

# Melodia: Training-Free Music Editing Guided by Attention Probing in Diffusion Models

Yi Yang\*, Haowen Li\*, Tianxiang Li, Boyu Cao, Xiaohan Zhang, Liqun Chen, Qi Liu†

School of Future Technology, South China University of Technology  
fity@mail.scut.edu.cn, fthaowen@mail.scut.edu.cn, drliuqi@scut.edu.cn

## Abstract

Text-to-music generation technology is progressing rapidly, creating new opportunities for musical composition and editing. However, existing music editing methods often fail to preserve the source music’s temporal structure, including melody and rhythm, when altering particular attributes like instrument, genre, and mood. To address this challenge, this paper conducts an in-depth probing analysis on attention maps within AudioLDM 2, a diffusion-based model commonly used as the backbone for existing music editing methods. We reveal a key finding: cross-attention maps encompass details regarding distinct musical characteristics, and interventions on these maps frequently result in ineffective modifications. In contrast, self-attention maps are essential for preserving the temporal structure of the source music during its conversion into the target music. Building upon this understanding, we present Melodia, a training-free technique that selectively manipulates self-attention maps in particular layers during the denoising process and leverages an attention repository to store source music information, achieving accurate modification of musical characteristics while preserving the original structure without requiring textual descriptions of the source music. Additionally, we propose two novel metrics to better evaluate music editing methods. Both objective and subjective experiments demonstrate that our approach achieves superior results in terms of textual adherence and structural integrity across various datasets. This research enhances comprehension of internal mechanisms within music generation models and provides improved control for music creation.

## Introduction

Text-to-music generation technology continues to evolve rapidly, opening new ways to edit music. Text-guided music editing enables modification of musical attributes through natural language instructions. As mentioned by MusicMagus (Zhang et al. 2024b), this field encompasses two categories: inter-stem editing and intra-stem editing. Inter-stem editing involves adding or removing separate instrumental tracks (e.g., “add drum” or “remove guitar”), while intra-stem editing focuses on modifying characteristics within a

\*These authors contributed equally.

†Corresponding author.

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

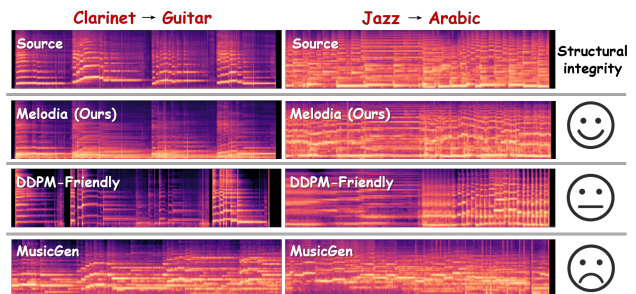


Figure 1: Spectrogram comparison of music editing results between different methods. The comparison reveals that existing methods struggle to preserve the original temporal patterns and rhythmic structures, while Melodia maintains better structural consistency with the source music.

single track, such as changing the timbre, style, or mood while maintaining its melody and structure.

However, existing text-guided music editing methods struggle with these issues: 1) Existing methods train specialized models from scratch (Agostinelli et al. 2023; Copet et al. 2023) or fine-tune pre-trained models (Ruiz et al. 2023; Zhang et al. 2024a), both requiring significant computational costs and training data. 2) Most existing methods require textual descriptions of the source music to guide the editing process. For example, MusicMagus (Zhang et al. 2024b) requires a descriptive word about the source music, which can be difficult for users to provide accurately. 3) Existing methods often struggle to maintain the temporal structure (including the melody and rhythm) of the source music during editing. Fig. 1 demonstrates this issue through spectrogram comparisons, where methods like DDPM-Friendly (Manor and Michaeli 2024) and MusicGen (Copet et al. 2023) show inadequate structural preservation. 4) While attention retention techniques have been applied in image editing (Hertz et al. 2022; Cao et al. 2023), how attention mechanisms function in music editing remains unexplored.

To address these issues, we conduct an in-depth probing analysis on different attention maps within AudioLDM 2 (Liu et al. 2024c), a diffusion-based model commonly used as the backbone for existing music editing methods. We reveal a key insight: cross-attention maps encompass details

regarding distinct musical characteristics, and interventions on these maps result in ineffective modifications. In contrast, self-attention maps are essential for preserving the temporal structure of the source music during its conversion into the target music. This finding suggests that manipulating self-attention maps may yield better editing results. Building on this insight, we introduce Melodia, a training-free and source-prompt-free approach that selectively manipulates self-attention maps in specific layers during the diffusion model’s denoising process. The source-prompt-free nature lowers the usage barrier for non-expert users. By building an attention repository to store self-attention information from the source music and applying it during editing, our method achieves a good balance between textual adherence and structural integrity without requiring any textual description of the source music. Fig. 1 demonstrates that our proposed method Melodia achieves superior structural preservation compared to existing approaches. Furthermore, to address the challenge of evaluating the balance between textual adherence and structural integrity, we propose two novel metrics, Adherence-Structure Balance Score (ASB) and Adherence-Musicality Balance Score (AMB), that comprehensively measure this balance in music editing tasks. Additionally, we construct MelodiaEdit, a music editing benchmark covering diverse editing scenarios.

In summary, our main contributions are as follows:

- We conduct a thorough analysis of attention maps within the diffusion-based model, revealing the unique functions of cross-attention and self-attention maps in music editing, with self-attention playing a critical role in preserving temporal structure.
- We propose two novel evaluation metrics, ASB and AMB, to assess the balance between textual adherence and structural integrity, and construct MelodiaEdit benchmark covering diverse editing scenarios.
- We propose Melodia, a novel training-free music editing technique that selectively manipulates self-attention maps during denoising, achieving an optimal balance between textual adherence and structural integrity without requiring textual descriptions of the source music.
- Through comprehensive subjective and objective evaluations across three datasets, our approach consistently outperforms existing methods for intra-stem music editing.

## Related Work

### Text-to-Music Generation

Text-to-music generation follows two main paradigms: autoregressive (AR) models and diffusion-based models. AR models like MusicLM (Agostinelli et al. 2023) introduces a music-text embedding space, while MusicGen (Copet et al. 2023) improves controllability through text and melody guidance. Diffusion-based models transform noise into musical structures through iterative denoising. AudioLDM (Liu et al. 2023) applies latent diffusion to text-conditioned generation, while Mousai (Schneider et al. 2023) and MusicLDM (Chen et al. 2024) further advance quality and length capabilities. Recently, AudioLDM 2 (Liu et al. 2024c) improves musical structure representation, MelodyFlow (Lan

et al. 2024) advances the field with flow matching techniques, and Stable Audio Open (Evans et al. 2025) scales diffusion transformers to generate extended stereo music. As these models have made significant progress in music generation, researchers have expanded their focus to music editing with its unique challenges.

### Text-to-Music Editing

**Specialized Training Methods.** MusicLM (Agostinelli et al. 2023) utilizes MuLan (Huang et al. 2022) embedding space for style editing, while MusicGen (Copet et al. 2023) facilitates editing by conditioning generation on an original audio’s chromagram with text prompts for desired changes. These two models offer limited editing capabilities as a secondary function. Models like AUDIT (Wang et al. 2023) and InstructME (Han et al. 2023) train diffusion models specifically for inter-stem editing. While effective, these methods require extensive training on text-audio pairs.

**Fine-Tuning Approaches** adapt pre-trained models through targeted optimization. Plitsis et al. (2024) demonstrate techniques adapted from image editing, such as DreamBooth (Ruiz et al. 2023) for audio personalization. Instruct-MusicGen (Zhang et al. 2024a) exemplifies this category by fine-tuning the pre-trained MusicGen model with an instruction-following strategy. While requiring less training data, these methods remain computationally intensive.

**Zero-Shot Editing Methods** manipulate music without additional training. MusicMagus (Zhang et al. 2024b) uses cross-attention constraint and word swapping but requires specific keywords describing the original music to find editing directions. MEDIC (Liu et al. 2024b) unifies mutual self-attention control and cross-attention control. Yet, these methods focus on manipulating attention mechanisms without providing interpretability insights into how these mechanisms function in music diffusion models. Additionally, existing methods including DDPM Friendly (Manor and Michaeli 2024) and SDEdit (Meng et al. 2021) lack explicit structure guidance from the original music, making it difficult to preserve temporal structure during editing.

## Method

### Preliminary

**Latent Diffusion Model (LDM)** (Rombach et al. 2022) is a form of diffusion model trained within low-dimensional latent space. Given data  $x$ , LDM employs a variational autoencoder (VAE) (Kingma, Welling et al. 2013) encoder  $\mathcal{E}$  to compress  $x$  into the latent  $z$  and a corresponding decoder  $\mathcal{D}$  to reconstruct the data. In the diffusion process, the initial latent  $z_0$  is converted to a sample  $z_T$  by adding random noise  $\epsilon_t \sim \mathcal{N}(0, 1)$  with  $T$  iterations. The denoising process recovers  $z_0$  from  $z_T$  utilizing a trained denoiser  $\epsilon_\theta$ .

**Cross-Attention (CA) Mechanism**, which establishes connections between inputs of different modalities, is crucial for implementing a conditional denoiser  $\epsilon_\theta(z_t, t, y)$  con-

Class	Layer 1	Layer 4	Layer 6	Layer 10	Layer 13	Layer 16	Avg.
piano	0.90	1.00	0.50	0.97	0.97	0.95	0.88
accordion	0.80	0.77	1.00	0.90	0.70	0.72	0.82
jazz	0.97	0.72	0.47	0.72	1.00	0.85	0.79
country	0.90	0.70	0.32	0.82	1.00	0.85	0.77
sad	0.95	0.90	0.57	0.95	0.77	0.82	0.83
happy	0.70	0.60	0.62	0.90	0.77	0.60	0.70

Table 1: Probing accuracy of CA map in different layers.

Class	Layer 1	Layer 4	Layer 6	Layer 10	Layer 13	Layer 16	Avg.
piano	0.15	0.37	0.12	0.25	0.70	0.02	0.27
accordion	0.22	0.05	0.17	0.45	0.65	0.40	0.32
jazz	0.32	0.17	0.05	0.37	0.22	0.07	0.20
country	0.10	0.02	0.05	0.05	0.67	0.10	0.17
sad	0.02	0.15	0.27	0.27	0.67	0.60	0.33
happy	0.05	0.65	0.05	0.27	0.30	0.17	0.38

Table 2: Probing accuracy of SA map in different layers.

ditioned by  $y$ . Output  $\phi^c$  of a CA layer is defined as:

$$\phi^c = \text{Cross-Attention}(Q^c, K^c, V^c) = M^c \cdot V^c \quad (1)$$

$$M^c = \text{Softmax}\left(\frac{Q^c K^{c\top}}{\sqrt{d^c}}\right) \quad (2)$$

$$Q^c = W_{Q^c} \cdot \varphi(z_t), \quad K^c = W_{K^c} \cdot \tau(y), \quad V^c = W_{V^c} \cdot \tau(y) \quad (3)$$

where  $\varphi(z_t) \in \mathbb{R}^{N \times d^e}$  denotes a flattened intermediate representation of hidden spatial features in the denoiser  $\epsilon_\theta$ , and  $\tau(\cdot)$  is a modality-specific encoder introduced to embed inputs  $y$  of various modalities into unified embeddings  $\tau(y) \in \mathbb{R}^{M \times d_\tau}$ .  $W_{Q^c} \in \mathbb{R}^{d^c \times d^e}$ ,  $W_{K^c}$ ,  $W_{V^c} \in \mathbb{R}^{d^c \times d_\tau}$  are learnable projection matrices. With  $y$ , conditional LDM is enabled to generate desired data with a conditional denoiser  $\epsilon_\theta(z_t, t, y)$ . Moreover, adopted from Ho et al. (Ho and Salimans 2022), conditional LDM uses Classifier-Free Guidance strength (CFG) to control the strength of conditions.

**Self-Attention (SA) Mechanism** focuses on information processing within the latent itself by establishing connections among spatial patches. Output  $\phi^s$  of SA is defined as:

$$\phi^s = \text{Self-Attention}(Q^s, K^s, V^s) = M^s \cdot V^s \quad (4)$$

$$M^s = \text{Softmax}\left(\frac{Q^s K^{s\top}}{\sqrt{d^s}}\right) \quad (5)$$

$$Q^s = W_{Q^s} \cdot \varphi(z_t), \quad K^s = W_{K^s} \cdot \varphi(z_t), \quad V^s = W_{V^s} \cdot \varphi(z_t) \quad (6)$$

Unlike CA, query  $Q^s$ , key  $K^s$  and value  $V^s$  of SA are all derived from the hidden spatial feature  $\varphi(z_t)$ . Therefore, it can capture dependencies within patches of the input latent.

**AudioLDM 2** (Liu et al. 2024c), an open-source conditional LDM-based audio generation model, is the foundation model of our work. It contains 16 layers with cross-attention and self-attention mechanisms in its UNet architecture.

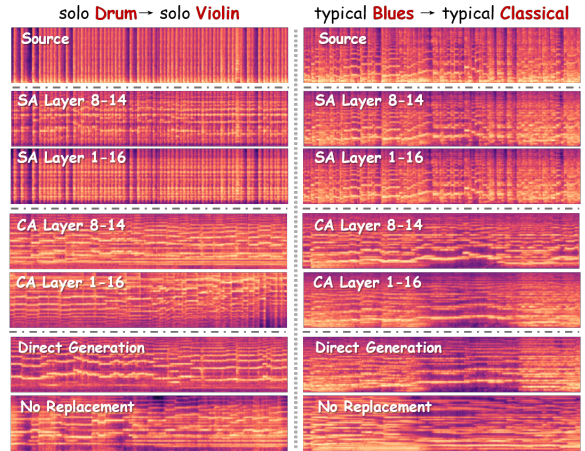


Figure 2: Results of cross-attention map and self-attention map replacement in different layers of the AudioLDM 2.

### Probing Analysis on Attention Maps

In this section, we analyze how cross-attention and self-attention maps in AudioLDM2 (Liu et al. 2024c) contribute to music editing effectiveness. Understanding these mechanisms is critical for developing methods that can transform musical attributes while preserving the original structure.

**Probing Methodology.** In AudioLDM 2 (Liu et al. 2024c), attention mechanisms are the key to its successful music generation. However, **it remains unclear whether these attention maps are merely weight matrices or contain rich semantic representations of music.** To address this, inspired by probing analysis methods in NLP (Clark et al. 2019; Liu et al. 2019), we build datasets and train classification networks to explore attention map properties. The principle is straightforward: If a classifier can accurately categorize attention maps into different classes, these maps must encode meaningful semantic information beyond mere weighting. We construct three prompt datasets targeting fundamental musical dimensions: instruments (16 types, *a solo [instrument] music*), styles (11 types, *a typical [style] music*), and moods (8 types, *a [mood] music*). For each prompt, we extract cross-attention maps corresponding to specific keywords ([instrument], [style], and [mood]) and self-attention maps from different layers during generation. Finally, we train a simple MLP classifier to determine if these maps encode specific musical attributes. High classification accuracy indicates that attention maps encode substantial information about musical attributes rather than functioning merely as weighting mechanisms.

**Probing Results on Cross-Attention Maps.** Tab. 1 shows our classification results on cross-attention maps across different layers. Our classifier achieves high accuracy across all three classification tasks, with average accuracy rates exceeding 70%, indicating that cross-attention maps contain rich semantic information about musical characteristics rather than simple weighting. This explains why direct manipulation of cross-attention maps causes unsuccessful music editing outcomes. Fig. 2 shows that when re-

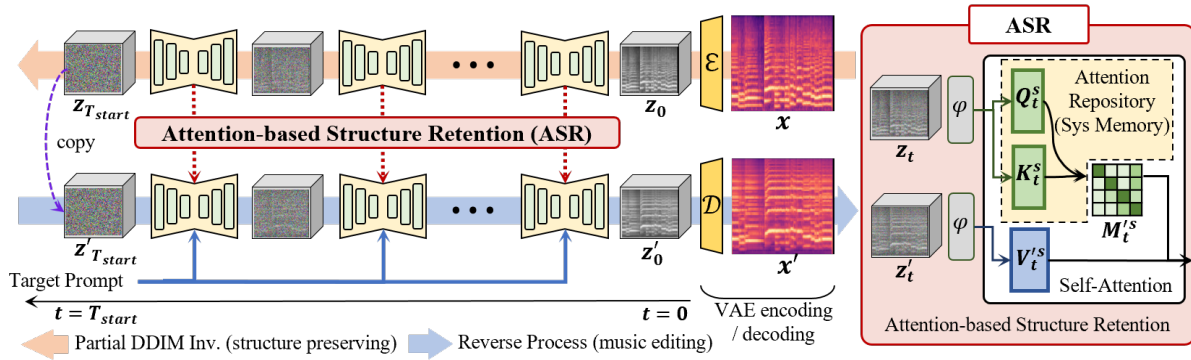


Figure 3: (Left) Overview of Melodia. (Right) Detailed explanation of Attention-based Structure Retention (ASR).

placing cross-attention maps (layers 8-14 or 1-16), the temporal structure of the edited music significantly differs from the source music, yet aligns closely with the music directly generated from the target prompt.

**Probing Results on Self-Attention Maps.** Tab. 2 shows that our classifier struggles with self-attention maps, achieving low average accuracy ( $<40\%$ ) significantly below cross-attention results. This suggests self-attention maps do not encode categorical information about musical attributes. Instead, our further analysis shows they capture temporal features critical for music coherence, such as melody and rhythm. This finding parallels observations in image diffusion models, where self-attention maps preserve the spatial structure of images (Liu et al. 2024a). To validate this, we conducted self-attention map replacement experiments. Fig. 2 shows that replacing self-attention maps in layers 8-14 successfully changes attributes while preserving original melody and rhythm—the edited violin music maintains the source drum’s beat pattern. Without replacement, the music loses source temporal structure and resembles direct generation from the target prompt. However, replacing all layers (1-16) partially preserves original timbre instead of complete transformation to violin. These experimental results support our idea that self-attention maps are essential for preserving temporal structure during music editing, with selective layer replacement showing promising results.

### Overview of Our Approach

Building on our exploration of attention layers, the fundamental innovation of our approach lies in preserving the temporal structure of the original music clip  $x_0$  within self-attention map manipulation. Based on this, we propose a straightforward yet effective approach named Melodia.

As shown in Fig. 3, in our approach, we first obtain the latent  $z_0$  of the original music clip via a VAE (Kingma, Welling et al. 2013) encoder  $\mathcal{E}$ , and then collect the SA features of  $z_0$  during the *Partial DDIM Inversion* (Song, Meng, and Ermon 2020; Manor and Michaeli 2024) process (hereafter DDIM inversion), a version that applies the inversion only up to  $T_{start} \in [0, T], T = 1000$  (Manor and Michaeli 2024). Specifically, for predefined time steps  $t \in \{0, \dots, T_{start}\}$ , we invert the original latent  $z_0$  into an editable intermediate latent  $z_{T_{start}}$  and store self-attention




Figure 4: Intuitive Illustration of DDIM Inversion and Reverse Process with Attention Repository based Structure Guidance. The orange and blue paths respectively refer to DDIM Inversion path and reverse path.

queries  $Q_t^s$  and keys  $K_t^s$  at each time step to build an attention repository in this process. The repository storing self-attention features of the original music clip is used to provide guidance for subsequent music editing. In the music editing process, we transform the features stored in the attention repository into original music clip structural guidance through the proposed Attention-based Structure Retention (ASR) method, which is based on manipulation of the self-attention mechanisms of pre-trained diffusion models, without extra fine-tuning or new blocks. Similarly to the guidance of the given target prompt, this guidance also modulates the denoising process. Finally, the output  $z'_0$  is decoded into a readable mel spectrogram of the desired music clip via the corresponding VAE decoder  $\mathcal{D}$ .

### DDIM Inversion with Attention Repository

Adopted from Content-Style Modeling in Multi-Domain Analysis (Sørensen, Kanatsoulis, and Sidiropoulos 2021; Shrestha and Fu 2024), we assume that any music sample  $x^{(n)}$  from domain  $n$  can be represented as a bijective func-

tion  $\mathbf{g}$  of content  $\mathbf{c}$  and style  $\mathbf{s}^{(n)}$  (i.e. timbre) components:

$$\mathbf{c} \sim \mathbb{P}_{\mathbf{c}}, \mathbf{s}^{(n)} \sim \mathbb{P}_{\mathbf{s}^{(n)}}, x^{(n)} = \mathbf{g}(\mathbf{c}, \mathbf{s}^{(n)}) \quad (7)$$

where  $\mathbb{P}_{\mathbf{s}^{(n)}}$  and  $\mathbb{P}_{\mathbf{c}}$  are distributions of the style components in  $n$ th domain and the content components respectively. Assuming that the VAE encoder  $\mathcal{E}$  is a bijective function and its inverse is the VAE decoder  $\mathcal{D}$ , the combination of two bijective functions  $\mathbf{g}$  and  $\mathcal{E}$  ensures that each  $z = \mathcal{E}(x) = \mathcal{E}(\mathbf{g}(\mathbf{c}, \mathbf{s}))$  corresponds to a unique style and content component pair  $(\mathbf{c}, \mathbf{s})$ .

Based on the above assumption, an intuitive illustration is presented in Fig. 4, showing a Structure-Timbre sampling space of the music latent diffusion model (Liu et al. 2023, 2024c). In our music editing process, the starting latent  $z'_{T_{\text{start}}}$ , copied from DDIM Inversion (Song, Meng, and Ermon 2020) prior  $z_{T_{\text{start}}}$ , provides implicit structure guidance for generation. However, as shown in Fig. 4 (a), relying on this implicit guidance leads to significant structure divergence in the editing process (Song, Meng, and Ermon 2020; Mou et al. 2023; Tumanyan et al. 2023). Since the target prompt  $y$  contains rich semantics of music, it can provide strong semantic guidance for editing, which is significantly stronger than the implicit structure guidance. Therefore, we build an attention repository to store SA features at each inversion time step in the system memory, and use the stored features to provide explicit structure guidance for the editing process, reducing divergence (shown in Fig. 4 (b)).

### Attention-Based Structure Retention

Attention-based Structure Retention (ASR) is the key to transforming the stored attention features into structure guidance. As shown in the right part of Fig. 3, to achieve this, the SA map  $M_t^s$  of the target latent  $z'_t$  is derived from the stored original SA queries  $Q_t^s$  and keys  $K_t^s$  at each reverse time step  $t \in \{0, \dots, T_{\text{start}}\}$  instead of attention features of  $z'_t$ . The manipulated SA can be defined as:

$$\phi_t^s = M_t^s \cdot V_t^s \quad (8)$$

$$M_t^s = \text{Softmax}\left(\frac{Q_t^s K_t^{s\top}}{\sqrt{d^s}}\right) \quad (9)$$

$$Q_t^s = W_{Q^s} \cdot \varphi(z_t), \quad K_t^s = W_{K^s} \cdot \varphi(z_t), \quad V_t^s = W_{V^s} \cdot \varphi(z_t) \quad (10)$$

where  $\phi_t^s$  is the output of a manipulated self-attention layer and  $V_t^s$  is the projected self-attention values of target latent  $z'_t$  at the specific time step  $t$ . In addition, we apply this manipulation to layers 8-14 of AudioLDM 2 (Liu et al. 2024c). Layer selection is analyzed in the Experiments section.

## Experiments

### Baselines

To comprehensively evaluate our approach, we conduct comparisons with several state-of-the-art music editing approaches, including DDPM-Friendly (Manor and Michaeli 2024), MusicMagus (Zhang et al. 2024b), SDEdit (Meng et al. 2021), DDIM Inversion (Song, Meng, and Ermon 2020) and MusicGen (Copet et al. 2023), for both our objective and subjective evaluations. For the music generation




Figure 5: Quantitative Comparison with methods over  $T_{\text{start}}$  range of 300-1000 on MelodiaEdit. The highlighted region is the optimal balance region where shows both text adherence and structural integrity. Our method outperforms other approaches across all  $T_{\text{start}}$  values.

model MusicGen (Copet et al. 2023), we utilize its melody-conditioned medium checkpoint to perform editing tasks.

### Metrics

For objective evaluation, we assess the results using four standard metrics. CLAP (Wu et al. 2023) measures the adherence of the result to the target prompt (higher is better). LPAPS (Paissan et al. 2023) measures the preservation of temporal structure and coherence between the edited audio and source audio (lower is better), while Chroma (Copet et al. 2023) quantifies the preservation of harmonic, melodic, and pitch elements (higher is better). FAD (Kilgour et al. 2018) measures the distributional difference between source and edited music (lower is better).

However, individual metrics can be misleading when evaluating editing performance. Methods that barely modify the source music achieve high structure-related scores but low CLAP scores, creating a false impression of success. To address this, we propose two composite metrics: Adherence-Structure Balance Score (ASB) and Adherence-Musicality Balance Score (AMB). We use harmonic mean (like F1-score) to ensure neither text adherence nor structure preservation dominates the assessment, thereby achieving balance in music editing evaluation.

For subjective evaluation, we introduce the Music Editing Balance (MEB) metric to capture human perception. Additionally, adapted from MusicMagus (Zhang et al. 2024b), we employ Relevance (REL) to assess how well the edited music aligns with the target prompt, and Structural Consistency (CON) to evaluate the consistency of the pitch contour and structural aspects. These metrics are rated on a five-point Likert scale (Likert 1932), with higher scores being better.

**Adherence-Structure Balance Score (ASB)** evaluates the balance between the adherence of the edited music to the target prompt and structural preservation to the source

Method	MusicDelta						ZoME-Bench						MelodiaEdit					
	CLAP↑	LPAPS↓	Chroma↑	FAD↓	ASB↑	AMB↑	CLAP↑	LPAPS↓	Chroma↑	FAD↓	ASB↑	AMB↑	CLAP↑	LPAPS↓	Chroma↑	FAD↓	ASB↑	AMB↑
SDEdit	0.17	6.82	0.20	1.47	0.00	0.00	0.12	6.85	0.19	1.27	0.00	0.00	0.34	5.46	0.49	1.15	0.29	0.29
MusicGen	0.22	6.27	0.24	0.98	0.24	0.31	<b>0.29</b>	6.42	0.22	0.92	0.25	0.46	<u>0.36</u>	5.11	0.59	1.00	0.25	0.54
MusicMagus	0.20	<u>5.06</u>	<u>0.29</u>	0.95	0.28	0.29	0.22	<u>4.82</u>	0.26	0.72	<u>0.63</u>	0.64	0.27	<u>3.32</u>	<b>0.73</b>	<b>0.57</b>	0.00	0.00
DDPM-Friendly	<b>0.35</b>	5.66	0.27	<u>0.88</u>	<u>0.58</u>	<u>0.74</u>	<u>0.23</u>	5.70	<u>0.27</u>	<u>0.68</u>	<u>0.49</u>	<u>0.72</u>	0.34	4.06	<u>0.70</u>	0.67	<u>0.59</u>	<u>0.70</u>
DDIM Inversion	0.30	5.93	0.26	1.03	0.45	0.60	0.22	5.82	0.25	0.77	0.44	0.59	0.35	4.58	0.65	0.90	0.48	0.67
<b>Melodia (Ours)</b>	<u>0.34</u>	<b>4.01</b>	<b>0.32</b>	<b>0.56</b>	<b>1.00</b>	<b>1.00</b>	<b>0.29</b>	<b>3.90</b>	<b>0.29</b>	<b>0.47</b>	<b>1.00</b>	<b>1.00</b>	<b>0.39</b>	<b>3.11</b>	0.68	<u>0.65</u>	<b>1.00</b>	<b>0.88</b>

Table 3: Objective evaluation results across three datasets.

music. ASB integrates CLAP and LPAPS:

$$\text{ASB} = \frac{2 \cdot \mathcal{S}(\mathcal{N}(s_{\text{CLAP}})) \cdot \mathcal{S}(\mathcal{N}(-s_{\text{LPAPS}}))}{\mathcal{S}(\mathcal{N}(s_{\text{CLAP}})) + \mathcal{S}(\mathcal{N}(-s_{\text{LPAPS}}))} \quad (11)$$

where  $s_{\text{CLAP}}$  is the CLAP score,  $s_{\text{LPAPS}}$  is the LPAPS score. We use the negative of  $s_{\text{LPAPS}}$  because lower  $s_{\text{LPAPS}}$  values indicate better preservation.

We apply Z-score normalization  $\mathcal{N}(\cdot)$  and Min-Max scaling  $\mathcal{S}(\cdot)$  to ensure scores fall within  $[0, 1]$ , with 1 indicating best performance and 0 the worst across evaluated methods.

**Adherence-Musicality Balance Score (AMB)** assesses the balance between the adherence of the edited music to the target prompt and musicality preservation of the source music. AMB combines CLAP with Chroma:

$$\text{AMB} = \frac{2 \cdot \mathcal{S}(\mathcal{N}(s_{\text{CLAP}})) \cdot \mathcal{S}(\mathcal{N}(s_{\text{Chroma}}))}{\mathcal{S}(\mathcal{N}(s_{\text{CLAP}})) + \mathcal{S}(\mathcal{N}(s_{\text{Chroma}}))} \quad (12)$$

where  $s_{\text{Chroma}}$  is the Chroma similarity. Both components undergo the same two-step normalization process as in ASB.

**Music Editing Balance (MEB)** is our newly proposed perceptual metric that evaluates how well edited music balances adherence to the target prompt while preserving the original music’s temporal structure and musicality. Higher MEB scores indicate better balance.

## Dataset

We evaluate our method using MusicDelta (Bittner et al. 2014), ZoME-Bench (Liu et al. 2024b), and MelodiaEdit, our newly constructed benchmark. MelodiaEdit comprises 2,015 music-prompt pairs designed for comprehensive intrastem editing evaluation. The dataset contains five subsets: three synthesized subsets with 1,090 pairs generated via AudioLDM 2 (Liu et al. 2024c), and two authentic subsets with 925 pairs reconstructed from Medley-solos-DB (Lostanlen and Cella 2016) and GTZAN datasets.

## Objective Evaluation

We evaluate our proposed method through a comparison with five competing approaches. All diffusion-based methods are implemented using the pre-trained AudioLDM 2 (Liu et al. 2024c) model with 200 inference steps. The target CFG of our method is set to 5.5.

### Quantitative Comparison with Baselines on Datasets.

Table 3 compares our method with baseline approaches. Our method achieves competitive CLAP and the lowest LPAPS across all datasets, indicating superior text adherence and

Method	MusicDelta			zoME-Bench			MelodiaEdit		
	REL↑	CON↑	MEB↑	REL↑	CON↑	MEB↑	REL↑	CON↑	MEB↑
SDEdit	1.24	1.18	1.15	1.73	1.35	1.78	2.32	2.58	2.52
MusicGen	2.43	1.92	2.18	<b>3.02</b>	1.59	2.04	2.87	1.75	<u>3.13</u>
MusicMagus	1.95	2.67	2.31	1.88	<u>2.65</u>	1.69	1.89	<b>3.84</b>	1.21
DDPM-Friendly	<u>3.09</u>	<u>2.88</u>	<u>3.02</u>	2.54	2.01	<u>2.65</u>	<u>2.58</u>	2.92	2.78
DDIM Inversion	2.76	2.34	2.49	2.17	1.68	1.89	3.02	2.67	2.92
<b>Melodia (Ours)</b>	<b>3.21</b>	<b>3.59</b>	<b>3.46</b>	<u>2.85</u>	<b>3.48</b>	<b>3.21</b>	<b>3.38</b>	<u>3.65</u>	<b>3.81</b>

Table 4: Subjective evaluation results across three datasets.

temporal structure preservation. Melodia also maintains good Chroma and FAD performance, effectively preserving harmonic elements and audio quality. While MusicMagus (Zhang et al. 2024b) shows high Chroma and FAD on MelodiaEdit, this indicates editing failure as outputs remain unchanged from the source music, explaining its poor textual adherence. The ASB and AMB results demonstrate the necessity of our proposed composite metrics. Several baselines achieve 0.00 on these metrics due to severe imbalances between text adherence and structure preservation, while Melodia achieves excellent performance on both composite metrics, demonstrating superior balance between editing effectiveness and structure preservation.

### Quantitative Comparison with Baselines across $T_{\text{start}}$ s.

To quantitatively evaluate the adherence of edited music to the target prompt and the fidelity to the original music, we plot the CLAP-LPAPS results on MelodiaEdit for all methods in Fig. 5. MusicMagus (Zhang et al. 2024b) is tested only at  $T_{\text{start}} = 1000$  following its original setting, while other diffusion-based methods are evaluated across multiple  $T_{\text{start}}$  values ranging from 300-1000. The results show that MusicGen (Copet et al. 2023) achieves good textual adherence but poor structural preservation. Notably, other diffusion-based methods also fail to precisely preserve the structure of the source music. In contrast, Melodia addresses this issue by introducing structure guidance and achieves better LPAPS scores for any level of CLAP.

**Qualitative Evaluation.** Fig. 6 compares editing spectrograms between Melodia and baseline methods. Our method successfully preserves the temporal structure of the original music while effectively transferring the target attributes, which can be verified by more consistent rhythmic patterns maintained from the source music and successful frequency pattern transfer. All methods use the same hyperparameter settings as in the quantitative evaluation.

Layers	Timbre Transfer						Genre Transfer						Mood Transfer					
	CLAP↑	LPAPS↓	Chroma↑	FAD↓	ASB↑	AMB↑	CLAP↑	LPAPS↓	Chroma↑	FAD↓	ASB↑	AMB↑	CLAP↑	LPAPS↓	Chroma↑	FAD↓	ASB↑	AMB↑
None	0.34	4.39	0.71	0.87	0.00	0.00	0.39	4.64	0.76	0.87	0.00	0.41	0.31	4.49	0.73	0.80	0.00	0.18
1-16	0.34	<b>2.65</b>	<b>0.81</b>	<b>0.55</b>	0.00	0.00	0.34	<b>2.54</b>	<b>0.84</b>	<b>0.54</b>	0.00	0.00	0.28	<b>1.07</b>	<b>0.96</b>	<b>0.11</b>	0.00	0.00
6-16	0.35	<u>2.96</u>	<u>0.80</u>	<u>0.57</u>	0.22	0.22	0.35	<u>2.59</u>	0.81	<b>0.54</b>	0.29	0.27	0.28	<u>1.08</u>	<u>0.95</u>	<b>0.11</b>	0.00	0.00
10-12	0.39	3.93	0.76	0.89	0.37	0.56	<b>0.40</b>	3.92	0.76	0.88	0.51	0.43	<b>0.35</b>	4.48	0.72	0.99	0.01	0.14
1-7	<u>0.40</u>	4.16	0.73	0.94	0.23	0.32	<b>0.40</b>	4.16	0.73	0.94	0.37	0.00	0.29	4.46	0.70	1.07	0.02	0.00
8-14	<b>0.42</b>	3.49	0.75	0.81	<b>0.68</b>	<b>0.57</b>	<b>0.40</b>	2.63	<u>0.82</u>	<u>0.56</u>	<b>0.98</b>	<b>0.90</b>	<u>0.34</u>	2.61	0.79	<u>0.46</u>	<b>0.67</b>	<b>0.49</b>

Table 5: Performance comparison across different layer selections for editing tasks.

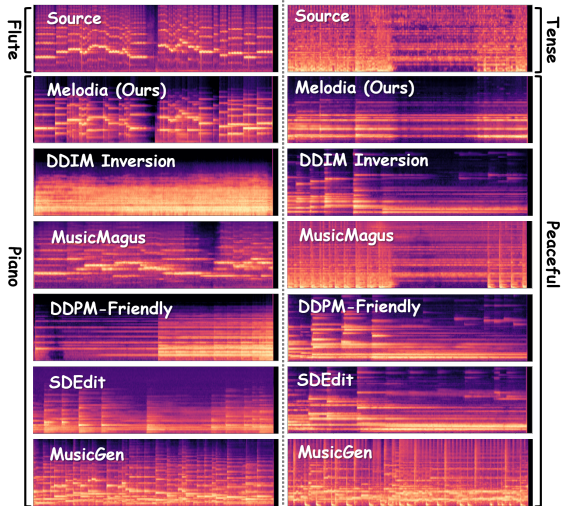


Figure 6: Spectrogram comparison of editing results between Melodia and baseline methods. Melodia successfully preserves the structure and achieves effective attribute transfer. Soundtracks can be found in our demo page.

## Subjective Evaluation

We implemented subjective listening evaluations across three datasets. A total of 100 participants were recruited from the MIR community and music forums, with 77 valid responses retained after screening. All participants had some degree of musical experience. The test is approved by the Institutional Review Board (IRB). We randomly selected 10 test samples from each dataset, and participants rated the edited results from each method. As shown in Table 4, Melodia substantially outperforms all baseline methods on the MEB metric, indicating that our method successfully achieves the desired balance between textual alignment and structural preservation. Melodia also obtained high scores on REL and CON metrics, confirming its superior performance in semantic alignment and structure preservation.

## Additional Analysis

**Effects of Layer Selection.** The selection of SA replacement layers is crucial in balancing structure preservation and attribute modification. Experimental results on MelodiaEdit demonstrate that replacing SA maps in layers 8–14 of AudioLDM 2 achieves optimal performance. As shown in

Method	CLAP↑	LPAPS↓	Chroma↑	FAD↓
Stable Audio Open	0.43	6.47	0.12	1.42
Stable Audio Open + Melodia	<b>0.44</b>	<b>6.19</b>	<b>0.21</b>	<b>1.19</b>

Table 6: Experimental results on Stable Audio Open.

Tab. 5, this layer selection yields the highest balance scores (ASB and AMB) while maintaining low structural disruption (LPAPS) and high textual alignment (CLAP). Visual evidence in Fig. 2 also confirms that SA manipulation in layers 8-14 successfully preserves rhythmic patterns and melodic contours during attribute transfer.

**Results on Another Pre-Trained Diffusion Model.** To evaluate the generalization capability of Melodia across different diffusion models, we conducted experiments with the *Stable Audio Open* (Evans et al. 2025) for timbre transfer on MelodiaEdit. As shown in Tab. 6, Melodia achieves improvements in both textual adherence (CLAP) and structural integrity (LPAPS and Chroma) compared to the baseline. These results confirm Melodia’s generalizability across different model architectures. Additionally, Melodia operates effectively across different sampling rates: 16kHz on AudioLDM2 (Liu et al. 2024c) and 44.1kHz on Stable Audio Open (Evans et al. 2025).

## Conclusion

We proposed Melodia, a training-free music editing method achieving optimal balance between textual adherence and structural integrity. Our approach stems from a key insight through attention probing analysis: cross-attention maps encode musical attributes, while self-attention maps preserve temporal structure. By selectively manipulating self-attention maps during denoising, we enable effective editing without requiring textual descriptions of the original music.

Our experimental results demonstrate that Melodia outperforms existing approaches across various datasets. Through our attention repository and Structure Retention technique, we achieved excellent performance in both objective and subjective evaluations. Melodia excels where previous approaches fail, simultaneously maintaining original melody while transforming to target attributes, without additional training or source music descriptions.

## Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 62202174, in part by the GJYC program of Guangzhou under Grant 2024D01J0081, and in part by the ZJ program of Guangdong under Grant 2023QN10X455, and in part by the Fundamental Research Funds for the Central Universities under Grant 2025ZYGXZR053.

## References

- Agostinelli, A.; Denk, T. I.; Borsos, Z.; Engel, J.; Verzetti, M.; Caillon, A.; Huang, Q.; Jansen, A.; Roberts, A.; Tagliasacchi, M.; et al. 2023. Musiclm: Generating music from text. *arXiv preprint arXiv:2301.11325*.
- Bittner, R. M.; Salamon, J.; Tierney, M.; Mauch, M.; Cannam, C.; and Bello, J. P. 2014. Medleydb: A multitrack dataset for annotation-intensive mir research. In *Ismir*, volume 14, 155–160.
- Cao, M.; Wang, X.; Qi, Z.; Shan, Y.; Qie, X.; and Zheng, Y. 2023. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF international conference on computer vision*, 22560–22570.
- Chen, K.; Wu, Y.; Liu, H.; Nezhurina, M.; Berg-Kirkpatrick, T.; and Dubnov, S. 2024. Musicldm: Enhancing novelty in text-to-music generation using beat-synchronous mixup strategies. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1206–1210. IEEE.
- Clark, K.; Khandelwal, U.; Levy, O.; and Manning, C. D. 2019. What does bert look at? an analysis of bert’s attention. *arXiv preprint arXiv:1906.04341*.
- Copet, J.; Kreuk, F.; Gat, I.; Remez, T.; Kant, D.; Synnaeve, G.; Adi, Y.; and Défossez, A. 2023. Simple and controllable music generation. *Advances in Neural Information Processing Systems*, 36: 47704–47720.
- Evans, Z.; Parker, J. D.; Carr, C.; Zukowski, Z.; Taylor, J.; and Pons, J. 2025. Stable audio open. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Han, B.; Dai, J.; Hao, W.; He, X.; Guo, D.; Chen, J.; Wang, Y.; Qian, Y.; and Song, X. 2023. Instructme: An instruction guided music edit and remix framework with latent diffusion models. *arXiv preprint arXiv:2308.14360*.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Huang, Q.; Jansen, A.; Lee, J.; Ganti, R.; Li, J. Y.; and Ellis, D. P. 2022. Mulan: A joint embedding of music audio and natural language. *arXiv preprint arXiv:2208.12415*.
- Kilgour, K.; Zuluaga, M.; Roblek, D.; and Sharifi, M. 2018. Fr’echet audio distance: A metric for evaluating music enhancement algorithms. *arXiv preprint arXiv:1812.08466*.
- Kingma, D. P.; Welling, M.; et al. 2013. Auto-encoding variational bayes.
- Lan, G. L.; Shi, B.; Ni, Z.; Srinivasan, S.; Kumar, A.; Ellis, B.; Kant, D.; Nagaraja, V.; Chang, E.; Hsu, W.-N.; et al. 2024. High Fidelity Text-Guided Music Editing via Single-Stage Flow Matching. *arXiv preprint arXiv:2407.03648*.
- Likert, R. 1932. A technique for the measurement of attitudes. *Archives of psychology*.
- Liu, B.; Wang, C.; Cao, T.; Jia, K.; and Huang, J. 2024a. Towards understanding cross and self-attention in stable diffusion for text-guided image editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7817–7826.
- Liu, H.; Chen, Z.; Yuan, Y.; Mei, X.; Liu, X.; Mandic, D.; Wang, W.; and Plumbley, M. D. 2023. Audioldm: Text-to-audio generation with latent diffusion models. *arXiv preprint arXiv:2301.12503*.
- Liu, H.; Wang, J.; Li, X.; Huang, R.; Liu, Y.; Xu, J.; and Zhao, Z. 2024b. MEDIC: Zero-shot Music Editing with Disentangled Inversion Control. *arXiv preprint arXiv:2407.13220*.
- Liu, H.; Yuan, Y.; Liu, X.; Mei, X.; Kong, Q.; Tian, Q.; Wang, Y.; Wang, W.; Wang, Y.; and Plumbley, M. D. 2024c. Audioldm 2: Learning holistic audio generation with self-supervised pretraining. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Liu, N. F.; Gardner, M.; Belinkov, Y.; Peters, M. E.; and Smith, N. A. 2019. Linguistic knowledge and transferability of contextual representations. *arXiv preprint arXiv:1903.08855*.
- Lostanlen, V.; and Cella, C.-E. 2016. Deep convolutional networks on the pitch spiral for musical instrument recognition. *arXiv preprint arXiv:1605.06644*.
- Manor, H.; and Michaeli, T. 2024. Zero-shot unsupervised and text-based audio editing using DDPM inversion. *arXiv preprint arXiv:2402.10009*.
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2021. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*.
- Mou, C.; Wang, X.; Song, J.; Shan, Y.; and Zhang, J. 2023. Dragondiffusion: Enabling drag-style manipulation on diffusion models. *arXiv preprint arXiv:2307.02421*.
- Paissan, F.; Della Libera, L.; Wang, Z.; Ravanelli, M.; Smaragdis, P.; and Subakan, C. 2023. Audio editing with non-rigid text prompts. *arXiv preprint arXiv:2310.12858*.
- Plitsis, M.; Kouzelis, T.; Paraskevopoulos, G.; Katsouros, V.; and Panagakis, Y. 2024. Investigating personalization methods in text to music generation. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1081–1085. IEEE.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22500–22510.
- Schneider, F.; Kamal, O.; Jin, Z.; and Schölkopf, B. 2023. Moüsai: Text-to-music generation with long-context latent diffusion. *arXiv preprint arXiv:2301.11757*.
- Shrestha, S.; and Fu, X. 2024. Content-Style Learning from Unaligned Domains: Identifiability under Unknown Latent Dimensions. *arXiv preprint arXiv:2411.03755*.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Sørensen, M.; Kanatsoulis, C. I.; and Sidiropoulos, N. D. 2021. Generalized canonical correlation analysis: A subspace intersection approach. *IEEE Transactions on Signal Processing*, 69: 2452–2467.
- Tumanyan, N.; Geyer, M.; Bagon, S.; and Dekel, T. 2023. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1921–1930.
- Wang, Y.; Ju, Z.; Tan, X.; He, L.; Wu, Z.; Bian, J.; et al. 2023. Audit: Audio editing by following instructions with latent diffusion models. *Advances in Neural Information Processing Systems*, 36: 71340–71357.
- Wu, Y.; Chen, K.; Zhang, T.; Hui, Y.; Berg-Kirkpatrick, T.; and Dubnov, S. 2023. Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Zhang, Y.; Ikemiya, Y.; Choi, W.; Murata, N.; Martínez-Ramírez, M. A.; Lin, L.; Xia, G.; Liao, W.-H.; Mitsufuji, Y.; and Dixon, S. 2024a. Instruct-MusicGen: Unlocking Text-to-Music Editing for Music Language Models via Instruction Tuning. *arXiv preprint arXiv:2405.18386*.
- Zhang, Y.; Ikemiya, Y.; Xia, G.; Murata, N.; Martínez-Ramírez, M. A.; Liao, W.-H.; Mitsufuji, Y.; and Dixon, S. 2024b. Musicmagus: Zero-shot text-to-music editing via diffusion models. *arXiv preprint arXiv:2402.06178*.