

1-1-2013

Using Multivariate Pattern Analysis to Investigate the Neural Representation of Concepts With Visual and Haptic Features

Laura Bradshaw Baucom
University of South Carolina - Columbia

Follow this and additional works at: <http://scholarcommons.sc.edu/etd>

Recommended Citation

Baucom, L. B.(2013). *Using Multivariate Pattern Analysis to Investigate the Neural Representation of Concepts With Visual and Haptic Features*. (Doctoral dissertation). Retrieved from <http://scholarcommons.sc.edu/etd/2540>

This Open Access Dissertation is brought to you for free and open access by Scholar Commons. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of Scholar Commons. For more information, please contact SCHOLARC@mailbox.sc.edu.

USING MULTIVARIATE PATTERN ANALYSIS TO INVESTIGATE
THE NEURAL REPRESENTATION OF CONCEPTS
WITH VISUAL AND HAPTIC FEATURES

by

Laura Bradshaw Baucom

Bachelor of Science
University of North Carolina at Charlotte, 2004

Submitted in Partial Fulfillment of the Requirements

For the Degree of Doctor of Philosophy in

Experimental Psychology

College of Arts and Sciences

University of South Carolina

2013

Accepted by:

Svetlana Shinkareva, Major Professor

Douglas Wedell, Committee Member

Christopher Rorden, Committee Member

John Rose, Committee Member

Lacy Ford, Vice Provost and Dean of Graduate Studies

© Copyright by Laura Bradshaw Baucom, 2013
All Rights Reserved.

DEDICATION

To my son, Maxwell. May he always know that he can accomplish whatever his heart desires, and that I will always be there to support him in his endeavors.

ACKNOWLEDGEMENTS

I would like to express my deep appreciation and gratitude to my advisor, Dr. Svetlana Shinkareva, for her steadfast guidance and mentorship over the course of my graduate studies. Her dedication to teach and inspire future scientists is matched by her enthusiasm for research. It has been an honor to work with her. Additionally, I would like to thank my colleagues Dr. Jing Wang, Matthew Facciani, Neha Jaggi, and Arjamand Sami for their help with data collection.

I would also like to thank my committee members, Drs. Douglas Wedell, Chris Rorden, and John Rose for the insightful suggestions, thought-provoking commentary, and encouragement.

Finally, I would like to acknowledge the people who hold my world together and have contributed to all my achievements, my family. I owe many thanks to my husband, Jeremy, for supporting me and being a great father to our son while I pursued this final degree. Additionally, I owe everything to my parents, John and Sandra, who always reinforced the importance of education and ensured I had every opportunity to succeed.

ABSTRACT

A fundamental debate in cognitive neuroscience concerns how conceptual knowledge is represented in the brain. Over the past decade, cognitive theorists have adopted explanations that suggest cognition is rooted in perception and action. This is called the embodiment hypothesis. Theories of conceptual representation differ in the degree to which representations are embodied, from those which suggest conceptual representation requires no involvement of sensory and motor systems to those which suggest it is entirely dependent upon them. This work investigated how the brain represents concepts that are defined by their visual and haptic features using novel multivariate approaches to the analysis of functional magnetic resonance imaging (fMRI) data.

A behavioral study replicated a perceptual phenomenon, known as the tactile disadvantage, demonstrating that that verifying the properties of concepts with haptic features takes significantly longer than verifying the properties of concepts with visual features. This study suggested that processing the perceptual properties of concepts likely recruits the same processes involved in perception. A neuroimaging study using the same paradigm showed that processing concepts with visual and haptic features elicits activity in bimodal object-selective regions, such as the fusiform gyrus (FG) and the lateral occipitotemporal cortex (LOC). Multivariate pattern analysis (MVPA) was successful at identifying whether a concept had perceptual or abstract features from patterns of brain activity located in functionally-defined object-selective and general

perceptual regions in addition to the whole brain. The conceptual representation was also consistent across participants. Finally, the functional networks for verifying the properties of concepts with visual and haptic features were highly overlapping but showed differing patterns of connectivity with the occipitotemporal cortex across people.

Several conclusions can be drawn from this work, which provide insight into the nature of the neural representation of concepts with perceptual features. The neural representation of concepts with visual and haptic features involves brain regions which underlie general visual and haptic perception as well visual and haptic perception of objects. These brain regions interact differently based on the type of perceptual feature a concept possesses. Additionally, the neural representation of concepts with visual and haptic features is distributed across the whole brain and is consistent across people. The results of this work provide partial support for weak and strong embodiment theories, but further studies are necessary to determine whether sensory systems are required for conceptual representation.

TABLE OF CONTENTS

DEDICATION	iii
ACKNOWLEDGEMENTS.....	iv
ABSTRACT	v
LIST OF TABLES	viii
LIST OF FIGURES	ix
LIST OF ABBREVIATIONS.....	xiii
CHAPTER 1 INTRODUCTION.....	1
1.1 THEORIES OF EMBODIED COGNITION	2
CHAPTER 2 VISUAL AND HAPTIC OBJECT PERCEPTION.....	19
CHAPTER 3 APPROACHES TO THE STUDY OF CONCEPTUAL REPRESENTATION.....	25
3.1 UNIVARIATE VS. PATTERN-BASED APPROACHES	25
3.2 UNIVARIATE APPROACHES TO THE STUDY OF CONCEPTUAL REPRESENTATION	39
3.3 PATTERN-BASED APPROACHES TO THE STUDY OF CONCEPTUAL REPRESENTATION	42
3.4 GOALS OF THE CURRENT WORK.....	44
CHAPTER 4 BEHAVIORAL EXPERIMENT 1.....	46
4.1 PURPOSE.....	46
4.2 MATERIALS & METHODS.....	46
4.3 RESULTS.....	49

4.4 SUMMARY	49
CHAPTER 5 BEHAVIORAL EXPERIMENT 2	51
5.1 PURPOSE.....	51
5.2 MATERIALS & METHODS.....	51
5.3 RESULTS	53
5.4 SUMMARY	55
CHAPTER 6 FUNCTIONAL LOCALIZER	56
6.1 PURPOSE.....	56
6.2 MATERIALS & METHODS.....	56
6.3 IMAGE ACQUISITION & PREPROCESSING	60
6.4 RESULTS	61
6.5 SUMMARY	66
CHAPTER 7 MAIN EXPERIMENT	69
7.1 PURPOSE.....	69
7.2 MATERIALS & METHODS.....	70
7.3 FMRI IMAGE ACQUISITION.....	72
7.4 DATA PROCESSING & ANALYSIS	73
7.5 UNIVARIATE ANALYSIS RESULTS	79
7.6 PATTERN CLASSIFICATION RESULTS	82
7.7 SUMMARY	98
CHAPTER 8 GENERAL DISCUSSION.....	100
8.1 SUMMARY & IMPLICATIONS	100
8.2 FUTURE DIRECTIONS	111

8.3 MERIT & CONTRIBUTION.....	113
REFERENCES	114
APPENDIX A – FUNCTIONAL LOCALIZER STIMULI	128

LIST OF TABLES

Table 5.1 Reaction times (in ms) for property verification task	54
Table 6.1 Brain regions displaying significant activation differences in functional localizer	61
Table 7.1 Brain regions displaying significant activation differences in main experiment	79

LIST OF FIGURES

Figure 2.1 The visual and haptic systems converge at the IPS LOC and FG.....	21
Figure 4.1 Experimental paradigm for behavioral experiment 1	48
Figure 4.2 Mean reaction times for verifying concepts with visual and haptic features.....	49
Figure 5.1 Experimental paradigm for behavioral experiment 2.....	54
Figure 6.1 Experimental paradigm for functional localizer scan.....	58
Figure 6.2 Setup of the scanner suite during the functional localizer.....	59
Figure 6.3 Activation for visual objects.....	63
Figure 6.4 Activation for haptic objects	64
Figure 6.5 Activation for visual or haptic objects.....	64
Figure 6.6 Activation for visual objects and textures	65
Figure 6.7 Activation for haptic objects and textures	66
Figure 7.1 Experimental paradigm for main experiment.....	73
Figure 7.2 Activation for visual and haptic features of concepts	80
Figure 7.3 Activation for visual features of concepts.....	81
Figure 7.4 Activation for haptic features of concepts.....	82
Figure 7.5 Accuracies for classification within regions selective for the visual and haptic features of objects.....	84
Figure 7.6 Correlation between VVIQ scores and classification accuracies	85

Figure 7.7 Accuracies for classifying within regions responsive to processing general visual, haptic, and visual and haptic perceptual features	87
Figure 7.8 Consistency of participant classification accuracies across classification problems and perceptual regions measured by correlation	88
Figure 7.9 Within-participant accuracies for classifying visual vs. haptic vs. abstract features from whole brain patterns of activity	89
Figure 7.10 Participant confusion matrices for classifying visual vs. haptic vs. abstract features from whole brain patterns of activity	90
Figure 7.11 Within-participant accuracies for classifying visual vs. abstract features from whole brain patterns of activity	91
Figure 7.12 Within-participant accuracies for classifying haptic vs. abstract features from whole brain patterns of activity	91
Figure 7.13 Within-participant accuracies for classifying visual and haptic vs. abstract features from whole brain patterns of activity	92
Figure 7.14 Cross-participant accuracies for classifying visual vs. abstract features from whole brain patterns of activity	94
Figure 7.15 Cross-participant accuracies for classifying haptic vs. abstract features from whole brain patterns of activity	94
Figure 7.16 Cross-participant accuracies for classifying visual and haptic vs. abstract features from whole brain patterns of activity	95
Figure 7.17 Thresholded probability maps of the most informative voxels consistently identified for each cross-participant classification problem	96
Figure 7.18 Condition-specific connectivity of all voxels with the occipitotemporal cortex	97
Figure 8.1 Valence ratings of stimuli from each experimental condition	109

LIST OF ABBREVIATIONS

BA.....	Brodmann Area
BOLD.....	Blood Oxygen Level Dependent
EPI.....	Echo Planar Imaging
FG.....	Fusiform Gyrus
fMRI.....	Functional Magnetic Resonance Imaging
GLM.....	General Linear Model
HO.....	Haptic Objects
HRF.....	Hemodynamic Response Function
HT.....	Haptic Textures
IPS.....	Intraparietal Sulcus
LOC.....	Lateral Occipitotemporal Cortex
LOOCV.....	Leave-one-out Cross-validation
MEPs.....	Motor Evoked Potentials
MNI.....	Montreal Neurological Institute
MVPA.....	Multivariate Pattern Analysis
PET.....	Positron Emission Tomography
PPC.....	Posterior Parietal Cortex
PSC.....	Percent Signal Change
ROI.....	Region of Interest
TE.....	Echo Time

TMSTranscranial Magnetic Stimulation
TRRelaxation Time
VO..... Visual Objects
VT..... Visual Textures
VVIQ..... Vividness of Visual Imagery Questionnaire

CHAPTER 1

INTRODUCTION

A fundamental debate in cognitive neuroscience concerns how conceptual knowledge is represented in the brain. Concepts are part and parcel of human cognition, as they serve a central role in various cognitive functions including thought and reasoning, language comprehension and production, action planning, and object recognition (Humphreys, Riddoch, & Quinlan, 1988; Kiefer & Pulvermüller, 2012; Solomon, Medin, & Lynch, 1999). Concepts are essential for human information processing, because they provide a link between action and perception. That is, they help to bridge the information gleaned from the environment through perception and the information dispersed to the environment through action (Kiefer & Pulvermüller, 2012).

A concept is a mental representation that integrates an individual's past sensory and motor experiences with his environment in order to categorize and provide information. For example, the concept "car" might include that a car is a mode of transportation for carrying people, has four wheels, seats, and must be steered with a wheel by a driver. While an individual encounters a variety of cars in his lifetime, the concept of "car" is a generalization across all the cars he has experienced. This aids the individual in identifying and responding appropriately to future instances of cars.

While most agree what constitutes a concept, how concepts are represented remains an important question. Prior to the Twentieth Century cognitive theorists

suggested that cognition was grounded in perception. That is, conceptual knowledge was believed to be represented in the same manner as mental images. Following the cognitive revolution, advancements in computer science, artificial intelligence, and statistics influenced modern theorists to turn away from theories of image-based cognition and to adopt theories of cognition that were inherently non-perceptual. These theories proposed that knowledge is represented in cognitive systems as abstract symbols that reside separately from perceptual systems (Barsalou, 1999).

Over the past decade, cognitive theorists have returned to favor explanations that suggest cognition is rooted in perception and action. This is called the embodiment hypothesis. Theories of conceptual representation differ in the degree to which representations are embodied. They fall along a continuum from “unembodied” to “strongly embodied” (see Meteyard, Cuadrado, Bahrami, & Vigliocco, 2012 for review). This chapter will review embodied theories, characterizing the degree to which sensory and motor representations are necessary in conceptual representation, predictions made by such theories, and evidence for and against them.

1.1 THEORIES OF EMBODIED COGNITION

1.1.1 Unembodied/secondary embodiment theories

Unembodied theories suggest that sensory and motor information is irrelevant for conceptual representation. That is, conceptual representations are entirely amodal. These unembodied theories propose that knowledge is represented in cognitive systems as abstract symbols that reside independently of perceptual systems (Barsalou, 1999; Meteyard et al., 2012). Additionally, these theories suggest that conceptual representations are formed by transforming, or transducing, the perceptual state elicited

by the experience of the concept's referent into an entirely new non-perceptual language, and the resulting abstract symbols are subsequently stored in long-term memory with arbitrary links to the precipitating perceptual state (Barsalou, 1999; Barsalou et al., 1993). Furthermore, there is no interaction between semantic information and sensory-motor systems. During semantic tasks, any activation of sensory-motor information occurs through an indirect route, such as when working memory processes engage sensory and motor processing (Meteyard et al., 2012).

According to unembodied theories, semantic processing is thought to occur in a conceptual "hub", which serves as a center for amodal conceptual representation (Kiefer & Pulvermüller, 2012). These theories predict that semantic processing should remain intact when sensory or motor systems are damaged or impaired. Only damage to the conceptual hub would result in deficits of semantic processing.

Secondary embodiment theories also propose that conceptual representations are amodal. They differ from unembodied theories, because they allow for non-arbitrary mappings between semantic representations and sensory and motor information. Sensory and motor information contribute to conceptual representations but are not essential. Mahon and Caramazza (2008) describe the role of sensory and motor information as "coloring", meaning that this information can enhance concepts but not change the "essence" of a concept (Pulvermüller, in press). It is the amodal system that gives concepts their meanings rather than the sensory and motor systems. Secondary embodiment theories would predict poorer conceptual representation with damage to sensory and motor systems, but semantic processing would remain largely intact. In imaging studies, secondary embodiment theories would predict activation across various

semantic tasks in regions outside of sensory and motor areas that do not correspond to task-related control processes (Meteyard et al., 2012).

Neuropsychological research in patients and healthy participants provides support for embodied theories of cognition. Patients who exhibit Semantic Dementia, characterized by a loss of conceptual knowledge across all conceptual domains, suffer from a neurodegenerative disease which attacks the temporal poles and surrounding areas. This condition provides evidence for amodal conceptual representation, because patients show deficits for concepts across semantic categories and feature types while sensory and motor systems remain intact. Additionally, stimulation of the temporal poles using transcranial magnetic stimulation (TMS) results in poor performance on various semantic tasks in healthy participants (Pobric, Lambon-Ralph, & Jeffries, 2009). As a result, the anterior temporal cortex has been proposed to be a hub for conceptual representation.

Pulvermüller (2013) argues that unembodied theories cannot completely explain conceptual representation, because grounding is paramount for semantics. As demonstrated by the classic thought experiment “Chinese Room” by Searle (1980), given rules for manipulating and combining the symbols of an unknown language, an individual can produce appropriate responses without understanding their meanings. The individual will only understand the meaning of the language once the symbols become grounded in perceptual and motor experiences. This implies that conceptual representation must involve interaction between amodal systems and sensory and motor systems. Given that amodal systems and sensory motor systems interact and exchange information, it is not prudent to argue that sensory and motor information is non-essential

for conceptual representation. Pulvermüller (in press, p. 3) argues this by providing the following analogy:

“It would obviously be wrong to state that the thrust pushing an airplane *occurs in* one of its three engines because two of them can optionally be switched off. There is reason to say that, if all three are at work, the airplane’s thrust in fact *occurs in* all three of them – even though one alone may do the job.”

Although amodal systems may contribute to the representation of concepts, sensory and motor systems must interact and provide information. The interactive nature of the two systems precludes amodal systems from providing the “essence” or meaning of a concept.

Based on the literature, it seems that unembodied and secondary embodiment theories of cognition cannot fully account for how the brain represents concepts. Theories accounting for a greater role of sensory and motor systems are necessary to explain how concepts are represented.

1.1.2 Weak/strong embodiment theories

Weak and strong embodiment theories propose that conceptual representations are modal and that sensory and motor information is essential, not secondary, to conceptual representation. Concepts are represented in distributed neural networks that overlap with the perceptual systems used to gain knowledge about a concept’s referent (Barsalou, 1999, 2003; Markman & Dietrich, 2000).

Weak and strong embodiment theories differ in the degree to which conceptual representations are dependent on sensory and motor systems as well as the nature of the interaction between the two. Weak embodiment theories suggest that secondary sensory

and motor regions are necessary for conceptual representation and that semantic information mediates early sensory and motor processing (Meteyard et al., 2012). Strong embodiment theories propose that conceptual representations are entirely dependent on primary sensory and motor regions and that semantic information directly modulates sensory and motor processing in order to fully simulate a concept (Meteyard et al., 2012). One of the most comprehensive embodied theories is perceptual symbol systems proposed by Barsalou (1999). This theory can be seen as weakly or strongly embodied based on whether one interprets a full simulation of a concept as necessary for conceptual representation (Meteyard et al., 2012). The theory of perceptual symbol systems proposes that concepts are represented as symbols that are records of the neural activation that occurs when perceiving the referent of the concept. These symbols can be consciously or unconsciously processed, where conscious processing produces mental imagery of the concept's referent. While the perceptual symbol is a record of the neural activation occurring at the time of perception, it is not a complete record of the entire cognitive state. The perceptual symbol captures a "schematic," or general representation, of the original cognitive state (Barsalou, 1999). Perceptual symbols are multimodal, in that they capture the perceptual experience of the referent of a concept through all sense modalities. When perceiving an apple, for example, perceptual symbols for the visual appearance, smell, taste, hand and mouth feel, and crunching sound during eating are formed and stored in their corresponding modality-specific brain regions. In addition to the five senses, symbols capture information about the proprioceptive and introspective experience. In the example of the apple, perceptual symbols of the emotional experience

of eating an apple and the motor movements associated with grasping and eating the apple are also formed and stored in their respective brain regions.

Perceptual symbol systems propose that symbols are a record of the neural activity that occurs when perceiving the referent of a concept, but how might these records be formed? The sensorimotor theory of conceptual processing suggests that sensory and motor features become attached to a symbol by correlation (Humphreys & Forde, 2001; Warrington, 1984). The neural representation of a concept with perceptual features becomes mapped onto neural activation in the perceptual regions originally active when experiencing the referent of a concept. The sensorimotor theory has also been used to explain how concepts and their meanings become linked to the word stimuli used to describe them. Pulvermüller (2001) proposes that language is represented by functional webs within the cortex that link word form with word meaning. These functional webs are formed and strengthened by a Hebbian learning process in which neurons firing in response to the perceptual and motor features of the word's referent become linked to neurons firing in response to word form. Thus, the functional web representing a single word includes both representations of its linguistic form as well as its perceptual and motor features. More recent evidence has elucidated the mechanism by which activation in the perceptual and motor regions becomes linked to neural activation in response to word form. Semantic-conceptual binding sites within the brain serve to bind perceptual, motor, and language-related information into one conceptual representation (Pulvermüller, 2005). Mirror neurons in the inferior frontal gyrus have been implicated in binding motor information and similar perceptual binding sites are hypothesized (see Aziz-Zadeh & Damasio, 2008 for review).

Once concepts are encoded, further conceptual processing requires that perceptual symbols be retrieved from memory. Unembodied theories propose that the perceptual state elicited by the experience of the concept's referent is transduced into an entirely new non-perceptual language. Subsequent conceptual processing involves retrieving a stored description of the concept in this non-perceptual language for use in cognitive processing, much like the way computer systems operate. In contrast, embodied theories, such as perceptual symbol systems, propose that the original cognitive state experienced during encoding of a concept becomes partially re-enacted when the concept is retrieved (Barsalou, 2003). The re-enactment of neural activity occurs in the sensory association areas of the modalities in which the referent of the concept was experienced. When the re-enactment is conscious, mental imagery occurs; however, conceptual processing is often unconscious and involves no mental imagery. A similar account of how concepts are processed has been proposed in the domain of language. The Language and Situated Simulation (LASS) theory proposes that word stimuli first activate linguistic areas needed to process word form and secondarily activate a "situated simulation" to represent word meaning. This simulation occurs in the perceptual, action-related and emotional neural systems activated when interacting with the referent of the word (Simmons, Hamann, Harenski, Hu , & Barsalou, 2008).

1.1.2.1 Evidence for weak/strong embodiment theories in language

Studies investigating how the brain processes words and sentences with perceptual and motor features have been implemented using perceptual features from all five sense modalities as well as actions involving multiple parts of the body to provide support for and against weak and strong embodiment theories. This section will present

the major findings of studies utilizing concepts containing information about the five sense modalities and motor activity.

1.1.2.2 Vision

The most often studied sense modality is vision, reflecting the overall importance and rich understanding of this sense modality. Of the studies using the visual modality, the perceptual feature of color has been well-studied, mainly due to its unimodal nature. Color is one of the few visual features perceived by vision alone. Pulvermüller and Hauk (2006) investigated how the brain processes words that describe the color and shape of objects during a passive reading task. This study demonstrated that color words elicit activation in the parahippocampus, and shape-related words elicit activation in the medial temporal gyrus, the fusiform gyrus, the inferior and middle frontal cortex, and the prefrontal cortex. The authors attributed activation of the parahippocampus to feature conjunction of color and activation of the fusiform gyrus to feature conjunction of form. In a similar study, Martin, Haxby, Lalonde, Wiggs, and Ungerleider (1995) demonstrated that generating color words produces activation in the ventral temporal lobe, which is anterior to a region involved in color perception. Tan et al. (2008) demonstrated that naming hard-to-name and easy-to-name color patches affects differently neural patterns in the visual cortex and bilateral frontal gyrus, which are activated during color perception. While all three studies implicate different regions in processing color-related words, they all agree that conceptual representation of color relies on perceptual areas. The inconsistency of brain regions may be attributed to the variability in task demands. Gerlach (2007) conducted a meta-analysis of fMRI studies comparing visual processing of natural objects and artifacts. Due to large variability in task demands, a lack of

consistent activation within categories suggests that activation is widely distributed and not organized by category. The author proposes that natural objects and artifacts are organized according to their sensory and functional features rather than category.

Conceptual representation studies have also been implemented with sentences that elicit visual imagery. In a study comparing sentences with high and low visual imagery, Just, Newman, Keller, McEleney, and Carpenter (2004) suggests that comprehension of sentences with high visual imagery produces greater activation in the intraparietal sulcus than sentences with low visual imagery. This region has been implicated previously in spatial processing. Based on subsequent studies, Just (2008) concludes that perceptual representations are not always necessary for sentence processing but become activated when perceptual information is useful for the task at hand. In contrast to studies of color concepts, these studies propose that task demands mediate whether perceptual representations become activated and that conceptual representation does not require perceptual systems. Seemingly, these studies provide evidence that weak/strong embodiment theories of cognition do not fully explain how concepts are represented in the brain, as amodal systems may be fully able to represent concepts.

Taken together, studies using concepts with visual features provide mounting evidence for weak embodiment theories. Overall, these studies have found that processing concepts with visual features involves brain regions anterior to primary visual areas, which is consistent with the predictions of weak embodiment theories.

1.1.2.3 Haptics

Similar studies have investigated how the brain represents concepts that contain haptic information. Due to the overlap between the visual and haptic systems, these

studies compare concepts with visual features to concepts with haptic features. One way to illustrate that conceptual representations rely on perceptual systems is to demonstrate a known perceptual phenomena in conceptual processing. Connell and Lynott (2010) replicated the perceptual phenomenon known as the “tactile disadvantage” for identifying the haptic properties of words in comparison to other perceptual properties. This study suggests that words with haptic properties are processed in a similar manner to objects with haptic properties. In contrast, Newman, Klatzky, Lederman, and Just (2005) found mixed results concerning similarity judgments of visual words describing shape and haptic words describing texture. Shape-similarity judgments activated the IPS, implicated in spatial processing, while texture-similarity judgments activated the inferior extrastriate. The inferior extrastriate has been implicated in semantic processing, which suggests that semantic representation of haptic words does not rely on perceptual systems. It should be noted that Newman et al. (2005) classifies shape as a visual feature only, when shape is perceived by the visual and haptic systems. This oversight may explain why regions involved in haptic perception of shape were not found when making texture-similarity judgments of haptic words. Finally, Goldberg, Perfetti, and Schneider (2006) demonstrates that retrieval of perceptual knowledge relies on the sensory brain regions necessary for obtaining that knowledge. Haptic knowledge retrieval activated somatosensory, motor and premotor areas, while visual knowledge retrieval activated the left ventral temporal lobe and superior parietal lobe.

Neuroimaging studies using concepts with haptic features provide evidence for both weak and strong embodiment theories. One study found that processing concepts with haptic features involves somatosensory association areas, which is consistent with

the predictions of weak embodiment theories. Another study implicated primary somatosensory and motor areas in processing concepts with haptic features, which is consistent with a full conceptual simulation predicted by strong embodiment theories.

1.1.2.4 Other senses

Fewer studies have investigated how the brain represents concepts containing perceptual information about the smell, taste and sound of objects. In the case of olfaction, Gonzalez et al. (2006) showed that reading words with strong associations to odor, such as “cinnamon” or “garlic,” elicits activation in the primary olfactory cortex, including the piriform cortex and amygdala. In a study designed to investigate the neural representation of concepts with acoustic features, Kiefer, Sim, Herrnberger, Grothe, and Hoenig (2008) demonstrated that words with acoustic conceptual features elicited activity in parts of the auditory association cortex, including the left posterior superior temporal gyrus and middle temporal gyrus. These same regions were activated when listening to corresponding real sounds. Similarly, verification of sound knowledge elicits activation in the left superior temporal sulcus (Goldberg et al., 2006). In the case of the gustatory modality, verification of taste knowledge elicits activation in the left orbitofrontal cortex (Goldberg et al., 2006), which is involved in representing taste and smell and becomes active when viewing pictures of food (Simmons, Martin, & Barsalou, 2005).

Neuroimaging studies using concepts with olfactory and gustatory features provide evidence for both weak and strong embodiment theories. One study found that processing concepts with gustatory features involves gustatory association areas, which is consistent with the predictions of weak embodiment theories. Another study implicated

primary olfactory areas in processing concepts with olfactory features, which is consistent with a full conceptual simulation predicted by strong embodiment theories.

1.1.2.5 Motor

Embodied theories of cognition propose that conceptual representations not only rely on perceptual systems but also rely on motor systems. Numerous studies have investigated how action concepts that involve bodily movement are represented in the brain. Desai, Binder, Conant, and Seidenberg (2009) demonstrated that comprehension of sentences describing an action involving hand and arm movements activates the inferior postcentral cortex, which is involved in hand movement control and planning. Similarly, several studies have shown that reading or listening to words and phrases about actions involving the body activate the corresponding region of the premotor cortex (Aziz-Zadeh & Damasio, 2008; Hauk, Davis, Kherif, & Pulvermüller, 2008; Tettamanti et al., 2005). Boronat et al. (2005) demonstrated that judging whether two objects are manipulated in the same way activates the left inferior parietal lobe when viewing object names or pictures. Hoenig, Sim, Bochev, Herrnberger, and Kiefer (2008) investigated conceptual flexibility of visual- and action-related attributes of artifactual and natural word categories to determine whether the conceptual attributes of words depends upon context or situation. They found that when probed with a non-dominant perceptual attribute, such as pairing a visual feature with an action-related word, activation in the modality-specific region was increased. Additionally, activation in the dominant modality always occurred even when probing with a non-dominant attribute. This suggests that conceptual representations are activated differently based on context.

Collectively, these neuroimaging studies utilizing concepts with motor features provide evidence for weak embodiment theories of cognition. These studies implicated motor association areas in processing concepts with motor features, which is consistent with weak embodiment theories. Studies investigating concepts with motor features also provide evidence for strong embodiment theories of cognition. Pulvermüller, Hauk, Nikulin, and Ilmoniemi (2005) used TMS to stimulate the hand and foot regions of the motor cortex while participants performed a recognition task with arm- and leg-related action words. Participants performed significantly better on the recognition task when the corresponding region of the motor cortex was stimulated. This study demonstrated that stimulation of the motor cortex directly influences semantic processing of concepts. Similarly, Buccino et al. (2005) found that passively listening to sentences about hand and foot actions results in motor evoked potentials (MEPs) in the hand and foot muscles respectively. In this study, semantic processing modulated activity within the motor cortex and muscles. Collectively, these studies indicate that semantic systems and motor systems are able to modulate one another and support strong embodiment theories.

The previous neuroimaging studies investigating the neural representation of concepts with perceptual and motor features have all provided evidence to support weak/strong embodiment theories of cognition based on the findings that conceptual processing activates regions that underlie perception and action; however, they do not demonstrate that these regions are required. Lesion studies are instrumental for testing hypotheses of embodied cognition, as they allow for making inferences as to whether an anatomical region is required for performing a particular task. If a region is required for a particular task, patients with lesions in that region will show severe deficits in

performing that task. Several studies have tested the embodiment hypothesis with patients displaying lesions in sensory and motor areas to determine whether these regions are required for representing concepts. Patients displaying damage to visual or auditory association areas show greater deficits in processing words that are visual or sound-related respectively (Neininger & Pulvermüller, 2006; Trumpp et al., 2013). Patients with motor deficits due to amyotrophic lateral sclerosis (ALS), a neurodegenerative disorder affecting the motor cortex, show more severe deficits in processing action words than object-related nouns. These studies suggest that sensory and motor regions are required for representing concepts with perceptual and motor features. In contrast, Arevalo et al. (2012) demonstrates that lesions to sensorimotor areas are not sufficient for producing deficits in processing motor-associated words, suggesting that these regions are required only when motor imagery must be used to represent a concept. Chatterjee (2010) speculates that the inconsistencies in findings may be due to individual differences, suggesting that motor simulation is not always necessary for understanding motor-associated words but influences our understanding when we have engaged in the action before. In line with this explanation is the finding that dancers show greater premotor and intraparietal sulcus activity when watching movements of their familiar style of dance versus another unfamiliar style (Calvo-Merion et al., 2005). This suggests our past motor experiences may enhance our understanding of motor-associated words but are not necessary. In summary, based upon current neuroimaging evidence, the lion share of research shows support for weak/strong embodiment theories; however, it is unclear whether sensory and motor regions are absolutely necessary for conceptual representation.

In contrast to studies supporting weak/strong embodiment theories, Grossman et al. (2002) found that abstract nouns and concrete nouns activate overlapping sensory-motor areas, suggesting that concepts are not organized by modality but rather a multimodal semantic organization. The authors propose that members of the animal category of concrete nouns recruit visual areas, not due to reliance of perceptual processing for comprehension, but because it is evolutionarily advantageous to be able to quickly discriminate predators by sight. While this study seemingly provides support for unembodied theories of cognition, weak/strong embodiment theories can explain how abstract concepts might be represented.

1.1.3 Abstract concepts

Abstract concepts, by definition, lack perceptual features and present a challenge for embodied theories of cognition. How can a system that relies on sensory processing represent a concept that is not defined by its perceptual features? When an abstract concept is considered in isolation, it seems embodied theories fail to explain how it may be represented. However, when an abstract concept is considered in context, embodied theories succeed.

Abstract concepts can be grounded in perception and action by viewing them as metaphorical extensions of concrete concepts (Lakoff, 1987). For example, it has been said that *life is a rollercoaster*. The conceptual representation of life is grounded in the experience of being on or passively viewing the nature of a rollercoaster. Similarly, Barsalou (1999) proposes that abstract concepts can be represented by perceptual symbols by framing them against simulated event sequences. This requires placing the abstract concept in a context that can be experienced perceptually. Selective attention

highlights the part of the simulation that gives the abstract concept its meaning while a perceptual symbol is formed that captures the focusing of selective attention.

Additionally, abstract concepts are associated more with internal affective states, whereas concrete concepts are associated more with external experience (Kousta, Vigliocco, Vinson, Andrews, & Del Campo, 2011). As noted previously, introspective states are also captured by perceptual symbols. Therefore, introspective symbols may be necessary for representing abstract concepts.

Vigliocco, Meteyard, Andrews, and Kousta (2009) proposes that all concepts, concrete and abstract, are represented by experiential and linguistic information. Experiential refers to sensory, motor and affective information, while linguistic refers to a concept's typical association with other concepts. This theory suggests that concrete and abstract conceptual representations differ in the amount that each type of information contributes. Concrete conceptual representations would tend to be dominated by sensory and motor experiential information, while abstract conceptual representations would be dominated by linguistic information with a relatively large contribution of affective experiential information.

Currently, few studies have investigated the representation of abstract concepts from the embodied cognition perspective. Pulvermüller and Hauk (2006) shows that moderately abstract words associated with color and form activate regions anterior to the pre-motor and visual cortices, suggesting abstract concepts are possibly grounded in action and perception. In line with Vigliocco et al. (2009), other neuroimaging studies have found activation in sensory and motor areas as well as regions associated with affective processing for abstract concepts (Pexman, Hargreaves, Edwards, Henry, &

Goodyear, 2007; Wilson-Mendenhall, Barrett, Simmons, & Barsalou, 2011). In contrast, a meta-analysis of 19 fMRI and positron emission tomography (PET) studies, indicates that abstract concepts elicit activity in regions associated with verbal processing (inferior frontal gyrus and middle temporal gyrus) while concrete concepts elicit activity in perceptual areas (Wang, Conder, Blitzer, & Shinkareva, 2010).

In summary, the embodiment hypothesis proposes that cognition is grounded in action and perception. Theories explaining how concepts are represented in the brain can be characterized by the extent to which sensory and motor representations are necessary for conceptual representation as well as how much interaction occurs between amodal and sensory and motor systems. Patient studies provide support for unembodied/secondary embodiment theories of cognition, which posit that sensory and motor representations are unnecessary for conceptual representation. In contrast, the bulk of neuroimaging studies of language support weak or strong embodiment theories, which suggest that conceptual representation is entirely dependent upon sensory and motor representations.

CHAPTER 2

VISUAL AND HAPTIC OBJECT PERCEPTION

The current work investigated how the brain represents concept with visual and haptic features. Embodied theories of cognition predict that conceptual representation is grounded in the sensory systems involved in perceiving the referent of a concept. Therefore, it is imperative to understand how objects are perceived through the visual and haptic senses.

Vision is perhaps the most important sense for object perception. Accordingly, studies of visual perception of objects greatly outnumber studies investigating object perception using other senses, and visual perception is relatively well-understood. Due to the heavy overlap in information acquired during visual and haptic perception of objects, vision and haptics have naturally been the focus of studies investigating multimodal representations of objects.

Visual perception provides rich information about object properties. Some information is exclusive to the visual modality, such as color, brightness and spatial pattern, but some object properties are shared across multiple senses. The geometric properties of objects, such as shape, size, and curvature can be perceived with both vision and haptics. For example, the curvature of a basketball can be seen with the eyes as well as felt with the hand. Therefore, geometric information is represented redundantly by these senses. Haptic perception can provide unique information regarding the material properties of objects that are unavailable to vision. Material properties include weight,

temperature, elasticity and texture. While visual cues may suggest which material properties an object has, haptic perception is often necessary to characterize an object's material properties.

To understand how the brain processes the material and geometric properties of objects, one must first consider how the visual and haptic systems are organized. The visual system can be divided into two separate pathways, the ventral and dorsal streams. The ventral stream originates in area V1 of the primary visual cortex and projects to the inferotemporal cortex, while the dorsal stream originates in area V1 and projects to the posterior parietal cortex (PPC). The visual system is hierarchical, in that information grows in complexity as it flows from V1 to its final destination in the parietal and temporal cortices. Ungerleider and Mishkin (1982) proposes a model in which the ventral and dorsal streams process different aspects of visual perception. The ventral stream processes information regarding the identity of objects, while the dorsal stream processes information regarding the spatial location of objects. As an alternative to this model, Goodale and Milner (1992) proposes a model in which the two pathways process the same perceptual information for different purposes. The ventral stream forms a perceptual representation of the object that captures its perceptual properties and relationship to its environment for the purpose of identification and extracting meaning, while the dorsal stream captures information regarding the location of the object in relationship to the body for the purpose of acting upon the object. The dual pathway model of the visual system has become widely accepted since the late 20th Century and has influenced the way in which other sensory systems are studied. As a result, similar models have been developed for the haptic system.

Like the visual system, the haptic system is organized hierarchically with ventral and dorsal streams of information flow. Three unique models have been proposed to explain how information is processed within the haptic system (see James & Kim, 2010 for review). For the purpose of this dissertation, only the third model will be presented, because it uniquely considers the convergence of the two visual and two haptic streams of information (James, Kim, & Fisher, 2007). In this model the ventral stream originates in the primary somatosensory cortex and projects to the lateral occipitotemporal cortex (LOC). The ventral stream is responsible for processing information regarding the material properties of objects, such as texture. In addition to these areas, material properties of objects, specifically texture and hardness, have been demonstrated to activate the parietal operculum, which contains the secondary somatosensory cortex (SII; Servos, Lederman, Wilson, & Gati, 2001). The dorsal stream also originates in the primary somatosensory cortex and projects to the intraparietal sulcus (IPS) of the PPC. The dorsal stream is responsible for processing information regarding the geometric properties of objects, such as shape. Figure 2.1 depicts the regions where the visual and haptic systems converge.

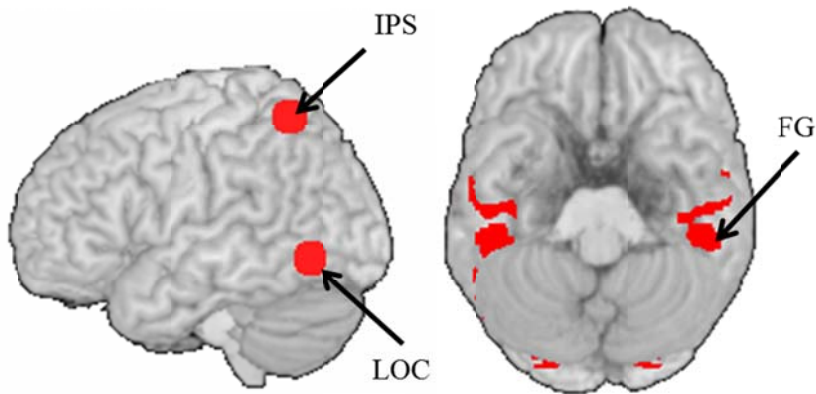


Figure 2.1 The visual and haptic systems converge at the IPS, LOC and FG.

Due to the heavy overlap in information processed in the visual and haptic systems, it stands to reason that information is shared between the two. Evidence suggests that the ventral and dorsal streams of the visual and haptic systems converge. The convergence of the corresponding streams of the visual and haptic systems occurs at the LOC and IPS, which are thought to be bimodal visuo-haptic processing centers (James et al., 2007). The LOC was once thought to be solely a visual processing area, as a lesion study of patient DF suggested the LOC is necessary for visual object recognition (James, James, Humphrey, & Goodale, 2005). However, recent evidence suggests the LOC is more than a visual processing area (Deshpande, Hu, Lacey, Stilla, & Sathian, 2010; James et al., 2005; Lacey, Flueckiger, Stilla, Lava, & Sathian, 2010; Lacey, Tal, Amedi, & Sathian, 2009). James et al. (2005) demonstrated that processing in the LOC can be driven by either visual or haptic exploration of an object's shape. It is possible; however, that LOC activation elicited by haptic processing of shape information occurs merely as a result of visual imagery of an object's shape. By manipulating the familiarity of objects, Lacey et al. (2010) found that the LOC is activated by visual imagery of shape only when the object is familiar. When an object is unfamiliar, LOC activation is driven by haptic input from exploration of the object's shape. Effective connectivity studies suggest the LOC is accessible by both top-down and bottom-up connections depending on the familiarity of the perceived object (Lacey et al., 2009; Deshpande et al., 2010). Bottom-up connections project from the somatosensory cortex and become activated during perception of unfamiliar objects. Top-down connections project from frontal areas and become activated during perception of familiar objects (Lacey et al., 2009; Deshpande et al., 2010). Familiar objects elicit activation in the LOC that is less

somatosensory driven, because global shape can be derived without spatial imagery. Therefore, it seems visual and haptic input activates the LOC directly, and activation is modulated by the familiarity of the object.

The existence of bimodal visuo-haptic areas raises the question of how perceptual information about objects is represented. When an object is perceived, is one integrated multimodal representation formed, or are multiple unimodal representations formed? The answer to this question can be discovered by examining the manner in which perceptual information is processed in these bimodal visuo-haptic areas during object perception. An early study suggests that visual and haptic representations of objects are modality-specific with cross-modal transfer of information, possibly through the insula claustrum (Hadjikhani & Roland, 1998). That is, visual and haptic information may be processed independently and become bound into a single percept through perceptual binding within this region (Crick & Koch, 2005). More recent evidence suggests otherwise, as the regions within the insula claustrum appear to be unimodal (Remedios, Logothetis, & Kayser, 2010). (Whitaker, Simões-Franklin, & Newell, 2008) suggests that information from visual and haptic perception of texture is processed in parallel and remains mostly independent. These studies suggest that multiple unimodal representations are formed during object perception, and visual and haptic information is merely processed within the same bimodal region but is not integrated.

While multiple unimodal representations cannot be ruled out, more evidence supports a single integrated multimodal representation for objects with visual and haptic properties (Helbig et al., 2012; James et al., 2005; Kim & James, 2010; Pietrini et al., 2004). By manipulating stimulus salience, Kim & James (2010) found evidence that

visual and haptic information is integrated in the LOC and IPS based on “enhanced effectiveness,” in which multisensory activation becomes enhanced with increasing effectiveness of unisensory stimuli. Similarly, Helbig et al. (2012) suggests that visual and haptic shape information is integrated as early as the primary somatosensory cortex. Taken together these studies indicate that an integrated visuo-haptic representation of objects is formed early on during object perception; however, the possibility of additional unimodal representations cannot be ruled out.

In summary, the visual and haptic systems are overlapping perceptual systems that contain dual pathways for processing different aspects of perceptual stimuli. The LOC, once thought to be a visual region, is bimodal, which both visual and haptic stimuli activate directly. Evidence suggests that these perceptual systems represent stimuli multi-modally rather than using multiple unimodal representations.

CHAPTER 3

APPROACHES TO THE STUDY OF CONCEPTUAL REPRESENTATION

3.1 UNIVARIATE VS. PATTERN-BASED APPROACHES

Traditional approaches to the analysis of fMRI data use univariate statistical methods to determine which brain regions are involved in the performance of a specific cognitive task. These methods seek to detect average activation differences in brain regions between experimental conditions. That is, the analysis asks which brain regions are on average activated to a greater extent during condition A in comparison to condition B. A significant difference in average regional brain activation in one condition over another suggests a brain region's involvement in a specific cognitive process.

Fundamentally, traditional approaches are advantageous, because they statistically link brain activity to the experimental conditions of interest; however, a major assumption of traditional approaches produces several disadvantages (O'Toole et al., 2007). Traditional approaches assume voxels are independent, when intercellular communication prevents this possibility. As a result, traditional approaches do not have the capacity to investigate the information present in the interaction between voxels. Furthermore, the assumption of independence necessitates measures to control for multiple comparisons. Since traditional approaches compare activity measured at every voxel between experimental conditions, the alpha level for statistical tests becomes inflated. The corrections made to counter the inflation of alpha lead to overly

conservative statistical tests, resulting in the possibility of experimenters falsely assuming the null hypothesis (O'Toole et al., 2007). Methodologically, traditional approaches utilize spatial smoothing of voxels within a region of interest (ROI) to reduce noise and increase sensitivity to activation in response to an experimental condition. However, spatial smoothing also reduces the sensitivity to detect fine-grained patterns of activation, which may discriminate between experimental conditions (Mur, Bandettini, & Kriegeskorte, 2009; Norman, Polyn, Detre, & Haxby, 2006). Another result of spatial smoothing is that traditional approaches can only detect situations when all voxels in an ROI display a signal change in the same direction. When voxels within an ROI exhibit signal changes in opposite directions, which may or may not change the spatial-mean activation, traditional approaches will not pick up the change.

In contrast, pattern-based approaches, such as multivariate pattern analysis (MVPA), use multivariate statistical methods to analyze the information content of fine-grained patterns of brain activity found in functional brain regions (Mur et al., 2009). These methods seek to detect differences in patterns of brain activity to infer how information is represented in the brain. Unlike traditional approaches, pattern-based approaches do not use spatial smoothing to increase sensitivity to activation in response to an experimental condition. Instead these approaches exploit the variation in brain activation across ROIs to investigate how patterns of brain activity discriminate between experimental conditions.

Fundamentally, pattern-based approaches possess the same advantage as traditional approaches but also overcome the disadvantages resulting from an assumption of independent voxels. Like traditional approaches, pattern-based approaches provide a

link between brain activity and experimental conditions presented during scanning (O'Toole et al., 2007). The ultimate goal is to use patterns of brain activity to predict the experimental condition being experienced by the participant. Rather than assuming voxels are independent, pattern-based approaches examine voxels jointly and detect patterns of brain activity resulting from interactions among voxels. While traditional approaches focus on answering the question of where information processing occurs in the brain, pattern-based approaches focus on explaining how the brain represents information while also revealing where information resides (O'Toole et al., 2007; Norman et al., 2006). Methodologically, pattern-based approaches have the advantage of detecting any activity pattern change within an ROI, even when the spatial-mean activity does not change (Mur et al., 2009). Finally, pattern-based approaches exhibit increased temporal resolution, as the experimental condition being experienced by the participant can be predicted from mere seconds of brain activity (Norman et al., 2006).

The following sections detail the steps involved in MVPA for extracting the fMRI signal and analyzing the observed patterns of brain activity. Typically, the procedure for MVPA entails preprocessing the data, dividing the data into training and test sets, selecting the features to be used to train the classifier, choosing an appropriate classifier, and cross-validating the results. Researchers must make choices at every step that impact the final result of pattern classification. These choices must be made in light of the experimental design and research question.

3.1.1 Preprocessing

The first step in pattern analysis is data preprocessing. For pattern-based approaches the data is preprocessed in a similar way as traditional approaches, including

slice timing correction, motion correction, and removal of trends. As mentioned previously, spatial smoothing is not employed for pattern-based approaches, as this removes the fine-grained patterns that carry informational content. Subsequently, the data must be transformed into examples, which entails extracting the relevant signal values to input into the classifier. Generating examples of experimental conditions can be done in many ways and largely depends on experimental design. One common way to create an example is to average multiple volumes of data from a single trial to approximate the peak of the hemodynamic response function (HRF; Pereira, Mitchell, & Botvinick, 2009). This creates a vector of average signal readings at each voxel, which is tied to the experimental condition presented in that trial. Alternative methods include using single volume measures as individual examples or averaging multiple trials of the same experimental condition (Mur et al., 2009; Pereira et al., 2009). Examples can also be created from estimates of predicted voxel activity derived using the General Linear Model (GLM; Mur et al., 2009; Pereira et al., 2009). In this case, the pattern of beta-values across voxels is used as an example for that condition. Regardless of the method chosen for creating examples, it is better to create more examples than fewer, as parameter estimates generated by the classifier become better with a larger input of examples. Additionally, patterns should not be averaged across participants to avoid averaging out the fine-grained informational content. All analysis should be performed in native subject space.

3.1.2 Data division

To ensure unbiased results, data should be divided into two sets, the training set and the test set. The training set refers to the examples used as input for the classifier,

from which the classifier learns a mapping from the experimental condition to the activity pattern. The test set refers to the examples whose class label is predicted from the mapping derived from the training set. It is important to choose the training and test sets carefully, so that the data are independent. This can be achieved by selecting examples that are created from blocks or trials that are not overlapping (Pereira et al., 2009; Mur et al., 2009). A violation of independence can cause an increase in accuracy estimates, as the example in the training and test sets are very similar.

3.1.3 Feature selection

Once the data has been preprocessed and split into independent training and test sets, the next step is feature selection. The number of features sampled in a typical fMRI study can reach into hundreds of thousands voxels. When using voxels as features, the number of features greatly surpasses the number of examples. It is advantageous to reduce the number of features used for classification due to issues of over-fitting (O'Toole et al., 2007). When there are too many free parameters relative to examples, the training data can be over-fit. This situation results in a solution that generalizes to any test set drawn from the same population. The solution to this problem is to select a subset of features to be used for classification. Feature selection should be performed on the training data only to maintain an assumption of independence between the training and test sets. Using the entire dataset for feature selection allows the test set to influence how well the classifier learns from the training set (Pereira et al., 2009).

A theory-driven approach to feature selection is to choose voxels located in a ROI to use for classification. For example, the primary somatosensory cortex could be used as an ROI for classifying whether an object is perceived with the visual or haptic modality,

as this region is well-known for processing information related to the sense of touch. However, ROIs chosen for feature selection must not necessarily be spatially contiguous (Mur et al., 2009). A localizer scan could be used to determine which areas of the brain are more responsive to a certain aspect of a task. A localizer task could be performed by comparing the presentation of a haptic stimulus to fixation, and those voxels displaying more activity for the haptic condition would be selected regardless of whether they reside in the primary somatosensory cortex or elsewhere in the brain.

Searchlight analysis is a classification method that uses a unique approach to feature selection. Rather than using functionally-defined ROIs, this analysis employs a spherical multivariate “searchlight” with a predefined search radius to scan an entire volume. The signals from all voxels falling within the searchlight region are combined using a multivariate statistic, such as the Mahalanobis distance, which compares the activity patterns between conditions for selected voxels (Kriegeskorte, Goebel, & Bandettini, 2006). The voxels within the searchlight are examined jointly with MVPA to determine whether information about the variables of interest is carried within the searchlight region (Chen et al., 2010). Computational expense depends on the size of the searchlight used, as the number of classifiers trained is equal to the number of searchlight regions. While possibly computationally expensive overall (when a small searchlight region is used), the searchlight analysis restricts the features examined during the training of each individual classifier, reducing the risk of over-fitting the data.

As an alternative to ROI-based approaches, feature selection can be done using inferential statistics to evaluate which features are most useful for classification (O’Toole et al., 2007; Pereira et al., 2009; Mur et al., 2009). Mitchell et al. (2004) demonstrated

the usefulness of feature selection methods which choose voxels that discriminate best between an experimental condition and fixation. Voxel discriminability is evaluated by computing a pairwise t-test between each voxel's activity level during the experimental condition and fixation condition. Voxels with the largest t-statistics are chosen for classification. Feature selection based on a measure of voxel stability has also been used successfully (Mitchell et al., 2008; Pereira et al., 2009; Shinkareva et al., 2008). Voxel stability is computed by averaging pairwise correlation coefficients between vectors of presentations of all conditions in the training set. Voxels with the largest t-statistics, reflecting more consistent variation in activity across conditions, are selected for classification. Both methods use inferential statistics to evaluate how each voxel responds across conditions to allow for reducing the overall number of features to those that will perform best for classification.

Dimensionality reduction techniques have also been used to select features for classification (O'Toole et al., 2007; Pereira et al., 2009; Mur et al., 2009). This type of feature selection involves finding a lower dimensional representation of the fMRI data by using multivariate statistical methods such as principal components analysis (PCA) or independent components analysis (ICA). In the case of PCA, the entire dataset is reduced to a set of orthogonal brain response patterns that capture as much of the variance in the data as possible. Components accounting for the most variance in the data are selected for classification, and the vectors of weights associated with the principal components can be used as input instead of the vectors of voxel readings (O'Toole et al., 2007; Pereira et al., 2009). Dimensionality reduction techniques are advantageous, because they reduce the number of features as well as reduce noise in the data (O'Toole et al., 2007).

However, unlike inferential statistical methods of feature selection, most dimensionality reduction techniques do not have the benefit of associating voxel readings with their corresponding experimental condition and may not improve classification results (Pereira et al., 2009).

3.1.4 Classification

The goal of classification algorithms is to discriminate between the patterns of brain activity elicited by each experimental condition. Classification is performed on the multivariate space derived from the fMRI signal readings of selected voxels at specific time points during the scan. Given N voxels are selected for classification, the pattern of brain activity is represented in an N -dimensional space with a single data point for every voxel reading (Tong & Pratte, 2012). The classification algorithm seeks to divide the representational space into classes of stimuli.

Two types of classification algorithms can be used to analyze fMRI data. The first and most simple is the linear classifier. Linear classifiers aim to find the most optimal separation of stimulus classes by dividing the representational space with a hyperplane (O'Toole et al., 2007). The second type of classifier is non-linear, which can achieve a more optimal separation of the representational space by bending the hyperplane in different ways (O'Toole et al., 2007). While non-linear classifiers can capture more complex relationships between stimulus classes and patterns of brain activity, it is suggested to start with the simpler linear classifier (O'Toole et al., 2007; Kriegeskorte, 2011; Tong & Pratte, 2012). Linear classifiers reduce the risk of overfitting the data, which occurs easily due to a greater number of voxels than signal readings (Kriegeskorte, 2011). Additionally, a linear relationship between stimulus

classes and patterns of brain activity is easier to interpret than a non-linear relationship. Finally, non-linear classifiers can capture relationships between stimulus classes and patterns of brain activity that reflect computations of the classifier itself rather than computations performed in the brain (Tong & Pratte, 2012). Kriegeskorte (2011) suggests that the benefits of linear classifiers outweigh the ability of non-linear classifiers to capture more complex relationships. O'Toole et al. (2007) suggests trying a non-linear classifier only after a linear classifier fails to achieve above chance accuracy and when there is a theoretical motivation to assume a more complex relationship, such as testing computational models of brain processing (Kriegeskorte, 2011).

3.1.5 Cross-validation

As mentioned previously, the data must be divided into training and test sets to get an unbiased estimate of how well the classifier learns the relationship between experimental conditions and patterns of brain activity. Additionally, classification algorithms benefit from having lots of examples from which to learn. As a result, the data must be divided in such a way that there are plenty of training examples available but there are enough examples on which to test. Cross-validation is a procedure for evaluating how well a classifier learns the identity of patterns of brain activity while optimizing the use of examples from the data. The most extreme version of cross-validation is the leave-one-out cross-validation (LOOCV) approach. LOOCV entails training the classifier on all examples except one and testing on the left-out example. Then, the procedure is repeated until each example serves as the test example once. The performance of the classifier is estimated by computing the percentage of correct classifications, also known as accuracy.

A disadvantage of LOOCV is its computational expense, as the number of classifiers needed equals the total number of examples in the data (Perreria et al., 2009). K -fold cross-validation is a method that reduces the computational expense by dividing the data into larger chunks or folds, where k is equal to the number of folds. The number of folds is typically dependent on the experimental design, which can provide natural folds in the data. For example, a fold could be equal to blocks in a blocked-design experiment or runs in an event-related experiment. The classifier is trained on all folds except one and tested on the left-out fold. The procedure is repeated until each fold serves as the test fold once. The performance of the classifier is estimated by averaging the percentage of correct predictions obtained at each fold.

3.1.6 Evaluating results

The ultimate goal of classification is to demonstrate that a classifier can predict which experimental condition elicited a pattern of brain activity better than a classifier that simply “guesses” at random. The classification accuracy obtained from cross-validation is an unbiased estimate of the true accuracy of the classifier. The true accuracy refers to how well the classifier would predict the identity of a new example drawn randomly from the distribution from which examples in the training set were drawn (Pereira et al. 2009). The classification accuracy estimate is said to be significant if it exceeds the accuracy expected if the classifier is simply guessing at random and the patterns of brain activity carry no information about the variables of interest (the null hypothesis). In the case of an experiment with two conditions, the classifier would have 50% chance of predicting the condition correctly given the null hypothesis is true. The significance of the classification accuracy estimate can be evaluated based on the

binomial distribution $B(n, p)$, where n is the number of trials of each classification computation and p is the probability of correct classification when the examples are randomly labeled (Pereira et al., 2009). An alternative to using the binomial distribution to evaluate the significance of classification accuracy is to utilize a permutation test. A permutation test simulates the results of a classifier that is randomly guessing by randomly assigning the condition labels of examples in the training set prior to training the classifier and testing on the test set (Pereira et al., 2009). This is done many times, each with a different random assignment of condition labels. The p -value computed from this test is the percentage of classification accuracies obtained from the permutation test that equal or exceed the observed classification accuracy (Pereira et al., 2009). A significant result suggests that patterns of brain activity contain information about the variables of interest.

3.1.7 Implications of pattern-based approaches

The primary goal of pattern-based approaches is to determine whether the fMRI signal contains information about the variable of interest (Pereira et al., 2009). That is, can we discriminate classes of the variable of interest based on patterns of brain activity? This question is answered by using classification algorithms to predict which stimuli a participant is experiencing from patterns of brain activity. Assuming a strong experimental design, accurate classification that is significantly above chance suggests that the patterns of observed brain activity contain information about the classes of the variables of interest. In addition to pattern discrimination, it is possible to determine in which areas of the brain this information is represented (Pereira et al. 2009; Tong & Pratte, 2012). Pereira et al. (2009) suggests a two-step approach for determining where

information is represented in the brain. First, one can determine which voxels are contributing to classification accuracy by examining the set of voxels selected by feature selection at each fold of cross-validation. Given that the voxels being selected contain sufficient information to discriminate classes, the overlap of voxels chosen at every fold can be viewed as the necessary set of voxels for accurate classification (Pereira et al., 2009). Examining the location of this necessary set may give insight as to where class information is represented in the brain. Next, one can evaluate which voxels in the subset affect classification the most. When using a linear classifier, this means simply examining the weight assigned to each voxel (Pereira et al., 2009). Voxels with the largest weights contribute more to accurate classification; therefore, these voxels more accurately discriminate class information. It follows that class information may reside in these voxels. An alternate method of examining which voxels contribute to classification performance is to selectively remove voxels to be used by the classifier based on a priori predictions (O'Toole et al., 2007). If the classification accuracy decreases, one can assume that these voxels contained information needed to discriminate between classes. If classification accuracy increases, one can assume these voxels contained mostly noise that impeded classification performance.

Once a subset of voxels has been identified, it is also possible to characterize how class information is represented within the region. The process of describing how information is represented requires characterizing the relationship between the observed patterns of brain activity and the stimuli presented to participants (Pereira et al., 2009; Tong & Pratte, 2012). This relationship is what the classifier learns, but it is up to the researcher to link this relationship back to the experimental design in order to understand

the structure of the class information. Characterization relies on strong experimental design and often multiple related experiments to eliminate confounds (Tong & Pratte, 2012). This can be achieved in various ways, such as correlating classifier performance with behavioral performance, comparing the similarity of classes with the similarity in observed patterns of brain activity, and generalizing classifier performance to new stimuli (Pereira et al. 2009; Tong & Pratte, 2012). In the first method of pattern characterization, classifier performance is compared with some behavioral measure to identify similarities. If a classifier makes similar mistakes in classification as a participant, one can infer that the participant and the classifier are using the same information for classification. For example, Raizada, Tsao, Liu, and Kuhl (2010) demonstrated that the neural representation of the sounds of syllables /ra/ and /la/ were most discriminable when the participant was better able to behaviorally discriminate between those syllables. When the participant made more mistakes in discriminating between those sounds, classification of the neural representation of those sounds was less accurate. Thus, the relationship captured by the classifier suggests information regarding the sound of syllables was present in the patterns of brain activity. The second method of pattern characterization involves relating the similarity of the classes of stimuli with the similarity of patterns of brain activity. For example, Weber, Thompson-Schill, Osherson, Haxby, and Parsons (2009) demonstrated that information about mammals is structured by category in the ventral visual pathway by comparing the computed similarity of brain responses to various categories of mammals with participants' subjective similarity ratings of the same stimuli. Since the brain responses and similarity ratings showed similar structure, it suggested that information carried in the neural patterns of activation

was organized by category. Finally, pattern characterization can be achieved by generalizing a classifier's performance to new stimuli. Mitchell et al. (2008) showed that a classifier trained on a subset of concrete nouns from a large corpus of text could predict the fMRI activation associated with thousands of novel words from the same corpus of text. This demonstrated that the classifier was able to learn a set of semantic features that make up the neural representation of concrete nouns.

Finally, pattern classification can be used to evaluate whether information is represented similarly in the brains of different people. Cross-participant classification refers to a classification method that trains classifiers across multiple participants in a study and predicts the class of variable a novel participant experienced. Given that the classifier can accurately predict which class of variable the novel participant experienced based on the patterns of brain activity of other participants, it follows that information regarding the classes of the variable is represented similarly across participants.

In summary, univariate and pattern-based approaches to the analysis of fMRI data ask different questions. Univariate approaches ask which brain regions are involved in a cognitive task, while pattern-based approaches seek to reveal the representational content of brain regions. Both have the ability to statistically link experimental conditions to neural activity, but pattern-based approaches are much more sensitive and consider the interactions among voxels. MVPA is a pattern-based approach that seeks to predict the experimental condition from observed patterns of brain activity. This powerful approach to analyzing fMRI data is data-driven and very flexible based on the experimental question.

3.2 UNIVARIATE APPROACHES TO THE STUDY OF CONCEPTUAL REPRESENTATION

Previous studies investigating how the brain processes concepts with perceptual and motor features have utilized univariate approaches. Univariate approaches ask the question of which brain regions are involved in a certain cognitive task. Results of the previous studies show which brain regions are involved in processing concepts with perceptual and motor features by examining which voxels show significantly greater activation in one condition over another. In most cases, these studies have implicated regions that underlie perceptual processing and motor movement in the processing of concepts with perceptual and motor features, providing support for embodied theories of cognition.

The strengths of univariate approaches stem from the simplicity of the questions they ask. When a region of the brain displays greater activation levels for one condition over another, it is inferred that the brain region is engaged by and involved in the cognitive state associated with the experimental condition. As a result, univariate models are easily interpretable. A brain region is either activated or not activated by an experimental condition. In previous studies examining concepts containing visual information, the left ventral temporal lobe was activated when concepts provided information about the property of color (Martin et al., 1995; Goldberg et al., 2006). In addition to perceptual studies showing that color perception also activates the left ventral temporal lobe, it can be concluded that the left ventral temporal lobe is involved in both color perception and the representation of concepts containing information about color. Univariate approaches have the ability to statistically link experimental conditions to regional brain activation while providing an easily interpretable model. These strengths

of univariate approaches have provided a large amount of evidence to link conceptual processing of concepts containing perceptual and motor features to brain regions previously implicated in processing perceptual experience and motor movement. This evidence is highly informative for further studies utilizing univariate approaches and, as will be demonstrated, pattern-based approaches.

While univariate approaches possess strong qualities, they are limited due to the assumption of voxel independence. Univariate approaches assume voxels are independent and evaluate each voxel in isolation to determine whether it shows greater activation in one condition over another. This produces a need for overly conservative statistical tests, which greatly diminishes the power to detect activation differences at the voxel level. Additionally, to increase the signal-to-noise ratio within a region of interest, spatial smoothing is utilized. This discards fine-grained patterns of information present within the region of interest. All three characteristics of univariate approaches result in a major loss of information. This necessitates the question of what information is being lost and how this information could provide insight into how concepts are represented. For instance, many more brain regions could be implicated in the representation of concepts containing perceptual information. In the case of a brain region that displays signal changes in opposite directions and does not achieve a change in spatial mean activation, univariate approaches will not be sensitive to the signal change. This brain region will not survive the statistical analysis and will, therefore, not be implicated in the representation of the concept. Meyer, Kaplan, Essex, Damasio, and Damasio (2011) demonstrates an example of this within the perception literature in a pair of studies examining cross-stimulus processing of tactile stimuli. A study utilizing single-cell

recordings, analogous to the univariate approach, failed to detect activity in the primary somatosensory cortex, because variations in the firing rates of individual neurons never reached significance (Lemus, Hernández, Luna, Zainos, & Romo, 2010). In an fMRI study with a similar experimental paradigm, a pattern-based approach was able to detect cross-stimulus processing in the primary somatosensory cortex, as variations in the firing rates of individual neurons were jointly analyzed as a neuronal population (Meyer et al., 2011). The difference in approach resulted in two different conclusions from similar studies. The univariate approach led to a conclusion that primary sensory cortices do not encode cross-modal stimuli, while the pattern-based approach led to a conclusion that primary sensory cortices do encode cross-modal stimuli.

The previous scenario demonstrates how both approaches are necessary in order to provide a clearer picture of how concepts are represented in the brain. However, univariate and pattern-based approaches can be complementary rather than contradictory. For instance, studies taking univariate approaches can be utilized by providing a set of core brain regions for analysis with pattern-based approaches. Activity in brain regions identified by univariate approaches have demonstrated a strong statistical link to experimental conditions and have survived highly conservative statistical tests. Therefore, studies employing univariate approaches suggest a core group of brain structures that may contribute to whole-brain patterns of activity. Additionally, they provide a great starting point for the ROI-based feature selection stage of MVPA. For these reasons, univariate approaches and pattern-based approaches should be considered complementary approaches to the study of conceptual representation.

3.3 PATTERN-BASED APPROACHES TO THE STUDY OF CONCEPTUAL REPRESENTATION

It is clear how univariate approaches measure up for the study of conceptual representation, but how are pattern-based approaches particularly well-suited to the study of concepts that contain perceptual information? Pattern-based approaches, such as MVPA, are beneficial for the study of concepts that contain perceptual information due to the unique questions pattern-based approaches ask as well as the nature of perceptual data. In contrast to univariate approaches, pattern-based approaches ask whether information about stimuli is present in a brain region. Pattern-based approaches answer this question by jointly examining voxels to detect patterns of brain activity resulting from interactions among them. Pattern-based approaches are well-suited for the study of how the brain represents concepts containing perceptual information, because perceptual data is inherently multivariate. It is thought that perceptual representations, as well as cognitive and motor representations, are encoded in groups of neurons through population coding (Kriegeskorte, 2011). For example, Groh (2000) demonstrated that the direction in which a stimulus is perceived to move is determined by the overall pattern of response rather than its peak in area MT. Given that perceptual representations are encoded in the activity of groups of neurons, pattern-based approaches are well-suited for studying such representations to reveal the informational content of the region containing those neurons.

In addition to its multivariate nature, perceptual data is inherently multi-modal. Findings from studies of visual and haptic object perception demonstrate that properties are represented both redundantly and in an integrated fashion within the visual and haptic systems. Additionally, information acquired through visual perception can be found in

patterns of brain activity in the primary somatosensory cortex (Meyer et al., 2011). The multi-modal nature of perceptual data makes pattern-based approaches appropriate for the study of concepts containing perceptual information, because conceptual representations may be spatially overlapping. Univariate approaches utilize spatial smoothing, which tends to blur the distinctions between spatially overlapping patterns (Raizada & Kriegeskorte, 2010). However, pattern-based approaches do not always use spatial smoothing in order to exploit the fine-grained patterns of brain activity. Furthermore, pattern-based approaches have been successfully used to investigate spatially overlapping neural representations. For example, Raizada et al. (2010) was able to discriminate between highly overlapping neural representations of the phonemes /ra/ and /la/ in the auditory cortex. Univariate approaches would not have been successful at making the distinction between the representations of the two phonemes, because the spatially smoothed average activation for each phoneme's representation was equal. Therefore, the activation difference between conditions was zero. In the case of visual and haptic object perception, many studies have suggested that the LOC is the site where visual and haptic information is either integrated or represented jointly (Deshpande et al., 2010; James et al., 2005; Lacey et al., 2010; Lacey et al., 2009). Univariate approaches to the study of concepts containing visual and haptic information may not be able to discriminate between spatially overlapping visual and haptic representations in this region. Perhaps, this is the reason no studies examining conceptual representation have implicated the LOC in the processing of concepts containing visual and haptic information. Pattern-based approaches may be able to demonstrate that visual and haptic information is indeed carried in the patterns of brain activity located in the LOC.

In summary, univariate and pattern-based approaches to the analysis of fMRI data ask different questions. Univariate approaches ask which brain regions are involved in a cognitive task, while pattern-based approaches seek to reveal the representational content of brain regions. Both have the ability to statistically link experimental conditions to neural activity, but pattern-based approaches are much more sensitive and consider the interactions among voxels. MVPA is a pattern-based approach that seeks to predict the experimental condition from observed patterns of brain activity. Because perceptual data is inherently multivariate and spatially-overlapping, pattern-based approaches are well-suited to study the representation of concepts with perceptual features.

3.4 GOALS OF THE CURRENT WORK

The current work investigated the neural representation of concepts with perceptual features, specifically visual and haptic, to understand how the modal aspects of concepts are represented. The purpose was to demonstrate that the representation of concepts with perceptual features is more consistent with weak and strong embodiment theories than unembodied and secondary embodiment theories; however, it is beyond the scope of the current work to provide evidence that rules out amodal conceptual representation. The central hypothesis was that the neural representation of concepts with perceptual features is distributed and includes brain regions in the perceptual systems activated when interacting with the referent of that concept. More specifically, concepts containing visual information should be represented in brain regions active when processing visual stimuli, while concepts containing haptic information should be represented in brain regions active when processing haptic stimuli.

The specific aims were as follows:

1. Determine which brain regions participate in processing concepts with perceptual features. The working hypothesis was that concepts with visual features are processed by regions known to be active when perceiving objects visually, while concepts with haptic features are processed by regions known to be active when perceiving objects haptically (Newman et al., 2005). Additionally, we examined the patterns of functional connectivity of these brain regions. We hypothesized that functional networks for processing concepts with visual and haptic features contain similar brain regions, but these brain regions are connected differently based on the type of stimulus being processed.
2. Determine if patterns of brain activity elicited by processing concepts can be used to predict the perceptual information content of a concept using MVPA within and between participants. Our working hypothesis was that the perceptual information content of a concept can be predicted from distributed patterns of brain activity as well as patterns of brain activity from *a priori*-defined regions of interest. Success with MVPA demonstrates that patterns of brain activity contain information pertaining to the perceptual features of concepts.

CHAPTER 4

BEHAVIORAL EXPERIMENT 1

4.1 PURPOSE

The purpose of this experiment was to demonstrate that the representation of concepts with visual and haptic features involves perceptual processing. One way to illustrate that conceptual representations rely on perceptual systems is to demonstrate a known perceptual phenomena in conceptual processing. Connell and Lynott (2010) replicated the perceptual phenomenon known as the “tactile disadvantage” for identifying the haptic properties of words in comparison to other perceptual properties. When participants were asked to respond to the arrival of a perceptual stimulus, they were slower to detect haptic stimuli than visual stimuli even though they were told which modality to expect. The current experiment intended to show a similar tactile disadvantage for making judgments about concepts with visual and haptic features. Given that conceptual processing relies on perceptual systems, we expected to find slower reaction times for processing concepts with haptic features than for concepts with visual features.

4.2 MATERIALS & METHODS

4.2.1 Participants

Participants were thirty-three (18 female) adults ranging in age from 18 to 29 years ($M = 21.1$). One participant was excluded from the behavioral analysis for low accuracy (less than 75% correct). Participants were native speakers of English with

normal or corrected-to-normal vision. All were recruited from the University of South Carolina Psychology Participant Pool. Informed consent was obtained from each participant prior to the experiment, in accordance with the protocol set forth by the University of South Carolina Institutional Review Board.

4.2.2 Stimuli

A set of 192 visual and haptic concept-property word pairings were selected from a database of 774 multi-modal concept-property items from Dantzig, Cowell, Zeelenberg, and Pecher (2011). Of the 192 visual concept-property pairings, 96 contained visual information, and 96 contained haptic information. Concept-property pairings were rated for how strongly each is experienced with five sensory modalities (sight, sound, touch, smell and taste) through a series of norming studies. The concept properties with the highest modality exclusivity ratings for vision and haptics were chosen to ensure stimuli were as unimodal as possible (threshold of 65% or higher for vision and 35% for haptics). Haptic stimuli are inherently more multi-modal, and the threshold for modality exclusivity reflects this. Words containing visual and haptic information did not differ significantly in length ($p = 0.11$) or familiarity ($p = 0.95$).

4.2.3 Experimental paradigm

Participants performed a perceptual property verification task similar to tasks used in behavioral and neuroimaging studies of conceptual processing (Goldberg et al., 2006; Pecher, Zeelenberg, & Barsalou, 2003). On any given trial, participants were asked to decide which of two properties best described a concept from either the visual or haptic categories. The two properties included perceptual features. For example, given the concept “ZEBRA” and the visual properties “STRIPED” and “RED,” the participant

would choose “STRIPED” as the applicable property, because a zebra can be striped but not red. This task was designed to prompt the participant to form a simulation of both the concept and its properties, which may involve sensory-motor processing (Dantzig et al., 2011). The number of times a property was used as the correct choice was balanced with the number of times it was used as the incorrect choice. Additionally, half of all trials had the correct choice listed on the right, while half had the correct choice listed on the left. The concept and property choices were presented for 3000 ms followed by a 1000 ms fixation cross using E-Prime software (Psychology Software Tools, Sharspburg, PA; Figure 4.1). Reaction times for property verification decisions were recorded from the onset of the presentation of concept and property choices.




Figure 4.1 Experimental paradigm for behavioral experiment 1.

4.3 RESULTS

The mean reaction times and error rates for verifying properties of concepts with visual and haptic features were compared using paired-samples t-tests. The mean reaction time for verifying properties of concepts with visual features ($M = 1440.58$ ms) was significantly shorter than the mean reaction time for verifying properties of concepts with haptic features ($M = 1488.55$ ms, $p < 0.001$, Figure 4.2). The mean number of correct responses for verifying properties of concepts with visual features ($M = 80.36$) was not significantly different than the mean number of correct responses for verifying properties with haptic features ($M = 79.09$, $p = 0.12$).

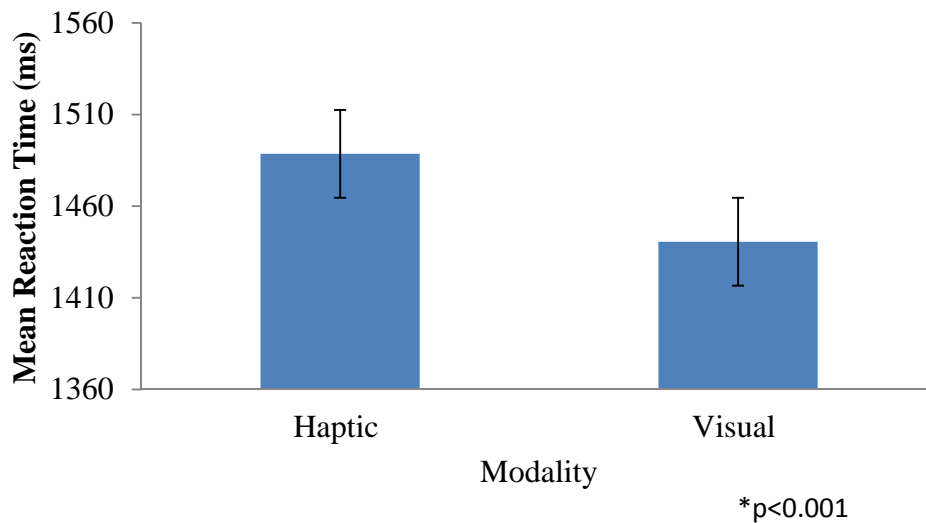


Figure 4.2 Mean reaction times for verifying concepts with visual and haptic features.

4.4 SUMMARY

The purpose of this study was to investigate whether the representation of concepts with visual and haptic features involves perceptual processing by demonstrating a perceptual phenomenon known as the “tactile disadvantage” in behavioral measures of conceptual processing. Given that conceptual processing relies on perceptual systems,

we expected to see a tactile disadvantage when participants verified properties of concepts with visual and haptic features, such that reaction times for verifying properties of concepts with haptic features would be significantly slower than reaction times for verifying properties of concepts with visual features.

A tactile disadvantage was found when participants verified properties of concepts with visual and haptic features. Participants were significantly slower to verify properties of concepts with haptic features than they were when verifying properties of concepts with visual features. No differences were found in the accuracy of responses for verifying properties of concepts with visual and haptic features, suggesting that the difference in reaction times was not due to a difference in task difficulty or a speed-accuracy trade off. The results suggest that conceptual processing indeed relies on perceptual systems, as a phenomenon specific to perception emerged during conceptual processing.

These findings further support modal theories of conceptual knowledge by demonstrating that conceptual representation involves perceptual processing. However, demonstrating that perceptual processing is involved in conceptual representation cannot rule out amodal representations. It is possible that perceptual processing is an emergent process that is unnecessary for the representation of concepts and that amodal representation is present. Further studies will need to be conducted to demonstrate the necessity of modal representations for conceptual processing.

CHAPTER 5

BEHAVIORAL EXPERIMENT 2

5.1 PURPOSE

The purpose of the second behavioral experiment was to validate stimuli chosen for the main fMRI experiment. In order to evaluate how perceptual features of concepts are represented, a baseline condition was needed to control for the perceptual features of concepts. Abstract concepts are defined by their lack of perceptual features, so a baseline condition utilizing abstract conditions was created. To ensure the task was equally as difficult across conditions, a behavioral experiment was conducted to compare the reaction times for making property verifications about concepts with visual, haptic and abstract features.

5.2 MATERIALS & METHODS

5.2.1 Participants

Participants were sixteen (12 female) adults ranging in age from 18 to 36 years ($M = 23.5$). Participants were native speakers of English with normal or corrected-to-normal vision. All were recruited from the University of South Carolina community. Informed consent was obtained from each participant prior to the experiment, in accordance with the protocol set forth by the University of South Carolina Institutional Review Board.

5.2.2 Stimuli

A set of 192 visual and haptic concept-property word pairings were selected from a database of 774 multi-modal concept-property items from Dantzig, Cowell, Zeelenberg, and Pecher (2011). Of the 192 visual concept-property pairings, 96 contained visual information, and 96 contained haptic information. Concept-property pairings were rated for how strongly each is experienced with five sensory modalities (sight, sound, touch, smell and taste) through a series of norming studies. The concept properties with the highest modality exclusivity ratings for vision and haptics were chosen to ensure stimuli were as unimodal as possible (threshold of 65% or higher for vision and 35% for haptics). Haptic stimuli are inherently more multi-modal, and the threshold for modality exclusivity reflects this. Additionally, 182 abstract stimuli were constructed by choosing frequently used abstract nouns and pairing these with commonly used descriptors from a thesaurus. Word stimuli were balanced for average length ($p = 0.351$) and average frequency ($p = 0.061$).

5.2.3 Experimental paradigm

Participants performed a perceptual property verification task similar to tasks used in behavioral and neuroimaging studies of conceptual processing (Goldberg et al., 2006; Pecher, Zeelenberg, & Barsalou, 2003). On any given trial, participants were asked to decide which of two properties best described a concept from either the visual, haptic, or abstract categories. In the visual and haptic conditions, the two properties included perceptual features. For example, given the concept “ZEBRA” and the visual properties “STRIPED” and “RED,” the participant would choose “STRIPED” as the applicable property, because a zebra can be striped but not red. In the abstract condition, the two

properties included non-perceptual features. For example, given the concept “LOSS” and the abstract properties “SAD” and “SECURE,” the participant would choose “SAD” as the applicable property, because loss can make one feel sad but not secure. This task was designed to prompt the participant to form a simulation of both the concept and its properties, which may involve sensory-motor processing (Dantzig et al., 2011). The number of times a property was used as the correct choice was balanced with the number of times it was used as the incorrect choice. Additionally, half of all trials had the correct choice listed on the right, while half had the correct choice listed on the left. The concept and property choices were presented for 3000 ms followed by a 1000 ms fixation cross using E-Prime software (Psychology Software Tools, Sharsburg, PA; Figure 5.1). Reaction times for property verification decisions were recorded from the onset of the presentation of concept and property choices.

5.3 RESULTS

The goal of the analysis of the behavioral data was to select 96 abstract stimuli to serve as a baseline condition in the main fMRI experiment. To ensure the chosen stimuli were logical concept-property pairings, the accuracy of property verification responses were analyzed. To be selected for further analysis, each abstract concept-property pairing had to receive a correct property verification response from at least 75% of participants. Of the 182 abstract concept-property pairings, 136 received correct property verification responses from at least 75% of participants. To ensure the property verification task was equally difficult across visual, haptic and abstract conditions, the reaction times for property verifications were analyzed. First, the mean reaction time across participants was computed for each abstract concept-property pairing. Next, the mean reaction times

for verifying properties of concepts with visual and haptic features were computed. Finally, 96 abstract stimuli were chosen, such that the mean reaction time difference between the visual, haptic and abstract conditions was not significant ($p = 1.00$). Table 5.1 shows the reaction times for the chosen stimuli.




Figure 5.1 Experimental paradigm for behavioral experiment 2.

Table 5.1 Reaction times (in ms) for property verification task.

Visual		Haptic		Abstract	
<u>M</u>	<u>SD</u>	<u>M</u>	<u>SD</u>	<u>M</u>	<u>SD</u>
1445.81	200.78	1446.43	210.31	1447.75	206.22

5.4 SUMMARY

The purpose of the second behavioral experiment was to validate stimuli chosen for the main fMRI experiment to provide a baseline condition to control for the perceptual features of concept-property pairings chosen for the visual and haptic conditions. Due to their lack of perceptual features, abstract concept-property pairings were created to be used in the baseline condition. A behavioral experiment was conducted to select abstract stimuli which ensured the fMRI task was equally difficult across visual, haptic and abstract conditions. The pool of 182 abstract concept-property pairings was narrowed down to a final set of 96 stimuli which received correct property verification responses from at least 75% of participants and whose mean reaction time across participants did not differ significantly from the mean reaction times of the visual and haptic conditions. Therefore, the stimuli selected for the main experiment were determined to be equally difficult for visual, haptic, and abstract conditions, and differences between conditions cannot be explained by differences in task difficulty.

CHAPTER 6

FUNCTIONAL LOCALIZER

6.1 PURPOSE

Embodied theories hypothesize that concepts are represented in the brain regions responsible for acquiring perceptual information about their referents. These brain regions include primary and secondary perceptual areas as well as more anterior object-selective regions. The purpose of the main experiment was to determine whether these perceptual areas, those underlying visual and haptic perception, contain information about word stimuli with perceptual features. Rather than define regions of interest by anatomy, which varies greatly across individuals, a functional localizer was designed to isolate regions functionally. The functional localizer task was designed to isolate regions of the brain which underlie visual and haptic perception in general (primary and secondary visual and somatosensory areas) as well as regions which are selective for perceptual information pertaining to objects (LOC, FG, and IPS).

6.2 MATERIALS & METHODS

6.2.1 Participants

Participants were 18 healthy adults (12 females) ranging in age from 18 to 33 years ($M = 23.6$). Participants were native speakers of English, right-handed with normal or corrected-to-normal vision and no history of neurological impairments. All were recruited from the University of South Carolina community. Informed consent was

obtained from each participant prior to the experiment, in accordance with the protocol set forth by the University of South Carolina Institutional Review Board.

6.2.2 Stimuli

A functional localizer was employed to localize visual and haptic object-selective regions. The protocol was similar to Kim and James (2010). Color photographs of 18 objects and 18 textures were used for the visual object localizer run (Appendix A). Objects and textures were photographed from the same visual angle on a plain white background. Texture photographs were cropped to display only the texture with no background. All photographs were sized to 640 x 480 pixels. Eighteen 3-dimensional objects encountered in everyday life (e.g., balloon, shoe, etc.) and eighteen 2-dimensional surface materials (e.g., sandpaper, bubble wrap, etc.) were used for the haptic object localizer run. All objects and surface materials were MR-compatible and selected such that they were able to be explored using two hands.

6.2.3 Experimental paradigm

In the functional localizer, participants were presented with an object or texture one at a time and asked to covertly name the object or texture. Prior to the day of the experiment, participants practiced the functional localizer task in a mock scanner using different objects and textures to familiarize the participants with the procedure and to ensure the participants could perform the task without excessive head motion. Participants received a list of the names of objects and textures to be used in the real functional localizer but were not allowed to interact with them until scanning. This ensured that participants could accurately name the objects and textures but would not rely on their memory of the objects and textures for the purpose of identification.

The functional localizer was presented in a blocked design with four conditions: visual objects (VO), visual textures (VT), haptic objects (HO), and haptic textures (HT). In the visual conditions, participants saw photographs of objects and textures, and in the haptic conditions, participants explored objects and textures with both hands. The run consisted of eight visual blocks (4 VO, 4 VT) and 12 haptic blocks (6 HO, 6 HT). Each block was followed by a 12 second rest period. During the visual blocks, participants were presented with nine photographs in succession for 1.33 seconds each. During the haptic blocks, participants were presented with three items in succession, each for four seconds. Figure 6.1 shows the experimental paradigm for the localizer scan. Fewer visual blocks were needed, because visual stimuli were presented more rapidly than haptic stimuli.

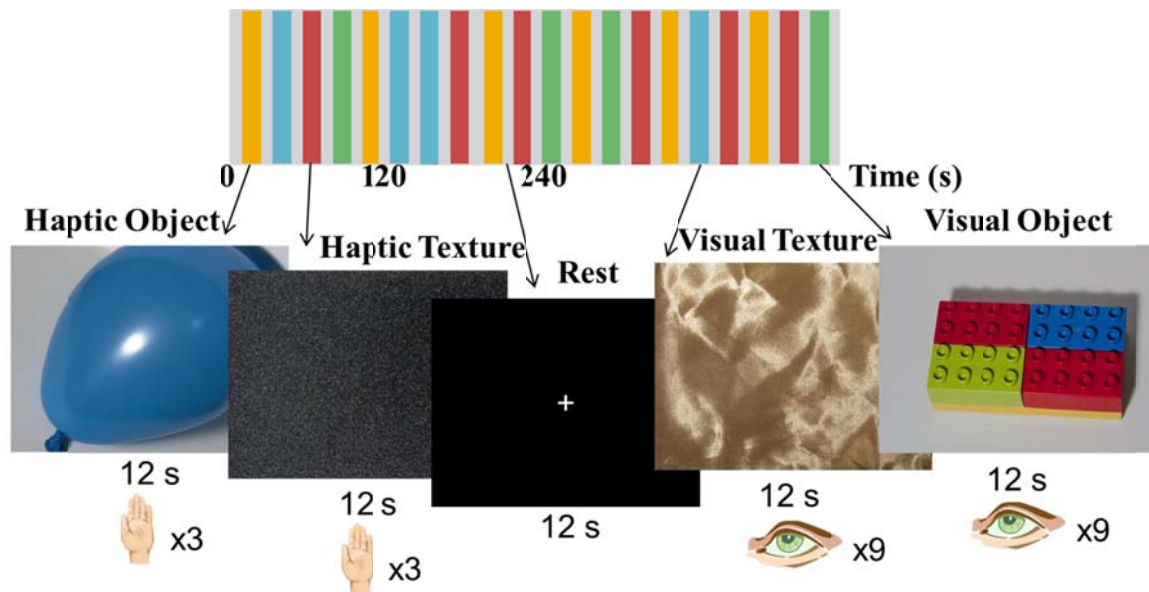


Figure 6.1 Experimental paradigm for functional localizer scan.

Two experimenters remained inside the scanner suite throughout the localizer run to present the haptic objects and textures. One experimenter was cued auditorily through headphones to hand participants objects and textures that were retrieved from a cart by

the second experimenter. Figure 6.2 shows the setup of the fMRI suite during the localizer scan. Participants were instructed to remain as still as possible while waiting for objects and textures to be placed in their hands. While viewing photographs and exploring objects and textures, participants covertly named the items. Photographs and auditory cues were presented using E-prime software (Psychology Software Tools, Sharspburg, PA).



Figure 6.2 Setup of the scanner suite during the functional localizer. Objects and textures were organized on a cart by experimental blocks (top). Two experimenters were positioned to relay objects and textures from the cart to the participants' hands while being cued auditorily of the timing of the experiment (bottom).

6.3 IMAGE ACQUISITION & PREPROCESSING

Functional images were acquired on a Siemens Magnetom Trio 3.0T scanner (Siemens, Erlangen, Germany) at the McCausland Center for Brain Imaging at the University of South Carolina. For the functional localizer, images were acquired using a gradient echo EPI pulse sequence (TR = 2200 ms, TE = 35 ms, flip angle = 90°). Thirty-six 3 mm thick oblique-axial slices were imaged with a 0.6 mm interslice gap, covering the whole brain, resulting in 3.0 x 3.0 x 3.0 mm voxels. Anatomical images of the entire brain were obtained using a standard T1-weighted 3D MP-RAGE protocol (TR = 2250 ms, TE = 4.15 ms, flip angle = 9°, voxel size = 1.0 x 1.0 x 1.0 mm).

Data preprocessing and the univariate statistical analysis was performed using Statistical Parametric Mapping 8 software (Wellcome Department of Cognitive Neurology, London, UK). The data was corrected for slice timing, motion, and linear trend, and a high-pass filter was applied (0.008Hz cut off). Functional images were spatially normalized to MNI space using a 12-parameter affine transformation and co-registered to the participant's anatomical image. Spatial smoothing was utilized for the univariate statistical analyses only with a Gaussian filter of 8 mm full-width-half-maximum.

For the univariate statistical analysis, a general linear model (GLM) was fit at each voxel using the canonical hemodynamic response function (HRF) convolved with onsets for each experimental condition, including six motion parameters as nuisance regressors. In order to isolate brain regions that process the visual and haptic features of objects, the following contrasts were computed: VO – VT (visual object-selective), HO – HT (haptic object-selective), and VO + HO – VT + HT (object-selective for either visual

or haptic), V – Fixation (visual objects and textures), and H – Fixation (haptic objects and textures). Additionally, for each participant, regions (cluster threshold of 5 voxels) showing significant activation differences ($p < 0.001$, unc.) for the five contrasts were used to create binary functional localizer masks. A sixth localizer mask was created for each participant containing regions that showed activation differences for both visual and haptic objects and textures greater than fixation. Functional localizer masks were used as regions of interest (ROIs) in subsequent analyses.

6.4 RESULTS

The purpose of the functional localizer task was to generate individual masks of functionally-localized regions for the main experiment; however, a group-level analysis was conducted to characterize which regions were represented. A group-level analysis of the fMRI results of the functional localizer task shows activation in many of the predicted visual and haptic perceptual regions found in previous neuroimaging studies. Table 6.1 shows the peak coordinates for regions showing activation differences for the contrasts of interest.

Table 6.1 Brain regions displaying significant ($p < 0.05$, FWE corrected) activation differences in functional localizer

Condition	Region	Talairach Coordinates				Voxel s	p value
		BA	x	y	z		
VO – VT	L Mid Occipital Gyrus	19	-42	-76	5	15	< 0.001
	L Inf Temporal Gyrus	37	-42	-68	-1		
HO – HT	R Precuneus	7	21	-48	56	148	< 0.001
	L Sup Parietal Lobule	7	-15	-50	61	53	< 0.001

	L Postcentral Gyrus	5	-24	-41	67		
	R Precuneus	19	36	-72	38	18	< 0.001
	R Mid Frontal Gyrus	6	27	-1	38	38	< 0.001
	L Mid Temporal Gyrus	37	-42	-59	-1	6	< 0.001
VO + HO - VT + HT	R Mid Temporal Gyrus	37	45	-62	-4	22	< 0.001
	L Postcentral Gyrus	2	-42	-36	58	21	< 0.001
	R Postcentral Gyrus	7	15	-50	64	19	< 0.001
	L Mid Temporal Gyrus	37	-42	-62	1	14	< 0.001
	L Postcentral Gyrus	7	-12	-50	64	11	< 0.001
	L Precuneus	7	-24	-75	43	10	< 0.001
	L Inf Occipital Gyrus	19	36	-77	0	8	< 0.001
V – Fixation	L Mid Occipital Gyrus	18	-21	-91	5	1704	< 0.001
	R Mid Occipital Gyrus	18	30	-80	0		
	R Lingual Gyrus	18	3	-80	2		
	L Sup Parietal Lobule	7	-30	-63	46	109	< 0.001
	R Sup Frontal Gyrus	6	3	19	55	73	0.001
	L Inf Frontal Gyrus	46	-48	29	8	36	0.005
	L Mid Frontal Gyrus	46	-50	24	27	32	0.001
	L Mid Frontal Gyrus	47	-45	42	-6	31	0.001
	L Precentral Gyrus	6	-50	1	48	30	0.001
	R Sup Parietal Lobule	7	33	-60	51	22	0.005
	L Mid Temporal Gyrus	21	-56	-33	-3	16	0.001

H – Fixation	L Postcentral Gyrus	3	-42	-22	49	2392	< 0.001
	R Postcentral Gyrus	4	36	-16	54		
	R Fusiform Gyrus	37	24	-49	-19	685	< 0.001
	L Insula	13	-39	-4	4	321	< 0.001
	R Insula	13	42	-1	4	97	< 0.001
	R Mid Frontal Gyrus	9	53	15	28	71	< 0.001
	L Inf Temporal Gyrus	19	-50	-59	-1	25	0.001

6.4.1 Visual Objects – Visual Textures

The VO – VT contrast is designed to isolate object-selective regions that process visual information. Differences in activation between VO and VT conditions were found in the left middle occipital gyrus (BA 19) and the left inferior temporal gyrus (BA 37).

Figure 6.3 depicts regions showing greater activation for visual objects than visual textures.




Figure 6.3 The visual features of objects were found to be processed in the left middle occipital gyrus (BA 19) and the left inferior temporal gyrus (BA 37).

6.4.2 Haptic Objects – Haptic Textures

The HO – HT contrast is designed to isolate object-selective regions that process haptic information. Differences in activation between HO and HT conditions were found in the right precuneus (BA 7 and 19), right middle frontal gyrus (BA 6), left superior

parietal lobule (BA 7), left postcentral gyrus (BA 5) and the left middle temporal gyrus (BA 37). Figure 6.4 depicts regions showing greater activation for haptic objects than haptic textures.




Figure 6.4 The haptic features of objects were found to be processed in the right precuneus (BA 7 and 19), right middle frontal gyrus (BA 6), left superior parietal lobule (BA 7), left postcentral gyrus (BA 5) and the left middle temporal gyrus (BA 37).

6.4.3 Visual Objects + Haptic Objects – Visual Textures + Haptic Textures

The VO + HO - VT + HT contrast is designed to isolate object-selective regions that process either visual or haptic information. Differences in activation were found in the right and left middle temporal gyrus (BA 37), the left and right postcentral gyrus (BA 2 and 7), the left inferior occipital gyrus (BA 19), and the left precuneus (BA 7). Figure 6.5 depicts regions showing greater activation for visual and haptic objects greater than visual and haptic textures.




Figure 6.5 Visual or haptic objects were found to be processed in the right and left middle temporal gyrus (BA 37), the left and right postcentral gyrus (BA 2 and 7), the left inferior occipital gyrus (BA 19), and the left precuneus (BA 7).

6.4.4 Visual – Fixation

The V – Fixation contrast is designed to isolate regions that process visual information, including visual information pertaining to both objects and textures. Differences in activation between the visual and fixation conditions were found in bilateral primary visual cortex (BA 18), bilateral superior parietal lobule (BA 7), bilateral premotor cortex (BA 6), left inferior and middle frontal gyri (BA 46/47), and left middle temporal gyrus. Figure 6.6 depicts regions showing greater activation for visual objects and textures greater than fixation.




Figure 6.6 Visual objects and textures were found to be processed in bilateral primary visual cortex (BA 18), bilateral superior parietal lobule (BA 7), bilateral premotor cortex (BA 6), left inferior and middle frontal gyri (BA 46/47), and left middle temporal gyrus.

6.4.5 Haptic – Fixation

The H – Fixation contrast is designed to isolate regions that process haptic information, including haptic information pertaining to both objects and textures. Differences in activation between the haptic and fixation conditions were found in bilateral somatosensory cortex (BA 3/4), bilateral insula (BA 13), right FG (BA 37), right middle frontal gyrus (BA 9), and left inferior temporal gyrus (BA 19). Figure 6.7 depicts regions showing greater activation for haptic objects and textures greater than fixation.




Figure 6.7 Haptic objects and textures were found to be processed in bilateral somatosensory cortex (BA 3/4), bilateral insula (BA 13), right FG (BA 37), right middle frontal gyrus (BA 9), and left inferior temporal gyrus (BA 19).

6.5 SUMMARY

The results of the functional localizer analysis replicate in part previous findings (Kim & James, 2010) and demonstrate that the visual and haptic features of objects activate the LOC, FG, and IPS, regions known to integrate visual and haptic perceptual information. Additionally, visual and haptic perceptual features, in general, elicit activation in the primary sensory cortices for visual and haptic perceptual information respectively.

The main findings of the functional localizer analysis show that visual and haptic features of objects activate the LOC, FG, and IPS. These results replicate previous findings in part but show some differences in the laterality of activation. Visual and haptic features of objects activated regions of the brain including the FG and LOC, which are known to be visual and haptic object-selective regions. The FG is likely a region that unifies object-specific information from auditory, visual, and haptic modalities into a trisensory representation (Kassuba et al., 2011) with visual information showing primacy over haptic information (Kassuba et al., 2013). The LOC is located at the convergence of visual and haptic streams of information and is thought to be bimodal visuo-haptic processing center (Amedi et al, 2005; Deshpande et al., 2010; James et al., 2005; James et

al., 2007; Lacey et al., 2010; Lacey et al., 2009; Kassuba et al., 2013). In addition to the LOC and FG, the visual and haptic features of objects were processed in BA 2, which includes the portion of the primary somatosensory cortex specializing in size and shape processing. Size and shape are two object features that are largely bimodal. One can see and feel the size and shape of an object.

In addition to the FG and LOC, haptic features of objects activated the IPS and motor regions. The IPS is a region bounded by BA 5 and BA 7, which is located at the convergence of visual and haptic streams of information and is thought to be a bimodal visual-haptic processing center (James et al., 2007; Kim & James, 2010). Reflecting an increased requirement for movement planning, haptic features of objects activated the pre-motor cortex and supplementary motor areas. This may be due to the fact that objects required more manipulation and rotation than textures to identify.

Lateralization differences were found in the current study in comparison to Kim and James (2010). Visual object processing was left-lateralized in the FG and LOC rather than bilateral. In contrary, haptic object processing was bilateral in the LOC and IPS rather than left-lateralized. However, activation in the motor areas for haptic objects was right-lateralized rather than bilateral. The differences in lateralization between this study and the previous study may be due to minor differences in stimuli and/or method of presentation. Stimuli were designed to be held comfortably in two hands and have discernible textures and shapes that could be easily recognized, but they were entirely different from the previous study. Objects and textures were presented for haptic exploration to both hands, which leaves questions as to why activation may have been right-lateralized in motor areas.

Processing visual stimuli, objects and textures, elicited activation in primary visual areas in addition to the IPS. The activation of the IPS may be the result of secondary activation of the haptic representation of the objects and textures being seen, possibly to create a unified experience of the object or texture by imagining other perceptual features of the stimulus. Other areas of the brain included the frontal eye fields, which may play a role in generating the contents of visual perception (Libedinsky & Livingstone, 2011).

Processing haptic stimuli, objects and textures, elicited activation in primary somatosensory areas in addition to the FG. The activation of the FG may be the result of secondary activation of the visual representation of the objects and textures being touched. Once again participants may have imagined the other perceptual features of stimuli to create a unified perceptual experience. Other areas of the brain included the bilateral insula, implicated as a non-primary motor area responsive to finger movements (Fink et al., 1997).

In conclusion, the results of the functional localizer analysis indicate that touching and seeing objects elicits activation in object-selective perceptual regions, such as the LOC, FG, and IPS. Touching and seeing objects and textures activates primary sensory areas in addition to some bimodal visual-haptic regions. The latter presumably reflects that objects and textures may be imagined in other sense modalities to create a unified and more complete perceptual experience.

CHAPTER 7

MAIN EXPERIMENT

7.1 PURPOSE

The purpose of the main experiment was to demonstrate that the representation of concepts with perceptual features is more consistent with weak and strong embodiment theories than unembodied and secondary embodiment theories. The central hypothesis was that the neural representation of concepts with perceptual features is distributed and includes brain regions in the perceptual systems activated when interacting with the referent of that concept. More specifically, concepts containing visual information should be represented in brain regions active when processing visual stimuli, while concepts containing haptic information should be represented in brain regions active when processing haptic stimuli.

The goals of the main experiment were two-fold. The first goal was to examine which brain regions participate in processing concepts with visual and haptic features. The second goal was to determine whether information about the perceptual content of concepts is present in patterns of brain activity elicited by processing concepts with visual and haptic features. To accomplish the first goal, we conducted a univariate analysis to investigate which brain regions respond more to processing concepts with visual or haptic features than concepts with more abstract features. Additionally, we examined the patterns of functional connectivity of these regions to characterize the functional networks recruited to process concepts with perceptual features. To

accomplish the second goal, we utilized MVPA to determine whether patterns of brain activity elicited by processing concepts can be used to predict the perceptual information content of a concept.

7.2 MATERIALS & METHODS

7.2.1 Participants

Participants were 18 healthy adults (12 females) ranging in age from 18 to 33 years ($M = 23.6$). Participants were native speakers of English, right-handed with normal or corrected-to-normal vision and no history of neurological impairments. All were recruited from the University of South Carolina community. Informed consent was obtained from each participant prior to the experiment, in accordance with the protocol set forth by the University of South Carolina Institutional Review Board.

7.2.2 Stimuli

A set of 192 visual and haptic concept-property word pairings were selected from a database of 774 multi-modal concept-property items from Dantzig, Cowell, Zeelenberg, and Pecher (2011). Of the 192 visual concept-property pairings, 96 contained visual information, and 96 contained haptic information. Concept-property pairings were rated for how strongly each is experienced with five sensory modalities (sight, sound, touch, smell and taste) through a series of norming studies. The concept properties with the highest modality exclusivity ratings for vision and haptics were chosen to ensure stimuli were as unimodal as possible (threshold of 65% or higher for vision and 35% for haptics). Haptic stimuli are inherently more multi-modal, and the threshold for modality exclusivity reflects this. Additionally, 96 abstract stimuli were constructed by choosing frequently used abstract nouns and pairing these with commonly used descriptors from a

thesaurus. Word stimuli were balanced for average length ($p = 0.351$) and average frequency ($p = 0.061$).

7.2.3 Questionnaire

Following the main experiment participants completed the Vividness of Visual Imagery Questionnaire (VVIQ; Marks, 1973) to evaluate imagery ability. This questionnaire consists of 16 questions with 5 response choices to evaluate the degree of clarity with which a participant is able to imagine a scenario. Lower scores on the VVIQ indicate more vivid visual imagery. Cui, Jeter, Yang, Montague, and Eagleman (2007) demonstrates that the vividness of mental imagery correlates with the activation levels in the visual cortex ($r = -0.73, p = 0.04$). The questionnaire was administered after completing the main experiment to avoid influencing the participants to imagine the stimuli presented in the main experiment.

7.2.4 Experimental paradigm

During scanning participants performed a perceptual property verification task similar to tasks used in behavioral and neuroimaging studies of conceptual processing (Goldberg et al., 2006; Pecher, Zeelenberg, & Barsalou, 2003). On any given trial, participants were asked to decide which of two properties best described a concept from either the visual (V), haptic (H), or abstract (A) categories. In the visual and haptic conditions, the two properties included perceptual features. For example, given the concept “ZEBRA” and the visual properties “STRIPED” and “RED,” the participant would choose “STRIPED” as the applicable property, because a zebra can be striped but not red. In the abstract condition, the two properties included non-perceptual features. For example, given the concept “LOSS” and the abstract properties “SAD” and

“SECURE,” the participant would choose “SAD” as the applicable property, because loss can make one feel sad but not secure. This task was designed to prompt the participant to form a simulation of both the concept and its properties, which may involve sensory-motor processing (Dantzig et al., 2011). The number of times a property was used as the correct choice was balanced with the number of times it was used as the incorrect choice. Additionally, half of all trials had the correct choice listed on the right, while half had the correct choice listed on the left. Property verification decisions were blocked by modality with four consecutive trials of each type. The concept and property choices were presented for 3000 ms followed by a 1000 ms fixation cross (Figure 7.1). Twenty-four blocks of each modality type, 16 s in duration, were presented over two sessions. This number of blocks per condition is recommended for use with blocked designs to ensure enough trials for MVPA when temporally averaging normalized signal intensity values (Kamitani & Tong, 2005). Fixation blocks were presented for 10 s each before and after each block to reduce overlap in the brain signal between experimental conditions.

7.3 FMRI IMAGE ACQUISITION

Functional images were acquired on a Siemens Magnetom Trio 3.0T scanner (Siemens, Erlangen, Germany) at the McCausland Center for Brain Imaging at the University of South Carolina. For the main experiment, images were acquired using a gradient echo EPI pulse sequence (TR = 1100 ms, TE = 35 ms, flip angle = 64°). Eighteen 5.4 mm thick oblique-axial slices were imaged with a 0.54 mm interslice gap, covering the whole brain, resulting in 3.3 x 3.3 x 5.4 mm voxels. Anatomical images of

the entire brain were obtained using a standard T1-weighted 3D MP-RAGE protocol (TR = 2250 ms, TE = 4.15 ms, flip angle = 9°, voxel size = 1.0 x 1.0 x 1.0 mm).




Figure 7.1 Experimental paradigm for main experiment.

7.4 DATA PROCESSING & ANALYSIS

Data preprocessing and the univariate statistical analysis was performed using Statistical Parametric Mapping 8 software (Wellcome Department of Cognitive Neurology, London, UK). The data was corrected for slice timing, motion, and linear trend, and a high-pass filter was applied (0.008Hz cut off). Functional images were spatially normalized to MNI space using a 12-parameter affine transformation and co-registered to the participant's anatomical image.

7.4.1 Univariate analysis

For the univariate statistical analysis, a general linear model (GLM) was fit at each voxel using the canonical hemodynamic response function (HRF) convolved with onsets for each experimental condition, including six motion parameters as nuisance regressors. Spatial smoothing was utilized for the univariate statistical analyses only with a Gaussian filter of 8 mm full-width-half-maximum. In order to isolate brain regions that process the visual and haptic features of word stimuli, the following contrasts were used: V+H – A (perceptual), V-A (visual), and H-A (haptic).

7.4.2 Pattern classification

The percent signal change (PSC) relative to the average activity in a voxel was computed for each voxel in every volume. The mean PSC of six volumes, offset 4.4 seconds (TR = 1.1 s) from the stimulus onset (to account for the delay in hemodynamic response), was used as the input for further analyses. Furthermore, the mean PSC data for each voxel was standardized to have a mean of zero and variance of one.

Classifiers were trained to identify cognitive states from the pattern of brain activity (mean PSC) elicited by verifying the properties of concepts from three categories. Two-category classification was performed to identify cognitive states associated with verifying concepts with visual or abstract, haptic or abstract, and visual or haptic, and visual and/or haptic or abstract features. For classification, classifiers were defined as a function $f: mean_PSC \rightarrow Y_j, j = \{1, \dots, k\}$, where k was the number of categories used for classification, Y_j were categories of visual, haptic, or abstract features and where $mean_PSC$ was a vector of mean PSC voxel activations.

Prior to classification, trials were divided into training and test sets, and relevant features (voxels) were extracted (see below for feature selection method) from the training set only. The classifier was constructed using the selected features from the training set. The classifier was applied subsequently to the unused test set and classification performance was evaluated with cross-validation.

To reduce the size of the data, a discriminative-based feature selection method was used. For each fold of the data, a classifier was trained using the data from one voxel at a time to obtain a classification accuracy for discriminating between the two conditions of interest. Voxels were ordered by classification accuracy, and the most discriminating voxels were chosen for classification. The impact of retaining different numbers of voxels on each analysis was explored, rather than deciding upon an arbitrary threshold.

A logistic regression classifier was used for classification (Bishop, 2006). Logistic regression is a widely used classifier that learns the function $f: P(Y|X)$, where Y is discrete dependent variable, and X is a vector containing discrete or continuous variables. By using the maximum likelihood estimation, this algorithm estimates the probability of the given data belonging to an output category and classifies the data into the most probable category. As a classifier, logistic regression directly estimates its parameters from the training data. Twenty-four fold cross-validation was used to evaluate classification performance, where each fold corresponded to one block of each of the conditions. Thus, the classifier was trained on 23 presentations and tested on one presentation. Classification was repeated iteratively until each presentation served as the test set once. Classification accuracies were computed based on the average classification accuracy across test folds. As a result, classification accuracy was always

based upon the test data only, which remained disconnected from the training data.

Classification procedures were conducted similarly to previous works investigating the neural representation of concepts (Baucom, Wedell, Wang, Blizter, & Shinkareva, 2012; Wang, Baucom & Shinkareva, 2012)

If classification is successful, accuracies should be significantly different from the chance level accuracy, i.e. the accuracy of guessing. The significance of classification accuracy was evaluated based on the binomial distribution $B(n, p)$, where n is the number of trials of each classification computation and p is the probability of correct classification when the exemplars are randomly labeled (Pereira et al., 2009).

To determine whether visual and haptic object-selective regions carry information about the visual and haptic features of concepts, an ROI-based classification analysis was also performed. A binary mask was generated for each participant by selecting for regions (cluster threshold of 5 voxels) showing significant activation differences for any of the three contrasts from the univariate analysis of the functional localizer data. The binary mask was applied to the main experiment data and used as input for classification. Classification, feature selection, and cross-validation were conducted in the same manner as the whole brain pattern classification. The significance of classification accuracy was evaluated based on the binomial distribution.

To establish commonalities between participants' neural representations of concepts with perceptual features, cross-participant classification was conducted. Data from all but one participant were used to train a classifier to distinguish cognitive states associated with each experimental condition. The classifier was then tested on the data of the left-out participant. Classification was repeated iteratively until each participant's

data served once as the test set. The significance of classification accuracy was evaluated based on the binomial distribution.

To investigate the consistency of informative voxels across individuals for cross-participant classification, a voxel location probability map was generated across participants after convolving each voxel with a 4 mm Gaussian kernel (Kober et al., 2008). The probability map was further thresholded by a simulated null hypothesis distribution to control for multiple comparisons (FWE = 0.05).

7.4.3 Functional connectivity

The task-related functional connectivity of brain regions was investigated in a similar manner to Rissman, Gazzaley, and D'Esposito (2004). Following the univariate analysis of the functional localizer data, a "seed" region was selected to investigate how other brain regions interact with it during each condition of the main experiment. The occipitotemporal cortex was selected to serve as the seed region, as it would be hypothesized to show differential activation for concepts with visual, haptic and abstract features based on the functional localizer. This brain region was shown to be selective for objects with either visual or haptic features. The seed region was identified separately for each participant in MNI space by masking the participant's data with a binary ROI mask of the bilateral occipitotemporal cortex based on the Talairach Daemon database (Lancaster et al., 2000), generated with the WFU Pickatlas (Maldjian, Laurienti, Kraft & Burdette, 2003). Next, the condition-specific beta values (or "beta series"; Rissman et al., 2004) of each voxel in the brain was computed for each trial to estimate the magnitude of the task-related BOLD response. The beta series averaged across the selected voxels in the seed region was correlated with the beta series of all other voxels in the brain to

quantify the extent that each pair of voxels interacted with each other during each condition of the task. The more highly correlated the voxels were, the greater the voxels interacted during the condition of the task. Finally, the correlation coefficients were transformed to Fisher's z-scores, mapped for each participant in MNI space, and submitted to a random effects group level analysis for each condition using Statistical Parametric Mapping 8 software (Wellcome Department of Cognitive Neurology, London, UK) to determine which correlation coefficients were significantly greater than zero.

7.4.4 Connectivity-based MVPA

Cross-participant MVPA was performed on the seed-based connectivity matrices using the occipitotemporal cortex as the seed region. Pattern classification was used to test for cross-participant consistencies of the patterns for visual, haptic and abstract conditions. A similarity-based classifier was trained on data from all but one participant to identify the connectivity matrices for the left out participant. Classification was performed iteratively until each participant's data served as the test set once. To reduce the size of the data, feature selection was used. To select connections that responded to the experimental conditions, matrices in the training set were first transformed to Fisher's z-scores. One sample t-tests against the null hypothesis of no response were then performed for each connection across all the participants in the training set for each condition separately. The connections with the highest t-values in either condition were selected jointly for both conditions, so that the feature selection was orthogonal to the classification categories.

For the training set, weighted average matrices for each condition were generated by weighting each participant's matrix by how similar they were to each other (Abdi,

Dunlop, & Williams, 2009; Shinkareva, Malave, Mason, Mitchell, & Just, 2011; Shinkareva, Ombao, Sutton, Mohanty, & Miller, 2006). Pairwise similarity between participants was measured by the RV coefficient (Robert & Escoufier, 1976), a multivariate generalization of the Pearson correlation coefficient to matrices. Each participant's data were scaled by the first eigenvector of the similarity matrix to sum up to one.

For each test matrix, the cosine similarity scores were computed, and the test matrix was labeled according to the training condition with the higher similarity score (Mitchell et al., 2008). When the hit score was higher than the miss score across the two conditions, classification was evaluated as successful. The overall classification accuracies were averaged across participants.

7.5 UNIVARIATE ANALYSIS RESULTS

The fMRI results of the main experiment show activation in many predicted brain regions found in previous neuroimaging studies of visual and haptic object perception and conceptual representation of words with visual and haptic features. Table 7.1 shows the peak coordinates for regions showing activation differences for the three contrasts of interest.

Table 7.1 Brain regions displaying significant ($p < 0.05$, FWE corrected) activation differences in main experiment

Condition	Region	<u>Talairach Coordinates</u>				Voxels	<i>p</i> value
		BA	x	y	z		
V + H – A	L Fusiform Gyrus	20	-29	-36	-33	19	< 0.001
	L Fusiform Gyrus	37	-51	-58	-12	5	0.003

V – A	L Fusiform Gyrus	20	-25	-36	-23	20	< 0.001
H – A	L Inf Frontal Gyrus	46	-45	33	12	11	0.001
	L Fusiform Gyrus	20	-29	-36	-33	8	0.002
	L Fusiform Gyrus	37	-48	-57	-12	5	0.005

7.5.1 Visual + Haptic – Abstract

The V + H – A contrast is designed to isolate regions of the brain that are involved in processing the perceptual (visual and haptic) features of word stimuli. Differences in activation between the perceptual and abstract conditions were found in the left FG (BA 20) and left LOC (BA 37). These regions are visual and haptic object-selective regions and fall in line with the results of previous studies investigating visual and haptic object perception (Deshpande et al., 2010; James et al., 2005; James & Kim, 2010; James et al., 2007; Kim & James, 2010; Lacey, Campbell, & Sathian, 2007; Lacey et al., 2009). Kim and James (2010) also found activation differences in the left IPS for visual and haptic object perception, but this region was absent in the current analysis. Figure 7.2 depicts the brain regions showing greater activation for perceptual than abstract features.




Figure 7.2 The visual and haptic features of concepts were found to be processed in the left FG (BA 20) and left LOC (BA 37).

7.5.2 Visual – Abstract

The V – A contrast is designed to isolate regions of the brain that are involved in the processing the visual features of word stimuli. Differences in activation between the visual and abstract conditions were found in the left FG (BA 20). Previous studies investigating visual and haptic object perception have found bilateral activation in the FG, LOC, and the IPS (Kim & James, 2010). Additionally, the FG (and the surrounding ventral temporal lobe) has been found to be involved in processing shape-related words (Pulvermüller & Hauk, 2006), color word generation (Martin et al., 1995) and color property verification (Goldberg et al., 2007). Figure 7.3 depicts the brain regions showing greater activation for visual features than abstract features.




Figure 7.3 The visual features of concepts were found to be processed in the left FG (BA 20).

7.5.3 Haptic – Abstract

The H – A contrast is designed to isolate regions of the brain that are involved in processing the haptic features of word stimuli. Differences in activation between the haptic and abstract conditions were found in the left inferior frontal gyrus (BA 46), left FG (BA 20), and left LOC (BA 37). Previous studies investigating visual and haptic object perception have found activation in the left FG and the left LOC (Kim & James,

2010). Conceptual representation of words with haptic features has previously implicated the motor cortex, premotor cortex, and the primary somatosensory cortex (Goldberg et al., 2007); however, these regions were not found to show activation differences in the current study. Finally, the left inferior frontal gyrus has been found previously to process action words and verbs (Martin et al., 1995). No other studies have implicated this region in the representation of words with haptic features. Figure 7.4 depicts the brain regions showing greater activation for haptic features than abstract features.

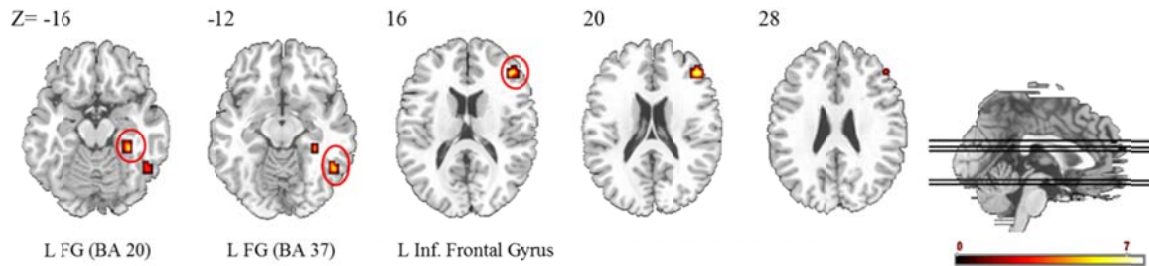


Figure 7.4 The haptic features of concepts were found to be processed in the left inferior frontal gyrus (BA 46), left FG (BA 20), and left LOC (BA 37).

7.6 PATTERN CLASSIFICATION RESULTS

7.6.1 ROI-based classification

A classifier was trained for each participant to determine if it was possible to identify whether a concept contained visual, haptic or abstract features based on the activation elicited by touching and feeling objects and textures during the functional localizer scan. Visual or haptic object-selective functionally-defined ROIs as well as general visual, haptic, and visual-haptic functionally-defined ROIs were used for classification. Feature selection thresholds were based on a percentage of the most discriminative voxels within the object-selective ROIs due to the high variability in the

number of active voxels across participants. Feature selection thresholds were based on a set number of discriminative voxels within general visual, haptic, and visual-haptic ROIs.

For regions which showed the greatest activation for visual or haptic objects, classification accuracies for classifying visual vs. abstract features exceeded chance level (0.50) for all levels of the most discriminative voxels ($p < 0.05$) for the majority of participants (Figure 7.5a). The highest classification accuracy obtained for a single participant was 0.83. Classification accuracies for classifying visual and haptic vs. abstract features exceeded chance level for the top 10% and 25% discriminative voxels only (Figure 7.5b). The highest classification accuracy obtained for a single participant was 0.80. Classification was unsuccessful in regions which showed the greatest activation for visual objects and haptic objects alone. Classification accuracies were consistent across people for classifying visual vs. abstract and visual and haptic vs. abstract concepts, such that participants with the highest and lowest classification accuracies for one classification problem had the highest and lowest classification accuracies on the other ($r = 0.811, p < 0.001$).

For regions which showed the greatest activation for visual or haptic objects, VVIQ scores and accuracies for classifying visual vs. abstract features showed a significant negative correlation, such that higher classification accuracies were associated with lower VVIQ scores ($r = -0.424, p < 0.05$; Figure 7.6). Lower VVIQ scores indicate a participant's ability to vividly imagine a scene.

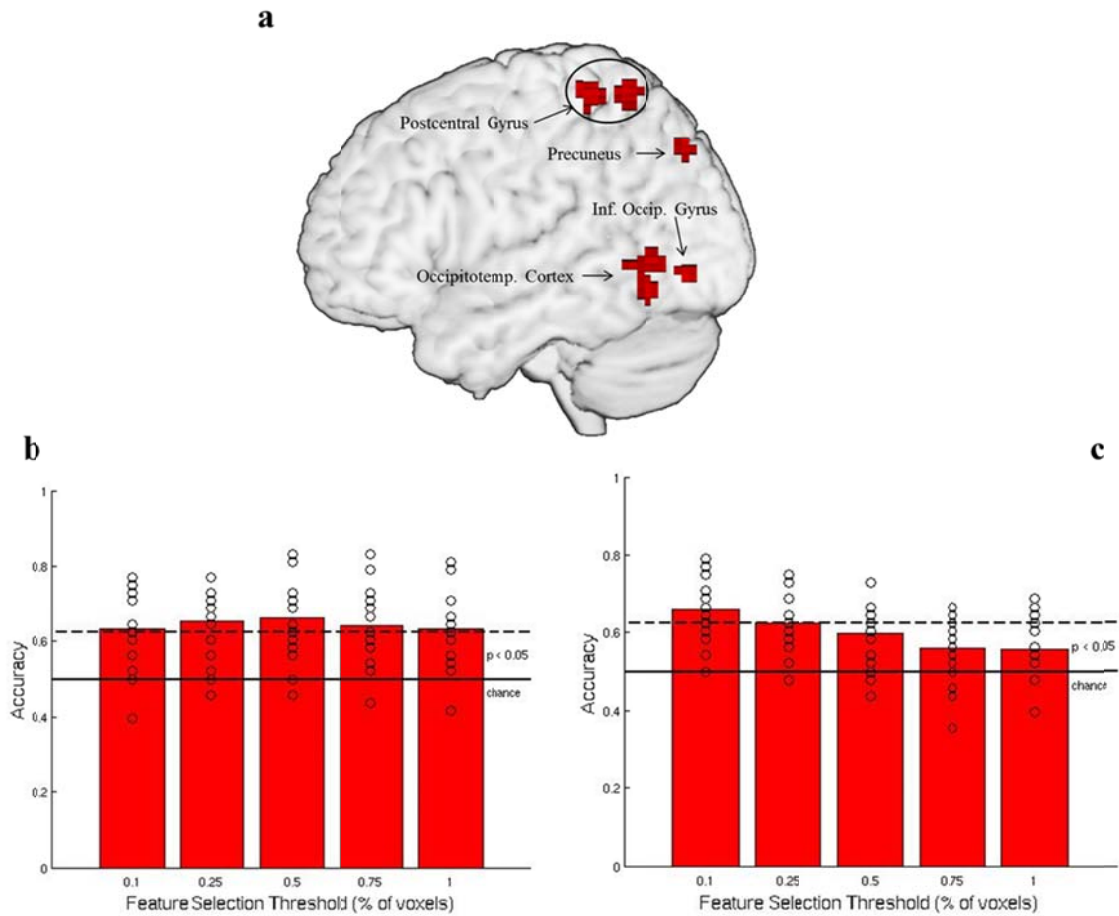


Figure 7.5 Accuracies for classification within regions selective for the visual and haptic features of objects. a) Regions in red showed significantly more activation when processing the visual and haptic features of objects. b) Classification accuracies for classifying visual vs. abstract concepts. c) Classification accuracies for classifying visual and haptic vs. abstract concepts.




Figure 7.6 Classification accuracies for classifying visual vs. abstract features within visual or haptic object-selective regions showed a significant negative correlation with VVIQ scores. Classification accuracies were higher for participants who reported the ability to imagine scenes more vividly.

For regions which showed the greatest activation for visual stimuli, objects and textures, classification accuracies for classifying visual vs. abstract features and visual and haptic vs. abstract features exceeded chance level (0.50) for all levels of the most discriminative voxels ($p < 0.05$) for the majority of participants (Figure 7.7). The highest classification accuracy obtained for a single participant was 0.84 for both classification problems. Accurate classification was robust across the range of voxels used (from 25 to 400). Classification accuracies for classifying haptic vs. abstract features exceeded chance level for most levels of the most discriminative voxels for the majority of participants (Figure 7.7). The highest classification accuracy obtained for a single participant was 0.82.

For regions which showed the greatest activation for haptic stimuli, objects and textures, classification accuracies for classifying visual vs. abstract features and visual

and haptic vs. abstract features exceeded chance level (0.50) for most levels of the most discriminative voxels ($p < 0.50$) for the majority of participants (Figure 7.7). The highest classification accuracy obtained for a single participant was 0.82 and 0.83 for visual vs. abstract and visual and haptic vs. abstract respectively.

For regions which showed the greatest activation for both visual and haptic stimuli, objects and textures, classification accuracies for classifying visual vs. abstract features and visual and haptic vs. abstract features exceeded chance level (0.50) for all levels of the most discriminative voxels ($p < 0.05$) for the majority of participants (Figure 7.7). The highest classification accuracy obtained for a single participant was 0.83 for both classification problems. Accurate classification was robust across the range of voxels used (from 25 to 400). Classification accuracies for classifying haptic vs. abstract features exceeded chance level for most levels of the most discriminative voxels for the majority of participants (Figure 7.7). The highest classification accuracy obtained for a single participant was 0.82.

Participant classification accuracies showed consistency across classification problems within and across general perceptual regions as measured by correlation (Figure 7.8). Classification accuracies for visual vs. abstract and visual and haptic vs. abstract were consistent across participants within visual, haptic, and visual-haptic perceptual regions, while haptic vs. abstract was consistent with visual vs. abstract and visual and haptic vs. abstract within the visual-haptic perceptual regions only. Classification accuracies for visual vs. abstract were consistent across all perceptual regions, while classification accuracies for visual and haptic vs. abstract were consistent across visual

and haptic perceptual regions only. Classification accuracies for haptic vs. abstract were inconsistent across all perceptual regions.

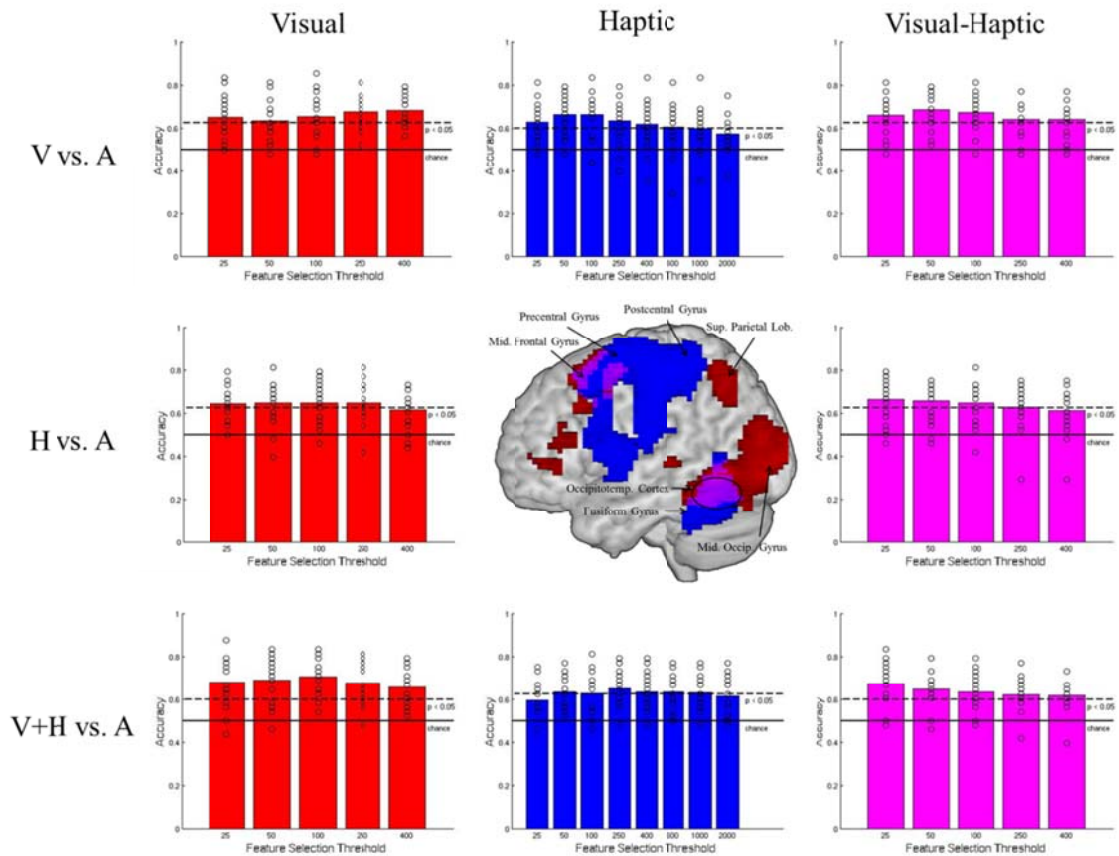


Figure 7.7 Accuracies for classifying within regions responsive to processing general visual (red), haptic (blue), and visual and haptic (magenta) perceptual features.

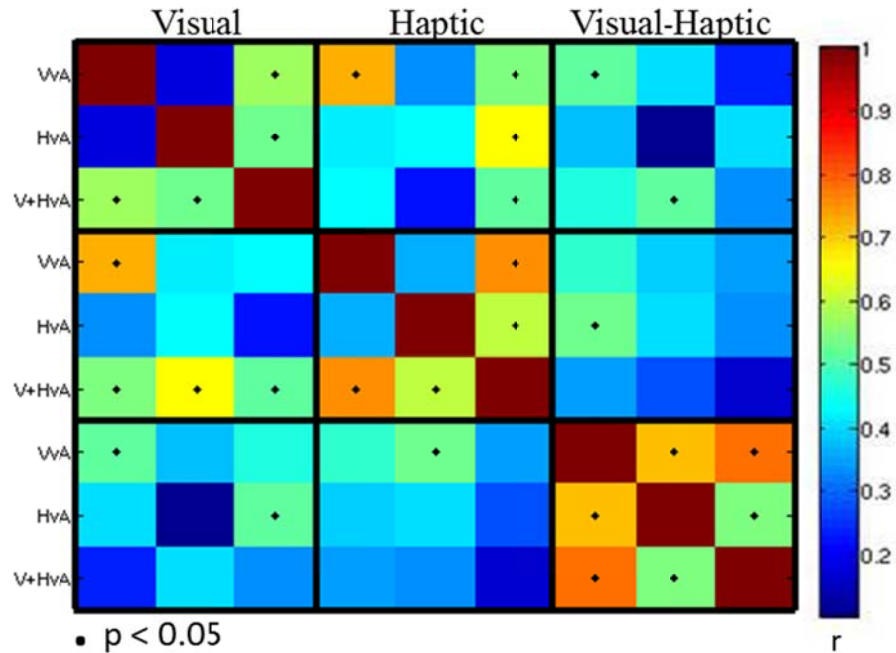


Figure 7.8 Consistency of participant classification accuracies across classification problems and perceptual regions measured by correlation.

Due to the large number of regions in which classification of the perceptual features of concepts was possible, a control region was tested to ensure significant classification accuracies were the result of the classifier detecting information about the perceptual features of concepts. BA 40 in the right hemisphere was chosen, because it has been demonstrated previously to not contain information about concrete or abstract concepts (Wang, Baucom & Shinkareva, 2012) while also being bounded by regions in which successful classification occurs. Classification of visual vs. abstract, haptic vs. abstract and visual and haptic vs. abstract was unsuccessful in right BA 40.

7.6.2 Whole brain classification

A classifier was trained for each participant to determine if it was possible to identify whether a concept contained visual, haptic or abstract features based on whole

brain activation elicited by verifying features of concepts. Feature selection thresholds were based on a set number of discriminative gray matter voxels.

Classification accuracies for classifying visual vs. haptic vs. abstract features exceeded chance levels (0.33) for the smaller levels (from 25 to 250) of the most discriminative voxels ($p < 0.05$) for the majority of participants (Figures 7.9). The highest classification accuracy for a single participant was 0.68. An examination of the confusion matrices, based on 100 of the most discriminative voxels, for each participant shows that the classifier most often made errors when classifying haptic features, confusing these with visual features (Figure 7.10).

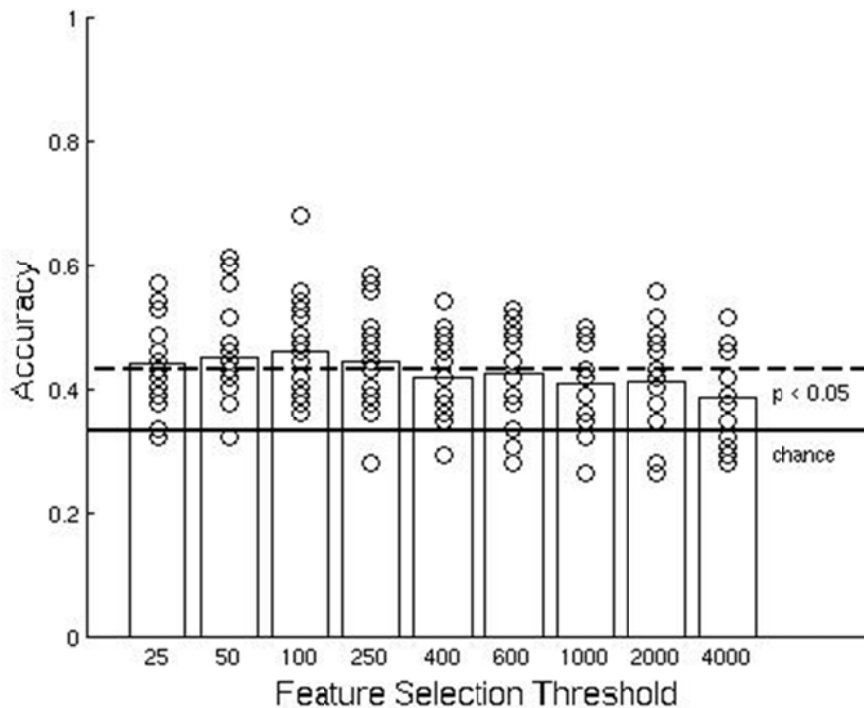


Figure 7.9 Within-participant accuracies for classifying visual vs. haptic vs. abstract features from whole brain patterns of activity. Classification accuracies across the 18 participants, mean accuracy summarized by bars with individual accuracies represented by open circles, are shown for different subsets of the most discriminative gray matter voxels.

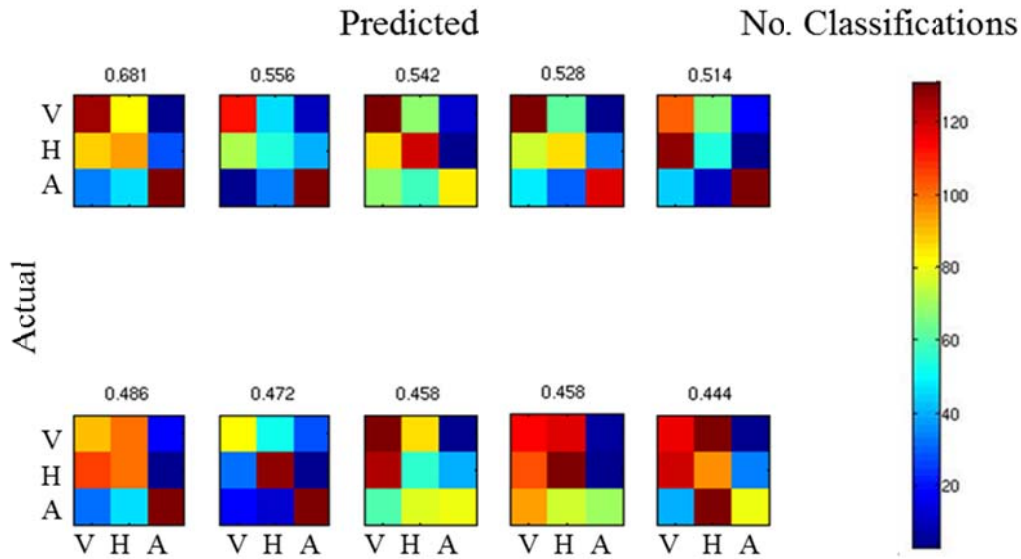


Figure 7.10 Participant confusion matrices for classifying visual vs. haptic vs. abstract features from whole brain patterns of activity, ordered by average accuracy across folds for 100 most discriminative voxels. The classifier most often confuses haptic features for visual features.

Classification accuracies for classifying visual vs. abstract features exceeded chance level (0.50) for all levels of the most discriminative voxels ($p < 0.05$) for the majority of participants (Figure 7.11). The highest classification accuracy obtained for a single participant was 0.85. Accurate classification was robust across the range of voxels used (from 25 to 4000).

Classification accuracies for classifying haptic vs. abstract features exceeded chance level (0.50) for moderate levels (from 100 to 400) of the most discriminative voxels ($p < 0.05$) for the majority of participants (Figure 7.12). The highest classification accuracy obtained for a single participant was 0.82.




Figure 7.11 Within-participant accuracies for classifying visual vs. abstract features from whole brain patterns of activity. Classification accuracies across the 18 participants, mean accuracy summarized by bars with individual accuracies represented by open circles, are shown for different subsets of the most discriminative gray matter voxels.




Figure 7.12 Within-participant accuracies for classifying haptic vs. abstract features from whole brain patterns of activity. Classification accuracies across the 18 participants, mean accuracy summarized by bars with individual accuracies represented by open circles, are shown for different subsets of the most discriminative gray matter voxels.

Classification accuracies for classifying visual and haptic vs. abstract features exceeded chance level (0.50) for all levels of the most discriminative voxels ($p < 0.05$) for the majority of participants (Figure 7.13). The highest classification accuracy obtained for a single participant was 0.84. Accurate classification was robust across the range of voxels used (from 25 to 4000).




Figure 7.13 Within-participant accuracies for classifying visual and haptic vs. abstract features from whole brain patterns of activity. Classification accuracies across the 18 participants, mean accuracy summarized by bars with individual accuracies represented by open circles, are shown for different subsets of the most discriminative gray matter voxels.

Classification accuracies were consistent across people for classifying visual vs. abstract and visual and haptic vs. abstract concepts from whole brain patterns of brain activity, such that participants with the highest and lowest classification accuracies for one classification problem had the highest and lowest classification accuracies on the other ($r = 0.715, p < 0.001$). Classification accuracies for haptic vs. abstract were not

statistically correlated with classification accuracies for visual vs. abstract or visual and haptic vs. abstract.

7.6.3 Cross-participant classification

To examine the consistency of the neural representations of concepts with perceptual features across participants, whole brain activation data from all but one participant were used to identify the category of stimuli presented to the left-out participant. A classifier was trained on the data from all but one participant and tested on the data from the left-out participant. Feature selection thresholds were based on a set number of discriminative gray matter voxels common to all participants. The highest accuracy for classifying visual vs. abstract features obtained for any voxel level was 0.70 (compared to 0.50 chance level). Classification accuracies for classifying visual vs. abstract features were significant for some levels of the most discriminative voxels ($p < 0.05$) for the majority of participants (Figure 7.14).

A classifier was trained on the combined data from all but one participant to identify haptic vs. abstract features for the left-out participant. The highest accuracy for classifying haptic vs. abstract features obtained for any voxel level was 0.75. Classification accuracies for classifying haptic vs. abstract features were significant for most levels of the most discriminative voxels ($p < 0.05$) for the majority of participants (Figure 7.15).

A classifier was trained on the combined data from all but one participant to identify visual and haptic vs. abstract features for the left-out participant. The highest accuracy for classifying visual and haptic vs. abstract features obtained for any voxel level was 0.78. Classification accuracies for classifying visual and haptic vs. abstract

features were significant for all levels of the most discriminative voxels ($p < 0.05$) for the majority of participants (Figure 7.16).




Figure 7.14 Cross-participant accuracies for classifying visual vs. abstract features from whole brain patterns of activity. Classification accuracies across the 18 participants, mean accuracy summarized by bars with individual accuracies represented by open circles, are shown for different subsets of the most discriminative gray matter voxels common to all participants.

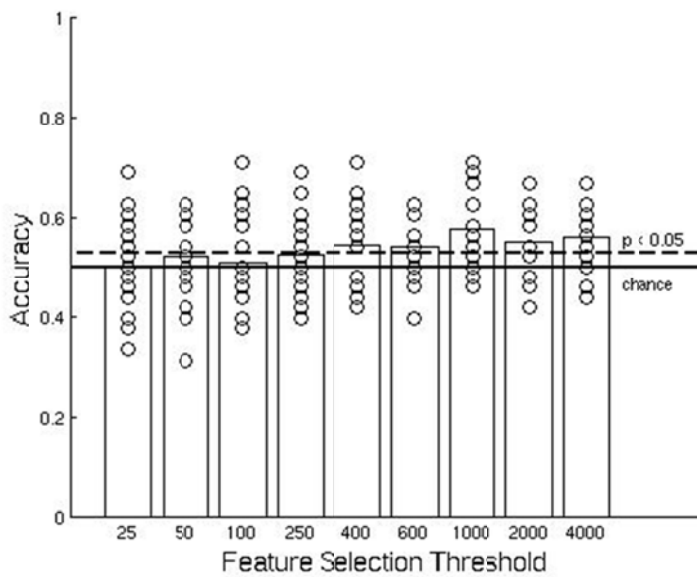


Figure 7.15 Cross-participant accuracies for classifying haptic vs. abstract features from whole brain patterns of activity.

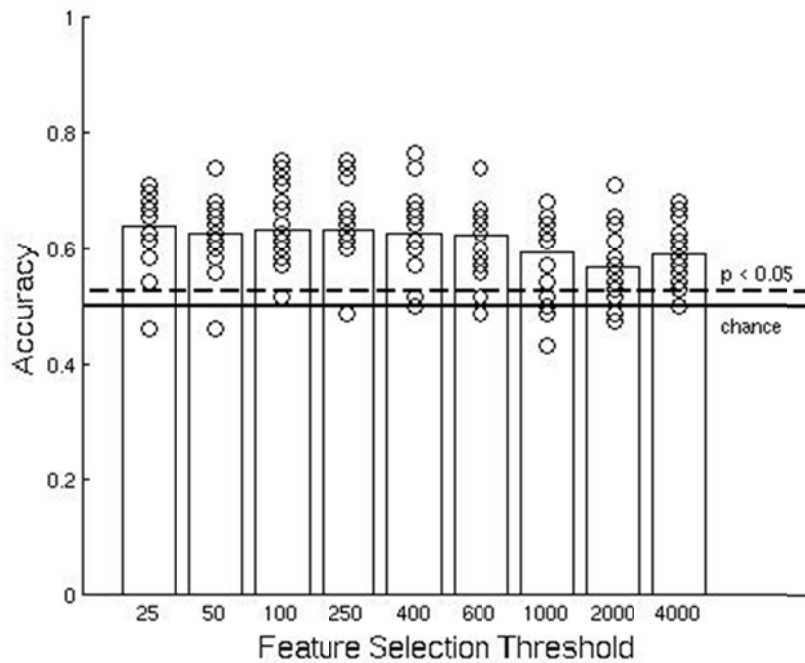


Figure 7.16 Cross-participant accuracies for classifying visual and haptic vs. abstract features from whole brain patterns of activity.

The locations of voxels with the largest classifier weights for identification of visual, haptic, and abstract features with cross-participant classification were distributed throughout the brain (Figure 7.17). The locations of informative voxels were similar across participants. Informative voxel location clusters that were robustly identified across participants (based on 400 voxels) and were critical for decoding visual vs. abstract features included the superior, middle, and inferior temporal gyri, right fusiform gyrus, right superior, medial, and inferior frontal gyri, left premotor cortex, and precuneus. Voxel locations specifically informative for decoding haptic vs. abstract features included superior, middle, and inferior temporal gyri, right fusiform gyrus, middle frontal gyrus, middle occipital gyrus, cuneus, and inferior parietal lobule. Voxel locations specifically informative for decoding visual and haptic vs. abstract features included bilateral inferior and superior temporal gyri, right middle temporal gyrus, right

medial frontal gyrus, bilateral fusiform gyrus, left parahippocampal gyrus, and middle occipital gyrus.




Figure 7.17 Thresholded probability maps (FWE = 0.05, height threshold) of the informative voxels that were consistently identified for each cross-participant classification problem.

7.6.4 Functional connectivity

To investigate the differences in condition-specific connectivity, we first examined whether connectivity between any voxel and the seed region of the occipitotemporal cortex differed between the visual, haptic and abstract conditions. Random-effects group level analyses were performed on the z-maps for each condition to show which connections differed significantly from the null hypothesis of no correlation between regions (Rissman et al., 2004). Regions that were significantly connected to the occipitotemporal area in the three conditions across participants were considerably overlapping (Figure 7.18).

The functional networks for visual, haptic, and abstract networks were highly overlapping but showed differences in connectivity between the visual and haptic networks. In comparison to the visual network, the haptic network showed greater connectivity between the occipitotemporal cortex and the premotor cortex (BA 6).

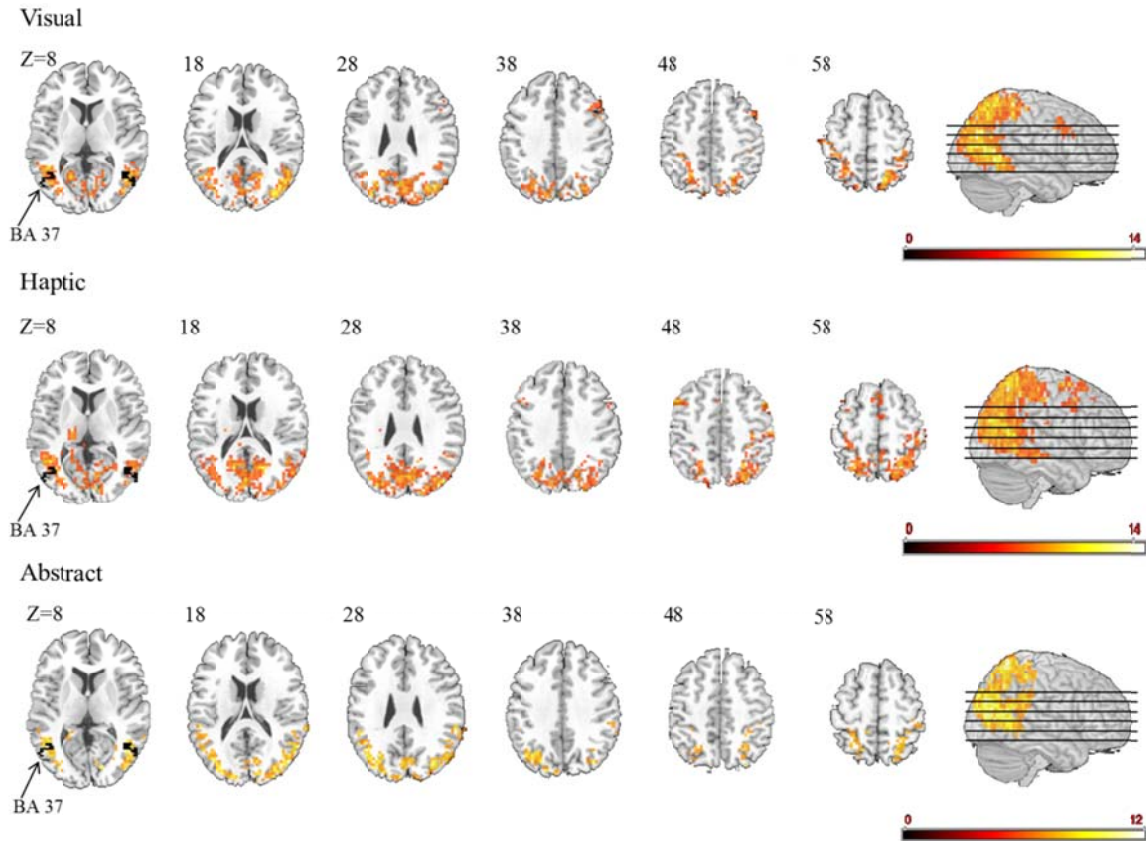


Figure 7.18 Condition-specific connectivity of all voxels with the occipitotemporal cortex (BA 37; in black).

The left primary motor cortex was used as a control region for the functional connectivity analysis to ensure that the network analysis reflected condition-specific differences related to the property-verification task. Since a button-press was required during trials of every condition, no condition-specific differences would be predicted for the connectivity between the left primary motor cortex and all other voxels in the brain. The seed-based condition-specific networks using the primary motor cortex as the seed region did not show any connections that significantly differed from the null hypothesis of no correlation.

To examine the consistency of the condition-specific seed-based functional networks across participants, the z-maps from all but one participant were used to

identify the functional network of the left-out participant. A classifier was trained on the data from all but one participant and tested on the data from the left-out participant.

Classification was significantly above chance for classifying visual vs. haptic functional networks with successful classification for 13 out of 18 participants ($p < 0.05$).

Classification was at chance levels for classifying visual vs. abstract and haptic vs. abstract networks.

7.7 SUMMARY

The first goal of the main experiment was to examine which brain regions participate in processing concepts with visual and haptic features. A univariate analysis indicated that the FG is activated when processing both visual and haptic concepts, while the LOC is activated when processing haptic concepts. These regions are known to be selective for processing the visual and haptic features of objects. Next, the condition-specific functional connectivity of the brain was investigated to characterize how brain regions interact when processing concepts with different types of features. Seed-based networks were constructed to show how brain areas interacted with the occipitotemporal cortex during the visual, haptic, and abstract conditions. The resulting functional networks were highly overlapping but showed differences in connectivity between visual and haptic networks across participants. In comparison to the visual network, the haptic network showed greater connectivity between the premotor cortex and the occipitotemporal cortex. The ability to classify the identity of functional networks across participants demonstrated that connectivity of the visual and haptic networks were quantitatively different as well.

The second goal was to determine whether information about the perceptual content of concepts is present in patterns of brain activity elicited by processing concepts with visual and haptic features. We utilized MVPA to determine whether patterns of brain activity elicited by processing concepts can be used to predict the perceptual information content of a concept. The results of classification demonstrated that information about the visual and haptic features of concepts was present in whole brain patterns of brain activity, regions selective for the visual and haptic features of objects, and regions involved in general visual and haptic perception. The conceptual representation of concepts with visual and haptic features was also consistent across people. Unexpectedly, the neural representation of concepts with visual features could not be distinguished from the neural representation of concepts with haptic features in any areas of the brain. Successful classification occurred only when decoding concepts with perceptual features versus abstract features.

CHAPTER 8

GENERAL DISCUSSION

8.1 SUMMARY & IMPLICATIONS

This work investigated the neural representation of concepts with perceptual features, specifically visual and haptic, to understand how the perceptual aspects of concepts are represented. The purpose was to demonstrate that the representation of concepts with perceptual features is more consistent with weak or strong embodiment theories than unembodied or secondary embodiment theories; however, it was beyond the scope of the current work to provide evidence that rules out amodal conceptual representation. The central hypothesis was that the neural representation of concepts with perceptual features is distributed and includes brain regions in the perceptual systems activated when interacting with the referent of a concept. More specifically, concepts containing visual information should be represented in brain regions active when processing visual stimuli, while concepts containing haptic information should be represented in brain regions active when processing haptic stimuli.

8.1.1 Which brain regions participate in processing concepts with visual and haptic features?

The first goal of this work was to determine which brain regions participate in processing concepts with visual and haptic features. Based on the literature examining visual and haptic object perception, we hypothesized that concepts with visual and haptic features elicit activity in regions known to be active when perceiving the visual and

haptic features of objects, such as the FG, LOC, and IPS, as well as general visual and haptic perceptual regions, such as the primary and secondary visual and somatosensory cortices. A univariate analysis was employed to show which brain regions were on average activated to a greater extent when verifying the properties of concepts with one feature type over another. A significant difference in average regional brain activation in one condition over another suggests a brain region's involvement in a specific cognitive process. The findings of the univariate analysis suggested two key brain regions were involved in processing the visual and haptic features of concepts, the FG and LOC. The FG was implicated in processing both the visual and haptic features of concepts. This area resides along the ventral stream of the visual system, which processes information regarding the identity of objects for the purpose of identifying and extracting meaning from stimuli (Ungerleider & Mishkin, 1982; Goodale & Milner, 1992). Additionally, the FG is likely a region that unifies object-specific information from auditory, visual, and haptic modalities into a trisensory representation (Kassuba et al., 2011) with visual information showing primacy over haptic information (Kassuba et al., 2013). Furthermore, the FG has been demonstrated to be active for processing concrete concepts consistently across studies investigating the differences between abstract and concrete words (Wang et al., 2010). The LOC was implicated in processing the haptic features of concepts. The LOC is located at the convergence of visual and haptic streams of information and is thought to be a bimodal visuo-haptic processing center (Amedi et al, 2005; Deshpande et al., 2010; James et al., 2005; James et al., 2007; Lacey et al., 2010; Lacey et al., 2009; Kassuba et al., 2013). Since the LOC is bimodal, it was expected that

both the visual and haptic features of concepts would activate this region. Univariate contrasts were constructed to compare perceptual features to abstract features, so this suggests that concepts with abstract features may have elicited activation in the LOC as well. It was also expected that processing the visual and haptic features of concepts would elicit activation in the IPS; however, the IPS was not implicated by the univariate analysis. The IPS is responsible for processing information regarding the geometric properties of objects, such as shape and size, which were under-represented by the stimuli used in the main experiment. Geometric properties tend to be bimodal, and stimuli were chosen to be as unimodal as possible. As such, texture and temperature features made up the bulk of the stimuli. As hypothesized, the univariate analysis implied that visual and haptic object-selective regions are important for the representation of concepts with perceptual features. The involvement of object-selective perceptual regions in conceptual representation provides support for weak embodiment theories, which predict regions anterior to primary perceptual systems underlie conceptual representation.

The univariate analysis implicated two key regions in the neural representation of concepts with visual and haptic features. Since brain regions do not act in isolation, an interesting question arises as to which other brain regions communicate with those identified as active during a cognitive task. Seed-based functional connectivity is a novel approach to characterize which brain regions interact during a cognitive task and how this interaction changes across different experimental conditions (Rissmann et al., 2004). This work examined how the brain regions involved in processing concepts with visual and haptic features were functionally connected. The hypothesis was that functional

networks for processing concepts with visual and haptic features contain similar brain regions, but these brain regions are connected differently based on the type of stimulus being processed. The occipitotemporal cortex was used as a seed region, because it was identified by the univariate analysis and contains the LOC. Seed-based functional networks were computed to examine which brain regions interacted with the occipitotemporal cortex during the visual and haptic conditions. The visual and haptic functional networks were highly overlapping but showed some qualitative differences in connectivity. An examination of the differences between the visual and haptic networks showed that the networks for verifying haptic features of concepts elicited stronger connections between the occipitotemporal cortex and the premotor cortex. Previously, the LOC and premotor cortex were demonstrated to be functionally connected during haptic shape and texture perception (Deshpande, Hu, Stilla & Sathian, 2008). In macaques, neurons in the premotor cortex show somatosensory responses characteristic of “mirror neurons,” which respond to both directing motor movements to explore by touch and watching others explore by touch (Rizzolatti, Luppino & Mattelli, 1998). The finding of the current work suggests that the conceptual representation of concepts with haptic features reflects some aspects of the functional connectivity that occurs during haptic perception. It is important to note that this finding is purely qualitative. Without direct interaction tests, the result must be interpreted with caution. To examine whether these networks were quantitatively different, a machine-learning algorithm was employed to classify the identity of connectivity maps across participants. The classifier was able to discriminate between the visual and haptic networks for the majority of participants,

demonstrating quantitative differences between the networks for verifying the visual and haptic features of objects. The results of the functional connectivity analysis are advantageous for characterizing how the interaction between brain regions changes across experimental conditions and provides a complementary approach to univariate analyses. Taken together, we can conclude that object-selective regions are involved in the neural representation of concepts with visual and haptic features, and the connectivity of the occipitotemporal cortex to other brain regions changes based on which concepts are represented. In the case of concepts with visual and haptic features, the neural representation of concepts with haptic features elicits stronger connectivity between the premotor cortex and the occipitotemporal cortex in comparison to the neural representation of concepts with visual features. This may be due to the importance of integrating motor representations for haptic exploration of objects when representing concepts with haptic features.

The findings of the univariate and functional connectivity analyses have important implications for weak embodiment theories, which suggest that conceptual representation is dependent on sensory and motor systems. Weak embodiment theories predict that processing concepts elicits activation in secondary perceptual areas rather than primary perceptual areas. The univariate analysis and functional connectivity show that, indeed, secondary perceptual regions, such as the FG and LOC, are activated by processing concepts with perceptual features. However, the results cannot speak to whether activity in these brain regions is required for the representation and understanding of concepts with perceptual features. Activation in sensory and motor areas might be epiphenomenal,

arising as feedback from semantic processes in language areas. The fMRI BOLD signal is too slow to characterize whether sensory regions receive input from or output to language processing areas. As a result, fMRI studies alone cannot provide complete support for weak, or strong, embodiment theories.

8.1.2 Do patterns of brain activity elicited by processing concepts carry information about their perceptual features?

As noted before, univariate analyses do not have the capacity to investigate the information present in the interaction between voxels. Pattern-based approaches are complementary to univariate approaches, as they can reveal where in the brain information is represented by predicting the identity of stimuli from distributed and regional patterns of brain activity elicited by those stimuli. The second goal of this work was to investigate where information about the visual and haptic features of concepts is represented. The hypothesis was that the perceptual information content of a concept can be predicted from patterns of brain activity within functionally-defined regions of interest, object-selective and general perceptual regions, as well as from distributed patterns of whole brain activity.

Using MVPA this work demonstrated patterns of brain activity located within regions functionally-defined as important for processing the visual and haptic features of objects as well as for regions which process general visual and haptic perception carry information about the perceptual features of concepts. Object-selective regions included the secondary somatosensory cortex, secondary visual cortex, and the LOC. General perceptual regions included the primary visual and somatosensory cortices. For all of

these regions, information about visual concepts and combined visual and haptic concepts could be discriminated from information about abstract concepts. Information about haptic concepts could not be discriminated from abstract concepts within these regions, which suggests visual information drove successful classification accuracies. The classifier tended to make errors by classifying concepts with haptic features as concepts with visual features, which may explain why classification accuracies for combined visual and haptic features were overall higher than visual features alone. Unexpectedly, information about concepts with haptic features was not present in regions functionally-defined for haptic perception. This could be explained by the bimodal nature of haptic features, as the conceptual representation of concepts with haptic features may have been dominated by visual information. This notion is supported by the fact that the conceptual representation of concepts with visual features was present in these haptic regions. The regions in which classification of perceptual features from abstract features was successful replicated the previous findings of Wang et al. (2012), which decoded concrete and abstract words using different stimuli and a different experimental paradigm.

Whole brain patterns of brain activity also carried information about the perceptual features of concepts. The conceptual representation of concepts with visual and haptic features was largely distributed throughout the cortex and was consistent across people. Consistencies in the locations of voxels identified as most informative for cross-participant classification provide some clues about the nature of conceptual representation. When classifying concepts with visual or haptic features alone from abstract concepts, voxels located in perceptual regions were consistently selected as most

informative across participants. However, when classifying concepts with visual and haptic features combined from abstract concepts, a large cluster of voxels in the temporal poles, in addition to perceptual regions, was consistently selected as most informative across people. The temporal poles have been suggested to be an amodal conceptual hub (Kiefer & Pulvermüller, 2012). This finding indicates that amodal linguistic representation may be important for discriminating between concepts with combined visual and haptic features and concepts with abstract features in whole brain patterns activity, whereas visual information was selected as most informative when classifying within perceptual regions. A limitation of the analysis of the consistency of informative voxels across participants is that within each participant, the most informative voxels are selected somewhat randomly due to the nature of logistic regression and is, therefore, not designed for speculating about the locations of selected voxels. The speculation that consistency of the most informative voxels selected across people reflects amodal conceptual representation must be made with extreme caution.

Unexpectedly, concepts with visual features could not be discriminated from concepts with haptic features in object-selective or general perceptual areas or whole brain patterns of activity. Bimodality can explain why visual and haptic representations are not differentiable. In normal individuals, haptic information is rarely experienced in the absence of visual information. As such, when haptic information is presented without visual information, individuals tend to imagine the corresponding visual information. This is supported by the results of the functional localizer utilized in this work, as visual areas were activated when perceiving haptic stimuli. Additionally, the modality ratings

of all stimuli in the database from which the stimuli were drawn for this work support that haptic stimuli are more bimodal than visual stimuli.

A difficulty of MVPA lies in the ability to link the structure of class information decoded back to the experimental design. In other words, are we really decoding the differences in the perceptual features of concepts? In this work the perceptual features of concepts were manipulated to determine whether information about the perceptual features of concepts can be decoded from patterns of brain activity elicited by making property-verifications. One way to establish a causal relationship is to correlate classifier performance with behavioral performance on a related measure. Within object-selective regions, classifier performance was significantly correlated with participants' ability to visually imagine a situation. This suggests that perceptual information was indeed captured by the classifier. Classification performance was not significantly correlated with mental imagery ability for general perceptual regions or whole brain; however, participants' classification accuracies were generally consistent across all classification problems. It cannot be ruled out that the classifier was capturing information regarding the lower-level features of stimuli. Although stimuli were balanced on word length and frequency, both measures were calculated for each triple of words rather than for single words. Additionally, word frequency was balanced across conditions as well as possible ($p = 0.145$), but haptic stimuli showed a trend for being less frequent. Additionally, it has been demonstrated that abstract concepts tend to be more emotionally-valenced than concrete concepts (Kousta et al., 2011; Vigliocco et al., 2013). The stimuli used in this work showed differences in mean emotional valence ratings, with visual stimuli showing

greater emotional valence ratings than abstract and haptic stimuli (Figure 8.1).

Classification accuracies may have been influenced by the differences in emotional valence, but it does not fully explain the classification results. We were able to classify visual vs. abstract but not visual vs. haptic trials. If the classifier was capturing only information about emotional valence, we would have been able to classify visual vs. haptic trials as well. However, emotional valence, in addition to perceptual content, may have contributed to successful classification of visual vs. abstract. Further research should investigate the effect of valence for the representation of concrete and abstract concepts.

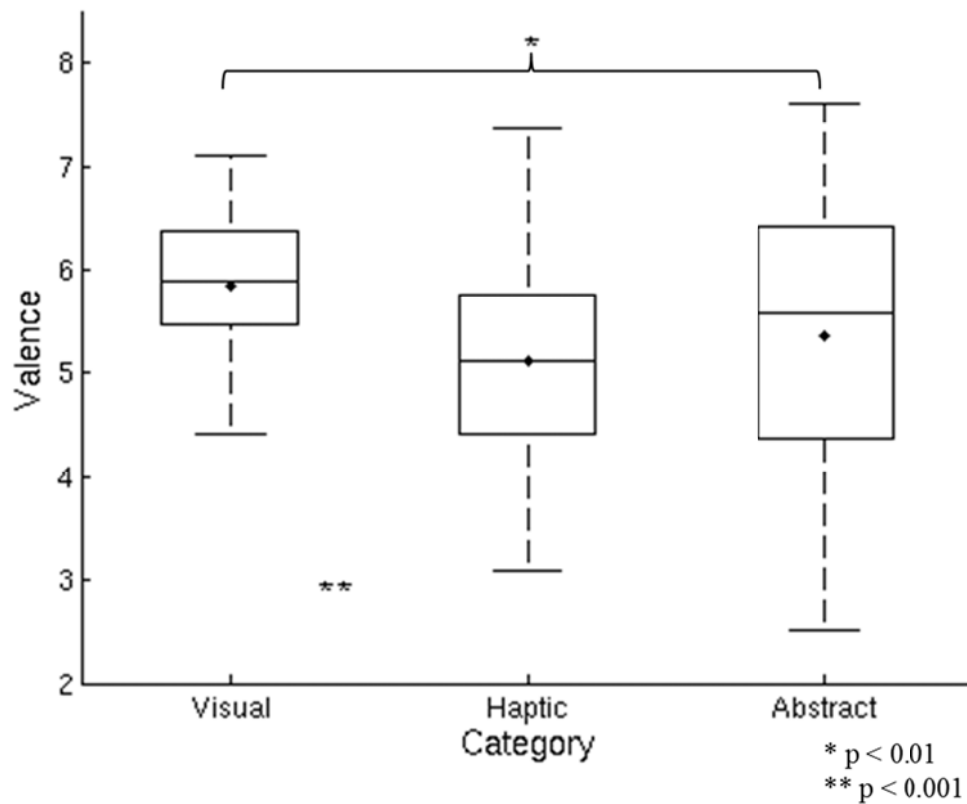


Figure 8.1 Valence ratings for stimuli from each experimental category.

Overall, the implications of the MVPA results support aspects of both weak and strong embodiment theories while also providing evidence of amodal conceptual representation. Weak embodiment theories are supported by the ability to classify concepts with perceptual features from regions involved in processing the perceptual features of objects. Strong embodiment theories are supported by the ability to classify concepts with perceptual features from regions involved in general perceptual processing, as this suggests a full simulation is elicited when processing concepts. A full simulation may be due to task-specific demands, which encourages participants to engage in mental imagery to complete a task. Meteyard et al. (2012) proposes that the depth of processing must be taken into account to determine whether task demands induce mental imagery. Deeper processing (i.e. narrative comprehension) would elicit greater mental imagery than superficial processing (i.e. lexical decision). This work utilized a property-verification task, which required participants to decide whether a concept has one of two properties. It has been demonstrated that reaction times for completing this task is influenced by factors that also influence perceptual processing (Dantzig et al., 2011). This suggests that participants may be engaging in a simulation of concepts. Whether or not task demands elicit a full simulation of concepts is unclear and poses a limitation for this work.

Several conclusions can be drawn from this work, which provide insight into the nature of the neural representation of concepts with perceptual features. The neural representation of concepts with visual and haptic features involves brain regions which underlie general visual and haptic perception as well visual and haptic perception of

objects. These brain regions interact differently based on the type of perceptual feature a concept possesses. Additionally, the neural representation of concepts with visual and haptic features is distributed across the whole brain and is consistent across people. The results of this work support aspects of weak/strong embodiment theories; however, the dependency of conceptual representation on these regions is beyond the scope of this work.

8.2 FUTURE DIRECTIONS

A limitation of this work was the inability to show full support for weak and strong embodiment theories. Strong embodiment theories cannot be fully supported by this work, because modulation of sensory representation must be shown in two directions. This work demonstrates that processing concepts with perceptual features modulates sensory representation by eliciting activation in primary sensory regions, but full support of strong embodiment requires showing that influencing sensory representation also modulates conceptual representation. All studies showing full support for strong embodiment theories have utilized action words to show that influencing the motor system modulates action word processing and vice versa (Buccino et al., 2005; Pulvermüller et al., 2005). Future studies will need to replicate this finding in sensory systems to demonstrate that conceptual representation is grounded in both sensory and motor systems. Both weak and strong embodiment theories propose that sensory and motor representations are required for conceptual representation; however, demonstrating this dependency was beyond the scope of this work. Due to the nature of the fMRI BOLD signal, fMRI evidence is not sufficient to make this determination. Neuroimaging

methods with higher temporal resolution, such as EEG, may be able to decouple feedforward and feedback effects to show whether sensory activation drives conceptual representation or is the output of semantic processing in language areas. Additionally, TMS and lesion studies may provide evidence that sensory areas are required by showing deficits in semantic processing of concepts with visual and haptic features when sensory areas are lesioned or temporarily inhibited. Previous studies have shown that lesions to visual and auditory association areas produce deficits in processing words with visual and auditory features (Neininger & Pulverüller, 2006; Trumpp et al., 2013); however no such study has investigated semantic processing of concepts with haptic features in patients with lesions to somatosensory areas.

Finally, it has been suggested that emotional valence may play an important role in the neural representation of concepts with perceptual and abstract features (Kousta et al., 2011; Vigliocco et al., 2013). Future work should aim to investigate how emotional valence contributes to the neural representation of concrete and abstract concepts, both across concepts and within sub-categories of concepts (i.e. visual, haptic, cognition or emotion).

This work was novel, because previously MVPA has not been used to investigate the neural representation of concepts with perceptual features. Future work should aim to replicate the current findings with stimuli of other modalities, such as auditory, olfactory, and gustatory, using MVPA. Cross-modality MVPA, discriminating between different types of visual features within haptic areas and vice versa, would also be an interesting approach to further characterize the nature of modal representations.

Finally, it would be interesting to conduct this line of research on participants with deficits in perception to see whether sensory information is necessary for conceptual representation. For example, one study investigating conceptual representation with sighted and congenitally blind participants showed that color knowledge contributes to similarity judgments for fruits and vegetables but not household objects in sighted participants (Connolly, Gleitman & Thompson-Schill, 2007). Future work with special populations could elucidate which information is absolutely necessary for representing different types of concepts.

8.3 MERIT & CONTRIBUTION

The current work was innovative, because no studies have examined how the brain represents concepts with visual and haptic features using MVPA. The research strategy of this work employed state of the art quantitative methods to explore the information content and functional connectivity of patterns of brain activity elicited by concepts with visual and haptic features for the first time. This strategy is more sensitive in comparison to the traditional univariate approach proposed in the first aim as it jointly investigates information in multiple voxels. The outcome of this work served to further our understanding of how the brain represents concepts and provides support for weak embodiment theories.

REFERENCES

- Abdi, H., Dunlop, J. P., & Williams, L. J. (2009). How to compute reliability estimates and display confidence and tolerance intervals for pattern classifiers using the Bootstrap and 3-way multidimensional scaling (DISTATIS). *NeuroImage*, 45, 89-95.
- Arevalo, A. L., Baldo, J. V., & Dronkers, N. F. (2012). What do brain lesions tell us about theories of embodied semantics and the human mirror neuron system? *Cortex*, 48(2), 242-254.
- Aziz-Zadeh, L., & Damasio, A. (2008). Embodied semantics for actions: findings from functional brain imaging. *Journal of Physiology - Paris*, 102, 35-39.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577-660.
- Barsalou, L. W. (2003). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society of London: Biological Sciences*, 358, 1177-1187.
- Barsalou, L. W., Hale, C. R., Van Mechelen, I., Hampton, J., Michalski, R. S., & Theuns, P. (1993). Components of conceptual representation: From feature lists to recursive frames *Categories and concepts: Theoretical views and inductive data analysis*. (pp. 97-144). San Diego, CA, US: Academic Press.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. New York: Springer.

- Boronat, C. B., Buxbaum, L. J., Coslett, H. B., Tang, K., Saffran, E. M., Kimberg, D. Y., & Detre, J. A. (2005). Distinctions between manipulation and function knowledge of objects: evidence from functional magnetic resonance imaging. *Brain Res Cogn Brain Res*, 23(2-3), 361-373.
- Buccino, G., Riggio, L., Melli, G., Binkofski, F., Gallese, V., & Rizzolatti, G. (2005). Listening to action-related sentences modulates the activity of the motor system: A combined TMS and behavioral study. *Cognitive Brain Research*, 24, 355-363.
- Calvo-Merino, B., Glaser, D. E., Grezes, J., Passingham, R. E., & Haggard, P. (2013). Action observation and acquired motor skills: An fMRI study with expert dancers. *Cerebral Cortex*, 15(1243-1249).
- Chatterjee, A. (2010). Disembodying cognition. *Language and Cognition*, 2(1), 79-116.
- Chen, Y., Namburi, P., Elliott, L. T., Heinzle, J., Soon, C. S., Chee, M. W. I., & Haynes, J. (2010). Cortical surface-based searchlight decoding. *Neuroimage*, 56(2), 582-592.
- Connell, L., & Lynott, D. (2010). Look but don't touch: Tactile disadvantage in processing modality-specific words. *Cognition*, 115, 1-9.
- Connolly, A. C., Gleitman, L. R., & Thompson-Schill, S. L. (2007). Effect of congenital blindness on the semantic representation of some everyday concepts. *PNAS*, 104(20), 8241-8246.
- Crick, F. C., & Koch, C. (2005). What is the function of the claustrum? *Philosophical Transactions of the Royal Society - Biological Sciences*, 360(127), 1-9.

- Cui, X., Jeter, C. B., Yang, D., Montague, P. R., & Eagleman, D. M. (2007). Vividness of mental imagery: Individual variability can be measured objectively. *Vision Research*, 47, 474-478.
- Dantzig, S., Cowell, R. A., Zeelenberg, R., & Pecher, D. (2011). A sharp image or a sharp knife: Norms for the modality-exclusivity of 774 concept-property items. *Behavioral Research*, 43, 145-154.
- Desai, R. H., Binder, J. R., Conant, L. L., & Seidenberg, M. S. (2009). Activation of sensory-motor areas in sentence comprehension. *Cerebral Cortex*, 20(2), 468-478. doi: 10.1093/cercor/bhp115
- Deshpande, G., Hu, X., Lacey, S., Stilla, R., & Sathian, K. (2010). Object familiarity modulates effective connectivity during haptic shape perception. *Neuroimage*, 49, 1991-2000.
- Fink, G. R., Frackowiak, R. S. J., Pietrzyk, U., & Passingham, R. E. (1997). Multiple nonprimary motor areas in the human cortex. *Journal of Neurophysiology*, 77.
- Gerlach, C. (2007). A review of functional imaging studies on category specificity. *Journal of Cognitive Neuroscience*, 19(2), 296-314.
- Goldberg, R. F., Perfetti, C. A., & Schneider, W. (2006). Perceptual knowledge retrieval activates sensory brain regions. *The Journal of Neuroscience*, 26(18), 4917-4921.
- Gonzalez, J., Barros-Loscertales, A., Pulvermuller, F., Meseguer, V., Sanjuan, A., Belloch, V., & Avila, C. (2006). Reading cinnamon activates olfactory brain regions. *Neuroimage*, 32(2), 906-912.

- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neuroscience*, 15, 20-25.
- Grossman, M., Anderson, C., Khan, A., Avants, B., Elman, L., & McCluskey, L. (2008). Impaired action knowledge in amyotrophic lateral sclerosis. *Neurology*, 71(1396-1401).
- Grossman, M., Koenig, P., DeVita, C., Glosser, G., Alsop, D., & Detre, J. (2002). The neural basis for category-specific knowledge: An fMRI study. *Neuroimage*, 15, 936-948.
- Hadjikhani, N., & Roland, P. E. (1998). Cross-modal transfer of information between the tactile and the visual representations in the human brain: A positron emission tomographic study. *Journal of Neuroscience*, 18(3), 1072-1084.
- Hauk, O., Davis, M. H., Kherif, F., & Pulvermüller, F. (2008). Imagery or meaning? Evidence for a semantic origin of category-specific brain activity in metabolic imaging. *European Journal of Neuroscience*, 27, 1856-1866.
- Helbig, H. B., Ernst, M. O., Ricciardi, E., Pietrini, P., Thielscher, A., Mayer, K. M., Noppeney, U. (2012). The neural mechanisms of reliability weighted integration of shape information from vision and touch. *Neuroimage*, 60, 1063-1072.
- Hoening, K., Sim, E., Bochev, V., Herrnberger, B., & Kiefer, M. (2008). Conceptual flexibility in the human brain: dynamic recruitment of semantic maps from visual, motor, and motion-related areas. *Journal of Cognitive Neuroscience*, 20(10), 1799-1814.

- Humphreys, G. W., & Forde, E. M. E. (2001). Hierarchies, similarity, and interactivity in object recognition: "Category-specific" neuropsychological deficits. *Behavioral and Brain Sciences*, 24, 453-509.
- Humphreys, G. W., Riddoch, M. J., & Quinlan, P. T. (1988). Cascade processes in picture identification. *Cognitive Neuropsychology*, 5, 67-103.
- James, T. W., James, K. H., Humphrey, G. K., & Goodale, M. A. (2005). Do visual and tactile object representations share the same neural substrate? In M. A. Heller & S. Ballesteros (Eds.), *Touch and blindness: psychology and neuroscience*. Mahwah, NJ: Lawrence Erlbaum.
- James, T. W., & Kim, S. (2010). Dorsal and ventral cortical pathways for visuo-haptic shape integration revealed using fMRI. In M. J. Naumer & J. Kaiser (Eds.), *Multisensory object perception in the primate brain*. New York: Springer.
- James, T. W., Kim, S., & Fisher, J. S. (2007). The neural basis of haptic object processing. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 61(3), 219-229.
- Just, M. A. (2008). What brain imaging can tell us about embodied meaning. In M. d. Vega, A. M. Glenberg & A. C. Graesser (Eds.), *Symbols, embodiment, and meaning: debates on meaning and cognition*. Oxford, UK: Oxford University press.
- Just, M. A., Newman, S. D., Keller, T. A., McEleney, A., & Carpenter, P. A. (2004). Imagery in sentence comprehension: An fMRI study. *Neuroimage*, 21, 112-124.

- Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nat Neurosci*, 8(5), 679-685.
- Kassuba, T., Klinge, C., Holig, C., Menz, M. M., Ptito, M., Roder, B., & Siebner, H. R. (2011). The left fusiform gyrus hosts trisensory representations of manipulable objects. *Neuroimage*, 56, 1566-1577.
- Kassuba, T., Klinge, C., Holig, C., Roder, B., & Siebner, H. R. (2013). Vision holds a greater share in visuo-haptic object recognition than touch. *Neuroimage*, 65, 59-68.
- Kiefer, M., & Pulvermüller, F. (2012). Conceptual representations in mind and brain: Theoretical developments, current evidence, and future directions. *Cortex*, 48, 805-825.
- Kiefer, M., Sim, E., Herrnberger, B., Grothe, J., & Hoenig, K. (2008). The sound of concepts: Four markers for a link between auditory and conceptual brain systems. *The Journal of Neuroscience*, 28(47), 12224-12230.
- Kim, S., & James, T. W. (2010). Enhanced effectiveness in visuo-haptic object-selective brain regions with increasing stimulus salience. *Human Brain Mapping*, 31, 678-693.
- Kober, H., Barrett, L. F., Joseph, J., Bliss-Moreau, E., Lindquist, K., & Wager, T. D. (2008). Functional grouping and cortical-subcortical interactions in emotion: A meta-analysis of neuroimaging studies. *Neuroimage*, 42, 998-1031.

- Kousta, S.-T., Vigliocco, G., Vinson, D. P., Andrews, M., & Del Campo, E. (in press).
The representation of abstract words: Why emotion matters. *Journal of Experimental Psychology, General*.
- Kriegeskorte, N. (2011). Pattern-information analysis: From stimulus decoding to computational-model testing. *NeuroImage*, 56(2), 411-421. doi: 10.1016/j.neuroimage.2011.01.061
- Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proc Natl Acad Sci U S A*, 103(10), 3863-3868.
- Lacey, S., Flueckiger, P., Stilla, R., Lava, M., & Sathian, K. (2010). Object familiarity modulates the relationship between visual object imagery and haptic shape perception. *Neuroimage*, 49, 1977-1990.
- Lacey, S., Tal, N., Amedi, A., & Sathian, K. (2009). A putative model of multisensory object representation. *Brain Topography*, 21, 269-274.
- Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago, IL: University of Chicago Press.
- Lancaster, J. L., Woldorff, M. G., Parsons, L. M., Liotti, M., Freitas, C. S., Rainey, L., . . . Fox, P. T. (2000). Automated Talairach atlas labels for functional brain mapping. *Human Brain Mapping*, 10, 120-131.
- Lemus, I., Hernández, A., Luna, R., Zainos, A., & Romo, R. (2010). Do sensory cortices process more than one sensory modality during perceptual judgments? *Neuron*, 67, 335-348.

- Libedinsky, C., & Livingstone, M. (2011). Role of prefrontal cortex in conscious visual perception. *The Journal of Neuroscience*, 31(1), 64-69.
- Mahon, B. Z., & Caramazza, A. (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology - Paris*, 102(1-3), 59-70.
- Maldjian, J. A., Laurienti, P. J., Kraft, R. A., & Burdette, J. H. (2003). An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fmri data sets. *Neuroimage*, 19, 1233-1239.
- Markman, A. B., & Dietrich, E. (2000). Extending the classical view of representation. *Trends Cogn Sci*, 4(12), 470-475.
- Marks, D. F. (1973). Visual imagery differences in the recall of pictures. *British Journal of Psychology*, 1, 17-24.
- Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C. L., & Ungerleider, L. G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. *Science*, 270, 102-105.
- Meteyard, L., Cuadrado, S. R., Bahrami, B., & Vigliocco, G. (2012). Coming of age: A review of embodiment and the neuroscience of semantics. *Cortex*, 48, 788-804.
- Meyer, K., Kaplan, J. T., Essex, R., Damasio, H., & Damasio, A. (2011). Seeing touch is correlated with content-specific activity in primary somatosensory cortex. *Cerebral Cortex*, 21(9), 2113-2121.

- Mitchell, T. M., Hutchinson, R., Niculescu, R. S., Pereira, F., Wang, X., Just, M. A., & Newman, S. (2004). Learning to decode cognitive states from brain images. *Machine Learning*, 57, 145-175.
- Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K., Malave, V. L., Mason, R. A., & Just, M. A. (2008). Predicting human brain activity associated with the meanings of nouns. *Science*, 320, 1191-1195.
- Mur, M., Bandettini, P. A., & Kriegeskorte, N. (2009). Revealing representational content with pattern-information fMRI - an introductory guide. *SCAN*, 4, 101-109.
- Neininger, B., & Pulvermüller, F. (2003). Word-category specific deficits after lesions in the right hemisphere. *Neuropsychologia*, 41, 53-70.
- Newman, S. D., Klatzky, R. L., Lederman, S. J., & Just, M. A. (2005). Imagining material versus geometric properties of objects: an fMRI study. *Cognitive Brain Research*, 23(2), 235-246.
- Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, 10, 424-430.
- O'Toole, A. J., Jiang, F., Abdi, H., Penard, N., Dunlop, J. P., & Parent, M. A. (2007). Theoretical, statistical, and practical perspectives on pattern-based classification approaches to the analysis of functional neuroimaging data. *Journal of Cognitive Neuroscience*, 19(11), 1735-1752.

- Pecher, D., Zeelenberg, R., & Barsalou, L. W. (2003). Verifying different-modality properties for concepts produces switching costs. *Psychological Science*, 14(2), 119-124.
- Pereira, F., Mitchell, T. M., & Botvinick, M. (2009). Machine learning classifiers and fMRI: a tutorial overview. *Neuroimage*, 45, S199-S209.
- Pexman, P. M., Hargreaves, I. S., Edwards, J. D., Henry, L. C., & Goodyear, B. G. (2007). Neural correlates of concreteness in semantic categorization. *Journal of Cognitive Neuroscience*, 19, 1407-1419.
- Pietrini, P., Furey, M. L., Ricciardi, E., Gobbini, M. I., Wu, W. H. C., Cohen, L., . . . Haxby, J. (2004). Beyond sensory images: Object-based representation in the human ventral pathway. *Proceedings of National Academy of Science USA*, 101, 5658-5663.
- Pobric, G., Lambon-Ralph, M., & Jeffries, E. (2009). The role of the anterior temporal lobes in the comprehension of concrete and abstract words: rTMS evidence. *Cortex*, 45, 1104-1110.
- Pulvermüller, F. (2001). Brain reflections of words and their meaning. *Trends in Cognitive Sciences*, 5(12), 517-524.
- Pulvermüller, F. (2005). Brain mechanisms linking language and action. *Nature Reviews: Neuroscience*, 6, 1-7.
- Pulvermüller, F. (in press). Semantic embodiment, disembodiment, or misembodiment? In search of meaning in modules and neuron circuits. *Brain & Language*.

- Pulvermüller, F., & Hauk, O. (2006). Category-specific conceptual processing of color and form in left fronto-temporal cortex. *Cerebral Cortex*, 16(8), 1193-1201. doi: 10.1093/cercor/bhj060
- Pulvermüller, F., Hauk, O., Nikulin, V. V., & Ilmoniemi, R. J. (2005). Functional links between motor and language systems. *European Journal of Neuroscience*, 21(3), 793-797.
- Raizada, R. D. S., & Kriegeskorte, N. (2010). Pattern-information fMRI: New questions which it opens up and challenges which face it. *International Journal of Imaging Systems and Technology*, 20(1), 31-41.
- Raizada, R. D. S., Tsao, F., Liu, H., & Kuhl, P. (2010). Quantifying the adequacy of neural representations for cross-language phonetic discrimination task: prediction of individual differences. *Cerebral cortex*, 20, 1-12.
- Remedios, R., Logothetis, N. K., & Kayser, C. (2010). Unimodal responses prevail in the multisensory claustrum. *Journal of Neuroscience*, 30(39), 12902-12907.
- Rissman, J., Gazzaley, A., & D'Esposito, M. (2004). Measuring functional connectivity during distinct stages of a cognitive task. *Neuroimage*, 23, 752-763.
- Robert, P., & Escoufier, Y. (1976). A unifying tool for linear multivariate statistical methods: the RV-coefficient. *Applied Statistics*, 25, 257-265.
- Searle, J. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417-457.
- Servos, P., Lederman, S., Wilson, D., & Gati, J. (2001). fMRI-derived cortical maps for haptic shape, texture, and hardness. *Brain Res Cogn Brain Res*, 12(2), 307-313.


- Shinkareva, S. V., Malave, V. L., Mason, R. A., Mitchell, T. M., & Just, M. A. (2011). Commonality of neural representations of words and pictures. *Neuroimage*, 54, 2418-2425. doi: 10.1016/j.neuroimage.2010.10.042
- Shinkareva, S. V., Mason, R. A., Malave, V. L., Wang, W., Mitchell, T. M., & Just, M. A. (2008). Using fMRI brain activation to identify cognitive states associated with perception of tools and dwellings. *PLoS ONE*, 3, e1394.
- Shinkareva, S. V., Ombao, H. C., Sutton, B. P., Mohanty, A., & Miller, G. A. (2006). Classification of functional brain images with a spatio-temporal dissimilarity map. *Neuroimage*, 33(1), 63-71.
- Simmons, W. K., Hamann, S. B., Harenski, C. L., Hu, X., & Barsalou, L. W. (2008). fMRI evidence for word association and situated simulation in conceptual processing. *Journal of Physiology - Paris*, 102, 106-119.
- Simmons, W. K., Martin, A., & Barsalou, L. W. (2005). Pictures of appetizing foods activate gustatory cortices for taste and reward. *Cerebral Cortex*, 15, 1602-1608.
- Solomon, K. O., Medin, D. L., & Lynch, E. (1999). Concepts do more than categorize. *Trends in Cognitive Science*, 3, 99-105.
- Tan, L. H., Chan, A. H. D., Kay, P., Khong, P.-L., Yip, L. K. C., & Luke, K.-K. (2008). Language affects patterns of brain activation associated with perceptual decision. *Proceedings of the National Academy of Sciences*, 105(10), 4004-4009. doi: 10.1073/pnas.0800055105

- Tettamanti, M., Buccino, G., Saccuman, M. C., Gallese, V., Danna, M., Scifo, P., . . . Perani, D. (2005). Listening to action-related sentences activates fronto-parietal motor circuits. *Journal of Cognitive Neuroscience*, 17, 273-281.
- Tong, F., & Pratte, M. S. (2012). Decoding patterns of human brain activity. *Annual Review of Psychology*, 63, 483-509.
- Trumpp, N. M., Kliese, D., Hoenig, K., Haarmeier, T., & Kiefer, M. (2013). Losing the sound of concepts: Damage to auditory association cortex impairs the processing of sound-related concepts. *Cortex*, 49, 474-486.
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual streams. In D. J. Ingle, M. A. Goodale & R. J. Mansfield (Eds.), *The analysis of visual behavior*. (pp. 549-586). Cambridge, MA: MIT Press.
- Vigliocco, G., Meteyard, L., Andrews, M., & Kousta, S.-T. (2009). Towards a theory of semantic representation. *Language and Cognition*, 1(2), 219-247.
- Wang, J., Baucom, L. B., & Shinkareva, S. V. (2012). Decoding abstract and concrete concept representations based on single-trial fMRI data. *Human Brain Mapping*, 34(5).
- Wang, J., Conder, J. A., Blitzer, D. N., & Shinkareva, S. V. (2010). Neural representation of abstract and concrete concepts: A meta-analysis of neuroimaging studies. *Human Brain Mapping*, 31, 1459-1468.
- Wang, J., Conder, J. A., Blitzer, D. N., & Shinkareva, S. V. (2010). Neural representation of abstract and concrete concepts: a meta-analysis of neuroimaging studies. *Human Brain Mapping*, 31, 1459-1468.

- Warrington, E. K., & Shallice, T. (1984). Category Specific Semantic Impairments. *Brain*, 107, 829-854.
- Weber, M., Thompson-Schill, S. L., Osherson, D., Haxby, J., & Parsons, L. (2009). Predicted judged similarity of natural categories from their neural representations. *Neuropsychologia*, 47, 859-868.
- Whitaker, T. A., Simões-Franklin, C., & Newell, F. N. (2008). Vision and touch: Independent or integrated systems for the perception of texture? *Brain Research*, 1242, 59-72.
- Wilson-Mendenhall, C. D., Barrett, L. F., Simmons, W. K., & Barsalou, L. W. (2011). Grounding emotion in situated conceptualization. *Neuropsychologia*, 49, 1105-1127.

APPENDIX A – FUNCTIONAL LOCALIZER STIMULI

Objects



Objects cont.



Textures



Textures cont.

