





Masked Generative Video-to-Audio Transformers with Enhanced Synchronicity

Santiago Pascual¹ , Chunghsin Yeh^{1*} ,
Ioannis Tsiamas^{1,2} , and Joan Serra¹ 

¹ Dolby Laboratories

² Universitat Politècnica de Catalunya

Abstract. Video-to-audio (V2A) generation leverages visual-only video features to render plausible sounds that match the scene. Importantly, the generated sound onsets should match the visual actions that are aligned with them, otherwise unnatural synchronization artifacts arise. Recent works have explored the progression of conditioning sound generators on still images and then video features, focusing on quality and semantic matching while ignoring synchronization, or by sacrificing some amount of quality to focus on improving synchronization only. In this work, we propose a V2A generative model, named MaskVAT, that interconnects a full-band high-quality general audio codec with a sequence-to-sequence masked generative model. This combination allows modeling high audio quality, semantic matching, and temporal synchronicity at the same time. Our results show that, by combining a high-quality codec with the proper pre-trained audio-visual features and a sequence-to-sequence parallel structure, we are able to yield highly synchronized results on one hand, whilst being competitive with the state of the art of non-codec generative audio models. Sample videos and generated audios are available at <https://maskvat.github.io/>.

Keywords: Video-to-Audio · Masked Token Generative Model

1 Introduction

Audio-visual cross-modal generation has gained a lot of traction in recent years, with the appearance of works for both audio-to-video (A2V) and video-to-audio (V2A) generation [9, 10, 14, 20, 25, 34, 48, 56]. V2A generation has some immediate and impactful applications for the media production industry. On the one hand, it promises to accelerate, improve, and/or simplify foley sound effect generation. On the other hand, tasks that feature both synchronization with respect to a visual input and also a textually guided conditioning, like automatic dubbing, can greatly benefit from a synchronized V2A generative model that features multi-modal conditioning.

Several proposals showed the potential of conditioning audio generative models on still images and video frame features to perform V2A [14, 23, 44, 48], with

* Corresponding author: cseh@dolby.com

a main focus on audio quality and semantic matching (i.e. how well the audio corresponds to the visible contents). Nevertheless, such audio quality focus is normally restricted to wide-band scenarios (i.e. up to 22.05kHz audio sampling rates), and due to their processing of videos as sequences of independent image frames, they also typically bypass the importance of audio-visual temporal synchronization. Although explicit efforts have been made to improve such synchronization [34, 36], these approaches have not particularly focused on also pushing the audio quality simultaneously [48]. A partial culprit to this may be the need for high-quality latent spaces to perform the V2A mapping efficiently and effectively, whilst offering a decoding between the generated latent representations and their full-band signal space counterparts.

In this work, we propose the Masked Generative Video-to-Audio Transformer (MaskVAT), a V2A system that generates high quality full-band audio with strong temporal alignment in the audio-visual correspondence. Our main ideas and contributions are summarized as follows: First, we leverage a state of the art neural audio codec to build the general audio synthesis model in its space, which allows us to generate high-quality audio at full-band resolution [40], unlike previous works that were limited to wide-band audio and thus constrained the quality of their general audio synthesis method [14, 23, 34, 36, 44, 48]. Secondly, we leverage a masked generative Transformer approach due to its inherently efficient sampling process [4, 6, 17], compared to autoregressive counterparts [23, 36, 44]. Thirdly, our proposal is especially motivated to explicitly enforce the temporal alignment between the generated audible events and the visual actions present in the video. We also refer to this temporal alignment as synchronicity. To achieve this, we explore the viability of using a novel combination of conditioning audio-visual features, architectural variations, and regularization losses that ensure a higher degree of semantic and temporal V2A alignment. Our dual focus on audio quality and temporal alignment addresses key limitations in previous research [14, 23, 34, 36, 44, 48], providing a more comprehensive solution for V2A mapping. Finally, in order to appropriately assess the importance of enhancing these dimensions of V2A generation, we propose a combination of objective metrics and subjective tests to compare the V2A mapping performance across three dimensions: (1) audio quality, (2) semantic alignment, and (3) synchronicity.

2 Related Work

Video to Audio Generation Early neural V2A approaches proposed sound synthesis from videos as a way to study physical interactions of materials within a visual scene of limited diversity [39]. Similarly, other works started tackling V2A inside a cross-modality generative framework, where both V2A and A2V were tackled as a joint problem [9, 20, 43]. A number of source-specific models (targeting specific video/sound classes) were also proposed [10, 56].

Motivated by the need to scale V2A as a source-agnostic problem, SpecVQ-GAN [23] was proposed as a first multi-class visually-guided sound generator model. Conditioned on a sequence of visual features like RGB frames or optical

flow, SpecVQGAN is built upon an autoregressive transformer [47] that learns to generate sequences of codewords that represent mel spectrograms through a VQGAN lossy compression [15]. Im2Wav [42] is another Transformer-based audio language model conditioned on image representation extracted from a pre-trained CLIP model. Similarly, CLIPsonic-IQ [14] leverages the CLIP features of individual visual frames to drive a diffusion sound generator.

The potential of multiple prior-mapping models has been recently investigated. In particular, V2A-mapper [48] bridges the domain gap between an average CLIP embedding, which summarizes the input video sequence, and a CLAP embedding [52], which drives an AudioLDM generative model [32]. Many works leverage visual encoders that were pre-trained for individual image recognition tasks [14, 23, 44, 48]. However, this usually hinders the process of modeling the visual dynamics intrinsic to the video scene and its audio-visual synchronicity. Diff-Foley [34] was proposed to improve this, by developing latent-diffusion generative model which is driven by a contrastive audio-visual pre-trained (CAVP) encoder. The CAVP explicitly learns to distill audio onset features into the video encoder through self-supervised training, fine-tuning a video encoder to extract alignment-sensitive visual cues to drive the V2A generation. On similar line, FoleyGen [36] proposes specific architectural attention patterns pre-designed to enforce audio-visual alignment in their generative model.

Audio-Visual Alignment Representations A crucial aspect of V2A is the synchronization (temporal alignment) between an input video and the generated audio. This is often achieved with the help of an audio-visual alignment representation model. AVST (Audio-Visual Synchronisation with Transformers) [7] detects audio-visual synchronisation in a self-supervised manner and predicts the class as either sync or off-sync. SparseSync [24] considers that the audio-visual correspondence may only be available at sparse events. The proposed SparseSelector compresses the audio and visual input tokens into two small sets of learnable selectors. These selectors form an input to a transformer which predicts the temporal offset between the audio and visual streams. It formulates audio-visual synchronisation as a classification task onto a set of offsets (e.g. 21 classes between $-2/+2$ sec.). As mentioned, Diff-Foley [34] adopts CAVP to learn more temporally and semantically aligned features, then it trains a latent diffusion model with CAVP-aligned visual features on spectrogram latent space. That is, it leverages CAVP for (1) generating audio that is temporally aligned with the visual events, and (2) deriving the Alignment Accuracy metric. The recently proposed sequential contrastive audio-visual (SCAV) encoder [46] contrasts audio-visual feature sequences based on their non-aggregated representations. Concretely, it maps CLIP and BEATs [11] feature sequences to a common sequential space, leveraging a distance-based contrastive approach. Contrasting in the sequential space takes advantage of the fine-grained local semantic information between the two modalities, thus surpassing other vanilla contrastive methods in audio-visual retrieval.

Autoregressive and Mask-based Audio Token Generation Early works proved that waveform-based generative modeling was possible with explicit maximum likelihood autoregressive (AR) strategies, as in WaveNet [38] or SampleRNN [35]. These proposals suffered from inefficiencies inherent to their AR nature, which was palliated by subsequent works like WaveRNN [26] or parallel WaveNet [37]. The advancement in neural audio codecs also facilitated the use of language modeling strategies for generative audio, and one of their strong advantages over previous models is the lower framerate featured in the codec spaces compared to the raw waveforms. Based on the SoundStream codec [53], AudioLM [3] is the first to take a language modelling approach to audio generation. Based on Encodec [13], AudioGen [29] is an AR generative model that generates audio samples conditioned on text inputs. Similarly, MusicLM [1] and MusicGen [12] are also proposed as music specific AR models.

Despite promising results are obtained in the aforementioned models, the AR sequence length grows quadratically, easily forming an extremely long sequence due to the temporally-dense nature of audio and the multiple levels of VQ codebooks. SoundStorm [4] is one of the first to adapt a parallel decoding scheme like MaskGIT [6] to predict masked audio tokens produced by SoundStream [53]. Based on DAC [30], VampNet [17] follows a similar approach for music audio generation. Through different prompting techniques, VampNet can operate in a continuum between compression and generation. Based on Encodec [13], MAGNet [57] proposes to further improve the efficiency and quality by means of predicting spans of masked tokens, scoring the prediction confidence with a pre-trained model, and fusing AR and non-AR generation.

3 Method

3.1 Audio Tokenizer

In this work, we consider full-band single channel audio sequences. This means that our model has to process waveforms of audio sampled at 44.1 kHz or more. In order to decouple audio quality from the scalability of our generative strategy, we choose to operate in a latent space of low framerate, and since our strategy follows a discrete masked-token framework, our latent encoder must feature some discretized bottleneck at its core. To this end, we leverage a state of the art pre-trained neural codec for general audio, the Descript Audio Codec³ (DAC) [30]. DAC takes an audio waveform of T samples $\mathbf{x}^a \in \mathbb{R}^T$ and returns a codegram, which is a tensor $\mathbf{C}^a = \text{DAC}(\mathbf{x}^a)$, where $\mathbf{C}^a \in \mathbb{R}^{L \times K}$. A strong convenience of DAC is the framerate reduction it features, converting the waveform at 44.1 kHz to K token sequences of 86 Hz. The number of parallel channels in \mathbf{C}^a refers to the amount of RVQ levels, which hierarchically increase the codec bitrate while maintaining the same sequence length L . In our case, we stick to the pre-trained DAC with $K = 9$.

³ <https://github.com/descriptinc/descript-audio-codec>




Fig. 1: Overview of the three MaskVAT structures proposed.

3.2 Masked Generative Video-to-Audio Transformer

Similarly to recent works [4, 17], our generative strategy follows the formulation introduced in masked generative token modeling from computer vision (MaskGIT [6]), which was adapted to perform masked acoustic token modeling [17]. Therefore we build a Transformer architecture that predicts a tokenized sequence of audio. A key difference with respect to the image domain is the usage of a hierarchical tokenizer to compress the raw audio through the RVQ neural codec [17]. Considering the codegram $\mathbf{C}^a \in \mathbb{R}^{L \times K}$ coming from the tokenizer (see Sec. 3.1), the output of our MaskVAT directly yields the probabilities for the whole codegram, spanning K levels and L time-steps, all in parallel. This is represented in the logits of the different explored models in Fig. 1. Nevertheless, since the summation of RVQ levels embeddings from the audio tokenizer is intrinsically representing a full-band acoustic composition [30], we follow this strategy of first embedding and then summing the codewords in order to obtain the input tokens for our MaskVAT. We also initialize the embeddings with the pre-trained RVQ embeddings from DAC, provided that during early experiments we found beneficial to leverage them for faster convergence. This codegram embedding per k -level and summation is depicted in Fig. 1, having as many embedding parallel layers as K codegram levels coming from the audio tokenizer. Once we obtain the embedded input token sequence, we inject it into a Transformer model, which is built based on two main possible blocks. On the one hand, we can use an adaptation of the AdaLN block proposed in Diffusion Transformers [41]. This modification adapts the AdaLN modulation to deal with conditioning sequences, hence featuring a temporal dimension of information. Both input tokens and conditioning sequence must feature the same length in this case. On the other hand, we can also consider the usage of cross-attention as a way of learning the alignment between the conditioning sequence and the audio token sequence. Fig. 1 shows the diagrams of three MaskVAT designs explored in this work.

MaskVAT models We propose the following three MaskVAT model variants. **MaskVAT_{AdaLN}** (see Fig. 1-a) stacks M AdaLN blocks to build the Transformer structure. The conditioning front-end outputs are adjusted to have the same length as the Transformer input token sequence through the length adapter, which is a nearest neighbor interpolation layer, and get concatenated channel-wise after-wards to be served to the AdaLN blocks. **MaskVAT_{Seq2seq}** (see Fig. 1-b) features a sequence-to-sequence model that first embeds visual embeddings through a Transformer encoder. Then, a stack of M cross-attention blocks acts as a parallel decoder in order to mix the conditioning with the main token sequence. An advantage of this approach is also the possibility of introducing auxiliary losses at the encoder output in order to enforce synchronicity and semantic proximity with respect to other audio features in an end-to-end fashion. As shown in Fig. 1-b, we enforce feature alignment at the Transformer encoder output upon linear projections of a pre-trained BEATs encoder [11]. BEATs is a state of the art self-supervised audio encoder used for large scale general audio classification [11], therefore a good semantic descriptor of the sequence of audio events. **MaskVAT_{Hybrid}** (see Fig. 1-c) is a hybrid approach that mixes the two previous approaches. Here we have the end-to-end learnable component of distilling BEATs into a Transformer encoder that processes the visual features, as well as the alignment enforcement of using AdaLN blocks depending on the most alignment-sensitive features, the ones coming from S3D (see Sec. 3.2). All these MaskVAT models end with a Linear head operator that yields a 3D grid of dimensions $L \times K \times D$ representing the logits over DAC codewords, where $K = 9$ and $D = 1024$.

Visual Conditioning The models we propose in Fig. 1 feature two possible conditioning front-ends, which extract video features from the RGB sequences $\mathbf{V}^{RGB} \in \mathbb{R}^{F \times 3 \times h \times w}$ to drive the V2A mapping, where F , h , and w are the number of frames, their height, and their width respectively. **CLIP_{seq}**: a pre-trained CLIP image encoder is used to process each video frame \mathbf{v}_f^{RGB} projected then through a time-independent MLP into the shared dimensionality of the subsequent Transformer, yielding $\mathbf{v}_{CLIP} \in \mathbb{R}^{F \times H_{trn}}$. The motivation to use CLIP as a video feature encoder in our case is twofold: (1) earlier works show its effectiveness in V2A already [14, 44, 48], and (2) its multi-modal nature expands the applicability of our proposal to text-driven video-editing applications. **S3D_{seq}**: a 3D convolutional video encoder [50]. We take the pre-trained version of S3D built in the SparseSync work⁴, for detection of audio-visual temporal offsets, i.e. detecting temporal shifts between the two modalities [24]. The authors of SparseSync originally took S3D pre-trained on the Kinetics 400 dataset for video activity recognition [27], and fine-tuned S3D for the aforementioned offset detection task on AudioSet [18]. This video encoder yields a spatio-temporal tensor of features $\mathbf{v}_{S3D} \in \mathbb{R}^{N_{S3D} \times 512 \times h_{S3D} \times u_{S3D}}$ which we average-pool spatially. We consider these features to be especially sensitive to alignment, since their pre-training task required synchronizing video activity events with the appearance

⁴ <https://github.com/v-iashin/SparseSync>

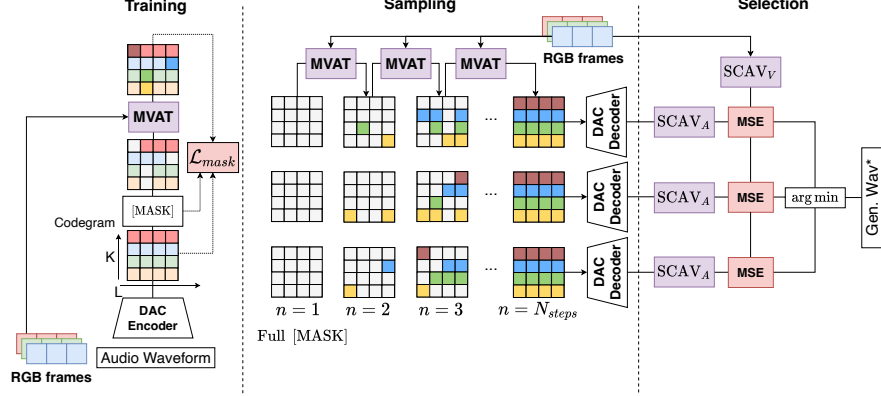


Fig. 2: Overview of the Training, Sampling and Selection parts involved in the MaskVAT framework.

of audio event onsets [24]. Then, following the same procedure as CLIP_{seq} , an MLP projects these features into the same dimensionality of the Transformer blocks H_{trn} , yielding $\bar{\mathbf{v}}_{\text{S3D}} \in \mathbb{R}^{N_{\text{S3D}} \times H_{trn}}$.

AdaLN blocks get a channel-wise concatenation of these feature sequences once they are resampled to have the same lengths, resulting in the visual conditioning tensor $\mathbf{V} = [\Phi_F^{N_{\text{DAC}}}(\mathbf{v}_{\text{CLIP}}); \Phi_{N_{\text{S3D}}}^{N_{\text{DAC}}}(\bar{\mathbf{v}}_{\text{S3D}})]$, where $\Phi_{N_{\text{in}}}^{N_{\text{out}}}$ is the nearest-neighbor resampling operator (i.e. frame repetition) between input N_{in} and output N_{out} lengths. On the other hand, the conditioning tensor for the sequence-to-sequence encoder is the channel-wise concatenation: $\mathbf{V} = [\mathbf{v}_{\text{CLIP}}; \Phi_{N_{\text{S3D}}}^F(\bar{\mathbf{v}}_{\text{S3D}})]$, where S3D features are adjusted to CLIP’s sequence length (F video frames).

Training Setup In a masked token modeling scenario like this, we have a codegram representation $\mathbf{C}^a \in \mathbb{R}^{L \times K}$ (introduced in Sec. 3.1), and a subset of these $L \times K$ tokens is masked with a special token [MASK], as shown in Fig. 2. The mask positions to be replaced by [MASK] in the codegram $M \in \{0, 1\}^{L \times K}$ is determined by a masking scheduler function. For this work, we chose the cosine scheduler due to its proven effectiveness [6], so the probability of each position to be masked is computed as $p = \cos(u)$, where $u \sim U[0, \frac{\pi}{2}]$, from which we obtain $M_{l,k} = \text{Bernoulli}(p)$. Let \mathbf{C}_M^a be the result of applying the mask M to the codegram \mathbf{C}^a and \mathbf{V} be the collection of conditioning features in either format of the three proposed in Fig. 1. The training objective is to minimize the negative log-likelihood, particularly through a cross-entropy loss, for the masked positions [6]:

$$\mathcal{L}_{\text{MaskVAT}} = -\mathbb{E} \left[\sum_{\forall l \in [1, L], \forall k \in [1, K], m_{l,k} = 1} \log p(c_{l,k} | \mathbf{C}_M^a, \mathbf{V}) \right].$$

As introduced earlier, in the sequence-to-sequence and hybrid setups of Fig. 1-b and Fig. 1-c respectively, we also feature a combination of regression and contrastive losses between the Transformer encoder outputs \mathbf{H} and the corresponding projected-BEATs audio embedding sequence \mathbf{B} , where $\mathbf{H}, \mathbf{B} \in \mathbb{R}^{N_{beats} \times H_{trn}}$, N_{beats} is the length of BEATs time-patch sequence, and H_{trn} the hidden size of the transformer encoder. Due to the contrastive loss, we prepend a [CLS] token before injecting the sequence into the Transformer encoder, and select that position as the pooled embedding representative $\bar{\mathbf{H}} \in \mathbb{R}^{H_{trn}}$, which is contrasted against the \mathbf{B} sequence average $\bar{\mathbf{B}} \in \mathbb{R}^{H_{trn}}$ [42]. The total loss to train MaskVAT seq2seq variants (Fig. 1-b and Fig. 1-c) then becomes:

$$\begin{aligned} \mathcal{L}_{\text{reg}} &= \mathbb{E} \left[\frac{1}{N_{beats}} \sum_{t=1}^{N_{beats}} \|\mathbf{H}_t^i - \mathbf{B}_t^i\|_2^2 \right] \\ \mathcal{L}_{\text{cont}} &= -\frac{1}{2} \mathbb{E} \left[\log \left(\frac{\exp(\sigma(\bar{\mathbf{H}}_i, \bar{\mathbf{B}}_i)/\tau)}{\sum_{j=1}^{N_S} \exp(\sigma(\bar{\mathbf{H}}_i, \bar{\mathbf{B}}_j)/\tau)} \right) + \log \left(\frac{\exp(\sigma(\bar{\mathbf{H}}_i, \bar{\mathbf{B}}_i)/\tau)}{\sum_{j=1}^{N_S} \exp(\sigma(\bar{\mathbf{H}}_j, \bar{\mathbf{B}}_i)/\tau)} \right) \right] \\ \mathcal{L}_{\text{MaskVAT}_{\text{seq2seq}}} &= \mathcal{L}_{\text{MaskVAT}_{\text{hybrid}}} = \mathcal{L}_{\text{MaskVAT}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}} + \lambda_{\text{cont}} \mathcal{L}_{\text{cont}}, \end{aligned}$$

where $\tau > 0$ is a learnable temperature parameter, $\sigma(\cdot)$ is the cosine similarity, N_S the contrastive batch size, and $\lambda_{\text{reg}} = \lambda_{\text{cont}} = 1$ by default in all experiments.

Sampling Once we have trained the model to perform unmasking given the \mathbf{C}_M^a tensor, we need a sampling scheme in order to generate new audio codegrams from an initial fully masked instance, as depicted in Fig. 2. Following the sampling process of the original MaskGIT [6], we first determine a number of sampling steps N_{steps} depending on the computational budget. Then, we begin estimating the probability distribution of each codegram position (l, k) over the codewords of the k -th codebook at each step $n \in [1, N_{steps}]$. While computing these probabilities, we also feature classifier-free guidance upon the logits, introducing the coefficient γ [5, 22]. This technique is known to improve generation quality at the expense of sample diversity. Let $l_n^c = \mathcal{M}(\hat{\mathbf{C}}_{M,n}^a, \mathbf{V})$ be the output logits of our MaskVAT \mathcal{M} in conditional form, and $l_n^u = \mathcal{M}(\hat{\mathbf{C}}_{M,n}^a)$ be the unconditional logits that only depend on the estimated and partially-masked codegram $\hat{\mathbf{C}}_{M,n}^a$, our guidance-weighted logits result in $l_n^g = (1 + \gamma)l_n^c - \gamma l_n^u$ [5], where $\gamma \geq 0$. When $\gamma = 0$, this is equivalent to a regular conditional mode in our predictions. Then, for each masked position (l, k) at step n , we sample from the multinomial distribution. With this we generate a candidate token $\hat{c}_{l,k,n}^g$ per masked position at step n . Then, we compute the confidence of each of these sampled tokens based on the log-probability of each position (l, k) . Following previous works [2, 17], we introduce a diversity term δ , which is linearly annealed throughout the N_{steps} as $\delta_n = \delta \cdot (1 - \frac{n+1}{N_{steps}})$. This is used to add noise into to the confidence computation:

$$\text{confidence}(\hat{c}_{l,k,n}^g) = \log p(\hat{c}_{l,k,n}^g | \hat{\mathbf{C}}_{M,n}^a, \mathbf{V}) + \delta_n \cdot \mathcal{N},$$

where $\hat{c}_{l,k,n}^g$ is a token estimate after applying guidance on its logits, at sampling step n , and \mathcal{N} is the i.i.d. noise sample drawn from Gumbel(0,1). This diversity technique has been proven to enhance the generation quality, especially when the number of N_{steps} is increased [2, 17]. In what follows, we select the next K number of tokens to mask at the next sampling iteration $n + 1$ (according to our selected mask scheduler), take the lowest K confidence positions of our estimates, and build a new mask by placing the [MASK] values in these low confidence positions. The remaining ones are kept as successfully unmasked in the estimated codegram $\hat{\mathbf{C}}_{M,n+1}^a$ at the $n + 1$ sampling step . This whole block of operations is repeated until $n = N_{steps}$ (as shown in Fig. 2), and once we get our fully-unmasked estimated codegram $\hat{\mathbf{C}}^a$, we run it through the DAC decoder [30] in order to obtain our generated waveform.

Beam-based selection The sampling process needs some tweaking of the diversity δ , N_{steps} , and guidance γ coefficients upon a validation set in order to produce good quality and diverse outcomes. Nonetheless, each sampling result can vary significantly, also in terms of semantic contents and/or alignment. In order to increase the semantic and time alignments with the input video, we first generate a beam-size B amount of audio instances $\hat{\mathbf{x}}_a^i$ (exemplified with $B = 3$ in Fig. 2). Then, we use a SCAV encoder (see Sec. 2), which we train on the MaskVAT trainset (see Sec. 4), to project our generated audio and its corresponding input video into SCAV feature sequences $\mathbf{E}_i^{\text{scav-a}} \in \mathbb{R}^{N_{\text{scav}} \times 8 \times H_{\text{scav}}}$ and $\mathbf{E}^{\text{scav-v}} \in \mathbb{R}^{N_{\text{scav}} \times H_{\text{scav}}}$ respectively, of common temporal length N_{scav} . Following the SCAV formulation, we use these two feature sequences to choose the generated audio that is closest to the input video $\hat{\mathbf{x}}_a^* = \arg \min_i \text{MSE}(\mathbf{E}^{\text{scav-v}}, \mathbf{E}_i^{\text{scav-a}})$ [46]. We only use this beam strategy with $B = 10$ for the subjective experiments, and provide further detail and evaluation in the Supplementary Material.

4 Experiments

Datasets We train our models on the VGGSound dataset [8], which contains videos curated to maximize the audio-visual correspondence in the videos while remaining unconstrained in the nature of their content. Originally, the dataset contained 200 k video clips in their training partition, but since many videos are not available anymore and we further filter videos based on quality heuristics, we end up with a copy of approximately 155 k video clips. The pre-processing heuristics involve removing videos with silent audio or whose audio length does not match a minimum of 10 seconds, as well as videos featuring less than 15 video frames per second (FPS). Each video ends up being 10 s long, with the audio sampled at 44 kHz, and we only use the audio-visual contents of the dataset and require no labels for the development of our work. In order to build the validation split, we selected 535 video clips re-purposed from the original train split, which amount to approximately 1.5 hours of content.

We use three test partitions to evaluate different aspects of performance for all models. First, we use a subset of the original VGGSound test split to evaluate the generated audio quality and semantic matching with the video (see Sec. 4). In this VGGSound-test split we end up with 12,639 video clips after following the same process as in train split. Secondly, we assess the temporal alignment only on a subset of VGGSound-test specifically filtered to contain only sparse in time-and-space synchronisation signals [24]. This subset contains videos whose audio events and their onsets exhibit strong alignments sparsely in time, like a dog barking in-camera, or a tennis player hitting the ball. This split is named VGGSound-test-sparse. Third, we also test each model’s capabilities out-of-distribution (OOD) on the music synthesis domain by leveraging the MUSIC dataset. The nature of these videos also requires strong audio onset detection from the close-up camera recording of someone playing a musical instrument, therefore we use this dataset to evaluate audio quality, semantic matching, and temporal alignment combined [14, 54, 55]. We extracted 1,908 test video clips, each spanning 10s duration, from a non-overlapped sliding window applied upon 103 test videos effectively downloaded from the MUSIC21-solo test partition ⁵.

Baselines The baselines we choose to compare against our MaskVAT variations are SpecVQGAN [23], Im2Wav [44], V2A-Mapper [48], and Diff-Foley [34], all of them introduced in Sec. 2. V2A-Mapper is considered a state of the art image/video-to-audio generator for VGGSound, hence we consider it our strongest competitor in quality terms. However, it does not model synchronicity explicitly. Therefore, we consider Diff-Foley as the strongest competitor in terms of alignment, since their work emphasizes a solution upon this problem. Since all baselines feature 16 kHz audio generations, except for SpecVQGAN with 22.05 kHz, we also run a band-width extension (BWE) algorithm based on AudioSR [31] upon them. This is done to compare fairly against our MaskVAT, which natively generates 44.1 kHz audio and would have a trivial advantage in quality/fidelity evaluation due to a wide-band vs. full-band comparison.

Implementation Details All models were trained until convergence, tracking the aggregated score of the masked token prediction accuracy, the average FD scores, and the WavCLIP score on the VGGSound validation partition. For MaskVAT_{AdaLN} model variations, we trained with an effective batch size of 200 (across 4 GPUs). For MaskVAT_{Seq2Seq} and MaskVAT_{Hybrid} approaches, we trained with an effective batch size of 400 (across 8 GPUs). Larger batches helped stabilizing convergence in this case, probably due to the contrastive component upon the CLIP+S3D encoder in \mathcal{L}_{cont} (see Sec. 3.2). We used AdamW [33], applying a weight decay of 10^{-5} , a learning rate warmup for the first 3,000 iterations and polynomial decay between $2 \cdot 10^{-4}$ and 10^{-6} for the rest. Additionally, all models were trained with 10% conditioning dropout that replaced visual conditionings by learnable [NULL] tokens representing the unconditional mode for classifier-free guidance [22] (see Sec. 3.2).

⁵ https://github.com/roudimit/MUSIC_dataset

For classifier free-guidance, we explored γ values within $[0, 16]$, and similarly, for diversity, we investigated δ within the same range. The number of sampling steps, N_{steps} , was varied from 8 to 128. As optimal settings across MaskVAT variations, we identified $N_{\text{steps}} = 32$, $\gamma \in [2, 4]$, and $\delta = 8$. We also found beneficial to feature the post-sampling selection with increasing beam size B through objective scans. More information about this hyper-parameter scans is available in the Supplementary Material.

Objective Metrics During the development of our V2A experiments we evaluated three axes of performance: (1) generated audio quality, (2) semantic matching between the generated audio and the original audio/video, and (3) temporal alignment between the generated audio and the original audio/video. For an objective measurement of quality, we rely on computing the Fréchet distance (FD) upon different audio feature extractors, as done by previous audio synthesis works [14, 19, 28, 32, 40, 48]. FD rates the trade-off between quality and diversity attained in the generated audio. Moreover, each audio feature extractor used to compute a different FD offers a different focus on aspects of the generated audio that fit those of the ground truth [19]. In this work, we leverage three types of embeddings to compute FDs. **FAD**: Fréchet audio distance [28] based on VGGish [21] is a popular metric operating on 16 kHz signals. We also propose two additional full-band FD metrics, provided the importance of analyzing how well a model can generate the full 44.1 kHz band-width for general audio [40]. **FDM**: This representation is frame-based, with each frame containing a window of 2048 samples and a shift of 512 samples. We extract 128 filter-banks and 64 MFCCs, so the embeddings to compute are 64-dimensional. **FDD**: We leverage the DAC codec 8-dimensional embeddings across the K RVQ levels after quantization ($8 \times 9 = 72$ dimensions per frame), prior to the residual summation at the input of the decoder. This is also a frame-based representation, with a wider receptive field than the MFCC one.

To assess the semantic matching, we propose two metrics that measure the proximity of the signals in the highly semantic CLIP space, as other generative works proposed [51]. Nonetheless, since we generate audio to be evaluated instead of images or text/labels, we leverage the audio waveform front-end Wav2CLIP [49] to project our generated outcomes into CLIP space. Wav2CLIP was precisely trained to project 16 kHz audio waveforms of variable length into a fixed embedding in the CLIP space from audio-video data [42]. Then, we measure the cosine similarity between the two projected embeddings. We implement two ways of projecting through Wav2CLIP to measure the proximity. **WaveCLIP (WC)**: we project both the generated audio and the ground truth audio that came originally with the video. Then, we measure the cosine similarity of both embeddings L2 normalized in CLIP space. **CyleCLIP (CC)**: we project only the generated audio and compare it against the average video CLIP embedding, both L2-normalized, which evaluates how aligned is the generated waveform with the original visual content.

Finally, we measure the degree of alignment of generated audios with two metrics. We compute the self-similarity-based audio novelty, reported as **Novelty Score** (NS) [16]. This is obtained as the Pearson correlation coefficient between the self-similarity audio novelty curves of the BEATs-encoded sequences [11] for the generated and ground truth audio signals. Note that video prompts are not involved in this metric, hence it is an audio-to-audio comparison. We also consider the **SparseSync** (SS) metric, based on the synchronization model proposed in [24] (see Sec. 2), as the mean offset prediction originally proposed in [24] between the prompted video with our generated audio: $(\text{video}_i, \text{genaudio}_i)$.

Subjective Evaluation We also set up a subjective test that features three sections explicitly asking 19 human subjects (with 11 audio processing experts) to rate. **Audio quality and relevance** (as semantic matching): we have selected samples from VGGSound-test containing both sparse events [24] (attack sounds with distinct onsets) and dense events (sustained sounds with temporal evolution). We take into account both 16 kHz and 44.1 kHz versions for all models and references for comparison. This was done by running the bandwidth-extension algorithm upon the baselines, or by resampling the references or MaskVAT generation down to 16 kHz. **Audio-video alignment**: we have selected samples from VGGSound-Test-Sparse containing only sparse events. In order to focus on the audio-visual synchronicity, we use samples of 16 kHz only. **Overall quality** (mix of audio quality, semantic matching, and temporal alignment): we select samples from the MUSIC dataset, because the audio is highly correlated with the video and the music audio is of high quality, which is challenging to generate. Users are asked to rate with all the criteria in mind (quality + semantic + alignment). Here we keep the original sample rate for all the samples such that the overall advantage of a model can be evaluated. More details about the subjective test setup can be found in the Supplementary Material.

5 Results

Tables 1 and 2 show the results evaluated with objective metrics for VGGSound and MUSIC test sets, respectively. Our proposed models beat all the baselines in FD terms across the full-band front-ends (FDD, FDM), exhibiting an advantage in natively modeling 44.1 kHz upon DAC. Nevertheless, when comparing the more prominent low-band content with the FAD metric, MaskVAT falls behind V2A-mapper. This may imply that the usage of a lossy codec is a quality upper bound (this is clear from Table 1, where the DAC reconstruction is already worse than V2A in FAD). Since SpecVQGAN generates 22.05kHz audio, it may make it advantageous in terms of bandwidth, compared to other baselines, which is reflected in the FDD and FDM scores. However, SpecVQGAN-BWE is not as advantageous, probably due to the difficulty of BWE given existing audio artifacts in the generated low-band content. In semantic terms, our MaskVAT_{Seq2Seq} and MaskVAT_{Hybrid} win over the baselines with quite a margin, indicating strong alignment of generated audio with respect to the input video, potentially due

Table 1: Objective Results on VGGSound-Test. Baselines featuring bandwidth extension to 44.1 kHz have a BWE suffix (e.g. V2A-Mapper-BWE). MaskVAT_{AdaLN-A}: only CLIP conditioning. MaskVAT_{AdaLN-B}: CLIP and S3D conditioning.

Model	Quality			Semantic		Alignment	
	FDD ↓	FDM ↓	FAD ↓	WC ↑	CC ↑	NS ↑	SS ↓
DAC reconstruct	0.04	0.10	1.06	0.90	0.126	0.97	0.46
Diff-Foley	1.09	30.8	8.60	0.35	0.087	0.07	0.57
Diff-Foley-BWE	1.22	22.4	7.54	–	–	–	–
Im2Wav	0.45	11.9	6.21	0.45	0.116	0.00	0.68
Im2Wav-BWE	0.45	7.24	7.89	–	–	–	–
SpecVQGAN	0.26	7.75	5.27	0.33	0.080	0.02	0.67
SpecVQGAN-BWE	0.42	8.10	5.75	–	–	–	–
V2A-Mapper	0.45	14.1	0.89	0.47	0.124	-0.01	0.68
V2A-Mapper-BWE	0.24	2.72	0.84	–	–	–	–
MaskVAT _{AdaLN-A}	0.06	1.21	3.83	0.48	0.123	0.05	0.60
MaskVAT _{AdaLN-B}	0.05	0.88	3.39	0.50	0.123	0.16	0.43
MaskVAT _{Seq2Seq}	0.06	0.60	1.51	0.55	0.140	0.05	0.63
MaskVAT _{Hybrid}	0.08	0.88	2.04	0.55	0.136	0.17	0.40

Table 2: Objective Results on MUSIC-Test. Same naming conventions apply as in VGGSound-Test Results.

Model	Quality			Semantic		Alignment	
	FDD ↓	FDM ↓	FAD ↓	WC ↑	CC ↑	NS ↑	SS ↓
DAC reconstruct	0.03	0.17	7.99	0.88	0.131	0.94	0.63
Diff-Foley	0.63	24.2	46.3	0.43	0.09	0.02	0.66
Diff-Foley-BWE	0.52	22.5	47.7	–	–	–	–
Im2Wav	0.49	14.1	38.4	0.38	0.08	0.00	0.69
Im2Wav-BWE	0.49	6.63	44.7	–	–	–	–
SpecVQGAN	0.27	7.06	43.2	0.29	0.07	0.01	0.68
SpecVQGAN-BWE	0.41	7.18	44.5	–	–	–	–
V2A-Mapper	0.55	14.4	12.8	0.56	0.124	0.01	0.68
V2A-Mapper-BWE	0.30	4.81	12.1	–	–	–	–
MaskVAT _{AdaLN-A}	0.08	1.60	22.8	0.53	0.123	0.02	0.67
MaskVAT _{AdaLN-B}	0.07	1.15	25.3	0.57	0.123	0.16	0.61
MaskVAT _{Seq2Seq}	0.07	1.02	15.8	0.63	0.137	0.06	0.66
MaskVAT _{Hybrid}	0.09	1.23	19.7	0.62	0.135	0.16	0.62

to the learnable intermediate features of the auxiliary losses. For alignment, our model wins when leveraging S3D features injected into AdaLN blocks in the architecture (MaskVAT_{AdaLN} and MaskVAT_{Hybrid}). Hence, leveraging a hybrid setup with MaskVAT_{Hybrid} results in a winner combination to best balance the

Table 3: Mean opinion score [45] results on VGGSound, VGGSound_{sparse}, and MUSIC test sets considering all subjects (top) and only audio processing experts (bottom, †).

Model	VGGSound		VGGSound _{sparse}	MUSIC
	Fidelity \uparrow	Relevance \uparrow	Alignment \uparrow	Overall \uparrow
Reference	4.09 \pm 0.17	4.21 \pm 0.18	4.82 \pm 0.08	4.3 \pm 0.21
Diff-Foley	1.58 \pm 0.16	2.43 \pm 0.21	2.14 \pm 0.22	1.49 \pm 0.19
V2A-Mapper	3.00 \pm 0.19	3.50 \pm 0.20	2.12 \pm 0.20	2.44 \pm 0.26
MaskVAT _{Hybrid}	2.76 \pm 0.20	3.33 \pm 0.21	3.81 \pm 0.21	2.91 \pm 0.27
Reference†	3.97 \pm 0.23	4.09 \pm 0.26	4.82 \pm 0.10	4.15 \pm 0.30
Diff-Foley†	1.47 \pm 0.19	2.41 \pm 0.28	2.15 \pm 0.29	1.48 \pm 0.24
V2A-Mapper†	2.92 \pm 0.22	3.41 \pm 0.25	2.13 \pm 0.23	2.27 \pm 0.31
MaskVAT _{Hybrid} †	2.88 \pm 0.25	3.39 \pm 0.27	3.76 \pm 0.28	2.91 \pm 0.33

overall trade-off in quality, semantic and temporal alignment results. This motivates our selection of MaskVAT_{Hybrid} for the subjective evaluation.

Table 3 shows the subjective evaluation results. We see that MaskVAT_{Hybrid} outperforms all other models in the specialized categories of Alignment and Overall, and that it is still competitive with V2A-Mapper in Fidelity and Relevance. Interestingly, the gap in the latter two categories considerably shrinks when only expert listeners are considered. That is not the case with the Alignment and Overall categories, where MaskVAT_{Hybrid} remains a clear winner. Of special mention is the Alignment category, which highlights the benefit of the proposed approach for synchronicity.

6 Conclusion

In this work we proposed a masked generative video-to-audio Transformer, a model that generates audio based on an input silent video. MaskVAT makes special emphasis on tackling temporal alignment between the generated audio and the input video. Our solution connects a state of the art full-band general audio codec to ensure high quality outcomes, with a sequence-to-sequence masked-token generative approach, which is driven by pre-trained semantic and alignment features. We find our MaskVAT_{Hybrid} to best balance the overall trade-off in quality, semantic and temporal alignment results. Moreover, we also leverage a post-sampling selection strategy that minimizes the distance between the generated audio and the source input video. Our model outperforms existing solutions, exhibiting strong temporal alignment in the audio generations, which are fundamental in the overall resulting quality of video-to-audio generation. Furthermore, MaskVAT shows competitive performance in terms of generated audio quality and semantic relevance against previously proposed systems that leverage the inter-connection of strong foundational models to perform V2A.

7 Acknowledgements

The authors want to thank Heng Wang, Xiaoyu Liu, Gautam Bhattacharya, Roy Fejgin, and David Cooper for insightful discussions about generative strategies and evaluation methodology.

References

1. Agostinelli, A., Denk, T.I., Borsos, Z., Engel, J., Verzetti, M., Caillon, A., Huang, Q., Jansen, A., Roberts, A., Tagliasacchi, M., et al.: Musiclm: Generating music from text. arXiv preprint arXiv:2301.11325 (2023) [4](#)
2. Besnier, V., Chen, M.: A pytorch reproduction of masked generative image transformer. arXiv preprint arXiv:2310.14400 (2023) [8](#), [9](#)
3. Borsos, Z., Marinier, R., Vincent, D., Kharitonov, E., Pietquin, O., Sharifi, M., Roblek, D., Teboul, O., Grangier, D., Tagliasacchi, M., et al.: Audioldm: a language modeling approach to audio generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* (2023) [4](#)
4. Borsos, Z., Sharifi, M., Vincent, D., Kharitonov, E., Zeghidour, N., Tagliasacchi, M.: Soundstorm: Efficient parallel audio generation. arXiv preprint arXiv:2305.09636 (2023) [2](#), [4](#), [5](#)
5. Chang, H., Zhang, H., Barber, J., Maschinot, A., Lezama, J., Jiang, L., Yang, M.H., Murphy, K., Freeman, W.T., Rubinstein, M., et al.: Muse: Text-to-image generation via masked generative transformers. arXiv preprint arXiv:2301.00704 (2023) [8](#)
6. Chang, H., Zhang, H., Jiang, L., Liu, C., Freeman, W.T.: Maskgit: Masked generative image transformer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11315–11325 (2022) [2](#), [4](#), [5](#), [7](#), [8](#)
7. Chen, H., Xie, W., Afouras, T., Nagrani, A., Vedaldi, A., Zisserman, A.: Audio-visual synchronisation in the wild. arXiv preprint arXiv:2112.04432 (2021) [3](#)
8. Chen, H., Xie, W., Vedaldi, A., Zisserman, A.: Vggsound: A large-scale audio-visual dataset. In: *Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 721–725. IEEE (2020) [9](#)
9. Chen, L., Srivastava, S., Duan, Z., Xu, C.: Deep cross-modal audio-visual generation. In: *Proceedings of the on Thematic Workshops of ACM Multimedia 2017*. pp. 349–357 (2017) [1](#), [2](#)
10. Chen, P., Zhang, Y., Tan, M., Xiao, H., Huang, D., Gan, C.: Generating visually aligned sound from videos. *IEEE Transactions on Image Processing* **29**, 8292–8302 (2020) [1](#), [2](#)
11. Chen, S., Wu, Y., Wang, C., Liu, S., Tompkins, D., Chen, Z., Wei, F.: Beats: Audio pre-training with acoustic tokenizers. arXiv preprint arXiv:2212.09058 (2022) [3](#), [6](#), [12](#)
12. Copet, J., Kreuk, F., Gat, I., Remez, T., Kant, D., Synnaeve, G., Adi, Y., Défossez, A.: Simple and controllable music generation. *Advances in Neural Information Processing Systems* **36** (2024) [4](#)
13. Défossez, A., Copet, J., Synnaeve, G., Adi, Y.: High fidelity neural audio compression. arXiv preprint arXiv:2210.13438 (2022) [4](#)
14. Dong, H.W., Liu, X., Pons, J., Bhattacharya, G., Pascual, S., Serrà, J., Berg-Kirkpatrick, T., McAuley, J.: Clipsonic: Text-to-audio synthesis with unlabeled

- videos and pretrained language-vision models. In: 2023 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA). pp. 1–5. IEEE (2023) [1](#), [2](#), [3](#), [6](#), [10](#), [11](#)
15. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12873–12883 (2021) [3](#)
 16. Foote, J.: Automatic audio segmentation using a measure of audio novelty. In: In Proc. of the IEEE International Conference on Multimedia and Expo (ICME). vol. 1, pp. 452–455. IEEE (2000) [12](#)
 17. Garcia, H.F., Seetharaman, P., Kumar, R., Pardo, B.: Vampnet: Music generation via masked acoustic token modeling. arXiv preprint arXiv:2307.04686 (2023) [2](#), [4](#), [5](#), [8](#), [9](#)
 18. Gemmeke, J.F., Ellis, D.P., Freedman, D., Jansen, A., Lawrence, W., Moore, R.C., Plakal, M., Ritter, M.: Audio set: An ontology and human-labeled dataset for audio events. In: Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP). pp. 776–780. IEEE (2017) [6](#)
 19. Gui, A., Gamper, H., Braun, S., Emmanouilidou, D.: Adapting frechet audio distance for generative music evaluation. arXiv preprint arXiv:2311.01616 (2023) [11](#)
 20. Hao, W., Zhang, Z., Guan, H.: Cmcgan: A uniform framework for cross-modal visual-audio mutual generation. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018) [1](#), [2](#)
 21. Hershey, S., Chaudhuri, S., Ellis, D.P., Gemmeke, J.F., Jansen, A., Moore, R.C., Plakal, M., Platt, D., Saurous, R.A., Seybold, B., et al.: Cnn architectures for large-scale audio classification. In: Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP). pp. 131–135. IEEE (2017) [11](#)
 22. Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022) [8](#), [10](#)
 23. Iashin, V., Rahtu, E.: Taming visually guided sound generation. arXiv preprint arXiv:2110.08791 (2021) [1](#), [2](#), [3](#), [10](#)
 24. Iashin, V., Xie, W., Rahtu, E., Zisserman, A.: Sparse in space and time: Audio-visual synchronisation with trainable selectors. arXiv preprint arXiv:2210.07055 (2022) [3](#), [6](#), [7](#), [10](#), [12](#)
 25. Jeong, Y., Ryoo, W., Lee, S., Seo, D., Byeon, W., Kim, S., Kim, J.: The Power of Sound (TPoS): Audio Reactive Video Generation with Stable Diffusion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7822–7832 (2023) [1](#)
 26. Kalchbrenner, N., Elsen, E., Simonyan, K., Noury, S., Casagrande, N., Lockhart, E., Stimberg, F., Oord, A., Dieleman, S., Kavukcuoglu, K.: Efficient neural audio synthesis. In: International Conference on Machine Learning. pp. 2410–2419. PMLR (2018) [4](#)
 27. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017) [6](#)
 28. Kilgour, K., Zuluaga, M., Roblek, D., Sharifi, M.: Fr`echet audio distance: A metric for evaluating music enhancement algorithms. arXiv preprint arXiv:1812.08466 (2018) [11](#)
 29. Kreuk, F., Synnaeve, G., Polyak, A., Singer, U., Défossez, A., Copet, J., Parikh, D., Taigman, Y., Adi, Y.: Audiogen: Textually guided audio generation. arXiv preprint arXiv:2209.15352 (2022) [4](#)

30. Kumar, R., Seetharaman, P., Luebs, A., Kumar, I., Kumar, K.: High-fidelity audio compression with improved rvqgan. *Advances in Neural Information Processing Systems* **36** (2024) 4, 5, 9
31. Liu, H., Chen, K., Tian, Q., Wang, W., Plumbley, M.D.: Audiosr: Versatile audio super-resolution at scale. *arXiv preprint arXiv:2309.07314* (2023) 10
32. Liu, H., Tian, Q., Yuan, Y., Liu, X., Mei, X., Kong, Q., Wang, Y., Wang, W., Wang, Y., Plumbley, M.D.: Audioldm 2: Learning holistic audio generation with self-supervised pretraining. *arXiv preprint arXiv:2308.05734* (2023) 3, 11
33. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017) 10
34. Luo, S., Yan, C., Hu, C., Zhao, H.: Diff-foley: Synchronized video-to-audio synthesis with latent diffusion models. *Advances in Neural Information Processing Systems* **36** (2024) 1, 2, 3, 10
35. Mehri, S., Kumar, K., Gulrajani, I., Kumar, R., Jain, S., Sotelo, J., Courville, A., Bengio, Y.: Samplernn: An unconditional end-to-end neural audio generation model. *arXiv preprint arXiv:1612.07837* (2016) 4
36. Mei, X., Nagaraja, V., Lan, G.L., Ni, Z., Chang, E., Shi, Y., Chandra, V.: Foleygen: Visually-guided audio generation. *arXiv preprint arXiv:2309.10537* (2023) 2, 3
37. Oord, A., Li, Y., Babuschkin, I., Simonyan, K., Vinyals, O., Kavukcuoglu, K., Driessche, G., Lockhart, E., Cobo, L., Stimberg, F., et al.: Parallel wavenet: Fast high-fidelity speech synthesis. In: *International conference on machine learning*. pp. 3918–3926. PMLR (2018) 4
38. Oord, A.v.d., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., Kavukcuoglu, K.: Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499* (2016) 4
39. Owens, A., Isola, P., McDermott, J., Torralba, A., Adelson, E.H., Freeman, W.T.: Visually indicated sounds. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2405–2413 (2016) 2
40. Pascual, S., Bhattacharya, G., Yeh, C., Pons, J., Serrà, J.: Full-band general audio synthesis with score-based diffusion. In: *Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 1–5. IEEE (2023) 2, 11
41. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4195–4205 (2023) 5
42. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PMLR (2021) 3, 8, 11
43. Ruan, L., Ma, Y., Yang, H., He, H., Liu, B., Fu, J., Yuan, N.J., Jin, Q., Guo, B.: Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10219–10228 (2023) 2
44. Sheffer, R., Adi, Y.: I hear your true colors: Image guided audio generation. In: *Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 1–5. IEEE (2023) 1, 2, 3, 6, 10
45. Taubert, S.: mean-opinion-score (Aug 2023). <https://doi.org/10.5281/zenodo.8238259>, <https://github.com/stefantaubert/mean-opinion-score> 14
46. Tsiamas, I., Pascual, S., Yeh, C., Serrà, J.: Sequential contrastive audio-visual learning. *arXiv preprint arXiv:2407.05782* (2024) 3, 9

47. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017) [3](#)
48. Wang, H., Ma, J., Pascual, S., Cartwright, R., Cai, W.: V2a-mapper: A lightweight solution for vision-to-audio generation by connecting foundation models. *arXiv preprint arXiv:2308.09300* (2023) [1](#), [2](#), [3](#), [6](#), [10](#), [11](#)
49. Wu, H.H., Seetharaman, P., Kumar, K., Bello, J.P.: Wav2clip: Learning robust audio representations from clip. In: *Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 4563–4567. IEEE (2022) [11](#)
50. Xie, S., Sun, C., Huang, J., Tu, Z., Murphy, K.: Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 305–321 (2018) [6](#)
51. Yariv, G., Gat, I., Benaim, S., Wolf, L., Schwartz, I., Adi, Y.: Diverse and aligned audio-to-video generation via text-to-video model adaptation. *arXiv preprint arXiv:2309.16429* (2023) [11](#)
52. Yusong Wu, Ke Chen, T.Z.Y.H.M.N.T.B.K., Dubnov, S.: Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. *arXiv preprint arXiv:2211.06687* (2024) [3](#)
53. Zeghidour, N., Luebs, A., Omran, A., Skoglund, J., Tagliasacchi, M.: Soundstream: An end-to-end neural audio codec. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **30**, 495–507 (2021) [4](#)
54. Zhao, H., Gan, C., Ma, W.C., Torralba, A.: The sound of motions. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1735–1744 (2019) [10](#)
55. Zhao, H., Gan, C., Rouditchenko, A., Vondrick, C., McDermott, J., Torralba, A.: The sound of pixels. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 570–586 (2018) [10](#)
56. Zhou, Y., Wang, Z., Fang, C., Bui, T., Berg, T.L.: Visual to sound: Generating natural sound for videos in the wild. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3550–3558 (2018) [1](#), [2](#)
57. Ziv, A., Gat, I., Lan, G.L., Remez, T., Kreuk, F., Défossez, A., Copet, J., Synnaeve, G., Adi, Y.: Masked audio generation using a single non-autoregressive transformer. *arXiv preprint arXiv:2401.04577* (2024) [4](#)