Limits of the "blowfish effect": Exemplar variability outweighs atypicality to support basic-level generalization during word learning

Kennedy Casey

In a Tortoiseshell: *In the concluding section of her final project for Cognitive Psychology, Kennedy Casey adeptly discusses her research on generalization during word learning. She clearly summarizes her findings and their limitations, while also defining her contribution to the* ***scholarly conversation*** *and calling attention to her* ***global motive****.*

Excerpt

Our findings demonstrate that exposure to variable atypical exemplars supports generalization during word learning. More specifically, upon encountering a new word in the context of multiple atypical exemplars from the same subordinate-level category (e.g., three blowfish), participants narrowly applied the label to only subordinate-level matches (i.e., *reng* = blowfish), thereby replicating the blowfish effect (Emberson et al., 2019) as well as the suspicious coincidence effect (e.g., Xu & Tenenbaum, 2007), outlined in past research. Conversely, upon learning a new word in the context of exposure to multiple atypical exemplars belonging to *different* subordinate-level categories (e.g., one blowfish and two other atypical fish), participants adopted a broader interpretation of the word's meaning and generalized the label to most, if not all, basic-level category members (i.e., *reng* = fish).

Thus, consistent with a broader literature suggesting that variability in training input facilitates generalization across myriad domains (e.g., Bambach et al., 2018; Higuchi et al., 2020; Perry et al., 2010; Vukatana et al., 2015), exemplar variability in the present study similarly supported abstraction of the newly learned label, broadly the basic-level category rather than narrowly to just the subordinate-level category. Critically, this generalization occurred even for exemplars that were less characteristic of the basic-level category (e.g., blowfish as an atypical fish), providing preliminary evidence to suggest that exemplar variability outweighs effects of atypicality to facilitate broader extension of a newly learned label to even typical category members. This finding is intriguing given that past research posits a bias in the directionality of generalization, where information learned about typical category members is extended to atypical members, but the reverse is less likely (e.g., Dunsmoor & Murphy, 2014). This begs the question of how far effects of variability might extend to diminish or entirely abolish typicality effects under certain learning conditions.

© Spring 2021 Tortoise

A notable caveat to the present study is that the visual stimuli were not normed for typicality due to the time constraints of the assignment. We relied on intuition to determine whether items were typical or atypical. However, given the graded nature of category membership (e.g., Rosch et al., 1976) as well as individual differences in conceptions of typicality (e.g., Hampton & Passanisi, 2016), this approach was necessarily imperfect. Thus, future experiments must first include a stimulus norming study as in similar previous investigations (e.g., Emberson et al., 2019; Emberson & Rubinstein, 2016).

Another limitation of the current work is the use of only atypical exemplars during exposure. Again, given the constraints of the assignment and the sample size of the class, we opted for a within-subjects design and dropped the typical condition to ensure sufficient statistical power to investigate the effect of interest (i.e., whether exemplar variability counteracts atypicality to support generalization). Given that basic-level abstraction has been shown for typical items, even without exemplar variability (e.g., Emberson et al., 2019), we predicted that this would have also been the case if the typical exemplar condition were included, though future replication studies are warranted to confirm this effect and obtain a full picture of the interaction between exemplar typicality and variability as it relates to word learning.

Although novel, the present findings are not all that surprising given input statistics. Intuitively, it seems unlikely that a competent speaker would simultaneously label a set of different items in an ostensive manner and intend to refer only to their subordinate-level categories. Future studies might compare children's and adults' performance on a similar word learning task to probe the question of how knowledge about the likelihood of encountering certain types of naming events influences meaning generalization. Although previous studies have found similar levels of generalization across child and adult populations (e.g., Emberson et al., 2019; Xu & Tenenbaum, 2007), it is nevertheless possible that performance could diverge when considering the interaction between exemplar typicality and variability.

Another dimension worth considering in future research is that of the temporal dynamics of naming events. While some studies suggest that results may not vary according to whether exemplars are presented simultaneously or interleaved across various trials (e.g., Lewis & Frank, 2018), others point to a key role of the temporal structure of naming events, with narrow subordinate-level generalization tending to occur only after simultaneous presentation of exemplars (e.g., Spencer et al., 2011). This provides yet another opportunity to examine potential developmental differences in the ability to track and exploit input statistics surrounding typical naming events to facilitate appropriate levels of generalization.

© Spring 2021 Tortoise

In sum, the present findings introduce an important new combination of factors to consider when reasoning about generalization during word learning – the interaction between exemplar variability and typicality. While these dimensions have been studied independently, this represents the first experiment to include both factors simultaneously. Only when examining the combination of factors do we see that exemplar variability seems to outweigh effects of exemplar atypicality to promote broad generalization during word learning. Thus, the current work introduces important potential implications for our understanding of how the generalization process may unfold in naturalistic word learning contexts.

© Spring 2021 Tortoise

Author Commentary

Kennedy Casey

The final project assignment for Cognitive Psychology gave students the opportunity to design and run an experiment (as a group) and write up the results in the style of a traditional journal article. The first step was to come up with a novel, well-motivated research question — and specifically, one that was reasonable for us to at least begin to answer with data from a short, online study conducted with participants from our class. Therefore, we needed to read up on past research in order to determine the focus of our project. Motive for scholarly work often comes from identifying a gap in the existing literature surrounding a specific phenomenon (in this case, word learning/generalization), but just because something is understudied does not necessarily provide motivation enough to research it. It's important to be able to articulate why your specific question is interesting and important. How does it relate to *and* extend previous research? How would getting an answer to your question advance our understanding of some phenomenon? These are considerations that should come up before conducting research and while writing.

Here, we were broadly interested in how people generalize word meanings. When they hear a word for the first time (e.g., *reng* being used to refer to a goldfish), how narrow or broad is their interpretation of its meaning, and how is this influenced by learning conditions? Past research on this topic indicated that, most of the time, people assume a 'basic-level' definition of a new word, which is not too general but also not too specific (e.g., *reng* = fish, rather than sea creature or goldfish). However, a recent study conducted at Princeton revealed that this is not always the case. Specifically, when people were taught that a new word corresponds to an atypical member of a familiar category (e.g., blowfish as an atypical fish), they instead formed a narrower interpretation of the new word and assumed that it referred only to blowfish and not to all types of fish. The researchers termed this phenomenon the 'blowfish effect', and we were interested in seeing how far this effect of atypicality could extend. We brought together this line of work with other research describing how exposure variability supports generalization, and we asked about the interaction between these two factors.

Description of past research came not only in the introduction of this paper to set up and motivate our study but also in the discussion section. I find that this is often the hardest section to write, but it's also the one that can be most exciting because it forces me to think big picture (or at least bigger picture) about my research and why it matters. Here, I started by summarizing the major findings from this experiment, while also relating them to past research

© Spring 2021 Tortoise

(i.e., explicitly stating when we had replicated previous studies). I then addressed the various limitations of our experiment. No study is ever perfect, so it's important to recognize any apparent limitations. This might seem like a negative way to end the paper, but this can actually give you the space to talk about open questions and ultimately help to motivate future research by you or others. Finally, ending with an explicit statement about your contribution to the scholarly conversation is critical. In this case, I re-emphasized our novel combination of exemplar typicality and variability in a single experiment and described how looking at this interaction is crucial for our understanding of how word learning/generalization happens not only in the lab, or in the context of this online study, but also potentially in the world, where learning usually takes place.

© Spring 2021 Tortoise

Editor Commentary

Annabelle Duval

A conclusion is a basic necessity for any paper, but writing one that does not merely summarize the ideas presented can be challenging. What should a conclusion aim to accomplish? After focused **analysis** of **evidence** in the body paragraphs, writers can use the conclusion to synthesize the key steps of their arguments and explicitly show how the paper builds to the **thesis** in its most nuanced form. A conclusion is also a return to **motive** — it looks at the bigger picture and reminds the reader why the **thesis** matters, both within the **scholarly conversation** and in a larger, global context. Finally, the conclusion may be a place to ask more questions: what has the research/**evidence/analysis** shown us and where do we go from here?

In the concluding section of her psychology paper, Kennedy skillfully achieves each of these goals. As she outlines in her commentary, Kennedy first provides a brief summary of her research findings and reminds the reader of the data she has carefully discussed in the results section. Kennedy then places these findings in context with previous studies and shows exactly why her research is interesting and different in comparison to the work that has come before it. Importantly, she acknowledges caveats, complications, and limitations to her research. As she writes in her commentary, Kennedy discusses the limitations of her study in a positive way — she turns these limitations into potential directions for future research. Acknowledging limitations also allows the reader to understand fully the complexity of the findings and see exactly how they build upon previous work. Finally, Kennedy leaves the reader with a paragraph on the importance of her research within the **scholarly conversation** and extends this **scholarly motive** to a **global motive**. Specifically, she connects studying exemplar variability and typicality together in a lab setting with the generalization process "in the real world."

While Kennedy's conclusion is an exemplar for papers in scientific disciplines, it can also help those writing in the humanities. Kennedy's directions for future research focus on the effects of different variables, but humanities writers can ask parallel questions in their conclusions. Instead of variables, what themes, language choices, or motifs appear that may be intriguing to track further or look into more? Rather than data, what does textual **evidence** tell us about continuities or disruptions within a piece of literature or a primary source document? No matter the discipline, every conclusion should strive to highlight a nuanced **thesis**, return to **motive**, and leave the reader curious about future research as Kennedy's does.

© Spring 2021 Tortoise

Works Cited

Bambach, S., Crandall, D. J., Smith, L. B., & Yu, C. (2018). Toddler-inspired visual object learning. *Advances in Neural Information Processing Systems, 31*, 1-10.

Dunsmoor, J. E., & Murphy, G. L. (2014). Stimulus typicality determines how broadly fear is generalized. *Psychological Science, 25*(9), 1816-1821.

Emberson, L. L., Loncar, N., Mazzei, C., Treves, I., & Goldberg, A. E. (2019). The blowfish effect: Children and adults use atypical exemplars to infer more narrow categories during word learning. *Journal of Child Language, 46*(5), 938-954.

Emberson, L. L., & Rubinstein, D. Y. (2016). Statistical learning is constrained to less abstract patterns in complex sensory input (but not the least). *Cognition, 153*, 63-78.

Hampton, J. A., & Passanisi, A. (2016): When intentions do not map onto extensions: Individual differences in conceptualization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 42*(4), 505-523.

Higuchi, Y., Ueda, Y., Shibata, K., & Saiki, J. (2020). Spatial variability induces generalization in contextual cueing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 46*(12), 2295-2313.

Lewis, M. L., & Frank, M. C. (2018). Still suspicious: The suspicious-coincidence effect. *Psychological Science, 29*(2), 2039-2047.

Perry, L. K., Samuelson, L. K., Malloy, L. M., & Schiffer, R. N. (2010). Learn locally, think globally: Exemplar variability supports higher-order generalization and word learning. *Psychological Science, 21*(12), 1894-1902.

© Spring 2021 Tortoise

Rosch, E., Simpson, C., & Scott, M. R. (1976). Structural bases of typicality effects. *Journal of Experimental Psychology: Human Perception and Performance, 2*(4), 491-502.

Spencer, J. P., Perone, S., Smith, L. B., & Samuelson, L. K. (2011). Learning words in space and time: Probing the mechanisms behind the suspicious-coincidence effect. *Psychological Science, 22*(8), 1049-1057.

Vukatana, E., Graham, S. A., Curtin, S., & Zepeda, M. S. (2015). One is not enough: Multiple exemplars facilitate infants' generalization of novel properties. *Infancy*, *20*(5), 548-575.

Xu, F., & Tenenbaum, J. (2007). Word learning as Bayesian inference. *Psychological Review, 114*(2)*,* 245-272.

© Spring 2021 Tortoise

Bios

Kennedy Casey ’21 is a senior in the Psychology Department, also pursuing certificates in Cognitive Science, Linguistics, and Neuroscience. At Princeton, Kennedy works at the Writing Center, conducts research at the Princeton Baby Lab, and serves as a PAA in Mathey College. She wrote this paper as a senior.

Annabelle Duval ’23 is a sophomore pursuing a concentration in History and is interested in Environmental Studies and Gender and Sexuality Studies. In addition to her work for *Tortoise*, Annabelle is an Associate Features Editor for *The Daily Princetonian* and President of the Club Field Hockey Team. She wrote this commentary as a sophomore.

© Spring 2021 Tortoise