Simulating Infant Visual Learning by Comparison: An Initial Model

Kezhen Chen (kezhenchen2021@u.northwestern.edu)
Department of Computer Science, Northwestern University
Evanston, IL 60208 USA

Kenneth D. Forbus (forbus@northwestern.edu)
Department of Computer Science, Northwestern University
Evanston, IL 60208 USA

Dedre Gentner (gentner@northwestern.edu)
Department of Psychology, Northwestern University
Evanston, IL 60208 USA

Susan J Hespos (hespos@northwestern.edu)
Department of Psychology, Northwestern University
Evanston, IL 60208 USA

Erin M. Anderson (e.m.anderson478@gmail.com)
Department of Psychology, Northwestern University
Evanston, IL 60208 USA

Abstract

Researchers have recently found that 3-month-old infants are capable of using analogical abstraction to learn the *same* or *different* relation, given the right conditions (Anderson et al. 2018). Surprisingly, seeing fewer distinct examples led to more successful learning than seeing more distinct examples. This runs contrary to the prediction of standard learning theories, which hold that a wider range of examples leads to better generalization and transfer, but is compatible with other findings in infant research (Casasola 2005; Maguire et al. 2008). Anderson et al. (2018) propose that this is due to interactions between encoding and analogical learning. This paper explores that proposal through the lens of cognitive simulation, using automatically encoded visual stimuli and a cognitive model of analogical learning. The simulation results are compatible with the original findings, thereby providing evidence for this explanation. The assumptions underlying the simulation are delineated and some alternatives are discussed.

Keywords: Analogy, Relational Learning, Cognitive Simulation, Infant Cognition

Introduction

Relational learning and reasoning are central in human cognition (Gentner, 2003, 2010; Gentner & Markman 1997; Hofstadter, 2001; Penn et al., 2008). How does this ability arise? Is analogical ability built up gradually via maturational change, or by combining other component processes? Or is structure-mapping ability an innate species-level adaptation? The first possibility may seem more plausible, given the abundant evidence that relational sophistication increases over development (Gentner & Rattermann, 1991). But recent findings suggest that analogical processing ability may be present early on, and that developmental gains in analogical fluency are due to increases in relational knowledge (Gentner, 2010; Gentner & Rattermann, 1991) and/or executive ability (Richland et al., 2006; Thibaut et al., 2010). For example, Ferry, Hespos and Gentner (2015) found evidence that 7-9-month-old infants can carry out analogical abstraction across a sequence of exemplars to derive an abstract *same* or *different* relation.

Anderson et al. (2018) recently reported that even 3-month-old infants can learn *same* or *different* relations via analogical abstraction. A surprising aspect of the research was that the infants learned better when given fewer examples. In the first experiment, infants failed to learn these relations after being shown six distinct examples of either *same* or *different* repeated until habituation. (The exact number of habituation trials varied, ranging from 6 to 9 trials until infants' looking times declined by 50% from the first three trials to the last three, or until infants had completed nine trials.) In the second experiment, infants succeeded after being given repeated exposure to only two examples of the relation. This result runs contrary to the predictions of standard learning theories, which predict that a wider range of examples leads to better generalization and transfer, but is compatible with some prior findings on infant relational learning (Casasola 2005; Maguire et al. 2008).

Anderson et al. (2018) proposed that these phenomena are due to interactions between encoding and analogical processing. This paper examines this proposal via cognitive modeling, using automatically encoded stimuli and a model of analogical learning. Specifically, we ask whether the 3-month-old pattern can be modeled by assuming that the infants have structure-mapping ability, but that they are limited by their encodings of examples. We lay out a set of assumptions that provide a possible processing account and show that these assumptions could explain the generalization pattern. The modeling enterprise also reveals other possible encoding assumptions, which can be explored in future work.

We first review prior research on analogical abstraction, then describe the Anderson et al. (2018) experiments to be modeled. Then we describe our model of the infants' encoding and learning process. To preview, the model is constructed from pre-existing components (described below). This includes automatic encoding of visual stimuli based on photos of the objects shown to the infants. We describe the processing performed by the model, laying out the assumptions we are making and noting where alternative explanations are feasible. Then we present the results of computational simulation based on the model. We end with a discussion of the implications and possible future work

Background

There is evidence of analogical ability in children from early preschool through adulthood (Gentner, 2003; Gentner & Rattermann, 1991; Richland et al., 2006). Two signatures of this ability are (1) the ability to perceive abstract relational matches can be enhanced by comparing instances of a relation, in both adults (Gick & Holyoak, 1983; Markman & Gentner, 1993) and children (Gentner, 2003; Kotovsky & Gentner, 1996); and (2) the presence of salient objects can interfere with relational mapping, especially early in development (Gentner & Toupin, 1986; Paik & Mix, 2008; Richland et al., 2006). These findings are consistent with other research suggesting that comparison entails a structural alignment process that highlights relational commonalities between the items compared (Markman & Gentner, 1993).

Recent research has explored relational learning in human infants (Anderson et al., 2018; Ferry et al., 2015; Gervain et al., 2012). Ferry et al. (2015) found evidence that 7 to 9-months-old infants can engage in analogical abstraction. When shown a series of *same* pairs (using the method described below), infants afterwards looked significantly longer at a novel *different* pair than at a novel *same* pair (and the reverse for habituation to *different*). This is evidence for the first signature of analogical processing—that comparing across examples promotes abstracting the common relational structure. They also found evidence for the second signature of analogical processing: that salient objects tend to distract from relational processing. When infants were shown a subset of objects prior to habituation, they performed poorly on test trials containing these objects, failing to distinguish *same* and *different*. Thus, Ferry et al. (2015) concluded that by 7-9-months, infants can use analogical generalization to form a relational abstraction.

Analogical Learning in 3-month-old Infants

To explore the origins of analogical ability, Anderson et al. (2018) asked whether 3-month-olds could abstract *same* and *different* relations. Infants were shown a series of pairs: half the infants saw *same* pairs and the other half saw *different* pairs[1]. The materials were pairs of colorful, distinctive objects (Fig. 1). In order to engage infants' attention, on each habituation and test trial, the pair was moved together

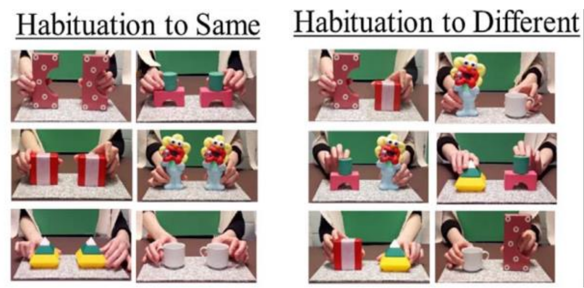

Figure 1: Examples of habituation pairs from Anderson et al. Expt. 1 (2018).

through a fixed motion path: up, then tilted left, then right, then down to the start point. This 8-second cycle was repeated continuously until the infant looked away for 2 seconds. Then the next pair was shown in the same way. The habituation trials continued until the infant's looking time declined by 50% from the first three trials to the last three, with a maximum of nine trials (range = 6 to 9 trials).

Both groups of infants then saw the same six test pairs—three depicting the *same* relation and three depicting the *different* relation. The pairs were shown one at a time, and the key dependent measure was how long the infant looked at each pair. The key test pairs had brand new objects instantiating either the *same* or *different* relation. If infants have abstracted the relation they saw during habituation, they should look longer at the novel relation. (This use of looking-time is commonly used with preverbal infants; the idea is that the familiar relation will fit their expectations, whereas the novel relation will be more surprising.)

In Experiment 1, infants were shown six distinct pairs (either all *same* or all *different*) during habituation. During test, the infants failed to look longer at novel pairs on the key trials. Thus, there was no evidence for analogical learning. Although this could mean that 3-month-olds lack this ability, the experimenters explored another possibility: that the

[1] To test for salient-object interference, the infants had previously seen some objects in the waiting room; this is not modeled here.

infants were overwhelmed by the variety of objects in the study, and thus failed to encode the relations between them (e.g., Casasola & Park, 2013). Consequently, in Experiment 2, only two distinct pairs were used during habituation (e.g., AA, BB, AA, BB...for *same*). The infants were then tested in the same manner as in Experiment 1. In this case, infants did indeed learn. They looked longer at pairs showing the novel relation, even with brand new objects—evidence that they had abstracted the relation.

Simulating the Infants' Learning

In order to abstract a *same* or *different* relation from a series of examples, two things must happen (not necessarily in a fixed order): (1) the learner must compare the objects *within* each pair to form some initial representation of the *same* (or *different*) relation within the pair; and (2) the learner must compare *across* the pairs to arrive at a more abstract encoding of the relation. Only if both those things happen will the learner experience a brand new *same* pair as familiar. Our simulation explores one path—by no means the only path—by which this could happen.

Simulation Design

Here we discuss our simulation. We begin by noting a critical point: in order to be informative about human cognition, a simulation must be constrained. Many simulations have used hand-coded representations to depict the learner's construal of a situation, and/or have implemented a simulation process specific to the situation being modeled. But this allows enormous latitude to tailor the representations and processes to fit whatever outcome is desired. To avoid this problem, (a) as input, the model is given representations that are automatically encoded from the visual stimuli given to the infants; and (b) our processing model is built out of pre-existing components that have successfully simulated prior findings in analogical processing.

We first describe the component models, then how they are combined.

Simulation of analogical processing

We use the Structure-Mapping Engine (SME, Forbus et al. (2016)) as a simulation of analogical mapping, and SageWM (Kandaswamy et al. 2014) as a simulation of analogical generalization in working memory. These models have been used to model a number of psychological phenomena already.

SME is based on Gentner's (1983) structure-mapping theory of analogy and similarity. Given two cases consisting of structured relational representations, SME computes one or more *mappings* between them, preferentially aligning common relational structure. A mapping includes a set of *correspondences* that align entities and statements in the base and target, a *similarity score* that indicates how similar the base and the target are, and *candidate inferences*, which are projections of additional structure from one case to the other, based on the aligned structure. SME also computes a structural evaluation score—a similarity score that takes into account the depth of the common structure as well as the among of overlap. Here SME is used both as a similarity metric and as a means of combining cases into generalizations in SageWM.

SageWM is the working-memory version of SAGE (McClure et al. 2015), the Sequential Analogical Generalization Engine. It provides a model of analogical abstraction. SageWM creates new generalizations from a series of examples, by iterative application of SME. When given a series of examples, SageWM stores the first example. When the next example arrives, SageWM compares it to the first one, using SME. If there is sufficient overlap (that is, if SME's score is above a pre-set assimilation threshold), the common structure is stored as a generalization. If the similarity to the abstraction is below threshold, the example will be stored separately. This process continues as new examples arrive. Thus, if new examples are sufficiently similar to the ongoing generalization, then the generalization will be updated to be somewhat more abstract. We use 0.95 as the assimilation threshold in these experiments[2] which is the default for SageWM.

Simulation of visual encoding

In order to avoid hand-coding the stimuli, we use CogSketch (Forbus et al. 2017), a pre-existing cognitive model of visual encoding and visual problem solving, to provide a vocabulary of visual representations. CogSketch has successfully modeled a variety of adult visual tasks, including Ravens' Progressive Matrices (Lovett & Forbus, 2017), an oddity task (Lovett & Forbus, 2011), and a paper-folding task (Lovett & Forbus, 2013).

The production of visual stimuli occurs via an automatic pipeline, starting with photographs the pairs of objects provided by the original experimenters. The photographs are blurred and the Canny edge detector is used to generate a set of initial edges describing each object. CogSketch decomposes these initial edges into segments based on discontinuities and intersections. CogSketch automatically computes a variety of information about each segment, e.g. its length, curvature, orientation, position, and topological relations with other segments. This graph of segments and junctions is also used to identify regions within an object (McLure et al. 2011). This includes the object's boundary, consisting of all exterior edges, which we assume is visually salient and likely to be encoded early in human processing. Several kinds of information are automatically encoded for regions as well, such as whether or not it has curved sides. CogSketch also estimates its closeness to a set of shape templates, e.g. spindles, triangles, rectangles, and ellipses. Since color is visually salient, we use a color extraction

[2] As descriptions are merged, frequency counts are kept for how often each statement is aligned. If the probability goes below a threshold (0.2 by default), the statement is eliminated.

library to extract up to eight of the most frequent colors for an object.

An important issue in this modeling effort is to consider the visual encoding processes available to 3-month-olds. In the first months of life, vision and attentional processes are becoming increasingly stable (Arterberry & Kellman, 2016; Colombo, et al., 1991). Visual acuity improves steadily through the first several months. Especially relevant here, infants' habituation and fixation periods decrease dramatically during the first 6 months (Bornstein, 1985; Colombo & Mitchell, 2009), suggesting that young infants' encoding is slower and more variable than that of older infants.

To capture young infants' relatively inefficient encoding processes, here we have assumed *slow encoding*—that is, that not all the available information is encoded on first exposure. (Other assumptions are possible, including variable extraction of information.) Specifically, we assume that the boundary of an object, its shape properties, and color are encoded early. When given more time, we assume infants compute higher-level representations of the shape, including internal properties and relations. We use the scheme from Chen et al. (2019), which is inspired by Biederman's (1987) Recognition by Components theory, which describes shapes as being made out of parts called *geons*. CogSketch identifies geons by using a medial axis transform, compatible with Biederman's original account and consistent with evidence from Lowet et al. (2018) regarding their use in human shape representations. Object-internal relationships between geons are computed in terms of positional relations and qualitative topological relations (Cohn et al.1997). Figure 2 shows examples of boundary and geon representations for one of the objects. We further assume that, given sufficient time, infants encode representations of both objects.

```
;; Overall properties
(WhiteColor A)
(RedColor A)
(highRectangularity R-1)
(ellipseSystemShape R-1)
(hasConcavedEdge R-1)
(someCurvedCycle R-1)
[…] ;; Internal regions
(LowRectangularity  R-34
(CycleSystemShape R-34)
(NonConcavedEdge R-34)
(AllCurvedCycle R-34)
(above R-29 R-34)
(rightOf R-34 R-29)
[…]
```

Figure 2: Example of boundary and Geon relational Representations. These relational representations are automatically generated. For the readers' convenience, we use English word to indicate relations and entities.

In the original experiments, the pairs were moved in a uniform way throughout the habituation and test trials. We assume that the infants encode these motions, since motion is extremely salient for them. While we could use qualitative spatial representations to automatically represent the specific motions of the stimuli as part of the encoding process, using techniques like (Chen & Forbus, 2018), this would involve considerable complexity to gather the video data. Thus, we do not explicitly encode such motions in the present model.

We hypothesize that the repeated motion influenced the infants' processing in two ways. First, within a trial, the two objects in a pair always move together. This gives rise to a perception of the unity of the pair and prompts the infant to compare the two objects in a pair. Over trials, as the object representations become more detailed, this will lead to perceiving many common attributes in a same pair (or few, in the case of a different pair). We call the representation of the two objects plus relations computed between them the *pair-level* description. We hypothesize that pair-level descriptions are only computed when both objects have been fully encoded. The second effect of the repeated motion is to invite comparison across trials: even though the individual pairs (say, AA and BB) are quite distinct, we hypothesize that the similarity in their motion leads the infant to compare them, as described below.

To represent the visual similarity of objects, we use one of two relations, depending on whether their similarity, as measured by SME, is above a particular threshold (here, 0.5). If their similarity is above the threshold, a statement using the **sameObject** relation is encoded, and otherwise, **differentObject**. We use these terms for convenience, but we do not assume that infants distinguish absolute sameness from high similarity (see Smith, 1993). It is also not clear whether infants are learning these relations de novo, or whether they already possess some kind of representations of *same* and *different*, either innately or through early learning. We return to this question in the Discussion.

Processing Assumptions

To recapitulate, we assume that infants encode the motion of the pair of objects and that this invites comparison both within and across trials. However, the comparison process become also requires that the object representations be sufficiently detailed. We do not assume that infants encode everything about the objects in a trial at first exposure. Here we assume that information about object boundaries and color are computed first, followed by information about the decomposition of the object into geons, and that these two levels of representation occur in that sequence. We assume that even partial object representations are stored in SageWM, and retrieved the next time they are exposed to the pair. This retrieval speeds up the initial encoding process, allowing processing to move on to the next level of encoding.

It is not clear whether infants are encoding both objects on first exposure to a pair. Here we assume that objects are encoded independently in parallel, but with the levels of representations outlined above. We assume that having the objects placed into correspondence causes them to be compared, once their encodings are complete. The result of this comparison results in the description of the pair being augmented with a **sameObject** or **differentObject** statement, depending on the outcome of that comparison.

Experiment Simulation

Now let us reconsider the experiments in Anderson et al. (2018) through the lens of cognitive simulation. We discuss each in turn. In both simulations, we did not simulate the infants' experience of some objects from waiting room.

Simulation of Experiment 1

Following the original experiment, we simulated two habituation sequences: one with a sequence of six pairs of objects satisfying the *same* relationship (<A,A>, <B,B>, <C,C>, <D,D>, <E,E>, <F,F>) and one with a sequence of six pairs of objects satisfying the *different* relationship (<A,B>, <C,D>, <E,F>, <B,C>, <F,A>, <D,E>). Given our assumption of parallel object encoding, in the *same* condition only the first level of encoding occurs for each object in the simulation, and hence the objects are not compared and no pair-level descriptions are generated. For the *different* condition, there are repeated exposures to particular objects, but another comparison involving them would be needed to generate pair-level representations. Since there are no pair-level examples, they cannot be compared and generalized, and hence no analogical learning takes place, compatible with the infant results.

Simulation of Experiment 2

Following the original experiment, two sequences of alternating pairs of objects were used. For the *same* habituation trials, these were (<A,A>, <B,B>, <A,A>, <B,B>, <A,A>, <B,B>, and for the *different* habituation trials, these were <A,B>, <C,D>, <A,B>, <C,D>, <A,B>, <C,D>. Thus for both habituation conditions, each pair was presented to the simulation three times, in alternation. In the first exposure to a pair, the first level of encoding occurs for its objects, which are stored in SageWM. In the second exposure, the second level of encoding occurs, building on the initial model stored in SageWM. In the third exposure, the fully-encoded objects retrieved are used to construct a pair description, including the cross-object comparison (because of the assumed common roles in the motion perceived by the infants). That pair description is also stored in SageWM. The pair representations are generalized by SageWM across pairs as they occur: that is, a generalization is formed that includes either a **sameObject** or a **differentObject** statement, depending on habituation condition. This new abstraction is relatively portable, since it has many fewer object details in common, and hence is retrieved when test pairs are presented. Even if these test pairs are not fully encoded (because of novel objects), alignment with the abstraction leads to a projection of a **sameObject** or **differentObject** statement as a candidate inference (depending on whether habituation was for *same* or *different*). When a test pair is compatible with the learned relation, the candidate inference fits. When a test pair is incompatible with the learned relation, the candidate inference is contradicted, and this novelty, we hypothesize, leads to greater looking times for the infant.

Discussion

The simulation captures the pattern of infant results across the two experiments: When given six different example pairs (Experiment 1), the simulation fails to form abstractions of *same* and *different* during habituation, and therefore fails to differentiate novel from familiar relations during test. When given two pairs (Experiment 2), the simulation forms abstractions of *same* and *different* during habituation, and therefore arrives at distinct matching scores for novel vs. familiar relations during test.

Thus, we have shown that a reasonable set of assumptions about the visual encoding of infants, along with pre-existing encoding algorithms and analogical process models, can be used to simulate Anderson et al.'s (2018) results on analogical learning in 3 month old infants. This provides evidence for their proposed explanation, in terms of partial infant encoding.

This simulation assumed that something like **sameObject** and **differentObject** were already available to infants. How might such relationships be learned, even perhaps during the experiment? It is not unreasonable, given how ubiquitous analogy and similarity appear to be in human cognition (Gentner 2003), that infants can remember the qualitative feeling of high-similarity or low-similarity for pairs that they have just seen. In other words, the alignments during analogical generalization could provide the basis for introducing a simple qualitative value on similarity, e.g. high or low (Forbus, 2019). For example, given habituation on *same* trials, these similarity scores will tend to cluster quite high, and given habituation on *different* trials, these similarity scores will tend to cluster quite low (see Figure 3). Seeing a score for a pair in the same role that is substantially different, i.e. a different qualitative value, could also predict looking times and reifying such a difference into a pair of relationships would then make such information accessible in future comparisons. This provides a possible explanation for how such relationships can be learned.

Our general assumption is that the rather surprising pattern—that 3-month-olds can from an abstraction from two alternating pairs over six pairs but not from six different pairs—results from inefficiencies in their visual encoding process. In this simulation, we have focused on slow encoding to capture this inefficiency. Another interesting possibility is variable encoding. For example, different subsets of geons might be computed over different exposures, so that the perceived similarity of a pair over time would depend on the particular orders in which geons were found. Such models will be explored in future work.

Despite the vast amount of research on analogical processing in children, there is very little research on how children learn relations in the first place. One exception is

DORA (Doumas et al. 2008). DORA begins with unstructured representations of objects as simple feature vectors. When DORA compares two or more objects, it forms explicit representations of any properties they share. These properties are then combined into relations. This contrasts with our model, in which the relations are formed from online differences in qualitative similarity.

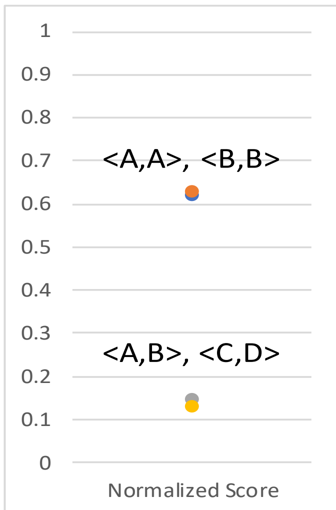

Figure 3: Each point represents the normalized similarity score for same vs. different pairs from Experiment 2. Distributional properties of numerical similarity estimates provide a possible signal for introducing qualitative values and learning *same/different* relationships.

Conclusion

Our results lend support to the idea that 3-month-old infants have structure-mapping ability, but are limited by their encodings of examples. Here we have shown that a reasonable set of assumptions about encoding and the use of analogical generalization within working memory simulate the experiments from Anderson et al. (2018). The simulation provides an explanation for why 3-month old infants are able to learn, or not learn, *same/different* relations.

We see a number of paths for future work. First, we think encoding variability may be an important factor in explaining the conditions under which infants can learn. Second, we want to simulate a wider range of experiments with this model, including experiments with older infants (e.g. Ferry et al. 2015). This will involve developing and testing plausible models for how encoding skills change across development with experience and building up models of long-term experiences and generalizations that infants accumulate.

Acknowledgements

This research was supported by the Machine Learning, Reasoning, and Intelligence Program of the Office of Naval Research awarded to Kenneth D. Forbus and by National Science Foundation Grant BCS-1729720 awarded to Susan J. Hespos and Dedre Gentner.

References

Anderson, E., Chang, Y., Hespos, H. & Gentner, D. (2018). Comparison within pairs promotes analogical abstraction in three-month-olds. *Cognition*, 176:74-86.

Arterberry, M. E., & Kellman, P. J. (2016) *Development of perception in infancy: The cradle of knowledge revisited*, 2nd Edition. New York: Oxford University Press.

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

Casasola, M. (2005). When less is more: How infants learn to form an abstract categorical representation of support. *Child Development*, 76(1), 279–290. http://dx.doi.org/10.1111/j.1467-8624.2005.00844.x.

Casasola, M., & Park, Y. (2013). Developmental changes in infant spatial categorization: When more is best and when less is enough. *Child Development*, 84(3), 1004–1019.

Chen, K. and Forbus, K.D. (2018). Action Recognition from Skeleton Data via Analogical Generalization over Qualitative Representations. *Proceedings of AAAI 2018.*

Chen, K., Rabkina, I., McLure, M., & Forbus, K. (2019) Human-like Sketch Object Recognition via Analogical Learning. *AAAI 2019*.

Cohn, A. G., Bennet, B., Gooday, J., and Gotts, N. M. 1997. Qualitative Spatial Representation and Reasoning with the Region Connection Calculus. *GeoInformatica* 1: 275-316.

Colombo, J., Mitchell, D. W., Coldren, J. T., & Freeseman, L. J. (1991). Individual differences in infant visual attention: Are short lookers faster processors or feature processors? *Child Development*, 62(6), 1247–1257. doi.org/10.2307/1130804

Doumas, L., Hummel, J., & Sandhofer, C. (2008). A theory of the discovery and prediction of relational concepts. *Psychological Review*, *111:*1-43.

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

Forbus, K. (2019). *Qualitative Representations: How People Reason and Learn about the Continuous World*, MIT Press

Forbus, K., Usher, J., Lovett, A., Lockwood, K., and Wetzel, J. (2011). CogSketch: Sketch understanding for Cognitive Science Research and for Education. *Topics in Cognitive Science*, 3(4), pp 648-666.

Forbus, K. D., Ferguson, R. W., Lovett, A., and Gentner, D. (2016). Extending SME to handle large-scale cognitive modeling. *Cognitive Science*, DOI: 10.1111/cogs.12377, pp 1-50.

Forbus, K. D., Chang, M., McLure, M. & Usher, M. (2017). The Cognitive Science of Sketch Worksheets. *Topics in Cognitive Science*. DOI: 10.1111/tops.12262

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

Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp. 195–235). 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., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10, 277-300.

Gervain, J., Berent, I., & Werker, J. F. (2012). Binding at birth: The newborn brain detects identity relations and sequential position in speech. Journal of Cognitive Neuroscience, 24(3), 564-574.

Bornstein, M. H. (1985). Habituation of attention as a measure of visual information processing in human infants: Summary, systematization, and synthesis. In G. Gottlieb & N. A. Krasnegor (Eds.), *Measurement of audition and vision in the first year of postnatal life: A methodological overview* (pp. 253–300). Norwood, NJ: Ablex.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.

Kandaswamy, S., Forbus, K., and Gentner, D. (2014). Modeling Learning via Progressive Alignment using Interim Generalizations. *Proceedings of the Cognitive Science Society*.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, 67, 2797-2822.

Lovett, A., and Forbus, K. (2011). Cultural commonalities and differences in spatial problem solving: A computational analysis. *Cognition* 121, pp. 281-287.

Lovett, A.; and Forbus, K. 2013. Modeling spatial ability in mental rotation and paper-folding. *Proceedings of CogSci 2013*.

Lovett, A.; Forbus, K. 2017. Modeling visual problem solving as analogical reasoning. *Psychological Review*, 124(1).

Lowet, A. S., Firestone, C., & Scholl, B. J. (2018). Seeing structure: Shape skeletons modulate perceived similarity. *Attention, Perception, & Psychophysics*, 80(5), 1278-1289.

Maguire, M. J., Hirsh-Pasek, K., Golinkoff, R. M., & Brandone, A. C. (2008). Focusing on the relation: Fewer exemplars facilitate children's initial verb learning and extension. *Developmental Science*, 11(4), 628–634.

Markman, A. B., & Gentner, D. (1993). All differences are not created equal: A structural alignment view of similarity. *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, 682-686.

McLure, M. D., Friedman, S. E., Lovett, A. and Forbus, K. D. (2011). Edge-cycles: A qualitative sketch representation to support recognition. *Proceedings of the 25th International Workshop on Qualitative Reasoning*. Barcelona, Spain.

McLure, M.D., Friedman S.E. and Forbus, K.D. (2015). Extending Analogical Generalization with Near-Misses. *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, Austin, Texas.

Paik, J. H., & Mix, K. S. (2008). It's all relative: Different levels of relational similarity used in children's comparisons. *British Journal of Developmental Psychology*, 26(4), 499–505.

Richland, L. E., Morrison, R. G., & Holyoak, K. J. (2006). Children's development of analogical reasoning: Insights from scene analogy problems. *Journal of Experimental Child Psychology*, 94, 249-271.

Smith, L. B. (1993). The concept of same. In H. W. Reese (Ed.), *Advances in child development and behavior* (Vol. 24, pp. 215-252). San Diego, CA: Academic Press.

Thibaut, J, French, R., & Vezneva, M. (2010). The development of analogy making in children: Cognitive load and executive functions, *Journal of Experimental Child Psychology 106*, 1–19.