

# SMITIN: Self-Monitored Inference-Time INtervention for Generative Music Transformers

Koo, Junghyun; Wichern, Gordon; Germain, François G; Khurana, Sameer; Le Roux, Jonathan

TR2025-012 January 23, 2025

## Abstract

We introduce Self-Monitored Inference-Time INtervention (SMITIN), an approach for controlling an autoregressive generative music transformer using classifier probes. These simple logistic regression probes are trained on the output of each attention head in the transformer using a small dataset of audio examples both exhibiting and missing a specific musical trait (e.g., the presence/absence of drums, or real/synthetic music). We then steer the attention heads in the probe direction, ensuring the generative model output captures the desired musical trait. Additionally, we monitor the probe output to avoid adding an excessive amount of intervention into the autoregressive generation, which could lead to temporally incoherent music. We validate our results objectively and subjectively for both audio continuation and text-to-music applications, demonstrating the ability to add controls to large generative models for which retraining or even fine-tuning is impractical for most musicians. Audio samples of the proposed intervention approach are available on our demo page

*IEEE Open Journal of Signal Processing 2025*



---

# SMITIN: Self-Monitored Inference-Time Intervention for Generative Music Transformers

Junghyun Koo<sup>1,2</sup>, Member, IEEE, Gordon Wichern<sup>1</sup>, Member, IEEE,  
François G. Germain<sup>1</sup>, Member, IEEE, Sameer Khurana<sup>1</sup>, Member, IEEE,  
and Jonathan Le Roux<sup>1</sup>, Fellow, IEEE

<sup>1</sup>Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139 USA

<sup>2</sup>Seoul National University, Seoul, South Korea

This work was performed while J. Koo was an intern at MERL.

---

**ABSTRACT** We introduce Self-Monitored Inference-Time Intervention (SMITIN), an approach for controlling an autoregressive generative music transformer using classifier probes. These simple logistic regression probes are trained on the output of each attention head in the transformer using a small dataset of audio examples both exhibiting and missing a specific musical trait (e.g., the presence/absence of drums, or real/synthetic music). We then steer the attention heads in the probe direction, ensuring the generative model output captures the desired musical trait. Additionally, we monitor the probe output to avoid adding an excessive amount of intervention into the autoregressive generation, which could lead to temporally incoherent music. We validate our results objectively and subjectively for both audio continuation and text-to-music applications, demonstrating the ability to add controls to large generative models for which retraining or even fine-tuning is impractical for most musicians.

Audio samples of the proposed intervention approach are available on our [demo page](#).

**INDEX TERMS** Classifier probes, music transformer, multi-head attention, music information retrieval.

---

## I. Introduction

Dynamic control of an audio signal can be considered a fundamental job of most audio professionals. Musicians must precisely control their instruments, and recording studio engineers manipulate various controls to achieve a desired result. Recent advances in the generative modeling of text and images have also been applied to audio signals, leading to an emerging body of literature on how to condition and control audio generative models [1]–[7].

Among possible conditioning inputs, describing the audio a user wishes to generate in terms of natural language arguably provides the most flexibility. Such text-to-music models, however, lack the fine-grained temporal control often desired by audio professionals. This has led to training text-to-music models that allow additional conditioning inputs, such as melody sequences [2], or other types of sequential input [8]–[11], but these models still suffer from the fact that, once trained, the type of conditioning input is fixed. To alleviate this limitation, some works have explored the addition of new types of control inputs without requiring retraining [4], [12]–[16]. The standard approach in most previous work trains supplementary adapters (e.g., LoRA [17]) atop the foundational generative model to introduce additional control, but once the adaptation is complete, it is often not straightforward to vary the strength of the control at inference time. Our goal is to build “knob-like” variable strength

controls given only a small number of audio examples both with and without a desired musical trait (i.e., binary labels).

Audio “language models” typically work by first tokenizing chunks of an audio waveform, for example by passing them through the encoder of a pre-trained residual vector quantized autoencoder [18]–[20], and then pass these audio tokens through an autoregressive transformer trained to predict the next token. Finally, a decoder converts the generated token back into audio. While not state-of-the-art in terms of generation quality compared to latent diffusion [21], [22] or masked token models [5], [23], autoregressive audio generative models are still useful for many real-time and interactive applications. Furthermore, because this class of text-to-music models shares many architectural characteristics with large language models (LLMs), it is interesting to explore whether techniques developed to provide inference-time control of LLMs [24]–[26] may also be effective for audio. In particular, we take inspiration from recent work attempting to make text language models more truthful by taking advantage of the fact that the learned internal representations of these models are able to represent the concept of truthfulness [27]. Classifier probes [28]–[30] applied to each attention head are used to determine whether or not that attention head has learned to represent truthfulness, and inference-time intervention (ITI) is then used to surgically modify the outputs of only the most truth-correlated attention heads, resulting in an LLM

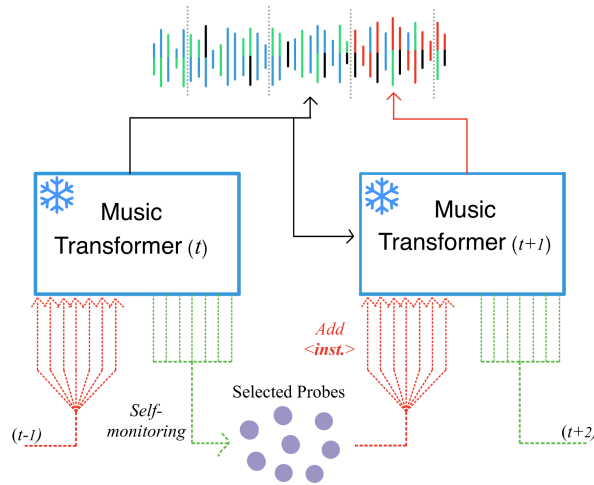


FIGURE 1: Overall pipeline of SMITIN for inference-time intervention on a pre-trained music generative transformer. The process attempts to enforce specific musical factors (e.g., presence of a particular instrument) during the generation process. SMITIN utilizes a self-monitoring technique to dynamically adjust the intervention strength at each generation step, enabling precise control over the inclusion of the target characteristic while preserving the musical integrity of the output.

that hallucinates less. Previous work [31] has shown that classifier probes trained on the internal representations of generative music transformers can lead to strong performance across a wide variety of music classification tasks. However, these probes were trained using the learned representation from entire layers, not individual attention heads.

We first show using classifier probes that individual self-attention heads in a pre-trained autoregressive music transformer, MusicGen [2], have indeed learned to represent aspects of music we may wish to build control knobs for, making it a priori suitable for ITI approaches such as in [27]. However, we find that direct application of this particular approach to a music transformer is sub-optimal because of the difficulty in generating long temporally coherent music samples, as too much intervention causes the generated music to quickly become incoherent. Furthermore, the method asks for the empirical tuning of the number of heads selected for intervention, impeding the ability to implement multiple intervention kinds at scale. To remedy this, we propose the inclusion of a self-monitoring process into the intervention operation, such that we only apply the intervention when the learned classifier probes tell us it is necessary based on the state of the generation network. Crucially, this self-monitoring technique enables real-time assessment of whether the current generated sample incorporates the target factor, allowing for the generation of musically aligned samples without a costly retraining or fine-tuning process. We also propose to weigh all heads as a function of their learned probe performance, removing the need for empirical head selection tuning. The overall idea of the self-monitoring process is outlined in Fig. 1.

The remainder of the paper is organized as follows. In Section II, we evaluate whether the internal representations from individual attention heads of the pre-trained MusicGen model are suitable for downstream music classification tasks by training simple linear classifier probes. Then, in Section III, we describe SMITIN, where these classifier probes are used to build custom “knobs” for inference-

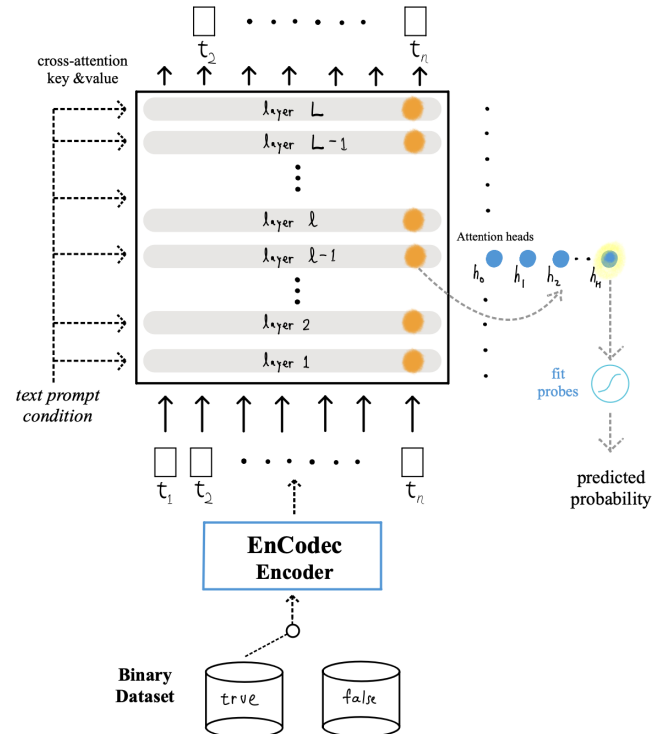


FIGURE 2: Overview of probing MusicGen. Each audio sample in a labeled dataset is converted to EnCodec tokens and input into MusicGen to predict the next token. The activations for the last time step (orange dots) for each attention head in each layer (blue dots) are used to train a logistic regression classifier (probe).

time control of MusicGen. In Section IV, we assess our proposed approach both objectively and subjectively in terms of intervention success and overall audio quality. Furthermore, we study in Section V several important practical aspects of our approach, such as the number of audio examples required to obtain classifier probes sufficient for intervention, whether multiple interventions can be used simultaneously, and the ability of the intervention to prevent the autoregressive generative process from diverging over time into unrealistic and temporally incoherent music. In the Appendix, we additionally provide probing experiments on more downstream tasks and additional details of our subjective and objective results. Our source code will be available online<sup>1</sup>.

## II. Understanding MusicGen

In this section, we seek to investigate and quantify the comprehension of music by each attention head within MusicGen [2]. MusicGen is a pre-trained and publicly accessible generative music transformer that uses an EnCodec encoder [19] to create discrete audio tokens, an autoregressive transformer to predict the next token, and an EnCodec decoder to output an audio signal. This analysis will provide insights into the model’s potential for fine-grained control via attention-head steering. We use a probing task designed to assess the model’s ability to distinguish musical pieces based on the presence or absence of specific instruments, and explore further downstream tasks in the Appendix.

<sup>1</sup><https://github.com/merlresearch/smitin>

## A. Autoregressive Transformer Models

Architectures such as MusicGen are characterized by the autoregressive generation of a sequence of audio frames by transformer models. They include a collection of  $L$  multi-head self-attention layers (residual connections, normalization layers, and fully-connected feed-forward layers are employed as usual and not described here). At current time step  $t$  in the generated sequence, the  $l$ -th self-attention layer computes  $H$  self-attention heads  $z_{l,h}(t) \in \mathbb{R}^D$  from an input vector  $x_l(t) \in \mathbb{R}^{DH}$  as

$$z_{l,h}(t) = \text{Att}(W_{l,h}^Q x_l(t), W_{l,h}^K x_l(1:t), W_{l,h}^V x_l(1:t)), \quad (1)$$

where  $x_l(1:t) = [x_l(1), \dots, x_l(t)]$ ,  $W_{l,h}^Q$ ,  $W_{l,h}^K$ , and  $W_{l,h}^V$  denote the head-specific query, key, and value projection matrices, all in  $\mathbb{R}^{D \times DH}$ , and  $\text{Att}$  denotes the attention operator [32]. The output vector  $y_l(t) \in \mathbb{R}^{DH}$  is obtained by projecting back each head into  $DH$ -dimensional space and summing:

$$y_l(t) = \sum_{h=1}^H W_{l,h}^O z_{l,h}(t), \quad (2)$$

where  $W_{l,h}^O \in \mathbb{R}^{DH \times D}$  is a projection matrix.

## B. Probing MusicGen

We describe the methodology of probing MusicGen by evaluating the capability of its self-attention heads in recognizing instruments (i.e., determining whether a target instrument is present in the audio stream). The overview of the probing procedure is shown as Fig. 2. We create a dataset by curating data from MUSDB [33] and MoisesDB [34], which offer multi-track recordings with isolated instrument stems. For a given target stem, we form a positive class of mixtures where the target stem is present, and a negative class of corresponding mixtures with the target stem removed as follows. First, we remove from every multi-track recording the time regions where the target stem is silent. Then, out of this pruned recording, the mixture of all of its stems is added to the positive class data, while the mixture of all of its stems except its target stem is added to the negative class data. Subsequently, we process 3-second-long segments of these tracks for training (and testing), passing them through MusicGen to extract the intermediate activation  $z_{l,h}(T)$  at the last time step  $T$  for every self-attention layer  $l$  and head  $h$ . This forms the basis for the training (and testing) sets of the probe classifier, wherein a simple logistic regression model is employed to distinguish the presence of the instrument.

The testing accuracy of probes from MusicGen<sub>large</sub> (3.3B parameters) across all self-attention layers  $l$  and heads  $h$  is illustrated in Fig. 3. We observe that specific subsets of heads outperform others in detecting the presence of each target stem. While certain attention layers show better performance, it is notable that not all heads within each layer result in uniform performance; rather, their effectiveness varies considerably. This variation underscores the utility of head-wise probing in achieving precise control over the transformer’s behavior. Furthermore, MusicGen’s proficiency varies across instruments; it demonstrates a strong understanding of drums and bass, whereas its accuracy on guitar and piano is comparatively lower. This discrepancy suggests a potential bias in MusicGen’s training dataset (see [2], Fig. A.3).

## III. Self-Monitored Inference-Time Intervention

### A. Inference-time Intervention (ITI)

Reference [27] first recognized that the output of the model could be somewhat controlled at inference by intervening in (i.e., modifying)




FIGURE 3: Instrument recognition performance of individual attention head probes from the MusicGen<sub>large</sub> model activations, sorted by accuracy, with all colorbars normalized to the same range. The values in brackets indicate the highest accuracy of the probe classifier for each respective instrument task, followed by the threshold value  $\tau$ , which is defined in Section III-D.

the computation of the output of the multi-head self-attention layer in Eq. (2). This intervention is done by adding a term to the heads  $z_{l,h}(t) \in \mathbb{R}^D$  before the projection:

$$y_l(t) = \sum_{h=1}^H W_{l,h}^O (z_{l,h}(t) + \alpha w_{l,h} \cdot \sigma_{l,h} \theta_{l,h}), \quad (3)$$

where  $\theta_{l,h} \in \mathbb{R}^D$  is a vector representing the head-specific ITI “direction” in the  $D$ -dimensional head-specific space,  $\sigma_{l,h} \in \mathbb{R}^+$  is a head-specific normalization of  $\theta_{l,h}$ ,  $\alpha \in \mathbb{R}^+$  is the system-wide ITI strength, and  $w_{l,h} \geq 0$  is a head-specific strength weighting. While the formulations are equivalent, we depart somewhat from [27] by introducing  $w_{l,h}$  as a separate term, whereas it was factored into  $\theta_{l,h}$  in their notations. This helps to disambiguate the dual function played by  $\theta_{l,h}$  in their formulation.

The head-specific ITI directions  $\theta_{l,h}$  are obtained through classifier probes. Following [27], we find an auxiliary probing dataset suited to train a classifier mirroring our ITI goal (e.g., drum classifier if the goal is to add drums). We then run the dataset through our (frozen) generation network and collect a set of heads  $z_{l,h}(t)$  for all  $l, h$ . For each  $l$  and  $h$ , we then train a distinct logistic regression classifier probe with parameters  $\hat{\theta}_{l,h}$ , whose prediction is obtained as  $\text{sigmoid}(\langle \hat{\theta}_{l,h}, z_{l,h}(t) \rangle)$ . Once probe training is complete, we set ITI direction  $\theta_{l,h}$  as the final  $\hat{\theta}_{l,h}$ , and  $\sigma_{l,h}$  as the standard deviation of  $\langle \hat{\theta}_{l,h}, z_{l,h}(t) \rangle$  for all  $z_{l,h}(t)$  obtained on the combined probing

training and testing data. We also take note of the final classifier accuracy  $\text{acc}_{l,h}$  on the probing test data. In the case of MusicGen, we apply the generative model in unconditional generation mode for probing, as our probing dataset generally lacks text queries for its audio samples.

In [27], head-specific weights  $w_{l,h}$  are set by finding the top- $K$  heads in terms of classifier probe accuracy  $\text{acc}_{l,h}$ , setting their  $w_{l,h}$  to 1 and the others to 0. The most effective  $\alpha$  and  $K$  are found by hyperparameter grid search.

### B. Sparse intervention

In music generation, the approach above may present some limitations. For example, for the task of audio continuation, we observe that ITI often leads to changes that are too abrupt to be musically plausible (see Section IV-B). As mitigation, we propose to diminish the ITI frequency across time steps, potentially allowing the generation process to better align with the underlying rhythmic structure of the generated music. This corresponds to replacing the ITI weights  $w_{l,h}$  in Eq. (3) by time-varying weights  $w_{l,h}(t)$ , which can only be non-zero for  $t = t_0 + is$ ,  $i \in \mathbb{N}$ , where  $t_0$  is an intervention start time and  $s$  represents the number of steps between each ITI (e.g.,  $s = 5$  to perform ITI every 5 time steps). The value of  $w_{l,h}(t)$  for  $t = t_0 + is$  can be set by another criterion, such as the original one based on top- $K$  heads.

### C. Soft-weighting

One crucial limitation of [27] is that  $K$ , the number of heads on which ITI is performed (i.e., the number of pairs  $(l, h)$  for which  $w_{l,h}(t) \neq 0$ ), is a hyperparameter that must be tuned. We additionally propose a hyperparameter-free soft-weighting approach based on the collected probe accuracies  $\text{acc}_{l,h}$ , and show that it is sufficient to perform effective ITI (see Appendix Table 15). In practice, we propose setting the weights as

$$w_{l,h}(t) = \left( \frac{\text{acc}_{l,h} - \text{acc}_{\min}}{\text{acc}_{\max} - \text{acc}_{\min}} \right)^c, \quad (4)$$

with  $\text{acc}_{\min}$  and  $\text{acc}_{\max}$  the minimum and maximum accuracies recorded across all  $l$  and  $h$ , and  $c$  a power factor allowing to modulate the relative weights of heads with accurate vs. inaccurate classifier probes (we use  $c = 3$ ). By construction,  $w_{l,h}(t)$  is guaranteed to fall between 0 and 1.

### D. Automated Intervention Modulation by Self-monitoring

Ideally, we would expect a system capable of modulating the ITI strength to be most effective, as systems with time-invariant  $w_{l,h}(t)$  make no use of (and, as such, cannot react to) the state of the inference model during generation.

As our core contribution, we propose to use the classifier probes to drive such a modulation. We first define as  $\mathcal{H}_K$  the set of top- $K$  heads  $(l, h)$  by probe classifier accuracy. For each new generation time step  $t$ , we collect the set  $\mathcal{C}(t)$  of output predictions of the classifier probes for all heads in  $\mathcal{H}_K$  before intervening on them, i.e.,

$$\mathcal{C}(t) = \{\text{sigmoid}(\langle \theta_{l,h}, z_{l,h}(t) \rangle) \mid (l, h) \in \mathcal{H}_K\}. \quad (5)$$

The statistics of  $\mathcal{C}(t)$  reflect the confidence of the probes regarding the success of the intervention. By comparing them to the baseline accuracies obtained on the probing training data, we can devise a modulation scheme to update  $w_{l,h}(t)$ . We compute the median  $\bar{\mathcal{C}}(t) = \text{med}(\mathcal{C}(t))$  and the change in this median since the prior generation step with intervention,  $\Delta(t) = \bar{\mathcal{C}}(t) - \bar{\mathcal{C}}(t-s)$ . We then

define a threshold based on the median and standard deviation of the set  $\mathcal{A}$  of probe accuracies  $\text{acc}_{l,h}$  on the probing training data over the same heads  $\mathcal{H}_K$ ,

$$\mathcal{A} = \{\text{acc}_{l,h} \mid (l, h) \in \mathcal{H}_K\}, \quad (6)$$

setting the threshold as  $\tau = \text{med}(\mathcal{A}) - \text{std}(\mathcal{A})$ .

The update algorithm is then defined following:

- For the first generation time step with ITI  $t_0$ ,  $w_{l,h}(t_0)$  is set following Eq. (4), and we set  $\Delta(t_0) \leftarrow 0$ ,
- For each subsequent generation time step with ITI  $t+s$ , we have 3 cases:

- 1) if  $\bar{\mathcal{C}}(t) < \tau$  and  $w_{l,h}(t) > 0$ , we set

$$w_{l,h}(t+s) \leftarrow w_{l,h}(t) \cdot (1 - \Delta(t)), \quad (7)$$

- 2) if  $\bar{\mathcal{C}}(t) < \tau$  and  $w_{l,h}(t) = 0$ , we reset to the initial value

$$w_{l,h}(t+s) \leftarrow w_{l,h}(t_0), \quad (8)$$

- 3) if  $\bar{\mathcal{C}}(t) \geq \tau$ , we set  $w_{l,h}(t+s) \leftarrow 0$ .

## IV. Controlled Experiments on Instrument Addition

Here, we validate the effectiveness of our proposed contribution by presenting a thorough investigation of ITI applied to the music generative task of instrument addition, a task for which evaluation is relatively straightforward. This section outlines methods to quantify the effectiveness of ITI techniques objectively, complemented by a subsequent subjective evaluation.

### A. Setup

For the evaluation of ITI on MusicGen, we utilize the *large mono* configuration, which is composed of 48 layers with 32 attention heads each, amounting to a total of 1,536 heads. Unless otherwise noted, the self-monitoring configuration of SMITIN is set to an intervention strength of  $\alpha = 5.0$ , a sparse intervention rate of  $s = 5$ , and  $K = 16$  for the selection of monitoring probes.

For the instrument addition task, we curate probing datasets for drums, bass, guitar, and piano detection from MUSDB [33] and MoisesDB [34], both of which provide multi-track recordings with isolated instrument stems. Stems from their training (resp. testing) partitions are used to form the probing training (resp. testing) partition. To create positive and negative classes, i.e., one including and one excluding the target instrument, we proceed as follows. We collect a first set consisting of the mixed audio from each track recording after removing the target instrument stem, and trim silent sections. We then collect a second set consisting of the target instrument stems, and trim silent sections. We then generate negative samples using audio segments randomly sampled from the first set, and positive samples by randomly mixing a mixture without the target instrument (from the first set) with a stem of the target instrument (from the second set). Following silence trimming, the datasets amount to 8.3, 8.0, 13.1, and 5.7 hours for drums, bass, guitar, and piano, respectively. These durations represent the total amount of paired data available for probing. Since the negative samples are more abundant, we balanced the dataset to ensure an equal number of positive and negative samples, based on the total duration of positive samples.

#### 1) Generation Approach

Our experimental setup includes two contexts in which to perform ITI: audio continuation and text-to-music. In both cases, the objective is to perform the task while adding the target instruments (drums, bass, guitar, and piano) into the music pieces through ITI.

*Audio continuation*: Given an input music signal, we continue to generate a music piece while trying to add the target instrument. The input segment is a 3-second-long music segment obtained from the MUSDB and MoisesDB test datasets, which does not contain the target instrument.

*Text-to-music*: This task uses text prompts from the MusicCaps [35] dataset to guide the generation process, allowing us to evaluate the effectiveness of ITI in the context of text-conditioned generation. We use the aspect list (e.g., “*pop, tinny wide hi hats, mellow piano melody*”) over the free text caption to further investigate the impact of adding simple instructions, e.g., “add inst.” to the tag list.

## 2) Objective Metrics

Our objective evaluation comprises three metrics designed to assess the effectiveness of the ITI process. We generate 1,000 music pieces, each 30 seconds long, for each experimental configuration.

**Success Rate**: We introduce Success Rate to quantify the efficacy of ITI in successfully adding the desired instrument to the music. This metric is derived from our self-monitoring technique from Section III-D, which measures the likelihood of the target musical factor based on the timewise probability inferred by the top- $K$  best-performing probes, where  $K = 16$ . For each audio, the success rate is calculated as  $n_{\text{success}}/N$ , where  $n_{\text{success}}$  is the number of samples  $t$  after the start of the intervention for which  $\bar{C}(t) > \tau$ , and  $N$  is the total number of time steps generated after the intervention. The reliability of this metric is contingent upon the top- $K$  probes’ test performance in the probing task.

Although the probing technique performs well on real-world audio samples, as demonstrated in Section II, there is some uncertainty regarding its effectiveness on generated samples. Also, because the Success Rate metric uses the same internal monitoring probes of MusicGen as the proposed intervention techniques, there is some concern that it may overestimate success in a way other instrument recognition models would not. Thus, to further validate the Success Rate, we compared its results on all the generated samples used for subjective evaluation in Section IV-D against decisions from other instrument recognition models. Specifically, we measured the Spearman correlation between the logits produced by our internal probes and those produced by external models such as YAMNet [36], PaSST [37], CLAP [38], and ConvNeXt [39]. YAMNet and PaSST were both trained on AudioSet [40], with PaSST further fine-tuned using the OpenMIC dataset [41]. For CLAP-based recognition, we trained a two-layer MLP to classify the embeddings generated by the frozen CLAP encoder, training on the same dataset as our internal probes. ConvNeXt was trained and tested using the same dataset as well (see Appendix A1 for detailed setup). We used publicly-available weights for YAMNet<sup>2</sup>, PaSST<sup>3</sup>, and CLAP<sup>4</sup>.

The results in Table 1 demonstrate a strong positive correlation, confirming that our internal probes align with external models in recognizing the desired musical traits. Furthermore, as we show in the analysis of subjective evaluation in Section IV-D, all the instrument recognition methods also correlate well with human decisions, providing additional validation of the models’ alignment with perceptual recognition.

A key advantage of using the Success Rate metric with MusicGen’s internal probes is that it allows for frame-wise recognition during generation, offering precise detection of musical traits

TABLE 1: Spearman correlation between the Success Rate and other instrument recognition methods. All p-values are below 1e-5.

Method	Correlation Coefficient $\rho$ ( $\uparrow$ )				
	drums	bass	guitar	piano	avg.
YAMNet [36]	0.71	0.50	0.58	0.33	0.53
PaSST [37]	0.69	0.73	0.63	0.40	0.63
CLAP [38]	0.82	0.74	0.68	0.66	0.75
ConvNeXt [39]	0.73	0.84	0.70	0.74	0.75

throughout the entire sample. For example, consider a 10-second music sample where drums are only present in the last 3 seconds. External instrument recognition models, which often operate with larger window sizes, may struggle to accurately identify that drums are only present in 30% of the sample. In contrast, our internal probes can precisely distinguish such temporal details, identifying the proportion of the sample that contains the target instrument. This level of granularity is particularly important in controlled generation tasks, where it is critical to accurately capture the presence and distribution of musical traits in the generated audio.

**Fréchet Audio Distance (FAD)**: FAD [42] is used to compare generated music with real music datasets, which helps ensure that the intervention does not cause an unrealistic shift in MusicGen’s output distribution. Instead of the conventional approach of using VGGish [43] as the deep encoder, we use L-CLAP mus [38] for dimension reduction, as it has been shown to offer a more accurate representation of music [44]. As the real music datasets, we adopt MUSDB (drums/bass) or MoisesDB (guitar/piano) for audio continuation and MusicCaps for text-to-music generation.

**Similarity Measurement in audio continuation**: For the audio continuation task, we measure the musical similarity between input and generated samples. We employ the Music Effects Encoder [45], a model trained with self-supervised contrastive learning designed to encode song identity. The encoder is used to embed both input and generated samples, and their cosine similarity is calculated. This metric assesses how well the identity of the input music is preserved in the generated output.

## B. Qualitative Evaluation

We qualitatively analyze the performance of top- $K$  probes’ monitoring capability through Fig. 4. The red dashed line in each plot indicates the threshold value  $\tau$ , as mentioned in Section III-D. From the left side of the figure, it is apparent that the median prediction (yellow line) more closely aligns with the actual presence of the target instrument than the mean prediction (green line). Consequently, we use  $\bar{C}(t) = \text{med}(C(t))$  to monitor during inference time.

The sub-plots on the right illustrate how the hyper-parameter  $s$  influences the outcome of the audio continuation task. We observe a relationship between the intervention frequency  $s$  and the rapidity with which the model aligns toward the target factor. This suggests that more frequent interventions cause the transformer to prioritize the target factor more heavily in subsequent token generation, potentially at the expense of referencing the similarity to previous musical context. We discuss this trade-off between intervention quantity and musical coherence throughout this section.

## C. Objective Evaluation

Utilizing the metrics introduced in Section IV-A2), we objectively evaluate our methods on *audio continuation* and *text-to-music*. We

<sup>2</sup><https://github.com/tensorflow/models/tree/master/research/audioset/yamnet>

<sup>3</sup><https://github.com/kkoutini/PaSST>

<sup>4</sup><https://github.com/LAION-AI/CLAP>

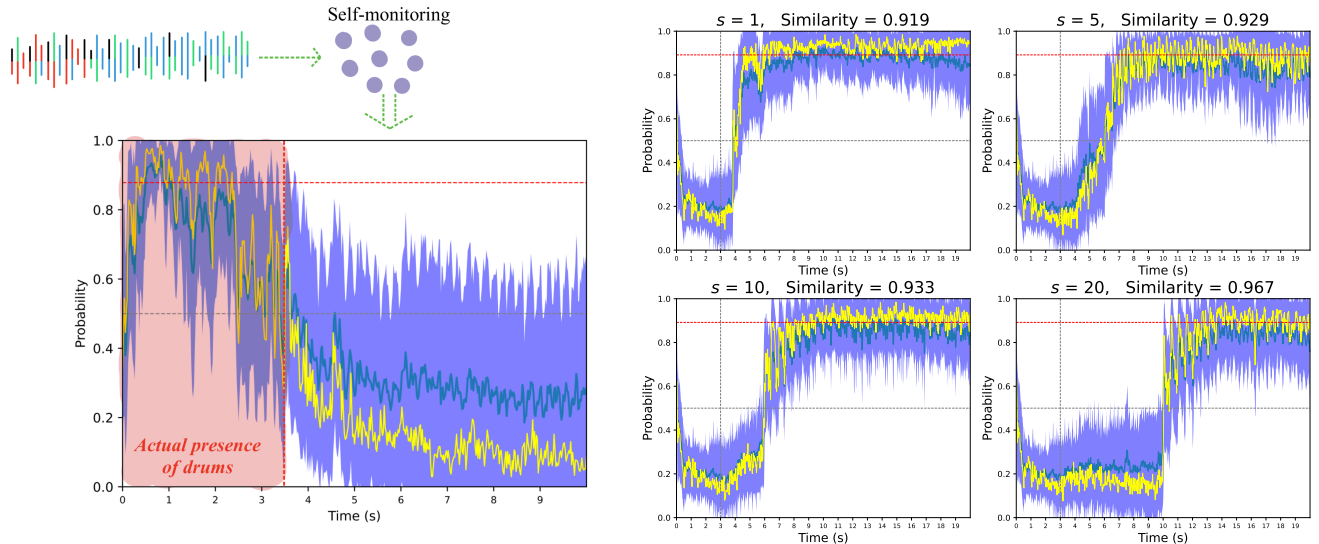


FIGURE 4: Inferred prediction of the top-K probes’ monitored decision along the time axis. The yellow line, green line, and the shaded blue region denote the median, mean, and standard deviation of inferred outputs by the probes, respectively. The red dashed line indicates threshold value  $\tau$  of the current monitoring probes. (Left) Monitored result on a real-world music sample. The high prediction (close to 1.0) until 3.5 seconds reflects the actual presence of drums, which aligns with the audio sample where drums are present only up to that point. (Right) The four sub-plots display the results of audio continuation on the same input music with varying ITI frequencies ( $s = [1, 5, 10, 20]$ ). These illustrate that more frequent intervention leads to a swifter convergence towards the target musical factor, at the expense of losing musical consistency with the input music.

TABLE 2: Objective evaluation on *audio continuation*.  $FAD_{L-CLAP_{mus}}$  is computed with MUSDB (drums/bass) and MoisesDB (guitar/piano).

Method		Success Rate [%] ( $\uparrow$ )					$FAD_{L-CLAP_{mus}}$ ( $\downarrow$ )					Similarity ( $\uparrow$ )				
ITI	Configuration	drums	bass	guitar	piano	avg.	drums	bass	guitar	piano	avg.	drums	bass	guitar	piano	avg.
$\times$	<b>unconditioned</b>	12.5	13.5	3.6	0.0	7.4	0.326	0.253	0.377	0.320	0.319	0.864	0.933	0.929	0.941	0.916
	<b>“add &lt;inst.&gt;”</b>	16.6	27.4	5.5	1.7	12.8	0.364	0.301	0.385	0.351	0.350	0.833	0.905	0.929	0.936	0.900
$\checkmark$	<b>original ITI</b>	21.3	76.0	17.2	85.2	49.9	0.358	0.510	0.447	0.368	0.420	0.868	0.858	0.919	0.939	0.896
	<b>weight decay ITI</b>	13.2	17.5	15.4	3.1	12.3	0.336	0.259	0.378	0.337	0.327	0.871	0.927	0.933	0.941	0.918
	<b>SMITIN</b>	20.6	30.4	33.8	8.7	23.3	0.346	0.267	0.397	0.337	0.336	0.857	0.922	0.935	0.938	0.913
	<b>SMITIN + “add &lt;inst.&gt;”</b>	16.9	28.0	16.9	17.4	19.8	0.381	0.303	0.391	0.374	0.362	0.839	0.909	0.929	0.942	0.904

compare with baseline methods that do not incorporate ITI, such as MusicGen’s standard unconditional and text-prompt conditioned generations. Additionally, we benchmark against the original ITI method, which applies a uniform intervention at every time step (i.e.,  $s = 1$ ) with a configuration of  $\alpha = 5.0$  and  $K = 64$  to apply the intervention with the equal amount  $\alpha$  to each of the  $K$  selected heads. For *audio continuation*, we examine another ITI approach, weight decay, where  $\alpha$  is progressively reduced across time steps during the generation process. For *text-to-music*, our evaluation includes both our default configuration ( $\alpha = 5.0$ ) and a stronger intervention level of  $\alpha = 10.0$  to observe the trade-off between the intervention strength and the preservation of musical fluency within the generated samples.

Table 2 presents the objective evaluation for the *audio continuation* task. The baseline methods reveal that MusicGen inherently tends to add certain instruments without explicit direction with varying success rates across instruments; drums and bass are more easily added compared to guitar and piano. Intriguingly, text-prompted generation does not consistently achieve the target instruction. Intervention methods, on the other hand, consistently outperform baseline approaches in successfully directing MusicGen toward adding the desired instruments, as indicated by the success

rate. The original ITI, applying a constant level of intervention at each generation step, leads to a higher success rate but significantly alters the distribution of generated content, which can be observed via the similarity score for maintaining musical consistency. In contrast, the weight decay ITI approach has a subtler influence on the quality but produces a lower success rate. Our proposed SMITIN shows a notable 10.5% jump over text-prompt conditioning and is better at retaining the model’s output distribution and generating consistent music. When SMITIN is deployed in conjunction with text conditioning, it outperforms text-based instructions, but not to the same extent as when used independently (likely because we do not intervene on the text cross-attention layers).

For the *text-to-music* task, we explore the interplay of ITI with text prompts, as shown in Table 3. In this context, we utilize MusicCaps text prompts as the basis of MusicGen’s performance. Consistent with earlier observations, the original ITI method, despite showing the highest success rate in instrument addition, performs the lowest in terms of FAD, indicating a more pronounced shift from the natural music distribution. SMITIN, in its base configuration ( $\alpha = 5.0$ ), achieves an average performance on par with additional text guidance. We further analyze SMITIN with an increased intervention strength of  $\alpha = 10.0$  and observe that a clear trade-off emerges between



TABLE 3: Objective evaluation on *text-to-music*.  $FAD_{L-CLAP\ mus}$  is computed with MusicCaps.

Method		Success Rate [%] ( $\uparrow$ )					$FAD_{L-CLAP\ mus}$ ( $\downarrow$ )				
ITI	Configuration	drums	bass	guitar	piano	avg.	drums	bass	guitar	piano	avg.
X	text	29.9	50.4	17.4	7.1	26.2	0.482 (indep. of inst.)				
	text + “add <inst.>”	34.4	56.1	22.5	10.5	30.8	0.474	0.481	0.473	0.488	0.479
-----		-----		-----		-----		-----		-----	
	text + original ITI	44.3	79.6	33.4	51.0	52.0	0.506	0.577	0.555	0.517	0.538
✓	text + SMITIN ( $\alpha = 5.0$ )	29.9	53.4	23.1	12.5	29.7	0.471	0.493	0.499	0.487	0.487
	text + SMITIN ( $\alpha = 10.0$ )	40.9	60.7	35.4	21.0	39.5	0.481	0.516	0.507	0.485	0.497

TABLE 4: Average and 95% confidence interval of subjective listening tests for overall audio quality across four instrument addition interventions for audio continuation and text-to-music.

Method	Continuation	Text-to-Music
No intervention	3.35 $\pm$ 0.10	3.67 $\pm$ 0.09
“add <inst.>”	3.27 $\pm$ 0.10	3.61 $\pm$ 0.09
original ITI	3.28 $\pm$ 0.11	3.55 $\pm$ 0.09
SMITIN	3.34 $\pm$ 0.10	3.80 $\pm$ 0.09
MoisesDB	3.98 $\pm$ 0.29	

TABLE 5: Subjective results across multiple algorithms, comparing human (*Hum.*) detection on the presence of instruments added by intervention, with Success Rate (*Alg.*). All Spearman correlation coefficients ( $\rho$ ) have p-values less than 1e-5.

Inst.	Audio continuation					Text-to-music				
	$\rho$	Top		Bottom		$\rho$	Top		Bottom	
		Hum.	Alg.	Hum.	Alg.		Hum.	Alg.	Hum.	Alg.
drums	0.76	0.88	0.81	0.19	0.13	0.80	0.95	0.95	0.16	0.18
bass	0.42	0.78	0.89	0.48	0.06	0.58	0.80	0.96	0.38	0.13
guitar	0.40	0.78	0.76	0.52	0.15	0.67	0.86	0.88	0.34	0.14
piano	0.32	0.60	0.72	0.33	0.32	0.55	0.70	0.79	0.31	0.23
avg.	0.48	0.76	0.80	0.39	0.17	0.67	0.83	0.90	0.30	0.17

the success rate and the naturalness of the generated music. This outcome highlights the flexibility of SMITIN, offering users a tunable “knob” to balance between precision in achieving specific musical characteristics and maintaining the authenticity of the musical piece.

#### D. Subjective Evaluation

To further validate our intervention techniques on MusicGen, we performed a subjective listening test through Amazon Mechanical Turk following the best practices from [46]. The purpose of this test was not to evaluate the quality and relevance of MusicGen as this was already explored in [2], but rather to provide evidence that: (1) our proposed *success rate* objective metric correlates with human perception, and (2) our proposed intervention techniques are not detrimental to overall audio quality. To do this, we assessed whether humans could detect the presence of the instruments added by our intervention technique, and asked them to rate samples in terms of overall quality on a scale of 1-5. We also included real music samples from MoisesDB to provide a performance baseline. In selecting files for each algorithm and instrument, we randomly selected 20 files above the median objective success rate and 20 files below the median. This leads to 1280 generated files to be rated (20 files x 2 top/bottom x 4 instruments x 4 algorithms x 2 audio continuation/text-to-music). All files in the listening test were ten seconds long and normalized to -12 loudness units full scale (LUFS) [47]. We obtain at least 3 ratings per file. For the audio continuation experiments, we ask raters to ignore the first 3 seconds

TABLE 6: Spearman correlation of instrument recognition between human decisions and recognition models. All correlations have p-values less than 1e-5.

Method	Correlation Coefficient $\rho$ ( $\uparrow$ )				
	drums	bass	guitar	piano	avg.
Success Rate	0.78	0.48	0.57	0.45	0.58
YAMNet	0.80	0.37	0.56	0.33	0.57
PaSST	0.79	0.49	0.59	0.43	0.57
CLAP	0.80	0.53	0.56	0.52	0.65
ConvNeXt	0.69	0.52	0.51	0.46	0.57

TABLE 7: Ablation study on multi-directional ITI (performance averaged over all instruments)

Method	Success Rate [%]			$FAD_{L-CLAP\ mus}$	Similarity
	indiv.	simult.			
unconditioned	11.1	0.1		0.419	0.781
text	27.8	0.4		0.369	0.664
SMITIN	25.3	1.2		0.445	0.734

of each audio file, as that conditioning signal intentionally does not include the target instrument.

Table 4 displays the overall objective quality results. In general, the performance of all algorithms is quite similar, with SMITIN and no intervention having the best overall quality. In addition to overall quality, we also ask whether listeners can detect the presence of the instruments targeted by our intervention. This serves to help validate our success rate objective metric. We also compare the average human score with the average success rate objective metric (*Alg.*) for the top-ranked and bottom-ranked files selected for the listening test on each generation task in Table 5. For the top-ranked files the scores appear to match quite well, while for the bottom-ranked files, it seems human raters tend to overestimate the presence of most instruments compared to the objective metric. We also hypothesize that the lower match for audio continuation compared to text-to-music may be due in part to the test being more subjectively difficult as raters have to focus on the end of the audio file under test.

Table 6 compares the Spearman correlation coefficient  $\rho$  computed between each of the instrument recognition systems discussed in Section IV-A2) and the average human rating computed across all processing algorithms. Notably, using CLAP embeddings yielded the highest  $\rho$  value with human decisions, while the other methods produced similar values. This result further supports the validity of using the Success Rate on generated samples. Not only does the Success Rate show a strong correlation with other recognition models, but the high correlation observed between the Success Rate and CLAP in Table 1 strengthens its reliability as an objective metric for controlling and evaluating generative music.

TABLE 8: Ablation on amount of probing data (performance is averaged over all instruments, \* shows full-set success rate).

Number of Paired Data	Success Rate* [%]	FAD <sub>L-CLAP<sub>mus</sub></sub>	Similarity	Probes Max. Acc. [%]
$n = 10$	31.0	0.357	0.909	79.1
$n = 100$	31.8	0.349	0.914	82.6
$n = 500$	27.3	0.341	0.915	84.4
$n = 1000$	25.5	0.340	0.914	84.8
$n = \text{full}$	23.3	0.336	0.913	85.1

## V. Ablations and Practical Considerations

**Intervention with Multiple Directions.** We explore ITI in scenarios where multiple musical elements are introduced simultaneously. We focus on the continuation of music that has one instrument present and assesses the addition of three others. For instance, the intervention task generates a continuation with drums, bass, and piano when the input music only has guitar. We consider two types of success rates, an individual one for each instrument as before, and a simultaneous one to assess how often all 3 target instruments were added together. We consider a successful case for a song if each instrument is deemed present more than 50% of the time.

The summarized results in Table 7 reveal that while text prompts yield higher success rates for individual instruments, their generated output diverges from the input music’s characteristics. This indicates that text conditioning may not adequately consider the input and opts to generate new content based solely on the given instruction. Without ITI, MusicGen tends to produce continuations focused on a single instrument, which deviates from the desired complex mixtures, as evidenced by the FAD scores. SMITIN, however, significantly outperforms text prompts in preserving input music similarity, and further enables more granular control over each musical aspect, leading to a higher success rate of jointly generating all desired instruments. This fine-grained control mechanism bolsters SMITIN’s potential as a robust tool for complex music generation tasks where maintaining the essence of the input is crucial.

**Generating “Realistic” Music.** To generate more realistic music through ITI, we fit probes to discern between real music (DISCO-10M dataset’s DISCO-200K-high-quality subset [48]) and synthetic outputs from MusicGen itself. Despite MusicGen being trained only on real music, the probing showed high performance with the best-performing probe achieving a 96.2% accuracy rate and an average accuracy of 77.5% in distinguishing real from synthetic music. Leveraging this result, we steer MusicGen in an audio continuation task, using real-world music inputs from the DISCO-10M dataset to generate continuations that maintain the “realistic” quality.

Analyzing the quantitative results depicted in Fig. 5, simple text prompts prove insufficient in preserving the realistic nature of the music; however, ITI emerges as a more effective approach. Regarding FAD scores, configurations of SMITIN with  $\alpha$  values of 1 and 5 yield smaller deviations in distribution compared to text-based methods. Even with an increased intervention strength ( $\alpha = 10$ ), SMITIN does not shift the distribution more than that of the text prompt “high-quality realistic music.”

**Effects of Number of Probing Data.** To investigate the influence of probing data quantity on SMITIN, we test the system with varying sizes of probing datasets randomly selected from the complete set. These subsets consist of  $n = 10, 100, 500,$  and  $1000$  data pairs, equivalent to music durations of 1, 10, 50, and 100 minutes respectively. The results are presented in Table 8, where we show the “full-set success rate” which is computed using the classifier probes trained with the full dataset.

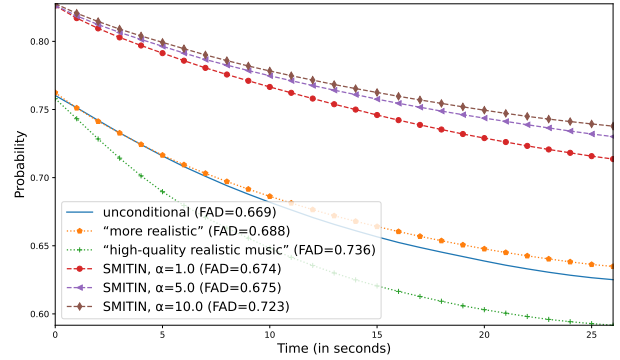


FIGURE 5: Temporal dynamics of maintaining “realism” in audio continuation of real music sample. The graph tracks the probability of generated music being classified as (real) over time, where all configurations of SMITIN demonstrate an enhanced capacity to preserve realistic music qualities over time.

The results indicate that the system demonstrates reasonable results even with a minimal dataset of just 1 minute. As the data size increases to  $n = 500$  (50 minutes), the objective metrics stabilize, indicating an optimal data quantity threshold for effective probe training. An interesting observation from our results is the higher full-set success rates associated with smaller datasets. This trend suggests that probes trained on limited data may have lower confidence in their decisions, prompting them to recommend more interventions during the ITI process to meet the success threshold. At the same time, we observe a clear correlation between better probes (represented here through the accuracy of the best probe) and better ITI performance, whether measured as a higher full-set success rate or similarity, or lower FAD.

This suggests that users can effectively leverage ITI with only a small amount of data to control generation. It means that effective ITI for controlling any musical trait can be accessible and achievable without the need for extensive datasets, which is particularly beneficial for individual artists or smaller studios. This broadens the potential for creative and personalized applications of ITI in music generation, making it a versatile tool for a wide range of users.

## VI. Conclusions and Future Work

We proposed a novel approach for inference-time control of generative music transformers, which self-monitors probe accuracy to impose desired musical traits while maintaining overall music quality. A limitation of our approach is the dependence of performance on probe accuracy, that is, if the pre-trained transformer has not accurately learned a concept, or if the probe training set is inadequate, the success rate of our intervention technique may suffer (although it still could be creatively useful). In the future, we plan to explore different generative model architectures [5], [21], [49] and interventions using musical traits such as genre, emotion, etc. Furthermore, we hope to investigate building “knobs” for interactively controlling the hyper-parameters necessary for inference-time control to enable new technical and creative interactions.

## REFERENCES

- [1] P. Dhariwal, H. Jun, C. Payne, J. W. Kim, A. Radford, and I. Sutskever, “Jukebox: A generative model for music,” *arXiv preprint arXiv:2005.00341*, 2020.
- [2] J. Copet, F. Kreuk, I. Gat, T. Remez, D. Kant, G. Synnaeve, Y. Adi, and A. Défossez, “Simple and controllable music generation,” in *Proc. NeurIPS*, 2023.

- [3] H. Liu, Y. Yuan, X. Liu, X. Mei, Q. Kong, Q. Tian, Y. Wang, W. Wang, Y. Wang, and M. D. Plumbley, "AudioLDM 2: Learning holistic audio generation with self-supervised pretraining," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 2024.
- [4] S.-L. Wu, C. Donahue, S. Watanabe, and N. J. Bryan, "Music ControlNet: Multiple time-varying controls for music generation," *arXiv preprint arXiv:2311.07069*, 2023.
- [5] H. F. F. Garcia, P. Seetharaman, R. Kumar, and B. Pardo, "VampNet: Music generation via masked acoustic token modeling," in *Proc. ISMIR*, 2023.
- [6] F. Kreuk, G. Synnaeve, A. Polyak, U. Singer, A. Défossez, J. Copet, D. Parikh, Y. Taigman, and Y. Adi, "AudioGen: Textually guided audio generation," in *Proc. ICLR*, 2022.
- [7] Z. Borsos, R. Marinier, D. Vincent, E. Kharitonov, O. Pietquin, M. Shariif, D. Roblek, O. Teboul, D. Grangier, M. Tagliasacchi *et al.*, "AudioLM: A language modeling approach to audio generation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 2023.
- [8] C. Hawthorne, I. Simon, A. Roberts, N. Zeghidour, J. Gardner, E. Manilow, and J. Engel, "Multi-instrument music synthesis with spectrogram diffusion," in *Proc. ISMIR*, 2022.
- [9] C. Donahue, A. Caillon, A. Roberts, E. Manilow, P. Esling, A. Agostinelli, M. Verzetti, I. Simon, O. Pietquin, N. Zeghidour *et al.*, "SingSong: Generating musical accompaniments from singing," *arXiv preprint arXiv:2301.12662*, 2023.
- [10] J. D. Parker, J. Spijkervet, K. Kosta, F. Yesiler, B. Kuznetsov, J.-C. Wang, M. Avent, J. Chen, and D. Le, "Stemgen: A music generation model that listens," in *Proc. ICASSP*, 2024, pp. 1116–1120.
- [11] J. Nistal, M. Pasini, C. Aouameur, M. Grachten, and S. Lattner, "Diff-a-riff: Musical accompaniment co-creation via latent diffusion models," *arXiv preprint arXiv:2406.08384*, 2024.
- [12] M. Levy, B. Di Giorgi, F. Weers, A. Katharopoulos, and T. Nickson, "Controllable music production with diffusion models and guidance gradients," *arXiv preprint arXiv:2311.00613*, 2023.
- [13] Z. Novack, J. McAuley, T. Berg-Kirkpatrick, and N. J. Bryan, "DITTO: Diffusion inference-time T-optimization for music generation," *arXiv preprint arXiv:2401.12179*, 2024.
- [14] L. Lin, G. Xia, J. Jiang, and Y. Zhang, "Content-based controls for music large language modeling," *arXiv preprint arXiv:2310.17162*, 2023.
- [15] L. Lin, G. Xia, Y. Zhang, and J. Jiang, "Arrange, inpaint, and refine: Steerable long-term music audio generation and editing via content-based controls," in *Proc. IJCAI*, 2024, pp. 7690–7698.
- [16] Y. Zhang, Y. Ikemiya, G. Xia, N. Murata, M. A. Martínez-Ramírez, W.-H. Liao, Y. Mitsufuji, and S. Dixon, "MusicMagus: Zero-shot text-to-music editing via diffusion models," in *Proc. IJCAI*, 2024.
- [17] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *Proc. ICLR*, 2022.
- [18] N. Zeghidour, A. Luebs, A. Omran, J. Skoglund, and M. Tagliasacchi, "SoundStream: An end-to-end neural audio codec," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 495–507, 2021.
- [19] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *Trans. Machine Learning Research*, 2023.
- [20] R. Kumar, P. Seetharaman, A. Luebs, I. Kumar, and K. Kumar, "High-fidelity audio compression with improved RVQGAN," in *Proc. NeurIPS*, 2023.
- [21] Z. Evans, J. D. Parker, C. Carr, Z. Zukowski, J. Taylor, and J. Pons, "Stable audio open," *arXiv preprint arXiv:2407.14358*, 2024.
- [22] G. Mariani, I. Tallini, E. Postolache, M. Mancusi, L. Cosmo, and E. Rodolà, "Multi-source diffusion models for simultaneous music generation and separation," in *Proc. ICLR*, 2023.
- [23] A. Ziv, I. Gat, G. L. Lan, T. Remez, F. Kreuk, A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "Masked audio generation using a single non-autoregressive transformer," in *Proc. ICLR*, 2024.
- [24] N. Subramani, N. Suresh, and M. E. Peters, "Extracting latent steering vectors from pretrained language models," in *Proc. ACL*, 2022, pp. 566–581.
- [25] A. Turner, L. Thiergart, D. Udell, G. Leech, U. Mini, and M. MacDiarmid, "Activation addition: Steering language models without optimization," *arXiv preprint arXiv:2308.10248*, 2023.
- [26] E. Hernandez, B. Z. Li, and J. Andreas, "Inspecting and editing knowledge representations in language models," *arXiv preprint arXiv:2304.00740*, 2023.
- [27] K. Li, O. Patel, F. Viégas, H. Pfister, and M. Wattenberg, "Inference-time intervention: Eliciting truthful answers from a language model," in *Proc. NeurIPS*, 2023.
- [28] G. Alain and Y. Bengio, "Understanding intermediate layers using linear classifier probes," in *Proc. ICLR*, 2017.
- [29] Y. Belinkov, "Probing classifiers: Promises, shortcomings, and advances," *Computational Linguistics*, vol. 48, no. 1, pp. 207–219, 2022.
- [30] I. Tenney, D. Das, and E. Pavlick, "BERT rediscovers the classical NLP pipeline," in *Proc. ACL*, 2019, pp. 4593–4601.
- [31] R. Castellon, C. Donahue, and P. Liang, "Codified audio language modeling learns useful representations for music information retrieval," in *Proc. ISMIR*, 2021.
- [32] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NeurIPS*, vol. 30, 2017.
- [33] Z. Rafii, A. Liutkus, F.-R. Stöter, S. I. Mimitakis, and R. Bittner, "MUSDB18-HQ - an uncompressed version of MUSDB18," Aug. 2019. [Online]. Available: <https://doi.org/10.5281/zenodo.3338373>
- [34] I. Pereira, F. Araújo, F. Korzeniowski, and R. Vogl, "MoisesDB: A dataset for source separation beyond 4-stems," in *Proc. ISMIR*, 2023, pp. 619–626.
- [35] A. Agostinelli, T. I. Denk, Z. Borsos, J. Engel, M. Verzetti, A. Caillon, Q. Huang, A. Jansen, A. Roberts, M. Tagliasacchi *et al.*, "MusicLM: Generating music from text," *arXiv preprint arXiv:2301.11325*, Oct. 2023.
- [36] Google, "YAMNet: A pretrained audio classifier," <https://tfhub.dev/google/yamnet/1>, 2021.
- [37] K. Koutini, J. Schlüter, H. Eghbal-Zadeh, and G. Widmer, "Efficient training of audio transformers with patchout," in *Proc. Interspeech*, 2022.
- [38] Y. Wu, K. Chen, T. Zhang, Y. Hui, T. Berg-Kirkpatrick, and S. Dubnov, "Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation," in *Proc. ICASSP*, 2023, pp. 1–5.
- [39] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A ConvNet for the 2020s," in *Proc. CVPR*, 2022, pp. 11976–11986.
- [40] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio Set: An ontology and human-labeled dataset for audio events," in *Proc. ICASSP*. IEEE, 2017, pp. 776–780.
- [41] E. Humphrey, S. Durand, and B. McFee, "OpenMIC-2018: An open data-set for multiple instrument recognition," in *Proc. ISMIR*, 2018, pp. 438–444.
- [42] K. Kilgour, M. Zuluaga, D. Roblek, and M. Sharifi, "Fréchet Audio Distance: A metric for evaluating music enhancement algorithms," in *Proc. Interspeech*, 2019, pp. 2350–2354.
- [43] S. Hershey, S. Chaudhuri, D. P. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold *et al.*, "CNN architectures for large-scale audio classification," in *Proc. ICASSP*, 2017, pp. 131–135.
- [44] A. Gui, H. Gamper, S. Braun, and D. Emmanouilidou, "Adapting fréchet audio distance for generative music evaluation," in *Proc. ICASSP*. IEEE, 2024, pp. 1331–1335.
- [45] J. Koo, S. Paik, and K. Lee, "End-to-end music remastering system using self-supervised and adversarial training," in *Proc. ICASSP*, 2022, pp. 4608–4612.
- [46] F. Ribeiro, D. Florêncio, C. Zhang, and M. Seltzer, "CrowdMOS: An approach for crowdsourcing mean opinion score studies," in *Proc. ICASSP*, 2011, pp. 2416–2419.
- [47] E. Grimm, R. Van Everdingen, and M. J. L. C. Schöpping, "Toward a recommendation for a european standard of peak and LKFS loudness levels," *SMPTE Motion Imaging Journal*, vol. 119, no. 3, pp. 28–34, 2010.
- [48] L. A. Lanzendorfer, F. Grötschla, E. Funke, and R. Wattenhofer, "DISCO-10M: A large-scale music dataset," in *Proc. NeurIPS*, 2023.
- [49] J. Hewitt, J. Thickstun, C. D. Manning, and P. Liang, "Backpack language models," in *Proc. ACL*, 2023, pp. 9103–9125.
- [50] J. Koo, Y. Chae, C.-B. Jeon, and K. Lee, "Self-refining of pseudo labels for music source separation with noisy labeled data," in *Proc. ISMIR*, 2023, pp. 716–724.
- [51] E. Law, K. West, M. I. Mandel, M. Bay, and J. S. Downie, "Evaluation of algorithms using games: The case of music tagging," in *Proc. ISMIR*, 2009, pp. 387–392.

TABLE 9: Comparison of multi-label instrument recognition for head-wise probing on MusicGen and a supervised model.

Model	Accuracy / F1 Score					Num. Param.
	<i>vocals</i>	<i>bass</i>	<i>drums</i>	<i>other</i>	<i>avg.</i>	
ConvNeXt <sub>tiny</sub> [39]	<b>97.8%</b> / <b>0.947</b>	<b>94.4%</b> / 0.891	95.1% / 0.914	<b>93.2%</b> / <b>0.880</b>	<b>95.1%</b> / 0.906	28.5M
MusicGen <sub>small</sub>	92.5% / 0.929	91.6% / 0.920	95.0% / 0.952	87.1% / 0.868	91.5% / 0.917	300M
MusicGen <sub>medium</sub>	92.9% / 0.934	91.8% / <b>0.923</b>	94.4% / 0.947	87.4% / 0.874	91.6% / 0.919	1.5B
MusicGen <sub>large</sub>	92.8% / 0.932	91.7% / 0.920	95.1% / 0.953	85.7% / 0.872	91.8% / 0.919	3.3B
MusicGen <sub>melody</sub>	94.1% / 0.945	91.8% / <b>0.923</b>	<b>95.8%</b> / <b>0.960</b>	87.9% / 0.876	92.4% / <b>0.926</b>	1.5B

- [52] G. Tzanetakis and P. Cook, “Musical genre classification of audio signals,” *IEEE Trans. Speech Audio Process.*, vol. 10, no. 5, pp. 293–302, 2002.
- [53] P. Knees, Á. Faraldo Pérez, H. Boyer, R. Vogl, S. Böck, F. Hörschläger, M. Le Goff *et al.*, “Two data sets for tempo estimation and key detection in electronic dance music annotated from user corrections,” in *Proc. ISMIR*, 2015.
- [54] M. Soleymani, M. N. Caro, E. M. Schmidt, C.-Y. Sha, and Y.-H. Yang, “1000 songs for emotional analysis of music,” in *Proc. ACM Workshop on Crowdsourcing for Multimedia (CrowdMM)*, 2013, pp. 1–6.

## Appendix

### A. Analytical Insights from Probing

Following the probing method outlined in the main paper, we investigate additional downstream tasks through probing to present more objective results and enhance our understanding. First, we assess MusicGen’s probing capabilities by comparing them to ConvNeXt [39], a model trained via supervised learning. Subsequently, we contrast our results with prior work on music probing [31] which uses entire attention layer outputs from Jukebox, another generative music transformer model.

#### 1) Multi-label Instrument Recognition

Following up on the performance of single-instrument recognition, we compare the capabilities of individual attention heads in MusicGen against a supervised model trained for the multi-label instrument recognition task, i.e., identifying the presence of all instruments of a given music clip. This is achieved by fitting multiple logistic regression classifiers (probes) for each instrument class using again the extracted intermediate MusicGen activations at the last time step. For the final evaluation, we select the best-performing probe for each instrument class, noting that these may originate from different attention layers ( $l$ ) and heads ( $h$ ). Following the methodology described in [50], we conduct recognition tasks on 3-second music segments from the MUSDB dataset [33]. We also report the performance of a model trained in a supervised manner specifically for instrument recognition [50], utilizing the ‘*tiny*’ configuration of ConvNeXt [39].

The objective results for multi-label instrument recognition are detailed in Table 9. We find that the performance of MusicGen’s probes is comparable to that of the supervised method. However, across all configurations of MusicGen, we note strong performance in terms of the F1 score. Specifically, MusicGen exhibits exceptional performance in recognizing *drums*, whereas its least effective performance is observed in the *other* category. This discrepancy likely arises from the ambiguous definition of the *other* category in the MUSDB18 dataset, with the supervised method gaining an advantage by explicitly training on the dataset’s label, thereby enhancing its ability to identify *other*.

#### 2) Various Music Downstream Tasks

To expand our understanding of MusicGen’s capabilities beyond instrument recognition, we conducted probing on a set of more general music understanding tasks. These include music tagging (MTT) [51], genre classification (GTZAN) [52], key detection (GS) [53], and emotion recognition (EMO) [54]. We adopt appropriate regression models for probing for these tasks: multinomial logistic regression for multi-class classification tasks, multiple logistic regressions for multi-label classification tasks, and linear regression for regression tasks. For these tasks, we set the input duration as 29.0 seconds, and let the probes take the last time step activation output for classification.

Table 10 presents the probing results for various configurations of MusicGen using our attention head probing, and a prior music probing work [31] that relies on a comparative model, Jukebox [1], but uses a different probing methodology:

- [31] probed Jukebox using entire intermediate attention layers, whereas our approach evaluates performance based on individual attention heads.
- [31] employed a one-layer MLP with 512 hidden units. In preliminary experiments, using a one-layer MLP for head-wise probes showed limited benefits, possibly because the attention dimension per head that we use is 75 times smaller than that of Jukebox.

Jukebox probing typically shows superior performance, maybe due to the model’s larger number of parameters and activation dimension. In music tagging, we present both the performance of the best-performing probe across all classes and the ensemble result of the best-performing probes for each class, indicated within parentheses. We observe an improved performance with the ensemble technique, though the improvement is typically not substantial. Interestingly, this means that we can find a single head capable of understanding multiple attributes simultaneously, a notable feature given the MTT dataset comprises 188 different tags. While a larger number of parameters generally correlates with enhanced performance, the melody configuration outperforms others in key detection. This superior performance is likely attributed to the model’s training with a melody conditioning objective [2].

Figure 6 provides a qualitative visualization of the performance of MusicGen<sub>large</sub> probes across various music understanding tasks, alongside histograms representing the distribution of head-wise performance. In tasks such as multi-label instrument recognition and music tagging, there is a noticeable similarity in the trends observed both in the histograms and the layer-wise performance distributions. For key detection, the results indicate that comprehension is limited to only a few heads. On the other hand, emotion recognition

TABLE 10: Performance comparison of layer-wise and head-wise probing on generative music transformers. For music tagging, we report the performance of the best-performing probe across all classes, with the ensemble result of the best-performing probes for each class indicated within parentheses.

Dataset Task	MTT Tagging		GTZAN Genre	GS Key	EMO Emotion			Dim.	Num. Param.
	AUC	AP	Acc	Acc <sup>ref.</sup>	R2 <sup>A</sup>	R2 <sup>V</sup>	Avg.		
Jukebox	91.5	41.4	79.7	66.7	72.1	61.7	69.9	4.8K	5B
MusicGen <sub>small</sub>	85.5 (86.7)	34.1 (37.5)	66.2	46.6	64.2	43.5	55.2	64	300M
MusicGen <sub>medium</sub>	85.9 (87.3)	33.9 (38.4)	69.7	57.4	65.3	51.6	59.6	64	1.5B
MusicGen <sub>large</sub>	85.1 (87.2)	32.9 (38.5)	71.0	58.5	69.1	49.3	60.3	64	3.3B
MusicGen <sub>melody</sub>	85.8 (87.1)	33.3 (38.1)	65.2	62.1	64.7	44.8	58.8	64	1.5B

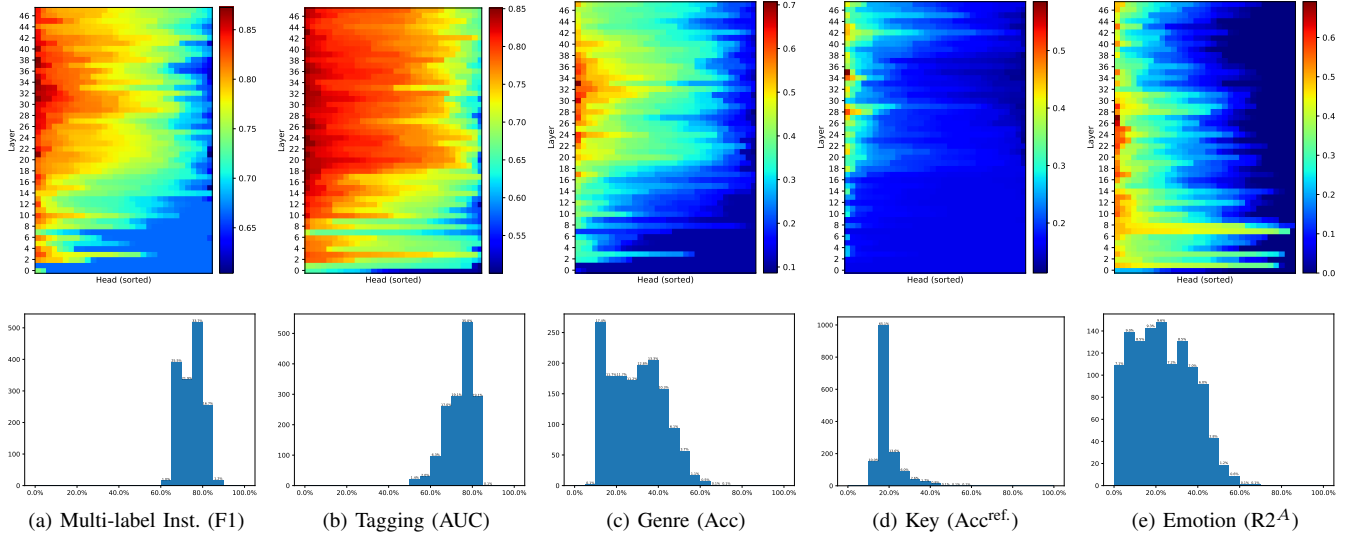


FIGURE 6: Probe performance on various music understanding tasks. Each set of figures shows the performance on each task with probes fitted with MusicGen<sub>large</sub> model output activations. Histograms in the bottom row show the distribution of head-wise probes according to their performance.

demonstrates a relatively uniform performance across the majority of heads in the earlier layers, particularly layer 7. However, the highest-performing head for emotion recognition is located in a middle layer, specifically at the 27<sup>th</sup> layer. This analysis underscores the nuanced role that different attention heads play in music understanding, highlighting the potential for precise control on each attention head to optimize performance for specific tasks within generative music models.

### B. Subjective Listening Test Setup and Detailed Results

In selecting files for the listening test discussed in the main paper, for each algorithm and instrument addition intervention, we randomly select 20 files above the median objective success rate, and 20 files below the median. This leads to 1280 generated files to be rated (20 files x 2 top/bottom x 4 instruments x 4 algorithms x 2 audio continuation/text-to music). We additionally include real music samples from MoisesDB as references. All audio files included in the listening test were ten seconds in length and normalized to -12 loudness units full scale (LUFS) [47]. We obtain at least 3 ratings per file. For the audio continuation experiments, we ask raters to ignore the first 3 seconds of each audio file, as that conditioning signal intentionally does not include the target instrument. In addition to following the best practices from [46], we limited participation to

“Mechanical Turk Masters” for which Amazon charged a \$0.005 fee per assignment, and paid \$0.10 per rated file.

Table 11 displays the overall objective quality results broken out by instrument and algorithm. In general, the performance of all algorithms are quite similar, with SMITIN an no intervention having the best overall quality.

In addition to overall quality, we also ask whether listeners can detect the presence of the instruments targeted by our intervention. This serves to help validate our success rate objective metric, which was shown and discussed in the main paper.

### C. Intervention with Multiple Directions

We explore the application of ITI in scenarios where multiple musical elements are introduced simultaneously. In our setup, we focus on the continuation of music that has one instrument present and assesses the addition of three others. For instance, the intervention task is to generate a continuation with drums, bass, and piano when the input music only has guitar. We consider two types of success rates, an individual one for each instrument as before, and a simultaneous one to assess how often all 3 target instruments were added together. We consider a successful case for a song if each instrument is deemed present more than 50% of the time.

The summarized results in Tables 12 and 13 reveal that while text prompts yield higher success rates for individual instruments, their

TABLE 11: Average and 95% confidence interval of subjective listening tests for overall audio quality across four instrument addition interventions for both audio continuation and text-to-music experiments.

Method	Audio continuation					Text-to-music				
	<i>drums</i>	<i>bass</i>	<i>guitar</i>	<i>piano</i>	<i>avg</i>	<i>drums</i>	<i>bass</i>	<i>guitar</i>	<i>piano</i>	<i>avg</i>
No intervention	3.22 ± 0.21	3.75 ± 0.18	3.26 ± 0.22	3.16 ± 0.18	3.35 ± 0.10	3.63 ± 0.18	3.66 ± 0.20	3.70 ± 0.19	3.70 ± 0.17	3.67 ± 0.09
“add <inst.>”	3.31 ± 0.21	3.41 ± 0.18	3.06 ± 0.21	3.31 ± 0.22	3.27 ± 0.10	3.78 ± 0.18	3.61 ± 0.18	3.54 ± 0.21	3.50 ± 0.19	3.61 ± 0.09
original ITI	3.42 ± 0.22	3.04 ± 0.21	3.54 ± 0.20	3.12 ± 0.22	3.28 ± 0.11	3.62 ± 0.19	3.61 ± 0.18	3.65 ± 0.19	3.33 ± 0.20	3.55 ± 0.09
SMITIN	3.40 ± 0.20	3.45 ± 0.18	3.27 ± 0.21	3.23 ± 0.19	3.34 ± 0.10	3.93 ± 0.18	3.75 ± 0.17	3.77 ± 0.18	3.73 ± 0.17	3.80 ± 0.09
<i>MoisesDB</i>	3.98 ± 0.29									

TABLE 12: Intervention with Multiple Directions (Individual Success Rate).

Method	Individual Success Rate [%]																
	start with <i>drums</i>				start with <i>bass</i>				start with <i>guitar</i>				start with <i>piano</i>				total
	<i>bass</i>	<i>guitar</i>	<i>piano</i>	<i>avg.</i>	<i>drums</i>	<i>guitar</i>	<i>piano</i>	<i>avg.</i>	<i>drums</i>	<i>bass</i>	<i>piano</i>	<i>avg.</i>	<i>drums</i>	<i>bass</i>	<i>guitar</i>	<i>avg.</i>	
unconditioned	29.0	4.3	0.2	11.1	17.7	10.6	1.4	9.9	18.8	18.4	1.9	13.0	11.9	10.4	9.1	10.4	11.1
text	58.4	2.6	5.2	22.0	45.7	2.8	2.9	17.1	45.5	67.0	4.8	39.1	42.9	50.4	6.6	33.3	27.8
SMITIN	48.3	16.1	1.6	22.0	25.4	23.6	4.1	17.7	41.1	49.0	3.3	31.1	30.2	35.7	26.2	30.7	25.3

TABLE 13: Intervention with Multiple Directions (Simultaneous Success Rate, FAD, and Similarity).

Method	Simultaneous Success Rate [%] (↑)					FAD <sub>L-CLAP<sub>mus</sub></sub> (↓)					Similarity (↑)				
	<i>drums</i>	<i>bass</i>	<i>guitar</i>	<i>piano</i>	<i>avg.</i>	<i>drums</i>	<i>bass</i>	<i>guitar</i>	<i>piano</i>	<i>avg.</i>	<i>drums</i>	<i>bass</i>	<i>guitar</i>	<i>piano</i>	<i>avg.</i>
unconditioned	0.0	0.0	0.2	0.2	0.1	0.442	0.517	0.280	0.440	0.419	0.884	0.580	0.844	0.817	0.781
text	0.3	0.0	0.5	0.9	0.4	0.364	0.426	0.376	0.312	0.369	0.847	0.565	0.589	0.655	0.664
SMITIN	0.1	0.4	0.5	3.9	1.2	0.467	0.562	0.300	0.452	0.445	0.837	0.549	0.790	0.763	0.734

TABLE 14: Objective evaluation on *Instrument Removal*.

ITI	Method Configuration	Success Rate [%] (↑)					FAD <sub>L-CLAP<sub>mus</sub></sub> (↓)					Similarity (↑)				
		<i>drums</i>	<i>bass</i>	<i>guitar</i>	<i>piano</i>	<i>avg.</i>	<i>drums</i>	<i>bass</i>	<i>guitar</i>	<i>piano</i>	<i>avg.</i>	<i>drums</i>	<i>bass</i>	<i>guitar</i>	<i>piano</i>	<i>avg.</i>
X	unconditioned	2.2	16.2	14.1	31.6	16.0	0.442	0.517	0.338	0.513	0.452	0.884	0.579	0.847	0.816	0.781
	“remove <inst.>”	1.8	4.8	19.2	23.2	12.2	0.568	0.634	0.420	0.789	0.602	0.871	0.545	0.795	0.808	0.754
	“no <inst.>”	7.3	10.8	17.0	32.5	16.9	0.267	0.267	0.307	0.289	0.282	0.934	0.889	0.937	0.943	0.925
✓	original ITI	11.5	0.0	84.9	95.3	47.9	0.508	0.598	0.686	0.565	0.589	0.833	0.835	0.702	0.867	0.809
	SMITIN	9.1	5.1	79.9	78.4	43.1	0.481	0.538	0.481	0.455	0.488	0.779	0.545	0.691	0.778	0.698

generated output diverges from the input music’s characteristics. This indicates that text conditioning may not adequately consider the input and opts to generate new content based solely on the given instruction. Without ITI, MusicGen tends to produce continuations focused on a single instrument, which deviates from the desired complex mixtures, as evidenced by the FAD scores. SMITIN, however, significantly outperforms text prompts in preserving input music similarity, and further enables more granular control over each musical aspect, leading to a higher success rate of jointly generating all desired instruments. This fine-grained control mechanism bolsters SMITIN’s potential as a robust tool for complex music generation tasks where maintaining the essence of the input is crucial.

#### D. Instrument Removal

We explore an additional application: *removing* the target instrument during audio continuation using the same probes employed for instrument addition. To accomplish this, we apply negative intervention strength  $\alpha$  to eliminate the target instrument from the given input music. The hyperparameter settings for this task are configured as  $\alpha = -10.0$  for both original ITI and SMITIN, with  $s = 1$  and

$\tau = 0.5$  for SMITIN. The success rate in this context is defined as  $\bar{C}(t) < \tau$ , where  $\tau = 0.5$ .

Table 14 reveals an inverse tendency for each instrument’s removal: drums and bass tend to be continuously generated, whereas guitar and piano are more likely to be omitted during the generation process. Moreover, we find that both text-conditioning prompts: “remove <inst.>”, and “no <inst.>” are ineffective in removing the target instrument. This outcome suggests another aspect of controllability that is challenging to achieve with text prompts alone.

#### E. Soft-weighting

This ablation study justifies the soft-weighting technique over selecting top- $K$  heads for ITI. The rationale for soft-weighting arises from the observation that the optimal set of top  $K$  heads varies depending on the task or instrument. By adopting soft-weighting, we can bypass the need for hyper-parameter tuning specific to each task, simplifying the ITI process.

To validate this approach, we compare the performance of the full SMITIN (with soft-weighting) against SMITIN where the initial weights  $w_{l,h}(t_0)$  are instead 1 for the top- $K$  most accurate heads according to probing (i.e.,  $(l, h) \in \mathcal{H}_K$ ), and 0 otherwise. Results

TABLE 15: Ablation study: top- $K$  head selection vs. head soft-weighting

Top-K Interventions	Success Rate [%] ( $\uparrow$ )					FAD <sub>L-CLAP mus</sub> ( $\downarrow$ )					Similarity ( $\uparrow$ )				
	drums	bass	guitar	piano	avg.	drums	bass	guitar	piano	avg.	drums	bass	guitar	piano	avg.
<b>K = 16</b>	14.9	24.8	11.9	3.9	13.8	0.327	<b>0.249</b>	<b>0.392</b>	0.335	<b>0.325</b>	0.864	<b>0.929</b>	0.931	0.941	0.916
<b>K = 32</b>	14.4	20.2	16.9	5.6	14.2	<b>0.325</b>	0.259	0.393	0.334	0.327	0.867	0.928	0.934	0.941	<b>0.917</b>
<b>K = 64</b>	15.3	26.5	17.4	11.0	17.5	0.335	0.267	0.390	0.336	0.332	<b>0.869</b>	0.923	0.934	0.942	<b>0.917</b>
<b>K = 128</b>	18.2	33.6	27.4	13.8	23.2	0.352	0.277	0.411	<b>0.327</b>	0.341	0.860	0.924	0.932	<b>0.944</b>	0.915
<b>K = 256</b>	23.9	36.0	36.8	16.8	28.3	0.370	0.283	0.405	0.332	0.347	0.848	0.920	0.933	<b>0.944</b>	0.911
<b>K = 512</b>	30.0	50.5	47.8	29.7	39.5	0.386	0.300	0.411	0.339	0.359	0.837	0.917	0.932	0.942	0.907
<b>K = 1024</b>	<b>32.9</b>	<b>79.0</b>	<b>61.1</b>	<b>43.5</b>	<b>54.1</b>	0.560	0.355	0.417	0.363	0.423	0.823	0.895	0.930	0.936	0.896
<b>soft-weighting</b>	20.6	30.4	33.8	8.7	23.3	0.346	0.267	0.397	0.337	0.336	0.857	0.922	<b>0.935</b>	0.938	0.913

TABLE 16: Comparison with different intervention directions on *Audio continuation*.

ITI Direction	Success Rate [%] ( $\uparrow$ )					FAD <sub>L-CLAP mus</sub> ( $\downarrow$ )					Similarity ( $\uparrow$ )				
	drums	bass	guitar	piano	avg.	drums	bass	guitar	piano	avg.	drums	bass	guitar	piano	avg.
$\tilde{\theta}_{l,h}$	20.6	30.4	33.8	8.7	23.3	0.346	0.267	0.397	0.337	<b>0.336</b>	0.857	0.922	0.935	0.938	<b>0.913</b>
<b>mass mean shift</b>	17.0	48.6	36.0	9.7	<b>27.8</b>	0.355	0.291	0.398	0.323	0.341	0.855	0.923	0.928	0.941	0.911

TABLE 17: FAD compared with the distribution of unconditioned generation on *audio continuation* and *text-to-music*.

ITI Configuration	Method	FAD <sub>L-CLAP mus</sub> ( $\downarrow$ ) w/. unconditioned									
		Audio continuation					Text-to-music				
		drums	bass	guitar	piano	avg.	drums	bass	guitar	piano	avg.
$X$	“add <inst.>”	0.042	0.026	0.026	0.064	0.030	0.013	0.009	0.011	0.009	0.010
	original ITI	0.039	0.206	0.066	0.039	0.087	0.018	0.056	0.031	0.017	0.030
$\checkmark$	SMITIN ( $\alpha = 5.0$ )	0.022	0.017	0.018	0.010	<b>0.016</b>	0.012	0.009	0.010	0.006	<b>0.009</b>
	SMITIN ( $\alpha = 10.0$ )	0.079	0.042	0.033	0.014	0.042	0.045	0.017	0.019	0.008	0.022

from Table 15 indicate that while top- $K$  selection can yield good performance, it often requires task-specific tuning of  $K$ . In contrast, soft-weighting demonstrates a well-balanced performance across various metrics without necessitating such tuning. This balance is crucial because it means soft-weighting can adapt to various musical tasks and preferences, making ITI more flexible and user-friendly. However, it is important to acknowledge that soft-weighting may not always be the optimal choice for every ITI application. Users might prefer the top- $K$  approach depending on their specific needs and the nature of their musical goals.

### F. Intervention with Different Directions

Following [27], we compare results using different intervention directions: logistic regression classifiers weights  $\tilde{\theta}_{l,h}$  and mass mean shift. The intervention direction for mass mean shift is determined during probing as a vector that directs from the centroid of negative activations to the centroid of positive activations. The results are shown in Table 16. Similar to [27], we observe that using mass mean shift as the intervention direction better achieves the target musical

factor (according to success rate). However, we again observe a trade-off between controllability and generation distribution shift. Considering that generation quality is crucial in the music generation task, we adopt  $\tilde{\theta}_{l,h}$  as the final intervention direction.

### G. FAD Compared with Unconditioned Generation

To support the claim that SMITIN does not significantly alter the original distribution of MusicGen, we compute the FAD score using the distribution of music outputs from unconditioned generation. For text-to-music, the comparison distribution is that of text-conditioned generation using MusicCaps text prompts, as discussed in the main paper. According to Table 17, the FAD score of “add <inst.>” is significantly lower in *text-to-music* than in *audio continuation*, indicating that the output distribution between unconditional and conditional generation experiences a significant shift. We observe that SMITIN with  $\alpha = 5.0$  exhibits the closest distribution distance to the original MusicGen. This supports the claim that SMITIN achieves the desired controllability while retaining MusicGen’s generation capability.