

Sentient Sketchbook: Computer-Aided Game Level Authoring

Antonios Liapis
Center for Computer
Games Research
IT University of Copenhagen
Copenhagen, Denmark
anli@itu.dk

Georgios N. Yannakakis
Institute of Digital Games
University of Malta
Msida, Malta
georgios.yannakakis
@um.edu.mt

Julian Togelius
Center for Computer
Games Research
IT University of Copenhagen
Copenhagen, Denmark
juto@itu.dk

ABSTRACT

This paper introduces the Sentient Sketchbook, a tool which supports a designer in the creation of game levels. Using map sketches to alleviate designer effort, the tool automates playability checks and evaluations and visualizes significant gameplay properties. This paper also introduces constrained novelty search via a two-population paradigm for generating, in real-time, alternatives to the author's design and evaluates its potential against current approaches. The paper concludes with a small-scale user study in which industry experts interact with the Sentient Sketchbook to design game levels. Results demonstrate the tool's potential and provide directions for its improvement.

Categories and Subject Descriptors

I.2.1 [Artificial Intelligence]: Applications and Expert Systems

General Terms

Algorithms

Keywords

Procedural content generation, Computer-aided design, Novelty search, Genetic algorithms, Level design

1. INTRODUCTION

Over the last twenty years, computer games have grown from a niche market targeting young adults to an important player in the global economy, engaging millions of people from different cultural backgrounds [3]. As the number and the size of computer games continues to increase, game companies handle rising demand by expanding their cadre, compressing development cycles and reusing code or assets. To limit development time and reduce the cost of content creation, commercial game engines and procedural

content generation are popular shortcuts among game companies. Low-level engine features such as physics simulation and game-independent content such as vegetation are often handled by third-party middleware such as *Havok Physics* (Havok, 2000) and *SpeedTree* (IDV, 2002) respectively. For content more contingent on a game's theme, mechanics and quests, most commercial game engines such as the *Unreal Development Kit* (Epic Games, 2009) assist in the fast and intuitive creation of game levels and scripts via game editors which automate mundane tasks such as pathfinding.

This paper presents a tool, named Sentient Sketchbook, which supports the design of map sketches via an intuitive interface, via real-time feedback regarding the map's playability or balance, and via the suggestion of alternative map designs. The tool assists the level designer as it automatically tests maps for playability constraints, calculates and displays navigable paths, evaluates the map on gameplay properties and adds details to the coarse map sketch. Additionally, the tool aims to inspire the designer through the suggestion of map alternatives generated via genetic search, which is novel as it enhances the constrained optimization capabilities of novelty search via the feasible-infeasible two-population paradigm. We assume that designers are driven by certain objectives pertaining to strategic gameplay (such as game pace and player balance) in their attempt to design a game level. However, it is likely that not all designers follow the goals or principles implemented by the set of objectives; we argue that, in such cases, the principles of novelty search could be useful in providing alternative suggestions. These hypotheses will be tested with the presented version of the Sentient Sketchbook.

The Sentient Sketchbook tool addresses certain limitations of mixed-initiative design by relying on simple map sketches to counter user fatigue and designer biases reported in the literature [9]. This paper extends previous work [11] which presented objective functions for evaluating strategy game levels and explored their optimization behavior. This paper focuses on the user interactions afforded by the tool and on the automatic generation of maps via feasible-infeasible novelty search. The former is evaluated via a small-scale user study with industry experts and the latter is evaluated via controlled experiments with no human interaction.

2. RELATED WORK

The Sentient Sketchbook combines computer-aided design with the suggestion of alternative game levels generated procedurally via novelty search. A short overview of related

work is presented below.

Computer-Aided Design

Since its early stages of development, the computer was expected to assist in solving engineering problems by being involved in the creative design process and by automating tedious tasks. Computer-aided design tools have often been identified by their dual role as “the designer’s slave” [16] — performing simulations, analysis and constraint satisfaction tests — and as advisors when certain requirements are not met. As computers are becoming efficient at performing the former role, more researchers focus on the latter: the most ambitious role for the computer is that of the “colleague” [14]. According to Lubart [14], as a colleague a computer should contribute equally to the design discourse but could also incite creativity via “semi-random search mechanisms to generate novel, unconventional ideas”.

In commercial games and game-like applications, computer-aided design speeds up the development process in the form of game editors. Game editors use an intuitive graphic interface, allowing designers with little programming experience to script behaviors and create content, usually as part of a game level. Many of these tools ship with the final game, allowing end-users to generate content which increases replayability and fan loyalty. *Modding* via the provided game editors has often transcended the original game’s concept and mechanics, leading to new game titles such as *Counter-Strike* (Valve, 2000) or new subgenres such as Multiplayer Online Battle Arenas. Over the years, game editors have become very sophisticated, driven by a desire to support the modding community or to reuse code across products. As an example, the *Unreal Development Kit* supports landscape sculpting, asset organization, scaling rendering accuracies and code-free visual scripting. On the other hand, game-specific editors such as the Creation Kit of *Skyrim* (Bethesda Softworks, 2011) allow less customization but offer game-tailored easy-to-use automations such as leveled item lists, navigation path generation and quest scripting.

Procedural Game Content Generation

The game industry has often relied on the procedural generation of game content during playtime to enhance the unexpectedness of the player’s experience and increase the game’s replayability value. From early games such as *Rogue* (M. Toy and G. Wichman, 1980) and *Elite* (Acornsoft, 1984) to contemporary titles such as *Torchlight 2* (Runic, 2012) and *Civilization V* (Firaxis, 2010), gameworld and level creation has been the principal application of procedural content generation (PCG); other applications include the creation of enemies as in *Darkspore* (Maxis, 2011) and weapons as in *Borderlands* (Gearbox, 2009). Although academic interest in PCG is relatively new [24], the majority of PCG researchers challenge the mostly random generative algorithms used in the game industry. Whether generating platformer levels [21], mazes [1], board games [2], racing tracks [22], weapons [4] or spaceships [13], most projects within academia attempt to control the algorithms’ stochastic processes via constraints [12], objective functions [23] and predicted or reported player experience [25]. In the domain of computer-aided level design, the most promising results have been accomplished through constraint solvers, where the constraints are explicitly described by the designer [19, 20].

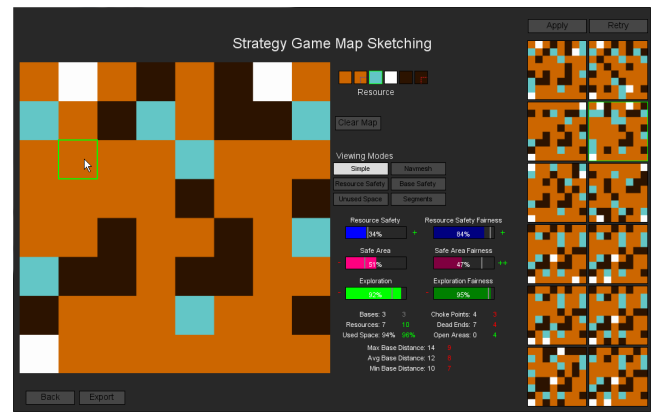


Figure 1: The Sentient Sketchbook tool during a design session. To the left is the sketch editor, far to the right are the automatically generated map suggestions; between these elements is the tile palette, the map display menu and an overview of map’s fitness dimensions and metrics.

Novelty Search and Constrained Optimization

Genetic algorithms traditionally optimize a population of individuals by selecting the fittest according to a *fitness function* [5]. The solutions discovered by a genetic algorithm are sensitive to the design of the fitness function, especially if this function is ill-defined, difficult to quantify, or subjective. Novelty search [8] is a recent approach to genetic search, replacing optimization towards a quantifiable objective with optimization towards the population’s diversity; it argues that by stimulating exploration of the search space, unforeseen high-quality solutions can be discovered. To better control exploration, the notion of *minimal criteria novelty search* was introduced in [7], where solutions which do not satisfy certain criteria are assigned a fitness of 0, severely limiting their chances of reproducing. In constrained evolutionary optimization, it is argued that infeasible solutions, especially those on the border of feasibility, contain valuable information which can guide search towards the discovery of feasible individuals [17, 15]. The feasible-infeasible two-population GA (FI-2pop GA) paradigm [6] optimizes infeasible individuals, stored in a different population than feasible ones, towards the border of feasibility in order to increase their likelihood of generating feasible offspring. While PCG projects use FI-2pop to optimize the population of feasible individuals according to an objective function [21, 12, 11], this paper contends that novelty search can be applied to the feasible population, generating diverse solutions which satisfy the subjective, unpredictable desires of the tool’s users.

3. SENTIENT SKETCHBOOK

The Sentient Sketchbook is a tool which allows a designer to create low-resolution map sketches. In this experiment the map sketches are abstractions of levels used in successful strategy games such as *Starcraft* (Blizzard, 1998), although other types of games can easily be represented in a similar sketch format. A map sketch consists of a small number of tiles (see Figure 1); tiles can be passable, impassable, player bases or resources. The map layout assumes that each player starts at a base and gathers resources to produce units which move through passable tiles to attack the opponent’s base.

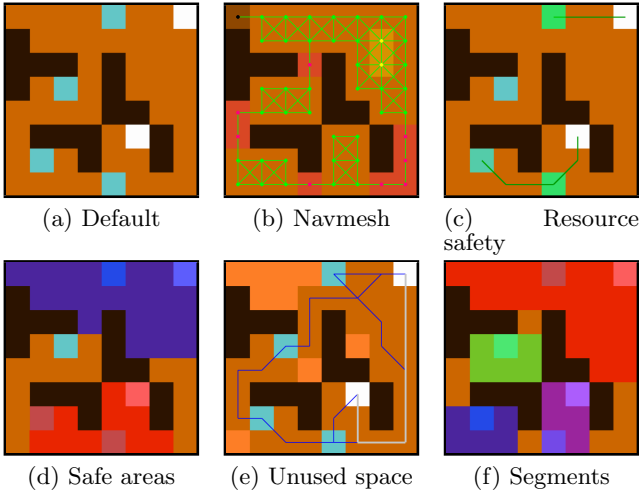


Figure 2: Different map displays on a small two-player map sketch: *Default* displays passable tiles (light brown), impassable tiles (dark brown), resources (cyan) and player bases (white). *Navmesh* displays the passable pathways on the map and the location of choke points (red), dead ends (black) and open areas (yellow); *Resource safety* displays the safety value of each resource (in shades of green) and connects bases to their safe resources. *Safe areas* shows the tiles close to each base (here in red and blue) which are considered safe. *Unused space* shows all shortest paths between bases and resources, and highlights leftover unused space (orange). *Segments* shows the passable areas (in different colors) which are surrounded solely by chokepoints.

The Sentient Sketchbook assists the design process as it tests maps for playability constraints, evaluates the map on gameplay properties, supports informative map displays and adds details to the coarse map sketch; these functionalities are presented in Section 3.1. The tool also aims to contribute to the design dialog by generating map suggestions for the user; the generative algorithms are presented in Section 3.2.

3.1 Map Editor

A graphical interface has been developed to allow a user to manually edit a map sketch (see Figure 1). The map editor allows the designer to place tiles on the map; while the designer edits the map, the tool tests it for *playability*. A map is playable (or feasible) if all resources and bases are reachable from any other base or resource.

Playable maps are evaluated on quantifiable measures of quality regarding game pace and player balance. Game pace is affected by the location of a player’s base: slow, defensive play is afforded if the base has many nearby resources, is surrounded by a large easily controllable area, and cannot be easily reached by other players. These game patterns are captured as *resource safety* (f_{res}) which measures the number of resources which are close to only one base, *safe area* (f_{saf}) which measures the number of tiles around a base which are close only to that base, and *exploration* (f_{exp}) which simulates how difficult every base is to find starting from all enemy bases. Exploration is simulated via a flood fill algorithm, which stops when an enemy base is discovered; the filled area is compared with the number of passable tiles on the map. Player balance is evaluated via *resource safety*

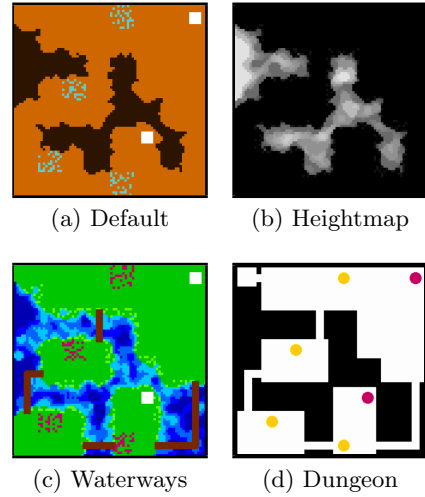


Figure 3: Visualizations of the final map, which offers a higher-detail view of the map sketch in Figure 2. The *Default* visualization adds detail to resource tiles and impassable tiles, while *Heightmap* elaborates the impassable regions of *Default*. *Waterways* treats impassable tiles and chokepoints as water and replaces chokepoints with bridges, while *Dungeon* divides the passable segments and dead ends into rooms and treats chokepoints as corridors.

balance (b_{res}), *safe area balance* (b_{saf}) and *exploration balance* (b_{exp}), which have high values if all bases have equally safe resources, a similar number of safe tiles and are equally difficult to find from enemy bases, respectively.

The six *fitness dimensions* presented above are re-calculated after every user interaction and displayed on-screen. If a map has less than two bases or unreachable bases, all fitness dimensions display “N/A” to indicate infeasibility; if resources are missing or unreachable, only the resource safety and resource safety balance dimensions display “N/A”.

In addition to these fitness dimensions, several other *metrics* are computed and displayed via the map editor, mostly pertaining to map navigation. The displayed metrics include the number of bases and resources on the map, the shortest paths’ length between bases, the percentage of used space¹ and the number of chokepoints, dead ends and open areas².

The editor allows the designer to switch between different map displays, which visualize navigational information and other properties (see Figure 2). At any time, the designer can switch to the final map view which displays the complete map on which a strategy game can be played (see Figure 3). The visualizations of the final map show how map sketches can be interpreted in different ways and how they can be applied to other game genres. These final maps (excluding Dungeon) are constructed via random processes and cellular automata to create an organic-looking map which retains all the properties (chokepoints, paths) of the coarse map sketch. Manual editing, evaluation and evolution are all done on the sketch level, making computations such as pathfinding easier and reducing the effort required to design a complete map.

¹ *Used Space* consists of passable tiles which are on a shortest path between any two bases or any base and any resource.

² *Chokepoints* are tiles with two neighboring (passable) tiles, *dead ends* are tiles with one neighboring tile and *open areas* are tiles with the maximum number of neighboring tiles.

3.2 Map Suggestions

In order for the Sentient Sketchbook to approximate the role of a colleague, it should not only evaluate the user’s current design but also provide new and unexpected alternatives to it. Through these *suggestions*, the tool intends to challenge the user’s current design focus and provide unforeseen alternatives to achieving the designer’s goals. The map suggestions are updated while the user edits the map, using the current map’s appearance as their starting point. Currently map suggestions are generated via genetic algorithms (GAs) performing constrained optimization to maximize either the maps’ scores in the fitness dimensions described in Section 3.1 or to maximize the maps’ diversity. Six genetic algorithms, running on separate threads, optimize maps in a single fitness dimension each; the best evolved individual of each GA is included in the map suggestions. A seventh genetic algorithm evolves maps via novelty search; the six most different evolved maps in this GA are also included in the map suggestions. Every genetic algorithm runs for 10 generations, and optimizes a total of 10 individuals which initially consist of mutations of the authored map sketch. The 12 map suggestions (six evolved via objective-based search and six evolved via novelty search) are displayed, as thumbnails, on the edge of the tool’s window; maps identical with the user’s map or with each other are omitted. The user can at any time select a map suggestion and replace their current sketch with it. While a suggestion is selected, its scores in the fitness dimensions and its other metrics are displayed along the relevant scores in the current sketch, clearly indicating which scores increase or decrease. This extra feedback mechanism is expected to reduce the designer’s cognitive load of closely inspecting each of the 12 map suggestions.

In order to evolve map suggestions, each map is encoded as an array of integers: each integer represents a tile’s type (passable, impassable, base or resource). All maps are tested for feasibility constraints: feasible maps must have an appropriate number of bases and resources and all resources and bases must be reachable from any other base or resource. The maps are adjusted according to the feasible-infeasible two-population GA (FI-2pop GA) paradigm [6] which evolves two populations, one with feasible maps and the other with infeasible maps. The infeasible population optimizes its members towards minimizing the distance from feasibility, which in this case is the number of unreachable bases and resources and the number of excess (or missing) bases or resources. As the infeasible population converges to the border of feasibility, the likelihood of discovering new feasible individuals increases. Feasible offspring of infeasible parents are transferred to the feasible population, boosting its diversity (and vice versa for infeasible offspring). The feasible population can be optimized to maximize different objective functions or to perform novelty search. Optimizing feasible individuals to maximize the fitness dimensions presented in Section 3.1 has been explored in [11]; the innovation in this paper is the evolution of the feasible population via novelty search. Following the novelty search paradigm [8], evolution optimizes feasible maps towards maximizing a function ρ corresponding to the average distance of the k most similar maps in the population and in an archive of novel maps. The archive stores the l highest scoring individuals in the population (in terms of ρ) and is reset at the start of every run of the evolutionary algorithm. Throughout this study, $k = 20$ and $l = 5$. The fitness score

$\rho(i)$ for individual i is calculated as:

$$\rho(i) = \frac{1}{k} \sum_{j=1}^k \text{dist}(i, \mu_j) \quad (1)$$

where μ_j is the j -th-nearest neighbor of i (within the feasible population and in the archive of novel individuals). Distance $\text{dist}(i, j)$ between maps i and j is calculated as the number of tiles which are not of the same type (passable, impassable, base or resource) for both maps at the same coordinates.

Both feasible and infeasible populations evolve via fitness-proportionate roulette-wheel selection of parents; parents are recombined via two-point crossover. Mutation may occur on parents (5% chance) or offspring (1% chance); during mutation, either the map is rotated by 180° (10% chance) or a portion of the map’s tiles (between 5% and 20%) are altered (swapped with a neighboring tile, transformed from impassable to passable or from passable to any other tile type). This aggressive mutation scheme is appropriate for the short evolutionary runs of map suggestions.

4. EXPERIMENTS

The Sentient Sketchbook tool introduces constrained novelty search via the FI-2pop paradigm, and uses Feasible-Infeasible Novelty Search (FINS) to generate map suggestions while a user edits the map. Before the generated map suggestions can be evaluated by human designers, the optimization behavior of FINS must be evaluated in a number of controlled experiments and compared against minimal criteria novelty search (MCNS) which has been used for the constrained optimization of novelty [7]. The experiments presented in this section use a random initial population of 100 individuals (including feasible and infeasible) and run for 100 generations. Such a large population size and number of generations can’t be afforded for evolving suggestions in real-time with the Sentient Sketchbook, but are used to test the robustness of the optimization behavior. Both MCNS and FINS are compared with objective-driven constrained optimization using a standard GA and FI-2pop GA; both MCNS and the standard GA assign a fitness score of 0 to infeasible maps. The argument for allowing infeasible individuals to survive and reproduce in a separate population is twofold: a) while no feasible individuals exist, infeasible parents move closer to the border of feasibility and the discovery of feasible individuals is much more likely than with random search, and b) infeasible parents can eventually generate unforeseen feasible offspring, increasing diversity in the feasible population. The performance metrics are therefore: a) the number of feasible individuals and b) the *diversity* of the feasible population, measured as the average tile difference between all pairs of feasible maps in the population. The experiments test the optimization of a small map where constraints can be easily satisfied, and that of a large map where the discovery of feasible individuals is unlikely.

The first experiment tests the optimization progress of a map sketch of 8 by 8 tiles, where 2 player bases and 4 to 10 resources must be present; due to the small map size and few bases, it is likely that such maps are feasible even when generated randomly. Figure 4 displays the progress of the feasible population’s diversity and size for MCNS, FINS, standard GA and FI-2pop GA; the standard GA and FI-2pop GA optimize the fitness dimension of resource safety, although their performance (in terms of the metrics being

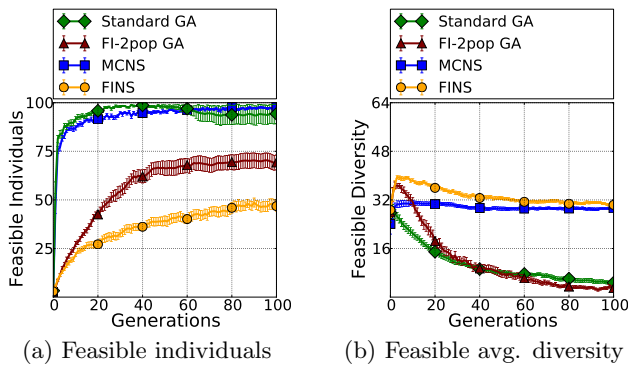


Figure 4: The progress of the number of feasible individuals and their average diversity when small, simple maps are optimized via different constrained optimization methods. Displayed values are averaged across 20 independent runs, and error bars denote the standard error.

evaluated) is similar for all fitness dimensions. The optimization behavior of FI-2Pop GA in terms of the fitness dimensions presented in Section 3.1 is explored in [11].

Observing the number of feasible individuals for all approaches in Figure 4a, it is immediately obvious that single-population approaches (MCNS and standard GA) swiftly generate a large number of feasible individuals (almost 100% of the population). Since the feasible space is relatively large, the occurrence of feasible individuals in the (random) initial population is very likely. Infeasible individuals have a fitness of 0 for MCNS and standard GA and are swiftly killed in favor of fitter feasible ones; on the other hand, FI-2pop GA and FINS retain the most promising infeasible individuals and therefore the number of feasible individuals is relatively small. For FINS and FI-2pop GA the number of feasible individuals grows slowly as infeasible individuals converge to the border of feasibility, increasing migration from the infeasible to the feasible population. Novelty search is more likely to create an infeasible offspring from feasible parents, since it encourages exploration of the search space which may lead back to infeasible solutions; this increases migration from the feasible population to the infeasible and explains the lower number of feasible individuals for FINS.

Observing the diversity of the feasible population for all approaches in Figure 4b, it is expected that the diversity of feasible individuals drops swiftly for objective-driven approaches (FI-2pop GA and standard GA), as the population converges to a few highly fit solutions. Unsurprisingly, novelty search approaches maintain or increase their feasible individuals’ diversity, as they optimize towards the same heuristic for distance used in the diversity metric. FINS maintains a higher diversity than MCNS, which can be attributed to the constant migration of novel feasible individuals from the infeasible population. This is most pronounced in early stages of evolution, when both FINS and FI-2pop GA have higher diversity values than their single-population counterparts due to infeasible individuals becoming feasible.

The second experiment tests the optimization progress of a map sketch of 16 by 16 tiles, where 8 player bases and 12 to 30 resources must be present; the numerous bases and resources are much more likely to be unreachable in this large map, and feasible individuals are less likely to be discovered

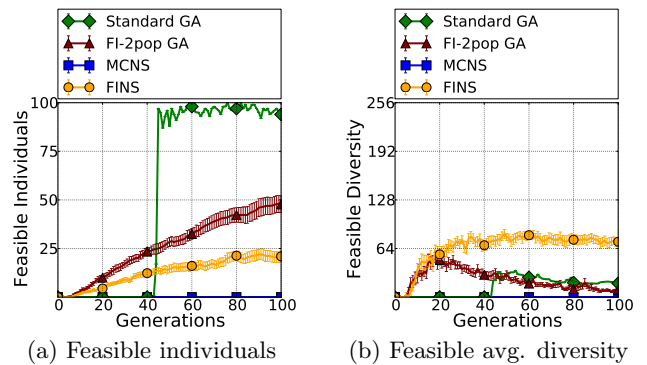


Figure 5: The progress of the number of feasible individuals and their average diversity when large, complex maps are optimized via different constrained optimization methods. Displayed values are averaged across 20 independent runs, and error bars denote the standard error. Note that for Standard GA, values refer to the single run on which a feasible individual was discovered.

Method	Runs	First Generation
Standard GA	1	44 (0.0)
FI2pop GA	20	10.9 (2.1)
MCNS	0	-
FINS	20	11.15 (1.6)

Table 1: The number of runs (out of 20) in which a feasible individual was discovered when optimizing large, complex maps, and the first generation in which this discovery occurred. The first generation is averaged across runs in which an individual was found, and the first generations’ standard error is shown in parentheses.

via random search. Figure 5 displays the progress of the feasible population’s diversity and size for MCNS, FINS, standard GA and FI-2pop GA — the latter two optimizing the dimension of resource safety. No feasible individuals were present in the initial population on any run and for any approach. For MCNS and standard GA, all members of the initial population have a fitness of 0 and the resulting random search makes discovery of a feasible individual both unlikely and unpredictable. Table 1 displays the number of runs that any feasible individual was discovered, and the average generation on which it occurred. Since both FI-2pop approaches (FINS and FI-2pop GA) perform an objective-driven optimization of the infeasible population towards the border of feasibility, random search is avoided and a feasible individual is consistently discovered in all runs within a few generations. Once a feasible individual is found within single-population approaches, the entire population converges to it and the number of feasible individuals remains high. For FI-2pop approaches, the fact that feasible parents are very likely to generate infeasible offspring results in few feasible individuals, even after extensive optimization. For FINS, migration from the feasible to the infeasible population is more pronounced, and feasible individuals rarely amount to more than 25% of the total population. As with the smaller map in the first experiment, FINS attains a higher diversity in the feasible population than objective-driven approaches.

While different from the Sentient Sketchbook generator




Figure 6: A sample of the final maps created by five participants interacting with the Sentient Sketchbook; the range of map sizes and visualizations is shown.

for their larger, randomly initialized populations and more generations, the controlled experiments presented in this section showcase several desirable properties of FINS which make it more appropriate than MCNS for the generation of map suggestions. As demonstrated by the experiment with large maps, FINS and all FI-2pop GAs are better at discovering feasible individuals in cases where the feasible space is small. Since FINS consistently finds feasible individuals within few generations, it complies to the real-time requirements of the map suggestion generator. On the other hand, the experiment with smaller maps demonstrated that FINS can achieve a higher diversity in the population than MCNS, especially in early stages of evolution. Given the short evolutionary sprint for generating map suggestions, FINS is preferable to MCNS for generating diverse maps — although in problems with easily satisfied constraints MCNS creates more feasible maps, which may also be desirable.

5. USER STUDY

The Sentient Sketchbook tool assists human designers in level creation via real-time map evaluation and playability checks, via a number of map displays and via generated suggestions. In order to assess its usability, a small-scale user study was conducted with 5 expert users consisting of independent game developers, game programmers and level designers. The tool was sent via e-mail to the participants, who were asked to interact with it as they saw fit on their spare time; feedback and interaction data was returned via e-mail. Each participant had several map design sessions; in each session a user creates a complete map starting from a blank canvas. Data is collected from 24 design sessions, with each participant contributing a minimum of 4 sessions.

A sample of the final maps created by users are displayed in Figure 6. Users could choose between small maps (8 by 8 tiles), medium maps (12 by 12 tiles) and large maps (16 by 16 tiles). Participants favored medium-sized maps in 42% of design sessions; remaining 29% of maps were small, and 29% large. Most maps had 2 bases (50%) or 4 bases (25%); all small maps had 2 bases. Unlike bases, resources vary — from 2 to 12 on Small maps, 6 to 27 on Medium maps and




Figure 7: A sample of user maps before (top) and after (bottom) a suggestion was selected, showcasing different instances where suggestions were deemed useful.

6 to 40 on Large maps. Regarding map appearance, participants predominantly favored symmetrical features, with bases placed at the corners of the map as in Figure 6d or following symmetries on non-Cartesian axes as in Figure 6b.

The number of user actions taken (placing tiles, clearing the map or applying suggestions) in each design session range from 13 to 168 (58 on average). Observing the changes in fitness scores after each action, it appears that safe area and resource safety balance are consistent with design patterns as user actions lead to increased scores in these fitness dimensions. Using the difference between positive and negative changes over the total number of fitness changes as the test statistic, obtained values (0.41 and 0.24, respectively) are significant, with $p < 10^{-6}$ for both fitness dimensions.

Across all design sessions, map suggestions were selected 22 times to replace the user’s sketch; although that averages to 0.92 suggestions selected per session, no map suggestions were selected in 9 sessions. Users offered a number of reasons why suggestions were not always useful. The simplest reason is that for some design sessions the user had a specific map in mind even before starting the tool; for instance, a user wanted to “re-create the map from the first Warlords game”. In such cases the design process did not include much map inspection or interaction with the suggestions. A related reason concerns cases where the number of bases differed between the user’s map and the map suggestions; since adding or removing bases (and therefore players) would need a thorough re-evaluation of the designer’s current ideas, such suggestions were not preferred. Another reason was the organic appearance of generated maps; as a user puts it, “when I do something very symmetrical, it feels like the generator just does a more messy version of the same map”. Finally, one user stated that “the suggestions were interesting, but it was not entirely clear how they related to what I was doing”.

Users often selected one map suggestion per session (in 12 sessions). In two sessions of different users, map suggestions were selected 4 times per session; in the first case, the user created an elaborate map (see Figure 7a), and then sequentially applied a map suggestion and saved the new map, thus creating four variations of their original map as a cheap shortcut for generating more content. In the other case, the user created a very simple symmetrical map (see Figure 7b) with all resources on two map edges, and successively used the map suggestions to add impassable regions and randomize tile placement. Overall, the selection of suggestions followed two usage patterns. In 3 sessions, map suggestions

Objective	f_{res}	f_{saf}	f_{exp}	b_{res}	b_{saf}	b_{exp}	NS
# Selected	3	3	5	2	2	1	6

Table 2: The objectives, including novelty (NS), towards which the five users’ selected suggestions were optimized.

were selected early in the design process to quickly generate a map “draft” which was then edited by the designer (see Figure 7c); such design sessions were short, since the suggestion was already playable and did not need significant changes. In 10 sessions, map suggestions were applied as one of the last steps before saving the final map; this last step either aimed at breaking the simple authored patterns and create more organic maps (see Figure 7d) or to increase the score in one or more fitness dimensions (see Figure 7e).

The type of map suggestions selected by users are shown in Table 2. Despite the fact that novelty search generated as many suggestions as all the objective-driven GAs together, its suggestions are selected less frequently than those generated to optimize specific objectives. Novelty search was preferred by users requesting large changes in the map, usually in early stages of the design process when user-created maps were still uninteresting; e.g. Figures 7b, 7c and 7d feature suggestions generated via novelty search. In late stages of the design process, when the authored maps were almost final, suggestions were used to fine-tune the map and were selected based on their improvements to the fitness dimensions. Surprisingly, suggestions targeting balance were not chosen often; a likely reason is that user-created maps were often optimal in most balance dimensions because of the strict map symmetries of human designs. Suggestions which increased the Exploration score, on the other hand, were more preferable as they create long, winding paths which would require significant effort to create manually.

Overall, the expert users’ reception of the Sentient Sketchbook was positive, especially regarding the different map displays; a user stated that “I used the alternative map displays quite often [...] especially the navmesh [...] and I checked at least once for each map the resource safety display to guarantee some safe resources for everyone”. The fitness dimensions were also mostly accurate, although a user commented on the safe area fitness that “I don’t consider empty space valuable to a player, safe or not, I only care about the resource points within it”. While one user initially commented on the suggestions’ apparent lack of visual structure and symmetry, they noted that after a while “the tool starts pushing you to places you didn’t really consider” such as “working with some high level abstract ideas for the maps [...] instead of doing something symmetrical and intricate”. Directions for improvement were also provided, such as speeding up the generation of suggestions for large maps, being “able to lock some features” in the map suggestions and improving the detection of chokepoints, which were deemed “confusing”.

6. DISCUSSION

The experimental results for the feasible-infeasible novelty search paradigm show that it is faster than minimal criteria novelty search at discovering feasible individuals in cases where feasibility is not easily achieved. Based on the small user study with industry experts, the rapid generation of map suggestions was identified as an important feature; since participants in the user study preferred larger

maps with a small feasible space, FI-2pop GA and FINS are ideal for finding feasible individuals quickly. The user study also demonstrated that maps optimized for novelty were selected when human designers needed inspiration; on the other hand, suggestions optimized for specific objectives were also popular when refining an elaborate human-authored map. Based on these results, it is clear that both objective-driven evolution via FI-2pop GA and novelty search via FINS are valuable for generating map suggestions.

The Sentient Sketchbook’s usability as a computer-aided level design tool was tested with five industry experts, who were overall positive about the simple interface, the different visualizations and the evaluated metrics and fitness dimensions which reduce user effort. Feedback from game developers with experience in game engines and computer-aided tools is important for planning future enhancements to the tool; however, their expert knowledge was sufficient for designing “good” maps, and generated suggestions were thus often overlooked. A future user study should invite novice designers and track their response to suggested alternatives; suggestions from the Sentient Sketchbook could be more useful for teaching level design rather than for creating levels in commercial games. Additionally, a more thorough evaluation of the effect on the user’s creativity from suggestions generated via novelty search and objective-driven search should test our assumptions that unexpected suggestions enhance creativity, and are thus desirable.

Future work for the Sentient Sketchbook should aim to improve the quality of generated map suggestions. An interface option allowing the user to “lock” certain map metrics could transform such “locked” values into additional constraints. Moreover, users expressed the need for more visually consistent maps, favoring patterns like symmetry, straight lines and placement of bases on map edges. Finally, the map generator could learn from data of previous design sessions to simulate humans’ design process and aesthetic preferences.

In its current form, the Sentient Sketchbook is more appropriate for generating maps for strategy games. The base and resource tiles and the fitness dimensions assume strategic gameplay; they can be used for authoring maps for different game genres only if creatively reinterpreted. For instance, First Person Shooter maps can be generated via the Dungeon visualization, treating player bases as the team’s spawn location and resources as weapons (with powerful weapons placed in unsafe areas). Future work should allow maps for different game genres to be created, either by removing the current fitness dimensions or by allowing the designer to create their own types of tiles and even their own objectives. Alternatively, a specific game can be targeted by the map sketches; this affords better context for the designer and allows the created maps to be playtested.

Future additions to the Sentient Sketchbook should explore ways to connect the user to a larger group of designers, encouraging the exchange of ideas. This collaboration can be achieved via an online, persistent database of map designs which can be viewed, edited and resubmitted by any user. The collaborative creation of content has shown potential in projects such as PicBreeder [18], enhancing creativity and limiting user fatigue by crowdsourcing design choices.

This paper introduces FINS and tests it against objective-driven GAs and MCNS. Further experiments assessing the performance of FINS in different setups and against more methods for constrained novelty search are included in [10].

7. CONCLUSION

This paper presented the Sentient Sketchbook, a tool which allows a designer to create game levels via a computer-aided sketching interface. The tool supports the designer by automating map evaluations, visualizing them on-screen and proposing alternative designs. The creation of map suggestions combines the feasible-infeasible two-population paradigm with novelty search, and is shown to be more robust than state-of-the-art constrained novelty search in cases where discovery of feasible individuals is difficult. A user study with industry experts demonstrated the tool's potential, attested to the different instances where map suggestions can prove useful and provided important feedback which will guide future additions to the Sentient Sketchbook.

8. ACKNOWLEDGMENTS

We would like to thank the participants of the user study for their valuable feedback. The research is supported, in part, by the FP7 ICT project SIREN (project no: 258453) and by the FP7 ICT project C2Learn (project no: 318480).

9. REFERENCES

- [1] D. Ashlock, C. Lee, and C. McGuinness. Search-based procedural generation of maze-like levels. *IEEE Transactions on Computational Intelligence and AI in Games*, 3(3):260–273, 2011.
- [2] C. Browne and F. Maire. Evolutionary game design. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(1):1–16, 2010.
- [3] Entertainment Software Association. Essential facts about the computer and video game industry, 2012. Retrieved December 12, 2012 from http://www.theesa.com/facts/pdfs/ESA_EF_2012.pdf.
- [4] E. J. Hastings, R. K. Guha, and K. O. Stanley. Evolving content in the galactic arms race video game. In *Proceedings of IEEE Conference on Computational Intelligence and Games*, pages 241–248, 2009.
- [5] J. H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. MIT Press, 1992.
- [6] S. O. Kimbrough, G. J. Koehler, M. Lu, and D. H. Wood. 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, 2008.
- [7] J. Lehman and K. O. Stanley. Revising the evolutionary computation abstraction: Minimal criteria novelty search. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 2010.
- [8] J. Lehman and K. O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2011.
- [9] A. Liapis, G. Yannakakis, and J. Togelius. Limitations of choice-based interactive evolution for game level design. In *Proceedings of AIIDE Workshop on Human Computation in Digital Entertainment*, 2012.
- [10] A. Liapis, G. Yannakakis, and J. Togelius. Enhancements to constrained novelty search: Two-population novelty search for generating game content. In *Proceedings of Genetic and Evolutionary Computation Conference*, 2013. In Print.
- [11] A. Liapis, G. Yannakakis, and J. Togelius. Generating map sketches for strategy games. In *Proceedings of Applications of Evolutionary Computation*, 2013.
- [12] A. Liapis, G. N. Yannakakis, and J. Togelius. Neuroevolutionary constrained optimization for content creation. In *Proceedings of IEEE Conference on Computational Intelligence and Games*, pages 71–78, 2011.
- [13] A. Liapis, G. N. Yannakakis, and J. Togelius. Adapting models of visual aesthetics for personalized content creation. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(3):213–228, 2012.
- [14] T. Lubart. How can computers be partners in the creative process: classification and commentary on the special issue. *International Journal of Human-Computer Studies*, 63(4-5):365–369, 2005.
- [15] Z. Michalewicz. Do not kill unfeasible individuals. In *Proceedings of the Fourth Intelligent Information Systems Workshop*, pages 110–123, 1995.
- [16] J. F. Reintjes. *Numerical control: making a new technology*. Oxford University Press, Inc., 1991.
- [17] M. Schoenauer and Z. Michalewicz. Evolutionary computation at the edge of feasibility. In *Proceedings of the 4th Parallel Problem Solving from Nature*, pages 245–254, 1996.
- [18] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, J. T. Folsom-Kovarik, and K. O. Stanley. Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evolutionary Computation*, 19(3):373–403, Sept. 2011.
- [19] A. M. Smith and M. Mateas. Answer set programming for procedural content generation: A design space approach. *IEEE Transactions on Computational Intelligence and AI in Games*, 3(3):187–200, 2011.
- [20] G. Smith, J. Whitehead, and M. Mateas. Tanagra: Reactive planning and constraint solving for mixed-initiative level design. *IEEE Transactions on Computational Intelligence and AI in Games*, (99), 2011.
- [21] N. Sorenson, P. Pasquier, and S. DiPaola. A generic approach to challenge modeling for the procedural creation of video game levels. *IEEE Transactions on Computational Intelligence and AI in Games*, 3(3):229–244, 2011.
- [22] J. Togelius, R. De Nardi, and S. Lucas. Towards automatic personalised content creation for racing games. In *Proceedings of IEEE Symposium on Computational Intelligence and Games*, pages 252–259. IEEE, 2007.
- [23] J. Togelius, G. Yannakakis, K. Stanley, and C. Browne. Search-based procedural content generation: A taxonomy and survey. *IEEE Transactions on Computational Intelligence and AI in Games*, (99), 2011.
- [24] G. N. Yannakakis. Game AI revisited. In *Proceedings of ACM Computing Frontiers Conference*, 2012.
- [25] G. N. Yannakakis and J. Togelius. Experience-driven procedural content generation. *IEEE Transactions on Affective Computing*, 99, 2011.