

# JEN-1 Composer: A Unified Framework for High-Fidelity Multi-Track Music Generation

Yao Yao, Peike Li\*, Boyu Chen, Alex Wang

Jen Music AI  
alex@jenmusic.ai

## Abstract

With rapid advances in generative artificial intelligence, the text-to-music synthesis task has emerged as a promising direction for music generation. Nevertheless, achieving precise control over multi-track generation remains an open challenge. While existing models excel in directly generating multi-track mix, their limitations become evident when it comes to composing individual tracks and integrating them in a controllable manner. This departure from the typical workflows of professional composers hinders the ability to refine details in specific tracks. To address this gap, we propose JEN-1 Composer, a unified framework designed to efficiently model marginal, conditional, and joint distributions over multi-track music using a single model. Building upon an audio latent diffusion model, JEN-1 Composer extends the versatility of multi-track music generation. We introduce a progressive curriculum training strategy, which gradually escalates the difficulty of training tasks while ensuring the model’s generalization ability and facilitating smooth transitions between different scenarios. During inference, users can iteratively generate and select music tracks, thus incrementally composing entire musical pieces in accordance with the Human-AI co-composition workflow. Our approach demonstrates state-of-the-art performance in controllable and high-fidelity multi-track music synthesis, marking a significant advancement in interactive AI-assisted music creation.

**Demo** — <https://www.jenmusic.ai/audio-demos>

## 1 Introduction

The rapid evolution of generative modeling has positioned AI-driven music generation as a prominent field, merging research innovation with practical applications in the music industry. Early systems like Music Transformer (Huang et al. 2018) and MuseNet (Payne 2019), which utilized symbolic representations (Engel et al. 2017), were pivotal in translating textual descriptions into MIDI-style outputs. Although these methods were groundbreaking, their dependence on predefined virtual synthesizers often compromised the audio quality and restricted the diversity of their musical outputs.

Recent advancements in text-to-music synthesis, as demonstrated by models like MusicGen (Copet et al. 2024),

MusicLM (Agostinelli et al. 2023), and Jen-1 (Li et al. 2024), represent a significant leap forward in directly generating authentic audio waveforms from textual prompts. These innovations have greatly expanded the versatility and diversity of music generation, bypassing the need for extensive musical theory knowledge and traditional symbolic representations. However, their focus on producing composite audio mixes rather than discrete, manipulable tracks limits the level of creative control necessary in professional music production environments. Similarly, the introduction of digital audio workstations and the expansion of available timbres have indeed revolutionized musical creativity—enabling composers to explore complex harmonies, melodies, and rhythms beyond the confines of physical instruments (Zhu et al. 2020)—but these tools still impose significant barriers. Despite their advancements, they require a deep understanding of music theory and proficiency in symbolic musical notation, which continue to pose challenges for many aspiring musicians and composers.

In response to these challenges, we propose JEN-1 Composer, a framework designed to democratize music production by streamlining the creative process. This framework employs end-to-end training to intuitively grasp the relationships between different tracks, enabling audio-to-audio orchestration that learns directly from waveform datasets. By allowing both direct audio and text input, JEN-1 Composer not only simplifies user interaction but also expands creative freedom through detailed manipulation of individual tracks.

JEN-1 Composer is a comprehensive generative framework designed to model the marginal, conditional, and joint distributions of multi-track music within a single model. By leveraging the Jen-1 (Li et al. 2024) audio latent diffusion model as a foundation, our approach effectively and efficiently manages these distributions concurrently. To extend the capabilities of Jen-1, we introduce several key enhancements: (a) We design specialized input-output configurations to handle latent representations of multiple music tracks, allowing the model to effectively capture the temporal relationships and harmonic coherence across these tracks. (b) We incorporate timestep vectors to govern the generation of individual tracks, providing the necessary flexibility for fine-grained control during the generation process. (c) We augment conventional text prompts with prefix prompts to clearly define generation tasks, reducing ambi-

\*Work done while was at Jen Music AI.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

guity and improving model performance. Additionally, we employ a curriculum training strategy that gradually introduces more complex tasks, from generating single tracks to orchestrating intricate combinations of multiple tracks.

To further bridge the gap between AI capabilities and human creativity, we introduce a Human-AI co-composition workflow, as illustrated in Figure 1. This approach enables iterative refinement of music tracks during the model’s inference phase, allowing producers to collaboratively adjust and blend AI-generated tracks to align with their creative visions. Through this workflow, users can directly influence the music generation process by providing feedback on textual prompts and previously generated music tracks, ensuring that all tracks meet their precise standards.

In summary, the contributions of this work are four-fold:

1. We introduce a collaborative music generation workflow that seamlessly integrates human creativity with AI, designed for the iterative creation of multi-track music within an audio-based framework.
2. We present JEN-1 Composer, a unified framework that effectively models marginal, conditional, and joint distributions for multi-track music generation using a single audio latent diffusion model.
3. We design an intuitive curriculum training strategy that progressively enhances the model’s capability to generate complex musical compositions.
4. We provide comprehensive quantitative and qualitative evaluations demonstrating that JEN-1 Composer achieves state-of-the-art performance in generating diverse track combinations, advancing the flexibility and creativity of music production.

## 2 Related Work

In this section, we discuss the research background relevant to our work, focusing on two key areas: conditional music generation and multi-track music generation.

### 2.1 Conditional Music Generation

Conditional music generation has become a cornerstone in the field, with models incorporating a range of inputs, from low-level control signals like lyrics (Yu, Srivastava, and Canales 2021) and MIDI sequences (Muhammed et al. 2021) to high-level abstract representations such as text (Kreuk et al. 2022; Agostinelli et al. 2023; Liu et al. 2023) and images (Huang et al. 2023b). While these conditioning methods enable tailored outputs, they also present challenges in data alignment and model training. The scarcity of well-aligned data has spurred advances in self-supervised learning (Marafioti et al. 2019; Borsos et al. 2023) to enhance model generalization across diverse musical contexts. Due to the complexity of raw audio waveforms, direct generation is often impractical (Gârbacea et al. 2019), leading researchers to develop feature extraction and representation strategies. Techniques like VQ-VAE and VQ-GAN, using mel-spectrograms (Van Den Oord, Vinyals et al. 2017; Creswell et al. 2018; Huang et al. 2023a), and quantization-based methods that transform waveforms into compact rep-

resentations (Zeghidour et al. 2021; Défossez et al. 2022), have been particularly influential.

The emergence of non-autoregressive models, particularly diffusion models (Ho, Jain, and Abbeel 2020), has significantly advanced the field. Models like MeLoDy (Lam et al. 2024) and Jen-1 (Li et al. 2024; Chen et al. 2024) have demonstrated exceptional capabilities in producing high-fidelity music. Building on these advancements, our JEN-1 Composer integrates high-level textual prompts and inter-track dependencies to enhance alignment and harmonic cohesion among tracks, allowing for sophisticated control and compositional flexibility. Unlike previous models that typically produce composite mixes, JEN-1 Composer supports track-wise generation and iterative refinement, aligning with real-world music production workflows and fostering collaboration between humans and AI.

### 2.2 Multi-track Music Generation

Multi-track music generation enables the simultaneous creation of interdependent tracks, offering greater complexity and cohesion in music composition. Early models, such as MuseGAN (Dong et al. 2018), leveraged GANs but faced issues with training instability, limited diversity, and suboptimal sound quality. Subsequent advancements, like MIDI-Sandwich2 (Liang, Wu, and Cao 2019) with hierarchical RNNs and VAEs, and transformer-based models such as MMM (Ens and Pasquier 2020) and MTMG (Jin et al. 2020), improved inter-track dependency modeling. More recently, MTT-GAN (Jin et al. 2022) integrated GANs with transformers to enhance adherence to musical rules. Traditional approaches primarily relied on symbolic representations like MIDI, which constrained their ability to capture nuanced audio textures. Recent methods shift to direct audio modeling to avoid fidelity losses from intermediate representations. For example, StemGen (Parker et al. 2024) uses transformers for single-stem generation, while Stable Audio (Evans et al. 2024) and Diff-A-Riff (Nistal et al. 2024) adopt cross-attention for prompt-based conditioning but lack multi-track alignment. Models like MSDM (Mariani et al. 2024) handle multi-track spectrogram generation but do not support prompt conditioning.

Our JEN-1 Composer addresses these challenges through a diffusion-based framework designed for multi-track audio generation in the audio latent space. By directly modeling audio, it overcomes the fidelity loss inherent in symbolic and spectrum-based methods, capturing richer and more nuanced sound textures. The model leverages concatenation-based alignment to ensure inter-track coherence and incorporates text-prompt conditioning, enabling a flexible and interactive human-AI co-creation workflow. This approach surpasses traditional symbolic and spectrum-based methods in both audio fidelity and creative versatility (Kong et al. 2020; Liu et al. 2023).

## 3 Preliminary

### 3.1 Diffusion Model

Diffusion models (Ho, Jain, and Abbeel 2020) are generative models that produce high-quality samples through iter-

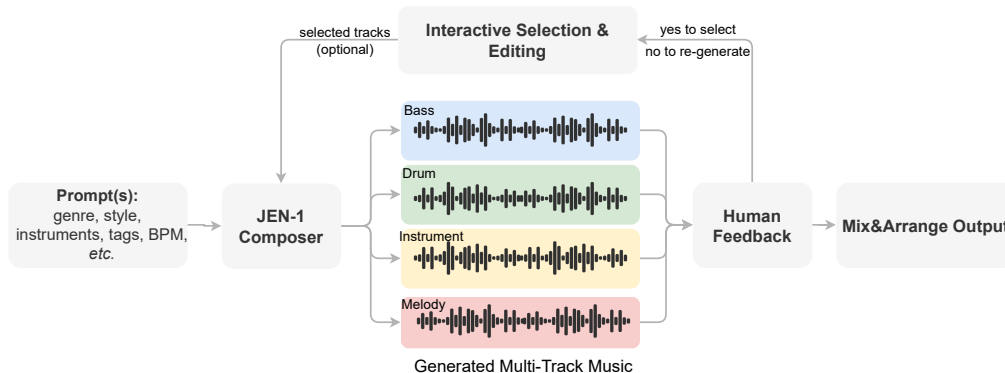


Figure 1: The Human-AI co-composition workflow of JEN-1 Composer. JEN-1 Composer generates multiple music tracks based on user-provided text prompts (specifying genres, eras, rhythms, *etc.*) and optional audio feedback, where users can select, edit, or upload tracks. The human feedback guides the generation of target tracks, ensuring temporal alignment and musical coherence. The iterative process of human feedback and AI generation continues until a harmonious and cohesive musical piece is achieved.

ative denoising. The process begins by gradually corrupting the original data  $x_0$  with Gaussian noise over a series of timesteps in a forward process, where each noisy sample  $x_t$  is generated as:

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon_t, \quad (1)$$

with  $\epsilon_t$  as the standard Gaussian noise,  $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ , and  $\alpha_t = 1 - \beta_t$ , where the sequence of  $\beta_t$  is the noise schedule that controls the level of corruption over time.

In the reverse process, the diffusion model aims to recover  $x_0$  by iteratively denoising  $x_t$ . A noise prediction model, parameterized by  $\theta$ , is trained to estimate the noise  $\epsilon_t$  in  $x_t$  at each timestep  $t$  by minimizing the following loss function:

$$\min_{\theta} \mathbb{E}_{t, x_0, \epsilon_t} \|\epsilon_t - \epsilon_{\theta}(x_t, t)\|_2^2, \quad (2)$$

where  $t$  is uniformly sampled from  $\{1, 2, \dots, T\}$ . With the optimized noise predictor,  $x_0$  can be approximated by sampling from a Gaussian model  $p(x_{t-1} | x_t) = \mathcal{N}(x_{t-1} | \mu_t(x_t), \sigma_t^2 I)$  in a stepwise manner (Bao et al. 2023). The optimal mean for this Gaussian, under maximum likelihood estimation, is:

$$\mu_t^*(x_t) = \frac{1}{\sqrt{\alpha_t}} \left( x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_{\theta}(x_t, t) \right). \quad (3)$$

By iteratively applying this process, the diffusion model refines the noise, progressively generating new samples that closely resemble the original training data.

### 3.2 Audio Latent Diffusion Model

Modeling raw audio waveforms directly poses challenges due to their high dimensionality, where  $x_0 \in \mathbb{R}^{C \times S}$  represents the waveform, with  $C$  denoting the number of channels and  $S$  indicating the sequence length. To address this, Jen-1 (Li et al. 2024) extends the Latent Diffusion Model (LDM) (Rombach et al. 2022) framework, originally formulated for images, to the domain of audio generation. In the

Jen-1 architecture, the audio waveform  $x_0$  is mapped to a lower-dimensional latent representation  $z_0 \in \mathbb{R}^{D \times S'}$  via an audio encoder  $f_{\phi}$ , and then reconstructed back to the original waveform  $\hat{x}_0$  through an audio decoder  $g_{\psi}$ . Here  $S' \ll S$  is the compressed sequence length and  $D$  is the latent dimension. This process is denoted as:

$$z_0 = f_{\phi}(x_0), \quad \hat{x}_0 = g_{\psi}(z_0) \approx x_0. \quad (4)$$

In Jen-1 (Li et al. 2024), the diffusion model  $\theta$  operates in the audio latent space, predicting the noise  $\hat{\epsilon}_t = \epsilon_{\theta}(z_t, e, t)$  and iteratively denoising from Gaussian noise to generate the final latent  $\hat{z}_0$ . The variable  $e$  represents the embedding of conditioning inputs, such as text prompts, guiding the generation process. Similar to LDM, Jen-1 utilizes a U-Net architecture (Ronneberger, Fischer, and Brox 2015) as the backbone for its diffusion process. Specifically adapted for audio data, Jen-1 replaces the 2D convolutions used in image processing with 1D convolutions tailored for audio latent representations. The model consists of a sequence of blocks, including AttnDownBlock1D, UNetMidBlock1D, and AttnUpBlock1D, which integrate residual 1D convolutional layers with cross-attention transformers (Vaswani et al. 2017). The overall architecture is depicted in Figure 2.

## 4 Method

### 4.1 Multi-track Music Generation

To enable JEN-1 Composer to handle multi-track input and output for joint modeling, we make several critical modifications to Jen-1’s single-track architecture. As elaborated below, the input-output representation, timestep vectors, and prompt prefixes are adapted to fit multi-track distributions efficiently using a single model.

**Multi-track Input-Output Representation.** We extend the single-track input paradigm of Jen-1 to accommodate multi-track inputs, denoted as  $\mathbf{X} = [x_0^1, \dots, x_0^K]$ , where  $x_0^i$  represents the  $i$ -th track and  $K$  denotes the total number of tracks. Each track undergoes encoding into the audio latent space,

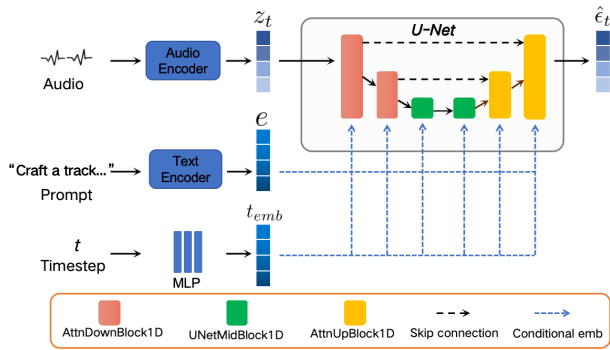


Figure 2: The U-Net architecture used in Jen-1.

yielding latent representations  $z_0^i = f_\phi(x_0^i) \in \mathbb{R}^{D \times S'}$  prior to being inputted to the audio latent diffusion model. These latent representations are concatenated along the channel dimension to form the input latent variables  $\mathbf{Z} \in \mathbb{R}^{KD \times S'}$ . During inference, the output of the audio latent diffusion model is split into  $K$  tracks, with each denoised latent variable reconstructed into waveform via the pre-trained audio decoder, denoted as  $\hat{x}_0^i = g_\psi(\hat{z}_0^i)$ . The extension of the input-output representation to multi-track enables explicit modeling of inter-track dependencies and consistency, crucial for high-quality multi-track generation that is absent in single-track models.

**Individual Timestep Vectors.** Introducing separate timesteps for each track not only provides precise control over the generation process but also enables unified distribution modeling. This is achieved by extending the scalar timestep  $t$  in Jen-1 to a multi-dimensional vector  $\mathbf{T} = [t_1, \dots, t_K]$ . Each  $t_i$  determines the corresponding latent variable  $z^i$  in  $\mathbf{Z} = [z^1, \dots, z^K]$  according to the diffusion forward process defined in Equation (1). In the diffusion model, these timesteps are independently learned for each track and concatenated to form the conditional embedding. The process is formalized as follows:

$$z^i = \begin{cases} z_T^i, & \text{if } t_i = T \\ z_0^i, & \text{if } t_i = 0 \\ z_{t_i}^i, & \text{if } 0 < t_i < T \end{cases} \quad (5)$$

For  $k \geq 2$  target generation tracks, we adopt a uniform timestep  $t$  across these tracks, mirroring the modeling of their joint probability distribution. Specifically, if  $k < K$ , two modes are considered for the remaining  $(K - k)$  tracks. Firstly, the conditional generation mode sets all timesteps to 0, representing latent variables corresponding to original waveforms, akin to conditional generation. Here, the corresponding latent variables and timestep vectors are denoted as  $\mathbf{Z}_c$  and  $\mathbf{T}_c$ , respectively. Secondly, the unconditional generation mode involves fixing all non-target timesteps to  $T$ , indicating perturbation of corresponding latent variables to approximate Gaussian random noise, akin to marginal generation. Correspondingly, the latent variables and timestep vectors are labeled as  $\mathbf{Z}_m$  and  $\mathbf{T}_m$ . During loss computation, emphasis is laid solely on the channels corresponding to the

target tracks, while inference benefits from the Classifier-Free Guidance (CFG) technique (Ho and Salimans 2022). Specifically, if  $k < K$ , the alignment across tracks and generation quality are enhanced via the expression:

$$\hat{e} = (1 - \lambda) \epsilon_\theta(\mathbf{Z}_m, e, \mathbf{T}_m) + \lambda \epsilon_\theta(\mathbf{Z}_c, e, \mathbf{T}_c), \quad (6)$$

where  $\lambda$  denotes the guidance scale. Illustrations presented in Figure 3 showcase diverse scenarios of a straightforward two-track generation task, providing additional clarity on the concept of achieving unified distribution modeling through the manipulation of timestep vectors.

## 4.2 Integrating Task Tokens as Prefix Prompts

We enhance traditional text prompts, which typically describe the musical content and style, by integrating task-specific tokens as prefix prompts. These tokens act as explicit directives, similar to command flags in programming, providing clear and concise instructions regarding the generation task at hand. By using specific prefixes like “[bass & drum generation]”, we direct the model’s focus to the production of target tracks, such as bass and drums. This approach not only specifies the generation objectives but also significantly diminishes ambiguity, thus improving both the fidelity and relevance of the generated content.

## 4.3 Progressive Curriculum Training Strategy

We introduce a progressive curriculum training strategy designed to systematically enhance the model’s capability to generate coherent multi-track audio sequences while accommodating varying levels of conditioning and noise injection. This strategy includes curriculum decay and task allocation, a strategic sampler for conditional and marginal generation modes, and self-bootstrapping training to improve model generalization.

The training begins with single-track text-to-music generation, establishing a strong foundation for the model. As the model advances, more complex multi-track tasks are gradually introduced. This progression is carefully managed by reducing the sampling probabilities of simpler tasks, allowing the model to develop the ability to generate harmonically aligned tracks across multiple channels. Each task involves multi-track audio input and output, with latent representations configured as described in Equation (5). During this phase, the model’s learning is focused on critical aspects by computing losses only for target tracks, while non-target tracks are masked. This structured approach facilitates efficient learning, enabling the model to generate high-fidelity audio compositions.

**Curriculum Decay and Task Allocation.** The curriculum starts with single-track ( $k = 1$ ) tasks, focusing on conditional generation using other tracks as signals or simpler marginal generation tasks. As training progresses, the focus shifts towards multi-track generation ( $2 \leq k < K$ ), with increased sampling probabilities for more complex tasks over time. Ultimately, the curriculum incorporates joint generation tasks ( $k = K$ ) driven solely by text prompts.

**Sampler for Conditional and Marginal Generation.** A strategic sampler is employed when fewer target tracks are

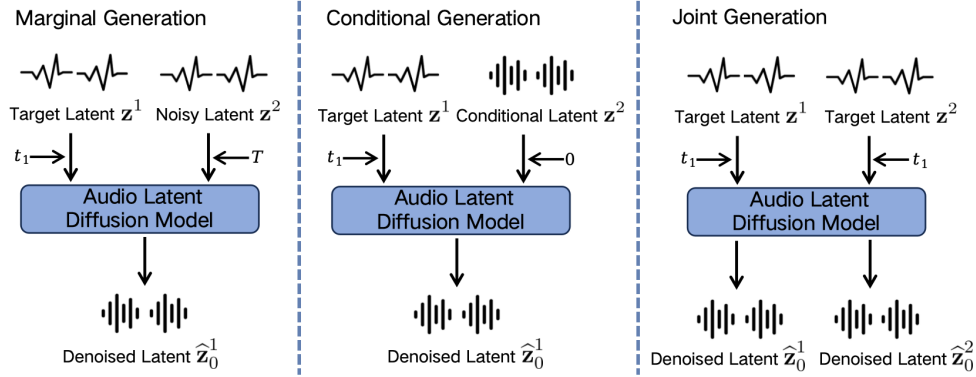


Figure 3: Illustration of three generation modes using independent timesteps as indicators. In Marginal Generation, non-target track latents are fixed as Gaussian noise (timestep  $T$ ) to minimize their impact on the target track’s latent. Conditional Generation designates a timestep of 0 for a conditional track, guiding the target track’s generation. Joint Generation synchronizes multiple target tracks by sharing the same timestep  $t$ , allowing for coordinated denoising from  $T$  to 0.

---

#### Algorithm 1: Human-AI Co-composition Workflow

---

```

1: Input: Text prompt, user-provided tracks  $\mathbf{S}$  (optional)
2: Output: Set of selected and refined tracks  $\mathbf{S}$ 
3:  $e \leftarrow$  Embedding of the given prompt
4: while  $\mathbf{S}$  is empty do
5:   # Joint Generation
6:    $(\hat{x}_0^1, \dots, \hat{x}_0^K) \leftarrow$  Model.GenerateTracks( $e$ )
7:    $\mathbf{S} \leftarrow$  User.selectAndRefineTracks( $\hat{x}_0^1, \dots, \hat{x}_0^K$ )
8: end while
9: while not all  $K$  tracks are satisfactory do
10:  # Using the CFG technique defined in Equation (6)
11:   $(\hat{x}_0^1, \dots, \hat{x}_0^K) \leftarrow$  Model.GenerateTracks( $\mathbf{S}, e$ )
12:  # Update  $\mathbf{S}$ 
13:   $\mathbf{S} \leftarrow \mathbf{S} \cup$  User.selectAndRefineTracks( $\hat{x}_0^1, \dots, \hat{x}_0^K$ )
14: end while

```

---

generated than available ( $k < K$ ). The sampler assigns non-target tracks a timestep of either 0 or  $T$ : with probability  $p_1$ , a timestep of 0 is chosen to encourage conditioning, and with probability  $1 - p_1$ , a timestep of  $T$  is selected for non-conditioned generation. This approach allows the model to effectively learn both conditional and marginal generation, preparing it for CFG technique implementation during inference and enhancing its overall performance in generating coherent music tracks.

**Incorporation of Self-Bootstrapping Training.** In later training stages, self-bootstrapping is introduced with a probability  $p_2$  to improve generalization and align with the Human-AI co-composition workflow. During this phase, tracks generated by a teacher model—using an exponential moving average of the model’s parameters—replace a portion of the ground truth-aligned conditional input tracks. This technique refines the model’s alignment and synchronization capabilities, expands the training dataset, and enhances generalization, which is crucial for performance in real-world, interactive environments.

## 4.4 Human-AI Co-composition Workflow

During inference, our model supports the conditional generation of multiple tracks given 0 to  $K - 1$  input tracks. To facilitate Human-AI collaborative music creation, we propose an interactive generation procedure, outlined in Algorithm 1. The proposed interactive inference approach effectively integrates human creativity with AI capabilities, enabling a collaborative music generation process. This workflow offers three primary benefits:

- **Enhanced Refinement.** The iterative feedback mechanism allows users to refine each track, enabling nuanced improvements that purely AI-driven generation may struggle to achieve. By selecting and refining satisfactory tracks, users steer the generation process toward desired outcomes, filtering out low-quality outputs.
- **Alignment with Human Aesthetics.** The interaction between human creators and the model enhances the AI’s understanding of human aesthetic preferences and sound quality standards. This ongoing collaboration ensures that the generated tracks align more closely with artistic intent and professional expectations.
- **Creative Control and Engagement.** The collaborative experience empowers human producers, providing a sense of control over the creative process. By balancing AI-driven generation with human input, the workflow ensures that both improvisation and structural coherence are maintained, enabling the realization of creative visions with AI assistance.

## 5 Experiment

### 5.1 Experimental Setting

We conducted extensive experiments to evaluate the capabilities of JEN-1 Composer, focusing on its performance across various dimensions to understand its potential in real-world applications.

**Dataset Setup.** Like Diff-A-Riff (Nistal et al. 2024), we prioritize high-quality audio generation suitable for professional use, which requires training on proprietary datasets

due to their superior sound quality compared to open-source MIDI-based data. For JEN-1 Composer, we utilized an 800-hour private studio recording dataset comprising five temporally aligned tracks: bass, drums, instrument, melody, and the final mix. Each track is annotated with metadata, including genres (*e.g.*, blues, folk), instruments (*e.g.*, guitar, piano), moods (*e.g.*, cheerful, romantic), tempo, keywords, and themes. The dataset is divided into 640 hours for training and 160 hours for testing. Tracks are randomly sliced into aligned segments of varying lengths to ensure the model learns inter-track dependencies, enabling it to generate cohesive, high-quality multi-track music guided by text prompts.

**Evaluation Metrics.** Our evaluation employs both quantitative and qualitative metrics to assess the model’s performance.

**Quantitative Evaluation.** We use the CLAP score (Elizalde et al. 2023) to evaluate how well the generated music aligns with the intended semantic content of the text prompts. Higher CLAP scores indicate better alignment. In our study, we compute these scores for both individual tracks and the aggregated mixed track (MIXED CLAP) to assess the effectiveness of JEN-1 Composer in adhering to textual descriptions. For comparison models, we apply Demucs (Défossez 2021; Rouard, Massa, and Défossez 2023) to separate tracks before calculating their individual CLAP scores. We also use Fréchet Audio Distance (FAD) (Roblek et al. 2019) as metric. We evaluated the quality of our generated mixed audio by comparing it to mixed audio from both our proprietary dataset and the public Slakh2100 dataset (Manilow et al. 2019) using the FAD metric (lower is better), computed with VGGish embeddings (Hershey et al. 2017). This comparison was done without fine-tuning, providing a zero-shot evaluation. This approach ensures a fair and consistent evaluation of the contextual relevance and musical fidelity of the generated outputs across different methods.

**Qualitative Evaluation.** We employ the Relative Preference Ratio (RPR) to capture human judgment of audio quality. Multiple raters evaluate pairs of audio samples—one generated by JEN-1 Composer and the other by a comparison model—without knowing the origin of each sample. Raters assess based on audio quality, coherence, harmony, and adherence to the text prompt. The RPR is recorded as a percentage, where 0% indicates no preference for JEN-1 Composer’s output, and 100% indicates complete preference. This metric captures subjective preferences, providing insight into the perceived quality and effectiveness of the generated music.

By integrating CLAP scores, FAD, and RPR, we offer a comprehensive evaluation framework that balances objective alignment with subjective human perception, ensuring a thorough assessment of the model’s strengths and areas for improvement.

**Implementation Details.** Our task involves generating four distinct audio tracks: bass, drums, instrument, and melody, alongside a synthesized mixed track. All tracks are high-fidelity stereo audio sampled at 48 kHz. We utilize the 48k version of the pre-trained EnCodec (Défossez et al. 2022), resulting in a latent space representation of 150 frames per

second, each with 128 dimensions. The volumes of individual tracks are adjusted to ensure consistent relative loudness before encoding. For text encoding, we employ the pre-trained Flan-T5-Large model (Chung et al. 2024), which provides robust capabilities for understanding and processing complex textual inputs. The architecture of JEN-1 Composer is built upon a 1D UNet backbone (Ronneberger, Fischer, and Brox 2015), with critical modifications to the Jen-1 model (Li et al. 2024), as outlined in Section 4.1. These key adjustments include the channel-wise concatenation of tracks and the expansion of the single-track timestep into a vector of four elements, enabling the model to effectively handle and generate cohesive multi-track music.

Training follows a progressive curriculum strategy, as detailed in Section 4.3. Initially, the probability for single-track generation tasks is set to  $1/K$ , with  $K = 4$ , where each track is independently considered as a target track. As training progresses, these probabilities gradually decay, allowing for the introduction of more complex multi-track generation tasks. Eventually, all task types are covered, with the probability for each generation scenario (whether a track is a target track or not) set to  $1/(2^K - 1)$ . Sampler settings for conditional and marginal generation are optimized with  $p_1 = 0.8$ . After 300 epochs, self-bootstrapping training is introduced with a probability  $p_2 = 0.5$ . We determined the optimal value for the guidance scale parameter  $\lambda = 7$  through a grid search. Training was conducted on two NVIDIA A100 GPUs, with hyperparameters including the AdamW optimizer (Loshchilov and Hutter 2018), a linear decay learning rate starting at  $3e-5$ , a batch size of 12 per GPU, and optimization settings of  $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ , weight decay of 0.1, and a gradient clipping threshold of 0.7.

## 5.2 Comparison with State-of-the-art Methods

We conduct a comprehensive comparison with leading text-to-music generation models, including MusicLM (Agostinelli et al. 2023), MusicGen (Copet et al. 2024), and Jen-1 (Li et al. 2024), which primarily focus on single-track generation. In contrast, JEN-1 Composer supports multi-track generation with flexible conditional control, enabling track-wise generation and enhanced inter-track alignment. As shown in Table 1, JEN-1 Composer achieves superior CLAP scores across individual tracks and the mixed track, highlighting its ability to generate music that closely adheres to text prompts while maintaining inter-track coherence. Additionally, the RPR metric results confirm a strong user preference for JEN-1 Composer’s outputs, demonstrating its effectiveness in delivering high-quality compositions.

To evaluate multi-track generation quality, we performed a zero-shot comparison on the Slakh2100 dataset (Manilow et al. 2019), without any fine-tuning. As shown in Table 3, JEN-1 Composer achieves competitive FAD scores compared to StemGen (Parker et al. 2024) and MSDM (Mariani et al. 2024). Notably, our method, which infers directly on the Slakh2100 test set, achieves a FAD of 4.04, outperforming StemGen (4.30) and significantly surpassing MSDM (6.55). This indicates the strong generalization capability of our model, which can be attributed to the high-quality proprietary dataset used for training. Moreover, JEN-1 Com-

METHODS	CLAP↑					RPR↑
	BASS	DRUMS	INSTRUMENT	MELODY	MIXED	MIXED
MusicLM	0.16	0.17	0.23	0.28	0.28	27%
MusicGen	0.17	0.15	0.25	0.33	0.35	36%
Jen-1	0.19	0.16	0.29	0.32	0.36	40%
<b>JEN-1 Composer</b>	<b>0.21*</b>	<b>0.18*</b>	<b>0.29</b>	<b>0.36**</b>	<b>0.39**</b>	–

Table 1: Comparison of mixed-track text-to-music generation. \*\* and \* represent significance level  $p$ -value  $< 0.01$  and  $p$ -value  $< 0.05$  of comparing JEN-1 Composer with Jen-1. Wherever possible, we use open-source models, and for MusicLM, we employ the publicly available API.

METHODS	CLAP↑					RPR↑
	BASS	DRUMS	INSTRUMENT	MELODY	MIXED	MIXED
baseline	0.20	0.18	0.20	0.28	0.28	16%
+ individual timestep vector	0.19	0.18	0.22	0.32	0.33	20%
+ curriculum training strategy	0.21	0.17	0.26	0.35	0.37	35%
+ interactive inference	<b>0.21</b>	<b>0.18</b>	<b>0.29</b>	<b>0.36</b>	<b>0.39</b>	–

Table 2: Ablation studies. Starting from the baseline configuration, we incrementally modify the model to investigate the impact of each component. The baseline model involves only minimal modifications to the input and output channels of Jen-1 and employs a unified timestep for joint generation training.

METHODS	DATASETS	FAD↓
StemGen	Slakh	4.30
MSDM	Slakh	6.55
JEN-1 Composer	Slakh	4.04
JEN-1 Composer	Proprietary	3.76

Table 3: Comparison of multi-track generation quality. Results for StemGen and MSDM are taken directly from their original papers.

poser achieves an even lower FAD of 3.76 on our proprietary dataset. This result is reasonable given that our dataset comprises high-fidelity studio recordings, as opposed to MIDI-rendered audio typically used in other studies. These findings emphasize the importance of high-quality training data and further demonstrate the robustness and audio fidelity of JEN-1 Composer.

### 5.3 Ablation Studies

Our ablation studies underscore the significance of each component within the JEN-1 Composer framework. As summarized in Table 2, we started with a baseline model featuring a four-track input-output configuration inspired by Jen-1 (Li et al. 2024), and incrementally introduced our proposed enhancements. Using individual timestep vectors for each track was pivotal in effectively modeling both marginal and conditional distributions, resulting in notably higher CLAP scores for individual tracks.

The progressive curriculum training strategy, which transitions from simpler conditional generation to complex joint modeling tasks, further improved performance, particularly for intricate tracks like melody and instrument. Moreover,

integrating an interactive Human-AI co-composition workflow yielded the highest mixing quality. This design allowed the model to alternate flexibly between generation modes and incorporate multiple human inputs as additional conditional signals. For example, the model could initially generate melody and instrument tracks, then leverage these as guidance to condition the subsequent generation of drums and bass, ensuring greater coherence across tracks.

In summary, the careful orchestration of these components enables JEN-1 Composer to achieve exceptional quality and flexibility in multi-track music synthesis.

## 6 Conclusion

In this study, we introduce JEN-1 Composer, a text-to-music framework for multi-track audio generation that enhances user interaction through intuitive, direct audio input. By moving beyond the limitations of symbolic and spectrum-based methods, JEN-1 Composer simplifies the creative process while preserving rich audio fidelity. The model not only leverages text prompts as conditional control signals but also integrates multi-track coherence as an additional guiding signal to ensure harmonious and well-aligned generation. Through curriculum training, it progressively learns to model marginal, joint, and conditional distributions, enabling the generation of both simple and complex multi-track compositions.

Despite these advancements, achieving professional music standards remains a challenge. Future improvements will require deeper integration of engineering, design, and the arts, alongside enhanced data collection and annotation processes, to further refine the model’s control, coherence, and output quality.

## Ethical Statement

In developing JEN-1 Composer, we strictly adhered to ethical standards, using only authorized data and complying fully with copyright and licensing agreements. JEN-1 Composer is designed to enhance, not replace, human creativity, promoting Human-AI collaboration to push artistic boundaries. We are committed to the ethical use of AI, opposing any misuse for unauthorized reproduction or plagiarism. Our approach prioritizes creators' rights and encourages responsible AI integration in the creative process. This manuscript was polished and proofread with the assistance of ChatGPT.

## 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*.
- Bao, F.; Nie, S.; Xue, K.; Li, C.; Pu, S.; Wang, Y.; Yue, G.; Cao, Y.; Su, H.; and Zhu, J. 2023. One transformer fits all distributions in multi-modal diffusion at scale. In *International Conference on Machine Learning*, 1692–1717. PMLR.
- Borsos, Z.; Marinier, R.; Vincent, D.; Kharitonov, E.; Pietquin, O.; Sharifi, M.; Roblek, D.; Teboul, O.; Grangier, D.; Tagliasacchi, M.; et al. 2023. Audioldm: a language modeling approach to audio generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Chen, B.; Li, P.; Yao, Y.; and Wang, A. 2024. JEN-1 DreamStyler: Customized Musical Concept Learning via Pivotal Parameters Tuning. *arXiv preprint arXiv:2406.12292*.
- Chung, H. W.; Hou, L.; Longpre, S.; Zoph, B.; Tay, Y.; Fedus, W.; Li, Y.; Wang, X.; Dehghani, M.; Brahma, S.; et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70): 1–53.
- Copet, J.; Kreuk, F.; Gat, I.; Remez, T.; Kant, D.; Synnaeve, G.; Adi, Y.; and Défossez, A. 2024. Simple and controllable music generation. *Advances in Neural Information Processing Systems*, 36.
- Creswell, A.; White, T.; Dumoulin, V.; Arulkumaran, K.; Sengupta, B.; and Bharath, A. A. 2018. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1): 53–65.
- Défossez, A. 2021. Hybrid Spectrogram and Waveform Source Separation. In *Proceedings of the ISMIR 2021 Workshop on Music Source Separation*.
- Défossez, A.; Copet, J.; Synnaeve, G.; and Adi, Y. 2022. High fidelity neural audio compression. *arXiv preprint arXiv:2210.13438*.
- Dong, H.-W.; Hsiao, W.-Y.; Yang, L.-C.; and Yang, Y.-H. 2018. Musegan: Multi-track sequential generative adversarial networks for symbolic music generation and accompaniment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Elizalde, B.; Deshmukh, S.; Al Ismail, M.; and Wang, H. 2023. Clap learning audio concepts from natural language supervision. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Engel, J.; Resnick, C.; Roberts, A.; Dieleman, S.; Norouzi, M.; Eck, D.; and Simonyan, K. 2017. Neural audio synthesis of musical notes with wavenet autoencoders. In *International Conference on Machine Learning*, 1068–1077. PMLR.
- Ens, J.; and Pasquier, P. 2020. Mmm: Exploring conditional multi-track music generation with the transformer. *arXiv preprint arXiv:2008.06048*.
- Evans, Z.; Carr, C.; Taylor, J.; Hawley, S. H.; and Pons, J. 2024. Fast timing-conditioned latent audio diffusion. *arXiv preprint arXiv:2402.04825*.
- Gărbacea, C.; van den Oord, A.; Li, Y.; Lim, F. S.; Luebs, A.; Vinyals, O.; and Walters, T. C. 2019. Low bit-rate speech coding with VQ-VAE and a WaveNet decoder. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 735–739. IEEE.
- Hershey, S.; Chaudhuri, S.; Ellis, D. P.; Gemmeke, J. F.; Jansen, A.; Moore, R. C.; Plakal, M.; Platt, D.; Saurous, R. A.; Seybold, B.; et al. 2017. CNN architectures for large-scale audio classification. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 131–135. IEEE.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Huang, C.-Z. A.; Vaswani, A.; Uszkoreit, J.; Shazeer, N.; Simon, I.; Hawthorne, C.; Dai, A. M.; Hoffman, M. D.; Dinulescu, M.; and Eck, D. 2018. Music transformer. *arXiv preprint arXiv:1809.04281*.
- Huang, Q.; Park, D. S.; Wang, T.; Denk, T. I.; Ly, A.; Chen, N.; Zhang, Z.; Zhang, Z.; Yu, J.; Frank, C.; et al. 2023a. Noise2music: Text-conditioned music generation with diffusion models. *arXiv preprint arXiv:2302.03917*.
- Huang, R.; Huang, J.; Yang, D.; Ren, Y.; Liu, L.; Li, M.; Ye, Z.; Liu, J.; Yin, X.; and Zhao, Z. 2023b. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. In *International Conference on Machine Learning*, 13916–13932. PMLR.
- Jin, C.; Wang, T.; Li, X.; Tie, C. J. J.; Tie, Y.; Liu, S.; Yan, M.; Li, Y.; Wang, J.; and Huang, S. 2022. A transformer generative adversarial network for multi-track music generation. *CAAI Transactions on Intelligence Technology*, 7(3): 369–380.
- Jin, C.; Wang, T.; Liu, S.; Tie, Y.; Li, J.; Li, X.; and Lui, S. 2020. A transformer-based model for multi-track music generation. *International Journal of Multimedia Data Engineering and Management (IJMDEM)*, 11(3): 36–54.
- Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; and Catanzaro, B. 2020. DiffWave: A Versatile Diffusion Model for Audio Synthesis. In *International Conference on Learning Representations*.

- Kreuk, F.; Synnaeve, G.; Polyak, A.; Singer, U.; Défossez, A.; Copet, J.; Parikh, D.; Taigman, Y.; and Adi, Y. 2022. AudioGen: Textually Guided Audio Generation. In *The Eleventh International Conference on Learning Representations*.
- Lam, M. W.; Tian, Q.; Li, T.; Yin, Z.; Feng, S.; Tu, M.; Ji, Y.; Xia, R.; Ma, M.; Song, X.; et al. 2024. Efficient neural music generation. *Advances in Neural Information Processing Systems*, 36.
- Li, P. P.; Chen, B.; Yao, Y.; Wang, Y.; Wang, A.; and Wang, A. 2024. JEN-1: Text-Guided Universal Music Generation with Omnidirectional Diffusion Models. In *2024 IEEE Conference on Artificial Intelligence (CAI)*, 762–769.
- Liang, X.; Wu, J.; and Cao, J. 2019. MIDI-Sandwich2: RNN-based Hierarchical Multi-modal Fusion Generation VAE networks for multi-track symbolic music generation. *arXiv preprint arXiv:1909.03522*.
- 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. In *International Conference on Machine Learning*, 21450–21474. PMLR.
- Loshchilov, I.; and Hutter, F. 2018. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- Manilow, E.; Wichern, G.; Seetharaman, P.; and Le Roux, J. 2019. Cutting music source separation some Slakh: A dataset to study the impact of training data quality and quantity. In *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 45–49. IEEE.
- Marafioti, A.; Perraudin, N.; Holighaus, N.; and Majdak, P. 2019. A context encoder for audio inpainting. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(12): 2362–2372.
- Mariani, G.; Tallini, I.; Postolache, E.; Mancusi, M.; Cosmo, L.; and Rodolà, E. 2024. Multi-Source Diffusion Models for Simultaneous Music Generation and Separation. In *The Twelfth International Conference on Learning Representations*.
- Muhamed, A.; Li, L.; Shi, X.; Yaddanapudi, S.; Chi, W.; Jackson, D.; Suresh, R.; Lipton, Z. C.; and Smola, A. J. 2021. Symbolic music generation with transformer-gans. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 408–417.
- Nistal, J.; Pasini, M.; Aouameur, C.; Grachten, M.; and Lattner, S. 2024. Diff-A-Riff: Musical Accompaniment Co-creation via Latent Diffusion Models. *arXiv preprint arXiv:2406.08384*.
- Parker, J. D.; Spijkervet, J.; Kosta, K.; Yesiler, F.; Kuznetsov, B.; Wang, J.-C.; Avent, M.; Chen, J.; and Le, D. 2024. StemGen: A music generation model that listens. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1116–1120. IEEE.
- Payne, C. 2019. MuseNet, 2019. URL <https://openai.com/blog/musenet>.
- Roblek, D.; Kilgour, K.; Sharifi, M.; and Zuluaga, M. 2019. Fréchet Audio Distance: A Reference-free Metric for Evaluating Music Enhancement Algorithms. In *Proc. Interspeech*, 2350–2354.
- 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.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III* 18, 234–241. Springer.
- Rouard, S.; Massa, F.; and Défossez, A. 2023. Hybrid transformers for music source separation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Yu, Y.; Srivastava, A.; and Canales, S. 2021. Conditional lstm-gan for melody generation from lyrics. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 17(1): 1–20.
- Zeghidour, N.; Luebs, A.; Omran, A.; Skoglund, J.; and Tagliasacchi, M. 2021. Soundstream: An end-to-end neural audio codec. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30: 495–507.
- Zhu, H.; Liu, Q.; Yuan, N. J.; Zhang, K.; Zhou, G.; and Chen, E. 2020. Pop music generation: From melody to multi-style arrangement. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 14(5): 1–31.