

**From the holdings of the Department of Special Collections , Princeton University Library,
Princeton, NJ**

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyright law. We do not charge any permission or usage fees for the publication of images of material in our collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff. For more information regarding copyright, credit, and citation please review the information on the Copyright, Credit and Citation Guidelines page of our website.

<https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines>. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. <https://library.princeton.edu/special-collections/ask-us>

Doctoral Dissertations, 2011-2026

Psychology

2013

THE COGNITIVE AND NEURAL BASIS OF LANGUAGE LEARNING: INVESTIGATIONS IN TYPICAL AND AUTISTIC POPULATIONS

Johnson, Matt A.

Princeton, NJ : Princeton University

<http://arks.princeton.edu/ark:/88435/dsp01gt54kn09s>

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. <https://www.copyright.gov/title17/title17.pdf>

THE COGNITIVE AND NEURAL BASIS OF LANGUAGE LEARNING:
INVESTIGATIONS IN TYPICAL AND AUTISTIC POPULATIONS

Matt A. Johnson

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF PSYCHOLOGY
ADVISERS: ADELE E. GOLDBERG & NICHOLAS B. TURK-BROWNE

JUNE, 2013

© Copyright by Matt A. Johnson, 2013. All rights reserved

Outline

Abstract	4 - 5
Chapter 1: What are phrasal constructions?	6 - 8
Chapter 2: Evidence for abstract argument structure knowledge	9 - 29
Experiment 1: Online evidence for constructional meaning	15 - 29
Chapter 3: Introduction to construction learning	30 - 35
Chapter 4: Construction learning and the Brain	36 - 82
Experiment 2: Neural systems supporting the learning of novel constructions	36 - 58
Experiment 3: Construction learning and prediction	59 - 83
Chapter 5: Construction learning in children with autism	84 - 103
Experiment 4: Construction learning in children with autism	84 - 103
Final Words: A general picture of language impairment in children with autism	103 - 110
Summary of Findings	111
References	112 - 120

Abstract

This dissertation will focus on the processing and learning of abstract, phrasal argument structure constructions. Chapter 1 provides the theoretical framework for abstract constructions, and illustrates the importance of such representations in speakers' linguistic knowledge. Chapter 2 reviews the evidence for meaning being associated with abstract phrasal forms, including experiment 1, which supports the idea that constructional meaning is accessed online and automatically.

The dissertation then focuses on how these stable representations are initially learned. This will focus on two avenues of research, which explore this issue in two different ways. First, chapter 3 will discuss the neural mechanisms that contribute to construction learning. This chapter will present evidence from two neuroimaging studies (experiments 2 and 3) – the first scientific investigations that probe the neuroscientific basis of abstract construction learning. These findings will be discussed in detail; they suggest that construction learning may rely heavily on a largely domain general learning system, including a prediction based mechanism. We also find that such abstract learning may alleviate low level visual processing in previously unexplored ways.

Next, chapter 4 reports research on construction learning in children with autism. This will include behavioral work (experiment 4) suggesting that children have difficulty learning abstract constructions, and discuss the implications of this finding. The chapter closes with final thoughts on our current understanding of language delay in autism, which includes both social and cognitive factors.

Overall this dissertation report provides several novel insights into the learning and processing of abstract phrasal correspondences, and will argue that they are 1) a crucial element of our linguistic knowledge, allowing language to be used flexibly 2) are accessed quickly and implicitly during online sentence comprehension 3) are likely learned at least in part on the basis of a prediction-based, domain general learning mechanism, by generalizing over discrete pairings of form and function, and 4) require a type of learning which the (locally biased) cognitive style of autism is ill-suited for.

Chapter 1

What are phrasal constructions? The place and importance of abstract linguistic knowledge

What does one's knowledge of a language actually consist of? As a starting point, we know that the linguistic knowledge of fluent speakers entails a memory and understanding of a vast amount of concrete, item specific information. Whether for specific adages, idioms, or the general meaning of specific words (estimated at ~60,000 for adult speakers) our concrete linguistic knowledge is impressive. In fact, we have some enduring (at least 1 week) memory for specific, verbatim utterances, even when such items are later retrieved implicitly (Gurevich, Johnson, and Goldberg 2010).

At the same time, it is also recognized that abstract linguistic representations must also exist. If we only memorized what we've already heard, we would only be capable of understanding and producing these particular units. But language is open-ended in that never before utterances can seamlessly be understood and produced. Abstract representations are what allow such flexibility to exist.

The constructionist approach (e.g. Fillmore, Kay and O'Connor 1988; Goldberg 1995; Langacker 1987) adopts the view that a network of both specific knowledge and generalizations form the foundation of our knowledge of language. That is, within a constructionist system, our linguistic knowledge is argued to consist of a rich and expansive network of form – meaning correspondences (e.g. 'constructions'), at varying levels of abstraction. For example, individual words can be understood as associations between a linguistic form (e.g. *dog*) and meaning (e.g. a domestic mammal that barks, has large ears, is loyal etc.). Constructions can also be *partially* abstract, as in the case of idioms with open slots. For example, the *pulling <one's> leg* idiom,

can take many different variations which convey the same idiomatic meaning, such as *he pulled my leg*, *we were just pulling their legs* etc. Its meaning is not tied with the individual words per se, but with the following abstract form: “pull [NP₂’s leg]”.

More abstract still are argument structure constructions (Goldberg 1995), which will be the focus of this dissertation. Argument structure constructions are skeletal sentence structures that constrain a sentence’s overall meaning. For example, the ditransitive construction, <Subj Verb Obj₁ Obj₂> conveys a meaning of “give” or “transfer” – the subject gives or transfers Obj₁ to Obj₂. Phrases like *He gave the ball to Ricky*, *He gifted them the goal* etc., are all examples of the ditransitive construction. The ditransitive construction is one but many argument structure constructions in the English language. Table 1 below lists 5 other phrasal level constructions.

Table 1: English Argument Structure Constructions

Caused-Motion: (Subj) V Obj Oblique _{path}	X CAUSES Y to MOVE Z
Resultative: (Subj) V Obj Pred	X CAUSES Y to BECOME Z
Transitive: (Subj) V Obj	X ACTS on Y; X EXPERIENCES Y
Removal: (Subj) V Obj Oblique _{source}	X CAUSES Y to MOVE from Z
Way construction: (Subj _i) V [poss _i way] Oblique _{path}	X CREATES PATH & MOVES Z _{path}

While all linguistic theories acknowledge the existence of phrasal level meaning, there is much debate on how just how *abstract* this meaning is, and the degree to which it depends on the meaning of a clause’s main verb. While the constructionist account argues that these form-meaning correspondences exist independently of the main verb (e.g. Goldberg 1995; Jackendoff 1997), others argue that this meaning is projected from the main verb’s specifications (via *linking rules*; e.g., Bresnan & Kanerva, 1989; Dowty, 1991; Grimshaw, 1990; Jackendoff, 1983;

Pinker 1989) or as *lexical templates* overlain on specific verbs (Rappaport Hovav & Levin, 1998).

In the following section, I will review the evidence for the constructionist position, and argue that such constructional meaning is truly abstract, in that the meaning is not specified by the main verb, but is wed to the abstract phrasal form in and of itself.

Chapter 2

Evidence for *abstract* argument structure knowledge

Adapted from:

Johnson, M., & Goldberg, A. (2012) On-line Evidence for Constructional Meaning:

Jabberwocky sentences prime associated verbs. *Language and Cognitive Processes*, 1: 1-

14

Introduction

Theoretical Evidence. There are good reasons to doubt that phrasal meaning is beholden to the main verb. I will first review the theoretical evidence for abstract constructions , and then review empirical evidence for speakers' knowledge of abstract constructions.

In frameworks *sans* phrasal meaning, such as in lexicosemantic theory (e.g. Lado 1990, Pinker 1989, 1994), the main verb of the sentence is generally taken to play a pivotal role in interpretation, by specifying the way that overt arguments are related to one another. Verbs are therefore classified according to how many arguments they “take.” It is generally assumed, for example that *sneeze* is intransitive, *kick* is transitive, and *give* requires an agent, a theme, and recipient arguments. In this way, basic sentence patterns of a language are believed to be determined by syntactic and semantic information specified by the main verb. For example, the sentence pattern in (2) appears to be due to the specifications of *put*:

1.Pat put the ball on the table.

That is, *put* is a verb that requires an agent, a theme and a location, and it appears overtly with the corresponding three complements.

If we are to believe that argument structure were projected exclusively from a verb's semantics, we would need special verb senses for each of the verbs in the expressions in (2) (Goldberg 1995; 2006; Jackendoff 2002):

- 2.a. If time is money then save yourself rich at Snyder's! (Mark Turner, personal communication)
- b. The people of this small town...have been unable to pray Mrs. Smith's two little boys home again. (Mark Turner, personal communication)
- c. "his thousands of travelling fans... had roared him into the Thomas and Mack Center ring" www.topix.net/wire/world-soccer/manchester-united
- d. She tried to avoid blinking the tears onto her cheeks (Anne Tyler, 1992, Dinner at the Homesick Restaurant, NY: Knopf)
- e. "Demi Moore thinks this will Halle Berry her back to the B List." <from Mr. Brooks movie> (Grush, personal communication 2007)
- g. "I actually had a moth go up my nose once. I...coughed him out of my mouth" (bikeforums.net/archive/index.php/t-292132)

That is, we would need a sense of *save* that meant roughly "to cause to become by saving"; a special sense of *pray* "to cause to move by praying," a special sense of *roar* that entails motion and so on. These senses are implausible in that one doesn't find languages that devote unique stems to these meanings. For example, it is unlikely that one would find a word *kamo*, meaning "to cause to move by coughing" because this is not a situation that is likely to occur regularly enough to warrant a lexical meaning (Goldberg, 2010).

Indeed, some verbs have a wide range of different meanings. Goldberg (1995) points out that "kick" is an especially promiscuous case, occurring in at least eight different constructions, assuming a different role in each, yet retaining its basic meaning:

1. Pat *kicked* the wall
2. Pick *kicked* Bob silly
3. Pat *kicked* the football into the stadium
4. Pat *kicked* at the football
5. Pat *kicked* his foot against the chair
6. Pat *kicked* Bob the football
7. The horse *kicks*
8. Pat *kicked* and screamed his way out of the room

The lexicosemantic position also puts a major burden on the speaker, compelling them to store multiple designations and stipulations in their knowledge of the same verb. Moreover, it is raises the issue of how speakers could learn the meaning of each verb in each frame, without an understanding of the argument structure initially (Gleitman 1990).

Empirical Evidence. In addition to these theoretical considerations, there have been numerous demonstrations that abstract phrasal forms (i.e. phrasal constructions), irrespective of the component lexical items, are strongly associated with a specific meaning. One example is that individuals overwhelmingly describe a novel verb as meaning “give” when it occurs within the ditransitive construction (e.g. *He **mooped** him the lorp*) (Goldberg 1992; Ahrens 1995). Other work has noted that learners use the semantics associated with syntactic patterns in order to figure out what new verbs mean (Fisher, 1996; Gillette, Gleitman, Gleitman, & Lederer, 1998; Landau & Gleitman, 1985); this idea presupposes the idea that the syntactic patterns are associated with meanings independently of the main verb.

Bencini & Goldberg (2000) conducted a sorting experiment with the aim of directly comparing the semantic contribution of the construction with that of the verb. The stimuli were

sixteen sentences created by crossing four verbs with four different constructions: A sample set of sentences for the verb *throw* is given in Table 2:

Table 2: Sample stimuli for Sorting Experiment

a. Pat threw the hammer	Transitive
b. Chris threw Linda the pencil	Ditransitive
c. Pat threw the key onto the roof.	Caused Motion
d. Lyn threw the box apart.	Resultative

Participants were asked to sort the sixteen sentences, provided in random order, into four piles based on “overall sentence meaning.” Participants could sort equally well by verb: e.g., all instances of *throw* (a-d) being grouped together, regardless of construction; or they could sort by construction: all instances of e.g., the ditransitive construction being grouped together. The stimuli were designed to minimize contentful overlap contributed by anything other than the lexical verb. No other lexical items in the stimuli were identical or near synonyms.

Results demonstrated that participants were just as likely to sort by construction as they were to sort by verb, providing evidence that the constructions were as strong a cue to sentence meaning as verbs.

These results are striking, as the only way for participants to discern which verb sense was involved was to recognize the argument structure pattern and its associated meaning. That is, the proposed different verb senses all look the same; the only way to determine that a particular sense is involved is to note the particular argument structure pattern that is expressed and infer which verb sense must have produced such a pattern. Therefore, at least from an off-line comprehension point of view, the pairing of argument structure pattern with meaning must be primary.

Kaschak and Glenberg (2000) also provide important evidence of constructional meaning

through a series of comprehension studies involving novel denominal verbs (Clark and Clark 1979). In particular, they asked participants to read passages that set up potential transfer contexts such as that in (3):

3. Tom and Lyn competed on different baseball teams. After the game, Tom, who had been pitching, was kidding her about striking out three times. Lyn said, “It was an aberration! I was distracted by your ugly face. I can hit anything to any field using anything!” To prove it, she took her apple over to manager who was recovering from a twisted ankle, and she grabbed his crutch.

Participants were asked to paraphrase sentences such as (4a) or (4b). Another group was asked to define the denominal verbs involved.

- 4. a. Lyn crutched Tom her apple to prove her point. (double object)
- b. Lyn crutched her apple to prove her point to Tom (transitive)

Results demonstrated that participants were more likely to decide that transfer had occurred in 4a than in 4b, and were more likely to decide that the novel verb (e.g., *crutch*) was a verb of transfer. In addition, Kaschak and Glenberg demonstrated that different aspects of the *affordances* of the denominal verb played a role in the sentences’ interpretations. For example, participants were faster to judge “the crutch is sturdy” as true after a passage like that in (4) than they were to judge “The crutch can help with injuries,” despite the fact that helping with injuries is more associated with crutches in general, as determined by various Latent Semantic Analysis measures (Landauer & Dumais, 1997). They conclude that “the syntax specifies a general scene, and the affordances of objects are used to specify the scene in detail sufficient” (pg 508).

Goldwater and Markman (2009) have likewise shown that instances of the middle construction (5) involving novel denominal verbs are more likely to be judged as being

nonsensical when followed by purpose clauses than passive constructions (6) involving the same novel verbs.

Middle:

5.??The ripe tomatoes had sauced well to complement the pasta at the gala dinner.

Passive:

6.The ripe tomatoes were sauced well to complement the pasta at the gala dinner.

They attribute the difference to the fact that only the passive construction requires a (possibly unspecified) agent argument. As in Kaschak and Glenberg (2000), Goldwater and Markman (2009) used novel denominal verbs in order to determine what role the construction played in assigning meaning. The meaning could not be ascribed to a preexisting denominal verb because the verbs were normally only used as nouns.

The studies discussed above collectively argue that argument structure constructions play a role in speaker's ultimate interpretations of sentences. But a critic might argue that all of these tasks lent themselves to strategic responding, since they are all either off-line tasks (Bencini and Goldberg 2000; Kaschak and Glenberg 2000), or tasks that require sensicality judgments (Goldwater and Markman 2009). Work demonstrating a role of constructional meaning in the acquisition of verbs ("syntactic bootstrapping") in younger children is compelling, but the aspects of meaning that have been demonstrated to date have focused primarily on the number of arguments involved (cf. also Goldwater and Markman 2009).

In fact, there have been virtually no experiments designed to determine whether contentful constructional meaning is accessed automatically and without reflection. Experiment 1 was motivated by this lacuna, as it investigates the possibility of automatically accessed constructional meaning, using in a speeded lexical decision task.

Materials

Four abstract constructions were used as primes. These are shown in Table 3. The experimental target words chosen for each construction are provided in Table 4. Each verb was presented in its past tense form.

Table 3: The four abstract phrasal constructions used as primes

Abstract Construction	Constructional frame used to create stimuli	Example
Ditransitive:	S/he nonseV-ed him/her the nonseN.	<i>He daxed her the norp</i>
Resultative:	S/he nonseV-ed it nonseAdj.	<i>She jorped it miggy</i>
Caused-motion:	S/he nonseV-ed it on the nonseN.	<i>He lorped it on the molp.</i>
Removal:	S/he nonseV-ed it from him/her.	<i>She vakoed it from her.</i>

We consider whether Jabberwocky type sentences such as those in Table 3 prime words related to their hypothesized meanings. Positive evidence of priming would be evidence that the constructions are *associated* with meaningful verbs. In particular, we consider whether argument structure constructions prime the following types of words: high frequency associates and lower-frequency associates. High frequency associates were chosen to be the verbs that (most) frequently occur in a particular construction. For example, *give* is the most frequent verb that occurs in the ditransitive, accounting for close to half of all tokens of the construction (Goldberg, Casenhiser, & Sethuraman, 2004; Stefanowisch and Greis 2003). Low frequency associates are verbs that appear in the construction, but markedly less frequently than the high frequency associates. For example, *hand* occurs in the ditransitive (e.g., *She handed him something*), but less frequently than *give* does. Statistics were gathered from the 400 million

word Corpus of Contemporary American English (COCA), and are provided in the Methods section (Table 5).

We also sought to investigate whether the constructions primed semantically related non-associates as a way to more directly determine whether semantic priming as well as associative priming would be evident. To this end, we also included semantically related non-associate target verbs. These were verbs that do not generally occur in the corresponding construction, but are semantically related to the meaning that has been hypothesized for the construction. For example, the meaning of *transfer* captures the hypothesized meaning associated with the ditransitive construction, and yet the verb itself rarely occurs in that frame (e.g., *?She transferred him something*). (quantitative data is given in Table 5). The particular verbs chosen for each of four verbs are provided in Table 4.

Table 4: Experimental target words for each construction

Constructions	High frequency associate	Low frequency associate	Semantically related-non associate
Ditransitive	<i>Gave</i>	<i>Handed</i>	<i>Transferred</i>
Resultative	<i>Made</i>	<i>Turned</i>	<i>Transformed</i>
Caused-motion	<i>Put</i>	<i>Placed</i>	<i>Decorated</i>
Removal construction	<i>Took</i>	<i>Removed</i>	<i>Ousted</i>

If results indicate evidence of priming, it would provide evidence that the “nonsense sentence” primes were not, after all, completely nonsensical. If the ditransitive, for example, devoid of a contentful verb or any open class lexical items, primes *gave*, it would provide evidence that the abstract pattern is associated directly with this verb. It is possible that constructions with nonse-words prime *only* the highest frequency verbs that can occur in them (such as *gave* in the ditransitive), because by hypothesis, the constructions are closely

associated with these verbs. Or it may be that only high or low frequency verbs that can appear in the construction are primed, but semantically related non-associates are not primed, because it may be that priming requires compatibility of the verb in the construction. The strongest finding would be that all three types of verbs are primed by their respective abstract constructions. The overarching goal of the study is to determine whether there exists on-line evidence in favor of abstract semantics being associated with syntactic frames that contain no open-class lexical items.

Methods

Participants. Forty Princeton undergraduate students, aged 18-23, participated in exchange for course credit. All participants were native English speakers.

Procedure. Participants performed lexical decisions following sentential primes. In particular, they were instructed that they would be presented with a phrase written in black, directly followed by a ‘word’ written in green. They were to read the sentence aloud, and then respond as quickly as possible, pressing ‘1’ if the green ‘word’ was a real word and ‘2’ if it was a non-word.

Construction primes. Stimuli included four different abstract constructions that included nonsense open-class items (See Table 1). The four constructions were completely abstract since all of the open class items involved nonsense words. These nonsense words were drawn randomly from pre-determined lists of twenty-five items that did not overlap with the

probe non-words. Prime sentences were created using all nonsense open class words. Nonsense words were chosen randomly from a set of 75 forms that had the typical morphophonological form of verbs (25), nouns (25), or adjectives (25). Each nonsense-word appeared randomly (with replacement) in the constructions. Example sentences for each construction type were given in Table 3.

Lexical targets. The target words of interest were high or low associates of one of the constructions, or words that were semantically related to one of the constructions, but which did not regularly occur in the construction. In order to classify verbs, we searched the 400 million word on-line Corpus of Contemporary American English. The frequencies of each of the target word in the relevant construction type are provided in table 6.


Table 6: Raw frequencies of target words in their congruent construction. Based on 400 million word COCA corpus

Construction (search string)	High frequency associate		Low frequency associate		Semantically related, non- associate	
Ditransitive (V pronoun the N)	<i>Gave</i>	2365	<i>Handed</i>	362	<i>Transferred</i>	0
Resultative (V it Adj)	<i>Made</i>	19,833	<i>Turned</i>	184	<i>Transformed</i>	1
Caused motion (V it on the N)	<i>Put</i>	591	<i>Placed</i>	224	<i>Decorated</i>	0
Removal (V proN from pronoun)	<i>Took</i>	394	<i>Removed</i>	17	<i>Ousted</i>	0

The general procedure was the same throughout the experiment (see Diagram 1). Participants were first presented with a fixation cross (3 seconds), then an abstract phrasal construction which they read aloud (1000 ms), then another brief fixation (300 ms). Lastly, they were presented with the target word to which they decided whether it was a word or non-word. Participants were instructed to read the prime sentence aloud (1000ms), then respond as quickly as possible to the target which directly followed (Mean RT = 609), indicating via button press whether it was a real word or non-word. Half of the target words were words, half non-words.

Participants received feedback on each trial on whether their response was correct or not, or whether they took too long (> 1000 ms.) to respond. Feedback screens were presented for 500ms each, after which a new trial would begin.

Figure 1: Basic structure for a priming trial



Within each different block type, half of the targets were real words and half were non-words. As explained below, different types of non-words and real words were presented during different blocks.

The dependent measure was reaction time, and the independent variable was prime type. In order to control carefully for length, frequency, and other factors, we compared reaction times to the *same* set of words, systematically varying whether the prime construction was either congruent or incongruent with the target word, across participants. For example, when *gave* was preceded by the ditransitive construction (e.g., *He jorped him the brap*), it is “congruent;” when *gave* is preceded by the removal construction it is incongruent.

Practice. After the task was explained to participants, they proceeded to a practice phase in which they were given 12 training trials, again half of which were real words and half of which were non-words (see stimuli). These training stimuli were only used during the practice phase. New items were used during the experimental trials. The purpose of the practice phase

was to get participants properly adjusted to the speed of the experiment. Although many participants found it difficult at first, by the end of the practice trials, all participants felt confident continuing.

Experimental Structure

Experimental trials were divided into 3 blocks based on the 3 priming categories: High Frequency (HF) Associative, and Low Frequency (LF) Associative and Semantic Non-associative (S). Four experimental words were seen for each experimental block. Participants were randomly divided into two groups such that one group saw two experimental words in a congruent context and the other two in an incongruent context; the other group saw the first two experimental words in an incongruent context, and the other two in a congruent context. In this way, each participant witnessed two words in each block in a priming context and two words without a relevant prime (see table 7).

All participants in both groups saw each experimental word exactly once. 20 participants (Group 1) saw and responded to *gave* and *took* after congruent constructions—i.e., after the ditransitive and the removal constructions, respectively, and responded to *made* and *put* after incongruent constructions—i.e., after the caused-motion and the resultative constructions, respectively. A second group of 20 participants (Group 2) saw the reverse, as shown in Table 7. Responses to each particular target word after congruent and incongruent primes were then compared between subjects.

Table 7: Target words and the prime construction given to each of two groups of participants

Prime Words	Group 1	Group 2
<i>Gave, Handed, Transferred</i>	Congruent (ditransitive)	Incongruent (removal)
<i>Made , Turned, Transformed</i>	Incongruent (caused-motion)	Congruent (resultative)
<i>Put, Placed, Decorated</i>	Incongruent (resultative)	Congruent (caused motion)
<i>Took, Removed, Ousted</i>	Congruent (removal)	Incongruent (ditransitive)

Again, the *same* target words were used following incongruent constructions to determine baseline RTs for each word (between subjects). Each block consisted of two phases: An ‘acclamation’ phase (12 trials), explained below, and an ‘experimental’ phase (8 trials). The overall structure of the experiment was as follows:

1. Practice (12)

-Break-

2. High Frequency – acclamation (12 trials: 6 words and 6 non-words)

3. High Frequency - experimental (8 trials: 4 experimental words and 4 non-words)

-Break-

4. Low Frequency – acclamation (12)

5. Low Frequency – experimental (8)

-Break-

6. Semantics – acclamation (12)

7. Semantics – experimental (8)

The reason for including the acclimation phase was to condition participants to attend carefully to target forms in order to avoid floor effects in response times. If words had been sufficiently distinct from non-words, the task would have been quite easy and might well have led to insufficient variability in response times. We therefore included target words and non-words that were minimally different from each other and from the experimental items in that block. For example, in the high frequency associates block, the experimental words were *gave*, *put*, *make*, and *took*. The words in the acclimation phase included verbs that are phonetically similar to those in the experimental phase (e.g. *have*, *met*, *stood*, *saw*), as well as non-words that are also phonetically similar (*stook*, *puv*, *goot*, *gade*). In the acclimation phase, participants saw 6 non-experimental real verbs and 6 non-words. While in the experimental phase, participants saw the 4 experimental words, as well as 4 non-words.

There was no transition from the acclimation phase to the experimental phase of each block. From the participant's perspective, the practice phase was simply followed by 3 seamless blocks.

After each block, participants were shown a break screen, telling them that they had finished block 1 or 2 of three. They were asked to press a space bar to continue whenever they felt ready. Participants typically only paused long enough to read the break screen before continuing on to the next block.

Results

Only response times of the 12 experimental words that were accurately responded to as words within the 1000 ms. time window were analyzed. This provided a total of 425 data points, with a range of 13-20 data ($M = 18$) for each cell (mean RT = 608 ms.).


By subtracting the reaction time for a given word after a congruent sentence from the reaction time for the same word after an incongruent sentence, we determine the extent of priming.

Table 8: Average difference in RT organized by construction and verb

Construction	High frequency associate		Low frequency associate		Semantically related, non-associate	
Ditransitive	<i>Gave</i>	137	<i>Handed</i>	135	<i>Transferred</i>	106
Resultative	<i>Made</i>	86	<i>Turned</i>	-35	<i>Transformed</i>	-7
Caused-motion	<i>Put</i>	104	<i>Placed</i>	82	<i>Decorated</i>	6
Removal	<i>Took</i>	120	<i>Removed</i>	128	<i>Ousted</i>	132
		111		77		59

We find an overall main effect of primed over unprimed by-subject, ($F(1,430) = 50.9$; $p < .001$). Considering each of the three conditions in turn, we find significant priming for strong and weak associates, $F(1, 136) = 43.5$, $p < .001$; $F(1, 153) = 21.7$; $p < .001$, and for semantically related non-associates, $F(1, 140) = 6.9$, $p = .009$ (see figure 2).

Figure 2: Reaction times for congruent and incongruent presentations, by category of target with example words given for the ditransitive construction.



Effects are expected to be weaker when considering performance by-item because we had relatively few items. Yet we still find an overall priming effect, $F(2,23) = 11.05$, $p < .01$. Considering each category, we find significant priming for both strong associates ($F(2, 7) = 125.8$, $p < .001$), and weak associates ($F(2, 7) = 7.2$, $p = .036$). However, we do not find a significant priming effect for the semantically related non-associates ($F(2, 7) = 1.6$, $p = .25$).

If we consider each of the four constructions in turn by-subject, we find significant priming of the ditransitive ($F(1, 106) = 44.6$, $p < .001$), the caused-motion ($F(1, 114) = 7.76$, $p = .006$) and the removal constructions ($F(1, 98) = 18.02$, $p < .001$). The only construction that does not show significant across the board priming is the resultative ($F(1, 106) = .79$,

$p=.364$). When considering the constructions by-item, the caused motion and the resultative do not reach significance ($F(2,2) = 4.5$, $p = .16$ and $F(2,2) = .163$, $p = .725$, respectively), but the ditransitive and the removal constructions do ($F(2,2) = 167.1$, $p = .006$ and $F(2,2) = 361.2$, $p = .003$, respectively).

Discussion

We compared the reaction times for target words following the construction that is expected to prime them (congruent) with reaction times for the same target words following one of the other three constructions (incongruent). Overall, the same words were recognized significantly faster after instances of semantically related constructions than after instances of unrelated constructions, when analyzed both by-items and by -subjects. The fact that the by-items analysis was significant overall is particularly striking since we used relatively few items (three for each of four constructions). More specifically, both strong and weak associates of constructions were primed by the Jabberwocky sentences, by-items and by-subjects.

The evidence for priming of semantically-related non-associates was mixed: priming was significant by-subjects only. We should also note that purely semantic priming is well-known to be weaker than associative priming (Lucas, 2000), and so the fact that the effect was more fragile for non-associates is not unexpected.

When considering individual constructions, the weakest results were found for one particular construction: the resultative construction which did not exhibit priming either by-subject or by-item. In order to understand why this might have been the case, recall that priming was determined by comparing reaction times after the congruent construction to

reaction times after an incongruent construction. The “incongruent” construction in this case was the caused motion construction (??*She made/turned /transformed it into the room*).

However, the target verbs can appear in the caused-motion construction with a metaphorical change of state interpretation as in 8:

7. She made/turned/transformed the clay into a vase.

Moreover, we used sentences such as *He lorped it on the molp*, using the preposition *on* which can readily be interpreted as a locational adjunct, and can in that capacity be used with transitive uses of the relevant verbs:

8. She made the model on the table.

9. She turned the car on the street.

10. She transformed the clay on the table.

Thus the null effect in the case of the resultative may have been due to the fact that the “incongruent” construction was not altogether incongruent.

It might be suggested that evidence of priming is not the same as demonstrating that the construction “has” meaning. This critique raises the thorny issue of the nature of meaning. At the least, we have demonstrated that constructions prime meaningful words that occur in them, both with high frequency and with markedly lower frequency. It has been argued that associative priming is a type of semantic priming (McRae & Boisvert, 1998). In fact, Thompson-Schill et al. (1998) found that associates only show priming if they are in fact semantically related. It is left for future work to determine conclusively whether semantically related non-associates are primed by constructions since the present evidence was mixed on this point: priming was significant by-subjects but not by-items.

In any case, it is sufficient to note that each construction primed associated meaningful words. If, for example, the ditransitive makes one activate *give*, then this provides a cue to interpretation. In fact it has been argued that words “have” meaning in much this way: they serve as cues that combine to direct a comprehender to an interpretation (Elman, 2009). The on-line nature of the task indicates that constructions provide cues to interpretation in real-time.

The present results do not speak to the question of whether argument structure constructions are necessarily phrasal or whether they may be conceived of as lexical templates (Hovav and Levin 1998; Mueller 2006). But the present work provides evidence that they are not exclusively tethered to particular verbs: meaningful argument structure constructions are evoked even by Jabberwocky sentences.

Conclusion

There is a growing trend toward distinguishing a verb’s inherent or “core” lexical semantics from the semantics associated with the grammatical structures in which the verb can occur (Goldberg, 1992, 1995; Jackendoff, 1997; Rappaport Hovav & Levin, 1998). At the same time, there is also much work that presupposes that all meaning comes from the lexicon. The present findings indicate that phrasal abstract constructions are associated with semantics even when they contain no open class lexical items, and that the meaning is accessed automatically. Constructions prime not only main verbs with which they regularly occur, but they also prime main verbs that are low-frequency associates, and may even prime semantically related non-associates in some cases. This evidence of semantic priming indicates a tight link between

syntax and semantics in the domain of argument structure. That is, argument structure constructions are associated with meanings.

We have seen so far that understanding these inherent patterns is key to being able to flexibly use and understand one's language. How these stable linguistic representations are initially learned will be the subject of the rest of this dissertation.


Chapter 3

Introduction to Construction Learning

How does one go from discrete pieces of information to broad abstract concepts? And how do these abstractions then allow us to intelligently categorize and assimilate new items and events? Until relatively recently it was inconceivable that such open, domain general questions could be asked with reference to language learning. Instead, traditional linguistic theory has held that language could only be learned by a specialized language-specific mechanism, and that general cognition alone could not account for such a complex task. However, converging evidence has begun to suggest that the language learning process may crucially involve such a broad learning mechanism. In this view, meaningful linguistic patterns are learned via generalization; by noticing and abstracting over the shared similarities across linguistic exemplars.

Put simply, such patterns are thought to be learned by generalizing over a structured set of exemplars, abstracting these form/meaning pairings from the utterances one hears (Tomasello 2003; Goldberg 2006). For example, the abstract ditransitive form, discussed above to be reliably associated with ‘give’, can be learned by witnessing the consistent pairing of linguistic form (Subj V Obj₁ Obj₂) and meaning (give) in the encountered examples of this construction (see figure 3). Thus, generalization is a vital skill in language learning, allowing for the compositionality of language to unfold.

Figure 3: The generalization of specific phrases, giving rise to an abstract form which can then be extended to new instances (“She mooped him something”)



Previous research has demonstrated that adults can learn a novel construction with a relatively small amount (3 mins/16 exemplars) of structured input. In the earliest incarnations of this work (Goldberg, Casenhiser, and Sethuraman, 2004; Casenhiser and Goldberg 2005), participants were exposed to video exemplars pairing the action of ‘appear’ with the novel phrasal form *NP NP verb*. During exposure phase, participants were simply instructed to ‘watch and pay attention to’ short (6-8 seconds) video clips depicting various scenes, which were each overlaid with a linguistic narration. In each of the scenes, an object would appear on or in a certain location. For example, a king would fall out of the sky into a chair, or a ball would roll onto a stage. The linguistic narrations took the following syntactic form:

<noun phrase_(theme) noun phrase_(location) nonsense verb+o>

The audio description was given in the simple present tense at the start of the scene and was repeated in the past tense at the end of the scene. For example, given a scene where a spot

appears on the king's nose, subjects heard, *The spot the king moopos... The spot the king moopoed*. The entity named by the first noun phrase appears in the place named by the second noun phrase (see figure 4).

Figure 4: Example trial of the ‘appear’ construction. A king following into a chair, narrated by the phrase “The king the chair moopoes .. the king the chair moopoed”



After this exposure, participants were given a test to gauge how well they were able to learn this new construction. This took the form of a forced choice task, in which two videos played side by side, with one overlapping narration. Participants were simply instructed to select which video the narration was referring to. The videos consisted of either a novel act of appearance (A balloon falling onto a monster) or a intransitive action (A monster biting a balloon). The narrations were either an intransitive description involving a novel verb (e.g. *the monster lems the balloon*), or the appear construction, also involving a novel verb (e.g. *the monster the balloon pookoes*). Thus, if the narration was the latter description (*the monster the balloon pookoes*), the video of appear would be the correct choice. Half of the videos were

narrated by the appearance construction, while the other half were narrated by the intransitive construction. In essence, this gauged how well participants were able to distinguish this new construction from with which they were already familiar. Results from these initial experiments show that both children and adult participants were able to readily generalize this novel construction, performing above chance and significantly better than a no-sound control group.

Additional follow up work replicated these results without the morphological –o marker that provided a stable, concrete cue. With order as the only cue to the meaning of the construction, participants were able to readily learn this construction. These initial results provided the first example that the meanings of abstract constructions could be readily learned.

An interesting question concerns how input shapes this generalization process. In extensions of this work, Goldberg and Casenhiser (2005) varied the structure of the input, and created two experimental groups. In one group, the input was ‘skewed’ to a single, prototypical exemplar; the first verb was heard 4 times, with each of the remaining 4 verbs being heard just once (4-1-1-1-1). Corpus data reveals that this type of input actually closely mirrors the type of input children receive – the majority of syntactic frames have a single, prototypical verb that appears much more frequently than other verbs. For example, *give* typifies the meaning of the ditransitive frame (He/She Verb NP₁ NP₂), and is the most frequent verb to appear within it.

Other examples can be seen in table 8:

Table 8: Total verb types and frequency of the most common verb within three common constructions (from Goldberg and Casenhiser, 2005)

Construction	Total verb types	Most Common Verb	Frequency of most common verb
Subj Verb Loc	39	<i>go</i>	39% (136/353)
Subj Verb Obj Obl	43	<i>put</i>	38% (99/256)
Subj Verb Obj ₁ Obj ₂	13	<i>give</i>	20% (11/54)

In order to see if this more naturalistic, skewed input facilitates learning, performance with this group was compared to performance after witnessing a more balanced array of exemplars. In the balanced condition, subjects heard five different novel verbs across the eight different instances, each heard with relatively low frequency (1-1-2-2-2). While participants in the balanced condition still showed evidence for learning the construction, they were significantly outperformed by the participants in the skewed condition. This suggests that having the overall input dominated by low variance, prototypical exemplar facilitates the generalization process. This effect has also been replicated in children aged both 5 years, and 7 years.

Intriguingly, this performance mirrors observations from non-linguistic categorization, where it has been long known that early instances are heavily weighted in category formation (Bruner, Goodnow and Austin 1956). Elio and Anderson (1984) conducted a generalization task in which participants learned to categorize ‘people’ into two clubs based on several different features. They exposed participants to two different categories: a ‘centered’ condition and a ‘representative’ condition. In the centered condition, participants were initially shown more prototypical exemplars which then gradually transitioned to include other category members. In the ‘representative’ condition, participants received exemplars in a random, unbiased ordering.

Importantly, each group saw exactly the same exemplars – the only difference was the order in which the prototypical exemplars were viewed.

Similar to the Goldberg and Casenhiser (2005) findings, categorization performance was significantly greater in the ‘centered’ condition, when early exemplars were more prototypical, than in the representative case. This finding supports the link between linguistic and non-linguistic categorization, suggesting that both processes are accomplished by, or rely on at least in part, a common mechanism.

While the previous investigations described above have shown that correspondences between phrasal forms and abstract meanings can be learned with relative ease, little is known about exactly *how* they are learned. And while the behavioral parallels with non-linguistic categorizations are intriguing, there has been little work directly investigating how these mechanisms may actually overlap. The purpose of the following study is to investigate the neural mechanisms that underlie this type of linguistic generalization, in order to provide initial insights into the processes involved. Specifically, we sought to better understand the dependence of construction learning on brain systems that mediate related but different forms of learning as well as the brain network relevant to language.

Chapter 4

Construction Learning and The Brain

This chapter is based closely on a draft by Johnson, Turk-Browne and Goldberg (submitted)

Language comprehension evokes activity throughout the brain insofar as it involves mental imagery and simulation (Barsalou, 2008; Pulvermüller, 2005). For example, if comprehending language that simulates visual imagery or motor actions, corresponding brain areas in occipital or motor cortex can become active. At the same time, there is also a more targeted language network activated regardless of the content of the language being comprehended. This language-relevant network contains a broad fronto-temporal system including traditional Broca's and Wernicke's areas (Fedorenko, Behr, & Kanwisher, 2011; Friederici, 2009; Geschwind, 1975; Grodzinsky & Friederici, 2006; Hickock, Bellugi, & Kilma, 1998; Kaan & Swaab, 2002; Kemmerer 2012; Osterhout, Kim, & Kuperberg, 2007; Stowe, Haverkort, & Zwarts, 2005; Ullman, 2001). It has been argued that such a network may be specialized for, or even devoted to, linguistic functions (Fedorenko, Nieto-Castañon, & Kanwisher, 2011; Optiz & Friederici, 2004; but see Kaan & Swaab, 2002).

The majority of work in neurolinguistics has focused on language processing, with relatively less research on language *learning* itself (Gernsbacher & Kaschak, 2003). Indeed, several domain-general processes have been implicated in language learning, suggesting that regions outside the classic language-processing network may be involved.

First, *statistical learning* (SL) contributes to language learning (for a review, see Saffran, 2003), and is recognized to be robustly domain-general in that it can operate over visual (Kirkham, Slemmer, & Johnson, 2002; Turk-Browne, Jungé, & Scholl, 2005; Fiser & Aslin,

2005), auditory (Saffran, Aslin, & Newport, 1996), and even tactile input (Conway & Christiansen, 2005). Correspondingly, at the neural level, SL is supported by multiple memory systems, including the hippocampus and striatum (Schapiro, Kustner, & Turk-Browne, 2012; Turk-Browne, Scholl, Chun, & Johnson, 2009), which lie outside of the classic language-processing network.

Second, language learning has much in common with *category learning* (Bates & MacWhinney, 1987; Bybee, 1985; Goldberg, 2006; Tomasello, 2003). Both language learning and category learning involve the acquisition of a structured set of exemplars, including prototypical and less prototypical instances (Lakoff, 1987; Medin, Altorn, & Murphy, 1984; Taylor, 1995). Productive extension of constructions to new instances is dependent on the *coverage* of witnessed instances (Suttle & Goldberg, 2011), as has been claimed for categorization more generally (Osherson, Smith, Ormond, Lopez, & Shafir, 1990). In both language learning and category learning, spaced exposure provides an advantage over massed exposure (e.g., Childers & Tomasello, 2002; Merriman, Rovee-Collie, & Wilk, 1997) and skewed (low variance) input provides an initial advantage over more balanced (high variance) input (Casenhiser & Goldberg, 2005; Elio & Anderson, 1984).

Finally, language learning may also depend on *prediction*. It has long been recognized that humans use partial cues across a wide variety of contexts in order to anticipate future events (Tolman, 1932). We know that language is comprehended incrementally, with listeners anticipating upcoming words (Allopenna, Magnuson, & Tanenhaus, 1998; Arnold, Losongco, Wasow, & Ginstrom, 2000; Marslen-Wilson, 1973; Tanenhaus, Spivey-Knowlton, Eberhard, & Eberhard, 1995). In fact, cloze probability shows a strong inverse correlation with the strength of the N400 effect, known to be involved in lexical access and semantic integration (King & Kutas,

1995; Kutas & Hillyard, 1984) — the less predictable an upcoming word, the stronger the N400. Prediction is possible in language because language involves probabilistic regularities that constrain our expectations and provide cues about upcoming words, phrases, and content (Bates & MacWhinney, 1987; Pickering & Garrod, 2004, 2007; Goldberg, Casenhiser, & Sethuraman 2005). While prediction in these cases is a *consequence* of knowing a language, it can also be an essential part of the learning process itself (see Rescorla & Wagner, 1972); prediction may thus contribute significantly to construction learning. For instance, the ability to use word predictability to facilitate comprehension under noisy input conditions is positively correlated with sensitivity to transitional probabilities (Conway, Bauernschmidt, Huang, & Pisoni, 2010). Relatedly, sensitivity to sequential expectations correlates positively with the ability to successfully process complex, long-distance dependencies in natural language (Misyak, Christiansen, & Tomblin, 2010).

Thus, construction learning may depend at least to some extent on brain systems that are more domain-general. Further support for this hypothesis comes from the field of artificial grammar learning (AGL) — the focus of much previous research on the neural bases of language learning (see also Abia, Katahira, & Okanoya, 2008; McNealy, Mazziotta, & Dapretto, 2006). In such tasks, learners are exposed to strings generated by a finite state grammar and then tested on whether they correctly recognize new instances generated by the same formal rules (Reber, 1967). During this latter test phase, both language-relevant and other regions have been implicated (Fletcher, Büchel, Josephs, Friston, & Dolan, 1999; Optiz & Friederici, 2004; Seger, Prahakharan, Poldrack, & Gabrieli, 2000; Skosnik et al., 2002). For example, Broca's area and ventral pre-motor regions are involved in *similarity*-based AGL (grammaticality based on

similarity of exemplars), whereas the hippocampus is engaged in *rule*-based AGL (grammaticality based on presence/absence of specific features) (Optiz & Freiderici, 2004).

Although there has been a growing focus on the neural underpinnings of AGL, there has been a notable absence of studies on construction learning, which critically involve generalizations about meaning as well as form. Moreover, most fMRI studies of AGL have scanned participants at test (e.g., Lieberman, Chang, Chiao, Bookheimer, & Knowlton, 2004; Skosnik et al., 2002), examining the processing of grammars after learning has occurred, rather than the online learning processes that support the acquisition of the grammars.

Experiment 2 seeks to address this lacuna. We employ fMRI during exposure to novel form-meaning correspondences to monitor construction learning as it unfolds in time. Our goal is exploratory, as we seek to investigate which brain regions are involved and how they participate in construction learning. In order to ensure our design represents a close analogue to natural learning, the present study satisfies three desiderata: First, participants are scanned as they are exposed to a novel grammatical pattern, and we focus our analysis on fMRI activity reflecting online learning of the patterns. Second, the grammatical patterns are paired with abstract meanings, since a crucial part of acquiring a language is learning how forms are used to convey functions. Finally, no explicit instructions are given to participants in order to make the acquisition task more ecologically valid, since language is naturally learned without explicit corrective feedback (e.g., Braine, 1971; Brown & Hanlon, 1970; Marcus, 1993; Pinker, 1989).

Methods

Participants. Sixteen, right-handed native English speakers (age 18-32) participated in this study in return for \$30 in compensation. The study protocol was approved by the Princeton University IRB. Informed consent was provided by all participants.


Procedure. The task consisted of two phases, an exposure phase and a testing phase. During the exposure phase, participants witnessed both an *approach* construction and an *appearance* construction.

Approach construction. Each exposure trial began with two characters (a construction worker and a doctor) standing roughly 10 feet apart. The character on the right approached the one on the left. The linguistic narration was “NP₁ NP₂ nonsense-verb”, with the first noun phrase approaching the second noun phrase in a particular way. For example, if a film clip showed a doctor hopping towards a construction worker, then the sentence “The doctor the construction worker vakoos” would be heard simultaneously with the hopping motion. Five different manners of motion (swimming motion, jumping jacks, crab walk, marching, and walking) were used, and five different novel verbs were used to describe these actions. Despite these differences in verbs and approach styles, the basic pattern held constant: NP₁ approached NP₂, narrated by the linguistic form “NP_{1(agent)} NP_{2(goal)} nonsense-verb” (Figure 5A).

Appearance construction. For the ‘appearance’ construction, the video exemplars were scenes of various objects appearing in or on various locations. The linguistic narration was

“nonsense-verb $NP_{1(\text{theme})}$ $NP_{2(\text{location})}$,” and the theme (named by NP_1) would appear on or in the location (named by NP_2). For example, participants would witness a king falling onto a chair and hear “moopoes the king the chair” (Figure 5B). The manners of motion were different for each of the 12 exemplars; falling from the sky, ‘appearing’ out of nowhere, rolling into view, etc.

Figure 5: Sample exemplar from each construction. (A) In the approach construction, a construction worker might march toward a doctor, narrated by “The construction worker the doctor vakoes”. (B) In the appearance construction, a king puppet might fall into a chair, narrated by “Moopoes the king the chair”.



These two constructions were chosen because previous behavioral work suggests they are easily learnable with brief (3 minutes) exposure. In addition, we deliberately chose two constructions that are quite different: the approach construction is very concrete and visually

predictable, while the appearance construction is more abstract and less visually predictable. These differences helped us investigate various aspects of the construction learning process.

Exposure phase. Participants were instructed to simply “watch the videos and pay attention to the language that describes them.” Importantly, no further explicit instruction was provided — the viewing was completely passive and thus the learning was unsupervised. Participants were given a button box upon entering the scanner, but they were told it would be used in a second study that was separate and different from the video-viewing task.

During exposure, participants were shown short (6-10 s) audio-visual clips that provided the opportunity to learn each of the two novel constructions. Each run consisted of 12 videos: eight clips of the novel construction and four filler clips. For the approach construction, the filler clips were intransitive scenes paired with conjoined intransitive sentences (e.g., “the princess and the basketball player are zekking.”). For the appearance construction, they were semantically transitive scenes paired with transitive sentences (e.g., “the monster lemmmed the balloon”). The filler constructions were included to ensure that not all of the exposure clips in one run involved scenes of approach (or appearance). This was important for the random control presentation as described below.

For each construction, there was a *patterned* condition and a *random* control condition. In the patterned condition, the pairing of the depicted actions and the linguistic narrations was consistent (as described above), which enabled the abstract generalization to be learned. In the random condition, the same videos were viewed in the same order as during the patterned condition; however, the linguistic narrations were scrambled. Specifically, the same individual words were witnessed, but the order of the words was randomized: For example, instead of

consistently hearing “NP₁ NP₂ nonsense-verb” for approach and “NP₁ and NP₂ are nonsense-verb-ing” for fillers, participants were equally likely to hear “nonsense-verb NP₂ NP₁” and “nonsense-verb and NP₁ NP₂”, respectively. Thus, it was impossible to learn a stable association between form and meaning in the random condition because no stable pattern existed in the input. The filler items served to prevent participants from inferring a generalization from the visual scenes alone since the scenes presented were not uniformly scenes of approach or appearance. Importantly, the perceptual features of the random condition were identical to the patterned condition, but only in the latter did the systematic language input allow participants to relate, for example, all “NP₁ NP₂ nonsense-verb” exemplars with scenes of one entity approaching another, and all “NP₁ and NP₂ nonsense-verb” exemplars with scenes involving a conjoined intransitive action.

Each participant completed four scanning runs corresponding to each of four conditions (Table 1). The 12 videos for each construction were randomized across participants, and this order was held constant for the patterned and random runs of each construction. The order of runs was counterbalanced across participants with the exception that the patterned and random conditions of a given construction never appeared back-to-back; this was done to avoid direct comparison of patterned and random versions of the same construction. Participants were given the option of taking a brief break between runs, and then the experimenter started the next run by saying “OK, here comes the next one”; this provided a cue that the next set of stimuli was distinct from previous set.

Table 9: Experiment 2's four exposure conditions.

	Approach Runs		Appearance Runs	
	Approach Construction (8)	Filler (4)	Appearance Construction (8)	Filler (4)
Patterned	“NP ₁ NP ₂ verb” NP ₁ approaches NP ₂	“NP ₁ and NP ₂ are verb-ing” NP ₁ and NP ₂ perform an intransitive action	“verb NP ₁ NP ₂ ” NP ₁ appears on/in NP ₂	“NP ₁ verb NP ₂ ” NP ₁ does something to NP ₂
Random	Shuffled narrations over the same videos		Shuffled narrations over the same videos	

Testing phase. Although we focused data analysis on the exposure phase, participants were also tested while in the scanner. Participants' knowledge of the two constructions was assessed by determining how accurately they could generalize beyond the exemplars witnessed during the exposure phase. Specifically, participants were given a forced-choice task, in which they were asked to identify new instances of the novel constructions (involving new nouns and new nonsense verbs, paired with new scenes) and distinguish them from new instances of familiar constructions.

On each trial, two videos played simultaneously while a single sentence was heard. The videos were similar in length to those in the exposure phase (6-10 s). After each video played, participants were shown a prompt on the screen: “Did the narration match the video on the left or on the right?” Participants signaled their responses with a button box. They had up to 10 s to respond to the prompt, after which the screen would turn black until the next trial commenced. There were eight trials in each testing run, with each trial beginning at 30 s intervals, and one

testing run for each construction. Foils included the same two arguments as in the target scene, but instead of being involved in approach or appearance scenes, the two arguments were performing either a semantically transitive action or a conjoined intransitive action.

fMRI Analysis

Data acquisition. Scanning was performed on a Siemens 3T Allegra head-only scanner using a Nova Medical NM-011 Head Transmit Coil with receive-only array system. Functional sequences were acquired using a T2*-weighted EPI sequence (echo time = 30 ms, repetition time = 2000 ms, flip angle = 64°, matrix = 64 x 64). Twenty-six interleaved oblique axial slices were aligned parallel to the anterior-commissure/posterior-commissure, and whole-brain coverage was obtained at a resolution of 3.5 x 3.5 x 3.5 mm. For alignment, two T1-weighted structural images were acquired: a coplanar FLASH image and a high-resolution MPAGE image.

Preprocessing. The first four volumes (8 s) were discarded from each run in order to allow for T1 equilibration. Data were preprocessed using FSL's FEAT (Woolrich et al., 2009; Smith et al., 2004; www.fmrib.ox.ac.uk/fsl), including: corrections for head motion and slice acquisition time, high-pass temporal filtering with a 128 s period cut-off, and spatial smoothing with a 5 mm kernel. Functional runs were registered to the coplanar (FLASH) structural image, and then to the high-resolution (MPAGE) structural image. Data were normalized to Montreal Neurological Institute (MNI) space and interpolated to 2 mm isotropic voxels.

Patterned vs. Random. Each run was fit using a general linear model (GLM) that contained one regressor modeling the average blood-oxygen level-dependent (BOLD) response across videos in that run. Each video was modeled with a boxcar function convolved with a canonical hemodynamic response function. Since the videos were of variable lengths and occurred in a random order across participants, unique regressors were created for every run of each participant. Since learning in this context requires noticing shared structure across exemplars, it can be reasonably assumed that learning can only begin after the first two videos have been witnessed. For this reason, we analyzed only the final 10 video exemplars from each run. To compare across conditions, we contrasted parameter estimates obtained in the GLM of each run in a higher-level GLM treating subject as a random effect. Specifically, we contrasted patterned vs. random separately for each construction and also collapsing across constructions. A pre-threshold mask was generated by averaging the overall patterned and random responses from each participant together. The resulting voxelwise statistical maps were thresholded at $p < 0.05$ corrected, using a cluster-mass correction and an initial Z threshold of 2.3.

Behavioral correlations. To examine brain-behavior relationships, we used individual test performance as a regressor for BOLD responses during the exposure runs. First, we demeaned across participants the behavioral scores for overall performance, as well as for each construction individually. A new higher-level GLM was then computed, where the parameter estimate for each participant and condition was weighted based on the behavioral score. In other words, this analysis produced voxel clusters in which there was a significant relationship between evoked BOLD responses and test performance (same threshold as above). These relationships were examined for both patterned and random runs, separately for each


construction as well as collapsing across constructions. The random runs provided a control in that no relationship was expected because no learning was possible.

Time course analyses. Because we expect learning to increase over time with more exposure, we modeled such a change using a parametric regressor. Specifically, within individual participants, we weighted the predicted BOLD response for each block based on the block number in the run: -5.5, -4.5, -3.5, -2.5, -1.5, -0.5, 0.5, 1.5, 2.5, 3.5, 4.5, 5.5. Critically, these weights summed to zero, and thus the parametric regressor only captured the linear change in response over the run (no main effects). This analysis was applied to all four run types. To compare these linear effects across conditions, we computed a higher-level GLM separately for each construction and collapsing across constructions. Furthermore, we examined the relationship between these linear changes and behavior, using test performance as a covariate for the higher-level GLM, as described above.

Results

Behavioral Results. Test performance confirmed that participants learned the constructions. Across participants, mean accuracy was 80% (SD = 23%) for the approach construction and 73% (SD = 17%) for the appearance construction. These levels were greater than chance (50%) for both constructions (approach: $t(15) = 5.18, p < .001$; appearance: $t(15) = 5.51, p < .001$), and did not differ significantly from each other ($t(31) = 1.68; p = .21$). Figure 6 depicts the distribution of participants' performance, averaged over the two constructions.

Figure 6: Distribution of overall test performance averaging across approach and appearance constructions, with each point representing one participant. Chance is 50%.




fMRI Results

Patterned vs. random. No regions showed activation in the contrast patterned > random. In the reverse contrast, random > patterned, there was significant activation (corrected, $p < .001$) in medial prefrontal cortex (mPFC) and left supramarginal gyrus (SMG) (see figure 7). Interestingly, both of these regions are also implicated in additional, independent analyses discussed below. The patterned vs. random contrasts were also examined separately for each construction. The effects in L SMG and mPFC were not present when the constructions were tested alone. For the approach construction, no significant differences were observed. For the appearance construction, the contrast of random > patterned revealed activation in left occipital cortex; no differences were observed for the patterned > random contrast.


Figure 7: Contrasting random and patterned runs, collapsed across both constructions.

Left supramarginal gyrus (A) and medial prefrontal cortex (B) were significantly more active in random presentations than in patterned presentations.



Behavioral correlations. Because there was variability in behavioral test performance (see Figure 2), we were able to examine brain-behavior relationships. For the patterned condition, no regions showed a positive correlation with behavior. However, we observed a robust negative correlation — i.e., lower activation predicts better test performance — in the ventral striatum (Figure 8A). As expected, there were no significant positive or negative correlations between activation in the random condition and behavior. Considering each construction in turn, we did not observe any significant correlations for the random condition of either construction, nor for the patterned condition of the approach construction (Figure 8C). However, for the patterned condition of appearance construction, there was a reliable ($p < .001$) negative relationship between behavioral test performance and activity in bilateral ventral striatum (Figure 8B). Thus, the overall negative ventral striatum correlation for patterned appears to be driven by the appearance construction.


Figure 8: Correlation of test performance with activation in the patterned condition. (A) Collapsing over constructions, bilateral ventral striatum was negatively correlated with behavior (left: $x = -18, y = 10, z = -8$; right: $x = 18, y = 10, z = -8$). (B) This effect was robust within the appearance construction condition, but (C) was absent in the approach construction condition.



Time course analyses. Learning of an abstract construction in the patterned runs can only occur after multiple exemplars have been witnessed. To test for learning effects that evolved during the exposure runs, we fit a parametric model in which we predicted that the BOLD response would change linearly as a function of the number of exemplars that had been witnessed. For the patterned condition, we observed activation in the precuneus and posterior cingulate (Figure 9A). For the random condition, we also observed activation in posterior cingulate (Figure 9B). Contrasting these two statistical maps revealed a greater linear change for patterned vs. random in the precuneus (Figure 9C); no regions showed a greater change for random vs. patterned. The specificity of the precuneus effect to the patterned condition suggests that the activation does not result from boredom or fatigue, and must instead reflect some

component of construction learning. We tested for these effects in patterned and random conditions separately within each construction and no regions reached significance, possibly due to reduced statistical power.

Figure 9: Linear trend as a function of the number of exemplars. (A) For the patterned condition, precuneus ($x = -2, y = -65, z = 20$) and posterior cingulate ($x = -1, y = -52, z = 18$) showed increasing activation. (B) For the random condition, only the posterior cingulate cluster remained. (C) The greater trend in the patterned condition was confirmed in a contrast of patterned vs. random.



Insofar as linear changes reflect a learning process, they might predict test performance. To explore this possibility, we correlated behavior with the amount of increasing or decreasing activation for the patterned and random conditions. For the patterned condition, better behavioral learning was predicted by increasing activation in mPFC and decreasing activation in left SMG over time. For the random condition, no regions showed a correlation between behavioral learning and positive or negative linear trends in activation. Interestingly, both mPFC and SMG were also observed in the contrast of average activation in random > patterned conditions. It is

important to note that these two effects are independent: The linear change in a region is orthogonal to its average activation, and the current analysis relies on an individual differences correlation with behavior rather than a main effect collapsed across individuals.


Discussion

Summary of findings. This study provides an initial investigation of the neural systems that support abstract construction learning. We found that mPFC and left SMG were less active during patterned exposure, in which there was a systematic relationship between form and meaning, relative to random exposure, in which there was no systematic relationship. The linear time course analysis further revealed that these regions were significant predictors of behavioral learning; the precuneus showed a greater change in activation over time in the patterned than random conditions, but this change did not correlate significantly with behavior. When using behavior as a covariate for overall activity rather than the time course of activity, we found a robust inverse relationship between the ventral striatum and test performance, which was limited to the patterned condition. Analyzing the constructions separately revealed that the appearance construction alone showed deactivation of left occipital cortex in the patterned vs. random conditions, and also revealed that the overall ventral striatal relationship was largely driven by the appearance construction. These results have several interesting implications.

Categorization. As noted earlier, the learning of linguistic constructions has been likened to a kind of similarity-based categorization. This idea is reinforced by two of our findings. First, the time course analysis found increasing engagement of the posterior precuneus

during learning. This engagement cannot be due to the mere presence of visual or linguistic input, or any general factors such as boredom or fatigue, since it was not observed in the random condition that controlled for these factors. Interestingly, the posterior precuneus has been linked to non-linguistic categorization. For example, in one study (Koenig et al., 2005) participants were shown novel creatures called ‘crutters’, comprised of various features, and asked to categorize each creature based on rules (Figure 10A) or overall similarity (Figure 10B). In the similarity-based condition but not the rule-based condition, the posterior precuneus was more active in later vs. earlier training trials. This effect was interpreted in terms of the need for retrieval of exemplars in the similarity-based condition, given that the posterior precuneus has been strongly linked to episodic memory retrieval (Cavanna & Trimble, 2006; Krause et al., 1999; Platel, Baron, Desgranges, Bernard, & Eustache, 2003; Tulving, 1983).

Figure 10: Involvement of the precuneus in non-linguistic categorization (adapted from Koenig et al., 2005). (A) In the rule-based condition, participants were asked to select which animal had three of the four features depicted. (B) In the similarity-based condition, participants were asked to sort based on the overall relationship of the animals to the target (top row). The posterior precuneus was more active late vs. early during learning in the similarity-based condition.



In the present study, posterior precuneus activity similarly increased with the number of encountered exemplars. Insofar as this activation reflects retrieval of past exemplars, retrieval may play an important role in linguistic generalization (Bates & MacWhinney, 1987; Bybee, 1985; Goldberg, 2006; Tomasello, 2003). It is worth noting that any such retrieval likely occurred incidentally, in the absence of explicit task requirements to retrieve/compare exemplars.

The second finding that reinforces the link between construction learning and categorization is the reduced activation of left occipital cortex in the patterned condition of the appearance construction. Such occipital deactivation has also been found in non-linguistic categorization of visually presented stimuli (Chao, Weisberg, & Martin, 2002; Koenig et al. 2005). More generally, the attenuation of neural responses in occipital cortex often reflects the

expression of implicit memory or priming for repeated experiences (Smith & Grossman, 2008; Wiggs & Martin, 1998). In particular, when a stimulus is similar to another recently experienced stimulus, regions of visual cortex that code for the features shared across the two stimuli respond less strongly than when the stimulus is novel (Grill-Spector, Henson, & Martin, 2006; Turk-Browne, Scholl, & Chun, 2008). The reduced occipital response may thus reflect the fact that the videos were processed as more similar to each other when accompanied by consistently patterned narrations, as compared to when the same videos were presented in the same order with randomized narrations. Together, these two results suggest that construction learning may rely in part on a general similarity-based categorization mechanism.

Prediction. Our results are also consistent with a role for prediction in the language learning process. This is supported by the fact that better discrimination accuracy at test correlated with less activity in the ventral striatum during exposure. This result can be interpreted in terms of prediction error, whereby discrepancies between predicted and actual outcomes elicit responses in the ventral striatum (Niv & Montague, 2008; O'Doherty et al., 2004; Pagnoni, Zink, Montague, & Berns, 2002). In our study, the narration in the patterned condition allowed participants to anticipate how the scene would unfold. Thus, prediction errors were less likely than in the random condition, in which there was no systematic relationship between the video and linguistic input. The idea that reduced VS activity was the result of increased predictability relative to the random condition is consistent with the fact that the effect was stronger for the appearance construction. Recall that scenes of appearance varied quite a bit in their visual details, and more than scenes of approach. Understanding the construction would allow participants to reduce some of this uncertainty by predicting that the first entity mentioned was to

appear on the second entity. In contrast, all approach scenes began with two participants on opposite sides of the visual frame, followed by one approaching the other. In this case, understanding the construction might not have provided as much of a predictive advantage.

If this interpretation is correct, our study expands the recognized role of prediction and negative feedback in language learning. For instance, prediction is understood to be involved in artificial grammar learning: When learning an ABA pattern, hearing *wo-fe* affords a prediction that the next syllable will be *wo* and not *fe* (Marcus, 2000; Seidenberg, Elman, & Negishi, 1999). Prediction may also play a key role in learning constraints on language, such as learning to use *feet* instead of *foots*. Speakers regularly predict or anticipate what others will say: If a learner repeatedly predicts *foots* and systematically hears *feet* instead, prediction error can be used to learn to correct this mistake (e.g., Aronoff, 1976). Finally, prediction is central to computational models of language learning (e.g., Alishahi & Stevenson, 2008; Elman, 1993; Pearl & Sprouse, forthcoming) and theories of language comprehension (Pickering & Garrod, forthcoming).

Spatial Learning. The involvement of the left SMG may suggest an interesting connection between construction learning and the learning of spatial relations. Recall that the left SMG was implicated in two independent analyses: it was less active in patterned than random conditions, and the degree to which its activity decreased was a significant predictor of behavioral accuracy at test. These effects are unlikely to result from purely visual or auditory processes, since these features were identical across conditions. Thus, this effect must be multimodal or even conceptual. The left SMG has been strongly linked to the conceptual representation of space, in both linguistic and non-linguistic domains. In particular, it is associated with spatial tasks (Mountcastle, Lynch, Georgopoulos, Sakata, & Acuna, 1975),

including tool use (Vingerhouts, 2008) and skilled action (Sirigu et al., 2004), as well as language (Noordzij, Neggers, Ramsey, & Postma, 2008; Tranel & Kemmerer, 2004). For example, the SMG may be involved in the processing of linguistic prepositions (Noordzij et al., 2008; Kremmer et al., 2008): the SMG is more engaged during the comprehension of phrases with a preposition (*triangle to the right of the circle*) than phrases without (*a triangle and a circle*). Furthermore, in addition to apraxia (Heilman & Valenstien, 2003), patients with damage to the left SMG have impaired knowledge of locative prepositions (Tranel & Kemmerer, 2004) and even impaired recognition of spatially rich metaphors, such as *He hit the nail on the head* (McGeoch, Brang, & Ramachandran, 2007). Together, this work suggests that the SMG may link a symbolic representation (i.e., linguistic form) with visuospatial information.

In our study, the attenuated response of the left SMG in the patterned condition may reflect primed processing of constructions with particular visuo-spatial meaning. That is, participants learned to associate each abstract form with a specific spatial relationship (i.e., approach towards; appearance in or on). The degree to which such representations are primed may be a significant predictor of behavior because the meanings are spatially rich. These results therefore extend previous findings (Noordzij et. al, 2008), implicating the left SMG in more than the processing of prepositions per se. Rather, the left SMG may be involved in the learning of ‘prepositional’ meaning — learning between any linguistic form and spatial meaning.

Conclusion

There is growing evidence that speakers’ knowledge of language includes the knowledge of systematic phrasal form-function pairings (constructions), and yet little is known about the

neural mechanisms involved in learning these pairings. We have identified several processes that might be involved in construction learning, including similarity-based categorization, error-driven learning, and spatial learning. The lack of noticeable language-network involvement further supports a role for domain-general processes like these. Clearly, more research is needed to clarify the respective roles of different processes in construction learning. The present study suggests that neural systems typically linked to other forms of learning and memory will be an important starting point.

Experiment 3:

Construction Learning and Prediction

In experiment 3, we sought to further explore the role of prediction and categorization in the construction learning process. Overall, the findings in experiment 2 suggest that construction learning is in part subserved by a prediction based, domain general learning mechanism. Two key findings from experiment 2 point to this conclusion: the inverse relationship between ventral striatal activation and task performance in patterned trials only, and the relative de-activation of occipital cortex during appearance learning trials. While such findings are intriguing, both effects are present only in the ‘appearance’ construction, by hypothesis because this construction was much more visually variable, and less visually predictable than the ‘approach’ construction.

One of the primary goals of experiment 3 is to attempt to replicate, and further explore these findings from experiment 2. First, we attempt to replicate the findings with a new “departure” construction, which was designed to have the features we hypothesize to be driving the two neural responses discussed above. Specifically, we designed the construction to be unpredictable on the basis of visual input alone, yet predictable with the addition of the auditory narration. In this way, the construction’s meaning, if learned, would provide a useful predictive cue for the displayed action.

A further goal of experiment 3 was attempting to compare construction learning with word learning. To this end, we also included a word learning condition. Here, we exposed participants to instances of a novel lexical category that they could learn and extend to new items

at test. Unlike the construction condition, which involved dynamic videos, the word learning category involved auditory labels over static images.

Whereas the emphasis in experiment 2 was on the initial learning of these correspondences, the focus of experiment 3 was on the process of evaluating new exemplars. The testing procedure for experiment 3 was altered in order to compare neural responses to video exemplars that were either congruent (fulfilled expectation) or incongruent (violated expectation) with the form of the construction. Overall, the exposure condition in experiment 3 served to teach a correspondence that can predict a visual action, and the testing phase systematically probed the neural responses to a violation of this expectation (as compared to a fulfillment). If construction learning involves gradual prediction based learning (e.g. accomplished through resolving conflicts between expected and revealed events), then we would expect to find a similar profile when it comes time to evaluate new exemplars. That is, we would expect to find a similar profile of activation in the testing phase of experiment 3, as in the learning phase of experiment 2.

Methodology

Participants. Sixteen, right-handed native English speakers (age 18-25) participated in this study in return for \$30 in compensation. The study protocol was approved by Princeton University's IRB, and informed consent was provided by all participants.


Stimuli. Experiment 3 consisted of an exposure phase, and a testing phase.

Exposure Phase

There were two types of runs during the exposure phase: Word learning runs which employed 12 image/narration pairings, and construction learning runs, which consisted of 12 exemplars of video/narration pairings. Each participant viewed both run types twice (2 word learning runs, 2 Departure Construction runs) in counterbalanced ordering.

Departure Construction. For the ‘departure’ construction, all of the videos involved a ‘gorilla’ and a ‘frog’, which were played by two actors in realistic animal costumes. In each video, the characters were initially either standing or sitting side by side. Each video’s narration took the form “NP₁ NP₂ nonsense-verb”. After the narration, the argument designated by NP₂ leaves or moves away from the argument designated by NP₁ (see figure 11). The manner of motion varied from video to video: e.g. the frog, for example, might push himself away on a chair, walk away through a door, or magically disappear from the scene completely. For example, one narration was, “The gorilla the frog jippoes,” and the frog would walk up stairs, away from the gorilla. All narrations were created using the same male voice.

Figure 11: Example run from the “departure” construction: The characters are together for 2 seconds, after which the narration plays (e.g. *The frog the gorilla vakoos*). Then at the 5 second mark the gorilla moves away from the frog.




The Departure Construction was explicitly designed in order to allow for the constructional meaning (if learned) to provide a predictive cue for the visual scene to come. In all examples, both characters appeared side by side, then after 4 seconds the narration would sound, and 1 second after that, the 2nd NP argument would move away from the 1st NP argument. There were 12 runs, which were displayed to participants in a randomized order.

Word Learning. In the word learning trials, participants were exposed to static images of various objects, each paired with a linguistic narration. There were two categories of objects, which could be defined by a single attribute. One category consisted of objects that are sharp (e.g. a porcupine’s quill, a broken wine bottle, a pizza cutter etc.), and were paired with the

nonse label *jorp*. The other category was of objects with a clear line of vertical symmetry (e.g. a fish from bird's eye view, an inkblot, a road), and were all given the narration *blick* (see figure 12).

Figure 12: The two stimuli types for the word learning task: On left, “Blicks”: objects with a clear line of vertical symmetry; on right, “Jorps”: objects which are sharp. Each static image was viewed for 10 seconds, with the narration (*this is a blick/jorp*), occurring at the 2 second mark.



As with the stimuli for the Departure Construction, there were 12 exemplars in total (6 ‘blicks’ and 6 ‘jorps’), and the order of presentation was randomized for each participant. Each participant saw 2 word learning runs, and 2 construction learning runs, with the order

counterbalanced across participants, with the constraint that two of the same type of run did not occur back to back.

Testing Phase

The testing phase consisted of two testing blocks: one for the Departure Construction, and one for Word Categories.

The general set up was similar for both test types: participants would see an image (in the case of the word test) or a video exemplar, accompanied by a narration. Participants witnessed 16 exemplars in each testing trial: For exactly half (8) of the exemplars, the image or video was narrated *correctly* (consistent with the learning phase), and the other half (8) of exemplars, it was be narrated *incorrectly* (in violation of the pattern seen during learning).

Participants witnessed each exemplar in its entirety, and then were asked with a prompt on the screen “Was the video/image narrated correctly?” Participants indicated their choice with a button press, and had 10 seconds to do so. After 10 seconds, the prompt would disappear and the next trial would commence, whether or not a response had been registered or not. There were 16 exemplars in each testing trial, and the order of presentation was randomized for each participant.

Departure Construction Test. The testing trials for the construction were new instances that had not been witnessed during the exposure phase. Each involved new characters and new manners of motions and corresponding new novel verbs. Testing scenes consisted of a penguin and chicken, or a dog and a bear. New manners of motion included “transporting” to another

location in the scene, climbing stairs away from the other character, and walking away in order to sit down on a chair. These exemplars were systematically different from exposure items in order to gauge how well participants were able to learn the abstract form/meaning correspondence, as opposed to stimulus specific associations. That is, the test trials tested how well the participants could *generalize* to these new exemplars.

On correct trials, NP₂ would leave or move away from NP₁ (consistent with the pattern witnessed during exposure). On incorrect trials, NP₁ would leave or move away from NP₂ (inconsistent with the pattern seen in exposure). The exemplars were similar to the exposure items, in that the narration would be presented 4 seconds in, and only afterwards would the scene's action take place. Likewise, test items were 10 seconds long in total.

Word Test. The word test stimuli were also different items than those seen during the exposure phase. For example, new sharp items included a crocodile tooth, a shaving razor etc., while new symmetrical items included a leaf, and an 8 ball. Trials were also 10 seconds long, with the narration occurring 4 seconds after the image was first shown.

Organization of testing versions. For both the word and construction tests, there existed two versions of test items. Each contained the same visual stimuli, but differed in which videos were narrated correctly and which were narrated incorrectly. In version 1, videos 1-8 are narrated correctly, while videos 9-16 are narrated incorrectly; version 2 was the reverse: videos 1-8 are narrated incorrectly and videos 9-16 are narrated correctly. Thus, for each test exemplar, there exists a correctly narrated version and an incorrectly version, and within any given version,

exactly half of the items were narrated correctly. The order of videos was randomized for each participant.

Each participant saw only one version of the word test, and one version of the construction test, and this distribution was counterbalanced across participants. Thus, each version of each test type was seen by 8 participants and overall, with each exemplar being viewed 8 times with a correct narration, and 8 times viewed with an incorrect narration.

fMRI Analysis

Data acquisition. Scanning was performed on a Siemens 3T Allegra head-only scanner using a Nova Medical NM-011 Head Transmit Coil with receive-only array system. Functional sequences were acquired using a T2*-weighted EPI sequence (echo time = 30 ms, repetition time = 2000 ms, flip angle = 64°, matrix = 64 x 64). Twenty-six interleaved oblique axial slices were aligned parallel to the anterior-commissure/posterior-commissure, and whole-brain coverage was obtained at a resolution of 3.5 x 3.5 x 3.5 mm. For alignment, two T1-weighted structural images were acquired: a coplanar FLASH image and a high-resolution MPAGE image.

Preprocessing. Data were preprocessed using FSL's FEAT (Woolrich et al., 2009; Smith et al., 2004; www.fmrib.ox.ac.uk/fsl), including: corrections for head motion and slice acquisition time, high-pass temporal filtering with a 128 s period cut-off, and spatial smoothing with a 5 mm kernel. Functional runs were registered to the coplanar (FLASH) structural image, and then to the high-resolution (MPAGE) structural image. Data were normalized to Montreal Neurological Institute (MNI) space and interpolated to 2 mm isotropic voxels.

Congruent vs. Incongruent. Within each testing run, half of the video exemplars correctly pair the audio and image, and the other half pair them incorrectly—that is, in a way that is inconsistent with the exposure. Each testing run was fit using a general linear model (GLM) that with two regressors: one which modeled the average blood-oxygen level-dependent (BOLD) for the 8 congruent videos, and the other for the average BOLD response of the 8 incongruent videos. Each video was modeled with a boxcar function convolved with a canonical hemodynamic response function. Since the videos were of variable lengths and occurred in a random order across participants, unique regressors were created for each run of each participant.

To compare across conditions, we contrasted parameter estimates obtained in the GLM of each run in a higher-level GLM, treating subject as a random effect in a non-paired t-test.

Specifically, we contrasted correct vs. incorrect separately for construction and also collapsing across constructions. In each analysis, all correct videos (across both testing versions) were compared to all incorrect videos (across both testing conditions). A pre-threshold mask was generated by averaging the overall patterned and random responses from each participant together. The resulting voxelwise statistical maps were thresholded at $p < 0.05$ corrected, using a cluster-mass correction and an initial Z threshold of 2.3.

Behavioral correlations. In order to examine brain-behavior relationships, we used individual test performance as a regressor for BOLD responses during the exposure runs. First, we z-scored across participants the behavioral scores for overall performance, as well as for each construction individually. A new higher-level GLM was then computed, where the parameter estimate for each participant and condition was weighted based on the behavioral score.

Essentially, this analysis produced voxel clusters in which there was a significant relationship between evoked BOLD responses and test performance (same threshold as above).


Averages for construction learning trials were created by collapsing across the two construction learning runs, and a similar procedure was done to create averages for the word learning trials. These behavioral relationships were examined for the average construction learning trials, as well as for the average word learning runs.

Time course analyses. Within individual participants, we weighted the predicted BOLD response for each block based on the block number in the run: -5.5, -4.5, -3.5, -2.5, -1.5, -0.5, 0.5, 1.5, 2.5, 3.5, 4.5, 5.5. These weights summed to zero, and thus the parametric regressor only captured the linear change in response over the run (no main effects). This analysis was applied to construction learning blocks, and word learning runs separately. To compare these linear effects across blocks, we computed a higher-level GLM separately for each construction and collapsed across constructions. In addition, we probed the relationship between these linear changes and behavior, using test performance as a covariate for the higher-level GLM, as described above.

Results

Behavioral Results. Participants' scores at test were in the same range as in experiment 2? On the construction test items, the average score was 76% (SD = 22.9). Participants averaged 78% (SD = 21) on the word learning test. As determined by t-tests against 50%, performance on both the tests were significantly above chance: ($t(15) = 4.85, p < .001$) and ($t(15) = 5.14, p < .001$), and not significantly different from each other ($t(31) = .35; p = .72$). Each participant's score is represented in Figure 12.


Figure 12: Participants' individual performance on construction learning (blue) and word learning (red). Chance is 50%.



fMRI Results

Congruent vs. Incongruent. We first contrasted both the word tests and construction tests together. No regions showed significance in either the congruent > incongruent or incongruent > congruent contrasts. We also examined these contrasts for the test trials of word and construction tests separately. In the construction analysis, we find significantly greater activity in incorrect over correct trials in the left occipital pole ($X = 2, Y = -94, Z = 8$; see figure 13). No areas are significantly more active in the congruent > incongruent contrast. In the word condition, we do not see any significant differences in either the congruent > incongruent or incongruent > congruent contrasts.

Figure 13: Left Occipital Cortex, less active in congruent testing trials than in incongruent testing trials.




Left Occipital Cortex

Correct Hit v Correct Rejections. In order to directly probe the processes of category assimilation and category rejection, we can further divide our testing trials into ‘correct hits’ (congruently narrated trials that were correctly recognized as congruent) and ‘correct rejections’ (incongruently narrated trials that were correctly recognized as incongruent). Essentially this analysis takes all of the accurate responses from the Congruent v. Incongruent trials. Not surprisingly then, we find similar results to the Congruent v. Incongruent condition: more activity in Left Occipital cortex ($X = 2$, $Y = -94$, $Z = 8$) in response to correct v. incorrect trials.

ROI Analysis: fROI from Appearance Contrast in Experiment 2. In addition to the whole brain analysis, we also conducted an ROI analysis using the L Occipital Deactivation from random appearance > patterned appearance condition in experiment 2. Here we find significant de-activation (Rejections > Hit) in left occipital cortex ($x = -38$, $y = -88$, $z = 6$; see figure 14). Using a probabilistic atlas based on previous functional topographies (Wang et al, in prep), this activation is most likely V2 or V3. There are no regions which approach significance in the Hit > Rejections contrast.

Figure 14: Left V2; this region represents the area which shows a pattern of relative de-activation in both studies.




Time Course Analysis. Experiment 2 also provided an opportunity to track the evolution of this learning mechanism over time. As in experiment 2, we fit a parametric model in which we predicted that the BOLD response would change linearly as a function of the number of exemplars that had been witnessed. We reasoned that the majority of learning takes place during the first exposure run, so for this reason this analysis was only done on the first functional run of word exposure and the first run of construction exposure. For both the construction and word analysis, we find similar patterns of activation as we find for the TCA analysis in experiment 2. For the construction time course analysis, we see increased activation in the posterior precuneus and posterior callosum, while the word construction has a more anterior pattern of activation, encompassing broader areas of the callosum.

Behavioral Covariate Analysis. There was significant variation in how well participants performed at test, which allowed us to use learning as a covariate for neural activation both while learning is taking place, and while it is being tested.

We regressed behavioral performance with neural activity from the first exposure run for both the word, and construction conditions. For the word condition, no regions showed a significant correlation with performance, but negative activity in Broca's Area ($x = -44, y = 16, z = 14$) during learning was predictive of performance (see figure 15A). For the construction condition, better performance was associated with activity in Wernicke's Area/Heschel's gyrus ($x = -38, y = -34, z = 14$) during learning the first exposure run (see figure 15B). There were no regions that displayed a significant inverse correlation.


Figure 15: Correlations with performance and activity in the first run of exposure. In the word condition (A), we see an inverse relationship with Broca's Area and performance, while in the Departure construction we see a positive correlation between activity in Wernicke's Area and performance.



In experiment 2, we saw an inverse relationship between performance at test, and activity in ventral striatum during learning – an effect largely driven by the visually varied appearance construction. In experiment 3, however, we did not see evidence of such a relationship in the whole brain, behavioral covariate analyses. In order to take a more direct approach, we used an ROI of the ventral striatum, and collected the overall activity here from each subject. We then correlated this with performance by each participant. In this analysis, an indication of the involvement of ventral striatum is evident

We correlated behavioral performance on the departure construction test with activity from the first learning run of the departure construction and find a moderate ($r^2 = -.23$, $p = .11$), inverse correlation (see figure 16).


Figure 16: Inverse correlation between performance on the construction test and activity in ventral striatum during learning.



We also correlated performance with the average activity in ventral striatum separately for both congruent and incongruent trials at test. Here, we find a negative correlation ($r^2 = -.47$,

$p = .038$) with congruent trials, but an positive correlation ($r^2 = .46$, $p = < .04$) with incongruent trials. That is, more activity in VS correlates with more accurate recognition of congruent trials, and less activity correlates with more accurate recognition of congruent trials (see figure 17).

Figure 17. The relationships between ventral striatal activity during departure testing trials (congruent and incongruent), and test performance.



Critically, when we run the same analyses in the word condition, we do not see any significant relationships emerge, either for first run of exposure ($r^2 = -.01$, $p > .05$), congruent testing trials ($r^2 = .18$, $p > .05$), or incongruent testing trials ($r^2 = .16$, $p > .05$).

Discussion

Summary of findings. This study brings new data into our understanding of construction learning, and helps shed light on the processes that are involved in this process. Two key findings from experiment 2 find support in this follow up study. First, we find similar occipital de-activation during runs of the departure construction (at test) in experiment 3 and during runs of the appearance construction (during exposure) in experiment 2. And secondly, while we do not see the same inverse correlation with performance in ventral striatum evident in experiment 2 with the whole brain analysis, a more targeted ROI analysis does reveal a related effect : activity in the ventral striatum is positively correlated with performance on incongruent trials and negatively correlated with performance on congruent trials.

We also see new findings emerge as well. In our behavioral covariate analysis, we see in the Departure construction that behavioral performance is significantly correlated with activity in the left superior temporal gyrus (Wernicke's Area) during learning, while in the word construction, activity in the LIFG (Broca's Area) during learning is inversely correlated with performance.

These results, in conjunction with the results from the first fMRI study, will be discussed in terms of three primary themes: Prediction and visual processing, violation of expectation, and the role of traditional language areas.

Visual Suppression. One consistent finding was some neural suppression in the visual area of the cortex when participants witnessed congruent trials as compared with incongruent trials. In experiment 2, this was evident during the exposure phase, in the random appearance > patterned appearance contrast. In study 3, this was found in examining the incongruent >

congruent trials in the departure construction. Interestingly, this effect is absent in the approach construction (study 2), as well as in the word condition (see table 10).

Table 10: Summary of all the findings that implement a de-activation of occipital cortex across the two studies, broken down by condition.

	Condition	Comparison	Finding	Coordinates
Study 3	Departure	Incongruent > congruent	Left Occipital Cortex	X = 2, Y = -94, Z = 8
	Word	-	<i>None in any contrast</i>	-
Study 2	Appearance	Random > Patterned	Left Occipital Cortex	X = -22, Y = -84, Z = 8 X = 8, Y = -36, Z = 8
	Approach	-	<i>None in any contrast</i>	-

Despite the visual information being identical across patterned and random presentations, there is a trend towards less evoked activity in the congruent condition. This pattern of activation has also been found in two mutually supportive processes: visual categorization (Reber 1998, 2000; Koenig 2005), and predicted visual stimuli (Kok et al 2011; Kok, De Lange, and Jehee 2010; Summerfield and Egnér 2009).

One paradigm commonly used in fMRI studies of visual categorization tasks is an adaptation of the Posner and Keele (1968) prototype extraction task. Here, participants are trained to categorize two distinct dot patterns, which each involve distortions from a prototype. Afterwards, participants are presented with 3 sets of stimuli, randomly presented: old category members previously seen, new category members, and random dot patterns (non-category members). On each trial, they are asked whether they have seen that pattern, or not. The fMRI analysis of interest is the contrast between new category members and random dot patterns. Despite the fact that neither pattern has been previously seen before, there is significantly less

visual activation in V1 and V2 in response to category members than for non-category members (Reber et al. 1998). This effect has also been shown to hold (albeit less dramatically) when instructions at test are to explicitly judge the category membership (and not the newness) of each exemplar pattern (Reber et al. 2000).


Similar findings have also been found using more abstract categories such as animals (Chao, Weisberg and Martin, 2003; Koenig et al. 2007). It is thought that such de-activation may in fact represent the process of learning itself, and not simply an alleviation of sensory processing. In their review of memory systems in the brain, Smith and Grossman (2008) state “Substantial research indicates that, relative to an appropriate baseline... de-activation of extrastriate cortex is a marker of implicit memory” (261). This is most often found, but not limited to, investigations of perceptual priming (see Buckner 2000 for a review), where a visually presented item can lead to the implicit and facilitated retrieval of related items.

It is important to note the conditions in which this effect has been found in the present experiments: only the departure construction and in the appearance construction. In these conditions (but not the word or approach construction), the linguistic construction, if learned, provides a predictive cue for the upcoming visual action. Thus, the comparison between predictability and non-predictability may be needed to see this de-activation effect.

Indeed, recent work by Kok, Jehee, and De Lange (2012) suggests that such predictability may be the crux of such lower level de-activation. While undergoing fMRI, healthy subjects viewed patterns of parallel lines (“grating”), and were asked, on separate trials, to respond whether the second grating witnessed was clockwise or counterclockwise to the first; or whether the second grating was of higher or lower contrast than the first. Each image was preceded by an

auditory tone (either high or low), which served as a probabilistic (75%) cue for the orientation of the grating pattern (either 45° or 135°) (see figure 18).

Figure 18: Example trial, borrowed from Kok, Jehee, and De Lange (2012). An auditory cue precedes the grating stimuli, and probabilistically predicts either its general orientation (either ~45° or ~135°).



The fMRI results showed that activation in response to expected stimuli was reduced in V1, V2 and V3 compared to unexpected stimuli. Furthermore, using MVPA, it was demonstrated that such a de-activation was associated with orientation selectivity in V1, and that this orientation classification was positively correlated with both behavioral performance and reduced activity in V1. Overall, the findings suggest that expected visual stimuli receive less low level processing than unexpected stimuli, and that this de-activation may sharpen the neural systems responsible for expected representations.

We interpret our findings as reflecting a similar expectation - driven de-activation. If this is the case, the present studies extend the current understanding of this phenomenon in an interesting way. Results indicate that occipital de-activation occurs even when the prediction

about the visual stimuli is quite modest and abstract. In the case of the Departure construction, participants could only predict that the departure would occur either on the left or right half of the screen, and in the case of the Appearance construction, participants could only predict that one entity would appear in or on another named entity. The direction and manner of departure or appearance were novel for each stimuli. That is, the retinal images could not themselves be predicted; the prediction involved in both cases was quite a bit more abstract.

Prediction. Our results provide evidence for the important role of prediction in this learning process, both in the occipital deactivation and in the correlations evident in the ventral striatum. In study 2, whole brain analyses during learning reveal an inverse relationship with discrimination performance in VS. In study 3, a more targeted ROI analysis of VS likewise yielded a significant inverse correlation between performance and activity in congruent test items, and a positive correlation between performance and incongruent items. That is, performance is better when ventral striatal activity is higher for incongruent items and lower for congruent items.

As in study 2, these results may be best interpreted in light of the ventral striatum's role in prediction error and learning. As we discussed in the discussion of experiment 3, the ventral striatum has been shown to encode reference-dependent reward information (e.g. Schultz et al. 1997; Niv and Montague 2008). While ventral striatum has been most classically linked to reward prediction, there is evidence that the dopaminergic response may also be involved in the processing of salient stimuli (Ungless 2004), and the more general prediction for information (Bromberg-Martin and Hikosaka 2009; Redgrave and Gurney 2006). For example, Bromberg-Martin and Hikosaka (2009) monitored dopaminergic responses in macaques while they

performed a simple decision task - choosing whether or not to receive information about an upcoming water reward. Despite the fact that the monkeys were aware that the size of the reward was independent of their choice, participants consistently sought this additional information. Consistent with previous investigations, there was a dopaminergic response to cues which predicted water rewards. More interestingly, there was a large increase in the dopaminergic response to trials offering a choice for information as compared to cues that merely predicted the amount to come. In this view, dopaminergic responses may be monitoring a more general predictive response involved in learning, and most optimally tuned for the learning of cues that predict rewarding stimuli.

Another, mutually supportive interpretation of these findings, is that the responses in ventral striatum also represent the intrinsic reward of learning. Although these findings are most often linked with paradigms involving the prediction and anticipation of extrinsic rewards, it is possible that the ventral striatum may also be involved in the monitoring of intrinsic rewards as well (e.g. Botvinick, Huffstetler and McGuire 2009). From this perspective, we see greater prediction error correlated with performance in incongruent items because there is a greater violation of this reward expectation. Similarly, this would also explain why participants who learned the construction well during exposure would have less of a prediction error in response to congruent items.

The current study does not disambiguate between these two interpretations: whether dopaminergic activity reflects the prediction of intrinsic rewards, the prediction of more general non-rewarding stimuli, or both. Overall, these findings do provide support for the idea that the striatum's response may aid in the learning of more general types of learning outside of classical, external reward driven learning.

Traditional Language Areas. It is difficult to know what to make of the contribution of traditional language areas (e.g. the left frontal-temporal network) to construction learning. There was scant evidence for their involvement in the first investigation, while in the second, we do see evidence for their engagement, albeit in a somewhat ambiguous way: while activity in Wernicke's Area during Departure learning was positively correlated with learning outcome, activity in Broca's Area was negatively correlated with learning outcome.

It is somewhat surprising that L Superior Temporal regions would relate to abstract phrasal meaning, while word learning would be linked with Broca's Area. Given many traditional accounts of these regions (Gernsbacher and Kaschak 2003), one might expect the opposite association, as Broca's Area has typically been associated with 'phrases' and sentence level meaning while Wernicke's Area has more often been associated with semantics and lexical meaning (Osterhour, Kim, and Kuperberg 2006).

By itself, the fact that the engagement of Wernicke's area is associated with positive performance is not unexpected. Wernicke's area, as well as neighboring regions of the left superior temporal lobe are strongly associated with the learning and processing of semantic content (Tzourio et al. 1998, Robertson et al. 2000). In fact, the left temporal lobe's well documented roles in sentence level processing (e.g. Catani et al. 2005), and semantic combination (Baron and Osherson 2011; Baron et al. 2010) make it a likely candidate for providing key contributions to construction learning. More unclear however, is the inverse relationship between Broca's activity and word learning performance. While there is precedence for this region to be involved in many aspects of semantics, and word-level semantics specifically (Demb et al. 1995; Desmond et al. 1995; Frost et al. 1999), it is unclear why the

absence of activity here would be associated with learning. One would expect to see more activity, not less, associated with performance. However, there is precedence in other types of learning where less activity is associated with better proficiency and learning. This type of relationship has been found, for example, in motor learning (Jenkins et al. 1994), associative learning (Zeineh, Engel, and Thompson 2003), and artificial grammar learning (Seger et al. 2000). In fact, exceptional abilities have been linked to brain damage to specific brain regions (Caron et al. 2006), and have been shown to be facilitated by artificial suppression of neural activity using TMS (Barhamali et al. 2006; Oliveri, and Papagno 2004; Snyder et al. 2003).

It is difficult to say whether the inverse response is related to this sort of phenomena or not, but if it does, it would beg the question why this type of relationship is not present for Wernicke's Area in the Departure construction. Ultimately, these findings are inconclusive with regards to the involvement of traditional, language relevant regions in construction learning.

Chapter 5: Construction Learning in Children with Autism

Experiment 4, adapted from:

Johnson, M., Boyd, J., and Goldberg, A. (2012). Language generalization in children with autism. Conference proceedings of the Boston University Child Language Conference 2011

A better knowledge for how abstract constructions are learned may also provide insight into certain language difficulties as well. The language difficulties of children with autism may be one such example. In Johnson, Goldberg, and Boyd, (2012), we explore the possibility that autistic language delay may be due, at least in part, to a deficiency in the ability generalize over linguistic exemplars. Children with autism tend to have an unusually narrow focus of attention, within which their focus may be supernormal (Allen and Courschene, 2001). Relatedly, autistic individuals appear to have a locally biased processing style which arguably comes at the expense of processing more holistic aspects of information (Frith 1989).

For example, children with autism have demonstrated a reduced ability to sort by gestalt principles (Brosnan, Scott, Fox and Pye, 2004), and display superior performances on embedded figure tasks, which require a neglect of contextual feature (Shah and Frith 1983; Jarrold et al 2000; Morgan et al, 2003).

Thus, children with autism have been noted to “neglect the forest for the trees” showing a strong focus and memory for the details of a given stimuli while showing deficits in the perception of holistic aspects. There is accumulating evidence that cuts across various domains, as it has been demonstrated in vision (Frith and Dakin 2005), audition (DeGelder, Vroomen, &

Van der Heide, 1991), and verbal memory (Hermelin and O’Conner, 1967; Happe 1997). This processing style has been contextualized in the weak central coherence model of autism (Happe, 1999; Happe and Frith 2006), where autistics have been described as having difficulty “drawing together diverse information to construct higher-level meaning in context” (Frith and Happe, 1994, p. 117).

It has also been suggested that this locally over-focused attentional style may be closely related to noted difficulties in abstraction (Minschew, Goldstein, Muenz, and Payton, 1992, Klinger and Dawson, 2001). In a series of studies, Plaisted et al.(1998) demonstrated that children with autism were unable to exploit the commonalities between the training and test phases of a perceptual dot pattern learning task, while typically developing children could; i.e., they successfully identified dot patterns that had been witnessed, but they were unable to extend the category to novel patterns that were systematically similar to those witnessed during training. In essence, they treated similar dot patterns as being completely novel. Observations such as these have led some researchers to make specific claims about the autistic profile. “Children with autism differ in their use of similarity: specifically..they represent individual stimuli with very steep generalization gradients and do not perceive stimuli as similar unless they are very close in stimulus space” (Molesworth, Bowler, & Hampton, 2005; p. 663). How this inability relates to the attentional profile is uncertain (a specific deficit in prototype formation or otherwise), but it does appear that children with autism show a specific deficit in the ability to generalize and make inferences beyond the specifics of their previous experiences.

It is difficult to gauge how difficulties in abstraction influence categorization and concept formation, and the research into this relationship has yielded mixed results. While it appears that children with autism show an ability to form categories in certain contexts (Ungerer and Sigman,

1987, Tager-Flusberg, 1985a, 1985b), this proficiency may be limited to categorizations based on one dimensional, perceptual stimuli. For example, autistic individuals show difficulties when sorting on the basis of more than one perceptual criteria simultaneously (i. e. red *and* round) (Rutherford and McCintosh, 2007). Autistic individuals also fare poorly when categorization is based on more abstract features (i.e “things that can fly”, “things that give light”) (Soulieres et al, 2007). However they have demonstrated the ability to successfully sort objects that are animate from those that are inanimate (Baron-Cohen, 1991) exhibiting *some* facility in abstraction.

It has been proposed that, due to their over-focused attention and inability to integrate different attributes across exemplars, children with autism do not appear to form prototypes, but instead, form categories based on rule based strategies (Klinger and Dawson, 1995, Klinger and Dawson, 2001, Rutherford and McCintosh, 2006). Although it appears that individuals with autism do rely more on rule based strategies, there is mixed evidence in favor of a specific deficiency in prototype formation. There is some support for the idea that difficulties in prototype formation may be due to methodological confounds in task difficulty (Molesworth et al. 2008), and it is also possible that any problems with prototype based categories may only be characteristic of low, but not high functioning autistic individuals (Molesworth, Bowlwe & Hampton 2005, Molesworth et al. 2008).

Thus, deficiencies in generalization are very well researched and oft-noted feature of autism. And while direct measures of cognitive style are not currently part of the diagnostic criteria, cognitive features are arguably the most robust and commonly noted features of the disorder (Frith and Happe, 2006) outside of the defining criteria. This is especially note-worthy considering the heterogeneity of the autistic profile, and the inconsistency found in noting other

non-diagnostic features (Moreno, Borjas, Arrieta, Saez, Prasad, Estevez and Bonilla 1992; Dawson, Mottron and Gernsbacher 2008; Happe, Roland and Plomin 2006).

We hypothesize that this generalization deficiency may have important consequences for language learning. Despite the crucial role that abstraction plays in language development, as outlined in Chapter 1, no investigations to date have examined how it may play a role in the language delay of autistics. The present paradigm addresses this issue head on, evaluating how children with autism learn and generalize from the language they are exposed to.

Methodology

A computer-based language learning task was used, which had been employed in past experiments with adults and typically developing children of varying ages (Casenhiser & Goldberg, 2005; Boyd, Gottschalk and Goldberg, 2009; Boyd and Goldberg, 2011; Wonnacott et al. 2012). Its computer-based nature is an important aspect of the design, as it reduces the likelihood that the child's social profile unduly influences performance, and therefore provides an opportunity to investigate how deficiencies in categorization may directly influence the learning task.

Participants. Eighteen typically developing 7 year olds, and 16 autistic children, ages 8-13, participated in the task. Seven year olds were recruited from nearby summer afterschool programs, while the autistic group was recruited from a nearby middle school. All of the autistic children had received a formal diagnosis of autism from a pediatric professional. All parents gave written consent prior to their child's meeting with experimenters, and all children gave

verbal consent at the time of participation. Each child was given a children’s book in exchange for their participation.

Pre-Test. The program began with a series of six trials testing basic language comprehension abilities. Participants were shown two similar videos side by side, with one overlapping auditory description and were asked to point to which video that the sentence described. For example, in one video a frog puppet jumps on a king puppet, while in the other the king jumps on the frog. Subjects heard “the frog is jumping on the king”, so the correct response is a point to the first video. Pre-trials were included in order to ensure participants had a basic understanding of both the task and of simple language. As will be discussed in more detail later, it was also used in part, to ‘match’ the two groups. A summary of all six of the pretest trials is given in Table 10:

Table 10: Pretest items of basic language comprehension. Correct items in boldface.

Left Video	Right Video
The frog is jumping on the prince	The prince is jumping on the frog
The frog moves around the car	The frog is moving around the car
The king bows at the queen	The king and queen are bowing
The monster is rolling on the ball	The monster rolls the ball
The queen waves the flower	The queen waves at the flower
The witch waves at the man	The witch and the man wave

Exposure to the novel, ‘approach’ construction. The goal of this task was to test children on their ability to generalize at the level of the phrasal construction. To this end, a novel construction was used pairing a novel word order with a novel abstract meaning; the form was <NP1 NP2 V>, and the abstract meaning was “approach”. In particular, NP1 designated the actor who moved toward the actor designated by NP2. Five different nonsense verbs ending in “o” were used, each denoting a particular manners of motion. This construction does not exist in English and thus provides the opportunity to investigate learning while minimizing the influence of previous linguistic knowledge.

Participants were exposed to 16 brief videos, in which scenes of approach were depicted and narrated by the novel construction. All videos involved the same two characters, a doctor and a construction worker, and lasted approximately 8 seconds each. For example, subjects saw a doctor crawling to a construction worker and heard “the doctor the construction worker moopoed”. Half of the time NP1 designated the doctor and NP2, the construction worker, and the other half of the time it was the reverse. While the same two characters were used across all trials, different verbs, depicting different types of approach motion were used. For example, on another trial, instead of crawling, the doctor *hopped* to the construction worker – this time narrated by “The doctor the construction worker vakoos”. Five different verbs/approach actions were used across the 16 trials. In line with previous work demonstrating a learning advantage for an input skewed towards a single verb (Casenhiser and Goldberg 2005; Goldberg et al. 2004), *moopoes* was employed 8 of the 16 times, while *vako*, *keybo*, *feigo*, and *sooto* were used twice each. While different approach actions and different verbs were used across the 16 trials, all of the clips retained the same abstract form <NP1 NP2 V-o> and meaning (NP1 approached NP2 in a V-like fashion).

Testing Trials. In order to determine what participants learned during exposure, they were given a forced choice comprehension task that tested their knowledge of this new “approach” construction. All testing trials involved participants watching two videos play simultaneously, side-by-side, while being narrated by an overlapping description. As in the pre-test, participants were simply instructed to point to which video the sentence referred to. Participants were given 16 test items: 6 holistic approach trials, 6 holistic intransitive trials and four linking trials.

The purpose of the holistic trials was to simply determine whether participants could discriminate the meaning of the novel construction from that of a common English construction: the intransitive. All of these testing trials tested whether participants could distinguish a scene of approach from a scene depicting the agents performing a repetitive intransitive action in synchrony (i.g rubbing their stomachs). On approach trials, the narration had the form <NP1 NP2 V>, thus making the correct answer the approach video. In contrast, intransitive trials were narrated by an intransitive form with a novel verb (e.g. *the doctor and construction worker are koobing*), and thus the intransitive video in this instance was the correct choice.

Of crucial interest was whether participants were able to learn the *abstract* form/meaning pairing; i.e., whether the constructional knowledge would transfer to items that differed systematically from the original exemplars. In order to gauge this, these holistic trials differed along three lines of novelty: *old items*, *new verb*, and *all novel*. On ‘old item’ trials, the approach video was simply repeated from one seen in exposure. Participants simply needed to recognize this old exemplar and its corresponding narration. In the ‘new verb’ condition, the target video retained the same characters as in exposure (the doctor and the construction worker),

but this time a new verb and a corresponding new action were used. In the all novel trials, both new characters (e.g boxer and princess) and new verbs/actions were used.

Despite manipulating these surface differences, success is still possible on holistic trials given a non-comprehensive understanding of this new construction. It is possible for participants to do well on holistic trials knowing only that two initial NPs are associated with a meaning of approach. Linking trials were used in order to gauge whether participants understood that NP1 was associated with the approaching agent, while NP2 referred to the person being approached. Unlike the holistic trials, linking trials included two videos of approach scenes. In one, the construction worker approached the doctor, while in the other, the doctor approached the construction worker in the same manner. The voiceover was an approach form corresponding to one of the videos. Thus, if the description was, *The construction worker the doctor fagoes*, the video of the construction worker approaching the doctor (and not vice versa) should be selected as being correct. A summary of all testing trials can be found below in Table 2

Table 11: Summary of the different testing trials in experiment 3. Bold items represent aspects of testing items that differed from the original exemplars

Holistic Trials	Old Items(4)	Narration: "The doctor the construction worker moopoes" Correct Video: the doctor approaches the construction worker Foil Video: Intransitive Action
	Intransitives – New Verb (2)	Narration: "The doctor and the construction worker are lorping " Correct Video: Intransitive action Foil Video: the doctor approaching the construction worker
	Intransitives – All new (2)	Narration: "The boxer and the princess are jorping " Correct Video: Intransitive Action Foil Video: The boxer approaches the princess
	Approach – New Verb (2)	Narration: "The construction worker the doctor vakoes " Correct Video: the construction worker approaching the doctor Foil Video: Intransitive action
	Approach – New Verb & Character (2)	Narration: "The boxer the princess kafoes " Correct Video: the boxer approaches the princess Foil Video: Intransitive action
	Linking – New Verb (4)	Narration: "The doctor the construction worker satoes " Correct Video: the doctor approaches the construction worker Foil Video: the construction worker approaching the doctor

Both groups saw identical exposure and testing videos. The only difference between the two groups was the order in which the different testing trials appeared. In the typically developing group, participants were first tested on ‘no novelty’ items, then on ‘new verb’, ‘all

new’ and then finally on linking trials. In order to avoid possible fatigue effects, the autistic group witnessed ‘new verb’ trials *first*, followed by ‘all new,’ then linking trials, and then ‘no novelty’ items. We recognize this is a potential confound, but it should if anything work against our hypothesis, since it should if anything lead to better performance on the earlier trials and yet we predict that autistic children will perform better on the later trials than on the earlier ones. In a second follow up study, detailed below, we avoided this potential confound by exposing a new set of autistic children to the same construction learning stimuli, and keeping the test order consistent with the typical group.

Within this computer-based learning paradigm, it has been previously demonstrated that after just 3 minutes of training, typically developing children can learn a novel phrasal construction, and can generalize over the construction to novel linguistic input (Goldberg, Casenhiser and Sethuraman 2004, Casenhiser and Goldberg 2005, Boyd and Goldberg, 2012). However, we might expect that autistic children may not perform comparably in this task, given their well-noted difficulties in generalization. We hypothesize that children with autism will perform significantly below their typically developing counterparts on new items requiring generalization.

Results

Comparison of group performances. An overall main effect of accuracy was found, $F(1, 196) = 29.6$; $p < .001$, with typically developing children being more accurate in their responses across all trials ($\bar{x} = .89$, $\sigma = .23$) than the autistic group ($\bar{x} = .68$, $\sigma = .31$). More interesting, however, is how the different groups compared on the two types of test items - those

that do not rely on generalization (pre-test, intransitive, and no novelty items; “O” trials), and those which do rely on generalization (new verb, all new and linking trials; “G” trials).

We compared performance across group on for each trial type using between subject ANOVAs. First, we found no significant difference between autistic ($\bar{x} = .93$, $\sigma = .10$) and typical developing ($\bar{x} = .98$, $\sigma = .05$) children on pretest items: $F(1, 31) = 2.86$, $p = .1$. This was also true of no novelty items, $F(1, 31) = 2.16$, $p = .152$, and intransitive trials, $F(1, 31) = 2.4$, $p = .72$.

Group differences only emerge on the performance of “G” trials, which require generalization. Here, we find that autistic performance is significantly worse, on novel verb items $F(1, 31) = 8.47$, $p = .005$, all novel items $F(1,31) = 6.62$, $p = .015$ and linking rule items $F(1,31) = 18.06$, $p < .001$ (see figure 19).

Comparisons against chance performance. In order to evaluate whether participants demonstrated learning (e.g. performed significantly above chance), we conducted one-tailed t-tests against a chance baseline for each trial type. The typically developing group performed significantly above chance on both “O” trials (pre-test = $t(17) = 15.35$, $p < .001$; intransitives = $t(15) = 4.04$, $p < .001$; old items = $t(17) = 10.29$, $p < .001$), and “G” trials (novel verb = $t(17) = 3.34$, $p = < .01$; all novel = $t(17) = 5.33$, $p < .001$ and linking trials = $t(17) = 6.89$, $p < .001$).


The autistic group also demonstrated learning on “O” trials, performing significantly above chance on pre-test items $t(15) = 14.45$, $p < .001$, intransitive items $t(15) = 4.04$, $p < .001$, and old items $t(15) = 5.46$, $p < .001$. In contrast, the autistic group performed at chance on all “G” trials: novel verb $t(15) = 1.48$, $p = 4$, all novel $t(15) = 1.03$, $p = 5.4$ and linking trials $t(15) = 1.26$, $p = 4.9$. See table 7 for a complete list of means.

Table 12: Means (S.D.) of typically developing and autistic groups across all trials

	<u>Pre-Test</u>	<u>Old-Items</u>	<u>Intransitives</u>	<u>NewVerb*</u>	<u>AllNew*</u>	<u>Linking*</u>
Typical	98(.05)	94(.11)	90(.19)	77(.35)	86(.28)	87(.22)
Autistic	93(.10)	83(.25)	78(.19)	58(.33)	50(.32)	53(.23)

*= significant difference ($p < .05$) between groups on trial type

Figure 19: Performance of typically developing and autistic groups across testing items



Follow up Investigation

In this follow up work, we aimed to overcome the primary weakness of experiment 4 – the fact the order of items used during testing differed between the autistic and typically developing groups. Although the fact that the autistic children received the items that required generalization first, therefore, possibly making the task easier, the difference in order is still a confound that needed to be addressed. Therefore, a second group of autistic children were enlisted to take part in the same order of test items as the typical developing group from experiment 3 (items requiring generalization last).

Another weakness of experiment 4 is that we lacked a formal assessment of our autistic sample. Because of the significant heterogeneity of autistic populations, it is important to locate where the participants fall along the spectrum. For this, all children in this follow up study were assessed on the Autism Diagnostic Observation Scale (ADOS) (Lord et al 1989).

Methodology

Participants. The subjects for this study were 7 children with autism, aged 8 – 13, recruited from The Douglas Center for Developmental Disabilities at Rutgers University. All parents gave written consent for the child prior to meeting with experimenters, and all children gave verbal consent at the time of participation. Each child had been formally given a diagnosis of autism by an outside agency. Each child was given a children's "eye spy" book in exchange for their participation. Diagnosis for each child was confirmed using the module 2 (limited language ability) of the ADOS (Lord et. al 2000).

ADOS. The ADOS is a semi-structured standardized interview using developmentally appropriate social and toy-based interactions in a 30–45-minute interview to elicit symptoms of autism in four areas: social interaction, communication, play, and repetitive behaviors. The ADOS consists of four different modules, each directed at a particular level of language ability. In the present study, all subjects received Module 2, for children with limited language ability.

Construction Learning. The construction learning task was the same as was used in the previous experiment, with the exception that items requiring generalization are now given last (which is the same order used with the typical group in experiment 4).

Results

ADOS. All 7 children met criteria for a diagnosis of ASD. A summary of their average scores (module 2) is below in table 13.

Table 13. Average scores (S.D.) on the four components of module 2 of the ADOS


Communication	Reciprocal Social Interaction	Play and Imagination	Stereotyped behaviors and Restricted Interests
3.12 (1.26)	4.5(2.64)	0.81(.55)	1.08 (.45)

Construction Learning. A summary of mean performance for each trial type can be found in table 13 below. Overall (see figure 20), we find a significantly $t(6) = 2.3$, $p = < .05$, greater performance on “O” items ($\bar{x} = .64$, $\sigma = .29$) than “G” items ($\bar{x} = .47$, $\sigma = .26$).

Table 14. Summary of mean performance (SD) on each trial type

Pretest	Old Items	Intransitives	New Verb	All New	Linking
0.64 (.26)	0.71(.36)	0.57(.26)	0.35 (.37)	0.57(.18)	0.5(.14)

Figure 20. Performance on “O” items (not requiring generalization), and “G” trials (requiring generalization)



Discussion

Summary of Results. In experiment 3, there exists a large main effect of accuracy between the two groups, but in examining the data closely it does not appear that this reflects global differences in ability. Importantly, both groups in experiment 3 performed comparably on all “O” trials (pretest, old items, and intransitives), which do not require generalization. Instead, these “O” trials collectively assess everything the task requires *except* for generalization: basic language comprehension (pre-test trials and intransitives) and a memory for the previous exemplars (old items). Children with autism show evidence of discerning the ‘approach’ meaning when the test exemplars did not differ from those seen in original exposure, which is especially impressive given in this follow up study, the autistic children saw them last (instead of first) during the testing phase.

Further, this follow up investigation demonstrates that the difference in test item orderings likely cannot account for these group differences. Through just seven autistic participants, we find a similar pattern of performance – a demonstrated proficiency on non generalization items, but a sharp deficiency in the ability to generalize from the witnessed exemplars.

Overall, neither autistic group appeared to form an abstract representation of the approach construction, as they were unable to show any evidence of productive transfer to test items with any degree of surface dissimilarity. Items that differed from the exposure items in any way were not seen as belonging to this same abstract ‘approach’ category. These novel items necessarily require generalization over the exemplars during exposure, in order form such an abstraction capable of assimilating and interpreting new items. This interaction in

performance suggests that the group difference does not result from potential global differences between groups, but reflects something more specific and more interesting. The interaction in performance suggests that the group difference does not result from potential global differences between groups, but reflects something more specific and more interesting. While memory for individual exemplars is reasonably strong, children with autism showed a specific deficit in generalizing over utterances.

General Discussion. We have provided evidence to suggest that children with autism do not readily generalize over the linguistic exemplars they are exposed to. Given the important role that generalization plays in language learning, there is good reason to believe that this may well play a role in hindering autistic language development. Knowledge of language does not consist of a set of unrelated separate exemplars, but rather a rich interrelated network of partially generalized information. Without being able to extract similarities and generalize across exemplars, children with autism are at a disadvantage when trying to learn the rich grammatical patterns of their ambient language.

In general, research into language development in autism has been product-oriented (documenting the peaks and troughs in the language profile) as opposed to process-oriented (looking at the mechanisms at work in the process of development) (Gernsbacher, Geye, & Ellis Weismer, 2005; Swensen, Kelley, Fein & Naigles, 2007; Karmiloff-Smith 1998; Annaz, Karmiloff-Smith and Thomas, 2008). The study of autistic language appears to be no exception to this general trend. While there have been numerous observations of peculiarities in autistic language, few studies to this point have provided insight about how these qualities may have come to exist. Noteworthy findings of autistic language are reviewed here, in light of the current

findings; linguistic under-generalization may help to explain why autistic language appears as it does.

Overall, autistic language has been characterized as involving stored verbatim utterances, exhibiting an overall lack of creativity and productivity (Prizant & Rydell 1984; Dobbison 2000; Perkins, Dobbison, Boucher, Bol and Bloom, 2006). Autistic language often uses the same precise words for general expressions (e.g., “in the middle afternoon” for anytime between 12pm and 7pm), and it appears that many individual words are learned only within strings of longer utterances, as they are not always used appropriately in other contexts (Perkins et al 2006). Relatedly, expressive language in autism is often noted to be better than comprehension (Groen, Zwiers, van der Gaag, and Buitelaar, 2008), which may also reflect an inability to generalize over individually stored chunks of language. If one’s base linguistic knowledge only consisted of a series of separate, stored phrases, an individual would be limited in their expressive ability, but they would still be able to use their stored phrases appropriately. However, this would prove much more costly for comprehending language, since the chance that they would happen to hear one of these phrases in the ambient language is lower.

Parallels appear in conversational speech as well, where much of the autistic individual’s contribution to the conversation appears to be borrowed directly from their interlocutor (Dobbinson et al.2003). While a certain degree of borrowing is typical in normal conversation (Garrod and Pickering, 2004; Pickering and Garrod 2004), with autistics this is often taken to an extreme, jeopardizing basic coherence. The following is a conversation (from Perkins et al, 2006) between an experimenter and an autistic adult named “Gary”:

Exp: And can you tell me Gary, in what direction does to the sun set?

Gary: That direction. Whats the sun set does?

Exp: That's when the sunset goes down, isn't it?

Gary: Tell me, isn't it the sunset does?

The underlined words and phrases are those that Gary borrowed directly from the experimenter. Gary's impoverished ability to communicate is typified by an over reliance on his interlocutor for specific linguistic formulations. It may be that autistics lack the abstract representation needed in order to creatively construct meaningful sentences on their own, and therefore resort to borrowing a great deal of the specific linguistic formulations directly from the conversation.

An extreme form of this is the well-noted autistic tendency of echolalia, in which (typically low functioning) individuals will repeatedly say the same verbatim phrase over and over without any apparent communicative purpose (Prizant and Rydell, 1984). It is uncertain even whether the speaker understands the basic meaning of the utterance. And while not all autistic children exhibit echolalia per se, it appears that a general tendency towards repetitive, unproductive speech does, with echolalia representing an extreme end of this continuum.

It has also been argued that this heavy reliance on recycled, unanalyzed forms may in turn help explain another oft-noted feature of autistic language – errors in pronoun reversal (Bartak and Rutter 1974; Tager-Flusberg, Lord and Paul 2002; Rapin and Dunn 2003). This is the common tendency for autistic children to use “I” to refer to “him” or “her”, and vice versa. Although it is also likely that social cues play a role here as well (Charney 1980), the role of “conversational recycling” has been strongly emphasized (Dabrowska 2004) as well. The

reasoning is relatively straightforward: if phrases are borrowed directly from conversational partners, and these phrases contain a pronoun, it is very likely that this pronoun will be mis-used. Thus, these abstraction difficulties may ultimately help explain somewhat specific features of autistic language.

This work suggests that the language development of children with autism may be hindered by their general cognitive style – a strong attention to details, at the expense of building abstractions. Furthermore, features of the language which do emerge in this population (inflexibility, over-reliance on verbatim utterances) are symptomatic of what one might expect from a population that struggles to learn abstractions. Thus, the ‘generalization deficit’ may play a large role in both hindering and shaping the autistic language profile.

Final words: Language Learning in Children with Autism


Of the many features of language, one of the most striking and impressive is how robust the mechanism of learning is. Provided children are given sufficient input from an early age, they will reliably learn language in a whole host of difficult or unconventional circumstances. Children of bilingual parents reliably learn both languages fluently (Pearson 1998), as do children with low IQs, limited short term phonological memory capacity, or sensory deficits (Snow 1995). In addition, young children who suffer damage to language relevant regions in the left hemisphere still acquire language within a normal range within a few years of the injury (Aram 1998; Bates 1999; Bates et al. 1997). In fact, children whose entire left hemisphere is completely removed can still achieve normal language by maturity (Curtiss and de Bode 1998).

As Ewa Dabrowska states, “non-linguistic cognition does matter – but good input can compensate for cognitive impairments (within limits); and, conversely, good language learners may be able to compensate for sub-optimal input”. (Dabrowska, 2004, pg. 35).

The brain’s ability to acquire functional language is extremely robust, likely because we have access to mechanisms that are somewhat redundant. These observations make the question of autistics’ delay all the more interesting, and may suggest that no single feature is wholly responsible for their difficulties in language. In considering the autistic language profile, we can think of the outcome as being determined by two broad factors: the input (which is primarily influenced by the social interactions and abilities), and the processing and learning based on the input (which is primarily influenced by cognitive factors); see figure 22.

The results discussed here suggest that generalization difficulties may play an important role in both hindering and shaping language development in children with autism. I have argued here that generalization difficulties play an important role, but that’s not to say that they fully explain language difficulties in this population. There is good reason to believe that social characteristics also play an important role in deterring language gains, and the generalization difficulties discussed here should be considered within a greater, all-inclusive understanding of language difficulties. As discussed above, language learning may still be possible if either input is low/impoverished OR if cognitive abilities are deficient. In the case of autism, there appear to be deficits in both: I am stipulating here that autistic language is delayed because input is both low and impoverished, and furthermore, that children with autism are deficient in their ability to generalize.

Figure 22: Summary representation for the primary factors responsible for language acquisition: linguistic input, and learning over that input.



Below, I will briefly detail how children with autism fair on each of these four basic factors: input volume, input richness, item specific knowledge, and generalization over witnessed linguistic exemplars.

The Social Domain

The social profile is clearly an important influence. There are likely (at least) two primary influences within the social profile which likely impede the progress of language development: 1) the quantity of input that one has the opportunity to learn over, and 2) the richness and informativeness of this input.

Input Volume. Children with autism receive less linguistic input than their typically developing counterparts. In general, children with autism have been described as being somewhat indifferent to the individuals in their immediate environment, and do not attend to people the same way as typically developing input. The individuals in our environment are a huge source of linguistic input, and attending infrequently to these sources of linguistic input surely reduces the amount of overall volume autistic children receive. Furthermore, there is evidence that there are key differences in the social and linguistic environments of children with autism, as opposed to typically developing children. In natural settings, parents of autistic children use less child directed speech than parents of typically developing children (Siller and Sigman, 2002). Thus, an initial deficit may be compounded by an altered socio-linguistic environment which the child lives.

Input Richness. At the same time, there is reason to believe that the linguistic input autistic is relatively impoverished, compared to the input which typically developing children receive. The second major detractors here are deficits in joint attention, the ability to engage with another individual in shared focus to an external referent. Children with autism have been noted to have strong difficulties in recognizing eye gaze, in following directed pointing, and overall, in detecting and following the attentional focus of others (Rogers 1999; Gopnick et al 2000; Carpenter, Pennington and Rogers 2002). This has important consequences for language development, as such an ability has been thought to underlie an important requisite for language (Tomasello and Farar 1986; Tomasello 1995; Morales et al 2000). Indeed, early ability in joint attention is a significant predictor of later language outcome (Loveland and Landry 1986; Mundy, Sigman, and Kasari 1990; Kuhl et al. 2005), and therapies directed towards remediating

joint attention difficulties are a primary aim of language remediation therapies (Schreibman 2004).

The Cognitive Domain

Item specific knowledge. As discussed above, item specific knowledge is extremely important aspect of our linguistic knowledge. In order to fully understand language, we must recognize and understand a countless amount of words and idioms. Item specific knowledge also forms a foundation for generalization and abstraction. We must retain specific, verbatim instances in mind at least long enough to encounter similar examples. Without a memory for the examples we've encountered, we would be unable to abstract and generalize. Indeed, we've found evidence that speakers do retain a strong, long-term (e.g. 1 week) memory for specific verbatim instances, even when they are not explicitly asked to remember such items, and even when these instances are non-emotion laden (Gurevich, Johnson and Goldberg 2010).

Item specific linguistic knowledge may be a relative strength of children with autism. Overall, children with autism have been noted to process language in an item-based style, with a relative disregard for greater semantic context. There is evidence that while children with autism have a poor understanding of a story's gist meaning, they have been noted to have a strong memory for verbatim instances (Scheuffgen 1998). A memory for specific exemplars was also a relative strength in experiment, with autistic individuals performing as well as the TD group on old item trials.

In addition, word knowledge may be a relative strength of high functioning children and adolescents with autism, who score well (better than expected for overall language ability) on

standardized vocabulary tests (Jarrold, Boucher, and Russell, 1997; Kjelgaard & Tager-Flusberg, 2001). And, while lexical *knowledge* does not appear to be heightened in non-high functioning autistic individuals, there still appears to be circumstances in which the autistic local bias may lead to superior lexical level performance. As mentioned earlier, and evidenced by the results of Hermelin and O'Connor's (1967) early work, memory is not facilitated by coherent meaning. In false recognition paradigms, this has been shown to yield superior performances. For normal participants, the prototype/misremembering effect is a very robust finding (Omohundro 1981, Solso & McCarthy, 1981, Cabeza, Bruce, Kato, & Oda, 1999): when an individual is shown a series of related words (cake, pie, muffin, donut etc.) and given a recognition test afterwards, they are much more likely to misremember a semantically related word (cupcake) than one which is unrelated (bicycle). However, autistic individuals are much less likely to misremember a semantically word, and are just as likely to make mistakes on related words as they are with unrelated words (Molesworth, Bowler, and Hampton 2005). It appears that individual words are retrieved (and perhaps stored) without the activation a greater semantic network which ordinarily influences (and in some cases misleads) subsequent recognition.

However, it should be pointed out that 'superior word knowledge' is not completely straightforward. While hyperlexia and precocious reading ability have long been associated with autism (Whitehouse and Harris 1984; Tirosh and Canby 1993), superior abilities seen with such decoding is not matched by a superior or even adequate level of understanding for these words (Groen, Zwiers, van der Gaag, Buitelaar 2008). Children with autism do not appear to have a heightened understanding for individual words, perhaps because semantic knowledge, even for a single word, may exist within largely distributed network, requiring at least a certain degree of integration in order to be elicited (Elman 2004). As a side note, deficiency in linguistic

integration is also seen in a reduced McGurk effect (DeGelder, Vroomen, & Van der Heide, 1991) in autistics, under which vision ordinarily has a strong influence verbal perception.

Overall, item specific knowledge seems to be an area of relative strength for children with autism.

Categorization over witnessed items. We have reviewed autistic performance on generalization in detail above, so they will not be repeated here. Suffice to say, children with autism have a strong local bias in non-linguistic cognition, which appears to have strong and largely negative consequences for language learning. Additionally, the observed language characteristics of children with autism (reliance on inflexible, rote memorized utterances, an unorthodox precociousness of expressive language compared to receptive language) are also symptomatic of difficulties in abstraction.

Overall. In summary, children with autism appear to have deficits effecting language in both the social and cognitive domain. Within the social domain, children with autism appear to suffer both from a lack of quality in their linguistic input, as well as a lack of overall volume. Cognitively, children with autism have been noted to have a strong local, item based information processing style, with associated difficulties in abstraction and generalization. Parallels of this are seen in linguistic processing as well, where children have a relatively strong memory and (sometimes) understanding for individual utterances. At the same time however, as we see in experiment 4, children with autism do not fair well in generalizing over these discrete instances to more abstract representations.

Future Directions within Autistic Language Learning

What we may be seeing in children with autism is an imbalance between item specific knowledge and generalization ability: while autistic children have a heightened attention and memory for specific linguistic inputs, they appear to be deficient in their ability to generalize over these instances. It remains to be seen whether these two forces represent a zero sum relationship. For example, does exceptional performance on item specifics necessarily lead to poor abilities in generalization? Or is it possible to have an enhanced ability to process and remember specific which does not infer a cost to generalization ability? Future work should aim to address the relationship between these two cognitive factors for language learning, both in autistic children, and in a normative view of language.

In order to study and discuss their respective influences, it has been useful to isolate social and cognitive factors. However, it is clear that these factors are not completely dissociable, and it may be fruitful to investigate how they interact. For example, difficulties in generalization may hinder one's ability to recognize new examples of common social routines (e.g. meeting somebody new) may detract from the potential richness of linguistic input in such situations. Future work should address such interactions, in order to elucidate a more encompassing understanding over language disturbances in autism.

Summary of Findings

This dissertation has broadly explored language learning. Experiment 1 paves the way for such an exploration, by designating an important aspect of linguistic knowledge to be learned: abstract phrasal constructions. Specifically, we find evidence here that abstract linguistic structures are strongly tied with semantic function, and convey a meaning that is relatively specific, and retrieved automatically. While experiment 1 provided evidence that language users have such stable representations, the rest of the dissertation focuses on how such representations are initially learned. Experiments 2 and 3 explored how such generalized pairings of function are learned at the level of the brain. In both studies, we found evidence that learning may be mediated (at least in part) through a domain general, prediction-based learning mechanism. We also found evidence for the consequences of prediction on neural processing; when constructional form correctly predicts an upcoming visual action, this evokes less overall activity in low-level visual areas. Our final study, experiment 4, examined construction learning in children with autism, a population which struggles in their acquisition of language. We found evidence to suggest that this language deficit may be due, in part, to a difficulty generalizing from discrete linguistic exemplars to more abstract representations.

References:

- Allen, G., & Courchesne, E. (2001). Attention function and dysfunction in autism. *Frontiers in Bioscience*, 37(3), 105-119.
- Baron-Cohen, S. (1991). Do People with Autism Understand What Causes Emotion? *Child Development*, 62(2), 385-395.
- Baron, S.G., Thompson-Schill, S., Weber, M., & Osherson, D. (2010). An early stage of conceptual combination: Superimposition of constituent concepts in left anterolateral temporal lobe. *Cognitive Neuroscience*, 1(1), 44—51. 63
- Baron, S.G., & Osherson, D. (2011). Evidence for conceptual combination in the left anterior temporal lobe. *Neuroimage*. In Press
- Bartak, L., & Rutter, M. (1974). The use of personal pronouns by autistic children. *Journal of Autism and Developmental Disorders*, 4(3), 217-222.
- Bates, E., Thal, D., Trauner, D., Fenson, J., Aram, D., Eisele, J., et al. (1997). From first words to grammar in children with focal brain injury. *Developmental Neuropsychology*, 13(3), 275-343.
- Boyd, J.; & Goldberg, A. Young children's failure to generalize when exposed to the same input as older learners: An advantage in language learning. *To appear*.
- Braine, M. (1971). The acquisition of language in infant and child. In C. Reed (Ed.), *The learning of language* (pp. 7-95). New York: Academic Press.
- Brennan, J., Nir, Y., Hasson, U., Malach, R., Heeger, D. J., & Pylkkanen, L. (n.d.). Syntactic structure building in the anterior temporal lobe during natural story listening. *Brain and Language*, *In Press*
- Bresnan, J., & Kanerva, J. M. (1989). Locative Intension in Chichewa : A Case Study of Factorization in Grammar. *Linguistic inquiry*, 20(1), 1-50. MIT Press.
- Briscoe, T. (2000). Grammatical Acquisition: Inductive Bias and Coevolution of Language and the Language Acquisition Device.
- Buckner, R. (2000). Neuroimaging of Memory. In M. Gazzaniga (Ed.), *The New Cognitive Neurosciences (2nd Edition)* (pp. 1013-1022). Cambridge, MA: MIT Press.
- Canivez, G. (1995). Validity of the Kaufman Brief Intelligence Test: Comparisons with the Wechsler Intelligence Scale for Children, third edition. *Assessment*, 2 (2): 101-111.

- Casenhiser, D., & Goldberg, A. (2005). Fast mapping between a phrasal form and meaning. *Developmental Science*, 8(6), 500-508.
- Cavanna, A. E., & Trimble, M. R. (2006). The precuneus: a review of its functional anatomy and behavioural correlates. *Brain*, 129(3), 564-583.
- Catani, M., Jones, D. K., & Ffytche, D. H. (2005). Perisylvian language networks of the human brain. *Annals of neurology*, 57(1), 8-16
- Chao, L. L., Weisberg, J., & Martin, A. (2002). Experience-dependent Modulation of Category-related Cortical Activity . *Cerebral Cortex* , 12 (5), 545-551.
- Charney, R. (1980). Pronoun errors in autistic children: Support for a social explanation. *International Journal of Language & Communication Disorders*, 15(1), 39-43.
- Chomsky, N. (2002). *On Nature and Language*. Cambridge: Cambridge University Press.
- Chong, S. C., & Treisman, Anne. (2005). Statistical processing: computing the average size in perceptual groups. *Vision Research*, 45(7), 891-900.
- Chong, S., & Treisman, Aanne. (n.d.). Attentional spread in the statistical processing of visual displays. *Attention, Perception, & Psychophysics*, 67(1), 1-13.
- Christiansen, M. H., & Chater, N. (1994). Generalization and connectionist language learning. *Mind & language*, 9(3), 273-287. Blackwell.
- Dabrowska, E. (2004). *Language, mind and brain: Some psychological and neurological constraints on theories of grammar*. Edinburgh: Edinburgh University Press.
- Dakin, S., & Frith, Uta. (2005). Vagaries of Visual Perception in Autism. *Neuron*, 48(3), 497-507.
- Dayan, P., & Niv, Y. (2008). Reinforcement learning: The Good, The Bad and The Ugly. *Current Opinion in Neurobiology*, 18(2), 185-196.
- Dobbinson, S. (2000). Repetitiveness and productivitiy in the language of adults with autism. *Unpublished PhD thesis*.
- Dobbinson, S., Perkins, M., & Boucher, J. (1998). Structural patterns in conversation with a woman who has autism. *Journal of Communication Disorders*, 31(2), 113-134.
- Dobbinson, Sushie, Perkins, M., & Boucher, J. (2003). The interactional significance of formulas in autistic language. *Clinical Linguistics & Phonetics*, 17(4-5), 299-307.

- Dronkers, N. F., Wilkins, D. P., Van Valin, R. D., Redfern, B. B., & Jaeger, J. J. (2004). Lesion analysis of the brain areas involved in language comprehension. *Cognition*, 92(1-2), 145-177.
- Elio, R., & Anderson, J. (1984). The effects of information order and learning mode on schema abstraction. *Memory and Cognition*, 12(1), 20-30.
- Fink, G. et al. (n.d.). A role for Broca's area beyond language processing: evidence from neuropsychology and fMRI. In Y. Amuntz, K and Grodzinsky (Ed.), *Broca's Area*. Oxford University Press.
- Fodor, J. (1983). *Modularity of Mind*. Cambridge: MIT Press/Bradford Books.
- Frith, U. (1989). *Autism: Explaining the Enigma*. Oxford: Oxford University Press.
- Frith, U. & Happe, F. (1994). Autism: Beyond "theory of mind". *Cognition*, 50(2), 115-132.
- Garrod, S., & Pickering, M. J. (2004). Why is conversation so easy? *Trends in Cognitive Sciences*, 8(1), 8-11.
- Gelder, B. de, Vroomen, J., & Heide, L. van der. (1991). Face recognition and lip-reading in autism. *Journal of Cognitive Psychology*, 3(1), 69-86. Psychology Press.
- Gentner, D., Loewenstein, J., & Hung, B. (2007). Comparison Facilitates Children's Learning of Names for Parts. *Journal of Cognition and Development*, 8(3), 285-307. Psychology Press.
- Goldberg, A., Casenhiser, D., & Sethuraman, N. (2004). Learning argument structure generalizations. *Cognitive Linguistics*, 15(3), 289-316.
- Goldberg, A., Casenhiser, D., & White, T. (2007). Constructions as categories of language. *New Ideas in Psychology*, 25(2), 70-86.
- Goldberg, A. (1992). The inherent semantics of argument structure: The case of the English ditransitive construction. *Cognitive Linguistics*, 3(1), 37-74.
- Goldberg, A. (2006). *Constructions at work: The nature of generalization in language*. Oxford: Oxford University Press.
- Goldberg, A. E., Casenhiser, D. M., & Sethuraman, N. (2004). Learning argument structure generalizations. *Cognitive Linguistics*, 15(3), 289-316. De Gruyter. Retrieved from <http://dx.doi.org/10.1515/cogl.2004.011>.
- Green, G. (1974). *Semantics and syntactic regularity*. Urbana-Champaign: UI Press.
- Grodzinsky, Y., & Friederici, A D. (2006). Neuroimaging of syntax and syntactic processing. *Current Opinion in Neurobiology*, 16(2), 240-246.

- Groen, W. B., Zwiers, M. P., Gaag, R.-J. van der, & Buitelaar, J. K. (2008). The phenotype and neural correlates of language in autism: An integrative review. *Neuroscience & Biobehavioral Reviews*, 32(8), 1416-1425.
- Happé, F. (1999). Autism: cognitive deficit or cognitive style? *Trends in Cognitive Sciences*, 3(6), 216-222.
- Happé, F., & Frith, Uta. (n.d.). The Weak Coherence Account: Detail-focused Cognitive Style in Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders*, 36(1), 5-25. Springer Netherlands.
- Henson, R. N. A., Rugg, M. D., Shallice, T., Josephs, O., & Dolan, R. J. (1999). Recollection and Familiarity in Recognition Memory: An Event-Related Functional Magnetic Resonance Imaging Study. *The Journal of Neuroscience*, 19(10), 3962-3972.
- Hill, E. L. (2004). Executive dysfunction in autism. *Trends in Cognitive Sciences*, 8(1), 26-32.
- Hoffman, J. E. (1980). Interaction between global and local levels of a form. *Journal of Experimental Psychology: Human Perception and Performance*, 6(2), 222-234.
- Iacoboni, M., & Wilson, S. M. (2006). Beyond a Single Area: Motor Control and Language Within a Neural Architecture Encompassing Broca's Area. *Cortex*, 42(4), 503-506.
- Jarrold, C., & Brock, J. (n.d.). To Match or Not to Match? Methodological Issues in Autism-Related Research. *Journal of Autism and Developmental Disorders*, 34(1), 81-86. Springer Netherlands.
- Jenkins, I. H., Brooks, D. J., Frackowiak, R. S. J., & Passingham, F. E. (1994). Motor Sequence Tomography Learning : A Study with Positron Emission Tomography. *Journal of Neuroscience*, 14(6): 3775 - 3790
- Karmiloff-Smith, A. (1998). Development itself is the key to understanding developmental disorders. *Trends in Cognitive Sciences*, 2(10), 389-398.
- Karmiloff-smith, A. (2010). A Developmental Perspective on Modularity. (B. Glatzeder, V. Goel, & A. Müller, Eds.) *Brain*, 179-187.
- Klinger, L., & Dawson, G. (2001). Prototype formation in autism. *Development and Psychopathology*, 13(01), 111-124.
- Klinger, L. & Dawson, G. (1997). Autistic Disorder. In A. Mash, J., & Barley (Ed.), *Childpsychopathology* (pp. 315-339). New York: Guilford Press.
- Koenig, P., Smith, E. E., Glosser, G., DeVita, C., Moore, P., McMillan, C., et al. (2005). The neural basis for novel semantic categorization. *NeuroImage*, 24(2), 369-383.

- Kuhl, Patricia K., & Rivera-Galxiola, M. (2008). Neural Substrates of Language Acquisition. *Annual review of neuroscience*, 31(10), 511-534.
- Liss, M., Fein, D., Allen, D., Dunn, M., Feinstein, C., Morris, R., et al. (2001). Executive Functioning in High-functioning Children with Autism. *The Journal of Child Psychology and Psychiatry and Allied Disciplines*, 42(02), 261-270.
- Lord, C., Rutter, M., Goode, S., Heemsbergen, J., Jordan, H., Mawhood, L., et al. (n.d.). Autism diagnostic observation schedule: A standardized observation of communicative and social behavior. *Journal of Autism and Developmental Disorders*, 19(2), 185-212. Springer Netherlands.
- Lovaas, O. I., Koegel, R. L., & Schreibman, L. (1979). Stimulus overselectivity in autism: a review of research. *Psychological bulletin*, 86(6), 1236-54.
- Marcus, G. (2000). Negative evidence in language acquisition. *Cognition*, 46(1), 53-85. Elsevier.
- Markman, A. B., & Maddox, W. T. (2003). Classification of exemplars with single- and multiple-feature manifestations: The effects of relevant dimension variation and category structure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(1), 107-117. doi: 10.1037/0278-7393.29.1.107.
- May, J. C., Delgado, M. R., Dahl, R. E., Stenger, V. A., Ryan, N. D., Fiez, J. A., et al. (2004). Event-related functional magnetic resonance imaging of reward-related brain circuitry in children and adolescents. *Biological Psychiatry*, 55(4), 359-366.
- Mills, D., Coffey-Corina, S. A., & Neville, H. (1993). Language acquisition and cerebral specialization in 20 month old infants. *Journal of Cognitive Neuroscience*, 5, 317-334.
- Mills, D., Prat, C., Zangl, R., Stager, C., Neville, H., & Werker, J. (2004). Language experience and the organization of brain activity to phonetically similar words: ERP evidence from 14 and 20 month olds. *Journal of Cognitive Neuroscience*, 16(1452-1464).
- Minschew, N. J., Goldstien, G., & Siegel, D. O. N. J. (1997). Neuropsychologic functioning in autism: Profile of a complex information processing disorder. *Journal of the International Neuropsychological Society*, 3(04), 303-316.
- Molesworth, C. J., Bowler, D. M., & Hampton, J. A. (2005). The prototype effect in recognition memory: intact in autism? *Journal of Child Psychology and Psychiatry*, 46(6), 661-672.
- Molesworth, C., Bowler, D., & Hampton, J. (n.d.). When Prototypes Are Not Best: Judgments Made by Children with Autism. *Journal of Autism and Developmental Disorders*, 38(9), 1721-1730.

- Molfese, D., Wetzel, W., & Gill, I. (1993). Known versus unknown word discrimination in 12 month-old infants: Electrophysiological correlates. *Developmental Neuropsychology*, 3(4), 241-258.
- Mottron, L., Burack, J. A., Iarocci, G., Belleville, S., & Enns, J. T. (2003). Locally oriented perception with intact global processing among adolescents with high-functioning autism: evidence from multiple paradigms. *Journal of Child Psychology and Psychiatry*, 44(6), 904-913.
- Mottron, L., Dawson, M., Soulières, I., Hubert, B., & Burack, J. (n.d.). Enhanced Perceptual Functioning in Autism: An Update, and Eight Principles of Autistic Perception. *Journal of Autism and Developmental Disorders*, 36(1), 27-43.
- Mundy, P., Sigman, M., & Kasari, C. (1990). A longitudinal study of joint attention and language development in autistic children. *Journal of autism and developmental disorders*, 20(1), 115-28.
- Namy, L., & Gentner, S. (2001). Making a silk purse out of two sow's ears: Young children's use of Comparison in category learning. *Journal of experimental psychology. General*, 131(1), 5-15. American Psychological Association.
- Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, 9(3), 353-383.
- Neville, H; Mills, D; Bellugi, U. (1994). Effects of altered auditory sensitivity and age of language acquisition on the development of language - relevant neural systems: Preliminary studies of Williams Syndrome. In G. J. Broman S (Ed.), *Atypical cognitive deficits in developmental disorders: implications for brain function* (pp. 67-83). Hillsdale, NJ.
- Opitz, B., & Friederici, Angela D. (2003). Interactions of the hippocampal system and the prefrontal cortex in learning language-like rules. *NeuroImage*, 19(4), 1730-1737.
- Opitz, B., & Friederici, Angela D. (2004). Brain correlates of language learning: the neuronal dissociation of rule-based versus similarity-based learning. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 24(39), 8436-40.
- Osterhout, L., Kim, A., & Kuperberg, G. (2006). The Neurobiology of Sentence Comprehension. *Psychiatry: Interpersonal and Biological Processes*, (27), 1-23.
- O'Connor, N., Hermeilin. (1967). Auditory and visual memory in typical and atypical children. *Journal of Intellectual Disability Research*, 11(2), 126-131. Blackwell Publishing Ltd.
- Perkins, M. R., Dobbinson, Sushie, Boucher, J., Bol, S., & Bloom, P. (n.d.). Lexical Knowledge and Lexical Use in Autism. *Journal of Autism and Developmental Disorders*, 36(6), 795-805.

- Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27(02), 169-190.
- Pinker, S. (1989). *Learnability and Cognition*. Boston: MIT Press.
- Pinker, S. (1994). *The language instinct: How the mind creates language*. New York: William Morrow and Company.
- Plaisted, K., Dobler, V., Bell, S., & Davis, G. (n.d.). The Microgenesis of Global Perception in Autism. *Journal of Autism and Developmental Disorders*, 36(1), 107-116.
- Plaisted, K., O'Riordan, M., & Baron-Cohen, S. (1998). Enhanced Visual Search for a Conjunctive Target in Autism: A Research Note. *Journal of Child Psychology and Psychiatry*, 39(5), 777-783. Blackwell Publishers Ltd.
- Platel, H., Baron, J.-C., Desgranges, B., Bernard, F., & Eustache, F. (2003). Semantic and episodic memory of music are subserved by distinct neural networks. *NeuroImage*, 20(1), 244-256.
- Poldrack, R A, Clark, J., Pare-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., et al. (2001). Interactive memory systems in the human brain. *Nature*, 414(6863), 546-550.
- Poldrack, Russell A, & Foerde, K. (2008). Category learning and the memory systems debate. *Neuroscience & Biobehavioral Reviews*, 32(2), 197-205.
- Pomerantz, J. R. (1983). Global and local precedence: Selective attention in form and motion perception. *Journal of Experimental Psychology: General*, 112(4), 516-540.
- Prizant, B. M., & Rydell, P. J. (1984). Analysis of Functions of Delayed Echolalia in Autistic Children. *Journal of Speech and Hearing Research*, 27(2), 183-192.
- Rutherford, M., & McIntosh, D. (n.d.). Rules versus Prototype Matching: Strategies of Perception of Emotional Facial Expressions in the Autism Spectrum. *Journal of Autism and Developmental Disorders*, 37(2), 187-196. Springer Netherlands.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80(1), 1-27.
- Seger, C., Prahakharan, V., Poldrack, R., and G. J. (2000). Neural activity differs between explicit and implicit learning of artificial grammar strings : An fMRI study. *Psychobiology*, 28(3), 283-292.
- Silva-Pereyra, J., Conboy, B. T., Klarman, L., & Kuhl, P. K. (2007). Grammatical Processing without Semantics? An Event-related Brain Potential Study of Preschoolers using Jabberwocky Sentences. *Journal of Cognitive Neuroscience*, 19(6), 1050-1065. MIT Press.

- Skosnik, P., Mirza, F., Gitelman, D., Parrish, T., Mesulam, M., & Reber, P. (2002). Neural Correlates of Artificial Grammar Learning. *NeuroImage*, 17(3), 1306-1314. doi: 10.1006/nimg.2002.1291.
- Smith, E. E., & Grossman, M. (2008). Multiple systems of category learning. *Neuroscience & Biobehavioral Reviews*, 32(2), 249-264.
- Snyder, A. W., Mulcahy, E., Taylor, J. L., Mitchell, D. J., Sachdev, P., & Gandevia, S. C. (2003). Savant-like skills exposed in normal people by suppressing the left fronto-temporal lobe. *Journal of integrative neuroscience*, 2(2), 149-58.
- Soulières, I., Mottron, L., Saumier, D., & Larochelle, S. (n.d.). Atypical Categorical Perception in Autism: Autonomy of Discrimination? *Journal of Autism and Developmental Disorders*, 37(3), 481-490.
- South, M., Ozonoff, S., & McMahon, W. (2007). The relationship between executive functioning, central coherence, and repetitive behaviors in the high-functioning autism spectrum. *Autism : the international journal of research and practice*, 11(5), 437-51.
- Spelke, E. (1998). Nativism, empiricism, and the origins of knowledge. *Infant Behavioral Development*, 21(2), 181-200.
- Swensen, L. D., Kelley, E., Fein, D., & Naigles, L. R. (2007). Processes of Language Acquisition in Children With Autism: Evidence from Preferential Looking. *Child Development*, 78(2), 542-557. Blackwell Publishing Inc.
- Tager-Flusberg, H. (1985). The Conceptual Basis for Referential Word Meaning in Children with Autism. *Child Development*, 56(5), 1167-1178. Blackwell Publishing on behalf of the Society for Research in Child Development.
- Tager-Flusberg, H. (2004). Strategies for Conducting Research on Language in Autism. *Journal of Autism and Developmental Disorders*, 34(1), 75-80.
- Tager-flusberg, H., Paul, R., & Lord, C. (2002). Language and Communication in Autism. *Language*, 335-364.
- Tanaka, D., Kiyokawa, S., Yamada, A., Dienes, Z., & Shigemasa, K. (n.d.). Role of selective attention in artificial grammar learning. *Psychonomic Bulletin & Review*, 15(6), 1154-1159.
- Thal, D. J., Marchman, V., Stiles, J., Aram, D., Trauner, D., Nass, R., et al. (1991). Early lexical development in children with focal brain injury. *Brain and Language*, 40(4), 491-527.
- Thierry, G., Giraud, A.-L., & Price, C. (2003). Hemispheric Dissociation in Access to the Human Semantic System. *Neuron*, 38(3), 499-506.

- Tomasello, M. (2005). *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Boston: Harvard University Press.
- Trabasso, T., & Bower, G. (1966). Presolution dimensional shifts in concept identification: A test of the sampling with replacement axiom in all-or-none models. *Journal of Mathematical Psychology*, 3(1), 163-173.
- Treisman, Anne, & Souther, J. (1985). Search Asymmetry: A Diagnostic for Preattentive Processing of Separable Features. *Journal of Experimental Psychology: General*, 114(3), 285-310.
- Tulvig, E. (1983). *Elements of episodic memory* (2nd ed.). Oxford and New York: Clarendon Press.
- Tulving, L. N. E. (1997). Searching for Memory Systems. *Journal of Cognitive Psychology*, 9(1), 121-125. Psychology Press.
- Turk-Browne, N. B., Jung, J. A., & Scholl, B. J. (2005). The Automaticity of Visual Statistical Learning. *Journal of Experimental Psychology: General*, 134(4), 552-564.
- Ungerer, J. A., & Sigman, Marian. (n.d.). Categorization skills and receptive language development in autistic children. *Journal of Autism and Developmental Disorders*, 17(1), 3-16. Springer Netherlands.
- Wang L., Mruczek REB, Arcaro MJ and Kastner S. (in prep) Visual topographic probability maps (VTPM) in standard MNI space.
- Wang, Y., Sereno, J. A., Jongman, A., & Hirsch, J. (2003). fMRI Evidence for Cortical Modification during Learning of Mandarin Lexical Tone. *Journal of Cognitive Neuroscience*, 15(7), 1019-1027. MIT Press.
- White, S., O'Reilly, H., & Frith, Uta. (2009). Big heads, small details and autism. *Neuropsychologia*, 47(5), 1274-1281.