Adaptive Game Level Creation
through Rank-based Interactive Evolution

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

Héctor P. Martínez
Institute of Digital Games
University of Malta
Msida, Malta
Email: hector.p.martinez
@um.edu.mt

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

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

Abstract—This paper introduces *Rank-based Interactive Evolution* (RIE) which is an alternative to interactive evolution driven by computational models of user preferences to generate personalized content. In RIE, the computational models are adapted to the preferences of users which, in turn, are used as fitness functions for the optimization of the generated content. The preference models are built via ranking-based preference learning, while the content is generated via evolutionary search. The proposed method is evaluated on the creation of strategy game maps, and its performance is tested using artificial agents. Results suggest that RIE is both faster and more robust than standard interactive evolution and outperforms other state-of-the-art interactive evolution approaches.

I. INTRODUCTION

The generation of personalized digital content for different stakeholders (e.g. player, designer, tester, artist) in the game development process is gaining importance in the game industry and within academia. Algorithmic approaches able to both automatically create game content and adapt the generation processes according to a user's preferences can offer a number of advantages: first, algorithmic processes can alleviate the enormous effort and cost of content creation and make it easier to tailor content to the player; second, content can automatically adapt the game to the needs and preferences of individual players and yield maximal game replayability.

This paper introduces a method for creating game content which is *rapidly* adapted to a user's preferences. The adaptive model of user taste is built upon user rankings of content in terms of domain-specific functional and aesthetic properties, while content iteratively evolves according to this changing model. The proposed method, named Rank-Based Interactive Evolution (RIE), couples computational models of user preference and interactive content generation. RIE considers the rankings (instead of simple selections) of users when interacting with content and employs preference learning [1] to infer selection patterns based on those rankings. In particular, in this study we apply ranking-based back-propagation [2] for inferring a computational model of the user's aesthetic preference represented as an artificial neural network. This model is used, in turn, as a fitness function to generate content tailored to each individual user. The proposed RIE approach incorporates adaptation modules in two layers: at the layer of tailoring user aesthetics models and at the layer of generating personalized content based on the tailored models of aesthetics.

Unlike standard interactive evolution, RIE uses a number of domain-specific predefined fitness dimensions whose impact is adjusted according to content rankings by users. In this paper, RIE is used to generate personalized strategy game maps; its optimization behavior is tested, in a controlled experiment with artificial users, against state-of-the-art algorithms such as Choice-Based Interactive Evolution (CIE) [3] and variants of traditional Interactive Evolution.

This work is novel in several distinct ways: first, computational models of user aesthetics are constructed based on functional and aesthetic criteria identified in the literature as significant for the design of strategy maps; second, the aesthetic models are tailored to the tastes and preferences of users based on rich information coming from user rankings of a number of strategy maps; third, new personalized maps are generated via the RIE algorithm.

II. RELATED WORK

This section reviews related studies in the three key areas that the presented work touches upon: procedural content generation, interactive evolution and user preference modeling.

A. Procedural Content Generation for Games

While the game industry has used procedurally generated content since the eighties in order to increase a game's unexpectedness and replayability, academic interest in procedural game content generation is relatively new [4]. Recent applications of *search-based procedural content generation* (SBPCG) [5] have shown promise in creating high-quality racing tracks [6], weapons [7], spaceships [8] or the gameworld [9]. SBPCG uses evolutionary algorithms or other stochastic search/optimization approaches to explore the space of possible artifacts for those that best satisfy one or several evaluation functions measuring novelty [10], aesthetic pleasantness [11], fairness [12] and interesting gameplay [13], to name a few. In the domain of strategy game maps, Frade et al. evolve game maps via genetic programming towards producing interesting terrain formations without gameplay considerations [14]. Togelius et al. evolved complete maps for the game *Starcraft* (Blizzard 1998) using a multiobjective evolutionary algorithm and a set of fitness functions focused on game balance [15]; the benchmark used to evaluate RIE in this paper is influenced by that set of experiments. Finally, maps

have been generated for the game *Warzone 2100* (Pumpkin Studios, 1999) using Answer Set Programming, a form of constraint solving [16]. These PCG systems allow minimal interactivity, since human preference needs to be expressed as constraints or objectives prior to the generative process.

B. Interactive Evolution

Interactive Evolutionary Computation (IEC) is "the technology in which EC optimizes the target systems based on subjective human evaluation as fitness values for system outputs" [17], and is often used in problems where the evaluation is too complicated or subject to individual taste. In IEC, a user selects (or rates) one or more candidates among presented content; the selected individuals receive preference for evolutionary selection. Since user input is a constant requirement for the success of the optimization algorithm, most IEC systems suffer from *user fatigue* due to the number of choices that human users have to make. IEC projects try to counter fatigue by reducing the number of choices imposed on their users, e.g. via discrete rating "levels" or by crowdsourcing content selection [18]. In game applications, content selection for IEC can be inferred based on gameplay traces [7]. However, the time investment on the part of the players remains considerable. Our approach aims to capture the preferences of the system user from a minimal set of human-computer interactions and use the preference model as a fitness function for EC, bypassing the requirement of constant user feedback. Our user preference modeling solution for accelerating IEC is based on user ranked selections rather than simple user choices [3], providing richer information.

C. User Preference Modeling

User modeling has been widely employed as an input attribute to content generation procedures following the *experience-driven procedural content generation* (EDPCG) paradigm [19]. Through EDPCG, user behavior and preferences drive the generation of various types of digital content including platform game levels [20], [21], weapons [7] and camera settings [22]. In this paper we follow the EDPCG approach for the iterative generation of strategy game maps driven by constructed user preference models; however, the EDPCG framework is enhanced as the user preference model is self-tailored based on subjective rankings of the maps made by the user. While previous studies [3], [23] adapted a user model via direct interaction in IEC applications, the proposed RIE method is novel as it considers the richer information derived from rankings for adapting the user preference model.

III. STRATEGY MAP GENERATION

For the purposes of this study, small-scale overland maps are being evolved to maximize several aesthetic and gameplay properties important to strategy games. These maps are abstractions of levels used in successful strategy games such as *Starcraft* (Blizzard 1998). Each map has a size of 64 by 64 tiles which can be either *passable* or *impassable*, two player bases, ten resources and ten regions of impassable tiles (see Fig. 1). The map layout assumes that each player starts at one of the bases and gathers resources to produce units; units move through passable tiles in order to attack the opponent's base.
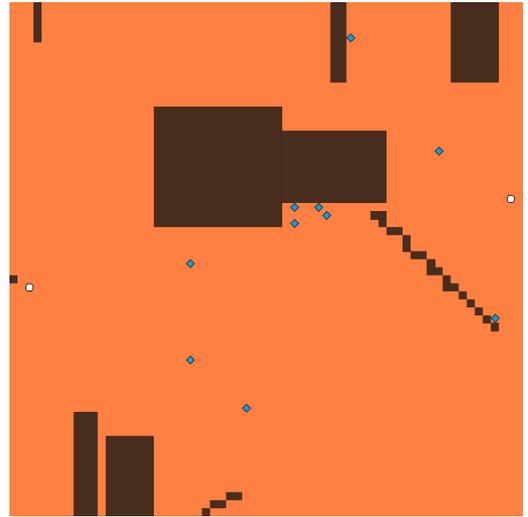

Fig. 1. A sample map generated by the algorithm: light areas represent passable tiles, dark areas represent impassable tiles, blue rhombi represent resources and the two white circles represent the players' bases.

Each map is encoded in an array of 74 real numbers within $[0, 1]$. Each player base or resource is encoded in two parameters, corresponding to the X and Y coordinate if multiplied by the map's width and height respectively. Each impassable region is encoded in five parameters and can be a line or a rectangle, diversified by a parameter t; a rectangle ($t < 0.5$) is represented by the coordinates of its diagonal corners, while a line ($t \geq 0.5$) by the coordinates of an origin point, an angle and a distance.

The 74 parameters which represent the game map are optimized by a genetic algorithm [24]. The gene with the highest fitness is transferred to the next generation, while suitable parents are selected using fitness-proportional roulette wheel selection; the parameters of selected parents are recombined via 2-point crossover. An offspring of two parents has a mutation probability of 1%, while a single parent has a 5% chance of being mutated and copied to the next generation. During mutation, the parameter array's order is reversed (15% chance) or 1 to 5 parameters are modified by a random number following a normal distribution with 0 mean and 0.33 standard deviation — results are bound within [0,1]. Mutation favors local search via small changes, but the occasional drastic change is permitted in order to hinder premature convergence.

The maps are evaluated on gameplay and aesthetic features; gameplay features are largely inspired by previous work on evolving strategy game maps [15], while aesthetic features are inspired by studies on visual perception [25], [26], previous work on computational models of visual aesthetics [3], [11], and popular strategy game design patterns such as "King of the Hill" map templates in *Age of Mythology* (Ensemble Studios, 2002). For the purposes of this study, ten fitness functions are defined in eq. (1)-(10) falling under the following categories:

1) Spatial Navigation: These fitness dimensions evaluate the passable space of the map, and focus on navigation around and between player bases: *base distance* fitness (f_{BD}) in eq. (1) rewards long distances between bases, *base space* fitness (f_{BS}) in eq. (2) rewards passable areas around a base

and *choke point* fitness (f_{CP}) in eq. (3) rewards narrow passes between bases.

$$f_{BD} = min\left\{1, \frac{d_B}{w_M + h_M}\right\} \tag{1}$$

$$f_{BS} = \frac{\bar{a}}{C_1^2} \tag{2}$$

$$f_{CP} = \begin{cases} \frac{1}{2}\left(1 - \frac{w_C}{C_1} + h_A\right) & \text{if } w_C < C_2 \\ C_3\left(1 - \frac{w_C}{d_M}\right) & \text{if } w_C \geq C_2 \\ 0 & \text{if } w_C = 0 \end{cases} \tag{3}$$

where d_B is the distance between the two player bases; w_M, h_M are the map's width and height in tiles, respectively; $\bar{a}$ is the average number of passable tiles within a C_1 by C_1 tile grid centered around every player base ($C_1 = 7$ in this study); w_C is the smallest width on the shortest path between the two player bases; $d_M = (w_M + h_M)/2$ is the map's diagonal; C_2 is the maximum width of a chokepoint and C_3 limits the reward of narrow paths which do not constitute choke points ($C_2 = 10$ and $C_3 = 0.01$ in this study); h_A is the heuristic of the alternate path, assuming the chokepoint was closed: $h_A = min\{1, (d_B - d'_B)/d_M\}$, where d'_B is the shortest path between bases if the chokepoint becomes impassable.

2) Resource Distribution: These fitness dimensions evaluate how resources are placed on the map, in terms of player bases and aesthetic criteria: *nearby resources* fitness (f_{NR}) in eq. (4) rewards resources near both player bases, *nearby resource balance* fitness (f_{RB}) in eq. (5) rewards equality between player bases' nearby resources and the aesthetically-oriented *king of the hill* fitness (f_{KH}) in eq. (6) rewards resources placed in the central area of the map.

$$f_{NR} = \frac{r_{N,1} + r_{N,2}}{2r} \tag{4}$$

$$f_{RB} = \frac{|r_{N,1} - r_{N,2}|}{r} \tag{5}$$

$$f_{KH} = \frac{r_M}{r} \tag{6}$$

where r the number of all resources in the map; $r_{N,i}$ is the number of resources within 16 tiles from the base of player i; r_M the number of resources in a square of width $w_m/2$ and height $h_m/2$ in the center of the map.

3) Aesthetic Appearance: These fitness dimensions evaluate aesthetic properties of the maps; since the maps are presented as 2D images, it is important to consider certain visual cues on which humans evaluate artifacts universally [25], [26], focusing on the *balance* and *division* of impassable regions. The *unconnected segments* fitness (f_{US}) in eq. (7) rewards the presence of distinct "chunks" of impassable tiles, *Symmetry (X and Y)* fitnesses (f_{Sx} and f_{Sy}) in eq. (8) and eq. (9) reward balance of impassable tiles between the left and right halves and between the top and bottom halves respectively, and *Quadrant Symmetry* (f_{QS}) in eq. (10) rewards

balance of impassable tiles in neighboring map quadrants.

$$f_{US} = \frac{I_U}{I} \tag{7}$$

$$f_{Sx} = 1 - \frac{1}{I}|I_T - I_B| \tag{8}$$

$$f_{Sy} = 1 - \frac{1}{I}|I_L - I_R| \tag{9}$$

$$f_{QS} = 1 - \frac{1}{I}\sum_{i=2}^{3}\left(|I_{Q_i} - I_{Q_1}| + |I_{Q_i} - I_{Q_4}|\right) \tag{10}$$

where I the number of impassable tiles in the map; I_U the number of unconnected impassable chunks calculated using a 4-direction flood fill algorithm on all impassable tiles; I_T, I_B, I_L, I_R the number of impassable tiles in the top, bottom, left and right halves of the map respectively; I_{Q_i} is the number of impassable tiles in quadrant i of the map.

IV. USER PREFERENCE MODELING

The fitness functions for map evaluation presented in Section III can be used individually to optimize a single property such as resource balance or symmetry on one axis; they can also be aggregated into a weighted sum in order to embody a more inclusive preference model. Using this weighted sum as the map's fitness score, the genetic algorithm can create content with high scores in many different properties. Also identified as preference score (F), the weighted sum of fitness scores ($\sum w_i f_i$ where f_i is the score of fitness i and w_i its corresponding weight) is normalized to $[0, 1]$.

Deriving a content quality score from a weighted sum allows for the adjustment of its contributing properties' weights based on user choices. As feedback from the user affects the preference score which, in turn, is used as an evolutionary algorithm's fitness function, the proposed interactive evolution process requires minimal human input and thus limits user fatigue. Such a preference model can be adjusted over a series of *iterations* of user interaction, weight adjustment and content evolution. In each iteration a number of maps is presented to the user who provides feedback based on their personal preference; this feedback is used to adjust the weights of the preference model, and the updated preference model is then used to evolve a population of maps, a sample of which is shown in the next iteration. This paper will present two methods of user feedback: the choice of a single map as favorite, or *choice-based* feedback (detailed in Section IV-A) and the ranking of maps in order of preference, or *rank-based* feedback (detailed in Section IV-B). While the former approach has already been successfully used for the generation of spaceships for a 2D game [3], we expect the rank-based approach to allow for more user control and provide a more accurate model of personal taste within fewer iterations.

A. Choice-based Interactive Evolution (CIE)

The simplest approach from a user's perspective is the selection of a single map as the best among those presented; the *choice-based interactive evolution* (CIE) method, introduced in [3], uses an adaptive model informed by the user's choice of one favorite map among a range of presented content. The goal of the adaptive model in such an approach is to reward

fitnesses with a higher fitness score in the selected map than in the unselected ones while penalizing fitnesses with a lower fitness score in the selected map than in the unselected ones. Towards that end, the weight w_k of fitness k is updated on each epoch t as follows:

$$w_k^{t+1} = w_k^t + \alpha(f_{k_S} - \bar{f}_{k_U}) \tag{11}$$

where α is a weight update step (0.01 for this study), f_{k_S} is the selected map's score for fitness k and $\bar{f}_{k_U}$ is the average score for fitness k across all unselected maps. Assuming we are adjusting the weight of the selected individual's fitness property k, eq. (11) follows the key principles of the Widrow-Hoff [27] weight update rule.

The weights are adjusted until the selected map has the highest preference score F among those presented; the adjustment process can be prematurely terminated if the preference score difference between the highest scoring map and the selected map starts to increase or after $3 \cdot 10^5$ weight updates. Since the preference score only measures the relative contribution of visual properties' fitness scores, the final adjusted weights are divided by $\sum_i |w_i|$ resulting in normalized weight values.

B. Rank-based Interactive Evolution (RIE)

In order to increase the amount of information that the user provides to the system on each iteration, we propose a method in which the user ranks the generated content in order of decreasing preference. It is expected that for a subset of the presented maps, users will not have a strong preference of one over the others; in such scenarios they are encouraged to rank those maps equally. Equal ranking is interpreted as an unknown relative preference between those maps; this option might lead to a certain loss of information but also lowers the task's cognitive load and reduces the amount of *noise* inherent to subjective reports [28]. Additionally, to facilitate the task users are allowed to rank only a subset of the presented maps; it is assumed that ranked maps are preferred over non-ranked maps, which would be equivalent to ranking the remaining maps last.

Reports based on comparisons between objects (i.e. maps in our study) have proven to be consistently 'better' than other methods of reporting such as ratings (i.e. assigning to each object a natural or real mark) [28]. Despite their advantages, comparative approaches have been neglected, in part, because popular computational methods for classification and statistical analysis are not directly applicable. Note that the position of objects in the ranking take ordinal values which would be misrepresented as items from a set of labels (order information is lost) and incorrectly analyzed as natural numbers (the distance between consecutive values in an ordinal scale is potentially variable). Fortunately there is a growing interest in *ordinal regression* or *object ranking* [1] and new methods are being developed to cater for such approaches.

The modeling task at hand can be formalized as finding a function $F(x)$ that satisfies:

$$\forall (x_i, x_j), r_{x_i} > r_{x_j} : F(x_i) > F(x_j) \tag{12}$$

where x_j is an object and r_{x_j} is its corresponding rank. As mentioned above, we define our function as a weighted sum of map's fitnesses which is similar to a Single Layer Perceptron [29] employing a linear activation function. Then we can train those weights combining a standard gradient descent algorithm [30] with a cost function defined in terms of pairwise preferences [2].

Without losing information of the objects' relative ordering, a ranking can be transformed into a set of pairwise relations P such as $\forall x_i, x_j : r_i > r_j \implies (x_i \succ x_j) \in P$, i.e. for all possible combinations of two objects with different ranks, we can say that the object with a higher rank is preferred over the object with lower rank. A cost function is then defined as:

$$C = \sum_{\forall (x_i \succ x_j) \in P} \max \{0, b - F(x_i) + F(x_j)\} \tag{13}$$

where $F(x_i)$ is the expected preference score for the preferred object and $F(x_j)$ is the expected preference score for the non preferred object. This cost function equals zero when the expected preference score for the preferred object is greater than the expected preference score for the non preferred object by at least b in every pair ($b = 0.1$ in this study).

Following a gradient descent algorithm, we can iteratively adjust the weights of the network to minimize that cost function. For the linear model used in this study, the update on each epoch t is:

$$w_k^{t+1} = w_k^t - \alpha \frac{\partial C}{\partial w_k} \tag{14}$$

$$\frac{\partial C}{\partial w_k} = \sum_{\substack{\forall (x_i \succ x_j) \in P: \\ F(x_i) - F(x_j) < b}} f_k^j - f_k^i \tag{15}$$

where α is the learning rate (0.01 for this study) and w_k^t is the weight of fitness f_k at epoch t.

This update is repeated until the cost function is equal to zero or a maximum number of epochs is reached ($3 \cdot 10^5$ in this study). Similarly to CIE, the weights are normalized by dividing with $\sum_i |w_i|$.

V. EXPERIMENTS

In order to assess the adaptive models' ability to predict a user's personal taste based on their selections, a number of artificial users are designed to simulate human selection of content. Artificial users provide consistent choices — which can be controlled up to a desired degree — offering the ideal benchmark for observing the preference models' behavior over a number of independent runs. Since the artificial users evaluate content in the same fitness dimensions which are integrated in the aesthetics model, this experiment does not assess the ability of the proposed framework to capture human taste. The experiment instead assesses the efficiency of capturing user criteria and of optimizing content towards them, assuming that the fitness functions included in such a model are representative of user criteria for evaluating content.

Using the criteria of artificial users, two measures of performance are considered for the two-step adaptation process: a) the *population's best individual* or real best individual (R) as evaluated by a user's taste, which demonstrates how well evolution can optimize to the user's taste, and b) the *estimated*

best individual (E) as evaluated by the preference model, which demonstrates how well the preference model can predict the user's taste. Experiments in this section compare the following algorithms:

UB: The *Upper Bound* (or Bound) of the adaptive approaches assumes full knowledge of the user's taste; it optimizes a population to the fitness used to select content like a standard genetic algorithm. For UB the real best individual is identical with the estimated best individual, so the latter is omitted.

RIE: The proposed approach to rank-based interactive evolution presented in Section IV-B, using an aesthetic model derived from the full ordering of eight sample maps according to the fitness used to select content.

CIE: The existing approach to choice-based interactive evolution presented in Section IV-A, using an aesthetic model derived from the choice of the single fittest map among eight sample maps according to the fitness used to select content.

IE: A traditional interactive evolution approach, driven by the rankings of eight sample maps according to the fitness used to select content. In line with literature on interactive evolution [17], artificial user ratings are discretized to fractions of 8: the fittest individual has a fitness of $\frac{8}{8}$ and the least fit individual a fitness of $\frac{1}{8}$. Since the eight maps presented to the artificial user are a small subset of those in the population, this IE approach predicts fitness scores [17] of unpresented individuals based on their genotypical distance from the ranked individual: any individual i in the population is assigned a fitness score f_i of $f_i = d_{i,p}q(p)$ where p is the presented individual with the smallest distance between genotypes ($d_{i,j}$) among presented individuals j and $q(p)$ is the discrete fitness assigned by the artificial user to p; $q(p) = \frac{k}{8}$ where k is the rank of individual p according to the fitness used to select content.

Presented experiments are made on a population of 100 individuals, which evolve for 100 generations divided into *iterations*. In each iteration eight maps are selected from the population and presented to the artificial user who either ranks them or selects their favorite among them. Presented maps include the estimated best individual and the seven individuals with the largest genotypical distance ($d_{i,j}$) from the estimated best individual. Once the preference model is adjusted, the current population is evolved for 10 generations and a new iteration begins. Every new iteration updates the preference model but continues evolution on the previous iteration's population. After a sequence of 10 iterations, content generation is expected to more closely match the user's taste; in the case of artificial users, their taste is the fitness used for selecting content. Mean values and standard error values are collected from 20 independent runs of each approach.

A. User preference based on a single fitness dimension

The simplest artificial user selects content according to one of the fitness dimensions defined in eq. (1)-(10); these users are identified as A_f where f the fitness notation (for instance, A_{BD} selects content according to f_{BD}). For space considerations, the optimization progress for three sample artificial users (A_{BD}, A_{NR}, A_{KH}) is included in Fig. 2; a

	UB	RIE	CIE	IE
A_{C1}				
f_{C1}	0.75 (0.02)	0.65 (0.02)	0.61 (0.02)	0.46 (0.01)
f_{BD}	0.88 (0.04)	0.66 (0.06)	0.57 (0.06)	0.61 (0.03)
f_{NR}	0.48 (0.03)	0.45 (0.06)	0.58 (0.06)	0.37 (0.01)
f_{KH}	0.81 (0.03)	0.85 (0.03)	0.71 (0.04)	0.22 (0.03)
A_{C2}				
f_{C2}	0.85 (0.01)	0.80 (0.01)	0.69 (0.02)	0.46 (0.02)
f_{BD}	0.28 (0.01)	0.23 (0.02)	0.37 (0.06)	0.26 (0.03)
f_{NR}	1.00 (0.00)	0.95 (0.03)	0.81 (0.05)	0.65 (0.03)
f_{KH}	0.99 (0.01)	0.94 (0.02)	0.75 (0.04)	0.44 (0.03)

TABLE I.   THE CONTRIBUTING FITNESS SCORES OF THE REAL BEST INDIVIDUAL AFTER 10 ITERATIONS FOR ARTIFICIAL USERS A_{C1} AND A_{C2}. STANDARD ERROR FOR 20 RUNS IS INCLUDED IN PARENTHESES.

comparison of the final best individuals of all artificial users can be found in Table II.

B. User preference based on composite fitness functions

In order to simulate a more complex model of human taste, artificial users in this experiment take into account multiple criteria, represented as a weighted sum of multiple fitness dimensions from eq. (1)-(10). For space considerations, two sample artificial users are tested (A_{C1}, A_{C2}), selecting content according to the fitness functions f_{C1} and f_{C2} respectively:

- $f_{C1} = 0.5f_{BD} + 0.3f_{NR} + 0.2f_{KH}$
- $f_{C2} = 0.2f_{BD} + 0.3f_{NR} + 0.5f_{KH}$

The choice of f_{C1} and f_{C2} is expected to highlight the limitations of the weighted sum approach for non-dominated multiobjective evolution as well as the impact, if any, of the adaptive models to this particular limitation. Both of these functions combine f_{BD} (which rewards distant player bases), f_{NR} (which rewards many resources near both player bases) and f_{KH} (which rewards resources in the map's middle): f_{BD} is in conflict with f_{NR} since player bases cannot simultaneously be far apart from each other and near all resources, while both f_{BD} and f_{NR} are independent of f_{KH}. These experiments also explore how different weights of the fitness dimensions of f_{C1} and f_{C2} affect the optimization process.

Figures 2(e) and 2(f) show the progress of the two adaptive models with artificial users selecting content based on a linear combination of three fitness scores; a comparison of the final best individuals of all artificial users can be found in Table II. Table I contains information about the individual fitnesses of the best individuals after 100 generations.

C. Dynamic user preference based on variant fitness values

In order to simulate the levels of dynamicity and stochasticity existent in human nature (and preference), an experiment was conducted with an artificial user that changes the fitness used to select content after a number of iterations. For space considerations, the presented artificial user (A_{alt}) selects content according to f_{BD} for the first five iterations and according to f_{NR} for the next five iterations. The two fitnesses are conflicting since maps with high f_{NR} values must have both player bases near every resource, meaning that the bases cannot simultaneously be far apart from each other. Figures 2(g) and 2(h) show the progress of the two fitnesses f_{BD} and f_{NR} both while they are the user's selection criterion and when they are not. A comparison of the best individuals according to f_{NR} after 10 iterations can be found at Table II.
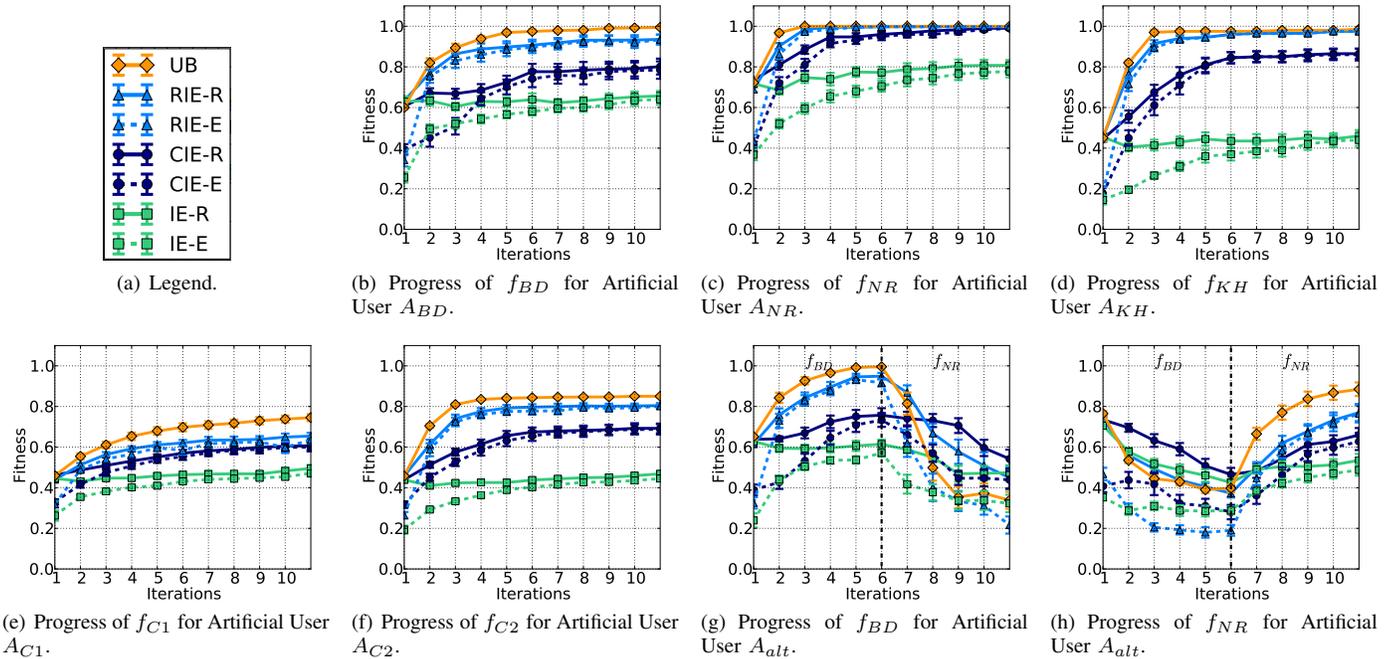
Fig. 2. Optimization progress for evolutionary runs with an artificial user selecting based on different fitnesses (single, composite and alternating). Solid lines represent the individual with the highest stated fitness in the population (R), and dotted lines represent the preference models' estimated best individual's stated fitness dimension (E). Error bars represent the standard error for the 20 individual runs.

	P_{th}	A_{BD}	A_{BS}	A_{CP}	A_{NR}	A_{RB}	A_{KH}	A_{US}	A_{Sx}	A_{Sy}	A_{QS}	A_{C1}	A_{C2}	A_{alt}
UB	Best	1.00	1.00	0.75	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.85	0.88	1.00
RIE	50%	20	20	7	20	20	20	19	20	20	20	20	20	19
	70%	19	20	4	20	20	20	9	20	20	20	17	20	13
	90%	15	20	1	20	20	20	2	20	20	20	1	12	8
	100%	12	20	0	19	20	15	2	11	12	0	0	0	3
CIE	50%	20	20	6	20	20	20	8	20	20	20	20	19	19
	70%	12	20	2	20	20	19	7	20	20	20	9	16	10
	90%	8	20	0	20	20	10	2	20	20	20	3	1	2
	100%	7	20	0	17	20	6	1	16	16	1	0	0	0
IE	50%	18	20	2	20	20	11	8	20	20	20	16	15	13
	70%	6	20	0	18	20	1	0	20	20	19	0	2	3
	90%	0	20	0	6	20	0	0	20	20	6	0	0	1
	100%	0	20	0	1	20	0	0	11	11	0	0	0	0

TABLE II. NUMBER OF RUNS (OUT OF 20 TRIALS) THAT THE POPULATION'S BEST INDIVIDUAL FOR THE COMPARED ALGORITHMS REACHED THE STATED PERFORMANCE THRESHOLD, P_{th}, DEFINED AS THE PERCENTAGE OF THE BEST FITNESS SCORE FOUND BY UB (SHOWN IN THE FIRST LINE).

D. Conclusions of the Experiment

Table II assesses the ability of the different algorithms to optimize their individuals over 10 iterations, using the UB algorithm to derive the algorithms' performance benchmarks.

As expected, the Upper Bound (with full knowledge of the agent's selection scheme) outperforms both CIE and RIE, especially in the early stages of evolution when the interactive approaches are still approximating the user's taste. Figures 2(b)-2(d) show how UB reaches optimal values within a few iterations. There are however fitnesses which rely on very specific map patterns (such as A_{CP}) or try to combine conflicting criteria (such as A_{C1} and A_{C2}): in such cases UB struggles to find fit individuals, and its optimization is slow (shown by the low best scores in Table II). On the other end, there are fitnesses (such as A_{BS} or A_{RB}) that have optimal individuals already in the initial population, and their optimization is effortless.

Interactive evolution does not perform well in most of these experimental settings, primarily because its optimization

progress is sensitive to the initial selection of presented maps. While CIE and RIE select their estimated best individual and seven maps of the most different genes to it, interactive evolution does not have an estimated best individual before the first iteration; the initial presented maps are selected based on their genetic distance from other individuals in the population. There is no guarantee, therefore, that the initial presented maps will have any of the features rewarded by the fitnesses of the artificial user. This often results in a drop in the maximum fitness score in the population after the first iteration, as the highest rated map may still be much worse than unpresented maps; such unpresented maps receive lower fitness scores and are not preferred for selection. In cases where a random mutation is likely to discover fit individuals (such as in f_{NR}), IE manages to overcome the initial drop in fitness scores and increase the population's fitness. This is more pronounced for f_{BS} and f_{RB}, since many optimal individuals exist already in the initial population and are likely to be presented; the population quickly converges to these optimal individuals, and IE performs as well as the other approaches. However, when

highly-fit solutions are difficult to find (such as for A_{CP}, A_{C1} and A_{C2}), IE hardly optimizes its initial population and the best fitness scores rarely reach over 50% of UB's performance. Finally, IE fails to handle shifting objectives in A_{alt}, primarily because when the shift in user taste changes, the population is already converged to maps with high fitness in f_{BD} but low fitness to f_{NR}. This accentuates the sensitivity of IE to presented individuals: at the 6th iteration for A_{alt} the presented individuals will inevitably have low fitness in f_{NR}, causing IE's sub-par performance with A_{alt} compared to A_{NR}.

Comparing the progress of the rank-based and choice-based approaches, the large amount of information contained in the absolute ranking of content for the former approach results in a much more accurate adaptation of the preference model in early stages of evolution. This is evident in the estimated best individuals' progress (dotted lines) in Fig. 2; it is even more pronounced in Fig. 2(h), where the switch in the user's selection criterion is registered much quicker for the estimated best individual of RIE than that of CIE. As evolution pushes towards the estimated best individual according to the preference model, the more accurate RIE model results in the faster optimization of the population's best individual. RIE is faster to optimize both its estimated and its real best individual than CIE for all artificial users, except in cases where optimization is equally effortless for all tested approaches (e.g. f_{BS} and f_{RB}). The models' efficiency varies depending on the fitness function and how easy it is to optimize: Table II demonstrates how, in cases where UB struggles to discover fit individuals (such as A_{US} and A_{CP}), RIE and (moreso) CIE fail to match the performance of UB. Finally, while composite functions are slow to optimize for UB, RIE manages to perform well; both CIE and RIE, however, put less emphasis on the fitness dimensions with the largest weight than UB, and their best individuals are often less dominated by it (especially for f_{BD} in A_{C1}). The large standard error in A_{C1} and A_{C2} for CIE and RIE approaches points to an unpredictable behavior when conflicting fitnesses are simultaneously optimized, as their weights are often readjusted in each iteration to give a single fitness prominence over others.

VI. Discussion

The highly controllable nature of artificial users allowed experiments presented in this paper to demonstrate the ability of RIE and CIE to capture an artificial user's preference in a few iterations. Indeed, the preference model is very quickly adjusted when an artificial user selects content according to one of the fitnesses included in the model and content is evolved to satisfy the artificial user's taste. Experiments with artificial users, however, assume that the fitness functions included in the adaptive model are indicative of human criteria of assessing strategy game maps; in order to validate the full potential of RIE, an experiment with human users is necessary to test whether the fitness functions included in the user model match human taste. Additionally, experiments with human users will assess any user fatigue caused by the cognitive load of relative comparisons when ranking multiple maps.

It should be noted that, unlike traditional interactive evolution, both CIE and RIE adaptive models are contingent on engineered fitnesses and are thus sensitive to bad fitness definitions. The fitnesses currently used often have very different score values for "typical" maps, such as f_{BS} which is optimal for most maps (even randomly generated ones) and f_{CP} which requires very specific map patterns to reach even average scores. In order for the adaptive models to more accurately detect all fitness dimensions as well as to avoid optimization being dominated by one dimension in the weighted sum, a more fair distribution of fitness scores should be in place. Future work will attempt to refine the current fitness scores, in order to balance fitnesses which are currently difficult to optimize with those that are effortless. More fundamentally, methods for automatically refining the current fitness functions and for automatically inventing new fitness functions (for instance, based on raw scores) should be researched.

The interactive evolution method used for comparative purposes follows several "standard" principles regarding quantized ratings, fitness prediction and presenting to the user a subset of a large population. There are several ways in which all three tested interactive evolution approaches (standard, CIE and RIE) could be enhanced, which may improve their performance: a better clustering method could increase the efficiency of IE's fitness prediction, which was currently distance-based, while island models could enhance the optimization behavior of CIE and RIE by increasing exploration of the search space.

The current map representation is straightforward both to describe and to implement, and creates patterns which can be immediately appraised by a user. Small changes in most alleles correspond to small changes in a single map feature's placement, which helps preserve map structure through the genetic operators used. Future work should address the impassable region's *type* parameter, as the change from line to rectangular area is abrupt at a specific value (0.5) and lines and rectangles are represented differently in the genotype. Future work should also explore different phenotypical and genotypical representations which create more visually interesting maps; since optimization of content is susceptible to representation [31], it would be useful to run similar comparisons between IE, CIE and RIE methods across representations.

The proposed framework and algorithms are applicable well beyond the domain of strategy game maps or even games in general. While the current benchmark used to evaluate RIE is tailored to strategic gameplay including base-building, unit production, and defensive tactics, the interactive evolution methodologies presented in the paper can be used for any content with multiple fitness dimensions. Choice-based interactive evolution has been used in previous studies for adapting the visual properties of 2D spaceships [3], showcasing the potential for adapting content such as computer graphics or other forms of evolutionary art based on a designer's or an end-user's taste. Any content which depends on user taste, from the pragmatic considerations of performance versus cost when evolving engineering schematics to the artistic impulses in evolutionary art and music, can be enhanced by RIE's ability to quickly capture user tastes. Beyond the procedural generation of new content, the preference model proposed in this paper can be used independently for evaluating author-created content or content generated by other algorithmic means, such as personalizing suggestions of search results and advertisements. Finally, the accurate estimations afforded by the rank-based preference model can be used to inform designers of their end-user preferences, complementing traditional market research.

VII. Conclusion

This paper introduced the rank-based interactive evolution (RIE) method for personalizing game content based on a two-step adaptation process: while content is evolved to maximize a set of fitnesses, the impact of each fitness is adjusted to approximate a user's preference (expressed in the form of ranking). The RIE framework enhances the capabilities of choice-based interactive evolution introduced in earlier work of the authors [3] as it offers more control to the users and allows for a faster converge to their personal taste. The search-based procedural generation of strategy game maps is used as a testbed for RIE, with maps being evaluated by heuristics inspired by gameplay properties such as spatial navigation and resource distribution as well as visual properties such as symmetry. Controlled experiments with artificial users established the power of RIE at predicting and accommodating a user's taste, assuming that taste can be approximated by the heuristics provided. However, the true potential of personalizing content with RIE should be evaluated via tests with human users. The approaches proposed and methods presented in this paper are directly applicable to other content generation problems where human evaluation is in short supply or expensive.

References

[1] J. Fürnkranz and E. Hüllermeier, "Preference learning: An introduction," *Preference Learning*, pp. 1–17, 2010.

[2] B. Bai, J. Weston, D. Grangier, R. Collobert, K. Sadamasa, Y. Qi, O. Chapelle, and K. Weinberger, "Learning to rank with (a lot of) word features," *Information Retrieval*, vol. 13, pp. 291–314, 2010.

[3] 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*, vol. 4, no. 3, pp. 213–228, 2012.

[4] G. N. Yannakakis, "Game AI revisited," in *Proceedings of ACM Computing Frontiers Conference*, 2012.

[5] 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*, vol. 3, no. 3, pp. 172–186, 2011.

[6] 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*. IEEE, 2007, pp. 252–259.

[7] E. J. Hastings, R. K. Guha, and K. O. Stanley, "Evolving content in the galactic arms race video game," *Proceedings of the 2009 IEEE Conference on Computational Intelligence and Games (CIG)*, pp. 241–248, 2009.

[8] A. Liapis, G. N. Yannakakis, and J. Togelius, "Neuroevolutionary constrained optimization for content creation," in *Proceedings of IEEE Conference on Computational Intelligence and Games*, 2011, pp. 71–78.

[9] ——, "Sentient world: Human-based procedural cartography," in *Proceedings of Evolutionary and Biologically Inspired Music, Sound, Art and Design Conference*, vol. 7834, LNCS.

[10] ——, "Enhancements to constrained novelty search: Two-population novelty search for generating game content," in *Proceedings of Genetic and Evolutionary Computation Conference*, 2013.

[11] ——, "Optimizing visual properties of game content through neuroevolution," in *Proceedings of the Artificial Intelligence for Interactive Digital Entertainment Conference*, 2011.

[12] ——, "Generating map sketches for strategy games," in *Proceedings of Applications of Evolutionary Computation*, vol. 7835, LNCS. Springer, 2013, pp. 264–273.

[13] C. Browne and F. Maire, "Evolutionary game design," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 2, no. 1, pp. 1–16, 2010.

[14] M. Frade, F. F. de Vega, and C. Cotta, "Evolution of artificial terrains for video games based on accessibility," in *Proceedings of EvoApplications 2010*, vol. 6024, LNCS. Springer, 2010, pp. 90–99.

[15] J. Togelius, M. Preuss, N. Beume, S. Wessing, J. Hagelback, and G. Yannakakis, "Multiobjective exploration of the starcraft map space," in *IEEE Symposium on Computational Intelligence and Games (CIG)*. IEEE, 2010, pp. 265–272.

[16] A. Smith and M. Mateas, "Answer set programming for procedural content generation: A design space approach," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 3, no. 3, pp. 187–200, 2011.

[17] H. Takagi, "Interactive evolutionary computation: Fusion of the capabilities of EC optimization and human evaluation," *Proceedings of the IEEE*, vol. 89, no. 9, pp. 1275–1296, 2001, (invited paper).

[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*, vol. 19, no. 3, pp. 373–403, 2011.

[19] G. N. Yannakakis and J. Togelius, "Experience-Driven Procedural Content Generation," *IEEE Transactions on Affective Computing*, vol. 2, pp. 147–161, 2011.

[20] N. Shaker, G. N. Yannakakis, and J. Togelius, "Towards automatic personalized content generation for platform games," in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 2010.

[21] C. Pedersen, J. Togelius, and G. N. Yannakakis, "Modeling Player Experience for Content Creation," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 2, pp. 54–67, 2010.

[22] G. N. Yannakakis, H. P. Martínez, and A. Jhala, "Towards Affective Camera Control in Games," *User Modeling and User-Adapted Interaction*, vol. 20, no. 4, pp. 313–340, 2010.

[23] P. Machado, J. Romero, A. Cardoso, and A. Santos, "Partially interactive evolutionary artists," *New Generation Computing – Special Issue on Interactive Evolutionary Computation*, vol. 23, no. 42, pp. 143–155, 2005.

[24] D. Goldberg, *Genetic algorithms in search, optimization, and machine learning*. Addison-wesley, 1989.

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

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

[27] B. Widrow and S. D. Stearns, *Adaptive signal processing*. Prentice-Hall, Inc., 1985.

[28] G. Yannakakis and J. Hallam, "Rating vs. preference: A comparative study of self-reporting," in *Affective Computing and Intelligent Interaction*. Springer, 2011, vol. 6974, LNCS, pp. 437–446.

[29] F. Rosenblatt, "The perceptron: A probabilistic model for information storage and organization in the brain." *Psychological review*, vol. 65, no. 6, pp. 386–408, 1958.

[30] D. Rumelhart, *Backpropagation: theory, architectures, and applications*. Lawrence Erlbaum, 1995.

[31] D. Ashlock, C. McGuinness, and W. Ashlock, "Representation in evolutionary computation," in *Advances in Computational Intelligence*. Springer Berlin Heidelberg, 2012, vol. 7311, LNCS, pp. 77–97.