

A Fluid Knowledge Representation for Understanding and Generating Creative Metaphors

Tony Veale

School of Computer Science
University College Dublin
Ireland
tony.veale@ucd.ie

Yanfen Hao

School of Computer Science
University College Dublin
Ireland
yanfen.hao@ucd.ie

Abstract

Creative metaphor is a phenomenon that stretches and bends the conventions of semantic description, often to humorous and poetic extremes. The computational modeling of metaphor thus requires a knowledge representation that is just as stretchable and semantically accommodating. We present here a flexible knowledge representation for metaphor interpretation and generation, called *Talking Points*, and describe how talking points can be acquired on a large scale from WordNet (Fellbaum, 1998) and from the web. We show how talking points can be fluidly connected to form a slipnet, and demonstrate that talking points provide an especially concise representation for concepts in general.

1 Introduction

Metaphor serves two important roles in language. The first of these is to make the unfamiliar and the strange seem more familiar and understandable (Indurkha, 1992). For instance, one might describe a burqa (a full body covering for Muslim women) as a suit of armor, as a shield against prying eyes or, depending on one's communication goal, as a wearable cage. The other role of metaphor is most often associated with the poetic and fanciful use of language, but is no less important: to make the familiar and mundane seem strange and unfamiliar. In this latter guise, metaphor allows us to view a commonplace idea from a new and revealing category perspective

(Camac and Glucksberg, 1984). For instance, one might describe makeup as “the Western burqa”, to communicate not just the idea that each involves a covering of the female form, but that each reflects a society-imposed expectation on the public presentation of women. Each of these roles is a manifestation of the same underlying mechanism for combining concepts, for understanding how they interact (Black, 1962) and for determining how they are connected (Fauconnier and Turner, 1998), even if those connections are tenuous, hidden or not always obvious (Collins and Loftus, 1975). For example, consider the connections needed to make and understand the above metaphors:

Burqa ⇒

≡ *for concealing a Muslim woman*
≈ *for protecting a Muslim woman*
≈ *protecting a woman*
≈ *for protecting a person*

Armor ⇐

Make-up ⇒

≡ *typically worn by women*
≈ *expected to be worn by women*
≈ *must be worn by women*
≈ *must be worn by Muslim women*

Burqa ⇐

In each case we see how metaphor draws out and highlights, in a modified or exaggerated form, an existing aspect of each target concept. In other words, metaphor does not indiscriminately transplant arbitrary aspects of a source concept onto a target, but accommodates a selective graft of the most salient aspects of this source concept onto those aspects of the target that can

© 2008. Licensed under the *Creative Commons Attribution-Noncommercial-Share Alike 3.0 Unported* license (<http://creativecommons.org/licenses/by-nc-sa/3.0/>). Some rights reserved.

be highlighted by the comparison (Ortony, 1979). This connection between concepts requires a flexible knowledge representation, one that allows the connections between non-identical source and target aspects to be recognized, reconciled and even compressed (Fauconnier and Turner, 1998). This fluid representation (Hofstadter *et al.*, 1995) defines the search space in which the processes of metaphor generation and metaphor understanding are cognitively situated (Veale and O'Donoghue, 2000): for generation, fluid connectivity allows a system to search outwards from a given target to find those potential source concepts that offer a new yet appropriate perspective; for understanding purposes, connectivity allows an agent to focus on those key aspects of a source concept that are most apt for a target because they can be linked to that target.

In this paper we describe the construction of a fluid knowledge representation for creative metaphor processing, one that is acquired automatically from WordNet (Fellbaum, 1998) and from the texts of the web. In section 2 we summarize related work in the field of metaphor as it pertains to flexible knowledge representation. In section 3 we describe two complementary means of acquiring the basic elements of this representation, from WordNet and from the web, before describing how these elements can be placed into a fluid network of connections - what Hofstadter (*ibid*) calls a *slipnet* - in section 4. We then present in section 5 some empirical evaluation of the acquired representation on an objective test of term categorization, before concluding with some consideration of future work in section 6.

2 Related Work

Since metaphor can be viewed as a stretching of linguistic conventions to cover new conceptual ground, the interpretation of metaphor crucially hinges on a system's ability to recognize these conventions and accommodate the exceptional meaning conveyed by each figurative expression. Indeed, most computational approaches embody a sense of what it means to be literal, and accommodate metaphoric meanings within this conventional scheme through a form of relaxation, mapping or translation. Wilks (1978) advocates that the typically *hard* constraints that define a literal semantics should instead be modeled as *soft* preferences that can accommodate the violations that

arise in metaphoric utterances, while Fass (1991) builds on this view to show how these violations can be repaired to thus capture the literal intent behind each metaphor. This repair process in turn relies on the availability of a concept taxonomy through which metaphoric uses can be mapped onto their literal counterparts; a car that "drinks gasoline" would thus be understood as a car that "consumes gasoline". Way (1991) emphasizes the importance of this taxonomy by positing a central role for a dynamic type hierarchy (DTH) in metaphor, one that can create new and complex taxonomies on the fly. For instance, Way's DTH would understand the "make-up as Western burqa" metaphor via a dynamically created taxonym like *things-women-are-expected-to-wear-in-public*, though Way offers no algorithmic basis for the workings of such a remarkable taxonomy. Another family of computational approaches combines explicit knowledge about certain metaphors with knowledge about the domains connected by these metaphors. Martin's (1990) *Midas* system encodes schematic knowledge about conventionalized metaphors such as "to kill a process" and "to open a program", and uses this knowledge to fit novel variations of these metaphors into the most appropriate schemas. Barnden and Lee (2002) focus on the role of inference in a metaphorically-structured domain, and describe a system called *ATTMeta* that contains sufficient knowledge about e.g., conventional metaphors of mind to reason about the mental states implied by these metaphors. Each of these approaches sees metaphor interpretation as a process of fitting what is said to what can meaningfully be represented and reasoned about. This fitting process is most explicitly modelled by Hofstadter *et al.* (1995), who focus on the slippage processes that are required to understand analogies in abstract domains that e.g., involve the mapping of letter sequences or the mirroring of actions in a highly stylized tabletop environment. Though simplified and toy-like, these are non-deterministic problem domains that are nonetheless shaped by a wide range of pragmatic pressures. Hofstadter and Mitchell (1994) model these pressures using a *slipnet*, a probabilistic network in which concepts are linked to others into which they can slip or be substituted with. In this view, deeply embedded concepts that are further removed from direct observation are less likely to engage in slippage than more superficial concepts.

To take a linguistic example, word choice in natural language generation is more susceptible to slippage (as influenced by synonym availability) than the concepts underlying the meaning of a sentence.

Slippage can be seen as a lossy form of conceptual re-representation: the greater the slippage, the more dramatic the re-representation and the greater the potential for loss of accuracy. For instance, a recent magazine cover proclaims the governor of California, Arnold Schwarzenegger, as “president of 12% of the United States”. This conceptualization can be viewed as an intermediate stage in a slippage path from Governor to President as follows:

Governor of California ⇒

≈ *governor of 12% of the United States*
 ≈ *leader of 12% of the United States*
 ≈ *president of 12% of the United States*
 ≈ *president of 100% of the United States*

President of the U.S. ⇐

This labeling is creative enough to grace a magazine cover because it involves an ambitious level of re-conceptualization, at least from a computational perspective. The pivotal insight comes from *California = 12% of the United States*, an ad-hoc synonym that one is unlikely to find in a dictionary or general-purpose resource like WordNet. While ultimately aiming for this kind of creative transformation, our goal in this paper is more modest: to build a slippage network of concepts that are connected via their most salient features, one that combines the principled flexibility of a Hofstadter-style slipnet with the comprehensive scale of a resource like WordNet.

3 Acquiring Conceptual Talking Points

We refer to the knowledge elements connected by this slipnet as conceptual *talking points*. We first describe the form of these talking points and how they are acquired, before describing in section 4 how slippage operates between these talking points. We discuss two complementary kinds of talking point here: objective descriptions, extracted from WordNet glosses, and informal, stereotypical descriptions, harvested from the text of the web via a search engine like Google.

3.1 Objective Talking Points

Objective talking points are aspects of conceptual description that contribute to the consensus definitional view of a concept. Though WordNet does not provide explicit semantic criteria for the definition of each lexical concept, many of these criteria can be gleaned from a shallow parse of the pithy dictionary gloss it associates with each (e.g., see Ahlswede and Evans, 1988). Thus, whenever the head phrase of a concept’s gloss has the form “ADJ+ NOUN” where NOUN can denote a hypernym of the concept, we can associate the talking point *is_ADJ:NOUN* with that concept. For example, the gloss of {*Hamas*} is “a militant Islamic fundamentalist political movement that ...”, which yields the talking points *is_militant:movement*, *is_islamic:movement*, *is_fundamentalist:movement* and *is_political:movement* for *Hamas*. When a WordNet concept has a hypernym of the form {ADJ_NOUN}, where NOUN can denote a hypernym of this concept, we likewise associate the talking point *is_ADJ:NOUN* with that concept. For example, {*Taliban*, *Taleban*} has {*religious_movement*} as a hypernym, which yields *is_religious:movement* as a talking point for *Taliban*.

Objective talking points can also be gleaned from the subject-verb-object structure of a WordNet gloss. For instance, the gloss for synset {*conductor*, *music_director*} is “the person who leads a musical group”, which yields the talking point *leads:musical_group*. The hypernym of this concept, {*musician*}, has the gloss “*artist who composes or conducts music ...*”, which yields the talking points *composes:music* and *conducts:music* that are then inherited by {*conductor*, ...} and other sub-types of *musician* in WordNet. A shallow parse will generally not lead to a complete understanding of a concept, but will typically produce some interesting talking points of the *predicate:object* variety that can be used to relate a concept to others that are analogically or metaphorically similar. Using WordNet’s noun and verb taxonomies, we can identify the following slippage paths between talking points:

composes:music → *composes:speech* → *writes:speech* → *writes:oration* → *writes:sermon*
 → *writes:law* → *writes:philosophy* → *writes:theorem* → *writes:plan* → ...

In all, we extract talking points of the form *is_adj:noun* for over 40,000 WordNet concepts, and talking points of the form *verb:noun* for over 50,000 concepts. However, the real power of talking points emerges when they are connected to form a slipnet, as we discuss in section 4.

3.2 Stereotypical Talking Points

The talking points we harvest from the web do not have the authoritative, definitional character we find in hand-crafted resources like WordNet, but they do reflect how people typically speak of (and, perhaps, actually think of) the world. Veale and Hao (2007) argue that similes present the clearest window into the stereotypical talking points that underpin everyday conversations, and collect from the web instances of the pattern “*as ADJ as a **” for thousands of WordNet adjectives. Though the simile frame is somewhat leaky in English, and prone to subversion by irony, Veale and Hao construct a comprehensive database of more than 12,000 highly stereotypical adjective:noun associations, such as *precise:surgeon*, *straight:arrow*, *balanced:pyramid* and *sharp:knife*. We use their data here, as the basis of an additional web harvesting process to gather stereotypical talking points of the form *has_ADJ:facet*. For every stereotypical association ADJ:NOUN in their database, we send the query “*the ADJ * of a|an|the NOUN*” to Google and collect noun values for the wildcard * from the first 200 hits returned for each query.

This pattern allows us to determine the conceptual attributes that are implicit in each stereotypical *adjective:noun* pairing. For instance, “the delicate hands of a surgeon” and “the inspiring voice of a preacher” reveal that *hand* is a salient attribute of surgeons while *voice* is a salient attribute of preachers. The frequency with which we find these attributes on the web also allows us to build a textured representation for each concept. So while these expanded web patterns also reveal that surgeons have a thorough *eye* and steady *nerves*, “the hands of a surgeon” are mentioned far more frequently and are thus far more salient to our understanding of surgeons. To avoid noise, the set of allowable attribute nouns, such as *hands*, *soul*, *heart*, *voice*, etc., is limited to the nouns in WordNet that denote a kind of trait, body part, quality, activity, ability or faculty. This allows us to acquire meaningful talking points like *has_magical:skill* for Wizard, *has_brave:spirit* for

Lion and *has_enduring:beauty* for Diamond, while avoiding dubious or misleading talking points like *has_proud:owner* for Peacock that lack either representational value or insight. In all, this process acquires 18,794 stereotypical talking points for 2032 different WordNet noun senses, for an average of 9 facet:feature pairs per sense. Specific senses are identified automatically, by exploiting WordNet’s network of hypernymy and synonymy relations to connect talking points that describe variations of the same concept.

4 Building a Slipnet of Talking Points

To construct a slipnet in the style of Hofstadter and Mitchell (1994), but on the scale of WordNet, we need to connect those talking points that express similar but different meanings, and to quantify the difference between these meanings. Issues of scale mean that we need only connect talking points that are close in meaning, since greater slippage can be achieved by following longer paths through the slipnet. This slippage can be based on semantic or pragmatic criteria. Thus, the talking points *has_sacred:authority* (for Pope) and *has_sacred:power* (for God) are semantically similar since the potency sense of “authority” is a specialization of the control sense of “power” in WordNet. Likewise, *writes:speech* and *composes:speech* are similar because “compose” and “write” are synonymous in the context of literary creation, and it is this particular linkage that supports a slippage pathway from *composes:music* to *writes:poetry*. In contrast, *is_political:movement* (for Hamas) and *is_religious:movement* (for Taliban) are pragmatically similar since movements that are religious often have a political agenda also. We can use WordNet to construct the semantic links of the slipnet, but pragmatic links like these require not just word senses but a sense of the world, of a kind we can distil from the text of the web.

Two talking points *is_ADJ₁:OBJ₁* and *is_ADJ₂:OBJ₂* should be connected in the slipnet if: OBJ₁ and OBJ₂ are semantically close (i.e., synonymous, or semantic siblings in WordNet); and ADJ₁ and ADJ₂ are synonymous, or ADJ₁ frequently implies ADJ₂ or ADJ₂ frequently implies ADJ₁. These implications are recognized and quantified using another web trawling process, in which the query “*as * and * as*” is used to harvest pairs of adjectives that are seen to

mutually reinforce each other in web comparisons. This search reveals that “religious” reinforces “superstitious” (5 times), “moral” (4), “political” (3), “conservative” (3), “intolerant” (2) and “irrational” (1). These slippage connections link *is_religious:movement* to *is_political:movement* (pragmatic) to *is_political:campaign* (semantic) to *is_military:campaign* (pragmatic), thereby connecting Taliban (*is_religious:movement*) to Crusade (*is_military:campaign*).

4.1 The Slipnet in Action

Slippage is a phenomenon best explained with an example, so consider again the task of creating metaphors for the concept Pope. We have already seen that slippage among talking points allows Pope to be linked to the concept God via Pope → *has_sacred:authority* → *has_sacred:power* ← God. Pope can also be linked to Rabbi via the path Pope → *has_sacred:words* → *has_wise:words* ← Rabbi and to Judge by extending this pathway: Pope → *has_sacred:words* → *has_wise:words* → *has_solemn:words* ← Judge. Black (1962) saw metaphor as an interaction between concepts, in which the interpretation of a particular source concept depends crucially on how it is able to interact with a specific target concept. This concept-sensitive interplay is clearly on display here. The Pope can be metaphorically viewed as a warrior not by considering what it means for a generic person to be a warrior, but by considering how the concept Pope actually interacts with the concept Warrior, e.g., Pope → *has_infallible:voice* → *has_powerful:voice* ← Warrior.

Consider the potential for slippage between objective talking points from WordNet:

Pope ⇒
 ≡ *leads:Roman_Catholic_Church*
 ≈ *leads:congregation*
 ≈ *leads:flock*
 ≈ *leads:mob*
 ≈ *leads:organized_crime*

Don (Crime Father) ⇐

Pope ⇒
 ≡ *leads:Roman_Catholic_Church*
 ≈ *leads:congregation*
 ≈ *leads:political_movement*
 ≈ *leads:gang*
 ≈ *leads:military_force*

Warlord (Military Leader) ⇐

One can typically terminate a slippage path at any point, to produce different metaphors with varying semantic similarity to the starting concept. Thus, at *leads:flock* one can reach Shepherd, and from *leads:political_movement*, one can reach Civil_rights_leader.

A lexicon alone, like WordNet, is generally insufficient for metaphor processing, but such a resource can still reveal useful lexical resonances that may enrich an interpretation. In the example above, we see a resonance between the Pope, which WordNet also lexicalizes as “holy father”, and a mafia Don, which WordNet also lexicalizes as “father”. Indeed, since WordNet conceptualizes Roman_Catholic_Church as a specialization of Organized_religion, the metaphor establishes a parallelism between crime and religion as organized activities.

5 Empirical Evaluation

To understand whether talking points are sufficiently descriptive of the concepts they are acquired for, we replicate here the clustering experiments of Almuhareb and Poesio (2004, 2005) which are designed to measure the effectiveness of web-acquired conceptual descriptions. Since Almuhareb and Poesio use WordNet as a semantic gold-standard, we consider here the effectiveness of stereotypical talking points alone; it would be circular to consider objective talking points, since these are extracted from WordNet.

Almuhareb and Poesio describe two different clustering experiments. In the first, they choose 214 English nouns from 13 of WordNet’s upper-level semantic categories, and proceed to harvest property values for these concepts from the web using the pattern “*a|an|the * C is|was*”. This pattern yields a combined total of 51,045 values for all 214 nouns; these values are primarily adjectives, such as *hot*, *black*, etc., but noun-modifiers of C are also allowed, such as *fruit* for cake. They also harvest 8934 attribute nouns, such as *temperature* and *color*, using the query pattern “*the * of the C is|was*”. These values and attributes are then used as the basis of a clustering algorithm to partition the 214 nouns back into their original 13 categories. Comparing these clusters with the original WordNet-based groupings, Almuhareb and Poesio report a cluster accuracy of 71.96% using just values like *hot* (all 51,045), an accuracy of 64.02% using just attributes like *temperature* (all 8934), and

Table 1: Experiment 1, accuracy for 214 nouns

<i>Approach</i>	<i>Values only</i>	<i>Attr's only</i>	<i>All (V + A)</i>
Almu. + Poesio	71.96% (51045 vals)	64.02% (8934 attr)	85.51% (59979 v+a)
Talking Points	70.2% (2209 vals)	78.7% (4974 attr)	90.2% (7183 v+a)

an accuracy of 85.5% using both together (59979 features).

In a second, larger experiment, Almuhareb and Poesio select 402 nouns from 21 different semantic classes in WordNet, and proceed to harvest 94,989 property values (again mostly adjectives) and 24,178 attribute nouns from the web using the same retrieval patterns. They then applied the *repeated bisections clustering* algorithm to this larger data set, and report an initial cluster purity measure of 56.7% using only property values like *hot*, 65.7% using only attributes like *temperature*, and 67.7% using both together. Suspecting that noisy features contribute to the perceived drop in performance, those authors then applied a variety of noise filters to reduce the value set to just 51,345 values and the attribute set to just 12,345 attributes, for a size reduction of about 50% in each case. This in turn leads to an improved cluster purity measure of 62.7% using property values only and 70.9% using attributes only. Surprisingly, filtering actually appears to reduce the clustering performance of both data-sets used together, to 66.4%.

We replicate here both of these experiments using the same data-sets of 214 and 402 nouns respectively. For fairness, we collect *raw* descriptions for each of these nouns directly from the web, and use no filtering (manual or otherwise) to remove poor or ill-formed descriptions. We thus use the pattern “*as * as a|an|the C*” to collect 2209 raw adjectival values for the 214 nouns of experiment 1, and 5547 raw adjectival values for the 402 nouns of experiment 2. We then use the pattern “*the ADJ * of a|an|the C*” to collect 4974 attributes for the 214 nouns of experiment 1, and 3952 for the 402 nouns of experiment 2; in each case, ADJ is bound to the raw adjectival values that were acquired using “*as * as a|an|the C*”. A comparison of clustering results is given in Tables 1 and 2. These tables illustrate that clustering is most effective when it

Table 2: Experiment 2, accuracy for 402 nouns

<i>Approach</i>	<i>Values only</i>	<i>Attr's only</i>	<i>All (V + A)</i>
Almu. + Poesio (no filtering)	56.7% (94989 vals)	65.7% (24178 attr)	67.7% (119167 v+a)
Almu. + Poesio (with filtering)	62.7% (51345 vals)	70.9% (12345 attr)	66.4% (63690 v+a)
Talking Points	64.3% (5547 vals)	54.7% (3952 attr)	69.85% (9499 v+a)

is performed on the basis of both values *and* attributes (yielding the highest scores, 90.2% and 69.85%, in each experiment respectively). These results thus support the combination of conceptual attributes with specific adjectival values into integrated talking points which reflect how people actually talk about the concepts concerned.

6 Conclusions

Metaphor is a knowledge-hungry phenomenon, so any computational treatment of metaphor will only be as good as the knowledge representation that supports it. Moreover, from a computational perspective, any theory of metaphor – cognitive, linguistic, or otherwise – is only as good as the algorithmic and representational insights that it provides, and the scale of the implementation that it ultimately allows us to realize. In this paper we have given computational form to some of the key insights in the metaphor literature, from the interaction theory of Black (1962) to the salience imbalance theory of Ortony (1979) to the theory of conceptual blending of Fauconnier and Turner (1998). We also employ a key insight from the work of Hofstadter and his fluid analogies group (1995), that robust reasoning on a conceptual level requires a degree of slippage that must be supported by the underlying knowledge representation.

Our knowledge base of talking points is derived from two complementary information sources: the objective definitions contained in WordNet (Fellbaum, 1998) and the stereotypical comparisons that pepper the texts of the web. These sources yield a knowledge-base that is neither small nor hand-crafted. While the knowledge-base needs to grow by at least an order of magnitude, slip-

page means that non-identical talking points can be treated as equivalent for purposes of robust processing, which in turn extends the *halo* of talking points that surrounds each concept in the knowledge-base (Hofstadter *et al.*, 1995). The experiments of section 5 also indicate that, in a pinch, new talking points for a previously under-represented concept can be acquired dynamically from the web with reasonable accuracy. As it currently stands, the talking points approach to metaphor is robust enough and scalable enough to generate simple but imaginative metaphors on demand for a wide range of user inputs.

But what does it mean to state, at a knowledge-representation level, that *lions* and *knights* both have a *brave heart*, that *wolves* and *tyrants* both have a *cruel face*, or that *eagles* and *warriors* have a *fierce expression*? Stereotypical talking points such as these can be poetic or metaphorical, and may express a viewpoint that is overly simplistic, subjective or even technically inaccurate. Nonetheless, our experiments suggest that the linguistic insights we acquire from non-literal descriptions strongly reflect our ontological intuitions about concepts and are more than mere linguistic decorations. Most significantly, we see from these experiments that stereotypical talking points yield an especially concise representation, since with no filtering of any kind, this approach achieves comparable clustering results with feature sets that are many times smaller than those used in previous work. We anticipate therefore that stereotypical descriptions will be a key growth area for the development of the talking points knowledge-base.

The Pope examples of section 4.1. exemplify the competence of the system as it is currently implemented, while the Burqa and Governor/President examples of sections 1 and 2 mark out our future directions. The Burqa examples demonstrate the need for a more complex representation of talking points that can accommodate nested propositions, while the Governor example demonstrates the need for more radical and ad-hoc slippage patterns in creative metaphors. Rather than add special rules to handle such individual cases (which are creative *because* of their one-off disposal nature), our ambition is to develop a general corpus-grounded mechanism for explaining *all* metaphor-related slippage. We remain a considerable distance from this goal, yet believe

it is best attained using the kind of robust and scalable approach described here.

References

- Almuhareb, A. and M. Poesio. 2004. Attribute-Based and Value-Based Clustering: An Evaluation In *proceedings of EMNLP*, the Conference on Empirical Methods on Natural Language Processing. Barcelona.
- Almuhareb, A. and M. Poesio. 2005. Concept Learning and Categorization from the Web. In *proceedings of the annual meeting of the Cognitive Science society*. Italy.
- Ahlsvede, T. and M. Evans. 1998. Parsing vs. Text Processing in the analysis of dictionary definitions. In *proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, 217–224.
- Barnden, J. A. and M. G. Lee. 2002. An Artificial Intelligence Approach to Metaphor Understanding. *Theoria et Historia Scientiarum*, 6(1):399–412.
- Black, M. 1962. *Models and Metaphor: studies in language and philosophy*. Ithaca, NY: Cornell University Press.
- Camac, K. and S. Glucksberg. 1984. Metaphors do not use associations between concepts, they are used to create them. *Journal of Psycholinguistic Research*, 13(6).
- Collins, A. and E. F. Loftus. 1975. A Spreading-Activation Theory of Semantic Processing. *Psychological Review* 82, 407–428.
- Fauconnier, G. and M. Turner. 1998. Conceptual Integration Networks. *Cognitive Science*, 22(2):133–187.
- Fellbaum, C. (ed.). 1998. *WordNet: An electronic lexical database*. The MIT Press. 1985 *A comprehensive grammar of the English*.
- Hofstadter, D. R. and M. Mitchell. 1994. The Copycat Project: A model of mental fluidity and analogy-making. In *Holyoak, K.J. & Barnden, J. A. (Eds.) Advances in Connectionist and Neural Computation Theory*, Vol. 2. Norwood, NJ: Ablex.
- Hofstadter, D. R. and the Fluid Analogy Research Group. 1995. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. NY: Basic Books.
- Fass, D. 1991. Met*: a method for discriminating metonymy and metaphor by computer. *Computational Linguistics*, 17(1):49–90.
- Indurkha, B. 1992. *Metaphor and Cognition: Studies in Cognitive Systems*. Kluwer Academic Publishers, Dordrecht: The Netherlands.

- Martin, J. H. 1990. *A Computational Model of Metaphor Interpretation*. NY: Academic Press.
- Ortony, A. 1979. Beyond literal similarity. *Psychological Review*, 86, 161–180.
- Veale, T. and Y. Hao. 2007. Making WordNet Functional and Context-Sensitive. In *proceedings of the 45th Annual Meeting of the Association for Computational Linguistics*. Czech Republic.
- Veale, T. and O'Donoghue. 2000. Computation and Blending. *Cognitive Linguistics*, 11(3–4), special issue on Conceptual Blending.
- Way, E. C. 1991. Knowledge Representation and Metaphor. *Studies in Cognitive systems*. Holland: Kluwer.
- Wilks, Y. 1978. Making Preferences More Active, *Studies in Cognitive systems. Artificial Intelligence* 11(3), 197–223.