Slip-Sliding Along in Linguistic Creativity: Building A Fluid Space for Connecting Disparate Ideas

Tony Veale[1], Yanfen Hao[1],

[1] School of Computer Science and Informatics, University College Dublin, Dublin D4, Ireland.
{Tony.Veale, Yanfen.Hao}@UCD.ie

Abstract. The art of linguistic creativity lies not in finding new truths to express in language (if there are any), but in finding new and more resonant ways to express truths that are already known or implicitly accepted. Creative expression thus requires that we take adopt a new and revealing perspective on a familiar idea, one that may prompt an audience to conceptually re-represent this idea to draw out new meaning or new relevance to a situation. As such, a computational model of linguistic creativity requires a knowledge representation that is as semantically agile and accommodating as the creative speakers who use it. We thus present a flexible and very concise knowledge representation for dealing with creative metaphors, called *Talking Points*, and show how talking points can be acquired on a large scale both from WordNet and from the web to form a fluid connected structure called a *slipnet*.

Keywords: linguistic creativity, metaphor, humour, re-expression, slippage.

1 Introduction

Linguistic creativity manifests itself in two key ways in language. In one guise, it makes the unfamiliar and the strange seem more familiar and understandable [1]. 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. In its other guise, the one most often associated with the poetic and fanciful use of language, it makes the familiar and mundane seem strange and unfamiliar, allowing us to view a commonplace idea from a new and revealing category perspective. For instance, one might describe make-up 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 [3] and for determining how they are connected [4], even if those connections are tenuous, hidden or not always obvious [5]. For instance, Burqa and Make-up are connected by a conceptual path that nudges the meaning "something that must be worn by Muslim women" into "something that must be worn by women" into "something that is conventionally worn by women" into "something that society expects women to wear".

Creative linguistic devices like metaphor allow a speaker to draw out and highlight, in a modified or exaggerated form, surprising connections between different concepts. A flexible knowledge representation is thus needed if a computational system is to identify, follow and reconcile these connections. In fact, the kind of fluid representation that is required has already been described by Hofstadter in [6], who emphasizes the role of slippage in semantic representation when dealing with creative phenomena such as formal analogies, linguistic parlor games and other playful uses of language. Such a fluid knowledge representation will define the search space in which creative language processes like metaphor generation and metaphor understanding can be cognitively and computationally situated [7]: for generation, fluid connectivity will allow a system to search outwards from a given target to find those source concepts that offer the newest yet most appropriate perspectives; for understanding, fluid connectivity will allow a system to reconcile those aspects of a source concept that are most relevant to the given target concept.

In this paper we describe the construction of a fluid knowledge representation for creative metaphor processing, one that is acquired automatically from WordNet [1] 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 [6] 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

Computational work in linguistic creativity has focused on the vexing problem of metaphor interpretation for two reasons: poetic metaphor is a creative phenomenon *par excellence* [3], while conventional metaphor is pervasive is everyday language and an important challenge for NLP [8,9,10,11]. Most 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 [8] advocates that the *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 [9] shows how these violations can be repaired to thus capture the literal intent behind each metaphor. The *Midas* system of [10] explicitly 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 apt schemas. The role of inference in metaphor understanding is instead emphasized by [11], who describe a system called *ATTMeta* that contains sufficient knowledge about e.g., conventional metaphors of mind to support complex inferences about the mental states implied by metaphors.

Hofstadter [6] considers a more formal and tightly-controlled kind of language creativity in the guise of abstract analogies. This toy-like format, which supports tasks as diverse as the mapping of letter sequences or the mirroring of actions in a highly stylized tabletop environment, allows these authors to focus on the slippage processes that are required to understand analogies whose interpretation is shaped by a wide range of pragmatic pressures. These pressures are modeled using a *slipnet*, a probabilistic network in which concepts are linked to others into which they can easily be transformed 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, the choice of word forms in a sentence is more susceptible to slippage (as influenced by e.g., synonym availability in [2]) than the choice of word meanings for that sentence.

Slippage can be seen as a potentially 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 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 insights are *Governor ≈ President* and *California ≈ 12% of the United States*. WordNet can be a rich source of insights like the former (since both presidents and governors are characterized as leaders in WordNet), but the latter is an entirely ad-hoc equivalence that one is unlikely to find in any general-purpose resource like WordNet. While ultimately aiming for this kind of creative transformation, our goal here is more modest: to build a network of concepts that are connected by incremental degrees of slippage along pathways of related facts and beliefs. We show how this network can combine the principled flexibility of a Hofstadter-style slipnet with the comprehensive scale of a resource like WordNet.

3 A Knowledge-base of *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 Acquiring Objective Talking Points from WordNet

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. 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 Harvesting Stereotypical Talking Points from the Web

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. It has been argued in [12] 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 shown to be somewhat leaky in English, and prone to subversion by irony, the authors of [12] 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 [6], but on the scale of [2], 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 $is\_ADJ_1{:}OBJ_1$ and $is\_ADJ_2{:}OBJ_2$ should be connected in the slipnet if: OBJ_1 and OBJ_2 are semantically close (i.e., synonymous, or semantic siblings in WordNet); and ADJ_1 and ADJ_2 are synonymous, or ADJ_1 frequently implies ADJ_2 or ADJ_2 frequently implies ADJ_1. 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* (a pragmatic shift) to *is_political:campaign* (a semantic shift) to *is_military:campaign* (another pragmatic shift), thereby connecting Taliban (*is_religious:movement*) to Crusade (*is_military:campaign*).

4.1 Creative Slippage 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 the path: Pope→*has_sacred:words*→*has_wise:words*→*has_solemn:words*←Judge. The concept-sensitive interplay predicted by Black's "interaction view" of metaphor [3] is clearly on display here, since the interpretation of a particular source concept depends crucially on how it is able to interact with a specific target concept. 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 interacts with the concept Warrior, e.g., Pope→*has_infallible:voice*→*has_powerful:voice*←Warrior.

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

Pope ⇒	**Pope** ⇒
≡ *leads:Roman_Catholic_Church*	≡ *leads:Roman_Catholic_Church*
≈ *leads:congregation*	≈ *leads:congregation*
≈ *leads:flock*	≈ *leads:political_movement*
≈ *leads:mob*	≈ *leads:gang*
≈ *leads:organized_crime*	≈ *leads:military_force*
Don (Crime Father) ⇐	**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 creative metaphors, 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 taxonomically organizes {Roman_Catholic_Church} as a specialization of {Organized_religion}, the metaphor creatively 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 [13,14] which are designed to measure the effectiveness of web-acquired conceptual descriptions. These authors use WordNet as a semantic gold-standard, so it would be circular to replicate their experiments on talking points that are extracted from WordNet. We consider here just the effectiveness of stereotypical talking points.

Almuhareb and Poesio describe two different clustering experiments. In the first [13], 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, [13] report a cluster accuracy of 71.96% using just values like *hot* (all 51,045), an accuracy of 64.02% using just attributes like *color* (all 8934), and 85.5% for both combined (59979 total).

In a second, larger experiment, Almuhareb and Poesio [14] 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 apply the *repeated bisections clustering* algorithm to this data set, and report an initial cluster purity measure of 56.7% using only property values like *hot*, 65.7% using only attributes like *color*, and 67.7% for both combined. 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%. This leads to an improved cluster purity measure of 62.7% using property values only and 70.9% using attributes only. However, this filtering reduces the clustering performance of both attributes and values combined, to 66.4%.

We replicate here both of these experiments using the same data-sets of 214 and 402 nouns. 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 attributes 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*". The combination of attributes and values yields a clustering accuracy of 90.2% for experiment 1 (compare with the 85.5% reported in [13]) and an accuracy of 69.85% for experiment 2 (compare with the 66.4% reported in [14]). Though just slightly superior, these clustering results are achieved with considerably smaller representations (at least seven times smaller) than that used in [13,14]. We therefore conclude that talking points capture not just meaningful aspects of a lexical concept, but the most salient aspects of a concept.

6 Conclusions

Creative linguistic devices like metaphor are knowledge-hungry in the extreme, since they exploit both a factual knowledge of the world and a knowledge of how these facts, or talking points, can be nudged into the realm of the colorful, the fanciful and the resonant. Any computational treatment of metaphor will thus only be as good as the knowledge representation that supports it. The representation described here – called *talking points* – is simple yet scalable, and lends computational substance to some key insights in the metaphor literature, from the interaction theory of Black [3] to the conceptual blending theory of [4] as computationally modeled by [7]. We also employ a key insight from the work of Hofstadter and his fluid analogies group [6], that creative reasoning on a conceptual level requires a degree of meaning 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 [2] and the stereotypical comparisons that pepper the texts of the web [12]. 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, slippage 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 [6]. Our replication of the experiments of [13,14] also indicates 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 and scalable enough to generate simple but imaginative metaphors on demand for a wide range of user inputs.

References

1. Indurkhya, B.: Metaphor and Cognition: Studies in Cognitive Systems. Kluwer. (1992)
2. Fellbaum, C. (ed.): WordNet: An electronic lexical database. MIT Press. (1998)
3. Black, M.: Models and Metaphor: studies in language and philosophy. Cornell Uni. (1962)
4. Fauconnier, G., Turner, M.: Conceptual Integration Networks. Cognitive Science, 22(2), 133–187. (1998)
5. Collins, A., Loftus, E. F.: A Spreading-Activation Theory of Semantic Processing. Psychological Review 82, 407-428. (1975)
6. Hofstadter, D.R., Fluid Analogy Research Group: Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought. Basic Books (1994).
7. Veale, T., O'Donoghue, D.: Computation and Blending. Cog. Linguistics, 11(3-4). (2000)
8. Wilks, Y.: Making Preferences More Active. Artificial Intelligence 11(3), 197-223. (1978)
9. Fass, D.: Met*: a method for discriminating metonymy and metaphor by computer. Computational Linguistics, 17(1), 49-90. (1991).
10. Martin, J.: A Computational Model of Metaphor Interpretation. Academic Press. (1990)
11. Barnden, J. A., Lee, M. G.: An Artificial Intelligence Approach to Metaphor Understanding. Theoria et Historia Scientiarum, 6(1), 399-412. (2002)
12. Veale, T., Hao, Y.: Comprehending and Generating Apt Metaphors: A web-driven, case-based approach to figurative language. Proc. of the 22nd AAAI conf. on AI, 1471-76. (2007)
13. Almuhareb, A., Poesio, M Attribute-Based and Value-Based Clustering: An Evaluation. Proc. of EMNLP, Emerging Methods in NLP, 158--165. (2004)
14. Almuhareb, A., Poesio, M.: Concept Learning and Categorization from the Web. Proc. of the annual meeting of the Cognitive Science society. (2005)