

Précis of *Learning and the language of thought*

Steven T. Piantadosi

ONE of the outstanding mysteries of language acquisition is how children learn the meaning of *function words*. Function words like “and,” “the,” and “when” are linguistic symbols that do not refer to any observable object, property, or action in the world. Instead, they convey semantic operations and reflect a deep cognitive process that allows language users to express—and potentially think—complex thoughts. For instance: *Most politicians who are successful are born near a town with either one factory or two malls*. Such utterances are possible because language users know the meaning of words such as “most,” “who,” “a,” “either,” and “or” that combine with other words to express abstract logical relationships.

Because function words are non-referential, they outside the domain of most contemporary word learning models (e.g. Frank, Goodman, & Tenenbaum, 2007a; Xu & Tenenbaum, 2007; Yu & Ballard, 2007), which typically only track co-occurrences between words and objects. As such, there are no working theories for how children might acquire the semantics of function words. This is unfortunate because the compositionality, abstraction, and structure behind these word meanings makes them some of the most compelling and unique features of human communication. Indeed, functions words motivate an entirely different perspective on language acquisition, where the key aspect of learning word meanings is inferring the unobserved semantic operations the word conveys, rather than the observed features of the world it maps to.

This thesis, *Learning and the language of thought*, combines computational, developmental, and experimental approaches in order to understand the type of inductive learning that supports acquisition of rich conceptual structures like function word meanings. The starting point for this approach is Fodor (1975)’s *language of thought* (LOT) hypothesis. This theory posits that a structured, compositional system supports mental representation (see also Boole, 1854; Frege, 1892). A LOT is much like a computer programming language for thinking: complex thoughts are represented by syntactically combining elements from a smaller set of primitive representations. A thought such as “All men are violinists” might be represented by a logical expression such as “ $\forall x.man(x) \rightarrow violinist(x)$ ”, where “ \forall ”, “*man*”, “*violinist*”, and “ \rightarrow ” are primitive cognitive symbols. Such a representation system has been argued to explain core properties of cognitive processes, including most notably their compositionality, systematicity, and productivity (Fodor & Pylyshyn, 1988).

This thesis extends theorizing about a LOT into the domain of learning and developmental change (see also Goodman, Tenenbaum, Feldman, & Griffiths, 2008; Goodman, Tenenbaum, Griffiths, & Feldman, 2008; Kemp, Goodman, & Tenenbaum, 2008). Learners may be genetically endowed with some elementary representations—perhaps for operations that manipulate sets (e.g. *union*, *intersection*, etc.) or logical relationships (e.g. *and*, *or*, *not*, etc.). These are the types of representations would likely be part of the conceptual *core* that all learners bring to development (Spelke, 2003, 2004; Spelke & Kinzler, 2007; Carey, 2009). Key aspects of learning can then be captured by statistical *inferences* over compositions of these primitives. That is, learners may observe word usages and infer the unobserved compositions of primitives that characterize the word’s truth conditions. For instance, the meaning of a quantifier like “most” is typically denoted by semanticists as a relation on sets (e.g. Heim & Kratzer, 1998): “most *X* are *Y*” if there are more *X*s that are *Y*s than *X*s that are not *Y*s. Formally: $|X \cap Y| > |X \setminus Y|$, where $|\cdot|$ is a cardinality operation, \cap is set intersection, and \setminus is set difference. A learner who brought knowledge of operations like \cap , \setminus , and $|\cdot|$ might construct this meaning by composing these operations, searching through combinations until they found one that could adequately explain utterances adults produce. They would have to use available data to rule out other potential compositions of primitives, such as $|X \setminus Y| > |X \cap Y|$ or $|X| > |Y \setminus X|$. As might be intuitively clear, such learning is not computationally trivial. However, we prove that such learners can always arrive at the correct (unobserved) semantic representations under reasonable statistical assumptions.

This approach to learning is distinguished from most contemporary theories of language acquisition. It differs from, for instance, “triggering” approaches (e.g., Gibson & Wexler, 1994) in that the specific adult representations need not be innately specified—the compositional structures themselves are what learners must create. Our approach is also distinguished from classically empiricist models that address learning using arguably plausible representational and implementational assumptions (e.g., Rumelhart & McClelland, 1986; Elman, 1997): we focus on computational-level accounts (Marr, 1982) of inductive phenomena without

addressing issues of neural implementation¹. Our work builds on previous computational work in language learning that has attempted to learn rich linguistic structures by “building in” relatively minimal structural components (Bod, 2009; O’Donnell, Tenenbaum, & Goodman, 2009; Perfors, Tenenbaum, & Regier, 2011). Our approach corresponds to empiricism about complex conceptual knowledge and representations, combined with nativism about scaffolding that allows for structured representations—the syntax and primitives of the language of thought.

The assumptions of this setup are simultaneously trivial and consequential. Trivially, any theory that posits that development is driven by combining and reusing early abilities must specify a means of combination. Here, this is chosen to be perhaps the simplest means of combination that is also powerful: function composition. Our formalism that captures this compositionality, *lambda calculus* (Church, 1936), builds in only a few simple—almost vacuous—rules for composition (see Hindley & Seldin, 1986). But including functional composition is tremendously consequential for a cognitive or developmental theory since it allows any computable function to be expressed. This approach gives a powerful tool for developmental theories, as it allows learners to consider hypotheses of arbitrary computational complexity. The right learning theory in such a system has the potential to be far-reaching, providing a computational framework for explaining people’s remarkably productive cognitive capacities. The idea of learning in computationally powerful systems builds off of recent work examining the learnability of language from the perspective of algorithmic information theory (Chater & Vitányi, 2007). The present work takes the theoretical idea of Chater and Vitányi (2007)—based on inferring bit strings that describe Turing machines—and implements it in a real *cognitive* theory that expresses computation using developmentally plausible primitives.

Three primary areas are studied by the chapters in this thesis: number words, quantifiers, and rule-based concept learning.

Learning number through bootstrapping

The first chapter in this thesis studies the case of number-word acquisition. Number words are interesting because their acquisition seems to involve a dramatic conceptual shift (Carey, 2009). Children initially learn number words as a counting routine, a memorized list much like the alphabet. They are able to recite the words sequentially, but at first do not know any of their meanings. Children then successively learn the meaning of the first several words, becoming *one-knowers*, then *two-knowers*, *three-knowers* and sometimes *four-knowers* (Wynn, 1990, 1992; Sarnecka & Lee, 2009; Lee & Sarnecka, 2010b, 2010a). A two-knower for instance is able to correctly count a set containing one or two objects, but when asked for three or more does not respond correctly. Around “three” or “four,” children stop learning the number word meanings sequentially, and appear to acquire all of the rest of the word meanings more or less at once, becoming “CP-knowers” (Cardinal-principle knowers). CP-knowers can use counting to determine cardinality for any number word they know, and thus appear to have induced more general knowledge about how counting, numerosity, and their memorized list of number words relate.

This developmental pattern is interesting because it has been argued to be a clear instance of developmental change (though see Gallistel & Gelman, 1992; Gelman & Gallistel, 1978; Leslie, Gelman, & Gallistel, 2008). Children in later stages of number learning appear to have qualitative knowledge—how the number system works—that is not possessed by younger learners. The puzzle of how learners come to this knowledge motivated Carey (2009) to formulate a bootstrapping theory of number word learning centered on elaboration of early representational systems. In Carey’s theory, learners acquire the first several number word meanings by mapping the words to memorized sets of objects. For instance, children might learn that “three” maps to $\{X, X, X\}$. To determine if a set has “three” elements, they must find a one-to-one correspondence with this memorized set. Carey argues that the transition to a CP-knower occurs when children notice a simple analogy between their memorized set-based meanings, and the number list: moving one more on the number list corresponds to “one more” element in the relevant set. Reasoning by analogy, they use this principle to determine all of the other number word meanings in the counting routine. However, Carey’s theory has been criticized for failing to justify the critical induction it requires, and indeed presupposing the conceptual system it aims to explain (Rips, Asmuth, & Bloomfield, 2006, 2008; Rips, Bloomfield, & Asmuth, 2008).

The first chapter of the thesis provides a formal model which is capable of learning the structure of number words as compositions of operations which manipulate sets and logical values. The conceptual representation that the model arrives at can naturally be interpreted as an algorithm for counting. This model implements a theory much like Carey (2009)’s, but uses several pieces of machinery, each of which has been independently proposed to explain cognitive phenomena in other domains. The representational system we use is lambda

¹Connectionist approaches have, however, attempted to understand how compositional, hierarchical structures similar to those in our LOT models might be encoded with neurally plausible representations (Smolensky & Legendre, 2006).

calculus, a formal language for compositional semantics (e.g., Heim & Kratzer, 1998; Steedman, 2000), computation more generally (Church, 1936), and other natural-language learning tasks (e.g. Zettlemoyer & Collins, 2005, 2007; Piantadosi, Goodman, Ellis, & Tenenbaum, 2008). The core inductive part of the model uses Bayesian statistics to formalize what inferences learners should make from data. This involves two key parts: a likelihood function which measures how well hypotheses fit observed data, and a prior which measures the complexity of individual hypotheses. We use simple and previously-proposed forms of both. The model uses a likelihood function that uses the *size principle* (Tenenbaum, 1999) to penalize hypotheses which make overly broad predictions. Frank, Goodman, and Tenenbaum (2007b) proposed that this type of likelihood function is important in cross-situational word learning and Piantadosi et al. (2008) showed that it could solve the *subset problem* in learning compositional semantics. The prior is from the *rational rules* model of Goodman, Tenenbaum, Feldman, and Griffiths (2008), which first linked probabilistic inference with formal, compositional, representations. The prior assumes that learners prefer simplicity and re-use in compositional hypotheses and has been shown to be important in accounting for human rule-based concept learning. We provide the model with naturalistic data—number words noisily paired with sets of objects, matched to the frequency distribution of number words in natural language (Dehaene & Mehler, 1992).

In this system, learners are assumed to know set-theoretic operations, logical connectives, and have representations for small set cardinalities. We show that when given idealized evidence, a compositional model over these connectives robustly gives rise to several qualitative phenomena: the model progresses through three or four distinct subset knower-levels (Wynn, 1990, 1992; Sarnecka & Lee, 2009; Lee & Sarnecka, 2010b, 2010a), does not assign specific numerical meaning to higher number words at each subset-knower level (Condry & Spelke, 2008), and suddenly infers the meaning of the remaining words on the count list after learning (Wynn, 1990, 1992; Sarnecka & Lee, 2009; Lee & Sarnecka, 2010b, 2010a). We also show that this model is also able to learn other kinds of conceptual systems that manipulate sets—for instance, singular/plural systems, and Mod-N systems like those suggested by (Rips et al., 2006).

This work provides the first fully formalized and implemented theory of conceptual change in number learning. Contra Rips et al. (2006), this work shows that bootstrapping can be formally implemented in a way that does not presuppose numerical knowledge. Our results are closely tied to developmental findings, showing that children’s progression through stages of numerical knowledge may be the result of statistical inference over a sufficiently rich representational space.

Quantifiers and learnability

The second chapter in this thesis extends our approach of statistical models over a LOT to the acquisition of quantifiers. As discussed above, quantifiers are often taken to denote relations between sets (see also Montague, 1973; Barwise & Cooper, 1981; Keenan & Stavi, 1986; Keenan & Westerstahl, 1997; Heim & Kratzer, 1998), and their meaning is expressed by semanticists in a logical representation language. Quantifiers are especially interesting for statistical learning theories because their meanings are abstract, and often involve subtle presuppositional and pragmatic content.

We show that a model much like the number model is capable of learning representations required for quantifiers, using only noisy positive evidence. As with the number model, this model assumes that these semantic representations are expressed by compositions of primitive logical and set-theoretic functions. Indeed, the probabilistic setup is nearly identical to that of the number model, meaning that the same representational and inferential machinery could be used in both domains. Critically, the model learns several abstract aspects of meaning, including literal meaning, presupposition, and probability of production. For instance, in a sentence like “The X is Y.” the word “the” presupposes that there is one (contextually salient) element of X and asserts that it is also an element of Y (e.g. “The boy is a fisherman.”). We show that an appropriate kind of statistical learning model can acquire logical representations for these aspects of meaning in a developmentally-plausible amount of data.

We compare errors made by the model to patterns observed developmentally in the learning of “the” (Wexler, 2011), and “every” (Roepfer, Strauss, & Pearson, 2004; Philip, 1995), and argue that the learning model makes similar patterns of mistakes. This potentially provides a domain-general account of these errors based on idealized statistical learnability. We additionally present a statistical proof that the model is always capable of arriving at the correct representations, even in the face of noisy data. This provides substantial progress beyond the previous state-of-the-art in quantifier learning, which was only able to prove learnability for a subset of quantifiers (Clark, 2010; Tiede, 1999; Florêncio, 2002), and did not begin to approach other aspects of meaning, like presupposition. This chapter articulates a version of the *subset problem* (Berwick, 1985; Wexler & Manzini, 1987; Smolensky, 1996; Hale & Reiss, 2003) in quantifier learning, where learners may incorrectly infer over-restricted meanings for quantifiers and never receive negative evidence that contradicts their hypothesis. Learners might incorrectly hypothesize, for instance, that “every” means

some. Because “every” logically implies *some*, learners would never be “wrong” with such a hypothesis. We show that, as in other models of word learning, this problem can be handled by using a generative statistical model with the *size principle* (Tenenbaum, 1999). Intuitively, if “every” meant *some*, one would expect to see it in many more contexts where it does not appear. This suggests that generative statistical models may solve the subset problem in many domains of language acquisition.

Our implementation further allows us to test various restrictions on the space of quantifier meanings, including “maximally nativist” theories in which only the correct set of quantifier meanings is innately specified, to learning in a space of only conservative quantifiers (Keenan & Stavi, 1986; Barwise & Cooper, 1981), and finally to a full, unrestricted hypothesis space. We show that learning in the full, unrestricted space is not substantially harder than the maximally nativist space, nor is it substantially harder than learning in the space of conservative quantifiers. These results illustrate the plausibility of learning complex function word semantics without requiring substantial domain-specific knowledge.

What is the language of thought?

One advantage of studying development as inductive learning in a LOT is that doing so allows for some flexibility in theorizing. One can write down any hypothetical LOT and see its consequences for learning. This has recently allowed LOT-learning theories to be applied to explain a wide range of developmental phenomena, including those outside of natural language, such as learning family-tree relations and theories like magnetism (e.g. Katz, Goodman, Kersting, Kemp, & Tenenbaum, 2008; Kemp et al., 2008; Goodman, Ullman, & Tenenbaum, 2009; Ullman, Goodman, & Tenenbaum, 2010). Most of this work shows how learning could proceed when assuming a *particular* set of primitive components. The third chapter extends this approach to quantitatively compare the performance of *different* LOTs—as representational theories—in the same domain.

In a massive, high-throughput concept-learning experiment, we taught subjects rule-based concepts on sets of objects, ranging from simple Boolean predicates (*circle or red*) to predicates involving quantification (*at least one other object in a set is the same color*). This work extends previous research on Boolean concept learning (Shepard, Hovland, & Jenkins, 1961; Feldman, 2000; Goodman, Tenenbaum, Feldman, & Griffiths, 2008) to the types of concepts that are likely necessary for natural language semantics—in particular function words which manipulate and quantify over sets, although the specific concepts learned did not correspond to function word meanings.

Each concept can be computed or represented in a huge number of ways; for instance, all Boolean concepts can be written using standard logical connectives (*and*, *or*, *not*), or using only a single universal logical connective (*nand*, or not-and). Or, one might imagine a system full of a rich set of logical connectives including, perhaps, logical implication (*implies*) or biconditional (*iff*). Similarly, expressions with quantifiers may employ variously rich or simplified systems, ranging from a single existential or universal quantifier, to rich types of first- or higher-order quantification. In each case, a set of logical operations represents a specific *representational theory* that the study of concept representation and learning should aim to discover.

By implementing a learning model and a Bayesian data analysis method, we take the learning experiment and produce a “score” for any hypothesized language, corresponding to its ability to predict human learning curves. This ability to work backwards from a learning curve and attempt to discover what type of compositional system likely gave rise to it is novel and potentially quite powerful, as it provides a solid empirical grounding for LOTs that were largely theoretical. Through this, we are able to provide evidence against intuitively implausible bases such as the *nand*-basis, and can test the cognitive plausibility of several interesting representation systems that have been suggested in cognitive science and AI. Moreover, we are able to use this technique to provide quantitative evidence that concept learning operates over richer structures than boolean logic, including interesting types of quantification.

The third chapter therefore presents experiments in adults which both lend credibility to the general framework of statistics-over-compositions proposed in earlier chapters, and attempts to discover what exactly the right compositional system might look like. As such, it relates concept learning to developmental models, and combines several novel statistical techniques. In general, we find that representational systems with non-restricted syntactic forms, a rich set of primitive connectives, and quantification can best explain human learning. This work moves the LOT from philosophical ground to a firm empirical basis, and the experiments provide a compelling data set for comparing different paradigmatic approaches in cognitive science.

Conclusion

This work shows how ideas about induction and representation can be combined into a coherent picture of developmental change. This thesis formalizes the idea that learners bring a capacity for conceptual combination to the problem of learning, and construct representations of the world much as programmers

write programs or scientists develop theories. It argues that this can explain (i) patterns of development in learning number words, (ii) the capacity of children to acquire abstract function word meanings like quantifiers, and (iii) adult learning curves in a massive concept learning experiment. To accomplish this, a learning system must combine statistical inference with rich structured representations. With this setup we have developed computational and empirical methods for studying the interaction of learning and the language of thought. Further, we have argued that this approach can help bring clarity to several deep issues in the study of cognition, including the cognitive problem of induction, questions of innateness and learnability, and the basic puzzle of how complex representations may arise throughout development.

References

- Barwise, J., & Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and philosophy*, 4(2), 159–219.
- Berwick, R. (1985). *The acquisition of syntactic knowledge*. Cambridge, MA: MIT Press.
- Bod, R. (2009). From exemplar to grammar: A probabilistic analogy-based model of language learning. *Cognitive Science*, 33(5), 752–793.
- Boole, G. (1854). *An investigation of the laws of thought: on which are founded the mathematical theories of logic and probabilities*. London, UK: Walton and Maberly.
- Carey, S. (2009). *The origin of Concepts*. Oxford: Oxford University Press.
- Chater, N., & Vitányi, P. (2007). Ideal learning³ of natural language: Positive results about learning from positive evidence. *Journal of Mathematical Psychology*, 51(3), 135–163.
- Church, A. (1936). An unsolvable problem of elementary number theory. *American journal of mathematics*, 58(2), 345–363.
- Clark, R. (2010). *Some computational properties of generalized quantifiers*. (Unpublished manuscript)
- Condry, K. F., & Spelke, E. S. (2008). The Development of Language and Abstract Concepts: The Case of Natural Number. *Journal of Experimental Psychology: General*, 137, 22–38.
- Dehaene, S., & Mehler, J. (1992). Cross-linguistic regularities in the frequency of number words* 1. *Cognition*, 43(1), 1–29.
- Elman, J. (1997). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: the MIT Press.
- Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, 407(6804), 630–633.
- Florêncio, C. (2002). Learning generalized quantifiers. In *Proceedings of Seventh ESSLLI Student Session*.
- Fodor, J. (1975). *The language of thought*. Cambridge, MA: Harvard University Press.
- Fodor, J., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: a critical analysis, Connections and symbols. *A Cognition Special Issue, S. Pinker and J. Mehler (eds.)*, 3–71.
- Frank, M., Goodman, N., & Tenenbaum, J. (2007a). A Bayesian framework for cross-situational word learning. *Advances in neural information processing systems*, 20.
- Frank, M., Goodman, N. D., & Tenenbaum, J. B. (2007b). A Bayesian framework for cross-situational word-learning. In *Advances in Neural Information Processing Systems 20*.
- Frege, G. (1892). Über sinn und bedeutung. *Wittgenstein Studien*, 1(1).
- Gallistel, C., & Gelman, R. (1992). Preverbal and verbal counting and computation. *Cognition*, 44, 43–74.
- Gelman, R., & Gallistel, C. (1978). *The Child's Understanding of Number*. Cambridge, MA: Harvard University Press.
- Gibson, E., & Wexler, K. (1994). Triggers. *Linguistic Inquiry*, 25(3), 407–454.
- Goodman, N., Tenenbaum, J., Feldman, J., & Griffiths, T. (2008). A Rational Analysis of Rule-Based Concept Learning. *Cognitive Science*, 32(1), 108–154.
- Goodman, N., Tenenbaum, J., Griffiths, T., & Feldman, J. (2008). Compositionality in rational analysis: Grammar-based induction for concept learning. *The probabilistic mind: prospects for Bayesian cognitive science*, 377.
- Goodman, N., Ullman, T., & Tenenbaum, J. (2009). Learning a theory of causality. In *Proceedings of the 31st annual conference of the cognitive science society* (pp. 2188–2193).
- Hale, M., & Reiss, C. (2003). The Subset Principle in phonology: why the tabula can't be rasa. *Journal of Linguistics*, 39(02), 219–244.
- Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Malden, MA: Wiley-Blackwell.
- Hindley, J., & Seldin, J. (1986). *Introduction to combinators and λ -calculus*. Cambridge, UK: Press Syndicate of the University of Cambridge.
- Katz, Y., Goodman, N., Kersting, K., Kemp, C., & Tenenbaum, J. (2008). Modeling semantic cognition as logical dimensionality reduction. In *Proceedings of Thirtieth Annual Meeting of the Cognitive Science Society*.
- Keenan, E., & Stavi, J. (1986). A semantic characterization of natural language determiners. *Linguistics and philosophy*, 9(3), 253–326.
- Keenan, E., & Westerståhl, D. (1997). Generalized quantifiers in linguistics and logic. *Handbook of logic and language*, 837–893.
- Kemp, C., Goodman, N., & Tenenbaum, J. (2008). Learning and using relational theories. *Advances in neural information processing systems*, 20, 753–760.
- Lee, M., & Sarnecka, B. (2010a). A Model of Knower-Level Behavior in Number Concept Development. *Cognitive Science*, 34(1), 51–67.

- Lee, M., & Sarnecka, B. (2010b). Number-knower levels in young children: Insights from Bayesian modeling. *Cognition*, 120(3), 391–402.
- Leslie, A. M., Gelman, R., & Gallistel, C. (2008). The generative basis of natural number concepts. *Trends in Cognitive Sciences*, 12, 213–218.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W.H. Freeman & Company.
- Montague, R. (1973). The Proper Treatment of Quantification in Ordinary English. *Formal Semantics*, 17–34.
- O'Donnell, T., Tenenbaum, J., & Goodman, N. (2009). *Fragment Grammars: Exploring Computation and Reuse in Language* (Tech. Rep.). MIT Computer Science and Artificial Intelligence Laboratory Technical Report Series, MIT-CSAIL-TR-2009-013.
- Perfors, A., Tenenbaum, J., & Regier, T. (2011). The learnability of abstract syntactic principles. *Cognition*, 118.
- Philip, W. (1995). *Event Quantification in the Acquisition of Universal Quantification*. Ph.D. thesis, University of Massachusetts, Amherst.
- Piantadosi, S., Goodman, N., Ellis, B., & Tenenbaum, J. (2008). A Bayesian model of the acquisition of compositional semantics. In *Proceedings of the Thirtieth Annual Conference of the Cognitive Science Society*.
- Rips, L., Asmuth, J., & Bloomfield, A. (2006). Giving the boot to the bootstrap: How not to learn the natural numbers. *Cognition*, 101, 51–60.
- Rips, L., Asmuth, J., & Bloomfield, A. (2008). Do children learn the integers by induction? *Cognition*, 106, 940–951.
- Rips, L., Bloomfield, A., & Asmuth, J. (2008). From numerical concepts to concepts of number. *Behavioral and Brain Sciences*, 31, 623–642.
- Roeper, T., Strauss, U., & Pearson, B. (2004). The acquisition path of quantifiers: Two kinds of spreading. *Current Issues in Language Acquisition, UMOP*, 34.
- Rumelhart, D., & McClelland, J. (1986). *Parallel distributed processing*. Cambridge, MA: MIT Press.
- Sarnecka, B., & Lee, M. (2009). Levels of number knowledge during early childhood. *Journal of Experimental Child Psychology*, 103, 325–337.
- Shepard, R., Hovland, C., & Jenkins, H. (1961). Learning and memorization of classifications. *Psychological monographs*.
- Smolensky, P. (1996). *The initial state and 'richness of the base' in Optimality Theory* (Tech. Rep.). Johns Hopkins University, Department of Cognitive Science. JHU-CogSci-96-4.
- Smolensky, P., & Legendre, G. (2006). *The Harmonic Mind*. Cambridge, MA: MIT Press.
- Spelke, E. (2003). What makes us smart? Core knowledge and natural language. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in Mind*. Cambridge, MA: MIT Press.
- Spelke, E. (2004). Core Knowledge. In N. Kanwisher & J. Duncan (Eds.), *Attention and performance, vol. 20: Functional neuroimaging of visual cognition*. Oxford: Oxford University Press.
- Spelke, E., & Kinzler, K. (2007). Core knowledge. *Developmental Science*, 10(1), 89–96.
- Steedman, M. (2000). *The syntactic process* (Vol. 131). Cambridge, MA: MIT Press.
- Tenenbaum, J. (1999). *A Bayesian Framework for Concept Learning*. Ph.D. thesis, Massachusetts Institute of Technology.
- Tiede, H. (1999). Identifiability in the limit of context-free generalized quantifiers. *Journal of Language and Computation*, 1(1), 93–102.
- Ullman, T., Goodman, N., & Tenenbaum, J. (2010). Theory Acquisition as Stochastic Search. In *Proceedings of thirty second annual meeting of the cognitive science society*.
- Wexler, K. (2011). Maximal trouble: cues don't explain learning. In E. Gibson & N. Pearlmuter (Eds.), *The Processing and Acquisition of Reference*. Cambridge, MA: MIT Press.
- Wexler, K., & Manzini, M. (1987). Parameters and learnability in binding theory. *Parameter setting*, 41–76.
- Wynn, K. (1990). Children's understanding of counting. *Cognition*, 36, 155–193.
- Wynn, K. (1992). Children's Acquisition of the Number Words and the Counting System. *Cognitive Psychology*, 24, 220–251.
- Xu, F., & Tenenbaum, J. (2007). Word learning as Bayesian inference. *Psychological Review*, 114(2), 245–272.
- Yu, C., & Ballard, D. (2007). A unified model of early word learning: Integrating statistical and social cues. *Neurocomputing*, 70(13-15), 2149–2165.
- Zettlemoyer, L. S., & Collins, M. (2005). Learning to Map Sentences to Logical Form: Structured Classification with Probabilistic Categorical Grammars. In *UAI* (pp. 658–666).

Zettlemoyer, L. S., & Collins, M. (2007). Online Learning of Relaxed CCG Grammars for Parsing to Logical Form. In *EMNLP-CoNLL* (pp. 678–687).