

A Service-Oriented Architecture for Computational Creativity

Tony Veale*

Web Science and Technology Division, Korean Advanced Institute of Science and Technology, Daejeon, Korea
tony.veale@gmail.com

Abstract

Creativity is a long cherished and widely studied aspect of human behavior that allows us to re-invent the familiar, and to imagine the new. Computational creativity (CC) is a newly burgeoning area of creativity research that brings together academics and practitioners from diverse disciplines, genres and modalities, to explore the potential of computers to be autonomously creative, or to collaborate as co-creators with people. We describe here an architecture for creative Web services that will act as a *force magnifier* for CC, both for academic research, and for the effective deployment of real CC applications in industry. For researchers, this service-oriented architecture supports the pooling of technologies in a robust interoperable framework, in which CC models are conceived, developed and migrated from lab settings to an industrial strength platform. Industry developers, for their part, will be able to exploit novel results of CC research in a robust, low-risk form, without having to re-implement algorithms from a quickly moving field. We illustrate the architecture with the first of a growing set of creative Web services that provide robust figurative language processing on demand.

Category: Smart and intelligent computing

Keywords: Creativity; Web services; Service-oriented architecture; Metaphor; Lateral thinking; Poetry

I. INTRODUCTION

The most popular software tools—in both social and commercial terms—are those that allow users to express themselves creatively, either by facilitating the construction of creative output (documents, images, videos, Web pages, etc.), or by facilitating the sharing of this output (via blogs, social networks, etc.). However, no matter how powerful or how rich in features the software happens to be, the user remains the “creator”, while the computer/software remains the “tool”. This is true whether one is using consumer-grade software like *Microsoft Word*, or professional-grade software like *Adobe Photoshop*. While these tools facilitate the creation of new digital artifacts, the creativity still resides entirely within the user.

Computational creativity (CC) is an emerging field

that studies the potential of computers to be more than feature-rich tools, and to be autonomous creators and co-creators in their own right [1]. In a CC system, the creative impetus comes from the machine, not the user, though in a hybrid CC system, a joint impetus may come from both together. As a field, CC draws on elements of artificial intelligence, philosophy, cognitive science, psychology and anthropology, and asks: What does it mean to be “creative”? Does creativity reside in the individual, in the process, in the product, or in a combination of all three together? How does creativity exploit norms, and subvert expectations? What cognitive paradigms—from search in a conceptual space, to conceptual blending—offer the most usable and explanatory theories of creativity?

Each of these questions is just as valid to the study of human creativity, as it is to the study of machine creativ-

Open Access <http://dx.doi.org/10.5626/JCSE.2013.7.3.159>

<http://jcse.kiise.org>

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/3.0/>) which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

Received 7 June 2013, Accepted 1 July 2013

*Corresponding Author

ity (e.g., [2-11]) What makes CC different is that it adopts an explicitly algorithmic perspective on creativity, and seeks to tie down the study of creative behavior to specific processes, algorithms and knowledge structures (e.g., [3, 12-34]). The goal of CC is not just to theorize about the generative capabilities of humans and machines, but to build working systems that embody these theoretical insights in engineering reality. As such, CC is an engineering discipline *and* an experimental science, in which progress is made by turning insights into applications that can be experimentally tested and evaluated. The purpose of these applications is to create novel artifacts—stories, poems, metaphors, theorems, riddles, jokes, paintings, scientific hypotheses, musical compositions, games, etc.—in which a large measure of the perceived creativity is credited directly to the machine. We believe that the future of intelligent computers lies in transforming our computers from passive tools into active co-creators, and that CC is the field that can make this transformation a reality.

II. CREATIVITY AS A SERVICE

Creativity is an elusive phenomenon that organizations put significant effort and resources into fostering, rewarding, retaining, and reproducing on demand. The systematic harnessing of creativity is complicated by the complex and definition-defying nature of the phenomenon, and the realization that it depends crucially on many different social, cultural and contextual factors [2, 4-11]. For these reasons, companies often out-source their creative needs to external agencies with a track record in the exploration, composition and framing of innovative solutions. Such agencies are not so much *problem solvers* as *option providers*, leaving the ultimate responsibility for choosing among this diversity of new options to the client. To out-source in this way is not to abdicate creative responsibility, but to broaden the range of choices one can choose from.

Complex software systems share many similarities with large organizations. Each must be well defined, operate in a predictable fashion, and facilitate an efficient and orderly flow of information. But like large organizations, software systems should continuously engage their users, and react with grace and agility when faced with unexpected situations. So imagine if systems could out-source their creative needs to an external service with a track record in CC. This service would not be a cadre of creative workers, but a suite of interoperable tools that provide, on demand, the processes and representations that are keys to creative thinking. Software systems, like organizations, could thus maintain their well-tested structures and disciplined information-flows, while appealing to outside creative services whenever they need to diversify the range of possibilities (both in form and content) that are available to choose from.

Whether a service or set of services is being offered by a company or a software system, it pays for either to follow the principles of a well-designed service-oriented architecture (SOA). Erl [35] defines the essence of a SOA platform as “an architectural model that aims to enhance the efficiency, agility, and productivity of an enterprise by positioning services as the primary means through which solution logic is represented.” Erl [35] further notes that well-designed services should be discoverable, autonomous and widely reusable, and should be flexible enough to compose in groups, while remaining loosely coupled to others. Services should also maintain minimal state information and use abstraction to hide the complexity of their inner workings and data.

CC is not a field that hinges on any one algorithm, process or data structure [1]. Rather, CC is a field defined by its goals—the selective generation of novel and task-appropriate artifacts—more than by any particular means of representation, generation or selection. A CC system may use any of a wide range of approaches, structures and processes to achieve its generative goals. A comprehensive SOA for CC is a SOA with many diverse and competing services, operating at different levels of specificity and scale. Services for music generation, say, will hide different complexities, and rely on different information and knowledge sources, than services for language or image generation; yet such services must ultimately work together, to allow users to create rich multimodal constructs. Even a single modality, such as language, will require a wide range of services—each perhaps operating at different levels of form and mean-

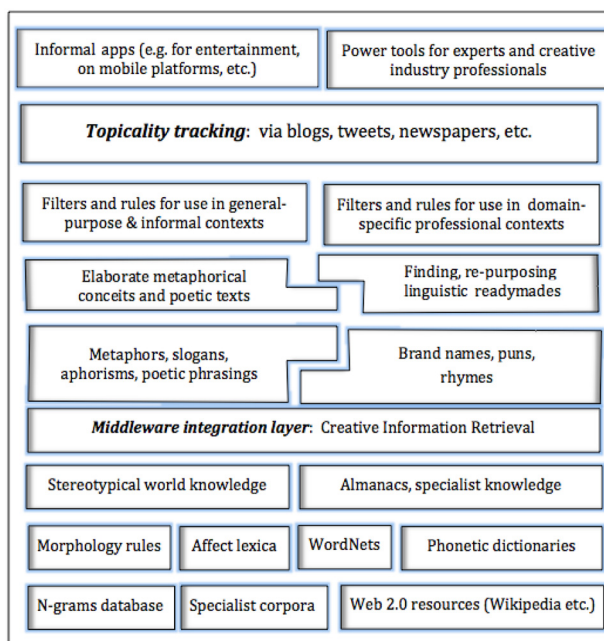


Fig. 1. A layer-cake of services for creative language processing, on a foundation of core information sources.

ing—to provide a generative capacity for metaphors, poems, jokes, stories, and persuasive descriptions. Fig. 1 shows how the necessary information sources, processes and client-oriented services can be stacked in a way that maximizes the reuse, abstraction, data-hiding, and composability demanded by SOA, while minimizing coupling between services and statefulness within services.

The layered architecture of Fig. 1 employs *Creative Information Retrieval* [28] to integrate a diverse set of resources into a single middleware component. This middleware hides the complexity of language resources that vary in organization (from structured to unstructured), scale (single files to massive corpora) and content (raw data, tagged information, or conceptual knowledge), while providing an expressive means of exploiting these resources to the SOA services perched above it.

This vision of a service-oriented Web architecture for CC proposes three kinds of service: *discovery & insight services*; *idea composition services*; and *framing services*. Each service may rely on different sources of knowledge, but each should use interoperable data structures, and so can call upon other services during its operation. The overall architecture is theory-neutral, yet will provide a rich ecology of theory-informed CC services that can be composed in any way that suits a client system's needs.

A. Discovery and Insight Services

Documents and domains are containers of knowledge, but this knowledge is more than a simple bag of true-or-false propositions. Rather, knowledge is textured, so that some elements are strongly explicit or foregrounded, while many others remain implicit, latent or presumed, in the conceptual background [30]. Knowledge that resides at the boundaries of two or more domains may only come to the fore—where it can appear surprising and insightful—when representations of these domains are studied in juxtaposition [9]. Discovery services will mine diverse corpora with bisociative tools (e.g., [19, 22]), to acquire emergent insights and novel perspectives on everyday concepts.

B. Composition Services

Creativity often arises from frame conflict, when one concept is incongruously viewed through the lens of another, very different idea [20, 21, 23]. The key to the fruitful exploitation of frame conflict is two-fold: one must first choose which concepts to juxtapose, and then formulate a resonant form for the resulting content. A SOA architecture for CC will provide services for suggesting, elaborating and comprehending conceptual metaphors [25], analogies [24], and blends [22], as well as services for accessing the large store of common-sense knowledge that these composition services will crucially rely upon.

C. Framing Services

The conceptual conceit that underpins a creative act must be packaged for an audience in a concise, easily appreciable and memorable form, such as a linguistic metaphor, simile, joke, name, slogan, short story, poem, picture, piece of music, or a mixture of these forms. Each of these forms may frame the same underlying conceit in very different ways to achieve competing goals (e.g., catchiness, brevity, resonance, wit) for diverse audiences [36]. A SOA architecture for CC must provide services for framing the outputs of the discovery and composition services in a variety of parameterized forms, from affective analogies to metaphors to poems to stories to pictures to music.

III. CREATIVE LANGUAGE RETRIEVAL

Though linguistic creativity primarily involves the generation of novel texts, or texts that achieve their communicative ends in novel or non-obvious ways, generation need not be computationally achieved from first principles, using explicit grammars and other formal machinery. Veale [28] argues that much linguistic creativity arises from the purposeful reuse of existing texts or language fragments. So an innovative text is not necessarily one that uses rare or fanciful words, but one that finds fresh and surprising uses for familiar forms [36]. Creativity resides in the gap between what is familiar, and what is obvious, to produce an *optimal innovation* [37] that knowingly plays with convention. The ability to retrieve familiar or existing phrases that can be aptly reused in resonant new ways is just as important then, if not more so, as the ability to generate completely novel phrasings ab initio. Given a large repository of common language fragments, such as the *Google n-grams* [38], one needs an especially expressive query language to retrieve “readymade” phrases, based not on their form (which cannot be known *before* retrieval), but on their meanings, connotations or resonances. *Creative Information Retrieval* [28] provides one such query language to support *non-literal* retrieval.

Conventional text search performs a *literal* matching of keywords to find relevant texts; texts are retrieved only if they are literally similar to the query, or an automatic expansion thereof. Literal matching works poorly for more creative, *non-literal* retrieval tasks, where the goal is to retrieve texts that convey a similar meaning to that of the query, while perhaps using a very different language, as in e.g., the retrieval of potential metaphors for a given topic. A non-literal query language must thus understand the meaning of keywords, in a way that a literal query language need not. Suppose one's goal is to find novel metaphors to highlight the *coldness* of a topic. We formu-

late a non-literal query with the keyword *cold*, and assume the system knows which nouns denote the kind of concepts that are stereotypically *cold*, such as *snow*, *ice*, *wind*, *January*, *winter* or *fridge*, and which adjectives denote properties that imply *coldness*, such as *wet*, *bitter* or *icy*. Veale [28] describes a retrieval system for the Google n-grams, in which the non-literal query **?cold @cold** retrieves any Google 2-gram phrase—such as *wet fish*, *rainy January*, *icy tomb*, *heartless robot*—in which the first word is an adjective reinforcing the property *cold* (e.g., *metallic*, *hard*, etc.), and the second word is a noun denoting a stereotype of *coldness*, such as *corpse*, or *icicle*. More generally, the query term **?ADJ** matches any adjective denoting a property that is known to evoke or reinforce **ADJ-ness**, while **@ADJ** matches any noun that denotes a stereotype with the property **ADJ**.

Conventional query operators are content-neutral, but non-literal operators, such as **?** and **@**, demand a rich model of stereotypical world knowledge to resolve possible matches. Such knowledge can be acquired on a large scale from the Web, by looking for instances of language use in which speakers expose their tacit expectations of the world, such as that *ice* is *cold*, *knives* are *sharp*, or *rockets* are *fast*. Veale and Hao [25] and Veale [28] show that Web similes open a revealing window onto our most useful stereotypes, arguing that similes are so revealing precisely because they themselves are important vectors for the transmission of cultural knowledge through language. The common simile “as *hot* as an *oven*” thus tells us that ovens are expected to be hot, while “as *sharp* as a *scalpel*” and “as *precise* as a *surgeon*” reveal our cultural expectations of surgeons and their tools. By harvesting large amounts of the “as *P* as *N*” pattern from the Web via a search engine, our computers can learn the mappings **@:P→N** and **@:N→P**. By also seeking out many instances of the co-description pattern “as *P_α* and *P_β* as” from the Web, a computer can likewise learn the mapping **?:P→P**. Veale and Hao [25] and Veale [28] provide empirical evidence that shows these mappings offer wide coverage of many common nouns, and also provide highly informative features with which to cluster nouns into semantic categories.

Veale [28] also shows that a diverse range of creative language applications, for the generation of metaphors, similes, analogies, and ironic descriptions, can be implemented in a lightweight manner, by formulating task-appropriate non-literal queries over Google n-grams. *Creative Information Retrieval* thus acts as an expressive middleware for building lightweight creative services in an agile manner, while hiding the complexity of the underlying data.

IV. LATERAL THINKING AS A SERVICE

Just as *observing* is more than just *seeing*, *comparing*

is much more than mere *matching*. It takes understanding and inventiveness to discern a useful basis for judging two ideas as similar in a particular context, especially when our perspective is shaped by an act of linguistic creativity, such as metaphor, simile or analogy. Lexical resources such as WordNet [39] offer a convenient hierarchical means for converging on a common ground for comparison, but offer little support for the divergent thinking that is needed to creatively view one concept as another. Fortunately, the Web can be used to harvest many divergent views for many familiar ideas. These lateral views complement the vertical views of WordNet, and support a Web service for idea exploration via lateral categorization.

Any measure that models similarity as an objective function of a conventional worldview employs a *convergent* thought process. Using WordNet, for instance, a similarity measure can vertically converge on a common superordinate category of both inputs, and generate a single numeric result based on their distance to, and the information content of, this common generalization [40]. To find the most conventional ways of viewing a concept, one simply ascends a narrowing concept hierarchy, using a process de Bono [41] calls *vertical thinking*. But to find novel, non-obvious and useful ways of looking at a lexical concept, one must use what Guilford [8] calls *divergent thinking*, and what de Bono [41] calls *lateral thinking*. de Bono [41] argues that vertical thinking is selective, while lateral thinking is generative. Whereas vertical thinking aims to select the “right” way or a single “best” way of looking at an object or situation, lateral thinking focuses on the generation of alternatives to the status quo. So that they can be as useful for creative tasks as they are for conventional tasks, we need to re-imagine our computational measures of similarity as generative services that are expansive, rather than reductive, divergent *and* convergent, and lateral *and* vertical. These processes should be able to cut across category boundaries, to simultaneously place a concept in many different categories at once, to see them in many diverse ways.

Hao and Veale [27] re-express each “as *P* as *C*” simile gathered in [25] in the form “*P* * *such as C*” (where * is a wildcard), and harvest all attested uses of this new form from the Web. Since each hit will also yield a value for a hypernym *H* of *C* via the wildcard *, each match will provide a fine-grained categorization *P-H* for *C*. For instance, given the simile “as *fizzy* as *cola*”, the Web harvester generates the new form “*fizzy* * *such as cola*”, and goes to the Web to find the fine-grained perspectives *fizzy-drink*, *fizzy-mixer*, and *fizzy-beverage* for *cola*.

Once *C* is seen to be an exemplary member of the category *P-H*, such as *cola* for *fizzy-drink*, a targeted Web search is used to find other members of *P-H*, via the anchored query “*P H such as * and C*”. For example, “*fizzy drinks such as * and cola*” will retrieve Web texts in which * is matched to *soda* and *lemonade*. Each new

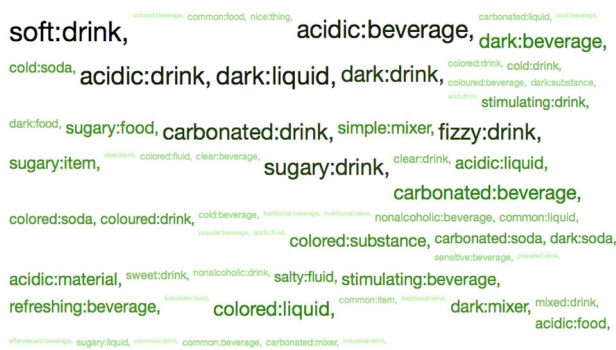


Fig. 2. Fine-grained perspectives for *cola* bootstrapped from the Web, using common similes as a starting-point.

exemplar can then be used to instantiate a further query, as in “*fizzy drinks such as * and soda*”, to retrieve other members of *P-H*, such as *champagne* and *root beer*. This bootstrapping process runs in successive cycles, using doubly anchored patterns that—following Kozareva et al. [42] and Veale et al. [26]—explicitly mention both the category to be populated (*P-H*), and a recently acquired exemplar of this category (*C*).

Five successive cycles of bootstrapping are performed, using the 12,000+ Web similes of Veale and Hao [25] as a starting point. Consider *cola*: after 1 cycle, the harvester acquires 14 new categories, such as *effervescent-beverage* and *sweet-beverage*. After 2 cycles it acquires 43 categories; after 3 cycles, 72; after 4 cycles, 93; and after 5 cycles, it acquires 102 diverse, fine-grained perspectives on *cola*, such as *stimulating-drink* and *corrosive-substance*. Fig. 2 presents a phrase cloud of the most frequently harvested perspectives on *cola* from the Web.

V. METAPHOR AS A SERVICE

Metaphor Magnet is a Web service for idea composition that understands and generates affective metaphors on demand. This service’s representations are harvested in bulk from the Web, and its linguistic outputs are evaluated and ranked according to the amount of evidence that can be found for them in a corpus of attested language use (the Google n-grams of [38]). *Metaphor Magnet* also employs an affective lexicon built from the same simile-derived knowledgebase of stereotypes that underpins *Creative Information Retrieval*. This lexicon allows the service to appreciate the affective subtleties of properties like *warm* versus *cold*, or *happy* versus *sad*, and of nouns like *hero* versus *villain*, *cult* versus *religion*, and *war* versus *warrior*. In turn, this affective knowledge allows *Metaphor Magnet* to predict the emotional resonances of a given candidate metaphor, and to decide whether these are aligned with the client’s stated communicative goals.

Metaphor Magnet is designed to be a lightweight Web service that provides both HTML output (for humans

menacing:bully

perilous:battle twisted:devil

menacing:tough menacing:lackey

isolated:exception bewildered:victim

overbearing:king bloated:dinosaur

menacing:threat

troubling:question menacing:evil

threatening:enemy arrogant:master

twisted:joke domineering:leader threatening:competitor

menacing:criminal

threatening:problem

menacing:giant

twisted:root conservative:corporation domineering:superior

Fig. 3. Screenshot of a phrase cloud for the perspective on *Microsoft* cast by “*Google is as -powerful as Microsoft*”.

interacting via browsers), and XML (for remote client applications). The service accepts affective inputs, such as *Google is like -Microsoft*, *life is a +game*, *Steve Jobs is Tony Stark*, or even *Rasputin is Karl Rove*, and generates a ranked list of apt metaphors in response. A minus sign indicates a negative spin is desired on the given concept, while a plus sign requests a positive spin. When provided with the input “*Apple is a -religion*”, for example, the service returns a list of apt religion metaphors that are appropriate to the target *Apple*, and which show religious topics in a negative light. It gathers negative metaphor candidates from the Google n-grams via *Creative Information Retrieval*, by asking for all copula 4-grams of the form *religion is a Y* where *Y* denotes a negative stereotype—such as *cult* or *virus*—with more negative than positive properties overall.

Each candidate is then considered in juxtaposition with the given target concept, to e.g., consider how *Apple might be a cult*, or *Apple might be a virus*. *Creative Information Retrieval* is again employed for each juxtaposition, to ask whether n-gram evidence can be marshaled to show e.g., Apple behaving like a cult, or exhibiting cult-like properties and associations. Finally, if so desired by the client, this evidence is used to generate a series of phrasal IR queries (such as “Apple’s mantra” and “Apple’s followers”) that are used to retrieve relevant evidence from the Web via the *Google* application programming interface. In effect, the system allows users to interface with a search engine like *Google*, using metaphor and other affective language forms.

Metaphor Magnet’s interpretation of the affective simile “*Google is as -powerful as Microsoft*” highlights a range



Fig. 4. Screenshot of a phrase cloud for the perspective on Google cast by "Google is as -powerful as Microsoft".

of affective viewpoints on the source concept, *Microsoft*, and projects a number of negative viewpoints (note the minus in *-powerful*) onto the target, *Google*. The *Metaphor Magnet* Web application displays both selections as phrase clouds, in which each hyperlinked phrase—a combination of a stereotype and a projected quality, such as "menacing giant"—is clickable, to yield linguistic evidence for the selection, and corresponding Web-search results (obtained via *Google*). The phrase cloud representing *Microsoft* in this simile is shown in Fig. 3, while the cloud for qualities projected onto *Google* is shown in the screenshot of Fig. 4.

VI. POETRY AS A SERVICE

An important advantage to the delivery of CC via a suite of Web services is that CC services can be interoperably paired with other services, which can, in turn, be dynamically added and opportunistically discovered by an inquisitive client application. *Metaphor Magnet*, for instance, is interoperable with a companion framing service for poetry generation, in which the space of possible interpretations of a given metaphor is crystallized into a single poetic form. In this service, poetry is viewed as a means of reducing information overload, of summarizing a rich metaphor, whose interpretation entails a rich space of affective possibilities. A poem can thus serve as a visualization device, offering a concise alternative to the clouds of Figs. 3 and 4.

Given the metaphor *Marriage is a Prison*, *Metaphor Magnet*'s companion service generates the following poem as a distillation of the space of feelings that arise from the metaphor's interpretation:

The legalized regime of this marriage

*My marriage is a tight prison
 The most unitary federation scarcely organizes so much
 Intimidate me with the official regulation of your prison
 Let your close confines excite me
 O Marriage, disgust me with your undesirable security*

Each time the poetry service dips into the space of interpretations of the metaphor, a new poem is generated. One can sample the space at will, hopping from one interpretation to the next, or from one poem to another. Here is an alternate rendition by the poetry service of the same metaphorical conceit:

The official slavery of this marriage

*My marriage is a legitimate prison
 No collective is more unitary, or organizes so much
 Intimidate me with the official regulation of your prison
 Let your sexual degradation charm me
 O Marriage, depress me with your dreary consecration*

In the context of Fig. 3, which samples the space of metaphors that negatively describe Microsoft's perceived misuse of power, consider the following poetic framing, which distills the assertion, *Microsoft is a Monopoly*, into a suitably aggressive ode:

No Monopoly Is More Ruthless

*Intimidate me with your imposing hegemony
 No crime family is more badly organized,
 or controls more ruthlessly
 Haunt me with your centralized organization
 Let your privileged security support me
 O Microsoft, oppress me with your corrupt reign*

This poetic companion service to *Metaphor Magnet* is a recent addition to the CC architecture, one whose rhetorical workings are beyond the scope of the current paper. In essence, the service combines a property-rich behavioral model of stereotypes with a linguistic understanding of rhetorical tropes, finding apt instances of the latter to frame poetic uses of the former. The approach is described in detail in [32]. For a related approach to poetry that also uses *Metaphor Magnet*'s inventory of similes, stereotypes and affective metaphors, readers are directed to [13].

VII. EVALUATION

An interactive demonstration of these exemplar CC services for divergent categorization (discovery), metaphor interpretation and generation (composition), and poetry

generation (framing) can each be accessed via Web applications hosted at the following URLs:

1) Metaphor Magnet (and poetry generation):

<http://boundinanutshell.com/metaphor-magnet>

2) Thesaurus Rex:

<http://boundinanutshell.com/therex2>

Thesaurus Rex allows users (via a browser) as well as client applications (via a Web service) to explore a diversity of fine-grained perspectives on thousands of everyday concepts. Such perspectives range from obvious to insightful to surprising; but just how valid are they? Can they be used to do more than to provoke, and to entertain? If the insights they provide are truly insightful, we should expect them to inform our computational metrics of similarity, so that their predictions of lexico-conceptual similarity can align more closely with human judgments. We test this possibility by enriching WordNet with compatible categorizations from *Thesaurus Rex*: simply, a *Rex* perspective $P-H$ on a concept C is added to WordNet, if H denotes some hypernym of C in WordNet. As shown in Veale & Li [31], when WordNet is augmented with perspectives from *Thesaurus Rex* in this way, a variant of the WordNet-based similarity metric presented in [40] can be used to yield similarity judgments on the 30 word-pairs of the Miller and Charles [43] test-set that align closely with mean average judgments for these pairs. Veale & Li [31] report a correlation of **0.93** with human judgments, equaling the best non-WordNet models, in an approach that is fast, transparent and explanatory.

What of *Metaphor Magnet*'s metaphorical outputs? As with *Thesaurus Rex*, this service returns results that can be both thought provoking, and entertaining; but, do they truly accord with human judgments? A comprehensive evaluation of *Metaphor Magnet* is presented in [29, 33], showing that this service is capable of accurately and intuitively partitioning a complex stereotype like *Baby* (to which the system ascribes 163 typical properties), into both positive perspectives (e.g., "you are my baby!"), and negative perspectives (e.g., "you are such a baby!"), on demand.

Moreover, Veale [29] shows that copula metaphors of the form T is M in the Google n-grams [38]—the basis of *Metaphor Magnet*'s ability to extrapolate from known metaphors to novel metaphors—are also broadly consistent with the properties and affective profile of each of its simile-derived stereotypes. In **87%** of cases, one can correctly assign the label *positive* or *negative* to a topic T , using only the copula metaphors for T in the Google n-grams as a guide. Furthermore, Veale [29] shows that the T is M copula metaphors in the Google n-grams provide enough representational coverage to figuratively highlight 99% of the stereotypical properties of a stereotypical target T .

Metaphor Magnet's companion service for framing metaphorical conceits in poetic form produces outputs that are a good deal more subjective. Once again, however, the Google n-grams are used to ensure that any outputs are grammatically and semantically sound. This framing service highlights the benefits of the SOA approach to CC: the complexity and scale of a service that involves many interacting components and large-scale data-resources is hidden from any end-user or client application. Moreover, the service continues to evolve with new features (and to reap the benefits of any new CC services that contribute to its poetical musings), while remaining in active use online.

Indeed, the outputs of this service need not be seen as complete poems in their own right, but as a source of poetic tropes that can be reused by other poem generation systems and services. Thus, for instance, poetry generators such as that described in [13]—which creates topical poems from fragments of newspapers and tweets—can use this service's rich inventories of similes, fine-grained categories and affective metaphors in their poetry.

VIII. CONCLUSIONS

Creativity does not arise from the simple application of rules, or even *meta*-rules, but from an insightful exploration of a conceptual space [3, 18, 30]. The possibility of insight, and the need to view familiar concepts from atypical perspectives, means that creativity is also a learning process. One learns from feedback as one creates for an audience, to develop an aesthetic sense of what works in which contexts. As creative agencies accumulate experience across successive commissions, they develop their own aesthetic filters, which allow them to present only the best options to a client. A SOA of creative Web services can likewise learn from its actions, to fine-tune its own aesthetic filters across the diverse requests it is tasked with. A set of creative services is ultimately an options provider, but a raw service that overwhelms with a generative barrage of unfiltered options is little better than one that does nothing at all.

Developers can use a SOA platform of CC services in a theory-neutral fashion, testing varying combinations of services, to see what works best for them. Client feedback will provide evidence as to the utility of competing theory-informed components, and may allow researchers to compose new hybrid theories from the most successful mixtures. For researchers and developers alike, a common service platform will allow the services themselves to learn, both from the information they are given, and from the knowledge they discover for themselves from the Web, as well as via feedback from client applications and their users, as to which outputs work best. In other words, a SOA platform for CC will combine the best

qualities of creative agencies in the real world, providing services that compete for our attention, while adapting to our needs.


ACKNOWLEDGMENTS

This research was partly supported by the 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 the Science Foundation Ireland, via the *Centre for Next Generation Localization*.

REFERENCES

1. A. Cardoso, T. Veale, and G. A. Wiggins, "Converging on the divergent: the history (and future) of the International Joint Workshops in computational creativity," *AI Magazine*, vol. 30, no. 15, pp. 15-22, 2009.
2. T. M. Amabile, "The social psychology of creativity: a componential conceptualization," *Journal of Personality and Social Psychology*, vol. 45, no. 2, pp. 357-376, 1983.
3. M. A. Boden, "Computational models of creativity," in *Handbook of Creativity*, R. J. Sternberg, editor, New York: Cambridge University Press, 1999, pp. 351-373.
4. J. S. Bruner, "The conditions of creativity" in *On Knowing: Essays for the Left Hand*, Cambridge: Belknap Press of Harvard University Press, 1962.
5. M. Csikszentmihalyi, *Creativity: Flow and the Psychology of Discovery and Invention*, New York: Harper Collins Publishers, 1996.
6. R. A. Finke, T. B. Ward, and S. M. Smith, *Creative Cognition: Theory, Research and Applications*, Cambridge: MIT Press, 1996.
7. J. P. Guilford, "Creativity," *American Psychologist*, vol. 5, no. 9, pp. 444-454, 1950.
8. J. P. Guilford, *The Nature of Human Intelligence*, New York: McGraw-Hill, 1967.
9. A. Koestler, *The Act of Creation*, London: Arkana/Penguin Books, 1964.
10. R. K. Sawyer, *Explaining Creativity: The Science of Human Innovation*, New York: Oxford University Press, 2006.
11. R. J. Sternberg, *Wisdom, Intelligence & Creativity Synthesized*, New York: Cambridge University Press, 2003.
12. S. Colton, "Creativity versus the perception of creativity in computational systems," in *Proceedings of the AAAI Spring Symposium on Creative Intelligent Systems*, Stanford, CA, 2008, pp. 14-20.
13. S. Colton, J. Goodwin, and T. Veale, "Full-FACE poetry generation," in *Proceedings of the 3rd International Conference on Computational Creativity*, Dublin, Ireland, 2012, pp. 95-102.
14. D. R. Hofstadter and Fluid Analogies Research Group, *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*, New York: Basic Books, 1995.
15. J. McCarthy, "Creative solutions to problems," in *Proceedings of the AISB'99 Symposium on Artificial Intelligence and Scientific Creativity*, Edinburgh, Scotland, 1999.
16. P. McCorduck, *Aaron's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen*, New York: W. H. Freeman, 1991.
17. M. L. Minsky, "Jokes and the logic of the cognitive unconscious," Massachusetts Institute of Technology Artificial Intelligence Laboratory, Cambridge, MA, *AI memo no. 603*, 1980.
18. A. Newell, J. C. Shaw, and H. A. Simon, "The process of creative thinking," in *Contemporary Approaches to Creative Thinking: A Symposium Held at the University of Colorado*, H. E. Gruber, Ed., New York: Atherton Press, 1962, pp. 63-119.
19. F. C. Pereira, *Creativity and Artificial Intelligence: A Conceptual Blending Approach*, Berlin: Walter de Gruyter, 2007.
20. V. Raskin, *Semantic Mechanisms of Humor*, Dordrecht: D. Reidel Publishing Co., 1985.
21. G. Ritchie, "Some empirical criteria for attributing creativity to a computer program," *Minds and Machines*, vol. 17, no. 1, pp. 67-99, 2007.
22. T. Veale and D. O'Donoghue, "Computation and blending," *Cognitive Linguistics*, vol. 11, no. 3/4, pp. 253-281, 2000.
23. T. Veale, "Incongruity in humor: root-cause or epiphenomenon?," *Humor*, vol. 17, no. 4, pp. 419-428, 2004.
24. T. Veale, "An analogy-oriented type hierarchy for linguistic creativity," *Journal of Knowledge-Based Systems*, vol. 19, no. 7, pp. 471-479, 2006.
25. T. Veale and Y. Hao, "Comprehending and generating apt metaphors: a Web-driven, case-based approach to figurative language," in *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, Vancouver, Canada, 2007, pp. 1471-1476.
26. T. Veale, G. Li, and Y. Hao, "Growing finely-discriminating taxonomies from seeds of varying quality and size," in *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, Athens, Greece, 2009, pp. 835-842.
27. Y. Hao and T. Veale, "An ironic fist in a velvet glove: creative mis-representation in the construction of ironic similes," *Minds and Machines*, vol. 20, no. 4, pp. 635-650, 2010.
28. T. Veale, "Creative language retrieval: a robust hybrid of information retrieval and linguistic creativity," in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, Portland, OR, 2011, pp. 278-287.
29. T. Veale, "A Context-sensitive, multi-faceted model of lexico-conceptual affect," in *Proceedings of the 50th Annual Conference of the Association for Computational Linguistics*, Jeju, Korea, 2012, pp. 75-79.
30. T. Veale, *Exploding the Creativity Myth: The Computational Foundations of Linguistic Creativity*, London: Bloomsbury Publishing, 2012.
31. T. Veale and G. Li, "Creating similarity: lateral thinking for vertical similarity judgments," in *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, Sofia, Bulgaria, 2013.
32. T. Veale, "Less rhyme, more reason: knowledge-based

- poetry generation with feeling, insight and wit,” in *Proceedings of the 4th International Conference on Computational Creativity*, Sydney, Australia, 2013, pp. 152-159.
33. T. Veale, “Once more, with feeling! using creative affective metaphors to express information needs,” in *Proceedings of the 4th International Conference on Computational Creativity*, Sydney, Australia, 2013, pp. 16-23.
 34. G. Wiggins, “A preliminary framework for description, analysis and comparison of creative systems,” *Journal of Knowledge-Based Systems*, vol. 19, no. 7, pp. 449-458, 2006.
 35. T. Erl, *SOA: Principles of Service Design*, Upper Saddle River: Prentice Hall, 2008.
 36. D. Fishelov, “Poetic and non-poetic simile: structure, semantics, rhetoric,” *Poetics Today*, vol. 14, no. 1, pp. 1-23, 1993.
 37. R. Giora, “Optimal innovation and pleasure,” in *TWLT20: the April Fools’ Day Workshop on Computational Humour: Proceedings of the Twentieth Twente Workshop on Language Technology*, Enschede: Universiteit Twente, 2002.
 38. T. Brants and A. Franz, *Web 1T 5-gram Version 1*, Philadelphia: Linguistic Data Consortium, 2006.
 39. C. Fellbaum, *WordNet: An Electronic Lexical Database*, Cambridge: MIT Press, 1998.
 40. N. Seco, T. Veale, and J. Hayes, “An intrinsic information content metric for semantic similarity in WordNet,” in *Proceedings of the 16th European Conference on Artificial Intelligence*, Valencia, Spain, 2004, pp. 1089-1090.
 41. E. de Bono, *Lateral Thinking: Creativity Step by Step*, New York: Harper & Row, 1970.
 42. Z. Kozareva, E. Riloff, and E. Hovy, “Semantic class learning from the Web with hyponym pattern linkage graphs,” in *Proceedings of the 9th SIGdial Workshop on Discourse and Dialogue*, Columbus, OH, 2008, pp. 1048-1056.
 43. G. A. Miller and W. G. Charles, “Contextual correlates of semantic similarity,” *Language and Cognitive Processes*, vol. 6, no. 1, pp. 1-28, 1991.



Tony Veale

Dr. Tony Veale is a Computer Scientist at University College Dublin, Ireland. He is currently a visiting professor in Web Science at KAIST, Korea Advanced Institute of Science and Technology, as part of the Korean government’s World Class University initiative. His research focuses on the computational treatment of creative linguistic phenomena, including metaphor, simile, analogy and irony. Dr. Veale leads the European Commission’s PROSECCO project, which is an international coordination action that aims to promote the scientific exploration of computational creativity.