

# Analogical Similarity: Performing Structure Alignment in a Connectionist Network

**Jonathan B. Handler**

The Institute for the Learning Sciences  
1890 Maple Ave  
Evanston, IL 60201  
jhandler@ils.nwu.edu

**Paul R. Cooper**

The Institute for the Learning Sciences  
1890 Maple Ave  
Evanston, IL 60201  
cooper@ils.nwu.edu

## Abstract

We describe a connectionist network that performs a complex, cognitive task. In contrast, the majority of neural network research has been devoted to connectionist networks that perform low-level tasks, such as vision. Higher cognitive tasks, like categorization, analogy, and similarity may ultimately rest on alignment of the structured representations of two domains. We model human judgments of similarity, as predicted by Structure-Mapping Theory, in the one-shot mapping task. We use a localist connectionist representation in a Markov Random Field formalism to perform cross-product matching on graph representations of propositions. The network performs structured analogies in its domain flexibly and robustly, resolving local and non-local constraints at multiple levels of abstraction.

## Introduction

The process of structure-matching may underlie a broad range of cognitive phenomena, ranging from analogy and metaphor through visual recognition [Gentner, 1983; Falkenhainer, Forbus & Gentner, 1986; Medin, 1989; Cooper, 1990]. The development of biologically plausible models of structure matching may offer additional insights into how this important process can be computed.

In this paper, we develop a structured connectionist

model that performs structure matching on a well-studied task – one-shot mapping under conditions of cross-mapping (Figure 1)[Markman & Gentner, in press; Gentner & Toupin, 1986; Goldstone, Medin & Gentner, 1991]. The psychological evidence shows both that structure alignment must play a role in the human performance of this task and that there is a subtle interaction of local and non-local constraints at multiple levels of abstraction intrinsic to its solution. These are the kinds of computational characteristics that lend themselves naturally to modelling by a connectionist network. Even so, the inherently non-local and relational nature of structure matching offered an interesting challenge for connectionist modelling.

Unlike most neural network research, our model performs a real high-level cognitive task. It is a localist, structured, connectionist model [Feldman, Fenty & Goddard, 1988; Feldman & Ballard, 1988; Hinton, McClelland & Rumelhart, 1986], where network entities represent natural problem constituents. The key constraints in the task, exemplified by the rules of structure mapping theory [Gentner, 1983; Falkenhainer, Forbus & Gentner, 1986] such as the necessity to construct a 1-to-1 mapping, are instantiated directly by the system of connections. The architecture closely follows that of earlier connectionist models that addressed the problem of structure recognition in vision [Cooper, 1990]. In particular, the matching is computed by a cross-product network that considers part-part and relation-relation matches in parallel. We adopted Markov Random Fields as the formal basis of our network,

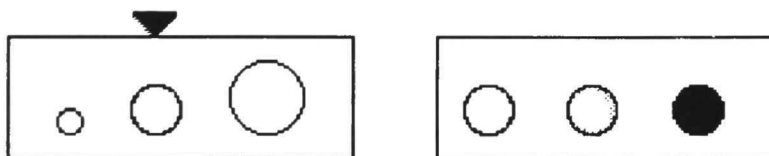


Figure 1: The one-shot mapping task. Which one from the set of the right best “goes with” the shape selected from the left set? Both literal and structural properties of the problem play a role in the answer.

which provides a robust and principled framework for resolving systems of simultaneous soft constraints [Cooper, 1992].

Our large network model successfully solves problems in the domain, and conforms to human performance on the task. Furthermore, it seems naturally suited to expressing the relevant constraints that determine task performance. As a result, we suspect that it could be easily and naturally modified to model related tasks and predict performance. In addition, the evidential nature of the underlying formal framework suggests that it could naturally incorporate real perceptual input, a possibly significant advantage when compared to other models [Falkenhainer, Forbus & Gentner, 1986; Holyoak & Thagard, 1989; Hofstadter & Mitchell, 1992].

## SMERF: Network Definition

### Overview

Structure-mapping theory predicts, and experiment has borne out, that people prefer mappings that preserve related sets of predicates over mappings that preserve only isolated predicates. This is the *systematicity principle*. There are two additional constraints on mapping. First, we want our mapping to be 1:1. Second, if we choose to map two predicates, then we must map their arguments. At a computational level, this implies that we must constrain the possible matches for a predicate such that it never matches two predicates from the other set (nor any from its own set). We must also find a way to propagate the information that a match has occurred at superordinate levels down to lower levels [Gentner, 1983; Falkenhainer, Forbus & Gentner, 1986].

We wanted a network formalism that simplified expressing these constraints. A distributed representation would have been inappropriate since it would not allow the direct representation of the interaction of differing values in a computation. Further, we wanted our network to be flexible in the type of mapping that it could make, which

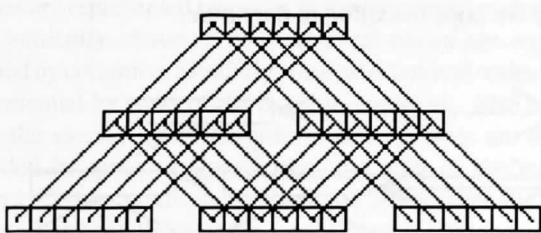


Figure 2: Relational Networks. Objects are represented as feature vectors at the bottom of the diagram. Relations are indicated by connections between levels.

required explicitly representing the potential matches and constraints between them. Thus, we chose a localist model – Markov Random Fields (MRF's) [Kindermann & Snell, 1980] – that can represent the structure-mapping constraints via their connection structure. MRF's also embody soft-constraint satisfaction, making them noise-resistant, and are stochastic, potentially allowing the integration of this model with perception.

In a Markov random field, each variable is assigned to one processor in the network, a *site*. The value for each site comes from a set of possible values, called *labels*. A site changes its label in response to pressure from its nearest neighbors, grouped into *cliques*. *N-cliques* are sets of *n* nodes that mutually constrain one-another. Constraints are implemented by assigning, prior to calculation, a numeric score to each constraint, called the *potential*. Input to the network is in the form of *evidence* that a set of label(s) should be applied to a set of site(s). Finally, each site is influenced by a *prior probability* – a measure for each label at each site that that site will take that label. SMERF is not currently implemented in parallel hardware (For a parallel implementation of Markov Random Fields, see [Swain, Wixson & Chou, 1990]). Instead, the network was simulated serially, using Chou's highest-confidence-first (HCF) algorithm [Chou & Brown, 1990]. HCF is guaranteed to converge to a local minimum; empirical observation has established its performance as excellent on a wide range of tasks [Cooper, 1990].

### Representing the Source and Target: Relational Networks

In the one-shot mapping task (Figure 1), subjects are presented with the two sets of figures and the experimenter points to the figure indicated by the arrow. Subjects are then asked which of the figures in set B "goes with" the one indicated. We used the structure implicit in the possible relations for the two domains to dictate the structure of the network representation. Thus, we had predicates on three different levels, or *orders*<sup>1</sup>.

Figure 2 shows the layout for nodes in the relational networks. They are organized by object, predicate order, and attribute dimension. In the diagram, each object is represented by a vector of six attributes, one for each of six attribute dimensions. Attribute dimensions were size, shape, color, orientation of short axis, orientation of long axis, and regularity of figure. In the diagram, each node represents

<sup>1</sup> Structure mapping theory defines *order* as follows. The order of objects is zero. The order of a predicate is one more than the highest order of any of its arguments. Attributes are defined as one-place, first-order predicates.

one attribute dimension. Objects are represented by the three full vectors at the bottom of the structure. The second level represents twelve possible second-order binary relations over the attributes. Relations are computed, within the same dimension, for left-middle and middle-right pairing of objects. Finally, the third level represents six binary relations over second-order predicates, for within-dimension pairings. The overall effect is of six unconnected tree-like structures, each one representing a possible relational hierarchy for one attribute dimension.

Label sets for the nodes depended on level, but, interestingly, not dimension. This is because the third-order predicates that we were modeling were ordinal predicates. We could thus allow attributes to range over a set of ordinal values that was the same for each attribute dimension. Second-order nodes calculated the ordinal difference between the attributes of the objects they compared. There were seven possible labels at this level, three-less-than-through three-greater-than (including equal). The third-order predicates were “Monotonically Increasing/Decreasing,” “Symmetric (two valences as well),” “Flat (no ordinal change),” and “No Relation.”

### Matching Networks

The matching networks in SMERF perform the work of aligning the relational structure between the two domains. They are based on the concept of the cross-product matching network [Cooper, 1990; Feldman, Fanty & Goddard, 1988]. There is one node for each potential match between predicates from the two domains. As with simple cross-product matching networks, we use inhibition between match nodes for objects that require a unique match.

Simple cross-product matching networks seek to place two sets of items in one-to-one correspondence. In contrast, SMERF seeks to match *relations on* three objects to relations on three objects. This results in a conceptual sub-division of each node of a simple cross-product matching network into a six-by-six matrix. All nodes have the label “Match/NoMatch.”

There are two types of connection in the matching networks – vertical and horizontal. Vertical constraints

express the structural-consistency constraint of structure-mapping theory and horizontal constraints express the 1:1 matching constraint.

There are two types of vertical connection – encouraging and discouraging. Encouraging connections act to strengthen relational matches that are consistent with decisions about higher-order relational matches. Discouraging vertical connections are connections between a single match node at one level and all nodes in each of the subgrids that correspond to matches between objects that are inconsistent with the matching node at the higher level. These connections act to discourage matches for these predicates.

Horizontal connections within the matching networks are also of two types. The first type are normal cross-product matching connections. They act within subgrids of the matching networks, insuring that within-object predicate matches are unique. The second type of horizontal connection acts between related predicates for different-object matches. Thus, if we match predicates for the size of an object and the color of an object, we do not want a match between the size of the same object and any predicate in any other object than the one to which it is already matched.

### Putting it All Together

We have described, so far, two networks that perform the tasks of structure instantiation and cross-product matching. Performing structure alignment with these tools is simple. We merely connect predicates from the relational networks to the appropriate match nodes in the cross-product matching networks, as shown in Figure 3. The relational networks are on the left and right of the figure, with the matching networks in the center. The arrows represent the most normal flow of information in the network. That is, information is received, as evidence, at the bottom of the relational networks. It flows both up – allowing the relational networks to instantiate the appropriate structure – and at the same time across to the matching networks. These arrows are not uni-directional – information is free to flow across *any* connection, a fact that we hope to exploit in the future.

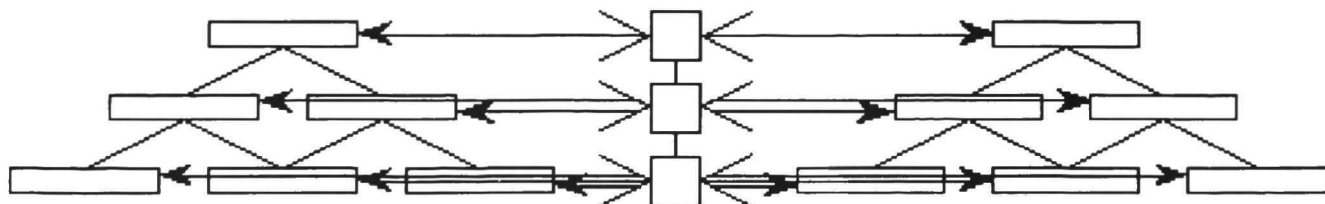


Figure 3: SMERF, an abstract view. Relational networks are to the left and right of the matching networks (center). Information flows according to the arrows.

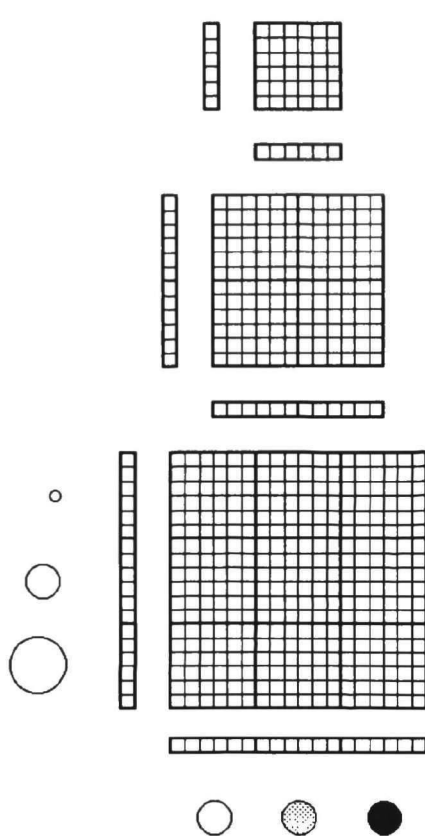


Figure 4a: SMERF, before running.

The empty network shows the layout of matching nodes into grids and sub-grids. Relational nodes are shown so that the matching nodes for those relations are in the same row/column.




Figure 4b: SMERF, final state.

Figure 4a (next page) shows a full view of the network before performing any calculation (see [Handler, 1992] for a fuller description). We can see that there are three levels of nodes, corresponding to the three predicate orders. These are arranged “top down” – highest order to lowest order. Arranged vertically on the left is the relational network for the top set of polygons in Figure 1. The relational network for the bottom set of polygons is represented horizontally, below each matching grid. Matching grids are represented in between the relational networks of the appropriate level. Note that matches for horizontal nodes are represented by nodes in the matching levels that are vertically above them (horizontal nodes are represented by columns). Similarly, vertical relational nodes are represented by rows of the matching network. Also note that the second- and first-order matching grids are subdivided by groups of predicates that refer to the same object. The six predicates are ordered down vertically (or across horizontally) at each level for each object as follows: Size; Shape; Color; Orientation of Short Axis; Orientation of Long Axis; and Regularity of Figure. Finally, we have placed the objects to the left of and below their

respective feature vectors.

Figure 4b shows the network in its final state for the problem given in Figure 1. For this experiment, the evidence for attribute labels amounted to certainty. Attributes are represented as squares proportional to their ordinal value. There are two different second-order relations in the example – all relations but size in the vertical set and color in the horizontal set are labeled “equal.” The size/color predicates are labeled “1-greater-than” for both sets. There are also two third-order relations – “flat” for all predicates but size and color; and “Monotonic Increase” for size/color. The matching network labels were “Match,” represented by white nodes and “No Match”, represented by black nodes. We can see that the network has chosen to match the small circle to the white circle, the medium circle to the grey circle, and the large circle to the black circle, corresponding to their places in the matching monotonic relations in each of the two domains.

There are two interesting facts to note. First, the match between “Monotonic Increase” predicates forces a match at lower levels for size and color predicates that would not match otherwise. Second, since the labels are the

same for predicates (within order) the network has no means of selecting which are the correct matches for predicates not constrained by higher-level matches. This was, in fact, a shortcoming in earlier versions of SMERF. However, discouraging vertical connections insure that matches for under-constrained predicates take place within object grids that are consistent with higher-level predicates. It is the unified functioning of these two abstract constraints that allows the network to come up with a good solution for the two sets, rather than a disconnected set of matches for underconstrained predicates along with correctly constrained matches for one predicate.

The network converges in time that is on the order of the total number of nodes. The simulation runs in approximately 40 seconds on a Sparcstation 1, about half of which is initialization of data structures. More importantly, the network grows only as  $O(n^2)$ , where  $n$  is the number of attributes.

## Related Work

The current model is a very close theoretical match to SME [Falkenhainer, Forbus & Gentner, 1986]. Like SME, the network attempts every match between predicates. However, as with all connectionist models, there is a trade-off made between time and space. Every match is "tried" only in the sense that there is a matching site for every possible combination of predicates. Space is used to represent matches, rather than time to compute them. Since the network converges without backtracking, we expect that the algorithm would complete in constant time in a fully parallel implementation. SME is input a set of propositional statements about the base and the target. Knowledge for the structure of the relations in both the base and target is provided explicitly. In contrast, the input to SMERF is in the form of feature vectors. The decision about labels for higher-order predicates is made during the calculation. This should allow for greater flexibility in modeling the analogy process in domains where a propositional representation is not available beforehand.

ACME [Holyoak & Thagard, 1989] is also an extension of structure-mapping theory. It too attempts to preserve relational structure and disregard attributional information. As in SMERF, all matches are explicitly represented. However, there is a deep difference in the method of specifying the constraints between the two systems. ACME takes its inputs in propositional form. Because of this, it constructs a different network for every problem. This seems to us to be a somewhat restrictive method for modeling analogy in humans. In a sense, ACME must solve the analogy problem – construct a network – in order to solve the analogy problem – draw an analogy. In contrast, SMERF attempts to express the constraints in the

*matching process* itself. We attempt to leave open the question of the type of input. In general, since SMERF is probabilistic in nature, it extends naturally to domains that behave well for evidential treatment.

## Discussion

The chief challenge addressed in this work was the development of a connectionist network that expressed the constraints of structure-mapping theory for a real task. There were three relevant constraints that we drew from structure-mapping theory. The cross-product matching networks enforce 1:1 mapping. SMERF's propagation of information on matches vertically in its matching networks works to encourage/discourage matches that involve the same/different predicates enforce systematicity and structural connectedness.

The second question was really part of the first: how can structure-mapping theory and parallel architecture be combined? Specifically, it is an assumption of structure-mapping theory that non-local, relational information is the basis of analogy. How could local information be unified with non-local information in a parallel framework so that it could override attribute similarity? The network provided a mechanism for the propagation of high-order information downward, so that high-order matches override attribute matches.

There are several directions for future work. It might be interesting to allow matching information to move upward through the matching levels and then out to the relational networks. We imagine that this might help solve traditional analogy problems of the form "a : b :: c : ?" It would also be interesting to provide conflicting evidence at the attribute level to see how the network would deal with it.

Prior probabilities for sites, especially the matching sites, provide a further arena for exploration. Consider the *relational shift*. Markman and Gentner found that the number of relational similarity responses rose as their experiment progressed (see also [Goldstone, Medin & Gentner, 1991]). This fact can be explained if we assume that the prior probability distribution is modified after each judgment, so that the network "expects" to make a relational response.

Processing symbolic information in parallel systems provides a powerful method for modeling high-level cognitive functions. This paper has described a network capable of modeling human performance on the one-shot mapping task. We hope that in the future the methods described here will provide leverage for giving computers the ability to solve these larger problems.

## References

- P. Chou and C. Brown. 1990. The Theory and Practice of Bayesian Image Labeling. *International Journal of Computer Vision*, 4:185-210.
- P. Cooper. 1990. Parallel Structure Recognition with Uncertainty: Coupled Segmentation and Matching. In *Proceedings of the Third International Conference on Computer Vision ICCV '90*. For a longer version, see Technical Report #5, Institute for the Learning Sciences, Evanston, IL 60201.
- P. Cooper. 1992. Structure Recognition by Connectionist Relaxation: Formal Analysis. *Computational Intelligence*, 8(1).
- B. Falkenhainer, K. Forbus, and D. Gentner. 1986. The Structure-Mapping Engine, Algorithm and Examples. *Artificial Intelligence*, 41:1-6.
- J. Feldman and D. Ballard. 1988. Connectionist Models and Their Properties. In D. Waltz and J. Feldman, editors, *Connectionist Models and Their Implications: Readings from Cognitive Science*, Chapter 2. Ablex Publishing Corporation.
- J. Feldman, M. Fianty and, N. Goddard. March 1988. Computing with Structured Neural Networks. *Computer*.
- D. Gentner. Structure Mapping: A Theoretical Framework for Analogy. 1983. *Cognitive Science*, 7:155-170.
- D. Gentner, and C. Toupin. 1986. Systematicity and Surface Similarity in the development of Analogy. *Cognitive Science*, 10:277-300.
- R. Goldstone, D. Medin, and D. Gentner. 1991. Relational Similarity and the Non-Independence of Features in Similarity Judgments. *Cognitive Psychology*, pages 222-264.
- J. Handler. 1992. Analogical Similarity: Structure Alignment in a Connectionist Network. Master's Thesis, The Institute for the Learning Sciences, Northwestern University.
- G. Hinton, J. McClelland, and D. Rumelhart. 1986. Distributed Representations. In D. Rumelhart and J. McClelland, editors, *Parallel Distributed Processing*, Volume 1, Chapter 3. The MIT Press.
- D. Hofstadter and M. Mitchell. 1992. An Overview of the Copycat Project. In K. Holyoak and J. Barnden, editors, *Connectionist Approaches to Analogy, Metaphor, and Case-Based Reasoning*. Ablex.
- K. Holyoak and P. Thagard. 1989. Analogical Mapping by Constraint Satisfaction. *Cognitive Science*, 13:295-355.
- R. Kindermann and J. Snell. 1980. *Markov Random Fields and Their Applications*. American Mathematical Society.
- A. Markman and D. Gentner. In Press. Structural Alignment During Similarity Comparisons. *Cognitive Psychology*.
- D. Medin. 1989. Concepts and Conceptual Structure. *American Psychologist*, 44(12):1469-1481.
- M. Swain, L. Wixson, and P. Chou. 1990. Efficient Parallel Estimation for Markov Random Fields. In M. Henrion, R. Shachter, L. Kanal, and J. Lemmer, editors, *Uncertainty in Artificial Intelligence 5*, pages 407-419. Elsevier Science Publishers.