

Your Gameplay Says It All: Modelling Motivation in *Tom Clancy's The Division*

David Melhart*, Ahmad Azadvar[†], Alessandro Canossa[†], Antonios Liapis* and Georgios N. Yannakakis*

**Institute of Digital Games, University of Malta*

{david.melhart, antonios.liapis, georgios.yannakakis}@um.edu.mt

†Massive Entertainment a Ubisoft Studio

{ahmad.azadvar, alessandro.canossa}@massive.se

Abstract—Is it possible to predict the motivation of players just by observing their gameplay data? Even if so, how should we measure motivation in the first place? To address the above questions, on the one end, we collect a large dataset of gameplay data from players of the popular game *Tom Clancy's The Division*. On the other end, we ask them to report their levels of *competence, autonomy, relatedness* and *presence* using the *Ubisoft Perceived Experience Questionnaire*. After processing the survey responses in an ordinal fashion we employ preference learning methods based on support vector machines to infer the mapping between gameplay and the reported four motivation factors. Our key findings suggest that gameplay features are strong predictors of player motivation as the best obtained models reach accuracies of near certainty, from 92% up to 94% on unseen players.

Index Terms—Self-determination theory, affective computing, digital games, player modelling, preference learning

I. INTRODUCTION

Fostering long term player engagement has been a traditional core challenge of game design and development. Players who are self-motivated to return to the game and keep playing are critical for a game's success [1]–[3]. The central role of *motivation* for the design of games, and the experiences they elicit, has been highlighted by a growing number of studies which adopt psychological theories of motivation within games [4]–[6]. Such studies, however, follow a top-down integration of phenomenological models of motivation, which aim to identify and explain stereotypical player behaviour. Over the last decade, games user research and industry-based game testing has shifted its focus towards quantitative approaches to understand player behaviour and experience, based on *player analytics* [7], [8]. These approaches focus mainly on either clustering players based on their behavioural patterns or predicting objectively-defined aspects of their gameplay behaviour (e.g. churn prediction) [3]. Despite these efforts, the majority of approaches that aim to capture aspects of player experience, such as engagement or motivation, remain qualitative in their analysis—even when quantitative data is involved—due to the complexity of measuring subjective notions of user experience in games [9].

Motivated by the lack of quantitative studies on the relationship between motivation and play, in this paper we introduce a data-driven *player modelling* approach [10] by assuming there is an unknown underlying function between what a player

does in the game *behaviourally*—as manifested through her gameplay data—and her motivation. In particular, we assume that solely behavioural data from a player's gameplay would yield accurate predictors of motivation in games. To define motivation we rely theoretically on *Self Determination Theory* [11] and examine four core factors: *competence, autonomy, relatedness* and *presence*, the latter of which is often associated with the theory in the domain of videogames [1].

To study motivation quantitatively we are grounded in recent developments in motivation measurement tools, namely the *Ubisoft Perceived Experience Questionnaire* (UPEQ) [12], which was developed as a game-specific tool to observe player motivation. To infer the relationship between player motivation and gameplay we collect data from more than 400 players of *Tom Clancy's The Division* (Ubisoft, 2016). We process and aggregate this data and collect surveys on the players' motivation in relation to the game independently. We use the UPEQ questionnaire to measure players' general levels of *competence, autonomy, relatedness* and *presence* in the game. Given the subjective nature of the reported notions we adopt a *second-order* data processing approach [13] and we process the reported UPEQ Likert-scale values of the players as ordinal data, and not as scores. We then apply simple statistical rank-based models and *preference learning* [14] methods based on support vector machines to infer the function between gameplay and reported factors of motivation. Our results suggest that factors of reported motivation can be predicted with high accuracy just relying on a few high-level gameplay features. In particular, the nonlinear machine learned models manage to predict the four motivation factors of unseen players with around 80% average accuracy; the best models are reaching an accuracy of 94% for *competence* and *autonomy* and 92% for *relatedness* and *presence*. The obtained results add to the existing evidence for the benefits of ordinal data processing on subjectively-defined notions [13] and validate that motivation can be accurately captured in the examined game based only on behavioural high-level data of playing.

This paper is novel in a number of ways. First, this is the first time player motivation (as in *self-determination*) is modelled computationally only through gameplay data in games. Second, we introduce a *second-order* [13] methodology for treating Likert-scale scores which are used frequently in game testing and games user research at large. The ordinal approach

we adopt compares the subjective scores of all players with each other and hence generates combinatorially very large datasets based only on small sets of participants. The approach is also effective in eliminating reporting biases of respondents, thereby better approximating the ground truth of reported motivation. Third, we model aspects of player motivation using preference learning based solely on a small number of key gameplay features. Finally, for the first time we evaluate the method in *Tom Clancy's The Division* (Ubisoft, 2016) on a large set of players and the predictive capacity of the motivation models for this game reach near certainty (i.e., over 80% of average accuracy).

II. BACKGROUND: MEASURING AND MODELLING MOTIVATION

This section gives an overview of related research on player modelling and motivation studies in games (Section II-A) and then it introduces the fundamental principles of *Self Determination Theory* (SDT) and the notion of *presence* within the framework of UPEQ (Section II-B). The section ends with a discussion on the strengths of preference learning for modelling subjectively-defined psychological constructs such as motivation (Section II-C).

A. Player Modelling and Player Motivation

Player modelling is one of the primary foci of AI research in the field of videogames [3]. It is generally concerned with the prediction of player behaviour or other cognitive and affective processes. Among their many uses, player models can inform and shape a game's monetization strategy, they can be directly applied as drivers of personalised content generation [15], [16], and they can equip agents for believable user testing [17].

Player modelling often relies on *clustering* or *prediction* [3]. Clustering is based on unsupervised learning methods with the aim to cluster players within groups of common behavioural patterns; approaches include *k*-means, self-organising maps [7], matrix factorisation and archetypal analysis [18], [19], and sequence mining [8], [20], [21]. Prediction uses supervised learning to predict patterns of playing such as completion time [22] and churn [23]–[25]. Predictive models can predict a player's *behaviour* (i.e., *what would a player do?*) [26] or the game *experience* (i.e., *what would a player feel?*) [3]. This type of modelling often uses behavioural data but also other modalities of player input such as physiological data [27], [28], independently or in a multimodal fashion [29], [30].

In the literature, motivation applied to games is built primarily on theoretical models that inform the design of experimental protocols instead of defining target outputs for predictive modelling. Indicatively, Borbora et al. [31] use Yee's player motivation [32] typology to enhance churn prediction, while Shim et al. [33] use motivational survey data to model player enjoyment. Birk et al. [34] also incorporated motivational survey data along with top-down personality and player profiles to model enjoyment and effort. Although Birk et al. include high-level telemetry of progression (game stage) as a control variable in their experiments, they rely mainly on

survey data. While the above studies use motivational survey data as top-down domain knowledge that forms the input of enjoyment models, this paper instead focuses on predicting motivation based on game metrics alone. Similarly to this study, Canossa et al. [4] applied regression models of telemetry to predict psychological survey data. However, that study was not aimed at player motivation modelling but instead used the Reiss Motivational Profile [35] as a tool for personality profiling based on personal motivational drives. This paper focuses on the structure of motivation by considering four of its core manifestations, but our models do not rely on personality profiles of the players; instead, we try to infer the unknown mapping between gameplay features and motivation.

B. From Theory to Measures of Motivation

Self-determination theory (SDT) is a well-established positive psychology theory of the facilitation of motivation [36] which has been adopted in a wide variety of domains, including videogames [1], [6], [37], [38]. SDT distinguishes between an *intrinsic* and *extrinsic locus of causality* behind motivation [39]. The latter is facilitated by external or internal rewards, pressures, and expectations, while the former is based on the intrinsic properties of the activity itself, namely how well it can support *competence*, *autonomy*, and *relatedness*. While videogames include pressures and rewards which can promote extrinsic motivation [40], they are generally regarded as good facilitators of intrinsic motivation [6], [38]. Ryan et al. [37] describe the basic psychological needs (and *presence*) underlying intrinsic motivation in videogames as:

1) *Competence*: a sense of accomplishment and a desire for the mastery of an action. It is tied to self-efficacy and a sense of meaningful progression. It is supported through the interactions that players must master to complete the game, but not completion in itself.

2) *Autonomy*: a sense of control and a desire for self-determined action. It manifests through meaningful choices, tactics, and strategic decisions that players can take. It is supported through rules and mechanics that structure the play experience but allow for a high degree of freedom and meaningfully different outcomes.

3) *Relatedness*: a sense of belonging and a desire to connect and interact with others. It manifests through interactions with other players and believable computer agents. It is supported by multilayer interactions, believable and rich non-player characters, narrative design, and even interactions with other players outside the game.

4) *Presence*: the feeling of a mediated experience is a main facilitator of both *competence* and *autonomy*. It can be viewed as having physical, emotional, and narrative components [1], [37]. While it is not considered as one of the basic psychological needs, presence or *immersion* can be a driving force behind gameplay motivation [6], [41], [42], and it is measured by both the *Player Experience of Need Satisfaction Questionnaire* [37] and UPEQ [12].

It is important to note that the above factors are not contributing equally to the formulation of intrinsic motivation;

while *competence* or *relatedness* are regarded as the core catalysts, *autonomy* generally plays a supporting role in the facilitation of motivation. Nevertheless, in absence of *autonomy*, motivation can only be considered *introjected* or compulsive [43]. Within games the main drive of intrinsic motivation is generally *competence* because of how the activity is structured, while *relatedness* contributes to enhancing the experience [1].

This paper measures the four factors of SDT as affected by the gameplay experience via UPEQ, a game-tailored questionnaire developed by researchers at *Massive Entertainment* [12] to predict gameplay outcomes relevant for industry designers and stakeholders. UPEQ measures the aforementioned factors of SDT through a 24 item survey using 5-point Likert-scales. From the 24 items, 21 measure the basic psychological needs of *competence*, *autonomy*, and *relatedness*, while 3 additional measure *presence*. Earlier work [12] has demonstrated that UPEQ is able to predict playtime, money spent on the game, and group playtime based on measured factors of SDT. UPEQ also addresses the limitations of prior domain-specific SDT questionnaires, such as the *Game Engagement Questionnaire* [44], *BrainHex* [45], and the *Player Experience of Need Satisfaction* [37], by focusing on the adaptation of the *Basic Need Satisfaction Scale(s)* [46] into a survey specific to videogame play. The result is a reliable and consistent assessment tool with a strong theoretical foundation in SDT.

C. The Ordinal Nature of Motivation

This study uses preference learning (PL) methods as a predictive player model, due to the assumption of an ordinal nature of player experience. PL focuses on the *differences* between occurrences instead of their *absolute values* [13], [47]. This approach falls much closer to the players’ cognitive processes—e.g. *anchoring-bias* [48], [49], *adaptation* [50], *habituation* [51], and other *recency-effects* [52]—that help them evaluate their own experience internally.

There is growing evidence supporting the strength of preference learning for modelling emotions and user experiences both on a conceptual [13], [47] and a technical basis [53]–[56]. Conceptually, treating subject-defined ground truth data as ordinal variables brings the representation of data closer to the players’ underlying true attitudes [47]. On the technical side, studies have compared processing of affective annotations as both ratings (e.g., Likert items) and rankings and found that 1) *first-order* data processing (i.e. ranks) yields higher reliability and inter-rater agreement and 2) *second-order* processing of the absolute rating values was also beneficial with regards to both reliability and validity [13], [54], [55], even when applied across different games [30] or dissimilar affective corpora [56].

Given the above theoretical framework on the ordinal nature of experience and the large body of recent empirical evidence on the benefits of the ordinal modelling approach [13], in this paper we view player motivation as an emotional construct [1] with ordinal properties. We compare player feedback on relative grounds and use PL to model the ranking between reported motivation in players as measured by the factors of UPEQ. While earlier work focused on the internal validity of

the survey data [12], this paper considers the UPEQ scores as the underlying *ground truth* to be predicted. After acquiring a general score for all the measured factors for each participant, we analyse and model the data as ordinal values, thereby following a *second-order* modelling approach [13].

III. PREFERENCE LEARNING FOR MODELLING MOTIVATION

Preference learning is a supervised machine learning technique [14], in which an algorithm learns to infer the preference relation between two variables. PL is a robust method which relies on relative associations instead of absolute values or class boundaries. PL applies a *pairwise transformation* to the original dataset, yielding a representation of differences between feature vectors in the query [57] which can be solved by a binary classifier. As an example, we observe the preference relation: $x_i \succ x_j \in X$ (x_i is preferred over x_j) based on their associated output: $y_i > y_j$. Through the pairwise transformation two new features are created: $x'_1 = (x_i - x_j)$, associated with $y'_1 = 1$ and $x'_2 = (x_j - x_i)$, associated with $y'_2 = -1$. This comparison between each pair of feature vectors provides $X' \subseteq X \cdot (X - 1)$ new datapoints. X' is a subset of all possible combinations because a clear preference relation can not always be inferred. During the pairwise transformation, a hyperparameter is applied to control the preference threshold (P_t) during comparisons. For instance, with a $P_t = 0.1$, only datapoints with a value difference above 10% of the value of the second datapoint are considered (i.e. a clear preference). The purpose of this threshold is to counter the noise in the ground truth data which can skew modelling results. Although increasing P_t eliminates insignificant differences and potential noise while increasing the accuracy of the model, sparsity of unique values in the ground truth can also lead to a rapid decrease in sample size, and thus hurt the robustness of the trained models.

This study uses ranking *Support Vector Machines* (SVM) [58] as they are implemented in the *Preference Learning Toolbox*¹ [59] which is based on the LIBSVM library [60]. We choose SVMs in this initial study as they can yield robust models even with a limited amount of data and input features. SVMs were originally employed to solve classification tasks by maximizing the margins of a hyperplane separating the datapoints projected into a higher dimensional feature space [61] but were later adopted to solve PL tasks as well [58]. We use both linear and non-linear SVMs with *radial basis function* (RBF) kernels. Unlike linear SVMs, which aim for a linear separation between datapoints, RBF SVMs emphasize the local proximity of datapoints, fitting the maximum-margin hyperplane in a transformed feature space [61]. For tuning our algorithms, we rely on the C regularization term which controls the trade-off between maximizing the margin and minimizing the classification error of the training set, and—in case of RBF kernels—the γ hyperparameter, which controls how each comparison is weighted in the non-linear topology.

¹<http://plt.institutedigitalgames.com/>



Fig. 1. An example of the gameplay of *Tom Clancy's The Division* (Ubisoft, 2016). Image taken from: store.steampowered.com/app/365590. No copyright infringement intended.

IV. THE GAME AND THE DATA

This section first presents the testbed game, and then discusses the collected data and steps taken to pre-process it.

A. *Tom Clancy's The Division*

This study uses in-game behavioural data (player metrics) and survey questionnaire responses from players of *Tom Clancy's The Division* (Ubisoft, 2016), hereafter *The Division*, collected between 2016 and 2018. *The Division* is an online multiplayer action-role playing game (Fig. 1) set in a post-apocalyptic New York. Players, as government agents, have to work together (and against each other) to scavenge and investigate the fallen city. Gameplay relies on third-person, cover-based, tactical shooting combat mechanics.

The core of the game is a progression system, in which players gain new levels by participating in different in-game activities including story-focused and optional missions. Levels unlock new abilities and activities offer new equipment (e.g. weapons, armour). The strength of a player can be measured by their level (up to 30) and by the quality of their equipment—expressed in *Gear Score* points. In the player versus environment (PvE) sections of the game, players can group up and complete missions together. The game also features a competitive player versus player (PvP) area—called the *Dark Zone*—which has its own progression system. In this special area players can still group up to complete missions for better equipment; however, they can also turn on each other and become *Rogue* by killing other players and taking their rewards. At the maximum level (30), players can participate in *Incursions*, which are particularly difficult missions for groups. Finally, downloadable content (DLC) adds new areas, equipment, and both PvE and PvP content to the game.

The game was well received (80/100 Metacritic score on consoles²) and was the best selling game of Ubisoft at the time of its release³. Due to a blend of PvP and PvE and support for different play styles and interaction modes, *The Division* is a rich and complex testbed for research on motivation. With the

second instalment of *The Division* set to release on March 15, 2019⁴, a study on how the gameplay of *The Division* shapes the motivation of its players is both timely and relevant.

B. Dataset

The collected data consists of aggregated information on the in-game activity of players over a long period of time and their corresponding UPEQ survey scores. These two types of data were collected independently, with the gameplay features recorded between 2016 and 2018 and the survey data collected through a web interface separately in 2018. As such, the survey data measures a general disposition of the players.

The dataset consists of one datapoint per player. In total, 443 players participated in the above-mentioned data collection process. The dataset is cleaned of datapoints with missing values, corrupted entries, and outliers to prevent skewing any statistical analysis process. Extensive pruning was necessary due to outliers distorting the distribution of general game metrics (see Section IV-C) and due to noise generated by the data logging service which inflated playtime. The clean dataset contains 298 players. As part of a preliminary step in the PL task, the dataset is converted through a pairwise transformation, using the P_t parameter (see Section III). We report our findings with $P_t = 0.1$, which preserves more than 60% of the possible training and test samples. The transformation is applied within cross-validation folds in a 10-fold setup. After applying the appropriate P_t , on average we are left with $44 \cdot 10^3$ training and 536 test samples for *competence*; $39.5 \cdot 10^3$ training and 471 test samples for *autonomy*; $56.5 \cdot 10^3$ training and 681 test samples for *relatedness*; $52 \cdot 10^3$ training and 632 test samples for *presence*.

C. Extracted Features

Player behaviour is measured through 30 high-level gameplay features. While most of these are simple aggregated *game metrics* describing the time allocation and progression of the player, 4 of them are exclusive categories of distinct *play styles* based on sequence-based profiling of the player's in-game activities [8]. Additionally, the dataset contains 4 Likert scores that represent the four *motivation factors* of each player as measured by the UPEQ survey. The three types of data considered in this study as detailed as follows:

1) *Game Metrics*: These features may relate to *general playtime* (Days Played, Days in Groups, Days in the Dark Zone, Sessions, Playtime, Group Playtime, Dark Zone Playtime, Playtime as Rogue); *completion* (Non-Daily Missions, Daily Missions, Side Missions, Days with Incursions, Incursions); *progression* (Gear-Score, Dark Zone Rank, Level, Early Level 30⁵, Reached Level 30); *early gameplay* (Level Below 30, Early Playtime, Early Group Playtime, Early Dark Zone Playtime, Early Playtime as Rogue); and *DLC gameplay* (Underground Playtime, Survival Playtime, Season-Pass).

⁴<https://news.ubisoft.com/en-us/article/342463/the-division-2-pc-features-specs-detailed/>

⁵Early Level 30 shows whether the player reached the maximum level faster than the average.

²<https://www.metacritic.com/game/xbox-one/tom-clancys-the-division>

³<https://news.ubisoft.com/en-us/article/313497/the-division-sets-new-sales-records/>

2) *Player Types*: The 4 different player types are named *Adventurer*, *Elite*, *PvE All-Rounder*, and *Social Dark Zone Player*. These types have been derived through a traditional k -means clustering and qualitative interpretation of highly aggregated data. Adventurers are focused on solo main and side missions, Elites have the top gear score, PvE All-rounders engage in cooperative activities, while the Dark Zone Players prefer competitive activities.

3) *Motivation Factors*: UPEQ scores the four factors of motivation in the form of averaged Likert-scale values. While computing the mean of ordinal data can be problematic conceptually [13], [55], average survey scores are a widespread method of using Likert-like data as they can still show certain tendencies within the scores (e.g., a higher score is assumed to correspond to an overall more positive response). As mentioned in Section II-C, we adopt a *second-order* modelling approach [13] and treat these scores as ordinal data through pairwise comparisons across all players.

V. MODELS OF PLAYER MOTIVATION

This section presents the results of the machine learned models of player motivation based on different feature sets. As the extracted play styles are more complex descriptors of the player behaviour than the aggregated game metrics, we examine their capacity of predicting motivation independently (see Section V-C and V-B) but also in fusion (Section V-D). Results presented in this section measure model accuracies in terms of predicting the higher *competence*, *autonomy*, *relatedness* and *presence* in pairwise comparisons of players. Before delving into the details of the obtained results, we outline the validation and parameter tuning process followed.

A. Validation and Parameter Tuning

All PL models are validated with 10-fold cross validation. To prevent data leakage, the training and test folds are separated before the normalisation and pairwise transformation of the data. After the pairwise transformation of the dataset, the training and test inputs are z -normalised. To preserve the independence of the test set, we assume that it is drawn from the same distribution as the training set and apply z -normalisation to it based on the properties of the training set.

The best C and $RBF\gamma$ hyperparameters are found in the $C \in \{1, 2, 3, 4, 5\}$ and $\gamma \in \{0.1, 0.5, 0.75, 1, 2\}$ parameter spaces. The best γ is 0.5 across all RBF models. For linear SVMs best performance is with $C = 2$ for competence; $C = 1$ for autonomy; $C = 4$ relatedness; $C = 3$ presence. For RBF SVMs best performance is with $C = 3$ for competence; $C = 4$ for autonomy; $C = 4$ for relatedness; $C = 2$ for presence.

B. Models Based on Play Styles

In the first round of experiments only the four play styles are used as input of the SVM model. Figure 2a shows the final results. Despite the low dimensionality of the input feature set, both linear and non-linear models are able to surpass the 50% baseline (with 57% on average). Introducing an RBF kernel does not improve the average accuracy and results

are comparable to linear models. Despite the low average performance, the best models reaching as high as 73.2% on average (*competence*: 73.2%; *autonomy*: 70.68%; *relatedness*: 65%; *presence*: 71%) Nevertheless, the low average score allows us to conclude that the four play styles on their own are not sufficient for building accurate models of motivation.

C. Models Based on Game Metrics

In the second round of PL experiments models are trained only considering the 26 game metrics as model inputs—as discussed in Section IV-C. As evidenced by the bar plots of Figure 2b, game metrics are quite successful in predicting the reported motivation factors. In particular, linear SVM models are successful with an average accuracy of 74.31% across all models while the best models for individual factors are performing at well above 80% accuracy on certain folds (*competence*: 89%; *autonomy*: 85.25%; *relatedness*: 83%; *presence*: 86.5%) Non-linear kernels further improve the model’s performance to a 79.3% accuracy on average across all models. The best individual models outperform the corresponding linear models reaching past 90% accuracy (*competence*: 94.3%; *autonomy*: 93.55%; *relatedness*: 92.6%; and *presence*: 91.4%). Compared to linear models, RBF SVMs appear to be more robust across all motivation factors (approx. 5% increase in accuracy). Unlike the poor performances of PL models based solely on play styles, models based on game metrics are very accurate and robust across all four factors.

D. Models Based on All Features

Even if under-performing on their own, play style information can enhance other PL models by adding domain-specific information. To test this assumption we pool play style profiles and game metrics together and train SVM models based on the combined input. Based on Fig. 2c, including play styles has only marginal improvement, if any. Linear Models perform 75% on average across all tests with comparable peak performance to the models based on game metrics with an average best accuracy over 85% (*competence*: 89.4%, *autonomy*: 84.33%, *relatedness*: 83.3%, and *presence*: 84.36%) for individual models. RBF SVMs show a marginal improvement over those which are based on game metrics alone with 79.88% average accuracy. The best models perform above 94% on average (*competence*: 94.3%; *autonomy*: 94.5%; *relatedness*: 92.3%; *presence*: 92.33%). Similar to models based on game metrics, models based on all features show very similar performance across all four factors.

E. Visualising the Motivation Models

To demonstrate the applicability of the models for game design, we re-examine the players’ features ordered by the predictions of each model. Using the predictions of the best SVM models, we create a global ordering of players on each of the four motivation factors and visualise the feature sets for the top 10 and the bottom 10 players (Fig. 3).

As evidenced by Figure 3, players with predicted high competence have more experience with the game and have

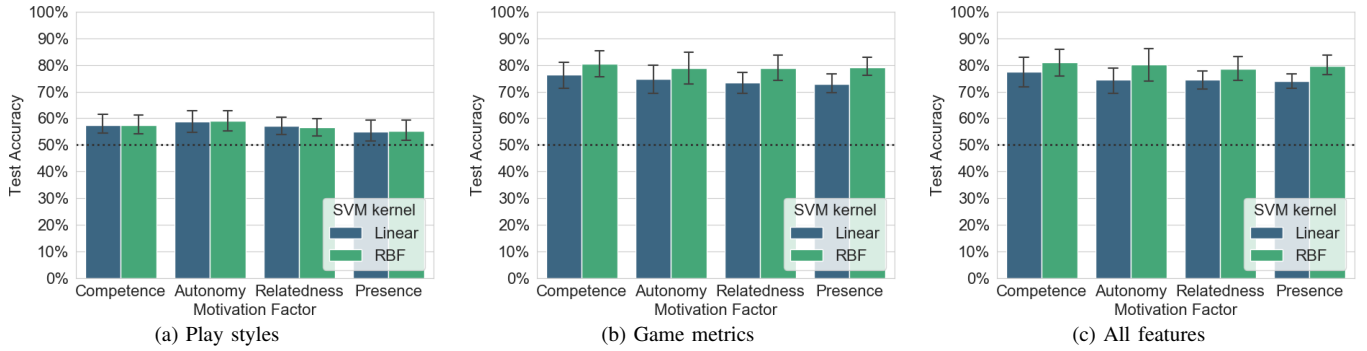


Fig. 2. Average accuracy of the best linear and RBF SVMs. The dotted line shows the baseline. The error bars represent 95% confidence intervals.

spent more time playing with various systems of the game in general. The high concentration of *Early Level 30* in the low competence group suggests that these players are rushing through the content. While there is disparity in terms of general playtime in the high competence group, high completion rates and/or *Dark Zone* and *Group Playtime* (coupled with no *Early Level 30*) indicate that *competence* aligns with a steady progression through the game.

In predicted *autonomy*, a similar effect can be observed in terms general playtime and completion. Additionally to high completion and progression rates, PvE and PvP success from an early level (Completed Incursions, and Dark Zone and Rogue Playtime) as well as DLC playtime also seem to be a deciding factor between high and low perceived *autonomy*.

In terms of *relatedness* the picture is more fuzzy. While both the top and bottom players share many similar features, the most prominent pattern is the high level of social interaction within the top group either as pro-social (e.g. *Group Playtime*) or antagonistic (e.g. *Playtime as Rogue*).

Finally, ordering by perceived *presence* shows a seemingly uniform distribution across different features with no discernible patterns between top and bottom players. Even so, the best models could still predict the higher ranking players in pairwise comparisons with 92.33% within given folds.

These observations largely match theories of SDT, further validating both the predictive models and the UPEQ questionnaire [12]. Across all factors, the most divergent features between top and bottom ranked players are associated with general playtime with an emphasis on mission completion, PvP activity, and playing with others in a group. Even in cases where no obvious linear relationship exists between individual features and motivation factors, non-linear PL techniques prove to be efficient methods for predicting motivation and offer an insightful qualitative tool for game design.

VI. DISCUSSION

This study applied preference learning via SVMs to model psychological constructs of the self determination theory and demonstrated the models' supreme efficiency and robustness in predicting players' motivation factors based solely on their in-game data. The SVM models can be applicable directly in a games user research context to predict factors of players' mo-




Fig. 3. Feature sets of the top and bottom 10 players ordered based on predicted motivation factors. Each column represents a player and each row is a different input feature. Darker cells show high values (normalised to each value range).

tivation based on aggregated data from multiple play sessions, which can be the subject of further qualitative analysis. In this study, we collected and processed the data of 298 players of *Tom Clancy's The Division* and defined motivation quantitatively using the UPEQ questionnaire. We converted the Likert scores to ordinal values using a *second-order* data processing approach [13] that yielded data corpora of over $4 \cdot 10^4$ samples, representing the differences between players' UPEQ Likert scores in pairs. The predicted pairwise relationships between players in terms of *competence*, *autonomy*, *relatedness* and *presence* with more than 80% accuracy on average, while the best models reach accuracies of at least 92%.

The presented research offers insights for game industry professionals and stakeholders, who aim to leverage and enhance the positive psychological effects of gameplay but also for researchers in affective computing and games user research. The paper offers a novel approach that utilises a machine (preference) learning approach to model psychological constructs such as motivation by treating the subjectively defined notions as an ordinal phenomenon. To the best of our knowledge, this is the first study which attempted to quantitatively model and predict constructs of SDT based on

player behaviour in a commercial-standard game and with that level of success.

A core limitation of this work is posed by the processing of the motivation factors (i.e., the ground truth). In particular, we opted for a simple experimental design and took the aggregated UPEQ scores at face value; aggregating Likert scale scores in this way, however, poses a number of methodological issues. The most important of these limitations is the processing of ordinal values as interval data [55]. As the distance between ordinal points is arbitrary, their mean value is not necessarily meaningful and introduces subjective reporting bias to the measurements [47]. To address this issue, future research should focus on new ways of transforming Likert-like data to preference relations. One possible approach would be to observe how an individual player's baseline shifts through the survey, scoring the four observed factors depending on significant changes in the responses.

Another limitation is posed by the nature of the dataset. As the dataset contains aggregated gameplay data, we cannot observe clearly how each player's motivation changes over time and thus the relationship between the player's "lifetime" and a one-time post-experience survey is rather broad. Although this type of data is easier to handle in a qualitative analysis, machine learning models are able to handle larger datasets with a higher dimensional feature space. Future research could focus on collecting a new dataset, recording multiple sessions and questionnaires per participant, and creating models which can predict temporal changes in the player's motivation.

It is important to note that the ordinal data processing method we propose in this paper yields large datasets which are generated through the pairwise comparison of all players. As such, the method is viable in current game development settings since it only requires a small sample of player experience annotations—such as those usually available in quality assurance departments of game studios. Although SVMs already showcased high levels of efficiency in predicting motivation on sets of several thousand datapoints, it can be argued that they might not be as robust when faced with much larger datasets of that type. In such instances alternative methods derived from *deep preference learning* [62], [63] are directly applicable to the modelling task.

VII. CONCLUSION

This paper attempted to find a computational mapping between gameplay data and aspects of player motivation, measured through the *Ubisoft Perceived Experience Questionnaire*. In this way, the *competence, autonomy, relatedness* and *presence* of almost 300 players of *Tom Clancy's The Division* was collected, as well as these players' in-game data. Experiments in this paper explored the degree to which such data can be a powerful predictor of players' survey responses. For that purpose, survey scores were converted to ordinal values and preference learning was applied to create efficient and robust support vector machine models of the four factors of motivation. The ordinal machine learning approach proved extremely successful in predicting such complex psychological

constructs by eliminating the reporting bias of the questionnaires. Core findings in this paper suggest that not only is it possible to infer the mapping between high-level gameplay metrics and survey-based annotations of complex emotional and cognitive states, but the inferred models have predictive capacities that reach certainty levels.

REFERENCES

- [1] S. Rigby and R. M. Ryan, *Glued to games: How video games draw us in and hold us spellbound*. Praeger, 2011.
- [2] G. N. Yannakakis and A. Paiva, "Emotion in games," *Handbook on affective computing*, pp. 459–471, 2014.
- [3] G. N. Yannakakis and J. Togelius, *Artificial Intelligence and Games*. Springer, 2018.
- [4] A. Canossa, J. B. Martínez, and J. Togelius, "Give me a reason to dig minecraft and psychology of motivation," *Proceedings of the Conference on Computational Intelligence and Games (CIG)*, vol. 2013, pp. 1–8, 2013.
- [5] A. Canossa, J. B. Badler, S. Tignor, and R. C. Colvin, "In your face(t) impact of personality and context on gameplay behavior," in *Proceedings of the International Conference on the Foundations of Digital Games (FDG)*, 2015.
- [6] D. Melhart, "Towards a comprehensive model of mediating frustration in videogames," *Game Studies*, vol. 18, no. 1, 2018.
- [7] A. Drachen, A. Canossa, and G. N. Yannakakis, "Player modeling using self-organization in Tomb Raider: Underworld," in *Proceedings of the Symposium on Computational Intelligence and Games (CIG)*. IEEE, 2009, pp. 1–8.
- [8] S. Makarovych, A. Canossa, J. Togelius, and A. Drachen, "Like a dna string: Sequence-based player profiling in tom clancys the division," in *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment Conference*. York, 2018.
- [9] M. S. El-Nasr, A. Drachen, and A. Canossa, *Game analytics*. Springer, 2016.
- [10] G. N. Yannakakis, P. Spronck, D. Loiacono, and E. André, "Player modeling," in *Dagstuhl Follow-Ups*, vol. 6. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2013.
- [11] R. M. Ryan and E. L. Deci, "Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being," *American Psychologist*, vol. 55, no. 1, p. 68, 2000.
- [12] A. Azadvar and A. Canossa, "Upeq: ubisoft perceived experience questionnaire: a self-determination evaluation tool for video games," in *Proceedings of the International Conference on the Foundations of Digital Games (FDG)*. ACM, 2018.
- [13] G. N. Yannakakis, R. Cowie, and C. Busso, "The ordinal nature of emotions: An emerging approach," *IEEE Transactions on Affective Computing*, 2018.
- [14] J. Fürnkranz and E. Hüllermeier, "Preference learning," in *Encyclopedia of Machine Learning*. Springer, 2011, pp. 789–795.
- [15] G. N. Yannakakis and J. Togelius, "Experience-driven procedural content generation," *IEEE Transactions on Affective Computing*, vol. 2, no. 3, pp. 147–161, 2011.
- [16] N. Shaker, G. N. Yannakakis, and J. Togelius, "Towards automatic personalized content generation for platform games," in *Proceedings of the Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 2010.
- [17] C. Holmgård, A. Liapis, J. Togelius, and G. N. Yannakakis, "Evolving personas for player decision modeling," in *Proceedings of the Conference on Computational Intelligence and Games (CIG)*, 2014.
- [18] C. Bauckhage, A. Drachen, and R. Sifa, "Clustering game behavior data," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 7, no. 3, pp. 266–278, 2015.
- [19] A. Drachen, R. Sifa, C. Bauckhage, and C. Thurau, "Guns, swords and data: Clustering of player behavior in computer games in the wild," in *Proceedings of the Conference on Computational Intelligence and Games (CIG)*, 2012, pp. 163–170.
- [20] H. P. Martínez and G. N. Yannakakis, "Mining multimodal sequential patterns: a case study on affect detection," in *Proceedings of the International Conference on Multimodal Interfaces*. ACM, 2011, pp. 3–10.

- [21] G. Wallner, "Sequential analysis of player behavior," in *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*. ACM, 2015, pp. 349–358.
- [22] T. Mahlmann, A. Drachen, J. Togelius, A. Canossa, and G. N. Yannakakis, "Predicting player behavior in Tomb Raider: Underworld," in *Proceedings of the Symposium on Computational Intelligence and Games (CIG)*. IEEE, 2010, pp. 178–185.
- [23] J. Runge, P. Gao, F. Garcin, and B. Faltings, "Churn prediction for high-value players in casual social games," in *Proceedings of the Conference on Computational Intelligence and Games (CIG)*, 2014, pp. 1–8.
- [24] Á. Periañez, A. Saas, A. Guitart, and C. Magne, "Churn prediction in mobile social games: towards a complete assessment using survival ensembles," in *Proceedings of the International Conference on Data Science and Advanced Analytics (DSAA)*, 2016, pp. 564–573.
- [25] M. Viljanen, A. Airola, J. Heikkonen, and T. Pahikkala, "Playtime measurement with survival analysis," *IEEE Transactions on Games*, vol. 10, no. 2, pp. 128–138, 2018.
- [26] S. C. Bakkes, P. H. Spronck, and G. van Lankveld, "Player behavioural modelling for video games," *Entertainment Computing*, vol. 3, no. 3, pp. 71–79, 2012.
- [27] H. P. Martínez, M. Garbarino, and G. N. Yannakakis, "Generic physiological features as predictors of player experience," in *Proceedings of the International Conference on Affective Computing and Intelligent Interaction (ACII)*. Springer, 2011, pp. 267–276.
- [28] V. Georges, F. Courtemanche, M. Fredette, P.-M. Lger, and S. Sncal, "Developing personas based on physiological measures," in *Proceedings of the International Conference on Physiological Computing Systems*, 2018.
- [29] N. Shaker, S. Asteriadis, G. N. Yannakakis, and K. Karpouzis, "Fusing visual and behavioral cues for modeling user experience in games," *IEEE Transactions on Cybernetics*, vol. 43, no. 6, pp. 1519–1531, 2013.
- [30] E. Camilleri, G. N. Yannakakis, and A. Liapis, "Towards general models of player affect," in *Proceedings of the International Conference on Affective Computing and Intelligent Interaction (ACII)*, 2017.
- [31] Z. Borbora, J. Srivastava, K.-W. Hsu, and D. Williams, "Churn prediction in mmorpgs using player motivation theories and an ensemble approach," in *Proceedings of the International Conference on Privacy, Security, Risk and Trust (PASSAT) & Proceedings of the International Conference on Social Computing (SocialCom)*, 2011.
- [32] N. Yee, "Motivations for play in online games," *CyberPsychology & Behavior*, vol. 9, no. 6, pp. 772–775, 2006.
- [33] K. J. Shim, K.-W. Hsu, and J. Srivastava, "An exploratory study of player performance, motivation, and enjoyment in massively multiplayer online role-playing games," in *Proceedings of the International Conference on Privacy, Security, Risk and Trust (PASSAT) & Proceedings of the International Conference on Social Computing (SocialCom)*, 2011.
- [34] M. V. Birk, D. Toker, R. L. Mandryk, and C. Conati, "Modeling motivation in a social network game using player-centric traits and personality traits," in *Proceedings of the International Conference on User Modeling, Adaptation, and Personalization*. Springer, 2015, pp. 18–30.
- [35] S. Reiss and S. M. Havercamp, "Toward a comprehensive assessment of fundamental motivation: Factor structure of the reiss profiles," *Psychological Assessment*, vol. 10, no. 2, p. 97, 1998.
- [36] E. Deci and R. M. Ryan, *Intrinsic motivation and self-determination in human behavior*. Springer Science & Business Media, 1985.
- [37] R. M. Ryan, C. S. Rigby, and A. Przybylski, "The motivational pull of video games: A self-determination theory approach," *Motivation and Emotion*, vol. 30, no. 4, pp. 344–360, 2006.
- [38] A. K. Przybylski, C. S. Rigby, and R. M. Ryan, "A motivational model of video game engagement," *Review of General Psychology*, vol. 14, no. 2, p. 154, 2010.
- [39] R. M. Ryan and J. P. Connell, "Perceived locus of causality and internalization: Examining reasons for acting in two domains," *Journal of Personality and Social Psychology*, vol. 57, no. 5, p. 749, 1989.
- [40] E. L. Deci, R. Koestner, and R. M. Ryan, "A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation," *Psychological Bulletin*, vol. 125, no. 6, p. 627, 1999.
- [41] M. Lombard and T. Ditton, "At the heart of it all: The concept of presence," *Journal of Computer-Mediated Communication*, vol. 3, no. 2, 1997.
- [42] G. Calleja, *In-game: From immersion to incorporation*. MIT Press, 2011.
- [43] R. M. Ryan and E. L. Deci, "Intrinsic and extrinsic motivations: Classic definitions and new directions," *Contemporary Educational Psychology*, vol. 25, no. 1, pp. 54–67, 2000.
- [44] J. H. Brockmyer, C. M. Fox, K. A. Curtiss, E. McBroom, K. M. Burkhardt, and J. N. Pidruzny, "The development of the game engagement questionnaire: A measure of engagement in video game-playing," *Journal of Experimental Social Psychology*, vol. 45, no. 4, pp. 624–634, 2009.
- [45] L. E. Nacke, C. Bateman, and R. L. Mandryk, "Brainhex: A neurobiological gamer typology survey," *Entertainment Computing*, vol. 5, no. 1, pp. 55–62, 2014.
- [46] B. Chen, M. Vansteenkiste, W. Beyers, L. Boone, E. L. Deci, J. Van der Kaap-Deeder, B. Duriez, W. Lens, L. Matos, A. Mouratidis *et al.*, "Basic psychological need satisfaction, need frustration, and need strength across four cultures," *Motivation and Emotion*, vol. 39, no. 2, pp. 216–236, 2015.
- [47] G. N. Yannakakis, R. Cowie, and C. Busso, "The ordinal nature of emotions," in *Proceedings of the International Conference on Affective Computing and Intelligent Interaction (ACII)*, 2017.
- [48] A. R. Damasio, *Descartes error: Emotion, rationality and the human brain*. New York: Putnam, 1994.
- [49] B. Seymour and S. M. McClure, "Anchors, scales and the relative coding of value in the brain," *Current Opinion in Neurobiology*, vol. 18, no. 2, pp. 173–178, 2008.
- [50] H. Helson, *Adaptation-level theory: an experimental and systematic approach to behavior*. Harper and Row: New York, 1964.
- [51] R. L. Solomon and J. D. Corbit, "An opponent-process theory of motivation: I. temporal dynamics of affect," *Psychological Review*, vol. 81, no. 2, p. 119, 1974.
- [52] S. Erk, M. Kiefer, J. Grothe, A. P. Wunderlich, M. Spitzer, and H. Walter, "Emotional context modulates subsequent memory effect," *Neuroimage*, vol. 18, no. 2, pp. 439–447, 2003.
- [53] H. Martinez, G. Yannakakis, and J. Hallam, "Don't classify ratings of affect; rank them!" *IEEE transactions on Affective Computing*, no. 1, 2014.
- [54] G. N. Yannakakis and H. P. Martinez, "Grounding truth via ordinal annotation," in *Proceedings of the International Conference on Affective Computing and Intelligent Interaction (ACII)*, 2015.
- [55] G. N. Yannakakis and H. P. Martínez, "Ratings are overrated!" *Frontiers in ICT*, vol. 2, p. 13, 2015.
- [56] D. Melhart, K. Sfikas, G. Giannakakis, G. N. Yannakakis, and A. Liapis, "A motivational model of video game engagement," *Proceedings of the IJCAI workshop on AI and Affective Computing*, vol. 86, pp. 26–33.
- [57] J. Fürnkranz and E. Hüllermeier, "Pairwise preference learning and ranking," in *Proceedings of the European Conference on Machine Learning*. Springer, 2003, pp. 145–156.
- [58] T. Joachims, "Optimizing search engines using clickthrough data," in *Proceedings of the SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2002, pp. 133–142.
- [59] V. E. Farrugia, H. P. Martínez, and G. N. Yannakakis, "The preference learning toolbox," *arXiv preprint arXiv:1506.01709*, 2015.
- [60] C.-C. Chang and C.-J. Lin, "LIBSVM: a library for support vector machines," *Transactions on Intelligent Systems and Technology (TIST)*, vol. 2, no. 3, p. 27, 2011.
- [61] V. Vapnik, "Constructing learning algorithms," *The Nature of Statistical Learning Theory*, pp. 119–157, 1995.
- [62] H. P. Martínez and G. N. Yannakakis, "Deep multimodal fusion: Combining discrete events and continuous signals," in *Proceedings of the International Conference on Multimodal Interaction*. ACM, 2014.
- [63] H. P. Martinez, Y. Bengio, and G. N. Yannakakis, "Learning deep physiological models of affect," *Computational Intelligence Magazine*, vol. 8, no. 2, pp. 20–33, 2013.