Available online at www.sciencedirect.com

ScienceDirect

Current Opinion in
Behavioral Sciences

Do graded representations support abstract thought?

Alexandra Carstensen and Michael C Frank

Relational reasoning requires the reasoner to go beyond her/his 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 toward a reconciliation of these contradictions by suggesting that same-different reasoning is supported by graded representations.

Address
Department of Psychology, Stanford University, United States

Corresponding author:
Carstensen, Alexandra (abcarstensen@stanford.edu)

Current Opinion in Behavioral Sciences 2020, **37**:90–97

This review comes from a themed issue on **Same-different conceptualization**

Edited by **Edward A Wasserman**, **Jean-Rémy Hochmann** and **Susan Carey**

https://doi.org/10.1016/j.cobeha.2020.10.009

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, that is, 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 toward 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 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

---

[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.

received far less attention in the literature (cf. Ref. [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 Ref. [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. Refs. [9–11,12••]). In their simplest form, such graded functions approximate 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 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).

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. Refs. [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. Ref. [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 and 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.

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 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

---

[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.

practice was a critical ingredient in solving the task [13]. But three other language-naive 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 naive 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. Ref. [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] Follow-up 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 this generalization decrement, 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 (Box 2, Figure 1).

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. Refs. [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 in infant same-different reasoning.[5] For example, Addyman and Mareschal [37] found that human infants' success in a

---

[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. Ref. [4] for an empirical case in favor of stronger interpretations of success with 2-item RMTS).

[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 versus 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.

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 would not 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. Refs. [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.

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. Ref. [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*.

Figure 1

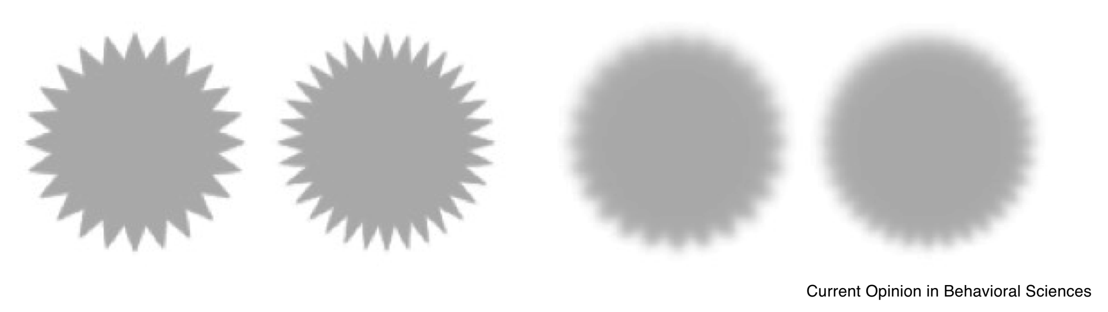

Current Opinion in Behavioral Sciences

(a) Two similar, but distinct, shapes. **(b)** A styled depiction of representational uncertainty about the shapes in (a) that could lead to confusion errors.

same-different task was modulated by the experimental setup (looking time with naturalistic picture pairs versus 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* versus *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.

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 and 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 uniquely[6] 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 [9–11,12••]. 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.

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.

Perhaps the clearest single case study for gradation comes from the wealth of same-different learning studies with pigeons (summarized in Ref. [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. 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

---

[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.

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-linguistic and non-linguistic experiences, leading to some of the incredibly fast and efficient generalization behavior shown by adults (e.g. Refs. [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] (Box 3).

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.

Conflict of interest statement

Nothing declared.

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 and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- • of special interest
- •• of outstanding interest

1. Gentner D: **Why we're so smart**. In *Language in Mind: Advances in the Study of Language and Thought.* Edited by Gentner D, Goldin-Meadow S. Cambridge, MA: MIT Press; 2003:195-235.

2. Marcus GF: *The Algebraic Mind*. Cambridge, Mass: MIT Press; 2001.

3. Giurfa M, Zhang S, Jenett A, Menzel R, Srinivasan MV: **The concepts of 'sameness' and 'difference' in an insect**. *Nature* 2001, **410**:930-933.

4. Hochmann JR, Tuerk AS, Sanborn S, Zhu R, Long R, Dempster M, Carey S: **Children's representation of abstract relations in relational/array match-to-sample tasks**. *Cogn Psychol* 2017, **99**:17-43 http://dx.doi.org/10.1016/j.cogpsych.2017.11.001.

5. Munakata Y, McClelland JL, Johnson MH, Siegler RS: **Rethinking infant knowledge: toward an adaptive process account of successes and failures in object permanence tasks**. *Psychol Rev* 1997, **104**:686.

6. •• Wasserman EA, Castro L, Fagot J: **Relational thinking in animals and humans: from percepts to concepts**. In *APA Handbook of Comparative Psychology: Perception, Learning, and Cognition.* Edited by Call J, Burghardt GM, Pepperberg IM, Snowdon CT, Zentall T. American Psychological Association; 2017:359-384.

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 GF, Vijayan S, Rao SB, Vishton PM: **Rule learning by seven-month-old infants**. *Science* 1999, **283**:77-80.

8. • Alhama RG, Zuidema W: **A review of computational models of basic rule learning: the neural-symbolic debate and beyond**. *Psychon Bull Rev* 2019, **26**: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 DG, Malinowski M, Pascanu R, Battaglia P, Lillicrap T: **A simple neural network module for relational reasoning**. *Advances in Neural Information Processing Systems*. 2017:4967-4976.

10. Alhama RG, Zuidema W: **Pre-wiring and pre-training: what does a neural network need to learn truly general identity rules?** *J Artif Intell Res* 2018, **61**:927-946.

11. Weyde T, Kopparti RM: **Modelling identity rules with neural networks**. *J Appl Logics* 2019, **6**:745-769.

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

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: **The codes of man and beasts**. *Behav Brain Sci* 1983, **6**:125-136.

14. Premack D: *Intelligence in Ape and Man*. Hillsdale, NJ: Erlbaum; 1976.

15. Thompson RKR, Oden DL: **A profound disparity revisited: perception and judgment of abstract identity relations by chimpanzees, human infants, and monkeys**. *Behav Process* 1995, **35**:149-161 http://dx.doi.org/10.1016/0376-6357(95)00048-8.

16. Pepperberg IM: **Acquisition of the same/different concept by an African Grey parrot (*Psittacus erithacus*): learning with respect to categories of color, shape, and material**. *Anim Learn Behav* 1987, **15**:423-432 http://dx.doi.org/10.3758/BF03205051.

17. Martinho A, Kacelnik A: **Ducklings imprint on the relational concept of "same or different"**. *Science* 2016, **353**:286-288.

18. Birman D, Gardner JL: **Parietal and prefrontal: categorical differences?** *Nat Neurosci* 2016, **19**:5-7.

19. Gillan DJ, Premack D, Woodruff G: **Reasoning in the chimpanzee: I. Analogical reasoning**. *J Exp Psychol Anim Behav Process* 1981, **7**:1-17 http://dx.doi.org/10.1037/0097-7403.7.1.1.

20. Thompson RKR, Oden DL, Boysen ST: **Language-naive chimpanzees (*Pan troglodytes*) judge relations between relations in a conceptual matching-to-sample task**. *J Exp Psychol Anim Behav Process* 1997, **23**:31-43 http://dx.doi.org/10.1037/0097-7403.23.1.31.

21. Thompson RK, Oden DL: **Categorical perception and conceptual judgments by nonhuman primates: the paleological monkey and the analogical ape**. *Cogn Sci* 2000, **24**:363-396.

22. Fagot J, Wasserman EA, Young ME: **Discriminating the relation between relations: the role of entropy in abstract conceptualization by baboons (*Papio papio*) and humans (*Homo sapiens*)**. *J Exp Psychol Anim Behav Process* 2001, **27**:316-328 http://dx.doi.org/10.1037/0097-7403.27.4.316.

23. Cook RG, Wasserman EA: **Learning and transfer of relational matching-to-sample by pigeons**. *Psychon Bull Rev* 2007, **14**:1107-1114.

24. Penn DC, Holyoak KJ, Povinelli DJ: **Darwin's mistake: explaining the discontinuity between human and nonhuman minds**. *Behav Brain Sci* 2008, **31**:109-130.

25. Wasserman EA, Young ME: **Same–different discrimination: the keel and backbone of thought and reasoning**. *J Exp Psychol Anim Behav Process* 2010, **36**:3.

26. Smirnova A, Zorina Z, Obozova T, Wasserman E: **Crows spontaneously exhibit analogical reasoning**. *Curr Biol* 2015, **25**:256-260.

27. Obozova T, Smirnova A, Zorina Z, Wasserman E: **Analogical reasoning in amazons**. *Anim Cogn* 2015, **18**:1363-1371.

28. Fagot J, Thompson RKR: **Generalized relational matching by guinea baboons (*Papio papio*) in two-by-two-item analogy problems**. *Psychol Sci* 2011, **22**:1304-1309 http://dx.doi.org/10.1177/0956797611422916.

29. Castro L, Kennedy PL, Wasserman EA: **Conditional same–different discrimination by pigeons: acquisition and generalization to novel and few-item displays**. *J Exp Psychol Anim Behav Process* 2010, **36**:23-38 http://dx.doi.org/10.1037/a0016326.

30. Walker CM, Bridgers S, Gopnik A: **The early emergence and puzzling decline of relational reasoning: effects of knowledge and search on inferring abstract concepts**. *Cognition* 2016, **156**:30-40.

31. Son JY, Smith LB, Goldstone RL: **Connecting instances to promote children's relational reasoning**. *J Exp Child Psychol* 2011, **108**:260-277.

32. Jamrozik A, Gentner D: **Relational labeling unlocks inert knowledge**. *Cognition* 2020, **196**:104146.

33. • Carstensen A, Zhang J, Heyman GD, Fu G, Lee K, Walker CM: **Context shapes early diversity in abstract thought**. *Proc Natl Acad Sci U S A* 2019, **116**: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 JR, Thiessen ED: **Pattern induction by infant language learners**. *Dev Psychol* 2003, **39**:484.

35. • Rabagliati H, Ferguson B, Lew-Williams C: **The profile of abstract rule learning in infancy: meta-analytic and experimental evidence**. *Dev Sci* 2019, **22**: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 AL, Hespos SJ, Gentner D: **Prelinguistic relational concepts: investigating analogical processing in infants**. *Child Dev* 2015, **86**:1386-1405.

37. Addyman C, Mareschal D: **The perceptual origins of the abstract same/different concept in human infants**. *Anim Cogn* 2010, **13**:817-833.

38. Anderson EM, Chang YJ, Hespos S, Gentner D: **Comparison within pairs promotes analogical abstraction in three-month-olds**. *Cognition* 2018, **176**:74-86.

39. Hochmann JR, Mody S, Carey S: **Infants' representations of same and different in match-and non-match-to-sample**. *Cogn Psychol* 2016, **86**:87-111.

40. Walker CM, Rett A, Bonawitz E: **Design drives discovery in causal learning**. *Psychol Sci* 2020, **31**:129-138 http://dx.doi.org/10.1177/0956797619898134.

41. Goddu M, Walker CM: **Toddlers and adults simultaneously track multiple hypotheses in a causal learning task**. In In *Proceedings of the 40th Annual Conference of the Cognitive Science Society.* Edited by Miyake N, Peebles D, Cooper RP. *Proceedings of the 40th Annual Conference of the Cognitive Science Society* Madison, WI: Cognitive Science Society; 2018.

42. Christie S, Gentner D: **Language helps children succeed on a classic analogy task**. *Cogn Sci* 2014, **38**:383-397.

43. Du C, Miyazaki Y, Cook M, Papadopoulos J, Ho Y: **Relational language improves preschool children's performance of analogical reasoning**. *Int J Psychol Stud* 2018, **10**:91-101.

44. Hoyos C, Shao R, Gentner D: **The paradox of relational development: could language learning be (temporarily) harmful?** In *Proceedings of the 38th Annual Conference of the Cognitive Science Society.* Edited by Grodner D, Mirman D, Papafragou A, Trueswell J, Novick J, Arunachalam S, Christie S, Norris C. *Proceedings of the 38th Annual Conference of the Cognitive Science Society* Austin, TX: Cognitive Science Society; 2016.

45. Masuda T, Nisbett RE: **Attending holistically versus analytically: comparing the context sensitivity of Japanese and Americans**. *J Pers Soc Psychol* 2001, **81**:922-934.

46. Kitayama S, Duffy S, Kawamura T, Larsen JT: **Perceiving an object and its context in different cultures: a cultural look at new look**. *Psychol Sci* 2003, **14**:201-206.

47. Imada T, Carlson SM, Itakura S: **East–West cultural differences in context-sensitivity are evident in early childhood**. *Dev Sci* 2013, **16**:198-208.

48. Waxman SR, Fu X, Ferguson B, Geraghty K, Leddon E, Liang J, Zhao MF: **How early is infants' attention to objects and actions shaped by culture? New evidence from 24-month-olds raised in the US and China**. *Front Psychol* 2016, **7**:97.

49. Gibson BM, Wasserman EA: **Time-course of control by specific stimulus features and relational cues during same-different discrimination training**. *Anim Learn Behav* 2004, **32**:183-189.

50. Endress AD, Dehaene-Lambertz G, Mehler J: **Perceptual constraints and the learnability of simple grammars**. *Cognition* 2007, **105**:577-614.

51. Frank MC, Fedorenko E, Lai P, Saxe R, Gibson E: **Verbal interference suppresses exact numerical representation**. *Cogn Psychol* 2012, **64**:74-92.

52. Frank MC, Everett DL, Fedorenko E, Gibson E: **Number as a cognitive technology: evidence from Pirahã language and cognition**. *Cognition* 2008, **108**:819-824.

53. Gallistel CR, Fairhurst S, Balsam P: **The learning curve: implications of a quantitative analysis**. *Proc Natl Acad Sci U S A* 2004, **101**:13124-13131.

54. Daw N, Courville A: **The pigeon as particle filter**. *Adv Neural Inform Process Syst* 2008, **20**:369-376.

55. Gick ML, Holyoak KJ: **Schema induction and analogical transfer**. *Cogn Psychol* 1983, **15**:1-38.

56. Mix KS: **How spencer made number: first uses of the number words**. *J Exp Child Psychol* 2009, **102**:427-444.

57. Haryu E, Imai M, Okada H: **Object similarity bootstraps young children to action-based verb extensions**. *Child Dev* 2011, **82**:674-686.

58. Gentner D, Hoyos C: **Analogy & abstraction**. *Top Cogn Sci* 2017, **9**:800-818.

59. Frank MC, Tenenbaum JB: **Three ideal observer models for rule learning in simple languages**. *Cognition* 2011, **120**:360-371.

60. Tenenbaum JB, Kemp C, Griffiths TL, Goodman ND: **How to grow a mind: statistics, structure, and abstraction**. *Science* 2011, **331**:1279-1285.

61. Cibelli E, Xu Y, Austerweil JL, Griffiths TL, Regier T: **The Sapir-Whorf hypothesis and probabilistic inference: evidence from the domain of color**. *PLoS One* 2016, **11**:e0158725.

62. Holmes KJ, Moty K, Regier T: **Revisiting the role of language in spatial cognition: categorical perception of spatial relations in English and Korean speakers**. *Psychon Bull Rev* 2017, **24**:2031-2036.

63. Kamp H, Partee B: **Prototype theory and compositionality**. *Cognition* 1995, **57**:129-191.

64. Potts C: **A case for deep learning in semantics: response to pater**. *Language* 2019, **95**:e115-e124.

65. Marcus GF, Fernandes KJ, Johnson SP: **Infant rule learning facilitated by speech**. *Psychol Sci* 2007, **18**:387-391.

66. Saffran JR, Pollak SD, Seibel RL, Shkolnik A: **Dog is a dog is a dog: infant rule learning is not specific to language**. *Cognition* 2007, **105**:669-680.

67. Kroger JK, Holyoak KJ, Hummel JE: **Varieties of sameness: the impact of relational complexity on perceptual comparisons**. *Cogn Sci* 2004, **28**:335-358.