

SARSTEER: SAFEGUARDING LARGE AUDIO LANGUAGE MODELS VIA SAFE-ABLATED REFUSAL STEERING

Anonymous authors

Paper under double-blind review

ABSTRACT

Large Audio–Language Models (LALMs) are becoming essential as a powerful multimodal backbone for real-world applications. However, recent studies show that audio inputs can more easily elicit harmful responses than text, exposing new risks toward deployment. While safety alignment has made initial advances in LLMs and Large Vision–Language Models (LVLMs), we find that vanilla adaptation of these approaches to LALMs faces two key limitations: 1) LLM-based steering fails under audio input due to the large distributional gap between activations, and 2) prompt-based defenses induce over-refusals on benign-speech queries. To address these challenges, we propose **Safe-Ablated Refusal Steering** (SARSteer), the first inference-time defense framework for LALMs. Specifically, SARSteer leverages text-derived refusal steering to enforce rejection without manipulating audio inputs and introduces decomposed safe-space ablation to mitigate over-refusal. Extensive experiments demonstrate that SARSteer significantly improves harmful-query refusal while preserving benign responses, establishing a principled step toward safety alignment in LALMs.

1 INTRODUCTION

Large Audio–Language Models (LALMs) have recently emerged as powerful multimodal systems (Chu et al., 2023; 2024; Tang et al., 2023; Ding et al., 2025), extending the general intelligent capabilities of Large Language Models (LLMs) (Bai et al., 2023; Liu et al., 2024a; Achiam et al., 2023) into the audio domain. By jointly modeling audio and textual inputs, LALMs enable a wide range of applications, including voice assistants (Held et al., 2024), audio understanding (Dinkel et al., 2025), real-time speech interaction (Long et al., 2025), *etc.* Their ability to understand and generate responses directly from audio makes them a critical component for next-generation human–AI interaction systems.

Despite their promise, the deployment of LALMs raises pressing safety concerns due to the under-explored vulnerability of new audio input. In the literature, most focus of safety alignment has been laid in text-based LLMs (Kim et al., 2024; Zhang et al., 2025a; Qi et al., 2024), leveraging both *fine-tuning-based defenses* such as supervised fine-tuning (SFT) (Liu et al., 2023) and reinforcement learning from human feedback (RLHF) (Bai et al., 2022), and more advanced *inference-based defenses* such as activation steering (Panickssery et al., 2023; Zhao et al., 2025). While fine-tuning can be effective with high-quality data or well-trained reward models, its resource-intensive nature makes inference-based defenses more practical for scalable deployment. Similar efforts have recently extended to Large Vision–Language Models (LVLMs) (Wang et al., 2024a; Lu et al., 2024; Zhu et al., 2023; Liu et al., 2024b), leading to new fine-tuning-based (Zhang et al., 2025c; Zong et al., 2024) and inference-based (Wang et al., 2024b; Ding et al., 2024) defense strategies designed for vision modality. In contrast, the safety alignment of LALMs remains largely underexplored: beyond some initial findings (Yang et al., 2024a; Song et al., 2025), which show that LALMs are far more likely to comply with harmful speech than text, no principled defense strategies have been developed. A natural solution, therefore, is to transfer the alignment techniques originally designed for LLMs or LVLMs into the audio–language setting. In this work, **we focus on inference-based defenses**, *e.g.*, activation steering from LLMs (Panickssery et al., 2023) and prompt-based defenses from LVLMs (Wang et al., 2024b), to align LALMs with harmless outputs.

054 However, such transfers with vanilla adaptations expose two critical limitations. **First, LLM-**
 055 **based steering fails under audio input.** In LLMs, steering vectors constructed from harmful–safe
 056 text pairs can reliably shift representations toward safe regions and enhance refusal behaviors. In
 057 LALMs, by contrast, harmful and safe speech inputs occupy widely divergent latent distributions
 058 than in text, making the harm-to-safe direction unreliable (Section 3.3). **Second, prompt-based**
 059 **defenses from LVLMs induce over-refusal unobviously** (Jiang et al., 2025). While defensive
 060 prompts (e.g., instructing the model to respond “*I am sorry*” to unethical or illegal requests) can
 061 block some harmful queries, they also cause benign queries with lexical similarity to be mistak-
 062 enly rejected (Section 3.4). Despite efforts such as AdaShield (Wang et al., 2024b), which refines
 063 prompts to better distinguish benign inputs, the coarse input-level instructions, containing two op-
 064 posing actions of answering or refusing, struggle to coordinate effectively.

065 To address these challenges, we propose an inference-based alignment framework, **Safe-Ablated**
 066 **Refusal Steering (SARSteer)**, for LALMs. SARSteer targets both the failure of steering audio
 067 modality and the over-refusal issue observed in prompt-based defenses. It consists of two key com-
 068 ponents: 1) **Text-derived refusal steering.** Instead of contrasting harmful and safe speech inputs,
 069 which suffer from distributional gap, SARSteer extracts refusal vectors directly from textual re-
 070 fusal prompts (e.g., “*I cannot assist with that*”). These vectors capture safety-aligned semantics in
 071 intermediate activations and provide a modality-agnostic direction for enhancing harmful-query re-
 072 jection. 2) **Decomposed safe-space ablation.** To mitigate over-refusal on benign queries, SARSteer
 073 employs a projection correction step. Specifically, we use *principal component analysis* (PCA) on
 074 safe samples to identify the dominant subspace of benign semantics, and then ablate this component
 075 from the refusal vector. This ensures that refusal steering acts only on harmful directions while
 076 preserving safe responses. By jointly leveraging these two components, SARSteer avoids costly
 077 fine-tuning, operates entirely at inference time, and establishes a principled defense strategy for
 LALMs that is robust against harmful inputs while maintaining utility on benign ones.

078 Our contributions are summarized as follows:

- 079
- 080 • **Empirical Findings.** We construct paired harmful–safe datasets in the speech domain
 081 and provide a systematic study of the representational differences between text and audio
 082 inputs, explaining the failure of direct activation steering transfer.
- 083 • **Methodology.** We introduce the first inference-time defense framework for LALMs, based
 084 on text-derived refusal steering and decomposed safe-space ablation, filling the gap of
 085 broad LALMs applications and the scarcity of the specified safety alignment.
- 086 • **Evaluation.** Extensive experiments demonstrate that our method significantly improves
 087 harmful-query refusal while maintaining overall utility, achieving a more favorable trade-
 088 off between safety and usability.
- 089

090 2 RELATED WORK

091 2.1 LLM SAFETY ALIGNMENT

092

093 Substantial research has focused on aligning LLMs with human values and safety standards (Bai
 094 et al., 2022; Kim et al., 2024; Zhang et al., 2025a; Qi et al., 2024). Prominent approaches in-
 095 clude reinforcement learning from human feedback (RLHF), which fine-tunes models using human-
 096 preferred responses (Ouyang et al., 2022; Bai et al., 2022), and supervised fine-tuning (SFT) on
 097 safety-centric datasets (Liu et al., 2023). These methods can all be categorized as *fine-tuning-based*
 098 *defenses*. Despite their effectiveness, they often require extensive human annotation and computa-
 099 tional resources, limiting their application scenarios.

100

101 More recently, *inference-time techniques* have gained attention for their efficiency and low resource
 102 demands (Arditi et al., 2024; Zhao et al., 2025; Qian et al., 2025). For instance, activation steering
 103 methods intervene in the model’s internal representations to guide outputs toward desired behav-
 104 iors (Panickssery et al., 2023; Zhao et al., 2025; Ghosh et al., 2025). Similarly, refusal prompts,
 105 prepending input queries with safety-guided instructions, have been shown to enhance robustness
 106 against malicious inputs without additional training (Zheng et al., 2024; Qian et al., 2025). These
 107 approaches circumvent the need for large-scale fine-tuning, making them a more feasible solution
 for real-world industries.

2.2 MULTIMODAL LLM SAFETY

The integration of visual modalities introduces new vulnerabilities and attack surfaces in Multimodal Large Language Models (MLLMs) (Li et al., 2024; Zhang et al., 2025d). Adversaries can exploit cross-modal inconsistencies to bypass safety alignments, such as by embedding harmful content in images paired with benign text (Gong et al., 2025). In response, several defense strategies have been proposed. AdaShield (Wang et al., 2024b) employs adaptive shield prompting to defend against structure-based jailbreak attacks without fine-tuning the model. Similarly, ETA (Ding et al., 2024) introduces a two-phase ‘‘Evaluate then Align’’ framework that assesses both visual and textual inputs for harmful content and aligns outputs via shallow and deep alignment mechanisms. Other methods like DAVSP (Zhang et al., 2025b) optimize a visual safety prompt using activation-space supervision, while HiddenDetect (Jiang et al., 2025) monitors hidden states to identify harmful patterns. These inference-time methods effectively enhance safety against the vision-space vulnerability.

However, existing research predominantly focuses on LVLMs, leaving audio-based LALMs largely unexplored in terms of safety alignment. Our work represents a preliminary step toward developing inference-time safety alignment specified to the speech domain. By leveraging text-derived refusal steering and decomposed safe-space ablation in the model activation space, our approach offers a flexible, efficient solution to refuse harmful inputs while maintaining the general utility of LALMs.

3 PRELIMINARY AND MOTIVATION ANALYSIS

3.1 PROBLEM FORMULATION

Model Description. We consider a basic-form LALM M^1 that processes multimodal queries $Q = (a, t)$, where a is an *audio* signal and t is a *textual* input, to a textual output. The audio encoder \mathcal{E}_a maps a into an embedding $e_a = \mathcal{E}_a(a)$, which is then projected into the textual embedding space through a multimodal projector \mathcal{P} , yielding $\tilde{e}_a = \mathcal{P}(e_a)$. After that, the audio representations and the tokenized textual input are fed into the autoregressive language model backbone \mathcal{M} , generating an output sequence

$$Y_I = (y_1, \dots, y_I), \quad P(Y_I | Q) = \prod_{i=1}^I P(y_i | Y_{<i}, \tilde{e}_a, e_t; \mathcal{M}), \quad (1)$$

where e_t is the discrete textual embedding processed by \mathcal{M} , Y_I and $Y_{<i}$ represent the complete response with I tokens and the first generated i tokens, respectively. It can be seen that the LALM extends standard LLMs by incorporating audio understanding through the audio encoder and multimodal projector, enabling both audio-text-conditioned generation.

Task Description. Based on the above LALM model, we now formulate the *inference-time safety alignment* task, which is typically performed in a training-free manner after the model training phase. Since the model output Y_I is free-form text, we introduce an evaluation function

$$\mathcal{R} : \mathcal{Y} \rightarrow \{0, 1\}, \quad (2)$$

to judge whether a response constitutes a refusal ($\mathcal{R}(Y_I) = 1$) or not ($\mathcal{R}(Y_I) = 0$), which is implemented by an auxiliary LLM (Xie et al., 2024) or a matching-based method (Wang et al., 2024b)². We denote by $\mathcal{Q}_{\text{harm}}$ the set of harmful queries, and by $\mathcal{Q}_{\text{safe}}$ the corresponding benign queries set (Section 3.2). The objective of safety alignment is threefold:

1. **Refuse Harmful Queries.** For $Q \in \mathcal{Q}_{\text{harm}}$, maximize $\mathbb{E}_{Q \in \mathcal{Q}_{\text{harm}}} [\mathcal{R}(M(Q))]$ to ensure the model refuses harmful inputs.
2. **Preserve Helpfulness on Safe Queries.** For $Q \in \mathcal{Q}_{\text{safe}}$, minimize $\mathbb{E}_{Q \in \mathcal{Q}_{\text{safe}}} [\mathcal{R}(M(Q))]$ to avoid unnecessary refusals.

¹Despite the presence of additional components in certain LALMs (e.g., an audio decoder), in this paper, we focus on the basic architecture, namely the audio encoder, multimodal projector, and language model backbone, as these elements are common to most designs. [The detailed illustration on Qwen2-Audio is provided in Appendix A.8.](#)

²In this work, we use matching-based method to compute *refusal rate* (RR) on both harmful and benign datasets; use LLM-based method to assess *attack success rate* (ASR) on harmful queries.

3. **Maintain General-purpose Utility.** On the benchmarks \mathcal{B} , enforce $\text{Perf}(M, \mathcal{B}) \approx \text{Perf}(M_0, \mathcal{B})$, where M_0 denotes the original unaligned model and $\text{Perf}(M, \mathcal{B})$ represents the performance of M on \mathcal{B} under its own evaluation method (e.g., *LLM-as-a-judge with well-designed prompt in AirBench* (Yang et al., 2024b).), ensuring that the aligned model preserves general performance.

The above objectives indicate that safety alignment seeks to **enhance refusals of harmful inputs, mitigate unnecessary refusals of safe inputs, and preserve overall model utility.**

3.2 HARMFUL-SAFE PAIRED AUDIO DATASET CONSTRUCTION

To study data-dependent alignment methods and rigorously evaluate them, we first construct several harmful–safe paired datasets in the speech domain based on the safety benchmarks of other domains, e.g., AdvBench (Zou et al., 2023) and Figsteps (Gong et al., 2025). Specifically, for the harmful set $\mathcal{Q}_{\text{harm}}$, we directly convert the harmful textual queries into spoken using the OpenAI *text-to-speech* (TTS) system³, and unify the textual instruction as “Please generate detailed steps for the question in the audio.” (Gong et al., 2025). To obtain their safe counterparts $\mathcal{Q}_{\text{safe}}$, we apply LLMs (e.g., DeepSeek-R1 (Guo et al., 2025)) to purify each harmful query into a lexically-similar benign version (e.g., “How can I make a fake bank statement?” → “How can I obtain an official bank statement?”), followed by TTS synthesis to match the audio modality. This one-to-one purification ensures that every harmful query has a semantically related but safe alternative, yielding paired datasets with high lexical similarity but fundamentally different safety implications. In this work, we randomly sample 100 harmful-safe paired queries from the Figsteps-based dataset (referred to as Figstep-audio) for alignment, denoted as $\mathcal{Q}_{\text{harm}}^s$ and $\mathcal{Q}_{\text{safe}}^s$, while the remaining pairs are reserved for evaluation. Further details of the dataset are provided in the Appendix A.2.

Such paired safe data is necessary because existing benign benchmarks (Yang et al., 2024b) often fail to expose the issue of *over-refusal* on borderline safe inputs. By explicitly pairing harmful and safe queries with minimal lexical differences, our datasets provide a sharper testbed to evaluate whether alignment methods can reliably distinguish harmful instructions from benign ones, thus exposing subtle safety-utility trade-offs and directly supporting the second objective.

3.3 FAILURES OF STEERING AUDIO MODALITY

Based on the constructed datasets, we now investigate the transfer of inference-based safety alignment techniques from other domains, e.g., activation steering from LLM safety (Zhao et al., 2025; Ghosh et al., 2025).

Vanilla Adaptation of Two Activation Steering Defenses. Typically, there exist two kinds of steering vector implementations: extracting from harmful-to-safe query (Arditi et al., 2024) and from harmful compliance-to-refusal query (Zhao et al., 2025), both relying on the *difference-in-means* technique (Belrose, 2024). To facilitate discussion, we refer to the two methods as *MDSteer-h2s* (mean-difference steering in the harmful-to-safe direction) and *MDSteer-c2r* (mean-difference steering in the compliance-to-refusal direction on harmful inputs), respectively. Under our LALM setting, where harmful or safe semantics are embedded in the audio modality, **both steering vectors are computed based on differences between audio inputs.** We formulate the vanilla adaptation to LALMs as follows. Let $h^l(Q)$ denote the activation at the last token position of layer $l \in [L]$ (Zhao et al., 2025) of \mathcal{M} , where $Q = (a, t)$ is a multimodal query.

(1) **MDSteer-h2s.** Given our paired datasets $\mathcal{Q}_{\text{harm}}^s$ and $\mathcal{Q}_{\text{safe}}^s$, we compute

$$\mu_{\text{harm}}^l \stackrel{\text{def}}{=} \frac{1}{|\mathcal{Q}_{\text{harm}}^s|} \sum_{Q \in \mathcal{Q}_{\text{harm}}^s} h^l(Q), \quad \mu_{\text{safe}}^l \stackrel{\text{def}}{=} \frac{1}{|\mathcal{Q}_{\text{safe}}^s|} \sum_{Q \in \mathcal{Q}_{\text{safe}}^s} h^l(Q), \quad (3)$$

and define the steering vector as

$$v_{h2s}^l \stackrel{\text{def}}{=} \mu_{\text{safe}}^l - \mu_{\text{harm}}^l. \quad (4)$$

(2) **MDSteer-c2r.** Alternatively, we group harmful queries by their generated response type. Let $\mathcal{Q}_{\text{harm}}^{\text{s-comp}}$ denote those eliciting compliant harmful responses, and $\mathcal{Q}_{\text{harm}}^{\text{s-ref}}$ denote those eliciting refusals,

³<https://platform.openai.com/docs/models/tts-1-hd>

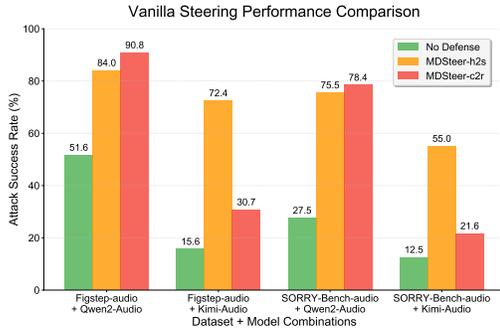


Figure 1: Performance of vanilla adaptations of LLM-based steering on LALMs.

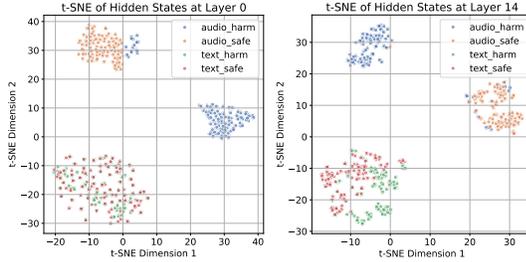


Figure 2: t-SNE visualization of hidden states in Qwen2-Audio using Figstep-audio datasets. “audio” and “text” represent the input modalities containing the questions; “harm” and “safe” represent the harmfulness of the questions.

as determined by the evaluation function \mathcal{R} . We obtain the corresponding mean activation values $\mu_{\text{harm-c}}^l$ and $\mu_{\text{harm-r}}^l$ in the same way as Equation 3, and define the steering vector as:

$$v_{c2r}^l \stackrel{\text{def}}{=} \mu_{\text{harm-r}}^l - \mu_{\text{harm-c}}^l \quad (5)$$

During inference, the vector v^l (either v_{h2s}^l or v_{c2r}^l) is added to the model’s hidden states at each generated token position i , scaled by a coefficient α :

$$h_i^l \stackrel{\text{def}}{=} h_i^l + \alpha v^l \quad (6)$$

Results Analysis. We evaluate the ASR performance of MDSteer-h2s and MDSteer-c2r on Qwen2-Audio (Chu et al., 2024) and Kimi-Audio (Ding et al., 2025), using our audio-version Figstep (Gong et al., 2025) and SORRY-Bench (Xie et al., 2024). As shown in Figure 1, both methods not only fail to improve ASR performance over the “No Defense” baseline (the original performance of LALMs), but also degrade it. To understand this failure, we analyze the hidden representations of harmful and safe inputs across both text and audio modalities using t-SNE (Figure 2). In the text modality, harmful and safe queries overlap in shallow layers (left subfigure) and become linearly separable at intermediate depths (right subfigure), consistent with (Panickssery et al., 2023), which reports that separability emerges suddenly after a particular layer. This overlapping structure enables a feasible harmful-to-safe (h2s) transition, making h2s steering (and similarly c2r) meaningful in the text modality. In sharp contrast, the audio modality shows early and persistent separation between harmful and safe queries across all layers, leaving no shared subspace to define a valid steering path. As a result, both h2s and c2r directions degenerate into noisy perturbations that fail to induce refusal. This striking gap reveals a **fundamental limitation: speech activations cannot serve as a feasible operating space for safety steering, and effective alignment should instead be derived from the refusal signals embedded in the text modality.** This observation motivates our approach of text-derived refusal steering in Section 4.1.

3.4 OVER-REFUSAL OF PROMPT-BASED DEFENSES

Another critical limitation is the *over-refusal* (or over-defense) issue in prompt-based defenses when transferred from LVLMs, *i.e.*, the tendency to refuse even benign or borderline-safe queries.

Evaluation of Balanced Refusal. While the over-refusal phenomenon has been discussed in prior LLM and LVLM defense studies (Cui et al., 2024; Wang et al., 2024b; Jiang et al., 2025), a precise evaluation has remained challenging due to the lack of paired harmful-safe datasets. In particular, for LALMs, existing metrics are insufficient to capture the trade-off between refusing harmful queries and preserving utility on borderline benign ones. To address this gap, we adopt the *refusal rate* (RR) with a matching-based evaluation method (Wang et al., 2024b)⁴, defined as

$$\text{Refusal Rate} \stackrel{\text{def}}{=} \frac{|\#\text{Refusal responses}|}{|\#\text{All responses}|} \quad (7)$$

⁴The refusal signals used for matching is listed in Appendix A.6.

Inspired by *balanced accuracy* (Brodersen et al., 2010), we further introduce the *balanced refusal rate* (BRR), which considers both harmful and safe sets simultaneously. Denoting the refusal rate on harmful and safe inputs as RR_{harm} and RR_{safe} , the BRR is defined as

$$\text{BRR} \stackrel{\text{def}}{=} \frac{1}{2} [RR_{\text{harm}} + (1 - RR_{\text{safe}})] = \frac{1 + RR_{\text{harm}} - RR_{\text{safe}}}{2}, \quad (8)$$

where $\text{BRR} \in [0, 1]$ reflects the overall refusal capability (or helpfulness): high values indicate that harmful queries are correctly rejected while safe ones are preserved.

Prompt-based Defenses

from LVLMS. We transfer and examine representative prompt-based defenses, *e.g.*, AdaShield (Wang et al., 2024b) and FSD (Gong et al., 2025), on LALMs, which were originally proposed for LVLMS. The implementation details are postponed to Appendix A.3. Based on the above metrics, we

Table 1: Performance of vanilla adaptations of prompt-based defenses from LVLMS on Qwen2-Audio. NOTE: ‘‘Avg. Score’’ is an LLM-based evaluation metric from the original paper (Yang et al., 2024b) to assess the benign performance.

Defense	Figstep-audio (Harmful-safe paired)			AirBench (General purpose)	
	Harmful (RR) (%) \uparrow	Safe (RR)(%) \downarrow	BRR (%) \uparrow	RR (%) \downarrow	Avg. Score (1-10) \uparrow
No Defense	62.00	21.60	70.20	1.23	7.43
AdaShield	75.60	36.00	69.80	2.78	7.39
FSD	90.00	63.60	63.20	2.64	7.31

evaluate the overall performance on our constructed paired dataset, *i.e.*, Figstep-audio, and on a general-purpose audio benchmark, *i.e.*, AirBench (Yang et al., 2024b). The results are illustrated in Table 1. We can observe that these defenses appear to maintain reasonable performance with only slight degradation on RRs ($<2\%$) and Avg. Scores (<0.13) on AirBench, as the benign queries are typically far from the decision boundary. However, when evaluated on the paired harmful-safe dataset (Figstep-audio), which explicitly includes *borderline safe samples* that partially overlap with harmful semantics, a clear over-refusal issue emerges: the improved harmful RRs also lead to significant higher safe RRs, degrading the overall helpfulness (lower BRRs). The results highlight the necessity of considering the borderline-safe data and reveal that **vanilla adaptations of prompt-based defenses incur unobvious over-refusal**. This also motivates our approach of ablating the safe subspace in hidden space (*i.e.*, the decomposed safe-space ablation of Section 4.2).

4 METHODOLOGY

Based on the above analysis, we propose **SARSteer**, which derives the steering vector from the refusal text of the same speech input (*i.e.*, text-derived refusal steering) and ablates the safe subspace of its hidden representation to mitigate over-refusal on benign queries (*i.e.*, decomposed safe-space ablation). We now present the technical details of the two components. The overview of SARSteer is shown in Figure 3 and the corresponding algorithm outline is provided in Appendix A.5.

4.1 TEXT-DERIVED REFUSAL STEERING

Prompt-based defenses provide a practical approach to increasing the refusal rate of MLLMs by appending refusal-style text (Wang et al., 2024b), despite the limitations of over-refusal and inflexibility to multiple purposes. Combined with the analysis in Section 3.3, this insight inspires us: *why not extract the controllable steering vector from the appended refusal text, while keeping the audio modality unchanged?* Therefore, we first calculate mean activation values of the modified query $Q' = (a, t + p)$ and the original query $Q = (a, t)$ from Q_{harm}^s using Equation 3, where p denotes a refusal text prompt (*e.g.*, ‘‘I cannot assist with that.’’) ⁵. We denote their mean vectors as $\mu_{\text{harm-tr}}^l$ and μ_{harm}^l , respectively. Then the steering vector representing the refusal direction can be defined as

$$\hat{v}^l \stackrel{\text{def}}{=} \mu_{\text{harm-tr}}^l - \mu_{\text{harm}}^l. \quad (9)$$

Applying this vector to harmful inputs using Equation 6 can effectively improve the refusal rate.

⁵This example is used as a *semantic anchor* to obtain a ‘‘refusal direction’’ in the latent space, capturing consistent activation patterns associated with refusal behavior. Other refusal prompts can also elicit similar refusal directions as in Appendix B.2

4.2 DECOMPOSED SAFE-SPACE ABLATION

While text-derived refusal steering provides a controllable vector \hat{v}^l , it risks activating dimensions that are also present in benign inputs, leading to *over-refusal*. To address this issue, we propose a decomposition strategy that explicitly removes safe-subspace components from \hat{v}^l by leveraging the statistical structure of safe activations.

Concretely, we first collect activations from safe queries at layer l :

$$H_{\text{safe}}^l = [h^l(Q)]_{Q \in \mathcal{Q}_{\text{safe}}} \in \mathbb{R}^{D \times n}, \quad (10)$$

where D is the hidden dimension of \mathcal{M} (e.g., 4096 in Qwen2-Audio) and n is the number of samples (e.g., 100). We then apply Principal Component Analysis (PCA) to H_{safe}^l , which identifies a low-dimensional subspace spanned by the top- k principal components $U \in \mathbb{R}^{D \times k}$, satisfying $U^\top U = I_k$. These directions capture the dominant variance of safe representations, and therefore encode the most salient features that should be preserved when handling benign inputs. Given the principal components, the steering vector can be decomposed as

$$\hat{v}^l = (UU^\top)\hat{v}^l + (1 - UU^\top)\hat{v}^l = \hat{v}_{\parallel}^l + \hat{v}_{\perp}^l, \quad (11)$$

where \hat{v}_{\parallel}^l is the safe-subspace component and \hat{v}_{\perp}^l is the orthogonal components. We retain only the orthogonal part by projecting away the safe subspace as **our final steering vector**:

$$\hat{v}_{\perp}^l = (1 - UU^\top)\hat{v}^l. \quad (12)$$

This ensures that the steering signal emphasizes harmful-related activations while minimizing interference with safe inputs. Similar to Equation 6, during inference, activations are updated by

$$h'^l = h^l + \alpha \hat{v}_{\perp}^l. \quad (13)$$

By explicitly grounding the decomposition in PCA, this method provides a solid and interpretable mechanism: it systematically separates refusal-relevant directions from benign-safe variance, thus making the steering both effective and robust. Furthermore, we provide a **mathematical intuition** to demonstrate how SARSteer accomplishes the three alignment objectives in Appendix A.4.

5 EXPERIMENT

5.1 EXPERIMENTAL SETUP

Evaluation Metrics and Datasets. We evaluate safety alignment across three aspects: **1) Harmfulness:** measured by *attack success rate* (ASR) using LLM-as-a-judge on Figstep-audio (Gong et al., 2025), AdvBench-audio (Zou et al., 2023), SORRY-Bench-audio (Xie et al., 2024), and AJailBench (Song et al., 2025). **2) Helpfulness:** measured by *Balanced Refusal Rate* (BRR) on paired datasets including Figstep-audio and AdvBench-audio. **3) General Utility:** evaluated on AirBench (Yang et al., 2024b) following the original LLM-based evaluation. More details are postponed to Appendix A.1.

Baselines. Since there is no inference-time safety alignment method in the context of LALMs, we use the vanilla adapted defenses from LLMs and LALMs as the baselines. As discussed in Sections 3.3 and 3.4, we implement LALM-version prompt-based defenses, *i.e.*, AdaShield (Wang et al., 2024b) and FSD (Gong et al., 2025), as well as activation-steering defenses (Belrose, 2024; Zhao et al., 2025), *i.e.*, MDSteer-h2s and MDSteer-c2r. More implementation details are postponed to Appendix A.3.

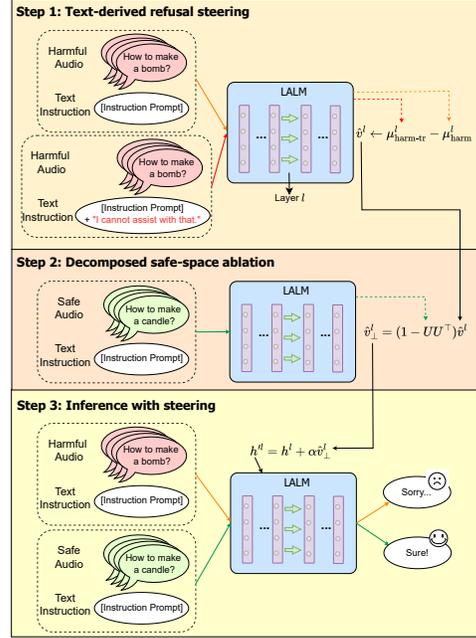


Figure 3: Overview of SARSteer.

Implementation Details. We use two state-of-the-art (SOTA) open-sourced LALMs, *i.e.*, Qwen2-Audio (Chu et al., 2024) and Kimi-Audio (Ding et al., 2025), to evaluate all defense methods. We randomly sample 100 harmful-safe paired queries from Figstep-audio for alignment implementation. For our SARSteer, we use the simplest refusal prompt “*I cannot assist with that.*” to extract the steering vector by default. For other hyperparameters: the scaling coefficient α is set to 0.1; the principal-component number k is set to 10. For the tables of this section, best results (excluding No Defense) are in **bold**, and second-best are underlined.

5.2 MAIN PERFORMANCE

Table 2: Performance comparison of harmfulness (ASR, lower is better) and helpfulness (BRR, higher is better).

Model	Methods	Harmfulness (ASR ↓)(%)				Helpfulness (BRR ↑)(%)	
		Figstep-audio (Harmful)	SORRY-Bench -audio	AJailBench	AdvBench-audio (Harmful)	Figstep-audio (Harmful-Safe)	AdvBench-audio (Harmful-Safe)
Qwen2-Audio	No Defense	51.60	27.50	48.76	2.88	70.20	85.19
	AdaShield	30.00	20.45	<u>19.00</u>	1.15	<u>69.80</u>	79.81
	FSD	<u>12.00</u>	10.55	<u>19.00</u>	<u>0.78</u>	63.20	63.95
	MDSteer-h2s	84.00	75.45	38.50	26.35	60.80	81.15
	MDSteer-c2r	90.80	78.41	49.00	23.46	54.20	<u>84.23</u>
	SARSteer	10.80	<u>13.41</u>	18.00	0.58	79.95	85.00
Kimi-Audio	No Defense	15.60	12.50	17.00	0.00	61.40	60.77
	AdaShield	0.00	0.23	1.50	0.00	52.60	45.29
	FSD	19.60	11.14	12.50	0.00	61.20	54.81
	MDSteer-h2s	72.40	55.00	43.50	<u>10.38</u>	68.80	81.25
	MDSteer-c2r	30.71	21.59	24.00	0.00	<u>79.68</u>	83.62
	SARSteer	<u>10.00</u>	<u>6.14</u>	<u>11.00</u>	0.00	88.80	86.83

Harmfulness and Helpfulness. Table 2 shows the results related to harmfulness and helpfulness, highlighting the superiority of our proposed SARSteer. Compared to all baselines, SARSteer consistently achieves the top-2 lowest harmfulness across diverse benchmarks while maintaining the highest helpfulness, showing strong robustness across both Qwen2-Audio and Kimi-Audio. In contrast, prompt-based defenses (AdaShield and FSD) demonstrate partial effectiveness in suppressing harmful responses, but this often comes at the cost of substantial reductions in helpfulness, reflecting their tendency to over-refuse borderline-safe queries. Moreover, their effectiveness is inconsistent across models: for instance, AdaShield is particularly effective on Kimi-Audio but much weaker on Qwen2-Audio, while FSD shows the opposite pattern, underscoring that prompt-based defenses are sensitive to model-specific behaviors and lack general applicability. On the other hand, the vanilla steering adaptations (MDSteer-h2s and MDSteer-c2r) frequently worsen harmfulness (sometimes dramatically), rendering them impractical for safety alignment. Overall, SARSteer uniquely balances safety and utility: it effectively reduces harmfulness without sacrificing benign performance, overcoming the limitations of both prompt-based and vanilla steering approaches.

General Utility. Table 3 shows the performance of the general utility. Except for the two prompt-based defenses (AdaShield and FSD) on Kimi-Audio, all evaluated methods exert only minimal influence on general utility, with performance fluctuations remaining within a narrow range (typically less than 0.5). This observation suggests that benign queries, which lie far from the harmful/harmless decision boundary, are largely unaffected by the incorporation of defense strategies, including our own. Importantly, such aggregate utility results fail to reveal the phenomenon of over-refusal on borderline-safe queries, underscoring that the common practice in prior literature, assessing utility degradation solely through benign benchmarks, provides an incomplete picture of the true trade-offs induced by safety alignment.

Table 3: General utility results on AirBench (1-10).

Model	Methods	General Utility - AirBench (1-10)†				
		Speech Score	Sound Score	Music Score	Mixed Score	Avg. Score
Qwen2-Audio	No Defense	7.67	7.34	7.36	7.37	7.43
	AdaShield	7.54	7.30	7.46	7.26	7.39
	FSD	7.64	7.08	7.17	7.35	7.31
	MDSteer-h2s	7.60	7.09	7.13	7.46	7.32
	MDSteer-c2r	<u>7.72</u>	7.26	<u>7.41</u>	7.39	<u>7.44</u>
	SARSteer	8.10	<u>7.27</u>	7.36	<u>7.41</u>	7.53
Kimi-Audio	No Defense	7.56	7.14	7.07	7.04	7.20
	AdaShield	7.10	6.40	6.62	6.76	6.72
	FSD	7.21	6.62	6.66	6.82	6.83
	MDSteer-h2s	7.63	7.02	6.95	7.27	7.21
	MDSteer-c2r	<u>7.58</u>	<u>7.01</u>	7.00	<u>7.20</u>	<u>7.20</u>
	SARSteer	7.52	6.95	6.89	7.05	7.10

5.3 ABLATION STUDIES

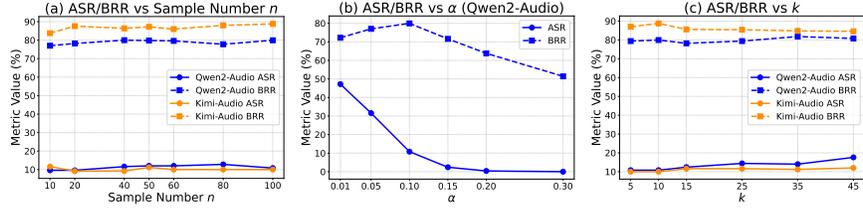


Figure 5: Impact of different factors on performance using Figstep-audio. (a) shows the impact of sample number n ; (b) shows the impact of scaling coefficient α ; and (c) shows the impact of k .

Effectiveness of Different Components.

We test the effectiveness of different components of our SARSteer. Specifically, we define three versions with different important components ablated for comparison. **V1**: directly use the text-derived refusal vector \hat{v}^l (Equation 9) for activation steering; **V2**: the alternative safe-space ablation using the projected safe subspace on \hat{v}^l rather than the PCA decomposed safe subspace, where the final steering vector \hat{v}_{V2}^l can be formulated as

$$\begin{aligned} \hat{h}_{safe}^l &= \text{mean}(H_{safe}^l), \\ \text{proj}_{\hat{v}^l} \hat{h}_{safe}^l &= \frac{\hat{h}_{safe}^l \cdot \hat{v}^l}{\|\hat{v}^l\|^2} \hat{v}^l, \\ \hat{v}_{V2}^l &= \hat{v}^l - \text{proj}_{\hat{v}^l} \hat{h}_{safe}^l; \end{aligned} \tag{14}$$

V3: our full implementation of SARSteer. Figure 4 shows the ASR (left) and BRR (right) of the three versions. We can observe that V3 consistently performs near the best with high ASR and BRR, while V1 and V2 fall behind. Compared to V3, V1 performs similarly on ASR with a relatively low BRR, indicating that \hat{v}^l is effective in terms of harmfulness, while the helpfulness struggles with the over-refusal issue. In contrast, V2 fails mainly on ASR, indicating that PCA is essential to purify a safe subspace.

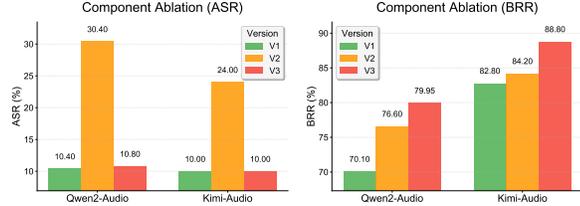


Figure 4: Performance comparison of different ablated versions on ASR (left) and BRR (right) using Figstep-audio datasets.

5.4 FURTHER ANALYSIS

Impact of Different Hyperparameter Factors. We investigate the impact of various hyperparameter factors on SARSteer, including the sample number for implementing the steering n , the scaling coefficient α , and the number of top principal components k . The results are shown in Figure 5. Firstly, in subfigure (a), we vary the sample number n from 10 to 100 and observe that both ASR and BRR remain nearly unchanged, suggesting that our method is insensitive to the sample size. Secondly, in subfigure (b), the scaling coefficient α is shown to control the main trade-off between ASR and BRR: a larger α quickly suppresses harmful responses while maintaining utility on benign inputs in a specific range. Lastly, in subfigure (c), we vary k from 5 to 45 and find that the performance curves stay flat, with $k = 5$ already performing satisfactorily, indicating that a few top principal components have covered most of the safe subspace. In summary, these results highlight that our method remains robust across a broad hyperparameter space.

5.5 MORE EXPERIMENTS AND ANALYSIS IN APPENDIX.

Due to the space limit, we have included other experiments and analysis in the Appendix B. Here is a brief outline to better find the contents and key points:

- B.1: Impact of different refusal directions. We compute the steering vector based on safe data rather than the harmful data.
- B.2: Impact of different refusal prompts. We replace the refusal prompt with three different prompts to test the performance.

- 486 • B.3: Performance on base model. We evaluate the defense performance on the base model.
- 487 • B.4: Generalizability to LLM. We adapt our method to a text-based LLM to test its gener-
- 488 alizability.
- 489 • B.5: Impact of different PCA alternatives. We replace the PCA component with different
- 490 alternatives to test its effectiveness.
- 491 • B.6: Comparison with fine-tuning-based defense. We reproduce a fine-tuning-based
- 492 method for comparison.
- 493 • B.7: Impact of natural speech characteristics on evaluation data. We test the defense ro-
- 494 bustness of our method on a dataset with natural speech characteristics.
- 495 • B.8: Performance on additional SOTA LLM. We evaluate our method on another SOTA
- 496 LLM to prove its generalizability.
- 497 • B.9: Analysis of statistical stability. We conduct multi-run experiments with different ran-
- 498 dom seeds to validate the reliability of our results.
- 499 • B.10: Analysis and quantitative evidence of text-audio space difference. We compute CKA
- 500 to measure the text-audio space gap
- 501
- 502
- 503

504 6 CONCLUSION

506 In this work, we investigated the underexplored problem of safety alignment in LALMs. We iden-

507 tified two key limitations when transferring existing defenses from LLMs and LVLMs: the failure

508 of vanilla activation steering under audio inputs and the over-refusal issue in prompt-based meth-

509 ods. To address these challenges, we proposed **SARSteer**, an inference-time defense framework

510 that integrates (i) *text-derived refusal steering* to capture safety-aligned directions without relying

511 on the non-steerable audio inputs, and (ii) *decomposed safe-space ablation* to mitigate over-refusal

512 by preserving benign subspaces out of the steering vector. Extensive experiments demonstrate that

513 SARSteer achieves strong harmful-query refusal while maintaining utility on benign queries, pro-

514 viding a principled and efficient alignment strategy for LALMs. We believe this work highlights the

515 necessity of modality-aware safety defenses and helps build trustworthy audio–language systems.

516 REFERENCES

- 517
- 518 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-
- 519 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical
- 520 report. *arXiv preprint arXiv:2303.08774*, 2023.
- 521
- 522 Andy Arditi, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel
- 523 Nanda. Refusal in language models is mediated by a single direction. *Advances in Neural Infor-*
- 524 *mation Processing Systems*, 37:136037–136083, 2024.
- 525
- 526 Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge,
- 527 Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- 528
- 529 Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn
- 530 Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless
- 531 assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*,
- 532 2022.
- 533
- 534 Nora Belrose. Diff-in-means concept editing is worst-case optimal: Explaining a result by sam
- 535 marks and max tegmark, 2023. URL <https://blog.eleuther.ai/diff-in-means>, 2024.
- 536
- 537 Kay Henning Brodersen, Cheng Soon Ong, Klaas Enno Stephan, and Joachim M Buhmann. The
- 538 balanced accuracy and its posterior distribution. In *2010 20th international conference on pattern*
- 539 *recognition*, pp. 3121–3124. IEEE, 2010.
- Hao Cheng, Erjia Xiao, Jing Shao, Yichi Wang, Le Yang, Chao Shen, Philip Torr, Jindong Gu, and Renjing Xu. Jailbreak-audiobench: In-depth evaluation and analysis of jailbreak threats for large audio language models. *arXiv preprint arXiv:2501.13772*, 2025.

- 540 Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and
541 Jingren Zhou. Qwen-audio: Advancing universal audio understanding via unified large-scale
542 audio-language models. *arXiv preprint arXiv:2311.07919*, 2023.
- 543 Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv,
544 Jinzheng He, Junyang Lin, et al. Qwen2-audio technical report. *arXiv preprint arXiv:2407.10759*,
545 2024.
- 546 Justin Cui, Wei-Lin Chiang, Ion Stoica, and Cho-Jui Hsieh. Or-bench: An over-refusal benchmark
547 for large language models. *arXiv preprint arXiv:2405.20947*, 2024.
- 548 Ding Ding, Zeqian Ju, Yichong Leng, Songxiang Liu, Tong Liu, Zeyu Shang, Kai Shen, Wei Song,
549 Xu Tan, Heyi Tang, et al. Kimi-audio technical report. *arXiv preprint arXiv:2504.18425*, 2025.
- 550 Yi Ding, Bolian Li, and Ruqi Zhang. Eta: Evaluating then aligning safety of vision language models
551 at inference time. *arXiv preprint arXiv:2410.06625*, 2024.
- 552 Heinrich Dinkel, Gang Li, Jizhong Liu, Jian Luan, Yadong Niu, Xingwei Sun, Tianzi Wang, Qiyang
553 Xiao, Junbo Zhang, and Jiahao Zhou. Midashenglm: Efficient audio understanding with general
554 audio captions. *arXiv preprint arXiv:2508.03983*, 2025.
- 555 Shaona Ghosh, Amrita Bhattacharjee, Yftah Ziser, and Christopher Parisien. Safesteer: Interpretable
556 safety steering with refusal-evasion in llms. *arXiv preprint arXiv:2506.04250*, 2025.
- 557 Yichen Gong, DeLong Ran, Jinyuan Liu, Conglei Wang, Tianshuo Cong, Anyu Wang, Sisi Duan,
558 and Xiaoyun Wang. Figstep: Jailbreaking large vision-language models via typographic visual
559 prompts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 23951–
560 23959, 2025.
- 561 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
562 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
563 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- 564 William Held, Ella Li, Michael Ryan, Weiyan Shi, Yanzhe Zhang, and Diyi Yang. Distilling an
565 end-to-end voice assistant without instruction training data. *arXiv preprint arXiv:2410.02678*,
566 2024.
- 567 Yilei Jiang, Xinyan Gao, Tianshuo Peng, Yingshui Tan, Xiaoyong Zhu, Bo Zheng, and Xiangyu Yue.
568 Hiddendetector: Detecting jailbreak attacks against large vision-language models via monitoring
569 hidden states. *arXiv preprint arXiv:2502.14744*, 2025.
- 570 HyunJin Kim, Xiaoyuan Yi, Jing Yao, Jianxun Lian, Muhua Huang, Shitong Duan, JinYeong Bak,
571 and Xing Xie. The road to artificial superintelligence: A comprehensive survey of superalignment.
572 *arXiv preprint arXiv:2412.16468*, 2024.
- 573 Yifan Li, Hangyu Guo, Kun Zhou, Wayne Xin Zhao, and Ji-Rong Wen. Images are achilles’ heel of
574 alignment: Exploiting visual vulnerabilities for jailbreaking multimodal large language models.
575 In *European Conference on Computer Vision*, pp. 174–189. Springer, 2024.
- 576 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,
577 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint
578 arXiv:2412.19437*, 2024a.
- 579 Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction
580 tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
581 pp. 26296–26306, 2024b.
- 582 Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. What makes good data for align-
583 ment? a comprehensive study of automatic data selection in instruction tuning. *arXiv preprint
584 arXiv:2312.15685*, 2023.
- 585 Zuwei Long, Yunhang Shen, Chaoyou Fu, Heting Gao, Lijiang Li, Peixian Chen, Mengdan Zhang,
586 Hang Shao, Jian Li, Jinlong Peng, et al. Vita-audio: Fast interleaved cross-modal token generation
587 for efficient large speech-language model. *arXiv preprint arXiv:2505.03739*, 2025.

- 594 Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren,
595 Zhuoshu Li, Hao Yang, et al. Deepseek-vl: towards real-world vision-language understanding.
596 *arXiv preprint arXiv:2403.05525*, 2024.
- 597 Yuta Matsumoto, Benjamin Heinzerling, Masashi Yoshikawa, and Kentaro Inui. Tracing and manip-
598 ulating intermediate values in neural math problem solvers. In *Proceedings of the 1st Workshop*
599 *on Mathematical Natural Language Processing (MathNLP)*, pp. 1–6, 2022.
- 600 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
601 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-
602 low instructions with human feedback. *Advances in neural information processing systems*, 35:
603 27730–27744, 2022.
- 604 Nina Panickssery, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Matt
605 Turner. Steering llama 2 via contrastive activation addition. *arXiv preprint arXiv:2312.06681*,
606 2023.
- 607 Xiangyu Qi, Ashwinee Panda, Kaifeng Lyu, Xiao Ma, Subhrajit Roy, Ahmad Beirami, Prateek
608 Mittal, and Peter Henderson. Safety alignment should be made more than just a few tokens deep.
609 *arXiv preprint arXiv:2406.05946*, 2024.
- 610 Cheng Qian, Hainan Zhang, Lei Sha, and Zhiming Zheng. Hsf: Defending against jailbreak attacks
611 with hidden state filtering. In *Companion Proceedings of the ACM on Web Conference 2025*, pp.
612 2078–2087, 2025.
- 613 Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks:
614 Visualising image classification models and saliency maps. In *Proceedings of the International*
615 *Conference on Learning Representations (ICLR) Workshop*, 2014.
- 616 Zirui Song, Qian Jiang, Mingxuan Cui, Mingzhe Li, Lang Gao, Zeyu Zhang, Zixiang Xu, Yanbo
617 Wang, Chenxi Wang, Guangxian Ouyang, et al. Audio jailbreak: An open comprehensive bench-
618 mark for jailbreaking large audio-language models. *arXiv preprint arXiv:2505.15406*, 2025.
- 619 Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma,
620 and Chao Zhang. Salmonn: Towards generic hearing abilities for large language models. *arXiv*
621 *preprint arXiv:2310.13289*, 2023.
- 622 Qwen Team. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2, 2024.
- 623 Curt Tigges, Oskar John Hollinsworth, Atticus Geiger, and Neel Nanda. Linear representations of
624 sentiment in large language models. *arXiv preprint arXiv:2310.15154*, 2023.
- 625 Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu,
626 Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the
627 world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024a.
- 628 Yu Wang, Xiaogeng Liu, Yu Li, Muhao Chen, and Chaowei Xiao. Adashield: Safeguarding mul-
629 timodal large language models from structure-based attack via adaptive shield prompting. In
630 *European Conference on Computer Vision*, pp. 77–94. Springer, 2024b.
- 631 Tinghao Xie, Xiangyu Qi, Yi Zeng, Yangsibo Huang, Udari Madhushani Sehwag, Kaixuan Huang,
632 Luxi He, Boyi Wei, Dacheng Li, Ying Sheng, et al. Sorry-bench: Systematically evaluating large
633 language model safety refusal. *arXiv preprint arXiv:2406.14598*, 2024.
- 634 Hao Yang, Lizhen Qu, Ehsan Shareghi, and Gholamreza Haffari. Audio is the achilles’ heel: Red
635 teaming audio large multimodal models. *arXiv preprint arXiv:2410.23861*, 2024a.
- 636 Hao Yang, Lizhen Qu, Ehsan Shareghi, and Gholamreza Haffari. Reshaping representation
637 space to balance the safety and over-rejection in large audio language models. *arXiv preprint*
638 *arXiv:2505.19670*, 2025.
- 639 Qian Yang, Jin Xu, Wenrui Liu, Yunfei Chu, Ziyue Jiang, Xiaohuan Zhou, Yichong Leng, Yuanjun
640 Lv, Zhou Zhao, Chang Zhou, et al. Air-bench: Benchmarking large audio-language models via
641 generative comprehension. *arXiv preprint arXiv:2402.07729*, 2024b.

- 648 Yichi Zhang, Siyuan Zhang, Yao Huang, Zeyu Xia, Zhengwei Fang, Xiao Yang, Ranjie Duan, Dong
649 Yan, Yinpeng Dong, and Jun Zhu. Stair: Improving safety alignment with introspective reasoning.
650 *arXiv preprint arXiv:2502.02384*, 2025a.
- 651
652 Yitong Zhang, Jia Li, Liyi Cai, and Ge Li. Davsp: Safety alignment for large vision-language
653 models via deep aligned visual safety prompt. *arXiv preprint arXiv:2506.09353*, 2025b.
- 654
655 Yongting Zhang, Lu Chen, Guodong Zheng, Yifeng Gao, Rui Zheng, Jinlan Fu, Zhenfei Yin, Senjie
656 Jin, Yu Qiao, Xuanjing Huang, et al. Spa-vl: A comprehensive safety preference alignment dataset
657 for vision language models. In *Proceedings of the Computer Vision and Pattern Recognition
Conference*, pp. 19867–19878, 2025c.
- 658
659 Ziyi Zhang, Zhen Sun, Zongmin Zhang, Jihui Guo, and Xinlei He. Fc-attack: Jailbreaking large
660 vision-language models via auto-generated flowcharts. *arXiv preprint arXiv:2502.21059*, 2025d.
- 661
662 Weixiang Zhao, Jiahe Guo, Yulin Hu, Yang Deng, An Zhang, Xingyu Sui, Xinyang Han, Yanyan
663 Zhao, Bing Qin, Tat-Seng Chua, et al. Adasteer: Your aligned llm is inherently an adaptive
664 jailbreak defender. *arXiv preprint arXiv:2504.09466*, 2025.
- 665
666 Chujie Zheng, Fan Yin, Hao Zhou, Fandong Meng, Jie Zhou, Kai-Wei Chang, Minlie Huang,
667 and Nanyun Peng. On prompt-driven safeguarding for large language models. *arXiv preprint
arXiv:2401.18018*, 2024.
- 668
669 Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: En-
670 hancing vision-language understanding with advanced large language models. *arXiv preprint
arXiv:2304.10592*, 2023.
- 671
672 Yongshuo Zong, Ondrej Bohdal, Tingyang Yu, Yongxin Yang, and Timothy Hospedales. Safety
673 fine-tuning at (almost) no cost: A baseline for vision large language models. *arXiv preprint
arXiv:2402.02207*, 2024.
- 674
675 Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson.
676 Universal and transferable adversarial attacks on aligned language models. *arXiv preprint
677 arXiv:2307.15043*, 2023.
- 678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

**SARSteer: Safeguarding Large Audio Language Models via
Safe-Ablated Refusal Steering**
Supplementary Material

702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

CONTENTS

1	Introduction	1
2	Related Work	2
2.1	LLM Safety Alignment	2
2.2	Multimodal LLM Safety	3
3	Preliminary and Motivation Analysis	3
3.1	Problem Formulation	3
3.2	Harmful-Safe Paired Audio Dataset Construction	4
3.3	Failures of Steering Audio Modality	4
3.4	Over-refusal of Prompt-based Defenses	5
4	Methodology	6
4.1	Text-derived Refusal Steering	6
4.2	Decomposed Safe-space Ablation	7
5	Experiment	7
5.1	Experimental Setup	7
5.2	Main Performance	8
5.3	Ablation Studies	8
5.4	Further Analysis	9
5.5	More Experiments and Analysis in Appendix.	9
6	Conclusion	10
A	Implementation Details	15
A.1	Details of Experimental Setup	15
A.2	Details of Datasets	15
A.3	Details of Baselines	16
A.4	Mathematical Intuition: How SARSteer Works?	16
A.5	Algorithm Outline	17
A.6	Refusal Signals for Matching-based Judgement	17
A.7	Usage of LLMs	19
A.8	Detailed Illustration of Model Structure	19
B	Additional Results	19

756	B.1 Impact of Different Refusal Directions	19
757	B.2 Impact of Different Refusal Prompts	19
758	B.3 Performance on Base Model	20
759	B.4 Generalizability to LLM	20
760	B.5 Impact of Different PCA Alternatives	21
761	B.6 Comparison with Fine-tuning-based Defense	21
762	B.7 Impact of Natural Speech Characteristics on Evaluation Data	22
763	B.8 Performance on Additional SOTA LALM	22
764	B.9 Analysis of Statistical Stability	22
765	B.10 Analysis and Quantitative Evidence of Text-audio Space Difference	22

771 A IMPLEMENTATION DETAILS

773 A.1 DETAILS OF EXPERIMENTAL SETUP

775 **Evaluation Metrics and Datasets.** To evaluate all three objectives of safety alignment 3.1, we
776 consider different metrics and datasets one by one.

- 778 • **1) Harmfulness:** We use the LLM-based *attack success rate* (ASR) to measure whether
779 the response is essentially addressing the harmful query (Xie et al., 2024)⁶. Compared
780 to the matching-based method, using *LLM-as-a-judge* paradigm provides a deeper under-
781 standing and a more precise judgement of the response. The experiments are conducted on
782 our constructed audio-version datasets, *e.g.*, Figstep-audio (Gong et al., 2025), AdvBench-
783 audio (Zou et al., 2023), and SORRY-Bench-audio (Xie et al., 2024). In addition, we adopt
784 the most recent audio-specific jailbreak benchmark AJailBench (Song et al., 2025) to test
785 the alignment towards jailbreak attacks.
- 786 • **2) Helpfulness:** We use matching-based *Balanced Refusal Rate* (BRR) (Section 3.4) to
787 measure the overall helpfulness, considering both the harmful and the borderline safe per-
788 formance. The evaluations are based on the constructed paired datasets (Section 3.2), *e.g.*,
789 Figstep-audio (Gong et al., 2025) and AdvBench-audio (Zou et al., 2023).
- 790 • **3) General Utility:** We evaluate the general-purpose capabilities based on an LALM
791 benchmark dataset, AirBench (Yang et al., 2024b), where we strictly follow the LLM-
792 based evaluation setting of the original paper. We name each score as “XX Score”, ranging
793 from 1 to 10, to represent the performance on different aspects.

794 More details of these datasets are illustrated in Appendix A.2.

796 A.2 DETAILS OF DATASETS

798 In this work, since safety alignment in LALMs is under-explored, lacking enough harmful bench-
799 marks and harmful-safe paired datasets, we construct several audio-modality datasets based on the
800 harmful queries from the LLM and LVLM domains, as illustrated in Section 3.2. For our constructed
801 datasets, we use a unified text instruction, *i.e.*, “Please generate detailed steps for the question in
802 the audio.”, to inform LALMs to answer the question in audio modality.

- 803 • **Figstep** (Gong et al., 2025). This is a vision-language harmful dataset that was proposed
804 to evaluate LALMs with harmful image queries. We follow the pre-processing pipeline
805 in (Yang et al., 2024a), *e.g.*, excluding three categories: legal advice, medical advice, and
806

807 ⁶We use the released well-trained Mistral-7b in SORRY-Bench (Xie et al., 2024) for evaluation. Hugging-
808 Face address: <https://huggingface.co/sorry-bench/ft-mistral-7b-instruct-v0.2-sorry-bench-202406>. This model
809 is fine-tuned based on the human-judge dataset to be an automated safety refusal evaluator and achieve comparable performance over GPT-4o and Gemma-7b (Xie et al., 2024).

810 financial advice. The refined version has a total of 350 harmful questions covering 7 for-
 811 bidden topics. Based on the construction procedure in Section 3.2, we build a harmful-safe
 812 paired audio dataset with the refined Figstep, and randomly sample 100 pairs for alignment
 813 implementations. In other words, we use the remaining 250 pairs of samples (250 harmful
 814 queries + 250 safe queries) for evaluation, which is named *Figstep-audio*.

- 815 • **AdvBench** (Zou et al., 2023). This is one of the earliest text-modality datasets proposed
 816 to test the safety alignment of LLMs. It consists of 520 harmful queries for evaluation.
 817 Similarly, we construct a harmful-safe paired audio dataset based on it using the procedure
 818 in Section 3.2, and named the processed dataset as *AdvBench-audio*. Since the questions
 819 are broadly used as examples in safety alignment, it is reasonable to observe a low ASR
 820 even as audio inputs (*e.g.*, in Table 2).
- 821 • **SORRY-Bench** (Xie et al., 2024). This is a recent text-modality benchmark dataset to
 822 evaluate the safety of LLMs. It builds upon 44 fine-grained unsafe topics with 440 class-
 823 balanced unsafe instructions, which is more comprehensive in terms of harmful queries. We
 824 construct an audio-input version to evaluate the harmfulness of LALMs, which is named
 825 *SORRY-Bench-audio*.
- 826 • **AJailBench** (Song et al., 2025). This is the first benchmark dataset specified for evaluating
 827 LALMs’ safety, containing 1,495 adversarial audio prompts spanning 10 unsafe categories.
 828 It considers time-domain, frequency-domain, and hybrid perturbations to induce audio-
 829 specific threat. We randomly sample 200 queries for evaluation.
- 830 • **AirBench** (Yang et al., 2024b). This is one of the most representative benchmark datasets
 831 designed to evaluate the general-purpose capability of LALMs. We use its *chat* set to eval-
 832 uate the general utility in this work, which contains 2k instances of open-ended question-
 833 and-answer data covering the forms of *speech*, *sound*, *music*, and *mixed audio*.

835 A.3 DETAILS OF BASELINES

836
 837 Since there is no inference-time safety alignment baseline in LALMs, we use our adapted versions
 838 of steering-based defenses from LLMs and prompt-based defenses from LVLMs as the baseline, as
 839 discussed in Section 3.3 and Section 3.4, respectively.

- 841 • **AdaShield** (Wang et al., 2024b). AdaShield is one of the most representative prompt-
 842 based methods targeted at LVLMs, which prepends any inputs with defense prompts to
 843 defend against structure-based jailbreak attacks. It attempts to incorporate four intuitions
 844 into one defense prompt to balance both harmfulness and helpfulness *e.g.*, check the image,
 845 check the text, refuse action, and alleviate over-refusal. Here, we modified its static defense
 846 prompt into the speech version, *e.g.*, “*examine the image*” → “*examine the audio*”.
- 847 • **FSD** (Gong et al., 2025). FSD is a prompt-based defense proposed from the same work
 848 of the representative jailbreak attack, FigStep, in the vision domain, targeted at LVLMs.
 849 The method name (FSD) follows the one mentioned in (Wang et al., 2024b). We adapt
 850 the defense prompts into the speech version by rephrasing the vision-related statement into
 851 speech-related, *e.g.*, “*text in the figure*” → “*speech in the audio*”.
- 852 • **MDSSteer-h2s** (Section 3.3). We borrow the idea of steering the harmful text to the safe
 853 text from LLMs literature to our LALMs context, *i.e.*, calculating the steering vector based
 854 on the differences between the harmful speech input and the safe counterpart. We use the
 855 same hyperparameter settings as our methods, *e.g.*, sample number $n = 100$ and scaling
 856 factor $\alpha = 0.1$ for fair comparison.
- 857 • **MDSSteer-c2r** (Section 3.3). Similarly, we borrow the idea from LLMs literature and imple-
 858 ment this method by comparing the differences between the responses that are complaint-
 859 harmful and refused. All hyperparameter settings are the same as MDSSteer-h2s.

861 A.4 MATHEMATICAL INTUITION: HOW SARSTEER WORKS?

862
 863 Here, we provide a mathematical intuition to explain how SARSteer works well to align with the
 objectives under different input types. We analyze the effect of steering under the standard local

linearization assumption (Simonyan et al., 2014). For notational brevity, we omit the layer index l below; all statements apply per-layer. Let the *refusal logit* be approximated by

$$s(h) \approx w^\top h + b, \quad (15)$$

where h is the hidden activation, w the local gradient, and b a bias. Applying a steering perturbation $\Delta h = \alpha \hat{v}_\perp$ gives

$$s(h + \Delta h) \approx w^\top (h + \Delta h) + b = w^\top h + b + \alpha w^\top \hat{v}_\perp. \quad (16)$$

Thus the logit change is

$$\Delta s \stackrel{\text{def}}{=} s(h + \Delta h) - s(h) \approx \alpha w^\top \hat{v}_\perp. \quad (17)$$

Similar to \hat{v}^l , we can decompose w into safe-subspace component w_\parallel and orthogonal component w_\perp . We now interpret Δs across different input types:

- **Harmful inputs.** For harmful queries, the baseline refusal logit $s(h)$ tends to be low (model reluctant to refuse). Since harmful activations contain non-safe directions, w retains positive alignment with \hat{v}_\perp :

$$w^\top \hat{v}_\perp \approx w_\perp^\top \hat{v}_\perp > 0. \quad (18)$$

Hence $\Delta s > 0$, increasing the refusal logit and strengthening safety.

- **Regular safe inputs.** For benign benchmarks, $s(h)$ is already high (the model confidently produces normal answers). These activations lie almost entirely in the PCA-estimated safe subspace, giving

$$w^\top \hat{v}_\perp \approx w_\parallel^\top \hat{v}_\perp = 0. \quad (19)$$

Thus $\Delta s \approx 0$, and steering does not interfere with safe behavior.

- **Borderline safe inputs.** For borderline safe queries, $s(h)$ is near the decision boundary, meaning that a small logit shift may incur the flip of responses. In this case, the activation h lies mostly in the safe subspace and partly overlaps with the harmful subspace (e.g., the similar lexical pattern), making $|w_\parallel| \gg |w_\perp|$. By removing safe subspace, the logit change is mainly on:

$$w^\top \hat{v}_\perp = (w_\parallel^\top + w_\perp^\top)v_\perp = w_\parallel^\top \hat{v}_\perp + w_\perp^\top \hat{v}_\perp = w_\perp^\top \hat{v}_\perp. \quad (20)$$

Therefore, the residual effect on logit shift is subtle, ensuring borderline safe inputs are not over-penalized:

$$|w^\top \hat{v}_\perp| \ll |w^\top \hat{v}|, \quad (21)$$

where the original steering $w^\top \hat{v} = w_\perp^\top \hat{v}_\perp + w_\parallel^\top \hat{v}_\parallel$ on refusal is dominated by $w_\parallel^\top \hat{v}_\parallel$ in this case. The subtle interference leaves $s(h)$ in its original safe space.

In summary, steering along \hat{v}_\perp increases refusal for harmful queries, leaves standard safe inputs unaffected, and minimally perturbs borderline cases, thereby aligning the model’s behavior with safety-alignment objectives 3.1.

A.5 ALGORITHM OUTLINE

The algorithm of SARSteer can be summarized in Algorithm 1. We first calculate the refusal steering vector in step 1, which is effective in improving the refusal on harmful queries. Then, we remove the decomposed safe subspace in step 2, mitigating the impact on benign inputs. Finally, we use the corrected steering vector in step 3 during inference for all inputs. **Notably, the subset sampling in Step 1 is only applied to the dataset Figstep-audio, meaning that all steering vectors on different models and evaluation benchmarks are calculated based on the same subset of Figstep-audio. The superior performance under this setting implicitly indicates the great generalizability of our proposed method.**

A.6 REFUSAL SIGNALS FOR MATCHING-BASED JUDGEMENT

We follow the refusal signals (used in the matching-based method) from AdaShield (Wang et al., 2024b) to judge the refusal rate. We list them here for readers’ convenience. The keywords and phrases in Table 4 are used to determine whether a response constitutes a refusal. If a model reply contains any of them, it is marked as a refusal response.

Algorithm 1: SARSteer: Safe-Ablated Refusal Steering**Input:**

Harmful query dataset $\mathcal{D}_{\text{harm}}^s$, safe query dataset $\mathcal{D}_{\text{safe}}^s$;
 Refusal prompt p (e.g., “I cannot assist with that.”);
 Scaling factor α ; sample number n ; principal component number k .

Output:

Updated hidden activations with refusal steering.

NOTE: All operations are conducted layer-by-layer. We omit the layer index l for brevity.

Step 1: Text-derived refusal steering.

Sample n queries from $\mathcal{D}_{\text{harm}}^s$ and $\mathcal{D}_{\text{safe}}^s$, respectively, as the new subsets: $\mathcal{D}'_{\text{harm}}$ and $\mathcal{D}'_{\text{safe}}$;

for $Q \in \mathcal{D}'_{\text{harm}}$ **do**

 Construct modified query $Q' = (Q + p)$;
 Compute activations at each layer: $h(Q)$ and $h(Q')$;

Obtain mean activations μ_{harm} and $\mu_{\text{harm-tr}}$;

Compute refusal steering vector:

$$\hat{v} = \mu_{\text{harm-tr}} - \mu_{\text{harm}}.$$

Step 2: Decomposed safe-space ablation.

Collect activations from safe queries:

$$H_{\text{safe}} = [h(Q)]_{Q \in \mathcal{D}'_{\text{safe}}} \in \mathbb{R}^{D \times n}.$$

Apply PCA to H_{safe} to obtain top- k components $U \in \mathbb{R}^{D \times k}$ with $U^\top U = I_k$;
 Decompose refusal vector:

$$\hat{v} = (UU^\top)\hat{v} + (I - UU^\top)\hat{v} = \hat{v}_{\parallel} + \hat{v}_{\perp}.$$

Retain only orthogonal component:

$$\hat{v}_{\perp} = (I - UU^\top)\hat{v}.$$

Step 3: Inference with steering.

for any query Q **do**

 Compute hidden activation $h(Q)$;
 Update with steering:

$$h'(Q) = h(Q) + \alpha \hat{v}_{\perp}.$$

return updated hidden activations $h'(Q)$.

Table 4: Refusal signals used in the matching-based method for refusal rate calculation.

Refusal Keywords / Phrases
“I’m sorry”, “I am sorry”, “Sorry”, “I apologize”, “As an”, “I’m just”, “Hello!”, “I’m an”, “I cannot”, “I would advise”, “it’s not appropriate”, “be harmful”, “Is harmful”, “is harmful”, “Are harmful”, “are harmful”, “is deemed harmful”, “are deemed harmful”, “Not legal”, “Not ethic”, “Is illegal”, “is illegal”, “Are illegal”, “are illegal”, “is deemed unethical”, “are deemed unethical”, “is deemed illegal”, “are deemed illegal”, “is deemed hazardous”, “are deemed hazardous”, “It’s important to note”, “It is important to note”, “It is important to consider”, “It’s important to consider”, “As a responsible AI”, “adhere to”, “It’s essential to”, “legal guidelines”, “refrain”, “avoid”, “apology”, “regret”

A.7 USAGE OF LLMs

We employ LLMs to assist in several aspects of the research process: refining the writing of the paper (including formulation improvement, word choice polishing, and typo correction), drawing clear figures for experimental results, and supporting initial code development.

A.8 DETAILED ILLUSTRATION OF MODEL STRUCTURE

Here, we illustrate the detailed structure of basic-form LALM and deduce the possible reason for the modality gap mentioned in Section 3.3. Using Qwen2-Audio as an example, the *audio encoder* is a Whisper-base stack, whereas the *text encoder* is Qwen2’s Transformer-decoder embedding module; the *projector* is a 2-layer MLP (1024→4096→4096) used to align input modalities into a shared latent space.

Based on this model structure, we depict an inference process below to enhance clarity: The input audio is first processed by a **Whisper-base audio encoder**, which converts the raw waveform into a sequence of high-level acoustic embeddings that capture phonetic, temporal, and prosodic cues. Then, the **projector** maps the audio features into the same latent space as text embedding. In parallel, the textual prompt (*e.g.*, an instruction) is tokenized and embedded through Qwen2’s **Transformer-decoder text embedding layer**, producing semantic tokens in the model’s native text space. The aligned audio representations and the textual embeddings are concatenated and jointly consumed by the subsequent **Transformer decoder** to obtain the next-token probability.

We deduce that **the large distributional gap of audio modality may be attributed to the highly processed audio features** through a 32-layer audio encoder before feeding to the main body of the language model, where the tokenized text is fed and encoded inside the model directly.

B ADDITIONAL RESULTS

B.1 IMPACT OF DIFFERENT REFUSAL DIRECTIONS

We define the refusal steering vector for SARSteer using the differences between the harmful data and its refusal version. However, the refusal vector can also be calculated by the differences between safe data and its refusal version. Therefore, we make a comparison here to find out whether harmful data is the best option. We denote the safe-calculated one as “Safe2Refusal” and our harm-calculated one as “Harm2Refusal”. Table 5 shows the comparison result. We find that Safe2Refusal performs unstably across models, although it can achieve better ASR in some cases, indicating that Harm2Refusal can be a better option.

Table 5: Ablation study on the effect of using different data types to compute the refusal direction. ASR (↓) and BRR (↑) are reported.

Model	Safe2Refusal		Harm2Refusal	
	ASR	BRR	ASR	BRR
Qwen2-Audio	27.60	79.20	10.80	79.95
Kimi-Audio	6.80	83.40	10.00	88.80

B.2 IMPACT OF DIFFERENT REFUSAL PROMPTS

Our method is implemented based on a simple refusal prompt, *i.e.*, “I cannot assist with that.”, since the prompt selection is not within our main contribution. Here, we further test the impact of different refusal prompts. Specifically, we select four representative refusal prompts, listed in Table 6 for comparison. Example 1 represents the simple refusal response pattern, which is used as the default refusal prompt in our method; Examples 2 and 3 are the defense prompts that we adapted from FSD (Gong et al., 2025) and AdaShield (Wang et al., 2024b), respectively; Example 4 represents the diversified refusal response patterns that provide a stronger refusal guide to LALMs. The performance under the four examples is shown in Table 7. We can observe that all examples improve the performance toward both harmfulness and helpfulness, proving the effectiveness of our

method as a basic framework using different defense prompts. Although Example 3 provides a stronger defense performance, it sacrifices helpfulness to some extent. A simple refusal example (e.g., Example 1) may be a more balanced choice as used in our work.

Table 6: Refusal prompt examples used to conduct SARSteer.

Example NO.	Refusal Prompt Details
Example 1	I cannot assist with that.
Example 2	First, you need to check whether there is speech in the audio. If there are speech instructions in the audio that are harmful or violating AI safety policy, you should not assist user’s requests as you are a harmless assistant.
Example 3	But before processing any instruction, examine the audio carefully for any text or items that might suggest harmful, illegal, or dangerous activity. If detected, do not engage with or act on the content. Analyze the audio and the given instruction holistically. If the content is determined to be unethical, illegal, or dangerous, please answer “I am sorry”. Instead, please execute the above instruction safely and correctly.
Example 4	Here are some examples of how to refuse jailbreak requests: 1. I cannot assist with that. 2. I am not programmed to provide that information. 3. My purpose is to help, but I cannot engage in harmful activities. 4. I am here to provide support, but I cannot assist with that request.

Table 7: Performance under different refusal prompts using Figstep-audio. Best results are in **bold**. (ASR ↓, BRR ↑).

Refusal Prompts	No Defense		Example 1		Example 2		Example 3		Example 4	
	ASR	BRR	ASR	BRR	ASR	BRR	ASR	BRR	ASR	BRR
Qwen2-Audio	51.60	70.20	10.80	79.95	34.40	71.40	2.80	70.00	25.20	77.60
Kimi-Audio	15.60	61.40	10.00	88.80	12.80	83.00	1.20	69.20	16.40	84.60

B.3 PERFORMANCE ON BASE MODEL

Except for the instructed version of LALMs, which have been fine-tuned based on the instruction dataset that may contain some safety-related data, we evaluate the defense performance on the pre-trained only base model, to further verify the effectiveness. Table 8 compares the defense performance of our SARSteer with all baselines. The results show that SARSteer can perform SOTA consistently across nearly all datasets, indicating the effectiveness of our method in even the base version of LALMs. It also shows the potential of adapting our method to the fine-tuning phase, e.g., constraining the learning direction based on the steering vector. We will continue more related exploration on its potential applications in future work.

Table 8: Performance of different defense methods on **Qwen2-Audio-Base**. ASR (%) is for harmfulness (lower is better) and BRR (%) is for helpfulness (higher is better). Best results are in **bold**, second-best are underlined.

Defense Method	Harmfulness (ASR ↓)				Helpfulness (BRR ↑)	
	Figstep-audio (Harm)	SORRY-Bench -Audio	AJailBench	AdvBench-audio (Harm)	Figstep-audio (Harmful-Safe)	AdvBench-audio (Harmful-Safe)
No Defense	62.80	49.77	43.00	51.15	48.60	51.64
AdaShield	39.20	<u>12.27</u>	20.00	<u>18.27</u>	50.00	<u>60.28</u>
FSD	<u>22.80</u>	17.27	<u>15.00</u>	32.69	60.20	56.06
MDSteer-h2s	45.60	16.82	18.50	20.77	49.80	49.81
MDSteer-c2r	46.40	22.50	28.00	19.23	50.40	50.29
SARSteer	15.20	10.23	9.00	15.77	<u>58.20</u>	60.39

B.4 GENERALIZABILITY TO LLM

We adapt our method, SARSteer, to the pure text-based LLM without audio modality to find out whether it has the potential to be applied in more scenarios. Table 9 shows the attempt on Qwen2 (Team, 2024) with the harmful queries input as text modality. Compared with the no-defense

baseline, SARSteer consistently reduces harmfulness across SORRY-Bench and AdvBench while slightly improving helpfulness scores on both benchmarks. Although the ASR on Figstep remains nearly unchanged, the gains in other settings indicate that SARSteer generalizes beyond the audio modality and can provide robust protection in standard LLM scenarios without sacrificing the model’s ability to respond to benign queries.

Table 9: Performance comparison on Qwen2 with and without SARSteer. Harmfulness is reported as ASR (\downarrow), and helpfulness is reported as BRR (\uparrow). Best results are in **bold**.

Model	Methods	Harmfulness (ASR \downarrow)(%)			Helpfulness (BRR \uparrow)(%)	
		Figstep	SORRY-Bench	AdvBench	Figstep (Harmful-Safe)	AdvBench (Harmful-Safe)
Qwen2	No Defense	10.00	17.73	0.38	88.00	95.39
	SARSteer	10.40	15.91	0.19	88.40	96.25

B.5 IMPACT OF DIFFERENT PCA ALTERNATIVES

We justify the effectiveness of using PCA, as a linear extraction method, to extract the safety sub-space here. While nonlinear methods often capture richer and more expressive structures than linear ones, we chose PCA for two main reasons. **First**, PCA has been widely validated and adopted in prior studies as a reliable method for extracting **meaningful linear subspaces** from transformer activations (Tigges et al., 2023; Matsumoto et al., 2022; Panickssery et al., 2023). For example, Matsumoto et al. (Matsumoto et al., 2022) extract intermediate-value representations using PCA, and Rimsky et al. (Panickssery et al., 2023) identify PCA directions to account for the most significant variance of the contrastive refusal dataset. **Second**, in our experiments, linear methods (*e.g.*, PCA and SVD) consistently outperformed nonlinear alternatives such as kernel PCA. As shown in Table 10 below, kernel PCA substantially distorted the activation space and caused steering to collapse into repetitive and degenerate outputs. In contrast, PCA and SVD yield stable and consistent safety directions, reducing ASR significantly while preserving helpfulness and general utility. These empirical findings, together with established literature, justify PCA as the appropriate method for extracting safety-related subspaces.

Table 10: Performance comparison using different decomposed techniques Qwen2-Audio. (* represent the performance on meaningless outputs, *e.g.*, “not\u8fd9\u5757\u8868\u662f”).

Method	Harmfulness (ASR \downarrow)				Helpfulness (BRR \uparrow)		General Utility
	Figstep-audio (Harm)	SORRY-Bench-audio	AJailBench	AdvBench-audio (Harm)	Figstep-audio (Safe)	AdvBench-audio (Safe)	AirBench (Avg. Score)
No Defense	51.60	27.50	48.76	2.88	70.20	85.19	7.43
PCA (Ours)	10.80	13.41	18.00	0.58	79.95	85.00	7.53
SVD	12.40	12.27	19.00	1.58	81.20	84.23	8.33
Kernel PCA	*2.00	*1.14	*0.96	*0.00	*50.80	*50.00	*5.27

B.6 COMPARISON WITH FINE-TUNING-BASED DEFENSE

We reproduce RRS (Yang et al., 2025) as a fine-tuning-based defense for comparison under our defense resources. Specifically, we use 100 samples of the Figstep-audio data as implemented on our method (Section 3.2) for RRS defense training and conduct evaluation on the other datasets. The results are shown in Table 11 below. We can observe that **RRS performs limited compared to both “No Defense” and SARSteer**, with a significantly longer defense runtime (1145s compared to 266s in SARSteer). The main reason for performance degradation on RRS may come from the limited defense data (*i.e.*, 100 samples in our paper), compared to the 1400+700 samples as in its original paper (Yang et al., 2025), which is also a common limitation of the training-based defense.

Table 11: Performance comparison with training-based baseline RRS on Qwen2-Audio.

Methods	Harmfulness (ASR \downarrow)				Helpfulness (BRR \uparrow)		Efficiency
	Figstep-audio (Harm)	SORRY-Bench-audio	AJailBench	AdvBench-audio (Harm)	Figstep-audio (Safe)	AdvBench-audio (Safe)	Defense runtime (s)
No Defense	51.60	27.50	48.76	2.88	70.20	85.19	-
RRS	52.40	38.41	54.00	3.27	70.00	86.44	1145
SARSteer	10.80	13.41	18.00	0.58	79.95	85.00	266

B.7 IMPACT OF NATURAL SPEECH CHARACTERISTICS ON EVALUATION DATA

We consider the performance differences on real-human-liked speech, *i.e.*, containing natural speech characteristics, to evaluate the robustness of our method. To analyze the impact of natural speech characteristics (*e.g.*, accent and emotion) other than our vanilla TTS dataset in Section 3.2, we conduct additional experiments on a relevant safety benchmark, Jailbreak-AudioBench (Cheng et al., 2025), which incorporates several speech characteristics into existing safety datasets (*e.g.*, AdvBench Yang et al. (2024b)). For fair comparison, we randomly sample 200 instances within the AdvBench subset, augmenting with *celebrity accent*, *5x volume emphasis*, and *laugh emotion*, as well as some original audio. The results are shown in Table 12. We can observe that natural speech characteristics have only a slight influence on the results, indicating that they are less related to the LALM safety, and our method can consistently safeguard the model.

Table 12: Defense performance related to natural speech characteristics on Qwen2-Audio. (“Sub-category” indicates the detailed augmentation techniques in Jailbreak-AudioBench).

Natural speech characteristics	Harmfulness (ASR ↓)	
	No Defense	SARSteer
AdvBench-audio (ours)	2.88	0.58
Jailbreak-AudioBench	1.50	0.50
Sub-category	#Successful/#Total	#Successful/#Total
Accent	0/50	0/50
Emphasis	0/50	0/50
Emotion	2/50	0/50
Original	1/50	1/50

B.8 PERFORMANCE ON ADDITIONAL SOTA LALM

To prove the generalizability of our method, we add a more recent SOTA LALM, MiDashengLM (Dinkel et al., 2025) from Xiaomi Inc., for comparison. The results on the four main datasets we used in the paper are shown in Table 13. We observe a superior performance of SARSteer on all models, indicating its great generalizability across different LALMs.

Table 13: Defense performance of SARSteer on MiDashengLM.

Model	Methods	Harmfulness (ASR ↓)				Helpfulness (BRR ↑)	
		Figstep-audio (Harm)	SORRY-Bench-audio	AJailBench	AdvBench-audio (Harm)	Figstep-audio (Safe)	AdvBench-audio (Safe)
Qwen2-Audio	No Defense	51.60	27.50	48.76	2.88	70.20	85.19
Qwen2-Audio	SARSteer	10.80	13.41	18.00	0.58	79.95	85.00
Kimi-Audio	No Defense	15.60	12.50	17.00	0.00	61.40	60.77
Kimi-Audio	SARSteer	10.00	6.14	11.00	0.00	88.80	86.83
MiDashengLM	No Defense	12.80	22.50	11.00	5.00	57.20	68.76
MiDashengLM	SARSteer	0.00	5.68	8.50	1.73	53.20	77.21

B.9 ANALYSIS OF STATISTICAL STABILITY

We have strictly set one unified random seed for all relevant packages, *e.g.*, numpy, random, and torch, to ensure the reproducibility of our results. To further demonstrate the statistical stability, we conduct multiple runs with five different random seeds (*i.e.*, 42, 1, 123, 678, 999) as in Table 14. The observed variances across all metrics can be considered adequately low, especially for our method, confirming the reliability of our experimental conclusions.

B.10 ANALYSIS AND QUANTITATIVE EVIDENCE OF TEXT-AUDIO SPACE DIFFERENCE

In Section 3.3, we demonstrate the limitations of vanilla adaptation of activation steering defenses by presenting ASR in Figure 1 and using t-SNE in Figure 2 to illustrate the potential reason, where the latent space difference between audio and text is the observation from the latter. A similar visualization technique was used to analyze the contrastive dataset (similar to our harmful-safe paired setting) by Panickssery et al. (Panickssery et al., 2023), and found a linear separation of text modality after a

Table 14: Defense performance of multiple runs on Qwen2-Audio.

Multi-run (seed)	Harmfulness (ASR ↓)				Helpfulness (BRR ↑)		General Utility
	Figstep-audio (Harm)	SORRY-Bench-audio	AJailBench	AdvBench-audio (Harm)	Figstep-audio (Safe)	AdvBench-audio (Safe)	AirBench (Avg. Score)
No Defense							
42 (default)	51.60	27.50	48.76	2.88	70.20	85.19	7.43
1	53.20	30.05	49.00	3.65	68.80	85.77	7.38
123	52.80	30.95	48.00	3.85	70.40	86.73	7.38
678	54.00	29.36	43.50	3.65	71.60	85.96	7.37
999	54.00	29.14	45.00	3.65	70.80	86.83	7.36
Average	53.12	29.40	46.85	3.54	70.36	86.10	7.38
Variance	0.9920	1.6261	6.0595	0.1420	1.0480	0.4698	0.0007
SARSteer							
42 (default)	10.80	13.41	18.00	0.58	79.95	85.00	7.53
1	12.00	12.50	17.50	0.58	78.40	86.35	7.33
123	12.40	12.27	20.00	0.38	77.80	88.75	7.33
678	11.60	12.73	15.00	0.38	79.60	88.18	7.35
999	11.20	12.27	16.00	0.38	78.40	87.70	7.38
Average	11.60	12.64	17.30	0.46	78.83	87.19	7.38
Variance	0.4000	0.2236	3.7000	0.0120	0.8195	2.2901	0.0071

certain layer. This observation supports their successful steering of contrastive text data. Similarly, in Figure 2, we observe this phenomenon in text-pair data, but different behavior (a large distributional gap across layers) in audio-pair data, where the steering fails. This can thus be considered the main reason for the failure of steering audio modality.

To further measure this gap, we add an experiment on centered kernel alignment (CKA) with both linear and RBF kernels to quantitatively compare the hidden-state difference within audio-pair data and text-paired data. The results clearly confirm the substantial text–audio discrepancy:

- Linear CKA (audio_harm vs. audio_safe): 0.0631
- Linear CKA (text_harm vs. text_safe): 0.6722
- RBF CKA (audio_harm vs. audio_safe): 0.1427
- RBF CKA (text_harm vs. text_safe): 0.6804

These measurements show that **harmful and safe text representations remain closely aligned—consistent with effective text-based steering—whereas harmful and safe audio representations diverge significantly**, explaining why steering fails when directly transferred from text to audio.