Cognition 177 (2018) 177–188

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

Fast mapping word meanings across trials: Young children forget all but their first guess

Athulya Aravind[a,*], Jill de Villiers[b,*], Amy Pace[c], Hannah Valentine[c], Roberta Golinkoff[d], Kathy Hirsch-Pasek[e], Aquiles Iglesias[f], Mary Sweig Wilson

[a] Department of Linguistics and Philosophy, MIT, United States
[b] Department of Psychology, Smith College, United States
[c] Department of Speech and Hearing Sciences, University of Washington, United States
[d] School of Education, University of Delaware, United States
[e] Department of Psychology, Temple University, United States
[f] Department of Communication Sciences and Disorders, University of Delaware, United States

ARTICLE INFO

Keywords:
Word learning
Cross-situational learning
Propose-but-verify
Fast mapping

ABSTRACT

Do children learn a new word by tracking co-occurrences between words and referents across multiple instances ("cross-situational learning" models), or is word-learning a "one-track" process, where learners maintain a single hypothesis about the possible referent, which may be verified or falsified in future occurrences ("propose-but-verify" models)? Using a novel word-learning task, we ask which learning procedure is utilized by preschool-aged children. We report on findings from three studies comparing the word-learning strategies across different populations of child learners: monolingual English learners, Spanish - English dual language learners, and learners at risk for language-delay. In all three studies, we ask what, if anything, is retained from prior exposures and whether the amount of information retained changes as children get older. The ability to make a good initial hypothesis was a function of various factors, including language ability and experience, but across-the-board, children were no better than chance after a wrong initial hypothesis. This suggests that children do not retain multiple meaning hypotheses across learning instances, lending support to the propose-but-verify models.

1. Introduction

A fundamental question in developmental research concerns how knowledge changes over time. Nowhere is this question more pertinent than in the domain of word learning: children go from having effectively no recognizable words in their productive vocabulary before the age of 1 year to having a few hundred just a year or so later. The remarkable ease and efficiency with which children learn the mappings between phonological forms and their meanings might mislead one into thinking that word-learning is a simple task. But as famously pointed out by Quine (1960), even in the simplest cases—involving words denoting concrete nouns and the objects they represent, which most researchers agree can be learned observationally—the task is highly non-trivial. Any naturalistic learning situation has the potential for infinite referential ambiguity in word-to-meaning mapping. Much research in early child language has tried to demonstrate that the situation is less extreme, for instance, children may be predisposed to favor certain mappings over others, thus imposing constraints on the possibilities (Golinkoff, Hirsh-Pasek, Bailey, & Wenger, 1992; Hollich, Golinkoff, & Hirsh-Pasek, 2007; Markman, 1990). Nevertheless, uncertainty is not eliminated altogether, and the question of how children learn the meaning of words given this state of affairs remains a core one in developmental research.

One response to the referential ambiguity problem is to say that the problem is only intractable if the learning situation is limited to one instance in which the word is uttered (Siskind, 1996; Smith & Yu, 2008; Yu & Smith, 2007 and many others). To illustrate the point, consider a child who hears the word 'rabbit' in a situation where a rabbit, a cat and a teddybear are present. Let us also assume that the myriad other possibilities imagined by Quine—the shadows, the lighting, the furriness, the ear length—are taken out of consideration by the constraints the child brings to the situation, such as the "whole-object" constraint (Markman, 1989). From this situation alone, it may be still be unclear what 'rabbit' means. But if the child encounters multiple situations where 'rabbit' is heard, and a rabbit is the only constant among the alternatives, then she may be able to identify the right meaning for the

* Corresponding authors at: Department of Linguistics & Philosophy, 77 Massachusetts Ave, 32D-808, Cambridge, MA 02139, United States (A. Aravind). Department of Psychology, 110 Elm Street, Bass Hall, 401, Northampton, MA 01063, United States (J. de Villiers).
E-mail addresses: aaravind@mit.edu (A. Aravind), jdevilli@smith.edu (J. de Villiers).

https://doi.org/10.1016/j.cognition.2018.04.008
Received 8 October 2017; Received in revised form 8 April 2018; Accepted 9 April 2018

word. This is the core idea behind *cross-situational word learning* (CSM) models: if the learner can track statistical co-occurrence rates across many situations, she is bound to gradually converge on the target mapping.

To successfully use this learning mechanism, the learner must store and compare multiple parallel hypotheses across several learning situations in order to accumulate enough evidence eventually to select the best hypothesis. But can human learners do this? This issue has been investigated experimentally with both adult and child populations. In a typical task, participants have to learn the meaning of several new words in ambiguous situations. Though each individual trial is uninformative, adult participants selected the correct referent at above-chance levels by the end of the study session (McMurray, Horst, & Samuelson, 2012; Vouloumanos, 2008; Yurovsky, Fricker, Yu, & Smith, 2014). The paradigm has also been adapted to test whether children use the same associative learning strategies as adults seemed to be using. Smith and Yu (2008) presented 12- to 14-month-olds with two novel words and two novel referents in an ambiguous training trial. In a later test trial, children saw one of the two referents along with a novel referent. Children preferentially looked at the target in the test trials, despite both the training and test trials being individually ambiguous. The difficulty arises in interpreting this finding as suggesting that the learners were indeed keeping track of information about multiple candidate referents for the same word. There are different possibilities for how experience can accumulate.

An alternative account to the cross-situational learning model posits that word-learning is a "one-track" process, where learners maintain a single hypothesis about the possible referent for any given novel word, which may be verified or falsified in future occurrences. Note that the only evidence that the guess is false is that the hypothesized referent does not appear on the next trial. In the event of falsification, a new candidate referent is hypothesized. In the kind of experiment described, this too would lead to increased success by the final trial, on average. *Propose-but-verify* (PBV) models of word learning posit that nothing is stored from one trial except the learner's latest guess. Supporting evidence comes in part from experiments with adults showing that participants maintain their hypotheses over time if they made an initial correct guess, but pick a new referent at chance among the available candidates if their initial guess was incorrect (Medina, Snedeker, Trueswell, & Gleitman, 2011; Trueswell, Medina, Hafri, & Gleitman, 2013). In other words, participants do not seem to retain any information about alternative targets from past trials. Only one study, by Woodard, Gleitman, and Trueswell (2016), has thus far tested predictions of propose-but-verify with children, and their results also support the model. In the study, 2- and 3-year old children were shown an initial trial involving two novel animals accompanied by a novel word and asked to select the referent for the word. In a second exposure, they were shown a new pair, one of which was either (i) the same animal they selected ("Same" condition) or (ii) the animal they had failed to select previously ("Switch" condition) (Fig. 1). Whereas children could remember their earlier hypothesis in a subsequent trial in the Same condition, they showed no indication of recalling the unselected alternative in the Switch condition. However, as Fig. 1 shows, the other choice was an unknown object that had been seen before, against a familiar alternative (Trial 1). Woodard et al. had designed the task this way to ensure that any above-chance responding observed would not be due to a simple familiarity effect. Nevertheless, in avoiding one confound, another one could have been introduced: the child on the test trial could have recalled both unselected alternatives, but lost track of the situation in which each was presented.

More recently, hybrid accounts have been proposed on which learners may strongly favor a single hypothesis, but also extract some information about alternative candidates (Yurovsky & Frank, 2015; Roembke & McMurray, 2016; Stevens, Gleitman, Trueswell, & Yang, 2017). Yurovsky and Frank's (2015) experimental results with adults were consistent with an integrative model, where learners allocated a fixed amount of attention to a single hypothesis but distributed the rest evenly among the remaining alternatives. The amount of attention allocated to each alternative candidate decreases as the complexity of the scene (i.e. number of potential candidates) increases, so whereas performance after failure was above chance in low-complexity trials, it did not differ from chance in high-complexity trials. There is further experimental evidence that adult participants consider information about previously seen alternatives and even extraneous information from previous exposures.[1] In an eye-tracking study by Roembke and McMurray (2016), participants looked more at previously seen competitors even when they would eventually choose the target. Dautriche and Chemla (2014) found that adults also retained additional, irrelevant information e.g. about the context in which the stimuli appeared and the spatial location of the target referent. The "Pursuit" model presented in Stevens et al. (2017) takes a more nuanced angle on Propose-but-verify. The critical way in which Pursuit differs from Propose-but-verify is in that on Pursuit, learners retain the disconfirmed hypotheses over the course of many learning instances and ultimately choose the best one, i.e., the one that was confirmed most often.

Overall, these hybrid accounts view word learning as a complex, multidimensional process that incorporates aspects of both associative, cross-situational word learning and Propose-but-verify. However, it is an open empirical question whether children's word learning involves both "one-trial" learning as well as associative mechanisms. The only previous study with children is the one we discussed above, and as we pointed out, we cannot be fully certain that the alternatives were not in fact remembered. The rest of the work target adult populations, but there are reasons to doubt the validity of inferences from findings in adult work to children: Adults and children differ substantially in general cognitive abilities such as memory and attention, and children may be less able to remember alternative hypotheses or other information from the initial exposure.

The present study uses a novel word-learning task to contribute to the question of how children learn new words in more complex settings. The task was not designed as an experiment, but rather, it was developed as a test of *fast mapping* ability (Carey & Bartlett, 1978) for a language screening assessment, QUILS (Golinkoff, de Villiers, Hirsh-Pasek, Iglesias & Wilson, 2017). The design is interestingly different from previous studies of novel word learning. The number of candidate-choices is greater, but a verbal prompt, providing semantic and contextual linguistic cues, was included to aid target selection. In both respects, the task more closely simulates a child's real-world experience of learning novel words and offers a more ecologically valid test of their learning mechanisms than the usual experimental set-up. Furthermore, on the test trial, all except the target represented objects the child had not seen previously. In Woodard et al., the alternatives had been presented before to control for familiarity.

We report on results from three large-scale studies that allow us to compare the various learning models across different populations of child learners. In all three studies, we ask what, if anything, is retained from prior exposures and whether the amount of information retained changes as children get older. The answers to these questions will help us arbitrate among cross-situational learning, propose-but-verify, and more hybrid, models of learning. We also ask whether the particular learning strategies employed vary across different types of learners. We compare the behavior of typically developing monolingual learners of English to typically developing dual language learners, as well as atypical learners to ascertain to what extent the core underlying learning mechanisms are shared across these different populations.

Study 1 presents results from typically developing monolingual

[1] Trueswell et al. (2013) criticize these findings on the grounds that the tasks themselves are highly unnatural: for instance, the set of stimuli are often "closed", revolving around a small set of objects participants see repeatedly over the course of the study. Consequently, they argue that although adults *can* in principle extract more than one meaning, they may only do it under very contrived experimental settings.