Cross-Situational Learning

Article · January 2012

DOI: 10.1007/978-1-4419-1428-6_1712

2 authors:

Andrew D M Smith
University of Stirling

Kenny Smith
University of Edinburgh

Philipona, D., O'Regan, K., & Nadal, J. P. (2003). Is there something out there? Inferring space from sensorimotor dependencies. *Neural Computation, 15*(9), 2029–2049.
Stein, B. E., & Meredith, M. A. (1993). *The merging of the senses.* Cambridge, MA: MIT Press.
Sloutsky, V. M. (2003). The role of similarity in the development of categorization. *Trends in Cognitive Sciences, 7*, 246–251.
Waxman, S. R., & Gelman, S. A. (2009). Early word-learning entails reference, not merely associations. *Trends in Cognitive Sciences, 13*(6), 258–263.

Cross-Sectional Research

► Longitudinal Learning Research on Changes in Learning of University Students

Cross-Situational Learning

Andrew D. M. Smith[1], Kenny Smith[2]
[1]Literature and Languages, School of Arts and Humanities, University of Stirling, Stirling, UK
[2]Linguistics and English Language, School of Philosophy, Psychology and Language Sciences, University of Edinburgh, Edinburgh, Scotland, UK

Synonyms

Associative learning

Definition

Cross-situational learning is a technique for learning the meanings of words across multiple exposures, despite uncertainty as to the word's meaning on each individual exposure. The cross-situational learner encounters a word in a number of different situations, each of which provides a set of multiple candidate meanings; the learner determines the word's meaning by selecting from those meanings which reliably recur across situations. Cross-situational learning is less cognitively demanding than many other models of word learning, because it does not require a learner to unambiguously identify a word's meaning on a single exposure.

Theoretical Background

During language acquisition, children learn a lexicon containing many thousands of associations between words and their meanings, at the rate of around ten new words a day. Children accomplish this task rapidly and remarkably successfully, overcoming potentially unlimited uncertainty about the meaning of every new word they encounter, and identifying some aspects of word meaning after only a very few exposures through so-called fast mapping. Quine (1960) famously illustrated the problem of referential uncertainty through the story of an imaginary anthropologist working with a speaker of an unfamiliar language: when a rabbit runs past, the speaker shouts "gavagai," and the anthropologist tentatively notes that this new word means "rabbit." Quine's insight was to point out, however, that the anthropologist can never be sure that "gavagai" means "rabbit," no matter how many future clarificatory tests are carried out; it could, after all, have an infinite number of possible meanings of varying plausibility, including "animal," "white," "undetached rabbit parts," "dinner," or "it will rain."

Yet despite the philosophical problem of unlimited referential uncertainty, children clearly *do* learn large lexicons, and the focus of much research into word learning has been on providing explanations for this. The dominant approach has been to identify mechanisms which allow the learner to exclude from consideration many meanings which are theoretically possible but in reality spurious, thus reducing the level of referential uncertainty in the input to a more manageable level, and simplifying the task of determining the word's true meaning. A number of heuristics have been put forward: interpreting behavioral cues in order to identify the speaker's focus of attention; assuming that novel words are more likely to refer to whole objects rather than their parts or properties; building on existing knowledge about the meanings of other words and assuming that new words will have different meanings; making use of the syntactic context in which the new word is presented to infer aspects of its meaning (see Bloom 2000, for review). While quantifying the impact of such heuristics is problematic, it is clear that some referential uncertainty is likely to remain even after the application of many or all of them; the utility of cross-situational learning stems from the fact that it allows words to be learnt despite the existence of residual uncertainty.

Cross-situational learning works by amalgamating information about the meaning of a word from across the various different situations in which that word occurs. Each separate context in which the word is used yields a (possibly infinite) set of possible candidate meanings, which is potentially reduced through the application of word-learning heuristics such as those described above to a finite set of candidate meanings (including the true meaning); the same word uttered in a different context may of course yield a different set of candidate meanings. Candidate meaning sets from different contexts can be combined, enabling the learner to identify the most likely correct meaning, for instance, by identifying the meaning which lies at the intersection of the sets, as shown in Fig. 1. Although each exposure to a new word may provide a large number of possible meanings, and thus a large degree of referential uncertainty, successive exposures in different contexts will gradually reduce the uncertainty, eventually eliminating it completely by winnowing the set of possible meanings down to the true meaning alone.

This eliminative approach to cross-situational learning illustrated in Fig. 1, however, is vulnerable to failure in a number of real-world circumstances (see Gleitman et al. 2005, for discussion): In noisy situations where the intended meaning is not suggested by the environment, it will be sifted out of the set of possible meanings; in homonymous or polysemous situations where the word has more than one intended meaning (e.g., the English word "bank"), none of the intended meanings will appear in all exposures, and thus the set of possible meanings will be empty (i.e., situations in which "bank" is used as a verb denoting turning will probably not feature financial institutions in their set of likely meanings; likewise, situations in which it is used as a noun will not reliably feature the act of turning).

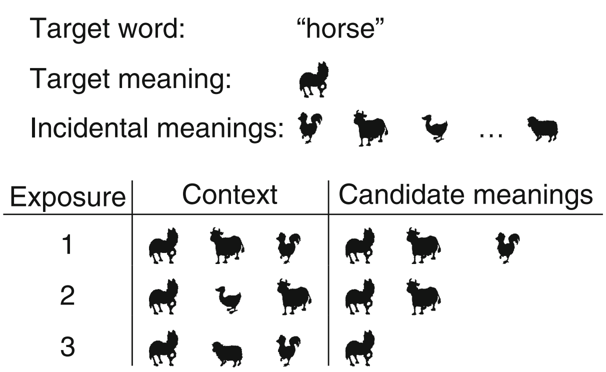

Cross-Situational Learning. Fig. 1 Cross-situational learning. Each time the word *horse* is used, the context provides a different set of candidate meanings. Uncertainty about the meaning of the word is gradually reduced and finally eliminated through its appearance in multiple exposures, as candidate meanings which are not suggested by each context are eliminated from consideration

These vulnerabilities stem from the pure cross-situational learner maintaining the *maximal* amount of cross-situational information, namely, an accurate set of candidate meanings which always occur with the word. At the other end of the spectrum, a learner could make *minimal* use of cross-situational information by simply remembering a single one of the meanings suggested in a previous exposure, and maintaining this as their preferred meaning so long as it is also suggested by the current context. Between these two extremes lie an infinite number of potential cross-situational learning strategies, much more resilient to noise, yet less powerful than pure cross-situational learning (Blythe et al. 2010). In particular, a *frequentist* strategy, where learners track the frequency with which candidate meanings co-occur with the target word, appears to match well the data from experimental tests of cross-situational learning (Yu and Smith 2007; Smith et al. 2011).

Important Scientific Research and Open Questions

Existing research into cross-situational learning can be grouped into two main approaches: formal computational and mathematical models examining the operationalization of cross-situational learning and its plausibility as a tool for language learning, and experimental work exploring the conditions under which humans use the different cross-situational learning strategies.

Siskind (1996) developed an early and influential computational implementation of cross-situational learning based on the eliminative process illustrated in Fig. 1, describing an algorithm which was capable of identifying word meanings after exposure to a synthesized corpus of utterances paired with both intended and spurious meanings. Siskind further demonstrated that his cross-situational learning procedures could be specified so that the algorithm could recover from errors originating from environmental noise and homonymy. More recent formal models (e.g., Yu et al. 2005) have successfully demonstrated that cross-situational learning can be used to infer the meanings of words from increasingly complex and realistic, though still small, corpora of natural language use. Mathematical models, meanwhile, have shown that cross-situational learning is viable not just with small corpora, but also scales up to the learning of large, human-size vocabularies within reasonable timescales (Blythe et al. 2010). Despite significant levels of referential uncertainty at each exposure, the relative learning speed disadvantage of cross-situational learning compared to an idealized fast-mapping learner is surprisingly small. There is, therefore, no necessary link between the ability to learn individual words rapidly and the ability to acquire large vocabularies.

A body of research has demonstrated that both adults and infants can effectively exploit cross-situational learning information when learning small numbers of words (e.g., Akhtar and Montague 1999; Yu and Smith 2007), using both naturalistic and more controlled (and therefore quantifiable) stimuli. The effectiveness of cross-situational learning in humans is affected by the degree of referential uncertainty: performance (in terms of number of words learnt) decreases as referential uncertainty increases. Furthermore, increasing referential uncertainty appears to change the mechanism by which cross-situational learning takes place, with increased referential uncertainty prompting a shift from a pure eliminative strategy to a less demanding, more nuanced, frequentist equivalent (Smith et al. 2011).

This recent emphasis on examining when alternative cross-situational strategies are employed by learners leads to a number of currently open questions. Increasing referential uncertainty naturally increases the time it takes to learn a lexicon, yet weaker forms of cross-situational learning (those which make less efficient use of cross-situational statistics) are disproportionately affected by increases in referential uncertainty than stronger forms.

At some point, therefore, increasing referential uncertainty will make a human-size lexicon impossible to learn by cross-situational learning in a reasonable timescale. Quantifying this critical point, however, is still problematic, not only because of the difficulties in accurately quantifying the referential uncertainty of naturalistic data, but also because the experimental evidence for when and how people shift learning strategies is currently rather minimal. Furthermore, existing research into different variants of cross-situational learning has primarily been carried out on adults, posing the question of whether children shift strategies in response to task demands in the same way as adults, or whether they even use the same cross-situational learning strategies at all.

Cross-References

► Associative Learning
► Embodied Cognition
► Heuristics and Problem Solving
► Matching
► Meaningful Verbal Learning

References

Akhtar, N., & Montague, L. (1999). Early lexical acquisition: The role of cross-situational learning. *First Language, 19*(57), 347–358.

Bloom, P. (2000). *How children learn the meanings of words*. Cambridge, MA: MIT Press.

Blythe, R. A., Smith, K., & Smith, A. D. M. (2010). Learning times for large lexicons through cross-situational learning. *Cognitive Science, 34*(4), 620–642.

Gleitman, L. R., Cassidy, K., Nappa, R., Papafragou, A., & Trueswell, J. C. (2005). Hard words. *Language Learning and Development, 1*(1), 23–64.

Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.

Siskind, J. M. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition, 61*, 39–91.

Smith, K., Smith, A. D. M., & Blythe, R. A. (2011). Cross-situational learning: An experimental study of word-learning mechanisms. *Cognitive Science, 35*(3), 480–498.

Yu, C., & Smith, L. B. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science, 18*(5), 414–420.

Yu, C., Ballard, D. H., & Aslin, R. N. (2005). The role of embodied intention in early lexical acquisition. *Cognitive Science, 29*(6), 961–1005.

Crosswise Research

► Longitudinal Learning Research on Changes in Learning of University Students