

# Creativity as a Web Service: A Vision of Human and Computer Creativity in the Web Era

Tony Veale

Web Science and Technology Division, KAIST, Daejeon, South Korea  
tony.veale@gmail.com

## Abstract

The marketplace for definitions and theories of creativity is crowded indeed. No hard consensus exists on the elements of an all-embracing theory, or on what specific sub-processes and representations are required to support creativity, either in humans or in machines. Yet commonalities do exist across theories: a search for novelty and utility is implied by most theories, as is the notion that an innovation can be considered creative only if it is not *too novel*, and can be adequately grounded in the familiar and the understandable. Computational creativity (CC) is the pursuit of creative behavior in machines, and seeks inspiration from both AI and from human psychology. As a practical engineering endeavor, CC can afford to adopt a *cafeteria approach* to theories of creativity, taking what it needs from different theories and frameworks. In this paper we present a vision for CC research in the age of the Web, in which CC is provided *on tap*, via a suite of Web services, to any third-party application that needs it. We argue that this notion of *Creativity as a Service* – which is already a popular business model for human organizations – will allow CC researchers and developers to build ad-hoc mash-ups of whatever processes and representations are most suited to a given application. By offering CC as a centralized service, we can collect statistics on the most useful mash-ups, and therein obtain a new empirical basis for theorizing about creativity in humans and in machines.

## The Creative Web

Creativity is an elusive phenomenon that organizations put significant effort and resources into fostering, rewarding, retaining, and reproducing on demand. But 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 so many different social, cultural and

contextual factors. For these reasons, organizations often out-source their creative needs to external agencies with a track record in creative exploration, idea composition, and in the optimal framing of creative outputs. Such agencies are not so much *problem solvers* as *option providers*, leaving the ultimate responsibility for choosing amongst this diversity of new options to the client. To out-source in this way is not to abdicate creative responsibility, but to merely broaden the range of choices one can choose from (see e.g. Veale, 2012).

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 computational creativity. 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 key to creative idea generation. 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 in content) that are available to choose from.

Our vision of a creative Web service imagines three kinds of sub-service: *discovery services*; *composition services*; and *framing services*. Each sub-service may rely on different sources of knowledge, but each will use inter-operable data structures and so each can call on other sub-services to achieve its goals. The overall architecture is theory-neutral, yet will provide a diversity of theory-informed sub-services that can be composed in any way that suits a client system's needs.

**Discovery 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 foregrounded while many others remain implicit, latent or presumed, in the conceptual background. 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 *bisociative* juxtaposition (Koestler, 1964). Discovery services will mine text corpora for implicit knowledge, and provide bisociative tools for acquiring emergent insights from the crossroads of diverse document sets.

**Composition services:** Creativity often arises from frame conflict, when one concept is incongruously viewed through the lens of another very different idea. The key to the fruitful exploitation of frame conflict is two-fold: one must first choose which concepts to place into juxtaposition, and then formulate a resonant form for the resulting content. The proposed Web architecture will provide services for suggesting, elaborating and comprehending conceptual metaphors, analogies and blends, as well as services for accessing the large store of common-sense world knowledge that these composition services will crucially rely upon.

**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, a simile, a joke, a name, a slogan, a short story, a poem, a picture, a 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. The proposed Web architecture will 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.

No architecture for providing creativity on the Web can be exhaustively complete, yet any that aims to be credible must provide enough sub-services to be initially useful, while demonstrating the interoperability and extensibility of the infrastructure as a whole. The proposed architecture will leave considerable room for future growth, but will be significantly usable and useful in its initial form, with the provision of core ideation functions such as bisociation, metaphor and blending, and popular framing devices such as poetry, painting and music.

## On Demand, But At Arm's Length

Creativity does not arise from the simple application of rules, or even *meta*-rules, but from an insightful exploration of a conceptual space. 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 creative Web service should likewise learn from its actions, to fine-tune its own aesthetic filters across the diverse requests it is tasked with. A creative service is an option provider, but a service that overwhelms with a combinatorial explosion of options is no better than one that does nothing at all.

For these reasons, the proposed architecture for creative Web services should act as a force magnifier, both for research in computational creativity (CC) and for the effective deployment of real applications of CC in industry. For researchers, the architecture will allow for a pooling of technologies in a robust inter-operable framework, in which CC models are conceived, developed and migrated from lab settings to an industry-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 the literature. These developers can work in a theory-neutral fashion, testing different combinations of services to see what works best in their applications. Client feedback on which services are used with which will then provide empirical evidence as to the utility of different theory-informed components, and may even allow theoreticians to compose new models of human creativity from the most successful mixtures. For researchers and developers alike, the use of a common service platform will allow the service components themselves to continuously learn, both from the information they are given and from the knowledge they induce for themselves, as well as from the feedback they receive from client applications about which options are most appreciated by end-users.

We envision two kinds of client applications for this framework: conventional applications that make no overt claim to creativity, but which need to diversify the options available to them, whether to overcome an impasse or to engage with a user via novel variation; and overtly creative systems that aid the user in a traditionally creative task, such as writing, picture generation, musical composition, etc.

## Innovation and Variation

Variety may be the spice of life, but moderate levels of variation are key to the enjoyment of creative stimuli. Too little variation leads to staid repetition and boredom, while too much can appear unintelligibly random. To produce an *optimal innovation* – an enjoyably novel variation that is recognized as a knowing twist on an accepted norm (Giora, 2002) – an application must walk a fine line between too much and too little variation. Moreover, variations must reflect a clear pragmatic intent to be appreciated as creative.

Many kinds of application can benefit from this ability to introduce variation on demand, to keep their interactions with the user constantly fresh. Games, for instance, or educational applications, or any program that interacts with the user via words, images or sounds, should marry a degree of predictability with a capability for surprising variation to keep their users engaged. Rather than hard-wire a weak generator of scripted variations into an application, developers can instead appeal to creative Web services to provide options for novel variation as they are needed. For truly intelligent variation requires a deep knowledge of conventions, and how they can be exploited (or indeed, subverted) relative to audience expectations. As the services framework grows in richness – of its services, its knowledge of convention, and its aesthetic filters – any client application that uses it can automatically benefit from its increasingly nuanced and clever variations.

### An Example Service: *Metaphor Magnet*

As an example service, consider *Metaphor Magnet*, an application on the Web that both understands and generates affective metaphors. Now, metaphor is a knowledge-hungry creative process *par excellence* (e.g. see Martin, 1990; Fass, 1996; Veale & Hao, 2007; Shutova, 2010; Veale, 2011, 2012), one that cannot be reduced to a single algorithm that operates over third-party representations. Rather, the algorithms and representations of a metaphor processing system are tightly co-dependent, so it makes even more sense to hide this complexity behind the simple interface of a remote Web service. This is all the more so when, as in *Metaphor Magnet*, the representations are dynamically acquired from large amounts of Web data, and where the resulting interpretations are evaluated and ranked according to the amount of evidence that can be found for them in a very large text corpus of attested language use. This is not the kind of process that one downloads as a library, or simply re-implements in one's own way.

*Metaphor Magnet* employs an affective lexicon that constructs for each entry a stereotypical representation of its most typical properties and behaviors. For

instance, the stereotype for *leader* notes that leaders are stereotypically *charismatic, dignified, honored* and *eminent*; their typical behaviors include *policymaking, planning, campaigning, commanding* and *domineering*. These stereotypes are extracted from similes that are in turn harvested in bulk from the Web. As Archer (1954) noted, similes are a rich and insightful source of cultural knowledge, though the Web makes it much easier to collect a large body of similes than in Archer's day.

*Metaphor Magnet* is designed to be a lightweight Web application that provides both HTML output (for humans) and XML (for client applications). The system allows users to enter affective statements such as *Google is like -Microsoft, life is a +game, Steve Jobs is Tony Stark*, or even *Rasputin is Karl Rove* (inputs are case-sensitive; a minus sign indicates a negative spin, and a plus sign indicates a positive affective spin). Each input is expanded into a set of apt metaphors using copula statements in the Google n-grams (Brants & Franz, 2006), and each metaphor is expanded into a set of contextually apt qualities. In turn, each quality is then expanded into an IR query that is used to retrieve relevant evidence from the Web via Google. In effect, the system also allows users to interface with a search engine like Google using metaphor and other affective language forms. A live demonstration of the service can be accessed using a standard browser at this URL:

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

*Metaphor Magnet* can exploit the properties and behaviors of its inventory of almost 10,000 stereotype representations, which have been mined from the Web. It can also acquire ad-hoc representations on demand. When accessed as a service, *Metaphor Magnet* returns either HTML or XML data, via simple *get* requests. For illustrative purposes, each HTML page also provides the URL for the corresponding XML-structured data set.

The system's interpretation of the affective simile "*Google is as -powerful as Microsoft*" highlights a range 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 an apt quality, such as "menacing giant" – is clickable, to yield linguistic evidence for the selection and corresponding Web-search results (obtained via a Google gadget). The phrase cloud representing *Microsoft* in this simile is shown in the screenshot of Figure 1, while the cloud for qualities projected onto *Google* is shown in Figure 2.

*Metaphor Magnet* demonstrates the potential utility of affective metaphors in human/computer linguistic interaction, and provides a Web service through which other natural language applications can remotely tap into the system's metaphorical competence as needed.

## Source Metaphors: Microsoft

menacing:bully  
perilous:battle twisted:devil  
menacing:tough menacing:lackey  
isolated:exception bewildered:victim  
overbearing:king bloated:dinosaur  
menacing:threat anxious:investor  
troubling:question menacing:evil  
threatening:enemy arrogant:master  
twisted:joke domineering:leader threatening:competitor  
menacing:criminal  
threatening:problem  
menacing:giant arrogant:intellectual  
twisted:root conservative:corporation domineering:superior

**Figure 1.** A screenshot of a phrase cloud for the perspective cast upon the source “Microsoft” by the simile “Google is as -powerful as Microsoft”.

An important advantage to the delivery of computational creativity as a suite of Web services is that these services can be updated with new features and sub-services that can, in turn, be automatically discovered by a client application. For instance, *Metaphor Magnet* continues to evolve as a service since its initial deployment. In addition to providing metaphors on demand, the service now provides a poetic framing facility, whereby the space of possible interpretations of a given metaphor is crystallized into a single poetic form. Poetry can be viewed as a means of reducing information overload, of summarizing a complex metaphor whose interpretation entails a rich space of affective possibilities. A poem can be seen in functional terms as a visualization device, one that offers a concise alternative to the phrase clouds in Figures 1 and 2. For example, *Metaphor Magnet* produces the following poem as a distillation of the space of feelings and associations that arise from the interpretation of *Marriage is a Prison*:

### *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, you disgust me with your undesirable security*

## Target Metaphors: Google

domineering:leader threatening:enemy  
menacing:tough threatening:bully  
menacing:lackey calculating:investor  
menacing:criminal  
dreaded:king  
domineering:monopoly  
complicated:web threatening:competitor arrogant:intellectual  
vain:celebrity bloated:dinosaur  
twisted:devil  
threatening:problem  
threatening:threat  
isolated:exception twisted:root  
dreaded:evil troubling:question

**Figure 2.** A screenshot of a phrase cloud for the perspective cast upon the target term “Google” by the simile “Google is as -powerful as Microsoft”.

Each time we dip into the space of possible interpretations, a new poem is produced. 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 of the same metaphor in poetic form:

### *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, you depress me with your dreary consecration*

In the context of Figure 1, which samples the space of metaphors that negatively describe Microsoft’s perceived misuse of power, consider the following, which distills the assertion *Microsoft is a Monopoly* into an 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, you oppress me with your corrupt reign*

Poetry generation in *Metaphor Magnet* is a recent addition to the service, and its workings are beyond the scope of the current paper (though they may be observed in practice by visiting the aforementioned URL). For details of a related approach to poetry generation – one that also uses the stereotype-bearing similes described in Veale (2012) – the reader is invited to read Colton, Goodwin & Veale (2012).

## Developmental Issues

Can the development of creative web services tell us anything at all about the development of creative abilities in humans? Conversely, can insights about human creative development be of any use in crafting the most effective ecology of web services, since these services are expected to provide the kind of on-demand creativity that we expect from human professionals? Though the latter route offers the most likely vector for insight, the creative web vision does offer a basis for addressing both questions at once.

On the first question, engineers are likely to exploit insights about human creative development only when these insights offer a practical means of achieving a concrete engineering goal. Engineers will be theory-driven only to the extent that these theories offer a modular blueprint for design, and suggest well-defined and reusable software components at appropriate levels of abstraction (Veale, 2006). Theories that offer the most engineering traction, such as conceptual blending theory (Fauconnier & Turner, 1998, 2002; Veale & O'Donoghue, 2000) – which is still vague in some important respects, but which offers a rich ontology of components, principles and constraints – will map most obviously onto the framework of creative web services. Likewise, key elements of contemporary metaphor theory – which views metaphor as a process of mapping between *spaces* (Fauconnier, 1994) of discrete knowledge representations – map naturally onto the web services architecture. Software engineering is a discipline that seeks the most effective levels of abstraction, and real experience with theory-driven modules may inevitably lead to the development of new abstractions that were not unanticipated by the original theories. The continuous refactoring of software services may thus lead to a similar refactoring of the cognitive theories that underpin them.

The answer to the first question then is a qualified 'yes'. Importantly, we do not expect a framework like this to be fully operational and full-featured from the get-go. Initial development of core services, to provide core competences such as metaphor, conceptual blending and bisociation, will support future developments that incrementally build on these middleware services. We have discussed one such higher-level service here, that of poetry generation. Other high-level services may include joke generation, story generation, music / lyric composition, slogan generation,

and so on. These services will themselves grow in ambition and sophistication over time, as their developers add new features and new capabilities. For instance, joke generation services will likely progress from simple pun-based humor to more complex idea-based and narrative-based humor, in much the same way that human joke generation progresses from schoolyard puns to full-fledged conceptual humor. So we have reason to expect the trajectory of a web services architecture to mirror that of human creative development.

The answer to the second question is also affirmative, but it is an answer that needs no qualification. It seems obvious that computational services that replicate the outputs of creative humans – if not the precise cognitive mechanisms – can still benefit from algorithmic insights into the workings of those mechanisms. These insights will be necessarily generic and simplifying, but reductiveness is often advantageous to practical software development. Consider that the richly-featured stereotype model mined from similes on the web in Veale & Hao (2007a) uses an approach that is, in essence, a turbo-charged version of how humans learn about the world via the language of other speakers. Dickens, in his opening to *A Christmas Carol*, notes that “the wisdom of our ancestors is in the simile”, and it does indeed seem that much of our stereotypical knowledge of the world – such as that sharks are ruthless, peacocks are proud and lawyers are unscrupulous – more often comes from proverbial similes (Taylor, 1954) than from first-hand experience. Veale & Hao (2007b) use this insight to build a model of knowledge acquisition that is a much simplified but immensely practical version of human acquisition via language, while Veale & Hao (2007c) uses this knowledge in a web-driven model of metaphor interpretation and generation.

In summary then, while a web services architecture for creativity on-demand is not explicitly built upon human cognitive principles, and is not a direct or obvious source of insight into the development of human creative abilities, it can certainly benefit from the former and furnish indirect empirical evidence for the latter.

## Concluding Thoughts

Creativity is such a complex phenomenon that computationalists often focus on the aesthetic appearance of creativity, e.g. to generate poem-shaped texts, painting-like pictures, or music-like sounds. However, while framing is an important dimension of the creative process, one must be careful not to embrace pastiche when aiming to mechanically reproduce the superficial characteristics of a creative product. Creativity should always respect meaning, and form should most definitely be shaped by intent. For this reason, the vision of the *Creative Web* outlined here assumes an overtly conceptual pathway from discovery to composition to framing: creative Web

services will enable applications to produce outputs that possess the necessary formal characteristics, but which are also meaningfully anchored in a conceptual representation of an idea, a message, or a discovery.

This vision offers an ambitious plan for the future development of computational creativity, yet it does not posit an all-or-nothing gamble. While significant research and development barriers must be hurdled, the proposed plan has a number of high-value milestones that can be credibly reached. As more services are implemented and published on the Web, the deployment of these services will provide de-facto standards for the future development of additional services, lending a critical mass to the CC field as a whole. Moreover, as more services are published, such as *Metaphor Magnet*, the greater the number of mash-ups of these services that developers can implement. The resulting churn of services and mash-ups should spur the development of new theories of creativity, theories that posit new combinations of processes and representations, which are perhaps suited to specific domains of creativity. The key research question concerns the over-arching vision of the *Creative Web*: can a Web services framework really provide computational creativity *on tap* to any application that requires it? Only by taking concrete steps to answer this question will we ever know for sure.

## Acknowledgements

This research was 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).

## References

- Brants, Y. and Franz, A. 2006. Web 1T 5-gram Version 1. *Linguistic Data Consortium*.
- Colton, S., Goodwin, J. and Veale, T. 2012. Full-FACE Poetry Generation. *Proceedings of ICCO-2012, the 3<sup>rd</sup> International Conference on Computational Creativity, Dublin, Ireland*.
- Dickens, C. 1845. *A Christmas Carol*. Bradbury & Evans.
- Fass, D. 1997. Processing Metonymy and Metaphor. *Contemporary Studies in Cognitive Science & Technology*. New York: Ablex.
- Fauconnier, G. 1994. *Mental spaces: aspects of meaning construction in natural language*. Cambridge University Press.
- Fauconnier, G. and M. Turner. 1998. Conceptual Integration Networks. *Cognitive Science*, 22(2):133–187.
- Fauconnier, G. and Turner, M. 2002. *The Way We Think. Conceptual Blending and the Mind's Hidden Complexities*. Basic Books.
- Giora, R. 2002. Optimal innovation and pleasure. In O. Stock, C. Strapparava and A. Nijholt (eds.). *The April Fools' Day Workshop on Computational Humour*, ITC-irst, Trento, 11-28.

- Koestler, A. 1964. *The Act of Creation*. Hutchinsons, London
- Martin, J. H. 1990. *A Computational Model of Metaphor Interpretation*. NY: Academic Press.
- Shutova, E. 2010. Metaphor Identification Using Verb and Noun Clustering. In *Proc. of the 23<sup>rd</sup> International Conference on Computational Linguistics*, 1001-1010.
- Taylor, A. 1954. Proverbial Comparisons and Similes from California. *Folklore Studies* 3. University of California Press.
- Veale, T. and D. O'Donoghue. 2000. Computation and Blending. *Cognitive Linguistics*, 11(3-4):253-281.
- Veale, T. 2006. Computability as a test on linguistic theories. In G. Kristiansen, M. Achard, R. Dirven and F. Ruiz de Mendoza Ibáñez (Eds.). *Cognitive Linguistics: Current Applications and Future Perspectives. Applications of Cognitive Linguistics*. Mouton de Gruyter.
- Veale, T. and Y. Hao. 2007a. Making Lexical Ontologies Functional and Context-Sensitive. *Proceedings of ACL 2007, the 45<sup>th</sup> Annual Meeting of the Association of Computational Linguistics*, 57–64.
- Veale, T. and Y. Hao. 2007b. Learning to Understand Figurative Language: From Similes to Metaphors to Irony. *Proceedings of CogSci-2007, the 29<sup>th</sup> Annual Meeting of the Cognitive Science Society*.
- Veale, T. and Hao, Y. 2007c. Comprehending and Generating Apt Metaphors: A Web-driven, Case-based Approach to Figurative Language. *Proceedings of AAAI- 2007, the 22<sup>nd</sup> AAAI conference on Artificial Intelligence, Vancouver, Canada*.
- Veale, T. 2011. Creative Language Retrieval: A Robust Hybrid of Information Retrieval and Linguistic Creativity. *Proceedings of ACL'2011, the 49<sup>th</sup> Annual Meeting of the Association of Computational Linguistics, Portland, Oregon*.
- Veale, T. 2012. *Exploding the Creativity Myth: The Computational Foundations of Linguistic Creativity*. Bloomsbury Academic, London.