Object Bias Disrupts Rule-Based Generalization in Adults Across Domains

Evan Orticio (eorticio@gmail.com)
Department of Psychology, Swarthmore College
500 College Ave., Swarthmore, PA 19081

Stella Christie (christie@tsinghua.edu.cn)
Tsinghua Laboratory of Brain and Intelligence and Department of Psychology, Tsinghua University
Tsinghua Laboratory of Brain and Intelligence, Tsinghua University, Beijing, China, 100084

Abstract

Humans are remarkably adept at abstract rule learning, but little is known about when learners apply this knowledge. We investigated a fundamental constraint in rule generalization: attention to featural similarity (object bias). Across two experiments in different domains, we asked whether adults' abstract rule generalization is constrained by superficial matches to the concrete exemplars present during learning, as is known to be the case for analogical reasoning (Gentner & Toupin, 1986). In the present studies, participants were exposed to a series of sequences following a simple rule and were asked to generalize to novel instances of either the same rule or a new rule. In one condition, an individual element present during initial learning was inserted into the new, unfamiliar pattern. Results showed that adults often chose this object match over the rule match, suggesting that abstract rule generalization, like analogical reasoning, is impacted by concrete features of the input.

Keywords: abstract rule learning; generalization; analogy; learning bias; object match

Introduction

Rule learning—the ability to extract abstracted, structural regularities from a set of exemplars and extend them to novel instances—is a core ability in the human cognitive toolkit. Such an ability is essential because it allows the learner to apply relevant existing knowledge to new stimuli and events whose particular surface characteristics are otherwise unfamiliar. Rule learning has been argued to play a critical role in a large breadth of domains, including categorization (Kruschke, 1992), object recognition (Biederman, 1987), and language acquisition (Pinker 1994). The developmental origins of rule learning are early: 3-month-old infants have been shown to learn abstract relations among sequences of shapes (Ferguson, Franconeri, & Waxman, 2018), and many other experiments have highlighted infants' rule learning abilities in a variety of domains (e.g., Marcus, Vijayan, Bandi Rao, & Vishton, 1999; Gomez & Gerken, 1999; Saffran, Pollak, Seibel, & Shkolnik, 2007).

While extensive work has shown humans' *ability* to learn abstract rules, few studies have looked at *how* rule abstraction actually operates in everyday learning. To do so, it is necessary to understand the constraints and limitations of rule learning—specifically, under what circumstances does rule learning fail? For example, limitations in attention and memory constrain the kinds of sequences that can be

abstracted throughout ontogeny (Endress, Nespor, & Mehler, 2009; Johnson et al., 2009). Schonberg, Marcus, and Johnson (2018) showed that 11-month-olds fail to learn an abstract repetition rule when the repetition occurs in an internal position (e.g., ABBC) or in variable positions (AABC, ABBC, or ABCC). Similar constraints also operate in adults' rule learning: in a sequence learning task, participants succeeded in generalizing a sequence-final repetition rule (ABCDE*FF*) but not a sequence-internal rule (ABC*DD*EF) (Endress, Scholl, & Mehler, 2005). Crucially, follow-up experiments showed that participants could discriminate between different sequences with repetitions (e.g., AB*CC*DEF vs. ABC*DD*EF), demonstrating that position salience was actually a constraint on generalization per se, and not the ability to perceive sequence-internal repetitions.

Additional constraints on the generalization of a rule arise from the fact that rule learning relies on the same process of abstraction that is involved in analogical reasoning. The acquisition of a simple rule like ABA requires the recognition of similarity between the internal relationships of multiple individual exemplars. Rule learning can thus be characterized as a process of iterative analogical reasoning between the previous input and a newly encountered exemplar. A consequence of this is that superficial properties of individual exemplars may interfere with rule generalization. The relational learning literature has observed a robust tension between object similarity and structural similarity. For example, given a sample AA, 4-year-olds were more likely to choose an object match AB over the relational match CC (Christie & Gentner, 2007).

It is well documented that focus on salient individual objects can directly interfere with relational abstraction (Gentner & Rattermann, 1991; Bulloch & Opfer, 2009; Paik & Mix, 2006). This object bias is present from infancy (Quinn, Polly, Furer, Dobson, & Narter, 2002; Ferry, Hespos, & Gentner, 2015) and persists throughout childhood. While the ability to ignore object matches in favor of relational matches improves with increasing age (Richland, Morrison, & Holyoak, 2006), adults are nevertheless susceptible to similar object biases in analogical tasks. Object matches become particularly tricky when items are cross-mapped, i.e., superficially similar objects occupy different roles in the relational structure across analogs (Gentner & Toupin, 1986). Because analogical reasoning operates as a function of the shared

activation of features, salient surface features influence the selection of an analog (Holyoak & Koh, 1987; Ross, 1989).

Given that an abstract rule is a set of learned relations between objects, there is good reason to suspect that the same constraints imposed on relational learning may also apply in rule learning. Indeed, the Marcus et al. (1999) findings have been simulated using an unsupervised model of analogical generalization (Kuehne, Gentner, & Forbus, 2000). Given the same data as in Marcus et al.'s study, the model generated incrementally abstracted representations via subsequent comparisons. Importantly, the model produced "rule-like" behavior despite retaining some concrete features in its final representations. The notion that a rule is achieved via structural alignment, without introducing fully abstracted algebraic variables, carries implications for its generalization. In particular, the learner should be drawn to featural similarities between novel stimuli and the concrete particulars still latent in her representation of the rule, leading to object match generalizations at the expense of abstract relational generalizations.

Do adults exhibit a similar object bias during abstract rule learning? Apart from the simulation study by Kuehne et al. (2000), to our knowledge no prior work has directly asked whether the object bias constrains rule learning. Given the ubiquity of object similarity—the world is full of items, events, and situations that share featural similarities— knowing about such a constraint provides a more nuanced picture of how rule generalization works (or fails) in everyday learning contexts. To test this, we adopted the standard rule learning methodology and stimuli from Marcus et al. (1999), but with the critical modification of pitting rule matches against object matches at test. That is, after being familiarized to exemplars of a certain rule, at test participants had to choose between the familiar rule and a novel rule that contained an object match. In two experiments using sound and visual stimuli, respectively, we asked whether adults' rule abstraction is affected by the presence of object matches.

Experiment 1

In Experiment 1 we tested whether adults' generalization of a rule is independent or constrained by the individual, concrete entities that were present during learning. To do this, we created an experiment closely modeled after that of Marcus et al. (1999). Participants were familiarized with an ABA (e.g., ga ti ga) or an ABB (e.g., li na na) sentence. In the Control condition, participants were tested with new sentences following either the familiar rule or a novel rule, similar to standard rule learning experiments. Crucially, however, participants in the Object Match condition had to choose between the familiar rule and a novel rule containing a familiar object (e.g., wo *la la*, where the participant was familiarized with ABA strings, some of which contained the syllable 'la').

Participants

Two hundred English-speaking participants were recruited through Prolific to participate in a 7-minute online experiment. Thirteen participants were excluded for failing one or more of three catch trials, described below. This yielded a final sample of 187.

Stimuli and Procedure

The experiment consisted of three catch trials, a familiarization phase, and an eight-trial test phase. Catch trials were created to assess the participants' attention and familiarize them with the procedure. Participants were shown two identical images of some location and listened as two different animal sounds played, each associated with the left or right image. The participants were instructed to select the location of a given animal based on the sounds that they heard. For example, one catch trial asked participants to find the location of a horse. Two barns were shown on the screen: for the left barn, a horse neigh played, and for the right barn, a pig oink played. These trials were intended to verify that the participant's device was playing sound and to elicit attention, as the sounds could not be replayed.

Participants were subsequently given a cover story in which they were tasked with leading a lost cartoon bunny, Bella, to her family. In order to identify Bella's family members, participants were instructed to listen to her family song, because "every family of bunnies has its own song." Relational language, e.g., "same song," was explicitly avoided in the cover story so as not to elicit additional bias toward relation-based analyses of the familiarization song.

The familiarization was a 2-minute speech sample closely modeled after that of Marcus et al. (1999). This sample contained three repetitions each of 16 three-syllable strings following either an ABA or ABB sequence, counterbalanced across participants. Participants heard only one of the patterns, ABA or ABB, during this familiarization phase. Strings following the ABA rule included "ga ti ga" and "li na li," for instance. In both samples, strings were separated by 1 second pauses.

After listening to this familiarization, participants completed eight test trials in which two bunnies were each shown to "sing" a single three-word sentence, one at a time. The participants were asked to generalize from the familiarization stimuli by selecting which of two otherwise identical bunnies belonged to Bella's family. The content of these choice sentences differed according to whether the participant was assigned to the Control condition or the Object Match condition. In the Control condition, all sentences consisted of novel syllables that did not appear in the familiarization. For each trial, one sentence followed the same rule as the sentences in the familiarization, and the other followed the opposite rule, which was novel for the participant. For instance, a participant in the ABA condition chose between "wo de wo," a "rule match" sentence, and "wo de de," an "opposite" sentence. In the Object Match condition, the sentences following the opposite pattern were changed to include one syllable from the familiarization, an

object match, alongside one novel syllable. The rule sentences remained the same, containing only novel syllables. Thus, a participant in this condition that was familiarized to an ABA pattern was given a choice between, e.g., "wo de wo" and "wo la la." The object match syllable was "la" for half of the trials and "li" for the other half. This syllable always occurred in the position that repeated within the sentence, i.e., A in ABA and B in ABB.

Results

Overall performance did not differ between counterbalance conditions (ABA vs. ABB rule), so we collapsed responses for the following analyses. On average, participants in the Control condition selected the rule sentence on 6.14 of 8 test trials (SD = 1.97, 95% CI: 5.73 – 6.55). This is significantly above than chance responding ($t(92)$ = 10.473, p < .0001). That is, after hearing 2 minutes of either ABA or ABB sentences, adults can easily abstract the rule and generalize it to sequences of novel syllables. In sharp contrast, participants in the Object Match condition chose the rule sentence on an average of 1.37 trials of 8 (SD = 2.55, 95% CI: 0.85 – 1.90), a rate significantly below chance ($t(93)$ = -9.98, p < .0001). Our primary interest is the effect of the presence of object match distractors on the rate of choosing the rule match. A one-way ANOVA revealed that participants in the Object Match condition indeed selected the rule match sentence on fewer trials than those in the Control condition, $F(1,185)$ = 204.1, p < .0001. These results strongly suggest that object matches were a preferred basis for generalization compared to rule matches.

To assess individual-level performance for consistency in response behavior, participants were sorted into three categories: rule choosers, opposite choosers, and mixed choosers. To be classified as a rule chooser, a participant had to choose the rule sentence on at least 7 of 8 trials (binominal p < .05); opposite choosers chose the rule on at most 1 trial. Participants with any other response pattern were classified as mixed choosers. Among those in the Control condition, 50 (53.8%) were rule choosers, 40 (43.0%) mixed choosers, and 3 (3.2%) opposite choosers. The Object Match condition consisted of 10 (10.6%) rule choosers, 13 (13.8%) mixed choosers, and 71 (75.5%) opposite choosers (i.e., object choosers). That is, in the Object Match condition, a majority of participants consistently selected the object match over the relational rule. A chi-square test comparing the number of rule choosers and a collapsed category of non-rule choosers between conditions confirmed that participants were less likely to consistently select rule sentences with the option of an object match sentence, $\chi_2(1)$ = 37.9, p < .0001.

Experiment 2

In Experiment 2 we seek to replicate the results of Experiment 1, in which adults generalized a simple ABA or ABB sequence based on surface-level object matches as opposed to a rule-based match. Additionally, we are testing this effect in a different domain using visual shapes, as some evidence indicates that speech may be a privileged domain with regard to abstract rule learning (Marcus et al. 2007; Rabagliati, Senghas, Johnson, & Marcus, 2012).

Participants

Two hundred participants were recruited to participate in an online experiment through Prolific. Participants from Experiment 1 were ineligible to participate. After excluding eight participants for failing one or more of three catch trials, the final sample was 192.

Stimuli and Procedure

Experiment 2 was created to be as alignable of a replication of Experiment 1 as possible. It consisted of three catch trials, familiarization, and 8 test trials, but now using visual stimuli. For each of the three catch trials, the participant was instructed to choose a location appropriate for a given action, e.g., swimming. Upon pressing a key, two images (e.g., a swimming pool and a soccer field) appeared one at a time on either side of the participant's screen. Each image appeared alone for 1700 ms, then disappeared. This is the same duration as the three-syllable strings in the familiarization from Experiment 1. After the second image had disappeared, the participant chose the appropriate image by pressing a key corresponding to the position in which the image had appeared. The catch trials were intended to train participants to attend to the ephemeral visual stimuli.
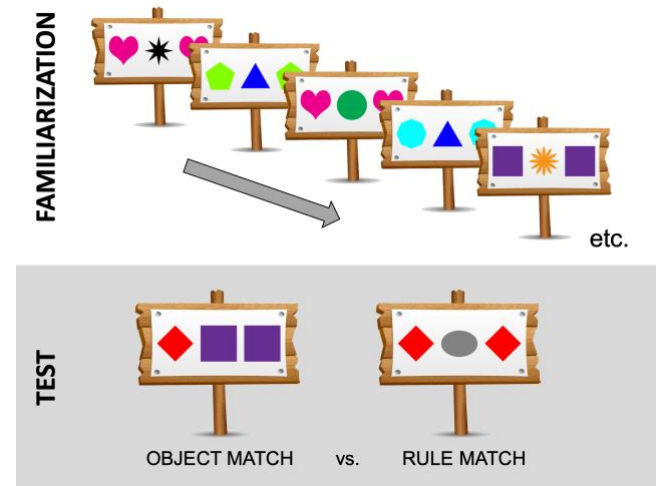

Figure 1: Schematic of familiarization and test phases of Experiment 2. Participants were exposed to a 2-minute stream of exemplars, and chose between a rule match sequence and object match sequence at test. While the full three-shape patterns are shown together here, shapes appeared one at a time during the experiment.

Participants read a cover story similar to that in Experiment 1. Participants were to help Bella, a bunny lost in the forest, find her way home by learning which signs represented safe paths. They were then familiarized with a 2-minute series of signs deemed safe for Bella. Each sign

consisted of three shapes, which followed either an ABA or ABB pattern, counterbalanced across participants. Each shape had its own color, so signs included patterns like "pink heart, blue triangle, pink heart" for the ABA rule or "purple square, green circle, green circle," for the ABB rule. Each pattern appeared on an otherwise blank signpost. The three shapes appeared and disappeared one at a time for 600 ms, roughly the duration of an average syllable from Experiment 1. This sequential presentation was designed to parallel the temporal structure of the sentences in Experiment 1. The first shape appeared on the left side of the signpost, followed by the second shape in the middle and the third shape on the right side. Thus, the signs had spatial structure as well as temporal structure, as seen in Figure 1. Each of the 16 signs was repeated three times in a pseudorandom order during familiarization, with a 1 second gap between each new sign.

In 8 test trials, participants were instructed to choose the sign that represented a safe path for Bella. Participants were assigned to either the Control or Object Match condition. For each trial, participants in the Control condition chose between two signs consisting entirely of novel shapes with novel colors. One sign, the rule match, followed the same pattern as familiarization; the other, an opposite match, followed the other pattern which was unfamiliar to the participant. In the Object Match condition, the opposite match was altered to include one shape taken from the participant's familiarization sequence, like in Experiment 1. The patterns were presented one at a time in the same sequential, disappearing pattern as the familiarization, again to align with the temporal structure of the sentences in Experiment 1. Participants made their choice by selecting a key corresponding to the location in which the target sign had appeared.

Results

Performance did not differ between counterbalanced conditions (ABA vs. ABB rules), so responses are again collapsed. One-sample t-tests revealed that the average number of trials in which participants in the control condition chose the rule match (M = 7.20 of 8, SD = 1.55, 95% CI: 6.88 – 7.51) was significantly above chance ($t(96)$ = 20.28, $p < .0001$), whereas the average for those in the object match condition (M = 4.26, SD = 3.36, 95% CI: 3.58 – 4.95) did not differ from chance level ($t(94)$ = 0.76, p = 0.45). A one-way ANOVA revealed that participants in the object match condition chose the rule match less frequently than those in the control condition ($F(1,190)$ = 60.81, $p < .0001$), replicating the primary effect of Experiment 1.

Individual participants were again classified as rule choosers, opposite choosers, or mixed choosers according to the criteria used in Experiment 1. In the Control condition, there were 79 rule choosers (81.4%), 17 mixed choosers (17.5%), and 1 opposite chooser (1.0%). Among those in the Object Match condition, 40 of 95 (42.1%) were rule choosers, 23 (24.2%) were mixed choosers, and 32 (33.7%) were opposite choosers. A chi-square test showed that there

was a significant association between the presence of object match distractors and being a non-rule chooser, $\chi_2(1)$ = 29.9, $p < .0001$.

Comparison of Experiments 1 and 2

Did the salience of object matches differ between the two experiments? To test this, we conducted a 2 (Experiment: 1 vs. 2) x 2 (Condition: Control vs. Object Match) ANOVA, with number of trials choosing the rule match as the dependent variable. This analysis revealed a main effect of Experiment ($F(1,375)$ = 63.17, $p < .0001$), with participants choosing the rule match at a higher rate overall in the visual domain. A main effect of Condition confirms that participants chose the rule match sentence less frequently in the object match condition, $F(1,375)$ = 232.31, $p < .0001$, whether they were presented with sound (Experiment 1) or visual stimuli (Experiment 2). However, there was also a significant interaction between Experiment and Condition, $F(1,375)$ = 13.27, $p < .001$, such that the effect of the object match condition was stronger for the syllable stimuli as opposed to the shape stimuli (Figure 2).
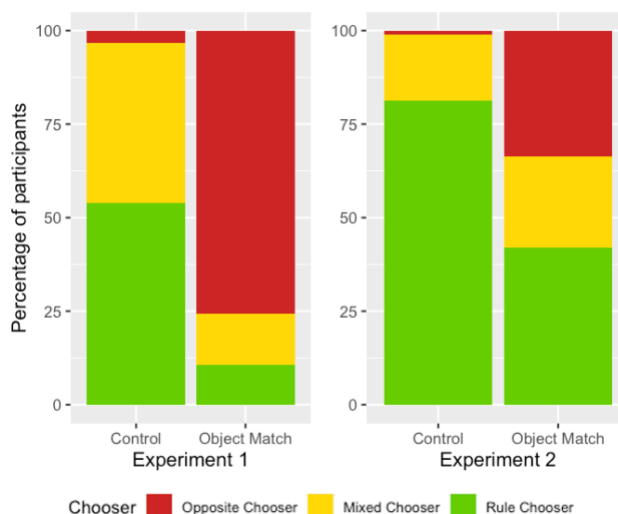

Figure 2: Stacked bar plots show proportions of rule choosers, mixed choosers, and opposite choosers in each condition of Experiment 1 and 2.

General Discussion

In two experiments, we investigated whether the generalization of an abstract rule is constrained by superficial matches to the concrete elements present during learning, in linguistic and visual domains, respectively. While the tension between object and structural similarity has been robustly demonstrated in relational learning, this is the first time that this issue has been investigated using a rule learning paradigm. Participants were familiarized to a 2-minute stream of sequences following an ABA or ABB rule, and were asked to generalize this rule to sequences composed of novel elements following either the same rule

or the unfamiliar, opposite rule. Critically, for half of the participants, the opposite rule sequence contained an element from the familiarization.

Our results underline two main conclusions. First, we replicated prior findings that adults, like infants, are capable of extracting abstract structural rules from a mere 2-minute exposure. Consistent with prior experiments (e.g., Marcus et al., 2007; Frank, Slemmer, Marcus, & Johnson, 2009; Rabagliati, Ferguson, & Lew-Williams, 2019), this ability is domain-general—we see this both with syllables (Experiment 1) and shapes (Experiment 2). Second, we obtained the novel finding that the process of rule learning can be constrained in a significant way by the presence of object matches. While participants in control conditions chose the rule match at high rates across the two experiments, when there were choices between object and rule matches, a plurality (if not a majority) of adults chose objects over rules. That is, adults are perfectly *capable* of acquiring these rules, but they do not always make use of this abstract knowledge under certain learning conditions. Our results suggest that object matches serve as a prominent basis for generalization across modalities, even overcoming rule-based generalizations in the case of syllables.

Why are object matches so alluring during generalization? One possibility is that the content of the "rules" participants learned during familiarization was not fully abstracted, but rather still contained traces of concrete features of the exemplars from the input, as in the simulation study (Kuehne et al., 2000). This predicts that object matches would be a salient basis for generalization due to the featural similarity to previously encountered input. Indeed, young children often display an object bias before undergoing a gradual relational shift (Gentner, 1988), and humans show a heightened sensitivity to object similarity compared to close ape relatives (Christie, Gentner, Call, & Haun, 2016). Consequently, because processing individual elements disrupts relational processing (Sloutsky & von Spiegel, 2004), rule-based generalization would suffer. Experimental evidence shows that exemplar-specific information plays a role in rule generalization. In an artificial grammar learning experiment, participants exhibited a decrease in performance when test items contained a different set of letters from the training set, compared to when test items were distinct from the training set but contained similar letters (Knowlton & Squire, 1996).

It is also possible that, regardless of the abstractness or exemplar-specificity of the participants' representations, object similarity per se was simply a preferred basis for generalization for many participants. The present data cannot distinguish between these two possibilities. However, given that adults overwhelmingly choose relations over object matches in relational reasoning studies with similar paradigms, it is unlikely that our experimental context was unique enough to elicit such elevated preferences for superficial featural similarities in the presence of a simple relation-based alternative. Regardless, these results represent a novel demonstration of the fallibility of adult rule generalization, which bears a previously underappreciated resemblance to a well-known constraint on analogical reasoning.

Although the primary effect of the object bias on rule-based generalization was replicated across the two experiments, it is worth noting that effect sizes differed. The interaction between condition and experimental domain (syllables vs. shapes) indicates that adults who heard syllables are relatively more likely to choose the object match and forgo the rule abstraction compared to those who saw shapes. There are at least two possible explanations for this unexpected result. First, one cover story may have encouraged more attention to surface features than the other. One may argue that the use of a rabbit "song" as a cover in Experiment 1 might have biased participants toward a concrete, sound-based generalization strategy, whereas signs are inherently symbolic. However, a majority of the participants in the control condition, who were given the same cover story, were still successful in making rule-based generalizations; there is no *a priori* reason for an abstract rule to be a dispreferred basis for generalization in what is ostensibly a "song matching" task.

The second, more interesting possibility is that while there is no reason to devote isolated attention to any particular shape within a sequence in Experiment 2, adults may have a bias to attribute meaning to individual syllables. Because syllables are part of our everyday linguistic system, they may automatically engage language comprehension networks in a way that invites more individualized attention than shapes. There is ample evidence that language processing is automatic and involuntary (Pulvermüller, Shtyrov, & Hauk, 2009; Shtyrov, 2010). In fact, brain responses to pseudowords are actually stronger than those to meaningful words under explicit instructions to listen, as in the familiarization stage of our task (Shtyrov et al., 2012). This suggests that object matches may be particularly salient for speech-based patterns, which has implications for word segmentation, grammar learning, and a variety of other tasks of language acquisition. For example, these results raise the possibility that for adult second language learners, strong attention to the semantics of individual words may in fact interfere with grammatical (rule) abstraction.

Two discussion points are noteworthy given this object bias. First, the object bias effect in our rule learning task is stronger than that has been typically found in prior work in analogical learning. In many analogical tasks (e.g., Markman & Gentner, 1993), adults are affected by object similarity, but rarely prefer it over relational similarity. We hypothesized that the presence of object matches would affect rule learning, but not that it could reverse the preference. The object match is much less obvious in the implicit rule learning task, in which the object match involves comparison to the declarative memory of a learned exemplar, than in a triad task, where object matches are visually co-present. Furthermore, the rule itself is arguably more obvious in the current task; there is only one rule to learn as opposed to many different structures within the

typical analogy paradigm. Given these considerations, it is surprising that the object bias had such a strong effect on adults' rule learning. We suspect that there are further implications that follow this tension between object and rule.

The second point concerns the role of salience. Prior work exploring the constraints on rule learning suggest that the *lack* of salience can limit rule learning (e.g., the low salience of medial positions; Endress et al., 2005). Our results suggest that this is not the full story: *too much* salience (resulting from an object match to prior input) can also have a direct impact on rule generalization. Again, this phenomenon is well known in the analogy literature, but is novel within the rule learning context. We suspect that there is a sweet spot between the salience and memory of individual items and the abstraction of a rule that binds them. The differential effects observed between the syllable and shape stimuli hint that this trade-off curve may differ across domains.

General Discussion

In sum, our results reveal that the generalization of an abstract rule is subject to the same fundamental constraint that shapes analogical reasoning: attraction to object similarity. Across auditory and visual domains, adults can be persuaded to generalize on the basis of superficial featural similarities despite the presence of an abstract, relation-based alternative that is demonstrably learnable. This finding is striking considering that even infants are able to abstract similar rules within the first few months of their life. Adults' susceptibility to object similarity during rule generalization demonstrates that the process of abstracting and generalizing a rule is deeply dependent on the concrete exemplars from which it arises. Our study highlights the importance of investigating not only the kinds of rules that can be learned, but the conditions under which learned rules are used and generalized.

References

Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological Review*, *94*(2), 115.

Bulloch, M. J., & Opfer, J. E. (2009). What makes relational reasoning smart? Revisiting the perceptual-to-relational shift in the development of generalization. *Developmental Science*, *12*(1), 114-122.

Christie, S. & Gentner, D. (2007). Relational similarity in identity relation: The role of language. In Vosniadou, S. & Kayser, D. (Eds.). *Proceedings of the Second European Cognitive Science Conference*.

Christie, S., Gentner, D., Call, J., & Haun, D. B. M. (2016). Sensitivity to relational similarity and object similarity in apes and children. *Current Biology*, *26*(4), 531-535.

Endress, A. D., Scholl, B. J., & Mehler, J. (2005). The role of salience in the extraction of algebraic rules. *Journal of Experimental Psychology: General, 134*(3), 406.

Endress, A. D., Nespor, M., & Mehler, J. (2009). Perceptual and memory constraints on language acquisition. *Trends in Cognitive Sciences, 13*(8), 348-353.

Ferguson, B., Franconeri, S. L., & Waxman, S. R. (2018). Very young infants learn abstract rules in the visual modality. *PloS one*, *13*(1).

Ferry, A. L., Hespos, S. J., & Gentner, D. (2015). Prelinguistic relational concepts: Investigating analogical processing in infants. *Child Development, 86*(5), 1386-1405.

Frank, M. C., Slemmer, J. A., Marcus, G. F., & Johnson, S. P. (2009). Information from multiple modalities helps 5-month-olds learn abstract rules. *Developmental Science*, *12*(4), 504-509.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*(2), 155-170.

Gentner, D. (1988). Metaphor as structure mapping: The relational shift. *Child development*, 47-59.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. *Perspectives on Language and Thought: Interrelations in Development*, *225*.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, *10*(3), 277-300.

Gomez. R. L., & Gerken, L. (1999). Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition, 70*, 109-135.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition*, *15*(4), 332-340.

Johnson, S. P., Fernandes, K. J., Frank, M. C., Kirkham, N., Marcus, G., Rabagliati, H., & Slemmer, J. A. (2009). Abstract rule learning for visual sequences in 8-and 11-month-olds. *Infancy, 14*(1), 2-18.

Knowlton, B. J., & Squire, L. R. (1996). Artificial grammar learning depends on implicit acquisition of both abstract and exemplar-specific information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(1), 169.

Kruschke, J. K. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychological Review, 99*, 22–44.

Marcus, G. F., Fernandes, K. J., & Johnson, S. P. (2007). Infant rule learning facilitated by speech. *Psychological science*, *18*(5), 387-391.

Marcus, G. F., Vijayan, S., Rao, S. B., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science, 283*(5398), 77-80.

Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology, 25*(4), 431-467.

Paik, J. H., & Mix, K. S. (2006). Preschoolers' use of surface similarity in object comparisons: Taking context into account. *Journal of Experimental Child Psychology*, *95*(3), 194-214.

Pinker, S. (1994). The language instinct. New York: Morrow.

Pulvermüller, F., Shtyrov, Y., & Hauk, O. (2009). Understanding in an instant: neurophysiological evidence for mechanistic language circuits in the brain. *Brain and Language*, *110*(2), 81-94.

Quinn, P. C., Polly, J. L., Furer, M. J., Dobson, V., & Narter, D. B. (2002). Young infants' performance in the object-variation version of the above-below categorization task: A result of perceptual distraction or conceptual limitation?. *Infancy*, *3*(3), 323-347.

Rabagliati, H., Senghas, A., Johnson, S., & Marcus, G. F. (2012). Infant rule learning: advantage language, or advantage speech?. *PLoS One, 7*(7).

Richland, L. E., Morrison, R. G., & Holyoak, K. J. (2006). Children's development of analogical reasoning: Insights from scene analogy problems. *Journal of Experimental Child Psychology*, *94*(3), 249-273.

Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(3), 456.

Saffran, J. R., Pollak, S. D., Seibel, R. L., & Shkolnik, A. (2007). Dog is a dog is a dog: Infant rule learning is not specific to language. *Cognition*, *105*(3), 669-680.

Schonberg, C., Marcus, G. F., & Johnson, S. P. (2018). The roles of item repetition and position in infants' abstract rule learning. *Infant Behavior and Development*, *53*, 64-80.

Shtyrov, Y. (2010). Automaticity and attentional control in spoken language processing: neurophysiological evidence. *The Mental Lexicon*, *5*(2), 255-276.

Shtyrov, Y., Smith, M. L., Horner, A. J., Henson, R., Nathan, P. J., Bullmore, E. T., & Pulvermüller, F. (2012). Attention to language: novel MEG paradigm for registering involuntary language processing in the brain. *Neuropsychologia*, *50*(11), 2605-2616.