Comprehending and Generating Apt Metaphors: A Web-driven, Case-based Approach to Figurative Language

Tony Veale and **Yanfen Hao**
School of Computer Science and Informatics, University College Dublin
Dublin, Ireland
{Tony.Veale, Yanfen.Hao}@ucd.ie

Abstract

Examples of figurative language can range from the explicit and the obvious to the implicit and downright enigmatic. Some simpler forms, like simile, often wear their meanings on their sleeve, while more challenging forms, like metaphor, can make cryptic allusions more akin to those of riddles or crossword puzzles. In this paper we argue that because the same concepts and properties are described in either case, a computational agent can learn from the easy cases (explicit similes) how to comprehend and generate the hard cases (non-explicit metaphors). We demonstrate that the markedness of similes allows for a large case-base of illustrative examples to be easily acquired from the web, and present a system, called Sardonicus, that uses this case-base both to understand property-attribution metaphors and to generate apt metaphors for a given target on demand. In each case, we show how the text of the web is used as a source of tacit knowledge about what categorizations are allowable and what properties are most contextually appropriate. Overall, we demonstrate that by using the web as a primary knowledge source, a system can achieve a robust and scalable competence with metaphor while minimizing the need for hand-crafted resources like WordNet.

Introduction

Language does not operate in a social vacuum, making knowledge of the world an essential part of the interpretation of most linguistic utterances. This dependence on knowledge is particularly true of figurative language. Metaphors, for instance, can often be enigmatic and challenging, conveying elaborate meanings that can only be unpacked by listeners that share the appropriate conceptual structures (Lakoff and Johnson, 1980). Most computational theories of metaphor are thus, by necessity, also theories of conceptual structure and knowledge representation (e.g., see Martin, 1990; Way, 1991; Fass, 1991; Veale and Keane, 1997; Barnden and Lee, 2002). Fortunately, figurative language can often be as revealing as it is enigmatic, allowing us to directly see its key mechanisms at work. For example, explicit similes of the form ”T is as P as [a|an] V” tell us not only that P is shared by T and V, but that P is an especially salient property of V, one that is most likely implicated in other similes or metaphors involving V as a vehicle.

For the purposes of constructing a case-base of figurative nouns and their most salient properties, similes, unlike metaphors, have the added advantage of explicit syntagmatic marking (Hanks, 2004). Explicit similes are thus the low hanging fruit of figurative language, and are easily identifiable in large corpora or on the worldwide web. In this paper we describe both how to acquire a large case-base of common similes from the web, and how to use the web as a source of world knowledge in the comprehension and generation of apt metaphors and similes using this case-base. Given the often dubious nature of most web content, great care must be exercised whenever mining knowledge from the web, and this is especially the case with examples of figurative language. As we shall show, many web similes are ironic in intent, to the extent that a CBR system with a tin ear for irony will exhibit a much reduced competence with figurative utterances.

The necessity of separating bona-fide similes (those that say what they mean) from ironic similes means that the computational system we describe here, named *Sardonicus*, effectively has two case-bases to draw upon, and two modes of operation, ironic and bona-fide. We focus primarily on the latter in this paper, and explain how new bona-fide similes can be learned on-the-fly for noun vehicles that are otherwise under-represented in the case-base. We begin in section 2 with a brief discussion of related computational work, before describing the case acquisition process in section 3. The case-based approach to comprehension and generation in section 4 is then motivated in terms of two psychological theories, Ortony's (1993) theory of salience-imbalance and Glucksberg's (2001) theory of category inclusion. In section 5 we demonstrate the descriptive sufficiency of the case-base, before concluding in section 6 with some empirical observations on the role of topical web-derived knowledge in simile and metaphor.

Related Work

Two general approaches dominate the computational treatment of metaphor: the taxonomic and the structural. The taxonomic approach has a classical provenance that derives from Aristotle and which finds expression in the computational work of (Fass, 1988) and (Way, 1991). This perspective views metaphor as a process of abstraction in which one concept is reconciled with another by virtue of a common taxonym. The problem with this approach is symmetry: if cigarettes can be described as time-bombs by virtue of belonging to a common category like *SlowDeadlyInstrument*, then it should be just as meaningful to describe time-bombs as cigarettes; however, such a reversal produces an entirely different, and often meaningless, proposition. Ortony (1993) argues that metaphors project high-salience properties of a descriptive source term (the vehicle) onto a target term (the tenor) for which those properties are not already salient. It is not enough then to simply identify a common class to which both vehicle and tenor belong, for a computational agent must also identify the salient properties of the vehicle that are ascribed to the tenor.

The second major paradigm is a structural approach, one that also explains the workings of scientific analogies. Structure-Mapping Theory (or SMT; see Falkenhainer *et al*., 1989) views the juxtaposition of semantic structures as a process of graph alignment, in which an isomorphic mapping is generated between the most systematic elements of both structures. Since this corresponds to the NP-hard problem of determining the largest common sub-graph among two representations, a sub-optimal heuristic approach is generally employed. The new information asymmetrically imparted by a metaphor is explained by SMT via a process of candidate inferencing: if a mapping based on structural systematicity reveals structural gaps in the topic structure that can filled by corresponding parts of the vehicle structure, then these parts are subsequently projected into the topic.

Other computational variations of metaphor-as-matching also exist. Martin (1990) describes a system (MIDAS) for comprehending novel instances of conventional metaphors (e.g., ”how do I kill emacs?”) by matching these to known precedents. Veale and Keane (1997) describe an approach to metaphoric structure-mapping that motivates the cross-domain alignment of entities (such as cigarettes to time-bombs, paper to fuses, nicotine to TNT, etc.) by first identifying any shared semantic properties of these entities. As a counter-balance to SMT, Barnden and Lee (2002) caution against the over-use of direct mapping in metaphor, emphasising instead the importance of reasoning within the literal domain of the source before mapping the results of this reasoning into the target.

The Sardonicus approach described here belongs fully to neither the taxonomic nor the structure-mapping schools, though it is compatible with each. As a case-based approach, Sardonicus is perhaps closest to the MIDAS system of Martin (1990), though our case-base comprises a large number of common similes rather than small number of highly productive conventional metaphors. Because these similes identify the most salient and diagnostic features of the categories most commonly used as figurative sources of comparison, psychologically, Sardonicus is closest in spirit to the category-inclusion model of Glucksberg (2001) and the salience-imbalance model of Ortony (1993). In fact, Glucksberg's theory is an ideal fit for a CBR approach. Case selection and adaptation are, after all, problems of category inclusion, wherein one must first determine which case best includes or covers a given situation, before subsequently determining which aspects of this selected case are most applicable to the situation.

Acquiring a Case-Base of Similes

Sardonicus employs the *Google* search engine as a retrieval mechanism for gathering explicit similes from the web, since the *Google* API and its support for the wildcard term * allows the simile gathering process to be fully automated. Since each case/simile should be applicable to a wide range of possible target concepts, we gather only partial similes that conform to the pattern ”*as ADJ as a|an NOUN*”. We do not expect to identify and retrieve all similes mentioned on the world-wide-web, but rather to gather a large, representative sample of the most commonly ascribed properties and their vehicle nouns.

Since we expect that explicit similes will tend to exploit properties that occupy an exemplary point on a scale, we first extract a list of antonymous adjectives, such as ”hot” or ”cold”, from the lexical database WordNet (Fellbaum, 1998). For every adjective ADJ on this list, we send the query ”*as ADJ as* *” to Google and scan the first 200 snippets returned to extract different noun values for the wildcard *. From each set of snippets we can also ascertain the relative frequencies of different noun values for ADJ. The complete set of nouns extracted in this way is then used to drive a second phase of the search, in which the query template ”*as * as a NOUN*” is used to acquire similes that may have lain beyond the 200-snippet horizon of the original search, or that may hinge on adjectives not included on the original list. Together, both phases collect a wide-ranging series of core samples (of 200 hits each) from across the web, yielding a set of 74,704 simile instances (of 42,618 unique types) relating 3769 different adjectives to 9286 different nouns.

Case Annotation

Many of these similes are not sufficiently well-formed to serve reliably as cases. In many cases, the noun value forms part of a larger noun phrase: it may be the modifier of a compound noun (as in ”bread lover”), or the head of complex noun phrase (such as ”gang of thieves” or ”wound that refuses to heal”). In the former case, the compound is used if it corresponds to a compound term in WordNet and thus constitutes a single lexical unit; if not, or if the latter case, the simile is rejected. Other similes are simply too contextual or under-specified to function well in a null context, so if one must read the original document to make sense of the simile, it is rejected. More surprisingly, perhaps, a substantial number of the retrieved similes are ironic, in which the literal meaning of the simile is contrary to the meaning dictated by common sense. For instance, ”as hairy as a bowling ball” (found once) is an ironic way of saying ”as hairless as a bowling ball” (also found just once). Many ironies can only be recognized using world knowledge, such as ”as sober as a Kennedy” and ”as tanned as an Irishman”. In addition, some similes hinge on a humorous deconstruction of the adjective, as in ”as fruitless as a butcher-shop” (the latter has no fruits) and ”as pointless as a beach-ball” (the latter has no points).

Given the creativity involved in these constructions, one cannot imagine a reliable automatic filter to safely identify bona-fide similes. For this reason, the filtering task was performed by a human judge, who annotated 30,991 of these simile instances (for 12,259 unique adjective/noun pairings) as non-ironic and meaningful in a null context; these similes relate a set of 2635 adjectives to a set of 4061 different nouns. In addition, the judge also annotated 4685 simile instances (of 2798 types) as ironic; these similes relate a set of 936 adjectives to a set of 1417 nouns. Perhaps surprisingly, ironic pairings account for over 13% of all annotated simile instances and over 20% of all annotated simile types.

While time-consuming (though not excessively so), the annotation of similes into bona-fide and ironic categories yields a set of positive and negative exemplars from which Sardonicus can learn how to automate the annotation of future similes. For instance, a newly acquired web simile might extend a known simile using hyponyms, antonyms or synonyms as found in WordNet. Alternately, Sardonicus can infer from existing similes that certain properties frequently co-occur in the same category, e.g., dead and cold, cold and stiff, stiff and slow, etc., allowing many new similes to be automatically recognized as inferentially related to existing similes for the same vehicle.

Case-based Comprehension and Generation

At its simplest, the Sardonicus case-base is a collection of basic similes that each identify a single salient property of a single vehicle noun, but these cases can be agglomerated to produce a richly composite picture of both adjectives and nouns. For instance, the case-base contains six basic similes for the noun ”funeral”, yielding the following properties:

{*sad(1), orderly(1), unfortunate(1), dignified(1), solemn(1), serious(1)*}

Associated with each feature in parentheses is the number of instances of the corresponding simile in the collected web data (recall that 12,259 unique as-adjective-as-noun pairings are derived from 30,991 bona-fide simile instances). Likewise, the case-base contains 14 basic similes for the adjective ”solemn”, yielding the following vehicles:

{*monument(3), owl(2), judge(2), funeral(1), temple(1), dowager(1), lighthouse(1), gravestone(1), pathologist(1), crow(1), funeral_march(1), oracle(1), prayer(1), wake(1)*}

A composite noun picture (as for ”funeral”) is constructed whenever Sardonicus must understand a novel simile or metaphor in which this noun is used as a vehicle. A composite adjective picture (as for ”solemn”) is constructed whenever Sardonicus must generate an appropriate metaphor (i.e., find an appropriate vehicle) for a user-supplied tenor in which this adjective is given as salient. While many metaphors undoubtedly have an additional structural dimension, as explained by existing theories of SMT, we believe a property-centric view is currently best positioned to deliver a genuinely useful, scalable and robust metaphor capability.

Comprehension

Consider the simile ”a wedding is like a funeral”. Similes are not categorization statements (Glucksberg, 2001), so this comparison simple enjoins us to look for salient properties of ”funeral” that are applicable to ”wedding”, in other words, to adapt the composite case structure describing funerals to our understanding of weddings. Our knowledge of funerals is given by property set given above, but there is nothing in the lexical semantics of ”wedding” (as found in WordNet, say) that tells us that weddings can be solemn, unfortunate or sad. Rather, the simile requires us to exploit our tacit, experiential understanding of such events, such as the fact that weddings occur in a (solemn) church, and are sometimes forced (unfortunately) for non-romantic (sad) reasons. A computer system like Sardonicus does not have an experiential grasp of events like weddings, but fortunately, it can tap into the written reflections of those that do by using the web as a corpus that can be queried on demand. Six queries of the form ”sad wedding”, ”orderly wedding”, etc. lead Sardonicus to produce the following adaptation:

{*serious(3746), sad(887), solemn(247), unfortunate(179), dignified(129), orderly(15)*}

The property list is ordered by the web frequency of each query, as given in parentheses. Now consider the reverse simile, ”funerals are like weddings”. The Sardonicus case-base contains the following bona-fide properties for the noun ”wedding”:

{*joyous(6), joyful(3), decisive(1), glorious(1), expensive(1), emotional(1)*}

Adaptation of this knowledge to the tenor ”funeral” via the web queries ”joyous funeral”, ”joyful funeral”, etc. leads to the following ordered property set:

{*emotional(2126), expensive(816), glorious(186), joyous(127), joyful(124), decisive(5)*}

The observation that funerals and weddings are typically expensive events is not a fact one expects to find in a typical knowledge representation. However, because Sardonicus gains its understanding of figurative categories from web similes, and again verifies the possible implications of a figurative interpretation via queries to the web, it easily side-steps the knowledge-bottleneck that so often plagues computational systems.

Now consider the metaphor ”weddings are funerals”. We can expect the property-based interpretation to be same as the equivalent simile, with one additional constraint: this metaphor involves a categorization, forcing us to confabulate a reason why a wedding can be considered as a kind of a funeral. There may be many such reasons (mourning the loss of bachelorhood?), some of them based on structural analogy, but Sardonicus can short-cut this reasoning process using the web, to determine with an appropriate query whether weddings can be included in the category Funeral (Glucksberg, 2001). In fact, the query ”funeral-like wedding” returns one exact match via Google, indicating that this is a valid category-inclusion. In contrast, the query ”wedding-like funeral” finds no exact matches via Google, suggesting this is either a poor metaphor or a wholly original one.

Generation

The set of metaphors for a given tenor is potentially infinite, so for any arbitrary term, any of the noun vehicles in the case-base may well be applicable. This range of possibilities not only explodes the search-space of the generation process, it can lead the user to be swamped with unwanted metaphors. For this reason, Sardonicus employs a process of goal-driven, targeted generation wherein the user specifies not just a tenor term, but a property of the tenor that should be highlighted by any generated metaphor.

To take a topical example, one not found in WordNet but nonetheless plentiful on the web, consider the tenor ”Paris Hilton” (the model, not the hotel, though the distinction is lost on Sardonicus), for which we want to generate apt metaphors that highlight the property ”skinny”. The Sardonicus case-base states that ”skinny” is salient for 35 different nouns, from ”pool cue” to ”cadaver”, but following Glucksberg's (2001) theory of category-inclusion, only those noun categories that can meaningfully include Ms. Hilton as a member should be considered as potential vehicles. Once again, the cognitive criteria governing category membership is beyond Sardonicus, but past acts of inclusion may well be available on the web. So for each of the 35 potential vehicles, a query of the form ”V-like T” (e.g., ”rake-like Paris Hilton”) is evaluated using the Google API. The following 5 metaphors are validated in this way, as ordered by web frequency:

{*post(46), pole(42), stick(38), miser(34), stick_insect(26)*}

Any one of these can be further analysed using the comprehension technique described earlier. For instance, ”Paris Hilton is a pole” yields the following interpretation set:

{*straight(387), skinny(369), thin(353), slim(204), stiff(20), scrawny(8)*}

Because Sardonicus has no prior knowledge of ”Paris Hilton”, it lacks the ability to abstract over this tenor and so generalize its web-search. The upside of this limitation is that the resulting interpretation is perfectly adapted to its target, but the obvious downside is that Sardonicus can only interpret metaphors that have already been expressed elsewhere on the web. Nonetheless, if presented with a tenor that can be found in WordNet, such as ”heiress”, Sardonicus can employ an abstraction of this tenor (such as person) when verifying the possibility of category inclusion via the web. Backing off to ”person” as an abstraction yields the following set of metaphoric category inclusions:

{*snake(58), noodle(57), post(46), miser(50), twig(49), stick(47), pencil(41), rope(34), supermodel(34), mosquito(26), scarecrow(23), thread(17), cadaver(17)*}

So it matters not then whether an heiress has ever been described as a ”pencil” on the web, since the query used to validate this possibility is the broader ”pencil-like person”.

Dynamic, Context-Situated Learning

When presented with a metaphor or simile containing an unknown vehicle noun, a small-scale version of the case acquisition process of section 3 can be launched on the fly. Consider the statement ”Atlantis is a myth”; as with the ”Paris Hilton” example, it matters not whether ”Atlantis” is known to Sardonicus, since case adaptation is performed using topical web knowledge. However, the noun ”myth” must have some properties ascribed to it in the case-base is an interpretation is to be generated. But even if the case-base contains no properties for ”myth”, these can be acquired dynamically using the Google query ”as * as a myth”. Yet this is not a full solution, since all acquired properties must still be hand-filtered for irony.

An alternate approach is to hypothesize what properties are actually possessed by the noun, before then verifying that these properties are salient by coining and then web-validating the corresponding similes. For example, if we hypothesize that timelessness is a salient property of myths, we can verify this hypothesis by generating the simile ”as timeless as a myth” and looking for this simile on the web. Such hypotheses can be derived from a simple collocation analysis of a large text corpus (such as the complete text of Wikipedia). The noun ”myth”, for instance, is used with 123 different adjectival modifiers across Wikipedia, such as ”urban” (61 times), ”popular” (53 times) and ”persistent” (10 times). When the corresponding 123 similes (such as ”as persistent as a myth”) are evaluated using the Google API, the following property set (with web frequencies) is validated:

{*religious(3), famous(3), strong(3), heroic(2), improbable(2), timeless(2), historical(1), inaccurate(1)*}

Now, when adapted to the tenor ”Atlantis”, the following interpretation is generated:

{*famous(1283), strong(178), historical(93), religious(10), inaccurate(6), timeless(5), heroic(5), improbable(3)*}

In contrast, the metaphor ”Hercules is a myth” yields an adaptation where, as one might expect, the properties ”strong” and ”heroic” are more prominent than they are for ”Atlantis”.

{*strong(295), famous(231), heroic(140), historical(69), religious(5), inaccurate(3), improbable(3), timeless(1)*}

An important advantage of using a corpus to hypothesize about the properties of noun vehicles is that every text corpus implicitly represents a different perspective on the world, and thus a different context for reasoning about the world. A user can thus influence the cultural mind-set of Sardonicus by choosing a learning corpus that best reflects that user's personal interests (e.g., pop-culture or sci-fi, via Wikipedia) and tacit beliefs.

In fact, because Wikipedia is a reasonably authoritative encyclopedia, it serves as an ideal reservoir of shallow knowledge, one that is constantly updated, into which Sardonicus can dip whenever its own structured knowledge proves deficient for a particular concept. To quantify its capacity to provide hypotheses about under-represented vehicle concepts, we considered a set of 568,400 unique adjective:noun collocations [1] that can be extracted from Wikipedia's textual content. For each collocation, such as "timeless myth", we automatically generated the corresponding simile, such as "as timeless as a myth", and searched for this simile on the web (again using the Google API). Promisingly, 28,400 of these manufactured similes, or 5%, are found to have attested uses on the web. Overall, the text of Wikipedia yields an average of 3.7 web-attested salient properties for 7638 different nouns.

Empirical Evaluation

As noted in section 3, the Sardonicus case-base currently contains 12,259 bona-fide similes describing 4061 noun vehicles in terms of 2635 different adjectival properties. The case acquisition process, abetted by Google's practice of ranking pages according to popularity, should contain the most frequently-used simile vehicles, and thus, we expect, the most useful source categories for figurative comparison. But the descriptive sufficiency of the case-base is not guaranteed until we can demonstrate that the properties ascribed to each vehicle term are collectively rich enough, and individually salient enough, to accurately predict how each vehicle category is actually understood by a typical language user.

If similes are indeed a good basis for constructing a case-base of figuratively useful categories and their most salient properties, we should expect the set of properties for each category to accurately predict how the category is perceived as a whole. For instance, humans – unlike computers – do not generally adopt a dispassionate view of ideas, but rather tend to associate certain positive or negative feelings, or affective values, with particular ideas. Unsavoury activities, people and substances generally possess a negative affect, while pleasant activities and people possess a positive affect. Whissell (1989) reduces the notion of affect to a single numeric dimension, to produce a *dictionary of affect* that associates a numeric value in the range 1.0 (most unpleasant) to 3.0 (most pleasant) with over 8000 words in a range of syntactic categories (including adjectives, verbs and nouns). So to the extent that the adjectival properties yielded by processing similes paint an accurate picture of each noun vehicle, we should be able to predict the affective rating of each vehicle via a weighted average of the affective ratings of the adjectival properties ascribed to these vehicles (i.e., where the affect of each adjective contributes to the estimated affect of a noun in proportion to its frequency of co-occurrence with that noun in our simile data). More specifically, we should expect that ratings estimated via these simile-derived properties should correlate well with the independent ratings contained in Whissell's dictionary.

To determine whether similes do offer the clearest perspective on a category's most salient properties, we calculate and compare this correlation using the following data sets:

A. Adjectives derived from annotated bona-fide (non-ironic) similes only.
B. Adjectives derived from all annotated similes (both ironic and non-ironic).
C. Adjectives derived from ironic similes only.
D. All adjectives used to modify a given noun in a large corpus. We use over 2-gigabytes of text from the online encyclopaedia Wikipedia as our corpus.
E. All adjectives used to describe a noun in any of the WordNet glosses for that noun (giving 63,935 unique adj-noun pairs in all). For instance, WordNet defines Espresso as "strong black coffee ..." so Espresso is seen as both *strong* and *black*.

Predictions of affective rating were made from each of these data sources and then correlated with the ratings reported in Whissell's dictionary of affect using a two-tailed Pearson test ($p < 0.01$). As expected, property sets derived from bona-fide similes only (A) yielded the best correlation (+0.514) while properties derived from ironic similes only (C) yielded the worst (-0.243); a middling correlation coefficient of 0.347 was found for all similes together, demonstrating the fact that bona-fide similes outnumber ironic similes by a ratio of 4 to 1. A weaker correlation of 0.15 was found using the corpus-derived adjectival modifiers for each noun (D); while this data provides far richer property sets for each noun vehicle (e.g., far richer than those offered by the simile database), these properties merely reflect potential rather than intrinsic properties of each noun and so do not reveal what is most salient about a vehicle category. More surprisingly, perhaps, property sets derived from WordNet glosses (E) are also poorly predictive, yielding a correlation with Whissell's affect ratings of just 0.278.

Concluding Remarks

The web is an essential resource for Sardonicus, serving as a vast and growing case-base of conceptual associations when a first-principles algorithmic account of category inclusion is lacking. Additionally, the web provides topical world knowledge of a character one cannot find in a resource like WordNet. Indeed, the category properties one *can* find in WordNet may not be the most diagnostic or salient for figurative purposes. To confirm this hypothesis, we estimated the likelihood that a property deemed salient by a WordNet gloss would also make a good simile. We thus generated

[1] These collocations use the same 3769 adjectives as the web acquisition process described earlier, and are extracted from a June 2005 text-dump of Wikipedia.

similes for each of the 63,935 unique adjective-noun associations extracted from WordNet glosses, e.g., ”as strong as espresso”, ”as Swiss as Emmenthal” and ”as lively as a Tarantella”, and counted how many of these manufactured similes can be found on the web, again using Googles API. Since we find that only 3.6% of these artificial similes have attested uses on the web, we conclude that: a) few nouns are considered sufficiently exemplary of some property to serve as a meaningful vehicle in a figure of speech; b) the properties used to define categories in hand-crafted knowledge resources like WordNet are not always the properties that best reflect how humans think of, and use, these categories.

References

Barnden, J. A. and Lee, M. G. 2002. An Artificial Intelligence Approach to Metaphor Understanding. *Theoria et Historia Scientiarum*, 6(1):399–412.

Falkenhainer, B. Forbus, K. and Gentner, D. 1989. Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence* 41, 1–63.

Fass, D. 1991. Met*: a method for discriminating metonymy and metaphor by computer. *Computational Linguistics*, 17(1):49–90.

Fellbaum, C. (ed.). 1998. *WordNet: An Electronic Lexical Database*. MA, Cambridge: The MIT Press

Glucksberg, S. 2001. *Understanding figurative language: From metaphors to idioms*. Oxford: Oxford University Press.

Hanks, P. 2004. The syntagmatics of metaphor. *The International Journal of Lexicography*, 17(3).

Lakoff, G. and Johnson, M. 1980. Metaphors we live by. Chicago University Press.

Martin, J. H. 1990. *A Computational Model of Metaphor Interpretation*. New York: Academic Press.

Ortony, A. 1993. The role of similarity in similes and metaphors. *Metaphor and Thought* (2^{nd} ed.), A. Ortony (Ed.). Cambridge University Press.

Veale, T. and Keane, M. 1997. The Competence of Sub-Optimal Structure Mapping on Hard Analogies. In *Proceedings of the 15th International Joint Conference on AI*. San Mateo, CA: Morgan Kauffman.

Whissell, C. 1989. The dictionary of affect in language. In R. Plutchnik & H. Kellerman (Eds.) *Emotion: Theory and research*, 113–131. New York: Harcourt Brace.