

A Connectionist Model of Metaphor by Pattern Completion.

Michael S.C. Thomas

Psychology Group, King Alfred's College
Sparkford Road, Winchester SO22 4NR, UK.
& Department of Experimental Psychology
University of Oxford, UK.
michael.thomas@psy.ox.ac.uk

Denis Mareschal

Department of Psychology
University of Exeter
Perry Road, Exeter EX4 4QG
UK.
dmaresch@singer.exeter.ac.uk

Abstract

In this paper we present a simple process model (based on connectionist pattern completion) of *A is B* metaphor comprehension. The Metaphor by Pattern Completion (MPC) model capitalizes on an existing semantic memory mechanism. Metaphorical enhancement is produced by presenting a semantic vector representation of the target word (A) to a connectionist network storing the knowledge base (B). Effects found in human data such as meaning enhancement, asymmetric processing, context sensitivity and compound indexing all fall naturally out of the pattern completion mechanism. The MPC model suggests a simple way of separating literal from metaphorical statements. It provides a means of predicting when a metaphor will appear to fail. Moreover, we suggest that the mechanism can form the basis of a comparison procedure that supports analogy. The MPC mechanism avoids the problem of identifying which features of a concept are relevant for similarity matching in analogies, because the prior metaphor stage naturally enhances relevant features and suppresses the irrelevant features. The MPC model is both domain general (in that it does not depend on the structure of the metaphor domain) and parsimonious (in that it does not posit metaphor-specific mechanisms).

1. Introduction

In this paper we describe a simple computational model of the processes involved in comprehending metaphors of the form *A is B* (e.g., "The Apple is a Ball"). These have been referred to as Image Metaphors (Lakoff, 1994) or simply Attribute Mapping (Holyoak and Thagard, 1995). There have been few attempts to build computational models of metaphor because it is assumed that a metaphor is equivalent to an analogy with the comparison made implicit (that is, *A is B* is just *A is like B* with the "like" removed). Although on grounds of parsimony it might be surprising if the processes underlying metaphor and analogy were radically different (Rumelhart, 1979), metaphors are often seen as producing a stronger and subtler effect than similes (Glucksberg and Keysar, 1993). This suggests that important differences may underlie the two processes.

Metaphors can imply a comparison but they are not reducible to comparisons (Black, 1979). It seems reasonable to say "I don't mean Richard is *like* a lion, I mean that Richard *is* a

lion." There is a sense in which the metaphorical comparison seems stronger than the analogical comparison. Metaphors can be viewed as an intermediate between literal attribution statements (e.g. Richard is brave) and similes (e.g. Richard is like a lion). By saying that "Richard is a lion", certain properties of a lion (such as bravery) are attributed to Richard. The process by which the features of Richard are modified through the use of a metaphor is still very much an open question.

Because of the conflation between metaphors and similes, existing computational models have primarily examined the processes involved in the formation of analogies and in similarity-based retrieval (e.g. ACME: Holyoak and Thagard, 1989; MAC/FAC: Forbus, Gentner, and Law, 1995). These computational models have proposed that analogical comparisons involve either: (a) forming mappings or links between static representations (e.g. ACME), or (b) a kind of "high level perception" in which representations are dynamically configured according to domain specific heuristics (e.g., Copycat: Hofstadter, 1984; Mitchell, 1993; Tabletop: Hofstadter and French, 1994). While most models seem to fall into either the mapping or high level perception camp, Burns and Holyoak (1994) have shown that if enough domain specific information is pre-wired into the systems, then both types of model can behave in a similar fashion.

The comprehension of complex metaphors, requiring the formation of mappings between the elements in two structured representations, may involve task-specific cognitive mechanisms. However, it is questionable how frequently such complex metaphors are understood 'on-line' by the operation of a single cognitive mechanism. Lakoff (e.g., Lakoff, 1994; Lakoff and Johnson, 1980) has proposed that the mappings involved in many "conventional" complex metaphors are derived and agreed in advance of usage, by members of a given linguistic community. Moreover, complex metaphor comprehension is likely to incorporate a range of processes and strategies. The model we present in this paper is intended to capture the 'on-line' comprehension of metaphors occurring at very short time scales (e.g., seconds). We believe that these simple on-line mechanisms may form the basis of (or contribute to) the

more complex comprehension strategies occurring at longer time scales (e.g., over minutes).

Simple *A is B* metaphors provide a way of exploring the basic mechanisms which underlie meaning enhancement independently of the need for any prior complex mappings. In particular, we suggest that meaning enhancement can be modeled by a simple domain general processing mechanism. For simple *A is B* metaphors, the Metaphor by Pattern Completion (MPC) model accounts for how the semantic features of a target word (A) are transformed by the semantic properties of a knowledge base (B). The model capitalizes on the properties of existing semantic memory mechanisms and does not posit "metaphor-specific" processes. In this sense it is both domain general (in that it does not depend on the structure of the specific metaphor domain) and parsimonious (in that it does not need to posit new mechanisms).

The heart of the model is based on Black's (1979) interaction theory of *A is B* metaphor genesis. Black's theory contains a number of abstract concepts which attempt to capture the complexity of the process of metaphor. According to Black, *A is B* metaphors gain their effect through an interaction between the target and the source concepts, whereby "associated implications" from the source concept are "projected upon" the target concept. The "associated implications" are derived from the source concept's "implicative complex" which is determined by the "current opinions shared by members of a certain speech-community" (Black, 1979, p. 28). The metaphor involves "a shift in the speaker's meaning - and the corresponding hearer's meaning - what both of them understand by the words as used on the particular occasion." (*ibid.*). The MPC model provides a more tangible expression of these abstract ideas.

The model captures four key phenomena of metaphorical comparisons: (1) the semantic effect of the juxtaposition of two concepts; (2) the directional asymmetry of such comparisons (that is, the metaphor *A is B* frequently has a quite different effect from the metaphor *B is A*, even though the similarities between the two concepts are the same in each case); (3) context sensitivity; and (4) compound indexing effects in metaphorical comparisons. Further, the model may provide some clues as to how metaphorical statements can be distinguished from literal statements. Finally it generates a hypothesis of the relation between metaphorical and analogical comparisons and provides an account of when metaphors will fail.

2. The Model.

The MPC model is based on the processing power of three-layered feed-forward connectionist networks. Figure 1 shows the complete network architecture.

Both the inputs and outputs to the network encode semantic features. There are 13 semantic features (though note Balls have 2 extra features, as described in Section 3.4 below).

Input information entering the network is vetted towards separate knowledge bases via a categorization mechanism. This mechanism must be able to separate inputs into appropriate categories as dictated by perceptual or linguistic contextual cues. Hence, the prior ability to categorize inputs is a necessary assumption of this model. The selector mechanism was not actually implemented in a connectionist form since it has no direct impact on the process through which metaphors emerge. However, it could be implemented as a feed-forward network with the same semantic inputs as the knowledge-base network but with a single category output acting as a shunting mechanism for redirecting information flow.

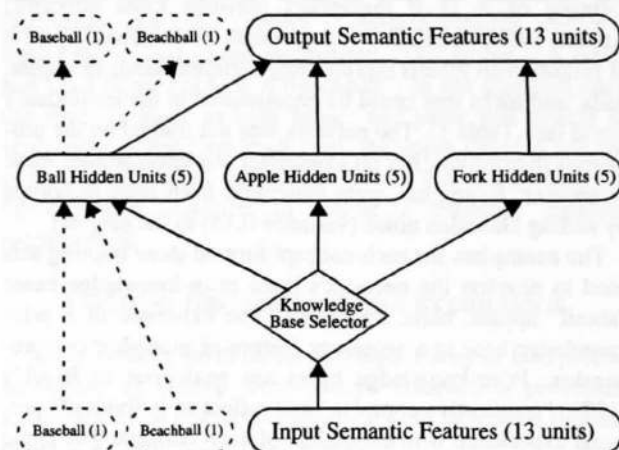


Figure 1. Architecture of the MPC Model.

Table 1: Semantic Features for Concept Prototypes.

	Index		Size		Texture		Shape		Actions		Color					
	Baseball	Beachball	Hand Sized	Lap Sized	Hard	Soft	Pointed	Irregular	Round	Thrown	Edible	White	Brown	Green	Red	
Apple			1.0	0.0	0.0	0.0	0.3	0.7	0.0	0.3	0.7	1.0	0.2	0.8	0.3	0.0
			0.0	1.0	0.0	0.0	1.0	0.0	0.3	0.7	1.0	0.0	0.2	0.8	0.3	0.0
	(Rotten)		0.0	0.0	1.0	0.0	0.0	0.2	0.8	0.3	0.0	0.0	0.2	0.8	0.3	0.0
Ball	Baseball		0.0	0.0	0.0	1.0	0.0	1.0	0.0	0.0	0.0	1.0	0.9	0.1	1.0	0.0
	Beachball		1.0	0.0	0.0	0.0	0.0	1.0	0.0	0.0	1.0	0.0	0.1	0.9	0.0	1.0
	Beachball		0.0	1.0	0.0	0.0	0.0	1.0	0.0	0.0	1.0	0.0	0.1	0.9	0.0	1.0
Fork			0.0	0.0	0.0	1.0	0.0	0.1	0.0	0.0	1.0	0.0	0.7	0.3		

In this limited example, we have chosen a small set of features with which to describe the objects. The intention here is to model an 'on-line' mechanism. If this mechanism fails, or the domain is too complex, other strategies can come into play. Because this is a default comprehension mechanism, we suggest that in a scaled up model, all concepts would be represented along a common set of default features (e.g., high frequency or high salience features). These may not represent all concepts sufficiently for all purposes but they should suffice as a default approximation. A concept that is not adequately represented will fail to be understood, both literally and metaphorically (see later), and will trigger a more complex strategy.

In order to generate knowledge bases for separate concepts, the network was trained to autoassociate exemplars of each concept. For simplicity, we restricted the model to the forming of *A is B* metaphors between three concepts: Apples, Balls, and Forks. The concepts were defined by a set of prototypical tokens representing different kinds of apples, balls, and forks that could be encountered in the individual's world (see Table 1). The network was not trained on the prototypes themselves, but on exemplars clustered around these prototypes. Exemplars were generated from each prototype by adding Gaussian noise (variance 0.15) to the original.

The exemplars for each concept formed three training sets used to develop the network's three prior-knowledge bases "about" apples, balls, and forks. The existence of a prior knowledge base is a necessary feature of metaphor comprehension. Prior-knowledge bases are analogous to Black's (1979) "implicative complex" and reflect an individual's personal experience with exemplars of each concept. The apple sub-network was trained to autoassociate patterns from 10 exemplars of each of three apple kinds (e.g., red, green, and rotten) for a total of 30 patterns. Similarly, the ball sub-network was trained to autoassociate 10 exemplars from three different kinds (for a total of 30 patterns). Finally, the fork sub-network was trained to autoassociate 10 exemplars from 1 kind (for a total of 10 patterns). Because there was only 1 kind of fork (as opposed to 3 kinds of both apples and balls), a single blank training pattern (zero input and output) was added to the fork training set to inhibit overlearning of the fork exemplars. All networks were trained with Backpropagation using the following parameter values: learning-rate: 0.1, momentum: 0.0, initial weight range: ± 0.5 . Each sub-network was trained for 1000 epochs. All reported results are averaged over $n=10$ replications.

In the MPC model, metaphorical interpretations arise naturally from the pattern completion properties of non-linear connectionist information processing. Pattern completion is often used to "clean up" noisy input patterns. That is, the network transforms the input to make it more consistent with the knowledge stored in the network. In the case of metaphors, however, the input is not a noisy version of a pattern on which the network has previously been trained, but an exemplar of another concept. For example, the sub-network

trained on Balls might be presented with an Apple pattern. The resulting output would be an Apple pattern transformed in such a way as to make it more consistent with the prototypical Ball representation stored in the network. The nature of the transformation will depend on the relationship between the Apple and Ball concepts. To cast this in Black's terms, a metaphor is produced when the source concept Ball "projects" its knowledge representation or "associated implications" onto the target concept Apple. The result is a new understanding of Apple in which features are selectively modified according to the Ball prototype.

Metaphors are achieved by the redirection of information flow into one knowledge base or another. The role of the "is" in the *A is B* metaphor is to trigger that redirection. If the information flow were not redirected, the result would be processing of the input by the knowledge base of which it is an exemplar (e.g., the Apple input would flow through the Apple sub-network). This would "clean-up" the Apple input to make it more apple-like and be akin to "recognizing the input as an apple". This type of semantic representation has been postulated before to account for prototype effects in semantic memory (McClelland and Rumelhart, 1986). Hence the MPC model can be seen as merely capitalizing on existing semantic memory mechanisms.

3. Model Performance.

Figures 2 to 5 show the results of various simulated metaphors. For each *A is B* metaphor, the top bar chart of each pair shows the activation of the semantic feature input representation of the target concept (A). The lower graph shows the semantic feature output representation of the same target (A) once it has been transformed by the source knowledge base (B). The metaphor represented is labeled at the top of the figure. This section reports on four metaphor effects found in humans that fall naturally out of the MPC mechanism.

3.1 Meaning Enhancement

Figure 2 shows the enhancement of the semantic features of an apple concept for the metaphor: "The Apple is a Ball". The input is an exemplar close to its prototype kind. The effect of this metaphor is to reduce the edible label, to suggest that this apple might be suitable for throwing, to increase the hardness and roundness labels, while introducing some ambiguity into the color features. This is the type of enhancement effect outlined by Black (1962, 1979) (see section 1). Note that despite the fact that 20/30 of the Ball exemplars are soft beachballs, the Apple is still made to look harder rather than softer by this metaphor. This is because the apple is closer in size to a hard baseball than it is to a soft beachball. Semantic enhancement is thus not a default imposition of ball features onto those of an apple, but an interaction between stored ball knowledge and the nature of the apple exemplar being presented to the ball sub-network.

3.2 Asymmetric Comparisons.

Real metaphors are rarely symmetrical (i.e., $A \text{ is } B$ is not equivalent to $B \text{ is } A$), though the similarities between A and B are the same in either case. For example, the metaphors "sermons are slipping pills" and "sleeping pills are sermons" have quite different implications (Glucksberg and Keysar, 1990). Figure 3 shows the result of the metaphor "The Ball is an Apple", the reverse of that shown in Fig.2. The effect of this metaphor is to: reduce the likelihood of being thrown, the size of the ball, and the roundness label, and increase the irregularity label, the softness label and the edibility label. The semantic effect of this metaphor is different from that in the previous case despite the fact that the distance between the sets of prototypes for apples and balls in n -dimensional similarity space remains the same. In fact, the asymmetry only seems problematic if one views the comparison of concepts as involving the measurement of static distances (e.g. Rumelhart and Abrahamson, 1973). In the current model, metaphors are based on dynamic transformations. There is no reason for these transformations to be symmetrical.

3.3 Context sensitivity.

Figure 4 shows the effect of changing the target context on the metaphor process. The effect of the metaphor: "The Ball is an Apple" is similar for both Red and Green balls: roundness and size are reduced, whereas irregularity, softness, and edibility labels are increased. However, a noticeably different enhancement is produced on Brown ball. Here, the softness label is very much increased, and the ball retains a low edibility label. This differential effect




Figure 2. Meaning Enhancement.




Figure 4. Context Sensitivity.

occurs because the Apple knowledge base is constructed around three prototypes (Table 1). Brown/Rotten apples have different properties from other apples. The target and source interact in producing the effect of the metaphor. The context of the target (e.g., brown vs. red) impacts on which aspects of the knowledge base (Apples) are relevant to the metaphor. This is also found in human metaphor interpretations (Black, 1979).

3.4 Compound indexing.

In the apple knowledge base, different kinds are implicitly represented by a distribution of exemplars around three prototypes. It is also possible to index separate kinds explicitly in metaphors (Malgady and Johnson, 1976). In the Ball sub-network, two semantic features were added to the input and output representations in order to code for Baseballs and Beachballs explicitly (Figure 1, dashed outline). Figure 5 shows the effects of the metaphors: "The Apple is a Baseball" and "The Apple is a Beachball". In the Baseball case, the apple is made to look harder and paler and retains its size, while in the Beachball case, the apple remains red and is seen as both softer and larger. Thus a knowledge base can be explicitly distorted to enable a different meaning enhancement.

4. Model Implications and Predictions.

In this section, we describe the implications of interpreting metaphors as arising from pattern completion processes. Each subsection discusses modeling work described in more detail in a forthcoming longer report.




Figure 3. Asymmetric Comparisons.




Figure 5. Compound Indexing.

4.1 Distinguishing metaphors & literal statements.

A key issue in metaphor research is how to detect that a statement is metaphorical and not literal (Ortony, 1993). In the MPC model, both types of statements are processed by the same mechanism. A statement is recognized as literal if only small semantic changes occur in the output. For example, in saying "this [apple-like] item is an apple", very little change is produced in the item's output features. However, saying "this [apple-like] item is a ball" results in a large amount of output feature enhancement.

In a given discourse, the listener may be expecting a certain amount of meaning change from a given communication. If the meaning change is greater than that which the listener expects given the context, then he or she may conclude that the communication is metaphorical. Thus the difference between a literal and metaphorical statement of the form *A is B* may be detected by matching the network error score against a criterion suitable for the current discourse. In our simulations, when testing the metaphors "An Apple is an Apple" and "An Apple is a Ball" using a novel Apple exemplar, the mean error scores were 0.07 and 2.54 respectively (sig. dif., related samples t-test, $t=8.04$, 9df, $p<0.0001$). Note that a low error score represents an accurate autoassociation, but also a metaphorically uninteresting juxtaposition: little meaning enhancement has taken place. This corresponds to the inverse relationship between the accuracy of a comparison and its aesthetic impact found in human subjects (Sternberg, Tourangeau, and Nigro; 1979).

4.2 Metaphors and Analogies.

Earlier, we suggested that a metaphor such as "Richard is a lion" can be viewed as intermediate between a literal attribution statement "Richard is brave" and a simile "Richard is like a lion". In the metaphorical case, we know Richard is *not* a lion, and thus must transfer some (relevant) features of the lion to Richard. We have seen how transforming the features for Richard using the knowledge we have about lions might achieve this enhancement. However, the problem of extracting the relevant features also exists for similes. When we are told that Richard is like a lion, how do we know the ways in which he resembles a lion?

One way to address this problem would be to perform the metaphor "Richard is a lion" using the MPC mechanism, as a first step to performing the simile "Richard is like a lion". This would produce an enhanced understanding of Richard, in which his bravery was enhanced, but not his possession of a mane (since Richard has no mane features to be enhanced). If we noted the ways in which Richard's features had been enhanced, we could then take these features to be the relevant similarities between Richard and the lion. These could then be output as the result of the comparison process implied by the simile.

The process of deriving the relevant similarities

would proceed as follows. For an *A is B* metaphor, the similarities between A and B would be computed based on those non-zero features in A that had been enhanced (or at least not suppressed) during the pattern completion process. Conversely, the differences would consist of the semantic features suppressed by the mapping, or previously zero features that had been enhanced. So, in the metaphor "The Apple is a Ball" (Figure 2), the similarities between apples and balls are: Round, Hard, and Red and Soft (to some extent), whereas the differences are Edibility, Thrown, Irregular, White, and Red and Soft (to some extent).


Analogy is then seen as a two stage process. The first stage comprises "seeing A as if it were B", which could be described as a form of high level perception. The second stage involves comparison between static representations to derive similarities. In short, this account of analogy incorporates elements of both the opposing views of the nature of analogy formation outlined in section 1, high level perception and representational mapping.

The MPC model suggests how metaphor and analogy may be related. Instead of being an implicit form of analogy, metaphor is seen as a communicative expression intended to enhance the meaning of a concept by borrowing information from another concept. Analogy on the other hand is a communicative expression intended to focus attention on the similarities between two concepts. Metaphor is the more primary process, highlighting any salient features 'on-line'. Analogy notes the resulting similarities between concepts.

4.3 Predicting when metaphors should fail.

The model suggests two situations under which metaphors will fail. First (and somewhat trivially), a metaphor *A is B* will fail if a subject knows nothing about B (i.e., if the knowledge base is not formed). One must have a representation of apples before apples can be used metaphorically. In other words, one must be able to recognize apples before one can see something else *as if it were an apple*. More interestingly, the process of metaphor will also be fruitless if the target and source have no positive features in common (i.e., if the semantic representation vectors are orthogonal). Pattern completion of target by source would produce no output, and thus no meaning. Figure 6 illustrates this in the case of the metaphor: "The Apple is a Fork". Here, few features show strongly on the output, since apples and forks have very little in common on this default feature set.

Figure 6. When Metaphors Fail.



5. Conclusion.

The model described in this paper provides a process account of key aspects of metaphor comprehension. It shows how semantic distortion, context sensitivity, and compound indexing fall naturally out of connectionist pattern completion mechanisms. Moreover, it provides a simple explanation of the directional asymmetry effects which have plagued models based on static similarity measures (e.g., Rumelhart and Abrahamson, 1973). The asymmetry effect arises from the fact that network transformations are not normally bi-directional.

All of this is accomplished through a simple and parsimonious mechanism. The MPC model capitalizes on an existing semantic memory mechanism and does not posit domain specific mechanisms. Existing models of analogy posit numerous complex mechanisms. It is worth noting how much can be achieved with so simple a mechanism. However, although compound indexing (Section 3.4) showed this limited case could be extended somewhat, it is not automatically clear how the current model could be extended beyond featural representations to incorporate the structure necessary to account for complex metaphorical comparisons.¹

We suggested that the pattern completion process could be used as the basis for an 'on-line' analogical comparison procedure which may begin to tackle the problem of determining which features of an object are relevant to a particular comparison. A prior metaphorical stage automatically identifies the relevant features as those that are modified by the metaphor. This account of analogy would incorporate elements of the two opposing views from current research into computational processes underlying analogical processing.

Lastly, we would not want to pretend that this model on its own can explain the full richness (or even mystery!) of metaphor. However, we do believe that it gives us some purchase on a complex and multi-faceted problem, a purchase which is gained at the expense of very few assumptions.

References.

- Black, M. (1962). *Models and Metaphors*. Ithaca, NY: Cornell University Press.
- Black, M. (1979). More about metaphor. In A. Ortony (Ed.), *Metaphor and Thought* (pp.19-43). Cambridge, England: Cambridge University Press.
- Burns, B. D. and Holyoak, K. J. (1994). Competing Models of Analogy: ACME Versus Copycat. *Proceedings of the 16th Annual Meeting of the Cognitive Science Society*. Erlbaum. 100-105.

1. However, as outlined in the Introduction, one might question whether the structural mappings necessary in complex metaphors are best described using a single, on-line process model of the type outlined here. For instance a multi-mechanism, iterative process may be more appropriate for stepping through the mappings in complex metaphors.

- Forbus, K. D., Gentner, D., and Law, K. (1995). MAC/FAC: A Model of Similarity-based Retrieval. *Cognitive Science*, 19, No. 2, 144-206.
- Glucksberg, S. and Keysar, B. (1990) Understanding Metaphorical Comparisons: Beyond Similarity. *Psychological Review*, 97, No. 1, 3-18.
- Glucksberg, S. and Keysar, B. (1993). How metaphors work. In A. Ortony (Ed.), *Metaphor and Thought 2nd Ed.* (pp. 401-424). Cambridge, England: Cambridge University Press.
- Hofstadter, D. R. (1984). The Copycat project: An experiment in non-deterministic and creative analogies. Cambridge, MA: MIT A.I. Laboratory Memo 755.
- Hofstadter, D. R. and French, R. M. (1994). Probing the Emergent Behavior of Tabletop: an Architecture Uniting High-level Perception with Analogy-making. *Proceedings of the 16th Annual Meeting of the Cognitive Science Society*. Erlbaum. 528-533.
- Holyoak, K. J. and Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.
- Holyoak, K. J. and Thagard, P. (1995). *Mental Leaps: Analogy in Creative Thought*. Cambridge, MA: MIT Press.
- Lakoff, G. (1994) What is Metaphor? In K. J. Holyoak and J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory, Vol.3* (pp. 203-258). Norwood NJ: Ablex.
- Lakoff, G. and Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- Malgady, R. G. and Johnson, M. G. (1976). Modifiers in metaphors: Effects of constituent phrase similarity on the interpretation of figurative sentences. *Journal of Psycholinguistic Research*, 5, 43-52.
- McClelland, J. L. and Rumelhart, D. E. (1986). A Distributed Model of Human Learning and Memory. In McClelland, Rumelhart, and The PDP Research Group (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol. 2: Psychological and Biological Models*. Cambridge, MA: MIT Press. pp. 170-215.
- Mitchell, M. (1993). *Analogy-making as perception*. Cambridge, MA: MIT Press.
- Ortony, A. (1993). The role of similarity in similes and metaphors. In A. Ortony (Ed.), *Metaphor and Thought 2nd Ed.* (pp. 342-356). Cambridge, England: Cambridge University Press.
- Rumelhart, D. E. (1979). Some problems with the notion of literal meanings. In A. Ortony (Ed.), *Metaphor and Thought* (pp. 71-82). Cambridge, England: Cambridge University Press.
- Rumelhart, D. E. and Abrahamson, A. A. (1973). A model for analogical reasoning. *Cognitive Psychology*, 5, 1-28.
- Sternberg, R. J., Tourangeau, R., and Nigro, G. (1979). Metaphor, induction, and social policy: The convergence of macroscopic and microscopic views. In A. Ortony (Ed.), *Metaphor and Thought* (pp. 277-306). Cambridge, England: Cambridge University Press.