Read Between the Tracks: Exploring LLM-driven Intent-based Music Recommendations

Anna Hausberger
Johannes Kepler University
anna.hausberger@jku.at

Petra Jósár
Johannes Kepler University

Markus Schedl
Johannes Kepler University
Linz Institute of Technology
markus.schedl@jku.at

Abstract

This paper evaluates the effectiveness of large language models (LLMs) on the task of context-aware music recommendation, specifically focusing on the alignment of music tracks with a listening intent, in addition to user preferences. We present a preliminary investigation in which five LLMs (variants of *LLama*, *Qwen*, and *Mistral*) are tasked with ranking a candidate set of tracks containing both ground-truth items (associated with specific user-intent pairs) and distractor items (containing user-relevant, intent-relevant, or non-user and non-intent relevant items). Our results show that LLMs rank intent-user-relevant items higher than the distractors, with *Llama-3.1-8B-Instruct* having the best performance (NDCG of $0.32_{0.20}$ vs. $0.20_{0.15}$). We further investigate whether performance differs when mentioning the listening intent explicitly in the prompt vs. implicitly given solely music preferences. Surprisingly, the LLMs achieved the best performance through an implicit indication of intent, versus explicitly adding it to the prompt, with *Mistral-7B-Instruct-v0.3* performing the best (NDCG of $0.37_{0.22}$ vs. $0.29_{0.18}$).

1 Introduction

Listening to music is a common part of everyday life, and for most listeners today this experience is shaped by large-scale streaming platforms. Within these mediated listening contexts, music serves purposes that extend beyond mere entertainment. Prior psychological work by Schäfer et al. (2013) identified 129 non-redundant functions of music listening, organized along three underlying dimensions: self-awareness, social relatedness, and arousal and mood regulation. These functions can be further classified into higher-level clusters that reflect a user's *listening intent*, such as relaxation or reflection (Hausberger et al., 2025).

Exploring the listening intent of a user raises multiple challenges, e.g., a listener might not be aware of how they use music or how music serves them in a specific listening session. Also, the availability of labeled data is scarce. Music listening interactions are often available without any further contextual information on the listening intent of the user. Detecting and identifying music tracks that are used for the same listening intent by the user is still an open problem. This work explores, in a preliminary study, if LLMs can detect and recommend music tracks that have been labeled for being listening to with the same listening intent. Therefore, we put forward two research questions:

- **RQ1:** To what extent can LLMs select and recommend music tracks for a specific listening intent and user from a constrained candidate set?
- **RQ2:** How does providing information about the user's music preferences and listening intent in the prompt affect the performance of the LLMs in the recommendation task?

The generated data and analysis code are released on GitHub.[1]

2 Related Work

The work at hand connects to two stands of research: context-aware and LLM-based recommender systems. As for the former, various approaches have been proposed to incorporate contextual information into music recommender systems (Adomavicius and Tuzhilin, 2010; Lozano Murciego et al., 2021). Methods vary with respect to observability, availability, and acquisition of context data as well as system design (Lozano Murciego et al., 2021).

An important category of contextual information is the user's listening intent. It has been studied in conversational music recommendation from several perspectives, including the specificity of user

[1] https://github.com/hcai-mms/intent-aware-llms

Proceedings of the 4th Workshop on NLP for Music and Audio (NLP4MusA 2026), pages 44–50

requests (Zhang et al., 2025), user interaction patterns within dialogue-based systems (Jannach et al., 2021), and the use of user queries to refine or filter recommendation results (Doh et al., 2024).

Recently, large language models (LLMs) have emerged as a promising paradigm for recommender systems (Jósár, 2025; Wang et al., 2025; Lin et al., 2025; Epure et al., 2025). Owing to their strong representation and reasoning capabilities, LLMs can be employed to enhance user modeling, extract high-level semantic features, and perform scoring, ranking, or recommendation in zero- or few-shot settings, either via ranking-based approaches (Hou et al., 2024) or generative methods based on token prediction (Doh et al., 2024). Moreover, LLM-based interfaces enable users to articulate their preferences, needs, and listening intents in a more natural and expressive manner (Yun and Lim, 2025), while simultaneously increasing user agency and control over the recommendation process (Friedman et al., 2023).

In this preliminary study, we extend these lines of research by *examining the capacity of LLMs to align users' music listening histories with their underlying listening intents* for the purpose of recommending additional, intent-consistent tracks. Specifically, we evaluate whether LLMs can identify and group music tracks associated with the same listening intent, based on explicitly or implicitly expressed intent cues provided in the prompts.

3 Methodology

In this section, we describe the setup of the experiment: what data was used (Section 3.1), how the listening intent was assigned (Section 3.2), how the prompts were formed, and the text was generated (Section 3.3), and finally, how we evaluate the generated recommendations (Section 3.4).

3.1 Datasets

For our study, we required information regarding both the tracks a user listened to and their corresponding listening intents. Listening intents have been previously studied (Hausberger et al., 2025) using Spotify's Million Playlist Dataset (Chen et al., 2018). This dataset consists of 1,000,000 playlists with title and constituting track names, and covering a total of 2,262,292 unique music tracks (described by title, artist, and Spotify URI). In addition, the LFM-2b Last.fm dataset (Schedl et al., 2022) provides rich information on users and their listening sessions. We created a subset of LFM-2b containing only songs that also occur in the Million Playlist Dataset. This subset contains $119,969$ unique users, $451,728$ unique music tracks, and $361,487,213$ listening events. We additionally dropped the top and bottom 25% of users based on their count of interacted songs, resulting in a minimum of 91 and a maximum of 858 unique tracks per user. The final dataset consists of $59,906$ users.

3.2 Listening Intent Assignment

An approach to assign listening intents to tracks in the Million Playlist Dataset has previously been proposed in (Hausberger et al., 2025). The authors use a text-similarity approach to map the playlist titles to 32 psychological listening intents (e.g., "Calming" or "Support"). We extend this approach by computing the mean textual similarity of the playlist titles to each listening intent name provided in (Hausberger et al., 2025) using cosine similarity over the embeddings created by the "google/embeddinggemma-300m" encoder model (Vera et al., 2025). For each track t in the LFM-2b subset that we created, we subsequently compute the list of playlists in which t occurs. For each track, the mean similarity vector to the listening intents over all occurring playlists is calculated. We further standardized the similarity scores using z-score scaling of each listening intent over all songs to zero mean and unit standard deviation. Z-score standardization makes similarity scores comparable across intents by removing intent-specific scale and variance biases, so that selecting the maximum reflects the strongest relative evidence for an intent rather than artifacts of differing score distributions. To each song, the listening intent with the highest score is assigned.

For each intent, we create a list of users that have listened at least 10 distinct songs with that intent assigned. On average, a user is assigned to 9.0 intents. To further reduce the set of users for computational costs, we randomly sample users equally from each list for each listening intent. In total, this results in $10,670$ unique users and on average 396.4 unique users per listening intent.

Ultimately, each track t is assigned 2 binary labels depending on user and listening intent: i) *User-relevance* (t_u), i.e., if the user u has interacted with the track t, and ii) *Intent-relevance* (t_i), i.e., if the track t is assigned to a specific listening intent i.

For simplicity, we then denote as U the set of all user-relevant tracks per user ($\bigcup t_u$) and $Intent$

describes the set of all intent-relevant tracks per intent ($\bigcup t_i$); also illustrated in Fig. 1.

For each pair of user u and intent i in the final dataset, we then create four sets of tracks without overlaps: i) $U \cap Intent$: all tracks that are user- and intent-relevant, ii) $U \setminus Intent$: user-relevant but not intent-relevant tracks, iii) $Intent \setminus U$: intent-relevant but not user-relevant tracks, iv) $\neg(U \cup Intent)$: tracks that are neither user- nor intent-relevant.

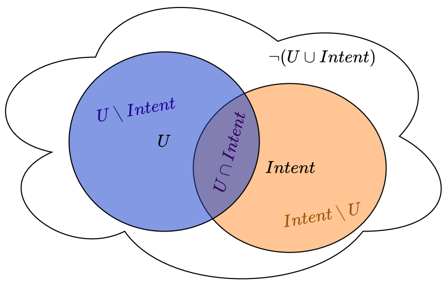

Figure 1: Sets of music track items per user.

3.3 Prompt Generation for Recommendation

The LLM under investigation was given one or two of the following information: i) 3 tracks that a user had listened to for a certain listening intent (further denoted as *music preference*), and ii) the name of the listening intent. Three variations of prompts were formed, in which the LLM got only the intent name (*explicit*), only the music preference of the user in an intent (*implicit*), or *both*. These information variations are further referred to *prompt types*. Then, it received a list of 40 candidate songs consisting of 10 randomly sampled songs from each set of tracks in the format "artist - track". The 40 songs are randomly shuffled for each prompt variation. The LLM is tasked to recommend the tracks sorted based on relevancy to the user and listening intent, and to return 10 tracks sorted in descending order. All prompt types and variations with examples can be found in the Appendix A.

For generating the recommendations, we use 5 different models ("Qwen2-7B-Instruct" (Team et al., 2024), "Mistral-7B-Instruct-v0.3" (Jiang et al., 2023), and models from the Llama-3 family (Dubey et al., 2024) ("Llama-3.2-1B-Instruct", "Llama-3.2-3B-Instruct", "Llama-3.1-8B-Instruct)), each with 3 different prompt types (explicit, implicit, and both) and 3 variations of each prompt. We set the parameters for the generation process to the following values for all LLMs: ($do\_sample = True$, $temperature = 1.0$, $top\_p = 0.9$, $repetition\_penalty = 1.0$, $num\_beams = 1$, $p = 0.9$).

The LLMs were run on four NVIDIA GeForce RTX 2080 Ti GPUs. For a batch size of 12, the average inference time per batch was highest for "Llama-3.1-8B-Instruct" (15.43 s), followed by "Mistral-7B-Instruct-v0.3" (14.91 s), "Qwen2-7B-Instruct" (12.28 s), "Llama-3.2-3B-Instruct" (10.09 s), and "Llama-3.2-1B-Instruct" (4.56 s).

3.4 Evaluation

We filter the set of recommendations for a given user to only include songs that have been given to the LLM as options to select from, and report the percentage of hallucinated tracks per model. Since the outputs may contain repeated tracks, we remove duplicates from predictions by retaining the first occurrence of each track, and finally consider the first 10 recommended tracks. We do not check if a LLM successfully recommends exact 10 music tracks. We evaluate the relevance of recommendations provided by the LLM through the Normalized Discounted Cumulative Gain metric computed on the top 10 recommended items (NDCG@10) over each of the four sets of tracks (see Section 3.2) defining the relevant items.

4 Results & Discussion

In this section, we analyze the results in two different ways. First, we look at the different sets of relevant items discussed before and evaluate how well the models ranked those individual sets (Section 4.1). Second, we examine to which extent explicitly mentioning the listening intent in the prompt influences the LLM's performance in recommending and ranking relevant versus non-relevant items (Section 4.2).

The LLMs produced on average 7.0% recommendations that did not match with any song provided in the candidate set, with the model "Llama-3.2-3B-Instruct" producing on average the most hallucinated recommendations (15.6%), followed by "Llama-3.2-1B-Instruct" (10.3 %), "Mistral-7B-Instruct-v0.3" (6.0%), "Llama-3.1-8B-Instruct" (2.0%), and the model "Qwen2-7B-Instruct" (1.0 %). 1.2 % of recommendation lists included only songs that did not match any item in the candidate list.

4.1 Ranking of Relevant Items (RQ1)

Fig. 2 shows the NDCG@10 scores of each LLM under investigation, computed using each set of tracks as relevant ones (see Section 3.2). As a baseline, we include results of a popularity-based recommender (Pop), which recommends the top k items of the candidate list with the highest count of interactions over all users, sorted, and a random recommender, which recommends the first 10 items of a random shuffled list of items. While the popularity based approach prioritizes generally popular and thus often user-relevant items, it largely fails to capture intent relevance. In contrast, all LLM-based methods consistently rank items that are both user- and intent-relevant higher than unrelated candidates, indicating their ability to model relationships between listening intent, previously consumed tracks, and candidate items. Larger models were able to prioritize user-relevant tracks higher than smaller ones, getting a better performance than random. This trend suggests that increased model capacity enhances the ability to leverage contextual and semantic information for intent-aware recommendation.

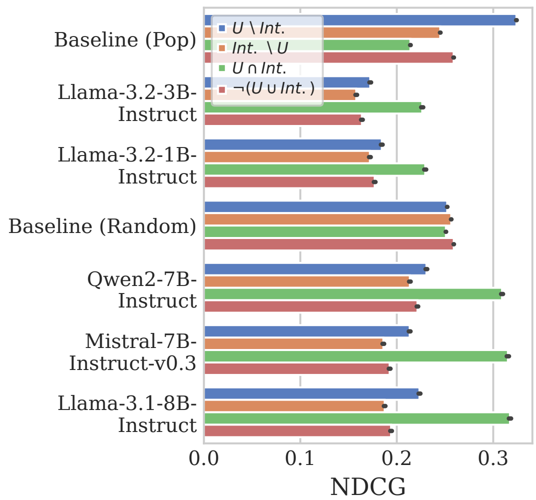

Figure 2: NDCG@10 for the baseline and 5 LLMs, for different user- and intent-relevant track sets. Error bars indicate 95% confidence intervals (CI).

4.2 Effect of Intent and Music Preference Inclusion in Prompt (RQ2)

Furthermore, we examine how the results vary with specific information provided in the prompt (explicit and/or implicit). In Fig. 3, the NDCG scores of the user- and intent-relevant items are shown when grouping by prompt type. It is seen that specifying the intent through intent name generally lowers the performance in ranking intent- and user-relevant items in comparison to only including the music tracks that have been listened to in the same listening intent. This performance gap is particularly noteworthy for larger models. It may seem counterintuitive, but it could indicate that some latent confound is picked up by the (larger) LLMs from the given music preferences alone, rather than by explicitly specifying the intent. This observation gives rise to studying other ways of classifying intent than through a limited set of playlist titles.

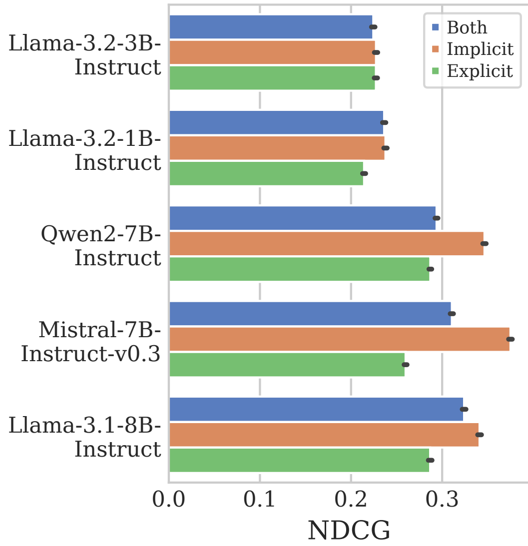

Figure 3: NDCG@10 scores computed on user- and intent-relevant item set ($U \cap Intent$), grouped by the prompt type. Error bars indicate 95% confidence intervals (CI).

5 Conclusion & Future Work

This work has explored whether LLMs are able to identify and rank relevant music tracks for specific users and listening intents. We found that LLMs are able to identify relevant items, and that implicitly indicating an intent alongside tracks the user has interacted with worked best for LLMs to rank intent- and user-relevant tracks. For future work, we will investigate other ways of identifying, encoding, and leveraging listening intent with LLMs to enhance intent-aware music recommendation systems.

6 Limitations

This work is limited by the small number of tested LLMs. We are aware that larger language models probably have a more advanced performance. Further, we acknowledge that text generation could vary depending on the prompt and run. Finally, the evaluation employs a strict string-matching heuristic for track identification, which may underestimate model performance by failing to account for minor formatting variations in the generated output.

Acknowledgments

This research was funded in whole or in part by the Austrian Science Fund (FWF): Cluster of Excellence *Bilateral Artificial Intelligence* (https://doi.org/10.55776/COE12), the doc.funds.connect project *Human-Centered Artificial Intelligence* (https://doi.org/10.55776/DFH23), and the PI project *Intent-aware Music Recommender Systems* (https://doi.org/10.55776/P36413).

References

Gediminas Adomavicius and Alexander Tuzhilin. 2010. Context-aware recommender systems. In *Recommender systems handbook*, pages 217–253. Springer.

Ching-Wei Chen, Paul Lamere, Markus Schedl, and Hamed Zamani. 2018. Recsys challenge 2018: Automatic music playlist continuation. In *Proceedings of the 12th ACM Conference on Recommender Systems*, pages 527–528.

SeungHeon Doh, Keunwoo Choi, Daeyong Kwon, Taesu Kim, and Juhan Nam. 2024. Music discovery dialogue generation using human intent analysis and large language models. *arXiv preprint arXiv:2411.07439*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.

Elena V Epure, Yashar Deldjoo, Bruno Sguerra, Markus Schedl, and Manuel Moussallam. 2025. Music recommendation with large language models: Challenges, opportunities, and evaluation. *arXiv preprint arXiv:2511.16478*.

Luke Friedman, Sameer Ahuja, David Allen, Zhenning Tan, Hakim Sidahmed, Changbo Long, Jun Xie, Gabriel Schubiner, Ajay Patel, Harsh Lara, et al. 2023. Leveraging large language models in conversational recommender systems. *arXiv preprint arXiv:2305.07961*.

Anna Hausberger, Hannah Strauss, and Markus Schedl. 2025. Exim: Exploring intent of music listening for retrieving user-generated playlists. In *Proceedings of the 2025 ACM SIGIR Conference on Human Information Interaction and Retrieval*, pages 348–357.

Yupeng Hou, Junjie Zhang, Zihan Lin, Hongyu Lu, Ruobing Xie, Julian McAuley, and Wayne Xin Zhao. 2024. Large language models are zero-shot rankers for recommender systems. In *European Conference on Information Retrieval*, pages 364–381. Springer.

Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2021. A survey on conversational recommender systems. *ACM Computing Surveys (CSUR)*, 54(5):1–36.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7b. *Preprint*, arXiv:2310.06825.

Petra Jósár. 2025. Large language models for intent-aware music recommendation/author petra jósár, bsc.

Jianghao Lin, Xinyi Dai, Yunjia Xi, Weiwen Liu, Bo Chen, Hao Zhang, Yong Liu, Chuhan Wu, Xiangyang Li, Chenxu Zhu, et al. 2025. How can recommender systems benefit from large language models: A survey. *ACM Transactions on Information Systems*, 43(2):1–47.

Alvaro Lozano Murciego, Diego M Jiménez-Bravo, Adrian Valera Roman, Juan F De Paz Santana, and María N Moreno-García. 2021. Context-aware recommender systems in the music domain: A systematic literature review. *Electronics*, 10(13):1555.

Thomas Schäfer, Peter Sedlmeier, Christine Städtler, and David Huron. 2013. The psychological functions of music listening. *Frontiers in psychology*, 4:511.

Markus Schedl, Stefan Brandl, Oleg Lesota, Emilia Parada-Cabaleiro, David Penz, and Navid Rekabsaz. 2022. Lfm-2b: A dataset of enriched music listening events for recommender systems research and fairness analysis. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval*, pages 337–341.

Qwen Team et al. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2(3).

Henrique Schechter Vera, Sahil Dua, Biao Zhang, Daniel Salz, Ryan Mullins, Sindhu Raghuram Panyam, Sara Smoot, Iftekhar Naim, Joe Zou, Feiyang Chen, et al. 2025. Embeddinggemma: Powerful and lightweight text representations. *arXiv preprint arXiv:2509.20354*.

Shijia Wang, Tianpei Ouyang, Yunfan Zhou, Qiang Xiao, Yintao Ren, Yifei Pan, Fangjian Li, and Chuanjiang Luo. 2025. Enhanced emotion-aware music recommendation via large language models. In *Proceedings of the 31st ACM SIGKDD Conference on*

Knowledge Discovery and Data Mining V. 2, pages 4986–4994.

Sojeong Yun and Youn-kyung Lim. 2025. User experience with llm-powered conversational recommendation systems: A case of music recommendation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–15.

Junjie Zhang, Ruobing Xie, Yupeng Hou, Xin Zhao, Leyu Lin, and Ji-Rong Wen. 2025. Recommendation as instruction following: A large language model empowered recommendation approach. *ACM Transactions on Information Systems*, 43(5):1–37.

A Prompts

This section presents the prompts used for text generation. In each prompt, the placeholder INTENT_SONGS is replaced with the three songs previously listened to by the user, CANDIDATE_LIST is replaced with the set of candidate songs, and INTENT_NAME is replaced with the corresponding intent label. Depending on the prompt type (explicit, implicit, or both), the prompt provided to the LLM includes the three listened songs (implicit information), the intent name (explicit information), or both.

A.1 System Prompt

"System: You are a music recommender. You recommend music tracks for a user and an intent from a set of music tracks. You get a user listening history, the intent and return a list of music tracks recommended for the intent and user. Recommend the best suitable track first. Return 10 recommendations. Return the track in the format: Recommendations: "artist - track" split with ",". You get the songs to sort in the format "ID": "artist - track", return 10 tracks recommended for the intent and user. Analyze the music tracks given based on how relevant they are for an intent and user and recommend 10 tracks best, most relevant tracks first. Recommend only songs before 2020. Recommend best suited songs first."

A.2 User Prompts

1. 'Intent: "INTENT_NAME"
   User Listened Songs in Intent: "INTENT_SONGS"
   Song Options: "CANDIDATE_LIST"'

2. 'User Listening History: "INTENT_SONGS"
   Intent: "INTENT_NAME"
   Song to rank: "CANDIDATE_LIST"'

3. 'Intent: "INTENT_NAME"
   Songs listened in Intent by User: "INTENT_SONGS"
   Song Options to rank: "CANDIDATE_LIST"'

A.3 Examples

A.3.1 Explicit

Intent: Fitness Motivation
Song Options: 0: Nonpoint - In The Air Tonight, 1: Omar Souleyman - Shift Al Mani, 2: Cavalera Conspiracy - Bloodbrawl, 3: Japandroids - The Nights of Wine and Roses, 4: Danzig - I Don't Mind The Pain, 5: Nine Inch Nails - The Day the World Went Away, 6: Misfits - Horror Hotel, 7: Social Distortion - Highway 101, 8: Airbourne - Raise the Flag, 9: Black Rebel Motorcycle Club - Spread Your Love, 10: Diablo - Icaros, 11: Social Distortion - Making Believe, 12: Winds of Plague - Legions, 13: A Day to Remember - We Got This, 14: We Came As Romans - To Plant a Seed, 15: Alexisonfire - We Are The End, 16: Kiss - Tough Love, 17: Sepultura - Subtraction, 18: Schiller - Solitude, 19: Kiss - 100,000 Years, 20: The Real McKenzies - Drink the Way I Do, 21: Ministry - Jesus Built My Hotrod, 22: Dropkick Murphys - It's A Long Way to The Top (If You Wanna Rock N Roll), 23: Jennifer Lopez - Papi, 24: Mad Sin - U.F.O., 25: Black Light Burns - Animal, 26: Cavalera Conspiracy - Black Ark, 27: Macklemore - Levitate (feat. Otieno Terry), 28: Rancid - 1998, 29: Twenty One Pilots - Heathens, 30: Tiger Army - As The Cold Rain Falls, 31: Linkin Park - Crawling, 32: The Clancy Brothers - Johnson's Motor Car, 33: Big Bad Voodoo Daddy - Minnie The Moocher, 34: The Gaslight Anthem - Handwritten, 35: Trapt - Headstrong, 36: Danzig - She Rides, 37: Sepultura - Desperate Cry, 38: Ramones - Substitute, 39: Karmin - Brokenhearted"

A.3.2 Implicit

'User Listened Songs in Intent: "Jenny Lewis - Pretty Bird", "Feist - Honey Honey", "The Format - Pick Me Up"
Song Options: 0: Love Unlimited Orchestra - Satin Soul, 1: Kent - Sverige, 2: Alkaline Trio - Nose Over Tail, 3: Zeigeist - Humanitarianism, 4: Old Crow Medicine Show - Cocaine Habit, 5: The Submarines - Maybe, 6: Teebs - Clapstick, 7: of Montreal - ErosÉntropic Tundra, 8: Matt Pond PA - I Want To See The Bright Lights Tonight, 9: Foxes - Devil Side, 10: Bhagavan Das - Hanuman Chalisa, 11: The Danse Society - Somewhere, 12: Slum Village - Players, 13: Jeff Beck - Over The Rainbow, 14: Neko Case - Fever, 15: Someone Still Loves You Boris Yeltsin - Yr Broom, 16: Andhim - Bermudachords, 17: Jenny Lewis - The Next Messiah, 18: Grouper - Headache, 19: of Montreal - Chrissy Kiss the Corpse, 20: Air - Missing the Light of the Day, 21: The Knife - Pass This On, 22: Starlight Mints - Brass Digger, 23: Mates of State - What I Could Stand For, 24: Robert Ellis - Perfect Strangers, 25: Grouper - Being Her Shadow, 26: Kesha - Finding You, 27: The Little Ones - Morning Tide, 28: Jenny Lewis - Bad Manś World, 29: Mates of State - Parachutes (Funeral Song), 30: Bright Eyes - Southern State, 31: Boys Noize - Drummer, 32: Paul Kalkbrenner - Torted, 33: Ted Leo and the Pharmacists - The Gold Finch and the Red Oak Tree, 34: Someone Still Loves You Boris Yeltsin - Anna Lee, 35: Regina Spektor - Human of the Year, 36: The Go! Team - Bottle Rocket, 37: Elvis Presley - Heartbreak Hotel, 38: Manowar - Animals, 39: Ani DiFranco - You Had Time'

A.3.3 Both

Intent: Nostalgia
User Listened Songs in Intent: "Live - I Alone", "Delerium - Paris", "When In Rome - The Promise"
Song Options: 0: Kim Carnes - Bette Davis Eyes, 1: Kut Klose - Surrender, 2: Moby - Hymn, 3: Britney Spears - E-Mail My Heart, 4: Counting Crows - Rain King, 5: Razed in Black - Am I 2 Blame?, 6: Joni Mitchell - Free Man In Paris, 7: Nice Smooth - Funky For You, 8: Marc Streitenfeld - Earth, 9: Lynyrd Skynyrd - Free Bird, 10: Michael Jackson - I just cant́ stop loving, 11: Orgy - Blue Monday, 12: Joss Stone - Sideway Shuffle, 13: Marvin Gaye - Distant Lover, 14: Depeche Mode - Personal Jesus (Telephone Stomp Mix), 15: Van Halen - Top of the World, 16: Beastie Boys - Triple Trouble, 17: Marc Streitenfeld - Friend From The Past, 18: Marconi Union - Breathing With Assistance, 19: The Temptations - Some Enchanted Evening, 20: Kelly Clarkson - Anytime, 21: Beastie Boys - Pass the Mic, 22: Bob Dylan - Subterranean Homesick Blues, 23: Jens Gad - Cape Blanc, 24: Seabird - Cottonmouth (Jargon), 25: Beastie Boys - Hey Ladies, 26: Lifehouse - First Time, 27: The Beautiful South - Everybodyś Talkin; 28: A Tribe Called Quest - Oh My God, 29: Chingy - Right Thurr, 30: Judas Priest - On the Run, 31: Enigma - Mea Culpa - Part II (Catholic Version), 32: The Cars - Drive, 33: Kelly Clarkson - A Moment Like This, 34: Beastie Boys - Shake Your Rump, 35: Audiomachine - House of Cards, 36: Razed in Black - Blush, 37: The Cranberries - How, 38: Marc Streitenfeld - Hyper Sleep, 39: Louie Culture - No Gal"