Picasso, Pato and Perro: Reconciling Procedure with Creativity

Patrick Summerhays McNally
Computer Science Department
Northwestern University
patrickmcnally2013@u.northwestern.edu

Kristian Hammond
Computer Science Department
Northwestern University
hammond@cs.northwestern.edu

Abstract

This paper presents and details 'Pato and Perro on the Movies,' a system that generates web comics about recently released movies. Information is extracted about movies from the internet and a series of panels are drawn with dialogue to set up a punchline in the comic's final frame. Definitions of creativity commonly used to examine computational processes are presented and used to examine this system. The system is used to discuss a common critique of creative systems, namely that procedural creation inherently limits the range of potential content produced. This paper argues that procedure and creativity can be reconciled, and that much of the content produced by humans is subject to similar critique. Finally, we discuss the implications of characterizing many human acts of creation as procedural.

Pato and Perro on the Movies

'Pato and Perro on the Movies' is a web comic written and drawn by a machine. Each comic gives bite sized impressions of a recently released film and closes with a punch line about one of the characters' mothers. Consider a comic produced by the system in Figure 1.

The system's procedure for creating a comic tries to meet these requirements:

-The subject must be a recently released film.
-Both a negative and a positive statement about the film should be presented if at all possible.
-The statement in the second panel should be fertile language for generating a punchline.
-That punchline should aim to be a humorous re-interpretation of language seen in the second panel.

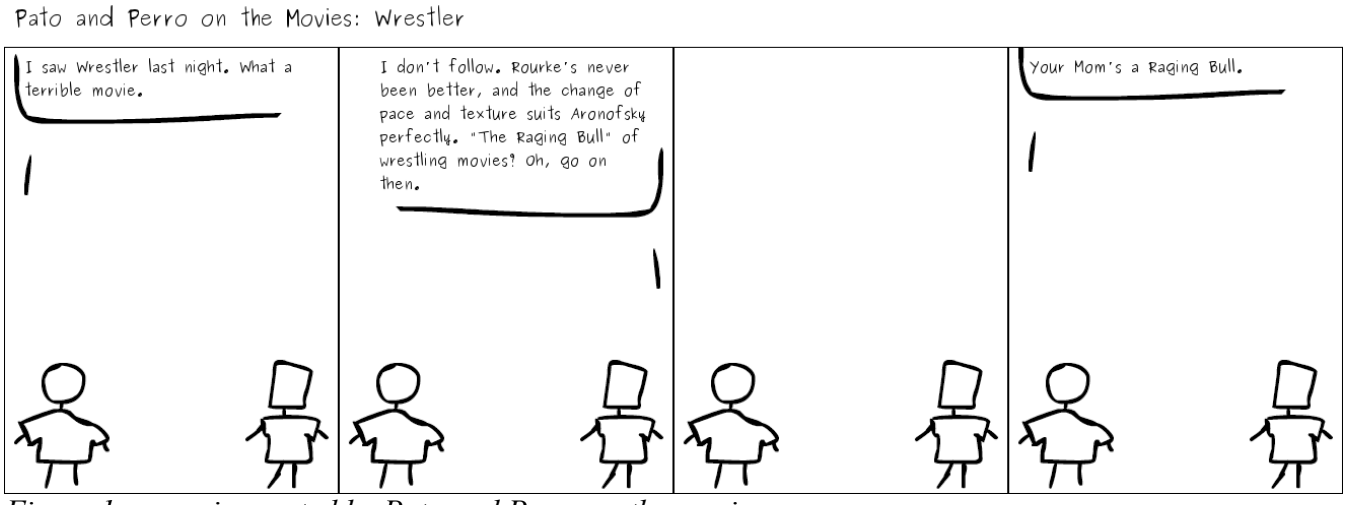

Figure 1: a comic created by Pato and Perro on the movies

Figure 2: On the left, the box office report on rottentomatoes.com. On the right, review snippets from Battlefield LA from the same site.

To accomplish these goals, the system mines movie reviews from rottentomatoes.com, a popular movie review site, for the week's top box office hits as well as bite-sized commentary of both positive and negative sentiment pertaining to these movies (see Figure 2). The bite-sized commentary is provided on each films' Rotten Tomatoes page. A section of each page is devoted to a wall of choice snippets from different reviews. The snippets are marked with a ripe tomato when the sentiment of the snippet is positive and a green splat when the sentiment is negative. This marker is used to determine the valence for a snippet when it is used to create dialogue in the comic.

The snippets are also analyzed using WordNet (Fellbaum 1998) for potential humorous meanings. The system's best pick is selected for the second panel, and another comment of opposing sentiment is selected to provide contrast in the first panel, or a flat statement of preference like the one seen in Panel 1 of Figure 1 is used. Tension is interesting and plays an important role in the build up to the punchline (Napier 2004) so it is important that the first two panels contain statements of opposing sentiment. The final panel receives the punchline.

To craft a punchline from text, the system follows this approach:

> -tokenize and tag the text that appears in the second panel of the comic
> -identify verb and noun phrases that occur in the text
> -compute the wordnet distance between the adjectives, verbs and nouns of a phrase to a set of humorous concepts pre-assembled by hand.
> -pick the verb or noun-phrase with the shortest distance to this target set of humorous concepts
> -form the punchline based on whether or not the phrase is a verb or noun phrase

The system relies on WordNet, a precompiled database of words and their semantic distances from one another. The method used to compute semantic distance was developed by Philip Resnik (Resnik 1995). This distance is used to determine how easily a phrase can be interpreted as something humorous. Put another way, the system uses a word's WordNet semantic similarity with another word as a proxy for how likely those words can be conflated in meaning.

For example, in the comic illustrated in Figure 1, a bull is semantically similar to a donkey, which is considered a humorous topic appropriate for a punchline. The system, in evaluating the words in panel two, selected 'raging bull' for this semantic similarity and crafted the punchline.

The result of the system's procedure is a comic strip. If a person had assembled this strip we would call it creative. Can we say the same of this system?

What is Creativity

Here is a general definition of creativity:

> *the production of something novel and useful.*

There is a certain amount of consensus around this definition (Mumford 2003) such that any theory or model of creativity needs to address these two tenants (novelty and utility).

Although consensus is good, the scope of this definition is quite broad and this generality comes with limitations. It becomes difficult to discuss how creative processes differ without a more precise definition. For example, how should the difference between a baker modifying a recipe and a baker working from scratch be characterized? This is where the consensus dissolves.

Given such a lack of consensus around core concepts, claims of creativity in computation have to be made with a spirit of exploration. Never-the-less, two definitions of creativity have been defined with a strong computational con-

text. Both are useful for discussing the creativity of a machine. They strike out in remarkably different directions. One is old and one is new. One is content-centric and the other is process-centric. Despite their apparent differences, both tell us fundamental things about our ideas of creativity and machines.

One of the earliest definitions used in the computational creativity field was put forth by Newell, Shaw and Simon in 1963. Their four criteria for a creative system are:

1) the solution is novel and useful
2) the solution demands the rejection of previous ideas
3) the solution occurs after much persistence
4) the solution should clarify an initially vague problem

Understandably, their criteria are framed with the task of searching in mind. Problem solving systems of the time were designed to strategically search a vast solution space. So, searching was a common paradigm upon which to base a system. These criteria were intended to judge when a solution found by such a system would have appeared creative had it been developed by a human mind.

To this end, Newell, Shaw and Simon developed this list (which was neither meant to be complete nor sufficient) to provide a scaffolding around which to organize an argument that a system exhibited creativity. The first criterion is simply a repetition of the general definition above. The second speaks to a sense that creative things often elicit surprise; their implications are unexpected in some way. The third criterion suggests that creative things are not easily arrived upon. The fourth criterion reflects the sense that creativity should not be obvious. In fact, creativity will often provide an 'ah ha!' moment, where the solution itself illuminates the structure of the very problem it solves.

Like the original definition, there is a tremendous amount of room to interpret these criteria. But they are meant to be interpreted. Each one indicates characteristics considered to be markers of creativity in human acts. Attributing each one to a system that produces content is a way to argue the system is creative.

One potential problem with these criteria is how content centric they are. They stipulate nothing about the architecture of a system or procedure making the content.

Fortunately, a second perspective on computational creativity is more process-centric. Margaret Boden has published ideas about different classes of creativity. And these categories pertain more to process than product (Boden 1998). Two important classes of creativity she discusses are exploratory creativity and transformational creativity. These distinguish between creativity that follows a procedure and creativity that defines a procedure. Put another way, Boden tells us that exploratory creativity navigates a defined space while transformational creativity redefines that space. A painter adhering to the process and techniques of watercolor, for example is different from Picasso establishing cubism.

A watercolorist will start in the background, and 'block' in their scene, moving from lightest to darkest colors. They will end with the foreground and employ a set of stroke techniques defined by the medium. For this reason a classic watercolor scene has a perspective and texture that is uniform across the medium. There is a recognizable proportion and boundary to watercolor as a space of potential content. Boden would describe the creativity of a painter following the classic watercolor procedure as 'exploratory.' The boundaries of the conceptual space are the medium and the procedure with which that space is explored are the stroke techniques and layering strategy. In contrast, any process that redefines an existing conceptual space or establishes a completely new one, like Picasso developing cubism, is an example of transformational creativity.

Transformational creativity is more complex. The minds we most celebrate as creative tend to be known for their grand acts of transformational creativity, but there can be minor acts too. If the watercolorist (noticing the way a heavy brush drips and splatters on the paper while traveling to begin the first stroke) decides to employ splatter as a technique, the space of potential content changes; A new technique has been added to the painter's procedure. This modification would also be transformational creativity.

These two characterizations of creativity prove useful in discussing whether or not (or to what degree) a machine exhibits creativity. Attempting to answer such a question forces clarifications about what it means to be creative. Which in turn may have important implications for our understanding of human creativity, and the further development of machines capable of entertaining and illuminating our lives.

Specifically, these two definitions of creativity can inform our understanding of Pato and Perro and help us to articulate the system's limitations as well as imagine beyond them.

Creativity in Pato and Perro

Pato and Perro produces content that is variable. Decisions are being made by the system as the space of possible constructions is examined. The case can be made that Pato and Perro meets three of the four criteria laid out above by Newell, Shaw and Simon.

Surely the comics could be considered novel, since they are new pieces of media. It can be argued that they have use as well, since humor often has a beneficial effect on mood, and each strip conveys some amount of information about a movie.

If the system succeeds, the reader will interpret the text in the second panel and then be forced to reinterpret the text again after reading the punchline, forcing a humorous rejection of the original meaning. This arguably satisfies the second criteria.

The third criteria seems to be rather subjective. Depending on the number of review snippets available, the system may select a comic it considers best from hundreds of potential candidates. If this isn't enough persistence, more

snippets could be retrieved and more potential comics could be evaluated.

The final criteria, which stipulates that the comic should clarify an initially vague problem, is perhaps inapplicable here. It is not clear what problem is being clarified.

Now, there is a lot of room for interpretation with these criteria. For example, consider the modern search engine. A results page for a given search term will be novel and useful because the page will present relevant, recent results catered to the user's language and geographic location. The results may demand the rejection of the user's previous ideas, be those ideas about the cheapest flights or the definitions of words. Certainly a good results page will present the user with varied perspectives on a topic. The results may arrive quickly, but vast amounts of computational power has gone into creating the indices that afford this speed, and these indices are constantly being revised. Finally, the best results set for a given search term is often unclear. The best engine should clarify this through its ranking mechanism; In other words, the best engine will have a strategy for determining why one set of results is superior to another.

A search engine appears, from this perspective, to be a creative system.

Rather than defend the previous statement, we will simply say that Newel, Shaw and Simon's criteria are open for interpretation. Furthermore it may not be appropriate to attribute creativity to a system in hindsight, looking only at it's solutions and our impressions of those solutions. Which means Pato and Perro's claim to creativity based on these criteria is debatable.

This is where Boden's conceptions of exploratory and transformational creativity can be of use. For example, if we examine Pato And Perro's process it becomes clear that the system exhibits exploratory creativity but lacks transformational creativity, since the system follows rules and has no mechanisms for breaking or amending those rules.

But what would it mean for the system to break it's own rules? The system has only a general idea of the pieces with which it works. Essentially, there is a plan (the overarching comic structure of contradicting opinions followed by a punchline) and potential chunks of data that could fit into that plan (the review-bites from rottentomatoes.com). In a general sense, the system ranks potential arrangements according to a fitness metric and arrives at the one these metrics deem best. Each review-bite has a sentiment valence and character length attached to it as the only orienting information the system uses to realize its plan. The most technical piece of the system is used to evaluate each chunk of review for potential humorous meanings.

In order to break with it's own procedure, the system would need a way to examine all these pieces and the motivations for how they currently fit together. But the system doesn't do this; The process described here is the procedure that Pato and Perro *always* executes.

The problem for many critics is the static nature of this procedure. There are plenty of procedural art forms, where the artist defines rules and then produces a work by following these rules, but even the art world feels vaguely uncomfortable with these (though they too have trouble putting their fingers on a why). There is a sense that creative processes should not be static, that creativity means being capable of stepping outside of one's process and implementing changes—what Boden would call transformational creativity because it would redefine the space of potential content.

This should be vaguely reminiscent of John Searle's thought experiment with the Chinese Room (Searle 1980). There is a strong parallel between following a procedure to create content (what we are calling exploratory creativity) and the Chinese Room producing believable conversation. For a system to exhibit transformational creativity it would likely need to actually understand the process it was executing. It may be that transformational creativity mirrors in difficulty what Searle called 'strong AI.' Which begins to suggest why following rules is so much easier than breaking or defining new ones for a system designed to create content.

In summary, examining only the content produced by Pato and Perro, it can be argued that the system is creative, but upon further examination, such hindsight assessments may be subject to the same sort of critique put forward for strong AI with the Chinese Room–meaning creativity has something to do with the process as well as the product. In fact, when we actually examine the system's process, it becomes clear that the creativity at work is perhaps the weaker of two forms defined by Boden.

Discussion

For now, if a machine is to produce compelling content the details of a compelling production process have to be determined. This means systems are built from their medium up; they are envisioned with their final products in mind and built to execute those products. This requires that the architect understand the nature–the what and why behind interestingness–of the content to be produced. So the architect of a creative system must understand the potential space of content to be explored and instill in the machine's procedure constraints or heuristics that will keep the machine within the interesting areas of that potential space.

For Pato and Perro, the procedure consists of building tension through disagreement in the first two panels and establishing an alternate interpretation for text appearing in the second panel. This implicitly says something interesting about the creative process behind the content: tension heightens effect and ambiguous meanings provide entertainment. The system doesn't have a macro-level process to examine this like a mind would, nonetheless these values are implicit in the system's procedure.

This often leads to the critique that the creativity happens outside the system. But it isn't clear why proceduralizing an act of creation negates, or rather separates it's creativity. If establishing a process abolishes creativity, many human pursuits generally considered to be creative should be reevaluated.

Almost every amateur taking art lessons has not yet been creative. Reporters adhering to the strictures of their form are not being creative. Students following a prescribed essay structure are not being creative. Programmers executing a well established architectural paradigm are not being creative.

Pato and Perro as a system, is like the novice painter practicing the process they've learned to paint with watercolors. To repeat, the medium provides a conceptual space and the learned strokes and techniques are the procedures with which the painter navigates that space. This paper is not claiming the person painting and the system creating comics are equivalent cognitively, only that if the painter sticks to a defined set of techniques their creations will be limited to a certain range of possible content in a fashion similar to that of systems like Pato and Perro. If Pato and Perro could examine its own process and establish new patterns of creation, it would exhibit the transformational creativity of Picasso inventing Cubism or the painter discovering a spatter technique; transformational creativity redefines the procedures and/or the conceptual space.

So why don't these systems exhibit transformational creativity? If the problem seems to be rooted in a sense that a static procedure isn't creative, why not make the procedure dynamic? Why are many 'exploratory' systems being built while so few 'transformational' systems are coming about, particularly if the latter seems to be more powerful or more highly regarded in some way?

The list of exploratory systems and their mediums is extensive. Jape builds linguistic puns (Binstead 1994). HAHACRONYM builds acronyms (Stock 2003). AARON builds paintings (McCorduck 1991). ASPERA builds poetry (Gervás 2001). Tale-spin builds stories (Meehan 1981). The explanation is likely that exploratory systems are just easier to build. It is easier to become a watercolorist adept at understood techniques than it is a Picasso capable of inventing new styles.

Returning to Searle's Chinese Room, the distinction between following procedure and understanding the procedure may actually be endemic to people as well as machines. A novice painter could study the stroke techniques of masters for a lifetime and they still might never recognize the rules being followed and transform them like Picasso. 'Understanding' the procedure, we argue, is often quite difficult even for humans.

Within their domains computationally creative systems have instructions and constraints that allow them to explore but they do not exhibit a mastery over their own process. They lack the higher perspective to modify their own production behaviors. They don't exhibit Boden's transformational creativity. Never-the-less, these system's successes suggest that for the purposes of creating compelling content of a fixed type for an audience, exploratory creativity may be sufficient. Exploratory creativity seems to excel when leveraging a specificity of domain regarding medium or form.

But critics might still claim exploratory creativity will never surprise and delight like transformational creativity. And they would be right. Systems exhibiting exploratory creativity may certainly provide useful, even valuable, material but the human mind, once it has grasped the system's process, will naturally imagine the boundaries and proportions that limit that process. And this will always feel disappointing.

But the fact remains that much of the content we produce is procedural in nature. Exploratory creativity is the most common type exhibited by people, day to day. Our most celebrated works are certainly the result of transformational creativity but our world spins by procedure.

Conclusion: Creative Systems in the World

The Pato and Perro system produces multimedia content that has never been seen before. It does this quickly and at scale. Most importantly, it is not alone. Automated content generation systems like it are beginning to emerge for public and commercial use.

Music generation systems are being seen on popular mobile devices like the iPad (Eno 2010). Systems are being built to craft sports narratives (Carr 2009). Investigative journalism is using systems to identify trends and soon perhaps even characterize them in language. All of these systems fit the general pattern of exploratory creativity, that is, they have a massive potential space for creation and instructions for how to position, navigate and orient themselves in this space.

Any category of content that is produced for a massive audience, and adheres to a procedure, has the potential to be produced by a machine. But it is unlikely that machines will be defining any new mediums or breaking the rules of old ones. For now it is unreasonable to expect our machines to be Picassos that transcend their procedures, but systems like Pato and Perro, Aarron, Jape and many more show us we can reasonably expect adepts of a specified medium.

Acknowledgements

We would like to acknowledge support from the National Science Foundation. Computational Creativity: Building a model of machine generated humor Grant Number: 0856058

References

Boden, M. 1992. *The Creative Mind*. London: Abacus.

Margaret A. Boden. 1998. Creativity and artificial intelligence, Artificial Intelligence, Volume 103, Issues 1-2, Artificial Intelligence 40 years later, August, Pages 347-356, ISSN 0004-3702

Carr, David (2009). “The Robots Are Coming! Oh, They’re Here.” the New York Times. http://mediadecoder.blogs.nytimes.com/2009/10/19/the-robots-are-coming-oh-theyre-here/

Christiane Fellbaum (1998). WordNet: An Electronic Lex-

ical Database. Bradford Books.

Cope, David (2006), *Computer Models of Musical Creativity*, Cambridge, MA: MIT Press

Eno, Brian (2010), Air, iOs application, http://itunes.apple.-com/app/air/id312163985?mt=8

Gervás, Pablo (2001), *An expert system for the composition of formal Spanish poetry*, Journal of Knowledge-Based Systems 14(3-4) pp 181–188

Kim Binsted, Graeme Ritchie (1994) "A symbolic description of punning riddles and its computer implementation." Research Paper 688, University of Edinburgh, Edinburgh, Scotland, 1994

McCorduck, Pamela (1991), *Aaron's Code*., W.H. Freeman & Co., Ltd.

Meehan, James (1981), TALE-SPIN, Shank, R. C. and Riesbeck, C. K., (eds.), *Inside Computer Understanding: Five Programs plus Miniatures*. Hillsdale, NJ: Lawrence Erlbaum Associates

Mumford, M. D. (2003). Where have we been, where are we going? Taking stock in creativity research. Creativity Research Journal, 15, 107–120.

Napier, Mick. 2004. *Improvise: Scene from the Inside Out*. Heinemann, Portsmouth, NH

Newell, Allen, Shaw, J. G., and Simon, Herbert A. (1963), *The process of creative thinking*, H. E. Gruber, G. Terrell and M. Wertheimer (Eds.), Contemporary Approaches to Creative Thinking, pp 63 – 119. New York: Atherton

Resnik, P. (1995a). Disambiguating noun groupings with respect to wordnet senses. In Third Workshop on Very Large Corpora. Association for Com- putational Linguistics.

Searle, John (1980), "Minds, Brains and Programs", Behavioral and Brain Sciences 3 (3): 417–457

Stock, Oliviero, Strapparava, Carlo (2003), *HAHAcronym: Humorous agents for humorous acronyms*, Humor: International Journal of Humor Research, 16(3) pp 297–314