

AudioGenX: Explainability on Text-to-Audio Generative Models

Hyunju Kang^{1*}, Geonhee Han^{1*}, Yoonjae Jeong², Hogun Park^{1†}

¹Department of Artificial Intelligence, Sungkyunkwan University, Suwon, Republic of Korea

²Audio AI Lab, NCSOFT, Seongnam, Republic of Korea

{neutor, gunhee8178}@skku.edu, hybris75@gmail.com, hogunpark@skku.edu

Abstract

Text-to-audio generation models (TAG) have achieved significant advances in generating audio conditioned on text descriptions. However, a critical challenge lies in the lack of transparency regarding how each textual input impacts the generated audio. To address this issue, we introduce AudioGenX, an Explainable AI (XAI) method that provides explanations for text-to-audio generation models by highlighting the importance of input tokens. AudioGenX optimizes an *Explainer* by leveraging factual and counterfactual objective functions to provide faithful explanations at the audio token level. This method offers a detailed and comprehensive understanding of the relationship between text inputs and audio outputs, enhancing both the explainability and trustworthiness of TAG models. Extensive experiments demonstrate the effectiveness of AudioGenX in producing faithful explanations, benchmarked against existing methods using novel evaluation metrics specifically designed for audio generation tasks.

Introduction

Text-to-audio generation models (TAG) (Kreuk et al. 2023; Ziv et al. 2024; Yang et al. 2023; Liu et al. 2023; Schneider et al. 2023) have emerged as a pivotal technology in generative AI, enabling textual content to be transformed into an auditory experience. Although models such as AudioGen (Kreuk et al. 2023) excel at generating high-quality audio based on textual prompts, a critical challenge remains: the lack of transparency in how each textual input affects the generated audio. Consequently, users may struggle to trust the model, making it essential to provide explanations for the TAG task. Explainability provides several key advantages. First, it enhances awareness of how input tokens affect the model’s outputs, enabling users to ensure that the model emphasizes the correct aspects of the text. Second, it provides actionable insights to support the decision-making about which elements to modify and to what extent in the audio editing process. Third, analyzing generated explanations can aid with debugging and identifying potential biases. Accordingly, this study argues that the ability to quantify the

*Equal contribution.

†Corresponding author.

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

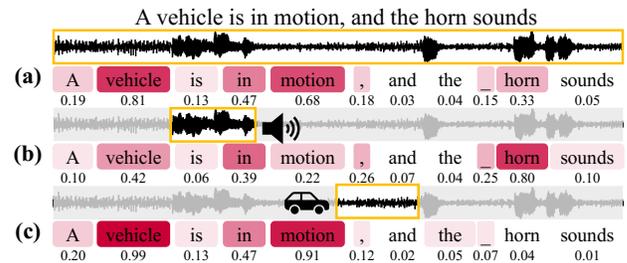


Figure 1: A comprehensive explanation provided by AudioGenX for the entire audio in (a). Granular explanations for the interval from 1 to 1.5 seconds in (b) and from 2.5 to 3 seconds in (c), respectively.

importance of textual inputs in TAG models is crucial to being able to unambiguously assess and communicate their value.

While approaches specifically tailored for explaining TAG models are limited, recent research has explored methodologies for calculating the importance of input tokens in large-scale transformer-based models. Cross-attention layers in multi-modal architectures, such as those in TAG models, are widely regarded as critical for integrating textual and auditory information, while also enhancing explainability by revealing how information from one modality influences another. A notable method (Abnar and Zuidema 2020) utilizes attention weights and aggregates them across all layers to approximate the importance of each input token. However, attention scores alone are not considered reliable for causal insights, as they do not directly indicate how perturbation to specific inputs influences the output. Recently, AtMan (Deiseroth et al. 2023) introduced a perturbation method that suppresses the attention score of one token at a time to observe the impact of each input on output prediction. This single-token perturbation approach, however, may overlook interactions between multiple tokens. Consequently, it provides less reliable explanations in scenarios where the model heavily relies on the contextual relationships between multiple tokens, leading to an oversimplification of the model’s behavior.

To address the challenge of faithful explanations, causal inference theory, encompassing factual and counterfactual

reasoning, is often utilized (Pearl 2009). These two approaches aim to identify impactful input information in different ways. Factual reasoning focuses on identifying critical input information that reproduces the original prediction, whereas counterfactual reasoning (Tan et al. 2022; Ali et al. 2023; Kenny et al. 2021) seeks to determine crucial input information that, if absent, would change the prediction. Given their differing assumptions, these reasoning approaches can be employed together as complementary frameworks to generate more faithful explanations. However, prior research has yet to investigate the feasibility of applying factual and counterfactual reasoning within TAG models.

To provide faithful explanations for the TAG model, we introduce AudioGenX, a perturbation-based explainability method leveraging factual and counterfactual reasoning. Our approach utilizes the latent representation vectors in TAG models to observe the effects of factual and counterfactual perturbations. These perturbations are applied in the cross-attention layer using a soft mask, enabling the simultaneous perturbation of multiple tokens’ attention scores. More importantly, the mask itself serves as an explanation, with its values quantifying the importance of the textual input. We optimize the mask through a gradient descent method guided by our proposed factual and counterfactual objective functions. To mitigate the high computational cost of calculating gradients for the entire sequential audio, we enhance efficiency by decomposing the explanation target into individual audio tokens. This approach enables us to customize the explanation range of generated audio interactively, providing comprehensive explanations for the entire audio or more granular explanations for specific segments of interest, depending on user demand. For instance, in Figure 1, (a) provides a comprehensive explanation for the entire audio, indicating a strong relation to vehicle motion. By focusing on a specific interval in (b) and (c), AudioGenX captures the different contexts of each audio segment and delivers contextually accurate explanations accordingly. Extensive experiments demonstrate the faithfulness of our explanations and benchmark their performance against recent techniques using proposed evaluation metrics for audio generation tasks.

Contributions. We summarize our contributions as follows: 1) We propose a faithful explanation method for text-to-audio generation models, grounded in factual and counterfactual reasoning to quantify the importance of text tokens to the generated audio. 2) We offer a framework that provides both holistic and granular audio explanations based on user requests, enabling tailored insights. 3) We introduce new evaluation metrics for text-to-audio explanations and demonstrate the effectiveness of AudioGenX through extensive experiments compared to existing methods. 4) We present case studies demonstrating how AudioGenX provides valuable insights to support the understanding of model behavior and editing tasks.

Related Work

Text-to-Audio Generation Models. Recent text-to-audio generation models can be categorized into two model architectures: Transformer-based (Kreuk et al. 2023; Ziv et al. 2024) and Diffusion-based (Yang et al. 2023; Liu et al. 2023;

Schneider et al. 2023). Transformer models, such as AudioGen (Kreuk et al. 2023), employ autoregressive Transformers to predict discrete audio tokens, while MAGNeT (Ziv et al. 2024) enhances efficiency through masked generative modeling in a non-autoregressive scheme. Diffusion-based approaches such as Diffsound (Yang et al. 2023) generate discrete mel-spectrogram tokens, whereas models like AudioLDM (Liu et al. 2023) and Moüsai (Schneider et al. 2023) directly predict continuous mel-spectrograms or waveforms. Despite architectural differences, these models commonly use cross-attention mechanisms, making AudioGenX a model-agnostic explainer for TAG models that use cross-attention in audio generation.

Explainable AI. Explainability involves methods that help to understand the importance of each input token with respect to output predictions. These methods generally fall into two categories: gradient-based methods (Selvaraju et al. 2017; Sundararajan, Taly, and Yan 2017; Nagahisarchoghaei et al. 2023) and perturbation-based methods (Ribeiro, Singh, and Guestrin 2016; Lundberg and Lee 2017). Gradient-based explanation methods trace gradients from the target layers to the predictive value, using the calculated gradients as a measure of importance. While effective, these methods require substantial memory resources to store the values of each targeted layer. In contrast, perturbation-based methods, such as SHAP (Lundberg and Lee 2017), are more memory-efficient, calculating feature importance by comparing predictions with and without specific features. Similarly, our method adopts a perturbation-based approach to effectively generate explanations.

Explainability on Audio Processing Models. Existing explainability approaches (Akman and Schuller 2024) on audio processing models have extended generic explanation methods. For instance, one study (Becker et al. 2018) employs Layer-wise Relevance Propagation (LRP) to explain the model trained on raw waveforms and spectrograms for spoken digit and speaker gender classification. Another study applied DFT-LRP (Frommholz et al. 2023) to audio event detection, assessing the significance of time-frequency components and guiding input representation choices. Similarly, audioLIME (Haunschmid, Manilow, and Widmer 2020) extends LIME (Ribeiro, Singh, and Guestrin 2016) to explain music-tagging models by perturbing audio components derived from source separation. However, since the above methods focus on explaining audio continuously and sequentially, they are not directly applicable to the unique challenges posed by TAG models, which require techniques that address the complex interactions between text inputs and generated audio outputs.

Explainability on Transformer. With the widespread use of Transformers, the demand for explainability has grown. Primarily, Rollout (Abnar and Zuidema 2020) primarily aggregates attention weights in all layers to track information flow but struggles to integrate cross-attention weights in multi-modal models with differing domain dimensionalities. Another recent work (Chefer, Gur, and Wolf 2021) leverages Layer-wise Relevance Propagation (LRP) (Samek, Wiegand, and Müller 2017) to calculate class-specific relevance scores based on gradients of attention weights in

self- and cross-attention layers. Nevertheless, AtMan (Deiseroth et al. 2023) raises the issue of excessive memory usage and introduces a scalable explanation method that employs single-token perturbation to observe the change of loss in the response. While intuitive and memory-efficient for large-scale models, this method is limited in its ability to account for the interrelationship of input tokens.

Preliminaries

AudioGen (Kreuk et al. 2023), a representative TAG model, consists of three key components: a text encoder (Raffel et al. 2020), an autoregressive Transformer decoder model (Vaswani et al. 2017), and an audio decoder (Défossez et al. 2023). The Transformer decoder serves as the core model responsible for generating the audio sequence, while the text encoder processes the input text and the audio decoder post-processes the generated audio token sequence into audio. Given a text prompt, it is converted into a tokenized representation vector, denoted as $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_L]$, $\mathbf{U} \in \mathbb{R}^{L \times d_u}$, where L denotes number of textual tokens and d_u represents a dimension of the textual token representation vector. The generated audio can be expressed in a discrete form, as EnCodec (Défossez et al. 2023) converts the audio into either discrete tokens or continuous token representations. The tokenized audio sequence is denoted as $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_T]$, $\mathbf{Z} \in \mathbb{N}^{T \times d_v}$, where T denotes the length of the audio sequence and d_v indicates the number of codebooks d_v . In detail, the codebook is a structured set of discrete audio tokens used in multi-stream audio generation to produce high-quality audio. For more comprehensive information on multi-streaming audio generation, we refer to the original AudioGen paper (Kreuk et al. 2023).

For the generation of an audio sequence, the Transformer-decoder model (Vaswani et al. 2017), denoted as h , generates \mathbf{z}_t as t -th order audio token in the sequence, following the formulation $h(\mathbf{U}, \mathbf{z}_{t-1}) = \mathbf{z}_t$. For brevity, we omit the detailed notation of other components and the top- p or top- k sampling process in the Transformer. Instead, we focus on the attention layers, including cross-attention, which are crucial components of the model, denoted as f . The computation within these layers is expressed in a simplified version as $f(\mathbf{U}, \mathbf{z}_{t-1}) = \mathbf{e}_t$, where \mathbf{e}_t represents the latent representation vector corresponding to the t -th audio token. In the absence of ground truth and class labels, the latent embedding vector \mathbf{e}_t in the audio token space provides information on how perturbation impacts subsequent generations. Particularly, the cross-attention layer is essential to fuse the textual information with auditory information in layers f , we denote the cross-attention layers as:

$$g(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \sigma \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}} \right) \mathbf{V}, \quad (1)$$

where σ indicates a softmax function, $\mathbf{Q}, \mathbf{K}, \mathbf{V}$, d_k refers to query, key, values, and the number of vector dimensions in the k -th layer, respectively. In detail, \mathbf{Q} refers to previously generated audio tokens, representing the query information, while \mathbf{K} and \mathbf{V} correspond to the textual tokens.

The Proposed AudioGenX

AudioGenX addresses the challenge of explaining TAG models, where the goal is to quantify the importance of textual input corresponding to the generated audio. To achieve this within a sequence-to-sequence framework, we decompose the explanation target, represented as sequential audio, into individually non-sequential audio tokens. Since the output is sequential data, calculating gradients across the entire sequence, from the first to the last token, is computationally expensive and time-consuming. To overcome these issues, we redefine the explanation target as individual audio tokens, rather than the entire sequence. This modification enables parallel computation of generating an explanation for each token, significantly speeding up the process. Finally, AudioGenX integrates these individual token-level explanations to provide a comprehensive understanding of the entire audio sequence. An overview of AudioGenX is illustrated in Figure 2.

Definition of Masks as Explanations

We quantify the importance of the t -th audio token \mathbf{z}_t within the audio sequence using a mask as the explanation. The soft mask is denoted as $\mathbf{M}_{\mathbf{U}, \mathbf{z}_t} \in \mathbb{R}^{L \times 1}$, where each element $\mathbf{m}_{\mathbf{u}_i, \mathbf{z}_t} \in \mathbf{M}_{\mathbf{U}, \mathbf{z}_t}$ represents the importance of the i -th textual token with respect to the t -th audio token \mathbf{z}_t . Each value lies in the range $[0, 1]$, where a value close to 1 indicates that the corresponding textual token is highly important for generating the target audio token, while a value closer to 0 indicates lower importance. To serve as a soft mask representing the importance of each text token, AudioGenX optimizes the *Explainer* to predict the mask $\mathbf{M}_{\mathbf{U}, \mathbf{z}_t}$ as the explanation. The *Explainer* consists of Multi-Layer Perceptrons (MLPs) with a sigmoid and gumbel-softmax (Jang et al. 2017) function to constrain values within the range $[0, 1]$ without additional scaling and to enforce the values close to either 0 or 1, thereby highlighting relatively distinguished contribution. Using the soft mask, we apply perturbation to modify the inner computational steps of the cross-attention layers, altering the attention score of the given textual input. Consequently, we measure the perturbation effect on the prediction at the layer $f(\mathbf{U}, \mathbf{z}_{t-1}) = \mathbf{e}_t$, observing how latent representation vector \mathbf{e}_t for the audio token \mathbf{z}_t changes under these perturbations. In the following section, we detail how we optimize *Explainer* to predict the mask as explanations based on both factual and counterfactual reasoning.

Formulating Factual Explanations

Factual reasoning (Tan et al. 2022; Ali et al. 2023; Kenny et al. 2021) aims to find sufficient input that can approximately reproduce the original prediction. To quantify the sufficiency of textual tokens, we employ a perturbation-based method using the soft mask, interleaving the computation to measure the impact of changes. Specifically, we mask out attention scores in the cross-attention layers where textual information is fed into the TAG model. We formulate the perturbation in factual reasoning as:

$$\tilde{g}(\mathbf{Q}, \mathbf{K}, \mathbf{V}, \mathbf{M}_{\mathbf{U}, \mathbf{z}_t}) = \left(\sigma \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}} \right) \odot \mathbf{M}_{\mathbf{U}, \mathbf{z}_t} \right) \mathbf{V}, \quad (2)$$

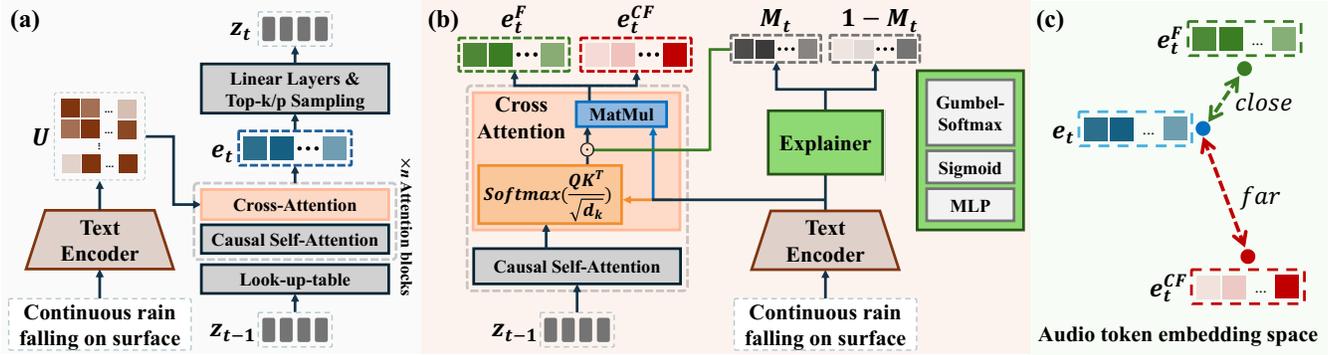


Figure 2: (a) The process by which AudioGen generates an audio. (b) AudioGenX’s procedure for generating and applying explanations, with the *Explainer* in the green box. (c) The method for calculating and applying the loss in AudioGenX.

where σ denotes the softmax activation function and the mask $\mathbf{M}_{\mathbf{U}, \mathbf{z}_t}$ controls the amount of information corresponding to the text token. When the mask value $\mathbf{m}_{\mathbf{u}_i, \mathbf{z}_t}$ approaches 0, the attention score is suppressed, meaning the information corresponding to the textual token is not fully propagated to the subsequent layer. Conversely, as the mask value approaches 1, the original value is fully preserved. To distinguish this process from the original layer $f(\mathbf{U}, \mathbf{z}_{t-1})$, we denote the layer applying perturbation with the factual mask as $\tilde{f}(\mathbf{U}, \mathbf{z}_{t-1}, \mathbf{M}_{\mathbf{U}, \mathbf{z}_t})$.

When the mask sufficiently serves as a factual explanation, the perturbed output remains approximately the same as the original prediction. To evaluate the impact of perturbation, we measure the resulting changes in the latent representation vector within the audio token space. Since the latent vector encodes rich and implicit information, we expect that two vectors close to each other indicate a similar auditory meaning, which is likely to result in similar audio generation. By leveraging this vector similarity, we can effectively measure the influence of perturbation and formulate the objective function for *Explainer* as:

$$\mathcal{L}_F = -\cos(f(\mathbf{U}, \mathbf{z}_{t-1}), \tilde{f}(\mathbf{U}, \mathbf{z}_{t-1}, \mathbf{M}_{\mathbf{U}, \mathbf{z}_t})), \quad (3)$$

where \cos function refers to cosine similarity, which measures how similar factual result $\tilde{f}(\mathbf{U}, \mathbf{z}_{t-1}, \mathbf{M}_{\mathbf{U}, \mathbf{z}_t})$ is to the original prediction $f(\mathbf{U}, \mathbf{z}_{t-1})$ in the audio token space. Since the objective function involves negative cosine similarity, minimizing the loss function corresponds to maximizing the similarity. Hence, following the objective function, the *Explainer* generates the factual explanation mask, ensuring that the two representations or predictions are as close as possible in the audio token space.

Formulating Counterfactual Explanations

Counterfactual reasoning (Tan et al. 2022; Ali et al. 2023; Kenny et al. 2021) aims to identify necessary inputs that can significantly alter the original prediction when it is perturbed or removed. This perturbation operates in the opposite direction of factual explanations, removing the important input to observe the counterfactual result. We formulate the pertur-

bation method in counterfactual reasoning as:

$$\tilde{g}(\mathbf{Q}, \mathbf{K}, \mathbf{V}, \mathbf{1} - \mathbf{M}_{\mathbf{U}, \mathbf{z}_t}) = (\sigma(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}})) \odot (\mathbf{1} - \mathbf{M}_{\mathbf{U}, \mathbf{z}_t})\mathbf{V}, \quad (4)$$

where $\mathbf{1} \in \mathbb{R}^{1 \times T_u}$ is a vector of ones and $\mathbf{1} - \mathbf{M}_{\mathbf{U}, \mathbf{z}_t}$ subtracts the importance of the corresponding textual tokens. Consequently, the more important a textual token is, the more its attention score is suppressed in proportion to its importance. This perturbation operates under a counterfactual assumption as the What-If scenario (Tan et al. 2022; Ali et al. 2023; Kenny et al. 2021): What happens if the important textual token does not exist? After applying the perturbation in Equation (4), the counterfactual result is observed as $\tilde{f}(\mathbf{U}, \mathbf{z}_{t-1}, \mathbf{1} - \mathbf{M}_{\mathbf{U}, \mathbf{z}_t})$. If the counterfactual result significantly differs from the original prediction, it indicates that the counterfactual mask is necessary to explain the original prediction. Conversely, if the change is trivial, the counterfactual mask is unnecessary to explain the causal relationship with the prediction.

Generally, counterfactual explanations in supervised settings aim to find the important inputs that change the prediction with minimal perturbation. However, no class labels or guidance are available in our task of audio generation. Instead, we measure the change of meaning in latent space leveraging the cosine similarity function after counterfactual perturbation. Thus, the counterfactual explanation objective function is formulated as:

$$\mathcal{L}_{CF} = \cos(f(\mathbf{U}, \mathbf{z}_{t-1}), \tilde{f}(\mathbf{U}, \mathbf{z}_{t-1}, \mathbf{1} - \mathbf{M}_{\mathbf{U}, \mathbf{z}_t})), \quad (5)$$

where \cos function measures how dissimilar counterfactual result $\tilde{f}(\mathbf{U}, \mathbf{z}_{t-1}, \mathbf{1} - \mathbf{M}_{\mathbf{U}, \mathbf{z}_t})$ is to the original prediction in latent space. As the cosine similarity decreases, the objective function minimizes the similarity. Consequently, the *Explainer* generates the counterfactual explanation mask to ensure that the two representations or predictions are as far as possible in the audio token space after counterfactual perturbation.

Objective Function for AudioGenX

Along with factual and counterfactual explanation objective functions, we add the regularization term to generate the explanation mask in a simple and efficient manner. Therefore,

Algorithm 1: AudioGenX

Input: Textual token representation vector \mathbf{U} , previously generated audio token vector \mathbf{z}_{t-1} , Transformer model f , audio generation length T , number of epochs K , learning rate λ , regularization coefficients α and β

for $t = 1$ **to** T **do**

Initialize *Explainer* with random parameters.

for $epoch = 1$ **to** K **do**

$\mathbf{M}_{\mathbf{U}, \mathbf{z}_t} = \text{Explainer}(\mathbf{U}, \mathbf{z}_{t-1})$

$L = L_F + L_{CF} + \alpha \cdot L_1(\mathbf{M}_{\mathbf{U}, \mathbf{z}_t}) + \beta \cdot L_2(\mathbf{M}_{\mathbf{U}, \mathbf{z}_t})$

$\theta := \theta - \lambda \nabla_{\theta} L$

end for

$\mathbf{M}_{\mathbf{U}, \mathbf{z}_t} = \text{Explainer}(\mathbf{U}, \mathbf{z}_{t-1})$

end for

Return $\mathbf{M}_{\mathbf{U}, \mathbf{z}} = \frac{1}{T} \sum_{t=1}^T \mathbf{M}_{\mathbf{U}, \mathbf{z}_t}$

we incorporate additional regularization in our final objective function for the *Explainer*, which is formulated as:

$$\mathcal{L} = \mathcal{L}_F + \mathcal{L}_{CF} + \alpha \cdot L_1(\mathbf{M}_{\mathbf{U}, \mathbf{z}_t}) + \beta \cdot L_2(\mathbf{M}_{\mathbf{U}, \mathbf{z}_t}). \quad (6)$$

Here, L_1 and L_2 represent the L_1 -Norm and L_2 -Norm, respectively, as regularization terms to minimize the mask size. This prevents a trivial solution where the *Explainer* generates an explanation mask assigning equal importance to all values. At the same time, adhering to Occam’s Razor principle, we favor simpler and more effective explanations (Tan et al. 2022; Blumer et al. 1987). Hence, according to the objective function in Equation (6), AudioGenX optimizes the *Explainer* generating faithful explanation masks in the audio-token level.

Providing Audio-Level Explanations

In this section, we aggregate audio token-level explanations to provide a comprehensive understanding of the entire audio sequence. The aggregation is performed by averaging the mask values across all audio tokens as follows:

$$\mathbf{M}_{\mathbf{U}, \mathbf{z}} = \frac{1}{T} \sum_{t=1}^T \mathbf{M}_{\mathbf{U}, \mathbf{z}_t}, \quad (7)$$

where t refers to the step, and T represents the total length of generated audio. Additionally, it is possible to focus on a specific interval of interest within the audio, defined between a starting step s and an ending step n . This is denoted as $\mathbf{M}_{\mathbf{U}, \mathbf{z}} = \frac{1}{|n-s|+1} \sum_{t=s}^n \mathbf{M}_{\mathbf{U}, \mathbf{z}_t}$, which provides granular explanations based on the user’s request. This flexible approach enables users to discover patterns within specific intervals, as AudioGenX can effectively capture and explain auditory content in targeted regions of the audio sequence.

Experimental Setup

Dataset. We use AudioCaps (Kim et al. 2019) as the source of textual prompts. For each prompt, we generate a 5-second

audio clip using AudioGen, pairing each prompt with its corresponding generated audio. For hyperparameter tuning, we select 100 validation captions, while the test dataset consists of 1,000 randomly selected captions.

Evaluation Metrics. We evaluate explanations based on two metrics: Fidelity and KL divergence, both derived from the classification probabilities of a pre-trained audio classifier. Specifically, we utilize PaSST (Cai et al. 2022), a classifier trained on the AudioSet dataset, which is also used in the evaluation of AudioGen. Its classification probabilities are likely to provide meaningful insights into the relationship between textual prompts and generated audio. Fidelity (Yuan et al. 2021; Ali et al. 2023), a core evaluation metric in the field of XAI, measures the change in top-1 label prediction probabilities of the generated audio after applying factual and counterfactual explanation masks, denoted as Fid_F and Fid_{CF} , respectively.

In addition, KL divergence (Kilgour et al. 2018), originally used to evaluate audio generative models (Kreuk et al. 2023; Yang et al. 2023; Huang et al. 2023), measures the differences of label distribution between generated and reference audio. For explanation evaluation, we introduce new metrics KL_F and KL_{CF} , which measure the conceptual change in the generated audio after applying explanation masks in factual and counterfactual reasoning, respectively. In factual evaluation, the generated audio should closely match the original audio, making lower values Fid_F and KL_F desirable. In contrast, in counterfactual evaluation, higher values of Fid_{CF} and KL_{CF} indicate a more effective explanation. Additionally, we include the average mask size as part of our evaluation.

Baselines. We compare our method with five baselines. Random-Mask is a mask with randomly assigned values ranging between 0 and 1. Grad-CAM (Selvaraju et al. 2017) is evaluated in two variations: Grad-CAM-a and Grad-CAM-e. Specifically, Grad-CAM-a computes the gradients of the latent representation vector of the t -th audio token \mathbf{e}_t with respect to the generated audio sequence \mathbf{z}_t , while Grad-CAM-e computes the gradients of the last cross-attention map to the \mathbf{z}_t . We also include the AtMan (Deiseroth et al. 2023) and the method proposed by Chefer et al. (Chefer, Gur, and Wolf 2021) as baselines.

Experimental Setting. The *Explainer* model includes a linear layer that reduces the text token embeddings from 1536 to 512 dimensions, followed by a PReLU activation function. The 512-dimensional text token embeddings are then mapped to a single value through another linear layer and a sigmoid function, producing a value in the $[0, 1]$ range. A Gumbel-Softmax function is subsequently applied to push values closer to 0 or 1, representing the importance of each text token. The *Explainer* is trained for 50 epochs with a learning rate as $\times 10^{-3}$ using the Adam optimizer. Hyperparameters are set as $\alpha = 1 \times 10^{-3}$ and $\beta = 1 \times 10^{-1}$ as coefficients for the explanation objective function. Hyperparameter sensitivity analysis and detailed experimental settings are both provided in the Appendix. Our code is available at the following link ¹.

¹<https://github.com/hjknng/audiogenX>

Method	$Fid_F \downarrow$	$Fid_{CF} \uparrow$	$KL_F \downarrow$	$KL_{CF} \uparrow$	Size \downarrow
$N_{audio} = 5$	0.128 \pm 0.004	-	1.318 \pm 0.030	-	-
Random-Mask	0.196 \pm 0.004	0.195 \pm 0.006	1.884 \pm 0.044	1.932 \pm 0.046	0.500
Grad-CAM-e	0.204 \pm 0.006	0.235 \pm 0.008	1.858 \pm 0.034	2.457 \pm 0.041	0.422
Grad-CAM-a	0.240 \pm 0.006	0.192 \pm 0.010	2.285 \pm 0.077	1.951 \pm 0.075	0.406
AtMan	0.195 \pm 0.008	0.222 \pm 0.008	2.010 \pm 0.049	2.198 \pm 0.048	0.497
Chefer et al.	0.198 \pm 0.003	0.229 \pm 0.004	1.899 \pm 0.025	2.348 \pm 0.040	0.441
AudioGenX w/ Eq. (3)	0.145 \pm 0.004	0.360 \pm 0.005	1.542 \pm 0.024	3.658 \pm 0.061	0.360
AudioGenX w/ Eq. (5)	0.143 \pm 0.004	0.385 \pm 0.005	1.514 \pm 0.043	3.977 \pm 0.044	0.385
AudioGenX w/ Eq. (7)	0.137 \pm 0.005	0.402 \pm 0.005	1.418 \pm 0.043	4.183 \pm 0.073	0.455
AudioGenX	0.132 \pm 0.004	0.405 \pm 0.004	1.416 \pm 0.029	4.259 \pm 0.039	0.455

Table 1: Evaluation of explanations generated by each method using factual and counterfactual reasoning. Five audio samples are generated and evaluated with different seeds based on the obtained explanations. The best results are highlighted in **bold**.

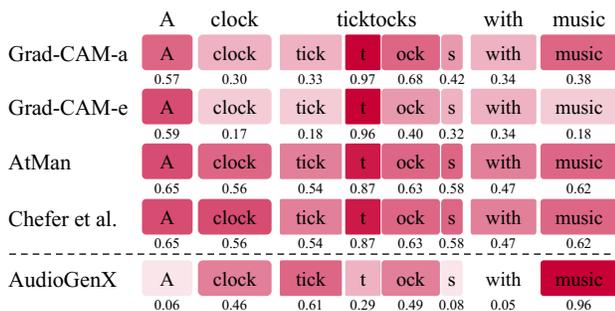


Figure 3: Visualization of AudioGenX and other methods.

Experimental Results

RQ 1: Does AudioGenX Generate Faithful Explanations?

We evaluate the generated explanations by AudioGenX based on factual and counterfactual reasoning, as presented in Table 1. AudioGenX achieves the best performance across the metrics Fid_F , Fid_{CF} , KL_F , and KL_{CF} , while also maintaining the smallest size ($Size$), demonstrating that our explanations are both simple and effective. The baseline, denoted as $N_{audio} = 5$, generates audio conditioned on the same textual input five times to observe the inherent variance, serving as the lower bound for Fid_F , KL_F . AudioGenX’s factual audio nearly reaches the lower bound, indicating high performance. Furthermore, significant changes in Fid_{CF} and KL_{CF} under counterfactual perturbations confirm that the explanations are both sufficient and necessary. The AudioGenX with factual and counterfactual losses in Eq.(6), outperforms the variants AudioGenX w/ Eq. (3) and AudioGenX w/ Eq. (5), which apply only factual or counterfactual loss with a regularization term. This indicates that the two losses complement each other, enhancing overall performance. Furthermore, we evaluate AudioGenX w/ Eq. (7) using an averaged explanation mask, showing the robustness of explainability in describing the entire audio. In contrast, other baselines fail to generate meaningful counterfactual audio, lacking the optimization properties needed to enforce counterfactual explanations.

The strong performance highlights the effectiveness of leveraging latent embedding vectors to generate explanations. While most baselines are designed to explain supervised learning models, they rely on vectors that represent the probability distribution of the final audio token. This approach, however, does not align well with the inference process of audio generation models. In extreme cases, such as top- k sampling ($k=250$), the 250-th audio token could be sampled, leading to significant discrepancies between the gradients or probability-related information the token most likely predicted by the model. In contrast, our approach avoids dependency on the sampling process, allowing the model to produce more faithful explanations.

RQ 2: How Well Do the Explanations from AudioGenX Reflect the Generated Audio?

We visualize the explanations generated by AudioGenX and other baselines, as shown in Figure 3. AudioGenX demonstrates a clear advantage in focusing on key audio elements. Unlike other baselines, which often assign relatively high importance scores to less important tokens like ‘A’ and ‘with’, AudioGenX consistently assigns higher importance scores to crucial tokens such as ‘ticktocks’ and ‘music’. For instance, AudioGenX assigns a notably high importance score of 0.96 to ‘music’, emphasizing its ability to focus on significant input tokens. In contrast, other models like Grad-CAM-e and AtMan distribute importance more broadly, including less relevant tokens. These results show that AudioGenX consistently provides faithful explanations, aligning the generated audio with the essential components of the input text.

Furthermore, when generating audio from a prompt containing multiple concepts, some words may be less prominently reflected. In such case, AudioGenX provides adequate explanations for each specific audio, indicating whether each word from the prompt has been incorporated into the generated audio. As illustrated in Figure 4, the difference between the two audios is that bird sounds are present in Figure 4-(a) but absent in Figure 4-(b). AudioGenX effectively describes the audios by assigning high importance scores of 0.98 and 0.99 to the token ‘Water,’ which is the primary sound in both audios. AudioGenX assigns a

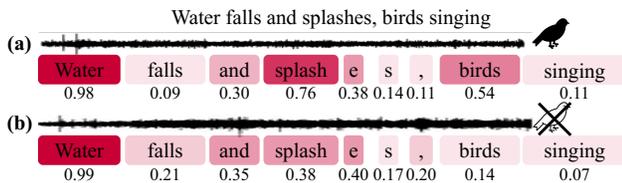


Figure 4: Explanation generated by AudioGenX for two audios created from a single prompt. (a) includes bird sounds, while (b) does not.

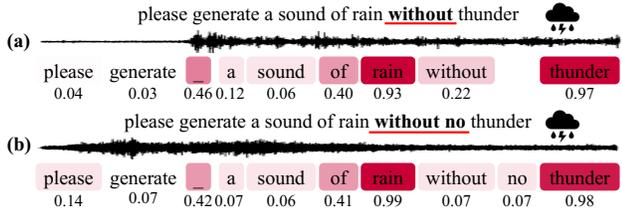


Figure 5: Explanations generated from negated prompts: (a) single negation, (b) double negation.

score of 0.54 to ‘birds,’ while it assigns a score of 0.14, accurately reflecting the different audio characteristics in each case. These results show that AudioGenX can provide explanations that are well-suited to the corresponding audio. Furthermore, these explanations serve as valuable insights for editing generated audio to better align with user intention.

RQ 3: How Can Explanations Help Understand AudioGen Behavior?

We explore the output patterns of AudioGen using the explanations generated by AudioGenX. First, we investigate whether AudioGen can effectively handle sentences containing negations and double negations, as shown in Figure 5. The explanations of the generated audios are presented in response to input prompts containing ‘without thunder’ and ‘without no thunder.’ In both cases, the generated audio includes the sound of thunder along with the rain. Using AudioGenX, we observe that ‘without’ and ‘without no’ have lower importance compared to ‘thunder’ in the explanations. We hypothesize that this occurs because the training dataset lacks sufficient examples of negation and double negation. An examination of the AudioCaps dataset reveals a scarcity of such cases. Additionally, by aggregating tokens from the explanations, we identify the top and bottom 50 tokens in Table 3 in the Appendix. Tokens with high importance are predominantly nouns, such as ‘thunder,’ while those with low importance include sound descriptors like ‘distant,’ as well as sequential expressions like ‘before.’ Such analyses could be used to debug TAG models or to identify potential inherent biases in their behavior.

RQ 4: Does AudioGenX Generate Explanations Efficiently?

We evaluate the efficiency of explanation methods based on the average time and total GPU memory usage per explana-

Method	Memory (MB)↓	Time (s)↓
Grad-CAM-e	8641.306	49.038
Grad-CAM-a	41655.848	62.276
AtMan	5081.957	7.295
Chefer et al.	41684.969	52.166
AudioGenX w/ Eq. (3)	11980.894	36.639
AudioGenX w/ Eq. (5)	11981.114	37.373
AudioGenX w/ Eq. (7)	12001.931	63.198
AudioGenX	12001.931	63.198

Table 2: Efficiency analysis of AudioGenX and other baseline methods. The best results are highlighted in **bold**.

tion, as shown in Table 2. For GPU memory efficiency, the results rank in the following order: AtMan, Grad-CAM-e, AudioGenX, Grad-CAM-a, and Chefer et al. For time efficiency, the order is AtMan, Grad-CAM-e, Chefer et al., Grad-CAM-a, and AudioGenX. Although AtMan is the most efficient, its performance remains subpar due to its simplistic approach. Grad-CAM-e demonstrates greater memory efficiency compared to Grad-CAM-a and Chefer et al., as it tracks a shallower layer. While AudioGenX requires additional computational time to train explanation masks, it achieves memory efficiency by reducing GPU storage and operates with $\mathcal{O}(Lk)$ complexity, ensuring linear scalability for large-scale tasks.

Conclusion

AudioGenX quantifies the importance of textual tokens corresponding to generated audio by leveraging both factual and counterfactual reasoning frameworks. This approach enables the generation of faithful explanations, providing actionable insights for users to edit audio and assisting developers in debugging. Consequently, AudioGenX enhances the transparency and trustworthiness of TAG models. For comprehensive details on experimental settings, additional results, and hyperparameter sensitivity analyses, please refer to the Appendix at <https://learndatalab.github.io/audiogenx.html>.

Acknowledgements

This work was supported by NCSOFT, the Institute of Information & Communications Technology Planning & evaluation (IITP) grant, and the National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIT) (RS-2019-II190421, IITP-2025-RS-2020-II201821, RS-2024-00438686, RS-2024-00436936, RS-2024-00360227, RS-2023-0022544, NRF-2021M3H4A1A02056037, RS-2024-00448809). This research was also partially supported by the Culture, Sports, and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture, Sports and Tourism in 2024 (RS-2024-00333068, RS-2024-00348469 (25%)).

References

- Abnar, S.; and Zuidema, W. 2020. Quantifying attention flow in transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4190–4197. Association for Computational Linguistics.
- Akman, A.; and Schuller, B. W. 2024. Audio Explainable Artificial Intelligence: A Review. *Intelligent Computing*, 2: 0074.
- Ali, S.; Abuhmed, T.; El-Sappagh, S.; Muhammad, K.; Alonso-Moral, J. M.; Confalonieri, R.; Guidotti, R.; Del Ser, J.; Díaz-Rodríguez, N.; and Herrera, F. 2023. Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence. *Information fusion*, 99: 101805.
- Becker, S.; Ackermann, M.; Lopuschkin, S.; Müller, K.-R.; and Samek, W. 2018. Interpreting and explaining deep neural networks for classification of audio signals. *arXiv preprint arXiv:1807.03418*.
- Blumer, A.; Ehrenfeucht, A.; Haussler, D.; and Warmuth, M. K. 1987. Occam’s razor. *Information Processing Letters*, 24(6): 377–380.
- Cai, J.; Fan, J.; Guo, W.; Wang, S.; Zhang, Y.; and Zhang, Z. 2022. Efficient deep embedded subspace clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1–10.
- Chefer, H.; Gur, S.; and Wolf, L. 2021. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 397–406.
- Défossez, A.; Copet, J.; Synnaeve, G.; and Adi, Y. 2023. High Fidelity Neural Audio Compression. *Transactions on Machine Learning Research*.
- Deiseroth, B.; Deb, M.; Weinbach, S.; Brack, M.; Schramowski, P.; and Kersting, K. 2023. AtMan: Understanding transformer predictions through memory efficient attention manipulation. *Advances in Neural Information Processing Systems*, 36.
- Frommholz, A.; Seipel, F.; Lopuschkin, S.; Samek, W.; and Vielhaben, J. 2023. XAI-based Comparison of Input Representations for Audio Event Classification. In *Proceedings of the International Conference on Content-Based Multimedia Indexing*, 126–132.
- Haunschmid, V.; Manilow, E.; and Widmer, G. 2020. audioLIME: Listenable Explanations Using Source Separation. *Proceedings of the International Workshop on Machine Learning and Music*. arXiv:2008.00582.
- Huang, R.; Huang, J.; Yang, D.; Ren, Y.; Liu, L.; Li, M.; Ye, Z.; Liu, J.; Yin, X.; and Zhao, Z. 2023. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. In *Proceedings of the International Conference on Machine Learning*, 13916–13932.
- Jang, Eric; Gu; Shixiang; Poole; and Ben. 2017. Categorical reparameterization with gumbel-softmax. In *Proceedings of the International Conference on Learning Representations*.
- Kenny, E. M.; Delaney, E. D.; Greene, D.; and Keane, M. T. 2021. Post-hoc explanation options for XAI in deep learning: The Insight centre for data analytics perspective. In *Proceedings of the ICPR International Workshops and Challenges*, 20–34.
- Kilgour, K.; Zuluaga, M.; Roblek, D.; and Sharifi, M. 2018. Fréchet Audio Distance: A Metric for Evaluating Music Enhancement Algorithms. *arXiv preprint arXiv:1812.08466*.
- Kim, C. D.; Kim, B.; Lee, H.; and Kim, G. 2019. Audiocaps: Generating captions for audios in the wild. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 119–132.
- Kreuk, F.; Synnaeve, G.; Polyak, A.; Singer, U.; Défossez, A.; Copet, J.; Parikh, D.; Taigman, Y.; and Adi, Y. 2023. Audiogen: Textually guided audio generation. In *Proceedings of the International Conference on Learning Representations*.
- 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 *Proceedings of the International Conference on Machine Learning*, 21450–21474.
- Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.
- Nagahisarchoghaei, M.; Nur, N.; Cummins, L.; Nur, N.; Karimi, M. M.; Nandanwar, S.; Bhattacharyya, S.; and Rahimi, S. 2023. An empirical survey on explainable ai technologies: Recent trends, use-cases, and categories from technical and application perspectives. *Electronics*, 12(5): 1092.
- Pearl, J. 2009. Causal inference in statistics: An overview. *Statistics Surveys*, 3: 96 – 146.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140): 1–67.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. “Why should I trust you?” Explaining the predictions of any classifier. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.
- Samek, W.; Wiegand, T.; and Müller, K.-R. 2017. Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *arXiv preprint arXiv:1708.08296*.
- Schneider, F.; Kamal, O.; Jin, Z.; and Schölkopf, B. 2023. Moûsai: Text-to-music generation with long-context latent diffusion. *arXiv preprint arXiv:2301.11757*.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, 618–626.

- Sundararajan, M.; Taly, A.; and Yan, Q. 2017. Axiomatic attribution for deep networks. In *Proceedings of the International Conference on Machine Learning*, 3319–3328.
- Tan, J.; Geng, S.; Fu, Z.; Ge, Y.; Xu, S.; Li, Y.; and Zhang, Y. 2022. Learning and evaluating graph neural network explanations based on counterfactual and factual reasoning. In *Proceedings of the ACM Web Conference*, 1018–1027.
- 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.
- Yang, D.; Yu, J.; Wang, H.; Wang, W.; Weng, C.; Zou, Y.; and Yu, D. 2023. Diffsound: Discrete diffusion model for text-to-sound generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31: 1720–1733.
- Yuan, H.; Yu, H.; Wang, J.; Li, K.; and Ji, S. 2021. On explainability of graph neural networks via subgraph explorations. In *Proceedings of the International Conference on Machine Learning*, 12241–12252.
- Ziv, A.; Gat, I.; Lan, G. L.; Remez, T.; Kreuk, F.; Défossez, A.; Copet, J.; Synnaeve, G.; and Adi, Y. 2024. Masked Audio Generation using a Single Non-Autoregressive Transformer. In *Proceedings of the International Conference on Learning Representations*.