



Predicting Music Relistening Behavior Using the ACT-R Framework

Markus Reiter-Haas
reiter-haas@tugraz.at
Graz University of Technology
Graz, Styria, Austria

Emilia Parada-Cabaleiro
Markus Schedl
emilia.parada-cabaleiro@jku.at
markus.schedl@jku.at
Johannes Kepler University Linz (JKU) and Linz Institute
of Technology (LIT)
Linz, Upper Austria, Austria

Elham Motamedi
Marko Tkalcic
elham.motamedi@famnit.upr.si
marko.tkalcic@famnit.upr.si
University of Primorska
Koper, Slovenia

Elisabeth Lex
elisabeth.lex@tugraz.at
Graz University of Technology
Graz, Styria, Austria

ABSTRACT

Providing suitable recommendations is of vital importance to improve the user satisfaction of music recommender systems. Here, users often listen to the same track repeatedly and appreciate recommendations of the same song multiple times. Thus, accounting for users' relistening behavior is critical for music recommender systems. In this paper, we describe a psychology-informed approach to model and predict music relistening behavior that is inspired by studies in music psychology, which relate music preferences to human memory. We adopt a well-established psychological theory of human cognition that models the operations of human memory, i.e., Adaptive Control of Thought—Rational (ACT-R). In contrast to prior work, which uses only the base-level component of ACT-R, we utilize five components of ACT-R, i.e., base-level, spreading, partial matching, valuation, and noise, to investigate the effect of five factors on music relistening behavior: (i) recency and frequency of prior exposure to tracks, (ii) co-occurrence of tracks, (iii) the similarity between tracks, (iv) familiarity with tracks, and (v) randomness in behavior. On a dataset of 1.7 million listening events from Last.fm, we evaluate the performance of our approach by sequentially predicting the next track(s) in user sessions. We find that recency and frequency of prior exposure to tracks is an effective predictor of relistening behavior. Besides, considering the co-occurrence of tracks and familiarity with tracks further improves performance in terms of R-precision. We hope that our work inspires future research on the merits of considering cognitive aspects of memory retrieval to model and predict complex user behavior.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

RecSys '21, September 27–October 1, 2021, Amsterdam, Netherlands

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8458-2/21/09.

<https://doi.org/10.1145/3460231.3478846>

CCS CONCEPTS

• Information systems → Music retrieval; Recommender systems; • Applied computing → Psychology.

KEYWORDS

music prediction, cognition-inspired retrieval, human cognition, relistening behavior, user modeling, psychology-informed recommender systems, adaptive control thought-rational (ACT-R)

ACM Reference Format:

Markus Reiter-Haas, Emilia Parada-Cabaleiro, Markus Schedl, Elham Motamedi, Marko Tkalcic, and Elisabeth Lex. 2021. Predicting Music Relistening Behavior Using the ACT-R Framework. In *Fifteenth ACM Conference on Recommender Systems (RecSys '21)*, September 27–October 1, 2021, Amsterdam, Netherlands. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3460231.3478846>

1 INTRODUCTION AND RESEARCH CONTEXT

Music recommender systems (MRS) [22] are nowadays employed in many use cases, such as automatic playlist generation, next-track recommendation, context-aware music recommendation, and even in the creative process of music production [23]. Providing useful music recommendations is a challenging task due to various aspects such as the variability in the purpose of music consumption, inconclusive and insufficient user feedback, or situational and contextual aspects that influence a user's current music preference [12, 24]. Besides, users of MRS show a tendency to *relisten to songs*, which means they appreciate recommendations of the same song multiple times [3, 25], which is in stark contrast to movie or product recommendation systems, where users typically do not prefer repeated recommendations [21, 24].

In this work, we introduce a novel psychology-informed approach to model and predict *music relistening* behavior of tracks. Our work is inspired by related work on psychology-informed recommender systems [13], and studies in music psychology, which find that music preferences are biased by human memory: user

studies showed that repeated exposure increases recognition and positive attitude towards music [18].

While several works have been proposed to model repeated consumption of items, such as temporal point processes [26], factorization models [20] or neural networks [21], to the best of our knowledge, they do not incorporate underlying cognitive processes of human behavior into their approaches.

In our work, we exploit a theory about human cognition, i.e., the cognitive architecture *Adaptive Control of Thought–Rational (ACT-R)* [1], and investigate the utility of its declarative memory module for the task at hand. The declarative memory module of ACT-R stores and retrieves information and consists of separate components. In contrast to prior work, which uses the base-level component to model predict user preferences in various domains (e.g., [9–12, 14]), we investigate the *utility of five components*, i.e., (i) base-level, (ii) spreading, (iii) partial matching, (iv) valuation, and (v) noise component, in a next track prediction scenario. These components enable us to investigate the effect of various factors on track relistening behavior, i.e., (i) recency and frequency of prior exposure to tracks, (ii) co-occurrence of tracks, (iii) the similarity between tracks, (iv) familiarity with tracks, and (v) randomness in behavior. To that end, we sequentially simulate track predictions for the next single and remaining track in user sessions considering their previous music consumption on a subset of a publicly available dataset, i.e., the LFM-2b dataset [15]. Our experiments show that recency of prior exposure to tracks is an effective predictor of relistening behavior in terms of R-precision; adding co-occurrence and familiarity in the model further improves the prediction performance.

Summing up, our main contributions are¹: (1) we propose a novel approach of using the declarative memory module within the ACT-R framework to model and predict music relistening behavior, (2) we investigate and discuss the merit of all five declarative memory components of ACT-R in light of what aspects of relistening behavior they let us model.

2 ACT-R COMPONENTS

Our work is based on the declarative memory part of the cognitive ACT-R framework [1]. In ACT-R, the activation function for a particular chunk i (which in the context of recommender systems represents an item) is defined by the sum of individual components [2] given by the equation

$$A_i = B_i + S_i + P_i + V_i + \epsilon_i \quad (1)$$

where B_i represents the base-level component, S_i the spreading component, P_i the partial matching component, and V_i the valuation component. In addition, ϵ_i accounts for the noise in the activation. The predictions consider all items in a user's music listening history ranked by the activation term A_i . We regularize the results using the softmax function

$$\text{softmax}(x_i) = \frac{\exp(x_i)}{\sum_{x_j \in I_u} \exp(x_j)} \quad (2)$$

¹ The code is openly available at <https://github.com/socialcomplab/recsys21-relistening-actr>

As a consequence, the sum of all values x_i assigned to candidate items I_u , i.e., items in the user history, is 1. This avoids one component dominating the others, while elevating each component's top items. In the following, we describe each component in more detail.

Base-level component. The base-level component models item occurrences, considering both their recency and frequency. The component requires a list of tuples (*item, timestamp*) extracted from the user history as input. Its activation is defined by the base-level learning equation [10] given by

$$B_i = \text{softmax} \left(\sum_{j=1}^n (t_{ref} - t_{ij})^{-d} \right) \quad (3)$$

where t_{ref} represents the reference timestamp, i.e., the timestamp of the prediction, and t_{ij} represents the timestamp of the j -th interaction out of n total interactions on item i by the user. The time decay factor d models the forgetting of item interactions and thus balances recency and frequency of occurrences.

The base-level component can approximate, depending on the setting of the parameter d , both a user-based most popular (in the case of frequency) or a user-based most recent (in the case of recency) algorithm. More specifically, a base-level component with no decay, i.e., $d = 0$, is equivalent to a user-based most popular recommender, while a most recent approach can be approximated arbitrarily close with a very high decay of the base-level component, i.e., $d \rightarrow \infty$.

Spreading component. The spreading component models item co-occurrences and considers (*item, session*) tuples from the user history as input. The component *spreads* its activation among items using contextual information, i.e., sessions. Thus, it learns to associate items with each other that appear in the same session, which is modeled using probabilities. Specifically, the spreading activation [7] is defined by

$$S_i = \text{softmax} \left(\frac{P(i \in C_j)}{P(i)} \right) \quad (4)$$

where $P(i)$ represents the probability of item i appearing over all contexts, i.e., past user sessions. $P(i \in C_j)$ represents the probability of item i appearing in context C_j . We define the last item in the user history as the context item j , and the context C_j as all sessions where the context item j appears. In short, the spreading component predicts items co-occurring often with the most recent item, while also penalizing overall frequent items. Since frequent occurrences of an item in sessions also lead to high co-occurrences with any other item, this penalty is necessary.

Partial matching component. The partial matching component captures similarities between the items using a list of features. The input tuples are a set (as multiple interactions on the same items do not change the result) and have the form (*item, f_1, f_2, \dots, f_m*), where each f_x is an entry of the feature vector of fixed length m . These features need to be extracted from the data associated with a given item. The activation is then simply defined by a similarity function $P_i = \text{softmax}(\text{sim}(i, j))$ between the candidate item i and the context item j , which again is the last item in the user history.

For the similarity function $sim()$, we use the dot product. This approach is equivalent to a content-based item-to-item recommender, i.e., predicting similar items to those in the target user's profile.

Valuation component. The valuation component was proposed by [8] and models core affect. Instead of assigning affect, e.g., emotions, to items directly, it defines the affective value as a learned parameter by a subject through interactions with stimuli [8]. The parameter uses previous interactions and their associated rewards to estimate the valuation of an item. Specifically, the valuation is defined by the update rule

$$V_i(n) = V_i(n-1) + \alpha(R_i(n) - V_i(n-1)) \quad (5)$$

where $V_i(n)$ specifies the valuation and $R_i(n)$ the associated reward at the n -th interaction with a learning rate of α . As a starting point, we set $V_i(0) = 0$. As input, the component uses an ordered list (due to the order affecting the update rule) in the form of (*item*, *reward*). Thus, we extract the reward directly from the input derived from other properties. For music preferences, we can derive the reward from the duration of a listening event, e.g., total time or ratio the track was listened to. The activation of the valuation is defined by $V_i = softmax(V_i(n))$ where $V_i(n)$ is calculated by iteratively applying the update rule.

In the simplest case, it applies the familiarity principle, i.e., exposure to an item leads to an increase of preference, which is achieved by setting $R_i(n) = 1$ for all interactions. Hence, the component equals an approach that predicts the most popular items of a user. This result could be achieved in a computationally less expensive way by just summing the rewards. However, if negative feedback is incorporated or the rewards differ, then this no longer holds, as the more recent rewards play a bigger role in the valuation term.

Noise. Noise accounts for randomness in behavior in the overall activation function. Moreover, random predictions are often used as the most basic baseline in recommender systems. The activation is given by $\epsilon_i = softmax(rng())$ with $rng()$ being a random number generator. On its own, it returns a random item from the list of possible items, i.e., candidate items from the user history. Hence, it takes just a set of (*item*) tuples as input.

3 EXPERIMENTAL SETUP

Dataset. Our dataset is a subsample of the LFM-2b dataset² [15] from 2019. To reduce the computational demand, we perform stratified sampling³ and create a subset of 150 users with 1,686,296 listening events (LEs). Thus, our dataset contains 11,242 LEs per user on average. The dataset contains 355,464 unique tracks, i.e., 5 LEs on average per track. This small number of average track LEs is explained by the skewed distribution with the most popular track being listened to 1,776 times, whereas most tracks appear in the long tail, e.g., 158,238 tracks only appear once in the dataset. This phenomenon outlines the utility of predicting relistening behavior in this dataset as users tend to listen to the same tracks often, i.e., the minority of tracks being responsible for the majority of LEs. The

²<http://www.cp.jku.at/datasets/LFM-2b/>

³Based on the distribution, we empirically excluded users with less than 1,000 and more than 30,000 listening events and assigned the remaining users to 10 equal sized bins depending on the number of listening events.

prevalence of relistening becomes even more evident considering individual users, who, on average, relisten to previous tracks 66% of the time and to a single track 111 times.

We enrich the dataset with session information derived from the timestamps of interactions, and with information of user-generated tags and duration of tracks. Sessions are inferred based on the gaps between LEs, with sessions corresponding to sequences of LEs containing gaps of at most 30 minutes. This is an established approach for session extraction [6, 17] and results in 96,156 user sessions with on average 18 LEs per session. The tags and duration of tracks are collected from the Last.fm API⁴. Note that not every track has tags or duration associated with it. As a result, 12,434 LEs are missing tag information and 108,210 LEs are missing duration information.

Evaluation Protocol. For each user, we shift a sliding window of one week considering past LEs over all the LEs (including LEs of non-relistened tracks) in the user's listening history, i.e., hop size equals 1 LE. Then, we predict the remaining tracks of the current session, i.e., the session from the most recent track. Using this procedure, we simulate $|Q| = 811,025$ predictions. Considering a fixed time interval instead of the whole user history has the benefit of avoiding scaling issues, while predicting only tracks in the current session provides a realistic prediction scenario. To evaluate the performance, we use the R-precision metric with R being the number of distinct tracks in the remainder of the session given by

$$R-prec = \frac{1}{|Q|} \sum_{q \in Q} \frac{r}{R} \quad (6)$$

with r being the number of relevant tracks evaluated over all generated predictions $q \in Q$. Hence, it captures the ratio of correctly predicted tracks in user sessions, i.e., the precision and recall at R .

To contrast this metric, we also evaluate the performance for the next track prediction using the hitrate metric. Given the top predicted track and the next track in the user session, a hit is recognized if the top and next coincide and a miss otherwise. The overall hitrate, i.e., *Next-HR*, is given by averaging the number of observed hits over all possible hits.

Algorithm Configurations. To configure the decay parameter of the *Base-level* component and in line with previous studies [12], we analyze the relistening behavior using a log-log plot as shown in Figure 1. We observe a good fit for a power law distribution of relistening behavior, i.e., linear fit in the log-log space. Thus, we conjecture that the *Base-level* component, which directly models the decay of such behavior, should already provide good results. We choose to evaluate the component on three configurations, i.e., using the *default* parameter of $d = 0.5$ of the ACT-R framework, the parameter fitted to the LEs of 2019 with $d = 1.737$, and the parameter fitted to the one *week* interval with $d = 0.86$.

The *Valuation* component is evaluated deriving the reward from the listening duration compared to the track duration by assigning either an always positive reward equal to the listening *ratio* or mapping it to *discrete* values using -1 for $\leq 33\%$, 1 for $\geq 66\%$, and 0 otherwise. Additionally, we configure it to be equivalent to a

⁴<https://www.last.fm/api/show/track.getTopTags> for tags and <https://www.last.fm/api/show/track.getInfo> for duration.




Figure 1: Relisting log-log plot of LEs in 2019. Blue dots correspond to value counts of relisting behaviors in hours. Relisting events follow a power law distribution, i.e., have an exponential drop-off in time. The solid orange line (longer, top) is fitted on all LEs in 2019. It provides a good fit of $R^2 = 0.801$ and has a slope of -1.737 . We observe an arch at the end of the scale, which we regard as an artifact of filtering by year (i.e., longer relisting behaviors are cut-off). Hence, we also fitted a line considering only shorter relisting behaviors of one week, which is displayed in red (bottom, dashed). This line serves as an estimate for our evaluation protocol and has a slightly better fit of $R^2 = 0.811$ and a more moderate slope of -0.860 . The spikes at certain intervals show that the relisting behavior follows some periodicity (e.g., after 1 day). One peak is right at the beginning as people tend to relisten the same track within a short interval (e.g., immediate relisting).

user-based most popular (*MP*) approach by setting the reward to 1 for all LEs.

For the *Partial Matching* component, we use the user-generated tags for a particular track from which a feature vector is extracted by computing the mean across the tags’ word embedding representation. To reduce computational power, we apply Principal Component Analysis as a method for dimensionality reduction to the pre-trained Word2Vec [16] from Google News, whose 300 dimensions were reduced to 100^5 . Finally, the *Spreading* component does not require any parameters. Similarly, there are no parameters for the simulation of *Noise*, as it just predicts random tracks.

For the combination of the *ACT-R* components, we use the most basic version of each component and denote it as letters, i.e., *B* for base-level with the default value, *S* for spreading, and *V* for valuation with the most popular configuration⁶. We also included two non-*ACT-R* approaches in our evaluation for comparison. To contrast the *Spreading* component, we use contiguous sequential pattern mining [19] to estimate the transition probability (*TransProb*) given a reference track. Furthermore, we use an algorithm that predicts the *MostRecent* tracks in the user history.

⁵We experimented with other content extraction methods with no avail.

⁶Due to the weak performance, we omitted including the partial matching component and noise.

Algorithm	R-prec	Next-HR
TransProb	.03839	.15907
Partial Matching	.03895	.01320
Noise	.03996	.00289
Valuation(discrete)	.04751	.00533
Valuation(ratio)	.05987	.01042
Valuation(MP)	.08436	.01477
Spreading	.09235	.02117
Base-level(2019)	.09903	.03200
ACT-R(B,V)	.10069	.02416
MostRecent	.10167	.05189
Base-level(default)	.10380	.02451
Base-level(week)	.10489	.02883
ACT-R(S,V)	.11009	.02998
ACT-R(B,S)	.11042	.02972
ACT-R(B,S,V)	.11119	.02961

Table 1: R-precision of 811,025 predictions per algorithms on a stratified sample of 150 users (*R-prec*) and next track hitrate (*Next-HR*). Base-level is the strongest individual component; combining components (*ACT-R* with base-level (*B*), valuation (*V*), spreading (*S*)) improves the results. Green shading indicates relative performance.

4 EXPERIMENTAL RESULTS

Prediction Performance. Table 1 reports the prediction performance sorted by R-precision. We refer to the algorithms by their names and their configurations (in parentheses). *Base-level(week)* is the best performing component in terms of R-precision, which we attribute to the good fit of the power law distribution to the relisting behavior in Figure 1. *Base-level(default)* is a close second, which demonstrates its applicability even without estimating the decay parameter. In comparison, *Base-level(2019)* performs worse, which hints at the discrepancy between the shorter interval of the evaluation protocol and all LEs in 2019. Notably, *MostRecent* provides a strong baseline, which is also indicated by the left-most data point in Figure 1 that displays an above-average relisting behavior for recent tracks. Note that neither a focus on short nor long-term memory has a clear advantage over more balanced approaches that fit the data.

The *Spreading* and *Valuation* components have a lower performance, whereas, *Partial Matching* and *Noise* seem unsuited for the task. While *TransProb* is unable to predict the majority of remaining tracks in a session, it has remarkably the best performance on predicting the next single track. Besides, most other algorithms perform similarly well on this task, with the notable exceptions of *MostRecent* performing better and *Noise* performing worst. Here, we also observe that *Partial Matching*, while still under-performing, performs better than random in this regard. We suspect that the

low performance of *Partial Matching* is partly explained due to missing tags in the dataset, but also the result of considering tags alone is insufficient for the task, especially when predicting longer sequences.

Comparing the configuration of *Valuation*, *MP* has the best performance of the three, while *discrete* has the lowest. This observation is more prominent considering the Next-HR metric and can be explained by the strong performance of *MostRecent* and the influence of the last, i.e., the most recent, update step on the final valuation term. Hence, assigning low or negative rewards on recent tracks has a detrimental effect when predicting the next track(s) in a sequence.

Considering the combined models, we observe that three of the four *ACT-R* combinations perform better in terms of R-precision than all individual components. This observation outlines the utility of incorporating multiple cognitive processes for predictions. Notably, *ACT-R(S,V)* outperforms the *Base-level* components without including the component itself. Hence, while *Valuation* and *Spreading* alone only provide mediocre performance, combining them leads to effective predictions. Similar to how *Base-level(2019)* is ill-fitted, *ACT-R(B,V)* leads to a drop that could be the result of a dominance of popular tracks, which are considered by both the *Base-level* and *Valuation* component. In comparison, *ACT-R(B,S)* does not suffer such issue, as the *Spreading* component explicitly penalizes occurrence alone. Finally, *ACT-R(B,S,V)* achieves the highest results as the three components seem to balance each other, i.e., the addition of the *Spreading* component counteracts the dominance of popularity present in *ACT-R(B,V)*.

Estimating Component Weights. In addition to naively combining the components by just adding their activation terms, we explore the applicability of weighted combinations. We evaluate the impact of the three best performing components, i.e., *Base-level(week)*, *Spreading*, and *Valuation(MP)*, on the overall outcome using linear models to weight their activation terms with respect to tracks in the remaining session⁷ on 10% of the users. Thus, the altered activation equation becomes $A_i = b \cdot B_i + s \cdot S_i + v \cdot V_i$ with the weights b , s , and v being estimated by the model. In line with the results in Table 1, linear regression models assign the highest weights to *Base-level(week)* and the lowest to *Valuation(MP)*. Considering for instance the linear regression model without fitting an intercept, the learned coefficients are *Base-level(week)* = 0.34, *Spreading* = 0.26, and *Valuation(MP)* = 0.18. Evaluating the performance of the weighted combinations does, however, not lead to an increase compared to *ACT-R(B,S,V)*. Moreover, the results are virtually identical to *ACT-R(S,V)*. This phenomenon can be explained when fitting an intercept, as this leads to a slightly negative factor of *Valuation(MP)* = -0.05 without altering the other factors substantially (i.e., *Spreading* = 0.26 and *Base-level(week)* = 0.35). Enforcing a positivity constraint would then nullify the valuation term. This suggests that the valuation term is mostly accounting for the bias, i.e., the fitted intercept, in the model and leads to intuitive results as *Valuation(MP)* scores the most popular tracks highest. Overall, the results indicate that Equation 1, i.e., without weights, is sufficient for predicting music relisting behavior.

⁷We assign a 1 to tracks in remaining session and 0 otherwise. Recall that activation terms lie also in the interval [0, 1].

5 CONCLUSION

In this paper, we described a psychology-informed approach to model and predict music relisting behavior based on the declarative module of the ACT-R theory. ACT-R consists of several components, i.e., base-level, spreading, partial matching, valuation, and noise, each corresponding to different aspects of memory processes such as recency prior exposure to an item or the familiarity of items. We evaluated our approach on a subset of the LFM-2b dataset by sequentially predicting track relisting from previous listening events of a user. Our results show that cognitive aspects of memory retrieval are effective predictors of relisting behavior. Specifically, the recency and frequency of prior exposure as modeled by the base-level component result in the highest predictive performance of individual components. Besides, we find even more effective combinations by adding spreading and valuation that model co-occurrence and familiarity, respectively. Thus, our work emphasizes the consideration of the whole complexity of user behavior and their underlying cognitive decision-making processes in music predictions.

As a limitation, we recognize that we consider user behavior predictions through the lens of a single yet popular framework from psychology. Also, we evaluate the framework on a single dataset from the music domain on a small number of users (i.e., 150 users with approx. 1.7 Mio. LEs).

In the future, we aim to expand the scope of the experiments by considering more datasets from different domains. Also, we plan to integrate content information such as lyrics via partial matching or other hybrid approaches [4, 5], rather than only considering interactions. Finally, we see our work not as a replacement, but as a complement for existing approaches, and we will research ways to integrate our findings into classic music recommender systems.

ACKNOWLEDGMENTS

We thank Stefan Brandl and Alessandro Melchiorre for their insights and support in data acquisition and preparation.

REFERENCES

- [1] John R Anderson, Daniel Bothell, Michael D Byrne, Scott Douglass, Christian Lebiere, and Yulin Qin. 2004. An integrated theory of the mind. *Psychological review* 111, 4 (2004), 1036.
- [2] Dan Bothell. 2020. *ACT-R 7.21+ reference manual*. Technical Report. Technical Report, Carnegie Mellon University.
- [3] Frederick Conrad, Jason Corey, Samantha Goldstein, Joseph Ostrow, and Michael Sadowsky. 2019. Extreme re-listening: Songs people love... and continue to love. *Psychology of Music* 47, 2 (2019), 158–172.
- [4] Yashar Deldjoo, Markus Schedl, Paolo Cremonesi, and Gabriella Pasi. 2020. Recommender systems leveraging multimedia content. *ACM Computing Surveys (CSUR)* 53, 5 (2020), 1–38.
- [5] Yashar Deldjoo, Markus Schedl, and Peter Knees. 2021. Content-driven Music Recommendation: Evolution, State of the Art, and Challenges. *arXiv preprint arXiv:2107.11803* (2021).
- [6] Carsten Eickhoff, Jaime Teevan, Ryan White, and Susan Dumais. 2014. Lessons from the journey: a query log analysis of within-session learning. In *Proceedings of the 7th ACM international conference on Web search and data mining*. ACM, New York, NY, USA, 223–232.
- [7] Danilo Fum and Andrea Stocco. 2004. Memory, Emotion, and Rationality: An ACT-R interpretation for Gambling Task results. In *Proceedings of the International Conference on Cognitive Modelling, ICCM 2004, Pittsburgh, Pennsylvania, USA, July 30 - August 1, 2004*. 106–111.
- [8] Ion Juvina, Othalia Larue, and Alexander Hough. 2018. Modeling valuation and core affect in a cognitive architecture: The impact of valence and arousal on memory and decision-making. *Cognitive Systems Research* 48 (2018), 4–24.

- [9] Dominik Kowald and Elisabeth Lex. 2016. The influence of frequency, recency and semantic context on the reuse of tags in social tagging systems. In *Proceedings of the 27th ACM Conference on Hypertext and Social Media*. 237–242.
- [10] Dominik Kowald, Subhash Chandra Pujari, and Elisabeth Lex. 2017. Temporal effects on hashtag reuse in twitter: A cognitive-inspired hashtag recommendation approach. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1401–1410.
- [11] Emanuel Lacic, Dominik Kowald, Markus Reiter-Haas, Valentin Slawicek, and Elisabeth Lex. 2018. Beyond accuracy optimization: On the value of item embeddings for student job recommendations. *WS on Multi-dimensional Information Fusion for User Modeling and Personalization* (2018).
- [12] Elisabeth Lex, Dominik Kowald, and Markus Schedl. 2020. Modeling Popularity and Temporal Drift of Music Genre Preferences. *Transactions of the International Society for Music Information Retrieval* 3, 1 (2020), 17–30.
- [13] Elisabeth Lex, Dominik Kowald, Paul Seitlinger, Thi Ngoc Trang Tran, Alexander Felfernig, and Markus Schedl. 2021. Psychology-informed Recommender Systems. *Foundations and Trends® in Information Retrieval* 15, 2 (2021), 134–242.
- [14] Leendert Van Maanen and Julian N. Marewski. 2009. Recommender systems for literature selection: A competition between decision making and memory models. In *Proceedings of the 31st Annual Meeting of the Cognitive Science Society*. Cognitive Science Society, 2914–2919.
- [15] Alessandro B. Melchiorre, Navid Rekabsaz, Emilia Parada-Cabaleiro, Stefan Brandl, Oleg Lesota, and Markus Schedl. 2021. Investigating Gender Fairness of Recommendation Algorithms in the Music Domain. *Information Processing & Management* 58 (2021), 102666.
- [16] Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2–4, 2013, Workshop Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1301.3781>
- [17] José Luis Ortega and Isidro Aguillo. 2010. Differences between web sessions according to the origin of their visits. *Journal of Informetrics* 4, 3 (2010), 331–337.
- [18] Isabelle Peretz, Danielle Gaudreau, and Anne-Marie Bonnel. 1998. Exposure effects on music preference and recognition. *Memory & cognition* 26, 5 (1998), 884–902.
- [19] Massimo Quadrana, Paolo Cremonesi, and Dietmar Jannach. 2018. Sequence-aware recommender systems. *ACM Computing Surveys (CSUR)* 51, 4 (2018), 1–36.
- [20] Dimitrios Rafailidis and Alexandros Nanopoulos. 2015. Repeat consumption recommendation based on users preference dynamics and side information. In *Proceedings of the 24th International Conference on World Wide Web*. 99–100.
- [21] Pengjie Ren, Zhumin Chen, Jing Li, Zhaochun Ren, Jun Ma, and Maarten De Rijke. 2019. Repeatnet: A repeat aware neural recommendation machine for session-based recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI Press, 4806–4813.
- [22] Markus Schedl, Emilia Gómez Gutiérrez, and Julián Urbano. 2014. Music information retrieval: Recent developments and applications. *Foundations and Trends in Information Retrieval*. 2014 Sept 12; 8 (2-3): 127-261. (2014).
- [23] Markus Schedl, Peter Knees, and Fabien Gouyon. 2017. New Paths in Music Recommender Systems Research. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. ACM, New York, NY, USA, 392–393. <https://doi.org/10.1145/3109859.3109934>
- [24] Markus Schedl, Hamed Zamani, Ching-Wei Chen, Yashar Deldjoo, and Mehdi Elahi. 2018. Current challenges and visions in music recommender systems research. *Int. J. Multim. Inf. Retr.* 7, 2 (2018), 95–116. <https://doi.org/10.1007/s13735-018-0154-2>
- [25] Kosetsu Tsukuda and Masataka Goto. 2020. Explainable recommendation for repeat consumption. In *Fourteenth ACM Conference on Recommender Systems*. ACM, New York, NY, USA, 462–467.
- [26] Chenyang Wang, Min Zhang, Weizhi Ma, Yiqun Liu, and Shaoping Ma. 2019. Modeling item-specific temporal dynamics of repeat consumption for recommender systems. In *The World Wide Web Conference*. ACM, New York, NY, USA, 1977–1987.