

CAN COMPUTERS FOSTER HUMAN USERS' CREATIVITY? THEORY AND PRAXIS OF MIXED-INITIATIVE CO-CREATIVITY

Antonios Liapis, Georgios N. Yannakakis,
Constantine Alexopoulos & Phil Lopes

Abstract: *This article discusses the impact of artificially intelligent computers to the process of design, play and educational activities. A computational process which has the necessary intelligence and creativity to take a proactive role in such activities can not only support human creativity but also foster it and prompt lateral thinking. The argument is made both from the perspective of human creativity, where the computational input is treated as an external stimulus which triggers re-framing of humans' routines and mental associations, but also from the perspective of computational creativity where human input and initiative constrains the search space of the algorithm, enabling it to focus on specific possible solutions to a problem rather than globally search for the optimal. The article reviews four mixed-initiative tools (for design and educational play) based on how they contribute to human-machine co-creativity. These paradigms serve different purposes, afford different human interaction methods and incorporate different computationally creative processes. Assessing how co-creativity is facilitated on a per-paradigm basis strengthens the theoretical argument and provides an initial seed for future work in the burgeoning domain of mixed-initiative interaction.*

Keywords: Computational creativity, human-computer interaction, computer-aided design, digital games, lateral thinking

Introduction

For over a decade, the use of digital computers (in the form of personal computers, smartphones, tablets, smart TVs, etc.) has become ubiquitous and indispensable for not only business people but also young adults, children, and the elderly. Digital technology offers diverse benefits to the lives of many; within formal and informal education, technology-enhanced learning encompasses digital systems which directly support learning activities, often existing online (Browne, Hewitt, Jenkins, & Walker, 2008). Given that creativity is increasingly being considered as an explicit educational objective within formal education (Sawyer, 2006), (Cachia, Ferrari, Kearney, Punie, & Van, 2009), it is imperative that the role of the computer in fostering human creativity is investigated. Digital technologies have demonstrated their capabilities in facilitating users to express their creativity (e.g. with intuitive photo editors) and to share it (e.g. via e-mail clients or social media). Instead, this article focuses on mixed-initiative computational tools which exhibit their own type of intelligence and creativity, and investigates how interaction with such tools influences the creativity both of the human user and of the computer.

Despite the lack of a concrete definition (Novick & Sutton, 1997), **mixed-initiative interaction** in this article refers to a computer and a human user both **proactively** contributing to the solution of a problem. In tasks involving computer-aided design, mixed-initiative interaction assumes a proactive computational initiative which is capable of a modicum of creativity in itself. However, mixed-initiative design does not necessitate an equal contribution from both the human and the computer. Drawing parallels between mixed-initiative interaction and conversation (Novick & Sutton, 1997), Novick and Sutton identify three types of initiative: **task initiative** (deciding the topic), **speaker initiative**

(deciding when each actor takes a turn), and **outcome initiative** (deciding when the problem is solved). With this type of initiative in mind, it is common for mixed-initiative tools (including the ones studied in this article) to allow the human user to take the task initiative, and usually the outcome initiative; most often, mixed-initiative tools take the role of an interlocutor, taking turns with the user in 'asking' or 'responding' to requests regarding the task and its outcome.

This article argues that interaction with a proactive computational initiative which is capable of its own creativity can foster the creativity of the human user. The **mixed-initiative co-creativity (MI-CC)** which emerges from this human-computer interaction cannot be ascribed either to the human or to the computer alone, and surpasses both contributors' original intentions. The human user is inspired by computational input, with optional suggestions or explicit changes to human creations acting as the stimulus for lateral thinking on the part of the designer. This process will be linked to theories of human creativity as well as computational creativity, with the focus on how human-computer interaction can affect and enhance both. This article is built upon the theoretical foundation of Yannakakis et al. (2014) which introduced the concept of mixed-initiative co-creativity. In this paper however, we largely extend previous work by investigating the potential of *collaborative* human and computational creativity and by exposing a number of case studies which realize different degrees of initiative and different ways that human (and computational) creativity can be fostered.

The article lays down the theoretical frameworks under which human and computational creativity is approached, linking them to the concept of mixed-initiative co-creativity. The theoretical argument for MI-CC is strengthened by four instances of design tools and games which incorporate algorithms in different proactive roles. The article concludes with a discussion on the possible extensions of the MI-CC paradigms shown in the presented MI-CC instances.

Human Creativity

The topic of creativity has always fascinated humanity at large, which has led to creativity theories formed around different academic fields and perspectives, such as philosophy (Wittgenstein, 2010), neuroscience (Damasio, 2001) or psychology (Sternberg, 1999). Several types of creative processes have been identified in the literature: examples include everyday, social (**little-c**) creativity (Jeffrey & Craft, 2001), (Craft, 2002), passive and active creativity (Beaney, 2005), exploratory, combinatorial and transformational creativity (Boden M. A., 2003). Investigating how mixed-initiative co-creativity occurs can therefore be pursued via several different lenses and theoretical frameworks. Due to the very nature of the mixed-initiative tools examined in this article (which focus on computational suggestions as stimuli to human creativity), creativity will be regarded primarily from the perspectives of lateral thinking (De Bono, *Lateral thinking: Creativity step by step*, 2010) and creative emotive reasoning (Scaltsas & Alexopoulos, 2013).

Lateral Thinking

In mixed-initiative interaction, a proactive computational initiative is aligned with the general principles of **lateral thinking** (De Bono, *Lateral thinking: Creativity step by step*, 2010) and **creative emotive reasoning** (Scaltsas & Alexopoulos, 2013), the latter being an instance and specialization of the former. Lateral thinking (De Bono, *Lateral thinking: Creativity step by step*, 2010) is the process of solving seemingly unsolvable problems or tackling non-trivial tasks through an indirect, non-linear, creative approach. According to De Bono, lateral thinking skills can be taught. MI-CC realizes the very nature of lateral thinking which, as a creativity process, is boosted through (increasingly) constrained spaces of solutions (De Bono, *Lateral thinking: Creativity step by step*, 2010). Co-creation

with computational creators of visual art and design (including game level design) encapsulates the very core principles of **diagrammatic reasoning** as human creativity, and especially lateral thinking creativity, is often associated with construction and the principles of customization (De Bono, *Lateral thinking: Creativity step by step*, 2010).

The **random stimulus** principle of lateral thinking (Beaney, 2005) relies on the introduction of a foreign conceptual element with the purpose of disrupting preconceived notions and habitual patterns of thought, by forcing the user to integrate and/or exploit the foreign element in the creation of an idea or the production of a solution. Randomness within lateral thinking is the main guarantor of foreignness and hence of stimulation of creativity (Beaney, 2005). According to creative emotive reasoning – which enriches the basic notions of lateral thinking with semantic, diagrammatic and emotive dimensions – the creative act is understood as an intervention that results in **re-framing**. Frames can be viewed as systems or established routes that divide the possibility space (e.g. the game design space) into bounded, meaning-bearing sub-areas. The disruption of an established routine is identified as a **lateral path**. More precisely a lateral path is a cognitive process that promotes deep exploration of a possibility space, whilst satisfying stated (or implicit) conditions, i.e. under constraints. On that basis, the random stimulus and the re-framing principles have one element in common: they are enablers of a change in the lateral path. The principles of re-framing and the random stimulus are embedded in the MI-CC paradigm as machine creativity offers heuristically-driven stimuli that are often altered through e.g. mutations within a genetic algorithm; that can, in turn, alter the user's framing on a particular task/problem. An artificial mutation to a visual diagram, an image, or a game map, resembles the random stimulus that can act as a potentiator of creativity and cause an alteration of lateral thinking.

MI-CC and Diagrammatic Reasoning

Diagrammatic reasoning can be defined as reasoning via the use of visual representations; a cognitive process which is enabled during game level design, interaction design and visual art. These representations can include all forms of imagery incorporating visual features (object shape, size, color, spatial orientation etc.) (Cheng, Lowe, & Scaife, 2001). Literature suggests that complex information processing is benefited by the use of diagrams, due e.g. to the fact that information in diagrams is indexed by spatial location, thus preserving explicitly the geometric and topological relations of the problem's elements (see e.g. (Larkin & Simon, 1987)). Diagrammatic reasoning is premised on the background knowledge of the relevant domain, as well as the specific nature of the diagram and its interconnections with the context within which one encounters it (Cheng, Lowe, & Scaife, 2001).

Diagrammatic Lateral Thinking (DLT) fuses the principles of diagrammatic reasoning and lateral thinking. Diagrammatic lateral thinking builds upon the **extended mind** theory (Clark, 1998): its core idea is that a diagram, through its use, serves as a vehicle of cognitive processes, embodying the various aspects of the problem. The user's mind is extended onto the diagram and reasoning proceeds through structural (rather than semantic or syntactical) entailment. One therefore **thinks** through the diagram rather than its use as a simple image. According to DLT, the process of constructing a diagram (an image, a map, or a character) is more important than the final product (Vile & Polovina, 1998). Moreover, the possibilities one sees for constructing, altering or transforming a given diagram are part of one's comprehension of the diagram itself; the functions of the diagram both on the semantic and pragmatic level are determined in part by these possibilities (Sloman, 2002).

MI-CC can not only be viewed as being closely related to lateral thinking but furthermore that it often constitutes a type of DLT: MI-CC occurring through diagrammatic representations (e.g. in game level design) offers diagrammatic alternative

paths that satisfy a number of conditions. These define non-linear **lateral paths** within the creative (possibility) space as they promote deep exploration of the space of possibilities which is, in turn, a core lateral thinking characteristic. DLT within MI-CC does not necessarily embed transformational creativity processes as identified by Boden (Boden M. A., 2003). The majority of MI-CC instances presented in this article realize DLT, as co-creativity occurs mainly on the visual (diagrammatic) level. MI-CC expands the very notion of DLT as it dichotomizes diagrammatic lateral thinking into two main creativity dimensions: one that is based on **analogical** thinking from diagrams and images and one that works purely on the **visual** level through imagistic lateral thinking pathways (Scaltsas & Alexopoulos, 2013). Details on the nature and impact of analogical DLT and visual DLT in the computer's suggestions during the design process are provided in the case study of the *Sentient Sketchbook* design tool.

Computational Creativity

Some of the fundamental questions within computational creativity research are “what does it mean to be creative?” and “does creativity emerge within the individual, the process, the product, or some combination of all three?”. The questions are as relevant to human as to machine creativity (Boden M. A., 2003), (Colton, 2008). Computational creativity, however, seeks creativity generated by, enhanced or fostered via algorithmic means.

Computational creativity literature suggests that **value** (or **usefulness**) and **novelty** are key elements characterizing a creative process (Boden M. A., 2003). An autonomous generative system is able to try out exhaustively many possible novel combinations of elements, often resulting in largely uninteresting outcomes or artifacts. For that very reason, computational creativity not only requires the generated artifacts to be **novel**, but also **valuable**. While other aspects of creativity have been discussed and proposed (such as surprise (Macedo & Cardoso, 2001)), novelty and value define the common denominators accepted by most theories within computational creativity. If the space of possibilities within MI-CC is constrained for both the machine and the human, the creative process is ultimately of **value** for both given the problem constraints set either by the human user or by an external observer (e.g. domain expert). Moreover, if the computer searches within a constrained space of possibilities for orthogonally possible solutions then the computer interacts with the human user by offering both **useful** and **novel** suggestions throughout the creative process (Boden M. A., 2003). The end outcome of MI-CC (both novel and useful) is ultimately a result of iterative co-creation. The autonomous creative system, in that case, finds novel ways to navigate a search space, by e.g. looking at orthogonal aspects of the human creative process; the computational discoveries from this search are suggested back to the human.

Computational creativity has been classified by (Boden M. A., 2003) in three types: **combinatorial**, **exploratory** and **transformational**. Combinatorial creativity revolves around the combination of different elements which is often trivially accomplished by a computer. Computers are also well suited for exploratory creativity, which involves traversing a well-defined search space. In contrast, transformational creativity requires the computer to ‘break the rules’ of that pre-existing conceptual space. Among the three types of computational creativity identified by Boden, MI-CC realizes mainly exploratory creativity. While it could potentially achieve transformational creativity, mere exploration of the solution space can often result in more creative outcomes than transformation (Bundy, 1994), (Pind, 1994). Pease et al. provide the example of an unusual but legal chess move as often being more creative than changing the rules of chess (Pease, Winterstein, & Colton, 2001). Ultimately, the borders between these types of creativity are unclear, as transformational creativity can also be viewed as exploration (Wiggins, 2006); the game

asset generator of (Liapis, Martínez, Togelius, & and Yannakakis, 2013), for instance, blurs the edges between transformational and exploratory creativity.

According to (Bundy, 1994) an outcome is considered creative if the possibility space in which it lies is large (and complex) and if it is generated from a less explored area. MI-CC tools that generate solutions which satisfy certain constraints (e.g. constraints on playability for generated game content) capture the complexity expressed by Bundy. The harder it is to find a solution within a constrained search space, the more novel it is deemed (Bundy, 1994). The notion of complexity has also been expressed via a number of alternative computational metrics including rarity and impressiveness (Lehman & Stanley, Beyond open-endedness: Quantifying impressiveness, 2012) that can be considered in a MI-CC tool which involves diagrammatic aspects of creativity.

Realizing Mixed-Initiative Co-Creativity

The previous sections examined how the mixed-initiative interaction between a human user and a proactive computational creator can result in the co-creativity of the human-machine ‘symbiotes’ – to use a term coined by (Licklider, 1960). The impact of a computer-generated stimulus to human creativity, and the impact of human design constraints imposed on computationally creative processes is largely dependent on the type of software, its goals, its interface, and the degree and type of initiative from human and computer. Below are short descriptions of a set of four selected design tools and games which make use of mixed-initiative interaction. The way in which co-creativity can emerge is also discussed for each system. This article focuses on games and game-specific design tools, although the principles described herein can be transferred to other domains (such as industrial schematic design or image/video editors). Games have two key attributes which make them ideal paradigms for mixed-initiative co-creativity: a) as games encompass many different facets (including audio, visuals, game design, narrative, game levels), the task of game development requires extensive human creativity (Liapis, Yannakakis, & Togelius, 2014) and benefits from computer-aided tools such as *Sentient Sketchbook* and *Sentient World*, while b) most digital games – especially freeform creation games such as *Iconoscope* and *4Scribes* – rely on their players’ imagination and have already shown considerable capacity in their use in classrooms (Pirius & Creel, 2010), (Watters, 2011).

Sentient Sketchbook

Sentient Sketchbook is a mixed-initiative tool for game level design (Liapis, Yannakakis, & Togelius, Sentient sketchbook: Computer-aided game level authoring, 2013). Via its user interface, the tool allows the user to draw game levels in the form of low-resolution, high-level map sketches. These map sketches are minimal abstractions of complete game levels, containing the absolutely necessary components for levels of this genre. The map sketches contain passable and impassable tiles (which allow and block movement respectively), as well as game-specific tiles such as weapon pickups for a first-person shooter level, player bases for a strategy game, or monsters and treasure for dungeon adventure games. The abstract map sketches can be automatically converted by the computer into high resolution, playable game levels (see Figure 1 for a strategy game level example). As the users draw on the abstract map sketch which contains only a handful of tiles, they can create complete game levels within minutes. The low effort of level design facilitated by *Sentient Sketchbook* enables novice users to create game levels without extensive experience, but also motivates experts and novices alike to attempt original, untried designs.




Figure 1: In *Sentient Sketchbook*, both initiatives contribute to creating the simple map sketch to the left, which however can be automatically converted to a detailed map on the right.



Figure 2: While the user draws the map sketch (left), multiple computational suggestions appear to the right. The user can select a suggestion at any time and replace their current sketch.

The role of the computational initiative in *Sentient Sketchbook* is three-fold. Firstly, the computer can automatically (and within seconds) convert the user's map sketch into a fully detailed game level; this alleviates the users' effort in managing the minutiae of the level's design and allows them to focus on the creative, high-level ideas. Secondly, the computer is able to discern when map sketches are not playable, either because they do not contain vital tiles (such as a maze level without an exit tile), or because some parts of the level are inaccessible (such as a treasure which cannot be reached from the player's starting location); the feedback from the computer allows users to correct their unplayable designs and ensures that even novices with no level design experience can create valuable results. Finally, the computer **proactively** contributes to the design process by creating suggestions for the human user to consider. These suggestions are map sketches, similar to what the human user is drawing, and they are generated by the computer and presented to the user in real-time, as the users are designing their own sketch (see Figure 2). At any time during the design process, the user can select a computer-generated suggestion, compare it to their current design, and replace their sketch with the suggestion. The suggestions are generated via evolutionary computation (De Jong, 2006), and take the user's current map sketch as inspiration. This ensures that the computer-generated suggestions will have many structural and visual similarities with what the user is currently designing; the suggestions will thus not alienate the user, but will appear as improvements of their current sketch. All computer-generated suggestions presented to the user are ensured to be playable, as the computer can test its creations against the playability constraints it applies on user sketches. Beyond this ensured playability, half of the computer-generated suggestions are evolved towards maximizing certain game-specific qualities (Kimbrough, Koehler, Lu, & Wood, 2008) which are modeled into the program by expert designers: for example a suggestion for a strategy game level will attempt to improve the game balance between players, the area that each player base can control at the start of the game, and the distribution of strategic resources. Suggestions which improve a user's sketch by maximizing some game-specific properties ensure that the computer's contribution to the design process is **valuable**, and are particularly helpful to novices which may lack the expert knowledge imparted to the computer. The other half of the computer-generated suggestions are evolved towards visual novelty (Liapis, Yannakakis, & Togelius, 2013), creating suggestions which are as visually different (in terms of tile placement) as possible from each other but also from the user's sketch. Suggestions which target visual novelty ensure that the computational input to the design process is **novel** to what the user is currently drawing, while also **valuable** since the generated sketches are ensured to be at least playable.

Sentient Sketchbook has been the first case study for mixed-initiative co-creativity (Yannakakis, Liapis, & Alexopoulos, 2014). On a theoretical level, computational suggestions in *Sentient Sketchbook* perform the role of stimuli which can lead to lateral

thinking. Since the design of game levels (as realized by *Sentient Sketchbook*) relies strictly on their diagrammatic representation, the type of creativity incited by the computational initiative is diagrammatic lateral thinking. More specifically, suggestions evolved to improve game-specific qualities prompt analogical diagrammatic lateral thinking, as game-specific tiles are treated differently than others (for instance, player bases are far more important than impassable tiles in a strategy game, as they determine the players' chances of winning). Suggestions evolved to create visually divergent suggestions from the user's sketch prompt visual diagrammatic lateral thinking, as the algorithm is agnostic of game properties and the suggestions appeal to the users' perception (instead of their level design experience). Beyond the effects of computational suggestions on human creativity, the algorithms used to generate them satisfy the requirements of computational creativity on **valuable** and **novel** output.

In order to evaluate the impact of the computational suggestions on the users' creative process, a study of five expert designers using *Sentient Sketchbook* for creating a total of 24 game levels was conducted. The study, which is detailed in (Yannakakis, Liapis, & Alexopoulos, 2014), investigated the degree of use (i.e. how often users selected computational suggestions, and reasons for cases where suggestions were not desirable), the qualitative evaluation of the creation paths (i.e. what design frames the users prioritized during the design and how the computational suggestions affected those), the quantitative evaluation of the creation paths (i.e. how the maps' appearance changed during users' drawing phases and computational suggestion phases) and the evaluation of creativity by a human audience (i.e. which steps of the creation path were considered creative milestones by designers other than the original user of *Sentient Sketchbook*). Results indicate that while computer-generated suggestions are not used often (and in some creation paths not used at all), they can result in major changes in the map sketches' appearance and often constitute creative milestones due to their ability to prompt diagrammatic lateral thinking (both in the tool's active user and in an inactive audience). Figure 3 shows an example of a creative milestone from (Yannakakis, Liapis, & Alexopoulos, 2014), where the designer's frame of reference (regarding the notion that symmetry on the visual level can ensure a fair gameplay between two competing players, whose bases are shown in white) is disrupted by the computational suggestion which was selected by the designer to replace their previous level. The computer-generated output breaks the visual patterns and introduces more imbalance (in the form of resource tiles in cyan closer to one player). Note, however, that much of the remaining level structure (such as the positions of white tiles) remains intact as the computer uses the designer's map as a starting seed. While the user's rationale for the level change is not known (as users were not asked to narrate their design process), 3 out of 4 audience members which evaluated this creation path identified the design step shown as a creative milestone.




Figure 3: An indicative creative milestone, tagged by 3 out of 4 audience members in the study detailed in (Yannakakis, Liapis, & Alexopoulos, 2014). The user chooses to break the symmetrical look of their designed level (left) in order to embrace the notion of a computer-generated 'asymmetrical' game level (right) which may be of interest to players of different playing skills.

Sentient World

Sentient World is a mixed-initiative tool for the task of designing gameworlds and more specifically their terrain (Liapis, Yannakakis, & Togelius, 2013). Terrain is important for large-scale Role-Playing Games, and can affect both the vegetation and climate but also civilization growth and types of goods produced in the region; however, *Sentient World* does not create terrain for a specific game and is decoupled by any game rules or playability constraints. The user begins drawing terrain in *Sentient World* on a very coarse map (i.e. nine tiles) and can only specify land or water tiles (see Figure 4). After drawing their low-level sketch, the user presses a “refine” button on the interface, at which point the computational initiative takes over and returns a higher-resolution version of the terrain, with nine times as many tiles and including details on hills, mountains and plains (see Figure 5). The user can select among the eight possible refined versions of their terrain, and edit it further if they wish. After this point, the computer can refine this further, creating an even larger map with details on shorelines, shallow seas, low hills etc.

The computational input of *Sentient World* in the creative process is **not optional** (contrary to the optional suggestions of *Sentient Sketchbook*) and takes the form of **turn-taking speaker initiative** (with the human user taking a turn editing the terrain and the computer taking a turn refining it). Unlike the suggestions of *Sentient Sketchbook*, moreover, the user and the computer have different tools at their disposal: the human user can only control the rough sketching process, while the computer can only control the refining process. The algorithms behind *Sentient World* combine novelty search (Lehman & Stanley, 2011), which creates visually divergent terrain from what the user has drawn, with backpropagation (Rumelhart, 1995), which attempts to fit the generated map to the lower-resolution user creation while extrapolating the higher-resolution terrain elevation details. The combination of these algorithms ensures an initial **novel** seed (which the user would find surprising) and then adapts it to become **valuable** by obeying some of the high-level user specifications.

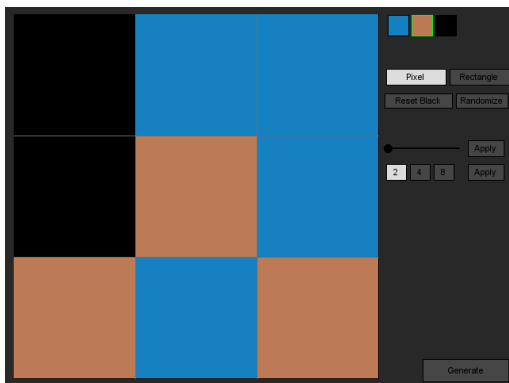


Figure 4: Initially the user of *Sentient World* paints (on 3 by 3 grid) a rough terrain sketch with water, land or blank tiles.




Figure 5: The computational initiative attempts to create higher-detail terrain (with mountains and hills) which conforms to the user's sketch patterns.

In terms of interaction between human and computational creativity, on the one hand computational creativity is stimulated by being constrained by the human rough terrain sketch, forcing it to both satisfy those constraints while also creating results which are not similar to each other. On the other hand, human creativity can be fostered by the computer-provided details to their rough sketch; having many alternative high-resolution terrain to choose from (as well as the option to edit them further), the generated terrain acts as a stimulus for visual diagrammatic lateral thinking (as it operates on the visual appearance of the terrain rather than any function it may serve in a game). Moreover, the human user can leave areas of their terrain sketch blank, letting the computer add details to those as it sees fit (without constraining its output). This allows the human user to

control the degree and freedom of the computational initiative, balancing between human authorial control (by specifying all elements of the rough terrain) and almost freeform, serendipitous co-creativity (by leaving most tiles of the terrain blank).

Iconoscope

Iconoscope is a creation game played on Android tablets, which revolves around the visual depiction of semantic concepts in a creative fashion (Liapis, Hoover, Yannakakis, Alexopoulos, & Dimaraki, 2015). The goal of the game is for players to create icons representing a concept (such as *heritage* or *dominate*) which they chose among three thematically or semantically linked concepts (e.g. *lead*, *govern*, *dominate*). The drawing interface (see Figure 6) allows only the use of simple shapes (e.g. circles, hearts, rhombi) and a few colors, constraining players to creatively combine them in meaningful ways but also abstract away from simply pictorial representations – which is enhanced by the semantically abstract concepts which must be represented. *Iconoscope* is played in a group of four or more players, with the winner of a game session determined by peer evaluation: each other player attempts to guess which of the three concepts the player’s icon represents. *Iconoscope* rewards high scores to icons which are ambiguous enough that the underlying concept is communicated to some but not all other players (i.e. some players guess the concept that the user chose to represent, and some others guess different concepts). The social component of observing each other’s creations and attempting to ‘trick’ the other players both influences the fun of gameplay and promotes community and shared values (Chappell, Craft, Rolfe, & Jobbins, 2012). The design of *Iconoscope* and its connection to both wise humanizing creativity and creative emotive thinking, is detailed in (Liapis, Hoover, Yannakakis, Alexopoulos, & Dimaraki, 2015).



Figure 6: The drawing interface of *Iconoscope*

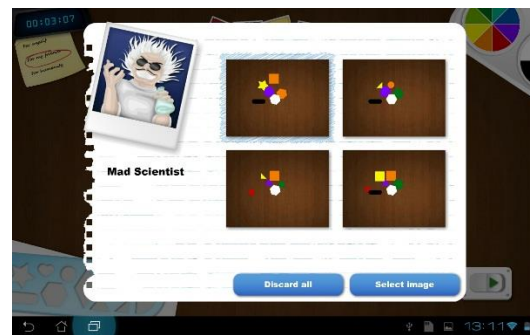


Figure 7: *Iconoscope* Assistant providing novel alternatives to the user’s icon

Besides the interaction among players, which takes place before and after a game session, computational suggestions are provided to each player in real-time as they draw their icon. These computational suggestions are provided by assistants, each with their own portrait, name and ‘personality’ (i.e. objective when generating suggestions). Similar to *Sentient Sketchbook*, most assistants change the player’s currently drawn icon by moving, recoloring and changing its shapes, or adding new shapes. The five assistants of *Iconoscope* each has different algorithmic goals, such as showing past users’ icons (rather than computer-generated ones), generating random permutations of the user’s icon, targeting visual novelty from the user’s current icon, or trying to diverge or converge towards a ‘typical’ icon for this concept specified by an expert (e.g. a red heart for the *love* concept). Users can request for an assistant’s suggestions by selecting its portrait on the drawing interface, and can choose one of the assistant’s suggestions to replace their current icon and continue drawing from there (see Figure 7). While the collaborative activity of guessing which concept is represented by which icon **after creation** is a stimulus for collaborative creativity and shared values, the role of assistants and computational

suggestions **during creation** acts as a stimulus for diagrammatic lateral thinking and prompts individual creativity.

4Scribes

4Scribes is a collaborative storytelling game played either digitally, on Android tablets, or as an analog game using special cards (Eladhari, Lopes, & Yannakakis, 2014). Both the digital and the analog version of *4Scribes* is played with four players, using cards which contain an evocative illustration (serving as a diagrammatic stimulus) and a caption of a few words (usually one). Examples of cards from the digital game are shown in Figure 8. *4Scribes* can be played fully collaboratively where players all decide on the story's ending, or competitively where players try to steer the story towards a specific ending described in a special "myth" card. Players begin the game by drawing five story cards and an additional "myth" card which is used for concluding the story. Players take turns playing a card and connecting it to the story being told thus far (writing down how the story progresses as this card enters play). Story cards can be characters which can be introduced to the story, or scene elements (e.g. emotions, events or items). Players do not gain new cards during play: their initial draw determines the entirety of their story contribution, which allows them to plan ahead accordingly. Once each player has placed 5 cards onto the story (thus leaving their hand empty save for the "myth" card), they choose to conclude the story using their "myth" card as inspiration (in the case of collaborative storytelling) or by revealing the ending they had preplanned with their "myth" card (in the case of competitive storytelling). All players then vote on which ending was the most appropriate (thematically and dramatically), and the winner resolves the story by applying their ending to it.



Figure 8: The 2nd player of *4Scribes* contributing to the story. At the bottom you can see the remainder of the 2nd player's hand (4 cards). The current assistant, who provided the players' initial cards, is shown as a bookmark (top right)

Unlike *Sentient Sketchbook* and *Iconoscope*, the computational initiative in the case of *4Scribes* does not contribute during play, while players put down story cards, but is used to determine each player's starting cards. Similar to *Iconoscope*, one among four different assistants can be chosen at the start of the game for allocating the players' cards: depending on which assistant is chosen, the cards may be chosen randomly (similar to a normal shuffle of the deck), chosen based on their semantic novelty (i.e. as different cards as possible among players), or based on their similarity or dissimilarity from an expert-defined 'typical' set of story cards. While most computational suggestions of *Iconoscope* rely

on visual difference (as the game relies on diagrammatic representations of concepts), the storytelling goal of *4Scribes* necessitates that the players' potential card sets are evaluated **semantically**, i.e. on the semantic difference between the cards' captions. The semantic difference in this case is based on the co-occurrence of the cards' words in a large corpus of texts; the less often these words co-occur in the same text, the larger their semantic difference. Beyond the differences in how artifacts are evaluated (semantically instead of diagrammatically), the computational initiative of *4Scribes* differs from that of *Sentient Sketchbook* and *Iconoscope* in that it specifies the **affordances** of the player's game (by choosing which cards are in play, and which players control them). Thus the computer constrains to a degree the possible stories that may emerge, but does not monitor or intervene during the periods of human play. Mixed-initiative co-creativity is achieved by a **computational task initiative**, as the computer specifies the 'topic' (story) of the play session, relinquishing speaker initiative (which card will be played) to individual human players and outcome initiative (how the story will be concluded) to the collaborative human creativity fostered by the group discussion and voting process.

Discussion

This article puts forth several arguments for the co-creativity potential of mixed-initiative interaction; the cases examined include both design tools for creative tasks (i.e. game level design) and game-based learning systems which incorporate a proactive, self-determining artificial intelligence. In this article, the potential of mixed-initiative interaction to foster human creativity is argued from the perspective of a computer-generated random stimulus which triggers the lateral thinking and re-framing of an individual human creator. Essentially, the creativity of the computer disrupts the idiosyncratic frame of an individual creator; this frame can be a certain routine for performing tasks, a lens through which the world is understood, or a pattern of associations between facts, emotions and actions. In order to understand (in the case of optional computer-generated suggestions as in *Sentient Sketchbook* and *Iconoscope*) or conform to (in the case of mandatory computational operations as in *Sentient World* and *4Scribes*) the computational initiative, the user must adjust their visual patterns, design goals, or gameplay preferences. On the other hand, the human initiative influences computational creativity primarily by constraining the possible output of the generator. With a human providing (as is often the case) the task initiative, the search of the system for valuable and novel solutions is limited by the user's specifications; thus, the exploratory creativity of the computer is bound by user intention. For instance, in *Sentient Sketchbook* the suggestions start with the user's current map sketch as a seed: while possibly better game levels could have been generated from an empty canvas, the fact that the computer must attempt to improve a potentially ill-fitted human design increases its creative potential (in finding shortcuts to correct what the human user has done). In *Sentient World*, the computational creator attempts both to create novel solutions which surprise the human user (via novelty search) but it also attempts to retain the human-provided patterns of the rough terrain sketch (via back-propagation); this process exemplifies the way in which computational creativity is both inspired and constrained by the human user while simultaneously attempting to surprise both the user and itself by discovering unexpected areas of the search space which contain valuable creative outcomes.

In the systems used here as case studies of mixed-initiative co-creativity, the human initiative primarily interacts with the computational initiative by inspiring (or seeding) the computational search (e.g. with *Iconoscope* assistants creating permutations of the user's icon), or by specifying features necessary in the final outcome (e.g. by explicitly fitting *Sentient World* generated terrain to conform to user-specified terrain patterns). However, constraining the possibility space of generators is not the only way in which human

initiative can influence computational creativity in a mixed-initiative tool. In particular, the human user can either explicitly or implicitly specify how the computer should evaluate its output. To a degree, this is the case in *Iconoscope* and *4Scribes* where the human user selects which computational assistant they prefer, thus explicitly choosing which heuristics will be used to evaluate the generated outcomes. More indirectly, human users could guide the computational initiative towards areas of the search space which they find (idiosyncratically) more desirable. Interactive evolution (Takagi, 2001) is an algorithm which allows users to evaluate the generated output; the computer performs evolutionary computation treating the user-preferred artifacts as the fittest, resulting in more and more artifacts which bear resemblance to those selected by users. Interactive evolution can be an inherently co-creative process, as the human user and the algorithm “cooperatively optimize target systems based on the mapping relation between physical and psychological spaces” (Takagi, 2001), i.e. the algorithm’s feature parameter space (physical) and the user’s preferences and intuition (psychological). Beyond the explicit selection of evaluation criteria (e.g. by selecting an assistant in *4Scribes*) and iteratively selecting preferred content among those generated (in interactive evolution), a less direct and less fatiguing way of adapting computational creativity to human desires is through designer modeling (Liapis, Yannakakis, & Togelius, 2013). Designer modeling refers to algorithmic methods (such as machine learning) for automatically recognizing the goals, preferences or process of a human designer based on their interactions with a mixed-initiative design tool. A designer model can therefore be useful for personalized, responsive computer-aided design tools; initial experiments of designer modeling with *Sentient Sketchbook* showed its potential at learning the user’s style from prolonged interactions as well as their current process based on their latest activities (Liapis, Yannakakis, & Togelius, 2014). The creativity of the computer can be more closely paired with (and more severely influenced by) the human user’s own creativity if the design process of the latter not only constrains where the computer should explore but also how (based on which criteria and goals). By using automated ways for the machine to learn user preferences, the human creator is not made aware of their preferences or cognitive associations (i.e. their **frames**), thus enhancing the re-framing potential of computational feedback which attempts to explicitly address these.

It should be noted that the majority of research in mixed-initiative interaction (e.g. the work of (Novick & Sutton, 1997)) assumed mixed-initiative interaction to take place between a single human user and a single computational process. Similarly, the case studies presented here largely follow this assumption. *Sentient Sketchbook* and *Sentient World* are standalone tools intended for a level designer working in isolation. *4Scribes* and *Iconoscope* are multi-player games focusing more on collaboration (*4Scribes*) and competition (*Iconoscope*), and thus the computer must accommodate multiple users. In *4Scribes*, the computational initiative must allocate cards to all players, taking into account the balance in each player’s cards (e.g. so that there is no player without a character card to play). *Iconoscope* does not directly account for opponents’ icons or concepts, but one of its assistants can present icons created by any player in the past as suggestions (accounting for the communal aesthetics of the *Iconoscope* player base). Communal and collaborative creativity (Chappell, 2008) are facilitated by the game design, but mostly targeting co-creativity between humans; the computer supports and motivates it (via e.g. starting card allocation and icon suggestions from a communal pool) but takes a less proactive role in those aspects. The role of the proactive computer in fostering co-creativity is more pronounced during periods where human users are pursuing individual creativity, e.g. during *Sentient Sketchbook* sessions or while they individually, secretly draw icons in *Iconoscope*. An argument can be made that computational creativity is more valuable during those tasks which involve individual creativity, acting as a human colleague would (Lubart T. , 2005); when multiple human creators work in a group (even as adversaries in a game),

collaborative creativity will de facto emerge. However, there is fertile ground for research in computers which can inspire a group of designers, players or learners: initial ideas include a computer which observes each group member's creative processes and pairs them with another group member with a conflicting frame (prompting re-framing during the collaboration between the two human users) or by providing conflicting goals or suggestions to each group member in order to encourage discussion and negotiation when human collaborators interact with each other. Beyond human creativity, the mixed-initiative co-creativity in cases where multiple computational processes are involved has not been investigated, but offers another interesting dimension for future research. Such collaborative computational creativity can emerge, for instance, when different systems used by (human) members of the same group are required to share information and coordinate for providing consistent suggestions to all group members (see Figure 9). The impact of this collaboration on computational creativity is likely to lead to transformational creativity as one computational process must change its objectives and preferences (i.e. "frames") when under the influence of another computational process.

Conclusions

This article has argued for the potential of computationally creative processes to foster human creativity in systems incorporating mixed-initiative interaction. Lateral thinking can be triggered by the stimuli of proactive computational creators, either from computer-generated suggestions or from necessary feedback during a creative process. Human creativity also affects the computational processes, as the computer must adapt its objectives and search directions to accommodate the human initiative. Four examples shed light on how different design tools and games can incorporate computationally creative processes and how the goals, algorithms and user interaction modalities affect how mixed-initiative co-creativity occurs. Finally, important future research both from a philosophical and from a technical point of view was identified for strengthening the potential of mixed-initiative co-creativity and broadening it to facilitate a more diverse set of creative tasks and processes.

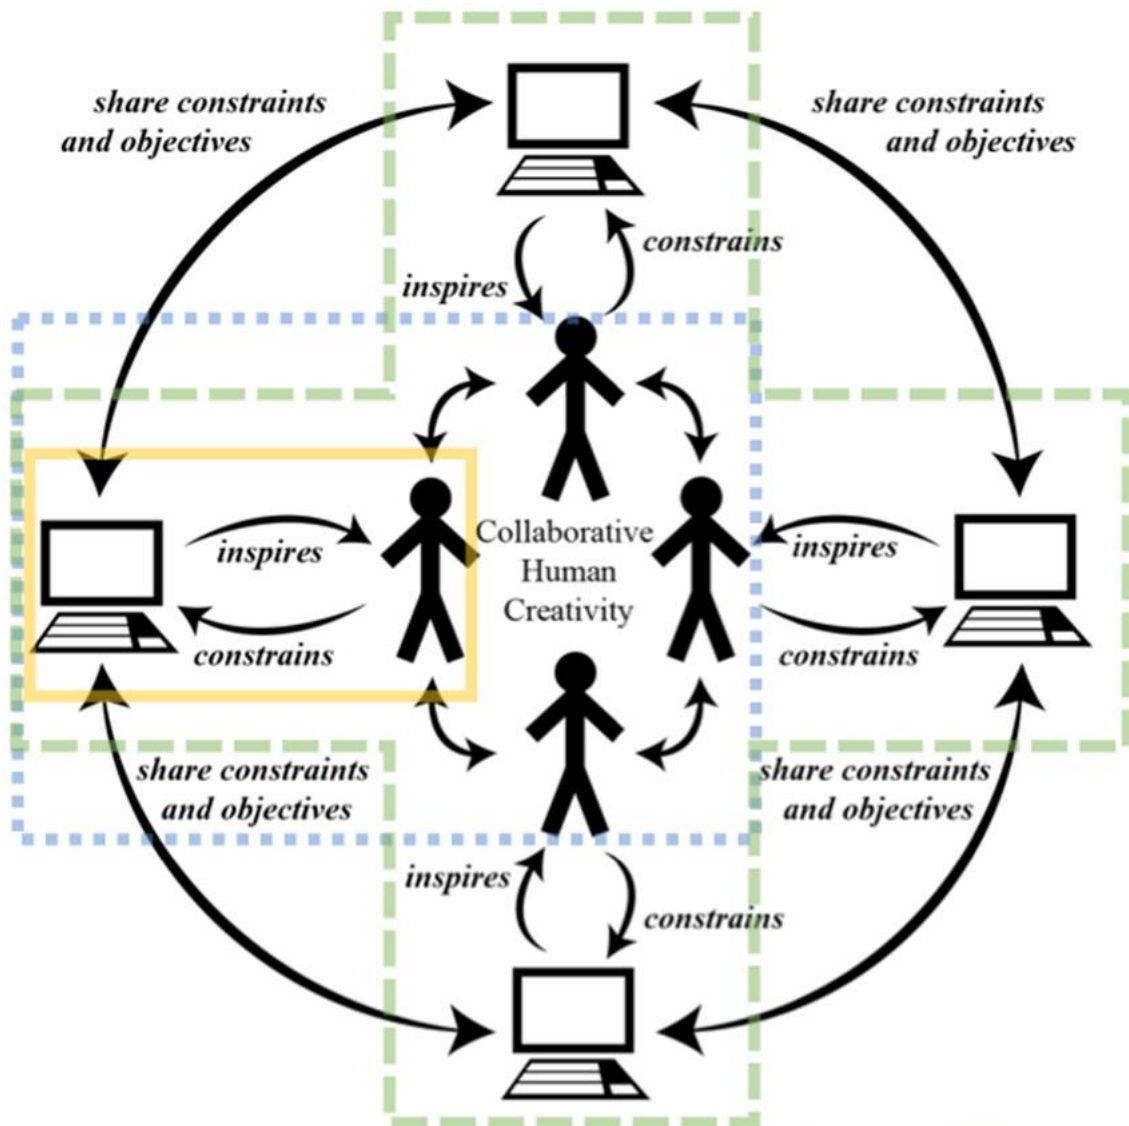


Figure 9: An overview of the potential for mixed-initiative co-creativity, as realized by the different tools enumerated in the article, and as an ideal collaborative mixed-initiative co-creativity. The full figure shows an ideal co-creativity tool where a group of human users is assisted by proactive computational initiatives which also influence each other, either by sharing each of the human creators' goals with each other or by coordinating the simultaneous generation of diverse stimuli for prompting lateral thinking. The current tools focus on smaller portions of this ideal interaction: in *Sentient Sketchbook* and *Sentient World* an individual human creator interacts with an individual computational creator (yellow frame), in *4Scribes* one computational creator defines the possibilities of a group of human creators as a whole (blue frame), while in *Iconoscope* independent computational initiatives interact with human creators (one each) as the latter compete in a group.

Acknowledgements

The authors would like to acknowledge and thank Amy K. Hoover, Evangelia Dimaraki, Pavlos Koulouris and Kerry Chappell for assistance in the design and test of the *Iconoscope* game; Mirjam P Eladhari for the design of *4Scribes*; Serious Games Interactive for the implementation of *Iconoscope* and *4Scribes*. The authors would like to thank the participants of the user study of *Sentient Sketchbook* for their feedback. The research is supported, in part, by the FP7 ICT project C2Learn (project no: 318480).

References

- Beaney, M. (2005). *Imagination and creativity*. Open University Milton Keynes.
- Boden, M. A. (2003). *The creative mind: Myths and mechanisms*. Routledge.
- Bundy, A. (1994). What is the difference between real creativity and mere novelty? *Behavioral and Brain Sciences*, 17 (03).
- Carroll, E. A., Latulipe, C., Fung, R., & Terry, M. (2009). Creativity factor evaluation: Towards a standardized survey metric for creativity support. *Proceedings of the seventh acm conference on creativity and cognition*.
- Chappell, K., Craft, A. R., Rolfe, L., & Jobbins, V. (2012). Humanizing Creativity: Valuing our Journeys of Becoming. *International Journal of Education & the Arts* 13 (8).
- Cheng, P. C.-H., Lowe, R. K., & Scaife, M. (2001). Cognitive science approaches to understanding diagrammatic representations. *Artificial Intelligence Review*, 15 (1-2).
- Clark, A. (1998). *Being there: Putting brain, body, and world together again*. MIT press.
- Colton, S. (2008). Creativity versus the perception of creativity in computational systems. *AAAI spring symposium: Creative intelligent systems*.
- Craft, A. (2002). *Creativity and Early Years Education*. Continuum.
- Damasio, A. R. (2001). Some notes on brain, imagination and creativity. In *The origins of creativity*. Oxford University Press.
- De Bono, E. (2010). *Lateral thinking: Creativity step by step*. Harper Collins.
- Hewett, T., Czerwinski, M., Terry, M., Nunamaker, J., Candy, L., Kules, B., & Sylvan, E. (2005). Creativity support tool evaluation methods and metrics. *Report of workshop on creativity support tools*.
- Hofstadter, D. R. (2008). *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought*. Basic Books.
- Jeffrey, B., & Craft, A. (2001). The Universalization of Creativity. In *Creativity in Education*. Continuum.
- Kimbrough, S. O., Koehler, G. J., Lu, M., & 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).
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive science*, 11 (1).
- Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19 (2).
- Lehman, J., & Stanley, K. O. (2012). Beyond open-endedness: Quantifying impressiveness. *Artificial Life* (13).
- Liapis, A., Hoover, A. K., Yannakakis, G. N., Alexopoulos, C., & Dimaraki, E. V. (2015). Motivating Visual Interpretations in Iconoscope: Designing a Game for Fostering Creativity. *Proceedings of the 10th Conference on the Foundations of Digital Games*.
- Liapis, A., Martínez, H. P., Togelius, J., & and Yannakakis, G. N. (2013). Adaptive game level creation through rank-based interactive evolution. *Proceedings of the IEEE Conference on Computational Intelligence and Games*.
- Liapis, A., Martínez, H. P., Togelius, J., & and Yannakakis, G. N. (2013). Transforming exploratory creativity with DeLeNoX. *Proceedings of the 4th International Conference on Computational Creativity*.
- Liapis, A., Yannakakis, G. N., & Togelius, J. (2013). Designer modeling for personalized game content creation tools. *Proceedings of the AIIDE workshop on artificial intelligence & game aesthetics*.
- Liapis, A., Yannakakis, G. N., & Togelius, J. (2013). Enhancements to constrained novelty search: Two-population novelty search for generating game content. *Proceedings of genetic and evolutionary computation conference*.

- Liapis, A., Yannakakis, G. N., & Togelius, J. (2013). Sentient sketchbook: Computer-aided game level authoring. *Proceedings of the 8th Conference on the Foundations of Digital Games*.
- Liapis, A., Yannakakis, G. N., & Togelius, J. (2013). Towards a generic method of evaluating game levels. *Proceedings of the AAAI artificial intelligence for interactive digital entertainment conference*.
- Liapis, A., Yannakakis, G. N., & Togelius, J. (2014). Designer Modeling for Sentient Sketchbook. *Proceedings of the IEEE Conference on Computational Intelligence and Games*.
- Licklider, J. C. (1960). Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics HFE-1*.
- Macedo, L., & Cardoso, A. (2001). Modeling forms of surprise in an artificial agent. *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*.
- Maher, M. L., Brady, K., & Fisher, D. H. (2013). Computational models of surprise in evaluating creative design. *Proceedings of the fourth international conference on computational creativity*.
- Novick, D., & Sutton, S. (1997). What is mixed-initiative interaction? *Proceedings of the AAAI spring symposium on computational models for mixed initiative interaction*.
- Pease, A., Winterstein, D., & Colton, S. (2001). Evaluating machine creativity. *Workshop on creative systems, 4th international conference on case based reasoning*.
- Pind, J. (1994). Computational creativity: What place for literature? *Behavioral and Brain Sciences 17 (03)*.
- Plato. (1961). *Ion*. In *Plato: The Collected Dialogues*. Princeton University Press.
- Ramachandran, V. S., & Hirstein, W. (1999). The science of art: a neurological theory of aesthetic experience. *Journal of consciousness Studies, 6*.
- Ritchie, G. (2007). Some empirical criteria for attributing creativity to a computer program. *Minds and Machines, 17 (1)*.
- Rumelhart, D. (1995). *Backpropagation: theory, architectures, and applications*. Lawrence Erlbaum.
- Scaltsas, T., & Alexopoulos, C. (2013). Creating creativity through emotive thinking. *Proceedings of the world congress of philosophy*.
- Slovan, A. (2002). *Diagrams in the mind?* Springer.
- Smelik, R. M., Tutenel, T., de Kraker, K. J., & Bidarra, R. (2011). A declarative approach to procedural modeling of virtual worlds. *Computers & Graphics, 35 (2)*.
- Smith, G., Whitehead, J., & Mateas, M. (2011). Tanagra: Reactive planning and constraint solving for mixed-initiative level design. *IEEE Transactions on Computational Intelligence and AI in Games(99)*.
- Sternberg, R. J. (1999). *Handbook of creativity*. Cambridge University Press.
- Takagi, H. (2001). Interactive evolutionary computation: Fusion of the capabilities of EC optimization and human evaluation. *Proceedings of the IEEE (9)*.
- Vile, A., & Polovina, S. (1998). Thinking of or thinking through diagrams? The case of conceptual graphs. *Thinking with diagrams conference*.
- Wiggins, G. A. (2006). A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems 19(7)*.
- Wittgenstein, L. (2010). *Philosophical investigations*. John Wiley & Sons.
- Yannakakis, G. N., & Togelius, J. (2011). Experience-driven procedural content generation. *IEEE Transactions on Affective Computing, 2 (3)*.
- Yannakakis, G. N., Liapis, A., & Alexopoulos, C. (2014). Mixed-Initiative Co-Creativity. *Proceedings of the 9th Conference on the Foundations of Digital Games*

Biographical information

Antonios Liapis is a Lecturer at the Institute of Digital Games, University of Malta (UoM). He received his 5-year Diploma (2007) in Electrical and Computer Engineering from the National Technical University of Athens and the M.Sc. (2011) and Ph.D. (2014) in Information Technology from the IT University of Copenhagen. He does research on the crossroads of game design, artificial intelligence and computational creativity. More specifically, he explores the limits of computational input to the human-driven design process in computer-aided design tools. Beyond AI-assisted game design, his research pursuits revolve around procedural content generation, digital aesthetics, evolutionary computation, neuroevolution and constrained optimization. He has published over 40 international journal and conference papers in the aforementioned fields, and has won several awards. Moreover, he has led or participated in the design and development of several games of varying scope and for different target audiences, including two FP7 ICT projects.

Contact: an.liapis@gmail.com

Georgios N. Yannakakis is an Associate Professor at the Institute of Digital Games. His research interests lie at the crossroads of AI (computational intelligence, preference learning), affective computing (emotion detection, emotion annotation), advanced game technology (player experience modeling, procedural content generation, personalization) and human-computer interaction (multimodal interaction, psychophysiology, user modeling). He pursues research concepts such as user experience modeling and procedural content generation for the design of personalized interactive systems for entertainment, education, training and health. Prof. Yannakakis is one of the leading researchers within player affective modeling and adaptive content generation for games. He has pioneered the use of preference learning algorithms to create statistical models of player experience which drive the automatic generation of personalized game content. He has published over 170 journal and conference papers in the aforementioned fields. His work has been cited broadly and has appeared in Science Magazine and New Scientist among other venues.

Constantine Alexopoulos has worked as a research assistant for the University of Edinburgh since 2012, where he is involved in international projects. Currently he is working on C2Learn, an EU funded project. Previously he worked as a research fellow in argumentation for the project Archelogos, and its subsidiary programs: LogAnalysis and Elenchos. He also worked part time for a digital production company called DG Dimension, taking part in the organization and execution of multiple national and international exhibitions with cultural and historical themes. He holds an MA in Mental Philosophy from the University of Edinburgh, and his research interests include: argumentation and justification dynamic structures, philosophy of language, epistemology and ethics.

Phil Lopes has a B.Sc. and M.Sc. in Computer Science from the Faculty of Science of the University of Lisbon, where he focused on human interactive procedural music generation systems. Currently he is a PhD student at the Institute of Digital Games of the University of Malta, where he investigates computational creativity and procedural content generation methodologies at the interplay between audio and level design, and how audio can be classified by the different kind of emotions they evoke in human players. Phil is currently working on the Sonancia system, which consists of procedurally generating and sonifying levels for the videogame horror genre.