

Annotation-free Audio-Visual Segmentation

Jinxiang Liu¹, Yu Wang^{1,2}, Chen Ju¹, Chaofan Ma¹, Ya Zhang^{1,2}✉, Weidi Xie^{1,2}✉

¹ Cooperative Medianet Innovation Center, Shanghai Jiao Tong University ² Shanghai AI Laboratory
{jinxliu, yuwangsJTU, ju_chen, chaofanma, ya_zhang, weidi}@sjtu.edu.cn

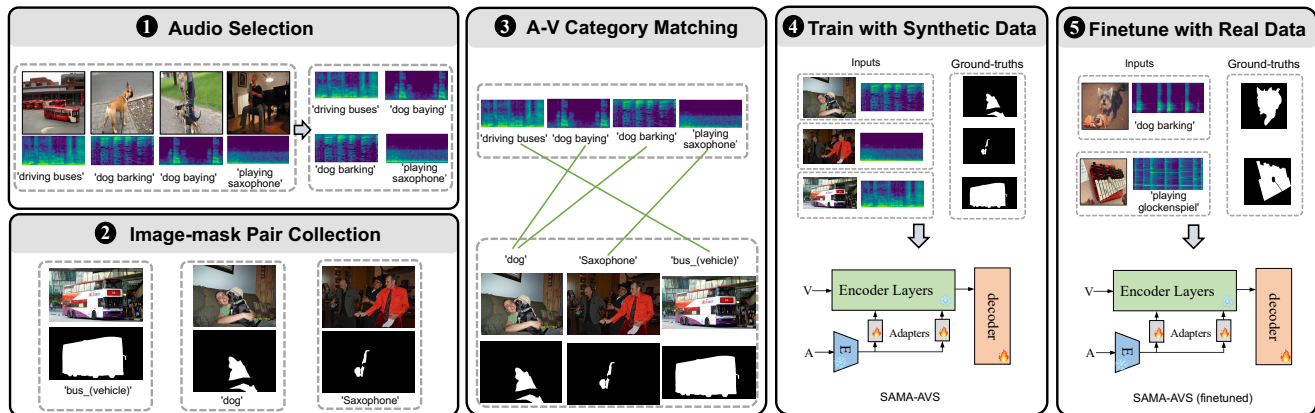


Figure 1. We propose an ingenious pipeline to synthesize datasets for audio-visual segmentation (AVS), by leveraging *off-the-shelf* large image segmentation datasets and audio corpus. The pipeline is *annotation-free* and can easily be *scaled up* to cover more categories. Besides, we develop an effective method SAMA-AVS to adapt *segmentation anything* (SAM) with Adapters to the AVS task. By introducing only a small number of model parameters, the model performance is substantially enhanced and achieves state-of-the-art performance.

Abstract

The objective of Audio-Visual Segmentation (AVS) is to localise the sounding objects within visual scenes by accurately predicting pixel-wise segmentation masks. To tackle the task, it involves a comprehensive consideration of both the data and model aspects. In this paper, first, we initiate a novel pipeline for generating artificial data for the AVS task without extra manual annotations. We leverage existing image segmentation and audio datasets and match the image-mask pairs with its corresponding audio samples using category labels in segmentation datasets, that allows us to effortlessly compose (image, audio, mask) triplets for training AVS models. The pipeline is annotation-free and scalable to cover a large number of categories. Additionally, we introduce a lightweight model SAMA-AVS which adapts the pre-trained segment anything model (SAM) to the AVS task. By introducing only a small number of trainable parameters with adapters, the proposed model can effectively achieve adequate audio-visual fusion and interaction in the encoding stage with vast majority of parameters fixed. We conduct extensive experiments, and the results show our proposed model remarkably surpasses other com-

peting methods. Moreover, by using the proposed model pretrained with our synthetic data, the performance on real AVSBench data is further improved, achieving 83.17 mIoU on S4 subset and 66.95 mIoU on MS3 set. The project page is <https://jinxiang-liu.github.io/anno-free-AVS/>.

1. Introduction

Audio-visual segmentation (AVS) task involves segmenting the sounding objects corresponding to the audio cues in the video frames. Previous research has explored this task under the umbrella of self-supervised learning by leveraging the occurrence of audio-visual signals [3, 19–21, 32, 37, 50, 51, 54]. However, these training objectives lack fine-grained pixel-wise supervision and results in coarse segmentations. This poses limitation to the applicability of AVS in real-world scenarios where accurate segmentations are required, such as video surveillance, multi-modal video editing and the robotics industry. Recently, Zhou *et al.* [70] tackled the AVS problem through supervised training, and they manually annotated a video dataset with pixel-level segmentations for sounding objects in the video frames. However, due to the expensive annotation costs, the dataset

has a deficiency in the limited sample quantity and diversity.

To mitigate this, we propose an *annotation-free* and *scalable* pipeline to construct artificial datasets for the AVS task. Strictly speaking, an ideal model for the task requires *dual-level* recognition ability: semantic level and instance level. For example, given the video frames with the sound of dog barking, the model should first distinguish the dogs from the cats with the dog barking sound (*semantic level*); and then further resolve the instance ambiguity by identifying the specific dog that barks when there are multiple dogs present (*instance level*). Yet in practice, the ability to distinguish semantic categories may have been sufficient for localising sounding objects in most daily scenarios, which can be achieved by training on artificially constructed data with semantic-identical image and audio pairs. Based on such observation, we present a practical pipeline to artificially synthesize AVS datasets by leveraging off-the-shelf image segmentation and audio datasets without additional manual annotations, resolving the issue of insufficient training data for the AVS task.

In Figure 1, we depict the proposed pipeline to construct a synthetic dataset by leveraging off-the-shelf image segmentation and audio datasets. For instance, we can combine an image-mask pair labeled as “dog” from the image segmentation dataset LVIS with an audio sample labeled as “dog barking” from the audio dataset VGGSound to obtain an (*image, mask, audio*) triplet, which serves as a training sample for the AVS task. In this triplet, *image* and *audio* are the model’s inputs, and *mask* provides supervision for model training. Comparing to the existing human-annotated dataset AVSBench [70], this dataset synthesis pipeline inherits existing free annotations from computer vision communities, such as LVIS [14], Open Images [28]; and most importantly, it requires zero extra annotation, thus can easily be scaled up to cover a large number of categories.

In addition, the foundation model *segment anything* [27] (SAM) for image segmentation has attracted great attention in computer vision recently. SAM is pretrained with large-scale data and has demonstrated strong generalisation ability. In this paper, we first evaluate two variants of SAM on the AVS task. One is directly applying SAM to the task, the other is the proposed Audio-Prompted SAM (AP-SAM) which encodes the audios as the prompts for the mapping inputs of the SAM decoder. However, both models performs suboptimally. We analyze that vanilla SAM does not associate the semantics between audios and images in the pretraining stage; and even incorporating the audio prompts in the later stage of decoder part (AP-SAM), the fusion of audio-visual features is too *shallow* to facilitate the AVS task due to the light-weight design of the SAM decoder. Given such finding, a natural question arise: *how can we exploit the powerful ability of SAM, and effectively adapt it to the AVS task?* In this paper, we propose SAMA-

AVS to employ the adapters to transfer the powerful knowledge from SAM for AVS task. In a nutshell, we utilise the adapters to perform audio-visual fusion in the early stage and only train the adapters with much fewer parameters in the encoding stage. Thereby, the knowledge of SAM is employed and the audios features can be injected and fused with visual features efficiently and effectively in the early stage while only introducing a small number of parameters (<0.4% of the entire model). And experimental results show SAMA-AVS outperforms other SAM-based methods by a large margin and achieves state-of-the-art performance on the AVSBench dataset, indicating the effectiveness of the proposed method.

To summarize, the contributions are as follows:

- We propose a novel annotation-free pipeline to synthesize datasets for the AVS task, by harnessing the existing image segmentation datasets and audio corpus. The pipeline circumvents the cost of manual annotation and thus easily scales up to cover a greater number of categories for the task.
- As an instantiation of the pipeline, we present the AVS-Synthetic dataset, which is significantly larger and more diverse than existing AVS datasets. We conduct experiments with the synthetic dataset and demonstrate its strong zero-shot and few-shot transfer ability to real-world scenarios.
- We devise a model termed SAMA-AVS for AVS task based on SAM foundation model. To realise the deep fusion of audio-visual representations, the method deploys simple adapters to inject the audio information into the pretrained SAM, and allows the deep audio-visual fusion in the encoding stage with the vast majority of parameters fixed. Extensive experiments demonstrate the effectiveness of our proposed method.

2. Related Work

Audio-visual segmentation. In the recent literature, researchers have explored a range of audio-visual tasks with the goal of getting a comprehensive understanding of multimedia resources. These tasks include audio-visual sound separation [10, 11, 60, 67, 68], visual sound source localisation [3, 19–21, 32, 37, 47, 51, 54] and audio-visual video understanding [26, 29, 33, 59, 59]. Visual sound source localisation (VSSL), or audio-visual segmentation (AVS), aims to localise and segment the sounding object regions based on its audio signals. Existing approaches [3, 37, 46, 50, 51, 69] to this problem are primarily based on the co-occurrence of audio and visual signals, which provides only weak instance-level supervision, making it difficult to predict fine-grained pixel-level segmentation masks. To overcome this issue, Zhou *et al.* [70] presented a dataset with mask annotations

for sounding objects in the video frames, which is the first benchmark for the supervised AVS task.

Segment Anything. SAM [27] is released as the first foundation model for image segmentation in computer vision community. The model is pretrained with over 11 million images associating with over one billion segmentation masks. It also has the flexibility of supporting diverse prompts including points, boxes, masks and texts. With the advantages of SAM, it is soon utilised to attack various visual problems including medical image segmentation [8, 12, 42, 52, 61, 63], weakly-supervised semantic segmentation [5, 15, 24], few-shot segmentation [40, 65], 3D vision [2, 39], shadow detection [6, 62, 66], camouflaged object segmentation [6, 22, 23] and so on. Currently, few works except [45] explores SAM for the AVS task but the results from [45] are less satisfying. In this paper, we examine the performance of vanilla SAM and describe our proposed SAM-based models for the AVS task.

Adapter. Adapter [18] has become one of the most prevalent methods in parameter-efficient transfer learning. By introducing only a small number of parameters, it can effectively leverage the knowledge within the pretrained models and transfer to the related problems, thus it is now widely utilised for both NLP tasks [1, 18, 25, 36, 55, 56] and CV tasks [6, 7, 31, 35, 38, 49]. Recently, it is also employed to solve multi-modal problems [30, 34, 57].

3. Dataset

Here, we start by detailing the dataset construction procedure in Sec. 3.1; then we present the statistics of the proposed AVS-Synthetic dataset in Sec. 3.2.

3.1. Dataset Synthesis Pipeline

Existing datasets for audio-visual segmentation, *e.g.*, AVSBench [70], are all collected by manually annotating frames with pixel-wise segmentation masks, which is often expensive and only covers a limited number of categories (only 23). To remedy this, we initiate a scalable pipeline for constructing a synthetic dataset, termed as **AVS-Synthetic** dataset. The dataset construction procedure leverages existing image segmentation datasets such as LVIS [14], Open Images [28], and audio classification dataset, such as VGGSound [4]. Therefore it is free of extra manual annotations and can easily be scaled up.

Figure 1 illustrates our proposed pipeline for building a synthetic dataset for semantic-level audio-visual segmentation. Generally speaking, it consists of three steps:

(1) Audio Selection. We select audios using the instance-level class labels from publicly-released audio dataset. In this paper, we employ VGGSound [4], which is a large-scale audio-visual dataset collected using computer vision techniques, as a data source for providing audio sample

clips in our proposed AVS-Synthetic dataset. However, other audio datasets such as Audio Set [13] can also be used.

(2) Image-mask Pair Collection. Next, based on the categories of selected audio clips, we pick the categorized image and segmentation mask pairs for the objects in the target categories. This can be easily achieved by filtering the existing popular image segmentation datasets such as LVIS [14] and Open Images [28].

(3) Audio-Visual Category Matching. Till here, we have separately collected the audio and visual samples. We match each image-mask pair with their corresponding audio clip, thus forming the $(image, mask, audio)$ triplet as a training sample for AVS task. For each triplet, $(image, audio)$ serve as model inputs and the $mask$ is the ground-truth segmentation mask to supervise the model training. Matching these pairs is done by using the class labels of both image-mask pairs and the audio labels. Notably, different datasets may have different class labels while referring to the same thing; such as *computer keyboard*: “computer_keyboard” (LVIS) and “typing on computer keyboard” (VGGSound); *dog*: “dog” (LVIS), “dog barking” (VGGSound) and “dog baying” (VGGSound). Minimal manual verification is involved in this matching procedure.

3.2. AVS-Synthetic Statistics

As an instantiation of our *annotation-free* dataset collection pipeline for AVS, **AVS-Synthetic** dataset is proposed that covers 62,609 sounding object instances spanning 46 common categories. The training set size is 52,609, and the sizes of validation set and test set are both 5,000. A detailed comparison between AVS-Synthetic and the human annotated AVSBench datasets [70] is listed in Table 1.

Dataset	Classes	Images	Masks	Data.Source	Human
AVSBench [70]	23	4932	10,852	[4]	✓
AVS-Synthetic (ours)	46	58,405	62,609	[4, 14, 28]	✗

Table 1. Comparison between our proposed AVS-Synthetic and AVSBench (S4) [70] datasets. AVSBench is collected with human annotations.

4. Method

In Sec. 4.1, we start by introducing the problem formulation. Then we describe the two SAM-based models for AVS task in Sec. 4.2 and Sec. 4.3. Next, we detail the proposed SAMA-AVS architecture in Sec. 4.4. Finally, we describe the model training in Sec. 4.5.

4.1. Problem Formulation

For the task of audio-visual segmentation (AVS), the input data consists of a sequence of video frames, $\mathcal{V} = \{v_i\}_{i=1}^T$, where $v_i \in \mathbb{R}^{3 \times H_0 \times W_0}$, and its corresponding

audios $\mathcal{A} = \{a_i\}_{i=1}^T$, where $a_i \in \mathbb{R}^{H_a \times W_a}$ refers to the corresponding audio spectrogram. The objective of AVS is to segment the sounding objects in each frame based on the acoustic cues. The target segmentation can be binary masks, *i.e.*, $\mathcal{M} = \{M_i\}_{i=1}^T$, and $M_i \in \{0, 1\}^{H_0 \times W_0}$.

4.2. Revisiting SAM as Baseline

Segment Anything (SAM) [27] is proposed as a foundation model for scene segmentation, which is pretrained on 11M images with 1 billion segmentation masks. In terms of the architecture design, SAM consists of three components: an image encoder, a prompt encoder and a light-weight mask decoder. The image encoder employ the Vision Transformer [53] pretrained with MAE [16] to extract the image features with the shape of $C \times H \times W$ from the input images. For the prompt encoder, its supported prompt inputs can be divided into two categories: sparse prompts including points, boxes, texts, and masks. The points and boxes are processed by a combination of positional encodings and learned embeddings. The texts are the embeddings from the text encoder of CLIP [48]. The mask are summed into the image embeddings with convolution operations. Note that, when no prompts are given, SAM can also generate the segmentation proposals by traversing the entire image.

In this paper, we first investigate the performance of the vanilla SAM model on the AVS task, by picking the predicted mask that best matches the ground-truth [22, 58]. Concretely, for an image I with the ground-truth mask G , SAM model predicts N masks $\{M_i\}_{i=1}^N$. Then by calculating the intersection over union (IoU) of the ground-truth mask with the predictions, we select the predicted mask with the highest IoU value as the final prediction of SAM for evaluation. This method serves as the baseline for our proposed SAM-based methods.

4.3. Audio-Prompted SAM

Inspired by the way of dealing with text prompts of vanilla SAM as elaborated in Sec. 4.2, we treat the audio as another type of prompt by extracting the audio embeddings and set up another SAM-based model named Audio-Prompted SAM (AP-SAM) for the AVS task. Concretely, given an audio clip, we first compute its features with the VGGish [17] model pretrained on AudioSet [13] dataset, then spatially pool the features into a vector embedding. Similar to the SAM design, we pose the audio features as the prompt tokens. And thus the image embedding ($256 \times 64 \times 64$) from the image encoder and the audio prompt tokens (1×256) are passed to the light-weight mask decoder. The mask decoder updates the image embedding and audio prompt tokens, and uses the output tokens to perform mask prediction, following the same design of vanilla SAM. At training phase, we fix the parameters of both im-




Figure 2. The architecture of the proposed SAMA-AVS.

age encoder and audio encoder to preserve their knowledge, and only update the weights within the mask decoder.

4.4. SAMA-AVS: SAM with Adapters for AVS

In the AP-SAM model, although the image embedding and the audio embedding tokens are interacting with each other in the mask decoder, we conjecture the fusion is not adequate to drive the AVS task. For the mask decoder of SAM, it is light-weighted with only two transformer decoder layers. The reason why it works with such simple design is that the prompts including points, boxes and masks generally only provide explicit location information of target segmentation. And the whole models are trained with millions of images with semi-supervised learning. Therefore, a two-layer mask decoder design is competent enough for the vanilla SAM pretraining. However, to complete the audio-visual segmentation (AVS) task successfully, it requires the model to understand and associate the semantics between both image and audio embeddings; but this audio-image semantic understanding are not explored in the vanilla pretrained SAM. And the fusion and interaction of audio-visual tokens using the two-layer mask decoder of SAM may be too limited.

One simple solution is to inject the audio information into the image features at the encoding process. However, the computation cost of finetuning the whole encoder weights would be prohibitively high. To tackle the problem, we propose to perform deep audio-visual fusion with the assistance of adapters. We term the proposed framework based on SAM with adapters for AVS task as SAMA-AVS. We elaborate the method shown in Figure 2 in more detail.

Image tokenization. Our model accepts the visual frames and audio as the inputs. For visual input, given a video frame $I \in \mathbb{R}^{3 \times H \times W}$, we first upsample the image to the shape of $3 \times 1024 \times 1024$ to meet the requirements of SAM. Then we patchify the image into non-overlapping patches and flatten them into visual embeddings. Before feeding the features into the network, we add the visual embeddings with learnable position embeddings initialized with zeros. Finally we obtain the visual embeddings $X_v \in \mathbb{R}^{N \times D}$.

Audio tokenization. For the audio, we obtain the audio spectrogram $A \in \mathbb{R}^{H_a \times W_a}$ of one second audio clip corre-

sponding to the video frame. we compute its features with the pretrained VGGish [17] model, then spatially pool the features into a vector embedding $X_a \in \mathbb{R}^\ell$.

Deploying adapters at encoding stage. The adapters are responsible for injecting the audio information into the feature encoding process. Specifically, the audio tokens are transformed with learnable parameters and then added to the outputs of each transformer encoder layer. Formally, we denote the outputs of the $(i - 1)$ th layer of transformer encoder as \mathbf{Y}^{i-1} and the outputs of i th adapter as \mathbf{T}^i , then the inputs for the i th encoder layer \mathbf{X}^i is obtained by:

$$\mathbf{X}^i = \mathbf{Y}^{i-1} + \mathbf{T}^i. \quad (1)$$

After N layers of transformer encoder with adapters, the obtained deeply-fused audio-visual embeddings are fed to the last neck layer to adjust the channel dimension. Finally the embeddings from the neck layer are fed to the mask decoder to perform mask prediction.

In terms of adapter design, it is simply implemented with two-layer MLPs to adjust the feature dimension and perform feature extraction. During the training phase, we fix the parameters of the transform encoder layers and the audio encoder, and only update the parameters of the adapters and the light-weight mask decoder. Note that, the parameters introduced by the adapters constitute only **0.4%** of the entire model, while yielding significant performance gains, as indicated in the experiments.

4.5. Training Objectives

We use a combination of two commonly used objectives at training time: the BCE (Binary Cross-Entropy) loss and the IoU (Intersection over Union) loss. The total loss for optimizing the model is the weighted sum of the BCE loss and IoU loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{BCE}} + \lambda_{\text{IoU}} \mathcal{L}_{\text{IoU}}, \quad (2)$$

where λ_{IoU} is weight to balance the losses.

5. Experiments

5.1. Datasets

AVSBench [70] is a newly released dataset for audio-visual segmentation that comes with human-annotated mask annotations. This dataset is divided into two subsets: the semi-supervised Single Sound Source Segmentation (S4) and the fully supervised Multiple Sound Source Segmentation (MS3).

AVSBench-S4 dataset. This subset contains videos with a maximum of one sounding object. During training, only the first frame of each video sequence is annotated with a ground-truth segmentation binary mask. At inference time, the goal is to segment the sounding objects across all frames

in the video. The subset includes 3,452/740/740 videos in the training, validation, and test sets, respectively, with a total of 10,852 annotated frames.

AVSBench-MS3 dataset. The MS3 subset consists of videos that may contain multiple sounding objects. All video frames are annotated with masks during both the training and evaluation stages. The subset contains 296/64/64 videos for training, validation and testing separately, with a total of 2,120 annotated frames.

AVS-Synthetic dataset. As detailed in Sec 3.2, AVS-Synthetic is a dataset collected using our proposed pipeline. The dataset includes 46 categories with over 60K (*image, audio, mask*) sample triplets. To facilitate training and evaluation, the dataset is divided into training, validation and testing splits, which includes 52,609/5,000/5,000 samples separately.

5.2. Evaluation Metrics

We quantify the segmentation quality by adopting standard segmentation metrics outlined in [70]. Specifically, we use $\mathcal{M}_{\mathcal{J}}$ to denote the mean Intersection-over-Union (mIoU) of ground-truth binary masks and predicted masks. We use $\mathcal{M}_{\mathcal{F}}$, known as the F-score, which is the harmonic mean of precision and recall. In both cases, higher values indicate better segmentation performance.

5.3. Implementation Details

We choose the released ViT-H SAM model with pre-trained ViT-H/16 image encoder with 32 Transformer encoder layers as weight initialization. The resolution of the input images are 1024×1024 . The weight to balance the two losses is set to 1. To optimize the model parameters, we employ the AdamW [41] optimizer with an initial learning rate of 2×10^{-4} with cosine decay. We implement the models with PyTorch. We train the models for 80 epochs with four NVIDIA RTX A6000 GPUs. The batch size is set to 8.

5.4. Effectiveness of Data Synthesis Pipeline

We train our proposed SAMA-AVS model on the training split of the AVS-Synthetic dataset and present results on three experimental settings: standard evaluation on the test split of AVS-Synthetic, zero-shot evaluation on the test split of AVSBench, and few-shot evaluation on the test split of AVSBench after real data finetuning.

Results on synthetic dataset. We trained the SAMA-AVS model using the training and validation splits of AVS-Synthetic dataset, and then evaluate on the test split. It achieves the results of 79.21 on $\mathcal{M}_{\mathcal{J}}$ and 0.855 on $\mathcal{M}_{\mathcal{F}}$. Figure 3 provides some qualitative results; and our model can accurately segment the objects semantically corresponding to the audios.



Figure 3. Qualitative results on the test split of AVS-Synthetic.

Trained with	ambulance	cat	dog	bus	horse	lion	bird	guitar	piano	computer_keyboard
Real	73.63	83.87	84.88	85.09	81.21	91.28	85.64	85.78	82.91	85.99
Synthetic	55.32	78.81	78.66	82.70	73.88	88.51	81.99	75.11	32.26	74.01

Table 2. Categorized results of direct zero-shot evaluations of the models trained with synthetic data on real test split of AVSBench (S4) [70]; The “Real” row means the model is trained with the real data of AVSBench training split; and the “Synthetic” row means the model is trained with AVS-Synthetic. Metric: $\mathcal{M}_{\mathcal{J}}$ (mIoU).

Real Data	w/o P.T.		w/ P.T.	
	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$
0%	17.65	0.227	64.05 (+46.1)	0.744
10%	72.76	0.813	76.22 (+3.46)	0.844
20%	75.93	0.840	79.05 (+3.21)	0.870
30%	78.20	0.859	80.41 (+2.21)	0.876
100%	81.53	0.886	83.17 (+1.64)	0.901

Table 3. Finetuning the model with different percentages of real training data from AVSBench dataset. “w/ P.T.” refers to the model pretrained on AVS-Synthetic. 0%, 10%, 20% and 30% are the ratios of real data from AVSBench S4 training set; while “w/o P.T.” means the normal training without using pretrained weights as initialization.

Zero-shot results on real dataset. In this section, we examine the zero-shot transfer setting. Concretely, we train the model on AVS-Synthetic dataset and test it on the AVS-Bench S4 test set. Table 2 quantitatively showcases the results for the overlapping categories between AVS-Synthetic and AVSBench. Surprisingly, the model achieves remarkably close results on certain categories, such as “cat”, “bus”, “bus” and “lion”, demonstrating the effectiveness of our proposed data synthesis procedure, that enables zero-shot transfer to real data. We also observe the model trained with synthetic data generally lags behind the one trained with real data for all categories, which may be attributed to the domain gap between different image/video datasets. The video of the AVSBench dataset are from VGGSound and tend to be object-centric, while the image sets of AVS-Synthetic comprise LVIS and Open Images which are scene-centric. In future work, further domain adaptation

techniques should be investigated to alleviate this issue.

Data-efficient adaptation on real data. To narrow the gap between synthetic and real data, we adopt a small number of real data from the AVSBench-S4 subset, and fine-tune the model trained with AVS-Synthetic, resembling the data-efficient few-shot learning scenario. The results in Table 3 demonstrate that the model trained on AVS-Synthetic achieves 64.05 mIoU ($\mathcal{M}_{\mathcal{J}}$) without using any real data. However, by using only 10% of real data for fine-tuning, the performance can be significantly improved to 76.22 mIoU ($\mathcal{M}_{\mathcal{J}}$). We observe *consistent* improvement when more real data (20%/30%/100%) is provided for training. And the model (w/ P.T.) trained using 10% real data outperforms the model (w/o P.T.) using 20% real data. Similarly, the model (w/ P.T.) trained using 20% real data outperforms the model (w/o P.T.) using 30% real data. The results show that synthetic dataset can effectively reduce the demand for annotating real data with its better performance. And when using 100% real data, the model using the weights of AVS-Synthetic pretrained model improves the performance by 1.64 $\mathcal{M}_{\mathcal{J}}$, achieving 83.17 $\mathcal{M}_{\mathcal{J}}$ and 0.901 $\mathcal{M}_{\mathcal{F}}$.

Overall, these results demonstrate that the model trained with synthetic data from our proposed pipeline brings two benefits: (i) improving real-data efficiency, (ii) boosting the model’s overall performance for the AVS task.

5.5. Effectiveness of SAMA-AVS

In this part, we conduct experiments to compare the performance of our proposed SAMA-AVS with other SAM-based methods in Sec. 5.5.1 and other state-of-the-art methods in Sec. 5.5.2. Ablation study is done in 5.5.3. We

Subsets Methods	S4		MS3	
	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$
SAM [27]	55.08	0.739	53.96	0.638
AV-SAM [45]	40.47	0.566	-	-
AP-SAM	69.61	0.796	51.58	0.578
SAMA-AVS	81.53	0.886	63.14	0.691

Table 4. Comparison between different SAM-based methods on the testsets of S4 and MS3 subsets of AVSBench dataset.

also present some segmentation results qualitatively in Sec. 5.5.4. To ensure fair comparison, we do not utilize synthetic data pretraining unless otherwise stated.

5.5.1 Comparison with SAM-based Baselines.

We evaluate the performance of different methods on AVSBench, including SAM, AV-SAM [45], AP-SAM and SAMA-AVS. For the SAM baseline, we load the released model weights of ViT-H SAM model without further training and adopt the maximum segmentation evaluation as described in Sec. 4.2. For the AV-SAM [45] method, we cite the results from its original paper. For the proposed AP-SAM and SAMA-AVS, we use the same experimental settings to train and evaluate the models on the two subsets of AVSBench separately.

The results are shown in Table 4. On both subsets, SAMA-AVS outperforms other SAM-based methods by a large margin across all metrics. On the S4 subset, we observe the AP-SAM achieves better performance than the vanilla SAM, and lags far behind SAMA-AVS. This may be attributed to the fact that, audio-visual fusion of AP-SAM at the decoding stage is insufficient, *i.e.*, decoder has only two cross-attention layers for fusion; while SAMA-AVS enables more effective audio-visual fusion at the encoding stage with simple adapters.

5.5.2 Comparison with Other SOTA methods.

We also compare our proposed SAMA-AVS with the state-of-the-art AVS method [70], which is based on an encoder-fusion-decoder framework. Additionally, we compare with other related audio-visual methods, including sound source localization (SSL) methods LVS [3] and MSSL [47], video object segmentation (VOS) methods 3DC [43] and SST [9], and salient object detection (SOD) methods iGAN [44] and LGVT [64].

According to the results presented in Table 5, our proposed approach outperforms existing methods in both subsets across all metrics, for example, achieving 81.53 for S4 subset and 63.14 for MS3 subset on $\mathcal{M}_{\mathcal{J}}$. Note that, the results indicate that the performance of all methods is notably inferior on the MS3 subset as compared to the S4 subset.

Subsets Methods	S4		MS3		
	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	
SSL	LVS [3]	37.9	.510	29.5	.330
	MSSL [47]	44.9	.663	26.1	.363
VOS	3DC [43]	57.1	.759	36.9	.503
	SST [9]	66.3	.801	42.6	.572
SOD	iGAN [44]	61.6	.778	42.9	.544
	LGVT [64]	74.9	.873	40.7	.593
AVS [70]	w/o F.T.	78.7	.879	54.0	.645
	w/ F.T.	-	-	51.45	.671
Ours	w/o F.T.	81.53	.886	63.14	.691
	w/ F.T.	-	-	66.30	.730
	w/ F.T. (Syn.)	83.17	.901	66.95	.754

Table 5. Comparison with the state-of-the-art method and other related audio-visual methods on test sets of two subsets of AVSBench dataset. (“F.T.” denotes finetuning on the pretrained weights on S4; “F.T.(Syn.)” denotes finetuning on the pretrained weights of model trained with AVS-Synthetic).

We conjecture that the causes stem from two aspects: (i) the inherent difficulty of multi-sound segmentation; (ii) the limited dataset size of MS3 subset. To address this issue, we fine-tune the models that have already been trained on the S4 subset for MS3 subset. And our method gives 3.16 *improvement* on $\mathcal{M}_{\mathcal{J}}$ while AVS [70] has 2.55 *drop*. This shows that our method can effectively leverages the knowledge from the single sound object segmentation to improve the performance for multi-sound objects.

Lastly, we train the models for both S4 and MS3 of AVSBench dataset, by finetuning the weights pretrained from AVS-Synthetic dataset. The models (w/F.T.(Syn.)) obtain further improvements which achieve 83.17 and 66.95 on $\mathcal{M}_{\mathcal{J}}$ for two subsets separately, outperforming current methods by a large margin. This shows the effectiveness of the proposed AVS-Synthetic dataset, implying the importance of data and model in tackling the AVS task.

5.5.3 Ablation Study

Here we perform an ablation study to investigate the training objectives and the choice for the number of adapters in our proposed SAMA-AVS method.

Training objectives. We train the models with either the BCE loss \mathcal{L}_{BCE} or the IoU loss \mathcal{L}_{IoU} on AVSBench S4 subset, and compare the results with the model trained with the combined loss. The results in Table 6 shows that model trained with only BCE loss can already obtain satisfactory performance; by contrast, using the IoU loss only lead to very poor performance. By combining the two objectives together, the model achieves the best performance.

Number of adapters. We perform the experiments by varying the number of adapters in the model design using AVSBench S4 subset. In the default SAMA-AVS design, the

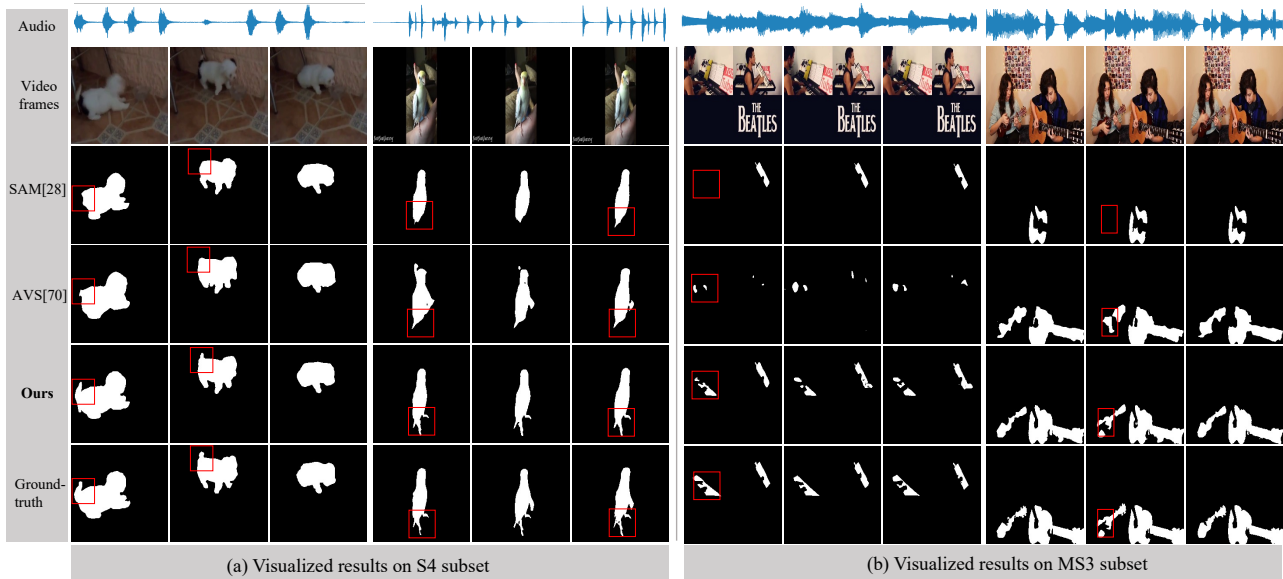


Figure 4. Qualitative results produced by our model and other methods on the test sets of S4 and MS3 subsets of AVSBench.

Model with	\mathcal{L}_{BCE}	\mathcal{L}_{IoU}	$\mathcal{L}_{BCE} + \mathcal{L}_{IoU}$
$\mathcal{M}_{\mathcal{J}}$	80.52	18.96	81.53
$\mathcal{M}_{\mathcal{F}}$	0.881	0.226	0.886

Table 6. Comparison between the models trained with different training losses.

image encoder has 32 transformer encoder layers thus we insert an adapter before feeding the inputs for each of the layer, ending up with a total of 32 adapters. We set the number of adapters to $\mathcal{N} \in \{2, 4, 8, 16, 32\}$, and evenly distribute \mathcal{N} adapters among 32 transformer encoder layers. The results in Table 7 show that the performance of the model consistently improves as the number of adapters increases, demonstrating the importance of audio-visual fusion that the adapters facilitate.

Adapter Num.	2	4	8	16	32
$\mathcal{M}_{\mathcal{J}}$	77.84	78.55	80.31	81.29	81.53
$\mathcal{M}_{\mathcal{F}}$.859	.857	.871	.880	.886

Table 7. Comparison between the models trained with different numbers of adapters.

5.5.4 Qualitative Examples

Figure 4 illustrates some qualitative results produced by different methods on AVSBench. The results show that on both subsets, the segmentations produced by our proposed SAMA-AVS best matches the ground-truths. Our model can precisely segment the whole body of the sounding objects without missing their parts; and it excels in handling the edges and details of the sounding target compared with

other competitors, such as the tail of the dog, and the tail and feet of the bird, marked with red boxes in Figure 4 (a). For the MS3 setting, our model can also localise the multiple targets from the sound and segment them accurately.

6. Conclusion

In summary, this paper addresses the audio-visual segmentation task from two perspectives: training data and computational model. For the former, we present a pipeline to synthesize audio-visual segmentation dataset by leveraging existing image segmentation datasets and audio dataset. The proposed pipeline is *annotation-free*, thus can be easily *scaled up*. And as an instantiation of the pipeline, we collect the AVS-Synthetic dataset to facilitate the study of AVS problem. For the latter, we propose an architecture SAMA-AVS that aims to harness the ability of SAM foundation model with minimal training. Our proposed model introduces very few *trainable* parameters ($<0.4\%$) with the design of adapters to perform audio-visual fusion in the encoding stage, which is highly *efficient* and *effective*. Extensive experimental results demonstrate the effectiveness of our proposed method, outperforming existing SAM-based methods and other competitors by a large margin. By pre-training the proposed model on the synthetic dataset, the model further attains higher performance on both subsets of real AVSBench dataset. We hope our proposed dataset and method will inspire more research on the audio-visual understanding.

Acknowledgements: This work was supported by Shanghai Science and Technology Committee under Grant (No. 21511101100).

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. [3](#)
- [2] Jiazhong Cen, Zanwei Zhou, Jiemin Fang, Wei Shen, Lingxi Xie, Dongsheng Jiang, Xiaopeng Zhang, and Qi Tian. Segment anything in 3d with nerfs. *CoRR*, abs/2304.12308, 2023. [3](#)
- [3] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. Localizing visual sounds the hard way. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16867–16876, 2021. [1](#), [2](#), [7](#)
- [4] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vggsound: A large-scale audio-visual dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 721–725. IEEE, 2020. [3](#)
- [5] Tianle Chen, Zheda Mai, Ruiwen Li, and Wei-Lun Chao. Segment anything model (SAM) enhanced pseudo labels for weakly supervised semantic segmentation. *CoRR*, abs/2305.05803, 2023. [3](#)
- [6] Tianrun Chen, Lanyun Zhu, Chaotao Ding, Runlong Cao, Shangzhan Zhang, Yan Wang, Zejian Li, Lingyun Sun, Papa Mao, and Ying Zang. Sam fails to segment anything?—sam-adapter: Adapting sam in underperformed scenes: Camouflage, shadow, and more. *arXiv preprint arXiv:2304.09148*, 2023. [3](#)
- [7] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. *arXiv preprint arXiv:2205.08534*, 2022. [3](#)
- [8] Dongjie Cheng, Ziyuan Qin, Zekun Jiang, Shaoting Zhang, Qicheng Lao, and Kang Li. SAM on medical images: A comprehensive study on three prompt modes. *CoRR*, abs/2305.00035, 2023. [3](#)
- [9] Brendan Duke, Abdalla Ahmed, Christian Wolf, Parham Aarabi, and Graham W Taylor. Sstvos: Sparse spatiotemporal transformers for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5912–5921, 2021. [7](#)
- [10] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T. Freeman, and Michael Rubinstein. Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation. *ACM Trans. Graph.*, 37(4):112, 2018. [2](#)
- [11] Ruohan Gao and Kristen Grauman. Visualvoice: Audio-visual speech separation with cross-modal consistency. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15490–15500. IEEE, 2021. [2](#)
- [12] Yifan Gao, Wei Xia, Dingdu Hu, and Xin Gao. Desam: Decoupling segment anything model for generalizable medical image segmentation. *CoRR*, abs/2306.00499, 2023. [3](#)
- [13] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 776–780. IEEE, 2017. [3](#), [4](#)
- [14] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5356–5364, 2019. [2](#), [3](#)
- [15] Chunming He, Kai Li, Yachao Zhang, Guoxia Xu, Longxiang Tang, Yulun Zhang, Zhenhua Guo, and Xiu Li. Weakly-supervised concealed object segmentation with sam-based pseudo labeling and multi-scale feature grouping. *CoRR*, abs/2305.11003, 2023. [3](#)
- [16] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [4](#)
- [17] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. Cnn architectures for large-scale audio classification. In *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*, pages 131–135. IEEE, 2017. [4](#), [5](#)
- [18] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019. [3](#)
- [19] Di Hu, Feiping Nie, and Xuelong Li. Deep multimodal clustering for unsupervised audiovisual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9248–9257, 2019. [1](#), [2](#)
- [20] Di Hu, Rui Qian, Minyue Jiang, Xiao Tan, Shilei Wen, Errui Ding, Weiyao Lin, and Dejing Dou. Discriminative sounding objects localization via self-supervised audiovisual matching. *Advances in Neural Information Processing Systems*, 33:10077–10087, 2020. [1](#), [2](#)
- [21] Di Hu, Yake Wei, Rui Qian, Weiyao Lin, Ruihua Song, and Ji-Rong Wen. Class-aware sounding objects localization via audiovisual correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [1](#), [2](#)
- [22] Ge-Peng Ji, Deng-Ping Fan, Peng Xu, Ming-Ming Cheng, Bowen Zhou, and Luc Van Gool. SAM struggles in concealed scenes - empirical study on "segment anything". *CoRR*, abs/2304.06022, 2023. [3](#), [4](#)
- [23] Wei Ji, Jingjing Li, Qi Bi, Wenbo Li, and Li Cheng. Segment anything is not always perfect: An investigation of SAM on different real-world applications. *CoRR*, abs/2304.05750, 2023. [3](#)
- [24] Peng-Tao Jiang and Yuqi Yang. Segment anything is A good pseudo-label generator for weakly supervised semantic segmentation. *CoRR*, abs/2305.01275, 2023. [3](#)
- [25] Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter

- layers. *Advances in Neural Information Processing Systems*, 34:1022–1035, 2021. 3
- [26] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5492–5501, 2019. 2
- [27] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 2, 3, 4, 7
- [28] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision*, 128(7):1956–1981, 2020. 2, 3
- [29] Jun-Tae Lee, Mihir Jain, Hyoungwoo Park, and Sungrack Yun. Cross-attentional audio-visual fusion for weakly-supervised action localization. In *International Conference on Learning Representations*, 2020. 2
- [30] Sangho Lee, Youngjae Yu, Gunhee Kim, Thomas Breuel, Jan Kautz, and Yale Song. Parameter efficient multimodal transformers for video representation learning. *arXiv preprint arXiv:2012.04124*, 2020. 3
- [31] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, pages 280–296. Springer, 2022. 3
- [32] Yan-Bo Lin, Hung-Yu Tseng, Hsin-Ying Lee, Yen-Yu Lin, and Ming-Hsuan Yang. Unsupervised sound localization via iterative contrastive learning. *CoRR*, abs/2104.00315, 2021. 1, 2
- [33] Yan-Bo Lin, Yu-Jhe Li, and Yu-Chiang Frank Wang. Dual-modality seq2seq network for audio-visual event localization. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2002–2006. IEEE, 2019. 2
- [34] Yan-Bo Lin, Yi-Lin Sung, Jie Lei, Mohit Bansal, and Gedas Bertasius. Vision transformers are parameter-efficient audio-visual learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2299–2309, 2023. 3
- [35] Ziyi Lin, Shijie Geng, Renrui Zhang, Peng Gao, Gerard de Melo, Xiaogang Wang, Jifeng Dai, Yu Qiao, and Hongsheng Li. Frozen clip models are efficient video learners. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 388–404. Springer, 2022. 3
- [36] Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin A Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35:1950–1965, 2022. 3
- [37] Jinxiang Liu, Chen Ju, Weidi Xie, and Ya Zhang. Exploiting transformation invariance and equivariance for self-supervised sound localisation. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 3742–3753, 2022. 1, 2
- [38] Weihuang Liu, Xi Shen, Chi-Man Pun, and Xiaodong Cun. Explicit visual prompting for low-level structure segmentations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19434–19445, 2023. 3
- [39] Youquan Liu, Lingdong Kong, Jun Cen, Runnan Chen, Wenwei Zhang, Liang Pan, Kai Chen, and Ziwei Liu. Segment any point cloud sequences by distilling vision foundation models. *CoRR*, abs/2306.09347, 2023. 3
- [40] Yang Liu, Muzhi Zhu, Hengtao Li, Hao Chen, Xinlong Wang, and Chunhua Shen. Matcher: Segment anything with one shot using all-purpose feature matching. *CoRR*, abs/2305.13310, 2023. 3
- [41] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [42] Jun Ma and Bo Wang. Segment anything in medical images. *CoRR*, abs/2304.12306, 2023. 3
- [43] Sabarinath Mahadevan, Ali Athar, Aljosa Osep, Laura Leal-Taixé, Bastian Leibe, and Sebastian Hennen. Making a case for 3d convolutions for object segmentation in videos. In *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, September 7-10, 2020*. BMVA Press, 2020. 7
- [44] Yuxin Mao, Jing Zhang, Zhexiong Wan, Yuchao Dai, Aixuan Li, Yunqiu Lv, Xinyu Tian, Deng-Ping Fan, and Nick Barnes. Transformer transforms salient object detection and camouflaged object detection. *CoRR*, abs/2104.10127, 2021. 7
- [45] Shentong Mo and Yapeng Tian. Av-sam: Segment anything model meets audio-visual localization and segmentation. *arXiv preprint arXiv:2305.01836*, 2023. 3, 7
- [46] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–648, 2018. 2
- [47] Rui Qian, Di Hu, Heinrich Dinkel, Mengyue Wu, Ning Xu, and Weiyao Lin. Multiple sound sources localization from coarse to fine. In *European Conference on Computer Vision*, pages 292–308. Springer, 2020. 2, 7
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 4
- [49] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Efficient parametrization of multi-domain deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8119–8127, 2018. 3
- [50] Andrew Rouditchenko, Hang Zhao, Chuang Gan, Josh McDermott, and Antonio Torralba. Self-supervised audio-visual co-segmentation. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2357–2361. IEEE, 2019. 1, 2

- [51] Arda Senocak, Tae-Hyun Oh, Junsik Kim, Ming-Hsuan Yang, and In So Kweon. Learning to localize sound source in visual scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4358–4366, 2018. 1, 2
- [52] Tal Shaharabany, Aviad Dahan, Raja Giryes, and Lior Wolf. Autosam: Adapting SAM to medical images by overloading the prompt encoder. *CoRR*, abs/2306.06370, 2023. 3
- [53] Gilad Sharir, Asaf Noy, and Lih Zelnik-Manor. An image is worth 16x16 words, what is a video worth? *arXiv preprint arXiv:2103.13915*, 2021. 4
- [54] Zengjie Song, Yuxi Wang, Junsong Fan, Tieniu Tan, and Zhaoxiang Zhang. Self-supervised predictive learning: A negative-free method for sound source localization in visual scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3222–3231, 2022. 1, 2
- [55] Asa Cooper Stickland and Iain Murray. Bert and pals: Projected attention layers for efficient adaptation in multi-task learning. In *International Conference on Machine Learning*, pages 5986–5995. PMLR, 2019. 3
- [56] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. Lst: Ladder side-tuning for parameter and memory efficient transfer learning. *arXiv preprint arXiv:2206.06522*, 2022. 3
- [57] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. VI-adapter: Parameter-efficient transfer learning for vision-and-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5227–5237, 2022. 3
- [58] Lv Tang, Haoke Xiao, and Bo Li. Can sam segment anything? when sam meets camouflaged object detection. *arXiv preprint arXiv:2304.04709*, 2023. 4
- [59] Yapeng Tian, Dingzeyu Li, and Chenliang Xu. Unified multisensory perception: Weakly-supervised audio-visual video parsing. In *European Conference on Computer Vision*, pages 436–454. Springer, 2020. 2
- [60] Efthymios Tzinis, Scott Wisdom, Tal Remez, and John R. Hershey. Audioscopev2: Audio-visual attention architectures for calibrated open-domain on-screen sound separation. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXVII*, volume 13697 of *Lecture Notes in Computer Science*, pages 368–385. Springer, 2022. 2
- [61] An Wang, Mobarakol Islam, Mengya Xu, Yang Zhang, and Hongliang Ren. SAM meets robotic surgery: An empirical study in robustness perspective. *CoRR*, abs/2304.14674, 2023. 3
- [62] Yonghui Wang, Wengang Zhou, Yunyao Mao, and Houqiang Li. Detect any shadow: Segment anything for video shadow detection. *CoRR*, abs/2305.16698, 2023. 3
- [63] Junde Wu, Rao Fu, Huihui Fang, Yuanpei Liu, Zhaowei Wang, Yanwu Xu, Yueming Jin, and Tal Arbel. Medical SAM adapter: Adapting segment anything model for medical image segmentation. *CoRR*, abs/2304.12620, 2023. 3
- [64] Jing Zhang, Jianwen Xie, Nick Barnes, and Ping Li. Learning generative vision transformer with energy-based latent space for saliency prediction. *Advances in Neural Information Processing Systems*, 34:15448–15463, 2021. 7
- [65] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. *CoRR*, abs/2305.03048, 2023. 3
- [66] Xiaofeng Zhang, Chaochen Gu, and Shanying Zhu. Sam-helps-shadow: When segment anything model meet shadow removal. *CoRR*, abs/2306.06113, 2023. 3
- [67] Hang Zhao, Chuang Gan, Wei-Chiu Ma, and Antonio Torralba. The sound of motions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1735–1744, 2019. 2
- [68] Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. The sound of pixels. In *Proceedings of the European conference on computer vision (ECCV)*, pages 570–586, 2018. 2
- [69] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016. 2
- [70] Jinxing Zhou, Jianyuan Wang, Jiayi Zhang, Weixuan Sun, Jing Zhang, Stan Birchfield, Dan Guo, Lingpeng Kong, Meng Wang, and Yiran Zhong. Audio-visual segmentation. In *European Conference on Computer Vision*, 2022. 1, 2, 3, 5, 6, 7