

# You Can't Know my Mind: A Festival of Computational Creativity

**Simon Colton**

Computational Creativity Group, Department of Computing, Goldsmiths, University of London, s.colton@gold.ac.uk

**Dan Ventura**

Computer Science Department, Brigham Young University, ventura@cs.byu.edu

[www.thepaintingfool.com/galleries/you\\_cant\\_know\\_my\\_mind](http://www.thepaintingfool.com/galleries/you_cant_know_my_mind)

## Abstract

We report on a week-long celebration of Computational Creativity research and practice in a gallery in Paris, France. The festival was called *You Can't Know my Mind*, and was intended to introduce to the public the idea that researchers such as ourselves are writing software to be surprisingly unpredictable and creative in nature. The festival included a traditional art exhibition with a vernissage, a live music evening, a poetry night coupled with a food tasting, and a week long demonstration of mood-driven portraiture from The Painting Fool software. Each of the events – which are described here for the first time – involved an element of creative responsibility taken on by various software systems. The success of the festival was demonstrated in terms of attendance and feedback, pieces written by journalists, and follow up events which have taken place in 2013 and 2014.

## Introduction

In addition to advancing scientific and philosophical understanding of creativity, a long-term aim of Computational Creativity research is to embed creative software into society. For the general public to accept software as being independently creative, they need exposure to such software in cultural settings. To this end, we held the first *Festival of Computational Creativity* in the Galerie Oberkampf, located in the 11th arrondissement of Paris, France, during the week of 12th to 19th July 2013. As described in the next section, the festival consisted of five elements: an art exhibition, a live music performance, poetry reading, food tasting and a portraiture demonstration. Each element showcased a different system/project contributing creatively to the event, and – as highlighted by the festival name – the overall purpose was to portray software as being possible of autonomous, unpredictable, yet interesting and creative behaviour.

Our aim with the festival was to expose audiences to the main ideas of Computational Creativity within a culturally relevant setting, rather than to study audience experiences. Hence, we did not undertake experiments to gauge reactions to the ideas, systems and outputs presented. As described in the discussion section below, we claim success for the event through the number of attendees, the informal feedback we gained, some attention from journalists and the invitations to demonstrate the portraiture system in further events. We conclude in the discussion section with a brief look at future directions, and end with a montage of images from the festival, which we refer to as images A to N throughout.

## Elements of the Festival

### Art Exhibition

The art exhibition ran for the duration of the festival in the Galerie Oberkampf and was open to the public for 10 hours each day. The curator was Blanca Pérez Ferrer, who, in collaboration with the authors, chose and arranged 42 pieces produced by The Painting Fool software (Colton 2012), which has a long history of involvement in Computational Creativity projects described at [www.thepaintingfool.com](http://www.thepaintingfool.com). The first 4 pieces (image G of the montage) came from a back-catalogue of pieces which have been previously exhibited (Colton and Pérez-Ferrer 2012). In addition, 14 new pieces from a series entitled *Concrete Nudes* were selected and arranged along the main wall of the gallery (images J and L) – see Figure 1 for examples. These were produced by The Painting Fool simulating handwriting onto digital photographs of concrete walls taken in Rio de Janeiro. The handwriting (of random words) picked out depictions of female and male bodies, via their silhouettes and the capturing of internal contours with breaks in the text. Examples were used in the publicity material for the festival, such as the poster (image H), and the frontage of the gallery (image E). Finally, two sets of 12 postcard sized prints from The Painting Fool's most recent projects were chosen and hung (image K). The vernissage for the exhibition was attended by around 70 people (image F). Around 50 people visited the exhibition from the street during the week, and there were 2 private viewings.



Figure 1: Examples from the *Concrete Nudes* series.

## Portraiture Demonstration

The Painting Fool is software that we hope will be taken seriously as a creative artist in its own right, one day. We have cultivated its image through web pages, exhibitions and papers, and given it certain behaviours with the hope that it becomes increasingly difficult for people to use the word ‘uncreative’ to describe what it does. Note that, given the philosophical standpoint presented in (Colton et al. 2014), we aim to avoid the uncreative label rather than to gain a label of ‘creative’. The central exhibit of the exhibition was also called You Can’t Know my Mind and involved The Painting Fool producing portraits, with the explicit purpose of modelling artistic behaviours onto which people can project the words: *skill, appreciation, imagination, learning, reflection* and most notably, *intentionality*. We argue in (Colton et al. 2014) that software lacking such behaviours is relatively easy to call uncreative. The exhibit works as follows:

(i) When a person sits down for a portrait, the software has been reading *The Guardian* newspaper articles for some time: performing sentiment analysis to determine whether an article is upbeat or downbeat relative to the corpus, and extracting key phrases with which to search for related articles. The average sentiment over 10 recent articles is used to simulate the software being in a very positive, positive, experimental, reflective, negative or very negative ‘mood’.

(ii) If the software is in a very negative mood, it essentially tells the sitter to go away, refusing to paint a portrait on the basis of having recently read too many downbeat articles. It chooses the most negative phrase in the most negative article, and uses this in a commentary for the sitter to take away, which explains why it couldn’t paint their portrait.

(iii) If in a positive/very positive mood, the software chooses one/two of nine upbeat adjectives (e.g. *bright, colorful, happy*) and directs the sitter to smile while it extracts their image from a video recording over a green-screen background (image N). If in a negative mood, the software chooses one of six downbeat adjectives (e.g. *bleary, bloody, chilling*) and directs the sitter to express a sad face. If in an experimental mood, it chooses one of 11 neutral adjectives (e.g. *glazed, abstract, calm*) and asks the sitter to pull an unusual face. If in a reflective mood, the software chooses an adjective for which it has previously had a failure (see later).

(iv) The chosen adjective is used to select a filter (from a set of 1,000 possibilities) that, when applied to an image of a face, is likely to achieve an appropriate visualisation for the adjective. Appropriateness is modeled using a set of visuo-linguistic association (VLA) neural networks (one per adjective) borrowed from the DARCI system (Norton, Heath, and Ventura 2013). These networks have learned correlations between visual features and semantic (adjectival) concepts, and high network output indicates high appropriateness for the adjective represented by the network. The background in the captured facial image is replaced with an arbitrary abstract art image, the chosen filter is applied to both the background and foreground (face) image, and edge detection is used to overlay edges from the face which pick out features of the sitter. The combined background+foreground+edge

image is taken as a ‘conception’ of what The Painting Fool aims to achieve in its rendering.

(v) One of seven rendering styles involving the simulation of paints (2 styles), pencils (3) and pastels (2) is chosen to produce the portrait. If a pairing of adjective/style hasn’t been attempted before, then that style is chosen, otherwise a style is chosen according to the probabilistic model it has learned for the adjective (see later), with better styles more likely. A hand appears on-screen, holding a pastel, pencil or paintbrush, and proceeds to render the image in a vaguely human-like fashion. This process (images D and N) takes from a few minutes for pastels to around 20 minutes for paints.

(vi) Once the rendering is complete, the VLA neural network for the chosen adjective is again used, this time to assess the appropriateness of the rendered image and hence whether it is actually appropriate to use the intended adjective to describe the final portrait. VLA outputs for the conception and the final rendered image are compared to assess whether the rendering technique has increased or decreased (relative to the conception) the appropriateness of the portrait (for conveying the adjective). This assessment determines whether the session has been a ‘great success’ (significantly increased) or a ‘miserable failure’ (significantly decreased) or something in between. To end the portraiture session, The Painting Fool prints the portrait with a commentary on the reverse, as per Figure 2. The commentary details the mood the software was in and what adjective it chose, shows the conception compared with the final portrait, and discusses whether it has achieved the aim of producing a portrait of a particular style and how the portrait compares with the conception in that respect. Finally, VLA neural networks for all negative/positive/experimental adjectives are opportunistically applied to the portrait, to see if it can be further described with additional pertinent adjectives.

(vii) Before returning to reading news articles, The Painting Fool scores how effectively the rendering method conveys the chosen adjective. In particular, if it has failed (by significantly reducing the VLA network output for that adjective), then, when in a reflective mood in the future, if this adjective is chosen again, the portrait will be attempted with a different rendering style. In this way, the system builds a probabilistic model of which rendering styles are likely to successfully convey which adjectives, e.g., it learns that pencils are better at producing monochrome or bleary portraits, while paints are better for busy or patterned portraits.

As an example, in Figure 2, while in a negative mood, the software chose the adjective ‘bleary’, which led to it selecting an image filter which desaturated the image, as depicted in the top conception image of Figure 2. It chose to simulate paints to produce the portrait, as depicted in the bottom rendered image of Figure 2, and then commented that (a) the final portrait is very bleary overall, and (b) it had achieved the same amount of blurriness in the rendered image as the conception, with which it was OK. As a final flourish, it also points out that the portrait is bleached, which fits its mood. Over the week of the festival, more than 100 portraits were produced, and we chose 60 to fill the back wall of the gallery towards the end of the festival (image M).

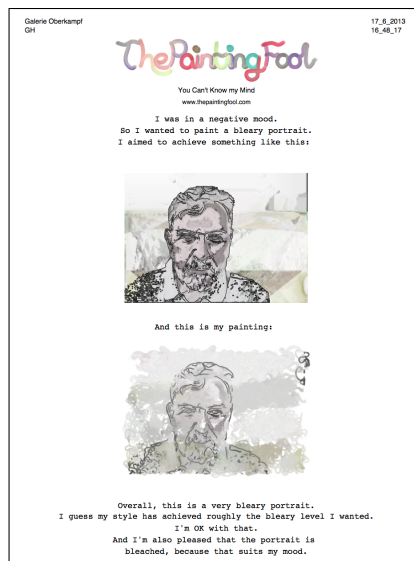


Figure 2: Example portraiture commentary.

### Moody Music Evening; Poems and Potage Night

On the first evening of the festival, musician Stéphane Bissières played live to an audience of around 50 people (image A). As part of the performance, The Painting Fool’s newspaper-reading mood model (described above) was adapted to inform software for performing affective, real-time sound design and rhythm construction. The software’s output was converted to MIDI and sent to Bissières’ music system, requiring him to react musically in real-time. Bissières and the system collaborated on three different musical sets, each approximately 20 minutes in length. Each set had a different musical feel, effected by three different algorithmic approaches to how different moods would affect composition and performance. Bissières was enthusiastic about collaborating with an autonomous system. He, and the audience, responded intuitively to the software’s mood changes, and the often unpredictable turns and reactions to them added energy to the performance. Moreover, the graphic visualisation of the mood on the monitor (image A) enabled the audience to appreciate the computer’s role in the composition/performance process.

On the fourth night of the festival, we presented computational poetry and computational cuisine, to an audience of around 60 people. In advance, seven automatically generated poems were selected from a much larger corpus by Russell Clark, then analysed as if they were required reading for an English exam. Two of the poems were generated by the system described in (Colton, Goodwin, and Veale 2012), and these were supplemented by more recent poems constructed using material from Twitter. On the poetry night, the poems were recited, along with their analyses during three sessions (image C). In each session, Clark complemented the computational poems with classical poems from Pope, Hulme and Eliot, and wove comparisons into his analysis.

Alongside this, following recipes created by a computational chef called PIERRE (Morris et al. 2012), Chef Sophie Grilliat prepared three soups for consumption between

the poetry sessions (image I). In practice, however, the soups were so popular that all three were eaten in the first break. As with the poems, the soups were presented in context, in this case French cuisine – Chef Grilliat prepared classical complementary finger food. A booklet of poems and recipes was handed out to audience members (image B).

### Discussion

The aim of the festival was to expose members of the public to the idea that software can be independently creative. With the art exhibition, we exposed the high quality of artefacts that can be produced by creative software; with the recipe generation, poetry and mood music, we highlighted the breadth of Computational Creativity systems, in terms of application domains and different human-computer interaction schemas; with the You Can’t Know my Mind exhibit, we demonstrated the intelligence, independence and unpredictability of creative software exhibiting behaviours onto which it might be appropriate to project words such as *intentionality*, *reflection* and *learning*.

To emphasise the behaviours exhibited by The Painting Fool, we put up posters explaining six of its behaviours in understandable terms, e.g., intentionality was addressed by the software being directed to choose an adjective by a mood, conceiving an image it wished to produce through simulation of artistic media, producing the rendering and then determining whether it had achieved its goals. We similarly explained how the software *reflected* on its failures and *learned* from its experience how to choose appropriate rendering styles for future portraits. We anthropomorphised the software ‘being in a mood’, ‘reading newspaper articles’ and ‘being happy’ to help explain to audiences what the software was doing. This was done in order to enable them to make an informed opinion about whether it was appropriate to call the software ‘uncreative’ or not. We asked dozens of audience members to give us a good reason why they felt it was appropriate to call the software uncreative, and we didn’t receive any salient answers in this respect, which we believe indicates how well we handled public perception of The Painting Fool during the festival.

Around 200 different people attended the events of the festival, which was covered by journalists writing for *Wired* and *Pacific Standard Magazine*, which in turn have led to the You Can’t Know my Mind project being covered by *Stuff* magazine, *The Smithsonian* magazine, and German and British radio shows. Naturally, this has led to much wider exposure of people to the notion of creative software. It has also led to invitations to demonstrate the exhibit at the London Science Museum, the Cité des Sciences in Paris, the AISB Convention and the American University in Paris. With each portrait painted, The Painting Fool becomes more aware of its abilities, and we plan to enhance these, for instance, with further machine vision techniques (to tell during a painting whether it is on the right track) and the ability to tweak its painting style (to try to get back on the right track). By further enhancing its artistic and creative abilities, and continuing to present the You Can’t Know my Mind exhibit as widely as possible, we hope to convince people that creative software is coming, and will enhance our lives.



## References

- [Colton and Pérez-Ferrer 2012] Colton, S., and Pérez-Ferrer, B. 2012. No photos harmed/growing paths from seed – an exhibition. In *Proc. of NPAR*.
- [Colton et al. 2014] Colton, S.; Cook, M.; Hepworth, R.; and Pease, A. 2014. On acid drops and teardrops: Observer issues in Computational Creativity. In *Proc. of the AISB symposium on AI and Philosophy*.
- [Colton, Goodwin, and Veale 2012] Colton, S.; Goodwin, J.; and Veale, T. 2012. Full-Face poetry generation. In *Proc. of the 3rd Int. Conference on Computational Creativity*.
- [Colton 2012] Colton, S. 2012. The Painting Fool: Stories from building an automated painter. In McCormack, J., and d’Inverno, M., eds., *Computers and Creativity*. Springer.

[Morris et al. 2012] Morris, R.; Burton, S.; Bodily, P.; and Ventura, D. 2012. Soup over bean of pure joy: Culinary ruminations of an artificial chef. In *Proc. of the 3rd International Conference on Computational Creativity*.

[Norton, Heath, and Ventura 2013] Norton, D.; Heath, D.; and Ventura, D. 2013. Finding creativity in an artificial artist. *Journal of Creative Behavior* 47(2):106-124.

## Acknowledgements

This project was funded by EPSRC grants EP/J001058 and EP/J004049 and by a sabbatical from Brigham Young University. We owe a great deal to our family and friends who indulged us and helped us hugely with the festival. To Blanca Pérez-Ferrer, Sophie Grilliat, Stéphane Bissières and Russell Clark: thank you so much.

