Chapter 5
Interaction Design for Metacreative Systems

Oliver Bown and Andrew R. Brown

Abstract In this paper, we examine digital creativity as a collective activity performed through socio-technological networks of agency. We introduce metacreation—the automation of creative tasks with machines—as a domain that is usefully examined from a 3rd wave HCI approach. We discuss four general human-computer interaction activities that commonly appear in metacreation: (1) metagenerating form; (2) searching/finding; (3) helping machines learn; and (4) evaluation/iteration. These are not necessarily specific to metacreation, but nevertheless point to particular design considerations in a metacreative context. Four creative interaction design themes are considered in their relation to metacreation: direct manipulation and real-time control; supporting playful interaction and divergent goals; the programmatic design of behaviours, and; managing distributed creativity. We then identify three paradigms of interaction design for metacreation: operation-based interaction, involving the direct manipulation of generative algorithms; request-based interaction, involving the submission of requests to a system that returns results; and ambient interaction, that involves the operation of autonomous metacreative processes in the background. Our discussion of these suggests possible trends for design: an increasingly complex and modular future for networked human-machine digital creativity; an increasing role for request-based metacreative systems where users specify, rather than construct, outcomes; the increasing role of metacreation in 'prosumer' content creation; and, consequently, the reduction of labour involved in creating media. The chapter makes clear, we hope, that metacreative practices present unique challenges and opportunities for interaction design.

O. Bown (✉)
University of New South Wales, Sydney, Australia
e-mail: o.bown@unsw.edu.au

A. R. Brown
Griffith University, Nathan, Australia

M. Filimowicz, V. Tzankova (eds.), *New Directions in Third Wave Human-Computer Interaction: Volume 1 - Technologies*, Human–Computer Interaction Series,
https://doi.org/10.1007/978-3-319-73356-2_5

5.1 Introduction

Over the past half century, the design of creative software systems that semi-automate creative tasks—referred to here as metacreative systems—has moved steadily from concept to reality. Such systems might generate melodic content for use by a film composer, automatically mix or master music recordings, create humorous tweets relevant to a particular topic, design buildings that satisfy structural and aesthetic constraints, or perform stylistic renderings of images. Although it remains a largely experimental field, in recent years real applications of metacreation[1] have necessitated the consideration of how we interact with these systems, raising issues of human computer interaction (HCI) and the adaptation of existing HCI practices to the field (Bown 2014; Kantosalo et al. 2015). Indeed it has become apparent that some of the substantial challenges in metacreation involve issues of the interaction between people and smart machines. Metacreation requires new interaction paradigms that involve semi-autonomous computational processes, and that may also be applicable to other areas of interaction design.

Metacreation[2] involves the use of computational methods to shift the balance of creative agency from people toward machines with the assistance of artificial intelligence (AI) or related methods. Metacreation has been conceptualized in various ways. The terms *algorithmic* or *procedural content creation* (e.g., Langston 1989; Togelius et al. 2011) tend to refer to the use of computers to procedurally generate patterns and variations - for example, using parametric and combinatoric methods. *Generative art* often refers to the creative practice of software-based artists in which the software has some degree of autonomy (Galanter 2003). *Computational creativity* (CC) is another catch-all term that some within the community have used to emphasise non-trivial system autonomy, which is distinguished from 'mere generation' because CC systems perform evaluation of the artefacts they produce (Ventura 2016). In other words, systems fulfilling this 'lofty' idea of CC have an 'aesthetic sense'—or, one could say—they exhibit something approximating 'taste'. Another term in use is *creative AI*, which is less strictly concerned with the creative agency of computational systems, and includes any application of AI in creative domains.

Interaction design for metacreative systems requires us to reframe HCI thinking to take into account the greater contextual awareness and decision making processes that devices can take on. 3rd wave HCI thinking is particularly well-suited to designing software that takes into account a modern understanding of the performance of co-creative tasks. These tasks include a focus on open-endedness, the collective construction of practices amongst a community of users, the creative re-appropriation

[1] From machine learning-based music composition (Ghedini et al. 2016) to distributed evolutionary computing websites (Secretan et al. 2008).

[2] The term *metacreation* itself can be attributed to Whitelaw (2004), and although his interest in using it was with respect to artificial life-based art, we interpret it as having a wider remit. Quoting Schöffer: "we are no longer creating a work, we are creating creation… we are able to bring forth… results… which go beyond the intentions of their originators" (Whitelaw 2004: 17).

of technologies, personalisation, and the configuration of complex hybrid networks of technologies (Bødker 2006, 2015).

Creativity theorists are largely agreed that creative problems are, by their nature, unmapped; we cannot follow a known path to a solution, but must conduct searches for solutions and employ heuristic methods of navigation (Wallas 1926; Perkins 1994; Simonton 1999; Amabile 1996). Exploratory creative strategies exist in creative behaviours at a multiplicity of levels. Classical thinking about individual creative cognition recognises 'incubation' as a key process, in which multiple candidate solutions are evaluated subconsciously before being presented to conscious awareness in a stage called 'illumination' (Wallas 1926). Over longer timescales, individual creative practitioners exhibit a tendency to speculatively spread their attention in order to conduct broad searches (Perkins 1994; Simonton 1999). At the super-individual level, we see undirected search operating as a social process, with social institutions performing heuristic functions: the open market, prizes, funding bodies, and the social prestige of creative achievement. A consequence of this is that creative success is unevenly distributed, although it is related to the intensity and directedness of activity and influenced by factors like education, intrinsic motivation and talent (Simonton 1999).

Because creativity occurs as a set of exploratory, cumulative, and stochastic processes, then creative agency—the attribution of influences on creative outcomes—is necessarily found to be distributed across groups of people and things, that together form complex and fluid *networks of agency* (Brown 2016). This distributed agency is illustrated in Malafouris' (2008) description of a potter shaping clay on a wheel, that involves a process of dynamic interaction where the potter's hands shape the clay but also respond to it. Achieving a creative goal requires the potter to be aware of the physical affordances of the clay and, in some sense, to 'negotiate' with it rather than merely dictating its form. Contrary to our instinctive view of humans as the sole agents in such processes, this example highlights the value of conceptualising the various actors involved as each exerting some agency, with *interactions* within the human+clay+wheel+culture hybrid being the source of creativity.

Likewise, in the authors' own practices, making electronic music involves creative input inherited from the makers of the electronic instruments, but also the instruments' own musical affordances which may transcend the intentions of their inventors, for example in the nonlinearities of electronic circuits.

These preliminary references set up the context in which we see the automation of creativity as meaningful. Making machines that are active contributors to the production of aesthetic artefacts seems at first glance—as was the case for many of the field's early experimenters—to mean making *virtual artists*. Such was the dominant nature of the discussion, surrounding the work of well-known generative practitioners such as Harold Cohen (Cohen 1995) in visual art and David Cope (Cope and Mayer 1996) in music. In contrast, we bypass the question of whether computers can be 'truly creative', can 'originate anything', or fulfil the role of an artist, and focus instead on how the types of action and interaction that occur in hybrid socio-technological networks lead to creativity.

In this chapter, we provide an overview of some of the interaction design issues in metacreation, then proceed to consider the impact for the design and evaluation of metacreative systems within the extended contextual frame of 3rd wave HCI. We base our analysis on examples from various areas of musical metacreation. We propose three metacreation design paradigms that can introduce some structure to the field. Finally, we reflect on how metacreation is a constituent part of an advanced AI future and consider its potential impact in the spheres of creative work and daily life.

5.2 Making Metacreative Systems Usable: Some Areas for Consideration

Creative practices involving semi-autonomous dynamic systems present new interaction design challenges—from traditional HCI issues surrounding perceived affordances and usability, to emergent issues of how users react to software systems that convey a sense of autonomy and authorship, or that perform unorthodox actions such as challenging or directing a co-creative[3] activity. The latter imply profound transformations to the practice of working with computers, and are exemplary of the need for 3rd wave HCI frameworks and methods.

Algorithms that support metacreation include complex hand-coded rule-based systems, machine learning algorithms, evolutionary and other types of search or optimisation algorithms, along with reasoning and concept-manipulation methods.[4] There are a range of different users, from expert programmers to end-user creative practitioners, for whom different interaction design considerations will apply. We are particularly interested in the challenges of designing metacreative tools for non-programmer end users as a way of amplifying professional practices.

5.3 Human-Computer Interaction Activities Common to Metacreation

In this section we discuss four general human-computer interaction activities that commonly appear in metacreation: (1) generating forms, (2) searching and finding, (3) helping machines learn, and (4) evaluation and iteration. These are not

[3] Co-creativity refers to a situation in which human user and computer system are interacting, active participants in a creative process.

[4] An indication of the range of algorithms used in computational creativity research can be found by looking through the free online proceedings of the International Conference on Computational Creativity (http://www.computationalcreativity.net) or of edited volumes such as McCormack and D'Inverno (2012) and Besold et al. (2015).

necessarily specific to metacreation, but nevertheless point to specific design considerations in a metacreative context.

Following this, four creative interaction design themes are considered in their relation to metacreation: (1) direct manipulation and real-time control, (2) supporting playful interaction and divergent goals, (3) the programmatic design of behaviours, and (4) managing distributed creativity.

5.3.1 Generating Forms

The simplest application of metacreation is the programmatic generation of variation, whereby a creator defines a method of *representation* and a *space* of transformational possibilities. This space can then be explored by the creator using the metacreative system, or by third party creative producers, or an audience. For example, Jon McCormack's (McCormack 2004) series of 3D virtual plant forms used a generative structure called an L-system to establish an infinite search space of possibilities which he then searched using an interactive genetic algorithm. Similarly, in architecture, it is common to create parametric forms that vary by simple adjustment of a set of predefined parameters (Hudson 2008). Variations of these forms can then be explored to optimise building performance using readily measurable factors.

As a general observation, rich metacreative systems are likely to be complex and convoluted, and hence tend to appear as opaque black boxes to users. This sets up a tension with a core principle in interaction design—that of the system's increased visibility to the user through perceived affordances and constraints, mappings and system feedback (Norman 1988). This can be resolved through designs that present relevant information to a user that helps them maintain a model of what the system is doing, as best it can. But with complex metacreative systems that cannot easily be modelled there is a tension between the power of the system and the capacity that a user has to control it. Further, many machine learning systems used in such processes do not clearly reveal, even to their expert operators, the logic of their pattern matching. For the creation of metacreative systems we prefer, instead, to understand how the system operates to empower a user to pursue a plan of action or, at least, to understand the limits of actions (Winograd and Flores 1986). Equally, we might consider how playful and engaging the system is, so as to improve engagement, even if the user cannot maintain a clear model of the system behaviour (Pask 1976; Gaver et al. 2004).

5.3.2 Searching and Finding

Perhaps the most obvious way in which metacreative tools can support creativity is through search. Consider, for example, an advanced search feature on a database of artworks. Such a feature might combine filters or use natural language processing to

enable a query like: *show me artworks produced in nineteenth century France that are colourful, and involve people and city scenes*. Interfaces that support such queries might be further be used to direct a search, not through existing content, but through spaces of generative possibilities looking for instances that meet the descriptive criteria. With the current search infrastructure in place it is potentially a small step for companies like Google or Spotify to insert generated content into their search results.

Target-based automated searches involve a user specifying a set of target properties (in evolutionary computing this is known as a fitness function, and targets can also be presented in terms of sets of constraints) and having the system find outcomes that satisfy these targets (Davis 1991). Alternatively, forms of *novelty search* can be employed to offer diversity instead of specificity (Lehman and Stanley 2008). For example, given a parametric system such as a parametrically controlled creature morphology for a game, a novelty search process would show a range of possibilities offered by a system: for example, a set of prototypical creatures each of which is different in its own way.

Because in creative search a person may not necessarily have a clear idea of what they are searching for, outcomes may be more or less precisely specified at the beginning of the process. A co-creative search process may begin with a more speculative set of suggestions led by the system, which are then reflected upon and refined by the user. The user may be led to attend to areas they were not originally interested in, i.e., to explore rather than to optimise to a target.

Many hybrids of these basic search paradigms can be constructed. An automated system might be used to generate a new set of parameters by which a user can conduct a manual search (a nested search). For example, in multiobjective search, the system is asked to satisfy multiple conflicting objectives. The result is not a single best individual, but a set of potentially optimal individuals, defined by a *Pareto front*: a surface of possible compromise solutions that the user might interrogate manually (Konak et al. 2006). Alternatively, a system might do its own filtering to generate a suitably small set of prototypes that a user can feasibly skim through (as in novelty search) (Woolley and Stanley 2014).

5.3.3 Helping Machines Learn

Machine learning systems are increasingly used for content generation. This can involve learning specific relations (supervised learning), reinforcement learning (in which only good/bad feedback is given), or data-driven (unsupervised) learning. The most common machine learning applications are in perception or data mining, but certain types of machine learning systems are inherently capable of generation. Sequence learning systems such as Markov models and LSTM networks, for example, learn to predict the next event in a sequence given previous events. This means that when seeded with a sequence they can iteratively generate new content. It is not obvious that such systems would do anything besides regurgitate existing content

with little creative value, but it is interesting to begin to see the extent to which such systems' blind generation can be the basis for producing creatively insightful and novel outputs. Some kinds of image-based neural networks can adapt or generate images, utilising techniques such as style transfer—applying the low level features of one set of images to the high-level features of an input, e.g. painterly rendering (Bruckner and Gröller 2007). Machine learning systems that are incapable of generation can still be used as evaluation nodes in hybrid systems—for example, playing the role of a fitness function in an evolutionary process.

The selection of training data is key to training a machine learning system. A music system might need an interface to specify what style to emulate: feed it Bach and it plays Bach-like music, feed it Gershwin and it plays Gershwin-like music. However, for good on-the-fly learning with current state-of-the-art systems, it has proven necessary to configure the system itself and possibly select the correct way to represent the data (Martin et al. 2011). This requires knowledge of machine learning principles and/or expert knowledge of the domain, with novel interface design challenges to overcome if we want to put such tools in the hands of non-expert end users.

5.3.4 Evaluation and Iteration

Given an algorithmic system that generates candidate outputs, we would generally expect a user to engage in a feedback process, typically over several iterations. This is analogous to a client-producer relationship, such as when a company director interacts with a graphic designer regarding a logo design. Alternatively, an audience might give feedback about the output of a system over longer time scales, as in the DarwinTunes system (MacCallum et al. 2009, 2012), where an audience listens to a live radio stream of evolved music and votes on preferred segments of the output, feeding back into an evolutionary algorithm that aims to improve the music.

An important aspect of aesthetic evaluation is that individual judgement is far from fixed, and can be influenced by any number of factors external to the thing being evaluated. For example, Salganik and Watts (2008) show how music preferences are influenced by knowledge about what *other* people think of the same music; a winner-takes-all effect where music that is perceived to be popular is subsequently more likely to be rated highly (this effect could potentially be reversed amongst those with contrarian inclinations, where 'popular' co-equates with 'lowest common denominator'). Meanwhile Bloom (2010) presents examples in which the value associated with something is influenced by beliefs about the artefact such as its provenance (for example that it was owned by someone famous and therefore once in physical contact with them). Indeed, furthermore, evaluation should be understood as multi-dimensional, a rich engagement with objects and people involving disparate factors, rather than as a fixed scalar 'score' given to aesthetic artefacts.

Taking a longer-term perspective, we note that the lifelong iteration of creative techniques can be seen as part of a more substantial reflexive development of a creative practice for an individual. Currently we adapt the computational systems around us by configuring them, and we adapt our working practices to their affordances. Software systems do not automatically coevolve with us in the development of our creative practice, but they could. This necessitates looking at long term patterns of software use, and processes of mutual adaptation and habituation between people and the systems they use, an area in which HCI has started to establish good strategies as it develops into the 3rd wave.

5.3.5 Creative Interaction Design Themes

The forms of interaction discussed above occur in a range of metacreation contexts and are core foci of HCI design for metacreative systems. However, a number of other themes cross the boundary between existing digital creative practice and metacreative practice.

5.3.5.1 Direct Manipulation and Real-Time Control

Directly controlling a metacreative system is similar to the control of any other device. While advanced computing techniques such as machine learning or evolution may be used to generate artefacts—or to generate systems that generate artefacts—thought is required when interfacing them with direct manipulation by a user; such as a creative practitioner employing generative techniques. It can be productive, therefore, if generative processes can apply to standard objects, such as images or musical scores, and that multiple generative and direct manipulation processes are interoperable. In other scenarios, live *configuration* of some more or less complex network of interactions becomes a core part of the activity. For example, the ecosystemic compositions of Agostino di Scipio are predicated on using a human-managed feedback loop between speakers, an acoustic space, a microphone, and digital signal processing to simulate the emergence of complex patterns from the background room acoustics (di Scipio 2003).

5.3.5.2 Supporting Playful Interaction and Divergent Goals

Analyses of creative tasks indicate that goal-oriented behaviour—the primary focus of much software interaction design—is less applicable in creative domains where many tasks are open-ended, exploratory or playful in nature (Perkins 1994). Correspondingly, interface guidelines emphasise predictability and reliability (Nielsen 1993). This presents a potential misalignment between the static or dynamic nature of a system and its user interface. Complicating interaction design further is the potential for users to vary their objectives, especially between aims

that are utilitarian and speculative. For example, every aspect of a professional creative program such as Adobe Photoshop is designed to be stable and predicated on the idea of being 'powerful', offering the user a large number of goal-directed actions. Yet the creative practitioner might use it in very playful ways, engaging in odd experiments, appropriating quirks and errors to creative effect, and imposing peculiar constraints (Ferguson and Brown 2016). The intended interface goals of software are typically to provide leverage and ease of use. By comparison, computer games offer an alternative paradigm of software that does not simply aspire to make a task easy—quite the opposite—and the unexpected is often part of the pleasure of game play. This diversity of objectives exemplifies the reasons 'user experience' exists as a separate category from 'usability', and why the design of interfaces for metacreative systems often needs to privilege exploratory power so such systems can be effective in stimulating idea generation and discovery.

Even in the most professional context, creative production involves an element of playful search. Professional creative software tools can support this search by being predictable and powerful. But generative creative software might support such search by being proactive, challenging and surprising.

5.3.5.3 The Programmatic Design of Behaviours

Typically, metacreative systems are complex, bespoke pieces of software, and thus their behavioural algorithms are programmed as software algorithms. Given this, and despite the potential of graphical user interfaces (GUIs) for controlling specific metacreative systems, it may be the case that the expert manipulation of metacreative systems requires programming. In such systems the user is effectively operating in an open-ended context. A direct-manipulation widget-based GUI does not easily allow such open-ended configurations.

Because the current development of metacreative systems often involves software development in the traditional sense, programming is an important form of interaction to consider. The work on end-user programming, such as that of Blackwell (2002), serves this discussion. He identifies practitioners, besides professional software developers, who have to make regular choices between programming as a creative act, and what as we have seen, is termed 'direct manipulation' (from Shneiderman 1993). Programming is defined by a set of core cognitive tasks: requirements gathering, specification, design, coding and debugging (Blackwell 2002). With generative systems, regardless of the interface, the designer/user may face such tasks—particularly those resembling debugging or structural change—as a natural part of the exploratory nature of creative practices. The use of code as an interface enables this flexibility, but is not practicable as an interface for all, perhaps not even most, metacreative interaction contexts.

Meanwhile, software languages are becoming increasingly expressive and fluid. Some practitioners have taken programming as the key mode of interaction with metacreative systems in a performative framework. This is particularly prevalent in the practice of live coding which "involves writing and modifying computer pro-

grams that generate music in real time" (Brown and Sorensen 2009: 17). This performative approach to software development involves the on-the-fly programming of system behaviors and the management of them through code modification in real time. A core theme here is that of switching to an improvisational form of dynamic program creation using code as the interface.

5.3.5.4 Managing Distributed Creativity

The design of metacreative systems necessitates a distributed view of creative agency involving networks of people and things (Suchman 1987; Latour 1995; Rammert 2008).

In such networks, intelligence is extended beyond its conventionally assumed place in the human mind to include objects, machines and environments. In this view, "the human organism is linked with an external entity in a two-way interaction, creating a coupled system that can be seen as a cognitive system in its own right" (Clark and Chalmers 1998: 12) and the interface between agents is critical to the effectiveness of interactions. Correspondingly from a HCI perspective, Bødker argues that frameworks that properly represent the multiplicity of interactions of real working environments, described as webs-of-technology, are increasingly needed: "we never design single, monolithic devices or systems but technology that must be seen and used in relation to many other devices, applications and systems." (Bødker 2006: 3).

Echoing Latour's concept of the *script* (Latour 2011) where individuals can be either independently (above the script) or passively (under the script) following the directions of their circumstances, Bødker also highlights how individual creative practitioners are themselves shaped in networked socio-technological interactions: "users' shared capacities and experiences are not just based on individual acting and learning in the world; rather, they are bound to shared practices, joint activities, and so on, in artifact ecologies. It is against this background that the relationship between the user and the artifact exists." (Bødker 2015: 27).

Such considerations, in our view, reinforce the sense that creativity occurs as a socio-technological process, understood by identifying and revealing the nature of various temporary people-object aggregates, and in particular the various relations between the functions of different subsystems. In our practice-based work in musical metacreative systems, we have both independently assumed a framework in which creativity occurs through these agency networks (Bown 2012, 2016).

A consequence of the evolution of digital systems is that the ecosystems of creativity are becoming more complex. A creative practitioner working with software may access multiple software tools, plugins, and assets created by others and distributed via the internet, involving increasingly digitally mediated collaborative practices such as remixes and mashups (Lessig 2008). Those assets might be made from other assets that have in turn been put together by the same process. The dynamic topologies of digital systems transform the network properties of creativity. As individual humans organise into fluid networks of distributed cognition, so

computational systems exhibit porous boundaries: machine-to-machine interactions such as web-services create new computer-agency hybrids with new affordances. Bringing these together into human-object agency networks amplifies the challenges of coordination through interaction design.

A final comment on this theme concerns attribution. In artistic domains, distributed creativity relates to issues of authorship, an ever-shifting and contended topic throughout the twentieth century and into the twenty-first. Sometimes we care very much who the author is; we care if a Renaissance painting is a forgery, even if it *looks* like the real thing. At other times authorship is less important – advertising jingles, for example. Metacreative systems may encounter different degrees of leeway in terms of how they can be used in place of human creative work, depending on the importance of authorship in any given context. For example, computer generated (authorless) townhouses and elevator music might be more permissible than computer generated war monuments or popular music hit singles.

5.4 Metacreative Design Paradigms

The features of metacreative systems described above are far from exhaustive, but they offer an insight into the issues surrounding metacreative interaction, pointing to design strategies for metacreation.

Based on the above, we can split interaction scenarios into three top-level categories, which cut across the distinction used by Blackwell (2002) and others between direct manipulation and programming. These top-level categories are:

1. Operation-based. The user operates the system (controlling parameters, interactively evolving, training, etc.). The system may have degrees of autonomy and computational intelligence but is presented as a production tool; the user uses the system to produce work.
   Operation-based interfaces are what we commonly deal with in most HCI. The idea behind operation-based paradigm for metacreation is that we operate complex metacreative algorithms in much the same way we operate any other digital tool. A neural network, for example, could have an interface allowing us to set its parameters, train it, and operate it generatively.
2. Request-based. The user puts requests to the system, such as "generate this" or "evaluate this". The system is a (likely remote) service that would resemble a search interface. The system would have a clear role that resembles the activities of a person, such as producing something or evaluating something, and would give a clear sense of autonomy, even if its output is predictable and directly manipulable. The interface can take the form of a text box, programming API or GUI and has a relative degree of open-endedness.
   In creative domains, we can imagine smart systems that respond to such requests: "*make 100 deep house tracks with a walking jazz bassline and cut up excerpts from Martin Luther King's speeches*". A search engine-type interface parses the request, converts it into a machine specification, and then seeks generative services which perform various subtasks. Upon receiving the results, further correc-

tive instructions may be issued, where some dialogue may be used to better identify the criteria.

3. Ambient. The system operates ambiently and proactively. The user does not directly manipulate it nor submits requests. In one scenario, the system could be an evolving creative tool that is constantly updating based on improvements, possibly adapted to an individual user's needs. In another scenario, the system would take the form of a human assistant and make suggestions via language. In a third scenario, the system would perform forms of background analysis which is presented to the user as ambient information (e.g., the user begins to draw an image and the assistant tiles the background with related or complementary images, or indicates suggested completions of the image).
   Content-aware fill functions in software such as Adobe Photoshop provide a glimpse of what is possible with ambient metacreation. With these tools, users can delete sections of an image and the software will fill them in by extrapolating from surrounding materials. For example, when removing an object from a photo the content-aware fill will fill in the background effectively, dealing with complex structures such as trees and clouds. In conjunction with learning and predictive algorithms, such ambient systems could have the capacity to not only complete actions for the user, but anticipate actions and take some appropriate complementary steps.

This analysis applies at a similar level to a categorisation offered by Lubart (2005), which is somewhat richer in metaphor—computer as nanny, as pen-pal, as coach and as colleague—but our list places the boundaries between categories differently. Current creative practice focuses largely on operation-based interaction, although we are familiar with request-based and ambient interaction scenarios in non-generative contexts. We also note that these paradigms should not be seen as absolute or clearly demarcated, and any allocation of a given system into any of these categories requires interpretation. For example, a user interacting with a request-based interface can still be seen as operating that system.

Existing design concerns related to operation-based interaction include the successful creation of modular systems that enable cumulative development and easy debugging, powerful visualisation of system dynamics, tools that support rapid exploration, systems that can be edited in multiple ways and interacted with from multiple viewpoints. The operation-based paradigm is more associated with older waves of HCI but has great potential to be rethought through the 3rd wave HCI focus on artefact ecologies and distributed creativity.

Existing design concerns related to request-based interaction include producing results that exhibit consistency, relevance, clarity, meaningfulness, and a form of manipulability. Generative systems should produce meaningful diversity. They should successfully handle ambiguity: rather than throw an error or return nothing, they should return something, along with ways to cater for misunderstanding. In the Google search engine this includes alternative suggestions and autocomplete to help narrow in on the 'correct' search term. Request-based systems should ideally produce outputs that can be manipulated in multiple ways, which can be supported by technology that enables manipulation of arbitrary material. For example, source-

separation algorithms in music promise to allow the manipulation of elements that have previously been mixed together. Above all, generative systems should produce 'good' results, the central challenge of metacreation.

Request-based systems offer great potential to scale, and can be rapidly and fluidly reconfigured (or self-reconfiguring). They can naturally occur in a distributed manner which truly breaks from a one-user-one-system perspective, to one of an ecosystem of many-to-many interactions between multiple machines and users. We see smart knowledge and computation engines/interfaces such as Wolfram Alpha[5] as indicative of what form such systems might take.

Existing design principles related to ambient interaction include not disrupting workflows (good timing, not creating distractions, relevance), and being sensitive to context, such as the current activity the user is engaging in. This in part involves a recognition of different phases of creative production, some of which are more open-ended and exploratory, whilst others are more objective and goal-driven. In the former case (where the overall style of a work might not yet have been determined), the generation of diverse ideas, or of thought-provoking reconfigurations of concepts may be appropriate, whereas in the latter case (where the style is locked down and the work needs to be completed), focused, goal directed support in the style of text autocomplete might be more valuable.

The request-based and ambient paradigms prompt consideration of how metacreative technologies could be interacted with effectively in these contexts. Such analysis must take into account the evolution of creative practice in distributed networks of users, including remixes and mashups, the interaction between multiple individuals and software components in the creation of works, tool configuration activities such as appropriation and personalisation, and open-ended, non-goal oriented exploration. Research must ultimately take place in the wild—looking at users in context, seeing how they use metacreative software to achieve real creative outcomes—and we point to our practice-based work in the context of digital music practice as a productive model. However, as Bødker also notes (2015), it is not always appropriate to 'dump' software on people and study them. Bridging studies and approaches to bootstrapping metacreative technologies needs to be considered.

5.5 Examples and Speculative Futures

In this section we discuss current examples of each of the three paradigms, and briefly speculate on how each paradigm might develop into the future. We draw our examples predominantly from our own area of expertise: the world of digital music.

[5] https://www.wolframalpha.com/

5.5.1 Operation-Based

5.5.1.1 Examples

Game-like mobile music apps such as those made by Brian Eno and Peter Chilvers (2008–2016) provide an example of the simplest kinds of operation-based metacreative interfaces, emerging in the context of smartphones and tablets. In these works, a touch interface is used to trigger sounds or manipulate icons that represent musical fragments. The apps include underlying generative processes, typically of repetition or gradual variation, on top of which the user can collage a range of preset sound and musical fragments. Thus, the evolving soundscape results from automation and manipulation. Digital music artists now commonly release such apps as extensions of their music production. These are clear examples of the metageneration of form in an operation-based paradigm. More professionally oriented tools for metagenerating musical form using rules include Noatikl[6] and Patter.[7] In these cases the user is generally not assumed to have a technical knowledge of the algorithms involved, which are packaged to have a user-friendly interface which reveals the system behaviour.

Operation-based control of search systems include interactive genetic algorithms such as Picbreeder (Secretan et al. 2008), where the user can interactively breed new image forms. With these and with learning tools such as Jnana,[8] which learns musical style in an interactive context, interfaces are usually very simple and the direct control of parameters (such as the genetic structure that is evolved, or the statistical model that is learnt) is obscured (see further discussion in Bray and Bown 2017).

5.5.1.2 Futures

We see the complexification of digitally-mediated creative practice in phenomena such as remix and mashup culture. Copying and reinterpreting is a critical component in all creative practice, but in remix and mashup practices, it is not just the ideas that are copied, but the artefacts themselves: for example a sound file, or source code fragment (Lessig 2008). We are seeing this integration of practice through digital artefacts intensifying, with more instances of people working together on the same material: a Wiki entry or a DAW project. We believe the growth of creative coding infrastructure and communities will play an increasingly important role in this area, as code continues to be one of the best ways to interact with advanced metacreative algorithms, and the sharing of code allows for unlimited combinatoric creativity in hybrid, distributed, human-computer agency networks.

[6] https://intermorphic.com/noatikl/

[7] http://playpatter.com/

[8] https://ccrma.stanford.edu/~colinsul/projects/jnana/

5.5.2 *Request-Based*

5.5.2.1 Examples

Request-based interactions with machines are increasingly prevalent. Common examples include internet search and recommendation systems. Common examples include internet search and recommendation systems and conversational userinterfaces. This interaction design paradigm requires a medium of enquiry - typically text - familiar to the user, and a computational system with domain knowledge and an ability to service the request. Through developments in data mining and machine learning, we see these request-based systems being increasingly utilised in the creation of artistic works.

Style Machine[9] is an example of a music generation system that produces complete musical works in a given style or combination of styles from a range of electronic dance music options. Like Jnana and Patter, the system is built directly into a DAW, Ableton Live, so that users can have immediate low-level access to the generated content. Startups that are using machine learning engines to drive request-based interaction in music generation are emerging. Aiva and Jukedeck are two recent arrivals applying this type of interaction design to two different use-cases, one that that is client focused and mediated by technicians and the other that is end-user focused, with a web interface.

Both of these examples show the hazy boundary between operation-based and request-based paradigms. Both have parametric controls that affect their behaviour (operation-based) but are also presented as black boxes that respond to requests (request-based).

5.5.2.2 Futures

Given recent developments, we foresee scenarios in which people will interact with request-based outputs from generative systems for their own pleasure: for example by using a service similar to Spotify to listen to original generative music. This might mean in turn that artists output generative systems rather than fixed media works. People may still be listening to *U2*, but via a *U2* generator that might be capable of generating an infinite number of *U2*-style songs.[10] Another alternative is that the generated artefacts may by authorless without need for the kudos of renowned creators.

[9] https://www.youtube.com/watch?v=Xw2I9B4yt7I

[10] One way to continue the fruitful collaboration between *U2* and *iTunes*, started with *"Songs of Innocence"*?

5.5.3 Ambient-Based

5.5.3.1 Examples

Like DarwinTunes discussed above, Draves' Electric Sheep (1999) is a well-known ambient generative system in the visual arts domain. It operates as a computer screen saver using a set of evolving algorithms to generate complex visual animations. The system runs 'ambiently' and autonomously during times of inactive computer use. Computers with Electric Sheep installed communicate with each other via the internet to share the load of creating new morphing abstract animations.

Another form of ambient interaction occurs where the system is simply conceived of as a contributor to a collaborative work. This is found to the greatest extent in the scenario of live improvisation. A common strategy here depends on corpus-based matching algorithms. A (typically short) musical query is submitted, and the system attempts to find a match, which could mean something similar to the query, or complementary to it either sequentially (call and response) or in parallel (harmonisation or counterpoint). For example, the interactive music system Frank (Plans and Morelli 2007) captures audio and uses gestural matching techniques (k-means clustering) in combination with genetic co-evolution. The system interacts with a human player by responding to audio queries (performed audio input) by evolving solutions and outputting the resulting audio in real time.

Interactive computer improvisation agents such as Frank, and our own work with systems such as CIM (Brown et al. 2013) and Zamyatin (Bown 2011) fit an ambient interaction paradigm in the sense that the primary form of interaction is to 'jam' with them, as you would with a fellow human musician. Musicians interacting with such systems do not need to operate or explicitly submit queries to the system. They just play.

5.5.3.2 Futures

Along with user-led control, metacreation seems likely to involve a move away from the staple interaction paradigm of direct manipulation toward a more passive use of computers where the computer is actively suggesting, with users benefitting from these suggestions being ready to hand when required. Such change can be framed in the context of Harrison et al.'s (2007) survey of 3rd wave HCI trends, including shifts from generalised to situated knowledge, from information to interpretation, from clean to messy formalisms, and from scientific to unscientific strategies. To this list we would add a shift in the role of the computer from passive tool to active collaborator.

5.5.4 Professional and Social Implications of Metacreative Systems

Beyond generation, metacreative systems often employ processes of selection and judgement over generated (or input) materials, leading to situations where the computer takes over increasingly evaluative roles in a client-producer relationship (for example, suggesting why an aesthetic decision is valid). Automated evaluation can lead to increased efficiencies in traditionally labour intensive creative domains. But what other effects will such changes in technology have on creative professions and on society?

5.5.4.1 In Professions

Creative teams and workflows will inevitably be transformed. We foresee a continuation of the trend towards smaller groups involved in the development of works, which will take place more rapidly through the speed-up enabled by automation. We speculate that this may lead to greater domain generalism, with algorithms managing to capture domain-specific knowledge well, and successful operators of these algorithms consequently becoming empowered to work across domain boundaries. For example, in advertising, a single person may be involved in the production of an entire animated infomercial with characters, music, storyboard and focus-group tested messaging. Similarly, a single individual may create the sound-design for a movie in a day, mostly playing the role of a client to the software-as-producer. An architect may present a myriad of design options to a client, all tested and simulated with building analytics.

Metacreation offers new ways to engage in micro-branding, communicate via rich media, develop and modify memes on demand, and generate digital content on-the-fly for rich media environments such a AR, VR and game-based worlds. How and when to intervene in these processes and what modes of interaction to use when doing so remain design challenges best resolved with knowledge of particular situations and domains.

5.5.4.2 In Society

Metacreative systems that automate content generation will impact the contexts in which creative content production is conducted, not just in professional contexts, but in the everyday ‘prosumer’ creation of content. With such generative power, the casual creation of artefacts is set to become easier. Use cases might resemble Bødker’s observations of individual iPhone use innovation, where she points to “both the shared and explorative phases that are part of this development, and how the iPhone moved from a fancy telephone to, for example, a highly individual ‘poetry machine’, in the hands of one of the interviewees.” (Bødker 2015: 26).

The ability for people to rapidly create content without specialised training using metacreative systems will have an effect on many forms of human-computer interaction. A vast amount of computer use in casual social interaction takes the form of micro-content creation, primarily text, photos, video, audio, and types of structured data such as playlists. In the smartphone era, creation tools that enable rapid on-the-fly creation via minimal and intuitive interfaces have proven to hold special value. Inviting someone out, or wishing them happy birthday, could be augmented with a contextually relevant generative artwork whipped up on a smartphone. We can imagine how our list of interaction activities, themes, and overarching paradigms (operation-based, request-based, and ambient) might fit with the design of rapid metacreation tools for such scenarios. For example, stark emotional markers may be key in constraining a search of generated possibilities (e.g. 'give me something happy'). Meanwhile, the novelty requirements of artefacts might be minimal in these scenarios, but a personal style signature developed by system creators or users as part of their digital identity may play an important role in making them feel that they are creating something that is not just generic.

Casual and social computer use also includes various forms of evaluation such as ratings (including of artists, aesthetic artefacts, or creative tools), along with a search for valued artefacts, which provides the big data required to build massive models of cultural phenomena that can be used to feed generative processes. We tend to think of metacreative technologies as fun and on the whole positive, or at least harmless. But recent global political events suggest that the combination of big data analytics of cultural trends, with the ability to generate rich emotionally salient media optimised for cultural impact could be a potent mix with dangerous uses. Metacreative researchers should be mindful of the deeper political implications when creative practices become a site for automation.

5.6 Conclusion

Metacreation technologies form part of the drive toward 3rd wave HCI not least because they privilege the agency of computing systems in interactive contexts. The consequences that come with this shift filter through many of the issues that this era of HCI is concerned with.

As in other areas of automation, metacreation will have an impact on how creative work is done. We are witnessing an ongoing complexification of socio-technological creative systems, and the evidence from our work suggests that metacreative systems will play a part in (and benefit from) the natural continuation of this complexification. Systems such as neural networks, generative processes, and web-services using big-data generation, depend increasingly on being situated in interactive contexts. Metacreative systems are often more complex than existing software in terms of their opacity and their potential for autonomy. We suggest that these new systems should coexist alongside existing tools to enrich the ecology of participants in human-object agency networks. Metacreative tools can be seen as

adding new dimensions to existing technological infrastructures and thus to human creativity.

Although the focus of our research is on making metacreative systems that work well, as this technology evolves we are also becoming increasingly aware of wider, potentially negative social implications, as alluded to in our comments above. Beyond the erosion of creative work, which we do not think is a significant threat, the affordances such tools might have in social control, through the creation of emotionally salient and attention grabbing content is now worthy of serious consideration. In light of the rapid changes the world is facing, such as recent trends in political manipulation, a field such as metacreation could lose its claims of innocence rapidly. We hope that research will be conducted top-to-bottom, and from a range of technical and analytical perspectives, to connect algorithmic and HCI innovation with these possible wider social implications.

References

Amabile T (1996) Creativity in context. Westview Press, Boulder

Besold TR, Schorlemmer M, Smaill A (eds) (2015) Computational creativity research: towards creative machines. Atlantis Press, Berlin

Blackwell AF (2002) First steps in programming: a rationale for attention investment models. In: Proceedings of the IEEE 2002 Symposia on human centric computing languages and environments, 2002. IEEE, Piscataway, pp 2–10

Bloom P (2010) How pleasure works: the new science of why we like what we like. Random House, New York

Bødker S (2006) When second wave HCI meets third wave challenges. In: Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles. ACM, New York, pp 1–8

Bødker S (2015) Third-wave HCI, 10 years later—participation and sharing. Interactions 22(5):24–31

Bown O (2011) Experiments in modular design for the creative composition of live algorithms. Comput Music J 35(3):73–85

Bown O (2012) Generative and adaptive creativity: a unified approach to creativity in nature, humans and machines. In: Computers and creativity. Springer, Berlin/Heidelberg, pp 361–381

Bown O (2014) Empirically grounding the evaluation of creative systems: incorporating interaction design. In: Proceedings of the fifth international conference on computational creativity, pp 112–119

Bray L, Bown O (2017) How can we deal with the design principle of visibility in highly encapsulated computationally creative systems? In: Proceedings of the international conference on computational creativity, June 19–23, Atlanta

Brown AR (2016) Understanding musical practices as agency networks. In: Proceedings of the international conference on computational creativity. [Online]. 2016 Paris, Association of Computational Creativity. Available from: http://www.computationalcreativity.net/iccc2016/wp-content/uploads/2016/01/paper_21.pdf

Brown AR, Sorensen A (2009) Interacting with generative music through live coding. Contemp Music Rev 28(1):17–29

Brown A, Gifford T, Voltz B (2013) Controlling interactive music performance (CIM). In: Proceedings of the fourth international conference on computational creativity, vol 221

Bruckner S, Gröller ME (2007) Style transfer functions for illustrative volume rendering. In: Computer graphics forum (Vol. 26, No. 3, pp. 715–724). Blackwell Publishing Ltd, Cambridge

Clark A, Chalmers DJ (1998) The extended mind. Analysis 58:10–23
Cohen H (1995) The further exploits of AARON, painter. Stanf Humanit Rev 4(2):141–158
Cope D, Mayer MJ (1996) Experiments in musical intelligence, vol 12. AR Editions, Madison
Davis L (1991) Handbook of genetic algorithms. Van Nostrand Reinhold, New York
Di Scipio A (2003) Sound is the interface: from interactive to ecosystemic signal processing. Organised Sound 8(03):269–277
Draves S (1999) Electric sheep [Online]. Spotworks LLC. Available from: https://electricsheep.org/
Eno B, Chilvers P (2008–2016) Bloom, trope, air, quarta & reflection. iPhone App Store, Opal Limited
Ferguson J, Brown AR (2016) Fostering a post-digital avant-garde: research-led teaching of music technology. Organised Sound 21(2):127–137
Galanter P (2003) What is generative art? Complexity theory as a context for art theory. In: GA2003–6th generative art conference. 10–12 December, 2003, Milan, Italy
Gaver WW, Bowers J, Boucher A, Gellerson H et al (2004) The drift table: designing for ludic engagement. In: CHI'04 extended abstracts on Human factors in computing systems. ACM, New York, pp 885–900
Ghedini F, Pachet F, Roy P (2016) Creating music and texts with flow machines. In: Multidisciplinary contributions to the science of creative thinking. Springer, Singapore, pp 325–343
Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. In Alt. Chi. Session at the SIGCHI conference on human factors in computing systems. San Jose, California, USA, pp 1–18
Hudson R (2008) Frameworks for practical parametric design in architecture. In: Proceedings of the 26th eCAADe conference. Antwerp, Belgium, pp. 847–854
Kantosalo A, Toivanen JM, Toivonen HT (2015) Interaction evaluation for human-computer co-creativity. In: Proceedings of the sixth international conference on computational creativity. June 29–July 2, Park City, Utah, USA
Konak A, Coit DW, Smith AE (2006) Multi-objective optimization using genetic algorithms: a tutorial. Reliab Eng Syst Saf 91(9):992–1007
Langston P (1989) Six techniques for algorithmic music composition. In: 15th international computer music conference. ACMA, Columbus, pp 2–5
Latour B (1995) Mixing humans and nonhuman together: the sociology of a door-closer. In: Ecologies of Knowledge. New York, State University of New York Press, pp 257–277
Latour B (2011) What's the story? Organizing as a mode of existence. In: Passoth JH, Peuker B, Schillmeier M (eds) Agency without actors. Routledge, London
Lehman J, Stanley KO (2008) Exploiting open-endedness to solve problems through the search for novelty. In: Proceedings of the Eleventh International Conference on Artificial Life (ALIFE XI). MIT Press, Cambridge, MA, pp 329–336
Lessig L (2008) Remix: making art and commerce thrive in the hybrid economy. Penguin, New York
Lubart T (2005) How can computers be partners in the creative process: classification and commentary on the special issue. Int J of Hum Comput Stud 63(4):365–369
MacCallum B, Leroi A, Mauch M, Welburn S et al (2009) DarwinTunes [Online]. London, Imperial College and Queen Mary University. Available from: http://darwintunes.org/
MacCallum RM, Mauch M, Burt A, Leroi AM (2012) Evolution of music by public choice. Proc Natl Acad Sci 109(30):12081–12086
Malafouris L (2008) At the potter's wheel: an argument for material agency. In: Knappett C, Malafouris L (eds) Material agency: towards a non-anthropocentric approach. Springer, New York, pp 19–36
Martin A, Jin CT, Bown O (2011) A toolkit for designing interactive musical agents. In: Proceedings of the 23rd Australian computer-human interaction conference. November 28–December 02, 2011, Canberra, Australia, pp 194–197

McCormack J (2004) Aesthetic evolution of L-systems revisited. In: Workshops on applications of evolutionary computation. Springer, Berlin/Heidelberg, pp 477–488
McCormack J, D'Inverno M (eds) (2012) Computers and creativity. Springer, Berlin
Nielsen, Jakob. 1993. Usability Engineering. New York: Morgan Kaufmann, pp 115–164
Norman D (1988) The psychology of everyday things. Basic books, New York
Pask G (1976) Conversation theory. Applications in education and epistemology. Elsevier, Amsterdam/New York
Perkins DN (1994) Creativity: beyond the darwinian paradigm. In Boden M (ed), Dimensions of creativity, Bradford Books, Cambridge, MA. pp 119–142
Plans Casal D, Morelli D (2007) Remembering the future: an overview of co-evolution in musical improvisation. In: Proceedings of the international computer music conference. Copenhagen, International Computer Music Association, pp 200–206
Rammert W (2008) Where the action is: distributed agency between humans, machines, and programs
Salganik MJ, Watts DJ (2008) Leading the herd astray: an experimental study of self-fulfilling prophecies in an artificial cultural market. Soc Psychol Q 71(4):338–355
Secretan J, Beato N, D Ambrosio DB, Rodriguez A, Campbell A, Stanley KO (2008) Picbreeder: evolving pictures collaboratively online. In: Proceedings of the SIGCHI conference on human factors in computing systems. April 27–May 02, 2013, Paris, France, pp 1759–1768
Shneiderman B (1993) 1.1 Direct manipulation: a step beyond programming languages. Sparks of innovation in human-computer interaction, 17
Simonton DK (1999) Origins of genius: Darwinian perspectives on creativity. Oxford University Press, New York
Suchman LA (1987) Plans and situated actions: the problem of human-machine communication. Cambridge university press, New York
Togelius J, Yannakakis GN, Stanley KO, Browne C (2011) Search-based procedural content generation: a taxonomy and survey. IEEE Trans Comput Intell AI Games 3(3):172–186
Ventura D (2016) Mere generation: essential barometer or dated concept. In: Proceedings of the seventh international conference on computational creativity, ICCC. June 27–July 1, Paris, France
Wallas G (1926) The art of thought. Harcourt Brace, New York
Whitelaw M (2004) Metacreation: art and artificial life. The MIT Press, Cambridge, MA
Winograd T, Flores F (1986) Understanding computers and cognition: a new foundation for design. Addison Wesley, Reading
Woolley BG, Stanley KO (2014) A novel human-computer collaboration: combining novelty search with interactive evolution. In: Proceedings of the 2014 annual conference on genetic and evolutionary computation, 12–16 July, 2014, Vancouver, BC, Canada, pp 233–240