

Using psychological principles of memory storage and preference to improve music recommendation systems

Anthony Chmiel and Emery Schubert

Empirical Musicology Laboratory, School of the Arts and Media,
The University of New South Wales, Sydney, New South Wales 2052, Australia
Correspondence: e.schubert@unsw.edu.au

NOTE:

This is an accepted manuscript (pre-print version) of the below article. This paper is not the copy of record and may not exactly replicate the authoritative document

Using psychological principles of memory storage and preference to improve music recommendation systems. (2018). *Leonardo Music Journal* 28, pp. 77-81.

DOI: 10.1162/LMJ_a_01045

Abstract

This paper proposes a novel approach to automated music recommendation systems. Current systems use a number of methods, although these are generally based on similarity of content, contextual information, or user ratings. These approaches therefore do not take into account relevant, well-established models from the field of music psychology. Given recent evidence of this field's excellent capacity to predict music preference, we propose a function based on both the Ebbinghaus *forgetting curve* of memory retention and Berlyne's *inverted-U* model to inform recommendation systems through "collative variables" such as exposure/familiarity. According to the model, an intermediate level of these variables should generate relatively high preference and therefore presents significant untapped data for music recommendation systems.

Using psychological principles of memory storage and preference to improve music recommendation systems

Automated music recommendation systems have become increasingly sophisticated in recent years. These systems employ a number of methods for selecting dynamic, personalized music playlists that generally rely on principles of similarity (Celma, 2010; Schedl, Knees, McFee, Bogdanov, & Kaminskis, 2015). If an individual selects a song, a recommendation can interrogate a number of features of that song—such as similarity of content, metadata, choice by other listeners, and demographic—to form a list of recommendations. However, these systems do not take into account relevant advances made in the field of music psychology. By employing a theory of memory retention and Berlyne's (1960, 1971) inverted-U model of preference, the present work contains a broad overarching thesis on recommendation that can be applied to a variety of music stimuli. First, we outline the inverted-U model and review its efficacy and follow this with a summary of the state of the art in music recommendation systems. Next, we discuss how recommendation systems might benefit from the implementation of this psychological knowledge and propose a novel function for recommendation. Such an implementation is not intended as a replacement for current approaches but rather as a supplementary method to aid in the accurate contextual recommendation of songs the listener will enjoy.

Explanation of the inverted-U model

First published in 1960, Berlyne's inverted-U model is one of the most tested hypotheses of music preference (Hargreaves & North, 2010). The model proposes that preference is primarily related to *collative variables*—variables that can be mentally categorized and compared with one another. The most frequently tested collative variables are *complexity* and *familiarity/exposure* (for an extended list of collative variables, see Berlyne, 1960, 1971). The model proposes that preference is related to collative variables in an inverted-U fashion, as outlined in Figure 1. An intermediate level of a collative variable produces the highest level of preference and, as this collative level is increased or decreased, preference decreases from the optimal point. In cases where the range covered for a particular collative variable is narrow, the observed relationship may only produce a segment of the entire inverted-U curve. Therefore, in addition to a complete inverted-U pattern, the model can be expressed as segments of the curve (Heyduk, 1975; Walker, 1973), depicted in Figure 1 with arrows. First, a positively correlated relationship representing the rising slope of the inverted-U curve can be expected in specific cases, such as when an unfamiliar stimulus is exposed only a handful of times. In this example, we might expect the optimal level of the collative variable not to be surpassed, whereas additional exposures could allow the entire inverted-U curve to appear. In contrast, decreasing preference could be expected for repeated exposures to a stimulus that is already well known.

Figure 1 about here

Berlyne's model has received considerable support in the literature. Recently, we analyzed the results of 57 studies spanning 115 years investigating the relationship between music preference and one or more collative variables (Chmiel & Schubert, 2017). The reported trajectories of preference of 88% of the investigated studies were compatible with the segmented inverted-U model. In addition to this, 56% of the studies exclusively use linear analysis methods, whereas only 25% of the studies incorporate nonlinear analysis methods; therefore, the studies applying linear-only analysis may have hidden inverted-U results. With this in mind, when other non-collative variables—such as those associated with meaningful and personal associations, so-called ecological variables (Berlyne, 1971, p. 69)—are held constant, the inverted-U model can be used to predict a considerable amount of variance in preference.

Description of existing music recommendation systems

There are five common approaches used by automated music recommendation systems: (1) *collaborative filtering*, in which recommendations are based on previous user-item relationships such as user ratings; this is the most common approach used by music recommendation systems (Schedl et al., 2015, p. 454); (2) *content-based filtering*, in which the system examines descriptive information on the music, such as metadata, and compares this information with the user's listed preferences to create a recommendation; other systems also analyze the audio signal, comparing the timbral, temporal, and tonal content with the signal content of other songs; (3) *demographic filtering*, in which the system creates demographic stereotypes for specific songs, artists, or styles (based on variables such as age, gender, and cultural or social traits) and recommends music based

on the preferences of other users of a similar demographic; (4) *context-based filtering*, in which information is gathered to characterize the listening situation, which may entail gathering information about user habits and preferences, such as by examining sequential patterns of usage in a playlist (Celma, 2010), or about the mood, occasion, social setting or specific task at hand (Cai, Zhang, Zhang, & Ma, 2007); (5) *hybrid approaches*, in which two or more approaches are combined in an effort to increase the scope and accuracy of predictions. A hybrid approach may, for example, combine audio analysis with content-based or contextual approaches (e.g., Bu et al., 2010; Shao, Wang, Li, & Ogihara, 2009; Vignoli & Pauws, 2005) or integrate user-rating data with contextual information (e.g., Yoshii & Goto, 2009; Yoshii, Goto, Komatani, Ogata, & Okuno, 2006). As noted in the literature (Celma, 2010; Celma & Cano, 2008; Schedl et al., 2015) hybrid approaches are effective at bypassing the various limitations to which singular approaches are prone. As one of many examples of this, collaborative filtering is regarded as being prone to popularity bias, or the “rich get richer” principle, in which the more an item is rated and purchased, the more it is recommended, and therefore additionally likely to continue being rated and purchased (Celma, 2010). Collaborative filtering is conversely unable to recommend items that have no user ratings.

Most approaches to music recommendation include *similarity* as a central component from which their recommendations are based (Schedl et al., 2015). Similarity-based systems are often used for recommendations of books, television shows and films, in which systems are generally only required to suggest a single item at a time due to their longer consumption time, and as they

are not typically repeatedly consumed. This approach is not ideal for music recommendation systems, in which songs can be repeatedly consumed—sometimes in multiple successive listenings. We therefore suggest that music recommendation systems may be overlooking an important variable. We specifically propose the consideration of collative variables, such as familiarity, by tracking the previous number of exposures and the length of time between exposures of the same song, artist or genre.

Application of the inverted-U model

While some research (e.g., North, Krause, Sheridan, & Ritchie, 2017) has investigated the inverted-U model in terms of music sales, none has been cited that explicitly use collative variables as a predictor of preference in a recommendation system. However, one existing music recommendation system, by Hu and Ogihara (2011), has inspired our own attempt to address this gap. Their system uses five “perspectives” (genre, recording year, time pattern, favor and freshness) upon which recommendations are based. *Genre* and *recording year* were determined by time series analysis based on the latest 16 songs that were played for at least half of their duration, whereas *time pattern* is informed by the user’s typical listening habits for that time of the day. A *favor* rating is also produced, with the intent of giving recommendation priority to tracks that have previously been played in full. The final perspective used by Hu and Ogihara is *freshness*, defined as the “strength of strangeness or the amount of experience forgotten” for a piece (Hu & Ogihara, 2011, p. 105). This variable is therefore related to memory as well as to familiarity. Their system applies the Ebbinghaus (1913) *forgetting curve* to calculate freshness, as depicted below:

$$R = e^{-\frac{t}{S}} \quad (\text{Equation 1})$$

in which R is memory retention, S is the relative strength of memory, and t is time elapsed. S can also be thought of as an exposure event to a song in an individual's personal, mental library; we therefore express it as a positive integer. We propose that freshness or favor can be interpreted as the "boost" the R retention curve receives when a new exposure occurs (which we notate as exposure number S , which is incremented by one from the previous exposure, $S-1$, for $S \geq 2$) for the piece in question. For a recommendation system, this would be the modeled boost if S were incremented at time ∂t . That is, it would be $R(\partial t, S) - R(\partial t, S-1)$. This can be represented in a simplified form as:

$$F(\partial t, S) = e^{-\frac{\partial t}{S}} - e^{-\frac{\partial t}{S-1}} \quad (\text{Equation 2})$$

Here, F is the "freshness" or "favor"¹ that would be gained if the song k in the library were to be played at time ∂t . An array of Fk is generated at any point in time for all k songs in the individual's library, with the items of the array scoring the highest F values being the ones that are more likely to be recommended, along with consideration of other approaches as noted above.

Note that S is an integer ≥ 2 because at least one exposure is required against which to compare the next proposed exposure, and t is the time elapsed since the last exposure to piece k . The calculation of a value for F at a particular point in

¹ Thus, freshness/favor can be considered as a recommendation strength, with higher recommendation strength for a song signifying increased likelihood of its recommendation.

time is demonstrated in Figure 2, and the consequent function that emerges in terms of ∂t and S is shown in two dimensions, for three specific values of ∂t in Figure 3, and in three dimensions in Figure 4. If a vertical plane parallel to the F - S plane is used to form two-dimensional slices of the function (of Figure 4), with $S \geq 2$, an inverted-U curve emerges, with longer time delays between exposures (slicing plane placed further along the time-axis) producing a “gentle,” more slowly falling F value over exposure number than at shorter time delays (slicing, vertical plane placed closer to the time-axis origin). The slice at short time delays is consistent with preference behavior for pieces undergoing massed exposure that reach a peak F quickly, but then soon go out of favor, as is the case with high-rotation music playlists (Martindale, 1984). Inverted-U results also occur when the slicing plane is tilted, emanating from the time-exposure origin (still parallel to the F axis), where an increase in ∂t and S still traces out an inverted-U shape in two-dimensions. For simplicity, we have not included any constants in the proposed equations, such as α and β in

$$\alpha \left(e^{-\frac{\partial t}{S}} \right) - \beta \left(e^{-\frac{\partial t}{S-1}} \right) \quad (\text{Equation 3})$$

and it is assumed that each piece already exists in the individual’s mental music library. In his review of the literature, Finnäs (1989) surmises that six to nine exposures may be sufficient to move the preference curve to its optimal point. However, these matters need to be dealt with using empirical data, should the system be implemented, and we therefore recommend further study in this regard.

Figure 2 about here

Figure 3 about here

Figure 4 about here

Additionally, future approaches may serve to extend our use of exposure to other collative variables. A brief discussion of this in terms of two variables follows (although this is a non-exhaustive list in terms of potential applications of collative variables). First, future systems may recommend music based upon a listener's perceived level of complexity, either through similarity recommendations based upon the variable or with S in the above Equation 2 being replaced by a "complexity rating". Second, a system could use familiarity for *styles* of music (stylistic familiarity) by storing a library of styles (e.g. rock, classical, jazz, etc.) and using the F value for those (call it array F_j) to provide an additional level of sophistication for inverted-U recommendation, building on systems such as that reported by Cai et al. (2007), in which contextual information such as tagged metadata is used to predict a sequence of songs to be recommended next.

In short, the proposed system can continue to create recommendations according to the dynamic consideration of collative variables such as familiarity/exposure in an effort to predict the setting of those variables that maximize enjoyment according to the psychological principles we have employed. Thus, we could expect music recommendations generated by a system based on collative information to be quite different from the recommendations made by current systems.

Conclusion

This paper proposed that automated music recommendation systems can be improved by incorporating well-grounded psychological principles of preference for optimal levels of collative properties of music (such as complexity and familiarity) and thus applying the inverted-U model, which describes the trajectory of preference for these variables. We implemented the inverted-U model through another well-established psychological principle from mental memory storage literature. In contrast to our approach, a large number of conventional systems repeatedly recommend similar songs and styles without explicitly taking additional variables, such as oversaturation (too many exposures over too small a time period), into account. We propose that systems informed by collative variables and the inverted-U model would therefore produce significantly different, and arguably more contextual, recommendations than existing systems. The advantages afforded by the use of these traditional, well-studied variables of music psychology cannot be overstated. We have already found elements of such an approach in existing systems (Hu & Ogihara, 2011; North et al., 2017), yet not in a manner that is theoretically informed by the inverted-U model. The inverted-U shape is more conceptual than mathematical. Our study may be the first psychologically plausible mathematical representation of the inverted-U. Of note is our proposal of collative variables as a new set of parameters for recommendation to coexist alongside existing parameters, such as similarity of genre and recording year. In the present work, we aimed to develop this new direction, from which engineers could implement the inverted-U model into future systems.

References

- Berlyne, D. E. (1960). *Conflict, arousal and curiosity*. New York: McGraw-Hill.
- Berlyne, D. E. (1971). *Aesthetics and psychobiology*. New York: Appleton-Century-Crofts.
- Bu, J., Tan, S., Chen, C., Wang, C., Wu, H., Zhang, L., & He, X. (2010). *Music recommendation by unified hypergraph: Combining social media information and music content*. Paper presented at the Proc. the 18th ACM International Conference on Multimedia, Firenze, Italy.
- Cai, R., Zhang, C., Zhang, L., & Ma, W.-Y. (2007). *Scalable music recommendation by search*. Paper presented at the Proc. the 13th ACM International Conference on Multimedia, Singapore.
- Celma, O. (2010). *Music recommendation and discovery: The long tail, long fail, and long play in the digital music space*. Berlin: Springer-Verlag.
- Celma, O., & Cano, P. (2008). *From hits to niches? Or how popular artists can bias music recommendation and discovery*. Paper presented at the 2nd Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition, Las Vegas, USA.
- Chmiel, A., & Schubert, E. (2017). Back to the inverted-U for music preference: A review of the literature. *Psychology of Music*, 45(6), 886-909.
- Ebbinghaus, H. (1913). *Memory: A contribution to experimental psychology*. New York: Teachers College.
- Finnäs, L. (1989). How can musical preferences be modified? A research review. *Bulletin of the Council for Research in Music Education*, 102, 1-58.
- Hargreaves, D. J., & North, A. C. (2010). Experimental aesthetics and liking for music. In P. N. Juslin & J. Sloboda (Eds.), *Handbook of music and emotion: Theory, research, applications* (pp. 515-546). New York: Oxford University Press.
- Heyduk, R. G. (1975). Rated preference for musical compositions as it relates to complexity and exposure frequency. *Perception and Psychophysics*, 17(1), 84-91.
- Hu, Y., & Ogihara, M. (2011). *Nexttone Player: A music recommendation system based on user behavior*. Paper presented at the Proc. the 12th International Society for Music Information Retrieval Conference, Miami, Florida.
- Martindale, C. (1984). The pleasures of thought: A theory of cognitive hedonics. *Journal of Mind and Behavior*, 5, 49-80.
- North, A. C., Krause, A. E., Sheridan, L. P., & Ritchie, D. (2017). Energy, typicality, and music sales: A computerized analysis of 143,353 pieces. *Empirical Studies of the Arts*, 35(2), 214-229.
- Schedl, M., Knees, P., McFee, B., Bogdanov, D., & Kaminskas, M. (2015). Music recommender systems. In F. Ricci, L. Rokach, B. Shapira, & P. B. Kantor (Eds.), *Recommender systems handbook* (pp. 453-492). New York: Springer US.
- Shao, B., Wang, D., Li, T., & Ogihara, M. (2009). Music recommendation based on acoustic features and user access patterns. *IEEE Transactions on Audio, Speech, and Language Processing*, 17(8), 1602-1611.

-
- Vignoli, F., & Pauws, S. (2005). *A music retrieval system based on user driven similarity and its evaluation*. Paper presented at the Proc. the 6th International Conference on Music Information Retrieval, London, UK.
- Walker, E. L. (1973). Psychological complexity and preference: A hedgehog theory of behavior. In D. E. Berlyne & K. B. Madsen (Eds.), *Pleasure, reward, preference* (pp. 65-97). New York: Academic Press.
- Yoshii, K., & Goto, M. (2009). *Continuous pLSI and smoothing techniques for hybrid music recommendation*. Paper presented at the Proc. the 10th International Society for Music Information Retrieval Conference, Kobe, Japan.
- Yoshii, K., Goto, M., Komatani, K., Ogata, T., & Okuno, H. G. (2006). *Hybrid collaborative and content-based music recommendation using probabilistic model with latent user preferences*. Paper presented at the Proc. the 7th International Conference on Music Information Retrieval, Victoria, Canada.




Figure 1. Appropriation of the inverted-U relationship described by Berlyne (1960, 1971). The two segments of the curve are depicted with arrows.




Figure 2. Visualization of $F(\Delta t, S)$ calculation given a stimulus exposure number S occurring Δt_{ab} seconds after an earlier exposure to the same stimulus. Note that time a to b ($\Delta t = \Delta t_{ab}$) is the time elapsed from exposure number $S-1$ to the next exposure (S). F is the “favor/freshness” output used for recommendation analysis. The greater the value of F , the greater the recommendation value/strength. The two curves apply the Ebbinghaus *forgetting curve* $R(t, S)$. Equations and plots are shown in simplified form.




Figure 3. Three visualizations of the $F(\partial t, S)$ calculation for specific values of ∂t . These visualizations characterize the changing inverted-U curve that is produced, depending on the value of ∂t between exposures. The top, green curve has a smaller ∂t value between exposures in comparison to the other two curves depicted; the recommendation value (recommendation strength) therefore peaks early, followed by a quick, sharp decrease in recommendation value. In contrast, the bottom, red curve produces a more gradual inverted-U in terms of recommendation value, and the peak of the curve occurs after additional exposures.




Figure 4. $F(\partial t, S)$. Scale values on the z-axis are arbitrary, based on the function.