

Cognitive Science (2016) 1–50
Copyright © 2016 Cognitive Science Society, Inc. All rights reserved.
ISSN: 0364-0213 print / 1551-6709 online
DOI: 10.1111/cogs.12377

Extending SME to Handle Large-Scale Cognitive Modeling

Kenneth D. Forbus,^a Ronald W. Ferguson,^b Andrew Lovett,^c
Dedre Gentner^c

^aEECS Department, Northwestern University

^bLeidos, Reston, VA

^cPsychology Department, Northwestern University

Received 25 November 2014; received in revised form 1 September 2015; accepted 25 November 2015

Abstract

Analogy and similarity are central phenomena in human cognition, involved in processes ranging from visual perception to conceptual change. To capture this centrality requires that a model of comparison must be able to integrate with other processes and handle the size and complexity of the representations required by the tasks being modeled. This paper describes extensions to Structure-Mapping Engine (SME) since its inception in 1986 that have increased its scope of operation. We first review the basic SME algorithm, describe psychological evidence for SME as a process model, and summarize its role in simulating similarity-based retrieval and generalization. Then we describe five techniques now incorporated into the SME that have enabled it to tackle large-scale modeling tasks: (a) *Greedy merging* rapidly constructs one or more best interpretations of a match in polynomial time: $O(n^2 \log(n))$; (b) *Incremental operation* enables mappings to be extended as new information is retrieved or derived about the base or target, to model situations where information in a task is updated over time; (c) *Ubiquitous predicates* model the varying degrees to which items may suggest alignment; (d) *Structural evaluation* of analogical inferences models aspects of plausibility judgments; (e) *Match filters* enable large-scale task models to communicate constraints to SME to influence the mapping process. We illustrate via examples from published studies how these enable it to capture a broader range of psychological phenomena than before.

Keywords: Analogy; Cognitive simulation; Symbolic modeling; Similarity; Analogical reasoning; Analogical learning; Artificial intelligence; Cognitive psychology

1. Introduction

Psychological studies of analogy and similarity suggest that there are core processes of comparison and analogical inference that enter into tasks ranging from visual perception to conceptual change (Gentner, Holyoak, & Kokinov, 2001; Hofstader & Sander, 2013; Holyoak & Thagard, 1995). These comparison processes can be characterized by the principles of structure-mapping theory (Forbus, Gentner, & Law, 1995; Gentner, 1983; Gentner & Markman, 1997). To be sure, some aspects of analogical processing seem to vary across different tasks: comparison in visual perception operates at a faster time scale than reasoning through an analogy in a debate, for example. However, many fundamental properties (for example, 1:1 mappings and systematicity) are preserved over a surprisingly broad range of cognitive processes (Gentner, 2003, 2010; Gentner & Markman, 1997; Krawczyk, Holyoak, & Hummel, 2004, 2005).

The Structure-Mapping Engine (SME; Falkenhainer, Forbus, & Gentner, 1986, 1989) embodies a process model of how structure-mapping takes place. A number of other computational models—including IAM (Keane & Brayshaw, 1988), ACME (Holyoak & Thagard, 1989), COPYCAT (Mitchell, 1993), TABLETOP (French, 1995), LISA (Hummel & Holyoak, 1997), DRAMA (Eliasmith & Thagard, 2001), AMBR (Kokinov & Petrov, 2001), CAB (Larkey & Love, 2003), and DORA (Doumas, Hummel, & Sandhofer, 2008)—also draw on structure-mapping theory, but use different processing algorithms. We briefly review some of these later in this paper. (For a more complete review of current simulations of analogical comparison, see Gentner & Forbus, 2011.)

SME's basic algorithm has held up quite well since its first publication in 1986, with a few important changes. First, SME no longer uses separate rules for analogy and similarity as in the earliest version; instead, it uses overall similarity—it seeks matches at every level, from object attributes to higher-order relations (Gentner, Rattermann, & Forbus, 1993). Because SME favors structural consistency and systematicity (further described below), relational matches will still tend to win out over shallower matches. The second change is that whereas the initial version computed all possible interpretations of a comparison, SME now uses a greedy merge process (described below) that finds one or a few best interpretations (Forbus & Oblinger, 1990). The third change (also described below) is that SME now carries out structural evaluation of candidate inferences, allowing it to estimate the potential soundness inferences. SME has successfully modeled a large range of comparison tasks, including both developmental (Gentner et al., 1993; Loewenstein & Gentner, 2005) and adult tasks (Gentner, Loewenstein, Thompson, & Forbus, 2009; Markman & Gentner, 1993; Sagi, Gentner, & Lovett, 2012).

This paper summarizes what is needed to go beyond modeling analogy in isolation, on local tasks. Our goal is to make good on the claim that analogical comparison is a core process that cuts across different domains and tasks. Thus, the main focus of this paper is to describe extensions that allow for what we term *large-scale* analogical processing—processing that operates in concert with other cognitive processes in the kinds of complex tasks that occur in everyday reasoning. We aim to capture the way

analogy is used in tasks such as visual reasoning, understanding and solving textbook problems, and moral decision-making. These kinds of tasks have been largely neglected in the modeling literature. In part, this is simply because modeling large-scale phenomena is hard—both because of the complexity of creating such software and because of the difficulty of gathering the necessary psychological data to guide the design and evaluate the results. But another reason is that most analogical matchers, including our own first-generation version of SME (Falkenhainer et al., 1986), cannot scale up to these challenges (Eliasmith & Thagard, 2001; Larkey & Love, 2003). Most work on analogical modeling, including some of our own prior work, suffers from three limitations: (a) restriction to small analogies; (b) the use of hand-coded representations; and (c) failure to integrate analogical processes with other parts of cognition. Our term “large scale” is meant to convey an approach to simulation that addresses these three points.

This paper describes a set of techniques we have incorporated into SME over the last two decades to improve its capacities and meet the challenge of simulating large-scale psychological phenomena involving comparison. The techniques are as follows:

- 1 *Greedy merging* enables SME to rapidly construct up to three near-optimal global interpretations, guaranteeing polynomial-time operation.
- 2 *Structural evaluation of candidate inferences* enables SME to model judgments of the plausibility and interestingness of an analogical inference.
- 3 *Incremental matching* enables SME to model tasks for which information is not all available at once.
- 4 *Ubiquitous predicates* enable SME to model the varying degrees to which items may suggest alignment.
- 5 *Match filters* enable task models to communicate constraints to SME to influence the mapping process.

Importantly, when modeling small-scale analogies in isolation (including most of the examples used in psychological studies), incremental matching, ubiquitous predicates, and match filters are unnecessary. SME always uses greedy merge and structural evaluation of inferences. The importance of incremental matching, match filters, and ubiquitous predicates is in allowing extension to large-scale tasks, where the nature of the tasks and their representations tightly constrain their usage.

Section 2 briefly reviews the principles of structure-mapping and provides a high-level overview of how SME works to provide the necessary context for what follows. Section 3 discusses the issue of scale in analogical processing and summarizes five investigations that provide evidence for the necessity of considering larger descriptions in analogical processing than most simulations have used. Section 4 discusses the techniques above in detail, showing how they work and why we believe they are psychologically plausible. Section 5 shows that the theoretical worst-case complexity of SME is $O(n^2 \log(n))$, where n is the number of items in the base or target.¹ We also show, via an empirical complexity analysis over a large number ($>5,800$) of examples from simulation studies, that the theoretical bound is a gross over-estimate of resource requirements in

realistic situations, although it does capture growth in resource usage appropriately. Section 6 discusses related work not already brought up in earlier sections, and Section 7 discusses some broader issues. The current SME algorithm and the complexity analysis are described in the supplemental material.

2. A brief summary of structure-mapping theory and SME

Structure-mapping theory (Gentner, 1983, 1989, 2010; Gentner & Markman, 1997) postulates that analogy and similarity operate via the same structural alignment process, operating over structured representations. That is, the descriptions being compared are symbolic, including entities, attributes of those entities, and relationships between entities and other relationships.

Formally, the elements of SME's representations are objects (or entities), object attributes (one-place predicates), relations (predicates that take two or more arguments), and functions. As predicates, attributes and relations express assertions with potential truth values, such as `HOT(sun)` or `REVOLVE-AROUND(earth, sun)`.² In contrast, functions map from a set of arguments onto another argument—typically, a dimension. Functions are chiefly used to express dimensional information, such as `DIAMETER(earth) = 12,742 km`. Unlike other predicates, functions can match nonidentically within a mapping: For example, in a spring-pendulum analogy (see below), `spring-constant` matches `length`. Psychologically, functions capture the phenomenon that people fluently map relational structures across dimensions. Extended mappings involving functional correspondences are common in everyday language as well as in scientific contexts; for example, vertical depth \rightarrow profundity, as in “The work is quite shallow/it needs to be deepened/try to get to the bottom of this” (Boroditsky, 2000; Gentner, Bowdle, Wolff, & Boronat, 2001; Lakoff & Johnson, 1980; Thibodeau & Durgin, 2011). A final distinction is the *order* of a relation. The order of a relation is 1 plus the order of its highest order argument. First-order relations are relations between objects. Higher-order relations are relations between other statements. Examples of higher-order relations include logical connectives (e.g., `IMPLIES`), causal relationships, and modal operators (e.g., `BELIEVES`).

The same structural alignment process is used whenever a comparison is to be done, whether it is an analogy or a similarity match. Moreover, the same process is used when contrasting two things, as discussed below. The results of a comparison can be classified based on the kind of overlap that the process finds. A match is considered to be *literal similarity* if both attributes and relations align, *analogy* if relations align but attributes do not, *surface similar* if attributes align but relations do not, and *anomaly* if neither align. These are graded concepts, of course, since matches are rarely perfect.

According to structure-mapping theory, structural alignment takes as input two structured representations (*base* and *target*) and produces as output a set of *mappings*. Each mapping consists of

1. A set of *correspondences* between items (i.e., entities and expressions) in the base and items in the target.
2. A set of *candidate inferences*—surmises about the target made on the basis of common structure plus the base representation. Reverse candidate inferences can also be computed, from target to base, and these can give rise to *alignable differences* (Lovett, Tomai, Forbus, & Usher, 2009; Markman & Gentner, 1996). These inferences can include *analogy skolems*, which represent a projected entity conjectured to exist in the other description. For example, in the historical heat/water analogy, a new entity, *caloric*, was conjectured to exist as one of the consequences of the analogy. Analogy skolems are defined as functions of the entity in the originating description, and can be read as “something like” the original entity.
3. A *structural evaluation score* indicating the overall quality of the match. This is a purely structural score; other considerations, such as the relevance of the inferences, are computed outside the mapping engine (i.e., outside of the structure-mapping process).

Mappings are governed by the following constraints:

1. *Structural consistency*: Structural consistency is defined by two constraints. The first, the *1:1 constraint*, requires that each item in the base maps to at most one item in the target and vice-versa. The second, the *parallel connectivity constraint*, requires that if a correspondence between two statements is included in a mapping, then so must correspondences between their arguments.
2. *Tiered identicality*: Identical matches between predicates (relations and attributes) and functions are preferred. By default, relations must match identically, but non-identical functions can be aligned if such alignments would support a larger overlapping structure. A classic analogy from the history of science, for example, aligns PRESSURE with TEMPERATURE. Many conventional metaphors involve aligning nonidentical functions, such as height to intelligence (“Her IQ is through the roof”) or space to time (“Winter is behind us”).

Depending on task demands, the identicality constraint can be relaxed further to allow nonidentical relations to correspond, if they are suggested by a larger structure and satisfy additional criteria. The most commonly used additional criterion is *minimal ascension* (Falkenhainer, 1987), whereby nonidentical predicates are required to share a close superordinate.

3. *Systematicity constraint*: Preference is given to mappings that align systems of relations in the base and target, especially including those involving nested expressions—that is, those involving higher-order relations.

Each of these theoretical constraints is motivated by the role analogy plays in cognitive processing. The 1:1 and parallel connectivity constraints ensure that the candidate inferences of a mapping are well defined. Tiered identicality is a strong semantic constraint, avoiding structurally isomorphic but nonsensical mappings. The systematicity

constraint reflects a (tacit) preference for coherence and inferential power in analogical reasoning.

There is now widespread agreement on the importance of structural consistency constraints in analogical reasoning (cf. Eliasmith & Thagard, 2001; Kokinov & French, 2003; Hummel & Holyoak, 1997; Krawczyk et al., 2004, 2005; Larkey & Love, 2003). However, some computational models of analogy do not utilize systematicity, and there is no consensus on how central other constraints might be (e.g., pragmatics; Holyoak, 1985) and on how these constraints should be expressed computationally (see Gentner & Forbus, 2011, for a review).

Our own simulation, the SME (Falkenhainer et al., 1986, 1989), works roughly as follows: Given base and target descriptions, SME finds globally consistent interpretations via a local-to-global match process, which can be divided into three phases (see Fig. 1).

Phase One: Constructing the match hypothesis network: SME begins by proposing local correspondences, called *match hypotheses*. All possible local identity matches between expressions in the two representations are made in parallel,³ regardless of whether they are mutually consistent. Additional matches are made via local parallel connectivity—for example, if two relations are matched, then SME attempts to match their arguments. No attempt is made at this stage to enforce global consistency. Thus, the initial network is inchoate, providing the material for potential mappings.

Phase Two: Parallel construction of structurally consistent kernels: At this point, SME starts building mappings by extracting structurally consistent sets from the forest of match hypotheses. To do so, SME does two things, again in parallel.

1. It marks as inconsistent those match hypotheses that violate parallel connectivity.
2. It marks as mutually inconsistent pairs of match hypotheses that would violate the 1:1 constraint.

It then coalesces the local matches into a set of structurally consistent connected structures, called *kernels*. Kernels are the grist from which mappings are created.

Each kernel receives a structural evaluation score. SME begins the structural evaluation process by first assigning a local score to each match hypothesis, then using a trickle-down process to propagate evidence downwards from a match hypothesis to match hypotheses between the arguments of the corresponding statements. This provides a local means of implementing the systematicity preference, since match hypotheses that participate in a large matching structure will receive a higher score.

Phase Three: Constructing mappings: SME uses a greedy merge algorithm to combine kernels into one or more global mappings. The basic idea is to start with the largest and deepest kernel (that is, the one with the highest structural evaluation) and serially add others that are structurally consistent with it. This results in one or a few large, structurally consistent *global mappings* between the representations. The global mapping reveals common structure between the two representations. At this stage, candidate inferences may be projected from one representation to the other and alignable differences emerge.

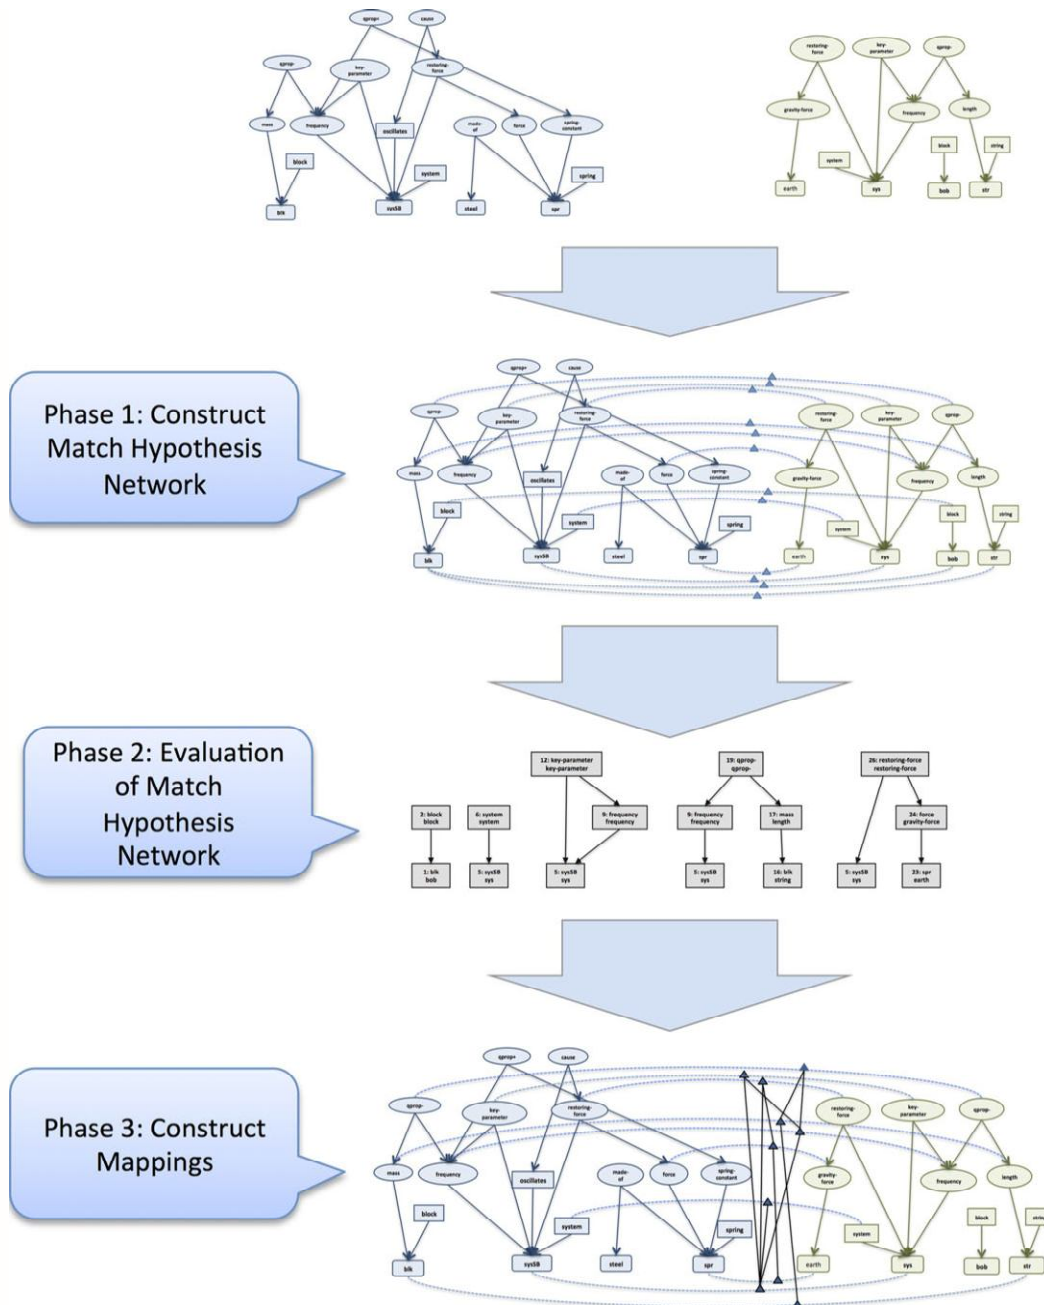


Fig. 1. The three phases of the SME algorithm.

SME computes only a handful of mappings, based on the settings of two parameters. The *Max Limit* is an upper bound on the number of mappings produced, and defaults to three. The *Score Cutoff* is a drop in score below which further mappings are ignored, and defaults

to 0.8; that is, any mappings produced must be within 20% of the best mapping to be output. These parameters help model capacity limits in analogical mapping. For each mapping, SME computes the candidate inferences for the mapping and its structural evaluation score.

For concreteness, consider the simple example in Fig. 2, based on a commonly used analogy in physics instruction. It describes a cross-domain analogy between two systems that oscillate. The base is a classic spring-block oscillator. The attributes `spring`, `block`, and `system` describe the types of entities involved, while relationships such as `made-of` and `restoring-force` describe some of the basic physical properties of the system. There are two causal statements that describe relationships between continuous parameters of the system, which are drawn from Qualitative Process theory (Forbus, 1984). The `qprop+` statement says that the frequency of oscillation increases if the spring constant is increased, all else being equal. The `qprop-` statement says that the frequency of oscillation decreases if the mass of the block is increased, all else being equal. Finally, the restoring force is specified as the cause of the system's oscillation, using the `cause` relationship. The target is a pendulum. It, too, has a restoring force, and like the base, has the frequency of oscillation identified as an important aspect of the system. There is one causal relationship specified about frequency, namely that if the length of the string is longer, the frequency will decrease (i.e., the `qprop-` statement). Fig. 3 also shows a graphical visualization of these representations, with the height of each predicate indicating what order it is. Fig. 4 shows the match hypothesis forest that SME generates for this example, and Fig. 5 shows the kernels. The mapping SME constructs is shown in Fig. 6.

While abstract, this description of SME's operation is enough to ground the discussion of the advances described in Section 4. These advances interact—for instance, greedy merge implies the need to sometimes use automatically derived filters in the matching

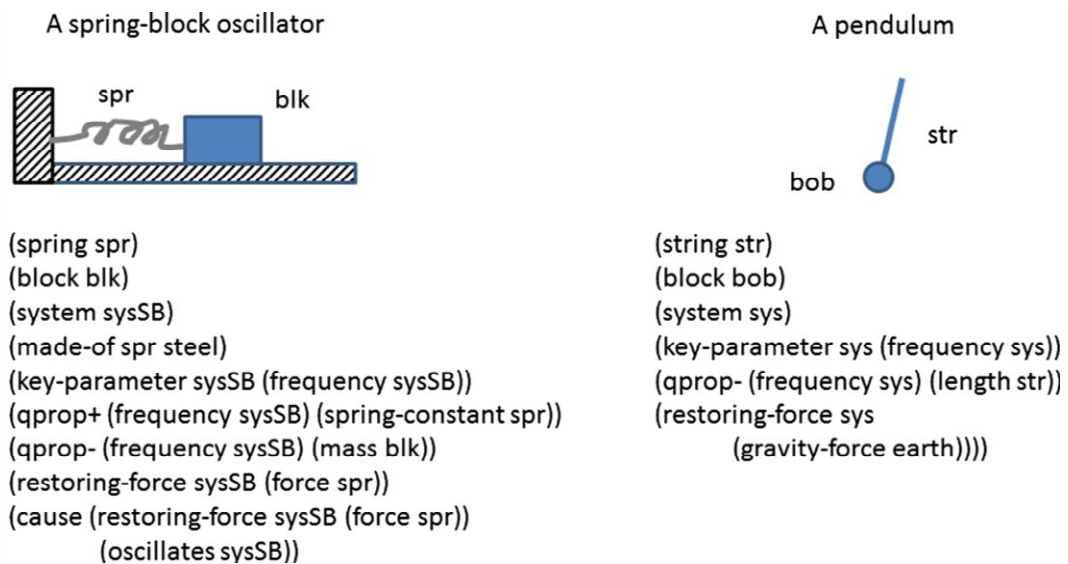


Fig. 2. Analogous oscillators.

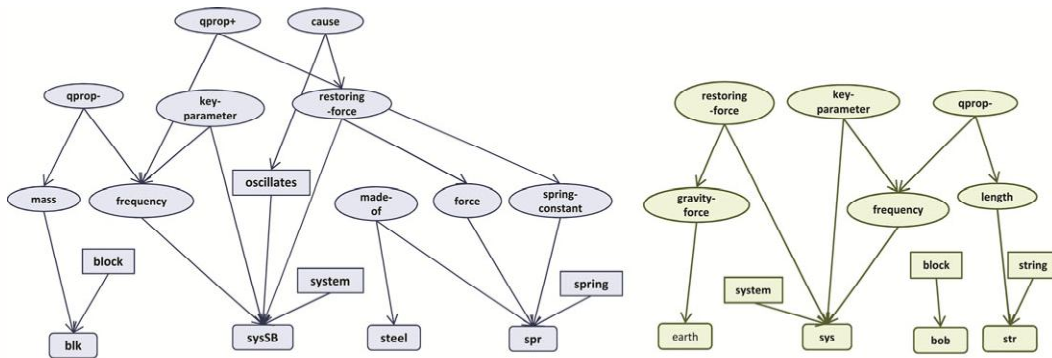


Fig. 3. Graphical depiction of the base and target for the oscillators example.

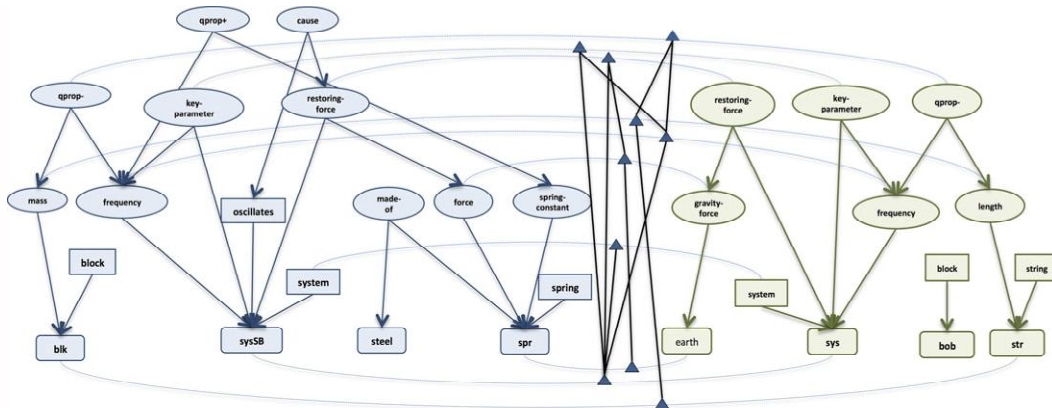


Fig. 4. Match hypothesis forest for the oscillators.

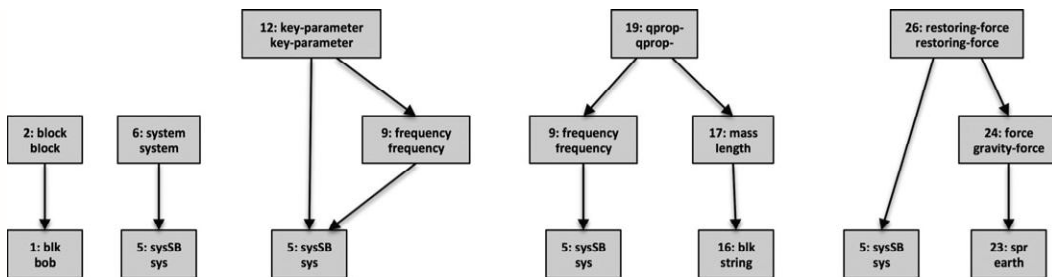


Fig. 5. Kernels for the oscillators comparison.

process—but to a first approximation they can be described independently. The supplemental material describes the new SME algorithm in detail, illustrating how these techniques work together and analyzing its overall complexity. Complexity is a critical question: A central operation in cognitive processes must work in polynomial time, in order to account for the rapidity and scalability of human processing. We return to this

Mapping 15				Expression Correspondences			Candidate Inferences		
SME #1				SME #2			SME #1		
[3 items]				[3 items]			[5 items]		
Score: 0.1791									
Base: spring-block-oscillator									
Target: pendulum									
Support	Base Item	Target Item	MH Score	Base Item	Target Item	Score	Inference	Support	Extrapolation
(5)	sys	sys	0.0560	spring part	gravity-force earth	0.0041	(spring earth)	0.0005	0.9000
(2)	disk	bob	0.0080	frequency sys	frequency sys	0.0041	(part-of sys earth)	0.0015	0.8500
(2)	sp	earth	0.0992	spring-constant part	gravitational-constant earth	0.0041	(spring- (frequency sys) (gravitational-constant earth))	0.0068	0.5556
				block bob	block bob	0.0041	(spring- (frequency sys) (mass bob))	0.0050	0.6296
				two-oscillator sys (frequency sys)	two-oscillator sys (frequency sys)	0.0041	(cause (restoring-force sys (gravity-force earth)) (oscillates sys))	0.0090	0.4857
				spring (force part) (spring-constant part)	spring (gravity-force earth) (gravitational-constant earth)	0.0041			
				restoring-force sys (force part)	restoring-force sys (gravity-force earth)	0.0041			
				part-of sys bob	part-of sys bob	0.0041			
				bob	bob	0.0041			
				bob	bob	0.0041			

Fig. 6. The mapping SME constructs between the oscillators.

issue in Section 5. The rest of this section summarizes some psychological evidence for SME and outlines how SME is used as a component in models for analogical retrieval and category formation, to illustrate the explanatory power of this model.

SME’s processing account has led to several predictions that have been borne out in psychology experiments. Table 1 shows a set of benchmark phenomena that characterize analogical mapping (Gentner & Markman, 1995; Markman & Gentner, 2000). For example, there is considerable evidence that people prefer structural consistency, including 1-1 mappings, in analogical processing (Gentner & Markman, 2006; Krawczyk et al., 2004, 2005; Markman & Gentner, 1993; Spellman & Holyoak, 1996). There is also evidence for systematicity: People prefer analogies that share deep systematic structure over matches that are otherwise identical, but that lack common higher-order relations linking the lower-order relations (Gentner et al., 1993). Systematicity also influences analogical inferences: People draw inferences from the more systematic to the less systematic of two analogs (Bowdle & Gentner, 1997) by projecting predicates connected to the common structure (Clement & Gentner, 1991).

There is also evidence supporting SME’s local-to-global matching algorithm—specifically, that the first stage of comparison processing is an initial symmetric alignment process. Wolff and Gentner (2011) tested comprehension of strongly directional metaphors, such as “Some suburbs are parasites.” In untimed tasks, people strongly prefer “Some suburbs are parasites” to “Some parasites are suburbs.” However, if participants had to answer under deadline pressure (within 600 ms), they did not exhibit a directional

Table 1
Benchmark phenomena of analogy

Relational similarity	Analogies involve relational commonalities; object commonalities are optional
Structural consistency	Analogical mapping involves one-to-one correspondence and parallel connectivity
Systematicity	In analogical mapping, connected systems of relations governed by higher-order constraining relations are preferred over isolated relations
Candidate inferences	Analogical inferences are generated via structural completion
Alignable differences	Differences that are connected to the commonalities of a pair (and not unconnected differences) are rendered more salient by a comparison
Interactive interpretation	Analogy interpretation depends on both terms; the same term yields different interpretations in different comparisons
Multiple interpretations	Analogy allows multiple interpretations of a single comparison
Cross-mapping	Though difficult, cross-mappings are generally interpreted relationally although the competing object similarities are perceived

Note. Adapted from Markman and Gentner (2000).

preference: they found both versions equally comprehensible (Importantly, the mean “comprehensible” response was significantly greater than for scrambled metaphors, indicating that meaningful processing had been initiated.) By 1,200 ms, there was a significant advantage for the forward direction. This is exactly what follows from an initially symmetric alignment process, assuming the deadline occurs during the first two phases of processing. During this early alignment stage, we conjecture that people can have the sense that something promising is happening, even in cases where they will ultimately reject the comparison.

Further psychological evidence for SME’s processing algorithm comes from investigations of difference processing. SME predicts an empirical disassociation in the response times for two seemingly related tasks: The same-different task, and a “name the difference” task. A long-established finding is that in same-different tasks, people are faster to respond “different” for very dissimilar pairs than for similar (but nonidentical) pairs (Goldstone & Medin, 1994; Luce, 1986; Posner & Mitchell, 1967). SME can capture this finding, as described below. But SME also makes a novel prediction: that if asked to state a difference between two things, people should be faster to do so for very similar pairs (Sagi, Gentner, & Lovett, 2012). This prediction rests on two prior findings. First, there is abundant psychological evidence that when people are asked to state differences between two things, they are likely to name *alignable differences*—differences that play the same role in the common structure⁴ (Gentner & Gunn, 2001; Gentner & Markman, 1994; Kurtz & Gentner, 2013; Markman & Gentner, 1993, 1996). By their nature, these differences emerge only after structural alignment is complete. Second, high-similarity pairs are faster to align than low-similarity pairs (Gentner & Kurtz, 2006). This also follows from SME’s process. In high-similarity pairs, since most of the matches are compatible, there will be one or two large, dominant kernels. This means that the greedy merge process generally only needs to run once, since SME does not bother producing more than one mapping when there is little left over. For low-similarity pairs, there are typically many small kernels, and the final step may require comparing two or more different merges. Thus, alignment takes longer for low-similarity than for high-similarity pairs (all else being equal). Thus, naming a difference should be faster for high-similarity pairs than for low-similarity pairs.

Now consider a same-different task. If two descriptions are completely different (i.e., the pair is very dissimilar), the size of the initial match hypothesis forest will be small compared to the size of the items. This means that the alignment process can be terminated at the first stage, resulting in an early “different” response. But if the pair is highly (or even moderately) similar, such that the initial stage feels promising (i.e., the match hypothesis forest is large), then the alignment process cannot be aborted at Phase 1 and must be carried to the end. SME thus predicts opposite patterns for these two tasks: faster responding for very similar pairs in a name-a-difference task, and faster responding for very dissimilar pairs in a same-different task. Sagi et al. (2012) found exactly this pattern. To our knowledge, SME is the only simulation of similarity or analogy that can predict this pattern.

SME has several attractive features as a cognitive model. As noted above, it operates in a parallel to serial manner,⁵ and it can capture some essential properties of the comparison process, including abstraction, inference projection, and alignable difference detection. It also matches findings on the ordinal time course of human comparison processes, as discussed above—such as the disassociation in the time course of difference processing (Sagi et al., 2012) and the fact that that high-similarity matches are processed faster than low-similarity matches (Gentner & Kurtz, 2006). Most important, SME operates without advance knowledge of the point of the comparison—capturing the fact that people can happen upon a comparison with no advance idea of what it will yield and discover a hitherto unnoticed common structure or a new inference.

2.1. The evolution of SME

SME has evolved considerably since 1986. For example, in the initial version of SME (Falkenhainer et al., 1989), we used different rule sets to process different match types—analogy, literal similarity, and mere-appearance (surface similarity). In terms of cognitive modeling, this amounts to the assumption that people process analogy with a different set (attending only to relations) than they use for literal similarity (in which attention goes to both relations and object attributes). But this has the disadvantage of having to postulate that people know in advance what kind of match they will be getting—violating a prime goal of capturing spontaneous discovery. Moreover, such rule sets turn out to not be necessary (Forbus, Ferguson, & Gentner, 1994). SME now runs in what used to be called literal similarity mode, in which it tries to match all types of predicates—attributes, functions, and relations. Depending on what it finds to match, the comparison will be characterized as literal similarity, analogy, surface match, or anomaly, or as something intermediate.

We further note that SME can be used with cases from a variety of different sources; it can process cases whether they were originally presented perceptually, given in text, retrieved from long-term episodic memory, produced dynamically from semantic memory (Mostek, Forbus, & Meverden, 2000), or derived through reasoning and problem solving (as the examples in Section 3 illustrate). In larger models, it tends to be used in a map/analyze cycle (Falkenhainer, 1990), in which the results of a comparison are evaluated in a task-dependent way. The models in Section 3 illustrate these ideas in the context of large-scale cognitive tasks.

A key principle in this work is that the basic similarity engine—SME—should be able to operate by default on its own, without external guidance. There are two reasons for this. First, comparison often occurs spontaneously. Although we sometimes compare things on command—for example, when we are told that two things are analogous—it is clear that the comparison process often occurs unbidden. For example, walking through a city, we generally are not looking for twins, but if twins should appear, we will notice their similarity. This kind of “comparison-based interrupt” also happens at the conceptual level; noticing that two ideas are similar can be a source of insight. Indeed, there is evidence that comparison-based inference can happen even without our noticing it (Day &

Gentner, 2007). A second argument for a relatively independent similarity engine is that comparison is ubiquitous in cognition, in arenas from decision-making to causal reasoning to categorization to perceptual learning. Although it is theoretically possible that different comparison processes are called on in each arena,⁶ the many phenomenological similarities across areas suggest that the same basic comparison process is involved across a wide swath of cognition. This argues for a relatively modular similarity engine that can be used as a subprocess within more complex cognitive processes.

2.2. *SME as a component in retrieval and generalization*

Another piece of evidence that something like SME may operate as a general-purpose similarity operation is that it can be productively viewed as a subprocess in analogical retrieval and generalization. We illustrate by briefly summarizing how SME is used in MAC/FAC, a model of analogical retrieval, and SEQL (and its newer incarnation, SAGE), a model of analogical generalization.

MAC/FAC (Forbus et al., 1995) models similarity-based retrieval—the phenomenon by which a currently active representation reminds us of some prior similar situation. We model similarity-based retrieval as a fast, relatively indiscriminate process. This choice is motivated in part by information-level considerations (in Marr’s [1983] sense): This process must by its nature operate over large swaths of LTM, suggesting that it should be computationally cheap to carry out. A second rationale is the empirical fact that similarity-based retrieval is a hit-or-miss phenomenon—the retrieved instances are sometimes relevant to the current situation, but often they simply share some surface features. Yet, once the instance is retrieved, people can often show considerable discernment, rejecting their own retrievals if they lack structural similarity (Gentner et al., 1993). MAC/FAC uses a two-stage retrieval process to capture these phenomena. The first stage (MAC) uses a vector representation automatically generated from the structured representations in working memory and in long-term memory as a crude, fast search. These content vectors consist of the relative frequency of occurrence of the predicates and functions in the structured representation. Psychologically, this captures the fact that memory retrieval is strongly influenced by content, and only weakly influenced by relational structure (Gentner et al., 1993; Holyoak & Koh, 1987; Trench & Minervino, 2014). It also captures the idea that deliberate indexing is not necessary in order for retrieval to occur. Computationally, these vectors have the property that their dot product provides an estimate for the size of the match hypothesis forest that SME would produce for the corresponding structural descriptions. We assume that the dot products are happening in parallel, and that the top few (up to three) structured representations corresponding to the winning vectors are passed on to the second stage. The second stage (FAC) runs SME in parallel, comparing these structured representations to the current situation, returning one or more mappings as the result. Studies comparing MAC/FAC’s retrieval patterns with those of humans have shown good ordinal matches (Gentner et al., 1993, 2009).

SME is also used as a component process in models of analogical generalization. SEQL (Kuehne, Forbus, Gentner, & Quinn, 2000) used SME to compare incoming

examples and assimilate them into ongoing generalizations. A later version, SAGE (McLure, Friedman, & Forbus, 2010), keeps track of frequency information about alignable structures, enabling it to produce probabilistic generalizations. For example, if given a series of examples with the same category label, SAGE begins by storing the first input example. When the next example arrives, SAGE compares it to the first one, using SME. If there is sufficient overlap (that is, if SME's score is above a pre-set 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; if new examples are sufficiently similar to the ongoing generalization, they are assimilated into it and the generalization is updated. New examples that cannot be assimilated into the main abstraction are compared to the set of examples; if a new example is very similar to a stored example, a new generalization is formed from their common structure. For example, suppose SAGE is given a set of items labeled "birds." If it receives the series *sparrow*, *thrush*, *starling*, *finch*, it will form an abstraction that we would recognize as *songbird*. If the next example of "bird"—say, a *stork*—is insufficiently similar to the abstraction, then it will be stored as a separate example. This process allows SAGE to form similarity-based subclusters; for example, if a *heron* and *egret* were added to the *bird* category, SAGE will form an abstraction of tall, long-legged birds subsuming *stork*, *heron*, and *egret*.

The goal of using SME in models of large-scale cognitive processes brings up the crucial issue of scale, to which we turn next.

3. Representation and scale in analogical processing: Evidence from cognitive models

One point of strong agreement among analogy researchers is the importance of representation in modeling analogy—in particular, the need for explicit representations of relations. Flat feature vectors, even if very large, have no way to capture the phenomena of relational matching and mapping (Gentner & Markman, 1993, 2006; Goldstone, Medin, & Gentner, 1991; Holyoak & Hummel, 2000; Markman, 1999; Sagi et al., 2012). In the early days of analogical research, this need was met by using hand-generated examples. Indeed, many prominent models of analogical processing have been tested only with hand-generated examples (e.g., ACME, LISA, DRAMA, CAB, DORA). SME has also been tested with such examples, especially in the early years (e.g., Gentner et al., 1993). While considerable insight can be gained from experiments using such materials, they raise significant methodological concerns. The most serious is *tailorability*—that is, the degree to which a model's results depends on representation choices that are not theoretically constrained, and chosen by the modeler (often unknowingly) to make the simulation come out in the way that they desire. Tailorability is difficult to avoid in hand-coded representations, because many questions about the exact formats of mental representations are not yet strongly constrained by data. For example, even though there is agreement that relations are important in human cognition, exactly *which* relations people use to

represent a given situation is still an open question. When hand-coded representations must be used, tailorability can be reduced by using independent evidence as to the likely representations when it is available, and by using uniform attested representational conventions otherwise.

However, by far the best way to reduce tailorability is to use representations that are independently generated—either automatically generated (e.g., Falkenhainer, 1990) and/or produced outside one's laboratory (Forbus et al., 1995). In the last 20 years, although we have sometimes used hand-coded representations (e.g., Gentner et al., 2009), we have placed a high priority on testing SME with independently generated representations. Automatically generated representations can be taken from other AI systems that are carrying out some task. In systems that have SME as a component, SME's representations can be produced by another part of the system itself, through interpreting natural language text (e.g., Dehghani, Tomai, Forbus, & Klenk, 2008) or sketches (e.g., Lovett et al., 2009). This has the advantage that in cognitive simulation, the materials given to the system are the same stimuli given to human subjects (Lovett et al., 2009; Sagi et al., 2012). Below we describe tests of SME using automatically generated representations, including representations derived from perceptual input (see Sections 3.1.1 and 3.1.2), representations taken from other AI systems, including a natural language system (Sections 3.1.3 and 3.1.5), and representations provided by another institution (Section 3.1.4).

A further methodological concern is the adequacy of the representation—that is, whether it actually contains enough information to carry out the task(s) it is intended for. When carefully evaluated, a running program demonstrates that there is at least one set of representations and processes that can be used to carry out the task—that is, it shows sufficiency (though not necessity). Not all AI systems are intended as cognitive simulations, of course. However, even systems that are not designed as simulations can still provide evidence as to what information content, and how much of it, is required for particular tasks—that is, they can bear on Marr's (1983) information processing level of explanation. Another criterion for an adequate representation (beyond being able to support simulations that carry out the task being modeled) is that it can be used in other tasks that could reasonably be supposed to draw on the same representation. For any single process, the representations can be chosen so as to create the desired outcome. Thus, we agree with Cassimatis, Bello, and Langley (2008), who argue that ability and breadth are crucial, but underutilized, criteria for evaluating cognitive simulations.

A third methodological point is the *Integration Constraint* (Forbus, 2001). This constraint arises from the claim that analogical mapping processes underlie many important cognitive processes. To make good on this claim, we need to test SME as a component of other large-scale cognitive processes, such as categorization or moral reasoning. The Integration Constraint states that a model of a cognitive process P should be usable as a component in models of larger-scale cognitive processes that are hypothesized to use P (Forbus, 2001). Embedding a model within a larger system also reduces tailorability; if the representations are automatically constructed by other processes, and the results of analogical comparison are used by yet other processes, there are many more constraints on the representations than would be found in isolation.

Thus, the properties of tasks impose constraints on models of component processes. In the case of analogical mapping, one implication of this principle is that a model of analogical mapping must be able to handle the kinds of representations that arise in the tasks in which people use analogy. Let us therefore look at representations used by large-scale cognitive models that have used SME as a subprocess.

3.1. Evidence from five computational investigations

Here, we briefly summarize five previously published computational investigations. In each case, SME was used as a component in a larger system. Data were collected by recording the contents of SME during the running of the larger model and saving this information to files.⁷ These investigations provide evidence about the kind of descriptions and the number of relationships that models of analogy must handle.

3.1.1. Geometric analogies

The first simulation of analogy was Evans's (1968) ANALOGY program. Fig. 7 shows two example problems from Evans's original corpus. These are problems of the form "A is to B as C is to . . .?" Running on an IBM mainframe, using punch cards as input, ANALOGY was able to automatically construct representations for half of the examples it operated over, a tour de force for that era. The program used a transformation-based model, with separate domain-specific comparison processes to compare stimuli and to compute transformations between them. However, despite multiple efforts to do so, until recently no model was built that could automatically encode these stimuli and successfully solve the same range of problems.

Lovett et al. (2009) showed that, using automatically constructed inputs, structure-mapping could be used to perform this task. The figures for the problems were drawn using PowerPoint, and copy/pasted into CogSketch (Forbus, Usher, Lovett, Lockwood, & Wetzel, 2011), an open-domain sketch understanding system. CogSketch automatically produces structured, relational representations from digital ink, and these were used as inputs to the system. CogSketch computes a variety of qualitative spatial relationships, including positional relationships (e.g., *above*, *leftOf*), topological relationships (e.g., *partiallyOverlapping*, *inside*), and relative sizes. Recognizing resizing and rotation of shapes was automatically performed by using a model of mental rotation, which in turn uses SME. That is, SME was used to perform a qualitative comparison between two shapes, using an automatically constructed edge level representation, from which quantitative scaling and rotation relations, if appropriate, could be derived. Thus, even the original encoding of the stimuli involves the use of SME over automatically constructed representations.

The original model (Lovett et al., 2009) used a two-stage comparison process. That is, SME was used to compare A to B, and then SME was used to compare C to each possible solution in turn. SME's mappings were then compared to each other, again using SME (i.e., a *second-order comparison*). This involved automatically reifying mappings and candidate inferences as assertions, for input to the second-order comparisons.




Fig. 7. Examples of geometric analogies. “A is to B as C is to . . .?”

This process also identifies reversals and dimensional changes (e.g., `leftOf` in the A/B pair, vs. `above` in a C/X pair). Candidate inferences are computed in both directions, with the presence of analogy skolems indicating that an extra entity is in one of the descriptions. A score is computed for each answer, based on SME’s structural evaluation score for the second order comparison, plus a penalty for extras, normalized to avoid size effects. The choice with the highest scoring comparison with C is selected as the answer.

An interesting complication is that SME’s initial mapping may not always be the best answer. An *executive* process is used to automatically look at mappings other than the best,⁸ if SME has produced more than one, and to consider whether to re-prioritize its interpretation of reflection versus rotation in comparing shapes. Its default strategy is to prefer identity matches, followed by reflections, and then rotations. But if a good solution is not found, it will also explore giving higher preference to reflections or rotations. Problems for which the default encoding strategy suffices to find an answer should be faster for people, since fewer comparisons are required, in contrast to problems which require backtracking and trying other strategies. As predicted, a laboratory study indicates that people do indeed take longer on problems where the model predicts that backtracking is required (Lovett et al., 2009).

An extended model (Lovett & Forbus, 2012) added a second strategy based on constructing an answer by projecting the differences in the A/B comparison onto C, and then comparing the constructed answer to the alternatives. Fig. 7(a) shows an example for which the projection strategy works, and Fig. 7(b) shows an example where second-order comparison is required (because the constructed answer is not one of the valid responses). The combination of the model’s predicted strategy shifts and working memory load together account for most of the variance in human reaction times on these problems (R^2 of 0.95). This model and the complete set of inputs is available for download.⁹ Working memory load was coded for by counting the number of elements involved in the difference computed for two geometric descriptions. We return to the issue of working memory below.

3.1.2. Visual oddity task

To explore possible cultural differences involving geometric reasoning, Dehaene, Izard, Pica, and Spelke (2006) used a visual oddity task (Fig. 8), in which participants are instructed to select the odd image from a group of six images. They gave this task to two distinct groups: North Americans and the Mundurukú¹⁰ (a South American indigenous group). Overall, there was a high correlation in the performance of the groups, suggesting that some aspects of geometric reasoning may be universal. However, there were open questions about how representations and reasoning might vary across the groups. To address these questions, Lovett and Forbus (2011) modeled human performance on this task, using a combination of SME and qualitative visual representations.

Again using stimuli copy/pasted from PowerPoint, CogSketch was used to automatically construct initial representations and do rerepresentation as required by the model. CogSketch incorporates three distinct levels of representation. The *object level* is the default, consisting of relationships between entities. The *group level* describes higher-level relationships between sets of objects, using gestalt principles of proximity and similarity (computed via SME). The *edge level* segments a shape into its component edges and computes qualitative spatial relationships between the edges. The model uses SME multiple times to determine what is in common with half of the images in a problem and looks at the others to find if there is a noticeable dip in similarity. This can require shifting levels of representation, if the default object level does not lead to an answer. Fig. 8(a) shows a problem that can be solved at the object level, while Fig. 8(b) shows a problem requiring an edge-level representation.

Out of 45 problems, the model correctly solved 39 of them, and the problems that it failed to solve were among those hardest for human participants. Moreover, an ablation study suggested reasons for the differences between the two cultural groups found in the original study. The behavior of the Mundurukú is consistent with a stronger focus on the edge level representation, whereas North Americans tended to focus on objects or groups of objects—perhaps an outcome of schooling.

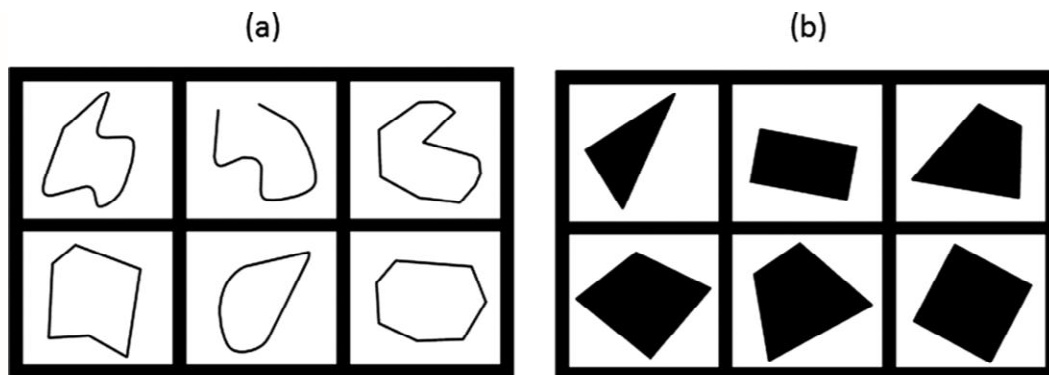


Fig. 8. Examples of a visual oddity task. “Pick the image that does not belong.”

3.1.3. Solving textbook thermodynamics problems

Problem solving is one of the signature roles of analogy in cognition. To explore the importance of incremental mapping in problem solving, we built a very simple problem solver (MARS, Forbus et al., 1994), which had no built-in knowledge of engineering thermodynamics but was capable of using analogies with previously solved problems to solve new problems. The worked solutions were automatically generated by an AI system, CyclePad (Forbus et al., 1999), which has an expert-level model of the domain. CyclePad solutions include explanations, in terms of how each value is derived in terms of others, including what equations were used. These explanations were automatically translated to predicate calculus to be used as analogs for solving new problems. Given a new problem, MARS compared the problem to the analog it was given, and performed an initial mapping. This initial mapping was used to import equations to be used in solving the new problem. As newly derived information is added to the problem being solved, the mapping is incrementally extended, providing new candidate inferences with further potentially relevant information.

We note that this is only one point in the spectrum of how analogy might be used in problem solving (Gentner et al., 1997; VanLehn, 1998). In this approach, the analogy is treated as a recipe for getting an onerous job done quickly. Although people do sometimes use analogy as a shortcut, this approach fails to take advantage of analogy as a learning mechanism. To capture the more ambitious use of analogy, we built a more extensive simulation of human textbook problem solving, which incorporated analogical retrieval to automatically find analogous worked solutions, first-principles qualitative reasoning to provide a more robust initial understanding of the problem and the ability to import control decisions as well as equations via analogy. Ablation experiments with this simulation that removed more expert-like aspects of performance (e.g., more rigorous encoding using qualitative models, validating analogy suggestions via reasoning, and using multiple retrievals) led to more novice-like behavior, providing evidence for these factors being among those underlying novice/expert differences (Ouyang & Forbus, 2006).

We include examples from the original model here, because it is sufficiently simple that its source code will run in any modern Common Lisp environment, enabling others to experiment with it more easily.¹¹ The more sophisticated model's worked solutions are more or less equivalent in structure.

3.1.4. Solving advanced placement physics problems

The Educational Testing Service (ETS) runs Advanced Placement exams for high-school students in the United States. These are high-stakes tests, since doing well on them can improve a student's chances of going to a top university. Fig. 9 shows a typical example. Students must ascertain which equations are relevant, solve them, and check whether their answer makes sense. As an experiment, ETS trained and tested a Companion (Forbus, Klenk, & Hinrichs, 2009) on part of the Dynamics component of the AP Physics examination. The relevance of this test for present purposes is that the Companion cognitive architecture¹² makes the structure-mapping processes of comparison,




Fig. 9. An example AP Physics problem.

retrieval, and generalization central to its operations. The goal of the ETS experiment was to see if a Companion could, from worked solutions, rapidly transfer knowledge across the following six near-transfer conditions:

- 1 *Varying parameters.* Changing the numerical values, but in small ways that do not affect the qualitative outcome (e.g., changing the height of the building to be 81 m).
- 2 *Extrapolation.* Changing the numerical values so much that the qualitative outcome changes (e.g., changing the height of the building to be 10 m).
- 3 *Restructuring.* Asking for a different parameter (e.g., how high will the ball be?).
- 4 *Distractors.* Adding additional events that are irrelevant.
- 5 *Restyling.* Changing the kinds of everyday objects involved in problems.
- 6 *Composing.* Compound problems whose solution requires combining the results of solving two simpler kinds of problems.

The problems involved were drawn from the types of problems found in the Dynamics portion of the AP Physics exam, such as deriving numerical values for a situation, producing the correct prediction of qualitative outcomes in a situation based on numerical values, and producing symbolic equations to characterize a situation.

The representations for problems and worked solutions were generated by ETS. ETS, with the help of Cycorp, modified their problem generation process to produce problems in predicate calculus, using the Cyc ontology (Matuszek, Cabral, Witbrock, & DeOliveira, 2006). In addition to producing problems, they also generated worked solutions, also in predicate calculus. The level of information in the worked solutions was similar to that found in textbook explanations of how to solve physics problems. Importantly, these explanations were not geared toward the workings of the Companions' problem-solving mechanisms—indeed, ETS had no knowledge of how those mechanisms worked. Each step of a worked solution used the same general set of relationships. For example, the relation `solutionStepUses` identifies assumptions and antecedents, `solutionStepResult` indicates the result of that step, and `priorSolutionStep` indicates the sequential relationship between two steps of the solution. Thus, each step requires multiple relations to encode, in addition to the relations involved in the

antecedents and the result. A typical problem can require eight or more steps. Thus, despite the relatively abstract nature of the ETS representations, the number of relationships needed to encode a worked solution is substantial.

To make this a clear test of transfer ability, the Companion had knowledge of how to solve equations and a set of problem-solving strategies, but no prior knowledge of the equations of physics. Its strategies included using MAC/FAC to retrieve prior analogs, looking first for a mapping between event structures similar to the problem it was facing. If a prior analog involved a different kind of event (e.g., throwing instead of dropping), it rejected that reminding and tried again. Candidate inferences were mined for equations and assumptions that could help it solve the current problem. If it could not retrieve a relevant example, it gave up, and did not guess.

All training and testing was done by ETS as follows. Two kinds of training sets were developed:

- 1 Initial training set: Five quizzes, consisting of four problems each.
- 2 Transfer training set: Four quizzes, consisting of four problems each, designed to vary from an initial training set according to which of the six transfer conditions was being tested.

After ETS quizzed a Companion, they gave it the worked solutions corresponding to those quiz problems. This built up the set of potential analogs that a Companion could draw upon. Learning curves across each set were constructed by measuring the number of problems correctly solved in each quiz. For each of the six conditions, the following protocol was used. First, they administered an initial training set followed by a transfer training set (i.e., nine quizzes, 36 problems). Then they wiped the Companion's memory of its recent experiences, and ran just the transfer training set, to provide a baseline. This was done five times for each transfer level, with novel problems used in each training set.

The results were encouraging (Klenk & Forbus, 2009), in that by using MAC/FAC and SME, Companions were able to quickly transfer knowledge across all six conditions. Moreover, their operation could be analyzed in terms of the analogy events (e.g., transferring a line from an example solution, checking transferred knowledge) that VanLehn (1998) found in human students (Klenk & Forbus, 2007).

3.1.5. *Moral decision making*

Cultural narratives have been identified as an important way in which people derive moral understanding about right and wrong behavior (Prasad, 2007; Weber & Hsee, 1999). The prevalence of such narratives suggests that they may serve as the basis for analogies to other moral situations. Indeed, Dehghani, Sachdeva, Ekhtiari, Gentner, and Forbus (2009) have shown that analogical retrieval affects how moral stories are used across cultures. Based partly on prior research suggesting that decision-making often involves analogy (Markman & Medin, 2002), Dehghani's MoralDM system (Dehghani, Tomai, Forbus, Iliev, & Klenk, 2008) provides a computational model of moral decision-making that relies heavily on analogical retrieval. The MoralDM model includes

protected values (Baron & Spranca, 1997), sometimes called sacred values, as well as utilitarian concerns. Protected values are modeled via a qualitative order of magnitude representation—that is, when protected values are present, normal utilitarian differences are lost in the noise. The stories used in psychological experiments were hand-translated into a simplified English syntax, which was then automatically translated into relational representations via a natural language system (Tomai & Forbus, 2009). While the stimuli are simplified in syntax, they involve quite subtle relationships. For example, one simplified English scenario (based on Ritov & Baron, 1999) reads as follows:

Because of a dam on a river, 20 species of fish will be extinct. You can save them by opening the dam. The opening would cause 2 species of fish to be extinct.

Notice that this involves several kinds of events (e.g., extinction, opening), numerical quantification (e.g., that there are two groups of species with different cardinalities) and a counterfactual (e.g., “would cause”). A strictly utilitarian view would lead to opening the dam, but if directly causing extinction of a species by one’s actions is a protected value, then inaction would be preferred. MoralDM uses a combination of first principles and analogical reasoning to work through moral dilemmas. If there are no protected values in a scenario, it looks at the relative utility between the two choices and makes the best choice (i.e., fewer people dying, fewer species of fish becoming extinct). On the other hand, if the highest utility choice involves taking an action that violates a protected value (i.e., taking an action that will directly cause an extinction), consistent with most participants in these experiments, it prefers inaction to action. In addition to reasoning from first principles, MoralDM also uses SME to look for actions that violate protected values. These methods complement each other, since the rules used for first-principles reasoning are incomplete. MoralDM makes choices that are compatible with the majority choices made in multiple experiments (Dehghani et al., 2008). Moreover, as the size of the case library grows, the proportion of correct responses improves (Dehghani et al., 2008), demonstrating that the use of analogy matters in the model’s performance.

3.2. Implications for scale of analogical reasoning

Table 2 summarizes the properties of the representations used in the experiments above. While this is only a small sample of the set of possible tasks, it includes very different kinds of tasks (visual reasoning, textbook problem solving, and moral reasoning), providing a look at how properties of representations might change across different domains. Table 3 summarizes the statistics of the representations.¹³

One thing that stands out in Table 3 is the large size of the representations. These statistics suggest that in order to model human performance over this range of tasks, the analogical mapping process must be able to handle between 10 and 100 relations. Of course, these representations may be larger than is needed. Our priority has been to use automatically generated representations, from systems developed both in our laboratory and by others. Such systems tend to optimize for compact representations, given that

Table 2
Experiments with sources of representations

Task	# Cases	# Comparisons	Source of Cases
Geometric analogies	299	873	CogSketch
Oddity task	446	3,409	CogSketch
Thermodynamics problems	8	4	AI System
Physics problems	253	1,137	Educational Testing Service
Moral decision-making	41	420	Natural language system

Notes. # Cases indicates the total number of cases generated, across all problems in that task.

Comparisons indicates the total number of comparisons made in each of those tasks. This includes comparisons made in service of analogical retrieval, by the second stage of MAC/FAC.

Table 3
Representation statistics from computational investigations

Task	Entities			Expressions			Relations		
	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max
Geometric analogies	2.5	1	8	24	3	136	16	0	112
Oddity task	3	1	16	29	2	216	20	0	168
Thermo problems	15	8	31	147	57	400	89	25	282
Physics problems	32	4	71	129	5	431	87	2	298
Moral reasoning	16	13	21	45	31	62	31	18	46

Notes. This table shows the mean, minimum, and maximum number of entities, expressions, and relations in the descriptions given to SME in the computational experiments summarized here. Expressions include both relations and non-atomic terms, that is, entities denoted via functional expressions.

large representations tend to raise the cost of other operations as well. However, there is no guarantee that they are the most compact possible representations. But even if our representations are two or three times larger than needed, this still leaves us with 20 or 30 relations for some kinds of problems. How do these figures square with current estimates of online processing capacity?

Let us first consider the case when external representations are present, such as in a geometric analogy problem. In these cases, people may match large descriptions incrementally, relying on the external representations to relieve the working memory burden. This would be consistent with the incremental matching techniques discussed above. Of course, even in this case, people would still have to internalize enough of the ongoing mapping to be able to accurately maintain structural consistency across the solution steps.

But what about cases in which no external representation is present, as in comparing a current situation to a representation in LTM? This brings us to the critical issue of working memory. In some accounts, working memory (WM) is synonymous with short-term memory (STM), as assessed in tasks such as digit span. STM is primarily concerned with short-term information storage and is subject to temporal decay and capacity limits of around 4 chunks at best (Cowan, 2001). Clearly, limits implied by equating WM with

STM are difficult to square with our representational assumptions. However, some recent theories have emphasized the role of WM in holding and manipulating information, including information from LTM (Baddeley, 2012; Ericsson & Kintsch, 1995). For example, Ericsson and Kintsch (1995) propose a distinction between ST-WM and LT-WM. The former is the traditional STM, characterized by strict temporal and chunk capacity limits. In contrast, LT-WM is concerned with accessing and manipulating information in LTM, and varies with people's knowledge and skill. Ericsson and Kintsch review evidence from studies of text comprehension that shows that readers are able to temporarily retain and use amounts of information that are far larger than typical STM limits. Moreover, this temporary store is durable in the face of brief interruptions. Baddeley (2012) agrees that interactions with LTM could boost WM performance, and maintains that this kind of interaction is compatible with current conceptions of WM.¹⁴

In any case, the important point is that Ericsson and Kintsch's comprehensive review shows that domain knowledge is an important determinant of effective WM capacity. They review a wide array of evidence from studies of experts (in chess, bridge, abacus calculation, and even waiting on tables), showing that the amount of material a person can retain and manipulate over short periods is far greater in their domain of expertise than in other arenas. They also review evidence concerning "everyday expertise"—that is, people's ability to encode and retain meaningful text—which shows that the amount of information people can temporarily retain when processing meaningful text is far larger than would be expected from most estimates of capacity limits (and much larger than when processing scrambled text). As they put it, "Domain knowledge provides the retrieval structures that give readers direct access to the information they need when they need it. Given a richly interconnected knowledge net, the retrieval cues in the focus of attention can access and retrieve a large amount of information" (Ericsson & Kintsch, 1995, p. 231).

Ericsson and Kintsch's point that domain knowledge vastly influences the effective capacity of WM accords with work in analogy, which has shown repeatedly that the nature and quality of domain representations is a key determinant of how people process analogies. In particular, Ericsson and Kintsch's emphasis on "richly interconnected knowledge" is consistent with studies showing the importance of systematicity—the presence of higher-order relations¹⁵ that connect lower-order relations—in allowing people to carry out large analogical mappings (e.g., Gentner & Toupin, 1986; Loewenstein & Gentner, 2005). A well-structured domain representation has at least two advantages for analogical processing: (a) it may be treated as a few higher-order nodes and unpacked into more detailed structures as needed; and (b) it may permit incremental processing, because the connecting relations allow the person to keep track of where they are in processing a large representation. This is compatible with the assumptions we made about working memory in the visual processing models described in Section 3. Specifically, the objects mentioned in differences were counted as working memory, rather than every relation involving them.

Such higher-order structure might operate much like the large-scope chunks in cognitive architectures such as SOAR (Laird, 2012) and ACT-R (Anderson, 2007). If so, then

the capacity limit may indeed be the number of coherent substructures, rather than the number of individual predicates.

4. Techniques for scaling up analogical processing

This section describes five techniques that we have found crucial for scaling up analogical processing, in the investigations above as well as others: Greedy Merge, incremental operation, ubiquitous predicates, structural evaluation of candidate inferences, and match filters. We discuss each in turn.

4.1. The GreedyMerge algorithm

Recall that in SME's local-to-global algorithm, local matches are first constructed in parallel, and kernels (structurally consistent combinations of matches starting from non-subsumed match hypotheses) are subsequently combined to form the correspondences of a global mapping. In our first two versions of SME, we used an exhaustive merge process to find all globally consistent solutions. While guaranteed to find optimal solutions, the worst-case performance of such an algorithm is factorial in the number of kernels. Clearly a more efficient search technique is necessary for structural alignment to be widely used as a process within large-scale tasks.

Our solution has been to trade off optimality for performance, by using a greedy algorithm for merging. The essence of a greedy algorithm is this: Suppose one has a set of partial solutions, only some of which are mutually consistent with each other, which must be combined into the best possible global solution. In general, this is an NP-hard problem, that is, requiring exponential growth in resources as the size of descriptions grows, because one must try all combinations of partial solutions together to guarantee optimality. What a greedy algorithm does instead is to select the best partial solution, add to that the next-best partial solution consistent with it, then the next-best partial solution consistent with those two, and so on, until nothing else can be added. This algorithm is linear in the number of partial solutions and yields optimal (or near optimal) answers surprisingly often, as discussed below.

The greedy idea is applied to building mappings in SME as follows: Kernels are the partial solutions, and are rank-ordered by their structural evaluation, that is, the sum of the structural evaluations of the local match hypotheses included in them. The first interpretation is created by starting with the structurally largest kernel and going down the list, merging kernels with it that are not structurally inconsistent with the solution being generated. SME can generate more than one global mapping by selecting the highest ranked remaining kernel and starting the process over again. The maximum number of interpretations returned is a parameter of the simulation, and it defaults to three. There is also a cutoff percentage, such that if a subsequent kernel is going to be significantly smaller than previous solutions it is not generated. Thus, when there is only one

“obvious” mapping, only one is produced. But if there are competing interpretations, up to two additional mappings will be produced as well.

We have subsequently improved on the simple greedy merge algorithm reported in Forbus and Oblinger (1990) in two ways. First, we divided the merge process into two steps. The insight is that, since candidate inferences only arise when base statements are imported in the target, merging kernels that project to a common base root increases the likelihood of finding productive matches. Thus, we first perform a greedy merge operation within kernels whose base projection shares a common root, and then a second round of greedy merging over those solutions to construct mappings. Second, we allow mappings to share kernels. The original greedy algorithm placed each kernel in at most one mapping. However, it is theoretically possible for distinct mappings to share some matches, and thus some kernels. To support this, after a mapping is greedily computed, the algorithm iterates over all kernels from previously computed mappings. If any kernel is consistent with the current mapping, it is added to it. This allows secondary mappings to be better filled out, and because it requires only a single pass through the previous mappings, it adds minimally to the computational cost. The supplemental material provides a complete description of the current GreedyMerge algorithm.

How good are the solutions produced by the greedy merge algorithm? As reported in Forbus and Oblinger (1990), a simple version of GreedyMerge was originally tested on 56 analogies, ranging from comparisons between physical phenomena, short stories, and object descriptions, drawn from a library of SME examples. We found that greedy merge produced the identical best mapping to the original exhaustive merge in 52 of these cases, that is, 93% of the time. We discuss some subsequent similar analyses by other researchers in Section 6.

Why does GreedyMerge normally do so well? Typically, analogies over large descriptions have a few large kernels, only some of which are mutually inconsistent, and a much larger set of small kernels. Thus, the first few decisions are the really critical ones, and they are relatively easy to make. When does GreedyMerge fail? There are two kinds of cases where it can fare poorly. The first is when there are many large kernels with a high degree of mutual inconsistency. In this case, a large number of decisions have to be correct, and hence the chance of error grows. This was the problem in the few cases in our original experiments (4 out of 56) in which a nonoptimal solution was generated. We have also seen nonoptimal solutions in practice, particularly when applying SME to large visual representations of diagrams (e.g., Ferguson & Forbus 2000). There are two solutions to this sort of problem. The first is to increase the number of interpretations one is willing to consider, or prune the set of possibilities via filter constraints, as discussed in Section 4.5. The second is to change the algorithms used to generate representations, to introduce more relevant structure.

Another kind of case where, in principle, Greedy Merge would do poorly is that in which an initial, large kernel is inconsistent with every member of a large set of small but mutually compatible kernels that together outweigh the initial one. We have not yet observed this phenomenon in natural representations. Fortunately, the ability to generate radically different interpretations provides the potential to recover from such problems.

The theory of *submodular functions* provides some useful insights. A submodular function can be thought of as an evaluation function such that the value of adding an additional component to a solution decreases as the size of the solution increases (Cormen, Leiserson, Rivest, & Stein, 2009). Greedy algorithms are optimal when their evaluation function is submodular, a global optimum can be arrived at by selecting a local optimum (the greedy choice property) and an optimal solution contains optimal solutions to sub-problems (the optimal substructure property). Unfortunately, as the cases above illustrate, structural evaluation is not always submodular; hence we know that GreedyMerge cannot always be optimal. However, this can be turned around. It seems possible that human representations are tuned such that structural evaluation tends to be submodular, and when it is not, that is a signal for rerepresentation. A related question is whether there is a set of broad representation conventions over which GreedyMerge is always optimal. These are currently interesting open questions.

4.2. Incremental operation

Many perceptual and cognitive tasks involving structural alignment, including metaphor understanding, problem solving, and learning, require the ability to process information incrementally. For example, when processing an extended analogy or metaphor, readers often build up correspondences across several sentences. This means that each sentence must be linked to the relational structure that has been built up so far. One indication that this happens is that people experience a “mixed metaphor” startle when the metaphoric mapping changes in midstream, as in “The ship of state is boiling over.” To test the idea that people often maintain consistent mappings, Gentner et al. (2001) gave people passages containing extended metaphors, one sentence at a time (see Table 4). The dependent measure was the time required to read each sentence. As predicted by the

Table 4

Sample materials from Gentner et al. (2001). Passages were presented sentence by sentence and time to read each sentence was timed. The key result was that people took longer to read the final sentence in the Inconsistent case than in the Consistent case

Consistent: A Debate is a Race

Dan saw the big debate as a *race*: he was determined to win it. *He* knew that he had to *steer his course* carefully in the competition. His strategy was to go *cruising through* the initial points and then make his move. After months of debating practice, Dan knew how to present his conclusions. If he could only *keep up the pace*, he had a good chance of winning. Before long, he felt the audience was receptive to his arguments. Then, he *revved up* as he made his last key points. His skill left his opponent *far behind him* at the *finish line*.

Inconsistent: A Debate is a War

Dan saw the big debate as a *war*: he was determined to be victorious. He knew that he had to use every *weapon* at his command in the competition. He mapped out his strategy to insure he established a *dominant position*. After months of debating practice, Dan knew how to present his conclusions. If he could only *marshal his forces*, he had a good chance of winning. Before long, he felt the audience was receptive to his arguments. Then, he *intensified the bombardment* as he made his last key points. His skill left his opponent *far behind him* at the *finish line*.

incremental mapping assumption, people took longer to read the final sentence in the Inconsistent case than in the Consistent case (see also Gentner & Boronat, 1991; Gentner, Imai, & Boroditsky, 2002; Thibodeau & Boroditsky, 2011; Thibodeau & Durgin, 2008).

In problem solving, students using a worked example to solve a related novel problem may go back and forth between them, seeking additional ways to interpret the new problem in light of the old. In conceptual change, new data can lead to analogies being modified or abandoned. Modeling these processes requires the ability to incrementally extend a match with new information. Burstein (1986) was the first to computationally model incremental processing in analogical learning, in a domain-specific system for modeling learning to program. Falkenhainer's (1987, 1990) PHINEAS demonstrated that SME could be used in a map/analyze cycle to model the incremental use of analogy in discovering physical theories, albeit with a number of external mechanisms.

The first general-purpose incremental analogical matcher was Keane's IAM (Keane & Brayshaw, 1988). A critical difference between IAM and the technique we have developed for SME concerns serial versus parallel processing. Recall that in SME, processing is essentially parallel within the first two stages: only the final step of constructing global mappings is serial. By contrast, IAM is serial throughout: Even decisions about local matches are made sequentially, so that exploring alternate interpretations requires backtracking. SME avoids backtracking by creating, in parallel, a network representing all local identity matches between items, followed by intermediate clusters (i.e., kernels, introduced in Section 2 and described formally in the supplemental material).

We believe that a combination of initial parallel processing and later serial processing will best model human structural alignment processes. Some serial processing is essential: One cannot combine all information in parallel when not all of it is yet available. However, we believe the fully serial approach of IAM would be difficult to scale up to cognitively plausible representations. Moreover, as noted above, there is now evidence that something like SME's initial symmetric alignment process occurs in human processing, even for strongly directional metaphors (Wolff & Gentner, 2011).

We suggest that the natural place for serial processing is in the Merge step. The kernels represent coherent, structurally consistent collections of local matches, and therefore form a more appropriate unit of analysis for limited-resource serial processing than the individual local matches themselves. When base and target share large systematic structures, the number of kernels is small. Serial, capacity-limited merging of kernels could thus provide a plausible explanation for the "More is Faster" phenomenon whereby additional shared knowledge can improve both the rapidity and the accuracy of mapping (Gentner & Rattermann, 1991; Loewenstein & Gentner, 2005).

The constraint of incremental operation changes the processes of a matcher in several ways. The default operation becomes extending the current set of mappings when new information is added to the base and/or target, instead of starting from scratch. However, one possibility raised by extending existing mappings is that what looked promising initially might turn out to be suboptimal as more information is known. Thus, the matcher must have the ability to remap, that is, to take a fresh look at a comparison. SME does this by discarding the existing mappings and redoing the Merge operation from the set of

kernels. We assume that the criterion for remapping is task-specific; hence, there is no automatic criterion for remapping built into SME. The problem solver described in Section 3.1.3 uses incremental mapping, for example, since the problems can be quite large.

4.3. Ubiquitous predicates

For purposes of matching, not all statements are created equal. Some researchers have suggested identifying specific “important” predicates, such as CAUSE, and focusing on those in matching (e.g., Winston, 1982). The problem with identifying certain predicates *a priori* as important is that it reduces the ability of a match process to be context-sensitive. While causality may be critical in some tasks, such as story understanding, other kinds of knowledge may be critical in different tasks, for example, spatial configurations in navigation. Moreover, if CAUSE is given extra weight, would this also apply to PREVENT and ENABLE, not to mention AID, ABET, HAMPER, and so on? Structure-mapping’s systematicity principle suggests that importance arises from consideration of the size and depth of the overlapping structure between two descriptions, not as a purely local judgment. Moreover, encoding processes that produce the inputs to matching are presumably tuned to what is important in the current context and what is stored in long-term memory and subsequently retrieved is also. Thus, we suggest that predicate-specific heuristics for judging a statement to be important are neither necessary nor sufficient to ensure optimal matching.

Perhaps surprisingly, the reverse turns out to sometimes be very useful: Using a local, predicate-level heuristic for judging a statement to be unimportant—that is, as unlikely to yield useful matches in itself—can greatly simplify the match process. In many domains, there are predicates that occur so frequently that the likelihood of any particular match involving them being useful is low. For example, using analogy to solve thermodynamics problems (Section 3.1.3) can involve matching descriptions that each contain many equations. Each equation could potentially match (because they all involve the predicate =), but most of them will be irrelevant. Another example is the conjunctive connective, and, which is crucial for bundling antecedents together, but is not, on its own, sufficient reason to attempt to form matches. Matches only based on = or and are unlikely to be useful. On the other hand, one cannot simply filter such statements out; they are central components of the domain knowledge. Indeed, matching the appropriate pair of equations can be essential for solving a problem.

In any local-to-global process, having a large number of irrelevant local matches reduces the likelihood of finding strong overall matches. The problem, of course, is that it is impossible to be certain when comparing two statements locally that they will not be useful, since (by systematicity) usefulness arises as a property of an overlapping system of relations.

Our solution to this dilemma is to declare certain predicates common in a domain to be *ubiquitous*. A *ubiquitous predicate* is a predicate which, by itself, does not constitute a sufficient condition to postulate a match between two statements that share it. Which predicates are to be treated as ubiquitous is provided as part of the input to the match process, automatically, by the larger-scale task model on a domain-wide basis. The only difference in how ubiquitous predicates are treated occurs in the initial match hypothesis

construction step. In that step, every pair of expressions is compared and, if they satisfy tiered identity, a match hypothesis is created—unless the predicates are ubiquitous. If the predicates are ubiquitous, the pair is ignored at this step. However, the pair can still appear in a match hypothesis if they are arguments of other statements that are matched. Thus, for example, the = relations in pairs of equations that play similar roles in an explanation or problem solution will have a match hypothesis created for them, by virtue of them being part of the aligned argument structures. Thus, statements involving ubiquitous predicates cannot be matched on their own, but they can be matched as part of a larger matching structure, thereby satisfying the demands of systematicity and parallel connectivity.

The choice of which predicates should be declared ubiquitous must be made with respect to the domain(s), independent of the demands of specific matching tasks or specific problems within the domain. This is an important point theoretically, since it reduces the potential for tailorability in modeling that would be introduced if such decisions were made on an example-specific basis. There are two general guidelines for using ubiquitous predicates in SME-based models:

- 1 Predicates that are basically “joints” in larger structures are not sufficient evidence, by themselves, to warrant matches. Examples include `and`, `TheSet`, and `TheList`. On the other hand, predicates like `cause` and `implies`, despite their high frequency, should never be declared ubiquitous because these form the potential backbone and connective tissue of relational structures. With them, the relevant joints can be found.
- 2 Attributes that are shared by most elements of two descriptions are also good candidates for being declared ubiquitous. For instance, in the representations for sketches used in (Forbus et al., 2011), every depicted object has an associated glyph. Consequently, the attribute of being a glyph is not diagnostic for matching two sketches, whereas spatial relationships involving the glyphs and domain attributes of the objects depicted typically are.

How important are ubiquitous predicates? We illustrate by examining the use of ubiquitous predicates in the five computational investigations outlined in Section 3. Table 5 describes what ubiquitous predicates were used in the SME comparisons performed in each domain. No ubiquitous predicates were used in the two visual tasks because the encoding processes for those models automatically filters out the `Glyph` attribute, which is reasonable because it plays no role in higher-order relations.

To see how ubiquitous predicates affect performance, Table 6 shows statistics on the increased size of the match hypothesis forest for each domain where ubiquitous predicates are used. While ubiquitous predicates had no effect in moral decision-making, they had significant impact on the size of the representations for the problem-solving experiments. This is because in the problem-solving domains most new steps introduce a new equation, so there are potentially many unproductive local matches there. On the other hand, in moral decision-making, the only ubiquitous predicate was `and`, which appears mostly in

Table 5
Ubiquitous predicates by domain

Domain	Ubiquitous predicates
Geometric analogies	None
Oddity task	None
Thermodynamics problems	=, and, nvalue, equation, the-set
Physics problems	and, multipleChoiceSingleOptionList, testAnswers-SingleCorrectMultipleChoice, TheList, TheSet
Moral decision-making	and

Table 6
Match hypothesis forest growth without ubiquitous predicates

Domain	% Matches Bloated	% Increase in No. of Match Hypotheses	
		Mean Bloat	Max Bloat
Thermodynamics problems	100	183	329
Physics problems	100	34	51
Moral decision-making	0	0	0

just one of the cases compared, and typically only once per case; hence, there is little chance for mismatches.

4.4. Structural evaluation of candidate inferences

Candidate inferences can be evaluated along several dimensions. One such dimension is structural quality, and since the comparison operation is based on structural properties, it is the natural place for such evaluations to be computed. (By contrast, utility or logical validity are properties that are task-specific and therefore outside the comparison operation.) This section describes how SME performs structural evaluation of candidate inferences.¹⁶ Structure-mapping theory defines analogical inferences as projections from the base to the target that are structurally supported by the correspondences of a mapping. It also defines reverse candidate inferences from the target to the base, which are also now supported by SME. The ideas in this section apply equally to forward and reverse candidate inferences.

To begin with, we need a little more terminology. Recall that candidate inferences for a mapping are generated by examining how either the base or target intersects the mapping. We call a statement a *root* in a description if it is not the argument of any other statement in that description. Recall that a match hypothesis indicates a potential correspondence between two statements or two entities in the descriptions being compared. The *arguments* of a match hypothesis are the match hypotheses that align the arguments of the statements which that match hypothesis aligns. Thus, if we had

MH1: $(\text{connectedTo } A \ B) \leftrightarrow (\text{connectedTo } C \ D)$

MH2: $A \leftrightarrow C$

MH3: $B \leftrightarrow D$

Then MH1 would have two arguments, MH2 and MH3. By analogy with roots in a description, a match hypothesis is a root if it is not an argument in another match hypothesis.

Consider a statement that is a root of its description. If it participates in the mapping, that is, there is a match hypothesis aligning it with a statement in the other description, then it is part of the overlap between the two descriptions and can provide no additional new information. But if a statement that is a root is not part of the mapping, but it has subexpressions that are, then a candidate inference is computed, to represent the projection of that potential new information into the other description. In Fig. 2, for example, there are two qualitative proportionality statements (q_{prop+} and q_{prop-}) which express causal relationships between continuous parameters. Both of these statements are roots of the base. The q_{prop+} statement indicating that the frequency of the oscillation is positively affected by the spring constant participates in the mapping. By contrast, the q_{prop-} statement indicating the causal connection between the frequency of oscillation and the mass of the block, being an unmapped root with sub-expressions covered in the mapping, serves as a starting point for generating a candidate inference. Similarly, since the *cause* statement, the *block* and *spring* statements, and one of the *part-of* statements are all roots, they too are used to create candidate inferences. The form of the inference is the root expression, with substitutions made as necessary from the correspondences, and with skolem functions introduced for base constants that do not have correspondences. This example has one skolem—namely, that the Earth is made of something like steel (i.e., $[:skolem \ steel]$, in Fig. 5).

Once candidate inferences are generated, how are they to be evaluated? One aspect of evaluating candidate inferences is validity—that is, using other types of reasoning to find out if they are in fact true or false in the target domain. However, this can be expensive, so it is worth providing some estimate of the properties of the inference based on the structural alignment, to serve as a guide to other processes.

The structural evaluation of a mapping provides an estimate of match quality, based on the nature of the overlap. We suggest that a similar structural evaluation occurs psychologically for candidate inferences. However, for candidate inferences we postulate two distinct dimensions:

- 1 *Support*: How much structural support does an analogical inference derive from the mapping that generated it? We estimate this by using the same trickle-down algorithm used in structural evaluation, but limited to the subset of the correspondences that support the inference in the mapping.
- 2 *Extrapolation*: How far does an analogical inference go beyond the support lent by the mapping? We estimate this by using the structural evaluation trickle-down algorithm within the candidate inference, taking the ratio of the score for the structure

outside the mapping over the sum of the score for entire inference (i.e., inside and outside).

The methods for computing these scores are described in more detail in the supplemental material.

We believe these two measures have significantly different functional roles. Support is like the structural evaluation of mappings: More is always better. Extrapolation is more complex: High extrapolation seems desirable in tasks like brainstorming or theory generation, but low extrapolation may be preferable for within-domain comparisons involving highly familiar situations.

That people are able to identify which inferences follow from a given set of correspondences has been demonstrated experimentally (Clement & Gentner, 1991; Spellman & Holyoak, 1996). For example, Markman (1997) found that analogical inferences follow structural consistency, even when there are multiple possible mappings (see also Clement & Gentner, 1991). Our model of candidate inferences, which only computes them from structurally consistent mappings, is consistent with these results.

Our definition of support score is consistent with several lines of evidence. Psychologically, matches involving larger systems of statements are viewed by subjects as more sound (Gentner et al., 1993). As noted above, Clement and Gentner (1991) showed that subjects made predictions based on statements connected to a common antecedent in the base, and that candidate inferences connected to systematic base structures are preferred to those which are not.

Similarity has been suggested as a central process in induction tasks (Heit & Rubinstein, 1994; Lassaline, 1996; Osherson, Smith, Wilkie, Lopez, & Shafir, 1990), so it is useful to see how this model fits with these studies. Lassaline (1996) asked subjects to rate the similarity of pairs of fictitious animals and the inductive strength of a property inference (i.e., if A has X, W, and Z, and B has X, Y, and Z, how likely is it that A has Y?, where A and B were fictitious animals and the rest of the variables were filled in with properties such as “dry flaky skin” or “attacks of paranoia”). Adding a relation in the base that explained the inferred property (i.e., telling the subjects that in B, X causes Y while leaving the description of A unchanged) increased inductive strength, but adding a relation that was not connected to the inference did not. A simple model of this task is to treat it as analogical mapping, with animal B serving as base and animal A as the target, and treating inductive strength as a function of the candidate inference support score. Using these assumptions, a simulation of her experiments using SME also yields this result (Forbus et al., 1997).

Because of the systematicity principle, when there are multiple possible inferences from the base to the target, SME predicts that the inferences will be governed by where the greatest structural commonality can be found. This fits with findings by Heit and Rubinstein (1994), who found that people make stronger inferences about whether one animal has a property based on another animal’s having it when the kind of property to be inferred (anatomical or behavioral) matches the kind of similarity between the animals (anatomical or behavioral). For instance, people judge the

likelihood that whales travel shorter distances in extreme heat to be higher when told that tuna do, relative to when they are told that bears do, presumably because whales and tuna have greater behavioral overlap (both swim) even though whales and bears match better anatomically (both mammals). These findings are consistent with the prediction that people prefer to make inferences with greater support from the common structure.

4.5. Match filters

One strength of SME is that it can detect unanticipated structural correspondences. This is essential for capturing the generativity of human analogical reasoning. However, some tasks impose particular constraints on analogies. When reading an explanation of the greenhouse analogy for warming in the Earth's atmosphere, for example, the reader's interpretation of that analogy must include a correspondence between the greenhouse glass and the atmosphere. Particular correspondences are often specified in instructional analogies (Barbella & Forbus, 2011; Richland, Zur, & Holyoak, 2007). For example, when learners are given the hydraulic analogy for electricity, they are typically told the correspondences (water flow \rightarrow electric current, pressure \rightarrow voltage, etc.) Further, in learning a new domain by analogy, for example, understanding how to solve rotational dynamics problems by analogy with linear dynamics problems, the mapping between the domains is often built up incrementally across multiple problems (Klenk & Forbus, 2013). These incrementally constructed domain mappings are then used in new analogies when solving problems that extend the system's understanding of the domain. In other words, the nature of some tasks require adding additional constraints that the matching process needs to respect.

Since SME is only generating one to three mappings, tasks need to be able to communicate constraints to SME, to influence it toward more useful matches. *Match filters* provide a restricted language for that communication. Match filters are local, based on structural properties of the representation system. By default, SME uses no filters. Importantly, when filters are used, they are automatically imposed by the task model, never by hand.

There are three kinds of match filters. The first is the *required correspondence* filter:

- 1 (`required Bi Ti`): Any mapping that includes a correspondence for either B_i or T_i must map them to each other. The `required` filter is useful when the task imposes correspondences, as in the electricity-water analogy above, or when a speaker declares that "fume hoods are like vacuum cleaners" in an analogy intended to be used by learners (Barbella & Forbus, 2011). They are also used to establish persistent cross-domain mappings (Klenk & Forbus, 2013). We note that `required` constraints are sometimes used in large-scale models, including some of those in Section 3.

Something similar to the `required` filter was first used in PHINEAS (Falkenhainer, 1987). PHINEAS learned new qualitative models for domains by first comparing a

novel behavior to an understood behavior. The analogy between the behaviors introduced correspondences that were used in a second mapping, with the explanation of the first behavior as the base, and the (initially empty) explanation of the novel behavior as the target. The correspondences found in the first mapping were required to hold in the second mapping, which provided the necessary translation of terms to enable the importation of the explanation (as a set of candidate inferences) into the new domain.

The second kind of filter is excluding a pair of correspondences:

2 (`excluded Bi Ti`): No mapping can include a correspondence between `Bi` and `Ti`.

In many tasks, SME is used iteratively—a mapping is generated, inspected by the larger model, and if the mapping is unsuitable, it must look for another. For example, solving geometric analogy problems sometimes requires backtracking and looking for an alternate mapping between two images (Lovett et al., 2009). Hence `excluded` constraints serve a role analogous to `nogoods` in truth-maintenance systems (Forbus & de Kleer, 1993), encoding information about correspondences that, while structurally sound, have been found to be inappropriate for other reasons.

The final kinds of filters are predicate filters:

3 (`identical-functions`): No mapping can include correspondences between non-identical functions. This overrides the usual policy of letting nonidentical functions match whenever suggested by a larger relational structure.

4 (`require-within-partition-correspondences Att1 Att2`): No mapping can include correspondences that map an entity with attribute `Att1` to an entity with attribute `Att2`.

The `identical-functions` filter supports a conservative strategy often used in routine problem solving, and by learners in an unfamiliar domain, when they may reject all but the most certain mappings. In solving physics or thermodynamics problems, for instance, the circumstances under which one can substitute one type of parameter for another are strictly limited. For example, the analogical problem solver in (Ouyang & Forbus, 2006) used the `identical-functions` constraint when using SME to retrieve and apply plans from previously solved problems to handling new thermodynamics problems.

The `require-within-partition-correspondences` filter supports another conservative strategy that is valuable for within-domain comparisons involving complex examples. For instance, people and mountains can both be used in representations as ways of denoting locations, but people can be given tasks whereas mountains cannot. (This example comes from using SME to reason about military tactics problems [Forbus, Usher, & Chapman, 2003].) This class of constraint was also used in solving geometric analogy problems (Lovett et al., 2009), to enforce human preferences for matching identical shape categories in this task.

5. Theoretical and empirical complexity analysis of SME

Recall that the initial phase of SME's processing is assumed psychologically to occur in parallel, constructing a forest of local hypotheses about matches and constraints between them. This includes a local structural evaluation process that uses a propagation algorithm to compute scores for match hypotheses consistent with the systematicity principle. The maximally consistent single-root subsets of this forest are the kernels of the match, which are then combined via a greedy-merge process to create mappings. Assuming that the number of items in base and target is both n , then the worst-case size of the match hypothesis forest is bounded by n^2 . This is also a worst-case bound on the number of kernels. As the supplemental material demonstrates, the theoretical worst-case complexity for SME on a serial machine is $O(n^2 \log[n])$.

Worst-case analyses are notorious for their pessimism, and this analysis is no exception. The best way to demonstrate this is empirically, using examples drawn from the modeling projects in Section 3. As noted earlier, we instrumented the simulations and dumped copies of the base, target, and enough of the vocabulary to enable that SME match to be duplicated later, apart from the rest of the simulation code, to provide more accurate timings. All runs were conducted on a cluster node running Linux, to minimize the number of other processes interfering. This gave us a total of 5,843 comparisons.

The size of the match hypothesis forest plays a central role in determining the overall effort needed in a comparison. If there is substantial overlap, the forest will be large. This may or may not lead to large matches, depending on the degree of structural consistency in the overlap, but certainly a large mapping cannot arise from a small forest. It is instructive to compare the actual size of match hypothesis forests to the theoretical worst-case. Since our worst-case analyses in the supplemental materials are all based on number of items (i.e., expressions plus entities), we can compare the square of this number, which is the worst-case complexity of this operation, to the actual number of match hypotheses and kernels in these mappings.¹⁷ Table 7 shows the results.

In general, the number of match hypotheses is typically well below the worst case, and the number of kernels (which governs the time for the serial processing phase) is even further below the worst case. In other words, on realistic descriptions, SME runs far faster than one would expect from the worst-case analysis. The extreme closeness of the statistics for the two visual reasoning tasks is almost certainly due to the use of CogSketch for automatic encoding in both tasks.

Can the size of descriptions be used to make any empirical prediction of run-time? The answer is no, because the size of the match hypothesis forest depends on the *overlap* between the base and target. There is more potential overlap for large descriptions than small ones, but two medium sized descriptions with many of the same predicates can have more overlap than two larger descriptions whose predicates have little overlap. Thus, there is no simple formula in terms of size alone for estimating run-time accurately. That said, the theoretical complexity analysis captures the growth in resource usage. Empirically, the correlation between $n^2 \log(n)$ and run time on a serial machine is quite

Table 7
 Statistics on sizes of match hypothesis forests and kernels, compared to worst-case analysis

Experiment	Actual Match Hypotheses Versus Worst Case		Actual Kernels Versus Worst Case	
	Mean %	Max %	Mean %	Max %
Geometric analogies	32	100	15	40
Oddity task	30	70	15	40
Thermodynamics problems	3	3.4	0.7	1
Physics problems	2	34	0.5	12
Moral decision-making	4	7	2	4

Notes. Mean % and Max % are the percentages of data structure sizes computed by SME over the computational experiments, compared to the theoretical worst-case sizes, to illustrate how pessimistic the worst-case analysis can be in practice.

high for four of the tasks ($R^2 > 0.9$ for geometric analogies, oddity task and solving problems in thermodynamics and AP physics), and positive albeit lower for moral decision-making ($R^2 > 0.36$). Why the difference for moral-decision-making? Moral decision-making is the only experiment in Section 3 involving cross-domain matches of simple stories. In any case, these correlations suggest that the complexity analysis in the supplemental materials is a reasonable approximation in terms of growth of resources. Just as important, as Table 7 indicates, the actual numbers rarely get close to the worst case (when they do, in geometric analogies, the cases involved are very small), meaning that empirical performance is better than the analysis suggests.

6. Related work

Here, we describe how our research fits into the broader picture of research on analogy and case-based reasoning. The success of the field has led to an increase in the number and variety of models of analogy, so this review is necessarily selective (for more comprehensive reviews, see Gentner & Forbus, 2011; Kokinov & French, 2003). Chronologically, perhaps the earliest computational model of analogy was Evans (1968) system for geometric analogy tests. Winston's (1980) pioneering work considered a broader range of analogies, including both perceptual and conceptual analogies. His work on "near miss" analogies dovetails with our work on alignable differences. While motivated by psychological concerns, these systems were rarely tested against empirical human data.

The original version of SME (Falkenhainer et al., 1986, 1989) was the first simulation of analogical matching shown to be consistent with human similarity ratings (Skorstad, Falkenhainer, & Gentner, 1987) while also being able to produce psychologically plausible candidate inferences. The current version of SME retains those abilities, while adding the ability to generate alignable differences and abstractions and to handle incremental inputs, match filters, and scaling behavior that makes it more plausible as a human model.

Holyoak and Thagard's (1989) ACME followed closely after SME and was partly based on SME. ACME was the first connectionist simulation of analogy. It utilized a multiconstraint approach to analogical mapping; the constraints were structural consistency; similarity between corresponding objects ("semantic similarity"); and pragmatic goal-consistency. ACME used a localist connectionist network to implement SME's match hypothesis forest, and a winner-take-all algorithm to compute the final mapping, taking all three constraints into account. However, ACME had some serious drawbacks (see Hummel & Holyoak, 1997). First, it allowed many-to-one matches, which led to structurally inconsistent mappings. Later studies confirmed structure-mapping's claim that many-to-one matches are psychologically implausible (Krawczyk et al., 2005; Markman, 1997). Second, people are capable of considering more than one mapping, and ACME's winner-take-all algorithm ruled this out. Third, ACME was not capable of generating novel candidate inferences. It could fill in a piece of structure given an explicit suggestion, but could not carry-over inferences spontaneously as people routinely do (Falkenhainer, 1990; Markman, 1997). Thus, it was unable to capture some of the key benchmark features of analogy summarized in Table 1.

Keane and Brayshaw's (1988) IAM operates in an incremental fashion. Like SME, it has explicit structural representations and maintains structural consistency. However, unlike SME, it can process "unnatural analogies" that do not involve semantic commonalities—for example, "Fido eats Kibbles" and "John loves New York." Because IAM hypothesizes matches sequentially, it is highly sensitive to the order in which information is presented. This strong dependence on the order of matches is problematic in light of evidence that some of the object matches made early in processing are later overruled by structural consistency (Goldstone, 1994).

Larkey and Love's (2003) CAB model provides a parsimonious localist connectionist model of matching. CAB uses a simple iterative computation that matches elements in one representation with elements in the other. As in Goldstone's (1994) SIAM, initial matches are governed by local object matches, with structural constraints becoming more important over processing. Larkey and Love (2003) showed that CAB can capture some of the benchmark phenomena for analogical processing (see Table 1). However, it is unable to produce candidate inferences, one of the key benchmark phenomena.

Several early models focused on particular domains, on the assumption that analogical matching is tightly integrated with domain-specific processing. For example, COPYCAT (Hofstadter & Mitchell, 1995; Mitchell, 1993) constructed representations of strings of letters and solved analogies involving them. Similarly, TABLETOP (French, 1995) used processes of encoding and matching to capture how place settings at a table placed in analogy with each other. These models have some interesting features, particularly the attempt to model interleaved processes of encoding and matching. However, their mapping algorithms rely on domain-specific processes. The psychological evidence to date suggests that analogical mapping processes are domain-general. The current preponderance of domain-independent computational models of analogy matching suggests that models that focus on capturing a broader range of behavior in particular domains might do well to use a general-purpose matcher.

The closest problem-solving model to ours is CASCADE (VanLehn & Jones, 1993), which used analogy in problem solving. When solving for a particular quantity in a new problem, CASCADE would look for a step in a prior solution that involved that quantity. We used the same type of strategy in the simple thermodynamics problem solver described earlier, but with a general-purpose model of analogical matching (SME) instead of the special-purpose matcher used in CASCADE. Another contrast is that in our AP Physics experiment, retrieval of prior problems was done using MAC/FAC, allowing us to capture the specific dynamics of analogical retrieval. Another interesting model was PRODIGY-ANALOGY (Veloso & Carbonell, 1993). This was the first use of analogy in a cognitive architecture. PRODIGY treated analogy as a means of replaying problem-solving traces (i.e., *derivational analogy*), and its matcher and retrieval mechanism were specialized for that purpose. By contrast, SME has been shown to be useful for derivational analogy, but also for decision-making and visual problem solving, as the experiments summarized in Section 3 illustrate.

Grootswagers (2013) took a different approach to analogical modeling by conducting explorations of optimality. He constructed a generator that would produce pairs of structures that varied in degree and kinds of match, and used this to look at when a greedy algorithm would produce optimal results. He implemented his own version of the 1990 version of SME as well as an exhaustive version. His greedy code found the same solution as his exhaustive algorithm 87.5% of the time (page 27, *ibid*) and found the same solution 99.89% of the time when run on pairs of plays from the original ACME datasets (page 34, *ibid.*). We find these results encouraging, since they support our claim that greedy merge performs well most of the time. However, there is an important caveat: Our criteria were generating results identical to the exhaustive algorithm, whereas his optimality criteria were to maximize the score. These are subtly different: Our greedy algorithm always produces structurally coherent candidate inferences, whereas Grootswagers' algorithm does not, since it uses as kernels nonmaximal collections of match hypotheses. If two maximal kernels cannot be merged, it is because they are structurally inconsistent. So while including subcomponents of them might increase the match score, the resulting structure used to generate candidate inferences will be structurally inconsistent, which is psychologically implausible (Krawczyk et al., 2005; Markman, 1997).

Grootswagers also suggested that two variants of the original exhaustive SME would be better models than our greedy algorithm. The specific models are (a) van Rooij, Evans, Muller, Gedge, and Wareham (2008), which searches every combination of sets of object matches; and (b) Wareham, Evans, and Van Rooij (2011), which searches every combination of statements. We disagree with this conclusion, since both of those algorithms have factorial complexity (the first in number of entities in the base and target, the second in the number of statements in the base and target). Those disagreements aside, we view Grootswagers' work as a valuable theoretical exploration of the computational properties of matching.

Like SME, the models discussed so far are focused on Marr's top two levels—the computational level (what does the system do) and the process level (what are the

representations and algorithms). Some simulations of analogy also focus on Marr's implementation level—how one might carry out such computations in neural systems. Two prominent models of this type are LISA (Hummel & Holyoak, 1997) and DORA (Doumas et al., 2008). Both integrate retrieval with matching and use a localist model of neural systems based on temporal binding. In both simulations, temporal binding imposes strong constraints on the number of relations that can be considered. LISA's initial representation scheme led to estimates of working memory involving at most two or three relations (Hummel & Holyoak, 1997; Hummel & Holyoak, 2003). This is much lower than that suggested by the task constraints described in Section 3. Even if our estimates turn out to be too high, we note that ordinary everyday analogies often require more than three relations. A further problem is that, as pointed out by Eliasmith and Thagard (2001), the original version of LISA was incapable of handling matches involving higher-order relational structures, one of the benchmark properties of analogy. Even simple causal analogies require dealing with higher-order relations. To overcome this problem, LISA has been extended with a new mechanism, group units. Group units provide a representation of relational structure that is not subject to LISA's normal working memory limitations (Hummel, Licato, & Bringsjord, 2014). We view this evolution of LISA as recognizing that its prior working memory limitations were too restrictive to capture human abilities. Finally, LISA relies on serial ordering of activation of its units, and this activation order can potentially be manually specified by the experimenters for each example. Thus, LISA allows for a high degree of experimenter intervention in the internal operation of the system, making it highly tailorable.

DORA (Doumas et al., 2008) is closely related to LISA, but it is aimed at modeling the discovery of relational representations during cognitive development. It represents binary relations via two unary predicates plus a linking connective. For example, *higher* is composed of the two unary predicates *high* and *low*, plus a unit that is active when both of them are active and which provides the connection between the two of them. DORA's process of connecting unary predicates into relations captures an important process in analogical development: namely, the relational shift, whereby children show the ability to match on object properties before the ability to match on relational structure (Gentner, 1988; Gentner & Rattermann, 1991; Gentner & Toupin, 1986; Richland, Morrison, & Holyoak, 2006).

DORA represents an important attempt to capture how relational representation begins, and it has been used to model a number of other developmental phenomena (Morrison, Doumas, & Richland, 2011). Still, there are some assumptions that can be questioned. First, DORA's constraint on the size of the representations seems unrealistic, even for children. In DORA, the capacity of WM is two or two-and-a-half role bindings, compared to LISA's original 4–5 role-filler bindings. Doumas et al. argue that this very low capacity fits with the finding that preschool children often fail to match on the basis of relational similarity. While this is an appealing feature of DORA, it is hard to see how to reconcile the assumption of a small fixed capacity with the many studies showing that children of the same age can show relational matching if they are led to have relational representations (Christie & Gentner, 2010; Gentner, Anggoro, & Klibanoff, 2011; Gentner

& Rattermann, 1991; Kotovsky & Gentner, 1996; Loewenstein & Gentner, 2005; Son, Doumas, & Goldstone, 2010).

DRAMA (Eliasmith & Thagard, 2001) constructs a localist network on top of fully distributed representations for base and target. DRAMA, unlike LISA, operates autonomously and can handle larger descriptions than LISA can—at least up to a dozen or so propositions. However, DRAMA currently has no mechanism for constructing candidate inferences—a critical benchmark feature for any model of analogy. Further, the distributed representations it uses require a localist network instead to implement match hypotheses and structural consistency relationships.

Finally, we note that to date none of the neurally inspired models has been used as a component in large-scale task models, as SME has, and all rely on hand-coded representations. Given their focus on attempting to work within constraints of biology as they are currently understood, this seems very reasonable. However, exactly how neural systems carry out their computations is still a matter of much debate. The information and process level constraints, on the other hand, are clearer, and provide evidence about what is sufficient and even to some degree what is necessary to carry out a range of human tasks. We hope that exploring how these two sets of constraints can be reconciled will be mutually productive.

7. Conclusions

We have argued that structure-mapping is a core cognitive mechanism, used in a vast range of processes, from categorization to learning and transfer to problem solving, and in everyday reasoning as well as scientific discovery. This means, first, that a simulation of analogical matching must be able to handle the size and scale of representations that are likely to be used in human processing. Second, the analogical matcher should be able to operate in concert with other processes to capture the use of structure-mapping in a range of cognitive arenas (the Integration Constraint; Forbus, 2001). This paper describes the changes we have made to SME to meet this challenge. We described five extensions to SME itself that enable it to deal with analogies found in human tasks and to integrate with other processes:

1. Greedy merge enables SME to rapidly construct one or two near-optimal global interpretations, making it a polynomial-time algorithm. This allows SME to operate successfully on representations that most other existing models cannot handle, and that are closer to the full range of representations likely to be used in human cognition.
2. Incremental matching enables SME to operate in models where information is not all available at once. This enables SME to better model tasks such as problem solving, where analogies are incrementally elaborated.
3. Ubiquitous predicates enable SME to model the varying degrees to which items may suggest alignment. By marking some predicates as ubiquitous, and thereby not

sufficient evidence by themselves to suggest a match, SME can handle complex representations, such as those found in problem solving.

4. Structural evaluation of candidate inferences provide quick plausibility estimates, enabling SME to model aspects of plausibility judgments in analogical inference, including aspects of category induction.
5. Match filters, automatically imposed by task models, enable SME to model the ability to respond to task demands in matching.

The examples described in Section 3 demonstrate that SME can be used as a module in building large-scale models that capture broader aspects of human behavior. We view task constraints as extremely important in evaluating cognitive models, since they bear directly on their ability to explain how people use comparison.

We see two main directions for future research. First, the large-scale simulations outlined in Section 3 are only a start: Exploring the roles of analogy and similarity in a broader range of cognitive processes via large-scale simulation is an enterprise that is just beginning. By making available a robust model of analogical matching, we hope we can encourage others to join us in these investigations.¹⁸ Second, the task demands on analogical matching pose a tough challenge for neural modeling, but such models must be constructed if we are to understand the phenomenon at all three of Marr's levels (computational, process, and implementation).

Acknowledgments

We thank Joe Blass, Morteza Dehghani, John Everett, Ben Jee, Mark Keane, Matthew Klenk, and Kate Lockwood for help with both experiments and the manuscript. This research was supported by the Office of Naval Research, the Defense Advanced Research Projects Agency, the Air Force Office of Scientific Research, and by an NSF Science of Learning Center grant, for the Spatial Intelligence and Learning Center.

Notes

1. Big-O notation is standard in computer science as a way of describing how some property of a computation grows as a function of the size of its inputs; that is, $O(n)$ time means that as the number of input items, here, n , doubles, the run time of the computation will double, up to some multiplicative constant.
2. For this paper, we use infix notation when discussing examples, and use Lisp-style notation when displaying representations used in working systems.
3. In the current implementation of SME, this process is serial (but very fast). However, we posit that it may proceed in parallel in the brain.
4. If the pair cannot be aligned, people will give nonalignable differences, typically by stating a fact about one item and denying it for the other. For example, for

shopping mall/traffic light, responses included “A traffic light tells you when to go, a shopping mall doesn’t” and “You can go inside a shopping mall, you can’t go inside a traffic light.” We suggest that nonalignable differences (unlike alignable differences) do not naturally spring to mind in the course of comparison.

5. The current implementation is serial, but parallel issues are also discussed in the supplemental materials. Lazy evaluation of commutative predicates is discussed in (Ferguson, 2003).
6. This view is not unanimous; for example, Lee and Holyoak (2008) argue that causal analogies require a different mode of processing than other analogies.
7. This kind of record (referred to as a “dehydrated SME file”) contains all of definitions of the predicate vocabulary, base, target, and SME correspondences and mappings needed to “reconstitute” the original results. This corpus of matches, along with SME source code, is included in the supplemental materials.
8. As discussed elsewhere, alternate mappings are only created if their structural evaluation is within 20% of that of the best mapping, with a maximum of three mappings.
9. The model is embedded in the CogSketch executable. Two sketches bundled with the distribution include the original Evans problems, and how to run the model is described in the CogSketch documentation.
10. The Mundurukú differ markedly in language and culture from Americans. For example, their language appears to have only a rudimentary system of numbers (Pica, Lemer, Izard, & Dehaene, 2004), and they have few terms for geometric figures (Dehaene et al., 2006).
11. The complete source code for MARS is included in the supplemental materials, as are the problems and solutions. The matches produced in this experiment, like the others, are in the corpus provided as well.
12. More generally, the Companion cognitive architecture is exploring the hypothesis that analogical reasoning is central to cognition (Gentner, 2003, 2010). It differs from other cognitive architectures in several ways, including that analogical matching, retrieval, and generalization are more primitive than back-chaining in its operations.
13. For the AP Physics experiment, we include only a subset of the problems, that is, those for one transfer condition. This is due to the large number of quizzes in that experiment.
14. “Ericsson and Kintsch (1995) proposed this concept [long-term working memory] in explaining the superior performance of expert mnemonists, going on to extend it to the use of semantic and linguistic knowledge to boost memory performance. They argue that these and other situations utilize previously developed structures in LTM as a means of boosting WM performance. I agree, but I cannot see any advantage in treating this as a different kind of WM rather than a particularly clear example of the way in which WM and LTM interact” (Baddeley, 2012, p. 18).

15. The exact set of linking relations that provides coherence is an open question. Ericsson and Kintsch (1995) discuss six classes of coherence elements proposed by Givón (1992): referents, temporality, aspectuality, modality/mood, location, and action/script. Gernsbacher (1997) proposes the coherence can arise from (at least) referential, temporal, locational, and causal links. Researchers in analogy have emphasized causal relations (including PREVENT and ENABLE) as well as logical relations such as IMPLIES and higher-order spatial relations such as MONOTONIC-INCREASE.
16. This summarizes and extends the presentation in Forbus, Gentner, Everett, and Wu (1997).
17. More precisely, we use [#Entities(Base)*#Entities(Target)]+[#Expressions(Base)*#Expressions(Target)], since entities and expressions can never align.
18. The source code for this version of SME is included in the supplemental materials for this article. There is also a system available, *CaseMapper*, which runs under Windows, Linux, and Macs which provides a cognitive-scientist-friendly interface for using SME and MAC/FAC, as well as a number of classic examples.

References

- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?*. New York: Oxford University Press.
- Baddeley, A. (2012). Working memory: Theories, models, and controversies. *Annual Review Psychology*, *63*, 1–29.
- Barbella, D., & Forbus, K. (2011). Analogical dialogue acts: Supporting learning by reading analogies in instructional texts. Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011), San Francisco, CA.
- Baron, J., & Spranca, M. (1997). Protected values. *Organizational Behavior and Human Decision Processes*, *70*, 1–16.
- Boroditsky, L. (2000). Metaphoric structuring: Understanding time through spatial metaphors. *Cognition*, *75* (1), 1–27.
- Bowdle, B., & Gentner, D. (1997). Informativity and asymmetry in comparisons. *Cognitive Psychology*, *34*, 244–286.
- Burstein, M. (1986). Concept formation by incremental analogical reasoning and debugging. In R. S. Michalski, J. G. Carbonell, & J. M. Mitchell (Eds.), *Machine learning II: An artificial intelligence approach* (pp. 351–369). Los Altos, CA: Morgan-Kaufmann.
- Cassimatis, N., Bello, P., & Langley, P. (2008). Ability, breadth, and parsimony in computational models of higher-order cognition. *Cognitive Science*, *32*, 1304–1322.
- 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.
- Cormen, T., Leiserson, C., Rivest, R., & Stein, C. (2009) *Introduction to algorithms* (3rd ed.). Cambridge, MA: MIT Press.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 1–185.

- Day, S., & Gentner, D. (2007). Nonintentional analogical inference in text comprehension. *Memory and Cognition*, *35*, 39–49.
- Dehaene, S., Izard, V., Pica, P., & Spelke, E. (2006). Core knowledge of geometry in an Amazonian indigene group. *Science*, *311*, 381–384.
- Dehghani, M., Sachdeva, S., Ekhtiari, H., Gentner, D., & Forbus, K. (2009). The role of cultural narratives in moral decision making. In N. A. Taatgen, & H. van Rijn (Eds.), *Proceedings of the 31st annual conference of the cognitive science society* (pp. 1912–1917). Amsterdam, Netherlands: Cognitive Science Society.
- Dehghani, M., Tomai, E., Forbus, K., Iliev, R., & Klenk, M. (2008). MoralDM: A computational model of moral decision-making. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th annual conference of the cognitive science society* (pp. 1410–1415). Washington, D: Cognitive Science Society.
- Dehghani, M., Tomai, E., Forbus, K., & Klenk, M. (2008). An integrated reasoning approach to moral decision-making. *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI)*. Chicago, IL.
- 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.
- Eliasmith, C., & Thagard, P. (2001). Integrating structure and meaning: A distributed model of analogical mapping. *Cognitive Science*, *25*(2), 245–286.
- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, *102*(2), 211.
- Evans, T. (1968). A program for the solution of a class of geometric-analogy intelligence-test questions. In M. Minsky (Ed.), *Semantic information processing* (pp. 271–353). Cambridge, MA: MIT Press.
- Falkenhainer, B. (1987). An examination of the third stage in the analogy process: Verification-based analogical learning. J. McDermott, (Ed.), *Proceedings of IJCAI-87* (pp. 260–263). Los Altos, CA: Morgan-Kaufmann.
- Falkenhainer, B. (1990). A unified approach to explanation and theory formation. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation* (pp. 157–197). San Mateo, CA: Morgan Kaufmann.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1986). The structure-mapping engine. *Proceedings of AAAI-86*. Los Altos, CA: Morgan-Kaufmann.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, *41*, 1–63.
- Ferguson, R. W. (2003). Mapping self-similar structure: Commutative expressions in structure mapping. R. Alterman, D. Kirsch. (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society*. (pp. 396–401). Boston, MA.
- Ferguson, R., & Forbus, K. (2000). GeoRep: a flexible tool for spatial representation of line drawings, *Proceedings of AAAI-2000*. Austin, TX. Available at: Files/QRG_Dist_Files/QRG_2000/Ferguson-Forbus_AAAI_2000.pdf.
- Forbus, K. D. (1984). Qualitative process theory. *Artificial intelligence*, *24*(1), 85–168.
- Forbus, K. (2001). Exploring analogy in the large. In D. Gentner, K. Holyoak, & B. Kokinov (Eds.), *Analogy: Perspectives from cognitive science* (pp. 25–38). Cambridge, MA: MIT Press.
- Forbus, K., & de Kleer, J. (1993). *Building problem solvers*. Cambridge, MA: MIT Press.
- Forbus, K. D., Ferguson, R. W., & Gentner, D. (1994). Incremental structure-mapping. In A. Ram, K. Eiselt (Eds.), *Proceedings of the sixteenth annual conference of the cognitive science society* (pp. 313–318). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Forbus, K. D., Gentner, D., Everett, J. O., & Wu, M. (1997). Towards a computational model of evaluating and using analogical inferences. In M. Shafto & P. Langley (Eds.), *Proceedings of the nineteenth annual conference of the cognitive science society* (pp. 229–234). Mahwah, NJ: Lawrence Erlbaum Associates.
- Forbus, K., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, *19*, 141–205.

- Forbus, K., Klenk, M., & Hinrichs, T. (2009). Companion cognitive systems: Design goals and lessons learned so far. *IEEE Intelligent Systems*, 24, 36–46.
- Forbus, K., Mostek, T., & Ferguson, R. (2002). An analogy ontology for integrating analogical processing and first-principles reasoning. Proceedings of IAAI-02, July.
- Forbus, K. D., & Oblinger, D. (1990). Making SME greedy and pragmatic. In M., Piattelli-Palmarini. (Ed.), *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 61–68). Hillsdale, NJ: Erlbaum.
- Forbus, K., Usher, J., & Chapman, V. (2003). Qualitative spatial reasoning about sketch maps. Proceedings of the Intelligence User Interfaces Conference, 2003, January. Miami, Florida.
- Forbus, K., Usher, J., Lovett, A., Lockwood, K., & Wetzel, J. (2011). CogSketch: Sketch understanding for cognitive science research and for education. *Topics in Cognitive Science*, 3(4), 648–666.
- Forbus, K. D., Whalley, P., Everett, J., Ureel, L., Brokowski, M., Baher, J., & Kuehne, S. (1999). CyclePad: An articulate virtual laboratory for engineering thermodynamics. *Artificial Intelligence*, 114, 297–347.
- French, R. M. (1995). *The subtlety of similarity*. Cambridge, MA: MIT Press.
- 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. (1989). The mechanisms of analogical learning. In S. Vosniadou, & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199–241). London: Cambridge University Press.
- 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. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, 34(5), 752–775.
- 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., & Boronat, C. (1991). Metaphors are (sometimes) processed as domain mappings. Paper presented at the Symposium on Metaphor and Conceptual Change, Meeting of the Cognitive Science Society of Chicago.
- Gentner, D., Bowdle, B., Wolff, P., & Boronat, C. (2001). Metaphor is like analogy. In D. Gentner, K. J. Holyoak, & B. N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 199–253). Cambridge, MA: MIT Press.
- Gentner, D., Brem, S., Ferguson, R. W., Markman, A. B., Levidow, B. B., Wolff, P., & Forbus, K. D. (1997). Analogical reasoning and conceptual change: A case study of Johannes Kepler. *The Journal of the Learning Sciences*, 6(1), 3–40.
- Gentner, D., & Forbus, K. (2011). Computational models of analogy. *WIREs Cognitive Science*, 2, 266–276.
- Gentner, D., & Gunn, V. (2001). Structural alignment facilitates the noticing of differences. *Memory and Cognition*, 29, 565–577.
- Gentner, D., Holyoak, K. J., & Kokinov, B. (Eds.) (2001). *The analogical mind: Perspectives from cognitive science*. Cambridge, MA: MIT Press.
- Gentner, D., Imai, M., & Boroditsky, L. (2002). As time goes by: Evidence for two systems in processing space-time metaphors. *Language and Cognitive Processes*, 17, 537–565.
- Gentner, D., & Kurtz, K. (2006). Relations, objects, and the composition of analogies. *Cognitive Science*, 30, 609–642.
- 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. (1993). Analogy–watershed or waterloo? Structural alignment and the development of connectionist models of cognition. In S. J. Hanson, J. D. Cowan & C. L. Giles (Eds.), *Advances in neural information processing systems*, Vol. 5 (pp. 855–862). Kaufmann: San Mateo, CA. Available at: <http://groups.psych.northwestern.edu/gentner/papers/GentnerMarkman93.pdf>.
- Gentner, D., & Markman, A. B. (1994). Structural alignment in comparison: No difference without similarity. *Psychological Science*, 5, 152–158.

- Gentner, D., & Markman, A. B. (1995). Analogy-based reasoning in connectionism. In M. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 91–93). Cambridge, MA: MIT Press.
- Gentner, D., & Markman, A. B. (1997). Reasoning and learning by analogy: Introduction. *American Psychologist*, *52*, 32–34.
- Gentner, D., & Markman, A. B. (2006). Defining structural similarity. *Journal of Cognitive Science*, *6*, 1–20.
- 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., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, *10*, 277–300.
- Gernsbacher, M. A. (1997). Two decades of structure-building. *Discourse Processes*, *23*, 265–304.
- Givón, T. (1992). The grammar of referential coherence as mental processing instructions. *Linguistics*, *30*(1), 5–56.
- Goldstone, R. L. (1994). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *20*(1), 3–28.
- Goldstone, R. L., & Medin, D. L. (1994). Time course of comparison. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *20*(1), 29–50.
- Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the non-independence of features in similarity judgment. *Cognitive Psychology*, *23*, 222–264.
- Grootswagers, T. (2013) Having your cake and eating it too; Toward a fast and optimal method for analogy derivation, MSc Thesis. Nijmegen, The Netherlands: Department of AI and Psychology, Radboud University Nijmegen.
- Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 411–422.
- Hofstadter, D., & Mitchell, M. (1995). The CopyCat Project: A model of mental fluidity and analogy-making. In D. Hofstadter and The Fluid Analogies Research Group (Eds.), *Fluid concepts and creative analogies* (pp. 205–267). New York: Basic Books (Chapter 5).
- Hofstadter, D., & Sander, E. (2013) *Surfaces and essences: Analogy as the fuel and fire of thinking*. Basic Books.
- Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), *The psychology of learning and motivation*. Vol. 1 (pp. 59–87). New York: Academic Press.
- Holyoak, K. J., & Hummel, J. E. (2000). The proper treatment of symbols in a connectionist architecture. In E. Dietrich and A. Markman (Eds.), *Cognitive dynamics: Conceptual change in humans and machines*, 229–263. Hillsdale, NJ: Erlbaum.
- Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition*, *15*(4), 332–340.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, *13* (3), 295–355.
- Holyoak, K. J., & Thagard, P. (1995). *Mental leaps: Analogy in creative thought*. Cambridge, MA: The MIT Press.
- Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*, 427–466.
- Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, *110*, 220–264.
- Hummel, J. E., Licato, J., & Bringsjord, S. (2014). Analogy, explanation, and proof. *Frontiers in Human Neuroscience*, *6*: 110. doi: 10.3389/fnhum.2014.00867.

- Keane, M. T., & Brayshaw, M. (1988). The incremental analogy machine: A computational model of analogy. In D. Sleeman (Ed.), *Third European working session on machine learning* (pp. 53–62). London: Pitman/San Mateo, CA: Morgan Kaufmann.
- Klenk, M., & Forbus, K. (2007). Cognitive modeling of analogy events in physics problem solving from examples. In D. S. McNamara, & J. G. Trafton (Eds.), *Proceedings of the twenty-ninth annual meeting of the cognitive science society* (pp. 1163–1168). Nashville, TN: Cognitive Science Society.
- Klenk, M., & Forbus, K. (2009). Analogical model formulation for transfer learning in AP Physics. *Artificial Intelligence*, *173*, 1615–1638.
- Klenk, M., & Forbus, K. (2013). Exploiting persistent mappings in cross-domain analogical learning of physical domains. *Artificial Intelligence*, *195*, 398–417.
- Kokinov, B., & French, R. M. (2003). Computational models of analogy-making. In L. Nadel (Ed.), *Encyclopedia of cognitive science*. Vol. 1 (pp. 113–118). London: Nature Publishing Group.
- Kokinov, B., & Petrov, A. (2001). Integrating memory and reasoning in analogy-making: The AMBR model. In D. Gentner, K. Holyoak, & B. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 161–196). Cambridge, MA: MIT Press.
- Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, *67*, 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., Forbus, K., Gentner, D., & Quinn, B. (2000). SEQL: Category learning as progressive abstraction using structure mapping. In L. R. Gleitman & J. K. Joshi (Eds.), *Proceedings of the 22nd annual meeting of the cognitive science society* (pp. 770–775). Austin, TX: Cognitive Science Society.
- Kurtz, K. J., & Gentner, D. (2013). Detecting anomalous features in complex stimuli: The role of structured comparison. *Journal of Experimental Psychology: Applied*, *19*(3), 219–232.
- Laird, J. E. (2012). *The soar cognitive architecture*. Cambridge, MA: MIT Press.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago, IL: University of Chicago Press.
- Larkey, L., & Love, B. (2003). CAB: Connectionist Analogy Builder. *Cognitive Science*, *27*, 781–794.
- Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 754–770.
- Lee, H. S., & Holyoak, K. J. (2008). The role of causal models in analogical inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34* (5), 1111–1122.
- Loewenstein, J., & Gentner, D. (2005). Relational language and the development of relational mapping. *Cognitive Psychology*, *50*, 315–353.
- Lovett, A., & Forbus, K. (2011). Cultural commonalities and differences in spatial problem solving: A computational analysis. *Cognition*, *121*, 281–287.
- Lovett, A., & Forbus, K. (2012). Modeling multiple strategies for solving geometric analogy problems. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the thirty-fourth annual meeting of the cognitive science society* (pp. 34–35). Sapporo, Japan: Cognitive Science Society.
- Lovett, A., Tomai, E., Forbus, K., & Usher, J. (2009). Solving geometric analogy problems through two-stage analogical mapping. *Cognitive Science*, *33*(7), 1192–1231.
- Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.
- Markman, A. B. (1997). Constraints on analogical inference. *Cognitive Science*, *21*(4), 373–418.
- Markman, A. B. (1999). *Knowledge representation*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Markman, A. B., & Gentner, D. (1993). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language*, *32*, 517–535.
- Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, *25*, 431–467.

- Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory & Cognition*, 24(2), 235–249.
- Markman, A. B., & Gentner, D. (2000). Structure-mapping in the comparison process. *American Journal of Psychology*, 113, 501–538.
- Markman, A. B., & Medin, D. L. (2002). Decision making. In D. L. Medin & H. Pashler (Eds.), *Stevens handbook of experimental psychology* (3rd ed). Vol. 2 (pp. 413–466). New York: John Wiley and Sons.
- Marr, D. (1983). *Vision*. San Francisco, CA: W. H Freeman.
- Matuszek, C., Cabral, J., Witbrock, M. J., & DeOliveira, J. (2006). An introduction to the syntax and content of Cyc. In C. Baral (Ed.), *AAAI spring symposium: Formalizing and compiling background knowledge and its applications to knowledge representation and question answering* (pp. 44–49). Menlo Park, CA: AAAI Press.
- McLure, M., Friedman, S., & Forbus, K. (2010). Combining progressive alignment and near-misses to learn concepts from sketches. Proceedings of the 24th International Workshop on Qualitative Reasoning. Portland, OR.
- Mitchell, M. (1993). *Analogy-making as perception: A computer model*. Cambridge, MA: MIT Press.
- Morrison, R. G., Doumas, L. A., & Richland, L. E. (2011). A computational account of children’s analogical reasoning: Balancing inhibitory control in working memory and relational representation. *Developmental Science*, 14(3), 516–529.
- Mostek, T., Forbus, K., & Meverden, C. (2000). Dynamic case creation and expansion for analogical reasoning. In H. Kautz and B. Porter (Eds.), *Proceedings of AAAI-2000* (pp. 323–329). Austin, TX: Menlo Park, CA: AAAI Press.
- Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, 97, 185–200.
- Ouyang, T. Y., & Forbus, K. D. (2006). Strategy variations in analogical problem solving. In *Proceedings of the National Conference on Artificial Intelligence*, 21, 446. Menlo Park, CA.
- Pica, P., Lemer, C., Izard, V., & Dehaene, S. (2004). Exact and approximate arithmetic in an Amazonian indigene group. *Science*, 306, 499–503.
- Posner, M. I., & Mitchell, R. F. (1967). Chronometric analysis of classification. *Psychological Review*, 74(5), 392–409.
- Prasad, L. (2007). *Poetics of conduct*. 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.
- Richland, L. E., Zur, O., & Holyoak, K. J. (2007). Cognitive supports for analogies in the mathematics classroom. *Science*, 316, 1128.
- Ritov, I., & Baron, J. (1999). Protected values and omission bias. *Organizational Behavior and Human Decision Processes*, 79, 79–94.
- van Rooij, I., Evans, P., Muller, M., Gedge, J., & Wareham, T. (2008). Identifying sources of intractability in cognitive models: An illustration using analogical structure mapping. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th annual conference of the cognitive science society* (pp. 915–920). Washington, D.C.: Cognitive Science Society.
- Sagi, E., Gentner, D., & Lovett, A. (2012). What difference reveals about similarity. *Cognitive Science*, 36(6), 1019–1050.
- Skorstad, J., Falkenhainer, B., & Gentner, D. (1987). Analogical processing: A simulation and empirical corroboration. In K. Forbus and H. Shrobe (Eds.), *Proceedings of the Meeting of the American Association for Artificial Intelligence* (pp. 322–326). Menlo Park, CA: AAAI Press.
- Son, J. Y., Doumas, L. A. A., & Goldstone, R. L. (2010). When do words promote analogical transfer? *Journal of Problem Solving*, 3, 52–92.
- Spellman, B. A., & Holyoak, K. J. (1996). Pragmatics in analogical mapping. *Cognitive Psychology*, 31, 307–346.

- Thibodeau, P. H., & Boroditsky, L. (2011). Metaphors we think with: The role of metaphor in reasoning. *PLoS ONE*, 6(2), e16782.
- Thibodeau, P., & Durgin, F. H. (2008). Productive figurative communication: Conventional metaphors facilitate the comprehension of related novel metaphors. *Journal of Memory and Language*, 58(2), 521–540.
- Thibodeau, P. H., & Durgin, F. H. (2011). Metaphor aptness and conventionality: A processing fluency account. *Metaphor and Symbol*, 26, 206–226.
- Tomai, E., & Forbus, K. (2009). EA NLU: Practical language understanding for cognitive modeling. *Proceedings of the 22nd International Florida Artificial Intelligence Research Society Conference*. Sanibel Island, FL.
- Trench, M., & Minervino, R. A. (2014). The role of surface similarity in analogical retrieval: Bridging the gap between the naturalistic and the experimental traditions. *Cognitive Science*, 39, 1292–1319.
- VanLehn, K. (1998). Analogy events: How examples are used during problem solving. *Cognitive Science*, 22(19), 347–388.
- VanLehn, K., & Jones, R. (1993). Better learners use analogical problem solving sparingly. *Proceedings of the 10th International Conference on Machine Learning*.
- Veloso, M., & Carbonell, J. (1993). Derivational analogy in PRODIGY: Automating case acquisitions, storage, and utilization. *Machine Learning*, 10, 249–278.
- Wareham, T., Evans, P., & Van Rooij, I. (2011). What does (and doesn't) make analogical problem solving easy? A complexity-theoretic perspective. *The Journal of Problem Solving*, 3(2), 30–71.
- Weber, E., & Hsee, C. (1999). Cross-cultural differences in risk perception, but cross-cultural similarities in attitudes towards perceived risk. *Management Science*, 44, 9.
- Winston, P. (1980). Learning and reasoning by analogy. *Communications of the ACM*, 23, 689–703.
- Winston, P. (1982). Learning new principles from precedents and exercises. *Artificial Intelligence*, 19, 321–350.
- Wolff, P., & Gentner, D. (2011). Structure-mapping in metaphor comprehension. *Cognitive Science*, 35, 1456–1488.

Supporting Information

Additional Supporting Information may be found online in the supporting information tab for this article:
Data S1. Appendix: SME Complexity analysis, plus source code and corpus.