Creating a Player-Centric Dynamic Game AI

I. Problem Description

In video games, AI is primarily used to simulate human-like intelligence while keeping within the restrictions of the game environment. Human beings playing against a computer game are often asked to select a "level" of difficulty that they feel would be challenging enough to be enjoyable but at the same time easy enough to not be overwhelming. In this paper, we have detailed our implementation of an AI game playing agent that uses the player's moves to estimate the skill level of the player. The agent then matches this skill level, ensuring that the player is challenged without being demoralized.

The ability to create intelligent game agents that can match the player's preferences and abilities without being explicitly asked to enables us to create more interesting, user-friendly games that are potentially addictive. This has applications not only in entertainment, but also in education. With an AI that adapts to user's skill level, players will always play against an opponent closely matched to their own skill level. In educational games that run the risk of being boring, too difficult or too easy, such a Game AI would keep people interested and improve their gameplay and learning experience.

II. Related Work

Dynamically catering the difficulty of a game to a player's skill level has been the subject of several recent research initiatives. Gratch and Marcella[1] demonstrated how emotions and personalities in the AI agent could impact the reasoning process deciding which actions are taken leading to a better game experience for the human player. Zook and Riedl [2] used tensor factorization to assess changes in a player's skill level and then adjust the game difficulty accordingly in a turn-based role-playing game. They were able to explain 86\% of the variance in player's skill with 27 players. However, this model did not predict future player performance. Tan Tan Tay[3] used two behavior-based controllers, an adaptive uni-chromosome controller (AUC) and an adaptive duo-chromosome controller (ADC), on a real-time car racing simulator. All these algorithms have the advantage that they could be trained online, but they only worked when there was a single player against a single AI.

We worked with multiple AI agent based game players in a turn-based game environment to predict future human player performance and adjust difficulty accordingly. We wanted the agents to choose strategies based on their evaluation of the human player's skill level. The AI was also developed to adapt to the most current skill level of the player, whether that be better or worse than previously assessed.

III. Method

A. Isolation

Isolation is a multiplayer, turn taking puzzle game played on an NxN grid. Each player takes a turn to select a position on the board such that there is a clear line (horizontal, vertical or diagonal) between the current position and the new position. The version of Isolation developed for this project was implemented on two platforms - Android and Mobile Web. The game features an 8x8 board and 3 players - one of whom was a human being while the other two were AI agents. The two agents play independently and do not attempt to compete or collaborate.

B. Implementation

The AI agents use the minimax algorithm to do a lookahead up to a depth of 2 and choose the best action. While this seems like a small number, the high branching factor and inability of pruning prevents the AI from evaluating the tree any further in the interest of providing a move to the player in a reasonable amount of time. Alpha-Beta pruning was not used since that technique is used in order to identify and prune the branches of the game tree whose utility is such that it ensures the branch doesn't lie on the optimal path. It makes the search along the tree faster. Since our game was not focussed on finding the optimal path along the tree but instead to maximize gameplay experience, we did not use pruning in the player modeling agent. Alpha beta pruning was, however, used to create an optimal AI agent to be used in the evaluations.

In order to evaluate the utility of a non-leaf node in the tree, we used a combination of the number of free spaces on the board and the number of legal moves available to each player. The number of free spaces on that board gives an indication of how much of the board is filled up. It is not always necessary that the legal moves would cover every square on the board and this metric is used to find how "isolated" a player is and whether there is opportunity to move to a more open area. The number of legal moves available to each player are used to not only maximize the possibilities of oneself, but also minimize the possibilities for others. The agents used (number of current player's moves - number of enemy's moves)/(num empty spaces +1).

In order to estimate the skill of the user, we used multiple features available during gameplay include the number of positions free on the board at any given point of time, the number of legal moves available to each user and the average utility (as returned by the evaluation function) of the moves chosen by the user. Each move the player makes is evaluated using the same evaluation function the AI agents used. The agent, during its turn, selects a move that evaluates to slightly greater than the average utility of all the moves made by the user. Using the same function to evaluate the utility of a position on the board allows us to compare the relative "goodness" factor of the positions.

IV. Evaluation

Two different user studies were conducted - one to gauge the user's response to the player modelling agent and another to map how the agent increases its own skill alongside that of the user.

Identify applicable sponsor/s here. If no sponsors, delete this text box (*sponsors*).

A. Evaluating Enjoyability

24 human players were divided into two groups. Each player was asked to play the game multiple times before the study so as to ensure that they were comfortable with the rules and understood what it would take to win. Afterward, one group was asked to play Isolation where the AI agent did not do player modelling but selected the next move randomly from all possible legal moves. The second group played Isolation against two AI agents that performed player modeling as described in the previous sections. Each user was then asked to rate the game in terms of enjoyability and simplicity. Figure 1. shows the average of how users from each group rated the game on enjoyability and simplicity on a scale of 1 to 5.

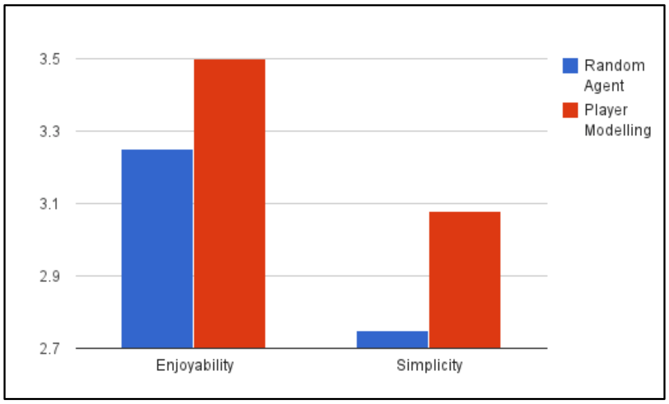

B. Evaluating Skill

Another user study was conducted with the Android agent in which 11 participants played 15 games: 5 with each of 3 AIs. The 3 AIs used were a random agent, the "best move" agent with alpha-beta pruning, and the dynamic AI. The percentage of players that won each game was recorded as a function of the game number to account for any experience the player gained after playing against the AI. Additionally, the order in which the users played against the AIs was randomly determined. The chart below shows the results of the user study.

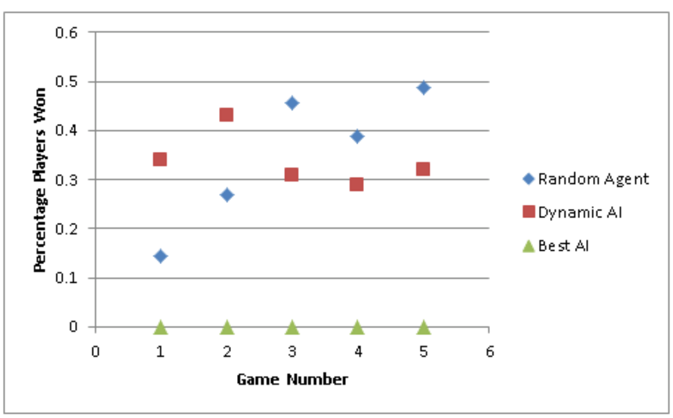

From this chart, we can determine that the Random agent was relatively easy to beat and got easier to beat the more familiarity people had with it. The best AI was too tough and not a single human could beat it in the small study. However the dynamic AI provided a relatively constant win-loss ratio that was not overly challenging, but not too easy either.

V. Discussion

Users described the player modeling agents as smart. We found that the moves made by the AI agents tended to be anthropomorphized by the human player into being "defensive" or "aggressive" as the game continued depending on whether each individual agent was seen to make more moves on the average to barricade the human into a bounded area or acquire squares around the board while dodging the human more during the duration of the game. The greedy algorithm would also occasionally produce mirroring moves to the human for some duration during the game. Human player received this behaviour favorably and described it as "sneaky".

The learning experience and competitiveness of the human players also improved over time. While players were, over time, able to consistently defeat the random AI, very few were able to defeat the dynamic-modeling AI. As the players played more games and gained confidence, their skill increased. They random agent, incapable of gauging skill, was defeated. The dynamic player modeling AI, however, increased it's own skill to match that of the player, causing the average percentage of wins to be nearly constant.

Additionally, many of the moves made by the agent seemed "illogical", degrading the user experience. They thus found the random AI less enjoyable. When people win they want to feel a sense of accomplishment. Beating a series of randomly generated moves does not provide this fulfillment in the same way beating a clever computer AI does.

VI. Future Work

The next step in perfecting and eventually publishing our implementation of the Isolation game would be to perform a more comprehensive user study. We propose to increase the size of our study group and to investigate the effect of our AI on more parameters, such as improvements in user performance over time, total time spent playing the game, and perceived self-confidence. The idea here would be to vary the difficulty ratio of the AI to the player and the AI's strategy (aggressive, passive, etc.) and observe the effects. A further advantage of a three player Isolation would be that with two of the players being AI agents we could attempt to create an awareness amongst the AI players so that they can choose to compete or collaborate with each other or the human player to win and improve gameplay. Depending on how individual users respond to each tactic, the AI should be able to choose for itself what strategy to implement.

Once the "best" version of the AI and it's corresponding evaluation function have been chosen based on the combinations that generate the most user enjoyment in the greatest number of people in these user studies, the subsequent step would involve improving upon the graphics and gameplay of the game in order to make it more marketable. Additionally, in order to provide a more educational look at the AI behind the isolation game, a visualization for the alpha-beta pruning would be paired with the game so that interested players can see how the AI makes its decisions. We intend for Isolation to be not only enjoyably addictive, but also secretly educational. Our long term goal is to garner the interest of a company making educational games and apply our dynamic player modeling approach to help them improve user learning and satisfaction.

REFERENCES

1. J. Gratch and S. Marsella. "A Domain Independent Framework for Modeling," Journal of Cognitive Systems Research, vol. 3, no. 4, pp. 269-306, 2004.
2. A. E. Zook and M. O. Riedl, "A Temporal Data-Driven Player Model for Dynamic Difficulty Adjustment," The Eighth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, 2012.
3. C. H. Tan, K. C. Tan and A. Tay, "Dynamic Game Difficulty Scaling Using Adaptive Behavior-Based AI," IEEE Transactions on Computational Intelligence and AI in Games, vol. 3, no. 4, pp. 289-301, 2011.
4. R. D. Laws, "Knowing Your Players," Robin's Laws of Good Game Mastering, 1st ed. Austin, Steve Jackson Games, ch. 2, sec. 1, pp. 4, 2002.
5. D. Ramirez-Cano, S. Colton, and R. Baumgarten, "Player classification using a meta-clustering approach." Proceedings of the 3rd Annual International Conference Computer Games, Multimedia & Allied Technology. 2010.