

# Twitter: The Best of Bot Worlds for Automated Wit

Tony Veale, Alessandro Valitutti, Guofu Li,

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

{Tony.Veale, Alessandro.Valitutti, Guofu.Li}@UCD.ie

**Abstract.** Language affords a great many opportunities for the intelligent reuse of linguistic content. Rather than always putting our own thoughts into our own words, we often convey feelings through the words of others, by citing, quoting, mimicking, borrowing, varying or ironically echoing what others have already said. Social networking platforms such as Twitter elevate linguistic reuse into an integral norm of digital interaction. On such platforms, *who* you follow and *what* you re-tweet can say as much about you as the clothes you wear or the art you hang on your walls. But not everyone that is worth following is human, and not everything that is worth re-tweeting was first coined by a real person. More and more of the witty and thought-provoking content on Twitter is generated by bots, artificial systems that write their own material and vie for our attention just as humans do. Real people knowingly follow artificial bots for reasons that are subtle and diverse, but a significant reason is surely Twitter itself. This paper explores Twitter as a smart environment for automated wit, and describes the mechanics of a wittily inventive new Twitterbot named [@MetaphorMagnet](#).

**Keywords:** Twitter, Twitterbot, Metaphor, Reuse, Irony, Social Networks, Wit.

## 1 We Can (Re)Tweet It For You Wholesale

The limitations on text length imposed by micro-blogging services such as Twitter do nothing to dampen our ardour for creative language. Indeed, such limitations further incentivize the use of creative devices such as metaphor, analogy and irony, as forms such as these allow us to interact in ways that are witty, memorable and concise. As a principally textual medium, Twitter supports all of the same compression strategies as written language, but also adds some that are uniquely its own. Hashtags, for instance, allow their originators to crystalize an emerging topic or movement into a single term, thus allowing followers to hop onto an ever-accelerating bandwagon by appending the hashtag *du jour* – such as [#CancelColbert](#) or [#GamerGate](#) – to their tweets. Twitter also encourages its users to reuse, re-purpose and disseminate the tweets of others via a simple re-tweeting mechanism. Re-tweeting is an action that creates added value for the originator of a tweet *and* those that pass it along: for the former, it allows their texts to reach a wider audience, and of latter it makes content intermediaries who – as self-appointed social sensors – interactively filter what is worthy of greater attention. More and more, however, the texts that are so anointed by successive re-tweeting are not the texts of human writers but of artificial content-producers called *Twitterbots*.

A *Twitterbot* is an autonomous software system (a *bot*) that generates and tweets messages of its own design and composition. Ironically, many of these Twitterbots are popular precisely because their followers know them to be automated bots, and value the sense of the uncanny (what Freud called the *Unheimlich*) that they engender via their tweets, especially on those rare occasions when their tweets communicate an apparent insight that seems witty, profound or just enigmatic. Indeed, when humans interfere in the operation of a Twitterbot – to manually filter its outputs to improve its quality, or to manually (and fraudulently) write its tweets for themselves – users feel cheated and quickly *unfollow* the bot. For users value the unusual perspective offered by non-human bots and are willing to tolerate large amounts of noise if a bot can occasionally generate a re-tweetable gem, even if these bots ultimately lack creative intent and cannot themselves tell the good from the bad from the unintelligible.

Twitterbots are an evolving technology and it is useful to distinguish the earliest or simplest exemplars from their more sophisticated and theory-guided successors. *First-generation* Twitterbots make little use of the rich techniques that linguistic theory has to offer, and rely instead on a combination of superficial language resources – such as word lists, rhyming dictionaries, thesauri – and recombinant aleatoric methods such as the *exquisite corpse* and *cut-up* techniques popularized by the early surrealists and by the beat poets William Burroughs and Brion Gysin ([12],[13]). Popular Twitterbots such as [@Pentametron](#) achieve a great deal with superficial resources; [@Pentametron](#) re-tweets pairings of random tweets that each have ten syllables (for an iambic pentameter reading) and that each end with a rhyming syllable (as in *Pathetic people are everywhere / Your web-site sucks*, [@RyanAir](#)). First-generation Twitterbots do not generate messages from the semantic-level up; rather, they manipulate texts at the word-level, and thus lack any sense of the meaning of a tweet, or any rationale for why one tweet might be better – more provocative, more apt, more re-tweetable – than another. A bot such as [@MetaphorMinute](#), which generates a random metaphor every two minutes (due to usage limitations imposed by the Twitter API) generates a great many outputs that are unintelligible for every one that a user might conceivably (with much effort) interpret as meaningful. In contrast, *next-generation* bots, such as the [@MetaphorMagnet](#) Twitterbot described here, use a panoply of linguistic and semantic techniques to craft their messages from the ground up. These theory-guided bots generate texts with specific rhetorical forms and semantics, to pithily reflect a bot's own semantic model of the world and to exploit its own inferential capabilities. Next-generation Twitterbots can thus generate observations, witticisms and metaphors that they themselves understand and recognize as interesting, surprising or ironic.

The simplest 1<sup>st</sup>-generation bots offer the clearest insights into why people actually follow mechanical content generators on Twitter. Consider [@everycolorbot](#), which simply generates a random six-digit hex-code every hour. As each code denotes a different color from the RGB color space, [@everycolorbot](#) attaches a swatch of the corresponding color to each tweet. The bot's followers, which number in the tens of thousands, favorite and re-tweet its outputs not because they prefer certain RGB codes over others but because they bring their own visual appreciation to bear on each color. Thus, they favorite or retweet a color because of what that color says about their own aesthetics. Or consider [@everyword](#), a bot which simply tweets the next word on its alphabetized inventory of English words every 30 minutes. ([@everyword](#) has since exhausted its word list, generating much media speculation as it neared the end of the

Z's.) The bot, which attracted 100 thousand followers at its peak, tweeted words, not meanings, yet followers brought their own context and their own meanings to bear on those tweets that occasionally (and quite accidentally) resonated with their times. For instance, the word “*woman*” – first tweeted on May 14, 2014 – was retweeted 243 times and favorited 228 times not because followers found the word new or unusual, but because the tweet coincided with the firing of the New York Times’ first female executive editor, in a decision that drew the ire of many for its apparent sexism. First-generation bots do not offer their own meanings or insights, but give us opportunities to see, impose and share meanings of our own. Timely bot tweets are conversational hooks, allowing us to show that we are in on the joke and part of the conversation.

Though metaphors can often be witty, and witticisms are often based on figurative conceits, one does not imply the other. Nonetheless, though [@MetaphorMagnet](#) is principally a generator of metaphors, analogies and similes, its figurative outputs exhibit many of the same characteristics and are shaped by many of the same constraints as witty observations. For instance, Twitter’s 140 character limitation on tweets leads [@MetaphorMagnet](#) to carefully ration its words, to favor brevity over verbosity and suggestiveness over detailed exposition. To attract the attention of new followers and encourage re-tweets from existing followers, the bot also aims to be provocative through its controlled use of semantic and pragmatic incongruity, realized at the textual level via semantic opposition (see [1],[3],[4],[8],[12]) and the violation of expectations. More specifically, [@MetaphorMagnet](#) views the schematic structures of Lakoff & Johnson’s Conceptual Metaphor Theory (or *CMT*; see [2]) as “*scripts*” in the vein of the *SSTH* (the *Semantic Script Theory of Humor*; see [1]) and the *GTVH* (the *General Theory of Verbal Humor*; see [3], [4]). In this paper we show how [@MetaphorMagnet](#) makes use of Twitter norms to turn metaphors into scripts, and to elevate a simple semantic opposition between scripts into a humorous social conflict.

## 2 Modular Concerns in Metaphor and Jokes

Metaphors and jokes share many interesting characteristics. At their best, each allows us to see a familiar situation or idea in a new and perhaps surprising light. Each involves a delicate balance of information, of what is explicitly said by the speaker and of what must be inferred by the listener. Each requires knowledge of words and of the world, and the careful packaging of ideas in a concise linguistic form. In the most thought-provoking instances, each sets out to surprise us by telling us what we already know, by spurring us to see the non-obvious ramifications of our knowledge of the familiar. And each derives a large measure of its success from its ability to evoke a palpable but ultimately resolvable semantic tension: jokes often peak with a closing incongruity that can only be resolved by an act of radical re-categorization, while metaphors present us with a demand for this re-categorization up front, by asking us to see deep similarities between ideas that are superficially very different.

Created to be a generator of novel and occasionally witty linguistic metaphors that are rich in semantic and pragmatic tension, [@MetaphorMagnet](#) relies on many of the same resources and processes that have been identified for joke generation. Applying the *GTVH* of Attardo and Raskin ([3],[4]), as a spirit guide if not a detailed blueprint,

it makes good engineering sense to use a similarly modular approach to the generation of metaphors. The GTVH identifies a variety of knowledge-based modular concerns in joke-generation, called *knowledge resources* (or KRs). The six KRs posited by the GTVH are: *Target* (TA), *Language* (LA), *Narrative Strategy* (NS), *Script Opposition* (SO), *Logical Mechanism* (LM) and *Situation* (SI). Each KR has its own significant role to play in packaging a novel metaphor as an eye-catching, retweet-worthy tweet. Since a metaphor will only strike us as apt if it tells us something of its target that we already feel to be true, the *Target* (TA) resource must ground a figurative comparison in facts or beliefs that most speakers will hold to be true of both the target and the source ideas of the metaphor. *Language* (LA) searches for the most judicious wording for this comparison, to express what explicitly needs to be said and to suggestively evoke the rest, all while staying within the limits of a 140-character tweet. *Narrative Strategy* (NS) gives a logical shape to the tweet, by deciding e.g. whether a figurative conceit should be expressed as a “*what if*” counterfactual, an alternate dictionary definition (in the style of Ambrose Bierce’s *The Devil’s Dictionary*), a “*would you rather*” analogy, an ironic observation, a rueful reflection on a changing world, a work of flash fiction, a sophistic argument, a clash of world views, and so on. In [@MetaphorMagnet](#), the LA and NS resources are tightly coupled, to ensure that the chosen narrative form can naturally be expressed in 140 characters. LA thus exploits Twitter norms such as hashtagging to squeeze maximum value from the medium, and will thus e.g. append the tag [#irony](#) rather than use the phrasing “*Isn’t it ironic.*”

[@MetaphorMagnet](#) employs a range of logical mechanisms (LMs) to juxtapose its knowledge so as to give rise to meaningful semantic oppositions. Consider this tweet:

[#Irony](#): *When the scientists that construct defined models propose the vague abstractions. #Defined=#Vague*

The pivotal opposition here is a semantic one, pitting the property *defined* of *model* against the property *vague* of *abstraction*. This kind of stereotypical association can be harvested automatically from similes such as “*as vague as an abstraction.*” Indeed, Web similes are shown in ([5]) to be an especially rich source of TA knowledge that can be reliably extracted in bulk from online texts. The relational statements *scientists construct models* and *models propose abstractions*, which comprise an important part of [@MetaphorMagnet](#)’s TA knowledge of *science*, are also extracted in bulk from the Web, specifically from the “*why do*” questions that naïve Web users frequently pose to search engines like Google; for example, “*why do scientists construct models?*” is a completion offered by Google for the partial query “*why do scientists.*” In the above observational tweet, an irony-seeking LM notes a potential opposition between the ideas *model* and *abstraction* as connected by the intermediary idea *scientist*. This leads the NS module to frame the opposition as a sardonic observation about *scientists* that the LA module then suggestively labels with an additional hashtag, [#Irony](#). Since [@MetaphorMagnet](#) assumes that connected facts typically belong to the same domain and to the same script (e.g. the domain *science*, or the script *scientist doing science*), the opposition above does not rise to the level of a true Script Opposition. A practical realization of SO that views scripts not as sequences of actions but as figurative world views, and which is thus more appropriate to metaphor generation, is presented in section 4. This SO will be further elaborated to bring in the additional participants and props of a narrative setting that will serve to anchor a tweet in its own Situation (SI).

### 3 Logical Mechanisms for Metaphorical Conceits

Big-budget movies hire specialized individuals to oversee every facet of a production, whereas those on a tight budget force a small number of people to wear multiple hats. So if it seems arbitrary to define a medium like Twitter by something so superficial as its constrictive 140-character limit on texts, this constraint truly does force modular concerns such as LA, NS, SI and LM to work so tightly together that they hardly deserve the label “*modular*”. Not only must NS work hand-in-hand with LA to squeeze its mini-narratives into the cramped confines of a tweet, or perhaps a pair of tweets linked by a shared hashtag that are issued in quick succession, but LM and NS must also be implemented as two sides of a very slim coin. It is not the case that any of [@MetaphorMagnet](#)'s LMs – each of which is designed to seek out a different kind of meaningful opposition in the system's knowledge of familiar topics (its TA) – can work with any of its NS forms. Rather, different LMs are designed to provide material for specific NS strategies that in turn employ specific LA rendering methods.

Consider the interaction of LM, NS and LA that produced the following tweet:

*[#Irony](#): When some high priests manage "welcoming" churches the way jailers manage stifling prisons. [#High\\_priest](#)=[#Jailer](#) [#Church](#)=[#Prison](#)*

The tweet is built upon an analogical chassis by a figurative LM that best corresponds to the GTVH's LM of *False Analogy* (see [4]). TA knowledge of *priests* and *jailers* indicates that each *manages* a very different kind of building, that each carries very different affect (*priests* are respected and carry positive affect, *jailers* are feared and carry negative affect), and that an interesting opposition exists between the property *welcoming* of churches and the property *stifling* of prisons. This opposition would undermine a conventional analogy, but here it offers a sound basis for the logic of false analogy, which in turn offers a sound basis for a narrative strategy that uses the opposition to rail against modern hypocrisy. NS is abetted in this gambit by LA, which affixes an [#Irony](#) tag to the tweet and puts the word *welcoming* in scare quotes. The combination suggests a failure of expectations, an indictment of those priests who should be welcoming congregants as guests but instead oppress them as sinners. That *welcoming* is the mutable property here, rather than *stifling*, is signaled by the use of scare quotes, a decision of the LA that is tightly managed by both the NS and the LM.

Consider the following [@MetaphorMagnet](#) tweet, the fruit of a very different LM:

*Love causes the arguments that create conflicts. Discord causes the confrontations that cause conflicts. Take your pick. [#Love](#)=[#Discord](#)?*

This LM might be called *causal and moral equivalence*: if TA knowledge of a target idea leads a system to conclude that this target is causally similar in some respect to a very different idea – such as one with a very different affective profile – this system might conclude (with a touch of sophistry) that the two ideas are morally equivalent. Thus, because *Love* (a typically positive idea) and *Discord* (a typically negative idea) each cause their own share of conflicts, they might well be considered the same thing. [@MetaphorMagnet](#) employs a simple logical calculus to reason about logical ends (see [6],[12]) in which semantic triples (such as *love causes arguments*) can be chained together to reveal the unexpected distal effects of a familiar idea. More

generally, the triples  $A \rightarrow r1 \rightarrow B$ ,  $B \rightarrow r2 \rightarrow C$  and  $C \rightarrow r3 \rightarrow D$  can be chained to yield the chain  $A \rightarrow r1 \rightarrow B \rightarrow r2 \rightarrow C \rightarrow r3 \rightarrow D$ . Causal propagation rules are used to reason about the effect of the head of a causal chain (e.g.  $A$ ) on the end of a chain (e.g.  $D$ ). For instance, if  $r1$  and  $r2$  have positive causality and  $r3$  has negative causality, a system can reason that more  $A$  causes more  $B$  which causes more  $C$  which causes less  $D$ , so more  $A$  ultimately causes less  $D$ . Though a TA's representation of an idea  $A$  may not directly link  $A$  to  $D$ , [@MetaphorMagnet](#) can infer the causal consequences of  $A$  on  $D$ .

What makes a chain humorous and/or thought-provoking? [@MetaphorMagnet](#) employs a simple but effective criterion: a provocative inference chain is one that links an idea  $A$  to another idea  $D$  by coherently chaining multiple triples together, where there is a bisociative tension between worlds with more  $A$  and worlds with more  $D$ . We expect a positive idea (such as *Love, Beauty, Romance, Art*, etc.) to have positive consequences on the world, by which we mean the proliferation of other positive ideas and the diminution of negative ideas. Likewise, we expect negative ideas (like *War, Hate, Jealousy, Pain*) to have negative consequences on the world, and to diminish the effect of positive ideas. So a chain  $A \rightarrow \dots \rightarrow D$  that shows how a positive idea  $A$  can have a positive causal effect on a negative idea  $D$  (so more  $A$  means more  $D$ ), or shows how a negative idea  $A$  can have a positive causal effect on a positive idea  $D$  (so less  $A$  means less  $D$ ), is considered interesting. A provocative chain will thus show how a target idea can reside in two mutually incongruous frames of reference – one that is desirable and one that is undesirable – thereby conforming to Arthur Koestler's definition of bisociation: "*the perceiving of a situation or idea in two self-consistent but habitually incompatible frames of reference*"([8]). The following tweet from [@MetaphorMagnet](#) illustrates just such a bisociation of views:

*Spouses embrace marriage. Prostitutes profit from the sex that nurtures marriages.  
Who is better? [#Prostitute](#)=[#Spouse](#)*

To support this degree of reasoning by LMs, [@MetaphorMagnet](#)'s TAs assign coarse +/- sentiment classes to individual ideas, and coarse +/- causal classes to individual relations, so that an LM can infer the broad causal effects of the idea at the head of a chain on the idea at its end. Moreover, we have empirically verified our hypothesis (in [6], [12]) that an inferential chain is more likely to be seen as surprising if there is a clear affective incongruity between the head and the tail of a causal chain.

## 4 Script Opposition As A Clash of World Views

The *script* is a necessarily elastic notion in humour research, one that stretches from the frame-like organization of case-roles and fillers in Raskin's SSTH ([1]) to the altogether more pliant graph-theoretic structures of the GTVH as re-imagined by Attardo, Hempelmann and Di Maio ([4]). There are few conventions in language or thought that cannot be subverted for humorous effect, and the notion of *script* in a theory of humour must accommodate them all. The central schematic structure in metaphor theory is the conceptual metaphor ([2],[7]), making this – the conceptual metaphor schema – the figurative equivalent of the *script*. These schemas, such as *Life Is A Journey* or *Politics is a Game*, serve not only as productive deep-structures



for the generation of whole families of linguistic metaphors, but also provide the conceptual mappings that shape our habitual thinking about such familiar concepts as *Life, Love, Emotion* and *Politics*. The SSTH and GTVH view jokes as carefully-crafted texts that set out to trick their audiences into applying a script that is only superficially appropriate, one that ultimately lacks enough explanatory power for subsequent developments in the joke. Politicians and philosophers employ conceptual metaphor schemas to frame an issue and shape our expectations; when a schema fails to match our own experience, we likewise reject it and switch to a more apt schema. So a metaphor-generating bot can seek out thought-provoking incongruity by pitting a metaphor schema against another that advocates a conflicting view of the world. The following tweet from [@MetaphorMagnet](#) contrasts two views on [#Democracy](#):

*To some voters, democracy is an important cornerstone. To others, it is a worthless failure. [#Democracy=#Cornerstone](#) [#Democracy=#Failure](#)*

The schema *Democracy is a Cornerstone* (of civilization) is frequently used to frame political discussions, and can be seen as an elaboration of the schema *Society is a Building*, which in turn elaborates the more primary schema *Organization is Physical Structure* ([7]). Yet the importance of cornerstones to the buildings they anchor finds a sharp contrast in the assertion that *Democracy is a Failure*. Each of these affective claims is so commonly asserted that it can be found in the Google n-grams ([9]), a large database of short fragments of frequent Web texts. Thus the 4-gram “*democracy is a cornerstone*” has a frequency of 91 in the Google n-grams while the 4-gram “*democracy is a failure*” has a frequency of 165. Once again, the stereotypical view of *cornerstones* as *important* and *failures* are *worthless* are themselves derived from Web similes (as in [5]). The following tweet employs a similar metaphorical LM, but renders the conflict of metaphor schemas using a different NS:

*Remember when tolerance was promoted by crusading liberals? Now, tolerance is violence that only fearful appeasers can avoid.*

The LM here – which is guided by the suggestive Google 3-gram “*Tolerance for Violence*” (freq=1353) does not directly contrast the ideas [#Tolerance](#) and [#Violence](#), but examines the juxtaposition at an analogical remove, to find an interesting double conflict, between advocates and opponents and between the advocates of [#Tolerance](#) (*crusading* liberals) and the opponents of [#Violence](#) (*fearful* appeasers). The LA module omits the hashtags [#Tolerance=#Violence](#) from this tweet as it lacks sufficient space to include them within Twitter’s 140-character limit. But LA chooses to split the following conceit across two successive tweets to create space for extra hashtags:

*Remember when research was conducted by prestigious philosophers?  
[#Research=#Fruit](#) [#Philosopher=#Insect](#)*

*Now, research is a fruit eaten only by lowly insects. [#Research=#Fruit](#)  
[#Philosopher=#Insect](#)*

Twitter offers other affordances that allow us to heighten the contrast in metaphorical tweets and to elevate this contrast into a dramatic social situation. So rather than talk of nameless voters or liberals or appeasers, we can give these straw men real names,

or at least invent names that look like the real thing and which, in their choice of Twitter handles, appear wittily apt. The reification of conceptual types into imaginary individuals turns an abstract metaphor into a concrete situation, with its own colorful participants. This is the role of the SI (Situation) KR in [@MetaphorMagnet](#): to bring a metaphor to life by imagining its central conceit as the subject of a vigorous debate by real people. Consider the imaginary debate in this tweet from [@MetaphorMagnet](#):

[.@war\\_poet](#) says history is a straight line  
[.@war\\_prisoner](#) says it is a coiled chain #History=#Line #History=#Chain

The handles [@war\\_poet](#) and [@war\\_prisoner](#) are invented by [@MetaphorMagnet](#)'s SI to suit, and thereby amplify, the metaphorical views that they are fictively advanced in the tweet, again by using a mix of TA knowledge and LA data (Web n-grams). Since poets write poems about the wars that punctuate history, and these poems contain lines, the 2-gram “war poet” is recognized as an apt handle for an imaginary Twitter user who would advance the view of *history as a line*. In this case the handle [@war\\_poet](#) actually denotes a real Twitter user, but this only adds to the sense that Twitterbot confections are a new kind of interactive theatre and performance art [10]. Note that the more profound aspects of this metaphorical contrast are not appreciated by [@MetaphorMagnet](#) itself, or at least not yet. For example, the system does not yet appreciate what it means for history to be a straight line, and while it knows enough to invent the intriguing handle [@war\\_prisoner](#), neither does it appreciate what it might mean to be a prisoner of history, enslaved in a repeating cycle of war. Our bots will always evoke in a human follower much more than they themselves can understand, but this, in the end, is a key ingredient of the allure of Twitterbots, smart or otherwise.

## 5 Evaluation

We argue that [@MetaphorMagnet](#) is a “next-generation” Twitterbot for a number of important reasons. Its actions are informed by, and grounded in, some well-developed theories, from Lakoff & Johnson’s *Conceptual Metaphor Theory* (CMT, ([2]) to the *Semantic Script Theory of Humour* (SSTH) of Raskin ([1]) and the *General Theory of Verbal Humour* (GTVH) of Attardo and Raskin ([1], [3], [4]). Since the bot aims to craft original tweets that are both metaphorically apt and humorously provocative, it represents a practical marriage of CMT and the SSTH/GTVH. Indeed, the bot draws on considerable semantic and linguistic resources to make this marriage work, from a large knowledge-base of conceptual relationships and stereotypical beliefs – which inform its TA (Target) KR – to the rich diversity of the Google n-grams which inform its LA (Language) KR. All of [@MetaphorMagnet](#)'s tweets – all its hits and its misses – are open to public scrutiny on Twitter. But to empirically evaluate the success of the bot as a generator of novel, meaningful and retweet-worthy metaphors, we turn to the crowdsourcing platform [CrowdFlower](#). To determine just how much of its success can be attributed to its use of CMT/GTVH mechanisms and knowledge resources, we perform a comparative analysis between this knowledge-based bot and a knowledge-free bot called [@MetaphorMinute](#) (designed by noted bot-maker Darius Kazemi) that uses a wholly aleatoric approach to metaphor generation. [@MetaphorMinute](#) crafts its



metaphors by filling a template with nouns and adjectives that are chosen more-or-less at random, to produce tweets such as “*a doorbell is a sportsman: fleetwide and infraclavicular.*” Though it generates inscrutable outputs such as these *every two minutes*, [@MetaphorMinute](#) is a popular bot that currently has over 500 followers.

We chose 60 tweets at random from the outputs of each Twitterbot. CrowdFlower annotators were not informed of the origin of any tweet, but simply told that each was collected from Twitter because of its metaphorical content. For each tweet, annotators were asked to rate its metaphor along three dimensions, *Comprehensibility*, *Novelty* and likely *Retweetability*, and to rate all three dimensions on the same scale, ranging from *Very Low* to *Medium Low* to *Medium High* to *Very High*. CrowdFlower was used to solicit ten annotations per tweet (and thus, per dimension), though scammers (non-engaged annotators) were later removed from this pool. Table 1 presents the distributions of mean ratings per tweet, along each dimension and for each Twitterbot.

**Table 1.** Comparative Evaluation of the [@MetaphorMagnet](#) and [@MetaphorMinute](#) bots.

<i>Rating</i>	<i>Comprehensibility</i>		<i>Novelty</i>		<i>Retweetability</i>	
<i>Metaphor System</i>	<i>Metaphor Magnet</i>	<i>Metaphor Minute</i>	<i>Metaphor Magnet</i>	<i>Metaphor Minute</i>	<i>Metaphor Magnet</i>	<i>Metaphor Minute</i>
<i>Very Low</i>	11.6%	<b>23.9%</b>	<b>11.9%</b>	9.5%	15.5%	<b>41%</b>
<i>Med. Low</i>	13.2%	<b>22.2%</b>	<b>17.3%</b>	12.4%	<b>41.9%</b>	34.1%
<i>Med High</i>	<b>23.7%</b>	22.4%	<b>21%</b>	14.9%	<b>27.4%</b>	15%
<i>Very High</i>	<b>51.5%</b>	31.6%	49.8%	<b>63.2%</b>	<b>15.3%</b>	9.9%

Note how more than half of [@MetaphorMagnet](#)’s tweets are ranked as very highly comprehensible, while less than a third of [@MetaphorMinute](#)’s tweets are so ranked. Even though only 1 in 4 of [@MetaphorMagnet](#)’s metaphors is rated as being hard or somewhat hard to comprehend, this is an area of performance that can be improved. More surprising is the result that raters found more than half of [@MetaphorMinute](#)’s wholly random metaphors to be of medium-high to very-high comprehensibility. The bot’s use of abstruse terminology, like *fleetwide* and *infraclavicular*, may be a factor here, as might the bot’s use of the familiar copula template for metaphors, which may well seduce raters into believing that an apparent metaphor really does have a comprehensible meaning, if only one were to expend enough effort to discern it.

The dimension *Novelty* yields results that are equally thought-provoking, for while one half of [@MetaphorMagnet](#)’s metaphors are ranked as very-highly novel, almost two-thirds of [@MetaphorMinute](#)’s metaphors are so ranked. Nonetheless, we should not be overly surprised that [@MetaphorMinute](#)’s bizarre combinations of rare words, as yielded by its unconstrained use of aleatoric techniques, are seen as more unusual than those word combinations arising from [@MetaphorMagnet](#)’s controlled use of Web n-grams and stereotypical knowledge. As demonstrated in [11], novelty is not in itself a source of pleasure or a reliable benchmark of creativity. Pleasurability derives from *useful* novelty, that is, novelty that can be understood and usefully exploited.

In this case of Twitter, useful exploitation is frequently a matter of social reach. A tweet is novel and useful to the extent that it attracts the attention of Twitter users and is deemed worthy of re-tweeting to others in their social circles. Our third dimension, *Re-Tweetability*, reflects the likelihood that an annotator would consider re-tweeting a given metaphor to others. Though we ask annotators to speculate here – neither bot has enough followers to perform a robust statistical analysis of actual retweet rates – the results largely conform to our expectations. Retweetability, it seems, is a matter of novelty *and* comprehensibility, and not novelty alone. Though raters are not generous with their *Very-High* ratings for either bot, [@MetaphorMagnet](#)'s tweets are deemed to be significantly more re-tweetable than the random offerings of [@MetaphorMinute](#).

This is just as well, given the considerable gap in complexity and sophistication that exists between the two bots. But this is an encouraging result not just for theory-informed Twitterbots like [@MetaphorMagnet](#) and their creators, but for Twitter itself. Twitter offers a compelling platform for research in interactive humour and metaphor, not least because its human users appreciate these phenomena when they see them.

## References

1. Raskin, V. *Semantic Mechanisms of Humor*. Dordrecht. Reidel. (1985)
2. Lakoff, G., Johnson, J. *Metaphors We Live By*. University of Chicago Press. (1980)
3. Attardo, S., Raskin, V. Script theory revis(it)ed: joke similarity and joke representational model. *Humor: International Journal of Humor Research*, vol. 4, no. 3, pp 293-347. (1991)
4. Attardo, S., Hempelmann, C.F., Di Maio, S. Script oppositions and logical mechanisms: Modeling incongruities and their resolutions. *Humor: International Journal of Humor Research*, vol. 15, no. 1, pp 3-46. (2002)
5. Veale, T., Hao, Y. Comprehending and Generating Apt Metaphors: A Web-driven, Case-based Approach to Figurative Language. In *Proceedings of AAAI- 2007, the 22<sup>nd</sup> AAAI conference of the Association for the Advancement of Artificial Intelligence*. (2007)
6. Veale, T., Valitutti, A. A World With or Without You\* (\*Terms and Conditions May Apply). In *Proceedings of AAAI-2014 Fall Symposium Series on Modeling Changing Perspectives: Re-conceptualizing Sensorimotor Experiences*. Arlington, VA. (2014)
7. Grady, J. *Foundations of Meaning: Primary Metaphors and Primary Scenes*. Berkeley, University of California. (1997)
8. Koestler, A. *The Act of Creation*. Penguin Books. (1964)
9. Brants, T., Franz, A. *Web 1T 5-gram Version 1*. Linguistic Data Consortium. (2006)
10. Dewey, D. What happens when @everyword ends? *Washington Post, Intersect* column, [May 23 edition](#). (2014)
11. Giora, R., Fein, O., Kronrod, A., Elnatan, I., Shuval, N., Zur, A. Weapons of Mass Distraction: Optimal Innovation and Pleasure Ratings. *Metaphor and Symbol* vol. 19, no. 2, pp 115-141. (2004)
12. Veale, T. Running With Scissors: Cut-Ups, Boundary Friction and Creative Reuse. *Proceedings of ICCBR-2014, the 22<sup>nd</sup> International Conference on Case-Based Reasoning*, Cork, Ireland, September. (2014)
13. Burroughs, W.S. *The Cut-Up Method*. LeRoi Jones (Ed.) *The Moderns: An Anthology of New Writing in America*. New York, Corinth Books. (1963)