FORGETTING WHO, PRESERVING WHAT: SPEAKER-IDENTITY UNLEARNING IN AUDIO-LANGUAGE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Audio-language models jointly encode *who* is speaking and *what* is said, creating a practical deletion request: remove a target speaker-identity association from a deployed model while preserving the linguistic value of the same recordings. Existing sequence-level unlearning objectives do not distinguish identity-bearing acoustic evidence from content-bearing evidence and can therefore trade stronger forgetting for broad representation damage. To address this challenge, we introduce CAUL, a selective model-side unlearning framework for AudioLLMs. It consists of three stages: (i) Intervention-Aligned Support Localization, which identifies acoustic channels whose neutralization suppresses target-identity evidence more than linguistic content; (ii) Cross-Utterance Evidence Aggregation, which aggregates channel-level evidence across utterances to obtain stable speaker-specific support; and (iii) Sparse Local Association Editing, which neutralizes the selected acoustic support toward a retain-derived reference and jointly optimizes target-identity unlikelihood with retain KL to preserve non-target behavior. Alongside CAUL, we establish VoxSU-Bench, a speaker-unlearning benchmark built from VoxCeleb datasets. Experimental results on VoxSU-Bench show that CAUL achieves strong speaker-identity forgetting while preserving high speech-content retention. To the best of our knowledge, this is the first study of speaker-identity association unlearning for AudioLLMs. We will release the code and benchmark in the near future.

1 INTRODUCTION

Audio-language models (AudioLLMs) support speech transcription, audio reasoning, and spoken interaction (Ghosh et al., 2026; Xu et al., 2025). Models such as AudioPaLM, LTU, SALMONN, Qwen-Audio, and Qwen2-Audio combine audio processing with language modeling to perform a range of speech and audio tasks (Rubenstein et al., 2023; Gong et al., 2024; Tang et al., 2024; Chu et al., 2023; 2024). Speech, however, conveys both linguistic content and speaker identity through cues such as timbre, pitch, prosody, accent, and habitual articulation. Recent studies show that AudioLLMs can identify public figures from voice recordings even after excluding samples whose transcripts alone allow correct identification (He et al., 2025). Related work also investigates the inference of private attributes from audio using multimodal language-model agents (Wang et al., 2025). The ability to identify speakers raises a selective deletion problem: once a model has acquired a voice–identity association, how can it forget *who* spoke without losing *what* was said?

This problem naturally arises in consent withdrawal, account deactivation, and personalization updates, and is closely related to the right to be forgotten (Zhang et al., 2025). We study this problem under a controlled identity-acquisition setting, which provides a reproducible proxy for voice–identity associations learned during prior training. We call this objective *content-preserving speaker-identity association unlearning*: the waveform itself remains unchanged, so the model must preserve the linguistic content of an utterance while breaking its association with the target identity.

Machine unlearning seeks to remove the influence of designated training data while reducing the cost of retraining from scratch (Cao & Yang, 2015; Ginart et al., 2019; Bourtoule et al., 2021; Sekhari et al., 2021; Xu et al., 2023). For language models, recent work studies both knowledge

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

suppression and multi-dimensional unlearning evaluation (Zhang et al., 2024; Fan et al.; Shi et al., 2025). Related efforts extend unlearning to multimodal models. SIU studies the unlearning of visual concept recognition (Li et al., 2024a), and MMUnlearner targets visual information associated with an entity while preserving its textual knowledge (Huo et al., 2025). MANU and MIP-Editor further investigate selective intervention through neuron pruning and influential neuron path editing, respectively (Liu et al., 2025; Li et al., 2026).

In speech and audio, existing work studies sample- and class-level unlearning (Cheng & Amiri, 2025), speaker-data removal in spoken-language understanding (Koudounas et al., 2025), voice-biometric unlearning (Pathak et al., 2025), and target-voice imitation in zero-shot text-to-speech (Kim et al., 2025). Other studies examine structural editing in audio models and unlearning in LLM-based speech recognition (Mason-Williams et al., 2025; Liu, 2025). However, these settings do not address a generative AudioLLM that must selectively remove a learned voice-to-identity mapping while keeping the linguistic content of the same, unchanged waveform usable. This couples the forgetting and preservation objectives within a single speech input.

AudioLLMs make this problem especially challenging for two reasons. First, phonetic and speaker-related information can be entangled in learned speech representations (Cheng & Amiri, 2025). Global behavioral suppression may therefore trade off identity removal against speech understanding. This motivates content-aware selection of identity-related acoustic support. Second, identity elicitation can depend on the query context (He et al., 2025). More broadly, unlearning evaluations show that suppressing an observed response need not establish that the underlying knowledge has been removed (Sun et al., 2025).

We introduce CAUL, a selective model-side unlearning framework with three components. (i) *Intervention-Aligned Support Localization* identifies acoustic channels whose neutralization suppresses target-identity evidence while limiting disruption to linguistic content. (ii) *Cross-Utterance Evidence Aggregation* combines channel-level evidence across utterances to obtain stable speaker-specific support. (iii) *Sparse Local Association Editing* neutralizes the selected acoustic support toward a retain-derived reference and jointly optimizes target-identity unlikelihood and KL regularization on retain data.

To make the deletion target observable and the evaluation reproducible, we establish VoxSU-Bench, a benchmark built from VoxCeleb1 (Nagrani et al., 2017), with controlled identity acquisition, source-video-disjoint splits, and held-out identity-query families. VoxSU-Bench evaluates speaker-identity association unlearning along four complementary dimensions: forgetting efficacy, generality to unseen elicitation patterns, linguistic-content preservation, and robustness to acoustic transformations.

Experiments on VoxSU-Bench show that CAUL achieves strong speaker-identity forgetting while maintaining high speech-content retention, yielding a favorable forgetting–content preservation trade-off across different AudioLLM backbones. Our contributions are fourfold:

❶ **Problem Formulation.** To the best of our knowledge, we are the first to formulate and systematically study *content-preserving speaker-identity association unlearning* in generative AudioLLMs, where the goal is to forget who spoke while preserving what was said.

❷ **Selective Acoustic Unlearning.** We introduce CAUL, which localizes identity-related acoustic support while accounting for linguistic content, aggregates evidence across utterances, and selectively edits the resulting support.

❸ **VoxSU-Bench.** We establish VoxSU-Bench, a reproducible VoxCeleb1-based benchmark with controlled identity acquisition, source-video-disjoint evaluation, and held-out identity-query families for evaluating forgetting efficacy, generality, content preservation, and robustness.

❹ **Experimental Evaluation.** We conduct systematic experiments across Qwen2-Audio and different scales of SALMONN, showing that CAUL achieves strong identity forgetting while preserving linguistic content across multiple AudioLLM backbones.

2 Related Work

Speech and Audio-Language Models. Audio-language models combine audio processing with language modeling to support speech and audio understanding. AudioPaLM integrates speech and

text modeling (Rubenstein et al., 2023), while LTU connects audio perception with language-based reasoning (Gong et al., 2024). SALMONN integrates speech and general-audio encoders through a Q-Former (Tang et al., 2024), and Qwen-Audio and Qwen2-Audio extend instruction following to diverse audio tasks (Chu et al., 2023; 2024). Beyond linguistic content, these models can exploit information about the speaker. He et al. (2025) show that AudioLLMs can identify public figures after excluding recordings whose transcripts alone permit correct identification. Wang et al. (2025) investigate private-attribute inference from audio, while selective-hearing work studies privacy-sensitive acoustic information in AudioLLMs (Zhan et al., 2026). Speaker anonymization addresses identity privacy by modifying the waveform while seeking to preserve speech utility (Fang et al., 2019; Tomashenko et al., 2022). Our setting keeps the waveform unchanged and updates the model to remove a learned voice–identity association.

Machine Unlearning. Machine unlearning aims to remove designated training influence while preserving model utility. Early work develops efficient deletion procedures, selective forgetting, sharded retraining, and certified removal (Cao & Yang, 2015; Ginart et al., 2019; Golatkar et al., 2020; Bourtoule et al., 2021; Guo et al., 2020; Sekhari et al., 2021). Approximate approaches include teacher–student optimization in SCRUB (Kurmanji et al., 2023), negative preference optimization in NPO and SimNPO (Zhang et al., 2024; Fan et al.), and representation intervention in RMU (Li et al., 2024b). Related work edits factual associations with ROME (Meng et al., 2022), studies unlearning from a mechanistic perspective (Cheng et al., 2026), and erases concepts from diffusion models (Gandikota et al., 2023). For multimodal language models, SIU studies visual concept recognition unlearning (Li et al., 2024a), while MMUnlearner removes entity-related visual information and preserves corresponding textual knowledge (Huo et al., 2025). MANU and MIP-Editor investigate selective neuron pruning and influential neuron path editing, respectively (Liu et al., 2025; Li et al., 2026). In the audio domain, existing studies investigate speech sample- and class-level unlearning, spoken-language understanding, speech recognition, audio-model components, speaker-related representations, and target-voice generation in zero-shot TTS (Cheng & Amiri, 2025; Koudounas et al., 2025; Liu, 2025; Mason-Williams et al., 2025; Pathak et al., 2025; Kim et al., 2025). Waveform anonymization instead modifies the input speech and evaluates privacy against downstream utility (Fang et al., 2019; Tomashenko et al., 2022). Nevertheless, existing work does not directly study a generative AudioLLM that removes an acquired *speaker-identity association* while preserving the linguistic content of the same, unchanged waveform.

3 Problem Formulation

In this work, we study the unlearning of the learned association between a target speaker's voice and identity in AudioLLMs, rather than removing the speaker's acoustic characteristics or all speaker-related representations. Such acoustic information may also contribute to linguistic-content understanding and general audio-language capabilities. Specifically, we consider a shared-model setting in which the model has acquired a voice–identity association through controlled identity acquisition. Let M_0 denote a pretrained AudioLLM and M_{id} the model after identity acquisition. All unlearning methods start from the same M_{id} and produce an unlearned model $M_{\hat{\theta}}$.

Formally, let $\mathcal{F}$, $\mathcal{R}$, and $\mathcal{N}$ denote the forget, retain, and neutral speaker sets, respectively. Given an utterance x spoken by speaker s, its linguistic content c, and an identity query q, our goal for $s \in \mathcal{F}$ is to remove the learned association between x and the target identity s, while preserving the linguistic content of the same utterance and the general audio-language capabilities of the model. At the same time, the non-target behavior of the model on retain speakers should remain largely unchanged. Training and evaluation utterances and identity queries are kept disjoint, so that successful unlearning should generalize to unseen utterances and query formulations.

We therefore define *speaker–identity association unlearning* as suppressing the support that a target speaker's voice provides for its learned identity while preserving the linguistic content of the same speech and non-target model behavior. Abstractly, we formulate this goal as

$$\boxed{M_{\hat{\theta}} = \arg\min_{M_\theta} \left[\mathcal{L}_{\mathrm{assoc}}\left(M_\theta; \mathcal{F}\right) + \alpha_{\mathrm{ct}}\mathcal{L}_{\mathrm{content}}\left(M_\theta; \mathcal{F}\right) + \alpha_{\mathrm{ret}}\mathcal{L}_{\mathrm{retain}}\left(M_\theta; \mathcal{R}\right)\right].} \tag{1}$$

Here, $\mathcal{L}_{\mathrm{assoc}}$ suppresses the target voice–identity association, $\mathcal{L}_{\mathrm{content}}$ preserves the linguistic content of the same speech, and $\mathcal{L}_{\mathrm{retain}}$ constrains behavioral drift on non-target speakers.

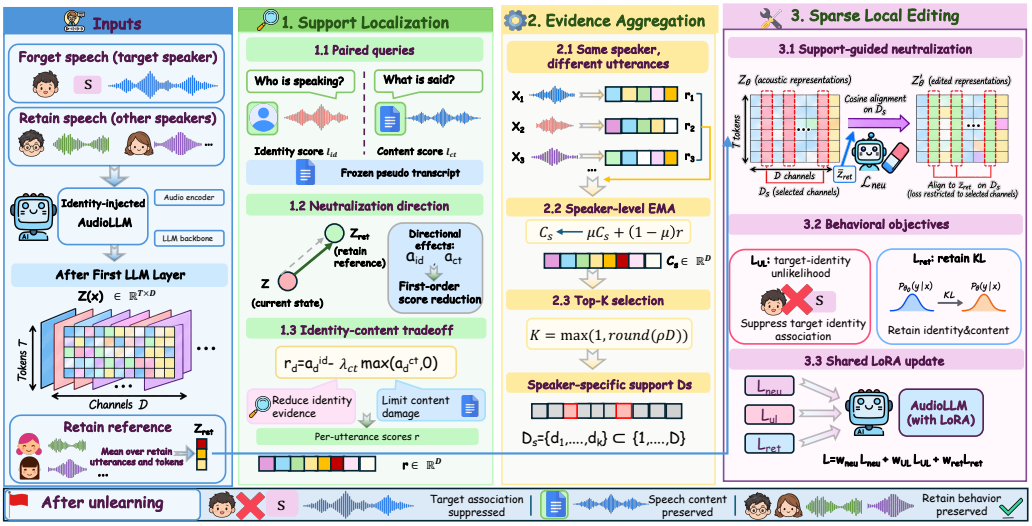


Figure 1: Overview of CAUL. The framework first localizes intervention-aligned acoustic support, aggregates evidence across utterances, and then performs sparse local association editing to remove the target speaker–identity association while preserving linguistic content and non-target behavior.

4 CAUL: Causal Acoustic Unlearning

4.1 Overview

We propose **CAUL**, an acoustic unlearning framework for selectively removing speaker–identity associations. As illustrated in Figure 1, CAUL consists of three stages. *Intervention-Aligned Support Localization* identifies acoustic channels whose neutralization preferentially reduces target-identity evidence while limiting its effect on linguistic content. *Cross-Utterance Evidence Aggregation* consolidates channel-level evidence across utterances into stable speaker-specific support. Finally, *Sparse Local Association Editing* neutralizes the selected support and enforces target-identity forgetting while preserving non-target behavior.

4.2 Intervention-Aligned Support Localization

Motivation. Speaker identity and linguistic content are jointly encoded in speech representations, such that removing a target identity association may also disrupt linguistic information that should be preserved. Selective unlearning therefore requires identifying acoustic evidence that more strongly supports the target identity while being less critical to linguistic content.

Acoustic representation and reference. Given speech x, we denote its audio-token states after the first language-model layer by

$$Z(x) = [z_1, \ldots, z_T] \in \mathbb{R}^{T \times D}, \qquad \bar{z}_{\text{ret}} = \frac{1}{|\mathcal{R}_p|} \sum_{x_r \in \mathcal{R}_p} \frac{1}{T_r} \sum_{t=1}^{T_r} z_t(x_r), \tag{2}$$

where each hidden dimension d is treated as a candidate acoustic channel and $\bar{z}_{\text{ret}}$ is a retain-derived reference computed from a fixed retain subset $\mathcal{R}_p$.

Identity–content directional effects. For each utterance, we use an identity query q for target speaker s and a transcript query q_{ct} with frozen pseudo transcript c. Let I_s and I_c denote their

answer-token spans. Using the identity-acquired model θ_0 before unlearning, we define

$$\ell_{\mathrm{id}} = \frac{1}{|I_s|} \sum_{j \in I_s} \log p_{\theta_0}(y_j \mid y_{<j}, Z, q),$$
$$\ell_{\mathrm{ct}} = \frac{1}{|I_c|} \sum_{j \in I_c} \log p_{\theta_0}(c_j \mid c_{<j}, Z, q_{\mathrm{ct}}). \tag{3}$$

Rather than ranking channels by gradient magnitude alone, we measure how each objective changes along a retain-directed neutralization direction. For channel d, we define the intervention

$$z_{t,d}^{(d)}(\alpha) = z_{t,d} + \alpha\,(\bar{z}_{\mathrm{ret},d} - z_{t,d}), \qquad \alpha \in [0,1], \tag{4}$$

while all other channels remain unchanged. We define its length-normalized first-order directional effect on $u \in \{\mathrm{id}, \mathrm{ct}\}$ as

$$a_d^u := -\frac{1}{T} \left. \frac{\partial \ell_u\big(Z^{(d)}(\alpha)\big)}{\partial \alpha} \right|_{\alpha=0} = \frac{1}{T} \sum_{t=1}^{T} \frac{\partial \ell_u}{\partial z_{t,d}} \left(z_{t,d} - \bar{z}_{\mathrm{ret},d}\right). \tag{5}$$

Positive a_d^u indicates a predicted reduction in the corresponding score under movement toward the retain-derived reference. The normalization by T makes channel effects comparable across utterances with different numbers of audio tokens. We rank channels by their identity–content tradeoff:

$$r_d = a_d^{\mathrm{id}} - \lambda_{\mathrm{ct}} \max\big(a_d^{\mathrm{ct}}, 0\big), \tag{6}$$

where λ_{ct} penalizes predicted content degradation. Larger r_d therefore prioritizes channels with a more favorable tradeoff between target-identity reduction and content preservation.

4.3 Cross-Utterance Evidence Aggregation

Motivation. Channel effects from a single utterance may be affected by linguistic content, recording conditions, or speaking variation. We therefore aggregate evidence across utterances to identify stable speaker-specific support.

Speaker-level aggregation. For the k-th utterance of speaker s, let $r^{(k)} = [r_1^{(k)}, \ldots, r_D^{(k)}]$ denote its channel-score vector. We aggregate these scores using an exponential moving average:

$$C_s^{(1)} = r^{(1)}, \qquad C_s^{(k)} = \mu C_s^{(k-1)} + (1-\mu) r^{(k)}. \tag{7}$$

We then retain the highest-scoring channels as speaker-specific support:

$$K = \max(1, \mathrm{round}(\rho D)), \qquad D_s^{(k)} = \mathrm{TopK}\Big(C_s^{(k)}, K\Big), \tag{8}$$

where μ controls cross-utterance aggregation and ρ controls support sparsity. The resulting $D_s^{(k)}$ favors channels whose intervention-aligned scores remain consistently high across utterances.

4.4 Sparse Local Association Editing

Given the speaker-specific support D_s, CAUL jointly performs localized representation editing, target-identity forgetting, and retain-behavior preservation. Let θ denote the model with the trainable unlearning LoRA and $Z_\theta(x)$ its current audio-token representations. We omit the step superscript of D_s for clarity.

Support-guided acoustic neutralization. We align the selected acoustic support with the retain-derived reference:

$$L_{\mathrm{neu}} = 1 - \frac{1}{T} \sum_{t=1}^{T} \cos\left(Z_{\theta,t,D_s}, \bar{z}_{\mathrm{ret},D_s}\right). \tag{9}$$

This restricts representation editing to the localized speaker-specific support.

Behavioral forgetting. To suppress the target identity at the output, we apply unlikelihood training to its answer tokens:

$$L_{\mathrm{UL}} = -\frac{1}{|I_s|} \sum_{j \in I_s} \log \left[1 - \mathrm{clip}\left(p_\theta(y_j \mid y_{<j}, x, q), 0, 1-\epsilon\right)\right], \tag{10}$$

where $\epsilon > 0$ ensures numerical stability.

Retain-behavior preservation. For each retain utterance x_r, let $P_\theta^{x_r,m}$ and $P_{\theta_0}^{x_r,m}$ denote the answer-token distributions of the updated and pre-unlearning models for mode $m \in \{\mathrm{id}, \mathrm{ct}\}$. We preserve non-target identity and content behavior through

$$L_{\mathrm{ret}} = \frac{1}{2} \sum_{m \in \{\mathrm{id},\mathrm{ct}\}} \mathrm{KL}\left(P_{\theta_0}^{x_r,m} \,\|\, P_\theta^{x_r,m}\right). \tag{11}$$

Overall objective. The final objective is

$$L = w_{\mathrm{neu}} L_{\mathrm{neu}} + w_{\mathrm{UL}} L_{\mathrm{UL}} + w_{\mathrm{ret}} L_{\mathrm{ret}}. \tag{12}$$

All three objectives jointly update the shared unlearning LoRA. Implementation details and hyperparameter settings are provided in Appendix B.2.

5 VoxSU-Bench

We establish **VoxSU-Bench** to evaluate speaker-identity association unlearning in AudioLLMs. Built upon VoxCeleb, VoxSU-Bench pairs identity and content queries on the same speech recordings. We first perform a controlled identity-acquisition stage and partition speakers into forget, retain, and neutral sets. All unlearning methods start from the same identity-acquired model. Details of benchmark construction and identity acquisition are provided in Appendix E.

VoxSU-Bench evaluates unlearning along four complementary dimensions: *Efficacy*, *Generality*, *Content Preservation*, and *Robustness*.

Efficacy. We evaluate whether the acquired voice–identity association is removed for target speakers. We report **Identity Leakage (IL)**, defined as the percentage of target-identity evaluation queries that still reveal the correct speaker identity after unlearning, and **Forget Rate (FR)**, defined as the percentage of pre-unlearning target-identity base hits that are no longer correctly identified after unlearning.

Generality. We evaluate whether forgetting generalizes to open-ended identity-query formulations unseen during unlearning. VoxSU-Bench includes held-out open-ended query families disjoint from those used for optimization. We report Unseen-Elicitation Forget Rate (UFR), defined as the percentage of pre-unlearning identity base hits under these unseen query formulations that are removed after unlearning.

Content Preservation. We evaluate speech-content understanding on held-out recordings spanning forget, retain, and neutral speakers using frozen pseudo-transcripts. We report Content Retention Rate (CRR), defined as the percentage of pre-unlearning successful transcription samples that remain successful after unlearning, where transcription is considered successful if word error rate (WER) ≤ 0.30. Accordingly, the headline CRR reported in the main paper reflects content retention over the global held-out evaluation set, while content retention specifically for forget speakers is reported separately in Appendix E.9.

Robustness. We evaluate whether forgetting persists under label-preserving acoustic transformations of held-out target-speaker recordings, including mild additive noise, speed perturbation, and volume variation. We report **Transformed-Audio Forget Rate (TFR)**, defined as the percentage of transformed inputs correctly identified before unlearning that are no longer correctly identified after unlearning.

Table 1: Main results on VoxSU-Bench. We report the mean and standard deviation over three independent runs. **Bold** and underlined values denote the best and second-best results, respectively.

Method	Efficacy		Generality	Content	Robustness
	IL ↓	FR ↑	UFR ↑	CRR ↑	TFR ↑
Qwen2-Audio$_{7B}$					
Base	$15.67_{\pm 0.00}$	–	–	$100.00_{\pm 0.00}$	–
GA	$\mathbf{0.00_{\pm 0.00}}$	$\mathbf{100.00_{\pm 0.00}}$	$\mathbf{100.00_{\pm 0.00}}$	$20.95_{\pm 36.28}$	$\mathbf{100.00_{\pm 0.00}}$
GA+KL	$11.41_{\pm 0.46}$	$38.89_{\pm 2.41}$	$37.59_{\pm 3.25}$	$\mathbf{93.27_{\pm 0.70}}$	$27.32_{\pm 1.89}$
NPO	$\mathbf{0.00_{\pm 0.00}}$	$\mathbf{100.00_{\pm 0.00}}$	$\mathbf{100.00_{\pm 0.00}}$	$20.03_{\pm 34.69}$	$\mathbf{100.00_{\pm 0.00}}$
SimNPO	$\mathbf{0.00_{\pm 0.00}}$	$\mathbf{100.00_{\pm 0.00}}$	$\mathbf{100.00_{\pm 0.00}}$	$19.27_{\pm 33.37}$	$\mathbf{100.00_{\pm 0.00}}$
SCRUB	$\underline{4.93_{\pm 1.66}}$	$\underline{83.33_{\pm 8.33}}$	$\underline{85.82_{\pm 5.35}}$	$83.49_{\pm 0.92}$	$\underline{83.06_{\pm 3.41}}$
CAUL	$\mathbf{0.00_{\pm 0.00}}$	$\mathbf{100.00_{\pm 0.00}}$	$\mathbf{100.00_{\pm 0.00}}$	$\underline{92.51_{\pm 3.05}}$	$\mathbf{100.00_{\pm 0.00}}$
SALMONN$_{7B}$					
Base	$55.63_{\pm 0.00}$	–	–	$100.00_{\pm 0.00}$	–
GA	$28.99_{\pm 10.02}$	$49.17_{\pm 15.28}$	$48.06_{\pm 18.02}$	$90.68_{\pm 0.62}$	$55.97_{\pm 17.52}$
GA+KL	$55.41_{\pm 0.22}$	$0.42_{\pm 0.72}$	$0.58_{\pm 0.58}$	$\mathbf{98.76_{\pm 0.00}}$	$0.00_{\pm 0.00}$
NPO	$21.93_{\pm 9.43}$	$60.83_{\pm 14.91}$	$60.66_{\pm 18.22}$	$91.20_{\pm 1.59}$	$65.84_{\pm 14.93}$
SimNPO	$26.49_{\pm 7.07}$	$53.75_{\pm 8.75}$	$52.91_{\pm 12.79}$	$91.72_{\pm 2.41}$	$62.55_{\pm 9.10}$
SCRUB	$\underline{19.28_{\pm 17.57}}$	$\underline{69.58_{\pm 31.38}}$	$\underline{70.54_{\pm 25.80}}$	$97.31_{\pm 0.18}$	$\underline{69.55_{\pm 27.78}}$
CAUL	$\mathbf{1.84_{\pm 0.92}}$	$\mathbf{97.08_{\pm 1.44}}$	$\mathbf{96.51_{\pm 2.01}}$	$\underline{97.52_{\pm 0.82}}$	$\mathbf{97.53_{\pm 1.23}}$
SALMONN$_{13B}$					
Base	$73.29_{\pm 0.00}$	–	–	$100.00_{\pm 0.00}$	–
GA	$62.33_{\pm 1.85}$	$20.13_{\pm 4.25}$	$18.29_{\pm 3.26}$	$91.21_{\pm 3.49}$	$18.48_{\pm 5.25}$
GA+KL	$73.22_{\pm 1.00}$	$1.26_{\pm 1.09}$	$1.03_{\pm 0.92}$	$98.18_{\pm 0.00}$	$0.61_{\pm 0.52}$
NPO	$59.97_{\pm 4.20}$	$22.33_{\pm 6.63}$	$22.71_{\pm 7.74}$	$91.01_{\pm 3.38}$	$21.21_{\pm 6.05}$
SimNPO	$60.26_{\pm 3.87}$	$22.01_{\pm 6.63}$	$21.39_{\pm 6.45}$	$93.74_{\pm 1.22}$	$19.09_{\pm 6.56}$
SCRUB	$\underline{30.68_{\pm 22.77}}$	$\underline{59.43_{\pm 31.40}}$	$\underline{65.49_{\pm 25.39}}$	$\underline{98.89_{\pm 0.35}}$	$\underline{60.61_{\pm 29.41}}$
CAUL	$\mathbf{23.18_{\pm 2.55}}$	$\mathbf{67.30_{\pm 3.31}}$	$\mathbf{74.19_{\pm 2.84}}$	$\mathbf{99.29_{\pm 0.17}}$	$\mathbf{66.36_{\pm 6.56}}$

6 EXPERIMENTS

6.1 EXPERIMENT SETUP

Models and Training. We evaluate Qwen2-Audio-7B-Instruct (Chu et al., 2024), SALMONN-7B, and SALMONN-13B (Tang et al., 2024). Each method starts from the corresponding identity-acquired checkpoint and trains an independent unlearning LoRA (Hu et al., 2022) for 50 optimizer steps. We report results over three random seeds and match the trainable module families within each backbone. Hardware, optimizer, adapter, and method-specific settings are detailed in Appendix B.2.

Baselines. We compare CAUL with (1) Gradient Ascent (GA) and (2) GA+KL (Maini et al., 2025), (3) Negative Preference Optimization (NPO) (Zhang et al., 2024), (4) SimNPO (Fan et al.), and (5) SCRUB (Kurmanji et al., 2023). All methods use the same forget/retain partitions within each backbone; their objectives and adaptations are detailed in Appendix B.1.

6.2 MAIN RESULTS

The experimental results in Table 1 provide a comprehensive evaluation of speaker-identity association unlearning across different AudioLLMs. The observations are as follows: (i) In terms of *Efficacy*, CAUL achieves strong identity forgetting across all three backbones. In particular, it reaches nearly complete forgetting on Qwen2-Audio and SALMONN-7B, demonstrating its effectiveness in removing learned voice–identity associations. (ii) Several sequence-level baselines can also achieve strong forgetting, but often at the cost of severe content degradation. On Qwen2-Audio, for example, GA, NPO, and SimNPO achieve complete forgetting while retaining only around 20%

Table 2: Objective ablation of Sparse Local Association Editing on Qwen2-Audio-7B. Results are mean ± standard deviation over three runs.

Variant	IL ↓	FR ↑	CRR ↑
w/o $\mathcal{L}_{\text{neu}}$	$2.58_{\pm 1.44}$	$88.89_{\pm 6.36}$	$94.19_{\pm 1.61}$
w/o $\mathcal{L}_{\text{UL}}$	$7.80_{\pm 1.90}$	$58.33_{\pm 8.33}$	$95.11_{\pm 1.32}$
w/o $\mathcal{L}_{\text{ret}}$	$0.44_{\pm 0.44}$	$98.61_{\pm 2.41}$	$38.69_{\pm 7.46}$
CAUL	$0.00_{\pm 0.00}$	$100.00_{\pm 0.00}$	$92.51_{\pm 3.05}$

Table 3: End-to-end ablation of acoustic support design on Qwen2-Audio-7B. All variants use the same editing objectives, support size, and optimization budget.

Support	IL ↓	FR ↑	CRR ↑
Random Top-K	$2.72_{\pm 1.00}$	$88.89_{\pm 6.36}$	$93.88_{\pm 1.15}$
Identity-only	$0.22_{\pm 0.22}$	$98.61_{\pm 2.41}$	$89.76_{\pm 1.15}$
Single-Utt.	$0.15_{\pm 0.13}$	$97.22_{\pm 2.41}$	$91.74_{\pm 1.21}$
CAUL	$0.00_{\pm 0.00}$	$100.00_{\pm 0.00}$	$92.51_{\pm 3.05}$

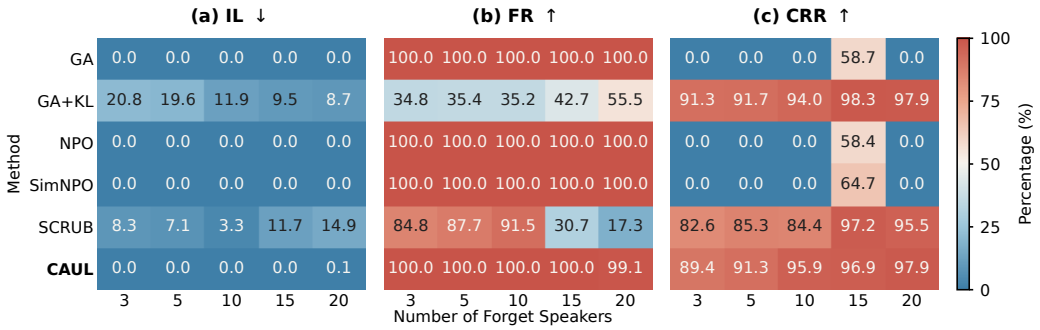


Figure 2: Fixed-budget scaling with the number of forget speakers on Qwen2-Audio-7B. We vary the forget-set size from 3 to 20 while fixing the total optimization budget to 50 steps. We report identity leakage (IL), forgetting rate (FR), and content retention rate (CRR). All results are from a single run.

CRR, whereas CAUL preserves 92.51% CRR under the same forgetting strength. (iii) In *Generality*, the UFR of CAUL remains close to its FR across all backbones, indicating that its forgetting effect generalizes beyond the identity-query formulations used during optimization. (iv) For *Content Preservation*, GA+KL generally retains strong speech content but exhibits much weaker forgetting. In contrast, CAUL maintains both strong identity forgetting and high content retention, achieving a more favorable forgetting–utility balance. Appendix E.9 further confirms that CAUL maintains high content retention specifically on forget-speaker recordings, with 91.87–98.77% CRR across backbones. (v) In *Robustness*, the TFR of CAUL closely follows its FR, showing that the forgetting effect remains stable under mild acoustic transformations such as noise, speed perturbation, and volume variation. (vi) As SALMONN scales from 7B to 13B, most methods exhibit weaker forgetting but improved content retention, suggesting that larger models are more resistant to unlearning while better preserving their original capabilities. CAUL still maintains a strong balance between forgetting and content preservation in this more challenging setting.

6.3 Ablation Study

Tables 2 and 3 examine the editing objectives and support design of CAUL, respectively. For the editing objectives, removing $\mathcal{L}_{\text{neu}}$ and $\mathcal{L}_{\text{UL}}$ reduces FR from 100.00% to 88.89% and 58.33%, confirming their complementary roles in localized and behavioral forgetting; removing $\mathcal{L}_{\text{ret}}$ preserves 98.61% FR but drops CRR from 92.51% to 38.69%, showing its importance for content preservation. For support design, Random Top-K lowers FR to 88.89%, while identity-only localization retains 98.61% FR but reduces CRR to 89.76%. Removing cross-utterance aggregation yields 97.22% FR and 91.74% CRR, whereas full CAUL achieves 100.00% FR with 92.51% CRR. Additional support diagnostics are reported in Appendix C.1.

6.4 Scaling with Forget-Set Size.

Figure 2 evaluates all methods as the number of forget speakers increases from 3 to 20 under the same 50-step optimization budget. CAUL maintains near-complete forgetting throughout the range, with IL remaining at 0.0–0.1% and FR at 99.1–100.0%, while CRR remains high at 89.4–97.9%. GA, NPO,

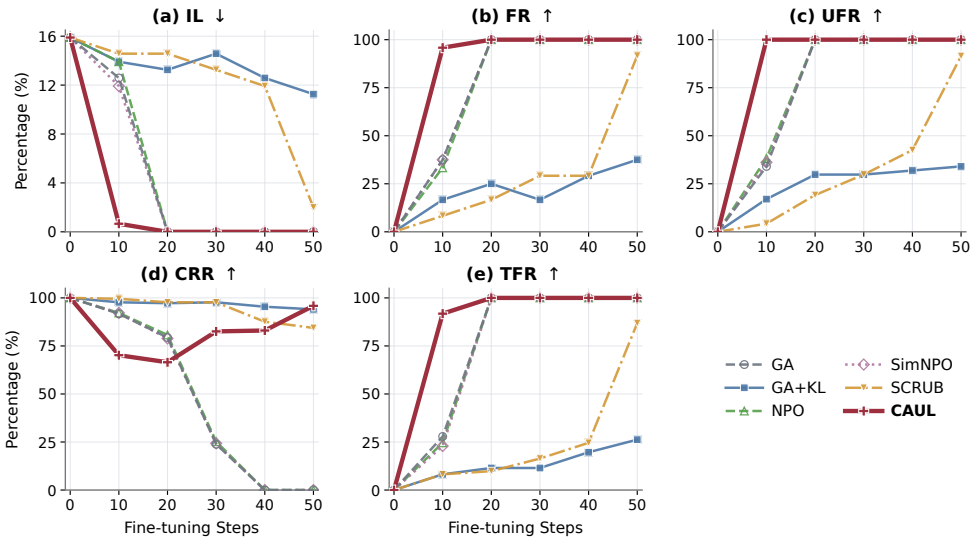


Figure 3: Impact of fine-tuning steps on unlearning performance for Qwen2-Audio-7B. We evaluate six unlearning methods at different optimization steps using identity leakage (IL), forgetting rate (FR), unseen-elicitation forgetting rate (UFR), content retention rate (CRR), and transformed-audio forgetting rate (TFR). All results are from a single run.

and SimNPO achieve similarly strong forgetting, but incur substantially larger content degradation. This is consistent with their sequence-level forgetting objectives, which directly suppress forget responses without explicitly separating speaker-identity evidence from linguistic-content evidence, and can therefore perturb shared acoustic representations more broadly. In contrast, GA+KL preserves content substantially better but leaves considerable identity leakage, while SCRUB exhibits degraded forgetting as the forget set grows. Overall, CAUL maintains a more favorable forgetting–retention balance as the number of simultaneous deletion targets increases.

6.5 Impact of Fine-tuning Steps

We analyze the impact of optimization steps in Figure 3. CAUL adapts rapidly, reaching 95.83% FR and 100% UFR after 10 steps and complete forgetting from step 20 onward. GA, NPO, and SimNPO exhibit similarly fast forgetting, but their CRR degrades sharply as training continues and approaches zero after 40 steps. In contrast, GA+KL preserves content well but achieves much weaker forgetting, while SCRUB adapts more gradually. CAUL avoids this progressive content collapse: after forgetting saturates, its CRR rises from 66.51% at step 20 to 95.87% at step 50, demonstrating a substantially more stable forgetting–retention trade-off under continued optimization. Notably, UFR and TFR closely track FR throughout optimization, indicating that query-level generalization and transformed-audio robustness emerge together with target-association forgetting.

7 Conclusion

In this paper, we introduce **CAUL**, a selective acoustic unlearning framework for removing acquired speaker–identity associations from AudioLLMs while preserving linguistic content. CAUL combines support localization, cross-utterance evidence aggregation, and sparse local association editing to identify, stabilize, and selectively suppress identity-related acoustic support. We further establish **VoxSU-Bench** for evaluating forgetting efficacy, generality, content preservation, and robustness. Experiments across Qwen2-Audio and SALMONN demonstrate strong speaker–identity forgetting with substantially better content retention than direct sequence-level unlearning baselines.

AI USE STATEMENT

We used generative AI tools to assist with searching for and summarizing relevant literature and to edit the manuscript for clarity and readability. During the construction of VoxSU-Bench, these tools also assisted in generating the data split design and identity-query prompts, including query formulations used for training and evaluation. The speech recordings in the benchmark come from VoxCeleb; the AI-assisted work concerned the organization of the data and the text of the prompts, rather than the generation of speech. The authors reviewed the AI-assisted materials and took full responsibility for the final benchmark configuration, manuscript text, and conclusions.

ETHICS STATEMENT

VoxSU-Bench is derived from VoxCeleb1, which contains recordings of public figures collected from publicly available online videos. We do not collect new recordings or interact with the represented speakers. The benchmark uses dataset-specific speaker identifiers and evaluates controlled voice–identity associations rather than disclosing additional personal information. Any released benchmark artifacts will provide derived manifests, prompts, and evaluation metadata without redistributing the original audio; users must obtain the recordings under the terms of the source dataset. Although speaker-identity unlearning may support privacy-oriented model updates, our results measure behavioral suppression under the evaluated queries and should not be interpreted as proof of complete information removal or as a guarantee of legal compliance.

REPRODUCIBILITY STATEMENT

Sections 3–5 define the task, CAUL, and the evaluation metrics. Appendix B reports the backbone-specific optimization settings, LoRA configuration, baseline objectives, and CAUL hyperparameters, while Appendix E documents speaker selection, source-video-disjoint splitting, query construction, identity acquisition, fixed pre-unlearning screening, and evaluation denominators. All methods within each backbone use the same identity-acquired starting checkpoint and frozen base-hit sets. The accompanying implementation provides the experiment configurations, training and evaluation scripts, benchmark manifests, and per-example outputs needed to reproduce the reported results.

REFERENCES

Lucas Bourtoule, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *2021 IEEE symposium on security and privacy (SP)*, pp. 141–159. IEEE, 2021.

Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015 IEEE Symposium on Security and Privacy*, pp. 463–480, 2015. doi: 10.1109/SP.2015.35.

Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, Wanxiang Che, Xiangzhan Yu, and Furu Wei. BEATs: Audio pre-training with acoustic tokenizers. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 5178–5193. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/chen23ag.html.

Jiali Cheng and Hadi Amiri. Speech Unlearning. In *Interspeech 2025*, pp. 3209–3213, 2025. doi: 10.21437/Interspeech.2025-2412.

Jiali Cheng, Ziheng Chen, Chirag Agarwal, and Hadi Amiri. A mechanistic perspective and circuit-guided difficulty metric for unlearning. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 10950–10964, San Diego, California, United States, July 2026. Association for Computational Linguistics. ISBN 979-8-89176-395-1. doi: 10.18653/v1/2026.findings-acl.532. URL https://aclanthology.org/2026.findings-acl.532/.

Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and Jingren Zhou. Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models. *arXiv preprint arXiv:2311.07919*, 2023.

Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, et al. Qwen2-audio technical report. *arXiv preprint arXiv:2407.10759*, 2024.

Chongyu Fan, Jiancheng Liu, Licong Lin, Jinghan Jia, Ruiqi Zhang, Song Mei, and Sijia Liu. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, pp. 1540–1567. doi: 10.52202/085713-0058.

Fuming Fang, Xin Wang, Junichi Yamagishi, Isao Echizen, Massimiliano Todisco, Nicholas Evans, and Jean-Francois Bonastre. Speaker Anonymization Using X-vector and Neural Waveform Models. In *10th ISCA Workshop on Speech Synthesis (SSW 10)*, pp. 155–160, 2019. doi: 10.21437/SSW.2019-28.

Rohit Gandikota, Joanna Materzynska, Jaden Fiotto-Kaufman, and David Bau. Erasing concepts from diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2426–2436, October 2023.

Sreyan Ghosh, Arushi Goel, Jaehyeon Kim, Sonal Kumar, Zhifeng Kong, Sang-gil Lee, Chao-Han Yang, Ramani Duraiswami, Dinesh Manocha, Rafael Valle, et al. Audio flamingo 3: Advancing audio intelligence with fully open large audio language models. *Advances in Neural Information Processing Systems*, 38:41819–41886, 2026.

Antonio Ginart, Melody Guan, Gregory Valiant, and James Zou. Making ai forget you: Data deletion in machine learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/cb79f8fa58b91d3af6c9c991f63962d3-Paper.pdf.

Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Eternal sunshine of the spotless net: Selective forgetting in deep networks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9301–9309, 2020. doi: 10.1109/CVPR42600.2020.00932.

Yuan Gong, Hongyin Luo, Alexander Liu, Leonid Karlinsky, and James R Glass. Listen, think, and understand. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun (eds.), *International Conference on Learning Representations*, volume 2024, pp. 18516–18545, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/510d0935b543a29d686f93fa52d1c288-Paper-Conference.pdf.

Chuan Guo, Tom Goldstein, Awni Hannun, and Laurens Van Der Maaten. Certified data removal from machine learning models. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3832–3842. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/guo20c.html.

Luxi He, Xiangyu Qi, Michel Liao, Inyoung Cheong, Prateek Mittal, Danqi Chen, and Peter Henderson. The model hears you: Audio language model deployments should consider the principle of least privilege. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 8, pp. 1231–1244, 2025.

Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.

Jiahao Huo, Yibo Yan, Xu Zheng, Yuanhuiyi Lyu, Xin Zou, Zhihua Wei, and Xuming Hu. MMUnlearner: Reformulating multimodal machine unlearning in the era of multimodal large language models. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 7190–7206, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.375. URL https://aclanthology.org/2025.findings-acl.375/.

Taesoo Kim, Jinju Kim, Dong Chan Kim, Jong Hwan Ko, and Gyeong-Moon Park. Do not mimic my voice : Speaker identity unlearning for zero-shot text-to-speech. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=1EgFJDjodU.

Alkis Koudounas, Claudio Savelli, Flavio Giobergia, and Elena Baralis. "Alexa, can you forget me?" Machine Unlearning Benchmark in Spoken Language Understanding. In *Interspeech 2025*, pp. 1768–1772, 2025. doi: 10.21437/Interspeech.2025-2607.

Meghdad Kurmanji, Peter Triantafillou, Jamie Hayes, and Eleni Triantafillou. Towards unbounded machine unlearning. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 1957–1987. Curran Associates, Inc., 2023. doi: 10.52202/075280-0095. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/062d711fb777322e2152435459e6e9d9-Paper-Conference.pdf.

Jiaqi Li, Qianshan Wei, Chuanyi Zhang, Guilin Qi, Miaozeng Du, Yongrui Chen, Sheng Bi, and Fan Liu. Single image unlearning: Efficient machine unlearning in multimodal large language models. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 35414–35453. Curran Associates, Inc., 2024a. doi: 10.52202/079017-1116. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/3e53d82a1113e3d240059a9195668edc-Paper-Conference.pdf.

Kunhao Li, Wenhao Li, Di Wu, Lei Yang, Jun Bai, Ju Jia, and Jason Xue. Cross-modal unlearning via influential neuron path editing in multimodal large language models. In *Proceedings of the Fortieth AAAI Conference on Artificial Intelligence and Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence and Sixteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'26/IAAI'26/EAAI'26. AAAI Press, 2026. ISBN 978-1-57735-906-7. doi: 10.1609/aaai.v40i42.40870. URL https://doi.org/10.1609/aaai.v40i42.40870.

Nathaniel Li, Alexander Pan, Anjali Gopal, Summer Yue, Daniel Berrios, Alice Gatti, Justin D. Li, Ann-Kathrin Dombrowski, Shashwat Goel, Long Phan, Gabriel Mukobi, Nathan Helm-Burger, Rassin Lababidi, Lennart Justen, Andrew B. Liu, Michael Chen, Isabelle Barrass, Oliver Zhang, Xiaoyuan Zhu, Rishub Tamirisa, Bhrugu Bharathi, Adam Khoja, Zhenqi Zhao, Ariel Herbert-Voss, Cort B. Breuer, Samuel Marks, Oam Patel, Andy Zou, Mantas Mazeika, Zifan Wang, Palash Oswal, Weiran Lin, Adam A. Hunt, Justin Tienken-Harder, Kevin Y. Shih, Kemper Talley, John Guan, Russell Kaplan, Ian Steneker, David Campbell, Brad Jokubaitis, Alex Levinson, Jean Wang, William Qian, Kallol Krishna Karmakar, Steven Basart, Stephen Fitz, Mindy Levine, Ponnurangam Kumaraguru, Uday Tupakula, Vijay Varadharajan, Ruoyu Wang, Yan Shoshitaishvili, Jimmy Ba, Kevin M. Esvelt, Alexandr Wang, and Dan Hendrycks. The wmdp benchmark: Measuring and reducing malicious use with unlearning, 2024b. URL https://arxiv.org/abs/2403.03218.

Zhe Liu. Unlearning LLM-Based Speech Recognition Models. In *Interspeech 2025*, pp. 3214–3218, 2025. doi: 10.21437/Interspeech.2025-287.

Zheyuan Liu, Guangyao Dou, Xiangchi Yuan, Chunhui Zhang, Zhaoxuan Tan, and Meng Jiang. Modality-aware neuron pruning for unlearning in multimodal large language models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5913–5933, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.295. URL https://aclanthology.org/2025.acl-long.295/.

Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary Chase Lipton, and J Zico Kolter. TOFU: A task of fictitious unlearning for LLMs. In *Red Teaming GenAI: What Can We Learn from Adversaries?*, 2025. URL https://openreview.net/forum?id=P8seBluN3c.

Israel Mason-Williams, Jing Han, Helen Yannakoudakis, and Cecilia Mascolo. Machine unlearning in audio: Bridging the modality gap via the prune and regrow paradigm. 2025.

Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 17359–17372. Curran Associates, Inc., 2022. doi: 10.52202/068431-1262. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/6f1d43d5a82a37e89b0665b33bf3a182-Paper-Conference.pdf.

Arsha Nagrani, Joon Son Chung, and Andrew Zisserman. VoxCeleb: A Large-Scale Speaker Identification Dataset. In *Interspeech 2017*, pp. 2616–2620, 2017. doi: 10.21437/Interspeech.2017-950.

Shreyansh Pathak, Sonu Shreshtha, Richa Singh, and Mayank Vatsa. Quantum-inspired audio unlearning: Towards privacy-preserving voice biometrics. In *2025 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–9, 2025. doi: 10.1109/IJCB65343.2025.11411246.

Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine Mcleavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 28492–28518. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/radford23a.html.

Paul K. Rubenstein, Chulayuth Asawaroengchai, Duc Dung Nguyen, Ankur Bapna, Zalán Borsos, Félix de Chaumont Quitry, Peter Chen, Dalia El Badawy, Wei Han, Eugene Kharitonov, Hannah Muckenhirn, Dirk Padfield, James Qin, Danny Rozenberg, Tara Sainath, Johan Schalkwyk, Matt Sharifi, Michelle Tadmor Ramanovich, Marco Tagliasacchi, Alexandru Tudor, Mihajlo Velimirović, Damien Vincent, Jiahui Yu, Yongqiang Wang, Vicky Zayats, Neil Zeghidour, Yu Zhang, Zhishuai Zhang, Lukas Zilka, and Christian Frank. Audiopalm: A large language model that can speak and listen, 2023. URL https://arxiv.org/abs/2306.12925.

Ayush Sekhari, Jayadev Acharya, Gautam Kamath, and Ananda Theertha Suresh. Remember what you want to forget: Algorithms for machine unlearning. In *Advances in Neural Information Processing Systems*, 2021.

Weijia Shi, Jaechan Lee, Yangsibo Huang, Sadhika Malladi, Jieyu Zhao, Ari Holtzman, Daogao Liu, Luke Zettlemoyer, Noah Smith, and Chiyuan Zhang. Muse: Machine unlearning six-way evaluation for language models. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 27797–27818, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/4556f5398bd2c61bd7500e306b4e560a-Paper-Conference.pdf.

Guangzhi Sun, Potsawee Manakul, Xiao Zhan, and Mark Gales. Unlearning vs. obfuscation: Are we truly removing knowledge? In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 11457–11467, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.577. URL https://aclanthology.org/2025.emnlp-main.577/.

Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun MA, and Chao Zhang. SALMONN: Towards generic hearing abilities for large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=14rn7HpKVk.

Natalia Tomashenko, Xin Wang, Emmanuel Vincent, Jose Patino, Brij Mohan Lal Srivastava, Paul-Gauthier Noé, Andreas Nautsch, Nicholas Evans, Junichi Yamagishi, Benjamin O'Brien, Anaïs Chanclu, Jean-François Bonastre, Massimiliano Todisco, and Mohamed Maouche. The voiceprivacy 2020 challenge: Results and findings. *Computer Speech Language*, 74:101362, 2022. ISSN 0885-2308. doi: https://doi.org/10.1016/j.csl.2022.101362. URL https://www.sciencedirect.com/science/article/pii/S0885230822000080.

Lixu Wang, Kaixiang Yao, Xinfeng Li, Dong Yang, Haoyang Li, Xiaofeng Wang, and Wei Dong. The man behind the sound: Demystifying audio private attribute profiling via multimodal large language model agents, 2025. URL https://arxiv.org/abs/2507.10016.

Heng Xu, Tianqing Zhu, Lefeng Zhang, Wanlei Zhou, and Philip S. Yu. Machine unlearning: A survey. *ACM Comput. Surv.*, 56(1), August 2023. ISSN 0360-0300. doi: 10.1145/3603620. URL https://doi.org/10.1145/3603620.

Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. Qwen2.5-omni technical report, 2025. URL https://arxiv.org/abs/2503.20215.

Xiao Zhan, Guangzhi Sun, Jose Such, and Phil Woodland. Protecting bystander privacy via selective hearing in audio llms. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15180–15192, 2026.

Dawen Zhang, Pamela Finckenberg-Broman, Thong Hoang, Shidong Pan, Zhenchang Xing, Mark Staples, and Xiwei Xu. Right to be forgotten in the era of large language models: Implications, challenges, and solutions. *AI and Ethics*, 5(3):2445–2454, 2025.

Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic collapse to effective unlearning. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=MXLBXjQkmb.

A Appendix

A.1 Notations

For clarity, we summarize the main notations used throughout the paper in Table 4.

Table 4: Summary of the main notations used in this paper.

Notation	Definition
M_0	Pretrained AudioLLM before controlled identity acquisition.
M_{id}	Identity-acquired AudioLLM used as the starting point for unlearning.
M_θ	AudioLLM with the trainable unlearning parameters θ.
$\mathcal{F}, \mathcal{R}, \mathcal{N}$	Forget, retain, and neutral speaker sets, respectively.
x, s, c	Speech waveform, its speaker identity, and linguistic content.
q, q_{ct}	Identity query and linguistic-content query, respectively.
$Z(x) \in \mathbb{R}^{T \times D}$	Audio-token representations with T positions and D channels.
$\bar{z}_{\mathrm{ret}}$	Retain-derived reference representation used for intervention and editing.
I_s, I_c	Answer-token spans for target identity and linguistic content.
$\ell_{\mathrm{id}}, \ell_{\mathrm{ct}}$	Identity and content log-scores used for support localization.
a_d^u	Intervention-aligned directional effect of channel d for $u \in \{\mathrm{id}, \mathrm{ct}\}$.
r_d	Identity–content localization score of acoustic channel d.
$C_s^{(k)}$	Aggregated channel-score vector for speaker s after the k-th utterance.
$D_s^{(k)}$	Top-K acoustic channels selected as speaker-specific support.
ρ, μ	Channel-selection ratio and cross-utterance EMA momentum.
λ_{ct}	Weight penalizing predicted content degradation during localization.
$\mathcal{L}_{\mathrm{neu}}$	Support-guided acoustic neutralization loss.
$\mathcal{L}_{\mathrm{UL}}$	Target-identity unlikelihood loss.
$\mathcal{L}_{\mathrm{ret}}$	Retain-behavior preservation loss.
$w_{\mathrm{neu}}, w_{\mathrm{UL}}, w_{\mathrm{ret}}$	Weights of the three objectives in the final CAUL loss.

B Experimental Details

B.1 Baseline Details

This section provides a comprehensive overview of the baseline approaches employed in our study.

Gradient Ascent (GA) (Maini et al., 2025). GA is a straightforward unlearning approach that reverses the standard maximum-likelihood training objective on forget examples. Instead of increasing the likelihood of the target responses, it performs gradient ascent on the corresponding training loss, thereby reducing the model's confidence in the designated knowledge. Despite its simplicity, GA provides a fundamental baseline for evaluating direct knowledge suppression in generative-model unlearning.

GA+KL (Maini et al., 2025). GA+KL extends gradient-ascent unlearning by introducing a retention constraint on non-forget data. While the GA objective reduces the likelihood of designated forget targets, a KL-divergence regularizer encourages the updated model to remain close to the original model on retain examples. The two objectives jointly balance knowledge removal against preservation of the model's original behavior, reducing the utility degradation commonly caused by unconstrained gradient ascent.

Negative Preference Optimization (NPO) (Zhang et al., 2024). NPO formulates machine unlearning as a preference-optimization problem with respect to a frozen reference model. Rather than directly maximizing the loss of forget examples as in GA, NPO applies a logistic objective to the likelihood ratio between the current and reference models. This objective progressively suppresses forget examples while controlling excessive parameter updates as their likelihood decreases, providing a more stable alternative to unconstrained gradient-ascent unlearning.

SimNPO (Fan et al.). SimNPO is a reference-free variant of NPO that simplifies its preference-based forgetting objective. It removes the dependence on reference-model likelihoods and constructs the logistic objective directly from the current model's likelihood on forget examples. A length-normalized formulation further reduces sensitivity to response length. Consequently, SimNPO retains the likelihood-suppression behavior of NPO while avoiding the computational and memory overhead associated with a reference model.

SCRUB (Kurmanji et al., 2023). SCRUB adopts a teacher–student framework that combines knowledge removal with preservation of model utility. On forget data, it encourages the student model to diverge from the original teacher, thereby suppressing information associated with the designated examples. On retain data, the student is trained to match the teacher's predictions while maintaining supervised task performance. Through these opposing objectives, SCRUB seeks to selectively forget designated information while preserving behavior on the remaining data.

B.2 Implementation Details

Environment. Experiments are conducted on NVIDIA A40 GPUs (48 GB each) with dual Intel Xeon Platinum 8558P CPUs (48 cores per socket). The implementation uses PyTorch 2.7.1 with a CUDA 11.8 build and Transformers 5.14.1. Unless otherwise stated, all methods are trained for 50 optimizer steps with three random seeds. We use AdamW, clip the gradient norm to 1.0, and sample forget examples in speaker round-robin order so that speakers with more utterances do not receive more updates. Each method starts from the same identity-acquired checkpoint for a given backbone and trains a separately initialized unlearning adapter.

Backbone-specific adaptation. For benchmark construction, Qwen2-Audio identity acquisition is implemented with a LoRA adapter, which is merged into the backbone before the unlearning stage. The acquisition adapter and its parameter update are not provided as inputs to any unlearning method. Starting from this common merged checkpoint, we apply a separately initialized rank-4 unlearning LoRA with alpha 8 and zero dropout to the audio projector and the language model's query and value projections. We use eager attention and a learning rate of 2×10^{-4}. For both SALMONN backbones, the unlearning stage uses separately initialized rank-4 LoRA modules with alpha 8 on the speech-to-language projector and the language model's query and value projections. The speech Q-Former, Whisper (Radford et al., 2023), and BEATs (Chen et al., 2023) encoders remain frozen. Both trainable unlearning LoRA branches use a learning rate of 5×10^{-5}.

Baseline configurations. GA minimizes the mean target-answer log-likelihood with unit weight and uses no retain regularizer. GA+KL uses the same forget objective with an answer-token retain KL weight of 5. NPO uses $\beta = 0.1$ together with answer-token retain KL regularization of weight 5. SimNPO uses $\beta = 0.1$, the same retain KL regularization, and mean answer-token log-likelihood, thereby normalizing each objective by the answer length. SCRUB alternates between maximizing teacher–student KL divergence on forget-answer tokens and minimizing KL divergence plus supervised cross-entropy on retain-answer tokens; its forget, retain-KL, and retain-CE weights are 1, 0.5, and 1, respectively, with the retain objective further multiplied by the common retain weight of 5. All methods use the same trainable module families and the same 50-step training schedule within each backbone.

CAUL configuration. At each training step, identity and content representation gradients update only the speaker evidence bank used to construct the sparse support mask; the localization probes do not update model parameters. CAUL localizes all audio-token positions at the first language-model layer in Qwen2-Audio and at the speech-to-language projector output in SALMONN. The neutral reference is fixed before unlearning by averaging acoustic states from five retained utterances. We set the channel ratio to $\rho = 0.10$, the speaker evidence momentum to $\mu = 0.5$, the content penalty to $\lambda_{\mathrm{ct}} = 1.0$, and $(w_{\mathrm{neu}}, w_{\mathrm{UL}}, w_{\mathrm{ret}}) = (1, 1, 5)$. For each paired retain utterance, identity and transcript answer-token KL divergences are averaged. Neutralization, target-identity unlikelihood, and retain preservation are computed in the same iteration and jointly produce one optimizer update; we do not alternate these objectives.

Table 5: Mechanistic diagnostics of acoustic support on Qwen2-Audio-7B. ID/CT Eff. denote the average changes in identity/content evidence induced by neutralizing the selected support. Sel. measures identity reduction relative to content disruption, and Jaccard measures overlap between consecutive Top-K support masks.

Support	ID Eff. ↑	CT Eff. ↓	Sel. ↑	Jaccard ↑
Random	-0.0018	0.0137	-0.0194	–
Identity-only	**0.1286**	0.0295	0.0948	–
Single-Utt.	–	–	–	$0.127_{\pm 0.131}$
CAUL	0.1163	**0.0124**	**0.0979**	$0.367_{\pm 0.135}$

C Additional Experimental Results

C.1 Support Selectivity and Stability Diagnostics

We complement the end-to-end support ablation in Table 3 with mechanism-level diagnostics of the localized acoustic support. Specifically, we measure the identity and content effects induced by neutralizing the selected channels, together with the overlap between consecutive Top-K support masks.

As shown in Table 5, random channels induce almost no identity reduction (-0.0018), indicating that the forgetting effect does not arise from arbitrary sparse intervention. Identity-only localization captures strong identity-related evidence (0.1286 ID Effect), but also incurs a substantially larger Content Effect (0.0295). In contrast, CAUL retains a comparable Identity Effect (0.1163) while reducing Content Effect to 0.0124, resulting in higher identity–content selectivity. Cross-utterance aggregation further increases the overlap between consecutive support masks from 0.127 to 0.367, confirming that speaker-level aggregation produces substantially more stable support across utterances.

D Complexity Analysis

In this section, we analyze the optimization complexity of CAUL. Let S denote the number of optimization steps, T the number of audio-token representations, D their hidden dimension, and $K = \rho D$ the number of selected acoustic channels. We denote the cost of one standard LoRA forward–backward update by C_{train} and one representation-gradient probe by C_{probe}. A conventional LoRA-based unlearning method therefore has complexity $O(SC_{\text{train}})$.

Compared with standard LoRA-based unlearning, CAUL introduces three additional components:

• **Intervention-Aligned Support Localization.** For each forget example, CAUL computes one identity and one content representation gradient, costing $2C_{\text{probe}}$. Computing the directional channel scores over T audio-token positions and D channels further requires $O(TD)$ operations.

• **Cross-Utterance Evidence Aggregation.** Updating the speaker-level score through exponential moving averaging costs $O(D)$, while selecting the top-K channels costs $O(D \log K)$.

• **Sparse Local Association Editing.** Neutralization operates on the selected $T \times K$ acoustic units, incurring $O(TK)$ additional computation. Neutralization, target-identity unlikelihood, and retain KL are jointly optimized within a single LoRA update.

Therefore, the total optimization complexity of CAUL is

$$O(S\,(C_{\text{train}} + 2C_{\text{probe}} + TD + D \log K))\,, \tag{13}$$

where the $O(D)$ aggregation cost and $O(TK)$ sparse-editing cost are absorbed by $O(TD)$ since $K \leq D$.

The auxiliary memory cost is $O(|\mathcal{F}|D)$ for maintaining speaker-level aggregated scores, with the retain-derived reference and selected support requiring only lower-order storage. Trainable parameters and optimizer states continue to scale with the LoRA parameters rather than the full backbone.

Table 6: Speaker and audio composition of VoxSU-Bench.

Split	Speakers	Train clips	Test clips
Forget	10	80	151
Retain	20	160	152
Neutral	10	80	82
Total	40	320	385

E VoxSU-Bench

Source Dataset. VoxSU-Bench is constructed from VoxCeleb1 (Nagrani et al., 2017), a large-scale audio-visual speaker-recognition dataset collected in the wild. VoxCeleb1 contains more than 150,000 speech utterances from 1,251 speakers, automatically extracted from interview and other public videos uploaded to YouTube. Unlike speech recorded under controlled studio conditions, these clips exhibit substantial natural variation in recording conditions, background noise, speaking style, and acoustic environment. Importantly, each speaker is represented by multiple utterances originating from multiple source videos, which allows us to separate training and evaluation recordings at the source-video level while preserving speaker identity. This property is particularly suitable for speaker-identity association unlearning: multiple natural utterances provide diverse evidence for learning a voice–identity association, while source-video-disjoint evaluation reduces the possibility that successful identification or forgetting is driven by video-specific recording artifacts.

E.1 Construction Objective

VoxSU-Bench is designed to compare speaker-identity association unlearning under a controlled and reproducible starting condition. Natural identity recognition in general-purpose AudioLLMs is concentrated among a small set of widely represented public figures, with measured accuracy shaped by name priors, recording context, and uneven pretraining exposure. We therefore construct the target associations through controlled identity acquisition. The benchmark fixes the speaker label space, acquisition examples, query families, and starting checkpoint. This isolates the effect of the unlearning method from variation in pretrained identity coverage.

The benchmark follows four stages. First, we construct source-video-disjoint audio splits from VoxCeleb1. Second, we freeze one pseudo-transcript for every selected clip and generate paired identity and content queries. Third, each backbone acquires the target voice–identity associations from the same training partition. Finally, we screen the acquired model before unlearning to establish fixed identity and content base-hit sets. All methods for a given backbone start from this screened identity-acquired model and share the same evaluation examples.

E.2 Speaker Sampling and Audio Splitting

We select speakers from the VoxCeleb1 development set. A speaker is eligible when at least eight source videos and 40 utterances are available. Deterministic hash-based ordering fixes speaker selection and utterance sampling for a given benchmark seed. The resulting benchmark contains 40 speakers divided into 10 forget, 20 retain, and 10 neutral speakers. Forget speakers define the requested associations to remove, retain speakers supply non-target preservation examples, and neutral speakers provide additional non-target acoustic variation.

We split audio at the source-video level before selecting utterances. Clips from a source video are assigned entirely to either training or evaluation, preventing adjacent segments and shared recording conditions from crossing the split. The final benchmark contains 705 unique clips: 320 training clips from 196 source videos and 385 held-out clips from 154 source videos. The train and test source-video sets have zero overlap. Table 6 summarizes the speaker and clip allocation.

E.3 Frozen Speech-Content Targets

We generate one pseudo-transcript for each of the 705 clips using the clean Qwen2-Audio backbone with deterministic decoding. The transcript manifest is produced before identity acquisition and

Table 7: Identity-query families in VoxSU-Bench. Training and evaluation use disjoint templates.

ID	Query template	Partition	Role
q1	Who is speaking in this audio?	Train	Controlled identity acquisition and unlearning
q2	Identify the speaker in this recording.	Train	Controlled identity acquisition and unlearning
q3	Which enrolled identity matches this voice?	Train	Controlled identity acquisition and unlearning
q4	Who produced this speech?	Test	Efficacy anchor
q5	Is the speaker {identity}?	Test	Auxiliary candidate-hint probe
q6	Does this voice match {identity}?	Test	Auxiliary candidate-hint probe
q7	Describe the speaker identity evidence.	Test	Unseen open elicitation
q8	What is the speaker's identity?	Test	Unseen open elicitation

remains fixed across backbones, unlearning methods, and random seeds. This gives every method the same content target for localization, retain supervision, and content evaluation.

E.4 Identity Query Construction

The 320 training clips are expanded with q1–q3, producing 960 training-query instances across all speaker roles. Identity acquisition itself uses the 240 instances associated with the 80 forget-speaker clips. The 385 held-out clips are expanded with q4–q8, producing 1,925 evaluation queries. Table 7 lists the fixed templates and their roles.

Queries q4, q7, and q8 require the model to generate the speaker identity without a candidate in the prompt. Across the 151 held-out clips of the forget speakers, they form 453 open-identity queries. Queries q5 and q6 introduce a candidate identity and use balanced positive and negative hints. For each family, 192 of the 385 prompts contain the correct identity and 193 contain a deterministically selected incorrect identity. We retain q5 and q6 as auxiliary candidate-hint probes, but do not include them in the primary identity metrics. The reported identity metrics therefore use the open-generation families q4, q7, and q8.

E.5 Content Query Construction

We construct content queries for the same 705 recordings and use frozen pseudo-transcripts as answer targets. Training uses three formulations: direct transcription, a request for what the speaker says, and a request to write down the spoken content. Evaluation uses five held-out formulations. For the primary CRR evaluation, we use `content_q4`, "What words are spoken in this recording?", once for each of the 385 held-out recordings spanning forget, retain, and neutral speakers. Each recording is paired with its frozen pseudo-transcript generated before unlearning, enabling consistent evaluation of speech-content retention across methods.

E.6 Controlled Identity Acquisition

For the reported experiments, identity acquisition uses the 80 training clips of the 10 target speakers. Combining each clip with q1–q3 produces 240 supervised voice–identity examples. The acquired model learns a common answer format, *The speaker is [identity]*, which enables exact and consistent screening across architectures. Retain and neutral clips are used to construct preservation examples and reference representations during unlearning.

Identity acquisition is performed independently for each backbone because their acoustic interfaces and adaptation modules differ. Within a backbone, the resulting identity model is frozen as the common starting point for GA, GA+KL, NPO, SimNPO, SCRUB, and CAUL. Each unlearning run attaches a newly initialized adapter to this same starting model.

E.7 Pre-Unlearning Screening

We screen each identity-acquired model before unlearning and freeze the resulting per-example labels. For the open-generation families q4, q7, and q8, a response is a hit when its normalized generation contains the exact target speaker ID. We use this conservative lexical criterion to measure explicit target-identity disclosure. Screening is model-specific because different backbones acquire different

Table 8: Split-wise content retention of CAUL. Values in parentheses denote the number of pre-unlearning content base hits in each subset. Results are mean ± standard deviation over three runs.

Backbone	Forget CRR ↑	Retain CRR ↑	Neutral CRR ↑	Global CRR ↑
Qwen2-Audio-7B	91.87±3.52 (82)	95.18±3.61 (83)	89.31±4.75 (53)	92.51±3.05 (218)
SALMONN-7B	96.88±1.35 (128)	97.71±0.00 (131)	98.41±1.59 (63)	97.52±0.82 (322)
SALMONN-13B	98.77±0.43 (135)	99.49±0.44 (131)	100.00±0.00 (64)	99.29±0.17 (330)

subsets of the association, and it is method-shared because every unlearning method for that backbone uses the same frozen subset.

We separate the q4 efficacy anchor from the q7/q8 unseen-elicitation families. Qwen2-Audio produces 24 q4 base hits and 47 q7/q8 base hits; SALMONN-7B produces 80 and 172; SALMONN-13B produces 106 and 226. Their pooled open-identity base-hit counts are therefore 71, 252, and 332 out of 453 queries, respectively. These counts yield the base Identity Leakage values reported in Table 1.

Content screening follows the same fixed-denominator principle. We generate one response for `content_q4` on every held-out clip and compute word error rate against the frozen transcript. A clip enters the content base-hit set when the identity-acquired model attains WER ≤ 0.30. This yields 218, 322, and 330 content base hits for Qwen2-Audio-7B, SALMONN-7B, and SALMONN-13B, respectively. Their forget/retain/neutral compositions are $82/83/53$, $128/131/63$, and $135/131/64$, respectively. CRR measures the fraction of these fixed base-hit sets that remains below the same WER threshold after unlearning.

E.8 Evaluation Sets and Metrics

Let $\mathcal{Q}_{\text{open}}$ contain all 453 q4/q7/q8 queries, let H_b^{q4} and $H_b^{q7,q8}$ denote backbone b's screened base-hit sets, and let H_b^{ct} denote its screened content set. VoxSU-Bench computes

$$\text{IL} = \frac{1}{|\mathcal{Q}_{\text{open}}|} \sum_{e \in \mathcal{Q}_{\text{open}}} \mathbf{1}[M_{\hat{\theta}} \text{ identifies the target in } e], \tag{14}$$

$$\text{FR} = \frac{1}{|H_b^{q4}|} \sum_{e \in H_b^{q4}} \mathbf{1}[M_{\hat{\theta}} \text{ no longer identifies the target in } e], \tag{15}$$

$$\text{UFR} = \frac{1}{|H_b^{q7,q8}|} \sum_{e \in H_b^{q7,q8}} \mathbf{1}[M_{\hat{\theta}} \text{ no longer identifies the target in } e], \tag{16}$$

$$\text{CRR} = \frac{1}{|H_b^{\text{ct}}|} \sum_{e \in H_b^{\text{ct}}} \mathbf{1}[\text{WER}(M_{\hat{\theta}}, e) \leq 0.30]. \tag{17}$$

IL measures residual leakage across all open formulations, FR measures forgetting on the q4 efficacy anchor, UFR measures transfer to the unseen q7/q8 elicitation families, and CRR measures generated-content preservation.

For transformed-audio evaluation, we start from q4 base-hit recordings and apply three deterministic, label-preserving transformations: additive noise at 20 dB signal-to-noise ratio, $0.9\times$ speed perturbation, and $+3$ dB gain. Each backbone is screened again on these transformed inputs to form a fixed transformed base-hit set. TFR is the fraction of that set on which the target identity is removed after unlearning.

E.9 Split-wise Content Retention

To examine content preservation specifically on the speakers targeted for unlearning, we partition the screened content set $\mathcal{H}_b^{ct}$ by speaker role and compute Eq. (17) independently on forget, retain, and neutral recordings. Table 8 reports the resulting split-wise CRR. CAUL maintains high content retention on forget-speaker recordings, showing that the high global CRR is accompanied by preservation of linguistic content on the speakers whose identity associations are removed.

E.10 Construction Checks

The released manifests store the speaker, source video, utterance, role, query family, transcript, expected answer, hint condition, and screening fields for every example. Construction verifies one transcript per audio path, zero source-video overlap between train and test, disjoint train/test query families, balanced verification hints, and fixed model-specific base-hit labels. Evaluation saves per-example generations in addition to aggregate scores, enabling per-speaker and per-query-family analysis under the same frozen protocol.

F Discussion on Limitations

Our study has several limitations. First, we investigate speaker–identity association unlearning under a controlled identity-acquisition setting, which enables reproducible deletion targets but differs from identity knowledge naturally acquired during large-scale pretraining. Second, our current evaluation mainly verifies behavioral suppression of target identities under open-ended queries and unseen query formulations; it therefore should not be interpreted as evidence that identity information has been completely removed from the model's internal representations. Finally, our experiments focus on a VoxCeleb-derived benchmark and the Qwen2-Audio and SALMONN model families. Future work can extend this setting to a broader range of AudioLLMs, naturally acquired identity knowledge, and more complex identity-elicitation scenarios.