Do graded representations support abstract thought?

Alexandra Carstensen & Michael C. Frank

Department of Psychology, Stanford University

Correspondence concerning this article should be addressed to Alexandra Carstensen,
Department of Psychology, Stanford University, Building 420, 450 Jane Stanford Way, Stanford,
CA 94305. Email: abcarstensen@stanford.edu.

Abstract

Relational reasoning requires the reasoner to go beyond their specific experience, abstracting from items to make inferences about categories and kinds on the basis of structural or analogical similarities. Reasoning about the relations *same* and *different* is one of the best-studied cases of relational reasoning, both across species and over human development, and has become a paradigm case study for abstract representation. However, decades of careful study have nonetheless produced seemingly contradictory findings—with surprising successes and puzzling failures—in both the comparative and developmental literatures. In this article, we review these literatures and suggest first steps towards a reconciliation of these contradictions by suggesting that same-different reasoning is supported by graded representations.

Introduction

Humans are a unique species, and one of our most unique traits may be our penchant for abstract thought. Concepts like *same* and *different* are prototypical examples of abstractions: their meaning transcends perceptual modalities (e.g., *same* in flavor, texture, or color, like blue sky and blue eyes) and domains of knowledge (*same* in kind, purpose, or number, like ten days and ten men). Words with this degree of abstraction are commonplace in human languages and are often learned early in childhood. While non-human animals may be able to represent abstract concepts like *same* and *different*—the evidence is mixed, as we will review below—there is a broad consensus that humans do so far more frequently and readily, especially in adulthood. Same-different reasoning is one of the best-studied cases of abstract thought, both across species and over human development, and has thus become a paradigm case study for abstract representation.

Questions about same-different reasoning are typically instantiated through disarmingly simple experimental paradigms. In a match-to-sample (MTS) task, participants learn to match a cue, like a blue square or a lemon scent, to an identical target instead of a distractor, i.e., another blue square or lemon scent, not a red triangle or mango scent. In more complex relational match-to-sample tasks (RMTS), participants are cued with a pair of stimuli that exhibit a relation, like *same* (AA), which they must match to a target pair that also exhibits the *same* relation (BB, not CD). While MTS tasks are analogous to a variety of matching behaviors, more interest has been focused on RMTS because it is arguably closer to some of the rich analogical inferences that characterize human cognition [1]. Despite the simplicity of these tasks, they have produced seemingly contradictory findings in both the comparative and developmental literatures. Our goal in this article is to review these literatures and suggest first steps towards a reconciliation of

these contradictions by suggesting that same-different reasoning is supported by graded representations.

A standard, symbolic account of same-different reasoning posits the existence of a mental representation of the operator $\text{SAME}(A,B)$ ¹ that produces a truth value if A and B are the same [2]. Such symbolic accounts are inflexible in the sense that this operator is either present or absent in any organism. The data make this kind of account hard to sustain; to take one example, bees succeed in a task logically equivalent to one in which school-aged children fail [3-4]². Accommodating such data in a purely symbolic account requires significant contortions to explain why the children fail. Further, a symbolic account loses explanatory power if it requires granting symbolic operators to every species that shows any relevant same-different performance.

Despite the weaknesses of binary symbolic accounts, and the success of graded accounts in other domains [5], graded accounts of same-different reasoning have received far less attention in the literature (cf. [6]). One major reason for this neglect has been the argument that a class of neural networks do not show convincing performance in approximating SAME functions ([7]; reviewed in [8]). But in the twenty years since this original argument, there has been substantial progress, and recent years have brought a range of neural network architectures that can learn graded relational functions (e.g., [9-12]). In their simplest form, such graded functions approximate $\text{SAME}(A,B)$ and can be applied successfully to novel, un-trained stimuli but with varying effectiveness for different types of input. We take some of the key features of graded representations to be that they (1) return values that are not binary but may be closer or farther

¹ The operator SAME could be represented equivalently as = or “identity”; we use SAME here for simplicity.

² NB: there is substantial variation between task designs that are behaviorally appropriate for bees and human children, which has fueled debate over the equivalence of these tasks; Giurfa et al. [3] maintain that their delayed MTS task is logically equivalent to RMTS.

from 0 and 1 depending on the similarity of the inputs, leading to “imperfect” task performance, (2) are at least partially bound to particular training stimuli, leading to graded generalization (i.e., poorer performance) with new stimulus sets, (3) show dose-response effects of training such that performance increases with more training, (4) generalize more robustly when trained on diverse exemplars, and (5) may maintain information about other representational dimensions (e.g., object identity) concurrent with relational information about sameness/difference. We argue that the comparative and developmental data strongly support gradedness of this sort (Box 1 provides some discussion of whether there are intermediate, graded symbolic accounts that can explain the data).

From a comparative perspective, the key puzzle is this: many non-human species succeed in simple tasks requiring a generalizable notion of *same*³ like MTS, and several have met and surpassed even the most strict criteria for RMTS. But these clear successes, particularly in RMTS, are usually seen in animals that have completed rigorous training with tens of thousands of trials or spent years learning in close contact with humans. Considering all of these successes, why don’t we see the quick and obvious success characteristic of human adults in any other species? Our answer is that all organisms need tens of thousands of training trials to increase the precision and flexibility of their representations of *same* and *different*, humans included. Humans learn more quickly because their experiences provide a type of pre-training that allows them to adapt their pre-existing representations easily to novel same-different tasks.

With respect to the developmental data, the puzzle concerns task performance. Children can abstract the *same* relation in early infancy (in tasks that are akin to MTS), yet they struggle

³ While success in MTS may suggest a representation of the relation *same*, Premack [13] argues that MTS and pairwise judgements of *same* and *different* measure different abilities, as animals that readily pass MTS do not readily distinguish between same and different pairs of objects. Accounts of the *same* relation that define it as operating over multiple items therefore caution that MTS success is not sufficient to imply representations of such a *same* relation.

in performing relational and analogical tasks into middle childhood, with performance even declining with age in some cases. Why is human relational reasoning so unreliable? We argue that children's representations are graded and not fully general and so they are undermined by a variety of conflicting cues (often to greater or lesser extents in different tasks and contexts).

In the remainder of this paper, we review each of these puzzles in more depth before turning back to the notion of graded representations in the conclusion.

Same-Different Abstraction Across Species

The simplest type of same-different task, MTS, can be passed with a small amount of training by members of a wide range of species (e.g., chimpanzees [14], honeybees [3], rhesus monkeys [15], parrots [16]). Even newly hatched ducklings spontaneously imprint on generalizable same-different patterns (following BB but not AB after imprinting on AA [17]). In contrast, the more complex RMTS task requires an ability to judge the relation between relations (e.g., to identify two matching objects, AA, as related *in the same way* as the pair of objects BB). Unlike MTS tasks, RMTS tasks have been passed by non-human species only after extensive training. Further, RMTS performance in non-human species typically asymptotes to levels far below ceiling performance [18].

Sarah, a language-trained chimpanzee, succeeded in the RMTS task [19], initially suggesting that symbolic practice was a critical ingredient in solving the task [13]. But three other language-naïve chimpanzees performed equivalently to Sarah in this task after training in which they learned symbols for same (AA) and different (CD) pairs [20]. These successes can be compared to failures in that study by a naïve chimpanzee and by a set of similarly token-trained rhesus monkeys [15, 21].

A set of experiments with baboons and pigeons that used extensive training regimes suggested that other types of training may be sufficient to achieve RMTS success, and that no symbolic training was needed ([22-23]; cf. [21]). The RMTS experiments for these species used arrays of objects (e.g., with 3 or 16 instead of 2 stimuli defining each relation), though, and could be solved using variability or entropy strategies in place of relational reasoning [24-25].⁴ Followup studies with both crows and parrots meet more stringent criteria: they show success in 2-item RMTS tasks, in conditions where only the kind and not the number of shared relations indicates the correct choice [26-27]. Further, animals in these experiments succeed spontaneously, without differential food reinforcement on any of the RMTS trials.

Even this brief review of a complex literature suggests that many of the obvious dividing lines between human and non-human performance have been gradually erased by clever experimental protocols. These protocols use a variety of strategies to promote performance, ranging from pre-training on specific feature comparisons [26] to sheer brute-force training quantity [28]. While these strategies may seem superficially different, they all appear to promote the creation of representations with generality sufficient to pass the relevant test trials. For example, both baboons [22] and pigeons [23] perform better in RMTS with familiar than with novel stimuli, despite the fact that familiarity is irrelevant for solving the RMTS problem. Indeed, the more closely this literature is examined, the more clearly evidence appears that abstract representations in these tasks are graded, rather than fully general, in that they are at least partially bound to particular training stimuli [6]. This evidence comes in the form of the

⁴ On the strongest version of this view, even success with 2-item RMTS is insufficient evidence for true relational reasoning, which must show sensitivity to the *kind* of same or different relation between pairs, not just the number of differentiating features [24]. By this strict account, only Sarah's success on functional analogies (not her RMTS performance or that of the token-trained chimpanzees) can be taken as support for a cognitive divide between great apes and other animals (but cf. [4] for an empirical case in favor of stronger interpretations of success with 2-item RMTS).

generalization decrement described above, but also in exemplar diversity effects during training, where training with a larger set of exemplars produces more accurate generalization to novel stimuli, attenuating the generalization decrement [29].

We see this body of evidence as providing support for the hypothesis that non-human species build up representations of *same* and *different* over training. These representations are graded in the following senses. First, they support levels of performance that are above chance but far from perfect (even with extensive training and within trained stimuli). Second, although representations may not be qualitatively different between simpler MTS and more complex relational tasks, the level of generality necessary for relational abstraction appears to require substantially more training. Finally, although this feature is not explicitly tested in most studies, learned representations are sufficiently general to be applied to novel stimuli outside of the training set, but not fully general—in the sense of being applicable to arbitrary new stimuli from different classes.

Same-Different Abstraction in Human Development

The key puzzle in understanding human same-different reasoning is one of robustness. Younger children sometimes outperform older children [30], children and adults can fail to generalize relations across varying stimuli (e.g., [31-32]), and unknown contextual factors play a determining role in successful relational reasoning [33]. We argue that ecological and contextual factors—relevance, salience, familiarity—account for variation in children (and adults). This explanation, while anomalous from a symbolic perspective, is natural from the graded perspective.

Human infants, like many non-human animals, have difficulty mastering experimental paradigms like RMTS that provide the strongest evidence for relational reasoning. However,

they have been shown to learn patterns predicated on repetitions in syllable sequences (“abstract rules” [7, 34]). A meta-analysis of abstract rule-learning experiments confirms that such patterns can be learned across a wide variety of stimulus types, but also that success depends on whether the stimuli have either communicative value or ecological familiarity to the infants [35]. We see this finding as parallel to the pattern of other species’ successes in MTS tasks in ecologically important contexts (e.g., duckling imprinting, bee navigation).

The pattern of successes and failures is similarly mixed for infancy.⁵ For example, Addyman and Mareschal [37] found that human infants’ success in a same-different task was modulated by the experimental setup (looking time with naturalistic picture pairs vs. anticipatory eye movement with simple geometric object pairs) and by perceptual features of the task (specifically, 4-month-old infants succeeded only in the latter task while 8-month-olds succeeded in both). Other studies with human infants have revealed successes in RMTS-derived tasks, but further highlighted substantial variation in performance dependent on specifics of the task, like the number of relational exemplars shown [36, 38], the salience of exemplars [36, 38], and the specific relation (*same* vs. *different*) to be learned [37, 39]. This literature converges with the abstract rule learning literature in showing clear evidence for competence with abstract relations in human infants. However, it also converges with the comparative literature in showing that, while the in-principle capacity for relational reasoning is present, success depends critically on appropriate task contexts.

⁵ Many infancy tasks are derived from RMTS. For example, Ferry, Hespos, and Gentner [36] habituate to pairs of objects (e.g., AA) and then examine dishabituation to relation-consistent or relation-violating pairs (e.g., BB vs. BC), a design they characterize as same-different discrimination. While Ferry et al. argue that this design provides evidence for abstract representations of *same* and *different* relations, linking this task to a classic RMTS requires additional assumptions. While infants could pass by generalizing a *same* relation and applying it, they could also be simply identifying the task of the paradigm as repeated matching (closer to training on an MTS task). We remain agnostic about whether success in this task is equivalent to RMTS or MTS; indeed, more generally it shows the weaknesses of our verbal task description in adequately characterizing the representations underlying performance.

In general, task performance tends to improve over childhood, but relational reasoning is sometimes an exception. Toddlers succeed in a causal variant of the RMTS task but performance decreases with age, declining to chance in the preschool years [30]. A range of minimal interventions can increase performance, however, including prompts to explain, the provision of evidence against alternative solutions, and changes to the context of the experiment that emphasize relational information [40]. In another variation on this paradigm, relational match-to-sample with arrays of objects, children succeed earlier, apparently using perceptual entropy to identify the solution, but fail to do so with pairs of items except in cases where the child employs the linguistic labels “same” and “different” [4]. These studies suggest that young children’s performance is heavily dependent on the specifics of the task, echoing the evidence for task-sensitivity in human infants and other species.

One specific obstacle to same-different reasoning is competition with alternative, non-relational hypotheses. For example, for any pair AA, cues in the form of specific A objects compete with the presence of the *same* relation. Forced-choice variations on RMTS show that toddlers and adults track evidence for both relational (same/different) hypotheses and plausible alternatives [41]. Preschoolers show a baseline bias toward object-based solutions over relational solutions in an ambiguous formulation of the RMTS task [33], which could account for the decline in performance between toddlers and preschoolers. Even adults are waylaid by object matches in comparisons between relational and object match-to-sample paradigms (Kroupin & Carey, unpublished), suggesting that baseline performance in relational reasoning may often be determined by participants’ bias on this dimension.

The use of language as a framing cue or element of the task can modulate biases toward object-based solutions. In adults, relational labels like “inoculation” or “positive feedback loop”

make retrieval of analogous relational structures more likely [32]. And relational language (e.g., learning a novel noun, “truffet,” that refers to pairs of objects) boosts children’s performance in same-different relational reasoning and analogy [42-43]. Conversely, priming object-focused language (e.g., nouns relative to verbs) hurts performance in same-different relational reasoning [44].

Culture, in addition to language, can modulate biases in relational reasoning. The bias toward objects over relations observed in previous studies may be characteristic of children and adults in the U.S., but flipped for their counterparts in Japan, who show greater attention to relations and context than individual objects [45-47]. Similar differences in bias appear between American and Chinese toddlers by 24 months, with American toddlers focusing more on objects in dynamic scenes and Chinese toddlers on actions [48]. Performance in relational reasoning tracks with these biases in China as in Japan. While relational reasoning declines between toddler and preschool-age in the U.S., Chinese children improve their performance over this period, and these cross-cultural differences in performance are accompanied by differences in baseline bias toward relational or object-based solutions in an ambiguous version of the RMTS task [33].

Taken together, these findings on the malleability of relational and analogical performance in children and adults, and the sensitivity of human relational reasoning to a diverse range of perceptual, linguistic, cultural, and contextual factors, suggest (1) that variation in human relational reasoning is driven by context rather than ability, and (2) that preferences and biases in reasoning play an additional role in relational thinking. A graded representation viewpoint naturally accommodates these generalizations, since on such a view, same or different is only one among many competing features of a particular stimulus on which to generalize, and

components of the task and context could naturally prime or bias toward one feature over another.

Conclusions

Same-different reasoning is an important case study for uniquely⁶ human abstract and relational reasoning. While the idea of a symbolic divide at first appears attractive as an account of these abilities, a closer examination of the evidence suggests that there is no “bright line” separating either humans from other species or human infants from adults. Instead, our review of the evidence suggests that in every case, same-different reasoning is dependent on careful training and appropriate stimulus selection, and that it often contends with competing superficial features or biases against relational features. These patterns suggest a graded account of relational reasoning, one that could in principle be supported by a new generation of neural network models [9-12]. On this account, human adults are the beneficiaries of years of pre-training on the relevant representations, such that they perform well by virtue of having practiced extensively across a range of different contexts.

Perhaps the clearest single case study for gradation comes from the wealth of same-different learning studies with pigeons (summarized in [6]). First, pigeons can succeed in completely relational same-different discrimination in an experiment where stimuli are never repeated, showing that they are truly performing relational learning: item information is not necessary for their learning of same-different categories. Second, pigeons show a generalization decrement: after reaching criterion performance in a same-different discrimination task, they perform worse on novel stimuli than familiar, even when the stimulus identity is irrelevant.

⁶ NB: We do not wish to imply that abstract and relational reasoning are uniquely human, rather to observe that the forms and frequency with which humans engage in these types of reasoning (e.g., keeping a social calendar, following a map, speaking metaphorically) can be seen as bizarre and unusual from a typological perspective on animal behavior.

Third, their learning is faster but the generalization decrement is stronger when the pool of training stimuli is small, suggesting a continuous process of abstraction, in which the more exemplars are seen, the more general the representations become. Finally, pigeons encode both item and relation information in an ambiguous same-different discrimination task [49], as do toddlers [41], but they show preferential learning of items relative to relations, similar to (American) preschoolers in ambiguous RMTS [33].

These patterns, which are so clearly attested in pigeons, appear to be present in humans as well. Infants show the same signs of continuous perception and conception (ecological familiarity effects, entropy) and appear to also require vast amounts of training (i.e., years of experience with linguistic forms “same” and “different”; Hochmann et al., unpublished) to achieve the fine-grained distinctions taken as evidence for symbolic representations. A logical step for the developmental literature should be to test for the effects of gradation that have been attested in pigeons.

Our proposal of graded representations of *same* and *different* in non-human animals and human children does not rule out the importance of symbols for *same* and *different* in older children and adults. Indeed, the accessibility of such symbols may complement their pre- and non-linguistic experiences, leading to some of the incredibly fast and efficient generalization behavior shown by adults (e.g., [18, 50]) in much the same way that the presence of explicit symbols for exact number can lead to fast and efficient reasoning about quantity information [51-52]. Our argument is that presence or absence of symbols is not itself a good account of comparative or developmental differences in same-different reasoning. Nevertheless, the relative accessibility of verbal symbols to adults could explain both why adults are so quick to recognize sameness and why they sometimes fail to do so [32].

In sum, we believe the evidence should lead us to reject discontinuous symbolic accounts of relational reasoning. The next challenge for graded accounts will be developing the specifics so that they can account for the wealth of empirical phenomena attested here and make clear novel predictions about new developmental and comparative studies.

Box 1: Are graded symbolic accounts possible?

Graded performance can be approximated within symbolic accounts via the averaging of discrete response functions [53-54]. One perspective on symbolic reasoning suggests that humans engage in a process of gradual abstraction over exemplars (e.g., [55-58]), building symbols that become increasingly abstract and general as a function of experience. Averaging data from these models would produce graded performance like that described in the main text. This kind of account is consistent with symbolic models of relational learning (e.g., [59]), which show many of the key features of graded representations that we review (imperfect performance, stimulus effects, dose response effects, effects of training diversity). Our dissatisfaction with these accounts is Fodorean: they provide no story about the origins of the primitives underlying learning. For example, the model proposed by Frank & Tenenbaum [59] simply asserts the presence of primitives at varying levels of generality (though it provides a principled framework for selecting between these). Further, for many of these models, the assumed trajectory of learning is from concrete to abstract but this trajectory may not be followed by learners: abstract relations could be entertained as early as featural relations in some cases (or even earlier, through a “blessing of abstraction” [60]). Finally, these accounts again presuppose the same kinds of hypotheses to be available in non-human animals, which leads us back to asserting the presence of symbolic predicates in any animal that can pass an RMTS task at any level of performance. Fully graded accounts in which representations themselves are continuous and inclusive of featural, relational, and even contextual information avoid these issues [12] and might provide a more parsimonious account of the full pattern of findings.

Box 2: What would constitute a graded sense of *same* or *different*?

If you were to encode the two shapes below in memory, your representation probably wouldn’t include the precise number of rays. If each were encoded separately, without an initial direct comparison between them, it would be difficult to decide whether they were the same shape or different shapes. When stimuli are available for comparison, it may be possible to use visual or counting strategies to make decisions that avoid conceptual fuzziness, but tasks that require delayed responding or limit visual access to stimuli suggest that human encoding schemes exhibit fuzziness (leading to confusion errors) that is often reduced through the use of linguistic symbols (e.g., [52, 61-62]). We suggest that relational labels like “same” and “different” may play a similar role to linguistic categories like “seven,” “green,” or “above” in signaling and promoting fine-grained distinctions between fuzzy concepts, which might otherwise be confused, in which case we would fail to discriminate *same* from same-ish.

Figure 1

A. Two similar, but distinct, shapes.

B. A styled depiction of representational uncertainty about the shapes in A that could lead to confusion errors.

A fuzzy or gradient notion of *same* could be formalized in many ways. Following a tradition of investigating neural nets as instantiations of this idea [8], Geiger et al. [12] provide a set of network models that learn such graded functions from data. But other formalisms are possible as well. For example, within fuzzy set theory, we could assign the pair of shapes above some probability of belonging to the *same* set, say 90%. However, this type of fuzzy set formalization provides a poor fit to notions involving compositional semantics in natural language [63], while neural network models for logical semantics are gaining increasing currency (see e.g., [64]). Our goal here is not to adjudicate between these approaches, but to highlight a range of perspectives that may hold promise for formalizing fuzzy notions of *same* and *different*.

Box 3: Using explicit symbols along with graded representations

Some of the variation that we observe in human behaviors may be due to the mixing of explicit symbols and graded representations in a single task. When you open your kitchen cabinet and see identical drinking glasses, you have available both the perceptual input and the mental algorithm necessary to arrive at a representation of *same* (we remain agnostic about whether this symbol is explicitly linguistic). Yet meeting these preconditions alone does not guarantee that you *will* engage the symbolic representation. Children and adults who have learned a linguistic symbol for *same* and a corresponding algorithm for assessing sameness may still show varied success in relational reasoning. They may be unable to engage *same* computations due to limitations on working memory or other components of executive function; these limitations could prevent *same* computations when the items to be compared are complex or other demands on processing are high. Further, if stimuli are less salient they may not trigger explicit *same* reasoning, while if they are more familiar, highly practiced, and/or meaningful they may lead to more *same* processing (e.g., as in studies of infant “rule learning” [35, 65-66]). Even adults making use of an explicit symbol for *same* may vary in their execution as a result of context-dependent construals [67]. In many real-world contexts, featural identity (e.g., both yellow) or equivalence (the same make of car) stand in for exact identity, and reasoners must first determine the relevant sense of *same* as a prerequisite to any evaluation of the relation.

Acknowledgements

We are grateful to Christopher Potts, Atticus Geiger, and an anonymous reviewer for helpful feedback on this manuscript. This work was funded in part by awards from the McDonnell Foundation and the Center for the Study of Language and Information at Stanford University supporting A. Carstensen.

References

- [1] Gentner, D. (2003). Why we're so smart. In D. Gentner and S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp.195-235). Cambridge, MA: MIT Press.
- [2] Marcus, G.F. (2001). *The Algebraic Mind*. Cambridge, Mass.: MIT Press.
- [3] Giurfa, M., Zhang, S., Jenett, A., Menzel, R., & Srinivasan, M. V. (2001). The concepts of 'sameness' and 'difference' in an insect. *Nature*, 410(6831), 930-933.
- [4] Hochmann, J.R., Tuerk, A. S., Sanborn, S., Zhu, R., Long, R., Dempster, M., & Carey, S. (2017). Children's representation of abstract relations in relational/array match-to-sample tasks. *Cognitive Psychology*, 99, 17–43. doi:10.1016/j.cogpsych.2017.11.001.
- [5] Munakata, Y., McClelland, J. L., Johnson, M. H., & Siegler, R. S. (1997). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, 104(4), 686.
- **[6] Wasserman, E. A., Castro, L., & Fagot, J. (2017). Relational thinking in animals and humans: From percepts to concepts. In J. Call, G. M. Burghardt, I. M. Pepperberg, C. T. Snowdon, & T. Zentall (Eds.). *APA handbook of comparative psychology: Perception, learning, and cognition* (p. 359–384). American Psychological Association.

A review of relational reasoning research in both human and non-human animals that examines the evidence for continuity between perceptual and relational reasoning in non-human animals and over human development.

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

*[8] Alhama, R. G., & Zuidema, W. (2019). A review of computational models of basic rule learning: The neural-symbolic debate and beyond. *Psychonomic Bulletin & Review*, 26 (4), 1174–1194.

A review of computational work on equality learning that argues that neural network approaches should include cognitively plausible pre-wiring and pre-training regimes, which would enable symbolic generalization without symbolic learning processes.

[9] Santoro, A., Raposo, D., Barrett, D. G., Malinowski, M., Pascanu, R., Battaglia, P., & Lillicrap, T. (2017). A simple neural network module for relational reasoning. In *Advances in Neural Information Processing Systems* (pp. 4967-4976).

[10] Alhama, R. G., & Zuidema, W. (2018). Pre-Wiring and Pre-Training: What does a neural network need to learn truly general identity rules?. *Journal of Artificial Intelligence Research*, 61, 927-946.

[11] Weyde, T., & Kopparti, R. M. (2019). Modelling Identity Rules with Neural Networks. *Journal of Applied Logics*, 6(4), 745-769.

**[12] Geiger, A., Carstensen, A., Frank, M., & Potts, C. (2020). Relational reasoning and generalization using non-symbolic neural networks. *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*, Austin, TX: Cognitive Science Society.

Computational modeling work demonstrating that neural networks can solve a range of equality-based reasoning tasks, including 2-item RMTS, suggesting that non-symbolic learning processes can support essential aspects of symbolic reasoning.

- [13] Premack, D. (1983). The codes of man and beasts. *Behavioral and Brain Sciences*, 6(1), 125-136.
- [14] Premack, D. (1976). *Intelligence in ape and man*. Hillsdale, NJ: Erlbaum.
- [15] Thompson, R. K. R., & Oden, D. L. (1995). A profound disparity revisited: Perception and judgment of abstract identity relations by chimpanzees, human infants, and monkeys. *Behavioural Processes*, 35(1-3), 149–161. [https://doi.org/10.1016/0376-6357\(95\)00048-8](https://doi.org/10.1016/0376-6357(95)00048-8)
- [16] Pepperberg, I. M. (1987). Acquisition of the same/different concept by an African Grey parrot (*Psittacus erithacus*): Learning with respect to categories of color, shape, and material. *Animal Learning and Behavior*, 15, 423–432.
<http://dx.doi.org/10.3758/BF03205051>
- [17] Martinho, A., & Kacelnik, A. (2016). Ducklings imprint on the relational concept of “same or different”. *Science*, 353(6296), 286-288.
- [18] Birman, D. & Gardner, J. L. (2016) Parietal and prefrontal: categorical differences? *Nature Neuroscience* 19: 5-7.
- [19] Gillan, D. J., Premack, D., & Woodruff, G. (1981). Reasoning in the chimpanzee: I. Analogical reasoning. *Journal of Experimental Psychology: Animal Behavior Processes*, 7, 1–17. <http://dx.doi.org/10.1037/0097-7403.7.1.1>
- [20] Thompson, R. K. R., Oden, D. L., & Boysen, S. T. (1997). Language-naïve chimpanzees (*Pan troglodytes*) judge relations between relations in a conceptual matching-to-sample

- task. *Journal of Experimental Psychology: Animal Behavior Processes*, 23 , 31–43.
<http://dx.doi.org/10.1037/0097-7403.23.1.31>
- [21] Thompson, R. K., & Oden, D. L. (2000). Categorical perception and conceptual judgments by nonhuman primates: The paleological monkey and the analogical ape. *Cognitive Science*, 24(3), 363-396.
- [22] Fagot, J., Wasserman, E. A., & Young, M. E. (2001). Discriminating the relation between relations: The role of entropy in abstract conceptualization by baboons (*Papio papio*) and humans (*Homo sapiens*). *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 316–328. <http://dx.doi.org/10.1037/0097-7403.27.4.316>
- [23] Cook, R. G., & Wasserman, E. A. (2007). Learning and transfer of relational matching-to-sample by pigeons.
- [24] Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, 31(2), 109-130.
- [25] Wasserman, E. A., & Young, M. E. (2010). Same–different discrimination: The keel and backbone of thought and reasoning. *Journal of Experimental Psychology: Animal Behavior Processes*, 36(1), 3.
- [26] Smirnova, A., Zorina, Z., Obozova, T., & Wasserman, E. (2015). Crows spontaneously exhibit analogical reasoning. *Current Biology*, 25(2), 256-260.
- [27] Obozova, T., Smirnova, A., Zorina, Z., & Wasserman, E. (2015). Analogical reasoning in amazons. *Animal Cognition*, 18(6), 1363-1371.

- [28] Fagot, J., & Thompson, R. K. R. (2011). Generalized relational matching by guinea baboons (*Papio papio*) in two-by-two-item analogy problems. *Psychological Science*, 22 , 1304–1309. <http://dx.doi.org/10.1177/0956797611422916>
- [29] Castro, L., Kennedy, P. L., & Wasserman, E. A. (2010). Conditional same–different discrimination by pigeons: Acquisition and generalization to novel and few-item displays. *Journal of Experimental Psychology: Animal Behavior Processes*, 36 , 23–38. <http://dx.doi.org/10.1037/a0016326>
- [30] Walker, C. M., Bridgers, S., & Gopnik, A. (2016). The early emergence and puzzling decline of relational reasoning: Effects of knowledge and search on inferring abstract concepts. *Cognition*, 156, 30-40.
- [31] Son, J. Y., Smith, L. B., & Goldstone, R. L. (2011). Connecting instances to promote children’s relational reasoning. *Journal of Experimental Child Psychology*, 108(2), 260-277.
- [32] Jamrozik, A., & Gentner, D. (2020). Relational labeling unlocks inert knowledge. *Cognition*, 196, 104146.
- *[33] Carstensen, A., Zhang, J., Heyman, G.D., Fu, G., Lee, K., & Walker, C.M. (2019). Context shapes early diversity in abstract thought. *Proceedings of the National Academy of Sciences*, 116(28), 13891-13896.
- Evidence that the early trajectory of same-different reasoning varies across cultural contexts in children between 18 and 48 months, implying that developmental changes in same-different reasoning are more sensitive to context than previously thought.
- [34] Saffran, J. R., & Thiessen, E. D. (2003). Pattern induction by infant language learners. *Developmental Psychology*, 39(3), 484.

*[35] Rabagliati, H., Ferguson, B., & Lew-Williams, C. (2019). The profile of abstract rule learning in infancy: Meta-analytic and experimental evidence. *Developmental Science*, 22(1), e12704.

A meta-analysis of 63 experiments on equality-based rule learning in infants and a novel study which show that infant equality rule-learning is not specific to speech, but dependent on ecological familiarity of the stimuli.

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

[37] Addyman, C., & Mareschal, D. (2010). The perceptual origins of the abstract same/different concept in human infants. *Animal Cognition*, 13(6), 817-833.

[38] Anderson, E. M., Chang, Y. J., Hespos, S., & Gentner, D. (2018). Comparison within pairs promotes analogical abstraction in three-month-olds. *Cognition*, 176, 74-86.

[39] Hochmann, J.R., Mody, S., & Carey, S. (2016). Infants' representations of same and different in match-and non-match-to-sample. *Cognitive Psychology*, 86, 87-111.

[40] Walker, C. M., Rett, A., & Bonawitz, E. (2020). Design drives discovery in causal learning. *Psychological Science*. doi: 10.1177/0956797619898134

[41] Goddu, M., & Walker, C.M. (2018). Toddlers and Adults Simultaneously Track Multiple Hypotheses in a Causal Learning Task. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society*. Madison, WI: Cognitive Science Society.

[42] Christie, S., & Gentner, D. (2014). Language helps children succeed on a classic analogy task. *Cognitive Science*, 38, 383–397.

- [43] Du, C., Miyazaki Y., Cook, M., Papadopoulos, J., & Ho, Y. (2018). Relational Language Improves Preschool Children's Performance of Analogical Reasoning. *International Journal of Psychological Studies*, 10(2), 91-101.
- [44] Hoyos, C., Shao, R., & Gentner, D. (2016). The paradox of relational development: Could language learning be (temporarily) harmful? D. Grodner, D. Mirman, A. Papafragou, J. Trueswell, J. Novick, S. Arunachalam, S. Christie, & C. Norris (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- [45] Masuda, T., & Nisbett, R. E. (2001). Attending holistically versus analytically: Comparing the context sensitivity of Japanese and Americans. *J. Pers. Soc. Psychol.* 81, 922–934.
- [46] Kitayama, S., Duffy, S., Kawamura, T., & Larsen, J. T. (2003). Perceiving an object and its context in different cultures: A cultural look at new look. *Psychological Science*, 14(3), 201-206.
- [47] Imada, T., Carlson, S. M., & Itakura, S. (2013). East–West cultural differences in context-sensitivity are evident in early childhood. *Developmental Science*, 16(2), 198-208.
- [48] Waxman, S. R., Fu, X., Ferguson, B., Geraghty, K., Leddon, E., Liang, J., & Zhao, M. F. (2016). How early is infants' attention to objects and actions shaped by culture? New evidence from 24-month-olds raised in the US and China. *Frontiers in Psychology*, 7, 97.
- [49] Gibson, B. M., & Wasserman, E. A. (2004). Time-course of control by specific stimulus features and relational cues during same-different discrimination training. *Animal Learning & Behavior*, 32(2), 183-189.
- [50] Endress, A. D., Dehaene-Lambertz, G., & Mehler, J. (2007). Perceptual constraints and the learnability of simple grammars. *Cognition*, 105(3), 577-614.

- [51] Frank, M. C., Fedorenko, E., Lai, P., Saxe, R., & Gibson, E. (2012). Verbal interference suppresses exact numerical representation. *Cognitive psychology*, 64(1-2), 74-92.
- [52] Frank, M. C., Everett, D. L., Fedorenko, E., & Gibson, E. (2008). Number as a cognitive technology: Evidence from Pirahã language and cognition. *Cognition*, 108(3), 819-824.
- [53] Gallistel, C. R., Fairhurst, S., & Balsam, P. (2004). The learning curve: implications of a quantitative analysis. *Proceedings of the National Academy of Sciences*, 101(36), 13124-13131.
- [54] Daw, N., & Courville, A. (2008). The pigeon as particle filter. *Advances in neural information processing systems*, 20, 369-376.
- [55] Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.
- [56] Mix, K. S. (2009). How Spencer made number: First uses of the number words. *Journal of Experimental Child Psychology*, 102(4), 427-444.
- [57] Haryu, E., Imai, M., & Okada, H. (2011). Object similarity bootstraps young children to action-based verb extensions. *Child Development*, 82(2), 674-686.
- [58] Gentner, D. & Hoyos, C. (2017). Analogy & abstraction. *Topics in Cognitive Science*, 9(3), 800-818.
- [59] Frank, M. C., & Tenenbaum, J. B. (2011). Three ideal observer models for rule learning in simple languages. *Cognition*, 120(3), 360-371.
- [60] 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.

- [61] Cibelli, E., Xu, Y., Austerweil, J. L., Griffiths, T. L., & Regier, T. (2016). The Sapir-Whorf hypothesis and probabilistic inference: Evidence from the domain of color. *PloS one*, 11(7), e0158725.
- [62] Holmes, K. J., Moty, K., & Regier, T. (2017). Revisiting the role of language in spatial cognition: Categorical perception of spatial relations in English and Korean speakers. *Psychonomic Bulletin & Review*, 24(6), 2031-2036.
- [63] Kamp, H., & Partee, B. (1995). Prototype theory and compositionality. *Cognition*, 57(2), 129-191.
- [64] Potts, C. (2019). A case for deep learning in semantics: Response to Pater. *Language*, 95(1), e115-e124.
- [65] Marcus, G.F., Fernandes, K.J., & Johnson, S.P. (2007). Infant rule learning facilitated by speech. *Psychological Science*, 18(5), 387–391.
- [66] 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.
- [67] Kroger, J. K., Holyoak, K. J., & Hummel, J. E. (2004). Varieties of sameness: The impact of relational complexity on perceptual comparisons. *Cognitive Science*, 28(3), 335-358.