A connectionist account of the emergence of the literal-metaphorical-anomalous distinction in young children

Michael S. C. Thomas (m.thomas@ich.ucl.ac.uk)
Neurocognitive Development Unit, Institute of Child Health, 30, Guilford Street
London, WC1N 1EH, UK

Denis Mareschal (d.mareschal@bbk.ac.uk)
Center for Brain and Cognitive Development, Birkbeck College, Malet Street
London, WC1E 7HX, UK

Andrew C. Hinds
Department of Psychology, King Alfred's College, Sparkford Road
Winchester, SO22 4NR, UK

Abstract

We present the first developmental computational model of metaphor comprehension, which seeks to relate the emergence of a distinction between literal and non-literal similarity in young children to the development of semantic representations. The model gradually learns to distinguish literal from metaphorical semantic juxtapositions as it acquires more knowledge about the vehicle domain. In accordance with Keil (1986), the separation of literal from metaphorical comparisons is found to depend on the maturity of the vehicle concept stored within the network. The model generates a number of explicit novel predictions.

Introduction

Despite the highly imaginative and figurative way in which children often describe the world, somewhat surprisingly it has been claimed that children are unable to understand figurative or metaphorical speech until they are quite old (Piaget, 1962; see Gibbs, 1994, for a complete review of this position). A likely explanation of this disparity is that adult usage of figurative devices such as metaphor involves several skills. For metaphor, these may include the perception of similarity and of anomaly in comprehending metaphors, the invention of similarities in generating metaphors, an understanding of the role of context in constraining possible meanings, an understanding of speaker intentions, and a metalinguistic ability to justify metaphor use based on specific cross-domain similarities (see e.g., Dowker, 1999). Moreover, it is possible that these skills have different developmental trajectories. Thus Dowker (1999) argues that age variations in similarity recognition and invention may be due to limited domain knowledge which serve to restrict the types of similarity employed by young children to mainly perceptual information. On the other hand, the temporary reduction found in the prevalence of metaphor in the language of children around the age of 6 to 7 (Gardner, Winner, Becchoffer, & Wolf, 1978; Winner, 1988) may be due to age variations in recognition of anomaly.

The idea that conceptual knowledge constrains the ability to use language figuratively is supported by evidence that metaphor usage in children is more prevalent in domains with which they are more familiar (Gottfried, 1997). Indeed Keil (1986) argued that metaphor usage closely shadows the development of conceptual categories. Similar arguments have been made in the related field of analogical reasoning, where it was also initially maintained that the relevant skills appear late in childhood (Piaget, 1962). However, when analogical reasoning was tested in more familiar domains, skills were found at a much earlier age. This implies that limitations in analogical reasoning arise from differences in the knowledge available to children as a basis for exercising this skill (Goswami, in press).

How, then, are we to interpret the apparent presence of metaphor in young children, for example, when a child aged 3 years and 5 months refers to a green carpet as 'grass' (Billow, 1981)? Putting aside the possibility of renaming in symbolic play (which need not involve any similarity between label and assigned referent), and the possibility that this is a case of over-extension (which can be ruled out by checking that the child knows the actual name for a carpet; see Gardner et al., 1978), the juxtaposition would qualify as metaphoric only under the following conditions: the child had not only spotted the similarity between the carpet and grass, but was also aware that carpet and grass fall into separate categories, so that the similarity between them was understood to be *non-literal*. Several authors have suggested that fuzziness in categorization could explain children's early use of apparently figurative language (Hakes, 1982; Marschark & Nall, 1985). If a child's conceptual knowledge has not formed into neat clusters, then there will be some overlap between categories. A sentence that appears figurative to adults may be interpreted as literal by the child.

Evidence to support this position can be found in a study by Vosniadou and Ortony (1983). In their investigations of

the emergence of the distinction between literal, metaphorical, and anomalous comparisons, these authors found evidence that, although 3-year-olds could produce metaphorical completions to target sentences, they were unable to reliably identify that the concepts juxtaposed in these sentences fell into separate categories. However, by four years of age, children who produced metaphors also showed an understanding that metaphorical statements involved concepts from different conventional categories. Both the 3- and 4-year-olds were able to identify anomalous from literal and metaphorical comparisons (see also Pearson, 1990). Vosniadou and Ortony interpreted these data as suggesting that children start with an undifferentiated notion of similarity, which at about the age of four becomes differentiated into literal and non-literal similarity. They suggested that the latter type forms the basis of metaphorical language comprehension.

In this paper, we describe the first computational model explaining the emergence of the distinction between literal and metaphorical similarity, based on an existing connectionist model of simple metaphor comprehension (Thomas & Mareschal, 2001). The importance of this model is that it directly relates the development of metaphor comprehension to the development of semantic representations. The structure of this paper is as follows. We begin by briefly reviewing connectionist approaches to metaphor comprehension. Second, we describe the main tenets of the Metaphor by Pattern Completion (MPC) model on which the developmental account is based. Third, we chart the development of category-specific representations that support metaphor comprehension and the distinction between literal and figurative statements within the MPC model. Finally, we discuss implications for the order of acquisition of such distinctions by young children.

Connectionist models of metaphor processing

First of all, it is important to point out that, although previous computational models have been proposed for the comprehension of metaphor, all of these models have related to the adult state, and none have contained a developmental component.

Previous models of metaphor comprehension have exploited the soft multiple constraint satisfaction abilities of connectionist networks to capture the interactions of conceptual domains when they are juxtaposed in comparisons. One class of models has focused on the potential of microfeature or vector representations of concepts to capture subtle interactions between knowledge bases (e.g., Chandler, 1991; Sun, 1995; Thomas & Mareschal, 2001). A second class of models has focused on structural mapping accounts of analogy formation, whereby target and vehicle domains are compared via the alignment of their relational structure, as well as evaluation of shared attributes (e.g., Holyoak & Thagard, 1989; Hummel & Holyoak, 1997). Why have computational models of metaphor comprehension been silent on developmental phenomena? The answer is that both classes of model have tended to include extensively pre-structured, domain-specific representations, which prevent them from exploring how representations (and their comparison) may emerge as a function of development.

In the present work, we will focus only on attribute mapping, which is readily captured by microfeature models, and put to one side problems of structural alignment. Although this limits the scope of the metaphors to which the model can be applied, it nevertheless makes the first initial steps towards exploring the developmental dimension of metaphor processing, and specifically, to investigating the ways in which metaphor comprehension can be linked to the development of semantic representations.

The MPC model

A full description of the MPC model can be found in Thomas and Mareschal (2001), along with an evaluation of its main assumptions. Here we provide a brief outline. In broad terms, the model suggests that, when presented with a metaphor such as *Richard is a lion*, the listener indeed attempts to fit the concept *Richard* into the category of *lion*; in so doing, an outcome of the categorization process is to alter the representation of *Richard* to make him more consistent with the features of a lion.

More specifically, metaphor comprehension is construed as a two-stage process. Consistent with Glucksberg and Keysar's (1990) view of metaphor comprehension as a type of categorization process, the first stage comprises *mis-classification* of a semantic input. A metaphor <A is B>, where A is the *topic* and B the *vehicle*, is comprehended by applying a representation of the first term (A) to a semantic memory network storing knowledge about the second term (B). Categorization is evaluated via the accuracy of reproduction of (A)'s representation in an autoassociator network trained on exemplars of (B). The degree of semantic distortion of (A) is a measure of the semantic similarity of concept A to domain B (Thomas & Mareschal, 2001).

However, the result of applying (A) to the network storing knowledge about (B) is a representation of (A) transformed to make it more consistent with the (B) knowledge base. In particular, there is an interaction in which features of (A) key into covariant structure between features in (B). If (A) shares some features of such covariant structures, it inherits further features by a process of *pattern completion*. Such feature inheritance depends on both terms, and provides an implementation of Black's (1979) well-known interaction theory of metaphor comprehension. However, enhancement of the features of (A) does not complete the process. In a second stage, the degree of meaning change of the topic is compared to the expected level of change given the current discourse context (Vosniadou, 1989). If the threshold is high, the statement is taken as literal and the full change in meaning is accepted. If it is at an intermediate level, only enhanced feature changes are accepted as the

communicative intent of a metaphor. If the threshold is at a low level, the sentence is rejected as anomalous.

Thomas & Mareschal (2001) evaluated the model's performance in comparing highly simplified domains to illustrate this process. Plausible metaphorical comparisons such as "the apple is a ball" were contrasted with anomalous comparisons such as "the apple is a fork". The model was able to account for a number of empirical phenomena, including the non-reversibility of comparisons and the predictability of interactions between topic and vehicle.

However, the degree to which metaphorical semantic transformations will occur depends not only on the similarity of (A) and (B), but also on the amount and quality of the knowledge stored in knowledge base B. In this way, metaphor comprehension can be linked to semantic development.

In the next section, we take a single vehicle knowledge domain and trace the development of metaphorical comprehension as the knowledge in the base network increases with learning. For simplicity, the sample knowledge base comprises information about *types of ball*, and performance is compared on literal comparisons ("the football is a ball") against metaphorical comparisons ("the pumpkin is a ball") and anomalous comparisons ("the kite is a ball").

The developmental model is intended to be illustrative: we make no claims about children's specific abilities to compare objects to balls at specific ages. Rather, we are interested in evaluating the effect of emerging semantic structure on the delineation of different types of similarity, and the consequent qualitative changes in the nature of metaphor comprehension during development.

Modeling the development of metaphor comprehension

Autoassociation is at the heart of the MPC mechanism. In the original model (Thomas & Mareschal, 2001), multiple parallel knowledge bases were available for different comparisons. However, in the present article and in the interest of clarity, we discuss only results obtained with a single autoassociator network.

A network with 16 input units, 16 output units, and 10 hidden units was trained to autoassociate a set of input patterns that defined the semantic knowledge of the vehicle domain. The number of hidden units was chosen to allow good training performance but also to encourage generalization. All units in the network used sigmoid activation functions.

The autoassociation network was trained for 500 presentations of the complete training set. At each epoch the training set was presented in a different random order. The learning rate and momentum were set to 0.05 and 0.0 respectively. Metaphor comprehension performance was evaluated at 0, 1, 2, 3, 4, 5, 7, 10, 15, 20, 30, 45, 70, 110, 200, and 500 epochs of training. The results reflect an average over n=12 replications with different initial random seeds.

The training set was constructed around 8 prototypes of various balls, constituting the 'ball' knowledge base. Prototypes were defined over 5 clusters of features: color (red, green brown, white), shape (round, irregular), consistency (soft, hard), size (small, large), weight (heavy, light), and associated action (thrown, kicked, hit, eaten), for a total of 16 semantic features. The last feature was included to permit anomalous and metaphorical comparisons. We assume that all concepts can be described by the same large feature set, and that the organization of knowledge into different categories happens within the hidden unit representations through learning. Feature values ranged between 0 and 1, so that the higher the activation, the more prominent the feature. Opposite feature values (e.g., small & large) were encoded on separate inputs to allow the coding of an absence of knowledge. From each prototype, 10 exemplars were generated by adding Gaussian noise (with standard deviation of 0.35) to the prototype pattern. The final training set thus constituted 8x10 = 80 exemplars of balls. The training prototypes are listed in Table 1, upper section.

Assessing different semantic comparisons

A comparison is evaluated by applying a novel input to the network and seeing how well it is reproduced on the output units. The more accurate the reproduction, the greater the similarity of the novel item to the knowledge stored within the network. Nine novel comparisons were created using the semantic features described above. These fell into three classes: (1) literal comparisons, (2) metaphorical comparisons, and (3) anomalous comparisons.

Literal comparisons involved novel exemplars of balls near the prototypical values. Metaphorical comparisons involved inputs that shared some properties with balls, but differed on other properties. Anomalous comparisons involved inputs constructed so that the inputs shared few features with balls in general.

The input vectors for the different classes of comparisons were constructed by comparing the novel input with the ball prototype vectors used to generate the knowledge training set. This was achieved by computing the angle between the two vectors in semantic space and selecting the closest match. For the literal comparisons, the angle had to be to be less than 10 degrees, for the metaphorical comparisons, it had to be between 40 and 45 degrees, and for the anomalous comparisons, it had to be between 60 and 66 degrees. (An angle of 90 degrees would constitute a novel pattern orthogonal to, or completely different from, all the prototypes used to generate the exemplars in the knowledge base.) Novel comparisons are shown in Table 1, lower section. A perfect reproduction of the input at the output indicates a similarity of 1.0 (self-similarity). The transformation similarity (S) of each novel comparison to the ball knowledge base was defined as:

$$S = 1 - RMS\ Error \quad (1)$$

An RMS error of 0 would give a similarity of S=1. High similarity implies low semantic distortion (as expected in a literal comparison), moderate similarity implies moderate semantic distortion (as expected in a metaphorical comparison), and low similarity implies high semantic distortion (as expected in an anomalous comparison). The similarity of novel comparisons was evaluated at different points during training. Principal Component Analyses of the hidden unit activations were also carried out during training to chart the development of the internal representations.

Results

Figure 1 shows *S* for each of the three types of comparison as learning progresses. Initially, there is little difference between literal, metaphorical and anomalous comparisons. However, even very early in learning a marked separation of the anomalous comparisons from the literal and metaphorical comparisons appears. The metaphorical and literal comparisons continue to be treated in a similar fashion for a further 5 presentations of the training set. At this point, metaphorical and literal similarities diverge. In the remaining epochs of training, the similarities from the three different types of comparisons separate into distinct bands. After an initial period of treating literal and metaphorical statements identically, the network has learnt to separate them out.

The process that underlies the development of this distinction can be better understood by examining the developing structure of the network knowledge base (Fig.2). Principal Components Analysis of the hidden unit activation space shows how the internal representations pull apart the different types of ball during training, according to their input characteristics.

In general, anomalous patterns fall in-between clusters, while metaphorical comparisons lie at the edge of clusters, and literal comparisons lie within the clusters. Once the clusters are sufficiently delineated from each other, an item that bears a metaphorical relation to a given category is distinguished from members of that category.

Novel inputs to the network are transformed in an attempt to classify them. Within the model, the transformed semantic representation corresponds to the meaning enhancement that is the outcome the comparison. Focusing on the metaphorical comparisons alone, examination of this enhancement yields three distinct phases during training. First, there is poor pattern completion, linked to an immature vehicle knowledge base. Next, with the initial emergence of semantic structure, metaphorical comparisons such as “the pumpkin is a ball” and “the apple is a ball” lead to enhancement of some of the target’s features. For example, ‘pumpkins’ and ‘apples’ are not associated with being ‘thrown’, ‘hit’, or ‘kicked’. The effect of each metaphor is to transfer such features from vehicle to topic. However, initially enhancement occurs according to an early, prototypical notion of ball, a notion that averages over all exemplars of balls, and corresponds to what one might call the *basic level* of the category. On average, most balls are ‘hit’ rather than ‘kicked’ or ‘thrown’. During this second phase, the ‘hit’ enhancement is inherited by all round, firm targets such as ‘apple’ and ‘pumpkin’. However, in the third phase, as further training produces delineation of the knowledge base, transfer now occurs according to the type of ball *most similar* to the particular target, according to

Table 1: Upper section: Prototypical patterns forming the ball knowledge base. Adding noise to the prototypes creates training sets. Lower section: Novel patterns used in literal, metaphorical, and anomalous comparisons.

	Color				Action				Shape		Consist-ency		Size		Weight	
Prototypes	Red	Green	Brown	White	Eaten	Thrown	Hit	Kicked	Round	Irregular	Soft	Hard	Small	Large	Heavy	Light
Football(white)	.00	.00	.00	.90	.00	.20	.00	.95	.90	.00	.00	.80	.00	.90	.90	.00
Football(brown)	.00	.00	.90	.00	.00	.20	.00	.95	.90	.00	.00	.80	.00	.90	.90	.00
Cricket ball	.90	.00	.00	.00	.00	.98	.97	.00	.90	.00	.00	.98	.80	.00	.80	.00
Ping-Pong ball	.00	.00	.00	.95	.00	.10	.98	.00	.98	.00	.00	.98	.95	.00	.00	.95
Tennis ball	.00	90	.00	.00	.00	.80	.98	.00	.90	.00	.80	.00	.80	.00	.00	.85
Squash ball (red)	.80	.00	.00	.00	.00	.50	.98	.00	.93	.00	.85	.00	.95	.00	.00	.90
Squash ball (green)	.00	.90	.00	.00	.00	.50	.98	.00	.93	.00	.85	.00	.95	.00	.00	.90
Beach ball	.98	.00	.00	.00	.00	.90	.90	.90	.90	.00	.98	.00	.00	.98	.00	.90
Novel comparisons																
Literal:																
Football (white)	.00	.00	.00	.85	.00	.20	.00	.98	.80	.00	.10	.80	.00	.90	.80	.00
Beach ball	.90	.00	.00	.00	.00	.80	.70	.90	.70	.00	.90	.00	.00	1.0	.00	.80
Ping-Pong ball	.00	.00	.00	.99	.00	.20	.99	.00	.95	.00	.00	.90	.95	.00	.00	.97
Metaphorical:																
Apple (red)	.80	.00	.00	.00	.95	.05	.00	.00	.75	.15	.70	.20	.70	.00	.00	.50
Pumpkin	.20	.00	.70	.00	.80	.00	.00	.00	.80	.50	.80	.60	.00	.80	.90	.00
Apple (green)	.00	.95	.00	.00	.95	.05	.00	.00	.75	.15	.70	.20	.70	.00	.00	.50
Anomalous:																
Kite	.99	.00	.00	.00	.00	.05	.00	.00	.00	.99	.00	.98	.00	.95	.20	.80
Spaghetti	.00	.00	.80	.20	.97	.00	.00	.00	.00	.70	.80	.20	.00	.70	.00	.60
Toast	.00	.00	.80	.10	.80	.00	.00	.00	.00	.80	.80	.00	.80	.00	.00	.90

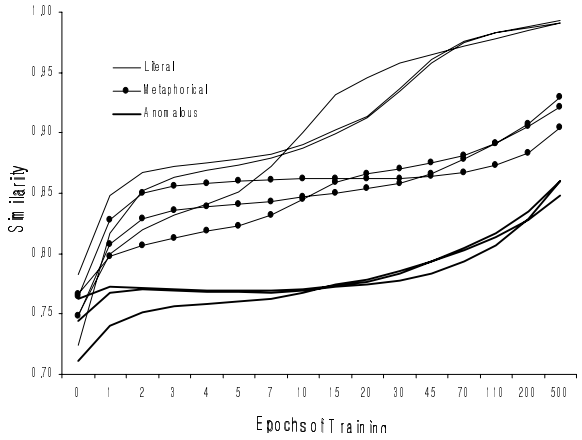

Figure 1. Similarity (*S*) of novel comparisons to the ball knowledge base during training. Three examples from each comparison type are plotted.

what one might call the *subordinate level* of the vehicle category. Table 2 shows that at 4 epochs 'apple' and 'pumpkin' have similar activation levels for the action features, loading maximally on 'hit', whereas at 500 epochs, 'apple' and 'pumpkin' now load on different features. Apples are now viewed as likely to be hit, and pumpkins to be kicked, according to their differing sizes. The model thus generates an explicit and testable prediction: *attribute inheritance will move from basic to subordinate level during development*.

Moreover, since there is variability within the internal structure of categories, not all literal comparisons will be equivalent. The more atypical the literal comparison, the more it will resemble a metaphor. This leads to a second explicit and testable prediction: *the recognition of atypical literal statements as distinct from metaphorical statements should lag behind the recognition of typical literal statements as distinct from metaphorical statements during development.*

Discussion

A common characterization of conceptual development views young children's knowledge as being assimilated into broad groups; as children develop, they make finer and finer distinctions until there are many different categories (e.g., Carey, 1985; Keil, 1986). Because the comprehension of metaphor requires the deliberate deconstruction of

Table 2: Attribute transfer from *basic* (4 epochs) and *subordinate* (500 epochs) category levels. Scores show the transformed feature values for action features Thrown (T), Hit (H) and Kicked (K) in the topic.

	Comparison <X is a Ball>					
	Red Apple			Pumpkin		
	T	H	K	T	H	K
4 epochs	.59	.75	.37	.50	.60	.30
500 epochs	.17	**.85**	.17	.18	.03	**.48**

a) 4 epochs

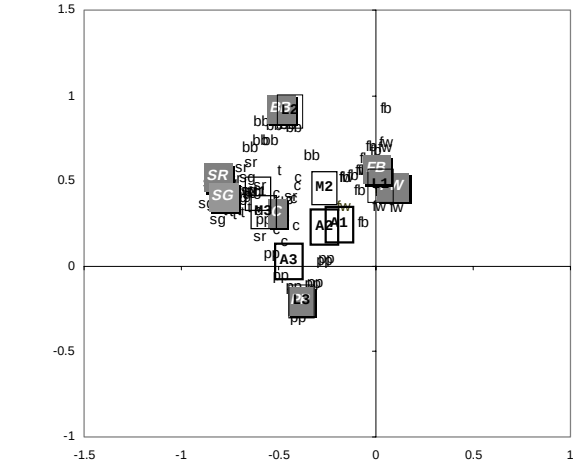

b) 15 epochs

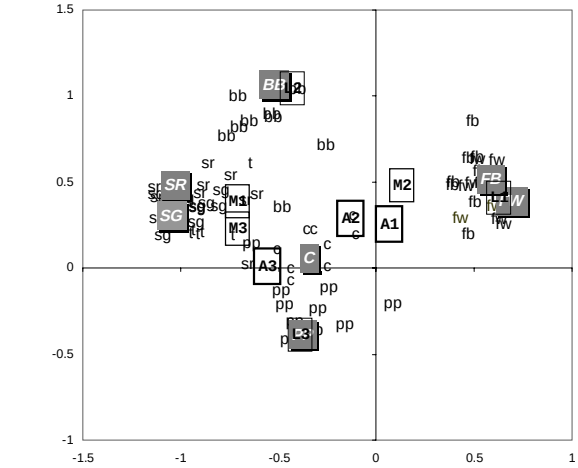

c) 500 epochs

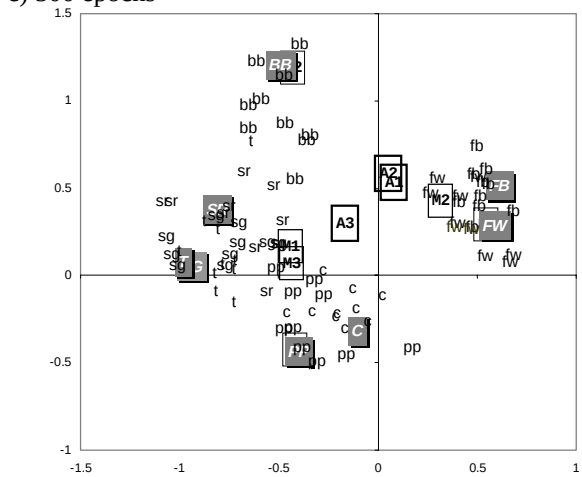

KEY

BB - Prototype (e.g. Beach Ball)
bb - Exemplar in training set
L1 - Literal comparison
M1 - Metaphorical comparison
A1 - Anomalous comparison

Figure 2. First two components of the hidden unit activations for training and test patterns of a representative network across training.

categories, the way knowledge is categorized will have a large effect on metaphor comprehension. The model we have described above provides a concrete implementation of Marschark and Nall's (1985) account of metaphor use in young children. Literal, metaphorical, and anomalous comparisons fall onto a conceptual space undergoing refinement. The process of refinement leads to the emergence of a notion of non-literal similarity.

Clearly this simple model does not capture all aspects of the development of metaphor comprehension. The metaphors we have dealt with are predominantly perceptual. Importantly, the model fails to capture the emerging use of structural information in children's metaphors (Gentner, 1988). However, existing computational models have not addressed develop-mental phenomena at all, let alone the relational shift. The next step for the MPC model will be an extension to structured representations, possibly via the inclusion of synchrony binding (see Hummel & Holyoak, 1997), while retaining the mechanism of pattern completion as a powerful tool for explaining the transfer of attributes in metaphorical comparisons. Despite its simplicity, the importance of the current model is its demonstration that the emergence of non-literal similarity can be driven by emerging semantic structure, and the explicit testable hypotheses it generates to progress our understanding of the development of metaphor comprehension in young children.

Acknowledgements

MT's work on this paper was funded by MRC Project Grant no. G9809880 and MRC Programme Grant no. G9715642; DM's work was funded in part by European Commission RTN Grant CT-2000-00065 and ESRC Grant no. R000239112.

References

Billow, RM (1981). Observing spontaneous metaphor in children. *Journal of Experimental Child Psychology*, *31*, 430-445.

Black, M (1979). More about metaphor. In A. Ortony (Ed.). *Metaphor and thought*. Cambridge, UK: Cambridge University Press.

Carey, S (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press.

Chandler, SR (1991). Metaphor comprehension: A connectionist approach to implications for the mental lexicon. *Metaphor & Symbolic Activity*, *6*, 227-258.

Dowker, A (1999). Metaphor: Is it the same for children and adults? In *Proceedings of the AISB'99 Symposium on Metaphor, Artificial Intelligence, and Cognition.* AISB: Brighton, UK.

Gardner, H, Winner, E, Becchoffer, R & Wolf, D (1978). The development of figurative language. In K. Nelson (Ed.) *Children's language Vol.1*. Cambridge, MA: Gardner Press Inc.

Gentner, D (1988). Metaphor as structure mapping: The relational shift. *Child Development, 59*, 47-59.

Gibbs, RW (1994). *The poetics of mind.* Cambridge, UK: Cambridge University Press.

Glucksberg, S & Keysar, B (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*, *97*, 3-18.

Goswami, U (in press). Analogical reasoning in children. In D. Gentner, K.J. Holyoak, Kokinov (Eds.). Analogy: Interdisciplinary perspectives.

Gottfried, GM (1997). Comprehending compounds: evidence for metaphoric skill? *Journal of Child Language*, *24*, 163-186.

Hakes, D (1982). The development of metalinguistic awareness: What develops? In S. Kuczaj (Ed.) *Language development Vol.2*. Hillsdale N.J.: Erlbaum

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

Hummel, JE & Holyoak, KJ (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*, 427-466.

Keil, FC (1983). On the emergence of semantic and conceptual distinctions. *Journal of Experimental Psychology: General*, *112*, 357-385.

Marschark, M & Nall, C (1985). Metaphor competence in cognitive and language development. *Advances in Child Development and Behavior, 26*, 49-78.

Pearson, B (1990). The comprehension of metaphor by preschool children. *Journal of Child Language*, *17*, 185-203.

Piaget, J (1962). *Play, dreams, and imitation in childhood*. New York, NY: Norton.

Sun, R (1995). A microfeature based approach towards metaphor interpretation. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence* (pp. 424-429).

Thomas, MSC & Mareschal, D (2001). A connectionist model of metaphor by pattern completion. *Metaphor & Symbol*, *16*, 5-27.

Vosniadou, S (1989). Context and the development of metaphor comprehension. *Metaphor & Symbolic Activity*, *4*, 159-171.

Vosniadou, S & Ortony, A (1983). The emergence of the Literal-Metaphorical-Anomalous distinction in young children. *Child Development*, *54*, 154-161.

Winner, E (1988). *Invented worlds: The psychology of the arts.* Cambridge, MA: Harvard University Press.