

SIFT-50M: A Large-Scale Multilingual Dataset for Speech Instruction Fine-Tuning

Prabhat Pandey^{*1} Rupak Vignesh Swaminathan^{*1} K V Vijay Girish^{*1}
Arunasish Sen¹ Jian Xie^{†2} Grant P. Strimel¹ Andreas Schwarz¹

¹Amazon AGI ²Apple Inc.

Correspondence: {panprabh, swarupak, kvvijayg}@amazon.com

Abstract

We introduce SIFT (Speech Instruction Fine-Tuning), a 50M-example dataset designed for instruction fine-tuning and pre-training of speech-text large language models (LLMs). SIFT-50M is built from publicly available speech corpora, which collectively contain 14K hours of speech, and leverages LLMs along with off-the-shelf expert models. The dataset spans five languages, encompassing a diverse range of speech understanding as well as controllable speech generation instructions. Using SIFT-50M, we train SIFT-LLM, which outperforms existing speech-text LLMs on instruction-following benchmarks while achieving competitive performance on foundational speech tasks. To support further research, we also introduce EvalSIFT, a benchmark dataset specifically designed to evaluate the instruction-following capabilities of speech-text LLMs.

1 Introduction

Recent years have witnessed significant advancements in integrating speech and audio capabilities into large language models (LLMs). A common approach involves coupling an audio encoder with an LLM by projecting audio embeddings into the model’s input space (Tang et al.; Yang et al., 2024a; Gong et al., 2023b; Das et al., 2024). This integration preserves rich prosodic and acoustic information beyond textual transcriptions, enabling speech-text LLMs to perform a broad range of speech understanding tasks. Another research direction focuses on extending LLM vocabulary to generate discrete audio tokens for speech synthesis (Rubenstein et al., 2023; Du et al., 2023).

Speech-text LLMs are typically trained on existing task-specific speech corpora, which are primarily designed for automatic speech recognition (ASR). These datasets predominantly consist of

paired audio and task-specific labels rather than natural language instructions, limiting their utility for instruction-following training. This lack of diverse, large-scale instruction datasets poses challenges in generalizing to broader speech understanding tasks. To address this gap, we introduce SIFT-50M¹ (Speech Instruction Fine-Tuning), a large-scale multilingual dataset covering five languages for instruction-based speech model training. SIFT-50M augments existing speech datasets with instruction-based question-answer (QA) pairs for speech understanding and includes approximately 5M examples for controllable speech generation.

In summary, our main contributions are as follows: (1) We release SIFT-50M, a 50M-example multilingual instruction dataset which, to the best of our knowledge, is the largest publicly available instruction dataset for speech understanding tasks. (2) We introduce EvalSIFT, a benchmark dataset designed for systematically evaluating speech-text LLMs across multiple dimensions. (3) We develop SIFT-LLM, a speech-text LLM trained on SIFT-50M, achieving state-of-the-art performance on instruction-following benchmarks compared to publicly available models of similar size. (4) We conduct controllable speech generation experiments using SIFT-50M, demonstrating the model’s ability to generate speech in instructed styles.

2 Related Work

Speech-text LLMs: Several speech-text LLMs have been introduced recently. Early works focused on specific tasks such as ASR (Fathullah et al., 2024; Yu et al., 2024) and second-pass ASR rescoring (Li et al., 2023b). SALMONN (Tang et al.) integrates Whisper (Radford et al., 2023) and BEATs (Chen et al., 2023) encoders with a pre-trained LLM using a window-level Q-former

^{*}Equal contribution.

[†]Work done while at Amazon.

¹The dataset can be accessed at <https://huggingface.co/datasets/amazon-agi/SIFT-50M>

(Li et al., 2023a). Qwen-Audio (Chu et al., 2023) and Qwen2-Audio (Chu et al., 2024) are trained on over 30 speech, audio, and music tasks. WavLLM (Hu et al., 2024) incorporates two speech encoders, Whisper-large-v2 and WavLM-base (Chen et al., 2022), while LTU-AS (Gong et al., 2023b) integrates Whisper-large with Llama-7B (Touvron et al., 2023) via a time and layerwise transformer (Gong et al., 2023a). These models process continuous audio representations, projecting encoded audio into the LLM’s input space.

To support both understanding and generation, AudioPaLM (Rubenstein et al., 2023) and LauraGPT (Du et al., 2023) extend LLM vocabulary with discrete audio tokens. AudioPaLM discretizes embeddings from USM (Zhang et al., 2023) and w2v-BERT (Chung et al., 2021) using k-means clustering, extending LLM vocabulary with discrete audio tokens. LauraGPT adopts a hybrid approach, utilizing a conformer-based encoder for continuous representation of input speech while generating discrete EnCodec (Défossez et al., 2022) codes for output speech.

Instruction Fine-Tuning Datasets: Most prior speech-text LLMs rely on task-specific datasets with instructions generated by LLMs. To improve task diversity, WavLLM (Hu et al., 2024) created a 3K-hour speech question-answering dataset using GPT-4 based on transcriptions from Librispeech (Panayotov et al., 2015), the AMI Meeting Corpus (Carletta et al., 2005), Fisher (Cieri et al.), and Switchboard (Godfrey et al., 1992). Additionally, they developed a multi-task dataset by merging independent training examples into single multi-task instructions.

Another relevant dataset is OpenASQA (Gong et al., 2023b), a publicly available instruction fine-tuning dataset with 2.7M examples for speech and audio tasks. Its construction methodology, closely related to ours, leverages datasets such as IEMO-CAP (Busso et al., 2008), LibriTTS (Zen et al., 2019), VoxCeleb2 (Chung et al., 2018), and MOSEI (Zadeh et al., 2018). OpenASQA augments existing annotations (e.g., transcription, emotion, gender) with metadata such as speed, pitch, and energy, and then uses this enriched metadata to prompt GPT-3.5-Turbo to generate QA pairs. In contrast, SIFT-50M is an order of magnitude larger than OpenASQA, covering a broader range of speech dimensions (e.g., accent, age, gender, pitch, speaking rate, word alignment, room characteristics, noise level, distortion) while also supporting instructions

in four additional languages beyond English.

Controllable Speech Generation: Prompt-based speech generation, enabling controllable synthesis of acoustic characteristics such as gender, speaking rate, pitch, and loudness, has become a prominent area of investigation (Guo et al., 2023; Shimizu et al., 2024). (Kawamura et al., 2024) annotated the LibriTTS corpus (Zen et al., 2019) with textual prompts to augment these specific speech properties. While (Lyth and King, 2024) detailed a method for annotating a substantial 45K-hour corpus for high-fidelity speech generation, these prompts are not publicly available, limiting reproducibility and benchmarking. Conversely, SIFT-50M offers speech generation instructions across five languages, enabling control over a broader spectrum of parameters, including gender, accent, age, speaking rate, pitch, and intensity.

3 SIFT Dataset Construction

To construct SIFT-50M, we utilize three publicly available multilingual speech data sources: Multilingual Librispeech (Pratap et al., 2020), Common Voice (Ardila et al., 2019), and VCTK (Yamagishi et al., 2019). These sources provide a diverse range of speakers, speech styles, languages, and accents. We augment the existing metadata with additional features extracted from the speech data and generate natural language instructions from the enriched metadata with the help of an LLM. Finally, we employ a feedback loop involving verification and iterative prompt refinement to improve the generation process.

3.1 Metadata Extraction

We augment the source speech datasets—which typically include content-level characteristics, with additional acoustic-level features such as pitch, speaking rate, intensity, duration, reverberation, and noise level. We extract pitch using PYYAPT (Kasi and Zahorian, 2002) and intensity (in decibels) using Parselmouth (Jadoul et al., 2018). Phonemes are derived from transcripts via a grapheme-to-phoneme (G2P) model (Park, 2019), while noise, speech quality, and reverberation features are computed following the method described in (Lacombe, 2024). Content-level attributes (e.g., transcription, gender, age, and accent) are provided in the source datasets, although accent and age information is available only for a random subset of utterances in the Common Voice and VCTK

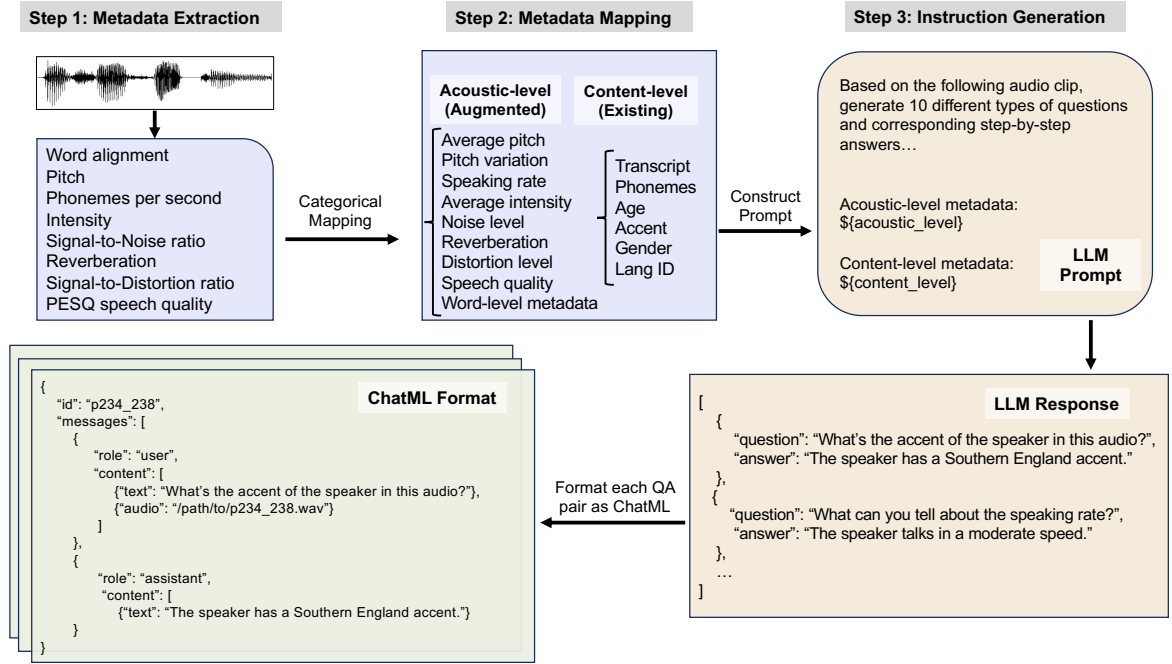


Figure 1: Block diagram showing the stages of SIFT-50M dataset construction. For non-English data generation, we substitute the metadata mapping with the respective language and prompt the LLM to generate responses in that language.

datasets. When a gender label is missing, we employ a gender classification model (Huh, 2024) to fill in the gap.

We then convert numerical utterance-level statistics into categorical ranges for pitch, intensity, age, reverberation, noise level, and speech quality based on the distribution of each variable within a dataset. For example, pitch is binned separately for male and female speakers due to their different typical ranges. Specifically, for a male speaker, an *average pitch* metadata is mapped to *slightly low pitch* when the median pitch falls within $125 < f_{\text{median}} < 138$ Hz, while for a female speaker the corresponding range is $166 < f_{\text{median}} < 182$ Hz. Age is similarly grouped into broad categories (e.g., twenties, thirties, . . . , nineties) to both de-identify and generalize the attribute. For a detailed list of categorical values, please refer to Appendix Section B.1.

To generate word-level metadata, we extract time alignments for each word using our in-house hybrid HMM-LSTM acoustic model (Sak, 2014) for forced alignment. The start and end times of each word are then used to derive word-level pitch and intensity statistics, which are subsequently converted to categorical values. For non-English languages (French, German, Spanish, and Italian), we map the metadata characteristics and their corresponding categorical values to the target language

using a phrase-level translation dictionary.

3.2 Natural Language Instruction Generation

Our instructions are organized into three main categories: closed-ended, open-ended, and controllable generation. To generate instruction-response pairs, we feed categorical metadata to the LLM and prompt it to produce up to 10 QA pairs per speech utterance, as shown in Figure 1. For multilingual instruction generation, the LLM is prompted to generate instructions in the specified language using the corresponding mapped metadata. Table 1 provides examples of QA pairs for different categories in the SIFT-50M dataset. High-level statistics of SIFT-50M are presented in Table 2.

Closed-ended Instructions: These instructions are subdivided into four sub-categories: (1) Acoustic-Level, (2) Content-Level, (3) Word-Align, and (4) Comparison types. The first three sub-categories, generated using Mixtral 8x7B (Jiang et al., 2024), are based on the respective metadata types. As reported in LTU (Gong et al., 2023c), closed-ended instructions are crucial for guiding the model to understand speech and follow instructions without hallucinating. Unlike LTU’s OpenASQA data, which relies on pre-processed questions and strict answer templates, our approach leverages careful prompt engineering to allow the

Closed-ended: Acoustic-Level
Question: How would you characterize the speaking rate in this audio?
Answer: The speaking rate in this audio is slightly slow.
Closed-ended: Word-Align
Question: What is the average pitch for the word ‘bent’ and how does it compare to the average pitch of the surrounding words?
Answer: The average pitch for the word ‘bent’ is quite low. The previous word ‘upon’ has a similar average pitch, while the following word ‘upon’ has a slightly lower pitch. Overall, there is not a significant change in pitch for these words.
Closed-ended: Comparison
Question: Which audio has a speaker with a Scottish accent?
Answer: The second audio has a speaker with a Scottish accent, whereas the first audio has a speaker with an English accent.
Open-ended
Question: Analyze the phonetic structure of the audio and discuss how the phonemes contribute to the overall perception of the speaker’s accent and age category.
Answer: The phonetic structure, including phonemes like ‘IH1’, ‘T’, ‘S’, ‘IY1’, ‘M’, ‘’, ‘AH0’, ‘W’, ‘EY1’, ‘S’, ‘T’, ‘AH1’, ‘V’, ‘T’, ‘AE1’, ‘L’, ‘AH0’, ‘N’, ‘T’, ‘S’, reflects an Irish accent characterized by distinct vowel sounds and consonant clusters. The precise articulation and the youthful intonation patterns suggest the speaker is in their twenties.
Controllable Generation
Question: Please synthesize a speech sample with a very low pitch and quite monotone pitch variation, spoken quite slowly and at a moderate volume. The speaker should be an adult woman in her fifties with a United States English accent. Given the transcription text as: In my book of memory, time writes tales of moments past.
Answer: <reference to original audio>

Table 1: Representative examples from SIFT-50M for different categories of instruction types.

model to construct diverse and generalizable QA pairs. In addition, a language field is incorporated into the metadata (by sampling examples from the top-20 locales in the Common Voice dataset) to facilitate the generation of instructions for the acoustic-based language identification (LID) task.

We introduce a novel task in the form of *Comparison* instructions, where the LLM compares two audio files based on their speech characteristics. The LLM is provided with metadata from two distinct audio files and is prompted to generate comparative questions and free-form answers rather than binary responses. Due to the increased complexity of this task, we employ Amazon Nova Pro (Intellicence, 2024), which has demonstrated stronger performance on text benchmarks compared to Mixtral 8x7B.

Open-ended Instructions: In this category, the LLM is prompted to generate more complex and

Category	#Samples (train/dev/EvalSIFT)
Closed-Ended	
– Acoustic-level	17.8M / 100K / 2.5K
– Content-level	14.5M / 80K / 2.5K
– Word-Align	9.8M / 40K / 2.5K
– Comparison	3.6M / 100K / 2.5K
Open-Ended	4.3M / 100K / 10K
Controllable Generation	5.6M / 50K / 10K
Total	55.6M / 470K / 30K

Table 2: High-level statistics of SIFT-50M dataset showing the distribution of categories pooled across the source speech datasets and languages.

diverse questions, with answers inferred from the metadata rather than provided directly by it. This approach encourages the model to produce thoughtful and detailed responses. We employ Amazon Nova Pro to generate open-ended dataset.

Controllable Generation: This category includes instructions for expressive speech synthesis from text with specific speech characteristics. These instructions are generated using Mixtral 8x7B based on the provided acoustic and content-level metadata.

EvalSIFT: We also release a benchmark, EvalSIFT, constructed in a manner similar to SIFT-50M. For each of the five languages, we generate 2K examples per category (closed-ended, open-ended, and controllable generation), yielding a total of 30K examples.

3.3 Quality Assurance

During the metadata extraction stage, we reject any values that fall outside the feasible ranges for the respective speech characteristics, ensuring accurate mapping from numerical to categorical values. High-quality, well-crafted LLM-generated data is essential for efficient model fine-tuning and improving the model’s ability to follow unseen instructions. We employ two types of quality control over the LLM-generated datasets. First, the authors perform an exhaustive review of the LLM responses across all categories and languages. We establish a feedback loop with iterative prompt refinement to ensure that the model’s responses are cogent, meet our expectations, and are free from hallucinations. Second, we conduct ablation studies on the generated data by evaluating performance on the SIFT-50M development sets, which allows us to identify and discard problematic data or even entire categories. For example, we found that instructions based on speech-text alignment that include numer-

ical values (e.g., “What words occur between the 3 and 5 second mark?”) significantly degraded model performance. Consequently, we retain only instructions with non-numerical values (i.e., instructions similar to the word-align category shown in Table 1). We release the rejected portion as a bonus research partition to support future research directions.

4 Experimental Setup

To evaluate the effectiveness of SIFT-50M, we train a speech-text LLM, called SIFT-LLM, on our dataset. In this section, we detail the model architecture and training configuration.

4.1 SIFT-LLM Architecture

Acoustic Encoders: We adopt a hybrid speech representation approach: input speech is encoded as continuous embeddings, while output speech is generated as discrete tokens. Specifically, we use Whisper-medium (Radford et al., 2023) as the encoder to process input audio. For discrete speech tokens, we experiment with two codec options: (1) k-means clustered HuBERT embeddings (Lakhotia et al., 2021), which primarily capture the semantic content of speech, and (2) X-codec2 (Ye et al., 2025), which fuses semantic and acoustic codes into a unified codebook. The HuBERT codes are subsequently converted to audio using a vocoder based on the Big-VGan architecture (Lee et al., 2023). Unless stated otherwise, the setup should be assumed to use HuBERT codes.

Large Language Model: Our core LLM is Qwen2.5-7B-instruct (Yang et al., 2024a). We add a linear layer on top of the Whisper-medium encoder to project the 1024-dimensional speech embeddings into the LLM’s 3584-dimensional space. Additionally, we expand the LLM’s vocabulary by the size of the codebook (2000 for HuBERT and 65536 for X-codec2). We employ low-rank adaptation (Hu et al., 2021) to train the LLM parameters.

4.2 Model Training

Continued Pre-training: To align the audio embedding space of the pre-trained Whisper encoder with the input representation of the Qwen2.5 LLM and to train the randomly initialized projection layer, we perform continued pre-training of SIFT-LLM. This phase involves training on a mixture of speech understanding and generation tasks—namely, ASR, emotion recognition (ER),

speaker-attributed ASR (SA-ASR), intent classification (IC), slot entity recognition (SER), speech-to-text translation (S2TT), speech-to-speech translation (S2ST), and text-to-speech (TTS). We convert the IC task into a multiple-choice question format by providing possible intent values in the instructions. All tasks are represented using natural language instructions. In addition to SIFT-50M, we also release the prompt templates used during this stage.

During pre-training, the acoustic encoder remains frozen, and we train only the linear connector layer and the LoRA parameters of the LLM. We set the LoRA rank to 16, resulting in 8.7M learnable parameters, and train for a total of 200K steps. The resulting pre-trained models are referred to as SIFT-LLM PT. For controllable generation experiments, we continue pre-training for an additional 200K steps, using a LoRA rank of 128 and making the core LLM’s embedding layer fully trainable. In this stage, the total number of learnable parameters increases to 590M for HuBERT codes setup and 830M for X-codec2 setup. We also increase the relative weights of the TTS and S2ST datasets.

Instruction Fine-tuning: Following continued pre-training, we perform an instruction fine-tuning stage using the SIFT-50M dataset. We conduct separate fine-tuning processes for speech understanding and controllable generation. For speech understanding, training is initialized from the 50Kth SIFT-LLM PT checkpoint and uses only the speech understanding instructions (closed- and open-ended) from SIFT-50M. The number of trainable parameters remains the same as during pre-training for the first 200K steps. We refer to this model as SIFT-LLM. For controllable generation, training is initialized from the 400Kth SIFT-LLM PT checkpoint and focuses exclusively on controllable generation instructions, with the number of trainable parameters unchanged from the last 200K steps of pre-training. This model is referred to as SIFT-LLM GEN. For further details on training datasets, prompts, and setup, we refer the reader to Appendix D.

5 Experimental Results

5.1 Evaluation Setup

We evaluate SIFT-LLM by benchmarking it against publicly available speech-text LLMs of similar sizes, including SALMONN-7B and Qwen2-Audio-7B-Instruct. For a direct comparison be-

Model	Closed-Ended		Open-Ended		Dynamic-Superb Tasks					
	DS-1	EvalSIFT	AB-Chat	EvalSIFT	Audio	PL	Semt.	Degrd.	Content	Speaker
SALMONN-7B	34.7	21.9	6.4	6.0	31.7	<u>30.5</u>	<u>47.5</u>	30.0	45.2	31.9
Qwen2-Audio-Inst.	<u>48.0</u>	<u>25.1</u>	<u>7.2</u>	<u>7.3</u>	53.5	28.9	40.3	43.9	70.6	<u>43.6</u>
O-ASQA-LLM	45.9	22.9	6.6	4.7	28.5	30.0	38.6	<u>45.9</u>	<u>72.3</u>	40.7
SIFT-LLM (ours)	57.4	46.1	7.3	7.8	<u>37.5</u>	42.8	51.3	63.6	75.6	47.7

Table 3: Evaluation results of speech-text LLMs on Dynamic-Superb (DS-1), AIR-Bench Chat (AB-Chat), and EvalSIFT (English). We report accuracy (in %) for closed-ended evaluations and LLM score (0 to 10) for open-ended evaluations. The adjacent table provides a breakdown by task categories in Dynamic-Superb – Audio, Paralinguistics (PL), Semantics (Semt.), Degradation (Degrd.), Content, and Speaker. **Bold** values indicate the best results, and underlined values indicate the second-best results.

Task (Metric)	Test Set	SALMN	QwA-Inst	OASQAL	SIFT-LLM	SIFT-LLM PT	
						50k-ckp	200k-ckp
ASR (WER ↓)	LS-Clean	2.5	4.8	3.8	<u>3.5</u>	2.5	2.3
	LS-Other	5.7	<u>7.4</u>	8.1	7.5	5.4	5.0
	PS	<u>22.2</u>	29.4	19.4	26.0	24.4	24.9
	FLEURS-en	9.2	<u>9.1</u>	11.3	8.1	6.6	6.4
	FLEURS-5	23.7	<u>13.2</u>	23.1	11.4	9.0	8.2
ER (Acc. ↑)	MSP-test1	38.5	<u>40.0</u>	38.6	53.6	54.3	54.3
	MSP-test2	28.4	32.1	<u>40.4</u>	50.2	52.4	52.7
IC (Acc. ↑)	SLURP	58.4	<u>86.0</u>	68.0	92.7	94.8	96.0
SER (Acc. ↑)	SLURP	26.4	<u>52.1</u>	48.3	71.3	73.1	72.7
S2TT (BLEU ↑)	de → en	21.3	31.4	28.9	<u>29.2</u>	34.4	35.2
	fr → en	20.6	34.2	29.8	<u>30.9</u>	36.3	37.0
	it → en	18.9	33.6	29.4	<u>31.7</u>	36.0	36.8
	es → en	21.6	36.3	33.1	<u>35.5</u>	39.7	40.4

Table 4: Evaluation results for SALMONN-7B (SALMN), Qwen2-Audio-Instruct (QwA-Inst), O-ASQA-LLM (OASQAL), and SIFT-LLM on foundational tasks, along with results for SIFT-LLM PT 50K and 200K checkpoints. Accuracies (Acc.) are reported as percentages. FLEURS-5 includes test sets for five languages: en, de, fr, it, and es. PS denotes the People’s Speech dataset. S2TT results are reported on CoVoST2. **Bold** values indicate the best results, while underlined values represent the second-best results among **instruction fine-tuned models**.

tween the SIFT-50M dataset and the OpenASQA dataset, we train a speech-text LLM using the same pre-trained model as SIFT-LLM but fine-tuned on OpenASQA; we refer to this model as O-ASQA-LLM. Across all evaluations, we use Temperature= 0.1, Top P= 0.95, a repetition penalty of 1.1 (Keskar et al., 2019), and set the no_repeat_ngram size to 3 (Wolf, 2020).

5.2 Instruction-following Evaluation

Benchmarks: Several benchmarks have been proposed for evaluating instruction-following speech-text LLMs. For example, Dynamic-Superb (Huang et al., 2024b) originally comprised 55 speech and audio tasks and was later expanded to 180 tasks in its Phase-2 release (Huang et al., 2024a). AIR-Bench (Yang et al., 2024b) is composed of two parts: *Foundation*, which covers 19 standard speech tasks, and *Chat*, which includes open-ended

questions. We use Dynamic-Superb (DS-1)² for closed-ended evaluations, and the *Chat* partition of AIR-Bench for open-ended evaluation, selecting only the speech subset (excluding sound and music). Additionally, we report metrics on both the closed- and open-ended categories of EvalSIFT.

SIFT-LLM and O-ASQA-LLM accept any number of text and audio segments in all possible permutations, whereas Qwen2-Audio-Instruct and SALMONN-7B accept only a single audio input. To address this limitation during evaluation for comparison tasks in EvalSIFT or for tasks in Dynamic-Superb that involve multiple audio segments, we concatenate the segments with a 1-second silence between each.

Metrics: While exact match accuracy is a common metric for closed-ended or classification tasks, it becomes less reliable for speech-text LLMs,

²Phase-2 data was not available at the time of writing.

which tend to produce verbose, open-ended responses. Therefore, we adopt an LLM-as-a-judge strategy (Zheng et al., 2023) similar to that used in DS-1 and AIR-Bench, employing Claude-3.5-Sonnet³ for evaluation. Using a chain-of-thought reasoning prompt (Wei et al., 2022), we provide Claude-3.5-Sonnet with both the question and the reference answer. For closed-ended and classification benchmarks, we prompt it to classify responses as correct or incorrect, while for open-ended benchmarks, we ask it to score the responses on a scale from 0 to 10 based on their alignment with the reference answer. We repeated LLM scoring multiple times and found that the width of 95% confidence interval for the mean accuracy (in %) was less than 0.1 and for the mean score was less than 0.01.

Results: Table 3 summarizes the performance of various models on instruction-following benchmarks. SIFT-LLM achieves competitive results across all benchmarks and outperforms all other models. Notably, it performs better than O-ASQA-LLM on every benchmark, underscoring the importance of the large-scale and diverse SIFT-50M dataset. Furthermore, SIFT-LLM is second-best on the Audio task in the Dynamic-Superb benchmark, despite not being explicitly trained on audio tasks. We observed that the performance gap between SIFT-LLM and other models is less for Content-level instructions in EvalSIFT, but significantly larger for Acoustic-level, Word-Align, and Comparison type instructions, highlighting the value of data addressing these dimensions in the SIFT-50M dataset. Finally, the model shows signs of reasoning abilities, as evidenced by the examples in Appendix H, where SIFT-LLM reasons through open-ended questions before arriving at its conclusions.

Results on non-English languages: Since both Dynamic-Superb and AIR-Bench Chat benchmarks provide instructions exclusively in English, we use EvalSIFT for multilingual evaluation. The results are reported in Table 5. SIFT-LLM outperforms SALMONN-7B and Qwen2-Audio-Instruct on closed-ended evaluations across languages. However, Qwen2-Audio-Instruct remains competitive on open-ended evaluations, as observed with English data as well. Overall, absolute accuracy is lower on non-English languages compared to English. Detailed evaluation results are provided in Appendix F.

Model	German	French	Italian	Spanish
SAL	15.0 4.3	16.3 5.0	14.3 5.0	16.7 5.4
QwA	18.6 6.0	18.8 6.8	18.2 7.2	21.2 7.3
SIFL	39.0 6.6	34.3 7.1	33.2 7.5	35.6 7.0

Table 5: Evaluation results for SALMONN-7B (SAL), Qwen2-Audio-Instruct (QwA), and SIFT-LLM (SIFL) on EvalSIFT for non-English languages. Accuracy (in %) for the closed-ended category and LLM score (0–10) for the open-ended category are separated by a pipe (|). **Bold** values indicate the best results.

5.3 Evaluation on Foundational Tasks

In addition to the instruction-following evaluation, we assess performance on standard speech understanding tasks using task-specific metrics: Word Error Rate (WER) for ASR, accuracy for ER, IC, and SER tasks, and BLEU for S2TT. Since model performance can be sensitive to the prompts used (Wang et al., 2024), we use the prompts recommended by the corresponding model wherever possible. For ER, IC, and SER tasks, we leverage Claude-3.5-Sonnet to classify responses as correct or incorrect, ensuring robustness against model-specific output formats. Sequence generation tasks, such as ASR and S2TT pose additional challenges due to some models producing extra text; to address this, we use regular expressions and Claude-3.5-Sonnet to parse model hypotheses. Furthermore, we apply Whisper text normalization⁴ before computing WER.

The results, presented in Table 4, reveal that no single model excels across all tasks. While SALMONN-7B achieves the best performance on ASR benchmarks for English, Qwen2-Audio-Instruct outperforms the others on S2TT, and SIFT-LLM significantly outperforms on ER and spoken language understanding tasks. On the multilingual ASR task, SIFT-LLM and Qwen2-Audio-Instruct show significantly better results than SALMONN-7B and O-ASQA-LLM. Compared to the pre-trained model from which SIFT-LLM is initialized (the 50Kth checkpoint), SIFT-LLM shows a slight decline in performance across tasks—an observation consistent with findings for Qwen2-Audio and its instruction fine-tuned variant (Wang et al., 2024). Furthermore, performance differences between the 50K and 200K pre-trained checkpoints are minimal across all tasks.

³<https://aws.amazon.com/bedrock/claude/>

⁴<https://github.com/openai/whisper/tree/main/whisper/normalizers>

Feature	HuBERT codes		X-codec2	
	MAE (\downarrow)	QWK (\uparrow)	MAE (\downarrow)	QWK (\uparrow)
Pitch variation	$0.99 \pm 1.05e-2$	$0.15 \pm 2.17e-2$	$0.50 \pm 0.83e-2$	$0.69 \pm 1.50e-2$
Speaking rate	$0.65 \pm 0.35e-2$	$0.46 \pm 0.49e-2$	$0.24 \pm 1.16e-2$	$0.73 \pm 1.67e-2$
Intensity	$0.18 \pm 0.16e-2$	$0.02 \pm 0.57e-2$	$0.11 \pm 0.68e-2$	$0.67 \pm 1.78e-2$

Table 6: Evaluation of SIFT-LLM GEN on the controllable generation subset of EvalSIFT. We report MAE and QWK metrics by comparing the speech characteristics of audio generated by SIFT-LLM GEN against the attributes specified in the instructions. We run 10 inference iterations of SIFT-LLM GEN with the temperature set to 0.8. For HuBERT codes, we synthesize audio using 10 different speaker combinations per iteration. Results are reported as the mean \pm standard deviation.

5.4 Evaluation on Controllable Generation

We provide baseline results on the controllable generation set of EvalSIFT for the SIFT-LLM GEN model. To assess controllability, we compare the acoustic metadata extracted from the generated speech with the controllability parameter specified in the instruction. We report two metrics: Mean Absolute Error (MAE) and Quadratic Weighted Kappa (QWK). The distance-based metrics are computed on values categorized into ordinal groups—for example, speaking rate is scored on a scale ranging from “very slowly” (0) to “very fast” (6). We set the sampling temperature to 0.8 during the inference to balance diversity and coherence in generation.

The results are presented in Table 6. We focus on stationary characteristics such as pitch variation, speaking rate, and intensity to isolate the controllability evaluation from the effect of the reference speaker signal used during waveform synthesis in the HuBERT-based setup. The results show that the MAE for each category is less than 1 across both setups, indicating that, on average, predictions deviate from the true category by less than one step. Notably, SIFT-LLM GEN with X-codec2 setup demonstrates significantly better controllability. This is attributed to X-codec2’s ability to encode richer speech representations that captures aspects beyond semantics, including prosody and timbre. It also enables more effective control over speaker-dependent characteristics; for example, we found accuracy for gender controllability to be 95.8% with the X-codec2 setup.

Further analysis of the limitations of semantic-only representations like HuBERT and the effect of speaker references as input to the vocoder, is provided in Appendix F.3.

5.5 Ablation Studies

For ablation studies, the default SIFT-LLM configuration employs a LoRA rank of 8 and is trained exclusively on the English subset of SIFT-50M dur-

Setup	DS-1	EvalSIFT (Closed)	EvalSIFT (Open)
Default	57.3	45.4	8.0
Init. from 200K ckpt	52.5	42.3	7.7
No Pre-training	56.7	43.7	7.9
LoRA rank=16	58.4	45.1	8.0
LoRA rank=32	55.7	43.9	7.8
No open-ended data	54.3	42.2	6.1
No word-align data	57.1	41.5	8.0
No comparison data	56.1	34.3	7.3

Table 7: Evaluation results on DS-1 and EvalSIFT (English) for various training setups. The *Default* setup is initialized from the 50Kth PT checkpoint and instruction-finetuned using LoRA (rank = 8). Metrics reported are accuracy (in %) for DS-1 and closed-ended EvalSIFT, and LLM score (0–10) for the open-ended EvalSIFT.

ing instruction fine-tuning. The remainder of the setup is identical to that described in Section 4.2.

Factors Influencing Instruction-Following: Table 7 summarizes the key factors impacting the model’s instruction-following abilities. Extended pre-training can cause the core LLM to lose some of its instruction-following capability. While eliminating pre-training leads to only a moderate degradation on instruction-following benchmarks, it results in a significant decline in performance on foundational tasks like ASR and S2TT. Increasing the LoRA rank from 8 to 16 yields mixed results, but raising it further to 32 results in degradation. Training exclusively on closed-ended datasets significantly impairs open-ended performance, and it also negatively affects accuracy on the closed-ended set. Additionally, omitting comparison data leads to a marked drop in performance on closed-ended EvalSIFT, especially for comparison instructions, since comparison data is the only subset of EvalSIFT that involves multiple audio inputs.

Effect of Data Volume: To assess the impact of data volume during instruction fine-tuning, we trained SIFT-LLM using varying fractions of the SIFT-50M dataset: half, one-quarter, and one-

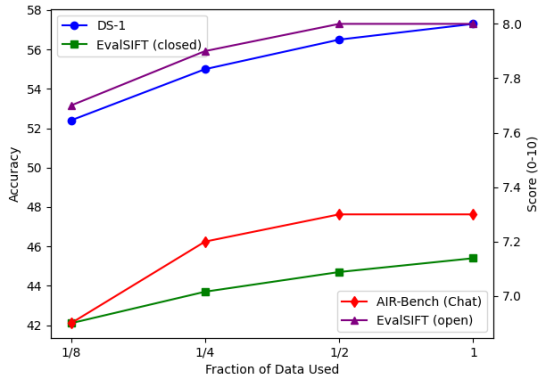


Figure 2: Effect of SIFT data volume used during instruction fine-tuning on SIFT-LLM’s performance, as measured on DS-1, AIR-Bench Chat, and EvalSIFT.

eighth of the total volume. The evaluation results, shown in Figure 2, indicate that even with just one-eighth of the data, the model achieves strong performance on open-ended benchmarks. However, increasing the data volume leads to further improvements on closed-ended instructions.

6 Conclusions

In this work, we introduce SIFT-50M, a large-scale multilingual dataset designed for instruction fine-tuning of speech-text LLMs. By integrating a diverse set of instructions spanning various speech understanding tasks, SIFT-50M not only facilitates instruction-following but also helps models generalize better to unseen instructions. Our experimental results show that SIFT-LLM, our model trained on SIFT-50M, achieves strong performance on instruction-following benchmarks and competitive results on foundational speech understanding tasks. Moreover, we present EvalSIFT, a benchmark dataset tailored for the systematic evaluation of speech-text LLMs for both speech understanding and controllable generation.

Limitations

While SIFT-LLM demonstrates strong performance on instruction-following benchmarks, it does not achieve state-of-the-art results on foundational speech tasks. Future work could explore the trade-off between enhancing instruction-following generalization and maintaining competitive performance on foundational tasks relative to task-specific models. Additionally, SIFT-LLM occasionally generates hallucinated responses when queried about content unrelated to the input audio. To mitigate this, we generated additional instruction data by

prompting the LLM to produce questions unrelated to the audio along with appropriate answers. However, assigning a higher weight to this auxiliary dataset during training resulted in reduced speech understanding capability.

As LLMs continue to improve, the LLM-as-a-judge approach has become increasingly common for subjective evaluations. However, the reported metrics can vary depending on the LLM used. We observed that the variation was significantly low when running evaluations multiple times with the same LLM. Furthermore, speech-text LLMs have been found to be sensitive to the prompts used for evaluating foundational tasks such as ASR and S2TT. Although we made every effort to use the prompts recommended by each respective model, the necessary information was not available for some tasks or was provided only for the pre-trained variants.

Ethics Statement

The proposed SIFT-50M dataset is constructed from three publicly available and widely used speech corpora: Common Voice, MLS, and VCTK. The SIFT-50M dataset is released under the CDLA-Sharing-1.0 license. We acknowledge the potential risks associated with the capability of controllable speech generation, such as misuse for deepfakes, voice impersonation, or the spread of misinformation. Although this task involves inherent complexities, our dataset is curated to limit control to generic, broad parameters. SIFT-LLM only extracts, and we only use non-specific features that cannot be used for identification.

Acknowledgements

We would like to thank Leif Rädcl, Hayoon Kim, Melanie McHale, Adam Kaplan, Akshat Jaiswal, Subhadeep Maji, Aparna Khare, Anna Piunova, Mohammed Hethnawi, Harish Mallidi, Haolin Chen, Mateusz Łajszczak, Roberto Barra-Chicote, Luyao Wang, Berke Toptas, Sunwoo Kim, Jing Liu, Yang Li, Ajay Srinivasamurthy, Nathan Susanj, Volker Leutnant, Sri Garimella, and Athanasios Mouchtaris for their support and valuable feedback.

References

Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber. 2019. Common voice: A massively-

- multilingual speech corpus. *arXiv preprint arXiv:1912.06670*.
- Emanuele Bastianelli, Andrea Vanzo, Pawel Swietojanski, and Verena Rieser. 2020. Slurp: A spoken language understanding resource package. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7252–7262.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeanette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42:335–359.
- Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, et al. 2005. The ami meeting corpus: A pre-announcement. In *International workshop on machine learning for multimodal interaction*, pages 28–39. Springer.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518.
- Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, Wanxiang Che, Xiangzhan Yu, and Furu Wei. 2023. Beats: audio pre-training with acoustic tokenizers. In *Proceedings of the 40th International Conference on Machine Learning*, pages 5178–5193.
- Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, et al. 2024. Qwen2-audio technical report. *arXiv preprint arXiv:2407.10759*.
- Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and Jingren Zhou. 2023. Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models. *arXiv preprint arXiv:2311.07919*.
- Joon Son Chung, Arsha Nagrani, and Andrew Senior. 2018. Voxceleb2: Deep speaker recognition. *arXiv preprint arXiv:1806.05622*.
- Yu-An Chung, Yu Zhang, Wei Han, Chung-Cheng Chiu, James Qin, Ruoming Pang, and Yonghui Wu. 2021. W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 244–250. IEEE.
- Christopher Cieri, David Miller, and Kevin Walker. The fisher corpus: A resource for the next generations of speech-to-text.
- Alexis Conneau, Min Ma, Simran Khanuja, Yu Zhang, Vera Axelrod, Siddharth Dalmia, Jason Riesa, Clara Rivera, and Ankur Bapna. 2023. Fleurs: Few-shot learning evaluation of universal representations of speech. In *2022 IEEE Spoken Language Technology Workshop (SLT)*, pages 798–805. IEEE.
- Nilaksh Das, Saket Dingliwal, Srikanth Ronanki, Rohit Paturi, David Huang, Prashant Mathur, Jie Yuan, Dhanush Bekal, Xing Niu, Sai Muralidhar Jayanthi, et al. 2024. Speechverse: A large-scale generalizable audio language model. *arXiv preprint arXiv:2405.08295*.
- Alexandre Défossez, Jade Copet, Gabriel Synnaeve, and Yossi Adi. 2022. High fidelity neural audio compression. *arXiv preprint arXiv:2210.13438*.
- Zhihao Du, Jiaming Wang, Qian Chen, Yunfei Chu, Zhifu Gao, Zerui Li, Kai Hu, Xiaohuan Zhou, Jin Xu, Ziyang Ma, et al. 2023. Lauragpt: Listen, attend, understand, and regenerate audio with gpt. *arXiv preprint arXiv:2310.04673*.
- Yassir Fathullah, Chunyang Wu, Egor Lakomkin, Jun-teng Jia, Yuan Shanguan, Ke Li, Jinxi Guo, Wenhan Xiong, Jay Mahadeokar, Ozlem Kalinli, et al. 2024. Prompting large language models with speech recognition abilities. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 13351–13355. IEEE.
- Daniel Galvez, Greg Diamos, Juan Ciro, Juan Felipe Cerón, Keith Achorn, Anjali Gopi, David Kanter, Maximilian Lam, Mark Mazumder, and Vijay Janapa Reddi. 2021. The people’s speech: A large-scale diverse english speech recognition dataset for commercial usage. *arXiv preprint arXiv:2111.09344*.
- John J Godfrey, Edward C Holliman, and Jane McDaniel. 1992. Switchboard: Telephone speech corpus for research and development. In *Acoustics, speech, and signal processing, ieee international conference on*, volume 1, pages 517–520. IEEE Computer Society.
- Yuan Gong, Sameer Khurana, Leonid Karlinsky, and James Glass. 2023a. Whisper-at: Noise-robust automatic speech recognizers are also strong general audio event taggers. *arXiv preprint arXiv:2307.03183*.
- Yuan Gong, Alexander H Liu, Hongyin Luo, Leonid Karlinsky, and James Glass. 2023b. Joint audio and speech understanding. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 1–8. IEEE.
- Yuan Gong, Hongyin Luo, Alexander H Liu, Leonid Karlinsky, and James Glass. 2023c. Listen, think, and understand. *arXiv preprint arXiv:2305.10790*.
- Zhifang Guo, Yichong Leng, Yihan Wu, Sheng Zhao, and Xu Tan. 2023. Promptpts: Controllable text-to-speech with text descriptions. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.

- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Shujie Hu, Long Zhou, Shujie Liu, Sanyuan Chen, Lingwei Meng, Hongkun Hao, Jing Pan, Xunying Liu, Jinyu Li, Sunit Sivasankaran, et al. 2024. Wavllm: Towards robust and adaptive speech large language model. *arXiv preprint arXiv:2404.00656*.
- Chien-yu Huang, Wei-Chih Chen, Shu-wen Yang, Andy T Liu, Chen-An Li, Yu-Xiang Lin, Wei-Cheng Tseng, Anuj Diwan, Yi-Jen Shih, Jiatong Shi, et al. 2024a. Dynamic-superb phase-2: A collaboratively expanding benchmark for measuring the capabilities of spoken language models with 180 tasks. *arXiv preprint arXiv:2411.05361*.
- Chien-yu Huang, Ke-Han Lu, Shih-Heng Wang, Chi-Yuan Hsiao, Chun-Yi Kuan, Haibin Wu, Siddhant Arora, Kai-Wei Chang, Jiatong Shi, Yifan Peng, et al. 2024b. Dynamic-superb: Towards a dynamic, collaborative, and comprehensive instruction-tuning benchmark for speech. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12136–12140. IEEE.
- Jaesung Huh. 2024. *Voice-gender-classifier*. Accessed: 2025-01-07.
- Amazon Artificial General Intelligence. 2024. *The amazon nova family of models: Technical report and model card*. Amazon Technical Reports.
- Yannick Jadoul, Bill Thompson, and Bart De Boer. 2018. Introducing parselmouth: A python interface to praat. *Journal of Phonetics*, 71:1–15.
- Adam Janin, Don Baron, Jane Edwards, Dan Ellis, David Gelbart, Nelson Morgan, Barbara Peskin, Thilo Pfau, Elizabeth Shriberg, Andreas Stolcke, et al. 2003. The icsi meeting corpus. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03)*, volume 1, pages I–I. IEEE.
- Ye Jia, Michelle Tadmor Ramanovich, Quan Wang, and Heiga Zen. 2022. Cvss corpus and massively multilingual speech-to-speech translation. *arXiv preprint arXiv:2201.03713*.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Kavita Kasi and Stephen A Zahorian. 2002. Yet another algorithm for pitch tracking. In *2002 IEEE international conference on acoustics, speech, and signal processing*, volume 1, pages I–361. IEEE.
- Masaya Kawamura, Ryuichi Yamamoto, Yuma Shirahata, Takuya Hasumi, and Kentaro Tachibana. 2024. Libritts-p: A corpus with speaking style and speaker identity prompts for text-to-speech and style captioning. *arXiv preprint arXiv:2406.07969*.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Yoach Lacombe. 2024. *Data-speech*. Accessed: 2024-12-23.
- Kushal Lakhota, Eugene Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Abdelrahman Mohamed, et al. 2021. On generative spoken language modeling from raw audio. *Transactions of the Association for Computational Linguistics*, 9:1336–1354.
- Sang-gil Lee, Wei Ping, Boris Ginsburg, Bryan Catanzaro, and Sungroh Yoon. 2023. *Bigvgan: A universal neural vocoder with large-scale training*. Preprint, arXiv:2206.04658.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Yuang Li, Yu Wu, Jinyu Li, and Shujie Liu. 2023b. Prompting large language models for zero-shot domain adaptation in speech recognition. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 1–8. IEEE.
- Reza Lotfian and Carlos Busso. 2017. Building naturalistic emotionally balanced speech corpus by retrieving emotional speech from existing podcast recordings. *IEEE Transactions on Affective Computing*, 10(4):471–483.
- Dan Lyth and Simon King. 2024. Natural language guidance of high-fidelity text-to-speech with synthetic annotations. *arXiv preprint arXiv:2402.01912*.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5206–5210. IEEE.
- Jongseok Park, Kyubyong & Kim. 2019. g2pe. <https://github.com/Kyubyong/g2p>.
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. Mls: A large-scale multilingual dataset for speech research. *arXiv preprint arXiv:2012.03411*.

- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR.
- Paul K Rubenstein, Chulayuth Asawaroengchai, Duc Dung Nguyen, Ankur Bapna, Zalán Borsos, Félix de Chaumont Quitry, Peter Chen, Dalia El Badawy, Wei Han, Eugene Kharitonov, et al. 2023. Audiopalm: A large language model that can speak and listen. *arXiv preprint arXiv:2306.12925*.
- H Sak. 2014. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. INTERSPEECH.
- Reo Shimizu, Ryuichi Yamamoto, Masaya Kawamura, Yuma Shirahata, Hironori Doi, Tatsuya Komatsu, and Kentaro Tachibana. 2024. Promptts++: Controlling speaker identity in prompt-based text-to-speech using natural language descriptions. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12672–12676. IEEE.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, MA Zejun, and Chao Zhang. Salmonn: Towards generic hearing abilities for large language models. In *The Twelfth International Conference on Learning Representations*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Bin Wang, Xunlong Zou, Geyu Lin, Shuo Sun, Zhuohan Liu, Wenyu Zhang, Zhengyuan Liu, AiTi Aw, and Nancy F Chen. 2024. Audiobench: A universal benchmark for audio large language models. *arXiv preprint arXiv:2406.16020*.
- Changhan Wang, Morgane Rivière, Ann Lee, Anne Wu, Chaitanya Talnikar, Daniel Haziza, Mary Williamson, Juan Pino, and Emmanuel Dupoux. 2021a. Voxpopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation. In *ACL 2021-59th Annual Meeting of the Association for Computational Linguistics*.
- Changhan Wang, Anne Wu, Jiatao Gu, and Juan Pino. 2021b. Covost 2 and massively multilingual speech translation. In *Interspeech*, pages 2247–2251.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Thomas Wolf. 2020. Transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Junichi Yamagishi, Christophe Veaux, and Kirsten Macdonald. 2019. Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92).
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024a. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Qian Yang, Jin Xu, Wenrui Liu, Yunfei Chu, Ziyue Jiang, Xiaohuan Zhou, Yichong Leng, Yuanjun Lv, Zhou Zhao, Chang Zhou, et al. 2024b. Airbench: Benchmarking large audio-language models via generative comprehension. *arXiv preprint arXiv:2402.07729*.
- Zhen Ye, Xinfu Zhu, Chi-Min Chan, Xinsheng Wang, Xu Tan, Jiahe Lei, Yi Peng, Haohe Liu, Yizhu Jin, Zheqi DAI, et al. 2025. Llasa: Scaling train-time and inference-time compute for llama-based speech synthesis. *arXiv preprint arXiv:2502.04128*.
- Wenqi Yu, Changli Tang, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. 2024. Connecting speech encoder and large language model for asr. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12637–12641. IEEE.
- AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2236–2246.
- Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu. 2019. Libritts: A corpus derived from librispeech for text-to-speech. *arXiv preprint arXiv:1904.02882*.
- Yu Zhang, Wei Han, James Qin, Yongqiang Wang, Ankur Bapna, Zhehuai Chen, Nanxin Chen, Bo Li, Vera Axelrod, Gary Wang, et al. 2023. Google usm: Scaling automatic speech recognition beyond 100 languages. *arXiv preprint arXiv:2303.01037*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623.

A SIFT-50M Dataset

<p>Closed-ended: Acoustic-Level Question: ¿Hay distorsión en el audio? Answer: Casi no hay ruido de distorsión en el audio.</p>
<p>Closed-ended: Acoustic-Level Question: What is the transcript of the audio, and what emotions can be inferred from it? Answer: The transcript of the audio is 'A girl smiles while interacting with medical personnel'. The emotions that can be inferred from the transcript are happiness and cooperation.</p>
<p>Closed-ended: Word-Align Question: Wie verändert sich die durchschnittliche Tonhöhe im Laufe des gesprochenen Satzes? Answer: Die durchschnittliche Tonhöhe des gesprochenen Satzes beginnt mit einer ziemlich tiefen Tonhöhe, sinkt dann zu einem sehr tiefen Ton und steigt dann allmählich zu einer mittleren Tonhöhe an. Am Ende des Satzes sinkt die Tonhöhe erneut auf eine sehr tiefe Tonhöhe ab.</p>
<p>Closed-ended: Comparison Question: Qual è la differenza di tonalità tra i due audio? Answer: Il primo audio ha una tonalità molto alta, mentre il secondo ha una tonalità abbastanza alta.</p>
<p>Open-ended Question: Dans quelle mesure la réverbération et le niveau de distorsion perçus dans l'audio affectent-ils la clarté de la transcription et la compréhension globale du message? Answer: La réverbération légèrement lointain et le niveau de distorsion bruyant rendent la transcription moins claire, compliquant la compréhension précise du message. Cela nécessite un effort supplémentaire de l'auditeur pour interpréter le contenu, ce qui peut entraîner des malentendus ou des interprétations erronées.</p>
<p>Controllable Generation Question: Synthetisieren Sie Sprache mit einer mäßig schnellen Geschwindigkeit, einer mittelmäßigen Lautstärke und einer sehr monotonen Tonhöhenvariation. Der Transkriptionstext lautet: 'Auf Kamelen ritten sie durchs Morgenland.' Answer: <reference to original audio></p>

Table 8: Representative multilingual examples from SIFT-50M for various instruction categories.




Figure 3: Dataset distribution showing the multi-lingual nature of SIFT-50M and the different categories within each language.




Figure 4: Distribution of the number of examples per language in the acoustic-based language ID (LID) task that is part of the closed-ended instructions.

B Metadata Extraction

B.1 Categorical Values of Acoustic-Level Metadata

Average Pitch	Pitch Variation	Speaking Rate	Average Intensity	Noise Level	Reverberation	Distortion Level	Speech Quality
Very low pitch	Very monotone	Very slowly	Softly	Very noisy	Very distant-sounding	Extremely noisy	Very bad speech quality
Quite low pitch	Quite monotone	Quite slowly	Moderate volume	Noisy	Distant-sounding	Very noisy	Bad speech quality
Slightly low pitch	Slightly monotone	Slightly slowly	Loudly	Slightly noisy	Slightly distant-sounding	Noisy	Slightly bad speech quality
Moderate pitch	Moderate intonation	Moderate speed		Balanced in clarity	Slightly close-sounding	Slightly noisy	Moderate speech quality
Slightly high pitch	Slightly expressive	Slightly fast		Slightly clean	Very close-sounding	Almost no noise	Great speech quality
Quite high pitch	Quite expressive	Quite fast		Clean		Very clear	Wonderful speech quality
Very high pitch	Very expressive	Very fast		Very clean			

Table 9: Categorical values of acoustic-level metadata

B.2 Sample Metadata

We show an example of metadata from the VCTK dataset after mapping to categorical values.

```
{
  "acoustic_level": {
    "average_pitch": "slightly low pitch",
    "pitch_variation": "quite monotone",
    "speaking_rate": "moderate speed",
    "average_intensity": "moderate volume",
    "noise_level": "slightly noisy",
    "reverberation": "very close-sounding",
    "distortion_level": "almost no noise",
    "speech_quality": "moderate speech quality"
  },
  "content_level": {
    "phonemes": "SH|IY1| |K|AE1|N| |S|K|UW1|P| |DH|IY1|Z| |TH|IH1|NG|Z| ...",
    "age_category": "twenties",
    "accent": "English",
    "gender": "Female",
    "transcript": "She can scoop these things into three red bags, and we will go meet her
Wednesday at the train station."
  }
}
```

B.2.1 Sample Word-level Metadata

```
[
  {
    "word": "nodded",
    "position": 1,
    "average_pitch": "slightly low pitch",
    "pitch_variation": "very monotone",
    "average_intensity": "moderate volume"
  },
  {
    "word": "his",
    "position": 2,
    "average_pitch": "slightly low pitch",
    "pitch_variation": "very monotone",
    "average_intensity": "moderate volume"
  },
  {
    "word": "head",
    "position": 3,
    "average_pitch": "quite low pitch",
    "pitch_variation": "very monotone",
    "average_intensity": "moderate volume"
  }
]
```

C Natural Language Instruction Generation

C.1 LLM Hyperparameters

We provide the LLM hyper-parameters used for generating instructions in Table 10. Mixtral 8x7B model was run on a g5.48xlarge⁵ instance (8 A10 GPUs), while Amazon Nova Pro was used on AWS Bedrock⁶.

Parameter	Mixtral 8x7B	Amazon Nova Pro
Temperature	0.7	0.7
Top p	0.9	-
Max. tokens	2048	2048

Table 10: LLM hyperparameters used during instruction generation

C.2 Data Format

Below is an example of how SIFT-50M data is organized in the ChatML⁷/Messages API⁸ format.

```
{
  "id": "<instance_id>",
  "messages": [
    {
      "role": "user",
      "content": [
        {"text": "Which audio has a speaker with a Scottish accent?"},
        {"audio_path": "<audio_path_1>"},
        {"audio_path": "<audio_path_2>"}
      ]
    },
    {
      "role": "assistant",
      "content": [{"text": "The second audio has a Scottish-accented speaker."}]
    }
  ]
}
```

⁵<https://aws.amazon.com/ec2/instance-types/g5/>

⁶<https://aws.amazon.com/bedrock/>

⁷<https://platform.openai.com/docs/api-reference/chat/create>

⁸<https://docs.anthropic.com/en/api/messages>

D Training Setup

D.1 Pre-training Datasets

Task	Task Wt.	Dataset	#Hours
Automatic Speech Recognition (ASR)	0.35	Librispeech (Panayotov et al., 2015)	960
		MLS (Pratap et al., 2020)	50K
		FLEURS (Conneau et al., 2023)	1.4K
		People’s Speech (Galvez et al., 2021)	30K
Emotion Recognition (ER)	0.05	MSP-Podcast (Lotfian and Busso, 2017)	237
Speaker-Attributed ASR (SA-ASR)	0.05	AMI Meeting Corpus (Carletta et al., 2005)	100
		ICSI Meeting Corpus (Janin et al., 2003)	70
		Fisher (Cieri et al.)	2K
Intent Classification (IC) Slot Entity Recognition (SER)	0.05	SLURP (Bastianelli et al., 2020)	58
		SLURP (Bastianelli et al., 2020)	58
Speech Translation (S2TT/ S2ST)	0.20	CoVoST (Wang et al., 2021b)	2.9K
		CVSS (Jia et al., 2022)	1.1K
		VoxPopuli (Wang et al., 2021a)	1.8K
		FLEURS (Conneau et al., 2023)	1.4K
Text to Speech (TTS)	0.30	MLS (Pratap et al., 2020)	50K
		FLEURS (Conneau et al., 2023)	1.4K
		People’s Speech (Galvez et al., 2021)	30K

Table 11: Details of the datasets used during the pre-training stage and the weights assigned to different tasks.

D.2 Training Hyper-parameters

We provide the training hyper-parameters used for SIFT-LLM in Table 12. All models were trained on four p4d.24xlarge⁹ instances (32 A100 GPUs in total, each with 40GB of memory), with training for 200K steps taking approximately 3 days.

Parameter	PT: Stage 1	IFT: SIFT-LLM	PT: Stage 2	IFT: SIFT-LLM GEN
Batch Size	256	128	256	128
# Steps	200K	200K	200K	200K
# Warmup Steps	1000	500	500	500
Max Learning Rate	1e-4	1e-4	1e-4	1e-4
LoRA rank	16	16	128	128
LoRA α	8	8	64	64
# Trainable Params	8.7M	8.7M	590M (HuBERT)	590M (HuBERT)

Table 12: Training hyperparameters used during pre-training (PT) and instruction fine-tuning (IFT).

D.3 Pre-Training Prompt Template Examples

Below are examples of instruction templates used during pre-training:

ASR:

Transcribe what the speaker is saying.
Decode this \$language speech.

ER:

Get the most probable emotion in the utterance.
Annotate the speech and obtain the most prominent emotion.

⁹<https://aws.amazon.com/ec2/instance-types/p4/>

SA-ASR:

Transcribe what each of the speakers says and assign speaker labels to transcribed segments. Format the output as text interspersed with speaker tags, where each utterance is followed by its speaker in JSON format (e.g., "speaker": "spk_id"). Use integer speaker ids starting with 0.

IC:

From this audio, what intent does the speaker have? Options include: joke, definition, takeaway_query, and social_query.

SER:

Can you identify the entity that fits the device_type slot?

S2TT:

Recognize and then translate this utterance to text in \$target_language.
Translate text from \$source_language to \$target_language for this audio.

S2ST:

Can you translate this recording to \$target_language

TTS:

Can you convert the text to speech: \$transcription?

E Evaluation Setup

E.1 Instruction-Following Benchmarks

Test Set	# Samples	Languages
Dynamic-Superb (Huang et al., 2024a)	10,400	English
AIR-Bench Chat (Yang et al., 2024b)	793	English
EvalSIFT (Closed-Ended)	10K (2K per lang)	English, German, French, Italian, Spanish
EvalSIFT (Open-Ended)	10K (2K per lang)	English, German, French, Italian, Spanish
EvalSIFT (Control. Generation)	10K (2K per lang)	English, German, French, Italian, Spanish

Table 13: Details of benchmarks used for instruction-following evaluation.

E.2 Foundational Tasks Benchmarks

Task	Test Set	# Samples
ASR	Librispeech Test-Clean (Panayotov et al., 2015)	2,620
	Librispeech Test-Other (Panayotov et al., 2015)	2,560
	People’s Speech (Galvez et al., 2021)	34,898
	FLEURS (en, de, fr, it, es) (Conneau et al., 2023)	350 * 5 = 1,750
ER	MSP-test1 (Lotfian and Busso, 2017)	30,647
	MSP-test2 (Lotfian and Busso, 2017)	14,815
IC	SLURP (Bastianelli et al., 2020)	13,078
SER	SLURP (Bastianelli et al., 2020)	11,585
S2TT	CoVost2 (de → en) (Wang et al., 2021b)	13,509
	CoVost2 (fr → en) (Wang et al., 2021b)	14,760
	CoVost2 (it → en) (Wang et al., 2021b)	8,945
	CoVost2 (es → en) (Wang et al., 2021b)	13,221

Table 14: Details of benchmarks used for foundational tasks evaluation.

E.3 Evaluation Prompt Templates

As model performance can be sensitive to the prompts used (Wang et al., 2024), we use the prompts recommended by the corresponding model wherever we can find for evaluation of SALMONN-7B¹⁰, Qwen2-Audio-Instruct¹¹ and O-ASQA-LLM¹².

Task	Model	Prompt Template
ASR	SALMONN-7B	[en] Recognize the speech and give me the transcription.
		[de] Hören Sie sich die Rede an und schreiben Sie ihren Inhalt auf.
		[fr] Écoutez le discours et écrivez son contenu.
		[it] Ascolta il discorso e scrivi il suo contenuto.
		[es] Escuche el discurso y escriba su contenido.
ASR	Qwen2-Audio-Instruct	[en/de/fr/it/es] Recognize the speech in \$language:
	O-ASQA-LLM	[en] Can you identify the spoken text?
		[de/fr/it/es] Can you identify the spoken text? Your output MUST be in \$language.
ASR	SIFT-LLM	[en] Transcribe what the speaker says.
		[de/fr/it/es] Transcribe this \$language speech.
ER	SALMONN-7B	Describe the emotion of the speaker in one word.
	Qwen2-Audio-Instruct	Recognize the emotion with keywords in English:
	O-ASQA-LLM	Identify the most likely emotion in the following speech.
	SIFT-LLM	Identify the most likely emotion in the following speech.
S2TT	SALMONN-7B	Listen to the speech and translate it into \$target_language.
	Qwen2-Audio-Instruct	Translate the speech into \$target_language:
	O-ASQA-LLM	Translate the audio to \$target_language, returning only the translated text.
	SIFT-LLM	Translate the audio to \$target_language, returning only the translated text.

Table 15: Prompts used to evaluate different speech-text LLMs on foundational tasks.

F Evaluation Results

F.1 EvalSIFT Closed-Ended Results

Model	Acoustic-Level	Content-Level	Word-Align	Comparison
SALMONN-7B	16.1	34.6	9.7	27.3
Qwen2-Audio-Instruct	19.1	37.5	16.8	27.3
O-ASQA-LLM	12.8	39.6	16.7	22.7
SIFT-LLM	58.4	41.9	28.6	55.5

Table 16: Breakdown of evaluation results, reported as accuracy (%) for different sub-categories of the closed-ended set in EvalSIFT (English).

F.2 Dynamic-Superb Detailed Results

Task	SALMONN	Qwen2-Audio-Instruct	O-ASQA-LLM	SIFT-LLM
AccentClassification_AccentdbExtended	13	14.5	17	50.5
DialogueEmotionClassification_DailyTalk	16.5	46	57.29	65.83
EmotionRecognition_MultimodalEmotionlinesDataset	32.32	38.5	51	60.41
HowFarAreYou_3DSpeaker	23.12	31	28	28.5
SpoofDetection_ASVspoof2015	61.5	20.1	14.07	5.5
SpoofDetection_ASVspoof2017	65	27.64	23.62	23
StressDetection_MIRSD	2	24.5	18.59	28.5
Paralinguistics	30.49	28.89	29.94	37.46
BirdSoundDetection_Warblrb10k	74	76.88	40.91	35.35
ChordClassification_AcousticGuitarAndPiano	7.5	40.2	49.5	46.5

¹⁰https://github.com/bytedance/SALMONN/blob/main/prompts/test_prompt.json

¹¹https://github.com/QwenLM/Qwen2-Audio/blob/main/eval_audio/EVALUATION.md

¹²<https://github.com/YuanGongND/ltu/blob/main/README.md#for-1tu-as-openasqa>

Task	SALMONN-7B	Qwen2-Audio-Inst.	O-ASQA-LLM	SIFT-LLM
EnvironmentalSoundClassification_ESC50-Animals	42	62.31	10.61	42.71
EnvironmentalSoundClassification_ESC50-ExteriorAndUrbanNoises	10	60.1	14	36.36
EnvironmentalSoundClassification_ESC50-HumanAndNonSpeechSounds	54.04	47.47	24.5	43.5
EnvironmentalSoundClassification_ESC50-InteriorAndDomesticSounds	5.56	31.5	34.5	48.72
EnvironmentalSoundClassification_ESC50-NaturalSoundscapesAndWaterSounds	28.5	55.84	25.5	46.23
Audio	31.66	53.47	28.5	42.77
DialogueActClassification_DailyTalk	42.35	33.17	23.62	40.7
DialogueActPairing_DailyTalk	53.77	40.1	41.41	47.24
SarcasmDetection_Mustard	46.5	47.5	50.75	66
Semantics	47.54	40.26	38.59	51.31
EnhancementDetection_LibriTTS-TestClean_WHAM	28.93	52.26	48.5	82.5
NoiseDetection_LJSpeech_MUSAN-Gaussian	49	49.5	52.5	72.5
NoiseDetection_LJSpeech_MUSAN-Music	49.75	46	51.28	86.93
NoiseDetection_LJSpeech_MUSAN-Noise	50	49	52.31	86.5
NoiseDetection_LJSpeech_MUSAN-Speech	43	46.5	52.55	87.94
NoiseDetection_VCTK-MUSAN-Gaussian	57	54.5	49	74
NoiseDetection_VCTK_MUSAN-Music	47.72	54.5	51.52	84.92
NoiseDetection_VCTK_MUSAN-Noise	45.5	58.29	47.98	89
NoiseDetection_VCTK_MUSAN-Speech	36.5	55.5	52.28	87.5
NoiseSNRLevelPrediction_VCTK_MUSAN-Gaussian	25.13	22.8	26.63	21.11
NoiseSNRLevelPrediction_VCTK_MUSAN-Music	26.74	16.67	30.15	25.5
NoiseSNRLevelPrediction_VCTK_MUSAN-Noise	29.29	18.18	28.28	23
NoiseSNRLevelPrediction_VCTK_MUSAN-Speech	23.5	28.21	28.43	25.13
ReverberationDetection_LJSpeech_RirsNoises-LargeRoom	9	47.5	45.5	65.5
ReverberationDetection_LJSpeech_RirsNoises-MediumRoom	8.54	48	51.5	54.82
ReverberationDetection_LJSpeech_RirsNoises-SmallRoom	3.5	48	50.51	57.79
ReverberationDetection_VCTK_RirsNoises-LargeRoom	12.5	46	54.5	67.84
ReverberationDetection_VCTK_RirsNoises-MediumRoom	15	46	48.24	60
ReverberationDetection_VCTK_RirsNoises-SmallRoom	10	46	50.51	56
Degradation	30.03	43.86	45.9	63.6
LanguageIdentification_VoxForge	31.5	89.95	50	88.5
SpeechCommandRecognition_GoogleSpeechCommandsV1	18.88	67.19	70.56	64.32
SpeechDetection_LJSpeech	57.87	40.7	67	62.28
SpeechDetection_LibriSpeech-TestClean	45.73	53.77	49.49	49.95
SpeechDetection_LibriSpeech-TestOther	46.5	48.5	50.75	45.73
SpeechTextMatching_LJSpeech	51.26	80.2	88.72	78.28
SpeechTextMatching_LibriSpeech-TestClean	52.26	86.36	83.94	72.96
SpeechTextMatching_LibriSpeech-TestOther	52.53	83.33	84.9	76.77
SpokenTermDetection_LJSpeech	51.27	81.87	87.05	97.97
SpokenTermDetection_LibriSpeech-TestClean	47.21	73.22	81.44	97.85
SpokenTermDetection_LibriSpeech-TestOther	41.75	71.2	81.35	97.3
Content	45.16	70.57	72.29	75.63
MultiSpeakerDetection_LibriSpeech-TestClean	13.2	44.95	47.24	63.13
MultiSpeakerDetection_VCTK	39.39	45.41	47.74	73.5
SpeakerCounting_LibriTTS-TestClean	9.18	24.37	21	17.17
SpeakerVerification_LibriSpeech-TestClean	44.72	51.5	41.71	46.73
SpeakerVerification_VCTK	52.76	51.5	45.73	38.19
Speaker	31.85	43.55	40.68	47.74

F.3 Limitations of HuBERT codes

To validate the efficacy of discretized HuBERT codes in preserving speech characteristics, we conducted an experiment in which we synthesized a waveform from the HuBERT codes extracted from an original audio waveform. We refer to the resulting audio as re-synthesized audio. In this experiment, we used a randomly selected reference speaker signal in the vocoder. Using the same metrics as in Table 6, we compared the speech characteristics of the re-synthesized audio with those of the original audio. Ideally, the re-synthesized audio should match the characteristics of the original audio. However, the results, shown in Table 18, indicate that while speaking rate exhibited some similarity, other characteristics differed considerably.

In a further experiment, we assessed the effect of the reference speaker signal used in the vocoder. First, we selected original audio samples from the source speech corpus, whose metadata was used to construct a given controllable generation instruction. We then compared two outputs: (1) the audio re-synthesized from HuBERT codes extracted from the original audio, and (2) the audio generated by SIFT-LLM GEN for the corresponding instruction. For each example, the same reference speaker signal was used in the vocoder. The results, reported in Table 19, show that pitch variation and intensity exhibit a much stronger correlation when the same reference speaker is used. We also observed that gender predictions

matched approximately 96% between the SIFT-LLM GEN generated audio and the re-synthesized audio, whereas the match was only around 50% when compared against the gender specified in the controllable instruction. We hypothesize this limitation is due to incomplete representation of speech in HuBERT codes—i.e, they primarily capture semantic content and limited stationary aspects such as speaking rate, pitch, and intensity.

Metric	Pitch variation	Speaking rate	Intensity
MAE (↓)	0.99	0.24	0.32
QWK (↑)	0.11	0.72	0.06
ρ (↑)	0.12	0.79	0.09

Table 18: Evaluation results on controllable generation, comparing the speech characteristics of resynthesized audio and original audio.

Metric	Pitch variation	Speaking rate	Intensity
MAE (↓)	0.43	0.95	0.06
QWK (↑)	0.62	0.23	0.10
ρ (↑)	0.60	0.48	0.15

Table 19: Evaluation results on controllable generation, comparing the speech characteristics of audio generated by SIFT-LLM GEN and resynthesized audio.

G Prompt Templates

G.1 Prompts used for Data Generation

Acoustic/ Content level

Based on the following audio clip, generate \$noques different types of questions and corresponding step-by-step answers. Questions should be about the audio (eg. asking about its properties). The question should be framed as if someone has uploaded an audio (and not its metadata) and asking question about it. It should be answered based on metadata values but answer should not mention about metadata being available or so. It should reflect as if it is answered after listening to the audio. The more complex and diverse the question, the better. Format each QA pair in a single line as a JSON dictionary (key “q” for question, and “a” for answer, wrapped with { and }). Do not include any other explanation.

Metadata: \$metadata

Word-Align

Based on the following audio clip, generate \$noques different types of questions and corresponding step-by-step answers in \$language language. The metadata is provided in \$language language. Questions should be about the audio (eg. asking about its properties). The question should be framed as if someone has uploaded an audio (and not its metadata) and asking question about it. It should be answered based on metadata values but answer should not mention about metadata being available or so. It should reflect as if it is answered after listening to the audio. The more complex and diverse the question, the better. The metadata is a list of word level characteristics like position and possible pitch and intensity categories. Questions can be framed on highlighting considerable changes in pitch and intensity at phrase level, each phrase is combination of consecutive words, determined by position info for each word. Format each QA pair in a single line as a JSON dictionary (key “q” for question, and “a” for answer, wrapped with { and }). Do not include any other explanation.

Metadata: \$metadata

Controllable Generation

You will be provided with metadata describing an audio sample of a person's speech. Your task is to generate a prompt in \$language language that instructs a model to synthesize speech based on a description derived from the given metadata.

Requirements:

- The generated prompt must be grammatically correct, concise, and easy to understand.
- It should not reference the transcription, which is provided separately.
- Aim for complexity and diversity by selecting random speech characteristics to introduce variability.
- Do not add details beyond what is provided in the metadata.

Return only the generated prompt and do not include any other explanation.

Metadata: \$metadata

G.2 Prompts used for Evaluation

Closed-ended Evaluation

Below is a question and the correct/reference answer about an audio:

```
{ "question": "$question", "answer": "$answer" }
```

Rate the following answer as correct or not. If the above answer is correct, "score" should be 1, else 0. Answer: \$hyp

Your output should be a JSON dictionary with keys as "reason" and "score". Do not include preambles, or any additional content.

Open-ended Evaluation

Below is a question and the correct/reference answer about an audio:

```
{ "question": "$question", "answer": "$answer" }
```

Rate the following answer on a scale of 0 to 10. Answer: "\$hyp"

Your output should be a JSON dictionary with keys as "reason" and "score". Do not include preambles, or any additional content.

H Examples of SIFT-LLM's reasoning abilities

Question 1: What might be the possible scenario in which this speech was delivered?

(Audio: We want to continue to be on the periphery of the European infrastructure system, besides the sins of our people.)

Reference Answer: One possible scenario in which this speech could have been delivered is during a conference or meeting where the speaker is discussing their organization's stance on the European infrastructure system.

Model Response: Given the formal tone, international accent, and specific vocabulary related to infrastructure and Europe, this speech could have been delivered in an academic or professional setting, possibly at a conference or seminar.

Question 2: What is the speaker's concern in the speech?

(Audio: How do I strengthen it all in an acceptable form that can be done in two or two and a half hours?)

Reference Answer: The speaker's concern is to figure out a way to achieve the task within the given time frame.

Model Response: The speaker is concerned about the effort and time required to achieve a certain form or appearance. Specifically, they are questioning how someone could accomplish this in two or two-and-a-half hours.

Question 3: Evaluate the effectiveness of the audio for use in a professional setting and propose modifications to enhance its suitability.

Reference Answer: The audio is ineffective for a professional setting due to its poor speech quality and slow delivery. Modifications such as re-recording with better equipment, increasing speaking rate, and reducing background noise would enhance its suitability.

Model Response: The audio is not suitable for a professional setting due to its very bad speech quality, noisy background, and distant-sounding reverberation. To enhance its suitability, the speaker should record in a quiet environment with minimal background noise. Using a high-quality microphone and ensuring proper distance from the microphone can reduce reverberations. Additionally, post-production editing could be used to clean up the audio and improve clarity.