Transactions of the SDPS:
Journal of Integrated Design and Process Science
Volume 17, Number 4 / 2013
DOI 10.3233/jid-2013-0012
http://www.sdpsnet.org

Linguistic Readymades and Creative Reuse

Tony Veale

Web Science and Technology Division, KAIST, Daejeon, South Korea

Abstract Creativity often arises from a process of *appropriation*, in which something is wrenched from its normative context of use and given new meaning in a new setting. In this vein, Marcel Duchamp popularized the notion of an artistic *ready-made* when his *Fountain* – a signed urinal – was presented with some controversy at a Dada exhibition in 1917. We normally think of readymades as physical objects whose artistic merit derives wholly from their selection by an artist, but language is also rich in linguistic readymades. Just think of how many movies, songs, novels and poems allusively borrow utterances and phrases from each other. The movie *The Usual Suspects*, for example, takes its name from a famous quote from the movie *Casablanca*, while the novel *All The King's Men* takes it title from a famous nursery rhyme; this title, in turn, inspired the title of Woodward and Bernstein's book *All The President's Men*. Large lexical resources, such as corpora and databases of Web n-grams, are a rich source of readymade phrases that can be reused in many different contexts. However, one must be careful in how these resources are used, and noted writers such as George Orwell have argued that the use of canned phrases encourages sloppy thinking and results in poor communication. Nonetheless, while Orwell prized home-made phrases over the readymade variety, Duchamp created a vibrant movement in modern art which shifted the emphasis of artistic creation from the production of novel artifacts to the clever reuse of readymades or *objets trouvés*. We describe here a system that makes creative reuse of the linguistic readymades in the *Google* n-grams. Our computational systems thus owe more to Duchamp than to Orwell, and harvest readymades on a large scale to support linguistic creativity.

Keywords: Readymades, Creativity, Reuse, Information Retrieval

1. Introduction

In a much-quoted essay from 1946 entitled *Politics and the English Language*, the writer and thinker George Orwell outlines his prescription for halting a perceived decline in the English language. He argues that language and thought form a tight feedback cycle that may be either virtuous or vicious. Lazy language can promote lazy thinking, and vice versa. Orwell pours scorn on two particular forms of lazy language: the expedient use of overly familiar metaphors merely because they come quickly to mind, even though they have lost their power to evoke vivid images; and the use of readymade turns of phrase as substitutes for individually crafted expressions. While a good writer bends words to his meaning, Orwell worries that a lazy writer bends his meaning to convenient words and phrases.

Orwell is especially scornful about readymade phrases which, when over-used, "are tacked together like the sections of a prefabricated henhouse." A writer who operates by "mechanically repeating the familiar phrases" and "gumming together long strips of words which have already been set in order by someone else" has, he argues, "gone some distance toward turning himself into a machine." Given his derogatory mechanistic view of the use of readymade phrases, Orwell would not be surprised to learn that

 Published by IOS Press

computers are highly proficient in the large-scale use of familiar phrases, whether acquired from large text corpora or from the Google n-grams (see Brants and Franz, 2006).

Though argued with passion, there are serious holes in Orwell's logic. If one should "never use a metaphor, simile or other figure of speech which you are used to seeing in print", how then are familiar metaphors ever to become *dead* metaphors and thereby enrich the language with new terms and new senses? And if one cannot use familiar readymade phrases, how can one make playful – and creative – allusions to the writings of others, or mischievously subvert the conventional wisdom of platitudes and clichés? Orwell's use of the term *readymade* is entirely negative, yet the term is altogether more respectable in the world of modern art, thanks to its use by artists such as Marcel Duchamp. For many artists, a readymade object is not a substitute, but a starting point, for creativity.

Also called an *objet trouvé* or *found object*, a readymade emerges from an artist's encounter with an object whose aesthetic merits are overlooked in its banal, everyday contexts of use. When this object is moved to an explicitly artistic context, such as an art gallery, viewers are better able to appreciate these merits. The artist's insight is to recognize the transformational power of this non-obvious context switch. Perhaps the most famous (and notorious) readymade in the world of art is Marcel Duchamp's *Fountain*, a humble urinal that becomes an elegantly curved piece of sculpture when viewed with the right mind set. Duchamp referred to his *objets trouvés* as "assisted readymades" because they allow an artist to remake the act of creation as one of pure insight and inspired recognition rather than one of manual craftsmanship (see Taylor, 2009). In computational terms, the Duchampian notion of a readymade allows creativity to be modelled not as a construction problem but as a decision problem. A computational Duchamp need not explore an abstract conceptual space of potential ideas, as in Boden (1994). However, a Duchampian agent must instead be exposed to the multitude of potentially inspiring real-world stimuli that a human artist encounters everyday.

Readymades represent a serendipitous form of creativity that is poorly served by exploratory models of creativity, such as that of Boden (1994), and better served by the investment models such as the *buy-low-sell-high* theory of Sternberg and Lubart (1995). In this view, creators and artists find unexpected or untapped value in unfashionable objects or ideas that already exist, and quickly move their gaze elsewhere once the public at large come to recognize this value. Duchampian creators invest in everyday objects, just as Duchamp found artistic merit in urinals, bottles and combs. From a linguistic perspective, these everyday objects are commonplace words and phrases which, when wrenched from their conventional contexts of use, are free to take on enhanced meanings and provide additional returns to the investor. The realm in which a maker of linguistic readymades operates is not the real world, and not an abstract conceptual space, but the realm of texts: large corpora become rich hunting grounds for investors in linguistic *objets trouvés*.

This perspective is given computational form in the following sections. We show how a rich vocabulary of commonplace cultural stereotypes can be acquired from the Web, and how this vocabulary facilitates the implementation of a decision procedure for recognizing potential readymades in large corpora – in this case, the Google database of Web *n-grams* (Brants and Franz, 2006). This decision procedure provides a robust basis for a creative simile-generation system called *The Jigsaw Bard*. The cognitive / linguistic intuitions that underpin the *Bard*'s concept of textual readymades are put to the empirical test in section 5. While readymades remain a contentious notion in the public's appreciation of artistic creativity – despite Duchamp's *Fountain* being considered one of the most influential artworks of the 20th century – we shall show that the notion of a linguistic readymade has significant practical merit in the realms of text generation and computational creativity.

2. Linguistic Readymades

Readymades are the result of artistic *appropriation*, in which an object with cultural resonance – an image, a phrase, a quote, a name, a thing – is re-used in a new context that prompts us to interpret it with a new and more figurative (perhaps even ironic) sense. As a fertile source of cultural reference points, language is an equally fertile medium for appropriation. Thus, in the constant swirl of language and culture, movie quotes suggest song lyrics, which in turn suggest movie titles, which suggest book titles, or restaurant names, or the names of racehorses, and so on, and on. The 1996 movie *The Usual Suspects* takes its name from a memorable scene in 1942's *Casablanca*, as does the Woody Allen play and movie *Play it Again Sam*. The 2010 art documentary *Exit Through the Gift Shop*, by graffiti artist Banksy, takes its name from a banal sign sometimes seen in museums: the sign, suggestive as it is of creeping commercialism, is the perfect readymade for a film that laments the mediocrity of commercialized art.

Appropriations can also be combined to produce novel mashups; consider, for instance, the use of tweets from rapper Kanye West as alternate captions for cartoon images from the *New Yorker* magazine (see hashtag *#KanyeNew-YorkerTweets*). Hashtags can themselves be linguistic readymades, and when free-speech advocates use the hashtag *#IAMSpartacus* to show solidarity with users whose tweets have incurred the wrath of the law, they are appropriating an emotional line from the 1960 film *Spartacus*. More recently, when British protesters wanted to show their disdain for the legacy of Margaret Thatcher upon the death of that polarizing politician, they encouraged others to buy the 1939 song "*Ding Dong, The Witch is Dead*" (from the *Wizard of Oz*) on *iTunes*, causing that song to climb up the music charts and gain a new political resonance for a new era. Whereas only a decade ago such a protest would require the cooperation of a music publisher to make the song available for purchase, the diverse products of bygone ages are readily available on the Web for opportunistic reuse whenever the occasion arises. Linguistic readymades, then, are fragments of text (or textual elements of multimedia objects) that are resonant and highly quotable because they carry a figurative content that is reusable in different contexts.

The quotes "*round up the usual suspects*" and "*I am Spartacus*" require a great deal of cultural knowledge to appreciate. Since literal semantics only provides a small part of their meaning, a computer's ability to recognize linguistic readymades is only as good as the knowledge at its disposal. We explore here a modest form of readymade phrases that can be used as evocative image builders in similes, as in:

a wet haddock
snow in January
a robot fish
a bullet-ridden corpse

Each phrase can be found in the Google 1T database of Web n-grams – snippets of Web text (of one to five words) that occur on the Web with a frequency of 40 or higher (Brants and Franz, 2006). Each is likely a literal description of a real object or event – even "robot fish", which describes an autonomous marine vehicle whose movements mimic real fish. But each exhibits figurative potential as well, providing a memorable description of physical or emotional coldness. Whether or not each was ever used in a figurative sense before is not the point: once this potential is recognized, each phrase becomes a reusable linguistic readymade for the construction of a vivid figurative comparison, as in "*as cold as a robot fish*". We now consider the building blocks from which these comparisons can be ready-made.

3. A Vocabulary of Cultural Stereotypes

How can a computer acquire the knowledge that *fish, snow, January, bullets* and *corpses* are all cultural signifiers of coldness? Much the same way that humans acquire this knowledge: by attending to how such signifiers are used by others, in cultural clichés and proverbial similes like "as cold as a fish".

In fact, folk similes are an important vector in the transmission of cultural knowledge: they point to, and exploit, the shared cultural touchstones that speakers and listeners alike can use to construct and intuit meanings. Taylor (1954) catalogued thousands of proverbial comparisons and similes from California, identifying just as many building blocks in the construction of new phrases and figurative meanings. Only the most common similes can be found in dictionaries, as shown by Norrick (1986), while Moon (2008) demonstrates that large-scale corpus analysis is needed to identify folk similes with a breadth approaching that of Taylor's study. However, Veale and Hao (2007) and Veale (2011, 2012a, 2012b, 2013) show that the World-Wide Web is the ultimate resource for harvesting similes.

Veale and Hao use the Google API to find many instances of the pattern "*as ADJ as a | an* *" on the Web, where ADJ is an adjectival property and * is the Google wildcard. WordNet (Fellbaum, 1998) is used to provide a set of over 2,000 different values for ADJ, and the text snippets returned by Google are parsed to extract the basic simile bindings. Once the bindings are annotated to remove noise, as well as frequent uses of irony, this Web harvest produces over 12,000 cultural bindings between a noun (such as *fish*, or *robot*) and its most stereotypical properties (such as *cold, wet, stiff, logical, heartless*, etc.). Stereotypical properties are acquired for approx. 4,000 common English nouns. This is a set of building blocks on a larger scale than even that of Taylor (1954), allowing us to build on Veale and Hao (2007) to identify readymades in their hundreds of thousands in the Google n-grams.

However, to identify readymades as resonant variations on cultural stereotypes, we need a certain fluidity in our treatment of adjectival properties. The phrase "*wet haddock*" is a readymade for coldness because "wet" accentuates the "cold" that we associate with "haddock" (via the web simile "*as cold as a haddock*"). In the words of Hofstadter (1995), we need to build a *SlipNet* of properties whose structure captures the propensity of properties to mutually and coherently reinforce each other, so that phrases which subtly accentuate an unstated property can be recognized. In the vein of Veale and Hao (2007), we use the Google API to harvest the elements of this *SlipNet*.

We hypothesize that the construction "*as* ADJ_1 *and* ADJ_2 *as*" shows ADJ_1 and ADJ_2 to be mutually suggestive properties, since they can be seen to work together as a single complex property in a single comparison. Thus, using the full complement of adjectival properties used by Veale and Hao (2007), we harvest all instances of the patterns "*as ADJ and * as*" and "*as * and ADJ as*" from Google, noting the combinations that are found and their frequencies. These frequencies provide link weights for the Hofstadter-style SlipNet that is then constructed. In all, over 180,000 links are harvested, connecting over 2,500 adjectival properties to one other. We put the intuitions behind this SlipNet to the test in section 6.

4. Creative Information Retrieval

In language, creativity is always a matter of construal. While conventional IR (information retrieval) queries articulate a need for information, creative IR queries articulate a need for expressions to convey the same meaning in a fresh or unusual way. A query and a matching phrase can be figuratively construed to have the same meaning if there is a non-literal mapping between the elements of the query and the elements of the phrase. In creative IR, this non-literal mapping – in which query words are mapped non-identically to target phrases on their basis of their potential to figuratively stand for those target words – is facilitated by the query's explicit use of semantic *wildcards* (see Veale [2011] and Mihalcea [2002]).

The conventional wildcard * is a boon for power-users of the Google search engine, precisely because it allows users to focus on the retrieval of matching phrases rather than relevant documents. For instance, * can be used to find alternate ways of instantiating a culturally-established linguistic pattern, or "*snowclone*": thus, the Google queries "*In * no one can hear you scream*" (from the movie *Alien*), "*Reader, I * him*" (from the novel *Jane Eyre*) and "*This is your brain on* *" (from a famous TV advert) find new ways in which old patterns have been instantiated for humorous effect on the Web. On a larger scale, Veale and Hao (2007) use the * wildcard to harvest web similes, but report that harvesting cultural data with wildcards is not a straightforward process. Google and other engines are designed to maximize document relevance and to rank results accordingly. They are not designed to maximize the diversity of results, or even to find the largest set of bindings for wildcards. Nor are they designed to find the most interesting bindings, or even the most commonplace.

Following Guilford's pioneering work (e.g. Guilford, 1967), a diversity and divergence of outputs is widely considered a key dimension in evaluating creative behavior (see also de Bono, 1970, 1971). By focusing on phrases rather than documents, and by returning phrase sets rather than document sets, creative IR maximizes diversity by finding as many bindings for its wildcards as a text collection will support. But we need more flexible and precise wildcards than the all-embracing *. We now consider four new semantic wildcards that enable the retrieval of candidate phrases for creative reuse as readymades.

4.1. The *Pragmatic Neighborhood* Wildcard ?X

Semantic query expansion expands a query term **X** into the set {X, X_1, X_2, …, X_n} where each X_i is related to **X** by a prescribed lexico-semantic relationship, such as synonymy, hyponymy or meronymy. A lightweight open-domain resource like WordNet can provide these relations, or a specific domain ontology can be used if one is available. Intuitively, each query term suggests other terms from its semantic neighborhood, yet there are practical limits to this intuition. X_i may not be an obvious or natural substitute for X. A neighborhood can be drawn too small, impacting recall, or too large, impacting precision.

Corpus analysis suggests an approach that is both semantic *and* pragmatic. As noted in Hanks (2005), languages provide constructions for building ad-hoc sets of items that can be considered comparable in a given context. For instance, a coordination of bare plurals suggests that two ideas are related at a generic level, as in "*priests and imams*" or "*mosques and synagogues*". More generally, consider the pattern "*X and Y*", where X and Y are proper-names (e.g., "*Zeus and Hera*"), or X and Y are inflected nouns or verbs with the same inflection (e.g., the plurals "*cats and dogs*" or the present-continuous forms "*kicking and screaming*"). Millions of matches for this pattern can be found in the Google 3-grams (Brants and Franz, 2006), allowing us to build a matrix of comparable terms by associating the root-forms of X and Y with a semantic similarity score obtained via a standard WordNet-based measure (e.g. see Veale and Li, 2013). The pragmatic neighborhood of a term X can now be defined as {X, X_1, X_2, …, X_n}, such that for each X_i, the Google 3-grams contains "*X+inf and X_i+inf*" or "*X+inf and X_i+inf*". These neighborhoods have pragmatic rather than semantic boundaries: if **?X** denotes the neighborhood of X, then **?artist** matches not just *artist*, *composer* and *poet*, but *studio*, *gallery*, *portfolio* and *exhibition,* and many other terms that are semantically dissimilar but pragmatically related to artist. Since each X_i in **?X** is ranked by similarity to **X**, matches for queries can also be ranked by similarity to the original term.

When X is an adjective, then **?X** matches any element of {X, ..,, X_i, …, X_n}, where each X_i pragmatically reinforces X and where X pragmatically reinforces X_i. To ensure that X and X_i really are mutually reinforcing adjectives, in the sense that X reinforces our impression of X_i–*ness* and vice versa, we use the double-ground simile pattern "*as* X *and* X_i *as*" to harvest co-descriptors for each X. Moreover, to maximize recall, we use the Google API (rather than Google n-grams) to harvest suitable bindings for X and X_i from the Web. For example, **?witty** = {*charming, clever, intelligent, entertaining, …, edgy, fun*}.

4.2. The *Cultural Stereotype* Wildcard @X

Dickens opens *A Christmas Carol* by claiming that "the wisdom of a people is in the simile". Similes exploit familiar stereotypes to describe a less familiar concept, so one can learn a great deal about a culture and its language from the similes that have the most currency (Taylor, 1954). The wildcard **@X** builds on the results of Veale and Hao (2007) to allow creative IR queries to retrieve matches on the basis of strong cultural expectations. This foundation provides a large set of adjectival features (over 2000 in total) for a larger set of nouns (over 4000) denoting stereotypes for which these features are highly salient.

If **N** is a noun, then **@N** matches any element of the set {A_1, A_2, …, A_n}, where each A_i is an adjective denoting a stereotypical property of **N**. For example, **@diamond** matches any element of {*transparent, immutable, beautiful, tough, expensive, valuable, shiny, bright, desirable, strong, …, hard*}. If **A** is an adjective, then **@A** matches any element of the set { N_1, N_2, …, N_n}, where each N_i is a noun denoting a stereotype for which **A** is a culturally established property. For example, **@tall** matches any element of {*giraffe, skyscraper, tree, redwood, tower, sunflower, lighthouse, beanstalk, castle … rocket*}. Stereotypes crystallize in a language as clichés, so one can argue that stereotypes and clichés are little or no use to a creative IR system. Yet, as demonstrated in Fishelov (1992), creative language is replete with stereotypes, not in their tired and clichéd guises, but in novel, imaginative and often incongruous combinations. The creative value of a stereotype lies in how it is used, as we'll see in section 6.

4.3. The *Ad-Hoc Category* Wildcard ^X

Barsalou (1983) introduced the notion of an *ad-hoc category*, a cross-cutting collection of often disparate elements that cohere in the context of a specific task. The ad-hoc nature of these categories is reflected in the difficulty we often have in naming them concisely: the cumbersome "things to take on a camping trip" is Barsalou's most cited example. But ad-hoc categories do not replace natural kinds; rather, they supplement our existing system of rigid categories that are optimized for convergent reasoning, such as the categories found in WordNet, with more flexible categories that support more divergent reasoning.

The semantic wildcard **^C** matches not just C but any element $\{C_1 \ldots C_n\}$ of the category *C*. **^C** may denote a fixed category in a resource like WordNet or Wikipedia; thus, **^fruit** will match any member of {*apple, orange, banana, ..., lemon*} while **^animal** matches any member of *{dog, cat, mouse, ..., fox}*. Veale and Li (2013) show how a very large and open-ended number of fine-grained categories – such as *traumatic-event* or *corrosive-substance* – can be acquired and populated from Web content. Using these categories with the ^ operator brings an element of divergent production (Guilford 1961) and lateral thinking (de Bono, 1970, 1971) to the retrieval of linguistic readymades. For instance, **^traumatic-event** will match both *war* and *divorce*, while **^corrosive-substance** will match both *acid* and *cola*.

4.4. The *Antonym* Wildcard -X

The playful conflict with expectations that one finds in ironic comments, or the creative tension one recognizes within a novel metaphor, often arise from the deliberate use of conceptual opposition. Though opposition can arise from a variety of sources and strategies, the use of antonyms is perhaps the easiest to model computationally. For instance, antonymous relationships can easily be extracted from the lexical resource WordNet. The resulting antonym operator is denoted **-**, so that **-P** will match any antonym of the adjective **P**. Thus, for example, -*soft* matches the word "hard", and -*strong* matches the word "weak".

4.5. Compound Operators

Each wildcard maps a single query term onto a set of expansion terms, thus expanding the number of retrievable phrases. The semantics of compound wildcards can be understood in simple set-theoretic terms. We consider here the most obvious and useful combinations of the **?**, **@** and ^ operators.

??

Neighbour-of-a-neighbour: if **?X** matches any element of $\{X, X_1, X_2, \ldots, X_n\}$ then **??X** matches any element of $\mathbf{?X} \cup \mathbf{?X_1} \cup \ldots \cup \mathbf{?X_n}$, where the ranking of X_{ij} in **??X** is a function of the ranking of X_i in **?X** and the ranking of X_{ij} in $\mathbf{?X_i}$. Thus, **??artist** matches a far larger set of terms than **?artist**, producing more noise but more potential for creativity.

@@

Stereotype-of-a-stereotype: if **@X** matches any element of $\{X_1, X_2, \ldots, X_n\}$ then **@@X** matches any of $\mathbf{@X_1} \cup \mathbf{@X_2} \cup \ldots \cup \mathbf{@X_n}$. For instance, **@@diamond** matches any stereotype noun that shares a salient property with *diamond* (such as *laser*, *scalpel* and *pearl*) and **@@sharp** matches any salient property of any stereotype noun for which *sharp* is a stereotypical property (such as *focused* and *lethal*).

?@

Neighborhood-of-a-stereotype: if **@X** matches any element of $\{X_1, X_2, \ldots, X_n\}$ then **?@X** matches any of $\mathbf{?X_1} \cup \mathbf{?X_2} \cup \ldots \cup \mathbf{?X_n}$. Thus, **?@cunning** matches any term in the pragmatic neighborhood of a stereotype for *cunning*, while **?@knife** matches any property that mutually reinforces any stereotypical property of *knife*.

@?

Stereotypes-in-the-neighborhood: if **?X** matches any of $\{X, X_1, X_2, \ldots, X_n\}$ then **@?X** matches any of **@X** $\cup$ **@**$\mathbf{X_1}$ $\cup \ldots \cup$ **@**$\mathbf{X_n}$. Thus, **@?corpse** matches any salient property of any stereotype in the neighborhood of *corpse*, while **@?fast** matches any stereotype noun with a salient property that is similar to, and reinforced by, *fast*.

?^

Neighborhood-of-a-category: if **^C** matches any of $\{C, C_1, C_2, \ldots, C_n\}$ then **?^C** matches any of **?C** $\cup$ **?**$\mathbf{C_1}$ $\cup \ldots \cup$ **?**$\mathbf{C_n}$.

^?

Categories-in-the-neighborhood: if **?X** matches any of $\{X, X_1, X_2, \ldots, X_n\}$ then **^?X** matches any of **^X** $\cup$ **^**$\mathbf{X_1}$ $\cup \ldots \cup$ **^**$\mathbf{X_n}$.

@^

Stereotypes-in-a-category: if **^C** matches any of $\{C, C_1, C_2, \ldots, C_n\}$ then **@^C** matches any of **@C** $\cup$ **@**$\mathbf{C_1}$ $\cup \ldots \cup$ **@**$\mathbf{C_n}$.

^@

Members-of-a-stereotype-category: if **@X** matches any element of $\{X_1, X_2, \ldots, X_n\}$ then **^@X** matches any of **^**$\mathbf{X_1}$ $\cup$ **^**$\mathbf{X_2}$ $\cup \ldots \cup$ **^**$\mathbf{X_n}$. Thus, **^@strong** matches any member of any category (such as *warrior*) that is stereotypically *strong*.

5. Harvesting Linguistic Readymades from Corpora

In the course of an average day, a creative writer is exposed to a constant barrage of linguistic stimuli, any small portion of which can strike a chord as a potential readymade. In this casual inspiration phase, the observant writer recognizes that a certain combination of words may produce, in another context, a meaning that is more than the sum of its parts. Later, when an apposite phrase is needed to strike a particular note, this combination may be retrieved from memory (or from a trusty notebook), *if* it has been recorded and suitably indexed.

Ironically, Orwell (1946) suggests that lazy writers “shirk” their responsibility to be “scrupulous” in their use of language by “simply throwing [their] mind open and letting the ready-made phrases come crowding in”. For Orwell, words just get in the way, and should be kept at arm’s length until the writer has first allowed a clear meaning to crystallize. This is dubious advice, as one expects a creative writer to keep an open mind when considering *all* the possibilities that present themselves. Nonetheless, Orwell’s proscription suggests how a computer should go about the task of harvesting readymades from corpora: by throwing its mind open to the possibility that a given n-gram may one day have a second life as a creative readymade in another context, the computer allows the phrases that match some simple image-building criteria to come crowding in, so they can be stored in a database.

Given a rich vocabulary of cultural stereotypes and their properties, computers are capable of indexing and recalling a considerably larger body of resonant combinations than the average human. The necessary barrage of linguistic stimuli can be provided by the Google database of Web n-grams (Brants and Franz, 2006). Trawling these n-grams, a modestly attentive computer can recognize well-formed combinations of cultural elements that might serve as a vivid vehicle of description in a future comparison. So in the following sub-sections we consider how one can computationally view the processes of metaphor and simile generation, as well as ironic description, as the retrieval of a readymade with the right resonances.

5.1. Metaphor Generation

For a term X, the wildcard **?X** suggests those other terms that writers have considered to be comparable to X, while **??X** extrapolates beyond the corpus evidence to suggest an even larger space of potential comparisons. A meaningful metaphor can be constructed for X by framing X with any stereotype to which it is pragmatically comparable, that is, any stereotype in **?X**. Collectively, these stereotypes can impart the properties **@?X** to X.

Suppose one wants to metaphorically ascribe the property P to X. The set **@P** contains those stereotypes for which P is culturally salient. Thus, close metaphors for X (what MacCormac (1985) dubs *epiphors*) in the context of P are suggested by **?X ∩ @P**. More distant metaphors (what MacCormac dubs *diaphors*) are suggested by **??X ∩ @P**. For instance, to describe a *scholar* as *wise*, one can use *poet*, *yogi*, *philosopher* or *rabbi* as comparisons. Yet even a simple metaphor will impart other features to a topic. Let $\mathbf{\hat{}P_S}$ denotes the ad-hoc set of additional properties that may be inferred for X when a stereotype S is used to convey property P, where $\mathbf{\hat{}P_S}$ = **?P ∩ @@P**. The query "$\mathbf{\hat{}P_S}$ **X**" identifies corpus-attested elements of $\mathbf{\hat{}P_S}$ that can meaningfully be used to modify X.

These IR formulations are used by *Aristotle*, an online metaphor generator, to generate targeted metaphors that highlight a property P in a topic X. *Aristotle* uses the Google n-grams to supply values for **?X**, **??X**, **?P** and $\mathbf{\hat{}P_S}$. The system can be accessed at:

www.educatedinsolence.com/aristotle

5.2. Simile Generation

The well-formed phrases of a large corpus can be viewed as the linguistic equivalent of *objets trouvés* in art: readymade or "found" objects that might take on fresh meanings in a creative context. The phrase "*robot fish*", for instance, denotes a literal object in its home context (autonomous robotic submersibles), but can be used to convey a figurative meaning as part of a creative comparison (e.g., "*he was as cold as a robot fish*").

Fishelov (1992) argues that poetic comparisons are most evocative when they combine mutually-reinforcing (if distant) ideas, to create more memorable images and more resonant feelings. Building on Fishelov's argument, creative IR can be used to turn the readymade phrases of the Google n-grams into vehicles for creative comparison. For a topic X and a property P, simple similes of the form "***X** is as **P** as **S***" are easily generated, where **S** ∈ **@P ∩ ??X**. Fishelov dubs these *non-poetic similes* (NPS). However, the query "***?P @P***" retrieves corpus-attested elaborations of stereotypes in **@P** to suggest similes of the form "***X** is as **P** as* $\mathbf{P_1}$ ***S***" where $\mathbf{P_1}$ ∈ **?P**. These similes exhibit elements of what Fishelov dubs *poetic similes* (PS). Why say "*as cold as a fish*" when you can say "*as cold as a wet fish*", "*as a dead haddock*", "*a wet January*", "*a frozen corpse*", or even "*a heartless robot*"? Complex queries can retrieve more creative combinations, so "***@P @P***" (e.g. "*robot fish*", "*snow storm*"), "**?P @P @P**" (e.g. "*creamy chocolate mousse*") and "***@P - ^pastparticiple @P***" (e.g. "*snow-covered graveyard*", "*bullet-riddled corpse*") each retrieve n-grams that juxtapose two different stereotypes of the same property, such as *cold*.

Blended properties also make for nuanced similes of the form "*as **P** and **?P** as **S***", where **S ∈ @P ∩ @?P**. While one can be "*as rich as a fat king*", something can be "*as rich and enticing as a chocolate truffle*", "*a chocolate brownie*", "*a chocolate fruitcake*", and stretching to a diaphor or two, "*a chocolate prince*" and "*a chocolate millionaire*".

The *Jigsaw Bard* is a web application that harnesses the readymades of the Google n-grams to formulate novel similes from existing phrases. By mapping blended properties to n-gram phrases that combine multiple stereotypes, the *Bard* expands its generative scope considerably, allowing this application to generate hundreds of thousands of evocative comparisons. The *Bard* is accessible online at:

www.educatedinsolence.com/jigsaw/

5.3. Ironic Descriptions

By using the antonym operator **-**, ironic similes can also be generated for the *P-ness* of a topic X using the Creative IR pattern "***X** is as **P** as* (**@-P ∩ ??X)**". In effect, (**@-P ∩ ??X)** finds counter-examples of *P-ness* that are comparable to X, rather than the stereotypical examples that we expect in the normative, non-ironic use of *as*-similes. However, adjectives can be ambiguous, and **-P** may not be an appropriate antonym for the intended sense of ***P***.

Since we so often seek to impress with irony, our goal is not merely to communicate an implicit negation, but to communicate an implicit negation in the most imaginative, memorable and quotable words we can muster. A vivid juxtaposition of ideas can help us to achieve this goal. We can thus use the retrieval pattern **@-***P* **@-***P* to eke out more elaborate phrases from the Google n-grams. For example, since a *wall* is typically *hard*, and a good basis for ironic descriptions of softness, the phrases "brick wall", "stone wall", "steel wall", "titanium wall", "oak wall", "granite wall" etc. will be retrieved. While phrases like "marshmallow bunny" and "snow baby" can seem decidedly odd, and thus fresh and imaginative when considered out of their original context, these can be retrieved as ironic descriptions of hardness. Likewise, the 3-gram pattern **(**^*group* ∩ **@***P*) "of " **@-***P* retrieves phrases that subvert a stereotypical grouping, so e.g. "army of dreamers", "army of civilians" and "army of irregulars" are all retrieved from the Google 3-grams as ironic vehicles for the property *disciplined*, while the 3-grams "army of cowards", "army of babies", "army of ants", "army of cripples", "army of kittens", "army of girls" and "army of worms" are retrieved for *strong*. Notice how this pattern retrieves descriptions that suggest *P-ness* right up to the last word, whereupon a final ironic reversal of meaning is delivered. Such patterns implement a more sophisticated ironic version of the adolescent strategy for generating sarcasm, where "Not!" is placed at the end of an otherwise affirmative utterance.

Generally speaking, the retrieval of elaborate readymades models Fishelov's (1992) view of poetic similes, by providing more elaborate and more vivid mental images than a single stereotype alone could do. A computer that uses a database of readymade phrases to suggest possible word/idea combinations for creative descriptions is thus tapping into the collective imagination of many different speakers at once.

5.4. A Readymade Architecture for Linguistic Creativity

Readymades do not present a complete solution to the question of creativity, human or computational, yet they do allow us to model a substantial part of that everyday creativity that alludes to a shared past. We thus see the creative information retrieval of linguistic readymades as serving a pivotal role in any comprehensive computational architecture. Creative IR will allow a system to seamlessly integrate the diverse language resource that are needed for linguistic creativity, from phonetic dictionaries to lexical ontologies to encyclopedias and large corpora. Above the creative IR layer, higher-level processes can exploit the finessed retrieval capabilities of the **@**, **-**, **^** and **?** wildcards (and arbitrary combinations thereof) to identify a pool of possible readymade solutions for a given communicative goal. In effect, we see creative IR as a flexible *middleware* for implementing systems that exhibit linguistic creativity. Fig. 1 organizes the various levels of a creative architecture – both above and below this middleware layer – into a services stack. The various functions of end-user applications will be provided by services that use the middleware layer to access word and world knowledge, to retrieve suitable linguistic readymades, to check the validity of certain linguistic sequences, and so on. Near the top of the services stack sits a module dedicated to topicality and trend tracking. It is this module that informs applications as to the topicality of certain phrases and constructions, arising from their trending uses in social media or on the Web. For example, given a commission to generate a witty book title to convey the properties *aging, mature, sexual* (e.g. erotic literature for older people), a system that employs this services stack may combine the trending title "Fifty shades of Grey" (used in many sexually-charged tweets and blogs) with the familiar title of Wilde's "The Picture of Dorian Gray" to suggest "Fifty Shades of Dorian Gray".

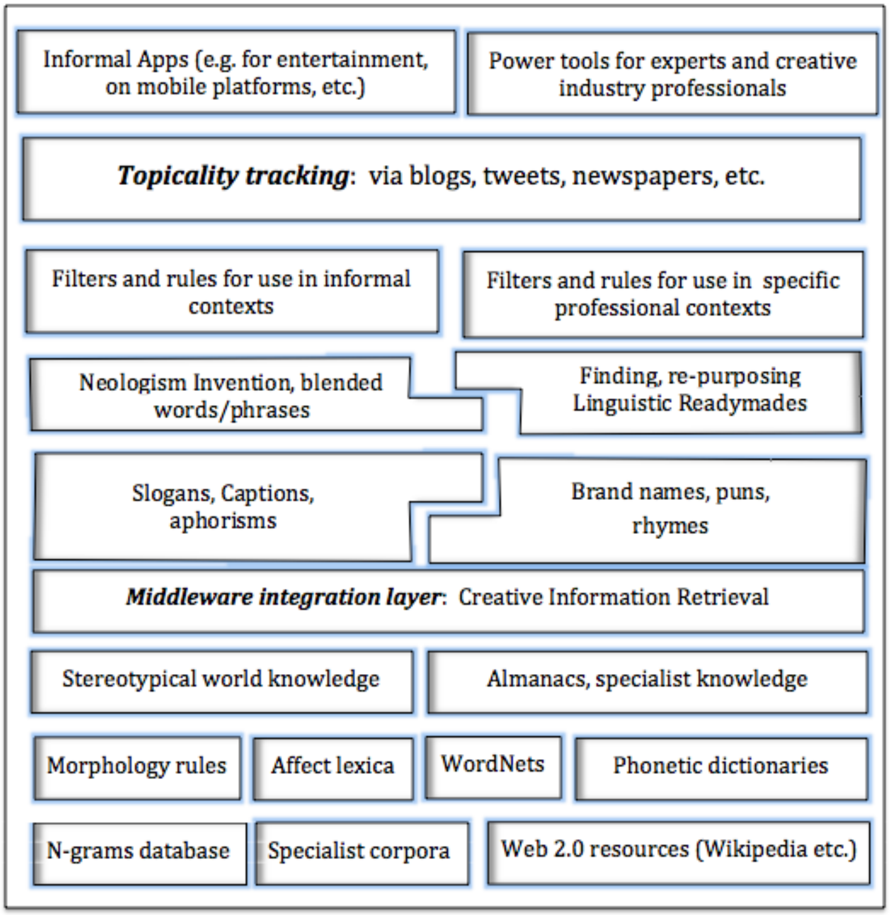

Fig. 1. A services stack for linguistic creativity that employs creative information retrieval as a middleware layer, to coordinate lower-level resources and higher-level processes.

6. Empirical Evaluation

Though ^ is the most overtly categorical of our wildcards, all four wildcards – **?**, **@**, **-** and ^ – are categorical in nature. Each has a semantic or pragmatic membership function that maps a term onto an expansion set of related members. The membership functions for specific uses of ^ are created in an ad-hoc fashion by the users that exploit it; in contrast, the membership functions for uses of **@, -** and **?** are derived automatically, via pattern-matching and corpus analysis. Nonetheless, ad-hoc categories in creative IR are often populated with the bindings produced by uses of **@** and **?** and combinations thereof. In a sense, **?X** and **@X** and their variations are themselves ad-hoc categories. But how well do they serve as categories? Are they large, but noisy? Or too small, with limited coverage? We can evaluate the effectiveness of **?** and **@**, and indirectly that of ^ too, by comparing **?** and **@** as category builders to a hand-crafted categorical gold standard like WordNet.

Other researchers have likewise used WordNet as a gold standard for categorization experiments, and we replicate here the experimental set-up of Almuhareb and Poesio (2004, 2005), which is designed to measure the effectiveness of Web-harvested conceptual descriptions. Almuhareb and Poesio chose 214 English nouns from 13 of WordNet's upper-level semantic categories, and proceeded to harvest property values for these concepts from the Web using the pattern "*a|an|the * C is|was*" (this general approach was first pioneered by Hearst, 1992). This pattern yielded a combined total of 51,045 values for all 214 nouns; these values were primarily adjectives, such as *hot, black*, etc., but noun-modifiers of C were also accepted, such as *fruit* for *cake*. They also harvested 8934 attribute nouns, such as *temperature* and *color*, using the Web query "*the * of the C is|was*". These values and attributes were used as descriptive features for a clustering algorithm to partition the 214 nouns back into their original 13 categories. Comparing these clusters with the original WordNet-based groupings, Almuhareb and Poesio reported a cluster accuracy of **71.96%** using just values like *hot* (all 51,045 values), an accuracy of **64.02%** using just attributes like *temperature* (all 8,934), and an accuracy of **85.5%** using both together (59,979 features).

How concisely and accurately does **@X** describe a noun X for purposes of categorization? Let **^AP** denote the set of 214 WordNet nouns used by Almuhareb and Poesio. Then **@^AP** denotes a set of 2,209 adjectival properties; this should be contrasted with the space of 51,045 adjectival values used by Almuhareb and Poesio. Using the same clustering algorithm over this feature set, **@X** achieves a clustering accuracy (as measured via cluster purity) of **70.2%**, compared to **71.96%** for Almuhareb and Poesio. However, when **@X** is used to harvest a further set of attribute nouns for X, via web queries of the form "*the **P** * of X*" (where P ∈ **@X**), then **@X** augmented with this additional set of attributes (like *hands* for surgeon) produces a larger space of 7,183 features. This in turn yields a cluster accuracy of **90.2%** which contrasts with Almuhareb and Poesio's **85.5%** for 59,979 features. In either case, **@X** produces comparable clustering quality to Almuhareb and Poesio, with just a small fraction of the features.

So how concisely and accurately does **?X** describe a noun X for purposes of categorization? While **@X** denotes a set of salient adjectives, **?X** denotes a set of comparable nouns. So this time, **?^AP** denotes a set of 8,300 nouns in total, to act as a feature space for the 214 nouns of Almuhareb and Poesio. Remember, the contents of each **?X**, and of **?^AP** overall, are determined entirely by the contents of the Google 3-grams; the elements of **?X** are not ranked in any way, and all are treated as equals. When the 8,300 features in **?^AP** are clustered into 13 categories, the resulting clusters have a purity of **93.4%** relative to the structure of WordNet. The pragmatic neighborhood of X, determined via **?X** , appears to be an accurate and remarkably concise proxy for the meaning of X.

It is worth noting that the Google 3-grams provide matches for 1,363,184 noun-to-noun mappings via **?** (using the patterns discussed in section 4.1), to enable **?X** to be determined for over 35,000 values of X. In total, the **?** wildcard represents a pragmatic sampling of just 0.1% of the space of possible semantic comparisons of these 35,000 nouns. So **?** captures a sweet-spot in the space of potential comparisons, only suggesting expansions that are corpus-attested by the way that ideas cluster together in ad-hoc lists. Yet, because this small sampling is pragmatically motivated, **?X** is expansive enough for most purposes, including the experiments of Almuhareb and Poesio. For other purposes, there is always **??X**.

What about adjectives? Almuhareb and Poesio's 214-word test set does not contain adjectives, and besides, WordNet does not impose a category structure on its adjectives. In any case, the role of adjectives in creative information retrieval is largely an affective one: if X is a noun, then one must have confidence that the adjectives in **@X** are consonant with our understanding of X, and if P is a property, that the adjectives in **?P** evoke much the same mood and sentiment as P. Our evaluation of **@X** and **?P** should thus be an affective one.

How well do the properties in **@X** capture our sentiments about a noun X? More empirically, how well can we estimate the pleasantness of X from the adjectives in **@X**? Whissell's (1989) dictionary of affect provides pleasantness ratings for a sizeable number of adjectives and nouns (over 8,000 words in total), allowing us to estimate the pleasantness of X as a weighted average of the pleasantness of each X_i in **@X** (where the weights are a function of the Web frequencies of the simile forms that underpin the **@** wildcard). We thus estimate the affect of all stereotype nouns for which Whissell also records a score. A two-tailed Pearson test ($p < 0.05$) shows a positive correlation of **0.5** between these estimates and the pleasantness scores assigned by Whissell. In contrast, estimates based on the pleasantness of adjectives found in corresponding WordNet glosses show a positive correlation of just **0.278**. Further, Veale (2012b) shows that when one determines a crude positive/negative polarity for a stereotype based on a weighted analysis of its properties, the correct polarity is assigned 96% of the time.

Finally, how well do the properties in **?P** capture our sentiments about the adjective P? Remember, we hypothesize that the adjectives in **?P** are highly suggestive of P, and vice versa. *Aristotle* and the *Jigsaw Bard* each rely on **?P** to suggest adjectives that evoke an unstated property in a metaphor or simile, or to suggest coherent blends of properties. When we estimate the pleasantness of each adjective P in Whissell's dictionary using a weighted average of the pleasantness of adjectives in **?P** (again using Web frequencies as weights), a two-tailed Pearson test ($p < 0.05$) shows a positive correlation of **0.7** between estimates and actual scores. It seems that **?P** does a good job of capturing the mood of P.

7. Applications to Computational Creativity

Productivity-enhancing software constitutes a major component of the commercial software market. Ranging from *Microsoft Office* to *Final Draft* (a screenwriting tool for film professionals) to *Adobe PhotoShop*, these tools make it easier for creative people to express themselves and to achieve their creative goals. Despite the market penetration of tools like *Microsoft Word*, it is fair to say that tools for creatively constructing and manipulating images are far more successful, and much more richly-featured, than the equivalent tools for manipulating language and texts. When Samsung entered the market for "creativity"-related tools and services with its Galaxy *Note* smartphone, the creative software on which this product-line was predicated was entirely image-based: the *Note* gives users a stylus and some feature-rich image-capture and manipulation software so they can be spontaneously creative in different social contexts, serendipitously grabbing and composing readymade images in situ to suit their needs. Tellingly, the *Note* and equivalent devices still possess very limited language and text manipulation capabilities.

There are good reasons for the language/image gap: One can manipulate an image with a wide range of continuous local transformations at a pixel level, and still obtain a valid image as output. Language, however, is a complex system of cultural symbols and rules – some explicit, many tacit – and simple manipulation of texts often yields a broken text (un-grammatical, unintelligible, meaningless) as output. Productivity tools for writers have made little headway beyond digitalized versions of *old reliables* such as the thesaurus, the style guide and the grammar book. These tools suggest simple word substitutions or word-order changes (e.g. to use the passive voice in formal documents). These tools cannot help us to be more creative, by suggesting changes that would make our texts wittier, or more persuasive, or more memorable. To be sure, software like *PhotoShop* is equally uncreative, yet it offers such a wide range of features that it encourages creative experimentation by its users. Current productivity software for language/text creation offers a much impoverished suite of features by comparison.

As presented here, *Creative Information Retrieval* is an architecture that is designed to support new and robust tools for linguistic creativity. The cornerstone of the architecture is the essential reusability of linguistic fragments in diverse contexts that could not have been imagined when the fragments were first composed. Creative IR is thus an attempt to efficiently systematize the discovery of readymades on a large scale. In the typical artistic scenario, creators encounter objects in the course of their daily lives that resonate with additional meanings and new artistic possibilities. Perhaps a creator is already in the *incubation* stage of creative problem solving (see Wallas, 1926) and is thus unconsciously attuned to the properties of any object that may serendipitously present itself as a possible solution. In Creative IR, however, this order of readymade invention, namely *object encounter* → *meaning perception* → *creative intent*, is reversed: instead, a user begins with a creative intent, expresses the desired meaning in formal, schematic terms, and then retrieves a wide range of readymade candidates from a large corpus. The candidates retrieved in this way are not finished products but the raw material from which a professional creator, wielding selectivity, personal aesthetic judgment and – in some cases – additional processes of editing, tweaking and elaboration, shapes a polished readymade result. This reversed process can hardly be considered cognitively plausible, but neither is Creative IR intended to serve as a cognitive model of human creativity. Nonetheless, in providing an integrated design process that serves to systematize human creativity on a commercial scale, this architecture can serve human professionals as a creative tool while also serving as a foundational component of future models of autonomous *machine* creativity.

8. Concluding Thoughts

Samuel Goldwyn, the co-founder of MGM studios, famously summed up Hollywood's attitude to creativity with the line "Let's have some new clichés". On the face of it, this is both a shocking indictment of the Hollywood "factory" and yet another of Goldwyn's many memorable misstatements (such as "include me *out*!" and "the atom bomb, it's *dynamite*!"): after all, it's hard to think of clichés as *new*, or as something that can be invented on demand. Yet, on closer analysis, one can find real insight about creativity in Goldwyn's remark. Clichés are considered anathema to the creative process because they represent everything that is conventional and jaded about the status quo. However, clichés become tired thru overwork, and are overworked precisely because they prove themselves so useful in so many different contexts. Few writers set out to create a new cliché, but most would like their efforts to become as much a part of the fabric of our linguistic culture as the most tenacious of clichés.

One productive form of *new cliché* is the humorously pithy comparison, as in "as durable as a chocolate teapot" or "as useful as a screen door on a submarine". Speakers recognize memorable comparisons when they hear them, and re-use them as eagerly as one retells a favorite joke. The most frequently reused comparisons can, in this way, acquire the clichéd status of a proverbial simile. When the folklorist Archer Taylor collected his corpus of proverbial similes in 1954, he observed not just a wide variety of humorous comparisons in American speech, but a wide variety of humorous forms for the same descriptive qualities, such as "durable" and "useful". Speakers are clearly drawn to popular comparisons of proven value, but are equally fond of coining their own, in the hope that their witty new descriptions will be widely reused by others in turn. This constant churn of re-invention keeps our language fresh, and ensures that comparisons retain their ability to challenge and to entertain, even as others – such as "crazy like a fox!" or "as clear as mud!" – acquire an idiomatic status which makes them effortlessly understood.

George Orwell echoed the conventional wisdom on cliché when he encouraged readers, in the strongest possible terms, to avoid any formulation or trope that they were used to seeing in print. This contradictory attitude to cliché, characterized by sharp insight in recognizing cliché but a plodding banality in describing what to do about it, lead the critic William Empson to memorably describe Orwell as "the *eagle eye with the flat feet"* (quoted in Ricks, 1980). Empson's put-down offers a marvelous repudiation of the whole anti-cliché movement, since it is itself a creative turn of phrase built from a blend of two rather hackneyed, even fly-blown, clichés: the cliché that an observant person has "the eyes of an eagle", and the cliché that a plodding thinker is "flat-footed". Likewise, Ricks (1995) notes with some humour that Orwell's listing of his least favorite clichés – "*jackboot*, *Achilles' heel, hotbed, melting pot, acid test, veritable inferno"* – has an accidental but vital poetry about it, as if clichés long to be grouped into creative combinations. Did Orwell succumb to a subconscious desire to use cliché creatively, Ricks wonders, since he so evocatively followed *jackboot* with *Achilles' heel* (what better to protect an exposed heel?), *hotbed* by *melting pot* (as if the same metaphoric heat is used to warm both), and *melting pot* by *acid test* and *veritable inferno* (as if this metaphorical heat offers a searing test of truth)?

Artist-*provocateur* Marcel Duchamp showed a more nuanced understanding of the creative potential of language when he encouraged speakers to constructively engage with, rather than outright disavow, the popular forms of a culture. Clichés and other tired forms may represent fossilized norms, but when used in the right way, with an open mind and a modicum of creative intent, they are readymades that can be given new purpose and new meaning in a new context, just as Duchamp's choice of a humble urinal underpinned the provocative work of art we know remember as *Fountain*. Between unthinking acceptance and high-minded disavowal of ossified norms lies a middle ground that is more creatively productive – the knowing exploitation and subversion of these norms. In this vein, Fishelov (1992) has shown that poetic similes represent a conscious deviation from the norms of non-poetic comparison. His analysis shows that poetic similes are longer and more elaborate, and are more likely to be figurative and to flirt with incongruity. Creative descriptions do not necessarily use words that are longer, or rarer, or fancier, but use many of the same cultural building blocks as non-creative similes. Armed with a rich vocabulary of building blocks, a great many readymade phrases can be retrieved on demand from the Google n-grams

– from the evocative "chocolate martini" (evoking smoothness) to the seemingly incongruous "robot fish" (evoking emotional coldness) – that can be used to suggest an equally wide range of properties.

The generativity of this retrieval mechanism for linguistic readymades is both scalable and robust. However, any creativity we may attribute to the use of these readymades comes not from the phrases themselves – they are *étant donnés* (already made), after all (Taylor, 2009) – but from the recognition of the subtle and often complex properties they evoke. Readymades shift the creative emphasis from construction to *selection*, from generation to *retrieval*, and computers – which excel at targeted retrieval – can creatively benefit from this change of emphasis. Duchamp memorably defended *Fountain* thusly: "whether [the artist] made the fountain or not has no importance. He CHOSE it. He took an ordinary article of life, placed it so its useful significance disappeared under the new title and point of view – created a new thought for that object". The retrieval approach to creativity presented here exploits a sweet-spot in our understanding of linguistic creativity just as much as Duchamp and his followers exploited a sweet-spot in the public's understanding of art and how it is practiced. Yet as presented here, this approach is merely a starting point on our exploitation of linguistic readymades, and not an end in itself. By harvesting more complex syntactic structures, and using more sophisticated techniques for analyzing the figurative potential of these phrases, this approach may gradually approach the levels of creativity exhibited by Duchamp and the levels of poeticism discussed by Fishelov.

Acknowledgments

This research was partly supported by the WCU (World Class University) program under the National Research Foundation of Korea (Ministry of Education, Science and Technology of Korea, Project no. R31-30007) and partly funded by Science Foundation Ireland via the *Centre for Next Generation Localization* (CNGL).

References

Almuhareb, A. and Poesio, M. (2004). Attribute-Based and Value-Based Clustering: An Evaluation. In *Proc. of EMNLP 2004*. Barcelona.

Almuhareb, A. and Poesio, M. (2005). Concept Learning and Categorization from the Web. In *Proceedings of the 27th Annual meeting of the Cognitive Science Society*.

Barsalou, L. W. (1983). Ad hoc categories. *Memory and Cognition*, 11:211–227.

Boden, M. (1994). Creativity: A Framework for Research, Behavioural and Brain Sciences 17(3):558-568.

Brants, T. and Franz, A. (2006). *Web 1T 5-gram Version 1*. Linguistic Data Consortium.

de Bono, E. (1970). *Lateral thinking: creativity step by step*. New York: Harper & Row.

de Bono, E. (1971). *Lateral thinking for management: a handbook for creativity*. New York: McGraw Hill.

Fellbaum, C. (ed.) (2008). *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge.

Fishelov, D. 1992. Poetic and Non-Poetic Simile: Structure, Semantics, Rhetoric. *Poetics Today*, 14(1), 1-23.

Hanks, P. (2005). Similes and Sets: The English Preposition 'like'. In: Blatná, R. and Petkevic, V. (Eds.), *Languages and Linguistics: Festschrift for Fr. Cermak*. Charles University, Prague.

Hofstadter, D. R. (1995). *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. Basic Books, NY.

Guilford, J.P. (1950). Creativity, *American Psychologist* 5(9):444–454.

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 14th Int. Conf. on Computational Linguistics*, pp 539–545.

MacCormac, E. R. (1985). *A Cognitive Theory of Metaphor*. MIT Press.

Mihalcea, R. (2002). The Semantic Wildcard. In *Proceedings of the LREC Workshop on Creating and Using Semantics for Information Retrieval and Filtering*. Canary Islands, Spain, May 2002.

Moon, R. (2008). Conventionalized as-similes in English: A problem case. *International Journal of Corpus Linguistics* 13(1), 3-37.

Norrick, N. (1986). Stock Similes. *Journal of Literary Semantics* XV(1), 39-52.

Orwell, G. (1946). Politics and The English Language. *Horizon* **13**(76), April issue.

Ricks, C. B. (1980). Clichés. In Leonard Michaels and Christopher B. Ricks (Eds.), *The State of the Language*. University of California Press.

Ricks, C. B. (1995). *The force of poetry*. Oxford University Press.

Sternberg, R. J. and Lubart, T. I. (1995). *Defying the crowd: Cultivating creativity in a culture of conformity*. Free Press, New York.

Taylor, A. (1954). Proverbial Comparisons and Similes from California. *Folklore Studies* 3. Berkeley: University of California Press.

Taylor, M. R. (2009). *Marcel Duchamp: Étant donnés* (Philadelphia Museum of Art). Yale University Press.

Veale, T. and Hao, Y. (2007). Making Lexical Ontologies Functional and Context-Sensitive. *In proc. of the 46th Ann. Meeting of the Assoc. of Computational Linguistics*.

Veale, T. and Li, G. (2013). Creating Similarity: Lateral Thinking for Vertical Similarity Judgments. In Proceedings of ACL'2013, the 51st *Annual Meeting of the Association for Computational Linguistics*, Sofia, Bulgaria.

Veale, T. (2011). Creative Language Retrieval: A Robust Hybrid of Information Retrieval and Linguistic Creativity. In *Proceedings of ACL'2011, the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, Oregon, USA.

Veale, T. (2012a). Exploding the Creativity Myth: The computational foundations of linguistic creativity. *London: Bloomsbury Academic*.

Veale, T. (2012b). Seeing the Best and Worst of Everything on the Web with a Two-level, Feature-rich Affect Lexicon. In Proceedings of WWW'2012, the 21st World-Wide-Web conference, Lyon France.

Veale, T. (2013). Humorous Similes. *Humor: The International Journal of Humor Research, **21**(1):3-22*.

Wallas, G. (1926). *Art of Thought*. London, UK: Watts and Co.

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

Author Biography

Dr. Tony Veale is a visiting professor in Web Science and Technology at KAIST, the Korean Advanced Institute of Science and Technology, Daejeon, South Korea, where he is funded by the government's World Class University program. His home institution is University College Dublin in Ireland, where his research interests are computational creativity and figurative language processing. He leads the EU-funded PROSECCO coordination action to promote the scientific exploration of computational creativity.