



Enhancing Low-Resource Language and Instruction Following Capabilities of Audio Language Models

Potsawee Manakul¹, Guangzhi Sun², Warit Sirichotedumrong¹, Kasima Tharnpipitchai¹, Kunat Pipatanakul¹

¹SCB 10X, SCBX Group, Thailand

²Department of Engineering, University of Cambridge, United Kingdom

{potsawee, warit, kasima, kunat}@scb10x.com, gs534@cam.ac.uk

Abstract

Audio language models process audio inputs using textual prompts for tasks like speech recognition and audio captioning. Although built on multilingual pre-trained components, most are trained primarily on English, limiting their usability for other languages. This paper evaluates audio language models on Thai, a low-resource language, and finds that they lack emergent cross-lingual abilities despite their multilingual foundations. To address this, we explore data mixtures that optimize audio language models for both a target language and English while integrating audio comprehension and speech instruction-following into a unified model. Our experiments provide insights into improving instruction-following in low-resource languages by balancing language-specific and multilingual training data. The proposed model, Typhoon-Audio, significantly outperforms existing open-source models and achieves performance comparable to state-of-the-art Gemini-1.5-Pro in both English and Thai.

Index Terms: Audio Language Model, Large Audio Model, Thai Language Model, Low-resource Language Training

1. Introduction

Recent advancements in multimodal large language models (LLMs) have involved augmenting audio encoders, allowing models to extend their capabilities beyond textual inputs to include the understanding of speech and audio events [1–7]. Audio language models are developed for different use cases with common types as follows. Audio-understanding models, e.g., Qwen-Audio [1], SALMONN [2], and LTU [3], can process various types of audio inputs, such as speech, environmental sounds, or music, and generate textual outputs. These models are primarily designed for tasks that require a deep understanding of audio content, prompted by human instructions such as speech recognition, speech translation, or audio captioning. Alternatively, models such as SpeechGPT [5] accept speech as input and produce speech as output, functioning as voice agents. These models are optimized for seamless voice interactions rather than complex audio understanding. This work focuses on audio-understanding and instruction-following models, which we refer to as audio language models.

Audio language models typically comprise three key components: an audio encoder backbone, an LLM backbone, and an adapter module, as outlined in Section 2. Despite leveraging multilingual backbones, most models are primarily trained on: (1) English data, and (2) only audio content understanding tasks, as seen in models like Qwen-Audio [1] and SALMONN [2], or only speech instruction understanding, such as AudioChatLlama [6]. Addressing these limitations, this work focuses on two goals. First, we examine model performance in

a low-resource language using Thai as a case study, and we provide a recipe to enhance the low-resource language ability while retaining the English performance. Second, we integrate *improved* audio-understanding and speech instruction understanding capabilities into one unified model.

2. Related Work

Audio Language Models: *SALMONN* integrates three primary components: an LLM based on Vicuna [8], a speech encoder based on Whisper-large-v2 [9], and BEATS [10] for audio events. The representations from Whisper and BEATS are concatenated and passed through an adapter (Q-Former) to obtain the audio representation as the input to the LLM. Trainable modules are the Q-Former and the LoRA weights of the LLM. Training data consists of speech recognition, translation, audio captioning, or spoken QA. *Qwen-Audio*, similarly, uses Whisper-large-v2 encoder, and its LLM is based on Qwen [11]. No connection module (adapter) is employed; however, a variety of special tokens are incorporated. The pre-training phase involves training the audio encoder while freezing the LLM. During the fine-tuning phase, the audio encoder is frozen, and only the LLM is trained. *LTU* (first version) [3] was developed to support audio events with Llama-7B (LLM) and an AST audio encoder [12] as backbones. The AST encoder maps audio into 32 tokens, with a connection module adjusting dimensions from 768 to 4096. Subsequently, *LTU-AS* (second version) [4] extends support to both speech and audio, using Whisper-encoder, similar to SALMONN and Qwen-Audio. The training involved a curated dataset, with around 5 million triples (audio, question, answer) in the first version, and expanded to 9 million triples in the second version. These training examples were generated using GPT-3.5 based on audio metadata. Alternatively, *AudioChatLlama* aligns the semantic space of speech and text for direct speech-to-response generation. Since AudioChatLlama is optimized for responding to speech input, it cannot be controlled to perform other tasks through text instructions like SALMONN, Qwen-Audio, or LTU.

LLMs for Low-Resource Language: Previous studies have adapted English-centric unimodal LLMs into bilingual unimodal LLMs, enabling them to perform effectively in both English and a target language [13–16]. For example, Typhoon [14] is a Mistral/Llama3-based LLM enhanced on Thai. Typhoon was continually pre-trained on 50% English and 50% Thai data where Thai texts were sourced from the MC4 and OSCAR datasets before supervised fine-tuning. This process showed an improvement in Thai evaluation benchmarks. We note that while adapting LLMs to a new language is well-explored in unimodal LLMs, research on adapting audio language models to a new language remains underexplored.

3. Methodology

Model Architecture (in Fig. 1): We follow the SALMONN architecture [2]. With Thai and English as target languages, our model is based on Typhoon-1.5-8B-Instruct [14] as the LLM, Whisper-large-v3 fine-tuned to Thai [17] coupled with BEATs [10] as the audio encoder, and Q-Former [18] trained from scratch as the adapter. Note that we examine other variants of LLM and audio encoder backbones in Section 5.2.




Figure 1: *The model architecture of Typhoon-Audio. The audio encoder consists of a Whisper encoder and a BEATs encoder. The adapter is based on a window-level Q-Former. The LLM is our Typhoon model.*

Training Strategies: The audio encoder maps the spectrogram into a representation, which is then transformed into the audio representation a in the text embedding space via an adapter. The model θ is trained to maximize the probability of the next word y_t in the textual response, conditioned on previous words $y_{1:t-1}$, textual prompt input x and the audio input a : $P(y_t|y_{1:t-1}, x, a; \theta)$. Training occurs in two phases:

- **Pre-training:** As the adapter is the only component initialized with random weights, this phase trains only the adapter to align audio and textual representations. We use ASR and audio captioning data shown in Tab. 1 in this phase.

- **Supervised Fine-Tuning (SFT):** This phase trains both the adapter and the LoRA weight [19] of the LLM ($r=8, \alpha=32$). During SFT, the model is trained on diverse tasks and instruction prompts to enhance its instruction-following capabilities. Tab. 2 presents the data used in our final model (Typhoon-Audio), and Section 5.3 studies the SFT data mixture.

4. Experimental Setup

Table 1: *Pre-training data – 1.82M examples in total*

Dataset	Task	Lang	#Examples
LibriSpeech [20]	ASR	En	281K
GigaSpeech-M [21]	ASR	En	900K
CommonVoice-Th [22]	ASR	Th	436K
Fleurs-Th [23]	ASR	Th	7.8K
Vulcan+Elderly+Gowajee	ASR	Th	65.1K
AudioCaps [24]	AudioCaption	En+Th	48.3+48.3K
Clotho [25]	AudioCaption	En+Th	19.2+19.2K

Data: Each example has {audio, textual_prompt}. For *pre-training* data (Tab. 1), a few prompts were pre-defined for each task, e.g., “Transcribe this audio” for ASR, and we sample a prompt language that matches the response language. Since there is no Thai audio-captioning data, we translate AudioCaps and Clotho to Thai. For *SFT* data (Tab. 2), we translate 10% of prompts/responses in existing QA data to Thai. To increase

Table 2: *SFT data of Typhoon-Audio – 640K examples in total*

Dataset	Task	New	#Ex
QA pairs from SALMONN used in SFT-v1, SFT-v2, SFT-v3			
LibriSpeech [20]	QA (Speech-En)	✗	40.0K
AudioCaps [24]	QA (Audio)	✗	30.0K
QA pairs from LTU-AS used in SFT-v1, SFT-v2, SFT-v3			
LibriTTS [26]	QA (Speech-En)	✗	21.1K
IEMOCAP [27]	QA (Speech-En)	✗	4.3K
FSD50K [28]	QA (Audio)	✗	11.5K
AudioSet [29]	QA (Audio-Speech)	✗	20.0K
AS20k [30]	QA (Audio-Speech)	✗	12.0K
ASR, Translation, Audio Caption, QA used in SFT-v2, SFT-v3			
LibriSpeech [20]	ASR (En)	✗	32.0K
CommonVoice-Th [22]	ASR (Th)	✗	52.0K
SelfInstruct-Th	ASR (Th)	✓	18.9K
AudioCaps(Gemini)	Audio Caption	✓	48.3K
Covost2 [31]	Translate (X2Th)	✗	30.0K
CommonVoice-Th [22]	Translate (Th2X)	✗	7.3K
VISTEC-SER [32]	QA (Emotion&Gender)	✓	18.0K
Yodas2-30S [33]	QA (Speech-Th)	✓	90.0K
Speech Instruction Following used in SFT-v3			
GigaSpeech [21]	SpeechIF-Type1 (En)	✓	20.0K
CommonVoice-Th [22]	SpeechIF-Type1 (Th)	✓	120.5K
jan-hq-instruction-v1 [34]	SpeechIF-Type2 (En)	✗	20.0K
Airoboros-Th	SpeechIF-Type2 (Th)	✓	5.7K
Alpaca-Th	SpeechIF-Type2 (Th)	✓	20.0K
SelfInstruct-Th	SpeechIF-Type2 (Th)	✓	18.9K

the diversity of prompts, we generate a prompt for each example in ASR, translation, and audio-captioning, using GPT-4o. For speech instruction following, the model is meant to listen to spoken instructions and respond, thus, the prompt is null. Next, we describe each SFT task in more detail as follows:

- **ASR:** Datasets are used as shown in Table 2. A prompt is, for example, “Transcribe this audio” as well as its paraphrases.

- **Audio Caption:** This task involves generating audio descriptions, and we use AudioCaps [24], with English references translated into Thai for Thai Audio Captioning. Although AudioCaps was used for pre-training, its short ground-truth captions limit detailed response generation during SFT. To address this limitation, we provide Gemini-1.5-Pro with both audio input and the short caption, prompting it to generate detailed responses. This augmented data is called AudioCaps (Gemini).

- **Speech Translation:** Thai-to-English is from CommonVoice (Thai), and target English texts are derived from translation. English/X-to-Thai is from Covost2, and target Thai texts are derived from translating English texts. X refers to a non-English audio language (Arabic, German, Spanish, French, Indonesian, Italian, Japanese, Chinese) taken from Covost2. The translation was performed using our internal system, which matches Google Translate API performance. Each setup includes 2000 examples. The evaluation metric is BLEU.

- **Spoken Document QA:** The QA examples in SALMONN and LTU are primarily based on short utterances, typically under 10 seconds. To achieve longer understanding capabilities, we chunked longer audio in Yodas2 [33] into 30-second segments and prompted GPT-4o to generate question-answer pairs from the reference texts. In addition to standard extractive QA pairs, we also generated multiple-choice questions (MCQ), as we found that this approach improved the SpokenQA performance, which will be discussed in Section 5.3. Furthermore, we focus on the Thai subset of Yodas2, as existing QA datasets predominantly consist of English spoken documents. Since Yodas2 emphasizes the content of speech, to capture the unique

characteristics of voices, we additionally generated QA pairs from the VISTEC-SER dataset [32], which includes metadata such as speaker gender and emotion. The evaluation metric is word overlap of generated response against reference using F1.

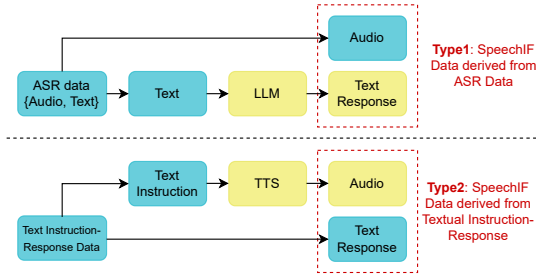


Figure 2: *Speech Instruction Following Data Creation Pipeline*

- *Speech Instruction Following (SpeechIF)*: This task requires models to listen to spoken instructions and directly respond. Current models like SALMONN lack specific data for this ability. We propose two methods for generating SpeechIF data (Fig. 2). *Type1* leverages ASR datasets to generate text responses from transcripts. However, since ASR data typically contains non-question utterances, LLMs often default to safe responses such as “I’m sorry, as an AI assistant I cannot...” in up to 30% of cases. While it offers voice diversity, it does not fully reflect real-world interactions. *Type2* synthesizes speech from instruction-response pairs (e.g., Alpaca, Airoboros), providing more practical commands but struggles with unsuitable instructions like math or coding. Though lacking voice diversity, *Type2* represents real interactions better. This work utilizes both types as shown in Table 2. For evaluation, we selected instructions from AlpacaEval (English) and SelfInstruct (Thai), creating a SpeechIF benchmark for both languages. The prompt for evaluating baseline models (e.g., SALMONN) is “Listen to the audio and answer the question”.

- *Complex Instruction Following (ComplexIF)*: We propose ComplexIF to assess models’ ability to follow unseen, compound instructions, where each instruction involves two to three audio tasks (such as transcribe, then translate). In ComplexIF, models have to respond in specific formats (e.g., JSON, XML), with format templates provided in the instruction prompt. As it evaluates the general instruction following ability, only English speech data is used. ComplexIF is used exclusively for evaluation, without additional training.

Evaluation: For existing tasks, we use standard metrics. For SpeechIF and ComplexIF, we follow the setup of MT-Bench in using an LLM judge (GPT-4o), and we adapt the single-turn evaluation prompt from MT-Bench and score responses on a scale from 1.0 to 10.0. For ComplexIF, we prompt the judge to evaluate the response on two aspects:

- *Quality* considers helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response.
- *Format* considers how well the response follows the format required by the user (e.g., JSON, XML, Markdown, etc).

Baselines: Competitive audio language models include,

- *Open-weights:* Qwen-Audio (Qwen-7B) [1], SALMONN (Vicuna-13B) [2], and DiVA (Llama3-8B) [35].
- *Proprietary:* Gemini-1.5-Pro (Audio)

For open models, we use available weights on Hugging-Face. For Gemini-1.5-Pro, we use `gemini-1.5-pro-001` through Google API with `{audio, text_instruction}` as input.

5. Results and Key Findings

5.1. Existing Audio Language Models on English vs Thai

The results in Table 4 demonstrate that: (1) Baseline models (Qwen-Audio, SALMONN, DiVA) despite using multilingual backbones exhibit significant performance degradation in Thai. Gemini-1.5-Pro maintains strong performance across both Thai and English. (2) Among the baselines, DiVA is the only model that performs well on the SpeechIF task, but it experiences a notable drop when tested on Thai. Thus, the subsequent experiments aim to develop a model that can effectively handle these tasks in both English and a low-resource language such as Thai.

5.2. Pre-training: Speech Encoder and LLM Backbones

This experiment focuses on selecting backbones, and comparing Whisper with its English+Thai fine-tuned variant. Similarly, Typhoon is a Llama-3 model fine-tuned to English+Thai. Our results (in Tab. 3) show that for ASR, models where both backbones are matched with the target language yield the best results. However, for audio captioning, the performance difference between these models is marginal. As a low-resource language such as Thai is our goal, Whisper-Th coupled with Typhoon-1.5 are selected for subsequent experiments.

Table 3: *Pre-training Results on ASR: LibriSpeech (other), CommonVoice (*subset-1K), AC: AudioCaps (En & Th). Whisper = Whisper-v3-large.*

SpeechEnc	LLM	ASR (WER↓)		AC (METEOR↑)	
		En	Th*	En	Th
Whisper	Llama-3	6.02	16.66	30.75	20.04
Whisper	Typhoon-1.5	7.76	20.01	29.56	20.62
Whisper-Th	Llama-3	7.35	15.68	29.52	19.94
Whisper-Th	Typhoon-1.5	9.15	13.52	30.83	20.55

5.3. Supervised Fine Tuning (SFT): Data Mixture

This experiment focuses on data mixture to enhance broad instruction-following abilities across tasks and languages. Training is initialized using the pre-trained model from the previous section. Our findings (results in Tab. 5) indicate that:

- *Pre-trained* model does not exhibit task ability and it simply provides transcriptions of speech regardless of instructions.

- *SFT-v1*: When fine-tuned on only English prompt-response pairs (a subset of around 600K pairs in total taken from SALMONN and LTU), the model achieves better performance on new tasks, but performs poorly on Thai ASR, showing similar characteristics to SALMONN shown in Table 4.

- *SFT-v2*: This setup translates 10% of SALMONN & LTU QA data into Thai and uses 10% Thai instructions for ASR, maintaining a 90/10 language ratio. Even with around 170K of mixed language QA pairs (due to a limited training budget) and 30K ASR examples, it improves SpQA and SpIF over the pre-trained model while achieving Thai ASR WER of 16.80. By adding around 20K of Thai speech QA, generated from 30-second chunks of Yodas2 and VISTEC-SER, the model achieves higher SpQA scores, and interestingly, it pro-

Table 4: *Audio LM Evaluation in English and Thai.* ASR: *En=LibriSpeech-other, Th=CommonVoice-17; Translation: Th-to-En=CommonVoice-17, En/X-to-Th=Covost2* where reference texts are derived from translation; *Gender Classification: En & Th = Fleurs; SpokenQA: En=SpokenSQuAD [36], Th=CommonVoice-17* where QA are generated from references using GPT-4o; *SpeechIF: En=AlpcaEval-TTS, Th=SelfInstruct-TTS.* *ComplexIF: Mixture of 5 other tasks in English.* *SpeechIF/ComplexIF $\in [1, 10]$.* *Typhoon2-Audio follows the same training methodologies of Typhoon-Audio, except for the change in the base LLM from Typhoon1.5 to Typhoon2. †We observed examples where the model did not provide any answer to nested commands; hence, receiving low scores.

Model	Size	ASR (WER↓)		Translation (BLEU↑)			Gender (Acc↑)		SpQA (F1↑)		SpeechIF (Judge↑)		CpxIF (Judge↑)	
		En	Th	Th2En	En2Th	X2Th	En	Th	En	Th	En	Th	Qual	Format
Qwen-Audio [1]	7B	6.94	95.12	0.00	2.48	0.29	37.09	67.97	25.34	0.00	1.07	1.03	3.13	1.68
SALMONN [2]	13B	5.79	98.07	14.97	0.07	0.10	95.69	93.26	52.92	2.95	2.47	1.18	4.10	5.09
DiVA [35]	8B	30.28	65.21	7.97	9.82	5.31	47.30	50.12	44.52	15.13	6.81	2.68	6.33	7.83
Gemini-1.5-Pro	-	5.98	13.56	22.54	20.69	13.52	90.73	81.32	74.09	62.10	3.24	3.93	7.25	8.99
Typhoon-Audio	8B	8.72	14.17	24.14	17.52	10.67	98.76	93.74	48.83	64.60	5.62	6.11	6.34	8.73
Typhoon2-Audio*	8B	5.83	14.04	33.25	27.15	15.93	76.51	75.65	69.22	70.01	6.00	6.79	5.35 [†]	9.01

Table 5: *SFT Results of different SFT recipes on various Thai Tasks and on English ComplexIF.* *ASR is evaluated on the subset-1K of CV17. †Avg. of Qual and Format.

Experiment	#Ex	ASR*↓	Th2En	SpQA	SpIF	CxIF [†]
Pre-trained	-	13.52	0.00	28.33	1.12	1.41
Typhoon-Audio (SFT-v1)						
100% En-Prompt	600K	80.86	6.01	36.88	1.48	6.35
Typhoon-Audio (SFT-v2)						
10% Th-Prompt	200K	16.80	0.00	35.26	3.72	5.08
+ QA	220K	16.93	0.02	46.82	4.29	5.33
+ QA+Trn	240K	18.33	21.53	44.93	4.25	5.97
+ 2*QA+Trn+MCQ	300K	19.84	22.04	61.63	4.60	6.31
Typhoon-Audio (SFT-v3)						
ScaledUp-v2+SpIF	620K	19.07	23.77	62.79	6.32	6.45
+ ASR (SelfInst-Th)	640K	16.89	24.14	64.60	6.11	7.54

vides a gain in SpIF as well. Speech translation ability is only acquired when explicitly adding around 20K translation examples in training (+Trn), though with minimal impact on other tasks. Next, we further add additional spoken QA and multiple-choice (+MCQ) of around 60K additional examples, showing further gains on SpQA, SpIF, CxIF tasks. Throughout these iterations, the improvement is also reflected in the CxIF score.

- **SFT-v3:** This configuration enhances SFT-v2 by incorporating speech instruction following data and increasing the total number of SFT examples by approximately twice. This approach improves performance across all tasks, especially SpeechIF. Due to limited gains from scaling, we did not construct more training examples. Compared to the pre-trained model, there is a large drop in ASR performance, with the WER increasing from 13.52 to 19.07. We observed that when the speech is a command, the model may unexpectedly give a response to the command even when tasked to perform ASR. To address this issue, we incorporated the SelfInstruct data as an ASR task in addition to the SpeechIF task. This approach lowered the WER to 16.89 while preserving other capabilities.

5.4. Typhoon-Audio vs Existing Audio Language Models

This experiment evaluates our best model, referred to as *Typhoon-Audio*, from the SFT section, benchmarking it against competitive models with results shown in Table 4. For ASR,

although Typhoon-Audio underperforms in English, it is one of only two models (along with Gemini) that achieves a WER under 15.0 on the Thai ASR benchmark. In translation, Typhoon-Audio surpasses SALMONN and Gemini-1.5-Pro in Thai-to-English, which requires Thai comprehension and English generation. Regarding voice characteristics, SALMONN’s English gender recognition transfers to Thai, with Typhoon-Audio showing comparable performance. In spoken document QA, Typhoon-Audio matches Gemini-1.5-Pro, making it the only open-source model capable of this task in Thai. For speech instruction following, while DiVA leads in English, it lacks Thai speech comprehension. Our Typhoon-Audio outperforms Gemini-1.5-Pro in both English and Thai for this task. Moreover, when presented with complex instructions, Typhoon-Audio performs the closest to Gemini-1.5-Pro in both quality and format adherence, making it the best open-source model.

Additionally, we explore the impact of upgrading the base LLM while using the same audio-text alignment (pre-training and SFT) methodologies. All of the experiments up to now were conducted using Typhoon1.5 as the base LLM. Given the new release of an improved version of Thai LLM (Typhoon2) at this stage of development, we investigate swapping the base LLM to Typhoon2 to develop **Typhoon2-Audio**. Results in Table 4 show improved performance of Typhoon2-Audio over Typhoon-Audio in ASR, speech translation, spoken QA, and speech instruction following. Our qualitative evaluation also shows that despite using the same audio-alignment data mixture, Typhoon2-Audio has a lower hallucination rate and reduced code-switching tendency than Typhoon-Audio, with hallucinations also lower than previous findings [37]. This suggests that the base LLM’s performance is crucial to the overall effectiveness of audio language models. However, both Typhoon-Audio versions struggle with background noise, likely due to training on mostly noise-free data, highlighting the need to incorporate noise during training for better robustness.

6. Conclusions

This paper demonstrates the limitations of low-resource language capabilities in open-source audio language models, using Thai as an example. Through our data mixtures which combine audio content understanding and speech instruction-following, we achieve performance on par with the proprietary Gemini-1.5-Pro in a range of audio tasks in both English and Thai with only around 1.82M pre-training and 0.64M SFT examples.

7. References

- [1] Y. Chu, J. Xu, X. Zhou, Q. Yang, S. Zhang, Z. Yan, C. Zhou, and J. Zhou, "Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models," *arXiv preprint arXiv:2311.07919*, 2023.
- [2] C. Tang, W. Yu, G. Sun, X. Chen, T. Tan, W. Li, L. Lu, Z. MA, and C. Zhang, "SALMONN: Towards generic hearing abilities for large language models," in *ICLR*, 2024.
- [3] Y. Gong, H. Luo, A. H. Liu, L. Karlinsky, and J. R. Glass, "Listen, think, and understand," in *ICLR*, 2024.
- [4] Y. Gong, A. H. Liu, H. Luo, L. Karlinsky, and J. Glass, "Joint audio and speech understanding," in *ASRU Workshop*, 2023.
- [5] D. Zhang, S. Li, X. Zhang, J. Zhan, P. Wang, Y. Zhou, and X. Qiu, "SpeechGPT: Empowering large language models with intrinsic cross-modal conversational abilities," *arXiv preprint arXiv:2305.11000*, 2023.
- [6] Y. Fathullah, C. Wu, E. Lakomkin, K. Li, J. Jia, Y. Shang-guan, J. Mahadeokar, O. Kalinli, C. Fuegen, and M. Seltzer, "AudioChatLlama: Towards general-purpose speech abilities for LLMs," in *NAACL*, 2024.
- [7] S. Hu, L. Zhou, S. Liu, S. Chen, H. Hao, J. Pan, X. Liu, J. Li, S. Sivasankaran, L. Liu *et al.*, "WavLLM: Towards robust and adaptive speech large language model," *arXiv preprint arXiv:2404.00656*, 2024.
- [8] W.-L. Chiang, Z. Li, Z. Lin, Y. Sheng, Z. Wu, H. Zhang, L. Zheng, S. Zhuang, Y. Zhuang, J. E. Gonzalez, I. Stoica, and E. P. Xing, "Vicuna: An open-source chatbot impressing GPT-4 with 90%* chatgpt quality," March 2023.
- [9] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *ICML*. PMLR, 2023, pp. 28 492–28 518.
- [10] S. Chen, Y. Wu, C. Wang, S. Liu, D. Tompkins, Z. Chen, W. Che, X. Yu, and F. Wei, "BEATs: audio pre-training with acoustic tokenizers," in *Proceedings of the 40th ICML*, 2023.
- [11] J. Bai, S. Bai, Y. Chu, Z. Cui, K. Dang, X. Deng, Y. Fan, W. Ge, Y. Han, F. Huang *et al.*, "Qwen technical report," *arXiv preprint arXiv:2309.16609*, 2023.
- [12] Y. Gong, Y.-A. Chung, and J. Glass, "AST: Audio Spectrogram Transformer," in *Proc. Interspeech 2021*, 2021, pp. 571–575.
- [13] A. Zeng, X. Liu, Z. Du, Z. Wang, H. Lai, M. Ding, Z. Yang, Y. Xu, W. Zheng, X. Xia, W. L. Tam, Z. Ma, Y. Xue, J. Zhai, W. Chen, Z. Liu, P. Zhang, Y. Dong, and J. Tang, "GLM-130b: An open bilingual pre-trained model," in *ICLR*, 2023.
- [14] K. Pipatanakul, P. Jirabovonvisut, P. Manakul, S. Sripaisarnmongkol, R. Patomwong, P. Chokchainant, and K. Tharnpipitchai, "Typhoon: Thai large language models," *arXiv preprint arXiv:2312.13951*, 2023.
- [15] Y. Cui, Z. Yang, and X. Yao, "Efficient and Effective Text Encoding for Chinese LLaMA and Alpaca," *arXiv preprint arXiv:2304.08177*, 2023.
- [16] Z. Csaki, B. Li, J. Li, Q. Xu, P. Pawakapan, L. Zhang, Y. Du, H. Zhao, C. Hu, and U. Thakker, "Sambalingo: Teaching large language models new languages," *arXiv preprint arXiv:2404.05829*, 2024.
- [17] A. Boribalburephan, Z. H. Aung, K. Pipatsrisawat, and T. Achakulvisut, "Thonburian Whisper: A fine-tuned whisper model for Thai automatic speech recognition," 2024. [Online]. Available: <https://huggingface.co/biodatlab/whisper-th-large-v3-combined>
- [18] J. Li, D. Li, S. Savarese, and S. Hoi, "Blip-2: bootstrapping language-image pre-training with frozen image encoders and large language models," in *Proceedings of the 40th International Conference on Machine Learning*, ser. ICML'23. JMLR.org, 2023.
- [19] E. J. Hu, yelong shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *ICLR*, 2022.
- [20] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [21] G. C. a Shuzhou Chai, G. Wang, J. Du, W.-Q. Zhang, C. Weng, D. Su, D. Povey, J. Trmal, J. Zhang, M. Jin, S. Khudanpur, S. Watanabe, S. Zhao, W. Zou, X. Li, X. Yao, Y. Wang, Y. Wang, Z. You, and Z. Yan, "GigaSpeech: An evolving, multi-domain asr corpus with 10,000 hours of transcribed audio," in *Proc. Interspeech 2021*, 2021.
- [22] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common Voice: A Massively-Multilingual Speech Corpus," in *Proceedings of the 12th Conference on Language Resources and Evaluation (LREC 2020)*, 2020.
- [23] A. Conneau, M. Ma, S. Khanuja, Y. Zhang, V. Axelrod, S. Dalmia, J. Riesa, C. Rivera, and A. Bapna, "Fleurs: Few-shot learning evaluation of universal representations of speech," *arXiv preprint arXiv:2205.12446*, 2022. [Online]. Available: <https://arxiv.org/abs/2205.12446>
- [24] C. D. Kim, B. Kim, H. Lee, and G. Kim, "Audiocaps: Generating captions for audios in the wild," in *NAACL-HLT*, 2019.
- [25] K. Drossos, S. Lipping, and T. Virtanen, "Clotho: an audio captioning dataset," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 736–740.
- [26] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "LibriTTS: A corpus derived from librispeech for text-to-speech," in *Proc. Interspeech 2019*, 2019.
- [27] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "IEMOCAP: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, 2008.
- [28] E. Fonseca, X. Favory, J. Pons, F. Font, and X. Serra, "Fsd50k: an open dataset of human-labeled sound events," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 829–852, 2021.
- [29] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 776–780.
- [30] S. Hershey, D. P. W. Ellis, E. Fonseca, A. Jansen, C. Liu, R. Channing Moore, and M. Plakal, "The benefit of temporally-strong labels in audio event classification," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 366–370.
- [31] C. Wang, A. Wu, J. Gu, and J. Pino, "Covost 2 and massively multilingual speech translation," in *Proc. Interspeech 2021*, 2021, pp. 2247–2251.
- [32] VISTEC, "Thai speech emotion dataset," 2021. [Online]. Available: <https://airesearch.in.th/releases/speech-emotion-dataset>
- [33] X. Li, S. Takamichi, T. Saeki, W. Chen, S. Shiota, and S. Watanabe, "Yodas: Youtube-oriented dataset for audio and speech," in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [34] A. Dao, D. B. Vu, and H. H. Ha, "Ichigo: Mixed-modal early-fusion realtime voice assistant," *arXiv preprint arXiv:2410.15316*, 2024.
- [35] W. Held, E. Li, M. Ryan, W. Shi, Y. Zhang, and D. Yang, "Distilling an end-to-end voice assistant without instruction training data," *arXiv preprint arXiv:2410.02678*, 2024.
- [36] C.-H. Lee, S.-L. Wu, C.-L. Liu, and H.-y. Lee, "Spoken SQuAD: A study of mitigating the impact of speech recognition errors on listening comprehension," *Proc. Interspeech 2018*, 2018.
- [37] G. Sun, P. Manakul, A. Liusie, K. Pipatanakul, C. Zhang, P. Woodland, and M. Gales, "CrossCheckGPT: Universal Hallucination Ranking for Multimodal Foundation Models," *arXiv preprint arXiv:2405.13684*, 2024.