TOPICS
TOPICS IN COGNITIVE SCIENCE

Topics in Cognitive Science 9 (2017) 672–693

ISSN:1756-8757 print / 1756-8765 online
DOI: 10.1111/tops.12278

This article is part of the topic "2016 Rumelhart Prize Issue Honoring Dedre Gentner," Jeffrey Loewenstein and Arthur B. Markman (Topic Editors). For a full listing of topic papers, see http://onlinelibrary.wiley.com/doi/10.1111/tops.2017.9.issue-2/issuetoc

Analogy and Abstraction

Dedre Gentner, Christian Hoyos

Department of Psychology, Northwestern University

Received 13 June 2016; received in revised form 13 March 2017; accepted 17 March 2017

Abstract

A central question in human development is how young children gain knowledge so fast. We propose that analogical generalization drives much of this early learning and allows children to generate new abstractions from experience. In this paper, we review evidence for analogical generalization in both children and adults. We discuss how analogical processes interact with the child's changing knowledge base to predict the course of learning, from conservative to domain-general understanding. This line of research leads to challenges to existing assumptions about learning. It shows that (a) it is not enough to consider the distribution of examples given to learners; one must consider the processes learners are applying; (b) contrary to the general assumption, maximizing variability is not always the best route for maximizing generalization and transfer.

Keywords: Analogy; Abstraction; Overhypotheses

1. Introduction

Abstract structured knowledge is a key feature of higher order cognition (Gentner & Medina, 1998; Hummel, 2011; Markman, 1999; Tenenbaum & Griffiths, 2001). The assertions that make up abstract knowledge are variously referred to as schemas, rules, abstractions, principles, or overhypotheses. Many such abstractions are expressed as relational categories—categories like *evidence*, *counterfactual*, and *proportion,* and on a more mundane level*, bargain, ally,* and *rescue.* These concepts cover a lot of ground; we can

Correspondence should be sent to Dedre Gentner, Department of Psychology, Northwestern University, 2029 Sheridan Road, Evanston, IL 60208, USA. Email: gentner@northwestern.edu

talk about the proportion of red balls in an urn, the proportion of females in the workforce, or the proportion of acres of rain forest remaining. Adults have access to a vast number of such abstractions. An adult faced with a new phenomenon can draw on any of a large number of known principles with which to see important patterns and make relevant inferences.

Because the term *abstraction* has been variously defined and used within cognitive science (Burgoon, Henderson, & Markman, 2013), we begin by giving our usage. We take the process of abstraction to be one of decreasing the specificity (and thereby increasing the scope) of a concept. For example, *causal system* is more abstract than *positive feedback system*, which in turn is more abstract than the specific positive feedback system by which the melting of polar ice causes lower reflectance of the sun's heat, leading in turn to more rapid melting. Our main focus is on *relational abstractions*, including principles, rules, and schemas, as well as abstract relational categories. Relational categories are categories for which the basis for membership is participation in a common relational structure; thus, they differ from the more studied entity categories, such as *tulip* and *spoon*, whose members share many intrinsic properties. For example, *carnivore* and *herbivore* are abstract relational categories, while *canine* and *feline* are abstract entity categories. Relational categories have been the focus of much recent research (Asmuth & Gentner, 2017; Gentner, 2005; Gentner & Kurtz, 2005; Goldwater & Markman, 2011; Markman & Stilwell, 2001; Ross & Murphy, 1999), in part because of their important role in conceptual learning and education (Goldwater & Schalk, 2016).

How are such abstractions acquired by children? Some of them, such as the knowledge that objects cannot pass through one another, may be present from birth, or it may require only minimal experience to release them. As another example, even young infants appear to interpret patterns of moving geometric shapes in terms of animate beings helping or hurting (Csibra, 2008; Hamlin, Wynn, & Bloom, 2007), suggesting that this abstract social construal is present from very early. However, it seems clear that much of our abstract knowledge is learned. For example, 7–10-year-olds require extensive training to learn that the idea of a *control variable* is an abstraction that applies across domains (Chen & Klahr, 1999); and even college students fail to notice and use important causal abstractions such as *positive* and *negative feedback* without concerted training (Goldwater & Gentner, 2015; Rottman, Gentner, & Goldwater, 2012).

How then are such abstractions acquired? For older children and adults, new abstractions can be learned via explanation and instruction from parents or teachers. But the earlier we go in development, the less able children are to comprehend verbal explanations of abstract ideas. In contrast, there is evidence that analogical comparison and abstraction processes are present in 7–9-month-old infants, and even earlier (Anderson, Chang, Hespos, & Gentner, under review; Ferry, Hespos, & Gentner, 2015). We propose that—both in the history of language and in children's learning—analogical processes are a major way in which new relational abstractions are acquired. We first review research on analogical processing and show that it can lead to relational abstractions. We illustrate our points with examples from adults and children, including examples from language evolution, and across both perceptual and conceptual domains. Then we show that comparison processes, given some reasonable assumptions about the state of early representations, can predict the course of children's experiential learning. Along the way, we will respond to some potential challenges to this kind of account.

1.1. Projection versus alignment

Analogy is often thought of chiefly as a way to transfer knowledge from one situation to another, and indeed, it often serves that function. Projecting information from a well-understood domain can lend structure to an unfamiliar domain, as in:

The mitochondria are the power supply for a cell.

In this analogy, knowledge about a power supply is projected to the mitochondria. But analogy can also operate in *mutual alignment*[1] analogies to reveal commonalities that were previously not obvious in either analog. For example, a related analogy between oxidation in cell biology and oxidation in a fire reveals new aspects of both phenomena:

A fire consumes fuel using oxygen, thereby producing energy in the form of heat; it releases carbon dioxide and water.

Likewise, a mitochondrium in a cell obtains energy from glucose using oxygen, thereby producing energy in the form of ATP; it releases carbon dioxide and water.

For the nonspecialist, this analogy offers new insights into both phenomena—for example, that cell metabolism can be seen as akin to burning, and that a fire releases water (a component of smoke).

Analogies typically involve both alignment and projection, but not always in equal degrees. We see the same contrast in the use of metaphoric language. In the strongly projective case, there is a clear asymmetry between the two terms: The base concept includes some prominent information that can be applied to the target. For example, a conventional metaphoric base (e.g., *goldmine*) conveys a standard abstraction (Bowdle & Gentner, 2005; Glucksberg & Keysar, 1990): If you hear that "X is (like) a goldmine," you understand that X must be the source of something valuable. At the other extreme, consider a comparison with a base term that lacks an established metaphorical abstraction (e.g., "X is (like) a snowflake"). In this case, the common structure emerges only through aligning the two analogs. For example, "A child is like a snowflake" might be interpreted to mean "Each one is unique." But given "A mayfly is like a snowflake," the commonality that might emerge might be that both are ephemeral.

This contrast is reflected in the way we express figurative comparisons. Figurative comparisons with novel bases are expressed as similes (An X is like a Y—inviting comparison between the two literal concepts); those with conventional bases are expressed as metaphors (An X is a Y—reflecting the fact that the base form has an established abstract category) (Bowdle & Gentner, 2005; Wolff & Gentner, 2011). In language evolution, this Career of Metaphor, from novel to conventional figurative, is a route for the emergence of new abstract concepts. For example, *sanctuary* once referred only to churches and other religious edifices; now it also has an established abstract meaning of "a place of safety."

Thus, analogical/metaphorical comparisons can be used to reveal common conceptual structure via alignment, or to project an established schema from one concept to the other, via alignment plus projection. Projective analogies are widely used in adult life, as can be seen by the ubiquity of conventional metaphoric systems in communication (e.g., *Love is a journey* or *Anger is heat*) (Lakoff & Johnson, 1980) and in reasoning (for example, in the use of spatial metaphors such as "The holidays are fast approaching" to reason about time [Boroditsky, 2000]).

But the situation is very different for young children. Projective analogy requires a known base situation to draw on. Because young children lack a large supply of stored well-understood and/or conventionalized base situations, they are far more reliant on mutual alignment as a source of insight. We next discuss the process of alignment, and then turn to its role in abstraction.

1.2. Analogical processing

We frame our account in terms of structure-mapping theory (Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1983, 2010; Gentner & Markman, 1997). The main tenets are broadly accepted in accounts of analogical processing (e.g., Doumas, Hummel, & Sandhofer, 2008; Hummel & Holyoak, 2003; see Kokinov & French, 2003; Gentner & Forbus, 2011, for reviews). Comparing two situations involves a process of establishing a *structural alignment* between the two representations—that is, an alignment based on matching like relations between the analogs.

However, the alignment process itself (as modeled by SME, the Structure-mapping Engine) is sensitive to both object matches and relational matches. It begins with a free-for-all, in which local matches at all levels are made, to one or a few global mappings that are maximal or near-maximal in terms of the size and depth of the common relational structure. (Thus, it is neither top-down nor bottom up, but local-global.) Matches that are connected to other matching elements are favored, so relational matches tend to win out in the final mapping. However, it is possible for object matches to win out. The richer and more distinctive the object matches, and the shallower the relational structure, the more likely it is that object matches will prevail (Gentner & Rattermann, 1991; Markman & Gentner, 1993; Paik & Mix, 2006). (For more details, see Gentner & Markman, 1997, and Forbus, Ferguson, Lovett, & Gentner, 2016).

People prefer alignments that are *structurally consistent* (Gentner, 1983; Krawczyk, Holyoak, & Hummel, 2004, 2005; Markman, 1997)—that is, mappings that follow *one-to-one correspondence* and *parallel connectivity* (meaning that if two relations are placed in correspondence, then their arguments must also be placed in correspondence according to role). People also prefer deeper mappings over shallow mappings (the *systematicity* principle)—that is, they prefer mappings in which the commonalities form a system connected by higher order constraining relations,[2] such as causality or symmetry (Clement & Gentner, 1991). This reflects a tacit preference for coherence and inferential power, and it means that the process of aligning two situations often leads to a focus on informative relational systems such as causal systems.

When structural alignment is achieved, the common structure becomes more salient and may be retained and reused. It will nearly always be more abstract than either of the analogs, since it is (in structural terms) the intersection. In addition to forming abstractions, achieving a structural alignment can lead to new learning in at least three other ways: potential *candidate inferences* are projected from the base to the target; *alignable differences*—differences that play the same role in the corresponding structures—will often pop out; and there may be *re-representation* of one or both analogs to improve the match.

2. Analogical comparison supports relational abstraction

There is abundant evidence that comparison can lead to relational abstraction in adults. Studies of problem solving and transfer show that comparing two analogous scenarios prompts people to form a more abstract schema, increasing the likelihood that the common principle will be transferred to a future item (relative to seeing just one exemplar, or even the same two items without encouragement to compare) (Catrambone & Holyoak, 1989; Gentner, Loewenstein, Thompson, & Forbus, 2009; Gick & Holyoak, 1983). For example, Loewenstein, Thompson, and Gentner (1999) gave business school students two analogous scenarios which differed in their surface content, but expressed the same negotiation strategy (e.g., a tradeoff strategy). One group was asked to write out commonalities between the two scenarios. The other group was given the same two scenarios sequentially and asked to give advice to the protagonists or to state the negotiation principle illustrated by the scenario. Those who compared the scenarios were far more likely to transfer their common strategy to an analogous test case than were those who processed the two scenarios separately. Examination of their written materials showed that participants who compared the two scenarios were far more likely to state an abstract schema expressing the negotiation principle than were those who wrote about each scenario separately.

Strikingly, students who processed the scenarios sequentially rarely commented on their parallels, even though the scenarios were presented in immediate succession. This illustrates two important points. First, even smart, motivated adults often miss potential comparisons that, if noticed, could lead to greater insight (and in this case, greater success in a competitive situation). Second, what people learn from a set of examples depends not just on properties of the distribution, but on how they *process* the examples.

Similar processes of alignment and abstraction occur in children. There is considerable evidence that comparison can prompt children to notice abstract conceptual patterns (Chen & Klahr, 1999; Kotovsky & Gentner, 1996). For example, Christie and Gentner (2010) taught 3- and 4-year-old children names for novel spatial relational configurations and asked them to extend the name to another instance—either one with a matching object (but a different relational pattern) or one with the same relational pattern, but new

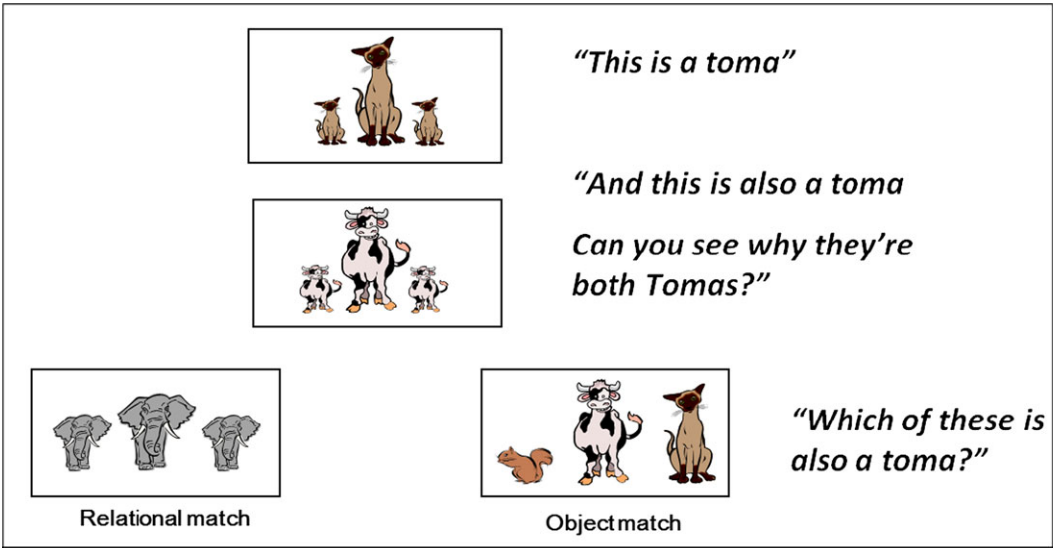

Fig. 1. Materials from Christie and Gentner (2010). When 3-year-olds were given either one of the standard "tomas" above and asked to choose which alternative was also a toma, 98% of the responses were to the Object match (and only 2% to the Relational match). However, when given both standards and encouraged to compare them (as shown here), 57% of the responses were to the Relational match.

objects. (See Fig. 1.) Children given one standard chose on the basis of matching objects, disregarding the relational configuration.[3] Those who saw two analogous standards and were encouraged to compare them were far more likely to choose the alternative with the same relational configuration.

2.1. An alternate interpretation

We have suggested that the process of comparison revealed a relational interpretation to the children. However, an alternate explanation is based on hypothesis search and selection. In the study above, the comparison group saw two standards, whereas the control groups saw only one. It could be that children formed a set of potential hypotheses when they saw the first standard, with object-based hypotheses having high prior likelihood. For example, when given the top standard in Fig. 1, the hypotheses might be

$$H1 : \text{toma} = \text{cats}$$

$$H2 : \text{toma} = \text{symmetric pattern}$$

When the next standard (with cows) is shown, children in the comparison group could reject H1 and accept H2. But the two solo groups are not given further evidence, and so retain their object-based H1. On this account, what differentiates the two groups is the set of exemplars they were given—not whether they engaged in comparison.

To test this possibility, Christie and Gentner (2010) conducted a further study in which children received the same two standards as the comparison group in Experiment 1 but were not encouraged to compare them. As in Experiment 1, the two standards were shown in immediate sequence, and each was labeled a *toma*. However, to discourage comparison, the first standard was removed before placing the second one on the table, and the invitation to compare (Can you see why these are both tomas?) was omitted. Children in this condition made object-based choices; they did not differ statistically from the solo groups in Experiment 1. We conclude that receiving the two labeled standards was not enough—the process of comparison was essential to arriving at the relational abstraction. As in the negotiation study with adults just above, knowing the distribution of examples is not enough to predict what hypotheses the learner will form. The moment-to-moment learning processes are critical in determining what is learned. In particular, online analogical comparison processes can generate new abstractions.

3. What prompts comparison?

In the studies just described, relational abstraction occurred only when participants were explicitly invited to compare the examples. This served to demonstrate the importance of analogical comparison processes—but it also raises a new challenge. If analogical comparison required a helpful tutor to encourage comparison, its value in accounting for children's gains would be seriously limited. To make the case that analogical processing is a major source of new abstractions, we need an account of what prompts children to compare things spontaneously in everyday life.

Based on prior research and theory, we know of factors that promote comparison: (a) spatiotemporal proximity, (b) high similarity, and (c) common labels (as discussed below). Even for adults, relational mapping is facilitated by simultaneous presentation; and relational reminding and transfer is far more likely when the pair is of high overall similarity (Gentner, Rattermann, & Forbus, 1993; Holyoak & Koh, 1987; Ross, 1987; Trench & Minervino, 2015). For children, there is the added issue that high-similarity comparisons are not only readily noticed, but also easy to align. In structural alignment, the corresponding objects do not have to be similar—but the alignment process is easier if they are. The easiest, most natural form of similarity to process is overall similarity, in which the object matches support the relational alignment (Gentner & Kurtz, 2006; Gentner & Toupin, 1986). Overall similarity is especially important for young children. Early in development, children's representations tend to be rich in information about individual objects and sparse in relational information. This means that young children are more likely to succeed in structural alignment when the objects match as well as the relations.

There is abundant evidence that early in learning, successful comparison depends on a high degree of overall similarity (Brown & Kane, 1988; DeLoache, 1989; Gentner & Rattermann, 1991; Gentner & Toupin, 1986; Kotovsky & Gentner, 1996; Smith, 1989). For example, Chen, Sanchez, and Campbell (1997) examined 10- and 13-month-olds' ability to combine steps (pulling a blanket and pulling a string) to obtain a toy. Although

13-month-olds could transfer a modeled solution strategy across isomorphic (but perceptually dissimilar) problems, 10-month-olds were dependent on high perceptual similarity to transfer the solution.

The same pattern of early reliance on close similarity occurs for older children in a conceptual task in which the key relations are social relations. For example, 6-year-old children are better able to retell stories with new characters if the characters involved in the retelling look similar and play the same roles as those in the original story. By 9 years of age, children are able to overcome this reliance on similar object matches. They can respond simply on the basis of the relational match, especially if there is a higher order plot structure to guide their mapping (Gentner & Toupin, 1986).

Fortunately, children are not entirely at the mercy of perceptual similarity. The "invitation to compare" can also come via language. When children hear a common label applied to two things, it signals that the two share some important commonality and leads them to compare the two (Gentner, 2010, 2016; Gentner & Namy, 1999). For example, in one study, children were shown two examples of a relational category (such as *cutter-of*) and were asked to generalize the common relation to a new item. Both 4- and 5.5-year-olds performed better when they were given a common relational term for the relation (e.g., "The knife is the blick for the watermelon, and the ax is the blick for the tree") than when they were simply told that the two situations were alike (e.g., "The knife goes with the watermelon, and the ax goes with the tree in the same way") (Gentner, Anggoro, & Klibanoff, 2011).

Common labels provide an invitation to compare, and thus they can compensate to some degree for lack of overall similarity. But in order to profit from the comparison, the child must be able to align the pair, and this is easiest for high-similarity matches. Thus, in early learning, structural alignment is both most likely to be *initiated* for highly similar pairs, and most likely to *succeed* for highly similarity pairs. This has important implications for the course of learning. Because high-similarity pairs (by their nature) share concrete as well as abstract commonalities, comparing them will yield a relatively concrete abstraction. This leads to a pattern of conservative learning, in which children initially tend to form fairly concrete abstractions, which become more abstract with further learning.

This pattern of early conservative learning can be seen in children's number learning. For instance, Mix (1999) examined preschoolers' ability to judge numerical similarity across different sets of items. Children were far more accurate when given highly similar sets (e.g., a set of three dots and a set of three discs) than when given low-similarity sets (e.g., a set of three dots and a set consisting of three different familiar objects). A similar pattern appears in studies by Huang, Spelke, and Snedeker (2010), in which *two-knowers* (children who knew the meanings of *one* and *two*[4]) were trained to learn *three*. In one study (Experiment 3), children were trained using only one animal category—for example, with repeated trials showing three dogs and labeling it ("This card has three dogs!"). When their understanding of the newly trained word was assessed, the results showed that they could apply the new number (*three*) successfully, but only to the category on which they had been taught. That is, they could apply *three* to three new instances of *dogs*, but not to three *cows*.

A natural question is whether the children could have learned the more abstract category if given a broader training set. Huang et al. (Experiment 1) tested this by training children using eight different categories of animals. As expected, three-knowers who received the broad training experience (aimed at teaching *four*) were able to generalize *four* to new animals (Experiment 1) and even to artifacts (Experiment 3). However, two-knowers given broad training on *three* failed to generalize *three*, even to novel animal categories.

Why were the two-knowers able to benefit from the narrow training set but not the broad set? We suggest that these children, who were only beginning to grasp relations among numbers, could successfully align the high-similarity pairs—for example, three spaniels followed by three retrievers. Thus, when given the narrow (all dogs) set, they achieved a rather concrete abstraction—that of *three dogs*. But when they received, say, three horses followed by three cats (in the broad training set), they failed to achieve a structural alignment—possibly because they simply saw no basis for comparison.

This study illustrates a recurring pattern: (a) early in domain understanding, children require overall similarity to succeed at structural alignment; (b) therefore, early generalizations are highly concrete; (c) with increasing understanding of the relations in the domain, learners can align surface-dissimilar pairs, as long as they match relationally—and this results in more abstract generalizations. Thus, there is a "rich get richer' pattern in children's learning: The better a child understands the relational patterns in a domain, the more that child can learn from receiving new examples in the domain.

This study also illustrates a more general point. A sturdy rule of thumb in studies of learning is that breadth of training predicts breadth of transfer. In analogical learning, this translates to the statement that for maximum generalization of a relational pattern, one should compare examples that are as different as possible in their surface content. This point is clear on both theoretical and empirical grounds (e.g., Barnett & Ceci, 2002; Halpern, Hansen, & Reifer, 1990). But to achieve a *relational* generalization, the learner must be able to align the examples. If the learner cannot align the examples, the training is useless or confusing. For relational learning, the rule is that breadth of *alignable* examples predicts breadth of transfer.

These conclusions might seem cause for pessimism. If early learners can only benefit from close comparisons, and therefore can only derive highly concrete generalizations, how can analogical comparison account for children's rapid gains in abstract knowledge? This takes us to the idea of *progressive alignment*. We now discuss evidence that although these early abstractions are limited in scope, they provide seeds for further abstraction.

4. Progressive alignment

Although early learners require concrete similarity to succeed at structural alignment, those comparisons can still act to advance children's understanding, via progressive alignment. *Progressive alignment* refers to the phenomenon whereby forming a relatively concrete abstraction increases the likelihood of making a match with a less similar example that shares the same relational structure. There is considerable evidence for the dual claims of conservative early learning and progressive alignment. For example, Kotovsky and Gentner (1996) showed 4-year-old children arrays illustrating abstract commonalities (e.g., *monotonic increase*) and asked them to choose a matching array. The standard and correct choice arrays either shared lower order commonalities and a common dimension (e.g., circles increasing in size: ●●●, and triangles increasing in size: ▲▲▲), making them easy to align, or the standard and choice arrays lacked this lower order support and crossed dimensional boundaries (e.g., circles increasing in size: ●●●, and circles increasing in color saturation: ○●●). When children saw trials in random order, they were successful only on the easier, within-dimension comparisons. However, when the trials were blocked so that children saw all the close, same-dimension trials before the cross-dimension trials, they were better able to select the choice with the same abstract relational structure in the subsequent cross-dimensional trials. Importantly, children received no feedback at any point. Yet carrying out close comparisons—on which they succeeded in choosing the relational match—increased the likelihood of success on subsequent more abstract relational matches.

This study exemplifies the key generalization of progressive alignment: Concrete alignments potentiate more abstract alignments and abstractions (for more detailed descriptions, see Gentner, Rattermann, Markman, & Kotovsky, 1995; Kandaswamy, Forbus, & Gentner, 2014, for a simulation). This study also illustrates our earlier point that early in learning, relational comparison is highly dependent on close spatiotemporal juxtaposition. When the within-dimension trials were presented in close succession, children could compare and learn from them, as evidenced by their improved performance on the subsequent cross-dimension trials. But when the same within-dimension trials were mixed with cross-dimension trials, no such benefit occurred.

This pattern of progressive alignment also shows up in language learning—for example, in learning verbs (Childers et al., 2016; Haryu, Imai, & Okada, 2011), part names (Gentner, Loewenstein, & Hung, 2007), and relational nouns (Gentner et al., 2011). For example, Gentner et al. (2007) taught new part terms to preschool children. Children were shown a novel "Martian creature" and were told that it had "a dibble." Then they were asked to choose which of two other creatures also had a dibble. Four- and 5-year-olds were successful even when the alternatives did not resemble the standard; but 3-year-olds succeeded only when the two alternatives were highly similar to the standard and were therefore easily aligned with it. However, when 3-year-olds were first given high-similarity trials (Experiment 3), and then went to receive low-similarity trials, they succeeded on the low-similarity trials (See Fig. 2) We suggest that aligning the high-similarity items led them to extract the common structure, and this helped them see the match with the low-similarity items.

Haryu et al. (2011) found that progressive alignment could aid children in verb learning. They taught 4-year-old children a verb for a novel action on a particular object and asked whether the children could apply the verb to the same action with a new object. The children extended the verb only when the new object was highly similar to the initial

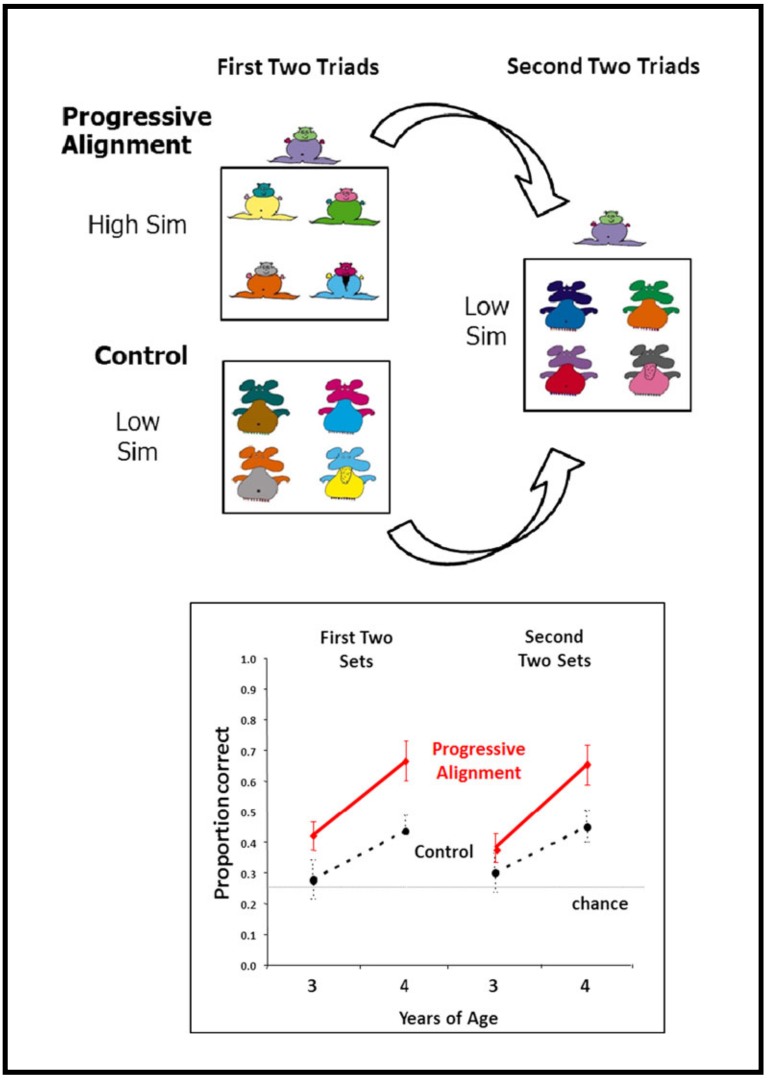

Fig. 2. Results of Experiment 3 (Gentner et al., 2007). The left half of the graph shows that 3-year-olds who received high-similarity (and therefore readily alignable) triads did better at identifying the creature with the same body part than did those who received low-similarity triads. The right half of the graph shows the effect of progressive alignment: When both groups went on to receive low-similarity triads, children who had initially received high-similarity triads did far better than those who had received low-similarity triads.

object; they failed when the object was dissimilar, even though the action remained constant. Thus, they showed conservative generalization, restricted to highly similar events. A second experiment showed that the children were able to benefit from progressive alignment. If they were first asked to extend the verb to an event with a high-similarity object (at which most children succeeded), they then went on to succeed with a low-similarity object. (In contrast, as in the part-learning study described above, children who received an equal number of low-similarity trials remained at chance.)

In the three studies just reviewed, children initially extended the new terms only to closely similar situations—fitting the pattern of conservative learning that has been found for verbs and other relational terms (Forbes & Farrar, 1993; Gentner, 2005; Huttenlocher, Smiley, & Charney, 1983; Imai, Haryu, & Okada, 2005; Tomasello, Akhtar, Dodson, & Rekau, 1997). Further, a salient point about these three studies is that children received no feedback as to their correctness on any trial—yet they performed better on difficult low-similar trials after receiving easily alignable high-similar trials (Gentner et al., 2007; Haryu et al., 2011; Kotovsky & Gentner, 1996). This means that the benefit they derived from successfully aligning the high-similarity matches was internally generated.

The implications of this pattern are important. First, progressive alignment emerges as a natural candidate for experiential learning. It does not require a caregiver to instruct children about word meanings, nor even to correct their errors. Indeed, error-correction is not required for progressive alignment.[5] Second, learning by progressive alignment will typically go unnoticed. A parent would simply observe their child using a word—say, a particular verb—more and more broadly. We suggest that this kind of invisible learning is a major driver of experiential learning, and of language learning in particular. We now turn to a further aspect of language learning—grammar learning.

5. Analogical processes in grammar learning

Learning the grammar of one's native language is one of the great achievements of childhood. The complexity of this task has led many theorists to the view that this learning must be based on a system of innate knowledge and/or a dedicated language acquisition device. Recently, however, a number of researchers have explored the role of general learning processes in grammar learning (Aslin & Newport, 2012; Chang, Dell, & Bock, 2006; Goldberg, 2006; Tomasello, 2000). In this spirit, we propose analogical learning as an important contributor to the acquisition of grammar (see Goldwater (this issue) for evidence and a more comprehensive review).

Over the last few decades, landmark research with artificial grammars has demonstrated that 7-month-old infants possess two key abilities: (a) the ability to learn statistical regularities within a particular input sequence (Saffran, Aslin, & Newport, 1996); and (b) the ability to abstract across sequences that have parallel structure (Marcus, Vijayan, Rao, & Vishton, 1999). This second ability, we suggest, is analogical processing. For example, in Marcus et al.'s (1999) study, after hearing repeated examples of a syllable pattern such as AAB, 7.5-month-olds were able to discriminate new instances of the AAB pattern from instances of an ABA pattern, even when all the specific syllables were new (see also Gomez & Gerken, 2000). Consistent with an analogical learning account, the ability to generalize across such patterns is not restricted to language-like materials; it can operate across a broad range of stimuli, including tones and visual stimuli (Marcus et al., 1999; Saffran, Pollak, Seibel, & Shkolnik, 2007).

Further support for the role of structure-mapping in grammar learning comes from a simulation of the Marcus et al. (1999) findings using a model of analogical generalization (Kuehne, Gentner, & Forbus, 2000). In this model, SEQL[6] (an offshoot of SME, the structure-mapping engine) was able to generate progressively more abstract representations via sequential analogical comparison and abstraction. The model was given the same input as the infants in Marcus et al.'s study: three repetitions of 16 ABA items (e.g., *ledile, widiwi, jileji).* Each syllable was encoded as 12 phonetic features, and symmetry and repetition within each string were encoded by Magi, which uses SME within patterns (Ferguson, 1994). When a new string was presented, SEQL compared it to the prior exemplar and stored the common structure. This resulted in an ongoing generalization. Each new string was compared to this generalization, and the common structure was retained. This resulted in incrementally stripping away nonshared phonetic features and retaining the common abstract structure. After receiving all the strings, SEQL was given two novel test strings with new syllables—one with the same ABA structure (e.g., *kogako*) and one with a different ABB structure (*kogaga*). The system showed two signs of having formed an abstract ABA pattern. First, it found the novel within-grammar (ABA) string more similar to its ongoing generalization than it did the novel out-of-grammar (ABB) string. Second, it made incorrect candidate inferences to the out-of-grammar string. For example, when given *kogaga*, it proposed that the third syllable should be *ko*—an inference that was quickly disconfirmed. One might think of this behavior as analogous to infants having their expectations disconfirmed.

Interestingly, this rule-like behavior emerged despite the fact that no purely abstract rules were derived. In each of the 20 runs of the simulation (with randomized item orders, and even with added noise), a few concrete features were still present at the end. Thus, it was not a fully abstract representation. Yet, in each case, the simulation responded as though it had an abstract rule. Analogical generalization can lead to "just-abstract-enough" abstractions—representations that are sufficiently abstract to behave like rules, even if they are not fully abstract. More generally, analogical generalization suggests a continuum of abstractness as learners gain knowledge in a domain.

We have stressed local analogical processes in artificial grammar learning. This approach contrasts with a highly prominent approach—the rational inference account of this learning. Across many cleverly designed studies, Gerken and colleagues have tested this account. For example, Gerken (2006) examined 9-month-olds' ability to extract abstract rules from an artificial grammar. In this study, infants listened to 2 minutes of four unique three-syllable strings. All four strings had an AAB pattern: *leledi, wiwidi, jijidi, dededi.* Gerken noted that these strings can support two types of abstractions: a broad generalization to the AAB pattern or a narrow generalization to the AA*di* pattern (because all the strings end with *di*) (Gerken, 2006; Xu & Tenenbaum, 2007). Gerken found that infants learned the narrow generalization. She argued that this was consistent with a Bayesian account of generalization in which infants consider both hypotheses, but select the narrow generalization because the four strings ending in *di* constitute a *suspicious coincidence* (Xu & Tenenbaum, 2007). (The idea is that if the broader generalization [AAB] were the rule, then it is unlikely that all the strings would belong to a narrower category.) In subsequent work, Gerken (2010) showed that when infants heard three additional examples that did not end in *di* (*wiwije, dedewe, jijili)*, they then selected the broader AAB rule, which fit the overall distribution.

We propose that analogical comparison can explain these patterns of generalizations quite naturally. According to the analogical account, the infants who heard only *leledi, wiwidi,* and so on would carry out comparisons across the items, resulting in the abstraction that preserves their commonalities—namely, AA*di.* Instead of considering two (or more) possible abstractions or hypotheses, there would be one natural abstraction. In Gerken's (2010) subsequent study, when examples with different endings followed later in the stream, the subsequent comparisons would yield a more abstract common structure, shorn of the nonshared *di* syllable. The diverse examples would thus give rise to the AAB abstraction—and indeed, the initial AA*di* pattern might serve as a seed for progressive alignment.

6. Higher order abstractions

One important kind of abstraction is higher order abstractions that constrain further more specific generalizations. This idea has been prominent in Bayesian accounts of learning (e.g., Kemp, Perfors, & Tenenbaum, 2007; Tenenbaum & Griffiths, 2001), under the term *overhypothesis,* a term taken from Nelson Goodman (1955). Goodman (1955; pp. 109–111) illustrated the idea of overhypotheses with the following example. Imagine there are several bags, each with marbles inside. A few red marbles are drawn out from the first bag, then a few blue marbles are drawn out from the second bag, and then a few green marbles are drawn out from the third bag. From the fourth bag, a single yellow marble is drawn out. Does the fourth bag therefore contain all yellow marbles? Many would infer that it does. The reason we can infer that the fourth bag has all yellow marbles is that we have formed a second-order abstraction, or overhypothesis, that the marbles in each bag are uniform in color.

In an elegant study, Dewar and Xu (2010) showed that 9-month-old infants could form this kind of higher order abstraction. In their study, designed along the lines of Goodman's marble problem, infants saw four boxes on a stage (see Fig. 3) and watched as the experimenter withdrew shapes from each box. Four circles were drawn from the first box and placed in front of the box. Then four squares were drawn from the second box and placed in view. Then four triangles were drawn from the third box and placed in view. At the fourth box, two possible events occurred: The experimenter drew out either two multicolored stars (the expected outcome) or a star and a circle (the unexpected outcome). Infants looked significantly longer at the unexpected outcome. Dewar and Xu concluded that the infants had formed an overhypothesis, based on the evidence from the first three boxes—that each box contained items that shared the same shape.

Building on Dewar and Xu's findings, we suggest that an analogical account of learning in this study would go as follows (see Fig. 3). Infants first compare the items from Box 1 and see that they are all circles. (The items are placed in a line in front of the box, and this high spatial juxtaposition invites comparison.) Likewise, they see that the

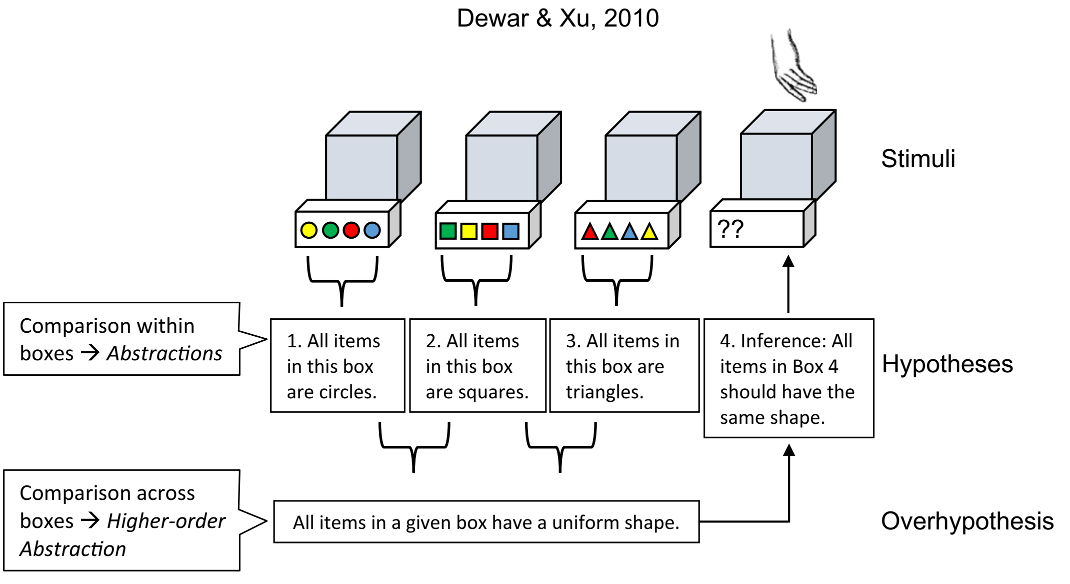

Fig. 3. Schematic of stimuli and procedure from Dewar and Xu (2010). We depict how comparison within boxes can give rise to an analogical abstraction for each box, and how comparison across the first three boxes can generate a higher level abstraction (the overhypothesis), which generates a candidate inference that the fourth box also contains uniform shapes.

items from Box 2 are all squares, and that the items from Box 3 are all triangles. These line of shapes are all visible at the same time, making it easy to compare them. By comparing the arrays, infants notice the commonality that *each box contains items of uniform shape*. This is the key higher order abstraction (or overhypothesis). When items of two different shapes are drawn from Box 4, this abstraction is violated—so infants look longer at this case.

This account generates testable predictions—notably, that if comparison processing is made difficult, learners will be less likely to arrive at the relational commonalities. Thus, if this method were altered to make comparison less likely—for example, if the objects from a given box were removed after presentation, rather than being arrayed together simultaneously—infants should be less likely to form the higher order abstraction[7] (see Christie & Gentner, 2010; Spencer, Perone, Smith, & Samuelson, 2011).

We suggest that analogical abstraction plays a key role in forming new hypotheses, including overhypotheses. If so, then analogical processes are a necessary adjunct to hypothesis-testing accounts. For example, Bayesian models can capture hypothesis selection and belief revision; but unless one adopts a radical nativist position that all hypotheses are innate, additional processes are needed to account for the formation of hypotheses. As Xu (2016) points out, processes of analogical comparison (along with self-explanation (Fonseca & Chi, 2011; Lombrozo, 2010) and other “learning by thinking” processes) can fill this gap.

7. Summary

Our goal here has been to make the case for a set of interrelated claims—to present evidence for these claims, to confront some challenges and limitations, and to lay out implications of those claims for cognitive development. Our central claim is that analogical comparison is a major driver of children's early cognitive development—and specifically, that it is the main driver of new relational abstractions.

This immediately leads to a conundrum. In analogical abstraction,

1. The more different the two comparands, the more abstract the generalization and the greater the scope of transfer
2. But, the more different the two comparands, (a) the less likely children are to spontaneously compare them; and (b) the less likely children [and other novices] are to be able to align them.

The implication for development is that children's early comparisons will be between highly similar situations, resulting in relatively concrete, context-specific abstractions. How, then, can such a process account for children's rapid early learning? Our response to this challenge is, first, that these early conservative abstractions potentiate further more distant comparisons, via progressive alignment. Thus, they form seeds for increasingly abstract generalizations. Second, the ambient language serves as an invitation to other, less obvious comparisons, via the use of common labels.

A further implication of this line of reasoning is that a key predictor of cognitive gains is the likelihood of making comparisons. And because hearing common labels is a key invitation to compare, children who are exposed to rich linguistic input have a royal road to the culture's standard abstractions. Further, the advantage of language input will be especially pronounced for relational categories, for which members lack obvious similarities. This is one reason that the "word gap" between high-SES and low-SES families matters for conceptual development.

Our claim that analogical processes are the key to relational learning also has implications for theories of learning. A standard rule of thumb is that

Breadth of training predicts breadth of transfer.

But when it comes to relational learning, we need to revise the standard rule to be

For relational learning, breadth of alignable training predicts breadth of transfer.

7.1. How learning changes

In predicting the developmental course of learning, we have emphasized factors such as spatiotemporal juxtaposition and high overall similarity. These are critical in determining whether early learners will initiate comparison, and whether they will arrive at a structural alignment. But these factors are of course far less important later in learning. Adults readily carry out analogies between surface-dissimilar pairs and can draw abstractions over widely spaced exemplars. Further, once learners have acquired multiple abstract generalizations, hypothesis selection becomes more prominent relative to hypothesis generation. Relatedly, learning by analogical projection from familiar base domains becomes more prominent relative to learning by mutual alignment.

7.2. Concluding remarks

There is a long-standing tradition in the philosophical and developmental literature that conceives of similarity as a deceiver (Goodman, 1972), or at best a distraction from the discovery of deep conceptual laws (Piaget, 1954). Indeed, the early beginnings of similarity processing may not inspire confidence. As Quine (1969) puts it, there is little reason to think that "the muddy old notion of similarity" (Quine, 1969, p. 172) has anything to contribute to the development of abstract capacities. But Quine goes on to propose that there are more sophisticated kinds of similarity, more in tune with the regularities of nature. "In this career of the similarity notion, starting in its innate phase, developing over the years in the light of accumulated experience, passing then from the intuitive phase into theoretical similarity, and finally disappearing altogether, we have a paradigm of the evolution of unreason into science" (Quine, 1969, p. 176). Structure-mapping, we suggest, is just the kind of process needed to accomplish this career.

Acknowledgments

This work was supported by the Office of Naval Research, ONR grant N00014-13-1-0470.

Notes

1. *Mutual alignment* analogies are those that are processed by structural alignment with few or no projected candidate inferences, According to structure-mapping, all analogies involve structural alignment. Even when the base offers a clear projection (as in the *goldmine* example), one still must align that structure with the target's representation in order to comprehend it. For example, if you heard, "That fork is a goldmine," you would probably be baffled, because it is unclear what part of the representation of *fork* could serve as the source of something valuable. Indeed, Wolff and Gentner (2011) found evidence for early alignment processing in metaphor comprehension, even for highly directional conventional metaphors. Thus, the difference here is between structural alignment only, and structural alignment plus projection.

2. The higher order relations must be *constraining* relations, such that changing one of the arguments will change the other. Thus, conjunctive relations such as AND (Event 1, Event 2), which do not constrain their arguments, do not contribute to the systematicity of a set of relations. There is also evidence that the *degree* of constraint matters: for example, that causal relations are more powerful than temporal relations (Lassaline, 1996).
3. A pattern of early focus on object matches followed by a later focus on relational matches has been widely found in children's similarity matching, and it is referred to as the *relational shift* (Gentner, 1988; Gentner & Toupin, 1986; Richland, Morrison, & Holyoak, 2006). Because this shift—or more accurately, shifts—occurs at different times in different domains, depending on domain learning, we maintain that it is largely driven by gains in domain knowledge (Gentner & Rattermann, 1991). However, there is evidence that growth of executive ability (Richland et al., 2006; Thibaut, French, & Vezneva, 2010) and/or processing capacity (Halford, 1992) also play a role.
4. The "give-n" task is used to assess children's number knowledge. A two-knower can respond correctly if given a set of objects and asked to "give me one" or "give me two," but will typically respond with a random handful when asked to give a number larger than two. Likewise, a three-knower is accurate up to three.
5. Of course, comparison processes are only part of the story. For example, there is evidence that children do benefit from adult responses to their utterances (Chouinard & Clark, 2003). Our point is simply that significant learning can occur without this help, as long as children are exposed to sufficient language.
6. The current version of our analogical generalization model, SAGE (Kandaswamy et al., 2014) incorporates some improvements, but utilizes a similar incremental generalization process.
7. We note that in Kemp et al.'s (2007) Hierarchical Bayesian Model (HBM), overhypotheses are seen as "constraints on the hypotheses considered by the learner" (p. 307), and are viewed as prior to the lower level hypotheses, rather than as arising from them. However, as Kemp et al. make clear, HBM is an information-level account (Marr, 1982), and does not address *how* overhypotheses are acquired. This kind of question must be asked at Marr's algorithmic level, where processes of analogy and explanation reside.

References

Aslin, R. N., & Newport, E. L. (2012). Statistical learning from acquiring specific items to forming general rules. *Current Directions in Psychological Science*, *21*(3), 170–176.

Asmuth, J., & Gentner, D. (2017). Relational categories are more mutable than entity categories. *Quarterly Journal of Experimental Psychology*, *70*(10), 2007–2025.

Barnett, S. M., & Ceci, S. J. (2002). When and where do we apply what we learn?: A taxonomy for far transfer. *Psychological Bulletin*, *128*(4), 612.

Boroditsky, L. (2000). Metaphoric structuring: Understanding time through spatial metaphors. *Cognition*, *75*(1), 1–27.

Bowdle, B., & Gentner, D. (2005). The career of metaphor. *Psychological Review*, *112*(1), 193–216.

Brown, A. L., & Kane, M. J. (1988). Preschool children can learn to transfer: Learning to learn and learning from example. *Cognitive Psychology*, *20*(4), 493–523.

Burgoon, E. M., Henderson, M. D., & Markman, A. B. (2013). There are many ways to see the forest for the trees: A tour guide for abstraction. *Perspectives on Psychological Science*, *8*(5), 501–520.

Catrambone, R., & Holyoak, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(6), 1147.

Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review*, *113*(2), 234.

Chen, Z., & Klahr, D. (1999). All other things being equal: Acquisition and transfer of the control of variables strategy. *Child Development*, *70*(5), 1098–1120.

Chen, Z., Sanchez, R. P., & Campbell, T. (1997). From beyond to within their grasp: The rudiments of analogical problem solving in 10- and 13-month olds. *Developmental Psychology*, *33*(5), 790–801.

Childers, J. B., Parrish, R., Olson, C. V., Burch, C., Fung, G., & McIntyre, K. P. (2016). Early verb learning: How do children learn how to compare events? *Journal of Cognition and Development*, *17*(1), 41–66.

Chouinard, M. M., & Clark, E. V. (2003). Adult reformulations of child errors as negative evidence. *Journal of Child Language*, *30*(03), 637–669.

Christie, S., & Gentner, D. (2010). Where hypotheses come from: Learning new relations by structural alignment. *Journal of Cognition and Development*, *11*(3), 356–373.

Clement, C. A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, *15*, 89–132.

Csibra, G. (2008). Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition*, *107*(2), 705–717.

DeLoache, J. S. (1989). Young children's understanding of the correspondence between a scale model and a larger space. *Cognitive Development*, *4*(2), 121–139.

Dewar, K. M., & Xu, F. (2010). Induction, overhypothesis, and the origin of abstract knowledge: Evidence from 9-month-old infants. *Psychological Science*, *21*, 1871–1877.

Doumas, L. A. A., Hummel, J. E., & Sandhofer, C. M. (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.

Ferguson, R. W. (1994). MAGI: Analogy-based encoding using symmetry and regularity. In A. Ram & K. Eiselt (Eds.), *Proceedings of the sixteenth annual conference of the cognitive science society* (pp 283–288). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Fonseca, B., & Chi, M. T. H. (2011). The self-explanation effect: A constructive learning activity. In R. E. Mayer, P. Alexander (Eds.), *The handbook of research on learning and instruction*, (pp. 296–321). New York: Routledge Press.

Forbes, J. N., & Farrar, M. J. (1993). Children's initial assumptions about the meaning of novel motion verbs: Biased and conservative? *Cognitive Development*, *8*, 273–290.

Forbus, K. D., Ferguson, R. W., Lovett, A., & Gentner, D. (2016). Extending SME to Handle Large-Scale Cognitive Modeling. *Cognitive Science*. doi:10.1111/cogs.12377

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. (2005). The development of relational category knowledge. In L. Gershkoff-Stowe, & D. H. Rakison (Eds.), *Building object categories in developmental time* (pp. 245–275). Hillsdale, NJ: Erlbaum.

Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, *34*(5), 752–775.

Gentner, D. (2016). Language as cognitive tool kit: How language supports relational thought. *American Psychologist.*, *71*(8), 650–657.
Gentner, D., Anggoro, F. K., & Klibanoff, R. S. (2011). Structure-mapping and relational language support children's learning of relational categories. *Child Development*, *82*(4), 1173–1188.
Gentner, D., & Forbus, K. (2011). Computational models of analogy. *WIREs Cognitive Science*, *2*, 266–276.
Gentner, D., & Kurtz, K. (2005). Relational categories. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. W. Wolff (Eds.), *Categorization inside and outside the lab* (pp. 151–175). Washington, DC: APA.
Gentner, D., & Kurtz, K. (2006). Relations, objects, and the composition of analogies. *Cognitive Science*, *30*, 609–642.
Gentner, D., Loewenstein, J., & Hung, B. (2007). Comparison facilitates children's learning of names for parts. *Journal of Cognition and Development*, *8*, 285–307.
Gentner, D., Loewenstein, J., Thompson, L., & Forbus, K. (2009). Reviving inert knowledge: Analogical abstraction supports relational retrieval of past events. *Cognitive Science*, *3*, 1343–1382.
Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, *52*, 45–56.
Gentner, D., & Medina, J. (1998). Similarity and the development of rules. *Cognition*, *65*, 263–297.
Gentner, D., & Namy, L. (1999). Comparison in the development of categories. *Cognitive Development*, *14*, 487–513.
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., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, *25*, 524–575.
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: Erlbaum.
Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, *10*, 277–300.
Gerken, L. A. (2006). Decisions, decisions: Infant language learning when multiple generalizations are possible. *Cognition*, *98*, B67–B74.
Gerken, L. A. (2010). Infants use rational decision criteria for choosing among models of their input. *Cognition*, *115*(2), 362–366.
Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15* (1), 1–38.
Glucksberg, S., & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*, *97*(1), 3.
Goldberg, A. E. (2006). *Constructions at work: The nature of generalization in language*. Oxford: Oxford University Press on Demand.
Goldwater, M. B. (2017). Grammatical constructions as relational categories. Topics in Cognitive Science.
Goldwater, M. B., & Gentner, D. (2015). On the acquisition of abstract knowledge: Structural alignment and explication in learning causal system categories. *Cognition*, *137*, 137–153.
Goldwater, M. B., & Markman, A. B. (2011). Categorizing entities by common role. *Psychonomic Bulletin & Review*, *18*(2), 406–413.
Goldwater, M. B., & Schalk, L. (2016). Relational categories as a bridge between cognitive and educational research. *Psychological Bulletin*, *142*(7), 729–775.
Gomez, R. L., & Gerken, L. (2000). Infant artificial language learning and language acquisition. *Trends in Cognitive Sciences*, *4*(5), 178–186.
Goodman, N. (1955). *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.

Goodman, N. (1972). *Problems and projects*. Indianapolis, IN: Bobbs-Merrill.

Halford, G. S. (1992). Analogical reasoning and conceptual complexity in cognitive development. *Human Development*, *35*(4), 193–218.

Halpern, D., Hansen, C., & Reifer, D. (1990). Analogies as an aid to understanding and memory. *Journal of Educational Psychology*, *82*, 198–305.

Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, *450*(7169), 557–559.

Haryu, E., Imai, M., & Okada, H. (2011). Object similarity bootstraps young children to action-based verb extension. *Child Development*, *82*(2), 674–686.

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

Huang, Y. T., Spelke, E., & Snedeker, J. (2010). When is four far more than three? Children's generalization of newly acquired number words. *Psychological Science*, *21*(4), 600–606.

Hummel, J. E. (2011). Getting symbols out of a neural architecture. *Connection Science*, *23*, 109–118.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review.*, *110*(3), 220–264.

Huttenlocher, J., Smiley, P., & Charney, R. (1983). Emergence of action categories in the child: Evidence from verb meanings. *Psychological Review*, *90*, 72–93. https://doi.org/10.1037/0033-295X.90.1.72.

Imai, M., Haryu, E., & Okada, H. (2005). Mapping novel nouns and verbs onto dynamic action events: Are verb meanings easier to learn than noun meanings for Japanese children? *Child Development*, *76*, 340–355. https://doi.org/10.1111/j.1467-8624.2005.00849.x.

Kandaswamy, S., Forbus, K., & Gentner, D. (2014). Modeling learning via progressive alignment using interim generalizations. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th annual conference of the cognitive science society* (pp. 2471–2476). Austin, TX: Cognitive Science Society.

Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical Bayesian models. *Developmental Science*, *10*(3), 307–321.

Kokinov, B., & French, R. M. (2003). Computational models of analogy-making. *Encyclopedia of Cognitive Science*, *1*, 113–118.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, *67*(6), 2797–2822.

Krawczyk, D., Holyoak, K., & Hummel, J. (2004). Structural constraints and object similarity in analogical mapping and inference. *Thinking and Reasoning*, *10*, 85–104.

Krawczyk, D. C., Holyoak, K. J., & Hummel, J. E. (2005). The one-to-one constraint in analogical mapping and inference. *Cognitive Science*, *29*, 797–806.

Kuehne, S. E., Gentner, D., & Forbus, K. D. (2000). Modeling infant learning via symbolic structural alignment. In L. R. Gleitman & A. K. Joshi (Eds.), *Proceedings of the twenty-second annual conference of the cognitive science society* (pp. 286–291) Ann Arbor, MI: Cognitive Science Society.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago, IL: University of Chicago Press.

Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(3), 754–770.

Loewenstein, J., Thompson, L., & Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin & Review*, *6*(4), 586–597.

Lombrozo, T. (2010). Causal-explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology*, *61*, 303–332.

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

Markman, A. B. (1997). Constraints on analogical inference. *Cognitive Science*, *21*, 373–418.

Markman, A. B. (1999). *Knowledge representation*. Mahwah, NJ: Lawrence Erlbaum Associates.

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

Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental & Theoretical Intelligence*, *13*, 329–358.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: Henry Holt.

Mix, K. S. (1999). Similarity and numerical equivalence: Appearances count. *Cognitive Development*, *14*, 269–297.

Paik, J. H., & Mix, K. S. (2006). Preschoolers' similarity judgments: Taking context into account. *Journal of Experimental Child Psychology*, *95*, 194–214.

Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.

Quine, W. V. (1969). *Ontological relativity and other essays*. New York: Columbia University Press.

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.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*(4), 629–639.

Ross, B. H., & Murphy, G. L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognitive Psychology*, *38*, 495–553.

Rottman, B. M., Gentner, D., & Goldwater, M. B. (2012). Causal systems categories: Differences in novice and expert categorization of causal phenomena. *Cognitive Science*, *36*(5), 919–932.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*, 1926–1928.

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

Smith, L. B. (1989). A model of perceptual classification in children and adults. *Psychological Review*, *96*(1), 125.

Spencer, J. P., Perone, S., Smith, L. B., & Samuelson, L. K. (2011). Learning words in space and time: Probing the mechanisms behind the suspicious-coincidence effect. *Psychological Science*, *22*(8), 1049–1057.

Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, *24*(04), 629–640.

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.

Tomasello, M. (2000). Do young children have adult syntactic competence? *Cognition*, *74*(3), 209–253.

Tomasello, M., Akhtar, N., Dodson, K., & Rekau, L. (1997). Differential productivity in young children's use of nouns and verbs. *Journal of Child Language*, *24*(02), 373–387.

Trench, M., & Minervino, R. A. (2015). The role of surface similarity in analogical retrieval: Bridging the gap between the naturalistic and the experimental traditions. *Cognitive Science*, *39*(6), 1292–1319.

Wolff, P., & Gentner, D. (2011). Structure-mapping in metaphor comprehension. *Cognitive Science*, *35*(8), 1456–1488.

Xu, F. (2016). Preliminary thoughts on a rational constructivist approach to cognitive development: Primitives, symbols, learning, and thinking. In D. Barner, & A. S. Baron (Eds.), *Core knowledge and concept change* (pp. 11–28). Oxford, UK: Oxford University Press.

Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, *114*(2), 245.