

METHODS

Profiling Television Watching Behavior Using Bayesian Hierarchical Joint Models for Time-to-Event and Count Data

RAFAEL A. MORAL¹, ZHI CHEN², SHUAI ZHANG², SALLY MCCLEAN², (Member, IEEE),
GABRIEL R. PALMA³, BRAHIM ALLAN⁴, AND IAN KEGEL⁴

¹Department of Mathematics and Statistics, Maynooth University, Maynooth, W23 F2H6 Ireland

²School of Computing, Ulster University, BT15 1ED Belfast, U.K.

³Hamilton Institute, Maynooth University, Maynooth, W23 F2H6 Ireland

⁴British Telecommunications, IP5 3RE Ipswich, U.K.

Corresponding author: Rafael A. Moral (rafael.deandrademoral@mu.ie)

This work was supported by the British Telecommunications (BT) Ireland Innovation Centre (BTIIC) funded by BT and Invest Northern Ireland.

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the YouView Data and Insights Team.

ABSTRACT Customer churn prediction is a valuable task in many industries. In telecommunications it presents great challenges, given the high dimensionality of the data, and how difficult it is to identify underlying frustration signatures, which may represent an important driver regarding future churn behaviour. Here, we propose a novel Bayesian hierarchical joint model that is able to characterise customer profiles based on how many events take place within different television watching journeys, and how long it takes between events. The model drastically reduces the dimensionality of the data from thousands of observations per customer to 11 customer-level parameter estimates and random effects. We test our methodology using data from 40 BT customers (20 active and 20 who eventually cancelled their subscription) whose TV watching behaviours were recorded from October to December 2019, totalling approximately half a million observations. Employing different machine learning techniques using the parameter estimates and random effects from the Bayesian hierarchical model as features yielded up to 92% accuracy predicting churn, associated with 100% true positive rates and false positive rates as low as 14% on a validation set. Our proposed methodology represents an efficient way of reducing the dimensionality of the data, while at the same time maintaining high descriptive and predictive capabilities. We provide code to implement the Bayesian model at https://github.com/rafamoral/profiling_tv_watching_behaviour.

INDEX TERMS Bayesian modelling, big data, churn prediction, clustering, dimensionality reduction, frustration signatures, machine learning.

I. INTRODUCTION

This paper presents the development and testing of a Bayesian hierarchical model-based approach to characterise the viewing behaviours of BT TV customers. A hierarchical model is appropriate here as each customer has multiple viewing sessions which can be quite diverse both in content and duration. Features extracted from this customer behaviour

The associate editor coordinating the review of this manuscript and approving it for publication was Ines Domingues^{id}.

model can then be used to profile customers in terms of their preferences and engagement. Features extracted from such profiles are then used to classify customers as either frustrated or satisfied, where we hypothesise that customer frustration is often an early indicator of later churn which could be driven by a perceived lack of content choice. The approach is validated using labelled data from customers who subsequently cancelled their subscription. User profiling is a key element in understanding user behaviours and relating these behaviours to subsequent actions.

BT TV is a subscription IPTV service provided by the BT Group, which uses the YouView television platform, to provide a range of channels and services. An anonymised dataset, based on YouView client logs, was prepared and shared by the YouView Data and Insights team for behavioural data analysis and modelling purposes, as described in this paper. As the aim here is to develop an initial model, the study was conducted on a subset of the data consisting of 20 active customers and 20 customers that subsequently cancelled the service. Overall, there were 32,556 journeys in the sample data from these customers that cancelled. The proposed model and associated results can then serve as a basis of transferring the knowledge to a much bigger dataset from BT customers.

The YouView data in this study are primarily time-stamped system logs concerning sequences of TV customer events of interest. A major focus in BT is to better understand how customer navigate their content and, from that, infer whether these customers are satisfied with the service and identify where they may be frustrated by their engagement. The analysis and modelling of such time-stamped event data has become increasingly popular over recent years, especially for business data where such work falls under the Process Mining umbrella [1]. Here, process mining can be defined as a data driven analytics approach used to extract knowledge and insights from the event logs, and use them to discover, monitor and improve the process [2]. In our case we use TV customer behaviour data provided by YouView in the form of records that capture the customer interactions with TV from the point when the TV is turned on, to the time when the TV is turned off, including all steps in the path to reach a particular channel; failures, reported as error messages, and changes of channel, are also recorded. This is a rich process, generating multiple and diverse possible trajectories for huge numbers of customers and, as such, is well suited to a mining and modelling approach.

More specifically, in the current research we use the YouView data to develop and implement a Bayesian hierarchical joint model for time-between-events and number of customer events. The model is implemented so as to profile typical journeys for each customer. It is based on covariates, such as number of viewing events and programme genre. We then use a combination of statistical machine learning methods to classify the customers into “frustration classes”, where the Bayesian hierarchical joint model parameters and associated covariates serve as additional features for the classification. The customers were subsequently labelled as either “active” or “cancelled”, regarding their subscription. While there is an extensive literature on predicting churn [3], particularly for telecoms, our approach is novel in that we provide a hybrid model to proactively predict subscription cancellation. The hierarchical Bayesian model proposed here handles heterogeneous and complex data efficiently. It is used to identify and extract concise latent customer-specific frustration signatures, which are then used for efficient churn classification,

thus paving the way for pre-emptive intervention as a strategy to improve customer retention.

II. PREVIOUS WORK

Customer Relationship Management (CRM) is an important aspect of business practice and focuses on developing a loyal and long-term customer base. An important aspect of CRM is customer retention which is often much more difficult and expensive than acquiring new customers. The burgeoning accumulation of massive business data collections, particularly in the telecoms industry, and increasingly mature data analytics technologies have served to accelerate the development and deployment of customer churn models to predict customer behaviour. In a recent survey, [3] demonstrate that data analytics techniques are becoming increasingly used for customer churn prediction, particular in the telecoms industry, where neural networks, decision trees, support vector machines and logistic regression have been most commonly used as the classification algorithms. [4] also observed that decision trees and logistic regression have both been used frequently in customer churn prediction. However, these algorithms fail to cope well with data complexities and interactions so, as a result, [4] improved on such approaches by introducing the logit leaf model (LLM) which segments the data, before prediction, thus better characterising the feature domain, while retaining previous advantages of strong predictive performance and good comprehensibility.

Since modern-day customers exhibit diverse and highly dynamic behaviours, retention is often highly important to businesses. In such dynamic situations, churn can be caused by customer dissatisfaction, rival products or services, regulations or the customer’s personal circumstances which the business cannot influence. The importance of dynamic behaviour-based churn prediction for telecoms was emphasised by [5] due to both the dynamic nature of the customer activities and their subsequent behaviours. [6] have recently highlighted the improving opportunities for real-time monitoring and resulting ability to discover changes in psychological patterns. They focus on Change Point Analysis (CPA) and describe a Pattern Transition Detection Algorithm for psychological time series data with pattern transitions. In particular, characterising human behaviour from online data has become increasingly feasible, particularly for Cybersecurity [7], where it is important to distinguish between normal and pathological user behaviours, in a timely manner. In our problem domain, IPTV customer logs capture comprehensive customer interactions and activities, facilitating the development of customer behaviour modelling. Such models may then be used to predict future behaviours of interests, such as churn. In such an approach, we can analyse the history of interactions and activities of dissatisfied customers who subsequently churn and compare this with the corresponding profile of a satisfied customer. For example, [8] analysed the activity data for Massively Multiplayer Online Role-Playing Games (MMORPGs) to identify three factors of engagement,

enthusiasm and persistence which proved useful features for churn prediction; these resonate well with our current focus on IPTV where the customer logs can track engagement with the service in terms of session viewing, enthusiasm as evidenced by volume of activity and persistence in terms of sustained viewing. For the telecoms and related online-based industries, the product is a subscription service [9], where prediction of subscription has been, and still is, a very active research area; in the current paper, churn presents as cancellation of the subscription, as in, for example, [10]. For such churn prediction, with advances in technology, (near) real-time data are becoming increasingly available and can serve to provide early indications of churn and learn when customers have an increased probability of churning [11]. The early detection of potential churners facilitates focussing on such customers for targeted interventions [12]. Customer churn management can be done in either reactive or proactive mode [13]. In the reactive mode, churn occurs first and the customer is then offered an incentive to return. In proactive mode, the probability of churn is predicted from available data, and, if churn is likely, incentives can be offered to increase the likelihood of retention. With the increasing availability of online, (near) real-time data, there are increasing possibilities for using dynamic features to characterise user behaviours and provide improved opportunities for proactive churn prediction.

The desirability of using hybrid/hierarchical approaches to classifying complex churn data from online domains has been highlighted by [14] where clustering is used to identify churners, and subsequently, a survival model is used to build relevant features into the classification model. [15] and [4], *inter alia*, provide further discussions of the advantage of such hybrid approaches for telecoms churn prediction.

III. METHODS

A. DATA PREPARATION

YouView Data and Insights team initially provided an anonymised dataset containing client event logs from BT YouView set-top boxes (STB) to aid our research. Informed consent was obtained by the YouView team for all participants. The data contains information about TV customer behaviour in terms of interactive activities between customers and the Electronic Programme Guide (EPG) system on the STB, such as powering on/off the box, tuning into IPTV channels, beginning/ending playback of Video-on-Demand content, etc. In addition to STB performance information, the data contains error/failure logs. Consequently, the raw data contains multiple streams of events, and we must extract the key data flow that illustrates customer journey paths. In this study, a customer journey refers to the steps a customer takes to navigate the EPG system from the time the STB is turned on until the customer reaches IPTV channel content.

Some columns in the raw dataset contain crucial information, such as “EVENT DT TM” displaying timestamps and “EVENT ID” displaying the event log category for each

record. And for other columns such as “EVENT SPEC 1”, include information regarding subcategories for specific main event categories. Using the domain knowledge shared by the YouView Data and Insights team, we create two separate reference tables for the original event logs in order to extract target data streams. Specifically, one table displays a list of unique event IDs for all events generated by customers’ “click” actions on the remote control. In addition to the event ID, the second table contains additional information for each target event log, including the IPTV system’s URL address, which indicates the UI context information. Using the two reference tables, we can extract a subset of the raw data comprising clickstream and UI context data. The target event log is then renamed by appending a shorter tag referencing the two tables. The general procedure is depicted in Figure 1(a).

Using the timestamp information from the raw dataset in conjunction with the newly generated labels, we can convert the initial event logs into event sequences for each customer journey on the STB. Each event sequence consists of a series of “click” events separated by intervals in the format “event-interval-event.” Using dummy data, Figure 1(b) provides an illustration of the procedure, while a snapshot of the sequences generated by the real data is shown in Figure 1(c).

As a result, we have reformatted the data by generating a new table of event sequences in accordance with the preceding step. In addition, each sequence’s start and end times, total duration, channel name, and genre name are included in the new table.

B. THE MODEL

The hierarchical model was developed with the objective of obtaining customer signatures, or profiles, based on their behaviour and then clustering them into similar groups. By comparing the observed groupings with the known cancellation variable, this can generate insight about a latent frustration signature.

Since journeys are nested within customers, and their features are likely to be correlated, we propose a hierarchical model. We attempt to jointly model the total number of events in a journey and the time between two consecutive events within the journey (see Figure 2). We assume that the time between two events depends on the previous time between two events, *i.e.* a first-order Markov assumption.

We assume that the total time per session is a random truncated Poisson sum of gamma random variables. In our notation, c is the index for customer, $c = 1, \dots, C$, s is the index for journey within customer, $s = 1, \dots, S_c$, i is the index for event within journey within customer, $i = 1, \dots, I_{sc}$, T_{sc}^* is the random variable representing the total time per journey, T_{isc} is the random variable representing the duration of an intermediate state (event), and N_{sc} is the number of events (states) per journey. We have that

$$T_{sc}^* = \sum_{i=1}^{N_{sc}} T_{isc}.$$




FIGURE 1. Illustration of how the YouView set-top box data was prepared for analysis. We illustrate event-labelling in (a), where the label 'A_1' is created by cross-referencing two tables. In (b), we provide an example of sequence generation from data on user IDs and timestamped events; using regular expression manipulation the table from the left-hand side is converted to a table where each row corresponds to a journey. Finally, in (c) a snapshot of generated journeys/sequences is displayed: in the first row, the customer starts with event 'K_385'; then 6.479 seconds later, event 'D_58' is triggered, and so on.




FIGURE 2. A diagram representing one customer journey (session) with n events, where t represents the time between two consecutive events.

We may write the model for the number of events in a journey as

$$N_{sc} | \mathbf{d}_c \sim \text{Poisson}^+(\lambda_{sc})$$

where the linear predictor depends on customer-level random effects \mathbf{d}_c , and may be written as

$$\log(\lambda_{sc}) = d_{0c} + d_{1c} \times (\text{predictor at journey level})_{sc},$$

with $\mathbf{d}_c \sim N(\boldsymbol{\delta}, \boldsymbol{\Sigma}_d)$ and $\boldsymbol{\Sigma}_d = \text{diag}\{\sigma_{d_0}^2, \sigma_{d_1}^2, \dots\}$ a diagonal variance-covariance matrix. The time in between events depends on the previous time, and is assumed to follow a gamma distribution, i.e.

$$T_{isc} | T_{i-1,sc}, \mathbf{b}_c, p_c \sim \text{Gamma}(\mu_{isc}, \psi_c)$$

with the linear predictor for the mean μ_{isc} written as the sum of two components:

$$\log(\mu_{isc}) = \eta_{isc} + \omega_{isc}.$$

The first, η_{isc} , includes the effects of covariates at the event level, modelled as random effects within customers, and may be written as

$$\eta_{isc} = b_{0c} + b_{1c} \times (\text{predictor at event level})_{isc},$$

with $\mathbf{b}_c \sim N(\boldsymbol{\beta}, \boldsymbol{\Sigma}_b)$, $\boldsymbol{\Sigma}_b = \text{diag}\{\sigma_{b_0}^2, \sigma_{b_1}^2, \dots\}$. The second component, ω_{isc} , is the autoregressive component, and models the dependence over time with the autoregressive parameter p_c , also a customer-level random effect, such that

$$\omega_{isc} = p_{sc} \times T_{i-1,sc},$$

with $p_{sc} \sim N(\phi_c, \sigma_p^2)$.

C. CASE STUDY

The original data provided by YouView contains approximately 100 million event logs from more than 10,000 customers over a three-month period (October to December 2019), which is large and complex data. The data also includes labelled customers, i.e., fifty percent of the customers cancelled their BT subscriptions while the other fifty percent remained with BT. For customers who cancelled their subscription the data represents the last journeys leading to cancellation within the time period of October to December 2019. To simplify our study, we chose a sample of 40 customers' journey data (20 from customers who cancelled and 20 from customers who did not), including only "click"-related records, which totalled over 30,000 journeys in the sample data.

For each customer we had variable numbers of journeys (ranging from about 500 to over a thousand), and within each journey a variable number of events. The number of events within journeys varied from only one to over 500. The median number of events per journey was small, but the distribution was highly skewed to the right, with many journeys totalling more than 100 events. The total number of events in the sample data was approximately half a million.

The size of the dataset caused a problem with respect to memory allocation when fitting the model. Therefore, we proceeded with a subset of this data, using the first 300 journeys per customer, and up to the first 300 events within each journey, which amounted to a little over a third of the available data.

D. MODEL FITTING AND ANALYSIS

We model the mean number of events with a random intercept per customer, and include the random effects of the eight types of channels watched in a particular event within session in the linear predictor for η_{isc} . The linear predictors we use for the sample data described above are

$$\log(\lambda_{isc}) = d_{0c}$$

and

$$\eta_{isc} = \sum_{k=1}^K b_{kc} \times \text{genre}_{k,isc},$$

where $\text{genre}_{k,isc}$ are dummy variables indicating whether event i within session s for customer c was related to programme genre k . More specifically, there are eight major genres of IPTV channels on the BT TV box. However, to avoid intellectual property violations against BT or YouView, we denote them by numbers only ($G^{(1)}$ to $G^{(8)}$). The linear predictor used can be written as

$$\begin{aligned} \eta_{isc} = & b_{1c} \times G_{isc}^{(1)} + b_{2c} \times G_{isc}^{(2)} + b_{3c} \times G_{isc}^{(3)} \\ & + b_{4c} \times G_{isc}^{(4)} + b_{5c} \times G_{isc}^{(5)} + b_{6c} \times G_{isc}^{(6)} \\ & + b_{7c} \times G_{isc}^{(7)} + b_{8c} \times G_{isc}^{(8)}, \end{aligned}$$

where $d_{0c} \sim \mathcal{N}(\delta, \sigma_d^2)$, $b_{kc} \sim \mathcal{N}(\beta_k, \sigma_{b_k}^2)$, $k = 1, \dots, 8$.

We estimate the model using the Bayesian framework, with 1,000 adaptation steps, 2,000 MCMC iterations as burn-in, and 20,000 MCMC iterations with a thinning of 10, for each of three MCMC chains. We use flat normal priors for δ , β_k and ϕ_c , since these parameters can take any real value; a flat gamma prior for ψ_c , since these parameters are strictly positive; and half-Cauchy priors for the variance components, as suggested by [16], i.e.

$$\begin{aligned} \delta & \sim \mathcal{N}(0, 1000) \\ \beta_k & \sim \mathcal{N}(0, 1000) \\ \phi_c & \sim \mathcal{N}(0, 1000) \\ \psi_c & \sim \text{Gamma}(0.001, 0.001) \\ \sigma_d^{-2} & \sim \text{Half-Cauchy} \\ \sigma_{b_k}^{-2} & \sim \text{Half-Cauchy} \end{aligned}$$

$$\sigma_p^{-2} \sim \text{Half-Cauchy}$$

Model implementation was carried out using JAGS [17] within R [18].

E. POST-PROCESSING AND BENCHMARKING

For each customer we compiled the estimates for the auto-correlation parameter ϕ (describing the degree of first-order dependence in terms of time taken between events), the gamma distribution dispersion parameter ψ (describing the degree of variation in the time between events), the random intercepts d_0 (average number of events per journey) and random genre effects b_1 to b_8 (effect of each genre on the average time between events). Therefore, each customer's profile can be described by these 11 variables. We then scaled each variable to have mean zero and unit variance. As exploratory analysis, we obtained a matrix of the Euclidean distances between customers and performed hierarchical clustering using Ward's method, so as to minimise the intra-cluster variability. We then produced a dendrogram.

We used six different statistical machine learning methods to assess how accurate we could classify customers as "active" or "cancelled": (1) logistic regression (LR) [19], (2) random forests (RF) [20], (3) support vector machines (SVM) with a third-degree polynomial kernel [21], (4) ν -support vector classifiers (ν -SVC) [22], (5) k -nearest neighbours (k -NN) [23], and (6) deep neural networks (DNN) [24]. For the logistic regression we used a LASSO-type penalty [25] to allow for mild regularisation of the predictor space, fixing the penalty parameter at 0.02. For the random forests, we used 1,000 trees and a number of possible predictors to choose from equal to the floor of the square root of the total number of predictors. For KNN, we experimented with a different number of nearest neighbours and used $k = 1$ in the end, which yielded the best performance. For DNN, after experimenting with different architectures, we used 17 hidden layers with 11, 12, 13, 14, 15, 14, 13, 12, 11, 10, 9, 8, 7, 6, 5, 4, and 3 neurons each, all using the rectified linear unit (ReLU) activation function; for the output layer we used the sigmoid activation function. For all methods, we used a 70:30 split for training and test data (therefore 28 customers for training the model and 12 for validation), and calculated the overall accuracy, and true and false positive rates as measures of predictive power of the machine learning methods.

We benchmarked the results from the machine learning methods against a naïve approach which consisted of using times between events within each journey for all customers as input for the machine learning algorithms described above. The times between events were used as predictors in the order journeys were recorded. Customers had an unequal number of recorded journeys, and journeys had unequal numbers of events. To circumvent the issue of incomplete cases, the minimum number of journeys per customer and minimum number of events per each journey were used. For example, if we had three customers with 20, 29 and 35 journeys each, we used data from the first 20 journeys for each customer.

The same applies to events within journeys. The predictors were therefore the times between events in order of event occurrence within each journey, with journeys stacked as a vector. We refer to this approach as naïve for two reasons. First, customers could display different frustration signals, and recordings could have taken place at different stages of their frustration, and therefore the order with which journeys occur might be uninformative. Second, a large amount of information was lost due to restricting observations to generate complete cases, rather than describing each customer with the same sets of variables that reflect their overall profile (which is what the hierarchical Bayesian model proposed here attempts to do). In total, the benchmarking predictor space included 1,012 variables.

The implementation of the classification methods was made in Python using the libraries TensorFlow and Scikit-learn [26]. The libraries NumPy and Pandas were also used for data manipulation and numerical analysis.

We provide code to implement the Bayesian model and machine learning methods at https://github.com/rafamoral/profiling_tv_watching_behaviour. However, we do not provide the raw data to avoid intellectual property violations against BT or YouView.

IV. RESULTS

A. CUSTOMER PROFILES

The model took approximately 49 hours to run using four cores on an R Studio Server with an Intel Xeon 4614 2.2 GHz processor with 576 GB RAM, using three cores (one for each MCMC chain). Table 1 displays the estimates for the δ and β model parameters and variance components, associated standard errors and 95% credible intervals, while Figure 3 displays box plots for the customer-level parameter estimates (ψ and ϕ) and random effects.

The means of the random effects corresponding to different programme genres are all close to zero, which means they play no significant role in predicting the time between events, apart from $\hat{\beta}_5$ (“Genre 5”) and $\hat{\beta}_8$ (“Genre 8”), which indicate that the times between events are $1 - e^{-0.48} = 38\%$ and $1 - e^{-0.94} = 61\%$ shorter, on average, when these programme genres are being watched, respectively. For the other genres, the main predictor of time until the next event is the time between the two previous events, modelled by the autoregressive parameter. There is, however, a degree of individual variability, modelled by the variance components (which are close to 1 for all genres). Customers whose random effects are at the lower end could indicate a “zapping” signature, which means they do not spend much time on a particular channel and instead, “zap” to other channels, making the time between events short and the total number of events large. On the other hand, customers whose random effects are at the higher end would indicate more time spent watching one particular programme. This was seen more for the customers who cancelled their subscription (see Figure 3).

TABLE 1. Parameter estimates, standard errors and 95% credible intervals for the model fitted to the 40 customer sample data, using the first 300 journeys per customer and up to the first 300 events within each journey, which accounts for approximately a third of the data. We omit the estimates for ϕ_c and ψ_c , since there are 40 of each (one for each customer). They are shown, however, in Figure 3, alongside predicted random effects.

Parameter	Estimate (std. error)	95% Credible interval	
		Lower bound	Upper bound
δ	2.49 (0.06)	2.37	2.61
σ_d	0.39 (0.04)	0.32	0.49
β_1	-0.19 (0.17)	-0.53	0.15
β_2	-0.15 (0.19)	-0.51	0.22
β_3	-0.24 (0.18)	-0.60	0.11
β_4	-0.06 (0.17)	-0.40	0.28
β_5	-0.48 (0.27)	-0.99	0.04
β_6	-0.31 (0.21)	-0.73	0.11
β_7	0.03 (0.20)	-0.46	0.41
β_8	-0.94 (0.45)	-1.80	-0.01
σ_{b_1}	1.10 (0.12)	0.89	1.37
σ_{b_2}	1.15 (0.13)	0.93	1.44
σ_{b_3}	1.11 (0.13)	0.89	1.39
σ_{b_4}	1.06 (0.12)	0.85	1.33
σ_{b_5}	1.50 (0.20)	1.15	1.93
σ_{b_6}	1.26 (0.15)	1.01	1.60
σ_{b_7}	1.13 (0.15)	0.89	1.45
σ_{b_8}	0.98 (0.41)	0.37	1.95
σ_p	0.23 (0.00)	0.23	0.24




FIGURE 3. Box plots for the customer-level parameter estimates (ψ , the dispersion of time-between-events, and ϕ , the mean autocorrelation between previous time-between-events and the time until the next event), and random effects (b_1, \dots, b_8 , corresponding to the effects of eight programme genres on time-between-events) for each of 40 customers, 20 having an active subscription to the TV service 20 users that cancelled the service.

The dispersion parameter ψ of the gamma distribution was small for all customers and did not vary much between them (Figure 3). Since the parameterisation of the gamma model used here implies that $\text{Var}(T) = \mu^2/\psi$, the model estimates indicate there was high variability in time between events. The variance components for the autoregressive parameter and number of events were small, indicating a smaller degree of individual variability between customers when looking at




FIGURE 4. Correlation matrix of the estimated parameters and random effects per customer for the active group (lower triangular part) and group that cancelled the service (upper triangular part).

the dependence on time between previous events. However, on average the autoregressive random effects were slightly higher for customers who had cancelled their subscription (Figure 3).

The mean number of events d was negatively correlated with all genre-related random effects b . This was expected, since a shorter time between events ultimately leads to more events taking place. The genre-related random effects b were highly correlated between themselves, with this correlation being slightly weaker for customers in the “cancelled” group. This suggests that television watching signatures did not change depending on programme genre. It is noteworthy to mention that the correlation between b_8 (“Genre 8”) and all other genre-related random effects was much weaker, for customers both in the “active” and “cancelled” groups (see Figure 4). The correlation is slightly smaller between b_5 (“Genre 5”) and all other genre-related random effects for customers in the “cancelled” group, which may help differentiate the two groups when applying machine learning methods.

Correlations between the dispersion parameter estimates and other customer-level parameter estimates and random effects are all negative (apart from d), indicating that lower times between events yielded higher dispersion estimates, ultimately translating into a smaller variance (Figure 4). Conversely, longer times between events yielded smaller dispersion estimates, which represent larger

variances. This is straightforward to explain, for when “zapping” variance will be smaller, since times are smaller and more consistent. If watching for longer times, they have scope to vary more. These correlations are much weaker for customers in the “cancelled” group, showing that their “zapping” behaviour was less intense. Overall variability remained unchanged when looking at customers that showed lower levels of activity in terms of triggering events versus those who spent more time watching offered content, and this can be helpful when attempting to identify an underlying frustration signature.

B. CUSTOMER CLUSTERING

When clustering customers based on their raw times between events, the dendrogram does not show a clear separation between customers in the “active” versus “cancelled” groups. Even though Ward’s method was used to minimise intra-cluster variability while maximising inter-cluster variability, the dendrogram seems to exhibit some level of chaining (see Figure 5, dendrogram on the left), and active customers are clustered alongside customers who cancelled their subscriptions at small Euclidean distances. The dendrogram produced using the customer-related estimated parameters and random effects from the Bayesian hierarchical model fit shows a much clearer two-cluster solution (see Figure 5, dendrogram on the right), and although there is no perfect separation between “active” and “cancelled” customer groups, at smaller Euclidean distances customers are generally clustered within their own groups.

C. CUSTOMER CLASSIFICATION

It is important to note that the classification results are subject to a dataset with a small number of individuals, and therefore some of the machine learning methods may not perform optimally due to the sample size.

When attempting to classify customer status (those who are “active” vs. those who had “cancelled” their subscription), the results obtained using the parameter estimates from the Bayesian hierarchical joint model were superior to the ones obtained using the raw data for all machine learning methods, in terms of accuracy, true and false positive rates (see Table 2). The SVM method yielded the best overall performance, with an accuracy of 92% associated with 100% true positive rate and 14% false positive rate. Moreover, DNN and RF also presented 100% true positive rates, however had a smaller accuracy due to mis-classification of customers in the “active” group. Finally, k -NN was also associated with a low false positive rate, but mis-classified customers in the “cancelled” group. The maximum accuracy obtained using the naïve approach was 58%, using the ν -SVC method; however, it was associated with a very low true positive rate (40%).

V. DISCUSSION

We proposed a novel Bayesian hierarchical joint model to analyse customer journeys. This model is able to characterise

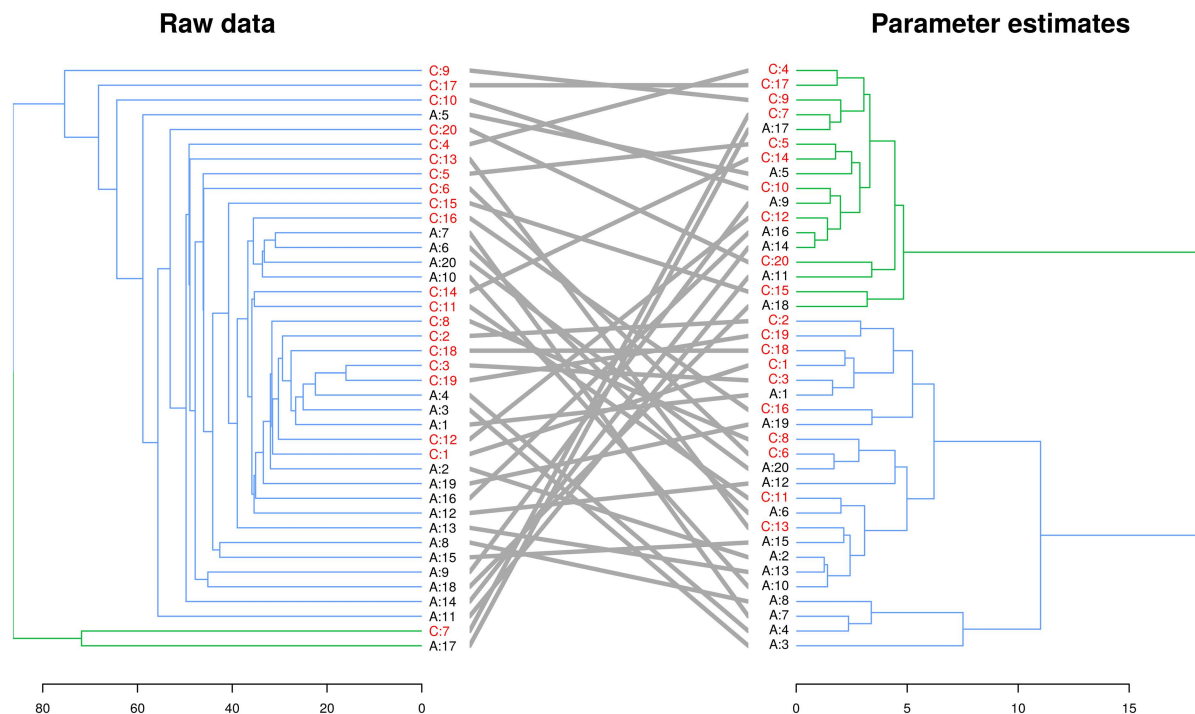


FIGURE 5. Dendrograms obtained from the hierarchical clustering applied to the raw data (left) and parameter estimates obtained from the modelling approach, indicating customer profiles (right). The x-axis represents the Euclidean distance between clusters. “A” means that customer had an “active” subscription and “C” means the customer had “cancelled” the service. The two-cluster solution is shown in different colours (green vs. blue), and the labels for customers in the “cancelled” group are displayed in red for better visualisation. The gray lines connect the same customers to show their different positions in the two separate dendrograms.

TABLE 2. Accuracy (Acc), true positive rates (TPR) and false positive rates (FPR) for each of six machine learning methods applied to the classification of customer status (“active” vs. “cancelled”) based on the parameter estimates obtained from the Bayesian hierarchical modelling approach (model-based) or the raw data (naïve). We used a 70:30 split between training:validation sets, i.e. the training set contained 28 customers and the validation set contained 12 customers. All measures reported on this table were obtained in the validation set. SVM: support vector machines with a third-degree polynomial kernel; DNN: deep neural networks; ν -SVC: ν -support vector classifier; k -NN: k -nearest neighbours; RF: random forests; LASSO-LR: logistic regression with LASSO-type penalty.

Method	Model-based			Naïve		
	Acc	TPR	FPR	Acc	TPR	FPR
SVM	0.92	1.00	0.14	0.41	1.00	1.00
DNN	0.83	1.00	0.29	0.42	1.00	1.00
ν -SVC	0.75	0.80	0.29	0.58	0.40	0.29
k -NN	0.75	0.60	0.14	0.33	0.80	1.00
RF	0.67	1.00	0.57	0.25	0.40	0.86
LASSO-LR	0.67	0.80	0.57	0.25	0.49	0.86

customer profiles based on how many events take place within different journeys and also the time taken between events. The model drastically reduces the dimensionality of the data from thousands of observations per customer to 11 customer-level parameter estimates and random effects. This allows for prompt application of many different machine learning methods for classification, ultimately allowing for churn prediction.

Our results show that our model reflects the underlying processes well, and is a sensible representation of a customer’s profile. Therefore it is able to retain the

most prominent features of the data, and is a better candidate for machine learning techniques rather than using the high-dimensional raw data. Multicollinearity between genre-related random effects may represent an additional challenge when employing machine learning methods, however that was circumvented by the classification tools used in different ways (e.g., regularisation when using LASSO-type penalties, random selection of predictors within a random forests framework, incorporation of non-linearities when splitting the parametric space within a support vector classifier framework, etc).

The strengths of our proposed methodology are evidenced by the high classification accuracy obtained. Obtaining customer profiles based on the Bayesian hierarchical joint model was therefore a successful way of dealing with the high-dimensional journey data. Not only does this represent promising results from a classification perspective, it also provides interpretable machine learning aimed at characterising customer churn. For instance, from a random forests model, we may look at different measures of variable importance, whereas in a logistic regression framework a LASSO penalty will automatically select variables that are most important. Moreover, the estimates from the Bayesian model themselves are characterisations of each customer’s profile, as well as how different covariates may affect the behaviour of the time between events and number of events in their journeys. This makes the novel model proposed here useful not only for post-processing aimed at churn prediction, but

also for using customer status as a predictor variable and studying, in hindsight, what are the most prominent features that may have led to churn. We highlight, however, that customers present very heterogeneous watching behaviour, and therefore it is difficult to identify potential indicators of frustration, especially given a small sample size. Complementary approaches are useful in this case, such as sequence mining [27].

We identify two main limitations to this study. First, the number of customers available in the dataset is small. A larger pool of customers would allow for better validation of the methods. However, this does serve as proof of concept and already shows great promise, for with a small sample the methodology showed a high predictive power. Second, fitting the Bayesian hierarchical joint model takes a high toll in terms of computational burden. Given the high-dimensionality of the data, allocating the necessary memory proves to be a difficult task, even when using a high-spec server. Nevertheless, even when using about a third of the data, we already obtained high predictive power, again demonstrating that our proposed framework is capable of accurately describing customer profiles.

Future work involves (1) overcoming memory limitations to run the model for a larger sample of customers, including more journeys and events; (2) carrying out simulation studies to assess the trade-offs between using more journeys versus more events within journeys; and (3) considering potential alternatives to speed up computation, allowing for scalability of the methodology. By understanding how these trade-offs work, we may use fewer journeys and/or events, therefore being able to analyse a larger pool of customers without losing much in terms of descriptive and predictive power.

ACKNOWLEDGMENT

The authors would like to thank the YouView team for providing the sample data and domain knowledge.

REFERENCES

- [1] W. M. Van der Aalst, "Process modeling and analysis," in *Process Mining: Data Science in Action*, 2nd ed. Berlin, Germany: Springer, 2016, ch. 3, pp. 55–88.
- [2] S. Sakr, Z. Maamar, A. Awad, B. Benatallah, and W. M. P. van der Aalst, "Business process analytics and big data systems: A roadmap to bridge the gap," *IEEE Access*, vol. 6, pp. 77308–77320, 2018.
- [3] Z. Tangyuan and S. Moro, "Research trends in customer churn prediction: A data mining approach," in *Proc. World Conf. Inf. Syst. Technol.*, 2021, pp. 227–237.
- [4] A. D. Caigny, K. Coussement, and K. W. D. Bock, "A new hybrid classification algorithm for customer churn prediction based on logistic regression and decision trees," *Eur. J. Oper. Res.*, vol. 269, pp. 760–772, Sep. 2018.
- [5] N. Alboukaey, A. Joukhar, and N. Ghneim, "Dynamic behavior based churn prediction in mobile telecom," *Expert Syst. Appl.*, vol. 162, Dec. 2020, Art. no. 113779.
- [6] K. Viol, H. Schöller, A. Kaiser, C. Fartacek, W. Aichhorn, and G. Schiepek, "Detecting pattern transitions in psychological time series—A validation study on the pattern transition detection algorithm (PTDA)," *PLoS ONE*, vol. 17, no. 3, Mar. 2022, Art. no. e0265335.
- [7] F. Amato, A. Castiglione, A. D. Santo, V. Moscato, A. Picariello, F. Persia, and G. Sperlì, "Recognizing human behaviours in online social networks," *Comput. Secur.*, vol. 74, pp. 355–370, May 2018.
- [8] Z. H. Borbora and J. Srivastava, "User behavior modelling approach for churn prediction in online games," in *Proc. Int. Conf. Privacy, Secur., Risk Trust Int. Conf. Social Comput.*, 2012, pp. 51–60, doi: 10.1109/SocialCom-PASSAT.2012.
- [9] V. Kavitha, G. H. Kumar, S. V. M. Kumar, and M. Harish, "Churn prediction of customer in telecom industry using machine learning algorithms," *Int. J. Eng. Res.*, vol. V9, no. 5, pp. 181–184, May 2020.
- [10] F. Altinisik and H. H. Yilmaz, "Predicting customers intending to cancel credit card subscriptions using machine learning algorithms: A case study," in *Proc. 11th Int. Conf. Electr. Electron. Eng. (ELECO)*, Nov. 2019, pp. 916–920.
- [11] T. Vafeiadis, K. I. Diamantaras, G. Sarigiannidis, and K. C. Chatzivasvas, "A comparison of machine learning techniques for customer churn prediction," *Simul. Model. Pract. Theory*, vol. 55, pp. 1–9, Jun. 2015.
- [12] J. Burez and D. Van den Poel, "CRM at a pay-TV company: Using analytical models to reduce customer attrition by targeted marketing for subscription services," *Expert Syst. Appl.*, vol. 32, no. 2, pp. 277–288, Feb. 2007.
- [13] P. Lalwani, M. K. Mishra, J. S. Chadha, and P. Sethi, "Customer churn prediction system: A machine learning approach," *Computing*, vol. 104, no. 2, pp. 271–294, Feb. 2022.
- [14] G. Mohammadi, R. Tavakkoli-Moghaddam, and M. Mohammadi, "Hierarchical neural regression models for customer churn prediction," *J. Eng.*, vol. 2013, pp. 1–9, Feb. 2013.
- [15] Y. Huang and T. Kechadi, "An effective hybrid learning system for telecommunication churn prediction," *Expert Syst. Appl.*, vol. 40, no. 14, pp. 5635–5647, Oct. 2013.
- [16] A. Gelman, "Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper)," *Bayesian Anal.*, vol. 1, no. 3, pp. 515–534, 2006.
- [17] M. Plummer, "JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling," in *Proc. 3rd Int. Workshop Distrib. Stat. Comput.*, vol. 124. Vienna, Austria, 2003, pp. 1–10.
- [18] *R: A language and environment for statistical computing*, R. Core Team, The R Foundation for Statistical Computing, Vienna, Austria, 2013.
- [19] *Generalized Linear Models ByP, McCullagh*, J. A. Nelder Ed. New York, NY, USA: Routledge, Jan. 2019.
- [20] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
- [21] M. A. Hearst, S. T. Dumais, E. Osuna, J. Platt, and B. Scholkopf, "Support vector machines," *IEEE Intell. Syst.*, vol. 13, no. 4, pp. 18–28, Jul./Aug. 1998.
- [22] B. Schölkopf, A. J. Smola, R. C. Williamson, and P. L. Bartlett, "New support vector algorithms," *Neural Comput.*, vol. 12, no. 5, pp. 1207–1245, May 2000.
- [23] B. D. Ripley *Pattern Recognition and Neural Networks*. Cambridge, U.K.: Cambridge Univ. Press, 1996, ch. 6, pp. 181–212.
- [24] B. D. Ripley *Pattern Recognition and Neural Networks*. Cambridge, U.K.: Cambridge Univ. Press, 1996, ch. 5, pp. 143–180.
- [25] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Roy. Statist. Soc., B, Methodolog.*, vol. 58, no. 1, pp. 267–288, 1996.
- [26] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, and J. Vanderplas, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, Oct. 2011.
- [27] Z. Chen, S. Zhang, S. McClean, B. Allan, and I. Kegel, "Sequence mining TV viewing data using embedded Markov modelling," in *Proc. IEEE SmartWorld, Ubiquitous Intell. Comput., Adv. Trusted Comput., Scalable Comput. Commun., Internet People Smart City Innov. (SmartWorld/SCALCOM/UIC/ATC/IOP/SCI)*, Oct. 2021, pp. 665–670.



RAFAEL A. MORAL received the B.Sc. degree in biology and the M.Sc. and Ph.D. degrees in statistics from the University of São Paulo, Brazil, in 2011, 2014, and 2017, respectively. He is currently an Assistant Professor in statistics at Maynooth University, Ireland. He has published 60 papers in varied fields of research, including theoretical and applied statistics, computer vision, agriculture, ecology, and psychology. His research interests include statistical modeling and software development applied to ecology and wildlife management.



ZHI CHEN received the B.Eng. degree in software engineering from the Nanjing University of Information Science and Technology, China, in 2014, the B.Sc. degree (Hons.) in software systems development from the Waterford Institute of Technology, Ireland, and the M.Sc. degree in computer science from the University of Edinburgh, U.K., in 2015. He is currently a Ph.D. Researcher with the School of Computing, Ulster University, and a part-time Research Associate at the BT Ireland Innovation Centre (BTIIC). He has been working as a Software Engineer in server development and sales data analysis. In 2018, he started his Ph.D. Research at Ulster University and BTIIC, working on big data TV analytics for service and customer improvement as the major research direction. His research interests include machine learning, data mining, big data analytics, and customer behavior analysis.



GABRIEL R. PALMA received the B.Sc. degree in biology from the University of São Paulo, Brazil, in 2021. He is currently pursuing the Ph.D. degree with the Hamilton Institute, Maynooth University, Ireland, funded by the Science Foundation Ireland's Centre for Research Training in Foundations of Data Science. His research interests include statistical and mathematical modeling, machine vision, machine learning, and applications to ecology and entomology.



SHUAI ZHANG received the Ph.D. degree in data analytics from Ulster University. She is currently a Lecturer in computing science with the School of Computing and the Course Director of M.Sc. Courses in the school. Her research interests include data analytics in the area of connected health applications, including activity and behavioral recognition in smart environments, change-point detection for sensor data annotation, and modeling user engagement with and adoption of assistive technologies for people with dementia, more recently on commercial customer and services analytics. She is a Grant Holder for projects funded from H2020, the Royal Society, Wellcome Trust-Wolfson, and Industrial funding.



BRAHIM ALLAN received the B.Eng. and M.Sc. degrees in telecommunication engineering from the King's College London, University of London, U.K., in 2007 and 2009, respectively. He joined the British Telecom Research Centre (BT), in 2011, as a Researcher, working in the future content and application services team. He holds four patents and has published eight papers. He is a member of BSI—British Standards Institution and Represent BT in Moving Picture Experts Group (MPEG)

Standard Meetings.



SALLY MCCLEAN (Member, IEEE) received the degree in mathematics from Oxford University, the M.Sc. degree in mathematical statistics and operational research from Cardiff University, and the Ph.D. degree in Markov and semi-Markov models from Ulster University. She is currently a Professor in mathematics at Ulster University. She has published over 300 research papers and was previously a recipient of Ulster University's Senior Distinguished Research Fellowship. Her research interests include stochastic modeling and optimization, particularly for healthcare planning and computer science, specifically databases, the Internet of Things, sensor technology, and telecommunications. Much of this work, she has been focused on healthcare applications, particularly with regard to patient modeling and pervasive technologies. She has been grant holder on over €13 million worth of funding, mainly from the EPSRC, industry, the EU, and charities. She is a fellow of the Royal Statistical Society, a fellow of the Operational Research Society, a fellow of IMA, and the Past President of the Irish Statistical Association.



IAN KEGEL is a Chartered Engineer who studied Electrical and Information Sciences at the University of Cambridge. He has worked in both the defence and telecommunications industries on projects ranging from radar signal processing to multimedia delivery and has spent over 18 years leading applications-related research within BT. He co-ordinates research in the fields of personalized media production and delivery, TV data analytics, and immersive communications, whose results are helping BT to develop differentiated products and improve customer experience for its consumer and enterprise customers. Open Innovation is core to his work, built around successful collaborations with leading industrial and academic organizations across Europe.

...