

# Extending Audio Context for Long-Form Understanding in Large Audio-Language Models

Yuatyong Chaichana<sup>1,2\*</sup>, Pittawat Taveekitworachai<sup>1†</sup>, Warit Sirichotedumrong<sup>1†</sup>,  
Potsawee Manakul<sup>1†</sup>, Kunat Pipatanakul<sup>1†</sup>

<sup>1</sup>SCB 10X, SCBX Group, <sup>2</sup>Chulalongkorn University

Correspondence: 6531337321@student.chula.ac.th

## Abstract

Large Audio-Language Models (LALMs) are often constrained by short audio context windows, even when their text backbones support long contexts, limiting long-form audio understanding. Prior work has introduced context-extension methods (e.g. YaRN) on unimodal LLMs, yet their application to LALMs remains unexplored. First, building on RoPE-based context extension, we introduce **Partial YaRN**, a training-free, modality-decoupled extension method that modifies only audio token positions, leaving text positions intact to preserve the base LLM’s text capabilities. Second, we propose Virtual Longform Audio Training (**VLAT**), a training strategy that extends Partial YaRN into a training-time positional augmentation. VLAT simulates diverse audio lengths during training, enabling generalization to inputs far longer than those seen in training. Our experiments on SALMONN and Qwen2-Audio confirm that Partial YaRN outperforms the original models across wide range of settings, and VLAT provides substantial performance improvement on long audio of unseen lengths.<sup>1</sup>

## 1 Introduction

The success of Large Language Models (LLMs) has spurred multimodal extensions, notably large audio-language models (LALMs)<sup>2</sup> that pair an audio encoder with a text backbone. By aligning audio and text into a shared representation space, LALMs can leverage the base LLM’s knowledge for complex audio understanding. However, practical use is constrained by short audio context windows: models such as SALMONN (Tang et al., 2024) and Qwen2-Audio (Chu et al., 2024) are typically trained on audio segments of 30s or less, and

thus generalize poorly to longer inputs (Guo et al., 2025).

To address this limitation, we study the application of LLM context extension methods such as Positional Interpolation (PI) (Chen et al., 2023) and YaRN (Peng et al., 2024) for audio context extension in LALMs. To our knowledge, the application of these techniques for this specific, modality-bound use case has not been systematically explored. Applying such *whole-context* techniques implicitly extends the audio context window as a side effect. However, as this straightforward approach alters the positional information of the entire sequence, including text tokens; it risks degrading the sophisticated language capability of the base LLM which was pretrained solely in text.

This potential drawback motivates a more targeted strategy. We introduce **Partial YaRN**, a *modality-decoupled*, audio-only context extension method that modifies only the audio tokens’ positional encodings. This design enables a direct comparative study against the whole-context approaches, allowing us to investigate the tradeoff between preserving position encodings of the text modality and maintaining a globally uniform positional space.

Our experiments on both training-free and fine-tuned audio context extensions reveal that extension-based methods outperform original models. However, the best choice between whole-context and modality-decoupled extensions is often model-dependent without an universally superior solution. This suggests that to build truly robust long-audio-capable models, we must address the core problem of generalization during the training process itself. Therefore, we extend Partial YaRN into a novel fine-tuning technique called *Virtual Longform Audio Training (VLAT)*. Acting as a positional augmentation technique, it simulates a wide range of audio lengths during fine-tuning, teaching the model to generalize beyond the lengths

\*Work done while interning at SCB 10X.

†Equal advising.

<sup>1</sup>Code and dataset are available at: <https://github.com/yophis/partial-yarn>.

<sup>2</sup>Also known by other names such as Speech-aware LLMs.

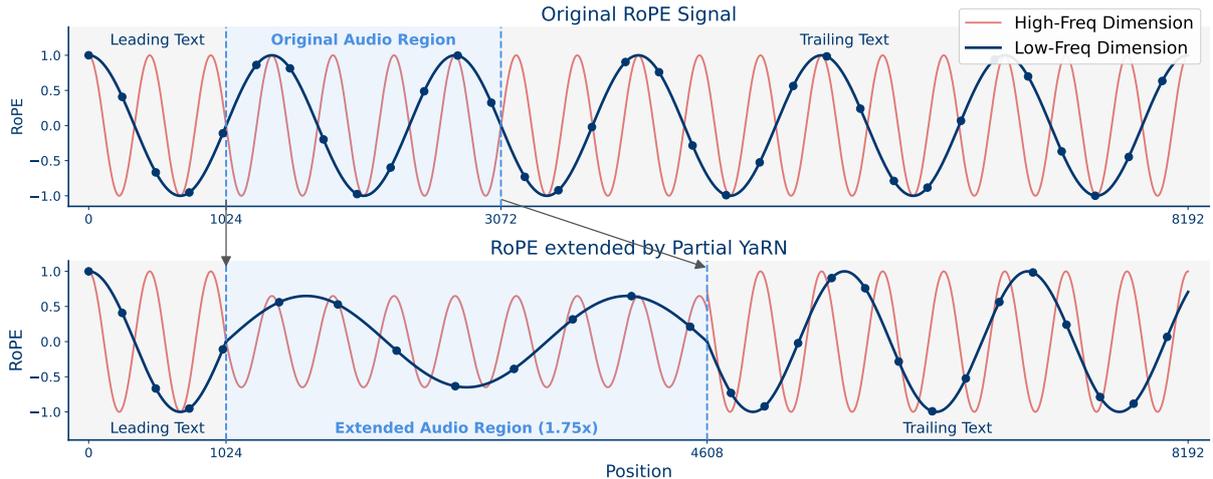


Figure 1: **An illustration of Partial YaRN.** The top plot displays low (blue) and high (red) frequency dimensions of RoPE. The bottom plot displays RoPE with Partial YaRN applied. In this example, the audio window of the low frequency dimensions are interpolated (stretched) by a factor of 1.75x to accommodate longer audio inputs, while the high frequency dimensions remain unaltered to preserve local distances and high frequency information (Peng et al., 2024). Within the *Extended Audio Region*, the magnitudes of every dimension is scaled to serve as a proxy for attention temperature scaling. Despite that the modality-decoupled scaling can create a minor discontinuity in the signal, information from other dimensions can help the model disambiguate this.

present in the training dataset. Through VLAT, we obtain substantially better generalization to long-form audio of unseen lengths.

Our main contributions are as follows:

- We propose *Partial YaRN*, a training-free audio-context extension for LALMs that preserves the base language model’s text capabilities by leaving text positions unaltered.
- We propose *Virtual Longform Audio Training (VLAT)*, a training strategy that extends Partial YaRN into a positional augmentation to simulate diverse audio lengths during training, for improved long-context generalization.
- We present a comparative study between existing context extension methods designed for unimodal LLMs and our proposed methods for LALMs, analyzing their performance tradeoff in both training-free and fine-tuning settings.

## 2 Preliminary

### 2.1 Rotary Positional Encoding

As Transformer (Vaswani et al., 2017) utilizes attention mechanism (Bahdanau et al., 2014), which is position-invariant, explicit positional information has to be provided for the models to differentiate tokens from different positions. Our study focuses on LALMs that use Rotary Positional Encoding (RoPE) (Su et al., 2024), a widely adopted positional encoding.

RoPE formulates relative positional information through the rotation of query and key vectors. Namely, given a query  $\mathbf{q}_m$  at position  $m$ , and key  $\mathbf{k}_n$  at position  $n$ , we have:

$$\begin{aligned} \mathbf{q}_m^T \mathbf{k}_n &= (\mathbf{R}_{m, \theta} \mathbf{q}_m)^T (\mathbf{R}_{n, \theta} \mathbf{k}_n) \\ &= \mathbf{q}_m^T \mathbf{R}_{n-m, \theta} \mathbf{k}_n, \end{aligned}$$

where  $\theta$  is the set of angular frequencies (step size) of the rotation.

### 2.2 Context Length Extension

Transformers are usually unable to generalize well past the context window length they have seen during the training (e.g. 2048 for LLaMA (Touvron et al., 2023), 8192 for Mistral (Jiang et al., 2023)). To permit longer context length, additional fine-tuning is often required.

**Position interpolation (PI).** Chen et al. (2023) proposes that instead of extrapolating RoPE further to the region unfamiliar for the trained model, interpolation is done inside the pretrained context window. Specifically, let  $L$  be the original context length of the model, and  $L'$  be the target length to be extended to. Define  $s = \frac{L'}{L}$  as the extension factor. PI simply downscales the base rotational frequency by the extension factor:  $\frac{\theta_i}{s}$ . For example, if originally each positional step is rotated by  $\theta_i = 10^\circ$ ; to double the context length, the frequency (step size) can simply be halved to  $5^\circ$ . This

way, the model will never have to attend to positions outside of its trained window, resulting in a more stable approach.

**YaRN.** Peng et al. (2024) identifies that interpolating every dimension equally leads to the loss of local distances and high frequency information. They propose to spread the interpolation pressure by partitioning RoPE dimensions into three frequency-based groups: (1) low frequency dimensions are interpolated, (2) high frequency dimensions are solely extrapolated without interpolation to preserve high frequency information, (3) dimensions in-between get a mix of both interpolation and extrapolation. Additionally, YaRN applies temperature scaling to the logits of attention softmax:

$$\text{AS}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax} \left( \frac{\mathbf{Q}^T \mathbf{K}}{t\sqrt{d_k}} \right) \mathbf{V},$$

where  $t$  is the attention temperature.

**Other methods.** Ding et al. (2024) proposes LongRoPE to account for non-uniformity across RoPE dimensions through a novel evolutionary search algorithm to determine dimension-wise interpolation factors. This method is further improved in LongRoPE2 (Shang et al., 2025).

### 2.3 Large Audio-language Models and Their Audio Contexts

One broadly adopted paradigm of large audio-language models (LALMs) is the *unified input space* (Tang et al., 2024; Chu et al., 2024; Ding et al., 2025; Ghosh et al., 2025a) which involves transforming audio inputs into the textual input space of a base LLM. By sharing the embedding space, the strong text capability and learned knowledge of the base LLM can be leveraged while simultaneously being augmented with audio understanding ability. There also exists other LALM paradigms such as *cross attention* (Kong et al., 2024; Ghosh et al., 2025b) where the audio and text modalities are fused through cross-attention modules. This work focuses on models under the unified-input-space.

As with text context window in LLMs, LALMs also do not generalize well to audio with lengths much longer than the window seen during training (e.g. 7s for Pengi (Deshmukh et al., 2023); 30s for Qwen2-Audio (Chu et al., 2024)).

## 3 Partial YaRN

To extend existing unified input LALMs to longer audio, we hypothesize that “LALMs already possess a sufficient general understanding of audio and text, the main bottleneck is instead their unfamiliarity to audio positions beyond the range seen during audio-text training.”<sup>3</sup> This hypothesis suggests a training-free solution: manipulating the backbone LLM’s positional encoding. The goal is to remap positions of a long audio input into the model’s familiar audio range.

First, we study the application of YaRN (Peng et al., 2024) to LALMs, as an example of whole-context extension methods. Since the audio window is a part of the total context window, performing whole-context extension extends the audio window as a byproduct. However, this alters positional information for all tokens, including text, risking the degradation of the base LLM, which was pretrained solely in text. Motivated by the aforementioned hypothesis, we conceive *modality-decoupled extension*, and propose **Partial YaRN**, an audio-only extension method designed to exclusively stretch the audio window. By leaving the text’s positional encodings unaltered, this approach aims to extend<sup>4</sup> audio context while preserving the base model’s text capability.

### 3.1 Methodology

Inspired by PI (Chen et al., 2023) and YaRN (Peng et al., 2024) for whole-context extension in unimodal LLMs, we adapt the interpolation technique to LALMs in modality-decoupled manner, applying it exclusively to the audio region of the base language models.

Concretely, assuming a single audio input, let  $L_{\text{audio}}$  be the original audio context length,  $L'_{\text{audio}}$  be the target length, and  $p$  be the position of the first audio token. We define the positional range  $[p, p + L_{\text{audio}})$  as the *original audio context window*. It can be either a predefined location for audio input, or a dynamic region enclosed by special tokens (e.g. <speech> and </speech>). We then apply our modified YaRN technique to exclusively

<sup>3</sup>This differs from whole-context extension in unimodal LLMs, where the total sequence length exceeds the pre-trained context window, causing an out-of-distribution problem. Here, the total sequence (audio + text) usually still fits within the original window, making the core challenge one of adapting to an unfamiliar audio length and positions, rather than extrapolating to completely unseen positions.

<sup>4</sup>We use the terms *extend* and *stretch* interchangeably.

stretch this region to the new length  $L'_{\text{audio}}$ . This creates a partially stretched positional encoding:

$$\underbrace{[0, p]}_{\text{Text (unaltered)}} \oplus \underbrace{[p, p+L_{\text{audio}}]}_{\text{Audio (stretched)}} \oplus \underbrace{[p+L_{\text{audio}}, L]}_{\text{Text (unaltered)}},$$

where  $\oplus$  denotes concatenation in the temporal axis. Note that while the final position index is denoted by  $L$ , the total token capacity now increases to  $L + L'_{\text{audio}} - L_{\text{audio}}$  to account for the stretched audio window with higher token frequency.

Furthermore, we only partition the RoPE dimensions into two frequency-based groups (instead of three as in the original YaRN formulation): a low frequency group that undergoes pure interpolation, and a high frequency group that undergoes pure extrapolation. The rationale for this design is twofold:

**1) Consistency:** A two-group partition ensures a consistent and uniform positional encoding across the entire audio stream. YaRN’s original “in-between” group receives a mix of interpolation and extrapolation, which would cause these RoPE frequencies to not be extended to fit the entire audio. Our approach avoids this potential distortion.

**2) Tuning Efficiency:** This simplification reduces the tuning space of the *cutoff* index from two coupled variables to a single parameter.

We provide empirical validation and a more in-depth discussion of the two-group partitioning in Section 7.2 and Appendix A respectively. See Figure 1 for an example depiction of Partial YaRN.

### 3.2 Hyperparameters

Partial YaRN employs two hyperparameters:

**1) Cutoff Dimension Index:** This defines the boundary separating the RoPE dimensions into two groups. Dimensions equal or below this cutoff (low-frequency) are interpolated to stably cover longer audio context. Dimensions above it (high-frequency) are extrapolated to preserve local positional distances and high-frequency information. We default this to 0 (interpolate every dimension).

**2) Attention Temperature:** This controls the sharpness of attention distribution within the audio context window. A higher temperature softens the distribution, preventing attention scores from collapsing to a few tokens over long sequences. We default this to 1.0 (no temperature scaling).

Instead of directly adding a new temperature term to the attention softmax, we follow Peng et al. (2024) and integrate the attention temperature into magnitudes of RoPE’s signals. Particularly, define

$m$  and  $n$  as two positions within the extended audio context window, we have the attention softmax:

$$\text{AS}(\mathbf{q}_m, \mathbf{k}_n) = \text{softmax} \left( \frac{(\mathbf{R}_m \mathbf{q}_m)^T (\mathbf{R}_n \mathbf{k}_n)}{t \sqrt{d_k}} \right).$$

Then, we invert  $t$  into the two rotation matrices by scaling them with  $1/\sqrt{t}$  each:

$$\text{AS}(\mathbf{q}_m, \mathbf{k}_n) = \text{softmax} \left( \frac{\left( \frac{\mathbf{R}_m}{\sqrt{t}} \mathbf{q}_m \right)^T \left( \frac{\mathbf{R}_n}{\sqrt{t}} \mathbf{k}_n \right)}{\sqrt{d_k}} \right).$$

This reparameterization also handles attention between the unaltered text regions and the scaled audio regions. For instance, when a text query (unscaled) attends to an audio key (scaled by  $1/\sqrt{t}$ ), the resulting logit naturally gets scaled by  $1/\sqrt{t}$ , creating a smoother temperature transition. Overall, we have temperature-scaled rotation matrices:

$$\tilde{\mathbf{R}}_m = \begin{cases} \mathbf{R}_{\theta, m} & \text{if } m \notin [p, p + L_{\text{audio}}) \\ \frac{1}{\sqrt{t}} \mathbf{R}_{\theta/s, m} & \text{if } m \in [p, p + L_{\text{audio}}) \end{cases}$$

Notice that under the default hyperparameter values, the audio-only extension reduces from YaRN to PI, we therefore specifically name Partial YaRN with default hyperparameters: *Partial PI*.

## 4 Experimental Setup

### 4.1 Data

Due to the limited availability of length-tailored long audio datasets, we construct custom datasets with specific audio lengths.

**YODAS2-MCQA** We source audio from the English subset of YODAS2 (Li et al., 2023). We segment audio samples into non-overlapping segments of 1, 2, 5, and 10 minutes. For each segment, we generate five multiple-choice question-answering (MCQA) pairs using Gemini 2.0 Flash (Team et al., 2023), with four choices each. We design the generation prompt to ensure that each question focuses on different portions of the audio, and that they collectively cover the entire audio segment. The test set of each audio duration (1, 2, 5, and 10 minutes) has 750 QA pairs.

Additional details about dataset construction are provided in Appendix D.3.

### 4.2 Models

We select two widely-used open-weights LALMs: SALMONN (Tang et al., 2024) and Qwen2-Audio

Method	YODAS2-MCQA			
	1 min	2 mins	5 mins	10 mins
<b>GPT 4o</b>	89.10	92.91	90.40	83.65
<b>Gemini 2.0 Flash</b>	94.70	95.86	92.80	89.87
<b>SALMONN</b>				
Vanilla	49.01	46.13	32.93	23.47
<i>Stretching from 30s (original context)</i>				
Whole PI	39.87 (-9.14)	5.73 (-40.40)	0.13 (-32.80)	23.20 (-0.27)
Partial PI	50.07 (+1.06)	45.60 (-0.53)	12.40 (-20.53)	1.20 (-22.27)
Whole YaRN	54.70 (+5.69)	<b>59.87 (+13.74)</b>	<b>47.20 (+14.27)</b>	30.13 (+6.66)
Partial YaRN	<b>57.35 (+8.34)</b>	<u>59.60 (+13.47)</u>	38.53 (+5.60)	32.93 (+9.46)
<i>Stretching from 2 mins (observed context)</i>				
Whole PI	49.01	46.13	13.07 (-19.86)	23.60 (+0.13)
Partial PI	49.01	46.13	28.27 (-4.66)	4.13 (-19.34)
Whole YaRN	49.01	46.13	<u>45.87 (+12.94)</u>	<b>38.27 (+14.80)</b>
Partial YaRN	49.01	46.13	41.47 (+8.54)	<u>33.45 (+9.98)</u>
<b>Qwen2-Audio 7B Instruct</b>				
Vanilla	72.71	<b>75.87</b>	55.33	22.00
<i>Stretching from 30s (original context)</i>				
Whole PI	64.50 (-8.21)	7.73 (-68.14)	5.87 (-49.46)	8.67 (-13.33)
Partial PI	72.45 (-0.26)	67.73 (-8.14)	40.13 (-15.20)	18.40 (-3.60)
Whole YaRN	<b>73.11 (+0.40)</b>	<u>74.27 (-1.60)</u>	50.93 (-4.40)	31.07 (+9.07)
Partial YaRN	<u>72.72 (+0.01)</u>	73.60 (-2.27)	48.53 (-6.80)	28.53 (+6.53)
<i>Stretching from 2 mins (observed context)</i>				
Whole PI	72.71	75.87	46.93 (-8.40)	21.20 (-0.80)
Partial PI	72.71	75.87	59.60 (+4.27)	45.73 (+23.73)
Whole YaRN	72.71	75.87	<u>60.27 (+4.94)</u>	<u>47.60 (+25.60)</u>
Partial YaRN	72.71	75.87	<b>60.40 (+5.07)</b>	<b>48.00 (+26.00)</b>

Table 1: **Training-free audio context extension performance (Accuracy)**. We compare extension methods under two strategies: stretching from the original 30s context and the empirically observed 2mins context. We see that stretching from the observed 2mins context proves significantly more effective. Comparing between the two top approaches: Whole and Partial YaRN, we find their performance to be closed, suggesting that preserving the positional information of text tokens might not be of utmost importance, at least in this MCQA setting of ours.

(Chu et al., 2024). Both models utilize Whisper-encoder (Radford et al., 2023) as audio encoder, which is limited to processing audio of up to 30s. We handle longer audio inputs by segmenting them into non-overlapping 30s chunks, where each chunk is encoded independently.

The number of audio representations differs considerably between these models. For each 30s chunk on 50Hz Whisper-encoder, SALMONN generates a sequence of 88 audio tokens (with Q-Former), whereas Qwen2-Audio produces a sequence of 750 tokens (with a 2x downsampling).

### 4.3 Baseline Methods

We evaluate Partial PI (default hyperparameters) and Partial YaRN (tuned hyperparameters) against the following methods:

**1) Vanilla:** Unmodified base models. Long audio inputs are passed directly without any manipulation to their RoPE.

**2) Whole Position Interpolation (PI) (Chen et al., 2023):** Every RoPE dimension is uniformly

interpolated across the whole context window. This serves as the primitive whole-context baseline.

**3) Whole YaRN (Peng et al., 2024):** Frequency-based RoPE interpolation with attention scaling, applied to the whole context window. This serves as the primary whole-context baseline.

These three baselines require no tuning as Vanilla and Whole PI don’t utilize any hyperparameter, and YaRN has predefined cutoffs and a closed-form formula for temperature.<sup>5</sup>

**Reference models.** We also report performance of two proprietary models capable of long audio context: GPT-4o (OpenAI, 2024) and Gemini 2.0 Flash (Team et al., 2023), in order to provide a broader perspective of LALMs and to validate the quality of our dataset.<sup>6</sup>

<sup>5</sup>Unlike Whole YaRN for unimodal LLMs, where information per token is relatively consistent, compression of audio tokens varied significantly across different LALMs as mentioned. This likely influences the optimal hyperparameter choice considerably, as discussed in Appendix E.2.

<sup>6</sup>The model versions are gpt-4o-mini-2024-07-18 and gemini-2.0-flash-001 respectively.

Training Method	Vanilla Inference			Partial PI Inference		
	2 mins	5 mins	10 mins	2 mins	5 mins	10 mins
Vanilla Fine-tuning	97.60	89.07	32.76	96.40	89.78	75.56
Virtual Longform	97.60 (+0.00)	<b>91.42 (+2.35)</b>	<b>75.11 (+42.35)</b>	<b>96.87 (+0.47)</b>	<b>91.91 (+2.13)</b>	<b>81.73 (+6.17)</b>

Table 2: **Virtual Longform Audio Training (VLAT) enables strong generalization to unseen audio lengths (Accuracy).** We compare standard Vanilla fine-tuning and our proposed VLAT on Qwen2-Audio, evaluated with (*Partial PI Inference* column group) and without (*Vanilla Inference* column group) Partial PI during inference.

## 5 Main Results

In this section, we comparatively study the effect of whole-context and modality-decoupled context extensions on both the training-free and fine-tuning settings for audio segments up to 10mins.

### 5.1 Training-free Audio Context Extension

First, we extend the audio context length without training. We report the results in Table 1, showing that: no single method—whole-context or modality-decoupled—is universally superior. Instead, the performance varies between the different models and the degree of extension. More importantly, we observe that *both models generally retain consistent performance up to 2mins*, suggesting that their innate audio context lengths are longer than just 30s, possibly a result from multi-audio training.

**Extension of SALMONN.** When stretching from the 30s window, extension methods provide substantial performance gains compared to the Vanilla baseline. We observe that Partial YaRN shows the strongest performance on the 1min and 10mins settings, achieving 57.35% and 32.93% accuracies. On 2mins, both Whole and Partial YaRN perform similarly. However, Whole YaRN outperforms Partial YaRN at 5mins by a large margin. Possible reasons for this performance flip are (1) higher expressivity and less compression pressure from Whole YaRN’s additional “in-between” dimension group, outweighs the preservation of the base language capability in this 5mins setting, or (2) noise in the hyperparameter tuning process of Partial YaRN, which is an inherent drawback.

**Extension of Qwen2-Audio.** In contrast to SALMONN, Qwen2-Audio demonstrates different characteristics. The Vanilla baseline is more robust, outperforming all stretch-from-30s methods at 2mins. Both Whole and Partial YaRN struggle to provide significant gains and even underperform the baseline at moderate lengths (2 and 5mins)

when stretching from the original 30s context; however, they manage to attain considerable gains on the 10mins setting. This indicates that Qwen2-Audio likely possesses a strong audio-length generalization capability or longer innate audio context window. For this model, simply applying an extension method from the base context does not guarantee an improvement. This suggests that the model’s high intrinsic audio-length generalizability is superior to the extension methods that inevitably cram the positional information as a side effect.

**Extending from observed audio context.** As mentioned in the beginning of this section that we can observe both model’s innate audio contexts to be around 2mins instead of just 30s, here we experiment with extending from this *observed audio context of 2mins*, to the target lengths. For both SALMONN and Qwen2-Audio, performance on the long audio improves substantially when the interpolation is anchored from the 2mins context instead of the original 30s. Notably on Qwen2-Audio, this strategy boosts Partial YaRN’s performance at 10mins from 28.53% to a 48.00%, an absolute improvement of almost 20%, yielding a final performance 26% higher than the Vanilla baseline. This finding suggests that determining the extension by observing the vanilla performance may be a more critical factor than the choice between whole-context or modality-decoupled methods, especially on large extension ratios.

### 5.2 Fine-tuning for Audio Context Extension

Next, we study training-based extension methods. Keeping other components frozen, we fine-tune the base language model of Qwen2-Audio using a *single epoch* LoRA (Hu et al., 2022) of rank 8 on the YODAS2-MCQA training set of the corresponding length. Two LoRA settings are explored: adapting only the query projection (q), and adapting query, key, value, and output projections (qkvo). Due to the high cost of tuning Partial YaRN’s hyperparameters during training, we only use the default con-

figuration (Partial PI) for the audio-only method. Whole PI is omitted due to its poor performance as previously observed.

The results are presented in Table 3, showing that fine-tuning with a context extension method dramatically outperforms the vanilla fine-tuning baseline, especially at longer audio lengths. In the qkvo setting, both Whole YaRN (83.47%) and Partial PI (83.07%) achieve an absolute performance gain of around 19% over the Vanilla baseline (64.93%) at 10mins. This highlights the benefit of integrating a context extension method into the fine-tuning process for better context extension.

Comparing the whole-context and modality-decoupled approaches, we find that their performance is generally competitive. While Whole YaRN triumphs when only adapting the query projection weights, suggesting its ease of adaptation. Their performance converges under standard LoRA practice which adapts qkvo. This result confirms that both whole-context and audio-only methods are highly effective strategies for fine-tuning.

Adapt	Method	2 mins	5 mins	10 mins
q	Vanilla	95.07	78.93	21.60
	Whole YaRN	<b>96.53</b>	<b>88.53</b>	<b>73.73</b>
	Partial PI	96.13	83.73	69.20
qkvo	Vanilla	96.13	85.20	64.93
	Whole YaRN	97.47	<b>93.60</b>	<b>83.47</b>
	Partial PI	<b>98.00</b>	93.33	83.07

Table 3: **Fine-tuned audio context extension performance (Accuracy).** We fine-tune Qwen2-Audio on YODAS2-MCQA, under two LoRA settings: adapting only  $q$ , and adapting  $qkvo$ . Each model is fine-tuned and evaluated using the same extension method.

## 6 Virtual Longform Audio Training: Generalizing Beyond the Training Data

Our previous results show that while the training-free methods are beneficial, their effectiveness is often model and length dependent. Based on this, we turn to address the core generalization problem during the training process, and propose to fine-tune LALMs with Partial YaRN repurposed as a positional augmentation technique we call “*Virtual Longform Audio Training*” (VLAT).

### 6.1 Methodology

The goal of VLAT is to simulate and expose models to audio context windows of diverse lengths during training. Let  $L_{\text{data}}$  be the actual length of a training audio sample. For each sample, we obtain

a “virtual” source length,  $L_{\text{virt}}$ , as the model’s default audio context length *times* a factor randomly sampled from the range  $[1, 5, 10, 15, 20, 25]x$ . We then apply Partial YaRN<sup>7</sup> to stretch or compress the  $L_{\text{virt}}$ -long positional window to  $L_{\text{data}}$ .

For instance, consider an audio sample with  $L_{\text{data}} = 2\text{mins}$ , and we draw a  $L_{\text{virt}}$  of 10mins, Partial YaRN will compress a 10mins context window to 2mins for this audio. This process effectively simulates the 2mins audio to be 10mins long virtually, thereby familiarizing the model to longer audio context and improving its ability to generalize to genuinely long audio at inference time.

Randomizing positional encoding has previously been explored in Ruoss et al. (2023), where positional indices are randomly downsampled to let the model see larger positional values. Rather than using  $\{1, 2, \dots, N\}$  for an length- $N$  input, they use a randomly subsampled values such as  $\{1, 4, \dots, L\}$  with  $L > N$  instead. In contrast, VLAT differs in employing interpolation inside the context window to create a dense, *continuous* space of positions, unlike sparse integer subsampling. It is also bidirectional, teaching the model through both compressed and stretched contexts. Furthermore, VLAT is a targeted, modality-decoupled method that modifies only the audio tokens, preserving the base LLM’s original textual space.

### 6.2 Evaluation

We fine-tune Qwen2-Audio on the 2mins YODAS2-MCQA training set<sup>8</sup> using two methods: a standard Vanilla fine-tuning and VLAT. We employ 10-epoch qkvo LoRA for the adaptations. We then evaluate the models from both training methods under two different extensions at inference: Vanilla and Partial PI. See Appendix E.1 for additional detail on VLAT configuration.

### 6.3 Results

As shown in Table 2, VLAT greatly improves generalization to unseen audio lengths. When evaluated with vanilla inference, the model trained with Virtual Longform shows a dramatic improvement on 10mins audio, increasing accuracy from 32.76% to 75.11%, closing in on the previous 1-epoch direct-fine-tuning result. This highlights the effectiveness

<sup>7</sup>Whole YaRN can’t be used here, as we observe that it diverges quickly when training under the VLAT framework.

<sup>8</sup>To prevent data leakage, the train/test splits were performed at the video level, ensuring no source material overlaps between the training set and any of the test sets.

of the training technique for training LALMs to generalize far beyond the audio lengths present in their training data, mitigating a key bottleneck in the development of robust long-audio models.

Furthermore, as can be seen in "Partial PI Inference" columns, VLAT’s benefit is complementary to inference-time extension. When Partial PI is applied during evaluation, the VLAT-trained model again outperforms its vanilla-trained counterpart, achieving the highest overall 10mins performance of 81.73%. This indicates that VLAT and inference-time extension are compatible strategies, and their combination yields the most robust long-audio-context performance.

We show in Appendix E.1 that VLAT’s performance can be even more enhanced by using a narrower range for virtual length  $L_{\text{virt}}$ .

## 7 Analysis

### 7.1 Ablation Study

Here, we conduct an ablation study to isolate and analyze the individual contributions of the two key components in Whole YaRN (Peng et al., 2024) and Partial YaRN: (1) frequency grouping and (2) attention temperature scaling. We note that by removing both of these components, the methods converge to Whole PI (Chen et al., 2023) and Partial PI. The results are presented in Table 4.

Method	Partial		Whole	
	2 mins	10 mins	2 mins	10 mins
<b>SALMONN</b>				
YaRN	<b>59.60</b>	<b>32.93</b>	<b>59.87</b>	<b>30.13</b>
- Freq grouping	56.93	4.53	23.33	29.20
- Attn temp	48.13	23.87	41.07	23.20
PI	45.60	1.20	5.37	23.20
<b>Qwen2-Audio</b>				
YaRN	<b>73.60</b>	<b>28.53</b>	73.47	<b>25.33</b>
- Freq grouping	66.00	28.13	17.07	9.87
- Attn temp	73.47	18.80	<b>75.07</b>	22.80
PI	67.73	18.40	7.73	8.67

Table 4: **Ablation of frequency grouping and attention temperature components (Accuracy)**. Removing either the frequency grouping or the attention temperature leads to a large drop in performance.

We can see that removing either the frequency grouping or the attention temperature generally leads to a large drop in performance, especially on higher extension ratios. Both the frequency partitioning and attention temperature scaling are crucial for robustly extending the audio context length of LALMs. Their combined application within

the Whole and Partial YaRN frameworks generally yield the most effective and stable performance.

### 7.2 Partial YaRN with Three-group Frequency Partitioning

For the Partial YaRN method, we modify the original 3-group frequency partition from original YaRN to a 2-group method as discussed in Section 3.1 and Appendix A. Here we conduct a direct empirical comparison between the 2-group and 3-group approaches.

For the 3-group baseline, we implement Partial YaRN using YaRN’s original partitioning scheme and its recommended hyperparameters. We compare this against our proposed 2-group Partial YaRN using the hyperparameters from our main experiments (Section 5.1). We present the results in Table 5.

Model	Freq Grouping	YODAS2-MCQA		
		2 mins	5 mins	10 mins
SALMONN	2-Group	<b>59.60</b>	38.53	<b>32.93</b>
	3-Group	56.00	<b>41.87</b>	25.46
Qwen2-Audio	2-Group	<b>73.60</b>	<b>48.53</b>	<b>28.53</b>
	3-Group	72.00	20.00	16.53

Table 5: **Performance Comparison of Two-Group vs. Three-Group Partitioning (Accuracy)**. We compare our Partial YaRN with 2-group frequency partitioning used throughout the work, against a 3-group version.

We see that our 2-group partitioning broadly outperform 3-group in our modality-specific scenario. This is most evident with the Qwen2-Audio model, where the 3-group method suffers a large performance drop on 5mins and 10mins audio, with an accuracy gap of 28.53% and 12.00% respectively. Overall, this result validates the effectiveness and reliability of our proposed 2-group scheme.

## 8 Related Work

This section discusses additional related works that aren’t addressed in Section 2.

**Positional encodings.** Absolute positional encoding was originally proposed in Vaswani et al. (2017). It utilizes sinusoidal waves to uniquely represent absolute position of each token. Later, (Devin et al., 2019; Lan et al., 2020; Clark et al., 2020) use trainable vectors for representing absolute positions, allowing the models to learn the positional encoding for themselves. However, it cannot extrapolate to sequences longer than the maximum

length it was trained on. [Shaw et al. \(2018\)](#) encodes the relative distances between the tokens instead of absolute positions. [Dai et al. \(2019\)](#) introduces additional trainable position-agnostic components to distinguish between content-based and location-based attention behaviors. [Raffel et al. \(2020\)](#) groups the relative distances into buckets, and associate each bucket with a learnable scalar for adding to attention logits. TUPE ([Ke et al., 2021](#)) uses separate linear projections for the token and positional encodings. ALiBi ([Press et al., 2022](#)) simply adds linear biases with a learnable slope into the attention scores, imposing a strong locality preference. Several later works extend RoPE for multimodal modeling, such as Multimodal RoPE ([Bai et al., 2025](#)), VideoRoPE ([Wei et al., 2025](#)), and VRoPE ([Liu et al., 2025b](#)).

**Large video-language models.** Video input is similar to the audio in that they're both naturally continuous and can be of arbitrary lengths. Many Large Video-Language Models (LVLMs) ([Zhang et al., 2023b](#); [Maaz et al., 2024](#); [Lin et al., 2024](#); [Jin et al., 2024](#); [Ren et al., 2024](#)), adopt architectural paradigm similar to the unified input space LALMs we studied, where video frames are embedded and projected into a base LLM's input space. However, the context length problem is even more exacerbated in the video domain due to the high dimensionality and large number of tokens needed to represent even short videos. To improve efficiency, the LVLM community has largely focused on input compression techniques like query aggregation ([Zhang et al., 2023b](#); [Li et al., 2025a](#); [Song et al., 2024](#)), frame pooling ([Luo et al., 2023](#); [Maaz et al., 2024](#)), and feature merging ([Weng et al., 2024](#)). These strategies primarily aim to reduce the number of tokens before they enter the base LLM. Our work offers a complementary perspective: instead of altering the input, we manipulate the positional space within the model to accommodate a longer sequence, similar to [Zhang et al. \(2025\)](#) in LVLMs. The principles of Partial YaRN and VLAT are likely applicable to the video domain as well, hopefully providing a distinct path to achieve long-video-context understanding and generative control in synergy with existing methods.

**Long-audio LALMs.** Concurrent works have proposed LALMs that have inherently long audio-context. Audio Flamingo 3 ([Ghosh et al., 2025a](#)) is capable of processing audio up to 10mins due long-audio training. Voxtral ([Liu et al., 2025a](#)) can

handle up to 40mins due to the long 32k context of its LLM backbone. FastLongSpeech ([Guo et al., 2025](#)) proposes Iterative Fusion which incrementally merges the audio representation based on their information density. This compressed representation facilitates efficient processing of long audio. Qwen3-Omni ([Xu et al., 2025](#)) is a multimodal model that can accept text, video, and audio as input. It employs a custom Audio Transformer which is trained in an encoder-decoder fashion similar to the Whisper model. Qwen3-Omni is also capable of processing audio of up to 40mins.

While these models achieve impressive long-audio performance, they do so by designing new architectures, curating large volumes of long-audio data, and undertaking extensive training. Our work investigates an orthogonal yet complementary direction of efficiently extending the context of existing LALMs. Rather than requiring costly data acquisition and retraining, we provide a lightweight, architecture-agnostic "drop-in" enhancement (Partial YaRN) and an efficient fine-tuning strategy (VLAT), offering a practical and accessible approaches for improving the long-audio capabilities of a wide range of existing models.

## 9 Conclusion and Future Work

This work addressed the challenge of extending audio context windows of LALMs by studying the application of LLM's whole-context extension methods, and proposing a training-free audio-only extension Partial YaRN and the VLAT training strategy. We demonstrated the effectiveness of extension-based methods over non-extension baselines across various long-audio settings. Leveraging our finding that SALMONN and Qwen2-Audio retain consistent performance up to 2mins audio window, we achieved an even stronger performance by extending from the 2mins window instead of 30s. Later, we showed that fine-tuning through VLAT allows the models to generalize to audio contexts far exceeding the lengths of the training data, offering a robust and data-efficient pathway to develop LALMs with better longform understanding.

Future work could apply Partial YaRN and VLAT to vision models. Beyond content understanding, these approaches could also be explored for controllable generation in multimodal models, such as multi-scale image synthesis, or even more finegrained decoupling for non-uniform speed control throughout audio or video generation.

## Limitations

**Evaluation scope.** Our experiments are conducted exclusively on the multiple-choice question answering tasks. We made this choice to ensure a precise and unambiguous measurement of the model performance under different audio lengths. This mitigates the confounding variables inherent in evaluating open-ended generation, such as judging the semantic equivalence or stylistic alignment. However, a key tradeoff of this focused evaluation is its limited ability to assess the nuanced generative and language modeling capabilities of the base LLM. Consequently, while our results validate the model’s ability to retrieve information from audio, they do not fully test the importance of text preservation of the audio-only methods like Partial YaRN. Furthermore, our study does not assess the performance across other long-audio task formats found in benchmarks such as [Guo et al. \(2025\)](#), [Ahia et al. \(2025\)](#), or [Ghosh et al. \(2025a\)](#).

**Reliance on synthetic data.** Our study utilizes YODAS2-MCQA, a custom dataset synthesized from the YODAS2 corpus using Gemini 2.0 Flash. While we carefully steered the synthesis process and validated a subset of the data, it’s still possible for the synthesized samples to carry noise, repetition, and hallucinations that may affect the fidelity of our performance estimates compared to real-world scenarios.

**Hyperparameter tuning burden.** Unlike the original YaRN which has generally good predefined cutoff dimensions and a closed-form formula for determining the attention temperature, our proposed Partial YaRN requires tuning of two hyperparameters, thereby introducing additional computational burden.

## Acknowledgments

YC would like to thank Ekapol Chuangsuwanich for insightful discussion on the earlier version of this work. He also thanks Thanapat Trachu for valuable guidance on the illustration of Partial YaRN, and Konpat Preechakul for illuminating discussion on extensions of this work to the vision domain. We gratefully acknowledge the compute resources provided by Hatari NEXT, on which the experiments for this research were run.

## References

- Orevaoghene Ahia, Martijn Bartelds, Kabir Ahuja, Hila Gonen, Valentin Hofmann, Siddhant Arora, Shuyue Stella Li, Vishal Puttagunta, Mofetoluwa Adeyemi, Charishma Buchireddy, and 1 others. 2025. Blab: Brutally long audio bench. *arXiv preprint arXiv:2505.03054*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibong Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-v1 technical report. *arXiv preprint arXiv:2502.13923*.
- Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. 2023. Extending context window of large language models via positional interpolation. *arXiv preprint arXiv:2306.15595*.
- Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, and 1 others. 2024. Qwen2-audio technical report. *arXiv preprint arXiv:2407.10759*.
- Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and Jingren Zhou. 2023. Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models. *arXiv preprint arXiv:2311.07919*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#). In *International Conference on Learning Representations*.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. Transformer-xl: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988.
- Tri Dao. 2024. [Flashattention-2: Faster attention with better parallelism and work partitioning](#). In *The Twelfth International Conference on Learning Representations*.
- Alexandre Défossez, Laurent Mazaré, Manu Orsini, Amélie Royer, Patrick Pérez, Hervé Jégou, Edouard Grave, and Neil Zeghidour. 2024. Moshi: a speech-text foundation model for real-time dialogue. *arXiv preprint arXiv:2410.00037*.
- Soham Deshmukh, Benjamin Elizalde, Rita Singh, and Huaming Wang. 2023. Pengi: An audio language model for audio tasks. *Advances in Neural Information Processing Systems*, 36:18090–18108.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep

- bidirectional transformers for language understanding. 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)*, pages 4171–4186.
- Ding Ding, Zeqian Ju, Yichong Leng, Songxiang Liu, Tong Liu, Zeyu Shang, Kai Shen, Wei Song, Xu Tan, Heyi Tang, and 1 others. 2025. Kimi-audio technical report. *arXiv preprint arXiv:2504.18425*.
- Yiran Ding, Li Lina Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, and Mao Yang. 2024. Longrope: Extending llm context window beyond 2 million tokens. In *International Conference on Machine Learning*, pages 11091–11104. PMLR.
- Sreyan Ghosh, Arushi Goel, Jaehyeon Kim, Sonal Kumar, Zhifeng Kong, Sang gil Lee, Chao-Han Huck Yang, Ramani Duraiswami, Dinesh Manocha, Rafael Valle, and Bryan Catanzaro. 2025a. [Audio flamingo 3: Advancing audio intelligence with fully open large audio language models](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Sreyan Ghosh, Zhifeng Kong, Sonal Kumar, S Sakshi, Jaehyeon Kim, Wei Ping, Rafael Valle, Dinesh Manocha, and Bryan Catanzaro. 2025b. [Audio flamingo 2: An audio-language model with long-audio understanding and expert reasoning abilities](#). In *Forty-second International Conference on Machine Learning*.
- Sreyan Ghosh, Sonal Kumar, Ashish Seth, Chandra Kiran Reddy Evuru, Utkarsh Tyagi, S Sakshi, Oriol Nieto, Ramani Duraiswami, and Dinesh Manocha. 2024. [GAMA: A large audio-language model with advanced audio understanding and complex reasoning abilities](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6288–6313, Miami, Florida, USA. Association for Computational Linguistics.
- Yuan Gong, Alexander H Liu, Hongyin Luo, Leonid Karlinsky, and James Glass. 2023. Joint audio and speech understanding. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 1–8. IEEE.
- Yuan Gong, Hongyin Luo, Alexander H. Liu, Leonid Karlinsky, and James R. Glass. 2024. [Listen, think, and understand](#). In *The Twelfth International Conference on Learning Representations*.
- Shoutao Guo, Shaolei Zhang, Qingkai Fang, Zhenrui Ma, Min Zhang, and Yang Feng. 2025. Fastlongspeech: Enhancing large speech-language models for efficient long-speech processing. *arXiv preprint arXiv:2507.14815*.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Ailin Huang, Boyong Wu, Bruce Wang, Chao Yan, Chen Hu, Chengli Feng, Fei Tian, Feiyu Shen, Jingbei Li, Mingrui Chen, and 1 others. 2025. Step-audio: Unified understanding and generation in intelligent speech interaction. *arXiv preprint arXiv:2502.11946*.
- Rongjie Huang, Mingze Li, Dongchao Yang, Jia-tong Shi, Xuankai Chang, Zhenhui Ye, Yuning Wu, Zhiqing Hong, Jiawei Huang, Jinglin Liu, and 1 others. 2024. Audiogpt: Understanding and generating speech, music, sound, and talking head. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 23802–23804.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L el io Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth ee Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint, arXiv:2310.06825*.
- Peng Jin, Ryuichi Takanobu, Wancai Zhang, Xiaochun Cao, and Li Yuan. 2024. Chat-univi: Unified visual representation empowers large language models with image and video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13700–13710.
- Guolin Ke, Di He, and Tie-Yan Liu. 2021. [Rethinking positional encoding in language pre-training](#). In *International Conference on Learning Representations*.
- Zhifeng Kong, Arushi Goel, Rohan Badlani, Wei Ping, Rafael Valle, and Bryan Catanzaro. 2024. [Audio flamingo: A novel audio language model with few-shot learning and dialogue abilities](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 25125–25148. PMLR.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [Albert: A lite bert for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario  s aško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, and 13 others. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2025a. Videochat: Chat-centric video understanding. *Science China Information Sciences*, 68(10):200102.

- Tianpeng Li, Jun Liu, Tao Zhang, Yuanbo Fang, Da Pan, Mingrui Wang, Zheng Liang, Zehuan Li, Mingan Lin, Guosheng Dong, and 1 others. 2025b. Baichuan-audio: A unified framework for end-to-end speech interaction. *arXiv preprint arXiv:2502.17239*.
- Xinjian Li, Shinnosuke Takamichi, Takaaki Saeki, William Chen, Sayaka Shiota, and Shinji Watanabe. 2023. Yodas: Youtube-oriented dataset for audio and speech. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 1–8. IEEE.
- Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. 2024. Video-llava: Learning united visual representation by alignment before projection. In *Proceedings of the 2024 conference on empirical methods in natural language processing*, pages 5971–5984.
- Alexander H Liu, Andy Ehrenberg, Andy Lo, Clément Denoix, Corentin Barreau, Guillaume Lample, Jean-Malo Delignon, Khyathi Raghavi Chandu, Patrick von Platen, Pavankumar Reddy Muddireddy, and 1 others. 2025a. Voxtral. *arXiv preprint arXiv:2507.13264*.
- Zikang Liu, Longteng Guo, Yepeng Tang, Tongtian Yue, Junxian Cai, Kai Ma, Qingbin Liu, Xi Chen, and Jing Liu. 2025b. VROPE: Rotary position embedding for video large language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 14460–14472, Suzhou, China. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Ruipu Luo, Ziwang Zhao, Min Yang, Junwei Dong, Da Li, Pengcheng Lu, Tao Wang, Linmei Hu, Minghui Qiu, and Zhongyu Wei. 2023. Valley: Video assistant with large language model enhanced ability. *arXiv preprint arXiv:2306.07207*.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. 2024. Video-ChatGPT: Towards detailed video understanding via large vision and language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12585–12602, Bangkok, Thailand. Association for Computational Linguistics.
- Potsawee Manakul, Guangzhi Sun, Warit Sirichotedumrong, Kasima Tharnpipitchai, and Kunat Pipatanakul. 2025. Enhancing Low-Resource Language and Instruction Following Capabilities of Audio Language Models. In *Interspeech 2025*, pages 2083–2087.
- Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, and Benjamin Bossan. 2022. PEFT: State-of-the-art parameter-efficient fine-tuning methods. <https://github.com/huggingface/peft>.
- OpenAI. 2024. Gpt-4o system card. *Preprint*, arXiv:2410.21276.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks. In *International conference on machine learning*, pages 1310–1318. PMLR.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, and 1 others. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. 2024. YaRN: Efficient context window extension of large language models. In *The Twelfth International Conference on Learning Representations*.
- Ofir Press, Noah A Smith, and Mike Lewis. 2022. Train short, test long: Attention with linear biases enables input length extrapolation. In *ICLR*.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Shuhuai Ren, Linli Yao, Shicheng Li, Xu Sun, and Lu Hou. 2024. Timechat: A time-sensitive multi-modal large language model for long video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14313–14323.
- Paul K Rubenstein, Chulayuth Asawaroengchai, Duc Dung Nguyen, Ankur Bapna, Zalán Borsos, Félix de Chaumont Quitry, Peter Chen, Dalia El Badawy, Wei Han, Eugene Kharitonov, and 1 others. 2023. Audiopalm: A large language model that can speak and listen. *arXiv preprint arXiv:2306.12925*.
- Anian Ruoss, Grégoire Delétang, Tim Genewein, Jordi Grau-Moya, Róbert Csordás, Mehdi Bannani, Shane Legg, and Joel Veness. 2023. Randomized positional encodings boost length generalization of transformers. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1889–1903, Toronto, Canada. Association for Computational Linguistics.
- Ning Shang, Li Lina Zhang, Siyuan Wang, Gaokai Zhang, Gilsinia Lopez, Fan Yang, Weizhu Chen, and Mao Yang. 2025. LongroPE2: Near-lossless LLM context window scaling. In *Forty-second International Conference on Machine Learning*.

- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468.
- Enxin Song, Wenhao Chai, Guan hong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Haozhe Chi, Xun Guo, Tian Ye, Yanting Zhang, and 1 others. 2024. Moviechat: From dense token to sparse memory for long video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18221–18232.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun MA, and Chao Zhang. 2024. SALMONN: Towards generic hearing abilities for large language models. In *The Twelfth International Conference on Learning Representations*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Xilin Wei, Xiaoran Liu, Yuhang Zang, Xiaoyi Dong, Pan Zhang, Yuhang Cao, Jian Tong, Haodong Duan, Qipeng Guo, Jiaqi Wang, Xipeng Qiu, and Dahua Lin. 2025. VideoPE: What makes for good video rotary position embedding? In *Forty-second International Conference on Machine Learning*.
- Yuetian Weng, Mingfei Han, Haoyu He, Xiaojun Chang, and Bohan Zhuang. 2024. Longvlm: Efficient long video understanding via large language models. In *European Conference on Computer Vision*, pages 453–470. Springer.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Jian Wu, Yashesh Gaur, Zhuo Chen, Long Zhou, Yimeng Zhu, Tianrui Wang, Jinyu Li, Shujie Liu, Bo Ren, Linqun Liu, and 1 others. 2023. On decoder-only architecture for speech-to-text and large language model integration. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 1–8. IEEE.
- Jin Xu, Zhifang Guo, Hangrui Hu, Yunfei Chu, Xiong Wang, Jinzheng He, Yuxuan Wang, Xian Shi, Ting He, Xinfu Zhu, Yuanjun Lv, Yongqi Wang, Dake Guo, He Wang, Linhan Ma, Pei Zhang, Xinyu Zhang, Hongkun Hao, Zishan Guo, and 19 others. 2025. [Qwen3-omni technical report](#). *Preprint*, arXiv:2509.17765.
- Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. 2023a. [SpeechGPT: Empowering large language models with intrinsic cross-modal conversational abilities](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15757–15773, Singapore. Association for Computational Linguistics.
- Hang Zhang, Xin Li, and Lidong Bing. 2023b. [Video-LLaMA: An instruction-tuned audio-visual language model for video understanding](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 543–553, Singapore. Association for Computational Linguistics.
- Peiyuan Zhang, Kaichen Zhang, Bo Li, Guangtao Zeng, Jing kang Yang, Yuanhan Zhang, Ziyue Wang, Hao-ran Tan, Chunyuan Li, and Ziwei Liu. 2025. [Long context transfer from language to vision](#). *Transactions on Machine Learning Research*.

## A Elaboration on Two-group Partitioning of Partial YaRN

The decision to simplify the original three-group frequency partition of YaRN (Peng et al., 2024) into a two-group system for Partial YaRN is a deliberate design choice aimed at ensuring positional consistency across the entire extended audio stream. The original YaRN formulation partitions RoPE dimensions into: 1) a low-frequency group that is purely interpolated, 2) a high-frequency group that is purely extrapolated, and 3) an "in-between" group that receives a mix of both, with a ramped interpolation factor.

While effective for whole-context extension, this "in-between" group becomes problematic in our partial, audio-only extension scenario. Consider a 2mins audio input, which requires a 4x extension of the model's 30s original audio context.

- The high-frequency dimensions, when extrapolated (a 1x factor), correctly preserve local relative distances, which is desirable.
- The low-frequency dimensions, responsible for quasi-absolute positioning, are fully interpolated by a 4x factor. This is crucial as it maps the entire 2mins audio into the model's familiar positional range (original audio context).
- The "in-between" dimensions, however, would receive interpolation factors between the 1x and 4x (e.g. 2x). This means that for these specific dimensions, the model's original context window is stretched to cover only the first minute of the 2mins audio. The second minute would lie in an extrapolated, out-of-distribution region.

This dimensional inconsistency risks creating a biased or distorted representation of the latter half of the audio. The model would receive conflicting signals about whether a token is "inside" or "outside" its familiar context.

By simplifying to a two-group partition, we circumvent this issue entirely. Every dimension is forced into a binary choice: either it is fully interpolated to cover the entire audio duration (low-frequency), or it is fully extrapolated to preserve local structure (high-frequency). This ensures that the model maintains a consistent positional understanding across all dimensions for the entirety of

the audio input, avoiding the representational bias that partial coverage would introduce.

See Section 7.2 for direct performance comparison between 2 and 3-group frequency partitioning versions of Partial YaRN.

## B Implementation of Partial YaRN via position\_ids Manipulation

A key implementation detail of our work is how we apply the positional stretching to only a subsection of the input sequence. While the original YaRN implementation modifies RoPE's frequencies ( $\theta$ ), this approach isn't straightforward for a *partial* modification. Instead, our implementation of Partial YaRN directly manipulates the `position_ids` and multiply it with the unmodified frequencies. This section details the implementation our method.

**Modifying positional indices.** In most Transformer architectures, including the models used in this work, the RoPE mechanism computes embeddings based on an input tensor called `position_ids`. This tensor is simply a sequence of integers representing the absolute position of each token (i.e.  $[\theta, 1, 2, \dots, N-1]$ ). Our method leverages this by modifying a section this tensor. The process is as follows:

1. **Identify the Audio Region:** First, we identify the start and end indices of the audio tokens within the `position_ids` tensor. This creates three distinct segments: the leading text positions (unaltered), the audio positions (to be stretched), and the trailing text positions (unaltered relative to each other).
2. **Stretch the Audio Positions:** We replace the original sequence of integer positions for the audio segment with a new sequence generated via `torch.linspace`. For example, if the original audio positions were  $[p, p+1, \dots, p+\text{base\_L\_audio}-1]$  and the new audio input has `L_audio` tokens, we generate `L_audio` new positions by interpolating within the original range: `torch.linspace(start=p, end=p+L_audio-1, steps=L_audio)`. This effectively "stretches" the familiar positional space to accommodate the longer audio input.
3. **Concatenate:** The final `position_ids` tensor is constructed by concatenating the leading text positions, the new stretched audio positions, and the trailing text positions. This

new `position_ids` is now ready to be use for subsequent construction of partially stretch RoPE signal for the relevant dimensions.

**Computational overhead.** This implementation introduces negligible computational overhead in practice. The only additional computations are incurred during the construction of the `position_ids` tensor, which involves a few tensor slicing operations, a single `torch.linspace` call, and a final concatenation. These are highly optimized, vectorized operations that constitute a minuscule fraction of the total computation in a full forward pass. Crucially, this precomputation is a one-time cost per input sample, applied only during the initial processing of the prompt that contains the long audio. During the subsequent autoregressive generation steps, where tokens are generated one at a time, our method adds no overhead whatsoever. The model reverts to the standard process of incremental positional stepping, just as it would without any context extension. Therefore, the cost of enabling partial context extension is a small, fixed precomputation that does not impact the per-token generation latency.

## C On Deriving Closed-Form Hyperparameter Formulae

A practical extension of this work is to derive closed-form formulae for determining the optimal cutoff dimension index and attention temperature for Partial YaRN akin to the solution provided in YaRN for unimodal LLMs (Peng et al., 2024), which would eliminate the cost of hyperparameter searches. We leave this for future work due to a key challenge not present in unimodal LLMs: the wide variance in audio representation lengths across different LALMs.

Unlike the relatively consistent token-to-information ratio in text across modern tokenizers, LALMs compress audio into discrete tokens at vastly different rates. For instance, a 30-second audio clip gets compressed into 88 tokens in SALMONN but 750 tokens in Qwen2-Audio. As our sensitivity analysis in Appendix E.2 shows, this disparity strongly influences the optimal hyperparameter choice, making closed-form formula non-trivial to derive.

## D Additional Experimental Details

### D.1 Training and Inference

**Environment.** All fine-tunings and inferences are conducted using PyTorch (Paszke et al., 2019) and HuggingFace’s Transformers (Wolf et al., 2020) libraries on a single Nvidia H100 GPU. When available, we utilize Flash-Attention 2 (Dao, 2024) to speed up the processes and reduce VRAM consumptions. For experiments on very long sequence, we use gradient checkpointing to circumvent out-of-memory issue.

**Fine-tuning hyperparameters.** We use LoRA (Hu et al., 2022) of rank 8 for fine-tuning through HuggingFace’s PEFT library (Mangrulkar et al., 2022). With alpha set to 16, and a small LoRA dropout of 5%. Models are trained with AdamW optimizer (Loshchilov and Hutter, 2019), learning rate is set to 5e-5, a batch size of 8 samples, and gradient norm clipped to 1.0 (Pascanu et al., 2013).

### D.2 Partial YaRN’s Hyperparameter Tuning

Recall that Partial YaRN has two hyperparameters, cutoff dimension and attention temperature. For the training-free experiment (Section 5.1), we jointly sweep the cutoff dimension from {56, 48, 40, 32, 24, 16, 8}, and the temperature from between 0.5 and 1.6 with a step size of 0.1. See Appendix E.2 for tuning results and sensitivity.

### D.3 Dataset Construction

Our dataset is derived from a portion of the *en129*<sup>9</sup> subset of YODAS2 (Li et al., 2023), a large-scale, multilingual corpus of YouTube-derived audio and speech data. The YODAS2 dataset is made available under the CC-BY-3.0 license. We accessed and processed the data using HuggingFace’s *Datasets* library (Lhoest et al., 2021).

To prevent video source leakage in the VLAT experiment, train-test split was performed at the video level; all dialogue segments from a single video were exclusively assigned to either the training set (2mins), or the test sets.

To obtain specific audio length, we chunk the long audio in the dataset to 1, 2, 5, and 10 minutes. Any remaining, incomplete chunk of each long-audio is also retained. Subsequently, we employed Gemini 2.0 Flash (Team et al., 2023) to gen-

<sup>9</sup><https://huggingface.co/datasets/espnet/yodas2/tree/main/data/en129>

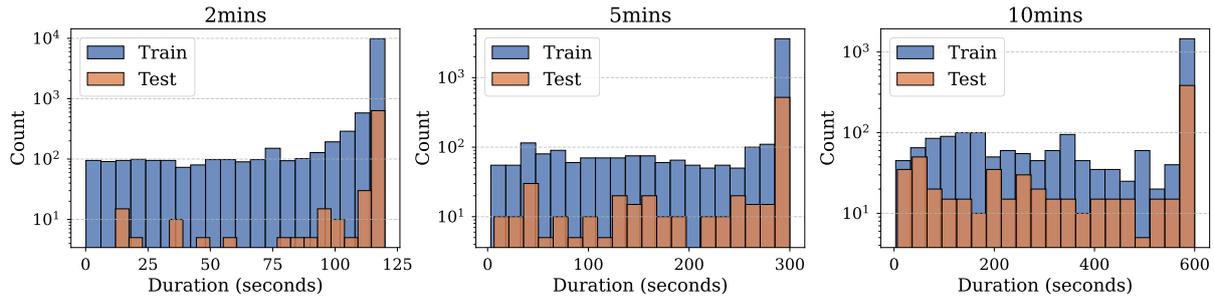


Figure 2: Distributions of audio segment lengths in the YODAS2-MCQA dataset, shown with *logarithmic counts* for improved visibility of the long-tail portions.

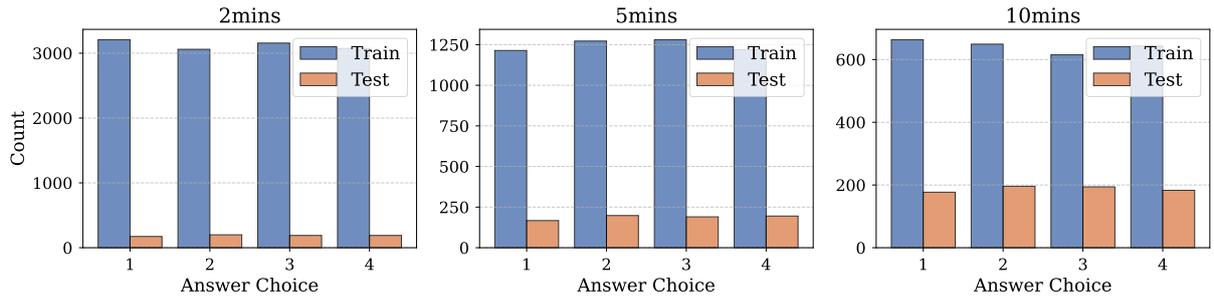


Figure 3: Distributions of correct answer choices for all question-answer pairs in the YODAS2-MCQA dataset.

**YODAS2-MCQA inference prompt**

Your answer **MUST** be in the format of "(NUMBER) STATEMENT". For example, if the answer was (4) A pen, you would **ONLY** output "(4) A pen". Do **NOT** include any other text. Using information from all the given audios. {question}

Figure 4: Prompt used for inferencing on YODAS2-MCQA.

erate five unique multiple-choice question-answer (MCQA) pairs for each audio chunk. The complete prompt utilized for this generation process is provided in Figure 9. The resulting dataset, which we have named **YODAS2-MCQA**, is publicly available at: <https://huggingface.co/datasets/yophis/yodas2-mcqa>.

For a more granular analysis of the YODAS2-MCQA dataset, we provide visualizations of the distributions of key characteristics. The distributions of audio segment lengths are presented in Figure 2. Additionally, the distribution of correct

answers across the multiple-choice options is presented in Figure 3.

## E Further Experimental Results

### E.1 Different VLAT's Virtual Length Sampling Strategies

The effectiveness of Virtual Longform Audio Training (VLAT) depends on the strategy used to sample the virtual length during training. For generality, here we define the virtual length as a produce of true audio context length (i.e. 30s) and the *virtual factor*. In our main experiment Section 6, we used a strategy of randomly sampling a factor from the range [1x, 5x, 10x, 15x, 20x, 25x]. To understand the sensitivity of our results to this choice, we compare our default strategy against four alternatives, training a Qwen2-Audio model on each strategy and evaluating on the 10mins audio task.

The virtual factor sampling strategies are:

1. **Default:** Uniform random sampling from {1.0, 5.0, 10.0, 15.0, 20.0, 25.0}. This cover the audio context length of 30s (1.0x) to 12.5mins (25.0x).
2. **Dense Sampling:** Random sampling from a discrete set of 100 uniformly spaced points within the default strategy's range.

3. **Very Dense Sampling:** Random sampling from a discrete set of 1000 uniformly spaced points within the default strategy’s range.
4. **Limited Range:** Uniform random sampling from  $\{1.0, 5.0, 10.0\}$ . This factor set covers audio context length of at most 5mins (10.0x).
5. **Fixed High-Factor:** No randomization; the scale factor is always fixed at 20.0x.

We report the result in Table 6.

**Analysis.** The experiment offers several insights into different VLAT sampling strategies: (1) *Fixed High-Factor* (always 20.0x) strategy performs the worst by a large margin, confirming that exposure to a variety of positional ranges is crucial for robust generalization, rather than training for just a single fixed target length. (2) *Default* strategy which uses a coarse set of 6 sampling points outperforms both the *Dense* (100 points) and *Very Dense* (1000 points) strategies. This suggests that finegrained sampling of virtual factors provides no additional benefit, and that a relatively small but diverse set of factors is sufficient for training length-robust models.

The most surprising result comes from the *Limited Range* strategy, which achieves the highest performance across all evaluation lengths, despite its virtual lengths not extending beyond the 5mins mark (10.0x factor). While this demonstrates a powerful capacity for extrapolation, the mechanism behind this superior performance is not immediately clear. It may suggest that randomization on extremely long virtual lengths could introduce instability to the training, which could in turn explain why Whole YaRN diverged quickly under VLAT framework; or that the specific distribution of factors in this limited set is coincidentally optimal for our test data. We left the more in-depth study of this behavior as future work.

Given this uncertainty, we retained the *Default* strategy for our main experiments in Section 6. Its factors ranges intuitively and completely covers the lengths seen during evaluation, making it a more principled and safer choice.

## E.2 Hyperparameter Sensitivity of Partial YaRN

Here we analyze the effect and sensitivity of different Partial YaRN’s hyperparameter configurations.

Virtual Length Strategy	YODAS2-MCQA		
	2 mins	5 mins	10 mins
Default	96.87	91.91	81.73
100 sampling points	95.91	89.25	79.78
1000 sampling points	96.18	87.69	79.60
Limited range	<b>98.00</b>	<b>93.55</b>	<b>89.20</b>
Fixed to 20.0x	89.60	68.71	50.31

Table 6: **VLAT performance with different sampling strategies.** All models are trained on 2mins audio and evaluated with Partial PI inference. The results are averaged across three random seeds.

We report validation accuracies w.r.t. different cut-off dimension indices and attention temperatures, for the Qwen2-Audio and SALMONN models.

### E.2.1 Sensitivity on Qwen2-Audio

Figure 5 displays the hyperparameter sensitivity of Partial YaRN on Qwen2-Audio when extending from the original 30s audio context window. Figure 6 shows the sensitivity when extending from the observed audio context window of 2mins.

**Sensitivity trend.** We observe the trend of optimal hyperparameter configurations to move from the bottom left (high cutoff dimension index, temperature close to 1.0) toward the top right of the heatmaps (lower cutoff dimension index, higher temperature) as the audio length increases.

**Hyperparameter recommendation.** This trend informs the following recommendation when using Partial YaRN. For moderate context extensions (e.g., 2x to 4x), it is advisable to start with a higher cutoff dimension index and a temperature value close to 1.0. For longer extensions, a lower cutoff dimension index with a considerably higher temperature is likely to yield better results.

**Plausible explanation of this trend.** First, the attention temperature controls the distribution of attention scores. In very long sequences with thousands of tokens, the standard softmax function can lead to an uneven attention score distribution, such as attending to only a few tokens while ignoring the vast majority of the context. Raising the temperature flattens this distribution, encouraging the model to maintain a broader attention pattern across the entire long audio stream and preventing "attention collapse".

Second, the cutoff dimension index directly controls the "compression pressure" of the position encoding extension. A high cutoff index is expres-

sive enough for shorter extensions while retaining majority of the original, detailed positional information. On the other hand, a lower cutoff index (higher compression) is required to map a much larger range of positions into the smaller familiar audio region.

### E.2.2 Sensitivity on SALMONN

Figure 7 displays the hyperparameter sensitivity of Partial YaRN on SALMONN when extending from the original 30s audio context window.

**Sensitivity trend.** For this model, we observe that the optimal hyperparameter sets are generally contained in the low temperature region (0.5 - 0.7), with the cutoff dimension index moving from low to higher ones as the audio length grows.

**Hyperparameter recommendation.** From the observed trend, it's advisable to set the attention temperature to  $0.6(\pm 0.1)$ , adopt a low cutoff index for shorter audio length, and higher index for longer audio inputs.

**Plausible explanation of this trend.** The vast distinction between SALMONN and the previously discussed Qwen2-Audio lies in their audio representation lengths—SALMONN only uses 88 tokens for every 30s audio vs. Qwen2-Audio's 750 tokens. This, coupled with the model's preference of low attention temperature, suggests that letting the model focus on a relevant local region outweighs the need to spreading the attention distribution as in Qwen2-Audio.

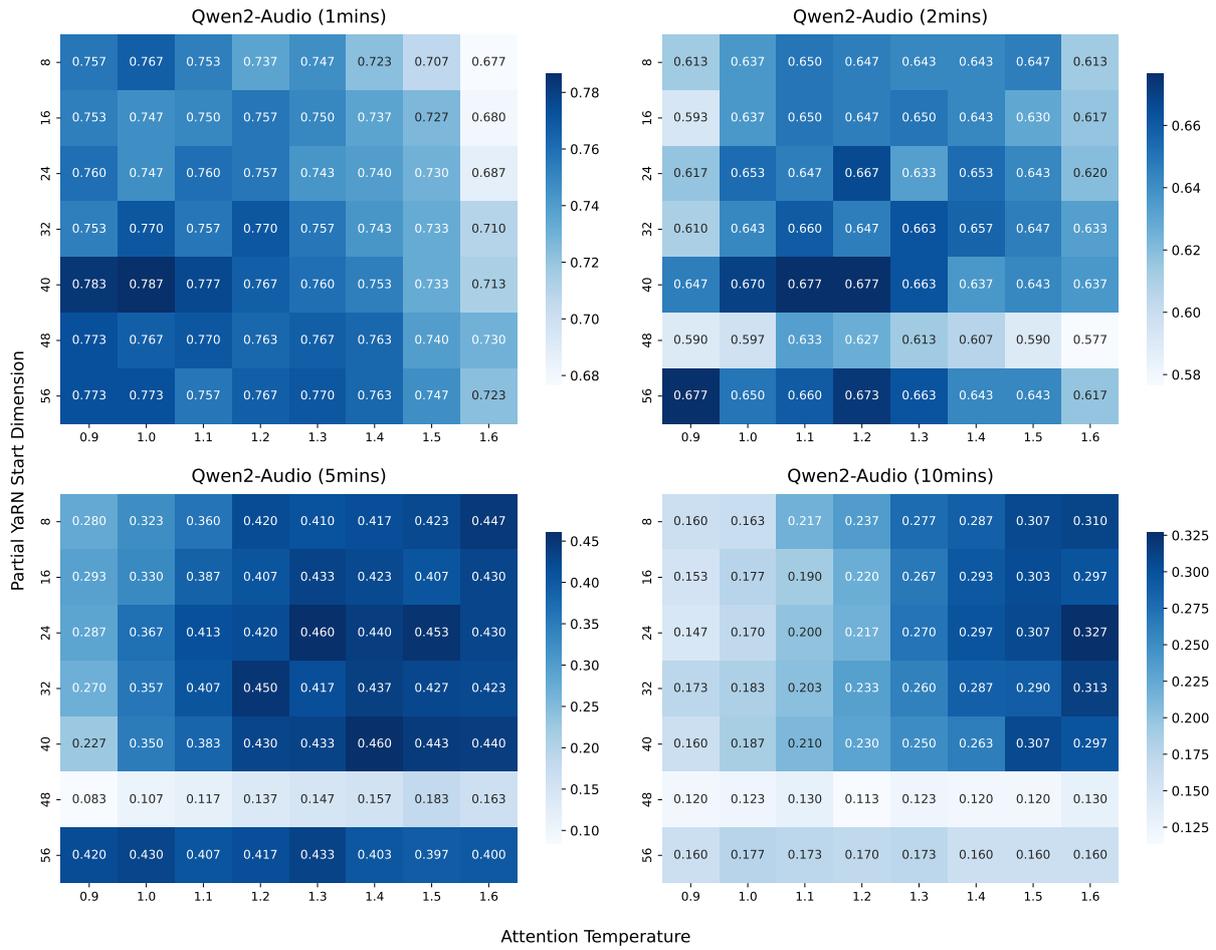


Figure 5: Hyperparameter Sensitivity of Partial YaRN on Qwen2-Audio (extending the original 30s audio context).

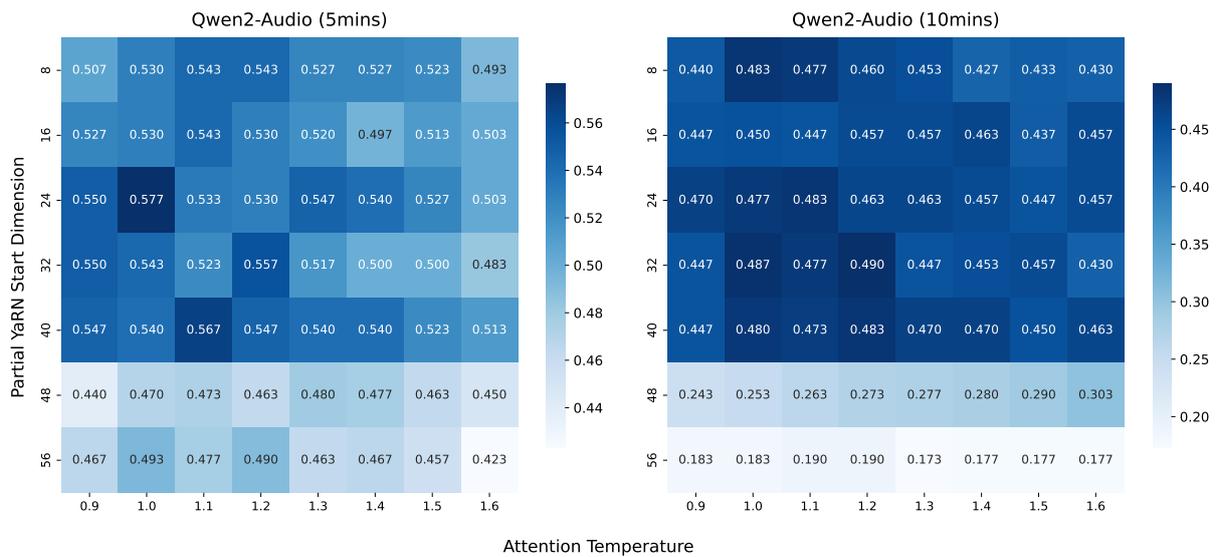


Figure 6: Hyperparameter Sensitivity of Partial YaRN on Qwen2-Audio (extending the observed 2mins audio context).

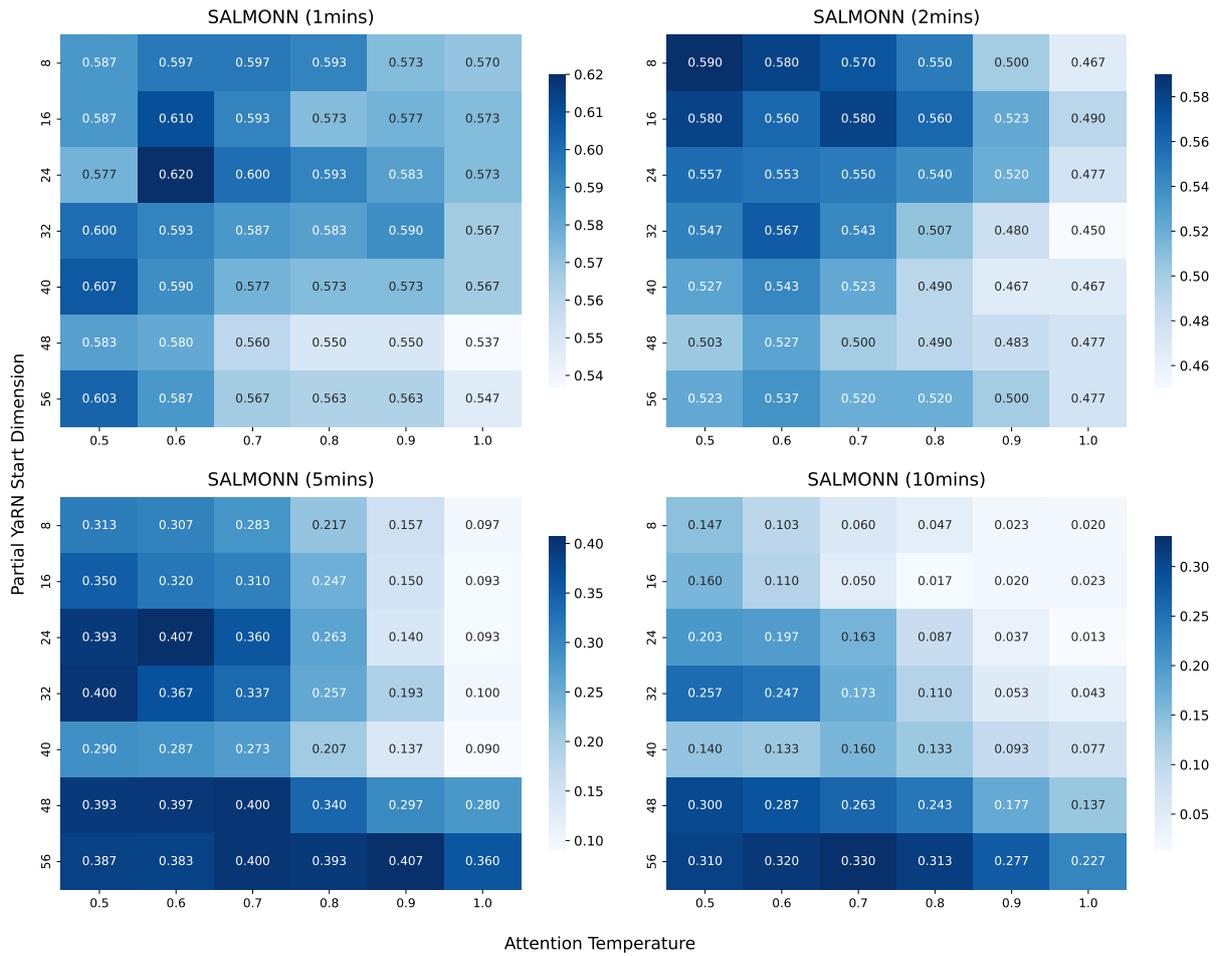


Figure 7: Hyperparameter Sensitivity of Partial YaRN on SALMONN (extending the original 30s audio context).

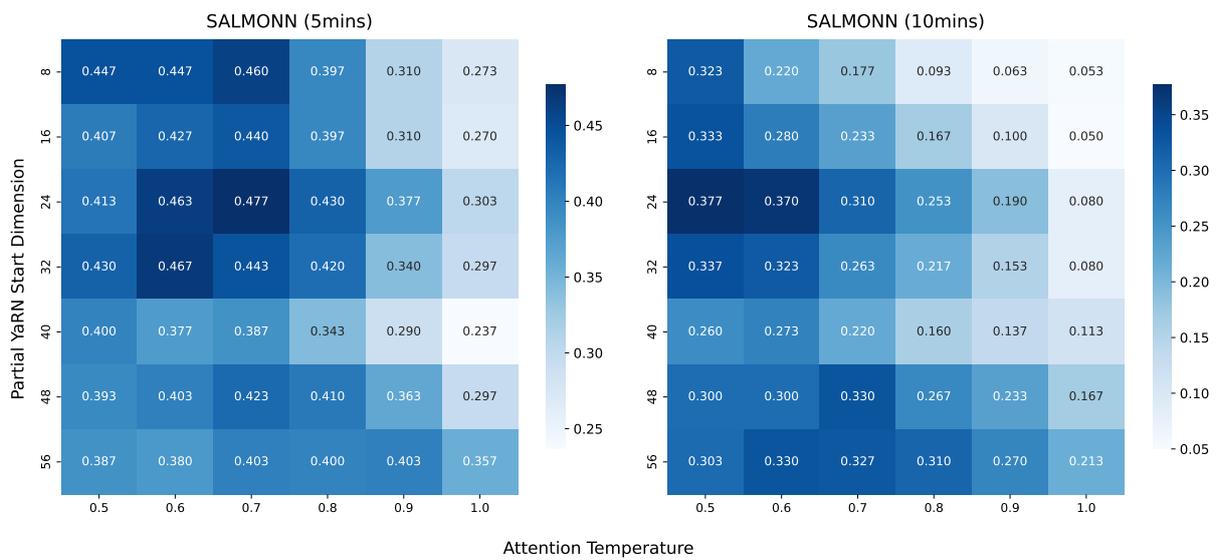


Figure 8: Hyperparameter Sensitivity of Partial YaRN on SALMONN (extending the observed 2mins audio context).

## YODAS2-MCQA generation prompt

You are an expert in creating high-quality Multiple-Choice Question Answering (MCQA) datasets for speech understanding. Your task is to generate QA pairs based on audio segments.

### \*\*Input:\*\*

You will be given a segment of an audio recording.

### \*\*Task:\*\*

For each given transcript, generate **5** distinct 4-Choice Question-Answer (QA) pairs that each has 1 correct and 3 incorrect choices (so four choices in total), and are factually grounded in the content of the transcript. The questions should cover different aspects of the audio segment, such as:

**Who:** Identify speakers or people mentioned.

**What:** Describe actions, objects, or events.

**When:** Determine times or dates mentioned.

**Where:** Locate places or locations.

**Why/How:** Explain reasons or processes.

**Details:** Ask for specific details mentioned in the audio segment.

**Purpose/Intention:** Infer the speaker's intent.

**Summary:** Condense information from the audio segment.

**Count:** How many people, men or women in the audio segment.

**Opinion/Sentiment:** What opinions or sentiment is expressed.

**Negation:** Understand what is NOT being said or done.

### \*\*Output:\*\*

Your output should be a JSON formatted string. The JSON should have the following structure:

```
"response": [ "question": "Question 1", "correct": "Correct answer", "incorrect1": "Incorrect answer1", "incorrect2": "Incorrect answer2", "incorrect3": "Incorrect answer3", "question": "Question 2", "correct": "Correct answer", "incorrect1": "Incorrect answer1", "incorrect2": "Incorrect answer2", "incorrect3": "Incorrect answer3", "question": "Question 3", "correct": "Correct answer", "incorrect1": "Incorrect answer1", "incorrect2": "Incorrect answer2", "incorrect3": "Incorrect answer3", "question": "Question 4", "correct": "Correct answer", "incorrect1": "Incorrect answer1", "incorrect2": "Incorrect answer2", "incorrect3": "Incorrect answer3", "question": "Question 5", "correct": "Correct answer", "incorrect1": "Incorrect answer1", "incorrect2": "Incorrect answer2", "incorrect3": "Incorrect answer3" ]
```

### Important Guidelines:

**Accuracy:** Ensure the answers are directly supported by content of the audio. Do not introduce information that is not present in the audio.

**Relevance:** The questions should be relevant to the content of the audio segment.

**Diversity:** The 5 QA pairs should be diverse in the type of question asked (who, what, when, where, why, how, detail, etc.). Try to use each question type at least once, if applicable to the transcript.

**Clarity:** The questions and answers should be clear, concise, and grammatically correct.

**Completeness:** The answers should fully address the question and provide sufficient information.

**Conciseness:** The answers should be as concise as possible while still being accurate and complete.

**Coverage:** The questions should collectively cover the entire audio segment. Aim for diverse coverage for each question. Prioritize the most salient and important information.

**Avoid Ambiguity:** Ensure the questions are not ambiguous and have a clear, single answer.

**Plausibility of Incorrect Answers:** The incorrect answers should be plausible but wrong, and related to the context of the transcript. Don't make them obviously incorrect or nonsensical. Use information or wording from the transcript to make them believable distractors.

**JSON Validity:** The output must be a valid JSON string. Double-check the formatting and escaping of special characters.

**Context-Awareness:** While each audio segment should be treated independently, the types of questions should logically relate to the context of human conversations. Try to create questions which someone would naturally ask.

**Answer Length Consistency:** Try to make the length of the correct and incorrect answers as similar as possible. This makes it more difficult to guess the answer based on length alone.

### Example:

**Input Transcript:** "Okay, so the meeting is scheduled for next Tuesday at 2 PM in Conference Room A. We'll be discussing the Q3 marketing report and the proposed budget changes."

### Output:

```
"response": [ "question": "When is the meeting scheduled?", "correct": "Next Tuesday at 2 PM", "incorrect1": "This Friday at 10 AM", "incorrect2": "Next Wednesday at 3 PM", "incorrect3": "Today at 4 PM", "question": "Where will the meeting take place?", "correct": "Conference Room A", "incorrect1": "Conference Room B", "incorrect2": "The Main Auditorium", "incorrect3": "The Executive Boardroom", "question": "What marketing report will be discussed?", "correct": "The Q3 marketing report", "incorrect1": "The Q2 marketing report", "incorrect2": "The annual marketing report", "incorrect3": "The preliminary marketing report", "question": "Besides the marketing report, what else will be discussed?", "correct": "The proposed budget changes", "incorrect1": "The sales figures", "incorrect2": "The employee satisfaction survey", "incorrect3": "The new product launch", "question": "What is the purpose of the meeting?", "correct": "To discuss the Q3 marketing report and budget changes", "incorrect1": "To plan the annual company picnic", "incorrect2": "To interview new job candidates", "incorrect3": "To celebrate the company's anniversary" ]
```

Now, generate 5 QA pairs in the specified JSON format for the input audio:

Figure 9: Prompt for generating multiple-choice question answering pairs from YODAS2.