

John Benjamins Publishing Company



This is a contribution from *Language in Interaction. Studies in honor of Eve V. Clark*.
Edited by Inbal Arnon, Marisa Casillas, Chigusa Kurumada and Bruno Estigarribia.
© 2014. John Benjamins Publishing Company

This electronic file may not be altered in any way.

The author(s) of this article is/are permitted to use this PDF file to generate printed copies to be used by way of offprints, for their personal use only.

Permission is granted by the publishers to post this file on a closed server which is accessible to members (students and staff) only of the author's/s' institute, it is not permitted to post this PDF on the open internet.

For any other use of this material prior written permission should be obtained from the publishers or through the Copyright Clearance Center (for USA: www.copyright.com).

Please contact rights@benjamins.nl or consult our website: www.benjamins.com

Tables of Contents, abstracts and guidelines are available at www.benjamins.com

Learning words through probabilistic inferences about speakers' communicative intentions

Michael C. Frank

Department of Psychology, Stanford University

How do children learn the meanings of words? This chapter presents a probabilistic, communicative view of word learning that synthesizes insights from work on statistical learning and social learning. By describing the formal characteristics of models, it is possible to differentiate communicative models that make inferences about the speaker's intentions from associative models that treat social information as a signal of salience. In addition, the probabilistic communicative framework can be integrated with models of pragmatic reasoning, leading to insights into how Gricean principles can facilitate word learning.

Acknowledgements: Thanks to the editors of this volume for the opportunity to contribute, to Dave Kleinschmidt, Molly Lewis, Dan Yurovsky, Chigusa Kurumada, and an anonymous reviewer for helpful comments, and also to Noah Goodman, my collaborator in much of the work described here.

Introduction

The linguistic world of young children is likely an overwhelming place. Even if they are not assaulted with James' (1890) "blooming buzzing confusion," it must be perplexing for infants to be surrounded constantly by sounds whose only interpretable meaning at first comes from the tone of voice in which they are uttered. Perhaps another metaphor is more apt: The infant is a traveler trying to negotiate a task in a foreign language, with the help of a sympathetic interlocutor (a parent or caregiver). The child and caregiver may share goals, or at least an understanding of the other's goals. But without knowing any words, only a few utterances can be decoded from context; and the fewer words that are known, the less leverage the

child has to infer the meanings of others. Even in the highly supportive contexts created by parents, some substantial portion of spoken language is likely to be incomprehensible to young infants.

How then do children begin to break into the vocabulary of their first language? Two broad proposals run throughout work on word learning, from historical sources to contemporary models: associative and intentional proposals. In associative accounts from Locke (1690/1964) onward, infants are hypothesized to match elements of their linguistic environment with the world around them, identifying the consistent mappings between words and other stimuli. In intentional accounts from St. Augustine (397/1963) onwards, in contrast, the attention of learners is on the speakers who produce words. These words are then mapped to the speakers' intended meanings – often instantiated by the speakers' intended referent in the current context. Although distinct evidence is often cited in support for one view or the other (e.g., Baldwin, 1993; Smith & Yu, 2008), these two views need not be in conflict. I will argue here that they can be integrated in a single framework.

I did not begin studying language acquisition with this view. When I entered graduate school, I was convinced that all of language acquisition, especially word learning, could be reduced to a process of pure statistical inference. This idea was more probabilistic than the classic generative story for syntactic acquisition, but otherwise quite similar: The child would passively observe some corpus of input. She would then carry out a defined learning procedure over this corpus, resulting in a grammar of the target language. Eve Clark's "First Language Acquisition" (2003) played a crucial role in my departure from this purely statistical viewpoint.

Reading Eve's account of early language development, illustrated with examples from the speech of actual children, made a substantial impression on me. Despite my interest in acquisition, and my presence at the same institution as Eve, I had unfortunately missed the opportunity to take her course on the topic. So it was an awakening to read how she wrote about children's learning situation – as consisting of a context, a communicative goal, and a creative drive to use language to accomplish that goal. This vision of children playing, exploring, and interacting at the same time as they learned also lined up nicely with new work suggesting this same sort of process is at work in children's causal learning (e.g., Schulz, Kushnir, & Gopnik, 2007). But most importantly, it felt true to the actual children I had met.

My goal in this chapter is to provide an overview of the broader framework that has emerged from my work on language learning in a communicative context. This framework is in some deep sense a computational instantiation of the view described in Eve's writings: It is fundamentally communicative in that it is oriented around representations of speakers' goals and intentions. But a further strength – I believe – is that it also captures a substantial amount of the core value of associative

accounts: In particular, it takes advantage of the graded, probabilistic nature of learning to aggregate information across multiple learning situations.

The probabilistic communicative framework fits well with our emerging understanding of children's cognitive development. The intentional aspects of the framework are congruent with research on children's social cognition (e.g., Csibra & Gergely, 2009; Onishi & Baillargeon, 2005; Vouloumanos, Onishi, & Pogue, 2012) and the framework supports graded probabilistic learning of the sort that is attested across development (Gopnik, 2012; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). Most importantly, however, the work I describe here is grounded in the literature on pragmatic inference (H. Clark, 1996; Grice, 1975) and its role in language acquisition (Bloom, 2002; E. Clark, 2003; Tomasello, 2003). In this chapter, for reasons of space I will focus primarily on the computational literature with only occasional references to empirical issues, but I hope the debt is clear.

I use the word "framework" in this chapter rather than "model" because a framework to me refers to a set of principles, while a model refers to a particular instantiation of those principles within a working, implemented system from which concrete, quantitative predictions can be made. Although my collaborators and I have made a number of such models (which I reference here), all are specific to particular applications or kinds of data. I believe that there are some more general conclusions that can be drawn from these individual systems when they are examined together, however; these are the conclusions I focus on in this chapter.

As I describe this framework, I will also try to lay out some theoretical distinctions in more depth than is allowed in a strictly empirical report. I will first describe a taxonomy of computational models of word learning. In the next section, I will specify formal details corresponding to this taxonomy. The goal of this section will be to distinguish between those cross-situational models that merely use social information and those that make an intentional assumption about the nature of the learning situation. Finally, the last section will focus on the links between our intentional framework for word learning and our work on modeling pragmatic inference.

A note on terminology. When discussing word learning, I will be discussing how children learn open-class words like "rabbit," "white," or "run" rather than closed-class words like "the" or "no." The problem of inferring meanings for open-class words can be broken into two separate tasks: the task of words to referents in the moment (reference assignment), and the task of identifying the meaning or concept corresponding to the indicated referent – or to the aspect or property of that referent that the speaker wishes to highlight (concept learning). The distinction between these two problems is clear, but our terms for discussing them are often clumsy. The issue is often confused further by the use of simple noun learning as a case study, since solving the reference problem in that case comes close to

giving away the concept-learning problem (modulo a bias for basic-level categories; Markman, 1991). I'll be focusing initially on word-object mapping but I will note when the equivalence between mapping and concept learning is broken.

A taxonomy of models of word learning

The strategy of learning words via co-occurrence has been labeled “cross-situational” word learning. Recent empirical work provides strong support for the idea that both adults and infants can learn word-object mappings by gathering consistent associations across multiple, ambiguous exposures (Smith & Yu, 2008; Vouloumanos, 2008; Vouloumanos & Werker, 2009; Yu & Ballard, 2007). A variety of theorists have asked about the utility of cross-situational learning in acquiring words of different types (Akhtar & Montague, 1999; Fisher, Hall, Rakowitz, & Gleitman, 1994; Gillette, Gleitman, Gleitman, & Lederer, 1999; Gleitman, 1990; Pinker, 1984). Though theoretically all are learnable (Siskind, 1996), there is likely to be a continuum from those words most learnable by co-occurrence, e.g., nouns and some property terms, to those least likely to be inferred from context alone. This latter category likely includes verbs – which refer to events that have multiple construals (e.g., “chase”/“flee”) – context-dependent adjectives, and function words. Because of the relative simplicity of learning basic-level object nouns from context, I begin by discussing word-object mapping models.

A taxonomy of models of word-object mapping is shown in Figure 1. All of these models are “cross-situational” in the sense that all consider evidence about the relationship between words and parts of the world across multiple observations. All of them also attempt to learn a lexicon: a set of consistent word-object mappings. They differ, however, in both the information sources they consider and the ways they use these information sources. Each of these models describes a set of variables, both observed (shaded) and unobserved (unshaded) as well as a set of causal dependencies between these variables. These dependencies define a generative process: a set of steps by which the learner assumes the observed data have been generated. Unobserved parts of this generative process can then be estimated using standard inference techniques.¹

In our taxonomy, we differentiate models on two dimensions. The first – represented by the vertical axis in Figure 1 – is the information sources considered by the model. “Pure” cross-situational models consider only the co-occurrence between words and objects in establishing links between words and objects in the

1. Perfors, Tenenbaum, Griffiths, and Xu (2011) give a very nice tutorial introduction to this general style of modeling as applied to developmental questions.

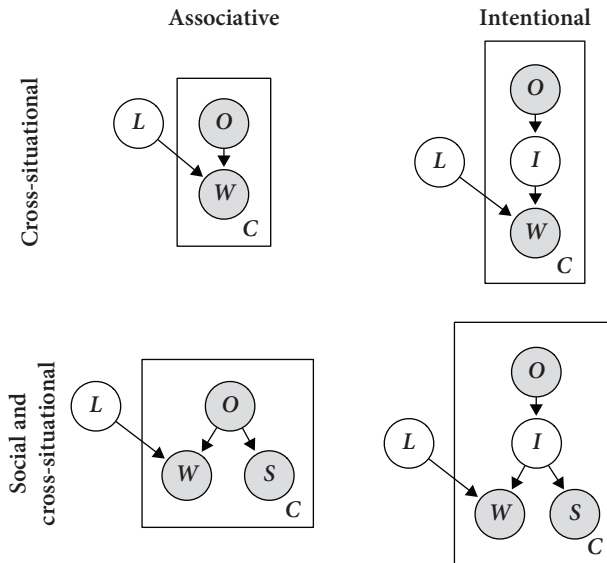


Figure 1. A progression of possible models of the basic challenge faced by early word learners. *L* is the child's lexicon, *C* refers to the contexts in which utterances containing words *W* are observed, accompanied by object referents *O* and social cues *S*, as well as (unobserved) communicative intentions *I*.

lexicon. Social and cross-situational models consider also a set of social cues (envisioned here as signals like a point, a gesture, or a gaze towards an object, with some temporal connection to a particular utterance).

The second dimension is the way models represent the relationship between words and the world – shown on the horizontal axis in Figure 1. Associative models describe a generative process in which words are assumed to be generated directly by the presence of their associated objects (without the presence of an intervening speaker). In contrast, intentional models, as we define them, assume that words are generated via the intention of a speaker to refer to an object. Thus, in the sense used by theorists of social cognition, intentional models are fundamentally triadic: they define a relationship between the child, the speaker, and objects in the context (Baldwin, 1995; Carpenter, Nagell, & Tomasello, 1998).

In the current analysis, models are ideal observers: realizations of the assumptions that learners make and information sources they use, rather than the mechanisms by which these assumptions and information sources are processed (Geisler, 2003). Another way of putting this is that our analyses are at the computational, rather than algorithmic level, describing models as they represent the task faced by the child rather than as the child solves them (Marr, 1982).

There is a rich literature describing how such algorithmic constraints should be implemented in word learning models (Fazly, Alishahi, & Stevenson, 2010; Yu & Smith, 2012), and my own work in other domains has investigated these question as well (Frank & Gibson, 2011; Frank, Goldwater, Griffiths, & Tenenbaum, 2010). In addition, there is a fascinating and growing literature investigating the ways that resource-bounded decision-making can relate to normative models (see e.g., Sanborn, Griffiths, and Navarro, 2010 for review and discussion). Yet I worry that implementational considerations often mask (rather than reveal) the learner's underlying assumptions about the learning situation, and I believe that there is value in considering the computational level independent from – and perhaps in parallel to – the algorithmic.

Related to this issue is a set of recent criticisms of the idea that cross-situational observation is a factor in learning word-object mappings (Medina, Snedeker, Trueswell, & Gleitman, 2011; Trueswell, Medina, Hafri, & Gleitman, 2013). The key empirical question underlying these criticisms is whether individual word learners represent multiple hypotheses about the meanings of words. The proposed alternative is that individuals make noisy, stochastic choices of individual hypotheses that they then test against data (a “propose but verify” strategy). The average of many such stochastic choices across items and individuals would then produce the pattern of gradual learning that is sensitive to the degree of cross-situational ambiguity that was observed in previous studies. This criticism is rooted in a long-standing debate about the nature of learning more generally, and whether it is typically gradual and associative in character or discrete, and hypothesis-based (Gallistel, 1990; Gallistel, Fairhurst, & Balsam, 2004).

Such concerns relate primarily to the mechanisms of learning applied by the learner in recovering the lexicon of their language, rather than to the underlying assumptions that guide this learning. Even if “propose but verify” learners do not retain previous data to test their hypotheses, they still must make the assumption that their hypotheses should in principle be consistent with previous data as well as future observations. In other words, even if learners are memoryless, their behavior is still consistent with an assumption of cross-situational statistical consistency. Therefore, in this chapter I will not discuss an important body of modeling work that investigates the consequences of algorithmic details of representation for modeling human performance in word learning (e.g., Li, Farkas, & MacWhinney, 2004; Regier, 2005; Yu & Smith, 2012).

In the next section, I describe computational details underlying the models in Figure 1. Providing these details allows for clarity about how an intentional assumption can be implemented, and additionally allows for comparison and integration with the models of Gricean pragmatics described below.

A formal framework for cross-situational learning

Formalization allows us to consider what has previously been a somewhat slippery distinction: between *social cross-situational models*, which consider information generated by other people, and *intentional models*, which are based on a stronger assumption about the generating source of this information. Following the taxonomy in Shafto et al. (2012), I take the fundamental assumption underlying a communicative model to be that language is produced as a rational action to accomplish a goal.² The “language as rational action” assumption allows for stronger inferences from data than those possible under a view that considers social data but does not consider the communicative intentions that lead to those data being generated. This section walks through the formal mechanics of these ideas, describing the basics of cross-situational learning under the general family of models, laying out the key differences between associative and intentional models, and describing how social information is used in associative and intentional models.

Basic cross-situational learning

Consider a schematic description of the child's learning problem, shown in Figure 2. The child finds herself in a set of contexts C . Each of these contexts contains possible referents $O_C = o_1 \dots o_n$. In each of these contexts, a speaker has an intention I_C , unobserved to the child, that captures the idea that the speaker would like to convey (described in greater detail below – at this point, the speaker's intention is simply a placeholder). On the basis of this communicative intention, the speaker utters words $W_C = W_1 \dots W_n$ and produces social gestures S_C . Having observed a set of these kinds of contexts, the learner's goal is to infer a set of correspondences between words and objects, which we denote L (the lexicon). This lexicon is assumed to be stable across contexts and individuals,³ and can be modeled either

2. Of course, the idea of language as a rational action does not originate in the computational literature; for example, Grice (1975) writes that “... one of my avowed aims is to see talking as a special case or variety of purposive, indeed rational, behavior...” (p. 47). Clark (1988, 1990) then describes the linkage of this idea to the acquisition context.

3. Does the child know that her lexicon should be stable across time and identical across individuals? There are several such foundational assumptions that are necessary for all of the models described in this chapter. For example, we assume that words are linked to a particular level of description (objects or object concepts in the models we consider), and not to some others (e.g., motor actions, other sensory stimuli). The basic learning frameworks we describe could in principle be applied to a scenario with many more targets for words, no stability of word meaning across time, or vast individual differences in language use (whether due to bilingualism or

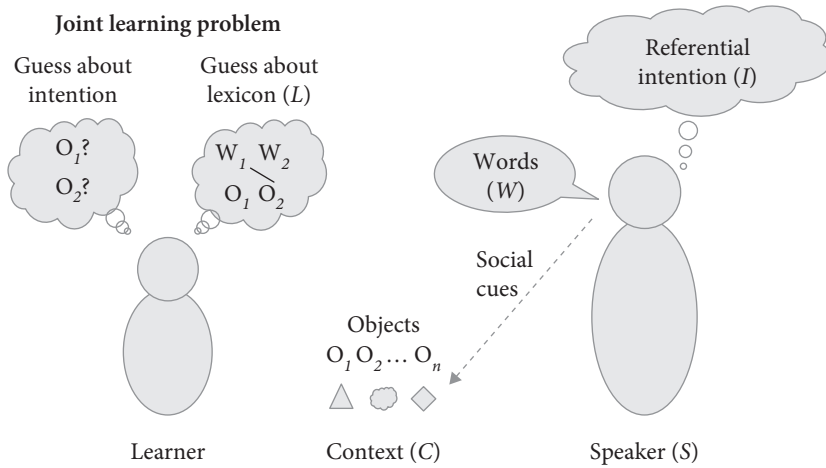


Figure 2. A schematic view of the intentional framework, linking learner and speaker via the two learning problems: guessing the speaker’s intended referent and guessing the lexicon of the language.

as a set of discrete links or continuous, probabilistic associations. The challenge for the child is that each individual context does not uniquely determine a set of lexical mappings. Which words go with which objects?

We can notate this problem of lexicon learning as a problem of Bayesian inference, that is, of inferring the most probable lexicon given the set of observed contexts:

$$P(L|C) \propto P(C|L)P(L) \tag{1}$$

The taxonomy in Figure 1 provides generative processes for four kinds of models that have been applied to this learning problem. The simplest approach to this problem is simply to estimate these probabilities, neglecting intentions or social cues. The model in the upper left corner, which assumes that words are generated via the observed objects in the context and the unobserved lexicon, can be written

$$P(L|C) \propto P(W|O, L)P(L). \tag{2}$$

If we represent the words in a context as the set W_C and the objects as O_C , we can expand this expression to

even simply random variation), but there is no guarantee that they would be sufficient. In practice, whether these assumptions are inborn or discovered, it is likely that they are necessary for learning to proceed.

$$P(L|C) \propto \prod_c P(W_c|O_c, L)P(L). \quad (3)$$

In other words, the probability of the lexicon under these models is the product across contexts of the probability of the words in the context, given the objects in the context and the lexicon. This “pure cross situational” approach is followed by a number of influential models (Fazly et al., 2010; Yu & Ballard, 2007).⁴

Differentiating associative and intentional models

In this section, we compare the assumptions made by associative and intentional models, differentiating between the upper left and upper right panels of Figure 1. Following the approach above, the next step in defining a word learning model is to define the likelihood of a word being uttered, given the presence of some object and its lexical entry: the term $P(W_c|O_c, L)$. There are two important sub-problems that arise in defining this term. First, we must define the alignment between particular words and objects (assuming that there are multiples of each in each situation). Second, we need to define the probability of a word being used with a particular object (given that they are aligned). I will discuss only the first of these here, since the details of assigning probabilities are covered in several of the source articles for these models.⁵

It is with respect to the problem of aligning words and objects that the differences between associative and intentional models becomes clear. Associative models typically make minimal assumptions about alignment and assume that any

4. This modeling work has built directly on work on machine translation that attempted this optimization problem for aligned corpora (e.g. where words and objects are actually words in a target language and words in a source language; Brown, Pietra, Pietra, & Mercer, 1993).

5. The intentional assumption (the focus of this section) is orthogonal to the precise mechanics of how a model assigns probabilities to particular lexical mappings. Our initial work used a fairly discrete likelihood function to determine the probability that a particular word was used, given that the speaker had an intention to refer to some object:

$$P(w|I=o_c, L) = \begin{cases} \propto 1 & \text{if } L(w_c, o_c) = 1 \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where $L(w_c, o_c) = 1$ indicated that w_c and o_c were linked in the lexicon. But it is equally possible to define a more clearly probabilistic function, using e.g. a multinomial distribution (which would then be conjugate to a dirichlet prior). More generally, it is likely very difficult to differentiate between a continuous lexical representation and a posterior distribution representing uncertainty over a discrete lexicon. Although our work on this topic is occasionally cited as providing a “hypothesis testing” (discrete) view of word learning, I see the discreteness of the lexicon in that particular model as an implementation choice rather than one that carries any particular theoretical weight.

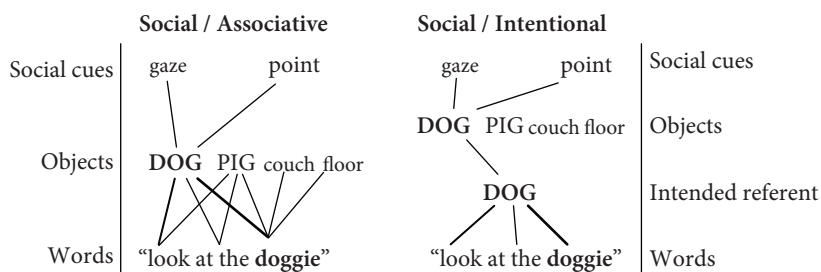


Figure 3. A schematic of a single situation for associative and communicative models that both use social and prosodic information. Gaze and pointing cues signal that a dog toy is more salient than a pig toy. Prosodic focus on the word “doggie” raises its salience. As a consequence, the strongest association is between the object dog and the word “doggie” for both models. The communicative model includes a filtering step in which dog is assumed to be the correct referent. Salience is shown by type weight and size, while associative weights are shown by line weight.

word can be aligned with any object. In some sense, this is the basic tenet of an associative model: that all words and objects present in the context are associated with one another to some degree.

In an intentional model, in contrast, we assume that the relationship between words and objects is mediated by the speaker’s intention to refer to some set of objects. This mediation relationship is shown graphically by the intervening node between O and W in the generative process for models on the right side of Figure 1. The concept of the speaker’s intention to communicate mediates the relationship between the physical context and the words produced by that speaker. This notion of intention in our framework also corresponds to the notion of an agent’s goal, which plays a central role in work on social learning and rational action inference (Baker, Saxe, & Tenenbaum, 2009; Gergely, Bekkering, & Király, 2002; Shafto, et al., 2012). This intuition is shown graphically in Figure 3, which shows the mediating relationship that the speaker’s intention can play in learning from a single situation.

Formally, this mediation relationship results in a revision to Equation 2, where we notate this mediating intention I_C :

$$P(L|C) \propto P(W_C|O_C I_C L)P(L). \quad (5)$$

The addition of this mediating variable affects the process of finding the alignment between words and objects: While associative models assume that all words are linked to all objects, the intentional models assume that there is an extra step that removes some of these associations from consideration. In our work on cross-situational learning to date, the representation of the speaker’s intention has been quite basic, representing the speaker’s menu of possible intentions as the set of objects in the

context (making these *referential intentions*).⁶ This restriction implicitly implemented a “whole object” assumption (Markman, 1991); but we reconsider this restriction in more recent work (described in the section on pragmatic inference, below).

Formally, in our initial model, we specified I_C as containing a subset of O_C , corresponding to the assumption that the speaker could talk about any subset of the objects in the context, including the empty set (Frank, Goodman, & Tenenbaum, 2009). Although this assumption allowed us to consider a wide variety of possibilities, if there were many objects present in a context it quickly became unwieldy, since it required considering the power set of objects in the context (which grows at 2^n where n is the number of objects in c). Hence, in more recent work we have begun using the simplifying assumption that I_C is a single object in O_C (Johnson, Demuth, & Frank, 2012), congruent with our empirical observation that most of the caregivers' utterances in a corpus referred to at most one object.

Under an intentional assumption, regardless of how the intended referent is chosen, we can define the likelihood of a word as the product of two terms: the probability of the words given the intention, and the probability of a particular intention:

$$P(W_c | O_c, I_c, L) = \sum_{I_c \in O_c} P(W_c | I_c, L) P(I_c | O_c). \quad (6)$$

Thus, intentional models describes a two-step process of uttering a word: first decide which object to refer to, then decide the words to use to refer to it.

The key difference between associative and intentional models on our account is this two-step process, separating the choice of what to talk about from the choice of how to refer. Both types of models allow for information to “weight” the learner's estimate of which objects are most salient, for example via social information. But in intentional models, this weighting influences the learner's guess about which object(s) are being referred to. The definitional assumption of intentional models is that speakers have a discrete intention for each utterance, even if the learner is uncertain of what this intention is. In contrast, in associative models, all words are associated with all objects. Implicitly, this assumption is tantamount to assuming that all objects are being referred to and all words have referential status (just to greater or lesser degrees).

The intentional assumption – that there is a discrete choice of intended referent by a speaker – implies that one major part of word learning is in-the-moment

6. This construct has the potential to be far more flexible and powerful than the use to which we have so far put it. Provided that this distribution over intentions is limited by context, discourse, and other pragmatic factors, models could consider a much wider variety of possible interpretations – including interpretations not involving grounded reference to the current context.

interpretation (as posited in many pragmatic accounts of language learning, especially E. Clark, 2003 and Tomasello, 2003). If a learner knows what object is being talked about by the speaker (their intention, in our loose terminology), there is no need to compute associations between the words that are heard and the other objects that are present. In the language of causal models, the intention “screens off” the physical context from the words: knowing the speaker’s intended referent is enough for learning. In this respect, the intentional framework is deeply related to recent computational work by McMurray, Horst, and Samuelson (2012), who emphasized the role of learners’ interpretations of reference in the moment in longer-term word learning.

In Frank, Goodman, and Tenenbaum (2009), we reported results based on running associative and intentional models on a small, hand-annotated corpus of infant-directed speech. The intentional model outperformed other models in the lexicon it learned. This success was due at least in part to its ability to “filter out” spurious associations between function words and objects. Under the intentional model, evidence that a word was mapped to an object in one situation could help constrain hypotheses about which object was being referred to in another situation. This mutual constraint meant that irrelevant co-occurrences could be explained away, rather than providing noise (as in the associative models).

Adding social information

It is relatively rare for a speaker to talk about a physically present referent without giving some signal – at least at some point in the conversation – that the referent is indeed the one being talked about. Speakers gaze to conversational referents during language production both as an explicit social signal (H. Clark, 1996) and as a consequence of processes underlying language production (Griffin & Bock, 2000). In addition, speakers often signal reference by pointing. “Social cues” like eye-gaze or points to conversational referents are often cited as an important source of evidence for learning word-object mappings (Baldwin, 1993; Bloom, 2002; Hollich, Hirsh-Pasek, & Golinkoff, 2000; St. Augustine, 397/1963).

Social cues can be represented in an associative model as cues to which objects are most salient in a particular situation (Yu & Ballard, 2007). The presence of a point or gaze on a particular object endows it with some additional salience, which in turn strengthens its associations with co-occurring words. In the generative process for such a model, social cues are a reflection of the underlying salience of individual objects (Figure 1, lower left). If social cues are present, object salience is inferred to be higher. Figure 3, left side, shows a caricature of what the data for a single situation might look under such a model. The associative weights between words and objects are determined by the social cues and perceptual salience of the objects in the scene (as well as the prosodic salience of the words, which we do not

discuss here). The result is that the same associative computation is performed, but over a word/object set whose weights are no longer uniform.

In contrast, in an intentional model (Figure 1, lower right), social information informs the process of interpretation (deciding which object, if any, is being referred to). Social cues are generated by the speaker's intention to refer and hence are signals to that underlying intention (Figure 3, right). This interpretive, inferential use of social cues can be implemented in a number of ways. In our early work on this topic, we assumed that each cue could be the consequence of a relevant intentional action or could be the result of baseline looking without an underlying intention, and estimated these probabilities via a "noisy-or" model. This formulation allowed the model to learn that, for example, even though speakers' gaze was a frequent cue, its overall reliability was low (Frank, Goodman, & Tenenbaum, 2007).

Our more recent work embeds these social cue probabilities within a probabilistic, grammar-based formalism. In Johnson et al. (2012), we used an adaptor grammar, a probabilistic context free grammar that allows common structures to be reused efficiently (Johnson, Griffiths, & Goldwater, 2007). Consider the example shown in Figure 4. The observed input representation of the situation (shown in blue italics at the bottom of the tree) specifies that the words "where's the piggie" are observed along with dog and pig toys, and the pig toy is marked by two social cues, the child's eyes and the mom's hands.⁷ The tree above it shows a possible parse of the situation. The extraneous dog toy is parsed as "non-topical" (e.g., not an intended

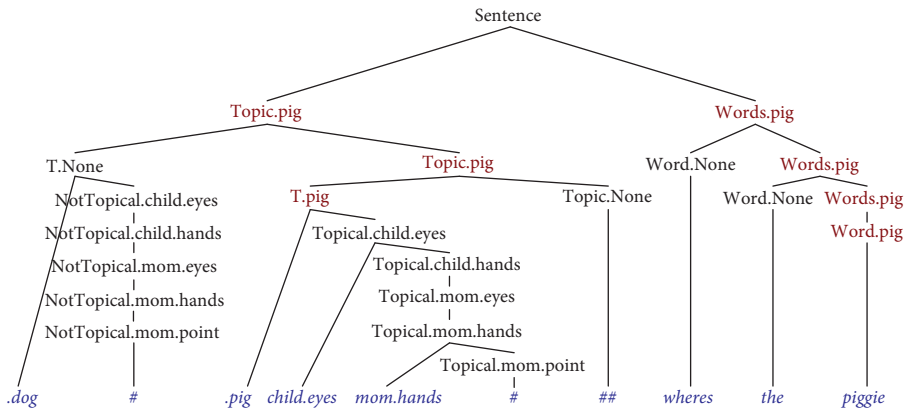


Figure 4. A parse tree for an entire situation, including a sentence along with its referential context and social cues. The sentence generates topical (referential) objects, the social cues that mark these objects, and the words that refer to them. Topic-specific words are marked in red and observed data are in blue. In this case, the referent is "pig" and referential words are propagated throughout the tree. Figure reprinted from Johnson et al. (2012).

7. We leave aside here the issue of whether the child's own eyes can be considered a "social cue" – this issue is discussed at length in Frank, Tenenbaum, and Fernald (2013).

referent, marked as *T.none*) and the pig toy and its accompanying social cues are generated on the basis of the topic. The words are also generated by the same topic, with several “non-topical” words (*Word.none*) followed by a word generated from the topical lexicon (marked as *Word.pig*).

This grammatical formalism, although distinct from our previous work in some of its details, still encodes the same two-stage computation we have associated with intentional models. The first decision is choosing the intention (topic, in the language of this model) underlying a sentence. This decision affects all other aspects of the sentence including the probabilities of both the social cues and the individual words, whose choice together constitutes the second decision: how to refer. The only difference is that these two “stages” are not broken out as separate computations but instead are instantiated as parts of the model’s grammatical representation. The intention or topic is represented through the topical productions in the parse tree, while the individual decisions about words in the sentence are represented through the choices of leaf nodes in the tree.

When we evaluated the grammatical model on a corpus tagged with social information, we found significant gains in the accuracy of both guessing the referents of utterances and the words associated with particular objects on the basis of adding the social cues. In addition, our analyses showed that the child’s own gaze on an object was the most predictive cue, suggesting that the corpus we used contained significant follow-in labeling by parents (replicating descriptive results from Frank et al., 2013).

To summarize: both associative and intentional models allow for the inclusion of social information, but associative models allow for social information to make particular referents more salient and bias the computation of associations between words and objects. In contrast, intentional models go beyond this interpretation of social cues as signals of salience and allow the social information to bias the computation of reference. Yet all of the models described here still treat only the mapping problem, implicitly equating referent identification with meaning learning. In this next section, we broaden the set of possible word meanings we consider, beginning the process of differentiating word meaning from reference.

Adding pragmatic inference to intentional models

In the Quinian (1960) framing of the word learning problem, even if a single word is heard alongside a single object, there are still an infinite number of possible interpretations for the word. Quine distinguished between interpretations for which co-occurrence or pointing – the information sources relied on in the models above – failed to provide any traction for learners and those for which these cues might

in principle be informative. This first class includes, for example, the ambiguity between “rabbit” and “undetached rabbit parts”; Quine argues that these meanings may be impossible to distinguish. On the other hand, the second class contains many word meanings that may be empirically distinguishable but are likely to be confounded in any given context. To take a small set of the many possible modes of reference, a particular rabbit might be talked about as a “rabbit,” but also as “white,” (when the contrast is a brown rabbit), “animal” (when pointing out something in the bushes), or “small” (when the contrast is a larger rabbit).

Since basic-level object names are common in speech to children (Callanan, 1985), it should not be difficult to learn a word like “rabbit” from co-occurrence. The prospects for noticing the co-occurrence between animacy and “animal” seem somewhat lower; they are likely lower still for a color like “white,” and close to nil for a gradable adjective like “small” (whose meaning changes from context to context). Information about the context of reference might provide a far more straightforward path to learning such terms (perhaps along with syntactic information, in the case of adjectives; Waxman & Booth, 2001). In this situation, a pragmatic word learner has an important advantage over any of the cross-situational learners described above: She can consider the context of use and the goal of the speaker in uttering a particular phrase, and crucially, she can consider why a term contrasting with the conventional descriptor is used (E. Clark, 1988).⁸ This is the intuition that I will follow in this last section.

On a standard cross-situational view, pragmatic inferences belong to an entirely different class than the associative inference that leads the child to consider that “rabbit” = RABBIT. Our communicative/intentional framework provides a way to integrate these two inferences, however. Because long-term learning is mediated by in-the-moment interpretation, learners can use pragmatic computations to inform their guesses about what words refer and what objects (and even aspects of these objects) are being referred to. In this section I will show an example of our recent work modeling pragmatic inference in context, and then demonstrate that this kind of model of pragmatic inference can be integrated with the intentional word learning framework described above. There is much work yet to do, but I believe that this combined framework provides the beginnings of the tools necessary to extend the cross-situational paradigm beyond simple word-object mapping.

8. Critically, in the models described here, there is no requirement that words *must* be learned by one route or another. Much of the vocabulary that children acquire is no doubt learned osensively or in socially-constrained situations that do not require inference. An important goal for future research is to quantify the contributions of different learning mechanisms to the growth of children’s vocabulary (see e.g., Mitchell & McMurray, 2009 for an example of this kind of “macro-economic” analysis).

Modeling pragmatic disambiguation of reference

Grice (1975) proposed a set of maxims for normative communication: Speakers should be truthful, relevant, clear, and informative. On Grice's account, listeners should interpret utterances as though these maxims are being followed, allowing them to go beyond the truth-functional meanings of the words in the sentences to derive richer meanings in context that he called "implicatures." In our work we have explored the idea that some of these implicatures can be captured in formal models (Frank & Goodman, 2012). We focus here on the maxim of informativeness and model listeners as doing statistical inferences about what a speaker's intended referent is, given a presumption of informativeness.

Our model assumes that in a context C , the listener is attempting to infer the speaker's intention I as one of a set of possible referents.⁹ The listener considers two factors: first, the relative informativeness of the speaker's utterance with respect to each of the referents, and second, the "contextual salience" of the referents. Contextual salience here refers to the relative probability of reference, given the conversational context, the shared knowledge between communicators, and other factors that jointly determine that a particular referent will be the object of the speaker's expression. Although this quantity could perhaps be derived *a priori*, we have measured it empirically in our previous work (Frank & Goodman, 2012).

This two-part model can be notated as follows:

$$P(I|w, L, C) \propto P(w|I, L, C)P(I, C) \quad (7)$$

with the two terms on the right corresponding to the two factors being considered (informativeness and contextual salience). We pursue the idea that the informativeness of a word in context is inversely proportional to its specificity, so that

$$P(w|I, L, C) = \begin{cases} \propto \frac{1}{|w|} & \text{if } w(I)=1 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

where $|w|$ notates the number of objects in the context that can be referred to using w and $w(I) = 1$ if word w is true of I .¹⁰ This inverse proportionality corresponds to

9. Note that in our other work we use R_s to notate the speaker's intended referent, rather than I – we use I here for consistency with previous sections. As before, intentions are cashed out as referents, though again we believe that this framework could in principle be extended to consider intended *meanings*.

10. This set of definitions implicitly assumes that words are no longer names of objects. Instead, they are functions that can be applied to a context and that return true or false for each object in

the “size principle” of Tenenbaum & Griffiths (2001): A word is more informative if its use is “less coincidental” because it better picks out the intended referent from the context. Note here the parallel with Equation 4: In some sense, we are just replacing an uninformative speaker (where probability is proportional to 1) with an informative speaker (where probability is proportional to extension in context).

An example is useful in clarifying how this model setup naturally leads to Gricean pragmatic implicatures. Consider the situation pictured in Figure 5 (originally from Stiller, Goodman, & Frank, 2011). There are three possible referents shown, referred to below as NULL, G, and H + G. Each has different features, which lead to different possible expressions that can be used to refer to it. For the purpose of this game, we assume that this set is limited to the expressions “hat” and “glasses.” A speaker utters the word “glasses”; the job of the listener is to decide which face is being referred to. Many people share the intuition that the intended referent is G, the face without a hat.

This simple pragmatic inference has many of the elements of scalar implicature (the often-studied inference that “some” typically is strengthened pragmatically to mean “some but not all”), so it is a useful case study of how our model can be applied. In order to simplify the computation, we assume that the contextual salience of the three faces is even, and that the NULL referent is not considered. We can compute the strength of the pragmatic inference that “glasses” refers to G using Equation 7 and expanding the proportionality by normalizing over all possible referents:




Figure 5. An example stimulus from our pragmatic inference experiments. Participants would be asked to identify the referent of a phrase containing “glasses” as the descriptor; given this message, the middle face (who has glasses but no hat) is the pragmatic implicature target, whereas the right-hand face (who has a hat and hence has a better potential descriptor) is the literal target.

the context. This truth-functional model of word meaning can easily be used to capture predicates like “white” or “furry” and is in principle extensible to context-dependent adjectives like “small” (Schmidt, Goodman, Barner, & Tenenbaum, 2009).

$$\begin{aligned}
 P(G | \text{"glasses"}, C) &= \frac{P(\text{"glasses"} | G, C)}{\sum_{I' \in C} P(\text{"glasses"} | I', C)} \\
 &= \frac{P(\text{"glasses"} | G, C)}{P(\text{"glasses"} | G, C) + P(\text{"glasses"} | H + G, C)}
 \end{aligned}
 \tag{9}$$

We can then expand this expression, using Equation 8 and notating the set of vocabulary items that can be used to describe e.g., item G as $w \in G$:

$$\begin{aligned}
 P(G | \text{"glasses"}, C) &= \frac{\frac{1}{|\text{"glasses"}|}}{\sum_{w' \in G} \frac{1}{|w'|}} \\
 &= \frac{\frac{1}{|\text{"glasses"}|}}{\frac{1}{|\text{"glasses"}|} + \frac{1}{|\text{"glasses"}|}} \\
 &= \frac{\frac{1}{|\text{"glasses"}|}}{\frac{1}{|\text{"glasses"}|} + \frac{1}{|\text{"glasses"}|} + \frac{1}{|\text{"hat"}|}} \\
 &= \frac{\frac{1/2}{1/2 + 1/2}}{\frac{1/2}{1/2 + 1/2} + \frac{1/2}{1/2 + 1}} = .75
 \end{aligned}
 \tag{10}$$

In other words, the probability of G , the face with glasses and no hat, given the descriptor “glasses,” is predicted to be .75 (and hence the probability of $H + G$ is .25). This computation encodes the intuition that, had the speaker been talking about $H + G$, he or she would have chosen the more specific descriptor “hat.” This prediction of an implicature in favor of G corresponds well with the judgments of both adults and preschoolers (Stiller et al., 2011).

As illustrated above, our pragmatics model provides a framework for quantifying the Gricean maxim “be informative” from the perspective of both a speaker and a listener. It and its extensions can provide an account for a wide variety of pragmatic phenomena (Frank & Goodman, 2012; Goodman & Stuhlmüller, 2013).

The pragmatic implicature under this model can be computed for any situation in which the set of contextual and vocabulary alternatives is known, although the problem of identifying relevant alternative is still an open research challenge. In the next section, we illustrate how this pragmatic framework can be integrated with the intentional approach described above.

Using informativeness to learn words

Our pragmatics model is deeply related to the intentional communication model described above. Recall that in Equation 6 of the cross-situational learning model, we defined the probability of a particular word being uttered, given some intention and context as a product of two terms: the probability of the word given the intention, and the probability of the intended referent given the objects in the context. This relation was stated as $P(w|O, I, L) = P(w|I, L)P(I|O)$. Intuitively, these two terms govern the probability of choosing a particular referent and choosing the proper referring expression. Note now that the pragmatic model described above uses the same breakdown of the process of inferring reference. The “contextual salience” term we described above maps directly onto the referent choice term $P(I|O)$, and the Gricean informativeness term in Equation 8 maps onto the term $P(w|I, L)$.

With this equivalence in hand, we can reverse the pragmatics model and derive a word-learning version (a version of this derivation is given in Frank, Goodman, Lai, & Tenenbaum, 2009 and Frank & Goodman, under review). In this version, we infer alternative meanings for a particular lexical item, given that a particular referent is known to be uttered. Consider the display in Figure 5. Imagine that a speaker pointed now to the literal target $H + G$, fixing the referent, but uttered a novel label, e.g., “fedora.” In this case, the referent is known, but the meaning of a novel element in the lexicon L is unknown. We notate the possibility that a particular word in L refers to a feature (e.g., having a hat) as $w = f$. Using the same formulation from the intentional model above, we can write

$$\begin{aligned} P(L|I, w, C) &\propto P(I|L, w, C)P(L) \\ &\propto P(I|w = f, w, C)P(w = f) \end{aligned} \quad (11)$$

and if we assume that there is no prior reason to prefer one meaning for w over another ($P(w = f) \propto 1$) and substitute from Equation 8, then we have

$$P(L|I, w, C) \propto \frac{1}{|f|}. \quad (12)$$

In other words, all else being equal, a word is likely to have the meaning that would be informative in this context. So “fedora” would be more likely to refer to the hat (the feature that is most informative in this display) than the glasses. Thus, the pragmatic model described above can be used to capture inferences about the likely meanings of words in context.

In Frank, Goodman, Lai, and Tenenbaum (2009), we presented data that this relation in fact fit adult participants’ judgments about novel adjective meanings with high accuracy. We used schematic displays of shapes that reproduced the same kind of game as shown in Figure 5 and varied the number of shapes with each property (in Figure 5 this would be the number of faces with hats vs. the number with glasses). As the relative extensions changed, participants altered their guesses about novel adjective meanings, suggesting that (at least as a group) they were sensitive to the relative informativeness of different possible word meanings. Although this work is still ongoing, we now have some data from young children that they are able to make such judgments as well (Frank & Goodman, under review).

Conclusions

Eve Clark’s (2003) book pushed me powerfully in the direction of considering communication as the grounding experience of language learning. The theoretical work that I have described here is my attempt to capture computationally some of the insights of Eve’s perspective, particularly the fundamental fact that children are not just passive absorbers of linguistic input but are participants in conversation. The consequence of this fact is that the task of understanding language in the moment – and ascertaining reference in particular – must interact with language learning.

The interaction between interpretation and learning leads to a class of models that I have referred to as communicative or intentional models. In the taxonomy described above, these models are dissociated from associative models not because of the information they include – both model classes can take advantage of social information – but because of the way they break down the learning task. While associative models use social information to bias associative learning, intentional models use social information (as well as pragmatic inference) to inform a guess about what speakers are trying to say. And of course, underlying this class of models is learners’ assumption that speakers are rational agents (Clark, 1988).

I have argued here that intentional models are a powerful framework for using social and pragmatic information in the service of learning the meanings of words. The evidence is strong that by the age of 18 months, children take advantage of this information and make inferences that go beyond the association of words and

objects (Baldwin, 1993; Hollich et al., 2000). But since recent evidence indicates the possibility of word knowledge in even younger children than had previously been suspected (Bergelson & Swingley, 2012), infants' assumptions during this early learning remain an important open question for both empirical and computational investigation.

Independent of whether infants begin life as probabilistic intentional learners, it seems likely that they converge to this position as their vocabulary expands to encompass terms that cannot be learned via contextual associations. For a learner who only need acquire basic level descriptors in a word of repeated exposures, the consequences of intentional learning are relatively modest. But for a learner of a language that contains a wide variety of complex, context-dependent predicates – in other words, a human child – it is essential to understand the contribution of the speaker's communicative intentions to the words they utter.

References

- Akhtar, N., & Montague, L. (1999). Early lexical acquisition: The role of cross-situational learning. *First Language*, 19, 347–358. DOI: 10.1177/014272379901905703
- Baker, C., Saxe, R., & Tenenbaum, J. (2009). Action understanding as inverse planning. *Cognition*, 113, 329–349. DOI: 10.1016/j.cognition.2009.07.005
- Baldwin, D. (1993). Early referential understanding: Infants' ability to recognize referential acts for what they are. *Developmental Psychology*, 29, 832–843. DOI: 10.1037/0012-1649.29.5.832
- Baldwin, D. (1995). Understanding the link between joint attention and language. In C. Moore & P. Dunham (Eds.), *Joint attention: Its origins and role in development* (pp. 131–158). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bergelson, E., & Swingley, D. (2012). At 6–9 months, human infants know the meanings of many common nouns. *Proceedings of the National Academy of Sciences*, 109, 3253–3258. DOI: 10.1073/pnas.1113380109
- Bloom, P. (2002). *How children learn the meanings of words*. Cambridge, MA: The MIT Press.
- Brown, P., Pietra, V., Pietra, S., & Mercer, R. (1993). The mathematics of statistical machine translation: Parameter estimation. *Computational Linguistics*, 19, 263–311.
- Callanan, M. (1985). How parents label objects for young children: The role of input in the acquisition of category hierarchies. *Child Development*, 56, 508–523. DOI: 10.2307/1129738
- Carpenter, M., Nagell, K., & Tomasello, M. (1998). Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the Society for Research in Child Development*, 63(4/Serial no. 255), 1–174. DOI: 10.2307/1166214
- Clark, E.V. (1988). On the logic of contrast. *Journal of Child Language*, 15, 317–335. DOI: 10.1017/S0305000900012393
- Clark, E.V. (1990). On the pragmatics of contrast. *Journal of Child Language*, 17, 417–431. DOI: 10.1017/S0305000900013842
- Clark, E.V. (2003). *First language acquisition*. Cambridge: CUP.
- Clark, H.H. (1996). *Using language*. Cambridge: CUP. DOI: 10.1017/CBO9780511620539

- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13, 148–153. DOI: 10.1016/j.tics.2009.01.005
- Fazly, A., Alishahi, A., & Stevenson, S. (2010). A probabilistic computational model of cross-situational word learning. *Cognitive Science*, 34, 1017–1063. DOI: 10.1111/j.1551-6709.2010.01104.x
- Fisher, C., Hall, D., Rakowitz, S., & Gleitman, L. (1994). When it is better to receive than to give: Syntactic and conceptual constraints on vocabulary growth. *Lingua*, 92, 333–375. DOI: 10.1016/0024-3841(94)90346-8
- Frank, M.C., & Gibson, E. (2011). Overcoming memory limitations in rule learning. *Language Learning and Development*, 7, 130–148. DOI: 10.1080/15475441.2010.512522
- Frank, M.C., Goldwater, S., Griffiths, T.L., & Tenenbaum, J.B. (2010). Modeling human performance in statistical word segmentation. *Cognition*, 117, 107–125. DOI: 10.1016/j.cognition.2010.07.005
- Frank, M.C., & Goodman, N.D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336, 998. DOI: 10.1126/science.1218633
- Frank, M.C., & Goodman, N.D. (Under review). Inferring word meanings by assuming that speakers are informative.
- Frank, M.C., Goodman, N.D., Lai, P., & Tenenbaum, J.B. (2009). Informative communication in word production and word learning. In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, (pp. 1228–1233). Amsterdam: Cognitive Science Society.
- Frank, M.C., Goodman, N.D., & Tenenbaum, J.B. (2007). A Bayesian framework for cross-situational word learning. *Advances in Neural Information Processing Systems*, 20.
- Frank, M.C., Goodman, N.D., & Tenenbaum, J.B. (2009). Using speakers' referential intentions to model early cross-situational word learning. *Psychological Science*, 20, 578–585. DOI: 10.1111/j.1467-9280.2009.02335.x
- Frank, M.C., Tenenbaum, J.B., & Fernald, A. (2013). Social and discourse contributions to the determination of reference in cross-situational word learning. *Language Learning and Development*, 9, 1–24. DOI: 10.1080/15475441.2012.707101
- Gallistel, C. (1990). *The organization of learning*. Cambridge, MA: The MIT Press.
- Gallistel, C., Fairhurst, S., & Balsam, P. (2004). The learning curve: Implications of a quantitative analysis. *Proceedings of the National Academy of Sciences*, 101, 13124–13131. DOI: 10.1073/pnas.0404965101
- Geisler, W. (2003). Ideal observer analysis. In L.M. Chalupa & J.S. Werner (Eds.), *The visual neurosciences* (pp. 825–837). Cambridge, MA: The MIT Press.
- Gergely, G., Bekkering, H., & Király, I. (2002). Rational imitation in preverbal infants. *Nature*, 415, 755. DOI: 10.1038/415755a
- Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition*, 73, 135–176. DOI: 10.1016/S0010-0277(99)00036-0
- Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition*, 1, 3–55. DOI: 10.1207/s15327817la0101_2
- Goodman, N.D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science*, 5, 173–184. DOI: 10.1111/tops.12007
- Gopnik, A. (2012). Scientific thinking in young children. Theoretical advances, empirical research and policy implications. *Science*, 337, 1623–1627. DOI: 10.1126/science.1223416
- Grice, H. (1975). Logic and conversation. *Syntax and Semantics*, 3, 41–58.

- Griffin, Z., & Bock, K. (2000). What the eyes say about speaking. *Psychological Science*, *11*, 274–279. DOI: 10.1111/1467-9280.00255
- Hollich, G., Hirsh-Pasek, K., & Golinkoff, R. (2000). Breaking the language barrier: An emergentist coalition model for the origins of word learning. *Monographs of the Society for Research in Child Development*, *65*, 1–135. DOI: 10.1111/1540-5834.00091
- James, W. (1890). *The principles of psychology, Vol. 1*. New York, NY: Henry Holt and Company. DOI: 10.1037/11059-000
- Johnson, M., Demuth, K., & Frank, M. (2012). Exploiting social information in grounded language learning via grammatical reductions. In *Proceedings of the Association for Computational Linguistics* (pp. 883–891).
- Johnson, M., Griffiths, T., & Goldwater, S. (2007). Adaptor grammars: A framework for specifying compositional nonparametric Bayesian models. *Advances in Neural Information Processing Systems*, *19*, 641–648.
- Li, P., Farkas, I., & MacWhinney, B. (2004). Early lexical development in a self-organizing neural network. *Neural Networks*, *17*, 1345–1362. DOI: 10.1016/j.neunet.2004.07.004
- Locke, J. (1690/1964). *An essay concerning human understanding*. Cleveland, OH: Meridian Books.
- Markman, E.M. (1991). *Categorization and naming in children: Problems of induction*. Cambridge, MA: The MIT Press.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York, NY: Henry Holt & Co.
- McMurray, B., Horst, J.S., & Samuelson, L.K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological Review*, *119*, 831–877. DOI: 10.1037/a0029872
- Medina, T., Snedeker, J., Trueswell, J., & Gleitman, L. (2011). How words can and cannot be learned by observation. *Proceedings of the National Academy of Sciences*, *108*, 9014–9019. DOI: 10.1073/pnas.1105040108
- Mitchell, C., & McMurray, B. (2009). On leveraged learning in lexical acquisition and its relationship to acceleration. *Cognitive Science*, *33*, 1503–1523. DOI: 10.1111/j.1551-6709.2009.01071.x
- Onishi, K., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, *308*, 255–258. DOI: 10.1126/science.1107621
- Perfors, A., Tenenbaum, J., Griffiths, T., & Xu, F. (2011). A tutorial introduction to bayesian models of cognitive development. *Cognition*, *120*, 302–321. DOI: 10.1016/j.cognition.2010.11.015
- Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.
- Quine, W. (1960). *Word and object*. Cambridge, MA: The MIT Press.
- Regier, T. (2005). The emergence of words: Attentional learning in form and meaning. *Cognitive Science*, *29*, 819–865. DOI: 10.1207/s15516709cog0000_31
- Sanborn, A., Griffiths, T., & Navarro, D. (2010). Rational approximations to rational models: alternative algorithms for category learning. *Psychological Review*, *117*, 1144. DOI: 10.1037/a0020511
- Schmidt, L., Goodman, N.D., Barner, D., & Tenenbaum, J. (2009). How tall is tall? Compositionality, statistics, and gradable adjectives. In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, (pp. 3151–3156).

- Schulz, L., Kushnir, T., & Gopnik, A. (2007). Learning from doing: Intervention and causal inference. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 67–85). Oxford: OUP.
- Shafto, P., Goodman, N.D., & Frank, M. (2012). Learning from others the consequences of psychological reasoning for human learning. *Perspectives on Psychological Science*, 7, 341–351. DOI: 10.1177/1745691612448481
- Siskind, J. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition*, 61, 39–91. DOI: 10.1016/S0010-0277(96)00728-7
- Smith, L.B., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, 106, 1558–1568. DOI: 10.1016/j.cognition.2007.06.010
- St. Augustine. (397/1963). *The confessions of St. Augustine*. Oxford: Clarendon Press.
- Stiller, A., Goodman, N.D., & Frank, M. (2011). Ad-hoc scalar implicature in adults and children. In L. Carlson, C. Hoelscher, & T.F. Shipley (Eds.), *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society*, (pp. 2134–2139).
- Tenenbaum, J., & Griffiths, T. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, 24, 629–640.
- Tomasello, M. (2003). *Constructing a language: A usage-based theory of language acquisition*. Cambridge, MA: Harvard University Press.
- Trueswell, J.C., Medina, T.N., Hafri, A., & Gleitman, L.R. (2013). Propose but verify: Fast mapping meets cross-situational word learning. *Cognitive Psychology*, 66, 126–156. DOI: 10.1016/j.cogpsych.2012.10.001
- Vouloumanos, A. (2008). Fine-grained sensitivity to statistical information in adult word learning. *Cognition*, 107, 729–742. DOI: 10.1016/j.cognition.2007.08.007
- Vouloumanos, A., Onishi, K., & Pogue, A. (2012). Twelve-month-old infants recognize that speech can communicate unobservable intentions. *Proceedings of the National Academy of Sciences*, 109, 12933–12937. DOI: 10.1073/pnas.1121057109
- Vouloumanos, A., & Werker, J. (2009). Infants' learning of novel words in a stochastic environment. *Developmental Psychology*, 45, 1611–1617. DOI: 10.1037/a0016134
- Waxman, S., & Booth, A. (2001). Seeing pink elephants: Fourteen-month-olds' interpretations of novel nouns and adjectives. *Cognitive Psychology*, 43, 217–242. DOI: 10.1006/cogp.2001.0764
- Yu, C., & Ballard, D. (2007). A unified model of early word learning: Integrating statistical and social cues. *Neurocomputing*, 70, 2149–2165. DOI: 10.1016/j.neucom.2006.01.034
- Yu, C., & Smith, L. (2012). Modeling cross-situational word-referent learning: Prior questions. *Psychological Review*, 119, 21. DOI: 10.1037/a0026182