

Dynamic Difficulty Adjustment

Olana Missura

Informatik III, Universität Bonn, Germany
olana.missura@uni-bonn.de

Abstract

One of the important questions that needs to be answered during a game development process is how to keep the game engaging for a wide spectrum of players and for a long time. The “interestingness” of a game depends on many different parts such as the storyline, the right amount of interaction, the attractive presentation, etc. The perceived difficulty of the game also contributes to the time that any given player spends playing it. Games that are too easy or too difficult become boring or frustrating fast.

To get the difficulty “right” the developer needs not only to balance all the game parameters, but also to cater for players with very different skill levels. Providing a game with an ability to automatically scale the difficulty depending on the current player would make the games more engaging over longer time. While a few commercial games boast about having such a system and the results of several machine learning techniques applied to dynamic difficulty adjustment (DDA) problem have been published, to the best of our knowledge it has not been researched as a general learning problem.

There is a belief that providing players with a personalised, adaptive experience can sustain their attention and keep their interest for longer [2, 3]. One useful way to create a personalized game is to implement a DDA system. We aim at creating a mechanism to allow developing games that provide challenges of the “right” difficulty, i.e., such that players are stimulated but not overburdened. For that, we investigate how DDA can be formulated and solved as a learning problem. A general technique for this problem has natural applications in the huge markets of computer and video games but can also be used to improve the learning rates when applied to serious games.

This work is about the design of a DDA system that is driven by the following requirements:

- **Universality:** The resulting system should be applicable to as many games as possible, independently of their type, structure, or features;
- **Non-intrusiveness:** If desired, the resulting system should work in a transparent way, in real-time, not requiring an interaction with players;
- **Feasibility:** There should be guarantees about the resulting system’s performance.

The main part of the thesis is organized into three chapters, each of these dealing with the same question but applying to different and progressively more difficult settings. Here in brief are some of the main contributions.

First, we start with a precisely defined world of board games and describe a DDA heuristic. Specifically, we propose that in a setting with opponents of the same kind and clearly defined ‘moves’, the DDA mechanism chooses the next move based on the quality of a player’s moves so far. We describe the experiments demonstrating the power of this heuristic, but also go in a particular detail about its limitations.

Next, we drop the assumptions about the nature of a game or opponents and require only that players have an ability to adjust the difficulty in real time. A much wider group of games fits this description. The main part of this chapter is about a learning algorithm that utilizes the data collected during the playtesting phases to construct a model or a set of models of players. After the models are created, any new player is provided with a difficulty adjustment that is specified by his or her model.

Finally, we formalise DDA as a sequential interaction between a *master* and a *player* in which the *master* tries to predict the most appropriate difficulty setting. As the *player* is typically a human with changing performance depending on many hidden factors as well as luck, no assumptions about the *player* can be made.

To be able to theoretically investigate DDA, we view it as a game between a *master* and a *player*, played on a partially ordered set modelling the ‘more difficult than’-relation. The game is played in turns where each turn has the following elements:

1. the game *master* chooses a difficulty setting,
2. the *player* plays one ‘round’ of the game in this setting, and
3. the game *master* experiences whether the setting was ‘too difficult’, ‘just right’, or ‘too easy’ for the player.

The *master* aims at making as few as possible mistakes, that is, at choosing a difficulty setting that is ‘just right’ as often as possible. We present an algorithm for the *master* with theoretical guarantees on the number of mistakes in the worst case while not making any assumptions about the *player*.

The crucial part of the analysis consists of representing a set of difficulty settings as a (finite) partially ordered set. As we can not make any assumptions about the *player*, we compare our algorithm theoretically and empirically with the best statically chosen difficulty setting, as is commonly the case in online learning [1]. This baseline has a natural relation to the real life: it reflects the best setting that the player would choose if she knew the game and her abilities.

While the theoretical results show that the *master’s* algorithm’s performance depends on the properties of the partially ordered set and in a good case is on par with a full information predictor (i.e. the one who at each round receives information about *all* states rather than one), the empirical studies we have conducted demonstrate even better results.

In sum this work presents several new methods for the DDA in order of increasing generality, finishing with a theoretically sound, online DDA algorithm that does not impose any assumptions on games or players that it can be applied to.

References

- [1] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [2] K. Gilleade, A. Dix, and J. Allanson. Affective videogames and modes of affective gaming: assist me, challenge me, emote me. In *Proc. DIGRA*, volume 2005. Citeseer, 2005.
- [3] P. Sweetser and P. Wyeth. GameFlow: a model for evaluating player enjoyment in games. *Computers in Entertainment (CIE)*, 3(3):3, 2005.