

# Fighting Words and Antagonistic Worlds

Tony Veale

School of Computer Science and Informatics  
University College Dublin  
Belfield, Dublin D4, Ireland.

Tony.Veale@UCD.ie

## Abstract

Metaphor is a fundamentally antagonistic way of viewing and describing the world. Metaphors ask us to see what is *not* there, so as to remake the world to our own liking and to suit our own lexicons. But if metaphors clash with the world as it is, they can also clash with each other. Each metaphor represents a stance from which to view a topic, and though some stances are mutually compatible, many more are naturally opposed to each other. So while we cringe at a clumsily mixed metaphor, there is real value to be had from a deliberate opposition of conceptual metaphors. Such contrasts reveal the limits of a particular worldview, and allow us to extract humorous insight from each opposition. We present here an automatic approach to the framing of antagonistic metaphors, embodied in a metaphor-generating Twitterbot named *@MetaphorMagnet*.

## 1 Two-Fisted Metaphors

The imagination often takes flight on the wings of metaphor. For metaphor allows us to make the fantastical seem real and the banal seem fresh and newly interesting. For example, consider this imaginary scenario, as packaged in a pithy tweet:

*What if [#TheXMen](#) were real? [#NoamChomsky](#) could be its [#ProfessorCharlesXavier](#): smart yet condescending, and scowling too*

This counterfactual injects some much-needed pizzazz into the banalities of modern politics and intellectual posturing, by reimagining a famously dour academic activist as the real-world equivalent

of a much-loved comic-book character. This counterfactual is, at its heart, a metaphor: we can construct a bridge from *Chomsky* to *Xavier* only because we believe them to share deep similarities. If the metaphor implies much more than this set of properties actually conveys, this is because it also sparks the imagination of its audience. We are led to imagine Chomsky as the cerebral hero of a battle between good and evil, in which he leads his own academic version of the X-Men, loyal students with a zealous sense of mission.

Now consider this follow-up tweet, which is designed to further stoke a reader's imagination:

*If [#NoamChomsky](#) is just like [#ProfessorCharlesXavier](#), smart yet condescending, then who in [#TheXMen](#) is [#GeorgeLakoff](#) most like?*

Metaphors are systematic, and lead us to project coherent systems of relational structure from one domain to another (see Lakoff & Johnson, 1980; Gentner *et al.*, 1989). In this way we invent hybrid worlds that combine elements of reality *and* fantasy, in which each mapping, such as Chomsky to Xavier, can prompt others, such as Lakoff to his mutant counterpart (Magneto, perhaps?).

The real world is not a comic book, and there is something mischievously silly about describing a serious scholar and activist as a fictional creation with super-powers. Yet metaphors work well as jokes when they make a virtue of the differences that separate ideas. As Pollio (1996) put it, “*split reference yields humour if the joined items (or the act joining them) emphasize the boundary or line separating them; split reference yields metaphor if the boundary between the joined items (or the act joining them) is obliterated and the two items fuse to form a single entity*.” So by dialing up the antagonism – between domains, between reality and

fantasy, or between people and ideas – a metaphor can yield a witty, eye-catching and thought-provoking text that is worth sharing on a platform such as Twitter. This point is worth stressing, as the above tweets were generated by an automated *Twitterbot*, named [@MetaphorMagnet](#), whose antagonism-stoking generative processes are the subject of this paper.

If metaphor can give you wings, it can also give you fists with which to pummel a contrary point of view. Every conceptual metaphor offers a potted world-view that encourages us to reason in certain ways, and thus speak in related ways, about our experiences. But like proverbs, or indeed ideologies, we can often pick and choose the ones that suit us best. Reasonable people can disagree about how best to categorize a situation, as no metaphor is ever objectively right or true, just potentially apt in a particular context. Thus, thinkers on different ends of the political spectrum offer antagonistic metaphors to frame the same goals, needs or problems, and by advancing their own conceptual frames they actively seek to undermine those of their opponents. Just as every proverb has a converse that is equally compelling (e.g., *many hands make light work* vs. *too many cooks spoil the broth*), there is conceptual sport to be had in finding the most apt *anti*-metaphor for a given figurative viewpoint. The following tweet thus frames two antagonistic views of *love*:

*To some beatniks, love is a sparkling rainbow. To others, it is a flat bed.*  
[#Love=#Rainbow](#) [#Love=#Bed](#)

This tweet nicely captures the antagonism that exists between competing perspectives on [#Love](#). The first is expansive, and views love as a many-splendored thing; the second is more reductive, and views love as just a means to an end: *sex*. By attributing these views to different members of the same category of person – *beatniks* – the tweet suggests that this conflict of ideas is also a conflict between otherwise similar people.

This paper explores the automated generation of antagonistic metaphors. By elevating simple contrasts into a contest of ideas, [@Metaphor Magnet](#) creates metaphors that also work as witty provocations to think differently, or at least to appreciate the limits of received wisdom. This automated system seeks its inspiration in attested usage

data and uses a variety of knowledge-rich services to produce elaborate, well-reasoned metaphors that hinge upon meaningful contrasts. In the sections that follow, we describe how this harmonious marriage of explicit knowledge and raw usage data is used to sow disharmony at the level of ideas and package the results as tweets.

## 2 Competing Points of View

A divergent problem is one that admits many potential solutions, each of them valid in its own way (Guilford, 1967). Though one may be privileged over others by its conventionality – e.g. the use of a brick as a building block, or of a paper clip to bind papers – there is no single, objectively *correct* answer. Conversely, a convergent problem is one that admits just one objectively-acceptable *correct* answer, relative to which all others are seen as deficient or just plain *wrong*. By this standard, metaphor is a divergent approach to the conveyance of meaning, while literal language – to the extent that any text can be truly literal – is considerably more convergent.

A cornerstone of divergent thinking is divergent categorization: this allows us to categorize a familiar object or idea in atypical ways that permit new and unusual uses for it (Torrance, 1980). Such categorization is, in turn, central to the act of figurative description. Consider the metaphor *divorce is war*, whose interpretation requires us to find a non-trivial category – one a good deal more specific than *event* – to embrace these very different-seeming concepts (Glucksberg, 1998). To see how people categorize, we need only see how they speak. On the Web, we see descriptions of both war and of divorce, in separate texts, as *traumatic events*, *serious conflicts*, *immoral acts*, and as *bad things* in general. Such descriptions often come in standardized linguistic containers, such as the “*A\_Bs such as Cs*” pattern of Hearst (1992), instances of which are easily harvested from the Web. The [Thesaurus Rex](#) Web service of Veale & Li (2013) offers up its resulting system of Web-harvested categorizations as a public service that can be exploited by 3<sup>rd</sup>-party metaphor systems. *Thesaurus Rex* can be used for the interpretation of metaphors by permitting another system to explore specific unifying categories for distant ideas, such as [divorce & war](#), but it can also be used in the generation of metaphors. So if looking for a meta-

phor for *creativity*, *Thesaurus Rex* suggests the category *special ability*, leading a metaphor generator to consider other members of this category as possible vehicles, such as *x-ray vision*, *superior strength*, *magic* or *prophecy*. @MetaphorMagnet thus uses *Thesaurus Rex* to package diverse ideas into a single tweet, as in:

#Take5 of the #Shallowest things:

1. Toilet Bowls
2. Rock Stars
3. Cookie Sheets
4. Soup Bowls
5. Rush Limbaugh

#TheRepublicans

Divergent thinking typically arises when we go *off-script* to imagine unconventional possibilities for a familiar object or idea. Raskin (1985) puts the concept of a script at the centre of his computational theory of jokes, the *Semantic Script Theory of Humour* (SSTH), arguing that most joke narratives are compatible with two competing scripts at once. The primary script, which listeners are lulled into applying based on a normative reading of a narrative, is activated as the result of convergent thinking; the secondary script, which the joker downplays at first and which listeners only perceive when a big “*reveal*” is delivered at the end, is a result of divergent thinking and an ability to find novel uses for familiar situations. Metaphors rely on categories the way jokes rely on scripts. Thus, while the category *immoral act* will embrace acts that are clearly immoral, such as *murder*, *torture*, *bribery* and *fraud*, in the right circumstances it can also be used to embrace the outlier ideas *divorce*, *drug use* and even *dancing*.

Nonetheless, the closest equivalent to a script in metaphor is the *Conceptual Metaphor* (CM). Conceptual Metaphors, as described in Lakoff & Johnson (1980), are the cognitive deep structures that underpin whole families of related linguistic metaphors. The *Life is a Journey* CM, for example, is the fountainhead of figures of speech such as “*go off the rails*”, “*hit the skids*”, “*crash and burn*”, “*smooth sailing*” and “*on the rocks*.” So just as trips to many kinds of restaurant can all be understood using a generic *Restaurant* script (i.e. *enter-sit-order-eat-pay-leave*), a CM such as *Life is a Journey* facilitates a generic level of reasoning about life’s events. And just as a script has slots for various roles, props and locations, a CM has its

own schematic structure with slots to fill, such as *Source*, *Path*, *Goal* and *Vehicle*. A CM such as *Life is a Journey* thus allows us to impose the schematic structure of a *Journey* onto our mental structure of a *Life*, to understand *Life* as something with a starting point, a destination, a path to follow and a means of conveyance.

Carbonell (1981), Martin (1990) and Barnden (2008) each build and exploit an explicit representation of conceptual metaphors, while Mason (2004) uses statistical methods to extract conventional metaphors – CMs that are so entrenched in the way we speak that their uses in language can often seem literal – from text corpora. Shutova (2010) uses statistical clustering to identify possible target ideas – such as *Democracy* and *Marriage* – for a given source idea such as *Mechanism*. This allows her system to recognize “*fix a marriage*” and “*the functioning of democracy*” (or vice versa) as figurative uses of a *Mechanism* schema because they each use verbs that typically take mechanisms as their objects. But whether one views CMs as real cognitive structures or as useful statistical generalizations, CMs serve as script-like bundles of norms and roles that shape the generation and interpretation of metaphors.

In any case, CMs are so often paraphrased in the metaphor literature using copula statements of the form *X is a Y* that candidate CMs are easily harvested from a source of Web n-grams, not just because the metaphor literature is itself part of the Web, but because lay speakers have over-used many of these forms to the point of cliché. So the Google n-grams (Brants & Franz, 2006) is not just a source of CM paraphrases such as “*Life is a Journey*” (freq=12,688) but of colorful variations on these themes as well, such as “*Life is a Highway*” (freq=2,443), “*Life is a Rollercoaster*” (freq=3,803), “*Life is a Train*” (freq=188), “*Life is a Maze*” (freq = 180), “*Life is a Pilgrimage*” (freq=178) and “*Life is a River*” (freq=119). If one doubts that metaphor is a divergent phenomenon, one need only look at the Google n-grams, which attests that people also speak “*Life is a Game*” (freq=8,763), “*Life is a Circus*” (freq=598), “*Life is a Banquet*” (freq=102), and even that “*Life is a Sitcom*” (freq=180).

These short linguistic expressions typically sit on the figurative continuum somewhere between proverbs and clichés, as such phrases must have a minimum Web frequency of 40 to ever find their

way into the Google n-grams. Like clichés, these phrases crystalize a wealth of received wisdom, but just like proverbs they offer just one potted view on a topic, one that is easily countered by an apt choice of counter-proverb or anti-metaphor, as we shall show in coming sections.

### 3 Grudge Matches

Google 4-grams are a rich source of copula metaphors such as “*Life is an Adventure*” (freq= 1,317) and “*Life is an Illusion*” (freq=95), while the 3-grams also offer up gems such as “*Life is Rubbish*” (freq=8,489), “*Life is Love*” (freq=889) and “*Life is War*” (freq=44,490). Many of these n-grams give linguistic form to established CMs, but many more occupy a questionable area between resonant metaphor and random, overheard phrase. So a computational system must exercise careful selectivity in deciding which n-grams are worthy of elaboration into a novel linguistic form and which are best discarded as unreliable noise.

A good starting point is affect, as those copula n-grams that assert the identity of polarized ideas with antagonistic sentiments, such as *faith* and *aggression*, make for provocative metaphors. So consider the 4-gram “*faith is an aggression*” (freq=44), whose frequency is high enough to suggest it is well-formed, but low enough to suggest it resides in the long-tail of public opinion. Most sentiment lexica will view *faith* as a strong positive idea and *aggression* as a strong negative, so these ideas make for a bold juxtaposition, as packaged in this tweet from @MetaphorMagnet:

*Remember when faiths were practiced by kind priests? Now, faith is an aggression that only unkind aggressors exhibit.*

Notice that the original motivating 4-gram “*faith is an aggression*” sits at the centre of the tweet. @MetaphorMagnet seeks its inspiration from the Google n-grams, to find some interesting snippet of text that may, with reasoned elaboration, blossom into a fuller form that is worthy of tweeting. Viewed in this way, an n-grams database is like a crowded railway station, buzzing with fleeting morsels of overheard conversations. When one’s interest is finally piqued by a particular fragment, one has no choice but to complete it oneself.

Yet reasoned elaboration demands knowledge over which a system can reason, and the tweet

above showcases several pieces of stereotypical knowledge: that priests are often kind and practice faiths, while aggressors are often unkind and exhibit aggression. Knowledge of stereotypical properties is sourced as needed from *Thesaurus Rex* and from a database of typical associations mined on the Web by Veale & Hao (2007), while relational knowledge – linking e.g. priests to their faiths via specific actions – is sourced from yet another public Web service, *Metaphor Eyes*, as presented in Veale & Li (2011). The relational triples provided by *Metaphor Eyes*, mined from WH-questions commonly found in Web query logs (e.g. “*why do priests wear white collars?*”), can also be used to generate simple analogies, though the most provocative analogies are often antagonistic disanalogies. Consider an analogical tweet that @MetaphorMagnet tags as an #irony:

*#Irony: When some anglers use "pointed" hooks the way salespersons use pointless gimmicks. #Angler=#Salesperson #Hook=#Gimmick*

Each of @MetaphorMagnet’s tweets strives for a balance of similarity and dissimilarity. The analogical similarity here derives from a parallelism in the action of two agents – each *use something* – while the dissimilarity derives from a specific contrast between the objects so used. Though the contrast of *pointed* and *pointless* is mere wordplay, it is may be enough to spark more profound processes of meaning construction in the reader. To spur the reader into engaging these processes, the system explicitly hashtags the tweet as *ironic*, and puts the positive side of the contrast, *pointed*, in scare quotes. The reader is thus prompted to view the dissimilarity as merely superficial, and to read a deeper meaning into what is essentially a superficial similarity. The reader, if not the system, is left with the image of a bad fisherman, for whom pointed hooks are just pointless gimmicks. The use of ironic scare quotes to signal fakeness or insincerity is made more explicit in this tweet:

*#Irony: When some jewelers sell "valuable" diamonds the way tinkers sell valueless junk. #Jeweler=#Tinker #Diamond=#Junk*

So @MetaphorMagnet strives to sow antagonism even in the presence of unifying similarity, by for example, choosing to mold this similarity into the most negative comparisons. Consider another of the system’s rendering strategies in this tweet:

*Tourist. noun. A creep who would rather enjoy bizarre excursions than bizarre perversions.*

[#Tourist](#)=[#Creep](#)

Once again the similarity here hinges on a rather generic shared relationship: tourists enjoy excursions and creeps enjoy perversions. The contrast is primarily one of affect: *tourist* has mildly positive sentiment as a lexical concept, while *creep* has an especially strong negative sentiment. And though *bizarre* is a stereotypical property of the concept *perversion*, the Google 2-gram “*bizarre perversion*” (freq=111) attests that speakers often apply the property *bizarre* to excursions too.

A system may go further and use hashtags to imply a similarity that borders on identity, as in:

*Would you rather be:*

1. A guardian supervising an innocent child?
2. A jailer supervising a culpable offender?

[#Guardian](#)=[#Jailer](#)

So while antagonistic views on the world stress the conflict between two opposing situations, we can provoke deeper antagonism still by asserting these situations to be almost identical beneath the surface. Yet the screenwriter’s maxim of *show, don’t tell* applies as much to tweets as it does to films, so it helps if we can do more than just *tell* of identity and actually *show* near-identity in action. This requires some imagination, and perhaps more space than a single tweet will permit. Fortunately, bots are not limited to single tweets, and can issue two in quick succession if need be:

*When it comes to the devotees they lead, some swamis can be far from mellow and can even seem authoritarian.*

[#Swami](#)=[#Warlord](#) [#Devotee](#)=[#Rebel](#)

*Authoritarian swamis lead hardened devotees the way warlords lead rebels.*

[#Swami](#)=[#Warlord](#) [#Devotee](#)=[#Rebel](#)

So tweets, like movies, can have sequels too.

#### 4 Counter-Punches and Anti-Metaphors

Metaphors are underspecified and often highly context-dependent, and so many of the potential CMs that are harvested from the Google n-grams are not amenable to computational interpretation. Indeed, many – though suggestive – are not truly CMs in any accepted sense, and the 4-gram “*love*

*is a bed*” is more Conceptual Metonymy than Conceptual Metaphor, a conflation of bed with sex that underpins euphemisms such as “*in the sack*”, “*between the sheets*” and “*sleep together*”. A CM-like paraphrase will always mean more to humans who experience the world first-hand than to machines with basic symbolic representations. So a possible CM in isolation, such as the 4-gram “*idea is a gift*” (freq=94) or “*idea is a contradiction*” (freq=72), may present few computational opportunities to provoke deep thoughts, but opportunities for meaning construction abound if candidate CMs are placed into antagonistic juxtapositions, as in this @*MetaphorMagnet* tweet:

*To some thinkers, every idea is a comforting gift. To others, every idea is a disturbing contradiction.*

[#Idea](#)=[#Gift](#) [#Idea](#)=[#Contradiction](#)

The ubiquity of most CMs makes them bland and uninteresting as linguistic statements to anyone but a metaphor theorist, and so they can resemble platitudes more than true insights. But computational systems like @*MetaphorMagnet* can make generic CMs seem interesting again, by undermining their generality and revealing their limits. The key is antagonistic contrast, either between rival CMs or between a CM and literal language. Consider the conceptual metaphor that underpins the expression “*pack of girls*.” The word “*pack*” is literally used to denote a group of animals, yet its figurative extension to people is so ubiquitous in speech that we often overlook the hidden slur. This tweet reminds us that it is, indeed, an insult:

*To join and travel in a pack: This can turn pretty girls into ugly coyotes.* [#Girl](#)=[#Coyote](#)

The Google n-grams furnish the 3-grams “*pack of coyotes*” (freq=2120) and “*pack of girls*” (freq=745”). This is as close as the system comes to the underlying CM, but it is enough to establish a parallel that facilitates a provocative contrast. Ultimately, the only pragmatics that @*MetaphorMagnet* needs is the pragmatics of provocation.

#### 5 And In The Red Corner ...

The notion that one CM can have an antagonistic relationship to another is itself just a metaphor, for antagonism is a state of affairs that can only hold between people. So to dial up the figurative antagonism to 11 and turn it into something approaching

the real thing, we might imagine the kinds of people that espouse the views inherent to conflicting CMs, and thereby turn a contest of ideas into an intellectual rivalry between people.

On Twitter, the handles we choose can be as revealing as the texts we write and re-tweet, and so the creation of an online persona often begins with the invention of an apt new name. For instance, we might expect a *beatnik* (to recall our earlier figurative tweet from @MetaphorMagnet) with the handle [@rainbow\\_lover](#) to agree with the general thrust of the CM *Love is a Rainbow*. Conversely, what better handle for an imaginary champion of the metaphor *Love is a Rainbow* than [@rainbow\\_lover](#)? To condense a CM into a representative Twitter handle such as this, we can look to the Google 2-grams for suggestions. Consider the CM *Alcohol is a Drug*; while many may see this as literal truth, it is mined as a likely CM by @MetaphorMagnet from the Google 4-gram “*Alcohol is a Drug*” (freq=337). The system learns from the *Metaphor Eyes* service that addicts abuse drugs, and finds the Google 2-gram “*alcohol addict*” (freq=1250) to attest to the well-formedness of the name [@alcohol\\_addict](#). It now has an imaginary champion for this CM, which it elaborates into the following tweet:

*I always thought alcohol was drunk by bloated alcoholics. But [@alcohol\\_addict](#) says alcohol is a drug that only focused addicts abuse.*

The same strategy – in which a CM is condensed into an attested 2-gram that integrates aspects of the source and target ideas of the metaphor – is used *twice* in the following tweet to name rival champions for two antagonistic views on life:

*[@life\\_lover](#) says life is a relaxing pleasure  
[@abortion\\_patient](#) says it is a traumatic suffering  
[#Life=#Pleasure](#) [#Life=#Suffering](#)*

Notice that in the examples above, [@life\\_lover](#) and [@alcohol\\_addict](#) turn out to be the names of real Twitter users, while no Twitter user has yet adopted the handle [@abortion\\_patient](#). When the system invents a plausible handle for the imaginary champion of a metaphorical viewpoint, we should not be surprised if a human has already taken that name. However, as the names fit the viewpoints, we do not expect an existing Twitter user such as [@alcohol\\_addict](#) to take umbrage at

what is a reasonable inference about their views. Indeed, names such as [@alcohol\\_addict](#) already incorporate a good deal of caricature and social pretense, and it is in this spirit of make-believe that @MetaphorMagnet re-uses them as actors.

## 6 The Judges’ Decision

Mark Twain offered this advice to other (human) writers: “*Get your facts first, then you can distort them as you please.*” It is advice that is just as applicable to metaphor-generating computational systems such as @MetaphorMagnet that seek to use their uncontentious knowledge of stereotypical ideas to generate provocative comparisons. Many of @MetaphorMagnet’s facts come from its various knowledge sources, such as the Web services *Thesaurus Rex* and *Metaphor Eyes*, as well as a large body of stereotypical associations. But many more are not “facts” about the world but observations of what people say on the Web. One might wonder then if a random sampling of @MetaphorMagnet’s outputs would yield tweets that are as comprehensible and interesting as the examples we have presented in this paper.

A notable benefit of implementing any metaphor-generating system as a Twitterbot is that all of the system’s outputs – its hits *and* its misses – are available for anyone to scrutinize on Twitter. Nonetheless, it is worth quantifying the degree to which typical users find a system’s outputs to be meaningful, novel and worth sharing with others. We thus sampled 60 of @MetaphorMagnet’s past tweets and gave these to paid volunteers on [CrowdFlower.com](#) to rate along the dimensions of *comprehensibility*, *novelty* and *retweetability*. Judges were paid a small fee per judgment but were not informed of the mechanical origin of any tweet; rather, they were simply told that each was taken from Twitter for its figurative content.

We solicited 10 ratings per tweet, though this number of ratings was eventually reduced once the likely scammers – unengaged judges that offer random or unvarying answers or which fail the simple tests interspersed throughout the evaluation – were filtered from the raw results set. For each dimension, judges offered a rating for a given tweet on the following scale: 1=*very low*; 2=*medium low*; 3=*medium high*; 4=*very high*. The aggregate rating for each dimension of each tweet

is then calculated as the mean rating from all judges for that dimension of that tweet.

For the dimension of *comprehensibility*, over half (51.5%) of tweets are deemed to have *very-high* aggregate comprehensibility, while 23.7% are deemed to have *medium-high* comprehensibility. Only 11.6% of the system's tweets are judged to have *very low* comprehensibility, and just 13.2% have *medium low* comprehensibility.

For the dimension of *novelty*, almost half of @MetaphorMagnet's tweets (49.8%) are judged to exhibit *very high* aggregate novelty, while only 11.9% are judged to exhibit *very low* novelty.

For the dimension of *retweetability*, for which judges were asked to speculate about the likelihood of sharing a given tweet with one's followers on Twitter, 15.3% of tweets are deemed to have *very high* retweet value on aggregate, while 15.5% are deemed to have *very low* retweet value. Most tweets fall into the two intermediate categories: 49.9% are deemed to have *medium low* retweet value, while 27.4% are deemed to have *medium high* retweet value. Though based on speculative evaluation rather than actual retweet rates, these numbers accord with our own informal experience of the bot on Twitter, as thus far its own designers have favorited approx. 27% of the bot's ~7500 tweets to date. It should also be noted that a 15.3% retweet rate would be considered rather ambitious for most Twitter users, and is thus perhaps an overstatement in the case of @MetaphorMagnet too. We thus see this as a speculative but nonetheless encouraging result.

@MetaphorMagnet currently has approx. 250 human followers (as of March 1<sup>st</sup>, 2015), though it has not yet attracted enough followers to facilitate a robust empirical analysis of their rates of *favoriting* or *retweeting*. If and when it attracts sufficient followers to permit such an analysis, we may no longer need to look to crowdsourcing platforms to evaluate the system's outputs, and may actually obtain a finer granularity of insight into the kinds of metaphors, oppositions and rendering strategies that humans most appreciate.

## 7 Lucky Punches

@MetaphorMagnet uses a variety of knowledge sources to formulate its observations and an even wider range of linguistic forms to package them

into pithy tweets. Yet in every case it employs the same core strategy: identify a semantic contrast in a knowledge-base; employ semantic reasoning to elaborate a plausible but antagonistic scenario around this central contrast; and use attested Web n-grams to render this scenario in a provocative linguistic form. Though each stage is distinct from an abstract design view, they are all conflated in practice, so that e.g. Web n-grams are also used to *inspire* the system by suggesting the contrasts, juxtapositions and conceptual metaphors that appear most worthy of elaboration.

The use of raw n-grams that a system can only superficially understand constitutes a leap of faith that often pays off but sometimes does not. Consider how the 4-gram "*design is the heart*" (freq=151) provides half of the following tweet:

.@design\_scientist says design is a united collaboration  
.@design\_lover says it is a divided heart  
[#Design=#Collaboration](#) [#Design=#Heart](#)

While a human reader might understand *divided heart* as a poetic allusion to *divided loyalties* – which is nicely antagonistic to the notion of a *united* collaboration – @MetaphorMagnet has an altogether more literal understanding of the stereotypical heart, which it knows to be *divided* into various chambers. That the above juxtaposition works well is thus as much a matter of raw luck as deliberate effort, though as the old saying puts it, "*the harder I work the luckier I get.*" @MetaphorMagnet works hard to earn its frequent good fortune, and so any risk that raw n-grams bring to the generation process is more than compensated for by the unforeseeable resonance that they so often bring with them.

For more detail on the internal workings of @MetaphorMagnet, readers are directed to the online resource to [RobotComix.com](#).

## Acknowledgements

The author is grateful to the European Commission for funding this work via the WHIM project (The [What If Machine](#); grant number 611560). The author is also grateful for the support of Science Foundation Ireland (SFI) via the *CNGL Centre for Global Intelligent Content*.

## References

- Barnden, J. A. (2008). Metaphor and artificial intelligence: Why they matter to each other. In R.W. Gibbs, Jr. (Ed.), *The Cambridge Handbook of Metaphor and Thought*, 311-338. Cambridge, UK: Cambridge University Press
- Brants, T. and Franz, A. (2006). *Web 1T 5-gram Ver. 1*. Linguistic Data Consortium.
- Carbonell, J. G. (1981). Metaphor: An inescapable phenomenon in natural language comprehension. *Report 2404*. Pittsburgh, PA. Carnegie Mellon Computer Science Dept.
- Torrance, E. P. (1980). Growing Up Creatively Gifted: The 22-Year Longitudinal Study. *The Creative Child and Adult Quarterly*, 3, 148-158
- Gentner, D., Falkenhainer, B. and Skorstad, J. (1989). Metaphor: The Good, The Bad and the Ugly. In *Theoretical Issues in NLP*, Yorick Wilks (Ed.) Hillsdale, NJ: Lawrence Erlbaum Associates.
- Glucksberg, S. (1998). Understanding metaphors. *Current Directions in Psychological Science*, 7:39-43.
- Guilford, J.P. (1967). *The Nature of Human Intelligence*. New York: McGraw Hill.
- Hearst, M. (1992). Automatic acquisition of hyponyms from large text corpora. In *Proc. of the 14<sup>th</sup> International Conference on Computational Linguistics*, pp 539-545.
- Lakoff, G. and Johnson, M. (1980). *Metaphors We Live By*. University of Chicago Press.
- Martin, J. H. (1990). *A Computational Model of Metaphor Interpretation*. New York: Academic Press.
- Mason, Z. J. (2004). CorMet: A Computational, Corpus-Based Conventional Metaphor Extraction System, *Computational Linguistics*, 30(1):23-44.
- Pollio, H. R. (1996). Boundaries in Humor and Metaphor. J. S. Mio and A. N. Katz (Eds.), *Metaphor: Implications and Applications*. Mahwah, New Jersey: Laurence Erlbaum Associates, 231-253.
- Raskin, V. (1985). *Semantic Mechanisms of Humor*. Dordrecht.
- Shutova, E. (2010). Metaphor Identification Using Verb and Noun Clustering. In *the Proc. of the 23<sup>rd</sup> COLING, the International Conference on Computational Linguistics*.
- Veale, T. and Hao, Y. (2007). Comprehending and Generating Apt Metaphors: A Web-driven, Case-based Approach to Figurative Language. In *Proc. of the 22<sup>nd</sup> AAAI conference of the Association for the Advancement of Artificial Intelligence*.
- Veale, T. and Li, G. (2011). Creative Introspection and Knowledge Acquisition: Learning about the world thru introspective questions and exploratory metaphors. In *Proc. of the 25<sup>th</sup> AAAI Conf. of the Assoc. for Advancement of A.I.*, San Francisco.
- Veale, T. and Li, G. (2013). Creating Similarity: Lateral Thinking for Vertical Similarity Judgments. In *Proceedings of ACL 2013, the 51<sup>st</sup> Annual Meeting of the Association for Computational Linguistics*, Sofia, Bulgaria, August 2013.