The Computational Creativity Continuum

Rafael Pérez y Pérez
Departamento de Tecnologías de la Información
División de Ciencias de la Comunicación y Diseño
Universidad Autónoma Metropolitana, Cuajimalpa
Av. Vasco de Quiroga 4871, Ciudad de México, C. P. 05348
rperez@correo.cua.uam.mx, www.rafaelperezyperez.com

Abstract

This paper argues that the construction of creative systems is motivated by –what sometimes seems to be–, diverse, even contradictory, viewpoints and understandings about the goals of computational creativity. To analyse those differences I introduce the Computational Creativity Continuum (CC-Continuum), which can be pictured as a line or band flanked by two poles; I refer to one of the poles as the engineering-mathematical approach and I refer to the opposite pole as the cognitive-social approach. Thus, creative agents are located along the Continuum based on their main goals as systems. Through the text I explain the general characteristics of each approach, how they complement each other, and some of the difficulties that arise when systems are misclassified. I finish pointing out the utility of frameworks like the CC-continuum.

Introduction

The capacity of developing artificial creative agents is an old dream. Some of the oldest systems were codified more than 58 years ago. For example, as part of his research in México City, the linguistic Joseph E. Grimes developed in 1960-1961 the first known plot generator (Ryan 2017). It took some time to the scientific community to start meeting regularly to discuss about the possibilities of this emergent field. The first International Conference on Computational Creativity was organised in 2010; it was preceded for 10 years of workshops. Through all these years, several systems have been developed, each one contributing to progress different aspects of the field. These works have played an important role in the development of theoretical ideas and practical perspectives about computers and creativity.

Some of such systems share important features while others employ methodologies and knowledge structures that, sometimes, appear to represent opposite views about computational creativity (CC). Naturally, this seemingly contrary perspectives are reproduced in some of the definitions that have been proposed recently. I would like to analyse two of them.

For Colton and Wiggins, computational creativity is the study and simulation, by computational means, of behaviour, natural and artificial, which would, if observed in humans, be deemed creative (Colton and Wiggins 2012). As Jordanous points out, from this perspective "the challenge is to engineer a system that appears to be creative to its audience, rather than engineering a system that possesses a level of creativity existing independently of an audience's perception" (Jordanous 2012). In general terms, this sort of approach employs mathematical models and engineering methods. In contrast, Pérez y Pérez defines computational creativity as the interdisciplinary study of the creative process employing computers as the core tool for reflection and generation of new knowledge (Pérez y Pérez 2015). This perspective accentuates the importance of contributing to the understanding of the creative process. In general terms, this approach is motivated by the work of philosophers, sociologists, cognitive psychologists, and so on. In this way, the engineering-mathematical perspective concentrates on the construction of products that are appealing for an audience while the cognitive and social point of view privileges the generation of models that produce insights about the phenomenon we are studying. I employ these two stances, generation vs. understanding, as the two poles of what I refer to as the Computational Creativity Continuum (CC-Continuum); (see figure 1); (I first published the idea of the continuum in Ackerman et al. 2017).

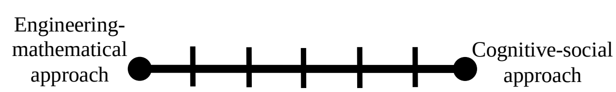

Figure 1. The CC-Continuum.

The CC-Continuum provides a framework that allows comparing creative agents. The descriptions of the engineering-mathematical and the cognitive-social approaches employed in this text should be understood in a broad sense; rather than providing precise definitions my goal is to create a context for discussion. I chose the names of the poles following Jordan and Russell (1999) description of AI:

> There are two complementary views of artificial intelligence (AI): one as an engineering discipline concerned with the creation of intelligent machines, the other as an empirical science concerned with the computational modelling of human intelligence. When the field was young, these two views were seldom distinguished. Since then, a substantial divide has opened up, with the former view dominating modern AI and the latter view characterizing much of modern cognitive science. (Jordan and Russell 1999, p. LXXIII).

The same Russell had previousy talked about this distinction in AI, i.e. the engineering-mathematical approach and the cognitive approach, in his famous book (Russell and Norvig 1995).

Engineering-Mathematical Approach

Traditionally, the engineering-mathematical approach uses optimization techniques like genetic algorithms; probabilistic techniques like DNN; logic and problem solving techniques; and so on. Usually, agents are built based on one of these procedures, although one can find programs that mix two or more of them. None of these methods have been developed with the explicit purpose of producing creative systems; they can be described as general purpose tools. Researchers have figured out how to manipulate them to develop computer programs that produce the desired results. For instance, the main challenge of those using genetic algorithms in the visual arts is to figure out a fitness function that drives the search into reaching interesting products.

Some researchers have come out with clever ways to exploit the existing resources. Heat and Ventura (2016b) report using in their system Darcy the gradient ascent method (Simonyan, Vedaldi, and Zisserman 2013). The gradient ascent employs, for example, a trained DNN for face recognition:

> [It] starts with a random noise image and tries to maximize the activation of the output node corresponding to the desired class to generate. The network then backpropagates the error into the image itself (keeping the network weights unchanged) and the image is slightly modified at each iteration to look more and more like the desired class (Heath and Ventura 2016b).

Because this is a general purpose tool, this technique seems useful in diverse creative domains. An important challenge that researchers using DNN face is the construction of a, sometimes, very complex process that needs to be applied in the training data before it can be useful.

Problem solving techniques characterize another popular approach which, in general, is described as goal-oriented reasoning. I have found that, when they are used in the context of CC, these type of methodologies tend to employ knowledge structures that somehow assure in advance the coherence of the final product. A typical example is the use of grammars (e.g. story-grammars, shape-grammars), although other types of predefined structures are also employed. For instance, one of the core problems in narrative generation is to progress coherence sequences of actions. Some researchers have faced this challenge using predefined structures like scripts, schemas, productions rules with elaborated preconditions, or even templates. Thus, the development of a plot consists in satisfying a set of characters' goals and/or authors' goals, which are represented by any of such structures. A goal is reached by instantiating partially filled schemas, finding actions that satisfy unfilled preconditions, and so on. Researchers that employ this kind of approach focuses on building schemas that represent core features of the work in progress, or in creating rules that chain with each other with the purpose of producing interesting outputs, and so on.

Thus, researchers working from the engineering-approach side of the CC-Continuum spend a vast quantity of time and energy performing technical tasks that allow facing research questions like, how can I develop mechanisms to produce pieces that are appealing for a given audience? How can I produce systems that explore unfamiliar domain spaces? And so on.

Cognitive-Social Approach

The cognitive and social inspired approach employs studies on human behaviour as basis to develop computer models of the creative process; such models are tested as running programs that works as prototypes. The main purpose of the cognitive and social inspired approach is to attempt to contribute to answer questions like: How do we get new ideas? How can we produce coherent sequences of actions during the creative act? How do we assess the quality of a piece? How does the collaboration of multiple agents shape the creative process? How can we represent in computer terms the role of the social-environment during the creative process? And so on. The systems included in this approach goes from those that reproduce the results of behavioural experiments performed by psychologists, e.g. tests that evaluate the subjects' responses to different stimuli, to those that are based on general cognitive theories or even cognitive accounts of the creative processes. For the last ones, it is the work of cognitive and social psychologists to test in humans how accurate the conclusions emerged from this programs are. In all cases, this approach only generates potential explanations about some aspects of how creativity works in humans.

From the algorithmic perspective, the discourse employed by philosophers, sociologists, cognitive psychologists, and so on is, in many cases, excessively general. That is, it lacks details about the processes and knowledge structures involved in the creative process that are necessary for the development of a computer model and its implementation.

Thus, one of the main tasks of researchers operating in this side of the CC-Continuum is to find ways of representing in computer terms relevant cognitive, cultural and social behaviours. This task is challenging for several reasons; the most important is that we hardly understand how many of such behaviours works in our mind. Social norms illustrate this condition. They dictate the acceptable ways of acting within a group; e.g. most societies classifies killing as a pursuit rejected by the community. However, the reality is that people's reaction to such a conduct change based on the circumstances; killing an individual might produce a hero, a villain, a hero than later is considered a villain or vice versa, or even divide people's opinion about the fact, i.e. the perpetrator might embody a hero and a villain at the same time. To design a computer representation of social norms that comprises all (or most of) these aspects is a complicated task. However, this type of information is needed by creative agents working individually or collectively, when social representations play an important role in the model (e.g. plot generators).

Another issue that those working in this approach pay attention to is how knowledge structures and cognitive process relate to each other and enforce creativity. For instance, I have claimed that, besides of being able to generate novel, coherent and interesting (or useful) products, a creative agent must be able to: 1) Employ a knowledge-base to build its outcomes; 2) Interpret its own outputs in order to generate novel knowledge that is useful to produce more original pieces; 3) Evaluate its own products; such an evaluation must influence the way the generation process works (Pérez y Pérez & Sharples 2004; Pérez y Pérez 2015).

I have expand these ideas to define cooperative creative systems: if a piece generated by collaborative agents cannot be developed by any one of them alone, and such a piece generates original structures within their knowledge base that can be employed by the contributors to produce new outputs, then it is referred to as a collectively-creative work (Pérez y Pérez 2015). In this way, the analysis and design of a cognitive model must be shaped by the necessity of producing plausible explanations about issues like how predictable is the outcome, how the system progress a piece, how the system maintains the coherence, interestingness and novelty of the piece in progress (Pérez y Pérez 2004)...

The next step is to figure out how to develop the algorithms and knowledge structures that represent all these processes; such representation should be as close as possible to the knowledge, theories or hypothesis we have about human behaviour. The researcher might develop new techniques, or employ those that already exist, to achieve this goal. In several occasions, the elaboration of the first prototypes makes evident the deficiency on the theoretical framework used to construct the system. Then, it is necessary to design new routines that fill those gaps in the theory.

Comparing both approaches

The CC-Continuum provides a reference for comparing systems. Base on its position in the continuum, one can infer the general purpose of a system, the kind of routines that it might perform, the type of features to be considered in order to assess the work and its results, the perspective that the creator has about the field… On the other hand, the Continuum does not reflect aspects like the technical complexity of the design and implementation of a prototype, the quality and originality of the program and its outputs, the impact of the system in the community and the general public, and so on. In this way, if an agent located close to the engineering approach uses the technique X for a given problem, one expects to learn why X is more efficient than techniques Y or Z. If an agent located close to the cognitive-social approach uses the technique X for a given problem, one expects to learn why X represents better a specific cognitive or social phenomenon than techniques Y or Z.

Sometimes, a system located towards the engineering approach requires to characterize some kind of cognitive or social behaviour. There are different ways to achieve this goal, e.g., employing productions to accomplish a particular behaviour: "If character A kills character B then character A is sent to jail." This rule does not embody the complications, explained earlier, contained by human social norms. However, it might help to provide the illusion that the system represents such a complexity and therefore to influence the audience's judgment about the output. The designer of this hypothetical system may perhaps decide to add more and more productions in order to attempt to build a more robust version of public conduct. In this case, the location of the system starts to move towards the right side of the Continuum.

In most cases, the implementation of systems located towards the cognitive-social approach requires the development of software that, has nothing to do with the purpose of the model but, it is necessary to run the program. I refer to it as *infrastructure for the program*. A typical example is the construction of a knowledge base. Many creative agents require the use of knowledge structures in order to produce their outputs; however, in several occasions, such systems do not attempt to represent how an agent acquires its beliefs and experience.

An analogous case arise when the researcher does not have the cognitive or social understanding about how one of the procedures that comprise the whole creative process that she is representing works. There is a gap in her knowledge that needs to be filled. Sometimes, the designers simply cannot work those problems out and employ solutions that might be considered as more related to the engineering approach. That is, they use procedures that do not represent a cognitive or social phenomenon, but that help the prototype to work. I refer to them as *routines that support the model*. However, overall, the system should still represent a cognitive or social phenomenon.

Therefore, it is important that the researcher clearly differentiates which part of her program characterises the model of the creative process, which other part works as infrastructure for the program to run and which other parts play the role of routines that support the model. As the number of routines that support the model increases, the location of the system starts to move towards the left side of the continuum.

The reason why I chose a two-pole band as a framework for this analysis is because the engineering and cognitive-social approaches complement each other. Thus, we have hybrid systems, located towards the centre of the Continuum, that allow exploring possibilities that otherwise would be complicated to study. In the same way, the experience and knowledge generated along the Continuum provide useful information for the rest of the systems.

Discussion

Figure 2 locates on the CC-continuum some systems that have been labelled as creative. This is not an exhaustive list; it only attempts to illustrate a possible classification. Some authors might disagree with the location of their programs and I am happy to modify their position. Figure 2 includes the following systems: MCMC for Story Generation (Harrison et al. 2017), DARCY (Heath & Ventura 2016a), The Painting Fool (Colton 2012), ALYSIA (Ackerman and Loker 2017), MARBLE (Singh et al. 2017), WASP (Gervás 2000), Scheherazade (Boyang 2015), Metaphor Magnet (Veale 2015), systems of social creativity (Saunders 2018), DIVAGO (Martins et al. 2015), MEXICA (Pérez y Pérez 2001) and Tlahcuilo (Pérez y Pérez et al. 2013).

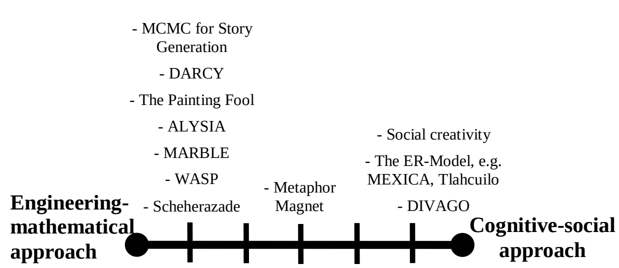

Figure 2. Systems on the CC-Continuum

Figure 2 suggests that the number of systems located towards the engineering-mathematical approach is much bigger than those located towards the centre or the right side of the Continuum. I believe that, in general terms, that information reflects what we see in the CC community. Thus, it is necessary to promote the development of systems that balance the distribution. The diversity of approaches and perspectives will enrich our field.

Can we compare systems that are positioned at different places along the CC-Continuum? It depends on the purpose of the comparison. Contrasting the qualities and limitations of programs along the Continuum would be beneficial, particularly for those new in the area, as long as the features and intentions behind each approach are clearly understood. Otherwise, confusions arise. The lack of knowledge about the Continuum might be the source of biased evaluations; the classic example is when a system that has been designed under the consideration of one of the approaches is assessed with the criteria of the other approach. As a result, a *Tower of Babel effect* is triggered, where researchers simply cannot communicate; they can only see a system from their own perspective without considering other possibilities.

I would like to illustrate this situation with a real experience. Some months ago, one of my students and I sent an expanded version of the Tlahcuilo system (Pérez y Pérez et al. 2013) for review to a journal. Tlahcuilo is a visual composer which is based on the ER-Model used to develop MEXICA (Pérez y Pérez & Sharples 2001), a system for plot generation. Its purpose is to contribute to the understanding of the creative process. Some reviewers' observations were very useful and constructive; nevertheless, there were others that clearly illustrated the Tower of Babel effect that I just mentioned. That is, Tlahcuilo, a model designed under the cognitive approach, was evaluated from an engineering perspective. In the following lines I reproduce some of such observations and add some comments about them.

> the paper must address in some detail the question of how their approach is different than evolutionary computation. The engagement seems very much like application of genetic operations and the reflection seems very much like computing fitness. Internal and external representations are very much like phenotypes and genotypes, respectively (anonymous reviewer).

It is hard for me to comprehend how a process that attempts to represent some of the core ideas described by some psychologists, philosophers.., about the creative process (i.e. the ER-Model) can be confused with a method for optimization (i.e. EC). I understand that, at the end, we are talking about computer programs, which are transformed into strings of zeros and ones. But systems based on evolutionary computing are hardly employed to represent cognitive processes (although, of course, EC can be used to develop routines that support the model). My interpretation is that the reviewer is trying to make sense of Tlahcuilo in terms of his/her area of expertise, i.e. EC, rather than from the authors' perspective.

> This model aims to be more interpretable than, for example, EC approaches, but that again limits it to interpretable actions and relations that must be supplied

> a priori, and that creates a tension between comprehensibility and utility (anonymous reviewer).

This comment probably represents the best example of the Tower of Babel effect. I am tempted to say that the reviewer, maybe without being aware of, disdains the importance of the cognitive-social approach and describes its main feature, the goal of offering explanations, as a drawback.

> The images demonstrated by the system do not compare well at all with images produced by other contemporary systems… Really, the system is still working at a very basic level, which is fine; there is value in looking at these kinds of toy examples, but it must also be able to demonstrate more interesting outputs if it is to support the claim of image composition (anonymous reviewer).

The reviewer focuses in comparing Tlahcuilo's outputs with the images of contemporary systems. There are two problems with this position. First, Tlahcuilo´s main goal is to produce insights about the creative process rather than producing astonishing illustrations. The reviewer clearly is not aware of it. He/she states that the system is "working at a very basic level" because, in his/her opinion, the outputs are "toy examples". However, the processes used to produce such outputs is never mentioned. Because of the emphasis on the output, this description illustrates an engineering perspective. Second, most "contemporary systems" can be located on the engineering approach; therefore, they spend lots of resources on producing spectacular outputs. It seems unfair to judge Tlahcuilo by only comparing its outputs.

> The E-R model has been applied mostly to natural language generation, e.g., storytelling (Pérez y Pérez and Sharples, 2001), and more recently to other areas, e.g., interior design (Pérez y Pérez, Aguilar and Negrete, 2010). The author's claim that "it makes sense to study its practicality for producing visual compositions", however, visual compositions are very different in nature from storytelling, so why does it make sense that a system developed for stories is applicable to visual compositions? (Anonymous reviewer).

It is obvious for me –clearly, it is not that evident for other colleagues–, that using the ER-Model in different domains will produce interesting information about its scope, strengths and limitations. Comparing the similitudes and differences between the prototypes of the model for storytelling, interior design and visual composition, will generate information that, hopefully, will helps us to understand better the common elements of the creative process between different domains.

There are much more examples that I can quote. My main point here is that, if we analyse the system from the Engineering approach, all the comments made by the reviewers make sense; if we analyse the system from the cognitive-social approach, the same observations are confusing or even senseless. Having a broad view of the possibilities of these type of systems would help to prevent this Tower of Babel problem.

The CC-Continuum is a work in progress. I am planning to incorporate branches along the way to illustrate the differences within the same approaches. Here is an example. On August 31, 2016, Mark d'Inverno published in the CC Forum (computational-creativity-forum@googlegroups.com) a post describing his point of views about the area.

> I find it hard to imagine a scenario where we could sustain interested in solely generated artificial content for very long. The times when something has sustained interest in me is in music performance situations because the human is put under new challenges to work with an autonomous system because it can take them out of their comfort zones and they have to work harder to make things work musically. And for the musical to work in this way they need to imbue the system with its own creative agency. They need to give it equal billing to get the best out of themselves and of the unfolding creative collaboration (Mark d'Inverno, CC Forum).

The idea of generating system that take humans out of their comfort zone, challenging their abilities, is really appealing. From the CC side, this is a complicated goal that requires time and effort. He continues his post as follows:

> So I think that where the future lies is exploring artificial creative agency. This is the idea that machines enable new kinds of creative partnerships for humans. That they stimulate, challenge, provoke us to work in new ways and to produce content that would not have been possible without the system. And, come to that, would not – or could not – have occurred working with any other human collaborator (Mark d'Inverno, CC Forum).

The last part resemble my definition of collaborative creative work mentioned earlier (see Pérez y Pérez 2015). Definitely, this is a very stimulating position that I locate on the Continuum at the engineering-mathematical approach; however, it does not really match the definition of Colton and Wiggins introduced at the beginning of this text. Mark d'Inverno emphasis is on producing artificial-human collaboration that results in products that would not have occurred working with a human partner. Thus, I suggest opening a new branch on the engineering side of the Continuum to allocate this kind of approach. The post continues:

> So we need to start with the human creative, and build systems that demonstrate this creative agency to creative. Systems which immediately – or at least quickly-

> open up new opportunities for collaboration where the human creative is happy for the system to take creative control at points in the dialogue. Such systems need agency, and this involves an awareness of the human creative, their goals, their previous content, the way they like to work, the artistic influences of the creative, and also - and this is where it starts to get interesting - influences (algorithmic or human) that could take the human creative into entirely unexplored territories. But I think we need to start with the creative and think about designing systems with the right kinds of agency and flow. Starting with the system and then trying to work out how a human creative might interact with it later seems the wrong way round (Mark d'Inverno, CC Forum).

In this last section, Mark d'Inverno acknowledges the importance of studying and incorporating in our systems knowledge about how human creative works. That is, in order to produce agents that perform at a level that makes a difference, we need to exploit all available resources. In this way, it seems clear that the engineering-mathematical and social-cognitive approaches are endlessly linked.

But things are more complicated. Counterpath Press has recently published a series of three computer generated books: The Truelist (Montfort 2017), MEXICA 20 years – 20 stories (Pérez y Pérez 2017) and Articulations (Parrish 2018). Nick Montfort, who might be called a "minimalist coder", has expressed that his goal with the The Truelist is to produce a text that never would have been written by a human (personal communication); the whole code that generated the book can be found in the last page. Alisson Parrish, who describes herself as an artist rather than a scientist, used statistical methods (in particular deep neural networks) and some other tools to generate Articulations. Although these two particular pieces might not have an explicit scientific goal –they are artistic works built employing algorithms–, I believe that they contribute to the field of computational creativity; the authors provide detailed explanations of the computational methods they develop, which therefore can be exploited by other creative agents, and their systems generates novel interesting outputs that invite to reflect about the creative process of writing and the role of computers in literature and art. However, this type of programs does not satisfy any of the definitions used to situate a system on the CC-Continuum. My provisional solution is to open a new branch towards the left side of it.

Similar problems arise when one analyses the cognitive-social approach. First, we need to include branches that allows differentiating systems based on their focus: social oriented or cognitive oriented; from there, it is possible to add sub-branches to differentiate, for instance, social cognition, embodied cognition, situated cognition, and so on.

One of the reviewers of this paper has expressed serious concerns because he/she feels that this author ignores the merits of those systems located towards the engineering-mathematical approach,

> regarding these as mere algorithmic efforts and disregarding the questions that system in this side of the spectrum try to tackle; for instance, the painting fool system uses framing as a mechanism to increase the perception of value of its output, searching for insights in research question such as "Does framing increases the perceived value of an automatically generated artefact?"' (Anonymous ICCC18-reviewer).

After carefully revising the original text I have not found any comments that suggest that one approach contributes to science more than the other; in fact, I have explicitly mentioned that this framework does not reflect those type of features. In any case, the scientific contribution depends on the characteristics of the project rather than in its position in the CC-Continuum. My main claim is that different systems pursuit different goals and therefore they are trying to answer different questions. The reviewer's example just illustrates my point: "Does framing increases the *perceived value* of an automatically generated artefact?" (Italics are mine). This is a typical research question of the engineering-mathematical approach, which is oriented towards finding mechanism to increase the *perceived value* of a computer generated product. At this point it is worth to remind the reader that I chose to use a continuum because none of this classifications should be considered definitive; as I mentioned earlier in the text, a system might even move through the Continuum depending on the interests of the researchers involved in its development.

A different reviewer pointed out that projects on the engineering-mathematical approached might also be interested in answering similar questions to those that I mentioned while describing the cognitive-social pole. I agree; nobody has the monopoly of research questions. Previously in this article I pointed out that most people working in plot generation, it does not matter the location of their systems on the Continuum, attempts to sort out how to automatically produce coherence sequences of actions. The difference relies on the kind of solutions that researchers are willing to undertake to solve a given problem.

A couple of reviewers found coincides between the CC-Continuum and previous well-known AI and CC debates: "I believe that this discussion [is] related to weak vs. strong creativity (like weak and strong AI)"; "I think that the CC-continuum is also about specific vs. general purpose"; "How does the CC-Continuum differentiate from the product vs process approach?" I agree that this work share some concerns with all these reflexions. However, I do not believe they are the same. One reviewer wrote:

> The engineering approach focuses on weak creativity, in which a system has to just look like creative. And in the strong creativity, the process for generating a creative

output has to be creative or show some general-purpose operations or mechanisms used in human creativity (Anonymous ICCC18-reviewer).

I would not claim that a cognitive-social oriented system is really creative or more creative than other type of systems. The reason is that we are far from understanding how creativity works. Furthermore, a cognitive-social oriented system might focus in very specific aspects of the phenomenon under study, rather than representing general-purpose operations. I would claim similar arguments regarding the specific vs. general purpose dispute: "Engineering based systems are more concerned about creating creative artefacts for a specific domain and cognitive-social to general domain" (Anonymous ICCC18-reviewer).
Concerning the product vs. process dispute, I see similar problems. Most cognitive-social oriented systems are designed to bring about an output. However, their focus is not on generating products to amaze an audience but rather in designing mechanisms that provide plausible explanations for the creative process; however, the generation of outputs that help to support the cognitive/social hypothesis underlying the model, is an essential part of such explanations. Thus, the process is essential but it cannot be separated from the product. In similar ways, some systems on the engineering-mathematical side of the Continuum might emphasise the technical value of their process. For instance, rather than the output, the most interesting aspect of the work previously mentioned by Heat and Ventura (2016b) is the use of the gradient ascent method (Simonyan, Vedaldi, and Zisserman 2013).

The Continuum is one of multiple possibilities of classifying systems in the area of computational creativity. I am incline to believe that, currently, there are more programs located close to the engineering-mathematical approach than to the other parts of the Continuum. Because cognition and society is at the core of creativity, I claim that the field would become stronger with a more balanced distribution. In the same way, it is important to be caution that one perspective does not rule over the other. The CC-Continuum is a useful tool to make sense of the different approaches that specialists and students might pursuit in this field, to study how different research interests might profit from each other, and, in summary, to provide a broad perspective of computational creativity.
Mark d'Inverno ends his post as follows: "Not sure I directly answered your questions Rafael but good to carry on the discussion!" Yes Mark, you did! Let us hope this text will encourage other people to continue the discussion.

References

Ackerman, M. and Loker, D. 2017. Algorithmic Songwriting with ALYSIA. *International Conference on Computational Intelligence in Music, Sound, Art and Design* (Evo-MUSART).

Ackerman, M., Goel, A., Johnson, C., Jordanous, A., León, C., Pérez y Pérez, R., Toivonen, H. and Ventura, D. 2017. Teaching Computational Creativity. In *Proceedings of the Eight International Conference on Computational Creativity*, Atlanta, pp. 9-16.

Boyang, L. 2015. *Learning knowledge to support domain-independent narrative intelligence*. PhD Dissertation, School of Interactive Computing, Georgia Institute of Technology.

Colton, S. 2012. The Painting Fool: Stories from Building an Automated Painter. In J. McCormack and M. d'Inverno (eds.), *Computers and Creativity*, Berling: Springer-Verlag, pp.3-38.

Colton, S., and Wiggins, G. A. 2012. Computational creativity: The final frontier? In *Proceedings of the 20th European conference on artificial intelligence*, 21–26. IOS Press.

Gervás, P. 2000. WASP: Evaluation of different strategies for the automatic generation of Spanish verse. In *Proceedings of the AISB-00 Symposium on Creative & Cultural Aspects of AI*, pp. 93–100.

Harrison, B., Purdy, C., and Riedl, M. O. 2017. Toward Automated Story Generation with Markov Chain Monte Carlo Methods and Deep Neural Networks. *Proceedings of the 2017 Workshop on Intelligent Narrative Technologies*, Salt Lake City, Utah.

Heath, D., and Ventura, D. 2016a. Creating images by learning image semantics using vector space models. In *Proceedings of The Thirtieth AAAI Conference on Artificial Intelligence,* pp. 1202-1208.

Heath, D. and Ventura, D. 2016b. Before a Computer Can Draw, It Must First Learn to See. In *Proceedings of the Seventh International Conference on Computational Creativity*, pp. 172-179.

Jordan, M. I. and Russell, S. (1999). Computational Intelligence. In R. A Wilson and F. C. Keil (edts.) *The MIT Encyclopedia of the Cognitive Sciences*, pp.LXXIII. Cambridge, MA: The MIT Press.

Jordanous, A. 2012. A standardized procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation*, 4(3):246-279.

Martins, P., Pereira, F. C., y Cardoso, A. 2015. Modelos computacionales para la integración conceptual. In Pérez y Pérez (ed.), *Creatividad Computacional*, pp. 51-64. Ciudad de México: Grupo Editorial Patria.

Montfort, N. 2017. *The Truelist*. Denver, CO: Counterpath Press.

Parrish, A. 2018. *Articulations*. Denver, CO: Counterpath Press.

Pérez y Pérez, R. 2017. *MEXICA: 20 years – 20 stories [20 años – 20 historias]*. Denver, CO: Counterpath Press.

Pérez y Pérez, R. 2015. A Computer-based Model for Collaborative Narrative Generation. *Cognitive Systems Research*, 36-37. pp. 30-48.

Pérez y Pérez R., González de Cossío, M. and Guerrero, I. 2013. A Computer Model for the Generation of Visual Compositions. In *Proceedings of the Fourth International Conference on Computational Creativity*, Sydney, Australia, pp.105-112.

Pérez y Pérez, R., Aguilar, A., and Negrete, S. 2010. The ERI-Designer: A Computer Model for the Arrangement of Furniture. *Minds and Machines*, Vol. 20, number 4, pp. 483-487.

Pérez y Pérez, R. & Sharples, M. 2004. Three Computer-Based Models of Storytelling: BRUTUS, MINSTREL and MEXICA. *Knowledge Based Systems Journal*, Vol. 17, number 1, pp. 15-29.

Pérez y Pérez, R. & Sharples, M. 2001. MEXICA: a computer model of a cognitive account of creative writing. *Journal of Experimental and Theoretical Artificial Intelligence*, Volume 13, number 2, pp. 119-139.

Russell, S. and Norvig. P. 1995. *Artificial intelligence a modern approach*. Prentice Hall.

Ryan, J. 2017. Grimes' Fairy Tales: A 1960s Story Generator. In *Proceedings of the10th International Conference on Interactive Digital Storytelling*, pp.89-103. Funchal, Madeira, Portugal.

Singh, D., Ackerman, M. and Pérez y Pérez, R. 2017. A Ballad of the Mexicas: Automated Lyrical Narrative Writing. In *Proceedings of the Eight International Conference on Computational Creativity*, Atlanta, pp. 229-236.

Simonyan, K.; Vedaldi, A.; and Zisserman, A. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*.

Saunders, R. Artificial Creative Systems. 2018. In T. Veale and A. Cardoso (eds.) *Computational Creativity: The Philosophy and Engineering of Autonomously Creative Systems*, Springer.

Veale, T. 2015. Hallando la creatividad en conceptos en conflicto. In Pérez y Pérez (ed.), *Creatividad Computacional*, pp. 65-76. Ciudad de México: Grupo Editorial Patria.