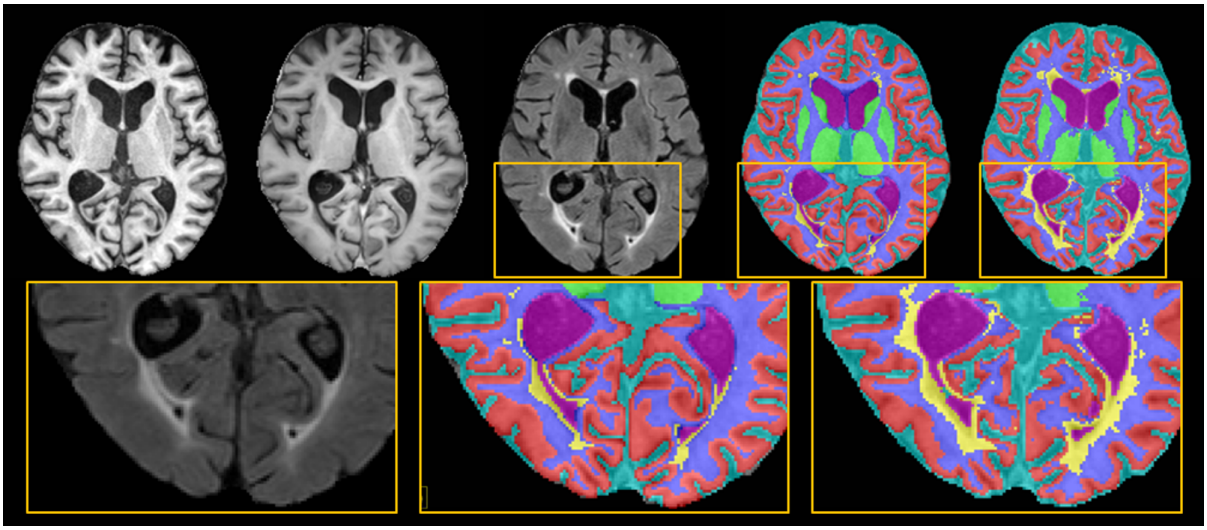

Figure 3.12: Worst case results for the MRBrainS database (top row: best case T1w image, T1IR image, T2FLAIR image, gold standard, JLF+IM+HMF. bottom row: enlarged ROI.)

of additional programming time previously required for hard-coded hierarchies as in Section 3.3.2.

One result of this investigation into HMF models via brain tissue segmentation is the improvement of the model compared to a similar Potts model. The accuracy improvements across almost all labels indicates that the improved expressivity in terms of regularization did have a noticeable effect. It is also notable from the OASIS experiments that improved data terms, in the form of the intensity model or JLF also have a significant effect on the segmentation accuracy regardless of the regularization used. This is understandable and offers a distinct mechanism for improving segmentation methods without more complicated regularization requirements.

3.7.1 Future Work

There are a number of areas of future work to pursue in terms of segmentation problems with hierarchical label orderings. Specifically:

- Investigate frameworks for generating better data terms and regularization terms. For improved data terms, these could include investigation into techniques such as artificial neural networks [83] and random forests [113]. In terms of learning regularization terms, techniques such as probabilistic edge maps [129] might be used to rigorously incorporate edge information based on probability theory similar to the data terms.
- Automatically determine weighting parameters which limit the general use and applicability of more complex hierarchical models.
- Develop interactive segmentation interfaces that would allow for rapid definition and comparison between multiple hierarchies. (This is the topic of Chapter 4.)
- Develop a segmentation framework that would allow for multiple hierarchies to be used simultaneously, even if the resultant label ordering is non-hierarchical. (This is the topic of Chapter 5.)
- Incorporate other topological notions outside of label orderings, such as shape information. (This is the topic of Chapter 6.)

Chapter 4

Optimization Based Interactive Segmentation with Anatomical Knowledge

This chapter is largely based on:

- John SH Baxter, Martin Rajchl, Terry M. Peters and Elvis C.S. Chen. "Optimization-based interactive segmentation interface for multi-region problems." SPIE Journal of Medical Imaging (2016).
- John SH Baxter, Eli Gibson, John M. Moore, Roy Eagleson and Terry M. Peters. "The Semiotics of Medical Image Segmentation." *submitted to* Medicine, Health Care and Philosophy (2017).

4.1 Introduction

As stated in Section 3.4, determining the largest hierarchy from a collection of regularization requirements, i.e. finding an HMF model which targets as many requirements as possible, is NP-hard. This implies that user knowledge is required to determine not only the components are in the segmentation ontology, but also the structure of said ontology as well. Because of the general HMF solver presented in Chapter 3, the hierarchical label orderings used to structure the segmentation ontology can be provided as input, rather than seen as inherent to the segmentation algorithm.

Altogether, this motivates the concept of ***interactive segmentation with anatomical knowledge*** where the user provides not only indications of object location and properties, but also how they relate to other objects. As interactive segmentation combines the positive aspects of both manual and automated segmentation, the ability to incorporate anatomical knowledge as a locus for interactivity can lead to better handling of imaging abnormalities/heterogeneities

and patient pathology.

This chapter begins with an exploration into the semiotics of segmentation including the Piercean notion of *signs* and how they contribute to interactive segmentation. This is followed by an overview of interactive segmentation interfaces for medical imaging and a technical description of the Segmentation User Environment (SegUE). This interface has been applied to cardiac segmentation across multiple modalities, and to neonatal cranial ventricle segmentation from MRI. This chapter concludes with areas of discussion and future work in the field of knowledge-based interactive segmentation.

4.2 Philosophy of Segmentation

When developing a general-purpose segmentation framework, it is important to consider the basic *philosophy* of segmentation and lay a firm conceptual foundation for how to categorize, compare, and contrast the large body of research into the subject.

The first important notion is that of *interactivity*. Briefly, interactivity is a blanket term that captures the difference between a segmentation framework that is fully manual and one that is fully automatic. The interactivity spectrum for segmentation is shown in Figure 4.1. On the left side are manual and interpolated segmentation in which the anatomical knowledge and effort are provided primarily by the user. (In the case of manual segmentation, all knowledge and activity is provided by the user with the computer acting as a mere receptacle for this activity.) For semi- and fully automatic segmentation, the anatomical knowledge is provided largely by the computer in terms of what has been encoded into the segmentation algorithms explicitly, or what knowledge is implicitly captured from training images, i.e. an atlas. The areas on this

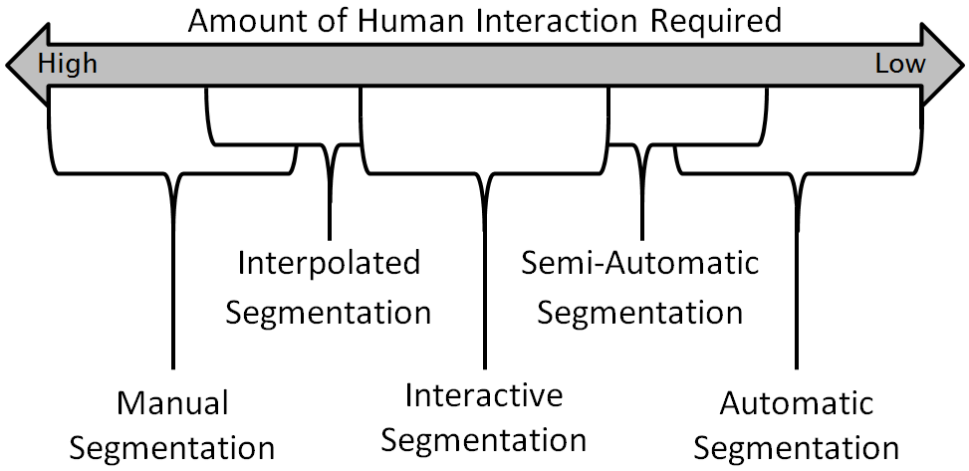

Figure 4.1: Segmentation Interactivity Spectrum

chart overlap in order to emphasize that these boundaries are far from objective or universally agreed upon, and many would dispute the binning of any non-trivial example into any category.

The primary concern of this chapter is *interactive segmentation*, that is, interfaces in which the user is the primary provider of significant anatomical knowledge, yet the computer plays an active role in determining the final segmentation result. (However, this semiotic analysis is valid for the manual and automatic extrema, however with less richness.) In the interactive segmentation paradigm, the problem of finding the correct segmentation is a dialogue between the clinical user and the computer, with the user providing information, structure, and possibly correction while the computer providing estimates of the segmentation and areas of uncertainty. In that sense, the computer is viewed as a *dialogue partner* [85]. In this sense, the other forms of segmentation could also been seen as limiting cases of a dialogue, particularly ones in which one party, either the clinical user or the computer, is largely passive.

There may be legal and regulatory ramifications to this approach, particularly because interactive segmentation lies in the intermediate space between interpolated and semi-automatic segmentation paradigms, both of which have well-defined regulatory approval requirements and best practices.

4.2.1 Semiotics in Segmentation

Olabarriaga & Smeulders [130] in their 2001 review of medical image segmentation from the perspective of interactivity identify three aspects for understanding the dialogue between the clinician and the computer. Specifically, it raises the questions:

1. What types of input signs are provided by the user?
2. How is the user input interpreted by the computer?
3. What is the purpose of user interaction in the segmentation process?

and asks what effect these answers would have on *accuracy*, *repeatability*, and *efficiency*. Their review article is important in that it focuses attention on other aspects of segmentation aside from computation time and accuracy. However, their analysis lacks the rigour of a formal philosophical analysis and thus many particulars have not stood the test of time well.

By motivating segmentation (and interactive segmentation in particular) as a dialogue between the clinical user and the computer, the philosophical focus of segmentation moves to the investigation of what constitutes communication, especially considering the lack of theoretical common ground between clinical and computational problems. For this, one can turn to the philosophical study of *semiotics*, that is, the study of signs, their structure and meaning. In particular, the theory of signs developed by Charles S. Pierce [133] offers a powerful categorization of signs into:

- *Icons* or *iconic signs* that represent an object by being similar to it in some way,
- *Indexes* or *indexical signs* that represent an object by correlating with said object or specifying a property thereof, and
- *Symbols* or *symbolic signs* that represent an object by virtue of some arbitrary or socio-cultural standard.

An important addition to this categorization is the historically preceding notion of a *part* or *partial sign* which represents an object by virtue of being a part of said object [162]. (These types of signs were originally referred to as *symptoms* but calling them partial signs limits confusion with the common-sense use of the former term.) The signs involved can also be either user-generated (*input sign*) or reflects the output of a particular computational process (*output sign*).

This classification gives a simple mechanism for interpreting any given sign, although at a more superficial level. However, it is not sufficient for understanding semiotic interaction in segmentation interfaces in which individual signs may be composed of multiple heterogeneous parts and interpreted on multiple levels. Lastly, this analysis should elucidate the theoretical basis used by the user and the computer to understand the sign underyling any given interaction, addressing this issue of heterogeneity and interpretation. This basis extends the *roles* played by the user and the computer described by Olabarriaga & Smeulders [130]. Thus, this analysis focuses on four aspects, extending the three presented by Olabarriaga & Smeulders [130]:

1. What information should constrain the problem, be provided as input, or result from the segmentation process?
2. What input signs are or should be provided by the computer?
3. What output signs are or should be provided by the computer?
4. How can these signs be designed to maximize coherence between the user and the computer?

The first question is the traditional fare of requirements engineering and can be investigated outside the scope of semiotics. The remaining three questions are addressed in the following subsections.

4.2.2 Input Signs

Table 4.1 presents some common input sign mechanisms with their corresponding semiotic type. Although the four types should be exhaustive, the list of examples is not.

In terms of input signs, one unintuitive action to consider is that the user selects the segmentation algorithm itself. Although this is largely implicit in many segmentation frameworks,

some image analysis interfaces such as Khoros [90] and MeVisLab [151] make this type of input sign explicit by allowing the user to create their own on-the-fly segmentation pipeline with basic processes such as thresholding and morphological operators, as connected boxes. These actions, the selection and modification of the segmentation pipeline as a whole, are largely symbolic in that the naming and graphical conventions dictate the segmentation pipeline. Similarly, user actions such as choosing an algorithm to immediately run, outside of the notion of a pipeline, could be viewed as symbolic. The selection of numeric parameters may also be considered symbolic to a clinician in that they control properties of the segmentation often in indirect or opaque ways. The number of segmentation interfaces that employ these signs are so numerous that it is difficult to find ones that are without opaque parameters.

Considering actions that don't modify the segmentation pipeline, arguably, the most common input signs used in modern segmentation frameworks are partial in that the user supplies a sign indicating a part of the segmentation. Common techniques include the picking of seed-points common in traditional region growing [2, 74] and fuzzy connectedness [70, 174] based frameworks. This single point methods can readily be extended to seed regions with higher density using paint-brush based approaches [23, 27, 63, 66, 150, 149, 198]. The previous signs discussed signified regions interior to the segmentation, but there are contour based approaches, such as LiveWire [15, 53], which rely on identifying portions of the region's boundary [151]. ITKSnap [198] is particularly illustrative in that it includes both a paint-brush tool and a polygonal outlining tool (contouring) more conceptually suitable for identifying edges. In the case

Partial Sign	**Iconic Sign**
• Centroid or seed-point selection [2, 70, 74, 174] • Paint-brush [23, 27, 63, 66, 149, 150, 198] • Contouring [15, 53, 138, 198] • Selection of a pre-computed region (i.e. a super-pixel) [58] • Full segmentation of an object in an individual slice [138, 159]	• Selecting or providing a prior 'rough' segmentation [53, 127, 187] • Selection of an image from an atlas [125]
Indexical Sign	**Symbolic Sign**
• Specifying topological / geometrical / volumetric properties of a single label [65, 140, 178] • Specifying topological properties between multiple labels [43]	• Algorithm selection • Selection of a simple image processing operator and composition of these into a pipeline [88, 90, 151]

Table 4.1: Classification with examples of different input sign types

of segmenting volumetric data, some segmentation methods rely on a full segmentation of individual 2D slices which are then propagated or interpolated between throughout the remainder of the volume. [159]

Iconic input signs are when the user provides an entire sample segmentation (as opposed to a partial one) in terms of either the same image or an atlas. In terms of most atlas-based segmentation paradigms, the atlas is a fixed component and not subject to modification by the user. One illustrative counter-example is that developed by Nowinski *et al.* [125] in which the user actively selects the atlas with which to segment the image. Separate from atlases is the concept of pre-segmentation in which the user provides a rough segmentation of the image which is then refined by the algorithm. This is often used as a pre-processing step for more complex methods designed to preserve the segmentation topology such as LOGISMOS [127, 187], although often these approaches are generated automatically rather than directly through user action.

The last category of input signs also happens to be the most conceptually diverse. Specifying properties of a segmentation such as topological [43], geometric [65, 178], or volumetric [140] give the clinician a more transparent mechanism to create large conceptual changes in the segmentation. Despite this diversity, indexical input signs appear to be relatively underutilized, with many algorithms enforcing these properties rather than allowing them to be user-specified.

These input signs should be designed in order for each to allow the user to fulfil a defined *role*. Olabarriaga & Smeulders [130] provide a list of five roles a user can play in a segmentation process:

- *Judge:* The user determines if a segmentation is adequate or inadequate;
- *Corrector:* The user informs the computer about which regions of the segmentation are incorrect, and corrects them;
- *Parameter-Selector:* The user selects the parameters involved in the segmentation process;
- *Composer:* The user composes a segmentation from regions in the image or rough segmentations; and
- *Builder:* The user builds a segmentation process out of more primitive operations.

(In some senses, the role of *composer* is redundant with that of *judge* or *corrector*.) As one may expect, certain input signs are particular well designed for particular roles. For example, due to the symbolic nature of defining image processing pipelines, the vast majority of signs that support the user's role as builder are symbolic. Similarly, as corrections are normally made to specific regions of the image, the majority of signs that support the user's role as corrector are partial.

But what does this classification tell us about segmentation? One thing that it clearly indicates is a preference in terms of the interactive segmentation community for the symbolic and the partial forms of interaction, with iconic and indexical filling niche roles. This is concerning as it is fundamentally opposed to the current paradigm in teaching anatomy which is almost solely dependent on the iconic (through diagrammatic/pictorial representations of anatomy in textbooks) and indexical (through the description of the properties of the anatomy/physiology). Exploring these forms of input sign may offer a lucrative area of research, bridging the gap between how anatomy is taught and how it is practised in the context of medical image segmentation [190].

4.2.3 Output Signs

One can apply the same classification to output signs based on the intent of the computer. The most common output signs are largely partial and iconic, providing cross-sections and renderings of the segmented anatomy respectively.

More interesting are the indexical and symbolic signs in which the computer does not try to represent the segmentation directly but instead signal some property belonging to the segmentation or indicate some process to the user. In terms of indexical output signs, the most common is the computer reporting some metrics regarding the segmentation such as its volume. Symbolic output signs can vary from specifying a potential diagnosis or engaging in active learning signalling to the user which slices would benefit the most from further user interaction [171].

Similar to the input side, Olabarriaga & Smeulders [130] also specify two roles that a computer can take in the segmentation process: segmentation process:

- the computer plays a *direct* role when it abstracts away the technical aspects of the underlying segmentation algorithms, allowing the user to focus on providing knowledge rather than having a full understanding of the computational processes; and

Partial Sign	**Iconic Sign**
• Coloured/labelled slices [23, 27, 63, 171, 198] • Contoured slices [15, 138]	• Segmentation volume rendering [91] • Segmentation isosurface rendering [138, 171, 198]
Indexical Sign	**Symbolic Sign**
• Energy/fitness functional values [172] • Property maps such as cortical thickness and curvature maps [56]	• Computer determined diagnosis • Active learning [171]

Table 4.2: Examples of different output sign types

- the computer plays an *indirect* role when it actively attempts to reduce the number and complexity of future interactions by learning or modelling the users intent.

In many aspects, the computer performs these roles in tandem, and individual computational processes may not comfortably sit in only one of those categories. For example, the computer may collect a intensity probability distribution from user provided seeds in order to predict the intensity probability distribution of the object being segmented. In this case, the computer can be thought of as acting both directly and indirectly, as it is abstracting the technical elements away from the user (i.e. by simplifying the user input to a partial input sign) and indirectly modelling the user's intent (i.e. that the object is bright/dark/heterogeneous/etc...).

4.2.4 Heterogeneity and Sign Graphs

Although in Peirce's original classification [133], the specified sign types are distinct, in practice the complexity of the underlying semiotic process and often the lack of in-depth computational knowledge of the users imply that elements in the taxonomy are not always mutually exclusive. In both computational and more general semiotic investigation, signs are often given multiple levels of meaning or mechanisms whereby this meaning is conveyed and this heterogeneity leads to a lack of distinctness in the typology of the overall sign. These considerations can be elucidated and clarified by explaining exactly how the signs are interpreted by both the computer and the user, but this can be very complex if left unstructured. Peirce recognized the problem of multiple levels of interpretation in his later semiotic accounts [134], but ultimately did not provide a singular framework suitable for modern field of human-computer interaction. Peirce's notion of infinite semiosis in which signs have an infinite constellation of meanings weakens its ability to be used in a human-computer stand-point where interpretations must happen in a finite time.

Thus, a finite description of interpretation, a finite semiosis, must be formulated which is sufficiently powerful to capture heterogeneity and multiple levels of meaning but compact enough to be quickly constructed and analyzed. The chosen representation is a *sign graph* in which a sign can be decomposed into a finite series of simpler connected signs. In these graphs, each directed edge represents a simple interpretant connecting a sign vehicle to an object. The benefit of this particular representation is, for the majority of computational roles, that the sign graph used is explicitly known, encoded in the implementation of the segmentation algorithm. The creation of sign graphs to structure new segmentation algorithms and interfaces display echoes of traditional object-oriented design with computation being represented as an abundance of interacting, yet distinct, objects.

These sign graphs can be used also to represent the user's interpretant as well as that of

the computer. Figure 4.2(a) demonstrates a novice user's possible interpretation of a single brush-stroke in which some action of the mouse (holding down a key and moving) is symbolically translated into the placement of a collection of seed points that are used to denote spatial parts of a single object in the segmentation. This may however be somewhat different to how a computer interprets the same action. For example, if the computer uses those seed points as spatial anchors (a partial sign interpretation) as well as samples of the single object's intensity distribution (a more heterogeneous sign with partial and indexical components), the interpretation used by the computer is more complex and can be represented as the sign graph in Figure 4.2b. As the user gains experience or knowledge of the system, the user's interpretant may grow in complexity, ideally aligning itself with that of the computer.

In terms of human-computer interaction, a fundamental aspect of design is the process of aligning interface metaphors which are mental models of how a computational system acts [52]. In this framework, all forms of communication between the human and the computer are seen as a metaphor in which the source and target domains are the elements of the user's and computer's sign graphs. These sign graphs thus capture a notion of the difference between novice and expert users by adding to or subtracting extraneous elements from the user's sign graph as said user intuits the behaviour and mechanisms of the computer. Similarly, if the computer employs mechanisms for learning the user's intention (as described by Olabarriaga and Smeulders [130] as an indirect role of the computer, opposed to providing a level of abstraction

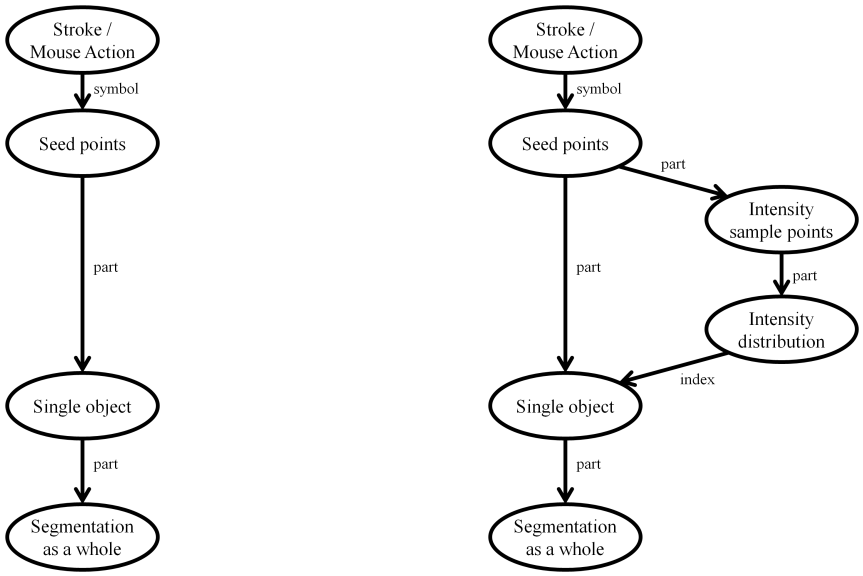

(a) Novice interpretation of a paintbrush stroke, interpreting each point as a spatial anchor for a single object.

(b) More advanced interpretation of a paintbrush stroke, interpreting each point as a spatial anchor and as providing statistical information about the intensity.

Figure 4.2: Example Sign Graphs in Segmentation Interfaces

over the underlying algorithms, which is described as a direct role), the computers sign graph may approach that of the user, although such a graph may not be explicitly represented in the computer.

The notion of a sign graph necessarily complicates the classification of input and output signs into parts, icons, indices, and symbols in that each sign graph may be heterogeneous and include multiple sign types. By varying these signs graphs with only slight steps, it is possible to blur the line between the classifications for the sign as a whole. This heterogeneity has its advantages; it subsumes the view of all interaction with a computer as being symbolic through the frequent occurrence of symbolic simple sign connecting the users physical action, such as moving the mouse, to an initial computational abstraction, the placement of seed points. However, unlike the purely symbolic approach, it encourages greater depth of explanation into how the computational processes are understood by the user.

So, how does an understanding of signs and sign graphs facilitate the design of segmentation interfaces? Aside from the informational concerns, which are the subject of traditional requirements engineering, semiotic concerns can also be used to guide development, specifically through the process of aligning interface metaphors [52]. For this, we can take advantage of Thagard and Holyoak's [75] constraint satisfaction theory, which includes metrics to evaluate the efficacy of metaphors, to motivate three ways in which semiotics should guide development, viewing the act of interpretation of any sign as similar to an analogy between the interpretants made by both the user and the computer. For current interactive segmentation interfaces, the computer's interpretant for any given sign is explicit, and is encoded in the interface's implement with the a priori knowledge of the interface designer.

The first major concern is that of structural consistency between the user and computer's interpretation of the sign. Note that this does not imply that the user or the computer has to have a full understanding of the myriad ways in which the sign can be interpreted; Olabarriaga and Smeulders' [130] notion of the computer's direct role in computational abstraction is still essential. However, both the user and the computer require a shared path through the sign graph that allows the sign to have a partially predictable effect. In the example given in Figure 4.2, all of the nodes and edges in Figure 4.2a also appear in Figure 4.2b indicating a relatively high degree of structural consistency though not the absolute highest, given that some nodes in Figure 4.2(b) lack a corresponding node in Figure 4.2a. Importantly, there is a shared path through the sign graph, e.g. that the seed points are considered inside the object, indicating some common ground for the user and computer to understand each other. If this was not the case for the computer, for example, and the seed points were only used to generate a probability distribution and their spatial locations were not respected in the algorithm, one can readily see users becoming frustrated, especially when using the paintbrush as a mechanism for

correcting erroneous regions in the segmentation. Similarly, an output sign should engender an interpretation in the user that incorporates key elements and relations used by the computer in the signs generation.

The second concern raised here is the notion of semantic similarity, that is, the similarity between the individual pairs of corresponding nodes or edges in the two sign graphs. For example, if both the user and computer interpret seed points as definitely part of the object of interest, that portion of the analogy has higher semantic similarity than if either the user or the computer viewed those points as only heuristic and not definitive. There is a wide body of thought that believes this is essential to the success of symbolic signs in human-computer interaction based of the interface metaphor and epitomized in common HCI constructs such as the desktop [60, 111]. In the context of medical image segmentation, the same considerations apply for symbolic input signs, but the notion of semantic similarity is less well-defined for abstract constructs (such as seed points, distributions, algorithms etc. . .) that form the majority of a sign graph in modern interactive segmentation interfaces. Semantic and structural consistency are different in that structural similarities refers to how well a mapping between the user's and computer's interpretants can be created, regardless of the similarity of corresponding nodes, which is the purview of semantic similarity.

The last concern is that of pragmatic centrality, that is, how well the designed sign can support a particular high-level action. Olabarriaga and Smeulders [130] provide some concept of these high-level actions, at least in terms of input signs, in their description of the role of the user. These roles center around a particular action, such as judging a segmentation as adequate or not, correcting a segmentation, building a segmentation process, etc The question of pragmatic centrality, from the perspective of input signs, is does the mapping between the user and computers interpretation of a sign support these actions well? If a user's intention is to place a dense series of seed points, a paintbrush mechanism is simply more pragmatic, more suitable for that use, than say, individually picking each pixel which would be more time-consuming. From the output sign perspective, one has to evaluate whether or not an output sign effectively expresses the computers goal. For example, displaying cortical thickness information is easier on a spherical map, where the entire cortex can be visualized without worrying about crevices or tortuosity, than it would if it were painted directly onto an image slice.

As stated earlier, the computer's interpretant of any given sign is encoded in the implementation of the interface and thus can be seen as explicitly known by the interface designer. Thus, modifying and improving the computer's interpretant is a natural place to begin the process of aligning the interpretants according to structural consistency, semantic similarity, and pragmatic centrality. In order to do so, one must elicit sign graphs from users, that is, get a detailed understanding of how the user interprets and understands the problem domain. In

knowledge communities, such as medicine, a large amount of domain information can be gathered from sources such as textbooks and clinical operating procedures, which can be used to obtain a general concept of elements in an idealized user's sign graph. This process may also furnish developers with the specialized terms and meanings associated with a particular field of medicine and what knowledge and tasks can be fundamentally support and be supported by medical image segmentation. It may even provide an understanding of the constant, input, and output information required in the segmentation problem. For understanding actual users, however, traditional human-computer interaction techniques which focus on an understanding of the user's mental state, such as the think aloud protocol [98] should be used. Think aloud protocols have been previous used to elucidate fundamental differences between idealized and actual clinical users in a variety of fields including oncology [82], pulmonology [105], and electronic health record management [54]. Note that this protocol on its own may not be sufficient for elucidating the entire sign graph under a variety of scenarios, but may have to be augmented with other techniques [81].

It is worth noting that modifying the interface and implementation is not the only mechanism for aligning sign graphs. Effective user training addresses the problem of aligning interpretants through a modification of the user's behaviour and mental state, rather than the computer's. Thus, in a semiotics driven approach, developing effective user training and explanatory manuals is an essential aspect of interface design that cannot be discarded or seen as separate. For many clinicians, continuing clinical education is a recognized portion of their practice, offering an avenue in which user training can be incorporated, provided it is recognized as efficiently and effectively aligning the user's and computer's interpretants.

By understanding the interaction as a co-ordinated series of sign graphs forming an analogy between the user and computer interpretants, conceptual tools from human-computer interaction and the cognitive science of analogies can be leveraged simultaneously with Peircean categorization. This combination subsumes several disparate views on segmentation design from the literature (specifically Olabarriaga and Smeulders [130]) and from traditional human-computer interaction (the notion of the interface metaphor) and offers a series of fundamental questions which elucidate and critique the design of segmentation interfaces.

4.2.5 Philosophical Call to Action

As becomes obvious in the relatively low number and high variability of the citations, the field of indexical input signs in interactive segmentation is nascent. Although developing computationally efficient solvers or approximations by far represents the more technically complex aspect of developing these input sign mechanisms, they are intuitively a wealth of novelty and

possible clinical applicability. Thus, my motivation is to take advantage of this and to develop an interactive segmentation interface that allows for an indexical input mechanism in the form of specifying hierarchical label ordering. By creating an interface with distinctly indexical input signs, we are demonstrating the capabilities, albeit in a limited way, of indexical input signs as a whole and where they can better address segmentation problems in a clinical context.

4.3 Interactive Segmentation Interface

4.3.1 Hierarchical max-flow segmentation

Hierarchical max-flow (HMF) models (see Chapter 3) extend the notion of orderings from the Ishikawa model [14, 78] to hierarchies. In this case, collections of labels can by unioned to create a super-label. This process repeats itself until the entire image is represented as a single label, denoted as S or the root label. (Such S labels are shown at the top of the hierarchies in Figures 4.7 and 4.8.) Alternatively, one can take a top-down interpretation, recursively splitting objects in the image into their constituent parts. The formula for these models (assuming isotropic Euclidean regularization) is:

$$\begin{aligned} &\min_{u} \sum_{L\in\mathbb{L}} \int_{\Omega} D_L(x)u_L(x)dx + \sum_{L} \int_{\Omega} R_L(x)|\nabla u_L(x)|dx \\ &\text{s.t. } \forall L \in \mathbb{L}, (u_L(x) \geq 0), \sum_{L\in\mathbb{L}} u_L(x) = 1 \\ &\qquad \forall L \notin \mathbb{L}, \left(\sum_{L'\in L.C} u_{L'}(x) = u_L(x) \right) \end{aligned} \tag{4.1}$$

This formula is similar to that of the continuous Potts model [193] in that it contains a series of unary data terms, $D_L(x)$, and a set of regularization or weighting terms, $R_L(x)$, on the gradient magnitude of each labelling function, $|\nabla u_L(x)|$, and that these terms are summed over each label (and super-label). These hierarchical models are strictly more expressive than both Potts [139] and Ishikawa [78] models together, allowing for a wider array of segmentation problems to be addressed. However, hierarchies are more difficult to specify. For more information about hierarchical models and HMF, see Chapter 3.

This tree structure has previously been considered a hard-coded part of the image segmentation algorithm, encouraging the use of Potts [139] or Ishikawa [78] models and their continuous max-flow counterparts [14, 193] to handle general-purpose segmentation. However, this poses fundamental limitations on what can be segmented. For example, in the Potts model, only a single smoothness parameter is assigned, which makes it difficult to simultaneously segment

smooth structures alongside irregular ones. Ishikawa models allow for more parametrization, but require the objects being segmented to satisfy a full ordering, which is not the case for complex anatomy.

The general HMF solver alleviates this problem by permitting any arbitrary hierarchy to be defined, allowing for more anatomical knowledge to be encoded. This intuitive form of anatomic knowledge can be readily incorporated into the optimization-based segmentation of multiple regions. Problems regarding constructing the largest meaningful hierarchy given label grouping information are NP-hard (See Appendix), meaning that interactive methods, at least for hierarchy definition may be required so as to make use of a user's anatomical knowledge.

4.3.2 Definition of Cost Terms

A crucial decision in optimization-based segmentation is the structure of the cost terms. Log-likelihood data terms, derived from Bayes' theorem, have been effective in interactive [27] and multi-region segmentation [176], taking the form:

$$D_L(x) = \begin{cases} \infty, & \text{if } x \text{ is a seed for a label other than } L \text{ or element of } L.P^* \\ -ln\left(P(I(x)|x\epsilon L)\right), & \text{else} \end{cases} \tag{4.2}$$

where $P(I(x)|x\epsilon L)$ is the likelihood of a voxel in label L having the same intensity as x, $I(x)$ and $L.P^*$ is the set of ancestors (parents, grandparents, etc...) of label L. The probability, $P(I(x)|x\epsilon L)$, is estimated from the histogram of the seeded voxels, which approximates the true value when a large number of seeds is used. The infinite cost ensures that any voxels used to seed a particular object remain a part of said object in the segmentation and that any voxel which has been seeded as a particular label can only be assigned to said label if it is an end-label or to its child labels otherwise.

Smoothness or regularization terms are non-negative costs used to both smooth the labelling and to align edges in the segmentation with those visible in the underlying image. The smoothness terms used were:

$$R_L(x) = \alpha_L * \exp(-\beta_L|\nabla I(x)|) + \gamma_L \,, \tag{4.3}$$

where the parameters α_L, β_L, and γ_L are specified by the user. The exponential term implies that lower costs are associated with label boundaries which occur when there is a high gradient magnitude, encouraging nearby edges in the segmentation to migrate to said areas similar to the contrast sensitive smoothness terms used by Boykov *et al.* [27].

4.3.3 Plane Selection

To improve efficiency and accuracy while encouraging interactivity, plane selection can be used [171]. Such planes can be used by the algorithm to inform the user as to which areas of the segmentation would benefit the most from user interaction. Top *et al.* [171] introduced a notion of active learning in which the segmentation algorithm identifies areas of maximum uncertainty, the uncertainty of a segmentation expressed as:

$$U(x,y) = \lambda_E U_E(x,y) + \lambda_B U_B(x,y) + \lambda_R U_R(x,y) + \lambda_S U_S(x,y) \tag{4.4}$$

where U_E is the entropy of the segmentation results, U_B the uncertainty associated with boundaries in the segmentation, U_R the the uncertainty associated with the regional intensity, and U_S the uncertainty associated with the tortuosity of the boundary around x. The λ's are constants with the majority (80%) of the weight given to λ_E [171]. Note that the U_R and U_B terms are explicitly handled by the segmentation algorithm itself by the definition of the cost functions. We assign all the weight to the U_E term and use only maximum axis-aligned planes. This ensures that the plane selection algorithm quickly produces planes in orientations to which the user is accustomed. The segmentation used in plane selection is the previous segmentation generated by the user. Thus, plane selection is only defined after the first segmentation is computed and remains available for all subsequent interactions.

4.3.4 Interface Description

The interface is implemented using Kitware's Visualization Tool-Kit (VTK) for image processing and visualization and the Qt framework for graphical user interface support. The generalized HMF solver was encapsulated into a VTK algorithm object and accelerated using NVIDIA's Compute Unified Device Architecture (CUDA). The sampling mechanism is brush-based similar to that used by Boykov *et al.* [27] and ITK-Snap [198] allowing for large portions of an object's interior to be covered with relative ease while not requiring strong boundary contrast. The user can place seeds for any label or super-label using the brush, creating the data model described in Eq. (4.2).

The interface is shown in Figure 4.3. Hierarchies are defined in a side bar as shown in Figure 4.3(c), which also acts as a widget for selecting the active label or super-label of the brush. This widget also allows the hierarchy to be restructured quickly, operating in a drag-and-drop manner. Lastly, the user can save the hierarchy along with smoothness term parameters and the initial user-defined samples for later use.

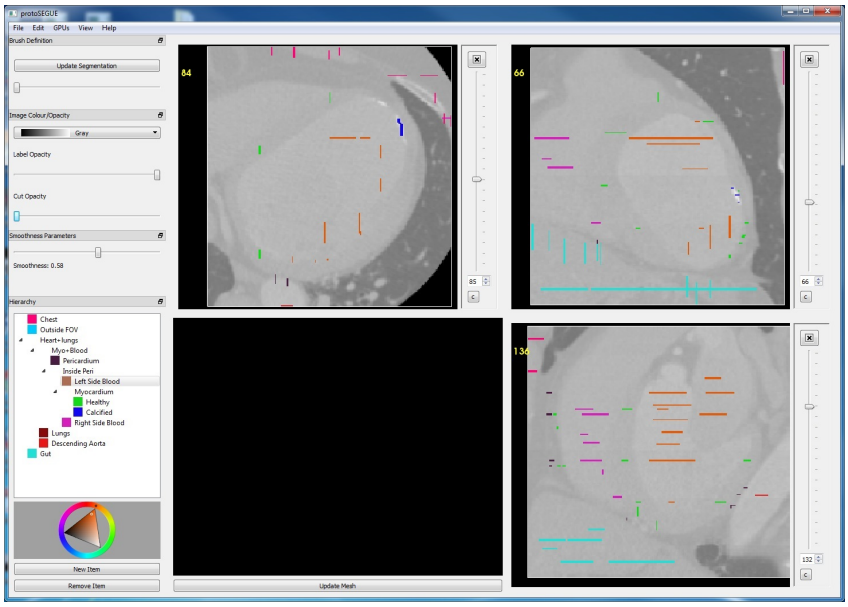

(a) Interface with seeding

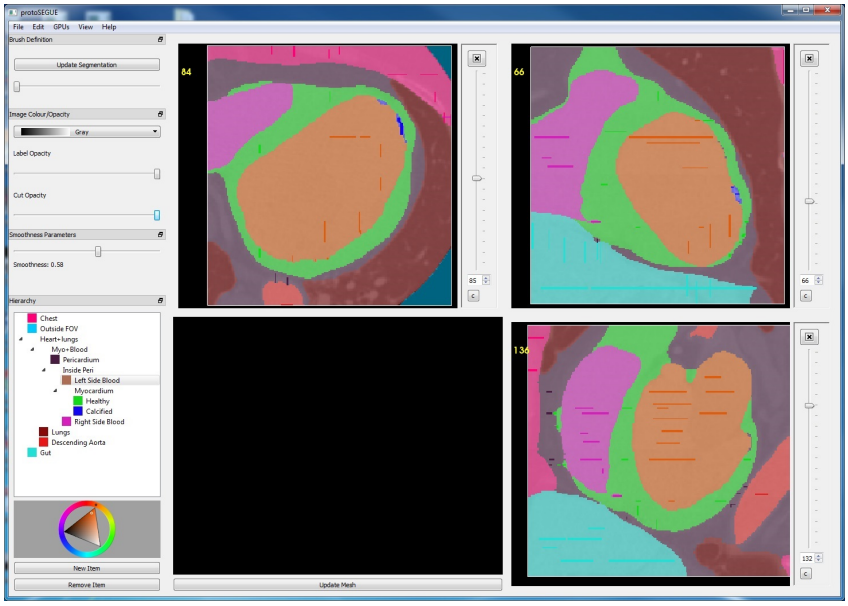

(b) Interface after labelling is performed with provided seeds

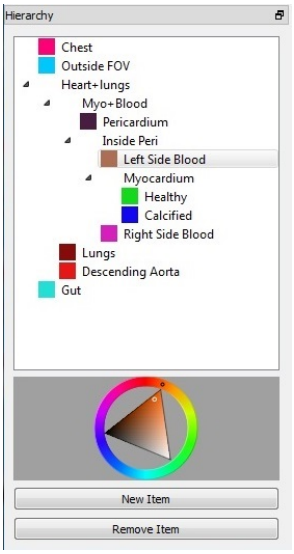

(c) Widget for hierarchy definition

Figure 4.3: Segmentation interface with user seeds before segmentation (a) and after segmentation (b). The hierarchy definition widget (bottom left corner of (a) and (b)) is shown enlarged in (c).

4.4 Example Applications of Interactive Segmentation

4.4.1 Cardiac Segmentation

Because of the generality of the algorithm and the interface, several existing continuous max-flow based methods, such as those developed by Rajchl *et al* [149, 150], can be easily replicated. We reproduced the experiments performed in [149], which included 3 cardiac volumes from computed tomography (CT), magnetic resonance angiography (MRA), and trans-esophageal echocardiography (TEE). Ethics approval for these images can be found in Appendix C Section C.3. An example of each modality and the corresponding segmentation is given in Figure 4.4.

Numerical results in terms of average volume difference (AVD), root mean squared distance error (rMSE), and Dice similarity coefficient (DSC) are recorded in Table 4.3. These results are very consistent with inter-operator variability above 90% and comparable with those presented by Rajchl *et al.* [149] illustrating that SEGUE, a general-purpose segmentation interface, can perform similarly to one designed specifically for cardiac segmentation. Interestingly, the results for TEE indicate that the proposed interface outperforms the previous interactive segmentation interface [149].

These results indicate that even though the proposed interface is general purpose, it has the capability of performing comparably to existing single purpose interfaces, at least to within manual segmentation variability.

(n = 3)		CT	MRA	TEE
Blood AVD	(%)	6.6 ± 6.6	6.2 ± 3.6	14.2 ± 6.2
Myocardium AVD	(%)	12.5 ± 11.3	16.7 ± 11.5	7.3 ± 4.5
Blood rMSE	(mm)	1.14 ± 0.64	0.70 ± 0.21	1.08 ± 0.27
Myocardium rMSE	(mm)	1.31 ± 0.24	0.71 ± 0.24	1.48 ± 0.57
Blood DSC	(%)	91.7 ± 2.6	94.3 ± 1.9	90.5 ± 4.3
Myocardium DSC	(%)	83.8 ± 3.9	82.1 ± 3.7	91.8 ± 2.7
Weighted DSC	**(%)**	**87.5 ± 2.0**	**89.8 ± 2.7**	**91.2 ± 3.2**
Inter-operator variability Weighted DSC	**(%)**	**92.7 ± 4.9**	**93.6 ± 2.5**	**92.0 ± 2.1**
Weighted DSC from [149]	**(%)**	**87.7 ± 3.7**	**89.3 ± 2.7**	**85.7 ± 2.0**

Table 4.3: Cardiac Segmentation Numerical Results. Results shown in **bold** indicate metrics that are common across SEGUE, manual segmentations, and [149] which can be used as a reference for comparison

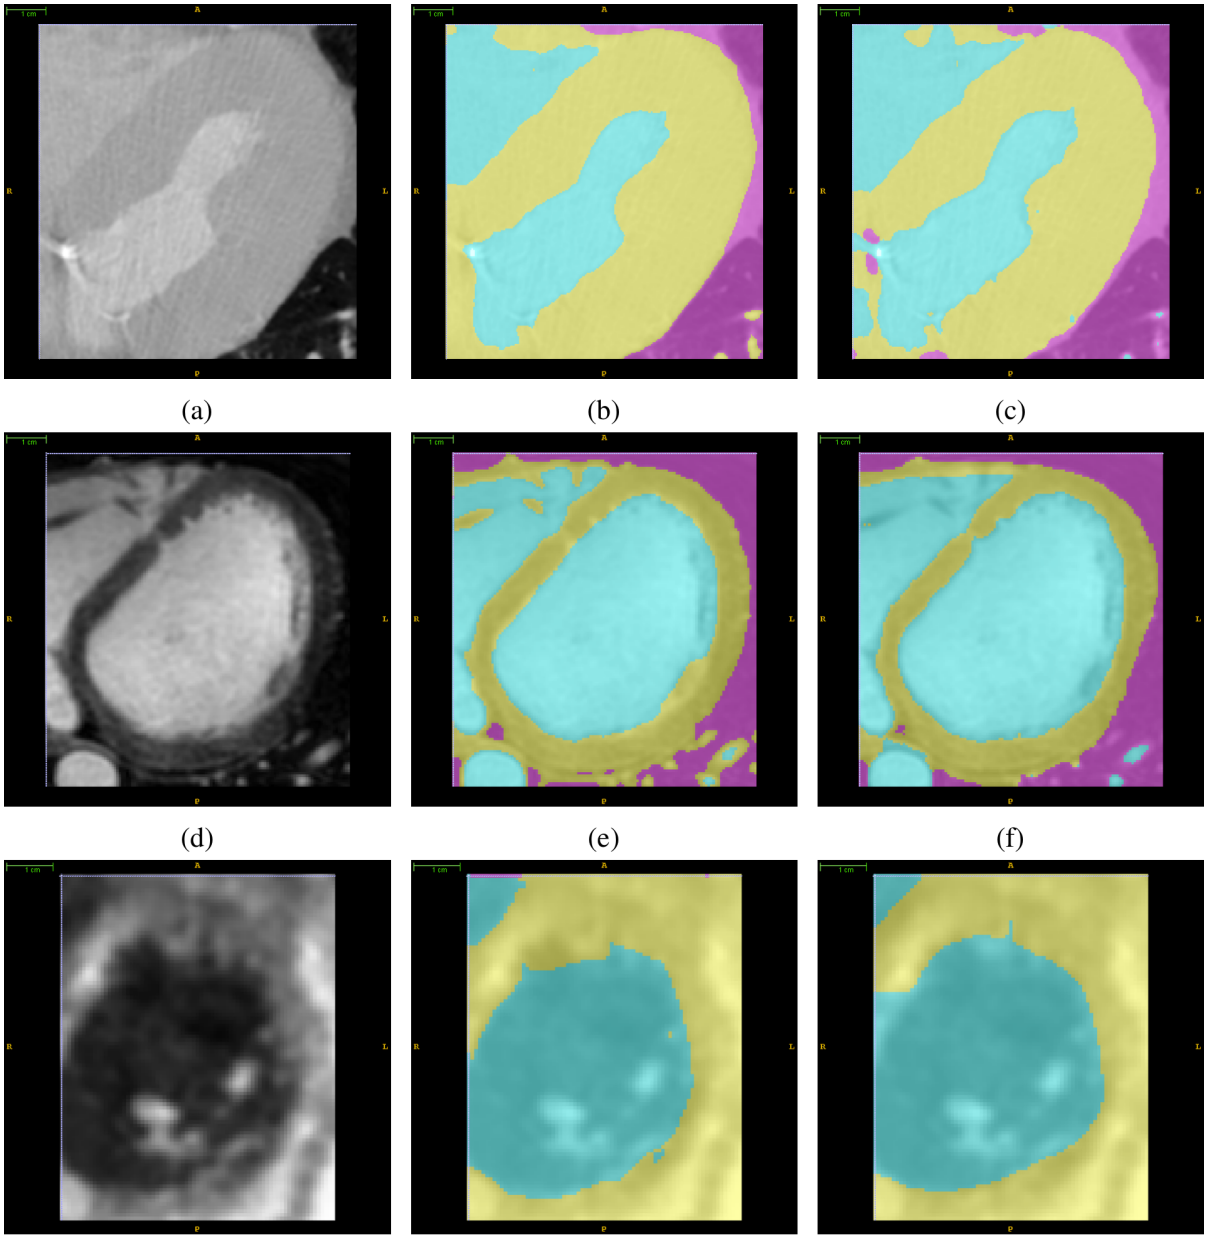

(a) (b) (c)

(d) (e) (f)

(g) (h) (i)

Figure 4.4: Cardiac segmentation with underlying (a) CT, (d) MRA, and (g) TEE. Manual segmentations are in (b), (e), and (h) respectively, and interactive segmentation results in (c), (f), and (i).

4.4.2 Neonatal Cranial MRI Segmentation

Neonatal brain images display some unique challenges for automated segmentation in that there are relatively few compared to adult brain images, making machine learning-based or atlas-based segmentation approaches infeasible. In addition, bleeds in the ventricular system further complicated segmentation. In this context, interactive interfaces can be extremely useful since manual segmentation or correction is largely unavoidable.

As part of a larger, ongoing, observational study on the potential for new technologies (near infrared spectroscopy, 3D ultrasound, and functional MRI) to allow for better monitoring of preterm neonates with post haemorrhagic ventricle dilatation, patients were enrolled after informed parental consent once a positive diagnosis of intraventricular haemorrhage (IVH) was made on routine clinical head ultrasound. Protocols used in this study were approved by the Research Ethics Board of Western University as given in Appendix C Section C.2. Figure 4.5 displays visual results of neonatal ventricle segmentation using this interface.

To demonstrate the interactive segmentation interface's robustness to pathology, the previous experiment was extended to a neonatal MR image in which a severe ventricular bleed changes the intensity distribution of the ventricle to an extreme degree. The segmentation re-

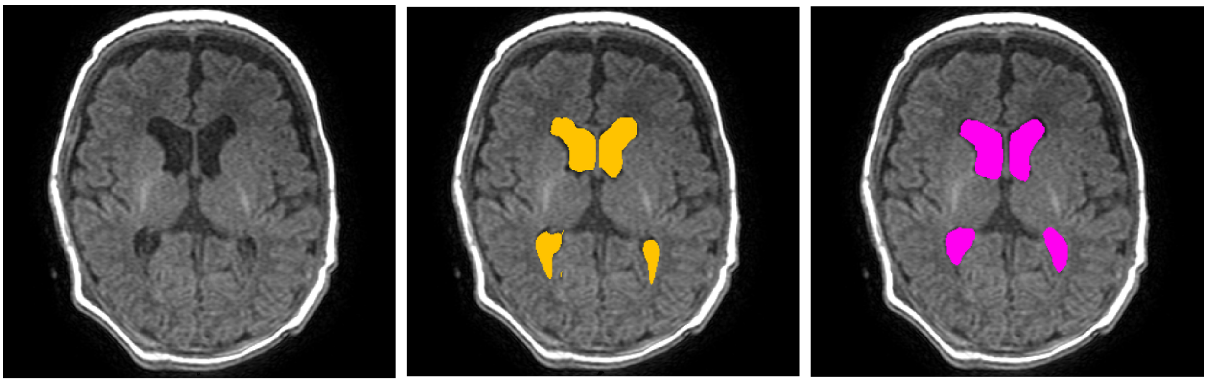

(a) Original MR Image (b) Manual Segmentation (c) Interactive Segmentation

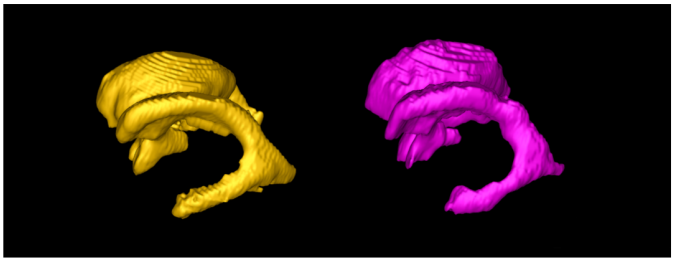

(d) Surface renderings

Figure 4.5: Neonatal Ventricle Segmentation with (a) the MR, (b) the manual segmentation, and (c) interactive segmentation results. (d) shows surface renderings of both the fully manual (left) and interactive (right) segmentation results.

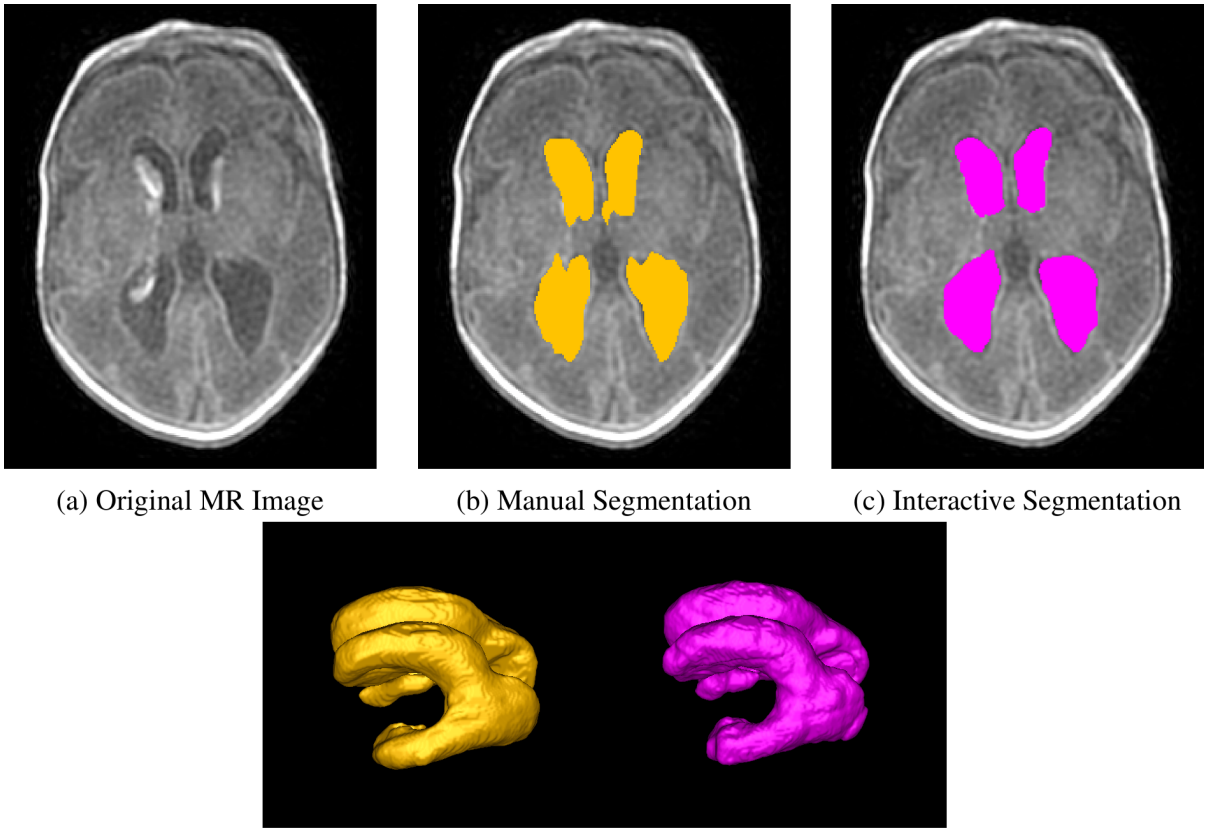

(a) Original MR Image (b) Manual Segmentation (c) Interactive Segmentation

(d) Surface renderings

Figure 4.6: Pathological Neonatal Ventricle Segmentation with (a) the MR, (b) the manual segmentation, and (c) interactive segmentation results. (d) shows surface renderings of both the fully manual (left) and interactive (right) segmentation results.

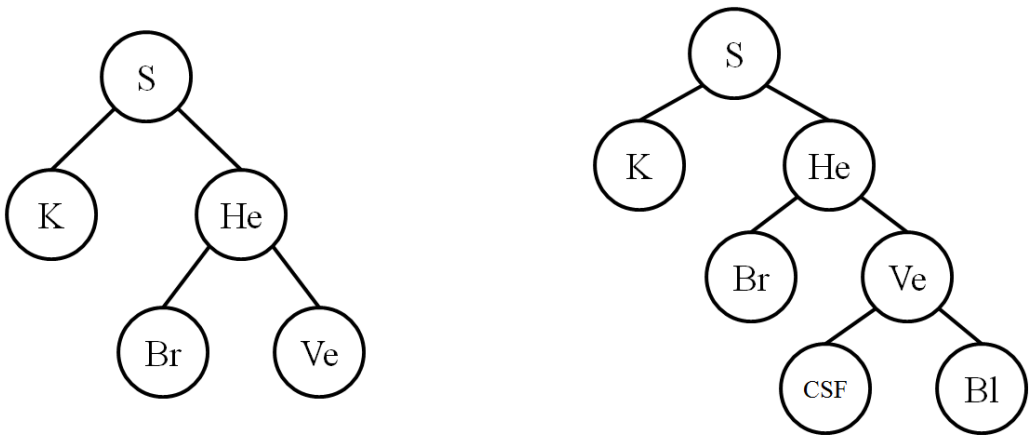

(a) Healthy neonate (b) Neonate with intraventricular haemorrhage

Figure 4.7: Hierarchies used in (a) healthy and (b) pathological neonatal ventricle segmentation.

sults are given in Figure 4.6. Note that the hyper-intense ventricular bleed is closer in intensity to white- and grey-matter than to the ventricles, and it's appearance on the boundary of the ventricles would likely cause severe registration artefacts The segmentation of the ventricle was achieved by partitioning it into two components; a healthy component (CSF) and the ventricular bleeding (Bl). The union of these components could then be regularized similar to the ventricle (Ve) in Figure 4.5. The Ve label (the union of the CSF and Bl labels) for the pathological case is given in Figure 4.6 In the hierarchies used in this segmentation problem, which are given in Figure 4.7, the remaining labels are K which refers to the background, He to the head, and Br to the brain.

4.5 Automatic Hierarchy Refinement

Although determining an appropriate hierarchy merely from grouping information is a computationally difficult problem, due to the mathematical formulation, it is possible to automatically

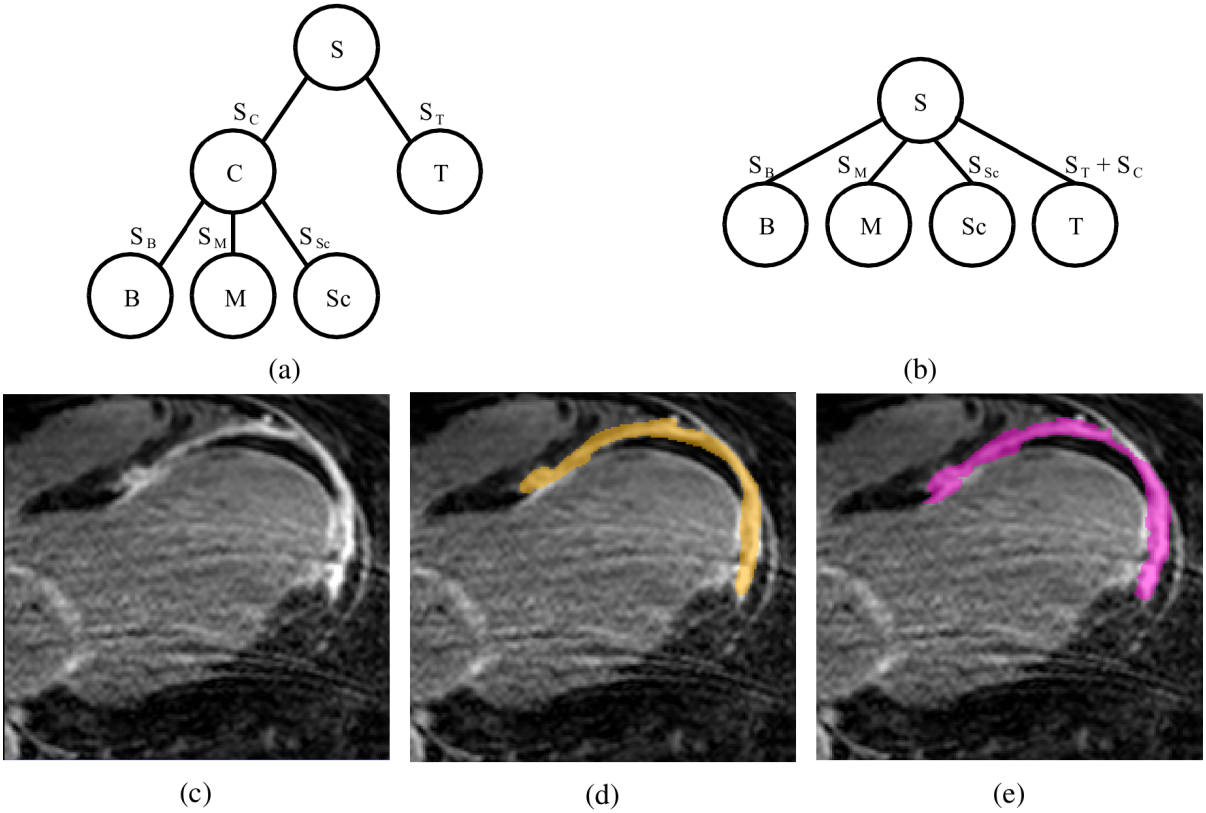

Figure 4.8: Example of automatic hierarchy segmentation. (a) is the original hierarchy reproduced from [150] and (b) the optimized version, (c) an LGE-MRI with (d) manual segmentation and (e) interactive segmentations results.

(n = 10)		Accuracy
Scar AVD	(%)	26.9 ± 15.6
Scar rMSE	(mm)	1.30 ± 0.32
Scar DSC	**(%)**	**74.1 ± 3.5**
Scar DSC from [150]	**(%)**	**76.0 ± 3.0**
Inter-operator variability from [150] - Scar DSC	**(%)**	**76.2 ± 2.6**
Intra-operator variability from [150] - Scar DSC	**(%)**	**75.2 ± 2.8**

Table 4.4: Scar Tissue Segmentation Results. Results shown in **bold** indicate metrics that are common across SEGUE, manual segmentations, and [150] which can be used as a reference for comparison

refine a user-provided hierarchy for improved computational efficiency without compromising segmentation quality. This involves the contraction and removal of vertices in the hierarchy with zero regularization or where zero regularization can be induced without changing the optimization functional. One specific example where automatic hierarchy refinement is possible is when the source node has only two children, whereby one can be contracted by transferring its smoothness value to the other. To demonstrate this, we performed automatic hierarchy optimization on the method presented by Rachl *et al.* [150] using Late Gadolinium Enhanced Magnetic Resonance Imaging (LGE-MRI). This segmentation problem involved partitioning the image into thoracic background (T) and cardiac (C) labels, the latter being subdivided into blood (Bl), healthy myocardium (M) and scar tissue (Sc) as shown in Figure 4.8a. The cardiac label, C, was automatically contracted, resulting in Figure 4.8b and an estimated 20% improvement in speed.

As with the previous experiments, the results (recorded in Table 4.4) were comparable to those presented by Rachl *et al.* [150] and was within the range of inter-operator variability. However, this level of accuracy was achieved without post-processing steps, such as connected components analysis, or other modifications that would make the interface specific to cardiac or LGE-MRI segmentation.

4.6 Discussion

Interactive segmentation helps bridge the gap between manual and automatic segmentation allowing each to address the weaknesses of the other. Improvements in interactive segmentation interfaces can have a distinct impact in clinical contexts in which automated segmentation is not feasible. Several clinical applications require manual segmentation due to pathology such as tumours in radio-oncological applications or bleeds in neonatal cranial imaging. These applica-

tions require a user to manually delineate some anatomy in order to perform relevant measurements such as tumour volume. In these applications, accurate segmentation may be necessary for robust, correct measurements, and the use of interactive segmentation can have a distinct benefit, conserving user time while encouraging accurate results, which will in turn improve patient outcomes by improving the diagnostic capabilities of these measurements (compared to manual segmentation) in single acquisition and longitudinal studies.

This interface allows for the user to define a segmentation hierarchy in run-time, taking advantage of a fast, GPU-accelerated general HMF solver, which in turn allows for more knowledge of spatial relationships between anatomical regions to be encoded. This encourages the use of optimization techniques and interactive interfaces in which a user can quickly define and correct a segmentation, and thereby increase the speed, quality, and robustness of general segmentation tasks. The ability to modify the hierarchy in run-time allows for the interactive segmentation interface to account for extreme deviations, such as ventricular bleeds, by the addition of multiple labels to account for them. This interface is the first to allow the user to modify the abstract anatomic knowledge, i.e. label ordering, provided to the computer in run-time.

The primary advantage of this interface over other interactive segmentation programs is that it allows the user to interactively specify both segmentation hierarchy and initial seeds. The former means that the interface is very general purpose, allowing for arbitrary regions to be defined, while incorporating anatomical knowledge in a direct manner. This gives it a distinct advantage over other interactive segmentation interfaces which either limit the number or type of regions, or do not allow the user to specify abstract anatomical knowledge. The latter takes advantage of a paint-brush mechanism which allows for large regions of the interior of the object to be seeded with minimal user effort thus improving the probabilistic data terms.

The second major advantage is that the algorithm is founded in optimization principles, ensuring robustness and repeatability across images. The formulation of the costs also allow for the regional and boundary uncertainty (U_R and U_B) identified in [171] to be actively addressed by the segmentation process, making plane selection simpler and more efficient. Plane selection is further improved by selecting only axis-aligned planes in which the user is accustomed.

4.6.1 Clinical Integration

In terms of integrating this type of interaction into clinical use, one likely candidate is radio-oncology. The nature of radio-therapy segmentation problems require multi-region segmentation in which there is a large degree of hierarchical structure to the anatomy. Nevertheless, user interaction may be required for initially locating/seeding tumours as well as correcting

the resultant segmentation. Some multi-region interactive segmentation frameworks for radiotherapy have been proposed [45] but lack the ability to apply structure to the anatomy being segmentation. The use of structured anatomy, known as *ontologies*, has recently gained support in cancer radiology and may provide initial segmentation hierarchies suitable for manipulation to suit the individual patient [132].

From the standpoint of traditional radiology, the goal of moving towards more user-friendly knowledge-centric interaction paradigms may lead to more divestment of responsibilities from the radiologist to the home department. For example, cardiologists may directly interact with cardiac imaging datasets without deferring to specialized radiologists. Thus, the goal of these interfaces is not to replace the radiologist, but to render radiological processes and decision-making more transparent to non-radiological clinicians. In addition, there is an increasing number of tasks involving medical image segmentation that are not the purview of radiology. Image guided surgical planning, for example, requires the segmentation of structures of interest either for targeting or for avoidance.

4.7 Future Work

There are several future directions in which to take this work aside from general improvements to computational resource usage and performance. Specifically:

- Incorporation of a more extensive model of label organization,
- Incorporation of geometric or shape constraints,
- Improvements to the definition of the smoothness model, and
- Improvements to the plane selection mechanism.

Recently, work has been performed which extends the possibility of label organization in continuous max-flow from hierarchical models to models that allow for any possible label ordering. (See Chapter 5) However, there remain issues in terms of how these structures can be specified by a user in run-time in an intuitive manner as they are defined using a constrained set of rooted, weighted directed acyclic graphs, which do not have a user-friendly tool already in place.

There has also been increasing interest in the use of generic geometric or shape constraints such as star-shaped priors in both graph-cuts [178] and max-flow image segmentation. [196] Shape complexes have already been proposed which combine the notions of label orderings and star-convex object constraints to develop complicated models of object geometry from

the union and disjunction of star-convex objects. (See Chapter 6.) Such frameworks can be readily incorporated into this interactive segmentation framework with minimal changes to the interface or usability, while contributing a significant improvement to the segmentation accuracy through the encoding of additional anatomical knowledge.

Currently, the interface allows the user to modify the parameters in the smoothness term, but does not permit any other manipulation. This could be incorporated through the addition of a contouring mechanism similar to that in Intelligent Scissors, TurtleSeg, and ITKSnap. These contours could supply specific information which can improve the smoothness terms, as well as give the user complimentary ways to sample regions.

In terms of plane selection, future work could include defining a sequence of planes sensitive to the distance between them, rather than a single set. This would allow the algorithm to intelligently inform the user of multiple areas of uncertainty without re-invoking the continuous max-flow segmentation algorithm and allow the user to provide feedback on multiple high uncertainty planes in a single interaction cycle.

Chapter 5

Directed Acyclic Graph Continuous Max-Flow Image Segmentation

This chapter is largely based on:

- John SH Baxter, Rajchl, Martin, A. Jonathan McLeod, Jing Yuan, and Terry M. Peters. "Directed Acyclic Graph Continuous Max-Flow Image Segmentation for Unconstrained Label Orderings." International Journal of Computer Vision (IJCV), (2017).

with additional material from:

- John SH Baxter, Martin Rajchl, Jing Yuan, and Terry M. Peters. "A continuous max-flow approach to multi-labelling problems under arbitrary region regularization." arXiv preprint arXiv:1405.0892 (2014).
- John SH Baxter, Martin Rajchl, Jing Yuan, and Terry M. Peters. "A proximal bregman projection approach to continuous max-flow problems using entropic distances." arXiv preprint arXiv:1501.07844 (2015).

5.1 Introduction

Chapter 3 motivated the development of a generalized solver for continuous max-flow problems in which the labels are equipped with a hierarchical topology. Specifically, the concept of incorporating a clinical user's understanding of the anatomy into the segmentation process through specifying a label ordering proved to be a powerful approach in interactive segmentation, which is then demonstrated in Chapter 4. However, certain immediate questions remain such as ***Can all anatomy be described hierarchically?*** and ***Can users reconcile two different anatomical models?*** These questions are related in that it is mathematically possible to construct two hierarchical models that are not reconcilable as a hierarchy. If these cases occur

in medical imaging, then it becomes necessary to further generalize what label orderings are soluble to non-hierarchical cases.

For an example of this, consider a segmentation problem in which Vessel A bifurcates into Vessels B and C as shown in Figure 5.1. In this imaging modality, there is high contrast between the vessel wall and the lumen implying that segmentation would benefit from modelling each vessel as having a wall component and a lumen component. If this problem was given to a clinician, they may chose to represent this anatomy in multiple ways. For example, one may chose to first model the bifurcation, that is, represent the greater vessel as first being decomposed into Vessels A, B, and C, followed by each vessel being represented by its respective wall and lumen components. (This hierarchy is shown in Figure 5.2a.) Another may chose to reverse this order, first representing the entire vessel as having a shared wall and a shared lumen, both of which are partitioned between Vessels A, B and C. (This hierarchy is shown in Figure 5.2b.)

One could readily argue that both orderings are valid, representing different facets of the segmentation problem. (The first encourages each vessel segment to be contiguous, and the second encourages the walls and lumens to smoothly transition between vessels.) However,

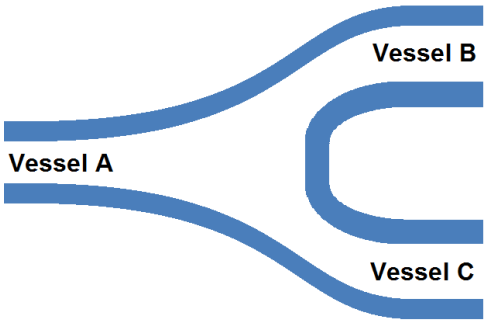

Figure 5.1: Schematic vessel bifurcation

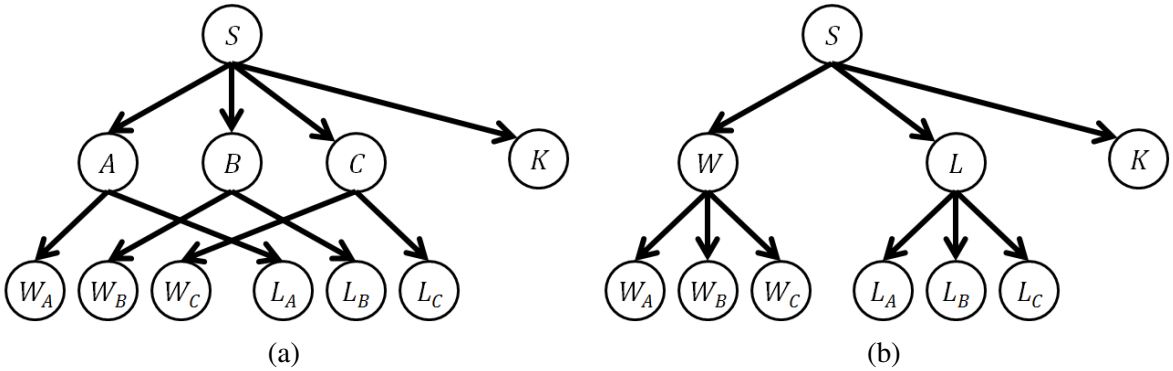

Figure 5.2: Two possible hierarchical label orderings for the vessel bifurcation segmentation example shown in Figure 5.1.

these models are irreconcilable; they cannot both be incorporated into a single larger hierarchy.

Given the description in Section 3.2.1 of a label ordering over a set of partition labels, $\mathbb{L}$, the only additional constraint placed on label orderings to make them hierarchical is that the parent operator, $L.P$, returns a singleton for every non-source label, $L \neq S$. Following that line of thought, if the constraint that $L.P$ is a singleton is relaxed, any label ordering should be possible to implement. Thus, any label ordering should be mathematically representable as a *rooted directed acyclic graph* with S as the root (only label with no parent) and $\mathbb{L}$ as the leaves (only labels with no children). Each intermediate node represents a group with its own distinct regularization requirements. The goal of this chapter is to build on that intuition to create a more generalized continuous max-flow solvable suitable for any arbitrary label ordering.

This chapter begins with a summary of graph-cuts and max-flow frameworks that allow for label ordering outside of discrete, linear, cyclic, and hierarchical topologies. Directed Acyclic Graph continuous max-flow (DAGMF) segmentation is presented as an extendable framework for incorporating label orderings defined over a directed acyclic graph. DAGMF is then validated on a variety of synthetic, natural, and medical image segmentation problems.

5.2 Previous Graph-Cuts and Max-Flow Approaches with More Complex Topologies

Aside from the constrained extendible topologies illustrated in previous chapters (See Chapter 2 Sections 2.2.3 to 2.2.5 and Chapter 3), there are a couple of max-flow/min-cut frameworks which have topological considerations in which the constraints are more difficult to describe.

5.2.1 Submodular Graph Construction

A conceptually important prior work in max-flow/min-cut segmentation approaches with more complex topology are the *submodular graph constructions* proposed by Delong *et al.* [43]. In these constructions, label orderings could be defined in terms of two relationships: *exclusion*, implying that the two labels are disjoint, and *containment*, implying that one label is a superset of the other. The label ordering could then be expressed as a graph with two edge types representing the two different types of relationships. The label ordering could be solved with global optimality provided that the graph could be two-coloured (white and black) with the colours being different across each *exclusion* edge and the same across each *containment* edge. (If this cannot be done, it is due to one or more *frustrated cycles* in the graph.)

Through the appropriate operations on this graph (extended over the number of voxels in the image) the entire segmentation problem can be represented as a single min-cut problem and

solved using traditional graph-cut solution algorithms. One of these operations, for example, is to *flip* the black labels, (i.e. explicitly optimizing $u_{\neg L}(x) = 1 - u_L(x)$ rather than $u_L(x)$) which ensures the entire graph construction remains submodular.

It is trivial to demonstrate that these submodular constructs strictly generalize Ishikawa models. However, it is worth noting that submodularity does constrain them in that Potts models are not strictly generalized. That is, all Potts models with more than two labels must contain a frustrated cycle.

5.2.2 Generalized Ordering Constraints in Continuous Min-Cut

Strekalovskiy *et al.* [168] developed a highly general variational minimization framework for a large number of topological considerations applied in a partitioning problem. These considerations are implemented in a single complex regularization function sensitive to both label change and directionality. This regularization function can be expressed as:

$$\begin{aligned} R(u) &= \sup_{q \in C} \sum_{L \in \mathbb{L}} \int_\Omega q_L(x) \cdot \nabla u_L(x) dx \\ \text{where } C &= \left\{ (q_L)_{L \in \mathbb{L}} : \Omega \to (\mathbb{R}^D)^{|\mathbb{L}|} \mid \forall i, j \in \mathbb{L}, \forall v \in \mathbb{R}^D,\ (q_i(x) - q_j(x)) \cdot v \leq d(x, i, j, v) \right\} \\ \text{such that } & \sum_{L \in \mathbb{L}} u_L(x) = 1 \end{aligned} \tag{5.1}$$

where the function $d(x, y, i, j, v)$ expresses any topological considerations. This function is subject to the constraints:

- $d(x, i, j, v)$ grows linearly with the size of the direction vector, i.e. $d(x, i, j, tv) = td(x, i, j, v)$ where $t \geq 0$
- $d(x, i, j, v)$ satisfies the triangle inequality with respect to direction, i.e. $d(x, i, j, v + w) \leq d(x, i, j, v) + d(x, i, j, w)$
- $d(x, i, j, v)$ is symmetric with respect to the labels, i.e. $d(x, i, j, v) = d(x, j, i, v)$
- $d(x, i, j, v)$ satisfies the triangle inequality with respect to labels, i.e. $\forall k \in \mathbb{L},\ d(x, i, k, v) \leq d(x, i, j, v) + d(x, k, j, v)$

which imply its solvability. This allows it to address a number of different problems, including incorporating topological considerations such as tiered and floating object representation [168]. However, it may be unintuitive to design this function as the label ordering is implicitly, rather than explicitly represented. This also may lead to difficulty in developing fully primal-dual

solution algorithms (rather than using duality solely for the regularization term) as it profoundly couples the continua representing each labelling in the spatial domain.

5.3 Directed Acyclic Graph Max-Flow Formulation

The Directed Acyclic Graph Max-Flow (DAGMF) segmentation model this article addresses is the minimization of the convex energy functional:

$$\begin{aligned} &\min_{u} \sum_{L} \int_{\Omega} \left(D_L(x) u_L(x) + |\mathbf{R}_L(x) \nabla u_L(x)|_p \right) dx \\ &\text{s.t. } u_S(x) = 1 \\ &\quad \sum_{L' \in L.C} w_{(L,L')} u_{L'}(x) = u_L(x) \text{ if } L \notin \mathbb{L} \\ &\quad u_L(x) \geq 0 \end{aligned} \tag{5.2}$$

Similar to hierarchical max-flow (HMF), the partition $\mathbb{L} = \{L | L.C = \emptyset\}$ is the set of leaf-labels or labels with no children. Unlike HMF, $L.P$ does not have to consist of a single element, but each label can have multiple parents. The non-negative weight terms $w_{(L',L)}$ determine what weight to give the labelling function of L when calculating the labelling function of its parent label, L'. These terms must sum to 1 over the parents of each label, that is:

$$\sum_{L' \in L.P} w_{(L',L)} = 1 \tag{5.3}$$

It is clear that HMF is a subset of these models as $L.P$ having a single element implies $w_{(L',L)} = 1$ showing that equation (3.1) is a special case of equation (5.2). Without loss of generality, assume that $D_L(x) = 0$ whenever $L \notin \mathbb{L}$. This is achievable by pushing down the data terms to the labels children, i.e. $D_L(x) u_L(x) = \sum_{L' \in L.C} D_L(x) w_{(L,L')} u_{L'}(x)$, similar to the procedure outlined in Chapter 3 Section 3.4. The constraint on the weight function is necessary for a valid partition of Ω into the leaf-labels, that is, $\sum_{L \in \mathbb{L}} u_L(x) = 1$. However, this does not limit the generality of the method as will be demonstrated in Section 5.3.1.

The modelling approach is derived from those presented by Yuan *et al.*[192][193] and follows the same format, showing the duality of a max-flow primal formation to this minimization problem through an intermediate primal-dual optimization problem. An augmented Lagrangian framework and a proximal Bregman framework are proposed for minimizing this intermediate representation.

5.3.1 Arbitrary Region Regularization

With the increasing complexity of part-whole and mutual exclusion relationships being integrated into optimization based segmentation, it is important to show the capability of DAGMF to address the regularization of arbitrarily defined label structures, that is, the incorporation of any label ordering. First, one must express the segmentation in terms of its partition set, $\mathbb{L}$, and a finite set of grouped-labels $\mathbb{G} \subset 2^{\mathbb{L}}$, each a union of a set of leaf-labels, defining the groups in which a regularization term should apply. For brevity, consider the segmentation problem to be stated in this form.

To show how arbitrary region regularization can be implemented with DAGMF, consider the construction of a DAG with associated transformations on smoothness parameters. First, create a graph with one vertex corresponding to each element of $\mathbb{L}$ and one for each element of $\mathbb{G}$ and one for the source node, S. Create an edge from the source vertex to each vertex corresponding to an element of $\mathbb{G}$. For each element $G \in \mathbb{G}$, create an edge from vertex corresponding to G to each vertex corresponding to an element $L \in G$. Now, find the vertex with the maximal in-degree which corresponds to an element in $\mathbb{L}$ and call this in-degree r. Add sufficient edges from the source vertex to each vertex in $\mathbb{L}$ to ensure that the in-degree for each is r. Figure 5.3a shows an example of this for the problem $\mathbb{L} = \{A, B, C, D, E\}$ with $\mathbb{G} = \{AB, BC, CD\}$.

This process yields an unweighted multi-graph in which the multiplicity of any given edge is indicate of the weight to be assigned to the corresponding edge in the weighted DAG proportional to adjacent incoming edges. Taking the example used in Figure 5.3a and applying the required normalization yields that shown in Figure 5.3b.

To determine the appropriate smoothness terms, note that for each vertex associated with

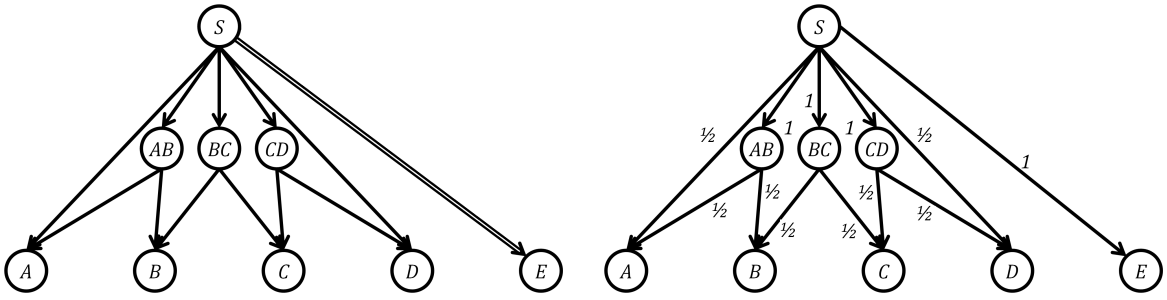

(a) DAG with edge multiplicity (b) DAG with weighted edges

Figure 5.3: DAG for segmentation into labels $\mathbb{L} = \{A, B, C, D, E\}$ in which label groups $\mathbb{G} = \{AB, BC, CD\}$ are regularized. Note that this would be impossible in a hierarchical model since the regularization groups conflict with each other. Figure 5.3a shows the intermediate multi-edged, unweighted DAG. Figure 5.3b shows this DAG with weights explicitly recorded rather than through multiplicity which is used by the solution algorithms.

element $G \in \mathbb{G}$ has the labelling constraint:

$$u_G(x) = \sum_{L \in G} \frac{1}{r} u_L(x) \tag{5.4}$$

from the construction of the graph. This means that whatever smoothness term that is desired for G must be multiplied by r to account for the factor of $1/r$.

This is not the only way to create a DAGMF model for a given problem, and many segmentation problems defined in such a way can be implemented by multiple DAGMF structures. In that respect, the DAGMF structure is not necessarily unique and optimization of the structure to improve performance may be possible.

5.4 Directed Acyclic Graph Max-Flow Algorithm

The DAGMF energy equation can be minimized using primal-dual optimization on a particular flow network which is derived from the label ordering diagram by replacing each label $L \neq S$ with a continuum, connecting those $L \in \mathbb{L}$ to a sink node. As with HMF, those sink node connections are then constrained by the data terms. (The HMF algorithm can be derived from the DAGMF algorithms under the assumption $w_{(L',L)} = 1$ and that the input flow to any node is the same as the output flow from a single parent node.) The algorithms presented in this section are based on the maximization of flow through these particular flow networks. Again, the main contribution of these *generalized* solution algorithms in particular is that the label ordering diagram, in the form of a rooted directed acyclic graph, can be given to them in run time and can reflect any label ordering.

The augmented Lagrangian solution algorithm to the DAGMF formulation is given in Algorithm 5.1. As with the previous augmented Lagrangian algorithms, it involves two positive parameters: τ, the gradient descent step size set to approximately $1/8$; and c, the quadratic penalty constant which has a default value of $1/4$.

The proximal Bregman solution algorithm to the DAGMF formulation Algorithm 5.4. Similar to the augmented Lagrangian algorithm, τ is the gradient descent step size used in the Chambolle iteration step which is set to approximately $1/8$. However, c is now the weight of the proximity term in each projection step which is set to be on the order of the difference in data terms.

Algorithm 5.1: Augmented Lagrangian solution algorithm for the DAGMF functional.

Topological sort the DAG into ordering $\mathbb{O}$ (begins with source label S) with reverse ordering $\mathbb{O}^{-1}$ (ends with source label S);
InitializeSolution() ;
while *not converged* **do**
  UpdateFlows() ;
  for $\forall L$ **do**
    $\forall x, u_L(x) \leftarrow u_L(x) - c\,(\operatorname{div} q_L(x) - \zeta_L(x) + p_L(x))$;
  end
end

Algorithm 5.2: **InitializeSolution()** subroutine in Algorithm 5.1.

InitializeSolution()
Clear $u_L(x), q_L(x)$ for all labels;
for *each L in order* $\mathbb{O}^{-1}$ **do**
  $\forall x, p_L(x) \leftarrow \min_{L'.C=\emptyset} D_{L'}(x)$;
  $\forall x, \zeta_L(x) \leftarrow \min_{L'.C=\emptyset} D_{L'}(x)$;
  if $L.C = \emptyset$ **then**
    if $L \in \operatorname{argmin}_{L'.C=\emptyset} D_{L'}(x)$ **then**
      $\forall x, u_L(x) \leftarrow 1/|\operatorname{argmin}_{L'.C=\emptyset} D_{L'}(x)|$;
    else
      $\forall x, u_L(x) \leftarrow 0$;
    end
  end
  for *each* $L' \in L.P/\{S\}$ **do**
    $\forall x, u_{L'}(x) \leftarrow u_{L'}(x) + w_{(L',L)} u_L(x)$;
  end
end

Algorithm 5.3: UpdateFlows() subroutine in Algorithm 5.1.

UpdateFlows()
for $\forall L \neq S$ **do**
  $\forall x, q_L(x) \leftarrow \text{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} \left(q_L + \tau \nabla \left(\text{div}\, q_L(x) + p_L(x) - \zeta_L(x) - u_L(x)/c \right) \right)$;
end
Clear $\zeta_L(x)$ for all labels ;
for *each L in order* $\mathbb{O}$ **do**
  for *each* $L' \in L.C$ **do**
    $\forall x, \zeta_{L'}(x) \leftarrow \zeta_{L'}(x) + w_{(L,L')} p_L(x)$;
  end
  if $L.C \neq \emptyset$ *and* $L.P \neq \emptyset$ **then**
    $\forall x, \sigma_L(x) \leftarrow \zeta(x) - \text{div}\, q_L(x) + u_L(x)/c$;
  else if $L = S$ **then**
    $\forall x, \sigma_S(x) \leftarrow 1/c$;
  end
end
for *each L in order* $\mathbb{O}^{-1}$ **do**
  if $L.C = \emptyset$ **then**
    $\forall x, p_L(x) \leftarrow \min\{D_L(x), \zeta_L(x) - \text{div}\, q_L(x) + u_L(x)/c\}$;
    for $L' \in L.P$ **do**
      $\forall x, \sigma_{L'}(x) \leftarrow \sigma_{L'}(x) + w_{(L',L)} \left(\text{div}\, q_{L'}(x) + p_{L'}(x) - \zeta_{L'}(x) + w_{L',L} p_L(x) \right)$;
    end
  else if $L = S$ **then**
    $\forall x, p_S(x) \leftarrow \frac{1}{\sum_{L' \in S.C} w^2_{(S,L')}} \sigma_S(x)$;
  else
    $\forall x, p_L(x) \leftarrow \frac{1}{1 + \sum_{L' \in L.C} w^2_{(L,L')}} \sigma_L(x)$;
    for $L' \in L.P$ **do**
      $\forall x, \sigma_{L'}(x) \leftarrow \sigma_{L'}(x) + w_{(L',L)} \left(\text{div}\, q_{L'}(x) + p_{L'}(x) - \zeta_{L'}(x) + w_{L',L} p_L(x) \right)$;
    end
  end
end

Algorithm 5.4: Proximal Bregman solution algorithm for the DAGMF functional.

Topological sort the DAG into the sorted list $\mathbb{O}$ (begins with source label S) with reverse list $\mathbb{O}^{-1}$ (ends with source label S);
$\forall L \in \mathbb{L},\ u_L(x) \leftarrow 1/|\mathbb{L}|$;
while *not converged* **do**
  $\forall L, \forall x,\ d_L(x) \leftarrow \operatorname{div} q_L(x)$;
  $\forall L \in \mathbb{L}, \forall x,\ d_L(x) \leftarrow d_L(x) + D_L(x)$;
  for L *in order* $\mathbb{O}/\{S\}$ **do**
    for $L' \in L.C$ **do**
      $\forall x,\ d_{L'}(x) \leftarrow d_{L'}(x) + w_{(L,L')} d_L(x)$;
    end
  end
  $\forall L \in \mathbb{L}, \forall x,\ u_L(x) \leftarrow u_L(x) \exp\left(-\frac{d_L(x)}{c}\right)$;
  $\forall L \in \mathbb{L}, \forall x,\ d_L(x) \leftarrow u_L(x)$;
  $\forall x,\ a(x) \leftarrow \sum_{L \in \mathbb{L}} u_L(x)$;
  $\forall L \in \mathbb{L}, \forall x,\ u_L(x) \leftarrow u_L(x)/a(x)$;
  $\forall L \notin \mathbb{L}, \forall x,\ d_L(x) \leftarrow 0$;
  for L *in order* $\mathbb{O}^{-1}/\{S\}$ **do**
    $\forall x,\ q_L(x) \leftarrow \operatorname{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} \left(q_L(x) - c\tau \nabla d_L(x)\right)$;
    for $L' \in L.P/\{S\}$ **do**
      $\forall x,\ d_{L'}(x) \leftarrow d_{L'} + w_{(L',L)} d_L(x)$;
    end
  end
end

5.5 Validation

5.5.1 Synthetic Image: Venn Diagram

To test the efficacy of the segmentation algorithm and to highlight the improvements of arbitrary region regularization made possible through DAGMF, a synthetic segmentation example, shown in Figure 5.4, was constructed. This image involves two overlapping objects, each of which is regularized by the proposed DAGMF model using the structure showing in Figure 5.5, but such a regularization scheme is not possible in HMF, as the dark yellow intersection has more than one parent region. (In HMF, the intersection can be denoted either as part of the red square or green circle but not both.) The result is that in HMF, there may be erroneous isolated regions as seen in Figure 5.5e for the red label. The data term for each leaf-label is:

$$D_L(x) = |I(x) - \bar{I}_L| \tag{5.5}$$

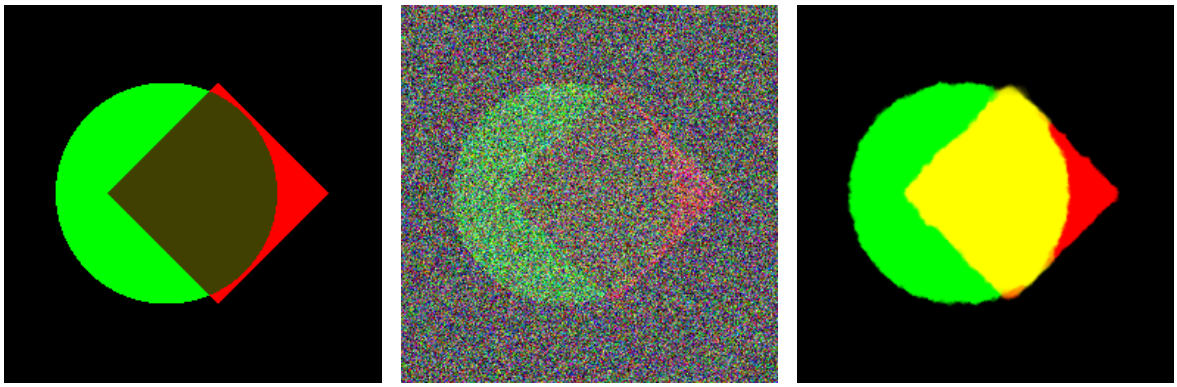

(a) Original (b) Noisy (c) Potts (DSC 94.6%)

(d) DAGMF (DSC 95.4%) (e) HMF #1 (DSC 95.0%) (f) HMF #2 (DSC 94.9%)

Figure 5.4: Synthetic image (a) polluted with noise (b) and reconstructed using a Potts model (c), DAGMF (d) and HMF models with either the red square (e) or green circle (f) regularized.Weighted DSC is given for each segmentation.

where $I(x)$ is the RGB intensity of the pixel and $\bar{I}_L$ is the RGB intensity of the label L prior to the addition of noise. The regularization terms $R_L(x)$ were all constants (times the identity matrix), tuned individually to improve the performance of each method. Note that the DAGMF reconstruction substantially improves the crispness of the area where all four regions meet which is heavily distorted in the other three segmentation techniques.

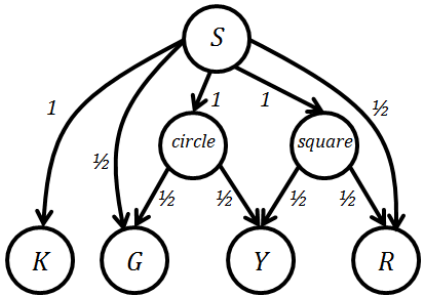

Figure 5.5: Directed acyclic graph and weights used for DAGMF segmentation shown in Figure 5.4. The nodes *circle* and *square* denote the labels associated with the union of green with yellow and red with yellow respectively.

5.5.2 Medical Images - Brain Tissue Segmentation

To demonstrate the applicability of this technique in medical image segmentation, consider segmentation of the brain into background (K), external cerebrospinal fluid ($eCSF$), cortical gray matter (cGM), white matter (WM), subcortical gray matter (sGM) and ventricles (V). This example uses the BrainWeb [37] database to provide a realistic digital phantom, providing an exact ground truth labelling. The data terms (shown in Fig. 5.7) were developed using a

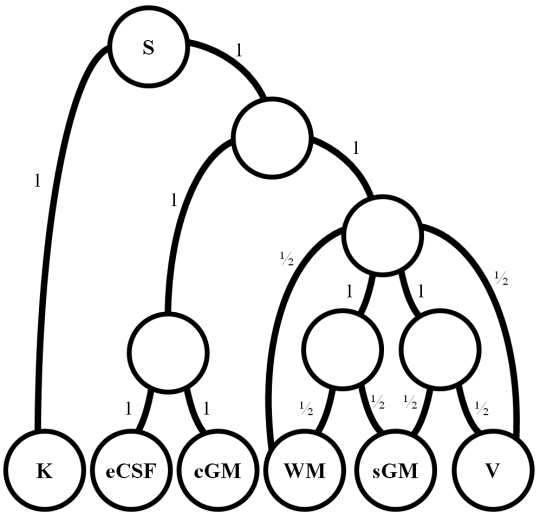

Figure 5.6: DAG representing the brain tissue segmentation problem in Fig. 5.8.

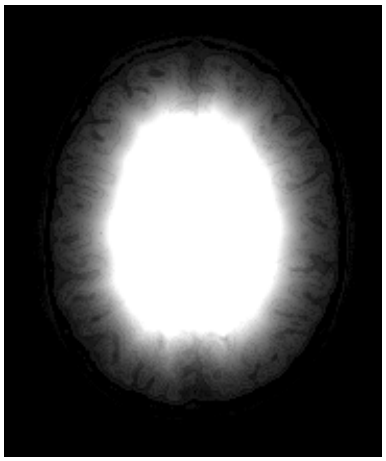

K Data Term

eCSF Data Term

cGM Data Term

WM Data Term

sGM Data Term

V Data Term

Figure 5.7: Bayesian data terms used in Fig. 5.8.

simple Bayesian framework which includes both intensity and localization components and are representative of those one would derive for this segmentation problem on clinical images [176]. The DAG used for segmentation is provided in Figure 5.6. The smoothness terms were all constants, that is, $R_L(x) = \alpha_L$, meaning that no contrast sensitive terms were used to localize edges, but only uniform regularization was used to penalize longer boundaries between regions. The data terms, shown in Figure 5.7, follow the Bayesian formula:

$$D_L(x) = -\ln P(I(x)|x \in L) + \text{dist}(x, R_L) \tag{5.6}$$

where $P(I(x)|x \in L)$ is the probability of voxel x having intensity $I(x)$ given it is in label L, and $\text{dist}(x, R_L)$ is an estimated distance between x and the region associated with label L. The segmentation is shown in Figure 5.8.

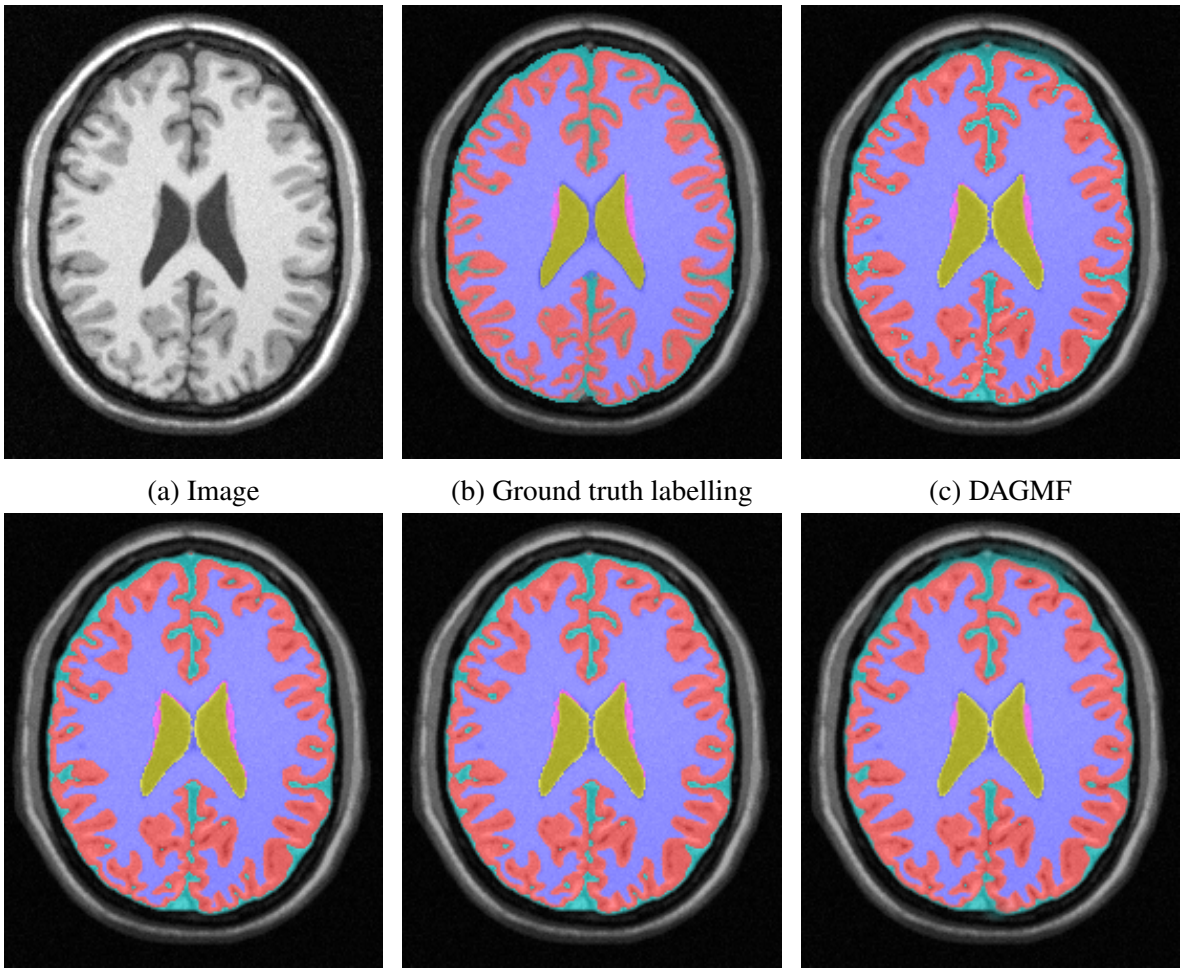

(a) Image (b) Ground truth labelling (c) DAGMF

(d) HMF (*sGM* grouped with *V*) (e) HMF (*sGM* grouped with *WM*) (f) Potts Model

Figure 5.8: Brain tissue segmentation using DAGMF using data terms in Fig. 5.7 and constant smoothness terms. Note the improvement in the pink subcortical gray matter region.

Region	DAGMF	HMF 1	HMF 2	Potts
K	99.2%	99.2%	99.2%	98.8%
eCSF	74.9%	76.5%	76.5%	72.5%
cGM	91.1%	91.0%	91.0%	90.6%
WM	95.9%	95.7%	95.7%	95.9%
sGM	**81.3%**	**76.1%**	**77.3%**	**76.6%**
V	95.7%	96.1%	95.9%	95.1%

Table 5.1: Dice coefficient for segmentations in Fig. 5.8. The results for the subcortical gray matter are shown in **bold** which reflect the quantitative improvement from using a more nuanced regularization model with DAGMF.

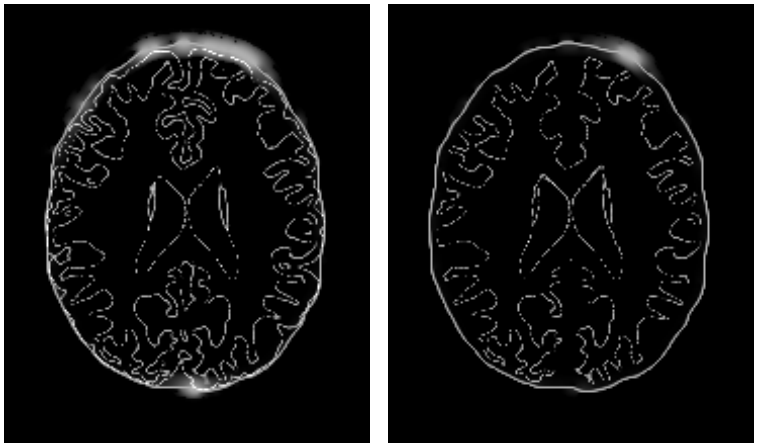

(a) Potts Entropy (b) DAGMF Entropy

Figure 5.9: Segmentation uncertainty (entropy) from Figure 5.8. The Potts model has much higher uncertainty in the background segmentation around the frontal lobe.

The construction of the DAG is motivated by the grouping of adjacent regions with similar regularization requirements (such as the grouping of the $eCSF$ and cGM) or by grouping regions that, when grouped, have a significantly less tortuous boundary (such as the grouping of all brain regions into a common superlabel opposed to the background or K label). Notably, a Venn diagram model (similar to Fig. 5.5) represents the subcortical labels, using the geometrical intuition that the subcortical gray matter could be interpreted spatially as the overlap of extended WM and V regions.

The primary benefit of DAGMF in the context of medical image segmentation in particular is that the regularization of any label group allows for the segmentation algorithm designers to incorporate multiple regularization schemes without prohibiting others. Analogous to Section 5.5.1, HMF could be used to regularize sGM with either V or WM, but not both. DAGMF allowed for both regularization schemes to be implemented simultaneously. This improved the segmentation accuracy of the subcortical gray matter over the Potts model and both HMF models as highlighted in Table 5.1, despite having the same data terms and uniform regularization.

In addition, the DAGMF result had a lower entropy, implying that the segmentation was more certain than that of Potts as shown in Figure 5.9. The entropy of the fuzzy labelling at each voxel can be used as a surrogate for the uncertainty of the binary labelling marginalized at each voxel. Note that this is an approximation as the theoretical basis of segmentation uncertainty in (both binary and fuzzy) continuous max-flow models has yet to be investigated.

Although applying only a single image, this experiment illustrates the increased capability of DAGMF in terms of more nuanced regularization and descriptive label orderings. This allows it to be readily expressed in problems where prior methods force fundamental tradeoffs between competing label orderings by combining their concerns.

5.5.3 Natural Images: Scene Decomposition

Natural image segmentation tasks that could benefit from unconstrained label orderings include scene decomposition [72]. As shown by Delong *et al.* [43], geometric context often requires non-hierarchical regularization schemes. This segmentation problem is composed of five distinct regions (F-front, T-top, B-bottom, L-left, R-right) as shown in Figure 5.10a. Strong priors are available for the T, B, L and R regions but often not the F region. For regularization purposes, one could consider F as the intersection of each of the four side regions, taking advantage of their priors and encourage a more central position in the segmented image. In previous approaches, such as HMF, such regularization would not be possible. The DAG representing this regularization structure is given in Figure 5.10b. Unlike previous discrete approaches [43], the image is partitioned into the label regions $\{T, B, L, R\}$ thus preventing errors such as labelling a single pixel as being in both L and R regions.

An example segmented image using this framework is shown in Figure 5.11. These images were collected from the Geometric Context dataset [72] and the data terms used were derived from the Surface Layout classification framework [73]. Specifically, the data terms for each region in the partition, i.e. labels $L \in \mathbb{L}$, are:

$$D_L(x) = -\ln P_L(x) \tag{5.7}$$

where $P_L(x)$ is the likelihood estimate of pixel x belonging to label L as determined by the Surface Layout classification framework [73]. The regularization terms, $R_L(x)$ were all the same constant (times the identity matrix) for the regions LF, RF, BF, and TF and were zero

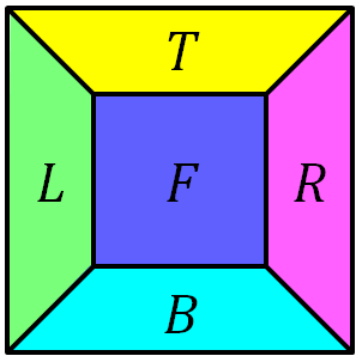

(a) Scene Decomposition

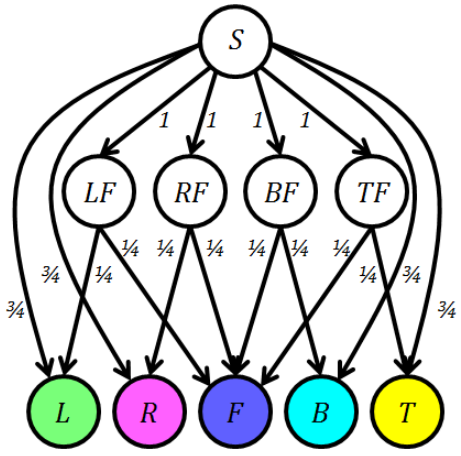

(b) Segmentation DAG

Figure 5.10: Segmentation structure used in scene decomposition into F-front, T-top, B-bottom, L-left, R-right. The color code corresponds to that used in Figure 5.11.

(a) Original Image (b) Original Image

(c) Segmentation (Accuracy rate = 93.0%) (d) Segmentation (Accuracy rate = 90.5%)

Figure 5.11: Example outdoor scene segmentation. Accuracy rate is given for each segmentation. The color code for the segmented images are shown in Figure 5.10.

Model - Label	Value
Potts	10
HMF 1 - TFB	5
HMF 1 - $L \in \mathbb{L}$	7.5
HMF 2 - LFR	5
HMF 2 - $L \in \mathbb{L}$	7.5
DAGMF - LF, RF, BF, TF	5
DAGMF - $L \in \mathbb{L}$	9

Table 5.2: Value of the constant regularization terms used in the various max-flow models.

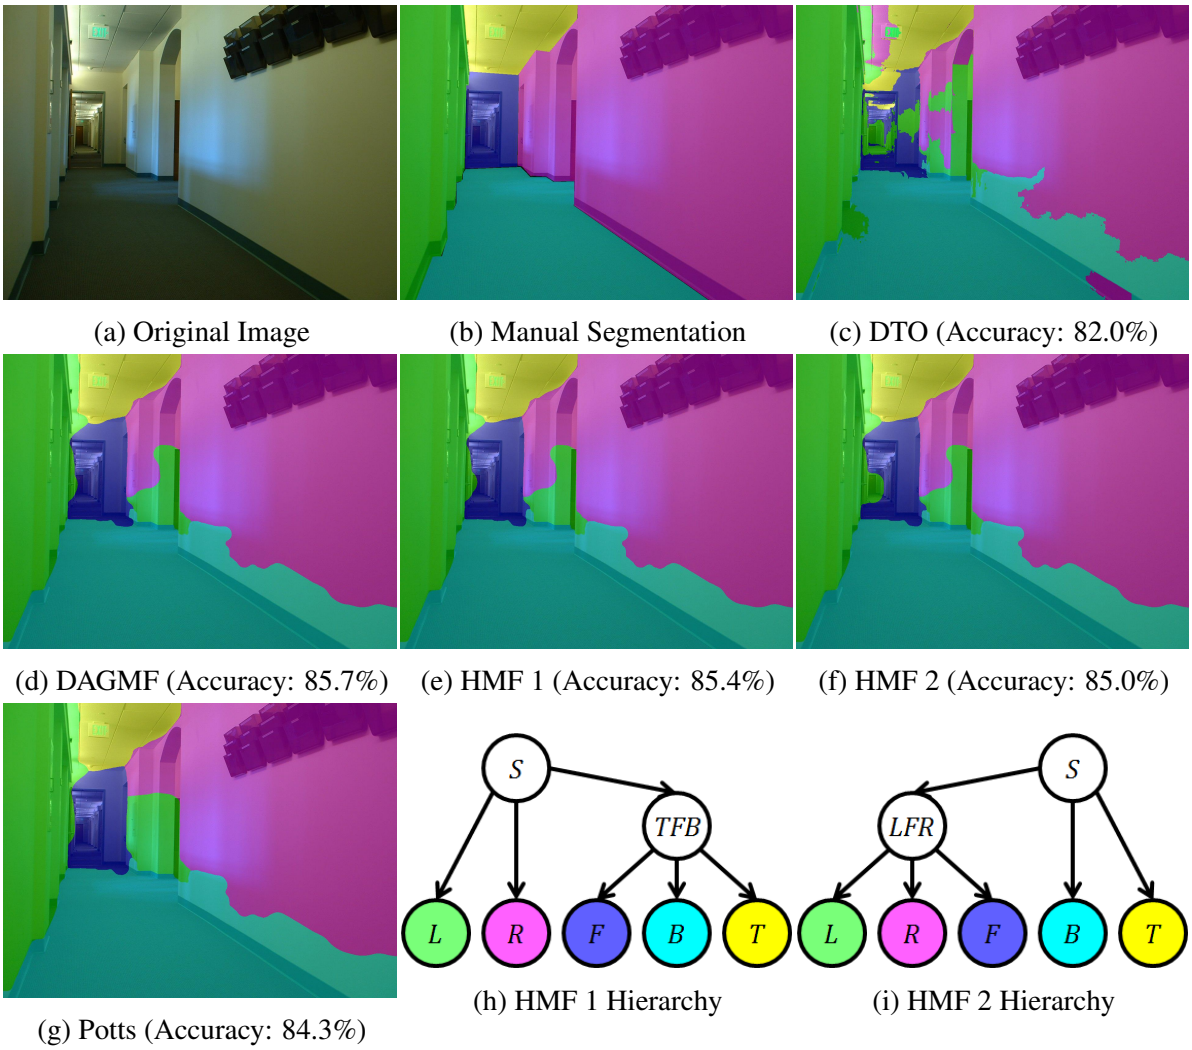

(a) Original Image (b) Manual Segmentation (c) DTO (Accuracy: 82.0%)

(d) DAGMF (Accuracy: 85.7%) (e) HMF 1 (Accuracy: 85.4%) (f) HMF 2 (Accuracy: 85.0%)

(g) Potts (Accuracy: 84.3%) (h) HMF 1 Hierarchy (i) HMF 2 Hierarchy

Figure 5.12: Example scene segmentation from the Stanford indoor dataset [42]. DTO refers to the "data term only" method. Accuracy rates are given for each segmentation. Label orderings used in the first and second HMF segmentation are shown in (f) and (i) respectively. The color code for the segmented images are shown in Figure 5.10 and in (f) and (i).

	Accuracy Rate	
(n=48)	Original	Improvement over DTO
DTO	84.4% ± 6.7%	
Potts	85.6% ± 7.1%	1.2% ± 1.2%
HMF 1	85.7% ± 6.9%	1.3% ± 1.1%
HMF 2	85.7% ± 7.0%	1.3% ± 1.0%
DAGMF	**85.9% ± 7.0%**	**1.5% ± 1.2%**

Table 5.3: Accuracy rates for segmentations in the Stanford indoor dataset such as that shown in Fig. 5.12. DTO refers to the "data term only" method. The results shown in **bold** represent those statistically significantly different from the DTO method under a paired t-test with Bonferroni correction.

for all end-labels.

For quantitative validation, this scene decomposition model was applied to the Stanford indoor image dataset [42] (48 images) using the classifiers developed by Hoiem *et al.* [73] as the basis for the data terms. The regularization terms, given in Table 5.2, determined by a exhaustive search on a subset of the database (4 images). Qualitative segmentation results for DAGMF as well as the continuous Potts model [193] and HMF models are shown in Figure 5.12 with quantitative results in Table 5.3. The "data term only" (DTO) results reflect the accuracy of the classifiers developed by Hoiem *et al.* [73]. The use of simple constant regularization improves the accuracy slightly overall, with the more representative DAGMF model having slightly higher accuracy than the previous max-flow models. The results for each max-flow model could readily be improved by created more complex regularization terms with appropriate parameterization.

5.5.4 Natural + Synthetic Images: Hue Reconstruction

Another application which is not currently possible with existing continuous max-flow methods is hue denoising in color images. The problem of hue reconstruction is of particular interest as the hue displays distinctly cyclic behaviour, meaning that it is especially ill-suited for Ishikawa models [78] more traditionally used for image reconstruction. In addition, hue offers a color property that is, in theory, invariant to lighting and atmospheric conditions but is sensitive to RGB additive noise especially at low brightness and saturation.

The first step in hue reconstruction with DAGMF is the construction of an appropriate DAG. In this scenario, the DAG can follow a relatively simple two-layer formula. The bottom layer consists of N nodes representing the discrete bins the hues are grouped into. The second layer is a regularization layer with N nodes, each with $M < N/2$ edges to the lower layer. These edges are arranged in a cyclic manner, each regularization node being connected to M consecutive hues. The result is truncated linear regularization. Between two hue nodes the regularization grows linearly with the distance between them if less than M, else, the regularization is constant. Fully linear regularization can be achieved by setting $M = \lfloor N/2 \rfloor$. An example of these types of DAGs is given in Figure 5.13. Note that in this case, using the Potts model for hue reconstruction can be expressed as a limiting case as this framework with $M = 1$.

In Figure 5.14, RGB additive noise was applied to a synthetic image with 50% saturation and 25-75% brightness. Reconstructing the hue using a DAGMF hue reconstruction model ($N = 36, M = 16$) allowed for linear regularization across hues resulting in a 79% decrease in the hue error, compared to 64% using a Potts model with the same data terms and regularization

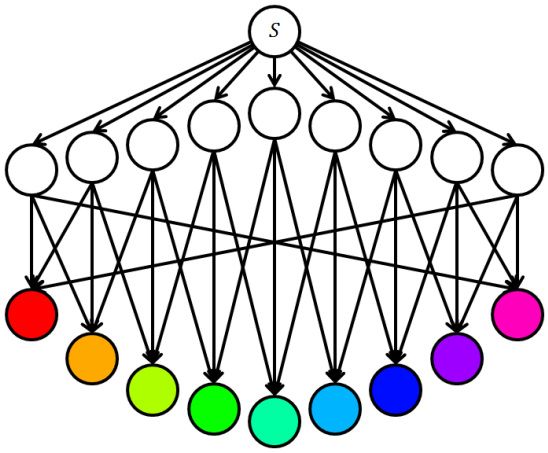

Figure 5.13: Example DAG for hue reconstruction with $N = 9$ discrete hues and a truncated linear model of width $M = 3$. Although not shown, the weight of each edge on the top level is 1, and $1/M$ on the bottom layer.

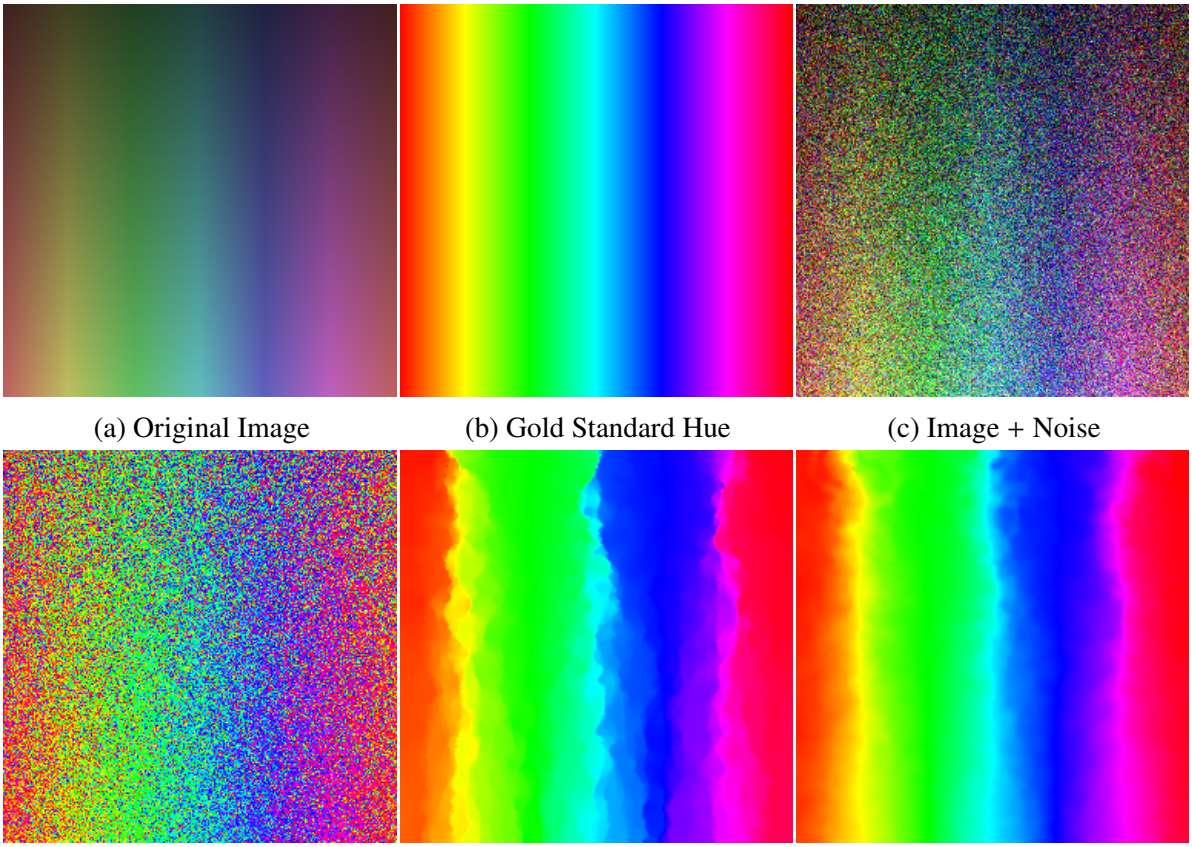

(a) Original Image (b) Gold Standard Hue (c) Image + Noise

(d) Hue of Image + Noise (Error: 32%) (e) Potts (Error: 5%) (f) DAGMF (Error: 3%)

Figure 5.14: Hue reconstruction on synthetic image with corresponding normalized hue error.

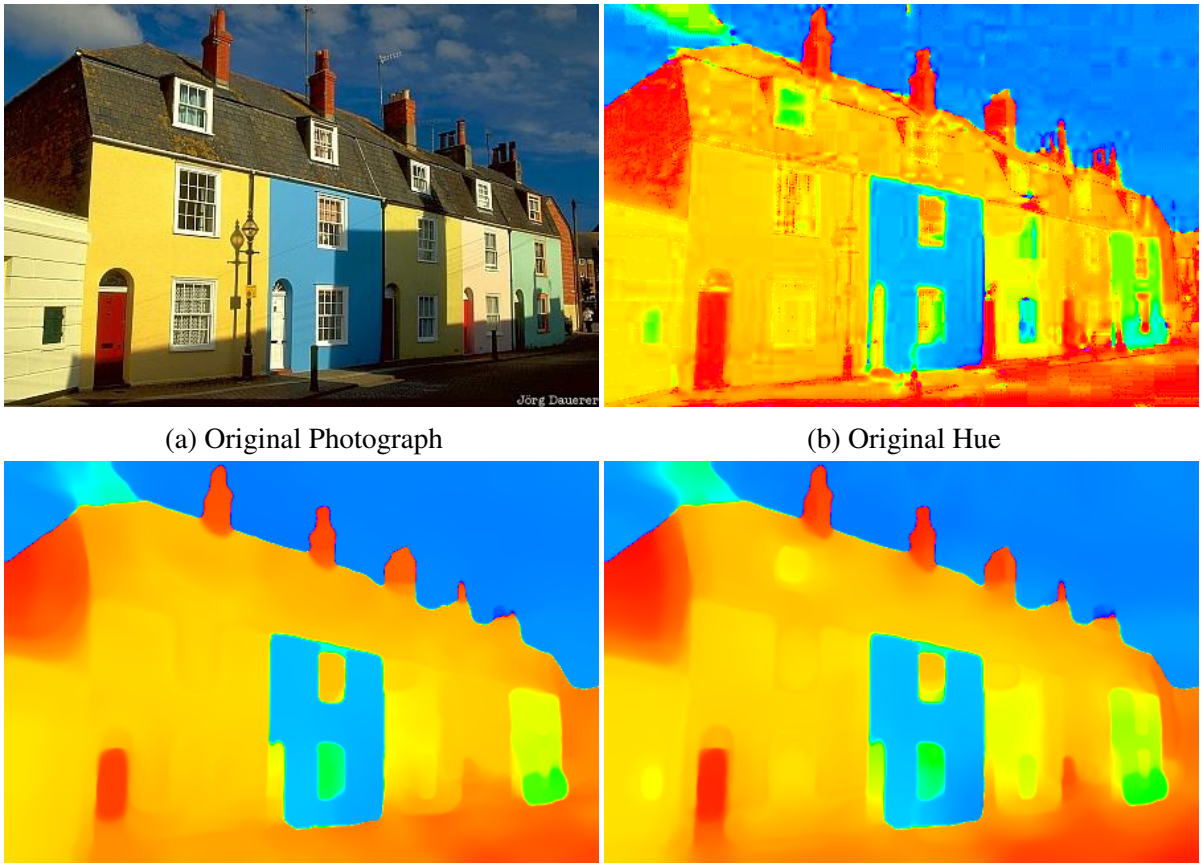

(a) Original Photograph (b) Original Hue

(c) Potts Model (d) DAGMF Model

Figure 5.15: Example hue reconstruction on natural images with DAGMF model ($N = 36$, $M = 16$).

values. The data term used was:

$$D_L(x) = |I(x) - \bar{I}_L| \tag{5.8}$$

where $I(x)$ is the RGB-valued intensity of pixel x and $\bar{I}_L$ is the RGB-valued intensity of the highest saturation colour corresponding to hue L.

The same pipeline can be used on natural images to robustly estimate hue in the presence of noise and lighting effects. Figure 5.15 displays this hue reconstruction on natural images in the presence of shadows and atmospheric perspective effects. Compared to the Potts model, the cyclic DAGMF model retains much more detail at the same level of regularization, preserving features such as smaller windows and doors.

This demonstrates that the cyclic range topology developed in Chapter 2 can also be developed via a more general framework. That being said, the framework developed using DAGMF is not as computationally efficient as cyclic continuous max-flow (CCMF) as it does not take advantage of the inherent symmetries expressed in the DAG or the more abstract regularization

structure. Thus, DAGMF offers a new tool for initially exploring image enhancement using different and possibly irregular range topologies, but other solvers need to be developed to ensure high computational performance.

5.6 Discussion

From a theoretical point of view, DAGMF illustrates that any label ordering is possible to incorporate into a continuous max-flow segmentation paradigm using primal-dual optimization with augmented Lagrangian multipliers or proximal Bregman projections. Thus, further work in continuous max-flow theory should focus on the development of more specific solvers for segmentation problem sub-classes reflecting particular topologies of interest with higher efficiency (such as in Chapter 2).

As stated in Section 5.3.1, there are multiple ways to construct a DAG that are mathematically equivalent but require different computation time and memory amounts. In addition, for certain subclasses of problems, a DAG may not be the most computationally efficient structure over which to perform max-flow. For problems in co-segmentation of multiple images simultaneously [66, 80] or segmentation based on a partially-annotated multi-atlas [87], other max-flow structures can take advantage of symmetry in the energy functional to reduce the number of, and simplify the interaction between, labellings ultimately creating more efficient algorithms. That being said, if that symmetry is removed or modified, such frameworks are no longer applicable. DAGMF provides, in such scenarios, an initial framework in which accuracy can be evaluated separate from computability concerns, decoupling the processes of model creation and solver optimization.

5.6.1 Future Work

There are a number of areas of future work to pursue in terms of segmentation problems with hierarchical label orderings. Specifically:

- Design and validate intuitive interfaces for defining DAGMF structures in the context of interactive medical image segmentation. Chapter 4 demonstrates the equivalent for HMF through the use of a common hierarchy widget often used to represent file systems. However, no such pre-designed widget exists for more general grouping structures in image segmentation.
- Automatically determine weighting parameters which limit the general use and applicability of more complex segmentation models.

- Derive procedures using DAGMF that allow for the creation of more extendable segmentation problems expressing different range topologies as in Chapter 2. This includes handling different types of symmetries in the regularization in a single coherent framework.
- Investigate the complexity of algorithms which take regularization descriptions similar to those given in Section 5.3.1 and automatically optimize the structure used for maximum computational efficiency while ensuring mathematical equivalence. As noted in Appendix B.3, this problem is NP-hard for hierarchical topologies, but has yet to be investigated rigorously for DAGMF graphs.
- Incorporate star convexity constraints [65, 178] into a subset of the labels, allowing for both topological and shape information to be optimized for simultaneously. (This is the topic of Chapter 6)

Chapter 6

Shape Complexes in Max-Flow Image Segmentation

This chapter is largely based on:

- John SH Baxter, Jiro Inoue, Maria Drangova and Terry M. Peters. "Shape Complexes: The Intersection of Label Orderings and Star Convexity Constraints in Continuous Max-Flow Medical Image Segmentation." SPIE Journal of Medical Imaging (2016).

with additional material from:

- John SH Baxter,Jing Yuan, and Terry M. Peters. 'Shape complexes in continuous max-flow hierarchical multi-labelling problems." arXiv preprint arXiv:1510.04706 (2015).

6.1 Introduction

Encoding shape information and constraining the shape of possible segmented objects has long been considered fundamental to incorporating anatomical knowledge in segmentation. The initial difficulty with this task is that there is no common definition of the word *shape* in the context of computer vision, leading to a plethora of methods addressing completely orthogonal concerns in un-related ways all under the heading of *shape information* [35].

Thus, the first step in any framework addressing shape information is to clarify the topological, geometrical and statistical information being encouraged or constrained through such a framework. As with previous chapters, the approach presented here is entirely non-statistical, in that the aim is to create descriptors of shape information that are simultaneously usable (requiring little to no training data), intuitive, and solvable. Along this line of thought, mechanisms already exist for addressing label orderings as one facet of range topology. Thus, the goal is to extend this, introducing another topological consideration coherent with the notion

of shape. To this end, star convexity, which entails the topological notion of simple connectedness, has been selected. These constraints also imply a certain level of prior knowledge in terms of location, as placement of vantage points inherently anchors the star convex object to those locations. Applying star convexity to individual labels is well-known in the literature; the novelty of this chapter is to simultaneously apply star convexity to a number of interrelated labels.

This chapter will develop a framework allowing *shape complexes*, which are *geodesic star convexity* constraints placed on a combination of multiple labels and a label ordering. This approach takes advantage of recent advances regarding multi-label topological constraints in continuous max-flow segmentation theory, and in directed acyclic graph max-flow (DAGMF) segmentation in particular. By using both star convexity constraints and label orderings in tandem, much more expressive, yet still general-purpose, shape information can be encoded. As with their predecessors, shape complexes require minimal training data, relying solely on the definition of the label ordering and the placement of vantage points.

This chapter begins with an overview of shape information in segmentation, narrowing down to the prior work in incorporating shape information into max-flow/min-cut based segmentation paradigms. This is followed by a description of *shape complexes* and the continuous max-flow algorithms used to address them. Lastly, these constructs are validated on a wide range of medical image segmentation problems in which there is complex anatomy (either in the object-of-interest or in adequately modelling the background objects) and anatomy with *walls*, such as vessels and cardiac structures.

6.2 Prior Approaches to Shape Information in Segmentation

There exist several paradigms in medical image segmentation for the incorporation of shape information. Active shape models [38, 39] and the general family of statistical shape models [69] use shape information to constrain or guide the evolution of segmentation contours to adhere to a pre-defined point distribution model for either boundary or skeleton points. Multi-level statistical shape models [128] embed hierarchical label orderings as a method of simplifying or sparsifying this information for multi-compartment objects or multiple objects in a single scene. Each of these approaches requires an intermediate representation both of the shape information and of the segmentation. That is, the segmentation being propagated is a polyhedron, which can make certain topological considerations difficult to implement.

Level-set based shape methods do not require an intermediate polyhedral representation of the segmentation, but constrain shape information directly on the labelling function [173]. These methods use energy functional imposed on the labelling function directly. The shapes

are encoded as a mean level set function and several variations. The mean level set is computed by averaging each of the signed distance functions for images in the training dataset, and the variations are the eigenmodes of the signed distance functions with this mean shape function subtracted [173]. However, encoding shape information in this manner has several drawbacks including the need for larger training datasets [40].

Alternatively, one does not have to use a sparse shape descriptor to take advantage of shape information from training data. In atlas-registration based segmentation approaches, the shape information is implicitly encoded in the atlas and constraints on the registration algorithm [9]. Atlas combination methods such as shape-based averaging [152] attempt to preserve shape information when multiple atlases are used for segmentation.

An important disadvantage common to all of these approaches is that the shape information, such as the pre-defined point distribution model or the shape atlas, is composed of a training set of segmented images. This training set must contain the necessary variability to capture pathology but limit excessive variability that would degrade performance. Shape information augmented with a more abstract anatomical model, such as hierarchical decomposed statistical shape models [189], can address some of the issues with smaller training sets by decoupling the variability across distinct anatomical regions, but ultimately cannot model the range of pathological variability seen in clinic in general.

6.3 Prior Work on Shape Information in Max-Flow

6.3.1 Discrete Domain

The development of efficient graph-cut based solvers [26] that guarantee optimality for MRFs with the property of submodularity [24] have made graph-cut techniques sufficiently flexible to encode shape information independent of training data. Of particular note are star-shape constraints or *simple star convexity constraints* [178] in which every point in an object must

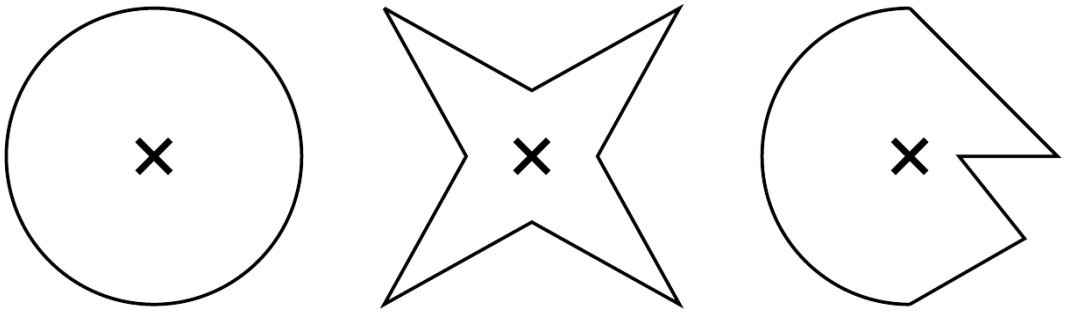

Figure 6.1: Simple star convex objects with vantage points indicated with an '**X**'.

be connected to a single vantage point using a linear path. (Examples of simple star convex objects are shown in Figure 6.1.) These are implemented in practice by the inclusion of infinite cost edges to local voxels which are sufficiently close to the line segment connecting each voxel to the vantage point.

Star convexity constraints can be extended to *geodesic star convexity constraints* [65], in which the path no longer has to be strictly linear, but follow a predefined geodesic. Similar to simple star convexity constraints, these too are implemented through the addition of infinite cost edges into the graph, but in a way that adheres to pre-defined geodesic path directions, rather than the straight line segment used by Veksler *et al.* [178]

A fundamentally different approach was taken by Egger *et al.* which perform a discrete analog of a co-ordinate system transformation rather than introduce infinite cost edges. These can be used to ensure a particular pre-defined 2D [49, 48] or 3D [47, 161] shape or extrapolate a shape from user interaction [50]. As these methods use minimal or no *a priori* training information, often relying solely on the identification of the vantage point, they are better suited for problems where a sufficient body of training segmentations cannot be collected but more geometric information about the shape is known *a priori*. In addition, they retain the guarantee of global optimality for foreground-background segmentation problems.

Statistical shape information derived from training images has been incorporated into graph-cuts [59, 117] using an iterative approach similar to that of its level-set predecessors [173]. These methods often guarantee local optimality with respect to a particular functional, but no longer have the global optimality guarantee from traditional graph-cuts.

6.3.2 Continuous Domain

Similarly, continuous max-flow segmentation has gained traction with the development of highly parallelizable duality-based optimization approaches [192]. Simple and geodesic star convexities in this continuous space have also been investigated. Yuan *et al.* [196] used an additional multiplier to allow for unconstrained flow along a predefined geodesic, allowing a single star convex label to represent the prostate. In this model, an additional set of multipliers

is used to permit flow in a predefined direction, yielding:

$$\begin{aligned} \min_{p_S, p_T, q, \lambda} & \int_\Omega p_S(x)dx \\ \text{s.t. } & |q'(x)| \leq R(x), \\ & \lambda(x) \geq 0, \\ & p_S(x) \leq D_S(x),\ p_T(x) \leq D_T(x), \\ & \operatorname{div}\left(q'(x) + \lambda(x)e(x)\right) - p_S(x) + p_T(x) = 0 \end{aligned} \tag{6.1}$$

where $D_S(x)$ and $D_T(x)$ are foreground and background data terms respectively, $R(x)$ is the isotropic regularization term, e is the local direction of the geodesic path, and q, p_S, and p_T are all flow variables. In this formulation, the foreground background labellings, u and $1-u$ respectively, are derived from the multipliers on the constraints $p_T(x) \leq D_T(x)$ and $p_S(x) \leq D_S(x)$ respectively. In the framework proposed by Yuan *et al.*, λ, the amount of flow along the predefined direction, is explicitly stored and optimized over. This explicit representation can be problematic in that it requires additional memory to store, and that, if not implemented in a similar approach as the spatial flow variables, determining the divergence of the field $q'(x) + \lambda(x)e(x)$ may be difficult. Ultimately, these limitations necessitate a framework in which the variable λ is represented and optimized over implicitly, rather than explicitly. Strekalovskiy *et al.* [168] have also proposed a form of directional regularization which can encode shape information between two interacting regions in a partitioning problem. Unlike the framework proposed by Yuan *et al.* [196], it does not explicitly represent shape constraints via an additional set of variables, but encodes them into the regularization constraints.

Ukwatta *et al.* [175] developed a max-flow propagated level-sets framework using a co-ordinate system warping approach to ensure star convexity for both the blood vessel as a whole and the lumen. (This co-ordinate system warping can be seen as a continuous analogue to Egger *et al.*'s [49, 50] discrete graph template approach.) Ukwatta's approach is conceptually important, in that it used a combination of geodesic star convex labellings to ensure a more complex ring-shaped vessel wall. However, using co-ordinate system warping to reformulate shape constraints into topological equivalents introduces a series of problems including computational expense and possible ill-definition and ambiguity when performed on multiple or branching vessels.

6.4 Shape Complexes

Shape complexes are simply segmentation structures defined as set theoretic operations on a series of labels in which some are subject to a shape constraint such as geodesic star convexity. For example, a circular ring can be described as one part of a circular super-label in which the other part is also circular. (That is, a ring is a circular object with another circular object removed from it.) Directed acyclic graph max-flow (DAGMF) from the previous chapter affords a representation of super-labels which, augmented with shape information (i.e. star convexity constraints), allows more complicated shapes to be represented.

In their current instantiation, shape complexes do not capture all possible shape information. For example, matching a segmentation to a particular *a priori* defined shape template, such as those presented by Egger *et al.* [47, 48, 49, 161, 50].

6.5 Shape Complexes Implementation

The shape complexes framework is derived from the geodesic star convexity constraint on a single label proposed by Yuan *et al.* [196]. However, a large number of these constraints may be simultaneously put in place in any given segmentation problem, and thus, memory consumption becomes a large concern. The additional optimization and storage of the λ variables is not required because $\lambda(x)$, assuming it is non-negative, can be determined exactly as:

$$\lambda(x)e(x) = \text{Proj}_{e(x)}\left(q'(x) + \lambda(x)e(x)\right) \text{ where } q'(x) \cdot e(x) = 0$$

given that all other variables are fixed. By defining a grouped spatial flow term, $q(x) = q'(x) + \lambda(x)e(x)$, (and extending the regularization term to the more general anisotropic L2 norm) the optimization problem expressed in Eq. 6.1 is equivalent to the more computationally efficient formula:

$$\begin{aligned} &\min_{p_S, p_T, q} \int_\Omega p_S(x)dx \\ &\text{s.t. } \left|\left(\mathbf{R}(\mathbf{x})^{-1}\right)^\top \left(q(x) - \lambda(x)e(x)\right)\right| \leq 1\,, \\ &\quad \lambda(x) = \max\left\{ {\left(\mathbf{R}(\mathbf{x})^{-1}\right)^\top q(x) \cdot \left(\mathbf{R}(\mathbf{x})^{-1}\right)^\top e(x)} \big/ {\left|\left(\mathbf{R}_\mathbf{L}(\mathbf{x})^{-1}\right)^\top e(x)\right|^2}, 0 \right\}\,, \\ &\quad p_S(x) \leq D_S(x),\ p_T(x) \leq D_T(x), \\ &\quad \text{div}\, q(x) - p_S(x) + p_T(x) = 0 \end{aligned} \tag{6.2}$$

assuming that the vector field $e(x)$ is normalized to unit length if non-zero, which can be achieved through an initialization step. This memory saving is crucial as multiple star con-

vexity constraints are likely to be present even in a simple shape complex.

The single-label formulation can be generalized to a framework in which labels are organized in a rooted, weighted directed acyclic graph as described in Chapter 5, necessitating a novel optimization framework extending DAGMF. This framework optimizes the primal-dual equation:

$$\begin{aligned}
&\min_{u}\max_{p,q} \int_{\Omega} p_S(x)dx + \sum_{L\neq S}\int_{\Omega} u_L(x)G_L(x)dx \\
&\quad \text{where } G_L(x) = \operatorname{div} q_L(x) + p_L(x) - \sum_{L'\in L.P} w_{(L',L)}p_{L'}(x) \\
&\qquad \text{s.t. } \left|\left(\mathbf{R_L}(\mathbf{x})^{-1}\right)^{\top}(q_L(x) - \lambda_L(x)e_L(x))\right| \leq 1, \\
&\qquad\quad \forall L \neq S,\ \lambda_L(x) = \max\left\{\left(\mathbf{R_L}(\mathbf{x})^{-1}\right)^{\top} q_L(x)\cdot\left(\mathbf{R_L}(\mathbf{x})^{-1}\right)^{\top} e_L(x) \Big/ |\left(\mathbf{R_L}(\mathbf{x})^{-1}\right)^{\top} e_L(x)|^2, 0\right\}, \\
&\qquad\quad \forall L \in \mathbb{L},\ p_L(x) \leq D_L(x)
\end{aligned} \tag{6.3}$$

This optimization can be addressed using augmented Lagrangian multipliers as shown in Algorithm 6.1 using the subroutines presented in Algorithms 6.2 and 6.3. It can also be addressed using a proximal Bregman projection based algorithm as shown in Algorithm 6.4. Both algorithms are trivially parallelizable, making them suitable for acceleration using general purpose programing on graphics processing units (GPGPU). More detailed technical information and a proof-of-correctness for these algorithms are provided in Appendix A.

6.6 Validation

In order to validate the shape complex framework, while maintaining a general focus, several distinct segmentation experiments are employed, including:

1. Synthetic images created to validate the basic properties of the shape complexes frame-

Algorithm 6.1: Augmented Lagrangian solution algorithm for Eq. (6.3).

```
Topological sort the DAG into ordering 𝕆 (begins with source label S) with reverse
ordering 𝕆^{-1} (ends with source label S);
InitializeSolution() ;
while not converged do
    UpdateFlows() ;
    for ∀L do
        ∀x, u_L(x) ← u_L(x) − c (div q_L(x) − ζ_L(x) + p_L(x)) ;
    end
end
```

work in comparison to the corresponding DAGMF models. These experiments demonstrate several important features of the algorithm such as improved accuracy and regularization parameter robustness in a controlled setting. The first of these synthetic experiments is designed to mimic the appearance of vessels in ultrasound;

2. Ultrasound images of the carotid artery were collected to verify that the behavior seen in the synthetic images is reproducible in a medical context. Thus, a similar accuracy and robustness evaluation is performed;

3. Synthetic images mimicking the mitral valve and corresponding trans-esophageal ultrasound images showing more complicated shape complexes; and

4. Cardiac CT images were collected to test an extreme-case of the algorithm with the presence of a very highly heterogeneous background with star convexity constraints applied to a very thin object-of-interest, specifically the left atrial wall.

6.6.1 Synthetic Image Segmentation

To demonstrate this approach, synthetic volumes where constructed, consisting of a medium intensity background with an embedded structure and white Gaussian noise. This structure has

Algorithm 6.2: InitializeSolution() subroutine in Algorithm 6.1.

InitializeSolution()
Clear $u_L(x), q_L(x)$ for all labels;
for *each L in order* $\mathbb{O}^{-1}$ **do**
  $\forall x, e_L(x) \leftarrow e_L(x) / \left| \left(\mathbf{R}_L^{-1}(x) \right)^\top e_L(x) \right|^2$;
  $\forall x, p_L(x) \leftarrow \min_{L'.C=\emptyset} D_{L'}(x)$;
  $\forall x, \zeta_L(x) \leftarrow \min_{L'.C=\emptyset} D_{L'}(x)$;
  if $L.C = \emptyset$ **then**
    if $L \in \operatorname{argmin}_{L'.C=\emptyset} D_{L'}(x)$ **then**
      $\forall x, u_L(x) \leftarrow 1/|\operatorname{argmin}_{L'.C=\emptyset} D_{L'}(x)|$;
    else
      $\forall x, u_L(x) \leftarrow 0$;
    end
  end
  for *each* $L' \in L.P/\{S\}$ **do**
    $\forall x, u_{L'}(x) \leftarrow u_{L'}(x) + w_{(L',L)} u_L(x)$;
  end
end

Algorithm 6.3: UpdateFlows() subroutine in Algorithm 6.1.

UpdateFlows()
for $\forall L \neq S$ **do**
  $\forall x, q_L(x) \leftarrow q_L + \tau\nabla\left(\operatorname{div} q_L(x) + p_L(x) - p_{L.P}(x) - u_L(x)/c\right)$;
  $\forall x, \lambda(x) \leftarrow \max\left\{0, \left(\mathbf{R}_L^{-1}(x)\right)^\top q_L(x) \cdot e_L(x)\right\}$;
  $\forall x, q_L(x) \leftarrow q_L - \lambda(x)e_L(x)$;
  $\forall x, q_L(x) \leftarrow \operatorname{Proj}_{|\left(\mathbf{R}_L^{-1}(x)\right)^\top q_L(x)|_{\tilde{p}} \leq 1}(q_L) + \lambda(x)e_L(x)$;
end
Clear $\zeta_L(x)$ for all labels ;
for *each L in order* $\mathbb{O}$ **do**
  for *each* $L' \in L.C$ **do**
    $\forall x, \zeta_{L'}(x) \leftarrow \zeta_{L'}(x) + w_{(L,L')}p_L(x)$;
  end
  if $L.C \neq \emptyset$ *and* $L.P \neq \emptyset$ **then**
    $\forall x, \sigma_L(x) \leftarrow \zeta(x) - \operatorname{div} q_L(x) + u_L(x)/c$;
  else if $L = S$ **then**
    $\forall x, \sigma_S(x) \leftarrow 1/c$;
  end
end
for *each L in order* $\mathbb{O}^{-1}$ **do**
  if $L.C = \emptyset$ **then**
    $\forall x, p_L(x) \leftarrow \min\{D_L(x), \zeta_L(x) - \operatorname{div} q_L(x) + u_L(x)/c\}$;
    for $L' \in L.P$ **do**
      $\forall x, \sigma_{L'}(x) \leftarrow \sigma_{L'}(x) + w_{(L',L)}\left(\operatorname{div} q_{L'}(x) + p_{L'}(x) - \zeta_{L'}(x) + w_{(L',L)}p_L(x)\right)$;
    end
  else if $L = S$ **then**
    $\forall x, p_S(x) \leftarrow \frac{1}{\sum_{L' \in S.C} w^2_{(S,L')}}\sigma_S(x)$;
  else
    $\forall x, p_L(x) \leftarrow \frac{1}{1+\sum_{L' \in L.C} w^2_{(L,L')}}\sigma_L(x)$;
    for $L' \in L.P$ **do**
      $\forall x, \sigma_{L'}(x) \leftarrow \sigma_{L'}(x) + w_{(L',L)}\left(\operatorname{div} q_{L'}(x) + p_{L'}(x) - \zeta_{L'}(x) + w_{(L',L)}p_L(x)\right)$;
    end
  end
end

Algorithm 6.4: Proximal Bregman algorithm for Eq. (6.3).

Topological sort the DAG into ordering $\mathbb{O}$ (begins with source label S) with reverse ordering $\mathbb{O}^{-1}$ (ends with source label S);
Clear $q_L(x)$ for all labels;
$\forall L \in \mathbb{L}, \forall x, u_L(x) \leftarrow 1/|\mathbb{L}|$;
$\forall L \in \mathbb{L}, \forall x, e_L(x) \leftarrow e_L(x) / \left| \left(\mathbf{R}_L^{-1}(x) \right)^\top e_L(x) \right|$;
while *not converged* **do**
  $\forall L, \forall x, d_L(x) \leftarrow \operatorname{div} q_L(x)$;
  $\forall L \in \mathbb{L}, \forall x, d_L(x) \leftarrow d_L(x) + D_L(x)$;
  for *L in order* $\mathbb{O}$ **do**
    for $L' \in L.C$ **do**
      $\forall x, d_{L'}(x) \leftarrow d_{L'}(x) + w_{(L,L')} d_L(x)$;
    end
  end
  $\forall L \in \mathbb{L}, \forall x, u_L(x) \leftarrow u_L(x) \exp\left(-\frac{d_L(x)}{c}\right)$;
  $\forall L \in \mathbb{L}, \forall x, d_L(x) \leftarrow u_L(x)$;
  $\forall x, a(x) \leftarrow \sum_{L \in \mathbb{L}} u_L(x)$;
  $\forall L \in \mathbb{L}, \forall x, u_L(x) \leftarrow u_L(x)/a(x)$;
  $\forall L \notin \mathbb{L}, \forall x, d_L(x) \leftarrow 0$;
  for *L in order* $\mathbb{O}^{-1}$ **do**
    $\forall x, q_L(x) \leftarrow (q_L(x) - c\tau \nabla d_L(x))$;
    $\forall x, \lambda_L(x) \leftarrow \max\{0, (\mathbf{R}_L(x))^\top q_L(x) \cdot e_L(x)\}$;
    $\forall x, q_L(x) \leftarrow q_L(x) - \lambda_L(x) e_L(x)\}$;
    $\forall x, q_L(x) \leftarrow \operatorname{Proj}_{|(\mathbf{R}_L^{-1}(x))^\top q_L(x)|_{\tilde{p}} \leq 1} (q_L(x)) + \lambda_L(x) e_L(x)$;
    for $L' \in L.P / \{S\}$ **do**
      $\forall x, d_{L'}(x) \leftarrow d_{L'}(x) + w_{(L,L')} d_L(x)$;
    end
  end
end

a slightly hypo-intense centre surrounded by a hyper-intense boundary. The contrast-to-noise ratio (CNR) between the background and the centre is $\approx 10\%$ and between the hyper-intense boundary and each other region is $\approx 100\%$. Due to the low CNR, segmentation of these images with minimal prior information can be challenging. These images were segmented with a simple DAGMF model and using shape complexes as shown in Figure 6.2, with a manually picked centroid or *vantage point* to define a simple star convexity constraint. The same intensity-based data terms and uniform regularization were used in both images.

This experiment was repeated for a range of regularization values between 10^{-1} and 10^{1}, and the Dice Similarity Coefficient (DSC) measured for each of the tree labels as shown in Figure 6.3. As expected, the segmentation with shape complexes consistently outperforms the

segmentation without shape complexes for all regularization values, and displays an additional degree of robustness to the regularization value chosen. This is especially important as star complexes allow for much lower regularization values to be used without sacrificing segmentation quality, which is desirable in terms of preserving less smooth portions of an object.

A second synthetic experiment, shown in Figure 6.4, was performed to illustrate the use of shape-complexes in a segmentation problem with a distinctly non-hierarchical label ordering. The *CNR* between regions in Figure 6.4b is 25%. As shown in Figure 6.4e, the regularization parameter was too low to enforce region contiguity under such high noise, which is readily addressed through the addition of a series of star-convexity constraints shown in Figure 6.4f. Again, keeping the regularization value low is essential to capturing more tortuous boundaries of a segmented object without over-smoothing. Efficiently determining the regularization parameter in these segmentation models which neither over-smooths or under-smooths is an area of open research [19].

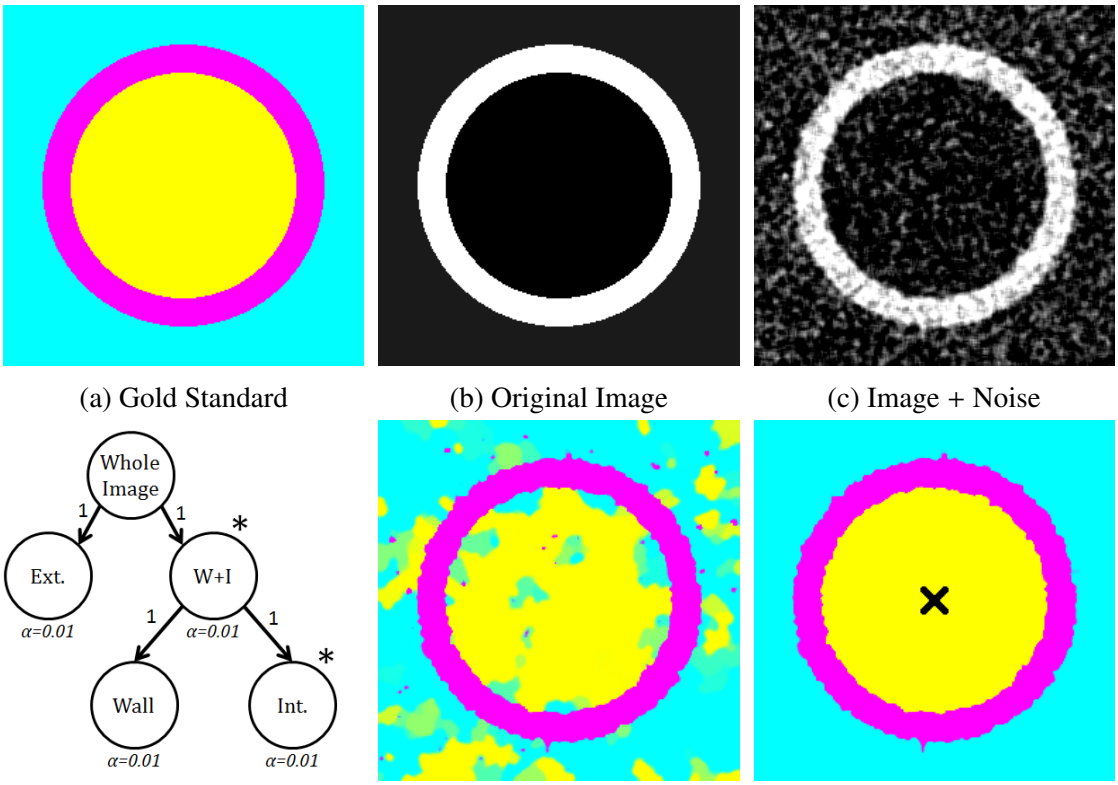

(a) Gold Standard (b) Original Image (c) Image + Noise

(d) Label Ordering (e) DAGMF Segmentation (f) DAGMF + Shape

Figure 6.2: Synthetic image segmentation problem using DAGMF (2e) and DAGMF augmented with shape complexes (2f) according to the label ordering in (2d) with α referring to the level of regularization. (* a simple star convexity constraint is applied to this label.) Any overlap between segmentations can cause false colors, e.g. green occurs when the result is 50% exterior (cyan) and 50% interior (yellow). The 'X' marks the vantage point for the simple star convexity constraint.

6.6.2 Ultrasound Vessel Segmentation

To demonstrate the applicability of these shape complexes to medical image segmentation, they were applied to the segmentation of an individual vessel of interest, in particular, the carotid artery, from an ultrasound image. This image was manually segmented into three regions, the background, vessel lumen, and vessel blood pool. Similar to the first synthetic experiment in Section 6.6.1, the shape complex applied creates a ring-shape prior on the vessel wall. However, this segmentation model takes into account a multi-component background, with both vessel wall and blood pool components, handling background heterogeneity. The overall model consists of labels for the vessel blood pool B_V, vessel wall W_V, background blood B_K, background hyper-intense tissue such as other vascular walls W_K, and other background tissue K. Segmentation results are shown in Figure 6.5.

The data terms used are derived from Bayes' theorem on the voxel intensity:

$$D_L(x) = -\ln P(I(x)|x \in L) \pm bias \tag{6.4}$$

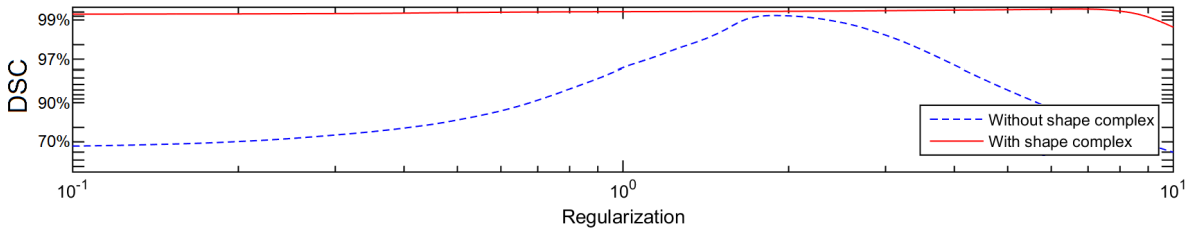

(a) External (cyan) region

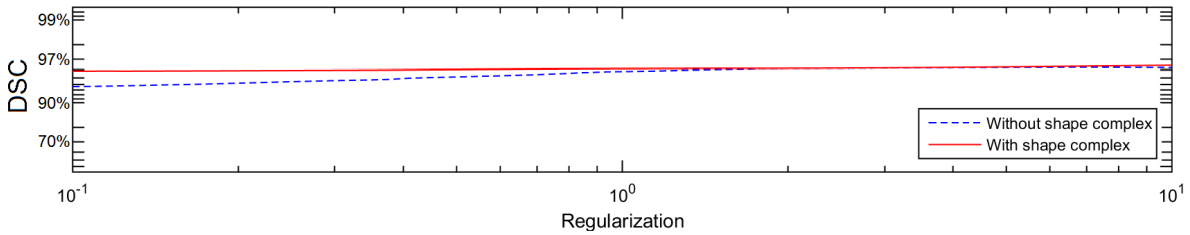

(b) Ring (magenta) region

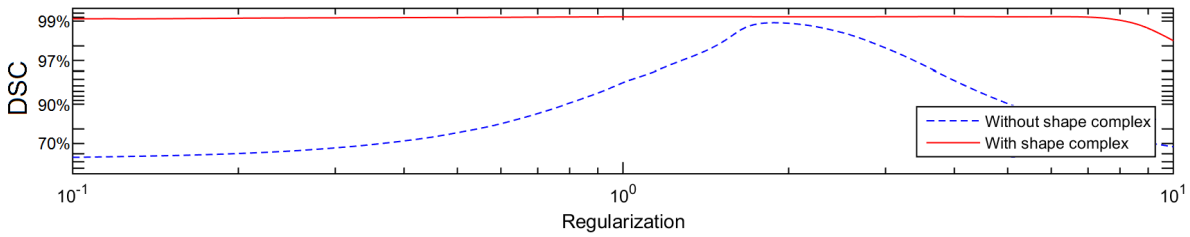

(c) Internal (yellow) region

Figure 6.3: Quantitative segmentation results for each region based on regularization strength. The Dice similarity coefficients are shown on a logarithmic scale approaching 100% DSC.

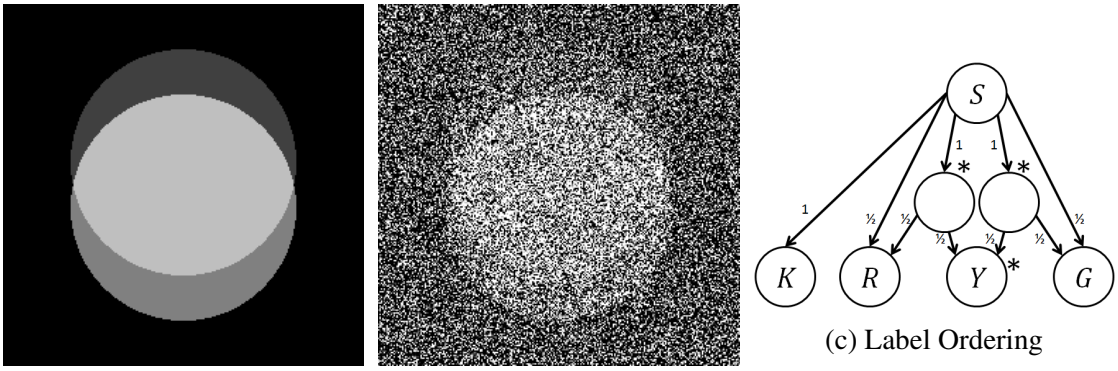

(a) Original Image (b) Image + Noise (c) Label Ordering

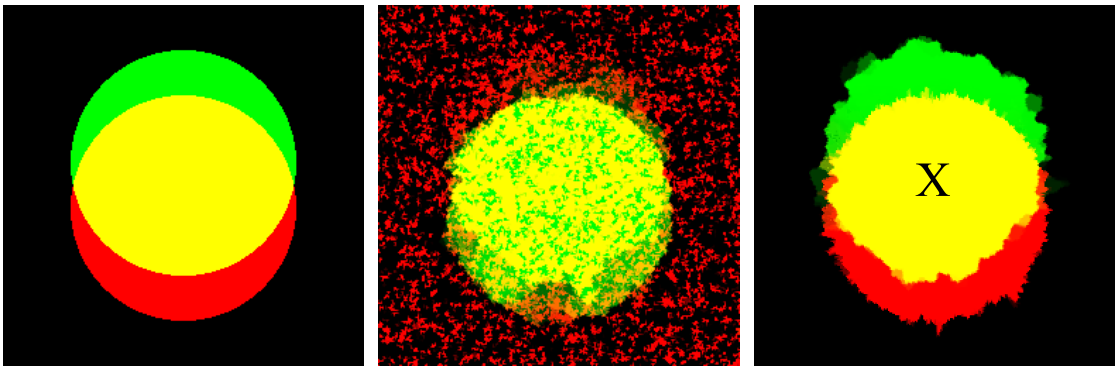

(d) Gold Standard (e) DAGMF Segmentation (f) DAGMF + Shape

Figure 6.4: Venn diagram segmentation with and without shape complexes. The label ordering is given in Figure 4c (* a simple star convexity constraint is applied to this label) with the vantage point for the shape complex was the centroid of the region. Similar to Figure 6.2, any overlap between segmentations can cause false colors.

where $P(I(x)|x \in L)$ is the probability of voxel x having intensity $I(x)$ given that it is a part of label L. The constant *bias* term (positive for background components and negative for foreground components) controls for the shrinking bias which is especially severe as the background super-label contains two components with the same intensity distribution as the vessel of interest, i.e. labels B_K and W_K with the same intensity as B_V and W_V respectively. The user provided seeds for the blood, vascular wall, and background tissue components are given in Figure 6.5b. The regularization is a constant applied to all labels and super-labels shown in Figure 6.5c with the exception of the vessel wall label, V, which has zero regularization to avoid shrinking bias.

An experiment similar to that shown in Figure 6.3 was performed on the vessel ultrasound dataset, varying both the regularization and bias parameters. Quantitative results are shown in Figure 6.6. Not only did including the star convexity constraint improve the DSC compared to unconstrained DAGMF at their respective optimal values, but the segmentation became more robust to parameterization, maintaining similarly high DSC over a much broader range of pa-

rameters, confirming the earlier observation on synthetic images. This is especially important for medical image segmentation problems in which an exhaustive search through or optimization of the parameter space is difficult to perform, or when parameters are selected interactively [17], as the addition of shape complexes makes selection easier and less sensitive to operator variability.

6.6.3 Cardiac Valve Segmentation from Ultrasound

Even more complicated shapes can be created by combining simpler shapes. For example, Figure 6.4 involved the overlapping of two star convex objects and Figure 6.2 showed how star shape priors could be nested to form rings. A shape complex combining these two ideas as shown in Figure 6.7a can eb created that describes the intersection of two ring-shaped objects. This shape occurs when segmenting structures such as cardiac valve annuli that are a

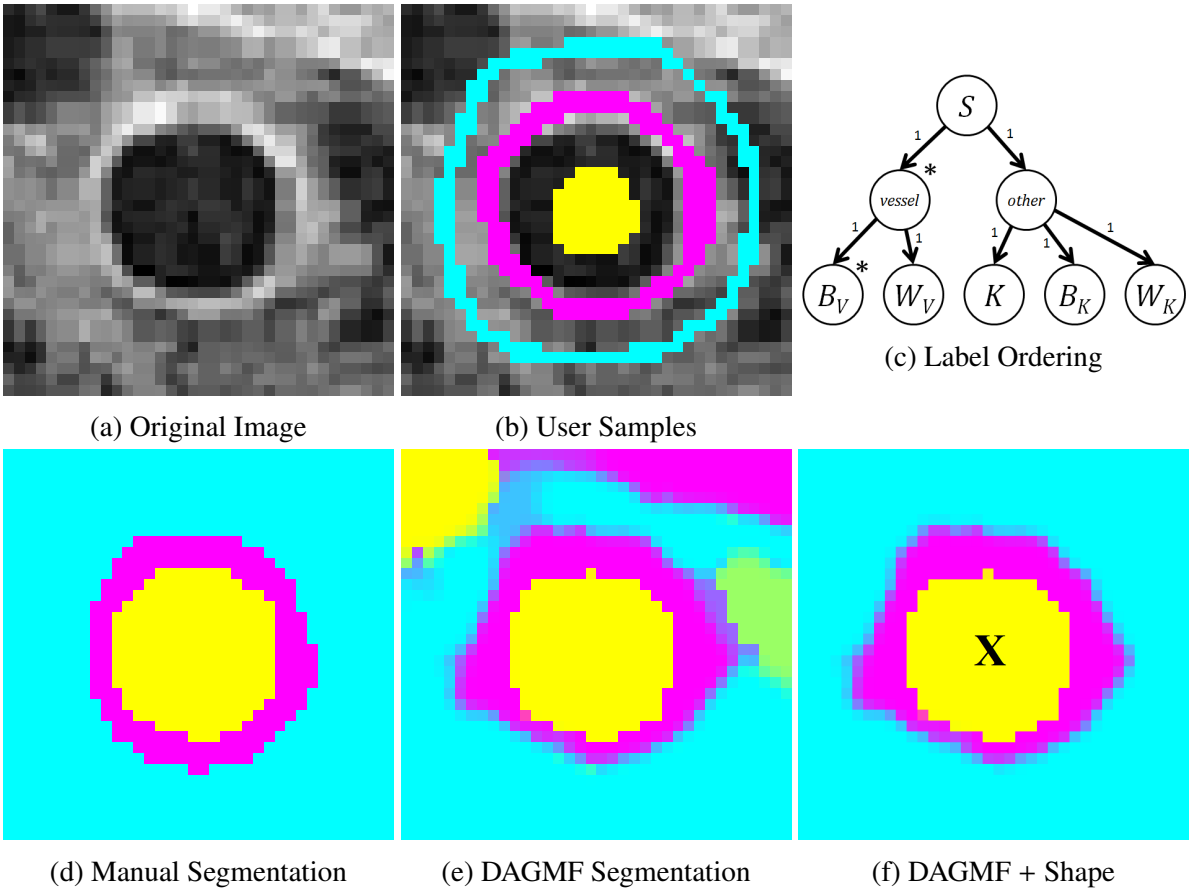

(a) Original Image (b) User Samples (c) Label Ordering

(d) Manual Segmentation (e) DAGMF Segmentation (f) DAGMF + Shape

Figure 6.5: Vessel segmentation in ultrasound with and without shape complexes. The label ordering is given in Figure 5c (* a simple star convexity constraint is applied to this label) with the vantage point for the shape complex is marked with an '**X**'. Similar to Figure 6.2, any overlap between segmentations can cause false colors.

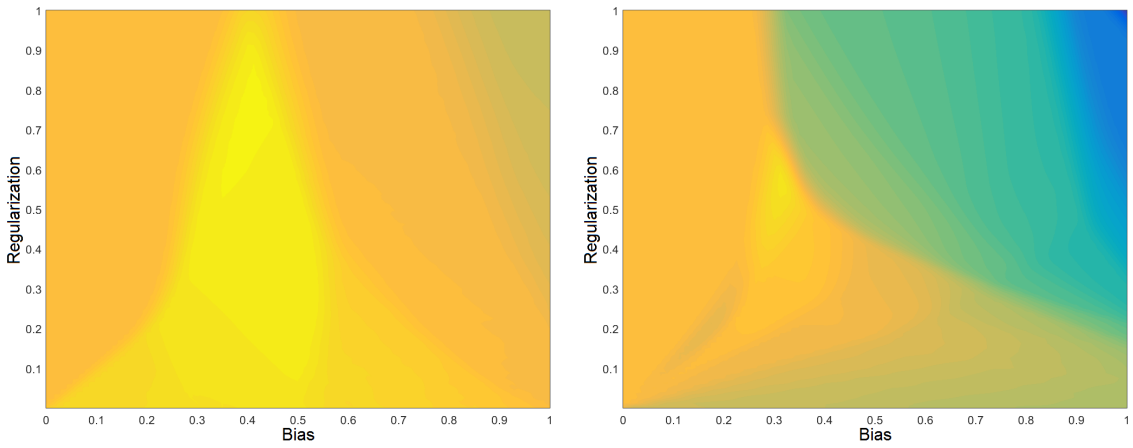

(a) DAGMF + Shape: Background

(b) DAGMF only: Background

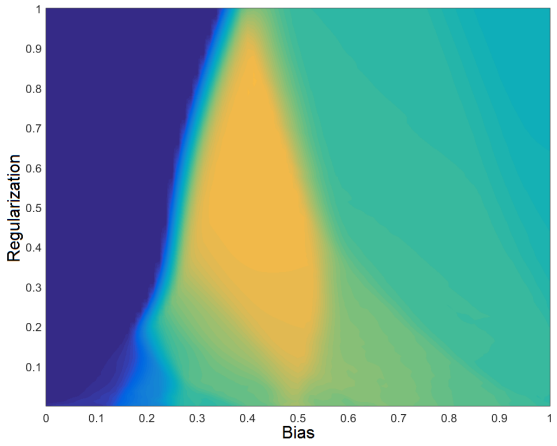

(c) DAGMF + Shape: Vessel wall

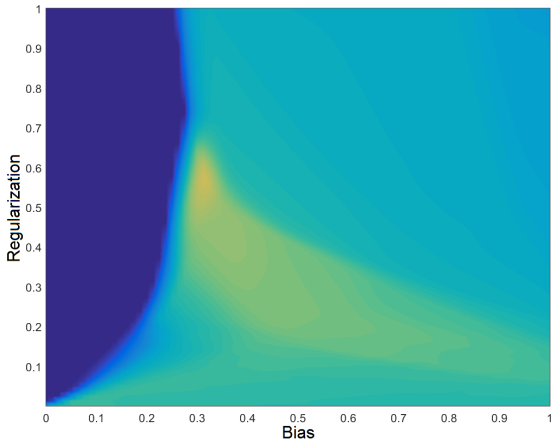

(d) DAGMF only: Vessel wall

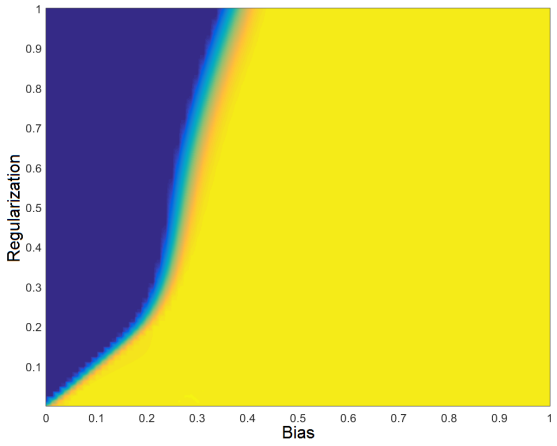

(e) DAGMF + Shape: Vessel lumen

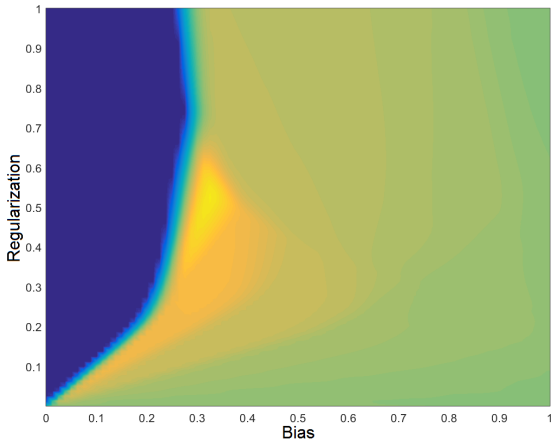

(f) DAGMF only: Vessel lumen

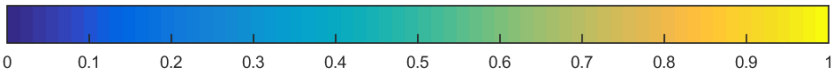

(g) DSC Color Scale

Figure 6.6: Quantitative results for the segmentation problem shown in Fig. 6.5 varying regularization and bias parameters. Blue indicates DSC ≈ 0 and yellow indicates DSC ≈ 1 as shown in Fig. 6g.

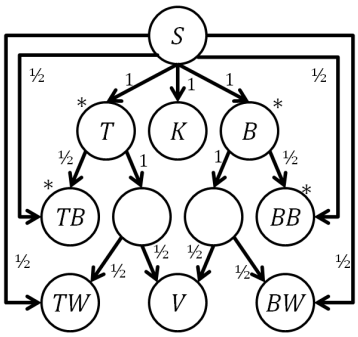

(a) DAG used for valve annulus segmentation

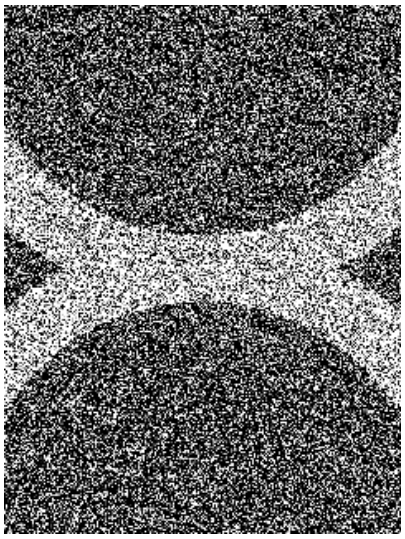

(b) Original noisy image

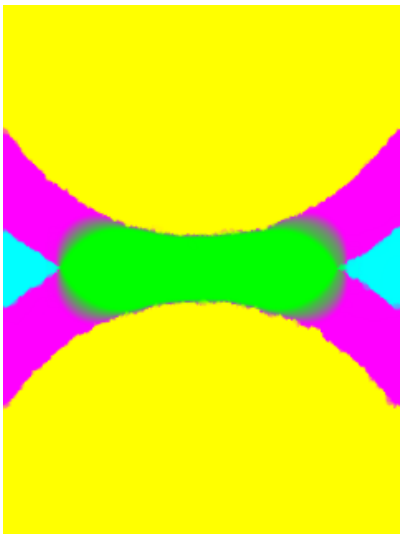

(c) Segmentation results

Figure 6.7: Synthetic valve annulus segmentation. The label K indicates the background (in cyan), TW and BW indicates the top and bottom walls respectively (in magenta), TB and BB indicate the top and bottom blood pools respectively (in yellow), and V indicates the valve annulus (in green).

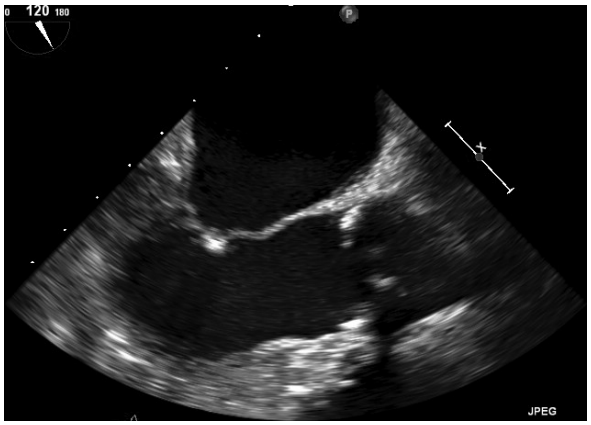

(a) Original Ultrasound Image

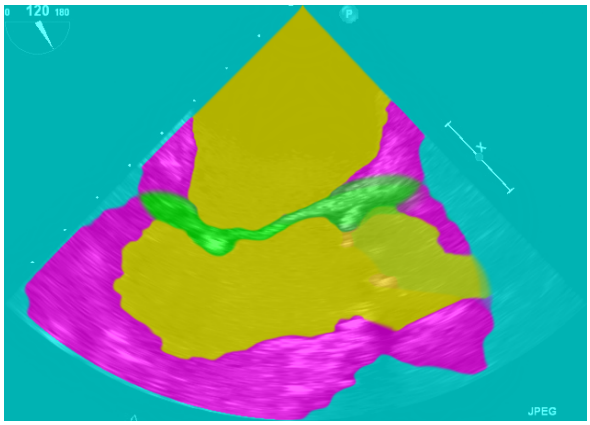

(b) Segmentation

Figure 6.8: Mitral valve labelling using trans-esophageal ultrasound images. The model (Figure 6.7a - previous figure) includes label K indicates the background (in cyan), TW and BW indicates the top and bottom walls respectively (in magenta), TB and BB indicate the top and bottom blood pools respectively (in yellow), and V indicates the valve annulus.

shared boundary between more readily segmented objects such as blood pools. This model contains six regions: K, the background region; TB and TW, representing the top blood pool and surrounding wall; BB and BW, representing the bottom; and V representing the intersection between them.

Figures 6.7b and c show a synthetic segmentation example for this shape complex. The image has $CNR = 100\%$ with zero contrast between the background and blood pools and between the walls and the valve. With no contrast between similar objects, the shape information is necessary for segmentation. The segmentation result demonstrates that the shape complex can adequately localize the synthetic valve, that is, a minimal region separating the blood pools. Without contrast, there are few defining features of the valve region defined against the walls, representing a limitation of the use of shape complexes alone.

Figure 6.8 displays a similar experiment using a trans-esophageal ultrasound image of the anatomy surrounding the mitral valve. Similar to the synthetic image example, the segmentation algorithm cannot accurately segment the valve annulus where it is adjacent to the myocardium. Additionally, the ultrasound image included a partial view of the aorta and aortic valve, which was not accounted for in the model. Because this was not included in the model, the segmentation algorithm had difficulty segmenting said region as the signal intensities contradicted the expectation of the shape term. (The hyper-intense valve did not allow the aortic blood pool to be easily grouped with the left ventricular blood pool yet its proximity and attached-ness discouraged associating it fully with the background.) Ambiguity in the segmentation result is evidenced by partial colouring.

Both experiments used uniform regularization and a relatively simple data term:

$$D_L(x) = \begin{cases} |I(x) - I_B|, & \text{if} L \in \{BB, TB, K\} \\ |I(x) - I_W|, & \text{if} L \in \{BWO, TWO\} \\ |I(x) - I_W| - bias_I, & \text{if} L \in \{I\} \end{cases} \tag{6.5}$$

where I_B and I_K were the average intensity of the blood pools and walls respectively, and $bias_I$ was a bias term to encourage the expansion of the valve annulus label and avoid partitioning the valve region between the BWO and TWO labels. Spatial seeding was provided for the background and two blood pool labels and, for the latter, the centroid of the seed locations was used as the vantage point for the star convexity constraints.

6.6.4 Atrial Wall Segmentation from Cardiac CT

This technique can be used to segment anatomy with boundary structures such as the atrial wall in contrast enhanced CT. (These images were collected as part of a larger study investigating

(n = 10)	Mean Distance Error (Operator 1)	Mean Distance Error (Operator 2)
Blood pool (inner wall)	0.76 ± 0.51 mm	0.59 ± 0.36 mm
Whole atrium (outer wall)	1.51 ± 0.55 mm	1.27 ± 0.29 mm

Table 6.1: Mean distance error results for the blood pool and whole atrium labels. These are reflective of the errors seen on the inner and outer boundary of the atrial wall label.

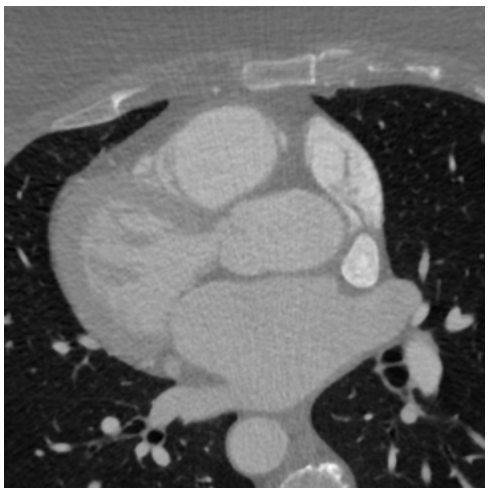

(a) Original Contrast-Enhanced CT Image (cropped to cardiac region)

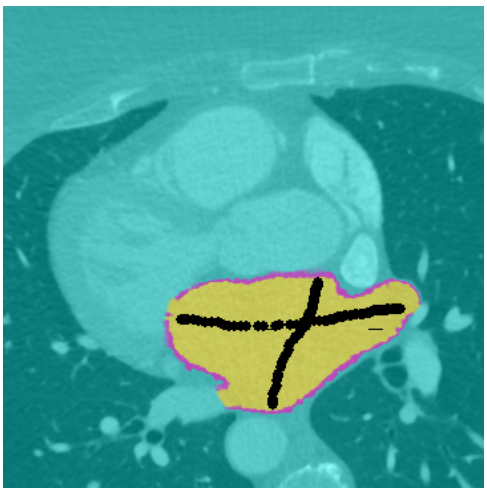

(b) Segmentation with user initialization for the atrial blood pool shown in black

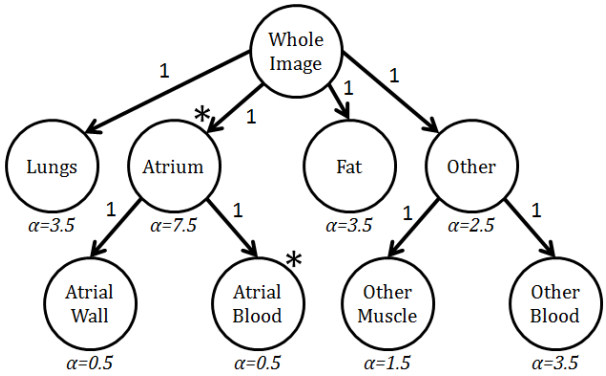

(c) Label Ordering

Figure 6.9: Atrial wall segmentation DAGMF augmented with shape complexes with α representing the regularization strength. (* a simple star convexity constraint is applied to this label.)

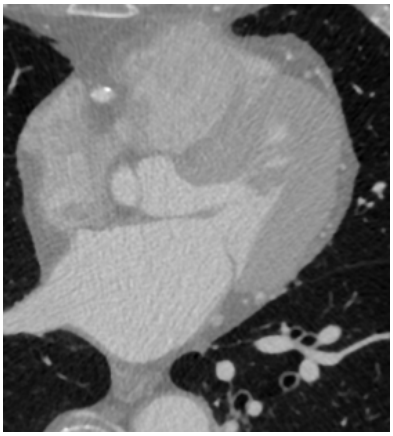

(a) Best CaseOriginal Contrast-Enhanced CT Image (cropped to cardiac region)

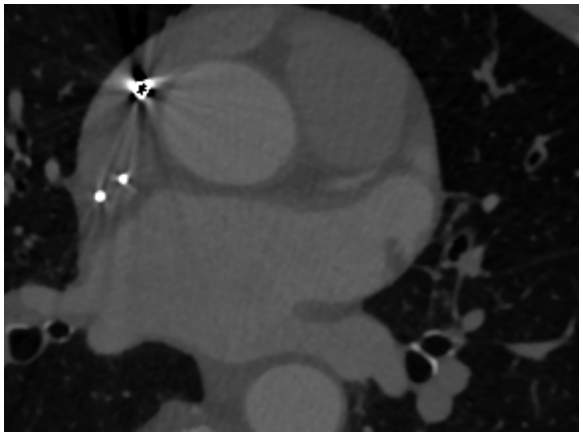

(b) Worst Case Original Contrast-Enhanced CT Image (cropped to cardiac region)

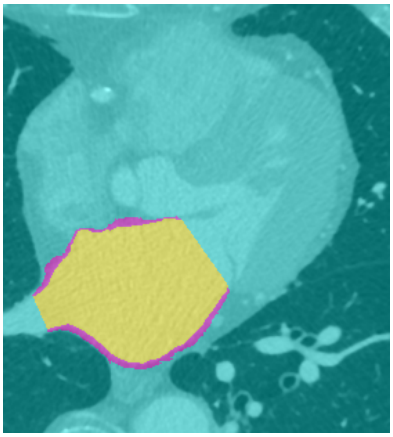

(c) Best Case Manual Segmentation

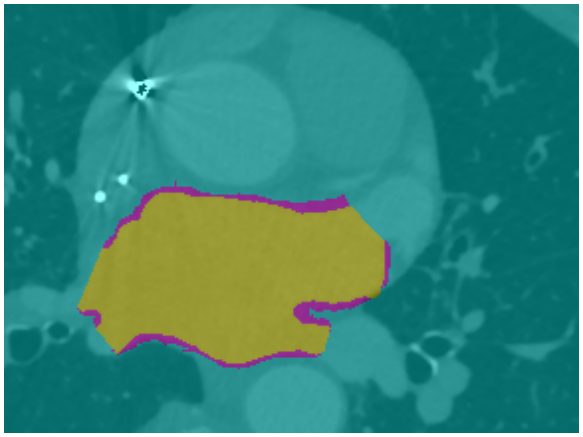

(d) Worst Case Manual Segmentation

(e) Best Case Shape Complex Result

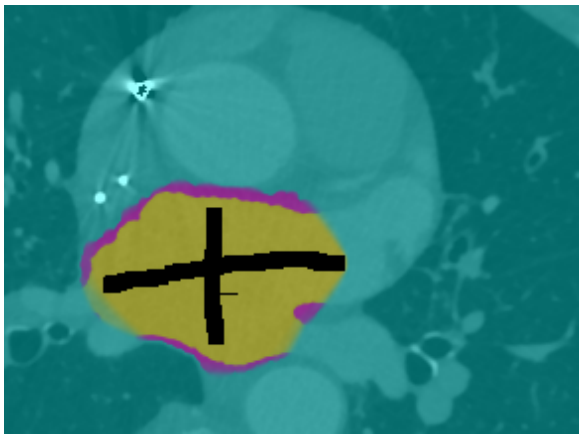

(f) Worst Case Shape Complex Result

Figure 6.10: Best and worst case atrial wall segmentation results. The atrial blood pool is shown in magenta and the atrial wall in cyan. The black regions are user-provided seed points for the atrial blood label.

the use of anatomic measurements of the heart for radio-frequency catheter ablation. The ethics approval for this data is given in Appendix C Section C.4.) A challenging aspect of this problem is that the atrial wall has little to contrast against the nearby muscle, thus requiring shape complexes to constrain it around the more distinctive atrial blood pool. The segmentation was semi-automated, whereby seeds placed by the user in the atrial blood pool, muscle, fat, and lungs were employed to fit a normal intensity distribution to each tissue type and define a geodesic star convexity prior for the atrium. A uniform smoothness term was created with a label-specific uniform regularization. The segmentation model is shown in Figure 6.9. Best and worst case results are shown in Figure 6.10a-c and d-e respectively. Quantitative results are provided in Table 6.1.

Despite having the same intensity distribution and therefore the same data terms, the segmentation successfully differentiated between the atrial and non-atrial blood pools, as well as atrial wall versus other muscular structures. Currently this segmentation protocol is semi-automatic, requiring some user initialization. In an automated protocol, the user initialization could be replaced by prior knowledge about the Hounsfield distribution of different tissue intensities with some mechanism for automatically estimating the centroid of the atrial blood pool. The smoothness terms allow the blood pools and walls to closely follow perceptible edges in the image even at high regulation values.

There is currently a non-negligible degree of user variability in the method as shown by the difference in accuracy results with respect to the manual segmentation outlined in Table 6.1. This is to be expected as the seeding not only provides a spatial anchor for the segmentation and defines the vantage points for the shape complex, but also affects the probabilistic data terms used in the optimization process.

6.7 Discussion

The addition of geodesic star convexity and related topological considerations in a general-purpose and application-agnostic manner improves the expressiveness of possible anatomical information that can be encoded in a segmentation problem. This encoded knowledge can greatly improve segmentation without requiring higher regularization which can obscure fine structures and detail on the segmented objects. The use of shape complexes to improve the robustness of the optimization algorithm to regularization parameter may be especially useful in these scenarios in which numeric parameterization may be opaque and unintuitive for clinical users. The mitigation of variations in accuracy may also be helpful in further validation and comparison with other algorithms, as less effort is required to generate acceptable performance.

Because of the emphasis on maintaining a single continuous-space, there is no additional

computational expense or ambiguity due to co-ordinate system warping, permitting the segmentation of multiple objects simultaneously and well as branching objects, both of which complicate co-ordinate system warping approaches. This is in stark contrast to prior approaches [175] in which such warping is necessary, preventing specific types of shape complexes to be solved due to co-ordinate system ambiguity. Generally speaking, these co-ordinate system difficulties are avoided by discrete approaches [65, 178], but by maintaining a continuous-space image domain, issues of metrication are avoided entirely, especially those resulting from the placement of infinite-cost edges associated with discrete domain star convexity constraints.

These priors however require some form of intelligent initialization to infer the geodesic direction field, $e_L(x)$. In both the synthetic and medical image segmentation experiments, this information was provided by the user through picking the centroid (as in Sections 6.6.2 and 6.6.1) or by seeding the region of interest (as in Section 6.6.4). Although these methods may be suitable for HMF-based interactive segmentation (as in Chapter 4), different approaches will be required for fully automated segmentation pipelines. Currently, the use of manual region-of-interest seeding is likely a large cause of user variability in complex problems such atrial wall segmentation.

6.7.1 Future Work

There are four immediate areas of future work for shape complexes:

- Incorporation of multiple star-convexity constraints into a single label, increasing the number of shape options for each label. For example, constraining a shape to be circular about a particular vantage point could be implemented with three geodesics, one pointing towards the centre of the circle and two tangent to it, pointing in the clockwise and counter-clockwise directions respectively. In theory, a similar combinations of geodesics could be used for arbitrary shapes provided that the boundary is parametrically defined or in interactive scenarios where edge information can be extracted and scale-invariance is assumed [50].

- Performance improvements including GPGPU acceleration and incorporation of C++ implementations into the ASETS library [1]. As stated in Section 6.5, shape complexes can be implemented in an inherently parallelizable manner suitable for GPGPU acceleration. Incorporation into C++ would more readily allow its integration into open source libraries for medical imaging processing and visualization, such as ITK and VTK.

- Incorporation into existing continuous max-flow based interactive segmentation interfaces that use input sampling mechanisms amenable to defining star convexity con-

straints. Alternatively, vantage point placement could be implemented as another secondary mechanism for user interaction.

- Investigation of automated approaches for defining geodesics based on atlas registration suitable for fully-automatic segmentation pipelines. By processing deformation field resulting from the registration, it may be possible to automatically derive geodesic, rather than simple, star convexity constraints.

While these initial results are promising, and the additional robustness to parameter selection widely desirable in an array of medical image segmentation tasks, more validation is required to establish clinical utility in a particular medical domain. For example, in order to be clinical applicable for atrial wall segmentation, shape complexes must be augmented with a cost estimation framework that is robust to CT artifacts, such as those generated by the presence of pacemakers, as well as variable contrast-to-noise ratios due to the variable dose rates of cardiac CT across clinical centres and scanners.

Chapter 7

Conclusions

Graph-cuts and continuous max-flow have had a profound effect on our understanding of image processing, emphasizing the solvability of simple, more-general purpose functionals over the increasing complexity of larger non-solvable functionals.

In terms of concrete contributions, the work presented in this thesis represents:

- the first instance in which continuous max-flow has been applied to MRI phase processing, as well as the first instance of a cyclic topology being used in image enhancement (Chapter 2),
- the first instances of continuous max-flow segmentation frameworks, HMF (Chapter 3) and DAGMF (Chapter 5), for very general classes of label orderings,
- the first interactive segmentation interface (SEGUE in Chapter 4) in which abstract anatomical knowledge is a locus for user interaction, and
- the first instance in which label orderings have been used to encode interacting shape constraints through shape complexes (Chapter 6), which allows it to represent a wide array of anatomical shapes in a optimizable manner.

Although these contributions are to distinctly different medical image processing problems, underlying them is the use of continuous max-flow theory and a structured approach to translating image labelling problems through image partitioning with indicator functions.

I would like to conclude this thesis with three points of discussion:

1. an investigation into the recurrent themes of the thesis,
2. a quick tour of the aspects of continuous max-flow theory in medical image processing that have not been expounded upon earlier in the thesis, and
3. some speculation over the future of continuous max-flow theory in medical imaging.

7.1 Recurrent Themes

Although this thesis is about the contributions of a particular mathematical formalism (that of continuous max-flow theory) to medical image processing, there are two deeper recurrent themes worth addressing:

- the principle of topology, and
- the principle of interactivity.

These two principles are important to medical image processing in a way that underlies any particular algorithmic or theoretical framework in which said processing is performed. Both principles are concerned with the epistemological question of ***what is knowledge in a medical image processing problem?*** This question and how it is addressed is fundamental to the underlying philosophy of medical image processing and thus how it is practised

In exploring these two principles, I hope to instil in the reader a new perspective on this thesis that complements the more down-to-earth emphasis on the medical image processing problems presented, or the more technical/theoretical emphasis on the algorithms and their development.

7.1.1 The Principle of Topology

The principle of topology answers the fundamental question of ***what is knowledge in a medical image processing problem?*** in a positive manner: ***knowledge is (at its least) the knowledge of topology***.

Topology provides a minimal working amount of anatomical knowledge to constrain a problem and it is often of the utmost importance. To provide an historical example, the movement from active contours to level-set segmentation was motivated almost solely by the latter's ability to incorporate topological change. In the opposite direction, graph-cut techniques were augmented by star convexity constraints for the purpose of constraining the topology to have simple connectivity. Interrogating and questioning the topological assumptions of a method has historically been a source of inspiration for new and improved methods.

In this thesis, four chapters were motivated technically by topological considerations:

- Chapter 2 by the lack of image enhancement frameworks incorporating a specific *cyclic range topology* found in MRI phase images,
- Chapter 3 by the lack of continuous max-flow frameworks for incorporating flexible label orderings, a purely topological notion,
- Chapter 5 by the desire for a singular framework capable of handling any label ordering as an extension of the existing HMF framework, and

- Chapter 6 concerned itself with equipping segmented objects with a particular spatial topology which includes notions of simple connectivity and basic shape constraints.

These topological considerations are representative of the most basic underlying assumptions about their corresponding medical image processing problem: assumptions about of the configuration of labelling functions both in terms of their domain and range.

7.1.2 The Principle of Interactivity

Chapter 4 was motivated by the desire to have the clinical user define segmentation label orderings and be able to modify them to suit particular clinical needs such as handling patient pathology and attempted to rigorously ground the importance of interactivity in a deep knowledge-driven sense. This interactivity is defined by the clinical user determining the fundamental knowledge about the image processing problem (specifically segmentation problems) and expressing that knowledge to the computer. This is in contrast to a more superficial view of interactivity in which the fundamental structure and assumptions of the medical image processing problem are both assumed and unassailable; the clinical user provides only minimal deep knowledge.

The notion of interactivity also arose in Chapters 3, 5 and 6, encouraging generality in order to incorporate (or at least not to constrain) knowledge-heavy interactivity. Chapter 6 also addressed the issue of improving the meaningfulness of interactivity with respect to parameter selection through demonstrating the increased parameter robustness resulting from incorporating other, more intuitive, forms of knowledge.

In response to the question: ***what is knowledge in a medical image processing problem?*** the principle of interactivity responds with a pragmatic constraint: ***knowledge cannot meaningfully exist in the void, but depends on communication with the clinical user***.

7.2 Unaddressed Aspects of Continuous Max-Flow

Continuous max-flow theory is a wide enough field that particular aspects of it must be left largely unaddressed. Because of their formulation as an optimization problem, continuous max-flow and graph-cuts have been used as a component in other medical image processing paradigms; two of particular interest, which have used max-flow and graph-cuts as a component, being:

- Level-set segmentation, and
- Multi-resolution deformable registration.

In addition, there are other cost terms that can encode different types of knowledge appropriate for different problems. These have come in the form of more complex regularization terms or existential priors which rigorously automate aspects of determining the relevant anatomical model in run-time.

7.2.1 Max-Flow and Graph-Cuts Propagated Level-Sets

Level sets are another optimization-based image labelling technique, but often incorporate highly non-convex energy functions. Thus, level sets are normally propagated through a gradient descent operator. [34, 40] Formally, a level set and its propagation can be expressed as:

$$\begin{aligned} &\min_{C} E(C) \text{ where } C \text{ is a curve or} \\ &\min_{\phi} E(\phi) \text{ where } \phi(x) \text{ is a function whose zero level-set, } \phi(x) = 0, \text{ is the curve } C \end{aligned} \tag{7.1}$$

which can be updated using the equation:

$$\phi(x) \leftarrow \phi(x) + \tau|\nabla\phi|F \text{ for a } \textit{speed} \text{ function } F \text{ such that } F \cdot n = \frac{\delta E(C)}{\delta C},\ n \text{ normal to } C. \tag{7.2}$$

The energy equation can be reformulated as a max-flow problem [191] with a discrete graph-cuts approximation [25] when linearised around a particular solution. This linearisation still maintains a regularization term giving it more expressiveness than the linearisation used in traditional gradient descent solvers. This allows for contours to evolve much more rapidly, solving a convex approximation of the energy optimally at each iteration. Additionally, this approach lends itself to having multiple coupled [175] or mutually-exclusive [144] contours, which would not be expressible in the traditional level-sets paradigm. Note however that the inclusion of these solvers into a level-sets paradigm does not change the local nature of the optimizers, and thus max-flow propagated level-sets are still vulnerable to performance variability based on initialization.

7.2.2 Multi-Resolution Deformable Registration

Multi-resolution approaches in image-based registration are popular in that they attempt to optimize the deformation field iteratively at finer and finer grained deformations. This *graduated non-convexity* approach allows for lower energy transformations to be found more readily despite only having locally optimal optimization methods available [164] and is widely used as an approximate solver for complex MRFs [101].

Graph-cuts have been used in non-rigid registration where the deformation field is assumed to be an integer vector with a finite range [170]. In this approach, the registration problem can be formulated as a partitioning problem suitable for graph-cut optimization and using a multi-resolution framework allows for larger ranges to be investigated while minimizing computational burden [163]. User-interaction in the form of landmarks have also been added to graph-cut based registration approaches [110]. Continuous max-flow has been used in a similar way, although the fuzzy labelling allowed by the continuous approach has higher interpolative potential [146, 169].

7.2.3 Other Regularization Functions

As suggested in Chapter 1 Section 1.4.1, the selection of the L2 norm in continuous max-flow was motivated not by theoretical constraints but by the desire for rotation invariance; that an L1 norm continuous max-flow was indeed possible as a limiting case (and a simpler one at that) of discrete graph-cuts.

As illustrated in Chapter 2, other norms can be used, specifically anisotropic directional norms. In fact, any p-norm magnitude of the gradient can be used with the algorithms as written in the thesis (with the exception of Chapter 6 which requires a anisotropic directional L2 norm). That is, the algorithms in this thesis can apply any regularization of the form:

$$|\nabla u|_p = \sqrt[p]{\sum_i \left|\frac{\delta u}{\delta x_i}\right|^p} \tag{7.3}$$

which includes the L1 and L2 norms examined earlier. In terms of variational optimization theory, this family of norms can be optimized for using the principle of duality as:

$$\int_\Omega |\mathbf{R}(x)\nabla u(x)|_p dx = \max_{|\mathbf{R}(x)^{-\top} q(x)|_{\tilde{p}} \leq 1} \int_\Omega u(x) \operatorname{div} q(x) dx \tag{7.4}$$

(shown in Appendix A.5) where $|\cdot|_{\tilde{p}}$ is the *dual norm* of $|\cdot|_p$. That is, another p-norm with the value $\tilde{p} = {}^{p}/_{p-1}$, which maps L1 to L$\infty$ and L2 to itself.

However, these norms are not the only regularization functions possible. Other functions such as the squared gradient magnitude $|\nabla u|^2$ have been proposed for continuous max-flow based registration [146]. The only requirement for the use of these functions in the current framework is that they are the dual of an easily computable function of the spatial flows. In the case of the squared gradient magnitude:

$$|\nabla u(x)|^2 = \max_q \int_\Omega \left(u(x) \operatorname{div} q(x) - {}^1/_4 |q(x)|^2 \right) dx \tag{7.5}$$

7.2.4 Existential Priors

There is another family of cost functions commonly used in discrete graph-cuts continuous max-flow segmentation techniques that this thesis does not discuss, which I shall refer to as *existential* priors. These are non-negative costs of the form:

$$h_\alpha = \begin{cases} c_\alpha, & \text{if } \exists x \in \Omega \text{ such that } \alpha \in u(x) \\ 0, & \text{else} \end{cases} \tag{7.6}$$

or the fuzzy equivalent (in terms of indicator functions):

$$h_\alpha = c_\alpha \max_{x \in \Omega} u_\alpha(x) \tag{7.7}$$

which penalize the occurrence (i.e. the existence) of a label in the segmentation. Thus, these priors encourage the labelling to use as few labels as possible. The benefit of such a structure is that it allows for the algorithm designer to incorporate labels and cost structures for objects that might be in the image. For example, in the segmentation of a chest CT, a cancer label may be attached with some existential prior. Thus, a clinician would not have to specify whether or not a patient had cancer, but leave it up to the computer to both detect the presence of cancer, adding it to the anatomical model of the patient, and further delineate it all as a single optimization task.

These types of priors have been previously investigated in a discrete context by Delong *et al.*[44] and in a continuous max-flow context by Yuan and Boykov [194] and Souiai *et al.* [165].

One philosophical issue one might have with these types of costs in medical image segmentation problems is they encourage the computer to determine the ontology of the segmentation, rather than taking advantage of the clinician's knowledge about the specific problem domain. Thus, existential priors err on the side of automaticity instead of interactivity.

7.3 The Future of Continuous Max-Flow in Medical Image Processing

There are several areas of medical image processing that are currently untouched by continuous max-flow theory. Thus, the future of continuous max-flow in medical image processing is bright with many *low hanging fruit* still available. Possible future applications and developments in continuous max-flow theory include:

- the use of continuous max-flow theory in image reconstruction problems with convex

constraints, in particular, undersampled MRI and low-dose CT. Both involve regularization terms with encourage image smoothness, and convex data constraints relating large number of voxels simultaneously to incompletely acquired or noisy data;

- the development of continuous analogues to higher-order clique terms traditionally used in more complex MRFs. This point is purposefully vague in that it is not currently clear what structure these terms might have and their connection to current and emerging medical image processing problems;
- the development of regularization terms that are well-founded in terms of the underlying imaging physics, especially in the context of image enhancement. For example, in QSM, the phase image away from susceptibility sources has a zero Laplacian, thus, distinguishing these sources may be possible from Laplacian-minimizing, rather than variation-minimizing, regularization;
- the development of algorithms for automatic, efficient and intelligent weighting parameter selection without requiring an *a priori* segmented dataset for training purposes; and
- the development of more intuitive and user-friendly interfaces for expressing anatomical knowledge in a way that is immediately translated into continuous max-flow objectives and constraints. These interfaces will meaningfully relate the medical and image processing aspects of medical image processing problems in a manner that is both general-purpose and intuitive, while minimizing or eliminating non-intuitive or opaque aspects such as parameter selection.

Thus, I expect this field to grow in both complexity and applicability but also in elegance as it matures.

Appendix A

Use of the Principle of Duality and Derivations of Algorithms

This appendix contains the derivations of the various variational optimization problems posed in this thesis using the principle of duality complemented by the Augmented Lagrangian and Proximal Bregman optimization paradigms.

A.1 CCMF Algorithm Derivation

The subsection concerns itself with the solution to the CCMF functional in terms of indicator functions as described in Eq. (2.23) and reproduced below:

$$\begin{aligned} &\min_{u} \int_{\Theta}\int_{\Omega} D_\theta(x)u_\theta(x)dxd\theta + \int_{\Theta}\int_{\Omega} |\mathbf{R}_\theta(x)\nabla u_\theta(x)|_p \, dxd\theta \\ &\text{s.t.} \int_{\Theta} u_\theta(x)d\theta = 1 \text{ and } u_\theta(x) \geq 0 \end{aligned} \tag{2.23}$$

A.1.1 Primal and Primal-Dual Models

The primal model suggested is:

$$\begin{aligned} &\max_{p,q} \int_{\Omega} p_S(x)dx \\ &\text{s.t. } p_\theta(x) \leq D_\theta(x) \text{ and } \left|\mathbf{R}_\theta^{-\top}(x)q_\theta(x)\right|_{\tilde{p}} \leq 1 \\ &\quad G_\theta(x) = \operatorname{div} q_\theta(x) + p_\theta(x) - p_S(x) = 0 \end{aligned} \tag{A.1}$$

with $G_\theta(x)$ is the *flow conservation constraint*. To create the primal-dual model, place a Lagrangian multiplier on said constraint, yielding:

$$\begin{aligned} \max_{p,q} \min_{u} \; & \int_\Omega p_S(x)dx + \int_\Theta \int_\Omega u_\theta(x) G_\theta(x) dx d\theta \\ \text{s.t.} \; & p_\theta(x) \leq D_\theta(x) \text{ and } \left| \mathbf{R}_\theta^{-\top}(x) q_\theta(x) \right|_{\tilde{p}} \leq 1 \end{aligned} \tag{A.2}$$

A.1.2 Equivalence to Dual Model

The primal and primal-dual models are trivially equivalent. In order to show the equivalence of these models to the desired dual model, one must optimize each variable showing that said optimization reconstructs the original minimization problems objective function and constraints. Isolating the optimization of q is equivalent to:

$$\begin{aligned} \max_{q} \; & \int_\Theta \int_\Omega u_\theta(x) \operatorname{div} q_\theta(x) dx d\theta \\ \text{s.t.} \; & \left| \mathbf{R}_\theta^{-\top}(x) q_\theta(x) \right|_{\tilde{p}} \leq 1 \end{aligned} \tag{A.3}$$

which, as shown in Appendix A.5, is equivalent to:

$$\int_\Theta \int_\Omega \left| \mathbf{R}(x) \nabla u_\theta(x) \right|_p dx d\theta \tag{A.4}$$

yielding the regularization portion of the objective function.

Isolating the outgoing flow, $p_\theta(x)$, yields the optimization:

$$\begin{aligned} \max_{p_\theta} \; & \int_\Theta \int_\Omega u_\theta(x) p_\theta(x) dx d\theta \\ \text{s.t.} \; & p_\theta(x) \leq D_\theta(x) \end{aligned} \tag{A.5}$$

which in order to be bounded requires $u_\theta(x) \geq 0$. Assuming non-negative $D_\theta(x)$ optimizes to:

$$\begin{aligned} & \int_\Theta \int_\Omega D_\theta(x) u_\theta(x) dx d\theta \\ \text{s.t.} \; & u_\theta(x) \geq 0 \end{aligned} \tag{A.6}$$

which reconstructs the data term portion of the objective function as well as the non-negativity constraint on the indicator functions.

Lastly, isolating $p_S(x)$ yields:

$$\max_{p_S} \int_\Omega p_S(x)dx - \int_\Theta \int_\Omega u_\theta(x) p_S(x) dx d\theta \tag{A.7}$$

which is only bounded if:

$$\int_{\Theta} u_\theta(x) d\theta = 1 \tag{A.8}$$

which reconstructs the last constraint on the indicator functions.

Thus, the primal-dual model is equivalent to Eq. (2.23) as maximizing the equation with respect to each variable yields both the objective function and constraints of the original formulation.

A.1.3 Augmented Lagrangian

The optimization problem addressed in the Augmented Lagrangian approach to the CCMF function is:

$$\begin{aligned} \max_{p,q} \min_{u} & \int_{\Omega} p_S(x) dx + \int_{\Theta} \int_{\Omega} u_\theta(x) G_\theta(x) dx d\theta - \frac{c}{2} \int_{\Theta} \int_{\Omega} G_\theta^2(x) dx d\theta \\ & \text{s.t. } p_\theta(x) \leq D_\theta(x) \text{ and } \left| \mathbf{R}_\theta^{-\top}(x) q_\theta(x) \right|_{\tilde{p}} \leq 1 \end{aligned} \tag{A.9}$$

which non-negative quadratic penalty parameter $c \geq 0$.

This problem is addressed by optimizing each variable iteratively while holding all others constant. This yields the following steps:

1. Maximize Eq. (A.9) with respect to $q_\theta(x)$ which can be accomplished through a Chambolle iteration [31]:

$$q_\theta(x) \leftarrow \text{Proj}_{\left|\mathbf{R}_\theta^{-\top}(x) q_\theta(x)\right|_{\tilde{p}} \leq 1} \left(q_\theta + \tau \nabla \left(\text{div}\, q_\theta(x) + p_\theta(x) - p_S(x) - u_\theta(x)/c \right) \right) \tag{A.10}$$

2. Maximize Eq. (A.9) with respect to $p_\theta(x)$ which can be done analytically, yielding:

$$p_\theta(x) \leftarrow \min \left\{ D_\theta(x),\ p_S(x) - \text{div}\, q_\theta(x) + u_\theta(x)/c \right\} \tag{A.11}$$

3. Maximize Eq. (A.9) with respect to $p_S(x)$ which can be done analytically, yielding:

$$p_S(x) \leftarrow \frac{1}{2\pi} (1/c + \int_{\Theta} (p_\theta(x) + \text{div}\, q_\theta(x) - u_\theta(x)/c) d\theta \tag{A.12}$$

4. Minimize Eq. (A.9) with respect to $u_\theta(x)$ which can be done analytically, yielding:

$$u_\theta(x) \leftarrow u_\theta(x) - c \left(\text{div}\, q_\theta(x) - p_S(x) + p_\theta(x) \right) \tag{A.13}$$

These steps yield Algorithm 2.6 reproduced below:

$\forall x, \theta,\ u_\theta(x) = {}^1/_{2\pi}$;
while *not converged* **do**
  $\forall x, \theta,\ q_\theta(x) \leftarrow \mathrm{Proj}_{\left|\mathbf{R}_\theta^{-\top}(x) q_\theta(x)\right|_{\tilde{p}} \leq 1} \left(q_\theta + \tau \nabla \left(\mathrm{div}\, q_\theta(x) + p_\theta(x) - p_S(x) - u_\theta(x)/c\right)\right)$;
  $\forall x, \theta,\ p_\theta(x) \leftarrow \min\{D_\theta(x), p_S(x) - \mathrm{div}\, q_\theta(x) + u_\theta(x)/c\}$;
  $\forall x,\quad p_S(x) \leftarrow \frac{1}{2\pi}(1/c + \int_\Theta (p_\theta(x) + \mathrm{div}\, q_\theta(x) - u_\theta(x)/c) d\theta$;
  $\forall x, \theta,\ u_\theta(x) \leftarrow u_\theta(x) - c\left(\mathrm{div}\, q_\theta(x) - p_S(x) + p_\theta(x)\right)$;
end

A.1.4 Proximal Bregman

In order to perform proximal Bregman optimization, one must first shown that the source and sink flows, $p_S(x)$ and $p_\theta(x)$, can be removed from the optimization problem by reintroducing the labelling function constraints. First, use the primal-dual model, Eq. (A.2):

$$
\begin{aligned}
&\max_{p,q} \min_{u} \int_\Omega p_S(x) dx + \int_\Theta \int_\Omega u_\theta(x) G_\theta(x) dx d\theta \\
&\qquad\qquad p_\theta(x) \leq D_\theta(x) \text{ and } \left|\mathbf{R}_\theta^{-\top}(x) q_\theta(x)\right|_{\tilde{p}} \leq 1 \\
&= \max_{p,q} \min_{u} \int_\Omega p_S(x) dx + \int_\Theta \int_\Omega u_\theta(x) \left(\mathrm{div}\, q_\theta(x) + p_\theta(x) - p_S(x)\right) dx d\theta \\
&\qquad\qquad p_\theta(x) \leq D_\theta(x) \text{ and } \left|\mathbf{R}_\theta^{-\top}(x) q_\theta(x)\right|_{\tilde{p}} \leq 1 \\
&= \max_{p,q} \min_{u} \int_\Theta \int_\Omega u_\theta(x) \left(\mathrm{div}\, q_\theta(x) + p_\theta(x)\right) dx d\theta \\
&\qquad\qquad p_\theta(x) \leq D_\theta(x) \text{ and } \left|\mathbf{R}_\theta^{-\top}(x) q_\theta(x)\right|_{\tilde{p}} \leq 1 \text{ and } \int_\Theta u_\theta(x) d\theta = 1 \\
&= \max_{q} \min_{u} \int_\Theta \int_\Omega u_\theta(x)\, \mathrm{div}\, q_\theta(x) + \int_\Theta \int_\Omega D_\theta(x) u_\theta(x) dx d\theta dx d\theta \\
&\qquad\qquad \left|\mathbf{R}_\theta^{-\top}(x) q_\theta(x)\right|_{\tilde{p}} \leq 1 \text{ and } u_\theta(x) \geq 0 \text{ and } \int_\Theta u_\theta(x) d\theta = 1
\end{aligned}
\tag{A.14}
$$

This formula can be optimized by iteratively replacing $u_\theta(x)$ with a proximal improved labelling for a fixed $q_\theta(x)$. This proximal Bregman projection addresses the optimization problem:

$$
u_\theta(x) \leftarrow \operatorname*{argmin}_{u'} \int_\Theta \int_\Omega u'_\theta(x)\, \mathrm{div}\, q_\theta(x) + \int_\Theta \int_\Omega D_\theta(x) u'_\theta(x) dx d\theta dx d\theta + c d_g(u', u) \tag{A.15}
$$

for any Bregman distance $d_g(u'(x), u(x))$ and positive c. To get a distance function, use the

function $g(u) = \int_\Theta \int_\Omega u_\theta(x) \ln u_\theta(x) dx d\theta$, which yields the distance:

$$d_g(u, v) = \int_\Theta \int_\Omega \left(u_\theta(x) \ln \frac{u_\theta(x)}{v_\theta(x)} - u_\theta(x) + v_\theta(x) \right) dx d\theta \tag{A.16}$$

this optimization problem (constrained by $\int_\Theta u_\theta(x) d\theta = 1$) can be solved analytically by:

$$u_\theta(x) \leftarrow \frac{u_L(x) \exp\left(-{}^{(D_\theta(x) + \operatorname{div} q_\theta(x))}/_{c}\right)}{\int_\Theta u_L(x) \exp\left(-{}^{(D_\theta(x) + \operatorname{div} q_\theta(x))}/_{c}\right) d\theta} \tag{A.17}$$

which also preserves the label function constraints $u_\theta(x) \geq 0$. Therefore, in order to address the optimization problem in Eq. (A.2), iterate between the steps:

1. Minimizing Eq. (A.14) with respect to $u_\theta(x)$ as described above, and
2. Maximizing Eq. (A.14) with respect to $q_\theta(x)$ using a Chambolle iteration [31]:

$$q_\theta(x) \leftarrow \operatorname{Proj}_{\left|\mathbf{R}_\theta^{-\top}(x) q_\theta(x)\right|_{\tilde{p}} \leq 1} \left(q_\theta - \tau \nabla \left(u_\theta(x) \exp\left(-{}^{(D_\theta(x) + \operatorname{div} q_\theta(x))}/_{c}\right)\right)\right) \tag{A.18}$$

These steps yield Algorithm 2.7 reproduced below:

$\forall x, \theta,\ u_\theta(x) = {}^{1}/_{2\pi}$;
while *not converged* **do**
  $\forall x, \theta,\ u_\theta(x) \leftarrow u_\theta(x) \exp\left(-{}^{(D_\theta(x) + \operatorname{div} q_\theta(x))}/_{c}\right)$;
  $\forall x, \theta, q_\theta(x) \leftarrow \operatorname{Proj}_{\left|\mathbf{R}_\theta^{-\top}(x) q_\theta(x)\right|_{\tilde{p}} \leq 1} \left(q_\theta - \tau \nabla u_\theta(x)\right)$;
  $\forall x,\quad a(x) \leftarrow \int_\Theta u_\theta(x) d\theta$;
  $\forall x, \theta,\ u_\theta(x) \leftarrow u_\theta(x) / a(x)$;
end

A.2 HMF Algorithm Derivation

The subsection concerns itself with the solution to the HMF functional as described in Eq. (3.1) and reproduced below:

$$\begin{aligned} &\min_{u} \sum_{L\in\mathbb{L}} \int_{\Omega} D_L(x)u_L(x)dx + \sum_{L} \int_{\Omega} \left|\mathbf{R}_L(x)\nabla u_L(x)\right|_p dx \\ &\text{s.t. } \forall L \in \mathbb{L},\ u_L(x) \geq 0 \\ &\qquad \forall L \notin \mathbb{L},\ u_L(x) = \sum_{L'\in L.C} u_{L'}(x) \\ &\qquad u_S(x) = 1 \end{aligned} \tag{3.1}$$

A.2.1 Primal and Primal-Dual Models

The primal model suggested is:

$$\begin{aligned} &\max_{p,q} \int_{\Omega} p_S(x)dx \\ &\text{s.t. } p_L(x) \leq D_L(x) \text{ for labels } L \in \mathbb{L} \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \leq 1 \\ &\qquad G_L(x) = \operatorname{div} q_L(x) + p_L(x) - p_{L.P}(x) = 0 \end{aligned} \tag{A.19}$$

in which $G_L(x)$ is the flow conservation constraint and the operators $L.C$ and $L.P$ encode the hierarchical label ordering. This is equivalent to a flow through a hierarchically constructed network with unconstrained inter-node flow (from parent to child) and constrained sink flow (from end-labels to the sink vertex). To create the primal-dual model, place a Lagrangian multiplier on the flow conservation constraint, yielding:

$$\begin{aligned} &\max_{p,q} \min_{u} \int_{\Omega} p_S(x)dx + \sum_{L} \int_{\Omega} u_L(x)G_L(x)dx \\ &\qquad \text{s.t. } p_L(x) \leq D_L(x) \text{ for labels } L \in \mathbb{L} \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \leq 1 \end{aligned} \tag{A.20}$$

A.2.2 Equivalence to Dual Model

The primal and primal-dual models are trivially equivalent. In order to show the equivalence of these models to the desired dual model, one must optimize each variable showing that said optimization reconstructs the original minimization problems objective function and constraints.

Isolating the optimization of q is equivalent to:

$$\begin{aligned} &\max_{q} \sum_{L} \int_{\Omega} u_L(x) \operatorname{div} q_L(x) dx \\ &\text{s.t. } \left| \mathbf{R}_L^{-\top}(x) q_L(x) \right|_{\tilde{p}} \leq 1 \end{aligned} \tag{A.21}$$

which, as shown in Appendix A.5, is equivalent to:

$$\sum_{L} \int_{\Omega} |\mathbf{R}(x) \nabla u_L(x)|_p dx \tag{A.22}$$

yielding the regularization portion of the objective function.

Isolating the outgoing flow, $p_L(x)$ at each end-label, yields the optimization:

$$\begin{aligned} &\max_{p_L} \int_{\Omega} u_L(x) p_L(x) dx \\ &\text{s.t. } p_L(x) \leq D_L(x) \end{aligned} \tag{A.23}$$

which in order to be bounded requires $u_L(x) \geq 0$. Assuming non-negative $D_L(x)$ optimizes to:

$$\begin{aligned} &\int_{\Omega} D_P(x) u_L(x) dx \\ &\text{s.t. } u_L(x) \geq 0 \end{aligned} \tag{A.24}$$

which reconstructs the data term portion of the objective function as well as the non-negativity constraint on the indicator functions.

Isolating the outgoing flow, $p_L(x)$ the other non-source labels, yields the optimization:

$$\max_{p_L} \int_{\Omega} u_L(x) p_L(x) - \sum_{L' \in L.C} \int_{\Omega} u_{L'}(x) p_{L'}(x) dx \tag{A.25}$$

which is only bounded when:

$$\sum_{L' \in L.C} u_{L'}(x) = u_L(x) \tag{A.26}$$

and thus reconstructs the label ordering constraint $\sum_{L' \in L.C} u_{L'}(x) = u_L(x)$.

Lastly, isolating $p_S(x)$ yields:

$$\max_{p_S} \int_{\Omega} p_S(x) dx - \sum_{L \in S.C} \int_{\Omega} u_L(x) p_S(x) dx \tag{A.27}$$

which is only bounded if:

$$\sum_{L \in S.C} u_L(x) = 1 \tag{A.28}$$

which reconstructs the last constraint, $u_S(x) = 1$.

Thus, the primal-dual model is equivalent to Eq. (3.1) as maximizing the equation with respect to each variable yields both the objective function and constraints of the original formulation.

A.2.3 Augmented Lagrangian

The optimization problem addressed in the Augmented Lagrangian approach to the HMF function is:

$$\begin{aligned} \max_{p,q} \min_{u} & \int_\Omega p_S(x)dx + \sum_L \int_\Omega u_L(x) G_L(x) dx - \frac{c}{2} \sum_L G_L^2(x) \\ \text{s.t. } & p_L(x) \le D_L(x) \text{ for labels } L \in \mathbb{L} \text{ and } \left| \mathbf{R}_L^{-\top}(x) q_L(x) \right|_{\tilde{p}} \le 1 \end{aligned} \tag{A.29}$$

which non-negative quadratic penalty parameter $c \ge 0$.

This problem is addressed by optimizing each variable iteratively while holding all others constant. This yields the following steps:

1. Maximize Eq. (A.29) with respect to $q_L(x)$ which can be accomplished through a Chambolle iteration [31]:

$$q_L(x) \leftarrow \text{Proj}_{\left|\mathbf{R}_L^{-\top}(x) q_L(x)\right|_{\tilde{p}} \le 1} \left(q_L + \tau \nabla \left(\text{div}\, q_L(x) + p_L(x) - p_{L.P}(x) - u_L(x)/c \right) \right) \tag{A.30}$$

2. Maximize Eq. (A.29) with respect to $p_L(x)$ at the end-labels which can be done analytically, yielding:

$$p_L(x) \leftarrow \min \{ D_L(x), \; p_{L.P}(x) - \text{div}\, q_L(x) + u_L(x)/c \} \tag{A.31}$$

3. Maximize Eq. (A.29) with respect to $p_L(x)$ at the other non-source labels which can be done analytically, yielding:

$$p_L(x) \leftarrow \frac{1}{|L.C| + 1} \left(p_{L.P}(x) - \text{div}\, q_L(x) + u_L(x)/c + \sum_{L' \in L.C} \left(p_{L'}(x) + \text{div}\, q_{L'}(x) - u_{L'}(x)/c \right) \right) \tag{A.32}$$

4. Maximize Eq. (A.29) with respect to $p_S(x)$ which can be done analytically, yielding:

$$p_S(x) \leftarrow \frac{1}{|S.C|}\left(1/c + \sum_{L \in S.C} (p_L(x) + \operatorname{div} q_L(x) - u_L(x)/c)\right) \tag{A.33}$$

5. Minimize Eq. (A.29) with respect to $u_L(x)$ which can be done analytically, yielding:

$$u_L(x) \leftarrow u_L(x) - c\,(\operatorname{div} q_L(x) - p_{L.P}(x) + p_L(x)) \tag{A.34}$$

This yields Algorithm 3.2 reproduced on the following page.

InitializeSolution(S) ;
while *not converged* **do**
  for $\forall L \neq S$ **do**
    $\forall x, q_L(x) \leftarrow \mathrm{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} \left(q_L + \tau \nabla \left(\mathrm{div}\, q_L(x) + p_L(x) - p_{L.P}(x) - u_L(x)/c \right) \right)$;
  end
  UpdateFlows(S) ;
  for $\forall L \neq S$ **do**
    $\forall x, u_L(x) \leftarrow u_L(x) - c \left(\mathrm{div}\, q_L(x) - p_{L.P}(x) + p_L(x) \right)$;
  end
end

InitializeSolution(L)
for $\forall L' \in L.C$ **do**
  InitializeSolution(L');
end
$\forall x,\ p_L(x) \leftarrow 0$;
if $L \neq S$ **then**
  $\forall x,\ u_L(x) \leftarrow 0$;
  $\forall x,\ q_L(x) \leftarrow 0$;
end

UpdateFlows(L)
for $\forall L' \in L.C$ **do**
  UpdateSinkFlows(L') ;
end
if $L \in \mathbb{L}$ **then**
  $\forall x, p_L(x) \leftarrow \min\{D_L(x), p_{L.P}(x) - \mathrm{div}\, q_L(x) + u_L(x)/c\}$;
else if $L = S$ **then**
  $\forall x, p_S(x) \leftarrow 1/c$;
  for $\forall L' \in S.C$ **do**
    $\forall x, p_S(x) \leftarrow p_S(x) + p_{L'}(x) + \mathrm{div}\, q_{L'}(x) - u_{L'}(x)/c$;
  end
  $\forall x, p_S(x) \leftarrow \frac{1}{|S.C|} p_S(x)$;
else
  $\forall x, p_L(x) \leftarrow p_{L.P}(x) - \mathrm{div}\, q_L(x) + u_L(x)/c$;
  for $\forall L' \in L.C$ **do**
    $\forall x, p_L(x) \leftarrow p_L(x) + p_{L'}(x) + \mathrm{div}\, q_{L'}(x) - u_{L'}(x)/c$;
  end
  $\forall x, p_L(x) \leftarrow \frac{1}{|L.C|+1} p_L(x)$;
end

A.2.4 Proximal Bregman

In order to perform proximal Bregman optimization, one must first shown that the source, sink, and inter-node flows, $p_S(x)$ and $p_L(x)$, can be removed from the optimization problem by reintroducing the labelling function constraints. First, use the primal-dual model, Eq. (A.20):

$$\begin{aligned}
&\max_{p,q}\min_{u}\ \int_\Omega p_S(x)dx + \sum_L \int_\Omega u_L(x)G_L(x)dx \\
&\qquad \text{s.t. } \forall L\in\mathbb{L},\ p_L(x)\le D_L(x) \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \le 1 \\
&= \max_{p,q}\min_{u}\ \int_\Omega p_S(x)dx + \sum_L \int_\Omega u_L(x)\left(\operatorname{div} q_L(x) + p_L(x) - p_{L.P}(x)\right)dx \\
&\qquad \text{s.t. } \forall L\in\mathbb{L},\ p_L(x)\le D_L(x) \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \le 1 \\
&= \max_{p,q}\min_{u}\ \int_\Omega p_S(x)dx - \sum_{L\in S.C}\int_\Omega u_L(x)p_S(x)dx + \sum_L\int_\Omega u_L(x)\operatorname{div} q_L(x)dx + \sum_{L\in\mathbb{L}}\int_\Omega u_L(x)p_L(x)dx \\
&\qquad \text{s.t. } \forall L\in\mathbb{L},\ p_L(x)\le D_L(x) \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \le 1 \text{ and } \sum_{L'\in L.C} u_{L'}(x) = u_L(x) \\
&= \max_{p,q}\min_{u}\ \sum_L\int_\Omega u_L(x)\operatorname{div} q_L(x)dx + \sum_{L\in\mathbb{L}}\int_\Omega u_L(x)p_L(x)dx \\
&\qquad \text{s.t. } \forall L\in\mathbb{L},\ p_L(x)\le D_L(x) \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \le 1 \text{ and } \sum_{L'\in L.C} u_{L'}(x) = u_L(x) \text{ and } u_S(x) = 1 \\
&= \max_{q}\min_{u}\ \sum_L\int_\Omega u_L(x)\operatorname{div} q_L(x)dx + \sum_{L\in\mathbb{L}}\int_\Omega u_L(x)D_L(x)dx \\
&\qquad \text{s.t. } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \le 1 \text{ and } \sum_{L'\in L.C} u_{L'}(x) = u_L(x) \text{ and } u_S(x) = 1 \text{ and } u_L(x)\ge 0
\end{aligned} \tag{A.35}$$

By introducing the variables $d_L(x)$ defined recursively as:

$$d_L(x) = \begin{cases} 0, & L = S \\ d_{L.P}(x) + \operatorname{div} q_L(x), & L\neq S \text{ and } L\notin\mathbb{L} \\ d_{L.P}(x) + \operatorname{div} q_L(x) + D_L(x), & L\in\mathbb{L} \end{cases}. \tag{A.36}$$

the former optimization can be written more simply as:

$$\begin{aligned}
&= \max_{q}\min_{u}\ \sum_{L\in\mathbb{L}}\int_\Omega u_L(x)d_L(x)dx \\
&\qquad \text{s.t. } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \le 1 \text{ and } \sum_{L\in\mathbb{L}} u_L(x) = 1 \text{ and } u_L(x)\ge 0
\end{aligned} \tag{A.37}$$

This formula can be optimized by iteratively replacing $u_L(x)$ with a proximal improved labelling for a fixed $q_L(x)$. This proximal Bregman projection addresses the optimization problem:

$$u_L(x) \leftarrow \underset{u'}{\operatorname{argmin}} \sum_L \int_\Omega u_L(x) d_L(x) dx + c d_g(u', u) \tag{A.38}$$

for any Bregman distance $d_g(u'(x), u(x))$ and positive c. To get a distance function, use the function $g(u) = \sum_{L\in\mathbb{L}} \int_\Omega u_L(x) \ln u_L(x) dx$, which yields the distance:

$$d_g(u, v) = \sum_{L\in\mathbb{L}} \int_\Omega \left(u_L(x) \ln \frac{u_L(x)}{v_L(x)} - u_L(x) + v_L(x) \right) dx \tag{A.39}$$

this optimization problem (constrained by $u_S(x) = 1$ and $\sum_{L'\in L.C} u_{L'}(x) = u_L(x)$ with a Lagrangian multipler on the former) can be solved analytically for the end-labels by:

$$u_L(x) \leftarrow \frac{u_L(x) \exp\left(-\frac{d_L(x)}{c}\right)}{\sum_{L\in\mathbb{L}} u_L(x) \exp\left(-\frac{d_L(x)}{c}\right)} \tag{A.40}$$

which also preserves the label function constraints $u_L(x) \geq 0$.

Lastly, one needs to determine how to optimize Eq. (A.37) with respect to $q_L(x)$ by find the gradient thereof with respect to $\operatorname{div} q_L(x)$. In order to simplify notation, introduce the variable $W_{(A,B)}$ with the recursive definition:

$$W_{(A,B)}(x) = \begin{cases} 1, & A = B \\ 0, & A \neq B \text{ and } A \notin B.P^* \\ \sum_{L\in A.C} W_{(L,B)}, & L \in B.P^* \end{cases} . \tag{A.41}$$

Following this gradient yields the Chambolle iteration [31]:

$$q_L(x) \leftarrow \begin{cases} \operatorname{Proj}_{\left|\mathbf{R}_L^{-\top}(x) q_L(x)\right|_{\tilde{p}} \leq 1} \left(q_L - \tau \nabla \left(u_L(x) \exp\left({}^{-(D_L(x) + \operatorname{div} q_L(x))}/_{c} \right) \right) \right), & L \in \mathbb{L} \\ \operatorname{Proj}_{\left|\mathbf{R}_L^{-\top}(x) q_L(x)\right|_{\tilde{p}} \leq 1} \left(q_L - \tau \nabla \left(\sum_{L'\in\mathbb{L}} W_{(L,L')} u_{L'}(x) \exp\left({}^{-(D_L(x) + \operatorname{div} q_L(x))}/_{c} \right) \right) \right), & L \notin \mathbb{L} \end{cases} \tag{A.42}$$

Therefore, in order to address the optimization problem in Eq. (A.37), iterate between:

1. Minimizing Eq. (A.37) with respect to $u_L(x)$ as described above, and
2. Maximizing Eq. (A.37) with respect to $q_L(x)$ using a Chambolle iteration [31] described above.

This yields Algorithm 3.2 reproduced on the following page.

$\forall L \in \mathbb{L},\ u_L(x) \leftarrow 1/|\mathbb{L}|;$
$\forall x,\ d_S(x) = 0;$
while *not converged* **do**
  PushDownCosts(S);
  $\forall x, L \in \mathbb{L},\ d_L(x) \leftarrow u_L(x) \exp\left(-\frac{d_L(x)}{c}\right);$
  $\forall x, L \in \mathbb{L},\ u_L(x) \leftarrow d_L(x) / \sum_{L' \in \mathbb{L}} d_{L'}(x);$
  $\forall x, L \notin \mathbb{L},\ d_L(x) \leftarrow 0;$
  PushUpCapacities(S);
end

PushDownCosts(L)
if $L \in L$ **then**
  $\forall x,\ d_L(x) \leftarrow d_{L.P}(x) + \operatorname{div} q_L(x) + D_L(x);$
else if $L \neq S$ **then**
  $\forall x,\ d_L(x) \leftarrow d_{L.P}(x) + \operatorname{div} q_L(x);$
end
for $\forall L' \in L.C$ **do**
  PushDownCosts(L');
end

PushUpCapacities(L)
for $\forall L' \in L.C$ **do**
  PushUpCapacities(L');
  if $L \neq S$ **then**
    $\forall x,\ q_L(x) \leftarrow \operatorname{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} (q_L(x) - c\tau \nabla d_L(x));$
    $\forall x,\ d_{L.P}(x) \leftarrow d_{L.P} + d_L(x);$
  end
end

A.3 DAGMF Algorithm Derivation

The subsection concerns itself with the solution to the DAGMF functional as described in Eq. (5.2) and reproduced below:

$$\begin{aligned} &\min_{u} \sum_{L\in\mathbb{L}} \int_{\Omega} D_L(x)u_L(x)dx + \sum_{L} \int_{\Omega} |\mathbf{R}_L(x)\nabla u_L(x)|_p \, dx \\ &\text{s.t. } \forall L \in \mathbb{L},\ u_L(x) \geq 0 \\ &\qquad \forall L \notin \mathbb{L},\ u_L(x) = \sum_{L'\in L.C} w_{(L,L')}u_{L'}(x) \\ &\qquad u_S(x) = 1 \end{aligned} \tag{5.2}$$

A.3.1 Primal and Primal-Dual Models

The primal model suggested is:

$$\begin{aligned} &\max_{p,q} \int_{\Omega} p_S(x)dx \\ &\text{s.t. } \forall L \in \mathbb{L},\ p_L(x) \leq D_L(x) \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \leq 1 \\ &\qquad G_L(x) = \operatorname{div} q_L(x) + p_L(x) - \sum_{L'\in L.P} w_{(L',L)}p_{L'}(x) = 0 \end{aligned} \tag{A.43}$$

in which $G_L(x)$ is the flow conservation constraint and the operators $L.C$ and $L.P$ encode the label ordering. (For the sake of simplicity, an additional set of variables is defined as:

$$\zeta_L(x) = \sum_{L'\in L.P} w_{(L',L)}p_{L'}(x) \tag{A.44}$$

which represent the incoming flow to any particular label, simplifying flow conservation to $G_L(x) = \operatorname{div} q_L(x) + p_L(x) - \zeta_L(x)$.) This is equivalent to a flow through a network with a over-arching DAG structure. The inter-node flow (from parents to children) are unconstrained with constrained sink flows (from end-labels to the sink vertex). To create the primal-dual model, place a Lagrangian multiplier on the flow conservation constraint, yielding:

$$\begin{aligned} &\max_{p,q}\min_{u} \int_{\Omega} p_S(x)dx + \sum_{L} \int_{\Omega} u_L(x)G_L(x)dx \\ &\qquad \text{s.t. } \forall L \in \mathbb{L},\ p_L(x) \leq D_L(x) \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \leq 1 \end{aligned} \tag{A.45}$$

A.3.2 Equivalence to Dual Model

The primal and primal-dual models are trivially equivalent. In order to show the equivalence of these models to the desired dual model, one must optimize each variable showing that said optimization reconstructs the original minimization problems objective function and constraints. Isolating the optimization of q is equivalent to:

$$\begin{aligned} &\max_{q} \sum_{L} \int_{\Omega} u_L(x) \operatorname{div} q_L(x) dx \\ &\text{s.t. } \left| \mathbf{R}_L^{-\top}(x) q_L(x) \right|_{\tilde{p}} \leq 1 \end{aligned} \tag{A.46}$$

which, as shown in Appendix A.5, is equivalent to:

$$\sum_{L} \int_{\Omega} \left| \mathbf{R}_L(x) \nabla u_L(x) \right|_p dx \tag{A.47}$$

yielding the regularization portion of the objective function.

Isolating the outgoing flow, $p_L(x)$ at each end-label, yields the optimization:

$$\begin{aligned} &\max_{p_L} \int_{\Omega} u_L(x) p_L(x) dx \\ &\text{s.t. } p_L(x) \leq D_L(x) \end{aligned} \tag{A.48}$$

which in order to be bounded requires $u_L(x) \geq 0$. Assuming non-negative $D_L(x)$ optimizes to:

$$\begin{aligned} &\int_{\Omega} D_L(x) u_L(x) dx \\ &\text{s.t. } u_L(x) \geq 0 \end{aligned} \tag{A.49}$$

which reconstructs the data term portion of the objective function as well as the non-negativity constraint on the indicator functions.

Isolating the outgoing flow, $p_L(x)$ the other non-source labels, yields the optimization:

$$\max_{p_L} \int_{\Omega} u_L(x) p_L(x) - \sum_{L' \in L.C} \int_{\Omega} w_{(L,L')} u_{L'}(x) p_{L'}(x) dx \tag{A.50}$$

which is only bounded when:

$$\sum_{L' \in L.C} w_{(L,L')} u_{L'}(x) = u_L(x) \tag{A.51}$$

and thus reconstructs the label ordering constraint $\sum_{L' \in L.C} w_{(L,L')} u_{L'}(x) = u_L(x)$.

Lastly, isolating $p_S(x)$ yields:

$$\max_{p_S} \int_\Omega p_S(x)dx - \sum_{L \in S.C} \int_\Omega w_{(S,L)} u_L(x) p_S(x)dx \tag{A.52}$$

which is only bounded if:

$$\sum_{L \in S.C} w_{(S,L)} u_L(x) = 1 \tag{A.53}$$

which reconstructs the last constraint, $u_S(x) = 1$.

Thus, the primal-dual model is equivalent to Eq. (5.2) as maximizing the equation with respect to each variable yields both the objective function and constraints of the original formulation.

A.3.3 Augmented Lagrangian

The optimization problem addressed in the Augmented Lagrangian approach to the DAGMF function is:

$$\begin{aligned} \max_{p,q} \min_{u} & \int_\Omega p_S(x)dx + \sum_L \int_\Omega u_L(x) G_L(x)dx - \frac{c}{2} \sum_L G_L^2(x) \\ \text{s.t. } & \forall L \in \mathbb{L},\ p_L(x) \le D_L(x) \text{ and } \left| \mathbf{R}_L^{-\top}(x) q_L(x) \right|_{\tilde{p}} \le 1 \end{aligned} \tag{A.54}$$

which non-negative quadratic penalty parameter $c \ge 0$.

This problem is addressed by optimizing each variable iteratively while holding all others constant. This yields the following steps:

1. Maximize Eq. (A.54) with respect to $q_L(x)$ which can be accomplished through a Chambolle iteration [31]:

$$q_L(x) \leftarrow \text{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \le 1} \left(q_L + \tau \nabla \left(\text{div}\, q_L(x) + p_L(x) - \zeta_{L'}(x) - u_L(x)/c \right) \right) \tag{A.55}$$

2. Maximize Eq. (A.54) with respect to $p_L(x)$ at the end-labels which can be done analytically, yielding:

$$p_L(x) \leftarrow \min \{ D_L(x),\ \zeta_L(x) - \text{div}\, q_L(x) + u_L(x)/c \} \tag{A.56}$$

3. Maximize Eq. (A.54) with respect to $p_L(x)$ at the other non-source labels which can be

done analytically, yielding:

$$p_L(x) \leftarrow \frac{1}{1 + \sum_{L' \in L.C} w^2_{(L,L')}} \Bigg(\zeta_L(x) - \operatorname{div} q_L(x) + u_L(x)/c + \sum_{L' \in L.C} w_{(L,L')} \left(p_{L'}(x) + \operatorname{div} q_{L'}(x) - \zeta_{L'}(x) + w_{(L,L')} p_L(x) - u_{L'}(x)/c \right) \Bigg) \tag{A.57}$$

4. Maximize Eq. (A.54) with respect to $p_S(x)$ which can be done analytically, yielding:

$$p_S(x) \leftarrow \frac{1}{\sum_{L \in S.C} w^2_{(S,L)}} \left(1/c + \sum_{L \in S.C} w_{(S,L)} \left(p_L(x) + \operatorname{div} q_L(x) - \zeta_L(x) + w_{(S,L)} p_S(x) - u_L(x)/c \right) \right) \tag{A.58}$$

5. Minimize Eq. (A.54) with respect to $u_L(x)$ which can be done analytically, yielding:

$$u_L(x) \leftarrow u_L(x) - c \left(\operatorname{div} q_L(x) - \zeta_L(x) + p_L(x) \right) \tag{A.59}$$

This yields Algorithm 5.1 reproduced on the following pages.

Topological sort the DAG into ordering $\mathbb{O}$ (begins with source label S) with reverse ordering $\mathbb{O}^{-1}$ (ends with source label S);
InitializeSolution() ;
while *not converged* **do**
  UpdateFlows() ;
  for $\forall L$ **do**
    $\forall x, u_L(x) \leftarrow u_L(x) - c\,(\operatorname{div} q_L(x) - \zeta_L(x) + p_L(x))$;
  end
end

InitializeSolution()
Clear $u_L(x), q_L(x)$ for all labels;
for *each L in order* $\mathbb{O}^{-1}$ **do**
  $\forall x, p_L(x) \leftarrow \min_{L'.C=\emptyset} D_{L'}(x)$;
  $\forall x, \zeta_L(x) \leftarrow \min_{L'.C=\emptyset} D_{L'}(x)$;
  if $L.C = \emptyset$ **then**
    if $L \in \operatorname{argmin}_{L'.C=\emptyset} D_{L'}(x)$ **then**
      $\forall x, u_L(x) \leftarrow 1/|\operatorname{argmin}_{L'.C=\emptyset} D_{L'}(x)|$;
    else
      $\forall x, u_L(x) \leftarrow 0$;
    end
  end
  for *each* $L' \in L.P/\{S\}$ **do**
    $\forall x, u_{L'}(x) \leftarrow u_{L'}(x) + w_{(L',L)} u_L(x)$;
  end
end

UpdateFlows()
for $\forall L \neq S$ **do**
  $\forall x, q_L(x) \leftarrow \text{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} \left(q_L + \tau \nabla \left(\text{div}\, q_L(x) + p_L(x) - \zeta_L(x) - u_L(x)/c\right)\right)$;
end
Clear $\zeta_L(x)$ for all labels ;
for *each L in order* $\mathbb{O}$ **do**
  for *each* $L' \in L.C$ **do**
    $\forall x, \zeta_{L'}(x) \leftarrow \zeta_{L'}(x) + w_{(L,L')} p_L(x)$;
  end
  if $L.C \neq \emptyset$ *and* $L.P \neq \emptyset$ **then**
    $\forall x, \sigma_L(x) \leftarrow \zeta(x) - \text{div}\, q_L(x) + u_L(x)/c$;
  else if $L = S$ **then**
    $\forall x, \sigma_S(x) \leftarrow 1/c$;
  end
end
for *each L in order* $\mathbb{O}^{-1}$ **do**
  if $L.C = \emptyset$ **then**
    $\forall x, p_L(x) \leftarrow \min\{D_L(x), \zeta_L(x) - \text{div}\, q_L(x) + u_L(x)/c\}$;
    for $L' \in L.P$ **do**
      $\forall x, \sigma_{L'}(x) \leftarrow \sigma_{L'}(x) + w_{(L',L)} \left(\text{div}\, q_{L'}(x) + p_{L'}(x) - \zeta_{L'}(x) + w_{L',L} p_L(x)\right)$;
    end
  else if $L = S$ **then**
    $\forall x, p_S(x) \leftarrow \frac{1}{\sum_{L' \in S.C} w^2_{(S,L')}} \sigma_S(x)$;
  else
    $\forall x, p_L(x) \leftarrow \frac{1}{1 + \sum_{L' \in L.C} w^2_{(L,L')}} \sigma_L(x)$;
    for $L' \in L.P$ **do**
      $\forall x, \sigma_{L'}(x) \leftarrow \sigma_{L'}(x) + w_{(L',L)} \left(\text{div}\, q_{L'}(x) + p_{L'}(x) - \zeta_{L'}(x) + w_{L',L} p_L(x)\right)$;
    end
  end
end

A.3.4 Proximal Bregman

In order to perform proximal Bregman optimization, one must first shown that the source, sink, and inter-node flows, $p_S(x)$ and $p_L(x)$, can be removed from the optimization problem by reintroducing the labelling function constraints. First, use the primal-dual model, Eq. (A.45):

$$\begin{aligned}
&\max_{p,q}\min_{u} \int_\Omega p_S(x)dx + \sum_L \int_\Omega u_L(x)G_L(x)dx \\
&\qquad \text{s.t. } \forall L \in \mathbb{L},\ p_L(x) \leq D_L(x) \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \leq 1 \\
&= \max_{p,q}\min_{u} \int_\Omega p_S(x)dx + \sum_L \int_\Omega u_L(x)\left(\operatorname{div} q_L(x) + p_L(x) - \zeta_L(x)\right)dx \\
&\qquad \text{s.t. } \forall L \in \mathbb{L},\ p_L(x) \leq D_L(x) \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \leq 1 \\
&= \max_{p,q}\min_{u} \int_\Omega p_S(x)dx - \sum_{L\in S.C}\int_\Omega u_L(x)p_S(x)dx + \sum_L \int_\Omega u_L(x)\operatorname{div} q_L(x)dx + \sum_{L\in\mathbb{L}}\int_\Omega u_L(x)p_L(x)dx \\
&\qquad \text{s.t. } \forall L \in \mathbb{L},\ p_L(x) \leq D_L(x) \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \leq 1 \text{ and } \sum_{L'\in L.C} w_{(L,L')}u_{L'}(x) = u_L(x) \\
&= \max_{p,q}\min_{u} \sum_L \int_\Omega u_L(x)\operatorname{div} q_L(x)dx + \sum_{L\in\mathbb{L}}\int_\Omega u_L(x)p_L(x)dx \\
&\qquad \text{s.t. } \forall L \in \mathbb{L},\ p_L(x) \leq D_L(x) \text{ and } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \leq 1 \text{ and } \sum_{L'\in L.C} w_{(L,L')}u_{L'}(x) = u_L(x) \text{ and } u_S(x) = 1 \\
&= \max_{q}\min_{u} \sum_L \int_\Omega u_L(x)\operatorname{div} q_L(x)dx + \sum_{L\in\mathbb{L}}\int_\Omega u_L(x)D_L(x)dx \\
&\qquad \text{s.t. } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \leq 1 \text{ and } \sum_{L'\in L.C} w_{(L,L')}u_{L'}(x) = u_L(x) \text{ and } u_S(x) = 1 \text{ and } u_L(x) \geq 0
\end{aligned} \tag{A.60}$$

By introducing the variables $d_L(x)$ defined recursively as:

$$d_L(x) = \begin{cases} 0, & L = S \\ \sum_{L'\in L.P} w_{(L',L)}d_{L'}(x) + \operatorname{div} q_L(x), & L \neq S \text{ and } L \notin \mathbb{L} \\ \sum_{L'\in L.P} w_{(L',L)}d_{L'}(x) + \operatorname{div} q_L(x) + D_L(x), & L \in \mathbb{L} \end{cases}. \tag{A.61}$$

the former optimization can be written more simply as:

$$\begin{aligned}
&= \max_{q}\min_{u} \sum_{L\in\mathbb{L}} \int_\Omega u_L(x)d_L(x)dx \\
&\qquad \text{s.t. } \left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_{\tilde{p}} \leq 1 \text{ and } \sum_{L\in\mathbb{L}} u_L(x) = 1 \text{ and } u_L(x) \geq 0
\end{aligned} \tag{A.62}$$

This formula can be optimized by iteratively replacing $u_L(x)$ with a proximal improved labelling for a fixed $q_L(x)$. This proximal Bregman projection addresses the optimization problem:

$$u_L(x) \leftarrow \underset{u'}{\operatorname{argmin}} \sum_L \int_\Omega u_L(x) d_L(x) dx + c d_g(u', u) \tag{A.63}$$

for any Bregman distance $d_g(u'(x), u(x))$ and positive c. To get a distance function, use the function $g(u) = \sum_{L\in\mathbb{L}} \int_\Omega u_L(x) \ln u_L(x) dx$, which yields the distance:

$$d_g(u, v) = \sum_{L\in\mathbb{L}} \int_\Omega \left(u_L(x) \ln \frac{u_L(x)}{v_L(x)} - u_L(x) + v_L(x) \right) dx \tag{A.64}$$

this optimization problem (using the constraints $u_S(x) = 1$ and $\sum_{L'\in L.C} u_{L'}(x) = u_L(x)$ with a Lagrangian multipler on the former) can be solved analytically for the end-labels by:

$$u_L(x) \leftarrow \frac{u_L(x) \exp\left(-\frac{d_L(x)}{c}\right)}{\sum_{L\in\mathbb{L}} u_L(x) \exp\left(-\frac{d_L(x)}{c}\right)} \tag{A.65}$$

which also preserves the label function constraints $u_L(x) \geq 0$.

Lastly, one needs to determine how to optimize Eq. (A.62) with respect to $q_L(x)$ by find the gradient thereof with respect to $\operatorname{div} q_L(x)$. In order to simplify notation, introduce the variable $W_{(A,B)}$ with the recursive definition:

$$W_{(A,B)}(x) = \begin{cases} 1, & A = B \\ 0, & A \neq B \text{ and } A \notin B.P^* \\ \sum_{L\in A.C} w_{(A,L)} W_{(L,B)}, & L \in B.P^* \end{cases} . \tag{A.66}$$

Following this gradient yields the Chambolle iteration [31]:

$$q_L(x) \leftarrow \begin{cases} \operatorname{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} \left(q_L - \tau \nabla \left(u_L(x) \exp\left({}^{-(D_L(x) + \operatorname{div} q_L(x))}/_{c} \right) \right) \right), & L \in \mathbb{L} \\ \operatorname{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} \left(q_L - \tau \nabla \left(\sum_{L'\in\mathbb{L}} W_{(L,L')} u_{L'}(x) \exp\left({}^{-(D_L(x) + \operatorname{div} q_L(x))}/_{c} \right) \right) \right), & L \notin \mathbb{L} \end{cases} \tag{A.67}$$

Therefore, in order to address the optimization problem in Eq. (A.62), iterate between:

1. Minimizing Eq. (A.62) with respect to $u_L(x)$ as described above, and
2. Maximizing Eq. (A.62) with respect to $q_L(x)$ using a Chambolle iteration [31] described above.

This yields Algorithm 5.4 reproduced on the following page.

Topological sort the DAG into the sorted list $\mathbb{O}$ (begins with source label S) with reverse list $\mathbb{O}^{-1}$ (ends with source label S);
$\forall L \in \mathbb{L},\ u_L(x) \leftarrow 1/|\mathbb{L}|$;
while *not converged* **do**
  $\forall L, \forall x,\ d_L(x) \leftarrow \operatorname{div} q_L(x)$;
  $\forall L \in \mathbb{L}, \forall x,\ d_L(x) \leftarrow d_L(x) + D_L(x)$;
  for L *in order* $\mathbb{O}/\{S\}$ **do**
    for $L' \in L.C$ **do**
      $\forall x,\ d_{L'}(x) \leftarrow d_{L'}(x) + w_{(L,L')} d_L(x)$;
    end
  end
  $\forall L \in \mathbb{L}, \forall x,\ u_L(x) \leftarrow u_L(x) \exp\left(-\frac{d_L(x)}{c}\right)$;
  $\forall L \in \mathbb{L}, \forall x,\ d_L(x) \leftarrow u_L(x)$;
  $\forall x,\ a(x) \leftarrow \sum_{L \in \mathbb{L}} u_L(x)$;
  $\forall L \in \mathbb{L}, \forall x,\ u_L(x) \leftarrow u_L(x)/a(x)$;
  $\forall L \notin \mathbb{L}, \forall x,\ d_L(x) \leftarrow 0$;
  for L *in order* $\mathbb{O}^{-1}/\{S\}$ **do**
    $\forall x,\ q_L(x) \leftarrow \operatorname{Proj}_{|\mathbf{R}_L^{-\top}(x) q_L(x)|_{\tilde{p}} \leq 1} \left(q_L(x) - c\tau \nabla d_L(x)\right)$;
    for $L' \in L.P/\{S\}$ **do**
      $\forall x,\ d_{L'}(x) \leftarrow d_{L'} + w_{(L',L)} d_L(x)$;
    end
  end
end

A.4 Discretization and Memory Consumption

All of the continuous max-flow algorithms presented are formulated in a way that is agnostic to how Ω is discretized. In our implementations, Ω is discretized into a grid similar to previous continuous max-flow approaches [14, 145, 192, 193].

The memory requirements for both algorithms are dominated by the space required buffers storing the primal/dual optimization variables, intermediate optimization variables, and input data and regularization terms. Each of these grows linearly with the size of the image in our discretization approach. Given a model with N_L leaf-labels and N_B branch-labels (labels that are not the source S or a leaf-label), the number of buffers required by the augmented Lagrangian algorithm (Algorithm 5.1) is $5N_L + 6N_B + 2$ buffers for 2D images and $6N_L + 7N_B + 2$ buffers for 3D volumes. The proximal Bregman (Algorithm 5.4) requires fewer buffers, specifically $4N_L + 3N_B + 1$ buffers for 2D images and $5N_L + 4N_B + 1$ for 3D volumes. The proximal Bregman approach therefore uses between 25% to 50% memory required for the augmented Lagrangian algorithm for segmentation problems with a large number of labels, depending on how many are leaf-labels versus branch-labels. This decrease can play a large role in the feasibility of large segmentation problems, especially when implemented using the limited memory space available to the GPU.

A.5 p-Norm Regularization Terms

The purpose of this section is to illustrate the duality relationship between particular spatial flow constraints and the corresponding norms placed on the labelling function.

Consider the following spatial flow maximization problem:

$$\max_{q(x)\in C} \int_\Omega u(x) \operatorname{div} q(x) dx \tag{A.68}$$

where $q(x) \in C$ represents some pointwise constraint on the spatial flow. In particular, if C is the constraint $|q|_p \leq 1$, then the flow maximization problem is equivalent to:

$$\int_\Omega |\nabla u(x)|_{\tilde{p}} dx \tag{A.69}$$

where $\tilde{p} = {}^{p}/_{p-1}$. In both cases, $|\cdot|_p$ is the p-norm:

$$|a|_p = \sqrt[p]{\sum_i |a_i|^p} \tag{A.70}$$

thus making $|\cdot|_{\tilde{p}}$ the dual norm.

This can be shown through the process:

$$\begin{aligned}
&\max_{|q(x)|_p \leq 1} \int_\Omega u \operatorname{div} q(x) dx \\
&= \max_{|q(x)|_p \leq 1} \int_\Omega (\nabla \cdot (uq) - \nabla u \cdot q(x))\, dx \\
&= \max_{|q(x)|_p \leq 1} \left(\int_{\delta\Omega} uq \cdot \mathbf{n} dS - \int_\Omega \nabla u \cdot q(x) \right) dx \text{ (Gauss-Ostrogradsky Theorem)} \\
&= \max_{|q(x)|_p \leq 1} - \int_\Omega \nabla u \cdot q(x) dx
\end{aligned} \tag{A.71}$$

Noting that $\nabla u \cdot q(x)$ is linear with respect to $q(x)$, the inequality constraint $|q(x)|_p \leq 1$ can be replaced with the equality constraint $|q(x)|_p = 1$ (and therefore $|q(x)|_p^p = 1$ holds as well) as the maximum and minimum must both occur at the boundary of the set. By introducing Lagrangian multipliers on said constraint:

$$\begin{aligned}
&= \max_{|q(x)|_p^p = 1} - \int_\Omega \nabla u \cdot q(x) dx \\
&= \max_q \min_\lambda - \int_\Omega \left(\nabla u \cdot q(x) dx + \lambda(x)(|q(x)|_p^p - 1) \right)
\end{aligned} \tag{A.72}$$

By differentiating with respect to $q_i(x)$ and setting it to zero:

$$\begin{aligned}
0 &= -\frac{\delta u}{\delta x_i} - \lambda(x) p q_i |q_i|^{p-2} \\
q_i |q_i|^{p-2} &= {}^{\frac{\delta u}{\delta x_i}}/_{\lambda(x) p} \\
q_i^{p-1} &\propto \frac{\delta u}{\delta x_i}
\end{aligned} \tag{A.73}$$

which combined with $|q(x)|_p = 1$ yields the maximum value at:

$$q_i(x) = -\frac{\left(\frac{\delta u}{\delta x_i}\right)^{1/p-1}}{\left(\sum_j \left(\frac{\delta u}{\delta x_i}\right)^{p/p-1}\right)^{1/p}} \tag{A.74}$$

which yields:

$$\begin{aligned}
&\max_{|q(x)|_p^p=1} -\int_\Omega \nabla u(x)\cdot q(x)dx \\
&= \int_\Omega \sum_i \frac{\delta u}{\delta x_i} \frac{\left(\frac{\delta u}{\delta x_i}\right)^{1/p-1}}{\left(\sum_j \left(\frac{\delta u}{\delta x_i}\right)^{p/p-1}\right)^{1/p}} dx \\
&= \int_\Omega \frac{\sum_i \left(\frac{\delta u}{\delta x_i}\right)^{p/p-1}}{\left(\sum_j \left(\frac{\delta u}{\delta x_i}\right)^{p/p-1}\right)^{1/p}} dx \\
&= \int_\Omega \left(\sum_i \left(\frac{\delta u}{\delta x_i}\right)^{p/p-1}\right)^{p-1/p} dx \\
&= \int_\Omega |\nabla u(x)|_{\tilde{p}} dx
\end{aligned} \tag{A.75}$$

where $\tilde{p} = {}^{p}/_{p-1}$. Thus showing the principle of duality holds between spatial flow constraints and p-norm regularization of $\nabla u(x)$ regardless of $p \geq 1$. ∎

A.5.1 Directional Regularization

One particularly important type of regularization is *direction-dependent* regularization. In this case, the directional p-norm $|\mathbf{R}(x)\nabla u(x)|_p$ can be used where $\mathbf{R}(x)$ is an invertible matrix.

The derivation is equivalent until step Eq. (A.72). Making the following substitutions $y(x) = \mathbf{R}(x)\nabla u(x)$, perform the following steps:

$$\begin{aligned}
-\int_\Omega \nabla u(x)\cdot q(x)dx = &-\int_\Omega \left(\mathbf{R}^{-1}(x)y(x)\right)\cdot q(x)dx \\
&-\int_\Omega y^\top(x)\left(\mathbf{R}^{-1}(x)\right)^\top q(x)dx \\
&-\int_\Omega y(x)\cdot q'(x)dx \text{ where } q'(x) = \left(\mathbf{R}^{-1}(x)\right)^\top q(x)
\end{aligned} \tag{A.76}$$

which has the same structure as the previous section, except with different variable names. Thus, if this is maximized subject to the constraint $|q'(x)|_{\tilde{p}} = 1$, the maximal value is $|y(x)|_p$. Therefore, the constraint $\left|\left(\mathbf{R}^{-1}(x)\right)^\top q(x)\right|_{\tilde{p}} = 1$ has the dual function $|\mathbf{R}(x)\nabla u(x)|_p$ allowing it to be used as a regularization term. ∎

It is worth noting that the weighted total variation regularization term $\int_\Omega R(x)|\nabla u(x)|_2 dx$ can be implemented with the dual $|q(x)|_2 \leq R(x)$.

A.5.2 Geodesic Star Convexity

The notion of star convexity falls under that of regularization in that the primal term associated with it is a spatial flow term. As shown by Yuan et al. [196], star convexity on a single label can be implemented through the addition of an unconstrained flow term following a particular direction, that is, the spatial flow can be reformulated as:

$$q_L(x) = q'_L(x) + \lambda_L(x) e_L(x). \tag{A.77}$$

Instead of optimizing $\lambda_L(x)$ numerically, one can find the optimal value by minimizing the norm of the residual (the portion that is magnitude constrained) assuming the constraint from the previous section.

$$\min_{\lambda(x)\geq 0} \left| \mathbf{R}_L^{-\top}(x)(q_L(x) - \lambda_L(x) e_L(x)) \right|_{\tilde{p}}^{\tilde{p}} \tag{A.78}$$

Noting that this is a 1D convex problem (each x is independent), one can solve it by finding the zero-valued derivative with respect to λ_L.

$$\begin{aligned}
0 =& \frac{d \left| \mathbf{R}_L^{-\top}(x)(q_L - \lambda_L e_L) \right|_{\tilde{p}}^{\tilde{p}}}{d\lambda_L} \\
0 =& \left(\tilde{p} \mathbf{R}_L^{-\top}(x) \left(q_L - \lambda_L e_L\right) \circ \left| \mathbf{R}_L^{-\top}(x)(q_L - \lambda_L e_L) \right|^{\tilde{p}-2} \right) \cdot \left(\mathbf{R}_L^{-\top}(x) e_L \right) \\
0 =& \sum_i \left(z_{qi} z_{ei} - \lambda_L z_{ei}^2 \right) |z_{qi} - \lambda_L z_{ei}|^{\tilde{p}-2} \text{ where } z_q = \mathbf{R}_L^{-\top} q_L \text{ and } z_e = \mathbf{R}_L^{-\top} e_L \\
\lambda_L =& \frac{\sum_i z_{qi} z_{ei} |z_{qi} - \lambda_L z_{ei}|^{\tilde{p}-2}}{\sum_i z_{ei}^2 |z_{qi} - \lambda_L z_{ei}|^{\tilde{p}-2}}
\end{aligned} \tag{A.79}$$

which can be analytically solved for $\tilde{p} = 2$ as $\lambda_L = \frac{z_q \cdot z_e}{|z_e|_2^2} = \frac{\left(\mathbf{R}_L^{-\top}(x) q_L\right) \cdot \left(\mathbf{R}_L^{-\top}(x) e_L\right)}{\left|\mathbf{R}_L^{-\top}(x) e_L\right|_2^2}$. For other values of $\tilde{p}$ this is more difficult to compute, so our discussion will be limited to the case $\tilde{p} = 2$. Implementing the non-negative constraint yields:

$$\lambda_L = \max \left\{ 0, \frac{\left(\mathbf{R}_L^{-\top} q_L\right) \cdot \left(\mathbf{R}_L^{-\top} e_L\right)}{\left|\mathbf{R}_L^{-\top} e_L\right|_2^2} \right\} \tag{A.80}$$

which allows λ_L to be optimized analytically and $\lambda_L(x) e_L(x)$ removed from the spatial flow prior to constraining it and added back in immediately afterwards. Thus, any Chambolle iteration

[31] on $q_L(x)$ under the constraint $\left|\mathbf{R}_L^{-\top}(x)q_L(x)\right|_2 \leq 1$ can be replaced by:

$$\begin{aligned} q_L(x) \leftarrow & \lambda_L(x)e_L(x) + \text{Proj}_{\left|\mathbf{R}_L^{-\top}(x)(q_L(x)-\lambda_L(x)e_L(x))\right|_2 \leq 1}\left(q_L(x) - \lambda_L(x)e_L(x) + \tau\nabla B_L(x)\right) \\ \text{where } \lambda_L(x) = & \max\left\{0,\ \frac{\left(\mathbf{R}_L^{-\top}(x)q_L(x)\right)\cdot\left(\mathbf{R}_L^{-\top}(x)e_L(x)\right)}{\left|\mathbf{R}_L^{-\top}(x)e_L(x)\right|_2^2}\right\} \end{aligned} \tag{A.81}$$

where $B_L(x)$ is whatever gradient term originally used in the max-flow problem.

Appendix B

Combinatorial and Complexity Analysis of Label Ordering Structures

This appendix contains the derivations of combinatorial results (such as the number of models given particular constraints) as well as a number of complexity results, namely the NP-hardness of hierarchy selection.

B.1 Combinatorics of Unconstrained Hierarchies

The number of hierarchies that can be formed given a set number of leaf nodes was first explored by Leclerc [95] and later generalized by Murtagh [123]. Specifically, the forms of hierarchies used in segmentation are *labelled, non-ranked, non-binary dendrograms* which satisfy the following recursive enumeration formula:

$$g(n,k) = kg(n-1,k) + (n+k+2)g(n-1,k-1) \tag{B.1}$$

where n is the number of leaf nodes and k is the number of levels in the hierarchy ranging between 1 and $n-1$. An important feature to note about this enumeration is that it is superfactorial. That is:

$$\sum_{k=1}^{n-1} g(n,k) = \Omega(n!) \tag{B.2}$$

However, what is more important in terms of hierarchies is the ability to construct one given knowledge of what labels should be co-regularized, placing constraints on the hierarchies available.

B.2 Grouping Graph

One useful structure in analyzing and describing the combinatorics and complexity of HMF and DAGMF results is the *grouping graph*, G. Each vertex of G correspond to an element of $2^{\mathbb{L}}/\{\mathbb{L}, \emptyset\}$, that is, $v \in 2^{\mathbb{L}}$ and $v \neq \emptyset$, $v \neq \mathbb{L}$. Each v indicates a potential or desired shared regularization term between the labels in the set it corresponds to. Edges in G refer to hierarchical incompatibility relations, that is, if it is impossible to apply both regularization terms simultaneously. Specifically, vertices v_1 and v_2 share an edge if $v_1 \cap v_2 \neq \emptyset$, $v_1 \not\subseteq v_2$, and $v_2 \not\subseteq v_1$.

B.3 NP-Hardness of Maximum Hierarchy Selection

One basic desirable operation is, given a collection of groups which one desires regularization for, can one automatically find the largest hierarchy implementing these groups. To consider this as a decision problem, does a hierarchy exist which contains a regularization term for at least k groups?

Theorem B.3.1 *The decision problem concerning the existence of a hierarchy which contains a regularization term for at least k groups is NP-complete.*

Proof

Consider constructing the grouping graph as described in Section B.2 which can constructed in polynomial time given a list of groups. Any valid hierarchy can be represented as an independent set of this graph and thus the existence of a k-specified-node hierarchy is equivalent to the existence of an independent set of size k and the problem of hierarchy selection can be reduced to that of finding the independent set.

Similarly, consider being given a (without loss of generality) connected graph with vertex set V and edge set E. Assign a unique identifier to each vertex and edge in said graph, a_v and $a_{(v_1,v_2)}$ respectively. Each vertex can be assigned a unique set composed of

$$\{a_v\} \cup \biguplus_{(v,v_2)\in E}. \tag{B.3}$$

Note that this graph with the assigned set at each vertex constitutes a valid grouping graph with $\#V$ groups. The existence of an independent set of size k, thus implies the existence of a k-specified-node hierarchy. Thus, the independent set problem can be reduced in polynomial time to the described decision problem. ∎

Appendix C

Ethics Approvals for Patient Images

Many chapters in this thesis make use of openly available medical imaging data sets such as the Open Access Series of Imaging Studies (OASIS) database [118] and Magnetic Resonance Brain Segmentation (MRBrainS) database [120] in Chapter 3 and the BrainWeb [37] dataset in Chapter 5.

Some chapters, however, used images from ongoing studies at the Robarts Research Institute. Ethics approval notices for those images are provided in this appendix in the order in which their respective imaging dataset appear in this thesis.

C.1 Magnetic Resonance Imaging of Multiple Sclerosis at 7 Tesla

Research Western

Use of Human Participants - Ethics Approval Notice

Principal Investigator: Ravi Menon
Review Number: 18752E
Review Level: Delegated
Approved Local Adult Participants: 85
Approved Local Minor Participants: 0
Protocol Title: Magnetic Resonance Imaging of Multiple Sclerosis at 7 Tesla
Department & Institution: Medical Biophysics, Robarts Research Institute
Sponsor: Canadian Institutes of Health Research

Ethics Approval Date: February 17, 2012 **Expiry Date:** February 28, 2017
Documents Reviewed & Approved & Documents Received for Information:

Document Name	Comments	Version Date
UWO Protocol		
Advertisement	Poster - healthy participants	
Letter of Information & Consent	Participants	2012/02/09
Letter of Information & Consent	Controls	2012/02/09
Advertisement	Poster - patients	

This is to notify you that The University of Western Ontario Research Ethics Board for Health Sciences Research Involving Human Subjects (HSREB) which is organized and operates according to the Tri-Council Policy Statement: Ethical Conduct of Research Involving Humans and the Health Canada/ICH Good Clinical Practice Practices: Consolidated Guidelines; and the applicable laws and regulations of Ontario has reviewed and granted approval to the above referenced revision(s) or amendment(s) on the approval date noted above. The membership of this REB also complies with the membership requirements for REB's as defined in Division 5 of the Food and Drug Regulations.

The ethics approval for this study shall remain valid until the expiry date noted above assuming timely and acceptable responses to the HSREB's periodic requests for surveillance and monitoring information. If you require an updated approval notice prior to that time you must request it using the UWO Updated Approval Request Form.

Members of the HSREB who are named as investigators in research studies, or declare a conflict of interest, do not participate in discussion related to, nor vote on, such studies when they are presented to the HSREB.

The Chair of the HSREB is Dr. Joseph Gilbert. The UWO HSREB is registered with the U.S. Department of Health & Human Services under the IRB registration number IRB 00000940.

[redacted]
Signature

Ethics Officer to Contact for Further Information

Janice Sutherland [redacted]	Grace Kelly [redacted]	Shantel Walcott [redacted]

This is an official document. Please retain the original in your files.

The University *of* Western Ontario
Office of Research Ethics
Support Services Building Room 5150 • London, Ontario • CANADA - N6G 1G9
PH: 519-661-3036 • F: 519-850-2466 • ethics@uwo.ca • www.uwo.ca/research/ethics

C.2 new technologies in the management of post-haemorrhagic hydrocephalus in preterm infants

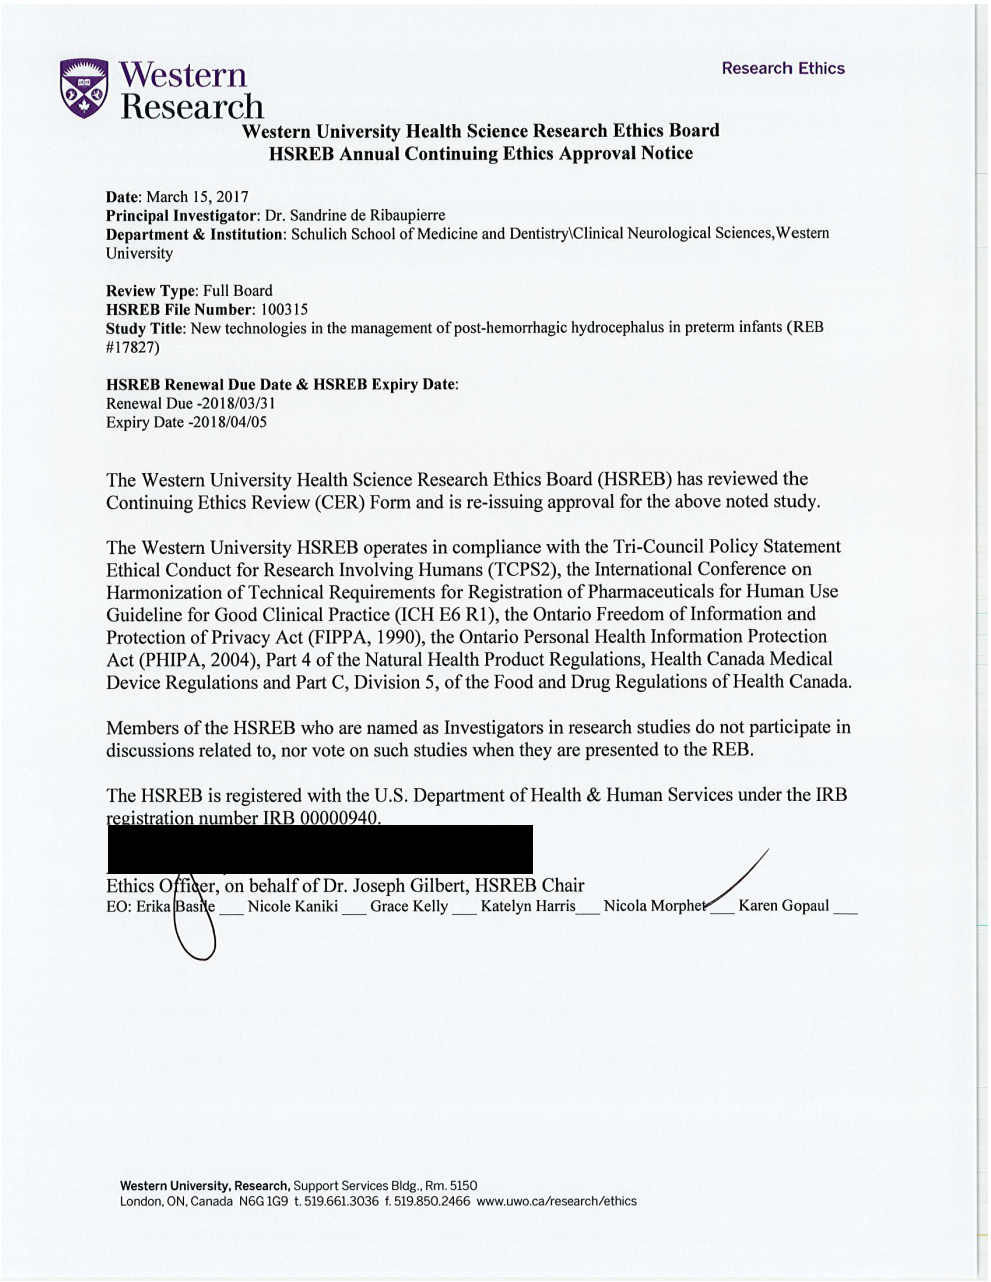

Western Research

Research Ethics

Western University Health Science Research Ethics Board
HSREB Annual Continuing Ethics Approval Notice

Date: March 15, 2017
Principal Investigator: Dr. Sandrine de Ribaupierre
Department & Institution: Schulich School of Medicine and Dentistry\Clinical Neurological Sciences,Western University

Review Type: Full Board
HSREB File Number: 100315
Study Title: New technologies in the management of post-hemorrhagic hydrocephalus in preterm infants (REB #17827)

HSREB Renewal Due Date & HSREB Expiry Date:
Renewal Due -2018/03/31
Expiry Date -2018/04/05

The Western University Health Science Research Ethics Board (HSREB) has reviewed the Continuing Ethics Review (CER) Form and is re-issuing approval for the above noted study.

The Western University HSREB operates in compliance with the Tri-Council Policy Statement Ethical Conduct for Research Involving Humans (TCPS2), the International Conference on Harmonization of Technical Requirements for Registration of Pharmaceuticals for Human Use Guideline for Good Clinical Practice (ICH E6 R1), the Ontario Freedom of Information and Protection of Privacy Act (FIPPA, 1990), the Ontario Personal Health Information Protection Act (PHIPA, 2004), Part 4 of the Natural Health Product Regulations, Health Canada Medical Device Regulations and Part C, Division 5, of the Food and Drug Regulations of Health Canada.

Members of the HSREB who are named as Investigators in research studies do not participate in discussions related to, nor vote on such studies when they are presented to the REB.

The HSREB is registered with the U.S. Department of Health & Human Services under the IRB registration number IRB 00000940.

[redacted]
Ethics Officer, on behalf of Dr. Joseph Gilbert, HSREB Chair
EO: Erika Basile ___ Nicole Kaniki ___ Grace Kelly ___ Katelyn Harris___ Nicola Morphet___ Karen Gopaul ___

Western University, Research, Support Services Bldg., Rm. 5150
London, ON, Canada N6G 1G9 t. 519.661.3036 f. 519.850.2466 www.uwo.ca/research/ethics

C.3 Image-Guidance in Cardiac Interventions

Western Research

Research Ethics

Western University Health Science Research Ethics Board
HSREB Annual Continuing Ethics Approval Notice

Date: July 13, 2015
Principal Investigator: Dr. Terry Peters
Department & Institution: Schulich School of Medicine and Dentistry\Medical Biophysics,Robarts Research Institute

Review Type: Expedited
HSREB File Number: 100617
Study Title: Image-Guidance in Cardiac Interventions
Sponsor: Canadian Institutes of Health Research

HSREB Renewal Due Date & HSREB Expiry Date:
Renewal Due -2016/05/31
Expiry Date -2016/06/16

The Western University Health Science Research Ethics Board (HSREB) has reviewed the Continuing Ethics Review (CER) Form and is re-issuing approval for the above noted study.

The Western University HSREB operates in compliance with the Tri-Council Policy Statement Ethical Conduct for Research Involving Humans (TCPS2), the International Conference on Harmonization of Technical Requirements for Registration of Pharmaceuticals for Human Use Guideline for Good Clinical Practice (ICH E6 R1), the Ontario Freedom of Information and Protection of Privacy Act (FIPPA, 1990), the Ontario Personal Health Information Protection Act (PHIPA, 2004), Part 4 of the Natural Health Product Regulations, Health Canada Medical Device Regulations and Part C, Division 5, of the Food and Drug Regulations of Health Canada.

Members of the HSREB who are named as Investigators in research studies do not participate in discussions related to, nor vote on such studies when they are presented to the REB.

The HSREB is registered with the U.S. Department of Health & Human Services under the IRB registration number IRB 00000940.

[redacted]
Ethics Officer, on behalf of Dr. Joseph Gilbert, HSREB Chair

Ethics Officer to Contact for Further Information

Erika Basile [redacted]	Grace Kelly [redacted]	Mina Mekhail [redacted]	Vikki Tran [redacted]

This is an official document. Please retain the original in your files.

Western University, Research, Support Services Bldg., Rm. 5150
London, ON, Canada N6G 1G9 t. 519.661.3036 f. 519.850.2466 www.uwo.ca/research/ethics

C.4 Anatomical measurements of the heart for radiofrequency cathetar ablation

Western Research

Research Ethics

Use of Human Participants - Initial Ethics Approval Notice

Principal Investigator: Dr. Allan Skanes
File Number:104640
Review Level: Delegated
Protocol Title: Anatomical measurements of the heart for radiofrequency catheter ablation
Department & Institution: Schulich School of Medicine and Dentistry\Medicine-Dept of,London Health Sciences Centre
Sponsor: Canadian Institutes of Health Research

Ethics Approval Date: November 27, 2013 **Expiry Date:**December 01, 2017
Documents Reviewed & Approved & Documents Received for Information:

Document Name	Comments	Version Date
Other	References for section 2.1	2013/11/04
Western University Protocol		

This is to notify you that The University of Western Ontario Research Ethics Board for Health Sciences Research Involving Human Subjects (HSREB) which is organized and operates according to the Tri-Council Policy Statement: Ethical Conduct of Research Involving Humans and the Health Canada/ICH Good Clinical Practice Practices: Consolidated Guidelines; and the applicable laws and regulations of Ontario has reviewed and granted approval to the above referenced revision(s) or amendment(s) on the approval date noted above. The membership of this REB also complies with the membership requirements for REB's as defined in Division 5 of the Food and Drug Regulations.

The ethics approval for this study shall remain valid until the expiry date noted above assuming timely and acceptable responses to the HSREB's periodic requests for surveillance and monitoring information. If you require an updated approval notice prior to that time you must request it using the University of Western Ontario Updated Approval Request Form.

Members of the HSREB who are named as investigators in research studies, or declare a conflict of interest, do not participate in discussion related to, nor vote on, such studies when they are presented to the HSREB.

The Chair of the HSREB is Dr. Joseph Gilbert. The HSREB is registered with the U.S. Department of Health & Human Services under the IRB registration number IRB 00000940.

[redacted]

Signature [redacted]

Ethics Officer to Contact for Further Information

Erika Basile [redacted]	Grace Kelly [redacted]	Mina Mekhail [redacted]	Vikki Tran [redacted]

This is an official document. Please retain the original in your files.

Appendix D

Copyright Transfers and Reprint Permissions

The reprint permissions for the following accepted articles are provided:

- Rajchl, Martin, John SH Baxter, A. Jonathan McLeod, Jing Yuan, Wu Qiu, Terry M. Peters, and Ali R. Khan. "Hierarchical max-flow segmentation framework for multi-atlas segmentation with Kohonen self-organizing map based Gaussian mixture modelling." Medical Image Analysis 27 (2016): 45-56.

- John SH Baxter, Martin Rajchl, Terry M. Peters and Elvis C.S. Chen. "Optimization-based interactive segmentation interface for multi-region problems." SPIE Journal of Medical Imaging (2016).

- John SH Baxter, Jiro Inoue, Maria Drangova and Terry M. Peters. "Shape Complexes: The Intersection of Label Orderings and Star Convexity Constraints in Continuous Max-Flow Medical Image Segmentation." SPIE Journal of Medical Imaging (2016).

- John SH Baxter, Rajchl, Martin, A. Jonathan McLeod, Jing Yuan, and Terry M. Peters. "Directed Acyclic Graph Continuous Max-Flow Image Segmentation for Unconstrained Label Orderings." International Journal of Computer Vision (IJCV), (2017).

All other copyright material is under sole ownership by the author, including arXiv pre-prints, and articles currently under submission or revision.

ELSEVIER LICENSE
TERMS AND CONDITIONS

Jun 06, 2017

This Agreement between Robarts Research Institute -- John Baxter ("You") and Elsevier ("Elsevier") consists of your license details and the terms and conditions provided by Elsevier and Copyright Clearance Center.

License Number	4123100283801
License date	Jun 06, 2017
Licensed Content Publisher	Elsevier
Licensed Content Publication	Medical Image Analysis
Licensed Content Title	Hierarchical max-flow segmentation framework for multi-atlas segmentation with Kohonen self-organizing map based Gaussian mixture modeling
Licensed Content Author	Martin Rajchl,John S.H. Baxter,A. Jonathan McLeod,Jing Yuan,Wu Qiu,Terry M. Peters,Ali R. Khan
Licensed Content Date	Jan 1, 2016
Licensed Content Volume	27
Licensed Content Issue	n/a
Licensed Content Pages	12
Start Page	45
End Page	56
Type of Use	reuse in a thesis/dissertation
Intended publisher of new work	other
Portion	full article
Format	electronic
Are you the author of this Elsevier article?	Yes
Will you be translating?	No
Order reference number	
Title of your thesis/dissertation	Contibutions of Continuous Max-Flow Theory to Medical Image Processing
Expected completion date	Jun 2017
Estimated size (number of pages)	220
Elsevier VAT number	GB 494 6272 12
Requestor Location	Robarts Research Institute 1151 Richmond Street North London, ON N6A 5B7 Canada Attn: John Baxter
Total	0.00 USD

INTRODUCTION

1. The publisher for this copyrighted material is Elsevier. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your Rightslink account and that are available at any time at http://myaccount.copyright.com).

GENERAL TERMS

2. Elsevier hereby grants you permission to reproduce the aforementioned material subject to the terms and conditions indicated.
3. Acknowledgement: If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source, permission must also be sought from that source. If such permission is not obtained then that material may not be included in your publication/copies. Suitable acknowledgement to the source must be made, either as a footnote or in a reference list at the end of your publication, as follows:
"Reprinted from Publication title, Vol /edition number, Author(s), Title of article / title of chapter, Pages No., Copyright (Year), with permission from Elsevier [OR APPLICABLE SOCIETY COPYRIGHT OWNER]." Also Lancet special credit - "Reprinted from The Lancet, Vol. number, Author(s), Title of article, Pages No., Copyright (Year), with permission from Elsevier."
4. Reproduction of this material is confined to the purpose and/or media for which permission is hereby given.
5. Altering/Modifying Material: Not Permitted. However figures and illustrations may be altered/adapted minimally to serve your work. Any other abbreviations, additions, deletions and/or any other alterations shall be made only with prior written authorization of Elsevier Ltd. (Please contact Elsevier at permissions@elsevier.com). No modifications can be made to any Lancet figures/tables and they must be reproduced in full.
6. If the permission fee for the requested use of our material is waived in this instance, please be advised that your future requests for Elsevier materials may attract a fee.
7. Reservation of Rights: Publisher reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.
8. License Contingent Upon Payment: While you may exercise the rights licensed immediately upon issuance of the license at the end of the licensing process for the transaction, provided that you have disclosed complete and accurate details of your proposed use, no license is finally effective unless and until full payment is received from you (either by publisher or by CCC) as provided in CCC's Billing and Payment terms and conditions. If full payment is not received on a timely basis, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and publisher reserves the right to take any and all action to protect its copyright in the materials.
9. Warranties: Publisher makes no representations or warranties with respect to the licensed material.
10. Indemnity: You hereby indemnify and agree to hold harmless publisher and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

11. No Transfer of License: This license is personal to you and may not be sublicensed, assigned, or transferred by you to any other person without publisher's written permission.
12. No Amendment Except in Writing: This license may not be amended except in a writing signed by both parties (or, in the case of publisher, by CCC on publisher's behalf).
13. Objection to Contrary Terms: Publisher hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and publisher (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.
14. Revocation: Elsevier or Copyright Clearance Center may deny the permissions described in this License at their sole discretion, for any reason or no reason, with a full refund payable to you. Notice of such denial will be made using the contact information provided by you. Failure to receive such notice will not alter or invalidate the denial. In no event will Elsevier or Copyright Clearance Center be responsible or liable for any costs, expenses or damage incurred by you as a result of a denial of your permission request, other than a refund of the amount(s) paid by you to Elsevier and/or Copyright Clearance Center for denied permissions.

LIMITED LICENSE

The following terms and conditions apply only to specific license types:
15. **Translation**: This permission is granted for non-exclusive world **English** rights only unless your license was granted for translation rights. If you licensed translation rights you may only translate this content into the languages you requested. A professional translator must perform all translations and reproduce the content word for word preserving the integrity of the article.
16. **Posting licensed content on any Website**: The following terms and conditions apply as follows: Licensing material from an Elsevier journal: All content posted to the web site must maintain the copyright information line on the bottom of each image; A hyper-text must be included to the Homepage of the journal from which you are licensing at http://www.sciencedirect.com/science/journal/xxxxx or the Elsevier homepage for books at http://www.elsevier.com; Central Storage: This license does not include permission for a scanned version of the material to be stored in a central repository such as that provided by Heron/XanEdu.
Licensing material from an Elsevier book: A hyper-text link must be included to the Elsevier homepage at http://www.elsevier.com . All content posted to the web site must maintain the copyright information line on the bottom of each image.

Posting licensed content on Electronic reserve: In addition to the above the following clauses are applicable: The web site must be password-protected and made available only to bona fide students registered on a relevant course. This permission is granted for 1 year only. You may obtain a new license for future website posting.
17. **For journal authors:** the following clauses are applicable in addition to the above:
Preprints:
A preprint is an author's own write-up of research results and analysis, it has not been peer-reviewed, nor has it had any other value added to it by a publisher (such as formatting, copyright, technical enhancement etc.).
Authors can share their preprints anywhere at any time. Preprints should not be added to or enhanced in any way in order to appear more like, or to substitute for, the final versions of articles however authors can update their preprints on arXiv or RePEc with their Accepted Author Manuscript (see below).

If accepted for publication, we encourage authors to link from the preprint to their formal publication via its DOI. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help users to find, access, cite and use the best available version. Please note that Cell Press, The Lancet and some society-owned have different preprint policies. Information on these policies is available on the journal homepage.
Accepted Author Manuscripts: An accepted author manuscript is the manuscript of an article that has been accepted for publication and which typically includes author-incorporated changes suggested during submission, peer review and editor-author communications.
Authors can share their accepted author manuscript:

- immediately
    - via their non-commercial person homepage or blog
    - by updating a preprint in arXiv or RePEc with the accepted manuscript
    - via their research institute or institutional repository for internal institutional uses or as part of an invitation-only research collaboration work-group
    - directly by providing copies to their students or to research collaborators for their personal use
    - for private scholarly sharing as part of an invitation-only work group on commercial sites with which Elsevier has an agreement
- After the embargo period
    - via non-commercial hosting platforms such as their institutional repository
    - via commercial sites with which Elsevier has an agreement

In all cases accepted manuscripts should:

- link to the formal publication via its DOI
- bear a CC-BY-NC-ND license - this is easy to do
- if aggregated with other manuscripts, for example in a repository or other site, be shared in alignment with our hosting policy not be added to or enhanced in any way to appear more like, or to substitute for, the published journal article.

Published journal article (JPA): A published journal article (PJA) is the definitive final record of published research that appears or will appear in the journal and embodies all value-adding publishing activities including peer review co-ordination, copy-editing, formatting, (if relevant) pagination and online enrichment.
Policies for sharing publishing journal articles differ for subscription and gold open access articles:
Subscription Articles: If you are an author, please share a link to your article rather than the full-text. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help your users to find, access, cite, and use the best available version.
Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.
If you are affiliated with a library that subscribes to ScienceDirect you have additional private sharing rights for others' research accessed under that agreement. This includes use for classroom teaching and internal training at the institution (including use in course packs and courseware programs), and inclusion of the article for grant funding purposes.
Gold Open Access Articles: May be shared according to the author-selected end-user license and should contain a CrossMark logo, the end user license, and a DOI link to the formal publication on ScienceDirect.
Please refer to Elsevier's posting policy for further information.
18. **For book authors** the following clauses are applicable in addition to the above:
Authors are permitted to place a brief summary of their work online only. You are not

allowed to download and post the published electronic version of your chapter, nor may you scan the printed edition to create an electronic version. **Posting to a repository:** Authors are permitted to post a summary of their chapter only in their institution's repository.
19. **Thesis/Dissertation**: If your license is for use in a thesis/dissertation your thesis may be submitted to your institution in either print or electronic form. Should your thesis be published commercially, please reapply for permission. These requirements include permission for the Library and Archives of Canada to supply single copies, on demand, of the complete thesis and include permission for Proquest/UMI to supply single copies, on demand, of the complete thesis. Should your thesis be published commercially, please reapply for permission. Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

Elsevier Open Access Terms and Conditions

You can publish open access with Elsevier in hundreds of open access journals or in nearly 2000 established subscription journals that support open access publishing. Permitted third party re-use of these open access articles is defined by the author's choice of Creative Commons user license. See our open access license policy for more information.
Terms & Conditions applicable to all Open Access articles published with Elsevier:
Any reuse of the article must not represent the author as endorsing the adaptation of the article nor should the article be modified in such a way as to damage the author's honour or reputation. If any changes have been made, such changes must be clearly indicated.
The author(s) must be appropriately credited and we ask that you include the end user license and a DOI link to the formal publication on ScienceDirect.
If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source it is the responsibility of the user to ensure their reuse complies with the terms and conditions determined by the rights holder.
Additional Terms & Conditions applicable to each Creative Commons user license:
CC BY: The CC-BY license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article and to make commercial use of the Article (including reuse and/or resale of the Article by commercial entities), provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by/4.0.
CC BY NC SA: The CC BY-NC-SA license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article, provided this is not done for commercial purposes, and that the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. Further, any new works must be made available on the same conditions. The full details of the license are available at http://creativecommons.org/licenses/by-nc-sa/4.0.
CC BY NC ND: The CC BY-NC-ND license allows users to copy and distribute the Article, provided this is not done for commercial purposes and further does not permit distribution of the Article if it is changed or edited in any way, and provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, and that the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by-nc-nd/4.0.
Any commercial reuse of Open Access articles published with a CC BY NC SA or CC BY NC ND license requires permission from Elsevier and will be subject to a fee.
Commercial reuse includes:

- Associating advertising with the full text of the Article
- Charging fees for document delivery or access

- Article aggregation
- Systematic distribution via e-mail lists or share buttons

Posting or linking by commercial companies for use by customers of those companies.

20. **Other Conditions**:

v1.9

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

TRANSFER OF COPYRIGHT TO SOCIETY OF PHOTO-OPTICAL INSTRUMENTATION ENGINEERS (SPIE)

Title of Paper	Optimization-based interactive segmentation interface for multiregion problems	**Journal**	*Journal of Medical Imaging*
Author(s)	John Baxter, Martin Rajchl, Terry Peters, and Elvis Chen	**Paper No.**	16025PR

This signed statement must be returned to SPIE prior to the scheduled publication of the Journal in which the Paper will be published. The intent of this Agreement is to protect the interests of both SPIE and authors/employers and to specify reasonable rights for both parties related to publication and reuse of the material.

The undersigned hereby assign(s) to Society of Photo-Optical Instrumentation Engineers (SPIE) copyright ownership in the above Paper, effective if and when the Paper is accepted for publication by SPIE and to the extent transferable under applicable national law. This assignment gives SPIE the right to register copyright to the Paper in its name as claimant and to publish the Paper in any print or electronic medium.

Authors, or their employers in the case of works made for hire, retain the following rights:

1. All proprietary rights other than copyright, including patent rights.
2. The right to make and distribute copies of the Paper for internal purposes.
3. The right to use the material for lecture or classroom purposes.
4. The right to prepare derivative publications based on the Paper, including books or book chapters, journal papers, and magazine articles, provided that publication of a derivative work occurs subsequent to the official date of publication by SPIE.
5. The right to post an author-prepared version or an official version (preferred version) of the published paper on an internal or external server controlled exclusively by the author/employer, provided that (a) such posting in noncommercial in nature and the paper is made available to users without charge; (b) a copyright notice and full citation appear with the paper, and (c) a link to SPIE's official online version of the abstract is provided using the DOI (Document Object Identifier) link.
   Citation format:
   Author(s), "Paper Title," Publication Title, Editors, Volume (Issue) Number, Article (or Page) Number, (Year).

   Copyright notice format:
   Copyright XXXX (year) Society of Photo-Optical Instrumentation Engineers. One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

   DOI abstract link format:
   http://dx.doi.org/DOI# *(Note: The DOI can be found on the title page or online abstract page of any SPIE article.)*

If the work that forms the basis of this Paper was done under a contract with a governmental agency or other entity that retains certain rights, this Transfer of Copyright is subject to any rights that such governmental agency or other entity may have acquired.

=By signing this Agreement, the authors warrant that (1) the Paper is original and has not previously been published elsewhere; (2) this work does not infringe on any copyright or other rights in any other work; (3) all necessary reproduction permissions, licenses, and clearances have been obtained; and (4) the authors own the copyright in the Paper, are authorized to transfer it, and have full power to enter into this Agreement with SPIE.

WHO SHOULD SIGN. This form must be signed by (1) at least one author who is not a U.S. Government employee and (2) the author's employer if the Paper was prepared within the scope of the author's employment or was commissioned by the employer. If not signed by all authors, the author(s) signing this Agreement represents that he/she is signing this Agreement as authorized agent for and on behalf of all the authors.

Print Name: John S.H. Baxter Date: June 10, 2016

U.S. GOVERNMENT EMPLOYMENT CERTIFICATION

A work prepared by a U.S. Government employee as part of his or her official duties is not eligible for U.S. Copyright. If all authors were U.S. Government employees when this Paper was prepared, and the authors prepared this Paper as part of their official duties, at least one author should sign below. If at least one author was not a U.S. Government employee, the work is eligible for copyright and that author should sign the Transfer of Copyright form above.

Print Name: John S.H. Baxter Date: June 10, 2016

Director of Publications, SPIE, P.O. Box 10, Bellingham, WA 98227-0010 USA
Phone: 360/676-3290 (Pacific Time) • Fax: 360/647 - 1445 • E-mail: journals@spie.org • *www.spie.org*
(revision 2 Feb 2009)

local_p_id: 328986
time: 1465580452
ip address: 67.204.217.179

TRANSFER OF COPYRIGHT TO SOCIETY OF PHOTO-OPTICAL INSTRUMENTATION ENGINEERS (SPIE)

Title of Paper	Shape complexes: the intersection of label orderings and star convexity constraints in continuous max-flow medical image segmentation	**Journal**	*Journal of Medical Imaging*
Author(s)	John Baxter, Jiro Inoue, Maria Drangova, and Terry Peters	**Paper No.**	16199PR

This agreement must be signed and returned to SPIE prior to publication of the Paper referenced above in an SPIE Journal. The intent of this Agreement is to protect the interests of both SPIE and authors/employers and to specify reasonable rights for both parties related to publication and reuse of the material.

The undersigned hereby assigns to Society of Photo-Optical Instrumentation Engineers (SPIE) copyright ownership in the above Paper, effective if and when the Paper is accepted for publication by SPIE and to the extent transferable under applicable national law. This assignment gives SPIE the right to register copyright to the Paper in its name as claimant and to publish the Paper in any print or electronic medium.

Authors, or their employers in the case of works made for hire, retain the following rights:

1. All proprietary rights other than copyright, including patent rights.
2. The right to make and distribute copies of the Paper for internal purposes.
3. The right to use the material for lecture or classroom purposes.
4. The right to prepare derivative publications based on the Paper, including books or book chapters, journal papers, and magazine articles, provided that publication of a derivative work occurs subsequent to the official date of publication by SPIE.
5. The right to post an author-prepared version or an official version (preferred version) of the published paper on an internal or external server controlled exclusively by the author/employer, provided that (a) such posting in noncommercial in nature and the paper is made available to users without charge; (b) a copyright notice and full citation appear with the paper, including the DOI.
   Citation format:
   Author(s), "Paper Title," Publication Title, Editors, Volume (Issue) Number, Article (or Page) Number, (Year); DOI.

   Copyright notice format:
   Copyright XXXX (year) Society of Photo-Optical Instrumentation Engineers. One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

If the work that forms the basis of this Paper was done under a contract with a governmental agency or other entity that retains certain rights, this Transfer of Copyright is subject to any rights that such governmental agency or other entity may have acquired.

By signing this Agreement, the authors warrant that (1) the Paper is original and has not previously been published elsewhere; (2) the Paper does not infringe on any copyright or other rights in any other work; (3) all necessary reproduction permissions, licenses, and clearances have been obtained; and (4) the authors own the copyright in the Paper, are authorized to transfer it, and have full power to enter into this Agreement with SPIE.

WHO SHOULD SIGN. This form must be signed by (1) one author who is not a U.S. Government employee or the author's employer if the Paper was prepared within the scope of the author's employment or was commissioned by the employer. The party signing this Agreement represents that he/she is signing this Agreement as authorized agent for and on behalf of all the authors.

Print Name: John S H Baxter Date: December 2, 2016

U.S. GOVERNMENT EMPLOYMENT CERTIFICATION

A work prepared by a U.S. Government employee as part of his or her official duties is not eligible for U.S. Copyright. If all authors were U.S. Government employees when this Paper was prepared, and the authors prepared this Paper as part of their official duties, at least one author should sign below. If at least one author was not a U.S. Government employee, the work is eligible for copyright and that author should sign the Transfer of Copyright form above.

Print Name: John S H Baxter Date: December 2, 2016

Journals Department, SPIE, P.O. Box 10, Bellingham, WA 98227-0010 USA
Phone: 360/676-3290 (Pacific Time) • Fax: 360/647-1445 • E-mail: journals@spie.org • *www.spie.org*
(revision 7 Oct 2016)

local_p_id: 328986
time: 1480701460
ip address: 129.100.44.77

SPRINGER LICENSE
TERMS AND CONDITIONS

Jun 06, 2017

This Agreement between Robarts Research Institute -- John Baxter ("You") and Springer ("Springer") consists of your license details and the terms and conditions provided by Springer and Copyright Clearance Center.

License Number	4123091298763
License date	Jun 06, 2017
Licensed Content Publisher	Springer
Licensed Content Publication	International Journal of Computer Vision
Licensed Content Title	Directed Acyclic Graph Continuous Max-Flow Image Segmentation for Unconstrained Label Orderings
Licensed Content Author	John S. H. Baxter
Licensed Content Date	Jan 1, 2017
Type of Use	Thesis/Dissertation
Portion	Full text
Number of copies	1
Author of this Springer article	Yes and you are the sole author of the new work
Order reference number	
Title of your thesis / dissertation	Contibutions of Continuous Max-Flow Theory to Medical Image Processing
Expected completion date	Jun 2017
Estimated size(pages)	220
Requestor Location	Robarts Research Institute 1151 Richmond Street North London, ON N6A 5B7 Canada Attn: John Baxter
Billing Type	Invoice
Billing Address	Robarts Research Institute 1151 Richmond Street North London, ON N6A 5B7 Canada Attn: John Baxter
Total	0.00 CAD
Terms and Conditions	

Introduction

The publisher for this copyrighted material is Springer. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions

established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your Rightslink account and that are available at any time at http://myaccount.copyright.com).
Limited License
With reference to your request to reuse material on which Springer controls the copyright, permission is granted for the use indicated in your enquiry under the following conditions:
- Licenses are for one-time use only with a maximum distribution equal to the number stated in your request.
- Springer material represents original material which does not carry references to other sources. If the material in question appears with a credit to another source, this permission is not valid and authorization has to be obtained from the original copyright holder.
- This permission
• is non-exclusive
• is only valid if no personal rights, trademarks, or competitive products are infringed.
• explicitly excludes the right for derivatives.
- Springer does not supply original artwork or content.
- According to the format which you have selected, the following conditions apply accordingly:
• **Print and Electronic:** This License include use in electronic form provided it is password protected, on intranet, or CD-Rom/DVD or E-book/E-journal. It may not be republished in electronic open access.
• **Print:** This License excludes use in electronic form.
• **Electronic:** This License only pertains to use in electronic form provided it is password protected, on intranet, or CD-Rom/DVD or E-book/E-journal. It may not be republished in electronic open access.
For any electronic use not mentioned, please contact Springer at permissions.springer@spi-global.com.
- Although Springer controls the copyright to the material and is entitled to negotiate on rights, this license is only valid subject to courtesy information to the author (address is given in the article/chapter).
- If you are an STM Signatory or your work will be published by an STM Signatory and you are requesting to reuse figures/tables/illustrations or single text extracts, permission is granted according to STM Permissions Guidelines: http://www.stm-assoc.org/permissions-guidelines/
For any electronic use not mentioned in the Guidelines, please contact Springer at permissions.springer@spi-global.com. If you request to reuse more content than stipulated in the STM Permissions Guidelines, you will be charged a permission fee for the excess content.
Permission is valid upon payment of the fee as indicated in the licensing process. If permission is granted free of charge on this occasion, that does not prejudice any rights we might have to charge for reproduction of our copyrighted material in the future.
-If your request is for reuse in a Thesis, permission is granted free of charge under the following conditions:
This license is valid for one-time use only for the purpose of defending your thesis and with a maximum of 100 extra copies in paper. If the thesis is going to be published, permission needs to be reobtained.
- includes use in an electronic form, provided it is an author-created version of the thesis on his/her own website and his/her university's repository, including UMI (according to the definition on the Sherpa website: http://www.sherpa.ac.uk/romeo/);
- is subject to courtesy information to the co-author or corresponding author.
Geographic Rights: Scope
Licenses may be exercised anywhere in the world.
Altering/Modifying Material: Not Permitted
Figures, tables, and illustrations may be altered minimally to serve your work. You may not

alter or modify text in any manner. Abbreviations, additions, deletions and/or any other alterations shall be made only with prior written authorization of the author(s).

Reservation of Rights

Springer reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction and (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

License Contingent on Payment

While you may exercise the rights licensed immediately upon issuance of the license at the end of the licensing process for the transaction, provided that you have disclosed complete and accurate details of your proposed use, no license is finally effective unless and until full payment is received from you (either by Springer or by CCC) as provided in CCC's Billing and Payment terms and conditions. If full payment is not received by the date due, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and Springer reserves the right to take any and all action to protect its copyright in the materials.

Copyright Notice: Disclaimer

You must include the following copyright and permission notice in connection with any reproduction of the licensed material:

"Springer book/journal title, chapter/article title, volume, year of publication, page, name(s) of author(s), (original copyright notice as given in the publication in which the material was originally published) "With permission of Springer"

In case of use of a graph or illustration, the caption of the graph or illustration must be included, as it is indicated in the original publication.

Warranties: None

Springer makes no representations or warranties with respect to the licensed material and adopts on its own behalf the limitations and disclaimers established by CCC on its behalf in its Billing and Payment terms and conditions for this licensing transaction.

Indemnity

You hereby indemnify and agree to hold harmless Springer and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

No Transfer of License

This license is personal to you and may not be sublicensed, assigned, or transferred by you without Springer's written permission.

No Amendment Except in Writing

This license may not be amended except in a writing signed by both parties (or, in the case of Springer, by CCC on Springer's behalf).

Objection to Contrary Terms

Springer hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and Springer (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.

Jurisdiction

All disputes that may arise in connection with this present License, or the breach thereof,

shall be settled exclusively by arbitration, to be held in the Federal Republic of Germany, in accordance with German law.

Other conditions:

V 12AUG2015

Bibliography

[1] http://www.advancedsegmentationtools.org/.

[2] Rolf Adams and Leanne Bischof. Seeded region growing. *IEEE Transactions on pattern analysis and machine intelligence*, 16(6):641–647, 1994.

[3] A Alansary, A Soliman, F Khalifa, A Elnakib, M Mostapha, M Nitzken, M Casanova, and A El-Baz. Map-based framework for segmentation of mr brain images based on visual appearance and prior shape. *Proceedings of the MICCAI Grand Challenge on MR Brain Image Segmentation (MRBrainS13)*, 2013.

[4] Asem M Ali and Aly A Farag. Automatic lung segmentation of volumetric low-dose ct scans using graph cuts. In *International Symposium on Visual Computing*, pages 258–267. Springer, 2008.

[5] Asem M Ali, Aly A Farag, and Ayman S El-Baz. Graph cuts framework for kidney segmentation with prior shape constraints. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 384–392. Springer, 2007.

[6] Li An, Qing-San Xiang, and Sofia Chavez. A fast implementation of the minimum spanning tree method for phase unwrapping. *IEEE transactions on medical imaging*, 19(8):805–808, 2000.

[7] Ben Appleton and Hugues Talbot. Globally optimal surfaces by continuous maximal flows. In *Digital image computing: techniques and applications*, pages 987–996, 2003.

[8] Ben Appleton and Hugues Talbot. Globally minimal surfaces by continuous maximal flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(1):106–118, 2006.

[9] Xabier Artaechevarria, Arrate Munoz-Barrutia, and Carlos Ortiz-de Solórzano. Combination strategies in multi-atlas image segmentation: application to brain mr data. *IEEE Transactions on Medical Imaging (TMI)*, 28(8):1266–1277, 2009.

[10] Andrew J Asman and Bennett A Landman. Non-local staple: An intensity-driven multi-atlas rater model. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 426–434. Springer, 2012.

[11] Ismail Ben Ayed, Kumaradevan Punithakumar, Shuo Li, Ali Islam, and Jaron Chong. Left ventricle segmentation via graph cut distribution matching. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 901–909. Springer, 2009.

[12] Ismail Ben Ayed, Kumaradevan Punithakumar, Rashid Minhas, Rohit Joshi, and Gregory J Garvin. Vertebral body segmentation in mri via convex relaxation and distribution matching. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 520–527. Springer, 2012.

[13] Egil Bae, Xue-Cheng Tai, and Jing Yuan. Maximizing flows with message-passing: Computing spatially continuous min-cuts. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 15–28. Springer, 2015.

[14] Egil Bae, Jing Yuan, Xue-Cheng Tai, and Yuri Boykov. A fast continuous max-flow approach to non-convex multi-labeling problems. In *Efficient Algorithms for Global Optimization Methods in Computer Vision*, pages 134–154. Springer, 2014.

[15] William A Barrett and Eric N Mortensen. Interactive live-wire boundary extraction. *Medical image analysis*, 1(4):331–341, 1997.

[16] George Bartzokis, Jeffrey L Cummings, Charles H Markham, Panos Z Marmarelis, Leo J Treciokas, Todd A Tishler, Stephen R Marder, and Jim Mintz. Mri evaluation of brain iron in earlier-and later-onset parkinsons disease and normal subjects. *Magnetic resonance imaging*, 17(2):213–222, 1999.

[17] John S H Baxter, Martin Rajchl, Terry M Peters, and Elvis CS Chen. Optimization-based interactive segmentation interface for multi-region problems. In *SPIE Journal of Medical Imaging*, pages 94133T–94133T. International Society for Optics and Photonics, 2015.

[18] John S H Baxter, Martin Rajchl, Jing Yuan, and Terry M Peters. A proximal bregman projection approach to continuous max-flow problems using entropic distances. *arXiv preprint arXiv:1501.07844*, 2015.

[19] John SH Baxter, Martin Rajchl, A Jonathan McLeod, Ali R Khan, Jing Yuan, and Terry M Peters. Smoothness parameter tuning for generalized hierarchical continuous max-flow segmentation. In *SPIE Medical Imaging*, pages 903410–903410. International Society for Optics and Photonics, 2014.

[20] Dimitri P Bertsekas. Nonlinear programming. 1999.

[21] O Beuf, A Briguet, M Lissac, and R Davis. Magnetic resonance imaging for the determination of magnetic susceptibility of materials. *Journal of Magnetic Resonance, Series B*, 112(2):111–118, 1996.

[22] Jos M Bioucas-Dias and Gonalo Valadao. Phase unwrapping via graph cuts. *IEEE Transactions on Image processing*, 16(3):698–709, 2007.

[23] Yuri Boykov and Marie-Pierre Jolly. Interactive organ segmentation using graph cuts. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 276–286. Springer, 2000.

[24] Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 26(9):1124–1137, 2004.

[25] Yuri Boykov, Vladimir Kolmogorov, Daniel Cremers, and Andrew Delong. An integral solution to surface evolution pdes via geo-cuts. In *European Conference on Computer Vision*, pages 409–422. Springer, 2006.

[26] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 23(11):1222–1239, 2001.

[27] Yuri Y Boykov and M-P Jolly. Interactive graph cuts for optimal boundary & region segmentation of objects in nd images. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 1, pages 105–112. IEEE, 2001.

[28] Lev M Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217, 1967.

[29] Sema Candemir, Stefan Jaeger, Kannappan Palaniappan, Sameer Antani, and George Thoma. Graph-cut based automatic lung boundary detection in chest radiographs. In *IEEE Healthcare Technology Conference: Translational Engineering in Health & Medicine*, pages 31–34, 2012.

[30] M Jorge Cardoso, Marc Modat, Sebastien Ourselin, Shiva Keihaninejad, and David Cash. Multi-steps: multi-label similarity and truth estimation for propagated segmentations. In *Mathematical Methods in Biomedical Image Analysis (MMBIA), 2012 IEEE Workshop on*, pages 153–158. IEEE, 2012.

[31] Antonin Chambolle. An algorithm for total variation minimization and applications. *Journal of Mathematical imaging and vision*, 20(1-2):89–97, 2004.

[32] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.

[33] Tony F Chan, Selim Esedoglu, and Mila Nikolova. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM journal on applied mathematics*, 66(5):1632–1648, 2006.

[34] Tony F Chan and Luminita A Vese. Active contours without edges. *IEEE Transactions on image processing*, 10(2):266–277, 2001.

[35] Guillaume Charpiat, Olivier Faugeras, and Renaud Keriven. Approximations of shape metrics and application to shape warping and empirical shape statistics. *Foundations of Computational Mathematics*, 5(1):1–58, 2005.

[36] Sofia Chavez, Qing-San Xiang, and Li An. Understanding phase maps in mri: a new cutline phase unwrapping method. *IEEE transactions on medical imaging*, 21(8):966–977, 2002.

[37] Chris A Cocosco, Vasken Kollokian, Remi K-S Kwan, G Bruce Pike, and Alan C Evans. Brainweb: Online interface to a 3d mri simulated brain database. In *NeuroImage*, 1997.

[38] Timothy F Cootes, Andrew Hill, Christopher J Taylor, and Jane Haslam. Use of active shape models for locating structures in medical images. *Image and Vision Computing*, 12(6):355–365, 1994.

[39] Timothy F Cootes, Christopher J Taylor, David H Cooper, and Jim Graham. Active shape models-their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.

[40] Daniel Cremers, Mikael Rousson, and Rachid Deriche. A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. *International journal of computer vision*, 72(2):195–215, 2007.

[41] R Cusack, JM Huntley, and HT Goldrein. Improved noise-immune phase-unwrapping algorithm. *Applied Optics*, 34(5):781–789, 1995.

[42] Erick Delage, Honglak Lee, and Andrew Y Ng. A dynamic bayesian network model for autonomous 3d reconstruction from a single indoor image. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 2418–2428. IEEE, 2006.

[43] Andrew Delong and Yuri Boykov. Globally optimal segmentation of multi-region objects. In *IEEE International Conference on Computer Vision (ICCV)*, pages 285–292. IEEE, 2009.

[44] Andrew Delong, Lena Gorelick, Olga Veksler, and Yuri Boykov. Minimizing energies with hierarchical costs. *International journal of computer vision*, 100(1):38–58, 2012.

[45] Jose Dolz, Hortense A Kirisli, Romain Viard, and Laurent Massoptier. Interactive approach to segment organs at risk in radiotherapy treatment planning. In *SPIE Medical Imaging*, pages 90343S–90343S. International Society for Optics and Photonics, 2014.

[46] Jack Edmonds and Richard M Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM (JACM)*, 19(2):248–264, 1972.

[47] Jan Egger, Miriam HA Bauer, Daniela Kuhnt, Barbara Carl, Christoph Kappus, Bernd Freisleben, and Christopher Nimsky. Nugget-cut: a segmentation scheme for spherically-and elliptically-shaped 3d objects. In *Joint Pattern Recognition Symposium*, pages 373–382. Springer, 2010.

[48] Jan Egger, Bernd Freisleben, Christopher Nimsky, and Tina Kapur. Template-cut: A pattern-based segmentation paradigm. *Nature Scientific Reports*, 2(420), 2012.

[49] Jan Egger, Tina Kapur, Thomas Dukatz, Malgorzata Kolodziej, Dženan Zukić, Bernd Freisleben, and Christopher Nimsky. Square-cut: a segmentation algorithm on the basis of a rectangle shape. *PloS One*, 7(2):e31064, 2012.

[50] Jan Egger, Tobias Lüddemann, Robert Schwarzenberg, Bernd Freisleben, and Christopher Nimsky. Interactive-cut: Real-time feedback segmentation for translational research. *Computerized Medical Imaging and Graphics*, 38(4):285–295, 2014.

[51] Ivar Ekeland and Roger Temam. *Convex Analysis and Variational Problems*. SIAM, 1976.

[52] Thomas D Erickson. Working with interface metaphors. *Baecker et al*, 11:147, 2000.

[53] Alexandre X Falcão, Jayaram K Udupa, Supun Samarasekera, Shoba Sharma, Bruce Elliot Hirsch, and Roberto de A Lotufo. User-steered image segmentation paradigms: Live wire and live lane. *Graphical models and image processing*, 60(4):233–260, 1998.

[54] Oladimeji Farri, David S Pieckiewicz, Ahmed S Rahman, Terrence J Adam, Serguei V Pakhomov, and Genevieve B Melton. A qualitative analysis of ehr clinical document synthesis by clinicians. In *AMIA Annual Symposium Proceedings*, volume 2012, page 1211. American Medical Informatics Association, 2012.

[55] Wei Feng, Jaladhar Neelavalli, and E Mark Haacke. Catalytic multiecho phase unwrapping scheme (campus) in multiecho gradient echo imaging: Removing phase wraps on a voxel-by-voxel basis. *Magnetic resonance in medicine*, 70(1):117–126, 2013.

[56] Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.

[57] Lester R Ford and Delbert R Fulkerson. Maximal flow through a network. *Canadian journal of Mathematics*, 8(3):399–404, 1956.

[58] Robin E Fredericksen, James M Coggins, Timothy J Cullip, and Stephen M Pizer. Interactive object definition in medical images using multiscale, geometric image descriptions. In *Visualization in Biomedical Computing, 1990., Proceedings of the First Conference on*, pages 108–114. IEEE, 1990.

[59] Daniel Freedman and Tao Zhang. Interactive graph cut based segmentation with shape priors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 755–762. IEEE, 2005.

[60] Chrysoula Gatsou, Anastasios Politis, and Dimitrios Zevgolis. Text vs visual metaphor in mobile interfaces for novice user interaction. *Information Services & Use*, 31(3-4):271–279, 2011.

[61] Andrew V Goldberg and Robert E Tarjan. A new approach to the maximum-flow problem. *Journal of the ACM (JACM)*, 35(4):921–940, 1988.

[62] Richard M Goldstein, Howard A Zebker, and Charles L Werner. Satellite radar interferometry: Two-dimensional phase unwrapping. *Radio science*, 23(4):713–720, 1988.

[63] Leo Grady. Random walks for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 28(11):1768–1783, 2006.

[64] Dorothy M Greig, Bruce T Porteous, and Allan H Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 271–279, 1989.

[65] Varun Gulshan, Carsten Rother, Antonio Criminisi, Andrew Blake, and Andrew Zisserman. Geodesic star convexity for interactive image segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition(CVPR)*, pages 3129–3136. IEEE, 2010.

[66] Fumin Guo, Jing Yuan, Martin Rajchl, Sarah Svenningsen, Dante PI Capaldi, Khadija Sheikh, Aaron Fenster, and Grace Parraga. Globally optimal co-segmentation of three-dimensional pulmonary 1 h and hyperpolarized 3 he mri with spatial consistence prior. *Medical image analysis*, 23(1):43–55, 2015.

[67] E Mark Haacke, Yingbiao Xu, Yu-Chung N Cheng, and Jürgen R Reichenbach. Susceptibility weighted imaging (swi). *Magnetic resonance in medicine*, 52(3):612–618, 2004.

[68] J. M. Hammersley and P. Clifford. Markov fields on finite graphs and lattices. 1971.

[69] Tobias Heimann and Hans-Peter Meinzer. Statistical shape models for 3d medical image segmentation: a review. *Medical Image Analysis*, 13(4):543–563, 2009.

[70] Gabor T. Herman and Bruno M. Carvalho. Multiseeded segmentation using fuzzy connectedness. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(5):460–474, 2001.

[71] Diego Hernando, P Kellman, JP Haldar, and Z-P Liang. Robust water/fat separation in the presence of large field inhomogeneities using a graph cut algorithm. *Magnetic resonance in medicine*, 63(1):79–90, 2010.

[72] Derek Hoiem, Alexei Efros, Martial Hebert, et al. Geometric context from a single image. In *IEEE International Conference on Computer Vision (ICCV)*, volume 1, pages 654–661. IEEE, 2005.

[73] Derek Hoiem, Alexei A Efros, and Martial Hebert. Recovering surface layout from an image. *International Journal of Computer Vision*, 75(1):151–172, 2007.

[74] SA Hojjatoleslami and Josef Kittler. Region growing: a new approach. *IEEE Transactions on Image processing*, 7(7):1079–1084, 1998.

[75] Keith J Holyoak and Paul Thagard. Analogical mapping by constraint satisfaction. *Cognitive science*, 13(3):295–355, 1989.

[76] Robyn Honea, Tim J Crow, Dick Passingham, and Clare E Mackay. Regional deficits in brain volume in schizophrenia: a meta-analysis of voxel-based morphometry studies. *American Journal of Psychiatry*, 162(12):2233–2245, 2005.

[77] Zahra Hosseini, Junmin Liu, Igor Solovey, Ravi S Menon, and Maria Drangova. Susceptibility-weighted imaging using inter-echo-variance channel combination for improved contrast at 7 tesla. *Journal of Magnetic Resonance Imaging*, 2016.

[78] Hiroshi Ishikawa. Exact optimization for markov random fields with convex priors. *IEEE transactions on pattern analysis and machine intelligence*, 25(10):1333–1336, 2003.

[79] Ernst Ising. A contribution to the theory of ferromagnetism. *Z. Phys*, 31(1):253–258, 1925.

[80] Jongseong Jang, Hyung Wook Kim, and Young Soo Kim. Co-segmentation of inter-subject brain magnetic resonance images. In *IEEE International Conference on Ubiquitous Robots and Ambient Intelligence (URAI)*, pages 80–84. IEEE, 2014.

[81] Monique WM Jaspers. A comparison of usability methods for testing interactive health technologies: methodological aspects and empirical evidence. *International journal of medical informatics*, 78(5):340–353, 2009.

[82] Monique WM Jaspers, Thiemo Steen, Cor Van Den Bos, and Maud Geenen. The think aloud method: a guide to user interface design. *International journal of medical informatics*, 73(11):781–795, 2004.

[83] Jianmin Jiang, P Trundle, and Jinchang Ren. Medical image analysis with artificial neural networks. *Computerized Medical Imaging and Graphics*, 34(8):617–631, 2010.

[84] Marie-Pierre Jolly. Automatic segmentation of the left ventricle in cardiac mr and ct images. *International Journal of Computer Vision*, 70(2):151–163, 2006.

[85] John Kammersgaard. Four different perspectives on human–computer interaction. *International Journal of Man-Machine Studies*, 28(4):343–362, 1988.

[86] Ali R Khan, Maged Goubran, Sandrine de Ribaupierre, Robert R Hammond, Jorge G Burneo, Andrew G Parrent, and Terry M Peters. Quantitative relaxometry and diffusion mri for lateralization in mts and non-mts temporal lobe epilepsy. *Epilepsy research*, 108(3):506–516, 2014.

[87] Lisa M Koch, Martin Rajchl, Tong Tong, Jonathan Passerat-Palmbach, Paul Aljabar, and Daniel Rueckert. Multi-atlas segmentation as a graph labelling problem: Application to partially annotated atlas data. In *Information Processing in Medical Imaging*, pages 221–232. Springer, 2015.

[88] Matthias Koenig, Wolf Spindler, Jan Rexilius, Julien Jomier, Florian Link, and Heinz-Otto Peitgen. Embedding vtk and itk into a visual programming and rapid prototyping platform. In *Medical Imaging*, pages 61412O–61412O. International Society for Optics and Photonics, 2006.

[89] Vladimir Kolmogorov and Ramin Zabin. What energy functions can be minimized via graph cuts? *IEEE transactions on pattern analysis and machine intelligence*, 26(2):147–159, 2004.

[90] Konstantinos Konstantinides and John R Rasure. The khoros software development environment for image and signal processing. *IEEE Transactions on Image Processing*, 3(3):243–252, 1994.

[91] Luis Landaeta, Alexandra La Cruz, Alexander Baranya, and María-Esther Vidal. Semvism: Semantic visualizer for medical image. In *Tenth International Symposium on Medical Information Processing and Analysis*, pages 928712–928712. International Society for Optics and Photonics, 2015.

[92] BA Landman and S Warfield. Miccai 2012 workshop on multi-atlas labeling. In *Medical Image Computing and Computer Assisted Intervention Conference 2012: MICCAI 2012 Grand Challenge and Workshop on Multi-Atlas Labeling Challenge Results*, 2012.

[93] Christian Langkammer, Stefan Ropele, Lukas Pirpamer, Franz Fazekas, and Reinhold Schmidt. Mri for iron mapping in alzheimer's disease. *Neurodegenerative Diseases*, 13(2-3):189–191, 2013.

[94] Curtis P Langlotz. Radlex: a new method for indexing online educational materials 1. *Radiographics*, 26(6):1595–1597, 2006.

[95] Bruno Leclerc. La comparaison des hiérarchies: indices et métriques. *Mathématiques et sciences humaines*, 92:5–40, 1985.

[96] Christian Ledig, Robin Wolz, Paul Aljabar, Jyrki Lötjönen, Rolf A Heckemann, Alexander Hammers, and Daniel Rueckert. Multi-class brain segmentation using atlas propagation and em-based refinement. In *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 896–899. IEEE, 2012.

[97] James J Levitt, Laurel Bobrow, Diandra Lucia, and Padmapriya Srinivasan. A selective review of volumetric and morphometric imaging in schizophrenia. In *Behavioral Neurobiology of Schizophrenia and Its Treatment*, pages 243–281. Springer, 2010.

[98] Clayton Lewis. *Using the" thinking-aloud" method in cognitive interface design*. IBM TJ Watson Research Center, 1982.

[99] Lin Li. Magnetic susceptibility quantification for arbitrarily shaped objects in inhomogeneous fields. *Magnetic resonance in medicine*, 46(5):907–916, 2001.

[100] Lin Li and John S Leigh. Quantifying arbitrary magnetic susceptibility distributions with mr. *Magnetic resonance in medicine*, 51(5):1077–1082, 2004.

[101] Stan Z Li. *Markov random field modeling in image analysis*. Springer Science & Business Media, 2009.

[102] Wei Li, Alexandru V Avram, Bing Wu, Xue Xiao, and Chunlei Liu. Integrated laplacian-based phase unwrapping and background phase removal for quantitative susceptibility mapping. *NMR in Biomedicine*, 27(2):219–227, 2014.

[103] Wei Li, Bing Wu, and Chunlei Liu. Quantitative susceptibility mapping of human brain reflects spatial variation in tissue composition. *Neuroimage*, 55(4):1645–1656, 2011.

[104] Xiang Lin, Brett Cowan, and Alistair Young. Model-based graph cut method for segmentation of the left ventricle. In *2005 IEEE Engineering in Medicine and Biology 27th Annual Conference*, pages 3059–3062. IEEE, 2006.

[105] Jeffrey A Linder, Alan F Rose, Matvey B Palchuk, Frank Chang, Jeffrey L Schnipper, Joseph C Chan, and Blackford Middleton. Decision support for acute problems: The role of the standardized patient in usability testing. *Journal of biomedical informatics*, 39(6):648–655, 2006.

[106] Chunlei Liu. Susceptibility tensor imaging. *Magnetic resonance in medicine*, 63(6):1471–1477, 2010.

[107] Junmin Liu and Maria Drangova. Intervention-based multidimensional phase unwrapping using recursive orthogonal referring. *Magnetic resonance in medicine*, 68(4):1303–1316, 2012.

[108] Junmin Liu, David A Rudko, Joseph S Gati, Ravi S Menon, and Maria Drangova. Inter-echo variance as a weighting factor for multi-channel combination in multi-echo acquisition for local frequency shift mapping. *Magnetic resonance in medicine*, 73(4):1654–1661, 2015.

[109] Tian Liu, Krishna Surapaneni, Min Lou, Liuquan Cheng, Pascal Spincemaille, and Yi Wang. Cerebral microbleeds: burden assessment by using quantitative susceptibility mapping. *Radiology*, 262(1):269–278, 2012.

[110] Herve Lombaert, Yiyong Sun, and Farida Cheriet. Landmark-based non-rigid registration via graph cuts. In *International Conference Image Analysis and Recognition*, pages 166–175. Springer, 2007.

[111] Anna Macaranas, Alissa N Antle, and Bernhard E Riecke. Bridging the gap: attribute and spatial metaphors for tangible interface design. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*, pages 161–168. ACM, 2012.

[112] Dwarikanath Mahapatra. Skull stripping of neonatal brain mri: using prior shape information with graph cuts. *Journal of digital imaging*, 25(6):802–814, 2012.

[113] Dwarikanath Mahapatra. Analyzing training information from random forests for improved image segmentation. *IEEE Transactions on Image Processing*, 23(4):1504–1512, 2014.

[114] Dwarikanath Mahapatra and Joachim M Buhmann. Prostate mri segmentation using learned semantic knowledge and graph cuts. *IEEE Transactions on Biomedical Engineering*, 61(3):756–764, 2014.

[115] Dwarikanath Mahapatra and Ying Sun. Orientation histograms as shape priors for left ventricle segmentation using graph cuts. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 420–427. Springer, 2011.

[116] JB Antoine Maintz and Max A Viergever. A survey of medical image registration. *Medical image analysis*, 2(1):1–36, 1998.

[117] James Malcolm, Yogesh Rathi, and Allen Tannenbaum. Graph cut segmentation with nonlinear shape priors. In *IEEE International Conference on Image Processing*, volume 4, pages IV–365. IEEE, 2007.

[118] Daniel S Marcus, Tracy H Wang, Jamie Parker, John G Csernansky, John C Morris, and Randy L Buckner. Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *Journal of cognitive neuroscience*, 19(9):1498–1507, 2007.

[119] Andrea Mechelli, Cathy J Price, Karl J Friston, and John Ashburner. Voxel-based morphometry of the human brain: methods and applications. *Current medical imaging reviews*, 1(2):105–113, 2005.

[120] Adrienne M Mendrik, Koen L Vincken, Hugo J Kuijf, Marcel Breeuwer, Willem H Bouvy, Jeroen De Bresser, Amir Alansary, Marleen De Bruijne, Aaron Carass, Ayman El-Baz, et al. Mrbrains challenge: online evaluation framework for brain image segmentation in 3t mri scans. *Computational intelligence and neuroscience*, 2015, 2015.

[121] Brandon Miles, Ismail Ben Ayed, Max WK Law, Greg Garvin, Aaron Fenster, and Shuo Li. Spine image fusion via graph cuts. *IEEE Transactions on Biomedical Engineering*, 60(7):1841–1850, 2013.

[122] Pim Moeskops, Max A Viergever, Manon JNL Benders, and Ivana Išgum. Evaluation of an automatic brain segmentation method developed for neonates on adult mr brain images. In *SPIE Medical Imaging*, pages 941315–941315. International Society for Optics and Photonics, 2015.

[123] Fionn Murtagh. Counting dendrograms: a survey. *Discrete Applied Mathematics*, 7(2):191–199, 1984.

[124] Cyrus MS Nambakhsh, Jing Yuan, Kumaradevan Punithakumar, Aashish Goela, Martin Rajchl, Terry M Peters, and Ismail Ben Ayed. Left ventricle segmentation in mri via convex relaxed distribution matching. *Medical image analysis*, 17(8):1010–1024, 2013.

[125] Wieslaw Lucjan Nowinski, Guo Liang Yang, and Tseng Tsai Yeo. Computer-aided stereotactic functional neurosurgery enhanced by the use of the multiple brain atlas database. *IEEE transactions on medical imaging*, 19(1):62–69, 2000.

[126] Masahiro Oda, Teruhisa Nakaoka, Takayuki Kitasaka, Kazuhiro Furukawa, Kazunari Misawa, Michitaka Fujiwara, and Kensaku Mori. Organ segmentation from 3d abdominal ct images based on atlas selection and graph cut. In *International MICCAI Workshop on Computational and Clinical Challenges in Abdominal Imaging*, pages 181–188. Springer, 2011.

[127] Ipek Oguz and Milan Sonka. Logismos-b: layered optimal graph image segmentation of multiple objects and surfaces for the brain. *IEEE transactions on medical imaging*, 33(6):1220–1235, 2014.

[128] Toshiyuki Okada, Ryuji Shimada, Masatoshi Hori, Masahiko Nakamoto, Yen-Wei Chen, Hironobu Nakamura, and Yoshinobu Sato. Automated segmentation of the liver from 3d ct images using probabilistic atlas and multilevel statistical shape model. *Academic Radiology*, 15(11):1390–1403, 2008.

[129] Ozan Oktay, Alberto Gomez, Kevin Keraudren, Andreas Schuh, Wenjia Bai, Wenzhe Shi, Graeme Penney, and Daniel Rueckert. Probabilistic edge map (pem) for 3d ultrasound image registration and multi-atlas left ventricle segmentation. In *International Conference on Functional Imaging and Modeling of the Heart*, pages 223–230. Springer, 2015.

[130] Sílvia Delgado Olabarriaga and Arnold WM Smeulders. Interaction in the segmentation of medical images: A survey. *Medical image analysis*, 5(2):127–142, 2001.

[131] Sébastien Ourselin, Radu Stefanescu, and Xavier Pennec. Robust registration of multimodal images: towards real-time clinical applications. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 140–147. Springer, 2002.

[132] James A Overton, Cesare Romagnoli, and Rethy Chhem. Open biomedical ontologies applied to prostate cancer. *Applied Ontology*, 6(1):35–51, 2011.

[133] Charles Sanders Peirce. On a new list of categories. In *Proceedings of the american academy of arts and sciences*, volume 7, pages 287–298, 1868.

[134] Charles Sanders Peirce. Collected papers, vol. 3: Exact logic (c. hartshorne and p. weiss, eds.), 1933.

[135] Jialin Peng, Jinwei Wang, and Dexing Kong. A new convex variational model for liver segmentation. In *Pattern Recognition (ICPR), 2012 21st International Conference on*, pages 3754–3757. IEEE, 2012.

[136] Sérgio Pereira, Joana Festa, José António Mariz, Nuno Sousa, and Carlos A Silva. Automatic brain tissue segmentation of multi-sequence mr images using random decision forests. *Proceedings of the MICCAI Grand Challenge on MR Brain Image Segmentation (MRBrainS13)*, 2013.

[137] Thomas Pock, Antonin Chambolle, Daniel Cremers, and Horst Bischof. A convex relaxation approach for computing minimal partitions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 810–817. IEEE, 2009.

[138] Miranda Poon, Ghassan Hamarneh, and Rafeef Abugharbieh. Efficient interactive 3d livewire segmentation of complex objects with arbitrary topology. *Computerized Medical Imaging and Graphics*, 32(8):639–650, 2008.

[139] Renfrey Burnard Potts. Some generalized order-disorder transformations. In *Mathematical proceedings of the cambridge philosophical society*, volume 48, pages 106–109. Cambridge Univ Press, 1952.

[140] Wu Qiu, Martin Rajchl, Fumin Guo, Yue Sun, Eranga Ukwatta, Aaron Fenster, and Jing Yuan. 3d prostate trus segmentation using globally optimized volume-preserving prior. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 796–803. Springer, 2014.

[141] Wu Qiu, Jing Yuan, Jessica Kishimoto, Jonathan McLeod, Yimin Chen, Sandrine de Ribaupierre, and Aaron Fenster. User-guided segmentation of preterm neonate ventricular system from 3-d ultrasound images using convex optimization. *Ultrasound in medicine & biology*, 41(2):542–556, 2015.

[142] Wu Qiu, Jing Yuan, Eranga Ukwatta, Yue Sun, Martin Rajchl, and Aaron Fenster. Dual optimization based prostate zonal segmentation in 3d mr images. *Medical image analysis*, 18(4):660–673, 2014.

[143] Wu Qiu, Jing Yuan, Eranga Ukwatta, Yue Sun, Martin Rajchl, and Aaron Fenster. Prostate segmentation: An efficient convex optimization approach with axial symmetry using 3-d trus and mr images. *IEEE transactions on medical imaging*, 33(4):947–960, 2014.

[144] Martin Rajchl, John SH Baxter, Egil Bae, Xue-Cheng Tai, Aaron Fenster, Terry M Peters, and Jing Yuan. Variational time-implicit multiphase level-sets. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 278–291. Springer, 2015.

[145] Martin Rajchl, John SH Baxter, A Jonathan McLeod, Jing Yuan, Wu Qiu, Terry M Peters, and Ali R Khan. Hierarchical max-flow segmentation framework for multi-atlas segmentation with kohonen self-organizing map based gaussian mixture modeling. *Medical Image Analysis*, 27:45–56, 2016.

[146] Martin Rajchl, John SH Baxter, Wu Qiu, Ali R Khan, Aaron Fenster, Terry M Peters, Daniel Rueckert, and Jing Yuan. Fast deformable image registration with non-smooth dual optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 25–32, 2016.

[147] Martin Rajchl, John SH Baxter, Wu Qiu, Ali R Khan, Aaron Fenster, Terry M Peters, and Jing Yuan. Rancor: non-linear image registration with total variation regularization. *arXiv preprint arXiv:1404.2571*, 2014.

[148] Martin Rajchl, Matthew CH Lee, Ozan Oktay, Konstantinos Kamnitsas, Jonathan Passerat-Palmbach, Wenjia Bai, Bernhard Kainz, and Daniel Rueckert. Deepcut: Object segmentation from bounding box annotations using convolutional neural networks. *arXiv preprint arXiv:1605.07866*, 2016.

[149] Martin Rajchl, Jing Yuan, Eranga Ukwatta, and PM Peters. Fast interactive multi-region cardiac segmentation with linearly ordered labels. In *2012 9th IEEE International Symposium on Biomedical imaging (ISBI)*, pages 1409–1412. IEEE, 2012.

[150] Martin Rajchl, Jing Yuan, James A White, Eranga Ukwatta, John Stirrat, Cyrus MS Nambakhsh, Feng P Li, and Terry M Peters. Interactive hierarchical-flow segmentation of scar tissue from late-enhancement cardiac mr images. *IEEE transactions on medical imaging*, 33(1):159–172, 2014.

[151] Felix Ritter, Tobias Boskamp, André Homeyer, Hendrik Laue, Michael Schwier, Florian Link, and H-O Peitgen. Medical image analysis. *IEEE pulse*, 2(6):60–70, 2011.

[152] Torsten Rohlfing and Calvin R Maurer. Shape-based averaging. *IEEE Transactions on Image Processing (TIP)*, 16(1):153–161, 2007.

[153] Daniel L Rubin. Creating and curating a terminology for radiology: ontology modeling and analysis. *Journal of digital imaging*, 21(4):355–362, 2008.

[154] Ashish K Rudra, Ananda S Chowdhury, Ahmed Elnakib, Fahmi Khalifa, Ahmed Soliman, Garth Beache, and Ayman El-Baz. Kidney segmentation using graph cuts and pixel connectivity. *Pattern Recognition Letters*, 34(13):1470–1475, 2013.

[155] Suresh A Sadananthan, Weili Zheng, Michael WL Chee, and Vitali Zagorodnov. Skull stripping using graph cuts. *NeuroImage*, 49(1):225–239, 2010.

[156] María F Salfity, Pablo D Ruiz, Jonathan M Huntley, Martin J Graves, Rhodri Cusack, and Daniel A Beauregard. Branch cut surface placement for unwrapping of undersampled three-dimensional phase data: application to magnetic resonance imaging arterial flow mapping. *Applied optics*, 45(12):2711–2722, 2006.

[157] Rares Salomir, Baudouin Denis de Senneville, and Chrit TW Moonen. A fast calculation method for magnetic field inhomogeneity due to an arbitrary distribution of bulk susceptibility. *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, 19(1):26–34, 2003.

[158] Sudeep Sarkar and Kim L Boyer. Perceptual organization in computer vision: A review and a proposal for a classificatory structure. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(2):382–399, 1993.

[159] Andrea Schenk, Guido Prause, and Heinz-Otto Peitgen. Efficient semiautomatic segmentation of 3d objects in medical images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 186–195. Springer, 2000.

[160] Dmitrij Schlesinger and Boris Flach. *Transforming an arbitrary minsum problem into a binary one*. TU, Fak. Informatik, 2006.

[161] Robert Schwarzenberg, Bernd Freisleben, Christopher Nimsky, and Jan Egger. Cube-cut: vertebral body segmentation in mri-data through cubic-shaped divergences. *PloS one*, 9(4):e93389, 2014.

[162] Thomas Albert Sebeok. *Signs: An introduction to semiotics*. University of Toronto Press, 2001.

[163] Ronald WK So and Albert CS Chung. Multi-level non-rigid image registration using graph-cuts. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 397–400. IEEE, 2009.

[164] Aristeidis Sotiras, Christos Davatzikos, and Nikos Paragios. Deformable medical image registration: A survey. *IEEE transactions on medical imaging*, 32(7):1153–1190, 2013.

[165] Mohamed Souiai, Claudia Nieuwenhuis, Evgeny Strekalovskiy, and Daniel Cremers. Convex optimization for scene understanding. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 9–14, 2013.

[166] Jean Stawiaski, Etienne Decenciere, and François Bidault. Interactive liver tumor segmentation using graph-cuts and watershed. In *Workshop on 3D segmentation in the clinic: a grand challenge II. Liver tumor segmentation challenge. MICCAI, New York, USA*, 2008.

[167] Gilbert Strang. Maximal flow through a domain. *Mathematical Programming*, 26(2):123–143, 1983.

[168] Evgeny Strekalovskiy and Daniel Cremers. Generalized ordering constraints for multilabel optimization. In *2011 International Conference on Computer Vision*, pages 2619–2626. IEEE, 2011.

[169] Yue Sun, Jing Yuan, Martin Rajchl, Wu Qiu, Cesare Romagnoli, and Aaron Fenster. Efficient convex optimization approach to 3d non-rigid mr-trus registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 195–202. Springer, 2013.

[170] Tommy WH Tang and Albert CS Chung. Non-rigid image registration using graph-cuts. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 916–924. Springer, 2007.

[171] Andrew Top, Ghassan Hamarneh, and Rafeef Abugharbieh. Active learning for interactive 3d image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 603–610. Springer, 2011.

[172] Thomas Torsney-Weir, Ahmed Saad, Torsten Moller, Hans-Christian Hege, Britta Weber, Jean-Marc Verbavatz, and Steven Bergner. Tuner: Principled parameter finding for image segmentation algorithms using visual response surface exploration. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):1892–1901, 2011.

[173] Andy Tsai, Anthony Yezzi, William Wells, Clare Tempany, Dewey Tucker, Ayres Fan, W Eric Grimson, and Alan Willsky. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging (TMI)*, 22(2):137–154, 2003.

[174] Jayaram K Udupa and Supun Samarasekera. Fuzzy connectedness and object definition: theory, algorithms, and applications in image segmentation. *Graphical models and image processing*, 58(3):246–261, 1996.

[175] Eranga Ukwatta, Jing Yuan, Martin Rajchl, Wu Qiu, David Tessier, and Aaron Fenster. 3-d carotid multi-region mri segmentation by globally optimal evolution of coupled surfaces. *IEEE Transactions on Medical Imaging (TMI)*, 32(4):770–785, 2013.

[176] Fedde van der Lijn, Tom den Heijer, Monique MB Breteler, and Wiro J Niessen. Hippocampus segmentation in mr images using atlas registration, voxel classification, and graph cuts. *Neuroimage*, 43(4):708–720, 2008.

[177] Annegreet van Opbroek, Fedde van der Lijn, and Marleen de Bruijne. Automated brain-tissue segmentation by multi-feature svm classification. *Proceedings of the MICCAI Grand Challenge on MR Brain Image Segmentation (MRBrainS13)*, 2013.

[178] Olga Veksler. Star shape prior for graph-cut image segmentation. In *European Conference on Computer Vision (ECCV)*, pages 454–467. Springer, 2008.

[179] Max A Viergever, JB Antoine Maintz, Stefan Klein, Keelin Murphy, Marius Staring, and Josien PW Pluim. A survey of medical image registration–under review. *Medical Image Analysis*, 2016.

[180] H Wang, B Avants, and P Yushkevich. A combined joint label fusion and corrective learning approach. In *MICCAI Workshop on Multi-Atlas Labeling*, 2012.

[181] Hongzhi Wang, Jung W Suh, Sandhitsu R Das, John B Pluta, Caryne Craige, and Paul A Yushkevich. Multi-atlas segmentation with joint label fusion. *IEEE transactions on pattern analysis and machine intelligence*, 35(3):611–623, 2013.

[182] Y Wang, Y Yu, D Li, KT Bae, JJ Brown, Weili Lin, and EM Haacke. Artery and vein separation using susceptibility-dependent phase in contrast-enhanced mra. *Journal of Magnetic Resonance Imaging*, 12(5):661–670, 2000.

[183] Yi Wang and Tian Liu. Quantitative susceptibility mapping (qsm): decoding mri data for a tissue magnetic biomarker. *Magnetic resonance in medicine*, 73(1):82–101, 2015.

[184] Simon K Warfield, Kelly H Zou, and William M Wells. Simultaneous truth and performance level estimation (staple): an algorithm for the validation of image segmentation. *IEEE transactions on medical imaging*, 23(7):903–921, 2004.

[185] Stephan Witoszynskyj, Alexander Rauscher, Jürgen R Reichenbach, and Markus Barth. Phase unwrapping of mr images using ϕun–a fast and robust region growing algorithm. *Medical image analysis*, 13(2):257–268, 2009.

[186] Hujun Yin and Nigel M Allinson. Self-organizing mixture networks for probability density estimation. *IEEE Transactions on Neural Networks*, 12(2):405–411, 2001.

[187] Yin Yin, Xiangmin Zhang, Rachel Williams, Xiaodong Wu, Donald D Anderson, and Milan Sonka. Logismoslayered optimal graph image segmentation of multiple objects and surfaces: cartilage segmentation in the knee joint. *IEEE transactions on medical imaging*, 29(12):2023–2037, 2010.

[188] Lei Ying, Zhi-Pei Liang, David C Munson, Ralf Koetter, and Brendan J Frey. Unwrapping of mr phase images using a markov random field model. *IEEE Transactions on Medical Imaging*, 25(1):128–136, 2006.

[189] Futoshi Yokota, Toshiyuki Okada, Masaki Takao, Nobuhiko Sugano, Yukio Tada, and Yoshinobu Sato. Automated segmentation of the femur and pelvis from 3d ct data of diseased hip using hierarchical statistical shape model of joint structure. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 811–818. Springer, 2009.

[190] Wan Chul Yoon. Task-interface matching: How we may design user interfaces. In *Proceedings of the 15th International Ergonomics Association Triennial Congress*, 2003.

[191] J Yuan, E Ukwatta, XC Tai, A Fenster, and C Schnoerr. A fast global optimization-based approach to evolving contours with generic shape prior. *IEEE TPAMI*, 12-38, 2012.

[192] Jing Yuan, Egil Bae, and Xue-Cheng Tai. A study on continuous max-flow and min-cut approaches. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2217–2224. IEEE, 2010.

[193] Jing Yuan, Egil Bae, Xue-Cheng Tai, and Yuri Boykov. A continuous max-flow approach to potts model. In *European conference on computer vision*, pages 379–392. Springer, 2010.

[194] Jing Yuan and Yuri Boykov. Tv-based multi-label image segmentation with label cost prior. In *BMVC*, volume 2, page 4, 2010.

[195] Jing Yuan, Brandon Miles, Greg Garvin, Xue-Cheng Tai, and Aaron Fenster. Efficient convex optimization approaches to variational image fusion. *Numerical Mathematics: Theory, Methods and Applications*, 7(02):234–250, 2014.

[196] Jing Yuan, Wu Qiu, Eranga Ukwatta, Martin Rajchl, Yue Sun, and Aaron Fenster. An efficient convex optimization approach to 3d prostate mri segmentation with generic star shape prior. *Prostate MR Image Segmentation Challenge, Conference on Medical Image Computing and Computer-Assisted Interventions (MICCAI)*, 7512, 2012.

[197] Jing Yuan, Juan Shi, Xue-Cheng Tai, and Yuri Boykov. A study on convex optimization approaches to image fusion. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 122–133. Springer, 2011.

[198] Paul A Yushkevich, Joseph Piven, Heather Cody Hazlett, Rachel Gimpel Smith, Sean Ho, James C Gee, and Guido Gerig. User-guided 3d active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *Neuroimage*, 31(3):1116–1128, 2006.

[199] Kun Zhou, Maxim Zaitsev, and Shanglian Bao. Reliable two-dimensional phase unwrapping method using region growing and local linear estimation. *Magnetic resonance in medicine*, 62(4):1085–1090, 2009.

[200] Mehrnaz Zouqi and Jagath Samarabandu. Prostate segmentation from 2-d ultrasound images using graph cuts and domain knowledge. In *Computer and Robot Vision, 2008. CRV'08. Canadian Conference on*, pages 359–362. IEEE, 2008.

Curriculum Vitae

Name:	John Stuart Haberl Baxter
Post-Secondary Education and Degrees:	University of Waterloo Waterloo, Ontario 2007 - 2012 B.S.E

Honours and Awards:

Western University Doctoral Excellence in Research Award (DERA)
2016-2017

Bayesian and Probabilistic Graphical Modelling Workshop - Best Paper Award
2016

NSERC Alexander Graham Bell Canada Graduate Scholarship (CGS D)
2015-2017

London Health Research Day - 1st Place Poster Award
2015

Imaging Network of Ontario - 1st Place Poster Award
2015

Ontario Graduate Scholarship - Queen Elizabeth II Graduate Scholarship in Science and Technology
2014-2015

London Health Research Day - 2nd Place Platform Presentation Award
2014

London Imaging Discovery Day - 3rd Place Student Research Award
2014

Ontario Graduate Scholarship
2013-2014

NSERC Alexander Graham Bell Canada Graduate Scholarship (CGS M)
2012-2013

University of Waterloo Sanford Fleming Foundation Award for Academic Excellence
2012

Sunnybrook Health Sciences Centre - Sunnybrook Prize Finalist
2012

Related Work Experience:
Teaching Assistant
The University of Western Ontario
Medical Imaging
2013-2016

Journal Manuscripts: (* - primary author)

Baxter, J.S.H.*, Rajchl, M., McLeod, A.J., Yuan, J. & Peters, T.M. "Directed Acyclic Graph Continuous Max-Flow Image Segmentation for Unconstrained Label Orderings" International Journal of Computer Vision, 2017. (accepted)

Ameri, G.*, **Baxter, J.S.H.**, McLeod, A.J., Peters, T.M. & Chen, E.C.S., "Effects of Line Fiducual Parameters and Beamforming on Ultrasound Calibration", Journal of Medical Imaging, 2017. (accepted)

Baxter, J.S.H.*, Inoue, J., Drangova, M., & Peters, T.M., "Shape Complexes: the intersection of label orderings and star convexity constraints in continuous max-flow medical image segmentation", Journal of Medical Imaging, 2016. (accepted)

Baxter, J.S.H.*, Rajchl, M., Peters, T.M. & Chen, E.C.S., "Optimization-based interactive segmentation interface for multiregion problems", Journal of Medical Imaging, 2015.

Rajchl, M.*, **Baxter, J.S.H.***, McLeod, A.J., Yuan, J. Qui, W., Peters, T.M. & Khan, A.R. "Hierarchical Max-Flow for Multi-Atlas Segmentation with Kohonen Self-Organizing Map Based Gaussian Mixture Modeling", Medical Image Analysis, 2015.

Cantor-Rivera, D.*, **Baxter, J.S.H.**, de Ribaupierre, S., Lau, J.C., Mirsattari, S.M., Goubran, M., Burneo, J.G., Steven, D.A., Peters, T.M. & Khan, A.R. "Individual feature maps: a patient-specific analysis tool with applications in temporal lobe epilepsy", International Journal of Computer Assisted Radiology, 2015.

McLeod, A.J,.*, **Baxter, J.S.H.**, Ameri, G., Ganapathy, S., Peters, T.M., & Chen, E.C.S. "Detection and Visualization of Dural Pulsation for Spine Needle Interventions ", International Journal of Computer Assisted Radiology, 2015.

Chen, E.C.S.*, McLeod, A.J., **Baxter, J.S.H.**, & Peters, T.M. "Registration of 3D Shapes Under Anisotropic Scaling: Anisotropic-Scaled Iterative Closest Point Algorithm", International Journal of Computer Assisted Radiology, 2015.

Kayvanrad, M.H.*, McLeod A.J., **Baxter, J.S.H.**, McKenzie C.A, & Peters, T.M., "Stationary wavelet transform for under-sampled MRI reconstruction", Magnetic Resonance Imaging, 2014.

Abhari, K.*, Baxter, J.S.H., Khan, A.R., Peters, T.M., de Ribaupierre, S., & Eagleson, R. "Visual Enhancement of MR Angiography Images to Facilitate Planning of Arteriovenous Malformation Interventions" ACM Transactions on Applied Perception, 2014.

Abhari, K.*, **Baxter, J.S.H.**, Chen, E.C.S., Khan, A.R., Peters, T.M., de Ribaupierre, S., & Eagleson, R. "Training for Planning Tumour Resection: Augmented Reality and Human Factors" IEEE Transactions in Biomedical Engineering, 2014.

Conference Manuscripts: (* - primary author)

Ameri, G.*, **Baxter, J.S.H.***, Bainbridge, D., Peters, T.M., & Chen E.C.S. (2017) "A Case Study on the Development of an Augmented Reality Ultrasound Guidance System", International Conference of Computer Assisted Radiology and Surgery, Barcelona, Spain.

Baxter, J.S.H.*, Hosseini, Z., Liu, J., Drangova, M., & Peters, T.M. (2016) "Cyclic Continuous Max-Flow: Phase Processing Using the Inherent Topology of Phase", Conference of the International Society of Magnetic Resonance in Medicine, Singapore.

Baxter, J.S.H.*, Inoue, J., Drangova, M., & Peters, T.M., (2016) "Shape Complexes in Continuous Max-Flow Hierarchical Multi-Labeling Problems". Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, San Diego, California, USA.

Pardasani, U.*, **Baxter, J.S.H.**, Peters, T.M., & Khan, A.R. (2016) "Single slice US-MRI registration for neurosurgical MRI-guided US". Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, San Diego, California, USA.

Rajchl, M.*, **Baxter, J.S.H.**, Qiu, W., Khan, A.R., Fenster, A., Peters, T.M., Ruecket, D., & Yuan, J. (2016) "Fast Defomrable Image Registration with Non-Smooth Dual Optimization", IEEE Conference on Computer Vision and Pattern Recognition.

Baxter, J.S.H.*, Rajchl, M., Yuan, J., & Peters, T.M., (2015) "Directed Acyclic Graphical Continuous Max-Flow Image Segmentation". MICCAI Workshop on Bayesian and Graphical

Models for Biomedical Imaging, Munich, Bavaria, Germany.

McLeod, A.J.*, **Baxter, J.S.H.**, Jayaranthne, U.., Pautler, S., Peters, T.M., & Luo, X. (2015) "Stereoscopic Motion Magnification in Minimally-Invasive Robotic Prostatectomy". MICCAI Workshop in Computer-Assisted and Robotic Endoscopy , Munich, Bavaria, Germany.

Inoue, J.*, **Baxter, J.S.H.**, & Drangova, M. (2015) "Left Atrial Wall Segmentation from CT for Radiofrequency Catheter Ablation Planning". MICCAI Workshop in Clinical Image Processing, Munich, Bavaria, Germany.

Ameri. G.*, **Baxter, J.S.H.**, McLeod, A.J., Peters, T.M. & Chen, E.C.S. (2015) "Augmented Reality Ultrasound Guidance for Central Line Procedures: Preliminary Results". MICCAI Workshop in Augmented Environments for Computer-Assisted Interventions, Bavaria, Germany.

Chen, E.C.S.* McLeod, A.J., **Baxter, J.S.H.**, & Peters, T.M. (2015) "An Iterative Closest Point Framework for Ultrasound Calibration". MICCAI Workshop in Augmented Environments for Computer-Assisted Interventions, Bavaria, Germany.

Baxter, J.S.H.*, Rajchl, M., McLeod, A.J., Khan, A.R., Yuan, J., & Peters, T.M., (2015) "Optimization Based Interactive Segmentation Interface for Multi-Region Problems". Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, Orlando, Florida, USA.

Ameri, G.*, McLeod, A.J., **Baxter, J.S.H.**, Chen, E.C.S., & Peters, T.M., (2015) "Line fiducial material and thickness considerations for ultrasound calibration". Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, Orlando, Florida, USA.

Rajchl. M.*, **Baxter, J.S.H.**, Bae, E., Xue-Cheng, T., Fenster, A., Peters, T.M. & Yuan, J. (2015) "Variational Time-Implicit Multiphase Level-Sets". Energy Minimization Methods in Computer Vision and Pattern Recognition, Hong Kong, China.

Baxter, J.S.H.*, Rajchl, M., McLeod, A.J., Khan, A.R., Yuan, J., & Peters, T.M., (2014) "Smoothness parameter tuning for generalized hierarchical continuous max-flow segmentation." Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, San Diego, California, USA.

Ameri, G.*, **Baxter, J.S.H.**, McLeod, A.J., Jayaranthe, U.L., Chen, E.C.S., & Peters, T.M., (2014) "Synthetic aperture imaging in ultrasound calibration." Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, San Diego, California, USA.

McLeod, A.J.*, **Baxter, J.S.H.**, de Ribaupierre, S., & Peters, T.M., (2014) "Motion magnification for endoscopic surgery." Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, San Diego, California, USA.

Abhari, K.*, **Baxter, J.S.H.**, Chen, E.C.S, Khan, A.R., Wedlake, C., Peters, T.M., de Ribaupierre, S., & Eagleson, R., (2013) "Use of a mixed-reality system to improve the planning of brain tumour resections: Preliminary results", Augmented Environments for Computer-Assisted Interventions, Lecture Notes in Computer Science, Springer Berlin / Heidelberg.

Abhari, K.*, **Baxter, J.S.H.**, Chen, E.C.S, Wedlake, C., Peters, T.M., Eagleson, R., & de Ribaupierre, S., (2013) "The Role of Augmented Reality in the Training and Planning of Brain Tumour Resection", Joint Medical Image and Augmented Reality (MIAR) Nagoya, Japan.

Rajchl, M.*, **Baxter, J.S.H.**, Yuan, J., Peters, T.M., & Khan, A.R., (2013) " ASeTs: MAP-Based Brain Tissue Segmentation using Manifold Learning and Hierarchical Max-Flow Regularization", Medical Image Computation and Computer-Assisted Interventions (MICCAI), Nagoya, Japan.

Abhari, K.*, **Baxter, J.S.H.**, Chen, E.C.S, Wedlake, C., Peters, T.M., Eagleson, R., & de Ribaupierre, S., (2013) "Development and Evaluation of an Augmented-Reality Training System for Planning Brain Tumour Resection Interventions", Medical Image Understanding and Analysis, Birmingham, UK.

Abhari, K.*, **Baxter, J.S.H.**, de Ribaupierre S., Peters T. M., & Eagleson R., (2012), "Perceptual Improvement of Volume-Rendered MR Angiography Data using a Contour enhancement Technique, Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, San Diego, California, USA.

Chen, E.C.S.*, Sarkar, K., **Baxter, J.S.H.**, Moore, J., Wedlake, C., & Peters, T.M., (2012) "An augmented reality platform for planning of minimally invasive cardiac surgeries", Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, San Diego, California, USA.

Buchanan, S.*, Moore, J., Lammers, D., **Baxter, J.S.H.**, & Peters, T.M., (2012) "Characterization of tissue-simulating phantom materials for ultrasound-guided needle procedures", Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, San Diego, California, USA.

Baxter J.S.H.*, Peters T. M & Chen E. C. S., (2011), "A unified framework for voxel classification and triangulation", Society of Photographic Instrumentation Engineers Proceedings in Medical Imaging, Orlando, Florida, USA.

Technical Reports: (* - primary author)

Baxter J.S.H.*, McLeod, A.J. & Peters, T.M., (2015) "A Continuous Max-Flow Approach to Cyclic Field Reconstruction", arXiv preprint arXiv:1511.03629.

Baxter, J.S.H.*, Yuan, J., & Peters, T. M. (2015). "Shape Complexes in Continuous Max-Flow Hierarchical Multi-Labeling Problems". arXiv preprint arXiv:1510.04706.

Baxter, J.S.H.*, Rajchl, M., Yuan, J., & Peters, T. M. (2015). "A Proximal Bregman Projection Approach to Continuous Max-Flow Problems Using Entropic Distances". arXiv preprint arXiv:1501.07844.

Baxter, J.S.H.*, Rajchl, M., Yuan, J., & Peters, T. M. (2014). "A Continuous Max-Flow Approach to Multi-Labeling Problems under Arbitrary Region Regularization". arXiv preprint arXiv:1405.0892.

Baxter, J.S.H.*, Rajchl, M., Yuan, J., & Peters, T. M. (2014). "A Continuous Max-Flow Approach to General Hierarchical Multi-Labeling Problems". arXiv preprint arXiv:1404.0336.

Rajchl, M.*, **Baxter, J.S.H.**, Qiu, W., Khan, A. R., Fenster, A., Peters, T. M., & Yuan, J. (2014). "RANCOR: Non-Linear Image Registration with Total Variation Regularization". arXiv preprint arXiv:1404.2571.

Reviewing Activities

Journals:

- SPIE Journal of Medical Imaging (1 paper)
- IEEE Signal Processing Letters (1 paper)
- Medical Image Analysis (MedIA) (1 paper)
- International Journal of Computer Assisted Radiology and Surgery (IJCARS) (1 paper)
- IET Journal of Computer Vision (2 papers)

Conferences:

- Information Processing in Computer Assisted Interventions (IPCAI) (4 papers)
- Workshop in Augmented Environments in Computer Assisted Interventions (AE-CAI) (2 papers)
- Medical Image Processing and Computer Assisted Interventions (MICCAI) (1 paper)