

Unsupervised Single-Channel Audio Separation with Diffusion Source Priors

Runwu Shi¹, Chang Li², Jiang Wang¹, Rui Zhang³, Nabeela Khan¹,
Benjamin Yen¹, Takeshi Ashizawa¹, Kazuhiro Nakadai¹

¹Department of Systems and Control Engineering, Institute of Science Tokyo

²University of Science and Technology of China

³University of Hong Kong

{shirunwu, nakadai}@ra.sc.e.titech.ac.jp

Abstract

Single-channel audio separation aims to separate individual sources from a single-channel mixture. Most existing methods rely on supervised learning with synthetically generated paired data. However, obtaining high-quality paired data in real-world scenarios is often difficult. This data scarcity can degrade model performance under unseen conditions and limit generalization ability. To this end, in this work, we approach this problem from an unsupervised perspective, framing it as a probabilistic inverse problem. Our method requires only diffusion priors trained on individual sources. Separation is then achieved by iteratively guiding an initial state toward the solution through reconstruction guidance. Importantly, we introduce an advanced inverse problem solver specifically designed for separation, which mitigates gradient conflicts caused by interference between the diffusion prior and reconstruction guidance during inverse denoising. This design ensures high-quality and balanced separation performance across individual sources. Additionally, we find that initializing the denoising process with an augmented mixture instead of pure Gaussian noise provides an informative starting point that significantly improves the final performance. To further enhance audio prior modeling, we design a novel time–frequency attention-based network architecture that demonstrates strong audio modeling capability. Collectively, these improvements lead to significant performance gains, as validated across speech–sound event, sound event, and speech separation tasks.

Project Page — <https://runwushi.github.io/unasdiff/>

Introduction

Single-channel audio separation aims to recover individual sources from a single-channel mixture and has long been an active research topic (Makino 2018). This problem is ill-posed with non-unique solutions and is typically addressed in a supervised manner, which requires large amounts of clean sources and synthetic mixtures for training. This limitation is particularly evident as recent research increasingly focuses on universal, multi-category audio separation tasks, encompassing speech, environmental sound events (e.g., keyboard keystrokes, animal vocalizations), and varied background sounds (Ma et al. 2024; Saijo et al. 2025).

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

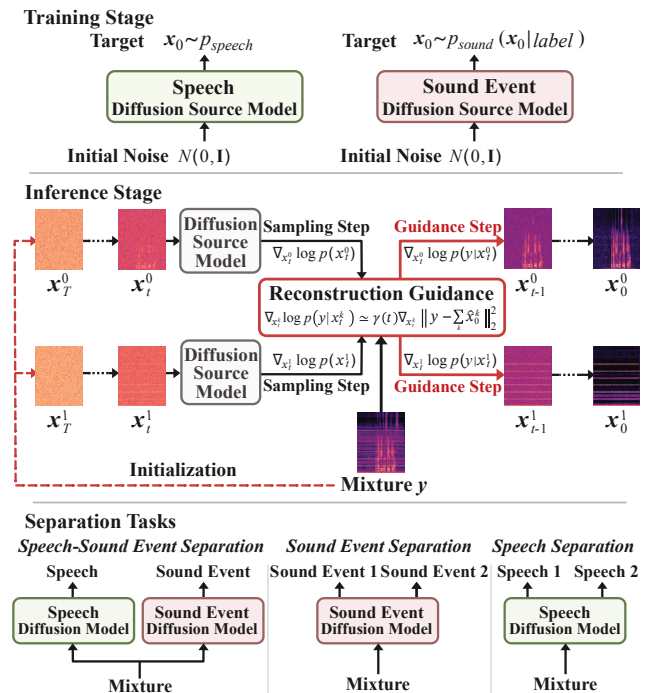


Figure 1: Diffusion source model-based audio separation.

Such tasks demand extensive training mixtures to represent diverse source combinations, which leads to a higher demand for data and training resources. To address this issue, we consider a more challenging setting without paired data, where the model learns from unpaired individual source data and separates audio signals through training-free guidance.

Traditional approaches for unpaired audio separation often model individual sound sources within the source model-based paradigm. Specifically, they characterized the intrinsic properties of each source signal to enable separation during inference. Such methods include Non-negative Matrix Factorization (NMF) (Le Roux, Hershey, and Wenginger 2015), Hidden Markov Model (HMM) (Wang, Woo, and Dlay 2014), and Bayesian models (Benaroya, Bimbot, and Gribonval 2005). Crucially, the priors inherent in such source modeling can be powerfully realized through modern deep generative models, including variational autoencoders

(VAE) (Karamatlı, Cemgil, and Kırılmaz 2019), Glow-based generative models (Zhu et al. 2022), and generative adversarial networks (GAN) (Narayanawamy et al. 2020). While promising, these techniques encounter performance ceilings inherent in the modeling capabilities and generative capacity of their underlying source models, significantly impacting separation performance. More recently, diffusion models (Cao et al. 2024) have shown unprecedented generative power in a wide range of audio tasks such as text-to-speech (Jung et al. 2025). Moreover, their iterative denoising inference supplies strong data priors, making diffusion models well-suited for inverse problems, which aim to recover an unknown signal x from noisy measurement y .

To this end, we introduce our unsupervised diffusion-based audio separation framework, as illustrated in Figure 1. We treat separation as a specific instance of the diffusion inverse problem, which involves two main considerations. On the one hand, when directly applying classical gradient descent based inverse-problem solvers such as Diffusion Posterior Sampling (DPS) (Chung et al. 2023) and Diffusion with Spherical Gaussian constraint (DSG) (Yang et al. 2024) to audio separation, we observe the diffusion prior $\nabla_{x_t} \log p(x_t)$ and the reconstruction guidance $\nabla_{x_t} \mathcal{L}_{\text{recon}}(y, \hat{y})$ suffer from severe gradient conflicts. These conflicts degrade separation quality and hinder accurate source recovery, often resulting in noisy or incomplete outputs. On the other hand, widely adopted representations like mel-spectrograms for audio diffusion models are nonlinear and therefore lack the additive property of time-domain waveforms. Existing waveform-based diffusion models (Kong et al. 2020), even those incorporating self-attention mechanisms (Ku et al. 2025), comparatively lack dedicated audio-specific designs for directly modeling fine-grained intra-frame and intra-subband features. This fundamentally limits their capacity to capture diverse frequency bands and long-range temporal dependencies in audio. Hence, a more powerful model that operates directly in the waveform space is urgently needed.

To address these issues, firstly, we propose a hybrid gradient guidance strength schedule during the reverse separation process, which mitigates gradient conflicts and substantially improves the separation quality. Secondly, we design a triple-path self-attention-based U-Net diffusion backbone as source models to perform guided sampling for the separation task, where the diffusion modeling process and separated audio are directly operated in the waveform space. In addition, we find that initializing the diffusion process with a noise-augmented mixture waveform, rather than pure Gaussian noise, markedly boosts separation fidelity. To sum up, the main contributions of this work are as follows:

- We are the first to address unpaired, general audio separation with diffusion models, framing the task as an inverse problem solved via flexible unconditional or conditional source models, supporting various separation tasks.
- We analyze gradient conflicts and suboptimal prior initialization in the inverse separation process, proposing an effective guidance strength schedule and noise-augmented mixture initialization for separation.

- We present a novel TF domain diffusion model design, employing a triple-path self-attention mechanism to effectively learn priors for diverse audio types like speech and sound events.
- Experiments across speech–sound event, sound event, and speech separation tasks on VCTK and FSD-Kaggle2018 datasets demonstrate the superior separation quality of our proposed method, even achieving performance comparable to fully supervised models.

Related Works

Single Channel Audio Separation

Audio separation can be broadly categorized into supervised and unsupervised methods. Supervised techniques typically learn to separate sources from paired mixture-source data, often employing Permutation Invariant Training (PIT) to handle output ambiguity (Yu et al. 2017). These methods can operate in either the time-domain (Luo and Mesgarani 2019; Li, Yang, and Hu 2022) or frequency-domain (Wang et al. 2023) to separate audio signals.

Many unsupervised approaches perform source model-based separation using deep generative models. These approaches typically model the prior distribution of individual source signals and then infer the posterior distribution of sources given an observed mixture. In (Narayanawamy et al. 2020), source-specific GAN generators served as audio source priors. Separation was achieved by iteratively optimizing latent codes via gradient descent on a waveform reconstruction loss. (Zhu et al. 2022) proposed music separation using source-specific Glow models, similarly employing iterative gradient descent to optimize latent codes and recover sources. (Jayaram and Thickett 2020) introduced a Bayesian approach for image separation using score-based generative models and Glow models as source priors. Source estimates were obtained by sampling from the posterior via noise-annealed Langevin dynamics. (Postolache et al. 2023) used pretrained VQ-VAEs to define source priors for image and audio separation, estimating posterior probability via direct statistical computation of dataset co-occurrences. (Mariani et al. 2023) applied diffusion models to music separation by modeling the joint prior distribution of sources and performing posterior sampling conditioned on the mixture.

Diffusion Inverse Problems and Audio Applications

Inverse problems can be modeled as $y = A(x) + n$, where y is the measurement, A is the degradation model, x is the unknown underlying signal, and n is measurement noise. These problems are fundamental across many fields, enabling tasks such as image restoration (Fei et al. 2023) and audio restoration tasks (Iashchenko et al. 2023). Many diffusion-based inverse problem solvers, such as DPS (Chung et al. 2023), work by iteratively refining samples through gradient descent. DPS is a well-known framework that tackles diverse inverse problems. Building upon this, DSG (Yang et al. 2024) refines DPS. DSG optimizes the guidance by choosing a theoretically optimal step size (Daras et al. 2024), combining it with projected gradient descent to achieve superior sample quality, particularly for

high-dimensional data. For audio inverse problems, (Moliner, Lehtinen, and Välimäki 2023) adopts a DPS-style method and applies it to bandwidth extension, audio inpainting, and declipping tasks. (Xu et al. 2025) solves multi-channel speech separation using Independent Vector Analysis (IVA) for initialization and then applies a DPS-style method to achieve the separation.

Methodology

Score-based Diffusion Models

Score-based diffusion models are used to learn audio source priors, which involves a forward noising process where a clean audio sample \mathbf{x}_t gradually becomes pure noise over time $t \in [0, T]$ (Song et al. 2020):

$$d\mathbf{x} = -\frac{\beta(t)}{2} \mathbf{x} dt + \sqrt{\beta(t)} d\mathbf{w}, \quad (1)$$

where $\beta(t)$ is the noise schedule and \mathbf{w} is the Wiener process. The corresponding reverse process is:

$$d\mathbf{x} = \left[-\frac{\beta(t)}{2} \mathbf{x} - \beta(t) \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \right] dt + \sqrt{\beta(t)} d\bar{\mathbf{w}}, \quad (2)$$

where $\bar{\mathbf{w}}$ is the Wiener process in backward. The score function $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$ is parametrically represented by a neural network s_θ , trained via score-based denoising objectives. This provides differentiable estimation for single-source distributions, forming the prior models for separation tasks.

Audio Separation as a Diffusion Inverse Problem

The audio separation task can be formulated as an inverse problem: a linear mixing operator A combines the K sources and additive noise to produce the observed mixture,

$$\mathbf{y} = A(\mathbf{x}^1, \dots, \mathbf{x}^K) + \mathbf{n} = \sum_{k=1}^K \mathbf{x}^k + \mathbf{n}, \quad (3)$$

and the goal is to recover the individual sources $\{\mathbf{x}^k\}_{k=1}^K$ from the single observation \mathbf{y} .

Throughout the diffusion process, we write \mathbf{x}_t^k for the state of source k at timestep t . Probabilistically, this corresponds to sampling from the posterior $p(\mathbf{x}^{1:K} | \mathbf{y})$. Within the score-based framework, we realize this by guiding the reverse diffusion process (2) with the conditional score $\nabla_{\mathbf{x}_t^k} \log p(\mathbf{x}_t^k | \mathbf{y})$, which is decomposed by Bayes' rule as:

$$\nabla_{\mathbf{x}_t^k} \log p(\mathbf{x}_t^k | \mathbf{y}) = \nabla_{\mathbf{x}_t^k} \log p(\mathbf{x}_t^k) + \nabla_{\mathbf{x}_t^k} \log p(\mathbf{y} | \mathbf{x}_t^k), \quad (4)$$

where $\nabla_{\mathbf{x}_t^k} \log p(\mathbf{x}_t^k)$ is obtained by the trained diffusion source model, and $\nabla_{\mathbf{x}_t^k} \log p(\mathbf{y} | \mathbf{x}_t^k)$ is the gradient of the log-likelihood. Following DPS, the likelihood term $\nabla_{\mathbf{x}_t^k} \log p(\mathbf{y} | \mathbf{x}_t^k)$ is approximated by replacing it with $\nabla_{\mathbf{x}_t^k} \log p(\mathbf{y} | \hat{\mathbf{x}}_0^k)$, where $\hat{\mathbf{x}}_0^k = \mathbb{E}[\mathbf{x}_0^k | \mathbf{x}_t^k]$ is the estimated clean sample computed via Tweedie's formula (Efron 2011):

$$\hat{\mathbf{x}}_0^k \simeq (\mathbf{x}_t^k + (1 - \bar{\alpha}(t)) s_\theta^k(\mathbf{x}_t^k, t)) / \sqrt{\bar{\alpha}(t)}, \quad (5)$$

where $\bar{\alpha}(t) = \prod_{j=1}^t (1 - \beta(j))$ is the cumulative product of the discrete variance schedule, s_θ^k is the k th source model.

Algorithm 1: Audio Separation with Diffusion Priors

Require: number of sound sources K , diffusion time step T , noise levels $\{\sigma_j\}_{j=1}^T$, gradient guidance schedule $\{\gamma_j\}_{j=1}^T$, initialization step t^* , diffusion source models $\{s_\theta^k\}_{k=1}^K$
Input: mixture \mathbf{y} , sound event label \mathbf{c} (optional)

```

1: Initialization from mixture:
2:  $\{\mathbf{x}_T^{\text{aug}}\}_{k=1}^K = \sqrt{\bar{\alpha}_{t^*}} \mathbf{y} + \sqrt{1 - \bar{\alpha}_{t^*}} \boldsymbol{\epsilon}$ ,  $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$ 
3: for  $t = T - 1$  to 0 do
4:   for  $k = 1$  to  $K$  do
5:      $\boldsymbol{\epsilon}_t \sim \mathcal{N}(0, \mathbf{I})$ 
6:      $\hat{\mathbf{x}}_0^k \leftarrow (\mathbf{x}_t + (1 - \bar{\alpha}_t) s_\theta^k(\mathbf{x}_t^k, t, [\mathbf{c}])) / \sqrt{\bar{\alpha}_t}$ 
7:     Prior sampling:
8:      $\mathbf{x}_{t-1}^{k'} \leftarrow \frac{\sqrt{\bar{\alpha}_t(1 - \bar{\alpha}_{t-1})}}{1 - \bar{\alpha}_t} \mathbf{x}_t^k + \frac{\sqrt{\bar{\alpha}_{t-1}\beta_t}}{1 - \bar{\alpha}_t} \hat{\mathbf{x}}_0^k + \sigma_t \boldsymbol{\epsilon}_t$ 
9:   end for
10:   $\hat{\mathbf{y}} = \sum_{k=1}^K \hat{\mathbf{x}}_0^k$ 
11:   $\mathcal{L}_{\text{recons}}(\mathbf{y}, \hat{\mathbf{y}}) = \lambda_{\text{time}} \mathcal{L}_{\text{time}} + \lambda_{\text{group}} \mathcal{L}_{\text{group}} + \lambda_{\text{stft}} \mathcal{L}_{\text{stft}}$ 
12:  for  $k = 1$  to  $K$  do
13:    Back propagation:
14:     $\nabla_{\mathbf{x}_t^k} \log p(\mathbf{y} | \mathbf{x}_t^k) \simeq \gamma(t) \nabla_{\mathbf{x}_t^k} \mathcal{L}_{\text{recons}}(\mathbf{y}, \hat{\mathbf{y}})$ 
15:     $\mathbf{x}_{t-1}^k \leftarrow \mathbf{x}_{t-1}^{k'} - \nabla_{\mathbf{x}_t^k} \log p(\mathbf{y} | \mathbf{x}_t^k)$ 
16:  end for
17: end for
18: return  $\{\mathbf{x}_0^k\}_{k=1}^K$ 

```

With the estimated clean sources $\hat{\mathbf{x}}_0^k$, and under a Gaussian assumption for $p(\mathbf{y} | \sum_{k=1}^K \mathbf{x}_0^k)$, the likelihood gradient for k th source can be approximated through backpropagation:

$$\nabla_{\mathbf{x}_t^k} \log p(\mathbf{y} | \mathbf{x}_t^k) \simeq \gamma(t) \nabla_{\mathbf{x}_t^k} \mathcal{L}_{\text{recons}}(\mathbf{y}, \hat{\mathbf{y}}), \quad (6)$$

where $\hat{\mathbf{y}} = \sum_{k=1}^K \hat{\mathbf{x}}_0^k$ is the reconstructed mixture, and $\gamma(t)$ is a step-dependent coefficient that controls the guidance strength. The design of this schedule is crucial to the final separation performance, as shown in our experiments. The loss function $\mathcal{L}_{\text{recons}}(\mathbf{y}, \hat{\mathbf{y}})$, utilized for computing guidance gradients, is a weighted sum of reconstruction errors in both the time and frequency domains, which is defined as:

$$\mathcal{L}_{\text{recons}} = \lambda_{\text{time}} \mathcal{L}_{\text{time}} + \lambda_{\text{group}} \mathcal{L}_{\text{group}} + \lambda_{\text{stft}} \mathcal{L}_{\text{stft}},$$

where $\mathcal{L}_{\text{time}} = \|\mathbf{y} - \hat{\mathbf{y}}\|_2^2$ measures the reconstruction error in time domain, $\mathcal{L}_{\text{group}} = (\sum_n \|\mathbf{y}_{(n)} - \hat{\mathbf{y}}_{(n)}\|_2^2) / N_g$ averages the same error over N_g non-overlapping segments indexed by n , and $\mathcal{L}_{\text{stft}} = \|\|\text{STFT}(\mathbf{y})\| - \|\text{STFT}(\hat{\mathbf{y}})\|\|_2^2$ compares magnitude spectrogram in the Short-time Fourier transform (STFT) domain. The algorithm is presented in Algorithm 1.

Design of Guidance Strength Schedule $\gamma(t)$

The performance of solving the separation problem is critically influenced by the design of the guidance schedule $\gamma(t)$, which weights the strength of gradient updates. While various guidance strategies exist as presented in Table 1, some are designed for general inverse problems and may not be optimal for the specific constraints of audio separation.

To gain deeper insights, we analyze the generative behavior of sources from the perspective of gradient updates. The

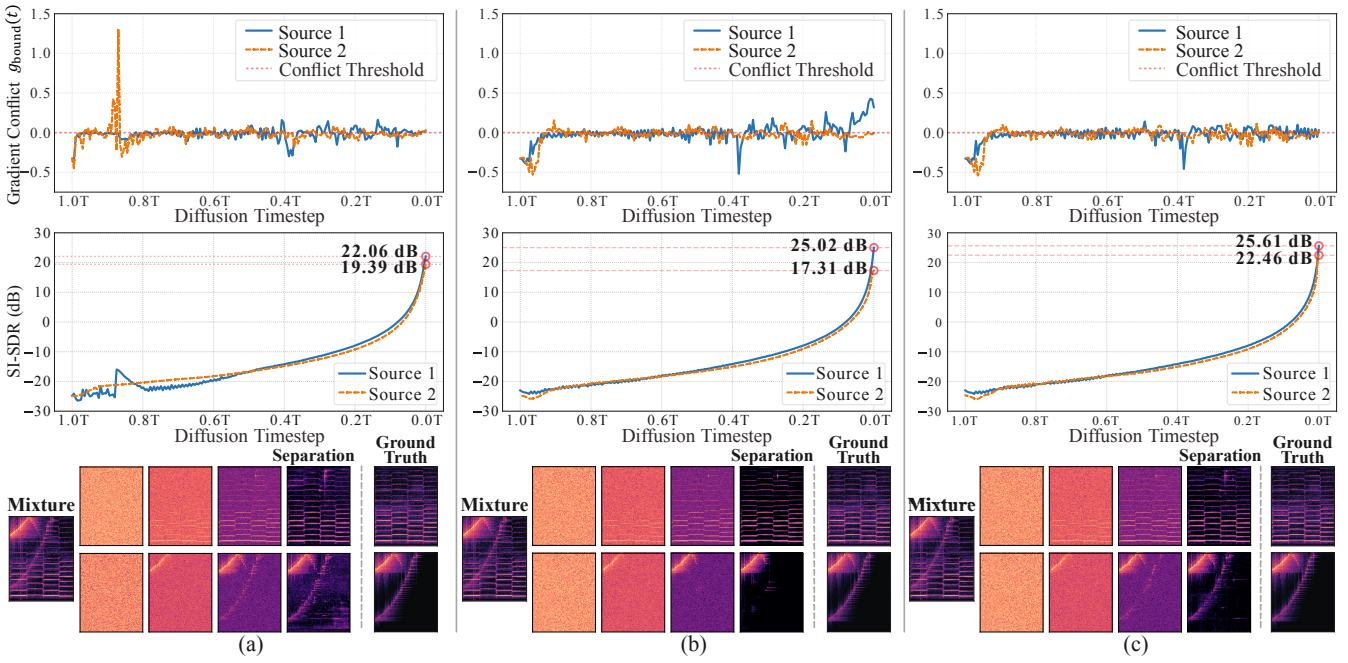


Figure 2: (a) Constant guidance of DPS. (b) Noise-proportional guidance of DSG. (c) Our proposed hybrid guidance schedule. The first row shows the gradient conflict metrics of all sources. The second row shows the SI-SDR of all sources. The third row visualizes the mixtures, intermediate states, and the ground truth of the sources.

overall inverse sampling process can be viewed as a multi-objective optimization problem, as discussed in (Yu et al. 2020). This involves two primary objectives:

1. **Prior Objective:** $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)$, aiming to maintain the learned structure and ensuring the generated sample remains on the manifold of natural data, leading to a prior gradient $g_{\text{prior}}(t)$.
2. **Likelihood Objective:** $\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t)$, ensuring the sample adheres to the physical constraints $\|\mathbf{y} - \sum_{k=1}^K \hat{\mathbf{x}}_0^k\|_2$ imposed by the measurement \mathbf{y} , resulting in a conditional guidance gradient $g_{\text{cond}}(t)$.

Consequently, the total gradient can be expressed as:

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t|\mathbf{y}) \propto g_{\text{prior}}(t) + \gamma(t)g_{\text{cond}}(t). \quad (7)$$

To minimize the reconstruction loss $\mathcal{L}_{\text{recons}}$, we require the update direction to have a positive projection onto the conditional gradient $g_{\text{cond}}(t)$:

$$(\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t|\mathbf{y}))^\top g_{\text{cond}}(t) > 0. \quad (8)$$

Substituting the update formula and simplifying yields the following condition for making progress:

$$\gamma(t) > -\frac{g_{\text{prior}}(t)^\top g_{\text{cond}}(t)}{\|g_{\text{cond}}(t)\|^2}. \quad (9)$$

We define the guidance bound metric $g_{\text{bound}}(t) := -\frac{g_{\text{prior}}(t)^\top g_{\text{cond}}(t)}{\|g_{\text{cond}}(t)\|^2}$. A positive $g_{\text{bound}}(t)$ indicates a gradient conflict between the prior gradient $g_{\text{prior}}(t)$ and the conditional gradient $g_{\text{cond}}(t)$, occurring when their dot product is negative, indicating opposing directions. For effective progress, the guidance strength $\gamma(t)$ should exceed this bound to overcome the conflict.

As illustrated in Figure 2, we evaluated the separation process of a 4-second audio mixture (fiddle and chime from the FSD test dataset) using different guidance strategies: DPS, DSG, and our proposed hybrid schedule. The top row of Figure 2 illustrates our guidance conflict metric $g_{\text{bound}}(t)$, while the second row shows the Scale-Invariant Signal Distortion Ratio (SI-SDR) for each source across diffusion steps. We observed that constant guidance DPS often faces significant gradient conflicts in early, high-noise stages (Figure 2(a)), requiring stronger initial guidance. Conversely, noise-proportional guidance DSG effectively mitigates these early conflicts, but its diminishing step size leads to rising gradient conflict and imbalanced separation in the final, low-noise regime (Figure 2(b)).

These complementary trade-offs, consistently observed across samples, directly inspire our novel hybrid guidance strength schedule. Our approach combines the early-stage adaptability of the noise-proportional schedule with the late-stage stability of the constant one. This is achieved using the SmoothMax function:

$$\text{SmoothMax}_c(a, b) = \frac{1}{c} \log(\exp(ca) + \exp(cb)), \quad (10)$$

which defines our proposed gradient strength schedule as:

$$\gamma(t) = \frac{\text{SmoothMax}_c(\sigma(t), s_{\text{floor}}) \sqrt{N}}{\|\nabla_{\mathbf{x}_t} \mathcal{L}_{\text{recons}}\|_2}, \quad (11)$$

where $\sigma(t)$ is the noise-dependent standard deviation, defined as $\sqrt{\beta(t)(1 - \bar{\alpha}(t - 1)) / (1 - \bar{\alpha}(t))}$. N is the signal length. s_{floor} is a constant that enforces a non-vanishing lower bound, and c controls the sharpness of the transition.

Guidance Strength Strategy	Formulation
DPS (Chung et al. 2023): Constant value.	$\gamma(t) = \text{const}$
DSG (Yang et al. 2024): Proportional to $\sigma(t)$.	$\gamma(t) = \frac{\sigma(t)\sqrt{N}}{\ \nabla_{\mathbf{x}_t} \mathcal{L}_{\text{recons}}\ _2}$
Proposed: Hybrid schedule.	$\frac{\text{SmoothMax}_c(\sigma(t), s_{\text{floor}})\sqrt{N}}{\ \nabla_{\mathbf{x}_t} \mathcal{L}_{\text{recons}}\ _2}$

Table 1: Comparison of guidance strength schedules $\gamma(t)$.

Different schedules without the normalization term and \sqrt{N} are presented in Figure 3. By integrating these complementary strengths, our hybrid approach can enhance separation quality and source balance, as shown in Figure 2(c).




Figure 3: Different guidance strength schedules.

Improve Separation Fidelity with Initialization

In diffusion-based inverse problems, most methods start sampling from pure Gaussian noise (Chung et al. 2023; Yang et al. 2024). However, this is suboptimal for audio separation, where sources exhibit strong structure and are constrained by their mixture composition. Starting from noise forces the model to explore a vast, unstructured space, essentially recovering everything from scratch, which often limits the separation quality.

To address this, we adopt a single-step initialization (Chung, Sim, and Ye 2022): instead of pure noise, each source is seeded by the same noised version of the mixture :

$$\mathbf{x}_T^{\text{aug}} = \sqrt{\alpha_{t^*}} \mathbf{y} + \sqrt{1 - \alpha_{t^*}} \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I}) \quad (12)$$

where t^* is a chosen time step that controls the initial noise variance. This single latent state $\mathbf{x}_T^{\text{aug}}$ thus serves as a common starting point for all source channels, grounding the subsequent separation process in a shared, coherent context derived directly from the measurement \mathbf{y} . The utilized single-step, analytical initialization is different from the DDPM Inversion (Liu et al. 2024), which requires an iterative sampling process and whose performance depends on the learned model.

Model Architecture Design

Current diffusion audio backbones typically lack fine-grained acoustic modeling with dedicated attention mechanisms, particularly in capturing inter-frame and sub-band interactions, for which we design a triple-path fully attention-based U-Net diffusion backbone that directly predicts

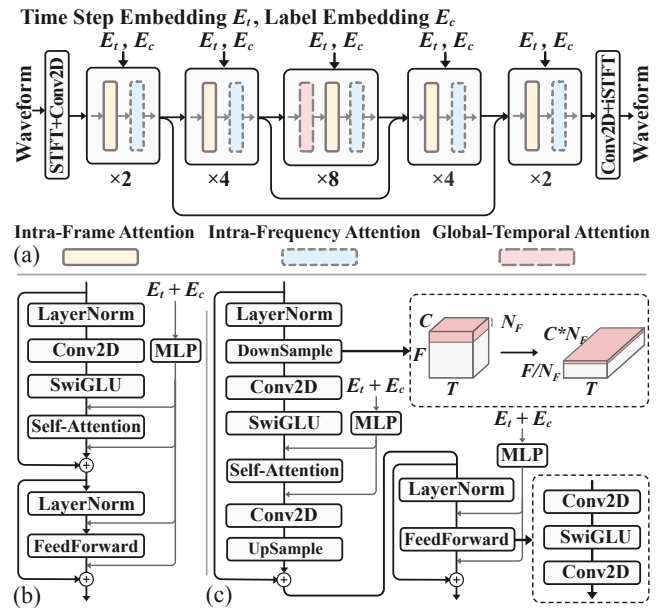


Figure 4: (a) Overall architecture of the proposed model. (b) Structure of the Intra-Frame Attention and Intra-Frequency Attention blocks. (c) Global-Temporal Attention block.

complex-valued noise from STFT spectrograms, as shown in Figure 4(a). Its core is three kinds of attention blocks that capture distinct spectro-temporal correlations. Each block uses learnable Swish Gated Linear Unit (SwiGLU) (Shazeer 2020) for both self-attention preprocessing and feed-forward layers. The entire network is conditioned using a timestep embedding E_t obtained via sinusoidal encoding followed by a multilayer perceptron (MLP) (Tian et al. 2024). For sound event modeling, an additional class vector E_c from a learnable embedding dictionary corresponding to sound class labels is also integrated. Both E_t and E_c are incorporated via the AdaLN-Zero mechanism (Peebles and Xie 2023).

Intra-Frame and Intra-Frequency Attention As shown in Figure 4(b), this attention strategy models inter-frequency correlations within each frame and temporal dynamics of individual frequency bands separately. Intra-Frame Attention scans across frequency bins within each time frame, while Intra-Frequency Attention focuses on temporal evolution within each frequency bin, collectively capturing fine-grained spectral and temporal features. These two blocks share the same structure while differing in scan direction. This approach has been effectively adopted in various speech processing models (Luo, Chen, and Yoshioka 2020; Shi, Yen, and Nakadai 2025; Xu et al. 2024).

Global-Temporal Attention The Global-Temporal Attention block, illustrated in Figure 4(c), is designed to capture temporal relationships across all frequency bins. To reduce computational cost, the input feature map is downsampled along the frequency dimension and processed by a SwiGLU module for channel projection. After attention calculation, a convolution layer projects the feature map back. This block is exclusively used in the latent stage of the U-Net.

Experiments and Results

Datasets and Preprocessing

We independently train diffusion source models for speech and sound event datasets. For speech modeling, we adopt the VCTK dataset, which contains approximately 110 speakers, using 100 speakers for training and 10 for testing. For sound event modeling, we utilize the FSD-Kaggle2018 dataset (Fonseca et al. 2018), which comprises around 11,000 clips across 41 event categories. We use the provided training set with 9,470 clips for training and the remaining clips for testing. All audio samples are resampled to 16 kHz and are cropped or padded to a fixed length of 4 seconds.

Configuration and Training

For the U-Net backbone, we adopt an STFT with a window size of 510 and a hop length of 255. The channel size is set to 72. For Global-Temporal Attention, the number of subbands is set to 4, and the suppressed channel is set to 16. Each attention block employs 4 heads, and the embedding dimension is 128. As illustrated in Figure 4, the U-Net adopts a five-stage block layout with repetitions of 2, 4, 8, 4, and 2, and the total number of parameters is 37M. Each model is trained for 400k steps with a batch size of 12 using the AdamW optimizer with a learning rate of $1e-4$ (Loshchilov and Hutter 2017). We train two diffusion source models on speech and sound event datasets. Denoising training is performed with $T = 200$ steps using a linear noise schedule $\beta(t) \in [10^{-4}, 2 \times 10^{-2}]$. s_{floor} is set to 0.002 and c is set to 10^3 . λ_{time} , λ_{group} , and λ_{sft} are set to 1, 0.05, and 0.1, respectively. The initialization step t^* is set to 150.

Experimental Implementation

Mixture Generation and Tasks For each of the three distinct tasks investigated: Speech-Sound Event Separation, Sound Event Separation, and Speech Separation, 400 unique mixtures were synthesized. These mixtures were created by overlapping single sources with random Root Mean Square levels (-25 to -20 dB) and random temporal offsets.

For Speech-Sound Event Separation, we tested separating speech from sound events in mixtures containing one speech source with either one or two sound event sources, utilizing an unconditional speech model and a label-guided conditional sound event model. Sound Event Separation focused on separating two or three different sound types within mixtures, employing a single conditional sound event model.

For speech separation, a homogeneous source separation problem involves two speech sources. This scenario is particularly challenging for source model-based methods, since it inherently presents permutation ambiguity, as local speech segments may separate, but consistent speaker assignment across time often fails because models trained on clean speech lack speaker clues and primarily learn local acoustic structures (Iashchenko et al. 2023). Our results suggest that a stronger model design can enhance such homogeneous separation without relying on explicit speaker conditioning.

Comparison Methodologies and Metrics Comparisons are made against two supervised separation models: the classical time-domain Conv-TasNet (Luo and Mesgarani 2019),

and the state-of-the-art frequency-domain audio separation model TF-LoCoformer (Saijo et al. 2024, 2025). We independently train the corresponding expert model with the SI-SDR loss function for 200 epochs for each task. Across all comparisons, we utilize identical mixture test datasets.

For unsupervised methods, we evaluate our proposed hybrid guidance schedule $\gamma(t)$ against DPS’s constant and DSG’s noise-proportional schedule, all within the gradient-descent framework and using the same initialization. Additionally, we compare with the Analytical Sampling method (Iashchenko et al. 2023). This approach is used for additive source separation, where the observed mixture \mathbf{y} is an additive combination of sources $\mathbf{y} = \mathbf{x}_1 + \mathbf{x}_2$. It leverages the key insight that the conditional likelihood $p_t(\mathbf{y}|\mathbf{x}_t^1, \mathbf{x}_t^2)$ simplifies to $p_t(\mathbf{y}|\mathbf{x}_t^1 + \mathbf{x}_t^2)$. This enables the analytical computation of the log-likelihood gradient $\nabla_{\mathbf{x}_t^k} \log p_t(\mathbf{y}|\mathbf{x}_t^1, \mathbf{x}_t^2)$.

For Speech Separation specifically, our full attention-based model is compared with large Diffwave (Iashchenko et al. 2023; Kong et al. 2020). This enhanced Diffwave model incorporates the Squeeze-Excitation mechanism, with 33M parameters and 48 layers. We use its publicly released checkpoint trained on 2-second VCTK segments to compare its speech modeling performance with ours.

Performance is evaluated using standard objective metrics: SI-SDR for all separated sources, and PESQ for separated speech in Speech-Sound and Speech separation tasks. For Speech Separation specifically, we also report the separation failure rate. Failure is defined as instances where the average SI-SDR of separated sources falls below 0 dB.

Results

Results for Speech-Sound Event Separation, Sound Event Separation, and Speech Separation are presented in Table 2. As anticipated, supervised methods achieve the highest separation quality, with TF-LoCoformer demonstrating superior performance across all scenarios. Among unsupervised approaches, our proposed hybrid guidance schedule consistently outperforms gradient-based alternatives across all tasks, demonstrating its robustness. Critically, for speech-containing tasks, our method yields significantly higher PESQ scores and lower failure rate, indicating enhanced speech naturalness. Notably, our proposed method yields highly competitive results when compared to supervised baselines, performing comparably to Conv-TasNet, and delivering significant improvements over the unprocessed baseline in all separation tasks. In general, gradient-based methods universally surpass Analytical sampling across tasks, particularly evident in the more challenging multi-source separation tasks. Overall, the results confirm that our schedule delivers superior performance in terms of both separation fidelity and speech naturalness.

Ablation Studies

Influence of Diffusion Model Backbone This ablation study investigates the impact of the underlying diffusion model architecture on speech separation performance, as presented in Table 3. We compare our proposed attention-based model backbone against Large Diffwave using our gradient-based method and Analytical sampling.

Method	1 Speech + 1 Sound		1 Speech + 2 Sound		2 Sound	3 Sound	2 Speech		
	SI-SDR (\uparrow)	PESQ (\uparrow)	SI-SDR (\uparrow)	PESQ (\uparrow)	SI-SDR (\uparrow)	SI-SDR (\uparrow)	SI-SDR (\uparrow)	PESQ (\uparrow)	Failure Rate (\downarrow)
Unprocessed	0.00	1.64	-3.15	1.22	-0.01	-3.13	0.01	1.29	-
<i>Unsupervised: Gradient+Initialization</i>									
DPS’s constant $\gamma(t)$	8.88	1.73	3.02	1.56	4.63	0.12	1.87	1.43	46.5%
DSG’s $\sigma(t)$ -prop. $\gamma(t)$	12.97	1.94	7.96	1.72	9.48	4.75	5.68	1.57	31.3%
Proposed $\gamma(t)$	<u>14.31</u>	<u>2.12</u>	<u>8.58</u>	<u>1.80</u>	<u>10.44</u>	<u>5.15</u>	<u>6.11</u>	<u>1.62</u>	<u>29.8%</u>
<i>Unsupervised: Analytical Sampling</i>									
(Iashchenko et al. 2023)	7.33	1.45	-0.04	1.19	1.01	-2.18	2.17	1.28	42.3%
<i>Supervised</i>									
Conv-TasNet	14.22	1.97	7.05	1.43	11.73	5.02	9.63	1.80	-
TF-LoCoformer	18.22	2.37	12.21	1.95	14.32	9.27	14.52	2.42	-

Table 2: Performance on three audio separation tasks: Speech-Sound Event Separation (1 Speech+1 Sound, 1 Speech+2 Sound), Sound Event Separation (2 Sound, 3 Sound), and Speech Separation (2 Speech). The failure rate is reported only for speech separation and is defined as the average SI-SDR of separated sources falling below 0 dB. Underline indicates unsupervised best.

Method	SI-SDR Succ (\uparrow)	Failure Rate (\downarrow)	SI-SDR All (\uparrow)
Unprocessed	-	-	0.00/0.01
<i>Unsupervised: Gradient+Proposed $\gamma(t)$</i>			
Large Diffwave	7.17/7.36	39%/44%	3.37/3.01
Proposed Model	9.34/9.68	28%/30%	6.04/6.11
<i>Unsupervised: Analytical sampling</i>			
Large Diffwave	5.12/4.69	44%/41%	2.01/1.99
Proposed Model	5.32/5.27	39%/42%	2.45/2.17

Table 3: Performance of Speech Separation using different backbones. All metrics report results for 2-second and 4-second audio segments, presented in a 2-second / 4-second format. SI-SDR Succ denotes successfully separated results.

Results are reported for both 2-second and 4-second speech segments. Our proposed model largely outperforms Large Diffwave in SI-SDR Success and Failure Rate across gradient-based and the Analytical sampling strategy. Our model was not specifically trained on 2-second segments, but its performance on these shorter durations remains superior to Large Diffwave. These findings indicate that our proposed diffusion model architecture is more effective at learning robust audio priors and facilitating superior separation, thereby significantly improving overall separation quality.

Ablation of Initialization and Guidance Parameters

Table 4 presents our ablation study on the key components for the Speech-Sound separation process. We first investigated the initialization time steps t^* . Initializing from pure noise without augmentation yielded the lowest performance, highlighting that an optimal intermediate t^* is crucial, as both excessively high and low noise levels are suboptimal. Next, we examine the impact of reconstruction loss $\mathcal{L}_{\text{recons}}$, showing that combining time-domain and

Parameters	SI-SDR (\uparrow)	PESQ (\uparrow)
<i>Initialization step t^*</i>		
200 (no x_T^{aug})	10.21 (-4.10)	1.90 (-0.22)
175	13.73 (-0.58)	2.10 (-0.02)
125	14.31 (-0.00)	2.11 (-0.01)
100	14.07 (-0.24)	2.10 (-0.02)
<i>Loss Type of $\mathcal{L}_{\text{recons}}$</i>		
$\mathcal{L}_{\text{time}}$	13.65 (-0.66)	1.71 (-0.41)
$\mathcal{L}_{\text{time}} + \mathcal{L}_{\text{group}}$	13.20 (-1.11)	1.65 (-0.47)
$\mathcal{L}_{\text{time}} + \mathcal{L}_{\text{stft}}$	13.03 (-1.28)	2.24 (+0.12)
<i>Final guidance s_{floor}</i>		
0.005	14.20 (-0.11)	2.05 (-0.07)
0.001	13.92 (-0.39)	2.08 (-0.04)

Table 4: Ablation study of Initialization and Guidance.

frequency-domain losses yields the best performance. Finally, we evaluated the influence of the guidance constant s_{floor} in the final guidance stage. We observed that while even a minimal floor value s_{floor} provides a clear gain, setting it to an intermediate value further enhances overall performance.

Conclusion

In this work, we address single-channel audio separation from an unsupervised perspective, framing it as a probabilistic inverse problem solvable with diffusion models. We observed and analyzed the phenomenon of gradient conflicts between prior and conditional updates. To effectively manage these conflicts, we proposed a hybrid guidance strength schedule. Furthermore, we introduce an effective noise-augmented mixture initialization strategy. These innovations significantly enhance separation quality and efficiency, built upon a novel TF domain model for robust audio priors learning. Our method demonstrates impressive effectiveness across three distinct audio separation tasks.

References

- Benaroya, L.; Bimbot, F.; and Gribonval, R. 2005. Audio source separation with a single sensor. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(1): 191–199.
- Cao, H.; Tan, C.; Gao, Z.; Xu, Y.; Chen, G.; Heng, P.-A.; and Li, S. Z. 2024. A survey on generative diffusion models. *IEEE transactions on knowledge and data engineering*, 36(7): 2814–2830.
- Chung, H.; Kim, J.; Mccann, M. T.; Klasky, M. L.; and Ye, J. C. 2023. Diffusion Posterior Sampling for General Noisy Inverse Problems. In *The Eleventh International Conference on Learning Representations, ICLR 2023*. The International Conference on Learning Representations.
- Chung, H.; Sim, B.; and Ye, J. C. 2022. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12413–12422.
- Daras, G.; Chung, H.; Lai, C.-H.; Mitsufuji, Y.; Ye, J. C.; Milanfar, P.; Dimakis, A. G.; and Delbracio, M. 2024. A survey on diffusion models for inverse problems. *arXiv preprint arXiv:2410.00083*.
- Efron, B. 2011. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496): 1602–1614.
- Fei, B.; Lyu, Z.; Pan, L.; Zhang, J.; Yang, W.; Luo, T.; Zhang, B.; and Dai, B. 2023. Generative diffusion prior for unified image restoration and enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9935–9946.
- Fonseca, E.; Plakal, M.; Font, F.; Ellis, D. P.; Favory, X.; Pons, J.; and Serra, X. 2018. General-purpose tagging of freesound audio with audioset labels: Task description, dataset, and baseline. *arXiv preprint arXiv:1807.09902*.
- Iashchenko, A.; Andreev, P.; Shchekotov, I.; Babaev, N.; and Vetrov, D. 2023. UnDiff: Unsupervised Voice Restoration with Unconditional Diffusion Model. In *Proc. Interspeech 2023*, 4294–4298.
- Jayaram, V.; and Thickstun, J. 2020. Source separation with deep generative priors. In *International Conference on Machine Learning*, 4724–4735. PMLR.
- Jung, J.; Ahn, J.; Jung, C.; Nguyen, T. D.; Jang, Y.; and Chung, J. S. 2025. Voicedit: Dual-condition diffusion transformer for environment-aware speech synthesis. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Karamatlı, E.; Cemgil, A. T.; and Kırılmaz, S. 2019. Audio source separation using variational autoencoders and weak class supervision. *IEEE Signal Processing Letters*, 26(9): 1349–1353.
- Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; and Catanzaro, B. 2020. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*.
- Ku, P.-J.; Liu, A. H.; Korostik, R.; Huang, S.-F.; Fu, S.-W.; and Jukić, A. 2025. Generative speech foundation model pretraining for high-quality speech extraction and restoration. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Le Roux, J.; Hershey, J. R.; and Wenginger, F. 2015. Deep NMF for speech separation. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 66–70. IEEE.
- Li, K.; Yang, R.; and Hu, X. 2022. An efficient encoder-decoder architecture with top-down attention for speech separation. *arXiv preprint arXiv:2209.15200*.
- Liu, G.; Sun, H.; Li, J.; Yin, F.; and Yang, Y. 2024. Accelerating diffusion models for inverse problems through shortcut sampling. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 1101–1109.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Luo, Y.; Chen, Z.; and Yoshioka, T. 2020. Dual-path rnn: efficient long sequence modeling for time-domain single-channel speech separation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 46–50. IEEE.
- Luo, Y.; and Mesgarani, N. 2019. Conv-tasnet: Surpassing ideal time–frequency magnitude masking for speech separation. *IEEE/ACM transactions on audio, speech, and language processing*, 27(8): 1256–1266.
- Ma, H.; Peng, Z.; Li, X.; Shao, M.; Wu, X.; and Liu, J. 2024. CLAPSep: Leveraging Contrastive Pre-trained Model for Multi-Modal Query-Conditioned Target Sound Extraction. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Makino, S. 2018. *Audio source separation*, volume 433. Springer.
- Mariani, G.; Tallini, I.; Postolache, E.; Mancusi, M.; Cosmo, L.; and Rodolà, E. 2023. Multi-source diffusion models for simultaneous music generation and separation. *arXiv preprint arXiv:2302.02257*.
- Moliner, E.; Lehtinen, J.; and Välimäki, V. 2023. Solving audio inverse problems with a diffusion model. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Narayananaswamy, V.; Thiagarajan, J. J.; Anirudh, R.; and Spanias, A. 2020. Unsupervised Audio Source Separation Using Generative Priors. In *Proc. Interspeech 2020*, 2657–2661.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- Postolache, E.; Mariani, G.; Mancusi, M.; Santilli, A.; Cosmo, L.; and Rodolà, E. 2023. Latent autoregressive source separation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 9444–9452.
- Saijo, K.; Ebberts, J.; Germain, F. G.; Wichern, G.; and Le Roux, J. 2025. Task-Aware Unified Source Separation. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.

Saijo, K.; Wichern, G.; Germain, F. G.; Pan, Z.; and Le Roux, J. 2024. TF-LoCoformer: Transformer with local modeling by convolution for speech separation and enhancement. In *2024 18th International Workshop on Acoustic Signal Enhancement (IWAENC)*, 205–209. IEEE.

Shazeer, N. 2020. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*.

Shi, R.; Yen, B.; and Nakadai, K. 2025. Distance based single-channel target speech extraction. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.

Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.

Tian, Y.; Tu, Z.; Chen, H.; Hu, J.; Xu, C.; and Wang, Y. 2024. U-dits: Downsample tokens in u-shaped diffusion transformers. *arXiv preprint arXiv:2405.02730*.

Wang, Q.; Woo, W. L.; and Dlay, S. S. 2014. Informed single-channel speech separation using HMM–GMM user-generated exemplar source. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22(12): 2087–2100.

Wang, Z.-Q.; Cornell, S.; Choi, S.; Lee, Y.; Kim, B.-Y.; and Watanabe, S. 2023. TF-GridNet: Integrating full-and sub-band modeling for speech separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31: 3221–3236.

Xu, M.; Li, K.; Chen, G.; and Hu, X. 2024. TIGER: Time-frequency Interleaved Gain Extraction and Reconstruction for Efficient Speech Separation. *arXiv preprint arXiv:2410.01469*.

Xu, Z.; Fan, X.; Wang, Z.-Q.; Jiang, X.; and Choudhury, R. R. 2025. Unsupervised Blind Speech Separation with a Diffusion Prior. *arXiv preprint arXiv:2505.05657*.

Yang, L.; Ding, S.; Cai, Y.; Yu, J.; Wang, J.; and Shi, Y. 2024. Guidance with spherical gaussian constraint for conditional diffusion. In *Proceedings of the 41st International Conference on Machine Learning*, 56071–56095.

Yu, D.; Kolbæk, M.; Tan, Z.-H.; and Jensen, J. 2017. Permutation invariant training of deep models for speaker-independent multi-talker speech separation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 241–245. IEEE.

Yu, T.; Kumar, S.; Gupta, A.; Levine, S.; Hausman, K.; and Finn, C. 2020. Gradient surgery for multi-task learning. *Advances in neural information processing systems*, 33: 5824–5836.

Zhu, G.; Darefsky, J.; Jiang, F.; Selitskiy, A.; and Duan, Z. 2022. Music source separation with generative flow. *IEEE Signal Processing Letters*, 29: 2288–2292.