
The Song Describer Dataset: a Corpus of Audio Captions for Music-and-Language Evaluation

Ilaria Manco^{*1,2}, Benno Weck^{*3}, Seungheon Doh⁴, Minz Won⁵,
Yixiao Zhang¹, Dmitry Bodganov³, Yusong Wu⁶, Ke Chen⁷, Philip Tovstogan³
Emmanouil Benetos¹, Elio Quinton², George Fazekas¹, Juhan Nam⁴
¹QMUL, ²UMG, ³UPF, ⁴KAIST, ⁵ByteDance, ⁶Mila, ⁷UCSD

Abstract

We introduce the Song Describer dataset (SDD), a new crowdsourced corpus of high-quality audio-caption pairs, designed for the evaluation of music-and-language models. The dataset consists of 1.1k human-written natural language descriptions of 706 music recordings, all publicly accessible and released under Creative Commons licenses. To showcase the use of our dataset, we benchmark popular models on three key music-and-language tasks (music captioning, text-to-music generation and music-language retrieval). Our experiments highlight the importance of cross-dataset evaluation and offer insights into how researchers can use SDD to gain a broader understanding of model performance.

1 Introduction

Multimodal approaches that jointly process audio and language are becoming increasingly important within music understanding and generation, giving rise to a new area of research, which we refer to as music-and-language (M&L). Several recent works have emerged in this domain, proposing methods to automatically generate music descriptions [21, 7, 9], synthesise music from a text prompt [2, 14, 30, 6], search for music based on language queries [8, 22, 13], and more [20, 17]. However, evaluating M&L models remains a challenge due to a lack of public and accessible datasets with paired audio and language, resulting in the widespread use of private data [21, 22, 23, 14, 2, 13] and inconsistent evaluation practices. To mitigate this, we release the Song Describer dataset (SDD), a new high-quality evaluation dataset of crowdsourced captions paired with openly licensed music recordings. Through our dataset, we allow for the first standardised comparison of models on realistic music-language data and expand the number and variety of datasets available to the research community. We propose SDD as an evaluation-only dataset to promote the development and use of out-of-domain data and counteract the tendency to overfit to a particular dataset [5, 29, 32].

Alongside SDD, there are currently two other human-curated datasets of music-caption pairs: MusicCaps [2] and YT8M-MusicTextClips [25], composed of 5.5k and 4k 10-second clips from AudioSet [11] and YouTube8M [1] respectively. More recently, Doh et al. [7] have released LP-MusicCaps, a dataset of 2.2M music descriptions generated by prompting a large language model with a set of instructions and ground-truth tags. While synthetic data presents an opportunity to artificially scale up training data, limitations such as LLM hallucinations still pose significant challenges to their use in evaluation, where reliable data plays a crucial role.

*Equal contribution

Table 1: An overview of the SDD compared to other music-caption datasets, MusicCaps (MC) [2] and YT8M-MusicTextClips (MTC) [25]. * denotes the validated subset.

Dataset	Annotators	Text			Audio			
		Caption #	Avg length	Vocab size	Audio #	Length	Source	Public
MC	10	5,521	54.8 ± 18.7	6,144	5,521	10 sec	[11]	✗
MTC	-	4,169	16.9 ± 4.4	2,599	4,169	10 sec	[1]	✗
SDD*	114	746	18.2 ± 7.6	1,942	547	2 min	[4]	✓
SDD	142	1106	21.7 ± 12.4	2,859	706	2 min	[4]	✓

Table 2: The main task instruction given to the annotators, together with an example of three captions produced by different annotators describing the same track in the dataset.

Instructions	Example caption
Write one sentence that describes the track you just listened to: focus on descriptive text which conveys the overall content of this track only.	<i>A folky song with a warm, organic and cosy sound, sung by a single male voice and accompanied by a number of acoustic guitars.</i>
	<i>This is a chill folk song with a simple rhythm and acoustic guitar, good for a slow morning of reading or working.</i>
	<i>A natural, singer-songwriter esque song driven by croaky male vocals, acoustic guitars, and found drum sounds; later a shaker emerges as the singer uplifts the melody with his voice.</i>

2 The Song Describer Dataset

SDD is a corpus of manually created descriptions for high-quality audio recordings from the MTG-Jamendo dataset [4], all available under Creative Commons licenses, and distributed as part of the dataset alongside the annotations, which we release under the CC BY-SA 4.0 license.² Captions in SDD contain syntactically and semantically rich single-sentence descriptions covering various musical features such as instrumentation, genre, style and mood, as shown in the examples in Table 2. Notably, recordings in SDD can also be connected to additional metadata from MTG-Jamendo, enriching this with tags of several musical categories (see Figure 3 in Appendix A.2).

Compared to MusicCaps and YT8M-MusicTextClips, SDD is characterised by (1) annotations of longer track segments (95% of the tracks are 2 minutes long); (2) openly accessible audio, persistently stored, and with explicit licenses; (3) a larger pool of annotators, composed of non-experts with diverse backgrounds and more representative of the general public (this is only known for MusicCaps); (4) a portion of recordings (25%) with more than one caption produced by different annotators, making it particularly suitable for evaluation via automatic metrics for captioning. In contrast, audio clips are not distributed as part of MusicCaps and YT8M-MusicTextClips and can only be indirectly obtained via YouTube, with no guarantees about data persistence. Moreover, audio in MusicCaps is derived from AudioSet, which is often used for pre-training large audio models, thus posing a risk of data leakage and in-distribution bias [33]. We summarise these and additional key differences in Table 1.

Data collection The captions making up SDD were collected through an open-source online annotation platform [24] aimed at non-experts and volunteers. Participants were invited via a public announcement circulated on mailing lists and online communities for researchers focused on audio, music, and machine learning, as well as university students more broadly. Most participants were from English-speaking countries (US, GB, IN) and Western Europe (ES, FR, DE, IT), as shown in Figure 1. Online participation required participants to be at least 18 years of age, with no additional requirements. We surveyed the participants for their background information for post-hoc quality control and analysis, by collecting answers to questions about their English writing ability (3-point Likert scale), age range, and country of residence. To evaluate the familiarity and music knowledge of the participants, we used selected questions from the Goldsmiths Musical Sophistication Index (Gold-MSI) [26] with 7-point Likert scale answers to the following statements: (1) *I spend a lot of my free time doing music-related activities [doing]*; (2) *I enjoy writing about music, for example on blogs*

²<https://doi.org/10.5281/zenodo.10072001>

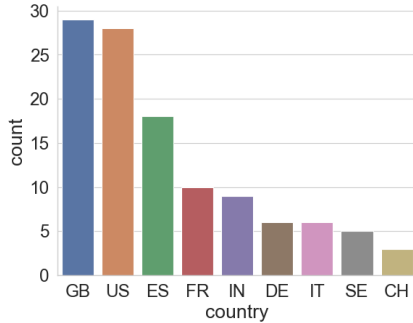


Figure 1: Distribution of the top nine countries of residence contributors reported in the on-boarding survey.




Figure 2: Distribution of music aspects covered by the most frequent word stems in the collected captions.

or social media [writing]; (3) I often read or search the internet for things related to music [reading]. The analysis of this data together with the code to reproduce the figures is available online.³

In the annotation process, annotators were asked to listen to a music recording (up to 2 minutes), write one sentence describing it, and indicate their familiarity with the music on a 3-point Likert scale. Music recordings were sampled at random from the test split *split-0* of MTG-Jamendo while giving preference to popular tracks as measured by the play count on the Jamendo platform. A recording could be captioned by up to five participants. Through this, we gathered 1106 captions from 142 volunteers. The main task instruction, together with an example of multiple captions of the same track provided by different annotators is shown in Table 2. Finally, additional details on the data collection process are given in the datasheet [10] in Appendix A.

Data quality validation All submitted captions were manually reviewed by the dataset creators to ensure consistency with the annotation guidelines and overall quality, leading to a smaller, curated subset. Reviewers were instructed to rate whether a submitted caption was valid, invalid or borderline, based on whether it adhered to the annotation guidelines provided: *Write one full sentence; use at least 8 words; include things like: genre, mood, instrumentation, ideal listening situation, emotions it evokes, etc.; don't simply provide a sequence of tags; don't use swear words or other offensive language; don't refer to other songs or artists you may know; don't mention personal events or experiences.* While we observed that, in most cases, captions were high-quality and followed the guidelines, a significant number of cases did not satisfy the single-sentence requirement. Given their high quality,⁴ these cases were marked as borderline. While all annotations are included in the full dataset, captions which fulfilled all the requirements are marked as part of the valid subset, to enable more reliable evaluation. The validated subset contains rich captions covering a wide range of music aspects. In an exploratory analysis, we labelled word stems appearing more than 10 times in the data (after removing stop words) with an aspect label. An overview of this analysis is presented in Figure 2, where we show the distribution of the most common aspects. These include *Instrumentation* (such as piano, guitar, vocals), *Structure* (e.g. song, accompaniment, chorus), *Genre/style* (e.g. rock, pop, ambient), alongside others such as *Acoustic qualities*, *Era*, *Language*, *Lyrical content*, *Rhythm*, *Tempo*, and *Usage* (denoted by the label *Other* in the plot).

3 Benchmarking Music-and-Language Models

In this section we demonstrate the use of SDD in benchmarking existing music-and-language models. We compare results to those obtained on MusicCaps, currently the most popular dataset in this domain, and highlight the insights that can be drawn from this cross-dataset evaluation.

³<https://github.com/mulab-mir/song-describer-dataset>

⁴An example of such a caption is: “A nice backbeat punctuates this upbeat, danceable track. The vocals make this sound like taking a walk on a summer afternoon, enjoying the scenery and saying hi to the neighbors. Synthesizers modernize the feel and give us a break from the vocals for a bit. A breakdown follows with piano.”. An example of an invalid caption is instead: “mood: uplifting, inspiring. vocals: male singer. instruments: guitar, electronic synthesizers”

Table 3: Baseline results on the **music captioning** task.

Model	MusicCaps						Song Descriptor					
	B ₁	B ₂	B ₃	M	R	B-S	B ₁	B ₂	B ₃	M	R	B-S
MusCaps	22.0	10.3	5.3	17.0	22.2	83.5	9.6	3.4	1.1	11.9	12.1	83.8
LP-MusicCaps	28.6	13.9	7.8	20.7	19.3	87.1	12.7	4.0	1.3	16.7	11.9	86.0

Table 4: Results of the two baselines we evaluate on the **text-to-music generation** task.

Model	FAD ↓	IS	KLD ↓	CCC
MUSICCAPS				
Riffusion	11.28	1.25	1.66	0.19
AudioLDM	31.55	1.02	2.80	0.21
SONG DESCRIPTER				
Riffusion	3.87	1.31	0.94	0.17
AudioLDM	20.74	1.01	2.28	0.15

Table 5: Results of the two baselines we evaluate on the **text-to-music retrieval** task.

Model	R@1	R@5	R@10	MedR ↓
MUSICCAPS				
TTMR	0.21	1.20	2.64	638
CLAP	3.81	10.95	17.18	82.5
SONG DESCRIPTER				
TTMR	2.95	10.19	17.43	51
CLAP	4.42	17.02	26.01	36

Tasks and baselines *Music captioning* is the task of generating natural language descriptions of a music audio recording and is typically evaluated through a set of automatic metrics such as BLEU (B) [27], METEOR (M) [16], ROUGE (R) [18] and, more recently, BERT-score (B-S) [36]. As our baselines for this task, we evaluate two encoder-decoder models, MusCaps [21] and LP-MusicCaps [7]. MusCaps combines a CNN as audio encoder and an LSTM as text decoder, whereas LP-MusicCaps utilizes a combination of a CNN and transformer as an audio encoder, and a transformer as a text decoder. We train both from scratch on the MusicCaps training split, using beam search decoding with a beam size of 4.

Text-to-music generation is the task of creating music audio conditioned on a text prompt. Performance on this task can be evaluated along two dimensions: (1) the quality and diversity of the generated audio, using metrics such as the Fréchet Audio Distance (FAD) [15], Inception Score (IS) [12] and Kullback–Leibler Divergence (KLD), and (2) the relevance between text and audio. Estimating the latter is more elusive and typically addressed via listening studies. Taking inspiration from the MCC metric proposed in [2], we define a variant of it based on the CLAP model [35], which we call CLAP Cycle Consistency (CCC). For our evaluation on this task, among the recent but growing pool of text-to-music generative models, we select Riffusion [31] and AudioLDM [19] based on the availability of pre-trained weights and diversity in their design.

In *text-to-music retrieval*, the goal is to use text inputs as queries to retrieve corresponding music items. The predominant approach to this is to use multimodal embedding models trained to map different modalities onto a shared space via contrastive learning [28, 34, 22]. Accordingly, for this task, we select two pre-trained state-of-the-art audio-text embedding models: TTMR [8], trained on the Million Song Dataset [3] and CLAP [35], trained on Audioset [11] and LAION-Audio-630k [35]. We evaluate their performance by computing standard information retrieval metrics on the top k retrieved results: Recall at k ($R@k$) and Median Rank (MedR).

Results For music captioning, in Table 3 we observe that, beyond inter-model differences ascribable to the different model architectures and pre-training, both systems display a sizeable gap between in-domain (MC) and out-of-domain (SDD) performance. Although not in itself surprising, this drop in performance underscores the importance of out-of-domain evaluation sets in providing a more realistic assessment of performance in the wild. From the text-to-music generation results in Table 4, we observe consistent trends between the two datasets on the audio generation metrics: Riffusion tends to produce higher-quality outputs, which may be explained by the fact that AudioLDM is not exclusively trained on music. However, the trend is inverted for semantic relevance. This implies that care should be taken in determining performance superiority based on a single dataset. Finally, we present the music-text retrieval results in Table 5. We observe that CLAP consistently outperforms TTMR, and retrieval scores are higher on SSD compared to MusicCaps, which is expected due to its smaller size ($\sim 4x$). Notably, this trend is less prominent for CLAP, and, when adjusting for dataset size, the difference between the two models becomes less stark. Since CLAP is partly trained on AudioSet, this also suggests that evaluating on in-domain distributions may lead to inflated results.

4 Conclusion

We have presented the Song Describer dataset, a corpus of audio-caption pairs for music-and-language evaluation, and shown how this can be used to perform a more complete assessment of model performance through cross-dataset evaluation. Given the small scale of the dataset, further data collection would be beneficial in strengthening its reliability. Additionally, this may be combined with an extension of the musical cultures represented in both the set of audio recordings and the pool of annotators, which remain limited in the current version.

References

- [1] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan. YouTube-8M: A Large-Scale Video Classification Benchmark, Sept. 2016. arXiv:1609.08675 [cs].
- [2] A. Agostinelli, T. I. Denk, Z. Borsos, J. Engel, M. Verzett, A. Caillon, Q. Huang, A. Jansen, A. Roberts, M. Tagliasacchi, M. Sharifi, N. Zeghidour, and C. Frank. MusicLM: Generating Music From Text, Jan. 2023. arXiv:2301.11325 [cs, eess].
- [3] T. Bertin-Mahieux, D. P. Ellis, B. Whitman, and P. Lamere. The million song dataset. In *Proceedings of the 12th ISMIR Conference*, 2011.
- [4] D. Bogdanov, M. Won, P. Tovstogan, A. Porter, and X. Serra. The MTG-Jamendo Dataset for Automatic Music Tagging. In *Machine Learning for Music Discovery Workshop, International Conference on Machine Learning (ICML)*, 2019.
- [5] S. R. Bowman and G. E. Dahl. What Will it Take to Fix Benchmarking in Natural Language Understanding?, Oct. 2021. arXiv:2104.02145 [cs].
- [6] J. Copet, F. Kreuk, I. Gat, T. Remez, D. Kant, G. Synnaeve, Y. Adi, and A. Défossez. Simple and Controllable Music Generation, June 2023. arXiv:2306.05284 [cs, eess].
- [7] S. Doh, K. Choi, J. Lee, and J. Nam. LP-MusicCaps: LLM-Based Pseudo Music Captioning. In *Proceedings of the 24th International Society for Music Information Retrieval (ISMIR) Conference*, Milan, July 2023. arXiv:2307.16372 [cs, eess].
- [8] 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)*. IEEE, 2022. arXiv:2211.14558 [cs, eess].
- [9] G. Gabbolini, R. Hennequin, and E. Epure. Data-Efficient Playlist Captioning With Musical and Linguistic Knowledge. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11401–11415, Abu Dhabi, United Arab Emirates, Dec. 2022. Association for Computational Linguistics.
- [10] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. D. III, and K. Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, Nov. 2021.
- [11] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter. Audio Set: An ontology and human-labeled dataset for audio events. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pages 776–780. Institute of Electrical and Electronics Engineers Inc., June 2017. ISSN: 15206149.
- [12] S. Gurumurthy, R. K. Sarvadevabhatla, and R. V. Babu. DeLiGAN: Generative Adversarial Networks for Diverse and Limited Data. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4941–4949, Honolulu, HI, July 2017. IEEE.
- [13] Q. Huang, A. Jansen, J. Lee, R. Ganti, J. Y. Li, and D. P. W. Ellis. MuLan: A Joint Embedding of Music Audio and Natural Language. In *23rd International Society for Music Information Retrieval Conference (ISMIR 2022)*, 2022.
- [14] Q. Huang, D. S. Park, T. Wang, T. I. Denk, A. Ly, N. Chen, Z. Zhang, Z. Zhang, J. Yu, C. Frank, J. Engel, Q. V. Le, W. Chan, Z. Chen, and W. Han. Noise2Music: Text-conditioned Music Generation with Diffusion Models, Mar. 2023. arXiv:2302.03917 [cs, eess].
- [15] K. Kilgour, M. Zuluaga, D. Roblek, and M. Sharifi. Fréchet Audio Distance: A Reference-Free Metric for Evaluating Music Enhancement Algorithms. In *Interspeech 2019*, pages 2350–2354. ISCA, Sept. 2019.
- [16] A. Lavie and A. Agarwal. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. *Proceedings of the Second Workshop on Statistical Machine Translation*, 2007.
- [17] M. E. Leszczynski, R. Ganti, S. Zhang, K. Balog, F. Radlinski, F. Pereira, and A. T. Chaganty. Conversational Music Retrieval with Synthetic Data. In *Second Workshop on Interactive Learning for Natural Language Processing at NeurIPS 2022*, 2022.

- [18] C. Y. Lin. Rouge: A package for automatic evaluation of summaries. *Proceedings of the workshop on text summarization branches out (WAS 2004)*, 2004.
- [19] H. Liu, Z. Chen, Y. Yuan, X. Mei, X. Liu, D. Mandic, W. Wang, and M. D. Plumbley. AudioLDM: Text-to-Audio Generation with Latent Diffusion Models, Feb. 2023. arXiv:2301.12503 [cs, eess].
- [20] S. Liu, A. S. Hussain, C. Sun, and Y. Shan. Music Understanding LLaMA: Advancing Text-to-Music Generation with Question Answering and Captioning, Aug. 2023. arXiv:2308.11276 [cs, eess].
- [21] I. Manco, E. Benetos, E. Quinton, and G. Fazekas. MusCaps: Generating Captions for Music Audio. In *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2021. arXiv: 2104.11984.
- [22] I. Manco, E. Benetos, E. Quinton, and G. Fazekas. Contrastive Audio-Language Learning for Music. In *23rd International Society for Music Information Retrieval Conference (ISMIR 2022)*, 2022. arXiv: 2208.12208.
- [23] I. Manco, E. Benetos, E. Quinton, and G. Fazekas. Learning music audio representations via weak language supervision. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022. arXiv: 2112.04214.
- [24] I. Manco, B. Weck, P. Tovstogan, M. Won, and D. Bogdanov. Song Descriptor: a Platform for Collecting Textual Descriptions of Music Recordings. In *Proceedings of the 23rd International Society for Music Information Retrieval Conference*, Dec. 2022. Conference Name: Ismir 2022 Hybrid Conference.
- [25] D. McKee, J. Salamon, J. Sivic, and B. Russell. Language-Guided Music Recommendation for Video via Prompt Analogies, June 2023. arXiv:2306.09327 [cs].
- [26] D. Müllensiefen, B. Gingras, J. Musil, and L. Stewart. The Musicality of Non-Musicians: An Index for Assessing Musical Sophistication in the General Population. *PLoS one*, 9(2), Feb. 2014. Publisher: Public Library of Science.
- [27] K. Papineni, S. Roukos, T. Ward, and W.-j. Zhu. BLEU : a Method for Automatic Evaluation of Machine Translation. *Computational Linguistics*, 2002.
- [28] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning*. PMLR, 2021.
- [29] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar. Do ImageNet Classifiers Generalize to ImageNet? In *Proceedings of the 36th International Conference on Machine Learning*, pages 5389–5400. PMLR, May 2019. ISSN: 2640-3498.
- [30] F. Schneider, Z. Jin, and B. Schölkopf. Mousai: Text-to-Music Generation with Long-Context Latent Diffusion, Jan. 2023. arXiv:2301.11757 [cs, eess].
- [31] H. M. Seth Forsgren. Riffusion - Stable diffusion for real-time music generation, 2022.
- [32] V. Vasudevan, B. Caine, R. Gontijo-Lopes, S. Fridovich-Keil, and R. Roelofs. When does dough become a bagel? Analyzing the remaining mistakes on ImageNet. In *ICML 2022 Shift Happens Workshop*, May 2023.
- [33] B. Weck and X. Serra. Data Leakage in Cross-Modal Retrieval Training: A Case Study. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, June 2023.
- [34] H.-H. Wu, O. Nieto, J. P. Bello, and J. Salamon. Audio-Text Models Do Not Yet Leverage Natural Language, Mar. 2023. arXiv:2303.10667 [cs, eess].
- [35] Y. Wu, K. Chen, T. Zhang, Y. Hui, T. Berg-Kirkpatrick, and S. Dubnov. Large-scale Contrastive Language-Audio Pretraining with Feature Fusion and Keyword-to-Caption Augmentation, Apr. 2023. arXiv:2211.06687 [cs, eess].
- [36] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations 2020*, 2020.

A Datasheet for the Song Describer Dataset (SDD)

A.1 Motivation

- **For what purpose was the dataset created?** *Was there a specific task in mind? Was there a specific gap that needed to be filled? Please provide a description.*

The dataset was created in order to provide the research community with a public dataset for music-and-language tasks containing audio with permissive licenses. The primary use case for the dataset is the evaluation of machine learning models on tasks such as music captioning, text-to-music generation and music-text retrieval.

- **Who created the dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?**

The dataset is a result of a research collaboration undertaken by Universal Music Group International Limited (part of Universal Music Group), Queen Mary University of London, and Music Technology Group (Universitat Pompeu Fabra).

- **Who funded the creation of the dataset?** *If there is an associated grant, please provide the name of the grantor and the grant name and number.*

The dataset creation was funded by:

- UK Research and Innovation [grant number EP/S022694/1]
- Universal Music Group
- The Musical AI project - PID2019-111403GB-I00/AEI/10.13039/501100011033, funded by the Spanish Ministerio de Ciencia e Innovación and the Agencia Estatal de Investigación.

- **Any other comments?**

No.

A.2 Composition

- **What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)?** *Are there multiple types of instances (e.g., movies, users, and ratings; people and interactions between them; nodes and edges)? Please provide a description.*

The dataset consists of music audio recordings (each recording of a different music track) associated with music captions (up to 5 captions per recording, written by different participants in the annotation platform).

- **How many instances are there in total (of each type, if appropriate)?**

There are 1106 captions of 706 recordings in total, written by 142 annotators. We also provide a cleaned subset containing manually validated annotations, consisting of 746 captions of 547 audio recordings, written by 114 annotators.

- **Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?** *If the dataset is a sample, then what is the larger set? Is the sample representative of the larger set (e.g., geographic coverage)? If so, please describe how this representativeness was validated/verified. If it is not representative of the larger set, please describe why not (e.g., to cover a more diverse range of instances, because instances were withheld or unavailable).*

The audio associated with music captions is a subset of the larger MTG-Jamendo dataset⁵ containing over 55k tracks. Due to the annotation effort, we sampled a random subset of tracks from the test set of the *split-0*. The representativeness of our random subset was validated by comparing the distribution of music tags in our subset to that of the full dataset (see Fig. 3).

- **What data does each instance consist of?** *“Raw” data (e.g., unprocessed text or images) or features? In either case, please provide a description.*

Each instance consists of an audio-text pair:

⁵<https://mtg.github.io/mtg-jamendo-dataset/>




Figure 3: Distribution of the top-10 tags per category taken from MTG-Jamendo in the data collected for SDD.

- Audio: one music audio segment per track, with segment duration up to 2 minutes, trimmed from the beginning of a music track (95% of the segments are 2 minutes). Each audio file is provided in 320kbps 44.1 kHz MP3 audio encoding.
 - Text: unprocessed music caption provided as a list of strings of text, associated with the audio segments.
- **Is there a label or target associated with each instance?** *If so, please provide a description.*
There is no ground-truth label. Depending on the task, either the audio or the text in a given instance constitutes the target (e.g. in music captioning, the target is the caption).
 - **Is any information missing from individual instances?** *If so, please provide a description, explaining why this information is missing (e.g., because it was unavailable). This does not include intentionally removed information, but might include, e.g., redacted text.*
There is no missing information.
 - **Are relationships between individual instances made explicit (e.g., users' movie ratings, social network links)?** *If so, please describe how these relationships are made explicit.*
Each instance is associated to a track ID which can be linked to the original MTG-Jamendo dataset. Relationships between instances can be easily explored from this through the MTG-Jamendo metadata, which contains additional information such as artist, album, genre etc.
 - **Are there recommended data splits (e.g., training, development/validation, testing)?** *If so, please provide a description of these splits, explaining the rationale behind them.*
There is no recommended split, as the dataset is intended to be used solely for evaluation.
 - **Are there any errors, sources of noise, or redundancies in the dataset?** *If so, please provide a description.*
The annotations may be prone to noise and bias due to the subjectivity of the task, different levels of familiarity with the kind of music being annotated by the participants, and differences in their English language level. The dataset creators manually reviewed all captions to ensure consistency with the annotation guidelines and overall quality, leading to a smaller, curated subset.
 - **Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)?** *If it links to or relies on external resources, a) are there guarantees that they will exist, and remain constant, over time; b) are there official archival versions of the complete dataset (i.e., including the external resources as they existed at the time the dataset was created); c) are there any restrictions (e.g., licenses, fees) associated with any of the external resources that might apply to a future user? Please provide descriptions of all external resources and any restrictions associated with them, as well as links or other access points, as appropriate.*
The dataset is self-contained (all audio and metadata are distributed as part of the dataset package, including CC licenses for the individual tracks). No fees or restrictions apply

to any of the material making up the dataset. The dataset can be additionally mapped to the metadata in the MTG-Jamendo dataset via track IDs. Additionally, the annotation platform used in creating the dataset is open-source and publicly accessible at <https://github.com/ilaria-manco/song-describer>.

- **Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor-patient confidentiality, data that includes the content of individuals' non-public communications)?** *If so, please provide a description.*
There is no confidential information in the dataset.
- **Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety?** *If so, please describe why.*
Individual songs might contain lyrics that could be considered offensive by listeners. Annotators were appropriately informed about this risk.
- **Does the dataset relate to people?** *If not, you may skip the remaining questions in this section.*
The dataset includes human annotations of music recordings and therefore relates to people. However, the participation in the annotation platform was anonymous and no personal data or personally identifiable information was collected.
- **Does the dataset identify any subpopulations (e.g., by age, gender)?**
No.
- **Is it possible to identify individuals (i.e., one or more natural persons), either directly or indirectly (i.e., in combination with other data) from the dataset?** *If so, please describe how.*
No.
- **Does the dataset contain data that might be considered sensitive in any way (e.g., data that reveals racial or ethnic origins, sexual orientations, religious beliefs, political opinions or union memberships, or locations; financial or health data; biometric or genetic data; forms of government identification, such as social security numbers; criminal history)?** *If so, please provide a description.*
No.
- **Any other comments?**
No.

A.3 Collection Process

- **How was the data associated with each instance acquired?** *Was the data directly observable (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indirectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for age or language)? If data was reported by subjects or indirectly inferred/derived from other data, was the data validated/verified? If so, please describe how.*
The data was reported by subjects in the form of free-form text annotations and manually validated by the authors following collection.
- **What mechanisms or procedures were used to collect the data (e.g., hardware apparatus or sensor, manual human curation, software program, software API)?** *How were these mechanisms or procedures validated?*
The data was collected via a self-developed online software platform.⁶ This was validated via two pilot studies prior to starting the data collection.
- **If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?**
The recordings were randomly sampled from the test set of split-0 of the MTG-Jamendo with weights proportional to play counts from the Jamendo platform.⁷

⁶<https://song-describer.streamlit.app/>

⁷<https://developer.jamendo.com/v3.0>

- **Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?**

Volunteers (above 18 years of age) were recruited via an open call circulated among relevant mailing lists and social networks. We hosted several rounds of annotation with a leaderboard showing (nicknames of) the participants and their number of annotations. The top-3 contributors for each round were rewarded with prizes. These consisted of online vouchers of a value ranging from £10 to £100.

- **Over what timeframe was the data collected? Does this timeframe match the creation timeframe of the data associated with the instances (e.g., recent crawl of old news articles)?** *If not, please describe the timeframe in which the data associated with the instances was created.*

The annotations were collected between 25th November 2022 and 14th April 2023. The instances are associated with music recordings created between 2004 and 2017.

- **Were any ethical review processes conducted (e.g., by an institutional review board)?** *If so, please provide a description of these review processes, including the outcomes, as well as a link or other access point to any supporting documentation.*

The project was approved by the Queen Mary Ethics of Research Committee (QMERC reference number: QMERC22.062). As part of the review, the following documentation was provided: a participant information sheet, a consent form, and an application form with detailed questions about the data collection procedure and potential risks. These were assessed by the committee through the low-risk approval route, according to the guidelines outlined on the website.⁸

- **Does the dataset relate to people?** *If not, you may skip the remaining questions in this section.*

Yes.

- **Did you collect the data from the individuals in question directly, or obtain it via third parties or other sources (e.g., websites)?**

We collected the data via the online annotation platform we developed for this purpose (see questions above).

- **Were the individuals in question notified about the data collection?** *If so, please describe (or show with screenshots or other information) how notice was provided, and provide a link or other access point to, or otherwise reproduce, the exact language of the notification itself.*

Yes, the platform was built for data collection, thus it includes explicit notice and instructions. The landing page with most of the information can be accessed at <https://song-describer.streamlit.app>. A more detailed description of the annotation platform can be found in [24]. Below is the text that informs user about data collection and terms of use.

Song Describer is a tool for crowdsourced collection of music captions, built by a team of researchers from Queen Mary University of London and Universitat Pompeu Fabra Barcelona.

You will be asked to listen to some music and write a new description for it (Annotation) or evaluate an existing one (Evaluation). Each task takes only a few minutes. You can complete either task as many times as you'd like.

By contributing to Song Describer, you are helping us release the first public dataset of music captions. We really appreciate your help!

We will only collect data you provide by submitting your answers and no other personal information. To find out more about the project and how the data will be used, read our FAQs below.

This survey is part of a research project being undertaken by Universal Music Group International Limited (part of Universal Music Group), Queen Mary University of London, and Universitat Pompeu Fabra (together the "Research Team").

By participating in this survey, you acknowledge and agree to the following:

⁸<http://www.jrmo.org.uk/performing-research/conducting-research-with-human-participants-outside-the-nhs/applications-and-approval/>

- You must be aged 18 or over
 - You are participating in an academic study, the results of which may or may not be published in an academic journal.
 - Your participation is voluntary and you are free to leave at any time
 - All intellectual property rights which may arise or inure to you as a result of your participation in this survey are hereby assigned jointly, in full and in equal proportion to the members of the Research Team.
 - By participating in this study, you agree to waive any moral rights of authorship that you may have in the responses that you provide in the survey to the extent permitted by law.
 - The data collected by this survey is intended to be published and shall be freely available to all. The responses submitted by you shall not be attributable to you and your participation in the survey shall remain confidential.
- **Did the individuals in question consent to the collection and use of their data?** *If so, please describe (or show with screenshots or other information) how consent was requested and provided, and provide a link or other access point to, or otherwise reproduce, the exact language to which the individuals consented.*
See the link and text in the previous question. At the bottom of the webpage is a button with "Get Started" text that shows a pop-up on hover: "By clicking here, you confirm that you agree with the statements above."
 - **If consent was obtained, were the consenting individuals provided with a mechanism to revoke their consent in the future or for certain uses?** *If so, please provide a description, as well as a link or other access point to the mechanism (if appropriate).*
Participants were made aware of the option to withdraw from the annotation procedure at any point.
 - **Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a data protection impact analysis) been conducted?** *If so, please provide a description of this analysis, including the outcomes, as well as a link or other access point to any supporting documentation.*
No analysis of the potential impact was conducted.
 - **Any other comments?**
No.

A.4 Preprocessing, Cleaning, and/or Labeling

- **Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)?** *If so, please provide a description. If not, you may skip the remainder of the questions in this section.*
As described in Sec. 2, the annotations were manually validated by the research team. This was done by evaluating each caption on a 3-point Likert scale based on its adherence to the annotation guidelines provided. Although we distribute all the annotations (without any preprocessing or cleaning), validated captions are also marked as part of a higher-quality subset.
- **Was there data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)?** *If so, please provide a link or other access point to the raw data.*
All data is provided as part of the dataset release.
- **Is the software used to preprocess/clean/label the instances available?** *If so, please provide a link or other access point.*
We release code to preprocess, clean and analyse the data in an accompanying GitHub repository.⁹
- **Any other comments?**
No.

⁹ <https://github.com/mulab-mir/song-describer-dataset/>

A.5 Uses

- **Has the dataset been used for any tasks already?** *If so, please provide a description.*
The dataset has been used for evaluating machine learning models on the following tasks: music captioning, text-to-music generation, and music-text retrieval.
- **Is there a repository that links to any or all papers or systems that use the dataset?** *If so, please provide a link or other access point.*
Links to papers and systems that use the dataset will be added to the GitHub repository.⁹
- **What (other) tasks could the dataset be used for?**
The dataset may additionally be used to study how people describe music. However we note that due to the small sample size and an unbalanced representation of annotators from different countries, the make-up of the annotator pool must be carefully considered before drawing any conclusions from the analysis of the text in the dataset.
- **Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses?** *For example, is there anything that a future user might need to know to avoid uses that could result in unfair treatment of individuals or groups (e.g., stereotyping, quality of service issues) or other undesirable harms (e.g., financial harms, legal risks) If so, please provide a description. Is there anything a future user could do to mitigate these undesirable harms?*
No.
- **Are there tasks for which the dataset should not be used?** *If so, please provide a description.*
We aim for this dataset to be used for evaluation and benchmarking of music-and-language models, thus we discourage using it for training.
- **Any other comments?**
No.

A.6 Distribution

- **Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created?** *If so, please provide a description.*
The dataset will be publicly available online.
- **How will the dataset be distributed (e.g., tarball on website, API, GitHub)?** *Does the dataset have a digital object identifier (DOI)?*
The dataset will be distributed via Zenodo (DOI: <https://doi.org/10.5281/zenodo.10072001>).
- **When will the dataset be distributed?**
The dataset will be available from 17th November 2023.
- **Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?** *If so, please describe this license and/or ToU, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms or ToU, as well as any fees associated with these restrictions.*
The dataset will be available under the CC BY-SA 4.0 license.¹⁰ The audio files, which are derived from the MTG-Jamendo dataset, are re-distributed as part of the dataset with the respective CC licenses. A list of individual licenses for all audio files is provided with the release.
- **Have any third parties imposed IP-based or other restrictions on the data associated with the instances?** *If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms, as well as any fees associated with these restrictions.*
No.

¹⁰<https://creativecommons.org/licenses/by-sa/4.0/>

- **Do any export controls or other regulatory restrictions apply to the dataset or to individual instances?** *If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any supporting documentation.*
No.
- **Any other comments?**
No.

A.7 Maintenance

- **Who will be supporting/hosting/maintaining the dataset?**
The dataset will be supported and maintained by Queen Mary University of London and the Music Technology Group (Universitat Pompeu Fabra). The dataset will be hosted on Zenodo and supporting code will be hosted on GitHub.
- **How can the owner/curator/manager of the dataset be contacted (e.g., email address)?**
Queries about the dataset can be submitted by opening an issue on GitHub or emailing the dataset curators (i.manco@qmul.ac.uk, benno.weck01@estudiant.upf.edu).
- **Is there an erratum?** *If so, please provide a link or other access point.*
An erratum will be provided in the GitHub repository as necessary.
- **Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)?** *If so, please describe how often, by whom, and how updates will be communicated to users (e.g., mailing list, GitHub)?*
The dataset will be updated to correct errors if necessary. Additionally, it may be updated with new instances if further annotations are provided by volunteers through our data collection platform. If created, future versions of the dataset will be released via Zenodo and updates will be communicated via both GitHub and Zenodo.
- **If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were individuals in question told that their data would be retained for a fixed period of time and then deleted)?** *If so, please describe these limits and explain how they will be enforced.*
No personal data was collected and therefore there are no applicable limits on the retention of the data.
- **Will older versions of the dataset continue to be supported/hosted/maintained?** *If so, please describe how. If not, please describe how its obsolescence will be communicated to users.*
Older versions of the dataset will continue to be hosted and maintained. However any code related to the dataset will be updated to support only the most recent version. This will be explicitly mentioned in the repository.
- **If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so?** *If so, please provide a description. Will these contributions be validated/verified? If so, please describe how. If not, why not? Is there a process for communicating/distributing these contributions to other users? If so, please provide a description.*
The website to collect captions is currently available online¹¹ and can be accessed for additional contributions. Its source code is also public.¹² Contributions are not directly integrated into the dataset, but will be considered for future version releases.
- **Any other comments?**
No.

¹¹<https://song-describer.streamlit.app/>

¹²<https://github.com/ilaria-manco/song-describer>