10-19-2022 11:00 AM

The Effect of Active Learning on Viewpoint Dependence for Novel Objects

Cassandra Wallace Bacher, *The University of Western Ontario*

Supervisor: Culham, Jody C., *The University of Western Ontario*
Co-Supervisor: Mur, Marieke, *The University of Western Ontario*
Co-Supervisor: Goodale, Melvyn A., *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Master of Science degree in Neuroscience

Abstract

Active learning of novel objects can facilitate subsequent object recognition and discrimination, but the reasons for its beneficial effects remain unclear. One potential explanation is that active learning enables the formation of a more detailed, realistic, or useful neural object representation than does passive learning. The current study addressed the question of whether active vs. passive learning of objects affects viewpoint discrimination. Participants learned novel wire-like objects either actively or passively and then completed a psychophysical task which they discriminated object orientation. This study did not find a significant difference in viewpoint discrimination between actively and passively learned object representations, which stands in contrast to earlier studies that found an effect of active learning on object recognition across different viewpoints. This suggests that viewpoint discrimination and viewpoint generalization rely on different mechanisms.

Keywords

visual perception, object recognition, active and passive learning, viewpoint dependence, viewpoint tuning, psychophysics

Summary for Lay Audience

People recognize new objects more quickly after moving and turning them around with their hands. However, scientists do not know exactly why handling objects is so helpful. Maybe looking at an object while controlling its movements is the best way to learn what the object looks like from every angle. This study explored whether people were better or worse at telling the difference between views of objects when they actively moved the objects than when they watched videos of objects being moved by someone else. Volunteers viewed four simple, wire-like objects on a computer screen. Each person learned two objects by moving them with a trackball (active learning) and two objects by watching a video (passive learning). After learning, the volunteers took tests that measured how good they were at telling whether the objects were facing left or right. This study found that volunteers' test performance was the same for all of the objects no matter how they were learned. These results suggest that handling a new object and watching someone else handling it are equally good ways of learning what the object looks like from different angles. There must be a reason why actively learning objects helps people recognize them more quickly, but does not help people tell the difference between views of the objects. Scientists will have to perform new experiments to discover that reason.

Acknowledgments

Support for this project was provided by a Natural Sciences and Engineering Research Council (Canada) Discovery Grant.

This research benefited from the OurBrainsCAN: University of Western Ontario's Cognitive Neuroscience Research Registry, which was funded by the Canada First Research Excellence Fund awarded to BrainsCAN at the University of Western Ontario.

I would like to thank Professor Jody Culham, Ph.D., Professor Marieke Mur, Ph.D., Professor Mel Goodale, Ph.D., Professor Stefan Köhler, Ph.D., and Professor Ryan Stevenson, Ph.D. for their strong support and valuable advice; they provided guidance in selecting this research topic and were very accessible mentors and supporters.

I would like to thank Marisa Donnelly for her guidance and assistance with conducting experiments; she was always there whenever I needed help or had a question.

I would like to thank Nishita Sharif for her valuable support and assistance with data collection.

Additional thanks are owed to Kevin Stubbs and Karsten Babin for their much-appreciated assistance with all of the computer hardware and software used for this research.

Last, but not least, thank you to my parents for their unending help and support throughout this project.

Table of Contents

List of Figures

Chapter 1

1 Introduction

1.1 Active Learning

Throughout their lives, people encounter thousands of physical objects varying in their appearance, function, and relevance. Babies and very young children must dedicate considerable time and effort to comprehending the multitude of novel objects surrounding them. Even adults, with their years of experience and vastly greater knowledge of the world, frequently encounter objects with which they are unfamiliar. In such situations, children and adults alike will often immediately reach out to grasp the strange object in their hands, simultaneously viewing and manipulating the item that has caught their interest. This behavior, second nature as it is to so many, may seem wholly unremarkable. However, self-produced movement appears to play a critical role in successfully understanding and interacting with our environment.

Past research suggests that active exploration promotes proper brain development early in life. In one classic experiment, Held and Hein (1963) placed pairs of kittens in an enclosure arranged so that each individual in the pair viewed the same stimuli at the same time with only a single difference in their experiences: the first kitten (active) had the freedom to move under its own power, while the second kitten (passive) sat in a gondola yoked to the movements of the first kitten. When the kittens were later given tests of visually guided behavior, the active kittens performed normally while the passive kittens showed impairments. Subsequent studies in other species yielded similar results. For example, active engagement with environmental enrichment such as plastic tubes, grids, and colorful cartons increased the survival of newborn mouse neurons and improved performance in a Morris water maze. However, passive viewing of other mice interacting with environmental enrichment failed to have the same effect (Iggena et al., 2015). Human babies also seem to benefit from active learning: three-month-olds who were allowed to independently pick up objects later showed more reaching behavior and more visual exploration of people and objects in their environment than infants who passively observed their parents' actions (Libertus et al., 2016; Libertus & Needham, 2010; Wiesen et al., 2016). In addition, other studies have found that "hands-on" experience with objects facilitates mental rotation in babies as young as six months old (Frick & Wang, 2014; Möhring & Frick, 2013).

Active learning is not only relevant in infancy, however; it continues to be an important component in perception and cognition well into adulthood. For example, evidence indicates that active exploration of a new environment aids spatial learning to a greater degree than does passive exposure, possibly due to the decision making involved in navigation or the motor and proprioceptive information acquired during locomotion (Chrastil & Warren, 2013, 2015; Peruch et al., 1995). Active learning seems to facilitate object recognition as well. Harman et al. (1999) conducted an experiment in which participants learned novel virtual objects either by using a trackball to rotate the object on a computer screen or by watching a video recording of a previous participant's exploration of the object. Object recognition was then assessed via an old/new discrimination task. While there was no effect of learning condition on the accuracy of subsequent object recognition, the actively learned objects were recognized significantly faster than the passively learned objects. These results were replicated within a virtual reality environment, and a follow-up study found that active learning resulted in faster performance on a mental rotation task (James et al., 2001, 2002). Indeed, other researchers later observed that active exploration of objects provided an advantage in a variety of different tasks involving view-matching, similarity judgement, and aperture-viewing (Craddock et al., 2011; Lee & Wallraven, 2013; Sasaoka et al., 2010). Moreover, active exploration was found to facilitate facial recognition, and it even resulted in faster access to conscious awareness during a continuous flash suppression paradigm (Liu et al., 2007; Suzuki et al., 2019). Although these behavioral effects have been well documented, a comprehensive explanation for the "active learning advantage" in object recognition is still being pursued.

One possible explanation is that actively learning an object provides more valuable visual information than does passively learning the object. This theory is supported by the finding that actions seemed to aid mental rotation in infants only when the actions provided useful visual cues (Antrilli & Wang, 2016). Furthermore, when given complete control over the movements made during object exploration, people tend to allocate their attention strategically, focusing heavily on certain aspects of the stimulus while showing comparative disregard for others (Craddock et al., 2011; Ernst et al., 2007; Harman et al., 1999; James et al., 2001, 2002, 2014). Nevertheless, this does not offer a sufficient explanation for the beneficial effects of active learning, which have still been observed even when the same visual information is provided in both active and passive learning conditions (Harman et al., 1999; James et al., 2001, 2002). The implication of such findings is that active learning, rather than simply supplying more useful visual information than passive learning, impacts the fundamental processes through which objects are perceived and remembered. Indeed, there is evidence that active exploration increases perceptual sensitivity, thereby improving the quality of visual memories of three-dimensional objects. This effect, which could not be explained solely by differences in attentional allocation, was observed for both simple and complex objects and was so robust that it compensated for perceptual degradation due to masking (Meijer & van der Lubbe, 2011). Perhaps this increased perceptual sensitivity enables people to form more comprehensive neural representations of objects. If so, that could at least partially explain the advantage that active learning provides. However, the majority of past studies examining active learning of objects focused only on the ways in which it impacts discrimination *between* multiple different objects. Very little research has explored the unique characteristics of actively learned representations of individual objects themselves. Thus, the current study aimed to begin investigating the effect of active learning on neural representations of objects by focusing specifically on one aspect of such representations: viewpoint dependence.

1.2 Viewpoint Dependence

One of the most remarkable aspects of the brain is the ease and speed with which it can recognize the thousands of individual objects that fill our environment despite their constant transformations in size, shape, and location on the retina. In a three-dimensional world, it is particularly important to be able to quickly and accurately identify a single object in any orientation. Viewpoint dependence refers to the extent to which changes in viewpoint or orientation affect object recognition.

Whether object recognition is viewpoint-dependent or viewpoint-independent has long been a subject of debate, and early studies on the subject yielded conflicting results. Findings from behavioral and electrophysiological experiments in both humans and monkeys provided support for viewpoint-dependent object representations (Lawson & Humphreys, 1996; Logothetis et al., 1994; Logothetis & Pauls, 1995). The presence of object-selective neurons tuned to specific views in the human brain was later confirmed by Fang and He (2005). In contrast, Biederman and Gerhardstein (1993) as well as Biederman and Bar (1999) found that participants were able to immediately form viewpoint-independent representations of objects rotated in depth, though only if the following three conditions were satisfied: (1) The objects must be capable of activating geon structural descriptions (GSDs), (2) the GSDs must be different for each object, and (3) the same GSD must be activated in both original and tested views. The authors theorized that the viewpoint dependence observed in other studies could be explained by the usage of study stimuli that did not meet these requirements. More recent evidence suggests, though, that complete viewpoint independence does not exist in the human visual system. Rather, object recognition seems to be viewpoint-dependent to a degree that is contingent on the nature of the stimuli, the demands of the current task, and the stage of visual processing (Andresen et al., 2009; Tarr & Hayward, 2017).

Another factor which may have an impact on the viewpoint dependence of object representations is the act of learning itself. For example, functional magnetic resonance imaging (fMRI) studies have found decreased sensitivity to changes in viewpoint for familiar faces versus unfamiliar faces (Eger et al., 2005; Ewbank & Andrews, 2008). In addition, researchers have observed that unsupervised learning improved participants' ability to recognize objects across rotation, suggesting that the learning resulted in broader viewpoint tuning (Tian & Grill-Spector, 2015). While the current study also examined the effect of learning on the viewpoint dependence of neural object representations, it was unique in addressing the potential difference between active and passive learning.

1.3 Current Study

Although previous studies have shown that active learning facilitates discrimination between different objects more than does passive learning, the precise reason for this phenomenon is still unknown. A logical first step in answering the question of why active learning expedites object recognition is to investigate how certain aspects of actively learned and passively learned object representations differ from one another. Viewpoint dependence is one such aspect which may be affected by active learning. Seeing the spatially and temporally correlated views of an object in motion seems to be critical for the formation of an accurate, three-dimensional representation enabling recognition across, or discrimination between, different views (Orlov & Zohary, 2018; Wallis, 2002; Wallis & Bülthoff, 2001). However, predictions about the relationship between the motion of the object and changes in viewpoint can be more easily tested during active learning than during passive learning, which could result in a richer representation of the various object views.

The current study assessed viewpoint dependence for four novel virtual objects, two of which were learned actively and two of which were learned passively. Immediately after learning each object, participants completed a psychophysical task requiring them to view images of the object from various viewpoints and indicate whether the object in each image was facing/rotated left or right. The participants' responses were recorded and used to calculate the just-noticeable difference (JND). The JND, which is the smallest difference between two stimuli that can be detected, reflected participants' sensitivity to changes in viewpoint. As such, it served as a quantifiable measure of the viewpoint dependence of object representations.

This study tested two opposing hypotheses. First, active learning improves our ability to discriminate one view of an object from another, leading to narrower viewpoint tuning for representations of actively learned objects than for representations of passively learned objects. This *narrower* viewpoint tuning would manifest as a *decreased* JND. Second, active learning improves our ability to recognize the same object across multiple views, leading to broader viewpoint tuning for representations of actively learned objects than for representations of passively learned objects. This *broader* viewpoint tuning would manifest as an *increased* JND.

Chapter 2

2 Materials and Methods

2.1 Participants

Data were analyzed from 22 right-handed participants (8 males and 14 females, ages 18-50), who were recruited using OurBrainsCAN: University of Western Ontario's Cognitive Neuroscience Research Registry as well as recruitment posters placed on campus. Data from four additional participants were collected, but discarded as outliers during the quality assurance process.

All participants had normal or corrected-to-normal vision. If a participant's handedness was not already recorded in the OurBrainsCAN registry, the Edinburgh Handedness Inventory was used prior to testing to ensure that the participant had a right-hand preference for at least 90% of the activities listed. Informed consent was obtained from all individual participants, who received financial compensation ($15) for their time. The study was approved by Western University's Non-Medical Research Ethics Board in accordance with the standards of the 1964 Declaration of Helsinki.

2.2 Stimuli and Apparatus

Five three-dimensional novel wire-like objects, including one practice object, were created using Blender (v. 2.83) 3-D rendering software (**Figure 1a**). Due to the nature of the psychophysical task that participants would perform, the objects were designed to be bilaterally symmetrical (the left and right halves of the object are mirrored) and non-bistable (the object is unambiguous, having only one possible perceptual interpretation). Eleven test images of each object (0° front views and 3°, 6°, 9°, 12° and 15° side views) were then generated by rotating the object in depth both clockwise and counterclockwise with the frontal view (as defined by Fang and He (2005)) as the initial position (**Figure 1b**).

Participants viewed all experimental stimuli on an Asus VE247H 24-inch computer monitor with a spatial resolution of 1920 x 1080 pixels at a viewing distance of 57 cm.

The stimuli, which consisted of both static images and dynamic renderings of three-dimensional models, were presented against a black background using the software programs Unity (v. 2019.4.18f1), MATLAB 2021a, and Psychtoolbox-3 (Brainard, 1997; Pelli, 1997). The test images, sized so that the objects they depicted extended no more than 3.2° of visual angle in width or height, were presented in randomly selected locations within the boundaries of a 5.7° x 5.7° area centered within the computer monitor. This random variation in image location prevented retinal adaptation to edges from becoming a confounding factor.

2.3 Procedure

Prior to beginning the experiment, all participants underwent a five-minute practice session designed to familiarize them with the procedure and tasks. The object featured in the practice session was never used in the experiment itself.

The experimental procedure lasted approximately 45 minutes and consisted of alternating learning and test phases. Each participant learned four objects, two of them actively and two of them passively. During the active learning condition, the participant explored an object for 60 seconds using a trackball to freely rotate it around the vertical axis. During the passive learning condition, the participant watched a 60-second recording of an object being actively learned by a previous participant (**Figure 1a**). In both conditions, the participant was instructed to focus on the object’s three-dimensional shape.

After learning each object, the participant completed a psychophysical task to assess viewpoint dependence for that object. Each of 11 test images was presented ten times in random order for a total of 110 trials; this technique of rapid and repetitive stimulus presentation is known as the method of constant stimuli. For the entire duration of the task, the participant was instructed to continuously fixate on a red cross presented in the center of the screen. Every trial began with a 2-second blank period followed by the 0.2-second presentation of a test image. After the test image presentation, the participant made a two-alternative forced-choice response, pressing one of two keyboard keys to indicate whether the object in the image had been facing to the left or to the right (**Figure 1b**). Participants had been encouraged in the training session to respond based on their initial visual impressions and to avoid using cognitive strategies to determine the direction in which objects were facing.

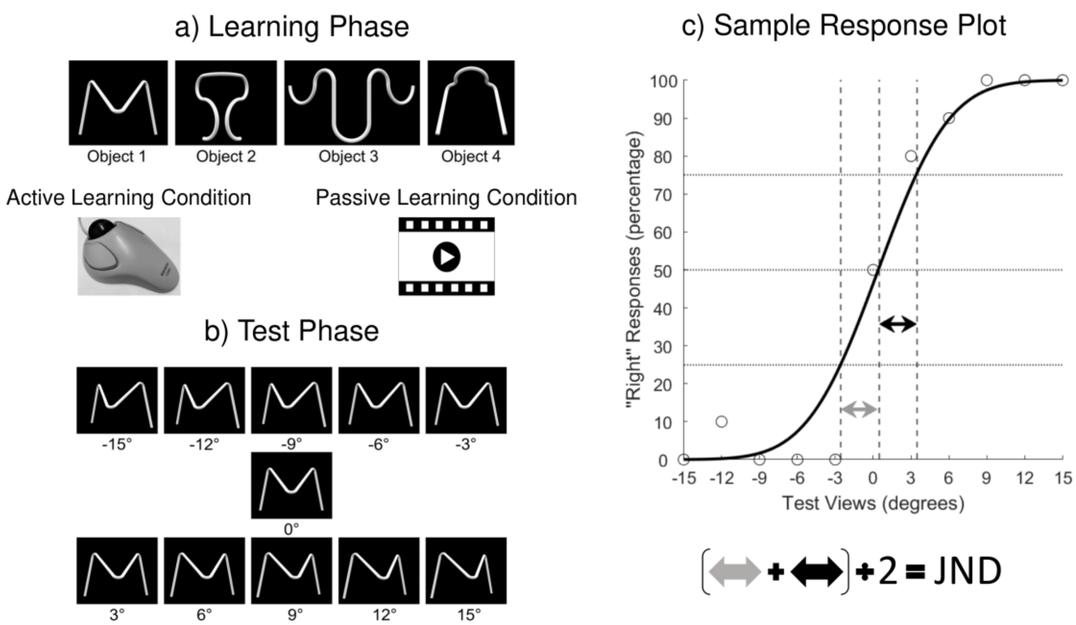

Figure 1: Experimental Stimuli and Design. a) Four wire-like novel objects were created for use in this study. A fifth object (not shown) was used as part of a practice session. Participants learned each novel object either by rotating it with a trackball (active learning condition) or by watching a recording of a previous participant's active learning session (passive learning condition). **b)** Immediately following learning, participants viewed test images of the object and indicated whether the object in each image appeared to be facing left (< 0°) or right (> 0°). Test images had been generated by rotating each object clockwise or counterclockwise from a frontal view. **c)** The fraction of trials in which participants indicated that the object was facing right was plotted as a function of test view. A psychometric function was then fitted to the data and used to estimate the just-noticeable difference (JND), which is the smallest difference between two stimuli that can be detected. The JND served as a measure of participants' sensitivity to changes in viewpoint. Participants who were better able to discriminate between orientations had a smaller JND (i.e., steeper slope in the middle of the S-shaped function) than participants with poorer discrimination ability.

2.4 Data Analysis

Data analysis was performed using scripts written in MATLAB 2021a. The fraction of trials in which participants indicated that the object in a test image was facing right was calculated and plotted as a function of test view. An S-shaped psychometric function was then fitted to the data points and used to estimate the JND (mathematically defined as the mean of the difference between the value of x at y = 75% and the value of x at y = 50% and the difference between the value of x at y = 50% and the value of x at y = 25%) (**Figure 1c)**.

Quality assurance was performed by first plotting each participant's data and examining the psychometric functions. With only a few exceptions, the functions were S-shaped, and responses spanned from 0% to 100% "right" responses (**Figure 2**). Then the JND distributions for each of the four wire objects were plotted separately. Participants with data lesser than Q1 – (1.5 x IQR) or greater than Q3 + (1.5 x IQR) for any of the objects were flagged as outliers and removed from subsequent analyses (**Figure 3**). After outlier removal, the JND distributions were approximately normal.

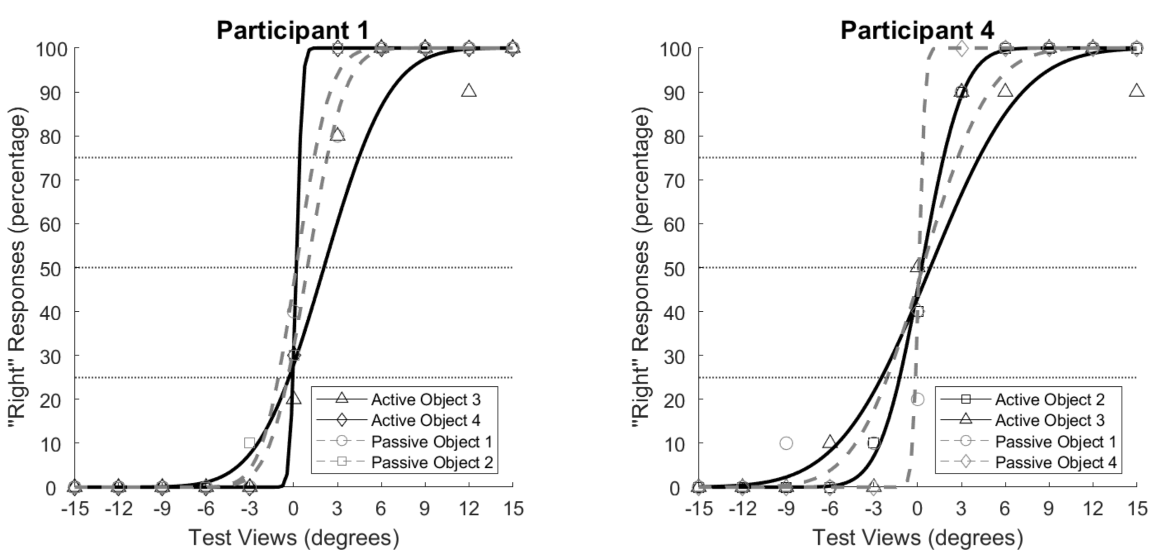

Figure 2: Sample Psychometric Functions. As part of data quality assurance, the psychometric functions for each object and participant were examined. The majority were visibly S-shaped and spanned from 0% to 100% "right" responses.

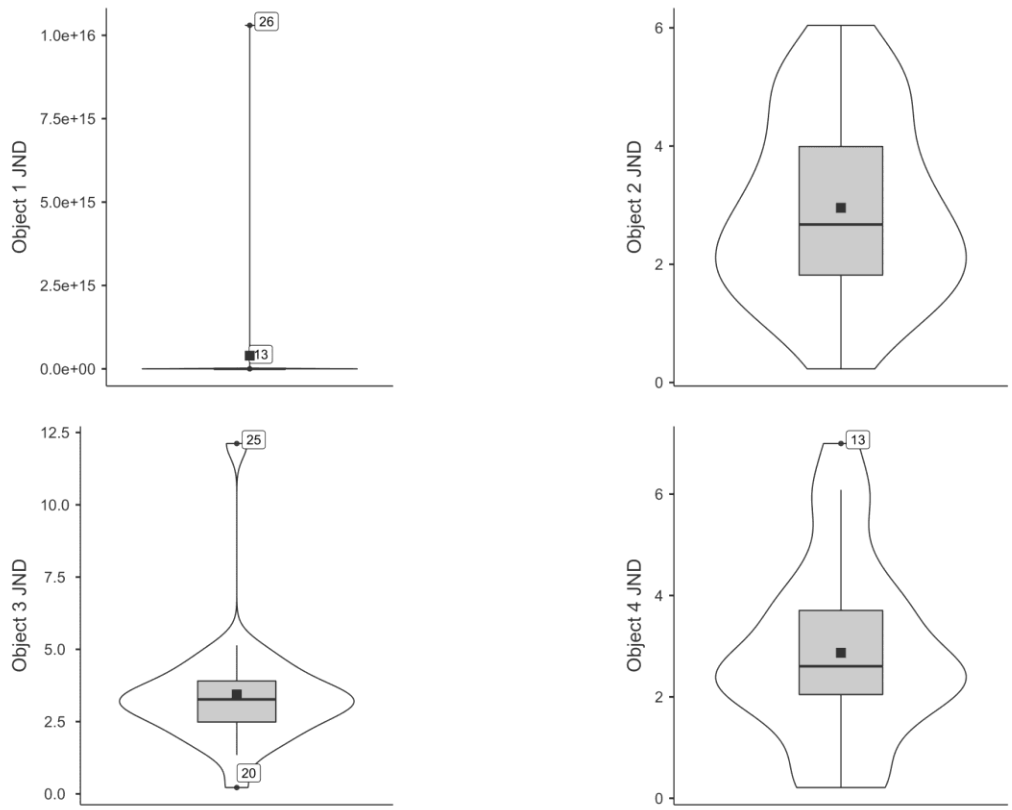

Figure 3: Outlier Identification. Data quality assurance also included the identification and removal of outliers. The JND distributions (*n* = 26) for the four wire objects were plotted and examined. Inside the outlines formed by violin plots are gray boxplots depicting the median as well as quartiles Q1 and Q3. The mean of each distribution is represented by a small black square. Outliers (defined as values lesser than Q1 – (1.5 x IQR) or greater than Q3 + (1.5 x IQR)) have been automatically marked and labeled with their participant IDs. The four participants with outliers were excluded from all further analyses.

Statistical tests were performed using jamovi (The jamovi project, 2021), an open-source statistical platform. A conventional (Frequentist) paired-samples *t* test was conducted to determine whether a significant difference in JND existed between actively and passively learned objects. A Frequentist linear mixed model was also utilized to this end because a simple *t* test was unable to take into account variation between participants and between objects.

An important limitation of Frequentist statistics is that they can only reject (or fail to reject) the null hypothesis. This limitation was overcome by performing a Bayesian paired-samples *t* test capable of estimating the extent to which the null hypothesis was more or less likely than the experimental hypothesis to be true.

Chapter 3

3 Results

3.1 JND

JND was the primary dependent variable of interest in this study. Descriptively, the mean JND for actively learned objects was minimally greater than the mean JND for passively learned objects (**Figure 4**), but the majority of participants displayed little to no difference in JND between actively learned and passively learned objects (**Figure 5**). Statistically, a Frequentist paired-samples *t* test found no significant difference in JND between the two learning conditions ($t(21) = 4.8$, $p = .35$). A Frequentist linear mixed model that took into account variation between participants and objects also found no significant difference in JND between the active learning condition and the passive learning condition ($F_{(1, 62.2)} = 1.1$, $p = .30$).

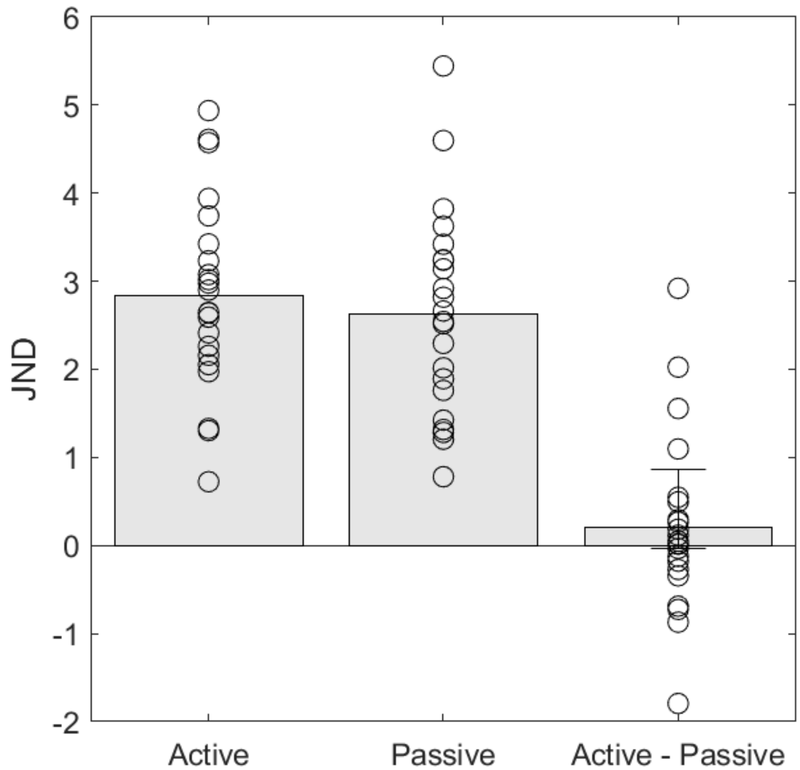

Figure 4: The Effect of Learning Condition on JND. No significant difference in mean JND between actively learned and passively learned objects was found (see text for details on statistical inference). Gray bars represent means, while values for individual participants (*n* = 22) are represented by black open circles. The error bar represents the 95% confidence interval for the "active – passive" difference. This error bar encompasses zero, indicating no significant difference between learning conditions.

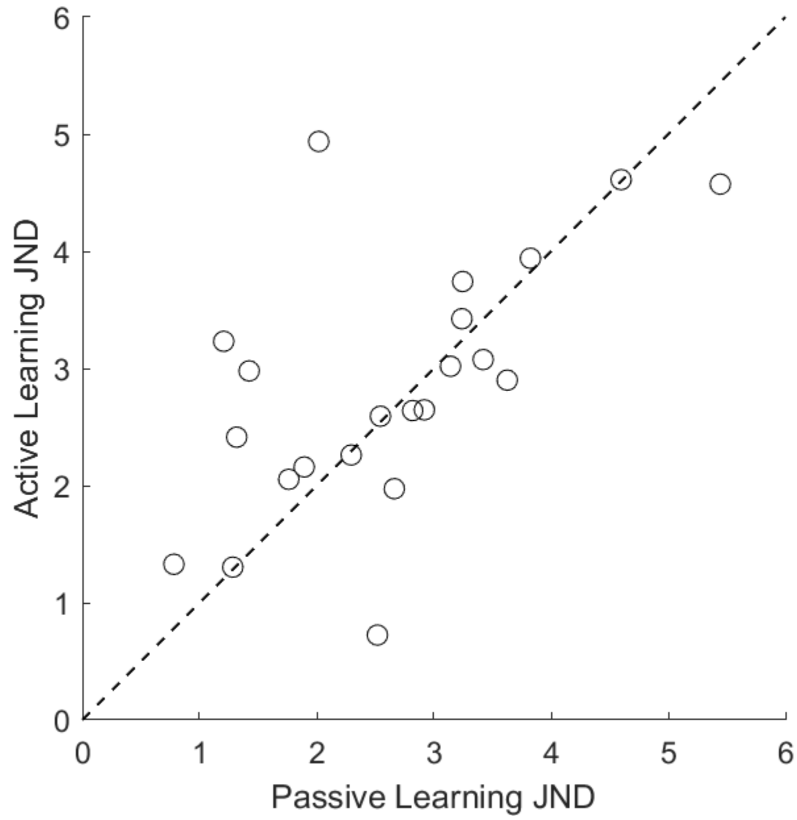

Figure 5: Active Learning vs. Passive Learning JND for Individual Participants (*n* = 22). Participants for whom the active learning JND exceeded the passive learning JND fall to the left of the dashed line; participants for whom the passive learning JND exceeded the active learning JND fall to the right of the dashed line. Note that most participants' data fall along the dashed line, indicating a tight correlation ($r(20) = .60$, $p = .003$) between the two measures with no clear differences between actively learned and passively learned objects.

In accordance with the results of the Frequentist paired-samples *t* test, a Bayesian paired-samples *t* test found that the null hypothesis (of no difference between active and passive learning) was nearly three times more likely (Bayes Factor, $BF_{01} = 2.96$) than the experimental hypothesis (which postulated a difference between active and passive learning). This is considered "anecdotal" evidence in favor of the null hypothesis. Furthermore, a plot of the strength of evidence as a function of the number of participants tested (**Figure 6**) shows that the level of evidence remained stable beyond 16+ participants, suggesting that the absence of a difference was not likely to be due to the limited sample size.

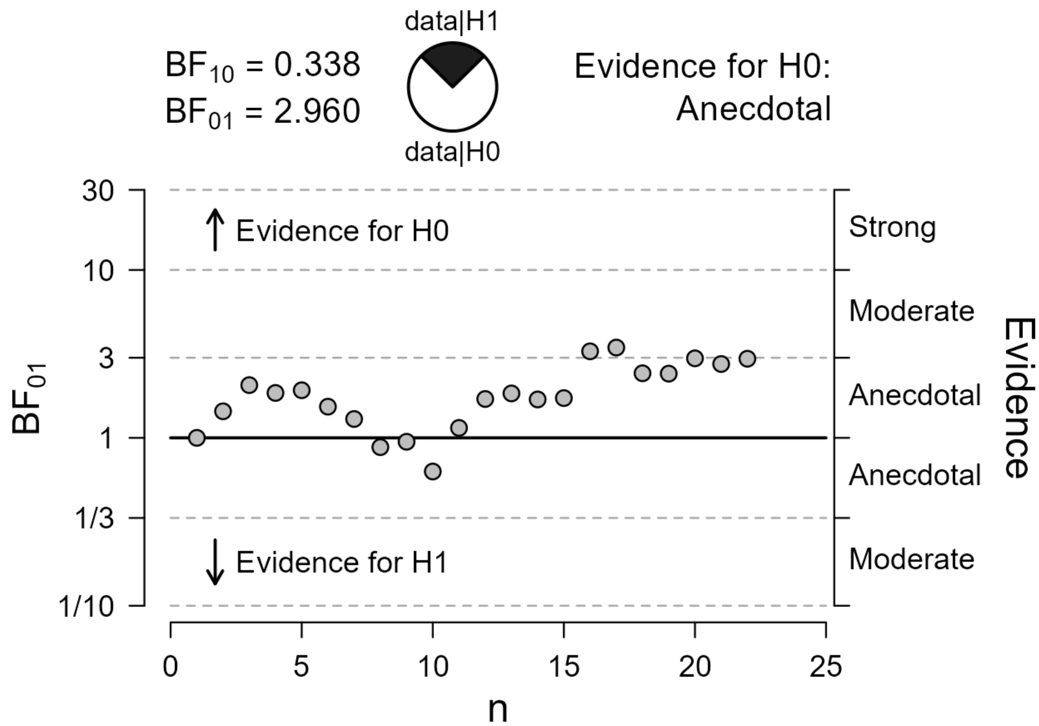

Figure 6: Bayesian *T* Test Inferential Plot. A Bayesian paired-samples *t* test was conducted to estimate the extent to which the null hypothesis (there is no difference between active learning and passive learning) is more or less likely than the experimental hypothesis (there is a difference between active learning and passive learning). The test revealed that the null hypothesis was almost three times more probable than the experimental hypothesis and that the evidence remained relatively stable once data from 16 participants had been collected.

Overall, both Frequentist and Bayesian statistics indicated that there was no significant difference in JND between actively learned and passively learned objects.

3.2 Viewing Time During Active Learning

Although the main focus of this study was to investigate the effect of active learning on viewpoint dependence, the strategies participants used during active learning were also examined and compared to the results of a past study that used different objects (Harman et al., 1999).

During the active learning condition, the software used to present the three-dimensional models continuously recorded the orientation of the object. The 360° rotation range was divided into 10° increments, and the percentage of total viewing time that participants spent studying the object at each orientation was calculated. Those values were then averaged across participants and objects.

As shown in **Figure 7**, participants focused on front views (350° - 10°), studying them for 16% of the learning phase. In comparison, they studied the right side (80° - 100°), back (170° - 190°), and left side (260° - 280°) views for, respectively, 4.6%, 4.8%, and 4% of the learning phase. This general pattern of exploration was relatively consistent across objects and participants.

Viewing patterns between the current study and Harman et al.'s (1999) study were compared. In Harman et al.'s (1999) study, participants could rotate the objects around two axes to see any possible view of the objects; thus, the exploration data were plotted as a two-dimensional heat map (viewing time for viewpoints along two axes). Qualitatively, their data showed that participants spent the most time looking at views rotated around the vertical axis, particularly from four cardinal directions (front, back, and two sides). In the current study, participants could only rotate objects around one axis (the vertical axis, corresponding to left-right rotation). To facilitate qualitative comparisons between data from the two studies, Harman et al.'s (1999) data (only the values for rotation around the vertical axis, where participants spent the most time) were replotted. This was accomplished by estimating relative viewing time percentages based on the intensity values of the lookup table for the heat map depicted in Figure 4 of the article. Although these inferred data were not perfectly accurate replications of the original data, they sufficiently depicted the trends found in the original data to enable qualitative (but not quantitative) comparisons. While the current study's participants focused primarily on front views of objects, Harman et al.'s (1999) participants concentrated on side views as much as they did front views (**Figure 8**).

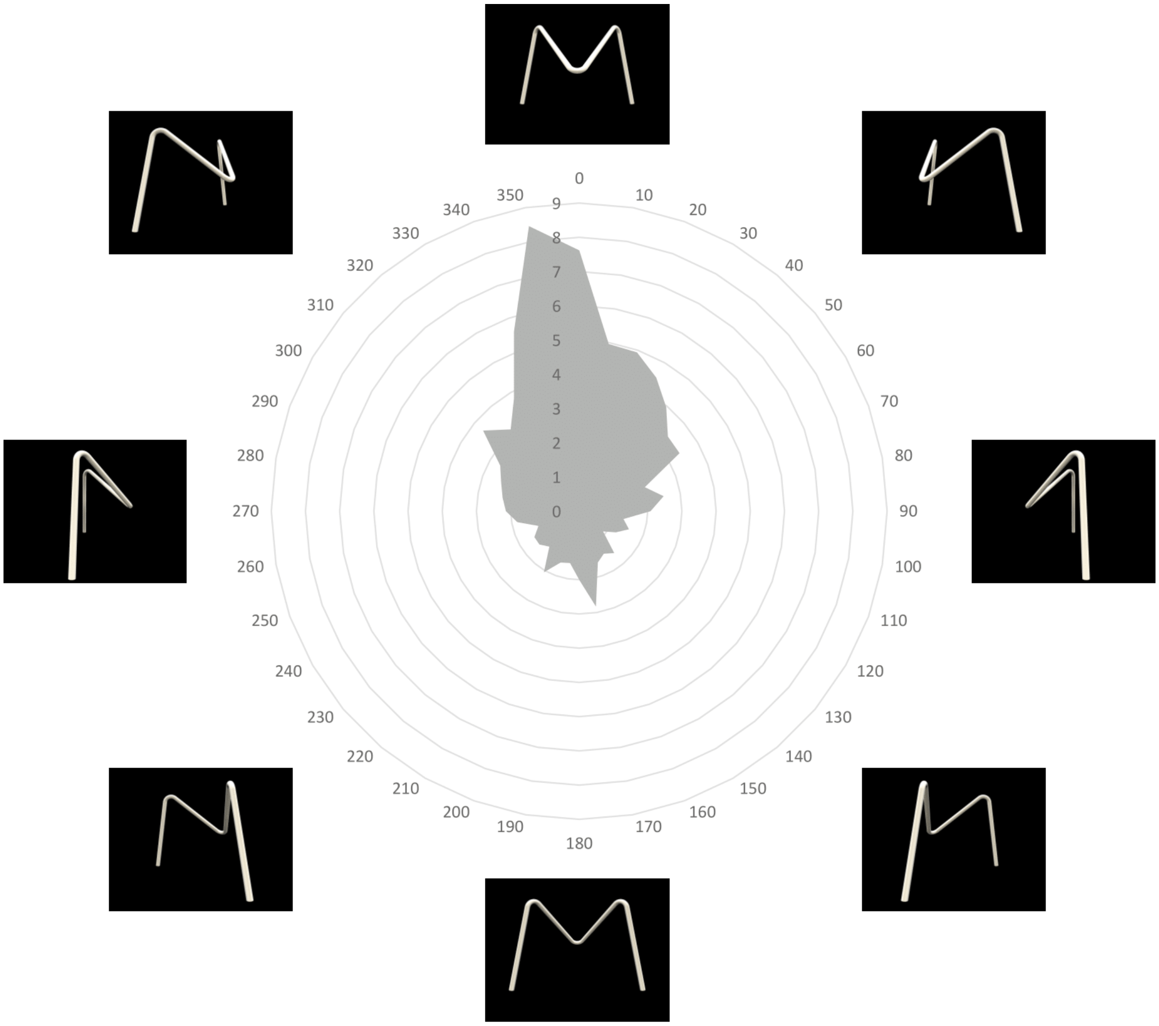

Figure 7: Viewing Patterns During Active Learning. Participants (*n* = 22) spent more time studying front views than they did studying any other views. Object orientation (°) was continuously recorded during every active learning session, and the percentage of total viewing time that participants spent observing the object at each orientation was calculated. Percentages were then averaged across all participants and objects before being plotted. This polar plot depicts the mean percentage of total viewing time (along the radius) as a function of object orientation (around the polar angle) with the initial orientation shown at 0° (12 o’clock position). The spatial resolution (bin size) of the viewing time calculation was 10°.

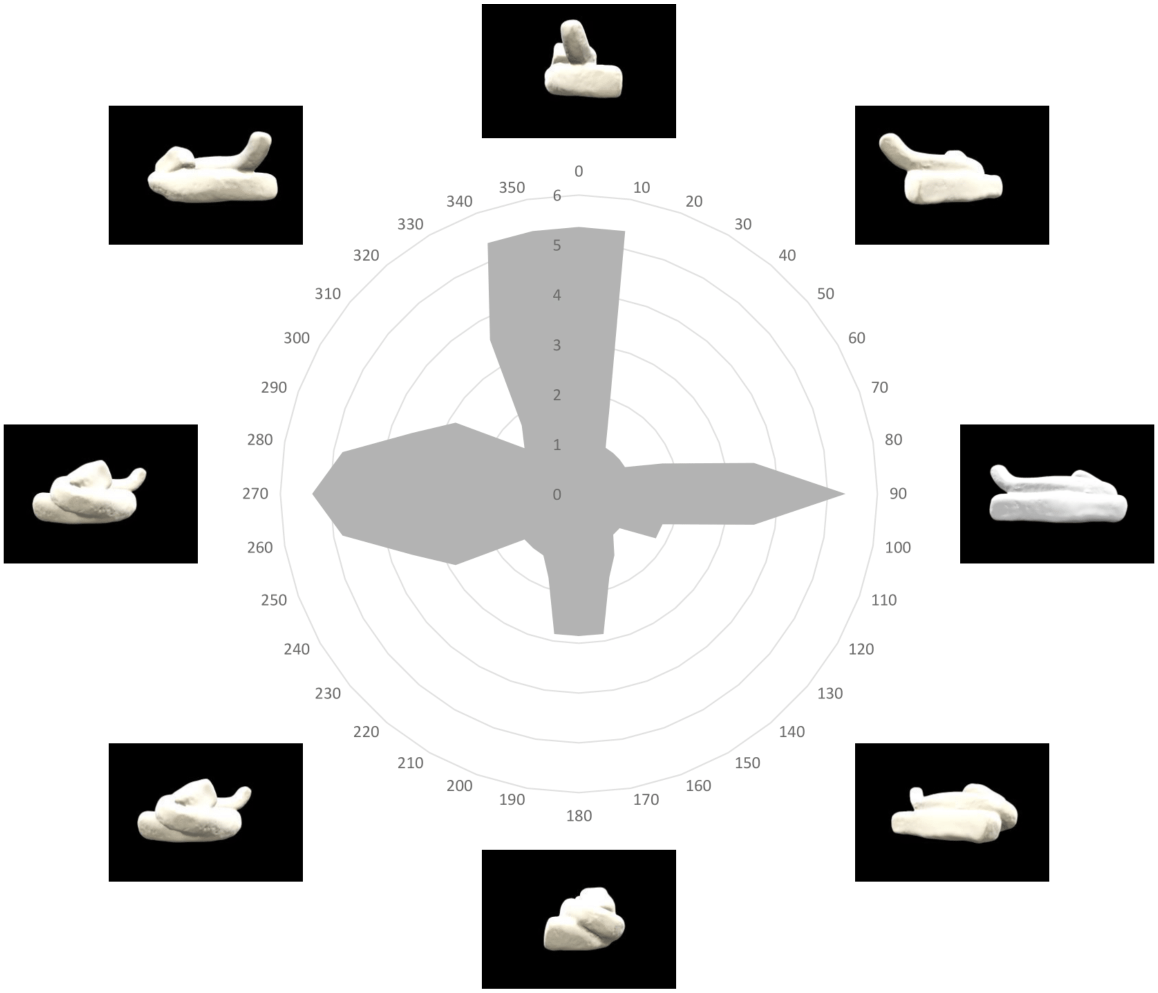

Figure 8: Viewing Patterns During Active Learning Estimated from Harman et al. (1999). Participants in this previous study spent more time observing front and side views of objects than they did observing any other views (Harman et al., 1999). This polar plot depicts the mean percentage of total viewing time during rotation around the vertical axis (corresponding to left-right rotation) as a function of object orientation (°) with the initial orientation shown at 0° (12 o'clock position). The spatial resolution (bin size) of the viewing time calculation was 10°.

Chapter 4

4 Discussion

4.1 Interpretation of Results

The current study tested the hypotheses that active learning of novel objects may improve our ability to discriminate between different views of an object or, alternatively, improve our ability to recognize the same object across multiple views. These effects would manifest as, respectively, a decreased or increased JND for actively learned objects relative to passively learned objects. However, this study found no significant difference in JND between actively learned objects and passively learned objects. Moreover, Bayesian statistical tests found "anecdotal" evidence in favor of the null hypothesis, providing additional support for the conclusion that active learning has no effect on viewpoint discrimination for novel objects.

The aim of the current study was to utilize viewpoint discrimination ability as a measure of the viewpoint dependence of the object representations most likely to be affected by active learning: those involved in object recognition across viewpoints (Harman et al., 1999; James et al., 2001, 2002; Sasaoka et al., 2010). It was assumed that discrimination between viewpoints and recognition across viewpoints would both involve the same object representations. However, this assumption did not take into account the probable existence of multiple object representations in the visual system which enable effective recognition across viewpoints without sacrificing sensitivity to changes in viewpoint (Andresen et al., 2009, Tarr & Hayward, 2017). Case studies of "object orientation agnosia" as well as experiments involving healthy participants suggest that recognition of an object's identity and recognition of an object's orientation are, to some degree, separate processes (Corballis et al., 2007; Fujinaga et al., 2005; Harris et al., 2001, 2020; Turnbull et al., 1996; Valyear et al, 2006). The current study found that active learning does not affect viewpoint discrimination as it does object recognition across changes in viewpoints. This provides further evidence that the two tasks do not rely on the same mechanisms. One potential explanation for the results could be that active and passive learning are equally sufficient for the formation of the neural object representations required for accurate viewpoint discrimination. Alternatively, viewpoint discrimination may not require extensive learning at all. Such hypotheses may be tested in future studies.

4.2 Methodology and Limitations

This study was the first to attempt to assess viewpoint dependence for novel objects using the method of constant stimuli. Even though many other researchers have investigated the effects of active and passive learning, it is important to understand that because of the current study's unique design, its findings cannot be directly compared to those of most previous studies on the subject. For instance, Sasaoka et al. (2010) observed that actively exploring a set of novel wire objects expanded participants' view generalization range. However, because the objects presented during the observation (learning) phase were different from those presented during the generalization (test) phase in their study, the researchers were able to conclude only that active learning facilitated the view matching process. In contrast, the current study examined the effect of active learning by utilizing the same objects in both the learning and test phases. Though these two studies are similar in several respects, a direct comparison is, nonetheless, difficult to make.

When interpreting the current study's results, certain important aspects of its methodology should be taken into account. Past studies that found an effect of active versus passive learning on object recognition utilized a greater number and variety of object stimuli, many of which were "solid" and not necessarily symmetrical (Harman et al., 1999; James et al., 2001; Meijer & van der Lubbe, 2011; Suzuki et al., 2019). This study, however, utilized wire objects that were bilaterally symmetrical. As such, the potential complexity and diversity of the objects were greatly restricted, and the experimenter was only able to create four that were suitable. Relatively little rotation was required to ascertain the wire objects' three-dimensional structures because very few, if any, of their components were ever occluded like those of "solid" objects are. Participants were also given 60 seconds to learn each object, while participants in previous studies were provided with only 20 - 30 seconds (Harman et al., 1999; James et al., 2001, 2002; Meijer & van der Lubbe, 2011; Sasaoka et al., 2010). Furthermore, the use of the method of constant stimuli necessitated over 100 presentations of each object: an additional 22 seconds for participants to view it during the test phase. The effect of active learning can be modulated by various factors, one of which could potentially be task difficulty. One study found that only participants with low visuo-spatial ability (VSA) benefited from active learning whereas participants with middle or high VSA did not (Meijer & van den Broek, 2010). If the facilitative effect of active learning is, in fact, only apparent when the task at hand is sufficiently challenging or demanding, then the simplicity of the current study's stimuli, the additional learning time with which participants were provided, and the sheer number of stimulus presentations during the task may have caused a ceiling effect, eliminating any "active learning advantage" that existed. Because the study also lacked a "no learning" control condition, it was unfortunately impossible to determine whether learning itself, active or passive, had any impact on task performance. Alternatively, active learning could have had an effect that just could not be detected. The range of stimuli viewpoints and the step size (+/- 15° in steps of 3°) may have been too large to estimate the JND with precision.

Participants exhibited preferences for frontal object views over others, which is consistent with the unequal distribution of viewing time observed in past studies while participants actively learned novel objects (Harman et al., 1999; James et al., 2001, 2002; Sasaoka et al., 2010). The specific patterns of exploration differ across studies, though: while participants in the current study spent more time examining frontal views of the objects, Harman et al. (1999) and James et al. (2002) found that participants focused on side views as well as frontal views of objects. Participants' viewing patterns may have been influenced by the unique characteristics of the novel objects themselves. When a "solid," non-wire object is observed from the front, its sides and back are almost, if not entirely, hidden from view. The front, sides, and back of the object must therefore be individually inspected to determine its appearance as a whole. Such thorough scrutiny of each side is not required when studying a wire object because the vast majority of its components are already visible from the front. As a result, the learning experience was simplified, potentially impacting the JND outcome of this study. Moreover, one must note that, at the start of every learning session in the current study, the front of the object was invariably oriented directly toward the participant. Thus, the apparent focus on frontal views exhibited by the current study's participants may, in fact, have been nothing more than an artifact. Harman et al. (1999) avoided an “artifact of starting position” by initially orienting each object such that it required a rotation before it was upright.

Other limitations of this study include its relatively small, non-representative convenience sample. Participants were recruited through the OurBrainsCAN: University of Western Ontario’s Cognitive Neuroscience Research Registry and through posters located on campus, so it can be assumed that many (if not most) of the participants were university students or employees. In addition, because only right-handed adults between the ages of 17 and 60 were eligible to participate, the study entirely excluded children, older adults, and the left-handed. These limitations are not unique to the current study: many similar studies within the fields of visual perception and object learning have used samples comparable in size and composition (Fang & He, 2005; Harman et al., 1999; James et al., 2001, 2002; Rice et al., 2007; Valyear et al., 2006). Nonetheless, it should not be assumed that these results can be generalized to the broader population.

While this study had many limitations, one strength lay in its already validated methodology: the tasks used in the learning and test phases were lightly modified versions of those employed successfully by Harman et al. (1999) and Fang and He (2005), respectively.

4.3 Future Directions

This study generated multiple new hypotheses and questions to be addressed in follow-up research. For example, it would be valuable to perform a direct comparison of discrimination between objects (as in studies conducted by Harman et al. (1999) and Sasaoka et al. (2010)) against discrimination between viewpoints for a single object (as in the current study) using the same objects and the same paradigm. This would provide a definitive method of determining whether active learning has genuinely different impacts on the two tasks or whether the different outcomes are due to methodological differences.

Another question that warrants further investigation is whether the nature of objects themselves has an impact on learning and its effect on object recognition. The hypotheses would be that (1) solid objects (like those used by Harman et al. (1999)) in which certain components are occluded by other components invoke different exploration strategies (particularly increased focus on side and back views from which all object components are visible) than do wire-like objects and (2) recognition of solid objects benefits more from active learning than does recognition of wire-like objects.

To test whether both active and passive learning both facilitate discrimination between viewpoints to an equal degree, the current study would be repeated with the addition of a third, “no learning” condition. Decreased performance in the “no learning” condition compared to the learning conditions would provide support for this hypothesis. Conversely, the study may find no difference between any of the three conditions, which would imply that either participants can form mental representations of objects without prior learning or the discrimination task doesn’t require knowledge of an object’s three-dimensional appearance at all. Those hypotheses could be tested by conducting a study in which participants are required, with and without learning, to discriminate between views of a target object as well as between the target object and distractor objects. Unimpaired performance on both tasks in both the learning and no-learning conditions would suggest that learning is not required for the creation of a three-dimensional object representation. This finding would be consistent with the immediate viewpoint invariance observed by Biederman and Gerhardstein (1993). Alternatively, impaired discrimination between target and distractor objects, but not between different views of target objects, in the “no learning” condition relative to the learning condition would suggest that learning is necessary to perform the former task, but not the latter.

In addition to testing the hypotheses generated by the results of the current study, subsequent research could address its limitations by employing a greater number of more complex and diverse novel objects as stimuli and achieving larger sample sizes. Recruiting a sufficient number of participants for a local, face-to-face experiment is often difficult, but online or virtual approaches would enable faster and more effective data collection for multiple studies.

This study investigated the ability to discriminate an object from different views and whether this ability is influenced by active versus passive learning. Other studies of object recognition have investigated viewpoint dependence, the ability to recognize an object regardless of changes in viewpoint. Of course, it is possible that active learning could have no effect on sensitivity to changes in orientation, even though it improves later object recognition. What is needed is measurements of sensitivity to changes in orientation and measurements of the viewpoint dependence of object recognition in the same participants in active and passive learning conditions.

Other methods of quantifying viewpoint dependence should be investigated in the future as well, especially because sensitivity to changes in viewpoint may not necessarily be a sufficiently accurate indicator of viewpoint dependence. In fact, it is possible that active learning has an effect on viewpoint dependence, but not on the ability to discriminate between viewpoints. A later behavioral study could assess the effect of active learning on both viewpoint dependence and discrimination between viewpoints simultaneously. Visual adaptation is one potential measure of viewpoint dependence. Fang and He (2005) found that adaptation to a particular view decreases the sensitivity of the neuronal populations tuned for that view. As a result, the viewer experiences a viewpoint aftereffect, the magnitude of which is positively correlated with viewpoint tuning width.

Although such psychophysical techniques are simple and fast to use, they provide only an indirect way of gaining insight into activity within the brain itself. Future studies could utilize fMRI to more directly examine neural representations of actively and passively learned objects. Researchers have already found evidence suggesting that sensorimotor experiences during active learning influence the formation of neural object representations. Study participants learned the functions of unfamiliar objects, actively manipulating one set of the objects and visually exploring (viewing without being allowed to touch) another set. Active manipulation elicited greater post-training activity in fronto-parietal areas, particularly the left inferior/middle frontal gyrus, than did visual exploration (Bellebaum et al., 2013). However, the effect of active learning on activity in object-selective visual areas such as the lateral occipital complex (LOC), has not yet been extensively investigated. The LOC, an area of the occipital lobe which responds preferentially to objects, may be a region in which the transition from low-level object representations to more abstract object representations takes place (Malach et al., 1995;

Vernon et al., 2016). Representational similarity analysis (RSA), a type of multivariate pattern analysis (MVPA) would be a highly effective method of examining these object representations, as it reveals subtle differences in patterns of activity that univariate analyses cannot detect. Activity patterns in the LOC seem to encode information about not only low-level image properties, but about object shape and category (Eger et al., 2008; Haushofer et al., 2008; Rice et al., 2014). Therefore, an inspection of the patterns elicited by an object after learning could reveal important differences between actively and passively learned object representations.

One more question that remains unanswered is how the learning of real, physical objects may differ from the learning of virtual objects on a computer screen. Human infants have demonstrated a preference for real objects over pictures of objects, and their knowledge of familiar size for real objects does not seem to generalize to pictures of the objects (Gerhard et al., 2021; Sensoy et al., 2021). Such a distinction between real objects and images of objects also exists in adults, who display better recall and recognition performance for real objects than for photographs or drawings of objects (Snow et al., 2014). An fMRI study even found that repetition effects that were robust for two-dimensional images of objects were attenuated or non-existent for real objects, suggesting that real objects and images of objects may in fact be processed using different neural mechanisms (Snow et al., 2011). Although the current study did not find any difference in viewpoint dependence for actively and passively learned virtual objects, its results cannot necessarily be generalized to real objects.

While humans take for granted their ability to successfully identify the thousands of objects that fill this world, creating computer algorithms that can do the same is still a formidable challenge. Computer vision is a rapidly growing field with numerous practical applications. The development of self-driving cars, a classic staple of science fiction that scientists and engineers are working to make a reality, is one example (Cervera-Uribe & Méndez-Monroy, 2022; Lu et al., 2022; Zhou et al., 2020). Computer vision algorithms can also assist visually impaired people with everyday tasks (navigating city streets, grocery shopping, etc.) and even identify handguns in airport baggage inspection (Franco et al., 2017; Kumar et al., 2022; Tapu et al., 2017). Creating and then training neural networks can be a difficult and time-consuming process that often requires massive datasets. Thus, developing more efficient methods of teaching them to recognize three-dimensional objects is a priority. Researchers have already built robots that, like humans, are capable of learning to recognize objects through active manipulation (Ivaldi et al., 2014; Schiebener et al., 2013). However, further research should be done to determine which aspects of object recognition are facilitated by active learning as well as how and why such facilitation takes place. A greater understanding of human object perception would expedite the development of better strategies for training neural network models to accurately process images.

References

Andresen, D. R., Vinberg, J., & Grill-Spector, K. (2009). The representation of object viewpoint in human visual cortex. *NeuroImage*, *45*, 522–536. https://doi.org/10.1016/j.neuroimage.2008.11.009

Antrilli, N. K., & Wang, S. (2016). Visual cues generated during action facilitate 14-month-old infants' mental rotation. *Journal of Cognition and Development*, *17*(3), 418–429. https://doi.org/10.1080/15248372.2015.1058262

Bellebaum, C., Tettamanti, M., Marchetta, E., della Rosa, P., Rizzo, G., Daum, I., & Cappa, S. F. (2013). Neural representations of unfamiliar objects are modulated by sensorimotor experience. *Cortex*, *49*, 1110–1125. https://doi.org/10.1016/j.cortex.2012.03.023

Biederman, I., & Bar, M. (1999). One-shot viewpoint invariance in matching novel objects. *Vision Research*, *39*, 2885–2899. https://doi.org/10.1016/S0042-6989(98)00309-5

Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance*, *19*(6), 1162–1182. https://doi.org/10.1037/0096-1523.19.6.1162

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433-436.

Cervera-Uribe, A. A., & Méndez-Monroy, P. E. (2022). U19-Net: A deep learning approach for obstacle detection in self-driving cars. *Soft Computing*, *26*, 5195–5207. https://doi.org/10.1007/s00500-022-06980-6

Chrastil, E. R., & Warren, W. H. (2013). Active and passive spatial learning in human navigation: Acquisition of survey knowledge. *Journal of Experimental Psychology: Learning Memory and Cognition*, *39*(5), 1520–1537. https://doi.org/10.1037/a0032382

Chrastil, E. R., & Warren, W. H. (2015). Active and passive spatial learning in human navigation: Acquisition of graph knowledge. *Journal of Experimental Psychology: Learning Memory and Cognition*, *41*(4), 1162–1178. https://doi.org/10.1037/xlm0000082

Corballis, M. C., Armstrong, C., & Zhu, Z. (2007). Orientation unbound: Dissociation of identity and orientation under rapid serial visual presentation. *Memory and Cognition*, *35*(6), 1518–1525.

Craddock, M., Martinovic, J., & Lawson, R. (2011). An advantage for active versus passive aperture-viewing in visual object recognition. *Perception*, *40*, 1154–1163. https://doi.org/10.1068/p6974

Eger, E., Ashburner, J., Haynes, J.-D., Dolan, R. J., & Rees, G. (2008). fMRI activity patterns in human LOC carry information about object exemplars within category. *Journal of Cognitive Neuroscience*, *20*(2), 356–370. https://doi.org/10.1162/jocn.2008.20.2.356

Eger, E., Schweinberger, S. R., Dolan, R. J., & Henson, R. N. (2005). Familiarity enhances invariance of face representations in human ventral visual cortex: fMRI evidence. *NeuroImage*, *26*, 1128–1139. https://doi.org/10.1016/j.neuroimage.2005.03.010

Ernst, M. O., Lange, C., & Newell, F. N. (2007). Multisensory recognition of actively explored objects. *Canadian Journal of Experimental Psychology*, *61*(3), 242–253. https://doi.org/10.1037/cjep2007025

Ewbank, M. P., & Andrews, T. J. (2008). Differential sensitivity for viewpoint between familiar and unfamiliar faces in human visual cortex. *NeuroImage*, *40*, 1857–1870. https://doi.org/10.1016/j.neuroimage.2008.01.049

Fang, F., & He, S. (2005). Viewer-centered object representation in the human visual system revealed by viewpoint aftereffects. *Neuron*, *45*, 793–800. https://doi.org/10.1016/j.neuron.2005.01.037

Franco, A., Maltoni, D., & Papi, S. (2017). Grocery product detection and recognition. *Expert Systems with Applications*, *81*, 163–176. https://doi.org/10.1016/j.eswa.2017.02.050

Frick, A., & Wang, S. H. (2014). Mental spatial transformations in 14- and 16-month-old infants: Effects of action and observational experience. *Child Development*, *85*(1), 278–293. https://doi.org/10.1111/cdev.12116

Fujinaga, N., Muramatsu, T., Ogano, M., & Kato, M. (2005). A 3-year follow-up study of “orientation agnosia.” *Neuropsychologia*, *43*, 1222–1226. https://doi.org/10.1016/j.neuropsychologia.2004.10.010

Gerhard, T. M., Culham, J. C., & Schwarzer, G. (2021). Manual exploration of objects is related to 7-month-old infants’ visual preference for real objects. *Infant Behavior and Development*, *62*, 1–9. https://doi.org/10.1016/j.infbeh.2020.101512

Harman, K. L., Humphrey, G. K., & Goodale, M. A. (1999). Active manual control of object views facilitates visual recognition. *Current Biology*, *9*(22), 1315–1318. https://doi.org/10.1016/S0960-9822(00)80053-6

Harris, I. M., Harris, J. A., & Caine, D. (2001). Object orientation agnosia: A failure to find the axis? *Journal of Cognitive Neuroscience*, *13*(6), 800–812. http://mitprc.silverchair.com/jocn/article-pdf/13/6/800/1759150/08989290152541467.pdf

Harris, I. M., Harris, J. A., & Corballis, M. C. (2020). Binding identity and orientation in object recognition. *Attention, Perception, and Psychophysics*, *82*, 153–167. https://doi.org/10.3758/s13414-019-01677-9

Haushofer, J., Livingstone, M. S., & Kanwisher, N. (2008). Multivariate patterns in object-selective cortex dissociate perceptual and physical shape similarity. *PLoS Biology*, *6*(7), 1459–1467. https://doi.org/10.1371/journal.pbio.0060187

Held, R., & Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, *56*(5), 872–876. https://doi.org/10.1037/h0040546

Iggena, D., Klein, C., Garthe, A., Winter, Y., Kempermann, G., & Steiner, B. (2015). Only watching others making their experiences is insufficient to enhance adult neurogenesis and water maze performance in mice. *Scientific Reports*, *5*, 1–10. https://doi.org/10.1038/srep14141

Ivaldi, S., Nguyen, S. M., Lyubova, N., Droniou, A., Padois, V., Filliat, D., Oudeyer, P.-Y., & Sigaud, O. (2014). Object learning through active exploration. *IEEE Transactions on Autonomous Mental Development*, *6*(1), 56–72. https://doi.org/10.1109/TAMD.2013.2280614

James, K. H., Humphrey, G. K., & Goodale, M. A. (2001). Manipulating and recognizing virtual objects: Where the action is. *Canadian Journal of Experimental Psychology*, *55*(2), 113–122. https://doi.org/10.1037/h0087358

James, K. H., Humphrey, G. K., Vilis, T., Corrie, B., Baddour, R., & Goodale, M. A. (2002). “Active” and “passive” learning of three-dimensional object structure within an immersive virtual reality environment. *Behavior Research Methods, Instruments, & Computers*, *34*(3), 383–390.

James, K. H., Jones, S. S., Swain, S., Pereira, A., & Smith, L. B. (2014). Some views are better than others: Evidence for a visual bias in object views self-generated by toddlers. *Developmental Science*, *17*(3), 338–351. https://doi.org/10.1111/desc.12124

The jamovi project (2021). Jamovi (Version 1.6) [Computer Software]. Retrieved from https://www.jamovi.org

Kumar, A., Saini, T., Pandey, P. B., Agarwal, A., Agrawal, A., & Agarwal, B. (2022). Vision-based outdoor navigation of self-driving car using lane detection.

International Journal of Information Technology, *14*(1), 215–227. https://doi.org/10.1007/s41870-021-00747-2

Lawson, R., & Humphreys, G. W. (1996). View specificity in object processing: Evidence from picture matching. *Journal of Experimental Psychology: Human Perception and Performance*, *22*(2), 395–416. https://doi.org/10.1037/0096-1523.22.2.395

Lee, H., & Wallraven, C. (2013). Exploiting object constancy: Effects of active exploration and shape morphing on similarity judgments of novel objects. *Experimental Brain Research*, *225*, 277–289. https://doi.org/10.1007/s00221-012-3368-1

Libertus, K., Joh, A. S., & Needham, A. W. (2016). Motor training at 3 months affects object exploration 12 months later. *Developmental Science*, *19*(6), 1058–1066. https://doi.org/10.1111/desc.12370

Libertus, K., & Needham, A. (2010). Teach to reach: The effects of active vs. passive reaching experiences on action and perception. *Vision Research*, *50*, 2750–2757. https://doi.org/10.1016/j.visres.2010.09.001

Liu, C. H., Ward, J., & Markall, H. (2007). The role of active exploration of 3D face stimuli on recognition memory of facial information. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(4), 895–904. https://doi.org/10.1037/0096-1523.33.4.895

Logothetis, N. K., & Pauls, J. (1995). Psychophysical and physiological evidence for viewer-centered object representations in the primate. *Cerebral Cortex*, *5*(3), 270–288. https://doi.org/10.1093/cercor/5.3.270

Logothetis, N. K., Pauls, J., Bülthoff, H. H., & Poggio, T. (1994). View-dependent object recognition by monkeys. *Current Biology*, *4*(5), 401–414. https://doi.org/10.1016/S0960-9822(00)00089-0

Lu, E. H.-C., Gozdzikiewicz, M., Chang, K.-H., & Ciou, J.-M. (2022). A hierarchical approach for traffic sign recognition based on shape detection and image classification. *Sensors*, *22*, 4768. https://doi.org/10.3390/s22134768

Malach, R., Reppas, J. B., Benson, R. R., Kwong, K. K., Jiang, H., Kennedy, W. A., Ledden, P. J., Brady, T. J., Rosen, B. R., & Tootell, R. B. H. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *92*, 8135–8139. https://doi.org/10.1073/pnas.92.18.8135

Meijer, F., & van den Broek, E. L. (2010). Representing 3D virtual objects: Interaction between visuo-spatial ability and type of exploration. *Vision Research*, *50*, 630–635. https://doi.org/10.1016/j.visres.2010.01.016

Meijer, F., & van der Lubbe, R. H. J. (2011). Active exploration improves perceptual sensitivity for virtual 3D objects in visual recognition tasks. *Vision Research*, *51*, 2431–2439. https://doi.org/10.1016/j.visres.2011.09.013

Möhring, W., & Frick, A. (2013). Touching up mental rotation: Effects of manual experience on 6-month-old infants' mental object rotation. *Child Development*, *84*(5), 1554–1565. https://doi.org/10.1111/cdev.12065

Orlov, T., & Zohary, E. (2018). Object representations in human visual cortex formed through temporal integration of dynamic partial shape views. *Journal of Neuroscience*, *38*(3), 659–678. https://doi.org/10.1523/JNEUROSCI.1318-17.2017

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437-442.

Peruch, P., Vercher, J.-L., & Gauthier, G. M. (1995). Acquisition of spatial knowledge through visual exploration of simulated environments. *Ecological Psychology*, *7*(1), 1–20.

Rice, G. E., Watson, D. M., Hartley, T., & Andrews, T. J. (2014). Low-level image properties of visual objects predict patterns of neural response across categoryselective regions of the ventral visual pathway. *Journal of Neuroscience*, *34*(26), 8837–8844. https://doi.org/10.1523/JNEUROSCI.5265-13.2014

Rice, N. J., Valyear, K. F., Goodale, M. A., Milner, A. D., & Culham, J. C. (2007). Orientation sensitivity to graspable objects: An fMRI adaptation study. *NeuroImage*, *36*, T87–T93. https://doi.org/10.1016/j.neuroimage.2007.03.032

Sasaoka, T., Asakura, N., & Kawahara, T. (2010). Effect of active exploration of 3-D object views on the view-matching process in object recognition. *Perception*, *39*, 289–308. https://doi.org/10.1068/p5721

Schiebener, D., Morimoto, J., Asfour, T., & Ude, A. (2013). Integrating visual perception and manipulation for autonomous learning of object representations. *Adaptive Behavior*, *21*(5), 328–345. https://doi.org/10.1177/1059712313484502

Sensoy, Ö., Culham, J. C., & Schwarzer, G. (2021). The advantage of real objects over matched pictures in infants' processing of the familiar size of objects. *Infant and Child Development*, *30*, 1–17. https://doi.org/10.1002/icd.2234

Snow, J. C., Pettypiece, C. E., McAdam, T. D., McLean, A. D., Stroman, P. W., Goodale, M. A., & Culham, J. C. (2011). Bringing the real world into the fMRI scanner: Repetition effects for pictures versus real objects. *Scientific Reports*, *1*(130), 1–10. https://doi.org/10.1038/srep00130

Snow, J. C., Skiba, R. M., Coleman, T. L., & Berryhill, M. E. (2014). Real-world objects are more memorable than photographs of objects. *Frontiers in Human Neuroscience*, *8*, 1–11. https://doi.org/10.3389/fnhum.2014.00837

Suzuki, K., Schwartzman, D. J., Augusto, R., & Seth, A. K. (2019). Sensorimotor contingency modulates breakthrough of virtual 3D objects during a breaking continuous flash suppression paradigm. *Cognition*, *187*, 95–107. https://doi.org/10.1016/j.cognition.2019.03.003

Tapu, R., Mocanu, B., & Zaharia, T. (2017). A computer vision-based perception system for visually impaired. *Multimedia Tools and Applications*, *76*, 11771–11807. https://doi.org/10.1007/s11042-016-3617-6

Tarr, M. J., & Hayward, W. G. (2017). The concurrent encoding of viewpoint-invariant and viewpoint-dependent information in visual object recognition. *Visual Cognition*, *25*(1–3), 100–121. https://doi.org/10.1080/13506285.2017.1324933

Tian, M., & Grill-Spector, K. (2015). Spatiotemporal information during unsupervised learning enhances viewpoint invariant object recognition. *Journal of Vision*, *15*(6), 1–13. https://doi.org/10.1167/15.6.7

Turnbull, O. H., Beschin, N., & della Sala, S. (1996). Agnosia for object orientation: Implications for theories of object recognition. *Neuropsychologia*, *35*(2), 153–163.

Valyear, K. F., Culham, J. C., Sharif, N., Westwood, D., & Goodale, M. A. (2006). A double dissociation between sensitivity to changes in object identity and object orientation in the ventral and dorsal visual streams: A human fMRI study. *Neuropsychologia*, *44*, 218–228. https://doi.org/10.1016/j.neuropsychologia.2005.05.004

Vernon, R. J. W., Gouws, A. D., Lawrence, S. J. D., Wade, A. R., & Morland, A. B. (2016). Multivariate patterns in the human object-processing pathway reveal a shift from retinotopic to shape curvature representations in lateral occipital areas, LO-1 and LO-2. *Journal of Neuroscience*, *36*(21), 5763–5774. https://doi.org/10.1523/JNEUROSCI.3603-15.2016

Wallis, G. (2002). The role of object motion in forging long-term representations of objects. *Visual Cognition*, *9*(1–2), 233–247. https://doi.org/10.1080/13506280143000412

Wallis, G., & Bülthoff, H. H. (2001). Effects of temporal association on recognition memory. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(8), 4800–4804. https://doi.org/10.1073/pnas.071028598

Wiesen, S. E., Watkins, R. M., & Needham, A. W. (2016). Active motor training has long-term effects on infants' object exploration. *Frontiers in Psychology*, *7*, 1–10. https://doi.org/10.3389/fpsyg.2016.00599

Zhou, S., Chen, Y., Li, X., & Sanyal, A. (2020). Deep SCNN-based real-time object detection for self-driving vehicles using LiDAR temporal data. *IEEE Access, 8*, 76903–76912. https://doi.org/10.1109/ACCESS.2020.2990416

Curriculum Vitae

Name:	Cassandra Bacher
Post-secondary Education and Degrees:	Knox College Galesburg, Illinois, United States of America September 2014 – June 2018 B.A.
	The University of Western Ontario London, Ontario, Canada September 2020 – present MSc. Program (in progress)
Related Work Experience:	Career Shadowing Klinikum Frankfurt a.d. Oder, Frankfurt a.d. Oder, Germany June – July 2014
	Summer Intern Loyola University Chicago, Maywood, Illinois, United States June – July 2015
	Summer Undergraduate Research Fellowship Loyola University Chicago, Maywood, Illinois, United States June – August 2017
	Animal Caretaker Knox College, Galesburg, Illinois, United States September 2017 – May 2018
	Research Assistant Loyola University Chicago, Maywood, Illinois, United States June 2018 – March 2020
	Graduate Teaching Assistant The University of Western Ontario, London, Ontario, Canada January – April 2021
	Graduate Teaching Assistant The University of Western Ontario, London, Ontario, Canada September 2021 – April 2022

Conference Posters:

Baig, N., Kouta, A., Siddiqui, F., Farooqui, A., Hoppensteadt, D., Jeske, W., Iqbal, O., Bacher, C., Krishnappan, S., & Fareed, J. (2019, December 7-10). *Potency equated porcine and bovine mucosal heparin are bioequivalent in terms of biochemical and*

pharmacological effects [Poster Presentation]. American Society of Hematology Annual Meeting, Orlando, FL, United States.

Walenga, J., Jeske, W., Bertini, S., Risi, G., Sung, M., Farooqui, A., Bacher, C., Hoppensteadt, D., Iqbal, O., Lewis, B., & Bakhos, M. (2019, December 7-10). *Bovine heparin demonstrates the same interaction with HIT antibodies as porcine heparin* [Poster Presentation]. American Society of Hematology Annual Meeting, Orlando, FL, United States.

Farooqui, A., Hoppensteadt, D., Iqbal, O., Bacher, C., Walenga, J., Kouta, A., Jeske, W., Lewis, B., & Fareed, J. (2019, December 7-10). *Comparative studies on the interaction of unfractionated heparin and sulodexide with functional anti-heparin platelet factor 4 antibodies* [Poster Presentation]. American Society of Hematology Annual Meeting, Orlando, FL, United States.

Bacher, C., Farooqui, A., Siddiqui, F., Daravathi, B., Hoppensteadt, D., Iqbal, O., Fareed, J., George, M., & van Thiel, D. (2020, April 4-7). *Liver diseases contribute to functional platelet aggregation defects* [Poster Abstract]. Experimental Biology, San Diego, CA, United States.

Bacher, C., Farooqui, A., Siddiqui, F., Daravati, B., Hoppensteadt, D., Iqbal, O., Fareed, J., George, M., & van Thiel, D. (2021, February). *Heptatic disorders contribute to platelet dysfunction* [Poster Abstract]. Annual Meeting of the Society of Thrombosis and Hemostasis Research.