

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 LLASO: A REPRODUCIBLE FOUNDATION FOR LARGE SPEECH-LANGUAGE MODELS

005 **Anonymous authors**

006 Paper under double-blind review

ABSTRACT

011 The development of Large Speech-Language Models (LSLMs) has been lim-
012 ited by fragmented architectures and poor transparency, making reproducibil-
013 ity and fair comparison difficult. In contrast to the vision–language domain,
014 where open resources have driven rapid progress, LSLMs are often released
015 only as model weights without their training data or configurations, leaving the
016 field without common baselines. We present LLaSO, the first fully open, end-
017 to-end framework for large-scale speech–language modeling. LLaSO consists
018 of three key components: (1) LLaSO-Align, a 12M-instance speech–text align-
019 ment corpus; (2) LLaSO-Instruct, a 13.5M-instance multi-task instruction-tuning
020 dataset for speech–text understanding; and (3) LLaSO-Eval, a standardized, re-
021 producible benchmark for cross-modal evaluation. To demonstrate its utility, we
022 train LLaSO-Base, a 3.8B-parameter reference model built entirely on public
023 data. LLaSO-Base achieves a normalized score of 0.72, outperforming compar-
024 able models and providing a strong, reproducible baseline. Our analysis fur-
025 ther shows that while broader training coverage improves performance, signif-
026 icant generalization gaps remain, especially in speech-only scenarios. By releasing
027 datasets, benchmarks, and models together, LLaSO establishes an open standard
028 for LSLMs, enabling unified research and faster community progress.

1 INTRODUCTION

032 The remarkable success of Large Language Models (LLMs) has established a powerful foundation
033 for multimodal AI OpenAI (2024); Yang et al. (2025). In the visual domain, this has led to the
034 rapid maturation of Large Vision-Language Models (LVLMs), where established paradigms, such as
035 leveraging CLIP-style encoders Radford et al. (2021), have enabled effective and scalable alignment
036 between vision and text Awadalla et al. (2023); Wang et al. (2024); Bai et al. (2025); Cocchi et al.
037 (2025). In contrast, the development of Large Speech-Language Models (LSLMs) remains in a more
038 nascent and fragmented stage. The field currently lacks consensus on fundamental architectural
039 principles, with competing approaches that include external feature fusion Radford et al. (2022);
040 Li et al. (2023b), dedicated cross-modal attention mechanisms Kong et al. (2024b); Elizalde et al.
041 (2024), and implicit alignment strategies Chu et al. (2024).

042 This architectural divergence is compounded by a lack of transparency in existing research. While
043 several open-source LSLM initiatives have emerged Chu et al. (2023); Défossez et al. (2024); Tang
044 et al. (2024), many are only partially open. Model weights may be released, but the underlying
045 training data and crucial configurations are often withheld. This opacity makes it difficult to con-
046 duct fair comparisons, as performance differences can be attributed as much to proprietary data or
047 undisclosed training strategies as to architectural merit, hindering systematic progress.

048 To address these challenges of fragmentation and opacity, we introduce LLaSO: a fully open, end-
049 to-end framework designed to establish foundational standards for LSLM research. LLaSO consists
050 of three core, publicly available components:

- 051 1. **LLaSO-Align:** A 12M-instance speech–text alignment corpus aggregated from diverse
052 sources, including conversational speech Chen et al. (2021), read narratives Panayotov et al.
053 (2015), audio books Ito & Johnson (2017); Pratap et al. (2020), and accented speech Veaux
et al. (2016).

054

055 **2. LLaSO-Instruct:** A 13.5M-instance instruction-tuning dataset covering 20 tasks across

056 linguistic, semantic, and paralinguistic domains. It supports three distinct modality config-

057 urations: audio instructions with audio inputs, textual instructions with audio inputs, and

058 audio instructions with textual inputs.

059 **3. LLaSO-Eval:** A reproducible benchmark of 15,044 stratified samples designed for com-

060 prehensive evaluation of instruction-following capabilities of LSLMs.

061 To validate our framework and provide the community with a strong, reproducible baseline, we

062 developed LLaSO-Base, a 3.8B-parameter reference model that adapts the successful LLaVA

063 architecture to the speech domain. Trained exclusively on LLaSO-Align and LLaSO-Instruct, and

064 evaluated on LLaSO-Eval, our model achieves a normalized score of 0.72, outperforming the next

065 best comparable model (0.65). As illustrated in Figures 1 (Middle and Right), LLaSO-Base is

066 designed not for state-of-the-art performance, but to demonstrate the power of an open, extensible, and

067 reproducible workflow.

068 Our evaluation shows that while broader training improves overall performance, models still struggle

069 with generalization, leaving substantial gaps on unseen tasks and pure audio settings. Investigating

070 potential solutions for this weakness, we found that models equipped with interleaving and parallel

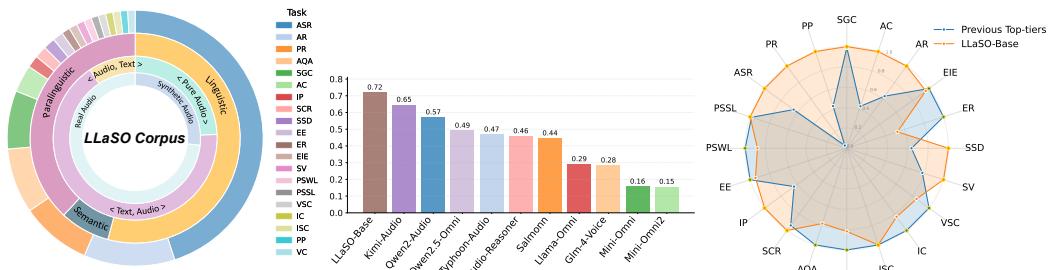
071 decoding mechanisms exhibit far greater robustness in these challenging scenarios.

072 In summary, LLaSO provides the first fully open, end-to-end stack for LSLM research, comprising

073 large-scale training datasets, a standardized benchmark, and a reference model. By releasing these

074 resources, we aim to lower the barrier to entry and foster a new wave of systematic, community-

075 driven progress in large-scale speech-language modeling.



085 Figure 1: (Left) LLaSO Corpus Overview: 25.5M audio-text pairs over 20 tasks (18 paralinguistic), integrating

086 LLaSO-Align, LLaSO-Instruct, and LLaSO-Eval, with 73% real and 27% synthetic audio (further statistics are

087 detailed in Appendix Q and R). (Middle) Overall model performance after min-max normalization for direct

088 comparison where higher bars indicate better overall performance. (Right) Normalized task-level results on

089 LLaSO-Eval: LLaSO-Base (orange) vs. leading baselines (blue) across 20 tasks, with scores scaled by setting

090 the best model to 1 (detailed results are provided in Appendix S).

091 2 RELATED WORK

093 **Vision-Language Modeling.** Vision-language modeling has rapidly advanced through a standard-
 094 ized two-stage paradigm: modality alignment followed by instruction tuning Brown et al. (2020);
 095 Bommasani et al. (2021); Li et al. (2023c). The rapid progress in this field has been facilitated by
 096 two essential types of open resources. First, public training datasets and standardized evaluation
 097 benchmarks Ma et al. (2023); Hsieh et al. (2023); Zeng et al. (2024b); Fu et al. (2024); Huang
 098 et al. (2025) have become widely adopted, enabling fair comparison and transparent reproducibil-
 099 ity across models and tasks. Second, open-source implementations with modular codebases such
 100 as LLaVA Lin et al. (2023) and OpenFlamingo Awadalla et al. (2023) have significantly lowered
 101 the technical barriers to development and fostered rapid iteration across the community. Together,
 102 these practices have fostered a shared research infrastructure where new models and tasks are often
 103 built upon existing resources Liu et al. (2023); Yin et al. (2024). This has allowed vision-language
 104 research to focus more on advancing scientific capabilities rather than reimplementing foundational
 105 components.

106 **Speech-Language Modeling.** Compared to vision-language modeling, progress in speech-
 107 language systems has been less cohesive Su et al. (2025); Ma et al. (2024). First, most leading
 108 models such as Audio Flamingo Kong et al. (2024a); Ghosh et al. (2025), Qwen-Audio Chu et al.

(2023; 2024), and Kimi-Audio KimiTeam et al. (2025) rely on proprietary data, limiting reproducibility Peng et al. (2025); Pandey et al. (2025). Second, most models support only narrow modality configurations (e.g., text-plus-audio), with few addressing more compositional tasks Tang et al. (2024); Chu et al. (2024); Chen et al. (2024). Third, existing datasets largely focus on semantic reasoning Fang et al. (2025); Wu et al. (2024); Mei et al. (2024), with limited coverage of prosody and emotion. Lastly, few open-source stacks unify models, datasets, and benchmarks; most systems (e.g., LauraGPT Du et al. (2023), Moshi Défossé et al. (2024), Westlake-Omni Xinchen-ai (2024)) lack full releases, hindering reproducibility and community development.

Name	Alignment Data	Alignment Tasks	Task Coverage	Modality Coverage	Audio Type	Sample Num.	Duration (Hours)
AVQA	✗	-	1	①	Collected	~57.3K	-
COTA	✗	-	5	①	Mixed	~1.2M	-
OpenAQA	✓	Multiple	4	①	Collected	~5.0M	-
OpenASQA	✓	Single	8	①	Collected	~9.6M	-
SIFT-50M	✓	Multiple	10	①	Collected	~55.6M	-
SALMONN	✓	Multiple	14	②	Collected	~2.3M	~4.4K
LLaSO Corpus	✓	Single	20	③	Mixed	~25.5M	~89.5K

Table 1: Comparison of public speech-language datasets and LLaSO Corpus. For “Modality Coverage,” ① means only text instruction with audio input, ② adds pure audio formats, and ③ indicates full support, including audio instruction with text input. “Audio Type” denotes real (“Collected”), synthetic, or mixed data. ✓/✗ show whether alignment data are presented.

3 LLaSO CORPUS

To support the development of LSLMs, we introduce the *LLaSO Corpus*, a comprehensive, modular benchmark suite.

3.1 CORPUS OVERVIEW

Inspired by practices in LVLMs, LLaSO comprises three tightly integrated components:

- *LLaSO-Align*: A large-scale corpus for aligning speech with semantic space through ASR-based supervision.
- *LLaSO-Instruct*: A multi-task instruction-tuning dataset spanning linguistic, semantic, and paralinguistic tasks.
- *LLaSO-Eval*: A stratified benchmark designed for consistent evaluation across tasks.

These components together support the full training pipeline of LSLMs, modality alignment, instruction tuning, and evaluation (see Figure 1 (Left)).

To advance LSLMs beyond vision-language paradigms, we anchor our benchmark design in two core properties of speech:

- *Inherent Paralinguistics*: Speech conveys rich, essential information beyond words such as speaker identity, accent, emotion, and prosody. These paralinguistic cues are omnipresent and crucial for natural human communication.
- *Flexible Modality Roles*: In LSLMs, both audio and text can serve as inputs or instructions, enabling diverse interaction patterns e.g., audio-instruction with text input, text-instruction with audio input, or audio-instruction with audio input.

To better reflect the needs of real-world systems, we adopt a balanced task weighting approach that corrects for limitations in existing corpora:

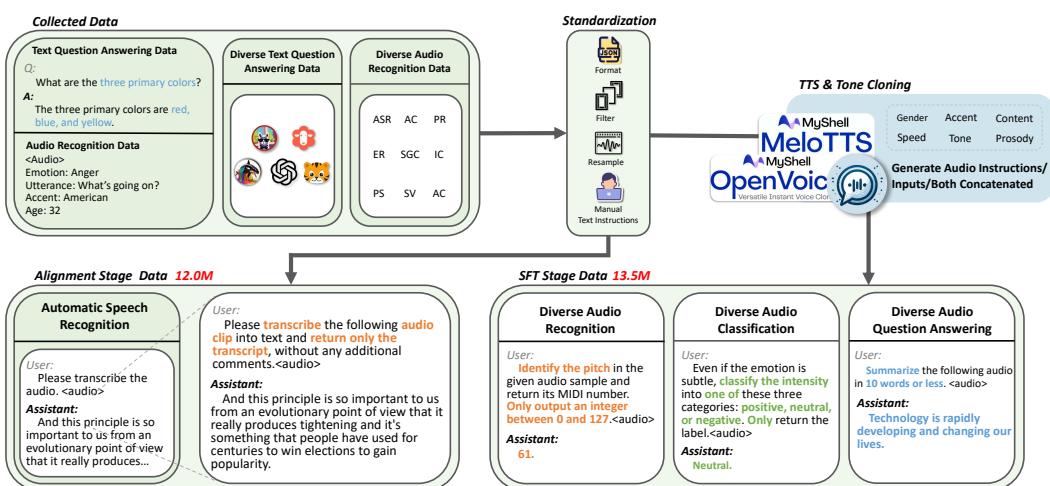
- *Semantic Tasks (8%)*: Intentionally underweighted, as their success often reflects language modeling capacity rather than speech understanding Rouditchenko et al. (2025), and they are already well-represented Gong et al. (2023b); Fang et al. (2025).
- *Paralinguistic Tasks (40%)*: Prioritized to address their underrepresentation in current resources Jiang et al. (2025); yu Huang et al. (2024). We ensure diversity by combining real-world metadata with synthetically generated variations.

162 • *Linguistic Tasks* (52%): Dominated by ASR, which remains foundational to grounding
 163 speech in linguistic structure and is critical for general performance.

164 The final LLaSO Corpus includes 71% real-world audio and 29% synthetic speech, and covers a
 165 broad range of modality configurations where both audio and text flexibly act as inputs and in-
 166 structions presented in Table 1. This design ensures robust coverage of the full speech landscape,
 167 supporting the development of unified and adaptable LSLMs.

169 3.2 LLaSO-ALIGN

170 To establish a robust semantic foundation for speech-language modeling, we adopt ASR as the
 171 core alignment task in LLaSO Corpus’s first stage. Following vision-language best practices,
 172 this approach grounds speech representations directly in textual semantic space through explicit
 173 instruction-response pairing. *LLaSO-Align* contains 12M instruction-formatted ASR samples, each
 174 including an audio input, a natural language instruction, and a reference transcript. Unlike traditional
 175 ASR datasets offering only raw audio-text pairs, we introduce 18 hand-crafted instruction
 176 templates that frame the task with varying specificity and constraints, e.g., ”Transcribe the audio
 177 precisely; return text only”, encouraging instruction adherence and realistic use. To ensure diver-
 178 sity in content and speaker profiles in LLaSO-Align, we aggregate wide-range public ASR corpora
 179 spanning conversational speech Chen et al. (2021), single-speaker narration Panayotov et al. (2015),
 180 audio books Ito & Johnson (2017); Pratap et al. (2020), and accented English Veaux et al. (2016),
 181 capturing a broad range of acoustic environments, accents, ages, and speech styles. All samples
 182 undergo a construction pipeline to ensure consistency and quality presented as Figure 2 where its
 183 standardization details in Appendix Q. By reframing ASR as an instruction-following alignment
 184 task and curating a diverse, high-quality dataset, LLaSO-Align lays the groundwork for downstream
 185 speech-language understanding across modalities.



202 Figure 2: LLaSO Corpus construction pipeline. We first aggregate heterogeneous sources, including text-
 203 based QA corpora and speech datasets covering acoustic, paralinguistic, and semantic tasks, then normalize
 204 format, sample rate, and instruction style, etc. We construct LLaSO-Align (12.0 M) for aligning speech and
 205 text modality via ASR, while LLaSO-Instruct (13.5 M) for multi-task instruction tuning including classification,
 206 recognition, and AQA. When synthesize audio, we use advanced audio synthesis as described in Appendix H
 207 for richer speaker variation, enabling *pure-audio*, *text plus audio*, and *audio plus text* formatted samples with
 208 diverse gender, accent, speed, and tone.

209 3.3 LLaSO-INSTRUCT

210 Building on the aligned speech-text representations from LLaSO-Align, we present *LLaSO-Instruct*,
 211 a multi-task instruction tuning dataset designed to advance speech-language modeling with greater
 212 task diversity and richer modality configurations. Unlike previous instruction datasets focused pri-
 213 marily on semantic tasks with limited input modalities, LLaSO-Instruct fully embraces the inher-
 214 ently multimodal and paralinguistic nature of speech, systematically expanding both task coverage
 215 and modality pairings, offering a comprehensive framework to instruction tuning LSLMs.

216 **Task Coverage.** LLaSO-Instruct spans **20 tasks** across linguistic, semantic, and paralinguistic
 217 categories. While linguistic tasks (e.g., ASR) and semantic tasks (e.g., audio-based QA) cover foun-
 218 dational capabilities, the majority of tasks are paralinguistic, including speaker-centric tasks and
 219 content-centric tasks, designed to capture speaker traits and contextual acoustic crucial for socially-
 220 aware interaction, with all included tasks presented in Appendix Q. To construct this wide range
 221 of tasks, firstly, we collect task-specific datasets with rich metadata, enabling reuse of the same
 222 audio sample across multiple tasks, with its associated labels such as accent and gender. When la-
 223 bel distributions are imbalanced, we implement targeted sampling strategies.¹ Secondly, for each
 224 task, we manually construct 20 text instructions across four prompt styles including standardized,
 225 contextualized, stylistic, and fine-grained (examples in Appendix L). For ASR and AQA tasks, in-
 226 structions are open-ended inviting free-form responses from the model, while paralinguistic tasks
 227 predominantly employ closed-ended instructions, requiring the model to select an answer from pre-
 228 defined options without additional analysis. To address diverse task requirements, we construct
 229 training samples at multiple levels of granularity, so that some paralinguistic tasks also include open-
 230 ended variants.² We display task prototypes in Figure 3 and the open/closed-ended mapping in Ap-
 231 pendix Q. As to linguistic and semantic task categories, we will discuss in the following paragraph.
 232

233 **Modality Coverage.** To reflect speech’s flex-
 234 ible role as both input and instruction, LLaSO-
 235 Instruct supports three core modality config-
 236 urations: **1)** Text instruction with audio input; **2)**
 237 Audio instruction with text input; **3)** Pure au-
 238 dio for both instruction and input. We provide
 239 three examples in Figure 4 and an overview
 240 in Table 11. Linguistic tasks retain their na-
 241 tive modality pairings, with one million ASR
 242 samples carried over from LLaSO-Align. Se-
 243 mantic QA tasks are derived from high-quality
 244 text datasets (e.g., OpenOrca and Alpaca) and
 245 converted into multimodal samples using audio
 246 synthesis (see Appendix H), thus each instance
 247 may yield multiple variants (e.g., text-with-
 248 audio, audio-with-text) to support cross-modal
 249 learning.³ Similarly, paralinguistic tasks are
 250 primarily configured as text-instruction with
 251 audio input, but where feasible, we also con-
 252 struct fully speech-driven formats by synthe-
 253 sizing both instruction and input as audio, en-
 254 abling training in pure audio scenarios better
 255 simulating human-human interaction.

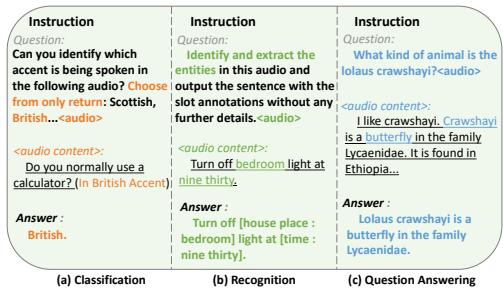
256 3.4 LLaSO-EVAL

257 To complete our data trio, we introduce LLaSO-Eval, a held-out, training-disjoint evaluation suite
 258 designed to accompany the LLaSO training set. Derived from the same underlying corpus but sepa-
 259 rate from the training split, LLaSO-Eval covers 15,044 samples across 20 tasks, categorized into
 260 linguistic, semantic, and paralinguistic categories and supports all three modality configurations to
 261 test both within- and cross-modal generalization. The suite includes open-ended prompts for free-
 262 form comprehension/reasoning and closed-ended prompts enabling quantify instruction following
 263 capability via abstention rate. A task-level breakdown is provided in Appendix Q and R.

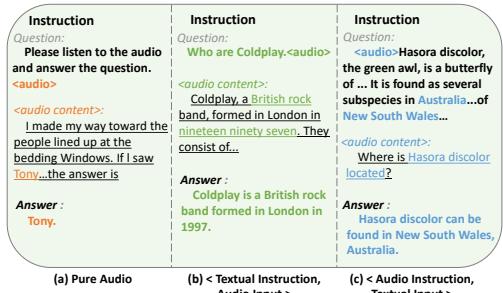
264 ¹For example, in the Meld accent dataset, we address the long-tail by removing rare accents and trimming
 265 dominant ones. Similarly, we repurpose VCTK’s gender metadata for speaker classification, balancing the
 266 dataset by downsampling the female subset to 1:1 ratio.

267 ²For example, in age classification, we use three levels: coarse-grained (e.g., “twenties”, “fifties”), medium-
 268 grained (e.g., “15-19”, “20-24”), and fine-grained where the model is required to predict the exact age as an
 269 integer between 18 and 80.

³When instruction and input segments are suitable for audio synthesis (English-only, properly normalized,
 269 and error-free), each textual QA instance yields both text-with-audio and audio-with-text variants.



264 **Figure 3: Task prototypes.** (a) closed-set *classifi-
 265 cation*; (b) multi-granularity/open-set *recognition*; (c)
 266 open-ended *AQA*.



267 **Figure 4: Interaction formats** in LLaSO Corpus.

270 4 MODEL
271

272 To validate the effectiveness of our LLaSO Corpus, we introduce LLaSO-Base, a reference model
273 in the speech-language domain that strictly aligns with the end-to-end instruction tuning paradigm
274 established in vision-language research Zhu et al. (2023); Cocchi et al. (2025); Li et al. (2023a).
275 Rather than pursuing new SOTA results, our objective is to offer the community a robust and exten-
276 sible baseline for systematic cross-modal instruction following.

277 4.1 MODEL ARCHITECTURE
278

279 Our model follows a simple yet proven three-
280 component design as illustrated in Figure 5
281 that uniformly supports text with audio, audio-
282 only, and audio plus text inputs via em-
283 bedding concatenation. For audio features,
284 we use the Whisper-large-v3 encoder Radford
285 et al. (2022); Zhang et al. (2024); Gong et al.
286 (2023a), retaining only the encoder ($\sim 640M$) to
287 leverage its strong representations while leav-
288 ing generation to the LLM. Audio is processed
289 by Whisper’s front end (16 kHz log-mel, stride-
290 2 ≈ 40 ms/frame) with SpecAugment Park et al.
291 (2019) during training. Final-layer features
292 $Z^a = F_{ae}(X^a)$ (where X^a may denote au-
293 dio instructions, content, or both) are projected
294 into the LLM embedding space by a two-layer
295 multi-layer perceptron (MLP) with Gaussian
296 Error Linear Unit (GELU) activation, $H^a =$
297 $F_{proj}(Z^a)$, chosen for its simplicity and effectiveness
298 over heavier alignment modules Tang et al.
299 (2024); Kong et al. (2024b); Lin et al. (2024).
300 The projected H^a is concatenated with text embed-
301 dings $H_{instruct/input}^t$ from the tokenizer, yielding a unified multimodal sequence with preserving order.
302 The sequence is subsequently processed by Llama-3.2-3B-Instruct Grattafiori et al. (2024), a main-
303 stream instructed backbone. With $\sim 3.8B$ total parameters, LLaSO-Base balances computational
304 efficiency and representational capacity.

305 4.2 TRAINING
306

307 We train the model in a single-turn dialogue setting, where each instance consists of audio X^a , its
308 paired text X^t , and the target response X_{answer}^t . To support different modality configurations, we
309 unify the query format as in Eq. 1.

$$X_{query}^{(t,a)} = [X_{instruct}^t, X_{input}^a], X_{query}^{(a,t)} = [X_{instruct}^a, X_{input}^t], X_{query}^{(a)} = [X_{instruct+input}^a] \quad (1)$$

310 Training optimizes parameters θ via next-token autoregressive prediction, maximizing the condi-
311 tional likelihood of the response given the query, as defined in Eq. 2. We adopt a proven two-stage
312 instruction-tuning paradigm, alignment followed by instruction tuning, with the set of trainable pa-
313 rameters θ varying by stage.

$$p(X_{answer}^t | X_{query}^{(*)}) = \prod_{i=1}^L p_{\theta}(x_i^t | X_{query,<i}^{(*)}, X_{answer,<i}^t), \quad (2)$$

314 where X_{answer}^t denotes the assistant’s text response, $X_{query}^{(*)}$ the input query under any modality
315 configuration in Eq. 1, and L the response length. Beginning with **Alignment Stage**, we use ASR as
316 the alignment objective on LLaSO-Align, where each example contains a text instruction $X_{instruct}^t$,
317 an audio input X_{input}^a , and its transcript X_{answer}^t , optimized with the objective in Eq. 2. During
318 this stage we freeze the speech encoder and the LLM, updating only the projector F_{proj} so that
319 $H^a = F_{proj}(Z^a)$ aligns with the pre-trained LLM word embedding space, thereby establishing
320 cross-modal semantic consistency for the next stage. In **Instruction Tuning Stage**, we then train
321 on LLaSO-Instruct to endow the model with compositional instruction-following across linguistic,
322 semantic, and paralinguistic tasks. The encoder remains frozen while we optimize F_{proj} and F_{llm}
323 under Eq. 2, using the unified query formats in Eq. 1 and always producing textual responses. We
324 provide training details in Appendix K.

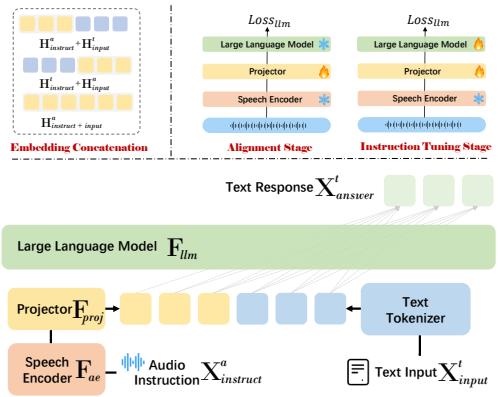


Figure 5: Overview of LLaSO-Base: model architecture and input flow (Bottom), three input layouts (Top Left), and the two-stage training recipe of alignment and instruction tuning (Top Right).

324	Linguistic Task Category										Semantic Task Category											
	325	Modality Format:	< Textual Instruction, Audio Input >					< Pure Audio >					< Textual Instruction, Audio Input >					< Audio Instruction, Textual Input >				
			326	Tasks	ASR					AQA												
327	Qwen2-Audio	0.22	0.12	2.41	2.42	2.73	2.78	2.59	2.56	3.49	2.13	3.14	3.13	2.20	2.82	3.47	3.62	3.29	1.29	3.14	2.52	2.89
328	Typhoon-Audio	0.11	0.06	1.76	1.77	2.16	2.22	1.98	1.87	3.14	1.61	2.83	3.04	2.36	2.60	2.69	2.91	2.47	1.68	3.04	1.91	2.45
329	Salmonn	0.86	0.69	1.47	1.41	1.41	1.72	1.50	2.05	3.13	1.42	2.96	3.12	2.37	2.60	2.04	3.03	2.42	1.83	3.19	1.58	2.35
330	Glm-4-Voice	0.93	0.79	2.22	2.34	3.29	2.93	2.70	2.49	3.21	2.51	3.11	2.82	1.97	2.72	3.09	4.06	1.68	1.03	3.10	1.98	2.49
331	Mini-Omni	0.95	0.81	1.42	1.47	1.75	1.45	1.52	1.63	1.54	1.22	2.34	1.33	1.41	1.57	1.42	1.32	1.17	1.21	1.27	1.20	1.27
332	Mini-Omni2	0.95	0.80	1.57	1.53	2.05	1.51	1.67	1.66	1.64	1.26	2.52	1.42	1.43	1.65	1.68	1.50	1.41	1.29	1.31	1.28	1.41
333	Llama-Omni	0.88	0.73	1.97	2.02	2.99	2.48	2.37	2.38	2.95	1.88	3.16	2.72	2.20	2.58	2.73	3.78	2.29	1.11	3.08	2.09	2.51
334	Audio-Reasoner	0.28	0.12	2.44	2.24	2.51	2.86	2.51	2.22	3.42	2.12	3.07	2.91	2.14	2.73	2.84	3.95	2.88	1.54	3.13	2.09	2.74
335	Kimi-Audio	0.14	0.05	2.94	2.70	3.22	3.45	3.08	3.28	3.77	3.35	3.38	3.27	2.71	3.35	3.69	4.01	3.38	1.16	3.16	2.77	3.03
336	Qwen2.5-Omni	0.40	0.26	2.94	3.09	3.22	2.63	2.97	2.99	3.80	3.20	2.96	3.19	2.12	3.05	3.46	3.88	3.58	1.19	3.15	2.42	2.95
337	LLaSO-Base (Ours)	0.08	0.03	2.06	1.80	2.39	1.46	1.93	2.57	2.48	1.71	2.74	3.05	2.90	2.58	2.72	2.62	2.28	2.23	3.74	2.60	2.70
338	Metrics	WER↓ CER↓		GPT-4o↑ Avg.GPT-4o↑		GPT-4o↑ Avg.GPT-4o↑		GPT-4o↑ Avg.GPT-4o↑		GPT-4o↑ Avg.GPT-4o↑		GPT-4o↑ Avg.GPT-4o↑		GPT-4o↑ Avg.GPT-4o↑		GPT-4o↑ Avg.GPT-4o↑		GPT-4o↑ Avg.GPT-4o↑		GPT-4o↑ Avg.GPT-4o↑		

Table 2: Comparison of 11 LSLMs on LLaSO-Eval linguistic (ASR) and semantic (AQA) tasks across three modality configurations. ASR is evaluated by WER/CER (lower ↓ is better); AQA is scored by GPT-4o (higher ↑ is better). Cell shading , , and no shading denotes each model’s relative ranking in a given modality (best → worst) by average GPT-4o score.

339	Paralinguistic Task Category										Paralinguistic Task Category																				
	340	Speaker-centric										Content-centric																			
		341	Modality Format: < Textual Instruction, Audio Input >					Modality Format: < Textual Instruction, Audio Input >					Modality Format: < Textual Instruction, Audio Input >					Modality Format: < Textual Instruction, Audio Input >													
342	Qwen2-Audio	1.00	0.95	0.67	0.99	0.16	0.12	0.05	0.23	0.52	18.69	0.54	0.31	0.24	0.30	0.43	0.25	1.86	1.95	0.17	1.19	2.60	2.40	3.04	2.52	2.73	0.85	0.60	0.60	19.02	0.02
343	Typhoon-Audio	0.85	0.77	0.59	0.67	0.21	0.14	0.11	0.10	0.12	20.47	0.40	0.24	0.28	0.12	0.46	0.20	2.04	1.71	0.33	3.08	0.98	0.85	3.13	1.86	2.85	0.49	0.16	0.16	36.83	0.17
344	Salmonn	0.59	0.44	0.13	0.18	0.22	0.32	0.10	0.26	0.06	11.24	0.31	0.24	0.30	0.21	0.50	0.19	1.32	1.38	0.13	1.82	1.09	0.75	4.07	1.88	3.29	0.61	0.16	0.16	41.92	0.22
345	Glm-4-Voice	0.11	0.12	0.04	0.07	0.07	0.09	0.03	0.02	0.01	15.35	0.13	0.08	0.14	0.02	0.10	0.04	1.62	1.84	0.24	0.90	1.00	0.98	1.85	1.78	2.34	0.32	0.00	0.03	40.20	0.08
346	Mini-Omni	0.14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	21.34	0.04	0.04	0.04	0.07	0.11	0.03	1.24	1.46	0.00	0.92	1.00	0.98	1.39	1.26	1.42	0.02	0.01	0.06	61.49	0.00
347	Mini-Omni2	0.11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	18.46	0.03	0.06	0.00	0.01	0.12	0.03	1.16	1.54	0.00	0.97	1.00	0.97	1.89	1.26	1.60	0.03	0.01	0.03	59.32	0.00
348	Llama-Omni	0.36	0.26	0.03	0.26	0.07	0.14	0.07	0.16	0.17	<i>Reject</i>	0.31	0.08	0.16	0.04	0.26	0.08	1.28	1.38	0.13	1.46	20.19	21.11	1.58	1.80	2.34	0.03	0.00	0.12	<i>Reject</i>	0.07
349	Audio-Reasoner	0.38	0.32	0.38	0.37	0.14	0.18	0.02	0.23	0.12	13.57	0.52	0.29	0.32	0.28	0.35	0.16	2.61	2.03	0.09	1.08	2.40	1.84	4.10	2.29	3.78	0.59	0.20	0.17	32.68	0.15
350	Kimi-Audio	0.98	0.97	0.66	0.81	0.38	0.31	0.20	0.17	0.12	12.07	0.65	0.34	0.52	0.32	0.63	0.22	3.30	2.76	0.20	1.58	1.00	0.31	4.56	2.05	3.57	0.84	0.26	0.38	31.64	0.19
351	Qwen2.5-Omni	0.53	0.41	0.40	0.35	0.06	0.19	0.02	0.11	0.08	10.31	0.52	0.27	0.29	0.33	0.42	0.15	1.25	2.12	0.27	1.28	3.53	3.52	3.91	2.00	2.78	0.92	0.51	0.44	18.37	0.12
352	LLaSO-Base	0.96	0.99	0.76	0.91	0.52	0.83	0.73	0.70	0.50	10.32	0.48	0.48	0.17	0.26	0.99	0.32	2.90	2.80	0.39	0.03	0.04	0.02	4.86	3.93	3.57	0.78	0.50	0.60	8.02	0.18
353	Metrics	ACC↑		MAE↓		ACC↑		GPT-4o↑		ACC↑		PER↓		WER↓		CER↓		GPT-4o↑		ACC↑		MAE↓									

abstention

Table 3: Performance of 11 LSLMs on LLaSO-Eval paralinguistic tasks, split by speaker-centric and content-centric groups. Cells are colored by abstention rate, as indicated by the color bar. Abstention rates were computed across all closed-ended tasks. Results are for the text instruction with audio input modality, reflecting the modality used for paralinguistic task training; our released datasets also include pure audio modality format samples for these tasks. *Reject* denotes 95% or more abstentions in a given task after manual inspection in open-ended settings/tasks.

5 EXPERIMENTS

5.1 SETUP

All experiments are conducted on LLaSO-Eval using the splits and task configurations defined in Section 3.4 and Appendix R. We benchmark LLaSO-Base against representative speech-language models, including Qwen2-Audio Chu et al. (2024), Typhoon-Audio Manakul et al. (2024), Salmonn Tang et al. (2024), GLM-4-Voice Zeng et al. (2024a), Mini-Omni Xie & Wu (2024b), Mini-Omni2 Xie & Wu (2024a), Llama-Omni Fang et al. (2025), Audio-Reasoner Xie et al. (2025), Kimi-Audio KimiTeam et al. (2025), and Qwen2.5-Omni Xu et al. (2025), running official checkpoints. Detailed versions and access methods for all baselines are provided in Appendix I.

5.2 METRICS

LLaSO-Eval spans a wide range of tasks, which we categorize into open-ended and closed-ended formats, defined as in Table 11. To support this diversity, we employ 7 evaluation metrics, with task-metric assignments summarized in Table 12. By default, metric scores are computed directly from the raw model outputs. For tasks where models generate free-form responses, we apply an answer-extraction step prior to evaluation. For open-ended tasks evaluation, we primarily use the GPT-4o Score, ranges from integer 1 to 5, providing a holistic assessment of response quality. In some cases, traditional metrics are also used to supplement GPT-4o. It is worth noting that certain tasks with well-established evaluation standards or clearly defined targets are better served by traditional metrics like PER or MAE. For closed-ended tasks, we use Accuracy and additionally track Abstention Rate for invalid or noncompliant answers. Complete metric details are provided in Appendix C.

378 5.3 RESULTS AND ANALYSIS
379

380 For cross-metric comparability, we report normalized
381 overall and per-task performance in
382 Figure 1 (Middle and Right), with full per-task
383 scores and closed-ended abstention rates de-
384 tailed in Table 2 and 3. LLaSO-Base attains the
385 highest normalized overall score (0.72 vs. 0.65
386 for the next best model) and performs better on
387 most individual tasks. The results establish the
388 LLaSO Corpus as a reliable source for building
389 reference models. LLaSO-Base, while not de-
390 signed for state-of-the-art performance, exem-
391 plifies the power of an open, extensible, and re-
392 producible workflow.

392 **Broader training coverage improves overall
393 quality.** For deeper analysis we visualize Fig-
394 ure 6 comparing the performance and absti-
395 ntion rates of 11 evaluated models as a function
396 of task coverage defined as the number of train-
397 ing tasks; for models with private, incomplete,
398 or ambiguously reported training data, we use
399 the number of evaluation tasks as a proxy. As
400 can be seen, models trained on broader task
401 sets consistently outperform task-focused sys-
402 tems in both overall and closed-ended perfor-
403 mance, while also exhibiting fewer abstentions.
404 This pattern highlights persistent difficulties on
405 unseen tasks and underscores the importance of
406 diversifying task coverage to improve perfor-
407 mance and lower abstention rates.

408 **LSLMs may prefer content related tasks.** To
409 further analyze model performance within par-
410 alinguistic tasks, we present a dumbbell plot
411 in Figure 7, contrasting content-centric and
412 speaker-centric results for each model. We ob-
413 serve that most models achieve higher perfor-
414 mance and lower abstention rates on content-
415 centric tasks than on speaker-centric ones. This
416 disparity likely arises because content-centric
417 tasks are more tightly linked to the semantic
418 content, which LLM-based decoders are nat-
419 urally equipped to process. In contrast, speaker-centric
420 tasks demand more nuanced inference of
latent speaker attributes, posing a greater challenge for current LSLMs and highlighting an impor-
421 tant area for future improvement.

422 **Generalization remains fragile, especially on unseen modalities.** Extending the observation from
423 task coverage to modality, we find that most models underperform when evaluated across the three
424 modality settings in Table 2. This weakness is unsurprising given that many baselines, as detailed
425 in Appendix I, were trained to support only one or two input-output formats rather than the full
426 spectrum. In particular, the `<audio instruction, text input>` configuration consistently lags behind
427 the more common `<text instruction, audio input>`, even though the former should in principle be
428 no harder for humans, as speech instructions are typically brief and the main content remains di-
429 rectly readable as text. A notable exception is Qwen2-Audio and its variant Audio-Reasoner, which
430 achieve comparable results across the two formats.

431 **Pure audio remains the most challenging modality.** To better visualize the tendency, we present
Figure 8 (Bottom). Even more striking, we observe that most models perform substantially *worse*
on `<pure audio>` inputs than on `<text instruction, audio input>`, even for systems explicitly trained

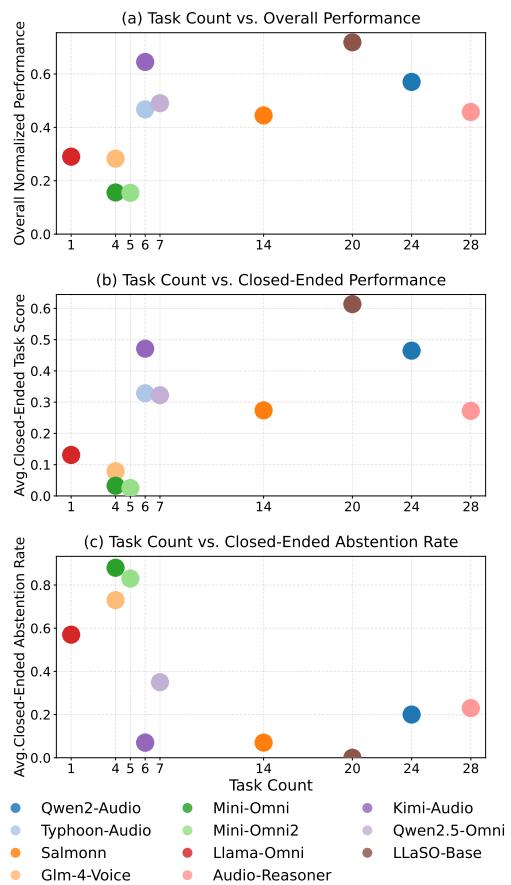


Figure 6: Task coverage vs. model performance and abstention. Each scatter plot shows 11 models by *Task Count*. (a) Overall performance (min–max normalized over all LLaSO-Eval tasks, cf. Figure 1 (Middle)). (b) Average Closed-ended task performance. (c) Average abstention rate on closed-ended tasks. Closed-ended scores and abstention rates are calculate only on tasks that require categorical selection. Higher scores indicate better performance; lower abstention rates indicate stronger instruction following performance.

8

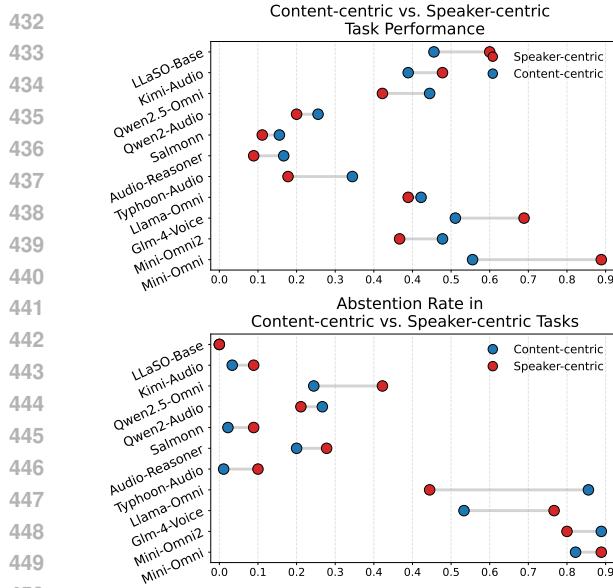


Figure 7: Comparison of LSLM performance on content-centric versus speaker-centric paralinguistic tasks. *Top*: For each model, min-max normalized performance scores are shown on content-centric (blue) and speaker-centric (red) tasks, with dumbbell lines indicating the magnitude and direction of intra-model performance differences. *Bottom*: Average abstention rates (lower is better) for closed-ended tasks in the same two centrics within the paralinguistic category. All evaluations are conducted under the text instruction paired with audio input configuration.

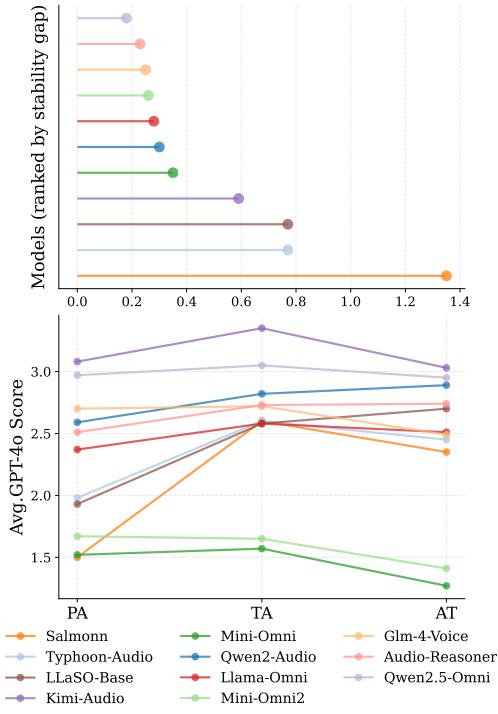


Figure 8: Stability and modality-wise performance of LSLMs. *Top*: Cross-modality stability, measured as the sum of absolute differences between a model’s GPT-4o score on text + audio (TA) and the other two formats, pure audio (PA) and audio + text (AT); lower values indicate greater robustness. *Bottom*: Average GPT-4o scores across the three configurations. Colors are consistent across plots.

on speech-to-speech or spoken-query QA, with drops sometimes exceeding those on unseen configurations. Interestingly, a few models such as Qwen2.5-Omni, GLM-4-Voice, and the Mini-Omni family achieve comparable performance across modalities. To quantify this, we measure cross-modality stability as the sum of absolute performance differences between the common text with audio setting and the other two formats, i.e., $|TA - PA| + |TA - AT|$, and report models in ascending order of stability in Figure 8 (Top). We find that **interleaving** and **parallel decoding** substantially reduce modality gaps, where the top 8 among 11 systems, excluding Qwen2-Audio and its variant, adopt these strategies and exhibit notably smaller disparities. Although the outliers likely reflect factors beyond modality combination design, These results highlight interleaving and parallel decoding as promising directions for improving cross-modal generalization. Further modality- and task-level analyses, case studies across different task and modality, and model-specific discussions are provided in Appendix D, F, and G.

6 CONCLUSION

Despite recent advances, progress in LSLMs has been constrained by fragmented resources, limited task diversity, and a lack of standardized evaluation. To address these challenges, we present LLaSO: the first fully open, end-to-end framework for LSLM development. It comprises 25.5M samples for alignment and instruction tuning (LLaSO-Align and LLaSO-Instruct), a stratified benchmark of 15K samples (LLaSO-Eval), and a reproducible 3.8B-parameter reference model (LLaSO-Base) successfully verified proven vision-language architecture to speech domain. Our evaluation further reveals that, although broader task coverage improves overall performance, current LSLMs still face notable generalization gaps on unseen tasks and pure-audio settings. Encouragingly, models employing interleaving or parallel decoding demonstrate improved robustness in these challenging scenarios, highlighting promising directions for future research. By releasing all data, benchmarks, and models, LLaSO lowers the barrier to entry and provides a foundation for systematic, community-driven progress in large-scale speech-language modeling. *See Supplementary Material for reproducibility.*

486 REFERENCES
487488 2noise. GitHub - 2noise/ChatTTS: A generative speech model for daily dialogue. — github.com.
489 <https://github.com/2noise/ChatTTS>, 2024. [Accessed 26-04-2025].490 Vaibhav Adlakha, Parishad BehnamGhader, Xing Han Lu, Nicholas Meade, and Siva Reddy. Evalu-
491 uating correctness and faithfulness of instruction-following models for question answering, 2024.
492 URL <https://arxiv.org/abs/2307.16877>.493 Afroz Ahamad, Ankit Anand, and Pranesh Bhargava. Accentdb: A database of non-native english
494 accents to assist neural speech recognition. *arXiv preprint arXiv:2005.07973*, 2020.495 Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer,
496 Reuben Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber. Common voice: A
497 massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670*, 2019.498 Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani
499 Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, Jenia Jitsev, Simon Kornblith, Pang Wei
500 Koh, Gabriel Ilharco, Mitchell Wortsman, and Ludwig Schmidt. Openflamingo: An open-
501 source framework for training large autoregressive vision-language models, 2023. URL <https://arxiv.org/abs/2308.01390>.502 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
503 Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan,
504 Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng,
505 Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report, 2025.
506 URL <https://arxiv.org/abs/2502.13923>.507 Emanuele Bastianelli, Andrea Vanzo, Paweł Swietojanski, and Verena Rieser. Slurp: A spoken
508 language understanding resource package. *arXiv preprint arXiv:2011.13205*, 2020.509 Mathieu Bernard and Hadrien Titeux. Phonemizer: Text to phones transcription for multiple lan-
510 guages in python. *Journal of Open Source Software*, 6(68):3958, 2021. doi: 10.21105/joss.03958.
511 URL <https://doi.org/10.21105/joss.03958>.512 Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx,
513 Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportu-
514 nities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.515 brianlan. This script calculates Phoneme Error Rate using package leven (edit distance algorithm).
516 Users can also choose whether to merge phonemes (refer to the paper Speaker-independent
517 phone recognition using hidden Markov models) during calculation. — gist.github.com/brianlan/954082454b25d7eb7abd612ba61c6a66, 2017.
518 [Accessed 15-07-2025].519 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
520 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
521 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.522 Huawei Cao, David G Cooper, Michael K Keutmann, Ruben C Gur, Ani Nenkova, and Ragini
523 Verma. Crema-d: Crowd-sourced emotional multimodal actors dataset. *IEEE transactions on
524 affective computing*, 5(4):377–390, 2014.525 Sahil Chaudhary. Code alpaca: An instruction-following llama model for code generation. <https://github.com/sahil280114/codealpaca>, 2023.526 Guoguo Chen, Shuzhou Chai, Guanbo Wang, Jiayu Du, Wei-Qiang Zhang, Chao Weng, Dan Su,
527 Daniel Povey, Jan Trmal, Junbo Zhang, et al. Gigaspeech: An evolving, multi-domain asr corpus
528 with 10,000 hours of transcribed audio. *arXiv preprint arXiv:2106.06909*, 2021.529 Stanley F Chen, Douglas Beeferman, and Roni Rosenfeld. Evaluation metrics for language models.
530 1998.

540 Ye Chen, Wei Cai, Liangmin Wu, Xiaowei Li, Zhanxuan Xin, and Cong Fu. Tigerbot: An open
 541 multilingual multitask llm. *arXiv preprint arXiv:2312.08688*, 2023.
 542

543 Yiming Chen, Xianghu Yue, Chen Zhang, Xiaoxue Gao, Robby T. Tan, and Haizhou Li. Voicebench:
 544 Benchmarking llm-based voice assistants, 2024. URL <https://arxiv.org/abs/2410.17196>.
 545

546 Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and
 547 Jingren Zhou. Qwen-audio: Advancing universal audio understanding via unified large-scale
 548 audio-language models, 2023. URL <https://arxiv.org/abs/2311.07919>.
 549

550 Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv,
 551 Jinzheng He, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen2-audio technical report, 2024.
 552 URL <https://arxiv.org/abs/2407.10759>.
 553

554 Federico Cocchi, Nicholas Moratelli, Davide Caffagni, Sara Sarto, Lorenzo Baraldi, Marcella Cor-
 555 nia, and Rita Cucchiara. Llava-more: A comparative study of llms and visual backbones for
 556 enhanced visual instruction tuning, 2025. URL <https://arxiv.org/abs/2503.15621>.
 557

558 Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick
 559 Wendell, Matei Zaharia, and Reynold Xin. Free dolly: Introducing the world's first truly open
 560 instruction-tuned llm, 2023.
 561

562 Alexandre Défossez, Laurent Mazaré, Manu Orsini, Amélie Royer, Patrick Pérez, Hervé Jégou,
 563 Edouard Grave, and Neil Zeghidour. Moshi: a speech-text foundation model for real-time dia-
 564 logue. *arXiv preprint arXiv:2410.00037*, 2024.
 565

566 Zhihao Du, Jiaming Wang, Qian Chen, Yunfei Chu, Zhifu Gao, Zerui Li, Kai Hu, Xiaohuan Zhou,
 567 Jin Xu, Ziyang Ma, et al. Lauragpt: Listen, attend, understand, and regenerate audio with gpt.
 568 *arXiv preprint arXiv:2310.04673*, 2023.
 569

570 Benjamin Elizalde, Soham Deshmukh, and Huaming Wang. Natural language supervision for
 571 general-purpose audio representations, 2024. URL <https://arxiv.org/abs/2309.05767>.
 572

573 Jesse Engel, Cinjon Resnick, Adam Roberts, Sander Dieleman, Douglas Eck, Karen Simonyan, and
 574 Mohammad Norouzi. Neural audio synthesis of musical notes with wavenet autoencoders, 2017.
 575

576 Qingkai Fang, Shoutao Guo, Yan Zhou, Zhengrui Ma, Shaolei Zhang, and Yang Feng. Llama-omni:
 577 Seamless speech interaction with large language models, 2025. URL <https://arxiv.org/abs/2409.06666>.
 578

579 Chaoyou Fu, Yi-Fan Zhang, Shukang Yin, Bo Li, Xinyu Fang, Sirui Zhao, Haodong Duan, Xing
 580 Sun, Ziwei Liu, Liang Wang, et al. Mme-survey: A comprehensive survey on evaluation of
 581 multimodal llms. *arXiv preprint arXiv:2411.15296*, 2024.
 582

583 Sreyan Ghosh, Zhifeng Kong, Sonal Kumar, S Sakshi, Jaehyeon Kim, Wei Ping, Rafael Valle, Di-
 584 nesh Manocha, and Bryan Catanzaro. Audio flamingo 2: An audio-language model with long-
 585 audio understanding and expert reasoning abilities. *arXiv preprint arXiv:2503.03983*, 2025.
 586

587 Yuan Gong, Jin Yu, and James Glass. Vocalsound: A dataset for improving human vocal sounds
 588 recognition. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and
 589 Signal Processing (ICASSP)*, pp. 151–155. IEEE, 2022.
 590

591 Yuan Gong, Sameer Khurana, Leonid Karlinsky, and James Glass. Whisper-at: Noise-robust
 592 automatic speech recognizers are also strong general audio event taggers. *arXiv preprint
 593 arXiv:2307.03183*, 2023a.
 594

595 Yuan Gong, Alexander H Liu, Hongyin Luo, Leonid Karlinsky, and James Glass. Joint audio and
 596 speech understanding. In *2023 IEEE Automatic Speech Recognition and Understanding Work-
 597 shop (ASRU)*, pp. 1–8. IEEE, 2023b.
 598

594 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 595 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan,
 596 Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Ko-
 597 reniev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava
 598 Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux,
 599 Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret,
 600 Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius,
 601 Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary,
 602 Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab
 603 AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco
 604 Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind That-
 605 tai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Kore-
 606 vaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra,
 607 Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Ma-
 608 hadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu,
 609 Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jong-
 610 soot Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasudevan Alwala,
 611 Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid
 612 El-Arini, Krishika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren
 613 Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin,
 614 Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi,
 615 Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew
 616 Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Ku-
 617 mar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoy-
 618 chev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan
 619 Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajwal Bhargava, Pratik Dubal, Praveen Krishnan,
 620 Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ra-
 621 mon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Ro-
 622 hit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan
 623 Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell,
 624 Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparth, Sheng
 625 Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer
 626 Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman,
 627 Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mi-
 628 haylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor
 629 Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei
 630 Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang
 631 Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Gold-
 632 schlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning
 633 Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh,
 634 Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria,
 635 Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein,
 636 Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, An-
 637 drew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, An-
 638 nie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel,
 639 Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leon-
 640 hardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu
 641 Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Mon-
 642 talvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao
 643 Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia
 644 Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide
 645 Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le,
 646 Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily
 647 Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smo-
 thers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni,
 648 Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia
 649 Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan,
 650 Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harri-
 651 son Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj,

648 Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James
 649 Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jen-
 650 nifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang,
 651 Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Jun-
 652 jie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy
 653 Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang,
 654 Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell,
 655 Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa,
 656 Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias
 657 Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L.
 658 Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike
 659 Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari,
 660 Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan
 661 Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong,
 662 Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent,
 663 Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar,
 664 Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Ro-
 665 driguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy,
 666 Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin
 667 Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon,
 668 Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ra-
 669 maswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha,
 670 Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal,
 671 Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satter-
 672 field, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj
 673 Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo
 674 Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook
 675 Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Ku-
 676 mar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov,
 677 Wei Li, Wencheng Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiao-
 678 jian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia,
 679 Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao,
 680 Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhao-
 681 duo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The llama 3 herd of models, 2024. URL
 682 <https://arxiv.org/abs/2407.21783>.

683 Or Honovich, Thomas Scialom, Omer Levy, and Timo Schick. Unnatural instructions: Tuning
 684 language models with (almost) no human labor. *arXiv preprint arXiv:2212.09689*, 2022.

685 Cheng-Yu Hsieh, Jieyu Zhang, Zixian Ma, Aniruddha Kembhavi, and Ranjay Krishna. Sugar-
 686 crepe: Fixing hackable benchmarks for vision-language compositionality, 2023. URL <https://arxiv.org/abs/2306.14610>.

687 Kaiyi Huang, Chengqi Duan, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. T2i-comptbench++:
 688 An enhanced and comprehensive benchmark for compositional text-to-image generation, 2025.
 689 URL <https://arxiv.org/abs/2307.06350>.

690 huggingface. Evaluation metrics for ASR - Hugging Face Audio Course — huggingface.co.
 691 <https://huggingface.co/learn/audio-course/chapter5/evaluation>,
 692 2023. [Accessed 15-07-2025].

693 Keith Ito and Linda Johnson. The lj speech dataset. <https://keithito.com/LJ-Speech-Dataset/>, 2017.

694 Feng Jiang, Zhiyu Lin, Fan Bu, Yuhao Du, Benyou Wang, and Haizhou Li. S2s-arena, evaluating
 695 speech2speech protocols on instruction following with paralinguistic information, 2025. URL
 696 <https://arxiv.org/abs/2503.05085>.

697 Jordan. Whitefox44/alpacagpt3.5customized · datasets at hugging face, Mar 2023. URL <https://huggingface.co/datasets/whitefox44/AlpacaGPT3.5Customized>.

702 KimiTeam, Ding Ding, Zeqian Ju, Yichong Leng, Songxiang Liu, Tong Liu, Zeyu Shang, Kai Shen,
 703 Wei Song, Xu Tan, Heyi Tang, Zhengtao Wang, Chu Wei, Yifei Xin, Xinran Xu, Jianwei Yu,
 704 Yutao Zhang, Xinyu Zhou, Y. Charles, Jun Chen, Yanru Chen, Yulun Du, Weiran He, Zhenxing
 705 Hu, Guokun Lai, Qingcheng Li, Yangyang Liu, Weidong Sun, Jianzhou Wang, Yuzhi Wang,
 706 Yuefeng Wu, Yuxin Wu, Dongchao Yang, Hao Yang, Ying Yang, Zhilin Yang, Aoxiong Yin,
 707 Ruibin Yuan, Yutong Zhang, and Zaida Zhou. Kimi-audio technical report, 2025. URL <https://arxiv.org/abs/2504.18425>.

709 Zhifeng Kong, Arushi Goel, Rohan Badlani, Wei Ping, Rafael Valle, and Bryan Catanzaro. Audio
 710 flamingo: A novel audio language model with few-shot learning and dialogue abilities. *arXiv*
 711 preprint [arXiv:2402.01831](https://arxiv.org/abs/2402.01831), 2024a.

713 Zhifeng Kong, Arushi Goel, Rohan Badlani, Wei Ping, Rafael Valle, and Bryan Catanzaro. Audio
 714 flamingo: A novel audio language model with few-shot learning and dialogue abilities, 2024b.
 715 URL <https://arxiv.org/abs/2402.01831>.

717 Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Nau-
 718 mann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large language-and-vision assis-
 719 tant for biomedicine in one day, 2023a. URL <https://arxiv.org/abs/2306.00890>.

720 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image
 721 pre-training with frozen image encoders and large language models, 2023b. URL <https://arxiv.org/abs/2301.12597>.

724 Lei Li, Yuwei Yin, Shicheng Li, Liang Chen, Peiyi Wang, Shuhuai Ren, Mukai Li, Yazheng Yang,
 725 Jingjing Xu, Xu Sun, Lingpeng Kong, and Qi Liu. M³it: A large-scale dataset towards multi-
 726 modal multilingual instruction tuning, 2023c. URL <https://arxiv.org/abs/2306.04387>.

728 Bin Lin, Yang Ye, Bin Zhu, Jiaxi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning
 729 united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*,
 730 2023.

732 Junyan Lin, Haoran Chen, Dawei Zhu, and Xiaoyu Shen. To preserve or to compress: An in-
 733 depth study of connector selection in multimodal large language models. In Yaser Al-Onaizan,
 734 Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical*
 735 *Methods in Natural Language Processing*, pp. 5666–5680, Miami, Florida, USA, November
 736 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.325. URL
 737 <https://aclanthology.org/2024.emnlp-main.325>.

739 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances*
 740 *in neural information processing systems*, 36:34892–34916, 2023.

741 Yinghao Ma, Anders Øland, Anton Ragni, Bleiz MacSen Del Sette, Charalampos Saitis, Chris Don-
 742 ahue, Chenghua Lin, Christos Plachouras, Emmanouil Benetos, Elona Shatri, et al. Foundation
 743 models for music: A survey. *arXiv preprint arXiv:2408.14340*, 2024.

745 Zixian Ma, Jerry Hong, Mustafa Omer Gul, Mona Gandhi, Irena Gao, and Ranjay Krishna.
 746 Crepe: Can vision-language foundation models reason compositionally?, 2023. URL <https://arxiv.org/abs/2212.07796>.

748 Potsawee Manakul, Guangzhi Sun, Warit Sirichotedumrong, Kasima Tharnpipitchai, and Kunat
 749 Pipatanakul. Enhancing low-resource language and instruction following capabilities of audio
 750 language models, 2024. URL <https://arxiv.org/abs/2409.10999>.

752 Xinhao Mei, Chutong Meng, Haohe Liu, Qiuqiang Kong, Tom Ko, Chengqi Zhao, Mark D. Plumb-
 753 ley, Yuexian Zou, and Wenwu Wang. Wavcaps: A chatgpt-assisted weakly-labelled audio caption-
 754 ing dataset for audio-language multimodal research. *IEEE/ACM Transactions on Audio, Speech,*
 755 *and Language Processing*, 32:3339–3354, 2024. ISSN 2329-9304. doi: 10.1109/taslp.2024.
 3419446. URL <http://dx.doi.org/10.1109/TASLP.2024.3419446>.

756 Andrew Cameron Morris, Viktoria Maier, and Phil D. Green. From WER and RIL to MER and WIL:
 757 improved evaluation measures for connected speech recognition. In *8th International Conference*
 758 *on Spoken Language Processing, INTERSPEECH-ICSLP 2004, Jeju Island, Korea, October 4-8,*
 759 *2004*, pp. 2765–2768. ISCA, 2004. doi: 10.21437/INTERSPEECH.2004-668. URL <https://doi.org/10.21437/Interspeech.2004-668>.

760

761 Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and
 762 Ahmed Awadallah. Orca: Progressive learning from complex explanation traces of gpt-4. *arXiv*
 763 *preprint arXiv:2306.02707*, 2023.

764

765 Arsha Nagrani, Joon Son Chung, and Andrew Zisserman. Voxceleb: a large-scale speaker identifi-
 766 cation dataset. *arXiv preprint arXiv:1706.08612*, 2017.

767 Gonzalo Navarro. A guided tour to approximate string matching. *ACM computing surveys (CSUR)*,
 768 33(1):31–88, 2001.

769

770 OpenAI. Hello gpt-4o, 2024. URL <https://openai.com/index/hello-gpt-4o>.

771

772 Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an asr corpus
 773 based on public domain audio books. In *2015 IEEE international conference on acoustics, speech*
 774 *and signal processing (ICASSP)*, pp. 5206–5210. IEEE, 2015.

775

776 Prabhat Pandey, Rupak Vignesh Swaminathan, KV Girish, Arunasish Sen, Jian Xie, Grant P Strimel,
 777 and Andreas Schwarz. Sift-50m: A large-scale multilingual dataset for speech instruction fine-
 778 tuning. *arXiv preprint arXiv:2504.09081*, 2025.

779

780 Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and
 781 Quoc V Le. Specaugment: A simple data augmentation method for automatic speech recognition.
 782 *arXiv preprint arXiv:1904.08779*, 2019.

783

784 Jing Peng, Yucheng Wang, Bohan Li, Yiwei Guo, Hankun Wang, Yangui Fang, Yu Xi, Haoyu Li,
 785 Xu Li, Ke Zhang, Shuai Wang, and Kai Yu. A survey on speech large language models for
 786 understanding, 2025. URL <https://arxiv.org/abs/2410.18908>.

787

788 Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada
 789 Mihalcea. Meld: A multimodal multi-party dataset for emotion recognition in conversations.
 790 *arXiv preprint arXiv:1810.02508*, 2018.

791

792 Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. Mls: A
 793 large-scale multilingual dataset for speech research. *ArXiv*, abs/2012.03411, 2020.

794

795 Zengyi Qin, Wenliang Zhao, Xumin Yu, and Xin Sun. Openvoice: Versatile instant voice cloning.
 796 *arXiv preprint arXiv:2312.01479*, 2023.

797

798 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 799 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 800 models from natural language supervision. In *International conference on machine learning*, pp.
 8748–8763. PMLR, 2021.

801

802 Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever.
 803 Robust speech recognition via large-scale weak supervision, 2022. URL <https://arxiv.org/abs/2212.04356>.

804

805 Ricardo Reimao and Vassilios Tzerpos. For: A dataset for synthetic speech detection. In *2019*
 806 *International Conference on Speech Technology and Human-Computer Dialogue (SpeD)*, pp. 1–
 807 10. IEEE, 2019.

808

809 Tiger Research. Tigerresearch/tigerbot-alpaca-en-50k · datasets at hugging face, May
 2023. URL <https://huggingface.co/datasets/TigerResearch/tigerbot-alpaca-en-50k>.

810

Andrew Rouditchenko, Saurabhchand Bhati, Edson Araujo, Samuel Thomas, Hilde Kuehne, Roge-
 811 rio Feris, and James Glass. Omni-rl: Do you really need audio to fine-tune your audio llm? *arXiv*
 812 *preprint arXiv:2505.09439*, 2025.

810 RuterNorway. RuterNorway/OpenOrcaNo-15k · Datasets at Hugging Face — hugging-
 811 face.co. <https://huggingface.co/datasets/RuterNorway/OpenOrcaNo-15k>,
 812 2023. [Accessed 17-04-2025].

813

814 Yi Su, Jisheng Bai, Qisheng Xu, Kele Xu, and Yong Dou. Audio-language models for audio-centric
 815 tasks: A survey. *arXiv preprint arXiv:2501.15177*, 2025.

816 Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and
 817 Chao Zhang. Salmonn: Towards generic hearing abilities for large language models, 2024. URL
 818 <https://arxiv.org/abs/2310.13289>.

819 Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy
 820 Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model.
 821 https://github.com/tatsu-lab/stanford_alpaca, 2023.

822

823 Christophe Veaux, Junichi Yamagishi, Kirsten MacDonald, et al. Supersedes-cstr vctk corpus: En-
 824 glish multi-speaker corpus for cstr voice cloning toolkit.(2016). URL <http://datashare.is.ed.ac.uk/handle/10283/2651>, 2016.

825

826 Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu,
 827 Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng
 828 Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-vl: Enhancing vision-language model's
 829 perception of the world at any resolution, 2024. URL <https://arxiv.org/abs/2409.12191>.

830

831 Pete Warden. Speech commands: A dataset for limited-vocabulary speech recognition. arxiv 2018.
 832 *arXiv preprint arXiv:1804.03209*, 1804.

833

834 Yusong Wu, Ke Chen, Tianyu Zhang, Yuchen Hui, Marianna Nezhurina, Taylor Berg-Kirkpatrick,
 835 and Shlomo Dubnov. Large-scale contrastive language-audio pretraining with feature fusion and
 836 keyword-to-caption augmentation, 2024. URL <https://arxiv.org/abs/2211.06687>.

837

838 Zhifei Xie and Changqiao Wu. Mini-omni2: Towards open-source gpt-4o with vision, speech and
 839 duplex capabilities, 2024a. URL <https://arxiv.org/abs/2410.11190>.

840

841 Zhifei Xie and Changqiao Wu. Mini-omni: Language models can hear, talk while thinking in
 842 streaming, 2024b. URL <https://arxiv.org/abs/2408.16725>.

843

844 Zhifei Xie, Mingbao Lin, Zihang Liu, Pengcheng Wu, Shuicheng Yan, and Chunyan Miao. Audio-
 845 reasoner: Improving reasoning capability in large audio language models, 2025. URL <https://arxiv.org/abs/2503.02318>.

846

847 Xinchen-ai. GitHub - xinchen-ai/Westlake-Omni — github.com. <https://github.com/xinchen-ai/Westlake-Omni>, 2024. [Accessed 21-06-2025].

848

849 Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang
 850 Fan, Kai Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. Qwen2.5-omni technical
 851 report, 2025. URL <https://arxiv.org/abs/2503.20215>.

852

853 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
 854 Gao, Chengan Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,
 855 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
 856 Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,
 857 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui
 858 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang
 859 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger
 860 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan
 861 Qiu. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.

862

863 Yi Yang, Wen-tau Yih, and Christopher Meek. WikiQA: A challenge dataset for open-domain
 864 question answering. In Lluís Márquez, Chris Callison-Burch, and Jian Su (eds.), *Proceedings
 865 of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 2013–2018,
 866 Lisbon, Portugal, September 2015. Association for Computational Linguistics. doi: 10.18653/v1/
 867 D15-1237. URL <https://aclanthology.org/D15-1237/>.

864 Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A sur-
 865 vey on multimodal large language models. *National Science Review*, 11(12), November 2024.
 866 ISSN 2053-714X. doi: 10.1093/nsr/nwae403. URL <http://dx.doi.org/10.1093/nsr/nwae403>.

867

868 Chien Yu Huang, Ke-Han Lu, Shih-Heng Wang, Chi-Yuan Hsiao, Chun-Yi Kuan, Haibin Wu,
 869 Siddhant Arora, Kai-Wei Chang, Jiatong Shi, Yifan Peng, Roshan Sharma, Shinji Watanabe,
 870 Bhiksha Ramakrishnan, Shady Shehata, and Hung Yi Lee. Dynamic-superb: Towards a dy-
 871 namic, collaborative, and comprehensive instruction-tuning benchmark for speech, 2024. URL
 872 <https://arxiv.org/abs/2309.09510>.

873

874 Aohan Zeng, Zhengxiao Du, Mingdao Liu, Kedong Wang, Shengmin Jiang, Lei Zhao, Yuxiao Dong,
 875 and Jie Tang. Glm-4-voice: Towards intelligent and human-like end-to-end spoken chatbot,
 876 2024a. URL <https://arxiv.org/abs/2412.02612>.

877

878 Yunan Zeng, Yan Huang, Jinjin Zhang, Zequn Jie, Zhenhua Chai, and Liang Wang. Investigating
 879 compositional challenges in vision-language models for visual grounding. In *Proceedings of the*
 880 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14141–14151, 2024b.

881

882 Junbo Zhang, Zhiwen Zhang, Yongqing Wang, Zhiyong Yan, Qiong Song, Yukai Huang, Ke Li,
 883 Daniel Povey, and Yujun Wang. speechocean762: An open-source non-native English speech
 884 corpus for pronunciation assessment. In *Proc. Interspeech 2021*, 2021.

885

886 Xin Zhang, Dong Zhang, Shimin Li, Yaqian Zhou, and Xipeng Qiu. Speechtokenizer: Unified
 887 speech tokenizer for speech large language models, 2024. URL <https://arxiv.org/abs/2308.16692>.

888

889 Wenliang Zhao, Xumin Yu, and Zengyi Qin. Melotts: High-quality multi-lingual multi-accent text-
 890 to-speech, 2023. URL <https://github.com/myshell-ai/MeloTTS>.

891

892 Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing
 893 vision-language understanding with advanced large language models, 2023. URL <https://arxiv.org/abs/2304.10592>.

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918 **A LIMITATION**
919920 While our work establishes a unified, open-source foundation for compositional speech-language
921 instruction tuning, several limitations remain. First, LLaSO Corpus is currently limited to English,
922 which constrains its direct applicability to non-English and low-resource languages. Extending the
923 dataset and benchmark to multilingual scenarios is an important direction for achieving broader im-
924 pact and inclusivity. Second, despite surpassing prior datasets in task diversity and modality cover-
925 age, the granularity and availability of source materials inherently influence our corpus composition,
926 particularly in the underrepresented paralinguistic categories and rare interaction scenarios. Third,
927 our reference model, LLaSO-Base, intentionally prioritizes reproducibility and extensibility over
928 achieving SOTA performance. Consequently, its architecture and model size (3.8 billion parameters)
929 are modest compared to larger models, and our evaluations have primarily included similarly sized
930 or smaller baselines. Assessing and benchmarking significantly larger LSLMs would provide further
931 insights into scaling behaviors and capabilities. Fourth, certain challenging multimodal interactions
932 such as open-ended dialogues involving overlapping speech, or zero-shot generalization to entirely
933 new domains are only partially addressed within our current benchmark and model architecture. We
934 encourage the research community to build upon our foundation to tackle these limitations, further
935 refining instruction-tuned speech-language models for diverse languages, scenarios, and real-world
936 applications.937 **B ETHICAL STATEMENT**
938939 **B.1 DATA PRIVACY AND CONSENT**
940941 All training and evaluation data are sourced solely from publicly available datasets, with no use of
942 private or personally identifiable information. Synthetic, TTS, and sound effect samples contain no
943 human-identifiable content. No re-identification or de-anonymization was performed at any stage.
944 All data handling complies with ethical standards and legal requirements.945 **B.2 LICENSING AND RESPONSIBLE USE**
946947 All data, code, and model weights are released under permissive open-source licenses, with explicit
948 terms governing use and redistribution. The resources are intended for academic, non-commercial
949 research, and must be used in accordance with ethical standards and applicable copyright laws.
950951 **B.3 DIVERSITY AND REPRESENTATIVENESS**
952953 We strive for diversity in gender, age, accent, language, and emotion across both collected and
954 synthetic data, employing balanced sampling where possible. Nonetheless, certain groups and lan-
955 guages remain underrepresented, and we acknowledge the risk of bias. We encourage the community
956 to further augment and improve coverage. All datasets and models are released without representing
957 the views or interests of any particular group or institution.958 **B.4 FAIRNESS AND MISUSE PREVENTION**
959960 Our models and datasets may exhibit uneven performance across different tasks, languages, or de-
961 mographic groups, and should not be considered universally fair or unbiased. We explicitly prohibit
962 the use of our work for surveillance, discrimination, harassment, or any activities that may harm
963 individuals or communities. We encourage responsible research and deployment that respects the
964 rights and dignity of all users.965
966
967
968
969
970
971

972 **C METRIC DETAILS**

973

974 Given the diversity of tasks and modalities in LLaSO-Eval, we define 7 metrics to ensure comprehensive evaluation. During evaluation, by default metric scores are computed directly from the 975 complete model outputs. However, certain tasks require an intermediate step of extracting structured 976 answers from the model outputs prior to evaluation; such cases are explicitly noted in their 977 respective metric descriptions below. Task-specific metric assignments are detailed in Table 12. 978

979

980 **WER and CER.** Word Error Rate (WER) and Character Error Rate (CER) Morris et al. (2004); 981 huggingface (2023); Chen et al. (1998) quantify transcription accuracy derived from Levenshtein 982 distance Navarro (2001) between the model prediction and the ground-truth transcript. WER 983 operates at the word level, while CER operates at the character level. Both metrics are employed for 984 ASR and SCR tasks. Lower values indicate better accuracy. Typically, WER and CER scores range 985 from 0 (perfect match) to 1, although values exceeding 1 can occur due to excessive insertions or 986 substitutions.

987

988 **PER.** Phoneme Error Rate (PER) is analogous to WER and CER but specifically measures the 989 Levenshtein distance between the predicted and ground-truth phoneme sequences with brianlan 990 (2017), providing a phoneme-level accuracy assessment. Similar to WER and CER, lower PER 991 values indicate superior performance, typically ranging from 0 upwards, with 0 representing a perfect 992 phoneme prediction. We apply this metric exclusively to the PR task.

993

994 **Accuracy.** Accuracy is defined as the proportion of exact matches between the model’s prediction 995 and the ground-truth label. This metric is applied to all closed-ended tasks, as specified in Table 11, 996 which also indicates which tasks are open- versus closed-ended. For closed-ended tasks, the model 997 must select a single answer from a predefined label set, and a response is marked correct only if it 998 precisely matches the reference label; predictions containing multiple candidate labels or irrelevant 999 content are treated as incorrect. Accuracy ranges from 0 to 1, with higher values indicating better 1000 performance. Additionally, the metric is also computed for the open-ended PSSL task, where 1001 models rate sentence-level pronunciation across three dimensions, accuracy, prosodic, and fluency. 1002 Given that model outputs are typically free-form, we use regular expressions to extract numeric 1003 scores from responses, accommodating variations such as “accuracy is 8”, “fluency: 7”, or “9 for 1004 prosodic”. Responses providing valid numeric scores for all three dimensions are retained; others 1005 are excluded. For each sample, we compute the average of the exact-match accuracies across these 1006 three dimensions, then report the overall accuracy averaged across all evaluated samples. Further, we 1007 complement this rule-based measure with an additional GPT-4o evaluation to ensure comprehensive 1008 assessment.

1009

1010 **MAE.** Mean Absolute Error (MAE) is adopted for tasks requiring numerical predictions, such 1011 as AR and PR. In these tasks, the model is explicitly instructed to generate a single numeric value. 1012 However, many LSLMs produce free-form textual outputs rather than direct numeric predictions Ad- 1013 lakha et al. (2024), necessitating an answer extraction procedure prior to metric calculation. For the 1014 AR task, the predicted numeric value represents age and thus must be extracted reliably from the 1015 model output. Employing a regular expression, our script initially attempts a direct integer 1016 conversion; if unsuccessful, it searches for numeric patterns, and if a numeric range like “40-45” is 1017 detected, it computes the rounded average of the two endpoints. For outputs containing descriptive 1018 keywords “adult” without numeric information, we substitute a canonical age value 22. The PR 1019 task requires evaluation of the model’s ability to predict MIDI note values ranging from 0 to 127. 1020 Specifically, our extraction function sequentially attempts integer conversion, rounded float 1021 conversion, and finally, averaging numeric ranges. Strings containing pitch-related keywords (e.g., “midi”, 1022 “pitch”, “hz”) but lacking numeric values are marked as invalid predictions. Only predictions within 1023 the MIDI range of 0 to 127 are considered valid for metric computation. In all cases, if a valid 1024 numeric value cannot be extracted from either the model’s prediction, that instance is omitted from the 1025 calculation. The final MAE is computed as the mean absolute difference between extracted predictions 1026 and ground-truth numeric values across all valid instances. Lower MAE values indicate better 1027 numerical prediction performance.

1026 **GPT-4o Score.** For AQA and other open-ended generative tasks, where model responses are un-
1027 constrained and may vary widely in form and content, thus we employ GPT-4o (OpenAI, gpt-4o-
1028 mini, Version 2024-07-18) as an automatic evaluator. Following a standardized evaluation template,
1029 GPT-4o assigns an integer score from 1 to 5, reflecting both the relevance and accuracy of the
1030 model’s response relative to the reference answer. Further details of the evaluation prompt are pro-
1031 vided in Appendix O, and task-specific metric assignments are summarized in Table 12.

1032 **Abstention Rate.** Some LSLMs may abstain from answering tasks involving unfamiliar modal-
1033 ity formats or instructions, fail to follow instructions, or explicitly state their inability to process
1034 audio. To quantify such behavior, we report the abstention rate for closed-ended tasks, defined as
1035 the proportion of responses in which the model either refuses to answer, returns irrelevant content,
1036 or fails to select a valid label from the predefined set. Higher scores indicate better performance;
1037 lower abstention rates indicate stronger instruction following. An abstention is counted whenever
1038 the model’s output does not comply with the task requirement to select a label. Abstention rate
1039 is not reported for open-ended tasks, as their free-form nature precludes a rule-based criterion for
1040 abstention.

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080 **D DETAILED ANALYSIS**
 1081

1082 We find that LSLMs perform poorly on unseen tasks and unfamiliar modality formats, with espe-
 1083 cially weak instruction-following in unseen tasks.

1084 *Broader task coverage leads to better performance and lower abstention rates.* We present the
 1085 overall performance, closed-ended task performance, and closed-ended task abstention rates for 11
 1086 models, alongside the number of tasks each model was trained on, in Figure 6. For models with
 1087 private, incomplete, or ambiguously reported training data, we use the number of evaluation tasks as
 1088 a proxy for task coverage. The results show that models exposed to a wider range of tasks achieve
 1089 higher performance in both overall and closed-ended tasks, and fewer abstentions. This finding
 1090 suggests that, in creating LSLMs for speech understanding, one should diversify the tasks as much
 1091 as possible, to improve the model performance and reduce the abstention rate.

1092 *LSLMs perform worse on unseen modality configuration.* Most existing LSLMs only support one or
 1093 two modality configurations. We evaluate their generalization across three different input formats.
 1094 Specifically, we select representative models and calculate their average performance on AQA task
 1095 across all major input modality configurations with results summarized in Table 2. We observe that
 1096 model performance in the audio instruction with text input setting drops consistently compared to the
 1097 familiar text instruction with audio input configuration. At the same time we find that Qwen2-Audio
 1098 is an outlier, showing that it and its variant Audio-Reasoner obtain similar results in both formats.
 1099 Notably, from a human perspective, audio instruction with text input should be no more difficult
 1100 and is arguably even simpler, since only the (typically brief) instruction needs to be heard, while
 1101 the main input remains directly readable as text, as illustrated in Figure 4 (b, c). Nonetheless, our
 1102 findings demonstrate that LSLMs still struggle with modality configurations outside their explicit
 1103 training coverage.

1104 *Pure audio modality configuration may still challenging.* We illustrate the performance of 11 models
 1105 three major modality formats in Figure 8 (Bottom). In most cases, models demonstrate substantially
 1106 *lower* performance on pure audio formats than on the more common text instruction with audio input
 1107 setting, even when they are explicitly trained to handle pure audio via speech-to-speech or spoken-
 1108 query-based QA (SQQA) tasks. Notably, for some of these models, the performance drop from text
 1109 with audio to pure audio is even greater than the decline observed on modality formats they have
 1110 never seen during training, such as audio instruction with text input. Interestingly, only a handful
 1111 of models such as Qwen2.5-Omni, GLM-4-Voice, and the Mini-Omni family achieve comparable
 1112 performance across pure audio and text + audio modalities. Nonetheless, for most current LSLMs,
 1113 the pure audio configuration remains a notably challenging setting.

1114 *Interleaving and parallel decoding strategies help bridge performance gaps across modality config-
 1115 urations.* As shown in Table 2 and Figure 8 (Bottom), nearly all models achieve their best results on
 1116 the common text instruction with audio input setting. To assess model robustness to modality shifts,
 1117 we compute the sum of absolute performance differences between this common configuration and
 1118 the other two input formats. We present the results in ascending order of stability in Figure 8 (Top).
 1119 Among the eleven models evaluated, the top eight with the exception of Qwen2-Audio and its vari-
 1120 ant employ interleaving or parallel decoding strategies (see Appendix I for benchmarking model
 1121 details), and exhibit notably reduced modality gaps. These outliers may reflect factors outside of
 1122 modality combination design. Overall, our results provide empirical evidence that interleaving and
 1123 parallel decoding can bridge the performance gap between text and audio modalities.

1124 *LSLMs may prefer content related tasks.* To further analyze model performance on paralinguistic
 1125 tasks, we present a dumbbell plot in Figure 7, contrasting content-centric and speaker-centric results
 1126 for each model. We observe that most models achieve higher performance and lower abstention
 1127 rates on content-centric tasks than on speaker-centric ones. This disparity likely arises because
 1128 content-centric tasks are more tightly linked to the semantic content, which LLM-based decoders are
 1129 naturally equipped to process. In contrast, speaker-centric tasks demand more nuanced inference of
 1130 latent speaker attributes, posing a greater challenge for current LSLMs and highlighting an important
 1131 area for future improvement.

1131
 1132
 1133

1134 **E ABLATION**
1135

1136 We conduct ablation experiments on different training strategies in LLaSO Corpus, as shown in Ta-
1137 ble 4 and 5. (i) *Alignment Robustness*. We evaluate ASR performance both immediately after the
1138 alignment stage and following the subsequent instruction-tuning phase. After alignment, the model
1139 achieves strong results (WER = 0.05, CER = 0.01). After multi-task instruction tuning, ASR per-
1140 formance declines slightly (WER = 0.08, CER = 0.03), yet remains competitive, due to we include
1141 ASR samples within the LLaSO-Instruct dataset for mitigating catastrophic forgetting. (ii) *Encoder*
1142 *Fine-tuning*. We ablate the effect of unfreezing the audio encoder during the instruction-following
1143 stage, comparing the results of freezing versus jointly training the encoder, projector, and LLM on
1144 LLaSO-Instruct. When the encoder is unfrozen, ASR performance drops more substantially (WER
1145 = 0.14, CER = 0.07), relative to the frozen configuration. In contrast, AQA (semantic) tasks see
1146 modest improvements (see Table 5), while paralinguistic tasks exhibit a slight decline (see Table 4).
1147 This suggests that while joint fine-tuning may benefit certain high-level reasoning tasks, it may
1148 compromise low-level speech recognition and nuanced paralinguistic abilities.

Paralinguistic Task Category																				
Speaker-Centric					Content-Centric															
Tasks	LLaSO-Base		LLaSO-Base (Unfrozen)		Δ (U-F)	Metrics	Tasks	LLaSO-Base		LLaSO-Base (Unfrozen)		Δ (U-F)	Metrics							
	0.96	0.88	-0.08 \downarrow			PR	SCR	0.03	0.03	0.00 $=$	PER \downarrow									
SGC	0.99	0.99	0.00 $=$					0.04	0.05	0.01 \downarrow	WER \downarrow									
	0.76	0.61	-0.15 \downarrow	ACC \uparrow				0.02	0.04	0.02 \downarrow	CER \downarrow									
	0.91	0.97	0.06 \uparrow					4.86	4.80	-0.06 \downarrow	GPT-4o \uparrow									
	0.91	0.86	-0.05 \downarrow	Avg.ACC \uparrow				3.93	3.90	-0.03 \downarrow										
	0.52	0.40	-0.12 \downarrow	EE				3.57	3.44	-0.13 \downarrow										
	0.83	0.86	0.03 \uparrow	ACC \uparrow				0.78	0.82	0.04 \uparrow										
AC	0.73	0.78	0.05 \uparrow	IP				0.50	0.55	0.05 \uparrow	ACC \uparrow									
	0.69	0.68	-0.01 \downarrow	Avg.ACC \uparrow				0.60	0.77	0.17 \uparrow										
	0.70	0.68	-0.02 \downarrow	ISC				8.02	10.55	2.53 \downarrow	MAE \downarrow									
	0.50	0.38	-0.12 \downarrow	PP				0.18	0.20	0.02 \uparrow	ACC \uparrow									
	0.60	0.53	-0.07 \downarrow	VC				Linguistic Task Category												
	10.32	8.78	-1.54 \uparrow	Avg.ACC \uparrow				< Text, Audio >												
EIE	0.48	0.45	-0.03 \downarrow	MAE \downarrow				Tasks	LLaSO-Base		LLaSO-Base (Unfrozen)		Δ (U-F)	Metrics						
	0.48	0.37	-0.11 \downarrow	ACC \uparrow					0.08	0.14	0.06 \downarrow	WER \downarrow								
	0.48	0.41	-0.07 \downarrow	Avg.ACC \uparrow					0.03	0.07	0.04 \downarrow	CER \downarrow								
	0.17	0.16	-0.01 \downarrow	ACC \uparrow				ER	LLaSO-Base		LLaSO-Base (Aligned)									
	0.26	0.30	0.04 \uparrow	ASR					0.05	0.08	0.03 \downarrow	WER \downarrow								
	0.22	0.23	0.01 \uparrow	Avg.ACC \uparrow					0.01	0.03	0.02 \downarrow	CER \downarrow								
SSD	0.99	0.99	0.00 $=$	ACC \uparrow					LLaSO-Base		LLaSO-Base		Δ (F-A)	Metrics						
	0.32	0.16	-0.16 \downarrow	ASR					Tasks		Tasks									
	0.99	0.99	0.00 $=$	ACC \uparrow					LLaSO-Base		LLaSO-Base									
	0.29	2.68	-0.22 \downarrow	GPT-4o \uparrow					Aligned		Base									
	0.39	0.24	-0.15 \downarrow	GPT-4o \uparrow					ACC \uparrow		Base									
	0.39	0.24	-0.15 \downarrow	GPT-4o \uparrow					ACC \uparrow		Base									

1178 Table 4: Ablation results for LLaSO-Base on paralinguistic (speaker-centric and content-centric) and linguistic
1179 tasks, all evaluated under the text instruction with audio input modality. Frozen (F) and unfrozen (U) refer to
1180 whether the audio encoder is fixed or updated during instruction tuning, respectively. Δ (U-F) reports the
1181 performance change between unfrozen and frozen encoder variants during finetuning, while Δ (F-A) compares
1182 results after finetuning with frozen encoder (F) and after the alignment stage (A) for ASR. \uparrow / \downarrow denote gains or
1183 drops; metrics follow previous tables.

Tasks	Semantic Task Category			Metrics
	LLaSO-Base	LLaSO-Base (Unfrozen)	Δ (U-F)	
< Pure Audio >				
AQA	2.06	2.27	0.21 \uparrow	
	1.80	2.27	0.47 \uparrow	
	2.39	2.23	-0.16 \downarrow	GPT-4o\uparrow
	1.46	1.98	0.52 \uparrow	
	1.93	2.19	0.26 \uparrow	Avg.GPT-4o\uparrow
< Text, Audio >				
AQA	2.57	2.54	-0.03 \downarrow	
	2.48	2.42	-0.06 \downarrow	
	1.71	1.96	0.25 \uparrow	GPT-4o\uparrow
	2.74	2.80	0.06 \uparrow	
	3.05	3.09	0.04 \uparrow	
	2.90	2.87	-0.03 \downarrow	
	2.58	2.61	0.03 \uparrow	Avg.GPT-4o\uparrow
< Audio, Text >				
AQA	2.72	3.22	0.50 \uparrow	
	2.62	2.84	0.22 \uparrow	
	2.28	2.47	0.19 \uparrow	GPT-4o\uparrow
	2.23	2.43	0.20 \uparrow	
	3.74	3.68	-0.06 \downarrow	
	2.60	3.26	0.66 \uparrow	
	2.70	2.98	0.28 \uparrow	Avg.GPT-4o\uparrow

Table 5: Ablation results for LLaSO-Base, comparing frozen (F) and unfrozen (U) audio encoder variants during instruction tuning. The table reports Δ (U-F) performance changes on AQA tasks across three major modality configurations; \uparrow / \downarrow denote relative gains or drops. Metrics follow earlier tables.

F CASE STUDY

We provide qualitative evidence of the compositional flexibility and unified modeling offered by our framework. Figure 4 demonstrates that LLaSO-Base seamlessly accommodates all three instruction-input modality pairings. In particular, the pure audio example highlights the system’s ability to disentangle instructions from content solely within the audio stream. For the other formats, LLaSO-Base reliably grounds reasoning and response generation in the correct modality, adapting to instructions and content presented in any combination. We present some task prototypes unified by our system in Figure 3 and present more cases across different instruction-input structures and tasks in Appendix J. These examples show that, unlike models limited rigid instruction-input structures, LLaSO-Base generalizes across both categorical and compositional tasks without requiring task-specific modules or post-processing.

To better understand the challenges across modality configurations and tasks, we further present cases under the same benchmark. As shown in Figure 9 (I), we sample three representative cases for Salmonn across the primary modality configurations. In case (I)(b) with textual instruction and audio input, the model’s familiar modality configuration, Salmonn correctly follows the user instruction. The textual instruction asks the model to determine whether the first sentence in the speech content can be used to define the term in the second sentence, and to answer “yes” or “no”. The audio provides a definition of “database” in the first sentence and describes “MySQL” in the second. The model correctly interprets both the instruction and the speech, and returns the correct answer “yes”. In contrast, when the modality configuration shifts outside the model’s primary training distribution, distinct failures emerge. Under the pure audio setting (I)(a) the model receives both the instruction and content as audio, yet responds with a counter-question: “What is the answer to the

1242	Instruction Please listen to the audio and answer the question. <audio>	Instruction You are given two sentences...to determine if the first sentence can be used to define the term in the second sentence...“Yes” or “No”. <audio>	Instruction Question: <audio> Just woke up had a really scary dream. It’s all rainy out. Describe the sentiment. OPT: 1). negative 2). Positive. I think the answer is
1243	<audio content> I made my way toward the people lined up at the bedding Windows. If I saw Tony...the answer is	<audio content> A database is... MySQL is a...database management system.	<audio content> You...will be given a task. You must...generate a...detailed and long...answer.
1244	Answer : What is the answer to the question?	Answer : A database is... MySQL is a...database management system.	Answer : 1. negative 2. positive

(a) Pure Audio

(b) < Textual Instruction, Audio Input >

(c) < Audio Instruction, Textual Input >

(I) Three evaluated cases from Salmonn across the primary modality configurations. The correct one corresponds to the model’s supported training format; the two errors are from modality formats with limited or no training exposure.

1252	Instruction Question: Make sure to interpret the speaker’s request correctly and reply accordingly. <audio>	Instruction Question: What is the aim of the speaker in this speech without any additional context? <audio>	Instruction Question: <audio> like dogs vs. I do not like cats. <audio content> Determine whether the following two sentences are equivalent or not.
1253	<audio content> Classify the following object as either living or non-living. A rock .	<audio content> I want to know more about this event.	Answer : I’m happy to help. However, I need more information about the speech, such...

(a) Pure Audio

(b) < Textual Instruction, Audio Input >

(c) < Audio Instruction, Textual Input >

(II) Three Llama-Omni responses sampled from the three major modality configurations. The correct example comes from its supported format; the two errors are from its unsupported modality configurations.

Figure 9: Case studies from LLaSO-Eval. (Left) Model behaviors under different modality configurations (Salmonn and Llama-Omni), highlighting the importance of supporting multiple modality formats for reliable instruction following. (Right) Model behaviors across tasks with different coverage (Typhoon-Audio and GLM-4-Voice), underscoring the necessity of broad task coverage for generalization.

question?” This indicates that while the model has some prior exposure to SQA tasks, it fails to correctly interpret or respond to this particular modality configuration where both instruction and content are delivered as audio. In Figure 9 (I)(c), where an audio instruction is followed by a textual input, the spoken instruction assigns a generic task, prompting the model to generate an answer, while the accompanying text presents a tweet and asks for its sentiment, providing two answer options. The relevant information for reasoning is contained within the text input, and the instruction directs the model to perform a classification. However, the model does not follow the instruction; instead, it merely repeats the sentiment options, “1. negative 2. positive”, without making a decision. We observe similar results in Llama-Omni across modality formats, illustrated in Figure 9 (II). This model is primarily trained on pure audio modality, and this is directly reflected in the sampled cases. In (II)(a) where both the instruction and input are delivered as audio, the model answers successfully classifying the object as non-living, demonstrating effective handling of its core modality. Nonetheless, when presented with configurations outside this primary distribution, the model fails to execute the intended tasks. In the text plus audio modality format (II)(b), it is unable to infer the speaker’s aim from the speech and instead requests further contextual details. Under (II)(c) the modality configuration of audio instruction paired with textual input, the model follows speech instruction but overlooks the explicit negation in the text input and incorrectly judges the two sentences as equivalent. Taken together, this observation highlights the importance of comprehensive modality coverage for multimodal instruction following.

To provide an intuitive comparison of model performance on covered versus uncovered tasks, we sample more cases for representative baselines in Figure 9. Figure 9 (III) illustrates this contrast for Typhoon-Audio model. In (III)(a) we present a sample from the entity extraction task, which

1245	Instruction Question: Find the entities in this audio and return the sentence with slot labels without any additional commentary. <audio>	Instruction Question: Listen to this audio and transcribe the speech, then score the sentence for accuracy, prosody, and fluency, followed by a total score, each on a scale of 1 to 10, as integers. <audio>	Instruction Question: Convert the speech in this audio file to its written form and return only the transcript, without any additional comments.
1246	<audio content> Put meeting with Pawel for tomorrow ten am.	<audio content> Mark is going to see elephant.	<audio content> Tell me the first day that you noticed something wrong.

(a) Entity Extraction

(b) Pronunciation Scoring Sentence Level

(c) Automatic Speech Recognition

(III) Three samples from Typhoon-Audio: one from a well-represented task with correct prediction, and two errors from tasks absent in its training.

1247	Instruction Question: Listen to the provided speech sample and classify it as either real human speech or a synthetic voice. Output only real or fake. <audio>	Instruction Question: Detect the emotional tone in this speaker’s voice for sentiment analysis. Return only one of these categories: neutrality, joy, surprise, sadness, anger, fear, or disgust. <audio>	Instruction Question: Can you answer the question by analyzing this audio? <audio>
1248	<audio content> A Woman Speaking (Real Audio)	<audio content> H~(Neutrality)	<audio content> The correct answer is no...

(a) Synthetic Speech Detection

(b) Emotion Recognition

(c) Audio Question Answering

(IV) Selected answers from GLM-4-Voice illustrate success on a task with ample training exposure and failure on two tasks that fall outside its main training coverage.

1296 is not included in the model’s training. Here, the query requests identification of entities from the
 1297 speech, but the model misinterprets the task as speaker gender classification, responding with “The
 1298 speaker is female.” We present a pronunciation scoring sentence-level (PSSL) task sample in (III)(b),
 1299 where the model is instructed to evaluate the speech for accuracy, prosody, and fluency. However,
 1300 it only provides a plain transcription and omits the required scoring. In contrast, when evaluated
 1301 on a task present in its training, the model demonstrates accurate performance. For the ASR task
 1302 presented at (III)(c) the model successfully transcribes the speech to text as instructed without addi-
 1303 tional information. Similar results are evident with GLM-4-Voice in Figure 9 (IV). In (IV)(a) tasked
 1304 with synthetic speech detection, the model avoids providing a categorical decision and instead pro-
 1305 duces an off-topic statement. When we prompt the model for emotion recognition, it generates a
 1306 generic conversational reply, neglecting to engage with the specified sentiment classification task as
 1307 in (IV)(b). Nevertheless, we can observe that the model successfully completes an audio question
 1308 answering task in (IV)(c), owing to the presence of this task in its training. These findings underscore
 1309 the essential role of comprehensive task coverage in building models across diverse speech-language
 1310 tasks.

1311 G DISCUSSION

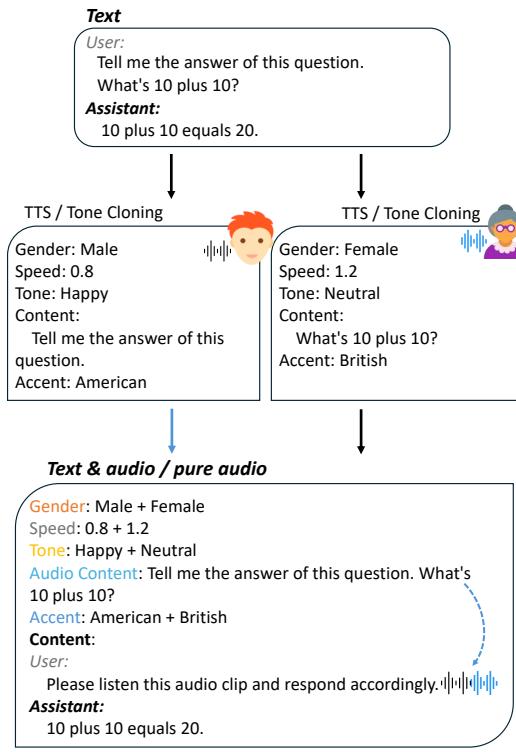
1312 In addition to our quantitative analysis, a closer manual inspection of model outputs reveals several
 1313 distinct patterns and recurring issues that warrant discussion. For example, Qwen2-Audio occasion-
 1314 ally misinterprets the SV task as SGC. On other tasks, Although this model achieves relatively strong
 1315 quantitative scores on the open-ended PP task, qualitative check reveals that a large proportion of
 1316 its outputs are empty strings (41 cases) or single periods (30 cases) among 112 samples in the text
 1317 instruction plus audio input configuration. For the VC task, the abstention rate is especially high. In
 1318 our manual review, 82 out of 100 samples in the same modality setting resulted in a single period
 1319 (“.”) as the response. As to Typhoon-Audio, it occasionally responds in Thai to English prompts,
 1320 which is likely attributable to the inclusion of Thai data during fine-tuning. Salmonn, when presented
 1321 with pure audio or audio plus text modality input, often refuses to answer, asks clarifying questions,
 1322 or claims the audio contains no content; this may stem from its limited exposure to pure audio in-
 1323 struction data (approximately 20K SQQA task samples from WikiQA Yang et al. (2015)) and to the
 1324 modality configuration of audio instruction with text input, restricting its ability to generalize. Addi-
 1325 tionally, in the ASR task, Salmonn’s outputs for most samples consist entirely of uppercase English
 1326 letters, which explains its poor quantitative performance on this task. GLM-4-Voice frequently gen-
 1327 erates responses in Chinese and at times misinterprets the audio input as part of the conversational
 1328 context, rather than as content to be analyzed. A similar pattern is observed in the Mini-Omni family,
 1329 which occasionally interprets the audio input, such as a speaker’s utterance to be classified, as either
 1330 an instruction or as primary content. For example, in ASR task with the text plus audio input config-
 1331 uration, approximately 43% of Mini-Omni and 53% of Mini-Omni2 responses begin with phrases
 1332 like “It sounds like...”, reflecting a tendency to treat the input as dialogue. Meanwhile Llama-Omni
 1333 exhibits a high rate of refusals in tasks beyond AQA, suggesting limited coverage of tasks. Manual
 1334 inspection of its ASR task outputs further reveals that over 10% of samples are explicit abstains,
 1335 likely because the model was not trained on ASR data. These phenomena further illustrate the ne-
 1336 cessity for broad and balanced coverage across both tasks and modality configurations in model
 1337 development. Audio-Reasoner presents a different set of challenges, often exhibiting hallucinated
 1338 completions such as appending “the answer is A” or fabricating multiple-choice options like “E,”
 1339 “F,” or “S.” Since it is trained on Qwen2-Audio-Instruct with chain-of-thought data, this tendency
 1340 may stem from exposure to reasoning-style supervision. Kimi-Audio, when performing the PR task,
 1341 sometimes treats the input as an ASR query or outputs a sequence of isolated phonemes, which
 1342 leads to lower evaluation scores. Qwen2.5-Omni occasionally shows similar confusion between PR
 1343 and ASR, and routinely appends conversational phrases like “feel free to ask me more.” While such
 1344 additions might be intended to improve user interaction, they can undermine instruction following
 1345 if overfitting, as some tasks in our benchmark explicitly requires models to return only the answer.
 1346 Taken together, these observations offer practical insights into persistent issues in instruction fol-
 1347 lowing, task differentiation, and output format consistency. We hope these manual inspections are
 1348 helpful for the community and can inform future model development and evaluation.

1349

1350 H VOCAL MIXING STRATEGY

1352 To generate multimodal variants from text-based semantic QA data, we employ advanced audio
 1353 synthesis tools such as MeloTTS (Zhao et al., 2023), ChatTTS (2noise, 2024), and OpenVoice (Qin
 1354 et al., 2023). These systems support controllable attributes (e.g., gender, speed, tone, accent), en-
 1355 abling the creation of audio segments with varied vocal styles. The synthesized audio is then com-
 1356 bined with either text or additional audio segments to form the three modality configurations de-
 1357 scribed in the main paper.

1358 For example, in a pure-audio sample shown as Figure 10, one speaker may voice the instruction in
 1359 a cheerful American accent, while another renders the task content in a neutral British tone, simu-
 1360 lating realistic multi-speaker interaction. Similarly, in <Text Instruction, Audio Input> or <Audio
 1361 Instruction, Text Input> formats, synthesized speech segments are paired with textual components
 1362 to enrich acoustic diversity. Using distinct voices for instructions and content helps clearly separate
 1363 roles within the dialogue and better approximates real-world speech-language interactions.



1390 **Figure 10: Illustration of vocal style mixing.** Utterances are synthesized with varied speaker traits and
 1391 applied across all three modality configurations, expanding acoustic diversity and simulating realistic multi-
 1392 speaker scenarios.

1394 I BENCHMARKING CANDIDATES

1396 **Qwen2-Audio.** A LSLM from Qwen Team designed for both audio analysis and voice chat. It
 1397 integrates a Whisper-large-v3 audio encoder with a Qwen-7B language model, enabling processing
 1398 of audio and text inputs for instruction following and conversational tasks. The model supports both
 1399 audio + text and pure audio modality configurations, automatically distinguishing between analysis
 1400 and dialogue modes without explicit prompts. It achieves state-of-the-art results such as AIR-Bench
 1401 and CoVoST2, with open-source demos, weights, and inference code.

1402 **Typhoon-Audio.** A LSLM from SCB 10X and the University of Cambridge supporting both En-
 1403 glish and Thai. It integrates Whisper-large-v3 (fine-tuned for Thai) and BEATs audio encoders, a Q-

1404
 1405 Former adapter, and a Typhoon-1.5-8B-Instruct LLM. The model supports both text-audio and pure
 1406 audio (namely speech instruction following in this paper) configurations. Demo, model weights, and
 1407 inference code are open-source.

1408
 1409 **Salmonn.** Salmonn is an unified LSLM with a dual-encoder architecture, Whisper and BEATs,
 1410 linked via a window-level Query Transformer to a Vicuna-based LLM. The model is trained in three
 1411 task levels using a two-stage alignment and instruction-tuning scheme, and further enhanced through
 1412 activation tuning to unlock emergent capabilities. Salmonn supports audio-plus-text and pure audio
 1413 (through SQQA) modality configurations and diverse task types. Demos, model checkpoints, training/inference
 1414 code, and training data are all publicly available.

1415
 1416 **Glm-4-Voice.** An end-to-end spoken chatbot supporting both Chinese and English from Zhipu.AI
 1417 and Tsinghua University. The model combines Gilm-4-9B-Base with a supervised speech tokenizer
 1418 and a flow-matching speech decoder, pre-trained on 1T tokens of speech-text and speech-only data.
 1419 Fine-tuned with a streaming-thoughts template, it alternates between text and speech tokens for
 1420 seamless, low-latency conversational output. GLM-4-Voice accepts speech or text inputs and pro-
 1421 duces simultaneous speech and text responses. Model weights, demo, and inference code are open-
 1422 source.

1423
 1424 **Mini-Omni.** Developed by Inspirai and Tsinghua University, Mini-Omni is a streaming speech-
 1425 to-speech conversational LLM integrating a Whisper-small encoder, modality adapters, a Qwen2-
 1426 0.5B transformer language model, and a TTS adapter. The system employs parallel decoding for
 1427 efficient, real-time, end-to-end speech input and streaming audio output. Model weights, inference
 1428 code, demo, and the VoiceAssistant-400K dataset are open-source.

1429
 1430 **Mini-Omni2.** An omni-interactive multimodal model, developed by Inspirai and Tsinghua Uni-
 1431 versity as an upgraded version of Mini-Omni, combining CLIP (ViT-B/32) for vision, Whisper-small
 1432 for audio, and Qwen2-0.5B for language. It enables real-time, end-to-end voice conversations with
 1433 users, supporting image, audio, and text inputs and text, audio outputs. The model, inference and
 1434 demo code are open-source.

1435
 1436 **Llama-Omni.** Developed by ICTNLPLab at CAS, this model integrates a frozen Whisper-large-
 1437 v3 encoder, a trainable speech adaptor, a Llama-3.1-8B-Instruct language model, and a streaming
 1438 speech decoder. Its key innovation is simultaneous generation of both text and speech responses
 1439 from spoken instructions, enabling low-latency, end-to-end speech-to-text and speech-to-speech
 1440 interaction. The model, along with its training data, weights, demo, and inference code, is open-
 1441 source.

1442
 1443 **Audio-Reasoner.** A reasoning-oriented LSLM developed by fine-tuning Qwen2-Audio with struc-
 1444 tured chain-of-thought (CoT) supervision on its 1.2M-sample CoTA dataset. Emphasizing complex
 1445 audio reasoning, it demonstrates the benefits of CoT-style instruction tuning, achieving competi-
 1446 tive results including MMAU-mini and AIR-Bench-Chat. It is open-source along with its model
 1447 checkpoint, demo, inference code, and dataset.

1448
 1449 **Kimi-Audio.** An audio foundation model developed by the Kimi Team featuring a hybrid archi-
 1450 tecture with an audio tokenizer, audio encoder, core audio LLM, parallel heads for both text and
 1451 audio generation, and an audio detokenizer, using continuous acoustic vectors and discrete semantic
 1452 tokens. Pre-trained on 13 million hours of diverse open and in-house audio, the model is fine-tuned
 1453 for multimodal comprehension and generation tasks involving speech, music, and sound effects,
 1454 including audio understanding, speech conversation, and audio-to-text chat. It demonstrates strong
 1455 performance on benchmarks such as VoiceBench, VocalSound, and MELD. The project is open-
 1456 source, providing demo data, fine-tuning and inference code, released model weights, and an audio
 1457 evaluation toolkit.

1458
 1459 **Qwen2.5-Omni.** A unified end-to-end real-time multimodal model developed by the Qwen team,
 1460 supporting text, audio, image, and video inputs, with both streaming text and speech outputs. Built
 1461 on the Thinker-Talker architecture, it enables flexible cross-modal interactions and streaming, facil-
 1462 itated by TMRoPE and block-wise encoding for efficient temporal alignment. The model achieves

strong performance on diverse multimodal benchmarks like VoiceBench and MMAU and open-source with released weights, APIs, and inference code.

Model Name in this Paper	Official Model Name	URL	#Params	Supported Modalities	Interleaving or Parallel Decoding
Qwen2-Audio	Qwen/Qwen2-Audio-7B-Instruct	[Model Card]	7B	$\langle T, A \rangle, \langle PA \rangle$	-
Typhoon-Audio	scb10x/llama-3-typhoon-v1.5-8b-audio-preview	[Model Card]	8B	$\langle T, A \rangle, \langle PA \rangle$	-
Salmonn	tsinghua-ee/SALMONN-7B	[Model Card]	7B	$\langle T, A \rangle, \langle PA \rangle$	-
Glm-4-Voice	THUDM/glm-4-voice-9b	[Model Card]	9B	$\langle T, A \rangle, \langle PA \rangle$	Interleaving
Mini-Omni	gpt-omni/mini-omni	[Model Card]	0.5B	$\langle T, A \rangle, \langle PA \rangle$	Parallel
Mini-Omni2	gpt-omni/mini-omni2	[Model Card]	0.5B	$\langle T, A \rangle, \langle PA \rangle$	Parallel
Llama-Omni	ICTNLP/Llama-3.1-8B-Omni	[Model Card]	8B	$\langle PA \rangle$	Parallel
Audio-Reasoner	zhifeixie/Audio-Reasoner	[Model Card]	7B	$\langle T, A \rangle, \langle PA \rangle$	-
Kimi-Audio	moonshotai/Kimi-Audio-7B-Instruct	[Model Card]	7B	$\langle T, A \rangle, \langle PA \rangle$	Interleaving and Parallel
Qwen2.5-Omni	Qwen/Qwen2.5-Omni-7B	[Model Card]	7B	$\langle T, A \rangle, \langle PA \rangle$	Interleaving and Parallel

Table 6: Details of Benchmarking Candidates in this work. Each row lists a tested model’s official HuggingFace repository (“Model Card” link), parameter size, supported modalities (“ $\langle T, A \rangle$ ” = Textual Instruction + Audio Input; “ $\langle PA \rangle$ ” = Pure Audio), and whether Interleaving or Parallel multi-modal decoding is applied.

1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511

J CASES

1512

1513

1514

1515

1516

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565

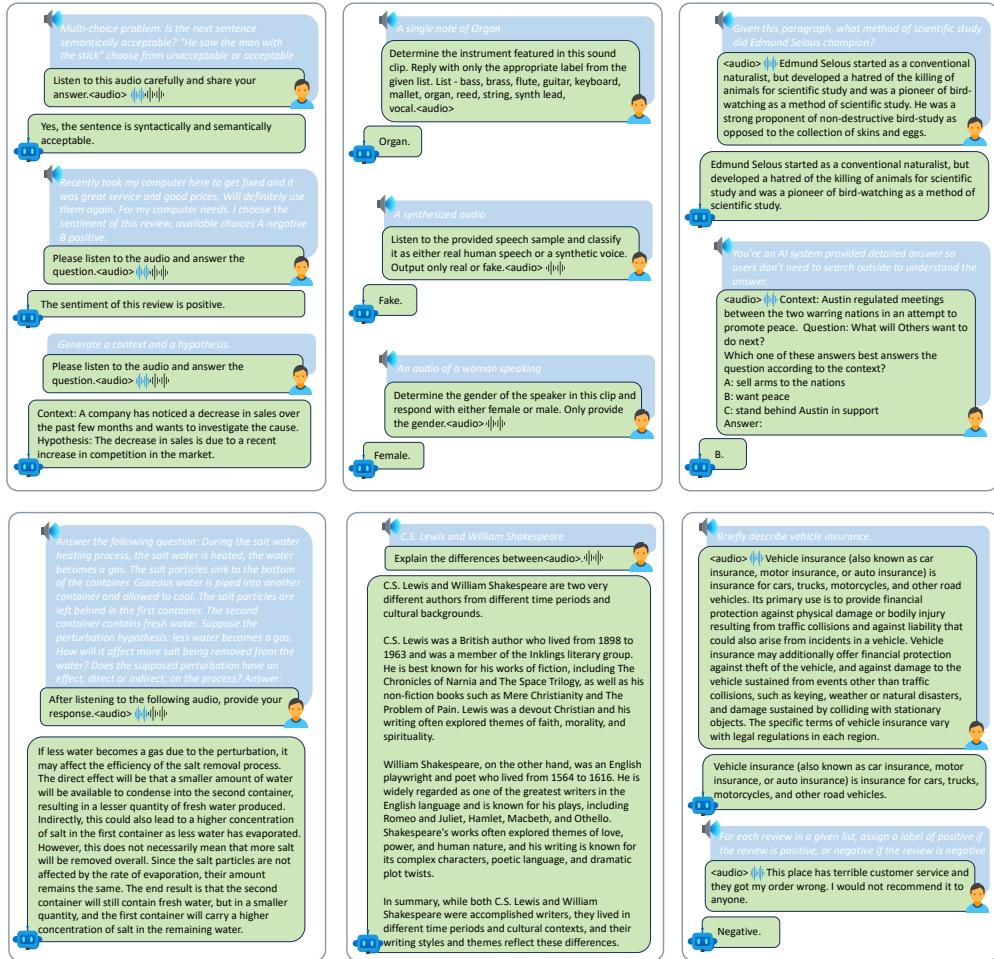


Figure 11: Representative case examples from LLaSO-Base demonstrating the three modality configurations in LLaSO-Eval: pure audio (left), text instruction with audio input (middle), and audio instruction with text input (right). Each column presents distinct tasks under its respective format.

1566 **K TRAINING DETAILS**
15671568 **K.1 SYSTEM PROMPT**
15691570 **System Prompt for LLaSO-Base**
15711572 A chat between a curious user and an artificial intelligence assistant. The assistant is able to un-
1573 derstand the audio content that the user provides, and assist the user with a variety of tasks using
1574 natural language. The audio content will be provided with the following format: <Audio>audio
1575 content</Audio>.
15761577 Box 12: Our system prompt for training and evaluation. <Audio>and </Audio>are added into the tokenizer
1578 vocabulary as special tokens.
15791580
1581 **K.2 PROMPT TEMPLATE**
15821583 **Chat Template for LLaSO-Base**
15841585 <|begin_of_text|><|start_header_id|>system<|end_header_id|>
1586
1587 $X_{system-prompt}^t$ <|eot_id|><|start_header_id|>user<|end_header_id|>
1588
1589 X_{query}^* <|eot_id|><|start_header_id|>assistant<|end_header_id|>
1590
1591 X_{answer}^t <|eot_id|>
15921593 Box 13: Illustration of the chat template used to construct every training example. We follow the official
1594 Llama-3.2 chat template for token ordering and special tokens, while inserting a custom system prompt (full
1595 text in Appendix K.1). The user request is encoded as X_{query}^* (see Eq. 1 for the three modality variants), and
1596 the model must generate the assistant reply X_{answer}^t followed by the end-of-turn token <|eot_id|>. During
1597 training, the loss is applied *only* to the assistant’s tokens (the last line in this box), teaching the network both
1598 the content of the response and where to terminate.
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

1620
1621 **K.3 TRAINING CONFIGURATION**

1622 1623 Parameter	1624 Stage 1: Modality Alignment	1625 Stage 2: Instruction Tuning
1624 Device	4 × NVIDIA A800	4 × NVIDIA A800
1625 Model Backbone	Llama-3.2-3B-Instruct	Llama-3.2-3B-Instruct
1626 Audio Encoder	Whisper-large-v3	Whisper-large-v3
1627 Audio Projector	MLP (2-layer, GELU)	MLP (2-layer, GELU)
1628 Pretrain Audio Aligner	—	Aligner Checkpoint (from Stage 1)
1629 Tune Audio Encoder	False	True/False (optional, see ablation)
1630 Tune Audio Projector	True	True
1631 Tune LLM	False	True
1632 Epochs	1	1
1633 Global Batch Size	256	128
1634 Learning Rate	1×10^{-3}	3×10^{-5}
1635 Weight Decay	0.0	0.0
1636 Warmup Ratio	0.01	0.01
1637 LR Scheduler	Cosine	Cosine
1638 Max Grad Norm	1.0	1.0
BF16	True	True
Model Max Length	2048	2048

1639
1640 Table 7: Training hyperparameters for LLaSO-Base. Stage 1 performs cross-modality alignment,
1641 while Stage 2 instruction-tunes the unified model.1642
1643 **K.4 TRAINING LOSS**1652
1653 Figure 14: Training loss visualization with Raw Loss and Smoothed Loss. From left to right: (1)
1654 alignment stage; (2) instruction tuning stage with frozen encoder; (3) instruction tuning stage with
unfrozen encoder.

1674
1675
1676
1677 **L FOUR STYLES INSTRUCTIONAL PROMPTS**
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698

Prompt Style	Closed-ended Instruction Examples
Standardized	Classify the instrument in this audio clip. Choose only from: bass, brass, flute, guitar, keyboard, mallet, organ, reed, string, synth lead, vocal. Output only the label.<audio>
Contextualized	For a music classification project, identify the primary instrument in this audio. Return only one of the following: bass, brass, flute, guitar, keyboard, mallet, organ, reed, string, synth lead, vocal.<audio>
Stylistic Variation	What is the primary instrument in this audio clip? Respond only with one of: bass, brass, flute, guitar, keyboard, mallet, organ, reed, string, synth lead, or vocal.<audio>
Fine-grained Task	Focus only on the instrumental characteristics and determine the correct classification. Output just one label from bass, brass, flute, guitar, keyboard, mallet, organ, reed, string, synth lead, vocal.<audio>
Prompt Style	Open-ended Instruction Examples
Standardized	Convert the speech in this audio file into an IPA phonemic sequence. Return phonemes only.<audio>
Contextualized	A linguist is analyzing speech samples. Your task is to transcribe the provided audio into an IPA phonemic sequence. Return phonemes only.<audio>
Stylistic Variation	Help build a pronunciation guide by converting this audio into IPA phonemes. Return only the phonemes.<audio>
Fine-grained Task	Phonetic decoding task: transcribe the provided speech into IPA phonemes and return them without any additional output.<audio>

1699
1700 Table 8: Representative prompts illustrating the four instruction styles used in our corpus. The
1701 closed-ended examples (top) are drawn from the Instrument Classification (IC) task, while the open-
1702 ended examples (bottom) are from the Phoneme Recognition (PR) task. Each style, Standardized
1703 (direct instructions), Contextualized (scenario-driven), Stylistic Variation (diverse linguistic formu-
1704 lations), and Fine-grained Task (specific sub-aspect focus), is designed to promote compositional
1705 generalization across tasks and formats.
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

1728
1729
1730
1731 **M MULTI-GRANULARITY SETTING DETAILS**
1732
1733
1734

Coarse-grained (10-year spans)	Medium-grained (5-year spans)	Fine-grained (exact age)
<i>Categories: eighties, fifties, forties, nineties, seventies, sixties, teens, thirties, twenties</i>	<i>Categories: 15-19, 20-24, 25-29, 30+</i>	<i>Range: integer between 18 and 80</i>
Analyze the speaker’s voice and determine their age category. Respond only with one of the following: eighties, fifties, fourties, nineties, seventies, sixties, teens, thirties, or twenties.<audio>	Based on the audio, identify the speaker’s age group. Select one of the following age groups only return: 15-19, 20-24, 25-29, 30+.<audio>	Estimate the age of the speaker from the human vocal sounds in this audio clip. Respond with the age only, between 18 and 80.<audio>
A speech-based recommendation system needs to identify user age. Analyze the voice and classify it into the correct age group from eighties, fifties, fourties, nineties, seventies, sixties, teens, thirties, or twenties.<audio>	Can you guess the age group of the speaker in this clip? Please select from the following age groups only return: 15-19, 20-24, 25-29, 30+.<audio>	Using this audio, which contains human vocalizations, estimate the speaker’s age. Respond with the age as an integer between 18 and 80, no extra information.<audio>
If the speaker’s age appears ambiguous, classify them into the closest matching age group. Select only one label - eighties, fifties, fourties, nineties, seventies, sixties, teens, thirties, or twenties.<audio>	Based on the audio, what age group is being used? Pick only return from: 15-19, 20-24, 25-29, 30+.<audio>	Listen to this sound sample of human vocalizations and predict the speaker’s age as a number between 18 and 80. Provide the age only.<audio>
Analyze the energy levels, speech rate, and vocal strain in the voice to determine the most accurate age category. Provide only the label from eighties, fifties, fourties, nineties, seventies, sixties, teens, thirties, or twenties.	From the following audio, can you determine the speaker’s age group? Options only return: 15-19, 20-24, 25-29, 30+.<audio>	Determine the speaker’s age based on this recording of human vocalizations. Respond with the age between 18 and 80, without any other explanation.

1761
1762 **Table 9:** Some of tasks in our data have granularity. We use Age Classification (AC) task as an ex-
1763 amples at three different granularity levels. Coarse-grained prompts elicit classification into decade-
1764 based age groups, medium-grained prompts target 5-year age spans, and fine-grained prompts re-
1765 quest exact age prediction within a specified range.
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

1782 **N PROMPTS FOR PURE AUDIO MODALITY FORMAT**
1783

Prompt Style	Closed-ended Instruction Examples
Standardized	Analyze the provided audio and complete the task mentioned in it.<audio> Based on the instruction in the audio, provide your response.<audio> Listen to the audio and respond accordingly.<audio> Carefully listen to the audio clip and perform the requested action.<audio> Follow the instruction given in the audio and provide an accurate response.<audio>
Contextualized	A voice assistant is asking you to do something. Carefully listen and respond.<audio> For a comprehension test, listen to the audio and answer the question presented in it.<audio> In this conversation, the speaker is giving you a directive. Listen and respond appropriately.<audio> In this experiment, you need to complete the task given in the audio. Provide your response accordingly.<audio> This is an interactive task. Listen to the speaker and follow their instruction.<audio>
Stylistic Variation	Can you understand and complete the request made in this audio?<audio> If the audio contains a question, answer it accurately. If it contains a command, follow it.<audio> What action is required in the audio? Complete it and provide your response.<audio> Make sure to interpret the speaker's request correctly and reply accordingly.<audio> The speaker in this audio needs a response. Listen and provide a relevant reply.<audio>
Fine-grained Task	After hearing the audio, provide your answer to the given task.<audio> Listen carefully and act according to the instruction in the recording.<audio> Pay attention to the details in the audio and respond exactly as instructed.<audio> Understand the content of the audio and give an appropriate response.<audio> Your task is to carefully analyze the instruction in the audio and execute it properly.<audio>

1814
1815 Table 10: Examples of text prompts used in the pure-audio modality format, where both the instruc-
1816 tion and content are embedded within a single audio stream. The textual cues only instruct the model
1817 to listen and respond, without specifying task details. All four prompt styles are included as Table 8
1818 - Standardized, Contextualized, Stylistic Variation, and Fine-grained Task.
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835

1836 **O EVALATION TEMPLATE**
18371838
1839 **Instructions:** You are evaluating the performance of an AI assistant in an audio question answering
1840 task.
1841Given a **Reference Answer** and a **Predicted Answer**, assign a score from **1 to 5** based on **Rele-**
1842 **vance and Accuracy.**1843 **Output Format (exactly, no other text):**1844 • Score: <integer 1--5>
1845 • Explanation: <concise justification focusing on both
1846 relevance and accuracy>
18471848 **Reference Answer:**
1849 {reference}
18501851 **Predicted Answer:**
1852 {predicted}
18531854 **Please produce the evaluation.**
1855
18561857 Figure 15: Evaluation template used for GPT-4o-based scoring of LSLMs’ responses. The model
1858 assigns an integer score (1–5) according to relevance and accuracy, accompanied by a concise ex-
1859 planation. All results were scored with OpenAI GPT-4o (gpt-4o-mini, Version 2024-07-18).
18601861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889

1890 P TASK CATEGORY DEFINITIONS

1891
 1892 To facilitate comprehensive and interpretable evaluation, both our training and evaluation datasets
 1893 are systematically organized into three principal categories: linguistic, semantic, and paralinguistic.
 1894 This categorization is designed to capture the spectrum of speech-language understanding, from core
 1895 speech processing and factual reasoning to the nuanced interpretation of speaker traits and acoustic
 1896 context. Next I will describe the definitions of each categories.
 1897

1898 P.1 LINGUISTIC CATEGORY

1899
 1900 Linguistic tasks are aimed at assessing models’ basic speech processing ability, primarily through
 1901 ASR. This foundational category evaluates how accurately a model can transcribe spoken language
 1902 into text, serving as the backbone for subsequent semantic or paralinguistic inference.
 1903

1904 P.2 SEMANTIC CATEGORY

1905 The semantic category tests a model’s ability to extract explicit meaning and perform higher-level
 1906 reasoning over audio input. In our benchmark, this is represented by the AQA task, which requires
 1907 models to interpret audio content, combine it with contextual cues, and deliver factual or reasoning-
 1908 based responses. Although limited to AQA, this category is critical for evaluating the transition from
 1909 basic perception to comprehension and inference.
 1910

1911 P.3 PARALINGUISTIC CATEGORY

1912 Paralinguistic tasks are structured to probe models’ sensitivity to information that lies beyond the
 1913 literal linguistic content. We further distinguish between speaker-centric and content-centric paralinguistic
 1914 tasks. Speaker-centric tasks focus on characteristics inherent to the speaker such as gender,
 1915 age, accent, emotion, and identity capturing traits that are independent of the message being de-
 1916 livered. In contrast, content-centric tasks emphasize cues embedded in the audio signal that reflect
 1917 content or context, such as phoneme recognition, intent prediction, or entity extraction, irrespective
 1918 of speaker identity.
 1919

1920
 1921
 1922
 1923
 1924
 1925
 1926
 1927
 1928
 1929
 1930
 1931
 1932
 1933
 1934
 1935
 1936
 1937
 1938
 1939
 1940
 1941
 1942
 1943

1944 Q DETAILS FOR LLASO-ALIGN AND LLASO-INSTRUCT

1945
 1946 We have a standardization step in the data construction process as presented in Figure 2. Corrupted
 1947 or unreadable files are removed, and valid audios are resampled to 16 kHz and stored as WAV/FLAC.
 1948 Transcripts are filtered to retain only English content with standard characters, then normalized to
 1949 follow conventional grammar and formatting (e.g., proper capitalization, spacing). Each cleaned
 1950 sample is paired with a randomly selected instruction template, and the final dataset is packaged in
 1951 a unified JSON format.

1952 Applying the constructing procedure as Figure 2 yields the finalized training corpora. We provide
 1953 Table 11 summarizing the resulting task-level composition of LLASO-Align and LLASO-Instruct,
 1954 covering the three categories, linguistic, semantic, and paralinguistic, their 20 sub-tasks, represen-
 1955 tative data sources, supported input formats, and sample statistics. For the held-out evaluation split,
 1956 see the stratified breakdown in Table 12.

1957 Tasks	1958 Descriptions	1959 Data Sources	1960 Modality Formats	1961 Sample Num.	1962 Hours	1963 Instr. Settings
Linguistic Task Category						
1964 ASR	1965 Automatic Speech Recognition	GigaSpeech LibriSpeech LJ Speech VCTK MLS	<Textual Instruction, Audio Input> & <Audio Instruction, Audio Input>	12M (LLASO-Align)&- 1M&0.2M	47K&- 4K&1K	Open-ended
Semantic Task Category						
1966 AQA	1967 Audio Question Answering	OpenOrca1M-GPT4 OpenOrca3.5M-GPT3.5 StanfordAlpaca CodeAlpaca AlpacaDan Dolly OpenOrcaNo Tigerbot_Alpaca Tigerbot_Multichat Unnatural	<Audio Instruction, Audio Input> & <Textual Instruction, Audio Input> & <Audio Instruction, Textual Input>	0.4M 0.8M 48K 9K 46K&46K <1K&3K 7K&10K 20K&20K 6K&33K 0.2M&0.2M	1.8K 3.6K <1K <1K <1K&<1K <1K&<1K <1K&<1K <1K&<1K <1K&<1K	Open-ended
Paralinguistic Task Category						
1968 SGC	1969 Speaker Gender Classification (Biologically)	VoxCeleb1 VCTK VocalSound CommonVoice VCTK	<Speaker-centric	35K&35K 71K&71K 20K&20K 0.7M&0.7M 71K&71K	<1K&<1K <1K&<1K <1K&<1K 2.3K&1.2K <1K&<1K	Closed-ended
1970 AC	1971 Accent Classification	AccentDB CommonVoice	<AccentDB	16K&16K 0.3M&0.3M	<1K&<1K 2.4K&<1K	Closed-ended
1972 AR	1973 Age Recognition (Three Granularities)	VCTK VocalSound CommonVoice	<Age	71K&71K 20K&20K <Textual Instruction, Audio Input>	<1K&<1K <1K&<1K 1.2M&1.2M	Open-ended
1974 EIE	1975 Emotion Intensity Estimation	MELD CREMA-D	<Emotion	11K&11K 1K&1K	<1K&<1K <1K&<1K	
1976 SSD	1977 Synthetic Speech Detection	For	<Speech	9K&9K 7K&7K	<1K&<1K <1K&<1K	Closed-ended
1978 SV	Speaker Verification	MELD	<Speaker	64K&64K 11K&11K	<1K&<1K <1K&<1K	
1979 PSWL	Pronunciation Scoring Word Level	speechocean762	<Pronunciation	4K&4K 4K&4K	<1K&<1K <1K&<1K	Open-ended
Content-centric						
1980 PR	1981 Phoneme Recognition	Phonemizer	<Content	1M&1M	5K&4K	
SCR	Speech Command Recognition	GeneratedSpeechCommands	<Content	68K&68K	<1K&<1K	
IP	1982 Intent Prediction	SLURP	<Content	71K&71K	<1K&<1K	Open-ended
EE	Entity Extraction	VocalSound	<Content	45K&45K 20K&20K	<1K&<1K <1K&<1K	
VSC	1983 Vocal Sound Classification		<Content	0.2M&0.2M	1.1K&<1K	Closed-ended
IC	Instrument Classification		<Content	0.3M&0.3M	<1K&<1K	
ISC	Instrument Source Classification	NSynth	<Content	0.3M&0.3M	<1K&<1K	Open-ended
PP	Pitch Prediction		<Content	1.1K&<1K	Closed-ended	
VC	Velocity Classification		<Content	~25.5M	~89.5K	
Total						

1984 Table 11: Overview of task-level composition in LLASO-Align and LLASO-Instruct, spanning three
 1985 core categories, linguistic, semantic, and paralinguistic, across 20 sub-tasks. Each entry summarizes
 1986 representative data sources, supported input formats, and sample-level statistics. LLASO-Eval is
 1987 constructed as a stratified evaluation set presented in Table 12.

1998 R DETAILS FOR LLASO-EVAL

1999
 2000 To complete our data trio, we introduce LLASO-Eval, a held-out evaluation suite designed to ac-
 2001 company the LLASO training set. Derived from the same underlying corpus but separate from the
 2002 training split, LLASO-Eval covers 15,044 samples across 20 tasks, categorized into linguistic,
 2003 semantic, and paralinguistic categories. Moreover, it supports all three major modality configura-
 2004 tions and tests both within- and cross-modal generalization. We provide task breakdown in Table 12.

2005 From a task perspective, LLASO-Eval enables comprehensive evaluation of model capabilities across
 2006 three major categories: linguistic, semantic, and paralinguistic tasks. Within the paralinguistic cate-
 2007 gory, tasks are further distinguished into speaker-centric (e.g., gender, age, accent) and content-
 2008 centric (e.g., intent prediction, phoneme recognition). This distinction enables fine-grained analysis
 2009 of how models handle both speaker identity and acoustic-semantic information. From a modality
 2010 perspective, by supporting three major configurations, LLASO-Eval not only evaluates model perfor-
 2011 mance on seen modalities configures but also tresses cross-modal generalization, testing robustness
 2012 to novel input combinations. Additionally, to evaluate instruction-following capabilities, LLASO-
 2013 Eval includes both open-ended tasks for free-form comprehension and reasoning, and closed-ended
 2014 tasks requiring predefined label selection. This allows for quantitative measurement of instruction
 2015 adherence through metrics such as abstention rate.

2016 Tasks	2017 Descriptions	2018 Data Sources	2019 Modality Formats	2020 Sample Num.	2021 Metrics	
<i>Linguistic Task Category</i>						
2022 ASR	2023 Automatic Speech Recognition	2024 GigaSpeech LibriSpeech LJ Speech VCTK MLS	2025 <Textual Instruction, Audio Input >	2026 4566	2027 WER&CER	
<i>Semantic Task Category</i>						
2028 AQA	2029 Audio Question Answering	2030 Open Orca 1M-GPT4 Mukherjee et al. (2023) Open Orca 3.5M-GPT3.5 Stanford Alpaca Taori et al. (2023) Code Alpaca Chaudhary (2023) AlpacaDan Jordan (2023) Dolly Conover et al. (2023) OpenOrcaNo RuterNorway (2023) Tigerbot_Alpaca Research (2023) Tigerbot_Multichat Chen et al. (2023) Unnatural Honovich et al. (2022)	2031 <Textual Instruction, Audio Input > and <Audio Instruction, Textual Input >	2032 100 100 100 100 100&100 100&100 100&100 100&100 100&100	2033 GPT-4o	2034
<i>Paralinguistic Task Category</i>						
2035 SGC	2036 Speaker Gender Classification (Biologically)	2037 VoxCeleb1 Nagrani et al. (2017) VCTK VocalSound Gong et al. (2022) Common Voice Ardila et al. (2019) VCTK	2038 <Textual Instruction, Audio Input >	2039 100&100 100&100 200&200 100&100 100&100 100&100	2040 ACC	2041
2039 AC	2040 Accent Classification	2041 AccentDB Ahamed et al. (2020) Common Voice VCTK VocalSound Common Voice VCTK	2042 <Textual Instruction, Audio Input >	2043 100&100 100&100 100&100 100&100 100&100	2044 ACC	2045
2041 AR	2042 Age Recognition (Three Granularities)	2043 MELD Portia et al. (2018) CREMA-D Cao et al. (2014) MELD CREMA-D	2044 <Textual Instruction, Audio Input >	2045 100&100 100&100 200&200	2046 ACC/MAE	2047
2042 EIE	2043 Emotion Intensity Estimation	2044 MELD Portia et al. (2018) CREMA-D Cao et al. (2014) MELD CREMA-D	2045 <Textual Instruction, Audio Input >	2046 100&100 100&100 100&100	2047 ACC	2048
2043 ER	2044 Emotion Recognition	2045 FoR Reimao & Tzepos (2019) MELD	2046 <Textual Instruction, Audio Input >	2047 100&100	2048 ACC	2049
2044 SSD	2045 Synthetic Speech Detection	2046 FoR Reimao & Tzepos (2019) MELD	2047 <Textual Instruction, Audio Input >	2048 100&100	2049 ACC	2050
2045 SV	2046 Speaker Verification	2047 MELD	2048 <Textual Instruction, Audio Input >	2049 100&100	2050 GPT-4o	2051
2046 PSWL	2047 Pronunciation Scoring Word Level	2048 speechcean762 Zhang et al. (2021)	2049 <Textual Instruction, Audio Input >	2050 200&200	2051 ACC&GPT-4o	
2047 PSSL	2048 Pronunciation Scoring Sentence Level	2049 <Content-centric	2050 <Textual Instruction, Audio Input >	2051 200&200	2052 ACC&GPT-4o	
2048 PR	2049 Phoneme Recognition	2050 Phonemizer Generated Bernard & Titoux (2021)	2051 <Textual Instruction, Audio Input >	2052 100&100	2053 PER	
2049 SCR	2050 Speech Command Recognition	2051 Speech Commands Warden (1804)	2052 <Textual Instruction, Audio Input >	2053 100&100	2054 WER&CER&GPT-4o	
2050 IP	2051 Intent Prediction	2052 SLURP Bastianelli et al. (2020)	2053 <Textual Instruction, Audio Input >	2054 85&858	2055 GPT-4o	
2051 EE	2052 Entity Extraction	2053 VocalSound	2054 <Textual Instruction, Audio Input >	2055 569&569	2056 GPT-4o	
2052 VSC	2053 Vocal Sound Classification	2054 NSynth Engel et al. (2017)	2055 <Textual Instruction, Audio Input >	2056 200&200	2057 ACC	
2053 IC	2054 Instrument Classification	2055 100&100	2056 <Textual Instruction, Audio Input >	2057 100&100	2058 ACC	
2054 ISC	2055 Instrument Source Classification	2056 100&100	2057 <Textual Instruction, Audio Input >	2058 112&112	2059 MAE	
2055 PP	2056 Pitch Prediction	2057 100&100	2058 <Textual Instruction, Audio Input >	2059 100&100	2060 ACC	
2056 VC	2057 Velocity Classification	2058 15044	2059 <Textual Instruction, Audio Input >	2060 -	2061 -	
2057 Total	2058 -	2059 -	2060 -	2061 -	2062 -	

2043 Table 12: Overview of LLASO-Eval composition. This stratified evaluation set, sampled from LLASO-Instruct,
 2044 includes 20 tasks across linguistic, semantic, and paralinguistic categories (sub-divided into speaker-centric
 2045 and content-centric). For each task, we provide data sources, modality formats, sample counts, and evaluation
 2046 metrics. Automatic metrics are used where applicable, with GPT-4o-based judgment for open-ended tasks.
 2047
 2048
 2049
 2050
 2051

2052 **S BASELINE PERFORMANCE DETAILS**

2053

Task	Qwen2- Audio	Typhoon- Audio	Salmonn	Glm-4- Voice	Mini- Omni	Mini- Omni2	Llama- Omni	Audio- Reasoner	Kimi- Audio	Owen2.5- Omni	LLaSO-Base	LLaSO-Base (Unfrozen)	Metrics	
ASR	0.22 0.12	0.11 0.06	0.86 0.69	0.93 0.79	0.95 0.81	0.95 0.80	0.88 0.73	0.28 0.12	0.14 0.05	0.40 0.26	0.08 0.03	0.14 0.07	WER \downarrow CER \downarrow	
2057	2.41	1.76	1.47	2.22	1.42	1.57	1.97	2.44	2.94	2.94	2.06	2.27		
2058	2.42	1.77	1.41	2.34	1.47	1.53	2.02	2.24	2.70	3.09	1.80	2.27		
2059	2.73	2.16	1.41	3.29	1.75	2.05	2.99	2.51	3.22	3.22	2.39	2.23		
2060	2.78	2.22	1.72	2.93	1.45	1.51	2.48	2.86	3.45	2.63	1.46	1.98		
AQA	2.56 3.47 3.49 3.62	1.87 2.69 3.14 2.91	2.05 2.04 3.13 3.03	2.49 3.09 4.06	1.63 1.42 1.54 1.32	1.66 1.68 1.64 1.50	2.98 2.73 2.95 3.78	2.22 2.84 3.42 3.95	3.28 3.69 3.77 4.01	2.99 3.46 3.80 3.88	2.57 2.72 2.48 2.62	2.54 3.22 2.42 2.84	GPT-4o \uparrow	
2061	2.13 3.29 3.14 1.29	1.61 2.47 2.83 1.68	1.42 2.42 2.96 1.83	2.51 1.68 3.11 1.03	1.22 1.17 2.34 1.21	1.26 1.41 2.52 1.29	1.88 2.29 3.16 1.11	2.12 2.88 3.07 1.54	3.35 3.38 3.53 1.16	3.20 3.58 2.96 1.19	1.71 2.28 2.74 2.23	1.96 2.47 2.80 2.43		
2062	3.13 3.14 2.20 2.52	3.04 3.08 2.36 1.91	3.12 3.19 2.37 1.58	2.82 3.10 1.97 1.98	1.33 1.27 1.41 1.20	1.42 1.31 1.43 1.28	2.72 3.08 2.20 2.09	2.91 3.13 2.14 2.09	3.38 3.16 2.90 2.60	3.19 3.15 2.71 2.77	3.05 3.74 2.12 2.42	3.09 3.68 2.87 3.26		
2063	1.00 0.95	0.85 0.77	0.59 0.44	0.11 0.12	0.14 0.00	0.11 0.00	0.36 0.26	0.38 0.32	0.98 0.97	0.53 0.41	0.96 0.99	0.88 0.99		
2064	0.67 0.99	0.59 0.67	0.13 0.18	0.04 0.07	0.00 0.00	0.02 0.00	0.03 0.26	0.38 0.37	0.66 0.81	0.40 0.35	0.76 0.91	0.61 0.97	ACC \uparrow	
2065	0.16	0.21	0.22	0.07	0.00	0.00	0.07	0.14	0.38	0.06	0.52	0.40		
2066	0.12	0.14	0.32	0.09	0.06	0.03	0.14	0.18	0.31	0.19	0.83	0.86	ACC \uparrow	
2067	0.05	0.11	0.10	0.03	0.00	0.00	0.07	0.02	0.20	0.02	0.73	0.78		
2068	0.23	0.10	0.26	0.02	0.00	0.00	0.16	0.23	0.17	0.11	0.70	0.68	ACC \uparrow	
2069	0.52 18.69	0.12	0.06	0.01	0.00	0.00	0.17	0.12	0.12	0.08	0.50	0.38		
2070	0.54	0.40	0.31	0.13	0.04	0.03	0.31	0.52	0.65	0.52	0.48	0.45		
2071	0.31	0.24	0.24	0.08	0.04	0.06	0.08	0.29	0.34	0.27	0.48	0.37	ACC \uparrow	
2072	0.24	0.28	0.30	0.14	0.04	0.00	0.16	0.32	0.52	0.29	0.17	0.16	ACC \uparrow	
2073	0.30	0.12	0.21	0.02	0.07	0.01	0.04	0.28	0.32	0.33	0.26	0.30	ACC \uparrow	
2074	SSD	0.46	0.50	0.10	0.11	0.12	0.26	0.35	0.63	0.42	0.99	0.99	ACC \uparrow	
2075	SV	0.25	0.20	0.19	0.04	0.03	0.08	0.16	0.22	0.15	0.32	0.16	ACC \uparrow	
2076	PR	1.19	3.08	1.82	0.90	0.92	0.97	1.46	1.08	1.58	1.28	0.03	0.03	PER \downarrow
2077	IP	2.60	0.98	1.09	1.00	1.00	1.00	20.19	2.40	1.00	3.53	0.04	0.05	WER \downarrow
2078	SCR	2.40	0.85	0.75	0.98	0.97	21.11	1.84	0.31	3.52	0.02	0.04	CER \downarrow	
2079	3.04	3.13	4.07	1.85	1.39	1.89	1.58	4.10	4.56	3.91	4.86	4.80	GPT-4o \uparrow	
2080	PSWL	2.52	1.86	1.88	1.78	1.26	1.26	1.80	2.29	2.05	2.00	3.93	3.90	GPT-4o \uparrow
2081	EE	2.73	2.85	3.29	2.34	1.42	1.60	2.34	3.78	3.57	2.78	3.57	3.44	GPT-4o \uparrow
2082	PSSL	1.86	2.04	1.32	1.62	1.24	1.16	1.28	2.61	3.30	1.25	2.90	2.68	GPT-4o \uparrow
2083	VSC	0.17	0.33	0.13	0.24	0.00	0.00	0.13	0.09	0.20	0.27	0.39	0.24	ACC \uparrow
2084	IC	1.95	1.71	1.38	1.84	1.46	1.54	1.38	2.03	2.76	2.12	2.80	2.66	GPT-4o \uparrow
2085	ISC	0.85	0.49	0.61	0.32	0.02	0.03	0.03	0.59	0.84	0.92	0.78	0.82	ACC \uparrow
2086	PP	0.60	0.16	0.16	0.03	0.06	0.03	0.12	0.17	0.38	0.44	0.60	0.55	ACC \uparrow
2087	VC	19.02	36.83	41.92	40.20	61.49	59.32	<i>Reject</i>	32.68	31.64	18.37	8.02	10.55	MAE \downarrow
2088		0.02	0.17	0.22	0.08	0.00	0.00	0.07	0.15	0.19	0.12	0.18	0.20	ACC \uparrow

Table 13: Comprehensive evaluation across multiple LSLMs in LLaSO-Eval. Blue highlights denote best performance per task. *Reject* indicates that, during manual inspection, for 95% or more of the responses in the corresponding open-ended setting/task, the model explicitly expresses inability to assist or process the task, states it is a text-only model unable to recognize audio, or behaves as a pure text model by asking the user to describe the audio, its content, or information therein. From SGC to VC we only tested the <Textual Instruction, Audio Input>format, because for these tasks we also used only the <Textual Instruction, Audio Input>format data of those tasks during training.

2106 T DISCLOSURE OF LLM ASSISTANCE
21072108 We used LLM-based assistants **only** to aid and polish writing. Assistance was limited to grammar
2109 and style edits such as tightening wording, improving flow and transitions, shortening captions,
2110 harmonizing terminology, and light LaTeX phrasing. All technical content, claims, experimental
2111 design, data processing, modeling, analysis, figures/tables, and conclusions were conceived and
2112 produced by the authors; LLMs did not contribute novel ideas, code, datasets, evaluations, or result
2113 interpretations and are not contributing authors.
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159