Proceedings, The Eleventh AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE-15)

Exploring Player Trace Segmentation for Dynamic Play Style Prediction

Josep Valls-Vargas[1], Santiago Ontañón[1] and **Jichen Zhu[2]**
[1]Computer Science, [2]Digital Media
Drexel University
Philadelphia, PA 19104
josep.vallsvargas@drexel.edu, santi@cs.drexel.edu, jichen.zhu@drexel.edu

Abstract

Existing work on player modeling often assumes that the play style of players is static. However, our recent work shows evidence that players regularly change their play style over time. In this paper we propose a novel player modeling framework to capture this change by using episodic information and sequential machine learning techniques. In particular, we experiment with different trace segmentation strategies for play style prediction. We evaluate this new framework on gameplay data gathered from a game-based interactive learning environment. Our results show that sequential machine learning techniques that incorporate predictions from previous segments outperform non-sequential techniques. Our results also show that too fine (minute-by-minute) or too coarse (whole trace) segmentation of traces decreases performance.

Introduction

The field of game analytics has received a lot of attention in recent years from both industry and academic research. Analyzing game data is nowadays a common practice widely used to validate level design or improve customer conversion rates. In general, the goal is to extract knowledge from the player's behavior in order to improve their experience. In order to achieve this goal, Artificial Intelligence (AI) techniques can be used to make interactive systems more adaptive, responsive and intelligent. Player modeling is a crucial component of adaptive computer games in the domains of digital entertainment (Togelius, De Nardi, and Lucas 2006; Yannakakis and Maragoudakis 2005) and education (Magerko, Heeter, and Medler 2010). Adaptive computer games can use a model of the players' skill or preferences to adapt the game in order to maximize player engagement and satisfaction (Riedl et al. 2008; Thue et al. 2007).

Currently, most research on player modeling assumes that the player property being modeled remains the same during gameplay. For instance, in learning sciences and educational games, researchers have traditionally relied on self-reported data to determine learner's characteristics, such as motivation (Magerko, Heeter, and Medler 2010). The data collected once is typically assumed to describe the learner over time. In research where actual gameplay behavior data are used for player modeling, the same assumption holds. For instance, procedural content generation systems such as those by Pedersen, Togelius, and Yannakakis (2009) or Togelius, De Nardi, and Lucas (2006) observe play style to build a model of players' preferences and generate levels to satisfy them, assuming players' preferences will not change within the play session. Similarly, interactive narrative systems such as Passage (Thue et al. 2007) and C-Dragger (Sharma et al. 2010), maintain a player model of play style used to select content to present to the player. Although models are updated on certain events they do not explicitly model the transitions between play styles over time.

However, in our recent study of how learners interact in an educational game (Valls-Vargas et al. 2015), evidence suggests that only 25% of our participants (n=55) adopted the same play style throughout a play session. In contrast, the majority of the players shifted between different play styles (such as exploring, goal-seeking or being uninterested) (as illustrated in Figure 1). *If we cannot safely assume that the players will play a game with a consistent play style, there is a need for more flexible player modeling approaches that can capture the temporal change in players' behavior.*

In this paper, we propose a novel player modeling framework that relaxes the assumption of a fixed style for each player and addresses the problem of predicting a dynamic play style. Our framework captures changes in a player's behavior by using episodic information and time interval models in a sequential machine learning approach that learns multiple models over time. In particular, we present our findings on the key problem of how to segment gameplay data for play style prediction — what is the appropriate level of granularity to build the dynamic player model. Our results show that best performance in predicting dynamic play style is obtained segmenting the trace at an intermediate granularity. If the trace segmentation is too fine-grained (minute-by-minute windows), then not enough information is present in each segment to make meaningful predictions. If the trace segmentation is too coarse-grained (whole game), then the dynamics of shifting play style are not captured properly. Additionally, we show that sequential machine learning techniques that incorporate predictions from previous segments outperform non-sequential techniques.

Our data is captured from a relatively small-scale educational game called *Solving the Incognitum* where players are required to complete a series of quests to win the game.

Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

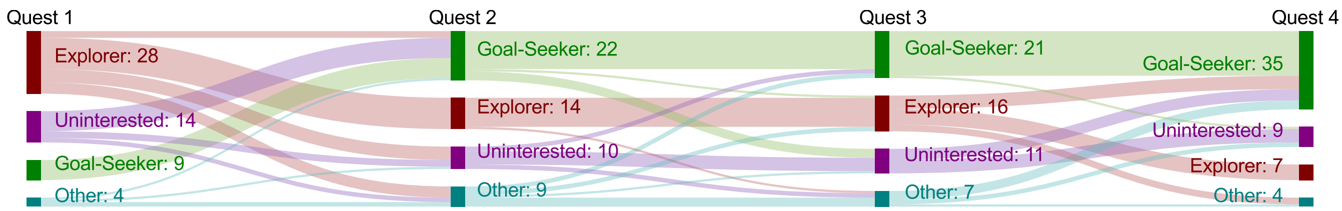

Figure 1: Illustration of the play style shifts observed among players in the user study (n=55). Each vertical block represents a segment (quest or episode) of the game and is divided into 4 play styles with their size proportional to the number of participants. The flow between blocks represent shifts on the observed play style over time.

The player property we model is their *play style*, derived from motivational research in education such as Achievement Goal Theory (Elliot and McGregor 2001). In particular, we use the play style classification of *goal seeker*, who primarily performs actions to win the game, and *explorer*, who primarily performs actions to explore their own goals in the game, not necessarily concerned about winning the game (Foster 2011; Heeter 2009). We believe that our framework for modeling dynamic play style can be generalized to other domains where the player property of interest may change.

In the rest of the paper, we first present our framework and describe how we extend the game analytics pipeline (Canossa 2013) to support a sequential machine learning approach to better capture dynamic play styles. After briefly describing the game *Solving the Incognitum* as our domain, we present our experiments on different strategies to segment gameplay data and report our findings.

Background

Player modeling is an active research topic in the field of digital entertainment. In a game environment, a player model is an abstracted description of a player capturing certain properties of interest such as preferences, strategies, strengths or skills (Van Der Werf et al. 2003). There has been numerous research on modeling player preferences in order to maximize engagement (Riedl et al. 2008; Thue et al. 2007), or provide better non-player-character AI (Weber and Mateas 2009). Player modeling has also been used to provide insights of player behavior patterns to game designers (Tychsen and Canossa 2008). For an overview on player modeling, the reader is referred to recent surveys of the area (Smith et al. 2011; Machado, Fantini, and Chaimowicz 2011).

Many approaches to player modeling build upon the *game analytics pipeline* (Canossa 2013). Within this framework, there is abundant published research regarding variable and feature selection but, to the best of our knowledge, there is little work studying segmentation strategies for game telemetry. The most common approach for segmentation strategies are time windows. For example, Bifet and Gavalda (2007) studied the use of window sizes of varying length for segmenting changing time-series. Another approach is to perform episodic segments. Gow et al. (2012) experimented segmenting gameplay traces using enemy encounters in an unsupervised framework for play style clustering. In our work, we will explore the use of episodic segments with known bounds to segment game telemetry, and we will also compare against time window-based strategies.

Modeling Dynamic Play Style Over Time

In our prior research (Valls-Vargas et al. 2015), we recorded gameplay data of 75 participants playing the educational game *Solving the Incognitum* (details in next section), both in the forms of screen recording and automatically logging gameplay traces. Two researchers independently annotated the video recordings with observed play style labels. Then they discussed their annotations to resolve disagreements in the play style. The guidelines used for annotation were developed from pilot gameplay data by an expert on the field.

A key observation emerged from the study: a single play style label is not sufficient to capture how most players interacted with the game. Over the average play time of 21 minutes, we observed many players exhibited different play styles. Thus, we divided the gameplay video based on four episodes (consistent with four quests in the game) and annotated them individually. From the 75 participants, we discarded 19 data points due to missing data (video recording or game logs) and 1 due to annotator disagreement. Figure 1 summarizes the annotation of the participants' play style (n=55) between the four quests and illustrates how they switch between the different play styles. In the data, only 25% of our participants exhibited a consistend play style throughout the play session.

To automatically construct a player model that can predict the play style of the majority of players at a given moment, we thus need to lift the assumption that they exhibit a static play style. Our proposed framework is based on two key ideas: *episodic segmentation* and *sequential episodic prediction*. Intuitively, our approach splits gameplay traces into segments representing self-contained parts of a game (such as *quests*), which we call *episodes*, and performs play style predictions at episode granularity. Episodes segment the entire game at a granularity that is not too coarse (entire game), as for not capturing shifting play styles, and not too fine (minute-by-minute), as for not having enough data to do accurate predictions. Then, our sequential machine learning prediction scheme employs the play style prediction from the previous episode to inform the prediction of the current episode. Thus, our proposed framework consists of three main steps: 1) Data acquisition, 2) Episodic segmentation of gameplay traces, 3) Sequential episodic prediction. We describe each of the steps in the following subsections.

Data acquisition. Since play style is time-dependent, we require that temporal information (e.g., time stamps) is recorded during game play. We will use the term *gameplay trace* T^u to refer to the data collected for one individual player u during one gameplay session.

Episodic segmentation of gameplay traces. The gameplay trace is then segmented into different pieces. We experimented with three different segmentation approaches:

- *Time windows*: each gameplay trace T^u is split into a collection of non-overlapping segments using a given granularity (e.g., time windows of 1 minute). The result is a sequence $W_1^u, ..., W_n^u$ of time windows.
- *Cumulative windows*: given a time granularity (e.g., 1 minute), this strategy generates a series of segments $C_1^u, ..., C_n^u$, where $C_i^u = W_1^u + ... + W_i^u$, i.e., the gameplay trace from the beginning and up to time i.
- *Episodes*: this strategy splits the trace into self-contained *episodes*. The concept of a self-contained episode might vary from domain to domain. For example, episodes may correspond to dramatic beats (Mateas and Stern 2005), enemy encounters (Gow et al. 2012), puzzles (Sharma et al. 2010), quests (Valls-Vargas et al. 2015) or levels (Drachen, Canossa, and Yannakakis 2009). In our game, players need to complete a tutorial, and then four quests, which naturally yields a split into six episodes $E_1^u, ..., E_6^u$ (tutorial, four quests, plus all the optional game time players spent after completing the last quest), although our experiments only employ the episodes corresponding to the four quests.

Each resulting segment is converted into an abstract representation used for learning and applying the player model. In our experiments, we use feature vectors as our data representation. In the rest of this paper, "train a model to predict label L from episode E" refers to using a supervised machine learning algorithm (such as a Decision Tree or a Support Vector Machine) to predict L from the features in the feature vectors computed from E.

Since our framework employs supervised machine learning to create player models, gameplay traces need to be annotated with play style labels. In our work, we only require the *episodes* to be annotated (not the time windows, nor the cumulative windows). Thus, each episode E_i^u is labeled with the observed play style L_i^u (*goal seeker*, *explorer*, etc.) by hand, as we describe later. We also assigned a *global* play style label L^u to each gameplay trace, containing the play style that most adequately describes the player overall.

Sequential episodic prediction. Based on the segmentation strategies described above, we experimented with a number of play style prediction strategies:

- *Cumulative Prediction (Single-Label)*: this is a baseline approach that does not take into account the fact that play styles change over time.

  Training: a model is trained to predict the global label L^u from feature vectors computed from the whole trace T^u.

  Prediction: at run time, a prediction is cast using a feature vector computed from the current trace from the beginning of the game and up to the current point in time in a similar fashion as the cumulative windows C_n^u work.
- *Cumulative Prediction (Episodic-Labels)*: this approach employs the episodic play style labels L_i^u and trains different models (one per episode), in order to take into account that different episodes of the game might exhibit different distributions of play style.

  Training: for each episode i (quest 1, quest 2, etc.) a model CM_i is trained to predict the episode label L_i^u from feature vectors computed from a time window from the beginning to the end of episode i.

  Prediction: at run time, a prediction is cast using a feature vector computed from the current trace from the beginning of the game and up to the current point in time (as in C_n^u), and then giving it to the model CM_i corresponding to the current episode i.
- *Episodic Prediction*: inspired by work on learning from demonstration (Ross and Bagnell 2010), we treat each episode separately and learn a different model to predict play style based only on information in the episode.

  Training: for each episode i, a model EM_i is trained to predict label L_i^u based on a feature vector computed from just episode E_i^u.

  Prediction: at run time, each time an episode i is complete, a feature vector is generated with information from only this episode and used to predict a new play style label with model EM_i. We experimented with this strategy using episodic, time window and cumulative segmentations.
- *Sequential Prediction*: this approach combines two main ideas: 1) use of the play style prediction from the previous episode as input to the prediction for the current episode, and 2) combination of models trained at episode granularity with models trained at more fine-grained time windows for accurate prediction.

  Training: two types of models are trained in this approach. First, for each episode i, a model SEM_i is trained to predict label L_i^u based on episode E_i^u and on the previous label E_{i-1}^u (for the first episode, a special empty label is provided). Second, for each time window j, a model STM_j is trained to predict the play style based on time window W_j^u and the label E_{i-1}^u, where i is the episode where time window j falls.

  Prediction: at run time, each time the player completes an episode i, an internal episodic play style prediction EL_i^u is generated using model SEM_i and the play style prediction from episode $i-1$ (i.e., the prediction from the previous episode is one of the features that is given as input to predict the next play style label). The final play style prediction is generated using the time window models. Assuming that the current episode the player is playing is i, the time window models are used to generate a play style prediction at regular intervals (same time granularity as the time windows). Using the data form the current time window j and EL_{i-1}^u, model STM_j is used to predict the play style of the player WL_j^u for the current time window. This is illustrated in Figure 2. For this approach,

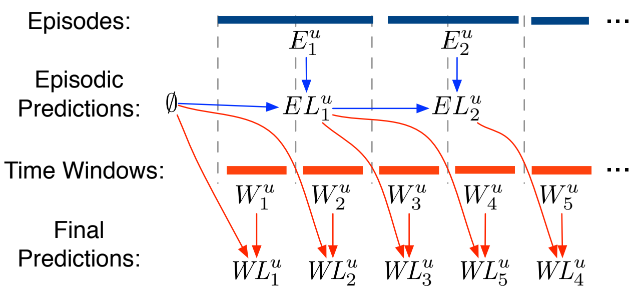

Figure 2: Illustration of the sequential prediction approach. Thick blue lines represent episodes, thick red lines represent time windows. Thin arrows indicate information used to generate a prediction.

Figure 3: Screenshot of *Solving the Incognitum* where the player can interact with exhibits related to earth sciences.

we report experiments using both time windows and cumulative windows.

Solving the Incognitum

Solving the Incognitum is a game-based interactive learning environment for teaching the relationships between geological time and the fossil record inspired in the historic Charles W. Peale's Museum of Art and Science. The environment and game mechanics have been designed to support different learning and play styles based on Achievement Goal Theory (AGT) (Foster 2011; Elliot and McGregor 2001). In the game, the player can interact with museum exhibits including fossils, minerals, strata deposits, and portraits of renowned historical figures related to the exhibits. A screenshot of the game is shown in Figure 3.

The game is designed to provide different gameplay options. After a brief tutorial, *Solving the Incognitum* provides goal-seekers with four main quests required to win the game. To complete each quest, the player needs to visit a certain set of exhibits, read the information cards associated with them, and apply the knowledge they learn to answer related questions. For explorers, the game contains different types of optional exhibits, grouped based on their types and associations to one another. The players can explore them based on their own interest and answer questions about certain exhibits. Although they do not contribute to winning the game, visiting optional exhibits can earn the player reward badges.

The design of the environment and the placement of the exhibits in the game are intended to highlight different play styles. For example, while explorers may spend time in a specific location of their interest, goal-seekers are more likely to move around tracking the necessary exhibits to complete the main quest. More details of the game can be found in our previous work (Valls-Vargas et al. 2015).

The game tracks and records a gameplay data. For example, we purposefully require the player to hover over exhibits to get related information, the game tracks the player's mouse location as an indicator for what she is paying attention to at the moment. The telemetry data captures movement and interaction variables and can be used to recreate the play session. There are 24 different events recorded during a play session which are converted to 60 features[1].

Experimental Evaluation

In this section we demonstrate the use of our proposed framework for predicting play styles from gameplay data from *Solving the Incognitum*. Here we describe our data collection, and experimental results with different segmentation and prediction approaches.

Data Acquisition

The data used in this paper was collected in a user study conducted on 75 college freshmen. The participants were asked to play the game individually for up to 60 minutes. We intentionally did not tell them they had to complete the game within this period of time. We used video capture software to record their screens while the game recorded telemetry data. From the initial 75 participants, we had to discard 20 data points. The experiments reported in this paper use the remaining 55 gameplay traces.

The gameplay is divided into 6 episodes: A tutorial, the 4 quests required for completing the main objective and, finally, after completion of the main goal, the players are allowed to continue the exploration of the game environment. For our experiments we use only the 4 quest (episodes). Additionally, the experiments are cropped after 20 minutes of gameplay once they complete the tutorial. Figure 4 illustrates the time spent by each player to complete each quest.

After the study, two researchers used the screen capture data to independently annotate play style for each of the 4 quests. The annotations consist of one of the following class labels: *explorer* play style, *goal-seeker* play style, *uninterested* or *other*. The latter two classes emerged from our observations. Our guidelines involve looking for behaviors regarding navigation, use of the quest tracking tools, items visited, order of actions completed and time spent reading and answering questions. For goal-seekers, the annotators looked for a strategy for completing the main goal. For explorers, the annotators looked at how participants interacted with different items based on their own interests neglecting

[1]More information: https://sites.google.com/site/josepvalls/home/taemile

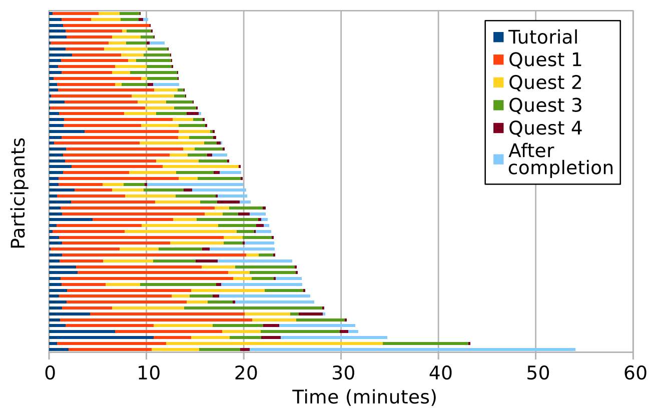

Figure 4: Distribution of the time spent in each of the 6 episodes. Each horizontal bar represents one player. The 55 gameplay traces range from 9 to 54 minutes of gameplay (mean=21.42 min., std=8.01 min.).

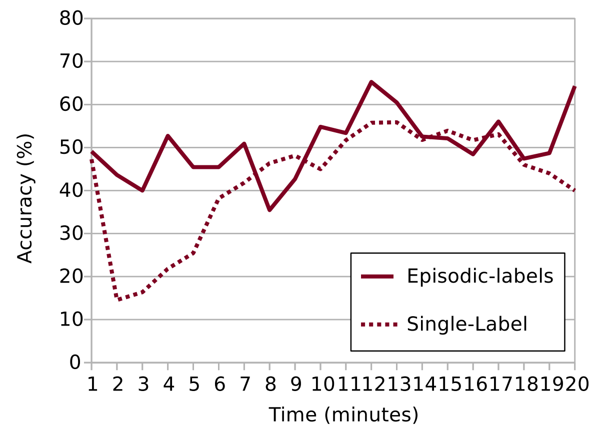

Figure 5: Prediction accuracy over time or Cumulative Prediction models.

the main goals of the game. Participants who were not interested in the game (did not read the material or did not make an effort to answer the questions) were annotated as *uninterested*. An additional *other* label was used for confused players or when the annotators were uncertain.

There are 87 episodes labeled as *goal-seeker*, 65 as *explorer*, 44 as *uninterested* and 24 as *other*. According to the annotations, only 14 people exhibited a consistent play style (i.e., had the same play style throughout the four quests), the rest of the participants all made at least one shift. Figure 1 illustrates the observed play style shifts. For instance, participants labeled as *goal-seekers* rarely switched to another play style. Notice that these trends could be captured by our proposed sequential prediction approach.

Results

Following the approach proposed by Mahlmann et al. (2010) we compared algorithms from the different families available in the WEKA machine learning environment (Hall et al. 2009). Although very simple, the algorithm that exhibited the best performance on our dataset was OneR (a method that employs only a single feature for prediction). It outperformed more complex algorithms: J48 decision trees, SVM, BayesNet and IBk (similar results have already been reported (Holte 1993)). All the experiments use a 5-fold cross-validation strategy and report the average accuracy in play style prediction. All results indicated as significant were tested for significance using the McNemar's test (McNemar 1947) on label predictions with $p < 0.01$.

Episodic Trace Labeling. Our first experiment was designed to validate the hypothesis that labeling play style episode by episode provides better information than just labelling whole traces with a single, static, play style. For this purpose, we compared the *Cumulative Prediction (Single-Label)* and *Cumulative Prediction (Episodic-Labels)* models described before.

As Figure 5 shows, the Episodic-Labels model significantly outperformed the Single-Label model at the beginning and end of the game, and performed similarly in the middle stages. In average, the prediction accuracies achieved by the models are 42.36% for Single-Label, and 55.45% for Episodic-Labels, resulting in a significant difference, which supports our hypothesis.

Episodic Prediction. We tested the episodic prediction model described before with several segmentation strategies: episodic segmentation (based on quests), time window segmentation (using 1 minute, 2 minute and 5 minute granularities), and cumulative windows (using 1 minute granularity). Notice that episodic prediction with cumulative segmentation is equivalent to the *Cumulative Prediction (Episodic-Labels)* model described before.

Figure 6 shows the prediction accuracy achieved over time for each segmentation strategy. Considering episodic granularity, the accuracies for predicting play style for each of the episodes were 52.73%, 39.99%, 60.00% and 65.45% (in the figure, we can see that the first episode is the longest, and the last episode only appears in minute 20). Two drawbacks of this approach are that predictions are only computed at the end of each episode (no intermediate prediction) and that the length of the episode is not know a priori.

The cumulative approach exhibits significant improvement on the classification accuracy (50.45% overall) over the 1 minute time windows (42.76%) and slightly underperforms when compared to the 2 minute window (51.31%) and 5 minute windows (52.49%). These experiments show that episodic prediction does not significantly improve performance with respect to Cumulative Prediction (Episodic-Labels), and when time windows are small (e.g. 1 minute) performance is lower.

Sequential Prediction. Finally, we evaluated the performance of the sequential prediction approach. As Figure 7 shows, results are significantly better than both cumulative models and episodic models. The average accuracy for 1 minute time window is 59.65% (compared to 42.76% for the episodic model), 61.06% for the 2 minute window (up from 51.31% for the episodic model), and 58.97 for the 5 minute window (up from 52.49% for the episodic model).

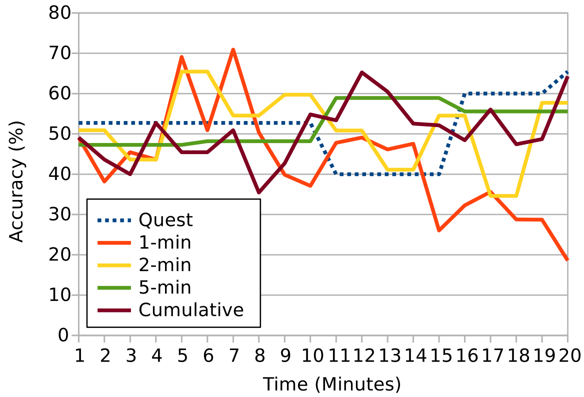

Figure 6: Prediction accuracy over time for the episodic prediction approach with different segmentation strategies. *Quest* is an episodic segmentation using the quests required to complete the game, 1, 2 and 5 minute time windows, and, cumulative windows sampling every minute.

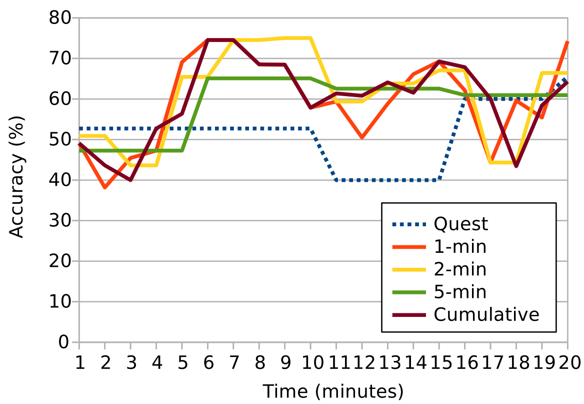

Figure 7: Prediction accuracy for the sequential approach with different segmentation strategies (time windows and cumulative). *Quest* is shown only for reference, to compare results with those from the episodic models from Figure 6.

The accuracy with cumulative windows was 59.85% (compared to 50.45% for the episodic model). Notice that these differences are significant, and show evidence that taking into account both the fact that player style changes over time, and the fact that the play style of a player in a given episode depends on her play style in the previous episodes are relevant contributions in our approach.

Discussion

Our results confirm that to model and predict dynamic play styles, performing an informed segmentation of a gameplay trace is necessary. By using an episodic segmentation for assigning play style annotations and for learning different models over time, we were able to better capture behavioral trends over the course of the play session.

The results also shows that our sequential prediction approach, which feeds forward information from episodic predictions as prior information for subsequent episodic or time interval predictions, substantially improves the accuracy of predictions for play style. The performance improvement can be attributed to the fact that feeding forward episodic information to subsequent episodes encodes conditional dependencies in play style tendencies (as seen in Figure 1).

The plots in the previous section illustrate two interesting phenomena related to our methodology and our dataset. We observed poor performance during the first few minutes of game play when there is not enough feature information in the time windows nor prior play style information. Thus, the predictive models basically cast predictions biased toward the explorer play style (exhibited by 54.54% of the players in the first episode). After the first few minutes, as more information becomes available, the prediction accuracy increases significantly and then, once predictions for previous episodes become available, the accuracy performance is sustained over 60%. The second phenomenon observed toward the end of the play session (after minute 15) is a drop in performance and unstable behavior, specially in the time models that do not use prior information. This is attributed to the fact that, as participants finish the game, less information is available for training and evaluating the models (39/55 traces at minute 15, 25/55 traces at minute 20).

Conclusions and Future Work

This paper presented an approach to dynamic play style prediction based on episodic segmentation of gameplay traces and sequential machine learning. Our approach is based on existing evidence showing that players shift play styles within a play session. This dynamic nature of play style undermines player modeling approaches that assume a static play style. The proposed sequential machine learning approach trains multiple models that include play style predictions from previous time intervals in order to consider how players change play style over time. We compared our proposed approach to a collection of other approaches assuming static play style or assuming changing play style but without taking into account previous play style predictions.

The results of our experimental evaluation show that fine-grained time windows do not provide enough information for casting meaningful predictions, and that casting predictions at the whole-trace level does not properly capture the dynamic nature of play style. Thus, we proposed an intermediate granularity episodic segmentation approach, which provides a good balance and results in better performance. Also, our experiments show that a sequential machine learning approach outperforms non-sequential, techniques.

As part of our future work we plan on incorporating this player modeling approach into an experience manager. We would like to continue our research on episodic segmentation and explore methodologies that automatically determine optimal segmentation. We are interested in issues related to change point detection that enable the identification of the exact moment when play style change happens.

Acknowledgement

This material is based upon work supported by the National Science Foundation under Grant Number IIS-1349082. Special thanks to Aroutis N. Foster, Glen Muschio, Justin H. Patterson, Daniel J. Newman, and Andrew G. Kahl whose outstanding efforts made this project possible.

References

Bifet, A., and Gavalda, R. 2007. Learning from time-changing data with adaptive windowing. In *SDM*, volume 7, 2007. SIAM.

Canossa, A. 2013. Meaning in gameplay: Filtering variables, defining metrics, extracting features and creating models for gameplay analysis. In *Game Analytics*. Springer. 255–283.

Drachen, A.; Canossa, A.; and Yannakakis, G. N. 2009. Player modeling using self-organization in tomb raider: Underworld. In *Proceedings of the IEEE Conference on Computational Intelligence and Games*, 1–8. IEEE.

Elliot, A. J., and McGregor, H. A. 2001. A 2× 2 achievement goal framework. *Journal of personality and social psychology* 80(3):501.

Foster, A. N. 2011. The process of learning in a simulation strategy game: Disciplinary knowledge construction. *Journal of Educational Computing Research* 45(1):1–27.

Gow, J.; Baumgarten, R.; Cairns, P.; Colton, S.; and Miller, P. 2012. Unsupervised modeling of player style with lda. *IEEE Transactions on Computational Intelligence and AI in Games* 4(3):152–166.

Hall, M.; Frank, E.; Holmes, G.; Pfahringer, B.; Reutemann, P.; and Witten, I. H. 2009. The weka data mining software: an update. *ACM SIGKDD Explorations Newsletter* 11(1):10–18.

Heeter, C. 2009. Play styles and learning. In *Handbook of Research on Effective Electronic Gaming in Education*. IGI Global. 826–846.

Holte, R. C. 1993. Very simple classification rules perform well on most commonly used datasets. *Machine learning* 11(1):63–90.

Machado, M. C.; Fantini, E. P.; and Chaimowicz, L. 2011. Player modeling: Towards a common taxonomy. In *Proceedings of the 16th International Conference on Computer Games*, 50–57. IEEE.

Magerko, B.; Heeter, C.; and Medler, B. 2010. Different strokes for different folks: Tapping into the hidden. *Gaming and Cognition: Theories and Practice from the Learning Sciences: Theories and Practice from the Learning Sciences* 255.

Mahlmann, T.; Drachen, A.; Togelius, J.; Canossa, A.; and Yannakakis, G. N. 2010. Predicting player behavior in tomb raider: Underworld. In *Proceedings of the IEEE Conference on Computational Intelligence and Games*, 178–185. IEEE.

Mateas, M., and Stern, A. 2005. Procedural authorship: A case-study of the interactive drama façade. *Digital Arts and Culture (DAC)*.

McNemar, Q. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika* 12(2):153–157.

Pedersen, C.; Togelius, J.; and Yannakakis, G. N. 2009. Modeling player experience in super mario bros. In *Proceedings of the IEEE Conference on Computational Intelligence and Games*, 132–139. IEEE.

Riedl, M. O.; Stern, A.; Dini, D.; and Alderman, J. 2008. Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning* 4(2):23–42.

Ross, S., and Bagnell, D. 2010. Efficient reductions for imitation learning. In *International Conference on Artificial Intelligence and Statistics*, 661–668.

Sharma, M.; Ontañón, S.; Mehta, M.; and Ram, A. 2010. Drama Management and Player Modeling for Interactive Fiction Games. *Computational Intelligence* 26(2):183–211.

Smith, A. M.; Lewis, C.; Hullet, K.; and Sullivan, A. 2011. An inclusive view of player modeling. In *Proceedings of the 6th International Conference on Foundations of Digital Games*. ACM Press.

Thue, D.; Bulitko, V.; Spetch, M.; and Wasylishen, E. 2007. Interactive Storytelling: A Player Modelling Approach. In *Proceedings of the 3rd Artificial Intelligence and Interactive Digital Entertainment Conference*, number July, 43–48.

Togelius, J.; De Nardi, R.; and Lucas, S. M. 2006. Making racing fun through player modeling and track evolution. In *Proceedings of the SAB Workshop on on Adaptive Approaches to Optimizing Player Satisfaction*.

Tychsen, A., and Canossa, A. 2008. Defining personas in games using metrics. In *Proceedings of the 2008 Conference on Future Play: Research, Play, Share*, 73–80. ACM.

Valls-Vargas, J.; Kahl, A.; Patterson, J.; Muschio, G.; Foster, A.; and Zhu, J. 2015. Designing and tracking play styles in solving the incognitum. In *Proceedings of the Games + Learning + Society Conference*, to appear. Games + Learning + Society.

Van Der Werf, E.; Uiterwijk, J. W.; Postma, E.; and Van Den Herik, J. 2003. Local move prediction in go. In *Computers and Games*. Springer. 393–412.

Weber, B. G., and Mateas, M. 2009. A data mining approach to strategy prediction. In *Proceedings of the IEEE Conference on Computational Intelligence and Games*, 140–147.

Yannakakis, G. N., and Maragoudakis, M. 2005. Player modeling impact on player's entertainment in computer games. In *User Modeling*. Springer. 74–78.