Proceedings of the Twenty-First AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2025)

Game Balancing via Procedural Content Generation and Simulations

Florian Rupp[1,2]

[1]Department of Computer Science
Technical University of Applied Sciences Mannheim, Germany
[2]University of Mannheim, Germany
f.rupp@hs-mannheim.de

Abstract

Balancing games requires extensive manual work and human playtesting during development. Existing research proposes using search-based optimization combined with game simulations to estimate balance. However, these approaches are tailored to specific environments, making them difficult to transfer to other games. Additionally, simulating a game is computationally intensive. My research aims to develop procedural content generation (PCG) methods to automate the generation of balanced content for competitive games. Unlike related work, domain-independent methods enable transferability to other games. I focus on two key aspects that significantly impact a game's overall balance: game levels and economies. Game level balancing is approached by framing it as both, a PCG task and a Markov decision process in order to apply reinforcement learning (RL). Therefore, I adapt and extend the PCGRL (PCG via RL) framework. Based on an existing formal definition, my research considers game economies as graph structures. In this context, I present two frameworks: G-PCGRL (Graph-PCGRL) and GEEvo (Game Economy Evolution).

Introduction

In order to ensure engagement in competitive games, they must be balanced for all players; otherwise, players will become frustrated or bored and quit playing (Becker and Görlich 2020). However, ensuring balance when creating new games is a challenging and time-consuming task (Schreiber and Romero 2021). For this reason, many works propose approaches to automate game balancing (see the related work section). These works, however, propose methods that use game-dependent information, making them difficult or impossible to transfer to other games. My goal is to address this issue by focusing on methods designed to be decoupled from a specific game.

In addition, game balancing is multifaceted. Based on the foundational books (Schreiber and Romero 2021; Adams 2014), I identify game level and economy design as crucial facets to target for automated game balancing. Game economies are abstract concepts that define how virtual resources are created and transferred in a game, offering a powerful system for controlling the game's overall balance (Schreiber and Romero 2021). However, there is little research on game economies in the literature (Rogers et al. 2023). Furthermore, I approach game balancing from a procedural content generation (PCG) perspective. This differs from traditional game balancing, which mainly involves adjusting the numerical values of game entities, such as health points. In my approach, e.g., in the context of levels, the spatial distribution of resources in relation to the players influences the balance. Therefore, I will define and answer four research questions (see the "Works to Date" section).

As a doctoral student nearing graduation, I seek feedback and discussions on my research findings from the scientific community at a glance – something not possible through the review process of individual papers. I am also looking to learn more about career opportunities in academia or industry at the AIIDE 2025 Doctoral Consortium (phase 2).

Related Work

Automated game balancing: Related works often use search-based methods considering game balancing as an optimization problem. For example, Hom and Marks (2007) use an evolutionary algorithm (EA) to balance board games, Volz, Rudolph, and Naujoks (2016) multi-object EAs to balance card game decks. In contrast, Pfau et al. (2020) introduce neural network-based deep player behavior models to replicate human in-game behavior based on existing player data in order to estimate game balance. Little research is present on using PCG directly for game balancing. Lara-Cabrera, Cotta, and Fernández-Leiva (2014) and Lanzi, Loiacono, and Stucchi (2014) generate balanced maps for a first-person shooter game with an EA. In contrast to the existing works, we do not incorporate any game-specific information and frame level balancing as a Markov decision process (MDP).

The PCGRL framework: The PCGRL framework by Khalifa et al. (2020) introduces several representations to frame the generation of tile-based levels as MDP to apply RL. We will use and extend this framework to define game level balancing as MDP. Since its introduction, it has been further extended also by other works. Earle et al. (2021) extended PCGRL to be controllable. Controllability is provided by including additional information in the observation space and adjusting the reward function dynamically. Jiang et al. (2022) further adapt PCGRL to generate 3D levels and demonstrate this in a Minecraft environment. Earle, Jiang, and Togelius (2024) try to address scaling issues and experiment with different definitions of the observation space like partial observability. Most recently, Baek et al. (2025) introduce PCGRLLM, an approach using LLMs to generate tailored reward functions for the RL.

Work to Date

I have published several peer-reviewed works that address automated game balancing in terms of game levels and economies by framing the problem as a PCG problem. In both cases, simulations are used to determine the balancing state. First steps towards this research direction were presented in a doctoral consortium paper, outlining the initial problem formulation and proposed methodology for level balancing (Rupp 2023). My research is structured around four research questions (RQs). RQ 1–3 cover automated level balancing and is based on the AIIDE paper PCGRL by Khalifa et al. (2020). For game economies (RQ 4), I build upon the work by Klint and van Rozen (2013).

RQ 1 **Quantification of balance:** *What is a reliable, data-driven foundation that enables automated, game-independent measurement and quantification of balance in competitive, two-player game levels?*

In my recent works (Rupp, Eberhardinger, and Eckert 2023, 2024) I successfully showed how the existing fairness metric *Statistical Parity* (Dwork et al. 2012) can be used to express a game's balance. The Statistical Parity is typically applied in the context of fair machine learning in order to estimate whether a trained classifier behaves fairly for different socioeconomic groups, for instance. Transferring to games, this formulation ensures that all players have an equal probability of winning. Using this underlying assumption, I derive a metric to express a game's balance numerically in the interval $[0, 1]$, where 0.5 indicates balance, and 0 and 1 indicate maximal imbalance. Moreover, by considering only which player wins, the metric remains game-independent and establishes the foundation for estimating balance when generating balanced levels.

Estimating a game's balance this way requires simulating it. Since games usually contain probabilistic elements, the simulation must be run multiple times to approximate the actual balance. However, this process is computationally intensive, which is why we pose RQ 2.

A Markov Decision Process for Generating Balanced Levels

RQ 2 **Accelerating automated level balancing:** *How can automated game level balancing be accelerated while maintaining content quality and diversity?*

The PCGRL framework was introduced to generate levels by framing the PCG problem of tile-based level generation as an MDP for applying RL. The model is trained in advance using RL, allowing for fast inference when generating content. To estimate levels' balance, I use the metric developed in RQ 1, which can be considered a static simulation-based evaluation function in this context according to Yannakakis and Togelius (2011). Since RL learns during the training phase, unnecessary simulations can be avoided during inference. This speeds up generation time and makes PCGRL suitable for this problem. To apply PCGRL to level balancing, I introduced an architecture in which level balancing is considered fine-tuning an already generated, yet unbalanced, level (Rupp, Eberhardinger, and Eckert 2023). Balance is thereby entirely ensured through the level design. Therefore, I presented a new approach to framing level balancing as an MDP with a swap-based pattern of the action space. In an unbalanced level, the RL agent can modify the level by swapping the positions of two tiles. The agent is rewarded for improving the game's overall balance. I evaluated the approach using a simplified version of the Neural MMO (NMMO) environment (Suarez et al. 2019) in a two-player setting. In my thesis, I will additionally demonstrate the easy transferability of this approach to a different tile-based environment. The swapping-based action space pattern outperforms the original PCGRL in direct comparison.

In an additional paper (Rupp, Eberhardinger, and Eckert 2024), I further investigated aspects such as improving the overall performance to consistently balance 89.7% of levels, enhancing the action and observation spaces, and examining content diversity. For instance, in comparison with hill climbing approaches, I demonstrated that my RL approach is superior.

In another paper (Rupp and Eckert 2025), I strengthened the previous results by investigating an approach that balances levels for asymmetric player archetypes. This refers to a setup in which one player can perform actions that another player cannot, yet we still want to ensure the level's overall balance. Previously, both player heuristics were identical when simulating the game to estimate balance. The main finding is that the approach from Rupp, Eberhardinger, and Eckert (2023) is feasible for solving this problem. However, performance and time to convergence depend on the initial imbalance of the player archetype setup. Games are, however, made for humans, and until now, I have only considered simulated balance using artificial heuristic agents. For this reason, we pose RQ 3.

RQ 3 **Human perception of balance:** *How do humans perceive and evaluate artificially quantified balance through simulations in game levels?*

We address RQ 3 by conducting an empirical evaluation of generated levels with human playtesters (Rupp et al. 2024). Participants were asked to play through different scenarios, each of which consisted of an unbalanced and a balanced version of a level that was generated using the method from RQ 2. To prevent bias, participants did not know which version they were playing, and the order was shuffled each time. In a subsequent survey, the participants were asked questions about their perception of balance. Since balance is multifaceted, the questions were tailored to different game elements. The descriptive results, combined with those from hypothesis testing, revealed that participants perceived the automatically balanced levels from RQ 2 as more balanced than the unbalanced ones in most scenarios. However, the perceived differences in balance depend on the levels and correspond to the changes made by the RL agent.

Game Economy Generation and Balancing

RQ 4 Automated game economy balancing: *Which approaches enable the integration of automation techniques into game economy generation and balancing?*

I model game economies as graphs, following the already introduced abstract representation of game economies as graphs by Klint and van Rozen (2013), which is also followed by the commercial tool Machinations[1]. Driven by previous results from RQ 2, I proposed an MDP for controllable graph data generation, further extending the PCGRL framework: G-PCGRL (Graph-PCGRL) (Rupp and Eckert 2024a). In G-PCGRL, the integer matrix of a level in PCGRL is considered the adjacency matrix of a graph. Thus, the MDP is modified to learn how to connect different node types in a graph by creating edges between them. Validity is specified by a set of constraints. Additionally, G-PCGRL is designed to be controllable; after training, it can generate graphs within a predefined size range and with a predefined set of allowed node types. Although G-PCGRL performs well with different sets of constraints during evaluation, its scalability is limited, resulting in poor performance with larger graphs.

To address this shortcoming, I introduce GEEvo (Game Economy Evolution (Rupp and Eckert 2024b)), a framework that uses EAs to generate and balance graph-based game economies. Although evolutionary optimization is computationally more intensive and highly dependent on randomness, it allows for the generation of larger and more complex graphs. As in RQ 2, GEEvo separates content generation from balancing, consisting of two EAs: one for generation and one for balancing. Experiments demonstrated its ability to generate economy graphs up to a tested size of 20 and balance arbitrarily configured economies 93.3% of the time. As with level generation and balancing, the generated content is simulated by executing the economy over several time steps to evaluate the balancing state. In a case study examining balancing, we found that it is important to configure the balancer to produce a balance within an acceptable range rather than a "perfect" balance. This is important to avoid weakening or mitigating any probabilistic mechanics a designer may have intended.

Limitations & Future Work

I have introduced several methods for automating game balancing. All these methods use simulations to estimate balance. While this can reduce manual effort, a final human evaluation is still necessary. Additionally, the RL-based methods suffer from limited scalability due to an exponentially growing action space based on level or graph size. All methods are tested and evaluated in game environments for research purposes only since the focus is on developing new algorithmic approaches, not game design itself.

My work contributes to game balancing by ensuring an equal chance of winning for all players. However, this must also fit into a game's narrative. My GEEvo framework for generating and balancing game economies is abstract and decoupled from any game narratives by design. In light of recent advances in applying LLMs (large language models) to PCG (Maleki and Zhao 2024), using LLMs to generate narrative settings for functional components of game economies could show promise. Moreover, recent work in the PCGRL context demonstrates how reward functions can be generated using LLMs in the context of level balancing (Baek et al. 2025).

Furthermore, the GEEvo simulation framework is intentionally lightweight and abstract. There is significant potential for further development and addition of functional components to create more complex economies. This could entail small scripts or trained models that trigger the actions of these components, thereby making the simulation more complex. Moreover, investigating how automatic balancing adjustments at the abstract economy level directly affect gameplay in (simulated) playtests would yield further insight. Lastly, another promising research direction to improve performance and achieve faster convergence in training is using transfer learning with RL for level balancing. Designing a curriculum that guides the agent to take actions with a greater impact on the overall balance when generating a level could also speed up the training process.

Conclusion

This work contributes to the academic fields of search-based optimization and deep RL, focusing on the application for automated game balancing – a niche academic field distinct from the broader, practice-driven domain of game balancing. I estimate balance using simulations and approach it from two angles: game levels and economies. Unlike other studies, my methods are designed to be independent of a particular competitive game, considering only which player won. Overall, my research yields three implications:

(1) Randomness is optimized out: My research shows that optimization approaches for both levels and economies try to eliminate probabilistic game elements to stabilize the balance. However, since randomness is important for making games interesting and replayable, I argue based on my observations, that a mechanism must be implemented to ensure randomness and balance coexist.

(2) PCGRL beyond just level generation: I adapted and extended the PCGRL framework by Khalifa et al., going far beyond its original use case of level generation. By introducing new representations, I was able to redefine the MDP to target game balancing as level fine-tuning, as well as to generate graph data.

(3) Games must be balanced, not fair: While balance seeks to achieve mathematical equilibrium, fairness considers a player's socioeconomic background. In a perfectly fair game, outcomes would then depend solely on luck. However, it is precisely the differences between human players that should influence winning or losing. If skill cannot be separated from randomness, the game becomes less engaging and ultimately unsatisfying.

[1]https://https://machinations.io/

References

Adams, E. 2014. *Fundamentals of Game Design*. New Riders Publishing, 3rd edition. ISBN 0321929675.

Baek, I.-C.; Kim, S.-H.; Earle, S.; Jiang, Z.; Jin-Ha, N.; Togelius, J.; and Kim, K.-J. 2025. PCGRLLM: Large Language Model-Driven Reward Design for Procedural Content Generation Reinforcement Learning. ArXiv:2502.10906 [cs].

Becker, A.; and Görlich, D. 2020. What is Game Balancing? - An Examination of Concepts. *ParadigmPlus*, 1: 22–41.

Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, 214–226. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-1115-1.

Earle, S.; Edwards, M.; Khalifa, A.; Bontrager, P.; and Togelius, J. 2021. Learning Controllable Content Generators. In *2021 IEEE Conference on Games (CoG)*, 1–9. Copenhagen, Denmark.

Earle, S.; Jiang, Z.; and Togelius, J. 2024. Scaling, Control and Generalization in Reinforcement Learning Level Generators. In *2024 IEEE Conference on Games (CoG)*, 1–8.

Hom, V.; and Marks, J. 2007. Automatic Design of Balanced Board Games. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 3(1): 25–30. Number: 1.

Jiang, Z.; Earle, S.; Green, M.; and Togelius, J. 2022. Learning Controllable 3D Level Generators. *Proceedings of the 17th International Conference on the Foundations of Digital Games*.

Khalifa, A.; Bontrager, P.; Earle, S.; and Togelius, J. 2020. Pcgrl: Procedural Content Generation via Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 16, 95–101.

Klint, P.; and van Rozen, R. 2013. Micro-Machinations. In *Software Language Engineering*, Lecture Notes in Computer Science, 36–55. Cham: Springer International Publishing. ISBN 978-3-319-02654-1.

Lanzi, P. L.; Loiacono, D.; and Stucchi, R. 2014. Evolving maps for match balancing in first person shooters. In *2014 IEEE Conference on Computational Intelligence and Games*.

Lara-Cabrera, R.; Cotta, C.; and Fernández-Leiva, A. J. 2014. On balance and dynamism in procedural content generation with self-adaptive evolutionary algorithms. *Natural Computing*, 13: 157–168.

Maleki, M. F.; and Zhao, R. 2024. Procedural Content Generation in Games: A Survey with Insights on Emerging LLM Integration. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 20(1): 167–178. Number: 1.

Pfau, J.; Liapis, A.; Volkmar, G.; Yannakakis, G. N.; and Malaka, R. 2020. Dungeons & Replicants: Automated Game Balancing via Deep Player Behavior Modeling. In *IEEE Conference on Games (CoG)*, 431–438.

Rogers, K.; Le Claire, V.; Frommel, J.; Mandryk, R.; and Nacke, L. E. 2023. Using Evolutionary Algorithms to Target Complexity Levels in Game Economies. *IEEE Transactions on Games*, 15(1): 56–66.

Rupp, F. 2023. Learning the Generation of Balanced Game Levels. DC@KI2023: Proceedings of Doctoral Consortium at KI 2023.

Rupp, F.; Eberhardinger, M.; and Eckert, K. 2023. Balancing of competitive two-player Game Levels with Reinforcement Learning. In *IEEE Conference on Games (CoG)*, 1–8. ISSN: 2325-4289.

Rupp, F.; Eberhardinger, M.; and Eckert, K. 2024. Simulation-Driven Balancing of Competitive Game Levels With Reinforcement Learning. *IEEE Transactions on Games*, 16(4): 903–913.

Rupp, F.; and Eckert, K. 2024a. G-PCGRL: Procedural Graph Data Generation via Reinforcement Learning. In *2024 IEEE Conference on Games (CoG)*, 1–8.

Rupp, F.; and Eckert, K. 2024b. GEEvo: Game Economy Generation and Balancing with Evolutionary Algorithms. In *2024 IEEE Congress on Evolutionary Computation (CEC)*, 1–8.

Rupp, F.; and Eckert, K. 2025. Level the Level: Balancing Game Levels for Asymmetric Player Archetypes With Reinforcement Learning. In *Proceedings of the 20th International Conference on the Foundations of Digital Games*, FDG '25. New York, NY, USA: Association for Computing Machinery. ISBN 9798400718564.

Rupp, F.; Puddu, A.; Becker-Asano, C.; and Eckert, K. 2024. It might be balanced, but is it actually good? An Empirical Evaluation of Game Level Balancing. In *2024 IEEE Conference on Games (CoG)*, 1–4.

Schreiber, I.; and Romero, B. 2021. *Game Balance*. Boca Raton, Florida, USA: CRC Press. ISBN 978-1-315-15642-2.

Suarez, J.; Du, Y.; Isola, P.; and Mordatch, I. 2019. Neural MMO: A Massively Multiagent Game Environment for Training and Evaluating Intelligent Agents.

Volz, V.; Rudolph, G.; and Naujoks, B. 2016. Demonstrating the Feasibility of Automatic Game Balancing. In *Proc. of the Genetic and Evolutionary Computation Conf. (GECCO)*, 269–276. ISBN 978-1-4503-4206-3.

Yannakakis, G. N.; and Togelius, J. 2011. Experience-Driven Procedural Content Generation. *IEEE Transactions on Affective Computing*, 2(3): 147–161.