

# Taming Data and Transformers for Audio Generation

Moayed Haji-Ali<sup>1,2†</sup>, Willi Menapace<sup>2</sup>, Aliaksandr Siarohin<sup>2</sup>,  
Guha Balakrishnan<sup>1</sup>, Vicente Ordonez<sup>1\*</sup>

<sup>1</sup>Rice University.

<sup>2</sup>Snap Inc.

\*Corresponding author(s). E-mail(s): [vicenteor@rice.edu](mailto:vicenteor@rice.edu);  
Contributing authors: [mh155@rice.edu](mailto:mh155@rice.edu); [wmenapace@snapchat.com](mailto:wmenapace@snapchat.com);  
[asiarohin@snapchat.com](mailto:asiarohin@snapchat.com); [guha@rice.edu](mailto:guha@rice.edu);

†Work partially done during an internship at Snap Inc.

## Abstract

The scalability of ambient sound generators is hindered by data scarcity, insufficient caption quality, and limited scalability in model architecture. This work addresses these challenges by advancing both data and model scaling. First, we propose an efficient and scalable dataset collection pipeline tailored for *ambient* audio generation, resulting in AutoReCap-XL, the largest ambient audio-text dataset with over *47 million* clips. To provide high-quality textual annotations, we propose AutoCap, a *high-quality* automatic audio captioning model. By adopting a Q-Former module and leveraging audio metadata, AutoCap substantially enhances caption quality, reaching a CIDEr score of **83.2**, a **3.2%** improvement over previous captioning models. Finally, we propose GenAu, a scalable transformer-based audio generation architecture that we scale up to 1.25B parameters. We demonstrate its benefits from data scaling with synthetic captions as well as model size scaling. When compared to baseline audio generators *trained at similar size and data scale*, GenAu obtains significant improvements of **4.7%** in FAD score, **11.1%** in IS, and **13.5%** in CLAP score. Our code, model checkpoints, and dataset are *publicly available*.

**Keywords:** audio-video dataset, audio generation, audio captioning

## 1 Introduction

Text-conditioned generative models have revolutionized the field of content creation, enabling the generation of high-quality natural images (Ramesh et al. 2022; Rombach et al. 2022; Podell et al. 2023; Haji-Ali et al. 2024a), vivid videos (Ho et al. (2022); Villegas et al. (2022); Wang et al. (2024c); Qiu et al. (2023); Menapace et al. (2024)), and intricate 3D shapes (Cheng et al. 2023). The domain of audio synthesis has undergone comparable advancement (Huang

et al. 2023b,a; Liu et al. 2023a; Xue et al. 2024; Guan et al. 2024; Saito et al. 2024; Niu et al. 2024; Yang et al. 2023a; Evans et al. 2024a; Liu et al. 2024b; Wang et al. 2024d; Guo et al. 2024), with three broad areas of study: speech, music and ambient sounds. The success in these domains rests on two key pillars: (i) the availability of high-quality large-scale datasets with text annotations, and (ii) the development of scalable generative models (Ho et al. 2020; Song et al. 2021). The objective of this work is to improve

audio generation quality by scaling *ambient* audio generators across both the data and model axes.

In the field of audio synthesis, ambient audio generation emerges as a critical domain. Unlike speech and music, ambient sound generation is particularly challenging due to the lack of extensive, well-annotated datasets (Kim et al. 2019; Drossos et al. 2020). Attempts to curate ambient audio from online videos predominantly failed, primarily due to the dominance of speech and music content in such videos. For instance, AudioSet (Gemmeke et al. 2017), the largest available audio dataset sourced from online videos, contains 99% speech or music clips. Previous efforts to filter ambient audio from similar datasets involved using expensive classifiers on the video or audio content, making it impractical to compile a large-scale dataset due to the high filtering rate. In this work, we propose a simple, yet scalable filtering approach that leverages existing automatic video transcription to identify segments with ambient sounds. This method is not only more efficient but also more feasible, as it eliminates the need to download audio or video content. Through this approach, we built AutoReCap-XL, a dataset containing 47 million ambient audio clips sourced from existing video datasets, representing a 75-fold increase over the size of previously largest available datasets.

Another challenge in compiling large-scale text-audio datasets is providing accurate textual descriptions. For visual modalities, such as images and videos (Xue et al. 2022; Miech et al. 2019), researchers often relied on the raw description and metadata to train strong visual-text models including reliable captioners (Chen et al. 2024b). For ambient sounds, however, the task is substantially more challenging as accompanying raw text tends to describe visual information or convey feelings, rather than detailing the audio content. Moreover, human-captioned audio datasets are limited, containing fewer than 51k text-audio pairs in total. This significantly impacts the training of current captioning models, making them more susceptible to overfitting. To address this, we introduce AutoCap, a high-quality audio captioner that leverages visual cues to enhance captioning.

AutoCap refines the commonly used encoder-decoder design based on a pretrained

BART (Lewis et al. 2020) by introducing a Q-Former (Li et al. 2023) module which learns an intermediate representation that better aligns the encoded audio and the original BART token representation. Second, we propose to remedy the data scarcity problem by using metadata and visual cues to aid the captioning process. Critically, we augment the encoder inputs with a set of descriptive textual metadata including audio title and a caption derived from the visual modality. This dual-input approach achieves improved performance over baselines on AudioCaps (Kim et al. 2019), marking a 3.2% improvement in CIDEr score. Using AutoCap, we provide textual descriptions for AutoReCap-XL and demonstrate the benefits of scaling audio generative models with synthetic captions.

Another axis for scaling generative models is model size Peebles and Xie (2023). While scaling diffusion backbones has shown significant benefits in image and video generation, ambient audio generation has shown poor scaling behavior. For instance, AudioLDM2 Liu et al. (2023b), reported worse metrics for their largest model compared to the smaller one. Similarly, EzAudio Hai et al. (2024), achieved only marginal improvements with scaling model size. In this work, we introduce GenAu, a *scalable* transformer-based architecture that achieves *significant* improvements over state-of-the-art models. We recognize that audio grows fast temporally, yet contains many silent and redundant segments. Therefore, an efficient architecture that can handle such properties is needed. In particular, we employ a transformer architecture in the denoising backbone where we modify the FIT transformer (Chen and Li 2023) to generate audio in the latent space. On AutoCaps dataset, GenAu achieves significant improvements over baselines, with 11.1% higher Inception Score, 4.7% better FAD, and 13.5% improvement in CLAP score, demonstrating superior audio-text alignment and generation quality. Moreover, GenAu shows promising scaling properties, with consistent improvements across all metrics as model size increases.

In summary, this work presents significant contributions in three areas: (i) AutoCap, a novel audio captioner tailored towards the annotation of data at a large scale that uses visual clues and

audio metadata to improve accuracy and robustness; (ii) AutoReCap-XL, a large scale ambient audio dataset, comprising 47M audio clips paired with synthetic captions, 75 times larger than available datasets (iii) GenAu, a novel audio generator based on a scalable transformer architecture specifically adapted to the audio domain, achieving significant improvements over previous state-of-the-art.

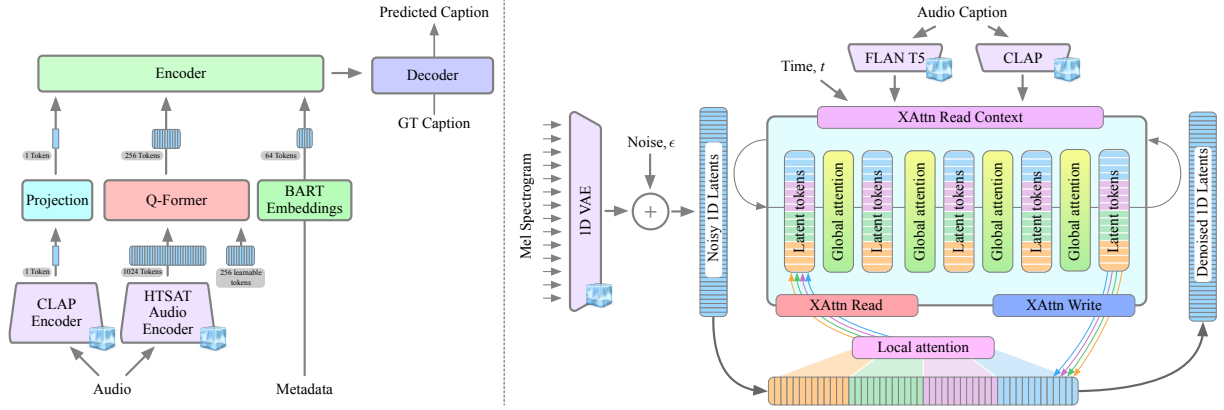
## 2 Related Work

**Automatic Audio Captioning (AAC).** The goal of AAC is to produce language descriptions for given audio content. Recent AAC methods (Deshmukh et al. 2024a; Gontier et al. 2021; Wu et al. 2024; Salewski et al. 2023; Sridhar et al. 2024; Kadlčík et al. 2023; Cousin et al. 2023; Labbé et al. 2023; Xu et al. 2024b; Zhang et al. 2024; Ghosh et al. 2024a; Deshmukh et al. 2024a; Chen et al. 2025; Rho et al. 2025) employ encoder-decoder transformer architectures, where an encoder receiving the audio signal produces a representation that is used by the decoder to produce the output caption. WavCaps (Mei et al. 2024a) employs HTSAT (Chen et al. 2022) audio encoder and uses a pretrained BART (Lewis et al. 2020) as the decoder. Similarly, EnCLAP (Kim et al. 2024c) uses a pretrained BART and improves on the audio representation. EnCLAP++ (Kim et al. 2024b) optimized the encoder selection of EnCLAP to further improve performance. Recognizing the limited data, CoNeTTE (Labb et al. 2024) proposes to train a lightweight vanilla transformer decoder (Vaswani et al. 2017) instead. Other work explored augmentation to counter data scarcity (Kim et al. 2022; Labb et al. 2024; Ye et al. 2022). Recent work (Liu et al. 2023c; Sun et al. 2024; Yuan et al. 2024) proposed to leverage visual information to address sound ambiguities, reporting improvements. More recent methods use audio large language model for zero-shot captioning (Kong et al. (2024); Deshmukh et al. (2023); Ghosh et al. (2024b)). Our method uses audio metadata and visual information as additional signals and leverages a *lightweight* Q-Former (Li et al. 2023) model to improve accuracy.

**Text-conditioned audio generation.** The current state-of-the-art text-to-audio generation methods widely adopt diffusion models (Yang et al. 2023b; Kreuk et al. 2023; Liu et al. 2023a,b;

Huang et al. 2023a; Ghosal et al. 2023; Evans et al. 2024a; Vyas et al. 2023; Kreuk et al. 2023; Hai et al. 2024; Evans et al. 2024b). AudioLDM 1 & 2 (Liu et al. 2023a,b) make use of a latent diffusion model and employ a UNet as the diffusion backbone. Recently, StableAudio Open (Evans et al. 2024b) introduced a 1.32B model that uses a DiT (Peebles and Xie 2023) to generate variable-length audio clips at 48kHz. Recent work also explored controllable audio generation (Shi et al. 2023; Xu et al. 2024a; Melechovsky et al. 2023; Paissan et al. 2023; Zhang et al. 2023b; Liang et al. 2024; Liu et al. 2024a), visual-conditioned audio generation (Wang et al. 2024d; Mei et al. 2024b; Wang et al. 2024a), and more recently joint audio-video generation (Tang et al. 2024, 2023b; Xing et al. 2024; Hayakawa et al. 2024; Tian et al. 2024; Vahdati et al. 2024; Chen et al. 2024a; Kim et al. 2024a; Wang et al. 2024b; Mao et al. 2024; Haji-Ali et al. 2024b). In this work, we propose a transformer architecture design that shows strong scalability properties.

**Text-Audio Datasets.** The performance of text-audio models (Zhu et al. 2024; Li et al. 2024; Deshmukh et al. 2023; Mahfuz et al. 2023; Deshmukh et al. 2024b; Shu et al. 2023; Elizalde et al. 2024; Liu et al. 2024c; Tang et al. 2023a; Gong et al. 2023; Cheng et al. 2024; Zhang et al. 2023a) is currently hindered by the lack of high-quality large-scale paired audio text data of ambient sounds. Existing human-captioned (AudioCaps (Kim et al. 2019) and Clotho (Drossos et al. 2020)), have in total only 51k audio-text pairs. Another challenge is the limited availability of audio clips from sound-only platforms. LAION-Audio (Wu et al. 2023a) relied on numerous sources of audio platforms such as BBC Sound Effects (BBC Sound Effects 2024), (Font et al. 2013) FreeSounds, and SoundBible (SoundBible 2024) to form a dataset consisting of 630k audio samples with highly noisy raw descriptions. Chen et al. (2020) attempted to extract audio clips from videos by employing classifiers to detect ambient audio, speech, and music. To annotate these clips, WavCaps (Mei et al. 2024a) proposes a filtering procedure based on ChatGPT (Achiam et al. 2023) to collect 400k audio clips and weakly caption them *based on the noisy descriptions alone*. Several subsequent work (Majumder et al. 2024; Sun et al. 2024) adopted similar strategies of using large language models to augment captions. While weak-captioning



**Fig. 1 (Left) Overview of AutoCap.** We employ a frozen HTSAT (Chen et al. 2022) encoder to produce an audio fine-grained representation of 1024 tokens. We then employ a Q-Former (Li et al. 2023) module to produce 256 tokens. These tokens, along with audio CLAP embeddings (Wu et al. 2023a) and 64 tokens of pertinent metadata, are processed by a pretrained BART to generate the final caption. **(Right) Overview of GenAu.** Following latent diffusion models, we use a frozen 1D-VAE to convert a Mel-Spectrogram into latent sequences, which are then divided into groups and processed using ‘local’ attention layers based on the FIT architecture (Chen and Li 2023). ‘Read’ and ‘write’ layers, implemented as cross-attention, facilitate information transfer between input latents and *learnable* latent tokens. Finally, ‘global’ attention layers on *latent tokens* facilitate global communication across groups.

improves downstream metrics, it is *suboptimal* because it fails to incorporate the audio signal itself. In this work, we introduce an efficient dataset collection pipeline that relies on video datasets to extract ambient audio clips and automatic captioners to provide textual descriptions. We collect 47M audio clips, marking the largest available text-audio dataset.

### 3 Method

In this section, we describe our approach to high-quality text-to-audio generation, starting with audio captioning using AutoCap in section 3.1, data collection in section 3.2, and ambient audio generation with GenAu in section 3.3

#### 3.1 Automatic Audio Captioning

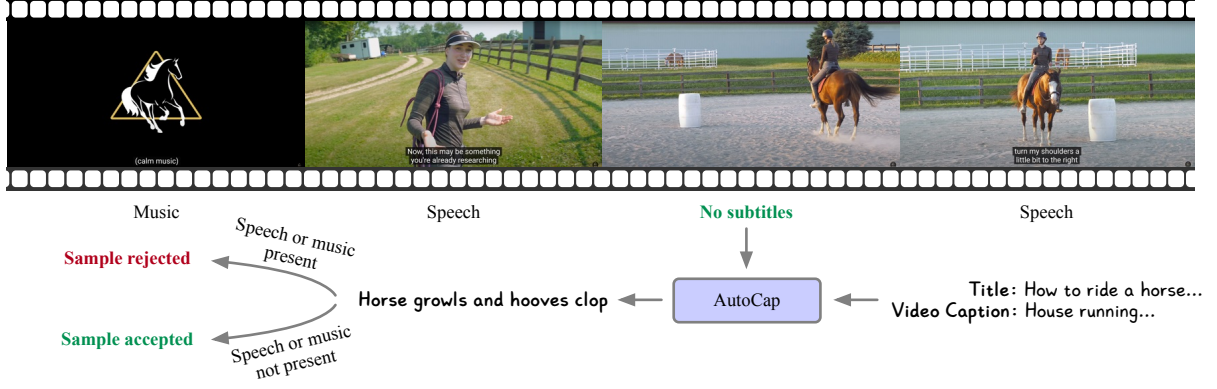
Recent state-of-the-art methods (Labbe et al. 2024; Kim et al. 2024c) generally employ an encoder-decoder transformer design where a pretrained audio encoder passes the audio representation to a pre-trained language model serving as the decoder. This language model (*e.g.* BART) is typically finetuned to adapt to the audio representation. However, due to the distribution mismatch between the pretraining data of the LLM and the audio embeddings produced by the encoder,

the decoder suffers from catastrophic forgetting. Furthermore, audio is an inherently ambiguous modality, as many events can produce similar sound effects—a phenomenon often leveraged in animation, where soundscapes are artificially constructed. Audio clips from many sources, however, are still commonly associated with metadata that might be relevant for captioning such as raw user descriptions, or related modalities (*i.e.* accompanied visual information). Motivated by these observations, we propose AutoCap, an audio captioning model that employs an intermediate audio representation to connect the pretrained encoder and decoder and uses metadata to aid with the captioning. Figure 1 (left) presents an overview of AutoCap.

We consider a dataset of audio-caption pairs  $\langle \mathbf{a}, \mathbf{y} \rangle$  and corresponding metadata represented as a set of token sequences  $\{\mathbf{m}_j\}_{j=1}^{j=M}$ . Inspired by state-of-the-art AAC methods (Mei et al. 2024a; Labbe et al. 2024; Kim et al. 2024c), we employ an encoder-decoder architecture. We first compute a global feature representation of the audio:

$$\mathbf{x}_{\text{clap}} = \mathcal{P}_{\text{clap}}(\mathcal{E}_{\text{clap}}(\mathbf{a})), \quad (1)$$

where  $\mathcal{P}_{\text{clap}}$  is a learnable projection layer and  $\mathcal{E}_{\text{clap}}$  is the audio encoder of a pretrained CLAP model (Wu et al. 2023a). We also compute local



**Fig. 2 Audio data collection pipeline.** We employ online video speech transcripts to identify audio segments without subtitles, which typically correspond to clips lacking speech or music. These are processed by AutoCap to generate captions. As a post-filtering technique for ambinet audio selection, we retain only clips whose captions lack music and speech keywords.

features as:

$$\mathbf{x}_{\text{audio}} = \mathcal{Q}(\mathcal{E}_a(\mathbf{a})), \quad (2)$$

where  $\mathcal{Q}$  is a Q-Former (Li et al. 2023) and  $\mathcal{E}_a$  is a pretrained HTSAT (Chen et al. 2022) audio encoder that produces a *time-aligned representation* (1024 tokens) following Mei et al. (2024a). The Q-Former efficiently learns 256 latent tokens, which serve as keys in cross-attention layers with the input features, thereby outputting 256 tokens. Metadata sequences  $\mathbf{m}_i$  are then embedded using the embedding layer of a pretrained BART to obtain embedding sequences  $\mathbf{x}_{\text{meta}_i}$ . For our experiments, we use video titles and captions as the metadata. We represent the input audio and metadata as the following input sequence:

$$\mathbf{x} = \mathbf{x}_{\text{clap}} [\text{boa}] \mathbf{x}_{\text{audio}} [\text{eoa}] [\text{bom}]_1 \mathbf{x}_{\text{meta}_1} [\text{bom}]_1 \dots [\text{bom}]_M \mathbf{x}_{\text{meta}_M} [\text{bom}]_M, \quad (3)$$

where  $[\text{boa}][\text{eoa}]$  represent beginning and end of audio sequence embeddings  $\mathbf{x}_{\text{audio}}$ , and  $[\text{bom}]_i, [\text{bom}]_i$  represent beginning and end of metadata embeddings  $\mathbf{x}_{\text{meta}_i}$ . The input sequence is passed to a pretrained BART (Lewis et al. 2020)  $\mathcal{D}_t$  to predict a caption as  $\hat{\mathbf{y}} = \mathcal{D}_t(\mathbf{x})$


**Training.** We train our model using a standard cross-entropy loss over next token predictions. To avoid degrading the quality of the pretrained BART and audio encoder models, we adopt a two-stage training procedure. In Stage 1, both the audio encoders and BART model are

kept frozen, thus allowing the Q-Former, projection layers, and newly introduced delimiter tokens to align to the pretraining BART representation. In this stage, we pretrain the model using a larger dataset of weakly-labeled audio clips. In Stage 2, we unfreeze all BART parameters apart from the embedding layer and finetune the model on the Audiocaps dataset at a lower learning rate to align the captioning style more with human style. This training strategy leverages the larger, weakly-labeled dataset while minimizing the knowledge drift in the pretrained BART. The use of Q-Former to learn an intermediate representation is *pivotal* for such a strategy.

### 3.2 Data Collection and Re-captioning Pipeline

Generative models in the image and video domains have shown benefits from increased quantities of data and improved quality of captions. In the audio domain, however, the major human-annotated audio-text datasets, namely AudioCaps (Kim et al. 2019) and Clotho (Drossos et al. 2020), provide only 51k audio clips combined. Previous methods attempted to extract additional ambient audio clips from existing video datasets using pretrained audio classifiers, but a high filtering rate marked this method impractical. Instead, we found that automatic transcripts offer reliable information about the segments containing ambient sounds. In particular, we propose to select the parts of the videos that contain no automatic transcription, suggesting the absence





**Fig. 3** Scaling analysis of model size (left) and data with synthetic captions (right) reveal consistent improvements in FD and IS.

of speech and music. Such an approach offers specific advantages over using pretrained classifiers. Automatic transcripts, readily available for most online videos, eliminate the need to download and process video and audio data before filtering. Additionally, as these transcripts provide precise time-aligned information, they facilitate the extraction of more segments per video. Subsequently, we leverage our AutoCap model to provide textual descriptions of the extracted audio clips. Despite the effectiveness of this method in collecting ambient sounds, some clips still inadvertently contain music or speech due to transcription errors, particularly with speech in less common languages. We address this by analyzing captions and filtering out clips with keywords related to speech or music. Finally, we filter all audio-text pairs with CLAP similarity less than 0.1.

We follow this process to extract 466k audio-text pairs from Audioset (Gemmeke et al. 2017) and VGGSounds (Chen et al. 2020). Additionally, we recaption audio-only datasets such as Freesound, BBC Sound Effects, and SoundBible. To provide metadata, we employ the captioning model of Chen et al. (2024b) to extract a caption whenever video content is available and pass an empty text otherwise. In total, we form AutoReCap, a large-scale dataset comprising of 761,113 audio-text pairs with precise captions. As an additional contribution, we introduce AutoReCap-XL, in which we scale our approach by analyzing four additional large-scale video dataset (Lee et al. 2021; Xue et al. 2022; Zellers et al. 2022; Nagrani et al. 2022) with a total of 71M videos and 715.4k hours. After filtering, we collect and re-caption 47M ambient audio clips spanning 123.5k hours from 20.3M different videos, forming by far the largest available dataset of audio with paired captions. Figure 2 summarizes our data collection pipeline, and Sec. A in

Appendix present more details about the dataset collection and processing and dataset statistics.

### 3.3 Scalable Text-2-Audio Generation

We design our audio generation pipeline, GenAu, as a latent diffusion model. Figure 1 (right) shows an overview of our proposed model. In the following section, we describe in detail the structure of our latent variational autoencoder (VAE) and the latent diffusion model.

**Latent VAE.** Directly modeling waveforms is complex due to the high data dimensionality of audio signals. Instead, we replace the waveform with a Mel-spectrogram representation and use a VAE to further reduce its dimensionality, following prior work (Melechovsky et al. 2023; Huang et al. 2023b). Once generated, Mel-spectrograms can be decoded back to a waveform through a vocoder (Lee et al. 2022). However, commonly-used 2D autoencoder designs (Liu et al. 2023a,b; Melechovsky et al. 2023), are not well suited to the Mel-spectrograms, as the separation between the Mel channels is non-linear, which is not well suited for 2D convolutions. We instead opt for a 1D-VAE design based on 1D convolutions similar to Huang et al. (2023a). We train the VAE following Esser et al. (2021).

**Latent diffusion model.** Following the latent diffusion paradigm, we generate audio by training a diffusion model in the latent space of the 1D-VAE. Transformer-based diffusion models currently attain state-of-the-art performance in audio generation (Huang et al. 2023a). However, both UNet and transformer-based baselines exhibited limited performance gains with increasing model size Liu et al. (2023b); Hai et al. (2024). We observe that ambient audio often contains extensive silent and redundant segments, which may

Model	Pretraining Data	BLEU1	BLEU4	ROUGE <sub>L</sub>	METEOR	CIDEr	SPICE	SPIDER
V-ACT	-	69.8	28.1	49.4	23.7	71.1	17.2	44.2
BART-tags	AS	69.9	26.6	49.3	24.1	75.3	17.6	46.5
AL-MixGEN	-	70.0	28.9	50.2	24.2	76.9	18.1	47.5
GAMA	-	20.4	12.2	19.1	16.3	64.8	5.4	35.1
ENCLAP-Large	-	-	-	-	25.5	80.2	18.8	49.5
HTSAT-BART	-	67.5	27.2	48.3	23.7	72.1	16.9	44.5
HTSAT-BART	AC+CL+WC	70.7	28.3	50.7	25.0	78.7	18.2	48.5
CNext-trans	-	-	-	-	-	-	-	46.6
CNext-trans	AC+CL+MA+WC	-	-	-	25.2	80.6	18.4	49.5
AutoCap (audio)	AC	70.0	28.0	51.7	24.6	77.3	18.2	47.8
AutoCap (audio+text)	AC	72.1	28.6	51.5	<b>25.6</b>	80.0	18.8	49.4
AutoCap (audio)	AC+CL+WC	<b>73.1</b>	28.1	<b>52.0</b>	<b>25.6</b>	80.4	<b>19.0</b>	49.7
AutoCap (audio+text)	AC+CL+WC	72.3	<b>29.7</b>	51.8	25.3	<b>83.2</b>	18.2	<b>50.7</b>

**Table 1** AutoCap results on AudioCaps test split for various models. AS: AudioSet, AC: AudioCaps, WC: WavCaps, CL: Clotho, MA: Multi-Annotator Captioned Soundscapes.

explain the poor scalability of UNet and DiT-based methods, as they distribute computation uniformly across the input. Therefore, we propose to use a more dynamic transformer architecture as a diffusion backbone (Chen and Li 2023; Menapace et al. 2024). In particular, we adopt the FIT architecture of Menapace et al. (2024) which was originally proposed to work in the *pixel space*, and revise it for the *latent space* of the audio modality.

Given a 1D input  $\mathbf{x}$ , we first apply a projection operation to produce a sequence of input patch tokens. We then apply a sequence of FIT blocks to the input patches where each block divides patch tokens into contiguous groups of a predefined size. A set of *local* self-attention layers are then applied separately to each group to avoid the quadratic computational complexity of attention computation. Differently from the video domain (Menapace et al. 2024) where the high input dimensionality makes the *local* layers excessively expensive, we found them to be beneficial for audio generation. To further reduce the amount of computation while maintaining long-range interaction, each block considers a small set of latent tokens. First, a *read* operation implemented as a cross-attention layer transfers information from the patches to the latent tokens. Later, a series of *global* self-attention operations are applied to the latent tokens, allowing information-sharing between different groups. Finally, a *write* operation implemented as a cross-attention layer transfers information from the latent tokens back to the patches. Due to the reduced number of latent tokens when performing

the global self-attention, computational requirements of the model are reduced with respect to a vanilla transformer design (Vaswani et al. 2017). Such a design is particularly suited for the audio modality, which contains mostly silent or redundant parts. Unlike DiT and UNet-based methods (Ronneberger et al. 2015; Peebles and Xie 2023) which allocate the computation resources uniformly across input tokens, the FIT architecture selectively focuses on the more informative parts, dedicating more compute for these parts as the model size scales.

To condition the generation on an input prompt, we use a pretrained FLAN-T5 model (Chung et al. 2024) and a CLAP (Wu et al. 2023a) text encoder to produce the their respective embeddings  $e_{\text{FLAN}}$  and  $e_{\text{CLAP}}$  following prior work Liu et al. (2023b), which we concatenate with the diffusion timestep  $t$  to form the input conditioning signal  $c$ . We insert an additional cross-attention operation inside each FIT block immediately before the ‘read’ operation that makes latent tokens attend to the conditioning. Moreover, we use conditioning on dataset ID to adapt the generation style to different datasets. We train the model using the epsilon prediction objective and follow a linear noise scheduler.

## 4 Experiments

In section 4.1, we evaluate AutoCap quantitatively. We then demonstrate the capabilities of GenAu in section 4.2 and discuss scaling trends with respect to data and model size. We also provide qualitative comparisons on the *Website*.

Model	METEOR $\uparrow$	CIDEr $\uparrow$	SPICE $\uparrow$	SPIDEr $\uparrow$
Ours	<b>25.3</b>	<b>83.2</b>	18.2	<b>50.7</b>
- w/o CLAP	<b>25.3</b>	80.7	<b>18.4</b>	49.6
- w/o Stage 2	24.2	75.6	17.3	46.5
- w/o Stage 1	22.6	59.6	15.4	37.5
- Unfreeze Word Embed	22.5	82.6	18.1	50.4

**Table 2** AutoCap ablation study on AudioCaps

## 4.1 Automatic Audio Captioning

**Training dataset and details.** We train AutoCap in two stages. During stage 1, we pretrain on a large weakly labeled dataset of 634,208 audio clips, constructed from AudioSet, Freesound, BBC Sound Effects, SoundBible, AudioCaps, and Clotho. We use ground truth captions from AudioCaps and Clotho datasets, WavCaps captions for Freesound, SoundBible, and BBC Sound Effects, and handcrafted captions through a template leveraging the ground truth class labels for AudioSet. As metadata, we use the title provided with each clip, and pre-compute video captions using a pretrained Panda70M model (Chen et al. 2024b) or pass an empty string when the video modality is unavailable. We pre-train the model for 20 epochs with a learning rate of  $1e-4$ , while keeping the audio encoder and pretrained BART frozen. In Stage 2, we fine-tune the model for 20 epochs on AudioCaps using a learning rate of  $1e-5$ . We use 10-second clips at 32KHz for all experiments.

**Baselines.** We compare with V-ACT (Liu et al. 2023c), BART-tags (Gontier et al. 2021), AL-MixGEN (Kim et al. 2022), ENCLAP (Kim et al. 2024c), HTSAT-BART (Xu et al. 2024b), CNext-trans (Labb et al. 2024) and GAMA Ghosh et al. (2024b). Among these, ENCLAP and CNext-trans have the best performance. ENCLAP benefits from a stronger audio encoder and a CLAP representation. CNext-trans trains a lightweight transformer instead of fine-tuning a pretrained language model to reduce overfitting.

**Metrics and evaluation.** We report results using the established BLEU1 and BLEU4 (Papineni et al. 2002), ROUGE (Lin 2004), Meteor (Lavie and Agarwal 2007), CIDEr (Vedantam et al. 2015), and SPIDEr (Liu et al. 2017) metrics. We evaluate our method on the AudioCaps test split using the last checkpoint of our trained model. We follow the same evaluation pipeline as baselines and include their reported results, except for GAMA which we evaluate using

their released checkpoint. Metrics unavailable in these publications are excluded from our analysis.

**Results.** Tab. 1 reports that our method outperforms baselines on all metrics, achieving notable improvements in CIDEr (83.2) and BLUE1 (73.1) scores. Notably, even without metadata (*i.e. audio only*), AutoCap surpasses baselines in most metrics. We found that incorporating metadata significantly enhances CIDEr but slightly reduces SPICE. This trade-off likely results from the enhanced descriptive detail brought by the metadata, which while enriching the content, introduces noise that may compromise the model’s semantic precision. In addition, AudioCaps is labeled based on audio alone. Thus, the evaluation penalizes the description of information that can not be deduced with certainty from the audio modality only, such as the specific type of object producing the sound. Qualitatively, our captions are more detailed and temporally accurate than baselines. ENCLAP-Large often misses key details. CNext-trans, while accurate, often produces short captions that lack details. We include qualitative comparisons in the *Website* and *Appendix*. Moreover, AutoCap is *four times* faster than ENCLAP, producing a caption for a 10-second clip in *0.28* seconds, compared to ENCLAP which takes *1.12* seconds. Furthermore, we observe consistent improvements when pretraining on weakly-labeled data, validating the effectiveness of our training strategy in benefiting from larger, weakly-labeled datasets.

**Ablations.** In Tab. 2, we ablate model design choices. Using CLAP embedding brings a 2.5 points increase in CIDEr. Omitting Stage 2 training, which involves finetuning BART (Lewis et al. 2020), results in performance degradation, likely due to the necessity of adapting BART’s decoder to the sentence structure typical of AudioCaps. A more severe degradation in performance is observed when Stage 1 is not performed, as the misaligned representation between the encoder and the decoder leads to catastrophic forgetting in the language model. Finally, finetuning BART word embeddings in Stage 2 reduces performance.

## 4.2 Text-2-Audio Generation

**Training dataset and details.** We follow baselines Liu et al. (2023b); Huang et al. (2023a); Majumder et al. (2024) and train on 10-second



Model	Prms	# Samples	FD ↓	IS ↑	FAD ↓	CLAP <sub>LAION</sub> ↑	CLAP <sub>MS</sub> ↑
GroundTruth	-	-	-	-	-	0.251	0.671
AudioLDM-L	739M	634k	37.89	7.14	5.86	-	0.429
AudioLDM 2-L	712M	760k	32.50	8.54	5.11	0.212	0.621
TANGO	866M	45k	26.13	8.23	1.87	0.185	0.597
TANGO 2	866M	60k	19.77	8.45	2.74	0.264	0.590
Make-An-Audio	453M	1M	27.93	7.44	2.59	0.207	0.621
Make-An-Audio 2	937M	1M	15.34	9.58	1.27	0.251	0.645
Stable Audio Open	1.32B	486K	21.23	10.48	2.32	0.246	0.584
GenAu-Large	1.25B	811K	16.51	11.75	<b>1.21</b>	0.285	0.668
GenAu-Large-Full	1.25B	19.6M	<b>14.48</b>	<b>12.56</b>	1.71	<b>0.318</b>	<b>0.677</b>

**Table 3** GenAu results on AudioCaps test split.

clips at 16kHz resolution. We use a patch size of 1 and a group size of 32. We use LAMB optimizer (You et al. 2019) with a LR of 5e-3. We train for 220k steps and choose the checkpoint with the highest IS.

**Baselines.** We compare with TANGO 1 & 2, (Ghosal et al. 2023), AudioLDM 1 & 2 (Liu et al. 2023a,b), and Make-An-Audio 1 & 2 (Huang et al. 2023b,a). Both AudioLDM and Make-an-Audio train a UNet-based latent diffusion model (Rombach et al. 2022) on Mel-Spectrogram representation, by regarding it as a single channel image, and use a pretrained CLAP encoder to condition the generation on an input prompt. TANGO proposed to use FLAN-T5 (Chung et al. 2024) as the text encoder and reported significant improvements. AudioLDM-2 and Make-an-Audio-2 proposed to use a dual encoder strategy of a T5 (Raffel et al. 2022) and CLAP encoder. Make-an-Audio-2 proposes to use a 1D VAE representation and employ a DiT as the diffusion backbone. Recently, Tango-2 proposed to use instruction fine-tuning on a synthetic dataset to enhance temporal understanding. In our experiments, we focus on text-conditioned natural audio generation.

**Metrics.** We compare the performance of our method with baselines using the standard Frechet Distance (FD), Inception score (IS), and CLAP score on the AudioSet test split. There is little consistency between baselines when computing the metrics. Some prior work reported the Fréchet distance results using the VGGish network (Hershey et al. 2017), denoted as (FAD) (Kilgour et al. 2018), while other uses PANNs (Kong et al. 2019). Additionally, to compute the CLAP score, some prior work (Liu et al. 2023b) used CLAP from LAION, which we denote as CLAP<sub>LAION</sub> (Wu

Ablation	Model	IS	FD	CLAP <sub>MS</sub>
Model Scale	GenAU-S	15.76	21.29	0.36
	GenAU-L	<b>18.98</b>	<b>20.81</b>	<b>0.38</b>
Data Scale	GenAU-L (AC)	12.14	25.82	0.30
	GenAU-L (AutoReCap)	<b>18.98</b>	<b>20.81</b>	<b>0.38</b>
Synthetic Captions	GenAU-S w/o Recap.	11.83	25.34	0.29
	GenAU-S	<b>15.76</b>	<b>21.29</b>	<b>0.36</b>
Backbone	GenAU-S w/ U-Net	13.55	32.51	0.28
	GenAU-S w/ DiT	13.96	23.32	0.33
	GenAU-S w/ FiT	<b>15.76</b>	<b>21.29</b>	<b>0.36</b>

**Table 4** GenAu ablation study on out-of-distribution dataset.

et al. 2023b), while others (Majumder et al. 2024; Huang et al. 2023b,a) used CLAP from Microsoft (Elizalde et al. 2023), which we denote as CLAP<sub>MS</sub>. Furthermore, some prior (Liu et al. 2023a,b) used CLAP re-ranking with 3 samples for computing the metrics. Due to such inconsistencies in evaluation pipelines and varying results for the same baselines reported in different studies, we recompute all metrics using the official checkpoints to ensure consistent comparisons. We follow the same evaluation protocols of AudioLDM (Liu et al. 2023a) without CLAP re-ranking and use the AudioLDM evaluation package to compute the metrics. Besides, to prevent biasing the evaluation based on the training data, we run our ablations on the Bigsoundbank split from Wav-Text5k (Deshmukh et al. 2022), which serves as an out-of-distribution evaluation for our models. Finally, to further validate our results, we run user studies. Details about the user study can be found in the Appendix.

**Comparison with baselines.** In Tab. 3, we report evaluation results. When trained with a *similar size and data scale* to state-of-the-art methods, GenAu achieves superior performance in most metrics, improving IS by 11.1%, FAD

by 4.7%, and 13.5% in  $CLAP_{LAION}$ . Using the full 10-seconds subset of AutoReCap-XL further enhances the results with over 21.8% in IS and 20.5% in CLAP score. Additionally, to isolate the impact of model architecture from data quality, we conduct a user study in Tab. 5, where GenAu is consistently preferred over Make-an-Audio-2, even when trained with comparable data quality.

**How GenAu scale with synthetic data?** To study this, we train GenAu-S (493M params) for 50k steps by fixing AudioCaps and Clotho in the training data and varying the amount of synthetic data from AutoReCap. As reported in Fig. 3 (right), increasing synthetic training data consistently improves both IS and FD. Similar improvements are evident in Tab. 4 where increasing the dataset size significantly boosts all metrics, improving IS by 56.3%. User studies in Tab. 5 further support these findings, where a model training with AutoReCap is consistently favored over a model trained only on AudioCaps. Finally, we report in Tab. 3 the scaling performance on AutoReCap-XL (47M). To ensure a fair comparison, we train on 10-second subset (19.7M) and observe improvements of 12.3% in FD and 11.6% in CLAP score compared to training with AutoReCap (800k).

**Is caption quality important?** We compare in Tab. 4 GenAu-S against the same model trained with WavCaps Mei et al. (2024a) captions. We observe gains across all metrics, confirming the importance of caption quality and the improvements brought by our high-quality captioner AutoCap. Interestingly, expanding data size with lower-quality captions offers no significant gains over training on AudioCaps alone, consistent with Liu et al. (2023b).

**Does GenAu benefit from model size scaling?** Similar to data scaling, increasing model size consistently enhances performance. As shown in Fig. 3 (left), larger models achieve better FD and IS. This is further confirmed in Tab. 4, where GenAu-L (1.25B) outperforms GenAu-S (493M) across all metrics with almost 20.5% increase in IS. User study in Tab. 5 shows a strong preference for GenAu-L over GenAu-S. Unlike previous methods which reported diminishing returns with model scaling Hai et al. (2024); Liu et al. (2023b), GenAu continues to improve as model size scale.

**How do FITs compare with other diffusion backbones?** We evaluate the impact of

Comparison	Realism	Quality	P-Align.	Pref.
<b>Model:</b> GenAu-L vs GenAu-S	61.20%	58.00%	61.20%	60.40%
<b>Data:</b> GenAu-L vs GenAu-L (AC)	60.40%	54.80%	60.40%	59.20%
<b>Backbone:</b> GenAu-S w/o Recap. vs MAD-2	64.40%	64.00%	63.20%	64.80%

**Table 5** User study between various baselines. % of votes in favor of the baseline to the left.

the diffusion backbone by replacing FIT with a UNet (Ronneberger et al. 2015), or a DiT Peebles and Xie (2023). In Tab. 4, we observe that GenAu with FIT outperforms these alternatives across all metrics. We infer that the FIT architect, with its read and write operations, allocates compute more efficiently to the key segments of the input, making it suitable for ambient audio clips which often include silent or redundant parts.

## 5 Conclusion

We take a holistic approach to improve the quality of existing audio generators. Starting by addressing the scarcity of large-scale captioned audio datasets, we propose a scalable and efficient dataset collection pipeline. We then build AutoCap, a strong audio captioner which leverages audio metadata to annotate a dataset of 47M annotated audio clips tailored for large-scale audio generation. We then built a latent diffusion model based on a scalable transformer architecture which we trained on our re-captioned dataset to obtain GenAu, a high-quality open-source model for audio generation. Our approach opens up possibilities for extending GenAu to other domains, such as speech and music. Additionally, AutoReCap-XL can serve as a joint text-audio-video dataset and broadens novel applications such as text to audio-video joint generation.

**Limitations and future work.** AutoCap was fine-tuned on AudioCaps, featuring 4,892 unique words, which limits the diversity of our generated captions. Consequently, GenAu may face challenges in accurately generating audio for detailed prompts. Additionally, while AutoReCap is extensive in size, it has only been validated for audio generation. We leave broader analysis on more tasks for future work.

**Supplementary information.** Please refer to the supplementary website for qualitative samples of the dataset, audio captioner and generator.

## Declarations

**Competing interests.** Authors Moayed Haji-Ali, Guha Balakrishnan, and Vicente Ordonez are affiliated with Rice University, and authors Willi Menapace and Aliaksandr Siarohin are affiliated with Snap Inc., where they are supervised by Sergey Tulyakov. The authors declare they have no financial interests.

**Funding.** This work was supported by Snap Inc and Rice University.

**Ethics approval and consent to participate.** All participants in the user studies provided gave consent and received financial compensation for their participation.

**Consent for publication.** All authors consent to the publication of this work.

**Data availability.** All audio data used in this study will be publicly available upon publication. Model checkpoints will also be released to facilitate reproducibility.

**Code availability.** All code used in this study will be publicly available upon publication.

**Author contributions.** All authors contributed to the conception of the method, its design, and manuscript writing. Authors Moayed Haji-Ali and Willi Menapace were responsible for the implementation of the method and the experiments. All authors read and approved the final manuscript.

# Appendix

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Related Work</b>	<b>3</b>
<b>3</b>	<b>Method</b>	<b>4</b>
3.1	Automatic Audio Captioning . . . . .	4
3.2	Data Collection and Re-captioning Pipeline . . . . .	5
3.3	Scalable Text-2-Audio Generation . . . . .	6
<b>4</b>	<b>Experiments</b>	<b>7</b>
4.1	Automatic Audio Captioning . . . . .	8
4.2	Text-2-Audio Generation . . . . .	8
<b>5</b>	<b>Conclusion</b>	<b>10</b>
<b>A</b>	<b>AutoReCap-XL Details</b>	<b>12</b>
A.1	Stage 1: Data Selection . . . . .	13
A.2	Stage 2: Speech and Music Filtering . . . . .	13
A.3	Stage 3: Post-filtering of Speech and Music. . . . .	13
<b>B</b>	<b>Architecture details</b>	<b>14</b>
B.1	HTSAT Embeddings Extraction . . . . .	14
<b>C</b>	<b>Limitations</b>	<b>15</b>
C.1	AutoCap . . . . .	15
C.2	GenAu . . . . .	16
C.3	AutoReCap-XL . . . . .	16
<b>D</b>	<b>Evaluation Details</b>	<b>16</b>
D.1	Audio Captioning . . . . .	16
D.2	Audio Generation . . . . .	16
D.3	User Study . . . . .	17
<b>E</b>	<b>Training and Inference Details</b>	<b>17</b>
E.1	AutoCap . . . . .	17
E.2	GenAu . . . . .	18
<b>F</b>	<b>Discussion with Concurrent work</b>	<b>18</b>
F.1	Text-conditioned audio generation . . . . .	18
F.2	Audio captioning . . . . .	19
<b>G</b>	<b>Additional Results</b>	<b>19</b>
G.1	Additional Audio Captioning Evaluation . . . . .	19
G.2	Additional Audio Generation Evaluation . . . . .	19
G.3	Additional HTSAT Embedding Extraction Evaluation . . . . .	19

## Appendix A AutoReCap-XL Details


This section outlines the collection and filtering processes for AutoReCap-XL.

### A.1 Stage 1: Data Selection

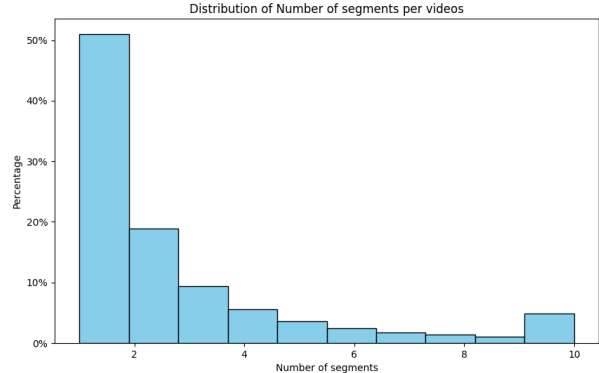
We selected existing video datasets primarily from YouTube for the ease of accessing automatic transcriptions. Specifically, we chose 73 million videos from the datasets AudioSet (Gemmeke et al. 2017), VGGSound (Chen et al. 2020), ACAV100M (Lee et al. 2021), VideoCC (Nagrani et al. 2022), YTTEMP1B (Zellers et al. 2022), and HDVila-100M (Xue et al. 2022). We select these datasets for their likelihood of containing videos with strong audio-video correspondence.

### A.2 Stage 2: Speech and Music Filtering

We downloaded English transcripts from YouTube and used automatically generated ones for videos without existing transcripts. However, we discard videos without any transcripts. While some datasets provide only video segments with specific timestamps, we processed the full videos, totaling around 73 million videos. We accepted audio segments longer than one second that lacked any corresponding subtitles, indicating the absence of speech and music. After filtering, we isolated approximately 327.3 million segments from 55.1 million videos. Fig. A1 displays the distribution of the number of segments per video. We denote this dataset as AutoReCap-XL-Raw. Subsequently, we use AutoCap to caption the audio segments. Fig. A3 shows the distribution of caption lengths. Given that AutoCap was trained for 10-second audio, we limited segments to this duration. Additionally, we concatenate consecutive segments yielding identical captions to form longer audio clips. Fig. A5 illustrates the audio length distribution, and a word cloud of the captions is shown in Fig. A7. Despite filtering, the dataset was still dominated by captions related to speech and music. We attribute this to the limitations of YouTube’s automatic transcription, particularly with certain types of music and less common languages.



**Fig. A1** Distribution of number of segments per-video in AutoReCap-XL-Raw




**Fig. A2** Distribution of the number of segments per-video in AutoReCap-XL

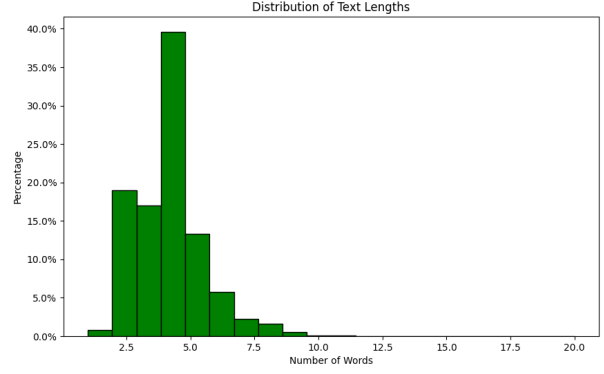
### A.3 Stage 3: Post-filtering of Speech and Music.

To further refine the dataset from speech and music, We follow a simple filtering approach. Specifically, we employed a large language model (LLM) to generate keywords associated with speech and music, such as "talking", "speaking", and "singing," and excluded all audio segments whose captions contained such keywords. Finally, we filter all audio-text pairs with CLAP score less than 0.1. A distribution of the CLAP scores for the collected dataset is present in Fig. A9. This process yielded 47 million audio-text pairs from 20.3 million videos. Fig. A2 shows the number of segments per video, Fig. A4 shows the caption length distribution, Fig. A6 shows the audio length distribution, and Fig. A8 presents a word cloud of the final captions. We outline the data sources for constructing this dataset in Tab. D1. Our proposed

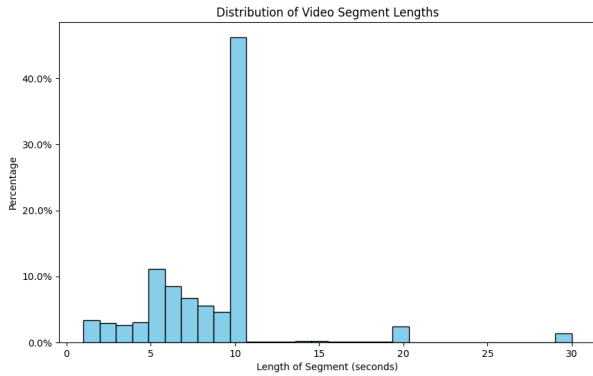





**Fig. A3** Distribution of caption length of AutoReCap-XL-Raw



**Fig. A4** Distribution of caption length of AutoReCap-XL



**Fig. A5** Distribution of audio duration of AutoReCap-XL-Raw





**Fig. A6** Distribution of audio duration of AutoReCap-XL

dataset is not only 75 times larger than the previously largest available dataset, LAION-Audio-630K [Wu et al. \(2023b\)](#) in terms of the number of audio clips, but also provides more accurate captions compared to existing datasets that rely on raw textual data. A comprehensive comparison with other datasets is detailed in Tab. [D2](#)

## Appendix B Architecture details

### B.1 HTSAT Embeddings Extraction

AutoCap uses HTSAT ([Chen et al. 2022](#)) embeddings to encode the input audio and follows the HTSAT-BART ([Mei et al. 2024a](#)) embedding extraction procedure, described in the following, to obtain “fine-grained” HTSAT embeddings. Given a 10-seconds single-channel input audio at 32Khz, HTSAT represents it as a mel-spectrogram using window size of 1024, 320 hop size, and 64 mel-bins, resulting in an input of shape  $(T = 1024, F = 64)$ . The spectrogram is then encoded as latent tokens of shape  $(\frac{T}{8P} = 32, \frac{F}{8P} = 2, 8D = 768)$  before the classification layer. HTSAT-BART [Mei et al. \(2024a\)](#), then averages over the frequency dimension to obtain a representation of shape  $(\frac{T}{8P} = 32, 1, 8D = 768)$  and replicates the latent token by a token replication factor of  $8P = 32$  to obtain a so-called “fine-grained” representation of shape  $32 \times 32 \times 768$ , which is flattened into a representation of shape  $1024 \times 768$ . We adopt this

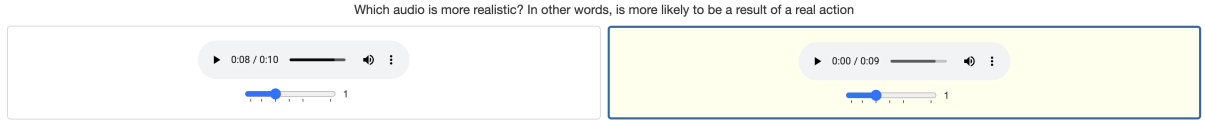
**Fig. A9** CLAP score distribution of AutoReCap-XL

representation throughout our work, and Appx. G.3 provides additional evaluation results showing the performance benefits of the token replication operation.

## Appendix C Limitations

### C.1 AutoCap

Sounds emitted by various objects can often sound similar, such as a waterfall compared to heavy rain, or a can versus a motorcycle engine. In scenarios where metadata lacks detail, our audio captioning model may struggle to disambiguate these sounds accurately. The model also tends to falter in capturing the temporal relationships between sounds and differentiating foreground from background noises. Additionally, since it is fine-tuned on AudioCaps, which contains a limited vocabulary of 4,892 unique words (excluding common stop words), the model frequently produces repetitive words and captions.



**Fig. D10** A screenshot of the user study interface.

## C.2 GenAu

Although our model is trained to generate natural sound effects, it underperforms in specialized areas like music generation or text-to-speech synthesis, where more targeted models are superior. Moreover, the limited vocabulary of the paired texts, even though extensive, hampers the model’s ability to accurately generate audio for long and detailed prompts.

## C.3 AutoReCap-XL

Our proposed dataset, AutoReCap-XL, is substantial in size but features a constrained vocabulary of only 4,461 unique words, excluding stop words, due to the vocabulary limitations of the AudioCaps-trained captioner. Furthermore, despite its potential as a significant contribution, this dataset has not yet been extensively analyzed for caption accuracy or performance in downstream tasks.

# Appendix D Evaluation Details

## D.1 Audio Captioning

While the established practice in the evaluation of audio captioning methods is to report the results on the test set using the checkpoint that performs best on the validation subset, prior work (Labb et al. 2024; Kim et al. 2024c) reported high instability of the metrics on the validation subset and weak correlation between the validation and test performance, making the model’s results vary significantly for different seeds. To alleviate this, ENCLAP (Kim et al. 2024c) selects around five best-performing validation checkpoints and reports their best results on the test set. CNext-trans (Labb et al. 2024) uses the FENSE score to pick the best validation checkpoint. This method of choosing the best checkpoint may produce misleading results and potentially disadvantage baselines. Our model, thanks to the two-stage training paradigm, significantly reduces this instability and we observe steady performance gains as training progresses. Therefore, we report the results at convergence, specifically after 20 epochs of pre-training and 20 epochs of fine-tuning.

## D.2 Audio Generation

There is a lack of consistency in the metrics used across text-to-audio generation baselines. Some baselines, such as Liu et al. (2023a) and Huang et al. (2023a), employ the VGGish network (Hershey et al. 2017) to compute the Fréchet Distance, while others, like Liu et al. (2023b), utilize the PANNs network (Kong et al. 2019), and still others rely on OpenL3 embeddings, such as Evans et al. (2024a). Additionally, some baselines use the LAION CLAP network (Wu et al. 2023b) to compute the CLAP score, whereas others use the Microsoft CLAP network (Elizalde et al. 2023). To further complicate matters, different baselines often report varying results in various publications. To address these discrepancies, we recalculated all metrics for the baselines using their publicly released checkpoints under identical evaluation configurations. Our method significantly outperforms the baselines across all metrics, except for the Fréchet Distance, where it is slightly behind Make-An-Audio 2 (Huang et al. 2023a). Nevertheless, our user study, detailed in the main paper, indicates that GenAu is generally preferred over Make-An-Audio 2.

Data Source	# pairs
AudioSet	339.4k
VGGSounds	126.9k
Freesounds	262.3k
BBC Sound Effects	31.2k
YouTube Videos	47.0M
ACAV-100M	
VideoCC	
YTTEMP1B	
HDVila-100M	
AutoReCap	761.1k
AutoReCap-XL	47.0M

**Table D1** Overview of the employed dataset sources and audio clips counts for each of them.

**Table D2** Comparative overview of the main audio-language datasets.

Dataset	# Text-Audio Pairs	Duration (h)	Text source
AudioCaps	52,904	144	Human
Clotho	5,929	37	Human
MACS	3,537	10	Human
WavText5K	4,072	23	Online raw-data
SoundDescs	32,979	1,060	Online raw-data
LAION-Audio-630K	633,526	4,325	Online raw-data
WavCaps	403,050	7,567	Processed raw-data
AutoReCap	761,113	8,763	Automatic re-captioning
AutoReCap-XL	47M	123,500	Automatic re-captioning
AutoReCap-XL-Raw	327.3M	-	Automatic re-captioning

### D.3 User Study

Each user study reported in this paper involved 5 different participants, yielding a total of 1000 responses per study. Samples were selected from the AudioCaps test split, specifically choosing the top 200 samples with the longest text prompts and sampling 50 for each study to enhance the likelihood of obtaining more complex audio scenarios. To minimize discrepancies between baselines, we fix the seed and other sampling parameters across all experiments.

During the user study, participants were initially presented with two audio clips from the compared baselines and asked to judge which one sounded more realistic. They were then prompted to choose the audio they believed had better quality. Next, after showing the prompt used to generate the audio, participants were asked to select the clip that most faithfully followed the prompt. Finally, they were asked to choose their overall preferred audio clip. A screenshot of the user study interface is included in Fig. D10, and the questions posed to the annotators are detailed in Tab. D3.

## Appendix E Training and Inference Details

### E.1 AutoCap

AutoCap introduces 6.2 million new parameters on top of the frozen HTSAT audio encoder and the base BART model. These parameters include 4.7M for the Q-Former, 0.9M for embedding layers, and 0.6M

**Table D3** Audio Evaluation Criteria

Criterion	Description
Realism	Which audio is more realistic? In other words, is more likely to be a result of a real action.
Quality	Which audio has better quality, regardless of the realism of the audio. Please note that some audio may have background noise, which should not be confused with low quality.
Prompt Alignment	Considering the prompt to generate the audio is "A sewing machine operating as a machine motor hisses loudly in the background", which audio better follows the given prompt?
Overall Preference	Considering the realism, quality, and prompt alignment of the audio, which audio do you prefer more overall? The prompt is: "A sewing machine operating as a machine motor hisses loudly in the background."

for projection layers. The Q-Former employs 256 learnable tokens, a hidden dimension of 256, 8 attention heads, and 2 hidden layers.

We train the audio captioning model using the Adam optimizer, starting with a learning rate of  $10^{-4}$  in stage 1, and reducing to  $10^{-5}$  in stage 2. The training was completed over 9 hours on eight A100 80GB GPUs. Although our model is training with 10-second audio clips, we observed qualitatively that it generalizes well to short audios, such as 1-2 second audio clips.

## E.2 GenAu

We employ the LAMB optimizer for our audio generation model, setting the learning rate at 0.005 with a cosine schedule, and incorporating a weight decay of 0.1 and a dropout rate of 0.1. The small model variant is trained for 210k steps with a batch size of 2,048, while the large model variant is trained for 220k steps with a batch size of 3,072. The large model is trained over 48 hours on 48 A100 80GB GPUs, and the small model on 32 GPUs. Ablation studies are conducted on eight A100 80GB GPUs using a batch size of 512. We further condition the model on the training dataset with a conditioning dataset ID. For generation, we utilize the AudioCaps dataset ID as it is the most reliable dataset.

# Appendix F Discussion with Concurrent work

## F.1 Text-conditioned audio generation

Recently, Stable Audio Open (Evans et al. 2024b) introduced a 1.32B-parameter model capable of generating variable-length stereo audio clips at 44.1 kHz. This model leverages a latent diffusion approach with a DiT (Peebles and Xie 2023) as its diffusion backbone, similar to prior work such as Make-An-Audio 2 (Huang et al. 2023a). In contrast, GenAu employs a FiT architecture. In Tab. 5, we show the superiority of our FiT-based approach over DiT by showing that GenAu-S is consistently preferred over a 937M-parameter DiT-based baseline (Make-An-Audio 2 Huang et al. (2023a)) when trained on comparable data settings (*i.e.* without recaptioning) at a smaller scale (493M parameters). Additionally, Stable Audio Open proposes directly encoding audio clips using a variational autoencoder (VAE) with a ResNet-like architecture, which is particularly effective for higher-resolution audio generation. In contrast, our work adopts previous approaches (Huang et al. 2023a; Liu et al. 2023b) and uses a Mel-spectrogram representation due to its simplicity. GenAu, being a latent model, can readily benefit from improved latent audio representations, such as those employed by Stable Audio Open.



## F.2 Audio captioning

A concurrent work, SOUND-VECAPS (Yuan et al. 2024), and Auto-ACD (Sun et al. 2024), propose prompting a pretrained large language model with multimodal information. SOUND-VECAPS utilizes visual captions generated by a pretrained visual captioner (Wang et al. 2025) alongside audio captions from a pretrained audio captioner, ENCLAP (Kim et al. 2024c), to produce more complex captions, showing significant improvements in the downstream task of audio generation. This aligns with our approach of incorporating visual captions in the audio captioning task. However, unlike these methods, which rely solely on pretrained models, we integrate visual information directly into the training process of the audio captioner. This enables a more dynamic and context-aware incorporation of visual information in the audio captioning task.

Additionally, there has been a recent trend toward training large audio-language models (Ghosh et al. 2024b; Kong et al. 2024; Gong et al. 2023; Deshmukh et al. 2023) and utilizing them for audio captioning in zero-shot settings. While promising in the pursuit of general-purpose models, their reported results on audio captioning remain inferior to state-of-the-art automatic audio captioning (AAC) methods. Consequently, we opt to train a dedicated AAC model, AutoCap, to achieve the highest-quality captions for our proposed dataset, AutoReCap.

## Appendix G Additional Results

In this section, we present additional results which are complemented by our *Website*.

### G.1 Additional Audio Captioning Evaluation

In Tab. D4 we show qualitative results of the captions produced by our method and compare them with state-of-the-art AAC methods. See the *Website* for qualitative results accompanied by the original audio. While ENCLAP (Kim et al. 2024c) and CoNeTTE (Labb et al. 2024) tend to produce short captions, our method produces the most descriptive captions, capturing the most amount of elements from the ground truth audio, an important capability to allow high-quality audio generation (Shi et al. 2020).

### G.2 Additional Audio Generation Evaluation

In this section, we report additional evaluation results and ablations on the task of audio generation.

In Tab. D5, we evaluate fundamental architectural choices in the design of our scalable FIT model. When removing either the Flan-T5 or CLAP encodings, we notice a steady reduction in all metrics. When increasing the number of latent tokens we also notice a steady improvement in performance as more compute is allocated to the model. Similarly, increasing the patch size to 2 results in a performance decrease under all metrics due to the reduced amount of allocated computation.

In Tab. D6, we ablate the 1D-VAE bottleneck size in terms of reconstruction loss and performance of a subsequently trained latent audio diffusion model, in terms of FAD, FD, and IS. Similarly to the phenomenon observed in the image and video generation domain (Gupta et al. 2024; Esser et al. 2024), we observe that a larger number of channels allocated to the latent space results in lower reconstruction losses, but making the latent space more complex, hindering generation quality. We adopt 64 1D-VAE channels for all our experiments.

### G.3 Additional HTSAT Embedding Extraction Evaluation

We perform a series of ablations on HTSAT-BART Mei et al. (2024a) employing different variants of the procedure of Mei et al. (2024a) for the extraction of HTSAT embeddings (see Appx. B.1). We consider HTSAT output tokens of shape  $32 \times 768$  after the averaging operation over the frequency dimension of Mei et al. (2024a), and apply different token repetition factors to produce embeddings with 32 tokens (no token repetition), 256 tokens (8x token repetition) and 1024 tokens (32x token repetition following Mei

et al. (2024a)). For completeness, we perform the same ablation on our AutoCap, using as input to the Q-Former 32 tokens (no token repetition) and 1024 tokens (32x token repetition). Training hyperparameters of AutoCap are modified to match HTSAT-BART Mei et al. (2024a) for the purpose of the ablation.

We followed the training procedure of Mei et al. (2024a) and report evaluation results on the AudioCaps test split for the last obtained checkpoint in Tab. G7 and Fig. G11. As the ablation shows, the token replication operation consistently improves model performance. We attribute this finding to the increased computation in the downstream model caused by it and consequently adopt the best performing 32x token replication embeddings extraction procedure of Mei et al. (2024a) throughout our work.

**Table D4** Qualitative comparison of captioning results on the AudioCaps dataset. See the *Website* for qualitative results accompanied by the respective audio.

Method	Caption
Ground Truth	<i>A man talking as ocean waves trickle and splash while wind blows into a microphone</i>
Ours	<i>A man speaks as wind blows and water splashes</i>
CoNeTTE	<i>A man is speaking and wind is blowing</i>
ENCLAP	<i>A man is speaking and wind is blowing</i>
Ground Truth	<i>An adult male speaks, birds chirp in the background, and many insects are buzzing</i>
Ours	<i>Birds chirp in the distance, followed by a man speaking nearby, after which insects buzz nearby</i>
CoNeTTE	<i>A man speaking with birds chirping in the background.</i>
ENCLAP	<i>Birds are chirping and a man speaks</i>
Ground Truth	<i>A telephone dialing tone followed by a plastic switch flipping on and off</i>
Ours	<i>A telephone dialing followed by a series of plastic clicking then plastic clanking before plastic thumps on a surface</i>
CoNeTTE	<i>A telephone ringing followed by a beep.</i>
ENCLAP	<i>A telephone dialing followed by a series of electronic beeps</i>
Ground Truth	<i>A running train and then a train whistle</i>
Ours	<i>A train moves getting closer and a horn is triggered</i>
CoNeTTE	<i>A train horn blows and a steam whistle is blowing</i>
ENCLAP	<i>A train running on railroad tracks followed by a train horn blowing as wind blows into a microphone</i>
Ground Truth	<i>A child is speaking followed by a door moving</i>
Ours	<i>A child speaks followed by a loud crash and a scream</i>
CoNeTTE	<i>A woman speaking followed by a door opening and closing.</i>
ENCLAP	<i>A young girl speaks followed by a loud bang</i>
Ground Truth	<i>Water splashing as a baby is laughing and birds chirp in the background</i>
Ours	<i>A baby laughs and splashes, and an adult female speaks</i>
CoNeTTE	<i>A baby is laughing and people are talking.</i>
ENCLAP	<i>A baby laughs and splashes in water</i>
Ground Truth	<i>Leaves rustling in the wind with dogs barking and birds chirping</i>
Ours	<i>Birds chirp in the distance, and then a dog barks nearby</i>
CoNeTTE	<i>A dog is barking and a person is walking.</i>
ENCLAP	<i>Birds chirp and a dog barks</i>
Ground Truth	<i>Tapping followed by water spraying and more tapping</i>
Ours	<i>Some light rustling followed by a clank then water pouring</i>
CoNeTTE	<i>A toilet is flushed and water is running.</i>
ENCLAP	<i>A faucet is turned on and runs</i>

**Table D5** Ablation of different FIT architectural variations in terms of patch size number of latent tokens and adopted text encoders on the AudioCaps dataset.

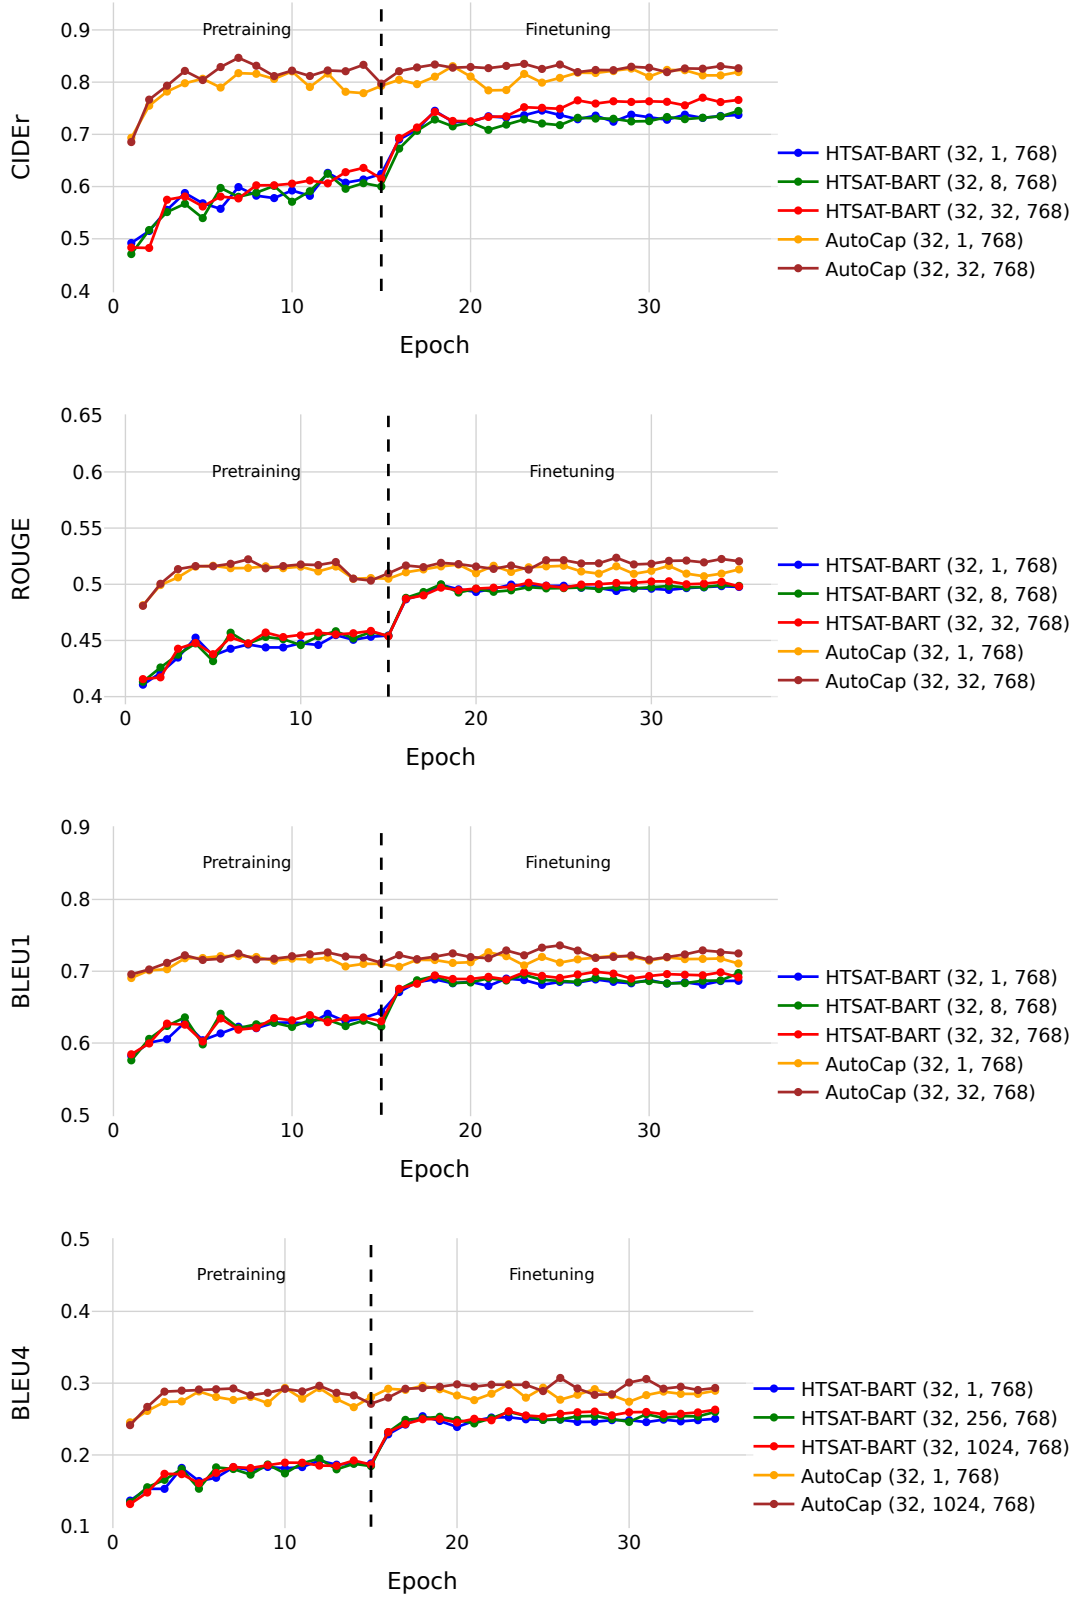
Tokens	Patch size	FLAN-T5	CLAP	FD ↓	FAD ↓	IS ↑
256	1	✓	✓	<b>16.45</b>	<b>1.29</b>	<b>10.26</b>
256	1		✓	17.41	1.39	10.0
256	1	✓		20.47	1.86	8.89
384	1		✓	17.41	1.39	10.0
192	1		✓	18.01	2.01	8.91
128	1		✓	25.56	1.77	7.49
256	2	✓	✓	18.53	1.70	9.0

**Table D6** Ablation of different 1D-VAE designs on audio generation on the AudioCaps dataset.

Channels	Recon. loss	FAD ↓	FD ↓	IS ↑
64	0.159	<b>1.29</b>	<b>16.45</b>	<b>10.26</b>
128	0.107	1.43	16.78	10.11
256	<b>0.064</b>	1.80	18.63	9.43

**Table G7** Ablation of token replication factors for the HTSAT embeddings extraction procedure of (Mei et al. 2024a) on the AudioCaps test split. Larger token replication factors consistently improve performance due to the related compute increase in the downstream model.

	Tokens Count	Replication Factor	CIDEr	BLEU1	BLEU4	ROUGE <sub>L</sub>
HTSAT-BART	32	1x	73.7	68.6	25.0	49.7
HTSAT-BART	256	8x	74.4	69.7	26.0	<b>49.8</b>
HTSAT-BART	1024	32x	<b>76.6</b>	<b>71.5</b>	<b>26.3</b>	<b>49.8</b>
AutoCap	32	1x	81.9	71.7	28.9	51.3
AutoCap	1024	32x	<b>82.7</b>	<b>72.5</b>	<b>29.3</b>	<b>52.0</b>



**Fig. G11** Ablation of token replication factors for the HTSAT embeddings extraction procedure of (Mei et al. 2024a) on the AudioCaps test split for the HTSAT-BART (Mei et al. 2024a) and our AutoCap model.



## References

- Achiam J, Adler S, Agarwal S, et al (2023) GPT-4 Technical Report. arXiv preprint arXiv:230308774
- BBC Sound Effects (2024) Bbc sound effects archive. <https://sound-effects.bbcrewind.co.uk/>, accessed: 2024-10-01
- Chen G, Wang G, Huang X, et al (2024a) Semantically consistent video-to-audio generation using multimodal language large model. arXiv preprint arXiv:240416305
- Chen H, Xie W, Vedaldi A, et al (2020) Vggsound: A large-scale audio-visual dataset. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)
- Chen K, Du X, Zhu B, et al (2022) Hts-at: A hierarchical token-semantic audio transformer for sound classification and detection. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)
- Chen T, Li L (2023) Fit: Far-reaching interleaved transformers. arXiv preprint arXiv:230512689
- Chen TS, Siarohin A, Menapace W, et al (2024b) Panda-70m: Captioning 70m videos with multiple cross-modality teachers. In: International Conference on Computer Vision and Pattern Recognition (CVPR)
- Chen W, Ma Z, Li X, et al (2025) Slam-aac: Enhancing audio captioning with paraphrasing augmentation and clap-refine through llms. In: ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 1–5
- Cheng YC, Lee HY, Tulyakov S, et al (2023) SDFusion: Multimodal 3d shape completion, reconstruction, and generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp 4456–4465
- Cheng Z, Leng S, Zhang H, et al (2024) Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. arXiv preprint arXiv:240607476
- Chung HW, Hou L, Longpre S, et al (2024) Scaling instruction-finetuned language models. Journal of Machine Learning Research 25(70):1–53
- Cousin M, Labbé E, Pellegrini T (2023) Multilingual audio captioning using machine translated data. arXiv preprint arXiv:230907615
- Deshmukh S, Elizalde B, Wang H (2022) Audio retrieval with wavtext5k and clap training. arXiv preprint arXiv:220914275
- Deshmukh S, Elizalde B, Singh R, et al (2023) Pengi: An audio language model for audio tasks. Advances in Neural Information Processing Systems 36:18090–18108
- Deshmukh S, Elizalde B, Emmanouilidou D, et al (2024a) Training audio captioning models without audio. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 371–375
- Deshmukh S, Singh R, Raj B (2024b) Domain adaptation for contrastive audio-language models. arXiv preprint arXiv:240209585

- Drossos K, Lipping S, Virtanen T (2020) Clotho: an audio captioning dataset. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)
- Elizalde B, Deshmukh S, Al Ismail M, et al (2023) Clap learning audio concepts from natural language supervision. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE
- Elizalde B, Deshmukh S, Wang H (2024) Natural language supervision for general-purpose audio representations. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 336–340
- Esser P, Rombach R, Ommer B (2021) Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 12873–12883
- Esser P, Kulal S, Blattmann A, et al (2024) Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first international conference on machine learning
- Evans Z, Carr C, Taylor J, et al (2024a) Fast timing-conditioned latent audio diffusion. In: International Conference on Machine Learning (ICML)
- Evans Z, Parker JD, Carr C, et al (2024b) Stable audio open. arXiv preprint arXiv:240714358
- Font F, Roma G, Serra X (2013) Freesound technical demo. In: Proceedings of the 21st ACM International Conference on Multimedia. Association for Computing Machinery, New York, NY, USA, MM '13, p 411–412, <https://doi.org/10.1145/2502081.2502245>, URL <https://doi.org/10.1145/2502081.2502245>
- Gemmeke JF, Ellis DPW, Freedman D, et al (2017) Audio set: An ontology and human-labeled dataset for audio events. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)
- Ghosal D, Majumder N, Mehrish A, et al (2023) Text-to-audio generation using instruction guided latent diffusion model. In: Proceedings of the 31st ACM International Conference on Multimedia, pp 3590–3598
- Ghosh S, Kumar S, Evuru CKR, et al (2024a) Recap: Retrieval-augmented audio captioning. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 1161–1165
- Ghosh S, Kumar S, Seth A, et al (2024b) Gama: A large audio-language model with advanced audio understanding and complex reasoning abilities. arXiv preprint arXiv:240611768
- Gong Y, Luo H, Liu AH, et al (2023) Listen, think, and understand. arXiv preprint arXiv:230510790
- Gontier F, Serizel R, Cerisara C (2021) Automated audio captioning by fine-tuning bart with audioset tags. In: DCASE 2021 - 6th Workshop on Detection and Classification of Acoustic Scenes and Events
- Guan W, Wang K, Zhou W, et al (2024) Lafma: A latent flow matching model for text-to-audio generation. arXiv preprint arXiv:240608203
- Guo Z, Mao J, Tao R, et al (2024) Audio generation with multiple conditional diffusion model. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp 18153–18161
- Gupta A, Yu L, Sohn K, et al (2024) Photorealistic video generation with diffusion models. In: European Conference on Computer Vision, Springer, pp 393–411

- Hai J, Xu Y, Zhang H, et al (2024) Ezaudio: Enhancing text-to-audio generation with efficient diffusion transformer. arXiv preprint arXiv:240910819
- Haji-Ali M, Balakrishnan G, Ordonez V (2024a) Elasticdiffusion: Training-free arbitrary size image generation through global-local content separation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp 6603–6612
- Haji-Ali M, Menapace W, Siarohin A, et al (2024b) Av-link: Temporally-aligned diffusion features for cross-modal audio-video generation. arXiv preprint arXiv:241215191
- Hayakawa A, Ishii M, Shibuya T, et al (2024) Discriminator-guided cooperative diffusion for joint audio and video generation. arXiv preprint arXiv:240517842
- Hershey S, Chaudhuri S, Ellis DPW, et al (2017) Cnn architectures for large-scale audio classification. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)
- Ho J, Jain A, Abbeel P (2020) Denoising diffusion probabilistic models. Advances in neural information processing systems 33:6840–6851
- Ho J, Chan W, Saharia C, et al (2022) Imagen video: High definition video generation with diffusion models. [arXiv:2210.02303](https://arxiv.org/abs/2210.02303)
- Huang J, Ren Y, Huang R, et al (2023a) Make-an-audio 2: Temporal-enhanced text-to-audio generation. arXiv preprint arXiv:230518474
- Huang R, Huang J, Yang D, et al (2023b) Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. Proceedings of the 40th International Conference on Machine Learning (ICML)
- Kadlčık M, Hájek A, Kieslich J, et al (2023) A whisper transformer for audio captioning trained with synthetic captions and transfer learning. arXiv preprint arXiv:230509690
- Kilgour K, Zuluaga M, Roblek D, et al (2018) Fr\’echet audio distance: A metric for evaluating music enhancement algorithms. arXiv preprint arXiv:181208466
- Kim CD, Kim B, Lee H, et al (2019) Audiocaps: Generating captions for audios in the wild. In: NAACL-HLT
- Kim E, Kim J, Oh Y, et al (2022) Exploring train and test-time augmentations for audio-language learning. arXiv preprint arXiv:221017143
- Kim G, Martinez A, Su YC, et al (2024a) A versatile diffusion transformer with mixture of noise levels for audiovisual generation. arXiv preprint arXiv:240513762
- Kim J, Jeon M, Jung J, et al (2024b) Enclap++: Analyzing the enclap framework for optimizing automated audio captioning performance. arXiv preprint arXiv:240901201
- Kim J, Jung J, Lee J, et al (2024c) Enclap: Combining neural audio codec and audio-text joint embedding for automated audio captioning. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)
- Kong Q, Cao Y, Iqbal T, et al (2019) Panns: Large-scale pretrained audio neural networks for audio pattern recognition. IEEE/ACM Transactions on Audio, Speech, and Language Processing

- Kong Z, Goel A, Badlani R, et al (2024) Audio flamingo: A novel audio language model with few-shot learning and dialogue abilities. arXiv preprint arXiv:240201831
- Kreuk F, Synnaeve G, Polyak A, et al (2023) Audiogen: Textually guided audio generation. In: The Eleventh International Conference on Learning Representations
- Labb E, Pellegrini T, Pinquier J, et al (2024) Conette: An efficient audio captioning system leveraging multiple datasets with task embedding. IEEE/ACM Transactions on Audio, Speech, and Language Processing
- Labbé E, Pellegrini T, Pinquier J (2023) Killing two birds with one stone: Can an audio captioning system also be used for audio-text retrieval? arXiv preprint arXiv:230815090
- Lavie A, Agarwal A (2007) Meteor: an automatic metric for mt evaluation with high levels of correlation with human judgments. In: Proceedings of the Second Workshop on Statistical Machine Translation
- Lee S, Chung J, Yu Y, et al (2021) Acav100m: Automatic curation of large-scale datasets for audio-visual video representation learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp 10274–10284
- Lee Sg, Ping W, Ginsburg B, et al (2022) Bigvgan: A universal neural vocoder with large-scale training. arXiv preprint arXiv:220604658
- Lewis M, Liu Y, Goyal N, et al (2020) BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics
- Li J, Li D, Savarese S, et al (2023) Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In: International Conference on Machine Learning (ICML)
- Li Y, Wang X, Liu H (2024) Audio-free prompt tuning for language-audio models. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 491–495
- Liang J, Zhang H, Liu H, et al (2024) Wavcraft: Audio editing and generation with large language models. arXiv preprint arXiv:240309527
- Lin CY (2004) Rouge: A package for automatic evaluation of summaries. In: Text summarization branches out, pp 74–81
- Liu H, Chen Z, Yuan Y, et al (2023a) Audioldm: Text-to-audio generation with latent diffusion models. Proceedings of the 40th International Conference on Machine Learning (ICML)
- Liu H, Tian Q, Yuan Y, et al (2023b) AudioLDM 2: Learning holistic audio generation with self-supervised pretraining. arXiv preprint arXiv:230805734
- Liu H, Chen K, Tian Q, et al (2024a) Audiosr: Versatile audio super-resolution at scale. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 1076–1080
- Liu H, Huang R, Liu Y, et al (2024b) Audioldm: Text-to-audio generation with latent consistency models. arXiv preprint arXiv:240600356

- Liu S, Zhu Z, Ye N, et al (2017) Improved image captioning via policy gradient optimization of spider. In: IEEE International Conference on Computer Vision (ICCV)
- Liu X, Huang Q, Mei X, et al (2023c) Visually-Aware Audio Captioning With Adaptive Audio-Visual Attention. In: Proc. INTERSPEECH 2023
- Liu X, Kong Q, Zhao Y, et al (2024c) Separate anything you describe. IEEE/ACM Transactions on Audio, Speech, and Language Processing
- Mahfuz R, Guo Y, Visser E (2023) Improving audio captioning using semantic similarity metrics. In: ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 1–5
- Majumder N, Hung CY, Ghosal D, et al (2024) Tango 2: Aligning diffusion-based text-to-audio generations through direct preference optimization. In: Proceedings of the 32nd ACM International Conference on Multimedia, pp 564–572
- Mao Y, Shen X, Zhang J, et al (2024) Tavgbench: Benchmarking text to audible-video generation. In: Proceedings of the 32nd ACM International Conference on Multimedia, pp 6607–6616
- Mei X, Meng C, Liu H, et al (2024a) Wavcaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research. IEEE/ACM Transactions on Audio, Speech, and Language Processing
- Mei X, Nagaraja V, Le Lan G, et al (2024b) Foleygen: Visually-guided audio generation. In: 2024 IEEE 34th International Workshop on Machine Learning for Signal Processing (MLSP), IEEE, pp 1–6
- Melechovsky J, Guo Z, Ghosal D, et al (2023) Mustango: Toward controllable text-to-music generation. arXiv preprint arXiv:231108355
- Menapace W, Siarohin A, Skorokhodov I, et al (2024) Snap video: Scaled spatiotemporal transformers for text-to-video synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp 7038–7048
- Miech A, Zhukov D, Alayrac JB, et al (2019) HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV)
- Nagrani A, Seo PH, Seybold B, et al (2022) Learning audio-video modalities from image captions. In: European Conference on Computer Vision, Springer, pp 407–426
- Niu X, Zhang J, Walder C, et al (2024) Soundlocd: An efficient conditional discrete contrastive latent diffusion model for text-to-sound generation. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 261–265
- Paissan F, Della Libera L, Wang Z, et al (2023) Audio editing with non-rigid text prompts. arXiv preprint arXiv:231012858
- Papineni K, Roukos S, Ward T, et al (2002) Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics
- Peebles W, Xie S (2023) Scalable diffusion models with transformers. In: Proceedings of the IEEE/CVF international conference on computer vision, pp 4195–4205



- Podell D, English Z, Lacey K, et al (2023) Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:230701952
- Qiu H, Xia M, Zhang Y, et al (2023) Freenoise: Tuning-free longer video diffusion via noise rescheduling. arXiv preprint arXiv:231015169
- Raffel C, Shazeer N, Roberts A, et al (2022) Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research (JMLR)*
- Ramesh A, Dhariwal P, Nichol A, et al (2022) Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:220406125 1(2):3
- Rho K, Lee H, Iverson V, et al (2025) Lavcap: Llm-based audio-visual captioning using optimal transport. arXiv preprint arXiv:250109291
- Rombach R, Blattmann A, Lorenz D, et al (2022) High-resolution image synthesis with latent diffusion models. In: *IEEE/CVF conference on computer vision and pattern recognition (CVPR)*
- Ronneberger O, Fischer P, Brox T (2015) U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*
- Saito K, Kim D, Shibuya T, et al (2024) Soundctm: Uniting score-based and consistency models for text-to-sound generation. arXiv preprint arXiv:240518503
- Salewski L, Fauth S, Koepke A, et al (2023) Zero-shot audio captioning with audio-language model guidance and audio context keywords. arXiv preprint arXiv:231108396
- Shi Y, Lan GL, Nagaraja V, et al (2023) Enhance audio generation controllability through representation similarity regularization. arXiv preprint arXiv:230908773
- Shi Z, Zhou X, Qiu X, et al (2020) Improving image captioning with better use of captions. arXiv preprint arXiv:200611807
- Shu F, Zhang L, Jiang H, et al (2023) Audio-visual llm for video understanding. arXiv preprint arXiv:231206720
- Song J, Meng C, Ermon S (2021) Denoising diffusion implicit models. In: *International Conference on Learning Representations (ICLR)*
- SoundBible (2024) Free sound effects. <https://soundbible.com/>, accessed: 2024-10-01
- Sridhar AK, Guo Y, Visser E, et al (2024) Parameter efficient audio captioning with faithful guidance using audio-text shared latent representation. In: *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, pp 1181–1185
- Sun L, Xu X, Wu M, et al (2024) Auto-acd: A large-scale dataset for audio-language representation learning. In: *Proceedings of the 32nd ACM International Conference on Multimedia*, pp 5025–5034
- Tang C, Yu W, Sun G, et al (2023a) Salmonn: Towards generic hearing abilities for large language models. arXiv preprint arXiv:231013289
- Tang Z, Yang Z, Zhu C, et al (2023b) Any-to-any generation via composable diffusion. *Advances in Neural Information Processing Systems* 36:16083–16099

- Tang Z, Yang Z, Khademi M, et al (2024) Codi-2: In-context interleaved and interactive any-to-any generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp 27425–27434
- Tian Z, Liu Z, Yuan R, et al (2024) Vidmuse: A simple video-to-music generation framework with long-short-term modeling. arXiv preprint arXiv:240604321
- Vahdati DS, Nguyen TD, Azizpour A, et al (2024) Beyond deepfake images: Detecting ai-generated videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp 4397–4408
- Vaswani A, Shazeer N, Parmar N, et al (2017) Attention is all you need. In: Advances in Neural Information Processing Systems (NeurIPS)
- Vedantam R, Zitnick C, Parikh D (2015) Cider: Consensus-based image description evaluation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
- Villegas R, Babaeizadeh M, Kindermans PJ, et al (2022) Phenaki: Variable length video generation from open domain textual descriptions. In: International Conference on Learning Representations
- Vyas A, Shi B, Le M, et al (2023) Audiobox: Unified audio generation with natural language prompts. arXiv preprint arXiv:231215821
- Wang H, Ma J, Pascual S, et al (2024a) V2a-mapper: A lightweight solution for vision-to-audio generation by connecting foundation models. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp 15492–15501
- Wang K, Deng S, Shi J, et al (2024b) Av-dit: Efficient audio-visual diffusion transformer for joint audio and video generation. arXiv preprint arXiv:240607686
- Wang W, Lv Q, Yu W, et al (2025) Cogvlm: Visual expert for pretrained language models. Advances in Neural Information Processing Systems 37:121475–121499
- Wang Y, Chen X, Ma X, et al (2024c) Lavie: High-quality video generation with cascaded latent diffusion models. International Journal of Computer Vision pp 1–20
- Wang Y, Guo W, Huang R, et al (2024d) Frieren: Efficient video-to-audio generation with rectified flow matching. arXiv e-prints pp arXiv–2406
- Wu SL, Chang X, Wichern G, et al (2024) Improving audio captioning models with fine-grained audio features, text embedding supervision, and llm mix-up augmentation. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 316–320
- Wu Y, Chen K, Zhang T, et al (2023a) Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)
- Wu Y, Chen K, Zhang T, et al (2023b) Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)
- Xing Y, He Y, Tian Z, et al (2024) Seeing and hearing: Open-domain visual-audio generation with diffusion latent aligners. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern

Recognition, pp 7151–7161

- Xu M, Li C, Su D, et al (2024a) Prompt-guided precise audio editing with diffusion models. arXiv preprint arXiv:240604350
- Xu Y, Chen H, Yu J, et al (2024b) Secap: Speech emotion captioning with large language model. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp 19323–19331
- Xue H, Hang T, Zeng Y, et al (2022) Advancing high-resolution video-language representation with large-scale video transcriptions. In: International Conference on Computer Vision and Pattern Recognition (CVPR)
- Xue J, Deng Y, Gao Y, et al (2024) Auffusion: Leveraging the power of diffusion and large language models for text-to-audio generation. IEEE/ACM Transactions on Audio, Speech, and Language Processing
- Yang D, Tian J, Tan X, et al (2023a) Uniaudio: An audio foundation model toward universal audio generation. arXiv preprint arXiv:231000704
- Yang D, Yu J, Wang H, et al (2023b) Diffsound: Discrete diffusion model for text-to-sound generation. IEEE/ACM Transactions on Audio, Speech, and Language Processing
- Ye Z, Wang Y, Wang H, et al (2022) Featurecut: An adaptive data augmentation for automated audio captioning. In: 2022 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC), IEEE, pp 313–318
- You Y, Li J, Reddi S, et al (2019) Large batch optimization for deep learning: Training bert in 76 minutes. arXiv preprint arXiv:190400962
- Yuan Y, Jia D, Zhuang X, et al (2024) Sound-vecaps: Improving audio generation with visual enhanced captions. arXiv preprint arXiv:240704416
- Zellers R, Lu J, Lu X, et al (2022) Merlot reserve: Neural script knowledge through vision and language and sound. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp 16375–16387
- Zhang H, Li X, Bing L (2023a) Video-llama: An instruction-tuned audio-visual language model for video understanding. arXiv preprint arXiv:230602858
- Zhang Y, Maezawa A, Xia G, et al (2023b) Loop copilot: Conducting ai ensembles for music generation and iterative editing. arXiv preprint arXiv:231012404
- Zhang Y, Xu X, Du R, et al (2024) Zero-shot audio captioning using soft and hard prompts. arXiv preprint arXiv:240606295
- Zhu G, Darefsky J, Duan Z (2024) Cacophony: An improved contrastive audio-text model. IEEE/ACM Transactions on Audio, Speech, and Language Processing