From Theory to Behaviour: Towards a General Model of Engagement

V. Bonometti†‡, C. Ringer†, M. Ruiz‡, A.R. Wade§, A. Drachen†
†Department of Computer Science, University of York, York, UK YO10 5DD
‡Square Enix Limited, London, UK SE1 8NW
§Department of Psychology, University of York, York, UK YO10 5DD
Email: {vb690, cr1116, anders.drachen, alex.wade}@york.ac.uk; analyticsl@eu.square-enix.com

***Abstract*—Engagement is a fuzzy concept. In the present work we operationalize engagement mechanistically by linking it directly to human behaviour and show that the construct of engagement can be used for shaping and interpreting data-driven methods. First we outline a formal framework for engagement modelling. Second we expanded on our previous work on theory-inspired data-driven approaches to better model the engagement process by proposing a new modelling technique, the Melchoir Model. Third, we illustrate how, through model comparison and inspection, we can link machine-learned models and underlying theoretical frameworks. Finally we discuss our results in light of a theory-driven hypothesis and highlight potential application of our work in industry.** [1]

***Index Terms*—Engagement Modelling, Player Modelling, Machine Learning, Artificial Neural Networks**

I. INTRODUCTION

Various attempts have been made in prior literature across several domains to describe and analyse the engagement process in a video-game context [1]. Engagement is an extremely useful construct because it allows us to understand how a user (e.g. a player) is interacting with a specific object or activity (i.e. a video game) and how this interaction can evolve over time. Here, we build upon our previous work [2] modeling engagement and strengthen its link with human behaviour. We expand our previous theoretical framework and use it to refine our modelling approach. In doing so we generate a set of theory-driven hypothesis that we investigate through model comparison and inspection. This was achieved by carrying out a series of three analyses over a 3.2 millions-entries cross-game data-set. Our goal was not just to evaluate the effectiveness of our new approach but also to ask how the assumptions generated by the theoretical framework compare to those learned by our modelling approach. In the final section we briefly discuss a series of practical application that our current approach can have in industry settings, ranging from estimation of engagement evolution and identification of engagement profiles. As far as we know this is one of the first works explicitly translating qualitative theories of engagement into machine-learned models. We argue that our refined theoretical framework allows for a re-framing of engagement in behavioural terms while also allowing us to formulate and test more precise hypotheses. The modelling approach we propose here extends and refines our previous model - maintaining the same advantages while also improving on some of its limitations. In particular, we explicitly consider the contribution of environmental variables, and we fully exploit the temporal nature of the engagement process allowing the model to work continuously with an arbitrarily long sequence of input features. Finally, the model is able to estimate a larger range of target metrics when compared with its predecessor. This work also presents a methodology for inspecting, analyzing and interpreting the user representation learned by the model, providing a generalizable way of producing understandable engagement profiles.

This work was supported by the EPSRC Centre for Doctoral Training in Intelligent Games & Games Intelligence (IGGI) [EP/L015846/1] and the Digital Creativity Labs (digitalcreativity.ac.uk), jointly funded by EPSRC/AHRC/Innovate UK under grant no. EP/M023265/1.

[1]Extensive ancillary information and results can be found at https://github.com/vb690/modelling_engagement_ammount/wiki.

II. STATE OF THE ART AND CONTRIBUTION

In this section we give a general overview on the state of the art in engagement modelling. Due to space constraints and the substantial literature on the matter we will focus on a restricted set of representative works in the area of large-scale behavioural modelling of engagement. The work on engagement modelling comes, generally, in two forms: estimation and profiling of in-game behaviour [3]. The estimation of in-game behaviour is usually formulated as a supervised machine-learning or more general statistical modelling problem [3]. Despite the literature on the topic often presenting compelling solutions for practical problems, it tends to follow a black box approach: a machine-learned solution is generated for solving a specific task but no attempts are made to inspect or interpret the model [4]–[7]. Moreover, when these attempts are made, the lack of a solid and predefined theoretical framework tends to lead to post-hoc interpretations which are sometimes difficult to verify or relate with actual human behaviour [8], [9]. When trying to estimate engagement profiles, the approach widely used in the literature is to adopt some form of unsupervised learning technique for individuating patterns of interaction with various in-game features [3], [9]. This however is usually done considering an unconstrained set of game-specific metrics. As a result, *a-posteriori* justifications for the characteristics of the individuated profiles are provided [10]–[12], which, without an overarching explanatory theoretical

framework, appear to be be very context-specific and difficult to interpret. What we see in the literature is that attempts are made to model a single behavioural manifestation of engagement rather than the construct in its entirety. A noticeable exception in this regard is the recent work by Reguera et al. [13], who adopt a complete data-driven approach managed to derive a general law for describing and quantifying the engagement process, similarly to what Bauckhage at all. did in [14]. However, neither group interpret their findings through the lens of existing human behaviour theory. We believe that a holistic model of engagement can be generated, constraining the great flexibility provided by data-driven approaches by employing solid and well established theoretical priors [15]. To do so, an *a-priori* theoretical framework which is guaranteed to generalise to different situations should be defined. Such a framework should clearly state what are the observable and measurable indicators of engagement and how they are expected to vary in relations with the construct's dynamics. In doing so the findings emerging from data-driven approaches can be compared with what the theoretical framework prescribes.

III. ENGAGEMENT AS A BEHAVIOURAL PROCESS

Although various attempts have been made to describe the construct of engagement, prior literature struggles to provide a formal definition [1]. In particular, we observed a lack of mechanistic explanations and parallelisms with human behaviour in favour of more holistic and phenomenological descriptions of the concept [1]. Although these phenomenological descriptions do provide qualitative insights, mechanistic explanations are necessary to generate and verify hypotheses as well as producing real-world applications for engagement modelling. For this reason, we extend the theoretical framework adopted in our previous work [2] to draw stronger connections between engagement and human behaviour. This allowed us to not only make more informed decisions when designing a strategy for modelling engagement in behavioural terms but also to use this approach for verifying a set of theory driven hypotheses.

A. The Engagement Process Model

When looking at the various formulation of engagement [1] a common denominator seems to emerge: from a behavioural point of view engagement can be represented with the amount, duration and frequency of interactions between an individual I and an object O. The Engagement Process Model proposed by O'Brien and Toms [16] perfectly summarizes this, describing these interactions in terms of a dynamic system. In their system, the ability of O to provide rewarding experiences to I, in conjunction with environmental factors Env, controls cycles of interaction between I and O. Figure 1 represents the process by which I engages with O inside Env. If I has an *a-priori* belief that O is able to provide rewarding experiences they will direct their attention towards it (1). If Env does not pose any constraints, I will interact with O for as long as O is able to provide rewarding experiences (2).

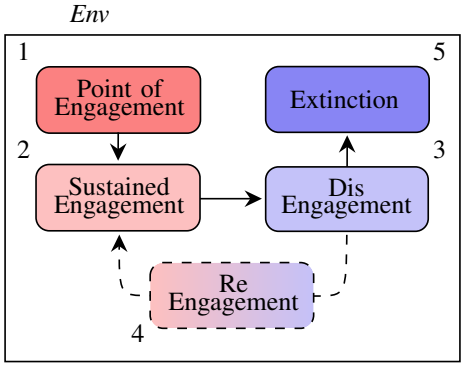

Fig. 1: **The Engagement Process Model.** Solid and dashed lines represent compulsory and optional paths.

However, if O fails to provide these experiences, or constraints from Env emerge, I will gradually begin to disengage from O (3). At this point I can either enter a cycle of re-engagement and disengagement (4) or reach an inevitable state of complete withdrawal from O (5).

B. Incentive Salience Attribution

The Engagement Process Model provides a good high-level description of the engagement process but fails to give a clear explanation of the dynamics controlling the system or explicitly draw connections to human behaviour. We believe that the Incentive Salience hypothesis formulated by Berridge and Robinson [17] offers a solution for both of these problems.

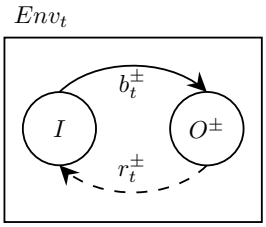

Fig. 2: **The process of incentive salience attribution**. Solid and dashed lines represent observable and latent variables.

From a behavioural point of view, the Incentive Salience hypothesis states that through repeated interactions (i.e. operant conditioning), objects, O, which elicit rewarding experiences, r, become valuable, i.e. they acquire salience, to the individual, I, interacting with them [17], [18]. The amount of salience then controls how likely and intense future interactions between I and O will be [17], [18]. Let $(b_t)_{t=1}^n$ represent the history of interactions between I and O, $(r_t)_{t=1}^n$ a measure of how rewarding an interaction O is perceived to be by I and $(Env_t)_{t=1}^n$ the changes in the environment in which I and O interact. Following Figure 2, we can imagine the intensity of each b_{t+1} increasing and decreasing according to the level of r_t. Changes in r_t alter the salience of O, which acts as an attracting force for future b. Reformulating this in a videogame context, I and O indicate a user and a specific videogame respectively, b_t is the amount of playing activity during a gaming session t, r_t stands for how rewarding the user found the gaming session. Finally Env are all those external factors interfering or promoting the gaming activity (e.g. school days or holidays) which therefore need to be taken into account for estimating an unbiased measure of r_t. For convenience, from

now on I and O will be used as synonyms for representing a user and a specific game the user is interacting with. The concept of attributed salience will be used interchangeably with that of level of engagement since the two are strongly interconnected and do not differ from a behavioural point of view: high levels of attributed salience pushes I to interact (i.e. engage) more with O [1], [16], [17]. By keeping this notation we also highlight how our current approach can in theory be extended to contexts other than video games.

C. From Theory to Modelling

We believe it is of core importance to estimate the level of salience I is attributing to O at a specific t. As we have seen before this, other than simply providing a way to assess the current state of the interaction between I and O, would allow us to perform informed estimations of all $b_{t+1:n}$. Following our theoretical framework we hypothesise that, being f, g and k a set of unknown arbitrarily complex functions:

$$b_{t+1:n} \sim f(O^{\pm}) \quad \text{(Eqn 1)}$$

This can be achieved by taking into consideration the full history of rewarding experiences provided to I by O:

$$O^{\pm} \sim g(r_t, r_{t-1}, \ldots, r_{t-n}) \quad \text{(Eqn 2)}$$

Because r_t is a latent variable it is necessary to find an observable outcome able to approximate it. Again, following our theoretical framework we hypothesize that r_t can be inferred from the full history of observed b between I and O weighted by the effect of Env

$$r_t \sim k((b \cdot Env)_t, (b \cdot Env)_{t-1}, \ldots, (b \cdot Env)_{t-n}) \quad \text{(Eqn 3)}$$

To summarize, according to our hypothesis the latent variable r_t (which control changes in the salience attributed to O) can be inferred from the intensity of observed behaviours produced by I when interacting with O.

D. Choosing the Right Modelling Approach

Following the formulations above we can see how a suitable approach for estimating attributed salience might be to adopt an auto-regressive-like (AR) model with order $p = n$. This model would take as input sequences of b, O and Env and attempt to estimate the intensity of all future b generated by I in response to O. The model should represent the temporal relationships between the elements in the input sequence. We showed in our previous work [2] that simply considering full history of past behaviours without explicitly modelling temporality lead to sub-optimal results. Finally, since our theoretical framework doesn't explicitly indicate which type of function best describes the relationship between the history of interactions and the intensity of future interactions, it is necessary for our modeling approach to be able to learn arbitrary complex functions. We argue therefore that the use of Artificial Neural Network (ANN), and in particular recurrent variants (RNN), are particularly well suited modelling techniques.

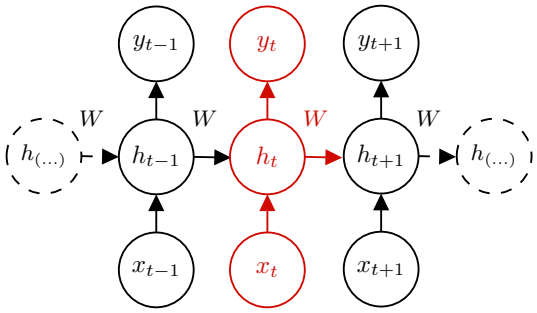

Fig. 3: **Many To Many Recurrent Neural Network**. Adapted from [19].

When using a RNN for sequence to sequence tasks, the neural network takes as input a sequence of vectors of arbitrary length $(x_t)_{t=1}^{n}$ and as target $(y_t)_{t=1}^{n}$, a sequence of vectors of the same length. A weight matrix W, with subsequent non-linear activation functions, is then applied recurrently to each x_t and to a latent variable h_{t-1}. This will produce the next h_t that is used for estimating y_t as well as for the subsequent operations. Figure 3 demonstrates an example of a 'many to many' model. We can see how an RNN satisfies the need for modelling the inputs temporally and in an arbitrarily complex manner. Moreover unlike traditional AR models it is not necessary to specify the order p, this is something W dynamically infers from the data, which allows to retain information from varying size sequences in a flexible manner. Following Figure 3, the input x_t is a vector carrying information about b_t, O and Env_t while y_t a vector of variables indicative of the intensity of $b_{t+1:n}$. The latent variable h_t would represent the level of salience I is attributing to O at time t. r_t is implicitly computed by the operations performed by W and it reflects in the changes of h_t. Given how an RNN computes its latent variables, we can see that at each t the associated h will represent the full history of information that have flowed through the model until then.

E. Manifold Learning

Despite the fact that ANNs seem to be a suitable approach for translating our theoretical framework in a machine-learned model, we need to assess if the learned model is able to represent a sensible approximation of the process of salience attribution. We believe that the concept of manifold learning can be used to demonstrate this. ANNs are known to be efficient universal function approximators but it is exceptionally hard to draw insights about the underlying learned function. However, one of the core concepts in deep learning, and machine learning in general, is that the data we observe lies on a manifold: a connected region where each point is surrounded by other extremely similar examples [19]. When training an ANN, we can imagine the operations performed by each layer as learning the coordinates of each input point on a manifold that holds a representation (i.e. an embedding) that is useful for subsequent layers. For example, in a supervised learning context, the last layer of an ANN is usually tasked with performing classification or regression while the layer before

that provides the best data transformation for that task. With this in mind we argue that if, through architectural choice, we can enforce an ANN to learn a representation z for which

$$b_{t+1:n} \sim f(z) \quad \text{(Eqn 4)}$$

where f is an arbitrarily complex function. This representation z should then be able to place individuals with similar characteristics, with respect to the objective, closer in the embedding space. With this in mind we argue that through the analysis of z it should be possible to inspect these characteristics and compare them with those predicted by the theoretical framework used for designing the model.

IV. METHODOLOGY

A. Data

To conduct our experiments, we gathered data from six games published by our partner company, *Square Enix Ltd.*: *Hitman Go* (hmg), *Hitman Sniper* (hms), *Just Cause 3* (jc3), *Just Cause 4* (jc4), *Life is Strange* (lis), and *Life is Strange: Before the Storm* (lisbf). The data-set contained data from 3,240,000 individuals, evenly distributed across the 6 games, and randomly sampled from all users who played between the games release and January 2020. Different from our previous work, [2], we included a series of metrics representing the Env as well as increased the number of target metrics to provide a better behavioural approximation of engagement. Moreover, in this work no set observation period was required since the modelling approach follows an online strategy: given the history of game sessions for a user the model is trained to perform estimation after each session. It was important, when deciding on which metrics should be adopted, to have a minimal and highly generalizable set of features. This feature selection improved the generalizability and usability of our methodology and allowed us to carry out analytical work on the representation learned by our model: relying on a limited set of of input metrics, selected on the basis of an underlying theoretical framework, makes it easier to generate and interpret human readable visualisations.

*1) **The behavioural and object metric**:* Given a set of $(b_t)_{t=1}^n$ between I and O (i.e. game sessions performed by a user within a specific game context), we needed a set of behavioural metrics able to represents the intensity of each b. Following the indications provided by our theoretical framework, we decided to employ the same behavioural metrics used in our previous work, see Table I, as they appear to be optimal candidates for representing the intensity of b in behavioural terms. For representing the O generating $(b_t)_{t=1}^n$ we simply retrieved a metric indicating the game context to which the behavioural metric are associated.

*2) **The environment metrics**:* To represent $(Env_t)_{t=1}^n$ we needed a set of metrics which provide a high-level description of the environment in which I and O are interacting. We chose the hour of the day, the day of the week and the day of the year to account for elements like weekends, working hours and bank holidays. We additionally consider a metric indicating the user's broad geographic area to account for regional variations.

TABLE I: **Considered Metrics over Sessions**. Taken from [2]

Metric	Description
Session Time	Overall session duration (minutes)
Play Time	Session Time spent actively playing (minutes)
Delta Session	Temporal distance between sessions (minutes)
Activity Index	Count of user initiated game-play-related actions. E.g. 'Talk to NPC' or 'Acquire Upgrade' were considered valid actions while 'Click Menu' or 'NPC Attacks You' were not.
Activity Diversity	Count of unique voluntarily initiated actions

*3) **The target metrics**:* The target metrics need to be a set of behavioural metrics summarising the intensity of all the $(b_t)_{t+1}^n$ after a specific b_t. These, following the intuitions from [17], can be used as a measure of attributed salience. We extend the traditional metrics of churn probability (ch) and survival time (st) to also include survival sessions (st) and the time the player is absent from the game between the end of one session and the start of the next (ab). The two survival target metrics were calculated following the formula:

$$survival_t = total\ Pt\ or\ Ps - \sum_{n=1}^{t} Pt\ or\ Pn_n \quad \text{(Eqn 5)}$$

with Pt and Ps being respectively $Played\ Time$ and $Played\ Sessions$. The absence metric is simply the time in minutes between the current session and the previous session. We choose a definition and encoding of the churn variable which is robust to outliers as well being able to represent uncertain cases. Given the two criteria:

(a) Completing the game
(b) Being inactive for a period equal or greater to:

$$inactivity = Q_3(x) + 1.5 \cdot IQR(x) \quad \text{(Eqn 6)}$$

where x is a vector of of inter-sessions distances for a specific game, we determined the probability of being in a churning state as follow:

$$churn\ probability = \begin{cases} 0.0 & if\ a \\ 1.0 & if\ \neg a \wedge b \\ 0.5 & otherwise \end{cases} \quad (7)$$

Summarizing, we argue that the intensity of all the future interactions between I and O can be expressed through the combinations of the four aforementioned metrics:

$$b_{t+1:n} \sim (ch, st_t, ss_t, ab_t) \quad \text{(Eqn 8)}$$

because all the four metrics are good quantifiers of the amount and frequency of future behaviour generated by I in response to O.

*4) **Data Preparation**:* For each user in our dataset we retrieve a single feature for the object metric and a temporal series for each other metrics, computing the input and target metrics for each recorded game session. We then split the data-set into training and testing subsets (80 and 20 % of the original dataset) and rescaled the behavioural metrics using the following formula:

$$scale(x) = \frac{x - \min(x)}{\max(x) - \min(x)} \quad \text{(Eqn 9)}$$

where x is the feature vector to be re-scaled. To avoid the risk of information leakage $\min$ and $\max$ are computed only on the training set. We then proceeded to numerically encode the variables indicative of O and Env because the categorical encoding employed by our models requires to transform all the unique categories in numerical indices [20].

B. Hypotheses and Models

To evaluate the assumptions of our theoretical framework we defined a series of hypotheses which we tested both through the comparison of a set of models as well as by inspecting the embedding space learned by our proposed approach. These hypotheses are: 1) Explicitly modelling temporality in the interactions between I and O results in better performance when estimating behavioural proxies of attributed salience. 2) Models able to learn arbitrarily complex functions will produce less error when estimating behavioural proxies of attributed salience. 3) Our modelling methodology is able to generate a latent representation which reasonably approximates the level of salience an individual has attributed to a game. 4) The encoding of the level of salience provided by the model will reflect increases and decreases in the behavioural metrics used for describing the strength of the interactions between I and O. For instance, individuals encoded has having had attributed high salience to O will show a history of more frequent and longer interactions. 5) While the patterns above will appear consistently across game contexts (i.e. the distinct O) various profiles will be observable within and between different O, because the behavioural manifestation of salience attribution partially depends on the nature of O. To test these hypotheses we designed and implemented three models with *ad-hoc* characteristics.

1) ***Autoregressive-like Models****:* Two AR models of order $p = 1$ were implemented for testing hypotheses 1 and 2. More specifically, we used an ElasticNet regression (TD ENet):

$$b_{t+1:n} \sim b_t + Env_t + O \qquad \text{(Eqn 10)}$$

and a Multi Layer Perceptron (TD MLP)

$$b_{t+1:n} \sim f(b_t, Env_t, O) \qquad \text{(Eqn 11)}$$

We can see that both models work under a Markovian assumption that intensity of future interactions with O is reliant only on the current state of I, but TD MLP is not constrained to learn only linear functions.

2) ***Melchior Model****:* Additionally, we propose the novel 'Melchior Model' (MM), which is specifically designed to implement the insights from our theoretical framework and to test all the aforementioned hypotheses, Figure 4. Our MM architecture models the contribution of each component, O, b, and Env, separately in a way which resembles the specifications of section III.D, each component is then pooled and used for estimating metrics representative of the intensity of future interactions between I and O. The guiding concept was multitask learning [19]: given a set of targets and the assumption that these share a common representation, explicitly modelling this last one allows to capture a collection of common factors shared among all targets. We hypothesised that this would have provided two benefits. Firstly, better generalization and superior performance in terms of goodness of fit. Secondly, the ability to inspect an overarching representation able to explain the variance in the target metrics. This representation would constitute an approximation of the concept of attributed salience expressed in section III.B.

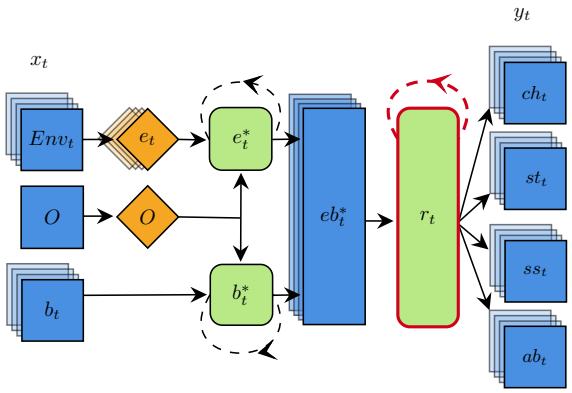

Fig. 4: **Melchior Model Core Architecture.** Blue squares represents inputs, outputs and fully connected operations. Yellow diamonds indicates categorical embedding operations. Green rounded rectangles are recurrent operations. Shapes with shaded copies represents time series. The red highlight indicates, according to what we specified in Eqn 4, the portion of the model that we hypothesize holding a representation of the salience I is attributing to O.

All three models take as inputs two series of vectors, $(b_t)_{t=1}^n$ and $(Env_t)_{t=1}^n$, and a single feature O and are trained to output four series of vectors corresponding to the target metrics. They were trained using Binary Cross Entropy for churn probability and the Symmetric Mean Absolute Percentage Error (SMAPE) for the other metrics. Both loss functions are bounded between 0 and 1 and lower values correspond to better performance. These loss functions were also used for computing the performance of the models on the test set[2].

C. Experiments

1) ***Experiment 1 - Model Comparison****:* The first experiment is designed to replicate the findings of our previous work [2] as well as test the assumption made by our theoretical framework that engagement needs to be modelled as a dynamic system. To achieve this we implemented the three models discussed above and compared their performance. We employ 20% of the training set to search for the best hyper-parameters while keeping the architecture of the three models fixed (i.e. only elements such as the number of layers and hidden units were tuned). This was achieved using the Hyperband algorithm [21] due to its capacity to converge to a good solution with relatively limited computational resources. After the tuning process we fit each model on the entirety of the training set and then compute the evaluation metrics on the test set.

2) ***Experiment 2 - Embedding Visualization****:* The second experiment aims to verify, through visual inspection, that MM learns a separate representation for each O that the model observes and that these representations reasonably encode the

different level of salience that various I attribute to O. To extract these representations we simply truncate the model's weights up to the point highlighted in Figure 4 and then transform the testing set through the usual forward pass. For generating the visualization necessary to test our hypothesis we needed a methodology able to represent in two dimensions the type of spatial properties discussed in section III.E For this reason we employed the Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP) technique [22]. UMAP is a manifold learning technique able to represent the global and local geometry of a set of data in a arbitrary number of dimensions, often 2 or 3 for visualization purposes. In general terms this means that, on a two dimensional plane, points are placed closer or further away to each other depending on their similarity (i.e. proximity) in the original space.

*3) **Experiment 3 - Embedding Partitioning**:* The final experiment analyzes the applicability of using the representation from experiment two to validate the hypothesis that behavioural traces that are coherent with the incentive salience framework can be extracted from the learned embedding. Following the manifold hypothesis we know that the learned representation should place points similar to each others close in the multidimensional space. Therefore partitioning this space, with distance-based techniques, should allow us to individuate different groups encoded with different levels of attributed salience. To achieve this we run a mini-batch k-means algorithm [23], [24] on the representation learned by the model for each O and then inspect the behavioural metrics associated to each partition. Since k-means requires a specific number of partitions, we used the elbow method for choosing the number of partitions that maximize the marginal gain in explained variance [2]. We decided to use the k-means algorithm despite its known limitations [24] because of its ability to scale well when $N \geq 10^5$. After performing the partitioning we visualized the temporal traces of the behavioural input for each partition and compared them with the distribution of target metrics for each partition. The code for all the experiments has been written in Python 3.6 relying on Keras [20] for the implementation of the models and on Scikit-learn [24] for the partitioning algorithm.

V. RESULTS

*1) **Experiment 1 - Model Comparison**:* Figure 5 shows the performance of each model collapsed over game context and play sessions. Each plot representing the performance for a specific objective. We can see how the performance of the three models follows the expected trend, TD ENet $<$ TD MLP $<$ Melchoir, in all four targets metrics. Additonally, performing this comparison on non-collapsed data [2] demonstrates a near-identical trend, strengthening the assumption that temporally modelling interactions is vital for modelling salience attribution.

[2]Omitted due to space concerns. Can be found at https://github.com/vb690/modelling_engagement_ammount/wiki

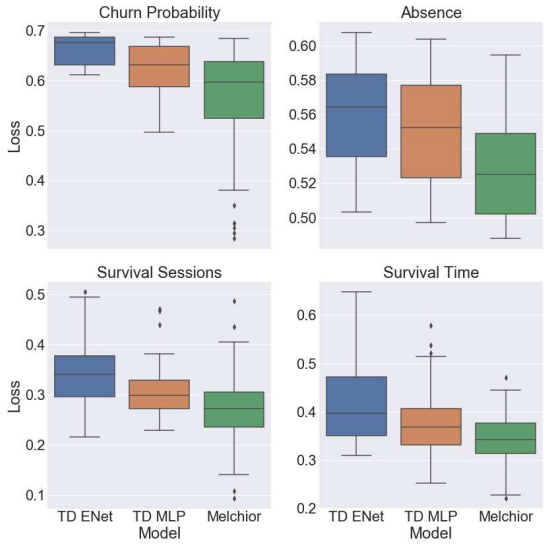

Fig. 5: **Models Comparison** The four plots report the performance of the three models

*2) **Experiment 2 - Embedding Visualization**:* Figure 6 reports the representation learned by MM on the test set.

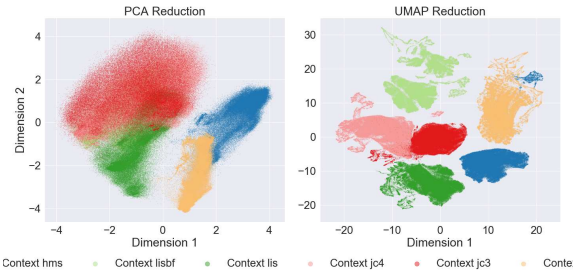

Fig. 6: **Learned representation for the six games contexts.**

We see that UMAP is able to more faithfully represents global geometry, placing points from different games within distinct regions, compared to traditional Principal Component Analysis (PCA). Focusing on jc4 and knowing that UMAP attempts to preserve the local geometry, we can see from Figure 7 that MM was able to generate a representation where users with similar behavioural measures of attributed salience (i.e. the four target metrics) are placed closer to each other in the embedding space. The fact that each metric is represented over different and partially non-overlapping areas of the embedding space indicates that the model learned the type of general representation that we illustrated with Eqn 4 and that is enforced by multitask learning. Such representation can be considered as an abstract and overarching factor, much like the concepts of engagement and attributed salience, able to explain the variations in the behavioural targets.

*3) **Experiment 3 - Embedding Partitioning**:* Figure 8 shows the results of K-Means partitioning on the embedding space. Looking at the characteristics of the partitions in Figure 9 we can see how each target metric develops differently

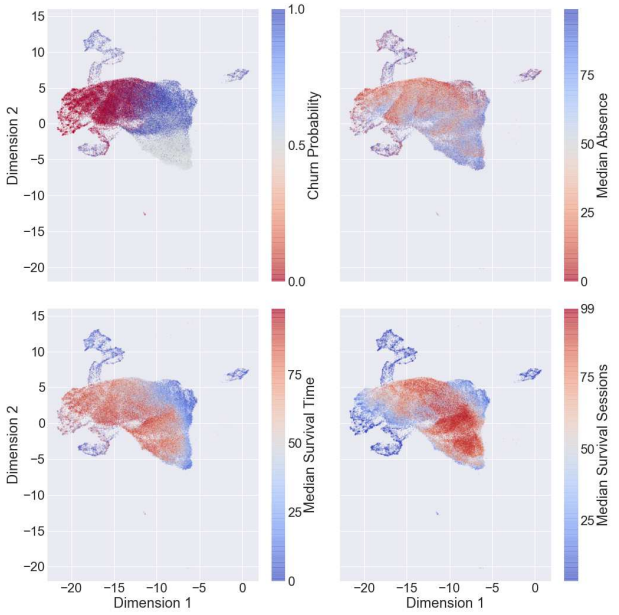

Fig. 7: **Visualization of the learned representation for the jc4 context.** Each plot report the UMAP reduction of the learned representation colour coded according to the median value of the four target metrics. From a behavioural point of view, blue colors might be interpreted as encoding low attributed salience while red colors high attributed salience. We advise to consult the ancillary results for examining the entire space.

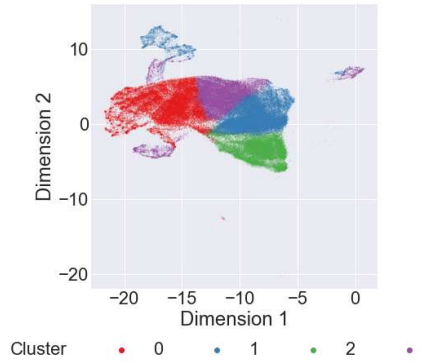

Fig. 8: **Visualization of the individuated partitions for the jc4 context.**

over time but that a general pattern seems to control each partition's trace. Partitions which have a low measurement in the target metrics representing attributed salience are shown to have shorter and less frequent interactions with the game, while the opposite appears to be the case for those partitions which have a high measurement in the same target metrics[3].

VI. DISCUSSION AND LIMITATIONS

*1) **Theoretical Implications**:* Our theoretical framework indicates that, from a behavioural point of view, I who attribute a high level of salience to a specific O will show a history of frequent and long interactions with O. Following the formulation by [16] this would mean that I stays in the sustained engagement state for longer, or frequently re-engage after periods of disengagement. According to our target metrics this would mean not churning, being absent for shorter and "surviving" for longer periods of time. Focusing on cluster number zero in figure 9 we see that it is the cluster that according to the distribution of its target metrics could be hypothesized having the highest level of attributed salience to playing the specific game. If we then look at the behavioural traces describing the interactions between the I in this cluster and the game, we can see that they perfectly match with what is prescribed by our theoretical framework (see sections III.a and III.b). We see that the same logic applies for all the other partitions but with slight differences indicative of the different manifestations of the engagement process. This, in addition to the results provided by experiment 1 and 2, shows that the theoretical constraints imposed on our model do not just lead to better performance when compared with unconstrained methodologies but also forces the model to learn a representations that appear to be coherent with our theoretical assumptions.

*2) **Practical Implications**:* As a byproduct of the present work we produced a scalable cross-game model for estimating metrics that are of core importance to industry applications (i.e. frequency and amount of future playing behaviour). The same model can be used as a feature extractor given its capacity to learn representations that are reasonable approximations of the quality of the interaction between a user and a game. Finally, the methodology we used for analyzing the learned embedding offers an efficient way to perform time series clustering, a task often crippled by prohibitive time and memory demands. This is done by first constructing a compact and static representation of the history of interactions between the user and the game and subsequently partitioning it. Moreover, the learned representation, when compared with the raw features traditionally used in the literature, has the advantage of having already been transformed by the network to best describe the target metrics.

*3) **Limitations and Future Work**:* This work has a series of limitations. 1) The analysis conducted for verifying the linkage between theory, model and observed behaviours is still preliminary and more careful and precise investigations would likely yield insights. 2) Due to the assumptions that k-means makes about the shape of the clusters it may be the case that more sophisticated approaches provide a higher-quality space partitioning. 3) For verifying that our modelling approach can be extended to general human behaviour we have to preform the same analysis done in this work with data coming from O that are not limited to video games. 4) Our approach appears to be suitable for modelling the amount of engagement or attributed salience but we can't say anything about the factors of I, O and Env which control the changes in the these constructs. 5) Differently to [13], [14] our approach doesn't explicitly produce a clear mathematical formulation explaining the observed changes in behaviour. This can at most

[3]Consulting the ancillary results we can see that this pattern is consistent through all the games.

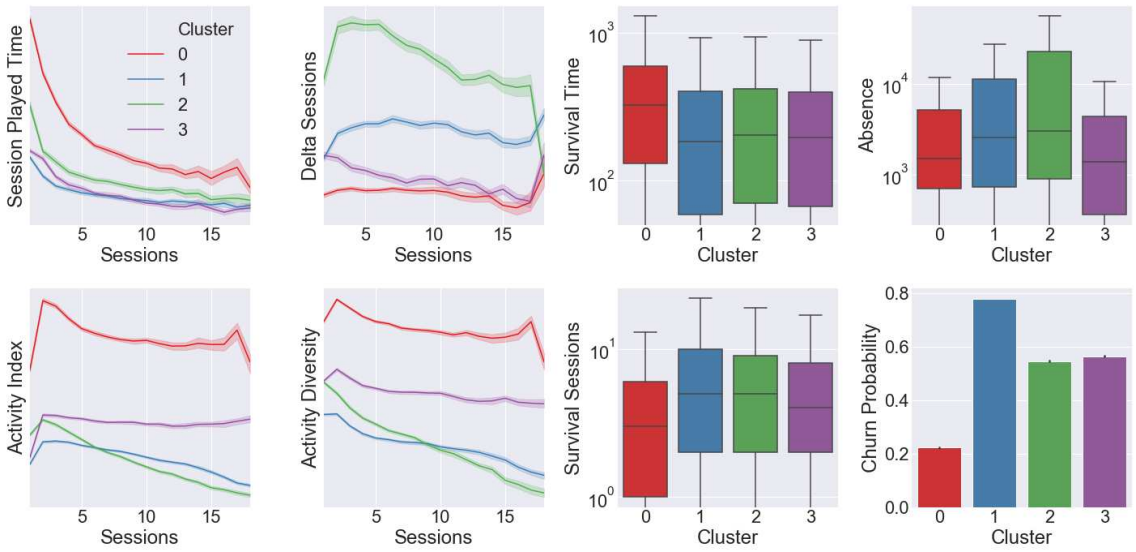

Fig. 9: **Characteristics of the individuated partitions.** The line-plots on the left show the progression over time of the input metrics, with the solid line indicating the average and the shaded area the 95% confidence interval. The plots on the right indicates the distribution of the target metrics.

be inferred or hypothesized observing the behavioral traces in Figure 9. All these constraints constitute venues for future works.

REFERENCES

[1] E. A. Boyle, T. M. Connolly, T. Hainey, and J. M. Boyle, "Engagement in digital entertainment games: A systematic review," *Computers in human behavior*, vol. 28, no. 3, pp. 771–780, 2012.

[2] V. Bonometti, C. Ringer, M. Hall, A. R. Wade, and A. Drachen, "Modelling early user-game interactions for joint estimation of survival time and churn probability," in *2019 IEEE Conference on Games (CoG)*. IEEE, 2019, pp. 1–8.

[3] M. S. El-Nasr, A. Drachen, and A. Canossa, *Game analytics*. Springer, 2016.

[4] E. Lee, Y. Jang, D.-M. Yoon, J. Jeon, S.-i. Yang, S.-K. Lee, D.-W. Kim, P. P. Chen, A. Guitart, P. Bertens *et al.*, "Game data mining competition on churn prediction and survival analysis using commercial game log data," *IEEE Transactions on Games*, vol. 11, no. 3, pp. 215–226, 2018.

[5] X. Liu, M. Xie, X. Wen, R. Chen, Y. Ge, N. Duffield, and N. Wang, "Micro-and macro-level churn analysis of large-scale mobile games," *Knowledge and Information Systems*, pp. 1–32, 2019.

[6] A. F. del Río, A. Guitart, and Á. Periáñez, "A time series approach to player churn and conversion in videogames," *arXiv preprint arXiv:2003.10287*, 2020.

[7] J. T. Kristensen and P. Burelli, "Combining sequential and aggregated data for churn prediction in casual freemium games," in *2019 IEEE Conference on Games (CoG)*. IEEE, 2019, pp. 1–8.

[8] A. Drachen, E. T. Lundquist, Y. Kung, P. Rao, R. Sifa, J. Runge, and D. Klabjan, "Rapid prediction of player retention in free-to-play mobile games," in *Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2016.

[9] A. F. del Río, P. P. Chen, and A. Periánez, "Profiling players with engagement predictions," in *2019 IEEE Conference on Games (CoG)*. IEEE, 2019, pp. 1–4.

[10] A. Drachen, R. Sifa, C. Bauckhage, and C. Thurau, "Guns, swords and data: Clustering of player behavior in computer games in the wild," in *2012 IEEE conference on Computational Intelligence and Games (CIG)*. IEEE, 2012, pp. 163–170.

[11] S. Makarovych, A. Canossa, J. Togelius, and A. Drachen, "Like a dna string: Sequence-based player profiling in tom clancys the division," in *Artificial Intelligence and Interactive Digital Entertainment Conference*. York, 2018.

[12] A. Drachen, A. Canossa, and G. N. Yannakakis, "Player modeling using self-organization in tomb raider: Underworld," in *2009 IEEE symposium on computational intelligence and games*. IEEE, 2009, pp. 1–8.

[13] D. Reguera, P. Colomer-de Simón, I. Encinas, M. Sort, J. Wedekind, and M. Boguñá, "Quantifying human engagement into playful activities," *Scientific Reports*, vol. 10, no. 1, pp. 1–7, 2020.

[14] C. Bauckhage, K. Kersting, R. Sifa, C. Thurau, A. Drachen, and A. Canossa, "How players lose interest in playing a game: An empirical study based on distributions of total playing times," in *2012 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 2012, pp. 139–146.

[15] 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.

[16] H. L. O'Brien and E. G. Toms, "What is user engagement? a conceptual framework for defining user engagement with technology," *Journal of the American society for Information Science and Technology*, vol. 59, no. 6, pp. 938–955, 2008.

[17] K. C. Berridge and T. E. Robinson, "What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience?" *Brain research reviews*, vol. 28, no. 3, pp. 309–369, 1998.

[18] S. M. McClure, N. D. Daw, and P. R. Montague, "A computational substrate for incentive salience," *Trends in neurosciences*, vol. 26, no. 8, pp. 423–428, 2003.

[19] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT press, 2016.

[20] F. Chollet *et al.*, "Keras," https://keras.io, 2015.

[21] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, "Hyperband: A novel bandit-based approach to hyperparameter optimization," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 6765–6816, 2017.

[22] L. McInnes, J. Healy, and J. Melville, "UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction," *ArXiv e-prints*, Feb. 2018.

[23] D. Sculley, "Web-scale k-means clustering," in *Proceedings of the 19th international conference on World wide web*, 2010, pp. 1177–1178.

[24] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikitlearn machine learning in python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.