

# The Interpretation Gap in Text-to-Music Generation Models

Yongyi Zang\*  
Independent Researcher  
zyy0116@gmail.com

Yixiao Zhang\*  
C4DM, Queen Mary University of London  
yixiao.zhang@qmul.ac.uk

## Abstract

Large-scale text-to-music generation models have significantly enhanced music creation capabilities, offering unprecedented creative freedom. However, their ability to collaborate effectively with human musicians remains limited. In this paper, we propose a framework to describe the musical interaction process, which includes expression, interpretation, and execution of controls. Following this framework, we argue that the primary gap between existing text-to-music models and musicians lies in the interpretation stage, where models lack the ability to interpret controls from musicians. We also propose two strategies to address this gap and call on the music information retrieval community to tackle the interpretation challenge to improve human-AI musical collaboration.

## 1 Introduction

In recent years, the field of human-AI music co-creation has experienced significant advancements (Huang et al., 2020; Zhang et al., 2021; Rau et al., 2022; Bougueng Tchemeube et al., 2022). The advent of large-scale text-to-music generation models has played a crucial role in this progress, enabling generating music with good sonic quality and well-defined musical structures (Copet et al., 2024; Evans et al., 2024; Agostinelli et al., 2023).

A primary focus of recent research has been to enhance these models through the incorporation of control signals (Lin et al., 2023; Tal et al., 2024; Wu et al., 2024; Lin et al., 2024; Nistal et al., 2024). This has led to significant success in manipulating dynamics, melody, and chord progressions in generated music contents. While precision in following these control signals can still be improved, these developments represent substantial progress.

Although extensive efforts are made to allow these models follow control signals precisely, misalignments between musicians' intent and model

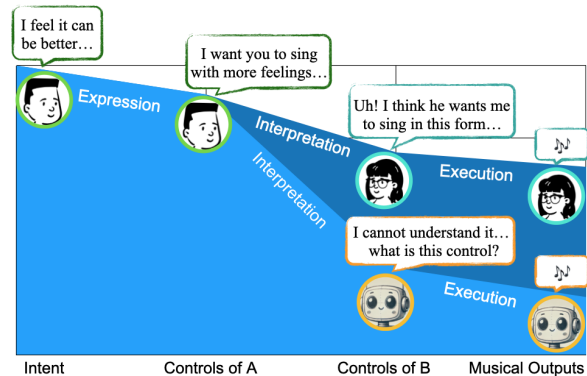


Figure 1: The comparison between human-human and human-AI interaction processes. We observe that the gap exists at both the interpretation stage and the execution stage, while the interpretation stage is often overlooked by current research.

output still exist, making effective collaboration with musicians challenging (Yakura and Goto, 2023; Newman et al., 2023; Ronchini et al., 2024; Majumder et al., 2024). In practice, we observe that musicians' control signals require interpretation before execution, and this process is often overlooked in current music information retrieval research. This oversight may hinder the practical applicability of these models in real-world musical settings. Figure 1 illustrates this issue through a single-round interaction among musician A and either musician B or a model. In this interaction, the control signals expressed by musician A are successfully interpreted by musician B before B generates the musical outputs. In contrast, the model fails to interpret these signals due to the neglected interpretation process in current text-to-music generation models.

In this paper, our contribution is threefold:

1. We propose a framework for the musical interaction process, consisting of three stages: expression, interpretation, and execution of control.
2. Our literature review identifies a communication

\*Both authors contributed equally.

gap in current models, which often fail to interpret controls in a way that aligns with human musicians’ natural communication methods.

- To address this gap, we propose two approaches: directly learning from human interpretation data or leveraging a strong prior understanding of human interpretation, such as that found in large language models (LLMs).

## 2 Interpretation of Controls

To begin with, we propose a general framework that conceptualizes the musical interaction procedure in three stages: the *expression*, *interpretation*, and *execution* of controls, as shown in Figure 2.




Figure 2: The proposed model that describes musical interaction process.

In an interaction between parties A and B:

- **Expression:** A’s intent is mapped to *Controls A*;
- **Interpretation:** B interprets *Controls A*, resulting in *Controls B*;
- **Execution:** B executes *Controls B*, producing the final musical output.

Table 1 provides several examples illustrating this framework. The framework encompasses both solo and multi-party musical interactions, with the interpretation stage becoming explicit in multi-party scenarios. Successful realization of the original intent hinges on effective mapping across all three stages of the process.

In this section, we examine the musical communication process following this framework. We observe that musical interactions often involve varying degrees of ambiguity in control expression, and skilled musicians can effectively interpret and execute these ambiguous instructions. In contrast, current text-to-music generation models struggle with this ambiguity, and can only understand highly semantical or highly precise instructions.

### 2.1 Musicians’ Interpretation of Controls

Musicians communicate through varying levels of ambiguity (Bishop, 2018). The most precise instructions often point to only one outcome (e.g.,

"Turn the bass 3 dB up") while the most abstract ones require much creative interpretation (e.g., "I want a *moody* synth"). Most communications, however, lie between these two extremes.

Consider this example of a producer addressing a vocalist: <sup>1</sup> "I want to try one where you just start this *chorus very soft*, and in that first phrase, like [*inaudible*]. You know what I mean? (*Sing to demonstrate*) just like, *get crazy* with it. Let’s start *quieter ... or softer*, or, *babier*. Just try it." This example showcases a wide range of communication types, from highly semantic descriptions (e.g., "very soft," "get crazy," "quieter," "softer," "babier") to performative instructions (e.g., "(Sing to demonstrate)"), and others that fall somewhere in between, requiring interpretation (e.g., "chorus," "first phrase," "[inaudible]").

Human musicians excel at interpreting musical instructions with varying ambiguity, a skill known as "musical taste" or "musicianship" (Sloboda, 1986). This ability enables jazz musicians to adapt improvisations (Berliner, 2009), film composers to modify scores for evolving narratives (Cooke, 2008), and orchestral conductors to guide an ensemble through gestures (Bishop et al., 2019) and verbal cues. This skill, which develops with experience (Lehmann et al., 2007), involves intuitive understanding of musical context, style, and intent (Meyer, 2008), allowing musicians to transform ambiguous directions into coherent expressions (Daniel et al., 2006).

### 2.2 Models’ Interpretation of Controls

While human musicians excel at interpreting ambiguous instructions, current music generation models struggle with this task. Traditional approaches to control often rely on disentangling representations in latent space (Luo et al., 2019; Wang et al., 2020). For music generative models, control mechanisms are typically implemented through various strategies. Some models integrate controls during initial large-scale pre-training, such as Mustango (Melechovsky et al., 2024) and MusicGen (Copet et al., 2024). Others employ post-training model augmentation, exemplified by Cocomulla (Lin et al., 2023), AIRGen (Lin et al., 2024), and Music ControlNet (Wu et al., 2024). Additionally, some approaches combine both stages’ efforts, as seen in MusicMagus (Zhang et al., 2024b),

<sup>1</sup>Billie Eilish In Studio Making Album "When We All Fall Asleep, Where Do We Go?", <https://www.youtube.com/watch?v=Sp-eNvKV0to>

Cases	Intent	Controls A	Controls B	Outputs
<i>Solo Interactions</i>				
Pianist	Light touch	→ Reduce finger force	→ N/A	→ Piano audio
Experienced Producer	Spacious sound	→ Reverb, cut lows	→ N/A	→ Natural result
Novice Producer	Spacious sound	→ Only adding reverb	→ N/A	→ Unnatural result
Composer	Modulate key	→ Write transition	→ N/A	→ Score
Experienced Guitarist	Emphasizing a chord	→ Use complex fingering	→ N/A	→ Clean strum sound
Novice Guitarist	Emphasizing a chord	→ Use complex fingering	→ N/A	→ Muffled strum sound
<i>Multi-Party Interactions</i>				
Producer & Experienced Vocalist	Emotive singing	→ "More feelings"	→ More dynamics & articulation	→ Emotional vocal track
Producer & Novice Vocalist	Emotive singing	→ "More feelings"	→ Sing closer to microphone	→ Unnatural vocal track
Experienced Rock Band	Guitar solo	→ Gesture	→ Drums and bass play fill; vocalist stop singing	→ Solo section
Novice Rock Band	Guitar solo	→ Gesture	→ Everyone ignores the guitarist	→ Solo fights with vocal, creating cacophony
Conductor & Orchestra	Crescendo	→ Rising arms	→ Gradually increasing dynamics	→ Balanced crescendo
DJ & Crowd	Build energy	→ Throwing hands up in the air	→ Crowd thinks it's peak	→ Premature movement

Table 1: Examples of solo and multi-party musical interactions.

Instruct-MusicGen (Zhang et al., 2024a), and Chat-Musician (Yuan et al., 2024a). Despite these advancements in control capabilities, current models still fall short of matching human-level interpretation of nuanced musical instructions.

We posit that the challenge lies not in control implementation methods, but in the nature of the controls themselves. Table 2 summarizes the controls offered by current models, typically either highly semantic (e.g., text descriptions) or highly specific (e.g., chords, melodies). These models struggle with both: for semantic inputs, they mainly interpret at keyword-level rather than understanding natural language (Wu et al., 2023), failing with concepts like negation and temporal order (Agostinelli et al., 2023; Yuan et al., 2024b); for specific inputs, they struggle with precise execution (Zhang et al., 2024a). When prompts combine semantic and specific instructions, models often fail to interpret the former and fail to execute the latter. The lack of support for other modalities, such as visual cues, also makes effective interpretation more difficult.

While resolving all these challenges is crucial, current research primarily focuses on improving execution ability, such as audio quality, while largely overlooking the interpretation stage. This oversight creates a significant gap in human-AI musical collaboration. Musicians are forced to adapt to the constrained and unnatural controls offered by these models, rather than the models adapting to musicians’ natural communication methods. We posit that this mismatch is a key factor in the limited adoption of these otherwise highly capable models

by musicians in practice.

### 3 Potential Solutions to Improve Interpretation of Controls

Addressing the interpretation gap between musicians and models is challenging due to the complex, multi-modal nature of musician communication, which includes visual cues, textual prompts, vocalizations, and musical references. No existing data sources comprehensively capture all modalities of music interactions, and creating such a dataset would be resource-intensive. Thus, we must approach the problem of learning interpretation under resource constraints. Given these limitations, two potential solutions emerge: directly learning from many aspects of human interpretation data, or leveraging a strong prior understanding of human interpretation, such as that encapsulated in large language models (LLMs). In the following sections, we explore these two avenues for enhancing AI models’ ability to interpret musical controls.

#### 3.1 Directly Learn from Human Interpretation Data

Previous research has explored many aspects of musical perception and interpretation, including auditory perception (Ananthabhotla et al., 2019; Wright and Välimäki, 2020; Manocha et al., 2020), emotion (Yang and Chen, 2012; Dash and Agres, 2023), song and artist similarity (Knees and Schedl, 2013; Allik et al., 2018), music discussions (Hauger et al., 2013), recommendation systems (Bertin-Mahieux et al., 2011), and non-verbal communications, such

Model	Semantic controls	Precise controls
<i>Integrated Controls in Foundation Models</i>		
Mustango (Melechovsky et al., 2024)	Text description, metadata	-
MusicGen (Copet et al., 2024)	Text description	melody spectrogram
Diff-A-Riff (Nistal et al., 2024)	Text description	Music audio mixture
Jen-1 Composer (Yao et al., 2023)	Text description	Other instrument tracks
GMSDI (Postolache et al., 2024)	Instrument name	Other instrument tracks
<i>Control Enhancement Modules</i>		
Coco-mulla (Lin et al., 2023)	Text description	Drum track, chord, melody
AIRGen (Lin et al., 2024)	Text description	Drum track, chord, melody
JASCO (Tal et al., 2024)	Text description	Drum track, chord, melody
Music ControlNet (Wu et al., 2024)	Text description	Dynamic, melody, rhythm
Jen-1 DreamStyler (Chen et al., 2024)	Text description	Reference music audio
<i>Music Editing Methods</i>		
MusicMagus (Zhang et al., 2024b)	Text swapping	Music audio mixture
InstructME (Han et al., 2023)	Edit instruction	Music audio mixture
Instruct-MusicGen (Zhang et al., 2024a)	Edit instruction	Music audio mixture
Loop Copilot (Zhang et al., 2023)	Edit instruction	Conversational context (music audio, text)
M <sup>2</sup> UGen (Hussain et al., 2023)	Edit instruction	Conversational context (music audio, text)
ChatMusician (Yuan et al., 2024a)	Edit instruction	Conversational context (symbolic music, text)

Table 2: List of representative text-to-music generation models with extra controls. Most controls can be classified into high-level semantic controls and low-level signal-level controls, while the exploration of intermediate-level musicians’ communication controls are limited.

as gesture and dance movements (Gillian, 2012; Fan et al., 2011). These studies often rely on crowd-sourced evaluations or public data, achieving good interpretations that can serve as control signals, as demonstrated by Huang et al. (2024) in music generation from dance movements.

Learning directly from these diverse sources require combining them into cohesive controls, which may be achieved through pseudo-description generation, an approach that has shown promise in music captioning (Mei et al., 2024; Doh et al., 2023) and understanding (Liu et al., 2024).

### 3.2 LLMs for Musical Interpretation

LLMs’ robust language understanding enables the decomposition of user queries into specialized tasks, an approach pioneered by Hugging-GPT (Shen et al., 2024). This method has inspired audio domain projects such as Loop Copilot (Zhang et al., 2023), WavJourney (Liu et al., 2023), WavCraft (Liang et al., 2024), and MusicAgent (Yu et al., 2023). Jiang et al. (2024) explores synthesizing natural language from control parameters for model training.

However, user studies (Gianet et al., 2024; Newman et al., 2023; Ronchini et al., 2024; Zhang et al., 2023) reveal that professional musicians often experience misalignment between model interpretations and their intentions, primarily due to LLMs’ lack of domain-specific musical knowl-

edge (Li et al., 2024). Research in other domains indicates that simply integrating domain knowledge can significantly enhance LLMs’ capabilities (Lee et al., 2024). Consequently, we posit that by collecting domain knowledge and natural music conversations incorporating this knowledge, we could effectively boost LLMs’ ability to execute music tasks. Furthermore, these enhanced LLMs could potentially generate synthetic training data for developing more compact interpretation models.

## 4 Conclusion

We identify a critical gap in text-to-music generation models: their inability to effectively interpret musicians’ controls. We propose a three-stage framework for musical interaction: expression, interpretation, and execution, and highlight how current AI models often struggle with the crucial interpretation stage. To address this gap, we suggest two potential solutions: directly learning from various sources of human interpretation data and leveraging large language models for musical interpretation. We call on the MIR community to prioritize research in this area, as improving the interpretation capabilities is crucial for their integration into creative workflows and for realizing their full potential as collaborative tools for musicians.

## Ethics Statement

Our work includes YouTube video transcript excerpts demonstrating artists' creative processes, used solely to illustrate our proposed framework. We thank these amazing artists for sharing their creative processes. All copyrights remain with the original video owners, and excerpts are included for research purposes only.

We acknowledge that musical communications and interpretations encapsulate diverse musicianship, tastes, and cultural nuances. While some aspects of musical communications may be universal, they are often influenced by social culture and individual experiences. We encourage the community to be mindful of this diversity when modeling musical interpretations, as capturing these nuances can enhance the music creation process with generative models.

## References

- Andrea Agostinelli, Timo I Denk, Zalán Borsos, Jesse Engel, Mauro Verzetti, Antoine Caillon, Qingqing Huang, Aren Jansen, Adam Roberts, Marco Tagliasacchi, et al. 2023. Musiclm: Generating music from text. *arXiv preprint arXiv:2301.11325*.
- Alo Allik, Florian Thalmann, and Mark Sandler. 2018. Musiclynx: Exploring music through artist similarity graphs. In *Companion Proceedings of the The Web Conference 2018*, pages 167–170.
- Ishwarya Ananthabhotla, Sebastian Ewert, and Joseph A Paradiso. 2019. Towards a perceptual loss: Using a neural network codec approximation as a loss for generative audio models. In *Proceedings of the 27th acm international conference on multimedia*, pages 1518–1525.
- Paul F Berliner. 2009. *Thinking in jazz: The infinite art of improvisation*. University of Chicago Press.
- Thierry Bertin-Mahieux, Daniel PW Ellis, Brian Whitman, and Paul Lamere. 2011. The million song dataset.
- Laura Bishop. 2018. Collaborative musical creativity: How ensembles coordinate spontaneity. *Frontiers in psychology*, 9:1285.
- Laura Bishop, Carlos Cancino-Chacón, and Werner Goebel. 2019. Moving to communicate, moving to interact: Patterns of body motion in musical duo performance. *Music Perception: An Interdisciplinary Journal*, 37(1):1–25.
- Renaud Bougueng Tchameube, Jeffrey John Ens, and Philippe Pasquier. 2022. Calliope: A co-creative interface for multi-track music generation. In *Proceedings of the 14th Conference on Creativity and Cognition*, pages 608–611.
- Boyu Chen, Peike Li, Yao Yao, and Alex Wang. 2024. Jen-1 dreamstyler: Customized musical concept learning via pivotal parameters tuning. *arXiv preprint arXiv:2406.12292*.
- Mervyn Cooke. 2008. *A history of film music*. Cambridge University Press.
- Jade Copet, Felix Kreuk, Itai Gat, Tal Remez, David Kant, Gabriel Synnaeve, Yossi Adi, and Alexandre Défossez. 2024. Simple and controllable music generation. *Advances in Neural Information Processing Systems*, 36.
- J Daniel et al. 2006. This is your brain on music: The science of a human obsession.
- Adyasha Dash and Kathleen Agres. 2023. Ai-based affective music generation systems: A review of methods and challenges. *ACM Computing Surveys*.
- SeungHeon Doh, Keunwoo Choi, Jongpil Lee, and Juhan Nam. 2023. Lp-musiccaps: Llm-based pseudo music captioning. *arXiv preprint arXiv:2307.16372*.
- Zach Evans, CJ Carr, Josiah Taylor, Scott H Hawley, and Jordi Pons. Fast timing-conditioned latent audio diffusion. In *Forty-first International Conference on Machine Learning*.
- Zach Evans, Julian D Parker, CJ Carr, Zack Zukowski, Josiah Taylor, and Jordi Pons. 2024. Long-form music generation with latent diffusion. *arXiv preprint arXiv:2404.10301*.
- Rukun Fan, Songhua Xu, and Weidong Geng. 2011. Example-based automatic music-driven conventional dance motion synthesis. *IEEE transactions on visualization and computer graphics*, 18(3):501–515.
- Eric Tron Gianet, Luigi Di Caro, and Amon Rapp. 2024. Music composition as a lens for understanding human-ai collaboration.
- Nicholas Edward Gillian. 2012. *Gesture recognition for musician computer interaction*. Ph.D. thesis, Cite-seer.
- Bing Han, Junyu Dai, Xuchen Song, Weituo Hao, Xinyan He, Dong Guo, Jitong Chen, Yuxuan Wang, and Yanmin Qian. 2023. Instructme: An instruction guided music edit and remix framework with latent diffusion models. *arXiv preprint arXiv:2308.14360*.
- David Hauger, Markus Schedl, Andrej Košir, and Marko Tkalcic. 2013. The million musical tweets dataset: What can we learn from microblogs. In *Proc. ISMIR*, pages 189–194.
- Cheng-Zhi Anna Huang, Hendrik Vincent Koops, Ed Newton-Rex, Monica Dinulescu, and Carrie J Cai. 2020. Ai song contest: Human-ai co-creation in songwriting. *arXiv preprint arXiv:2010.05388*.



- Jiang Huang, Xianglin Huang, Lifang Yang, and Zhulin Tao. 2024. Dance-conditioned artistic music generation by creative-gan. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 107(5):836–844.
- Atin Sakkeer Hussain, Shansong Liu, Chenshuo Sun, and Ying Shan. 2023. M<sup>2</sup>ugen: Multi-modal music understanding and generation with the power of large language models. *arXiv preprint arXiv:2311.11255*.
- Xilin Jiang, Cong Han, Yinghao Aaron Li, and Nima Mesgarani. 2024. Listen, chat, and edit: Text-guided soundscape modification for enhanced auditory experience. *arXiv preprint arXiv:2402.03710*.
- Peter Knees and Markus Schedl. 2013. A survey of music similarity and recommendation from music context data. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 10(1):1–21.
- Younghun Lee, Dan Goldwasser, and Laura Schwab Reese. 2024. Towards understanding counseling conversations: Domain knowledge and large language models. *arXiv preprint arXiv:2402.14200*.
- Andreas C Lehmann, John A Sloboda, and Robert H Woody. 2007. *Psychology for musicians: Understanding and acquiring the skills*. Oxford University Press.
- Jiajia Li, Lu Yang, Mingni Tang, Cong Chen, Zuchao Li, Ping Wang, and Hai Zhao. 2024. The music maestro or the musically challenged, a massive music evaluation benchmark for large language models. *arXiv preprint arXiv:2406.15885*.
- Jinhua Liang, Huan Zhang, Haohe Liu, Yin Cao, Qiuqiang Kong, Xubo Liu, Wenwu Wang, Mark D Plumbley, Huy Phan, and Emmanouil Benetos. 2024. Wavcraft: Audio editing and generation with large language models. In *ICLR 2024 Workshop on Large Language Model (LLM) Agents*.
- Liwei Lin, Gus Xia, Junyan Jiang, and Yixiao Zhang. 2023. Content-based controls for music large language modeling. *arXiv preprint arXiv:2310.17162*.
- Liwei Lin, Gus Xia, Yixiao Zhang, and Junyan Jiang. 2024. Arrange, inpaint, and refine: Steerable long-term music audio generation and editing via content-based controls. *arXiv preprint arXiv:2402.09508*.
- Shansong Liu, Atin Sakkeer Hussain, Chenshuo Sun, and Ying Shan. 2024. 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)*, pages 286–290. IEEE.
- Xubo Liu, Zhongkai Zhu, Haohe Liu, Yi Yuan, Meng Cui, Qiushi Huang, Jinhua Liang, Yin Cao, Qiuqiang Kong, Mark D Plumbley, et al. 2023. Wavjourney: Compositional audio creation with large language models. *arXiv preprint arXiv:2307.14335*.
- Yin-Jyun Luo, Kat Agres, and Dorien Herremans. 2019. Learning disentangled representations of timbre and pitch for musical instrument sounds using gaussian mixture variational autoencoders. *arXiv preprint arXiv:1906.08152*.
- Navonil Majumder, Chia-Yu Hung, Deepanway Ghosal, Wei-Ning Hsu, Rada Mihalcea, and Soujanya Poria. 2024. Tango 2: Aligning diffusion-based text-to-audio generations through direct preference optimization. *arXiv preprint arXiv:2404.09956*.
- Pranay Manocha, Adam Finkelstein, Richard Zhang, Nicholas J Bryan, Gautham J Mysore, and Zeyu Jin. 2020. A differentiable perceptual audio metric learned from just noticeable differences. *arXiv preprint arXiv:2001.04460*.
- Xinhao Mei, Chutong Meng, Haohe Liu, Qiuqiang Kong, Tom Ko, Chengqi Zhao, Mark D Plumbley, Yuexian Zou, and Wenwu Wang. 2024. Wavcaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Jan Melechovsky, Zixun Guo, Deepanway Ghosal, Navonil Majumder, Dorien Herremans, and Soujanya Poria. 2024. Mustango: Toward controllable text-to-music generation. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8286–8309.
- Leonard B Meyer. 2008. *Emotion and meaning in music*. University of chicago Press.
- Michele Newman, Lidia Morris, and Jin Ha Lee. 2023. Human-ai music creation: Understanding the perceptions and experiences of music creators for ethical and productive collaboration. In *ISMIR*, pages 80–88.
- Javier Nistal, Marco Pasini, Cyran Auameur, Maarten Grachten, and Stefan Lattner. 2024. Diff-a-riff: Musical accompaniment co-creation via latent diffusion models. *arXiv preprint arXiv:2406.08384*.
- Emilian Postolache, Giorgio Mariani, Luca Cosmo, Emmanouil Benetos, and Emanuele Rodolà. 2024. Generalized multi-source inference for text conditioned music diffusion models. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6980–6984. IEEE.
- Simeon Rau, Frank Heyen, Stefan Wagner, and Michael Sedlmair. 2022. Vi-sualization for ai-assisted composing. In *Proc. of the 23rd Int. Society for Music Information Retrieval Conf.*
- Francesca Ronchini, Luca Comanducci, Gabriele Perego, and Fabio Antonacci. 2024. Paguri: a user experience study of creative interaction with text-to-music models. *arXiv preprint arXiv:2407.04333*.

- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2024. Hugging-gpt: Solving ai tasks with chatgpt and its friends in hugging face. *Advances in Neural Information Processing Systems*, 36.
- John A Sloboda. 1986. *The musical mind: The cognitive psychology of music*. Oxford University Press.
- Or Tal, Alon Ziv, Itai Gat, Felix Kreuk, and Yossi Adi. 2024. Joint audio and symbolic conditioning for temporally controlled text-to-music generation. *arXiv preprint arXiv:2406.10970*.
- Ziyu Wang, Dingsu Wang, Yixiao Zhang, and Gus Xia. 2020. Learning interpretable representation for controllable polyphonic music generation. *arXiv preprint arXiv:2008.07122*.
- Alec Wright and Vesa Välimäki. 2020. Perceptual loss function for neural modeling of audio systems. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 251–255. IEEE.
- Ho-Hsiang Wu, Oriol Nieto, Juan Pablo Bello, and Justin Salamon. 2023. Audio-text models do not yet leverage natural language. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Shih-Lun Wu, Chris Donahue, Shinji Watanabe, and Nicholas J Bryan. 2024. Music controlnet: Multiple time-varying controls for music generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32:2692–2703.
- Hiromu Yakura and Masataka Goto. 2023. Iteratta: An interface for exploring both text prompts and audio priors in generating music with text-to-audio models. In *Ismir 2023 Hybrid Conference*.
- Yi-Hsuan Yang and Homer H Chen. 2012. Machine recognition of music emotion: A review. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(3):1–30.
- Yao Yao, Peike Li, Boyu Chen, and Alex Wang. 2023. Jen-1 composer: A unified framework for high-fidelity multi-track music generation. *arXiv preprint arXiv:2310.19180*.
- Dingyao Yu, Kaitao Song, Peiling Lu, Tianyu He, Xu Tan, Wei Ye, Shikun Zhang, and Jiang Bian. 2023. Musicagent: An ai agent for music understanding and generation with large language models. *arXiv preprint arXiv:2310.11954*.
- Ruibin Yuan, Hanfeng Lin, Yi Wang, Zeyue Tian, Shangda Wu, Tianhao Shen, Ge Zhang, Yuhang Wu, Cong Liu, Ziya Zhou, et al. 2024a. Chatmusician: Understanding and generating music intrinsically with llm. *arXiv preprint arXiv:2402.16153*.
- Yi Yuan, Zhuo Chen, Xubo Liu, Haohe Liu, Xuenan Xu, Dongya Jia, Yuanzhe Chen, Mark D Plumbley, and Wenwu Wang. 2024b. T-clap: Temporal-enhanced contrastive language-audio pretraining. *arXiv preprint arXiv:2404.17806*.
- Yixiao Zhang, Yukara Ikemiya, Woosung Choi, Naoki Murata, Marco A Martínez-Ramírez, Liwei Lin, Gus Xia, Wei-Hsiang Liao, Yuki Mitsufuji, and Simon Dixon. 2024a. Instruct-musicgen: Unlocking text-to-music editing for music language models via instruction tuning. *arXiv preprint arXiv:2405.18386*.
- Yixiao Zhang, Yukara Ikemiya, Gus Xia, Naoki Murata, Marco Martínez, Wei-Hsiang Liao, Yuki Mitsufuji, and Simon Dixon. 2024b. Musicmagus: Zero-shot text-to-music editing via diffusion models. *arXiv preprint arXiv:2402.06178*.
- Yixiao Zhang, Akira Maezawa, Gus Xia, Kazuhiko Yamamoto, and Simon Dixon. 2023. Loop copilot: Conducting ai ensembles for music generation and iterative editing. *arXiv preprint arXiv:2310.12404*.
- Yixiao Zhang, Gus Xia, Mark Levy, and Simon Dixon. 2021. Cosmic: A conversational interface for human-ai music co-creation. In *NIME 2021*. PubPub.