

Acoustic Prompt Tuning: Empowering Large Language Models with Audition Capabilities

Jinhua Liang, *Student Member, IEEE*, Xubo Liu, *Member, IEEE*, Wenwu Wang, *Senior Member, IEEE*, Mark D. Plumbley, *Fellow, IEEE*, Huy Phan, *Senior Member, IEEE*, Emmanouil Benetos, *Senior Member, IEEE*

Abstract—The auditory system plays a substantial role in shaping the overall human perceptual experience. While prevailing large language models (LLMs) and visual language models (VLMs) have shown their promise in solving a wide variety of language and vision understanding tasks, only a few of them can be generalised to the audio domain without compromising their domain-specific capability. In this work, we introduce Acoustic Prompt Tuning (APT), a new adapter extending LLMs and VLMs to the audio domain by injecting audio embeddings to the input of LLMs, namely soft prompting. Specifically, APT applies an instruction-aware audio aligner to generate soft prompts, conditioned on both input text and sounds, as the inputs to the language model. To mitigate data scarcity in the audio domain, a curriculum learning strategy is proposed by formulating diverse audio tasks in a sequential manner. Moreover, we improve the audio language model by using interleaved audio-text embeddings as the input sequence. In this improved model, zero constraints are imposed on the input format, thus it is capable of tackling diverse modelling tasks, such as few-shot audio classification and audio comparison. To further evaluate the advanced ability of the audio networks, we introduce natural language audio reasoning (NLAR), a new task that analyses two audio clips by comparison and summarisation. Experiments show that APT-enhanced LLMs (namely APT-LLMs) achieve competitive results compared to the expert models (i.e., the networks trained on the target datasets) across various tasks. We finally demonstrate APT’s ability in extending frozen VLMs to the audio domain without fine-tuning, achieving promising results in audio-visual question and answering. Our code and model weights will be released at <https://github.com/JinhuaLiang/APT>.

Index Terms—Audio understanding, large language model, audio-language learning, audio recognition, automated audio captioning, natural language audio reasoning.

I. INTRODUCTION

AUDIO plays a critical role in shaping human perception, enriching our understanding of the world beyond visual cues alone. Due to the prevalence of large language models (LLMs) and the advancement of visual language models (VLMs) [1]–[5], recent research primarily focus on visual and language modalities. Despite of the notable advances that have been made in image and video understanding, only a few of them [4]–[6] have been extended to the audio domain.

To embrace more diverse modalities, more recent works have explored multi-modal tasks and heterogeneous data types. UniVAL [4] unified input/output format, model architecture, and training objective, and therefore, learned a

shared encoder-decoder LLM with multi-modal curriculum learning. ImageBind-LLM [5] adopted ImageBind, a cross-modal encoder bundling six modalities (including images) to a shared embedding space, and adapted the LLM with a frozen image encoder. While both approaches broaden the scope of visual LLMs, their reliance on specific architectures limits their adaptability for new modalities.

Following the VLM paradigm, several studies [7], [8] have introduced audio-exclusive LLMs where a pair of an audio clip and a text token is used as input for text generation. LTU [9] bridged audio with language modalities by end-to-end fine-tuning on an instruction-based dataset. Qwen-Audio [7] and SALMONN [10] expand such a framework into the speech and audio modalities via additional feature extraction using Whisper and BEATs encoders [11], [12]. Pengi [13] learned from a large number of off-the-shelf datasets by using a multi-task framework, but is restricted to two domains (audio and language). To identify audio with a few labelled examples, Audio flamingo [14] learned to predict the next text token from contextual information. However, these approaches are often restricted to fixed input formats, lacking flexibility for tasks beyond [audio, question, answer] structures, addressing task like few-shot audio classification [15]. This constraint leads us to question: *Can we adapt LLMs/VLMs to the audio domain by simply encoding sound clips as acoustic prompts?*

In this work, we introduce Acoustic Prompt Tuning (APT), an acoustic adapter that extends LLMs and VLMs to audio understanding and reasoning tasks by learning additional acoustic tokens (i.e., soft prompts) for the language model. Specifically, APT encodes audio clips into audio feature maps and employs an audio aligner (in Section III-A) to generate acoustic prompts. The acoustic prompts are conditioned on both the input instructions and the audio feature maps, aligning audio embeddings with the language modality. When training APTs, a curriculum learning strategy is adapted by formulating diverse audio tasks in a sequential manner. In addition to popular audio tasks, such as audio tagging (AT) and automated audio captioning (AAC), APT makes full use of publicly-available datasets where the model training is performed on three new tasks, namely query-based sound event detection, temporal event retrieval, and sound event counting, to learn fine-grained audio features. Moreover, we improve the audio language model framework by juxtaposing acoustic prompts with text embeddings. Rather than applying soft prompts as prefixes to the input texts, the improved framework exerts no constraints on the format of the input sequence. The APT-enhanced LLMs (APT-LLMs) interleave audio and text tokens to enhance in-

J. Liang and E. Benetos are with Centre for Digital Music (C4DM), Queen Mary University of London, UK. X. Liu, W. Wang, and M. D. Plumbley are with Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey, UK. H. Phan is with Meta Reality Labs, Paris, France.

The work does not relate to H. Phan’s position at Meta.

Email correspondence: jinhua.liang@qmul.ac.uk

context learning, preserve multimodal capabilities, and enable unified generative tasks. We have designed this approach to better align with our goals of making effective use of LLM capabilities while maintaining a generalizable and extensible framework. To further evaluate the advanced ability of our model, inspired by natural language visual reasoning [16], we propose a new task referred to as natural language audio reasoning (NLAR) which is devised to distinguish, compare, and summarise two audio clips. Experiments on existing audio understanding tasks, including audio tagging, automated audio captioning, and few-shot audio classification, show that APT-LLMs achieve performance on par with those obtained by audio language models or domain-expert models. APT also yields promising performance on the proposed NLAR task, indicating its ability to compare and summarise over two different audio clips. Finally, quantitative studies are conducted to demonstrate the improved performance of APT over a VLM in an audio-visual question and answering (AVQA) task. Our contributions are summarised as below:

- An acoustic adapter is introduced to extend LLMs and VLMs to the audio modality by soft prompting. To mitigate data scarcity issue, we introduce curriculum learning through a diverse set of tasks and a designed multi-stage training recipe. Leveraging the annotations in off-the-shelf databases, APT-LLM learns fine-grained acoustic embeddings from multiple audio tasks.
- APT formulates diverse audio tasks as a sequential task where generated text is conditioned on interleaved audio-text tokens. Without any constraints on the input format, APT-LLMs are not only able to solve different tasks according to the diverse instructions, but also to exploit the correlation among different audio clips in the same sequence. To our knowledge, APT-LLM is the first audio-language model analysing beyond a single audio clip.
- A new audio analysis task, natural language audio reasoning, is proposed to distinguish, compare, and summarise two audio clips. Compared to existing audio tasks, this new task not only evaluates the ability of the model to understand an audio clip, but also requires the models to analyse the content of two recordings by comparison and summarisation. APT-LLM is benchmarked on this task.
- A pretrained vision-language model, BLIP-2 [1], coupled with APT (namely APT-BLIP-2) is studied qualitatively and quantitatively on the audio-visual question and answering task [17]. APT-BLIP-2 can directly work with the visual modality without further fine-tuning, showcasing an efficient approach for extending multi-modal LLMs to a new modality.

II. RELATED WORKS

A. Multimodal language models

In terms of recent advances, LLMs [18]–[20] have exhibited excellent comprehension and reasoning ability. Driven by the open-world knowledge in LLMs, a variety of visual language models have been proposed with different alignment methods to integrate image/video data with text tokens [1], [3], [21], [22]. However, most of them are restricted to the visual domain.

In other domains (such as audio), the development of such models is limited, largely due to the lack of training data and the presence of modality discrepancies. Recently, ImageBind-LLMs [5] bridged the image encoder of ImageBind [23], a six-modality language model, with an LLM, and used visual tokens as soft prompts within the language model. UniVAL [4] unified the input/output, the architecture, and the training objective of multimodal LLMs and then devised a curriculum learning method for gradual exposure to the new modality. While both works adapted VLMs to other domains, they demand massive amount of multimodal data to train the overall networks from scratch. Instead, our proposed work investigates a domain-specific adapter that can extend any existing VLM/LLM to an additional modality (such as audio).

B. Audio language models

Following VLMs, some works built large audio pretrained models¹ for sound-only tasks. SpeechGPT [24] learned from a curated speech-text instruction dataset to analyse and generate spoken content in the audio domain. LTU [9] was trained upon an open-ended dataset that contains 3.7M [audio, question, answer] tuples, with a perception-to-understanding curriculum. While LTU achieved a good audio analysis ability, it required a uniform input format as a triplet tuple. To work around this issue, Pengi [13] proposed a multi-task framework where an audio language model is trained with off-the-shelf audio datasets prompted with different predefined questions. There are also audio language agents that apply LLM for audio manipulation [25], [26]. Our work differs from these prior works in three ways: 1) Rather than being an audio-only language model, APT explores how to adapt existing VLMs and LLMs to the audio domain; 2) APT-LLM improves the curriculum learning by designing three new training tasks. By accessing existing datasets from different aspects, APT-LLM learns a fine-grained audio representation; and 3) APT-LLM re-frames the present input format, namely [audio, question, answer], to allow audio and text to alternate in a sequence. In this way, APT-LLM is able to ingest more than one audio clip in a single feed-forward step, applying it to more audio tasks. To the best of our knowledge, APT-LLM is the first model that integrates in-context learning [27] with curriculum learning.

III. ACOUSTIC PROMPT TUNING

Current audio LLMs [9], [13] learned to bridge audio with language by framing popular audio tasks (e.g., classification [15], [28] and captioning tasks [29]) to the audio-conditioned text generation problem. Going beyond the [audio, question, answer] format, APT-LLM encodes multiple audio clips in one feed-forward process and juxtaposes them with text embeddings without any order constraint. This flexible training paradigm mitigates the need for high-quality annotations and massive databases, and thus reduces required computations. Moreover, juxtaposing audio clips with text enables APT-LLM to address more comprehensive audio analysis tasks, such as natural language audio reasoning. We first discuss the overall

¹These are also referred to as “audio foundation models” in the literature

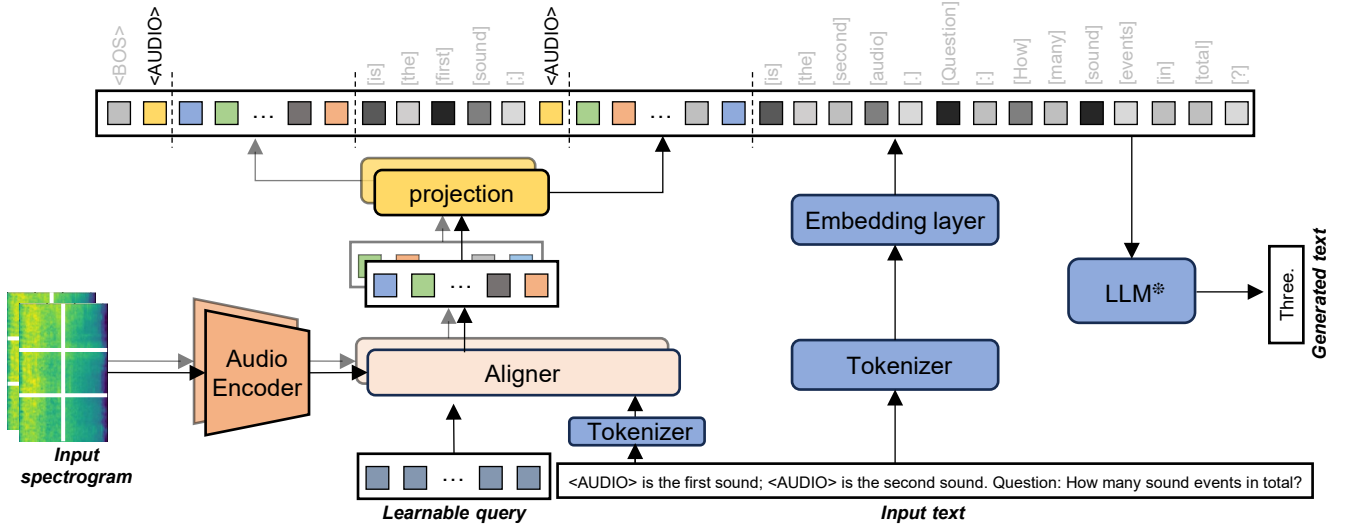


Fig. 1. Illustration of the proposed APT-LLM. APT-LLM includes three components: an audio encoder, an audio aligner, and an LLM. The audio encoder extracts audio feature maps from input spectrograms. The audio aligner then projects each audio feature map to 32 acoustic embeddings according to the input text. These acoustic embeddings, together with the added embeddings of the audio token “[AUDIO]”, are juxtaposed with text embeddings. The interleaved audio-text embeddings are fed into the LLM to generate the output text. APT-LLM can ingest multiple audio clips in a sequence and thus benefit from diverse tasks during training.

architecture of APT-LLM in Section III-A, and then elaborate the APT-LLM learning objective in Section III-B and the training recipe in Section III-C. In Section IV, we define the natural language audio reasoning task, a new task to evaluate the audio analysis ability of the proposed models.

A. Architecture

The overall structure of APT-LLM is illustrated in Figure 1, with the main components being an audio encoder, an audio aligner, and a large language model. APT-LLM alternates audio clips with text tokens without any format constraints and thus benefits from task diversity and large-scale pretraining.

Audio encoder: from spectrograms to feature maps. We use Audio-MAE [30], a vanilla 12-layer transformer encoder that learns to reconstruct randomly-masked spectrogram patches during training, as the audio encoder. Rather than using the last layer fine-tuned for classification tasks, we apply the output feature map from the penultimate block of the Audio-MAE to encode the fine-grained patterns in audio. To obtain audio features, the input audio is processed in segments of 10 seconds. We zero-pad the clip when the input audio is shorter than 10 seconds. When the input is longer than 10 seconds, we clip the audio into 10-second segments and pad the last segment to 10 seconds if necessary. The Audio-MAE encoder takes the audio input in segments of 10 seconds and produces feature maps of shape (1024, 512). The feature maps are then sent into the audio aligner and then concatenated before fed into the large language model.

Audio aligner: from feature maps to audio tokens. Audio aligner bridges the audio encoder to the frozen language model as shown in Figure 1. Figure 2 illustrates the structure of the audio aligner in APT. The audio aligner is constituted by N transformer blocks, each of which contains self-attention, cross-attention, and feed-forward layers. The aligner processes a text prompt alongside acoustic features of shape (1024,

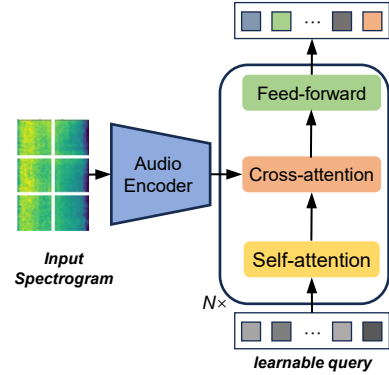


Fig. 2. Structure of audio aligner in APT.

512) from the audio encoder to produce a fixed number of acoustic embeddings. Learnable query tokens and text embeddings from a BERT tokenizer are taken as inputs, with the query tokens extracting auditory information from audio feature maps via cross-attention. For audio exceeding 10s, segments from audio embeddings are concatenated before being input into the large language model. Inspired by [1], we adopted a Query-transformer (Q-former) architecture with four transformer layers. Unlike prior work [1], our aligner specifically resamples variable-length audio embeddings into fixed acoustic embeddings, reducing computational complexity in downstream attention while filtering irrelevant information. The use of a BERT tokenizer aligns with the BERT-like transformer structure of our aligner. We set the number of learnable query tokens to 32 without additional specification.

Large language model: from interleaved audio-text tokens to generated text. As shown in Fig. 1, after encoding the whole text input with a dedicated text tokenizer, such as a LLAMA tokenizer [18], the large language model extract audio and text embeddings separately. The language model generates text tokens sequentially by leveraging both the extracted input

TABLE I

CURRICULUM LEARNING ADOPTED BY APT-LLMS. APT-LLMS ARE TRAINED WITH MULTIPLE TASKS IN DIFFERENT STAGES. “#AUDIO SAMPLES” DENOTE THE NUMBER OF AUDIO CLIPS IN THE DATASET. STAGE 0-2 DENOTES AUDIO-TEXT ALIGNMENT, LEARNING FROM SINGLE AUDIO CLIPS, AND LEARNING FROM MULTIPLE CLIPS, SEPARATELY.

Task	Training stages			Dataset	#Audio samples	Durations	Setup
	0	1	2				
Audio tagging	✓	✓	✓	AudioSet	2M	5.8kh	train/test
Audio captioning	✓	✓	✓	Wavcaps	400k	7.6kh	train/test
				AudioCaps	39k	108h	
				Clotho v2	7k	31h	
Audio question and answering		✓	✓	Clotho AQA	2k	12h	train
Query-based sound event detection		✓	✓	AudioSet-SL	81k	226h	train
Temporal event retrieval		✓	✓	AudioSet-SL	81k	226h	train
Sound event counting		✓	✓	AudioSet-SL	81k	226h	train
Few-shot audio classification			✓	AudioSet	2M	5.8kh	train/test
Natural language audio reasoning			✓	NLAR	0.2k	1.2h	train/test

audio-text embeddings and its previously generated tokens to inform each prediction step. This allows the model to maintain position-aware contextual information throughout the generation process. We freeze all parameters in the language model during training. Different from existing works [1], [9], we train a learnable audio token, “[AUDIO]”, to identify audio modality from other modalities. A learnable token is prepended to each audio clip as a special token to indicate the position of audio tokens in the interleaved sequence. We find this token helps the language model to distinguish audio tokens from text tokens when interleaving them together.

B. Learning objective

In order to motivate our training framework, we first present the learning objective used in existing work [13]. Let an audio-text pair in [audio, question, answer] format be referred to as (a, t, g) , where a , t , g are the audio clip, input text, and output text, respectively, and \mathbf{X} be the input sequential embeddings to the language model. To align the audio modality to the language modality, an audio encoder \mathcal{A} and an audio aligner \mathcal{M} project the audio a into a sequence $\mathbf{X}_{\text{audio}}$:

$$\mathbf{X}_{\text{audio}} = \mathcal{M}_{\theta}(\mathcal{A}_{\phi}(a), t), \quad (1)$$

where ϕ and θ are the parameters of the audio encoder \mathcal{A} and the aligner \mathcal{M} . The audio embeddings are used as a prefix to the input text embeddings as

$$\mathbf{X}_{\text{audio;text}} = \mathcal{C}(\mathbf{X}_{\text{audio}}, \mathbf{X}_{\text{text}}), \quad (2)$$

where \mathcal{C} is a concatenating function. By combining Eqn. 1 and Eqn. 2, the audio embeddings can be obtained by

$$\mathbf{X}_{\text{audio;text}} = \mathcal{C}(\mathcal{M}_{\theta}(\mathcal{A}_{\phi}(a), t), \mathcal{W}_{\psi}(t)), \quad (3)$$

where ψ denotes the parameters of the word embedding layer \mathcal{W} in the language model. Assuming the length of the concatenated embeddings $\mathbf{X}_{\text{audio;text}}$ be L , the parameters of

the audio LLM are optimised by measuring the probability distribution of the next token conditioned on its previous tokens:

$$p(\mathbf{X}_{\text{pred}} | \mathbf{X}_{\text{audio;text}}) = \prod_{i=L+1}^{L+|g|} p_{\phi, \theta, \psi}(\mathbf{x}_i | \mathbf{X}_{\text{audio;text}, < i}; \mathbf{X}_{\text{pred}, < i}), \quad (4)$$

where \mathbf{X}_{pred} is the predicted sequence, and \mathbf{x}_i is the i -th token of \mathbf{X}_{pred} . In this way, prevailing LLMs are able to unify many audio-to-text tasks in a sequential manner.

However, the [audio, question, answer] tuple format struggles to fully leverage rich context and structured information, hindering learning from contextual information. In contrast, the interleaved format improves in-context learning by exploiting structured information. For instance, a sequence like ‘(Input: a^0 , Output: t^0), (Input: a^1 , Output: t^1), (Input: a^2 , Output: g)’ allows audio features to be directly associated with their corresponding text descriptions, facilitating more effective learning and comprehension. We thereby propose a new learning framework in which interleaved audio-text embeddings are used as the LLM’s input such that the model is able to leverage and learn from more diverse tasks during training. Let \mathbf{a} and \mathbf{t} be audio clips and input texts, and g still be the output text. Assuming both \mathbf{a} and \mathbf{t} have N different elements, the input audio-text pairs are denoted as $[(a^i, t^i)]_{i=1}^N$ where a^i and t^i are the i -th audio clip and input text, respectively. Eqn. (2) can be re-written as

$$\mathbf{X}_{\text{audio;text}} = \mathcal{I}(\mathbf{X}_{\text{audio}}, \mathbf{X}_{\text{text}}), \quad (5)$$

where \mathcal{I} is the function that alternates acoustic embeddings with text embeddings. More specifically, the interleaved embeddings can be obtained by

$$\mathbf{X}_{\text{audio;text}} = [(S(a^1, t^1)), T_{\psi}(t^1), \dots, (S(a^N, t^N)), \mathcal{W}_{\psi}(t^N)], \quad (6)$$

where $S(\cdot, \cdot) = \mathcal{M}_{\theta}(\mathcal{A}_{\phi}(\cdot), \cdot)$. In this way, APT-LLM can integrate multiple audio clips in the input sequence, enabling itself to learn from more audio recognition tasks.

C. Curriculum learning strategy

With the uniform input/output format, APT-LLM is able to learn from a large variety of audio tasks and thus benefit from diverse training datasets. As shown in TABLE I, instead of passing through all training data directly, APT-LLM is trained through a three-stage curriculum:

Audio-text alignment. Before being coupled with an LLM, we pretrain the APT audio aligner to bridge the audio modality and the text modality. To this end, we freeze the audio encoder in the pipeline and optimise the parameters of the audio aligner with the audio-text pairs from AudioSet [31] and WavCaps [32]. During training, a fixed number of acoustic embeddings are learnt to extract relevant information from the audio feature maps according to the input text tokens. Following [1], the audio aligner is learned with the triplet objectives: audio-text matching, audio-grounded text generation, and audio-text contrasting (see more details in Appendix A).

Learning from a single audio clip. After APT has extracted acoustic embeddings according to the input text, the succeeding LLM learns to predict the next text tokens conditioned on both acoustic embeddings and the input text tokens. APT-LLM is thus trained with multiple tasks using various prompts (see more in Appendix B). In addition to existing audio tasks, namely audio tagging, automated audio captioning, and audio question and answering, we design three new tasks: (1) *Query-based sound event detection* that aims to train a model to predict the onset and offset time of a specific sound event; (2) *Temporal event retrieval* that is to recognise sound events occurred in a specific period, and (3) *Sound event counting* that requires a model to count the frequency of a specific sound event in a recording. Instead of rendering datasets, we exploit the publicly-available AudioSet with strong labels [33] using different prompts (see more in Appendix B). This curriculum learning strategy facilitates APT-LLM’s learning from diverse datasets, including AudioSet [31], WavCaps [32], AudioSet with strong labels [33], Clotho [34], AudioCaps [35], and Clotho-AQA [36].

Learning from multiple audio clips. In addition to the aforementioned tasks, APT-LLM learns from two additional tasks by juxtaposing more than one audio clips with the input text. Specifically, few-shot audio classification and natural language audio reasoning are added to the curriculum learning framework in this stage. On the one hand, for the few-shot audio classification, APT-LLM predicts labels of sound events by exploiting the correlation between input audio clips. On the other hand, APT-LLM is required to compare and summarise two different sounds in the natural language audio reasoning task (see the following Section IV). Trained on these two tasks, APT-LLM learns to analyse beyond a single recording and answer questions as per input questions. We adopt AudioSet [31] and the proposed datasets for few-shot audio classification and natural language audio reasoning, respectively.

D. Extending visual-language model to the audio domain

APT was also investigated as an audio adapter for an existing VLM, BLIP-2 [1]. BLIP-2 consists of a frozen image encoder, a 4-layer transformer, a projection layer, and a frozen Vicuna v1.1.

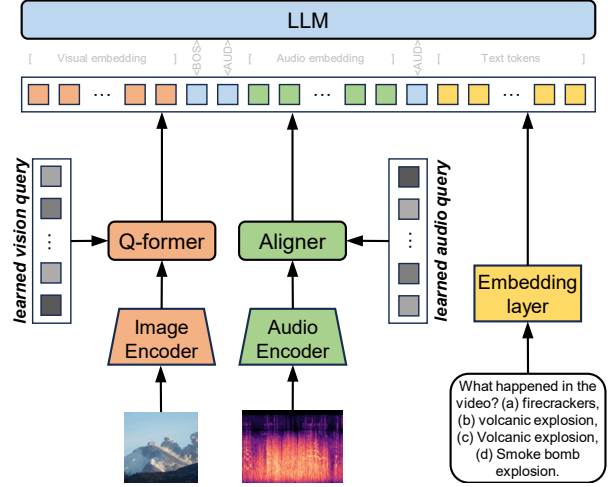


Fig. 3. Illustration of APT-BLIP-2 for the audio-visual question answering task. Text prompts for the Q-former and Aligner are omitted for simplicity.

Fig. 3 illustrates an example of APT-BLIP-2 for the audio-visual question answering task. Given raw audio a , an image input v , and a text query t , APT-BLIP-2 obtains audio embeddings $\mathbf{X}_{\text{audio}}$ and visual embeddings $\mathbf{X}_{\text{image}}$ with the audio aligner \mathcal{M} and the Q-former \mathcal{Q} , respectively. In particular, the audio aligner downsamples audio features to create a sequence of audio embeddings by $\mathbf{X}_{\text{audio}} = \mathcal{M}(\mathcal{A}(a), t)$. Meanwhile, the Q-former downsamples image features to a sequence of visual embeddings by $\mathbf{X}_{\text{image}} = \mathcal{Q}(\mathcal{V}(v), t')$ where \mathcal{V} denotes the image encoder and t' is a fixed text prompt as in [1]. Before incorporating $\mathbf{X}_{\text{image}}$, we concatenated audio and text embeddings in the format of “<BOS> <AUDIO> $\mathbf{X}_{\text{audio}}$ \mathbf{X}_{text} ”. Finally, following BLIP2’s approach [1], we prepended visual embedding $\mathbf{X}_{\text{image}}$ with the combined embeddings of audio and text, leading to the final format “ $\mathbf{X}_{\text{image}}$ <BOS> <AUDIO> $\mathbf{X}_{\text{audio}}$ \mathbf{X}_{text} ”.

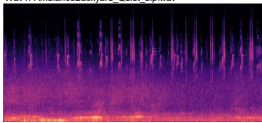
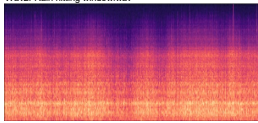
We refer to the APT-enhanced BLIP-2 as APT-BLIP-2. Of note, although we selected BLIP-2 as our backbone model, APT can be easily adapted to another vision language model.

IV. NATURAL LANGUAGE AUDIO REASONING

A. Task formulation

One of the complex reasoning abilities of humans is to learn from different sounds, understanding what is happening in each audio and analysing the content of different audio by comparison and summarisation. However, existing audio tasks focus on analysing acoustic scenarios in individual recordings by recognising the sound events [37] contained and/or retrieving their spatio-temporal information [38]. We thus propose natural language audio reasoning (NLAR), a new task where the model is required to answer questions by explicitly comparing or summarising two different audio recordings. TABLE II showcases an example of the NLAR task. A system takes two audio clips together with a free-form text query as input and is expected to answer the question by taking into consideration both audio recordings. Compared to existing audio tasks, the proposed audio comparison task features three notable differences:

TABLE II
AN EXAMPLE DEMONSTRATING APT-LLM’S ABILITY FOR AUDIO REASONING. THIS TASK REQUIRES AUDIO NETWORKS TO COMPREHEND RECORDINGS AND REASON ACROSS MULTIPLE RECORDINGS.

Natural Language Audio Reasoning (NLAR) example: “Where is the sudden sound?”	
User	
	<div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> <p>Wav1: “AmbianceBackyard_Quiet_bip.wav”</p>  </div> <div style="text-align: center;"> <p>Wav2: “Rain hitting window.wav”</p>  </div> </div>
	Question: “Which recording has a more sudden and startling sound event?”
APT-LLM	“First.”
Ground truth	“first”

Comparison and summary between two audio clips: This task requires a model to answer open-ended questions by comparing or summarising the content of two different audio clips. The model must first analyse the two audio clips as per the raised question separately and answer the question by taking into account the two audio inputs. An example of the audio comparison task can be found in TABLE III.

Diverse question types: Questions for the natural language audio reasoning task assess diverse auditory aspects, such as the presence, the frequency, and acoustic features of sound events. Therefore, the model should not only identify the sound events in the recordings, but also retrieve relevant information as per the input question.

Effects of the chronological order: Compared to existing audio tasks, e.g., [1] and [9], the proposed audio comparison task emphasises the order of the audio recordings in a pair. In other words, the answer associated with the audio pair “[Audio A, Audio B]” could be different with the answer associated with “[Audio B, Audio A]” when the questions are the same. In this way, we expect audio understanding models to be able to attend to different portions of the input sequence when the question is varied.

By evaluating audio language models on the natural language audio reasoning task, we achieve more comprehensive assessment of the audio language models.

B. Data curation

Natural language audio reasoning (NLAR) is proposed to evaluate a model’s ability in comprehending across audio clips as per the text instruction. We progressively render the NLAR dataset using the Clotho-AQA dataset [36], and every stage is described in the following:

Data collection: We filter out some audio samples from Clotho-AQA in terms of the quality of audio and the reliability of the annotations. To avoid data leakage, the test split of NLAR is collected from the test split of Clotho-AQA only.

Data cleaning: The annotation in the Clotho-AQA is noisy. For instance, annotators might not reach an agreement on the number of bird chirps in a recording. Therefore, we re-annotate the audio files by focusing on annotations where there is disagreement between annotators. We notice that it is infeasible to calculate the frequency of sound events in some cases, such as raining. In this case, we will annotate the presence of these activities throughout the recordings by using

TABLE III
AN EXAMPLE OF RENDERING THE NLAR DATASET

Based on the following two audio clips, generate 5 different questions that must be derived by summarising or comparing both audios and is prohibited to contain the information indicating its answer. The following information is provided: the sound events appear in the audio clip, together with its acoustic features, and corresponding onset and offset time stamps or frequency in the recordings. The answer should be either a binary one that can be responded by ‘yes’ or ‘no’, or an open-ended one that can be reply by a number or a single word.

Audios:

- First:
wood creaking (rustic, rhythmic, and creaky) [6 times];
second:
thunder (explosive, rumbling, and reverberating) [1 times];
rain (gentle, pitter-patter, rhythmic) [throughout the recording].

Questions and answers:

- are there 6 wood creaking and 2 thunder sounds in total? - no
- how many wood creaking or thunder are there in total? - 7
- does the second recording more likely to make one feel calm? - yes
- in which recording the events are more frequent? - first
- the frequency of wood creaking in the first recording is 6 times more than the frequency of thunder in the second one. - yes

Audios:

- First:
tapping glass (crisp, clear, and tingling sound) [9 times];
second:
water (rapid, and draining water sound) [throughout the recording].

Questions and answers:

- are there 9 tapping glass sounds and 3 water sounds in total? - no
- how many tapping glass or water sounds are there in total? - 10
- does the second recording create a continuous sound throughout? - yes
- in which recording are the sound events more repetitive? - first
- does the second recording sound more dynamic compared to the first recording? - yes

Audios:

- First:
wood creaking (rustic, rhythmic, and creaky) [6 times];
second:
shower (droplets, soothingly cascading) [throughout the recording].

Questions and answers:

tags such as “in the beginning”, “in the middle”, “in the end” and “throughout the recording”.

Describing acoustic features: We applied ChatGPT-turbo [20] to describe the acoustic feature of each sound event using prompt “describe the acoustic characteristic of {SOUND} precisely with a sentence less than 10 words” where {SOUND} refers to the name of the sound event. We inspect each audio feature description to ensure they are comprehensible.

Create audio-text pairs: We applied ChatGPT-turbo to generate five question-answer pairs according to a) their associated sound events; b) the temporal information, and c) audio description of the two audio recordings.

Inspect the rendered data: We checked the generated question-answer pairs using the annotations of the recordings and called ChatGPT again if the number of qualified pairs is lower than five. We dropped the generated pairs only if they contain some factual errors. We noticed that during data rendering, some answers to the questions could be unclear. We did not exclude them from the dataset as it is also important for models to learn what cannot be perceived from the audio. The curated NLAR dataset consists of 1722 audio-text pairs

TABLE IV

ZERO-SHOT AUDIO UNDERSTANDING PERFORMANCE (%) COMPARISON WITH AUDIO LANGUAGE MODELS. WE GROUP THE METHODS IN TERMS OF THEIR TRAINING STRATEGY. “#PARAMS.” DENOTES THE NUMBER OF TRAINABLE PARAMETERS AND “#PAIRS” REPRESENTS THE NUMBER OF AUDIO-TEXT PAIRS. \uparrow INDICATES THE HIGHER NUMBER, THE BETTER PERFORMANCE.

Model	#Params.	#Pairs	AudioSet (mAP \uparrow)	AudioCaps (SPICE \uparrow)	Clotho (SPICE \uparrow)
<i>Audio-language models trained with the contrastive loss</i>					
AudioCLIP [39]	30M	2M	25.9	-	-
CLAP [40]	190M	128k	5.8	-	-
<i>One-for-all models for various audio tasks</i>					
Qwen-Audio [7]	640M	-	18.5	-	13.6
LTU [9]	96M	5.7M	18.5	17.0	11.9
Pengi [13]	>191M	3.4M	-	18.2	12.6
APT-LLM	101M	2.6M	14.7	17.1	11.6

TABLE V

PERFORMANCE (%) COMPARISON IN AUTOMATED AUDIO CAPTIONING TASKS. \uparrow INDICATES THE HIGHER NUMBER, THE BETTER PERFORMANCE.

Model	AudioCaps		Clotho		Weighted average	
	SPICE \uparrow	SPIDEr \uparrow	SPICE \uparrow	SPIDEr \uparrow	SPICE \uparrow	SPIDEr \uparrow
<i>Specialised systems trained with task-specific examples</i>						
AT+CNN10 [41]	16.8	-	11.5	-	15.1	-
CNN-GPT2 [42]	16.7	43.8	11.1	21.5	14.9	36.7
WSAC+PD [43]	17.3	40.3	12.3	24.7	15.7	35.3
WavCaps [32]	18.2	48.5	13.3	29.7	16.6	42.5
<i>One-for-all models for various audio tasks</i>						
APT-LLM	19.1	40.2	13.2	24.8	17.2	35.3

in the training split and 1640 pairs in the test split.

V. EXPERIMENTS

APT was first coupled with LLMs (i.e., APT-LLM) and evaluated as a general-purpose audio learner on a variety of existing audio-related benchmarks, including audio tagging, automated audio captioning, and few-shot audio classification. To further assess its ability in analysing two audio clips, APT-LLM was further benchmarked on the natural language audio reasoning task. In addition to audio analysis, we tested and analysed (quantitatively and qualitatively) APT as a zero-shot adapter to BLIP-2 [1], [3], a state-of-the-art VLM.

A. Experimental setup

Our models were implemented based on the BLIP-2 framework [1]. We used Audio-MAE [30] as the audio encoder in all APT models we developed. Considering that Audio-MAE only contains 100M parameters, we used a four-layer transformer as the aligner to bridge the audio and text domains. Without an explicit statement, we coupled APT with a pretrained LLM named Vicuna 7B v1.1 [19] for evaluation. We tested APT-LLM with two close-ended datasets: AudioSet [31] and ESC-50 [44]; and four open-ended datasets: Clotho [34], AudioCaps [35], NLAR (in Section IV), and AVQA [17]. We did not evaluate APT models on new designed tasks, because there is no proper baseline to compare. Notably, we still trained APT-LLMs with those diverse objectives to build up a versatile audio learner.

Adam optimiser [45] was used for model training, with a warm-up strategy in the first 2000 steps and a cosine linear learning rate in the succeeding steps. We trained the models

using three NVIDIA A100 (40G) GPUs. Audio-text alignment pretraining and curriculum learning took 5 days each.

We used accuracy and mean average precision (mAP) as performance metrics to evaluate the audio systems on classification and QA datasets ² while using SPICE [46] and SPIDEr [47] on automated audio captioning datasets.

B. Comparison with existing approaches

We compare APT-LLM against the state-of-the-art specialised systems, i.e., the networks trained with task-specific data [15], [29], [32], [42], [43], [48], and audio-language models [7], [9], [13], [39], [40] on existing tasks, including audio tagging, automated audio captioning, and few-shot audio classification.

Audio tagging requires models to predict classes of test samples from a predefined label set. We evaluated the models on the AudioSet dataset [31]. During inference, APT-LLM was prompted using the sentence “*Summarize the audio with key words.*” Since APT generates free-form texts directly, we used the APT text encoder pretrained in the stage 1 to encode generated answers and the given class names to text embeddings. Afterwards, cosine similarity is calculated as the classification probability. Consistent with the findings in previous work [9], TABLE IV shows a performance gap between audio language models and task-specific models. This is expected since the latter model address the classification task as a close-set problem (as opposite to an open-set problem),

²The current audio question and answering datasets include single words as the ground truth so we use classification metrics in the evaluation following previous works [17], [36].

TABLE VI
ACCURACY (%) OF METHODS ON ESC-50 IN THE FEW-SHOT SETTINGS.

	Accuracy↑	
	5-way	12-way
<i>Specialised systems trained with task-specific examples</i>		
ProtoNet [49]	88.2	77.7
MatchNet [50]	86.8	71.8
HPN [51]	88.7	78.7
<i>Audio language models trained with contrastive learning</i>		
TIP-adapter [52]	97.5	95.6
Treff adapter [15]	98.5	96.3
<i>One-for-all models for various audio tasks</i>		
APT-LLM	91.0	54.2

with much lower complexity than open-set problem where models need to search across the entire word embedding space. In addition, we found that the performance of the text encoder greatly impacts the classification result when evaluating the generated answers. This finding can be attributed to the idea that word embeddings for different classes should remain sufficiently sparse when calculating their distance from the embeddings of the generated answer.

Audio captioning is the task where models are supposed to generate free-form description according to an input recording. The sentence “Describe the audio clip concisely.” is applied as the input prompt. To mitigate the vocabulary gap between the datasets and the language model, which leads to a lower score of pretrained audio-language models on the language-related metrics, we fine-tune APT-LLM for two epochs on the training split of AudioCaps [35] and Clotho [34] datasets and compare it with the captioning models trained on both tasks. As shown in TABLE V, APT-LLM achieves the best weighted average SPICE score of 17.2% while yielding a comparable SPIDeR score of 35.3% against task-specific systems on automated audio captioning datasets.

Few-shot audio classification aims to classify test audio clips using labelled audio examples. Models were evaluated in the N -way K -shot problem [48], [53] where: (1) there are N classes in the classification task, and (2) each class contains K different audio examples. Following previous works [50], in this task, we tested APT-LLM in the 5/12-way 5-shot settings where there are 5/12 types of sound events to classify and each type contains 5 audio examples. In our free-form query design, we prompt the few-shot classification question by adding the query audio clip together with the input text “This is a sound of” to the end of the sequence of labelled audio examples and their corresponding label texts. We implemented the same evaluation protocol to all few-shot learners for a fair comparison. As shown in TABLE VI, APT-LLM outperforms the task-specific models [49]–[51] in the 5-way 5-shot setting while having a competitive performance compared to CLAP Adapters [15], [52]. In the 12-way 5-shot problem, however, we can observe a performance degradation of APT.

We have observed some similar experimental findings in previous LLM research [54], [55]. One plausible hypothesis, as per the conclusion in [55], is that the distribution of examples within prompts impacts model performance. In contrast, the gradient-based optimization methods can avoid such issues

TABLE VII
BENCHMARKING APT FOR ACCURACY (%) ON THE NATURAL LANGUAGE AUDIO REASONING TASK.

Model	Accuracy↑
AAC+ChatGPT	27.9
APT-LLM	62.9
APT-LLM _{1.5}	63.8

TABLE VIII
ACCURACY (%) COMPARISON BETWEEN DIFFERENT MODALITIES IN AUDIO-VISUAL LEARNING.

Model	Modal	Accuracy↑
BLIP-2 [1]	Video-only	42.9
APT-LLM	Audio-only	27.7
APT-BLIP-2	Audio-video	59.7

in the few-shot setting. Additionally, it should be noted that while APT-LLM was trained with 4-way 1-shot tasks, it can generalise to other few-shot settings, suggesting that APT-LLM learns to act as a few-shot classifier rather than memorising the expected answers. More discussion on few-shot learning can be found in Appendix C.

C. Evaluation on natural language audio reasoning

Since APT-LLM is able to ingest multiple audio clips in a single feed-forward process, we investigated APT-LLM with natural language audio reasoning for which a model is expected to distinguish, compare, and summarise two audio clips (see Section IV). To the best of our knowledge, there is no previous work evaluating model ability to analyse more than one recording. We thus contrast APT-LLM to a cascaded model which connects an automated audio captioning model [29] with ChatGPT-4o [20]. TABLE VII compares different methods on the NLAR dataset. APT-LLM_{1.5} denotes the combination of our APT models and Vicuna v1.5. It can be observed that APT-LLM_{1.5} achieves 63.8% accuracy score, outperforming the baseline by a large margin. This result suggests that APT-LLM is able to not only recognise the content in an audio clip but also analyse more than one audio recording by comparison and summarisation. It is worth noting that there is marginal improvement when upgrading Vicuna from v1.1 to v1.5, indicating the performance of language models is not the bottleneck in this task, at least for the two used in our study.

D. Evaluation on zero-shot audio-visual tasks

APT-BLIP-2, together with other multimodal language models, was investigated on an audio-visual question and answering dataset where models are expected to choose one out of four options by using both audio and video modalities. We tested APT-BLIP-2 on the subset of the AVQA dataset [17] as many video links associated with the AVQA test segmentation were no longer available on the internet at the time of the experiment. As shown in TABLE VIII, APT-BLIP-2 yielded a better performance than video-only and audio-only counterparts, indicating the adaptation to the audio domain benefits models’ learning from the content of video.

E. Ablation study

TABLE IX compares the performance of the proposed audio aligner against a simple linear layer on AudioCaps and

TABLE IX
SPICE SCORE (%) COMPARISON BY ARCHITECTURE OF AUDIO ALIGNER.

	AudioCaps	Clotho
Linear	6.54	7.49
Transformer	17.1	11.6

TABLE X
SPICE SCORE (%) COMPARISON BY CHOICE OF AUDIO ENCODER.

	w/ aligner	AudioCaps	Clotho
CLAP		12.93	7.78
MAE		6.01	3.06
MAE	✓	15.83	11.30

Clotho datasets. For the comparison method, following [9] implementation, we used a linear layer to connect audio encoder and LLM and trained this method with the datasets in Table I. This result indicates that the transformer architecture significantly outperforms the linear architecture on both tasks, achieving a SPICE score of 17.1% on AudioCaps and 11.6% on Clotho, compared to the linear model’s scores of 6.54% and 7.49%, respectively. This experiment indicates that the transformer structure is highly effective for capturing complex audio patterns, resulting in substantially improved performance.

TABLE X compares the performance of two different audio encoders, MAE and CLAP, on the AudioCaps and Clotho datasets. All systems are trained on automated audio captioning and audio tagging datasets for simplicity. Since the MAE model was trained for spectrogram reconstruction, the length of output features is much longer than that of our audio aligner and the CLAP model. We evaluate two choices of downsampling approaches, mean average pooling in [9] and learnable downsampling, to match the input requirement of the subsequent language model. The results show that the CLAP encoder significantly outperforms the MAE encoder with mean average pooling on both tasks, with SPICE scores of 12.93% on AudioCaps and 7.78% on Clotho while underperforms that with the learnable downsampling module. One possible reason for this disparity is that the downsampling approach will lead to a loss of audio information, thus impacting MAE’s performance. We applied MAE as the audio encoder for APT because MAE preserves fine-grained information from audio and the audio aligner is expected to learn to extract the useful information by downsampling the audio features. We will leave the choice of the audio encoder for future exploration.

TABLE XI and TABLE XII assess the impact of curriculum learning by ablating the diverse set of tasks and multi-stage training recipe separately. TABLE XI shows that multi-task learning, training on both AAC and AT, yields the best performance on the downstream task, achieving 14.38% of weighted average SPICE score on AudioCaps and Clotho datasets, outperforming single-task training on either AAC or AT alone. Specifically, the AT-only model exhibits significant performance degradation, with scores of 5.27% weighted average SPICE score on AudioCaps and Clotho, suggesting the limitations of single-task pretraining on a task that may not align well with other downstream objectives, leading

TABLE XI
SPICE SCORE (%) COMPARISON BY ABLATING TRAINING TASKS INVOLVED.

AAC	AT	AudioCaps	Clotho	Weighted average
✓		13.69	11.37	12.95
	✓	7.64	2.89	6.12
✓	✓	15.83	11.30	14.38

TABLE XII
ZERO-SHOT SPICE SCORE (%) OF APT-LLM ON AUTOMATED AUDIO CAPTIONING DATASETS BY ABLATING THE CURRICULUM LEARNING DURING TRAINING.

	AudioCaps	Clotho
Single-stage	14.5	10.7
Multi-stage	17.1	11.6

TABLE XIII
ACCURACY (%) COMPARISON BY NLAR TASKS DURING TRAINING.

	w/o NLAR	w/ NLAR
Accuracy	16.4	63.8

to a less effective learning. In contrast, multi-task learning enhances the generalisation ability of APT-LLM by leveraging the heterogeneity across different tasks.

TABLE XII shows the zero-shot performance of APT-LLM on automated audio captioning datasets by ablating the curriculum learning during training. Multi-stage learning improves the performance of APT-LLM by 2.6% of SPICE on AudioCaps and 0.9% of SPICE on Clotho, which indicates the effectiveness of the proposed curriculum learning.

TABLE XIII compares model performance, measured by accuracy, when trained with and without NLAR tasks. The results show a substantial performance improvement when NLAR is included during training, with accuracy increased from 16.4% (without NLAR) to 63.8% (with NLAR). The result suggests that incorporating NLAR tasks is essential for achieving high accuracy, especially when the model is evaluated on tasks closely related to NLAR.

We evaluated how the choice of token length impacts on the performance of the proposed method. The results are shown in Fig. 4. We can see that APT with 32 learnable acoustic tokens achieved the best performance compared to configurations with 16 and 64 acoustic tokens.

VI. LIMITATIONS

The proposed method could be further improved in the following aspects. First, we devised APT to align acoustic embeddings with text embeddings obtained by language models. Word embeddings may change when a different language model is used, even if their architectures remain the same. As a result, each language model may need a dedicated APT aligner for adaptation. Second, APT-LLM was not trained with instruction-based datasets. Consequently, its ability is limited when responding to out-of-domain questions, which are not included in the training set. Finally, we focused the training and experimentation of APT-LLM on general-purpose audio

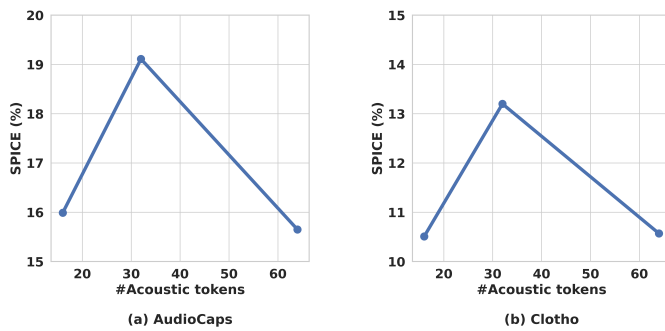


Fig. 4. Performance comparison on automated audio captioning datasets, AudioCaps and Clotho, by varying the number of acoustic tokens.

understanding tasks; thus its capability has yet to be expanded to the realm of speech and music audio.

VII. CONCLUSIONS

We proposed APT, a general-purpose acoustic adapter that extends LLMs/VLMs to the audio domain. We showed that LLMs coupled with APT are versatile audio learners that not only achieved a competitive performance across various audio analysis tasks but also are capable of in-context learning [56] when fed with a few labelled examples. We also benchmarked APT-LLM’s ability of audio comparison and summary using the natural language audio reasoning task, a new task that requires a model to distinguish, compare, and summarise two different audio clips. Finally, it is evident from our study on audio-visual learning that encoding sound clips as word tokens is helpful to adapt LLMs/VLMs to the audio domain.

Future works can extend audio language models to jointly model the mixture of audio, music, and speech as they are complimentary for audio content analysis, especially in complicated real-world scenarios [57]. This could involve the choices of the audio encoder in an audio-language model. Also, as suggested in [9], instruction tuning is promising to align audio analysis with human perception by exploiting the capacity of large language models. It is interesting to investigate how audio language models handle errors from in-context learning. Finally, scaling up audio reasoning and summary datasets is essential to benchmark audio language models on more complex tasks. One intriguing direction is to generate reliable annotations by using the existing fine-grained dataset, such as AudioSet strong-label, and large language models, e.g., ChatGPT. Another possible approach is to synthesise high-quality audio conditioned on target question-answer pairs by using generative models.

ACKNOWLEDGMENT

This work was supported by the Engineering and Physical Sciences Research Council [grant numbers EP/T518086/1 and EP/T019751/1]. E. Benetos was supported by a RAEng/Leverhulme Trust Research Fellowship [grant numbers LTRF22223-19-106]. X. Liu is supported by a Research Gift from Google. This research utilised Queen Mary’s Apocrita HPC facility, supported by QMUL Research-IT. <http://doi.org/10.5281/zenodo.438045>. For the purpose of open access, the authors have applied a Creative Commons attribution

(CC BY) license to any Author Accepted Manuscript version arising. The authors would like to thank the associate editor and the reviewers for their valuable comments.

REFERENCES

- [1] 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 (ICML’23)*. JMLR, 2023.
- [2] H. Liu, C. Li, Q. Wu, and Y. J. Lee, “Visual instruction tuning,” in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [3] W. Dai, J. Li, D. Li, A. Tiong, J. Zhao, W. Wang, B. Li, P. Fung, and S. Hoi, “InstructBLIP: Towards general-purpose vision-language models with instruction tuning,” in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [4] M. Shukor, C. Dancette, A. Rame, and M. Cord, “UnIVAL: Unified model for image, video, audio and language tasks,” *Transactions on Machine Learning Research*, 2023.
- [5] J. Han, R. Zhang, W. Shao, P. Gao, P. Xu, H. Xiao, K. Zhang, C. Liu, S. Wen, Z. Guo, X. Lu, S. Ren, Y. Wen, X. Chen, X. Yue, H. Li, and Y. Qiao, “Imagebind-llm: Multi-modality instruction tuning,” 2023. [Online]. Available: <https://arxiv.org/abs/2309.03905>
- [6] L. Meng, S. Hu, J. Kang, Z. Li, Y. Wang, W. Wu, X. Liu, and H. Meng, “Large language model can transcribe speech in multi-talker scenarios with versatile instructions,” *arXiv:2409.08596*, 2024.
- [7] 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.
- [8] S. Hu, L. Zhou, S. Liu, S. Chen, L. Meng, H. Hao, J. Pan, X. Liu, J. Li, S. Sivasankaran, L. Liu, and F. Wei, “WavLLM: Towards robust and adaptive speech large language model,” in *Findings of the Association for Computational Linguistics: EMNLP 2024*, Y. Al-Onaizan, M. Bansal, and Y.-N. Chen, Eds. Miami, Florida, USA: Association for Computational Linguistics, Nov. 2024, pp. 4552–4572.
- [9] Y. Gong, H. Luo, A. H. Liu, L. Karlinsky, and J. R. Glass, “Listen, think, and understand,” in *The Twelfth International Conference on Learning Representations*, 2024.
- [10] 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 *The Twelfth International Conference on Learning Representations*, 2024.
- [11] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *Proceedings of the 40th International Conference on Machine Learning*, ser. ICML’23. JMLR.org, 2023.
- [12] 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 International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds., vol. 202. PMLR, 23–29 Jul 2023, pp. 5178–5193.
- [13] S. Deshmukh, B. Elizalde, R. Singh, and H. Wang, “Pengi: An audio language model for audio tasks,” in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [14] Z. Kong, A. Goel, R. Badlani, W. Ping, R. Valle, and B. Catanzaro, “Audio Flamingo: A novel audio language model with few-shot learning and dialogue abilities,” in *Forty-first International Conference on Machine Learning*, 2024.
- [15] J. Liang, X. Liu, H. Liu, H. Phan, E. Benetos, M. D. Plumbley, and W. Wang, “Adapting language-audio models as few-shot audio learners,” in *Proc. INTERSPEECH*, 2023, pp. 276–280.
- [16] A. Suhr, S. Zhou, A. Zhang, I. Zhang, H. Bai, and Y. Artzi, “A corpus for reasoning about natural language grounded in photographs,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, A. Korhonen, D. Traum, and L. Márquez, Eds. Florence, Italy: Association for Computational Linguistics, Jul. 2019, pp. 6418–6428.
- [17] P. Yang, X. Wang, X. Duan, H. Chen, R. Hou, C. Jin, and W. Zhu, “AVQA: A dataset for audio-visual question answering on videos,” in *Proceedings of the 30th ACM International Conference on Multimedia*, ser. MM ’22. New York, NY, USA: Association for Computing Machinery, 2022, pp. 3480–3491, event-place: Lisboa, Portugal.
- [18] Llama team, “Llama 2: Open foundation and fine-tuned chat models,” Jul. 2023, arXiv:2307.09288.

- [19] 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," Mar. 2023.
- [20] OpenAI, "GPT-4 technical report," Mar. 2023, arXiv:2303.08774.
- [21] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, Y. Hasson, K. Lenc, A. Mensch, K. Millican, M. Reynolds, R. Ring, E. Rutherford, S. Cabi, T. Han, Z. Gong, S. Samangooei, M. Monteiro, J. Menick, S. Borgeaud, A. Brock, A. Nematzadeh, S. Sharifzadeh, M. Binkowski, R. Barreira, O. Vinyals, A. Zisserman, and K. Simonyan, "Flamingo: A visual language model for few-shot learning," in *Advances in Neural Information Processing Systems*, 2022.
- [22] R. Zhang, J. Han, C. Liu, A. Zhou, P. Lu, Y. Qiao, H. Li, and P. Gao, "LLaMA-Adapter: Efficient fine-tuning of large language models with zero-initialized attention," in *The Twelfth International Conference on Learning Representations*, 2024.
- [23] R. Giridhar, A. El-Nouby, Z. Liu, M. Singh, K. V. Alwala, A. Joulin, and I. Misra, "ImageBind: One embedding space to bind them all," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2023, pp. 15 180–15 190.
- [24] 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," in *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.
- [25] J. Liang, H. Zhang, H. Liu, Y. Cao, Q. Kong, X. Liu, W. Wang, M. D. Plumbley, H. Phan, and E. Benetos, "WavCraft: Audio editing and generation with large language models," in *ICLR 2024 Workshop on Large Language Model (LLM) Agents*, 2024.
- [26] X. Liu, Z. Zhu, H. Liu, Y. Yuan, Q. Huang, J. Liang, Y. Cao, Q. Kong, M. D. Plumbley, and W. Wang, "WavJourney: Compositional audio creation with large language models," arXiv:2307.14335, 2023.
- [27] J. Von Oswald, E. Niklasson, E. Randazzo, J. a. Sacramento, A. Mordvintsev, A. Zhmoginov, and M. Vladymyrov, "Transformers learn in-context by gradient descent," in *Proceedings of the 40th International Conference on Machine Learning (ICML'23)*. JMLR.org, 2023.
- [28] B. Ding, T. Zhang, C. Wang, G. Liu, J. Liang, R. Hu, Y. Wu, and D. Guo, "Acoustic scene classification: A comprehensive survey," *Expert Syst. Appl.*, vol. 238, Feb. 2024.
- [29] X. Xu, Z. Xie, M. Wu, and K. Yu, "The SJTU system for DCASE2022 challenge task 6: Audio captioning with audio-text retrieval pre-training," DCASE2022 Challenge, Tech. Rep., 2022.
- [30] P.-Y. Huang, H. Xu, J. B. Li, A. Baevski, M. Auli, W. Galuba, F. Metze, and C. Feichtenhofer, "Masked autoencoders that listen," in *Advances in Neural Information Processing Systems*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022.
- [31] 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)*. New Orleans, LA: IEEE, Mar. 2017, pp. 776–780.
- [32] X. Mei, C. Meng, H. Liu, Q. Kong, T. Ko, C. Zhao, M. D. Plumbley, Y. Zou, and W. Wang, "WavCaps: A ChatGPT-assisted weakly-labelled audio captioning dataset for audio-language multimodal research," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 32, p. 3339–3354, Jun. 2024.
- [33] S. Hershey, D. P. Ellis, E. Fonseca, A. Jansen, C. Liu, R. C. 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)*. IEEE, 2021, pp. 366–370.
- [34] K. Drossos, S. Lipping, and T. Virtanen, "Clotho: An audio captioning dataset," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 736–740.
- [35] C. D. Kim, B. Kim, H. Lee, and G. Kim, "AudioCaps: Generating captions for audios in the wild," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 119–132.
- [36] S. Lipping, P. Sudarsanam, K. Drossos, and T. Virtanen, "Clotho-AQA: A crowdsourced dataset for audio question answering," in *30th European Signal Processing Conference (EUSIPCO)*, 2022, pp. 1140–1144.
- [37] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, "PANNs: Large-scale pretrained audio neural networks for audio pattern recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [38] A. Politis, K. Shimada, P. Sudarsanam, S. Adavanne, D. Krause, Y. Koyama, N. Takahashi, S. Takahashi, Y. Mitsufuji, and T. Virtanen, "STARSS22: A dataset of spatial recordings of real scenes with spatiotemporal annotations of sound events," in *Proceedings of the 7th Detection and Classification of Acoustic Scenes and Events 2022 Workshop (DCASE2022)*, Nancy, France, November 2022.
- [39] A. Guzhov, F. Raue, J. Hees, and A. Dengel, "Audioclip: Extending clip to image, text and audio," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2022, pp. 976–980, iSSN: 2379-190X.
- [40] B. Elizalde, S. Deshmukh, M. A. Ismail, and H. Wang, "CLAP: Learning audio concepts from natural language supervision," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [41] X. Xu, H. Dinkel, M. Wu, Z. Xie, and K. Yu, "Investigating local and global information for automated audio captioning with transfer learning," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Jun. 2021, pp. 905–909.
- [42] M. Kim, K. Sung-Bin, and T.-H. Oh, "Prefix tuning for automated audio captioning," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Jun. 2023, pp. 1–5.
- [43] T. Kouzelis and V. Katsouros, "Weakly-supervised automated audio captioning via text only training," in *Proceedings of the 8th Detection and Classification of Acoustic Scenes and Events 2023 Workshop (DCASE2023)*, Tampere, Finland, September 2023, pp. 81–85.
- [44] K. J. Piczak, "ESC: Dataset for environmental sound classification," in *Proceedings of the 23rd ACM International Conference on Multimedia*. Brisbane Australia: ACM, Oct. 2015, pp. 1015–1018.
- [45] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.
- [46] P. Anderson, B. Fernando, M. Johnson, and S. Gould, "SPICE: Semantic propositional image caption evaluation," in *Computer Vision – ECCV 2016*. Cham: Springer International Publishing, 2016, pp. 382–398.
- [47] S. Liu, Z. Zhu, N. Ye, S. Guadarrama, and K. P. Murphy, "Improved image captioning via policy gradient optimization of SPIDER," *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 873–881, 2016.
- [48] J. Liang, H. Phan, and E. Benetos, "Learning from taxonomy: Multi-label few-shot classification for everyday sound recognition," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 771–775.
- [49] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," *Advances in Neural Information Processing Systems*, vol. 30, pp. 4077–4087, 2017.
- [50] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra, "Matching networks for one shot learning," in *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [51] J. Liang, H. Phan, and E. Benetos, "Leveraging label hierarchies for few-shot everyday sound recognition," in *Proceedings of the 7th Detection and Classification of Acoustic Scenes and Events 2022 Workshop (DCASE2022)*, Nancy, France, Nov. 2022.
- [52] R. Zhang, W. Zhang, R. Fang, P. Gao, K. Li, J. Dai, Y. Qiao, and H. Li, "Tip-Adapter: Training-free adaption of CLIP for few-shot classification," in *Computer Vision – ECCV 2022*, ser. Lecture Notes in Computer Science, S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds. Cham: Springer Nature Switzerland, 2022, pp. 493–510.
- [53] J. Liang, I. Nolasco, B. Ghani, H. Phan, E. Benetos, and D. Stowell, "Mind the domain gap: a systematic analysis on bioacoustic sound event detection," in *32th European Signal Processing Conference (EUSIPCO)*, 2024.
- [54] B. Jin, J. Yoon, J. Han, and S. O. Arik, "Long-Context LLMs meet RAG: Overcoming challenges for long inputs in RAG," arXiv:2410.05983, 2024.
- [55] T. Li, G. Zhang, Q. D. Do, X. Yue, and W. Chen, "Long-context LLMs struggle with long in-context learning," arXiv:2404.02060, 2024.
- [56] D. Dai, Y. Sun, L. Dong, Y. Hao, S. Ma, Z. Sui, and F. Wei, "Why can GPT learn in-context? Language models secretly perform gradient descent as meta-optimizers," in *Findings of the Association for Computational Linguistics: ACL 2023*, A. Rogers, J. Boyd-Graber, and N. Okazaki, Eds. Toronto, Canada: Association for Computational Linguistics, Jul. 2023, pp. 4005–4019.
- [57] H. Zhang, V. Cheung, H. Nishioka, S. Dixon, and S. Furuya, "LLaQo: Towards a query-based coach in expressive performance assessment," *Proceedings of International Conference on Acoustics, Speech, and Signal Processing (ICASSP) 2025*, 2025.

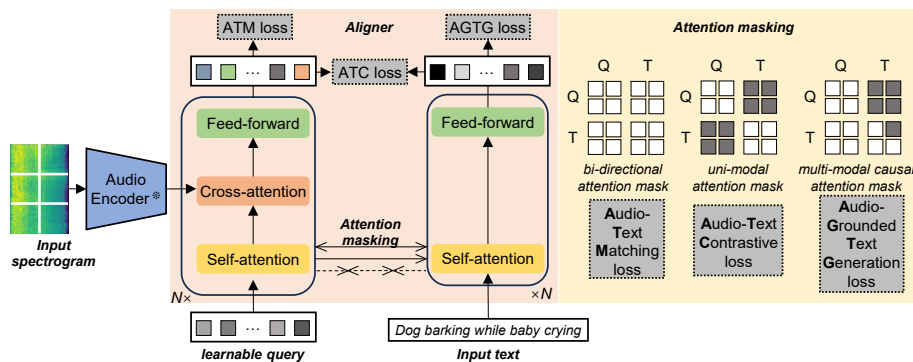


Fig. 5. Illustration of APT audio aligner pretrained with audio-text pairs. The rest components are frozen in the pretraining. Using different self-attention masks, the parameter of the audio aligner is optimised with triplet learning objectives: Audio-Text Matching (ATM), Audio-Grounded Text Generation (AGTG), and Audio-Text Contrastive (ATC).

TABLE XIV

TEMPLATE OF AUDIO TAGGING, AUTOMATED AUDIO CAPTIONING, LANGUAGE-QUERIED SOUND EVENT DETECTION, TEMPORAL EVENT RETRIEVAL, AND SOUND EVENT COUNTING.

Audio Tagging	Audio Captioning	Language-Queried Sound Event Detection	Temporal Event Retrieval	Sound Event Counting
Summarize the audio with key words.	Summarize the audio succinctly.	Pinpoint the presence of {LABEL} with the time stamps.	Summarize the audio with key words in the interval of {STT} seconds to {EDT} seconds.	How many times can the sound {LABEL} be heard?
What sound events can be heard in the audio clip?	Present a short overview of the provided audio samples.	Indicate the start and end time of the audio event {LABEL}.	What sound events can be heard from {STT} seconds to {EDT} seconds?	How many instances of the sound {LABEL} can be perceived?
What auditory incidents can be recognized in the recording?	Provide a compact summary of the auditory content.	Document the exact times the sound {LABEL} taking place.	What auditory incidents can be recognized in the recording from {STT} seconds to {EDT} seconds?	What is the number of times the sound {LABEL} is detectable?
Which auditory occurrences can be detected?	Offer a brief outline of the audio clips that have been given.	Specify the time stamps for {LABEL} occurrence.	Which auditory occurrences can be detected during {STT} seconds to {EDT} seconds?	How frequently can one hear the sound {LABEL}?
Which sound occurrences can be perceived?	Render a compressed version of the audio's main points.	When the sound {LABEL} happens?	Which sound occurrences can be perceived between {STT} seconds and {EDT} seconds?	How often can the sound {LABEL} be perceived?
Present a concise breakdown of the given audio clips.	Describe the audio clip concisely.	Capture the exact times when {LABEL} is happening.	Present a concise breakdown of the recording from {STT} seconds to {EDT} seconds.	
List the sound events in the audio clip.	Explain the audio clip in a brief and straightforward manner.	Describe the time intervals during which {LABEL} takes place.	List the sound events in the interval of {STT} seconds to {EDT} seconds.	
Describe the recording with names of sound events.	Write a terse but informative summary of the sound.	State the precise moment at which {LABEL} occurs.	Name the auditory incidents within the {STT} to {EDT} seconds timeframe.	
Enumerate the audio events present in the audio.	Give a quick overview of the provided audio excerpts.	What time does the sound event {LABEL} take place?	Enumerate the audio events present between {STT} seconds and {EDT} seconds.	
Name the auditory incidents in the audio sample.	Outline the given audio samples briefly.	Capture the beginning and end time of the sound {LABEL}.	Describe the recording with names of sound events within the {STT} to {EDT} seconds timeframe.	
#Output: {LABEL}	#Output: {CAPTION}	#Output: {STT}s- {EDT}s	#Output: {LABEL}	#Output: {NUMBER}




Fig. 6. APT-LLM performance in few-shot audio classification when: (a) the number of classes increases, and (b) the number of examples per class increases. In both cases, the increase in the number of classes/examples leads to the rapid increase of the input sequence's length.

APPENDIX A

AUDIO-TEXT ALIGNMENT

Before coupling with an LLM, we pretrain the APT audio aligner to bridge the audio modality and the text modality. Fig. 5 shows the implementation of audio-text alignment pretraining. Following Q-former pretraining in the previous work [1], we freeze the other components and optimise the parameters of the audio aligner. During training, a set of acoustic embeddings are learnt to extract relevant information from the audio feature maps using the input text tokens. In the pretraining, the audio aligner learns with triplet training objectives:

Audio-text matching (ATM). The goal of ATM is to learn a token-level alignment between acoustic and text tokens. APT

is trained to distinguish whether a pair of an audio recording and a text sentence are from the same source. A bi-directional self-attention mask is applied to the input sequence such that acoustic and textual tokens can attend to each other. APT’s output embeddings Z thus carry both audio and text information. Each vector of its embeddings Z is then fed into a binary classifier, and the prediction is produced by averaging the binary logits over all vectors.

Audio-grounded text generation (AGTG). AGTG shapes APT’s learning to generate texts, given the input audio tokens. The information for generating the text is first extracted from the audio tokens, and then passed to the text tokens via self-attention layers. In APT, a multimodal causal self-attention mask is used to control audio-text interaction, where each text token can attend to all audio tokens and its previous text tokens.

Audio-text contrasting (ATC). The ATC enforces alignment between the audio representation and the text representation such that their mutual information is maximized. It achieves so by contrasting the audio-text similarity of a positive pair against those of negative pairs. APT aligns acoustic embeddings output by the audio encoder with text tokens. To avoid information leakage, we employ a unimodal self-attention mask, where the acoustic and text tokens are not allowed to see each other.

APPENDIX B


PROMPTS FOR VARIOUS TASKS

TABLE XIV detail the prompts we used to prompt APT-LLMs during training and inference. For each task, we hand-crafted three prompts as seeds and used ChatGPT to generate the rest of prompts. We checked the generated prompts and rejected those of ambiguity. We iterated the process till we got ten prompts per task.

APPENDIX C

ANALYSIS ON FEW-SHOT AUDIO CLASSIFICATION

Figure 6 (a) and (b) show APT-LLM performance in various few-shot settings. In both settings, the input sequence becomes longer when the number of shots/classes increases. It is interesting to observe that as the number of shots increases, the overall performance becomes worse. This is in contrast to the behavior of few-shot learners based on back propagation [48]. One possible reason is that increasing the number of classes/shots rapidly increases the length of the input sequence, which is detrimental to APT-LLM global attention mechanism.



Jinhua Liang received the B.Sc. degree in 2018 and M.E degree in 2021, both from Tianjin University, China. He is a Ph.D. student at Centre for Digital Music, Queen Mary University of London. His current research interests focus on multimodal learning for sound, involving audio understanding and generation using language and visual cues. He co-organise the challenge of the “Few-shot Bioacoustic Event Detectio” on DCASE 2024.



Xubo Liu received the B.Eng. degree from Queen Mary University of London, UK in 2020. He is a Ph.D. student at the Centre for Vision, Speech, and Signal Processing (CVSSP) at the University of Surrey, UK. His PhD is working on multimodal learning for audio and language, focusing on the understanding, separation, and generation of audio signals in tandem with natural language. He organized the “Multimodal Learning for Audio and Language” special session at EUSIPCO 2023 and the “Generative AI for Media Generation” special session at MLSP 2024. He led the organization of the “Language-Queried Audio Source Separation” challenge on DCASE 2024.



Wenwu Wang (M’02–SM’11) received the B.Sc. degree in 1997, the M.E. degree in 2000, and the Ph.D. degree in 2002, all from Harbin Engineering University, China. He is a Professor of Signal Processing and Machine Learning within the Centre for Vision Speech and Signal Processing and an AI Fellow within the Surrey Institute for People Centred Artificial Intelligence, University of Surrey. His current research interests include signal processing, machine learning and perception, artificial intelligence, machine audition (listening), and statistical anomaly detection. He has (co)-authored over 300 papers in these areas. The works he co-authored have received over 15 awards including the 2022 IEEE SPS Young Author Best Paper Award, ICAUS 2021 Best Paper Award, DCASE 2020, 2023 and 2024 Judges’ Award, DCASE 2019 and 2020 Reproducible System Award, and LVA/ICA 2018 Best Student Paper Award.



Mark D. Plumbley (S’88–M’90–SM’12–F’15) received the B.A.(Hons.) degree in electrical sciences and the Ph.D. degree in neural networks from University of Cambridge, Cambridge, U.K., in 1984 and 1991, respectively. He is Professor of Signal Processing at the Centre for Vision, Speech, and Signal Processing at the University of Surrey, Guildford, U.K. His current research concerns AI, machine learning and signal processing for analysis, recognition and generation of sound. He led the first international data challenge on Detection and Classification of Acoustic Scenes and Events, and currently holds an Engineering and Physical Sciences Research Council Fellowship on “AI for Sound”. He is a Member of the IEEE Signal Processing Society Technical Committee on Audio and Acoustic Signal Processing, and a Fellow of the IET and IEEE.



Huy Phan (Senior Member, IEEE) received the Dr.-Ing. degree (summa cum laude) in computer science from the University of Lübeck, Lübeck, Germany, in 2017. From 2017 to 2018, he was a Postdoctoral Research Assistant with the University of Oxford, Oxford, U.K. From 2019 to 2020, he was a Lecturer with the University of Kent, Canterbury, U.K. From 2020 to 2022, he was a Lecturer of Artificial Intelligence with Queen Mary University of London, London, U.K., and a Turing Fellow with The Alan Turing Institute, London. In 2023, he joined Amazon, Cambridge, MA, USA, where he was a Senior Research Scientist. In 2025, he joined Meta, where he is currently a Research Scientist. His research interests include machine learning and signal processing, with a focus on audio and biosignal analysis. Dr. Phan was the recipient of the Bernd Fischer Award for the best Ph.D. thesis from the University of Lübeck in 2018 and the Benelux’s IEEE-EMBS Best Paper Award 2019–2020.



Emmanouil Benetos (Senior Member, IEEE) received the Ph.D. degree in electronic engineering from the Queen Mary University of London (QMUL), London, UK, in 2012. He is currently Reader in Machine Listening with the School of Electronic Engineering and Computer Science, QMUL. He has authored or coauthored more than 200 peer-reviewed papers spanning several topics in audio and music signal processing, with applications to sound scene analysis and music information retrieval.