Cognitive Science 27 (2003) 781–794

COGNITIVE SCIENCE

http://www.elsevier.com/locate/cogsci

CAB: Connectionist Analogy Builder

Levi B. Larkey*, Bradley C. Love

Department of Psychology, The University of Texas at Austin, 1 University Station A8000, Austin, TX 78712-0187, USA

Received 14 July 2002; received in revised form 22 May 2003; accepted 20 June 2003

Abstract

The ability to make informative comparisons is central to human cognition. Comparison involves aligning two representations and placing their elements into correspondence. Detecting correspondences is a necessary component of analogical inference, recognition, categorization, schema formation, and similarity judgment. Connectionist Analogy Builder (CAB) determines correspondences through a simple iterative computation that matches elements in one representation with elements playing compatible roles in the other representation while simultaneously enforcing structural constraints. CAB shows promise as a process model of comparison as its performance can be related to human performance (e.g., solution trajectory, error patterns, time-on-task). Furthermore, CAB's bounded working memory allows it to account for the inherent capacity limitations of human processing. CAB's strengths are its parsimony, transparency of operations, and ability to generate performance predictions. In this paper, CAB is evaluated against benchmark phenomena from the analogy literature.

Keywords: Cognitive modeling; Analogy; Connectionism; Psychology; Artificial Intelligence; Knowledge representation

1. Introduction

The ability to make informative comparisons is central to human cognition (James, 1985). This competence supports tasks ranging from simple perceptual judgments to complex reasoning. This paper presents a model of how humans detect correspondences when making comparisons. The model is intended to be applicable to a broad spectrum of comparisons,

* Corresponding author. Tel.: +1-512-232-5481; fax: +1-512-471-6175.
E-mail address: larkey@mail.utexas.edu (L.B. Larkey).

doi:10.1016/S0364-0213(03)00066-1

though we focus on analogical comparisons here as these comparisons are the most demanding in terms of representation and processing.

The proposed model, Connectionist Analogy Builder (CAB), determines which elements of two representations play compatible roles and places them into correspondence. CAB is a connectionist system that respects the compositionality of structured representations. Like other connectionist systems (e.g., Rumelhart, Hinton, & Williams, 1986), CAB uses a learning rule to continuously update weights among homogenous nodes. CAB displays desirable properties stereotypically associated with both connectionist and symbolic approaches. This synergy is achieved by pairing a constraint satisfaction approach to comparison with directed graph representations that capture variable bindings.

CAB places compatible elements across representations into correspondence using a simple iterative computation that reflects the inherent capacity limitations of the human brain. CAB makes a wealth of processing predictions (e.g., solution trajectory, error patterns, time-on-task, and the effect of working memory capacity), yet is parsimonious (i.e., CAB consists of a small set of equations and parameters, and makes minimal representational assumptions) and its operation is highly interpretable.

The remainder of this paper is organized as follows. First, we discuss some characteristics of human comparison. We then consider CAB's general account of comparison and follow this with a detailed description of CAB's operation. Then, we describe a set of simulations that demonstrate that CAB is consistent with Markman and Gentner's (2000) benchmark phenomena of human analogy. We conclude by comparing CAB to existing models of analogical comparison and by considering how CAB will be further developed.

2. Taxonomy of comparisons

Comparisons vary in their degree of complexity. Simple perceptual comparisons probably constitute the most basic of comparisons. For such comparisons, correspondences between stimulus dimensions are largely predetermined. For example, when comparing simple stimuli, size is compared to size. The size of one stimulus is not put into correspondence with the color of another stimulus. The basis for aligning two representations is predetermined by the dimensional structure of the stimuli. In such cases, stimuli can be represented and compared within metric frameworks like multidimensional scaling (MDS; Shepard, 1962). MDS representations have proven useful in modeling a number of tasks that rely on comparing stimuli, such as category learning (Nosofsky, 1986). Multidimensional representations, in which each stimulus is a point in a metric space, have even been applied to modeling analogy. For example, Rumelhart and Abrahamson's (1973) model completes analogies of the form "*A* is to *B* as *C* is to *blank*" by calculating the vector difference between *A* and *B* and adding it to *C* to compute what *blank* should be.

One problem for MDS techniques is that there is more than one way for two objects to differ (Markman & Gentner, 1993a). Alignable differences are differences that are related to commonalities (e.g., common dimensions). For example, a car and a motorcycle both have wheels but differ in the number of wheels. Non-alignable differences are differences in which a dimension is present for one representation but not the other. For example, the dimension of "restraining device" with value "seat belt" is present for the representation of a car while no corresponding dimension exists in a motorcycle's representation. Alignable differences and non-alignable differences are psychologically distinct. Objects with alignable differences tend to be remembered better, are rated as being more similar, and show advantages in feature listing tasks (Markman & Gentner, 1996, 1997). These results are difficult for MDS models to address. How would an MDS model explain why objects with many alignable differences tend to be similar when such differences should increase the distance between the two representations and therefore reduce similarity? Although feature set approaches (e.g., Tversky, 1977) were developed to address the shortcomings of metric models, these critiques also apply to such models.

These observations suggest that representations do not reside in a common space. Instead, common stimulus dimensions are aligned during processing. However, in some tasks even the proper dimensional correspondences are unclear. Such difficulties are not restricted to abstract representations. In perceptual domains, people often put into correspondence dimensions from different modalities, like pitch and brightness (Marks, 1989). Such non-trivial correspondences are common in analogies involving relations. For example, the solar system and the atom can be seen as similar because the Sun can be put into correspondence with the nucleus because electrons revolve around the nucleus as the planets revolve around the Sun. The relation "revolves" assists in establishing these mappings, which allow for other correspondences to be drawn such as attraction due to difference in charge corresponding to attraction due to difference in mass (Gentner, 1983). Systematic mappings between relational systems can lead to objects being placed into correspondence that are actually quite dissimilar (Gentner & Toupin, 1986; Markman & Gentner, 1993b).

In the remainder of this paper, we discuss CAB. CAB is an account of how people put representational elements into correspondence. Putting elements into correspondence is a necessary component of analogical inference, recognition, categorization, schema formation, and similarity judgment.

3. Description of CAB

CAB's output is a set of correspondences between two representations or analogs. Before discussing how CAB arrives at this output, we will discuss CAB's input. Then, CAB's operation will be qualitatively discussed, followed by CAB's formal description.

3.1. Knowledge representation in CAB

As in previous accounts of comparison (e.g., Tversky, 1977), we assume that other processes have compiled CAB's input. While the construction of analogs is outside CAB's scope, we have developed an approach to representation that conforms to what is currently known, is as simple as possible, makes minimal assumptions, and follows the logic of frame systems (Minsky, 1981). Analogs are represented as directed graphs that can be translated from and to predicate calculus. Fig. 1 depicts an analog in which Jim (a man) loves Betty (a woman). In predicate calculus we can represent this situation as *loves*(*Jim*, *Betty*), *gender*(*Jim*, *male*), *gender*(*Betty*, *female*). Translating between predicate calculus and the graph structure that CAB operates

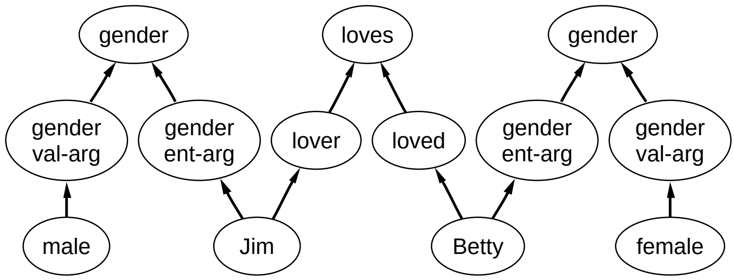

Fig. 1. The analog "Jim loves Betty" can be represented in predicate calculus as *loves*(*Jim*, *Betty*), *gender*(*Jim*, *male*), *gender*(*Betty*, *female*) or by the directed graph shown in the figure.

over is accomplished by first taking the predicates (in this case *loves*, *gender*, and *gender*) and replacing each with one node that denotes the predicate's name. Then, additional nodes are created for each argument of a predicate, and links point from these argument nodes towards the associated predicate node. In the case of *loves* and *gender*, each predicate has two arguments (see Fig. 1). Finally, each entity and value is instantiated as a node in the graph, and links point from these nodes to the arguments they bind to. In the present case, the nodes *Jim*, *Betty*, *male*, and *female* are created. Even though *Jim* appears in two predicates, only one *Jim* node is created because it is the same Jim that is a lover and a male. In contrast, there are two *gender* nodes because one *gender* node describes Jim's gender while the other *gender* node describes Betty's gender. Although it does not occur in this example, predicates can bind to arguments in other predicates. For example, in the analog "John knows that Jim loves Betty," the predicate *loves* binds to the argument *known* in the *knows* predicate.

Our approach to representation captures the distinction between alignable and non-alignable differences. An alignable difference arises from mismatched values on a common dimension, whereas a non-alignable difference arises from a dimension that is present for only one analog. For example, male versus female is an alignable difference arising from the shared predicate of gender (e.g., *gender*(*Jim*, *male*) and *gender*(*Betty*, *female*)). A non-alignable difference between two analogs occurs when one analog lacks a predicate that the other analog has (e.g., a glass of water does not have a gender). Another virtue of this representational approach is that it allows for knowledge to be represented unambiguously. For example, the graph in Fig. 1 distinguishes between the analogs "Jim loves Betty" and "Betty loves Jim." To represent "Betty loves Jim," *Jim* would point to the *loved* node and *Betty* would point to the *lover* node. The ability to explicitly encode relations allows for unambiguous representations, which are necessary for analogical comparison.

The directionality of the links is critical to representing knowledge as it distinguishes between arguments and predicates. This directional information will also prove critical in establishing analogical mappings between two analogs (as will be discussed shortly). The directionality of the links specifies paths between nodes within an analog. For example, the chain (or path) from the node *male* to the node *Jim* is $(+, +, -, -)$, where $+$ and $-$ denote whether a "hop" in the traversal from *male* to *Jim* is with or against the direction of the arrow. The length of this chain is four. The directed graph representation of each analog (e.g., Fig. 1), along with all acyclic chains between node pairs within an analog, compose CAB's input.

3.2. Processing in CAB

CAB iteratively constructs correspondences between nodes that play compatible roles within the two analogs (referred to as A and B). The role of a node is determined relative to the other nodes within its analog through consideration of chaining information (discussed above). This information, along with the mapping weights from the current iteration, is used to determine which nodes from A and B play compatible roles. Like other models of analogy such as SME (Falkenhainer, Forbus, & Gentner, 1989), CAB strongly prefers mappings that are one-to-one (i.e., one node in analog A corresponds to one node in analog B).

Early in processing, CAB is sensitive to any commonalities between the two analogs (i.e., nodes that have identical names). This semantic influence on mapping is captured by establishing mapping weights of size κ (a model parameter) between identical nodes at the start of the mapping process. For example, if both analogs A and B have a *loves* node, then an initial mapping weight would link these nodes. If each analog had two *loves* nodes, then four initial mapping weights would inter-link the four nodes. From this starting point, CAB iterates and learns and unlearns mapping weights.

Mapping weights tend to increase for interconnected node pairs participating in parallel structures. Such corresponding nodes play similar roles in their analog's representation. For example, in Fig. 2 the mapping weight between a_2 and b_2 votes for increasing the mapping weight between a_3 and b_3 because a_2 bears the same relationship to a_3 (i.e., $(+)$) as b_2 does to b_3 (i.e., $(+)$). In this case, the "child" maps to the "child" because the "parent" maps to the "parent." In contrast, the mapping between a_4 and b_5 does not vote for increasing the weight between a_3 and b_3 because the relationship of a_4 to a_3 (i.e., $(-)$) is different than the relationship of b_5 to b_3 (i.e., $(-, -)$). CAB also considers all other relationships (continuing the family analogy) including brother/sister, grandchild/grandparent, nephew/aunt, second cousins, etc. Evidence for increasing a mapping weight depends on the distance of such relationships. For example, the mapping weight between a_1 and b_1 also votes for increasing the weight between a_3 and b_3, but not as strongly as the more immediate pairing of a_2 and b_2.

The ability to appreciate distant relationships allows CAB to display a preference for mappings that exhibit systematicity (Gentner, 1983, 1989), i.e., mappings that preserve deep

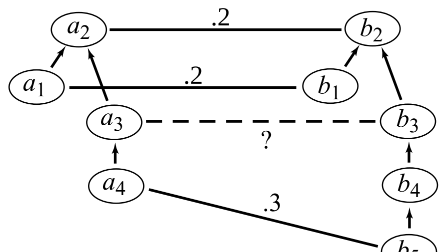

Fig. 2. Nodes a_1 through a_4 belong to analog A, and nodes b_1 through b_5 belong to analog B. There are mapping weights of sizes .2, .2, and .3 between nodes a_1 and b_1, a_2 and b_2, and a_4 and b_5, respectively. The correspondences between a_1 and b_1 and a_2 and b_2 vote for increasing the mapping weight between a_3 and b_3, but the correspondence between a_4 and b_5 does not.

systems of relations. However, as discussed above, relationships that are more immediate (i.e., that involve short chains) are more heavily weighted. The extent to which distant relationships are considered is governed by a model parameter. One way to view this parameter is as specifying working memory capacity because it governs how much information is considered simultaneously when establishing mapping weights (see Hummel & Holyoak, 1997, for a similar view of working memory and its effects on mapping). Considering immediate relationships establishes parallel connectivity (see Gentner, 1983, 1989).

In addition to this weight growth process, weights can also shrink through competition with other mapping weights. Each iteration, evidence for the growth process is filtered through a one-to-one constraint. In addition to this evidence constraint, mapping weights compete with one another in order to establish one-to-one mappings, with the smaller weights bearing the brunt of the competition. A final mapping is established when mapping weights stabilize to either 0 or 1. These final mapping weights constitute CAB's output. Because the mapping process is incremental, partial results can be inspected at any stage in processing.

3.3. CAB's formalism

This section describes the equations that follow from the qualitative description of CAB provided in the previous sections. The input to CAB is the directed graph representation of each analog along with all acyclic chains between node pairs within each analog, which specify the relative roles of each node. The nodes forming the two analogs are denoted as $A = \{a_1, a_2, \ldots, a_M\}$ and $B = \{b_1, b_2, \ldots, b_N\}$. $C(a_i, a_k)$ denotes the set of all acyclic chains of bindings from a_i to a_k and $C(b_j, b_l)$ denotes the set of all acyclic chains of bindings from b_j to b_l. Initially, all mapping weights from nodes in A to nodes in B are set to 0, except for nodes that are identical, which are set to an initial value determined by the parameter κ.

Analogy involves placing into correspondence elements that play similar roles. Because the role of a node is determined relative to other nodes within its analog, CAB first computes the compatibility or pertinence of each node pair in analog A to each node pair in analog B. The pertinence of nodes a_i and a_k in analog A to nodes b_j and b_l in analog B is given by

$$f(a_i, b_j; a_k, b_l) = \sum_{c_m \in C(a_i, a_k)} \sum_{c_n \in C(b_j, b_l)} s(c_m, c_n)\, \mathrm{e}^{-\gamma(L(c_m)-1)}, \tag{1}$$

where the function $s: C(a_i, a_k) \times C(b_j, b_l) \to \{0, 1\}$ is defined such that for all $c_m \in C(a_i, a_k)$ and $c_n \in C(b_j, b_l)$, $s(c_m, c_n)$ is 1 if the directions of the bindings in c_m match the directions of the bindings in c_n, and 0 otherwise. In other words, node pairs are pertinent to each other when they have identical chains (e.g., $(+, -, +)$ and $(+, -, +)$). When node pairs define several chains, each pair of identical chains contributes to the pertinence of the node pairs. $L(c_m)$ is the length of the chain c_m and γ is a parameter related to working memory capacity. As γ increases, only node pairs with relatively short connecting chains are considered pertinent to each other. The exponential function dictates that proximal relations weigh more heavily than distant relations. This exponential term bears a strong resemblance to work in stimulus generalization (Shepard, 1987).

The following three equations iterate until mapping weights stabilize to either the minimal value of 0 or the maximal value of 1. Raw evidence for correspondence between a_i and b_j is collected according to

$$R(a_i, b_j) = \sum_{a_k \in A} \sum_{b_l \in B} f(a_i, b_j; a_k, b_l) m(a_k, b_l), \tag{2}$$

where $m(a_k, b_l)$ is the mapping weight from node a_k to node b_l. That is, raw evidence is collected using all pertinent node pairs weighted by the current mapping weights.

The following equation applies a strong version of the one-to-one constraint to the raw evidence:

$$E(a_i, b_j) = \begin{cases} R(a_i, b_j), & \text{if } R(a_i, b_j) = \max(\{R(a_i, b_l) : b_l \in B\} \\ & \qquad \cup \{R(a_k, b_j) : a_k \in A\}) \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

In other words, the raw evidence for mapping $m(a_i, b_j)$ is filtered (set to 0) unless it is dominant from both a_i's and b_j's perspectives.

Finally, correspondences between the analogs evolve according to

$$\begin{aligned} \Delta m(a_i, b_j) = \alpha E(a_i, b_j) - \beta \max(&\{m(a_i, b_l) : b_l \in B, b_l \neq b_j\} \\ &\cup \{m(a_k, b_j) : a_k \in A, a_k \neq a_i\}), \end{aligned} \tag{4}$$

where α is the growth rate and β is a parameter governing the severity of weight competition arising from the one-to-one constraint. All mapping weights are simultaneously updated and then truncated to be between 0 and 1. CAB repeatedly cycles through Eqs. (2)–(4) until all mapping weights stabilize to either 0 or 1.

The computational complexity of CAB depends on the structure of the analogs as well as their size. If each node binds to every other node within an analog, then the number of chains increases factorially with the number of nodes; however, such analogs are not psychologically tenable. In practice, analogs are sparsely connected and the actual number of chains is usually quadratic in the number of nodes in the analog. Interestingly, the chains do not need to be encoded during the comparison process since they involve only node pairs within an analog, rather than across analogs. Thus one could view the encoding process as part of knowledge acquisition occurring on a longer timescale than comparison processes.

Calculating f for all pairs of node pairs is $O(M^2N^2)$, assuming the number of chains between a node pair is $O(1)$.[1] Each iteration, calculating R for all mapping weights is $O(M^2N^2)$, although this can be substantially reduced by taking advantage of the sparsity of f. Filtering the raw evidence and updating the mapping weights (see Eqs. (3) and (4)) is $O(MN \max(M, N))$.

4. Simulations

The simulations described here showcase CAB's behavior and promise as a process model of human comparison. The simulations also verify that CAB can account for Markman and Gentner's (2000) benchmark phenomena of human analogy.[2] The following parameter values were chosen for all simulations: $\alpha = .001$, $\beta = .001$, $\gamma = .1$, and $\kappa = .1$.

Table 1
Matching by feature similarity

Analog *A*	Analog *B*
gender(Jim, male)	gender(Bill, male)
gender(Betty, female)	gender(Cindy, female)

The analogs for the first simulation are shown in Table 1. In this simulation, CAB places the two males and two females into correspondence, along with their associated gender information. This simulation demonstrates that CAB can process simple comparisons that do not involve relations between entities.

The first simulation demonstrated that CAB is sensitive to features shared by entities. The second simulation demonstrates that CAB is sensitive to relational roles and that this relational information can override feature similarities. In this simulation, the relation *loves* governs the mapping between the analogs "Jim loves Betty" and "Cindy loves Bill." Table 2 shows the predicate calculus representation of both analogs, while Fig. 1 shows the equivalent directed graph representation for analog A. In contrast to the first simulation, the correspondences reverse and *Jim* maps to *Cindy* and *Betty* maps to *Bill*. The respective gender related nodes also map consistently (e.g., Jim's *gender* node maps to Cindy's *gender* node). This simulation highlights Markman and Gentner's (2000) benchmarks of relational similarity and structural consistency.

Interestingly, when the γ parameter is raised, a more circuitous solution trajectory is observed. The γ parameter can be viewed as a proxy for working memory as it governs how many pieces of information are simultaneously considered (i.e., the degree to which distant relations are considered). When γ is raised, the influence of distant relations is reduced and CAB iterates through an inconsistent state in which the entity correspondences respect the relational structure, but the gender related nodes map according to superficial similarity (e.g., analog A's *female* node maps to analog B's *female* node, as do the respective *gender* nodes). After some time, CAB arrives at a consistent, relationally driven solution. The above result suggests that limiting working memory resources should increase the likelihood of inconsistent solutions and lead to longer response times.

The previous simulation involved bringing into correspondence entities that were somewhat dissimilar to one another. In cross-mappings, objects that are similar play different roles in matching relational structure (Gentner & Toupin, 1986; Markman & Gentner, 1993a). For example, Markman and Gentner (1993a) asked subjects to form correspondences between two

Table 2
A simple relationally driven comparison

Analog *A*	Analog *B*
gender(Jim, male)	gender(Bill, male)
gender(Betty, female)	gender(Cindy, female)
loves(Jim, Betty)	loves(Cindy, Bill)

pictures, one which depicted a car towing a boat and another which depicted a truck towing a car. People generally prefer the relationally consistent interpretation (e.g., the car in the first analog matching the truck in the second analog), but can also appreciate the alternative mapping that places the two cars into correspondence (Markman & Gentner, 1993b). Markman and Gentner (2000) suggest this ability to entertain multiple interpretations is a benchmark of analogy.

CAB generates either the relation driven or the object similarity driven mapping between these analogs depending on differences in knowledge representation. When the features of the objects are stressed, CAB generates the object similarity driven mapping, but when the object encodings are not rich, CAB maps objects based on relational roles. The bifurcation point for the chosen parameter values is three or more common feature dimensions for each type of object (car, boat, and truck) with each object displaying a unique value on each dimension (e.g., *color*(*car*, *red*), *color*(*truck*, *green*), and *color*(*boat*, *blue*)). This interpretation of cross-mapping is consistent with Stilwell and Markman's (2001) recent work on packing and unpacking of mental representations—when objects are construed as symbols (i.e., features are not stressed), relation driven mappings abound. CAB also respects the distinction between alignable and non-alignable differences. In the above simulation, three or more alignable differences between objects resulted in object correspondences based on object similarity, while a follow-up simulation revealed that just one non-alignable difference per object (e.g., *bilge-pump*(*boat*, *electric*), *hoist*(*truck*, *pneumatic*), and *child-restraint*(*car*, *backseat*)) yielded the same object similarity interpretation. This result is consistent with the idea that alignable differences are rooted in commonalities.

One challenge for any model of higher-order cognitive processing is scaling to the pinnacles of human performance. The following simulation, borrowed from Falkenhainer et al. (1989), is a complex analogy between the solar system and the Rutherford model of the atom. The analogs (detailed in Table 3) require 42 and 33 nodes to represent the solar system and atom, respectively. CAB correctly maps the Sun to the nucleus and the planets to the electrons. CAB uses higher-order relations to disambiguate the possible correspondences. For example, CAB appreciates that differences in mass (which attract the planets towards the Sun) are analogous to differences in charge (which attract the electrons towards the nucleus), whereas differences in temperature between the Sun and the planets is irrelevant to the analogy. This example displays Markman and Gentner's (2000) benchmark of systematicity. Mappings that preserve large systems of interconnected relations are preferred over mappings that put small, disjointed systems of relations into correspondence. CAB's ability to appreciate distant relations (when γ is not too restrictive) leads to its preference for systematic mappings.

One long-term goal for CAB is to enable progress in understanding how comparisons are processed on-line. In the final set of simulations, we evaluate the possibility of using CAB to make response-time predictions. CAB is applied to three versions of Falkenhainer et al.'s (1989) analogy between water flow and heat flow. The water flow analog involves a scene in which water flows from a large beaker filled with water through a pipe into a smaller vial because of a pressure difference. The heat flow analog involves a melting ice cube attached to a silver bar resting in a cup of hot coffee. The heat flows from the coffee through the bar to the

Table 3
A complex analogy between the solar system and the Rutherford model of the atom

Solar system
mass(Sun, mass-of-Sun)
mass(planet, mass-of-planet)
causes(and(greater(mass-of-Sun, mass-of-planet), attracts(Sun, planet)), revolve(planet, Sun))
causes(gravity(mass-of-Sun, mass-of-planet), attracts(Sun, planet))
temperature(Sun, temperature-of-Sun)
temperature(planet, temperature-of-planet)
greater(temperature-of-Sun, temperature-of-planet)
Rutherford model of the atom
charge(nucleus, charge-of-nucleus)
charge(electron, charge-of-electron)
causes(opposite-sign(charge-of-nucleus, charge-of-electron), attracts(nucleus, electron))
revolve(electron, nucleus)
mass(nucleus, mass-of-nucleus)
mass(electron, mass-of-electron)
greater(mass-of-nucleus, mass-of-electron)

ice cube because of the temperature difference. Table 4 shows the exact representation of the analogs for the baseline version.

Two other versions of these analogs were created. In one version, distracting information is added (i.e., *form*(*water*, *liquid*) and *form*(*coffee*, *liquid*)) that incorrectly suggests that *water* should map to *coffee*. One sensible prediction is that this analogy should take longer to interpret. Indeed, CAB requires more computation to correctly solve this analogy. The third version of the analogs involves removing the causal relations from the baseline analogs. Causal information can be helpful in integrating information about new domains (Murphy & Allopenna, 1994), and in this case the causal relations yield more systematic mappings. In the absence of causal relations, CAB again requires more computation than in the baseline case to correctly form correspondences between the two domains. In the case of adding distracting information, more information slows convergence, while in the case of adding causal information that highlights the common structures, more information speeds convergence.

Table 4
Water/heat flow baseline analogs

Water flow
pressure(beaker, pressure-of-beaker)
pressure(vial, pressure-of-vial)
cause(greater(pressure-of-beaker, pressure-of-vial), flow(beaker, vial, water, pipe))
Heat flow
temperature(coffee, temperature-of-coffee)
temperature(ice-cube, temperature-of-ice-cube)
cause(greater(temperature-of-coffee, temperature-of-ice-cube), flow(coffee, ice-cube, heat, bar))

5. Model comparisons

In this section, we compare CAB to existing models of comparison. Space requirements prohibit consideration of all extant models. In its current state, CAB is solely a model of detecting correspondences. Many of the models discussed below are more mature than CAB, have been applied to other tasks, and have incorporated other task constraints. These differences will not be discussed, but it should be noted that CAB's design does not prohibit it from being extended in these directions.

Like CAB, Hummel and Holyoak's (1997) LISA model of analogy is connectionist and stresses the importance of working memory capacity limitations. LISA utilizes the synchronous firing of nodes to encode node bindings and only forms mapping weights for propositions currently active in working memory. In terms of Marr's (1982) levels of analysis, LISA straddles the algorithmic and implementational levels, whereas CAB is squarely positioned at the algorithmic level. In CAB's favor, CAB is much more parsimonious and transpicuous than LISA. LISA has an elaborate control structure, complex dynamics, numerous parameters, types of nodes, and special conditions. LISA's description includes 21 equations and 22 free parameters. In contrast, CAB is specified by four equations, has four parameters, and utilizes a common node type. Furthermore, Hummel and Holyoak (1997) state that it is unclear whether LISA can scale to processing analogies involving large representations such as Gentner, Ratterman, and Forbus's (1993) "Karla the hawk" analogy.[3] Though not discussed in Section 4, CAB can successfully simulate the "Karla the hawk" analogy using the analog descriptions in Falkenhainer et al. (1989).

The original version of SME (Falkenhainer et al., 1989) constructs the set of all possible structurally consistent correspondences and chooses the solution that displays the greatest degree of systematicity based on an evaluation function. More recent versions of SME (Forbus & Oblinger, 1990) approximate this goal in a less computationally costly fashion. In terms of Marr's levels, SME is somewhat more abstract than CAB and straddles the computational and algorithmic levels. Unlike CAB, SME does not form tentative mappings at every time step but instead merges all results in the final phase of computation. SME does not include capacity constraints and may not be applicable to predicting response time or dual-task data. While SME and CAB are motivated by similar ideas about the output of human comparison (e.g., Gentner, 1983, 1989), we hope that CAB can distinguish itself by accounting for processing data and motivating future experiments investigating processing.

Like CAB, Holyoak and Thagard's (1989) ACME uses a parallel constraint satisfaction method to iteratively construct a mapping between analogs. Although CAB and ACME posit similar mapping processes, the models differ in how they incorporate structural constraints. ACME treats structural constraints, such as one-to-one mapping, as soft constraints that may be violated. While ACME can form many-to-one mappings, humans (and CAB) do not (Markman, 1997). CAB and ACME also differ at the semantic level. ACME uses a supplied similarity table to determine initial node compatibilities, whereas CAB uses strict identicality. With respect to structural constraints and semantics, CAB has more in common with SME than ACME. Unlike CAB, ACME has no mechanism for limiting working memory capacity.

Keane and Brayshaw's (1988) IAM is intended to address the incremental nature of comparison. It holds that people hypothesize matches and attempt to build a consistent set of mappings from these initial starting points. CAB also incrementally develops a set of mappings, but not in the same fashion as IAM. CAB operates over the entire structure instead of selecting substructures. It would be trivial to modify CAB to operate like IAM by seeding CAB with initial match hypotheses corresponding to IAM's starting points. Nonetheless, data suggests that the mapping process is driven by semantic commonalities early in processing and reflects structural constraints in time (Goldstone, 1994). IAM may also be too powerful in its ability to map unnatural analogies that do not involve semantic commonalities and that require psychologically unrealistic working memory capacity (cf., Hummel & Holyoak, 1997). Keane (1997) has used IAM to investigate the effects of order and causal structure on analogical mapping. IAM, like CAB, predicts that causal structure will facilitate performance.

CAB bears a resemblance to Goldstone's (1994) SIAM model in that both CAB and SIAM iterate towards a final mapping solution and that initial processing is guided by superficial similarities. However, SIAM was intended as a performance model of simple perceptual comparisons and cannot be extended to handle more complex cases involving relations (e.g., the solar system/atom analogy). A recent model by Goldstone and Rogosky (2002) named ABSURDIST bears a stronger resemblance to CAB. ABSURDIST attempts to map points in one space to points in another space by considering pairwise distances between points in each space. The spirit of this procedure is commiserate with CAB's search for nodes playing compatible roles across analogs. CAB differs from ABSURDIST in that complex comparisons involve structured representations which are not metric. ABSURDIST is not equipped to process the directed graph representations demanded by comparisons that involve parts playing roles in a larger structure. Nevertheless, the similarities of the algorithms suggest a continuity between the models that we take as a favorable sign for the general approach.

6. Summary, conclusion, and future directions

Comparison involves determining correspondences. In some cases, such as simple perceptual comparisons, determining correspondences can be relatively straightforward. However, many other comparisons, especially those that involve matching relational structures, can be non-trivial.

CAB determines correspondences through an iterative computation that incrementally adjusts mapping weights by considering each node's role within an analog and the current mapping weights. CAB has a compact formalism consisting of a rule for updating mapping weights. CAB's knowledge representation is general and makes few assumptions. Analogs are represented as directed graphs that capture the binding relationships between representational elements. All nodes are of the same type. Other models have special types of nodes (e.g., entity, value, function, relation, binding, semantic, proposition, etc.) that demand special consideration during processing.

CAB can account for the key analogy phenomena identified by Markman and Gentner (2000). CAB also shows promise as a process model of comparison. More research in comparison is needed to address process as opposed to product (cf., Love, Rouder, & Wisniewski, 1999). CAB's performance can be mapped onto response-time measures by examining the amount of computation required for convergence. CAB also produces sensible intermediate mapping results during the course of processing. Furthermore, CAB's γ parameter, which is related to working memory, allows for consideration of dual-task manipulations, individual differences data, and group differences data. CAB's ability to make sensible processing predictions, while remaining simple, may make it an ideal tool for advancing empirical research in comparison.

Notes

1. Although CAB considers all acyclic chains between node pairs, this bound can be guaranteed by considering only the k (some constant) shortest chains between node pairs.
2. Modeling inference is beyond the scope of this paper. CAB's design is appropriate for such an extension.
3. LISA can map Spellman and Holyoak's (1992) large analogy between the Persian Gulf crisis and World War II if provided the mapping between Saddam Hussein and Hitler as well as a specific order in which to map subsequent propositions one at a time (Holyoak & Hummel, 2001).

Acknowledgments

This work was supported by AFOSR grant F49620-01-1-0295 to B.C. Love. We would like to thank Art Markman for his helpful comments.

References

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure mapping engine: Algorithm and examples. *Artificial Intelligence*, *41*, 1–63.

Forbus, K. D., & Oblinger, D. (1990). Making SME greedy and pragmatic. In *Proceedings of the 12th Annual Conference of the Cognitive Science Society* (pp. 61–68). Mahwah, NJ: Lawrence Erlbaum Associates.

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

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical processing* (pp. 199–241). Cambridge: Cambridge University Press.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analog. *Cognitive Science*, *10*, 277–300.

Gentner, D., Ratterman, M., & Forbus, K. (1993). The role of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, *25*, 524–575.

Goldstone, R. L. (1994). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *20*, 3–28.

Goldstone, R. L., & Rogosky, B. J. (2002). Using relations within conceptual systems to translate across conceptual systems. *Cognition*, *84*, 295–320.

Holyoak, K. J., & Hummel, J. E. (2001). Toward an understanding of analogy within a biological symbol system. In D. Gentner, K. J. Holyoak, & B. N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 161–195). Cambridge, MA: MIT Press.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, *13*, 295–355.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*, 427–466.

James, W. (1985). *Psychology: The briefer course*. Notre Dame, IN: University of Notre Dame Press.

Keane, M. T. (1997). What makes an analogy difficult? The effects of order and causal structure on analogical mapping. *Journal of Experimental Psychology: Learning, Memory,* & *Cognition*, *23*, 946–967.

Keane, M. T., & Brayshaw, M. (1988). The incremental analogy machine: A computational model of analogy. In *Third European Working Session on Machine Learning* (pp. 53–62). San Mateo, CA: Morgan Kaufmann.

Love, B. C., Rouder, J., & Wisniewski, E. J. (1999). A structural account of global and local processing. *Cognitive Psychology*, *38*, 291–316.

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

Markman, A. B., & Gentner, D. (1993a). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language*, *32*, 517–535.

Markman, A. B., & Gentner, D. (1993b). Structural alignment during similarity comparisons. *Cognitive Psychology*, *23*, 431–467.

Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory and Cognition*, *24*, 235–249.

Markman, A. B., & Gentner, D. (1997). The effects of alignability on memory. *Psychological Science*, *8*, 363–367.

Markman, A. B., & Gentner, D. (2000). Structure mapping in the comparison process. *American Journal of Psychology*, *113*, 501–538.

Marks, L. E. (1989). On cross-modal similarity: The perceptual structure in pitch, loudness, and brightness. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 586–602.

Minsky, M. (1981). A framework for representing knowledge. In J. Haugland (Ed.), *Mind design* (pp. 95–128). Cambridge, MA: MIT Press.

Murphy, G. L., & Allopenna, P. D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning, Memory,* & *Cognition*, *20*, 904–919.

Nosofsky, R. M. (1986). Attention, similairty, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.

Rumelhart, D. E., & Abrahamson, A. A. (1973). A model for analogical reasoning. *Cognitive Psychology*, *5*, 1–28.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*, 533–536.

Shepard, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. Part 1. *Psychometrika*, *1*, 125–140.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317–1323.

Spellman, B. A., & Holyoak, K. J. (1992). If Saddam is Hitler then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology*, *62*, 913–933.

Stilwell, C. H., & Markman, A. B. (2001). The fate of irrelevant information in analogical mapping. In *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 988–993). Mahwah, NJ: Lawrence Erlbaum Associates.

Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327–352.