Proceedings of the Sixteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE-20)

Personalized Procedural Content Generation for Increased Player Agency

Kristen K. Yu
University of Alberta
116 St & 85 Ave
Edmonton, Alberta, Canada, T6G 2R3

Abstract

My interest is in the intersection between AI directors, player modeling, and procedural content generation. Using a combination of these tools, I believe that we can expand the space of tellable stories to create a story that is completely driven by the player's actions in the game. There are two use cases for this work. The first is adding additional content to an open-world game by letting the player influence the world after the main story is completed, which would help address the problem of replayability. The second is to allow the player to create a completely personalized story from the beginning of the game, using the starting situation of the game as if it was the start of a simulation.

Introduction

Players enjoy games that provide the player with a high sense of agency with in a narrative, and a player feels a high sense of agency when they believe that the choice that they took in a story made a significant difference in the outcome of that story (Cardona-Rivera et al. 2014). In industry published games, developers provide a sense of agency by using a branching narrative. The choices that the player makes throughout the game influence which branch the player is experiencing. This kind of storytelling provides an interesting narrative for many players but is also restrictive in where the player's choices actively affect the overall narrative. This approach is also limited to the number of branches that the developers implement in the game. The branching nature sometimes creates a low sense of agency for the player. For example, in the game Infamous Second Son (Productions 2014) there are clearly two paths for the player to choose, the good path and the bad path. Each path makes unique changes to the gameplay, but the agency for the player is low because the only choice is to decide which path to play on during the game. One solution to this problem is to create more situations where the player's choice will have a meaningful effect on the story, which means that each of those choices needs to have a planned consequence within the game. Traditionally, each of the consequences are planned out by a developer, which limits the number of choices that can be made in the story to the amount of work that the developer can accomplish during the development period. Another possible way to generate more consequences is through procedural content generation. This brings up the second question of what kind of procedurally generated content would be best. I believe that a combination of a player model and AI director can be used to direct the content generation towards a story that the player will find the most engaging. The combination of these three systems can address the problem of limited player agency by allowing for more choices to affect the narrative, and can provide a personalized open-ended story to the player through the AI director.

Related Work

There has been a push in academia to provide more open-ended stories in interactive narratives through the use of drama managers (Riedl and Young 2010). These can be focused on authorial intent, virtual character autonomy, or player modeling to create the story (Riedl and Bulitko 2013). This changes the kinds of content that the drama manager will present to the player. There has been some work which attempts to capture the play style of the player and use that to influence the content that the player will receive next, which is the player modeling approach. Initial work in this area was done by Thue et al. (2007), where the quests were not generated automatically for the player, but instead different quests were assigned to the player based off of the current model of the player. These techniques often propose preferred plot points to direct a player through to the next point of the story (Giannatos et al. 2011). Other data-driven techniques support providing the player with a plot choice that they will most likely enjoy in a multi-branched story (Yu and Riedl 2013). Alternatively, plot points can be determined by a developer, and the quest generation system will provide content that the player engages with the most (Khaliq and Watson 2018).

Player modeling has been done in various ways. One way is to aggregate player data and classify the players into different archetypes. These clusters can be associated with personality traits (Levy et al. 2015), or with general play styles (Normoyle and Jensen 2015). The second way is to do individual player profiling (Thue et al. 2007, Thue 2006, de Lima, Feijó, and Furtado, 2016). These approaches focus on identifying the play style of the player. Additionally there has been work which approximates the skill level of the player, and provides dynamic difficulty adjustment (Zook

et al. 2012). In general, these player modeling approaches have the drawback of being genre dependent. To address this, there has been some work on a general player modeling technique (Shaker, Shaker, and Abou-Zleikha 2015).

Previous Work

My previous work focused on an application of AI to fighting games to better understand game balancing and game design (Yu and Sturtevant 2019). The goal of the AI is not to play the game the best, but instead is an AI tool for the developer to be able to identify moves that could be considered over- or under-powered. This AI simulates the game by formalizing the problem as a matrix game and calculating the Nash Equilibrium from the matrix game. If a move is overpowered, there will be a pure strategy and there is no variation in the action that the player should take. If there is a mixed strategy then the game can be considered more balanced because there are more options for the player.

Additionally, I have recently been working on a specific sub-problem of quest generation, where the goal is to generate quests in response to unexpected game states reached by the player. In large games, players can make unforeseen choices which causes the game to break and prevent the player from progressing or locking them out of sections of the game. I proposed an adaptive system based on a novel quest definition which can plan around the current game state and still provide quests which meet a target goal. The quest system will make changes to the current game state when the preconditions for actions cannot be met by the player, in order to repair the game state and allow for the player to progress.

Future Work

There has been limited work on the combination of an AI director, player model and a quest generator. All three of these pieces are necessary to allow for the player to have full agency in a game. When a player makes a meaningful choice, it should directly affect the current game state, which would affect the kinds of quests that could be proposed to the player next.

The first piece of work to be done is to determine a correct player model to use for the AI director. Thue et al. (2007) used a questionnaire at the beginning of the game to model a player. This approach is also used by Ramirez and Bulitko (2015), where they use a vector to model the player and use the vector to influence the choices which will be presented to the player in an interactive narrative. The authors determined that the results of using a player model in conjunction with an AI director is inconclusive. I think that this approach is still promising, but the player model should instead be determined from the player actually playing the game rather than a set of questions presented to the player before they have a chance to interact with the story. I propose that the best way to model a player is to build upon the goal-directed modeling approach, which models the player based off of the goal that the player is trying to achieve (Thue 2006). In simple games such as Rubiks cube, where the only goal is to solve the puzzle, the player model can be easily determined. However, in a fully realized game world, the goal of the player becomes less clear if they are not trying to accomplish an objective already present in the game. This idea creates prediction problem, where a system will have to predict what the goal of the player is based off of the previous interactions the player has had in the game. I believe this prediction problem can be solved using a combination of the objectives that the player has already completed and the actions that the player is currently taking.

The second piece of work is to determine the correct system for an AI director. The AI director I am proposing would be part of the player modeling category, which focuses on personalizing the experience to the player. I believe that the personalization aspect is the correct focus for the AI director because it will allow for the player to have better agency in the narrative of the story. If the AI director is focused on authorial intent, then it could prevent the player from making meaningful choices within the game in an effort to provide the authored narrative. If the AI director is focused on virtual character autonomy, then it could prevent player choice by requiring that the nonplayer characters have a cohesive narrative. AI directors use a wide range of approaches but I believe that a simple branching technique will be the most powerful, similar to the one proposed by Thue et al. (2007). The AI director can simply be aligning the player types to the types of high-level interactions that a player likes to engage with. However, I believe that the branching will divide the player into clusters of actions, instead of always providing a single, best action. This will allow for the AI director to provide variation to the player, while still proposing actions that would be best suited for the player. To chose which branch would be the best, I would try a Monte Carlo tree search approach, and score the outcomes of each potential outcomes based off the preference of the player.

The last piece of work is to dynamically generate a quest or objective for the player to complete. This provides the player with a in-game reward, and allows the player to keep engaging with the game. The AI director determines the main action of quest that the player will most likely enjoy, but the rest of the quest must be determined. The quest that is generated from this recommendation should be able to affect the state of the game. If the quest actively allows the player to change information about the world, then it allows the player to start building a custom narrative and increases player agency. To start out, I would try simulation based approach. These approaches prove popular as they allow the quest to be more narratively cohesive (Samuel et al. 2016, Breault, Ouellet, and Davies 2018). Combining the procedural content generation with the AI director will create a system similar to the one proposed by Robertson and Young (2015), which will allow the AI director to modify the game directly to generate quests.

Overall, I would like to focus on these three areas over the course of my PhD. I believe that a combination of these techniques can create a more powerful narrative experience than any of them can independently through the allowance of higher player.

References

Breault, V.; Ouellet, S.; and Davies, J. 2018. Let CONAN tell you a story: Procedural quest generation. *arXiv:1808.06217 [cs]*. arXiv: 1808.06217.

Cardona-Rivera, R. E.; Robertson, J.; Ware, S. G.; Harrison, B.; Roberts, D. L.; and Young, R. M. 2014. Foreseeing Meaningful Choices. In *Proceedings of the Tenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 7.

de Lima, E. S.; Feijó, B.; and Furtado, A. L. 2016. Player Behavior Modeling for Interactive Storytelling in Games. In *In the Proceedings of SBGames 2016*, 11.

Giannatos, S.; Nelson, M. J.; Cheong, Y.-G.; and Yannakakis, G. N. 2011. Suggesting New Plot Elements for an Interactive Story. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 6.

Khaliq, I., and Watson, Z. 2018. The Omni Framework: A Destiny-Driven Solution to Dynamic Quest Generation in Games. In *2018 IEEE Games, Entertainment, Media Conference (GEM)*, 306–311. Galway: IEEE.

Levy, L.; Solomon, R.; Gandy, M.; Moore, J.; Way, J.; and Liu, R. 2015. Actions Speak Louder Than Words: An exploration of game play behavior and resutls from traditional assessments of individual differences. In *Proceedings of the 10th International Conference on the Foundations of Digital Games*.

Normoyle, A., and Jensen, S. T. 2015. Bayesian Clustering of Player Styles for Multiplayer Games. In *Proceedings of the Eleventh AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 7.

Productions, S. P. 2014. Infamous second son. Playstation 4.

Ramirez, A., and Bulitko, V. 2015. Automated Planning and Player Modeling for Interactive Storytelling. *IEEE Transactions on Computational Intelligence and AI in Games* 7(4):375–386.

Riedl, M. O., and Bulitko, V. 2013. Interactive Narrative: An Intelligent Systems Approach. *AI Magazine* 34(1):67.

Riedl, M. O., and Young, R. M. 2010. Narrative Planning: Balancing Plot and Character. *Journal of Artificial Intelligence Research* 39:217–268.

Robertson, J., and Young, R. M. 2015. Automated Gameplay Generation from Declarative World Representations. In *In the Proceedings of the Eleventh AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 7.

Samuel, B.; Ryan, J.; Summerville, A. J.; Mateas, M.; and Wardrip-Fruin, N. 2016. Bad News: An Experiment in Computationally Assisted Performance. In *9th International Conference on Interactive Digial Storytelling*, volume 10045, 108–120. Cham: Springer International Publishing. Series Title: Lecture Notes in Computer Science.

Shaker, N.; Shaker, M.; and Abou-Zleikha, M. 2015. Towards Generic Models of Player Experience. In *Proceedings of the Eleventh AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 7.

Thue, D.; Bulitko, V.; Spetch, M.; and Wasylishen, E. 2007. Interactive Storytelling: A Player Modelling Approach. In *In the Proceedings for the Third Artificial Intelligence for Interactive Digital Entertainment Conference*, 6.

Thue, D. 2006. Modelling Goal-directed Players in Digital Games. 6.

Yu, H., and Riedl, M. O. 2013. Data-Driven Personalized Drama Management. In *Proceedings of the Ninth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 191–197.

Yu, K., and Sturtevant, N. R. 2019. Application of Retrograde Analysis on Fighting Games. In *2019 IEEE Conference on Games (CoG)*, 1–8. London, United Kingdom: IEEE.

Zook, A.; Lee-Urban, S.; Drinkwater, M. R.; and Riedl, M. O. 2012. Skill-based Mission Generation: A Data-driven Temporal Player Modeling Approach. In *Proceedings of the The third workshop on Procedural Content Generation in Games - PCG'12*, 1–8. Raleigh, NC, USA: ACM Press.