

# MoE Adapter for Large Audio Language Models: Sparsity, Disentanglement, and Gradient-Conflict-Free

Yishu Lei<sup>1\*</sup>, Shuwei He<sup>1,2\*</sup>, Jing Hu<sup>1,3</sup>, Dan Zhang<sup>1</sup>, Xianlong Luo<sup>1</sup>, Danxiang Zhu<sup>1</sup>, Shikun Feng<sup>1†</sup>, Rui Liu<sup>2</sup>, Jingzhou He<sup>1</sup>, Yu Sun<sup>1</sup>, Hua Wu<sup>1</sup>, Haifeng Wang<sup>1</sup>

<sup>1</sup> ERNIE Team, Baidu    <sup>2</sup> College of Computer Science, Inner Mongolia University

<sup>3</sup> Tsinghua Shenzhen International Graduate School, Tsinghua University  
{leiyishu, heshuwei, zhangdan20, luoxianlong, zhudanxiang, fengshikun01, hejingzhou, sunyu02, wu\_hua, wanghaifeng}@baidu.com  
cminuser@gmail.com, imucslr@imu.edu.cn

## Abstract

Extending the input modality of Large Language Models (LLMs) to the audio domain is essential for achieving comprehensive multi-modal perception. However, it is well-known that acoustic information is intrinsically *heterogeneous*, entangling attributes such as speech, music, and environmental sounds. Existing methods typically rely on dense, parameter-shared adapters to model these diverse patterns, which induces *gradient conflict* during optimization, as parameter updates required for distinct attributes contradict each other. To address this limitation, we introduce the **MoE-Adapter**, a sparse Mixture-of-Experts (MoE) architecture designed to decouple acoustic information. Specifically, it employs a dynamic gating mechanism that routes audio tokens to specialized experts that capture complementary feature subspaces while retaining shared experts for global context, thereby mitigating gradient conflicts and enabling fine-grained feature learning. Comprehensive experiments show that the MoE-Adapter achieves superior performance on both audio semantic and paralinguistic tasks, consistently outperforming dense linear baselines with comparable computational costs. To facilitate future research, our code are publicly available at <https://github.com/Alittleegg/Eureka-Audio>.

## 1 Introduction

Large Language Models (LLMs) have revolutionized cognitive intelligence through their exceptional text-processing capabilities (Achiam et al., 2023; Guo et al., 2025; Yang et al., 2025; Wang et al., 2026; Zhang et al., 2026). However, limiting these models to the textual modality constrains their ability to perceive and interact with the dynamic physical world. As audio serves as a primary channel for both human communication and environmental perception, extending LLMs to the audio

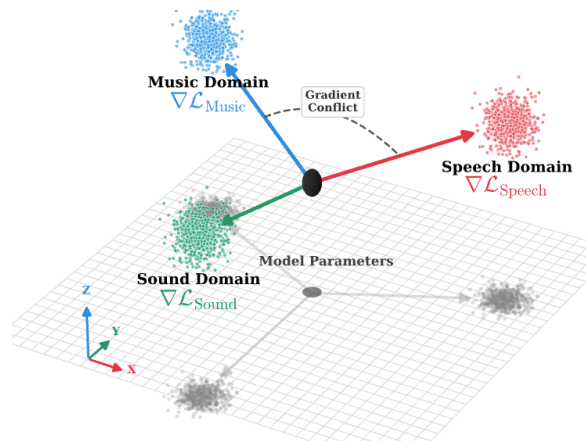


Figure 1: **Visualization of Gradient Conflict.** The naturally disjoint manifolds of Speech, Music, and Sound induce divergent gradient update directions ( $\nabla\mathcal{L}$ ), creating an optimization conflict that pulls the shared adapter in contradictory ways.

domain is a critical research frontier (Rubenstein et al., 2023; Ji et al., 2024; Shu et al., 2025). The prevailing paradigm involves inserting an adapter to project acoustic features into the textual semantic space. However, recent approaches predominantly rely on a dense, parameter-shared adapter for this purpose (Tang et al., 2023; Chu et al., 2024; Xie and Wu, 2024; Team et al., 2025; Xu et al., 2025). For instance, Kimi-Audio (Ding et al., 2025) utilizes a fixed MLP-based projector, while Audio Flamingo 3 (Goel et al., 2025) employs temporal compression layers with globally shared parameters. While straightforward, these designs implicitly assume that a single unified projection can uniformly capture the intrinsic heterogeneity of complex audio signals.

However, as illustrated in Figure 1, this assumption conflicts with the intrinsic topology of the acoustic modality: rather than conforming to a uniform distribution, human speech, music, and environmental sounds occupy distinct manifolds with divergent statistical properties (Gemmeke et al.,

\*Equal contribution.

†Corresponding author.

2017; Yang et al., 2021). Constraining a static, parameter-shared adapter to map these heterogeneous inputs into a unified space precipitates significant gradient conflicts (Chen et al., 2018; Yu et al., 2020). Specifically, parameter updates required for semantic speech often contradict those for paralinguistic cues, engendering destructive interference within the shared weights (Liu et al., 2021). We empirically substantiate these conflicts in Section 5. Consequently, a monolithic adapter is ill-suited for simultaneous alignment and disentanglement. This raises a fundamental question: *How can an architecture be designed to explicitly accommodate the heterogeneous structure of audio data rather than suppressing it?*

Mixture-of-Experts (MoE) architectures offer a principled inductive bias for managing such intrinsic distributional heterogeneity. By conditionally activating sparse subsets of experts, MoE enables the decomposition of complex data manifolds into manageable subspaces, allowing distinct experts to specialize in specific acoustic attributes (Shazeer et al., 2017; Fedus et al., 2022). This specialization effectively orthogonalizes optimization directions, mitigating gradient conflicts arising from entangled features. While the efficacy of this paradigm has been rigorously validated in vision and multi-modal domains (Riquelme et al., 2021; Mustafa et al., 2022; Zhang et al., 2025a; Li et al., 2025b; Su et al., 2025), its potential remains largely untapped in audio-text alignment. To date, prevalent Large Audio Language Models (LALMs) persist in their reliance on static, monolithic adapters, leaving the challenge of acoustic heterogeneity unaddressed.

To bridge this gap, we introduce the **MoE-Adapter**, a sparse and dynamic architecture explicitly designed to resolve optimization bottlenecks induced by acoustic heterogeneity. Departing from the monolithic paradigm that subjects all parameters to uniform global updates, our approach leverages a learnable gating mechanism to dynamically route input segments to specialized experts based on their intrinsic acoustic attributes. This mechanism orchestrates a decoupled optimization process: gradient updates from conflicting signal types are isolated to specific experts, significantly reducing interference while retaining shared experts for common feature extraction. Consequently, the MoE-Adapter achieves effective disentanglement of diverse acoustic factors and alignment with the textual space, all while maintaining inference-time computational costs comparable to dense baselines.

In summary, the main contributions of this paper are as follows:

- **Architectural Innovation:** We propose the MoE Adapter, a sparse, dynamically routed projection layer that seamlessly replaces static adapters, explicitly targeting gradient conflicts via expert specialization.
- **Empirical Insight:** We substantiate the existence of severe gradient interference across semantic, paralinguistic, and environmental factors in parameter-shared adapters, and demonstrate how our method effectively alleviates these optimization bottlenecks.
- **Performance & Efficiency:** We validate the proposed architecture on a 1.7B-parameter LLM backbone. The MoE-Adapter consistently outperforms dense baselines in both semantic understanding and paralinguistic tasks, demonstrating a more favorable performance–efficiency trade-off under a fixed model capacity.

## 2 Related Works

### 2.1 Large Audio Language Models

LALMs have evolved from cascaded pipelines to unified end-to-end frameworks (Ji et al., 2024). Early cascaded models (Zhang et al., 2023; Huang et al., 2024) relied on Automatic Speech Recognition (ASR) transcripts, suffering from error propagation and the loss of non-verbal cues. In contrast, end-to-end models (Rubenstein et al., 2023; Shu et al., 2025) bridge the modality gap by mapping acoustic features to the textual space via learnable adapters. These adapters typically fall into two categories: Q-Former-based architectures, exemplified by SALMONN (Tang et al., 2023), and Linear Projectors, as adopted by Qwen2-Audio (Chu et al., 2024), GLM-4-Voice (Zeng et al., 2024), LLaMA-Omni (Fang et al., 2024), Step-Audio2 (Wu et al., 2025) and Kimi-Audio (Ding et al., 2025). While simpler and more efficient, current state-of-the-art models largely favor the latter, utilizing a dense, global projection layer. However, this prevailing design implicitly assumes a uniform distribution across diverse audio types. As a result, it struggles to resolve the distributional shifts between semantic speech and acoustic events, leading to gradient conflicts (Yu et al., 2020; Zhang et al., 2025b; Li et al., 2025a). In this work, we challenge

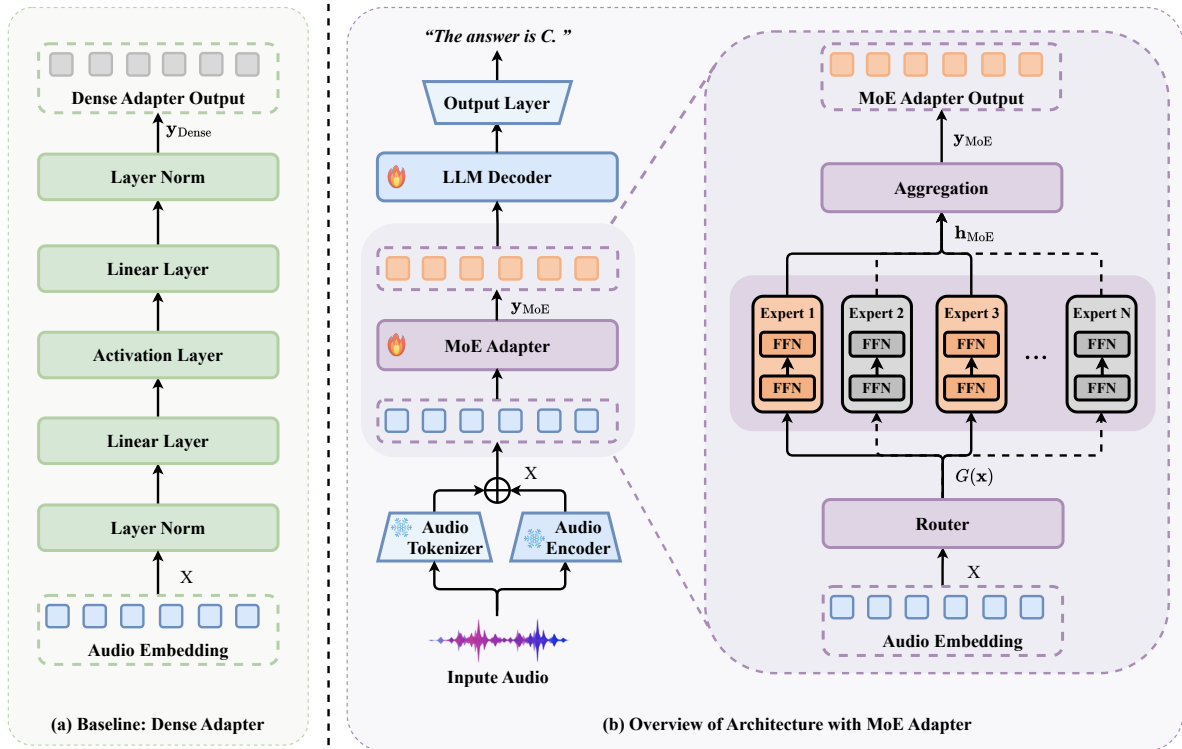


Figure 2: **The overall architecture of the MoE Adapter compared with the Dense Adapter baseline.** Unlike the monolithic, parameter-shared design in (a), the MoE-Adapter in (b) employs a sparse gating mechanism to dynamically route audio tokens to specialized experts. This architecture facilitates fine-grained feature learning and effectively mitigates gradient conflicts by disentangling heterogeneous acoustic attributes during the projection phase.

this monolithic convention by introducing a sparse MoE mechanism.

## 2.2 Mixture-of-Experts Architectures

The MoE paradigm has evolved from decomposing complex problems (Jacobs et al., 1991) into a standard for efficient training that handles data heterogeneity and mitigates gradient conflicts via sparse gating (Shazeer et al., 2017). Recent studies in multimodal domains (Chen et al., 2023; Li et al., 2025b; Zhang et al., 2025a) validate this capability, demonstrating that routing conflicting tasks to orthogonal experts effectively resolves optimization bottlenecks where different modalities require contradictory parameter updates (Huang et al., 2025; Xin et al., 2025). In the audio domain, the adoption of MoE is similarly expanding. Recent works have successfully adapted this paradigm for specific applications, exemplified by speech and music generation (Liu et al., 2025) and robust feature selection in healthcare (Shao et al., 2025). However, despite these advancements, the potential of MoE for general-purpose audio–text alignment remains largely untapped. Prevalent LALMs persist

in relying on static, parameter-shared adapters—a monolithic design that forces a single set of parameters to model intrinsically diverse acoustic attributes. We address this limitation by introducing the MoE-Adapter, leveraging dynamic expert routing to explicitly disentangle conflicting acoustic patterns during the alignment phase.

## 3 Method

### 3.1 Overall Framework

As depicted on the left side of Figure 2 (b), our model leverages the dual-stream architecture of Kimi-Audio (Ding et al., 2025). The audio front-end employs two pathways: a frozen tokenizer extracts *discrete semantic tokens*, while a speech encoder captures *continuous acoustic features*. These heterogeneous representations are subsequently aligned via a projection layer and fused through element-wise summation. Instead of using a standard dense adapter, we employ a MoE-Adapter to dynamically map the fused features into the LLM’s embedding space. Finally, the adapted audio embeddings are concatenated with text token embed-

dings and serve as input to the backbone for standard autoregressive next-token prediction (NTP).

### 3.2 Sparse MoE Adapter

**Preliminary: Dense Adapter.** As shown in Figure 2 (a), the LALMs typically employ a monolithic Feed-Forward Network (FFN) as the modality projector. Formally, given an input audio token  $\mathbf{x} \in \mathbb{R}^d$ , the projected embedding  $\mathbf{y} \in \mathbb{R}^D$  is formulated as:

$$\mathbf{y} = \mathcal{N}\left(\mathbf{W}_{d2} \cdot \sigma\left(\mathbf{W}_{d1} \cdot \mathcal{N}(\mathbf{x})\right)\right), \quad (1)$$

where  $\mathcal{N}(\cdot)$  denotes Layer Normalization,  $\sigma(\cdot)$  represents the SiLU activation function, and  $\mathbf{W}_{d1}, \mathbf{W}_{d2}$  are the learnable weight matrices specific to the dense projector. This output  $\mathbf{y}$  is subsequently utilized to substitute audio placeholder tokens within the LLM input sequence. However, this monolithic design imposes a rigid inductive bias: a single set of parameters is forced to accommodate heterogeneous acoustic modalities, inevitably inducing optimization interference.

**Sparse Adapter with MoE Mechanisms.** Figure 2 (b) illustrates the overall architecture of the MoE Adapter. To facilitate fine-grained feature specialization, we replace the monolithic architecture with a bank of  $N$  experts, denoted as  $\{E_i\}_{i=1}^N$ . Each expert  $E_i$  is parameterized as a lightweight FFN utilizing the SiLU activation  $\phi(\cdot)$ :

$$E_i(\mathbf{x}) = \mathbf{W}_{e2}^{(i)} \cdot \phi\left(\mathbf{W}_{e1}^{(i)} \cdot \mathcal{N}(\mathbf{x})\right), \quad (2)$$

where  $\mathbf{W}_{e1}^{(i)}$  and  $\mathbf{W}_{e2}^{(i)}$  denote the unique weight matrices for the  $i$ -th expert. To orchestrate expert selection, a learnable router  $G(\cdot)$  computes sparse gating probabilities based on the logits  $\mathbf{s} = \mathbf{x}\mathbf{W}_g$ , where  $\mathbf{W}_g$  is the gating matrix:

$$G(\mathbf{x}) = \text{softmax}\left(\mathcal{T}_k(\mathbf{s})\right). \quad (3)$$

Here,  $\mathcal{T}_k(\cdot)$  represents the Top- $k$  operator, which retains the  $k$  largest values while masking the remainder to  $-\infty$ . The expert outputs are synthesized via weighted aggregation to yield the intermediate representation  $\mathbf{h}_{\text{MoE}}$ :

$$\mathbf{h}_{\text{MoE}} = \sum_{i \in \mathcal{I}} G(\mathbf{x})_i \cdot E_i(\mathbf{x}), \quad (4)$$

where  $\mathcal{I}$  denotes the set of indices corresponding to the active experts selected by the router.

**Output Projection and Aggregation.** To align the aggregated features with the LLM’s embedding space,  $\mathbf{h}_{\text{MoE}}$  is processed by an aggregation block consisting of layer normalization followed by a two-layer feed-forward network:

$$\mathbf{y}_{\text{MoE}} = \mathbf{W}_{a2} \cdot \sigma\left(\mathbf{W}_{a1} \cdot \mathcal{N}(\mathbf{h}_{\text{MoE}})\right), \quad (5)$$

where  $\mathbf{W}_{a1}$  and  $\mathbf{W}_{a2}$  are the learnable weight matrices of the aggregation MLP. This adapted output  $\mathbf{y}_{\text{MoE}}$  is used to replace the pre-defined audio placeholders in the input sequence, thereby conditioning the next-token prediction task on the optimized, disentangled acoustic context.

### 3.3 Training Objectives

The entire model is optimized end-to-end using a joint objective:

$$\mathcal{L} = \mathcal{L}_{\text{NTP}} + \lambda \mathcal{L}_{\text{aux}}. \quad (6)$$

**Next-Token Prediction Loss.** The primary training objective is next-token prediction conditioned on the audio context:

$$\mathcal{L}_{\text{NTP}} = - \sum_{t=1}^T \log P(y_t | y_{<t}, \mathbf{X}; \theta), \quad (7)$$

where  $\theta$  denotes the model parameters and  $\mathbf{X}$  refers to the fused audio representation before the adapter, whose adapted form is used to condition the LLM.

**Auxiliary Load-Balancing Loss.** To mitigate expert collapse, we adopt a standard load-balancing loss (Lepikhin et al., 2021) over the routed experts  $\mathcal{E}_R$ . Let  $B$  denote the number of routed tokens in a batch. For each token  $\mathbf{x}_b$ , the router outputs a routing probability  $p_{b,e} = G(\mathbf{x}_b)_e$  for expert  $e \in \mathcal{E}_R$ . We define the *expert importance* as the average routing probability,

$$\bar{P}_e = \frac{1}{B} \sum_{b=1}^B p_{b,e}, \quad (8)$$

and the *expert load* as the fraction of tokens routed to expert  $e$ ,

$$\bar{f}_e = \frac{1}{B} \sum_{b=1}^B r_{b,e}, \quad (9)$$

where  $r_{b,e} = 1$  if expert  $e$  is selected for token  $\mathbf{x}_b$  by Top- $k$  routing, and 0 otherwise. The auxiliary loss is then computed as:

$$\mathcal{L}_{\text{aux}} = |\mathcal{E}_R| \sum_{e \in \mathcal{E}_R} \bar{P}_e \cdot \bar{f}_e. \quad (10)$$

Benchmark	Text Accuracy ( $\uparrow$ )	Audio Accuracy ( $\uparrow$ )			Modality Gap ( $\uparrow$ )	
	Backbone LLM	Baseline	Ours	Improv.	Baseline	Ours
<i>Task: Audio Knowledge Reasoning</i>						
MMSU (Chen et al., 2024)	52.86	35.03	<b>38.19</b>	<b>+3.16</b>	-17.83	<b>-14.67</b>
OBQA (Chen et al., 2024)	68.35	50.10	<b>53.85</b>	<b>+3.75</b>	-18.25	<b>-14.50</b>
<i>Task: Audio Perception</i>						
MMAU (Sakshi et al., 2024)	-	59.79	<b>61.50</b>	<b>+1.71</b>	-	-

Table 1: **Main results on audio understanding and cross-modal alignment.** Compared to a dense baseline under a controlled budget, MoE-Adapter consistently improves audio accuracy and reduces the modality gap (the performance disparity between audio and text inputs), demonstrating superior representational capacity.

## 4 Experiments

### 4.1 Experimental Setups

**Model Architecture & Training Settings.** The LLM backbone utilizes Qwen3-1.7B<sup>1</sup>. The audio frontend employs the Whisper-VQ tokenizer (Zeng et al., 2024) and the Whisper Encoder (Radford et al., 2022). Training is conducted on a high-quality 40B-token corpus using AdamW ( $\beta_1 = 0.9, \beta_2 = 0.95$ ) and a Warmup-Stable-Decay scheduler (Wen et al., 2024) (peak LR  $1 \times 10^{-5}$ , 20 warmup steps). For fair comparison<sup>2</sup>, we constrain the total parameter budget of both the dense adapter and the MoE-Adapter to 94.4M parameters. Notably, due to sparse activation, the MoE-Adapter activates only approximately 70.8M parameters during inference, corresponding to about 75% of the baseline. In addition, we fix random seeds across all experiments to control stochasticity and ensure experimental fairness.

**Evaluation.** We conduct few-shot evaluations on widely adopted benchmarks. Specifically, we leverage MMAU (Sakshi et al., 2024)—a large-scale benchmark covering speech, sounds, and music—for *perception-oriented audio understanding*, focusing on acoustic and paralinguistic reasoning. For world knowledge reasoning, we utilize the MMSU and OpenBookQA (OBQA) subsets of VoiceBench (Chen et al., 2024), which are audio adaptations of the text-only MMLU-Pro (Wang et al., 2024) and require *higher-level semantic reasoning* beyond low-level perceptual cues. All evaluations are performed using greedy decoding to ensure deterministic and comparable results.

<sup>1</sup><https://huggingface.co/Qwen/Qwen3-1.7B>

<sup>2</sup>Additional training hyperparameters and implementation details are provided in Appendix.

### 4.2 Main Results

Table 1 presents the performance comparison between our method and the baseline across all benchmarks. Here, we highlight how the MoE-Adapter drives consistent improvements across three key dimensions: knowledge reasoning, audio perception, and cross-modal alignment.

**Knowledge Reasoning Gains.** Our method consistently outperforms the baseline across audio knowledge benchmarks. Specifically, on MMSU and OBQA, we achieve gains of **3.16%** (35.03%  $\rightarrow$  38.19%) and **3.75%** (50.10%  $\rightarrow$  53.85%), respectively. These results suggest that expert specialization enables the model to effectively capture knowledge-centric semantic information, thereby enhancing both commonsense reasoning and cross-domain knowledge integration.

**Robust Audio Perception** Beyond knowledge-centric tasks, our method also demonstrates robust performance on paralinguistic understanding. Specifically, it achieves a **1.71%** absolute improvement on the MMAU benchmark (59.79%  $\rightarrow$  61.50%). This result indicates that the MoE mechanism effectively captures diverse acoustic and paralinguistic cues, leading to improved audio perception and demonstrating robustness across heterogeneous audio domains.

**Reduced Modality Gap.** We further quantify audio-text consistency using the Modality Gap, where our method shows consistent reductions. For instance, on MMSU, the gap narrows significantly from  $-17.83$  to  $-14.67$  (an improvement of **3.16**). These gains confirm that the dynamic routing mechanism maps heterogeneous acoustic representations into subspaces that are more compatible with the LLM of text embedding space, yielding superior alignment compared to static dense adapters.

### 4.3 Ablation Studies

In this section, we conduct ablation studies to examine the effectiveness of key architectural and routing design choices in MoE-Adapter. We analyze how expert configuration, routing sparsity, and training objectives influence audio understanding and reasoning performance across different benchmarks.

#### 4.3.1 Expert Configuration and Sparsity

We analyze how expert configuration and routing sparsity affect audio understanding and reasoning under the same training token budget. Results on MMAU, MMSU, and OBQA are summarized in Table 2.

Configuration	Dim	Accuracy ( $\uparrow$ )		
		MMAU	MMSU	OBQA
<b>8 choose 4 (Ours)</b>	<i>D</i>	<b>61.50</b>	<b>38.19</b>	<b>53.85</b>
16 choose 4	<i>D</i>	60.80	34.77	48.35
8 choose 1	<i>D</i>	59.89	32.95	43.95
4 choose 2	<i>D</i>	61.10	38.12	52.74
8 choose 4	<i>D/2</i>	60.60	37.28	54.72
8 choose 4	<i>2D</i>	59.69	35.94	50.54

Table 2: **Ablation on expert configurations.** We analyze the impact of expert count, routing sparsity, and dimension capacity. The default setting “8 choose 4” (highlighted) achieves the best overall balance.

Notably, increasing the total number of experts does not necessarily improve performance: expanding from “8 choose 4” to “16 choose 4” consistently degrades performance across all benchmarks. Regarding routing sparsity, stronger constraints are not always beneficial; activating too few experts, such as “8 choose 1”, significantly harms audio reasoning ability, while moderate sparsity (specifically “4 choose 2” or “8 choose 4”) yields stronger results. Furthermore, larger expert capacity is not inherently better: increasing expert dimensionality (“8 choose 4,  $2D$ ”) leads to consistent performance drops, whereas reduced capacity can remain competitive on certain reasoning benchmarks. Crucially, under the same training budget, configurations with larger effective capacity through more active experts (exemplified by “4 choose 2” vs. “8 choose 4”) can achieve comparable or improved performance, highlighting the importance of balancing model capacity and routing sparsity.

In conclusion, effective audio reasoning demands a careful balance of expert count, sparsity, and capacity, rather than simply scaling any single dimension in isolation.

#### 4.3.2 Impact of Expert Balance Loss

To isolate the impact of the expert balance loss ( $\mathcal{L}_{\text{aux}}$ ) on expert utilization and downstream performance, we evaluate a variant trained without this auxiliary term (w/o EBL). Table 3 summarizes the results across MMAU, MMSU, and OBQA.

Method	MMAU	MMSU	OBQA
<b>Ours (w/ EBL)</b>	61.50	<b>38.19</b>	<b>53.85</b>
w/o EBL	<b>63.01</b>	37.37	52.31

Table 3: **Ablation on Expert Balance Loss.** We report accuracy (%) on audio reasoning and understanding benchmarks.

Removing the balance loss yields divergent outcomes across benchmarks. On semantic reasoning tasks, the model incorporating EBL consistently outperforms the ablated counterpart by **0.82** and **1.54** points, respectively. This validates that load balancing effectively mitigates expert collapse, ensuring diverse experts are sufficiently trained—a prerequisite for handling complex semantic reasoning grounded in world knowledge.

Conversely, the w/o EBL variant achieves higher accuracy on MMAU. We attribute this behavior to the intrinsic heterogeneity of MMAU, which spans multiple audio modalities (speech, sound, and music) while being dominated by perceptual and low-level acoustic patterns in a large portion of instances. In the absence of explicit load balancing, the router naturally concentrates updates on a subset of “dominant experts” that are well aligned with these frequently occurring perceptual and low-level acoustic patterns, resulting in improved performance on perception-oriented audio understanding tasks. However, this concentration of expert utilization reduces the effective diversity of experts participating in optimization, which limits the model’s ability to support systematic generalization and higher-level semantic reasoning. This effect is reflected in the degraded performance observed on MMSU and OBQA, where successful prediction requires grounding audio inputs in world knowledge and performing more complex semantic inference.

In summary, the expert balance loss imposes a trade-off: it slightly constrains perception-oriented performance to significantly enhance semantic reasoning by ensuring balanced expert utilization. We next analyze the routing dynamics to elucidate this mechanism.

## 5 Analysis of Expert Specialization and Optimization Dynamics

The MoE-adapter is motivated by the heterogeneous nature of audio data, where different acoustic patterns and reasoning requirements benefit from specialized representations. In this section, we analyze the internal behaviors of the MoE-adapter to explain its performance gains over dense adapters, as well as the trade-offs observed when introducing expert balance regularization. Our analysis focuses on two complementary aspects: (i) how routing-based expert specialization emerges across heterogeneous audio modalities, and (ii) how such specialization reshapes the optimization dynamics by mitigating gradient conflicts.

### 5.1 Effect of Expert Balance on Routing and Specialization

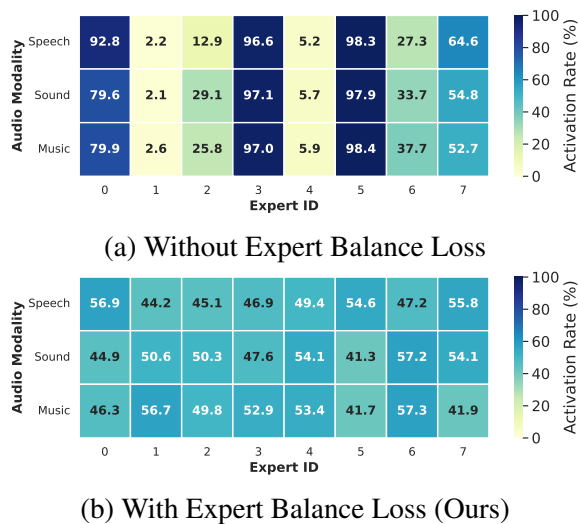


Figure 3: **Expert Activation Heatmaps.** The expert balance loss prevents collapse, fostering balanced utilization while maintaining modality-aware specialization. Notably, Sound serves as an “Acoustic Bridge” sharing experts with both Speech and Music.

To understand how the MoE adapter realizes feature disentanglement in practice, we analyze expert activation patterns on the MMAU, where samples are categorized into Speech, Sound, and Music. Figure 3 visualizes the token-normalized activation rates of routable experts, aggregated over forward-pass routing decisions. This comparison directly relates to the ablation results in Section 4.3.2, where the model without expert balance loss achieves higher performance on the heterogeneous MMAU benchmark. We observe clear expert specialization across audio modalities. Certain experts are pre-

dominantly activated by a single category, while others are shared between Sound and either Speech or Music. Notably, no expert is jointly specialized for both Speech and Music. This pattern reflects the intrinsic heterogeneity of audio data: although Sound shares low-level acoustic characteristics with both domains, Speech and Music differ substantially in temporal structure and semantic organization. As a result, Sound-centric experts naturally emerge as an acoustic bridge, enabling controlled parameter sharing without forcing incompatible representations to be jointly optimized.

Importantly, expert balance loss does not suppress specialization, but rather regularizes its extent. Compared to the unbalanced setting, it prevents extreme expert dominance while preserving modality-aware routing behavior. This regularization introduces a clear trade-off: while stronger expert dominance in the w/o EBL setting can benefit broad and heterogeneous benchmarks such as MMAU by concentrating capacity on a small subset of dominant experts, it reduces expert diversity and negatively impacts reasoning-intensive tasks. This emergent expert decomposition provides concrete evidence that our design motivation—explicitly accounting for audio heterogeneity—is well-founded. By allowing modality-aware specialization, the MoE adapter aligns heterogeneous acoustic inputs into a representation space that is more compatible with the underlying language model.

Beyond routing behavior, such specialization also manifests in the optimization dynamics. In the following section, we analyze how expert-based routing mitigates gradient conflicts that commonly arise when learning from heterogeneous audio sources.

### 5.2 Gradient Conflicts and Mitigation Mechanisms

The adoption of the MoE architecture is primarily motivated by the need to resolve gradient conflicts—a phenomenon frequently observed in dense models when learning heterogeneous tasks. In the context of LALMs, the adapter must simultaneously encode divergent modalities: Speech (semantic-dominant) versus Music and Sound (paralinguistic-dominant). To rigorously quantify how the proposed MoE Adapter mitigates these optimization bottlenecks compared to a dense FFN baseline, we analyze the training dynamics using two complementary metrics: Gradient Cosine Similarity and the Gradient Influence Score.

**Gradient Cosine Similarity.** To investigate optimization interference, we analyze the geometric alignment of gradient vectors across audio categories. Specifically, for any two categories  $i$  and  $j$ , we compute the pairwise cosine similarity of their gradients:

$$\text{Sim}(g_i, g_j) = \frac{g_i \cdot g_j}{\|g_i\| \|g_j\|} \quad (11)$$

As illustrated in Figure 4 (Left), the monolithic adapter exhibits persistent negative similarities (e.g.,  $-0.020$  between Music and Sound), implying that the optimization direction for one modality is geometrically opposed to another. This negative correlation confirms the presence of destructive interference, where minimizing the loss for one task inadvertently hinders the optimization of others.




Figure 4: **Gradient Cosine Similarity.** The dense FFN adapter (left) exhibits negative correlations between heterogeneous audio tasks, whereas the MoE adapter (right) demonstrates improved gradient alignment, indicating reduced optimization conflicts.

In contrast, the MoE-Adapter effectively resolves these conflicts, as evidenced by the shift towards positive gradient similarities in Figure 4 (Right) (e.g., Music-Speech improves to  $+0.023$ ). This emergent orthogonality indicates that the dynamic router successfully decouples the parameter space, directing conflicting gradients to distinct experts. This “subspace specialization” enables different experts to capture specific acoustic features independently, thereby mitigating the oscillatory dynamics observed in the baseline.

**Gradient Influence Analysis.** To move beyond geometric alignment and measure the causal impact of a gradient update from one task on the loss of another, we employ the Influence Score  $\mathcal{I}_{i,j}$ . Let  $\theta_t$  denote the current model parameters. We simulate a hypothetical update using only the gradient from task  $j$ , resulting in  $\theta_{t+1}^{(j)} = \theta_t - \lambda \frac{\nabla_{\theta} \mathcal{L}_j}{\|\nabla_{\theta} \mathcal{L}_j\|}$ . The resulting reduction in loss for task  $i$  is denoted as  $\Delta_j \mathcal{L}_i = \mathcal{L}_i(\theta_t) - \mathcal{L}_i(\theta_{t+1}^{(j)})$ . The normalized




Figure 5: **Gradient Influence Analysis.** Values represent the impact of updating task  $j$  (x-axis) on task  $i$  (y-axis). Unlike the FFN baseline (left) which suffers from harmful interference (negative values), the MoE adapter (right) promotes constructive transfer (positive values), effectively mitigating cross-task conflicts.

influence score is given by:

$$\mathcal{I}_{i,j} = \mathbb{E}_{\mathcal{D}} \left[ \frac{\Delta_j \mathcal{L}_i}{\Delta_i \mathcal{L}_i} \right] \quad (12)$$

where the expectation is taken over the validation batch. Intuitively, a positive score ( $\mathcal{I}_{i,j} > 0$ ) indicates constructive synergy (updating on  $j$  benefits  $i$ ), while a negative score implies destructive interference.

The results, presented in Figure 5, reveal a sharp contrast. The Dense Adapter (Left) suffers from significant negative interference; notably, updates derived from Speech data actively degrade performance on Music ( $\mathcal{I}_{\text{Music,Speech}} = -0.117$ ) and Sound. Conversely, the MoE Adapter (Right) exhibits predominantly positive influence scores. For instance, updates on Speech induce a strong positive transfer to Sound, while interference on Music is substantially mitigated. This suggests that the MoE architecture does not merely isolate conflicts; by retaining shared experts for common acoustic features, it fosters a selective positive transfer mechanism, allowing the model to leverage commonalities between Speech and Sound while insulating the distinct attributes of Music.

## 6 Conclusion

We introduce the MoE Adapter to resolve gradient conflicts in LALMs caused by audio heterogeneity. By leveraging dynamic expert specialization, our sparse architecture consistently outperforms dense baselines across diverse benchmarks while maintaining comparable latency. Our analysis confirms that this mechanism transforms destructive gradient interference into constructive updates, underscoring the pivotal role of sparsity as a scalable pathway for future multimodal foundation models.

## 7 Limitations

Although the proposed MoE-Adapter demonstrates superior representational capacity compared to the dense FFN baseline and effectively mitigates gradient conflicts, this work has limitations regarding scalability and generalization. First, our evaluations are currently confined to the Qwen3-1.7B backbone; the universality of our method across different LLM families and larger parameter scales (e.g., 70B) remains to be empirically verified. Second, we have not explicitly explored the scaling laws of the sparse routing mechanism with respect to training data volume, leaving the performance trajectory of expert specialization under massive-scale training for future investigation. Finally, we focus on audio understanding and reasoning tasks, and have not yet extended the framework to generative audio applications.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Yiming Chen, Xianghu Yue, Chen Zhang, Xiaoxue Gao, Robby T Tan, and Haizhou Li. 2024. Voicebench: Benchmarking llm-based voice assistants. *arXiv preprint arXiv:2410.17196*.
- Zeren Chen, Ziqin Wang, Zhen Wang, Huayang Liu, Zhenfei Yin, Si Liu, Lu Sheng, Wanli Ouyang, Yu Qiao, and Jing Shao. 2023. Octavius: Mitigating task interference in mllms via lora-moe. *arXiv preprint arXiv:2311.02684*.
- Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. 2018. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International conference on machine learning*, pages 794–803. PMLR.
- 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*.
- 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*.
- Qingkai Fang, Shoutao Guo, Yan Zhou, Zhengrui Ma, Shaolei Zhang, and Yang Feng. 2024. Llama-omni: Seamless speech interaction with large language models. *arXiv preprint arXiv:2409.06666*.
- William Fedus, Barret Zoph, and Noam Shazeer. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39.
- Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 776–780. IEEE.
- Arushi Goel, Sreyan Ghosh, Jaehyeon Kim, Sonal Kumar, Zhifeng Kong, Sang-gil Lee, Chao-Han Huck Yang, Ramani Duraiswami, Dinesh Manocha, Rafael Valle, and 1 others. 2025. Audio flamingo 3: Advancing audio intelligence with fully open large audio language models. *arXiv preprint arXiv:2507.08128*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, and 1 others. 2025. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638.
- Rongjie Huang, Mingze Li, Dongchao Yang, Jiaotong 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.
- Runhan Huang, Shaoting Zhu, Yilun Du, and Hang Zhao. 2025. Moe-loco: Mixture of experts for multi-task locomotion. *arXiv preprint arXiv:2503.08564*.
- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. 1991. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87.
- Shengpeng Ji, Yifu Chen, Minghui Fang, Jialong Zuo, Jingyu Lu, Hanting Wang, Ziyue Jiang, Long Zhou, Shujie Liu, Xize Cheng, and 1 others. 2024. Wavchat: A survey of spoken dialogue models. *arXiv preprint arXiv:2411.13577*.
- Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2021. Gshard: Scaling giant models with conditional computation and automatic sharding. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Chengao Li, Hanyu Zhang, Yunkun Xu, Hongyan Xue, Xiang Ao, and Qing He. 2025a. Gradient-adaptive policy optimization: Towards multi-objective alignment of large language models. *arXiv preprint arXiv:2507.01915*.
- Yunxin Li, Shenyuan Jiang, Baotian Hu, Longyue Wang, Wanqi Zhong, Wenhan Luo, Lin Ma, and Min Zhang. 2025b. Uni-moe: Scaling unified multimodal llms

- with mixture of experts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Bo Liu, Xingchao Liu, Xiaojie Jin, Peter Stone, and Qiang Liu. 2021. Conflict-averse gradient descent for multi-task learning. *Advances in Neural Information Processing Systems*, 34:18878–18890.
- Zhenyu Liu, Yunxin Li, Xuanyu Zhang, Qixun Teng, Shenyuan Jiang, Xinyu Chen, Haoyuan Shi, Jinchao Li, Qi Wang, Haolan Chen, and 1 others. 2025. Unimoe-audio: Unified speech and music generation with dynamic-capacity moe. *arXiv preprint arXiv:2510.13344*.
- Basil Mustafa, Carlos Riquelme, Joan Puigcerver, Rodolphe Jenatton, and Neil Houlsby. 2022. Multi-modal contrastive learning with limoe: the language-image mixture of experts. *Advances in Neural Information Processing Systems*, 35:9564–9576.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust speech recognition via large-scale weak supervision](#). *arXiv preprint*.
- Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. 2021. Scaling vision with sparse mixture of experts. *Advances in Neural Information Processing Systems*, 34:8583–8595.
- 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*.
- S Sakshi, Utkarsh Tyagi, Sonal Kumar, Ashish Seth, Ramaneswaran Selvakumar, Oriol Nieto, Ramani Duraiswami, Sreyan Ghosh, and Dinesh Manocha. 2024. Mmau: A massive multi-task audio understanding and reasoning benchmark. *arXiv preprint arXiv:2410.19168*.
- Yongqi Shao, Bingxin Mei, Cong Tan, Hong Huo, and Tao Fang. 2025. Motas: Moe-guided feature selection from tts-augmented speech for enhanced multimodal alzheimer’s early screening. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 8497–8505.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Fangxun Shu, Lei Zhang, Hao Jiang, and Cihang Xie. 2025. Audio-visual llm for video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4246–4255.
- Yiyang Su, Yunping Shi, Feng Liu, and Xiaoming Liu. 2025. Hamobe: Hierarchical and adaptive mixture of biometric experts for video-based person reid. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11525–11536.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. 2023. Salmonn: Towards generic hearing abilities for large language models. *arXiv preprint arXiv:2310.13289*.
- Meituan LongCat Team, Bairui Wang, Bin Xiao, Bo Zhang, Bolin Rong, Borun Chen, Chang Wan, Chao Zhang, Chen Huang, Chen Chen, and 1 others. 2025. Longcat-flash-omni technical report. *arXiv preprint arXiv:2511.00279*.
- Haifeng Wang, Hua Wu, Tian Wu, Yu Sun, Jing Liu, Dianhai Yu, Yanjun Ma, Jingzhou He, Zhongjun He, Dou Hong, and 1 others. 2026. Ernie 5.0 technical report. *arXiv preprint arXiv:2602.04705*.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024. [Mmlu-pro: A more robust and challenging multi-task language understanding benchmark](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Kaiyue Wen, Zhiyuan Li, Jason Wang, David Hall, Percy Liang, and Tengyu Ma. 2024. Understanding warmup-stable-decay learning rates: A river valley loss landscape perspective. *arXiv preprint arXiv:2410.05192*.
- Boyong Wu, Chao Yan, Chen Hu, Cheng Yi, Chengli Feng, Fei Tian, Feiyu Shen, Gang Yu, Haoyang Zhang, Jingbei Li, and 1 others. 2025. Step-audio 2 technical report. *arXiv preprint arXiv:2507.16632*.
- Zhifei Xie and Changqiao Wu. 2024. Mini-omni: Language models can hear, talk while thinking in streaming. *arXiv preprint arXiv:2408.16725*.
- Jiayi Xin, Sukwon Yun, Jie Peng, Inyoung Choi, Jenna L Ballard, Tianlong Chen, and Qi Long. 2025. I2moe: Interpretable multimodal interaction-aware mixture-of-experts. *arXiv preprint arXiv:2505.19190*.
- 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. *arXiv preprint arXiv:2509.17765*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Shu-wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhota, Yist Y Lin, Andy T Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, and 1 others. 2021. Superb: Speech processing universal performance benchmark. *arXiv preprint arXiv:2105.01051*.

Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. Gradient surgery for multi-task learning. *Advances in neural information processing systems*, 33:5824–5836.

Aohan Zeng, Zhengxiao Du, Mingdao Liu, Kedong Wang, Shengmin Jiang, Lei Zhao, Yuxiao Dong, and Jie Tang. 2024. Glm-4-voice: Towards intelligent and human-like end-to-end spoken chatbot. *arXiv preprint arXiv:2412.02612*.

Dan Zhang, Yishu Lei, Jing Hu, Shuwei He, Songhe Deng, Xianlong Luo, Danxiang Zhu, Shikun Feng, Rui Liu, Jingzhou He, and 1 others. 2026. Eureka-audio: Triggering audio intelligence in compact language models. *arXiv preprint arXiv:2602.13954*.

Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. 2023. Speechgpt: Empowering large language models with intrinsic cross-modal conversational abilities. *arXiv preprint arXiv:2305.11000*.

Jihai Zhang, Xiaoye Qu, Tong Zhu, and Yu Cheng. 2025a. Clip-moe: Towards building mixture of experts for clip with diversified multipler upcycling. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 5406–5419.

Zongmeng Zhang, Wengang Zhou, Jie Zhao, and Houqiang Li. 2025b. Robust multimodal large language models against modality conflict. *arXiv preprint arXiv:2507.07151*.

## A Hyperparameter Configuration

To ensure a rigorous evaluation of the architectural benefits of our proposed method, we conduct a controlled comparison between the **Dense FFN Baseline** and the **MoE Adapter**. We explicitly align the total parameter budget of both models to approximately 94M to isolate the gains attributed to sparse expert specialization. The detailed specifications are summarized in Table 4.

**Common Audio Frontend.** Both models utilize an identical frontend configuration to process acoustic signals. The continuous representations extracted by the speech encoder ( $d_{enc} = 5, 120$ ) are projected into the model’s audio hidden space ( $d_{model} = 2, 560$ ) via a shared *Feature Fusion* module. This module consists of a Layer Normalization (LN) layer, a linear projection ( $5, 120 \rightarrow 2, 560$ ),

Hyperparameter	Specification
<b>I. Common Audio Frontend</b>	
Speech Encoder Dim	5, 120
Audio Hidden Size	$d = 2, 560$
Feature Fusion	LN $\rightarrow$ Linear <sub>5120<math>\rightarrow</math>d</sub> $\rightarrow$ SiLU $\rightarrow$ LN
<b>II. Baseline: Dense FFN Adapter</b>	
Architecture	Monolithic MLP
Projection Structure	LN $\rightarrow$ MLP <sub>d<math>\rightarrow</math>20480<math>\rightarrow</math>d</sub>
Total Parameters	<b>94.4 M</b> (100% Active)
<b>III. Ours: MoE Adapter</b>	
Architecture	Sparse Mixture-of-Experts
Expert Config	$N = 8$ (Total), $k = 4$ (Active)
Expert Structure	LN $\rightarrow$ Experts <sub>d<math>\rightarrow</math>1280<math>\rightarrow</math>d</sub> $\xrightarrow{\text{Sum}}$ Out
Gate Network	Linear <sub>d<math>\rightarrow</math>8</sub>
Aggregation Block	LN $\rightarrow$ MLP <sub>d<math>\rightarrow</math>10240<math>\rightarrow</math>d</sub>
<b>Active Parameters</b>	<b>70.8 M</b> ( $\sim 75\%$ of Baseline)

Table 4: **Detailed Hyperparameter Configuration.** Both architectures are constrained to the same parameter budget ( $\approx 94\text{M}$ ). Note that  $d = 2560$  denotes the audio hidden size.

and a SiLU activation function, followed by a final LN layer.

**Adapter Architecture Specifics.** The core divergence lies in the parameter allocation strategy within the projection block:

- **Dense FFN Baseline:** This model employs a monolithic Multi-Layer Perceptron (MLP). To utilize the full parameter budget, we scale the intermediate expansion dimension to 20, 480. Consequently, all 94.4M parameters are active for every input token.
- **MoE Adapter (Ours):** We decompose the dense projection into  $N = 8$  specialized experts, each with a smaller intermediate dimension of 1, 280. A learnable gating network routes tokens to the Top- $k = 4$  most relevant experts. Distinctively, we introduce a post-routing **Aggregation Block**—an MLP with an intermediate dimension of 10, 240—to align the aggregated expert outputs with the LLM’s embedding space.

**Computational Efficiency.** As demonstrated in the statistics section of Table 4, although both architectures maintain a consistent storage footprint ( $\sim 94.5\text{M}$ ), the MoE Adapter significantly reduces computational overhead. During inference, only

the routed experts are activated, resulting in an effective parameter count of **70.8 M**. This represents a  $\sim 25\%$  reduction in active parameters compared to the dense baseline, achieving a superior trade-off between model capacity and inference latency.