New Theses

Personalized Procedural Map Generation in Games via Evolutionary Algorithms

Doctoral Thesis by William L. Raffe

In digital games, the *map* (sometimes referred to as the *level*) is the virtual environment that outlines the boundaries of play, aids in establishing rule systems, and supports the narrative. It also directly influences the challenges that a player will experience and the pace of gameplay, a property that has previously been linked to a player's enjoyment of a game [1]. In most industry leading games, creating maps is a lengthy manual process conducted by highly trained teams of designers. However, for many decades *procedural content generation* (PCG) techniques have posed as an alternative to provide players with a larger range of experiences than would normally be possible. In recent years, PCG has even been proposed as a means of tailoring game content to meet the preferences and skills of a specific player, in what has been termed *Experience-driven PCG* (EDPCG) [2].

This thesis contributes to the growing EDPCG research field with a focus on personalizing maps. Here, the EDPCG techniques are used within a Search-based PCG (SBPCG) [3] framework, utilizing evolutionary algorithms to search for maps that are appropriate for an individual player. Evolution is a common strategy in SBPCG as it provides a logical means of evaluating map candidates and iteratively improving them over multiple generations of recombination and mutation. Furthermore, this thesis investigates a decomposed approach to map generation, using separate evolutionary cycles, genetic representations, and fitness evaluators for two aspects of a map: the *geometry* and *content layout*. The geometry of a map defines the boundaries of play and the location of static virtual objects. Meanwhile, the content layout describes the location and quantity of interactive game assets, such as enemies and pick-ups. Both of these components affect a player's experience to varying degrees in different game genres but are typically related in that the content layout must be within the geometry.

Finally, as the maps should be appropriate for an individual player, both direct interactive evolutionary computing (IEC) and player preference modeling are investigated as methods of collecting, interpreting, and utilizing knowledge about the player's desires.

The thesis starts with a brief foray into evolutionary terrains. This work was conducted as an initial study into optimizing the most common type of base map geometry and ignoring the content layout completely. Terrains were generated by extracting uniform patches from user-provided sample terrains, recombining them in a grid based genetic representation, and rendering a larger terrain by stitching the patches back together again. The result was a content authoring tool that utilized IEC with both parent selection and gene selection to expedite the terrain creation process for novice designers [4].

However, the centrepiece of the thesis is an unsupervised public experiment on an online map personalization solution. In this solution, the geometry and content layout are linked through hierarchical optimization; first optimizing the geometry of a map and then using that as input to the content layout optimization. The geometry is represented as a custom made fixed n-ary tree that connects pre-made room and corridor templates together into a constrained tree structure and is again evaluated through IEC. The density of content within each room (node) of the geometry tree is then calculated by using the tree coordinates of each node as input to a Compositional Pattern-Producing Network (CPPN) [5] and translating the output as quantities of the various enemies and pick-ups within that room. IEC was deemed inappropriate as a fitness evaluation mechanism here because a content layout candidate could not be as easily visualized to the player. Instead, for each player-selected geometry, roughly 200 generations of NeuroEvolution of Augmenting Topologies (NEAT) [6] (with 50 CPPN candidates in each generation) are evaluated through a learned per-player preference model that is based around the paradigm of model-based and content-based recommender system (RS).

This RS player model requires that players provide a single rating value on how much they enjoyed a map after they play it. This rating is then combined with extracted map features to train a Naive Bayes classifier. That classifier then predicts the probability that a player will enjoy a map that results from a CPPN candidate and thus CPPN-NEAT is being used to reduce the number of RS evaluations need to provide a good recommendation. Further details of the system are presented in an initial report that analyzes the experiences of three sample players [7].

In summary, the primary contributions of this thesis are: 1) the use of decomposed optimization to personalize the geometry and content layout of a map in two separate, yet integrated, evolutionary processes; 2) two geometry generation solutions, one for outdoor terrains and the other of interior spaces, that both function by re-combining pre-made map segments and are controlled via IEC; 3) the application of CPPN-NEAT to the map generation process, specifically for determining the density of content given a map location; 4) an initial exploration into using knowledge from the RS field to create a player preference model to be used during the fitness evaluation of a SBPCG system; 6) an example of rigorous statistical analysis of a personalized PCG system deployed in an unsupervised public experiment; and 5) the introduction of the learning trend metrics, which give a clearer indication of RS performance in environments with noisy user data over traditional accuracy measurements.

References

[1] T. W. Malone, “What makes things fun to learn? a study of intrinsically motivating computer games.” *Pipeline*, vol. 6, no. 2, pp. 50–51, 1981.

[2] G. N. Yannakakis and J. Togelius, “Experience-driven procedural content generation,” *IEEE Transactions on Affective Computing*, vol. 2, no. 3, pp. 147–161, Jul 2011.

[3] J. Togelius, G. N. Yannakakis, K. O. 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, Sep 2011.

[4] W. L. Raffe, F. Zambetta, and X. Li, “Evolving patch-based terrains for use in video games,” in *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. ACM, 2011, pp. 363–370.

[5] K. O. Stanley, “Compositional pattern producing networks: A novel abstraction of development,” *Genetic Programming and Evolvable Machines*, vol. 8, no. 2, pp. 131–162, Jun 2007.

[6] K. O. Stanley and R. Miikkulainen, “Evolving neural networks through augmenting topologies,” *Evolutionary computation*, vol. 10, no. 2, pp. 99–127, 2002.

[7] W. L. Raffe, F. Zambetta, and X. Li, “Neuroevolution of content layout in the PCG: Angry Bots video game,” in *Evolutionary Computation (CEC), IEEE Congress on*. IEEE, 2013.

William L. Raffe received a B.CompSci, B.CompSci (Hons) in Computer Science in 2008 and 2009, respectively, from the School of Computer Science and Information Technology at RMIT University, Melbourne (Australia). He completed his Ph.D. in Computer Science in 2014, having been supervised by Dr. Fabio Zambetta and Dr. Xiaodong Li, both of RMIT University. William is currently a post-doctoral research fellow at the same institution, investigating player modeling for adaptive virtual-physical play in water environments. He is also a sessional lecturer, a member of the Golden Key International Honor Society, and a regularly attending member of the RMIT Evolutionary Computing and Machine Learning Group. His research interests include personalized procedural content generation, adaptive game A.I., player preference modeling, and exergaming, to all of which he focuses on applying machine learning and metaheuristic optimization techniques.

Email: william.raffe@rmit.edu.au
Homepage: WWW
Dissertation: pdf