Fostering creativity in learning through digital games

C²Learn creative learning

Creative Emotional Reasoning Computational Tools Fostering Co-Creativity in Learning Processes

www.c2learn.eu

MIXED-INITIATIVE PROCEDURAL CONTENT GENERATION

C²LEARN PROJECT DELIVERABLE NO. D4.3.3 (PROTOTYPE)

Author(s): Georgios N. Yannakakis, Antonios Liapis, Phil Lopes, University of Malta, Malta.

Dissemination level: Public

The C²Learn project has been supported by the European Commission through the Seventh Framework Programme (FP7), under grant agreement no 318480 (November 2012 – October 2015). The contents of this document do not represent the views of the European Commission and the Commission cannot be held responsible for any use which may be made of the information contained therein. Responsibility for the information and views set out in this document lies entirely with the authors.

COOPERATION

© C²Learn Consortium, 2013. Reproduction is authorised provided the source is acknowledged.

DOCUMENT IDENTITY

Project category	Details
Deliverable code	D4.3.3
Full title	Mixed-Initiative Procedural Content Generation
Work package	WP4
Task	T4.3
Consortium partners leading	UoM
Consortium partners contributing	

DOCUMENT HISTORY

Version	Date	Handling partner	Description
1.0	13/10/2015	UoM	Initial draft
1.1	16/10/2015	UoM	Second revised draft
1.2	25/10/2015	UoM	Third revision
Final	30/10/2015	UoM	Final

Table of Contents

1 INTRODUCTION 4
2 The Mixed-Initiative Co-Creation: Background and Definition 4
2.1 Mi-CC in the literature 4
3 MI-CC within The C²Learn theory: the Human Initiative 5
3.1 Lateral Thinking and CER 5
3.2 CER and Diagrammatic Reasoning 6
3.3 Wise Humanising Creativity 7
4 Mixed-initiative co-creativity: The Computational Initiative 7
5 General aspects of mixed-initiative co-creativity 8
5.1.1 Dissimilarity Metrics 9
5.1.2 Computational Suggestions 11
5.1.3 MI-CC as a webservice 11
6 MI-CC principles and uses within the C²Space 12
6.1 C²Assistants 12
6.2 Creative Stories 13
6.3 4Scribes 13
6.4 Iconoscope 13
6.4.1 Content Representation 14
6.4.2 Fitness Functions (Aspects of MI-CC) 15
6.4.3 Suggestions (C²Assistants) 15
6.4.4 Iconoscope C²Assistants Webservice 16
7 Conclusions 17
References 18

1 INTRODUCTION

This document describes the **final** version of the mixed-initiative procedural content generation prototype.

This document briefly covers work related to mixed-initiative co-creation (MI-CC) and mixed-initiative procedural content generation (MI-PCG). It then places MI-CC both within the general context of fostering human creativity and the principles of the C^{2}Learn theory and places the use of the MI-CC prototype within the overall C^{2}Learn environment. The report proceeds to describe the MI-PCG prototypes developed in the up until month 36 of the project and details the basic functionalities of the prototypes that encapsulate the MI-PCG principles.

The final MI-PCG webservices are available here:
http://idg-c2learn-games.info:8080/IDGWeb/

2 THE MIXED-INITIATIVE CO-CREATION: BACKGROUND AND DEFINITION

Creating and designing with a machine: do we merely create together (co-create) or can a machine truly *foster* our creativity as human creators? When does such co-creation foster the *co-creativity* of both humans and machines? This deliverable focuses on the simultaneous and/or iterative process of human and computational creators in a mixed-initiative fashion within the context of games and gameful activities, aimed at fostering human creativity and attempts to draw from both theory and praxis towards answering the above questions. For this purpose, we first discuss the strong links between mixed-initiative co-creation and theories of human and computational creativity.

Computer-aided design (CAD) tools have introduced new creation practices through which the computer and the human user collaborate to create new artefacts – be they architectural designs, industrial components, toys or computer games. This deliverable identifies **mixed-initiative co-creation** (MI-CC) (Yannakakis et al., 2014) as the task of creating artefacts via the interaction of a human initiative and a computational initiative. In tune with the definition of mixed-initiative design (Carbonell, 1970), in MI-CC both the human and the computer *proactively* "make *complete* contributions to the problem solution", although the two initiatives do not need to contribute to the same degree. MI-CC, however, differs from other forms of co-creation, such as those occurring with the collaboration of multiple human creators or with the interaction between a human and non-proactive computer support tools (e.g. spell-checkers or image editors) or non-computer support tools (e.g. artboards or idea cards).

2.1 MI-CC IN THE LITERATURE

C^{2}Learn focuses on MI-CC within game and playful activities. Level editors such as the *Garden of Eden Creation Kit* (Bethesda, 2009) or game engines such as the *Unreal Development Kit* (Epic Games 2009) limit the computer's initiative to interpolations, pathfinding and rendering; while they are very efficient at speeding up game development tasks, human initiative is the sole driver in the creative process. On the other end of the scale, procedural content generators specialized to a type of artefact such as trees with *SpeedTree* (IDV 2002) or First Person Shooter levels with *Oblige* (Apted 2007) can create large amounts of game content but limit the human's initiative to choosing parameters for the generation algorithms; granted that the user has no control during the computer's generative process except before it starts (customizing its parameters) or after it concludes (editing the generated artefact), there is no actual co-creation between human and machine. For the task of game development (and level design in particular), mixed-initiative tools include *Tanagra* (Smith et al., 2011) which allows the human designer to specify the position of key platforms in a platformer level with the computational designer generating the remaining level topology, and *Sentient Sketchbook* (Liapis et a. 2013a) which allows human designers to edit a strategy game level while computational creators are simultaneously creating variations of the user's level.

3 MI-CC WITHIN THE C²LEARN THEORY: THE HUMAN INITIATIVE

The understanding of human creativity has relied on diverse philosophical, neuroscientific and psychological perspectives (see Chappell, K., 2011 and Craft, A., 2011) (see also D2.3.1.). While MI-CC can potentially be linked to several theories of human creativity, this deliverable focuses on aspects of lateral thinking and attempts to draw the direct connections between principles of diagrammatic cognitive reasoning and MI-CC and, to a lesser and rather indirect degree, wise humanising creativity.

3.1 LATERAL THINKING AND CER

Mixed-initiative co-creation is aligned with the general principles of *lateral thinking* (De bono, 1970) and *creative emotive reasoning* (Scaltsas, 2013) (see also D2.1.1 and D2.2.1), the latter being an instance and specialization of the former. Lateral thinking (De bono, 1970) is the process of solving seemingly unsolvable problems or tackling nontrivial tasks through an indirect, non-linear, creative approach. MI-CC realizes the very nature of lateral thinking which, as a creativity process, is boosted through (increasingly) constrained spaces of solutions (De bono, 1970). Co-creation with computational creators of visual art, content design, and visual concepts 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, 1970) (D2.1.1 and D4.1.1).

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. 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 re-framing and the random stimulus principles 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 an icon resembles the random stimulus that can act as a potentiator of creativity and cause an alteration of lateral thinking.

3.2 CER AND DIAGRAMMATIC REASONING

Diagrammatic reasoning can be defined as reasoning via the use of visual representations; a cognitive process which is enabled during visual map design, interaction design and visual art. These representations can include all forms of imagery incorporating visual features (object shape, size, colour, spatial orientation etc.) (Cheng, 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, 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, 2001).

Diagrammatic Lateral Thinking fuses the principles of diagrammatic reasoning and lateral thinking. Diagrammatic lateral thinking builds upon the extended mind theory (Clark, 1998) and 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 (e.g. designer'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 diagrammatic lateral thinking, the process of constructing a diagram (an image, a map, or an icon) is more important that the final product (Vile, 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 diagrammatic lateral thinking: MI-CC occurring through diagrammatic representations (e.g. in level design) offers visual (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. Diagrammatic lateral thinking, as MI-CC, does not necessarily embed transformational creativity processes as identified by (Boden, 2003). The MI-CC instance presented here realizes diagrammatic lateral thinking since co-creativity in game asset design and icon or map creation occurs mainly on the visual (diagrammatic) level, at least in the way images, shapes and maps are presented in the C^{2}Learn games considered. MI-CC expands the very notion of diagrammatic lateral thinking 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, 2013). In the case of mixed-initiative iterative design as realized within C^{2}Learn, MI-CC encapsulates both **analogical** and **visual** diagrammatic lateral thinking: the first by constraining the possibility space to artefacts of high quality and value for the given problem (as that is either defined by the teacher of the context under investigation) allowing learners to make analogies to context-specific qualities via diagrams; the latter by targeting visual diversity in the C^{2}Assistants suggestions it provides to the learner.

3.3 WISE HUMANISING CREATIVITY

Complementary to LT and DLT a case can be made for aspects of *wise humanising creativity* (WHC) and the 4Ps (see D.2.1.1). Within WHC, MI-CC arguably may put users within *engaged action* and improve the *participation* level through intuitive user interfaces. Beyond user participation (and within the theory of 4Ps), MI-CC may enhance the *possibility* space and *playfulness* of the human creator as it, respectively, offers opportunities for possibility thinking (via the suggestions offered by the computational creator) and provides an arena for self-creation.

4 MIXED-INITIATIVE CO-CREATIVITY: THE COMPUTATIONAL INITIATIVE

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, 2003), (Colton, 2008). Computational creativity, however, seeks creativity generated by, enhanced or fostered via algorithmic means. The computational creativity literature suggests that *value* (quality, usefulness or goodness) and *novelty* are the key elements characterizing a creative process (e.g. see (Boden, 2003)). An autonomous generative system is able to try out exhaustively many possible novel combinations of elements, often resulting in largely uninteresting outcomes or artefacts. For that very reason, computational creativity not only requires the generated artefacts to be novel, but also valuable. While other aspects of creativity have been discussed and proposed (such as surprise (Machado and 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 (e.g. set by the teacher), the creative process is ultimately of value for both given the problem constraints as those are set by either the human or an external observer (e.g. domain expert). Moreover, if the generative process of the machine searches within the 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, 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, which are suggested back to the human.

Computational creativity has been classified by Boden (1999) 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. (2001) provide the example of an unusual but legal chess move as often being more creative than changing the rules of chess.

5 GENERAL ASPECTS OF MIXED-INITIATIVE CO-CREATIVITY

Following Boden's (1999) and Ritchie's (2007) approach on taxonomies of software creativity which is inspired by aspects of human creativity we adopt three key classes of creativity aspects to assess the creativity of an artefact (map, diagram, icon) that is co-designed in a mixed-initiative fashion. The three aspects of creativity considered by Ritchie (2007) as utilised for artefact evaluation and computational search in MI-CC are as follows: (NB. The details of the specific metrics to be used under each C²Learn game activity where MI-CC is employed are discussed in the corresponding subsections below.)

Novelty: *To what extent is the produced item dissimilar to existing examples of its genre?*

Novelty in C²Learn is **assessed by the machine** (computationally) via various metrics across all artefacts generated in its genre, game context and (possibly) semantic context. Various expressions of **novelty drive the search** for new artefacts. The novelty, *n* (of artefact *i*) is calculated as follows:

$$n(i) = \frac{1}{k}\sum_{j=1}^{k} d_s(i, D_j)$$

Equation 1 - Novelty Function

where *k* is the number of visual artefacts (e.g. maps, graphs, drawings, icons) considered in the domain D, D_j the *j*-th artefact from domain D, and $d_s(i, j)$ is the domain-specific heuristic for calculating the "difference" (or dissimilarity) between artefacts i and j.

Quality: *To what extent is the produced item a high quality example of its genre?*

Quality (or value or goodness) in C²Learn is **assessed by other students and/or the teacher** – e.g. the artefacts they are liked/preferred given as a solution to a problem/dilemma or as an artefact that is generated out of a specific context. All artefacts generated in a diagrammatic C²Learn game are ranked by the students in terms of likeliness/preference. The same procedure can be followed by the teacher. **Quality does not necessarily drive** the search for new content but it can certainly do that if machine learning is added as an intermediate process which learns the mapping between peer-evaluations and generated artefacts (more on this on the corresponding sections of each game activity described in the following section). Quality may alternatively be defined by **constraints in search** as those are set by teachers *a priori* to the creation (e.g. the word *Love* should be represented by maximum of three shapes).

Typicality: *To what extent is the produced item an example of the artefact class in question?*

Typicality in C²Learn is **assessed by the machine** provided a **typical set of artefact(s) by the teacher**. Typicality **drives** the search of new content as long as the teacher has defined a typical set for the problem (and/or the semantic context) under investigation. The typicality, *t* (of artefact *i*) is calculated as follows:

$$t(i) = \frac{1}{k}\sum_{j=1}^{k} d_s(i, T_j)$$

Equation 2 - Typicality Function

where *k* is the number of individuals (visual artefacts) in the typicality set T, T_j the *j*-th typical individual and $d_s(i, j)$ is the domain-specific heuristic for calculating the “difference” between two individuals i and j.

According to Ritchie (2007) both typicality and quality will usually be assessed by human judgement and may therefore be partly or wholly subjective. We follow his suggestion and consider typicality as a metric derived from a teacher’s example set and quality as assessed through annotations or ranks of produced items (by learners and/or teachers). Note also that “novelty and typicality may well be related, since high novelty may raise questions about, or suggest a low value for, typicality” (Ritchie, 2007).

5.1.1 DISSIMILARITY METRICS

A number of relevant heuristics have being developed in order to calculate the difference (“visual dissimilarity”) between two distinct generated visual artefacts (d_s - see earlier section) such as icons. We are considering both top-down (i.e. ad-hoc designed metrics) and bottom-up (i.e. machine learned metrics) approaches, while also experimenting with variant combinations of the metrics presented below. The metrics presented below we are inspired by theories of visual perception (Arnheim, 2004) for the generation of aesthetically pleasing artefacts as introduced by (Liapis et al., 2012).

Please note that **we hereby present the final list of heuristics tested and used by MI-PCG procedures in C²Learn**. For a complete list of initial metrics the reader is referred to earlier versions of the deliverable (D4.3.1 and D4.3.2). The initial metrics were tested and refined, with the priority of creating more visually diverse suggestions. In that regard, the individual perimeter or concentration metrics (e.g. narrowness) were largely omitted due to Iconoscope's focus on multiple shapes; instead, the color, shape type and spatial arrangement of these component shapes is considered instead. The final choice of metrics was also driven by computational limitations of the quick response time requirement for Iconoscope suggestions, as pixel-to-pixel comparisons would render real-time calculation of icon difference impossible. It should also be noted that, in order to speed up the response of the webservice, each evolutionary process (each request for suggestions) chooses a single metric (randomly) as its target rather than combining them together into a weighted sum, which would lead to longer computation times for difference evaluation but also sub-par optimization behavior due to multiple objectives in the fitness score.

Colour Metric

The colour metrics compares the shapes of two different icons based on whether they share the same colours. For each different colour possessed by any shape in one icon which does not exist in any shape in the other icon, the difference metric is increased by one. The final score is normalized to the number of unique colours possessed by both icons' shapes.

Shape Metric

Similarly to the colour metric, the shape metric compares the shapes of two different icons based on whether they share the same shape types (for instance, whether there is any number of star shapes in both icons). For each different shape type possessed by any one icon which does not exist in the other icon, the difference metric is increased by one. The final score is normalized to the number of unique shape types possessed by both icons.

Distance Metric

The distance metric compares the shapes of two different icons based on whether shapes of one icon are in a similar position to any shapes in the other icon. To evaluate this, each shape in one icon evaluates the Euclidean distance with all shapes of the other icon: a measure of proximity is calculated from those distances. The total proximity of all shapes of one icon to the other is used as the distance score, which rewards shapes which share similar coordinates across the two icons.

Grouping Metric

The grouping metric is inspired from Gestalt theories of visual perception (Arnheim, 2004), and it compares whether the shapes in both icons are similarly grouped together (or dispersed). Grouping is calculated on each icon individually, by calculating the average distance between its shapes' coordinates. The difference metric then amounts to comparing the grouping scores of each icon: low scores means that both icons have shapes grouped together or dispersed on the canvas, while high scores means that one icon has shapes clustered in one part of the canvas while the other has shapes dispersed around the canvas.

Shape and Colour Metric

The shape and colour metric combines the notion of shape and colour, and compares the shapes of two different icons on the basis of whether there are shapes in both icons with the same shape type and colour. This provides more specific information on the appearance of each icon. For instance, this metric would evaluate whether both icons contain at least one yellow star, while the shape metric and colour metric would evaluate whether both icons contain a star (of any colour) or yellow shapes (of any type) respectively. Similar to the shape and colour metrics, for each shape in one icon which does not have any instances (of the same colour and shape type) in the other icon, the difference metric is increased by one. The final score is normalized to the number of shapes with unique shape/colour combinations possessed by both icons.

5.1.2 COMPUTATIONAL SUGGESTIONS

During the creation process the MI-PCG generates suggestions kick-started by the user's current interactions (icons, drawings) on the canvas/game etc. Depending on the C^{2}Assistant, the suggestions may be generated either through Novelty (i.e. Mad Scientist) or Typicality (i.e. Typical Tom or Progressive Petra). Note that quality suggestions (i.e. through Wise Oracle) are not generated but, instead, selected from a pool of previous icons created by any Iconoscope player.

Suggestions for novelty and typicality (or atypicality) are evolved artificially through a mutation-based genetic algorithm (no crossover is implemented in this prototype). The algorithm includes elitism by keeping half of the fittest shapes from the previous population into the next generation's population. Parent selection is performed via a fitness-proportionate roulette wheel scheme. The initial population of the genetic algorithm is seeded from the player's current icons. The shapes of the icons in each genotype of the population are mutated by changing their coordinates (moving them on the canvas), rotating or scaling them, changing their colour (to another colour among those allowed in Iconoscope), changing the shape type (e.g. from square to star), deleting the shape or cloning it (creating an identical shape and moving, rotating or scaling it). Evolution is carried out via a standard genetic algorithm, favouring the fittest individuals (where fitness is dependent on the assistant chosen) for selection and replacement.

Through empirical analysis it was observed that most metrics considered in section 6.2.3 converge at approximately 10-15 generations with a population size of 10 individuals which requires several seconds of computational effort. In order to speed up the evolutionary process (and due to the fact that the genetic algorithm provides suggestions for the user's next steps rather than final results), the final population size is 10 individuals which evolve for 5 generations. Needless to mention that smaller populations provide suggestions at a much faster rate, which can be considered if more rapid human interaction is required. On the other hand, such small population sizes yield solutions that generally suffer from lack of diversity. To provide an appropriate balance between computational effort and divergence (novelty) or convergence (typicality/atypicality) all metrics of 6.2.3 have been empirically tested thoroughly and optimized in order to improve the performance of the final version of MI-CC tools.

5.1.3 MI-CC AS A WEBSERVICE

The webservice is built around the **Representational State Transfer** (**REST**) architecture. The webservice communicates with other applications (e.g. with the Iconoscope and 4Scribes android applications) using data structures in the JSON (JavaScript Object Notation) format, which is a lightweight data-interchange format easily understandable by both humans and machines. All method calls from C^{2}Learn games to the MI-CC webservice are performed via POST requests. More details and examples of the method calls can be found in 6.4.5.

6 MI-CC PRINCIPLES AND USES WITHIN THE C^{2}SPACE

In the subsequent subsections we describe the use of MI-CC principles in C2Learn games considered in the Game design deliverable (D4.1.2).

6.1 C^{2}ASSISTANTS

Metrics for evaluating mixed-initiative co-creativity and driving the search of content to be presented to learners are of direct use to the C^{2}Assistants described in the Game Design Deliverable (4.1.2). The table below offers a reminder and an overview of the use of computational metrics (creativity aspects) from the various C^{2}Assistant personas. The detailed equations that drive the suggestions in each C^{2}Learn game activity are provided in the corresponding sections below.

MI-CC (Diagrammatic) Creativity Aspect	C^{2}Assistant (see also D4.1.2)
Novelty (of artefacts such as diagrams, maps, or icons produced via MI-CC)	**Mad Scientist**: The Mad Scientist would be the assistant that always proposes artefacts that maximize the (diagrammatic) novelty value (or sets of novelty values) of the artefact.
Quality (or **Value**) (of artefacts such as diagrams, maps, or icons produced via MI-CC)	**Wise Oracle:** The Wise Oracle shows students earlier highly-valued artefacts (from students and/or the teacher) under a specific context (game and semantic context). Artefacts are evaluated via ranked/rated/like annotations.
Typicality (of artefacts such as diagrams, maps, or icons produced via MI-CC)	**Typical Tom** and the **Progressive Petra**: We see two key C^{2}Assistants in relation to diagrammatic typicality a *conservative* **(Typical Tom)** C^{2}Assistant which proposes maximally-typical suggestions to the learner and a *progressive* **(Progressive Petra)** C^{2}Assistant which suggests maximally atypical content (maximum divergence from the typical set).
Other aspects/objectives	**Chaotic Kate:** This C^{2}Assistant offers either completely random diagrammatic suggestions or suggestions driven by

	randomly selected heuristics of diagrammatic difference, maps and icons such as e.g. balance and symmetry. These will be defined in each C^{2}Learn game activity.

6.2 CREATIVE STORIES

MI-CC has not a direct or indirect use in Creative Stories according to D4.1.2.

6.3 4SCRIBES

MI-CC is used *indirectly* in the 4Scribes game solely during the element/card distribution phase. 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. C^{2}Assistants are chosen at the start of the game for creating 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. The essential nature of the computational initiative of 4Scribes differs from that of Iconoscope (where asking C^{2}Assistants or using their suggestions is optional). Moreover, it is different in that it specifies the affordances of the players' 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. Acting as an initial disruptor, it provides the canvas for collaborative human creativity to draw upon.

As MI-CC in C^{2}Learn is *tightly coupled with diagrammatic reasoning* (and games/activities related to that CER component) we refrain from providing further details of MI-CC for 4Scribes in this deliverable. Instead the details of the indirect use of MI-CC in 4Scribes are provided in the Game Design Document (D4.1.2).

6.4 ICONOSCOPE

MI-CC is *directly* applicable to the Iconoscope game as it belongs to the class of diagrammatic reasoning games. A detailed description of the Iconoscope game can be found in D4.1.2. Herein we provide an overview of the basic game rules and functionalities.

In Iconoscope the educator picks a set of three concepts from a pre-defined set of concepts/ideas/words existent in the game as the input to the learners' tablets. Pre-defined terms may include anything from abstract concepts such as *love* and *freedom* to more specific properties such as *house* and *storm*. Each member of the group chooses **in secret** which part of the concept input to use in order to produce a new diagram out of the initial one (or its subcomponents), which expresses (communicates) the concept input, albeit with the above evaluation constraints in mind.

Each player (or group of collaborating players) can choose from a predefined palette of shapes and icons existent in the game. They can drag and drop, rotate, resize, colour existing shapes as well as add new shapes to the shapes suggested by the teacher (see Figure 16). . After a period of time has passed, the game is over and the players show their icon to the group for the purposes of voting. Passing the tablets around, other players (opponents) take turns observing the icon and choosing which of the three initial concepts it represents. Once each player has voted for each other player's icon (and thus each tablet reached the icon's creator), the voting phase is complete. Based on the number of opponents and their votes, a score is given to each player's icon. The scoring system rewards ambiguous icons which are however specific enough to be correctly guessed by at least one opponent. If all opponents guess the concept correctly, or if no opponent guesses the concept then the player loses and receives no points.

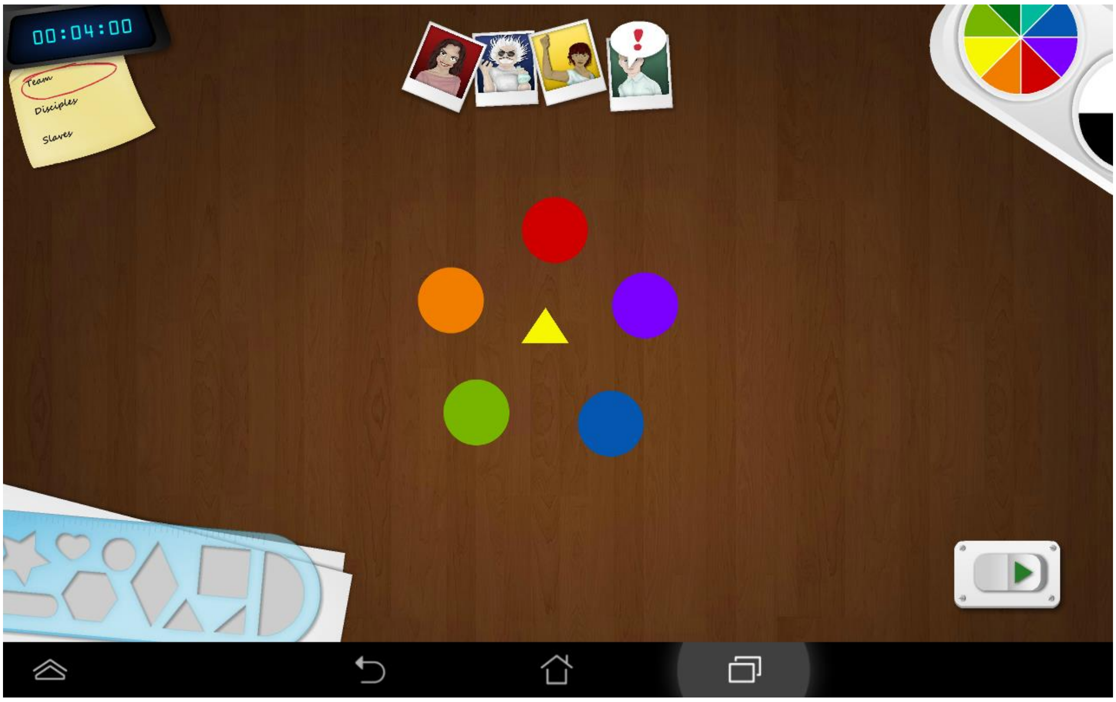

Figure 1 – A Screenshot of Iconoscope

6.4.1 CONTENT REPRESENTATION

The diagram represented in Iconoscope is composed of a number of shapes as chosen by the learner. Shapes are represented as follows:

- Shape type (e.g. left arrow, circle)
- Shape size
- X and Y coordinates of the shape's centre
- Rotation angle with respect to the X axis
- Colour of the shape in RGB
- Shape description: text tags which help designate properties of the shape:

examples include “jagged”, “square”, “curved”

6.4.2 FITNESS FUNCTIONS (ASPECTS OF MI-CC)

Fitness functions are used as heuristics for search towards novelty and typicality and the resulted generated content is suggested by the C²Assistants. Several fitness functions have been considered for searching for new content to be presented to learners as suggestions (following the principles of MI-CC). Iconoscope borrows from the fitness functions described above and considered several distance metrics presented in section 5.1.1. In addition, Iconoscope considered, examined and tested the following: *Colour of the map (diagram), Screen coverage, Size of shapes, Use of provided shapes Concentration of shapes (top, bottom, left, right, or centre), Jaggedness/Curvature/Squareness of shapes (on screen or in available objects display.* The final set of fitness functions is described in 5.1.1 and was determined after the game has been prototyped and tested with artificial and human users. The design decisions for the final fitness functions chosen were driven by the need for real-time response needs of Iconoscope and the visual diversity of the resulting concepts they produce.

6.4.3 SUGGESTIONS (C²ASSISTANTS)

During the game C²Assistants suggest novel (*Mad Scientist*), typical (*Typical Tom*), atypical (*Progressive Petra*), or valuable (*Wise Oracle*) diagrams to the learners to consider. Chaotic Kate is suggesting random mutations of the user’s diagram. An evolutionary computation process is followed where either existing diagrams are altered (shape type, size, colour, location) or new shapes are added to the diagram: both are achieved through variant mutation operators.

Novelty (Mad Scientist): the Mad Scientist C²Assistant optimizes for diagrammatic novelty and suggests corresponding diagrams for it. The Mad Scientist suggests a number of new diagrams that maximize the novelty score *n(i)* for icon *i*:

$$n(i) = \frac{1}{k}\sum_{j=1}^{k} d_m(i, \mu_j)$$

Equation 3: Novelty in Iconoscope

where d_m a difference heuristic among those in 5.1.1 (chosen randomly), μ_j is the *j*-th-nearest neighbour of *i* (within the current population), and *k* is the number of nearest neighbours considered (*k=5* for the iconoscope MI-CC). It is important to note that due to the low number of generators and the small population, novelty search does not retain previous novel individuals in a novel archive (due to the limited effect this would have in search).

Quality (Wise Oracle): All artefacts generated in Iconoscope ultimately receive a score at the end of a game session. The Wise Oracle C²Assistant suggests among the most valuable (highest scoring or winning diagrams) under a semantic theme (e.g. Love).

Typicality (Typical Tom and Progressive Petra): We see two ways C²Assistants affect diagram suggestions with regards to typicality. The Typical Tom C²Assistant offers suggestions that maximise the typicality from a given, typical, icon under a set of word contexts whereas the

Progressive Petra C²Assistant offers suggestions that maximize the divergence from the typical set (atypicality). Typicality considers the diagram originally provided by a human author (a teacher or the game's designers) under a word (semantic) context. Given that ground truth of diagrams, the algorithm attempts to minimize (for Typical Tom) or maximize (for Progressive Petra) the fitness score *t(i)* of icon *i*:

$$t(i) = d_m(i, T)$$

Equation 4 - Typicality Iconoscope

where d_m a difference heuristic among those in 5.1.1 (chosen randomly), and T is the typical icon for the concept selected by the player.

6.4.4 ICONOSCOPE C²ASSISTANTS WEBSERVICE

As described in 5.1.4, Iconoscope assistants are handled via a webservice using the REST architecture, which allows for inter-operability with most types of software. When requesting C²Assistants' suggestions, the Iconoscope application must provide as input (via a POST request) a JSON object containing information needed for generating the output. The C²Assistant chosen is determined by the web address where the request is sent. The webservice processes this input and generates output in the form of up to *four* suggestions (in different canvas objects) which is returned to the calling Iconoscope application as a JSON object.

An example JSON request from Iconoscope to the MI-CC webservice is:


```
{"App": {"AllWords": ["war","cunning", "threat"],"AppWidth":
1200,"DeviceWidth": 1920,"SelectedWord": "cunning","DeviceHeight":
1200,"QuestID": 1,"Scale": 1.3333333333333333,"MissionID":
1,"AppHeight": 900,"ChallengeText": "Dare you look deeper into War,
Cunning and Threat? Prove it, outsmart the others!","Points":
123456},"Canvas": {"ID": "user1_2015-01-07 11:18:39","Shapes": [
{"PosY": 600,"Rotation": 0,"ScaleX":
1.333333333333333,"BoundingBoxGlobal":
"x:708,y:508,w:184,h:184","ScaleY": 1.333333333333333,"ID":
"Square","Depth": 0,"Color": "0xcccccc","PosX":
800,"BoundingBoxLocal": "x:-69,y:-69,w:138,h:138"},
{"PosY": 600,"Rotation": 0,"ScaleX":
1.333333333333333,"BoundingBoxGlobal":
"x:712,y:441.35,w:176,h:317.3","ScaleY": 1.333333333333333,"ID":
"Diamond","Depth": 1,"Color": "0xf27e00","PosX":
800,"BoundingBoxLocal": "x:-66,y:-119,w:132,h:238"}
]}}
```


And the response of the Mad Scientist C²Assistant (within the MI-CC webservice) back to Iconoscope is:


```
[{"fitness":0.6666666666666666,"ID":"Canvas 0","shapes":[
{"Color":"0xcccccc","Depth":0,"ID":"Square","PosX":900,"PosY":604,"Ro
tation":0,"ScaleX":1.333333333333333,"ScaleY":1.333333333333333}]},
{"fitness":0.6666666666666666,"ID":"Canvas 1","shapes":[
{"Color":"0xcccccc","Depth":0,"ID":"Square","PosX":900,"PosY":604,"Ro
tation":0,"ScaleX":1.333333333333333,"ScaleY":1.333333333333333}]},
{"fitness":0.5333333333333333,"ID":"Canvas 2","shapes":[
```

{"Color":"0xf27e00","Depth":0,"ID":"Diamond","PosX":800,"PosY":600,"Rotation":0,"ScaleX":1.3333333333333333,"ScaleY":1.3333333333333333}]},
{"**fitness**":0.5333333333333333,"ID":"Canvas 3","shapes":[
{"Color":"0xf27e00","Depth":0,"ID":"Diamond","PosX":800,"PosY":600,"Rotation":0,"ScaleX":1.3333333333333333,"ScaleY":1.3333333333333333},
{"Color":"0xf27e00","Depth":1,"ID":"Diamond","PosX":800,"PosY":600,"Rotation":0,"ScaleX":1.4588637592871418,"ScaleY":1.4588637592871418}]}
]

7 CONCLUSIONS

The intention of this document has been to describe the core components of the final version of the C²Learn mixed-initiative procedural content generation prototype. The principles of mixed-initiative co-creation are directly applicable to any C²Learn game activity that involves diagrammatic CER. Mixed-initiative co-creation can potentially lead to skill acquisition by providing indications (such as novelty, typicality and quality) that could potentially inform the presence of creative patterns, allowing to consistently shift the tool's output to something that is new and surprising for the users. Tools as such revolutionize the use of creative tools within an educational setting and have already been integrated to the final game design of C²Learn (D4.1.2). MI-PCG forms a necessary input to the final game design (D4.2.1) and the game prototypes (D4.4.x) of C²Learn.

Versions of the MI-PCG tool can be instantiated by the C²Learn games that make use of it. The current tool can be generic across any image-based creation task within the C²Learn game environment either realising aspects of independent diagrammatic construction games or forming part of C²Learn games through C²Assistants.

References

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

Beaney, M. 2005. Imagination and creativity. *Open University Milton Keynes, UK*.

Blizzard Entertainment. 1998. StarCraft. [Computer Game]

Blizzard Entertainment. 1996. Diablo. [Computer Game]

Boden, M. 1999. Computational models of creativity. In Robert J. Sternberg (Ed.), Handbook of Creativity, 351–373.

Boden., M. 2003. The creative mind: Myths and mechanisms. *Routledge*.

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

Bundy, A. 1994. What is the difference between real creativity and mere novelty? *Behavioural and Brain Sciences*, 17(3):533 – 534. Open peer commentary on (Boden, 1990).

Bethesda Softworks. 2011. Creation Kit. [Game Content Editor]

Carbonell J. R.. 1970 Mixed-Initiative Man-Computer Instructional Dialogues. PhD thesis, Department of the Navy, 1970.

Chappell, K., & Craft, A. (2011). Creative learning conversations: producing living dialogic spaces, Educational Research, 53:3, 363-385;

Cheng, P., Lowe, R., and Scaife, M. 2001. Cognitive science approaches to understanding diagrammatic representations. *Artificial Intelligence Review, 15(1-2):79–94*.

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. *Proceedings of the AAAI Spring Symposium on Creative Systems*, Stanford University.

Craft, A. (2012) Childhood in a digital age: creative challenges for educational futures, London Review of Education, 10:2, 173-190;

De Bono, E. 1970. Lateral thinking: creativity step by step. Harper & Row.

Epic Games. 2009. The Unreal Development Kit. [Game Engine]

Hofstadter, D. 1994. Fluid Concepts and Creative Analogies. HarperCollins, New York, USA.

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.

Larkin, J., and Simon, H. 1987. Why a diagram is (sometimes) worth ten thousand words. Cognitive science, 11(1):65–100.

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

Lehman, J., and Stanley, K.O. 2012. Beyond open-endedness: Quantifying impressiveness. In: *Proceedings of Artificial Life Thirteen (ALIFE XIII)*.

Liapis, A., Yannakakis, G. N. and Togelius, J. 2013a. Sentient sketchbook: Computer-assisted game level authoring. In *Proceedings of the 8th International Conference on Foundations of Digital Games*.

Liapis, A., Yannakakis, G.N., and Togelius, J. 2013b. Towards a Generic Method of Evaluating Game Levels, in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*.

Liapis, A., Yannakakis, G. N., and Togelius, J. 2013c. Enhancements to constrained novelty search: Two-population novelty search for generating game content. In *Proceedings of the fifteenth annual conference on Genetic and evolutionary computation conference* (pp. 343-350). ACM.

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), pp. 213-228.

Lubart, T. 2005. How can computers be partners in the creative process: classification and commentary on the special issue, International Journal of Human-Computer Studies, vol. 63, no. 4-5, pp. 365—369.

Machado, L., and Cardoso, A. 2001. Modeling forms of surprise in an artificial agent. *Structure, 1(C2):C3*.

Machado, L., and Cardoso, A. 1999. Towards artificial forms of surprise and curiosity. In *Proceedings of the European Conference on Cognitive Science, S. Bagnara, Ed* (pp. 139-144).

Pease, A., Winterstein, D., and Colton, S. 2001. Evaluating machine creativity. In *Workshop on Creative Systems, 4th International Conference on Case Based Reasoning*, pp. 129-137.

Pind, J. 1994. Computational creativity: What place for literature? Behavioural and Brain Sciences, 17(3):547 – 548. Open peer commentary on (Boden, 1990)

Prosecco Network Coordination action. Available at: http://prosecco-network.eu/

Ramachandran, V. S., and Hirstein, W. 1999. The science of art: a neurological theory of aesthetic experience. *Journal of consciousness Studies* 6, 15–51.

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

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

Smelik, R. M., Tutenel, T., de Kraker, K. J., and Bidarra, R. 2011. A declarative approach to procedural modeling of virtual worlds. *Computers & Graphics* 35(2):352–363.

Smith, G., Whitehead, J., and Mateas, M. 2011. Tanagra: Reactive planning and constraint solving for mixed-initiative level design. *IEEE Transactions on Computational Intelligence and AI in Games.*

Sloman, A. 2002. Diagrams in the Mind? *Springer*.

Vile, A., and Polovina, S. 1998. Thinking of or thinking through diagrams? The case of conceptual graphs. *In Thinking with Diagrams Conference, The University of Wales, Aberystwyth*.

Yannakakis G. N., Liapis, A., and Alexopoulos C. 2014. Mixed-Initiative Co-Creativity. In *Proceedings of the 9th International Conference on Foundations of Digital Games*.

Valve Corporation. 2000. Counter Strike. [Computer Game]