Journal of Memory and Language 87 (2016) 105–127

Contents lists available at ScienceDirect

Journal of Memory and Language

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

Observational word learning: Beyond propose-but-verify and associative bean counting

Tanja C Roembke [a,*], Bob McMurray [a,b,c]

[a] Dept. of Psychological and Brain Sciences, University of Iowa, United States
[b] Dept. of Communication Sciences and Disorders, University of Iowa, United States
[c] Dept. of Linguistics, University of Iowa, United States

ARTICLE INFO

Article history:
Received 16 May 2014
revision received 23 September 2015
Available online 4 November 2015

Keywords:
Observational learning
Cross-situational learning
Associative learning
Word learning
Statistical learning
Eye movement

ABSTRACT

Learning new words is difficult. In any naming situation, there are multiple possible interpretations of a novel word. Recent approaches suggest that learners may solve this problem by tracking co-occurrence statistics between words and referents across multiple naming situations (e.g. Yu & Smith, 2007), overcoming the ambiguity in any one situation. Yet, there remains debate around the underlying mechanisms. We conducted two experiments in which learners acquired eight word–object mappings using cross-situational statistics while eye-movements were tracked. These addressed four unresolved questions regarding the learning mechanism. First, eye-movements during learning showed evidence that listeners maintain multiple hypotheses for a given word and bring them all to bear in the moment of naming. Second, trial-by-trial analyses of accuracy suggested that listeners accumulate continuous statistics about word–object mappings, over and above prior hypotheses they have about a word. Third, consistent, probabilistic context can impede learning, as false associations between words and highly co-occurring referents are formed. Finally, a number of factors not previously considered in prior analysis impact observational word learning: knowledge of the foils, spatial consistency of the target object, and the number of trials between presentations of the same word. This evidence suggests that observational word learning may derive from a combination of gradual statistical or associative learning mechanisms and more rapid real-time processes such as competition, mutual exclusivity and even inference or hypothesis testing.

Introduction

Observational learning and referential ambiguity

Early in language acquisition, children are often assumed to learn the mapping between words and objects largely from observation (Gleitman, 1990) without reliable feedback. However, a fundamental problem for observational learning is referential ambiguity (Quine, 1960): In any naming event, there is a vast array of possible interpretations for a novel word. Consequently, learners may require strategies or biases to cope with this ambiguity (Golinkoff, Hirsh-Pasek, Bailey, & Wenger, 1992; Markman, 1990). Recently, Yu and Smith (2007; see also Siskind, 1996) argued the problem of referential ambiguity may in part be an artificial consequence of restricting the analysis of word learning to one encounter with a word. Across multiple situations, there may be sufficient statistical information to support learning. For example, many words (e.g. objects) are more likely to co-occur with their referents than with other objects.

* Corresponding author at: E11 SSH, Dept. of Psychological and Brain Sciences, University of Iowa, Iowa City, IA 52242, United States. Fax: +1 319 335 0191.
E-mail address: tanja-roembke@uiowa.edu (T.C Roembke).

http://dx.doi.org/10.1016/j.jml.2015.09.005
0749-596X/Published by Elsevier Inc.

Yu and Smith (2007) tested this in adults (and later in infants, Smith & Yu, 2008): On each trial, participants saw a number of novel objects and heard novel names for each of them, creating considerable ambiguity. Across multiple trials, a word and its referent always co-occurred while its co-occurrence with other objects was lower. After a short training, participants showed above-change accuracy for selecting the words' referents, suggesting statistics were sufficient to support learning. This raises the possibility that learners have powerful mechanisms for inferring the words' meanings across multiple situations, even if any given situation is ambiguous.

How do people learn words in the cross-situational paradigm?

There has since been a large number of experiments examining how mostly adults learn words in observational paradigms (Medina, Snedeker, Trueswell, & Gleitman, 2011; Trueswell, Medina, Hafri, & Gleitman, 2013; Vouloumanos, 2008; Yurovsky, Yu, & Smith, 2013). This has led to a debate over the mechanism underlying such learning.

Originally, Yu and Smith (2007, 2012) described cross-situational learning as a process of tracking co-occurrence statistics between words and objects across many situations. This is a form of statistical or associative learning in which the word–object pairs with the highest co-occurrence are the correct mapping. However, more recent accounts suggest people could harness cross-situational information using propositional logic (Medina et al., 2011; Trueswell et al., 2013): The most prominent theory of this sort is "propose-but-verify", in which learners form a single explicit hypothesis after encountering a novel word, which is carried forward unless disconfirmed by later encounters.

Others have proposed hybrid accounts: For example, there are memory-based accounts in which such inferences are made over stored episodes of situations in long-term memory (Dautriche & Chemla, 2014). Bayesian accounts take a hypothesis-testing approach, but evaluate multiple probabilistic hypotheses simultaneously to find the most likely mapping given the data (Frank, Goodman, & Tenenbaum, 2009). Finally, McMurray, Horst, and Samuelson (2012) propose that gradual associative learning may be buttressed with real-time decision making to account for both cross-situational learning and other developmental phenomena. These real-time processes may allow the system to engage in more inferential processes in the moment (e.g. mutual exclusivity), while long-term statistics are tracked via associations.

These theories are still developing with newer iterations of purely statistical accounts (Yu & Smith, 2012), propose-but-verify (Koehne, Trueswell, & Gleitman, 2014) and the dynamic associative account (McMurray, Zhao, Kucker, & Samuelson, 2013). While these theories may exhibit stark differences in their core commitments (e.g. whether learning is propositional or associative), they appear flexible in how these commitments get implemented. Consequently, it may be premature to experimentally disentangle them.

However, there are crucial open questions about the basic properties of observational learning, which may constrain how these theories are developed. Thus, we identified four such questions that have played (or may play) a crucial role in these debates and critically evaluated them across two experiments. These questions include the issues of (1) whether participants maintain multiple hypotheses for a given word[1]; (2) whether information is gradually accumulated; (3) the role of context, and (4) other factors that may shape learning.

Do learners maintain multiple hypotheses about the meaning of a word?

The first question is how many hypotheses learners maintain for a given word. For example, in a dinner table event, when *fork* is heard for the first time, do learners form a single hypothesis for *fork* (positing that it refers to *either* the fork or the spoon), or do they note that this word co-occurred with both objects (but not with a car or boat)? In an associative account, learners track the co-occurrence of multiple objects with a word (e.g. Yu & Smith, 2007), relying on the accumulation of data to resolve any ambiguity. At the dinner table, for example, the learner will eventually encounter the word *fork* without a spoon, pushing its statistical co-occurrence with fork above that with spoon. Consequently, learners must maintain multiple hypotheses with different degrees of strength. In contrast, early versions of propose-but-verify suggested learners posit a *single* hypothesis about a word, which can be updated on future encounters. However, more recent propositional accounts also admit multiple hypotheses: For example, learners may recall previously considered hypotheses in the face of memory failure or disconfirming evidence (e.g. Koehne et al., 2014).

As an empirical issue, whether learners track one or many hypotheses remains unresolved. This is largely because most studies address this issue indirectly using trial-by-trial autocorrelation analyses. Such analyses infer what a learner may have learned about a word from previous trials' accuracy, and measure how it predicts performance on subsequent encounters (Trueswell et al., 2013): In propositional accounts, if learners previously selected the correct object, they must have arrived at the right hypothesis and should continue to select the correct object on present trials. However, if they were incorrect on a previous trial, they likely had the wrong hypothesis, and should now be at chance. In contrast, in statistical accounts, even on an incorrect trial, they accumulate more "data" and could show a benefit on subsequent trials. Autocorrelation analyses conducted by Trueswell et al. (2013) supported a single-hypothesis account, and even an analysis of participants' eye-movements (a potentially more sensitive measure) showed little evidence for any learning after an incorrect trial.

Dautriche and Chemla (2014) pointed out that prior incorrect trials may function differently depending on the information on the current trial: If the prior incorrect

[1] We here use the term *hypothesis* to refer to any knowledge structure mapping a word to potential referents, including both abstract knowledge and associative links.