Measuring Spontaneous Devaluations in User Preferences

Komal Kapoor, Nisheeth Srivastava, Jaideep Srivastava
Department of Computer Science
University of Minnesota
200 Union St SE
Minneapolis, MN 55455
kapoo031@umn.edu, nisheeths@gmail.com, srivasta@cs.umn.edu

Paul Schrater
Department of Psychology and Computer Science
University of Minnesota
75 East River Road
Minneapolis, MN 55455
schrater@umn.edu

ABSTRACT

Spontaneous devaluation in preferences is ubiquitous, where yesterday's hit is today's affliction. Despite technological advances facilitating access to a wide range of media commodities, finding engaging content is a major enterprise with few principled solutions. Systems tracking spontaneous devaluation in user preferences can allow prediction of the onset of boredom in users potentially catering to their changed needs. In this work, we study the music listening histories of Last.fm users focusing on the changes in their preferences based on their choices for different artists at different points in time. A hazard function, commonly used in statistics for survival analysis, is used to capture the rate at which a user returns to an artist as a function of exposure to the artist. The analysis provides the first evidence of spontaneous devaluation in preferences of music listeners. Better understanding of the temporal dynamics of this phenomenon can inform solutions to the similarity-diversity dilemma of recommender systems.

Categories and Subject Descriptors

I.6 [**Simulation and Modeling**]: Applications

General Terms

Human Factors

Keywords

Dynamic Preferences, Recommender Systems, Temporal Models, User Behavior Modeling

1. INTRODUCTION

Recommendation systems have become a popular means of suggesting relevant content to the user. Methods in recommendations have focused on constructing estimates of user preferences based on their history of choices. These preference estimates are then used to suggest new content to the user using content-based or collaborative methods. Content-based methods use a user's preference estimates to find similar content, while collaborative methods use a user's preference estimates to identify similar users (neighborhood) and recommend content popular in the identified neighborhood. But, it's not sufficient for a recommender agent to only estimate a user's past preferences; it's also important to predict their future preferences given past experiences. This makes the task of a recommender even more challenging by requiring it to predict when and how a user's preferences will change in the future. The recommendations community, however, lacks models which can predict changing preferences of users and doing so is generally accepted as a hard problem. On the other hand, user's recent choices have been found to be a good predictor of their future behavior. Efforts in modeling temporal recommendations have exploited this aspect of user choices by designing recommendation systems which systematically emphasize recency with good results. The critical shortcoming of this formulation is that such a system merely reacts to preference *changes* rather than trying to predict them.

While little work has been done on predicting changes in user preferences in the recommendation literature, psychologists and behaviorists have long studied the dynamics of individual preferences. Several theories have been proposed to explain why individuals seek out new content (novelty seeking, exploratory and information seeking behavior) [2]. Other studies talk about individuals making choices to actively seek an optimal level of stimulation in their environment [21]. The theory of flow [20] suggests that an environment which provides an optimal level of challenge for a given level of skill leads to a desirable state of flow. Despite such theoretical developments, it has been difficult to operationalize these aspects of individual choices to solve real world problems. However, modeling properties of individidual behavior is critical for advancing designs of automated agents which interact with individuals on a daily basis.

In this work, we study one aspect of dynamic individual preferences. Individuals are often found to develop disinterest and even dislike for their dearly preferred content both temporarily and lastingly. It's common to find that one's clothes, food, entertainment, jobs etc. have grown boring despite being enjoyable in the past. We call this phenomenon a spontaneous devaluation of one's preferences or boredom for a stimulus. Spontaneous devaluation is seen to arise when repeated exposure to a stimulus creates a feeling of satiation towards it leading to a loss in interest [7]. Alternatively, spontaneous devaluation has been linked to lost opportunity for novel experiences when similar experiences are repeated too often [18]. Both theories concur in suggesting that, in contrast to recency-based expectations, *repeated exposure to familiar choices spontaneously devalues one's preference for them.*

Human behavior driven by these dynamics could be modeled as systematically alternating between one's set of choices, assuming that the time spent in experiencing other stimuli is sufficient to mitigate the effects of boredom for a particular stimulus. Several studies on user purchase behavior have found buyers to alternate among their preferred alternatives [11, 18, 9] etc. However, in practice users have a non-uniform liking for different alternatives in their choice space. Furthermore, users have a pronounced tendency to stick to their recent choices [11] which has been responsible for the success of the previously proposed recommender models. We call this behavior the *'sticky'* behavior in users. This phenomenon has also been called reinforcement or inertial behavior. Such behavior can be explained to arise due to an actual increase in liking on exposure [9] or a tendency to avoid switching costs.

The presence of both stickiness and devaluation effects in user preferences make predicting the temporal choices of a user non-trivial. In this paper, we analyze user music listening behavior to extract signals of stickiness and boredom. Our analysis is limited to the music domain due to availability of public datasets, nevertheless, we expect our results to generalize to other items like movies, videos, books, vacation packages, shopping etc. which are fairly susceptible to boredom effects. We demonstrate the use of hazard functions for measuring these phenomena. Our work provides the first proof of spontaneous devaluation in music listening preferences of users and its impact on user choices. This work can inform design of future methods that incorporate these dynamics, producing agents that can cater to new needs of users suffering from boredom.

The rest of the paper is organized as follows: Section 2 provides a summary of the related work. Section 3 gives an overview of the dataset and pre-processing details. Section 4 lays out terminology relevant to our analysis. Section 5 provides details of our methodology. Our results are summarized in Section 6. We end with a discussion of the contributions of this work and possible future extensions in Section 7.

2. RELATED WORK

2.1 Dynamic Preferences

Stimulus satiation was initially used by researchers to explain spontaneous alternation in rats [7]. Rats were placed in a T-shaped maze and provided an unlimited supply of food at the left and the right corners of the maze at equal distances. The experiment was set up such that that the rat had to return to the starting point before each trial. It was seen that rats chose to alternate between the left and the right ends on repeated trials. Glanzer [10] suggested that such a behavior arose due to stimulus satiation such that each time the organism was exposed to the stimulus, satiation for the stimulus increased causing the rat to switch directions. Further, satiation for the stimulus diminished when the organism could no longer perceive the stimulus and the rat returned back to the same direction.

Researchers have found individuals to engage in more complex forms of variety seeking behavior while making choices. McAlister proposed a taxonomy of factors responsible for varied behavior in individuals [18]. These were classified into two categories based on whether they arose due to external factors (such as unavailability of a product, launch of new products etc.) or due to internal motivations. When arising out of internal motivations, variety seeing behavior was suggested to manifest in two forms; a desire for unfamiliar alternatives or a desire to alternate among familiar alternatives. The former was linked to individuals seeking an optimal level of stimulation [2, 21], while, the latter was seen as a weak form of exploratory behavior. It was also linked to devaluation in preferences due to satiation. A single peaked preference function was proposed to characterize the attractiveness of a stimulus on repeated exposure [6]. McAlister also proposed a dynamic attribute satiation model [17] which assumed an ideal level of inventory for different attributes of the items. The inventory was designed to dwindle over time to incorporate the effects of forgetting.

Researchers have subsequently focused on modeling the choice probabilities of consumers directly given their past choices. Consumers were found to exhibit either a short term loyalty for their last purchased brand (inertia) or devaluation for the last purchased brand (variety seeking) [11, 9]. Kahn [12] compared seven models for user choice behavior with similar results. Bawa et al [1] used a single peaked function, to model the conditional probability of repeat purchase given the number of times the brand was re-purchased since user's last switch (run length). Chintagunta [5] used hazard rates to model the level of inertia and variety seeking as a function of time between purchases. Recent efforts have expanded these models to incorporate heterogeneities between consumers and external environment variables affecting user choices [13].

Most of the research in this area, however, has been limited to panel datasets and analysis of user surveys and questionnaires. In this work we have adopted a data driven approach to elicit changes in user preferences towards a stimulus as a function of their past exposure to it. Our efforts do not look at variety seeking or inertial behavior in users in general, but at changes in choice probabilities with respect to particular stimulus, grounding ourselves in psychological theories of boredom and novelty seeking, which provides a causal explanation for the existence of these patterns.

2.2 Recommender Systems

State-of-the-art methods in recommender systems have assumed a static view of human preferences. Ding et al. [8] showed that the static view of user preferences used while generating recommendations was flawed as it did not take changing user interests into account. They used a decay function to gradually devalue the impact of a user's past history while making prediction of his future likings. Recently, a temporal model of recommendation was developed [15, 14] which was an important part of the solution to the KDD Cup on Yahoo Music dataset and the Netflix challenge. The model incorporated several time-sensitive user and item biases in the standard factor model. Gradual changes in user preferences over time were captured using a linear function. Their model showed that modeling temporal dynamics in user choices was essential for improving the performance of the recommender. Sahoo [22] has proposed a dynamic model of blog reading behavior in employees. He used a Hidden Markov Model to predict future interests of employees based on their previous choices. However, user transitions are assumed to be driven by a static transition matrix. At present, the recommendation community lacks models that predict *changes* in user preferences.

Also related to our work are methods to introduce diversity and novelty in recommendations. Lathia et al. [16] showed that popular recommendations methods such as kNN and SVD produced recommendations which were very similar (low in temporal diversity) on iterated train-test experiments on temporally ordered data. Many methods that systematically introduce diversity in the recommendations have been proposed [19, 23, 24, 3]. However, these methods focus on jointly optimizing both similarity and diversity indices described on the space of items being recommended rather than predicting changes in user preferences.

3. DATA

Our analysis is based on complete temporal music listening histories of users provided by Last.fm. Last.fm is a popular music website with millions of active users. It allows users to purchase tracks, listen to online radios and playlists etc. and has additional social networking features as well. Recently, Last.fm made available a dataset of complete music listening histories of around 1000 users as recorded till May 2009 [4]. This is the only publicly available dataset, to our knowledge, to provide complete temporal records of user choices. Because Last.fm hosts several online radios, it is quite probable that parts of the user histories capture radios, and playlists rather than active user choices. We filtered these effects by using the time gap between two consecutive tracks played by the user. Last.fm has a generous list of API's available to developers. The API, track.getInfo, was used to retrieve the duration of most of the songs in our dataset. We compared the time gap between song 1 and song 2 in that temporal order in the user history with the length of song 1. If the time gap was found to be more than the length of song 1 by less than 5 seconds, song 2 was identified to belong to an automated play list. All tracks 'not on *auto-play*' were assumed to be *active* user choices. We could not remove auto-play effects for the songs whose lengths were unavailable through the API. This corresponded to 0.05% of the songs. We only considered the first 1 year of each user history in our analysis. All the users which had less than 30 records of activity were eliminated from the dataset. Also, we only kept those artists in the user history which the user had listened to 15 or more times in that period of 1 year. We summarize some important statistics about the dataset in Table 1.

Property	**Value**
# unique tracks	1,084,872
# unique artists	174,091
# Users	957
Mean history length - # songs heard	6716
Mean history length - # active days	177
Mean # unique artists heard	37

Table 1: Statistics from the Last.fm dataset

4. TERMINOLOGY

Based on both the novelty-seeking and stimulus satiation theories of devaluation of preferences, repeated exposure to a stimulus causes devaluation in one's preferences towards it. Additionally, devalued preferences can get reinstated after a period of reduced or no exposure. A music piece can stimulate the listeners because of the combined effect of its multiple features (artist, genre, tempo, strong female vocals, etc.). For simplicity and ease of access, we use the artist of the songs as our basic stimulus. More sophisticated stimulus definitions that model the interaction between multiple features of a song can enhance our method.

Preferences have been linked to choice probabilities in the past. It is only a logical extension to relate changes in preferences to changes in choice probabilities, and in our case conditional choice probabilities. We suspect that the phenomenon of devaluation produces two different patterns in the choice probabilities of users for an artist.

Hypothesis 1: The probability that a user will listen to an artist again will decrease after he has listened to the artist some number of times. When this happens, we say that the user's preferences for the artist have devalued.

Hypothesis 2: Devalued preferences can get reinstated after a sufficient period of non/reduced exposure to the artist.

Through our experiments, we look for signals suggestive of spontaneous devaluation in choices probabilities of Last.fm users. By doing so, we establish a methodology for detecting this phenomenon and analyzing its properties.

We consider the state of the user at some time t to be defined by the artist of the song the user was listening to at that time. The temporal history of the user comprises the sequence of states visited by him as a function of time; i.e. $H^u(t) = s_a$ if user u was listening to artist a at time t. User u is said to *enter* a state a at time t if $H^u(t) = s_a$ and $H^u(t-1) \neq s_a$. A user u is said to *exit* a state a at time t if $H^u(t) \neq s_a$ and $H^u(t-1) = s_a$. We can now define the following conditional choice probabilities:

1. **Conditional probability of exit**: This is the conditional probability of a user u exiting state a at time t given that he last entered state a at time $t-r$ and has not exited state a yet. Formally, the probability is equal to $P(H^u(t) = s_a | H^u(t-1) = s_a, \ldots, H^u(t-r) = s_a, H^u(t-r-1) \neq s_a)$. Here, r is the time spent listening to the artist and corresponds to the idea of a run length in Bawa's model [1]. We make the simplifying assumption that this probability depends only on r. Hence, we can also represent the conditional prob-

ability of exiting state a by user u when time spent in state is r as $P^{ua}(\text{exit}|\text{time spent in state } a = r)$.

2. **Conditional probability of entry**: This is the conditional probability of user u entering a state a at time t given that the user last exited state a at time $t-(o+1)$. Formally, this corresponds to $P(H^u(t) = s_a | H^u(t-1) \neq s_a, \ldots, H^u(t-o) \neq s_a, H^u(t-o-1) = s_a)$. Here, o is the time spent not listening to the artist a. Again, for simplicity, we assume that this probability depends only on o. We later relax this assumption with interesting effects, described in Section 6.3. Thus, this probability can also be represented as the conditional probability of entering state a after having exited it o units of time ago or $P^{ua}(\text{entry}|\text{time spent out of state } a = o)$.

The definition of time has been kept ambiguous in the definitions above. We now define it more formally. Time can be defined in terms of the order in which songs are heard by the user such that $H^u(t)$ refers to the t-th song heard by user u. Such a definition, however, does not take the actual time gap between consecutive listenings into account. It is important to consider the actual time gap between user choices. This is because a user satiated with an artist can get unsatiated both by listening to other artists or due to forgetting if he returns to the system after a long time. To analyze the impact of actual clock time on the satiation level, we define time in terms of days since the first historical record of the user. Accordingly, $H^u(t)$ refers to the state of the user on t-th day since day 1. For simplicity, the state of the user on a day is defined by the artist listened to most frequently by him on that day.

5. METHODOLOGY

Survival Analysis is a statistical method commonly used for modeling time-to-event data. The purpose of this kind of analysis is to model the probability of survival (where the occurrence of the event corresponds to death) beyond a certain point in time. For simplicity, we use a discrete measures of time $t \in \mathbb{N}$. The survivor function at time t is defined as:

$$S(t) = P(T > t) \tag{1}$$

Where, T is a random variable denoting the time of death. The instantaneous rate of occurrence of the event at time t, conditioned on having survived up to time t, is captured using the hazard function. The hazard function is also called the conditional failure rate and is defined as:

$$\lambda(t) = \lim_{\Delta t \to 0|} \frac{P(t \leq T < t + \Delta t | T \geq t)}{\Delta t} = -S'(t)/S(t) \tag{2}$$

We use the hazard rate function to compute the exit and entry conditional probabilities defined in the previous section. We set $\Delta t = 1$. This allows us to use the terms hazard rate and conditional probability of death interchangeably. We can construct the two different hazard curves based on how we define our events.

1. **Exit Hazard Rate:** Here, we measure time from the point when a user u entered a state a. The event corresponds to his 'exit' from the state. The random variable T^{ua}_{exit} denotes the time of exit or death. This hazard rate captures the conditional probability of exiting the state at time $t+1$ having survived in the state for time t or greater; $\lambda^{ua}_{\text{exit}}(t) = P^{ua}(T^{ua}_{\text{exit}} = t | T^{ua}_{\text{exit}} \geq t)$.

2. **Entry Hazard Rate:** Here, we measure time from the point when a user u exited a state a. The event corresponds to his 'entry' back into the state. The random variable T^{ua}_{entry} denotes the time of entry or death. This hazard rate captures the conditional probability of entering a state at time t having survived outside the state for time t or greater;
$\lambda^{ua}_{\text{entry}}(t) = P^{ua}(T^{ua}_{\text{entry}} = t | T^{ua}_{\text{entry}} \geq t)$.

An exit and entry hazard rate can be defined for each artist a user listens to. For our analysis, we pool across the different users and the artist choices to compute an average exit and entry hazard rate for the entire dataset. We normalize the time of entry and exit variables to mitigate the effects of differences in a user's preferences for different artists and differences across users. The time of event variable is log transformed as well as it becomes harder to exactly predict the time of an event as time for which the event has not happened increases. In other words, this means that if a user has not returned to an artist in a month, its more difficult to predict the exact day of his return, than, when he has has not returned to the artist for a day. The log transform accommodates this non-linearity in the predictability of return time.

$$T^{\mathrm{N}}_i = \frac{\log_2(T^{ua}_i)}{\log_2(\frac{1}{P^u(a))}} \tag{3}$$

for a user u and artist a and $i \in \{\text{'}entry\text{'}, \text{'}exit\text{'}\}$. $P^u(a)$ is the prior probability of user u being in state a.

$$P^u(a) = \frac{N^u(a)}{L^u} \tag{4}$$

where, $N^u(a)$ is the number of times user u was in state a and L^u is the length of user u's history. The average hazard rates for the normalized time of event variable can then be computed across users and artists:

$$\lambda_i(t) = P(T^{\mathrm{N}}_i = t / T^{\mathrm{N}}_i \geq t) \tag{5}$$

We discretize t into intervals $(0, 0.1], (0.1, 0.2]$ and so on. The hypothesis presented by us in section 4 can now be represented using the hazard rates.

1. **Hypothesis 1** The exit hazard rate for an artist should be an increasing function of time. This indicates that a user's preferences for an artist decrease with increased exposure to the artist.

2. **Hypothesis 2** The entry hazard rate for an artist should be an increasing function of time. This indicates that user preferences for the artist are reinstated after sufficient time gap.

The *sticky* or inertial view of user choices, on the other hand, suggest that a user's probability of visiting a state would increase on having visited it. Contrary to the devaluation hypothesis, the conditional probability of visiting a state again would increase as time spent in the state increases. This implies that the exit hazard rate for an artist is a decreasing function of time for sticky users. The entry hazard rate, would also be a decreasing function of time as a user would be less likely to visit a state which they has not visited for long periods of time.

A common analysis methodology is to compare the hazard rate of interest in an analysis with that generated from a

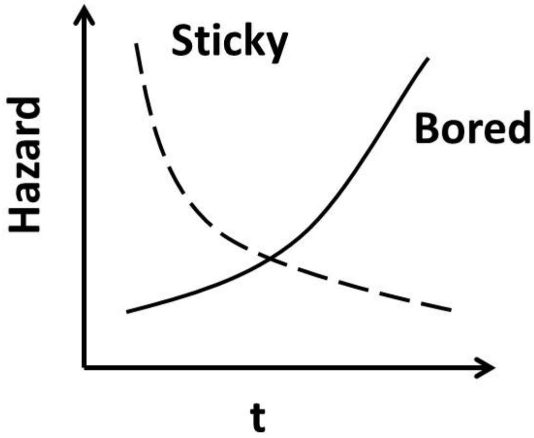

(a) Expected hazard rate for a sticky and boredom-prone user

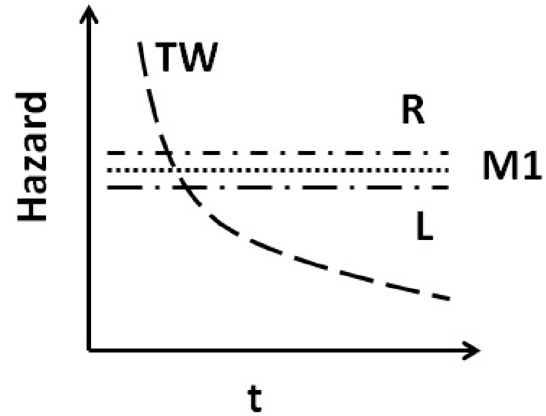

(b) Expected hazard rates for the baseline models

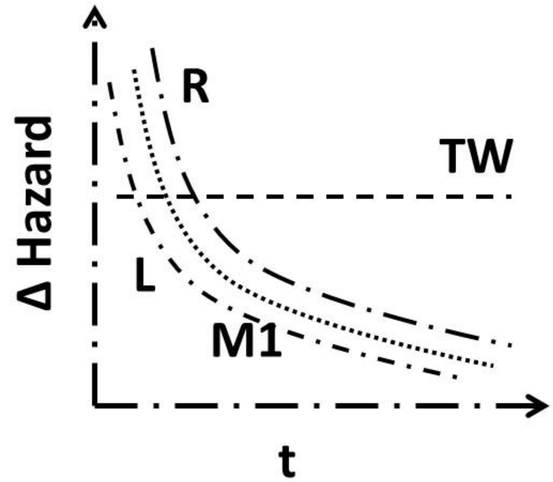

(c) Expected Δ Hazard Rates for sticky users

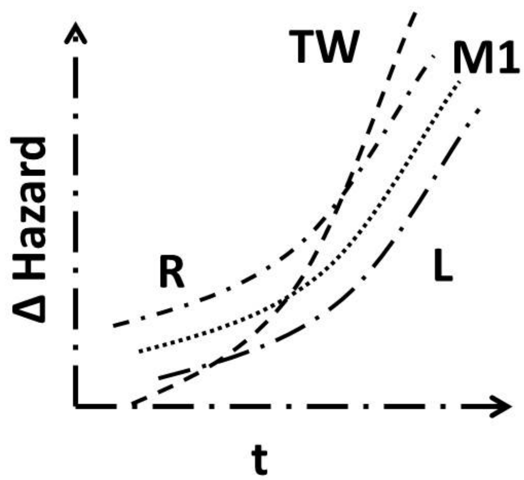

(d) Expected Δ Hazard Rates for boredom-prone users

Figure 1: Figure (a) and (b) depicts the expected hazard rates for sticky and boredom-prone users and the baseline models. Both the entry and exit hazard rates should decrease with time for sticky users and increase with time for uses susceptible to boredom. Figure (c) and (d) shows the expected Δ hazard rates computed against each baseline model for sticky and boredom-prone users.

control experiment. This is done to remove the effects of covariates not being considered in the analysis. We define four baseline models to serve as controls. We constructed listening sequences by simulating user histories using each of the baseline models for every user. The user histories were simulated by sampling randomly from the temporal preference vector (Pref) generated by each of the model. In order to make the baseline models as close to the real data as possible, the parameters of the models were fitted to the actual user histories.

1. **Random (R)** The user is assumed to sample states randomly from his average preference vector (P^u). $Pref^u(t) = P^u$

2. **1st order Markov (M1)** A user's switching probability from one state to the other is assumed to be controlled by a 1st order Markov model. The dynamics of the Markov model are controlled by a static transition matrix (T^u) which is learnt for each user u's history using maximum likelihood estimation. $Pref^u(t) = Pref^u(t-1) * T^u$

3. **Time weighted (TW)** We use a recency based model for generating user histories. $Pref^u(t) = \alpha^u * Pref^u(t-1) + c^u(t-1)$, where, $c^u(t-1)$ is $1 * |A|$ choice vector, which is set to 1 at index i if $H^u(t-1) = s_i$, and is 0 otherwise. The parameter α^u is a |A|*1 vector which was fit to the user u's history using stochastic gradient descent. We introduced a small exploratory component to this model to prevent extremely long lengths of continuous listening of the same artist. Therefore, our modified preference vector is computed as $Pref'^u(t) = 0.95 * Pref^u(t) + 0.05 * P^u$

4. **Linearly increasing or decreasing (L)** We used the temporal model of user preference used by Koren [14]. $Pref^u(t) = P^u + \text{sign}(t - L^u/2) * (t - L/2)^{\beta^u}$. The parameter β^u is a |A|*1 vector and was fitted to the user u's history using stochastic gradient descent.

The Log-Rank test can be used to test whether the survival distributions generated by the simulated models are sufficiently different from that of the real data. The hypothesis test is defined as:

H_o**:** The real data and the simulated data have different survivor function

H_a**:** The real data and the simulated data have the same survivor function

The Log-Rank test on the real and the simulated survival functions rejects the null hypothesis with a p-value $< 10^-6$. The discrepancy between the real data and the baseline model predictions can be quantified using a Δ hazard rate obtained by subtracting the simulated hazard rates from the hazard rates computed on real data.

$$\lambda^{\Delta}(t) \;=\; -\frac{S'^{\text{real}}(t)}{S^{\text{real}}(t)} \;-\; -\frac{S'^{(\text{simulated})}(t)}{S^{(\text{simulated})}(t)} \qquad (6)$$

We generate four Δ hazard rates for both the entry and exit time events for our analysis, namely real vs. random (λ_i^{A-R}) , real vs. Markov (λ_i^{A-M1}), real vs. time weighted (λ_i^{A-TW}) and real vs. linear (λ_i^{A-L}), where $i \in \{`entry', `exit'\}$.

We display the entry and exit hazard rates expected for the event times obtained from the 'sticky' and 'boredom-prone' models and those expected from the baseline models in Figure 1. The entry and the exit hazard rates for a random, markovian and linear model should be independent of time spent in the state. A TW model on the other hand, is essentially a sticky model. Hence, the exit and entry hazard rates for TW model would decrease with time. The objective of this study is to understand the form of the exit and entry hazard rates for the real data. Figure 1 displays the expected Δ hazard rates if the real data follows the sticky and the boredom-prone model, respectively.

6. RESULTS

In this section we examine the obtained Δ exit and Δ entry hazard rates in close detail.

6.1 Δ Exit Hazard Rates

Figure 2 displays the survivor functions for the exit time for the real data and data generated by each simulated model. It also depicts the obtained Δ exit hazard rates. The changes in $\lambda_{\text{exit}}^{A-R}$, $\lambda_{\text{exit}}^{A-M1}$ and $\lambda_{\text{exit}}^{A-L}$, directly represent changes in the λ_{exit} for the real data. Changes in $\lambda_{\text{exit}}^{A-TW}$ would depict changes in the exit hazard rate for real data against a decreasing baseline.

1. Real Vs.Random, Markov and Linear models: The $\lambda_{\text{exit}}^{A-R}$ and $\lambda_{\text{exit}}^{A-M1}$ are negative throughout suggesting that the exit rate for the real data is lower than that expected for the baseline models. This supports the sticky view of user preferences suggesting that a user has a lower rate of exiting a state after having visited it. However, contrary to what is expected for the sticky model, the Δ exit hazard rate increases with time after a point. We expect the Δ hazard rate to eventually flatten out, becoming uninformative. The survival function for R, M1 and L models drops sharply indicating a lower probability for large sequences than those observed in the real data. The L model has the sharpest drop in survival probability, such that we did not enough samples of exit times greater than 0.1.

2. Real vs. Time-Weighted model: $\lambda_{\text{exit}}^{A-TW}$ is negative for low values of t, suggesting larger stickiness in users than generated by the TW model. However, the Δ exit rate increases thereafter, becoming positive after some time. Since, the exit hazard rate for the TW model is expected to decrease with time, this suggests that the exit hazard rate for real data increases more than the decrease observed in the TW model.

From these observations we can conclude that users have high stickiness towards the state on entering the state. However, the stickiness for a state reduces with time and the dynamics driven by boredom start dominating as time spent in the state increases. A user is thus likely to stick to his previous state at a higher rate initially and a decreased rate as time in the state increases.

6.2 Δ Entry Hazard Rates

Figure 3 displays the survivor functions computed for the entry time variable for real and simulated data and the obtained Δ entry hazard rates. Similar to the Δ exit hazard rates, the changes in $\lambda_{\text{entry}}^{A-R}$, $\lambda_{\text{entry}}^{A-M1}$ and $\lambda_{\text{entry}}^{A-L}$ functions would depict changes in the entry hazard rate for the actual data. The TW model is expected to have a declining entry hazard rate, being a sticky model. The changes in $\lambda_{\text{entry}}^{A-TW}$ should reflect changes in the entry hazard rate for the real data against a decreasing baseline.

1. Real Vs.Random, Markov and Linear models: The $\lambda_{\text{entry}}^{A-R}$, $\lambda_{\text{entry}}^{A-M1}$ and $\lambda_{\text{entry}}^{A-L}$ functions are positive initially suggesting that the users have a higher rate of entry than that expected from the baseline models. This again can be attributed to the sticky nature of user choices, such that users have a high rate of returning to the artists they had listened to recently. The Δ hazard rates decrease for intermediate values of t suggesting a prominent devaluation in preferences. The Δ hazard rates eventually increase for larger values of t. However, they do not cross the 0-line again suggesting that a user always has a lower rate of return than that generated by the baseline models. This can be attributed to phasing out of an artist who is not being actively sampled.

2. Real vs. Time-Weighted model: The $\lambda_{\text{entry}}^{A-TW}$ function is slightly negative at the beginning suggesting that the actual entry hazard rate is lower than that of a TW model. Our TW model is seen to pull back users which have just left an artist at a higher rate than observed in real data. The hazard rate increases thereafter indicating the actual data seems to have a larger rate of return than that of the TW model.

The analysis on the Δ entry hazard rates reveals aspects of sticky behavior in users which produces quick switches in and out of the artist. Also, we find indicators of devalued preference for intermediate values of time spent out of

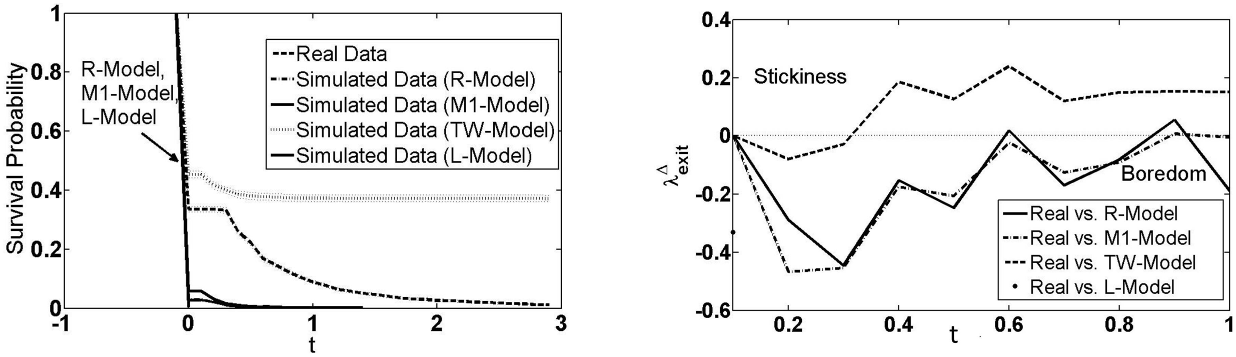

(a) Kaplan-Meier survival functions and 95% confidence interval

(b) Nelson-Aalen Δ exit hazard functions

Figure 2: The figure illustrates the survival and the hazard functions computed for the exit time variable. The negative Δ exit rates for low values of t are indicative of sticky behavior, while the increase in Δ exit hazard rate indicate a devaluation in preference.

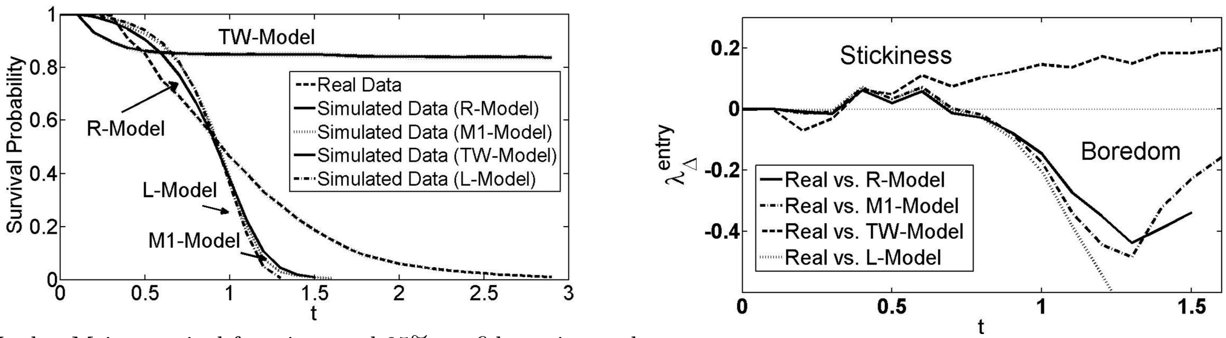

(a) Kaplan-Meier survival functions and 95% confidence interval

b) Nelson-Aalen Δ exit hazard functions

Figure 3: This figure illustrates the survival and the hazard functions computed for the entry time variable. The Δ hazard rates are positive for all the model for low values of t which is indicative of sticky behavior. A decline in the Δ entry hazard rates corresponding to the R, M1 and L models for intermediate values of t indicate that the preferences were temporally devalued. The increase in the Δ entry hazard rates corresponding to all the models for larger values of t suggest that preferences were reinstated

the state. Preferences are reinstated after longer periods of time spent away from the artist, however, the rate of return eventually flattens out becoming uninformative.

6.3 Previous Return Time

In our previous analyses, we found evidence suggesting that users quickly switch in and out of an artist in a short span of time. Such a characteristic of user temporal choices suggest that a user's level of exposure to an artist is not completely defined by the 'in time'. A user who has just switched out of the artist and has switched back in almost immediately after, somewhat continues to be in state a. Therefore, we suspect that the previous return time (PRT) $T^{N,P}_{\text{entry}}$ also indicates how much a user has been exposed to the artist recently. A low PRT indicates higher exposure to the artist than a larger PRT. A corollary to hypothesis 1 in terms of the $T^{N,P}_{entry}$ for the artist follows:

Corollary 1' The probability that a user listens to an artist again will depend on his PRT to the artist. We suspect that of if the user has returned to the artist quite quickly previously, he will have a lower rate of returning quickly to the artist in the future.

In order to test this hypothesis we generate two conditional entry hazard rates.

1. $\lambda^{LP}_{\text{entry}}$ Entry Hazard Rate given a low PRT, $T^{N,P}_{\text{entry}} < 1$
2. $\lambda^{HP}_{\text{entry}}$ Entry Hazard Rate given a high PRT, $1 < T^{N,P}_{\text{entry}} < 1.5$

We compute the Δ hazard rate for the two conditional entry hazard rates.

$$\lambda^{\text{LP-HP}}_{\text{entry}} = \lambda^{LP}_{entry} - \lambda^{HP}_{entry} \tag{7}$$

$\lambda^{\text{LP-HP}}_{\text{entry}}$ function is computed for the real data and data simulated using a Markov model. The simulated data serves as a comparison. Figure 4 displays the obtained $\lambda^{\text{LP-HP}}_{\text{entry}}$ functions and the survival functions for $\lambda^{\text{LP}}_{\text{entry}}$ and $\lambda^{\text{LP}}_{\text{entry}}$ for the real data and simulated data. The log rank test is rejected with a p-value of less than $10^{-}4$ on the conditional

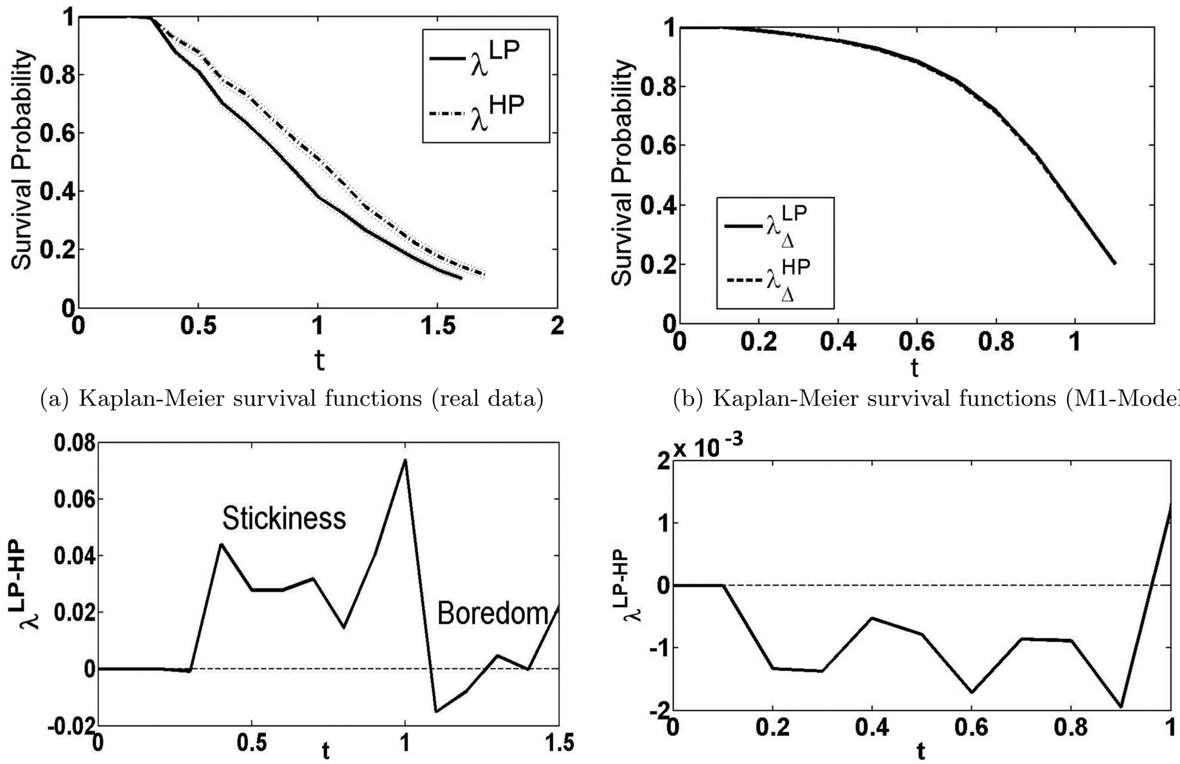

(a) Kaplan-Meier survival functions (real data)

(b) Kaplan-Meier survival functions (M1-Model)

(c) Nelson-Aalen Δ exit hazard functions (real data)

(d) Nelson-Aalen Δ exit hazard functions (M1-Model)

Figure 4: This figure illustrates the survival and the hazard functions computed for the entry time variable conditioned on the PRT. The conditioned survival function for the simulated data are coincident but vary significantly for the real data. The positive values of the Δ hazard function $\lambda_{\mathbf{entry}}^{LP-HP}$ for low values of t indicate an increased stickiness when conditioned on lower values of PRT. The negative values of the Δ hazard function for larger values of t are indicative of increased boredom effects when conditioned on lower values of PRT.

survival functions of the simulated and the real data. However $\lambda_{\text{entry}}^{\text{LP-HP}}$ varies by very small amounts. On the contrary, $\lambda_{\text{entry}}^{\text{LP-HP}}$ on the real data varies in an interesting way. We see that $\lambda_{\text{entry}}^{\text{LP-HP}}$ is highly positive initially, which indicates increased stickiness when PRT is low. However, $\lambda_{\text{entry}}^{\text{LP-HP}}$ decreases and becomes negative eventually which indicates a lower rate of return for larger values of t when PRT is low than when PRT is high. Hence, once a user is out of the state he has a lower rate of returning back to the state when previous return time is low than rate of return for a user-artist pair for whom previous return time was high.

7. DISCUSSION

In this work we have outlined a methodology for analyzing music listening histories of Last.fm users for studying the phenomenon of spontaneous devaluation in user preferences or boredom. We constructed hypothesis about boredom-prone behavior in Last.fm users and tested them through experiments on real and simulated data. Exploratory analysis of dynamic hazard rates computed on both the real and simulated data suggest that real data has strong evidence of spontaneous devaluation of preferences, as hypothesized. We also found strong evidence suggesting stickiness or reinforcement nature of past choices in users. Crucially, stickiness and boredom effects on user choices were found to be spaced out in time suggesting that methods can be designed to systematically appease the two driving forces effecting user temporal needs. The results obtained from this analysis motivate the design of sophisticated dynamic models of user choices impacting recommendation methods, product design and advertising.

Our findings suggest that methods which only focus on maximizing similarity, or focus on maximizing both similarity and diversity at all times, accommodate only some aspects of user behavior, leaving useful *temporal* information on the table. Sophisticated temporal models of individual preferences, well grounded in cognitive and psychological analysis of the dynamics of their choices, are required for the design of automated methods that can predict user temporal needs well.

Being able to say *when* a user is likely to be bored should yield considerably more responsive and accurate product recommendations. However, the gap between this exploratory analysis and usable applications, while bridgeable, is non-trivial. We suspect heterogeneities to exist among users and their behavior towards different items, which this analysis has not considered. This is principally because extricating good estimates of dynamic hazard rates for different user-item pairs requires large amounts of historical data, while we were limited in our analysis to the Last.fm publicly release dataset. Unavailability of datasets providing complete temporal histories of users makes procurement of data a challenge. While gaining access to more data would be the best solution, clustering methods can reduce the data scarcity problem in the interim. Additionally, for simplicity, we have assumed the user behavior for an item is independent of the other items experienced by him. However, one can expect similar/dissimilar items to increase/decrease one's level of satiation with an item. Extending our approach into a full-fledged recommendation system would require us to address user and item level heterogeneities and similarities between items in a single framework. Potential solutions can benefit from hierarchical approaches to cluster items using multiple features allowing estimation of the impact of history on the hazard rates for similar items.

Our work constitutes the first study on dynamics of preferences of online music listeners, and demonstrates that there is significant value in trying to study the temporal browsing history of users along the lines we have suggested. We hope our work will motivate further studies on this topic in the future. Also, larger datasets would be made accessible for studying aspects of user choices, allowing advancement in the design of predictive agents of temporal user choices.

8. ACKNOWLEDGEMENT

The research reported herein was supported by the Social Media and Business Analytics Collaborative (SOBACO) at the University of Minnesota. We would like to thank the funding agency for their support.

9. REFERENCES

[1] K. Bawa. Modeling inertia and variety seeking tendencies in brand choice behavior. *Marketing Science*, 9(3):263–278, 1990.

[2] D. E. Berlyne. Conflict, arousal, and curiosity. 1960.

[3] O. Celma. *Music Recommendation and Discovery in the Long Tail*. Springer, 2010.

[4] O. Celma. Music recommendation datasets for research. 2010.

[5] P. K. Chintagunta. Inertia and variety seeking in a model of brand-purchase timing. *Marketing Science*, 17(3):253–270, 1998.

[6] C. H. Coombs and G. S. Avrunin. Single-peaked functions and the theory of preference. *Psychological Review*, 84(2):216, 1977.

[7] W. N. Dember and H. Fowler. Spontaneous alternation behavior. *Psychological Bulletin*, 55(6):412, 1958.

[8] Y. Ding and X. Li. Time weight collaborative filtering. In *Proceedings of the 14th ACM international conference on Information and knowledge management*, pages 485–492. ACM, 2005.

[9] M. Givon. Variety seeking through brand switching. *Marketing Science*, 3(1):1–22, 1984.

[10] M. Glanzer. The role of stimulus satiation in spontaneous alternation. *Journal of experimental psychology*, 45(6):387, 1953.

[11] A. P. Jeuland. Brand choice inertia as one aspect of the notion of brand loyalty. *Management Science*, 25(7):671–682, 1979.

[12] B. Kahn, M. Kalwani, and D. Morrison. Measuring variety-seeking and reinforcement behaviors using panel data. *Journal of Marketing Research*, pages 89–100, 1986.

[13] T. Kamai and Y. Kanazawa. The latent class model of brand choice behaviors incorporating variety-seeking and state dependence. 2012.

[14] N. Koenigstein, G. Dror, and Y. Koren. Yahoo! music recommendations: modeling music ratings with temporal dynamics and item taxonomy. In *Proceedings of the fifth ACM conference on Recommender systems*, pages 165–172. ACM, 2011.

[15] Y. Koren. Collaborative filtering with temporal dynamics. *Communications of the ACM*, 53(4):89–97, 2010.

[16] N. Lathia, S. Hailes, L. Capra, and X. Amatriain. Temporal diversity in recommender systems. In *Proc. of the 33rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIRâĂŹ10)*, pages 210–217, 2010.

[17] L. McAlister. A dynamic attribute satiation model of variety-seeking behavior. *Journal of Consumer Research*, pages 141–150, 1982.

[18] L. McAlister and E. Pessemier. Variety seeking behavior: An interdisciplinary review. *Journal of Consumer research*, pages 311–322, 1982.

[19] S. McNee, J. Riedl, and J. Konstan. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI'06 extended abstracts on Human factors in computing systems*, pages 1097–1101. ACM, 2006.

[20] J. Nakamura and M. Csikszentmihalyi. The concept of flow. *Handbook of positive psychology*, pages 89–105, 2002.

[21] P. S. Raju. Optimum stimulation level: Its relationship to personality, demographics, and exploratory behavior. *Journal of Consumer Research*, pages 272–282, 1980.

[22] N. Sahoo, P. V. Singh, and T. Mukhopadhyay. A dynamic model of employee blog reading behavior.

[23] T. Zhou, Z. Kuscsik, J. Liu, M. Medo, J. Wakeling, and Y. Zhang. Solving the apparent diversity-accuracy dilemma of recommender systems. *Proceedings of the National Academy of Sciences*, 107(10):4511–4515, 2010.

[24] C. Ziegler, S. McNee, J. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*, pages 22–32. ACM, 2005.