Boosting Computational Creativity with Human Interaction in Mixed-Initiative Co-Creation Tasks

Antonios Liapis and Georgios N. Yannakakis
Institute of Digital Games, University of Malta, Msida, Malta
{antonios.liapis@um.edu.mt, georgios.yannakakis@um.edu.mt}

Abstract

Research in computational creativity often focuses on autonomously creative systems, which incorporate creative processes and result in creative outcomes. However, the integration of artificially intelligent processes in human-computer interaction tools necessitates that we identify how computational creativity can be shaped and ultimately enhanced by human intervention. This paper attempts to connect mixed-initiative design with established theories of computational creativity, and adapt the latter to accommodate a human initiative impacting computationally creative processes and outcomes. Several case studies of mixed-initiative tools for design and play are used to corroborate the arguments in this paper.

Introduction

For over two decades, the study of computational creativity has focused on "building software that exhibits behavior that would be deemed creative in humans" (Colton, De Mantaras, and Stock 2009). It is not surprising, therefore, that the grand challenges which are stressed in this research domain focus on fully automated systems "which learn to do what they do, before attempting to do it creatively" (Cardoso, Veale, and Wiggins 2009). Learning to be creative, for a computer system, can be achieved from large corpus of human-authored data — such as Wikipedia articles (Barros, Liapis, and Togelius 2015) or search engine queries (Veale 2014) — or from observing their own previous experiences in exploring the space of possible outcomes (Liapis et al. 2013; Correia et al. 2013). Human-based creative artifacts such as news articles can seed the creativity of the machine which transforms them into collages (Krzeczkowska et al. 2010) or playable games (Cook and Colton 2012). In such cases, human creativity initializes the search for creative outcomes (either by providing an initial seed for search, or by affecting the evaluation of creative outcomes), but does not actively affect the systems' exploration while that occurs.

However, computational creativity need not be entirely autonomous; systems which rely on interaction with a user during their creative process can still possess and express creativity. Moreover, this paper argues that software can foster and enhance their computational creativity potential through interactions with human users. The case is made for computer-aided design tools where the role of the software is not merely "the designer's slave" (Reintjes 1991), but is a proactive co-creator which actively contributes to the design process. This paper uses the term *mixed-initiative design* tools to differentiate such software, pointing to a design dialogue where both the human and the computational creator exhibit an initiative to the creative discourse (compared to software which merely reacts to a human command by e.g. performing simulations or constraint satisfaction tests). *Initiative* is traditionally considered under the prism of a dialogue between man and machine (Novick and Sutton 1997), and can refer to the *task initiative* (who decides the topic of the conversation), the *speaker initiative* (who decides when each actor speaks) or the *outcome initiative* (who decides when the problem has been solved). Previous work by the authors has argued that mixed-initiative design tools are capable of fostering the creativity of their human users, by disrupting both their creative processes and their aesthetic criteria (Yannakakis, Liapis, and Alexopoulos 2014). This paper, instead, focuses on how computational creativity is affected by prolonged interaction with creative human users.

Admittedly, the definition and distinction of mixed-initiative design processes (and their distinction from other forms of computer-aided design) is not clear-cut (Novick and Sutton 1997). Moreover, several other terms have been used to describe similar co-creative processes, including human-computer creativity (Kantosalo et al. 2014), AI-assisted design, or casual creators (Compton and Mateas 2015) for more playful design work. We follow the terminology used in previous papers, identifying mixed-initiative co-creation (MI-CC) "as the task of creating artifacts via the interaction of a human initiative and a computational initiative" (Yannakakis, Liapis, and Alexopoulos 2014). This distinguishes MI-CC from collaboration between humans (no computational initiative) and from tools with no proactive role (e.g. spell-checkers). We focus on the final mixed-initiative tool as software, rather on the priorities and design decisions that went into its design — a topic covered by Kantosalo et al. (2014). Moreover, we focus on the interaction between software and a human end-user (e.g. a player in a creation game, a designer in a task-driven game development task, etc.) rather than on the interaction between software and its developer (as the latter could identify bugs in the system, or directly affect it via e.g. code changes).

The paper starts by connecting mixed-initiative co-creation with some of the most prevalent theories of computational creativity, identifying which aspects of the creative process can be influenced by human interaction. Following this, several examples of mixed-initiative interaction tools (for different purposes and with different degrees of computational initiative) are analyzed in light of these theories.

Exploratory Creativity and MI-CC

In an attempt to formalize the model of creativity introduced by Boden (1992), Wiggins (2006) represents an exploratory creative system as a septuple: $\langle \mathcal{U}, \mathcal{L}, [[.]], \langle\langle ., ., . \rangle\rangle, \mathcal{R}, \mathcal{T}, \mathcal{E} \rangle$. $\mathcal{U}$ represents all possible outcomes of the creative system, $\mathcal{R}$ are rules (in language $\mathcal{L}$) of membership in a (target) conceptual space, $\mathcal{T}$ are the rules of traversal (search) of this space and $\mathcal{E}$ are the rules for evaluating the outcomes. The conceptual space follows the terminology of Boden (1992), acting as the mental representation of what a possible and appropriate outcome is, with regards to the current context (e.g. valid chess moves, jazz melodies etc.). $[[.]]$ is a function generator which maps elements of $\mathcal{U}$ to a real number, while $\langle\langle ., ., . \rangle\rangle$ is a function generator which creates new elements of $\mathcal{U}$ from existing elements of $\mathcal{U}$ (using $\mathcal{R}$,$\mathcal{T}$,$\mathcal{E}$).

Human creativity introduced via mixed-initiative interaction can influence several elements of Wiggin's septuple. Human initiative acting as $\mathcal{E}$ can take the form of direct evaluation of elements in $\mathcal{U}$, as is often the case in interactive evolutionary systems such as *PicBreeder* (Secretan et al. 2011) and *MaestroGenesis* (Hoover et al. 2012) where users select favorite outcomes to evolve or rate outcomes in terms of preference. Another option for human initiative acting as $\mathcal{E}$ is the possibility of customizing the evaluation method. For instance one may choose which fitness dimensions to use via an interface in an aggregated or multi-objective fitness function. More ambitiously, the impact of such fitness dimensions can be learned indirectly from human choices via user modeling (Liapis, Yannakakis, and Togelius 2013a). User modeling allows the computationally creative system to adapt its own $\mathcal{E}$ to match that of the human user without replacing it entirely with a human-authored one.

Human initiative can also influence $\mathcal{T}$, by specifying algorithmic parameters such as mutation rate (as in *PicBreeder*), or $\mathcal{R}$ by narrowing the conceptual space of the generator to only include e.g. jazz melodies of less than 1 minute. Changing $\mathcal{T}$ in terms of genetic operators and parameters requires direct human intervention (replacing the system's $\mathcal{T}$ with a human-authored one), and a degree of technical knowledge that is closer to that of the system's developer than that of its end user. As this paper focuses on interaction with end-users, such changes in $\mathcal{T}$ are out of scope. However, $\mathcal{T}$ can be indirectly affected by e.g. setting the starting point for exploration by seeding the initial population with human creations, as will be discussed in the examples. Similarly changes in $\mathcal{R}$ can be made indirectly if the system learns from, or is forced to follow, user creations which have those desirable properties. As the conceptual space is introduced as much by the machine (which limits what the user can or cannot do) as the human user (who narrows down the machine's conceptual space to their own frame of reference), the computational creator must identify and respect the boundaries ($\mathcal{R}$) of the human user's conceptual space.

It should be noted here that transformational creativity in mixed-initiative interaction can occur if the machine, through its own initiative and suggestions to the human user, manages to change the boundaries of the human user's conceptual space (without necessarily changing its superset, i.e. the machine's conceptual space). Treating human-machine as a single entity under the prism of the extended mind (Clark 1998) or as a 'symbiote' (Licklider 1960), transformation occurs when the human user's frame of reference (Scaltsas and Alexopoulos 2013) is disrupted and a designer's fixation is challenged, thus resulting in the transformation of that user's design/interaction process.

Quality, Novelty and Typicality in MI-CC

In order to be able to attribute creativity to a computer program, Ritchie (2007) proposed several criteria for a domain to be considered creative, as well for the artifacts within it. Since human interaction does not affect the domain itself, it is worthwhile to observe how the criteria of Ritchie regarding the resulting artifacts of a process must be reconsidered if the creation process is not purely computational but involves human interaction throughout. Ritchie (2007) identifies three essential properties of the final results of a process "for deciding whether creativity has occurred":

- **Novelty** To what extent is the produced item dissimilar to existing examples of its genre?
- **Quality** To what extent is the produced item a high quality example of its genre?
- **Typicality** To what extent is the produced item an example of the artefact class in question?

While these criteria are fairly general and can be used for any artifact regardless of the process used to create it (computer-generated, human-authored, or anything in between), it is worthwhile to refine them in order to consider the human user and their interaction with the software. In that regard, when dealing with human end-users interacting with a mixed-initiative tool, the criteria of novelty, quality and typicality can be adapted as such:

Novelty To what extent is the produced item dissimilar from what is currently created by the human user?

Quality To what extent would the produced item be of use to end-users?

Typicality How does the produced item match the human user's frame?

Obviously, the proposed criteria include a human user (or a broader human audience) in their formulation. For novelty, it is assumed that the human user is creating something alongside, in parallel, or by taking turns with the computational creator. The artifact (partial or complete) produced by the creative software must be dissimilar to that of the human author, in order to act as a disruptor of the human designer's routine and fixations. On the other hand, typicality requires that the created artifact is still recognizable as a member of the user's conceptual space or *frame* (Scaltsas and Alexopoulos 2013); this ensures that the user will not discard the computational output as an error in the system (causing them to remain fixated on their current frame). It should be noted that the way novelty and typicality are currently framed, attaining both novelty and typicality requires pushing against the boundary of the user's conceptual space (novelty) while remaining 'close' to those boundaries (typicality) in order for the computational output to be recognized as a viable alternative to the human creation. Finally, quality in MI-CC is not particularly different from the original notion of Ritchie, as in both cases quality refers to human (subjective) evaluation. In MI-CC, one can argue that quality can be entirely subjective as the only 'audience' evaluating the output is the human user interacting with the system; however, in certain cases of mixed-initiative interaction (e.g. in multiplayer games), quality is assessed by a larger audience.

MI-CC and FACE

Inspired by notions of exploratory and transformational creativity (Boden 1992; Wiggins 2006), Colton, Charnley, and Pease (2011) put forth "a plausible way in which creation by software could occur" using the FACE and IDEA models. The FACE model describes "creative acts performed by software" while the IDEA model describes "the impact of creative acts performed by software". We focus on the FACE model in this paper, as the most clearly defined and the most likely to be affected by human interaction of the two. The FACE model is an acronym for the possible generative acts in a system: F stands for the framing information (i.e. a description of the generative acts), A stands for the aesthetic evaluation, C stands for the concept (i.e. an executable program which produces an expression from input) and E stands for expression (i.e. a single outcome of a particular input). Each of these generative acts can be a singular instance (denoted with a g), e.g. E^g is a single expression, or a method for generating instances of this type (denoted with a p), e.g. E^p is a method for generating expressions.

Similarly to the septuple of Wiggins (2006), human initiative in a mixed-initiative tool often performs evaluation of the output, acting as A in the FACE model. The human user can replace the computational aesthetics of the system (acting as A^g); however, the system can also learn the preferences of the human user via user modeling (Liapis, Yannakakis, and Togelius 2013a) in which case the human acts as a meta-evaluator guiding the search of a computational process which generates aesthetic evaluations (A^p).

Obviously, the human user can also provide the framing information (F), identifying and rationalizing the intelligent (or less so) processes of the software. However, this is not interesting from a creativity perspective — especially regarding the creativity of the software. However, a system able to describe its own framing information to the user can avoid its output from being considered erroneous by the human evaluator, especially in case they are pushing against the user's frame. Such framing information helps the end-users of a mixed-initiative tool *perceive* the creativity of the software (Colton 2008), and draws attention to it.

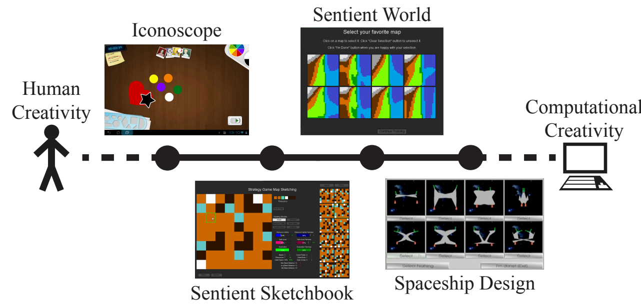

Figure 1: The MI-CC examples in this paper explore the spectrum between human-led and computer-led creativity.

Instances of MI-CC

In order to demonstrate the potential of mixed-initiative co-creation in enhancing both computational and human creativity, we examine software that realizes MI-CC. The four tools and games outlined in this section have been developed over the last few years in our attempt to further explore the capacity of mixed-initiative processes for co-creativity. These examples range from a predominant human initiative with optional computational suggestions (in Iconoscope and Sentient Sketchbook) to a computer-driven creative task guided by — and learning from — a human user (in Sentient World and the Spaceship Design interface). Unlike Iconoscope, in Sentient Sketchbook the suggestions are created without the express request of the user (thus exhibiting more computational initiative). On the other hand, Sentient World requires a human to provide the initial creative input, unlike the Spaceship Design interface (thus requiring more human initiative). Figure 1 ranks each initiative's contribution in the examined cases. The description of the tools themselves remains at a high-level in this paper, as the focus of this study is on how computational creativity is affected by interacting with human creativity.

Iconoscope

Iconoscope is a game designed for use in classrooms, in order to prompt creative thinking in young learners (Liapis et al. 2015). Along with other similar games, it has been developed for the purposes of the FP7 ICT funded project C2Learn. Iconoscope is a multi-player game, played on Android tablets by 4 or more players in the vicinity of each other: the goal is for each player to create an icon depicting a concept which confuses the other players. All players choose one among three concepts which are abstract themselves (e.g. "tolerance", "acceptance" and "solidarity") and attempt to depict it using simple shapes and colors (see Fig. 2). Once all players are finished, players vote for other players' icons by attempting to guess which of the three concepts is depicted. The most ambiguous icon (collecting as many incorrect as correct guesses) is the winner.

Computational creativity is an additional, optional module in the creative process of the user. While the user is drawing their icon by adding, moving, rotating, scaling or re-coloring shapes, they can select computational assistants which provide suggestions for alternatives to the user's

Figure 2: A player in Iconoscope is drawing an icon for the concept “solidarity” (attempting to confuse other players in guessing “tolerance” or “acceptance”). The portraits at the top of the screen are computational assistants which can provide up to four suggested alternatives to the user’s icon.

(current) icon. In Iconoscope, each computational assistant (C2Assisant) has a portrait, name and personality trait, and a different way of searching for alternative icons. There are five assistants, four of which perform evolutionary search: (a) *Chaotic Kate* performs random mutations to the user’s original icon, (b) *Mad Scientist* performs novelty search (Lehman and Stanley 2011), attempting to diverge from the user’s icon as well as from other icons in the same population, (c) *Typical Tom* attempts to evolve the user’s icon so that it more closely matches a human-authored archetype for this concept[1] stored in the game’s database, (d) *Progressive Petra* evolves the user’s icon to increase visual difference from the human-authored archetype, (e) *Wise Oracle* selects among previously user-created icons for the same concept and shows them as suggestions. While the Wise Oracle ensures the quality criterion of creative outcomes (via human-evaluated icons), we will be omitting it from further discussion as it does not generate its own outputs and is not affected by human interaction. Once the selected C2Assistant finishes its evolutionary search (which lasts a few generations), the four fittest individuals are shown to the user, who can select to replace their current icon with one of them or discard all of them and continue refining their own icon.

The computational processes of the C2Assistants are inherently tied to user interaction, since they *de facto* need to be initiated by the user selecting an assistant. More interestingly, however, the initial population of each assistant (ignoring the Wise Oracle) is seeded from the user’s current icon (i.e. all initial individuals are mutations of the user’s icon). By constraining exploration to start from a specific area of the search space, the human user indirectly provides the framing information (F^g in the FACE model) as

[1]Example archetypical icons include a red heart for the concept “love” or several green triangles for “forest”.

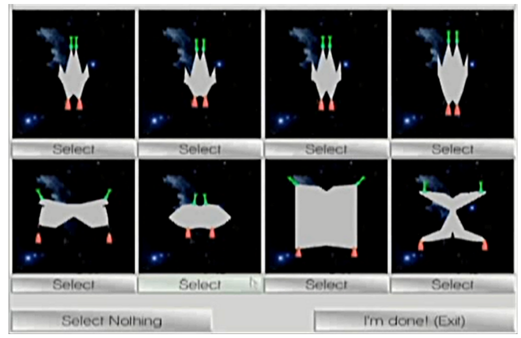

Figure 3: Interface for evolving spaceships. The aesthetic model learned that bottom-heavy spaceships are preferable.

the search can only discover nearby artifacts due to the few generations allocated for evolving suggestions. By specifying the area where search can take place, the user affects the traversal ($\mathcal{T}$ in Wiggins’ model). The user also directly specifies a traversal method when choosing a C2Assistant (e.g. random walk, novelty search), although arguably this amounts to initializing (not influencing) search parameters.

Spaceship Design

In earlier work, interactive evolution was enhanced with a model of aesthetics and used for creating 2D spaceships (Liapis, Yannakakis, and Togelius 2012). This mixed-initiative design tool allows human designers to finetune a spaceship design, i.e. its hull’s geometry, its weapons and its thrusters. The interaction paradigm is interactive evolution: a user chooses a single favorite among eight shown spaceships, sampled from an evolving population (see Figure 3). The creative process is enhanced via an aesthetic model acting as the fitness function which drives the search before presenting the next batch of spaceships to the user. The aesthetic model combines, in a weighted sum, ten different fitness dimensions of visual quality; inspired by cognitive psychology theory (Arnheim 2004), the balance (concentration of mass, symmetry) and shape (perimeter, jaggedness, size) of spaceships is evaluated. The weights of the model are adapted from the human user’s choices, increasing the weight of features in the chosen spaceship which are missing in the unselected spaceships. Through this process the user refines their preferences, which are used to evolve new content that the user is likely to find appealing without the need to constantly evaluate every individual in every generation. The model is also used to choose which spaceships in the population are shown to the user: e.g. in (Liapis, Yannakakis, and Togelius 2012) shown spaceships range from best to worst based on the aesthetic model (and an even distribution in-between).

The mixed-initiative tool for spaceship generation has the strongest computational initiative: the computer initializes the population and affects what is shown to the user at the start of co-creation, which can affect the user’s ideas for a spaceship and the progress of search. In that regard, the computational initiative unilaterally determines the rules and traversal of the space ($\mathcal{R}$ and $\mathcal{T}$ in Wiggins’ model) as well as the notion of typicality (through constraints on feasible spaceships). Human initiative only *indirectly* affects the aesthetic model used to evaluate the spaceships based on the user's selections; as they make no explicit choices in terms of aesthetic labels, users often fail to notice differences in balance and shape between spaceships and are thus surprised by the evolved outcomes. In this case, the creative system includes a generator of aesthetic evaluations (A^p in the FACE model), which uses the user's interaction data to refine the evaluation until it corresponds to users' choices.

Sentient World

Sentient World is a mixed-initiative tool which allows a user to define the terrain of a gameworld (Liapis, Yannakakis, and Togelius 2013c). In order to do so, the user begins by operating on a very low resolution of the terrain (i.e. 3 by 3 tiles) and define very high-level terrain properties (i.e. whether it is land or water). The software returns several higher resolution versions of this map with more details (i.e. water, plains, hills or mountains); the user can select their preferred one and (optionally) edit it further (see Fig. 4). At the end of this iterative process, the software can create a terrain of any resolution as the final outcome of the co-creation process.

Looking in more detail at the computational processes of Sentient World, the map is created from an artificial neural network (ANN). The height of each tile on the map is the output of the ANN where its input is the x, y coordinates of the tile's center; the canvas dimensions are normalized to [0,1]. As the ANN can return output for any coordinate pair at any numerical precision, Sentient World can create terrain at infinite resolution. This same property allows the software to increase the resolution of the user-created terrain. In order to create an ANN which conforms to the user's terrain, Sentient World performs backpropagation attempting to match the outputs of the ANN to the height ranges of the user's terrain (e.g. water tiles have a height range of $[0, 0.5]$). While backpropagation attempts to match the user-specified points in the low-resolution sketch, it has free reign on points between those specified by a user, as well as on the exact height of each point (thus land can be turned into plains, hills or mountains). In order to enhance the expressivity of the ANNs and in order to create more interesting alternatives to the user's map, a brief sprint of neuroevolution is applied before backpropagation. Using neuroevolution of augmenting topologies (Stanley and Miikkulainen 2002), the ANNs can increase the number of their nodes and connections and thus become able to capture more elaborate patterns. Moreover, neuroevolution is carried out by novelty search (Lehman and Stanley 2011) which rewards those low-resolution maps which possess different tiles from others in the same population; this results in backpropagation starting from different starting points (in terms of both ANN topology and weight values) and thus is more likely to result in visually different maps which still conform to the high-level patterns provided by the user (see Fig. 4b).

Sentient World is very much a mixed-initiative tool, as human and computational initiatives take turns contributing to the design: the human provides the high-level patterns and curates (via selection and minor edits) the computational output, while the computer attempts to diversify its results while conforming to the user's high-level guidelines. As in many of the mixed-initiative tools in this paper, the user specifies the aesthetics (A^g in the FACE model) by providing the initial low-resolution terrain and by curating the software's outputs. However, the software has leeway in specifying the aesthetic details of the higher-resolution terrain that it generates (A^g in the FACE model) while still conforming — in a soft manner — to the user's aesthetics. It should be noted that in Sentient World both user and computer attempt to refine an expression by observing it at different (and progressively higher) levels of detail: in a sense, the design dialog between man and machine acts as a generator of concepts which correspond to expressions at different levels of detail (C^p in the FACE model) although this generation ultimately settles into one final terrain (E^g in the FACE model). Regarding the properties of the final terrain, the human user asserts that the outcome is of high quality either directly (via curation) or indirectly (by providing the high-level terrain which the generator attempts to match); on the other hand, novelty is specified by the computer without human intervention (during neuroevolution) but then indirectly controlled by the user during backpropagation which must match the user's guidelines. Using the vocabulary of Wiggins to describe creativity in Sentient World, finally, evaluation ($\mathcal{E}$) is indirectly controlled by the human user as the system rewards conformity with human directives. More interestingly, however, one can argue that both the computer and the human user affect the traversal ($\mathcal{T}$) as the human specifies the high-level goal and the computer specifies how to reach that goal (i.e. during backpropagation). This is especially true when considering the stopping criteria for backpropagation, which are a maximum number of epochs, a small error value (i.e. all low-resolution tiles are correctly tagged water or land), or if the error does not decrease for several epochs. The patterns provided by the human user affect the performance of backpropagation (i.e. traversal of the space of possible outcomes) and may, in cases of extremely difficult patterns, result in high-resolution terrain which do not have all — or any — of the features specified by the user.

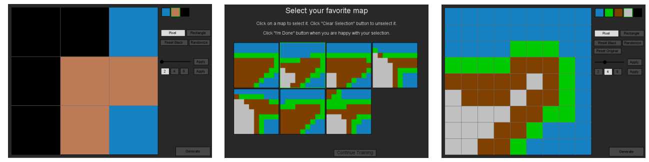

(a) 1st phase of user drawing (b) 1st phase of system's suggestions (c) 2nd phase of user editing

Figure 4: Sentient World interface and design process: the user starts by drawing a rough sketch (Fig. 4a) which the computer refines, offering alternative refinements (Fig. 4b) which the user can edit further (Fig. 4c).

Sentient Sketchbook

Sentient Sketchbook is a mixed-initiative tool for the design of game levels, where several computational designers create their own alternatives to the user's level in real-

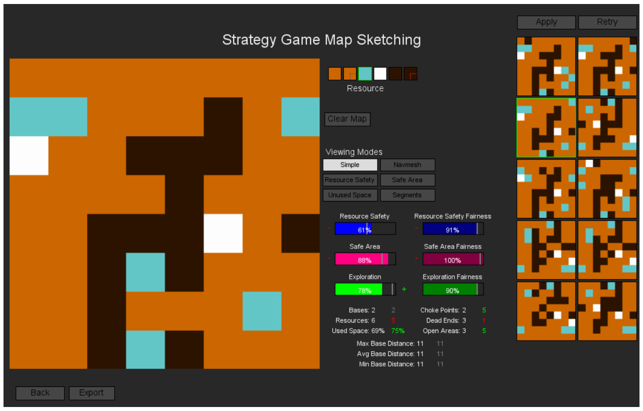

Figure 5: Sentient Sketchbook user interface: the user can choose the computer-generated suggestions (far right) to replace the map sketch they were drawing on the canvas (left).

time, presenting the results to the user who can choose to replace their own level with a computational suggestion (Liapis, Yannakakis, and Togelius 2013b). The interface of Sentient Sketchbook (see Fig. 5) operates on low-resolution, high-level abstractions of levels which contain only the absolute minimal details which define this game: in this case, a strategy game with impassable regions (dark), resources (cyan) and player bases (white). These low-resolution *map sketches* can be refined after the design process and can represent strategy games, dungeons, shooter levels (Liapis, Yannakakis, and Togelius 2013d), etc. Inspired by general game design patterns (Björk and Holopainen 2004) of *safety*, *exploration* and *balance*, six fitness dimensions have been identified for these map sketches which are usable across game genres. In real-time while the user draws their map sketch, each of the computational designers in Sentient Sketchbook uses one of these six fitness dimensions to evolve a population consisting of mutations of the user's current map sketch. The computational designer ensures that the resulting map sketches are playable by integrating playability constraints and evolving via a feasible-infeasible two-population genetic algorithm (Kimbrough et al. 2008). While most computational designers evolve towards a certain fitness dimension of map quality, one of them evolves to maximize the novelty of individuals via feasible-infeasible novelty search (Liapis, Yannakakis, and Togelius 2015).

Sentient Sketchbook has several similarities with Iconoscope, both in the interaction paradigm (optional suggestions) and in the computational creativity included (evolution of the user's creation). Therefore, the fact that the shown suggestions are evolved from the user's current map sketch means that the user's creation provides the framing information (F^g in the FACE model) and affects the traversal ($\mathcal{T}$ in Wiggins' model) of the computational designers. It should be noted that while selecting the suggestions in Sentient Sketchbook is optional, the suggestions are always generated and presented in real-time and do not require the user's request as in Iconoscope (where the C2Assistant must be clicked to create suggestions). This means that computational creativity occurs alongside human creativity at all times, and the human user determines when the *quality* (e.g. via objective-driven search) or *novelty* (e.g. via novelty search) of computational creators is appropriate to consider and make use of. This interaction paradigm, where the user has several alternatives to their own design to choose from is termed *mutant shopping* (Compton and Mateas 2015).

Of particular interest to the arguments in this paper, however, is the integration of designer modeling in Sentient Sketchbook (Liapis, Yannakakis, and Togelius 2014). Designer modeling refers to special cases of user modeling where interaction data between a designer and a computer-aided design tool are used to model the preferences, process, style and goals of the user (Liapis, Yannakakis, and Togelius 2013a). Three different types of designer models can be integrated in Sentient Sketchbook: (a) a model of the designer's style, learned from selected suggestions over a long period of interaction, (b) a model of the designer's process, derived by comparing the user's current level and comparing it to the previous state before the user's last action, and (c) a model of the designer's visual goals by assessing whether the user's level has certain symmetries which should also exist in the suggestions. In the case of the first two models, the model adjusts the weights of the fitness dimensions to better match the user's style or process, then evolves suggestions using the adjusted weighted sum. In the case of the model of visual goals, if symmetries are found then the mapping between genotype and phenotype forces that particular symmetry in the computational suggestions. We will focus on the first two models as they learn the style or process from the user (rather than applying rules for symmetry). When these designer models are in place, the computational initiative adapts its own aesthetics to match those indirectly specified by the user by generating a number of aesthetic models (A^p in the FACE model) and through gradient search choosing the one which best conforms to user choices. From the perspective of the created artifacts, the human indirectly specifies their *quality*: more accurately, the computer applies an interpretation of the user's notion of quality for its internal processes while the human user ultimately assesses the quality of the suggestion by selecting it or ignoring it.

Discussion

This paper identified the core aspects of computational creativity, in terms of process or outcomes, which can be boosted via the interaction with human users in mixed-initiative tools. Several MI-CC design tools and creation games were presented in order to highlight where and how the creative processes of the computational initiative were prompted, enhanced or facilitated from interaction with a human initiative. Design interfaces such as those used for spaceship generation can enhance the assessment of *quality* by observing human behavior and adapting to it indirectly — learning an aesthetic model which can be re-used in future creative tasks, be they autonomous or mixed-initiative. Design tools such as Sentient World can provide a user's *frame* and high-level goal for computational creativity to strive towards while retaining its own creative potential in interpreting this frame. In Sentient Sketchbook and Iconoscope, finally, human creativity provides the seed (as the human creation) for computational creativity, and inadvertently binds the area of the search space which the software can explore.

The paper attempted to posit mixed-initiative co-creativity (Yannakakis, Liapis, and Alexopoulos 2014) in the context of several computational creativity theories: the exploratory creativity of Wiggins (2006), the criteria of Ritchie (2007) and the FACE model of Colton, Charnley, and Pease (2011). As noted by Ventura (2008), however, "a great deal has been written about the nature of creativity in a computational setting and how we might characterize it or measure it or detect it or justify it". These three theoretical frameworks were chosen primarily due to the considerable attention they have received (also from the perspective of human-computer creative systems (Kantosalo et al. 2014)), but also due to the fact that they attempt to formalize (and isolate) the aspects of creative systems; thus it is easier to argue for specific aspects which are affected by human interaction. That said, there is an abundance of frameworks for creativity (Colton 2008; Jordanous 2012; Grace et al. 2014) which can also be connected with MI-CC; indeed, several such frameworks address the issue of computational creativity via human interaction (Bown 2014) and the principles of designing human-computer creative software (Kantosalo et al. 2014).

The examples covered in this paper were limited to software designed and developed by the authors, which target specific types of creative tasks: game or level design and creative play. We focus on these MI-CC tools because they largely share a design philosophy (multiple suggestions, evolutionary computation, quasi-real-time computational response) while also having several differences which highlight different ways in which computational creativity is boosted. It would be worthwhile, however, to examine mixed-initiative tools for purposes beyond games; examples include human-computer interfaces for generating poetry (Kantosalo, Toivanen, and Toivonen 2015), jokes (Ritchie et al. 2007), music (Hoover et al. 2012) or visuals (Secretan et al. 2011). In that regard, the core arguments put forth in this paper can be connected to similar positions regarding human interaction and computational creativity (Kantosalo et al. 2014; Bown 2014; Maher 2012).

When examining the different mixed-initiative tools in terms of both computational and human creativity, it becomes obvious that the frame (F^g) of the creative process is provided by the human user — when one is necessary. This is done by specifying high level goals in Sentient World, or by seeding evolution in Sentient Sketchbook and Iconoscope. In the current examples the computational initiative communicates with the user via its outputs (as optional or non-optional suggestions) without *framing information* as "a piece of natural language text that is comprehensible by people" (Colton, Charnley, and Pease 2011). It is worthwhile, however, to investigate such a possibility as it is expected to substantially enhance the human experience when interacting with such a tool, and drive home the notion of human-computer interaction as a dialogue (Novick and Sutton 1997) in natural language. The framing information provided by the creative software can highlight the changes it made and argue for its reasons for making such changes[2]. Investigating how computational framing information affects human interaction can verify hypotheses regarding the impact of the human user's perception of creativity in mixed-initiative software (Yannakakis, Liapis, and Alexopoulos 2014) and computational creativity in general (Colton 2008). Moreover, it can provide fertile ground for examining how framing information itself can be mediated between human and computer in a mixed-initiative fashion.

Conclusion

This paper investigated how mixed-initiative interaction can be considered under the prism of computational creativity research, and provided several examples of design tools and creation games which use human creativity to influence, spark and boost the creative capacity of the software. The paper focused on software where human creativity actively contributes to a design dialog with a computational initiative via persistent interaction throughout the creative process, as opposed to initializing the system's variables pre-generation, or evaluating and curating the output post-generation. For such software, the paper argued that computational creativity theories must be adapted to integrate human input in the system's decisions on how to traverse the search space, how to perceive typicality, or how to adapt its aesthetics (among others). Different instances of human creativity contributing to the creativity of the human-computer 'symbiote' (Licklider 1960) — and in different capacities — were highlighted in several instances of mixed-initiative interaction software. The analysis of the chosen software also underlined promising areas for future work in mixed-initiative interaction, especially in providing the computational initiative with the ability to frame its contributions to the creative process by presenting them in natural language.

Acknowledgments

This work has been supported in part by the FP7 Marie Curie CIG project AutoGameDesign (630665) and the FP7 ICT project C2Learn (318480).

References

Arnheim, R. 2004. *Art and visual perception: a psychology of the creative eye*. University of California Press, revised & expanded edition.

Barros, G. A. B.; Liapis, A.; and Togelius, J. 2015. Data adventures. In *Proceedings of the FDG workshop on Procedural Content Generation in Games*.

Björk, S., and Holopainen, J. 2004. *Patterns in Game Design*. Charles River Media.

Boden, M. 1992. *The Creative Mind*. London: Abacus.

Bown, O. 2014. Empirically grounding the evaluation of creative systems: Incorporating interaction design. In *Proceedings of the International Conference on Computational Creativity*.

Cardoso, A.; Veale, T.; and Wiggins, G. A. 2009. Converging on the divergent: The history (and future) of the international joint workshops in computational creativity. *AI Magazine* 30(3).

[2]For instance, a suggestion in Sentient Sketchbook can highlight that it moved one base behind impassable tiles because it made the exploration effort between bases/players more balanced.

Clark, A. 1998. *Being there: Putting brain, body, and world together again*. MIT press.

Colton, S.; Charnley, J.; and Pease, A. 2011. Computational creativity theory: The face and idea descriptive models. In *Proceedings of the International Conference on Computational Creativity*.

Colton, S.; De Mantaras, R. L.; and Stock, O. 2009. Computational creativity: Coming of age. *AI Magazine* 30(3).

Colton, S. 2008. Creativity vs. the perception of creativity in computational systems. In *Papers from the AAAI Spring Symposium on Creative Intelligent Systems*.

Compton, K., and Mateas, M. 2015. Casual creators. In *Proceedings of the International Conference on Computational Creativity*.

Cook, M., and Colton, S. 2012. Aesthetic considerations for automated platformer design. In *Proceedings of the AAAI Artificial Intelligence for Interactive Digital Entertainment Conference*.

Correia, J.; Machado, P.; Romero, J.; and Carballal, A. 2013. Feature selection and novelty in computational aesthetics. In *Proceedings of Evolutionary and Biologically Inspired Music, Sound, Art and Design (EvoMusArt)*, volume 7834, LNCS. Springer.

Grace, K.; Maher, M. L.; Fisher, D.; and Brady, K. 2014. Data-intensive evaluation of design creativity using novelty, value, and surprise. *International Journal of Design Creativity and Innovation*.

Hoover, A. K.; Szerlip, P. A.; Norton, M. E.; Brindle, T. A.; Merritt, Z.; and Stanley, K. O. 2012. Generating a complete multipart musical composition from a single monophonic melody with functional scaffolding. In *Proceedings of the International Conference on Computational Creativity*.

Jordanous, A. 2012. A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation* 4(3).

Kantosalo, A.; Toivanen, J. M.; Xiao, P.; and Toivonen, H. 2014. From isolation to involvement: Adapting machine creativity software to support human-computer co-creation. In *Proceedings of the International Conference on Computational Creativity*.

Kantosalo, A.; Toivanen, J. M.; and Toivonen, H. 2015. Interaction evaluation for human-computer co-creativity: A case study. In *Proceedings of the International Conference on Computational Creativity*.

Kimbrough, S. O.; Koehler, G. J.; Lu, M.; and Wood, D. H. 2008. On a feasible-infeasible two-population (fi-2pop) genetic algorithm for constrained optimization: Distance tracing and no free lunch. *European Journal of Operational Research* 190(2):310–327.

Krzeczkowska, A.; El-Hage, J.; Colton, S.; and Clark, S. 2010. Automated collage generation – with intent. In *Proceedings of the International Conference on Computational Creativity*.

Lehman, J., and Stanley, K. O. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation* 19(2):189–223.

Liapis, A.; Martínez, H. P.; Togelius, J.; and Yannakakis, G. N. 2013. Transforming exploratory creativity with DeLeNoX. In *Proceedings of the International Conference on Computational Creativity*.

Liapis, A.; Hoover, A. K.; Yannakakis, G. N.; Alexopoulos, C.; and Dimaraki, E. V. 2015. Motivating visual interpretations in iconoscope: Designing a game for fostering creativity. In *Proceedings of the Foundations of Digital Games Conference*.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2012. Adapting models of visual aesthetics for personalized content creation. *IEEE Transactions on Computational Intelligence and AI in Games* 4(3):213–228.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2013a. Designer modeling for personalized game content creation tools. In *Proceedings of the AIIDE Workshop on Artificial Intelligence & Game Aesthetics*.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2013b. Sentient sketchbook: Computer-aided game level authoring. In *Proceedings of the Foundations of Digital Games Conference*, 213–220.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2013c. Sentient world: Human-based procedural cartography. In *Proceedings of Evolutionary and Biologically Inspired Music, Sound, Art and Design (EvoMusArt)*, volume 7834, LNCS. Springer. 180–191.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2013d. Towards a generic method of evaluating game levels. In *Proceedings of the AAAI Artificial Intelligence for Interactive Digital Entertainment Conference*.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2014. Designer modeling for sentient sketchbook. In *Proceedings of the IEEE Conference on Computational Intelligence and Games (CIG)*.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2015. Constrained novelty search: A study on game content generation. *Evolutionary Computation* 23(1):101–129.

Licklider, J. 1960. Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*.

Maher, M. L. 2012. Computational and collective creativity: Whos being creative? In *Proceedings of the International Conference on Computational Creativity*.

Novick, D., and Sutton, S. 1997. What is mixed-initiative interaction? In *Proceedings of the AAAI Spring Symposium on Computational Models for Mixed Initiative Interaction*.

Reintjes, J. F. 1991. *Numerical control: making a new technology*. Oxford University Press, Inc.

Ritchie, G.; Manurung, R.; Pain, H.; Waller, A.; Black, R.; and O'Mara, D. 2007. A practical application of computational humour. In *Proceedings of the International Joint Workshop on Computational Creativity*.

Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17:76–99.

Scaltsas, T., and Alexopoulos, C. 2013. Creating creativity through emotive thinking. In *Proceedings of the World Congress of Philosophy*.

Secretan, J.; Beato, N.; D'Ambrosio, D. B.; Rodriguez, A.; Campbell, A.; Folsom-Kovarik, J. T.; and Stanley, K. O. 2011. Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evolutionary Computation* 19(3):373–403.

Stanley, K. O., and Miikkulainen, R. 2002. Evolving neural networks through augmenting topologies. *Evolutionary Computation* 10(2):99–127.

Veale, T. 2014. Coming good and breaking bad: Generating transformative character arcs for use in compelling stories. In *Proceedings of the International Conference on Computational Creativity*.

Ventura, D. 2008. A reductio ad absurdum experiment in sufficiency for evaluating (computational) creative systems. In *Proceedings of the International Joint Workshop on Computational Creativity*.

Wiggins, G. A. 2006. Searching for computational creativity. *New Generation Computing* 24(3):209–222.

Yannakakis, G. N.; Liapis, A.; and Alexopoulos, C. 2014. Mixed-initiative co-creativity. In *Proceedings of the Foundations of Digital Games Conference*.