

# CMI-BENCH: A COMPREHENSIVE BENCHMARK FOR EVALUATING MUSIC INSTRUCTION FOLLOWING

Yinghao Ma<sup>1</sup> Siyou Li<sup>1</sup> Juntao Yu<sup>1</sup> Emmanouil Benetos<sup>1</sup> Akira Maezawa<sup>2</sup>

<sup>1</sup> Queen Mary University of London, London, UK

<sup>2</sup> Yamaha Corporation, Hamamatsu, Japan

emmanouil.benetos@qmul.ac.uk, akira.maezawa@music.yamaha.com

## ABSTRACT

Recent advances in audio-text large language models (LLMs) have opened new possibilities for music understanding and generation. However, existing benchmarks are limited in scope, often relying on simplified tasks or multi-choice evaluations that fail to reflect the complexity of real-world music analysis. We reinterpret a broad range of traditional MIR annotations as instruction-following formats and introduce CMI-Bench, a comprehensive music instruction following benchmark designed to evaluate audio-text LLMs on a diverse set of music information retrieval (MIR) tasks. These include genre classification, emotion regression, emotion tagging, instrument classification, pitch estimation, key detection, lyrics transcription, melody extraction, vocal technique recognition, instrument performance technique detection, music tagging, music captioning, and (down)beat tracking — reflecting core challenges in MIR research. Unlike previous benchmarks, CMI-Bench adopts standardized evaluation metrics consistent with previous state-of-the-art MIR models, ensuring direct comparability with supervised approaches. We provide an evaluation toolkit supporting all open-source audio-textual LLMs, including LTU, Qwen-audio, SALMONN, MusiLingo, etc. Experiment results reveal significant performance gaps between LLMs and supervised models, along with their culture, chronological and gender bias, highlighting the potential and limitations of current models in addressing MIR tasks. CMI-Bench establishes a unified foundation for evaluating music instruction following, driving progress in music-aware LLMs.

## 1. INTRODUCTION

The emergence of large language models (LLMs) has reshaped the landscape of natural language processing by enabling general-purpose models to solve a wide variety of tasks through instruction following. This paradigm—where models are trained not just on pre-text corpora but instruction-response pairs—has unlocked new possibilities in model generalization, few-shot learning, and cross-domain reasoning. Supervised fine-tuning (SFT), also

known as instruction finetuning, and reinforcement learning from human feedback have further strengthened LLMs’ ability to align with human intent [1].

In the context of music, the instruction-following paradigm holds particular promise. Many music-related tasks are naturally multimodal and domain-specific and often lack large-scale annotated data. Instruction-tuned models can generalize to previously unseen problems such as chord generation under rhythmic constraints or personalized music recommendation based on context. Besides, by supporting in-context learning, LLMs offer a flexible path to interact with world music traditions, rare genres, and diverse user preferences—all without explicit retraining [2].

Recently, a growing number of audio LLMs [3–6], extended LLMs with audio encoders and instruction-following capabilities. However, these models have so far been evaluated on limited tasks, relying on caption similarity metrics on datasets like [3, 4], single-choice protocols [7–9], or multiple-choice question (MCQ) protocols [10]. Despite these successes, such evaluations fail to capture the complexity of core music information retrieval (MIR) tasks and offer limited insight to real-world performance.

This work makes three key contributions: First, we reinterpret a broad range of core MIR annotations as instruction-following tasks as illustrated in Figure 2, enabling the use of a wide range of MIR datasets, including sequential tasks, not only for evaluation but for training and SFT audio-text LLMs. Second, we provide a standardized benchmarking framework that includes implementations of major open-source audio-text LLMs, along with evaluation metrics aligned with prior MIR literature. Unlike earlier MCQ protocols, CMI-bench adopts open-ended, task-specific metrics, allowing more rigorous comparisons. Last, we present an initial analysis of generalization outside training data, along with cultural and gender bias across models, uncovering potential limitations in their generalization and pointing future directions for culturally inclusive music AI. Together, these contributions lay groundwork for systematic progress in music instruction following and its intersection with traditional MIR. The code <sup>1</sup> and testset audio <sup>2</sup> are available.

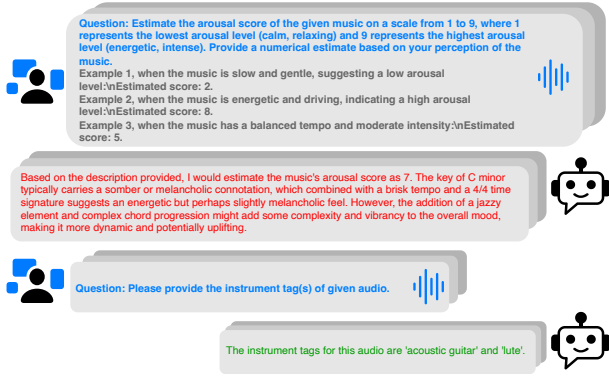
This paper is organized as follows. Section 2 reviews related work. Section 3 introduces the CMI-Bench design and instruction construction procedure, Section 4 describes the



© Y. Ma, S. Li, J. Yu, E. Benetos, and A. Maezawa. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Y. Ma, S. Li, J. Yu, E. Benetos, and A. Maezawa, “CMI-Bench: A Comprehensive Benchmark for Evaluating Music Instruction Following”, in *Proc. of the 26th Int. Society for Music Information Retrieval Conf.*, Daejeon, South Korea, 2025.

<sup>1</sup> <https://github.com/nicolaus625/CMI-bench/>

<sup>2</sup> <https://huggingface.co/datasets/nicolaus625/CMI-bench>



**Figure 1.** Instruction-Following Format Data in CMI-Bench and example response from Qwen2-audio.

experimental setup, including evaluation protocol and open-source models. Section 5 presents benchmarking results. Finally, section 6 draw the conclusion.

## 2. RELATED WORK

### 2.1 Instruction Following Dataset

**Instruction following** refers to the ability of LLMs to perform tasks based on natural language prompts that describe the task itself. This paradigm has become central to recent advances in NLP, where SFT models are trained on a wide range of instruction-response pairs. Super-NaturalInstructions [11] rewrote annotations from over 1,600 diverse NLP tasks into instruction-following formats, showing that models can generalize to unseen tasks given clear instructions. Self-Instruct [12] further advanced this approach by automatically generating diverse instruction-response pairs using the model’s own outputs, while Instructional-GPT [13] aligned models with human intent through SFT and reinforcement learning.

These techniques have recently been extended to music, enabling instruction-following models to engage with multimodal and domain-specific tasks. MusicQA [3] and MusicInstruct [4] repurpose descriptions and tags from MIR dataset to generate Q&A pairs. Such a dataset does not distinguish subtask on instrument, emotion, genre, and caption Q&A pairs, and the evaluating metrics are BERT-score, overestimate model’s music understanding capability without equally compared with traditional MIR algorithms. Finally, Audio-FLAN [14] presents a large-scale instruction-tuning corpus across 80 tasks, unifying understanding and generation in audio, music, and speech. Yet, many tasks are paraphrased to an MCQ format, significantly smaller than the range of pre-defined classes or labels. Furthermore, these works do not provide a model performance benchmark on such tasks, and the evaluation metrics are not compatible with traditional MIR method.

### 2.2 Instruction-Following Benchmarks for Music

While instruction-following has shown great promise in natural language and vision tasks, its application in the music domain remains underexplored. ZIQI-Eval [15] is an

instruction following the benchmark on textual symbolic music. AIR-Bench [7] covers a broader range of audio types, including music, but emphasizes low-level tasks such as pitch and instrument recognition and relies primarily on MCQ formats. MMAU [10] includes music reasoning, yet only covers six MIR datasets, lacks alignment with MIR-specific evaluation metrics, and reports only average scores across tasks. MuCho Music [8] evaluates music understanding in multimodal models through 1,187 MCQs. AudioBench [16] and MusicBench [17] primarily target audio and text-to-music generation respectively, without addressing MIR tasks. MuChin [18], while valuable for colloquial descriptions, is tailored to Chinese pop song generation. Across these efforts, most benchmarks omit key MIR tasks popularized by MIREX, few support sequential tasks, and evaluation protocols often rely on multiple-choice questions rather than the task-specific metrics used in supervised MIR literature.

### 2.3 Audio-Textual Large Language Models

Current audio-textual LLMs typically consist of an encoder of speech, audio or music, an intermediate architecture and an LLM backbone. LTU [19] and LTU-AS [20] focus on general audio comprehension and reasoning, combining whisper speech encoder [21] w. MU-LLaMA [3], MusiLingo [4] and Lark [22] are tailored for music-related tasks, leveraging audio encoders and instructional datasets to support captioning and open-ended music question answering. Pengi [23] frames all audio tasks as text generation, unifying audio perception with LLM-based reasoning via a simple prefix-tuning strategy. GAMA [24] and GAMA-IT [24] integrate multi-layer audio features and instruction tuning (CompA-R) to support complex reasoning over general audio, including music. SALMONN-Audio [6] introduces a Q-former window architecture for sequential speech and sound understanding. Qwen-Audio [5] and Qwen2-Audio [25] scale instruction tuning across over 30 audio tasks with hierarchical or natural prompts. Audio-Flamingo [26] and Audio-Flamingo2 [27] incorporate in-context learning and retrieval-based adaptation for audio-text interaction and dialogue. Beyond open-source models, proprietary systems Gemini-2.5 Pro and GPT-4o may represent the state-of-the-art (SOTA).

## 3. CMI-BENCHMARK

With CMI-Bench, we aim to address the following limitations in evaluating music understanding capabilities of audio-text LLMs. Previous benchmarks often cover only a narrow range of tasks, and no benchmark supports sequential tasks, overlooking many classic challenges which are central to MIR research. Moreover, evaluation protocols are typically inconsistent with standard MIR metrics, difficult to compare against traditional supervised models. To address these issues, we reformulate annotations from widely-used MIR datasets into instruction-following prompts and process model outputs into formats compatible with standard MIR Python library `mir_eval` [28].

### 3.1 Overview

Tasks	Dataset	Metrics	#Test Samples
Key detection	GS [29]	Gmean score	2406
Emotion Regression	EMO [30]	$R^2$	125
Music tagging	MagnaTagATune [31]	ROC-AUC, PR-AUC	5329
	MTG-Top50 [32]	ROC-AUC, PR-AUC	11356
Instrument Classification	MTG-Instrument [32]	ROC-AUC, PR-AUC	5115
	Nsynth-Instrument [33]	Accuracy	4096
Genre classification	MTG-Genre [32]	ROC-AUC, PR-AUC	11479
	GTZAN [34]	Accuracy	290
Emotion tagging	MTG-Emotion [32]	ROC-AUC, PR-AUC	4231
Pitch Estimation	Nsynth-Pitch [33]	Accuracy	4096
Singing Techniques	VocalSet [35]	Accuracy	1140
Music Captioning	SDD [36]	BL., ME., RO., Bert-Score	1106
	MusicCaps [37]	BL., ME., RO., Bert-Score	2813
Lyrics Transcription	DSing [38]	WER, CER	482
Beat tracking	GTZAN-Rhythm [34]	F_measure	290
	ballroom [39, 40]	F_measure	685
DownBeat tracking	GTZAN-Rhythm	F_measure	290
	ballroom	F_measure	685
Melody Extraction	MedleyDB v2 [41]	Melody Accuracy	618
Performance Technique	GuZheng_99 [42]	frame-level micro/macro-f1	94

**Table 1.** Overview of tasks, datasets, evaluation metrics, and the number of test samples in the CMI-Bench.

The CMI-Benchmark encompasses 14 tasks spanning multi-class, multi-label, regression, captioning, and sequential prediction challenges, evaluated across 20 diverse datasets. This benchmark integrates traditional MIR tasks with emerging music-and-language objectives, providing a robust platform to assess computational music intelligence. The tasks and datasets used in the benchmark are shown in Table 1. By standardizing splits and metrics, CMI-bench ensures reproducibility and fair comparisons.

### 3.2 Self-Instruction of MIR Annotations

In this subsection, we introduce the self-instruction framework for CMI-Bench designed to unify diverse MIR tasks under a consistent NLP paradigm, outlining the design of instructions and input tailored to tasks such as key estimation, genre classification, emotion regression, instrument tagging, and temporal sequence annotations. Our approach leverages structured prompts with multi-class, regression, and sequence-based outputs, enriched with few-shot examples to guide annotation generation.

For multi-label tasks, we allow flexible outputs without providing pre-defined tags, reflecting real-world complexity. For clip-level multi-class tasks with a manageable number of categories, such as musical key estimation and genre and vocal techniques classification, instructions explicitly list all possible choices. For instance, key estimation requires selecting one of 24 major and minor keys, with few-shot examples like "Bb major" to clarify the format. In cases with larger class sets, such as pitch classification on short excerpts across MIDI numbers 9 to 119, we provide a definition of MIDI standard alongside examples (e.g., "A4: 69", "Middle C (C4) = 60") to anchor the task.

Regression tasks, such as arousal estimation, adopt a numerical scale (1 to 9) with descriptive anchors—1 for "calm, relaxing" and 9 for "energetic, intense." To better utilize LLM-s in-context learning capability, we include examples for few-shot learning on tie scores to musical characteristics (e.g., "slow and gentle: 2," "energetic and driving: 8"), enabling precise emotional annotation.

Temporal tasks, such as beat tracking and instrument performance technique detection, require structured outputs. Beat tracking outputs timestamps in a comma-separated format (e.g., "0.1s, 1.19s, 2.25s"), while Guzheng (traditional Chinese Kyoto) technique detection uses a Python-style list of tuples (e.g., "[('70.8086', '71.4817', 'Tremolo')]")., covering techniques like Vibrato and Glissando. Default outputs "[('0.0', '10.0', 'No Tech')]" handle cases with no detections. Melody extraction follows similar principles, balancing specificity and clarity. We forbid tuples to have time overlapping on melody and (down)beat tracking, but allow for playing technique detections.

Inputs are uniformly represented as audio placeholders (e.g., "<ISOA><AUDIO><IEOA>"), paired with metadata such as audio paths and time segments. This ensures compatibility with NLP models while preserving MIR task diversity, offering a scalable framework for future efforts.

## 4. EXPERIMENTS

### 4.1 Evaluation Protocol

To enable rigorous and fair comparison with traditional MIR systems, we design an evaluation pipeline that closely follows the original task definitions and metrics. All model outputs are automatically post-processed to conform to each task’s expected format, ensuring compatibility with MIR evaluation tools such as `mir_eval`. Below, we detail the evaluation strategies used for each task category.

#### 4.1.1 Classification Tasks

**Multi-Class Classification.** Tasks include short-clip monophonic pitch estimation, instrument classification, singing technique classification, and genre classification. We evaluate using strict string matching: a model’s response is considered correct if it contains only the correct label (case-, space-, and punctuation-insensitive) and no others. For pitch classification, we additionally require the model to follow the instruction format and return MIDI numbers. Accuracy is used as the metric.

**Multi-Label Classification** Tasks include music tagging, genre labelling, emotion tagging, and instrument recognition. As model responses may include synonyms or free-form text, we embed both the predicted and ground truth tag sets using the BGE encoder [43], a model optimized for retrieval and multi-label matching. Cosine similarity scores are then used to compute ROC and PR, providing a soft but semantically aligned evaluation quality.

#### 4.1.2 Clip-level MIR Tasks

**Key Detection.** We adopt standard weighted score metric from `mir_eval.key`, which accounts for musically reasonable errors, such as relative minor or parallel key.

**Regression.** Model outputs are constrained to integers in the range [1, 9]; if a float is returned, we take the floor. Outputs are then z-score normalized to zero mean and unit variance. If a model fails to return a value, we assign the model’s mean value. The coefficient of determination ( $R^2$ )

is computed between predictions and annotations on arousal and valance.

**Music Captioning.** We assess caption quality using four standard NLP metrics: BLEU [44, 45], METEOR [46], ROUGE [47], and Bert-Score [48].

#### 4.1.3 Sequential MIR Tasks

**Lyrics Transcription.** We extract lyrics from model outputs by removing typical prefixes (e.g., “lyrics is as follows:”). Word Error Rate (WER) and Character Error Rate (CER) are computed against ground-truth lyrics.

**(Down)Beat Tracking.** Model are expected to return a list of time points for (down)beat events. We filter non-numeric outputs, sort the list by time, and apply F-measure metric from `mir_eval.beat`, with a 20ms tolerance.

**Melody Extraction** is treated as a sequential regression task on the fundamental frequency of notes calculated by `mir_eval.melody.evaluate` with 50 music cents tolerance. Models are instructed to return a list of (time, pitch) tuples. We discard invalid tuples (e.g., missing pitches, or improperly formatted entries, etc.). If multiple pitches are predicted for the same timestamp, we use only the first. Evaluation is based on frame-level accuracy.

**Instrument Playing Technique Detection.** For the GuZheng\_99 dataset, we evaluate frame-level predictions using macro- and micro-F1 scores, allowing for overlapping techniques. Invalid predictions (e.g., incorrect tuple formats) are filtered out. Empty responses are interpreted as a “no technique” prediction covering the full time range.

## 4.2 Models

Model	#Params	Sound Encoder	Music Architecture	Speech Decoder
Pengi [23]	323M	✓	✓	
Audio-Flamingo [26]	2.2B	✓	✓	
LTU [19]	7B	✓	✓	
LTU-AS [20]	7B	✓	✓	✓
MusiLingo-long [4]	7B		✓	
MuLLaMa [3]	7B		✓	
GAMA [24]	7B	✓	✓	
GAMA-IT [24]	7B	✓	✓	
Qwen-Audio-Chat [5]	8.4B	✓		
Qwen2-Audio-Instruct [25]	8.4B	✓	✓	✓
SALAMONN-Audio [6]	13B	✓	✓	✓

**Table 2.** Comparison of audio-textual LLMs by training domains. ✓denotes coverage or presence; ✗absence.

To provide a broad and representative evaluation, we implement and benchmark 11 audio-text LLMs with publicly available weights demonstrated in table 2. Our selection covers a wide spectrum of model designs and training corpus, enabling a comprehensive comparison of instruction-following capabilities across various music-specific tasks.

## 5. RESULTS AND DISCUSSION

### 5.1 Benchmarking Results

Experiment results reveal several important observations about the current state of audio-text LLMs on MIR tasks.

#### 5.1.1 LLMs Underperform Traditional MIR Baselines.

Despite LLMs have achieved excellent results on music captions and multi-choices QA, [4, 8–10, 37], all models in our study fall significantly short of the performance achieved by task-specific supervised systems when evaluated using standard MIR metrics besides music captioning. This is consistent across classification, regression, and sequential tasks. These findings suggest that instruction-following LLMs still lack the specialized precision and inductive bias of MIR models trained explicitly for each task.

#### 5.1.2 Best Performance May Skew toward Training Set

Interestingly, the peak performance on each task is typically achieved by models whose datasets overlap significantly with their training corpus, revealing limited generalization. Qwen2-Audio performs best on MTG-Jamendo-related tasks such as MTG-top50, MTG-Emotion, and SDD captioning, while common on other tagging and caption datasets. This aligns with its use of MTG-Jamendo and FMA during model development via AIR-Bench, suggesting unsatisfying generalization capability. Besides, MusiLingo performs best on MusicCaps, the same dataset it was trained on for captioning and Q&A. Lastly, GAMA shows the best on MTT and NSynth-instrument and comparative on MusicCaps, while common on other datasets on same tasks, reflecting bias in its SFT corpus. These demonstrate that supervised instruction-tuned models can capture task-specific patterns well when training data is directly aligned, but their generalization to unseen or structurally different tasks remains limited.

#### 5.1.3 All Models Perform Poorly on DSing Transcription

Despite the absence of instrumental accompaniment and use of English lyrics, none of the models reach usable performance levels on DSing for lyrics transcription, though it is relatively clean. This result is particularly striking for models like LTU and SALMONN, which include Whisper as their audio encoder and could theoretically benefit from ASR capabilities. LyricWhiz [64] utilizes GPT-4 to post-process whisper ASR output on DSing dataset, providing results similar to SOTA without training.

#### 5.1.4 Prompting Format May Impacts Performance.

Prompting without task-specific tokens used during training significantly degrades performance. Qwen-Audio performs far worse on Nsynth-Pitch than reported in its original paper. This is likely due to the absence of structured task tokens (e.g., “<|pitch|><|midi\_pitch|>piano”) in our prompt. Instead, CMI-bench relies on general natural language instructions. This highlights a critical gap in current audio LLMs: without clearly defined prompting schemas, their ability to interpret instructions can be fragile and fail to generalize. While different prompts for MusiLingo do not provide a significant difference on MusicCaps.

#### 5.1.5 Sequential Tasks Remain Challenging for All.

Tasks involving structured sequence-based outputs—such as melody extraction, instrument performance technique

		Qw2.	Qw.	Salm.	MusiL.	LTU	LTU-AS	MU-L.	auFla.	Gama	GamaI	Pengi	SOTA
GS-K	GES ↑	8.28	6.51	7.70	<b>9.50</b>	7.61	1.42	7.56	8.21	7.69	7.70	0.00	74.3 [49]
EMO	aR2 ↑	-0.75	-0.44	-0.51	-0.68	-1.14	-1.27	-0.03	-0.85	-1.08	-0.29	<b>0.00</b>	0.62 [50]
	vR2 ↑	-0.84	-0.78	0.0	-0.60	-1.13	-0.78	-0.12	-0.60	-1.30	<b>-1.19</b>	<b>0.00</b>	0.76 [51]
MTT	ROC ↑	66.78	66.00	59.07	63.39	65.75	65.83	68.32	68.68	<b>81.21</b>	78.32	66.75	92.0 [52]
	PR ↑	19.15	16.99	15.08	12.25	17.78	15.72	18.65	20.16	<b>34.26</b>	27.53	17.82	41.4 [50]
M-G	ROC ↑	64.44	<b>66.39</b>	57.71	57.48	52.22	57.14	57.36	62.83	52.50	62.49	58.23	88.0 [51]
	PR ↑	<b>9.23</b>	8.07	5.62	4.99	3.62	4.98	4.97	6.85	3.90	6.01	5.47	20.5 [51]
M-E	ROC ↑	<b>60.89</b>	59.06	50.69	53.07	51.41	52.02	54.40	55.80	51.97	58.84	53.88	78.6 [53]
	PR ↑	<b>7.85</b>	6.09	3.65	3.95	3.98	3.72	4.35	4.60	4.07	5.27	3.93	16.1 [53]
M-I	ROC ↑	<b>58.90</b>	56.95	48.78	55.63	55.34	53.02	50.81	56.99	51.15	55.16	56.09	78.8 [54]
	PR ↑	<b>12.41</b>	11.35	7.44	9.24	10.98	8.90	8.24	10.71	9.01	10.69	9.36	22.0 [51]
M-50	ROC ↑	64.64	63.00	53.46	<b>57.58</b>	<b>53.86</b>	54.11	54.88	60.96	52.01	60.68	57.22	84.3 [53]
	PR ↑	16.54	14.45	9.49	9.68	8.30	8.67	9.11	12.16	8.10	11.72	10.19	32.1 [53]
GTZ.	Acc. ↑	<b>72.07</b>	71.38	32.76	7.24	2.76	16.90	8.97	50.34	21.38	42.41	6.21	83.9 [55]
VS-T	Acc. ↑	14.91	15.18	<b>15.61</b>	1.23	7.11	0.53	4.56	11.32	7.72	7.89	0.00	76.9 [56]
NI	Acc. ↑	37.62	4.13	0.15	0.00	0.49	6.88	0.00	15.80	<b>58.37</b>	39.36	42.26	78.2 [57]
NP	Acc. ↑	1.51	0.37	0.00	0.00	0.73	0.05	0.00	0.73	0.20	0.00	<b>5.74</b>	89.2 [53]
SDD	BL. ↑	<b>23.40</b>	11.95	16.41	8.14	11.54	9.72	15.55	15.14	15.96	20.93	15.47	-
	ME. ↑	<b>23.21</b>	9.35	18.45	14.32	8.51	7.49	13.89	11.81	13.81	16.41	9.98	16.7 [58]
	RO. ↑	<b>28.47</b>	12.35	28.12	30.15	9.33	9.42	15.28	12.92	18.35	20.07	11.45	111.9 [58]
	BS. ↑	<b>87.44</b>	84.79	86.68	85.28	84.44	83.62	86.38	85.75	85.89	86.21	82.90	86.0 [58]
MC	BL. ↑	14.76	2.98	1.23	<b>21.50</b>	5.24	4.22	3.48	2.25	7.57	14.53	16.52	21.7 [4]
	ME. ↑	12.47	5.55	4.60	<b>22.49</b>	8.55	7.01	8.01	5.97	10.07	10.98	14.77	22.4 [58]
	RO. ↑	12.35	6.68	6.26	<b>30.29</b>	9.39	7.51	8.58	6.94	11.38	12.46	12.64	30.8 [4]
	BS. ↑	84.38	82.37	82.98	<b>85.75</b>	83.84	83.59	83.00	83.43	84.30	84.57	83.22	87.8 [58]
DS	WE. ↓	793.0	<b>115.7</b>	816.1	2019	235.5	191.7	191.9	275.7	225.4	152.6	343.2	12.99 [59]
	CE. ↓	818.6	<b>96.2</b>	760.00	2311	210.8	185.5	168.3	262.6	201.3	165.2	368.0	-
G-B	FM. ↑	7.50	<b>23.69</b>	11.49	0.04	0.10	0.00	0.71	3.96	0.00	1.49	0.00	88.3 [56]
G-D	FM. ↑	5.97	<b>10.21</b>	8.62	0.18	0.86	0.00	0.17	3.06	0.05	0.54	0.00	54.1 [60]
BR-B	FM. ↑	7.12	<b>21.96</b>	14.97	0.01	0.15	0.00	0.22	4.69	0.02	1.02	0.00	96.8 [61]
BR-D	FM. ↑	5.69	<b>10.68</b>	9.40	0.06	2.29	0.00	0.15	3.47	0.14	0.68	0.00	94.1 [61]
MDB	Acc. ↑	<b>5.06</b>	0.08	0.00	0.00	0.00	0.00	0.00	0.01	0.66	0.00	0.00	72.3 [62]
GZ	maF1 ↑	<b>3.18</b>	1.66	0.03	0.00	0.04	0.00	0.00	0.00	0.00	0.00	0.00	90.0 [63]
	miF1 ↑	<b>0.89</b>	0.44	0.01	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	80.4 [63]

**Table 3.** Performance of 11 open-source audio-text LLMs on CMI-Bench. Models: Qwen2-Audio (Qw2.), Qwen-Audio (Qw.), SALMONN-Audio (Salm.), MusiLingo (MusiL.), LTU, LTU-AS, MU-LLaMA (MU-L.), Audio-Flamingo (auFla.), GAMA, GAMA-IT (GamaI), Pengi. Tasks include key detection (GS-K), emotion regression (EMO), tagging (MTT, M-50), genre (M-G, GTZ.), emotion/instrument tagging (M-E, M-I), captioning (SDD, MC), lyrics transcription (DS), beat/downbeat tracking (G-B/G-D, BR-B/BR-D), melody (MDB), and Guzheng techniques (GZ). Metrics: GES, R<sup>2</sup>, ROC-AUC, PR-AUC, Accuracy, BLEU (BL.), METEOR (ME.), ROUGE (RO.), BERTScore (BS.), WER/ CER, FM(F-Measure), Macro-F1 (maF1), Micro-F1 (miF1). Best scores are in bold.

detection, and (down)beat tracking—are poorly handled by all evaluated models. Even Qwen-Audio, which shows relatively strong performance in genre and beat tracking, falls far short of MIR baselines, sometimes copying the input examples. We hypothesize two key reasons: For one thing, the diversity and ambiguity in how sequence tasks are phrased (e.g., timestamps, tuple formats) reduces consistency in model outputs. For another, many models have only limited exposure to audio tasks with dense temporal supervision. If pretraining data includes timestamped output and matched decoding formats, performance may improve.

#### 5.1.6 Emotion Regression Fails for All Models.

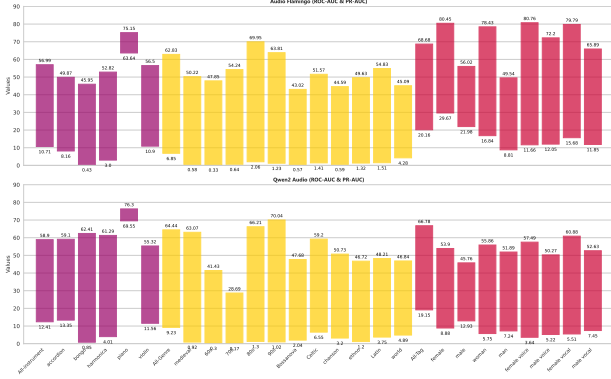
Despite clear instructions, carefully designed scales (1–9), and contextual music descriptions, and few-shot examples, all models fail to provide usable predictions for

arousal and valence. In fact, model outputs often cluster around meaningless values, sometimes performing worse than simply predicting the mean. Our post-processing rules convert empty or invalid outputs to dataset means, which often lead to better R<sup>2</sup> scores than the models themselves—highlighting the severe limitations in mapping continuous perceptual attributes from music using current audio-text LLMs.

These findings emphasize the gap between current SFT multimodal LLMs and traditional task-specific MIR systems. While open-source audio LLMs show promise in isolated tasks with aligned training data, substantial challenges remain in terms of generalization, structured output generation, and adaptation to real-world settings.

## 5.2 Culture and Gender Bias

We further analyze the performance of two top-performing models—Qwen2-Audio and Audio-Flamingo—on fine-grained instrument, genre, and music tag categories. While both models show competitive results overall, our breakdown highlights notable performance disparities across instrument types, cultural genres, and voice-related tags.



**Figure 2.** Fine-grained evaluation of Qwen2-Audio and Audio-Flamingo on instrument (purple), genre (yellow), and vocal (red) tag classification. The upper extremity represents the ROC-AUC value, and the lower is PR-AUC.

### 5.2.1 Instrument Bias on MTG-Instrument

Both models achieve high scores on piano, reflecting the strong representation of piano in most training datasets. Western instruments such as violin and accordion perform close to the average, suggesting moderate robustness across common musical timbres. However, performance drops significantly on bongo and harmonica — commonly associated with world music. These results point to a persistent bias toward Western instruments and limited generalization to underrepresented timbres in current pre-training corpora.

### 5.2.2 Cultural Genre Imbalance on MTG-Genre

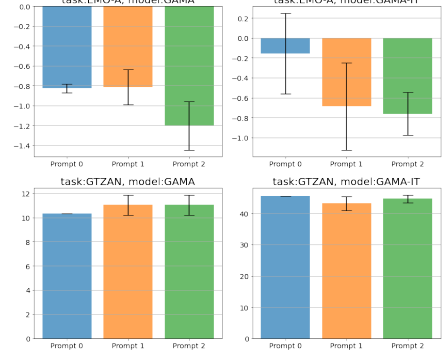
Genre classification results similarly reveal systematic disparities. Both models show relatively strong performance on mainstream Western pop genres (e.g., 80s, 90s), while genres associated with world music (e.g., Bossanova, Celtic, Chanson, Ethno, Latin) and music traditions (e.g., Medieval) consistently fall below average. For example, Audio-Flamingo’s performance on Bossanova and Chanson drops severely. Qwen2-Audio performs slightly better on some long-tail genres, but still shows considerable degradation. These highlight a lack of cultural and historical diversity in the data used for instruction tuning and model pretraining.

### 5.2.3 Voice Tag Differences on MTT

A detailed comparison on vocal tags reveals an interesting divergence. Audio-Flamingo is consistently better at identifying \*female\* voices than male voices, indicating a possible gender-related acoustic or annotation bias. In contrast, Qwen2-Audio achieves higher ROC-AUC for \*female\* tags but lower PR-AUC, suggesting that while the

model ranks positive examples correctly, its absolute predictions remain sparse or overconfident. This mismatch implies that Qwen2-Audio is sensitive to class ranking but may lack calibration in estimating tag presence probabilities, an issue worth investigating for fairness and reliability in music model deployment.

## 5.3 Ablation Study on Different Prompts and Trials



**Figure 3.** Ablation Study on Prompt Sensitivity for Genre Classification and Arousal Regression

We conduct an ablation study on prompt design using GAMA and GAMA-IT models across two representative tasks: GTZAN genre classification and EMO arousal regression. Variant Prompts 1 and 2 are evaluated over three runs, and the bars report mean performance with standard variant as error bars. GTZAN results (bottom row) are relatively stable across prompts and have small variance for each prompt in multi-trials, indicating that most genre-related instructions are consistently followed. The low variance suggests robustness to prompt changes. In contrast, EMO-A results (top row) show relative sensitivity to prompt variation, particularly under the GAMA-IT model. This instability stems from a higher rate of invalid or non-responsible generations, which are scored as mean values during evaluation. Consequently, differences in prompt phrasing might lead to large deviations, especially when valid predictions diverge significantly from the mean score.

## 6. CONCLUSION

We introduce CMI-Bench, a comprehensive benchmark for evaluating audio-text LLMs across diverse MIR tasks. Our results highlight a significant performance gap between LLMs and supervised MIR systems, with best models like Qwen2-Audio and GAMA also struggling with generalization. Sequence-based tasks, such as melody extraction and beat tracking, pose particular challenges, likely due to limited timestamped pretraining and prompt sensitivity. Fine-grained analysis also reveals cultural and gender biases tied to training data imbalances. By offering a standardized evaluation framework and toolkit, CMI-Bench bridges NLP and MIR research, providing a foundation for future advancements. Progress will hinge on improved pretraining, sequential output handling, and bias mitigation, and we hope this work spurs collaboration toward more capable music-aware LLMs.

## A. ACKNOWLEDGMENTS

Yinghao Ma is a research student at the UKRI Centre for Doctoral Training in Artificial Intelligence and Music, supported by UK Research and Innovation [grant number EP/S022694/1].

Siyu Li is a research student at the Computational Linguistics Lab at Queen Mary University of London, funded by the QMUL-CSC PhD scholarships.

Yinghao Ma would also like to express heartfelt gratitude to the Student Philharmonic Chinese Orchestra at the Chinese Music Institute, Peking University (abbreviated as CMI, unrelated to the paper title). We warmly congratulate the orchestra on its 20th anniversary.

## B. ETHICS STATEMENT

CMI-Bench repurposes existing publicly available datasets in the MIR domain by reformatting their annotations into instruction-following formats. No new human annotations were collected, and no human participants were involved in the creation of this benchmark. All data used in the project are licensed under terms that permit non-commercial research use. In compliance with these terms, we license CMI-Bench under a Creative Commons Attribution-NonCommercial-ShareAlike (CC BY-NC-SA) license. To promote long-term accessibility, we host the audio test set on Hugging Face with clear usage restrictions for non-commercial purposes.

The dataset primarily consists of Western, English-language popular music, with limited inclusion of instrumental tracks and non-English songs. Transcription tasks are restricted to English lyrics, and world music instruments besides Guzheng are underrepresented. We acknowledge this cultural and linguistic skew and encourage future extensions to improve global diversity and representation.

This work involves no safety, security, or environmental risks. The benchmark does not require high-compute model training or deployment of potentially harmful generative models. We release CMI-Bench and its evaluation toolkit to foster responsible and reproducible research in audio-language modeling.

## C. APPENDIX

Due to the limitation of the ISMIR proceeding, please refer to our arxiv version for more information on instruction examples, error case analysis and more discussion.

## D. REFERENCES

- [1] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong *et al.*, “A survey of large language models,” *arXiv preprint arXiv:2303.18223*, vol. 1, no. 2, 2023.
- [2] Y. Ma, A. Øland, A. Ragni, B. M. Del Sette, C. Saitis, C. Donahue, C. Lin, C. Plachouras, E. Benetos, E. Shatri *et al.*, “Foundation models for music: A survey,” *arXiv preprint arXiv:2408.14340*, 2024.
- [3] 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)*. IEEE, 2024, pp. 286–290.
- [4] Z. Deng, Y. Ma, Y. Liu, R. Guo, G. Zhang, W. Chen, W. Huang, and E. Benetos, “Musilingo: Bridging music and text with pre-trained language models for music captioning and query response,” in *NAACL-HLT (Findings)*, 2024.
- [5] Y. Chu, J. Xu, X. Zhou, Q. Yang, S. Zhang, Z. Yan, C. Zhou, and J. Zhou, “Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models,” *arXiv preprint arXiv:2311.07919*, 2023.
- [6] C. Tang, W. Yu, G. Sun, X. Chen, T. Tan, W. Li, L. Lu, M. Zejun, and C. Zhang, “Salmonn: Towards generic hearing abilities for large language models,” in *The Twelfth International Conference on Learning Representations*.
- [7] Q. Yang, J. Xu, W. Liu, Y. Chu, Z. Jiang, X. Zhou, Y. Leng, Y. Lv, Z. Zhao, C. Zhou *et al.*, “Air-bench: Benchmarking large audio-language models via generative comprehension,” in *ACL (1)*, 2024.
- [8] B. Weck, I. Manco, E. Benetos, E. Quinton, G. Fazekas, and D. Bogdanov, “Muchomusic: Evaluating music understanding in multimodal audio-language models,” in *Proceedings of the 25th International Society for Music Information Retrieval Conference, ISMIR 2024, San Francisco, California, USA and Online, November 10-14, 2024*, B. Kaneshiro, G. J. Mysore, O. Nieto, C. Donahue, C. A. Huang, J. H. Lee, B. McFee, and M. C. McCallum, Eds., 2024, pp. 825–833. [Online]. Available: <https://doi.org/10.5281/zenodo.14877459>
- [9] Y. Li, G. Zhang, Y. Ma, R. Yuan, K. Zhu, H. Guo, Y. Liang, J. Liu, Z. Wang, J. Yang *et al.*, “Omnibench: Towards the future of universal omni-language models,” *arXiv preprint arXiv:2409.15272*, 2024.
- [10] S. Sakshi, U. Tyagi, S. Kumar, A. Seth, R. Selvakumar, O. Nieto, R. Duraiswami, S. Ghosh, and D. Manocha, “Mmau: A massive multi-task audio understanding and reasoning benchmark,” in *The Thirteenth International Conference on Learning Representations*.
- [11] Y. Wang, S. Mishra, P. Alipoormolabashi, Y. Kordi, A. Mirzaei, A. Naik, A. Ashok, A. S. Dhanasekaran, A. Arunkumar, D. Stap *et al.*, “Supernaturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks,” in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 2022, pp. 5085–5109.
- [12] Y. Wang, Y. Kordi, S. Mishra, A. Liu, N. A. Smith, D. Khashabi, and H. Hajishirzi, “Self-instruct: Aligning

- language models with self-generated instructions,” in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 2023.
- [13] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray *et al.*, “Training language models to follow instructions with human feedback,” *Advances in neural information processing systems*, vol. 35, pp. 27 730–27 744, 2022.
- [14] L. Xue, Z. Zhou, J. Pan, Z. Li, S. Fan, Y. Ma, S. Cheng, D. Yang, H. Guo, Y. Xiao *et al.*, “Audio-flan: A preliminary release,” *arXiv preprint arXiv:2502.16584*, 2025.
- [15] J. Li, L. Yang, M. Tang, C. Chenchong, Z. Li, P. Wang, and H. Zhao, “The music maestro or the musically challenged, a massive music evaluation benchmark for large language models,” in *Findings of the Association for Computational Linguistics ACL 2024*, 2024, pp. 3246–3257.
- [16] B. Wang, X. Zou, G. Lin, S. Sun, Z. Liu, W. Zhang, Z. Liu, A. Aw, and N. F. Chen, “Audiobench: A universal benchmark for audio large language models,” *arXiv preprint arXiv:2406.16020*, 2024.
- [17] J. Melechovsky, Z. Guo, D. Ghosal, N. Majumder, D. Herremans, and S. Poria, “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)*, 2024, pp. 8286–8309.
- [18] Z. Wang, S. Li, T. Zhang, Q. Wang, P. Yu, J. Luo, Y. Liu, M. Xi, and K. Zhang, “Muchin: a chinese colloquial description benchmark for evaluating language models in the field of music,” in *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 2024, pp. 7771–7779.
- [19] Y. Gong, H. Luo, A. H. Liu, L. Karlinsky, and J. Glass, “Listen, think, and understand,” in *International Conference on Learning Representations*, 2024.
- [20] Y. Gong, A. H. Liu, H. Luo, L. Karlinsky, and J. Glass, “Joint audio and speech understanding,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [21] A. Radford, J. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision (arxiv: 2212.04356). arxiv,” 2022.
- [22] J. P. Gardner, S. Durand, D. Stoller, and R. M. Bittner, “Llark: A multimodal instruction-following language model for music,” in *International Conference on Machine Learning*. PMLR, 2024, pp. 15 037–15 082.
- [23] S. Deshmukh, B. Elizalde, R. Singh, and H. Wang, “Pengi: An audio language model for audio tasks,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 18 090–18 108, 2023.
- [24] S. Ghosh, S. Kumar, A. Seth, C. K. R. Evuru, U. Tyagi, S. Singh, O. Nieto, R. Duraiswami, and D. Manocha, “Gama: A large audio-language model with advanced audio understanding and complex reasoning abilities,” *CoRR*, 2024.
- [25] Y. Chu, J. Xu, Q. Yang, H. Wei, X. Wei, Z. Guo, Y. Leng, Y. Lv, J. He, J. Lin *et al.*, “Qwen2-audio technical report,” *arXiv preprint arXiv:2407.10759*, 2024.
- [26] Z. Kong, A. Goel, R. Badlani, W. Ping, R. Valle, and B. Catanzaro, “Audio flamingo: a novel audio language model with few-shot learning and dialogue abilities,” in *Proceedings of the 41st International Conference on Machine Learning*, 2024, pp. 25 125–25 148.
- [27] S. Ghosh, Z. Kong, S. Kumar, S. Sakshi, J. Kim, W. Ping, R. Valle, D. Manocha, and B. Catanzaro, “Audio flamingo 2: An audio-language model with long-audio understanding and expert reasoning abilities,” *arXiv preprint arXiv:2503.03983*, 2025.
- [28] C. Raffel, B. McFee, E. J. Humphrey, J. Salamon, O. Nieto, D. Liang, D. P. Ellis, and C. C. Raffel, “Mir\_eval: A transparent implementation of common mir metrics,” in *ISMIR*, 2014, pp. 367–372.
- [29] P. Knees, Á. Faraldo Pérez, H. Boyer, R. Vogl, S. Böck, F. Hörschläger, M. Le Goff *et al.*, “Two data sets for tempo estimation and key detection in electronic dance music annotated from user corrections,” in *Proceedings of the 16th International Society for Music Information Retrieval Conference (ISMIR); 2015 Oct 26-30; Málaga, Spain.[Málaga]: International Society for Music Information Retrieval, 2015. p. 364-70. International Society for Music Information Retrieval (ISMIR)*, 2015.
- [30] M. Soleymani, M. N. Caro, E. M. Schmidt, C.-Y. Sha, and Y.-H. Yang, “1000 songs for emotional analysis of music,” in *Proceedings of the 2nd ACM international workshop on Crowdsourcing for multimedia*, 2013, pp. 1–6.
- [31] E. Law, K. West, M. I. Mandel, M. Bay, and J. S. Downie, “Evaluation of algorithms using games: The case of music tagging,” in *ISMIR*. Citeseer, 2009, pp. 387–392.
- [32] D. Bogdanov, M. Won, P. Tovstogan, A. Porter, and X. Serra, “The mtg-jamendo dataset for automatic music tagging,” in *International Conference on Machine Learning*. ICML, 2019.
- [33] J. Engel, C. Resnick, A. Roberts, S. Dieleman, M. Norouzi, D. Eck, and K. Simonyan, “Neural audio synthesis of musical notes with wavenet autoencoders,” in *International Conference on Machine Learning*. PMLR, 2017, pp. 1068–1077.

- [34] G. Tzanetakis and P. Cook, “Musical genre classification of audio signals,” *IEEE Transactions on speech and audio processing*, vol. 10, no. 5, pp. 293–302, 2002.
- [35] J. Wilkins, P. Seetharaman, A. Wahl, and B. Pardo, “Vocalset: A singing voice dataset,” in *ISMIR*, 2018, pp. 468–474.
- [36] I. Manco, B. Weck, S. Doh, M. Won, Y. Zhang, D. Bogdanov, Y. Wu, K. Chen, P. Tovstogan, E. Benetos *et al.*, “The song describer dataset: a corpus of audio captions for music-and-language evaluation,” *arXiv preprint arXiv:2311.10057*, 2023.
- [37] A. Agostinelli, T. I. Denk, Z. Borsos, J. Engel, M. Verzetzi, A. Caillon, Q. Huang, A. Jansen, A. Roberts, M. Tagliasacchi *et al.*, “Musiclm: Generating music from text,” *arXiv preprint arXiv:2301.11325*, 2023.
- [38] G. Roa Dabike and J. Barker, “Automatic lyric transcription from karaoke vocal tracks: Resources and a baseline system,” in *Proceedings of the 20th Annual Conference of the International Speech Communication Association (INTERSPEECH 2019)*, 2019.
- [39] F. Gouyon, A. Klapuri, S. Dixon, M. Alonso, G. Tzanetakis, C. Uhle, and P. Cano, “An experimental comparison of audio tempo induction algorithms,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 5, pp. 1832–1844, 2006.
- [40] F. Krebs, S. Böck, and G. Widmer, “Rhythmic pattern modeling for beat and downbeat tracking in musical audio,” in *Ismir*, 2013, pp. 227–232.
- [41] R. Bittner, J. Salamon, M. Tierney, M. Mauch, C. Cannam, and J. Bello, “Medleydb: A multitrack dataset for annotation-intensive mir research,” 10 2014.
- [42] D. Li, M. Che, W. Meng, Y. Wu, Y. Yu, F. Xia, and W. Li, “Frame-level multi-label playing technique detection using multi-scale network and self-attention mechanism,” in *IEEE International Conference on Acoustics, Speech and Signal Processing ICASSP 2023, Rhodes Island, Greece, June 4-10, 2023*. IEEE, 2023, pp. 1–5.
- [43] S. Xiao, Z. Liu, P. Zhang, N. Muennighoff, D. Lian, and J.-Y. Nie, “C-pack: Packed resources for general chinese embeddings,” in *Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval*, 2024, pp. 641–649.
- [44] K. Papineni, S. Roukos, T. Ward, and W. jing Zhu, “Bleu: a method for automatic evaluation of machine translation,” in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002, pp. 311–318.
- [45] C.-Y. Lin and F. J. Och, “ORANGE: a method for evaluating automatic evaluation metrics for machine translation,” in *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics*. Geneva, Switzerland: COLING, aug 23–aug 27 2004, pp. 501–507. [Online]. Available: <https://www.aclweb.org/anthology/C04-1072>
- [46] S. Banerjee and A. Lavie, “Meteor: An automatic metric for mt evaluation with improved correlation with human judgments,” in *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005, pp. 65–72.
- [47] C.-Y. Lin, “ROUGE: A package for automatic evaluation of summaries,” in *Text Summarization Branches Out*. Barcelona, Spain: Association for Computational Linguistics, Jul. 2004, pp. 74–81. [Online]. Available: <https://www.aclweb.org/anthology/W04-1013>
- [48] 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. [Online]. Available: <https://openreview.net/forum?id=SkeHuCVFDr>
- [49] F. Korzeniowski and G. Widmer, “End-to-end musical key estimation using a convolutional neural network,” in *2017 25th European Signal Processing Conference (EUSIPCO)*. IEEE, 2017, pp. 966–970.
- [50] R. Castellon, C. Donahue, and P. Liang, “Codified audio language modeling learns useful representations for music information retrieval,” *arXiv preprint arXiv:2107.05677*, 2021.
- [51] R. Yuan, Y. Ma, Y. Li, G. Zhang, X. Chen, H. Yin, Y. Liu, J. Huang, Z. Tian, B. Deng *et al.*, “Marble: Music audio representation benchmark for universal evaluation,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 39 626–39 647, 2023.
- [52] Q. Huang, A. Jansen, J. Lee, R. Ganti, J. Y. Li, and D. P. Ellis, “Mulan: A joint embedding of music audio and natural language,” *arXiv preprint arXiv:2208.12415*, 2022.
- [53] M. C. McCallum, F. Korzeniowski, S. Oramas, F. Gouyon, and A. F. Ehmann, “Supervised and unsupervised learning of audio representations for music understanding,” *Ismir 2022 Hybrid Conference*, 2022.
- [54] P. Alonso-Jiménez, X. Serra, and D. Bogdanov, “Music representation learning based on editorial metadata from discogs,” 2022.
- [55] D. Niizumi, D. Takeuchi, Y. Ohishi, N. Harada, and K. Kashino, “Masked modeling duo: Learning representations by encouraging both networks to model the input,” in *ICASSP 2023-2023 IEEE International Conference On Acoustics, Speech And Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.

- [56] Y. Li, R. Yuan, G. Zhang, Y. Ma, X. Chen, H. Yin, C. Lin, A. Ragni, E. Benetos, N. Gyenge, R. Dannenberg, R. Liu, W. Chen, G. Xia, Y. Shi, W. Huang, Y. Guo, and J. Fu, “Mert: Acoustic music understanding model with large-scale self-supervised training,” 2023.
- [57] L. Wang, P. Luc, Y. Wu, A. Recasens, L. Smaira, A. Brock, A. Jaegle, J.-B. Alayrac, S. Dieleman, J. Carreira *et al.*, “Towards learning universal audio representations,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 4593–4597.
- [58] S. Doh, K. Choi, J. Lee, and J. Nam, “Lp-musiccaps: Llm-based pseudo music captioning,” in *Ismir 2023 Hybrid Conference*, 2023.
- [59] L. Ou, X. Gu, and Y. Wang, “Transfer learning of wav2vec 2.0 for automatic lyric transcription,” in *Ismir 2022 Hybrid Conference*, 2022.
- [60] M. Heydari, F. Cwitkowitz, and Z. Duan, “Beatnet: CRNN and particle filtering for online joint beat, downbeat and meter tracking,” in *Proceedings of the 22nd International Society for Music Information Retrieval Conference, ISMIR 2021, Online, November 7-12, 2021*, J. H. Lee, A. Lerch, Z. Duan, J. Nam, P. Rao, P. van Kranenburg, and A. Srinivasamurthy, Eds., 2021, pp. 270–277. [Online]. Available: <https://archives.ismir.net/ismir2021/paper/000033.pdf>
- [61] J. Zhao, G. Xia, and Y. Wang, “Beat transformer: Demixed beat and downbeat tracking with dilated self-attention,” in *Proceedings of the 23rd International Society for Music Information Retrieval Conference, ISMIR 2022, Bengaluru, India, December 4-8, 2022*, P. Rao, H. A. Murthy, A. Srinivasamurthy, R. M. Bittner, R. C. Repetto, M. Goto, X. Serra, and M. Miron, Eds., 2022, pp. 169–177. [Online]. Available: <https://archives.ismir.net/ismir2022/paper/000019.pdf>
- [62] S. Yu, X. He, K. Chen, and Y. Yu, “Hkdsme: Heterogeneous knowledge distillation for semi-supervised singing melody extraction using harmonic supervision,” in *ACM Multimedia 2024*.
- [63] D. Li, Y. Ma, W. Wei, Q. Kong, Y. Wu, M. Che, F. Xia, E. Benetos, and W. Li, “Mertech: Instrument playing technique detection using self-supervised pretrained model with multi-task finetuning,” in *ICASSP*, 2024.
- [64] L. Zhuo, R. Yuan, J. Pan, Y. Ma, Y. Li, G. Zhang, S. Liu, R. Dannenberg, J. Fu, C. Lin *et al.*, “Lyricwhiz: Robust multilingual lyrics transcription by whispering to chatgpt.” International Society for Music Information Retrieval Conference (ISMIR), 2023.