

Computers and Creativity: a Roadmap

Jon McCormack¹ and Mark d’Inverno²

Abstract. In this paper we introduce and reflect on current research questions regarding computers and creativity. Creativity is an enigmatic yet widely discussed phenomena. With the now widespread adoption of computers and information technologies, the nature of creativity and how we think about it has changed significantly. We argue for a shift in thinking about computers from tools to creative agents and collaborative partners. We present 21 questions we think are crucial to understanding this new relationship and begin to offer answers, or pathways to answers for a selective subset.

1 Introduction

It is almost impossible now not to hear of some new discourse, activity or emerging project that links computers with creativity. Either through re-thinking the way that computing is taught in Schools and Universities as a creative and practice-based subject, or through creating novel kinds of computer systems than can enhance our own creativity, or through the development of innovative software or robotics that can produce works that – if they were produced by a human – would be described as art (such as a painting, a poem, or a piece of music for example). For many years we have both been involved in research trying to understand how computing and creativity are connected. Furthermore, we are both active creative practitioners as a professional artist and musician respectively, often investigating how we can use technology to enhance our own creative practice and performance.

As a result of editing the book “Computers and Creativity” [29] over the course of three years, we posed 21 questions that we believe are central to research linking computers and creativity, and which map out areas of investigation we believe will be the focus of attention in the coming years and decades. In this paper we repose these questions and then offer answers or at least outline promising research that makes progress towards answering some of them. By doing so we hope to provide an introduction to the issues, themes and research areas that are emerging when embarking on any activity attempting to understand or exploit relationships between computing and creativity.

Before reviewing the questions and proposing some responses to them in detail, we first provide some context and background to the emerging impact of computers on creativity and society.

2 The Creative Computing Landscape

Computing has rapidly grown from its origins as a selectively focused, highly specialised technical discipline and profession to one

of the most important drivers of modern economies and cultures. Networked computer technologies have facilitated seismic changes to industry, government, education, business, culture, science and society. At least in the developed world, we find ourselves deeply engaged in activities that are completely reliant on technology and which impact on almost every aspect of contemporary living.

Computers have become an extension of ourselves [9], mediating how we communicate and think, even changing the way in which we think. They control a complex network of dependencies between us, and are constantly and rapidly developing, ever expanding in their potential to become dynamic cultural and creative partners for us all.

While schools are increasingly using technology in education, ironically a 2010 ACM study [45] found that secondary (K-12) US education in computer science had fallen dramatically, with advanced placement courses declining by 35 % over the years 2005 to 2009. The report also suggested educators confused the use of technology and teaching of technology literacy with computer science education as a core discipline. Similarly, in September 2011, Google chairman Eric Schmidt criticised UK education, claiming: “Your IT curriculum focuses on teaching how to use software, but gives no insight into how it is made. That is just throwing away your great computing heritage.” Fortunately this helped in changing the way in which the computing-related syllabus will now be taught in schools in the UK, but there is still a long way to go before we think about teaching programming in the same way we would teach music, art, design or poetry.

Schmidt is not alone in his concern about a lack of basic programming literacy. Douglas Rushkoff’s book “Program or Be Programmed” prophetically argues that “In the emerging, highly programmed landscape ahead, you will either create the software or you will be the software.” [42] The concept of unknown technological forces watching us, analysing us, categorising us and predicting us is increasingly written about in social sciences and often adds to a general malaise for ever wanting to know the full extent and the underlying agency behind much of the software widely in use.

Why does this stark disparity exist between computing’s astonishing impact and our cultural desirability to see it as a creative discipline or even our ability to recognise the computer as a genuine creative partner?

Our claim – a view of course which is shared by many though seldom achieved in practice – is that Art and Science need to be brought back together if we are to better appreciate and engage with the challenges and opportunities this new dependence on technology brings. Creativity is critical for our ability to develop as a society, yet the mainstream practice of computing has not formally situated itself around the exploration of creativity and creative ideas. Rather it has been approached from a scientific and engineering perspective that aims to represent aspects of the world as data, and then has the task of manipulating that data in order to solve problems or gain under-

¹ Centre for Electronic Media Art, Faculty of Information Technology, Monash University, Caulfield East, VIC 3145, Australia, email: Jon.McCormack@monash.edu

² Department of Computing, Goldsmiths, University of London, UK, email: dinverno@gold.ac.uk

standings.

Unsurprisingly, it is artists who most quickly recognised the potential for exploring new kinds of creative computational processes, engaging with computers since their earliest days. Indeed, the concept of “Creative Coding” has emerged as an artistic practice of rising popularity [40]. Software is not just for problem solving: it is a new medium for creativity and creative expression that has been enthusiastically embraced by many artists, designers and musicians. Software undergoes development at a pace and complexity that far exceeds all prior technologies we have thus far developed, so these practitioners see the computer as something more than a benign artistic tool such as a paintbrush. Nonetheless, many artists find their artistic expression limited by a lack of knowledge in how to program intuitively and creatively, hence the increasing number of courses that specifically cater for artists wanting to learn more about programming in order to meet their creative needs (e.g. the Programming for Artists course which is part of the MFA in Computational Studio Arts at Goldsmiths the second author’s home institution).

Our previously mentioned book, “Computers and Creativity” [29], presented computer programming as a creative practice, describing how the computer can augment human creativity and help us to understand different kinds of creativity formally. As noted in the book’s foreword by Margaret Boden (a celebrated UK cognitive scientist and philosopher who perhaps more than anyone has brought the discussion of creativity into the mainstream of scientific investigation), understanding programming as an artistically creative act offers a relatively new perspective on working creatively with computers [6].

Boden uses the example of the painter David Hockney’s 2012 exhibition at the Royal Academy of Arts in London, (“A Bigger Picture”) [20]. Hockney is widely acknowledged as one of Britain’s greatest living painters, with a career spanning more than fifty years. He has recently turned to using computers, enthusiastically adopting the Apple iPad and its “Brushes” painting app as his painting tools. Boden argues that many people will see these digital paintings as exemplars of art made by computers (“computer art”). But the novelty, skill, and creativity that are evident in these works are not due to the software anymore than we would attribute them to a physical paintbrush. Used in this way, the intent of the software’s design is simply to support an individual artist’s creativity, not to take an active role or be responsible for it in any way.

The artistic practice of writing computer code is fundamentally different than writing a software program such as Photoshop or Brushes, because the primary artistic decisions are made by the program’s author, not by subsequent users of the program. This requires the programmer to understand how the choices that are made within the algorithmic structure of the programming code determine the artistic meaning of what the software produces. In this sense, computer code is an expressive medium, intricately tied to the artwork it generates. With this mode of practice the artist is increasingly giving the computer more and more *creative agency*, as we discuss below.

In contrast, general-purpose “creative” software outwardly seeks to minimise the creative judgments of its programmers, deferring those decisions to the artists using it. As with any human-developed medium, software is influenced by the culture of its users and programmers [25] and this implicitly directs many aspects of the software design process (for example choosing to simulate the physical media of paint and canvas in a software painting program according to the Western tradition).

2.1 Creative Agency

In research that spans over 20 years, Mark d’Inverno (collaborating ostensibly with Michael Luck) has worked on defining the nature of agency in computational systems, in what is called the “SMART agent framework” [14, 24, 15]. In this framework, any object can be considered as an agent as long as it is serving a purpose for another agent. For example, a paintbrush would be seen as an agent whilst it was serving the goal of the artist using it to produce an artwork. 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 – within the SMART framework – it reverts to being an object.

Naturally 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 that are able to create their own goals. Entities that do not rely on others for purposeful goals and which can create their own goals are called *autonomous agents* whether they are human, animal or computational.

This is why the concept of agency is important in understanding how computational systems can create new kinds of agency that are not evident in non-computational tools. In [22], Jones, Brown and d’Inverno unpicked the notion of agency in creative partnerships, examining why relationships with computational systems provide a step-wise greater opportunity for collaborative engagement than with non computational agents such as a paintbrush. A computational system’s capacity for autonomy can allow it to operate with distinct agency in the creative process: the artist sets up a system with a given set of rules for its ongoing interaction with the artist or with others. For the system and artist to work together, the system cannot disregard this input and follow only its internal logic. Its output is mediated through some predetermined structure or ruleset, but it should always follow somehow from the previous actions of the artist. This forms the basis of a human-machine “creative conversation” or improvisation.

In any creative activity we can assign varying degrees of *creative agency* amongst the participants and tools [7], see Fig. 1. Typically as degree of creative agency increases in any participant, the less control the other has over it, i.e. the computer as tool gives us good control over it, but it exhibits little creative agency. The computer a collaborator exhibits much more, but in general this involves a greater degree of autonomy and independence, hence we have less direct control over it.

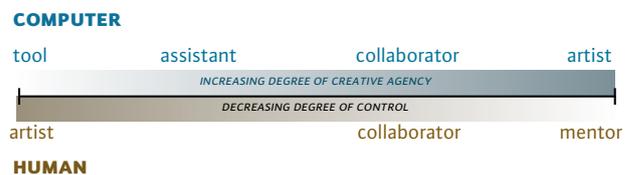


Figure 1. Types of creative agency and the relation between participants

In some sense both agency and creativity are in the eye of the beholder. If we see someone do something that surprises us through its originality or novelty then we think of that person or that act as a creative one. Furthermore, it is not unusual to hear of musicians talking after a good concert about instruments as having personali-

ties and even intentions. In the mind's eye of the performer there is a dialogue with the instrument as a collaborative agent. And with the sophistication afforded by state-of-the-art technology it will increasingly become possible to design systems that enable new kinds of creative collaboration in the minds eye of the user.

The reason this is potentially exciting from a social as well as a creative standpoint is because it can extend our collective creativity. People collaborate in constructive partnerships that may exceed their typical creativity as individuals (e.g. Lennon and McCartney). Artists such as Sol LeWitt and Bridget Riley made extensive use of assistants who created their paintings by interpreting the artistic instructions given to them. Partnerships and assistants typically involve different distributions of creative agency, but we would attribute some creativity to each human participant. Traditionally, the majority of creative agency originates from the person using the tool or playing a musical instrument, not from the tool or instrument itself.

However, we are now seeing new kinds of human-machine creative collaboration, such as where human artists improvise or jam with machines [3, 35] even research projects that aim to devise independent³ creative artists [27, 10, 11]. This provokes a shift in attribution of creative agency towards the computer, whose role is less like a conventional tool and much more like a creative assistant or partner and perhaps one day, an independent artist.

2.2 Defining Creativity

Thus far we have used the term “creativity” informally, without detailed explanation of its various interpretations and definitions. “Computers and Creativity” included many different understandings and definitions of creativity, as an act, artefact or outcome. These interpretations included formal mathematical definitions [43] and practice-based explanations [13]. A number of researchers downplay the concept of creativity itself, preferring terms such as “virtuosity” [34], focusing on skill, learnt behaviour and the intimate relationship between a person and their creative tool or instrument. There is also debate as to the significance of individuals, groups and environments in supporting human creativity [22]. For example it may be that creativity is purely a social construct in that it only makes sense to an observer, or the subsequent impact of a creative act on a community or society. Another view often attributed to Dewey is that all thinking is necessarily creative [19]. And yet other views include understanding classes of processes as creative (e.g. [38, 16, 28]), leading to broader appreciation of creativity in physical, biological, social and artificial systems.

This diversity of interpretations is strongly influenced by discipline and background, and in relation to computers and algorithmic definitions, the author's stance on whether autonomous creativity can exist in artificial systems (either in theory or practice) plays a crucial role. We will explore this issue in more detail in Section 4.1.

Like many others before us, we use as our starting point Boden's definition that creativity involves the generation of ideas or artefacts that are *new*, *surprising* and *valuable* [5]. Human creativity is commonly discussed in relation to artistic practices, often with an implicit sense of it being at its most “raw” and visceral. As artists, musicians and computer scientists this intersection is our main topic here, but we certainly view creativity as a general phenomena evident in nature, biology, society, and many different human activities.

Instead of focusing exclusively on creativity per-se, in this paper we are addressing the creative interaction of people and computers,

³ Independent in the sense of being autonomously creative and distinct from direct human intervention.

which necessarily encompasses a wide variety of activities and ideas. Our interest is not in how to automate or replicate human creativity on a computer, rather to examine how the computer can become a useful and stimulating creative partner, and taken to the fullest extent, become a creative entity in its own right through the on-going nature of that partnership.

2.3 Designing Creative Software

One of the great challenges for computing is to achieve a fuller understanding of processes and representations beyond what is easily computable or fully comprehensible by humans. Necessarily, human design of software requires reducing difficult and complex phenomena to simpler abstractions that can be practically implemented, in some cases even ignoring aspects that are too complex to express directly in a program.

Rather than trying to manually design creative software, one alternative is to design programs that are capable of initiating their own creativity to increase their complexity and discover ways of interacting independently of human design. But paradoxically we don't yet sufficiently understand human creativity in terms of formal algorithms, leading to a perceived gap between natural creative human expression and computation.

Enhancing creativity through more computer power does not necessarily follow either. When Seymour Cray was told that Apple had just purchased one of his CRAY supercomputers to design the next Macintosh computer, Cray responded that he had just purchased a Mac to design the next CRAY [2, p. 225]. At the same time, computers are undoubtedly capable of assisting, enhancing and augmenting our creativity, thus changing the nature of what it means to “be creative” in a connected, technological world.

Certainly there are critics of the study of computational creativity. Will it suffer the same philosophical and technical limitations that have been levelled at artificial intelligence or intelligent agents? Is it just another term in the history of this brand of computer science (artificial intelligence, then intelligent agents and multi-agent systems and now creativity) that will be replaced by something else in the coming decades?

An obvious objection that was levelled at artificial intelligence that can be equally said of computational creativity is that it does not make sense to even think about simulating or emulating creativity in machines for something that is a human, or biological, or a social concept [17, 37]. Some argue that creativity is such a nebulous concept that it would be hard to define anything that would really count as evidence that a computer was being creative. This question, and many related ones, will arise in the list of questions that we describe in the next section.

Whatever view we take about creativity one thing is clear: the discourse around creativity is on the rise. We hear it mentioned in school and higher education programs continually (e.g. creative musicianship and creative programming are just two examples). The “Creative Industries” have become a vital and growing part of major economies. A critical demand of nation-states is to become more creative in order to remain competitive [39]. The recent EU research-funding framework made creativity research a major funding priority, investing millions of euros on such research. We increasingly hear of the desire of many governments to invest heavily in their creative industries, seeing not only the transformative economic potential but also the social, political and cultural value of a highly creative population.

3 21 Questions: Computers and Creativity and the Road Ahead

We concluded “Computers and Creativity” with a final chapter that proposed twenty-one critical and important questions raised by the book [31]. We divide our questions and commentary into four topics: (i) how computers can enhance human creativity, (ii) whether computer art can ever be properly valued, (iii) what computing and computer science can tell us about creativity, and finally (iv) how creativity and computing can be brought together in learning.

How Can Computers Enhance Human Creativity?

- CC1: What are the kinds of interactions with computational systems that will inspire, provoke, and challenge us into meaningful creative dialogues with our machines?
- CC2: Relatedly, how can we remain mindful about the ways in which new technology can limit or defer creativity? Should we be concerned when creative decision-making is implicitly transferred to software at the expense of human creative exploration?
- CC3: Can we re-conceptualise current methods of interaction between computers and people so as to better encourage creative flow and feedback?
- CC4: How can our developing relationship with computers encourage new opportunities for experiencing both human- and computer-generated creative artefacts?
- CC5: Is there a point at which individual human creativity can no longer be enhanced by technology, no matter how sophisticated?

Could Computer Art Ever Be Properly Valued?

- CC6: When is the computer considered to have had “too much” involvement in the process of making art? To what extent is the produced artefact then devalued as a potential work of art because of the amount of automation?
- CC7: What are the implications of being clearer and bolder about just how much computing is impacting on any creative output?
- CC8: Are there ways of revealing the computational process to provide an alternative or additional aesthetic to the completed artefact or of the developing partnership between computers and artists in producing their art?
- CC9: Are current value systems that humans use to experience and evaluate art suitable for computer-created art, or are new value-systems required?
- CC10: What creative authorship should we attribute to an artist’s work that is assembled from existing code written by others (who may be anonymous)?

What Can Computing Tell Us About Creativity?

- CC11: Is autonomous creative thinking beyond the capacity of any machine that we can make now or in the future?
- CC12: Does creativity necessarily involve the creation of useful or appropriate novelty and how relevant is “value” to the definition of creativity?
- CC13: Broadly, the humanist view values what humans produce above what all other things produce. Does the ability of software to produce unusual and potentially non-human work mean that it can ever be given equal or even greater value?

- CC14: When building creative systems should we aim to mimic our own creative behaviour, the behaviours we find in nature, or design completely new mechanisms?
- CC15: If we could ever devise a formal description of everything we do as artists, then do we necessarily become limited within that description?
- CC16: The concept of creativity has changed significantly over the years. How will the increasing use of computers for creative applications change the concept of creativity further?

How Does Creativity and Computing Matter to Education?

- CC17: How then can we change the perception of computing so that programming is seen as an engaging creative subject in the same way as science, music and the arts? How can we then inspire students to develop their creativity through computing?
- CC18: How can we persuade people in education and the arts that programming can be a creative act, with its own creative practice?
- CC19: What kinds of environments provide the right level of feedback, intuition and control to inspire the idea of programming as a creative act in early learning?
- CC20: Can we find new ways of revealing and explaining computational processes where the flow of computation is more readily accessible to an audience, particularly students?
- CC21: Many companies are now beginning to recognise that they want technologists who can think like artists, yet mainstream computing education focuses mainly on engineering-based problem solving. How can we design new university computing programs to ensure that graduates have the necessary knowledge and skills that allow them to achieve their creative potential?

It is perhaps obvious but important to emphasise that many of these questions pose enormous challenges, some of which may not be overcome for many decades, even centuries, if at all. We are especially cautious about speculating on time-frames, or being naïvely optimistic about technological or scientific progress in this area. Nonetheless, we think the challenge is as exciting as the outcome, and the implications are profound if answers are to be found to the questions we now discuss.

4 Some Responses

In this section, our main concern is to set out a roadmap for research into how computers can enhance or shed light onto creativity with a bias towards encouraging and stimulating human creativity. We set out trying to respond to a handful of the questions above with a view to pump-priming a dialogue about some of these issues.

4.1 Automating Creativity

Considering questions CC1 and CC2 above, it is clear that no one likes software that makes simplistic assumptions about what we mean or are trying to do (think of the failed Microsoft Word paperclip, or early forays into automated typing correction, or the beeping we get if we don’t put our seatbelt on immediately). This raises an additional question: what are the kinds of responses and interactions we desire of computational systems so as to inspire, provoke,

and challenge us to develop meaningful creative dialogues with machines, and to have both the confidence in the system and in ourselves? Much of this becomes a personal choice, just as in the way we choose our friends and collaborators. Hence the question can be perhaps be phrased as: exactly what kind of collaborative agency do we want from our computational systems?

Relatedly, how can we remain mindful about the ways in which new technology can limit or defer creativity? We are increasingly seeing software developed which is intended to make creative decisions on our behalf. For example, modern digital cameras now take responsibility for many aspects of the creative photographic process, automatically adjusting numerous dependent properties in order to give the “best” picture, even deciding for us when to take the picture based on the human subject(s) smiling for example.

Automating creativity by engineering software and technology with a singular pre-defined definition or interpretation of what constitutes “good” results homogenises the results that it can produce. It sacrifices thought, experimentation and learning for simplistic cultural acceptability, diminishing individual creative decision-making under the mantra of “ease of use”. This kind of automated creativity is often justified as allowing non-experts the possibility of getting acceptable results without having to spend the necessary time it takes (typically around 10,000 hours) to become expert [18, 21].

There are a number of arguments against this kind of “creativity automation”. One relates to the attribution of creative agency (Section 2.1) to a mass-produced device. Each device is essentially identical in function and behaviour to any other of the same type: they are interchangeable. They do not learn or adapt to their users, nor do they offer any behavioural differences between instances. Hence any creative agency they have is fixed at the time of manufacture and is the same for all instances, meaning the only source of variation comes from their users. However the user’s role is limited by deferring aspects of the creative decision-making process to the machine, a machine that has no possibility of being novel along with a fixed and singular concept of value.

Another argument against creativity automation is the appropriation of ideas from minimalism as being appropriate for hiding complexity (e.g. [26]), a manta that has become particularly popular with technology companies. The argument goes along the lines that technology, the world, life, etc. are increasingly complex and the technology designer’s role is to hide or minimise any complexity through simple and intuitive interfaces. The best tools are simple tools, like the pencil and the hammer, and our technological tools should be like them. But as we discussed in Section 2.1, computational systems are not static tools like pencils or hammers. Simplifying a system typically involves minimising choice, hence possibilities, which as we have outlined above discourages creativity. For us the issues are those of promoting variation and experimentation in computational systems, minimising assumptions about use, allowing systems to learn and adapt to individual users, their environments and their ways of working. We don’t think of creative collaboration with people in terms of being simple interfaces to complexity. Creative relationships with computers should be viewed similarly.

So we should be concerned when creative decision-making is implicitly transferred to software at the expense of human creative exploration. Computer science educators have a responsibility to encourage understanding of this concern in our students and peers. Automation arose from a desire to replace tedious and repetitive human tasks with machines for reasons of scale, accuracy and efficiency. Computing grew from this idea and the culture that surrounds it. But creativity is not a tedious and repetitive process that can be mecha-

nised by a deterministic process of production, so we should begin by designing creative machines with a different conceptual model.

There is also something else worth noting here: that as soon as a computer can produce something “creative” it often becomes less interesting the more it produces that creative output, i.e. it loses creative agency. Researchers in artificial intelligence often claim that as soon as they have demonstrated that a particular method works (take for example, expert systems) then it simply becomes software engineering. From another perspective, there was a lot of media speculation about whether or not a computer could ever beat a human at chess. But after the initial excitement of Deep Blue beat chess champion Gary Kasparov, the problem was not so interesting. A computer could beat the best player in the world and so there was little point in designing more sophisticated programmes that could potentially pit computer algorithm against computer algorithm. The problem is that there is no sense that the computer is winning because it is being creative or subtle or ingenious in anyway but simply because it can search millions of data structures in fast and mechanistic ways. If Deep Blue’s success told us anything about the way the way the human chess mind works then further research avenues may have continued.

However, the idea that as soon as the computer can do something it is by definition not creative is an opportunity too because it can then identify that which is specifically human and creative. So that we continue to distinguish human creativity precisely by what cannot be achieved mechanically. Then as soon as something new can be achieved mechanically then perhaps that simply provides new opportunities for human creativity in collaboration with technology?

4.2 Flow and Feedback

Here we turn our attention to question CC3 and consider the increasingly important concepts of flow and feedback. The idea of “flow”, proposed by the psychologist Mihaly Csikszentmihalyi [12], is becoming extremely influential in the design of new kinds of computer systems. It is that state of mind we all can find ourselves in when completely consumed by an activity. We feel energised, engaged, egoless and enraptured – a total immersion in what we are doing. Flow states are exactly what many performers aim to achieve in every performance.

So can we re-conceptualise the methods of interaction between computers and people so as to better encourage creative flow and feedback? We have had many years of the mouse, keyboard and screen as the primary interface, but we have now entered the era of networked mobility and surface touch interfaces, where simple hand or body gestures form the locus of interaction. New ways of enhancing creative exchange are possible if we move beyond the standard mass-market paradigms and consumer technologies.

In recent work, researchers from Goldsmiths and IRCAM have been developing the concept of “fluid gesture recognition” systems. Such systems exhibit certain properties which enable a non-technical user to interact with a computer system through individual and personalised gestures [46]. The four properties that were described to build effective gestural interface systems (systems that genuinely go beyond what is possible with the traditional mouse and keyboard devices),’ are defined as follows:

1. *Continuous Control*, meaning that users can continuously and synchronously control the target software application moment to moment through their gestures.

2. *Tailorable for specific context*, which essentially means that users can define their own personal gestures suitable to whatever their local environment happens to be.
3. *Allow expert and non-expert use* so that end users rather than system designers can define their own gestures.
4. *Meaningful feedback* should be provided so as to precisely explain how the gesture recognition system is interpreting the gestures. Moreover, this feedback needs to be “moment to moment” so that the user has a real time sense of how their gestures are being interpreted.

When these four qualities are present they support the opportunity for a state of flow to exist for users interacting with this software system. Indeed all of these qualities are present in another system developed at Sony Computer Science Laboratories, called “Reflexive Loopers”. Solo performers are able to build interactive performances with musical copies of themselves [35]. (see <http://www.youtube.com/watch?v=yWb1myVSu2I> and <http://www.youtube.com/watch?v=Xp8tixrPM1U>). The system responds to the current performance by selecting the best fit from earlier in the same performance to accompany the soloist. The user is then able to build interactive musical performances with themselves and so perform as a jazz group would be expected to perform. What is interesting in this system is the concept of creative agency and where the agency is attributed. In one sense you are playing with historical copies of yourself so the agency can be attributed to you as a musician, but as the system is choosing which fragments of your previous musical self, it also has a degree of agency within the creative process.

In a final project which is also a collaboration between Pachet and d’Inverno amongst others in the FP7 Project called “PRAISE” that explicitly facilitates feedback in communities of learning (see e.g. <https://ai.vub.ac.be/cortona-2013/praise-workshop>). The project is specifically concerned with building computational systems that enable students to develop their own creative practice in music together in communities, and centres on *feedback* as the key mechanism. Whereas traditionally creativity is a rather nebulous and difficult concept to pin down (having its routes in divinity rather than human endeavour), feedback is a much clearer concept to focus on and as the project develops may give us a better understanding of how we *learn our creative practice*.

One of the ways in which the system operates is the ability for musicians to upload their performance and then to comment on it as can be seen in Figure 2.

This is important because when a musician is practicing and performing they are often in a state of flow and not in a position to reflect on the performance and articulate clear feedback to themselves. By uploading it to the PRAISE system, the musician can then more clearly and precisely comment on their own performance, listen perhaps in the same way that others would hear them, and so provide the right critical distance from their own creative practice. Not only in time (because the performance happened earlier), in space (the user is in a different environment than attached to a musical instrument), and in role (as critic and not creator).

4.3 What are the Limits of Creativity?

In response to Question CC5, is there a point at which individual human creativity can no longer be enhanced by technology or society, no matter how sophisticated? Cultural evolution or drift happens at much faster rates than genetic change. We “stand on the shoulders



Figure 2. The Music Circle System at Goldsmiths

of giants” to incrementally build on past knowledge, discoveries and achievements. However, we have physical, cognitive and practical limitations as social agents acting in a technological environment. Most technology is developed with the goal of enhancing or supporting human thinking – overcoming our limitations.

The last few centuries have seen accelerating progress in technologies, leading to exponential increases in the human population. But while its relatively easy to observe technological progress, as a species we still have effectively the same biology as our distant ancestors from the African savannah where our species originated. If little has changed in our basic genetic makeup, any (supposed) changes in our creative ability must come from elsewhere.

A number of recent computational systems have demonstrated a “counterintuitive” design logic that exceeds human designs significantly in terms of performance (see e.g. [23]). These designs were possible for a computer to find, but seemingly impossible for human designers to discover. Will the goal of augmenting or enhancing *human* creativity always be limited by our cognitive capacity and inherent genetically and socially conferred biases?

It is also interesting to ask whether computers face different limitations, or can they exceed areas of human creativity independently as they have begun to do in limited areas of human endeavour (such as the perviously mentioned chess playing example)? While beginning from a low starting point, the evolution of computer technology has shown roughly constant acceleration, and happens much faster than biological, or arguably, social evolution. One (especially speculative) point of view suggests that once computers exceed human capabilities they will have little need for us or our creative limitations.

Another more modest one suggests that once a computer can do something that we no longer think about it as a creatively challenging activity and simply move our attention as human creators elsewhere.

4.4 Could Computer Art Ever Be Properly Valued?

Considering Question CC9, does it make sense to ask if the same value system that humans use to experience art can be applied to art made by a computer? The philosopher Anthony O’Hear has argued that no matter how sophisticated or independent, machines cannot

originate art because art “in the full sense is based in human experience” and requires a communication between artist and audience drawn from that shared experience [33].

Computer works that mimic this communication are only parasitically meaningful as they derive their meaning from an analysis of existing art-objects, not directly from human experience.

This argument against any value for computer-originated art relies on computers having no basis in human experience. However we can see no reason to dismiss outright the possibility of a machine and a human sharing experiences that result in something meaningful and worth communicating. Would we be so dismissive of a hypothetical encounter with an alien intelligence?

A broader point is that value is very much a relative proposition. Many creative artworks, compositions or discoveries were not recognised at the time of their initial creation. Moreover, social confirmation plays an important role in value and acceptability, which shift with cultural evolution.

While we may see some value in animal art (such as that made by primates or elephants, for example), this is largely a curiosity – one that we must necessarily interpret through the lens of our anthropomorphism, seeing communication perhaps were no such communication really exists. Similar arguments apply to computer art.

So if our existing values do not apply to computer-originated art, then is there another value system that we can use to interact more richly and less dismissively with computer generated artefacts? Again we refer back to our objection to O’Hear’s argument – if we cannot dismiss outright a computer being able to communicate something meaningful to people (and potentially for people to reciprocate), then the possibility of a new or extended system of value remains a possibility.

The natural follow-up question (Question CC10) is what creative authorship can we attribute to a work that is assembled from existing code that has been written by others (who may be anonymous)? There is clearly creativity in a remix or mash up (where different musical fragments are bought together for a new project), even though we know the person doing the remix was not the original composer of each musical phrase or fragment. With software things are different because the code is generally hidden and is not so distinctively familiar as it is with music or collage, for example. This creates a new and challenging perspective about the ambiguity of authorship in art that is partially or completely produced by software. It points to a need for a clearer understanding of what we mean by the concept of *creative agency* and to whom or what we attribute it.

4.5 What can Computing tell us about Creativity?

Question CC11 asks if machines can ever be capable of autonomous creative thinking, an issue we have already touched upon in parts of this paper. It is also a question that has, more-or-less, been posed in many different guises well before the advent of digital computers (see e.g. [32, 44, 1, 36]) and is well discussed in the literature. Similarly, arguments about machine autonomy are well explored in both computer science and philosophy (see e.g. [4]).

Recognising software or a robot as autonomously creative requires some formal definitions of autonomy and creativity. While definitions already exist, e.g. [43], such definitions are bound by what is computable, in addition to their axioms, mappings and homomorphisms. Certainly it is possible to construct a definition of autonomous creativity and then develop a system that satisfies it within some limited context. Yet the practical value of such a system may be limited.

Under the basic definition of creativity used in Section 2.2, a creative software system must be able capable of generating valuable and surprising novelty in some domain. Certainly, some systems have already demonstrated valuable and surprising novelty in specific contexts. Evolutionary methods, for example have been able to produce designs that humans were incapable of designing [23]. Such designs appear strange and counterintuitive to conventional human design logic, yet perform far better than conventional designs.

Collaborative creative practices with machines also tell us that what creative systems are capable of producing depends heavily on the interactions and capabilities of *all* the participants, both human and machine. Moreover, the nature of human creativity and what it can produce is changed by such interactions. The “Live Algorithms” framework [3] and “Creative Ecosystems” approach [8, 28] for example illustrate how thinking of the computer as an improvisational collaborative partner or as a tightly coupled component in a feedback network can produce worthwhile creative outcomes.

Finally, systems such as that proposed in [41] consider both human and artificial agents as both creative producers and critics within a “hybrid society” with the aim of improving both through evolution. Each of these “systems” approaches see creativity as an emergent property of a network of interacting components (agents, producers, critics, software, instruments, etc.), rather than the exclusive result of a single human creator.

4.6 How Does Creativity and Computing Matter to Education?

Clearly it does and there is now a growing movement to teach computing as a creative subject in schools and universities (one pioneered at Goldsmiths for the last decade).

Despite this, it is still the case that computing is not seen as a creative subject by the general public or even at schools and universities in many countries around the world. How then can we change the perception of computing, especially in early learning, so that programming is seen as an engaging creative subject in the same way as science, music and the arts? How can we then inspire students to develop their creativity through computing? This project is being undertaken at Goldsmiths as part of the PRAISE research understanding how feedback can be used in communities of learners to recognise, articulate and feed-back to each other when creativity in programming has taken place. Moreover, Goldsmiths will in 2014 introduce a new Digital Arts Programme that is shared between the Art Department and the Computing Department clearly signifying how the practice of computing is being recognised as a creative discipline in its own right. Indeed this is why we need to build education programmes that bring technology, cultural theory and the social sciences together so that we can have a deeper and more widespread understanding of how technology is impacting us. It’s also interesting to note patterns of recruitment in large companies such as Google, Apple and Microsoft who are increasingly looking to graduates who have an arts, humanities or social science background rather than exclusively a computer science one.

5 Concluding Comments

The words “creative” and “creativity” are all around us and we must both be mindful and clear about their (over)use as well as receptive to the research and learning opportunities that come as a result of investigating creativity and particularly through the computational lens.

Computing presents an opportunity to understand what processes might be involved in creativity generally, but even more importantly it enables us as to build new kinds of systems that can support individual and collective human creative thinking and practice.

All of us enjoy the belief that we have done something creative, or that our work life is creative, or that we are working or collaborating or socialising with creative people. The challenge to computer science therefore is to build systems that give us all more opportunities to believe that we are being creative, either individually or with other creative agencies.

Acknowledgements

Supported by the FP7 project in the Technology-Enhanced Learning Program called Practice and Performance Analysis Inspiring Social Education (PRAISE). The first author was supported by an Australian Research Council Discovery Projects grant DP1094064. The second author would like to acknowledge Sony Research Labs in Paris who supported an extended six month sabbatical where he was involved in the development of some of the systems described in this paper. The authors would also like to thank Arthur Still, François Pachet, Robert Zimmer and Matthew Yee-King for their illuminating conversations over the years on some of these issues.

REFERENCES

- [1] William Ross Ashby, 'Can a mechanical chess-player outplay its designer?', *Br J Philos.*, **3**(9), 44–57, (May 1952).
- [2] Paul Atkinson, *Computer*, Objekt Series, Reaktion Books, 2010.
- [3] Tim Blackwell, Oliver Bown, and Michael Young, 'Live algorithms: Towards autonomous computer improvisers', In McCormack and d'Inverno [30], chapter 6, 147–174.
- [4] Margaret A. Boden, *Artificial intelligence*, number xviii, 376 in Handbook of perception and cognition, 2nd ed., Academic Press, San Diego, 1996.
- [5] Margaret A. Boden, *Creativity and Art: Three Roads to Surprise*, Oxford University Press, 2010.
- [6] Margaret A. Boden, 'Foreword to computers and creativity', In McCormack and d'Inverno [30], v–vi.
- [7] Oliver Bown and Jon McCormack, 'Creative agency: A clearer goal for artificial life in the arts', in *ECAL (2)*, eds., George Kampis, István Karsai, and Eörs Szathmáry, volume 5778 of *Lecture Notes in Computer Science*, pp. 254–261. Springer, (2009).
- [8] Oliver Bown and Jon McCormack, 'Taming nature: tapping the creative potential of ecosystem models in the arts', *Digital Creativity*, **21**(4), 215–231, (2010).
- [9] Andy Clark, *Natural-born cyborgs: minds, technologies, and the future of human intelligence*, Oxford University Press, New York, 2003.
- [10] Harold Cohen, 'The further exploits of Aaron, painter', *Stanford Humanities Review*, **4**, 141–158, (Jul 1995).
- [11] Simon Colton, 'The painting fool: Stories from building an automated painter', In McCormack and d'Inverno [30], chapter 1, 3–38.
- [12] M. Csikszentmihalyi, *Flow*, HarperCollins, 2009.
- [13] Palle Dahlstedt, 'Between material and ideas: A process-based spatial model of artistic creativity', In McCormack and d'Inverno [30], chapter 8, 205–233.
- [14] Mark d'Inverno and Michael Luck, *Understanding agent systems*, Springer, 2004.
- [15] Mark d'Inverno and Michael Luck, 'Creativity through autonomy and interaction', *Cognitive Computation*, **4**(3), 332–346, (2012).
- [16] Alan Dorin and Kevin B. Korb, 'Creativity refined: Bypassing the gatekeepers of appropriateness and value', In McCormack and d'Inverno [30], chapter 13, 339–360.
- [17] Hubert L. Dreyfus, Stuart E. Dreyfus, and Tom Athanasiou, *Mind over machine: the power of human intuition and expertise in the era of the computer*, Free Press, New York, 1986.
- [18] K.A. Ericsson, R. Krampe, and C. Tesch-Römer, 'The role of deliberate practice in the acquisition of expert performance', *Psychological Review*, **100**, 363–406, (1993).
- [19] J.W. Garrison, *The New Scholarship on Dewey*, number nos. 3-4; nos. 1994-1995, Springer Netherlands, 1995.
- [20] Martin Gayford, *A Bigger Message: Conversations with David Hockney*, Thames & Hudson, 2011.
- [21] M. Gladwell, *Outliers, the Story of Success*, Allen Lane, 2008.
- [22] Daniel Jones, Andrew R Brown, and Mark d'Inverno, 'The extended composer', in *Computers and Creativity*, 175–203, Springer Berlin Heidelberg, (2012).
- [23] A. J. Keane and S. M. Brown, 'The design of a satellite boom with enhanced vibration performance using genetic algorithm techniques.', in *Conference on Adaptive Computing in Engineering Design and Control 96*, ed., I. C. Parmee, pp. 107–113. P.E.D.C., (1996).
- [24] Michael Luck and Mark d'Inverno, 'Structuring a z specification to provide a formal framework for autonomous agent systems', *ZUM'95: The Z Formal Specification Notation*, 46–62, (1995).
- [25] Peter Lunenfeld, *Snap to grid: a user's guide to digital arts, media, and cultures*, number xxv, 226, Mit, Cambridge, MA, 2000.
- [26] John Maeda, *The laws of simplicity*, The MIT Press, Cambridge, MA, 2006.
- [27] P. McCorduck, *Aaron's Code: Meta-art, Artificial Intelligence and the Work of Harold Cohen*, W.H. Freeman, New York, 1990.
- [28] Jon McCormack, 'Creative ecosystems', In McCormack and d'Inverno [30], chapter 2, 39–60.
- [29] *Computers and Creativity*, eds., Jon McCormack and Mark d'Inverno, Springer, 2012.
- [30] *Computers and Creativity*, eds., Jon McCormack and Mark d'Inverno, Springer, Berlin; Heidelberg, 2012.
- [31] Jon McCormack and Mark d'Inverno, 'Computers and creativity: The road ahead', In *Computers and Creativity* [30], chapter 16, 421–424.
- [32] L. F. Menabrea, *Sketch of the analytical engine invented by Charles Babbage*, number 82, Bibliothèque Universelle de Genève, Oct 1842.
- [33] Anthony O'Hear, 'Art and technology: An old tension', *Royal Institute of Philosophy Supplement*, **38**, 143–158, (1995).
- [34] François Pachet, 'Musical virtuosity and creativity', In McCormack and d'Inverno [30], chapter 5, 115–146.
- [35] 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*, pp. 2205–2208. ACM, (2013).
- [36] G. H. R. Parkinson, 'The cybernetic approach to aesthetics', *Philosophy*, **36**, 49–61, (1961).
- [37] Roger Penrose, *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics*, Oxford University Press, Oxford, England, 1989.
- [38] David N. Perkins, 'Creativity: Beyond the Darwinian paradigm', in *Dimensions of Creativity*, ed., Margaret Boden, chapter 5, 119–142, MIT Press, (1996).
- [39] Daniel H Pink, *A whole new mind: Why right-brainers will rule the future*, Penguin, 2006.
- [40] Casey Reas, Chandler McWilliams, and LUST, *Form+Code in Design, Art, and Architecture*, Princeton Architectural Press, Inc., 2010.
- [41] Juan Romero, Penousal Machado, Antonino Santos, and Amílcar Cardoso, 'On the development of critics in evolutionary computation artists', in *Applications of Evolutionary Computing, EvoWorkshop 2003: EvoBIO, EvoCOP, EvoIASP, EvoMUSART, EvoROB, and EvoSTIM*, Essex, UK, April 14-16, 2003, *Proceedings*, eds., Günther R. Raidl, Jean-Arcady Meyer, Martin Middendorf, Stefano Cagnoni, Juan J. Romero Cardalda, David Corne, Jens Gottlieb, Agnès Guillot, Emma Hart, Colin G. Johnson, and Elena Marchiori, volume 2611 of *Lecture Notes in Computer Science*, pp. 559–569. Springer, (2003).
- [42] Douglas Rushkoff, *Program or Be Programmed: Ten Commandments for the Digital Age*, OR Books, New York, 2010.
- [43] Jürgen Schmidhuber, 'A formal theory of creativity to model the creation of art', In McCormack and d'Inverno [30], chapter 12, 323–337.
- [44] Alan M. Turing, 'Computing machinery and intelligence', *Mind*, **59**, 433–460, (1950).
- [45] Cameron Wilson, Leigh Ann Sudol, Chris Stephenson, and Mark Stehlik, 'Running on empty: The failure to teach k-12 computer science in the digital age', Report, ACM and The Computer Science Teachers Association, (2010).
- [46] Bruno Zamborlin, Frédéric Bevilacqua, Marco Gillies, and Mark d'Inverno, 'Fluid gesture interaction design: applications of continuous recognition for the design of modern gestural interfaces', *ACM Transactions on Interactive Intelligent Systems*, **3**(4), 30–45, (2014).