

# A Data-Centric Approach to Generalizable Speech Deepfake Detection

Wen Huang<sup>‡, b</sup>, Yuchen Mao<sup>‡, b</sup>, Yanmin Qian<sup>‡, b\*</sup>

<sup>‡</sup>Auditory Cognition and Computational Acoustics Lab  
MoE Key Lab of Artificial Intelligence, AI Institute  
School of Computer Science, Shanghai Jiao Tong University, China  
<sup>b</sup>VUI Labs

## Abstract

Achieving robust generalization in speech deepfake detection (SDD) remains a primary challenge, as models often fail to detect unseen forgery methods. While research has focused on model-centric and algorithm-centric solutions, the impact of data composition is often underexplored. This paper proposes a data-centric approach, analyzing the SDD data landscape from two practical perspectives: constructing a single dataset and aggregating multiple datasets. To address the first perspective, we conduct a large-scale empirical study to characterize the data scaling laws for SDD, quantifying the impact of source and generator diversity. To address the second, we propose the Diversity-Optimized Sampling Strategy (DOSS), a principled framework for mixing heterogeneous data with two implementations: DOSS-Select (pruning) and DOSS-Weight (re-weighting). Our experiments show that DOSS-Select outperforms the naive aggregation baseline while using only 3% of the total available data. Furthermore, our final model, trained on a 12k-hour curated data pool using the optimal DOSS-Weight strategy, achieves state-of-the-art performance, outperforming large-scale baselines with greater data and model efficiency on both public benchmarks and a new challenge set of various commercial APIs.

## 1 Introduction

Speech Deepfake Detection (SDD) has emerged as a critical research area in recent years. As speech synthesis technologies become increasingly sophisticated, the resulting deepfakes are often indistinguishable from authentic speech to the human ear, which poses significant security risks. This rapid progress introduces the primary unsolved challenge for detection: generalization. Detectors frequently fail when confronted with unseen forgery methods

\*Corresponding author

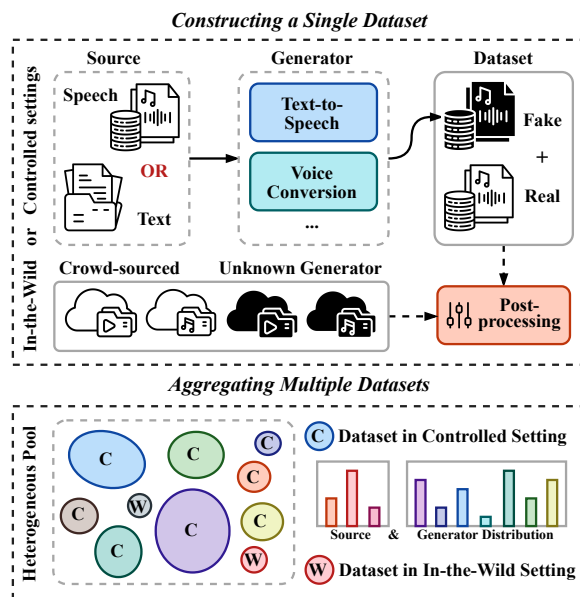


Figure 1: **The data landscape for speech deepfake detection**, from constructing a single dataset to aggregating multiple heterogeneous domains.

from ever-evolving generative systems. Furthermore, their performance can degrade significantly due to shifts in acoustic conditions, such as channel effects or background noise. This creates a critical need for detectors that are not only accurate but also robustly generalizable.

The community has largely pursued model-centric and algorithm-centric solutions to improve generalization. These include exploring novel architectures (Jung et al., 2022), leveraging large self-supervised (SSL) pretrained models (Tak et al., 2022b), and adopting advanced training algorithms (Huang et al., 2025c) or data augmentation techniques (Tak et al., 2022a). A common limitation of these works, however, is their reliance on fixed and limited training benchmarks. In parallel, a diverse array of new datasets for SDD has emerged (Zhao et al., 2024; Bhagtani et al., 2025); yet, these valuable resources are typically used in

isolated train/test settings or are simply held out as unseen test sets. We argue that a crucial yet often underexplored factor is data composition.

This paper aims to achieve generalizable SDD through a data-centric perspective. Unlike traditional audio tasks that rely on collected real-world data, the main portion of data for SDD is synthetically generated. This provides a unique opportunity to control and systematically study the data composition process, as a theoretically infinite number of samples can be generated.

To formalize our study, we analyze the SDD data landscape from two practical perspectives, as illustrated in Figure 1. First, we examine the process of constructing a single dataset. This can occur in two primary settings. In a controlled setting, developers typically begin with a known source (often real speech datasets that provide speech or text) and then select a specific generator (e.g., text-to-speech or voice conversion) to synthesize deepfakes. In an in-the-wild setting, audio is often crowdsourced from online platforms, where the underlying source and generator are typically unknown. Despite these differences in methodology, we identify two fundamental components that define any synthetic sample: the *source* it was derived from and the *generator* used to create it. Second, we address the common task of aggregating multiple datasets. Given that each dataset possesses a distinct combination of sources, generators, and acoustic conditions, this aggregation inevitably creates a heterogeneous data pool. This resulting pool is characterized by significant variance in data size, source types, and synthesis methods, making it a complex target for training. This framework motivates two central research questions:

- **RQ1:** When constructing a dataset, what principles determine how to allocate resources between source/generator diversity and data volume?
- **RQ2:** When aggregating multiple datasets, what is the most effective and efficient data mixing strategy to maximize generalization?

Our investigation into these questions leads to the following key contributions:

- **Discovering Data Scaling Laws.** We conduct the first large-scale empirical study to characterize the data scaling laws for SDD, quantifying the distinct impacts of source and generator diversity on model generalization (Section 3).

- **A Principled Data Mixing Strategy.** Based on these scaling laws, we propose the Diversity-Optimized Sampling Strategy (DOSS), a principled framework for efficiently training on heterogeneous data mixtures (Section 4).
- **State-of-the-Art Generalization.** We validate our data-centric approach by applying the DOSS framework to a large-scale, heterogeneous data pool, achieving state-of-the-art generalization performance (Section 5).

## 2 Related Work

**Generalization in SDD.** Generalization remains a core challenge in SDD. A large proportion of research has sought to improve this by focusing on model design. Initial attempts included specialized, compact models designed to effectively extract temporal and spectral information from audio (Tak et al., 2021; Jung et al., 2022). More recently, self-supervised learning (SSL) pre-trained models have become a common approach. SSL front-ends like Wav2Vec2 (Baeovski et al., 2020), XLS-R (Babu et al., 2022), and WavLM (Chen et al., 2022) are used to extract rich, intrinsic speech representations, which are then fine-tuned with a classifier backend. This method has shown strong performance and improved generalization (Tak et al., 2022b; Guo et al., 2024). A parallel direction focuses on algorithmic improvements to the training strategy. This includes advanced data augmentation techniques (Tak et al., 2022a; Wang and Yamagishi, 2024), novel training objectives (Zhang et al., 2021; Huang et al., 2025a), and refined optimization processes (Huang et al., 2025c).

While these attempts have advanced the field, most rely on fixed and limited training benchmarks, largely overlooking the impact of data composition. A few recent exceptions have begun to explore a data-centric paradigm to address these limitations. For instance, Combei et al., 2025 bridge the gap between scientific and real-world deepfakes using dataset selection and sample-level pruning, such as margin-based selection, to remove redundant or noisy data. In parallel, the work by Ge et al., 2025 demonstrated that post-training SSL models on 74k hours of speech improve generalization to unseen deepfakes; yet, this result relied on naive data aggregation which combined sources without regarding their variations. This paper builds upon this line of inquiry, shifting the focus from the sheer volume of data to its principled composition.

**Scaling Laws.** Recent research has found that neural model performance scales predictably as a power-law of data size, model size, or computation (Kaplan et al., 2020). This principle has been shown to apply broadly across diverse domains, including computer vision (Zhai et al., 2022), multi-modal learning (Aghajanyan et al., 2023), robotic manipulation (Lin et al., 2025), and speech recognition (Chen et al., 2025). Understanding these laws is crucial for informing training decisions (Hoffmann et al., 2022) and enabling more effective resource allocation during model development (Achiam et al., 2023). In this work, we extend this paradigm to SDD, examining how generalization performance scales with source and generator diversity, as well as sample volume, to inform principled data collection strategies.

**Data Mixing.** A highly relevant line of research is data mixing, which explores how to combine data from different sources for training large-scale models. Early approaches relied on manual heuristics, such as applying sampling caps to prevent domain dominance (Raffel et al., 2020) or explicitly oversampling domains perceived to be of high quality (Brown et al., 2020). More recent work has automated this process. These methods are typically either offline, using a proxy model to find a static set of optimal weights before training (Xie et al., 2023; Fan et al., 2024; Liu et al., 2025), or online, dynamically adjusting sampling weights during training based on the model’s real-time state (Albalak et al., 2023; Chen et al., 2023). A key limitation of many of these automated methods is their reliance on optimizing performance for a fixed set of known domains. We argue that performance on a specific validation set does not necessarily reflect true generalization to completely unseen data. In contrast, our approach avoids proxy models and their reliance on fixed validation sets. We instead use a more fine-grained domain definition and a sampling strategy derived from fundamental scaling principles, making it inherently designed for robust out-of-domain generalization.

### 3 Discovering Data Scaling Laws

To answer RQ1, this section aims to empirically discover the data scaling laws for SDD by investigating two fundamental factors: *generator diversity* and *source diversity*. Our objective is to model the relationship between generalization performance and these key data composition variables:

- $N_S$ : The number of distinct sources.
- $N_G$ : The number of distinct generators.
- $V$ : The volume of samples per unit of diversity.

#### 3.1 Experimental Setup

To isolate the effects of data composition, all scaling law experiments share a common methodology for model training and evaluation, with only the training data varying. Full implementation details for our data, model, and evaluation sets can be found in Appendix A and B.

**Data Generation.** To create controlled training sets for our experiments, we created two distinct pools of synthetic data:

- For the source diversity experiments, we selected 8 distinct source datasets. For each source, we selected 10k real samples and applied 4 generators to synthesize 10k fake samples each, creating a pool of 40k fake samples per source.
- For generator diversity experiments, we selected 16 distinct generators. For each generator, we selected a pool of 40k fake samples derived from two fixed source datasets.

**Model Training.** To ensure our results are robust and not dependent on a specific random selection of sources or generators, we performed 3 independent runs for each experimental condition (e.g., for a specific value of  $N_S$  and  $V$ ). Each run used a different random combination of the available sources or generators.

**Evaluation Protocol.** We measure generalization on 10 out-of-domain test sets. For each model, we compute the macro-average Equal Error Rate (EER) and Accuracy (ACC) across all 10 sets. Since ACC relies on a fixed threshold (0.5) while EER finds an optimal one, the two metrics capture different aspects of model performance. To provide a single, robust metric that balances both, we introduce the Calibrated Detection Error (CDE), defined as the harmonic mean of EER and (1-ACC):

$$CDE = \frac{2 \cdot EER \cdot (1 - ACC)}{EER + (1 - ACC)} \quad (1)$$

#### 3.2 Empirical Analysis

Figure 2 summarizes the results of our experiments on scaling with source and generator diversity. Our analysis of these results reveals three key findings that characterize the relationship between data composition and model generalization.

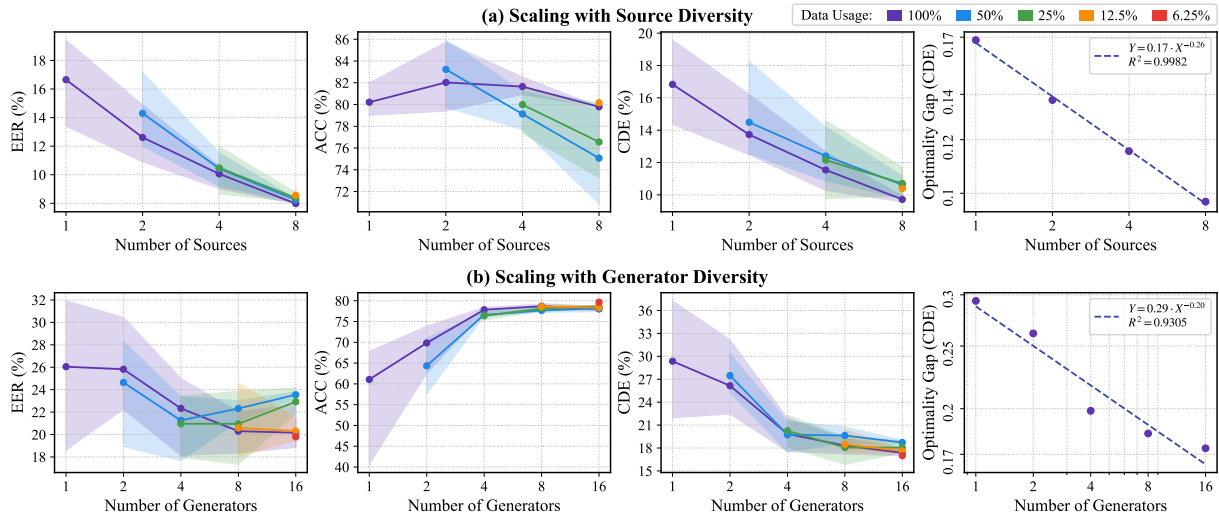


Figure 2: **Scaling with source and generator diversity.** Columns 1-3 plot generalization performance (EER, ACC, CDE) versus the number of training sources (a) or generators (b) on a logarithmic scale. Colored lines and shaded regions represent the average and min/max range, respectively, across experimental trials for different data usage percentages. Column 4 confirms a power-law fit for the Optimality Gap (CDE) at full data usage, with the fit equation and correlation coefficient provided in the legend.

**Finding 1: Diversity is the primary driver of generalization.** Our first key finding is that for a fixed data budget, increasing diversity both in sources and generators yields substantially greater performance gains than increasing the data volume from fewer diverse units. A model trained on 8 sources with 12.5% data usage consistently outperforms a model trained on 1 source with 100% usage. A similar trend is observed for generator diversity, where a model trained on 16 generators with just 6.25% of the data from each (2,500 samples per generator) achieves the best overall performance. This demonstrates the principle of diminishing returns. While increasing data volume is beneficial when diversity is low, its impact plateaus quickly beyond a saturation point. For data-efficient training, acquiring new information (increasing diversity) is a more effective strategy than reinforcing old information (increasing volume).

**Finding 2: The divergent roles of source and generator diversity.** While both forms of diversity improve overall performance (CDE), a closer analysis of EER and ACC reveals that they play surprisingly different and complementary roles.

- **Source diversity improves discrimination.** Increasing the number of sources consistently lowers EER but can degrade ACC. We hypothesize this is because a variety of sources helps the model build a more robust manifold of genuine speech, enhancing the fundamental separability

of the real and fake score distributions.

- **Generator diversity improves identification.** Conversely, increasing the number of generators strongly improves ACC, while its effect on EER is less consistent. We hypothesize this is because exposure to diverse forgery artifacts makes the model a better identifier, learning to confidently classify fakes but potentially at the cost of a less distinct real/fake decision boundary.

**Finding 3: Generalization follows a predictable power Law.** Finally, we find that despite the complex behaviors of individual metrics, overall generalization performance follows a predictable pattern. As shown in Figure 2, the optimality gap (CDE) scales as a power law with both source and generator diversity, confirmed by high correlation coefficients (Pearson’s  $R^2$ ). This result demonstrates that generalization in SDD is not a random process but a principled, modelable phenomenon. We conclude that investing in data diversity represents an efficient and effective strategy for improving detector robustness.

## 4 Principled Data Mixing with DOSS

Our scaling law analysis provides a clear prescription: maximize data diversity. In practice, the most straightforward path to achieving this is to aggregate the numerous datasets. This action, however, creates a heterogeneous data pool with significant imbalances, directly raising the challenge posed in

---

**Algorithm 1** DOSS-Select

---

**Require:** Sets  $\mathcal{F}, \mathcal{R}$ ; Size  $n_d$  for each domain  $d$ ;  
Parameters  $N_c, \rho$ ; Function  $\text{base}(f)$   
**Initialize:**  $s$  as an empty map  
// **Step 1:** Compute counts for fake domains  
1: **for**  $f \in \mathcal{F}$  **do**  
2:  $s[f] \leftarrow \min(n_f, N_c)$   
3: **end for**  
// **Step 2:** Compute counts for real domains  
4: **for**  $r \in \mathcal{R}$  **do**  
5:  $\Sigma_r \leftarrow \sum_{f \in \mathcal{F}: \text{base}(f)=r} s[f]$   
6:  $s[r] \leftarrow \min(n_r, \Sigma_r \times \rho)$   
7: **end for**  
8: **return**  $s$  // Return final domain counts

---

RQ2: what is the most effective and efficient strategy for mixing this data to best utilize its diversity?

Inspired by the principle of maximum entropy, we argue that for robust generalization against unpredicted future attacks, there is no prior evidence that any single known attack is more important than another. The most robust and least-biased assumption, therefore, is that all attacks should be treated as equally important. This leads us to propose the *Diversity-Optimized Sampling Strategy (DOSS)*, a framework designed to apply this principle to the realistic challenges of heterogeneous data mixtures.

#### 4.1 DOSS in Practice

The DOSS framework operates on granular domains. We first index our entire heterogeneous data pool and define these domains as follows:

- For real data, a domain is its source. We denote the set of all real domains as  $\mathcal{R}$ .
- For fake data, a domain is the combination of its source and generator. We denote the set of all fake domains as  $\mathcal{F}$ . We use a function  $\text{base}(f)$  to obtain the original source and to map the domains in  $\mathcal{F}$  with the domains in  $\mathcal{R}$ .

DOSS approximates the uniform distribution through two strategies: a pruning method (DOSS-Select) and a re-weighting method (DOSS-Weight).

**DOSS-Select.** DOSS-Select is a data pruning strategy that creates a smaller, more balanced, and data-efficient training subset. This method is directly motivated by our finding in Section 3 that data volume provides diminishing returns beyond a saturation point. We denote this saturation cap as  $N_c$ . As detailed in Algorithm 1, the process first

---

**Algorithm 2** DOSS-Weight

---

**Require:** Sets  $\mathcal{F}, \mathcal{R}$ ; Size  $n_d$  for each domain  $d$ ;  
Parameters  $N_c, \tau, \rho$ ; Function  $\text{base}(f)$   
**Initialize:**  $s, w$  as empty maps  
// **Step 1.** Compute weights for fake domains  
1: **for**  $f \in \mathcal{F}$  **do**  
2:  $s[f] \leftarrow \min(n_f, N_c)$   
3:  $w[f] \leftarrow (s[f])^{1/\tau}$   
4: **end for**  
// **Step 2.** Compute weights for real domains  
5: **for**  $r \in \mathcal{R}$  **do**  
6:  $\Sigma_r \leftarrow \sum_{f \in \mathcal{F}: \text{base}(f)=r} s[f]$   
7:  $s[r] \leftarrow \min(n_r, \Sigma_r \times \rho)$   
8:  $w[r] \leftarrow (s[r])^{1/\tau}$   
9: **end for**  
// **Step 3.** Enforce global real-to-fake ratio  
10:  $W_{\mathcal{F}} \leftarrow \sum_{f \in \mathcal{F}} w[f]$ ;  $W_{\mathcal{R}} \leftarrow \sum_{r \in \mathcal{R}} w[r]$   
11:  $\alpha_{\text{adj}} \leftarrow (W_{\mathcal{F}} \times \rho) / W_{\mathcal{R}}$   
12: **for**  $r \in \mathcal{R}$  **do**  
13:  $w[r] \leftarrow w[r] \times \alpha_{\text{adj}}$   
14: **end for**  
15: **return**  $w$  // Return final domain weights

---

determines the number of samples to select from each fake domain by capping its original size at  $N_c$ . Then, for each real domain, it selects a proportional number of real samples based on the aggregated fake count and ratio  $\rho$ , ensuring that the volume of real speech scales dynamically to preserve a consistent local class balance. The output is a list of counts for each domain, which guides the construction of the final pruned dataset.

**DOSS-Weight.** DOSS-Weight is a re-weighting strategy that uses the entire data pool but adjusts the sampling probability of each domain at training time. This allows the model to see data from smaller domains more frequently without discarding any samples. The process, detailed in Algorithm 2, involves three steps. First, it calculates an initial weight for each fake domain by capping its size at  $N_c$  and applying a diversity temperature  $\tau$ . Second, it calculates weights for real domains using the same proportional logic as DOSS-Select. Finally, it computes a global adjustment factor to scale all real domain weights, ensuring that the total sampling probability across all real and fake domains strictly adheres to the target ratio  $\rho$ . The output is a list of final domain weights used to guide a weighted random sampler during training.

Table 1: **Generalization performance in EER% ( $\downarrow$ ) for different data selection and mixing strategies.** Within each strategy section, the best and second-best results per column are in **bold** and underline. The overall top two results across all experiments are highlighted with **darker** and **lighter** grey backgrounds.

Strategy	#Hours	AVG	ASV19	DECR	ITW	SC	FOR	EF	ADD22	ADD23	CFAD	ODSS
<i>Traditional Benchmarks: Training on single datasets</i>												
ADD22	24	28.23	32.73	37.98	20.63	<u>15.77</u>	44.60	1.90	15.17 <sup>†</sup>	45.68	25.19	42.61
ASV19	25	<b>12.06</b>	<b>0.26</b> <sup>†*</sup>	<u>9.19</u>	<b>7.50</b>	20.24	<u>4.59</u>	<b>0.14</b>	11.80	<b>16.27</b>	<u>21.01</u>	29.63
FOR	48	<u>13.53</u>	3.24	31.49	15.39	<b>14.68</b>	<b>0.38</b> <sup>†</sup>	2.43	<b>4.19</b>	20.23	<b>15.23</b>	<u>28.07</u>
DECR	58	19.25	<u>0.39</u>	<b>0.02</b> <sup>†*</sup>	<u>7.85</u>	47.93	65.06	<u>1.74</u>	<u>9.58</u>	<u>20.03</u>	22.49	<b>17.43</b>
<i>Baseline: Naive aggregation with <math>k</math> datasets</i>												
$k = 2$	0.9k	8.78	<u>0.16</u> *	11.29	4.32	9.76	3.04	<u>0.03</u>	8.77	19.55	23.65	7.27
$k = 6$	1.3k	6.45	<b>0.10</b> *	<b>0.13</b> *	1.71	14.32	1.35	<b>0.01</b>	2.09	<u>12.57</u>	20.73	11.49
$k = 8$	3.3k	<u>4.37</u>	<u>0.37</u> *	<u>0.19</u> *	<u>1.57</u>	<u>0.05</u>	<u>0.93</u>	0.57	<u>1.94</u>	16.18	<u>15.72</u>	<b>6.20</b>
$k = 12$	6.4k	<b>3.29</b>	<u>0.19</u> *	<u>0.31</u> *	<b>1.21</b>	<b>0.05</b>	<b>0.72</b>	0.29	<b>1.94</b>	<b>6.44</b>	<b>15.14</b>	<u>6.57</u>
<i>DOSS-Select: Pruning via saturation cap (<math>N_c</math>)</i>												
$N_c = 100$	40	3.43	<u>1.31</u> *	<u>0.67</u> *	1.63	2.08	0.93	<u>0.11</u>	1.54	6.70	14.49	4.88
$N_c = 500$	0.2k	2.77	<u>0.32</u> *	<u>0.35</u> *	1.34	0.83	0.17	<b>0.11</b>	1.63	<u>4.42</u>	14.69	<b>3.85</b>
$N_c = 2500$	0.8k	<b>2.69</b>	<b>0.18</b> *	<b>0.17</b> *	1.24	<u>0.47</u>	<b>0.06</b>	0.15	<b>1.26</b>	<b>4.22</b>	15.27	3.90
$N_c = 12500$	2.9k	<u>2.72</u>	<u>0.21</u> *	<u>0.36</u> *	<b>1.23</b>	<b>0.21</b>	<u>0.13</u>	0.15	<u>1.40</u>	4.90	<b>14.20</b>	4.39
<i>DOSS-Weight: Re-weighting via saturation (<math>N_c</math>) &amp; temperature (<math>\tau</math>)</i>												
$N_c = 50000, \tau = 1$	6.4k	2.81	<u>0.20</u> *	<u>0.50</u> *	1.23	<b>0.16</b>	0.13	0.17	1.52	4.84	14.69	4.68
$N_c = 2500, \tau = 1$	6.4k	2.51	<u>0.12</u> *	<u>0.15</u> *	<u>1.22</u>	0.29	0.17	<u>0.09</u>	<u>1.31</u>	3.39	<b>13.00</b>	5.35
$N_c = 2500, \tau = 5$	6.4k	<b>2.34</b>	<b>0.09</b> *	<b>0.13</b> *	<b>1.16</b>	0.24	<b>0.08</b>	0.11	<b>1.25</b>	<u>2.97</u>	13.91	<b>3.44</b>
$N_c = 2500, \tau = 100$	6.4k	<u>2.41</u>	<u>0.10</u> *	<u>0.16</u> *	1.23	<u>0.21</u>	<u>0.08</u>	<b>0.07</b>	1.40	<b>2.56</b>	<u>13.89</u>	<u>4.42</u>

\*: in-domain test set; †: test set corresponding to the original traditional benchmark; all other columns are out-of-domain.

## 4.2 Experimental Validation

To validate the proposed DOSS framework, we conduct a comparative study using a fixed model architecture across different data selection and mixing strategies. Generalization performance is evaluated on 10 distinct test sets, with the primary results shown in Table 1. Full details of the experimental setup can be found in Appendix A and B.

**Traditional Benchmarks.** We first examine the performance of models trained under the traditional paradigm: using a single dataset for training and evaluating on both in-domain and out-of-domain test sets. The results clearly show that models trained this way are brittle and fail to generalize effectively. While they often achieve excellent performance on their corresponding in-domain test sets, their performance degrades significantly on out-of-domain data. For instance, the model trained only on DECR achieves a near-perfect 0.02% EER on its own test set but has a poor average EER of 19.25% across all ten sets. This demonstrates that single-dataset training encourages the model to overfit to dataset-specific biases rather than learning a universal representation of artificiality.

**Baseline.** Next, we establish a baseline by naively aggregating an increasing number of train-

ing datasets ( $k$ ). Simply increasing the size and diversity of the training pool leads to a substantial improvement in overall generalization, with the average EER dropping from 8.78% for  $k=2$  to 3.29% for  $k=12$ . However, this ‘brute force’ approach has a notable flaw. In such a heterogeneous and imbalanced pool, the training process is naturally dominated by the larger datasets, causing the model to prioritize the most common features. This can lead to negative transfer, where the model de-prioritizes artifacts from smaller domains, degrading performance on specific test sets as more data is added. For instance, performance on the ADD23 test set worsens from an EER of 12.57% when using  $k=6$  datasets to 16.18% with  $k=8$  datasets. This instability demonstrates that while naive aggregation is an improvement over single-dataset training, it is an unpredictable and suboptimal strategy, highlighting the need for a principled approach to data mixing.

**DOSS-Select.** Our evaluation of DOSS-Select demonstrates that principled data pruning is highly efficient and outperforms naive aggregation. With a saturation cap of  $N_c = 100$  (40 hours of data), the model achieves a 3.43% EER, which is better than all traditional benchmarks and most naive aggregation settings. Furthermore, increasing the cap to  $N_c = 500$  uses just 0.2k hours ( $\approx 3\%$  of the data)

but yields a 2.77% EER. This result is particularly striking as it surpasses the best naive aggregation baseline that required the full 6.4k hours. This confirms that effective generalization is driven less by raw data volume and more by diverse, well-balanced composition.

The choice of  $N_c$  is critical. Performance improves as the cap increases, reaching the overall best EER of 2.69% at  $N_c = 2500$  (0.8k hours). However, increasing the cap further to  $N_c = 12500$  provides no additional benefit, with performance plateauing. We attribute this to two factors. First, the model has likely reached the data saturation point for learning from the available domains. Second, Figure 3(a) in the Appendix shows that the domain distribution for  $N_c = 12500$  is less uniform than for  $N_c = 2500$ , as the high cap re-introduces the natural imbalance of the original data pool.

**DOSS-Weight.** In contrast to pruning, our DOSS-Weight strategy re-weights the entire data pool and achieves the most effective results overall. The optimal setting ( $N_c = 2500, \tau = 5$ ) yields an average EER of 2.34%, the lowest across all tested methods. This corresponds to a relative error reduction of approximately 29% compared to the best naive aggregation baseline and 13% compared to the best DOSS-Select result.

The choice of saturation cap  $N_c$  remains critical. A poorly chosen  $N_c$  can undermine the strategy’s effectiveness. For instance, a large cap ( $N_c = 50000$ ) with no temperature balancing ( $\tau = 1$ ) creates a distribution similar to naive pooling (Figure 3(b)) and performs worse than DOSS-Select. In contrast, a more reasonable cap like  $N_c = 2500$  creates a more uniform base distribution for the temperature to act upon.

The temperature parameter  $\tau$  is key to refining this distribution. At  $N_c = 2500$ , increasing  $\tau$  from 1 to 5 improves performance from 2.51% to 2.34% EER. The temperature primarily balances the sampling probability among the real domains, which have more varied initial weights. However, increasing  $\tau$  further to 100 provides no additional benefit, as the change in distribution is minimal.

Ultimately, the performance advantage of DOSS-Weight over DOSS-Select suggests that re-weighting is a more powerful strategy than pruning. Instead of discarding potentially redundant data, it is better to retain it as intra-domain variations and down-weight its influence during training.

## 5 Scalability and Application of DOSS

This section validates the robustness and scalability of the DOSS framework. We expand our data pool, apply DOSS to train final models, and evaluate them against state-of-the-art methods.

### 5.1 Data Pool Curation

To build a comprehensive training set and test our framework’s scalability, we expanded our data pool using the following data curation pipeline:

1. **Collection:** We aggregated 17 publicly available datasets, unified their audio formats, and gathered metadata on their respective sources and generators.
2. **Reorganization:** We de-duplicated the pool by identifying fake datasets that shared common real source corpora (e.g., VCTK, LibriTTS) and replaced the redundant real audio with canonical source versions.
3. **Enrichment:** We analyzed the pool for gaps in generator diversity and synthesized new data from recent models to ensure our final pool reflects the current state of speech synthesis.

This curation process resulted in a new 12k-hour data pool with enhanced diversity and reduced redundancy. Full details are in Appendix A.3.

### 5.2 Final Model Performance

Using the new data pool, we train final models (based on XLS-R (Babu et al., 2022), details in B) using the DOSS-Weight strategy. To validate its performance, we conducted a comprehensive evaluation against both established public benchmarks and a new set of commercial APIs.

**Performance on Public Benchmarks.** Table 2 presents the out-of-domain evaluation on 8 public test sets. Our results highlight two findings regarding data efficiency and model scaling.

First, we compare our data-centric approach against the large-scale baselines established by Ge et al., 2025. Their best-performing system utilizes a massive XLS-R-2B backbone trained on 74k hours of data to achieve an average EER of 3.94%. In contrast, our DOSS-trained model, using a smaller XLS-R-300M model and our 12k-hour curated data pool, achieves a significantly lower average EER of 2.14%. This result highlights a clear efficiency advantage. By prioritizing generator diversity over raw data volume, we surpass the performance of

Table 2: **Out-of-domain EER% ( $\downarrow$ ) comparison with prior works.** The State-of-the-Art (SOTA) results are collected from different systems in the literature. The best and second-best results are in **bold** and underline.

System	#Params	#Hours	AVG	ITW	FOR	EF	ADD22	ADD23	ODSS	DV	FSW
<i>Existing Benchmarks (Collected in Table 9)</i>											
SOTA	–	–	–	1.23	0.97	<u>0.20</u>	<u>1.05</u>	4.67	<b>1.13</b>	2.27	11.58
<i>Training on Naive Aggregation (Ge et al., 2025)</i>											
MMS-300M	317M	74k	6.61	2.90	6.08	1.58	2.64	7.96	13.34	2.27	16.15
MMS-1B	965M	74k	7.24	1.82	1.73	0.34	2.76	9.05	5.49	2.47	23.81
XLS-R-1B	965M	74k	4.47	1.37	12.15	0.24	1.68	5.39	1.53	2.35	21.45
XLS-R-2B	2.2B	74k	3.94	1.23	1.73	<u>0.20</u>	<u>1.05</u>	4.67	<b>1.13</b>	2.35	19.14
<i>Training with DOSS-Weight (Ours)</i>											
XLS-R-300M	317M	12k	<u>2.14</u>	<u>0.81</u>	<u>0.25</u>	0.29	<b>0.82</b>	<u>3.63</u>	1.65	<u>0.97</u>	<u>8.70</u>
XLS-R-1B	965M	12k	<b>1.65</b>	<b>0.80</b>	<b>0.13</b>	<b>0.10</b>	1.40	<b>2.25</b>	<u>1.23</u>	<b>0.86</b>	<b>6.47</b>

Table 3: **Detection ACC% ( $\uparrow$ ) on commercial APIs.** Results are evaluated on 5,000 synthetic samples per API (2,500 English + 2,500 Chinese). The best and second-best results are in **bold** and underline.

System	AVG	Google	Microsoft	OpenAI	11Labs	Alibaba	Baidu	iFlytek	MiniMax	Qwen3
<i>Training on Naive Aggregation (Ge et al., 2025)</i>										
XLS-R-2B	86.31	99.54	85.42	<b>99.20</b>	81.90	96.06	98.96	99.34	76.48	39.88
<i>Training with DOSS-Weight (Ours)</i>										
XLS-R-300M	<u>92.81</u>	<u>99.96</u>	<u>99.84</u>	<u>98.04</u>	<b>92.12</b>	<u>98.92</u>	<b>100.00</b>	<b>99.98</b>	<u>90.50</u>	<u>55.94</u>
XLS-R-1B	<b>96.01</b>	<b>99.98</b>	<b>99.92</b>	<u>98.92</u>	<u>86.30</u>	<b>99.78</b>	<u>99.98</u>	<u>99.96</u>	<b>91.94</b>	<b>87.32</b>

a model with roughly 7 times more parameters trained on 6 times more data.

Second, we observe robust scaling behavior within our framework. Scaling the backbone from 300M to 1B parameters yields further performance gains and reduces the average EER by approximately 23%. Notably, our 1B model establishes a new state-of-the-art on 6 out of the 8 individual benchmarks, surpassing the aggregated best results from prior literature. This suggests that the increased capacity allows the model to better capture the diverse artifacts present in the training data.

**Performance on Commercial APIs.** To further evaluate our model’s performance against current, real-world threats, we created a new challenge set using 9 different commercial APIs. This complements our previous results on public, open-source test sets, which may not represent the latest generation of synthesis technology. We generated 2,500 synthetic samples from the latest version of each API in both English and Chinese (full details in Appendix A.4) and evaluated our model alongside the large-scale baseline from Ge et al., 2025.

Table 3 reports the overall detection accuracy. Our DOSS-trained models exhibit remarkable generalization compared to the naive baselines. While the massive XLS-R-2B baseline achieves an average accuracy of 86.31%, our XLS-R-1B model im-

proves this by nearly 10 absolute percentage points to 96.01%. This advantage is most visible on the most challenging, high-fidelity synthesis systems. For instance, on the Qwen3 TTS where the baseline accuracy collapses to 39.88%, our 1B model maintains a high accuracy of 87.32%. Similarly, on MiniMax, our model improves detection from 76.48% to 91.94%. This suggests that the wild diversity in our training set transfers effectively to unseen, advanced commercial systems. Note that performance varies by language for certain providers; a detailed breakdown of these linguistic differences is provided in Table 10 in Appendix C.

## 6 Conclusion

This paper proposes a data-centric approach to generalization in speech deepfake detection. We first conducted a large-scale empirical study demonstrating that source and generator diversity are more impactful for generalization than raw data volume. Guided by these scaling laws, we introduced the Diversity-Optimized Sampling Strategy (DOSS) to effectively manage heterogeneous data mixtures. Our experiments validate that this principled approach enables both high data efficiency through pruning and superior representation through reweighting, proving that smart data composition is superior to naive aggregation. Consequently, our final model establishes a new state-of-the-art, sur-

passing large-scale baselines on both public benchmarks and a commercial API challenge set with significantly greater data and model efficiency.

## Limitations

Our work focuses on generalization from a data-centric perspective. To isolate the effects of data composition, we primarily utilized a fixed model architecture and training configuration. Consequently, while we validated the scalability of our approach by extending from 300M to 1B parameters, this study does not explore the exhaustive interplay with other factors like massive model scaling or varied computational budgets. Investigating how these data scaling laws and mixing strategies co-adapt across a broader spectrum of model sizes remains a valuable direction for future work.

Furthermore, our curated data pool is linguistically concentrated on English and Chinese. While this reflects the composition of most publicly available datasets, it could limit our model’s generalization to other languages. Building a multi-lingual detector is a significant future challenge that would require not only the collection of rare, diverse-language datasets but also principled methods to manage the resulting linguistic imbalance.

## Ethical Considerations

This work focuses on developing generalizable defensive countermeasures against the security risks posed by speech deepfakes.

- **Licenses & Usage:** All scientific artifacts, including public datasets and pretrained models, were used in strict accordance with their original licenses (e.g., CC-BY 4.0, non-commercial research licenses). Use of commercial APIs for generating the challenge set adhered to their respective Terms of Service for research evaluation.
- **Intended Use:** Our research is consistent with the intended use of the original datasets: advancing the field of audio forensics and speech security. The methodologies presented are designed to protect individuals and organizations from malicious deepfake applications.
- **Data Privacy & Safety:** The curated data consist of speech from established public corpora that is pre-anonymized and contains no personally identifiable information.

## Acknowledgements

This work was supported in part by China NSFC project under Grants No. U25A20409, and in part by SJTU Med-X (Medicine & Engineering) Translational Research Grant (YG2025LC09).

The authors would like to thank Prof. Shinji Watanabe from CMU for his insightful discussions on this work.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Armen Aghajanyan, Lili Yu, Alexis Conneau, Wei-Ning Hsu, Karen Hambardzumyan, Susan Zhang, Stephen Roller, Naman Goyal, Omer Levy, and Luke Zettlemoyer. 2023. Scaling laws for generative mixed-modal language models. In *International Conference on Machine Learning*, pages 265–279.
- Alon Albalak, Liangming Pan, Colin Raffel, and William Yang Wang. 2023. Efficient online data mixing for language model pre-training. *arXiv preprint arXiv:2312.02406*.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis Tyers, and Gregor Weber. 2020. Common voice: A massively-multilingual speech corpus. In *Proceedings of the twelfth language resources and evaluation conference*, pages 4218–4222.
- Zhongjie Ba, Qing Wen, Peng Cheng, Yuwei Wang, Feng Lin, Li Lu, and Zhenguang Liu. 2023. Transferring audio deepfake detection capability across languages. In *Proceedings of the ACM Web Conference 2023*, pages 2033–2044.
- Arun Babu, Changan Wang, Andros Tjandra, Kushal Lakhota, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Pino, Alexei Baevski, Alexis Conneau, and Michael Auli. 2022. Xls-r: Self-supervised cross-lingual speech representation learning at scale. In *Interspeech*, pages 2278–2282.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460.
- James Betker. 2023. Better speech synthesis through scaling. *arXiv preprint arXiv:2305.07243*.
- Kratika Bhagtani, Amit Kumar Singh Yadav, Paolo Bestagini, and Edward J Delp. 2025. Diffssd: A

- diffusion-based dataset for speech forensics. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.
- Jordan J Bird and Ahmad Lotfi. 2023. Real-time detection of ai-generated speech for deepfake voice conversion. *arXiv preprint arXiv:2308.12734*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Hui Bu, Jiayu Du, Xingyu Na, Bengu Wu, and Hao Zheng. 2017. Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline. In *20th conference of the oriental chapter of the international coordinating committee on speech databases and speech I/O systems and assessment (O-COCOSDA)*, pages 1–5.
- Mayee Chen, Nicholas Roberts, Kush Bhatia, Jue Wang, Ce Zhang, Frederic Sala, and Christopher Ré. 2023. Skill-it! a data-driven skills framework for understanding and training language models. *Advances in Neural Information Processing Systems*, 36:36000–36040.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, and 1 others. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518.
- William Chen, Jinchuan Tian, Yifan Peng, Brian Yan, Chao-Han Huck Yang, and Shinji Watanabe. 2025. Owls: Scaling laws for multilingual speech recognition and translation models. In *Forty-second International Conference on Machine Learning*.
- Yushen Chen, Zhikang Niu, Ziyang Ma, Keqi Deng, Chunhui Wang, Jian Zhao, Kai Yu, and Xie Chen. 2024. F5-tts: A fairytaler that fakes fluent and faithful speech with flow matching. *arXiv preprint arXiv:2410.06885*.
- David Combei, Adriana Stan, Dan Oneata, Nicolas Müller, and Horia Cucu. 2025. Unmasking real-world audio deepfakes: A data-centric approach. In *Interspeech*, pages 5343–5347.
- Wei Deng, Siyi Zhou, Jingchen Shu, Jinchao Wang, and Lu Wang. 2025. Indextts: An industrial-level controllable and efficient zero-shot text-to-speech system. *arXiv preprint arXiv:2502.05512*.
- Jiawei Du, I-Ming Lin, I-Hsiang Chiu, Xuanjun Chen, Haibin Wu, Wenze Ren, Yu Tsao, Hung-yi Lee, and Jyh-Shing Roger Jang. 2024a. Dfadd: The diffusion and flow-matching based audio deepfake dataset. In *2024 IEEE Spoken Language Technology Workshop (SLT)*, pages 921–928. IEEE.
- Zhihao Du, Yuxuan Wang, Qian Chen, Xian Shi, Xiang Lv, Tianyu Zhao, Zhifu Gao, Yexin Yang, Changfeng Gao, Hui Wang, and 1 others. 2024b. Cosyvoice 2: Scalable streaming speech synthesis with large language models. *arXiv preprint arXiv:2412.10117*.
- Sefik Emre Eskimez, Xiaofei Wang, Manthan Thakker, Canrun Li, Chung-Hsien Tsai, Zhen Xiao, Hemin Yang, Zirun Zhu, Min Tang, Xu Tan, and 1 others. 2024. E2 tts: Embarrassingly easy fully non-autoregressive zero-shot tts. In *IEEE Spoken Language Technology Workshop (SLT)*, pages 682–689.
- Simin Fan, Matteo Pagliardini, and Martin Jaggi. 2024. Doge: Domain reweighting with generalization estimation. In *International Conference on Machine Learning*, pages 12895–12915.
- Wanying Ge, Xin Wang, Xuechen Liu, and Junichi Yamagishi. 2025. Post-training for deepfake speech detection. In *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*.
- Sang gil Lee, Wei Ping, Boris Ginsburg, Bryan Catanzaro, and Sungroh Yoon. 2023. BigVGAN: A universal neural vocoder with large-scale training. In *The Eleventh International Conference on Learning Representations*.
- Yinlin Guo, Haofan Huang, Xi Chen, He Zhao, and Yuehai Wang. 2024. Audio deepfake detection with self-supervised wavlm and multi-fusion attentive classifier. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12702–12706.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, and 1 others. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Rongjie Huang, Max WY Lam, Jun Wang, Dan Su, Dong Yu, Yi Ren, and Zhou Zhao. 2022a. Fastdiff: A fast conditional diffusion model for high-quality speech synthesis. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI)*.
- Rongjie Huang, Zhou Zhao, Huadai Liu, Jinglin Liu, Chenye Cui, and Yi Ren. 2022b. Prodiff: Progressive fast diffusion model for high-quality text-to-speech. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2595–2605.
- Wen Huang, Yanmei Gu, Zhiming Wang, Huijia Zhu, and Yanmin Qian. 2025a. Generalizable audio deepfake detection via latent space refinement and augmentation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.
- Wen Huang, Yanmei Gu, Zhiming Wang, Huijia Zhu, and Yanmin Qian. 2025b. SpeechFake: A large-scale multilingual speech deepfake dataset incorporating

- cutting-edge generation methods. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9985–9998.
- Wen Huang, Xuechen Liu, Xin Wang, Junichi Yamagishi, and Yanmin Qian. 2025c. From sharpness to better generalization for speech deepfake detection. In *Interspeech*, pages 5338–5342.
- Keith Ito and Linda Johnson. 2017. The lj speech dataset. <https://keithito.com/LJ-Speech-Dataset/>.
- Ziyue Jiang, Yi Ren, Ruiqi Li, Shengpeng Ji, Boyang Zhang, Zhenhui Ye, Chen Zhang, Bai Jionghao, Xiaoda Yang, Jialong Zuo, and 1 others. 2025. Megatts 3: Sparse alignment enhanced latent diffusion transformer for zero-shot speech synthesis. *arXiv preprint arXiv:2502.18924*.
- Jee-weon Jung, Hee-Soo Heo, Hemlata Tak, Hye-jin Shim, Joon Son Chung, Bong-Jin Lee, Ha-Jin Yu, and Nicholas Evans. 2022. Aasist: Audio anti-spoofing using integrated spectro-temporal graph attention networks. In *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 6367–6371.
- Jee-weon Jung, Yihan Wu, Xin Wang, Ji-Hoon Kim, Soumi Maiti, Yuta Matsunaga, Hye-jin Shim, Jinchuan Tian, Nicholas Evans, Joon Son Chung, and 1 others. 2025. Spoofceleb: Speech deepfake detection and sasv in the wild. *IEEE Open Journal of Signal Processing*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Jaehyeon Kim, Sungwon Kim, Jungil Kong, and Sungho Yoon. 2020. Glow-tts: A generative flow for text-to-speech via monotonic alignment search. *Advances in Neural Information Processing Systems*, 33:8067–8077.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in neural information processing systems*, 33:17022–17033.
- Kundan Kumar, Rithesh Kumar, Thibault De Boissiere, Lucas Gesteira, Wei Zhen Teoh, Jose Sotelo, Alexandre De Brebisson, Yoshua Bengio, and Aaron C Courville. 2019. Melgan: Generative adversarial networks for conditional waveform synthesis. *Advances in neural information processing systems*, 32.
- Yuang Li, Min Zhang, Mengxin Ren, Xiaosong Qiao, Miaomiao Ma, Daimeng Wei, and Hao Yang. 2024. Cross-domain audio deepfake detection: Dataset and analysis. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4977–4983.
- Fanqi Lin, Yingdong Hu, Pingyue Sheng, Chuan Wen, Jiacheng You, and Yang Gao. 2025. Data scaling laws in imitation learning for robotic manipulation. In *The Thirteenth International Conference on Learning Representations*.
- Qian Liu, Xiaosen Zheng, Niklas Muennighoff, Guangtao Zeng, Longxu Dou, Tianyu Pang, Jing Jiang, and Min Lin. 2025. Regmix: Data mixture as regression for language model pre-training. In *The Thirteenth International Conference on Learning Representations*.
- Songxiang Liu, Dan Su, and Dong Yu. 2022. Diffgan-tts: High-fidelity and efficient text-to-speech with denoising diffusion gans. *arXiv preprint arXiv:2201.11972*.
- Haoxin Ma, Jiangyan Yi, Chenglong Wang, Xinrui Yan, Jianhua Tao, Tao Wang, Shiming Wang, and Ruibo Fu. 2024. Cfad: A chinese dataset for fake audio detection. *Speech Communication*, 164:103122.
- Nicolas Müller, Pavel Czepin, Franziska Diekmann, Adam Froggyar, and Konstantin Böttinger. 2022. Does audio deepfake detection generalize? In *Interspeech*, pages 2783–2787.
- Arsha Nagrani, Joon Son Chung, and Andrew Senior. 2017. Voxceleb: a large-scale speaker identification dataset. In *Interspeech*, pages 2616–2620.
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. Mls: A large-scale multilingual dataset for speech research. In *Interspeech*, pages 2757–2761.
- Zengyi Qin, Wenliang Zhao, Xumin Yu, and Xin Sun. 2023. Openvoice: Versatile instant voice cloning. *arXiv preprint arXiv:2312.01479*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Ricardo Reimao and Vassilios Tzerpos. 2019. For: A dataset for synthetic speech detection. In *2019 International Conference on Speech Technology and Human-Computer Dialogue (SpeD)*, pages 1–10.
- Davide Salvi, Brian Hosler, Paolo Bestagini, Matthew C Stamm, and Stefano Tubaro. 2023. Timit-tts: A text-to-speech dataset for multimodal synthetic media detection. *IEEE access*, 11:50851–50866.
- Yao Shi, Hui Bu, Xin Xu, Shaoji Zhang, and Ming Li. 2021. Aishell-3: A multi-speaker mandarin tts corpus and the baselines. In *Interspeech*, pages 2756–2760. ISCA.
- Hemlata Tak, Madhu Kamble, Jose Patino, Massimiliano Todisco, and Nicholas Evans. 2022a. Rawboost:

- A raw data boosting and augmentation method applied to automatic speaker verification anti-spoofing. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6382–6386.
- Hemlata Tak, Jose Patino, Massimiliano Todisco, Andreas Nautsch, Nicholas Evans, and Anthony Larcher. 2021. End-to-end anti-spoofing with rawnet2. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6369–6373.
- Hemlata Tak, Massimiliano Todisco, Xin Wang, Jee-weon Jung, Junichi Yamagishi, and Nicholas Evans. 2022b. Automatic speaker verification spoofing and deepfake detection using wav2vec 2.0 and data augmentation. In *The Speaker and Language Recognition Workshop (Odyssey 2022)*, pages 112–119.
- Christophe Veaux, Junichi Yamagishi, and Simon King. 2013. The voice bank corpus: Design, collection and data analysis of a large regional accent speech database. In *International conference oriental COCOSDA held jointly with 2013 conference on Asian spoken language research and evaluation (O-COCOSDA/CASLRE)*, pages 1–4.
- Xin Wang, Hector Delgado, Hemlata Tak, Jee-weon Jung, Hye-jin Shim, Massimiliano Todisco, Ivan Kukanov, Xuechen Liu, Md Sahidullah, Tomi Kinnunen, and 1 others. 2024a. Asvspoof 5: Crowdsourced speech data, deepfakes, and adversarial attacks at scale. In *The Automatic Speaker Verification Spoofing Countermeasures Workshop (ASVspoof 2024)*, pages 1–8.
- Xin Wang and Junichi Yamagishi. 2024. Can large-scale vocoded spoofed data improve speech spoofing countermeasure with a self-supervised front end? In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 10311–10315.
- Xin Wang, Junichi Yamagishi, Massimiliano Todisco, Héctor Delgado, Andreas Nautsch, Nicholas Evans, Md Sahidullah, Ville Vestman, Tomi Kinnunen, Kong Aik Lee, and 1 others. 2020. Asvspoof 2019: A large-scale public database of synthesized, converted and replayed speech. *Computer Speech & Language*, 64:101114.
- Yuancheng Wang, Haoyue Zhan, Liwei Liu, Ruihong Zeng, Haotian Guo, Jiachen Zheng, Qiang Zhang, Xueyao Zhang, Shunsi Zhang, and Zhizheng Wu. 2024b. Maskgct: Zero-shot text-to-speech with masked generative codec transformer. *arXiv preprint arXiv:2409.00750*.
- Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du, Hanxiao Liu, Yifeng Lu, Percy S Liang, Quoc V Le, Tengyu Ma, and Adams Wei Yu. 2023. Doremi: Optimizing data mixtures speeds up language model pretraining. *Advances in Neural Information Processing Systems*, 36:69798–69818.
- Yuankun Xie, Ruibo Fu, Xiaopeng Wang, Zhiyong Wang, Ya Li, Zhengqi Wen, Haonnan Cheng, and Long Ye. 2025a. Fake speech wild: Detecting deepfake speech on social media platform. *arXiv preprint arXiv:2508.10559*.
- Yuankun Xie, Yi Lu, Ruibo Fu, Zhengqi Wen, Zhiyong Wang, Jianhua Tao, Xin Qi, Xiaopeng Wang, Yukun Liu, Haonan Cheng, and 1 others. 2025b. The codec-fake dataset and countermeasures for the universally detection of deepfake audio. *IEEE Transactions on Audio, Speech and Language Processing*.
- Junichi Yamagishi, Xin Wang, Massimiliano Todisco, Md Sahidullah, Jose Patino, Andreas Nautsch, Xuechen Liu, Kong Aik Lee, Tomi Kinnunen, Nicholas Evans, and 1 others. 2021. Asvspoof 2021: accelerating progress in spoofed and deepfake speech detection. In *ASVspoof 2021 Workshop-Automatic Speaker Verification and Spoofing Countermeasures Challenge*.
- Ryuichi Yamamoto, Eunwoo Song, and Jae-Min Kim. 2020. Parallel wavegan: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6199–6203.
- Geng Yang, Shan Yang, Kai Liu, Peng Fang, Wei Chen, and Lei Xie. 2021. Multi-band melgan: Faster waveform generation for high-quality text-to-speech. In *IEEE Spoken Language Technology Workshop (SLT)*, pages 492–498.
- Zehui Yang, Yifan Chen, Lei Luo, Runyan Yang, Lingxuan Ye, Gaofeng Cheng, Ji Xu, Yaohui Jin, Qingqing Zhang, Pengyuan Zhang, and 1 others. 2022. Open source magicdata-ramc: A rich annotated mandarin conversational (ramc) speech dataset. In *Interspeech*, pages 1736–1740.
- Artem Yaroshchuk, Christoforos Papastergiopoulos, Luca Cuccovillo, Patrick Aichroth, Konstantinos Votis, and Dimitrios Tzovaras. 2023. An open dataset of synthetic speech. In *2023 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6.
- Jiangyan Yi, Ruibo Fu, Jianhua Tao, Shuai Nie, Haoxin Ma, Chenglong Wang, Tao Wang, Zhengkun Tian, Ye Bai, Cunhang Fan, and 1 others. 2022. Add 2022: the first audio deep synthesis detection challenge. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 9216–9220.
- Jiangyan Yi, Chu Yuan Zhang, Jianhua Tao, Chenglong Wang, Xinrui Yan, Yong Ren, Hao Gu, and Junzuo Zhou. 2024. Add 2023: Towards audio deepfake detection and analysis in the wild. *arXiv preprint arXiv:2408.04967*.
- Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J. Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu. 2019. Libritts: A corpus derived from librispeech for text-to-speech. In *Interspeech*, pages 1526–1530.

- Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. 2022. Scaling vision transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12104–12113.
- You Zhang, Fei Jiang, and Zhiyao Duan. 2021. One-class learning towards synthetic voice spoofing detection. *IEEE Signal Processing Letters*, 28:937–941.
- Yan Zhao, Jiangyan Yi, Jianhua Tao, Chenglong Wang, and Yongfeng Dong. 2024. Emofake: An initial dataset for emotion fake audio detection. In *China National Conference on Chinese Computational Linguistics*, pages 419–433.
- Kun Zhou, Berrak Sisman, Rui Liu, and Haizhou Li. 2022. Emotional voice conversion: Theory, databases and esd. *Speech Communication*, 137:1–18.

## A Dataset Details

### A.1 Training Data for Section 3

**Source Diversity.** To investigate the impact of source diversity, we first selected 8 distinct source speech datasets, comprising 4 English (EN) and 4 Chinese (ZH):

- **EN:** VCTK (Veaux et al., 2013), LibriTTS (Zen et al., 2019), MLS (Pratap et al., 2020), and CommonVoice (Ardila et al., 2020).
- **ZH:** Aishell1 (Bu et al., 2017), Aishell3 (Shi et al., 2021), MagicData (Yang et al., 2022), and CommonVoice (Ardila et al., 2020).

From each of these 8 datasets, we selected 10k audio samples to serve as the bonafide (real) data and as the basis for the generation process. Next, to create the synthetic counterparts, we chose 4 zero-shot TTS models: MaskGCT (Wang et al., 2024b), F5TTS (Chen et al., 2024), E2TTS (Eskimez et al., 2024), and CosyVoice2 (Du et al., 2024b). For each real source dataset, we randomly selected speech prompts and text transcripts from it, using each of the 4 TTS models to generate 10,000 fake samples. This process yielded a total of 80k real samples (8 sources  $\times$  10k) and 320k fake samples (8 sources  $\times$  4 generators  $\times$  10k).

To form the training sets, we maintained a fixed real-to-fake ratio ( $\rho = 0.25$ ). We defined two scaling variables: the number of source datasets ( $N_S \in \{1, 2, 4, 8\}$ ) and a data usage percentage ( $V \in \{100\%, 50\%, 25\%, 12.5\%\}$ ). For any given configuration ( $N_S, V$ ), the training set was constructed by aggregating  $10k \times V$  real samples and  $4 \times 10k \times V$  fake samples (from the 4 generators) from each of the  $N_S$  sources. By varying  $N_S$  and  $V$ , we created 10 distinct experimental settings (as shown in Figure 2(a)), resulting in a minimum training set size of 50k samples.

**Generator Diversity.** To investigate the impact of generator diversity, we leverage the data provided in SpeechFake (Huang et al., 2025b). We first select 2 source datasets, VCTK (Veaux et al., 2013) and LibriTTS (Zen et al., 2019), and establish a real data pool of 80k samples from each. Next, we select a total of 16 generators from SpeechFake that were built from these two sources, including models from three categories: 7 TTS (text-to-speech with fixed voice), 3 VC (voice clone), and 6 NV (neural-vocoded speech):

Table 4: Overview of the publicly available datasets composing the training pool for the DOSS experiments (Section 4).

ID	Dataset	#Hours
T01	ASVspoof2019 (Wang et al., 2020)	48
T02	ASVspoof5 (Wang et al., 2024a)	878
T03	ADD2022 (Yi et al., 2022)	54
T04	ADD2023 (Yi et al., 2024)	49
T05	DFADD (Du et al., 2024a)	227
T06	DECRO (Ba et al., 2023)	94
T07	DiffSSD (Bhagtani et al., 2025)	58
T08	SpoofCeleb (Jung et al., 2025)	1925
T09	SpeechFake (Huang et al., 2025b)	2837
T10	EmoFake (Zhao et al., 2024)	28
T11	FoR (Reimao and Tzerpos, 2019)	48
T12	CFAD (Ma et al., 2024)	109

- **TTS:** GlowTTS (Kim et al., 2020), ProDiff-TTS (Huang et al., 2022b), DiffGAN-TTS (Liu et al., 2022), TorToiSe (Betker, 2023), MeloTTS<sup>1</sup>, ChatTTS<sup>2</sup>, CosyVoice (Du et al., 2024b)
- **VC:** CosyVoice (Du et al., 2024b), OpenVoice (Qin et al., 2023) and its variant.
- **NV:** MelGAN (Kumar et al., 2019), ParallelWaveGAN (Yamamoto et al., 2020), HiFiGAN (Kong et al., 2020), FullBandMelGAN (Yang et al., 2021), FastDiff (Huang et al., 2022a), BigVGAN (gil Lee et al., 2023)

From each of these 16 generators, we formed a data pool of 40k fake samples (20k sourced from VCTK and 20k from LibriTTS). To form the training sets, we maintained a fixed real-to-fake ratio ( $\rho = 0.25$ ). We defined two scaling variables: the number of source datasets ( $N_G \in \{1, 2, 4, 8, 16\}$ ) and a data usage percentage ( $V \in \{100\%, 50\%, 25\%, 12.5\%, 6.25\%\}$ ). For any given configuration ( $N_G, V$ ), the training set was constructed by aggregating  $40k \times V$  fake samples from each of the  $N_G$  generators, along with a corresponding  $10k \times V \times N_G$  real samples drawn from the real data pool. By varying  $N_G$  and  $V$ , we created 15 distinct experimental settings (as shown in Figure 2(b)), resulting in a minimum training set size of 50k samples.

### A.2 Training Data for Section 4

Section 4 includes four parts, all conducted on the publicly available datasets detailed in Table 4. We take the training and development splits from these

<sup>1</sup><https://github.com/myshell-ai/MeloTTS>

<sup>2</sup><https://github.com/2noise/ChatTTS>

Table 5: **Overview of the datasets used in the data pool curation** (Section 5). #Dom., #Src., and #Gen. denote the number of domains, sources, and generators, respectively. \* represent unknown generator.

ID	Datasets	Association	Language	#Dom.	#Src.	#Gen.	#Hours
<b>Fake Datasets</b>				<b>332</b>	<b>18</b>	<b>175</b>	<b>9717</b>
F01	ADD2022 (Yi et al., 2022)	R02	zh	1	1	1*	44
F02	ADD2023 (Yi et al., 2024)	R02	zh	7	1	7	49
F03	ASVspoof2019 (Wang et al., 2020)	R11	en	19	1	19	105
F04	ASVspoof2021 (Yamagishi et al., 2021)	R11	en	17	1	17	620
F05	ASVspoof5 (Wang et al., 2024a)	R10	en	32	1	32	1825
F06	CDADD (Li et al., 2024)	R07	en	5	1	5	268
F07	CFAD (Ma et al., 2024)	R02	zh	12	1	12	132
F08	CodecFake (Xie et al., 2025b)	R02,R11	en,zh	14	2	7	885
F09	DECRO (Ba et al., 2023)	R04	en, zh	26	2	16	99
F10	DFADD (Du et al., 2024a)	R11	en	5	1	5	180
F11	DiffSSD (Bhagtani et al., 2025)	R07-08	en	7	2	7	58
F12	EmoFake (Zhao et al., 2024)	R05	en	1	1	1*	17
F13	FoR (Reimao and Tzerpos, 2019)	R08	en	1	1	1*	29
F14	FakeSpeechWild (Xie et al., 2025a)	R06	zh	4	4	1*	41
F15	TIMIT-TTS (Salvi et al., 2023)	R08,R11	en	12	1	12	17
F16	SpoofCeleb (Jung et al., 2025)	R12	en	23	1	23	1816
F17	SpeechFake (Huang et al., 2025b)	R01,R02,R07,R11	en,zh,etc	102	7	33	3035
F18	Self-Generated	R01-03,R07,R09-11	en,zh	48	8	7	590
<b>Real Datasets</b>				<b>18</b>	<b>18</b>	<b>-</b>	<b>2356</b>
R01	Aishell1 (Bu et al., 2017)	F17-18	zh	1	1	-	39
R02	Aishell3 (Shi et al., 2021)	F01-02,F07,F17-18	zh	1	1	-	86
R03	CommonVoice (Ardila et al., 2020)	F17-18	en,zh,etc	3	3	-	1023
R04	DECRO (Ba et al., 2023)	F09	en,zh	2	2	-	39
R05	ESD (Zhou et al., 2022)	F12	en	1	1	-	11
R06	FSW (Xie et al., 2025a)	F14	zh	4	4	-	33
R07	LibriTTS (Zen et al., 2019)	F06,F11,F17-18	en	1	1	-	245
R08	LJSpeech (Ito and Johnson, 2017)	F11,F13,F15	en	1	1	-	24
R09	MagicData	F18	zh	1	1	-	256
R10	MLS (Pratap et al., 2020)	F05	en	1	1	-	350
R11	VCTK (Veaux et al., 2013)	F03-04,F10,F15,F17-18	en	1	1	-	83
R12	Voxceleb (Nagrani et al., 2017)	F16	en	1	1	-	167

datasets to form the data pool. The setup for each part is as follows:

- **Traditional Benchmarks:** We trained separate models on the individual training splits of four datasets: ASVspoof2019 (ASV19), ADD2022 (ADD22), DECRO (DECR), and FoR (FOR).
- **Baseline:** This experiment (and the following two) utilizes the full 12-dataset pool. We cumulatively aggregated the first  $k$  datasets (from T01 to T $k$ ) to form training sets of increasing size:
  - $k = 2$ : T01–T02 (926 hours)
  - $k = 6$ : T01–T06 (1,351 hours)
  - $k = 8$ : T01–T08 (3,335 hours)
  - $k = 12$ : T01–T12 (6,357 hours, the full pool)
- **DOSS-Select:** This strategy uses the full training pool ( $k = 12$ ) but selects a subset of samples from each domain based on the algorithm, resulting in a smaller, pruned dataset.
- **DOSS-Weight:** This strategy also uses the full training pool ( $k = 12$ ), but all samples are kept

and assigned different sampling probabilities.

### A.3 Training Data for Section 5

The final data pool used was curated through a three-step process, with full details in Table 5.

First, we collected 17 publicly available SDD datasets. We unified their audio formats and gathered metadata on their respective sources and generators. This initial analysis revealed that many datasets, while large in volume, were limited in diversity (e.g., fewer than 20 domains).

Second, we reorganized and de-duplicated the pool. We traced the associations between fake datasets and their real source corpora, consolidating the shared real audio into a set of 12 datasets.

Third, we enriched the pool. An analysis of the aggregated generators revealed a gap in recent synthesis methods. We filled this gap by synthesizing new data from 7 recent generators: MegaTTS3 (Jiang et al., 2025), CosyVoice2 (Du et al., 2024b), Chatterbox<sup>3</sup>, MaskGCT (Wang et al.,

<sup>3</sup><https://github.com/resemble-ai/chatterbox>

Table 6: **Overview of the publicly available evaluation datasets used in experiments.** The ‘Notes’ column specifies the exact subset (e.g., LA\_eval) or version (e.g., normed) used for evaluation.

Abbr.	Dataset	Notes
ASV19	ASVspoof2019 (Wang et al., 2020)	LA_eval
DECR	DECRO (Ba et al., 2023)	eval
ITW	InTheWild (Müller et al., 2022)	-
SC	SpoofCeleb (Jung et al., 2025)	eval
FOR	FoR (Reimao and Tzerpos, 2019)	normed_test
EF	EmoFake (Zhao et al., 2024)	eval
ADD22	ADD2022 (Yi et al., 2022)	Track3_test2
ADD23	ADD2023 (Yi et al., 2024)	Track1_testR2
CFAD	CFAD (Ma et al., 2024)	test_unseen
ODSS	ODSS (Yaroshchuk et al., 2023)	-
DV	DeepVoice (Bird and Lotfi, 2023)	segmented
FSW	FakeSpeechWild (Xie et al., 2025a)	eval

2024b), F5-TTS (Chen et al., 2024), E2-TTS (Eskimez et al., 2024), and IndexTTS (Deng et al., 2025). Guided by our scaling law findings, we generated this new data using 8 different sources to ensure high acoustic diversity.

This curation process resulted in our final 12k-hour data pool, comprising 18 real domains and 332 fake domains.

#### A.4 Evaluation Datasets

**Public Benchmarks.** Our experimental results are primarily reported on the public benchmarks detailed in Table 6, which also lists the abbreviations used in our results. These evaluation sets can be broadly categorized by their relationship to the training data:

- **In-domain (partially seen):** ASV19, the most common benchmark for SDD, shares some similar synthesis algorithms with its own training set. Similarly, the DECR evaluation set contains the same generators present in its training set.
- **Out-of-domain (unseen):** The other test sets represent more challenging, unseen scenarios. They consist of either entirely unknown generators (e.g., ITW, ADD22) or generators that were not part of their respective training data (e.g., SC, FOR, EF).

Notably, while most of these datasets consist of English or Chinese audio, ODSS provides multilingual evaluation, as it includes English, German, and Spanish samples.

The experiments in Section 3 and Section 4 use the first 10 test sets for evaluation. For Section 5, we modify this benchmark by excluding in-domain

sets and complementing it with two others (e.g., DV and FSW).

**Commercial APIs.** To complement the generator types found in public benchmarks, we curated a new challenge set using 9 distinct commercial text-to-speech (TTS) APIs, as detailed in Table 7. These platforms range from established cloud providers (e.g., Google, Microsoft) to cutting-edge generative engines (e.g., OpenAI, ElevenLabs, Qwen), ensuring broad coverage of modern acoustic qualities and synthesis architectures. We generated a total of 5,000 synthetic samples per API, split evenly between English and Chinese contexts (2,500 samples each). To maximize diversity, we utilized the full range of available voices on each platform. Furthermore, for APIs that support parameter customization, we enhanced variability by randomly altering attributes such as pitch, timbre, volume, and speaking rate. The source text was drawn from established real datasets as in A.1.

## B Experimental Details

### B.1 Model Architecture

For all experiments, we employ the self-supervised XLS-R (Babu et al., 2022) architecture as our model backbone. This cross-lingual model was pre-trained on approximately 436k hours of unlabeled, publicly available speech data spanning 128 languages. This extensive, multilingual pre-training provides a powerful and generalized foundation for modeling speech representations and has been shown to achieve state-of-the-art performance in SDD (Tak et al., 2022b).

We adapt this backbone for detection by adding a temporal average pooling layer and an MLP classifier head, fine-tuning the entire network on our training data. To balance computational efficiency with performance, we utilize the 300M-parameter version for the extensive empirical studies on scaling laws and data mixing (Sections 3 and 4). For the final large-scale validation and robustness analysis (Section 5), we evaluate both the 300M and 1B-parameter versions to demonstrate the scalability of our approach.

### B.2 Training Configuration

For all training, input audio is first resampled to 16kHz and then processed into 4-second segments; utterances shorter than this length are repeatedly padded, while longer utterances are randomly chunked. We apply two data augmentation techniques:

Table 7: Summary of commercial text-to-speech APIs used for generating synthetic speech data.

API	Provider	URL	Voices
Google Cloud TTS	Google	<a href="https://cloud.google.com/text-to-speech">https://cloud.google.com/text-to-speech</a>	274
Azure Speech Service	Microsoft	<a href="https://learn.microsoft.com/azure/ai-services/speech-service">https://learn.microsoft.com/azure/ai-services/speech-service</a>	116
GPT-4o mini TTS	OpenAI	<a href="https://platform.openai.com/docs/guides/text-to-speech">https://platform.openai.com/docs/guides/text-to-speech</a>	10
ElevenLabs TTS	11Labs	<a href="https://Elevenlabs.io/">https://Elevenlabs.io/</a>	10
Aliyun TTS	Alibaba	<a href="https://ai.aliyun.com/nls/tts">https://ai.aliyun.com/nls/tts</a>	90
Baidu TTS	Baidu	<a href="https://ai.baidu.com/tech/speech/tts">https://ai.baidu.com/tech/speech/tts</a>	5
Xfyun TTS	iFlytek	<a href="https://www.xfyun.cn/services/online_tts">https://www.xfyun.cn/services/online_tts</a>	5
MiniMax TTS	MiniMax	<a href="https://www.minimaxi.com">https://www.minimaxi.com</a>	41
Qwen3 TTS Flash	Qwen3	<a href="https://help.aliyun.com/zh/model-studio/qwen-tts">https://help.aliyun.com/zh/model-studio/qwen-tts</a>	17

1) Rawboost (Tak et al., 2022a): Applied with a probability of 0.5, this method introduces convolutional and additive noise directly to the raw waveform. 2) Codec Augmentation: Applied with a probability of 0.3, this simulates real-world compression artifacts by converting the audio to various formats (e.g., FLAC, MP3, AAC, Opus), improving robustness to different audio inputs. The total effective batch size is set to 128.

We use the AdamW optimizer with a weight decay of  $1e-4$ . The learning rate (LR) schedule was scaled by data volume: for our largest experiments (>1k hours), the LR was held constant at  $1e-6$  for the first 50k steps, then decayed exponentially to  $1e-7$  until 200k steps. For medium-scale experiments (100-1k hours), training was run for 100k steps. For small-scale experiments (<100 hours, e.g., experiments in Sec 3), we trained for a fixed 50k steps to maintain a consistent training budget. We employ a weighted cross-entropy loss for training. The model is trained with a weighted cross-entropy loss, where the class weights are set based on the real-to-fake ratio to balance the dataset.

### B.3 Inference and Evaluation

During inference, test audio is processed using the same 4-second segmentation strategy as in training. We evaluate performance using two primary metrics: Equal Error Rate (EER) and Accuracy (ACC). The EER is computed as a threshold-independent metric from the distribution of the raw real class scores. The ACC is calculated by taking the argmax of the final output predictions, which is equivalent to applying a fixed decision threshold of 0.5 to the softmax-normalized real score.

For our comparison with models from Ge et al., 2025, we used their publicly released checkpoint. Notably, to ensure a fair comparison, we followed its original inference protocol, which processes the full-length audio input directly, rather than apply-

ing our 4-second chunking strategy to avoid performance degradation on their specific models.

## C Extended Experimental Results

### C.1 Domain Distribution under DOSS

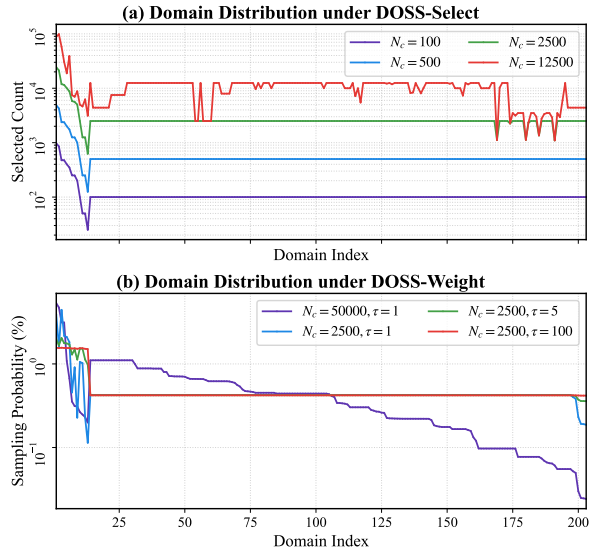


Figure 3: Comparison of domain distributions under DOSS-Select and DOSS-Weight strategies. Each colored line represents a distinct hyperparameter configuration as defined in Table 1.

Figure 3 illustrates the domain sampling distributions resulting from our two DOSS strategies. For DOSS-Select, the fake domain distribution is primarily determined by the saturation cap ( $N_c$ ). By setting  $N_c$  to a small value, we ensure that most fake domains are capped, which results in a near-uniform sampling distribution across them. However, since we use the proportional strategy, a real domain’s sampling weight is derived from the sum of its associated fake domains. This leads to a non-uniform distribution for the real domains, as some are associated with many more fake domains than others. For DOSS-Weight, the final distribution is controlled by both the saturation cap ( $N_c$ ) and

Table 8: **Out-of-domain ACC% ( $\uparrow$ ) comparison with prior works.** The best and second-best results are in **bold** and underline.

System	#Params	#Hours	AVG	ITW	FOR	EF	ADD22	ADD23	ODSS	DV	FSW
<i>Training on Naive Aggregation (Ge et al., 2025)</i>											
MMS-300M	317M	74k	89.60	97.31	87.21	98.39	89.58	94.08	86.59	90.49	73.13
MMS-1B	965M	74k	93.15	96.66	<u>98.05</u>	98.22	97.55	93.11	95.01	96.44	70.21
XLS-R-1B	965M	74k	92.58	98.62	<u>82.65</u>	98.29	98.26	95.04	<u>98.46</u>	92.43	76.92
XLS-R-2B	2.2B	74k	94.62	<u>98.70</u>	96.48	98.45	<u>98.97</u>	96.65	<b>98.92</b>	92.95	75.83
<i>Training with DOSS-Weight (Ours)</i>											
XLS-R-300M	317M	12k	<u>96.88</u>	<b>98.82</b>	97.84	<u>99.01</u>	<b>99.38</b>	<b>97.09</b>	97.01	<b>96.54</b>	<u>89.31</u>
XLS-R-1B	965M	12k	<b>97.40</b>	98.58	<b>99.78</b>	<b>99.51</b>	98.90	<u>96.93</u>	96.34	<u>96.42</u>	<b>92.77</b>

Table 9: **State-of-the-Art (SOTA) Benchmarks for Out-of-Domain Generalization.** This table lists the best-performing models from the literature for each test set, serving as the benchmark references for Table 2.

Test Set	SOTA System	Reference	EER(%)
ITW	XLS-R-2B	Ge et al., 2025	1.23
FOR	W2V-Large	Ge et al., 2025	0.97
EF	XLS-R-2B	Ge et al., 2025	0.20
ADD22	XLS-R-2B	Ge et al., 2025	1.05
ADD23	XLS-R-2B	Ge et al., 2025	4.67
ODSS	XLS-R-2B	Ge et al., 2025	1.13
DV	MMS-300M	Ge et al., 2025	2.27
FSW	XLSR-AASIST	Xie et al., 2025a	11.58

the temperature ( $\tau$ ). If  $N_c$  is set to a large value, the initial base importance of each domain remains highly varied. By setting a proper (smaller)  $N_c$ , we first create a near-uniform base importance for the fake domains, similar to DOSS-Select. Applying the temperature ( $\tau$ ) then primarily serves to flatten the distribution of the real domains, allowing us to balance their weights.

## C.2 Detailed Performance Analysis

**Convergence of Evaluation Metrics.** Table 8 provides the full Accuracy (ACC) results. Unlike the distinct scaling behaviors observed in Section 3 where EER and ACC occasionally diverged, the final results demonstrate that these metrics converge as overall model performance improves. This reflects a general phenomenon where, as the separation between real and fake distributions becomes robust, the performance disparity between an optimal threshold (EER) and a fixed threshold (ACC) naturally diminishes. Consistent with this trend, our final DOSS-trained systems achieve high average accuracy, with the 300M model reaching 96.88% and the 1B model improving further to 97.40%, mirroring the strong generalization observed in the EER analysis.

## Fine-Grained Analysis of Commercial APIs.

Table 10 provides a granular breakdown of performance across 9 commercial providers, separated by language (English and Chinese). Analyzing the results reveals distinct tiers of difficulty among the APIs. Established providers such as Google, Microsoft, Baidu, and iFlytek appear to use synthesis methods that are readily detectable by all evaluated models, with accuracy scores consistently exceeding 98% regardless of the training strategy. In contrast, the newer generation of generative TTS systems—specifically ElevenLabs, MiniMax, and Qwen3—poses a significantly greater challenge, causing performance drops across all models. Qwen3 proves to be the most difficult attacker in the English set, with the best-performing model (XLS-R-1B) achieving only 76.40%, compared to  $>99\%$  on the easier APIs.

Furthermore, we observe a notable dependency on language, where detection performance for the same API can fluctuate drastically between English and Chinese test sets. For the baseline models trained on naive aggregation, this variance is extreme on ElevenLabs, where the XLS-R-2B model detects Chinese fakes with 99.84% accuracy but fails on English fakes, dropping to 63.97%. A similar language gap exists for our DOSS-trained models on the hardest benchmarks; for Qwen3, our XLS-R-1B model detects Chinese samples with 98.24% accuracy but achieves only 76.40% on the English samples. This confirms that synthesis differences exist between languages within these APIs, suggesting that variations in linguistic features or language-dependent vocoder artifacts can impact detector robustness.

## C.3 Latent Space Analysis

To investigate the internal representations learned by our final model (using XLS-R-300M as an example), we conduct an analysis of its latent space. We

Table 10: **Detection ACC% ( $\uparrow$ ) on commercial APIs.** Results are evaluated on 2,500 synthetic samples per API. The table is vertically split by language test set (English vs. Chinese). The best and second-best results within each language section are in **bold** and underline.

Model	AVG	Google	Microsoft	OpenAI	11Labs	Alibaba	Baidu	iFlytek	MiniMax	Qwen3
— English (EN) Test Set —										
<i>Training on Naive Aggregation</i>										
MMS-300M	76.36	98.40	83.92	91.40	45.00	90.36	99.40	98.52	69.00	11.28
MMS-1B	79.40	99.60	90.36	<b>98.72</b>	38.76	97.48	99.96	99.92	68.16	21.64
XLS-R-1B	82.79	99.16	91.00	97.20	65.08	93.00	99.88	98.44	68.96	32.36
XLS-R-2B	83.25	99.12	80.32	<u>98.52</u>	63.97	96.44	99.76	99.64	74.36	37.12
<i>Training with DOSS-Weight (Ours)</i>										
XLS-R-300M	<u>90.62</u>	<u>99.92</u>	<u>99.80</u>	97.32	<b>89.96</b>	<u>98.00</u>	<b>100.00</b>	<b>99.96</b>	<b>89.32</b>	<u>41.28</u>
XLS-R-1B	<b>93.79</b>	<b>99.96</b>	<b>99.84</b>	98.16	<u>82.60</u>	<b>99.56</b>	<b>100.00</b>	<u>99.92</u>	<u>87.68</u>	<b>76.40</b>
— Chinese (ZH) Test Set —										
<i>Training on Naive Aggregation</i>										
MMS-300M	84.85	97.80	89.56	97.44	79.48	74.36	98.12	92.60	49.48	7.20
MMS-1B	90.47	99.84	85.47	99.40	96.80	89.91	94.92	<u>99.64</u>	57.80	22.28
XLS-R-1B	94.63	99.92	96.52	99.60	<u>99.80</u>	89.56	98.92	99.36	73.40	53.84
XLS-R-2B	95.21	<u>99.96</u>	90.52	<b>99.88</b>	<b>99.84</b>	95.68	98.16	99.04	78.60	42.64
<i>Training with DOSS-Weight (Ours)</i>										
XLS-R-300M	<u>98.05</u>	<b>100.00</b>	99.88	98.76	94.28	99.84	<b>100.00</b>	<b>100.00</b>	<u>91.68</u>	<u>70.60</u>
XLS-R-1B	<b>98.23</b>	<b>100.00</b>	<b>100.00</b>	<u>99.68</u>	90.00	<b>100.00</b>	<u>99.96</u>	<b>100.00</b>	<b>96.20</b>	<b>98.24</b>

hypothesize that despite being trained on a simple binary (Real/Fake) objective, the model implicitly learns and encodes separable representations of intrinsic attributes, such as data source and generator.

**Methodology** We test this hypothesis on a subset of our training data composed of 8 distinct sources and 8 distinct generators. Our analysis is two-fold:

- **Quantitative Probing:** We extract embeddings from the temporal pooling layer. To ensure a fair, balanced comparison, we create two 10k-sample testbeds: one real (for source probing) and one fake (for source and generator probing). On each testbed, we train a simple linear classifier (Logistic Regression) as a "probe" to predict the (1) source ID or (2) generator ID from the frozen embeddings. The probe is trained on 80% of the data and evaluated on a held-out 20% test set over 5 random seeds.
- **Qualitative Visualization:** We use t-SNE to visualize the embeddings in 2D, coloring the points by their source or generator labels.

**Results and Analysis** The quantitative probing results, shown in Figure 4, reveal three key findings. **First**, the Initial Model (pre-trained backbone) already exhibits high probing accuracy for both source and generator identification, far exceeding the random chance baseline. This confirms that self-supervised pre-training provides a

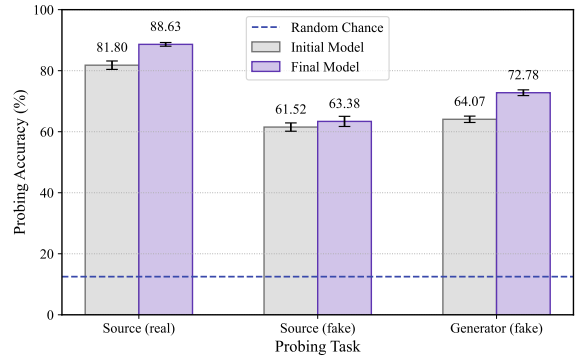


Figure 4: **Linear probing accuracy of model embeddings.** We compare the initial pre-trained model against our final DOSS-trained model.

powerful representation that already contains rich information about audio attributes. **Second**, the training actively sharpens this representation. The final model shows a clear performance increase in identifying both real-world sources and, critically, generator artifacts. This verifies our hypothesis that the model learns and potentially leverages these attributes during binary classification. **Finally**, the results reveal two informative asymmetries. We observe that probing accuracy for real sources is significantly higher than that for fake sources. This suggests that the generation process might partially obscure the original source information. Furthermore, on the fake data, the model is consistently better at identifying the generator than the under-

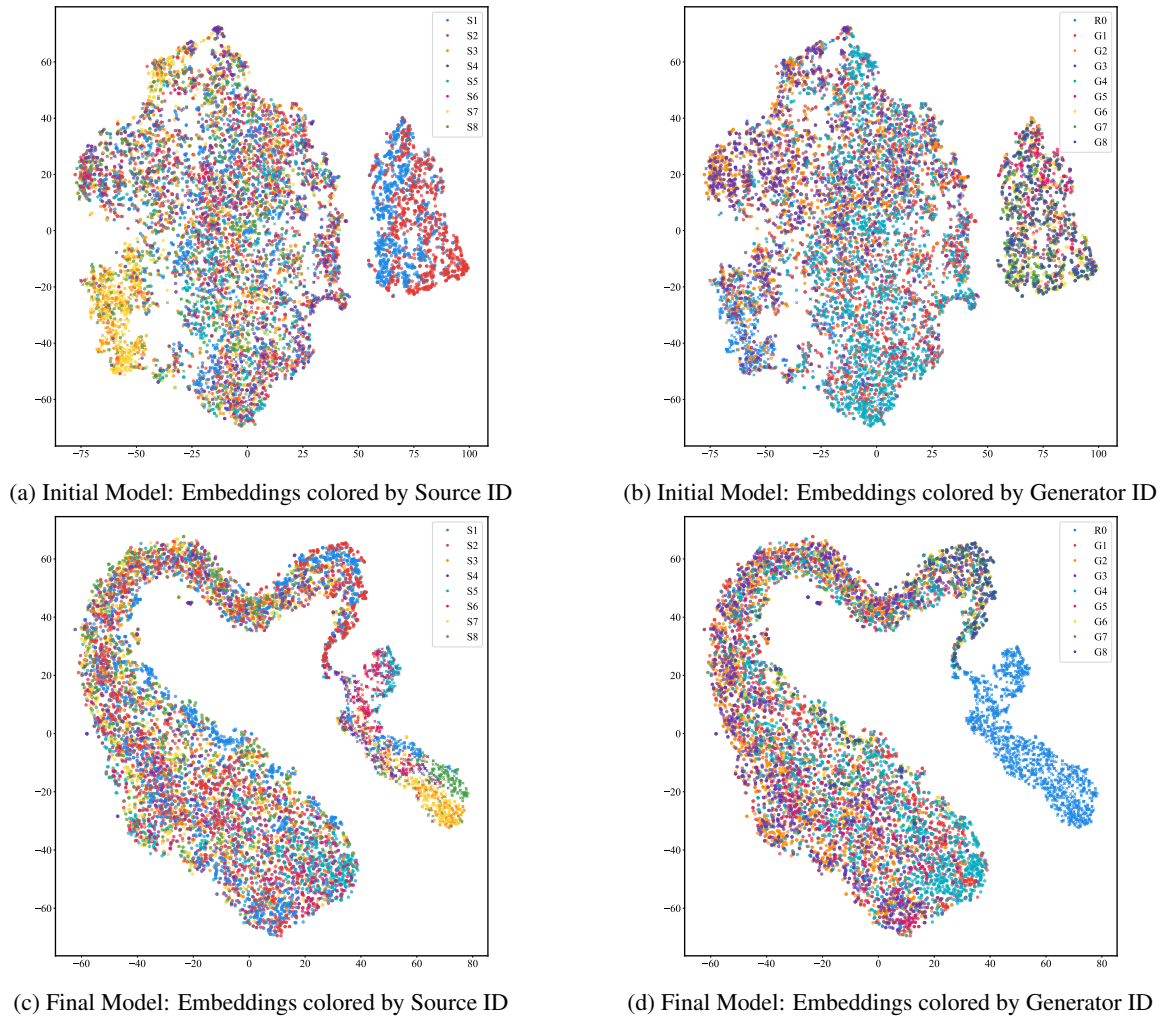


Figure 5: **t-SNE visualization of the initial and final model’s latent space.** (a-b) show the initial pre-trained model’s embeddings. (c-d) show the final DOSS-trained model’s embeddings. Across all subplots, crosses represent the real class and circles represent the fake class. In (b) and (d), R0 represents the real class.

lying source, implying that it could prioritize the generator’s artifact as the most dominant signal.

The t-SNE visualizations in Figure 5 confirm our quantitative probing results. Figure 5a and 5b show the latent space of the initial pre-trained model, where embeddings already form nascent clusters based on source and generator. However, the real and fake samples are highly overlapping and difficult to distinguish. In contrast, Figure 5c and 5d show that the final DOSS-trained model’s latent space is highly structured. While the source and generator clusters are preserved, the real and fake classes are now well-separated. This visual evidence directly supports the high probing accuracies reported in Figure 4. It confirms that the model learns and refines structured representations for these attributes, which it may leverage during the binary classification task.