



A connectionist account of the relational shift and context sensitivity in the development of generalisation

Paul H. Thibodeau^a, Aviva Blonder^{a,b} and Stephen J. Flusberg^c

^aDepartment of Psychology, Oberlin College, Oberlin, OH, USA; ^bUniversity of California, Davis, USA;

^cDepartment of Psychology, Purchase College, SUNY, Purchase, NY, USA

ABSTRACT

Similarity-based generalisation is fundamental to human cognition, and the ability to draw analogies based on relational similarities between superficially different domains is crucial for reasoning and inference. Learning to base generalisation on shared relations rather than (or in the face of) shared perceptual features has been identified as an important developmental milestone. However, recent research has highlighted the context-sensitivity of generalisation: children and adults use perceptual similarity to make inferences in some cases and relational similarity in others, a finding that suggests people track the predictive validity of different types of inferences. Here we demonstrate that this pattern of behaviour naturally emerges over the course of development in a domain-general statistical learning model that employs distributed, sub-symbolic representations. We suggest that this model offers a parsimonious account of the development of context-sensitive, similarity-based generalisation and may provide several advantages over other popular structured or symbolic approaches to modelling relational inference.

ARTICLE HISTORY

Received 26 August 2019

Accepted 8 February 2020

KEYWORDS

Analogy; similarity; relational shift; distributed connectionist model; generalisation; statistical learning

Introduction

Is a lemon more similar to a small yellow balloon or a green grape? The answer, it turns out, is not so straightforward. All three objects are small and round(ish), but the lemon and balloon are somewhat larger than the grape and both of them are yellow. On the other hand, the lemon and grape are filled with juice, grow on trees, and belong to the same basic category (fruit), while the balloon is man-made and filled with air. Your response, therefore, may depend on what type of similarity (you believe) the questioner has in mind; the lemon *looks* more similar to the yellow balloon but is *structurally* (and functionally) more similar to the grape.

Without any additional information, most adults would probably say that the lemon is more similar to the grape (Deng & Sloutsky, 2016; Gentner, 1988; Goswami & Brown, 1990). The shared taxonomic and structural elements of the lemon and grape trump the

CONTACT Paul H. Thibodeau pthibode@oberlin.edu Department of Psychology, Oberlin College, 120 W. Lorain St., Oberlin, OH 44074 USA

An earlier version of this work was presented in the proceedings of the 34th annual meeting of the Cognitive Science Society.

superficial similarity of the lemon and balloon. However, this *relational match* requires relatively sophisticated knowledge of lemons and grapes; without it, the lemon will seem more similar to the balloon: a great deal of empirical research has found that young children typically base similarity judgments on perceptual features before they have the relevant domain knowledge to make relational matches (Badger & Shapiro, 2012; Gentner & Rattermann, 1998; Hayes & Thompson, 2007; Sloutsky, Deng, Fisher, & Kloos, 2015). In other words, until young children gain sufficient knowledge of fruit, they are likely to say that a lemon is more similar to a yellow balloon than a grape. This developmental change in similarity matching—from an early reliance on surface-level, perceptual features to a later reliance on structural or relational properties—is known as the perceptual-to-relational shift (Gentner, 1988; Goswami, 1996; Piaget, 1952; Rattermann & Gentner, 1998).

Several properties of relational knowledge differentiate relational reasoning from non-analytic processes (e.g. simple association; Halford, Wilson, & Phillips, 2010). Relational knowledge is *structured* (if A is *larger* than B and B is *larger* than C, then A must be *larger* than C), *compositional* (A and B retain their identity in the compound representation, A is *larger* than B), and *systematic* (understanding John *loves* Mary implies the capacity to understand Mary *loves* John). Thus, relational knowledge provides a foundation for complex cognitive processes that are a hallmark of human intelligence (Gentner, 1983; Penn, Holyoak, & Povinelli, 2008).

Computational models have been instrumental in helping us understand the mechanistic underpinnings of relational reasoning (e.g. Falkenhainer, Forbus, & Gentner, 1989; Hummel & Holyoak, 1997), and the emergence of relational reasoning in children (Doumas, Hummel, & Sandhofer, 2008; Gentner, Rattermann, Markman, & Kotovsky, 1995; Leech, Mareschal, & Cooper, 2008; Lu, Wu, & Holyoak, 2019; Morrison, Doumas, & Richland, 2011; Rogers & McClelland, 2004; Thibodeau, Flusberg, Glick, & Sternberg, 2013). Notably, proponents of two modelling approaches that have been at the forefront of the field (*Structure Mapping Engine* – SME – proposed by Falkenhainer et al., 1989; and *Learning and Inference with Schemas and Analogies* – LISA – proposed by Hummel & Holyoak, 1997) have offered somewhat different accounts of the development of relational reasoning and the developmental trajectory of similarity-based generalisation.

Gentner et al. (1995) used SME to show how conceptual change and knowledge accretion could give rise to the relational shift. On this account, relational reasoning emerges as domain-specific knowledge increases (Gentner, 1988; Gentner & Rattermann, 1991; but see Goswami, 1995). In SME, concepts are coded in a predicate calculus that represents both objects and their relations in a structured, symbolic fashion. Knowledge accretion is achieved in the model by manually re-coding representations (and not, e.g. through experiential learning). While this model can accurately capture the perceptual-to-relational shift in this fashion (i.e. by using “object-centered” representations to model the performance of younger children and “relation-centered” representations to model the performance of older children and adults), it leaves open the question of how conceptual re-representation emerges as people acquire domain knowledge through everyday experience (for an extended discussion of related issues see Thibodeau et al., 2013).

Morrison et al. (2011) used LISA to show how the development of inhibitory control mechanisms could support a shift in attention from perceptual to relational structure during generalisation. On this alternative account, the development of flexible cognitive control resources is crucial for being able to inhibit the allure of a superficial perceptual

match. Importantly, and in contrast to SME, the basic principles of LISA have been extended in an attempt to explain how explicitly structured conceptual representations might be learned from experience (*Discovery of Relations by Analogy – DORA* – proposed by Doumas et al., 2008; see also Doumas, Morrison, & Richland, 2018).

There are clear advantages to both of these modelling frameworks, especially since SME and LISA have been used to simulate such a wide range of findings relating to knowledge representation and reasoning (Gentner & Forbus, 2011; Hummel & Holyoak, 2005). Using these models to explain the developmental trajectory of relational reasoning, therefore, represents a parsimonious extension of each approach that helps explain several key empirical findings.

However, the reliance on formal representational structure in these models may also represent a limitation. Computational models that lack formally structured representations (e.g. Leech et al., 2008; Rogers & McClelland, 2004; Thibodeau et al., 2013) may be better suited to explain the context-sensitivity of certain similarity judgments. For instance, recent findings call into question whether similarity-based generalisation follows a universal, across-the-board, perceptual-to-relational shift (Bulloch & Opfer, 2009; Opfer & Bulloch, 2007; Tarlowski, 2018). According to the *predictive validity* view, children do not *necessarily* proceed from generalisation by perceptual features to generalisation by relational structure. Instead, they generalise flexibly over different types of similarity depending on the context of their judgment. In certain domains, children (and adults) will have *learned* that inferences based on relational similarity are more reliably predictive of success, while in other domains inferences based on perceptual similarity may actually be more successful.

Data supporting the predictive validity view come from studies in which children and adults are asked to make inferences about a novel object in different contexts (Bulloch & Opfer, 2009; Opfer & Bulloch, 2007). Consider the triad of insects in Figure 1. In each of the three insect triplets, there are two adults (AA, BB, and TT) and one juvenile (*a*, *b*, and *t*). The triads were designed such that the insects on the top row (the “samples”: AA, *a*; BB, *b*) represent potential matches for the insects on the bottom (the “target”: TT, *t*). In every case, the target juvenile looked similar to the juvenile from one of the samples (in this case both *b* and *t* are light whereas *a* is dark) and the target adults looked similar to the adults in the other sample (in this case both AA and TT are light whereas BB is dark).

Bulloch and Opfer (2009) designed two different conditions to examine whether they could influence how people would generalise about the target juvenile: one in which the relational information was relevant (the juvenile is the *offspring* of the co-occurring adults, and so should they should share similar properties) and another in which the relational information was irrelevant (the juvenile is the *prey* of the co-occurring adults, so they need not have anything in particular in common). Then they had participants make inferences about the target juvenile, asking about category membership (is *t* the same kind as *a* or *b*?), an unobservable property (does *t* have “gogli” inside its blood similar to *a* or *b*?), and future appearance (will *t* look like *a* or *b* in the future?).

According to the *predictive validity* perspective, in the condition where the relation is relevant (i.e. when the participant is told that the juveniles are the *offspring* of the co-occurring adults), participants should choose the sample in which the adults look like the target adults (i.e. AA). That is, they should make an inference based on *relational* similarity. In the context where the relation is irrelevant (i.e. when the participant is told that the juveniles are

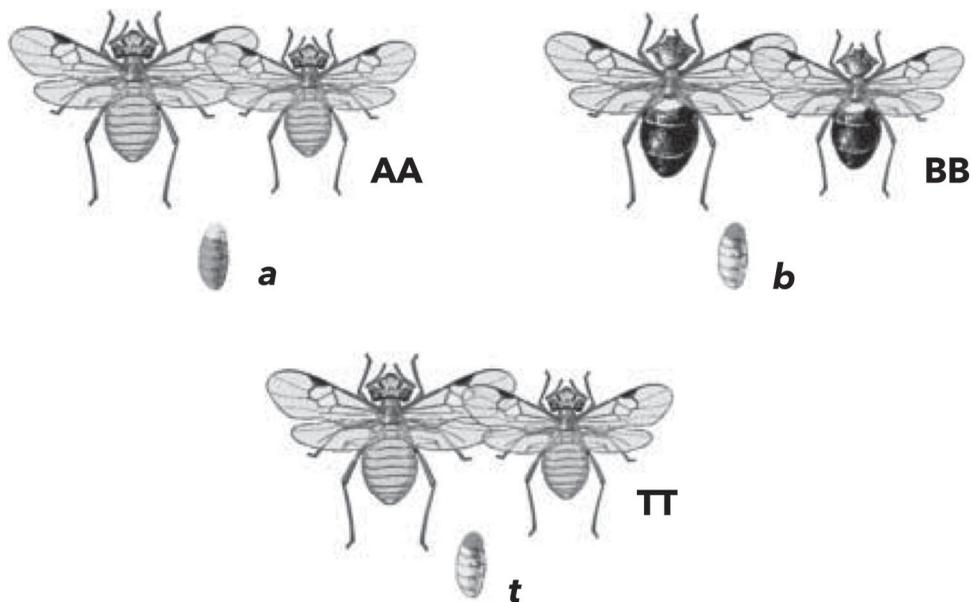


Figure 1. An example trial from Bulloch and Opfer (2009). The target juvenile (*t*) is perceptually more similar to (*b*) but is sometimes presented in a relational context that makes it more similar to (*a*).

the *prey* of the co-occurring adults), participants should choose the sample in which the juvenile looks like the target juvenile (i.e. *b*). That is, they should make an inference based on the *perceptual* similarity of the juveniles.

As expected, Bulloch and Opfer (2009) found that adults based their inferences about the target juvenile on perceptual properties of the juveniles in the prey context and relational properties (i.e. the similarity of the adults) in the offspring context. Patterns of results from three-, four-, and five-year-old children looked increasingly like those of the adults, supporting the view that there is not a universal trend from generalising by perceptual features to generalising by relational structure. Instead, these findings suggest that children and adults flexibly generalise using features or relations when contextually appropriate, based on their prior knowledge.¹

The present study

While the data provided by Bulloch and Opfer (2009) complicate the traditional picture of the emergence of relational reasoning over the course of development, we present a series of neural network simulations that spontaneously capture these findings and help explain the development of context-sensitive, similarity-based generalisation. The model architecture and simulated environment build on previous work that has explored the capacity of certain connectionist networks to capture and explain the development of semantic knowledge (Rogers & McClelland, 2004) and relational reasoning (e.g. Flusberg, Thibodeau, Sternberg, & Glick, 2010; Kollias & McClelland, 2013; Lampinen, Hsu, & McClelland, 2017; Leech et al., 2008; Thibodeau et al., 2013). This research has shown how and why higher-level cognitive abilities like analogical reasoning could spontaneously emerge over

the course of development based on domain-general principles of statistical learning and distributed representation.

The present simulations advance this work by focussing specifically on the relational shift and the mechanisms that support context-sensitive inferences. The modelling approach addresses limitations of classical structured and symbolic models like SME while retaining important insights from the empirical literature (e.g. the causal role that language seems to play in driving the development of relational reasoning; see Flusberg et al., 2010; Gentner & Rattermann, 1991; Thibodeau et al., 2013). In particular, our model is naturally context-sensitive (a well-known strength of connectionist networks; Flusberg & McClelland, 2014; Rogers & McClelland, 2004) and embodies the key principles underlying the predictive validity account of similarity-based reasoning.

Simulation

Methods

The environment and structure of our model were designed to replicate some of the essential features of Bulloch and Opfer's (2009) study (see Figure 2 for network architecture and simulation parameters). As input, the model takes a juvenile insect, presented as distributed patterns over a 15-unit array (of 0s and 1s), and a relational context, presented symbolically with a localist representation in a 5-unit array. The patterns that represent the juveniles were designed to operationalise Bulloch and Opfer's (2009) manipulation of perceptual similarity. The 15-unit array allowed for the creation of six "training juveniles" that were equally different from one another, with slightly negative pairwise correlations ($r = -0.2$), and a "test juvenile" that was perceptually similar to one pair of "training juveniles" ($r = 0.4$) but not the others ($r = -0.2$; see Figure 3). The symbolic representation of the relational

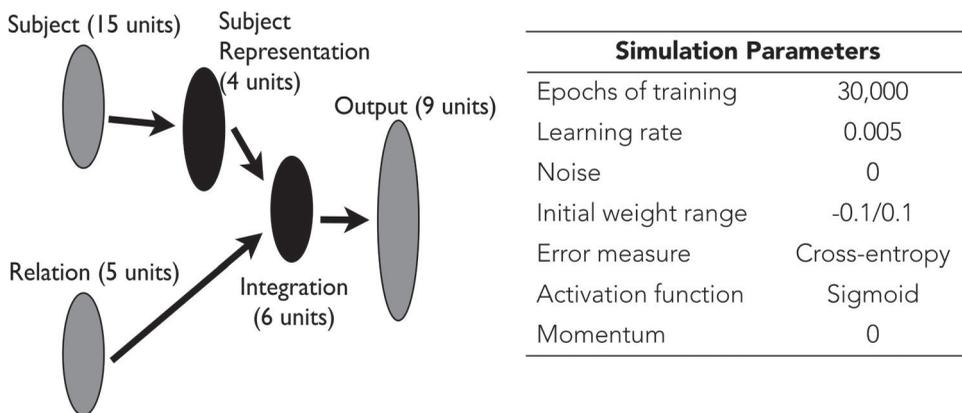


Figure 2. The network architecture and simulation parameters for the feedforward connectionist model, which is an adaptation of the Rumelhart network (Rumelhart, 1990). As input the model takes a distributed representation of a juvenile insect (*Subject*) in a relational context (*Relation*). These inputs feed forward to two hidden layers: in one (*Subject Representation*) the model learns an internal representation of the juvenile insects; in the other (*Integration*) the model learns to combine the two streams of input. The pool of *Output* units in the model include (a) three adult insects who the juvenile could be "born to," "eaten by," and "look like"; (b) three "bug types"; and (c) three "bug properties."

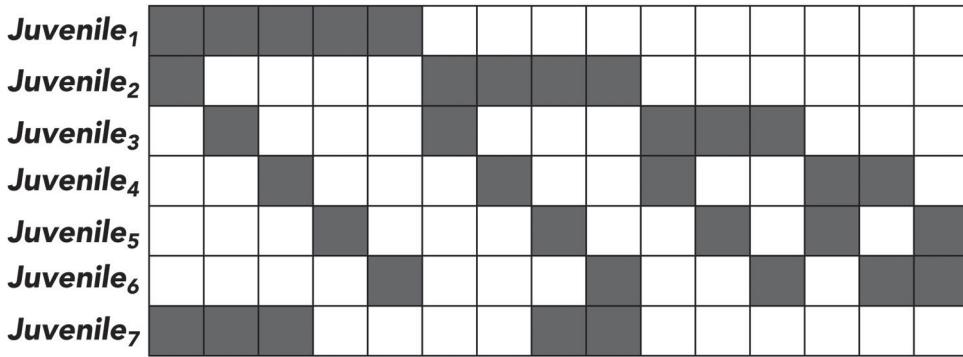


Figure 3. Visualisation of the input representation of the juveniles. Pairwise “perceptual” similarity (overlap of distributed representation) of the six training juveniles is equivalent ($r = -.2$). Test-juvenile₇ is “perceptually” similar to training-juvenile₁ and -juvenile₂ ($r = .4$) but equally dissimilar to the others (juveniles 3-6; $r = -.2$).

contexts corresponds to the linguistic information that the children were given about the juveniles or to the questions that the children were asked about the juveniles (e.g. *born to, eaten by, will look like, is, has*) by Bulloch and Opfer (2009).

As output, the model learns to complete the inputs with the appropriate adult, category, or property, which are also represented symbolically, consistent with the verbal output that the children provided in the empirical study. Three of the nine output units represent adult insects and inform who the juveniles are *born to, eaten by, and will look like*. The other six output units represent category membership and internal properties of the juveniles (3 each; see Table 1).

During training, the model learns about six juveniles in each of the five relational contexts. That is, the model learns that a given juvenile is *born to* a pair of ($adults_1, adults_2$, or $adults_3$) adults, is *eaten by* a pair of (the 3 possible) adults, *will look like* a pair of (the 3 possible) adults, *is* one of three types of bugs ($type_1, type_2$, or $type_3$), and *has* one of three types of specific properties ($property_1, property_2$, or $property_3$). It learns these relationships through

Table 1. Training and Test Patterns.

Pattern Type	Subject	Relation				
	Juvenile	<i>Born to</i> (adults)	<i>Eaten by</i> (adults)	<i>Looks like</i> (adults)	<i>Is</i> (bug type)	<i>Has</i> (property)
Train	1	1	2	1	1	1
Train	2	1	3	1	1	1
Train	3	2	1	2	2	2
Train	4	2	3	2	2	2
Train	5	3	1	3	3	3
Train	6	3	2	3	3	3
Test-relational	7	2		2	2	2
Test-perceptual	7		2	1	1	1

The top six rows represent training patterns and the bottom two represent test patterns. In training, the network learns about six juvenile bugs in each of five relational contexts for a total of 30 training patterns. At test, the model is given partial information about a novel juvenile (either who the juvenile was *born to*, for the “relational” problems or *eaten by*, for the “perceptual” problems) and is asked to make inferences about the future appearance, category membership, and internal properties of that juvenile (predicted inferences shown for *test* patterns; the model was not given feedback on the test juvenile’s future appearance, category membership, or internal properties).

experience: Initial errors are corrected through a supervised learning algorithm (backpropagation of error; Hinton, 1986). As connection weights between layers are adjusted, the model produces the “correct” output for the given input (subject-relation pairing). As a result, the model’s representations of the juveniles in the two hidden layers (Subject Representation and Integration) come to reflect the structure of the environment (e.g. the relationships between the juveniles and the adults). One of the key features of the model is that the size of hidden layers is smaller than the size of input and output layers, thereby forcing the model to compress what it knows about various juveniles in these overlapping, distributed representations (see Rogers & McClelland, 2004).

Of note, the model behaves similarly if the Subject Representation hidden layer is removed and the two streams of input flow directly to the Integration hidden layer (as in a more straightforward 3-layer feedforward architecture). Including the Subject Representation hidden layer, however, provides additional information about what the model knows about the juvenile bugs (see, e.g. Flusberg et al., 2010; Rogers & McClelland, 2004, 2008; Thibodeau et al., 2013). Whereas the Integration layer reflects what the model knows about a particular juvenile bug in a particular relational context, the Subject Representation layer reflects the model’s context-independent knowledge of the juvenile bugs.

Thirty training patterns were created: one pattern for each of the six “training juveniles” in each of the five *relational* contexts. The model was trained for 30,000 epochs with the 30 training patterns. Although this amount of training may seem substantial, it is important to note that the model starts with absolutely no prior experience. People, on the other hand, almost always have relevant prior knowledge to build on. Recent work has shown that when similar models are given relevant experience to build on, they learn much more quickly (Thibodeau et al., 2013).

Importantly, there is *coherent covariation* (Rogers & McClelland, 2004) between the *born to*, *will look like*, *is*, and *has* relations. Juveniles *will look like*, *belong to* the same category, and *have* the same property as the adults that they are *born to*. In contrast, knowing that a given juvenile is *eaten by* a particular pair of adults does not license inferences about future appearance, category membership, or internal properties.

To test the network’s ability to generalise, it is given partial information about a novel juvenile after it has learned about the six training juveniles. The pattern that represents this “test juvenile” was designed to be perceptually similar to one pair of juveniles that the network learned about in training and relationally similar to another. Perceptual similarity is operationalised as overlap in the distributed input representations ($r = 0.4$ between the novel juvenile and the two perceptually similar juveniles and $r = -0.2$ between each of the other juveniles; see Figure 3). That is, *juvenile₇* was perceptually similar to *juvenile₁* and *juvenile₂* (i.e. in terms of its distributed representation) but relationally similar to *juvenile₃* and *juvenile₄* in the sense that it might be *born to* the same adults as *juvenile₃* and *juvenile₄* (see the bottom rows of Table 1).

We presented the network with one of two kinds of inference problems after the training phase. In one (denoted as “test-*relational*” Table 1), the network was given the novel juvenile and information about who that juvenile was *born to*. In the other (denoted as “test-*perceptual*” in Table 1), the network was given the novel juvenile and information about who that juvenile was *eaten by*. In neither case was the network told what the novel juvenile *will look like*, *is*, or *has*. These were inferences that the network was asked to make. In other words, the model was given feedback on who the “test juvenile” was *born to* or *eaten*

by in test phase, but not who the “test juvenile” *looks like*, what type of bug the juvenile *is*, or what properties it *has*.

We presented the novel information (a single pattern) to the network for 1,000 epochs and monitored the trajectory of its inferences. Simulations were run ten times in each condition, initialised with different random starting weights to ensure that results were not the product of an idiosyncratic network initialisation and to allow for statistical tests.

Our prediction was that the network would initially make inferences about the novel juvenile that were consistent with the perceptually similar juveniles (i.e. that the network would infer *juvenile*₇ will look like, is of the same type as, and has the same properties as *juvenile*₁ and *juvenile*₂), since the distributed patterns representing these juveniles had significant overlap. However, we expected that the network would change what it thought about the novel juvenile in the *born to* condition (i.e. to infer that *juvenile*₇ is actually more similar to *juvenile*₃ and *juvenile*₄ because it is also born to *adults*₂); we expected no such change in the *eaten by* condition. In other words, we expected the network to behave flexibly, learning to use the relational information when it was predictive (based on its own prior experiences during training) and to ignore it when it was not.

Results

Before looking at the inferences the model made to the test patterns, we investigated how well the model learned the training patterns. Averaging across the 10 simulations, we found that cross-entropy error at the end of the training phase was low, $M = 16.1$, $SD = 9.7$, compared to the cross-entropy error at the beginning of the training phase, $M = 262.0$, $SD = 3.9$, $t(9) = 76.23$, $p < .001$. This indicates that the model learned to respond accurately to the training patterns during the training phase.

Then we turned to how the model behaved in the test phase. As predicted, the network initially made perceptual matches when presented with the novel juvenile. Learning who the novel juvenile was *born to*, however, led to a shift in the inference patterns of the model, consistent with a perceptual-to-relational shift. Such a shift did not occur when the model learned who the novel juvenile was *eaten by*, since there was no coherent covariation between *eaten by* and the inferential relational contexts (see Figure 4).

To statistically analyze the inferential tendencies of the model, we conducted three repeated measures ANOVAs. The first contrasted pre- and post-learning in the offspring context and found a main effect of perceptual inferences, $F(1,35) = 12.61$, $p < .01$, and an interaction between learning and inference type, $F(1,35) = 74.54$, $p < .001$. Before learning (epoch 0 of test), the model was strongly biased toward making perceptual inferences. After learning (epoch 1,000 of test), the network showed a dramatic shift towards relational inferences (see the first and third pairs of bars in Figure 5). That is, the model initially treated the novel juvenile like the previously learned, perceptually similar juveniles. But this changed when it was told that the novel juvenile was *born to* a different set of parents. Over time, it automatically re-conceptualized this juvenile to make inferences that were consistent with the juveniles that were *born to* the same adults.

The second ANOVA contrasted pre- and post-learning in the prey context and found a main effect of perceptual inferences, $F(1,35) = 60.00$, $p < .001$ and a relatively smaller but reliable interaction between learning and inference type, $F(1,35) = 7.15$, $p < .05$. As in the offspring condition, the model first made perceptual inferences. Unlike the offspring

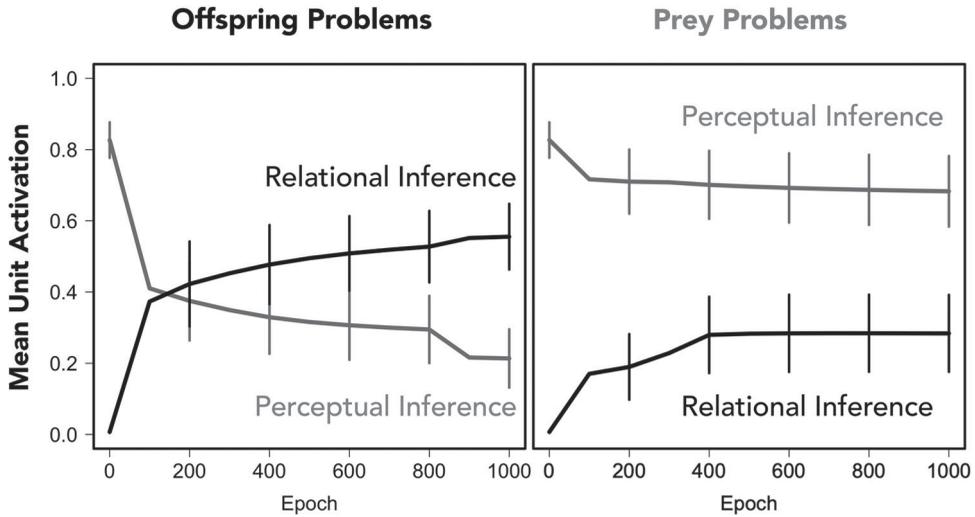


Figure 4. Mean activation of units that represent the “relational” and “perceptual” inference by context. Error bars denote the standard errors of the means. The left panel shows results from the offspring problems and shows a relational shift: Initially the network makes perceptual inferences and later makes relational inferences. The right panel shows results from the prey problems and does not show a relational shift: It makes the same initial inferences as in the offspring condition, which do not change qualitatively over time.

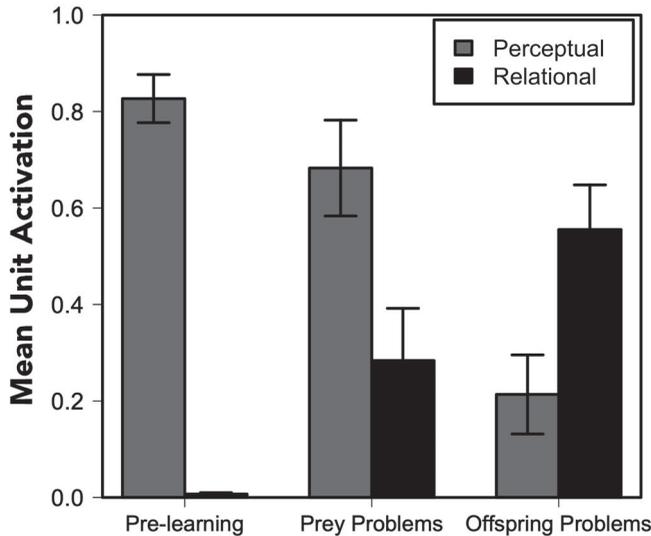


Figure 5. Mean activations of units that reflect “perceptual” and “relational” inferences before learning (left), after learning for the prey problems (middle), and after learning for the offspring problems (right). Error bars denote standard errors of the means.

condition, we did not see a crossover after learning, although it did become slightly more likely to make a relational inference (see the first two pairs of bars in Figure 5).

Finally, the third ANOVA contrasted the post-learning inferences across the two conditions and found a significant interaction, $F(1,35) = 14.50, p < .001$. Whereas the network

made more perceptual matches in the prey condition, it made more relational matches in the offspring condition (see the second and third pairs of bars in Figure 5).

After analyzing the model's inferences, we checked whether the learning that occurred in over the 1,000 epochs of the test phase affected how the model responded to the training patterns. We wondered if, for example, the learning that occurred in the test phase may have interfered with what the model had learned in the training phase. We found no major changes in how the model responded to the training patterns after the test phase. For each of the training patterns, the output unit that was most strongly activated before the test phase was also the most strongly activated after the test phase. However, in several cases, the activation strength of the relevant output unit had decreased to some extent (e.g. from .97 to .88) and the strength of an irrelevant output unit had increased (e.g. from .02 to .19). These shifts were reflected in a higher measure of cross-entropy error for the training patterns after the test phase, which rose to $M = 42.3$, $SD = 16.7$, $t(9) = 2.79$, $p = .02$, after being tested on *offspring problems*; and to $M = 32.3$, $SD = 11.1$, $t(9) = 2.67$, $p = .02$, after being tested on *prey problems*.

Discussion

Consistent with Bulloch and Opfer (2009)'s study of similarity-based generalisation, the model's learning trajectory moves towards predictive accuracy, rather than showing a wholesale shift from relying exclusively on perceptual information to relying exclusively on relational information. The empirical and modelling results therefore align with respect to the arc and endpoint of this learning trajectory. The arc is gradual and emerges from experience. The endpoint is skilful similarity-based generalisation that takes multiple sources of information into account.

However, the empirical and modelling results differ with respect to the starting point of this developmental phenomena. Bulloch and Opfer (2009) found that the youngest participants in their study (3-year-olds) were more ambivalent in their patterns of similarity-based inference, choosing the relational match 61% of the time in offspring problems and 56% of the time in prey problems. The model, on the other hand, relies on perceptual information to make inferences in the early portion of the test phase before learning to make different inferences for offspring and prey problems.

One reason for this difference is prior knowledge. Three-year-old children came to the task with relevant experience of the contexts that were tested, whereas the model came to the task as a blank slate. As Bulloch and Opfer (2009) acknowledge, "children came to our task knowing the value of the parent-offspring relation" (p. 120), which suggests that their participants may, at least in the offspring context, experience a perceptual-to-relational shift before they turn three. Prior empirical and theoretical work supports the view that perceptual information is primary in similarity-based generalization—until, at least, people have sufficient knowledge about the relevant categories and roles to understand the relational context (Badger & Shapiro, 2012; Gentner & Rattermann, 1998; Hayes & Thompson, 2007; Sloutsky et al., 2015).

Another reason for the difference is the design of the task. Participants in the study were presented with a target and two samples. One of the samples was coded as a purely perceptual match to the target, and the other was coded as a purely relational match to the target. However, close inspection of the stimuli reveals that this distinction is somewhat fuzzy.

Notice that in Figure 1 the target juvenile (t) is a better perceptual match to the juvenile on the right (b) but the target adults (TT) are a better perceptual match to the adults on the left (AA). Since the inference questions focussed on the target juvenile, it was argued that attending to the perceptual similarity of the adults represented a relational inference. However, it is unclear if children who chose the relational option did so because of the relational condition or because of the salient perceptual similarity between the sample and target adults. This latter possibility seems especially likely since there was an overall preference for the “relational” option (even 5-year-olds in the *prey* condition chose the relational match over 45% of the time).

Nevertheless, the primary goal of the behavioural study is to show that, as children develop, they become increasingly sensitive to contextual information when making inferences about category membership. This perspective challenges the view that there are two distinct stages in the development of generalisation: one governed by perceptual similarity and another governed by relational similarity. Instead, both the behavioural and modelling results suggest that children learn to generalise flexibly and skilfully as they develop.

General discussion

The results of our simulations provide evidence that certain important phenomena in the development of similarity-based generalisation can be captured by a general-purpose model of semantic learning (Leech et al., 2008; Rogers & McClelland, 2004; Thibodeau et al., 2013). In this case, relational knowledge was learned from experience and represented sub-symbolically (rather than explicitly coded) in the connection weights and hidden layers. Relational reasoning emerged from the learning algorithm (backprop) that extracted coherent covariation from stimuli in the environment. Although the structural similarities between the juveniles were not directly perceptible (i.e. as measured in the overlap between input vectors representing the insects), the inferences that the model made in the offspring condition were consistent with what one would expect from a more structured model. Thus, our approach captures the documented primacy of perceptual information (Gentner, 1988) and the context-flexibility of relational and perceptual generalisation (Bulloch & Opfer, 2009).

While popular models like SME and LISA can likely accommodate these findings, to do so might require ad-hoc changes to existing processing algorithms in order to account for the role of context and predictive validity. The model used for the current study showed this behaviour without positing analogy-specific machinery or structured, symbolic representations (as with SME, Falkenhainer et al., 1989; LISA, Hummel & Holyoak, 1997; and DORA, Doumas et al., 2008).

On this view, the primacy of perceptual information and context-flexibility emerge naturally from learned distributed representations of objects and relations. The model provides an account of how conceptual knowledge is re-organized through experience as it acquires domain-specific knowledge (Gentner et al., 1995) and how this re-representation gives rise to relational reasoning. It does not require the concurrent development of working memory or inhibitory control (as was the case in Morrison et al., 2011; although see Kollias & McClelland, 2013 for a fully connectionist account that considers these important cognitive mechanisms).

With this said, it is important to be clear that we are not claiming that our model can account for all facets of human analogical reasoning. Many of the tasks that SME and LISA and DORA model so well (e.g. explicit analogical mapping of features and relations between domains) rely on processes that we purposefully did not try to simulate (e.g. Bowdle & Gentner, 1997; Morrison et al., 2004). Tasks that involve maintaining representations in working memory for explicit mapping or sophisticated mechanisms of inhibitory control are likely beyond the scope of the current approach (Gentner & Markman, 1995; Halford et al., 2010). For example, in one classic study of analogical reasoning, adults were first presented with a complex linguistic description of a military problem and its solution before reading about a relationally similar medical problem (Gick & Holyoak, 1980). Results indicated that participants who consciously recognised the structural similarity between the two problems were more likely solve the medical problem because they could draw on an analogy between them. The full range of cognitive processing for such a task, which involves building mental representations from language, maintaining representations of the two problems in working memory, and mapping between them is beyond the scope of the current modelling endeavour. The present work does, however, raise questions about when formal structure is necessary and provides a parsimonious account of context-sensitivity in the development of generalisation.

Conclusion

Similarity-based generalisation is fundamental to human cognition, and the ability to draw analogies based on abstract relational connections between superficially different domains is crucial for reasoning and inference (Gentner, 1983, 2010; Hofstadter, 2001; Penn et al., 2008). Learning to base generalisation on shared relations rather than (or in the face of) shared perceptual features has been identified as an important developmental milestone (Gentner, 1988; Leech et al., 2008; Piaget, 1952; Rattermann & Gentner, 1998). Unlike many other approaches to analogical reasoning that use symbolic representations and analogy-specific mapping mechanisms, we have shown that context-sensitive perceptual and relational reasoning can emerge over the course of development in a domain-general learning model that employs distributed, sub-symbolic representations.

Note

1. Nevertheless, we would argue that the nature of Bulloch and Opfer (2009)'s task does not provide strong evidence against the primacy of perceptual information. We further explore this issue in more detail in the Discussion.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Aviva Blonder  <http://orcid.org/0000-0002-0859-139X>

References

- Badger, J. R., & Shapiro, L. R. (2012). Evidence of a transition from perceptual to category induction in 3- to 9-year old children. *Journal of Experimental Child Psychology*, *113*, 131–146.
- Bowdle, B., & Gentner, D. (1997). Informativity and asymmetry in comparisons. *Cognitive Psychology*, *34*, 244–286.
- 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*, 114–122.
- Deng, W., & Sloutsky, V. M. (2016). Selective attention, diffused attention, and the development of categorization. *Cognitive Psychology*, *91*, 24–62.
- Doumas, L. A. A., Morrison, R. G., & Richland, L. E. (2018). Individual differences in relational learning and analogical reasoning: A computational model of longitudinal change. *Frontiers in Psychology*, *9*, 1235.
- Doumas, L., Hummel, J., & Sandhofer, C. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review*, *115*, 1–43.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, *41*, 1–63.
- Flusberg, S. J., & McClelland, J. L. (2014). Connectionism and the emergence of mind. *Oxford University Press Handbook of Cognitive Science (online version)*.
- Flusberg, S. J., Thibodeau, P. H., Sternberg, D. A., & Glick, J. J. (2010). A connectionist approach to embodied conceptual metaphor. *Frontiers in Psychology*, *1*, 12.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.
- Gentner, D. (1988). Metaphor as structure mapping: The relational shift. *Child Development*, *59*, 47–59.
- Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, *34*, 752–775.
- Gentner, D., & Forbus, K. D. (2011). Computational models of analogy. *WIREs Cognitive Science*, *2*, 266–276.
- Gentner, D., & Markman, A. B. (1995). Analogy-based reasoning in connectionism. In M. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 91–93). Cambridge, MA: MIT Press.
- Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on thought and language: Interrelations in development* (pp. 225–277). London: Cambridge University Press.
- Gentner, D., & Rattermann, M. J. (1998). Deep thinking in children: The case for knowledge change in analogical development. *Behavioral and Brain Sciences*, *21*(6), 837–838. doi:10.1017/S0140525X9829176X
- Gentner, D., Rattermann, M. J., Markman, A. B., & Kotovsky, L. (1995). Two forces in the development of relational similarity. In T. J. Simon, & G. S. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling* (pp. 263–313). Hillsdale, NJ: LEA.
- Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, *12*, 306–355.
- Goswami, U. (1995). Transitive relational mappings in three and four year olds: The analogy of Goldilocks and the three bears. *Child Development*, *66*, 877–892.
- Goswami, U. (1996). Analogical reasoning in cognitive development. In H. Reese (Ed.), *Advances in child development and behavior* (pp. 92–135). San Diego, CA: Academic Press.
- Goswami, U., & Brown, A. L. (1990). Melting chocolate and melting snowmen: Analogical reasoning and causal relations. *Cognition*, *35*(1), 69–95.
- Halford, G. S., Wilson, W. H., & Phillips, S. (2010). Relational knowledge: The foundation of higher cognition. *Trends in Cognitive Sciences*, *14*(11), 497–505.
- Hayes, B. K., & Thompson, S. P. (2007). Causal relations and feature similarity in children's inductive reasoning. *Journal of Experimental Psychology: General*, *136*, 470–484.
- Hinton, G. (1986). Learning distributed representations of concepts. In *Proceedings of the eighth annual conference of the cognitive science society* (pp. 1–12). Hillsdale, NJ: Erlbaum.
- Hofstadter, D. (2001). Analogy as the core of cognition. In D. Gentner, K. Holyoak, & B. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 499–538). Cambridge, MA: MIT Press.

- Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*, 427–466.
- Hummel, J. E., & Holyoak, K. J. (2005). Relational reasoning in a neurally plausible cognitive architecture. *Current Directions in Psychological Science*, *14*, 153–157.
- Kollias, P., & McClelland, J. L. (2013). Context, cortex, and associations: A connectionist developmental approach to verbal analogies. *Frontiers in Psychology*, *4*, 857.
- Lampinen, A., Hsu, S., & McClelland, J. L. (2017). Analogies emerge from learning dynamics in neural networks. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th annual Conference of the cognitive science Society* (pp. 2512–2517). Austin, TX: Cognitive Science Society.
- Leech, R., Mareschal, D., & Cooper, R. (2008). Analogy as relational priming: A developmental and computational perspective on the origins of a complex cognitive skill. *Behavioral and Brain Sciences*, *31*, 357–378.
- Lu, H., Wu, Y. N., & Holyoak, K. J. (2019). Emergence of analogy from relation learning. *Proceedings of the National Academy of Sciences*, *116*, 4176–4181.
- Morrison, R. G., Dumas, L. A. A., & Richland, L. E. (2011). A computational account of children’s analogical reasoning: Balancing inhibitory control in working memory and relational representation. *Developmental Science*, *14*, 516–529.
- Morrison, R., Krawczyk, D., Holyoak, K., Hummel, J., Chow, T., Miller, B., & Knowlton, B. (2004). A neurocomputational model of analogical reasoning and its breakdown in frontotemporal lobar degeneration. *Journal of Cognitive Neuroscience*, *16*, 260–271.
- Opfer, J. E., & Bulloch, M. J. (2007). Causal relations drive young children’s induction, naming, and categorization. *Cognition*, *105*, 206–217.
- Penn, D., Holyoak, K., & Povinelli, D. (2008). Darwin’s mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, *31*, 109–130.
- Piaget, J. (1952). *The child’s concept of number*. New York: Norton.
- Rattermann, M., & Gentner, D. (1998). More evidence for a relational shift in the development of analogy: Children’s performance on a causal-mapping task. *Cognitive Development*, *13*, 453–478.
- Rogers, Timothy T., & McClelland, James L. (2008). Précis of *Semantic Cognition: A Parallel Distributed Processing Approach*. *Behavioral and Brain Sciences*, *31*(6), 689–714. doi:10.1017/S0140525X0800589X.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition*. Cambridge, MA: MIT Press.
- Rumelhart, D. (1990). Brain style computation: Learning and generalization. In S. Zornetzer, J. L. Davis, & C. Lau (Eds.), *An introduction to neural and electronic networks* (pp. 405–420). San Diego, CA: Academic Press.
- Sloutsky, V. M., Deng, W., Fisher, A. V., & Kloos, H. (2015). Conceptual Influences on Induction: A case for a Late Onset. *Cognitive Psychology*, *82*, 1–31.
- Tarlowski, A. (2018). Ontological constraints in children’s inductive inferences: Evidence from a comparison of inferences within animals and vehicles. *Frontiers in Psychology*, *9*, 520.
- Thibodeau, P. H., Flusberg, S. J., Glick, J. J., & Sternberg, D. A. (2013). An emergent approach to analogical inference. *Connection Science*, *25*, 27–53.