

# Reading Between the Lines – Towards an Algorithm Exploiting In-game Behaviors to Learn Preferences in Gameful Systems

Enrica Loria  
Fondazione Bruno Kessler  
University of Trento  
Trento, Italy  
eloria@fbk.eu

Annapaola Marconi  
Fondazione Bruno Kessler  
Trento, Italy  
marconi@fbk.eu

## ABSTRACT

Players' retainment can be fostered by investigating whether the game elements players are interacting with are to their liking and tailoring game dynamics to meet their preferences. Thus, adaptive gameplay is a widely interesting topic in both the Game User Research field and the game industry. Considering that explicit information on players' preferences often lacks, alternative approaches are needed. This task becomes even more challenging when the gameplay data available is limited due to the simplicity of the system employed, as it occurs in gameful systems in contrast to complex entertainment games or serious games. In this work, we propose an algorithm that exploits user behaviors as an implicit component to compute players' preferences by measuring their level of activity. The application domain is a persuasive gameful system, and the customizable game elements are single-player challenges. The proposed algorithm uses offline gameplay data to compute a preference score for every viable option. The outcomes are then compared against a ground truth calculated from players' in-game choices. Our findings suggest that players' behaviors can be used to inform the generation of tailored game elements.

## CCS CONCEPTS

• **Human-centered computing** → **User studies; Field studies; Empirical studies in HCI**; • **Computing methodologies** → **Model development and analysis**;

## KEYWORDS

player behaviors, player experience, adaptive gameplay, game analytics, persuasive games

### ACM Reference Format:

Enrica Loria and Annapaola Marconi. 2020. Reading Between the Lines – Towards an Algorithm Exploiting In-game Behaviors to Learn Preferences in Gameful Systems. In *International Conference on the Foundations of Digital Games (FDG '20)*, September 15–18, 2020, Bugibba, Malta. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3402942.3403016>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

FDG '20, September 15–18, 2020, Bugibba, Malta

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8807-8/20/09...\$15.00

<https://doi.org/10.1145/3402942.3403016>

## 1 INTRODUCTION

Games and game-like experiences seek to engage the user in (semi-) virtual worlds, which are an augmentation of reality. However, perception is highly subjective, and what may immerse and engage some users could also completely discourage or bother others [4, 20, 59] - such heterogeneity claims for customized and adaptable solutions.

The research conducted around player modeling techniques [4, 20, 43, 60], information extraction from telemetry data [7, 17, 63], and procedural generation of game content [23, 67, 68] is growing bigger by the second. The vast majority of studies have been conducted on entertainment games, which, thanks to the high amount of in-game interactions allowed, produce rich and varied datasets. The features extracted from those datasets can be used to gain insights on players' experience. Examples of applications range between estimating the likelihood of churn [12, 22, 58] to building customized strategies [65].

Studying behavioral patterns can bring to a further understanding of the motivations that lead a user to play [24]. Since gameful systems aim at making every-day actions and tasks enjoyable [15], deploying customizable content is an extremely relevant topic [32, 48]. This also applies to the field of gamification as a Persuasive Technology. These systems foster a positive behavioral change and keeping the player engaged for longer may result in a more permanent and consistent change [5, 25, 28, 61]. Thus, considering players' traits and preferences becomes particularly relevant [46]. Players' retention is also likely to be improved by adapting the experience to the player preferences [40]. Such an adaptation can be implemented by letting players express their preferences, e.g., by evaluating the content through an embedded rating system. However, only a part of the players enjoy personalizing their experiences actively [4, 20, 40, 43]. Others tend to ignore those mechanics, but still, expect the gameplay to be engaging to their tastes. For the latter class of players, there is a need to transparently understand how to adapt the system by relying on other sources.

In gameful systems, the game mechanics implemented are usually simpler than in *traditional* games. Besides, the features that can be analyzed to learn players' preferences are much more limited. This limitation is even more problematic for the second class of players, which lack in expressing their preferences explicitly.

Although in-game choices can be used to deduce player preferences [38], it requires players to have explicitly manifested preferences through in-game interactions. With our work, we want to move a step forward and understand whether such preferences can be inferred from implicit in-game behaviors. We measured players' in-game behaviors by analyzing players' levels of activity in the

game. Then, we used those values to deduce the suitability of the personalized game elements for each player. Using those values throughout the gameplay, we computed preference scores for each viable game element. Then, we used players in-game choices [38] as ground truth to evaluate the performance of our algorithm. We conducted a feasibility study, in which we explored the idea of iteratively using players' activity level as feedback for the generated game element in an algorithm inspired by the Reinforcement Learning (RL) paradigm. The goal was to learn a generation strategy tailored towards each player. Towards this, we exploited implicit information, and validated it with explicit information: we used in-game choices to assess the reliability of using in-game activity behaviors as feedback. Although the algorithm works with offline gameplay data, the aim is to set the ground for introducing it as a content generator integrated into the system. Thus, we aim at transforming an RL-inspired approach into an actual RL algorithm.

The results show that indeed in-game behaviors are informative. Moreover, we found that our approach works better on a specific subclass of players. Although the number of players eligible for the study was modest, we believe that the outcomes show a promising methodological direction. Other researchers could be inspired to investigate the topic further and to move towards a real adaptable experience.

The remainder of the paper is structured as follows. We present some relevant studies in the literature in Section 2. In Section 3, we describe our use case scenario, which is the source of the data analyzed. Then, we present the analysis settings, the algorithm, and the evaluation process. The Discussion Section concludes the paper.

## 2 RELATED WORKS

In information science, decisions in marketing and product design have been widely guided by customers' information. Profiling customers is a common practice in any user-centered systems. Therefore, this certainly applies to games and gameful environments. The immersive nature of games allows easy extraction of hundreds of features, by inspecting user-game interactions. Cluster analysis can be used to reduce the dimensionality in behavioral data [44], to eliminate noisy and redundant entries. This information can also be integrated with data coming from play-testing, social networks, and marketing.

Game Analytics gathers the emerging techniques aimed at defining and finding patterns in the behaviors of the player [19]. Scientists and designers have the opportunity to enhance player experience by exploiting those methodologies tracking users' experiences in real-time [31] and improving them [55]. Research swarms of data-driven methods employing player telemetry data to improve gameplay. Among the most prolific analyses, there are techniques to dynamically adjust in-game tasks' difficulty, to keep players in a state of flow [9]. For instance, within Dynamic Difficulty Adjustment, Artificially Intelligent (AI) agents ensure that the level of difficulty is tailored to the players' abilities and that it is adjusted as the player increases her experience [14]. Data-driven approaches are also used to inform more specific strategies. For instance, players' behaviors are used to decide the type of hints more suitable for each player [65], to intervene when the gameplay is thwarted by

other external psychological problems [51], and to provide crucial indicators for learning [30]. The analysis of gameplay data also aims at better understanding players. Examples are studies measuring self-esteem through in-game behaviors [49], and using this data as rapid feedback while designing game content [8]. Other works investigated how players' performance is affected by the network of players [63], and how to cluster players according to their behavioral profiles [17]. Sifa et al. suggested means for building behavioral profiles to interpret the player experience [2]. Previous works looked at creating behavioral profiles via clustering [17]. Others studies dealt with self-organizing networks, to inspect how profiles mutate with progress in the game [54].

Another class of problems catching a lot of interest is the prediction of player churn [22, 45]. Researchers are also active in investigating techniques to foster players' retainment. Player retention in games is a core concern in entertainment games. Therefore, player churn is also relevant across the game industry. This resulted in a consistent amount of concrete approaches to address this problem. Those methods range from the exploitation of specific behavioral characteristics to assess the likelihood of a player churn, to the investigation of the effect of different win/loss ratios [13]. Several of those papers originate in Game AI [66] and business intelligence requirements [18]. Examples of prediction techniques used are Hidden Markov models [58] and survival analysis via Mixed Effects Cox Regression [12]. Others have looked at the problem as a marketing model of customer retention [11] by using metrics such as playtime, stop rate, and the number of years players have been playing the game. Many of such studies provide accurate predictions on both whether and when each player would leave the game [6]. However, successfully predicting players' abandonment is only the first stage of the problem. Prompt strategies to prevent such abandonment are needed. Although this is partially achieved through difficulty-adjustment algorithms, it is equally important to tailor game content also according to players' preferences.

In the context of adaptable game content, in-game recommender systems remain relatively unexplored. In one of the few works available, Joshi et al. [27] built a team-based recommender system for player-versus-player (PvP) content in MMOGs, which merges clustered behaviors and recommendation systems. The idea is to combine both knowledge on individual players and teams. A recommender system for individual players in *Destiny* was studied by Sifa et al. [53], by considering the learning rate or the rate of skill acquisition of players and/or teams. Other works focused on giving players a visual representation of how they have been profiled to motivate the suggestions of game elements [34]. However, all those studies see their application domain limited to the entertainment commercial game titles [16], with a particular focus on free-to-play and MMORPG games.

### 2.1 Adaptive Gameful Systems

Gamification, defined as the use of game elements in non-game contexts [15], has been used for about ten years in educational settings to increase learner performance, motivation, or engagement [1, 10], and to make users intrinsically motivated in performing a certain task, i.e., crowdsourcing [42]. Recent studies conducted on the effects of gamification show that to be effective in the domain

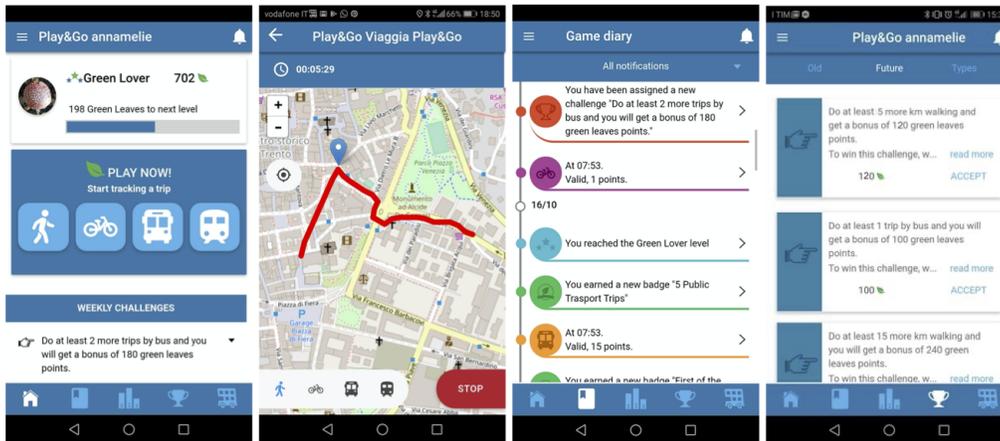


Figure 1: Screenshots. Basic interactions in Play&Go.

of health [47], sport [36] and learning [10, 29, 35], gamification should be tailored to the users. Adaptive gamification in education is a novel and cutting edge research field, that has been gaining in popularity in the past few years. Such emergent field benefits of a theoretical base build upon several studies conducted in real-world learning settings. In static adaptation approaches profiles are identified, often relying on theoretical models (e.g., the Bartle Player types [4], the Brainhex player satisfaction model [43], the Hexad player types [60], or the categories of players described by Ferro et al. [20]). Users are sorted into different categories based on these profiles; different game elements are given to each of the different categories of players [56]. In dynamic adaptation, the deployment of the game elements are informed by players' activity in the game. A user profile is built by monitoring her in-game behaviors. In the work of Paiva et al. [48], users' actions are divided into either collaborative, gamification, individual, or social interactions. The adapted element is the player's goal, which is assigned according to the type of actions performed. Another example is Jagust et al.'s [26] study, in which two dynamic adaptation scenarios are presented. In the first situation, the players' task is timed. When the task is completed, players are given less time for the following one. In the second situation, the difficulty is defined by the completion time, instead of the target score. The score increases as they carry out the task correctly. In another work, the customization is in the types of badges presented to and in the feedback given [29]. Other studies use a hybrid approach in which they define static adaptation rules for given player profiles. When the profile of a user changes, different game elements are given to her [41]. The Brainhex model inspires the profiles. In the systems suggested by Knutas et al. [32], an algorithm uses players' profiles and interactions. In this case, the authors use the Hexad player profile and users' skills. They recorded students during their project meetings to classify their interactions. Then, they proposed the game elements according to those profiles and interactions. However, in this work, a method to detect these actions at real-time lacks [32].

All in all, the field of adaptive gamification is still young and growing, especially for what concerns dynamic adaptation [56]. Besides, gamification research still misses coherence in research models, as well as theoretical foundations [33]. Moreover, the motivational affordances implemented in the majority of the studies remain 'goals and objectives', 'multimedia feedback', and 'metaphorical/fictional representations' [64], neglecting others such as the expression of self-identity [69]. Since different things motivate different people, it can be expected that personalizing the incentives and the way the rewards are presented to the individual, would be beneficial [61]. In particular, tailoring at the level of social influence strategies may increase the effects of the persuasive technologies [28]. Gamification studies would benefit from wider use of theories to account for the complexity of human behavior, and a more thorough exploration of the many opportunities coming from the world of games [52].

This work contributes to making a small step towards adaptive gamification in that the approach proposed can be used to learn players' preferences from gameplay data. Besides, it can be extended to inform an online procedural content generation algorithm that exploits the learned players' preferences.

### 3 USE CASE SCENARIO

*Play&Go* (Figure 1) is a mobile gameful system whose aim is to incentivize citizens to assume sustainable mobility behavior. The game is active in the city of Trento (Italy), and, with the sponsorship of the local municipalities, a 6-month edition is conducted every year. In this work, we considered gameplay data retrieved from the 4th edition (October 2018 - April 2019).

The game allows players to track their sustainable movement by specifying the transportation mean employed (walk, bike, bus, or train). An automatic system validates the transportation mode declared for the trip. It analyzes acceleration and speed of the player in the duration of the journey, and by checking whether the trip was feasible - e.g., checking the bus route and timetable. For each

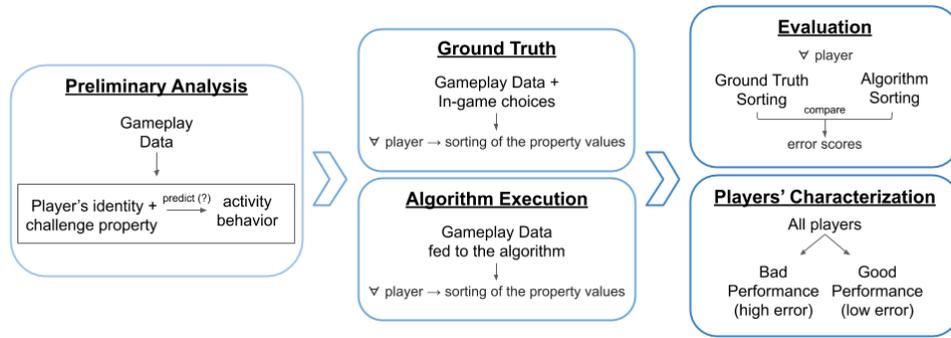


Figure 2: A visual representation of the analysis workflow.

validated trip, Green Leaves points are awarded based on the length of the journey and the level of sustainability of the mean used.

In the game, particular attention is paid to keep users in a state of flow [9]. To this extent, players are motivated to improve (or maintain) their performance through weekly challenges. Challenges are generated by a Recommender System (RS) to be tailored to each user, considering player history, preferences, and skills. The rewards for such challenges are computed in function of the estimated challenge difficulty. The difficulty refers to the effort required for the player to win it. An example of an individual challenge is the following: "Do at least 3 kilometers by bike to win 120 Green Leaves points." For the more inexperienced players, the weekly challenge is automatically assigned. When players reach Level 2, they have a small window of time (a couple of days), at the beginning of every week, when they can program the challenge for the following week. Programming a challenge means selecting the challenge they want to participate in. They can choose among 2 (after Level 2) or 3 options (after Level 3). In case of missing choice, the system assigned to the player the challenge more aligned to the player's history. Figure 1c shows an example of the UI presented to the users to program a challenge. The options available varied in the transportation mean upon which the target is evaluated and in the effort required for the player, estimated according to their previous performance. *Play&Go* has obtained a growing positive response from *Trento's* citizens edition after edition. In our last edition of the game, we observed a more committed audience of players. We counted 834 registered players (590 performed at least one game action). In total, players tracked about 64k trip and about 250k kilometers. Of those trips, the vast majority were tracked by foot, followed by train, bus, and bike. Therefore, transportation means we call "zero impact" were the favored ones. Interviews and surveys showed that players were particularly motivated by the challenges [21]. Since programming a challenge is optional, about 147 players tried to program at least one challenge. Players involved in at least 10 comparisons were 94. Being involved in 10 comparisons means to have made between 5 to 10 choices. Every time that a choice is made 1 or 2 comparisons occur: the selected challenge against the other(s) option(s) available.

In terms of demographics, we observed that our players are quite heterogeneous in terms of gender (57% female), which is probably due to the variety of game elements and mechanics implemented.

As for the distribution of the age ranges: 3% of our users are younger than 20 years old, 40% are between 20 and 35 years old, 25% are between 35 and 50 years old, 19% are between 50 and 70 years old and 2% are older than 70. This age range distribution is in-line with similar initiatives. All age ranges are well represented, exception made for elderly users. A low participation rate can be related to their resistance to smartphone applications. Even when owning a smartphone, older adults mainly exploit basic functionalities (e.g., calls and messages).

## 4 ANALYSIS WORKFLOW

In this work, we have investigated whether in-game behaviors, more specifically the evolution and modification of such behaviors, can be used to inform an algorithm learning how to customize the available game elements. In the following section, we will describe the analysis setting used to answer our research questions.

- RQ1. Can we learn players' preferences by studying in-game behaviors measuring their level of activity?
- RQ2. Does this approach work better for a subset of players? If so, how are those players characterized?

We formalized the concept of customization as follows. The game is built of game elements, which represents its low-level implementations [56], examples beings challenges, quests, or other simple game actions. Some of those game elements, which are the ones we are interested in, can be customized. Such customization can take place by modifying one, or more, aspects of the game elements (or properties). The properties can assume a value in a set (*PARAM*) a priori defined, which is finite and reasonably small.

In our use case scenario, the game elements to be personalized are individual challenges and the property that can be customized is the counter the challenge is evaluated upon (walk kilometers, bike kilometers, bus trips, train trips, green leaves). In the remainder of the paper, we will use the term counters, properties, and challenge types interchangeably.

In the following sections, we present the algorithm employed to compute the preferences' scores and its evaluation. As shown in Figure 2, in a **preliminary analysis**, we investigated whether players' activity level was influenced by different types of challenges. Then, we built our **ground truth** by calculating player preferences

expressed through their explicit in-game choices. We run the **algorithm** on each of our eligible players and **evaluated** the outcomes against the ground truth. Finally, we **characterized** a group of players for which the algorithm was particularly accurate.

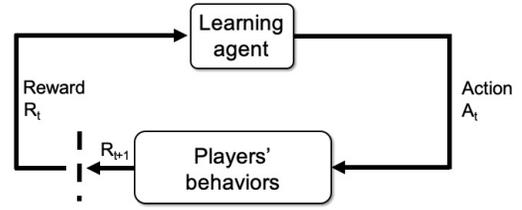
To achieve this goal we instantiated the active and committed behaviors proposed by Loria and Marconi [37], measuring the intensity and the constancy of players' participation. Those values have been used to evaluate the impact of the custom generated game elements on players. The outcome will be a sorting of the challenge types available, from the most preferred to the least appreciated. Then, we compare this sorting with our ground truth, produced by analyzing the preferences - indirectly - expressed by the players when choosing the game elements in the programming phase of the game.

We conducted a temporal analysis in contrast to an aggregated analysis to assess and learn player preferences for two reasons. The main motivation is that, although the current algorithm analyzes offline data and does not take part in the process of generating customized content, the final goal is to produce an online algorithm that, iteration after iteration, learns the best game element for each user. The second motivation is that behaviors, as preferences, can vary over time. Research on player experience has repeatedly highlighted the importance of studying the temporal aspects of game-play [31] through sequential pattern mining techniques [62]. Likewise, this view is supported by behavioral scientists arguing that, to understand the dynamic process, a sequential view is essential [3]. The system is already set to be adaptable in case players progressively shift from a preferred game element to another.

We structured our gameplay data in timeframes of 1 week each. The rationale lies in the definition of the game itself and in the game elements that we are interested in analyzing. The challenges, which represent the customizable items, have a validity of one week; and every week players can submit their choice - if they want - for the challenge that will be activated in the following week. Timeframes, however, are not required to be as strict as ours. In a gameful system where the customizable game elements are not constrained to a precise period, timeframes can also be adaptable. For example, if the customizable element is a quest, which is active for the time the player needs to complete it, the concept of timeframe could be linked to the life of the quest itself. From now on, the terms timeframes and weeks will be used interchangeably. However, it should be remarked that employing them as synonyms is acceptable only in this context, since the concept of timeframe can be adapted to every game and need.

#### 4.1 Offline Gradient Bandit

The intuition behind reinforcement learning algorithms has inspired the algorithm we designed to compute the players' preference score. In Reinforcement Learning (RL), the agent learns from experience through trial-and-error as it engages with the task. The goal is to maximize a long-term reward [57]. We saw a similarity between learning from experience and using in-game behaviors throughout the gameplay to iteratively learn player preferences. In our case, the experience is given by the level of activity fostered by the choice made (i.e., the challenge selected). We analyzed in-game behavior measuring players' activity levels at precise timeframes, where the timeframes were built to contain one customized game



**Figure 3: Algorithm Model.** The action that can be taken by the agent  $A_t$  represents the choice for the customized challenge, which differs on the counter the challenge is evaluated upon (green leaves points, walk kilometer, bike kilometers, bus trips and train trips). The reward  $R_{t+1}$  represents the player's activity level when the challenge at time  $t$  was active.

---

#### Procedure The Offline Gradient Bandit Algorithm

---

**Input:**  $B, R, k, n, \alpha$

```

1:  $H(a) \leftarrow 0, \forall 1 \leq a \leq k \in \mathbb{N}$ 
2:  $H(a) \leftarrow 1/k, \forall 1 \leq a \leq k \in \mathbb{N}$ 
3:  $avg\_reward \leftarrow 0$ 
4:  $i \leftarrow 1$ 
5: while  $i \leq n$  do
6:    $A_t = B_i$ 
7:    $r = R_i$ 
8:    $avg\_reward = avg\_reward + (1/i)(r - avg\_reward)$ 
9:    $H(A_t) = H(A_t) + \alpha(r - avg\_reward)(1 - Pr(A_t))$ 
10:  for all  $a \neq A_t, 1 \leq a \leq k \in \mathbb{N}$  do
11:     $H(a) = H(a) + \alpha(r - avg\_reward)Pr(a)$ 
12:  end for
13:  for all  $1 \leq a \leq k \in \mathbb{N}$  do
14:     $Pr(a) = exp(H_t(a))/sum(exp(H))$ 
15:  end for
16: end while
17: return  $H$ 
  
```

---

element at a time. Considering that the challenge is the primary driver of players' engagement [39], our intuition was that the activity level is an indicator of how much that game element had been appreciated. In other words, the more active the player was, the more the game element was of her liking. Thus, the gameplay was divided into timeframes and the reward was the activity level, with the underlying long-term reward being increasing retention and avoiding churn (Figure 3).

This interaction protocol goes hand in hand with RL algorithms, in which the data observed up to time  $t$  is considered to decide which action to take at time  $t + 1$ . The information is used to **evaluates** the actions performed, instead of providing the correct answer.

The simplest group of problems that can be tackled with RL methods are known as the *bandit problems*, which belong to the class of tabular solution methods. As a quick refresher on bandit problems in general, we have one or more "bandit" which is often represented as a slot machine. Pulling an arm results in a different reward depending on which one has been selected. The goal is to maximize our reward in the long run, by determining what the best action to take is. The faster we can find the optimal choice, the higher

the gain. The expected reward we get from each action is denoted by a Q-value, which is learned through interacting with the bandit. Those interactions increase the experience of the agent, which updates the Q-values after every interaction. The assumption is that every machine has its - unknown - distribution, from which it picks its outcome. A variant of the multi-armed bandit problem is called the gradient bandit. Rather than finding the optimal choice, the goal is to estimate a preference score for every action (bandit). Gradient algorithms are used to measure the value of an action relative to the other possible actions. In these algorithms, we learn a preference score  $H_t(a)$ , used in the selection of the action to take. The action with the highest preference score is chosen. Initially,  $H_0(a) = 0, \forall a$ , and the probability of choice is equal for every available action. In contrast to the traditional bandit algorithms, gradient algorithms update the  $H_t(a)$  value for each action using a stochastic gradient ascent, instead of updating the average of the rewards. Every time that an action is taken, the new value for  $H_t(A_t)$  is computed by weighting the reward returned by the probability of the action and the learning rate.

A widely employed formula (Formula 1) to update the preference scores is based on the idea of stochastic gradient descent (SDG).

$$\begin{aligned} H_{t+1}(A_t) &= H_t(A_t) + \alpha(R_t - \bar{R}_t)(1 - \pi_t(A_t)), \text{ and} \\ H_{t+1}(a) &= H_t(a) + \alpha(R_t - \bar{R}_t)\pi_t(a), \forall a \neq A_t \end{aligned} \quad (1)$$

$\bar{R}_t \in \mathbb{R}$  is the baseline and is the average of all the rewards up through and including  $t$ . If the actual reward is lower than the baseline, in the following iteration the probability of choosing that action decreases; increases otherwise.  $\alpha > 0$  is the learning step-size parameter.

The larger the preference, the more often that action is taken. The probability of taking the action  $a$  at time  $t$  follows the soft-max distribution.

$$\pi_t(a) = Pr\{A_t = a\} = \frac{e^{H_t(a)}}{\sum_{b=1}^n e^{H_t(b)}} \quad (2)$$

In this set of tasks, there is no need to associate different actions in different situations. The learner either tries to find a single best action (if stationary) or the best action as it changes over time (otherwise). The latter sentence shows the origin of the inspiration. In our problem, we had no additional data regarding players. The only information we had was the type of challenge active for them every week. The different types of challenges, here, represents the bandits (e.g., green leaves, walk kilometer, bus trips, etc.). The reward is represented by the level of activity the user had in the week that challenge was active (Figure 3). We run the algorithm for every player and every week. We selected the active challenge and updated the preference scores using the behavior values as a reward. The following intuition motivated this choice. If the challenge met players' preferences, then they would have been more engaged in the game, thus a higher retainment.

We measured the retainment level by monitoring the level of participation of players considering the amount of in-game actions and the length of gameplay during the week. Before conducting the core of the analysis, we validated this hypothesis in a preliminary analysis (see Section 5).

It must be noted that the algorithm is **inspired by** an RL algorithm,

but the problem we tackled is not yet an RL problem itself. We operated with offline gameplay data, while RL problems are online problems. Besides, due to the offline nature of the data, we already knew which challenge was active (the bandit), instead of having to choose it. Our algorithm computes the preference scores knowing, at each iteration, the current bandit (Line 7 in Algorithm ), and the value of the activity behavior (Line 8 in Algorithm ). Such values are used to adjust the preferences for each bandit (Lines 10-13 in Algorithm ). In this phase of the study, we used an offline algorithm to assess the feasibility of using in-game behaviors to learn players' preferences. Therefore, the algorithm does not take part in the choice of the property of the challenge. Rather, it is a silent spectator, which analyzes the player's activity to learn whether such property was suitable for them.

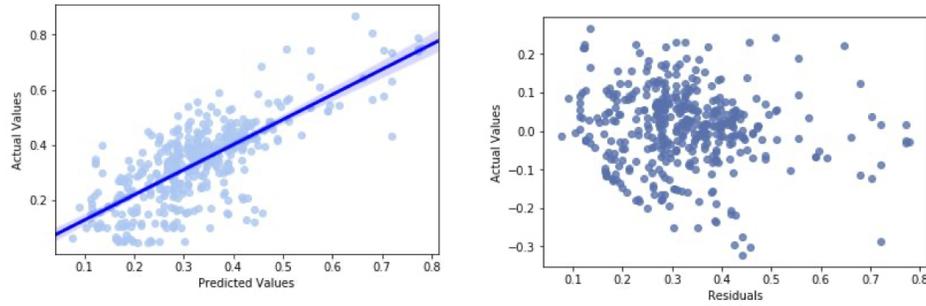
In a future step, to convert the algorithm in an actual RL algorithm Line 7 (Algorithm ) should be changed to select the option according to the probabilities computed upon the preferences scores (Equation 2).

The Algorithm is run for one player at a time, and takes in input the list of active challenges during his or her  $n$  weeks of gameplay ( $B$ ), the values of the activity behaviors for those weeks ( $R$ ), the number of types of challenges they have engaged with ( $k$ ), and the learning step-size ( $\alpha$ ). It is important to note that  $k$  may vary among players, since they may have been presented with a subset of the 5 available types of challenges.

## 5 ALGORITHM EVALUATION

The analysis was conducted following the work-flow hereby presented (Figure 2). In the first instance, we studied whether the challenge type was a predictor of the activity behavior, given a player. Then, we built our ground truth by considering the in-game choices made by players when programming the weekly challenge. As the next step, we run the Algorithm for each player to compute the preference scores, considering all available challenges. An evaluation of the algorithm's outcomes against the ground truth values followed. Finally, we characterized players for whom the algorithm performed particularly well.

**Preliminary Analysis.** Before diving into the evaluation of the core contribution of this work, we conducted a preliminary analysis. We studied whether the types of challenges and the identity of the users were predictors of the level of activity they manifested in each timeframe. We considered a subset of our registered players, who played for at least 5 weeks (timeframes). This filtering resulted in a population of 156 players. We built a dataset using as features the player (identified by the nickname, as a categorical variable) and the challenge property. We, then, predicted the value of the players' behavior within the week in which the challenge with that property was active. The behavior's value (from now on referenced as activity behavior) is computed as the mean of the active and committed behaviors. The *committed behavior* is measured considering the time the player spent in the game over the week, as a value in the  $[0, 1]$  interval. The *active behavior* is measured considering the number of game actions computed in each timeframe, normalized by the maximum amount of actions performed within that timeframe. The normalization is required to bring such values in the same range of the committed behavior values before the



**Figure 4: Linear Regression** The test-set values of the test plotted against the predicted values in (a) and the residuals plotted against the predicted values in (b).

computation of the mean.

We conducted a regression analysis (Figure 4), by building a Multiple Linear Regression model, to verify whether given the player identity and the type of challenge it was possible to predict the level of activity in the game, evaluated through the behaviors mentioned above. Such behavior value represents the impact that the challenge had in terms of repercussions on the player’s participation level. Our features were the players’ nicknames (converted in 155 dummy variables) and the types of counters (converted in 4 dummy variables); resulting in 159 features. We used a dataset of about 2k entries and 156 players, which had participated in between 5 and 23 challenges. The results showed that indeed those features were informative [(F(155,1943) = 161,  $p < 0$ ), with an  $R^2$  of 0.93]. Thus, we proceeded with our experimentation.

**Ground Truth.** In Play&Go, users, after having reached Level 2, can choose among 2 proposed challenges. When they reach Level 3 they can choose among 3 options. Challenges vary in the counter upon which the challenge is evaluated (green leaves points, walk kilometers, bike kilometer, bus trips, and train trips). Once a challenge is accepted, the other options are automatically refused. Thus every choice leads to 1 or 2 comparisons: the selected option over the other(s) suggestion(s). We used this information to build our ground truth for every player. For every parameter, we have the number of accepted and the number of refused challenges (which can be zero if the player never chose that counter) [38]. In this phase we used players that were involved in at least 10 in-game comparisons, resulting in 94 eligible players. We analyzed the choices made during the gameplay to infer player preferences. To do so, we used the paired comparison protocol to compute players’ preference score, an important tool in multi-attribute decision making. As classically done in paired comparison studies [50], we build a squared matrix of dimension  $k$  for every player, with  $k$  the number of properties available. In such a matrix, a value of  $x$  for the cell  $c_{ij}$  means that the counter  $i$  was chosen  $x$  times over the counter  $j$ . To compute the preference scores, we can then consider the number of votes received by each counter, which may then be divided by the number of comparisons per counter for normalization purposes [50]. At the end of the process, for every player, we obtained a vector  $w_c$  of dimension  $k$ , with the preference scores for every counter.

**Algorithm Execution.** The second step was to learn the preference scores of each player by feeding players’ activity data and the chosen challenge type for every week of their gameplay to Algorithm . We considered all available challenges, regardless of whether they were chosen by the player or automatically assigned. For each of the 94 players, the outcome was a vector of preferences  $w_p$  of dimension  $k$ . Each element of the vector refers to one of the counters available for the player, reflecting the order used while calculating the ground-truth values for the same player. Thus, the first element of both  $w_c$  and  $w_p$  referred to the first counter in *counters*, and so on.

Even though the counters available were 5, the dimension of the vectors  $k \in [2, 5]$ . Some players may have interacted with fewer challenges types during their gameplay. This is because the suggestions are produced by exploiting players’ in-game history. Thus, if a player never traced a train trip, it is unlikely that a Train Challenge was suggested. Thus, for each player, we considered only the subset of properties used to produce the challenges’ suggestions. However, the number of counters per player - i.e., the dimension of the vectors - are well distributed ( $M = 4.48$ ,  $SD = 0.63$ ). We will discuss it in more detail in the *Players Characterization* paragraph.

INITIAL SORTING			GROUND TRUTH SORTING			ALGORITHM SORTING		
?	1	Bike Km	0.47	2	Bike Km	-2	5	Bike Km
?	2	Bus Trips	0.6	1	Bus Trips	0.7	2	Bus Trips
?	3	Green Leaves	0.01	5	Green Leaves	1	1	Green Leaves
?	4	Train Trips	0.12	4	Train Trips	-0.1	4	Train Trips
?	5	Walk Km	0.3	3	Walk Km	0.2	3	Walk Km

**Figure 5: Example of properties sorting for a player, resulting from their in-game choices (ground truth) and the algorithm’s outcomes.**

**Evaluation.** The values of the two vectors  $w_c$  and  $w_p$  are not comparable, being computed in completely different ways and belonging to different ranges: the first is in the range  $[0, 1]$  and the second is in the range  $[-\infty, +\infty]$ . However, we are interested in

identifying and validating the options' sortings; in particular, the elements at the extremes (the most and the least liked). To make the sortings comparable, the indexing of both vectors  $w_c$  and  $w_s$  is mapped to the actual counters (e.g., 1 = Bike Km and 2 = Bus Trips, as shown in Figure 5), for every player. In other words, the vector containing the possible values of the challenge property is set a priori and then the sorting for both the ground truth and the algorithm are produced (Figure 5). The sortings  $s_c$  and  $s_p$  are a representation of where the counters are positioned in the ranking for both the ground truth and the learned preference scores respectively.

Figure 5 shows the following example.

Given  $w_c = [0.47, 0.6, 0.01, 0.12, 0.3]$  and  $w_p = [-2, 0.7, 1, -0.1, 0.2]$ , being the vector of the scores in the ground truth and the scores computed by the algorithm respectively, their decreasing sortings are the following:  $s_c = [2, 1, 5, 4, 3]$  and  $s_p = [5, 2, 1, 4, 3]$ . Thus, the actual ranking is:

- *BikeKm* in position 2 for the ground-truth with value 0.47 and position 5 for the algorithm with value -2;
- *BusTrips* in position 1 for the ground-truth with value 0.6 and position 2 for the algorithm with value 0.7;
- *GreenLeaves* in position 5 for the ground-truth with value 0.01 and in position 1 for the algorithm with value 1;
- *TrainTrips* in position 4 for the ground-truth with value 0.12 and position 4 for the algorithm with value -0.1;
- *WalkKm* in position 3 for the ground-truth with value 0.3 and position 3 for the algorithm with value 0.2.

We compared the sorting obtained from the programmed challenges (ground truth) and the one computed from the preference scores produced by the Offline Gradient Bandit Algorithm. To assess the similarity of the two sortings, and hence the performance of the designed algorithm, we used the following measures. As an error measure, we used *NRMSE* (Equation3), which is easier to interpret respecting to *RMSE* in that it gives a relative percentage on the error obtained. To assess the similarity between the two (non-zero) vectors we used cosine similarity. Mathematically, it measures the cosine of the angle between two vectors projected in a multi-dimensional space. The smaller the angle, the higher the cosine similarity. Finally, we computed the  $R^2$ , which is a statistical measure that represents the proportion of the variance for a dependent variable that's explained by an independent variable or variables in a regression model. The  $R^2$  is the square of the correlation. It measures the proportion of variation in the dependent variable that can be attributed to the independent variable.

$$NRMSE = \frac{\sqrt{\frac{1}{k} \sum_i^k (x_i - \bar{x}_i)^2}}{x_{max} - x_{min}} \quad (3)$$

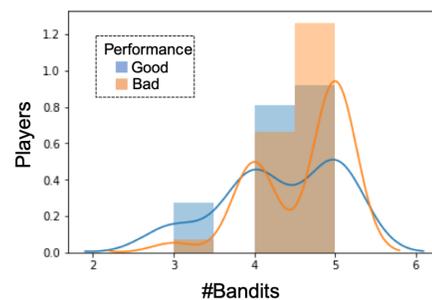
For how it is defined, *NRMSE* further penalizes very different sortings, in which the actual position - ground truth - of a counter in the ranking is far from the position computed from the learned preference scores. Considering the previous example (Figure 5),  $NRMSE = 0.57$ . This number well represents what is happening. The counters *TrainTrips* and *WalkKm* have been successfully positioned in the ranking (in both orderings they have positions 4 and 3 respectively). However, *BikeKm* and *BusTrips* and *GreenLeaves* have been ranked poorly. *BikeKm* should be in position 2, but in

the algorithm's sorting  $s_s$  is in position 5; *BusTrips* in the algorithm's outcome is in position 2, but should be in position 1; and *GreenLeaves* should be in the last (position 5), but in  $s_s$  assumes the first position.

	<i>min</i>	<i>mean</i>	<i>median</i>	<i>max</i>
<i>NRMSE</i>	0	0.27	0.24	0.8
$R^2$	-3	0.014	0.6	1
Similarity	0.2	0.83	0.93	1

**Table 1: Evaluation of the Algorithm.** Distribution of the error and similarity measures over the 94 eligible players.

**Players Characterization.** As Table1 shows, the evaluation of our algorithm had very variable outcomes. We found that for 51% of the players the prediction was very accurate (on average  $NRMSE = 0.13$ ,  $R^2 = 0.81$ , and  $sim = 0.97$ ), while for the remaining users the algorithm worked poorly. Thus, we investigated the causes of such variability. We formulated and tested the following hypotheses. To test the hypothesis, we divided the initial set of 94 players in the ones achieving good performance and in the ones for which the algorithm failed in learning their preferences. We split the dataset considering first the *NRMSE* value and then the  $R^2$  value. We did not consider the cosine similarity because it proved to be an optimistic measure (see Table1). We set two thresholds. For the *NRMSE* the worst performing players were the ones for which the error was  $> 0.2$ . Thus, we admit an error of 20%. For the  $R^2$  values we set a threshold of 0.5 for the best performing players. Before any analysis, we conducted the D'Agostino-Pearson  $K^2$  Test to determine whether the data were normally distributed. Then, to compare the difference between the two groups, we run (one-tailed) t-test on data normally distributed, the Mann-Whitney W test otherwise.



**Figure 6: Distribution of the number of bandits over the two group of players.**

**H1.** *The number of options (bandits) influenced the goodness of the learning.*

For how the challenge suggestion is designed, players may have been exposed to a smaller set of challenge types. We tested whether a smaller number of bandits (smaller  $k$ ) affected the performance.

Since both groups (players for which the algorithm had good performance and players for which the algorithm had bad performance) were not normally distributed (for the 1<sup>st</sup> group  $K^2 = 3.9$ , p-value > .05; and for the 2<sup>nd</sup> group  $K^2 = 8.9$ , p-value < .05), we computed the Mann-Whitney W test. The difference between players in the first group ( $M = 4.32$ ,  $SD = 0.70$ ) and players in the second group ( $M = 4.60$ ,  $SD = 0.56$ ), in terms of number of counters, was significant ( $U = 840.5$ , p-value < .05). Although the number of bandits impacted the algorithm's performance, the number of choices per player in the two groups are still well-distributed (Figure 6).

**H2.** *The number of challenges - i.e., iterations - influenced the quality of the outcomes.*

The number of challenges available for learning players' preferences depended on players' gameplay length. More challenges meant more iterations of the algorithm, and thus, we hypothesized that the number of challenges affected the learning accuracy. Since both of our group were not normally distributed (for the 1<sup>st</sup> group  $K^2 = 2.7$ , p-value > .05; and for the 2<sup>nd</sup> group  $K^2 = 3.2$ , p-value > .05), we computed the Mann-Whitney W test. The difference between players in the first group ( $M = 14.78$ ,  $SD = 4.40$ ) and players in the second group ( $M = 15.26$ ,  $SD = 04.79$ ), in terms of number of counters, was not significant ( $U = 975$ , p-value = .27). Therefore, the number of iterations per players did not significantly affect the general performance of the algorithm.

**H3.** *Players that manifested a much stronger preference were easier to model.*

The hypothesis was that it was difficult to model players manifesting shallow preferences, in that they explored several challenge types. We measured the strength of players' preferences by calculating the relative standard deviation (*RSD*) of the preference scores from the ground truth. *RSD* is often expressed as a percentage and is defined as the ratio of the standard deviation to the mean. We opted for the relative version of the SD to make it comparable among the entries. Since both for our group were not normally distributed (for the 1<sup>st</sup> group  $K^2 = 4.1$ , p-value > .05; and for the 2<sup>nd</sup> group  $K^2 = 0.9$ , p-value > .05), we computed the Mann-Whitney W test. Players in the first group ( $M = 104$ ,  $SD = 34.49$ ) manifested a higher relative variance in the preference scores ( $U = 1339$ , p-value < .01) than players in the second ( $M = 88.5$ ,  $SD = 22.6$ ). Thus, the algorithm succeeded in learning preferences for players that manifested stronger preferences - i.e., higher scores' variances.

Player Type	U	p-value	mean <sub>1</sub>	std <sub>1</sub>	mean <sub>2</sub>	std <sub>2</sub>
Free Spirit	889	0.4	23.12	3.15	23.3	2.87
Achiever	989	0.7	22.56	3.75	22	3.63
Player	865	0.33	17.82	5.98	18.6	5.47
Disruptor	931	0.56	14.74	5	14.26	4.7
Philanthropist	1060.5	0.89	24	2.9	23	3.67
Socializer	816.5	0.19	21.36	4.39	21.92	5.11

**Table 2:** Results of the test of hypothesis *H4*, investigating whether player types characterize the two groups of players (good performance of the algorithm vs bad performance of the algorithm).

**H4.** *Player personality represented by their player types scores affected the modeling power of the algorithm.*

We also researched the causes of such heterogeneity in players' identity, represented by the player type scores computed from the survey submitted at the beginning of the gameplay. In the registration phase players were asked to fill out a survey, containing the items of the Hexad Player types [59]. Thus, we analyzed the scores of the 6 player types (Socializer, Player, Free Spirit, Disruptor, Philanthropist and Achiever) to explain such variability. We computed the Mann-Whitney W test since the normality requirements were not satisfied we computed the Mann-Whitney W test:

- **Free Spirit.** For the 1<sup>st</sup> group  $K^2 = 0.54$ , p-value > .05; and for the 2<sup>nd</sup> group  $K^2 = 1.38$ , p-value > .05.
- **Player.** For the 1<sup>st</sup> group  $K^2 = 1.7$ , p-value > .05; and for the 2<sup>nd</sup> group  $K^2 = 1.6$ , p-value > .05.
- **Achiever.** For the 1<sup>st</sup> group  $K^2 = 1.5$ , p-value > .05; and for the 2<sup>nd</sup> group  $K^2 = 7.5$ , p-value < .05.
- **Disruptor.** For the 1<sup>st</sup> group  $K^2 = 3.4$ , p-value > .05; and for the 2<sup>nd</sup> group  $K^2 = 1.8$ , p-value > .05.
- **Philanthropist.** For the 1<sup>st</sup> group  $K^2 = 3$ , p-value > .05; and for the 2<sup>nd</sup> group  $K^2 = 6.2$ , p-value < .05.
- **Socializer.** For the 1<sup>st</sup> group  $K^2 = 3.7$ , p-value > .05; and for the 2<sup>nd</sup> group  $K^2 = 14.8$ , p-value < .05.

Table 2 shows the outcomes of the analyses, which point that Hexad Player Type scores cannot be used as indicators to identify players for which the algorithm would perform well.

## 6 DISCUSSION

In the last few years, the body of research around customizable and self-adapting systems is rapidly growing. With the enhances in technology and AI, users expect the system to be omniscient in delivering a highly tailored experience. Although users do not always explicitly manifest their needs, if the system fails in delivering what they expect, a churn may occur. Retaining users is also complicated by a seemingly an endless market's offer, and hence, an increasing competition. Since keeping users is easier than obtaining new ones, retention is crucial for the life of the system. To avoid churn, the system should deliver engaging experiences. Besides, since humans have heterogeneous needs, the system should be adaptable and intelligent in learning how to meet the users' requirements.

In our work, we moved a step towards this direction by testing our approach in a persuasive gameful system. We conducted a feasibility study by investigating the usage of in-game behaviors to evaluate player preferences. Instead of considering specific features to characterize players, we focused on in-game behaviors monitoring players' activity levels. This choice was moved by an urgent requirement coming from the gamification application domain since gameful systems tend to implement simple game mechanics. The only feature inherently implemented is the activity level, which emerges the very moment that interactions are allowed. Therefore, having only activity-related data was set as a constraint for our study.

Our use case scenario was a persuasive gamified mobile application, promoting positive behavior in terms of sustainable mobility. The app uses customized weekly challenges to keep the player engaged by fostering a sense of flow. The gamified system allows players

to select a challenge among two or three choices, to be activated in the following game week. In the case of no choice, one of the options was activated automatically.

To answer our research questions, we feed to our algorithm players' data, separately, and used the values of the activity behaviors as rewards. To evaluate the algorithm's performance, we used the players' in-game choices to build the ground truth.

#### **RQ1. Can we learn players' preferences by studying in-game behaviors measuring their level of activity?**

We conducted a preliminary analysis in which we verified whether there was a correlation among the participation level players manifested and the challenge they were engaging with. Towards this, we built a dataset enclosing information on players' identity and the type of challenge active for that timeframe. We obtained a  $R^2$  value of 0.93, which indicates a high degree of correlation. Thus, the outcomes suggested that players' activity behaviors were influenced by the type of challenge active per player.

Then, we fed players' gameplay data to our algorithm to obtain an ordered list of the challenge types, for each player. To properly answer the research question, we evaluated the algorithm's outcomes. We compared the learned preferences with players' real preferences, expressed through in-game choices. We computed three types of metrics: NRMSE, which represents how different the two sorting are; and  $R^2$  and Similarity, which measure the goodness of the outcome.

We found that for a portion of the users, we obtained very accurate outcomes. While, for others, the accuracy was poor. More specifically, for 51% of the population the algorithm succeeded in learning their preferences (on average NRSME = 0.13 and  $R^2 = 0.81$ ).

#### **RQ2. Does this approach work better for a subset of players? If so, how are those players characterized?**

We researched why the performance of our algorithm was very variable. Towards this, we formulated and verified 4 hypotheses. In particular, we divided players in two groups, according to how accurate the learned preferences were. Then we studied the differences among the two.

We found that the number of challenge types and the number of challenges available did not impact the goodness of the results (H1 and H2). We also found that the accuracy of outcomes were independent to the player types characterizing the player (H4).

We measured whether players that manifested a strong preference of challenge type affected the accuracy of the learned preferences (H3). Towards this, we compared the relative standard deviation of the preference scores (in the ground truth) of players. We found that the algorithm behaved considerably better for players with a higher variance. This finding highlights an aspect proper of the problem itself. The purpose of the algorithm is to find preferences. However, if there is a lack of thereof, it fails in finding them. In other words, the algorithm succeeds in identifying players' preferences, assumed that players actually enjoyed some game element more than others.

## **6.1 Limitations**

The impact of our results is weakened by the modest amount of players eligible for the study. Such a limited sample of data impacted the outcome, in that (1) we had fewer data to learn from, and (2) some evaluation measures penalize smaller samples with stricter and pessimistic values. Besides, although using player choices as their ground truth in this particular work may be influenced by the time pressure of making a selection, filtering out players making incoherent decisions over time attenuates this limitation.

Nevertheless, we believe that this study can be of inspiration for future works. We believe that addressing those issues would lead to insightful and relevant knowledge applicable in on-the-field applications.

## **7 CONCLUSIONS AND FUTURE WORKS**

The work hereby presented contributes to the state of the art by moving towards an algorithm that learns, from in-game behaviors inherent to player activity, whether they enjoyed the customized game element deployed. In this study, we investigated whether such a simple measure as the level of activity is informative in assessing player preferences. The constraint of using a very limited source of knowledge links back to the nature of the use case employed: a gameful system, which rarely implements complex game mechanics. Nevertheless, being able to adapt the gameplay, even if simple, is essential for those systems to keep their players engaged. The outcomes show that this is a promising direction and that there still is plenty of work to be done. We hope that this work will also stimulate other researchers' curiosity and foster promising collaboration aimed at solving the issue of self-adapting game-based systems.

This study has been designed to be the ground for more research on the topic. In the future, we plan to upgrade the algorithm from an offline to an online algorithm. Having assessed the feasibility of using participation data to inform players' preferences, we can go back to the initial inspiration: a reinforcement learning algorithm. Therefore, the goal is to learn the players' preferences online and to deploy the best choice for every player. This would result in a modification of Line 7 of Algorithm , where, instead of already knowing the bandit/option active, the option would be picked according to the probabilities built upon the preference scores (Equation 2).

Following interesting investigations would be studying the time needed for the algorithm's convergence, and its ability to adapt to preferences' shifting. Finally, an interesting perspective could be comparing the quality of the results also in a more complex system, such as entertainment games, to understand whether the increased complexity facilitates or hinders the applicability of the approach.

## **REFERENCES**

- [1] Yigal Attali and Meirav Arieli-Attali. 2015. Gamification in assessment: Do points affect test performance? *Computers & Education* 83 (2015), 57–63. <https://doi.org/10.1016/j.compedu.2014.12.012>
- [2] Myat Aung, Simon Demediuk, Yuan Sun, Ye Tu, Yu Ang, Siva Nekkanti, Shantanu Raghav, Diego Klabjan, Rafet Sifa, and Anders Drachen. 2019. The trails of Just Cause 2: spatio-temporal player profiling in open-world games. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*. ACM, 41.
- [3] Roger Bakeman and John M. Gottman. 1997. *Observing Interaction: An Introduction to Sequential Analysis*.
- [4] Richard Bartle. 1996. Hearts, Clubs, Diamonds, Spades: Players Who Suit MUDs. 1 (1996), 19.

- [5] Shlomo Berkovsky, Mac Coombe, Jill Freyne, Dipak Bhandari, and Nilufar Baghaei. 2010. Physical Activity Motivating Games: Virtual Rewards for Real Activity. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 243–252. <https://doi.org/10.1145/1753326.1753362> event-place: Atlanta, Georgia, USA.
- [6] P. Bertens, A. Guitart, and Á Periañez. 2017. Games and big data: A scalable multi-dimensional churn prediction model. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*. 33–36. <https://doi.org/10.1109/CIG.2017.8080412>
- [7] Alessandro Canossa, Sasha Makarovych, Julian Togelius, and Anders Drachenn. 2018. Like a DNA String: Sequence-Based Player Profiling in Tom Clancy's The Division. In *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.
- [8] Alessandro Canossa and Gillian Smith. 2015. Towards a Procedural Evaluation Technique: Metrics for Level Design. In *Proceedings of the 10th International Conference on the Foundations of Digital Games*. Pacific Grove, CA, USA.
- [9] Mihaly Csikszentmihalyi. 1997. *Finding flow: The psychology of engagement with everyday life*. Basic Books, New York, NY, US.
- [10] Luis de Marcos, Eva Garcia-Lopez, and Antonio Garcia-Cabot. 2016. On the effectiveness of game-like and social approaches in learning: Comparing educational gaming, gamification & social networking. *Computers & Education* 95 (2016), 99–113. <https://doi.org/10.1016/j.compedu.2015.12.008>
- [11] Thomas Debeauvais, Bonnie Nardi, Diane J. Schiano, Nicolas Ducheneaut, and Nicholas Yee. 2011. If You Build It They Might Stay: Retention Mechanisms in World of Warcraft. In *Proceedings of the 6th International Conference on Foundations of Digital Games (FDG '11)*. ACM, New York, NY, USA, 180–187. <https://doi.org/10.1145/2159365.2159390> event-place: Bordeaux, France.
- [12] Simon Demediuk, Alexandra Murrin, David Bulger, Michael Hitchens, Anders Drachen, William L. Raffe, and Marco Tamassia. 2018. Player retention in league of legends: a study using survival analysis. In *Proceedings of the Australasian Computer Science Week Multiconference*. ACM, 43.
- [13] Simon Demediuk, Marco Tamassia, Xiaodong Li, and William L. Raffe. 2019. Challenging AI: Evaluating the Effect of MCTS-Driven Dynamic Difficulty Adjustment on Player Enjoyment. In *ACSW*. 43–1.
- [14] Simon Demediuk, Marco Tamassia, William L. Raffe, Fabio Zambetta, Xiaodong Li, and Florian Mueller. 2017. Monte carlo tree search based algorithms for dynamic difficulty adjustment. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 53–59.
- [15] Sebastian Deterding, Dan Dixon, Rilla Khaled, and Lennart Nacke. 2011. From Game Design Elements to Gamefulness: Defining "Gamification". In *Proceedings of the 15th International Academic MindTrek Conference: Envisioning Future Media Environments (MindTrek '11)*. ACM, New York, NY, USA, 9–15. <https://doi.org/10.1145/2181037.2181040> event-place: Tampere, Finland.
- [16] A. Drachen, A. Canossa, and G. N. Yannakakis. 2009. Player modeling using self-organization in Tomb Raider: Underworld. In *2009 IEEE Symposium on Computational Intelligence and Games*. 1–8. <https://doi.org/10.1109/CIG.2009.5286500>
- [17] Anders Drachen, James Green, Chester Gray, Elie Harik, Patty Lu, Rafet Sifa, and Diego Klabjan. 2016. Guns and guardians: Comparative cluster analysis and behavioral profiling in destiny. In *2016 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 1–8.
- [18] A. Drachen, R. Sifa, C. Bauckhage, and C. Thurau. 2012. Guns, swords and data: Clustering of player behavior in computer games in the wild. In *2012 IEEE Conference on Computational Intelligence and Games (CIG)*. 163–170. <https://doi.org/10.1109/CIG.2012.6374152>
- [19] Magy Seif El-Nasr, Anders Drachen, and Alessandro Canossa. 2016. *Game analytics*. Springer.
- [20] Lauren S. Ferro, Steffen P. Walz, and Stefan Greuter. 2013. Towards Personalised, Gamified Systems: An Investigation into Game Design, Personality and Player Typologies. In *Proceedings of The 9th Australasian Conference on Interactive Entertainment: Matters of Life and Death (IE '13)*. ACM, New York, NY, USA, 7:1–7:6. <https://doi.org/10.1145/2513002.2513024> event-place: Melbourne, Australia.
- [21] Michela Ferron, Enrica Loria, Marconi Annapaola, and Massa Paolo. 2019. Play&Go, an Urban Game Promoting Behaviour Change for Sustainable Mobility. *Interaction Design and Architecture(s) Journal - IxD&A* (2019), 24 – 45.
- [22] F. Hadiji, R. Sifa, A. Drachen, C. Thurau, K. Kersting, and C. Bauckhage. 2014. Predicting player churn in the wild. In *2014 IEEE Conference on Computational Intelligence and Games*. 1–8. <https://doi.org/10.1109/CIG.2014.6932876>
- [23] Danial Hooshyar, Moslem Yousefi, and Heuseok Lim. 2018. Data-Driven Approaches to Game Player Modeling: A Systematic Literature Review. *ACM Comput. Surv.* 50, 6 (2018), 90:1–90:19. <https://doi.org/10.1145/3145814>
- [24] Danial Hooshyar, Moslem Yousefi, and Heuseok Lim. 2018. A Procedural Content Generation-Based Framework for Educational Games: Toward a Tailored Data-Driven Game for Developing Early English Reading Skills. *Journal of Educational Computing Research* 56, 2 (2018), 293–310. <https://doi.org/10.1177/0735633117706909>
- [25] Ruud S. Jacobs. 2016. Play to Win Over: Audiences and Effects of Persuasive Games.
- [26] Tomislav Jaguš, Ivica Botički, and Hyo-Jeong So. 2018. Examining competitive, collaborative and adaptive gamification in young learners' math learning. *Computers & Education* 125 (2018), 444–457. <https://doi.org/10.1016/j.compedu.2018.06.022>
- [27] Rishabh Joshi, Varun Gupta, Xinyue Li, Yue Cui, Ziwen Wang, Yaser Ravari, Rafet Sifa, Azita Parsaianen, Anders Drachen, and Simon Peter Demediuk. 2019. A Team Based Player Versus Player Recommender Systems Framework For Player Improvement. In *Proceedings of ACSW*. York.
- [28] Maurits Kaptein, Boris De Ruyter, Panos Markopoulos, and Emile Aarts. 2012. Adaptive Persuasive Systems: A Study of Tailored Persuasive Text Messages to Reduce Snacking. *ACM Trans. Interact. Intell. Syst.* 2, 2 (2012), 10:1–10:25. <https://doi.org/10.1145/2209310.2209313>
- [29] Michael D. Kickmeier-Rust, Eva-C. Hillemann, and Dietrich Albert. 2014. Gamification and Smart Feedback: Experiences with a Primary School Level Math App. *International Journal of Game-Based Learning (IJGBL)* 4, 3 (2014), 35–46. <https://doi.org/10.4018/ijgb.2014070104>
- [30] Kristian Kili, Korbinian Moeller, and Manuel Ninaus. 2018. Evaluating the effectiveness of a game-based rational number training - In-game metrics as learning indicators. *Computers & Education* 120 (2018), 13–28. <https://doi.org/10.1016/j.compedu.2018.01.012>
- [31] Jun H. Kim, Daniel V. Gunn, Eric Schuh, Bruce Phillips, Randy J. Pagulayan, and Dennis Wixon. 2008. Tracking Real-time User Experience (TRUE): A Comprehensive Instrumentation Solution for Complex Systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 443–452. <https://doi.org/10.1145/1357054.1357126> event-place: Florence, Italy.
- [32] Antti Knutas, Rob van Roy, Timo Hynninen, Marco Granato, Jussi Kasurinen, and Jouni Ikonen. 2019. A process for designing algorithm-based personalized gamification. *Multimedia Tools and Applications* 78, 10 (2019), 13593–13612. <https://doi.org/10.1007/s11042-018-6913-5>
- [33] Jonna Koivisto and Juho Hamari. 2019. The rise of motivational information systems: A review of gamification research. *International Journal of Information Management* 45 (2019), 191–210. <https://doi.org/10.1016/j.ijinfomgt.2018.10.013>
- [34] Johannes Kunkel, Benedikt Loepp, and Jürgen Ziegler. 2017. A 3D item space visualization for presenting and manipulating user preferences in collaborative filtering. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*. ACM, 3–15.
- [35] Nikoleta Zampeta Legaki, Nannan Xi, Juho Hamari, and Vassilios Assimakopoulos. 2019. Gamification of The Future: An Experiment on Gamifying Education of Forecasting. <https://doi.org/10.24251/HICSS.2019.219>
- [36] C. Lopez and C. Tucker. 2018. Towards Personalized Adaptive Gamification: A Machine Learning Model for Predicting Performance. *IEEE Transactions on Games* (2018), 1–1. <https://doi.org/10.1109/TG.2018.2883661>
- [37] Enrica Loria and Annapaola Marconi. 2018. Player Types and Player Behaviors: Analyzing Correlations in an On-the-field Gamified System. In *Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts - CHI PLAY '18 Extended Abstracts*. ACM Press, Melbourne, VIC, Australia, 531–538. <https://doi.org/10.1145/3270316.3271526>
- [38] Enrica Loria and Annapaola Marconi. 2019. A Data-Driven Approach to Deduce Players' Preferences from In-Game Interactions in Gameful Systems. In *Proceedings of The First Game Analytics Workshop*, Vol. In Press. CEUR Workshop Proceedings, Atlanta, GA, USA.
- [39] Annapaola Marconi and Enrica Loria. 2020. Sustainable Mobility in Smart Cities: The Key Role of Gamified Motivational Systems for Citizens' Engagement and Behavior Change. In *Implications of Mobility as a Service (MaaS) in Urban and Rural Environments: Emerging Research and Opportunities*, IGI Global (Ed.). Hershey, PA, 211–246.
- [40] Andrzej Marczewski. 2018. *Even Ninja Monkeys Like to Play: Unicorn Edition*.
- [41] Baptiste Monterrat, Élise Lavoué, and Sébastien George. 2017. Adaptation of Gaming Features for Motivating Learners. *Simulation & Gaming* 48, 5 (2017), 625–656. <https://doi.org/10.1177/1046878117712632>
- [42] Benedikt Morschheuser and Juho Hamari. 2019. The Gamification of Work: Lessons From Crowdsourcing. *Journal of Management Inquiry* 28, 2 (2019), 145–148. <https://doi.org/10.1177/1056492618790921>
- [43] Lennart E. Nacke, Chris Bateman, and Regan L. Mandryk. 2014. BrainHex: A neurobiological gamer typology survey. *Entertainment Computing* 5, 1 (2014), 55–62. <https://doi.org/10.1016/j.entcom.2013.06.002>
- [44] Aline Normoyle and Shane T. Jensen. 2015. Bayesian Clustering of Player Styles for Multiplayer Games. In *Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*.
- [45] Dmitry Nozhnin. 2012. Predicting Churn: When Do Veterans Quit?
- [46] Rita Orji, Regan L. Mandryk, and Julita Vassileva. 2017. Improving the Efficacy of Games for Change Using Personalization Models. *ACM Trans. Comput.-Hum. Interact.* 24, 5 (2017), 32:1–32:22. <https://doi.org/10.1145/3119929>
- [47] Rita Orji, Lennart E. Nacke, and Chrysanne Di Marco. 2017. Towards Personality-driven Persuasive Health Games and Gamified Systems. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New

- York, NY, USA, 1015–1027. <https://doi.org/10.1145/3025453.3025577> event-place: Denver, Colorado, USA.
- [48] Ranilson Paiva, Ig Ibert Bittencourt, Thyago Tenório, Patricia Jaques, and Seiji Isotani. 2016. What do students do on-line? Modeling students' interactions to improve their learning experience. *Computers in Human Behavior* 64 (2016), 769–781. <https://doi.org/10.1016/j.chb.2016.07.048>
- [49] Carlos Pereira Santos, Kevin Hutchinson, Vassilis-Javed Khan, and Panos Markopoulos. 2017. Measuring self-esteem with games. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*. ACM, 95–105.
- [50] Maria Perez-Ortiz and Rafal K. Mantiuk. 2017. A practical guide and software for analysing pairwise comparison experiments. *arXiv:1712.03686 [cs, stat]* (2017). arXiv: 1712.03686.
- [51] Felix Putze, Mazen Salous, and Tanja Schultz. 2018. Detecting memory-based interaction obstacles with a recurrent neural model of user behavior. In *23rd International Conference on Intelligent User Interfaces*. ACM, 205–209.
- [52] Amon Rapp, Frank Hopfgartner, Juho Hamari, Conor Linehan, and Federica Cena. 2019. Strengthening gamification studies: Current trends and future opportunities of gamification research. *International Journal of Human-Computer Studies* 127 (2019), 1–6. <https://doi.org/10.1016/j.ijhcs.2018.11.007>
- [53] Rafet Sifa, Anders Drachen, and Christian Bauckhage. 2018. Profiling in Games: Understanding Behavior from Telemetry. *Social Interactions in Virtual Worlds: An Interdisciplinary Perspective* (2018).
- [54] R. Sifa, A. Drachen, C. Bauckhage, C. Thureau, and A. Canossa. 2013. Behavior evolution in Tomb Raider Underworld. In *2013 IEEE Conference on Computational Intelligence in Games (CIG)*. 1–8. <https://doi.org/10.1109/CIG.2013.6633637>
- [55] Tom Stafford and Erwin Haasnoot. 2017. Testing sleep consolidation in skill learning: A field study using an online game. *Topics in cognitive science* 9, 2 (2017), 485–496.
- [56] HALLIFAX Stuart, Audrey Serna, Jean-Charles Marty, and Elise Lavoué. 2019. Adaptive gamification in education: A literature review of current trends and developments. In *European Conference on Technology Enhanced Learning (EC)*. Delft, Netherlands.
- [57] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement learning: an introduction* (second edition ed.). The MIT Press, Cambridge, Massachusetts.
- [58] M. Tamassia, W. Raffe, R. Sifa, A. Drachen, F. Zambetta, and M. Hitchens. 2016. Predicting player churn in destiny: A Hidden Markov models approach to predicting player departure in a major online game. In *2016 IEEE Conference on Computational Intelligence and Games (CIG)*. 1–8. <https://doi.org/10.1109/CIG.2016.7860431>
- [59] Gustavo Fortes Tondello, Deltcho Valtchanov, Adrian Reetz, Rina R. Wehbe, Rita Orji, and Lennart E. Nacke. 2018. Towards a Trait Model of Video Game Preferences. *International Journal of Human-Computer Interaction* 34, 8 (2018), 732–748. <https://doi.org/10.1080/10447318.2018.1461765>
- [60] Gustavo F. Tondello, Rina R. Wehbe, Lisa Diamond, Marc Busch, Andrzej Marczewski, and Lennart E. Nacke. 2016. The Gamification User Types Hexad Scale. In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play (CHI PLAY '16)*. ACM, New York, NY, USA, 229–243. <https://doi.org/10.1145/2967934.2968082>
- [61] Julita Vassileva. 2012. Motivating participation in social computing applications: a user modeling perspective. *User Modeling and User-Adapted Interaction* 22, 1 (2012), 177–201. <https://doi.org/10.1007/s11257-011-9109-5>
- [62] Guenter Wallner. 2015. Sequential Analysis of Player Behavior. In *Proceedings of the 2015 Annual Symposium on Computer-Human Interaction in Play (CHI PLAY '15)*. ACM, New York, NY, USA, 349–358. <https://doi.org/10.1145/2793107.2793112> event-place: London, United Kingdom.
- [63] Günter Wallner, Christopher Schinnerl, Michael Helfried Schiller, Alexander Monte Calvo, Johanna Pirker, Rafet Sifa, and Anders Drachen. 2019. Beyond the individual: Understanding social structures of an online player matchmaking website. *Entertainment Computing* 30 (2019), 100284.
- [64] Harald Warmelink, Jonna Koivisto, Igor Mayer, Mikko Vesa, and Juho Hamari. 2018. Gamification of the work floor: A literature review of gamifying production and logistics operations. *Hawaii International Conference on System Sciences 2018 (HICSS-51)* (2018).
- [65] Helen Wauck and Wai-Tat Fu. 2017. A Data-Driven, Multidimensional Approach to Hint Design in Video Games. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*. ACM, 137–147.
- [66] Geogios N. Yannakakis. 2012. Game AI Revisited. In *Proceedings of the 9th Conference on Computing Frontiers (CF '12)*. ACM, New York, NY, USA, 285–292. <https://doi.org/10.1145/2212908.2212954> event-place: Cagliari, Italy.
- [67] G. N. Yannakakis and J. Togelius. 2011. Experience-Driven Procedural Content Generation. *IEEE Transactions on Affective Computing* 2, 3 (2011), 147–161. <https://doi.org/10.1109/T-AFFC.2011.6>
- [68] G. N. Yannakakis and J. Togelius. 2015. Experience-driven procedural content generation (Extended abstract). In *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)*. 519–525. <https://doi.org/10.1109/ACII.2015.7344619>
- [69] Ping Zhang. 2008. Motivational Affordances: Reasons for ICT Design. *Technical Opinion* (2008).