

Design Criteria for Challenge Balancing of Personalised Game Spaces

Sander Bakkes
University of Amsterdam
Intelligent Systems Laboratory Amsterdam
The Netherlands
s.c.j.bakkes@uva.nl

Shimon Whiteson
University of Amsterdam
Intelligent Systems Laboratory Amsterdam
The Netherlands
s.a.whiteson@uva.nl

ABSTRACT

This short paper focuses on games that can tailor the provided game experience to the individual player (personalised games), typically by effectively utilising player models. A particular challenge in this regard, is utilising player models for assessing online (i.e., while the game is being played) and unobtrusively which game adaptations are appropriate. We propose design criteria for approach for personalising the space in which a game is played (i.e., levels) – to the end of tailoring the experienced challenge to the individual player *during actual play* of the game. Our approach specifically considers two persisting design challenges, namely *implicit user feedback* and *high risk of user abandonment*. Our contribution to this end is proposing a clear separation between (intelligent) offline exploration and (safety-conscious) online exploitation. We are presently assessing the effectiveness of the developed approach in an actual video game: INFINITE MARIO BROS. [5]. To this end, we have enhanced the game such that its process for procedural-content generation allows the game spaces (i.e., levels) to be personalised during play of the game. We use data from intelligent offline exploration to determine both a model of experienced challenge as well as safe level design parameters for use on new players. Online, we use a gradient ascent algorithm with designer-specified domain knowledge to select the next set of level design parameters.

Keywords

Personalisation, game spaces, challenge balancing, video games, implicit user feedback, user abandonment

1. INTRODUCTION

Ideally, artificial intelligence (AI) in games would provide satisfactory and effective game experiences for players regardless of gender, age, capabilities, or experience [3]; ideally, it would allow for the creation of personalised games, where the game experience is continuously tailored to fit the individual player. However, a prevailing limitation in the context

of game personalisation, is that learning effective behaviour online (1) requires an inconveniently large number of trials, (2) is generally highly resource-intensive, and (3) usually generates many inferior solutions before coming up with a good one [2].

As such, achieving the ambition of creating personalised games requires the development of novel techniques for assessing online and unobtrusively which game adaptations are required for optimizing the individual player’s experience. To this end, we have proposed an approach for online game personalisation that reasons about uncertainty about the player probabilistically, specifically using Gaussian Process optimisation. It considers game personalisation as an optimisation problem, and details design considerations for effectively personalising the space of a game, while the human user is interacting with it. Specifically, the proposed approach considers two persisting design challenges, namely *implicit user feedback* and *high risk of user abandonment*. Our goal is to demonstrate how to personalise game spaces *during play of the game*, such they that are appropriately challenging to the individual user. To this end, our approach is implemented in the actual video game INFINITE MARIO BROS.

2. DESIGN CRITERIA

The goal of our approach is to online generate game spaces such that the spaces optimise player challenge for the individual player. Our contribution to this end – derived from the design criteria discussed below – is proposing a clear separation between (intelligent) offline exploration and (safety-conscious) online exploitation.

Our approach to challenge balancing for personalised game spaces consists of three phases. In Phase 1, we learn a *global safe policy* across users for initialising the game, while generating a set of training instances (offline). In Phase 2, we learn a *feedback model* across users from the generated set of training instances (offline). In Phase 3, we *personalise the game space* to the individual player (online)¹.

Two main challenges persist for the task of challenge balancing in personalised game spaces, namely (1) only *implicit feedback* on the appropriateness of the personalisation actions are available, and (2) there is a *high risk of user*

¹We kindly refer the interested reader to a detailed description of the approach, published in the proceedings of the Fifth workshop on Procedural Content Generation in Games (PCG 2014).

abandonment when inappropriate game personalisation actions are performed. We propose the following four design criteria to address these challenges.

2.1 Intelligently generated set of training instances

To provide an effective basis for online personalisation, it is necessary to intelligently generate a set of training instances (Phase 1). Indeed, given a typically multi-dimensional parameter space, it is infeasible to acquire training instances that cover the entire space of the investigated personalisation problem (particularly when employing human participants to generate the instances). As such we propose to sequentially acquire instance data from those data points in the parameter space that maximise the upper-confidence bound on the expected experienced challenge across users. This is achieved via the Gaussian Process Upper Confidence Bound (GP-UCB) algorithm [7], a Bayesian approach to global optimization that captures its uncertainty about the fitness function in the form of a distribution, represented using a Gaussian process, over the space of possible fitness functions. This distribution is used to select the next point to evaluate based on the principle of Upper Confidence Bounds [1], which focuses the search on points that are either highly promising or highly uncertain. The result of evaluating the selected point is then used to update the distribution using Bayes rule, and the process repeats.

2.2 Implicit feedback model

To address the problem of implicit feedback, we propose to learn a *feedback model* from the intelligently generated set of training instances (Phase 2). In a typical video gaming domain, the only input which the personalisation agent receives, are observations on the player interacting with the generated game space (e.g., how successful is the player in dealing with a specific game element, such as jumps). As a result, for game personalisation, a mapping from gameplay observations to player experience needs to be learned. To this end, while generating the set of training instances, we gather labels on the player's gameplay experience with regard to the dependent variable (DV) 'experienced challenge'². On the basis the labelled training instances, we train a random forest decision-tree classifier [4] for classifying the experienced challenge of an observed gameplay observation O . The random forest classifier consists of a combination of tree classifiers where each classifier is generated using a random vector sampled independently from the input vector, and each tree casts a unit vote for the most popular class to classify an input vector. The advantage of using random forest classification, is that it does not solely output a classification, but returns a probability distribution of the classification, which we will employ as a reward signal for online game-space personalisation (Phase 3).

2.3 Smart cold-start initialisation

To address the problem of user abandonment in actual play of the game, we propose to initialise a new game with a challenge level that is most appropriate in expectation across users. We achieve smart cold-start initialisation by learning

²In accordance with the work of Shaker *et al.* [6], we also gather labels on the DV's engagement and frustration. These DV's will be employed for multi-objective learning in future work.

a *global safe policy* while intelligently generating the set of training instances, with the GP-UCB algorithm (in Phase 1). When the global safe policy has been learned offline, is applied at the start of an online game personalisation session (Phase 3).

2.4 Guided exploration

Also, to address the problem of user abandonment in actual play of the game, we propose to guide state-space exploration by using *expert domain knowledge* and a *model of user abandonment* (Phase 3). With regard to a model of user abandonment, one may consider a personalisation action's expected reward (derived from the feedback model's classification of appropriateness of the action to the individual user), as inversely correlated to the probability of the player abandoning the game. This design decision has as consequence that with a high probability of abandonment one permits the personalisation method to perform a relatively exploratory action (assuming that if the inappropriate state persists, the user will abandon in any case). On the other hand, with a low probability of abandonment one restricts the algorithm to performing a relatively exploitative action (assuming that the algorithm may be approaching a global maximum). With regard to guiding the exploration process by domain knowledge, we consider that if *a priori* it is known what effect certain action features have on the overall player experience, then this knowledge can be exploited intelligently. As such, one should ideally let expert domain knowledge *guide the exploration* of the state-space so that it is much more effective. Specifically, the domain knowledge can be formalised as a set of rules that define how, and under which conditions, a specific action feature should be adapted. Here, we posit that such expert domain knowledge is available by the designer of a game, or can be computationally derived from gamelogs.

We are presently performing experiments that validate the approach in studies with human game players.

3. REFERENCES

- [1] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *The Journal of Machine Learning Research*, 3:397–422, 2003.
- [2] S. C. J. Bakkes, P. H. M. Spronck, and H. J. Van den Herik. Rapid and reliable adaptation of video game AI. *IEEE Transactions on Computational Intelligence and AI in Games*, 1(2):93–104, 2009.
- [3] D. Charles, M. McNeill, M. McAlister, M. Black, A. Moore, K. Stringer, J. Kucklich, and A. Kerr. Player-centered design: Player modeling and adaptive digital games. In *DiGRA*, pages 285–298, 2005.
- [4] T. K. Ho. Random decision forests. In *Proc. of the 3rd Int. Conf. on Document Analysis and Recognition*, pages 278–282, 1995.
- [5] M. Persson. Infinite mario, 2009. [Online] Available: <http://www.mojang.com/notch/mario>.
- [6] N. Shaker, S. Asteriadis, G. N. Yannakakis, and K. Karpouzis. A game-based corpus for analysing the interplay between game context and player experience. In *Affective Computing and Intelligent Interaction*, pages 547–556. Springer, 2011.
- [7] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.