# AudioScopeV2: Audio-Visual Attention Architectures for Calibrated Open-Domain On-Screen Sound Separation

Efthymios Tzinis<sup>1,2\*</sup>, Scott Wisdom<sup>1</sup>, Tal Remez<sup>1</sup>, and John R. Hershey<sup>1</sup>

<sup>1</sup> Google Research <sup>2</sup> University of Illinois Urbana-Champaign etzinis20illinois.edu, {scottwisdom,johnhershey}0google.com

Abstract. We introduce AudioScopeV2, a state-of-the-art universal audio-visual on-screen sound separation system which is capable of learning to separate sounds and associate them with on-screen objects by looking at in-the-wild videos. We identify several limitations of previous work on audio-visual on-screen sound separation, including the coarse resolution of spatio-temporal attention, poor convergence of the audio separation model, limited variety in training and evaluation data, and failure to account for the trade off between preservation of on-screen sounds and suppression of off-screen sounds. We provide solutions to all of these issues. Our proposed cross-modal and self-attention network architectures capture audio-visual dependencies at a finer resolution over time, and we also propose efficient separable variants that are capable of scaling to longer videos without sacrificing much performance. We also find that pre-training the separation model only on audio greatly improves results. For training and evaluation, we collected new human annotations of onscreen sounds from a large database of in-the-wild videos (YFCC100M). This new dataset is more diverse and challenging. Finally, we propose a calibration procedure that allows exact tuning of on-screen reconstruction versus off-screen suppression, which greatly simplifies comparing performance between models with different operating points. Overall, our experimental results show marked improvements in on-screen separation performance under much more general conditions than previous methods with minimal additional computational complexity.

Keywords: Audio-visual sound separation, self-attention

## 1 Introduction

Humans are able to effortlessly perceive sounds in a noisy scene, and associate them with any corresponding visible objects. In audio processing, a corresponding challenge is to isolate sound sources from a mixture waveform and identify the associated visual appearance of each sound source. In this paper, we target the task of *universal audio-visual on-screen sound separation*, where the goal is to

<sup>\*</sup> Work done during an internship at Google Research.



(a) General audio-visual on-screen sound separation task.

Fig. 1: Left: illustration of the task where our models make no assumptions about the existence, type, or count (up to a maximum) of on-screen and off-screen sources. Right: illustration of our proposed calibration procedure for on-screen separation models, which allows setting a specified tolerance level for off-screen sound suppression which allows more user control and easier model comparison.

recover only the sounds that originate from on-screen objects, regardless of the types of on-screen and off-screen objects as illustrated in Figure 1a.

This is a difficult task for numerous reasons. In stark contrast to visual objects that generally occupy distinct regions of pixels, sound sources are superimposed in the time domain. This imposes a challenge for unsupervised learning of audiovisual separation, because unlike their visual counterparts, the audio sources in a scene cannot be easily selected and aligned with video objects. Therefore, separating the constituent sources is needed, by conditioning separation on selected video objects and/or separating the audio *a priori* before associating the sounds with video objects. The *a priori* separation of the sounds, which we pursue here, has a few advantages. Thanks to recent work [47], it can be learned in an unsupervised way, and it can handle an unknown number of sounds, including those that do not appear on-screen. Also, the individual separated sounds are available to downstream processes, in addition to the on-screen estimate.

Despite remarkable progress in the field of on-screen sound separation, most of these works are constrained to isolating only a specific set of sound classes that can appear on-screen such as speech [1, 10, 14] or music [11, 12]. Although this strategy works well under a restricted domain, where such labeled data are available, the reliance on human labels precludes scaling to large open-domain data. Recent works have started to expand beyond music and speech to a wider variety of classes, such as using visual scene graphs to model audio-visual relationships [7], but this approach still requires labeled data to train a supervised object detector. Although the seminal works in on-screen sound separation proposed models that were somewhat invariant to the types of sources [12, 32], those systems were unable to be trained with real world videos mainly because they needed labeled videos in which the sources always appeared on-screen during training.

Recently, AudioScope [43] addressed several of the aforementioned problems using *mixture invariant training* (MixIT) [47] with synthetic mixtures of video soundtracks. The derived sources-to-soundtracks assignments from MixIT were used as pseudo-targets to self-supervise the training of an audio-visual sound separation model from in-the-wild videos, without requiring object detection modules or assuming that all sources have to be on-screen. However, AudioScope still suffers from generalization issues since it relies on training data filtered by an unsupervised audio-visual coincidence model [22], which limits its generalization ability, as we show in our experiments. We also hypothesize that AudioScope's performance is limited by the simplicity of its visually guided spatio-temporal attention layer [5], and the low temporal resolution (one frame per second) of its visual model. These factors may prevent AudioScope from capturing synchronization features which can be crucial for detecting the audio-visual interplay [2, 19, 26]. Another limitation of AudioScope is the lack of ability to trade off between reconstruction of on-screen sounds and suppression of off-screen sounds. These models achieve an arbitrary operating point during training, which makes comparing performance between different models difficult (see Figure 1b).

We propose solutions for all of the aforementioned problems and limitations:

- 1. AudioScopeV2 leverages richer cross-modal and self-attention network architectures that capture audio-visual dependencies at a finer time resolution, as well as efficient separable variants that are capable of scaling to longer videos without sacrificing much performance (Section 3). We also find that pre-training the audio separation model using MixIT greatly improves results.
- 2. We provide a new dataset, for which we collected new human annotations of onscreen sounds from a large database of in-the-wild videos (YFCC100M [38]), described in Section 4.1. We show that our new proposed models both generalize and perform better on the more diverse and challenging evaluations sets compared to previous state-of-the-art methods.
- 3. We propose a novel calibration procedure (Section 4.4) that allows precise tuning of on-screen reconstruction versus off-screen suppression, that can also be used to greatly simplify model comparison across different models that each have their own operating point.

Dataset recipes, demos, and other supplementary material are available online google-research.github.io/sound-separation/papers/audioscope-v2.

# 2 Relation to Prior Work

Joint perception of audio and video modalities is not trivial, in part due to the problems of alignment between corresponding representations in each modality. Nevertheless, a variety of works have shown promising results using multimodal neural network architectures [2, 6–8, 13, 43, 48]. Audio-visual sound separation [1,10,20], and specifically separation of on-screen versus off-screen sounds [32], has enjoyed remarkable performance improvements since the initial works. Important innovations have included using localization of objects [29, 48, 51], forc-

ing consistency between audio and visual representations [3, 13, 14, 28], weaklysupervised [34] and self-supervised [2, 26, 35, 43] approaches.

Recent work has shown that it is possible to train an open-domain audio-only universal sound separation model using a mask-based convolutional architecture regardless of the category of the sound [23, 41]. A related direction is to extract sources of interest by conditioning separation networks using identity or multimodal cues. This has yielded performance improvements for speech [45] as well as universal source separation [16, 24, 30, 31, 42]. However, these experiments relied on having sufficient supervised training data and were evaluated only on test sets with similar environmental conditions and sound distributions. In order to extend the reach of this approach, methods have been proposed to train separation models with no access to ground truth clean sources by utilizing weak class labels [33], the spatial separability of the sources [9, 36, 40] and self-supervision in the form of MixIT [47]. These methods make it possible to learn separation of signals well outside the domains for which isolated source databases exist.

An open question in audio-visual correspondence models concerns the level of processing at which audio and video objects can be aligned. Typically audio-visual models have used high-level features at the output of neural networks to estimate correspondence between audio and video signals [7, 22, 26, 28, 39, 43]. Such high-level representations may tend to focus on semantic information about the class of objects and sounds, especially when the features are computed at low video frame rates. Such methods may work well for single instances of a class of object or sound, but may struggle with identification for multiple instances of a class, or for classes not seen during training. In contrast, there may be significant information in the correspondence between lower-level features. Mutual information between low-level features was used for audio-visual localization [19], and several more recent works have shown promising results for self-supervised audio-visual learning using low-level motion [51] and optical flow [2] features. Such features may help with generalization and instance-level correspondence by detecting synchronous dynamics of the audio and video, regardless of their semantic class.

Attention mechanisms can align representations across modalities, both at the level of semantic association and in terms of low-level correspondences. An attention-based framework was recently used to modulate audio representations using motion-based visual features [51] for separation and localization. Conversely, modulating video features based on audio embeddings has also been used for speech separation [28] as well as in AudioScope's spatio-temporal attention module [43]. Other works combined self-attention layers [44] for modeling intermodality temporal patterns, as well as cross-modal attention modules for intramodality associations [8,48,49], for sound localization and representation learning. One issue with self-attention is that its complexity grows quadratically with the dimensionality of the input length. We therefore propose a separable variants of our proposed architectures that factorize attention across different dimensions and modalities. This strategy allows us to achieve similar performance to full selfattention with a much lower computational footprint. Other separable attention



mechanisms have emerged recently [6,50], but our approach differs in that we process and capture intra-modality patterns from both audio and video features.

Fig. 2: Illustration of AudioScopeV2's architecture and its training procedure.

## 3 Model Architecture

Figure 2 illustrates the architecture and training procedure of AudioScopeV2. The run-time architecture takes as input the video frames and a mixture waveform x. The audio separation module estimates M = 4 sources. An audio embedding network is run on each of the estimated sources, producing audio embeddings  $Z_{\rm A}$ . In parallel, an image embedding network processes each input video frame independently and produces the visual embeddings  $Z_{\rm V}$ . The audio and video embeddings are fed to an audio-visual attention network, for which we propose a family of attention-based architectures. The output of the audio-visual attention network is passed to a final on-screen classifier head which produces a probability  $\hat{y}_m$  corresponding to the event that the source m originates from an on-screen object. Finally, the probabilities  $\hat{y}_m$  are used as weights to mix the separated sources  $\hat{s}_m$  together, producing an estimate of on-screen audio. The training procedure for this model is also illustrated, where we create mixtures of mixtures (MoMs) for audio by adding the soundtrack from another random video. A MixIT SNR loss is computed by finding the best combination of estimated sources to approximate each one of the reference mixtures in terms of SNR. The assignments of these best combinations are used as pseudo-labels in the classification loss. We describe each of these components in more detail below.

## 3.1 Separation module

The separation module  $\mathcal{M}^{S}$  uses a dilated convolutional architecture [23] with learnable encoder and decoder, which was also used in [43]. This module takes as

input a mixture waveform  $x \in \mathbb{R}^{T'}$ , estimates M masks in the encoded latent space, and outputs M estimated source waveforms  $\hat{s} \in \mathbb{R}^{M \times T'}$ . The latter are forced to add up to the input mixture through a mixture consistency layer [46]. However, in contrast to AudioScope [43], our separation model is not conditioned on global visual embeddings, for two reasons. First, visual conditioning was not shown to be effective [43], and second, this allows us to use MixIT to pre-train the separation module on all YFCC100M [38] audio tracks to provide a better initialization for training the audio-visual model. Thus, AudioScope<sup>\*</sup> refers to our AudioScope implementation with the aforementioned source separation module.

#### 3.2 Audio and video embedding networks

Features are extracted for the M estimated sources  $\hat{s}$  and the corresponding T input video frames (128 × 128 pixels), using a MobileNetV1 architecture [21], as in [43]. The audio encoder takes as input the log mel-scale spectrogram of the separated source waveforms  $\hat{s}$ , and audio features are extracted from the 23rd layer. The visual embedding network is applied to each of the T input video frames independently and extracts features with  $8 \times 8$  spatial locations. Audio and video features are converted to a common depth D = 128 with a dense layer.

#### 3.3 Audio-visual attention

We propose attention mechanisms to identify dependencies across the M estimated audio sources, space, and time between the audio features  $Z_A \in \mathbb{R}^{M \times T \times D}$  and the video features  $Z_V \in \mathbb{R}^{G \times T \times D}$ . Our video encoder provides  $G = 8^2$  spatial locations, and the time dimension T is shared across both tensors. Specifically, we propose to use *audio-visual self-attention* (SA) [44], which treats the audio sources and visual locations as a joint attention space (see Section 3.3), and *cross-modal attention* (CMA) layers, which perform attention between audio sources and visual locations, but avoid uni-modal attention between sources and between spatial locations (see Section 3.3). For both the SA and CMA attention we consider two settings: a *joint attention* setting, in which attention operates jointly over time and space / sources, and a *separable attention* setting, in which attention across time is interleaved with attention across spatial locations and sources. The joint attention scales quadratically with the product of dimensions from all of the axes which is computed over whereas the more efficient separable variation factorizes the operation across each axes individually, which makes an important difference in practice (see Section 5.2).

In the following formulations we use a slightly more general version of an attention layer [5] to show how attention operates across different axes of the corresponding tensors. Attention computes similarities between a packed tensor of queries  $Q \in \mathbb{R}^{X_Q \times T_Q \times D}$  w.r.t. some packed keys  $K \in \mathbb{R}^{X_K \times T_K \times D}$ , where D is the depth dimensionality of the tensors. The similarities are computed using a generalized version of the typical tensor inner product  $\langle Z_1, Z_2 \rangle_A$ , which reduces across the specified dimensions  $\mathcal{A}$  of the second tensor  $Z_2$ . Note that we assume that  $Z_1$ 's dimensions are a subset of  $Z_2$ 's. By using a scaled tensor



(a) Self-Attention (SA). (b) Cross-modal attention (CMA). (c) Separable.

Fig. 3: AudioScopeV2's attention architectures for audio-visual alignment and feature extraction. Tensors are depicted by omitting batch and depth dimensions.

inner-product [44] and a softmax<sub> $\mathcal{A}$ </sub> activation that averages over the dimensions specified by  $\mathcal{A}$  at the output of the tensor product  $\langle K, Q \rangle_{\{D\}}$ , we produce the resulting similarity tensor which modulates the values  $V \in \mathbb{R}^{X_V \times T_V \times D}$ :

$$\operatorname{Att}_{\mathcal{A}}(Q, K, V) = \langle \alpha, f_{V}(V) \rangle_{\mathcal{A}}, \ \alpha = \operatorname{softmax}_{\mathcal{A}}\left(\frac{1}{\sqrt{D}} \langle f_{K}(K), f_{Q}(Q) \rangle_{\{D\}}\right), \ (1)$$

where Q, K, V, and  $\alpha$  are the query tensor, the key tensor, the value tensor, and the attention weight distribution tensor across the set of specified axes  $\mathcal{A}$  of the value/key tensors. For example, for an input query Q of shape  $X_Q \times T_Q \times D$  and value V of shape  $X_V \times T_V \times D$ ,  $\operatorname{Att}_{\{X_V, T_V\}}(Q, V, V)$  performs attention over the first and second axes of V, yielding an output tensor of shape  $X_Q \times T_Q \times D$ . The dense layers  $f_Q, f_V, f_K$  are trainable and applied to the depth dimension D.

We utilize the multi-head attention (MHA) layer [44]. Each one of the H heads performs attention over some low-dimensional embeddings derived from the tensors Q and V, with the output embedding depth reduced to D/H. These independent attention heads have the capability to focus on different semantics of the input tensors. After performing attention across the specified axes  $\mathcal{A}$ , the final output is given by aggregating across the head outputs  $o^{(h)}$ :

$$o^{(h)} = \text{Att}_{\mathcal{A}}(f_Q^{(h)}(Q), f_V^{(h)}(V), f_V^{(h)}(V)),$$
  
MHA\_{\mathcal{A}}(Q, V) = f(Concat(o\_1, \dots, o\_H)), (2)

where f denotes a dense layer  $\mathbb{R}^{X_Q \times T_Q \times D} \to \mathbb{R}^{X_Q \times T_Q \times D}$  and the dense layers  $f_Q^{(h)}$  and  $f_V^{(h)}$  are linear maps  $\mathbb{R}^{X_Q \times T_Q \times D} \to \mathbb{R}^{X_Q \times T_Q \times D/H}$  and  $\mathbb{R}^{X_V \times T_V \times D} \to \mathbb{R}^{X_V \times T_V \times D/H}$ , respectively for each one of the  $h \in 1, \ldots, H$  heads, where for our purposes we always assume that the keys and values tensors are the same size. Using these definitions, we now formulate our proposed attention methods.

**Self-attention (SA)** First, we concatenate the audio tensor  $A_{\text{SA}}^{(0)} = Z_{\text{A}}$  and the video  $V_{\text{SA}}^{(0)} = Z_{\text{V}}$  tensors across the first axis to form the  $(M + G) \times T \times D$  tensor  $Z_{\text{SA}}^{(0)}$ , the input to the first self-attention layer.

**Joint SA**: Attention is performed jointly across space, time, and sources (see Figure 3a). We express the l-th layer of a joint self-attention module as follows:

$$b^{(l)} = \mathrm{MHA}_{\{\mathrm{M+G,T}\}}(Z_{\mathrm{SA}}^{(l-1)}, Z_{\mathrm{SA}}^{(l-1)}) + Z_{\mathrm{SA}}^{(l-1)},$$
  

$$Z_{\mathrm{SA}}^{(l)} = \mathrm{LN}(f^{(l)}(\mathrm{Dropout}(b^{(l)})) + b^{(l)}),$$
(3)

where  $f^{(l)}$  is a dense layer, LN is layer normalization [4] and Dropout denotes a dropout layer [37]. We define the sequence of operations in (3) as  $Z_{\rm SA}^{(l)} =$  ${\rm SA}_{\{M+G,T\}}(Z_{\rm SA}^{(l-1)})$  where the self-attention is performed across the joint sourcesand-spatial dimension M + G and the time axis T. The final representation z, after the repetition of L self-attention blocks, for all M sources, is obtained by slicing and performing attentional pooling [43] across the time axis of  $\hat{z}$ :

$$z = \mathrm{MHA}_{\{\mathrm{T}\}}(\sum_{t}^{T} \hat{z}_{t}, \hat{z}) \in \mathbb{R}^{M \times D}, \ \hat{z} = A_{\mathrm{SA}}^{(L)} = Z_{\mathrm{SA}}^{(L)}[1:M] \in \mathbb{R}^{M \times T \times D}.$$
 (4)

**Separable SA**: The l-th separable self-attention block can be expressed using the SA module defined in (3):

$$a^{(l)} = \mathrm{SA}_{\{\mathrm{T}\}}(Z_{\mathrm{SA}}^{(l-1)}[1:M]), \ v^{(l)} = \mathrm{SA}_{\{\mathrm{T}\}}(Z_{\mathrm{SA}}^{(l-1)}[M:M+G]),$$
  
$$Z_{\mathrm{SA}}^{(l)} = \mathrm{SA}_{\{\mathrm{M}+\mathrm{G}\}}\left(\mathrm{Concat}(a^{(l)},v^{(l)})\right).$$
(5)

The final audio-visual representation is obtained through attentional pooling and slicing as before (see also Figure 3c).

**Cross-modal attention (CMA)** In this attention layer we keep the audio and the video modality tensors separate, and we perform queries from one modality to another. Formally, the input to the stacked CMA blocks is a pair of an audio  $A_{\text{CMA}}^{(0)} \in \mathbb{R}^{M \times T \times D}$  and a video  $V_{\text{CMA}}^{(0)} \in \mathbb{R}^{G \times T \times D}$  feature tensors.

**Joint CMA**: We perform a directional attention from the audio (video) modality tensor to the video (audio) tensor, attending across both sources and time (space and time) axes. Formally, at the *l*-th layer we have the following sequence of operations for the directional attention where we use as a query the audio modality, also illustrated in Figure 3b:

$$a_{1}^{(l)} = \mathrm{MHA}_{\{\mathrm{G},\mathrm{T}\}}(A_{\mathrm{CMA}}^{(l-1)}, V_{\mathrm{CMA}}^{(l-1)}), \ a_{2}^{(l)} = \mathrm{LN}(a_{1}^{(l)} + A_{\mathrm{CMA}}^{(l-1)}), A_{\mathrm{CMA}}^{(l)} = \mathrm{LN}(f(\mathrm{Dropout}(a_{2}^{(l)})) + A_{\mathrm{CMA}}^{(l-1)}).$$
(6)

For the other direction, we modulate the video features  $v^{(l)}$  using the audio features  $a^{(l)}$  by swapping  $A_{\text{CMA}}$  and  $V_{\text{CMA}}$  and attending over {M, T} in (6).

We define (6) as  $A_{\text{CMA}}^{(l)}, V_{\text{CMA}}^{(l)} = \text{CMA}_{\{M|G,T\}}(A_{\text{CMA}}^{(l-1)}, V_{\text{CMA}}^{(l-1)})$ , where each cross-modal attention is performed across the dimension of audio sources M or spatial locations G (denoted in our notation as "M|G") and the time axis T. The output audio-visual embedding z contains information for all M sources and is obtained after the repetition of L cross-modal attention blocks via attentional pooling across time (4) on the output audio tensor  $\hat{z} = A_{\text{CMA}}^{(l)} \in \mathbb{R}^{M \times T \times D}$ . Separable CMA: Similar to Section 3.3, we can reduce the space complexity

**Separable CMA**: Similar to Section 3.3, we can reduce the space complexity of the proposed CMA layer by first performing self-attention across the time axis for each modality separately, then performing CMA across the remaining axis (i.e. sources or spatial locations) as shown next, also illustrated in Figure 3c:

$$a^{(l)} = \mathrm{SA}_{\{\mathrm{T}\}}(A^{(l-1)}_{\mathrm{CMA}}), \ v^{(l)} = \mathrm{SA}_{\{\mathrm{T}\}}(V^{(l-1)}_{\mathrm{CMA}}), A^{(l)}_{\mathrm{CMA}}, V^{(l)}_{\mathrm{CMA}} = \mathrm{CMA}_{\{\mathrm{M}|\mathrm{G}\}}(a^{(l)}, v^{(l)}).$$
(7)

#### 3.4 Audio-visual on-screen sound classifier

For each estimated source  $\hat{s}_m$ , AudioScopeV2 predicts the probability  $\hat{y}_m$  that it originates from an on-screen object. These probabilities are computed using the extracted audio-visual representation from the output z of our attention-based models for the self-attention and cross-modal attention encoders. Specifically, for each source m, we feed the audio-visual embedding  $z_m \in \mathbb{R}^D$  through a dense layer  $f_z$  tied across sources to produce logits  $\hat{\ell}_m$ , and then apply an element-wise sigmoid activation  $\sigma$  to compute the audio-visual coincidence probability  $\hat{y}_m$ . The final on-screen waveform estimate  $\hat{x}^{\text{on}}$  is produced using these probabilities as soft weights and multiplied with the corresponding estimated sources:

$$\hat{\ell}_m = f_z(z_m), \ \hat{y}_m = \sigma(\hat{\ell}_m) \in [0,1], \ \hat{x}^{\text{on}} = \sum_{m=1}^M \hat{y}_m \hat{s}_m.$$
 (8)

#### 3.5 Training procedure

To train the separation model, we use MixIT [47]. Given two reference mixtures  $r_1, r_2 \in \mathbb{R}^{T'}$ , M separated sources  $\hat{s}$  predicted from the MoM  $x = r_1 + r_2$ , and a signal-level training loss  $\mathcal{L}$ , MixIT infers an optimal  $2 \times M$  binary mixing matrix A where each column sums to one. This mixing matrix assigns each estimated source  $s_m$  to one of the reference mixtures  $r_1$  or  $r_2$ :

$$\mathcal{L}_{\text{MixIT}}(r, \hat{s}) = \min_{A \in \mathbb{B}^{2 \times M}} \sum_{n=1}^{2} \mathcal{L}(r_n, [A\hat{s}]_n).$$
(9)

For  $\mathcal{L}$ , we use the negative thresholded SNR loss [43, 47].

We use purely unsupervised training, which uses batches composed of *noisy-labeled on-screen* (NOn) examples. Each NOn example consists of the video frames and audio for a primary input 5-second video clip, where additional audio from another random 5-second video clip serves as synthetic off-screen

audio that is mixed with the primary video soundtrack. These examples provide noisy pseudo-labels, because even after optimal MixIT combination, the primary soundtrack may contain off-screen background noise.

For training AudioScopeV2's audio-visual on-screen classifier with NOn examples, we use the *active combinations* loss  $\mathcal{L}_{AC}$  [43] computed between the pseudo-label assignments y provided by MixIT and the predictions  $\hat{y}$  of the classifier.  $\mathcal{L}_{AC}$  corresponds to the minimum cross-entropy loss  $\mathcal{L}_{CE}$  over all settings  $\mathscr{P}_{\geq 1}(\mathbb{B}^M)$  of the labels such that at least one label is 1 (equivalent to at least one source appearing on-screen):

$$\mathcal{L}_{\mathrm{AC}}(y,\hat{y}) = \min_{\ell \in \{\mathcal{O}_{\geq 1}(\mathbb{B}^M)} \sum_m \mathcal{L}_{\mathrm{CE}}(\ell_m, \hat{y}_m).$$
(10)

## 4 Experimental Framework

#### 4.1 Data preparation and labeling

For our open-domain experiments we use Flickr Creative Commons 100 Million Dataset (YFCC100M) [38] CC-BY videos respecting published train/validation/test splits [18]. Instead of using only a filtered subset of the data, as previous state-of-the-art methods suggested [43], our proposed training recipe is able to leverage unrestricted open-domain datasets by using better pre-trained separation models.

The unfiltered training data consists of  $\approx 1600$  hours, and we extract 5second clips with a hop of 1 second (4.85M total clips). We also gathered human annotations for 5-second clips from unfiltered videos from the train, validation, and test splits. The count of total clips rated, unanimously-rated on-screen-only clips, and unanimously rated off-screen clips were 20000/480/4664 for training, 6500/109/1421 for validation, and 3500/43/762 for test. Notice that during training we dynamically create mixtures of pairs of these clips, so the effective number of unique examples is  $\mathcal{O}(10^{13})$ . We also experiment with faster video frame rate of 16 frames per second (FPS), instead of 1 FPS used by AudioScope [43].

#### 4.2 Training details

Both audio and visual embedding networks were pre-trained on AudioSet [15] for unsupervised coincidence prediction [22]. We also found that freezing these networks during training leads to better results. Also, instead of training the separation model from scratch, we pre-train the separation model with MixIT on unfiltered audio-only MoMs drawn from YFCC100M for 3.6M steps, which also significantly boosted the performance of our models. We use L = 4 stacked proposed layers of joint/separable SA/CMA (ablation studies can be found in the supplementary material). All models were trained on 32 Google Cloud TPU v3 cores with the Adam optimizer [25], batch size 128, and learning rate  $10^{-4}$ .

11

#### 4.3 Evaluation datasets

In order to compare with the current state-of-the-art, we use the AudioScope dataset splits [43] provided online [17], which were drawn from a subset of YFCC100M filtered by an unsupervised coincidence model [22]. These datasets (we refer to them as *filtered off-screen background*) contain two kinds of examples: on-screen MoMs, where additional off-screen audio was injected into the soundtrack of a unanimously-rated on-screen-only video (i.e. input audio  $x = x^{\text{on}} + x^{\text{off}}$ ), and off-screen MoMs, where additional audio was mixed into the soundtrack of a unanimously-rated off-screen-only video (input audio  $x = x^{\text{off}}$ ).

Similarly, we construct new evaluation sets of MoMs using videos from unfiltered validation and test splits of YFCC100M, which we call *unfiltered random background*. Instead of using off-screen-only audio, our dataset uses randomly sampled audio from all unfiltered videos as synthetic off-screen background for unanimously-rated on-screen-only or off-screen-only videos (see Section 4.1).

#### 4.4 Evaluation metrics and calibration

For on-screen examples with input audio  $x = x^{\text{on}} + x^{\text{off}}$ , we report the reconstruction fidelity of on-screen estimates  $\hat{x}^{\text{on}}$  (8) using *signal-to-noise ratio* (SNR) in dB. For off-screen examples with input audio  $x = x^{\text{off}}$  where we know that no audio originates from on-screen objects, we measure the ability of our models to suppress off-screen sources using off-screen suppression ratio (OSR).

$$SNR(x^{\text{on}}, \hat{x}^{\text{on}}) = 20 \log_{10} \frac{\|x^{\text{on}}\|}{\|x^{\text{on}} - \hat{x}^{\text{on}}\|}, OSR(x, \hat{x}^{\text{on}}) = 20 \log_{10} \frac{\|x\|}{\|\hat{x}^{\text{on}}\|}.$$
 (11)

OSR measures the power reduction of the on-screen estimate  $\hat{x}^{\text{on}}$  relative to the input audio, and is only measured on examples where the input audio  $x = x^{\text{off}}$  is entirely off-screen. Prior work [43] has also used scale-invariant signal-to-noise ratio (SI-SNR) [27] instead of SNR; however, this allows a model to obtain large OSR, without sacrificing SI-SNR, by scaling down its estimates.

Both OSR and SNR are important, but there is an inherent trade-off between them, since OSR can always be increased by scaling down the output, at the expense of SNR. This makes it difficult to compare models that have different operating points in this trade-off, a problem not addressed by [43]. To illustrate this, Figure 4 plots OSR versus SNR for the models we consider in this paper. Notice that each model achieves a different operating point. Because of these differences, SNR cannot be meaningfully compared across models without considering OSR. For example, one of our proposed models achieves a lower SNR of about 5.3 dB at 10.8 dB OSR, compared to an AudioScope<sup>\*</sup> model that achieves a higher SNR of 5.5 dB at 8.9 dB OSR. It is difficult to say which model is better.

To solve this problem, we propose a novel calibration method for these onscreen separation models, where we adjust a bias in the classifiers to achieve a target average OSR. This is akin to choosing a threshold for a detector to achieve a target false positive rate. For our procedure, we define a calibrated on-screen



Fig. 4: OSR versus SNR curves when calibration offset  $\theta$  in (12) is varied, on unfiltered random background test set for all models in Table 1 (except for AudioScope [43], which achieves 23.4 dB OSR / 1.1 dB SNR for 1 FPS).

estimate given by adding a global scalar offset  $\theta$  to the on-screen logits  $\hat{\ell}_{1:M}$  (8):

$$\tilde{x}^{\rm on}(\theta) = \sum_m \hat{s}_m \sigma(\hat{\ell}_m + \theta). \tag{12}$$

The offset  $\theta$  is tuned such that the median OSR (11) across all  $N_{\text{off}}$  off-screen examples,  $\text{med}_{j=1}^{N_{\text{off}}} \text{OSR}[x_j, \tilde{x}_j^{\text{on}}(\theta)]$ , is equal to a desired  $\text{OSR}_{\text{target}}$ . The curves in Figure 4 illustrate the effect on OSR and SNR when  $\theta$  is varied, and are akin to receiver operating characteristic (ROC) or precision-recall curves.

In practice,  $\theta$  has a monotonic relationship to OSR. As  $\theta$  tends towards 0 (inversely  $\infty$ ), the on-screen probabilities  $\hat{y}_m$  tend towards 0 (inversely 1), and thus OSR approaches  $\infty$  (inversely 0) dB. Because of this property, optimization of  $\theta$  is very simple, and can be accomplished efficiently via binary search.

For our results, we use  $OSR_{target} = 6dB$ , which corresponds to all off-screen sources sounding as if they are twice as far away. To choose the early stopping point, we evaluate all models on the *unfiltered random background* validation data. Running calibration on each step of model training would be prohibitively expensive. Thus, we choose the point which maximizes the minimum of SNR and OSR, favoring models which are not strongly biased towards any metric.

We also report the performance of classification using weighted area under the curve of the ROC (AUC-ROC), where the weight for each probability  $\hat{y}_m$  is the normalized power of the corresponding source  $\hat{s}_m$ :  $\|\hat{s}_m\|_2^2 / (\sum_{m'} \|\hat{s}_{m'}\|_2^2)$ .

#### 5 Results

#### 5.1 Open-domain on-screen separation

Results are shown in Table 1 for the filtered off-screen background dataset from [43], and our new and challenging unfiltered random background dataset. We include two "no processing" baselines, with the on-screen estimate equal to the input audio  $\hat{x}^{\text{on}} = x$  (0 dB OSR), or half the input  $\hat{x}^{\text{on}} = x/2$  (6 dB OSR).

Table 1: Evaluation results for filtered off-screen background (from [43]) and unfiltered random background (our new proposed) test sets at 1 and 16 FPS. Separate models are trained for 1 FPS and 16 FPS. For each model, calibration to 6 dB OSR is performed separately on the filtered and unfiltered test sets. "PT" indicates separation model pre-training, "Filt." means training on filtered data, and "Complexity" is the theoretical complexity of each AV-alignment module.

				Filtered [43]				Unfiltered (new proposed)			
				$1 \; \mathrm{FPS}$		$16 \ \mathrm{FPS}$		$1 \ \mathrm{FPS}$		$16 \ \mathrm{FPS}$	
AV alignment		Complexity	PT Filt.	SNR	AUC	SNR	AUC	SNR	AUC	SNR	AUC
No processing $(\hat{x}^{on} = x \text{ with 0dB OSR})$			R)	4.4	_	4.4	_	2.5	_	2.5	_
No processing $(\hat{x}^{o})$		$x^n = x/2$ with 6dB 0	OSR)	4.7	-	4.7	-	4.1	_	4.1	-
AudioScope [43]		$\mathcal{O}(TMG)$	$\checkmark$	6.0	0.79	_	-	2.7	0.69	-	-
AudioScope*		$\mathcal{O}(TMG)$	$\checkmark$	8.2	0.80	5.9	0.77	5.8	0.78	5.2	0.71
SA	Joint	$\mathcal{O}(T^2[M+G]^2)$	$\checkmark$	10.0	0.84	9.9	0.86	7.2	0.82	7.7	0.83
	Sep.	$\mathcal{O}(T^2 + [M+G]^2)$	) 🗸	9.6	0.84	8.2	0.83	6.6	0.78	6.6	0.80
CMA	Joint	$\mathcal{O}(T^2 M G)$	$\checkmark$	10.0	0.88	10.0	0.85	7.3	0.83	7.7	0.84
	Sep.	$\mathcal{O}(T^2 + MG)$	$\checkmark$	9.5	0.83	9.3	0.82	6.4	0.78	7.1	0.80

On the filtered test set, our proposed models significantly outperform the previous state-of-the-art AudioScope model [43] trained on filtered data, by more than 4dB in terms of SNR and 0.09 in AUC-ROC. Training on unfiltered data, including pre-training the separation model on this data (AudioScope\*), improves over the baseline AudioScope model trained on filtered data (6.0 dB  $\rightarrow$  8.2 dB).

On our newly introduced unfiltered random background dataset, AudioScope [43] trained on mismatched filtered data suffers from poor generalization, achieving only 2.7 dB SNR, which is worse than doing no processing (4.1 dB SNR). This demonstrates the limitation of training on filtered data. For our proposed models, joint CMA yields the best improvements over AudioScope\* trained with matched unfiltered data (5.8 dB  $\rightarrow$  7.7 dB SNR and 0.78  $\rightarrow$  0.84 AUC-ROC). Also, notice that the much more efficient separable versions of SA and CMA only suffer minor degradation compared to the joint versions (7.2 dB  $\rightarrow$  6.6 dB for SA, 7.3 dB  $\rightarrow$  6.4 dB for CMA). Our proposed models can perceive and leverage higher-frequency dynamics from the audio-visual scene and thus, using a higher frame rate of 16 FPS provides improvement of up to 0.7 dB SNR. On the other hand, AudioScope<sup>\*</sup> seemingly cannot scale to higher frame-rates (8.2 dB  $\rightarrow$  5.9 dB for filtered and 5.8 dB  $\rightarrow$  5.2 dB for unfiltered, going from 1 FPS to 16 FPS). presumably because its audio-visual alignment is limited by only using shallow spatio-temporal attention. It is also possible that filtering with the audio-visual coincidence model [22], which averaged coincidence scores over frames at 1 FPS, may have biased towards video clips with always-visible sounding objects.

#### 5.2 Computational efficiency

In theory, joint SA has complexity  $\mathcal{O}(T^2[M+G]^2)$ , which scales poorly relative its separable version with complexity  $\mathcal{O}(T^2+[M+G]^2)$ . Joint CMA has a slightly

better complexity  $\mathcal{O}(T^2MG)$  but still lags behind its separable counterpart  $\mathcal{O}(T^2 + MG)$ . However, it is important to evaluate how the computation scales in practice. The execution time versus input video length for all models is depicted in Figure 5. The efficiency of the proposed separable self- and cross-modal attention becomes apparent for longer videos. These separable variants perform comparably to the much more complex joint CMA and joint SA models (see Table 1) but have computational requirements on par with the simpler AudioScope [43].



Fig. 5: Computation requirements AudioScope [43] and our proposed models with separable and joint SA and CMA for different video input lengths at 16 FPS. All measurements were taken on a machine with 16 GB of RAM and 2 Intel Xeon CPU @ 2.20GHz cores. Input duration is increased until out-of-memory error.

# 6 Conclusion

We identified several issues with the previous state-of-the-art open-domain audiovisual on-screen sound separation model, AudioScope [43]. These issues include oversimplicity of shallow attention used for audio-visual alignment, lack of generalization due to filtering video data with an unsupervised audio-visual coincidence model [22], and inability to specify the trade off between on-screen reconstruction and off-screen suppression. We proposed more sophisticated self- and cross-modal attention architectures, along with more efficient separable versions, that improve performance and are able to leverage higher video frame rates. To address lack of generalization, we provide annotations for a new dataset constructed from YFCC100M [38] that is unfiltered by the audio-visual coincidence model. As a result, this dataset is more diverse and representative of in-the-wild video data. Finally, we proposed a calibration procedure that allows any on-screen separation model to be tuned for a specific level of off-screen suppression, which allows more control over model behavior and much easier comparison between different models. Using our calibration procedure, our results show that our proposed architecture is able to generalize to our more challenging test set and achieve clear improvements over the previous state-of-the-art AudioScope model [43].

15

## References

- 1. Afouras, T., Chung, J.S., Zisserman, A.: The Conversation: Deep audio-visual speech enhancement. Proc. Interspeech pp. 3244–3248 (2018)
- Afouras, T., Owens, A., Chung, J.S., Zisserman, A.: Self-supervised learning of audio-visual objects from video. In: Proc. European Conference on Computer Vision (ECCV) (2020)
- Arandjelovic, R., Zisserman, A.: Objects that sound. In: Proc. European Conference on Computer Vision (ECCV). pp. 435–451 (2018)
- 4. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
- Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. In: Proc. International Conference on Learning Representations (ICLR) (2015)
- Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: Proc. International Conference on Machine Learning (ICML). vol. 2, p. 4 (2021)
- Chatterjee, M., Le Roux, J., Ahuja, N., Cherian, A.: Visual scene graphs for audio source separation. In: Proc. IEEE International Conference on Computer Vision (ICCV). pp. 1204–1213 (2021)
- Cheng, Y., Wang, R., Pan, Z., Feng, R., Zhang, Y.: Look, listen, and attend: Co-attention network for self-supervised audio-visual representation learning. In: ACM MM. pp. 3884–3892 (2020)
- Drude, L., Hasenklever, D., Haeb-Umbach, R.: Unsupervised training of a deep clustering model for multichannel blind source separation. In: Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP). pp. 695–699 (2019)
- Ephrat, A., Mosseri, I., Lang, O., Dekel, T., Wilson, K., Hassidim, A., Freeman, W.T., Rubinstein, M.: Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation. ACM TOG 37(4), 1–11 (2018)
- Gan, C., Huang, D., Zhao, H., Tenenbaum, J.B., Torralba, A.: Music gesture for visual sound separation. In: Proc. IEEE International Conference on Computer Vision (CVPR). pp. 10478–10487 (2020)
- Gao, R., Feris, R., Grauman, K.: Learning to separate object sounds by watching unlabeled video. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 35–53 (2018)
- Gao, R., Grauman, K.: Co-separating sounds of visual objects. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3879–3888 (2019)
- Gao, R., Grauman, K.: VisualVoice: Audio-visual speech separation with crossmodal consistency. In: Proc. IEEE International Conference on Computer Vision (CVPR). pp. 15490–15500 (2021)
- Gemmeke, J.F., Ellis, D.P.W., 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. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP). pp. 776–780 (2017)
- Gfeller, B., Roblek, D., Tagliasacchi, M.: One-shot conditional audio filtering of arbitrary sounds. In: Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP). pp. 501–505 (2021)

- 16 E. Tzinis et al.
- Google Research Sound Separation Team: AudioScope YFCC100M clip lists (2021), https://github.com/google-research/sound-separation/tree/ master/datasets/audioscope
- Google Research Sound Separation Team: YFCC100M clip lists (2021), https://github.com/google-research/sound-separation/tree/master/ datasets/yfcc100m
- Hershey, J., Movellan, J.: Audio vision: Using audio-visual synchrony to locate sounds. Advances in neural information processing systems 12, 813–819 (1999)
- Hershey, J.R., Casey, M.: Audio-visual sound separation via hidden markov models. In: Advances in Neural Information Processing Systems. pp. 1173–1180 (2002)
- Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)
- 22. Jansen, A., Ellis, D.P., Hershey, S., Moore, R.C., Plakal, M., Popat, A.C., Saurous, R.A.: Coincidence, categorization, and consolidation: Learning to recognize sounds with minimal supervision. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 121–125 (2020)
- Kavalerov, I., Wisdom, S., Erdogan, H., Patton, B., Wilson, K., Le Roux, J., Hershey, J.R.: Universal sound separation. In: Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA). pp. 175–179 (2019)
- 24. Kilgour, K., Gfeller, B., Huang, Q., Jansen, A., Wisdom, S., Tagliasacchi, M.: Text-driven separation of arbitrary sounds. In: Proc. Interspeech (2022)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proc. International Conference on Learning Representations (ICLR) (2015)
- Korbar, B., Tran, D., Torresani, L.: Cooperative learning of audio and video models from self-supervised synchronization. In: Advances in Neural Information Processing Systems. pp. 7763–7774 (2018)
- Le Roux, J., Wisdom, S., Erdogan, H., R. Hershey, J.: SDR-half-baked or well done? In: Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP). pp. 626–630 (2019)
- Lee, J., Chung, S.W., Kim, S., Kang, H.G., Sohn, K.: Looking into your speech: Learning cross-modal affinity for audio-visual speech separation. In: Proc. IEEE International Conference on Computer Vision (CVPR). pp. 1336–1345 (2021)
- Lin, Y.B., Tseng, H.Y., Lee, H.Y., Lin, Y.Y., Yang, M.H.: Unsupervised sound localization via iterative contrastive learning. arXiv preprint arXiv:2104.00315 (2021)
- Liu, X., Liu, H., Kong, Q., Mei, X., Zhao, J., Huang, Q., Plumbley, M.D., Wang, W.: Separate what you describe: Language-queried audio source separation. In: Proc. Interspeech (2022)
- Ochiai, T., Delcroix, M., Koizumi, Y., Ito, H., Kinoshita, K., Araki, S.: Listen to What You Want: Neural Network-Based Universal Sound Selector. In: Proc. Interspeech. pp. 1441–1445 (2020)
- Owens, A., Efros, A.A.: Audio-visual scene analysis with self-supervised multisensory features. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 631–648 (2018)
- Pishdadian, F., Wichern, G., Le Roux, J.: Finding strength in weakness: Learning to separate sounds with weak supervision. IEEE/ACM Transactions on Audio, Speech, and Language Processing 28, 2386–2399 (2020)
- Rahman, T., Sigal, L.: Weakly-supervised audio-visual sound source detection and separation. In: Proc. IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6 (2021)

- Rouditchenko, A., Zhao, H., Gan, C., McDermott, J., Torralba, A.: Self-supervised audio-visual co-segmentation. In: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 2357–2361 (2019)
- 36. Seetharaman, P., Wichern, G., Le Roux, J., Pardo, B.: Bootstrapping singlechannel source separation via unsupervised spatial clustering on stereo mixtures. In: Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP). pp. 356–360 (2019)
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. The journal of machine learning research 15(1), 1929–1958 (2014)
- Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., Li, L.J.: Yfcc100m: The new data in multimedia research. Communications of the ACM 59(2), 64–73 (2016)
- Tian, Y., Li, D., Xu, C.: Unified multisensory perception: Weakly-supervised audiovisual video parsing. In: Proc. European Conference on Computer Vision (ECCV). pp. 436–454 (2020)
- 40. Tzinis, E., Venkataramani, S., Smaragdis, P.: Unsupervised deep clustering for source separation: Direct learning from mixtures using spatial information. In: Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP). pp. 81–85 (2019)
- Tzinis, E., Venkataramani, S., Wang, Z., Subakan, C., Smaragdis, P.: Two-step sound source separation: Training on learned latent targets. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 31–35 (2020)
- Tzinis, E., Wisdom, S., Hershey, J.R., Jansen, A., Ellis, D.P.W.: Improving universal sound separation using sound classification. In: Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP). pp. 96–100 (2020)
- Tzinis, E., Wisdom, S., Jansen, A., Hershey, S., Remez, T., Ellis, D.P., Hershey, J.R.: Into the wild with audioscope: Unsupervised audio-visual separation of on-screen sounds. In: Proc. International Conference on Learning Representations (ICLR) (2021)
- 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**, 5998–6008 (2017)
- Wang, Q., Muckenhirn, H., Wilson, K., Sridhar, P., Wu, Z., Hershey, J., Saurous, R.A., Weiss, R.J., Jia, Y., Moreno, I.L.: Voicefilter: Targeted voice separation by speaker-conditioned spectrogram masking. In: Proc. Interspeech. pp. 2728–2732 (2019)
- 46. Wisdom, S., R. Hershey, J., Wilson, K., Thorpe, J., Chinen, M., Patton, B., A. Saurous, R.: Differentiable consistency constraints for improved deep speech enhancement. In: Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP). pp. 900–904 (2019)
- 47. Wisdom, S., Tzinis, E., Erdogan, H., Weiss, R.J., Wilson, K., Hershey, J.R.: Unsupervised sound separation using mixtures of mixtures. In: Advances in Neural Information Processing Systems. vol. 33, pp. 3846–3857 (2020)
- Wu, Y., Zhu, L., Yan, Y., Yang, Y.: Dual attention matching for audio-visual event localization. In: Proc. IEEE International Conference on Computer Vision (CVPR). pp. 6292–6300 (2019)
- Yu, J., Cheng, Y., Feng, R.: Mpn: Multimodal parallel network for audio-visual event localization. In: Proc. IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6 (2021)

- 18 E. Tzinis et al.
- Zhang, Y., Li, X., Liu, C., Shuai, B., Zhu, Y., Brattoli, B., Chen, H., Marsic, I., Tighe, J.: Vidtr: Video transformer without convolutions. In: Proc. IEEE International Conference on Computer Vision (ICCV). pp. 13577–13587 (2021)
- Zhu, L., Rahtu, E.: Visually guided sound source separation and localization using self-supervised motion representations. In: Proc. IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 1289–1299 (2022)