

# Running with Scissors: Cut-Ups, Boundary Friction and Creative Reuse

Tony Veale \*

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

[Tony.Veale@UCD.ie](mailto:Tony.Veale@UCD.ie)

<http://Afflatus.UCD.ie>

**Abstract.** Our experience of past problems can offer valuable insights into the solution of current problems, though since novel problems are not merely re-occurrences of those we have seen before, their solutions require us to integrate multiple sources of inspiration into a single, composite whole. The degree to which the seams of these patchwork solutions are evident to an end-user offers an inverse measure of the practical success of the reuse process: the less visible its joins, the more natural a solution is likely to seem. However, since creativity is neither an objective nor an intrinsic property of a solution, but a subjective label ascribed by a community, the more perceptible the tensions between parts, and the more evident the wit that one must employ to ameliorate these tensions, then the more likely we are to label a solution as creative. We explore here the conceit that creative reuse is more than practical problem-solving: it is reuse that draws attention to itself, by reveling in *boundary friction*.

**Keywords:** Boundary friction, incongruity resolution, semantic tension, computational creativity, creative reuse

## 1 Introduction

Robert Altman's movie *The Player* opens with a glorious tracking shot that swoops through the offices of a fictitious Hollywood film studio. Joining the camera on its sweeping arcs, the audience is allowed to eavesdrop on insider conversations that illuminate the Hollywood philosophy of filmmaking. In one office we hear a group of technocrats discussing the longest and most impressive tracking shots in the history of cinema, while in another we hear a writer pitching a new idea to a studio executive. The writer frames his pitch as a high-concept blend; this new film, he suggests, is best imagined as a cross between *Ghost* and *The Manchurian Candidate*. Altman wants us to laugh at these scenes, to set the tone for the sharp-edged parody to follow. For we can only marvel at why a writer should want to marry elements from a then-recent blockbuster romance about

---

\* This work forms part of the EC-funded WHIM project (*The What-If Machine*). The author would like to acknowledge the valuable contribution of his WHIM collaborator Alessandro Valitutti to the work described in this paper.

a lovelorn ghost with aspects of a grim political thriller about brainwashing and McCarthyism. Altman seems to be suggesting that Hollywood's love of the high-concept pitch, combined with a ruthless determination to recycle past successes, has led to a cynical dumbing-down of its creative processes.

Hollywood has long sought to create familiar surprises by recombining past successes in high-concept way ([?]). Consider Akira Kurosawa's 1954 Japanese film *The Seven Samurai*, which became the 1960 American Western *The Magnificent Seven* when sword-wielding samurai were re-imagined as heroic gun-slingers. Kurosawa's own *Ran* can also be seen as a Japanese reworking of Shakespeare's *King Lear*, though Shakespeare was hardly the first to repackage the Celtic legend of Lir and his daughters. The 1961 musical *West-Side Story* is a Hollywood reworking of Shakespeare's 1594 play *Romeo and Juliet*, in which rival Verona families the Montagues and the Capulets become the rival street-gangs of the Jets and the Sharks. Baz Luhrman's movie *Romeo + Juliet* is also a modern musical set amongst crime families in Verona beach, with most of its dialogue taken directly from Shakespeare's play. Surprisingly, the re-combination not only works but excels, and big risks often produce the biggest rewards. Though cinematic titans like Orson Welles had struggled and failed to film a worthy treatment of Joseph Conrad's novel *Heart of Darkness*, the director Francis Ford Coppola succeeded with his 1979 movie *Apocalypse Now*, which transposed the action from the Belgian Congo to Vietnam and Cambodia so as to critique the U.S. role in the Vietnam war. In its own way, the 1994 film *Forrest Gump* also found success by making Vietnam the centerpiece of its rather loose re-combination of elements from Voltaire's 1759 novel *Candide*. The Coen brothers depression-era comedy *O Brother, Where Art Thou* was also a rather loosely-based reworking its classical source, Homer's *Odyssey*, but the comedic re-combination works so well that the film won an Academy Award for best adapted screenplay.

When stripped of their commercial motivation, the reuse mechanisms favored by Hollywood studios bear a surprising similarity to the recombinant techniques pioneered by the surrealist school of experimental art. The surrealists sought more than a conscious disavowal of cliché; they sought a radical means of escaping the deep-hewn ruts that unconsciously inhibit our spontaneity and creativity. As William Burroughs put it, one “cannot will spontaneity into being”, but one can “introduce the unpredictable spontaneous factor with a pair of scissors”. The scissors here alludes to the cut-up technique pioneered by Burroughs and artist/writer Brion Gysin, in which a linear text is sliced, diced and randomly re-spliced to form new texts that give rise to new and unexpected interpretations ([?]). The purpose of the cut-up technique is actually two-fold: not only does it aim to create new combinations from old, much like the Hollywood system, it also consciously aims to disrupt the mind's attempts to automatically group commonly co-occurring words and ideas into familiar gestalts. Unlike the Hollywood approach, the technique embraces uncertainty and incongruity, and aims to challenge rather than to comfort, to unsettle rather than to placate.

The cut-up technique was originally inspired by the collage movement in visual art, in which fragments of images and texts are re-combined to form a

novel patchwork whole. When the cut-up technique is applied to a linear text, the text is segmented into short strands of contiguous words that do not necessarily respect either phrase or sentence boundaries. These strands are then randomly recombined, to form a new text that uses the same words in different linear juxtapositions, to facilitate – if one charitably overlooks the inevitable bad grammar and illogical punctuation – very different global interpretations. Gysin originally applied the technique to layers of newsprint, which he sliced into linguistic chunks with a razor, and Burroughs later extended the technique to audio tapes. In principle, any linear source of information, from text to audio to video and even DNA, can be sliced and re-spliced using the cut-up technique to deliberately subvert familiar patterns and spontaneously suggest new meanings. Note, however, that the cut-up technique does not actually create new meanings, and is “*merely generative*” in the purest sense. Rather, the goal of the cut-up technique – and of related techniques developed by Burroughs, such as the *fold-in* and the *drop-in* – is to generate candidates for interpretation that an artist then evaluates, filters and ranks according to their creative potential. By automatizing the production process in a way that frees it from the often baleful effects of cliché, such techniques also free a creator to step back, observe, and focus instead on the crucial acts of selection and evaluation. In other words, these techniques are the arts-and-crafts equivalent of a generative algorithm

### 1.1 Boundary Friction

The cut-up technique has intriguing resonances with automatic approaches to language processing, which also aim to understand complex texts by carving them into smaller and more manageable chunks. For instance, *boundary friction* is a problematic phenomenon observed in a case-based approach to machine translation called “*Example-Based Machine Translation*” (or EBMT; see [?], [?], [?]). In a typical EBMT system, a target-language translation of a text is stitched together from a combination of pre-translated linguistic chunks that were previously extracted from a parallel bilingual corpus. While these chunks are internally coherent, the recombination of chunks can often give rise to cross-chunk incoherencies. For example, a system may translate a novel source-language sentence by combining a target-language chunk for its subject noun-phrase with a target-language chunk for its verb-phrase. But if these NP and VP chunks, derived from different translation contexts, now disagree on the grammatical number of their head noun and head verb, this friction must be ameliorated with an appropriate word-level change to achieve native fluency. Commercial EBMT systems succeed or fail on their ability to reduce boundary frictions and thus produce a fluent, natural translation. In contrast, the creation of boundary friction is the very *raison d’être* of the cut-up technique, whose explicit goal is the generation of texts whose juxtapositions are both challenging and surprising. It is rare that the cut-up technique generates fluent texts straight out of the gate; though such outputs are possible, albeit with careful curation, the goal is not simply to replace one natural text with another.

EBMT is most successfully applied in domains where there is a logical progression from past to present to future texts, as in the evolving documentation of a maturing software product. Since present texts are known to the system and exist in multiple translations, an EBMT system can carve out a collection of reusable translation chunks by aligning parallel translations of the same source texts. Though future texts are unknown and unseen, their subject matter is easy to predict, so in an important sense the present allows the future to be predicted. The general principle at work here – which sees the future as a functional re-arrangement of the present – also applies to artistic cut-ups. By cutting and re-arranging a text on a given subject, one is likely to produce another text on a similar subject, albeit one with some provocative juxtapositions. Occasionally, it is even possible that a cut-up produces a text that is not just a logical progression of the original, but a testable prediction of some future state of affairs. Burroughs describes one such experiment with cut-ups, in which a text about the oil billionaire John Paul Getty was randomly re-arranged to produce the sentence “*It’s a bad thing to sue your own father.*” A year later, Burroughs claims, the billionaire was indeed sued by one of his own sons. This is a hand-picked example, one that shows the necessity of carefully filtering the generative products of the cut-up technique, but it does show the capacity of cut-ups to bring into focus the latent possibilities that implicitly reside in a text. Burroughs described the relation between combinatorial generation and prediction thusly: “*when you cut into the present, the future leaks out.*” In truth, cut-ups are no more predictive of the future than fortune cookies or horoscopes, yet if creativity is our goal, spontaneity and unpredictability are often the better parts of generation when placed in the hands of a selective creator.

## 1.2 Forgiven, Not Forgotten

An EBMT system that fails to resolve all of the frictions that arise at the boundaries of a patchwork translation is effectively creating more work for a human post-editor. Though such frictions are primarily grammatical in nature, they undermine a reader’s faith in the competence of a translation even if the gist of its meaning is still apparent. In contrast, the cut-technique often produces piecemeal texts that ask for a suspension of superficial grammatical rigour from their readers, by posing an implicit bargain that promises conceptual rewards at a deeper level of meaning. Whereas an EBMT system strives to produce a target-language translation that would not look out of place in its original corpus of target-language texts, the cut-up technique strives to produce novel composites that reside comfortably in *none* of its contributing corpora. EBMT systems thus seek to obliterate the joins in its composite outputs, while instances of the cut-up technique wear these joins proudly, as signifiers of creativity.

A cut-up text is a text of divided references, for it simultaneously evokes different viewpoints on the world as captured in different contributory fragments. This rift may prompt a subtle convergence of views that goes unnoticed as a creative act, or it may force an unsubtle divergence that calls attention to itself as a creative provocation. As Howard Pollio puts it, “*split reference yields humour*

if the joined items (or the act joining them) emphasize the boundary or line separating them; split reference yields metaphor if the boundary between the joined items (or the act joining them) is obliterated and the two items fuse to form a single entity” [?]. While metaphor is so pervasive in language that speakers are often unaware of just how much they rely on common figures of speech, the spell of metaphor is easily broken, and with jarring effect, when metaphors are injudiciously mixed so as to clash at ill-fitting boundaries. Creative combinatorial reuse resolves the tensions that arise between the components that contribute to a solution, so that the resulting whole can achieve its intended goals, but it does not obliterate all traces of these tensions. Rather, though the whole works, it continues to surprise, by drawing attention to itself as an unnatural – albeit useful – marriage of strange bedfellows. In humorous wit, as in ostentatious creativity more generally, one must not fail to address the internal tensions in a juxtaposition of forms and ideas, but neither must one resolve them so utterly that an audience does not recognize their influence, nor appreciate the means of their resolution ([?], [?]). Such tensions are a vital part of how creativity announces itself to an audience

### 1.3 Structure Of This Paper

EBMT systems suffer from *superficial* boundary friction, which can diminish our faith in the quality of a translation if not fully resolved. Creative applications of the cut-up technique revel in *deep* boundary friction, at the semantic and pragmatic levels of meaning rather than at the surface levels of language. In this paper we will computationally explore the latter, deeper source of friction, to generate cut-ups of our knowledge-representations that are non-obvious and challenging, yet ultimately meaningful and occasionally even humorous, insightful and profound.

In section two we introduce the notions of a *knowledge cut-up* and a *conceptual mash-up*, to demonstrate that the surrealist cut-up technique can be applied as much to our logical representations of knowledge as to our linguistic realizations of this knowledge in text. To a computer, after all, a symbolic knowledge representation is merely another kind of symbolic text, albeit one that is unambiguous and logically-privileged. In section three we describe a novel, inferential use of these new techniques, and describe, in section four, how inferential cut-ups form the basis of a creative, fully-automated Twitterbot. Section five then uses the outputs of this system to obtain some initial experimental findings regarding the kinds of boundary friction that are most satisfying to human audiences. The paper concludes in section six with some closing observations about the creativity of compositional reuse.

## 2 Knowledge Cut-Ups and Conceptual Mash-Ups

The mechanics of the cut-up technique are easily automated, though superficial boundary friction abounds when the technique is randomly applied to raw texts.

The more informative the mark-up that can be extracted from a text, the more well-formed the outputs can be. Consider a twitterbot called *Two Headlines* by bot-designer Darius Kazemi (Twitter handle *@twoheadlines*), which does exactly what its name suggests. Two random headlines from today’s news, as provided by a popular news aggregator, are mashed together to yield a headline that is well-formed yet factually provocative. Because the source headlines contain additional mark-up that identifies the named-entities in the text – whether *Justin Bieber* or *Barack Obama* or *Google* or *Italy* – it is a simple matter to mix-and-match the entities from two headlines into one without introducing any boundary friction at the grammatical level. The goal, of course, is to introduce friction at the conceptual level by, for example, associating Justin Bieber with momentous international political decisions or scientific discoveries, or Barack Obama with the laughable or criminal indiscretions of pampered celebrities. The resulting headline often produces what Koestler ([?]) terms a “*bisociation*”, a clash of two overlapping but conflicting frames of reference that is jarring but meaningful.

The more mark-up present in the input texts, the more insight that can be gleaned about their structure and their meaning, and so the more informed the decisions that a computational cut-up agent can make when creating a new text. The most informative texts are those that explicitly define the knowledge possessed by a computer for a particular domain: that is, the computer’s own symbolic knowledge representations. For example, these representations are commonly structured around the notion of a semantic triple, in which a logical subject is linked to a logical object by an explicit semantic relation. These triples carve up the world at its joints, but it is an easy matter to further cut-up these triples at their joints, to mix-and-match subjects, relations and objects from different triples into provocative new hypotheses, in much the same way that *@twoheadlines* serves up a bisociative mix of today’s news headlines. The triples in question can come from any generic triple-store on the Semantic Web, or they may form part of the bespoke knowledge-base for an AI/NLP reasoning system. Consider the triples harvested by Veale and Li ([?], [?]) from the texts of the Web. These authors glean completions to common why do questions from the Google search engine, such as “*why do dogs chase cats*”, “*why do philosophers ask questions*” and “*why do poets use rhyme*”. They then extract the implicit consensus presupposition at the core of each question (such as e.g. that most poets use rhyme, imagery and metaphor) and automatically convert each presupposition into an explicit semantic triple (such as  $\langle poet, use, metaphor \rangle$ ).

These triples are grist for a metaphor interpretation and generation system named *Metaphor Eyes*. Suppose *Metaphor Eyes* were to consider the metaphorical potential of viewing philosophers as poets. The system generates knowledge-level cut-ups from the triples it possesses for *poet* by replacing the subject *poet* with the subject *philosopher* in each *poet* triple. This simple, generative action is then followed by a validation phase that automatically evaluates each of the newly-produced cut-up triples, by looking to the Web for evidence that the new triple captures a relationship that is attested in one or more Web texts. Thus, the cut-up that “*philosophers use metaphors*” is attested by a Web search that

finds hundreds of supporting documents. But *Metaphor Eyes* does more than generate and test cut-ups for a given pairing of concepts: it suggests interesting pairings for itself, by examining how words and ideas are clustered in a large text corpus such as the Google n-grams ([?]). Thus, since *poets* and *philosophers* are seen in sufficiently frequent proximity to each other, and each is found in proximity to ideas that are related to the other (such as *metaphor*, or *idea*, or *argument*), *Metaphor Eyes* actively suggests the pairing of poets and philosophers to the user, and proceeds to generate a Web-validated knowledge cut-up as its output. *Metaphor Eyes* fluidly plays with the boundaries between knowledge and text, moving between one and the other to extract maximum reuse value from each.

Veale and Li also explore the potential of metaphorical cut-ups to plug gaps in a system’s knowledge representation of the world. Suppose a system possesses little or no knowledge about philosophers, but possesses knowledge about concepts that are not so far removed, such as scholars, theologians, scientists and poets. *Metaphor Eyes* combines triples from neighboring concepts to mash together a proxy representation for *philosopher*, by combining the Web-validated triples for the cut-ups *philosopher as scholar* (each accumulates and applies knowledge, each performs research), *philosopher as theologian* (each spreads doctrines), *philosopher as scientist* (each develops and explores ideas) and *philosopher as poet* (each nurtures ideals). The resulting set of cut-up triples from diverse sources is called a “*conceptual mash-up*”, and Veale and Li rank the triples in a mash-up according to how much Web evidence can be found for each. They demonstrate empirically that the top-ranked triples in a mash-up have a high probability of being attested in real Web texts, and show that a mash-up is a good substitute for human-engineered knowledge in an AI system. Like human users of the cut-up technique, they explore the space of ideas with a scissors.

### 3 Inferential Cut-Ups Lead to Surprising Conclusions

Creative producers are masters of ambiguity. They make the most of the ambiguity in their inputs, and induce ambiguity in their outputs to foster indeterminism and the emergence of new, unexpected meanings. The cut-up technique is designed to unleash the latent ambiguity in an otherwise business-as-usual text, such as a news story or a well-thumbed novel. In a knowledge-based computational setting, ambiguity allows a system to transcend one of its most vexing limitations: the knowledge that comprises most AI systems is *safe* knowledge, facts and rules that most informed users would consider to be true if they ever gave the matter any thought. So how does one go from safe premises to conclusions that can genuinely surprise us, and that force us to view the world anew from a very different perspective? The ancient sophists were masters of using familiar knowledge to reach desirable ends, and even Socrates, who openly disdained the sophists, practiced similar techniques for nobler ends: by linking familiar assumptions in the right order, a sophisticated thinker can lead an audience from their safe moorings in received wisdom into new, uncharted territories.

Suppose we ponder the dystopian scenario of a “*world without beauty*”? Since *Beauty* is an unqualified positive in most world-views, and most knowledge-bases that address the matter at all will axiomatize it using other positive qualities such as *Art* and *Love*, most of us (and our systems) will initially conclude that a world without beauty would be a much-diminished place to live. Yet a good repository of common-sense knowledge will also capture the relationship between *Beauty* and *Jealousy*, noting that the former often causes the latter. *Jealousy*, in turn, promotes *Hate* and *Conflict*, with the latter sometimes escalating to the level of *War*. A good sophist might thus conclude – and show how others might reasonably reach this conclusion for themselves – that a world without beauty might be a world with less jealousy, less hate, and less war. A world without beauty might, in fact, be a desirable place to live.

The elements of this causal chain are unassuming semantic triples of a knowledge base, and none have been subjected to the dislocations of the cut-up technique. Yet the chain as a whole can be viewed as a cut-up in its own right, akin to a newspaper cut-up that reuses whole sentences as its building blocks rather than the phrasal sub-components of these sentences. The guiding principle of such a cut-up is to produce a surprising conclusion from the most unsurprising premises, much as Socrates and his followers, and their rivals the sophists, did in ancient Athens. Such a process can also be modeled on a computer, as computers excel at pursuing inferential chains to their logical ends. Each successive link in a chain may exhibit little or no boundary friction with the link that goes before or the link that comes after, but the cumulative effect will be to produce a very evident friction between the first and last links in the chain.

The potential for semantic and pragmatic friction between the end-points of a chain is increased if one is willing to use a sophist’s disregard for logical rigor when adding new links to the chain. Consider  $\langle \textit{dictator}, \textit{suppress}, \textit{critic} \rangle$ , a triple which captures the widespread belief that dictators censure their critics, or worse. A system that also believes that  $\langle \textit{critic}, \textit{criticize}, \textit{artist} \rangle$  and  $\langle \textit{artist}, \textit{produce}, \textit{art} \rangle$  may well construct an inferential chain from dictators to art via critics and artists, to infer that dictators indirectly promote art by thwarting the critics that impede its producers. Such a chain embodies a surprising claim, that more dictators lead to more art, but it is predicated on several acts of sophistry. Firstly, critics come in multiple guises, and the political critics that decry a regime are not typically the same as those that criticize art. Secondly, to criticize is not always to hinder or deter. Thirdly, art critics do more than criticize, and often encourage artists too. Lastly, even if the inferential chain can be taken at face value, it says nothing about whether dictators promote more *good* art. These may be considered the weak points of the causal argument, but they may also be seen as the resolution of the semantic friction – between bad dictators and good art – that gives the argument its shock value. One man’s sophistry may well be another’s *semantic slippage* ([?]). In either case, inference chains that exhibit deep friction often prove to be surprising in ways that make us think about what we know.



This notion of deep friction in an inferential chain may be operationalized in different ways, but the simplest and most effective employs the affective profile of the concepts concerned. *Art* and *Beauty* are positive concepts, and are denoted by words with a strong positive sentiment. *Dictators* and *Critics* are negative concepts, and are denoted by words with a strong negative sentiment. It follows that more art, more beauty and fewer critics should be considered positive outcomes, just as less art, less beauty and more dictators should be considered a negative outcome. An inferential chain that shows how a positive outcome can be derived from a negative cause (such more art from more dictators), or how a negative outcome can be derived from a positive cause (such as more war from more beauty) thus exhibits a deep friction between its initial premises and its final conclusion. This friction poses a challenge to an audience – *how can this be so?* – that must be resolved meaningfully, either by accepting the conclusion at face value or by identifying the sophistry at its heart. In either case, the audience is aware of both the friction and its resolution; indeed, the resolution actively draws our attention to the friction, and draws us into its worldview.

#### 4 Creative Twitterbots

Twitterbots are fully-automated, generative systems whose outputs are designed to be distributed via the micro-blogging service *Twitter*. Though easily dismissed as trivial, Twitterbots embody all of the anarchic spirit of the cut-up technique and of the conceptual artists that created it. Consider again the *@twoheadlines* bot, which uses the cut-up technique just as Burroughs and Gysin imagined it, to “*introduce the unpredictable spontaneous factor*” into our readings of the daily news. However, as a simple, uncurated system, *@twoheadlines* captures only the generative part of the cut-up process, and omits the vital reflection phase in which candidate outputs are ranked, evaluated and then chosen or discarded by an intelligent agent with creative intent. Though a productive system, *@twoheadlines* is ultimately a “*merely generative*” system that lacks any self-appreciation, and cannot thus be considered a computationally creative agent.

In contrast, *@MetaphorMagnet* is a Twitterbot that employs the knowledge cut-ups of section 2 (via the computational service *Metaphor Eyes*) and the inferential cut-ups of section 3 to produce metaphorical insights that are carefully evaluated and filtered by the system itself. *@MetaphorMagnet* is not merely a generative system; rather, it is a system that only tweets metaphors that it considers to be well-formed and thought-provoking, the latter quality arising out of its deliberate use of resolvable, causal boundary friction. The bot generates a new metaphorical insight every hour, based on its explorations of its own sizable knowledge-base of semantic triples. Because these explorations allow it to detect deep frictions between its own knowledge of the world – knowledge that it believes most humans will also possess – *@MetaphorMagnet* has a sharp eye for hypocrisy and disappointment that is well-suited to ironic commentary. Consider these representative tweets:

#Irony: When the initiates that learn about engaging mysteries participate in boring rituals. #EngagingOrBoring #Initiate

Writers write metaphors. Dictators suppress the critics that criticize the authors that write metaphors. Who is better? #DictatorOrWriter

Slavery imposes the shackles that impose bonds. Marriage creates the attachments that create bonds. Take your pick. #MarriageOrSlavery?

Truth provides the knowledge that underpins arts. Illusion enhances the beauty that underpins arts. Take your pick. #TruthOrIllusion?

#Irony: When the tests that are conducted in tidy laboratories are conducted by rumpled professors. #TidyOrRumpled #Test

Twitter’s 140-character size limitation per tweet imposes obvious limitations on the length of the inferential chains that can be entertained by the system, though its biggest challenge is not the size of its search horizon, nor the size of its knowledge-base, but the breadth of its repertoire for expressing new insights in suggestive forms. The best Twitterbots exhibit an identifiable aesthetic and an identifiable world-view, in much the same way that human Twitter users exude an identifiable personality. *@MetaphorMagnet* exudes a distinctly hard-boiled personality through its causal linking of positive and negative concepts. As shown in the *#MarriageOrSlavery* and *#TruthOrIllusion* tweets above, *@MetaphorMagnet* is a sophist that assumes moral equivalence wherever it can demonstrate an apparent causal equivalence. We leave it to other Twitterbots to present a more positive and uplifting view of the world.

## 5 Half-Baked vs. Deep-Fried: Comparing Friction Types

We have hypothesized that certain forms of boundary friction are inherently valuable, insofar as they signal not so much a lapse in combinatorial finesse by a computational system, but a conceptual challenge that draws an audience into the creative act itself. Superficial friction typically arises either as a loose-end – a flaw overlooked by a careless or inept creator that detracts from his or her achievement – or as a weak attempt at generating a mere frisson of semantic tension. Deep friction, like a gnashing of conceptual gears, is both surprising yet appropriate, a piquant insight emerging from a plate of bland generalizations. We test this claim in the context of inferential cut-ups by using a crowd-sourcing platform to elicit the feedback of human judges. But we do not present judges with the polished outputs of *@MetaphorMagnet*, as this twitterbot carefully packages its tweets in a variety of rhetorical guises that may influence a judge’s perception of the underlying reasoning. Rather, judges are presented with inferential chains in unpolished, quantitative forms, so e.g. the chain *dictators suppress the critics that criticize the artists that produce art* is presented simply as *more dictators → fewer critics → more artists → more art*.

Inferential chains are randomly generated from the underlying knowledgebase to support three different test conditions. In the first condition, the *random* condition, inferential chains of connected generalizations are produced without regard to whether they contain any kind of friction at all. These chains are logically well-formed, but no interestingness criterion is applied as a filter. In the second condition, the *surface* condition, chains of connected generalizations are produced so that the first and last concepts in the chain exhibit a polar opposition in sentiment, such as *love vs. war* or *art vs. death*. Any friction in these chains is most likely the product of superficial differences in sentiment between connected concepts, rather than of a surprising twist of causality. In the third condition, the *deep* condition, inference chains are produced that exhibit precisely this kind of causal friction, so that a positive idea (such as *more beauty*) is shown to indirectly have negative consequences (such as *more war*), or a negative idea (such as *more war*) is shown to indirectly have positive consequences (such as *more patriotism* or *more prosperity*)

A pool of 80 inferential chains was randomly generated for each condition, 30 of which (per condition) were manually annotated as a gold standard to detect scammers, and 50 of which (per condition) were finally annotated by independent judges. The crowd-sourcing platform *CrowdFlower* was used to recruit a panel of 70 human judges to estimate, for each of these 3x50 inference chains, the degree of surprise exhibited by each. The full inferential pathway was presented in each case, so that judges could see not only its conceptual end-points, but the coarse logic at work in each link of the chain. The gold standard paths were used to detect unengaged scammers, resulting in 2.5% of judgments overall being discarded. Ultimately, 50 chains for each condition were judged by 15 or more judges, producing 765 judgments for the *random* condition, 751 for the *surface* condition, and 750 for the *deep* condition. Each elicited judgment provided a measure of surprise on a scale from 0 (no surprise) to 3 (very surprising) for a given inference chain.

The mean surprise value for the 765 judgments of the *random* condition is 1.06, that for the 751 judgments of the *surface* condition is 0.96, and that for the 750 judgments for the *deep* condition is 1.44. There is little here to distinguish the chains of the random condition from those of the surface condition, suggesting that combinations that rely only on superficial differences in form, rather than on deep differences in causality, fail to reliably elicit any interest. However, as we have hypothesized, there is a statistically significant difference in surprisingness between, on one hand, the chains of the *surface* and *random* conditions, and on the other, those of the *deep* condition. A one-sided Wilcoxon rank-sum test verifies that the increase in mean surprisingness from the *surface* and *random* conditions to the *deep* condition is significant at the  $p < .001$  level. To be surprised one must have prior expectations that are thwarted in some way. Surprise turns to fascination when initial dissonance gives way, after some consideration, to a deeper resonance of ideas. The inventors of the cut-up technique did not want to generate jarring combinations for their own sake, but to produce meaningful bisociations that break with banality. So to be appreciated as “cre-

*ative*”, combinatorial reuse systems must do more than seek out combinations that work; they must seek out combinations that work in spite of themselves.

## 6 Conclusions

Creative thinkers must constantly question received wisdom, and look beyond the superficial, isolated meaning of consensus beliefs. Creative thinkers actively seek out boundary friction whenever different texts, ideas, viewpoints or rules are combined, because this friction allows one to see wherever convention has fixed the acceptable boundaries of everyday categories. By identifying the boundaries of the conventional mindset, a creative thinker can *break set* and search for untapped value beyond those boundaries. Though certain forms of boundary friction are vexing issues to be resolved away, others offer a map to what is novel, unexpected and interesting.

The recombinant elements of past solutions rarely click together as cleanly as Lego bricks, but for that we should be thankful. Combinatorial reuse is a well-proven strategy for problem-solving that offers many opportunities for creative friction, at various levels of resolution. Reusable solution elements can range from individual words or concepts to heftier chunks of meaning or text to physical artefacts that exhibit different functionalities in different contexts. The artist provocateur Marcel Duchamp introduced the notion of a creative readymade to art when, in 1917, when he displayed a signed urinal (christened “*Fountain*”) at a Dadaist art exhibition in New York. Duchamp shifted the emphasis of artistic production from the act of original generation to the act of aesthetic selection, arguing that the artifacts at the core of the artistic process need not be constructed by the artists themselves ([?]). Duchamp used physical objects, constructed by skilled artisans, as his “*readymades*”, but Burroughs and Gysin were to show that any object, physical or textual or conceptual, could be creatively re-purposed and re-used as a readymade in a new context. By viewing the phrasal elements of a text database such as the Google n-grams ([?]) as a database of reusable elements that gain new meanings in new contexts, algorithms can create and exploit linguistic readymades of their own ([?], [?]).

Reusable elements may themselves be complex software components that strive for creativity in their own right. The creative friction that spurs innovation from antagonistic partners, of a kind that is well-attested by famous creative human partnerships, may also be evident in the qualities that arise from productive mash-ups of independent Web-services. A service oriented architecture ([?]) supports combinatorial reuse at multiple levels of resolution, and allows boundary friction to subtly influence the creativity of the end-result at each of these levels. For instance, other creative systems may use the outputs of the *@MetaphorMagnet* twitterbot as inputs to their own generative processes, or they may call directly upon the corresponding *Metaphor Magnet* Web-service to ask for specific kinds of metaphors for specific topics ([?], [?]). Web-services allow us to convert the generative components of past successes into stand-alone, modular, reusable, discoverable, generative engines in their own right. These

services will exploit their own internal forms of boundary friction, much as the *Metaphor Magnet* Web-service exploits its own, intra-service notion of semantic tension to generate exciting new metaphors. But they will also give rise to higher-level, inter-service frictions that are harder to predict but just as useful to exploit. In a thriving ecosystem of competing Web services, creativity will arise not just from the carefully planned interactions of these services – whether for metaphor generation, language translation, story-telling, affective filtering, multimodal conversion, and so on – but from the unplanned frictions that emerge from the less-than-seamless integration of these services. In such complex systems of imperfectly described, underspecified modules, the potential for boundary friction abounds at every level, offering opportunities for creativity all the way down. William Burroughs viewed his scissors as a deliberate instigator of boundary friction. We too can harness boundary friction to “*introduce the unpredictable spontaneous factor*” into our combinatorial reuse systems to make them more genuinely creative.

## 7 Acknowledgements

This research was supported by the EC project WHIM: The What-If Machine. <http://www.whim-project.eu/>

## References

1. Veale, T.: Creativity as pastiche: A computational treatment of metaphoric blends, with special reference to cinematic “*borrowing*.” In: Proceedings of the Mind II: Computational Models of Creative Cognition, Dublin, Ireland. (1997).
2. Robinson, E. S.: Shift Linguals: Cut-Up Narratives from William S. Burroughs to the Present. Rudolphi, Amsterdam, The Netherlands. (2011).
3. Sumita, E., Iida, H., Kohyama, H.: Translating With Examples: A new approach to Machine Translation. In: Proceedings of the 3rd International Conference on Theoretical and Methodological Issues in Machine Translation of Natural Language. Austin, Texas, USA. (1990).
4. Veale, T., Way, A.: Gaijin: A template-driven bootstrapping approach to example-based machine translation. In: Proc. NeMNLN Conference on New Methods in Natural Language Processing. Bulgaria. (1997).
5. Way, A., Gough, N.: wEBMT: developing and validating an example-based machine translation system using the world wide web. Computational Linguistics 29 (3), pp 421–457. (2003).
6. Pollio, H. R.: Boundaries in humor and metaphor. In: Mio, Jeffery Scott and Katz, Albert N. (eds.) Metaphor, Implications and Applications. pp 231–253. Mahwah: Lawrence Erlbaum Associates (1996).
7. Raskin, V.: Semantic Mechanisms of Humor. Dordrecht: Reidel. (1985).
8. Attardo, S., Hempelmann, C. F., Di Maio, S.: Script oppositions and logical mechanisms: Modeling incongruities and their resolutions. Humor: International Journal of Humor Research, 15(1), pp 3–46. (2002).
9. Koestler, A.: The Act of Creation. Hutchinsons, London (1964).

10. Veale, T., Li, G.: Creative Introspection and Knowledge Acquisition: Learning about the world thru introspective questions and exploratory metaphors. In: Burgard, Wolfram and Roth, Dan (eds.) Proc. of 25th AAAI International Conference of the Association for the Advancement of AI. San Francisco, California (2011).
11. Veale, T.: Exploding the Creativity Myth: The computational foundations of linguistic creativity. London: Bloomsbury Academic. (2012).
12. Brants, T, Franz, A.: Web 1T 5-gram Ver. 1. Linguistic Data Consortium (2006).
13. Hofstadter, D. R.: Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought. Basic Books, New York, NY, USA. (1995).
14. Taylor, M. R.: Marcel Duchamp: Étant donnés. (Philadelphia Museum of Art). Yale University Press. (2009).
15. Veale, T.: Linguistic Readymades and Creative Reuse. Transactions of the SDPS: Journal of Integrated Design and Process Science, 17(4), pp 37–51 (2013).
16. Erl, T.: SOA: Principles of Service Design. Prentice Hall (2008).
17. Veale, T.: A Service-Oriented Architecture for Computational Creativity. Journal of Computing Science and Engineering, 7(3), pp 159–167 (2013).
18. Veale, T.: A Service-Oriented Architecture for Metaphor Processing. In: Proc of the 2nd ACL Workshop on Metaphor in NLP, at ACL 2014, the 52nd Annual Meeting of the Association for Computational Linguistics, Baltimore, USA. (2014).

## Appendix: Online Resources

*Metaphor Magnet* exists as both a Twitterbot, which pushes metaphors of its own invention, and as a public Web-Service, which generates specific families of metaphors for specific topics, on demand, via a HTML and an XML Web interface. It also provides other creative products, such as poems and blends, for these metaphors. This service can be accessed at: <http://boundinanutshell.com/metaphor-magnet-acl>