Special Issue: Against stored abstractions: A radical exemplar model of language acquisition

FIRST LANGUAGE

First Language
2020, Vol. 40(5-6) 581–584

DOI: 10.1177/0142723720915402
journals.sagepub.com/home/fla
SAGE

The many blessings of abstraction: A commentary on Ambridge (2020)

Joshua K. Hartshorne
Boston College, USA

Abstract

Ambridge argues that the existence of exemplar models for individual phenomena (words, inflection rules, etc.) suggests the feasibility of a unified, exemplars-everywhere model that eschews abstraction. The argument would be strengthened by a description of such a model. However, none is provided. I show that any attempt to do so would immediately run into significant difficulties – difficulties that illustrate the utility of abstractions. I conclude with a brief review of modern symbolic approaches that address the concerns Ambridge raises about abstractions.

Keywords

Bayesian modeling, exemplars, language, language acquisition, symbolic systems

Ambridge (2020) asks us to imagine language without abstractions. This is motivated in part by empirical phenomena, but also by theoretical work. Ambridge reviews a plethora of models that replace common linguistic abstractions like the words *dog* and *cat* or the phoneme /t/ with real-time generalization over exemplars. Each of these typically only eliminates a single abstraction. For instance, Regier (2005) has no words, Nosofsky (1990) has no inflectional rules, and Walsh et al. (2010) has no syntactic constructions. Unfortunately, each of these models uses many other abstractions: Regier (2005) uses numerous abstract phonetic features, Nosofsky (1990) employs both phonetics and abstract word classes, Walsh et al. (2010) uses words, and so on. Ambridge suggests that if every abstraction can be eliminated individually, perhaps they can all be eliminated simultaneously.

This is easier said than done. Indeed, Ambridge provides little guidance as to how this might work. One option might be to treat each utterance–meaning pair as its own

Corresponding author:
Joshua K. Hartshorne, Department of Psychology and Neuroscience, Boston College, 140 Commonwealth Ave., McGuinn 527, Chestnut Hill, MA 02467, USA.
Email: joshua.hartshorne@bc.edu

exemplar, with no internal structure, as Ambridge appears to suggest in footnote 19. The linguistics literature provides numerous reasons to be skeptical. For instance, it is trivial to construct utterances that are arbitrarily dissimilar from previously-encountered utterances (*Arctic parrots cook seven-legged trucks*), are of arbitrary length (*Bill said that Sally said that Horace said that. . .*), or involve mostly novel words (*'Twas brillig and the slythy toves. . .*). An American, having only just been informed that the British say 'elk' to mean MOOSE and 'torch' to mean *flashlight* can readily use her new cultural competency to correctly interpret a British utterance of *John saw the elk by the light of the torch*. Similarly, people can quickly adapt to novel accents based on systematic phonetic transformations (Maye et al., 2008) – an ability that MacWhinney (2020) discusses in his commentary as well. It is unclear how any of this could be explained by analogy over previously-encountered utterance–meaning pairs, particularly when one realizes that Utterance is itself exactly the sort of abstraction we are trying to eliminate. Indeed, these are exactly the kinds of reasons linguists hypothesize structured represention such as phonemes, words, morphemes, and syntactic constructions.

The other option is to chain exemplar models together: utterances are real-time generalizations over syntactic structure exemplars, syntactic structures are real-time generalizations over exemplars made of words, words are real-time generalizations of exemplars made of phonemes, phonemes are real-time generalizations over phonetic features, etc. Leaving aside whether this would work in principle – many of the problems raised above still apply – it faces serious practical difficulties. Even very conservative estimates put the number of real-time analogies needed to compute a single new utterance at over 1 trillion, and we have not yet considered abstractions such as parts of speech or inflection. This problem may be solvable, but the solution cannot be assumed.

The best-known solution to this sort of combinatorial explosion is to approximate these calculations using stored abstractions. A well-designed abstraction is a simplification that, at some minimal cost to precision, provides computational tractability. For instance, one could replace the abstract concept *money* with a bag of exemplars (coins, bills, rubles, yen, credit, bitcoin, etc.) with some arguable gain to precision, but at the cost of making the entire discipline of economics unwieldy (Fodor, 1974). Abstractions are *useful*.

This usefulness is illustrated by the already-mentioned fact that all the reviewed models make extensive use of abstractions in defining their exemplars. Indeed, Ambridge employs abstractions even in 'toy examples.' In the following example, he makes extensive use of discourse-level and conceptual abstractions, despite elsewhere arguing that linguistic meaning does not make use of conceptual abstractions:

> The child generates *She's tamming* to express the message 'Discourse-old female undergoes spinning action' by analogy across utterances [such as] 'Discourse-old female undergoes bouncing action' = *She's bouncing*; 'Discourse-old female undergoes spinning actions' = *She's spinning* [and] 'Discourse-new female undergoes spinning action' = *Sue's spinning* [and] the utterance in which the novel verb was trained ('Toy block undergoes bouncing+spinning action'=*Look, tamming!. . .*)

This brings us to our central point. It *might* be possible to design a mind free of abstractions, but would we want to? Abstractions – particularly symbols and rules – are phenomenally powerful computational devices. That is why mathematics and computer science depend on them. Even developers of non-symbolic deep neural networks define these networks using highly abstract languages, interpreted by symbol-processing machines. Abstractions also allow for rapid learning and generalization (Goodman et al., 2011; Tenenbaum et al., 2011). Moreover, abstract symbols allow for straightforward conceptual combination. In contrast, consider that *pet fish* are atypical exemplars of pets and atypical exemplars of fish, making it unclear how any exemplar theory could make sense of that phrase (cf. Fodor & Lepore, 1996). The situation is even more pronounced for *pet vampires* or *pet Martians*.

Ambridge raises concerns about symbolic approaches: they are brittle, do not capture graded human behavior, cannot instantiate 'islands of productivity,' require intricate innate specification, etc. However, these concerns do not apply as readily to modern work as to the 20th-century approaches Ambridge focuses on.[1] Bayesian approaches have endowed symbolic systems with the flexibility and graded behavior that were once the domain of purely non-symbolic approaches (Tenenbaum et al., 2011). Models of morphology have pushed beyond 'single default rule' paradigms to employ myriad productive rules (O'Donnell, 2015). Bayesian priors and program induction allow symbolic models to flexibly modify their own structure, learning the model from the input rather than requiring precise innate specification (Lake et al., 2015; Overlan et al., 2017; Perfors et al., 2010; Tenenbaum et al., 2011). For instance, it is now possible to induce phrase-structure grammars directly from corpora (Pate & Johnson, 2016). Other work has augmented symbolic approaches with aspects of non-symbolic computation, opening up exciting new research directions and suggesting new hypotheses about brain function (Battaglia et al., 2018; Bingham et al., 2019; Yildirim et al., 2019).

We will not know what approach to modeling language acquisition will work until we have one that does. However, as work on exemplar theories shows, it is literally difficult to imagine cognition without abstractions. Given the utility of abstractions, one may wonder whether we should even try.

Acknowledgements

The author thanks Nathan Hartshorne for a timely discussion of deer classification.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by National Science Foundation [1918813].

ORCID iD

Joshua K. Hartshorne https://orcid.org/0000-0003-1240-3598

Note

1. While I focus on symbolic approaches, it is also true that modern prototype accounts are far more powerful than the ones Ambridge critiques (Battleday et al., 2017).

References

Ambridge, B. (2020). Against stored abstractions: A radical exemplar model of language acquisition. *First Language 40*(5-6): 509–559.

Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., . . . Pascanu, R. (2018). Relational inductive biases, deep learning, and graph networks. arXiv preprint. arXiv:1806.01261.

Battleday, R. M., Peterson, J. C., & Griffiths, T. L. (2017). Modeling human categorization of natural images using deep feature representations. arXiv preprint. arXiv:1711.04855.

Bingham, E., Chen, J. P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., . . . Goodman, N. D. (2019). Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research*, *20*(1), 973–978.

Fodor, J. A. (1974). Special sciences (or: The disunity of science as a working hypothesis). *Synthese*, *28*, 97–115.

Fodor, J. A., & Lepore, E. (1996). The red herring and the pet fish: Why concepts still can't be prototypes. *Cognition*, *58*(2), 253–270.

Goodman, N. D., Ullman, T. D., & Tenenbaum, J. B. (2011). Learning a theory of causality. *Psychological Review*, *118*(1), 110–119.

Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, *350*(6266), 1332–1338.

MacWhinney, B. (2020). The role of competition and timeframes: A commentary on Ambridge (2020). *First Language 40*(5-6): 604–607.

Maye, J., Aslin, R. N., & Tanenhaus, M. K. (2008). The weckud wetch of the wast: Lexical adaptation to a novel accent. *Cognitive Science*, *32*(3), 543–562.

Nosofsky, R. M. (1990). Relations between exemplar-similarity and likelihood models of classification. *Journal of Mathematical Psychology*, *34*(4), 393–418.

O'Donnell, T. J. (2015). *Productivity and reuse in language: A theory of linguistic computation and storage*. The MIT Press.

Overlan, M. C., Jacobs, R. A., & Piantadosi, S. T. (2017). Learning abstract visual concepts via probabilistic program induction in a language of thought. *Cognition*, *168*, 320–334.

Pate, J. K., & Johnson, M. (2016). Grammar induction from (lots of) words alone. In Y. Matsumoto & R. Prasad (Eds.), *Proceedings of COLING 2016, the 26th international conference on computational linguistics: Technical papers* (pp. 23–32). COLING 2016 Organizing Committee.

Perfors, A., Tenenbaum, J. B., & Wonnacott, E. (2010). Variability, negative evidence, and the acquisition of verb argument constructions. *Journal of Child Language*, *37*(3), 607–642.

Regier, T. (2005). The emergence of words: Attentional learning in form and meaning. *Cognitive Science*, *29*(6), 819–865.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, *331*(6022), 1279–1285.

Walsh, M., Möbius, B., Wade, T., & Schütze, H. (2010). Multilevel exemplar theory. *Cognitive Science*, *34*(4), 537–582.

Yildirim, I., Wu, J., Kanwisher, N., & Tenenbaum, J. (2019). An integrative computational architecture for object-driven cortex. *Current Opinion in Neurobiology*, *55*, 73–81.