

Generalizable Prompt Tuning for Audio-Language Models via Semantic Expansion

Jaehyuk Jang*

Wonjun Lee*

Kangwook Ko*

Changick Kim

KAIST

{jhyuk, dpenguin, kw.ko, changick}@kaist.ac.kr

Abstract

Prompt tuning has achieved remarkable progress in vision–language models (VLMs) and is recently being adopted for audio–language models (ALMs). However, its generalization ability in ALMs remains largely underexplored. We observe that conventional prompt tuning for ALMs also suffers from the Base–New Tradeoff, and we identify that this issue stems from the disrupted semantic structure of the embedding space. To address this issue, we propose **Semantically Expanded Prompt Tuning (SEPT)**—a plug-and-play framework that explicitly regularizes the prompt embedding space by incorporating semantic neighbors generated by large language models. SEPT introduces a novel semantic expansion loss with margin constraints that promote intra-class compactness and inter-class separability, thereby enhancing the semantic structure of the prompt embedding space. For comprehensive evaluation, we establish the first benchmark setup for prompt generalization in ALMs, covering both base-to-new generalization and cross-dataset transferability. Extensive experiments demonstrate that SEPT consistently improves generalization performance across multiple prompt tuning baselines, while maintaining computational cost during inference.

1 Introduction

Prompt tuning has emerged as an effective technique for adapting pretrained vision–language models (VLMs) such as CLIP (Radford et al., 2021) to downstream tasks. Rather than relying on hand-crafted templates, it learns continuous prompt embeddings optimized for the task. CoOp (Zhou et al., 2022b) first showed that optimizing prompts on a small set of base classes can significantly improve performance. However, conventional prompt tuning often overfits to base (seen) classes, resulting in

*Equal contribution

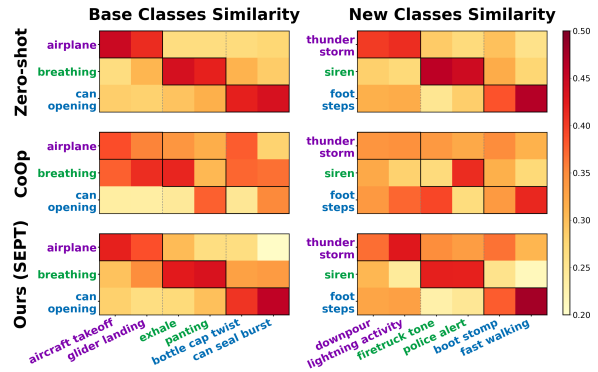


Figure 1: Motivation. The y -axis corresponds to class names, and x -axis lists semantic neighbors for each class. Hand-crafted zero-shot prompts show strong similarity between each class and its neighbors. In contrast, conventional prompt tuning disrupts this alignment. SEPT preserves this semantic similarity, keeping strong alignment between each class and its neighbors.

poor generalization to new (unseen) ones—a phenomenon known as the Base–New Tradeoff (BNT). BNT has become a central challenge in prompt tuning, motivating a wide range of solutions, including architectural enhancements (Zhou et al., 2022a; Zhang et al., 2024; Khattak et al., 2023; Gao et al., 2024) and regularization strategies (Yao et al., 2023; Zhu et al., 2023). These efforts underscore the importance of learning prompts that generalize well across categories and domains.

The success of prompt tuning in VLMs has recently motivated its extension to audio–language models (ALMs) (Elizalde et al., 2023, 2024; Wu et al., 2023; Deshmukh et al., 2023). Audio-free prompt tuning methods (Li et al., 2024; Shi et al., 2025) have been proposed for sound event classification and emotion recognition. These approaches leverage task-specific textual descriptions of audio instead of raw audio inputs, aiming to improve generalization to unseen classes. Meanwhile, PALM (Hanif et al., 2024) introduced a few-shot

prompt tuning method for ALMs, showing strong performance across a range of audio classification tasks, including instrument, acoustic scene, and vocal sound classification. Despite these advances, important generalization challenges—such as base-to-new generalization and cross-dataset transferability—remain insufficiently addressed in the context of prompt tuning for ALMs.

The generalization ability of VLMs and ALMs comes from a semantically well-structured text embedding space (Frome et al., 2013; Jia et al., 2021; Radford et al., 2021). When text embeddings of semantically related classes, such as words with similar meanings or usage contexts, are organized as coherent clusters, the model is better equipped to generalize to unseen classes. We conduct a preliminary study to examine how prompt tuning in ALMs affects the structure of the text embedding space. As shown in Fig. 1, hand-crafted zero-shot prompts (“This is a sound of {class}”) naturally preserve semantic structure: classes (y-axis) maintain strong similarity with their semantically related neighbors (x-axis), (e.g., synonyms or paraphrases), as indicated by the high similarity scores for both base and new classes. In contrast, CoOp (Zhou et al., 2022b) tends to disrupt this desirable structure when applied to ALMs. Due to the overfitting of learned prompt embeddings only with seen class names, the similarity between a class and its semantic neighbor weakens, indicating that the semantic relationship is not adequately preserved. We attribute the severity of this vulnerability in ALMs to the inherent semantic sparsity of audio benchmarks (Salamon et al., 2014; Gong et al., 2022; Cao et al., 2014), which are constrained to only tens of classes. In such a sparse semantic space, the learned prompts lack sufficient support to maintain geometric cohesion, degenerating into isolated prototypes rather than robust semantic anchors (Harun et al., 2024; Wu et al., 2018).

To address the limitations of prompt overfitting and poor semantic alignment, we propose **Semantically Expanded Prompt Tuning (SEPT)**—a novel and plug-and-play prompt tuning framework that explicitly regularizes the structure of the text embedding space by leveraging semantic neighbors. Our approach utilizes the large language models to enrich the semantic coverage of each class. For every seen class, we generate a diverse set of semantically related neighbor classes, and these neighbors are incorporated into the prompt tuning process to encourage each class and its se-

mantic variants to form a coherent cluster in the shared text embedding space. We explicitly enforce a robust geometric structure within this enriched space by introducing a semantic expansion loss that encourages the embedding of each class to be close to its own semantic neighbors (i.e., positive neighbors), while maintaining sufficient distance from the neighbors of all other classes (i.e., negative neighbors). This pull–push mechanism promotes intra-class compactness and inter-class separability in the embedding space, thereby encouraging a semantically coherent structure that is critical for generalization to unseen classes. In addition, to avoid over-compression of positives and excessive repulsion of negatives, we propose margin constraints for both intra-class and inter-class relationships.

Our proposed SEPT is model-agnostic and plug-and-play, and can be seamlessly integrated into a wide range of existing prompt tuning approaches. To demonstrate its versatility, we apply our method to several representative baselines. To the best of our knowledge, we are the first to establish a comprehensive generalization evaluation setup for ALMs and to rigorously evaluate prompt generalization across diverse datasets. Through comprehensive experiments on multiple audio classification benchmarks, we demonstrate that our method consistently improves the performance of all these baselines, both base-to-new generalization and cross-dataset transferability.

2 Related Work

Prompt Tuning in Vision-Language Models.

Prompt tuning has emerged as a promising technique for vision-language models (VLMs), with the introduction of CLIP (Radford et al., 2021), which aligns vision and language modalities via contrastive learning. To adapt CLIP to downstream tasks with minimal resources, CoOp (Zhou et al., 2022b) proposed optimizing a set of learnable context tokens instead of relying on handcrafted textual prompts while keeping encoders frozen. However, the learned context tokens tend to overfit the classes seen during training, which weakens the general knowledge of CLIP and leads to a significant performance drop on unseen data, which is also known as Base-New Tradeoff (BNT).

To alleviate this issue, CoCoOp (Zhou et al., 2022a) introduced an input-conditional context token by leveraging an additional neural network called Meta-Net. Extending prompt tuning beyond

the language branch, MaPLe (Khattak et al., 2023) jointly learns prompts for both vision and language encoders, with coupled prompt representations that explicitly model cross-modal interactions and improve generalization. KgCoOp (Yao et al., 2023) addresses the problem by adding a regularization term that minimizes the Euclidean distance between learned prompt embeddings and handcrafted, knowledge-rich general prompt embeddings. Similarly, LASP (Bulat and Tzimiropoulos, 2023) also mitigates base-class overfitting by text-to-text objective that encourages soft prompts to remain close to handcrafted textual prompts. DePT (Zhang et al., 2024) analyzes the BNT as a channel bias problem, proposing a decoupling strategy that isolates base-specific knowledge to preserve task-shared knowledge, significantly enhancing generalization to unseen tasks. More recently, DPC (Li et al., 2025) further alleviates BNT by using two collaborating prompts, and TAC (Hao et al., 2025) improves generalization by injecting task-aware tokens derived from class clustering. Collectively, these works demonstrate the efficacy of prompt tuning for few-shot generalization in VLMs.

Prompt Tuning in Audio-Language Models. The success of prompt tuning for VLMs has recently inspired its application in audio-language models (ALMs). Test-time domain adaptation for ALMs (Deshmukh et al., 2024) adapts learnable prompts by learning domain representations directly from test audio, requiring neither extra annotations nor access to training data. PT-Text (Li et al., 2024) and CLEP-DG (Shi et al., 2025) propose audio-free prompt tuning strategies, where learnable prompts are trained using task-specific audio descriptions instead of raw audio, thereby improving generalization beyond seen classes. PALM (Hanif et al., 2024) explores few-shot prompt tuning by adapting learnable prompts in feature space, enhancing performance under limited supervision. However, the generalization problem, which is a critical research focus of prompt tuning for VLMs, remains underexplored in the context of ALMs across diverse audio classification datasets.

3 Method

We present Semantically Expanded Prompt Tuning (SEPT), a novel and plug-and-play framework for prompt tuning that enhances the structure of the text embedding space by incorporating semantic neighbors. Our method explicitly regularizes the

structure of the text embedding space to form semantically meaningful clusters, as shown in Fig. 2.

3.1 Preliminaries

Audio-Language Models (ALMs) (Elizalde et al., 2023, 2024; Wu et al., 2023; Deshmukh et al., 2023) align audio and text modalities via contrastive learning on large-scale paired datasets. By optimizing contrastive loss, ALMs bring matching pairs closer in a shared embedding space while separating unrelated ones, effectively bridging the modality gap.

Zero-shot audio classification uses the pre-trained ALM without any additional fine-tuning on the target task. Let \mathbf{x} be an audio embedding obtained from the audio encoder, and $\{\mathbf{w}_i\}_{i=1}^K$ be text embeddings generated from class-descriptive prompts (e.g., “This is a sound of {class}”) by the text encoder, using the class $\{c_i\}_{i=1}^K$ as class name. The prediction probability is defined as:

$$p(y|\mathbf{x}) = \frac{\exp(\text{sim}(\mathbf{x}, \mathbf{w}_y)/\tau)}{\sum_{i=1}^K \exp(\text{sim}(\mathbf{x}, \mathbf{w}_i)/\tau)}, \quad (1)$$

where $\text{sim}(\cdot, \cdot)$ is the cosine similarity and τ is a temperature parameter.

Prompt-based tuning replaces fixed hand-crafted templates with learnable prompts—continuous context vectors that are optimized alongside the downstream task objective. Specifically, for each class name c_i with token embedding function \mathcal{E} , M learnable context vectors $[\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_M]$ are concatenated with token embedding $\mathcal{E}(c_i)$ to form the prompt: $\mathbf{t}_i = [\mathbf{v}_1, \dots, \mathbf{v}_M, \mathcal{E}(c_i)]$. These prompts are then passed to the frozen text encoder $g(\cdot)$ to produce the class prompt embedding. Given an audio input \mathbf{x} and class prompt embeddings $g(\mathbf{t}_i)$, the prediction is made using:

$$p(y|\mathbf{x}) = \frac{\exp(\text{sim}(\mathbf{x}, g(\mathbf{t}_y))/\tau)}{\sum_{i=1}^K \exp(\text{sim}(\mathbf{x}, g(\mathbf{t}_i))/\tau)}, \quad (2)$$

where y is the corresponding label. The prompt vectors are updated using a cross-entropy loss, while the audio and text encoders are kept frozen, thus maintaining the generalization capacity of the pre-trained ALM. The cross-entropy loss is given by:

$$\mathcal{L}_{ce} = - \sum_{\mathbf{x} \in \mathbf{X}} \log p(y|\mathbf{x}). \quad (3)$$

3.2 Semantic Neighbors

Unlike image classification benchmarks with thousands of categories (Deng et al., 2009), most audio datasets contain only tens of classes (Engel

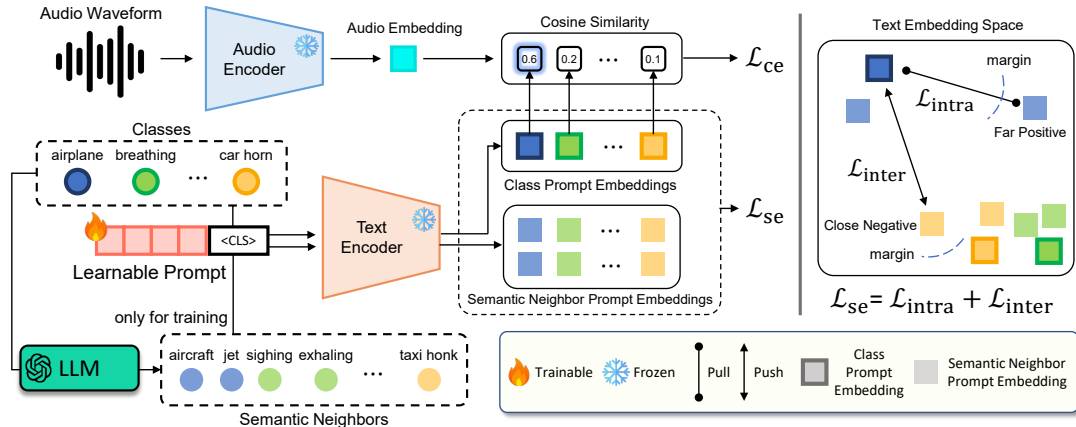


Figure 2: Framework of proposed Semantically Expanded Prompt Tuning (SEPT), applied to CoOp (Zhou et al., 2022b). During training, we use LLM to generate Semantic Neighbors—words that are semantically related to each class name. These neighbors are encoded into semantic neighbor embeddings using the same learnable prompt and frozen text encoder. To preserve the semantic structure of the original text embedding space, we introduce a semantic expansion loss \mathcal{L}_{se} . Specifically, as illustrated in the figure, \mathcal{L}_{intra} pulls class embeddings closer to far semantic neighbors, and \mathcal{L}_{inter} pushes away from overly close unrelated ones. Note that our method can be integrated into other prompt tuning methods and does not affect inference efficiency.

et al., 2017), providing far narrower semantic coverage. This scarcity amplifies the overfitting issue observed in CoOp, causing class embeddings to become overly isolated and weakening their similarity to semantic neighbors, as illustrated in Fig. 1. Such degradation of semantic structure makes it substantially more difficult for the model to generalize to unseen classes.

To address this, we introduce a semantic expansion strategy that leverages semantic neighbors. Specifically, for each class name c_i , we generate a set of N semantically related neighbors $\{p_i^1, \dots, p_i^N\}$ using a large language model (LLM). These neighbors capture fine-grained acoustic variants and natural language expressions related to c_i , effectively forming a dense semantic cluster around each original class in the text embedding space. We structure the semantic expansion loss such that the prompt embedding for each class is pulled closer to its own neighbors (i.e., positive neighbors), while being pushed away from the neighbors of all other classes (i.e., negative neighbors). This dual mechanism encourages both intra-class compactness and inter-class separability, thereby expanding the robust semantic structure of the embedding space.

3.3 Margin Constraint

While pulling positive neighbors close and pushing negative neighbors away in embedding space can clarify the semantic structure, applying this mechanism naively—forcing all positives to be maxi-

mally close and all negatives maximally far—can paradoxically degrade the embedding space. Over-compressing positives suppresses the natural diversity of language and sound, while excessive separation of negatives can distort legitimate semantic relationships across classes.

To mitigate this, we introduce distance margin constraints $m_{i,j,n}$ that reflect their original semantic distance. To ensure robustness, we compute this margin by averaging statistics across a diverse pool of T prompts $\{\pi_t\}_{t=1}^T$ automatically generated using LLM (e.g., "a sound of", "a recording of"). The margin constraint between the c_i and n -th neighbor of c_j , p_j^n , is defined as the average L_2 distance of their embeddings across all T prompts:

$$m_{i,j,n} = \frac{1}{T} \sum_{t=1}^T \left\| g(\mathcal{E}(\pi_t; c_i)) - g(\mathcal{E}(\pi_t; p_j^n)) \right\|_2. \quad (4)$$

Here, $g(\mathcal{E}(\pi_t; c_i))$ denotes the prompt embedding of the i -th class name c_i under prompt π_t , $g(\mathcal{E}(\pi_t; p_j^n))$ denotes the embedding of the n -th neighbor of class j , p_j^n , under the same prompt, and $\|\cdot\|_2$ represents the Euclidean (L_2) norm.

Building on this margin, we introduce a novel semantic expansion loss that pulls class prompts towards positive neighbors and pushes them away from negative neighbors with a margin constraint.

3.4 Semantic Expansion Loss

We define two losses to regularize the prompt embedding space via semantic expansion: (i) an intra-class alignment loss $\mathcal{L}_{\text{intra}}$, which encourages each class embedding to be close to its positive neighbors, and (ii) an inter-class separation loss $\mathcal{L}_{\text{inter}}$, which pushes it away from negative neighbors.

As defined in the preliminaries, each class prompt embedding for the class c_i is given by

$$\mathbf{z}_i = g(t_i) = g([\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_M, \mathcal{E}(c_i)]), \quad (5)$$

where $\{\mathbf{v}_m\}_{m=1}^M$ are the learnable context vectors and $\mathcal{E}(c_i)$ is the token embedding of the class c_i . We also define $\{\mathbf{p}_i^n\}_{n=1}^N$ as prompt embeddings of semantic neighbors of class c_i , which can be expressed as

$$\mathbf{p}_i^n = g([\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_M, \mathcal{E}(p_i^n)]). \quad (6)$$

To encourage tighter semantic cohesion within each class, we pull the class prompt embedding \mathbf{z}_i closer to its own neighbors \mathbf{p}_i^n in the embedding space, if the distance is larger than the margin. The intra-class alignment loss is computed as a margin-based hinge loss:

$$\mathcal{L}_{\text{intra}}^{(i)} = \frac{1}{N} \sum_{n=1}^N \max(0, \|\mathbf{z}_i - \mathbf{p}_i^n\|_2 - m_{i,i,n}), \quad (7)$$

where $m_{i,i,n}$ is the precomputed intra-class margin between embeddings of the c_i and its neighbor p_i^n as defined in Eq. 4.

Furthermore, to preserve inter-class distinctiveness, we push the class embedding \mathbf{z}_i away from the neighbors $\{\mathbf{p}_j^n\}_{n=1}^N$ of all other classes $j \neq i$, by enforcing a minimum semantic distance via:

$$\mathcal{L}_{\text{inter}}^{(i,j)} = \frac{1}{N} \sum_{n=1}^N \max(0, m_{i,j,n} - \|\mathbf{z}_i - \mathbf{p}_j^n\|_2), \quad (8)$$

where $m_{i,j,n}$ reflects the average margin distance between embeddings of c_i and the n -th neighbor of class j , p_j^n .

The semantic expansion loss across all classes is computed by averaging intra- and inter-class losses:

$$\mathcal{L}_{\text{se}} = \frac{1}{K} \sum_{i=1}^K \left(\mathcal{L}_{\text{intra}}^{(i)} + \frac{1}{K-1} \sum_{j=1, j \neq i}^K \mathcal{L}_{\text{inter}}^{(i,j)} \right). \quad (9)$$

This loss encourages a structured expansion of the text embedding space, where each class forms a

compact cluster of semantically related variants while remaining discriminative from other classes (e.g., “bell” and “chime” remain relatively close, while “explosion” and “chirp” are farther apart).

The total training loss combines the cross-entropy loss \mathcal{L}_{ce} with the semantic expansion loss \mathcal{L}_{se} , enabling the model to learn from both supervised signals and semantic structure regularization:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{ce}} + \lambda \cdot \mathcal{L}_{\text{se}},$$

where λ is a balancing hyperparameter. Note that our \mathcal{L}_{se} can be incorporated in a compatible manner alongside other auxiliary losses designed to enhance prompt tuning.

4 Experiments

4.1 Experimental Setup

Datasets. Following PALM (Hanif et al., 2024), we evaluate our method on a diverse collection of audio classification datasets, covering a wide range of speech and non-speech domains. We used Beijing-Opera (Tian et al., 2014) and NS-Instruments (Engel et al., 2017) for instrument classification, ESC50 (Piczak, 2015), ESC50-Actions, and UrbanSound8K (Salamon et al., 2014) for sound event classification, CREMA-D (Cao et al., 2014) and RAVDESS (Livingstone and Russo, 2018) for emotion recognition, SESA (Spadini, 2019) for surveillance sound classification, TUT2017 (Mesaros et al., 2017) for acoustic scene classification, GT-Music-Genre (Sturm, 2012) for music analysis, and VocalSound (Gong et al., 2022) for vocal sound classification.

Model and Baselines. We utilize the audio and text encoders of the Pengi (Deshmukh et al., 2023) for audio classification tasks following PALM. Although Pengi was originally designed for multimodal generation tasks, we adopt its audio and text encoders for classification, similar to prior work in prompt tuning. To demonstrate the compatibility of SEPT with existing prompt tuning methods, we apply it to several representative textual prompt tuning methods: CoOp (Zhou et al., 2022b), Co-CoOp (Zhou et al., 2022a), KgCoOp (Yao et al., 2023), and DePT (Zhang et al., 2024), all of which are built upon CoOp. Note that SEPT is designed as a model-agnostic and plug-and-play module and can be seamlessly integrated into other ALM or prompt tuning methods as well.

Method	Avg. over 11 datasets			Beijing-Opera			NS-Instruments			ESC50		
	Base	New	H	Base	New	H	Base	New	H	Base	New	H
Pengi	40.68	38.21	38.46	64.16	68.55	66.15	43.32	30.01	35.46	15.70	13.70	14.33
CoOp	65.00	34.09	42.83	97.27	61.38	74.87	41.11	23.88	30.08	61.97	11.83	19.46
+ SEPT	64.36	42.98	49.70	97.88	71.87	82.45	43.33	37.78	40.29	59.00	15.80	24.70
CoCoOp	69.13	36.83	46.26	97.86	70.80	81.84	47.42	37.66	41.84	70.13	13.63	22.78
+ SEPT	68.63	42.59	50.65	97.85	71.06	81.60	52.15	29.96	37.95	69.10	17.33	27.62
KgCoOp	37.99	37.42	36.39	67.26	61.40	62.96	39.80	39.27	39.11	14.10	10.33	11.76
+ SEPT	58.92	45.28	49.79	94.44	67.99	78.28	51.24	40.86	45.28	47.60	18.33	26.33
DePT	63.86	39.91	46.79	97.26	62.25	75.32	43.75	28.21	34.01	60.80	14.00	22.53
+ SEPT	64.57	41.63	49.06	96.32	62.74	75.00	47.66	37.90	41.60	56.80	16.03	24.91

Method	ESC50-Actions			UrbanSound8k			CREMA-D			RAVDESS		
	Base	New	H	Base	New	H	Base	New	H	Base	New	H
Pengi	25.50	28.00	26.44	30.68	39.51	34.14	52.25	32.36	39.96	25.38	32.16	28.37
CoOp	85.33	49.33	61.97	61.86	31.22	40.94	52.50	9.01	15.29	51.01	27.75	35.88
+ SEPT	82.50	52.67	63.90	61.22	36.34	44.95	52.75	42.06	44.17	54.55	33.92	41.39
CoCoOp	89.17	41.50	56.15	64.25	30.77	41.05	56.80	26.07	35.62	56.82	26.87	36.39
+ SEPT	87.50	51.50	64.58	61.06	25.93	36.26	54.06	19.72	28.21	63.55	37.57	46.77
KgCoOp	36.50	30.67	32.88	26.14	26.15	25.40	42.70	55.43	42.48	27.53	27.46	27.42
+ SEPT	67.83	51.50	57.97	56.96	38.77	45.32	53.06	69.53	59.92	45.58	32.31	37.79
DePT	81.33	44.67	57.12	63.21	37.28	46.05	51.56	33.19	33.54	52.65	32.01	39.28
+ SEPT	79.67	52.67	63.02	61.11	36.22	44.79	51.44	35.60	38.93	52.90	29.07	37.49

Method	SESA			GT-Music-Genre			VocalSound			TUT2017		
	Base	New	H	Base	New	H	Base	New	H	Base	New	H
Pengi	75.00	89.19	81.48	35.64	15.15	21.26	59.28	53.01	55.97	20.52	18.65	19.53
CoOp	85.78	74.77	79.33	57.76	20.54	29.76	75.26	45.01	56.00	45.11	20.25	27.53
+ SEPT	78.92	83.78	81.22	56.11	23.23	32.15	74.11	54.14	62.38	47.55	21.14	29.05
CoCoOp	85.78	72.97	78.68	60.73	24.24	32.68	82.47	40.03	53.19	49.01	20.55	28.61
+ SEPT	88.73	86.49	87.42	49.80	46.32	47.99	79.57	61.06	68.84	51.60	21.54	29.94
KgCoOp	68.14	79.28	72.37	27.39	15.82	19.77	52.50	49.91	51.10	15.83	15.90	15.10
+ SEPT	79.90	82.88	81.23	42.90	20.88	28.08	70.12	56.05	62.23	38.48	18.94	25.30
DePT	80.39	91.89	85.72	50.17	22.90	31.08	72.24	49.00	58.28	49.12	23.67	31.75
+ SEPT	87.25	82.88	84.90	55.78	23.91	33.13	75.34	57.70	65.24	46.05	23.25	30.70

Table 1: Comparison on the base-to-new generalization. Prompts are learned using 16-shot samples per class. 'Base' reports accuracy on seen classes, 'New' on unseen classes, and 'H' is the harmonic mean.

Implementation Details. For training, we randomly sample 16 instances per class from the training set to learn generalizable prompts, while keeping both the audio and text encoders frozen. The number of semantic neighbors per class N is set to 10, and the number of hand-crafted prompts used for margin computation T is 100. We employ ChatGPT-4o as the primary LLM in all experiments. The balancing hyperparameter λ is set to 3 by default. All results are reported as the average over three random seeds. Additional experimental details, including the prompts for generating semantic anchors and the hand-crafted prompt pool constructed by LLM, are provided in the Ap-

pendix A.2.

4.2 Comparative Study

Base-to-New Generalization. The base-to-new generalization setting evaluates a model's ability to transfer knowledge from base classes to previously unseen classes, reflecting a category shift between training and evaluation. Following the previous works on visual-language prompt tuning (Zhou et al., 2022a; Yao et al., 2023), we divide the classes of each dataset equally into two splits, designated as base and new classes. To evaluate the generalization ability of the prompt tuning methods, all the compared methods are trained using only the base

Method	Sound Event Classif.		Emotion Recog.		Instrument Classif.	
	ESC50-A. (Source)	UrbanS. (Target)	RAV. (Source)	CREM. (Target)	NS-Inst. (Source)	Beijing. (Target)
CoOp	70.08	19.02	30.68	15.74	35.30	31.92
+ SEPT	69.42	24.21	31.70	23.89	37.21	41.66
CoCoOp	77.83	16.43	33.81	8.98	37.45	36.44
+ SEPT	78.17	23.18	31.50	19.14	39.63	37.99
KgCoOp	15.92	13.34	14.19	14.08	16.55	29.96
+ SEPT	54.50	21.84	25.05	20.60	39.16	50.01
DePT	70.83	21.86	30.96	17.53	36.89	37.72
+ SEPT	68.08	24.46	27.83	27.07	36.30	44.49

Table 2: Comparison on cross-dataset evaluation. Prompts are trained only on the source dataset with 16-shot samples per class, and then directly applied to both the source and a different target dataset. Reported values are accuracy (%).

classes. During inference, we measure the accuracy separately for the base and new class splits, and report the harmonic mean (H) of these two accuracies as the primary evaluation metric.

Table 1 presents a detailed comparison across 11 diverse audio classification benchmarks. We evaluate four representative prompt tuning—CoOp, CoCoOp, KgCoOp, and DePT—both with and without our SEPT. Across all datasets and methods, incorporating our semantic expansion strategy markedly improves the harmonic mean and new class accuracy, while maintaining competitive performance on base classes. In particular, CoOp and CoCoOp without our approach achieve high accuracy on base classes, but suffer from significant drops on new classes, reflecting the well-known base-to-new trade-off. By integrating SEPT, new class performance increases substantially (e.g., CoOp: 34.09% \rightarrow 42.98%, CoCoOp: 36.83% \rightarrow 42.59% on average), resulting in clear gains in the harmonic mean. Collectively, results demonstrate that our framework can serve as a compatible module for previous prompt tuning approaches, consistently improving base-to-new generalization and mitigating the overfitting in base classes commonly observed in conventional prompt optimization for audio classification tasks. Note that the comparatively lower performance of KgCoOp in ALM, in contrast to VLM, is analyzed in detail in Appendix B.3.

Cross-Dataset Evaluation. To verify transferability beyond category shift and address dataset shift, we conducted cross-dataset evaluations. In this setup, prompts are trained on a source dataset and deployed for inference on a different target dataset without additional fine-tuning.

	$\mathcal{L}_{\text{intra}}$	m^{intra}	$\mathcal{L}_{\text{inter}}$	m^{inter}	Base	New	H
1					65.00	34.09	42.83
2	✓				58.31	33.58	41.17
3	✓	✓			64.99	36.90	44.52
4			✓		62.01	36.68	43.48
5			✓	✓	62.84	38.79	46.04
6	✓		✓		61.82	32.89	41.47
7	✓	✓	✓	✓	64.37	42.93	49.70

Table 3: Ablation of semantic expansion and margins. m^{intra} and m^{inter} indicate whether a margin constraint is applied to $\mathcal{L}_{\text{intra}}$ and $\mathcal{L}_{\text{inter}}$, respectively.

As summarized in Table 2, SEPT substantially improves cross-dataset generalization across diverse audio classification tasks, including sound event classification, emotion recognition, and instrument recognition. Baseline methods such as CoOp, CoCoOp, and DePT achieve decent performance on their source datasets but suffer notable accuracy drops when transferred to unseen target datasets. In contrast, incorporating our semantic expansion strategy yields significant gains in transfer accuracy while maintaining competitive source-domain performance. These results demonstrate that our semantic expansion loss serves as a versatile, plug-and-play regularizer that enhances the robustness of the prompt space, which is crucial for successful cross-dataset transfer in ALMs.

4.3 Analysis

Semantic Expansion with Margin Constraints.

Table 3 presents an ablation study on the semantic expansion components of SEPT built on CoOp, specifically the intra- and inter-class losses ($\mathcal{L}_{\text{intra}}$, $\mathcal{L}_{\text{inter}}$) and their margin constraints (m^{intra} , m^{inter}). m^{intra} and m^{inter} indicate whether margin constraints are applied to $\mathcal{L}_{\text{intra}}$ in Eq. 7 and $\mathcal{L}_{\text{inter}}$ in Eq. 8, respectively. The first row reports baseline CoOp performance without semantic expansion. When applying naive expansion losses without margin constraints (rows 2, 4, 6), performance does not improve over the baseline and often deteriorates on both base and new classes. This shows that indiscriminately pulling prompts towards positives or pushing them away from negatives—without considering their semantic relationships—disrupts the embedding space, leading to over-compression or fragmentation of class clusters. In contrast, introducing margin constraints (rows 3, 5, 7) substantially mitigates these issues. In particular, the full model with both intra- and inter-class margins (row 7) improves the harmonic mean with a large margin, from 42.83 to 49.70, demonstrating a clear

	# Params	Training Time (s)	Inference Time (ms)	H
CoCoOp	107,072	681.5	40.9	46.26
KgCoOp	8,192	96.0	2.8	36.39
DePT	114,738	108.1	2.8	46.79
CoOp	8,192	94.0	2.8	42.83
+ SEPT	8,192	170.7	2.8	49.70

Table 4: Comparison of computational costs. # Params means the number of trainable parameters. The training time is calculated for 50 epochs, while the inference time is measured as the time taken per sample.




Figure 3: Ablation studies of the hyperparameter λ (left) and the number of semantic neighbors N (right).

gain in generalization. These results show that margin constraints play a critical role in structuring the embedding space by controlling intra-class compactness and inter-class distinctiveness.

Complexity. We evaluate the computational efficiency of SEPT by comparing the number of trainable parameters, training/inference time against baseline methods on the ESC50 dataset (Table 4). Crucially, SEPT operates as a plug-and-play module that introduces zero additional parameters and no inference latency overhead, maintaining the exact model size and high inference speed of the CoOp. While the incorporation of semantic expansion incurs an increase in training time compared to CoOp, it remains significantly more efficient than architecture-heavy methods such as CoCoOp. Ultimately, SEPT achieves a superior trade-off, delivering the highest generalization performance without compromising inference efficiency.

Hyperparameter λ and N . Figure 3 presents the ablation studies of two hyperparameters: the balancing parameter λ and the number of semantic neighbors N . Optimal performance is achieved at $\lambda = 3$, which effectively balances supervised classification and structural regularization, yielding the highest harmonic mean. Moreover, SEPT consistently outperforms the baseline across a wide range of λ values, demonstrating its robustness to hyperparameter variations. Meanwhile, increasing the number of semantic neighbors N tends to improve




Figure 4: t-SNE visualization of prompt embeddings of new classes and their semantic neighbors from ESC50-Actions. \diamond denotes the prompt embedding of a new class, while \circ represents its corresponding semantic neighbors.

Method	Avg over 11 datasets		
	Base	New	H
CoOp	65.00	34.09	42.83
+ SEPT (Qwen3 Max)	64.70	40.19	47.71
+ SEPT (Gemini 2.5 Pro)	64.61	41.26	48.32
+ SEPT (ChatGPT-4o)	64.36	42.98	49.70

Table 5: Comparison of different LLMs for neighbor generation when applying SEPT to CoOp.

the harmonic mean. This indicates that incorporating a more diverse set of semantic neighbors provides richer semantic information, allowing the prompt embedding space to be better structured, boosting generalization performance.

Qualitative Analysis. To qualitatively assess whether semantic expansion with base class leads to better structured representations for new (unseen) classes, we visualize the learned text embeddings for new classes and their semantic neighbors using t-SNE, as shown in Fig. 4. The scatterplots compare CoOp with and without our proposed SEPT framework on the ESC50-Actions. Without SEPT (left), the text embedding space lacks structure—semantic neighbors are scattered and do not form coherent clusters. In contrast, with SEPT (right), each new class forms a compact cluster with its corresponding semantic neighbors, illustrating improved intra-class alignment and inter-class separation. This result suggests that the semantic structure induced during training transfers effectively to unseen classes. Additional qualitative examples and neighbor lists are provided in Appendix B.4.

4.4 Robustness Across Different LLMs.

We verify the robustness of our approach by employing various LLMs for neighbor generation. As shown in Table 5, SEPT consistently achieves substantial gains over the CoOp baseline, regardless

of the underlying LLM backbone. Notably, even with Qwen3 Max (Yang et al., 2025) and Gemini 2.5 Pro (Comanici et al., 2025), our method yields impressive H-scores of 47.71 and 48.32, respectively. This demonstrates that the effectiveness of SEPT stems from the proposed semantic expansion strategy itself, rather than being dependent on a specific language model.

5 Conclusion

In this paper, we identify that prompt tuning in Audio-Language Models (ALMs) also suffers from the Base–New Tradeoff and disrupts the semantic structure of the prompt embedding space, leading to poor generalization to unseen classes. To address this, we propose Semantically Expanded Prompt Tuning (SEPT), a simple yet effective framework designed to enhance the generalization capability of prompt tuning in ALMs. SEPT leverages large language models to generate semantically related neighbors for each class, thereby enriching the prompt embedding space with meaningful semantic structure. Concretely, we introduce a semantic expansion loss with margin constraints that jointly promote intra-class compactness and inter-class separability. To rigorously evaluate SEPT, we introduce the first evaluation protocols for generalization in ALM prompt tuning. Extensive experiments demonstrate that SEPT consistently improves existing methods in both base-to-new generalization and cross-dataset transferability. Results highlight the versatility and effectiveness of SEPT as a plug-and-play module, paving the way for more robust and generalizable prompt tuning in ALMs.

6 Limitations

While the proposed SEPT framework achieves significant improvements by explicitly regularizing the textual prompt space, its underlying principle of semantic expansion holds potential for broader multimodal adaptation. Currently, our implementation focuses on text-modality regularization to address the Base–New Trade-off. Given the recent emergence of learnable multimodal prompts as a powerful technique for downstream tasks, a promising avenue for future research is to extend SEPT to the audio prompt domain. Adapting our semantic expansion strategy to audio prompts could foster semantically enriched joint audio–text embeddings, potentially unlocking further generalization capabilities in fully multimodal ALM setups.

Acknowledgments

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2020-II200153, Penetration Security Testing of ML Model Vulnerabilities and Defense).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, and Shyamal et al. Anadkat. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Adrian Bulat and Georgios Tzimiropoulos. 2023. Lasp: Text-to-text optimization for language-aware soft prompting of vision & language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 23232–23241.
- Houwei Cao, David G Cooper, Michael K Keutmann, Ruben C Gur, Ani Nenkova, and Ragini Verma. 2014. Crema-d: Crowd-sourced emotional multimodal actors dataset. *IEEE transactions on affective computing*, 5(4):377–390.
- Ke Chen, Xingjian Du, Bilei Zhu, Zejun Ma, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. 2022. Hts-at: A hierarchical token-semantic audio transformer for sound classification and detection. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 646–650. IEEE.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee.
- 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.
- Soham Deshmukh, Rita Singh, and Bhiksha Raj. 2024. [Domain adaptation for contrastive audio-language models](#). In *Interspeech 2024*, pages 1680–1684.
- Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang. 2023. Clap learning audio concepts from natural language supervision. In *IEEE*

- International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Benjamin Elizalde, Soham Deshmukh, and Huaming Wang. 2024. Natural language supervision for general-purpose audio representations. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 336–340. IEEE.
- Jesse Engel, Cinjon Resnick, Adam Roberts, Sander Dieleman, Mohammad Norouzi, Douglas Eck, and Karen Simonyan. 2017. Neural audio synthesis of musical notes with wavenet autoencoders. In *International conference on machine learning*, pages 1068–1077. PMLR.
- Li Fei-Fei, Rob Fergus, and Pietro Perona. 2004. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pages 178–178. IEEE.
- Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. 2013. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems*, 26.
- Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. 2024. Clip-adapter: Better vision-language models with feature adapters. *International Journal of Computer Vision*, 132(2):581–595.
- Yuan Gong, Jin Yu, and James Glass. 2022. Vocal-sound: A dataset for improving human vocal sounds recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 151–155.
- Asif Hanif, Maha Tufail Agro, Mohammad Areeb Qazi, and Hanan Aldarmaki. 2024. **PALM: Few-shot prompt learning for audio language models**. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18527–18536, Miami, Florida, USA. Association for Computational Linguistics.
- Fusheng Hao, Fengxiang He, Fuxiang Wu, Tichao Wang, Chengqun Song, and Jun Cheng. 2025. Task-aware clustering for prompting vision-language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 14745–14755.
- Md Yousuf Harun, Kyungbok Lee, Gianmarco Gallardo, Giri Krishnan, and Christopher Kanan. 2024. What variables affect out-of-distribution generalization in pretrained models? *Advances in Neural Information Processing Systems*, 37:56479–56525.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR.
- Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. 2023. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19113–19122.
- Haoyang Li, Liang Wang, Chao Wang, Jing Jiang, Yan Peng, and Guodong Long. 2025. Dpc: Dual-prompt collaboration for tuning vision-language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 25623–25632.
- Yiming Li, Xiangdong Wang, and Hong Liu. 2024. Audio-free prompt tuning for language-audio models. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 491–495. IEEE.
- Steven R Livingstone and Frank A Russo. 2018. The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english. *PloS one*, 13(5):e0196391.
- Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen. 2017. Tut acoustic scenes 2017, development dataset. (*No Title*).
- Karol J Piczak. 2015. Esc: Dataset for environmental sound classification. In *Proceedings of the 23rd ACM international conference on Multimedia*, pages 1015–1018.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, and Jack et al. Clark. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763.
- Justin Salamon, Christopher Jacoby, and Juan Pablo Bello. 2014. A dataset and taxonomy for urban sound research. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 1041–1044.
- Jiacheng Shi, Yanfu Zhang, and Ye Gao. 2025. Clep-dg: Contrastive learning for speech emotion domain generalization via soft prompt tuning. *arXiv preprint arXiv:2507.04048*.
- Tito Spadini. 2019. **Sound events for surveillance applications**.
- Bob L Sturm. 2012. An analysis of the gtzan music genre dataset. In *Proceedings of the second international ACM workshop on Music information retrieval with user-centered and multimodal strategies*, pages 7–12.
- Mi Tian, Ajay Srinivasamurthy, Mark Sandler, and Xavier Serra. 2014. A study of instrument-wise onset detection in beijing opera percussion ensembles. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2159–2163. IEEE.

- Yusong Wu, Ke Chen, Tianyu Zhang, Yuchen Hui, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. 2023. Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Zhirong Wu, Alexei A Efros, and Stella X Yu. 2018. Improving generalization via scalable neighborhood component analysis. In *Proceedings of the european conference on computer vision (ECCV)*, pages 685–701.
- 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*.
- Hantao Yao, Rui Zhang, and Changsheng Xu. 2023. Visual-language prompt tuning with knowledge-guided context optimization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6757–6767.
- Ji Zhang, Shihan Wu, Lianli Gao, Heng Tao Shen, and Jingkuan Song. 2024. Dept: Decoupled prompt tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12924–12933.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022a. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16816–16825.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022b. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348.
- Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. 2023. Prompt-aligned gradient for prompt tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15659–15669.

Appendix

A Detailed Experimental Settings	12
A.1 Dataset overview.	12
A.2 Implementation Details.	12
A.3 Details of Pengi.	12
B Additional Experiments	14
B.1 Ablation of the number of shots.	14
B.2 Ablation of Margin Computation Strategy.	14
B.3 Analysis of KgCoOp.	15
B.4 Additional Qualitative Results.	16
B.5 Quality of Semantic Neighbors.	16
B.6 Comparison with Alternative Topology-Preserving Mechanisms	17
B.7 Applicability on VLMs.	17
C Additional Analysis	18
C.1 Analysis of embedding space structure collapse in ALMs and VLMs	18
C.2 Neighbor Quality.	18
C.3 Performance Variations per Dataset.	19
D AI Assistant Usage Statement	19

A Detailed Experimental Settings

A.1 Dataset overview.

Table 6 summarizes the information about the datasets used in this study. Following PALM (Hanif et al., 2024), we evaluate prompt tuning methods on various downstream tasks, including instrument classification, sound event classification, emotion recognition, vocal sound classification, surveillance sound classification, acoustic scene classification, and music analysis. For datasets with multiple folds, we conduct experiments on all folds, and the average accuracy across them is reported. Each dataset is split into base and new classes with equal proportions, except for TUT2017 (Mesaros et al., 2017), which consists of 15 classes divided into eight base classes and seven new classes.

A.2 Implementation Details.

All learnable prompts were trained for 50 epochs using a learning rate of 0.0125 and the SGD optimizer with a momentum of 0.9. In all experiments, the length of the learnable context vector M was set to 16, with each vector having a dimensionality of 512. Since no separate validation set was used, performance was evaluated using the model checkpoint with the lowest training loss. All reported base accuracy, new accuracy, and harmonic mean (H) values are averaged over three random runs. For a fair comparison, DePT was implemented without test-time knowledge fusion. Additionally, we excluded the batch normalization layers of the pretrained Pengi audio encoder during our experiments. All experiments were conducted on a machine equipped with an Intel(R) Xeon(R) Silver 4210R CPU @ 2.40GHz and a single NVIDIA RTX 3090 GPU. To compute a robust margin, we generate diverse prefixes using the prompts in Fig. 5. The prompt for generating semantic neighbors for semantic expansion is in Fig. 6.

A.3 Details of Pengi.

We utilize Pengi (Deshmukh et al., 2023), a multimodal model originally designed for audio-conditioned text generation. Structurally, it consists of three branches: an audio encoder, a text encoder, and a causal language model. In line with prior work (Hanif et al., 2024), we exclusively adopt the audio and text encoder branches to extract feature embeddings. The audio branch is built upon the Hierarchical Token-Semantic Audio Transformer (HTSAT) (Chen et al., 2022), adopting the archi-

Task	Dataset	Base Classes	New Classes	# Classes	Split
Instrument Classification	Beijing-Opera	bangu, daluo	naobo, xiaoluo	4	Five Fold
	NS-Instruments	bass, brass, flute, guitar, keyboard	mallet, organ, reed, string, vocal	10	Train-Test
Sound Event Classification	ESC50-Actions	breathing, brushing teeth, clapping, coughing, crying baby	drinking sipping, footsteps, laughing, sneezing, snoring	10	Five Fold
	ESC50	airplane, breathing, brushing teeth, can opening, car horn, cat, chainsaw, chirping birds, church bells, clapping, clock alarm, clock tick, coughing, cow, crackling fire, crickets, crow, crying baby, dog, door wood creaks, door wood knock, drinking sipping, engine, fireworks, footsteps	frog, glass breaking, hand saw, helicopter, hen, insects, keyboard typing, laughing, mouse click, pig, pouring water, rain, rooster, sea waves, sheep, siren, sneezing, snoring, thunderstorm, toilet flush, train, vacuum cleaner, washing machine, water drops, wind	50	Five Fold
	UrbanSound8K	air conditioner, car horn, children playing, dog bark, drilling	engine idling, gun shot, jackhammer, siren, street music	10	Ten Fold
	CREMA-D	anger, disgust, fear	happy, neutral, sad	6	Train-Test
Emotion Recognition	RAVDESS	angry, calm, disgust, fearful	happy, neutral, sad, surprised	8	Train-Test
Vocal Sound Classification	VocalSound	Cough, Laughter, Sigh	Sneeze, Sniff, Throat clearing	6	Train-Test
Surveillance Sound Classification	SESA	casual, explosion	gunshot, siren	4	Train-Test
Acoustic Scene Classification	TUT2017	beach, bus, cafe/restaurant, car, city center, forest path, grocery store, home	library, metro station, office, park, residential area, train, tram	15	Four Fold
Music Analysis	GT-Music-Genre	blues, classical, country, disco, hiphop	jazz, metal, pop, reggae, rock	10	Train-Test

Table 6: Summary of datasets, including task, class split, and number of classes.

I am working on an audio classification task using CLAP (Contrastive Language-Audio Pretraining). In this setup, I embed both the audio signal and a text template like "a recording of" and then compute cosine similarity to classify the audio.

However, I suspect that the default template "a recording of" may not be optimal for all audio types. For example, "a sound of", "the noise made by", or "an audio clip of" could potentially lead to better representations depending on the nature of the sound.

I want to explore a diverse and creative list of candidate text templates that describe sounds, voices, audio events, or audio sources. Please generate at least 100 diverse, natural-sounding prompt templates without duplication that could plausibly describe a sound class, where <class> is the audio class label (e.g., dog barking, thunder, doorbell, car, air conditioner, blues, etc.).

Please vary:

- the descriptive style (e.g., technical, casual, poetic)
- the use of verbs and modifiers (e.g., "produced by," "emitted from," "generated by")

Figure 5: Prompt for diverse prefixes.

For each class name, I would like you to generate a list of 10 semantically similar terms that are specifically relevant in the context of audio, sound, or voice characteristics.

The output must adhere to the following criteria:

1. Each set of suggestions should be highly specific to the corresponding class name, reflecting similarity in terms of acoustic properties, usage contexts, or auditory perception.
2. The suggestions for each class must be semantically distinct from those of the other classes in the list, so that the resulting sets are clearly differentiated and non-overlapping.
3. Within each set, the suggestions should be diverse yet representative of the broader category or concept that unites them, ensuring a well-rounded and meaningful expansion.

Before generating the similar terms, please briefly analyze the given class list to identify any overarching themes or semantic groupings, and use this analysis to ensure the relevance and distinctiveness of the expanded sets.

The output should be returned in the following Python dictionary format:

```
'word1':[suggestion1, suggestion2,...],
```

```
'word2':[suggestion1, suggestion2,...],
```

Please ensure that the suggestions for each word reflect its unique role within the shared conceptual domain, and the suggested words for each provided word must not include any of the other words from the original list.

Figure 6: Prompts for semantic neighbors.

Method	# shots = 2			# shots = 4			# shots = 8			# shots = 16		
	Base	New	H	Base	New	H	Base	New	H	Base	New	H
CoOp	47.95	33.29	37.69	53.98	33.95	39.34	60.62	36.28	43.70	65.00	34.09	42.83
+ SEPT	50.89	37.94	41.83	52.45	37.82	42.26	60.48	40.12	46.27	64.36	42.98	49.70
CoCoOp	50.91	38.10	42.06	54.73	36.55	41.96	62.30	37.94	45.02	69.13	36.83	46.26
+ SEPT	54.18	41.11	45.04	56.40	41.99	46.30	63.38	39.90	46.96	68.63	42.59	50.65
KgCoOp	35.24	37.69	35.22	35.31	38.76	35.75	36.17	38.33	36.08	37.99	37.42	36.39
+ SEPT	50.92	41.49	44.65	54.19	43.81	47.14	56.67	42.66	47.37	58.92	45.28	49.79
DePT	49.58	37.71	40.91	54.91	41.44	45.47	61.18	39.28	45.73	63.86	39.91	46.79
+ SEPT	50.28	41.82	44.15	55.29	41.45	45.98	61.73	41.03	47.56	64.57	41.63	49.06

Table 7: Ablation of the number of shots per class. H denotes the harmonic mean between Base and New classes. Reported values are averaged across 11 datasets.

texture of CLAP (Elizalde et al., 2024). The input processing pipeline involves specific preprocessing steps: raw audio waveforms are first resampled to 44.1 kHz and truncated or padded to a fixed duration of 7 seconds. These waveforms are then converted into log-Mel spectrograms before being mapped to the embedding space by the HTSAT encoder. In parallel, the text branch utilizes the CLIP (Radford et al., 2021) transformer to generate corresponding linguistic embeddings.

B Additional Experiments

B.1 Ablation of the number of shots.

To investigate the robustness of our method with respect to varying levels of supervision, we conduct an ablation study under different few-shot settings, where the number of training examples per seen class is set to 2, 4, 8, and 16. As shown in Table 7, the proposed SEPT consistently improves generalization performance across all prompt tuning baselines and all shot settings. Specifically, SEPT achieves notable gains in harmonic mean (H) under low-resource regimes. For instance, with only two shots per class, CoOp (Zhou et al., 2022b) improves from 37.69% to 41.83%. This demonstrates that incorporating semantic neighbors provides a strong regularization effect even when labeled data is scarce. Similar trends are observed for 4-shot and 8-shot settings, where SEPT leads to stable improvements in both base and new class accuracy. Even in the 16-shot setting, where models already achieve relatively strong performance, SEPT yields consistent gains. For example, DePT (Zhang et al., 2024) improves from 46.79% to 49.06% in H, and CoCoOp (Zhou et al., 2022a) reaches 50.65%, up from 46.26%. These results highlight that SEPT

Method	Avg over 11 datasets		
	Base	New	H
Pengi	40.68	38.21	38.46
CoOp	65.00	34.09	42.83
+ SEPT⁻	63.37	36.58	45.09
+ SEPT	64.36	42.98	49.70

Table 8: Ablation study on margin computation strategies. **SEPT⁻** computes margins using a fixed prefix (“This is a sound of {class}”), while **SEPT** averages over a diverse set of $T = 100$ hand-crafted prefixes. H denotes the harmonic mean between Base and New accuracy. Reported values are averaged across 11 datasets.

enhances generalization not only in extremely low-shot regimes but also in moderately supervised scenarios, validating its effectiveness as a robust, plug-and-play augmentation for prompt tuning.

B.2 Ablation of Margin Computation Strategy.

Table 8 presents an ablation study comparing two margin computation strategies in our semantic expansion loss. In **SEPT⁻**, we compute all semantic margins using a single fixed prefix template: “This is a sound of {class}”. This approach already brings noticeable improvements over the CoOp baseline, demonstrating the effectiveness of semantic alignment through even a single prompt template. However, relying solely on one fixed prefix may introduce bias or noise, especially if the template happens to be semantically ambiguous or suboptimal for certain classes. To mitigate this, we propose averaging distances across a diverse pool of $T = 100$ hand-crafted prefixes. This strategy smooths out template-specific noise and produces a more robust

Method	Avg over 11 datasets			Beijing-Opera			NS-Instruments			ESC50		
	Base	New	H	Base	New	H	Base	New	H	Base	New	H
CoOp	65.00	34.09	42.83	97.27	61.38	74.87	41.11	23.88	30.08	61.97	11.83	19.46
KgCoOp	37.99	37.42	36.39	67.26	61.40	62.96	39.80	39.27	39.11	14.10	10.33	11.76
KgCoOp [†]	59.02	38.76	44.76	97.53	72.94	82.55	48.79	44.35	46.41	48.77	14.63	22.39
+ SEPT	58.34	40.52	46.31	97.23	71.22	81.41	49.66	42.04	45.48	46.47	15.77	23.39

Method	ESC50-Actions			UrbanSound8K			CREMA-D			RAVDESS		
	Base	New	H	Base	New	H	Base	New	H	Base	New	H
CoOp	85.33	49.33	61.97	61.86	31.22	40.94	52.50	9.01	15.29	51.01	27.75	35.88
KgCoOp	36.50	30.67	32.88	26.14	26.15	25.40	42.70	55.43	42.48	27.53	27.46	27.42
KgCoOp [†]	66.00	50.17	56.68	55.86	35.32	42.23	54.99	9.56	16.22	46.09	34.21	39.25
+ SEPT	65.33	53.17	58.37	55.50	36.48	43.10	51.12	19.34	26.15	43.56	34.65	38.21

Method	SESA			GT-Music-Genre			VocalSound			TUT2017		
	Base	New	H	Base	New	H	Base	New	H	Base	New	H
CoOp	85.78	74.77	79.33	57.76	20.54	29.76	75.26	45.01	56.00	45.11	20.25	27.53
KgCoOp	68.14	79.28	72.37	27.39	15.82	19.77	52.50	49.91	51.10	15.83	15.90	15.10
KgCoOp [†]	78.43	77.48	77.94	50.83	19.53	28.18	68.30	52.06	59.08	33.62	16.11	21.50
+ SEPT	79.41	78.38	78.74	49.50	21.89	30.04	69.79	55.73	61.96	34.20	17.06	22.50

Table 9: Comparison of base-to-new generalization across different methods with and without SEPT. KgCoOp uses a single prompt template (“This is a sound of”) as in the original paper, whereas KgCoOp[†] averages predictions over 100 hand-crafted prompt templates to enhance robustness and generality. All methods are evaluated using 16-shot learning.

estimation of the semantic margin. As shown in the table, SEPT achieves a substantial performance boost, improving the harmonic mean from 45.09% to 49.70%. This result validates our hypothesis that averaging over multiple semantic views leads to a more stable and generalizable prompt embedding space.

B.3 Analysis of KgCoOp.

Contrary to the performance improvement trend of KgCoOp (Yao et al., 2023) in vision-language models (VLMs), our experiments in the main paper show that it performs significantly worse than CoOp. To further investigate this discrepancy, we conduct a detailed analysis using Table 9.

Following the original implementation, we adopt a single hand-crafted prompt template—“This is a sound of class”—to compute the zero-shot prompt embedding, which serves as the knowledge-guided regularizer for the learned prompt. However, across all datasets, KgCoOp underperforms not only relative to other prompt tuning methods such as CoOp, but in some cases even falls below the zero-shot baseline. This performance drop can be attributed to KgCoOp’s heavy reliance on the quality of the

zero-shot prompt embedding. Since the method is designed to preserve knowledge from the fixed prompt, it implicitly assumes that the hand-crafted prompt yields a robust and semantically rich representation. In our setting, however, the fixed prompt often fails to capture the diverse acoustic semantics of different audio classes, resulting in poor zero-shot performance. While this assumption may hold in VLMs—where strong pretrained models like CLIP and rich textual supervision lead to generally reliable zero-shot regularizer—it becomes less valid in the ALM setting. Compared to vision-language tasks, audio-language models often operate in a narrower semantic space with fewer class labels, and many audio classes are not easily described by generic prompts such as “This is a sound of”. As a result, a poorly aligned base class regularizer can negatively affect transferability and generalization in ALMs.

To address this issue, we construct an enhanced version, KgCoOp[†], which averages the zero-shot embeddings obtained from 100 diverse hand-crafted prompt templates. This ensembling strategy yields a more robust and semantically stable regularizer by mitigating the bias introduced by

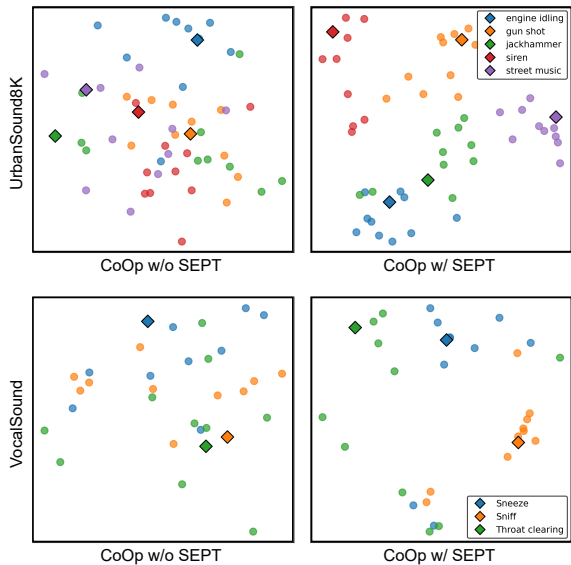


Figure 7: t-SNE visualization of prompt embeddings of new classes and their semantic neighbors from UrbanSound8K dataset (top), VocalSound dataset (bottom). \diamond denotes the prompt embedding of a new class, while \circ represents its corresponding semantic neighbors.

any single prompt. As a result, KgCoOp[†] achieves substantial performance improvements across most datasets, confirming that the diversity and quality of prompts used for knowledge anchoring are critical to the success of knowledge-guided prompt tuning. Furthermore, we apply our proposed SEPT framework on top of KgCoOp[†], further enhancing the generalization ability of the model. SEPT regularizes the prompt embedding space by leveraging semantic neighbors, which complements the stability provided by prompt ensembling.

B.4 Additional Qualitative Results.

We provide additional t-SNE visualizations of prompt embeddings for new classes and their semantic neighbors from the UrbanSound8K (Salamon et al., 2014) and VocalSound (Gong et al., 2022) datasets in Fig. 7. Importantly, the new classes and their corresponding neighbors were not used during training to enforce any semantic structure. Nonetheless, consistent with the observations in the main paper, SEPT produces a well-structured embedding space in which new classes are closely clustered with their semantic neighbors. These results further demonstrate the strong generalization capability of our method across diverse audio classification tasks.




Figure 8: Distribution of similarity scores between classes and semantic neighbors across all datasets. The majority of similarity scores fall between 0.6 and 0.9, indicating a tendency toward moderately high semantic similarity among features.

Dataset	Mean Diversity
Beijing-Opera	0.4482
NS-Instruments	0.4964
ESC50	0.4296
ESC50-Actions	0.4141
UrbanSound8K	0.5657
CREMA-D	0.4480
RAVDESS	0.4618
SESA	0.4382
GT-Music-Genre	0.3491
VocalSound	0.5217
TUT2017	0.3294
Average	0.4457

Table 10: Mean diversity for each dataset.

B.5 Quality of Semantic Neighbors.

Our framework relies on semantic neighbors generated by large language models (LLMs), such as ChatGPT-4o (Achiam et al., 2023), to enhance the semantic structure of the prompt embedding space. Ideally, these semantic neighbors should remain semantically close to the original class while being sampled with slight variations to ensure diversity, which is crucial for achieving effective semantic expansion. To assess the quality of the generated neighbors, we evaluated both their similarity to the original class and their internal diversity using a text encoder with a hand-crafted prompt (“This is a sound of”).

In Fig. 8, we visualize the distribution of cosine similarity scores between each class and its semantic neighbors across all datasets. The majority of semantic neighbors exhibit high similarity scores, indicating that most are well-aligned and semantically relevant. However, a small number show notably low similarity, suggesting the presence of noisy or overly generic neighbors. Furthermore, we present the diversity score of each dataset to

quantify how diversely the semantic neighbors are sampled within the original classes in Table 10. The diversity of class c_i is formulated as:

$$\text{Diversity}_i = 1 - \frac{1}{NC_2} \sum_{n < m} \text{sim}(\mathbf{p}_i^n, \mathbf{p}_i^m), \quad (10)$$

where $\text{sim}(\cdot, \cdot)$ is the cosine similarity. For each dataset, we computed the diversity score for each class and then averaged them across all classes. On average, a diversity score of approximately 0.45 was achieved, indicating a moderate level of semantic variability among neighbors within each class. Overall, the results indicate that the semantic neighbors maintain strong relevance to their original classes while providing enough variation to enhance prompt tuning performance.

B.6 Comparison with Alternative Topology-Preserving Mechanisms

To further validate the necessity of our proposed localized semantic expansion, we investigate whether alternative topology-preserving mechanisms or global regularization strategies could achieve similar improvements in base-to-new generalization. Specifically, we implement and compare two additional baselines. First, inspired by CLIP-Adapter, we introduce Text-Adapter, a learnable linear layer (adapter) on top of the frozen text embeddings. This topology-preserving mechanism adapts the representation without explicitly reshaping local neighborhood relations via loss constraints. Second, we explore Global Regularization, where instead of using targeted local semantic neighbors, we regularize the text embeddings against large, global concept pools via an L2 loss. For this approach, we evaluate two variants, each utilizing a pool of 200 concepts: (1) Global Reg. (FSD50K-200), which consists of classes sampled from the FSD50K dataset (a subset of PENGII’s pretraining data), and (2) Global Reg. (LLM-random-200), which consists of random audio-related concepts generated by an LLM.

As shown in Table 11, both the topology-preserving Text-Adapter and the global regularization strategies improve new-class generalization compared to the standard CoOp baseline. This indicates that preventing the structural collapse of the pre-trained embedding space—either via architectural constraints or global anchors—is indeed beneficial. However, our proposed SEPT still consistently outperforms these alternatives, achieving

Method	Base	New	H
CoOp	65.00	34.09	42.83
Text-Adapter	60.85	41.30	46.67
Global Reg. (FSD50K-200)	64.74	39.21	46.81
Global Reg. (LLM-random-200)	63.32	41.66	47.50
SEPT (Ours)	64.36	42.98	49.70

Table 11: Comparison with alternative topology-preserving and global regularization mechanisms. Performance is averaged across 11 datasets.

the highest harmonic mean (49.70%). This demonstrates that the gains from alternative topology-preserving methods are complementary to our approach. The targeted, task-relevant inductive bias provided by localized semantic neighbor expansion in SEPT proves to be more effective for base-to-new generalization than relying solely on linear adapters or broad global concept pools.

B.7 Applicability on VLMs.

To further support the plug-and-play property of SEPT beyond ALMs, we additionally evaluate SEPT on a VLM prompt tuning setup. Specifically, we apply SEPT on CoOp, using CLIP (Radford et al., 2021) as a backbone, and evaluate the base-to-new generalization on Caltech101 (Fei-Fei et al., 2004). As shown in Table 12, SEPT improves performance on both base and novel classes, indicating that it transfers effectively to VLM prompt tuning. These results validate the modality-agnostic nature of SEPT, confirming its potential as a universal regularizer that enhances prompt tuning performance across both audio and vision domains. However, we emphasize that the impact and practicality of SEPT are particularly pronounced in the audio domain. Unlike vision benchmarks, which often contain dense category sets (e.g., ImageNet), audio datasets frequently suffer from class sparsity. In such sparse semantic spaces, the prompt embedding structure is more prone to degradation, making the semantic expansion provided by SEPT even more critical for preserving generalization capabilities. Furthermore, since the computational cost of neighbor-based regularization scales with the number of categories, SEPT remains highly efficient in these sparse audio settings, offering a favorable balance between structural robustness and training overhead compared to dense vision tasks.

Method	Base	New	H
CoOp	98.03	90.10	93.90
+ SEPT	98.23	93.40	95.76

Table 12: Results on Caltech101 under the base-to-new generalization protocol with CLIP/CoOp. SEPT consistently improves base and novel accuracies and increases the harmonic mean.

Domain	Dataset	Cos. Sim.		Mat. Ent.		# Base
		Zeroshot	CoOp	Zeroshot	CoOp	
Audio	CREMA-D	0.583	0.921	0.911	0.664	3
Audio	ESC50	0.363	0.505	0.942	0.882	25
Image	Caltech101	0.696	0.404	0.907	0.953	50
Image	ImageNet	0.544	0.538	0.892	0.883	500

Table 13: Comparison of global geometry changes in the base-class text embedding space after prompt tuning. In ALMs, CoOp increases global clustering and decreases effective dimensionality, whereas VLMs show much weaker or no collapse-like behavior.

C Additional Analysis

C.1 Analysis of embedding space structure collapse in ALMs and VLMs

To clarify why SEPT is particularly beneficial in the audio domain, we analyze how prompt tuning changes the global structure of the base-class text embedding space in both ALMs and VLMs. Starting from the zeroshot text prototypes and then applying CoOp, we measure two complementary indicators of global geometry: (i) the mean pairwise cosine similarity between class embeddings, which reflects how strongly class prototypes become globally clustered, and (ii) the normalized matrix entropy of the singular-value spectrum, which captures how evenly the embedding energy is distributed across dimensions. A lower matrix entropy indicates a more anisotropic and lower-dimensional geometry, i.e., stronger collapse.

Let $Z \in \mathbb{R}^{K \times d}$ denote the matrix of base-class text embeddings, where K is the number of base classes and d is the embedding dimension. Given the singular values $\{s_i\}_{i=1}^r$ of Z , we define $p_i = s_i / \sum_{j=1}^r s_j$ and compute the normalized matrix entropy as

$$H_{\text{norm}}(Z) = \frac{-\sum_{i=1}^r p_i \log p_i}{\log r},$$

where r is the rank of Z .

Table 13 summarizes the results. In ALMs, prompt tuning consistently increases the mean pairwise cosine similarity while decreasing matrix entropy, indicating that class prototypes become more

Method	Avg over 11 datasets		
	Base	New	H
CoOp	65.00	34.09	42.83
+ SEPT	64.36	42.98	49.70
+ SEPT[‡]	64.37	42.93	49.64

Table 14: Impact of new class name overlap in generated semantic neighbors. Filtering these neighbors (SEPT[‡]) yields nearly identical performance over 11 datasets, indicating minimal effect from the overlap.

globally clustered and concentrated in a lower-dimensional subspace. This trend is especially pronounced in semantically sparse audio benchmarks such as CREMA-D and ESC50. In contrast, VLMs do not exhibit the same collapse-like behavior: the cosine similarity remains stable or even decreases, and the matrix entropy is largely preserved. These observations suggest that prompt tuning distorts the pretrained text-space geometry much more severely in ALMs than in VLMs. We attribute this difference to the fact that audio benchmarks typically contain fewer and sparser semantic categories, making the prompt space more vulnerable to structural degradation. This also explains why preserving local semantic structure is particularly important in ALM few-shot adaptation, while SEPT remains applicable to VLMs as a general regularization strategy.

C.2 Neighbor Quality.

To assess the reliability of the generated semantic neighbors, we quantify their alignment with class concepts by computing the average cosine similarity across 11 datasets. The results—0.64 for positive pairs versus 0.34 for negative pairs—demonstrate that the generated neighbors achieve a balance between semantic coherence and diversity. Qualitative inspection confirms that these neighbors capture diverse facets of a class (e.g., neighbors for footsteps include specific variations like running steps or gravel crunch).

As shown in Fig. 8 and Table 10, certain semantic neighbors exhibit low similarity to their target classes or lack sufficient diversity, which may introduce noise or redundancy into the embedding space. Crucially, SEPT exhibits robustness to potential generation noise because neighbors are employed exclusively for text-space regularization, distinct from the cross-entropy supervision used for base class classification. Consequently, noisy neighbors do not directly corrupt the primary de-

cision boundaries. The core safeguard lies in our margin constraint, which is precomputed based on the text-space geometry of the pretrained ALM—a space already aligned with audio semantics via contrastive pretraining. In this framework, an acoustically distant positive neighbor is assigned a large margin, effectively preventing the loss term from activating unless the embedding violates this geometric prior. This mechanism acts as a gate for gradient updates, naturally capping the influence of outliers and mitigating the impact of neighbor noise.

We further analyze whether the generated neighbors contain terms that overlap with the held-out new class names, to address the potential concern that such overlap could inadvertently introduce new class information. Across the 11 datasets, only 1.7% of the neighbors generated for base classes overlap with any new class name, indicating that such cases are rare. More importantly, this overlap does not imply using any new class audio samples: SEPT never accesses new class audio samples during training, and the neighbors are not used in the cross-entropy objective for base class classification. Instead, neighbors are used solely in the text-space regularization term, which shapes the prompt embedding geometry rather than directly supervising class decision boundaries. To empirically assess the impact of this overlap, we define **SEPT[‡]**, which filters out all neighbors that overlap with any new class name before training. As shown in Table 14, SEPT[‡] achieves performance nearly identical to SEPT over 11 datasets, suggesting that the observed gains are not driven by the small fraction of overlapping neighbors.

C.3 Performance Variations per Dataset.

While SEPT consistently improves generalization across the majority of benchmarks, we observe marginal gains or slight performance drops in specific dataset–method combinations. We attribute these variations to two primary factors:

1. Semantic Ambiguity: In domains with high abstraction, such as music genres or emotional states, semantic boundaries are inherently ambiguous. In such cases, generated neighbors may overlap between classes or lack distinctiveness (e.g., exhibiting small similarity gaps between positive and negative neighbors). This can introduce noise into the embedding space rather than structural clarity, potentially hindering the precise optimization of base classes.

2. Compound Regularization Effects: SEPT functions as an explicit regularizer to structure the embedding space. When integrated with baseline methods that already impose strong constraints (e.g., via knowledge distillation or channel bias correction), the addition of SEPT may result in an over-regularized objective. This can excessively constrain the optimization landscape, limiting the model’s ability to fully fit the base classes. Nevertheless, even in these instances, SEPT typically preserves or improves the harmonic mean, demonstrating that the trade-off favors enhanced generalization to unseen classes.

D AI Assistant Usage Statement

We utilize Large Language Models (LLMs) for two distinct purposes in this study:

Model-based Experimentation. We employ ChatGPT-4o as the primary engine to generate semantic neighbors for the proposed SEPT framework. Additionally, to verify the robustness of our method across different architectures (as detailed in Sec. 4.4), we utilize Gemini 2.5 Pro and Qwen3 Max to generate neighbor candidates.

Writing Assistance. We use Gemini 2.5 Pro exclusively for grammatical error correction and stylistic polishing of the manuscript. We confirm that no AI tools were used to generate scientific claims, experimental results, or the core intellectual contributions of this paper.