

Casual Creators in the Wild: A Typology of Commercial Generative Creativity Support Tools

Elena Petrovskaya
Computer Science,
University of York, York
ep1167@york.ac.uk

Sebastian Deterding
Digital Creativity Labs,
University of York, York
sebastian.deterding@york.ac.uk

Simon Colton
School of Electronic Engineering,
Queen Mary University London.
SensiLab, Faculty of IT,
Monash University, Australia
s.colton@qmul.ac.uk

Abstract

Casual creators are a genre of creativity support tool that use generative systems to empower autotelic and enjoyable amateur creativity. They have been posited as a unique means of democratising creativity through the support of user exploration via system generativity, yet little is known about what casual creators are actually available to wider audiences. We conducted a qualitative analysis of currently available casual creators on the App Store. We found three categories of interaction techniques across 89 commercially available apps that qualified as casual creators, which we describe in their exploration potential, feedback speed, and user autonomy.

Introduction



Figure 1. *Kaleidoscope Drawing Pad*, a casual creator which involves creation of abstract art through touching the screen.

Casual creators are a genre of software that playfully facilitates autotelic creativity – for the sake of its own enjoyment (Compton and Mateas 2015). Unlike professional computer-aided design or creativity support systems (such as Photoshop or the Unreal Engine), the primary purpose of casual creators is not to create a professional end product; nor do they require a particular starting level of professional skill, creativity, or even a specific creative intent. Their main aim is affording the enjoyable experience of engaging in creativity, placing focus on the creative process over the creative product.

Examples include applications for generating visual imagery, combining basic musical elements into music compositions, or choosing an image from a set to generate similar art. For example, in the app *Kaleidoscope Drawing Pad*

(Bejoy Mobile 2012) (Figure 1), touching the screen generates kaleidoscopic images for amusement.

Casual creators have seen a quick uptake among researchers and practitioners in computational creativity, who have investigated new casual creators across domains, from making live music to writing stories to making whole mobile games (Samuel, Mateas, and Wardrip-Fruin 2016; Kreminski and Wardrip-Fruin 2019; Nelson et al. 2017; Lorway et al. 2019). This is motivated by the promise that casual creators could help democratise creative practice (Nelson et al., 2017) as part of the general rise of amateur (digital) making, which is seen to improve both individual and community wellbeing (Compton 2019, 61-66; Gauntlett 2013).

Compton (2019) proposes a collection of design patterns for casual creators to successfully support autotelic creativity. The key premise of these patterns is to support user exploration by reducing error and offering entertaining feedback, allowing users to approach the creative process with confidence and pride in their work. These patterns work thanks to the incorporation of *generativity* into the casual creator system, in that the system will alter and enhance the initial user input to produce a more complex output.

However, while generativity is a defining feature of *intentionally designed* casual creators, Compton (2019) acknowledges that users have adopted a range of systems for autotelic creativity that lack generativity; Compton labels such systems as casual creators ‘*in some way*’. An example is the app *Let’s Create! Pottery* (Infinite Dreams 2011), in which the user creates digital pottery in real-time. This system does not involve a generative element, yet is used for autotelic creativity.

Existing work on casual creators has to date focused on such pre-existing exemplars *avant la lettre* and prototype systems by researchers and artists, with quite small audiences of lab participants, drawing on social networks of researchers and artists. Initial field work on other computational creativity systems (Ackerman and Pérez Y Pérez 2019) has drawn attention to the potential discrepancy between research-based

designs of such systems and their use in the real world, as well as the potential to tap into user feedback to help further design.

If casual creators indeed aim to reduce the barriers of creative practice for broad, general audiences, this raises the question of whether and how casual creators can be designed to actually be broadly accessible and engaging. Do intentionally designed casual creators, with generativity at their core, survive ‘in the wild’ (Rogers and Marshall 2017) in the same way as systems which have been naturally adapted by users to be casual creators ‘in some way’?

One way of approaching this question is to see whether there are already commercially successful casual creators beyond research labs and art exhibitions, and analyse what design characteristics (if any) they have in common. The logic here is that market pressures are likely to have spurred the evolution and spread of designs for casual creation that ‘work’ for broad audiences (Gee 2003). To be sure, commercial *availability* does not equate public *appeal* – the majority of games available on Steam or itch.io see little if any uptake. But there is some information to be gleaned from what kinds of applications of a particular genre actually exist in an open marketplace.

In this paper, we report the results from a study following this logic. Specifically, we were interested in two questions: (1) *Are there already applications commercially offered to general audiences that can be counted as casual creators?* And if so: (2) *How are their interactions designed?* More precisely, how does their design realise key aspects research has stipulated as essential to supporting casual creation, and can we find specific reoccurring *interaction techniques* that do so? Interaction techniques (Hinckley et al. 2014) describe the particular arrangements of input and output that allow a user to perform a particular task.

To answer these questions, we conducted a qualitative review and analysis of creativity support applications commercially available on the Apple App Store in 2019. We found numerous already-existing applications that qualify as casual creators, which clustered into three categories of interaction techniques, and differ in the user autonomy, exploration potential, and feedback they afford. Thus, our paper sketches a first landscape of currently commercially available casual creators, which provides a broader empirical grounding for research. Those looking to develop and deploy their own casual creation systems might benefit from understanding the current landscape, and how their design might be best adapted to fit into this. It is also interesting to see how generative systems are being used in the real world, rather than in content generation for research or artistic purposes, which may provide future insight into user interaction with these systems.

The rest of this paper is organised as follows. We first introduce the concept of casual creators and detail four key

design dimensions of casual creators which structure our empirical work: possibility space, feedback, and user/system autonomy. Next, we report our research rationale and design. The Results section details the three types of commercial casual creators we found along the three design dimensions. We close with a discussion of our findings, their limitations and ramifications for future work.

Background

Casual creators are designed with the goal of providing a *fun* and *pleasurable* experience of the creative process for creative amateurs (Compton 2019). Interaction with casual creators has even been speculated to be *playful*, placing them on a spectrum with digital games (Compton 2019).

As mentioned above, the key computational enabler underlying the design of casual creators is *generativity*. A generative method consists of a function which takes initial input and creates a different – often bigger or more complex – output without any additional contribution (Compton, Osborn, and Mateas 2013).

Generative algorithms are often used for art. The key feature of a generative art system is that the user lets a computer system take over some of the decision-making. This is useful for creativity support in several ways: because art is an iterative process, incorporating a computer may help with time-based work by tightening and quickening the iterative cycle. Moreover, it can help with making the decision space smaller and more manageable (Boden and Edmonds 2009).

Generative systems are useful for supporting amateur creativity as they take away some of the responsibility of user creation and add complexity to the final product. Where creativity support tools merely automate or digitally mediate certain steps in the wholesale production of a creative artefact (like copy-pasting or moving a string of notes in a digital music composition), casual creators afford enjoyable amateur creativity by letting users engage with a generative system – e.g. feeding it particular inputs, then observing and selecting outputs from various choices.

Beyond the defining features of generative systems supporting autotelic amateur creativity, Compton and Mateas (2015, Compton 2019) identify a range of design patterns for casual creators. Across these patterns, two key functional aspects or dimensions are repeatedly highlighted and discussed: (1) their possibility space and (2) their fast, entertaining feedback. A further quality discussed or emphasised throughout relates to (3) the relative autonomy of the user over the creative output, also referred to as the power-control trade-off. We will discuss each of these patterns below.

Explorability: Meaningful-yet-Limited Possibility Space

Generative systems can be described in terms of their explorability or *possibility space*: The range of possible

outputs they can produce. As Compton and Mateas (2015; Compton, 2019) stress, for a creativity support tool to ‘feel’ truly creative, their possibility space must be large enough to continually produce novel, surprising outcomes. On the other hand, the almost limitless possibility space of professional creativity support tools can be quickly overwhelming and frustrating for amateur creators. Casual creators offer a comparatively limited possibility space, e.g. by providing starting points, or drastically limiting potential inputs. This guides user exploration of the possibility space and minimises room for error, creative blocks, and anxieties, enabling users to create products which may not have been possible purely on the basis of their own ability.

Feedback: Fast and Entertaining Evaluations

As users interact with a generative system underlying a casual creator, they start to build an understanding or mental model of how the system works – how different kinds of input and input dimensions shape the output. Similarly, in creative processes, creators compare their creative vision or idea with its materialised execution.

Compton (2019) argues that feedback in casual creators for any produced artefact should be fast and entertaining. In particular, Compton and Mateas (2015) highlight the theory of reflection-in-action, which argues that people learn from reciprocal interaction with an artefact and then reflecting on said interaction (Schon and Wiggins 1992). Fast and entertaining feedback speeds up learning about the generative system and the creative material one is working with, but also makes the overall experience pleasurable for the users and allows them to feel progress in their creative activity.

Autonomy: Limited-yet-Meaningful Control

Compton (2019) puts forward a particular trade-off between power and control as another characteristic of casual creators. Software for creative professionals aims to give the user full, detailed control over the system and end product. Such control is not necessarily essential for amateur creators or autotelic creation: if a user is less concerned with the outcome, they do not need full control over its every last detail. Hence, casual creators shift focus from control to support: they *empower* the user to produce relatively ‘polished’ outcomes relatively quickly by taking over a large portion of the creative process. This is their *power-control trade-off*: users sacrifice part of their control over the creative process and product in exchange for increased aesthetic ‘power’ of the overall human-computer system. Transferring some control from the generative system to the user can accelerate the exploration of possibility spaces and learning for the human user, can ensure the end product fits certain aesthetic qualities and requirements, and can make the process more accessible and enjoyable for the amateur user. Again, casual creators achieve this thanks to their generative systems

producing rapid and varied outputs, with minimal user input required.

An analytically useful way of translating the trade-off between two aspects (power and control) into one aspect is the relative *autonomy* of user and computing system, defined as the extent to which an agent has independence over their choices and actions (Barber and Martin 1999). In fact, casual creators have been characterised as *mixed-initiative creative interfaces* – systems in which human and computer users interact as creative collaborators in a feedback cycle (Deterding et al. 2017; Yannakakis, Liapis, and Alexopoulos 2014). Such systems lie on the midpoint of a spectrum of user and system autonomy between strong computational creativity systems, in which the computer is a fully autonomous creator and humans are merely the audience, and creativity support tools, in which the computer is a tool for the support of fully autonomous human activity. In mixed-initiative systems, neither side has full autonomy over the creative process and outcome. Due to the generative nature of casual creators, user interaction with these creative systems can even be conceptualized as *discovery* rather than *making* of creative outputs.

Study Aims and Method

As noted, work on casual creators has been chiefly concerned with defining the genre and deriving characteristics and design patterns for ‘good’ casual creators from select case studies that span research prototypes and artistic creations (Compton and Mateas 2015, Compton 2019). Little is known about how many and what kinds of casual creator applications are already broadly available and used ‘in the wild’, what design features they share, and how they realise the three characteristics of casual creators – a limited-yet-meaningful possibility space, limited-yet-meaningful user autonomy, and fast feedback. Specifically, do existing casual creators form some kind of meaningful types, categories, or sub-genres? To answer these questions, we set out to conduct a qualitative review of publicly available casual creators.

Data collection and sample generation

To establish a sample of casual creator applications in a reliable and replicable fashion, we broadly followed systematic review procedures, akin to Lister et al. (2014) review of game design elements in mHealth and fitness apps. Between September 26th and October 3rd, 2019, we conducted a search of the Apple UK App Store, running separate searches each for the keywords ‘creative’, ‘create’, ‘creativity’, ‘make’, ‘draw’, ‘art’, and ‘generative’. The Apple App Store was specifically targeted as casual creators are generally developed as mobile applications.

We then defined inclusion/exclusion criteria for counting an application as a ‘casual creator’. We used the definition of casual creators by Compton and Mateas (2015), outlined

above, as our inclusion criterium. The exclusion criterium was “professional creativity tools and apps that don’t involve a generative element”. This criterium was selected because, as noted above, we were interested in analyzing the type of software which can be classified as casual creators *by design*, rather than applications which have been adapted by users to be casual creators *in some way*. Generativity in the current paper was conceptualized on the basis of the definition used by Compton and Mateas (2015), who note that generative algorithms ought to produce a ‘wide and interesting space of possible valid artifacts’ (p. 4). Based on this, we included only generative applications which featured an element of randomness in the output, thus being technically able to inspire surprise in the creator.

Our initial search yielded a total 1,121 applications, of which 89 were taken for analysis after applying the inclusion and exclusion criteria.

Coding and evaluation

The applications of the included sample were downloaded onto an iPad and interacted with. This interaction allowed the main *interaction technique* (Hinckley et al. 2014) of the applications to be determined based on how the user interacted with the system throughout the creative process. Coding and analysis followed qualitative content analysis (Mayring 2004) to identify high-level recurring features, types, or genres. (In the context of qualitative analysis, *coding* refers to the process of identifying qualitative elements which represent concepts of interest).

Coding was carried out by the same researcher who conducted the initial search and inclusion/exclusion of apps. To ensure the validity of this analysis, coding was initially done at one point, and then repeated after a period of time at a second point. The inter-rater reliability between these time points was calculated as being 0.75, which according to Cohen (1968) constitutes substantial agreement. In situations where there was a disagreement between the two time points, the coding from the second time point was used. Repeated coding at separate time points is a common way of testing inter-coder reliability when only one coder is available (Mackey and Gass 2005).

All codes used were ones which were discovered through the analysis. Three categories naturally emerged from this coding based on shared interaction techniques. The applications in each category were then examined according to the three characteristics of casual creators outlined above: the exploration potential/possibility space, speed of feedback, and user autonomy.

Exploration potential was assessed descriptively based on the size of the possibility space available for exploration within the casual creator: what kind and how many different possible outputs can be produced after one user input?

Feedback speed was assessed by how quickly the user sees any system output after they complete their part of the creative act. Because none of the apps across categories featured any direct delays to system output, assessing feedback speed in time units would not provide a meaningful picture of the differences between apps and categories. Instead, we found it helpful to describe feedback speed in terms of whether the system presents output *while* or *after* the user creates.

Lastly, *autonomy* was assessed using Barber and Martin’s proposed scale of agent autonomy (Barber and Martin 1999), as follows:

- *Command-driven*: the agent does not make any of its own decisions about how to pursue a goal.
- *True consensus*: the agent works as a team member, sharing decision control with other agents.
- *Local autonomy*: the agent makes their decisions alone.
- *Supervised autonomy*: between command-driven and true consensus.
- *Supervisory autonomy*: between true consensus and local autonomy.

Results

Table 1 provides a brief summary of the three categories. We then present each category in turn, led by their shared interaction technique. For each category, we give an overview and examples, and then assessments of exploration potential, feedback speed and autonomy.

Category	Autonomy	Possibility space	Feedback speed
One-touch creativity	Supervised	Largest – infinite range.	Moderate
Vague creation	True consensus	Restricted by user input.	Slowest
Mutant shopping	True consensus	Restricted by starting variants	Quickest

Table 1. Overview of the categories and their features.

Category 1: One-touch creativity

Interaction technique

The interaction technique of this genre is *the repetition of one type of touch gesture*, such as tapping, swiping, or flicking the screen to interact with the interface, which leads to the generation of an output.

The user would use one finger to perform a gesture. Touching the screen with a certain gesture yields a generative output which is reflective of this gesture, but otherwise entirely

random – although, in some instances, the user has the option to specify initial parameters, such as colour, meaning the algorithm is parametrized. For example, in *Silk 2* (Vishnevsky 2010) (Figure 2), the user specifies colour and type of line, and the system takes these parameters to create a more complex output. The generative algorithm has stochastic elements, so the output is likely to surprise the user, as they could not fully preconceive the output based on their input. Finally, the generative output is constructive: the user only sees one version of the output, rather than several iterations.




Figure 2. Creation in the art system *Silk2*. The user draws and taps the screen, in response to which the system generates aesthetic outputs. The user also has the option to specify some parameters they would like to see in the output, such as line style.

Some one-touch creativity creators involve independent agents whose behavior is influenced by the gesture and makes up the output. For example, in the app *abcdefghijklmnopqrstuvwxyz* (joerg piringer 2010), the user directs the movement of agents visualised as letters of the alphabet (with musical accompaniment) (Figure 3).



Figure 3. A user swipes the screen in the casual creator *abcdefghijklmnopqrstuvwxyz* to move the letters.

One-touch creativity is the most frequent type of casual creator in our sample, spanning 60% of apps across visual art, music, and text. Notable examples include *Uzu*, a visual

interactive light show (Smith 2010), *Figure* (Reason Studios AB 2012), which does so with music by setting up tones and patterns and then creating music simply by sliding a finger, and a series of apps developed by Brian Eno which incorporate both art and music for multimedia creations – e.g. *Trope, Quarto* ('GenerativeMusic.Com | Apps by Brian Eno and Peter Chilvers' n.d.).

Visual casual creators of this category often focus on repetition: kaleidoscopes, fractals, and mandalas, e.g. the app *iOrnament* (science-to-touch 2012), which generates mandalas. Often, this is supplemented with a science theme, such the contextualisation of such art as the creation of particles or molecules.

Possibility space

One-touch creativity apps feature the largest possibility space out of all the casual creators. Each gesture yields a stochastically generated element of art from an essentially infinite range.

User autonomy

In this way, one-touch creativity casual creators allow only for supervised autonomy: there is an element of choice in how the user executes the gesture and some initial parameters, yet the user does not have much control over the output.

Feedback speed

Feedback provision is slower than mutant shopping but faster than vague creation.

Category 2: Vague creation

Interaction technique

Vague creation creators are less common, accounting for only around 25% of our sample. In these applications, users interact primarily through *drawing* something on the screen with their fingers. The user implements some vague or unfinished shape, and the generative system completes this into something advanced. Sometimes this is manifested through the user creating parts, which the program combines to make a whole creation. A distinction between this category and the previous is that the output produced through vague creation is of a *complete* piece.

The underlying algorithms of creators in this category appear to be parametrized: they produce exactly one deterministic output in response to the user input. Thus, the user experiences less surprise than one-touch creativity.

This type of casual creator spans the creation of visual art and also music. For example, in the app *PendantMaker* (Compton and Mateas 2015), the user makes vague patterns and the app makes a pendant from them. Similarly, in *Scape*, the program recombines musical elements to create new compositions (Figure 4).




Figure 4. *Scope*, a musical app in which the user combines shapes to create musical compositions.

User autonomy

There is true consensus in autonomy – while the user does not have full control over the output, the decisions they make on input largely constrain the finished product.

Possibility space

Because the generative algorithm is parametrized, the generative output of this category is therefore based on initial ideas from the user, and the possibility space is also restricted by user input.

Feedback speed

Feedback is slower than the other two categories.

Category 3: Mutant shopping

Interaction Technique

Compton and Mateas (2005) originally identified mutant shopping as a design pattern, which consists of using suggested alternatives to browse the parameter space. We found that it constitutes an interaction technique genre of its own, although this is the least common category, covering only 10% of our sample. This type of casual creator is interacted with by *choosing* from a selection of starting variants, which serve as the starting parameters. The generate-and-test system then produces a new line-up of variants based on the user's choice. In some cases, the user has the option to edit the 'mutant' before it is reproduced.



Figure 5. *Kandinsky.io* (Khosravi 2019) – users select their favourite wallpaper variant.

In the casual creator *Kandinsky.io* (Figure 5), users can generate wallpaper art for their phones, and select their preferred variant out of several others. The element of randomness is quite low in the output, meaning the user is likely to have a clear understanding of what they can expect to see, and is unlikely to experience a high level of surprise. This category is akin to evolutionary art (Romero 2008).

User autonomy

The system and user have true consensus – equally split control over the outcome: the user has autonomy but is constrained by the starting variants.

Possibility space

Exploration potential is very low here: users are only able to create an image by selecting an existing option to 'mutate'. Given this parametrised nature of the generative algorithm, the possibility space is therefore heavily constrained by the starting variants.

Feedback speed

Feedback provision by the system occurs after every choice the user makes. In this way, feedback is immediate after every choice, and assists with the iterations of the generative system.

Discussion

In this paper, we provided a review and typology of commercially available casual creators. To our knowledge, this is the first broad, data-based assessment of the prevalence of casual creators 'in the wild': Compton (2019, p. 6) stipulates that "hundreds (or thousands) of casual creator systems are already part of people's lives", but does not provide any evidence for this claim. We found that 89 of 1,121 or about 8% of applications on the Apple App Store findable with search terms associated with artistic creativity qualified as a casual creator, meaning casual creators are indeed available in the wild.

In addition, we employed a qualitative methodology to analyse the main interaction techniques used by said apps, and three distinct categories of casual creators emerged: one-touch creativity, vague creation, and mutant shopping. Again, this presents to our knowledge the first typology of common casual creators. While we cannot claim that these three are the only 'working' interaction techniques for casual creators, we do think they may provide insight into what kinds of interaction techniques application developers consider to be proven and working.

We analyzed these categories according to three core characteristics of casual creators: size of the possibility space, speed of feedback, and user autonomy. We found that there were distinct differences in the levels of each characteristic between the three categories. We particularly found that stochastic versus deterministic outputs

An initial overview of the different types of generative algorithms employed in the studied casual creators also lays down some understanding of the optimal level of generativity to benefit users. It is interesting that the most widely available category of casual creators – one-touch creativity – is also the type of casual creator with the biggest potential to surprise the user with the output. The implications of surprise in generative art creation are something we would like to consider in future work.

Limitations

The most pertinent limitation to this work is that it is based on a limited sample, assessing only applications available on one particular digital marketplace (the Apple App store) at one point in time (2019). Also, the number of available applications says nothing about their actual uptake in terms of e.g. number of installs, users, or time spent on app.

Furthermore, coding and analysis were conducted qualitatively by a single researcher. While re-coding at a separate time point ensured some degree of inter-coder reliability, we acknowledge that re-analysis with several coders and a pre-defined coding handbook would produce more reliable results. This holds especially for categorizing applications in terms of what experiences of autotelic creativity they afford.

Future work

This work is a first step in assessing the evolving landscape of creativity support tools and casual creators. We lay down foundations of how different levels of generativity between casual creators may affect the way users interact with the systems, and in our future work we will explore this through a more user-centered approach.

Next steps should also look further into the general user experience of casual creators, and how the described design features of possibility space size, feedback speed, and user autonomy contribute to this experience. We are already working on the construction of a grounded theory of what might motivate the construction and sustainment of user engagement with these apps. Initial findings show the distinction between casual creators as interfaces and casual creation as a mindset. We are interested in seeing what the contribution to this mindset is of the various features outlined in the current work, and how the design of casual creators can be further improved to facilitate autotelic creativity.

We hope that this paper helps build an initial understanding of the user side of casual creators, and patterns of interaction with casual creators in the wild.

Conclusion

This study conducted a qualitative review of applications designed to support casual creativity on the AppStore and found 89 commercially available apps that fit the definition of casual creators. Qualitative analysis revealed three

distinct categories of interaction techniques. Those categories were analyzed based on core design characteristics of casual creators: size of the possibility space, feedback speed, and the level of user autonomy in the creative act. The categories of casual creators differed in their levels of these characteristics and the level of generativity provided by the systems. This work outlines an ‘in the wild’ landscape of casual creators, and points towards directions for further work into the user experience of amateur digital creativity, as well as support for those looking to commercially release their own casual creators.

Acknowledgements

This work was supported by the EPSRC Centre for Doctoral Training in Intelligent Games & Games Intelligence (IGGI) [EP/L015846/1 and EP/S022325/1] and the Digital Creativity Labs (digitalcreativity.ac.uk), jointly funded by EPSRC/AHRC/Innovate UK under grant no. EP/M023265/1.

References

- Ackerman, Margareta, and Rafael Pérez y Pérez. 2019. ‘Field Work in Computational Creativity.’ In *Proceedings of the International Conference on Computational Creativity*, 236-242.
- Barber, K. S., and C. E. Martin. 1999. ‘Agent Autonomy: Specification, Measurement, and Dynamic Adjustment’. In *Proceedings of the Autonomy Control Software Workshop at Autonomous Agents*, 1999:8–15. Citeseer.
- Bejoy Mobile. 2012. *Kaleidoscope Drawing Pad*. Mobile application. <https://apps.apple.com/gb/app/kaleidoscope-drawing-pad/id525904070>
- Boden, Margaret A., and Ernest A. Edmonds. 2009. ‘What Is Generative Art?’ *Digital Creativity* 20 (1–2): 21–46.
- Cohen, Jacob. 1968. ‘Weighted Kappa: Nominal Scale Agreement Provision for Scaled Disagreement or Partial Credit.’ *Psychological Bulletin* 70 (4): 213.
- Compton, Kate, and Michael Mateas. 2015. ‘Casual Creators.’ In *Proceedings of the International Conference on Computational Creativity*, 228–235.
- Compton, Kate, Joseph C. Osborn, and Michael Mateas. 2013. ‘Generative Methods’. In *Proceedings of The Fourth Procedural Content Generation in Games Workshop, PCG*. Vol. 1.
- Davis, Nicholas Mark, Chih-Pin Hsiao, Kunwar Yashraj Singh, and Brian Magerko. 2016. ‘Co-Creative Drawing Agent with Object Recognition’. In *Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference*.
- Deterding, Sebastian, Jonathan Hook, Rebecca Fiebrink, Marco Gillies, Jeremy Gow, Memo Akten, Gillian Smith, Antonios Liapis, and Kate Compton. 2017. ‘Mixed-Initiative Creative Interfaces’. In *Proceedings of the 2017 CHI*

- Conference Extended Abstracts on Human Factors in Computing Systems*, 628–635. ACM.
- Gauntlett, David. 2013. *Making Is Connecting*. John Wiley & Sons.
- Gee, James Paul. 2003. ‘What Video Games Have to Teach Us about Learning and Literacy’. *Computers in Entertainment (CIE)* 1 (1): 20–20.
- ‘GenerativeMusic.Com | Apps by Brian Eno and Peter Chilvers’. n.d. Accessed 17 November 2019. <http://www.generative-music.com/>.
- Hinckley, Ken, Robert JK Jacob, Colin Ware, Jacob O. Wobbrock, and Daniel Wigdor. 2014. ‘Input/Output Devices and Interaction Techniques.’ *Computing Handbook, 3rd Ed. (1)* 21: 1–54.
- Infinite Dreams. 2011. *Let’s Create! Pottery*. Mobile application. <https://www.potterygame.com/>
- joerg piring. 2010. *Abcdefghijklmnopqrstuvwxyz*. Mobile application. <https://apps.apple.com/us/app/abcdefghijklmnopqrstuvwxyz/id358824534>
- Khosravi, Bardia. 2019. *Kandinsky.Io*. Mobile application. <https://kandinsky.io/>
- Kreminski, Max, and Noah Wardrip-Fruin. 2019. ‘Generative Games as Storytelling Partners’. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*, 103. ACM.
- Lister, Cameron, Joshua H. West, Ben Cannon, Tyler Sax, and David Brodegard. 2014. *Just a Fad? Gamification in Health and Fitness Apps*. *JMIR Serious Games* 2 (2), E9.
- Lorway, Norah, Edward Powley, Arthur Wilson, John A. Speakman, and Matthew Jarvis. 2019. ‘Autopia: An AI Collaborator for Live Coding Music Performances’ at ICLC 2020: *International Conference on Live Coding*. Limerick, Ireland, 5-7 Feb 2020.
- Mackey, Alison, and Susan Gass. 2005. ‘Common Data Collection Measures’. *Second Language Research: Methodology and Design*. Mahwah: Lawrence Erlbaum, 43–99.
- Mayring, Philipp. 2004. ‘Qualitative Content Analysis’. *A Companion to Qualitative Research* 1: 159–176.
- Nelson, Mark, Swen Gaudl, Simon Colton, Edward Powley, Blanca Perez Ferrer, Rob Saunders, Peter Ivey, and Michael Cook. 2017. ‘Fluidic Games in Cultural Contexts’. In *Proceedings of the International Conference on Computational Creativity*, 175–183.
- Reason Studios AB. 2012. *Figure*. Mobile application. <https://www.reasonstudios.com/mobile-apps>
- Rogers, Yvonne, and Paul Marshall. 2017. ‘Research in the Wild’. *Synthesis Lectures on Human-Centered Informatics* 10 (3): i–97.
- Romero, Juan J. 2008. *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*. Springer Science & Business Media.
- Samuel, Ben, Michael Mateas, and Noah Wardrip-Fruin. 2016. ‘The Design of Writing Buddy: A Mixed-Initiative Approach towards Computational Story Collaboration’. In *Proceedings of the International Conference on Interactive Digital Storytelling*, 388–396. Springer.
- Schon, Donald A., and Glenn Wiggins. 1992. ‘Kinds of Seeing and Their Functions in Designing’. *Design Studies* 13 (2): 135–156.
- science-to-touch. 2012. *IOrnamet*. Mobile application. <http://www.science-to-touch.com/en/iOrnamet.html>
- Smith, Jason. 2010. *Uzu, An Interactive Light Show*. Mobile application. <https://apps.apple.com/gb/app/uzu-an-interactive-light-show/id376551723>
- Togelius, Julian, Georgios N Yannakakis, Kenneth O Stanley, and Cameron Browne. 2011. ‘Search-Based Procedural Content Generation: A Taxonomy and Survey’. In *IEEE Transactions on Computational Intelligence and AI in Games* 3 (3): 172–86.
- Vishnevsky, Yuri. 2010. *Silk2*. Mobile application. <http://weavesilk.com/>
- Yannakakis, Georgios N, Antonios Liapis, and Constantine Alexopoulos. 2014. ‘Mixed-Initiative Co-Creativity’. In *Proceedings of the 9th International Conference on the Foundations of Digital Games*, 1–8.