Twenty-Five Years of The Painting Fool Project

Simon Colton and Blanca Pérez Ferrer

School of Electronic Engineering and Computer Science, Queen Mary University of London, UK

www.thepaintingfool.com

Abstract

The aim of The Painting Fool project is to build an automated artist with cultural impact. For twenty-five years, it has helped shape the growing field of Computational Creativity, and has likewise been shaped by it. We describe here timelines of technical development and public engagement activities undertaken with the software. We also describe how philosophical discourse in Computational Creativity has paid into the project and been influenced by it, and we offer some observations from the long-term interaction with visual arts culture and communities afforded by the project.

Introduction

The development and deployment of the *The Painting Fool* software has been a long-term project of AI researcher Simon Colton and colleagues for 25 years. The overall aim of the project has remained the same throughout, namely that the software itself be taken seriously as a creative artist in its own right, one day (Colton 2012). In addition to driving numerous technical projects to simulate visual art production processes, there have been many public engagement activities, including live demonstrations, art exhibitions, festivals, residencies, companies and commissions. Moreover, reference to The Painting Fool itself and the idea of people taking seriously an AI system as being creative in a human-centric cultural domain such as the visual arts has provided a focal point and provocation for numerous philosophical contributions, also published in the Computational Creativity literature.

As The Painting Fool implementation has reached 25 years as a software entity, it seems sensible to record details of the project so far. To do so, we describe below the origins of the software, the implementation of the system which evolved in epochs, and the technical projects that filled gaps in the software's abilities and behaviours. We also describe the public engagement activities that Colton and others have undertaken with The Painting Fool, and some of the ways in which philosophical ideas drawing on the project have been developed. We end with some conclusions drawn from this long-term study and detail some of the current projects and future directions for The Painting Fool.

Technical Development Timeline

Colton describes the pre-history of the project in terms of the revelation of buying his first digital camera in 1999. Submitting a roll of film and waiting two weeks for a chemist to develop them into print phototraphs, was frustrating for amateur photographers. Digital cameras meant that hundreds of photos could be taken in one session and viewed/printed immediately. Colton used Adobe Photoshop extensively to make the photos he took of the city of Edinburgh look more artistic, which eventually led to him writing graphics software to perform image filtering more to his liking.

During post-doc work on generative maths, in 2001 Colton found he was spending significant time on graphics coding, and decided to include this in his research, broadening his domains of interest in Computational Creativity to include the visual arts. The first public mention of The Painting Fool was at the AISB'01 symposium on AI and Creativity at the University of York in 2001. Colton gave a talk on Meta-Theory Formation, mentioning details of his graphics work at the end.

Colton splits description of the first decade of implementation of The Painting Fool into three overlapping epochs, where behaviours relating to *skill*, *appreciation* and *imagination* in the (simulated) painting process were implemented. During the skill period, the initial graphics software was extended to perform edge detection and image segmentation in order to produce abstracted versions of digital photographs. By 2004, this had been extended with non-photorealistic rendering (NPR) techniques, so that colour regions could be filled or outlined with simulated paint/pastel/pencil strokes. Colton set up web pages at `thepaintingfool.com` displaying generated images from the *City* series and *Amelie's Progress* series (see figure 1 of the appendix).

Reading in an NPR textbook (Strothotte and Schlechtweg 2002) about simulating artistic media and techniques, Colton found the following proclamation:

> Simulating artistic techniques means also simulating human thinking and reasoning, especially creative thinking. This is impossible to do using algorithms or information processing systems.

Colton saw this as denying the potential of AI in gen-

eral, and Computational Creativity in particular, providing a rallying point to make simulation of artistic behaviours a focal point for his research. An important element implemented during the skill period was live painting demonstrations. Inspired by the robotic painting skills of Harold Cohen's AARON system and by embodied Computational Creativity in general (Guckelsberger et al. 2021) (Moruzzi 2022), this feature has been the most influential in enabling people understand the technologies and theories of The Painting Fool project. Accordingly, it has been the foundation of many public engagement activities, as described below.

In 2006, Colton worked with colleagues Michel Valstar and Maja Pantic at Imperial College London on enabling The Painting Fool to employ their machine vision approaches to guide the artistic process. As described below, this work led to a winning entry in the 2007 BCS Machine Intelligence Competition. It also kick-started the *appreciation* period where sub-projects and code developments were prioritised which led to The Painting Fool having more understanding of its subject material, its output and the artworks of others. During this period, Colton began to seriously hand over creative responsibility to the software, with technical projects testing how independent the software could be.

At the first art exhibition where The Painting Fool's images were displayed (in 2008, also detailed below), another exhibiting artist, namely James Faure-Walker, mentioned in conversation with Colton that artists rarely start projects with an image, more commonly beginning with an idea, emotion or technique that they want to explore (Faure-Walker 2006). This inspired Colton and led to the third period of development where behaviours that might be related to notions of imagination were implemented and tested in The Painting Fool. This involved further handing over creative responsibility in ideation, process and curation.

In particular, the software was enhanced to generate starting imagery to replace digital photographs, to be passed through various pipelines to produce images, videos and framing information (Cook et al. 2019). This was mainly achieved in two ways, namely producing abstract images using evolutionary art techniques and representational art using context-free design grammars. Colton also experimented with other AI techniques such as constraint solving (Colton 2008c) to get The Painting Fool to construct what he called *scenes from a perceived imagination*. This led to images such as those in the *PresiDENTS* series, where photographic elements (of Democratic US presidents), were painted as part of a scene (of generated trees), as portrayed in figure 1 of the appendix. Colton also experimented with fitness function invention, using his HR software (Colton 2008a), leading to surprising prepainting cityscape scenes with buildings on top of each other and other unexpected anomalies.

A suitable bookend for the first decade of development of The Painting Fool was the exhibition entitled *No Photos Harmed* in 2011, described in (Colton and Pérez Ferrer 2012) and below, with details of the project so far written up in a book chapter in 2012 (Colton 2012). The following decade was a period where skilful, appreciative, imaginative and other types of behaviour were implemented in an ad-hoc fashion in response to philosophical developments and technical advances in Computational Creativity research. During this time, Colton undertook various public engagement activities and technical projects with the software, spinning off sub-projects such as image filtering in the *Filter Feast* software (Colton and Torres 2009), and abstract art generation in software eventually called *Art Done Quick* (Colton 2020), with a particular emphasis on casual creation applications (Compton and Mateas 2015).

During this period of exploitation, Colton added what he calls *accountable unpredictability*, where it may not be possible to predict where The Painting Fool will go with a painting, but if something interesting happens (or is produced), one can investigate the source of the event, which is strictly not random-number generation. To do this, he enabled the software to use APIs connected to Twitter, Flickr and The Guardian newspaper, so unpredictable trending information or content could be used as starting points for artistic production.

As an example, MSc. student Anna Krzeczkowska used such functionality to produce collages in response to Guardian news articles, where image retrieval from Flickr was used with The Painting Fool's NPR abilities (Krzeczkowska et al. 2010). One notable output is given in figure 1 of the appendix, where a collage involving babies, a graveyard, a bomber plane and an explosion was generated in response to an article on the war in Iraq. This project, and work in a follow up paper (Cook and Colton 2011), led to questions around where – if anywhere – intentionality lay in producing such images. Ultimately, this work was a pre-cursor to more extensive development and engagement in the *You Can't Know my Mind* portraiture project, described below. This involved working with Dan Ventura to add further machine vision abilities (Colton and Ventura 2014) to enable The Painting Fool to aim for a particular visual style, and assess success in this respect.

In around 2020, Colton realised that generative deep learning models were being trained that could undertake much of what The Painting Fool did via more traditional graphics and AI techniques. He wrote code to interact with pre-trained neural models for image generation (via GAN and diffusion models) and image analysis (via machine vision models such as ResNet and vision-language models such as CLIP). As part of this, Colton worked with PhD students Amy Smith and Sebastian Berns, to explore the creative potential of early text-2-image generation (Colton et al. 2021), producing time-lapse images with generative adversarial networks (GANs) (Colton and Pérez Ferrer 2021), blending neural style transfers for casual creation (Colton 2021) and evolving prompts for diffusion models (Colton et al. 2023). The code from these projects was incorporated into The Painting Fool for future use.

Public Engagement Activities in the Field

The notion of Computational Creativity is controversial, divisive and nebulous (Colton et al. 2014), so it is important to sample opinions from sectors of the general public in order to guide practical and theoretical developments in the field. Colton has used The Painting Fool project to do this in the following ways. Note that images from some public engagement activities are portrayed in figures 3 and 4 of the appendix.

• Art Exhibitions. Artwork from The Painting Fool has appeared in two solo and nine group exhibitions. Along with John Cass, in 2006, Colton organised an art exhibition entitled *Computer Generated Artworks*, held in the business school of Imperial College, fittingly on Exhibition Road in London. This was the first time that The Painting Fool's artworks were exhibited, and Colton gained much feedback from established digital artists, including James Faure-Walker, William Latham, Patrick Tresset and Pensoual Machado. In 2011, along with Eileen Chen, Colton organised an exhibition in Paris entitled *No Photos Harmed*, which highlighted that AI systems can generative artworks which are not abstract, but also not produced from a digital photograph (Colton and Pérez Ferrer 2012). The centrepiece of this was a 3m by 1m piece entitled *The Dancing Salesman Problem*, produced by collaging 300 photo-sized printouts. In the same year, a pop-up retrospective called *Growth*, showcased more than 50 of The Painting Fool's outputs at the Maison Rouge gallery in Paris, as part of a Sony CSL open-house event.

Along with Dan Ventura and Blanca Pérez Ferrer, in 2013, Colton organised a Festival of Computational Creativity entitled *You Can't Know my Mind* in Paris, which included a generative music concert, a reading of generated poetry, a night of generated hot-pots and an exhibition of The Painting Fool's work (Colton and Ventura 2014). A central part of this was a demonstration where The Painting Fool painted portraits to express a simulated mood gained from reading newspaper articles, and learned from the experience, described in (Colton et al. 2015). More recently, Colton employed the new generative AI technologies in The Painting Fool for an exhibition of *GANLapse photography* at the Etopia Centre for Art and Technology in Spain (Colton and Pérez Ferrer 2021).

• Demonstrations. The first major demonstration of The Painting Fool's abilities was as a competition entry for the annual 2007 British Computer Society. The team of Maja Pantic, Michel Valstar and Colton won the competition with a demonstration of *emotionally aware portraiture painting* (Colton, Valstar, and Pantic 2008) where The Painting Fool was informed by software detecting the emotion of a sitter, in order to choose a painting style accordingly. Colton and neuroscientist Derek Jones exhibited art at a Brussels exhibition kicking off the S+T+ARTS (starts.eu) programme. This led in 2018 to Jones commissioning The Painting Fool to be artist in residence at the CUBRIC centre for brain imaging, which was a fruitful collaboration, e.g., producing art to illustrate dreams. In 2016, The Painting Fool appeared at a New Scientist Live event entitled *Anything You Can Do, AI Can Do Better* to paint comedian Dara O'Briain. However, the most substantial demonstration was held in the Science Museum London in 2015: five days and one evening of the You Can't Know My Mind demonstration painting portraits as part of the Museum's Antenna Live and Lates programme. The Painting Fool produced more than 300 portraits in continuous generation projected onto large screens. This demonstration was also undertaken at other venues, including the London Natural History Museum and the Paris Cité des Sciences.

• Media Coverage. Colton estimates that around 100 third-party print/web articles, news items and TV/radio documentaries have covered The Painting Fool project, including every major UK newspaper, science publications like the New Scientist and many international outlets. Of particular note, in 2012, mathematician and science communicator Marcus du Sautoy profiled The Painting Fool and Colton in an episode on Learning and Creativity for a flagship BBC Horizon documentary on AI. This brought the ideas of AI autonomy in the arts to a relatively large public audience. Similarly, in 2016 Mrs. Yang Lan – a well-known personality on Chinese television – profiled The Painting Fool as part of a CCTV documentary on AI, had her portrait painted by the software, and purchased a printed piece. A member of the production team told Colton that they hoped the worldwide audience for the documentary would be 1.2 billion. Most of the media coverage of the project has been positive, but there has been negative feedback, for instance a writer for the Observer Magazine was particularly critical of the project.

• Commercial activities. Along with business partner Glen Pearson, Colton set up UK company Machine Creations Ltd., in 2005, which produced paint-by-numbers-kits for customers who uploaded an image to the website. This used The Painting Fool's image segmentation software with enhanced abilities to choose paint mixes and render a painting kit. The company lasted two years and was not hugely successful, but provided an important learning experience in terms of commercial testing of the software. In 2014, behavioural information data analysis company UBIC commissioned The Painting Fool to produce printed images for their offices in Tokyo, as described in (Colton et al. 2015). This resulted in a series called *I Can See Unclearly Now* and tested the flexibility of the software under the scrutiny of a paying patron. More recently, in late 2025, AI-art collector *Delronde* (see eng.art) contacted Colton to purchase pieces from The Dancing Salesman Problem series and others, as both physical prints and NFT digital artworks. This led to Colton and The Painting Fool being represented by the Fellowship.xyz gallery and an online exhibition of old and new work being made available for purchase as NFTs (see figure 4).

Philosophical Provocations

Recalling the aim for the software is to be taken seriously as a creative artist in its own right one day, Colton has proceeded by trying to minimise reasons people could cite for why The Painting Fool is not creative, hoping that eventually – without such reasons – they would begrudgingly admit that the software *is* creative. On many occasions, this approach has led to new implementations, experimentation and outreach, e.g., in response to comments about lack of imagination and intentionality. On other occasions, working closely with colleagues such as Alison Pease, Teresa Llano, Christian Guckelsberger, Michael Cook and Berker Banar, Colton was able to generalise issues with The Painting Fool and write in a more philosophical way to add to the theory of Computational Creativity.

The discourse in (Colton 2008b) was a response to early (well-meaning) criticisms of The Painting Fool as an independent artist. Here, the idea was proposed that if a creative AI agent was lacking in abilities onto which people could project notions of skill, appreciation and imagination, then it could be easily dismissed as being uncreative. As mentioned above, this formulation of behaviour types guided the first decade of projects with The Painting Fool, and was later expanded to include more behaviour types covering innovation, intentionality, learning, accountability and reflection. In addition, constant suggestions that the only way to project artistic creativity onto the software was to compare its output to art produced by people led to a critique in (Pease and Colton 2011) of Turing-Style blind comparison studies and a suggestion of alternatives. On two different occasions, outputs from The Painting Fool were shown to artists (one in Paris and another in Beijing) by television crews, to get their reaction to the images and the project as a whole. On both occasions, the artists said they felt the images were "cold", lacking empathy and understanding of the human condition, clearly judging the artworks in terms of the process, personality and press (Jordanous 2016) rather than product. This led to special consideration in (Colton et al. 2014) of observer issues in Computational Creativity.

Ultimately, autonomy in process became less of a criticism through projects such as *You Can't Know my Mind*. This left difficult questions of autonomy with respect to motivation, and led to theoretical papers highlighting issues with respect to intrinsic motivations (Guckelsberger et al. 2017) and authenticity (Colton, Pease, and Saunders 2018) in computational systems. One solution to this was proposed in (Colton et al. 2020), where the authors suggest that creative AI systems like The Painting Fool can record aspects of their daily existence to become the focus of artistic creations within a framework they call the Machine Condition, in order to express what it is like to be the AI system. This idea was pushed further, by looking at how AI systems could add to artistic cultures (Colton and Banar 2023) and what it would mean for them to gain creative personhood (Pease, Colton, and Banar 2023).

Conclusions and Future Work

With such a volume of longitudinal technical research, philosophical discourse and public engagement in the field, a number of conclusions can be drawn from The Painting Fool project that might transfer to other Computational Creativity projects. These include:

- Complexity in backstory, motivation and process is far more interesting to artistic communities than sophistication, novelty or value of the artefacts that a generative system produces.
- In general, people still find it difficult to project notions of artistic creativity onto software systems, in both principle and practice (Hertzmann 2018). It is still reasonable to say that the origin of creativity lies more with the programmer of such systems than the system itself.
- With long-term projects, it is beneficial for practical and theoretical work to drive each other forward in the study of Computational Creativity.
- Handing over creative responsibilities to software is a common mantra in Computational Creativity research, but has rarely been studied anywhere else in AI or the arts. Moreover, the position that AI systems should only ever be tools for people is even more entrenched now, given the rise of issues around livelihood, legacy and copyright affecting artists and musicians.
- Public engagement has three facets: projects such as an art exhibition which showcases Computational Creativity work; engagement indicators, such as a newspaper article written about the exhibition, and actual engagement, such as conversations in the exhibition or comments left online. In general, engagement begets engagement and *relatable behaviours* (Colton 2026) in creative production and assessment can be key in bringing attention to projects involving creative AI systems.

Historically, The Painting Fool project sits between Cohen's AARON (McCorduck 1991) and Klingemann's BOTTO projects (Klingemann, Hudson, and Thompson 2023), both of which allude to some independence in art creation by an AI system. While it is clear that The Painting Fool is not yet taken seriously as a creative artist in its own right, Colton cites moments when he felt the software was acting as an artist, albeit momentarily and not in a hugely creative way. The Painting Fool project continues and the recent interest in its back-catalogue of art has prompted work towards a retrospective exhibition. There is more than enough theoretical work from the Computational Creativity community to guide the project into later stages. Moreover, the global tech industry regularly drops generative AI models which could be employed by The Painting Fool in the quest to endow it with creative autonomy and send it out into the artistic world. It's not impossible to think that after another 25 years, the project – or one inspired by it – will have succeeded, and the art world will be enriched with computationally creative artists.

Acknowledgements

We would like to wholeheartedly thank the many computational creativity researchers, students, artists and journalists who have contributed to The Painting Fool project, as well as the anonymous reviewers of the many papers written about it, including this one.

References

Colton, S. 2026. Experiments in Text-2-Tool-Use for Relatable Computational Creativity In *Proceedings of the 17th Int. Conference on Computational Creativity.*

Colton, S., and Banar, B. 2023. Automatically adding to artistic cultures. In *Proc. 12th Int. Conf. on Artificial Intelligence in Music, Sound, Art and Design.*

Colton, S., and Pérez Ferrer, B. 2012. No photos harmed/growing paths from seed: an exhibition. *Proc. Symp. on Non-Photorealistic Animation & Rendering.*

Colton, S., and Pérez Ferrer, B. 2021. GANlapse generative photography. In *Proceedings of the 12th International Conference on Computational Creativity.*

Colton, S., and Torres, P. 2009. Evolving approximate image filters. In *Workshops on Applications of Evolutionary Computation*, 467–477. Springer.

Colton, S., and Ventura, D. 2014. You Can't Know my Mind: A festival of computational creativity. In *Proc. of the 5th Int. Conference on Computational Creativity.*

Colton, S.; Cook, M.; Hepworth, R.; and Pease, A. 2014. On acid drops and teardrops: Observer issues in computational creativity. In *Proceedings of the 7th AISB Symposium on Computing and Philosophy.*

Colton, S.; Halskov, J.; Ventura, D.; Gouldstone, I.; Cook, M.; and Pérez Ferrer, B.. 2015. The Painting Fool sees! New projects with the automated painter. In *Proc. 6th Int. Conference on Computational Creativity.*

Colton, S.; A, P.; Guckelsberger, C.; McCormack, J.; Llano, T.; and Cook, M. 2020. On the machine condition and its creative expression. In *Proc. 11th ICCC.*

Colton, S.; Smith, A.; Berns, S.; Murdock, R.; and Cook, M. 2021. Generative search engines: Initial experiments. In *Proceedings of the 12th International Conference on Computational Creativity.*

Colton, S.; Smith, A.; Pérez Ferrer, B.; and Berns, S. 2023. Artist discovery with stable evolusion. In *Proc. 14th Int. Conference on Computational Creativity.*

Colton, S.; Pease, A.; and Saunders, R. 2018. Issues of authenticity in autonomously creative systems. In *Proc. 9th Int. Conference on Computational Creativity.*

Colton, S.; Valstar, M.; and Pantic, M. 2008. Emotionally aware automated portrait painting. *3rd Int. Conf. on Digital Interactive Media in Entertainment & Arts.*

Colton, S. 2008a. Automatic invention of fitness functions with application to scene generation. In *Workshops on Applications of Evolutionary Computation.*

Colton, S. 2008b. Creativity versus the perception of creativity in computational systems. In *AAAI spring symposium: creative intelligent systems*, volume 8.

Colton, S. 2008c. Experiments in constraint-based automated scene generation. In *Proc. of the 5th international joint workshop on computational creativity.*

Colton, S. 2012. The Painting Fool: Stories from building an automated painter. In McCormack, J., & d'Inverno, M., *Computers and creativity.* Springer.

Colton, S. 2020. Casual creation, curation, captioning, clustering and crossover in the Art Done Quick app. In *Proc. of the Workshop on Casual Creators at ICCC.*

Colton, S. 2021. Evolving neural style transfer blends. In *Proc. International Conference on Computational Intelligence in Music, Sound & Art (EvoMusArt).*

Compton, K., and Mateas, M. 2015. Casual creators. In *Proc. 6th Int. Conference on Computational Creativity.*

Cook, M., and Colton, S. 2011. Automated collage generation-with more intent. In *Proceedings of the 2nd International Conference on Computational Creativity.*

Cook, M.; Colton, S.; Pease, A.; and Llano, M. 2019. Framing in computational creativity-a survey and taxonomy. *Proc. 10th Int. Conf. on Comp. Creativity.*

Faure-Walker, J. 2006. *Painting the Digital River.* Prentice Hall.

Guckelsberger, C.; Salge, C.; and Colton, S. 2017. Addressing the "why?" in computational creativity. In *Proc. 8th Int. Conf. on Computational Creativity.*

Guckelsberger, C.; Kantosalo, A.; Negrete-Yankelevich, S.; and Takala, T. 2021. Embodiment and computational creativity. In *Proceedings of the 12th International Conference on Computational Creativity.*

Hertzmann, A. 2018. Can computers create art? *Arts* 7(2), MDPI.

Jordanous, A. 2016. Four PPPPerspectives on computational creativity in theory and in practice. *Connection Science* 28(2), 1-23.

Klingemann, M.; Hudson, S.; and Thompson, J. P. 2023. Cluster #069. In *SIGGRAPH 2023 Art Gallery.*

Krzeczkowska, A.; El-Hage, J.; Colton, S.; and Clark, S. 2010. Automated collage generation-with intent. In *Proc. 1st Int. Conf. on Computational Creativity.*

McCorduck, P. 1991. *AARON's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen.* W.H. Freeman and Company.

Moruzzi, C. 2022. The (artificial) physicality of creativity : How embodiment influences perceptions of creativity. In *Proceedings of the 13th International Conference on Computational Creativity.*

Pease, A., and Colton, S. 2011. On impact and evaluation in computational creativity: A discussion of the Turing test and an alternative proposal. In *Proceedings of the AISB symposium on AI and Philosophy.*

Pease, A.; Colton, S.; and Banar, B. 2023. On the notion of creative personhood. In *Proceedings of the 14th International Conference on Computational Creativity.*

Strothotte, T., and Schlechtweg, S. 2002. *Non-Photorealistic Computer Graphics.* Morgan Kaufmann.

Appendix

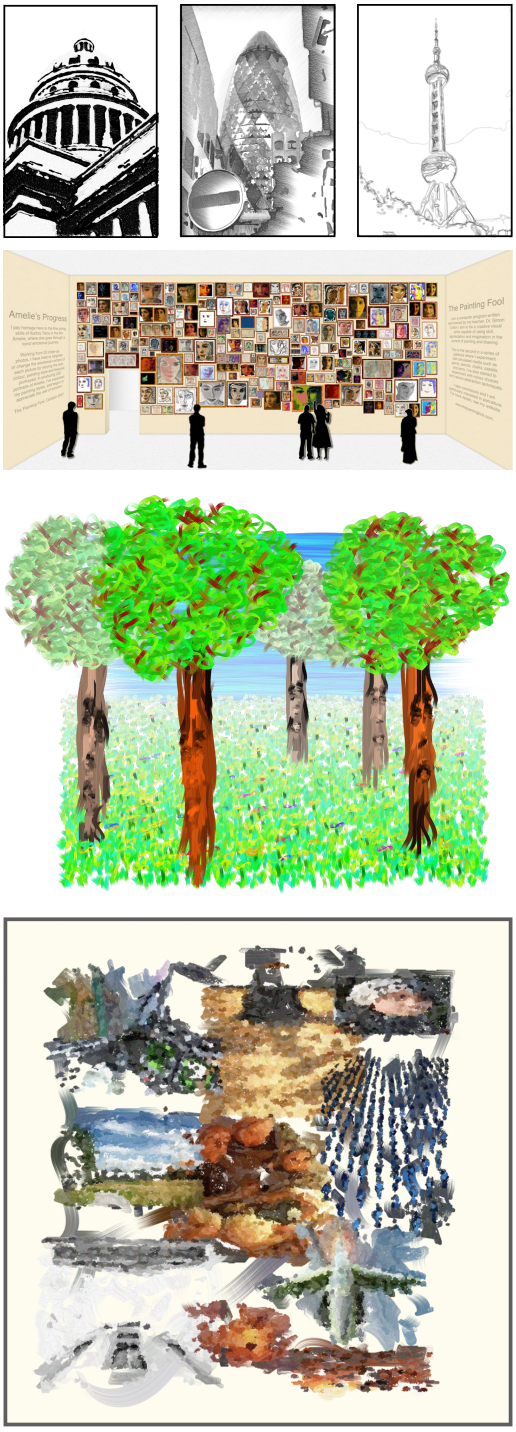

Figure 1: (a) images from the *City Series* (2003) and (b) *Amelie's Progress* (2005), highlighting The Painting Fool's non-photorealistic rendering abilities (c) image from the PresidENTS series (2009/2010), where scene invention was facilitated by constraint solving, and (d) an image from a news-based collage series, derived from an article on the war in Afghanistan, 2011.

Figure 2: Images from public engagement activities with The Painting Fool: (a) the Computer Generated Art exhibition of 2006 (b) first mainstream news article about The Painting Fool, 2006 (c) wall text for the *Growth* retrospective exhibition, Maison Rouge, Paris 2011, and (d) a child sits for a portrait at a portraiture demonstration at Cité des Science, Paris 2014.

Figure 3: Further images from public engagement activities with The Painting Fool: (a) Colton with The Dancing Salesman Problem piece at the *No Photos Harmed* exhibition of 2011 (b) Mrs. Yang Lan purchasing artwork by The Painting Fool in 2016 (c) The Painting Fool's walrus-like portrait of Dara O'Briain at a New Scientist Live event, 2016, and (d) visitors at the *You Can't Know my Mind* festival, 2013.

Figure 4: New image styles added in 2026 to the Dancing Salesman Problem series: (a) *deconstructed* style (b) *pop* style (c) *ghosts* style (d) *watercolour* style, and (e) *fuego* style.