

Heroic vs Collaborative AI for the Arts

Mark d’Inverno

Goldsmiths, University of London
London, UK
dinverno@gold.ac.uk

Jon McCormack

Monash University
Caulfield East, Australia
jon.mccormack@monash.edu

Abstract

This paper considers the kinds of AI systems we want involved in art and art practice. We explore this relationship from three perspectives: as artists interested in expanding and developing our own creative practice; as AI researchers interested in building new AI systems that contribute to the understanding and development of art and art practice; and as audience members interested in experiencing art. We examine the nature of both art practice and experiencing art to ask how AI can contribute. To do so, we review the history of work in intelligent agents which broadly speaking sits in two camps: autonomous agents (systems that can exhibit intelligent behaviour independently) in one, and multi-agent systems (systems which interact with other systems in communities of agents) in the other. In this context we consider the nature of the relationship between AI and Art and introduce two opposing concepts: that of “Heroic AI”, to describe the situation where the software takes on the role of the lone creative hero and “Collaborative AI” where the system supports, challenges and provokes the creative activity of humans. We then set out what we believe are the main challenges for AI research in understanding its potential relationship to art and art practice.

1 Introduction

This paper sets out to specifically address the theme of this special track by considering the possible relationships between AI and the Arts. We will consider three related questions.

1. What kinds of systems should we as AI researchers interested in supporting and understanding art practice be designing?
2. What kind of AI systems should we as artists and musicians want in order to challenge, explore and expand our own creative practice?
3. What kinds of AI involvement in the process of producing art would satisfy current or new audiences?

In a general sense, AI research attempts to design, build and understand computational systems that undertake activities which are normally thought to require human intelligence. In a similarly general sense, “Art” can be defined as the application of *human skill and imagination* to produce work that has *beauty or emotional power* [def, 2015]. Using these descriptions, exploring the relationship between AI and Art would seem to present us with at least two high-level objectives described below.

1. Does considering art and art practice as a specific aspect of human intelligence help us drive AI forward as a discipline?
2. How can AI assist in the application of “human skill and imagination”, and in the production of objects of “beauty and emotional power”?

In this paper we set out to explore the relationship between AI, Art and Art Practice that can help guide AI research in the future. We do this by considering this relationship from three different viewpoints: (i) as artists interested in the possibility of using AI systems to produce art, (ii) as AI researchers interested in developing novel AI systems that can be used in the practice of producing art, and (iii) as potential audiences for art produced (partly or fully) using AI.

This leads us to the following three broad sets of questions:

- **As artists.** What kinds of AI systems do artists and musicians *want* to use in their art making? How might artists want to interact with AI systems? How much autonomy would artists want to defer to an AI system in the production of art? Would artists be happy to set up autonomous AI systems that would then go on to create art on behalf of the artist? And to do so repeatedly without further intervention? Can we envisage artists wanting to bestow a sense of creativity on the AI system itself? Is there a greater opportunity for artists to reflect on and develop their own creative practice through interaction with AI systems? How might AI systems best challenge, provoke, and stimulate artists in the future?
- **As researchers.** On what kinds of systems and techniques should we in the AI community be focussing our efforts? What can we learn about both the potential and limitations of AI when trying to build systems that can produce artefacts that if produced by humans might

be considered art? How can we build systems that enable artists to invest human artistic sensibilities into algorithms so that the outputs of the system have the qualities of art and that reflect us as artists in style, aesthetic or process? How does thinking specifically about art, rather than other aspects of intelligence, drive AI forward as a discipline?

- **As audiences.** Can we envisage scenarios where audiences would sustain interest in art produced solely by an autonomous AI system? What kinds of evolving artefacts might sustain our interest for an extended period of time? What kinds of involvement by AI in art practice would audiences be interested in understanding or knowing more about? Could audiences become interested in the process of an artist interacting with AI systems to produce art? What are the new opportunities for revealing the nature of artistic process itself through interacting with AI systems? And could that revelation be used to heighten the audience experience of art?

In order to consider these three different perspectives we introduce two distinct views on the role of AI systems in producing art. This distinction will help us answer these different sets of questions, and identify the range of potential AI interventions or collaborations in art. We label these different approaches to design and intention in AI systems as *Heroic AI* and *Collaborative AI*.

The first approach, *Heroic AI*, can be summarised as systems designed to produce art autonomously without the direct intervention of the artist or programmer. In instances where the software is written by a researcher or programmer, no artist may be involved at all. In recent years, a significant amount of AI research has been undertaken to develop such systems where the software assumes the role of a lone “creative hero”. We introduce the term “heroic” to emphasise the intended or implicit role of such systems (and their creators) being close to that of the heroic artist who struggles (often in adversity) to create great works of art. The kinds of challenges here are often described in terms of producing novelty or value, or in producing *objects* that if they had been created by humans would have been described as art. Whilst there have been notable successes, researchers remain fixated on the challenge of such systems creating artefacts that can sustain the interest or engagement of human audiences.

The second approach, *Collaborative AI*, encompasses those systems designed to augment or support the individual or collective practice of human artists. The challenges here are understanding issues such as interaction, flow and feedback [Csikszentmihalyi, 1997; Jones *et al.*, 2009] when working with technology. For example, how can an artist interact with a system so that they can offer the kinds of feedback and sense of collaboration that occurs when working with human collaborators? In general, can we build systems that can positively influence the artist’s approach as the work progresses, without “getting in the way”?

We will next explore these classes in the context of artistic AI system in more detail.

2 Heroic AI

When Deep Blue beat world champion chess grandmaster Gary Kasparov in 1997, many saw it as a turning point for both AI and human intelligence [Campbell *et al.*, 2002]. In the years prior, most scoffed at the idea of a computer playing at even the level of an average human player. But incremental improvements to chess-playing algorithms gradually delivered breakthroughs, leading eventually to the bettering of the best human player. And if a machine could beat the best human chess player, the implication was that it could better the best humans in other activities that require intelligence as well.

Today, any number of computer chess playing programs are regularly able to beat almost any human player and beating a grandmaster does not make front page news. Learning and playing chess has certainly been enhanced by AI, yet watching two chess playing programs battle each other isn’t interesting to the public. Nor is there huge interest in watching humans play against computers. Why? Because chess is much more than just the description of chess encapsulated by the computer.

Undoubtedly, if researchers did develop an “autonomous computer artist” that could produce work comparable with the best human art, it would be an extraordinary achievement. The researchers would be lauded as AI heroes for an amazing breakthrough. Others would no doubt lament that perhaps the last bastion of what is uniquely human had finally been toppled by a machine, albeit one designed by humans. But after this celebration and lament, how differently would we really view human creativity? Perhaps one way to identify what we mean by human creativity is precisely that which automation cannot do, so whenever AI systems improve it shifts our interpretation to whatever is not currently possible by machine.

While trying to create a machine that makes art as competently as the best human artists is a goal which is likely to move forward the discipline of AI, perhaps the greater challenge is to consider how it changes our view of art and art practice as artists and audiences. Unlike chess, art is not a fixed game with a formal set of rules and defined outcome. Art evolves and changes with culture, it is a complex system of social dialogues and inter-relationships between people, objects and environments. That environment includes technoscience (and AI technoscience) which continues to have important impacts on art itself.

It is relevant here to consider how we assign little value to what we might call “art” made by other species, except perhaps for its immediate novelty. While “elephant art” or “primate art” occasionally piques our interest, clearly the value we assign to it does not approach that of human art (we don’t build museums for animal art, for example) [Weesatchanam, 2006]. Could we ever imagine a gallery of “autonomous computer art” sustained by public or private funding? Even for our own species, children’s art is not valued in the same way as that of adult artists. We might expect an even more difficult situation for “heroic computer artists” because they are not even mammals.

Heroic AI attends on an anthropocentric creativity, one

where the software or system is conceptualised as the lone “hero”, a personification of certain romantic and Byronic notions of art and artists that find their origins in 19th century Romanticism [Cookson, 1937]. With AI on an heroic mission to bring a new form of creativity into the world, this tradition is renewed, albeit with a new context and a greater ironic twist. Heroic AI systems focus on a fixed notion of what their creators believe to be art and creativity, one that privileges the aesthetic art object. The real hero is the programmer, not the software or the artefacts it produces. It is not surprising that competent technical researchers are often drawn to Heroic AI, because it is seen as a grand technical challenge.

In reality, art is a continuing process of negotiation. It can ignite the broadest palette of human emotion and motivation: joy, sadness, excitement, discomfort, reflection, disgust, wonderment, pleasure, despair, collaboration and action. Art reflects something of what it is to be alive and to be able to communicate this shared experience with others. Philosophers such as O’Hear dismiss the idea of machines originating art entirely, because art “in the full sense is based in human experience” and requires a communication between artist and audience drawn from that shared experience [O’Hear, 1995]. Any “communication” of human experience by machine can only be parasitically meaningful, being a pastiche of real human creativity.

However, any concerns we have to the Heroic AI approach are not due to objections such as those of O’Hear (one can at least imagine the possibility of having some meaningful communication with prior hominoid species or alien intelligences for example). Nor do past concerns of AI (issues of embodiment, turing completeness, symbol grounding, etc.) determine our position. We see no in-principle reason to rule out the possibility of Heroic AI being able to create outputs that meet or exceed human capabilities, even without “being human” in any sense.

Our main concern of Heroic AI approaches is an inability to consider wider social and cultural context in which we experience art. Even if computers could offer something new or different to art, could it ever be properly valued in a social context? Being an artist, making art – is a choice. The fact that human artists intentionally set out to make work, and underwent a process to do so, is an important aspect of the way in which we engage with art. A computer or machine programmed to make art cannot make that choice.¹ It might be argued that biology has “programmed” us to make art, just as the male Bower bird is genetically disposed to constructing a beautiful nest to attract a mate. But becoming a practicing artist and understanding the social context of art is more nuanced than a binary decision or even a genetic predisposition.

2.1 Automating the creative process

Thinking about creative activity more generally, we should be concerned when technology assumes the main creative role and limits the potential for human creative expression. Criticisms of the automation of creativity are not new and we dis-

¹Programming the machine to make this choice does not circumvent the fundamental problem either.

cussed them elsewhere [McCormack and d’Inverno, 2012b], highlighting automation’s reductive effects on personal creative development. Deferring creativity to a machine encourages a creative lethargy on behalf of human users. In current incarnations (e.g. Camera smile detection, “retro” image filters, or “Content-aware” filters) automating creativity homogenises the creative artefact to a singular concept of value.

A primary impetus for machine automation is timesaving: by freeing us from tedious or repetitive task we will have more “spare” time (for leisure, for example). Despite increasing automation in many aspects of working life, the average time spent at work in developed countries has not changed significantly in the last fifty years [Gershuny *et al.*, 2003]. Yet many people report subjective time-shortening, or “time-famine” effects brought about by the always-on, networked society and often blame this “famine” on technology [Wajcman, 2015].²

3 Collaborative AI

We take Collaborative AI as research which is concerned with designing AI systems that provoke, challenge and support human art practice. As AI researchers we are interested in building such systems because AI has much more to offer art than attempts at mimicking the creative process itself. A wide range of AI research opportunity opens up if we set ourselves the challenge of building computers that can become our creative partner or where new forms of creative activity are developed by meaningful collaboration between human artist and machine agent.

Artists often require reflection, feedback, reminding, connections, stimulation and interaction; all of which can potentially be offered by AI systems. Artists need sources of new ideas, feedback on ongoing ideas, as well as a sense of one’s own creative practice. AI systems can help the artist develop a greater understanding of their creative process and how they work best with others. Collaborative AI improves *everyone’s* creative activity, empowering people to be at their creative best. This potential seems a much more worthwhile goal than that of the Heroic AI approach and opens up clear avenues for AI research.

One important example is the Flow Machines project [Ghedini *et al.*, 2013], which specifically asks the question: “Can machines make us more creative?”. This research is inspired by the insight that the future of content generation is in ourselves [Pachet, 2008] and that we can use AI technology to bring out the best in ourselves, being better than we could ever be without it.

In other words the idea is one of building Collaborative AI agents that can help us develop our own artistic and musical practice. One example of such a system is known as the Reflexive Looper [Pachet *et al.*, 2013] which is shown in Figure 1. In this system a pianist or guitarist plays in a whole range of different patterns but over a specific known chord sequence such as that of a standard. The system then starts to “jam” with the musician, deciding what parts of the input source material to use to perform in response to the live

²Although such effects are not uniformly accepted.

musician. In this way a solo guitarist or pianist can start to *collaborate* with the agent looper.

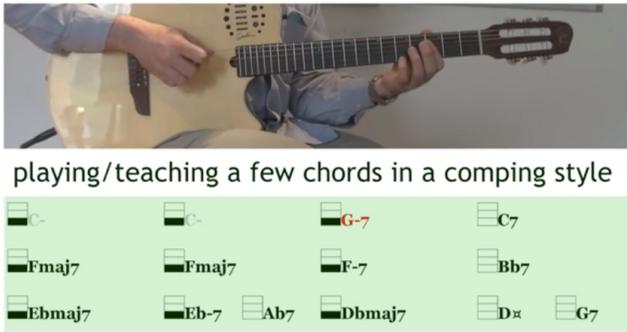


Figure 1: *The Reflexive Looper in Action.* The guitars first plays a bass line over a chord sequence from a jazz standard. In this case the Miles Davis Tune “Solar”. The guitarist inputs chord shapes as the looper plays back a bass line. Eventually the looper has enough material that the musician can start to play with several copies of himself.

The looper can take on one, two or more instantiations of a guitarist (for example) playing a bass line, a chord line and a solo, with each of these instantiations responding to the live guitarist. So the system shares the goal of creating great music, and it has the goal of trying to respond to the live guitarist in order to create the best “ensemble” sound possible.

The musician’s creativity is challenged and stimulated by playing with responsive copies of themselves. Sometimes they come up with fantastic musical creations that would not have been possible with the musician playing alone without the looper. However, this works because the agent looper is not expected or desired to be autonomous in any way. It does not have any of its own motivations and musical intentions. Instead it is designed to serve the musical goals and motivations of the autonomous musician. Of course the looper could be programmed to be more autonomous, to have its own sounds, its own performance ideas but it would not then be able to collaborate with a human user in such a successful way.

In summary, we see Art as a social enterprise that encompasses human communication and experience. Imitations or parodies only lessen the human condition (consider a film such as *The Truman Show* for example). Collaborative AI on the other hand, presents a broad set of possibilities for AI and Art, including expanding the creative activity of everyone. In order to distinguish these different types of system within a traditional AI framework we next look at the concepts of agency and autonomy that have been present in AI for several decades. Understanding these concepts will be crucial to exploring Collaborative AI relationships from the perspective of artist, researcher and audience, that we will explore in the final part of the paper.

4 Revisiting Agency and Autonomy

Here we revisit notions of agency and autonomy in order to ground our ideas in mainstream AI research. There are es-

entially two main streams of agent work over the last 30 years: autonomous agent architectures and multi-agent systems. These are the direct pre-cursors to our heroic and collaborative agents in art.

AI has its origins in the ambitious idea of building a single intelligent autonomous agent that equalled or bettered human intelligence. Systems such as PRS [d’Inverno *et al.*, 2004], Soar [Li and Laird, 2014] and Cyc [Lenat and Guha, 1993] are typical examples from the early days of AI that tried to achieve this goal. Applying this classification to art-producing systems we would see these agents as heroic lone artists capable of independently producing a series of artworks over a sustained period.

As opposed to single-agent architectures the other focus of research has been into multi-agent systems. Research has addressed issues of interaction, communication, negotiation and collaboration [d’Inverno *et al.*, 2012]. Extending this to the world of art, this approach is about building systems that could collaborate with other agents (typically human) in order to produce work that the human would not be able to do (or not so easily) without this relationship.

From the paintbrush to the piano, artists have always used tools (or “agents”) to create their art and AI offers a whole new range of interactions that are possible because of computing. Systems that can respond to ongoing creative acts in an artistic process can propose alternatives, can put together previous content in new ways and can provide feedback or analysis of our ongoing work and so reshape our activity – even as we are undertaking it. Thus AI systems can provide a much wider range of collaborative possibilities because of the analytical, generative and adaptive features of such systems. They offer new creative routes based on a dynamic awareness of context and past history. Computational systems can give us a much wider palette of opportunity than traditional non-digital tools for making art and in order to understand this relationship more clearly we next identify the distinction between agency and autonomy.

4.1 Distinguishing Agency and Autonomy

The SMART Agent Framework [Luck and d’Inverno, 1995] was the first attempt to define agency and autonomy in computational systems and explicate the distinction and relationship between them. We summarise the model briefly here by outlining only the basic components informally (Figure 2).

Attributes are the fundamental component and represent anything that can be perceived. An *environment* consists of a set of attributes. Some of these attributes can be grouped together to form *entities*. Entities therefore also comprise a set of attributes that can be perceived as a whole (such as a tree) by others. *Objects* are then special kinds entities that have the ability to change the state of the environment. Objects can include a paint brush because it has the ability to place paint on a piece of paper for example.

Any object can be considered as an *Agent* as long as it is serving a purpose for another agent. In other words, that the agent is serving a *goal* where a goal is some description of the state in the environment. (“brush is clean”, “brush has red paint”, “room is dry”, and so on, are examples of goals that could be ascribed to a paint brush). In this model, the paint-

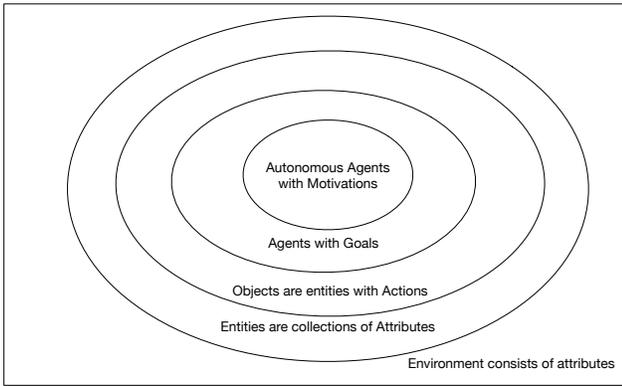


Figure 2: *The SMART Framework for Agency and Autonomy.* The environment consists of attributes which can be perceived. Some of these attributes may be grouped together into entities. Entities that have the ability to change the environment by adding and removing attributes through actions are called objects. Those objects which are serving a purpose in the world are called agents. And those agents who do not rely on other agents for purposeful behaviour and can generate their own goals through higher-level intrinsic motivations are autonomous agents.

brush is seen as an agent whilst it is serving the goal of the artist using it to produce an artwork. The fact that the paintbrush is serving the goals of the artist sets up an important social relationship (if anyone else tried to remove the paintbrush whilst it was serving as an agent for the artist then it would clearly be met with resistance). However, as soon as the paintbrush is discarded by the artist as no longer being useful, then the agency of the paintbrush is lost and reverts to being an object.

Here the paintbrush is heavily reliant on the goals of others to become an agent because it cannot create its own goals. It therefore requires more complex agents (an elephant, a human, or even perhaps an AI system) in order for those goals to be generated and then given to it. This facilitated the description of autonomy – or autonomous agents – to be defined as those entities that could create their own goals. Such agents did not rely on other agents for their own self-determined purposeful existence. Clearly a paintbrush cannot create its own goals, but elephants and humans can. This led to the definition of intrinsic higher-level motivations – which did not relate to descriptions of the world-like goals – to define autonomous agents. So, from the perspective of the artist, the researcher and the audience, what do the issues become about the relationship between agency and autonomy? This we consider in the next section.

5 How much Autonomy do we want in Collaborative AI?

We explore this from the point of view of our three different perspectives.

As Artists. How much autonomy would we actually want to

give to a machine about the artistic choices that are made in any process? Harold Cohen – an artist who decided to use technology to build his own AI system called AARON [Sundararajan, 2014] – is happy to speak informally of AARON in the sense of having its own intrinsic motivations and in the way it approaches a new piece of work. And once the system starts to create an image there is no intervention from a human until that image has been created. So in one sense the system is autonomous in its action but on the other hand the system is clearly vehicle for the artistic motivations of Cohen, an established and successful artist. It would be very hard in the goal/motivation model of agency and autonomy to think of the system as having its own artistic motivations.

It is worth pointing out that though every output is different, Cohen is always re-coding the system to produce new kinds of behaviours and aesthetics. He also heavily edits the voluminous output down to only what he considers the very best works. This is not the autonomous journey of a human artist, but a new set of criteria that come from the autonomous human artists that essentially instantiates a new kind of agent each time the code is edited. Indeed, as artists it is almost impossible to conceive of a situation where we would have any need for a Heroic AI system generating its own goals from its higher-level motivations. The system will always embody the goals of the artist, just as Damien Hirst’s team of autonomous human workers will always embody his.

What about if we think of the system not as an artist in its own right but as a potential collaborator? Many of us have shared the joy of collaborating and it will often stretch and challenge us in ways that would not have been possible through a lone artistic journey. The Lennon in all of us wants a McCartney for collaboration. Even when we think of the system as a collaborator the question is still how much autonomy do we want in a computational collaborator? Do we want it to have its own artistic goals and sensitivities that may be quite different from our own? Would we really respect the artistic sensitivities and motivations of the system sufficiently that and see them as an equal in the same way that Lennon and McCartney clearly saw each other?

What we suspect here is that artists would be keen to collaborate with AI systems as long as it was clear that the collaboration in the end was simply about serving the goals of the artists. We want systems with sufficient agency to provoke, challenge and collaborate but in such a way that all of the artistic motivations, artistic choices and sensitivities belonged to us the artists. Again there is no sense that the artificial collaborator would rely entirely on itself for a purposeful existence. We would see it as an agent whose goal was to serve the artistic ambitions of the artist. This is not to say of course that designing such systems is any easier or harder than Heroic AI architectures: only that the ambition underlying the design of the different classes are quite different.

As Researchers. It is clearly an attractive and seductive goal to produce a system capable of producing work that – if produced by a human – could or would be seen as “art”. This embodies and refreshes the original exciting ideas of artificial intelligence as set out by McCarthy and friends in the 1950s, but now the focus is on “human creativity” as opposed to “human

intelligence". This has been a successful field of AI in the last 10 years thanks to the efforts of a handful of key researchers in the field of computational creativity [Boden, 1994; 2009; Colton, 2012; Colton *et al.*, 2009; Cardoso *et al.*, 2009; Wiggins, 2003; 2006]. There has been much media, press, TV, grant-funding opportunities related to the development of AI activity opening up all kinds of debates about the nature of creativity and the nature of art. Examples include the BBC's "Can a computer write Shakespeare" [bbc, 2014] and another the Wired article: "Artificial artists: when computers become creative" [wir, 2013].

What is interesting about this work are the related goals that underpin their design. Is it a demonstration of what is possible by bringing together a range of different techniques from AI and showing how different aesthetic and stylistic concerns can be embodied in algorithms? Or is the goal to design a system to autonomously build music, text or images? Or lastly, is the goal to develop notions of individual creativity which challenge the idea that only humans can be creative?

Despite this initial success, the challenge of building a Heroic AI design that will sustain interest to audiences now seems harder than it did even a few years ago. On the other hand, there will always be a desire from artists and creatives to use Collaborative AI systems that can stimulate and promote creative activity. Systems that we can collaborate with to extend our own creative activity, that can remind us of past processes, approaches or outputs at the right moment, could sustain and develop our own creative practice. In this view we can be clear that the motivations driving the artistic process belong solely to that artist and there is no claim from anyone that the artificial system plays an equal or greater role.

As Audiences. Does art exist without an audience that is willing to experience it as art? One of the fascinating things about the experience of art is to unpick the emotional, motivational and cognitive state that the artist was in when they were producing the art in question. Art leaves traces of that. Even if we could create an AI system with motivations including a desire to please, a desire to be respected, a desire to mate, desire to keep learning to a be better at something and so on – would we really ever be that interested in the output? There is also the question of social, political, economic contexts as well as the personal and emotional ones. Could an audience empathise with the plight of the AI system? Would they be open to understanding the nature of the process? The process of the algorithm and the challenges that were met when trying to build the system and to see this as an important part of appreciation the final output as art?

We can imagine various Turing-type tests where an audience is made to believe that the performer or artists is human. But that simply tells us more about how technology is improving and what is possible functionally. It tells us little about the human condition in the way allows us to explore what it is to be human. No doubt software systems can do wonderful things too, but can we ever be interested in the outputs of these systems as art if we can't relate it to understanding the programmer as an artist themselves?

6 Conclusions

In this paper we have tried to set out how AI can play a part in art and art practice by categorising effort into two distinct camps. In the blue corner we have the autonomous AI agent as the heroic artist, an autonomous intelligent creative system with its own sense of aesthetics able to make its own artists choices and produce works of "art". In the red corner we have the collaborative computational system as an agent of the artist. Helping the artist move through the myriad of imaginative, artistic, evaluative, and editorial choices that unfold in every moment.

We have tried to set out what research questions the Heroic AI approach could be addressing by considering the perspective of artist, audience and researcher. We have argued that a stronger and more coherent focus comes from a Collaborative AI approach, which aims to build systems that give human artists new modes of artistic activity, interacting in ways not available through non-digital tools such as paintbrushes and pianos. In this way AI can be driven by the demands of artists who are open to working in new ways, and this focus can enliven and challenge AI as a research discipline.

One might think of this as the "participatory design of AI research" where the needs and desires of artists help map out clear avenues of AI research.

Art and Music are currently evaluated in the context of human experience and the effect they leave on their audience. Art broadens and expands the human condition like no other human activity, so being able to assist humans in continuing this creative journey through Collaborative AI is a challenging and credible goal for AI and the benefit of humanity. And whilst AI can enrich the creative process of artists, the field of AI as a research discipline can itself be enriched by the ongoing challenges that artists bring to the table.

Finally, it is worth pointing out that the biggest recent successes of AI have not been in heroic anthropomorphic intelligences. Rather, this has occurred in our networked society, bringing an era of "big data" and "deep learning", where global mediated social interactions provide the empirical data for machine prediction. Largely in the service of corporate or government knowledge wealth, these systems consider humans from the perspective of data points clustered on a landscape. This approach is more about prediction or classification than supporting creative activity and is most commonly used for marketing purposes. What the long term repercussions of this "datafication" of human desire and creativity are remains an open question. Collaborative AI presents a wonderful opportunity to develop AI research to enable, support and enhance social human creative activity.

Acknowledgments

This research was supported by an Australian Research Council Discovery Projects grant DP1094064 and by the FP7 Practice and Performance Analysis Inspiring Social Education (STREP FP7- 318770).

References

- [bbc, 2014] Can a computer write shakespeare? <http://www.bbc.co.uk/programmes/b0435kdk>, 2014. Accessed: 2015-02-12.
- [Boden, 1994] Margaret A. Boden. *Dimensions of Creativity*, chapter What is Creativity?, pages 75–117. MIT Press, Cambridge, MA, 1994.
- [Boden, 2009] Margaret A. Boden. Computer models of creativity. *AI Magazine*, 30(3):23–34, Fall 2009.
- [Campbell *et al.*, 2002] Murray Campbell, A. Joseph Hoane, Jr., and Feng-hsiung Hsu. Deep blue. *Artif. Intell.*, 134(1-2):57–83, January 2002.
- [Cardoso *et al.*, 2009] Amílcar Cardoso, Tony Veale, and Geraint A Wiggins. Converging on the divergent: The history (and future) of the international joint workshops in computational creativity. *AI Magazine*, 30(3):15–22, 2009.
- [Colton *et al.*, 2009] Simon Colton, Ramon L. De Mantaras, and Oliviero Stock. Computational creativity: Coming of age. *AI Magazine*, 30(3):11–14, Fall 2009.
- [Colton, 2012] Simon Colton. The painting fool: Stories from building an automated painter. In McCormack and d’Inverno [2012a], chapter 1, pages 3–38.
- [Cookson, 1937] George Cookson. *The Decline and Fall of the Romantic Ideal*, volume 1. Cambridge University Press, 1937.
- [Csikszentmihalyi, 1997] Mihaly Csikszentmihalyi. *Finding flow: The psychology of engagement with everyday life*. Basic Books, New York, N.Y., 1997.
- [def, 2015] Definition of art. <http://www.oxforddictionaries.com/definition/english/art>, 2015. Accessed: 2015-02-12.
- [d’Inverno *et al.*, 2004] Mark d’Inverno, Michael Luck, Michael Georgeff, David Kinny, and Michael Wooldridge. The dmars architecture: A specification of the distributed multi-agent reasoning system. *Autonomous Agents and Multi-Agent Systems*, 9(1):5–53, 2004.
- [d’Inverno *et al.*, 2012] Mark d’Inverno, Michael Luck, Pablo Noriega, Juan A. Rodriguez-Aguilar, and Carles Sierra. Communicating open systems. *Artificial Intelligence*, 186(0):38 – 94, 2012.
- [Gershuny *et al.*, 2003] Jonathan Gershuny, Kimberly Fisher, and Ather Akbari. Statistics on working time arrangements based on time-use survey data. Research Report Conditions of Work and Employment Series No. 3, International Labour Office – Geneva, 2003.
- [Ghedini *et al.*, 2013] F. Ghedini, P. Roy, and François Pachet. The flow machines project, 2013.
- [Jones *et al.*, 2009] Daniel Jones, Oliver Bown, Jon McCormack, François Pachet, Michael Young, Rodney Berry, Iris Asaf, and Benjamin Porter. Stimulating creative flow through computational feedback. In Jon McCormack, Margaret A. Boden, and Mark d’Inverno, editors, *Dagstuhl Seminar 09291: Computational Creativity: An Interdisciplinary Approach*, volume 09291, pages 1–10, Germany, 2009. LZI (Leibniz-Zentrum für Informatik GmbH).
- [Lenat and Guha, 1993] DB Lenat and RV Guha. Building large knowledge-based systems: Representation and inference in the cyc project. *Artificial Intelligence*, 61(1), 1993.
- [Li and Laird, 2014] J Li and JE Laird. Spontaneous retrieval for prospective memory: Effects of encoding specificity and retention interval. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*. AAAI, 2014.
- [Luck and d’Inverno, 1995] M. Luck and M. d’Inverno. A formal framework for agency and autonomy. In *Proceedings of the First International Conference on Multi-Agent Systems*, pages 254–260. AAAI Press/MIT Press, 1995.
- [McCormack and d’Inverno, 2012a] Jon McCormack and Mark d’Inverno, editors. *Computers and Creativity*. Springer, Berlin; Heidelberg, 2012.
- [McCormack and d’Inverno, 2012b] Jon McCormack and Mark d’Inverno. Computers and creativity: The road ahead. In *Computers and Creativity [2012a]*, chapter 16, pages 421–424.
- [O’Hear, 1995] Anthony O’Hear. Art and technology: An old tension. *Royal Institute of Philosophy Supplement*, 38:143–158, 1995.
- [Pachet *et al.*, 2013] François Pachet, Pierre Roy, Julian Moreira, and Mark d’Inverno. Reflexive loopers for solo musical improvisation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2205–2208. ACM, 2013.
- [Pachet, 2008] François Pachet. The future of content is in ourselves. *Computers in Entertainment*, 6(3), 2008.
- [Sundararajan, 2014] L Sundararajan. Mind, machine, and creativity: An artist’s perspective. *Creative Behaviour*, 48, 2014.
- [Wajcman, 2015] Judy Wajcman. *Pressed for Time: The Acceleration of Life in Digital Capitalism*. University of Chicago Press, 2015.
- [Weesatchanam, 2006] Aum-Mon Weesatchanam. Are paintings by elephants really art. *The Elephant Art Gallery*, 31, 2006.
- [Wiggins, 2003] Geraint A Wiggins. Characterising creative systems. In *Proceedings of the IJCAI*, volume 3, 2003.
- [Wiggins, 2006] Geraint A. Wiggins. A preliminary framework for description, analysis and comparison of creative systems. *Know.-Based Syst.*, 19(7):449–458, November 2006.
- [wir, 2013] Artificial artists: when computers become creative. <http://www.wired.co.uk/news/archive/2013-08/07/can-computers-be-creative>, 2013. Accessed: 2015-05-02.