

Applying Core Interaction Design Principles to Computational Creativity

Liam Bray

Art and Design

The University of New South Wales
Cnr Oxford St and Greens Rd
Paddington NSW 2021 Australia
liam.bray@sydney.edu.au

Oliver Bown

Art and Design

The University of New South Wales
Cnr Oxford St and Greens Rd
Paddington NSW 2021 Australia
o.bown@unsw.edu.au

Abstract

If we understand computational creativity (CC) as ultimately leading to useful interactive systems, then interaction design (ID) is a relevant body of theory with which to develop and test systems. Yet by engaging with complex and opaque systems, CC appears to break core ID wisdom, which preferences the comprehensibility of the system to users. We discuss core ID principles and ask how we can bring together ID and CC towards a better understanding of interaction in CC, whether in ‘merely’ generative art, human-computer co-creativity or full blown automated creativity. We look at ID issues surrounding creative processes of playful and non-objective search and consider how a more developed form of ID theory could work in these contexts.

Introduction

Recent work in computational creativity (CC) has begun to look at applying interaction design (ID) principles to CC systems, with the intention of advancing the usability and experience of these systems. But to date, no detailed discussion of the application of ID principles to CC has been had.

ID theory is a rich and diverse body of knowledge which extends the ability of designers to address the behavioural and experiential, whilst potentially being inclusive of computationally complex systems Gero (1990). If CC is to truly embrace the interaction between people and CC systems, then it follows that core interaction design issues should be explored.

CC designers intend the goals of their systems to be clear, but a means by which to determine the users’ perceived success in achieving these goals is not, as the evaluation of creative success is not empirically grounded in an objective methodology. In a previous paper we (Bown, 2014) speak to this:

“Terms such as ‘creativity’ and ‘imagination’ do not describe things that we can readily measure or objectively identify, they are concepts that frame other kinds of measurable and objectively identifiable things, as part of a loose theoretical framework.”

This is echoed by Carroll (2013) “It is critical to look beyond traditional time, error, and other productivity measure-

ments that are commonly used in HCI because these measures do not capture all the relevant dimensions of creativity support” and again by Shneiderman (2007): “The complex nature of human discovery and innovation cannot be studied like pendulums or solid-state materials”.

By contrast, if we consider CC from an ID point of view we are able to engage with the challenge of evaluation of creative systems in a meaningful way. Both enabling systems to be more effectively designed for use by creative practitioners, and genuinely resolving dilemmas of empirical grounding (Bown, 2014).

In this paper we take a more detailed look at key principles from ID, and how they might apply to CC systems, in order to develop a more holistic means of evaluating and designing CC systems from a user’s perspective. We also suggest a simple framework that describes potential visibility concerns in CC systems by defining the behaviour of a system in terms its structure and its trajectory.

Our comments apply most readily to more traditional creative tool use cases, and in this sense are focused on supporting creative users (Candy and Edmonds, 1997). These same comments might not carry so well into the relationship between audiences and creative art machines, but we nevertheless pursue the possible value of this ID approach in such cases. We take the view that there is always an interface of some description, that warrants a discussion about the design of that interface. At the same time, we realise that different interaction scenarios will have very different conditions, and although we aim for general principles, we do not expect to be able to find too narrow a set of principles that is applicable such a wide set of cases.

Application of ID to CC

Norman’s conceptual model approach (Norman, 1988) popularised several key principles for the design of ‘everyday things’. His principles were rapidly applied to interactive technologies.

One of Norman’s most influential usability principles is *perceived affordance*. This describes a person’s conception of the various things you can do with a given object. This encompasses the heuristic experience of working with a system and broadly outlines the ability of the user to perceive and recognise a system’s interface. A common example of this is a door handle; a door handle affords pulling, as its physical

properties constrain what can be done with it in relation to its environment (Rogers, Preece, and Sharp, 2007). This is the same for a mouse button, which has a physical relationship to the digital interface it controls. The digital interface itself also offers perceived affordances as it too can be described as having constraints, and intuitive heuristic methodologies can be applied to it. For example a user can use past experience or common sense about what a digital button might do when they click on it.

Norman's principle of *visibility* is the simple idea that the more visible the operations of the system are, the more likely users will be able to know what to do next (Rogers, Preece, and Sharp, 2007). The complexity of CC systems often requires that functions are simplified or hidden from the user. This can lead to a conceptual black box. A user sees an input and receives an output, but the extent and nature of what happened in-between can be hard to understand.

The lack of visibility of the process can also scale with complexity. A system which appears simple at first can, in a CC process, become complex and unmanageable for a user to effectively make decisions. For example, a user may be able to manage a simple 2D physical model such as balls bouncing around in a 2D environment, which are easy to recognise and mentally model. But if the environment contains any more than a few interacting agents the ability of the user to make meaningful and effective decisions decreases. This has a downward-spiral effect for users; as interactions become more complex their ability to maintain and develop a clear conceptual model decreases along with the systems visibility.

Mapping, the direct relationship between controls and their effect on a system, is closely related to visibility, contributing further to the intelligibility of the system. A user's ability to interpret the affordances of an interface element depends in part on the arrangement of interface elements as they are presented to the user. In CC systems that are designed to enable users to manage computationally complex scenarios, it becomes paramount to the intelligibility of the system that a coherent mapping remains visible and intuitive.

At this point, conventional wisdom might say that if the complexity and opacity of CC systems are completely at odds with these very foundational principles of ID, then perhaps ID principles are simply not relevant to CC.

We contend that instead ID and CC should evolve together to develop a rich model of ID that is specifically suited to CC (as well as a wealth of other situations involving rich interaction with AI systems that are likely in the near future). Part of the argument for this is that it is hard to think of CC systems in the absence of some form of interaction. Instead, despite the isolated lab-based nature of much CC research, the majority of CC researchers do take care to emphasise the essential embeddedness of art in a complex of human social behaviour, and ultimately aspire to create work that interconnects with this complex, whether in the form of simulated artist agents, creativity support tools, Twitter bots, multi-agent simulations or other types of interactive systems. Work in CC that displays ignorance of this inherent network complexity has not generally been widely accepted.

In short, all CC research requires the developers to 'design' (at least establish and observe) interactions at some point along the way.

Lubart states that computers can facilitate (a) the management of creative work, (b) communication between individuals collaborating on creative projects, (c) the use of creativity enhancement techniques, and (d) the creative act, through integrated human-computer cooperation during idea production (Lubart, 2005). If we consider different CC goals according to Lubart's classification of the different ways in which computers can act as creative partners, then ID clearly plays a role in each of these forms of interaction, almost by definition. The clear application of the ID principles discussed so far becomes harder as we work our way through this list. In extension of Lubart's list, we could add (e) the complete artistic autonomy of the system, interacting with others only as an artist interacts with her audience. One contention of this paper is that even the latter should be subjected to ID thinking, and that ID principles need to be modified to extend that far.

One practical systematic approach to this conundrum is to distinguish between areas where (or levels at which) transparency is needed and where opacity can be allowed, extracting the former into what both programmers and designers understand as an 'interface'. This is only to reiterate conventional ID thinking in a way that might be more palatable for the above concerns in CC.

Dennett's (1989) *intentional stance* offers one well-known strategy for interaction with a certain group of complex systems – other humans and animals. We 'model' (i.e., intuitively understand) these systems not in terms of their physics or mechanical design, but in terms of their thoughts, intentions and goals. This reduces the complexity involved in predicting the system's behaviour, and we do this innately because our brains have evolved to do so. It would be useless trying to use a 'physics stance' to model what an adversary was going to do next, even though it will help predict the swing of their arm in a fist fight. In particular we have specifically evolved to 'model' the minds of other humans, the cognitive product of a competitive coevolutionary race (Whiten and Byrne, 1997; Dunbar, 2004; Boyd and Richerson, 1985) that some think is key to our sense of consciousness. Making systems that behave exactly like humans might be a good strategy in CC, but it is interesting to note that this would not make them necessarily easy to model.

Gaver argues that as computing has become progressively more ubiquitous, it has brought with it the values of the workplace. Concerns for clarity, efficiency, productivity and a preoccupation with finding solutions to problems have been imposed on digital devices as if they are limited purely to mirroring the work required to achieve an ordinary life, such as the completion of everyday chores (Gaver, 2002). He suggests that the idea of *homo ludens*, a term taken from Huizinga, humans defined first as playful creatures, brings our curiosity, our love of diversion, our explorations, inventions and wonder to the fore of designing interactive technologies. Gaver is intentional in his definition of play, and diverges from Huizinga's definition, preferring

Kaprow's (Gaver, 2002) definition of play as distinct from games. Kaprow acknowledges that while games and play:

“both involve free fantasy and apparent spontaneity, both may have clear structures, both may (but needn't) require special skills that enhance the playing. Play, however, offers satisfaction, not in some stated practical outcome, some immediate accomplishment, but rather in continuous participation as its own end. Taking sides, victory, and defeat, all irrelevant in play, are the chief requisites of game. In play one is carefree; in a game one is anxious about winning(Gaver, 2002).”

Gaver's application of homo ludens comes to bear on CC in that, if we are to leave work behind and design systems that embrace human creativity, then we need to intentionally seek play as a form of engagement. “This is an engagement that has no fixed path or end, but instead involves a wide-ranging conversation with the circumstances and situations that give it rise.” (Gaver, 2002), it is important that open-ended and self motivated forms of interaction are employed. This enables users to find new perspectives and new ways to create, “through ambitions, relationships, and ideals” (Gaver, 2002).

Gordon Pask, an early proponent and practitioner of cybernetics sought to build machines that coexisted in a mutually constructive relationships with users (Negroponte, 1975). Pask defined this process as conversation theory. He was specifically interested in how human-machine interactions could be subject to context and interpretation as an additional way of locating meaning in the interaction with the machine. Recent practitioners of conversation theory include Haque (2007), who argues that creative use of computers needs to incorporate these mutually constructive relationships as a means of expanding creative potential.

Likewise, many CC researchers have attempted to engage with the open-ended nature of creative discovery (Saunders and Gero, 2002), building on creativity research (Csikszentmihalyi and Sternberg, 1988; Boden, 1990) to design systems that exhibit these properties. Biological evolution has been one source of inspiration here. Whilst building emergent complexity into closed computer systems has proven difficult (Bown and McCormack, 2010), several researchers have reported moderate success with interactive genetic algorithms (IGAs) as human-computer collaborative tools for open-ended search. Stanley and Lehman Stanley and Lehman (2015) have been notable advocates for the open-ended nature of creativity following the observation that a distributed IGA system, *Picbreeder*, built by their team, was used by participants in a way that clearly demonstrated an absence of preconceived goals. Users were observed selecting images to evolve and then allowing the image evolution process to lead them towards recognisable shapes. Highly recognisable images emerged, such as faces and cars, but not because users set out to draw faces and cars. In their book Stanley and Lehman develop such observations, as well as their research in novelty search as a form of optimisation, into a general theory that attempts to define objectives as more of an impediment than a help to true discovery. By definition, they argue, a hard creative problem does not in-

dicates the direction in which you should head to discover the solution, so setting out in the apparent direction of the objective is a flawed approach.

Stanley and Lehman are in a sense restating a well-known principle of creativity theorists, with added evidence from computer science. Perkins (1996), for example, examines successful creative individuals and identifies their most common strategy as being one of spreading their bets across a wide range of solutions. Others in design creativity have identified a form of reverse creative discovery where problems are found to suit existing solutions (the story of the Post-It Note is one of the best known examples).

In this discussion we have encountered a series of ideas around the intersection of ID and CC: that Norman's principles of visibility and a clear conceptual model have limits in the context of the complex and opaque nature of CC systems; that, according to Gaver (2002), opacity is acceptable in the context of playful interaction, and that according to Stanley, Lehman and many others, open-ended search, which we closely associate with playful interaction, is critical to true creative search; and that for any system or use-case we should attempt to identify where transparency is needed (i.e., in the context of goal-directed functional behaviour) and where opacity can be accepted (i.e., in the context of open-ended search, with conditions attached), in the creation of CC interfaces.

Visualising Structure and Trajectory

The above formulation still does not give much insight into specific methods for breaking down CC-ID problems to find suitable balances between opaque and transparent aspects of interaction. Our claim is only that it reframes a common problem from ID in a way that is palatable for CC. In our previous work studying popular end-user generative music composition tools (Bray and Bown, 2014), we have sometimes found it useful to think about how users attempt to understand the behaviour of the system in terms of a breakdown between its structure and its trajectory. Most systems can be easily decomposed into these two parts: a structure that is generally assumed to be fixed, but may be mutable to some minor extent, and a set of ongoing movements or state changes around that structure. For example, dropping a pinball into a pinball machine, we think of the fixed structure of the pinball machine layout dictating the trajectory of the pinball. We think of all traditional acoustic instruments as having specific fixed structures around which a musician defines a trajectory. The instruments can be (imperfectly) parameterised, as is often seen performed in the creation of virtual instruments.

Our suggestion is that Norman's visibility or clear system model occurs wherever the user can clearly perceive the system's structure, and get a handle on how this structure dictates the trajectory.

By contrast, more complex generative systems can get harder to model because, we suggest, it is harder to see the structure and pull it apart from the complex movement of the system. McCormack and McIlwain's *Nodal* (McCormack et al., 2007) is an example we explicitly looked at in this way

(Bray and Bown, 2014). Whilst there is always a clear structure in Nodal, it is hard to tell how it will influence the system's unfolding trajectory just by looking at it. In the case of Nodal, the user is expected to build the networks by hand, so this kind of opacity could be seen as an impediment. But it may not be: another view is that the user develops strategies for progressing their work, and heuristics for thinking about what is going on, even though they struggle to develop a clear model of the system behaviour. This type of user behaviour would seem to make a clear break into the domain discussed above as more playful open-ended search.

Another possibility, alluded to here, is that users, drawing on their general intelligence, are able over time to better model the system, becoming experts. This expertise might be equivalent in ways to the species-specific expertise we have discussed in the case of Dennett's intentional stance. Whether or not these derived models have any common abstract properties would be of great interest.

Specific CC scenarios

We now briefly work through how these ideas might be applied to specific CC scenarios, and look at the different ways in which we might apply ID concepts to these different areas.

The first case we consider is already introduced above: generative tools such as *Nodal* that employ different generative paradigms with diverse approaches to user interaction. In this case, the distinctions between opaque and transparent approaches are applied straightforwardly, as described above. If a system can be transparent, then it could be potentially used in a more goal-directed manner. Opaque systems can often only be used in a goal directed way if you are focusing on process-based creativity, otherwise they require a search-based approach. However, as we suggest, there may be strategies for making opaque systems less opaque, through their representation, learnability and so on.

Another case we have already discussed are IGAs, which, as we have seen, seem to lend themselves to open-ended search more readily than to goal directed search. This is perhaps due to their randomness – it would be frustrating to aim for goals because you'd be forever looking for the next link in the chain, much better to respond to the available options. Conventional GA theory does however require certain conditions of transparency, for example in that mutated objects should be similar to their parents; there must be smoothness and consistency which we can think of as something clearly modellable.

Corpus-based learning approaches are interesting from a transparency point of view because they rarely offer any intuitive way to understand what the system has learnt. Such systems also tend to be self-contained processes, transforming an 'inspiring set' into new candidate outputs. There are rarely coherent ways for users to get involved in this process except in the tweaking of parameters, although this is hugely important for successful results, and ID research has been conducted in this area. Martin, Jin, and Bown (2011) observe that this imperviousness to user input has been a key problem to making usable systems.

More recently, work such as that of Pachet and Roy (2014) involves the use of corpus-based systems that 'mash-up' musical styles, where there is plausibly more involvement of the user. Here we may approach something akin to an intentional stance approach, where we might ask for, say, a performance of a Beatles song in the style of Wagner. Here the user can clearly engage in tasks in goal-directed or open-ended ways.

User interfaces that allow users to specify target goals are now common across a range of application areas, and is becoming an increasingly active area in architecture, where we need to reach multi-objective targets of, for example, structural stability, temperature regulation and visual criteria all at once. Some interfaces consist of a simple bank of sliders, whilst others must be programmed. In other producer-critic models, we might interactively evolve the fitness function that is used to do targeted evolution of an outcome. In other cases we might train a neural network to learn a preference. Veale (2015) argues that these types of CC systems go beyond 'mere generation' and take on artistic responsibility for selection or evaluation. Users become meta-creators, creating with and through CC processes.

Looking at the bigger picture, Plotkin (2009) discusses how automated discovery methods transform computers from machines that we instruct to perform specific tasks, to *genies* that respond to specific requests for outcomes. Such systems may therefore have wide reaching implications for how we interact with computers on a daily basis.

Lastly, we have recently seen work in art-making systems that explain themselves to their audience in natural language (Colton and Ventura, 2014), as a form of interactive experience that, in Colton's terms, 'frames' the artwork with additional relevant information. This is an approach that very much places ID at the centre of the design of CC systems, both by holistically considering the user experience associated with evaluating art, and more specifically by breaking from the unidimensional approach to aesthetic evaluation just mentioned, considering instead a rich multimodal set of possible interactions and judgements. What has yet to be elaborated on in theoretical terms is how we might frame these interactions between a machine artist and its audience in terms of a set of goals. The makers of the system invariably have goals when they place the system in front of people, just as other software developers do, and indeed, individual artists do when they interact. If the goal is open-ended co-creative search then the ID issues will be framed by this, and if the goal is to produce entertaining artworks for the home, then the ID issues will be different.

Conclusion

In this paper we have presented a series of ideas that can be summarised as follows:

- Opacity is inherent to CC but generally problematic in ID, except in the context of playful interaction.
- Open-ended search, associated with much creativity, can be stimulated through playful interaction.
- An ID approach to CC would be to attempt to work out what can be made visible, and what cannot, and work

out how the opaque elements can still be usable given the above assumptions.

- It may be possible to make some opaque aspects of systems more visible by considering how we mentally model these systems. We pose a distinction between structure and trajectory as one way this might be handled.

As each of these areas within CC research matures and starts to be applied in real software, the ID issues become more relevant, apparent, and better understood. As this happens we have the opportunity to build ID techniques specific to advanced CC.

References

- Boden, M. 1990. *The Creative Mind*. George Weidenfeld and Nicholson Ltd.
- Bown, O., and McCormack, J. 2010. Taming nature: tapping the creative potential of ecosystem models in the arts. *Digital Creativity* 21(4):215–231.
- Bown, O. 2014. Empirically grounding the evaluation of creative systems: incorporating interaction design. In *Proceedings of the Fifth International Conference on Computational Creativity*.
- Boyd, R., and Richerson, P. J. 1985. *Culture and the Evolutionary Process*. Chicago, IL, US: University of Chicago Press.
- Bray, L., and Bown, O. 2014. Linear and non-linear composition systems: User experience in nodal and pro tools. In *Proceedings of the Australian Computer Music Association Conference*.
- Candy, L., and Edmonds, E. A. 1997. Supporting the creative user: a criteria-based approach to interaction design. *Design Studies* 18(2):185–194.
- Carroll, E. A. 2013. *Quantifying the personal creative experience: evaluation of digital creativity support tools using self-report and physiological responses*. Ph.D. Dissertation.
- Colton, S., and Ventura, D. 2014. You can't know my mind: A festival of computational creativity. In *Proceedings of ICCO 2014 (International Conference on Computational Creativity)*.
- Csikszentmihalyi, M., and Sternberg, R. 1988. The nature of creativity: Contemporary psychological perspectives. *Society, culture, and person: A systems view of creativity* 325–339.
- Dennett, D. C. 1989. *The intentional stance*. MIT press.
- Dunbar, R. 2004. *Grooming, Gossip and the Evolution of Language*. London: Faber and Faber.
- Gaver, B. 2002. Designing for Homo Ludens, Still. *Interaction Research Studio, Goldsmiths, University of London, I3 Magazine No. 12* 163–178.
- Gero, J. S. 1990. Design prototypes: a knowledge representation schema for design. *AI magazine* 11(4):26.
- Haque, U. 2007. The Architectural Relevance of Gordon Pask. *Architectural Design* 77(4):54–61.
- Lubart, T. 2005. How can computers be partners in the creative process: classification and commentary on the special issue. *International Journal of Human-Computer Studies* 63(4):365–369.
- Martin, A.; Jin, C.; and Bown, O. 2011. A toolkit for designing interactive musical agents. In *Proceedings of the 23rd Australian Computer-Human Interaction Conference*. ACM.
- McCormack, J.; McIlwain, P.; Lane, A.; and Dorin, A. 2007. Generative composition with nodal. In Miranda, E., ed., *Workshop on Music and Artificial Life*.
- Negroponte, N. 1975. *Soft architecture machines*. MIT press Cambridge, MA.
- Norman, D. 1988. *The Design of Everyday Things*. New York: Basic Books.
- Pachet, F., and Roy, P. 2014. Non-conformant harmonization: The real book in the style of take 6. In *Proceedings of ICCO 2014 (International Conference on Computational Creativity)*.
- Perkins, D. N. 1996. Creativity: Beyond the darwinian paradigm. In Boden, M., ed., *Dimensions of Creativity*. MIT Press. chapter 5, 119–142.
- Plotkin, R. 2009. *The genie in the machine: how computer-automated inventing is revolutionizing law and business*. Stanford University Press.
- Rogers, Y.; Preece, J.; and Sharp, H. 2007. Interaction design.
- Saunders, R., and Gero, J. S. 2002. How to study artificial creativity. In *Proceedings of the 4th conference on Creativity & cognition*, 80–87. ACM.
- Shneiderman, B. 2007. Creativity Support Tools: Accelerating Discovery and Innovation. *Communications of the ACM* 50(12).
- Stanley, K. O., and Lehman, J. 2015. Why greatness cannot be planned. *Springer Science Business Media. doi* 10:978–3.
- Veale, T. 2015. Creativity ex machina. *Routledge Handbook of Language and Creativity* Rodney Jones:353–366.
- Whiten, A., and Byrne, R. W. 1997. *Machiavellian Intelligence II: Extensions and Evaluations*. Cambridge, UK: CUP.