

MUSIC DISCOVERY DIALOGUE GENERATION USING HUMAN INTENT ANALYSIS AND LARGE LANGUAGE MODELS

SeungHeon Doh^b Keunwoo Choi[‡] Daeyong Kwon^b
 Taesu Kim[‡] Juhan Nam^b

^b Graduate School of Culture Technology, KAIST, South Korea

[‡] Prescient Design, Genentech, New York, NY, USA

[‡] Department of Industrial Design, KAIST, South Korea

{seungheondoh, ejmj63, tskind77, juhan.nam}@kaist.ac.kr, choi.keunwoo@gene.com

ABSTRACT


A conversational music retrieval system can help users discover music that matches their preferences through dialogue. To achieve this, a conversational music retrieval system should seamlessly engage in multi-turn conversation by 1) understanding user queries and 2) responding with natural language and retrieved music. A straightforward solution would be a data-driven approach utilizing such conversation logs. However, few datasets are available for the research and are limited in terms of volume and quality. In this paper, we present a data generation framework for rich music discovery dialogue using a large language model (LLM) and user intents, system actions, and musical attributes. This is done by i) dialogue intent analysis using grounded theory, ii) generating attribute sequences via cascading database filtering, and iii) generating utterances using large language models. By applying this framework to the Million Song dataset, we create – **LP-MusicDialog**, a Large Language Model based Pseudo Music Dialogue dataset, containing over 288k music conversations using more than 319k music items. Our evaluation shows that the synthetic dataset is competitive with an existing, small human dialogue dataset in terms of dialogue consistency, item relevance, and naturalness. Furthermore, using the dataset, we train a conversational music retrieval model and show promising results.¹

1. INTRODUCTION

In recent years, conversational systems have emerged as a promising solution to enhance user experience in various domains [1–4], including conversational music retrieval and recommendation [5, 6]. The goal of a conversational

¹ Our code is available at <https://github.com/seungheondoh/lp-music-dialog/>

Step1. Intent Analysis from Human Dialogue



Step2. Cascading Music Database Filtering



Step3. Music Discovery Dialogue Generation

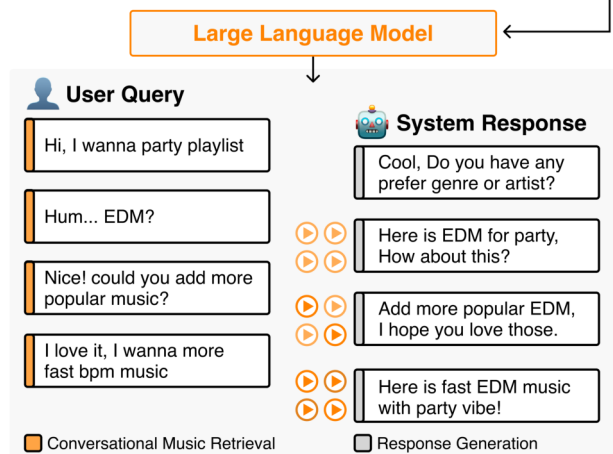


Figure 1. The generation process of pseudo musical dialogue.

music system is to assist users in finding their desired music through dialogues. Such a system should possess three key capabilities: i) to understand the *intents* and *musical needs* of users from their queries expressed in natural language, ii) to generate responses and facilitate human-like interaction, iii) to find music that aligns with the user’s preferences by taking previous dialogues into account.

Currently, the primary challenge of developing a conversational music retrieval system is the scarcity of large-scale public datasets. Chaganty et al. [5] introduce the Conversational Playlist Curation Dataset (CPCD). This crowd-sourced dataset comprises human-to-human dialogues that simulate the process of music discovery. How-



ever, as it relies on a manual process, the dataset is small and exhibits biases from the music streaming platforms used by the recommenders. To address this problem, Leszczynsk et al. [6] propose a dialogue generation framework through random walks in the music-text joint embedding space and dialogue inpainting [7]. However, this approach requires a high-quality music-text joint embedding and needs to use template-based system responses as input. As a result, the system’s responses are always composed of limited format utterances, leading to low naturalness in human evaluation.

In this paper, we introduce a framework for generating *human-like* music discovery dialogues using intent and a large language model (LLM). The proposed framework is based on the existing method [6], but we address their limitations by employing *cascading music database filtering* instead of a joint embedding and extensive *intent analysis* for naturalness. Using the grounded theory approach [8], we analyze a dataset of human music discovery dialogues and develop taxonomies for user intents, system actions, and musical attributes relevant to the task of music discovery. Furthermore, we introduce a model-free attribute sequence generation by applying cascading filtering to a multi-label music annotation database. Finally, we synthesize music discovery dialogues through an LLM using the created attribute sequences and human intents/actions.

Our contributions are threefold: First, we analyze music discovery dialogues and propose a taxonomy. Second, we introduce the LP-MusicDialog dataset, a large-scale synthetic dialogue dataset created using human intent and an LLM. Third, we present extensive objective and subjective evaluations to demonstrate the effectiveness of LLM-based pseudo-music dialogues.

2. RELATED WORK

Recently, there has been some progress in conversational music systems including language-based music retrieval [9–11]. However, existing systems are often limited to single-turn tasks. This means users are not able to refine their queries to obtain a highly satisfactory outcome. Towards multi-turn dialogues, Chaganty et al. [5] released the Conversational Playlist Curation Dataset (CPCD), which comprises 917 dialogues averaging 5.7 turns each. Unlike single-turn retrieval, conversational retrieval takes into account previous chat history to find relevant items. The model in [5] aggregates the context of history embedding and the current query embedding using average pooling, then uses contrastive loss to maximize the similarity between them. However, their model showed limited performance due to the small size of CPCD.

A solution to the data scale issue is synthesizing data using existing datasets and language models. Recently, synthetic datasets that bridge natural languages and music have been proposed to enhance music understanding [12, 13], captioning [13, 14], reasoning [13], and retrieval [11]. For the conversational music retrieval, Leszczynski et al. [6] proposed a two-stage data synthesis framework: musical attribute sequence generation via

random walk and utterance generation through dialogue inpainting [7]. The musical attribute sequence represents the evolution of user queries over turns (e.g., ask for workout music in the first turn, then refine the results to be also pop music). During utterance generation, a language model creates user queries using system responses as input, which include sampled musical attributes and a static template.² As a result, they created one million multi-turn music discovery dialogues: TtW Music, leveraging a private playlist dataset. However, this approach encounters issues with model errors in the music-text joint embedding space used for the random walk method and faces challenges with response consistency due to the reliance on manually created templates for system responses.

Deep understanding in music query has to be preceded aforementioned data generation. So far, query understanding has primarily focused on describing musical needs. Downie and Cunningham [15] analyzed 161 music queries and categorized them into 1) information needs, 2) desired outcomes, 3) intended uses for the information, and 4) social and contextual elements. Bainbridge et al. [16] utilized the grounded theory approach to analyze 502 real-world music queries, expanding upon prior research with 10 types of need descriptions. Lee [17] analyzed 1,705 Google Answers queries to propose a refined taxonomy for information needs and searching behavior. Despite these efforts, previous studies have been limited to the information needs (musical attributes) contained in queries, and intent in multi-turn queries has not received significant attention.

3. DIALOGUE INTENT ANALYSIS

3.1 Taxonomy Development

For the dialogue-specific music discovery taxonomy, we analyze the existing human-to-human music dialogue dataset (CPCD [5]) using the grounded theory approach [8]. Grounded theory is a qualitative approach that creates refined theory from unstructured real-world data. In detail, (1) we adopt the taxonomy from previous research as our initial taxonomy. For user intent and system action, we use the *conversational movie intent* taxonomy [18], and for musical needs, we use *music feature* taxonomy [17] as the initial taxonomy. (2) Three authors annotate ten dialogues using the initial taxonomy and discuss the limitations of the existing taxonomy. (3) We update the taxonomy and annotate a new randomly sampled five dialogues. This cycle of proposing, refining, and annotating was completed three times to ensure our taxonomy could capture the full range of scenarios present in the dialogue samples.

3.2 Analysis Results

3.2.1 Taxonomy for User Intents

In Table 1, we categorize user intents into four main categories, with eight detailed sub-intents. The **Initial Query**

² For example, “Of course! Let me add some songs described as <musical attributes>. What else?”

User Intent	Description	Example	%
<i>Start Dialogue</i>			
Initial Query	User initiates the inquiry with a specific request.	"Hi, I want to create a playlist for hiking."	18.5
Greeting	User initiates the dialog, often with greeting words.	"Hello, I would like... / "Good Morning! Let's start..."	12.7
<i>Item Discovery Query (Retrieval / Recommendation)</i>			
Positive Filter	User requests to include an additional criterion.	"I would like to add a bit of Rihanna."	76.7
Negative Filter	User requests to negative criterion in this turn.	"I do prefer them to not have any lyrics"	3.7
Continue	User requests for more songs with the current criteria.	"Those are good songs. More like these would be great."	6.8
<i>Item Understanding Query (Question Answering)</i>			
Item Attribute Question	User questions attributes of music.	"Do you know where Samer (Artist) is from?"	0.2
<i>Feedback Response</i>			
Accept Response	User responds positively to the recommendations.	"Thank you, they are perfect"	44.4
Reject Response	User responds negatively to the recommendations.	"I still didn't get any song suggestions..."	4.8

Table 1. Taxonomy for user intents. % represents the percentage of occurrences within the user query.

System Action	Description	Example	%
<i>Request</i>			
Feedback Request	System requests the user to evaluate recommendations.	"What about these ones?"	12.8
Detail Attribute Request	System requests for the user's needs or desires for recommendations.	"Are there any particular artists you want to see?"	20.8
<i>Item Discovery Response (Retrieval / Recommendation)</i>			
Passive Recommendation	System recommends music based on user's preferences.	"Here are some pop songs for kids."	71.4
Active Recommendation	System proactively recommends music without being explicitly asked.	"Nice! I added a couple more."	14.9
<i>Item Understanding Response (Question Answering)</i>			
Item Attribute Answer	System answers to user's question with musical attributes.	"Frederic Chopin was a Polish composer..."	0.1
<i>General Response</i>			
Parroting Response	System responds to the user's inquiry by mirroring.	"Here's some picks from The Who."	25.4
Sympathetic Response	System responds to the user's inquiry with human-like sympathy.	"Excellent! Looks like a great Kickstart!"	57.2

Table 2. Taxonomy for system actions. % represents the percentage of occurrences within the system response.

is part of the 'start-dialogue' category, serving as the kick-off point for recommendation dialogues. At times, users start interactions with **Greetings** to initiate the system into a more engaging mode of communication. The 'item-discovery-query' captures user preferences for music retrieval and recommendations, subdivided into **Positive Filter** for add preferences, **Negative Filter** for discarding existing ones, or **Continue** to sustain the current preferences. **Item Attribute Question**, where users inquire about the precise attributes of music tracks, such as their genre, mood, tempo, and key/mode. Feedback responses are outlined, with users either expressing satisfaction (**Accept Response**) or dissatisfaction (**Reject Response**) to the recommended music.

3.2.2 Taxonomy for System Actions

Table 2 shows the system action taxonomy, structured into four primary categories, and seven specific intents. The 'request' category enhances the search experience, either through post-recommendation **Feedback Request** to gauge satisfaction or **Detail Attribute Request** to clarify vague or incomplete user queries. To address user queries for music discovery, the system can adopt a **Passive Recommendation** approach to comply with user requests. This approach is similar to a retrieval task because there is an explicit query from the user. The system also proactively engage through **Active Recommendation**. It is similar to a recommendation task because it implicitly uses the context of the dialogue even without an explicit query from the user.

In **Item Attribute Answer**, the system responds to users' inquiries about specific musical attribute questions

such as genre, mood, tempo, and key/mode. For general responses, the system mimics user requests in its recommendations (**Parroting Response**) to affirm that user preferences are being considered, or it may adopt a more empathetic stance (**Sympathetic Response**) to foster a more human-like interaction.

3.2.3 Taxonomy for Musical Attribute

Musical attributes such as genre, mood, and artist are closely related to user preferences. They are categorized into (1) objective metadata produced when a track is registered on the platform, (2) subjective similarity with music entities, (3) user & listening context, and (4) music content information (Table 3). Metadata is mainly associated with entity recognition [19], where **Track** refers to requests for a single music recording entity, **Artist** denotes user requests for tracks released by a specific artist, **Year** reflects the era/year in which a piece of music was released, and **Popularity** indicates the level of attention a piece of music has received. **Culture** is related to the national or regional style of the music, often linked to the artist's nationality. Unlike objective metadata, **Similar Track** and **Similar Artist** are subjective musical attributes, representing connections to other track or artist entities. The user % listening context consist of **User**, which is related to the demographics of the listener, and **Theme**, which encompasses the location, time, usage, and activities associated with listening. **Mood** refers to the emotional tone conveyed by the music or the listener's emotional state. Lastly, the categories tied closely to music content itself include **Genre**, which relates to high-level music form and style, and aspects associated with timbre such as **Instru-**

Musical Attribute	Description	Example	%
<i>Metadata</i>			
Track	A single musical work, recording, performance	Can you add montero from lil nas?	4.5
Artist	A creator or performer of music.	How about some more Justin Timberlake?	43.1
Year	The time of music’s creation or release	Could you throw in some 90s hip hop?	6.6
Popularity	The widespread acclaim of music.	Awesome! How about more male disco hits?	0.7
Culture	The national or regional influences on music.	i like Nigerian songs	0.9
<i>Similar with Music Entity</i>			
Similar Track	The song is similar to the specific track.	Can I have more songs similar to Jessie’s Girl?	1.1
Similar Artist	The song is similar to the specific artist.	Adele and Sia type of music	5.0
<i>User & Listening Context</i>			
User	The listeners characterized by demographics.	I want to create a fun playlist for the kids	1.9
Theme	A context for music listening related to location, time, usage, and activity.	hi please i need a playlist to workout with	15.2
Mood	The emotional tone conveyed by music, or the emotional state of user.	I want to create playlist for when I’m sad	6.7
<i>Music Content Information</i>			
Genre	A category of music characterized in form, style, or subject matter.	I would love to start with classical music	17.2
Instrument	A tool or device designed to produce musical sounds	What other quartet music do you suggest?	1.1
Vocal	The singing voice in music, including styles, techniques, and expressions	try some female vocalists too	1.2
Tempo	The speed or pace at which a piece of music is played	Add me some slow slow songs from juice wrld	1.3
Key / Mode	The tonal information of music	(Not appear)	0.0

Table 3. Taxonomy for musical attributes. % represents the percentage of occurrences within the user query.

ment and **Vocal**, as well as **Tempo** related to rhythm, and **Key/Mode** for tonal information.

3.2.4 Dialogue intent annotation

After the development of the taxonomy, the three annotating authors proceeded to annotate user intent, system action, and musical attributes according to the proposed taxonomy within a new sampled 30 CPCD [5] dialogues. To measure the level of agreement among the annotators for the degree of concurrence, we employed Krippendorff’s alpha [20]. The results showed high agreement levels, with Krippendorff’s alpha scores of 0.83 for user intent, 0.85 for system action, and 0.71 for musical attributes. Following this agreement phase, each annotator independently annotated a portion of the total 888 CPCD dialogues, excluding 29 error samples out of the total 917 dialogues.

3.3 Intent Analysis and Findings

The proportion of categories in each taxonomy is on the rightmost columns of Tables 1 – 3. For user intent (in Table 1), the majority of item discovery queries progress through the *Positive Filter* (76.7%), whereas the opposite, *Negative Filter*, constitute a smaller portion (3.7%). *Item Attribute Questions* are almost non-existent (0.2%). The majority of recommendation are accepted (44.4% of *Accept Response* vs 4.8% of *Reject Response*). Regarding system actions, a notable observation is that the system predominantly provides recommendations passively, rather than actively offering suggestions (71.4% vs 14.9%). In the case of general responses, the system favors eliciting human-like conversations through sympathetic responses over merely parroting back information (57.2% vs 25.4%). In the case of music, categories such as artist (43.1%), genre (17.2%), and theme (15.2%) occupy a significant portion of the user queries. This may suggest that in a *Wizard of Oz* setting [21], where the recommender utilizes music streaming platforms, the platform’s search support may be limited to these categories. Conversely, similarity queries or music content queries beyond genre are rarely employed.

4. MUSIC DISCOVERY DIALOGUE GENERATION

In this section, we describe the framework for generating music discovery dialogue consisting (1) attribute sequence generation and (2) utterance generation. This two-step approach is an improved version of previous research [6] with model-free attribute sequence generation by *Cascading music database filtering*, and utterance generation by a LLM and *Intent Analysis*.

4.1 Musical Attribute Sequence Generation

For attribute sequence generation, we employ a series of cascading filters on a multi-label annotated dataset to extract samples with overlapping semantics. Inspired by [22], which utilized *functional programs* to construct a visual reasoning dataset, we apply this approach to the domain of music discovery dialogue using user intents such as add filter, remove filter, and continue. Figure 2 illustrates an example of a cascading data filter. For example, a user may initially request songs in the *EDM* genre, narrow it down to a *party* theme, and finally specify a need for *fast tempo* music. This sequence can be represented by a series of connected functional program filters: *filter(tempo:fast, filter(theme:party, filter(genre:edm, database)))*. This method requires the types of filters and musical attributes. We derive filter types from the annotated intents (Section 3.2). For musical attributes, we initially conducted random sampling. Subsequently, to ensure diverse sampling with high co-occurrence with previous attributes, we employed top-*k* sampling ($k=20$), which involves randomly selecting a word from the top *k* words with the highest frequency. During top-*k* sampling, we include attributes from the *metadata* and *similar with music entity* categories.

4.2 Utterance Generation via Language Model

The sequence of musical attributes and annotated intent becomes prompts for LLMs to generate user and system

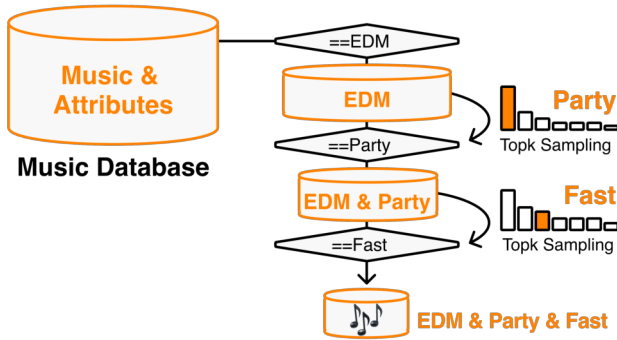


Figure 2. Cascading database filtering.

utterances. Our proposed utterance generation follows the formulation: $x_{text} = f_{LLM}(\mathcal{P}, y_{intent}, y_{music})$, where y_{intent} and y_{music} refer to the annotated intent/action (Sec 3.2) and the sampled musical attribute (Sec 4.1), respectively, and \mathcal{P} is the prompt for dialogue generation. We assumed that dialogue intent and musical attribute can serve as effective conditions for human-like dialogue generation.

5. DATASET: LP-MusicDialog

5.1 Data Source

To construct our synthetic dialogue dataset, we utilize the Million Song Dataset (MSD) [23], which has rich metadata including track details, artist information, release year, and artist familiarity. We quantize year and artist familiarity into decades and popularity, respectively.³ To cover a wide range of musical attributes as listed in Table 3, we interlink multiple annotation datasets using track IDs. For culture, mood, theme, genre, instrument, and vocals, we utilize tags from Tagtraum [24], Last.fm [25, 26], and AllMusic [27].⁴ For similar track attribute, we incorporate merged data from the Art of the Mix playlist [28] and EchoNest taste profiles. We utilize a weighted matrix factorization technique [29] to create a similarity matrix of item vectors. We then annotate the top- k similar tracks for each track ($k=128$). For artist similarity, we use cultural similarity annotation from the OLGA [30] dataset. Finally, for key/mode and tempo, we extract beats per minute (BPM), 24 key/mode using the pretrained classifier [31, 32].⁵

5.2 Creation Process

Based on the proposed pseudo dialogue generation method, we created LP-MusicDialog, an LLM-based Pseudo Music Dialogue dataset. We integrate user intents and system actions annotated in CPCD dialogues (Sec.3.2) with musical attributes obtained through cascading music database filtering (Sec.4.1) as inputs for GPT 3.5-turbo. At each cascading filtering step, we randomly sample 10

³ We categorize the top 10% of artist familiarity as high popularity, 10-30% as mid popularity, and the lower 30% as low popularity.

⁴ As the ‘style’ category in AllMusic is structured as sub-genres, we merged it with ‘genre’ category.

⁵ BPM was quantized into three text label: songs below 70 BPM were classified as slow, those between 70 BPM and 130 BPM as moderate, and those above 130 BPM as fast.

Data Source	# of Track	Musical Attributes
Million Song Dataset [23]	1,000,000	Track, Artist, Year, Popularity
TagTraum [24]	280,831	Genre, Instrument, Vocal, Mood, Theme, Culture
Last.fm [25, 26]	344,865	
AllMusic [27]	507,435	
Art of the Mix [28]	119,686	Similar Track
TasteProfile [23]	380,462	
OLGA [30]	542,364	Similar Artist
Madmom Key/Mode [32]	992,525	Key/Mode
Madmom Tempo [31]	978,759	Tempo

Table 4. Data sources for the dialogue generation

	CPCD [5]	TtW [6]	LPMD (Ours)
# of Dialog	917	1,037,701	287,675
# of Tracks	106,736	332,594	391,465
# of Vocab	9872	N/A	105,832
Avg.# of turns	5.7	5.6	4.97
Avg. query len.	54.4	80.3	63.8
Avg. response len.	45.8	N/A	87.8
Public Available	Yes	No	Yes

Table 5. Statistics of conversational music retrieval datasets. LPMD stands for LP-MusicDialog.

tracks to link with the dialogue turn. As a result, we acquire 287,675 user query, system response, music item triplets for each turn of the dialogue.

As detailed in Table 5, LP-MusicDialog is significantly larger and more diverse than existing datasets. Compared to the human dialogue dataset CPCD [5], it contains $\times 313$ larger dialogues, a more than $\times 10$ diverse vocabulary, and nearly four times more tracks. While remaining open to the public, our dataset is on par with a private dataset, TtW [6], in many aspects, with plans to expand further upon confirming active usage. In contrast, LP-MusicDialog is not only publicly available but also offers an extensive collection of connected tracks. Figure 3 shows the musical attribute ratio in dialogue. Unlike CPCD [5], where a significant portion is occupied by artist and genre due to platform bias, LP-MusicDialog shows a higher proportion of dialogues concerning music content and user listening context. Although the proportion of queries for entity recognition such as track and artist has decreased, the percentage of similarity queries has increased.

5.3 Human Evaluation

Following previous work [6], we assess the quality of our generated data through human evaluation, focusing on three key aspects: 1) Consistency - evaluating if the user preferences are coherent across dialogue turns; 2) Relevance - determining the alignment between the retrieved music items and the user query; 3) Naturalness - assessing the likelihood of such a conversation occurring in real life. Unlike previous work, we adhere to mean opinion score (MOS) that uses a 5-point Likert scale instead of a 3-point scale. A total of 26 raters evaluated 10 randomly sampled dialogues each. Within 260 total ratings, we only reported dialogues assessed by three or more raters.

Table 6 presents the MOS evaluations for dialogues from CPCD [5] and LP-MusicDialog. To understand the




Figure 3. The ratios of musical attributes in music discovery dialogues.

	Consistency	Relevance	Naturalness
LP-MusicDialog	3.79 ± 0.56	4.04 ± 0.74	3.87 ± 0.57
+ Intent / Action	3.88 ± 0.87	4.05 ± 0.53	4.24 ± 0.38
Human Dialogue [5]	3.90 ± 0.70	4.16 ± 0.68	4.29 ± 0.43

Table 6. Mean opinion scores of the generated dialogues (LP-MusicDialog) and human dialogues (CPCD [5])

impact of user intent and system action, we conducted an ablation study synthesizing dialogues with only musical attributes and prompts. Comparing the first and second rows, we find minimal differences in consistency and relevance, as both sets of generated dialogues utilize identical musical attributes. However, a notable distinction arises in naturalness, suggesting that LLMs can foster more human-like dialogue synthesis by incorporating intents and actions. In comparing human dialogues with our generated dialogues, we found that the generated dialogues perform comparably to the human dataset across all three metrics, within the standard deviation.

6. CONVERSATIONAL MUSIC RETRIEVAL

In this section, we present a benchmark of conversational music retrieval models. Unlike prior studies [5, 6] that have music embeddings solely relying on metadata-based text modality,⁶ we expand it to include audio modality. We use a pre-train audio-text joint embedding model (TTMR++ [11]), that consists of a text encoder that handles user queries and system responses and an audio encoder that takes music tracks. We freeze TTMR++ and add a trainable MLP layer for both text and audio encoder. To handle chat history, we use a chat embedding created by average-pooling the current query, previous queries, responses, and music embeddings. Two encoders are trained to maximize the cosine similarity between the chat embedding and target music embedding using the InfoNCE [33] loss. In the inference stage, we extract chat embeddings for each turn in the same way as in the training stage and measure the similarity score with all music embeddings in the train and test splits. Based on the similarity score, we retrieve the most similar k items by nearest neighbor search.

We chose the CPCD dataset as the compared dataset and report Hit@ K as evaluation metric ($k=\{10,20,100\}$).⁷ Our baseline models are as follows: (1) BM25 [34], as a sparse retrieval baseline; (2) Contriever [35], an unsu-

⁶ {title} by {artist} from {album}

⁷ Excluding the lost audio due to YouTube crawling, we use 98,738 out of 106,736 tracks for evaluation. We use the official evaluation codebase: <https://github.com/google-research-datasets/cpcd>

Model	Type	Hit@10	Hit@20	Hit@100
BM25 [34]	-	0.180	0.251	0.433
Contriever [35]	Zeroshot	0.176	0.255	0.344
TTMR++ [11]	Zeroshot	0.201	0.275	0.505
+ LPMD only	Finetune _{OD}	0.173	0.253	0.479
+ CPCD only	Finetune _{ID}	0.209	0.295	0.530
+ LPMD & CPCD	Finetune _{OD+ID}	0.219	0.304	0.533

Table 7. Conversational music retrieval performance on CPCD Dataset. **OD** stands for out of domain. **ID** stands for in domain.

pervised dense retrieval baseline; and (3) TTMR++ [11], an audio-text joint embedding baseline. Table 7 shows the performance of conversational retrieval. Among the baselines, TTMR++ shows superior performance over BM25 and Contriever, highlighting the importance of the audio modality in the music domain. Furthermore, training TTMR++ with only LP-MusicDialog (i.e., inter-dataset evaluation) somewhat leads to a performance decrease. This is presumably due to the musical entity difference between the two datasets. Specifically, the LP-MusicDialog dataset derives from the MSD, which contains music up to the year 2010, while the CPCD dataset includes music extending up to the year 2023. Nonetheless, training with both LP-MusicDialog and CPCD results in performance improvement over training only with CPCD, suggesting the usefulness of the proposed dataset.

7. CONCLUSION

We proposed a musical dialogue generation approach with i) dialogue intent analysis using the grounded theory, ii) generating attribute sequences via cascading database filtering, and iii) generating utterances using a large language model. Our intent analysis underpins the synthesis of human-like conversations, demonstrating strengths in naturalness. Cascading filtering allows us to utilize music attributes from external music databases to generate new dialogues. The outcome, the proposed **LP-MusicDialog** covers a broader range of musical attributes and aids in the conversational music retrieval task.

However, the proposed methods have several limitations. The first is that cascading filtering is sensitive to annotation errors [36]. The second is that top-k sampling, by following the tag distribution, inevitably leads to a data imbalance problem. We hope that these limitations of cascading filtering will be addressed in future research by incorporating balanced sampling.

8. REFERENCES

- [1] K. Christakopoulou, F. Radlinski, and K. Hofmann, "Towards conversational recommender systems," in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016.
- [2] R. Li, S. Ebrahimi Kahou, H. Schulz, V. Michal-ski, L. Charlin, and C. Pal, "Towards deep conversational recommendations," *Advances in neural information processing systems (NeurIPS)*, 2018.
- [3] D. Jannach, A. Manzoor, W. Cai, and L. Chen, "A survey on conversational recommender systems," *ACM Computing Surveys (CSUR)*, 2021.
- [4] Z. He, Z. Xie, R. Jha, H. Steck, D. Liang, Y. Feng, B. P. Majumder, N. Kallus, and J. McAuley, "Large language models as zero-shot conversational recommenders," in *Proceedings of the 32nd ACM international conference on information and knowledge management*, 2023.
- [5] A. T. Chaganty, M. Leszczynski, S. Zhang, R. Ganti, K. Balog, and F. Radlinski, "Beyond single items: Exploring user preferences in item sets with the conversational playlist curation dataset," in *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2023.
- [6] M. Leszczynski, S. Zhang, R. Ganti, K. Balog, F. Radlinski, F. Pereira, and A. T. Chaganty, "Talk the walk: Synthetic data generation for conversational music recommendation," *arXiv preprint arXiv:2301.11489*, 2023.
- [7] Z. Dai, A. T. Chaganty, V. Y. Zhao, A. Amini, Q. M. Rashid, M. Green, and K. Guu, "Dialog inpainting: Turning documents into dialogs," in *International Conference on Machine Learning (ICML)*, 2022.
- [8] B. G. Glaser, A. L. Strauss, and E. Strutzel, "The discovery of grounded theory; strategies for qualitative research," *Nursing research*, 1968.
- [9] I. Manco, E. Benetos, E. Quinton, and G. Fazekas, "Contrastive audio-language learning for music," *International Society for Music Information Retrieval (ISMIR)*, 2022.
- [10] S. Doh, M. Won, K. Choi, and J. Nam, "Toward universal text-to-music retrieval," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [11] S. Doh, M. Lee, D. Jeong, and J. Nam, "Enriching music descriptions with a finetuned-llm and metadata for text-to-music retrieval," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024.
- [12] S. Liu, A. S. Hussain, C. Sun, and Y. Shan, "Music understanding llama: Advancing text-to-music generation with question answering and captioning," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024.
- [13] J. Gardner, S. Durand, D. Stoller, and R. M. Bittner, "Llark: A multimodal foundation model for music," *arXiv preprint arXiv:2310.07160*, 2023.
- [14] S. Doh, K. Choi, J. Lee, and J. Nam, "LP-musiccaps: LLM-based pseudo music captioning," *International Society for Music Information Retrieval (ISMIR)*, 2023.
- [15] J. S. Downie and S. J. Cunningham, "Toward a theory of music information retrieval queries: System design implications," 2002.
- [16] D. Bainbridge, S. J. Cunningham, and J. S. Downie, "How people describe their music information needs: A grounded theory analysis of music queries," 2003.
- [17] J. H. Lee, "Analysis of user needs and information features in natural language queries seeking music information," *Journal of the American Society for Information Science and Technology*, 2010.
- [18] W. Cai and L. Chen, "Predicting user intents and satisfaction with dialogue-based conversational recommendations," in *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, 2020.
- [19] E. Epure and R. Hennequin, "A human subject study of named entity recognition in conversational music recommendation queries," in *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, 2023.
- [20] K. Krippendorff, "Testing the reliability of content analysis data," *The content analysis reader*, 2009.
- [21] E. Dinan, S. Roller, K. Shuster, A. Fan, M. Auli, and J. Weston, "Wizard of wikipedia: Knowledge-powered conversational agents," *arXiv preprint arXiv:1811.01241*, 2018.
- [22] J. Johnson, B. Hariharan, L. Van Der Maaten, L. Fei-Fei, C. Lawrence Zitnick, and R. Girshick, "Clevr: A diagnostic dataset for compositional language and elementary visual reasoning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017.
- [23] T. Bertin-Mahieux, D. P. Ellis, B. Whitman, and P. Lamere, "The million song dataset," in *International Society for Music Information Retrieval (ISMIR)*, 2011.
- [24] H. Schreiber, "Improving genre annotations for the million song dataset," in *International Society for Music Information Retrieval (ISMIR)*, 2015.

- [25] M. Won, S. Oramas, O. Nieto, F. Gouyon, and X. Serra, "Multimodal metric learning for tag-based music retrieval," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [26] M. Won, K. Choi, and X. Serra, "Semi-supervised music tagging transformer," in *International Society for Music Information Retrieval (ISMIR)*, 2021.
- [27] A. Schindler and P. Knees, "Multi-task music representation learning from multi-label embeddings," in *Proc. International Conference on Content-Based Multimedia Indexing (CBMI)*, 2019.
- [28] B. McFee and G. R. Lanckriet, "Hypergraph models of playlist dialects," in *International Society for Music Information Retrieval (ISMIR)*, 2012.
- [29] Y. Hu, Y. Koren, and C. Volinsky, "Collaborative filtering for implicit feedback datasets," in *2008 Eighth IEEE international conference on data mining*, 2008.
- [30] F. Korzeniowski, S. Oramas, and F. Gouyon, "Artist similarity with graph neural networks," in *International Society for Music Information Retrieval (ISMIR)*, 2021.
- [31] S. Böck, M. E. Davies, and P. Knees, "Multi-task learning of tempo and beat: Learning one to improve the other," in *International Society for Music Information Retrieval (ISMIR)*, 2019.
- [32] F. Korzeniowski and G. Widmer, "Genre-agnostic key classification with convolutional neural networks," *arXiv preprint arXiv:1808.05340*, 2018.
- [33] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv:1807.03748*, 2018.
- [34] S. Robertson, S. Walker, and S. Jones, "M. hancock-beaulieu, m., and gatford, m.(1995). okapi at trec-3," in *The Third Text REtrieval Conference (TREC-3)*, 1994.
- [35] G. Izacard, M. Caron, L. Hosseini, S. Riedel, P. Bojanowski, A. Joulin, and E. Grave, "Unsupervised dense information retrieval with contrastive learning," *arXiv preprint arXiv:2112.09118*, 2021.
- [36] K. Choi, G. Fazekas, K. Cho, and M. Sandler, "The effects of noisy labels on deep convolutional neural networks for music tagging," *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2018.