

Knowledge-Level Creativity in Game Design

Adam M. Smith and Michael Mateas

Expressive Intelligence Studio
University of California, Santa Cruz
{amsmith,michaelm}@soe.ucsc.edu

Abstract

Drawing on inspirations outside of traditional computational creativity domains, we describe a theoretical explanation of creativity in game design as a knowledge seeking process. This process, based on the practices of human game designers and an extended analogy with creativity in science, is amenable to computational realization in the form of a discovery system. Further, the model of creativity it entails, creativity as the rational pursuit of curiosity, suggests a new perspective on existing artifact generation challenges and prompts a new mode of evaluation for creative agents (both human and machine).

Introduction

Paintings (Colton 2008), melodies (Cope 2005), and poems (Hartman 1996) are familiar domains for artifact generation in computational creativity (CC), and much established theory in the field is focused on evaluating such artifacts and the systems that produce them. In this paper, we draw inspiration for a new understanding of creativity from the less familiar (but no less creative) domain of game design. In its full generality, game design overlaps visual art, music, and other areas where there are many existing results, but where it stands apart is in its unavoidably deep, active interaction with the audience: in gameplay.

Crafting gameplay is the central focus of game design (Fullerton 2008). Play, however, is not an artifact to be generated directly. Instead, it is a result that emerges from the design of the formal rule system at the core of every game (Salen and Zimmerman 2004, chapter 12), a machine driven by external player actions.

Where, in visual art, we might judge the creativity (as *novelty* and *value*) of an artifact on the basis of the work's similarity to known pieces and its affective qualities (Pease, Winterstein and Colton 2001), it is not so easy to make direct statements about the properties of the artifacts in game design. Desirable games are celebrated for their *innovative gameplay* or the *fun* experiences they enable—these are properties of the artifact's interaction with the audience, not of the artifact itself. The focus on predominantly passive artifacts in CC, those which can be appre-

ciated via direct inspection rather than through interactive execution, has masked what is obvious in game design: that the desirability of artifacts is in their relationship to their environment.

Armed with such an understanding, we seek a theoretical explanation of creativity in game design—not the engineering application of established design knowledge, but the rarer experimentation that realizes new forms of gameplay and original player experiences. This theory should speak to both the artifacts and processes of game design, and do so in a way that meaningfully explains game design as done by humans as well as computational means. Towards capturing the richness of existing human design activity, we are most interested in a theory of *transformational* creativity (Boden 2004) that explains how designers build new conceptual spaces of game designs and reshape them in response to feedback experiences observing play.

We introduce a new theoretical model that is amenable to computational realization which describes creative game design as a knowledge-seeking process (a kind of active learning). Our broader contribution, creativity as the rational pursuit of curiosity, can provide an explanation of and suggest new questions for applications in traditional CC artifact generation domains.

In the following sections we will review established game design practices, draw an analogy between game design and scientific discovery, review and apply Newell's concept of the knowledge level, and then introduce our model of creativity. Finally we will conclude with a discussion of the implications of this theory for game design and the larger CC context.

Game Design Practices

In a standard text, Salen and Zimmerman (2004, p. 168) introduce the “second-order” problem of game design bluntly:

“The goal of game design is meaningful play, but play is something that emerges from the functioning of the rules. As a game designer, you can never directly design play. You can only design the rules that give rise to it. Game designers create experience, but only indirectly.”

Play includes the objective choices made by a player and the conditions achieved in the game, along with the play-

er's subjective reactions and expectations. At this point, it is straight forward to adopt the first tenet of our theory of creative game design: *game designers are really designers of play*.

The idea of adopting an iterative, "playcentric" (Fullerton 2008) design process, in which games are continually tested to better understand their emergent (play) properties, is corroborated by others like Schell (2008), who further describes the supreme importance of "listening" in the design process (being able to process feedback from the player's experience of candidate designs). Beyond initial conceptualization of a game idea and tuning and polish of the final product, the two most important practices of game design are *prototyping* and *playtesting*, both of which are intentionally focused on providing the designer with a better understanding of play.

Prototypes are playable artifacts, working models of a game idea that permit asking and answering questions about how a game will interact with its environment without requiring the effort to create a complete, polished game. The aesthetics of prototypes are very different than for complete games: most artwork and sound is stripped away from a design idea to produce an artifact that most effectively elicits feedback on a designer's current focus of interest (often how the interaction of game mechanics affects the trajectory of play).

Prototypes must be set to interact with an audience to gain the answers they are designed to provide. Playtesting is the practice of playing a game or gameplay prototype while observing the choices made, actions taken, or reactions expressed by sample players (the designer, a friend, a dedicated tester, or even a member of the target audience). Observations made during playtesting can reveal objective properties of a game such as unwritten-but-implied mechanics, exploits, and alternative puzzle solutions, or subjective properties such as the level of engagement, fun, or hesitation expressed by the sample players (Smith, Nelson, and Mateas 2009).

Despite the ostensible purpose of game design being the production of complete, desirable games for play by end-users, the practices of playtesting and prototyping are centered on providing feedback to the designer. Through the rough generate-and-test process of iterative game design, where several prototypes are created and playtested during a single project, the underlying goal is to build up sufficient skill and understanding to later produce the high-quality, final game artifact. Such a self-affecting process is exactly what McGraw and Hofstadter call the "central loop of creativity" (1993).

Beneath the surface, the practices of game design are almost exclusively about the collection of design knowledge, knowledge regarding the relationship between the component elements of a game system and that game's potential execution in interaction with a player. Such design knowledge spans what design patterns to employ, how to assemble them, and why such an assembly will produce a certain play experience.

Existing design knowledge can be applied to realize familiar, well-understood play experiences, but creative game design demands a continuous source of new design knowledge. Thus, the second tenet of our theory is this: *creative game design is about seeking design knowledge*.

An Analogy with Science

To expand on knowledge-seeking in game design, we want to draw an extended analogy between game design and science. Doing so will allow us to connect the creative activity in the game design process with the activity carried out by scientific discovery systems in CC.

For design generally, Dasgupta claims "design problem solving is a special instance of (and is indistinguishable from) the process of scientific discovery" (1991, p. 353). While Dasgupta focuses on explaining design activity specifically in terms of finding a confirmable theory which resolves a particular unexplained phenomena, our analogy is intentionally softer, to enable applications in a variety of CC domains that do not immediately appear as "design problem solving" domains.

Scientific Practices

Roughly, the scientific method is a closed loop with the following phases: A hypothesis is generated from a working theory, the hypothesis drives the design of an experiment (usually realized with a physical apparatus), data from executing this experiment is collected, and conclusions are drawn which can be integrated into the working theory. A scientist will design an experimental setup, despite already possessing a theory which makes prediction about the situation, precisely because there is some uncertainty about the result. This result, whether matching the prediction or not, should provide informative detail about the natural laws at play in the experiment's environment.

In our analogy, experiment design is prototyping; experiment execution and subsequent analysis is playtesting. The combination of the declarative knowledge of natural laws and the procedural knowledge of operating laboratory equipment is game design knowledge. Finally, the closed loop of the overall scientific method corresponds to iterative game design.

Making the parallel clearer, Gingold and Hecker (2006) talk about how gameplay prototypes should be informative, answer specific questions, and be falsifiable. In the philosophy of science, the notion of informative content (and its relation to falsifiability) guides the evaluation of theories and experimental designs.

In their capacity to design and execute informative experiments and produce coherent and illuminating explanations of the anomalous results, scientists are clearly creative. In Colton's terms (2008), we can easily *perceive* this creativity in the *skill* of precise experimental design, the *appreciation* of unexpected result in the context of a working theory, and the *imagination* of previously difficult to consider alternative theories and the invention of new instruments. By the analogy above, a scientist's kind of creativity can apply to the game designer as well, prompting

our third tenet: *designers act as explorers in a science of play*, an “artificial science” in Simon’s terms (1996, Ch. 5).

Automated Discovery

Though largely distinct from artifact generation, automating discovery in science and mathematics is an established CC tradition (Langley et al. 1984; Lenat 1976). Within these systems it is common to find subprocesses which generate artifacts as part of the larger discovery process.

The GT system (Epstein 1988), an automated graph theorist, would periodically “doodle” random graphs within a specific design space as a means of generating relevant data which might spark a new conjecture about the desired area of focus. With our analogy in mind, such doodling is reminiscent of the exploratory, rapid prototyping process sometimes used in game design (in what Gingold and Hecker call the “discovery” phase of development). A heuristic to generate new concrete examples of abstract concepts was even present in the original automated mathematician, AM, working with number theory (Lenat 1976).

Beyond generating artifacts with only the indirect intention of knowledge gain, more recent discovery systems internally optimize expected knowledge gain when deciding which experimental setup to test next, realizing an active learning process (Bryant et al. 2001). Where artifact generation provides opportunistic benefits in mathematical domains (in which graphs and conjectures are non-interactive, static artifacts), discovery systems working in the physical sciences fundamentally cannot avoid artifact generation (as experimental design) during active exploration of their domain.

A notion of “interestingness” is the glue that binds the various subprocesses of automated discovery (artifact generation included) together into an overall control flow (Colton and Bundy 2000). In many cases, interestingness measures the likelihood or quality of knowledge expected to be discovered by taking a particular action (e.g. searching for a counterexample) or focusing on a particular concept (e.g. looking for new examples of special graphs). A system’s overall notion of interestingness can be used to induce a measure of value for artifacts generated in its artifact generation subprocesses, a measure related to potential for knowledge gain as opposed to aesthetic value.

Returning to game design, our fourth tenet holds that *automated discovery systems inspire a computational model of creative game design* that explains the prototypes produced in exploratory game design as the doodles produced trying to flesh out design theories residing in the designer’s head, motivated by an *interest* in designs that have the potential to reveal new patterns.

Newell’s Knowledge Level

To complete the image of the creative game designer as a discoverer, we need some better vocabulary for talking about a designer’s knowledge, around which the entire discovery process revolves.

Newell (1982) describes the “knowledge level” as a systems level set above the symbolic, program level. At the

knowledge level, we find *agents* with bodies of *knowledge* that can take *actions* in some environment to make progress towards *goals*. The actions taken by an agent are said to be governed by a principle of rationality that states “If an agent has knowledge that one of its actions will lead to one of its goals, then the agent will select that action.” This sense of rationality is distinct from decision theoretic rationality in that it does not necessarily imply that an agent must optimize anything.

While this radically underspecifies an agent’s behavior from a computational perspective, the constellation of concepts at the knowledge level is useful for making statements about our game designer. The intention of knowledge-level modeling is to explain the behavior of knowledge-bearing agents (be they human or machine) without reference to how that knowledge is represented or access to an operational model of the agent’s mode of processing.

Understanding the game designer, at the knowledge level, starts with making an assumption about what is known and what is sought. But from our theory so far, we can safely assume an important body of knowledge possessed is that of tentative game design knowledge. This knowledge permits the designer the use of tools such as paper and trinkets for physical gameplay prototypes (often styled after board games) and programming languages and compilers for more detailed, computational prototypes. This same knowledge permits understanding to be gained from the observation of game artifacts in play, and suggests a tentative vocabulary for composition of those artifacts (i.e. knowledge of design patterns for game rules). The creative designer’s goal, per our analogy with science, is clearly to gain more design knowledge.

Given this, we expect the designer to rationally (specifically in the knowledge level sense) go about the practices of game design as part of taking actions that lead towards the gain of design knowledge. That is, *game design activity can be explained as the rational pursuit of design knowledge gain*.

Creativity as Rational Curiosity

The knowledge level lets us talk about a kind of rationality, one that gives an explanation for why game designers take the actions they do. But not all game design activity is creative, no more than all of science being creative. So where does creativity come in?

The most creative parts of game design, we claim, are the ones where the designer’s behavior is best explained by the direct intention to gain new knowledge, to satisfy *curiosity*. The bulk efforts of game production are a kind of engineering which applies the knowledge gained in the curiosity-driven creative mode.

As the motivation to reduce uncertainty and explore novel stimuli (Berlyne 1960), curiosity has long been known to be intertwined with the judgment of aesthetics (Berlyne 1971). Saunders’ “curious design agents” (2002) generate aesthetic artifacts according to their potential to satisfy an internal measure of curiosity, doing so *in order to learn* about an outside environment. This framework has

also been used to drive the behavior of a simulated society of curious visual artists and even a flock of curious sheep in a virtual world (Merrick and Maher 2009).

How curiosity-driven behavior can explain the various processes and artifacts of human creativity at a high level has been demonstrated at great length (Schmidhuber 2010). Our unique claim, that *creativity is a knowledge level phenomenon*, gains similar explanatory power without reference to algorithmic details (such as the use of reinforcement learning or optimization procedures) or human psychology, as in Loewenstein's comprehensive review of research on human curiosity (1994).

Looking concretely at curiosity in the domain of game design, consider the example of speed runs, gameplay traces that demonstrate a way of doing something in a game (completing a level or collecting certain items) much faster than a designer previously expected. Speed runs are interesting to game designers, from a curiosity perspective, because they often represent a novel stimulus and quickly increase uncertainty about what is possible in a game, creating an urge to seek out related gameplay traces that would illustrate the general pattern by which the run was achieved. With additional experience, the designer can learn to either design-out such speed runs by adjusting the game's rules, or create new mechanics that reward them.

Putting together curiosity about design knowledge with knowledge-level rationality, we have our complete theory of creativity in game design:

Creativity is the rational pursuit of curiosity.

This claim applies to human and machine design agents and gives a goal-oriented explanation to sequences of design activity that result in design knowledge gain (clearly including prototyping and playtesting). A creative game designer makes games, not because that is their function, but because they want to learn things about play that require experimentation with certain artifacts to illuminate.

We call this theory *rational curiosity* because it is a knowledge-level treatment of the concept of curiosity that focuses on how curiosity explains the selection of actions towards a known goal. It claims that curiosity, applied rationally, will result in behavior recognizable as creative design activities.

Transformational Creativity in Game Design

Consider our model creative game designer, over time, producing increasingly complex and refined playable artifacts that are in line with the complexity of their currently operating design theory. That playable artifacts are produced is just an externally visible byproduct of the more interesting process going on in the designer's thoughts, the growth and refinement of design knowledge.

Taking a snapshot at any one time, the designer's knowledge is fixed. The present knowledge describes a "conceptual space", in Boden's terms, of game designs and play possibilities. *Combinational* creativity within this pace would entail the generation of artifacts from known structures and construction constraints or, perhaps, the enumeration of explanations of a player's behavior with respect

to known patterns. Taking a series of steps in terms of the current design theory, producing a new game using a design pattern of interest, producing a prediction of player behavior, and then performing a playtest and comparing the results with the prediction is an example of *exploratory* creativity in this space. These activities are weakly creative in the rational curiosity view because, though they might indeed be motivated by potential knowledge gain, neither realizes an actual change to the designer's personal theory.

Transformational creativity in game design, then, is design activity which results in a redefinition of design theories. In iterative game design, where many prototypes are produced in succession in response to feedback from playtesting, is an intensely transformative process. Such transformations can include the definition of a new design pattern which simplifies the explanation of how another designer's game was constructed, a constraint which limits the use of two patterns together, or a rule which predicts a certain kind of player's behavior when a certain combination of pattern is present.

Discussion

Having proposed a theoretical explanation of creativity in game design, let's look at what it entails.

Computational Creativity in Game Design

The theory of rational curiosity in game design can be realized computationally along two major paths: the development of a game design discovery system, and new, knowledge-oriented creativity support tools. More generally, however, it suggests new elements that need to be modeled computationally in support of either path.

Game Design Discovery Systems Recalling the analogy with scientific and mathematical discovery systems, we can imagine the design of a new kind of discovery system that would work in the domain of game design knowledge. This discovery system would produce games as part of its experiments in exploring play, but it would also allocate significant attention to decomposing games made by other designers and producing explanations of observed human player actions.

The notion of interestingness in this discovery system would correspond to a symbol-level realization of the agent's knowledge-level goal: the satisfaction of curiosity about design knowledge. By selecting actions (such as the construction of a prototype, the simulation of a playtest using a known player model, or the analysis of a previously produced game in light of a refined theory) according to their calculated prospects for improvement of the working design theory (a library of design patterns, predictive rules for player behavior, and constraints on play-model construction), the system's behavior would implement a rational pursuit of curiosity, creativity in game design.

Constructing such a system would require new research into adapting symbol-level *representations of design knowledge* for use in game design, the development of a *task decomposition* of creative game design into subgoals and

actions (such a concrete design methodology which would be of interest to human designers as well), and the identification of the relevant *external tools* of game design (certainly paper prototyping materials and programming environments are some of these tools, but where are the CAD systems for games?).

Knowledge-oriented Creativity Support Tools

With recognition that design knowledge gain is the designer's goal, creativity support tools in game design should focus on easing this process. In terms of Yeap's desiderata for creativity support tools, these tools should focus on *ideation* and *empowerment* (2010). In game design, these translate to the generation of candidate design knowledge for the designer to consider and then leaving the adoption of the new knowledge up to the designer (without undue interference). Knowledge-oriented creativity support tools should attempt to *remove bottlenecks in the discovery process*.

We created a gameplay pattern language and corresponding search tool which is intended to accelerate the extraction of feedback about game designs from pre-recorded traces of play (Smith and Mateas 2011). The system is also capable of compiling patterns into a lower-level form that can be used to search for additional evidence of gameplay patterns with the machine playtesting tools included in the BIPED early-stage computational prototyping tool (Smith, Nelson and Mateas 2009). Neither of these systems are themselves creative, but they are designed to provide new actions to the creative designer for rational selection in the service of knowledge gain.

In another project (Smith and Mateas 2010), we captured a design space of mini-games in a logic program and used model-finding techniques to automatically generate artifacts from the described space. Use of the logic programming tools automates a small slice of game design (the literal construction of artifacts) and provides a convenient symbolic representation for some types of design knowledge. While this project automates combinational creativity in game design, the design space representation and sampling tools are intended to support an external designer's transformational creativity in the realization of new forms of gameplay through rapid exploration and expressive redefinition of the mini-game design space.

New Perspective for Computational Creativity

We have mostly focused on game design, but rational curiosity is intentionally worded so as to apply to other CC domains. In fact, it should apply even to domains with apparently non-interactive artifacts. Where, in game design, we were concerned with the implications of game rule system on player actions and reactions, in the domain of music we should explore the implications for sound patterns on audience anticipation and mood, in visual art the implications for perceptual details on where the viewer's eye lingers or flees, and in sculpture the implications of geometric arrangements on audience interest from particular viewpoints. Such domains are not as interactive as game design, but they could be equally deep in the subtlety of

how an audience reacts to an artifact – depth enough to keep the rationally curious artist busy producing experiments for quite some time.

The knowledge-level analysis of creativity suggests new questions to ask of CC systems: What does this system want to learn? How is knowledge represented in this domain? Is the system experimenting with the affordances of the raw medium or focusing on audience reactions achievable through it? (Both are equally creative assuming the desired kind of knowledge is gained.)

Consider *NEvAr* (Machado and Cardoso 2002), a creative system in the domain of visual art. Unlike the straightforward interactive genetic algorithm in *PicBreeder* (Secretan et al. 2008), *NEvAr* does not ask its audience for feedback on every artifact it internally considers. The system summarizes sparse feedback from its human audience in the form of a neural network which becomes a proxy for their ratings in an internal evolutionary process. Rational curiosity would describe *NEvAr* as a creative system, not because it produces novel and valuable images, but because, at the knowledge level, the system appears to be rationally soliciting reactions on believed high-valued images and incorporating the responses in a way that transforms the space of images that the system will next produce—it behaves consistently with the explanation that it is rationally pursuing its curiosity (albeit with limited design knowledge storage capabilities). If redesigned from scratch with rational curiosity in mind, the system might incorporate a more interpretable representation of learned knowledge (for which it is easier to read as a design theory that improves over time) and put more computation into experimental design, reasoning over the expected knowledge gain from enticing the human audience to provide feedback on a particular work rather than always trying to display the estimated-best available artifacts. Orienting the system around active learning, we predict, would improve the system's apparent creativity.

From our perspective, it is natural to ask what a system learns as it runs. Though established techniques in computational visual art such as design grammars and iterated function systems can, in some cases, produce very interesting (valuable) images, the static nature of these techniques in isolation implies that, over time, our sense of novelty of the kinds of artifacts these techniques produce will necessarily wane because these techniques do not learn. While rational curiosity would deem a technique that merely samples a fixed generative space uncreative, these techniques are still valuable to us—they encode very rich design spaces that, upon gaining experience through experimentation, a creative agent can alter as part of large-scale experiments in the design of these generative spaces.

Conclusion

We have followed the clues embedded in the practices of human game designers to a set of building blocks for a theory of creative game design. To recap:

- 1) Game designers are really designers of play.
- 2) Creative game design is about seeking knowledge.

- 3) Designers act as explorers in a science of play.
- 4) Automated discovery systems inspire a computational model of creative design.
- 5) Game design activity can be explained as the rational pursuit of design knowledge gain.

This has led us to a new statement about creativity that can apply to human or machine design agents in any artifact generation domain: **creativity is the rational pursuit of curiosity**, a knowledge-level phenomenon.

We hope this model of creativity will inspire the exploration of discovery system architectures for artifact generation systems and the development of a new space of knowledge-oriented creativity support tools.

Acknowledgements

This work was supported in part by the National Science Foundation, grant IIS-1048385. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Berlyne, D. E. 1960. *Conflict, arousal and curiosity*, McGraw-Hill.
- Berlyne, D. E. 1971. *Aesthetics and psychobiology*, Appleton-Century-Crofts.
- Boden, M. A. 2004. *The creative mind: myths and mechanism*, 2nd ed., Routledge.
- Bryant, C. H.; Muggleton, S. H.; Oliver, S. G.; Kell, D. B.; Reiser, R. D.; and King, R. D. 2001. Combining inductive logic programming, active learning and robotics to discover the function of genes. *Electronic Transactions on Artificial Intelligence* 5(B1):1-36.
- Colton, S. and Bundy, A. 2000. On the notion of interestingness in automated mathematical discovery. *International Journal of Human Computer Studies*.
- Colton, S. 2008. Creativity versus the perception of creativity in computational systems. *Creative Intelligent Systems: Papers from the AAI Spring Symposium*, 14-20.
- Cope, D. 2005. *Computer models of musical creativity*, MIT Press.
- Dasgupta, S. 1991. *Design theory and computer science*, Cambridge University Press.
- Epstein, S. L. 1988. Learning and discovery: one system's search for mathematical knowledge. *Computational Intelligence* 4(1):42-53.
- Fullerton, T. 2008. *Game design workshop: a playcentric approach to creating innovative games*, 2nd ed., Morgan Kaufmann.
- Gingold, C. and Hecker, C. 2006. *Advanced prototyping*. GDC 2006 lecture.
- Hartman, C. O. 1996. *Virtual muse: experiments in computer poetry*. Wesleyan University Press.
- Langley, P.; Bradshaw, G.; and Simon, H. A. 1984. Rediscovering chemistry with the bacon system. In *Machine Learning: An Artificial Intelligence Approach*, 307-329.
- Lenat, D. 1976. *AM: an artificial intelligence application to discovery in mathematics*. PhD Thesis.
- Loewenstein, G. 1994. The psychology of curiosity: a review and reinterpretation. In *Psychological Bulletin*, 116(1), 75-98.
- Machado, P. and Cardoso, A. 2002. All the truth about NEvAr. *Applied Intelligence* 16, 101-118.
- McGraw, G. and Hofstadter, D. 1993. Perception and creation of diverse alphabetic styles. *AISBQ*, (85):42-49.
- Merrick, K. and Maher, M. L. 2009. *Motivated reinforcement learning: curious characters for multiuser games*, Springer.
- Newell, A. 1982. The knowledge level. *Artificial Intelligence*, 18(1):87-127.
- Pease, A.; Winterstein, D.; and Colton, S. 2001. Evaluating machine creativity. In *Workshop of Creative Systems*.
- Salen, K. and Zimmerman, E. 2004. *Rules of play*.
- Saunders, R. 2002. *Curious design agents and artificial creativity*. PhD Thesis.
- Schell, J. 2008. *The art of game design: a book of lenses*, Morgan Kaufmann.
- Schmidhuber, J. Formal theory of creativity, fun, and intrinsic motivation (1990-2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230-247.
- Secretan, J.; Beato, N.; D'Ambrosio, D. B.; Rodrigues, A.; Campbell, A.; and Stanley, K. O. 2008. Picbreeder. In *SIGCHI Conf. on Human Factors and Computing Systems*.
- Simon, H. 1996. *Sciences of the Artificial*, 3rd ed., MIT Press.
- Smith, A. M.; and Mateas, M. 2010. Variations Forever: flexibly generating rulesets from a sculptable design space of mini-games. In *Proc. of the 2010 IEEE Conf. on Computational Intelligence and Games (CIG 2010)*.
- Smith, A. M.; and Mateas, M. 2011. Towards knowledge-oriented creativity support in game design. In *Proc. of the 2nd International Conference on Computational Creativity (ICCC 2011)*.
- Smith, A. M.; Nelson, M. J.; and Mateas, M. 2009. Computational support for play testing game sketches. In *Proc. of the 5th Annual AI and Interactive Digital Entertainment Conference (AIIDE2009)*.
- Yeap, Wai K.; Opas, Tommi; and Mahyar, Narges. 2010. On two desiderata for creativity support tools. In *Proc. of the Intl. Conference on Computational Creativity*, 180-189.