
Continual Audio-Visual Sound Separation

Weiguo Pian¹ Yiyang Nan² Shijian Deng¹ Shentong Mo³ Yunhui Guo¹ Yapeng Tian¹

¹ The University of Texas at Dallas ² Brown University ³ Carnegie Mellon University

Abstract

In this paper, we introduce a novel continual audio-visual sound separation task, aiming to continuously separate sound sources for new classes while preserving performance on previously learned classes, with the aid of visual guidance. This problem is crucial for practical visually guided auditory perception as it can significantly enhance the adaptability and robustness of audio-visual sound separation models, making them more applicable for real-world scenarios where encountering new sound sources is commonplace. The task is inherently challenging as our models must not only effectively utilize information from both modalities in current tasks but also preserve their cross-modal association in old tasks to mitigate catastrophic forgetting during audio-visual continual learning. To address these challenges, we propose a novel approach named ContAV-Sep (**Continual Audio-Visual Sound Separation**). ContAV-Sep presents a novel Cross-modal Similarity Distillation Constraint (CrossSDC) to uphold the cross-modal semantic similarity through incremental tasks and retain previously acquired knowledge of semantic similarity in old models, mitigating the risk of catastrophic forgetting. The CrossSDC can seamlessly integrate into the training process of different audio-visual sound separation frameworks. Experiments demonstrate that ContAV-Sep can effectively mitigate catastrophic forgetting and achieve significantly better performance compared to other continual learning baselines for audio-visual sound separation. Code is available at: https://github.com/weiguopian/ContAV-Sep_NeurIPS2024.

1 Introduction

Humans can effortlessly separate and identify individual sound sources in daily experience [25, 7, 64, 33]. This skill plays a crucial role in our ability to understand and interact with the complex auditory environments that surround us [34]. However, replicating this capability in machines remains a significant challenge due to the inherent complexity of real-world auditory scenes [7, 77]. Inspired by the multisensory perception of humans [62, 60], audio-visual sound separation tackles this challenge by utilizing visual information to guide the separation of individual sound sources in an audio mixture.

Recent advances in deep learning have led to significant progress in audio-visual sound separation [84, 23, 21, 67, 14, 65, 81, 63, 11, 71]. Benefiting from more advanced architectures (*e.g.*, U-Net [84, 23], Transformer [14], and diffusion models [27]) and discriminative visual cues (*e.g.*, grounded visual objects [67], motion [83], and dynamic gestures [21]), audio-visual separation models are able to separate sounds ranging from domain-specific speech, musical instrument sounds to open-domain general sounds within training sound categories. However, a limitation of these studies is their focus on scenarios where all sound source classes are presently known, overlooking the potential inclusion of unknown sound source classes during inference in real-world applications. This oversight leads to the *catastrophic forgetting* issue [32, 3], where the fine-tuning of models on new classes detrimentally impacts their performance on previously learned classes. Despite Chen et al. [14] demonstrating that their iQuery model can generalize to new classes well through simple fine-tuning, it still suffers from the catastrophic forgetting problem on old classes. This prevents the trained models from

continuously updating in real-world scenarios, impeding their adaptability to dynamic environments. The question *how to effectively leverage visual guidance to continuously separate sounds from new categories while preserving separation ability for old sound categories* remains open.

To bridge this gap, we introduce a novel *continual audio-visual sound separation* task by integrating audio-visual sound separation with continual learning principles. The goal of this task is to develop an audio-visual model that can continuously separate sound sources in new classes while maintaining performance on previously learned classes. The key challenge we need to address is catastrophic forgetting during continual audio-visual learning, which occurs when the model is updated solely with data from new classes or tasks, resulting in a significant performance drop on old ones. We illustrate our new task and the catastrophic forgetting issue in Fig. 1.

Unlike typical continual learning problems such as task-, domain-, or class-incremental classification in visual domains [2, 57, 38, 53, 85], which result in progressively increasing logits (or probability distribution) across all observed classes at each incremental step, our task uniquely produces fixed-size separation masks throughout all incremental steps. In this context, each entry in the mask does not directly correspond to any specific classes. Additionally, the new task involves both audio and visual modalities. Therefore, simply applying existing visual-only methods cannot fully exploit and preserve the inherent cross-modal semantic correlations. Very recently, Pian *et al.* [53] and Mo *et al.* [44] extended continual learning to the audio-visual domain, but both focused on classification tasks.

To address these challenges, in this paper, we propose a novel approach named ContAV-Sep (**Continual Audio-Visual Sound Separation**). Upon the framework, we introduce a novel *Cross-modal Similarity Distillation Constraint (CrossSDC)* to not only maintain the cross-modal semantic similarity through incremental tasks but also preserve previously learned knowledge of semantic similarity in old models to counter catastrophic forgetting. The CrossSDC is a generic constraint that can be seamlessly integrated into the training process of different audio-visual sound separators. To evaluate the effectiveness of our proposed ContAV-Sep, we conducted experiments on the MUSIC-21 dataset within the framework of continual learning, using the state-of-the-art audio-visual sound separation model iQuery [14] and a representative audio-visual sound separation model Co-Separation [23], as our separation base models. Experiments demonstrate that ContAV-Sep can effectively mitigate catastrophic forgetting and achieve significantly better performance than other continual learning baselines. In summary, this paper contributes follows:

- (i) To explore more practical audio-visual sound separation, in which the separation model should be generalized to new sound source classes continually, we pose a *Continual Audio-Visual Sound Separation* task that trains the separation model under the setting of continual learning. To the best of our knowledge, this is the first work on continual learning for audio-visual sound separation.
- (ii) We propose ContAV-Sep for the new task. It uses a novel cross-modal similarity distillation constraint to preserve cross-modal semantic similarity knowledge from previously learned models.
- (iii) Experiments on the MUSIC-21 dataset can validate the effectiveness of our ContAV-Sep, demonstrating promising performance gain over baselines.

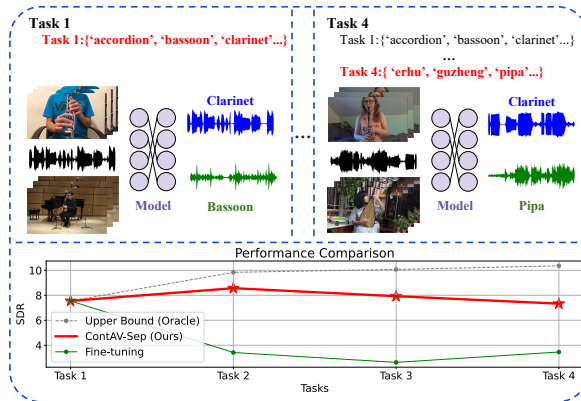


Figure 1: **Top:** Illustration of the continual audio-visual sound separation task, where the model (separator) learns from sequential audio-visual sound separation tasks. **Bottom:** Illustration of the catastrophic forgetting problem in continual audio-visual sound separation and its mitigation by our proposed method. Fine-tuning: Directly fine-tune the separation model on new sound source classes; Upper bound: Train the model using all training data from seen sound source classes.

2 Related Work

Audio-Visual Sound Separation. Audio-visual sound separation aims to separate individual sound sources from an audio mixture guided by visual cues. A line of research emerges under various scenarios, such as separating musical instruments [23, 84, 79, 21, 83, 67], human speech [20, 1, 18, 49, 15], or sound sources in in-the-wild videos [22, 70]. Many frameworks and methods have been proposed to address challenges within specific problem settings. For instance, the extraction of face embeddings proves beneficial for speech audio separation [18]. Moreover, incorporating object detection can provide an additional advantage [23, 22]. The utilization of trajectory optical flows to leverage temporal motion information in videos, as demonstrated by [83], also yields improvements. In this work, not competing on designing stronger separators, we would advance the exploration of the audio-visual sound separation within the paradigm of continual learning. We investigate how a model can learn to consistently separate sound sources from sequential separation tasks without forgetting previously acquired knowledge.

Continual Learning. The field of continual learning has drawn significant attention, especially in visual domains, with various approaches addressing this challenge. Notable among these are regularization-based methods, exemplified in works such as [32, 3, 31, 39]. These approaches involve applying regularization to crucial parameters associated with old tasks to maintain the model’s capabilities and during incremental steps, less important parameters are given higher priority for updates compared to important ones. Conversely, several works [57, 9, 5, 26, 12, 54, 8, 10, 40] applied rehearsal-based pipelines to enable the model review previously learned knowledge. For instance, Rebuffi *et al.* [57] proposed one of the most representative exemplars selection strategy Nearest-Mean-of-Exemplars (NME), selects the most representative exemplars in each class based on the distance to the feature center of the class. Meanwhile, pseudo-rehearsal [47, 48, 66] employs generative models to create pseudo-exemplars based on the estimated distribution of data from previous classes. Moreover, architecture-based/dynamic architecture methods [52, 4, 46, 24, 28, 37, 40] proposed to modify the model architecture itself to enable the model to acquire new knowledge while mitigating the forgetting of old knowledge. Specifically, Pham *et al.* [52] proposed a dual network architecture, in which one is to learn new tasks while the other one is for retaining knowledge learned from old tasks. Wang *et al.* [72] combined the dynamic architecture and distillation constraint to mitigate the issue of continual-increasing overhead problem in dynamic architecture-based continual learning method. However, above studies mainly concentrate on the area of continual image classification. Recently, researchers also explored other continual learning scenarios beyond image classification. For instance, Park *et al.* [50] extend the knowledge distillation-based [2, 17] continual image classification method to the domain of video by proposing the time-channel distillation constraint. Douillard *et al.* [16] proposed to tackle the continual semantic segmentation task with multi-view feature distillation and pseudo-labeling. Xiao *et al.* [78] further addressed the continual semantic segmentation problem through weights fusion strategy between old and current models. Wang *et al.* [75] addressed the continual sound classification task through generative replay. Furthermore, continual learning has also been explored in the domain of language/vision-language learning tasks [30, 43, 59, 61, 19, 86], self-supervised representation learning [19, 42, 55, 80, 35, 74], audio classification [6, 73] and fake audio detection [41, 82], etc. Despite the success of existing continual learning methods in various scenarios, their applicability in the domain of continual audio-visual sound separation is still unexplored. Although Pian *et al.* [53] and Mo *et al.* [44] proposed to tackle the catastrophic problem in audio-visual learning, their studies mainly concentrated in the area of audio-visual video classification. In contrast to existing works in continual learning, in this paper, we delves into the continual audio-visual sound separation, aiming to tackle the challenge of catastrophic forgetting specifically in the context of separation mask prediction for complicated mixed audio signals within joint audio-visual modeling.

3 Method

3.1 Problem Formulation

Audio-Visual Sound Separation. Audio-visual sound separation aims to separate distinctive sound signals according to the given associated visual guidance. Following previous works [14, 23, 67, 21, 79], we adopt the common “mix-and-separation” training strategy to train the model. Given two videos $V_1(s_1, v_1)$ and $V_2(s_2, v_2)$, we can obtain the input mixed sound signal S by mixing two

video sound signals s_1 and s_2 , and then we can have the ratio masks $M^1 = s_1/S$ and $M^2 = s_2/S$ ¹. The goal of the task is to utilize the corresponding visual guidance v_1 and v_2 to predict the ratio masks for reconstructing the two individual audio signals. This process can be formulated as:

$$\begin{aligned}\hat{M}^1 &= \mathcal{F}_\Theta(\mathcal{S}, v_1), \\ \hat{M}^2 &= \mathcal{F}_\Theta(\mathcal{S}, v_2),\end{aligned}\tag{1}$$

where \mathcal{F}_Θ is the separation model with trainable parameters Θ . And then, the original sound signals s_1 and s_2 are used to calculate the loss function for optimizing the model:

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \mathbb{E}_{(v_1, v_2) \sim \mathcal{D}} \left[\mathcal{L}(\hat{M}^1, M^1) + \mathcal{L}(\hat{M}^2, M^2) \right],\tag{2}$$

where \mathcal{D} denotes the training set, and \mathcal{L} is the loss function between the prediction and ground-truth.

Continual Audio-Visual Sound Separation. Our proposed continual audio-visual sound separation task aims to train a model \mathcal{F}_Θ continually on a sequence of T separation tasks $\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_T\}$. For the t -th task \mathcal{T}_t (incremental step t), we have a training set $\mathcal{D}_t = \{\mathbf{V}^i(s^i, v^i), y_t^i\}_{i=1}^{n_t}$, where i and n_t denote the i -th video sample and the total number of samples in \mathcal{D}_t respectively, and $y_t^i \in \mathcal{C}_t$ is the corresponding sound source class of video \mathbf{V}^i , where \mathcal{C}_t is the training sound class label space of task \mathcal{T}_t . For any two tasks \mathcal{T}_{t_1} and \mathcal{T}_{t_2} and their corresponding training sound class label space \mathcal{C}_{t_1} and \mathcal{C}_{t_2} , we have $\mathcal{C}_{t_1} \cap \mathcal{C}_{t_2} = \emptyset$. Following previous works in continual learning [57, 2, 53, 44, 29, 76], for a task \mathcal{T}_t , where $t > 1$, holding a small size of memory/exemplar set \mathcal{M}_t to store some data from old tasks is permitted in our setting. Therefore, with the memory/exemplar set \mathcal{M}_t , all available data that can be used for training in task \mathcal{T}_t ($t > 1$) can be denoted as $\mathcal{D}'_t = \mathcal{D}_t \cup \mathcal{M}_t$. Finally, the training process of Eq. 2 in our continual audio-visual sound separation setting can be denoted as:

$$\begin{aligned}\Theta_t &= \underset{\Theta_{t-1}}{\operatorname{argmin}} \mathbb{E}_{(v_1, v_2) \sim \mathcal{D}'_t} \left[\mathcal{L}(\hat{M}^1, M^1) + \mathcal{L}(\hat{M}^2, M^2) \right], \\ \text{s.t. } \hat{M}^1 &= \mathcal{F}_{\Theta_{t-1}}(\mathcal{S}, v_1), \hat{M}^2 = \mathcal{F}_{\Theta_{t-1}}(\mathcal{S}, v_2),\end{aligned}\tag{3}$$

which means that the new model Θ_t is obtained by updating the old model Θ_{t-1} which was trained on the previous task, using the current task’s available data \mathcal{D}'_t . After the training process for task \mathcal{T}_t with \mathcal{D}'_t , the updated model will be evaluated on a testing set which includes video samples from all seen sound source classes up to continual step t (task \mathcal{T}_t). And the evaluation also follows the common “mix-and-separation” strategy. During this continual learning process, the model’s separation performance on the previously learned tasks drops significantly after training on new tasks. This learning issue is referred to as the *catastrophic forgetting* [32, 38, 3] problem, which poses a considerable challenge in continual audio-visual sound separation.

3.2 Overview

To address the challenge of catastrophic forgetting in continual audio-visual sound separation, we introduce **ContAV-Sep**. This new framework, illustrated in Fig. 2, consists of three key components: a separation base model, an output mask distillation module, and our proposed *Cross-modal Similarity Distillation Constraint (CrossSDC)*. We use a recent state-of-the-art audio-visual separator: iQuery [14] as the base model of our approach, which contains a video encoder to extract the global motion feature, an object detector and image encoder to obtain the object feature, a U-Net [58] for mixture sound encoding and separated sound decoding, and an audio-visual Transformer to get the separated sound feature through multi-modal cross-attention mechanism and class-aware audio queries. For the object detector, we follow iQuery [14] and use the pre-trained Detic [87], a universal object detector, to detect the sound source objects in each frame. For the video encoder and the image encoder, inspired by the excellent generalization ability of recent self-supervised pre-trained models, which has been proven to be effective and appropriate in continual learning [53], we apply two self-supervised pre-trained models VideoMAE [69] and CLIP [56] as the video encoder and the image encoder, respectively. Note that, during the training process, the object detector, video encoder, and image encoder are frozen.

¹In practice, the audio signal is first processed using the Short-Time Fourier Transform (STFT) to generate a spectrogram. For brevity, we will denote spectrogram magnitudes as s_1 , s_2 , and S .

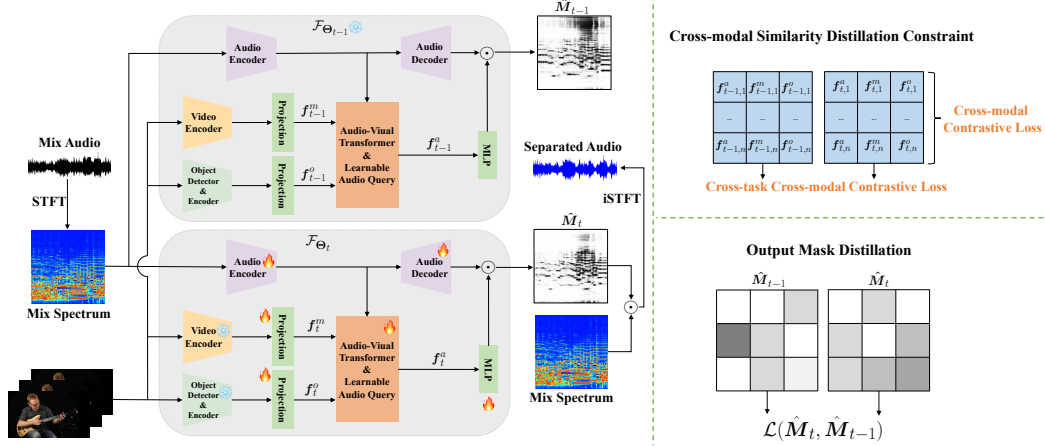


Figure 2: Overview of our proposed ContAV-Sep, which consists of an audio-visual sound separation base model architecture, an Output Mask Distillation, and our proposed Cross-modal Similarity Distillation Constraint. The fire icon denotes the module is trainable, while the snowflake icon denotes that the module is frozen. The (i)STFT stands for (inverse) Short-Time Fourier Transform. Please note that, the old model $\mathcal{F}_{\Theta_{t-1}}$ is frozen during training.

Given a pair of videos $\mathbf{V}_1(s_1, v_1)$ and $\mathbf{V}_2(s_2, v_2)$, at incremental step t (task \mathcal{T}_t), the U-Net audio encoder \mathcal{F}_t^{AE} takes the mixed audio signal \mathbf{S} obtained by mixing s_1 and s_2 as input, and generates the latent mixed audio feature. This process can be expressed as:

$$\mathbf{f}_t^{lat.} = \mathcal{F}_t^{AE}(\mathbf{S}), \quad (4)$$

Then, the audio-visual Transformer $\mathcal{F}_t^{Trans.}$ is employed to generate the separated sound feature by taking the latent mixed audio feature and visual features as inputs:

$$\begin{aligned} \mathbf{f}_t^{a,1} &= \mathcal{F}_t^{Trans.}(\mathbf{f}_t^{lat.}, \mathbf{f}_t^{o,1}, \mathbf{f}_t^{m,1}), \\ s.t. \quad \mathbf{f}_t^{o,1} &= \mathbf{U}_t^o(\mathbf{Obj}^1), \quad \mathbf{f}_t^{m,1} = \mathbf{U}_t^m(\mathbf{Mo}^1), \end{aligned} \quad (5)$$

where $\mathbf{f}_t^{a,1}$ denotes the separated sound feature of video \mathbf{V}_1 ; \mathbf{Obj}^1 and \mathbf{Mo}^1 denote the object and motion features extracted by the frozen pre-trained image and video encoders respectively from the visual signal v_1 of video \mathbf{V}_1 , $\mathbf{U}_t^o(\cdot)$ and $\mathbf{U}_t^m(\cdot)$ are learnable projection layers to map the object and motion features into the same dimension. Similarly, we can also obtain the separated sound feature of \mathbf{V}_2 guided by the associated visual features.

The extracted separated sound feature and the latent mixed audio feature are combined to generate a mask. This mask is subsequently applied to the mixed audio, leading to the reconstruction of the separated sound spectrogram.

$$\begin{aligned} \hat{M}_t^1 &= \mathcal{F}_t^{AD}(\mathbf{f}_t^{lat.}) \odot MLP_t(\mathbf{f}_t^{a,1}), \\ \hat{M}_t^2 &= \mathcal{F}_t^{AD}(\mathbf{f}_t^{lat.}) \odot MLP_t(\mathbf{f}_t^{a,2}), \end{aligned} \quad (6)$$

where \hat{M}_t^1 and \hat{M}_t^2 denote the predicted masks for audio signals of video \mathbf{V}_1 and \mathbf{V}_2 , respectively; \mathcal{F}_t^{AD} is the U-Net decoder at incremental step t ; $MLP_t(\cdot)$ denotes a MLP module; and \odot denotes channel-wise multiplication. The sound s_1 at this incremental step can be reconstructed by applying $\mathbf{S} \odot \hat{M}_t^1$ and then performing an inverse STFT to obtain the audio waveform.

3.3 Cross-modal Similarity Distillation Constraint

Recent studies [53, 45] have highlighted the importance of cross-modal semantic correlation in audio-visual modeling. However, this correlation tends to diminish during subsequent incremental phases, which leads to catastrophic forgetting in our continual audio-visual sound separation task. To address this challenge, we propose a novel Cross-modal Similarity Distillation Constraint (CrossSDC)

that serves two crucial purposes (1) maintaining cross-modal semantic similarity through incremental tasks, and (2) preserving previous learned semantic similarity knowledge from old tasks.

CrossSDC preserves cross-modal semantic similarity from two perspectives: instance-aware semantic similarity and class-aware semantic similarity. Both similarities are enforced by integrating contrastive loss and knowledge distillation. Instead of exclusively focusing on the similarities within current and memory data generated by the current training model, CrossSDC incorporates the cross-modal similarity knowledge acquired from previous tasks into the contrastive loss. This integration not only facilitates the learning of cross-modal semantic similarities in new tasks but also ensures the preservation of previously acquired knowledge. In the incremental step t ($t > 1$), the instance-aware part of our CrossSDC can be formulated as:

$$\mathcal{L}_{inst.} = -\mathbb{E}_{\mathbf{V}^i \sim \mathcal{D}'_t} \left[\frac{1}{\sum_j \mathbb{1}[i = j]} \sum_j \mathbb{1}[i = j] \log \frac{\exp(\text{sim}(\mathbf{f}_{\tau_1, i}^{mod_1}, \mathbf{f}_{\tau_2, j}^{mod_2}))}{\sum_k \exp(\text{sim}(\mathbf{f}_{\tau_1, i}^{mod_1}, \mathbf{f}_{\tau_2, k}^{mod_2}))} \right], \quad (7)$$

where $\mathbb{1}[i = j]$ is an indicator that equals 1 when $i = j$, denoting that video samples \mathbf{V}^i and \mathbf{V}^j are the same video; The sim function represents the cosine similarity function with temperature scaling; The modalities mod_1 and mod_2 , where $(mod_1, mod_2) \in \{(a, o), (a, m), (m, o)\}$, denote different pairs of features to be compared: separated sound and object features, sound and motion features, and motion and object features. Here, τ denotes the incremental step, for which we have:

$$\tau_1, \tau_2 \in \mathbf{T}, \text{ where } \mathbf{T} = \begin{cases} \{t, t-1\}, & \text{if } \mathbf{V} \in \mathcal{M}_t, \\ \{t\}, & \text{if } \mathbf{V} \in \mathcal{D}_t, \end{cases} \quad (8)$$

which means that, for current task’s data \mathcal{D}_t , we calculate the contrastive loss using features from the current model ($\tau_1 = \tau_2 = t$), while for memory set data \mathcal{M}_t , we use features from *both the old and current models* (e.g., $\tau_1 = t$ and $\tau_2 = t-1$). In this way, knowledge distillation would be integrated into the cross-modal semantic similarity constraint for the current task, which ensures better preservation of learned cross-modal semantic similarity from previous tasks.

While the instance-aware similarity provides valuable semantic correlation modeling, it does not account for the class-level semantic correlations, which is also crucial for audio-visual similarity modeling. To capture and preserve the semantic similarity within each class across incremental tasks, we also incorporate a class-aware component specifically designed for inter-class cross-modal semantic similarity, which can be formulated as:

$$\mathcal{L}_{cls.} = -\mathbb{E}_{(\mathbf{V}^i, y^i) \sim \mathcal{D}'_t} \left[\frac{1}{\sum_j \mathbb{1}[y^i = y^j]} \sum_j \mathbb{1}[y^i = y^j] \log \frac{\exp(\text{sim}(\mathbf{f}_{\tau_1, i}^{mod_1}, \mathbf{f}_{\tau_2, j}^{mod_2}))}{\sum_k \exp(\text{sim}(\mathbf{f}_{\tau_1, i}^{mod_1}, \mathbf{f}_{\tau_2, k}^{mod_2}))} \right]. \quad (9)$$

In this context, visual and audio features from two videos are encouraged to be close when they belong to the same class. The overall formulation of our CrossSDC is as follows:

$$\mathcal{L}_{CrossSDC} = \lambda_{ins} \mathcal{L}_{ins} + \lambda_{cls} \mathcal{L}_{cls}, \quad (10)$$

where λ_{ins} and λ_{cls} are two scalars that balance the two loss terms. In this way, the model captures and preserves semantic correlations not just between instances but also within the same classes.

3.4 Overall Loss Function

In the previous subsection, we introduced our proposed CrossSDC constraint. To effectively combine CrossSDC with the overall objective, we incorporate it alongside output distillation and the main separation loss function.

Output distillation is a widely used technique in continual learning [38, 2, 53] to preserve the knowledge gained from previous tasks while learning new ones. In our approach, we utilize the output of the old model as the distillation target to preserve this knowledge. Note that we only distill knowledge for data from the memory set, as represented by:

$$\mathcal{L}_{dist.} = \mathbb{E}_{(\mathbf{V}_1^i, \mathbf{V}_2^i) \sim \mathcal{M}_t} \left[\|\hat{\mathbf{M}}_t^1 - \hat{\mathbf{M}}_{t-1}^1\|_1 + \|\hat{\mathbf{M}}_t^2 - \hat{\mathbf{M}}_{t-1}^2\|_1 \right], \quad (11)$$

where $\hat{\mathbf{M}}_{t-1}^1$ and $\hat{\mathbf{M}}_{t-1}^2$ are predicted masks generated by the old model that is trained at incremental step $t-1$. For the loss function here, we follow [84, 14] and adopt the per-pixel L_1 loss [84]. For

the main separation loss function, we also apply the per-pixel L_1 loss:

$$\mathcal{L}_{main} = \mathbb{E}_{(\mathbf{V}_1^i, \mathbf{V}_2^i) \sim \mathcal{M}_t} \left[\|\hat{\mathbf{M}}_t^1 - \mathbf{M}^1\|_1 + \|\hat{\mathbf{M}}_t^2 - \mathbf{M}^2\|_1 \right], \quad (12)$$

Finally, our overall loss function is denoted as:

$$\mathcal{L}_{ContAV-Sep} = \mathcal{L}_{main} + \lambda_{dist} \mathcal{L}_{dist} + \mathcal{L}_{CrossSDC}. \quad (13)$$

3.5 Management of Memory Set

In alignment with the work of [76], our framework maintains a compact memory set throughout incremental updates. Each old class is limited to a maximum number of exemplars. After completing training for each task, we adopt the exemplar selection strategies in [2, 53] by randomly selecting exemplars for each current class and combining these new exemplars with the existing memory set.

4 Experiments

In this section, we first introduce the setup of our experiments, *i.e.*, dataset, baselines, evaluation metrics, and the implementation details. After that, we present the experimental results of our ContAV-Sep compared to the baselines, as well as ablation studies. We also conduct experiments on the AVE [68] and the VGGSound [13] datasets, which contain sound categories beyond the music domain. We put the experimental results on the AVE and the VGGSound datasets, the comparison to the uni-modal semantic similarity preservation method, the performance evaluation on old classes in incremental tasks, and the visualization of separating results in the Appendix.

4.1 Experimental Setup

Dataset. Following common practice [83, 88, 14], we conducted experiments on *MUSIC-21* [83], which contains solo videos of 21 instruments categories: accordion, acoustic guitar, cello, clarinet, erhu, flute, saxophone, trumpet, tuba, violin, xylophone, bagpipe, banjo, bassoon, congas, drum, electric, bass, guzheng, piano, pipa, and ukulele. In our experiments, we randomly selected 20 of them to construct the continual learning setting. Specifically, we split the selected 20 classes into 4 incremental tasks, each of which involves 5 classes. The total number of available videos is 1040, and we randomly split them into training, validation, and testing sets with 840, 100, and 100 videos, respectively. To further validate the efficacy of our method across a broader sound domain, we conduct experiments using the AVE [68] and the VGGSound [13] datasets in the appendix.

Baselines. We compare our proposed approach with vanilla Fine-tuning strategy, and continual learning methods EWC [32] and LwF [38]. As we mentioned before, typical continual learning methods, *e.g.*, class-incremental learning methods, which yield progressively increasing logits (or probability distribution) across all observed classes at each incremental step and design specific technique in the classifier, we consider that these methods are not an optimal choice for our proposed continual audio-visual sound separation problem. Thus, considering that continual semantic segmentation problem has a more similar form compared to conventional class-incremental learning, we also select two state-of-the-art continual semantic segmentation methods PLOP [16] and EWF [78] as our baselines. Moreover, we compare our method to the recently proposed audio-visual continual learning method AV-CIL [53], in which we adapt the original class-incremental version to the form of continual audio-visual sound separation by replacing their task-wise logits distillation with the output mask distillation. Further, we also present the experimental results of the Oracle/Upper Bound, which means that using the training data from all seen classes to train the model. **For fair comparison, all compared continual learning methods and our ContAV-Sep use the same state-of-the-art separator, *i.e.* iQuery [14], as the base separation model.** Further, we also incorporate our proposed and baseline methods into another representative audio-visual sound separation model Co-Separation [23]. Notably, the Co-Separation model does not utilize the motion modality. Therefore, when CrossSDC is applied to Co-Separation, the (mod_1, mod_2) in Eq. 7 and 9 is constrained to $(mod_1, mod_2) = (a, o)$. For baselines that involve memory sets, we ensure that each of them is allocated the same number of memory as our proposed method for fair comparison.

Implementation Details. Following [14], we use a 7-layers U-Net [58] as the audio net, and sub-sample the audio at 11kHz, each of which is approximately 6 seconds. We apply the STFT with

Table 1: Main results of different methods on MUSIC-21 dataset under the setting of Continual Audio-Visual Sound Separation with base separation models of iQuery [14] and Co-Separation [23], respectively. The bold part denotes the best results. Our proposed ContAV-Sep achieves the best performance among all baselines.

Method	SDR↑	SIR↑	SAR↑	Method	SDR↑	SIR↑	SAR↑
<i>w/o memory</i>				<i>w/o memory</i>			
iQuery [14] + Fine-tuning	3.46	9.30	10.57	Co-Sep. [23] + Fine-tuning	1.93	8.75	9.75
iQuery [14] + LwF [38]	3.45	8.78	10.66	Co-Sep. [23] + LwF [38]	2.32	7.84	10.28
iQuery [14] + EWC [36]	3.67	9.58	10.30	Co-Sep. [23] + EWC [36]	2.01	8.36	9.61
iQuery [14] + PLOP [16]	3.82	10.06	10.22	Co-Sep. [23] + PLOP [16]	3.24	9.17	9.59
iQuery [14] + EWF [78]	3.98	9.68	11.52	Co-Sep. [23] + EWF [78]	2.61	7.77	10.85
<i>w/ memory</i>				<i>w/ memory</i>			
iQuery [14] + LwF [38]	6.76	12.77	12.60	Co-Sep. [23] + LwF [38]	3.85	9.62	10.74
iQuery [14] + EWC [36]	6.65	13.01	11.73	Co-Sep. [23] + EWC [36]	3.31	9.55	9.80
iQuery [14] + PLOP [16]	7.03	13.30	11.90	Co-Sep. [23] + PLOP [16]	3.88	9.92	9.99
iQuery [14] + EWF [78]	5.35	11.35	11.81	Co-Sep. [23] + EWF [78]	3.63	9.07	10.58
iQuery [14] + AV-CIL [53]	6.86	13.13	12.31	Co-Sep. [23] + AV-CIL [53]	3.61	9.76	9.68
ContAV-Sep (with iQuery [14])	7.33	13.55	13.01	ContAV-Sep (with Co-Sep. [23])	4.06	10.06	11.07
Upper Bound (with iQuery)	10.36	16.64	14.68	Upper Bound (with Co-Sep.)	7.30	14.34	11.90

the Hann window size of 1022 and the hop length of 256, to obtain the 512×256 Time-Frequency representation of each audio signal, followed by a re-sampling on the log-frequency scale to generate the magnitude spectrogram with $T, F = 256$. We set the video frame rate (FPS) to 1, and detect the object using the pre-trained universal detector Detic [87] to detect the sound source object on each frame, and then, each detected object is resized and randomly cropped to the size of 224×224 . For the image encoder and the video encoder, we apply the self-supervised pre-trained CLIP [56] and VideoMAE [69] to yield the object feature and motion feature, respectively. For the audio-visual Transformer module, we follow the design in [14]. For all the baseline methods, we apply the same model architecture and modules with ours for them, including the mentioned Detic, CLIP, VideoMAE, audio-visual Transformer, etc. Please note that, during our training process, the pre-trained Detic, CLIP, and VideoMAE are frozen. In our proposed Cross-modal Similarity Distillation Constraint (CrossSDC), the balance weights λ_{ins} and λ_{cls} are set to 0.1 and 0.3, respectively. And the balance weight λ_{dist} for the output distillation loss is set to 0.3 in our experiments. For the memory set, we set the number of samples in each old class to 1, so as other baselines that involve the memory set. All the experiments in this paper are implemented by Pytorch [51]. We train our proposed method and all baselines on a NVIDIA RTX A5000 GPU. We follow previous works [67, 14] in sound separation, and evaluate the performance of all the methods using three common metrics in sound separation tasks: Signal to Distortion Ratio (SDR), Signal to Interference Ratio (SIR), and Signal to Artifact Ratio (SAR). The SDR measures the interference and artifacts, while SIR and SAR measure the interference and artifacts, respectively. In our experiments, we report the SDR, SIR, and SAR of all the methods after training at last incremental steps, *i.e.*, testing results on all classes. For all these three metrics, higher values denote better results.

4.2 Experimental Comparison

The main experimental comparisons are shown in Tab. 1. Our proposed method, ContAV-Sep, outperforms the state-of-the-art baselines by a substantial margin. Notably, compared to baselines using state-of-the-art audio-visual sound separator iQuery [14] as the separation base model, ContAV-Sep achieves a 0.3 improvement in SDR over the compared best-performing method. Additionally, our method surpasses the top baseline by 0.25 in SIR and 0.41 in SAR. Furthermore, compared to continual learning baselines with Co-Separation [23], our ContAV-Sep still outperforms other approaches. This consistent superior performance across different model architectures highlights not only the effectiveness but also the broad applicability and generalizability of our proposed CrossSDC.

Our observations further demonstrate that retaining a small memory set significantly enhances the performance of each baseline method. For instance, for the iQuery-based continual learning methods, equipping LwF [38] with a small memory set results in improvements of 3.31, 3.99, and 1.94 on SDR, SIR, and SAR, respectively. Similarly, the addition of a small memory set to EWC [32] leads to enhancements of 2.98, 3.43, and 1.43 in the respective metrics. The memory-augmented version of PLOP [16] exhibits superior performance with margins of 3.21, 3.24, and 1.68 for SDR, SIR, and SAR, respectively. Finally, incorporating memory into EWF [78] results in improvements of 1.37,

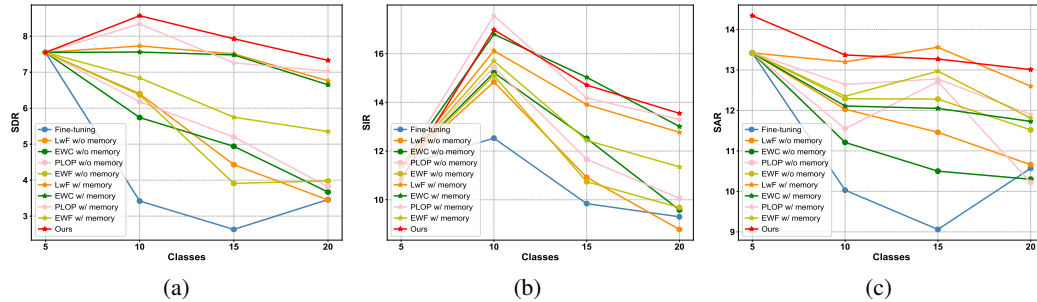


Figure 3: Testing results of different continual learning methods with iQuery [14] on the metrics of (a) SDR, (b) SIR, and (c) SAR at each incremental step.

1.67, and 0.29 for the three metrics. This phenomenon can be attributed to the inherent nature of the sound separation training process. In training, the audio signal from each sample mixes with others, giving a composite audio signal. This mixed audio signal, coupled with the corresponding visual data pair for each separated audio, constitutes the actual training sample for the separation task. As a result, even a single memory sample can be associated with multiple samples from the current training set, generating a diverse array of effective training pairs.

We also present the testing results of SDR, SIR, and SAR at each incremental step in Figures 3a, 3b, and 3c, respectively. Our method is consistently observed to outperform others in terms of SDR at all incremental steps. While our approach may not always produce the best SIR and SAR results at the intermediate steps (specifically, steps 2 and 3 for SIR, and step 3 for SAR), it ultimately achieves the highest performance at the final step. This demonstrates the robustness of our method, indicating minimal forgetting throughout the incremental learning process.

4.3 Ablation Study on CrossSDC and Memory Size

In this subsection, we conduct an ablation study to investigate the effectiveness of our proposed CrossSDC. By removing single or multiple components of the CrossSDC, we evaluate the impact of each on the final results. The results of the ablation study are presented in Tab. 2. From the table, we can see that our full model achieves the best performance compared to the variants, which further demonstrates the effectiveness of our proposed CrossSDC.

Moreover, we also discuss the effect of memory size on our proposed ContAV-Sep. In our main experiments, the default setting of the memory size is 1 sample per old class. In this subsection, we conduct experiments by increasing the memory size from 1 sample per old class to 30 samples per old class. The experimental results are shown in Tab. 3 and Figure 4. Observations from the table indicate a positive correlation between the size of the memory and the overall performance metrics. As the memory size increases, there is a discernible trend of improvement in the results.

Table 2: Ablation study on our proposed ContAV-Sep. Our full approach achieves best results compared to the variants.

	$\mathcal{L}_{dist.}$	$\mathcal{L}_{inst.}$	$\mathcal{L}_{cls.}$	SDR \uparrow	SIR \uparrow	SAR \uparrow
ContAV-Sep	✓	✗	✗	6.32	12.99	11.82
	✓	✓	✗	6.01	11.92	11.74
	✓	✗	✓	6.86	13.12	12.25
	✓	✓	✓	7.33	13.55	13.01

4.4 Limitation and Discussion

Our experimental findings reveal that the utilization of a small memory set, even a single sample per old class, markedly improves the performance of each baseline method. This improvement is attributed to the ability of a single memory sample to pair with diverse samples from the current training set, thereby generating numerous effective training pairs. Consequently, this process enables the model to acquire new knowledge for old classes in subsequent tasks, as the memory data can be

Table 3: Experimental results of our proposed ContAV-Sep with different memory size from 1 to 30 samples per memory class.

	# of samples per class	SDR \uparrow	SIR \uparrow	SAR \uparrow
ContAV-Sep	1	7.33	13.55	13.01
	2	7.26	13.10	12.65
	3	7.88	13.66	13.43
	4	8.16	14.16	13.21
	10	8.97	15.16	13.72
	20	9.39	15.93	13.69
	30	10.09	16.34	14.10




Figure 4: Testing results with different memory size (number of samples per class in the memory) on the metrics of (a) SDR, (b) SIR, and (c) SAR at each incremental step.

paired with data from previously unseen new classes — this is different from conventional continual learning tasks, where old classes do *not* acquire new knowledge in new tasks. This could be a potential reason why the baseline continual learning methods do not perform well in our continual audio-visual sound separation problem. In this work, our method also mainly focuses on preserving old knowledge of old tasks, which may prevent the model from acquiring new knowledge of old classes when training in new tasks. Recognizing this, we identify the exploration of this problem as a key avenue for future research in this field.

Additionally, the base model architectures used in our approach and baselines require object detectors to identify sounding objects. Although iQuery [14] can supplement object features with global video representations, it may still suffer from undetected objects. It is a fundamental limitation of the **object-based audio-visual sound separators** [23, 14]. While our work, unlike previous efforts, does not compete on designing a stronger audio-visual separation base model, enhancing the robustness of sounding object detection presents a promising direction for future research.

5 Conclusion

In this paper, we explore training audio-visual sound separation models under a more practical continual learning scenario, and introduce the task of continual audio-visual sound separation. To address this novel problem, we propose ContAV-Sep, which incorporates a Cross-modal Similarity Distillation Constraint to maintain cross-modal semantic similarity across incremental tasks while preserving previously learned semantic similarity knowledge. Experiments on the MUSIC-21 dataset demonstrate the effectiveness of our method in this new continual separation task. This paper opens a new direction for real-world audio-visual sound separation research.

Broader Impact. Our proposed continual audio-visual sound separation allows the model to adapt to new environments and sounds without full retraining, which could enhance efficiency and privacy by reducing the need to transmit and store sensitive audio data.

Acknowledgments. We thank the anonymous reviewers and area chair for their valuable suggestions and comments. This work was supported in part by a Cisco Faculty Research Award, an Amazon Research Award, and a research gift from Adobe. The article solely reflects the opinions and conclusions of its authors but not the funding agents.

References

- [1] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. The conversation: Deep audio-visual speech enhancement. *arXiv preprint arXiv:1804.04121*, 2018.
- [2] Hongjoon Ahn, Jihwan Kwak, Subin Lim, Hyeonsu Bang, Hyojun Kim, and Taesup Moon. SS-IL: separated softmax for incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 824–833, 2021.
- [3] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European conference on computer vision (ECCV)*, pages 139–154, 2018.
- [4] Elahe Arani, Fahad Sarfraz, and Bahram Zonooz. Learning fast, learning slow: A general continual learning method based on complementary learning system. *arXiv preprint arXiv:2201.12604*, 2022.
- [5] Eden Belouadah and Adrian Popescu. Il2m: Class incremental learning with dual memory. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 583–592, 2019.
- [6] Ruchi Bhatt, Pratibha Kumari, Dwarikanath Mahapatra, Abdulmotaleb El Saddik, and Mukesh Saini. Characterizing continual learning scenarios and strategies for audio analysis. *arXiv preprint arXiv:2407.00465*, 2024.
- [7] Albert S Bregman. Auditory scene analysis: Hearing in complex environments. *Thinking in Sounds*, pages 10–36, 1993.
- [8] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020.
- [9] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 233–248, 2018.
- [10] Sungmin Cha, Sungjun Cho, Dasol Hwang, Sunwon Hong, Moontae Lee, and Taesup Moon. Rebalancing batch normalization for exemplar-based class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20127–20136, 2023.
- [11] Moitreyra Chatterjee, Narendra Ahuja, and Anoop Cherian. Learning audio-visual dynamics using scene graphs for audio source separation. *Advances in Neural Information Processing Systems*, 35:16975–16988, 2022.
- [12] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019.
- [13] 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.
- [14] Jiaben Chen, Renrui Zhang, Dongze Lian, Jiaqi Yang, Ziyao Zeng, and Jianbo Shi. iquery: Instruments as queries for audio-visual sound separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14675–14686, 2023.
- [15] Soo-Whan Chung, Soyeon Choe, Joon Son Chung, and Hong-Goo Kang. Facefilter: Audio-visual speech separation using still images. *arXiv preprint arXiv:2005.07074*, 2020.
- [16] Arthur Douillard, Yifu Chen, Arnaud Dapogny, and Matthieu Cord. Plop: Learning without forgetting for continual semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4040–4050, 2021.

- [17] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 86–102. Springer, 2020.
- [18] 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. *arXiv preprint arXiv:1804.03619*, 2018.
- [19] Enrico Fini, Victor G Turrisi Da Costa, Xavier Alameda-Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-supervised models are continual learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2022.
- [20] Aviv Gabbay, Asaph Shamir, and Shmuel Peleg. Visual speech enhancement. *arXiv preprint arXiv:1711.08789*, 2017.
- [21] Chuang Gan, Deng Huang, Hang Zhao, Joshua B Tenenbaum, and Antonio Torralba. Music gesture for visual sound separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10478–10487, 2020.
- [22] Ruohan Gao, Rogerio Feris, and Kristen Grauman. Learning to separate object sounds by watching unlabeled video. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 35–53, 2018.
- [23] Ruohan Gao and Kristen Grauman. Co-separating sounds of visual objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3879–3888, 2019.
- [24] Siavash Golkar, Michael Kagan, and Kyunghyun Cho. Continual learning via neural pruning. *arXiv preprint arXiv:1903.04476*, 2019.
- [25] Simon Haykin and Zhe Chen. The cocktail party problem. *Neural computation*, 17(9):1875–1902, 2005.
- [26] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 831–839, 2019.
- [27] Chao Huang, Susan Liang, Yapeng Tian, Anurag Kumar, and Chenliang Xu. Davis: High-quality audio-visual separation with generative diffusion models. *arXiv preprint arXiv:2308.00122*, 2023.
- [28] Ching-Yi Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [29] Minsoo Kang, Jaeyoo Park, and Bohyung Han. Class-incremental learning by knowledge distillation with adaptive feature consolidation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16050–16059, 2022.
- [30] Zixuan Ke, Yijia Shao, Haowei Lin, Tatsuya Konishi, Gyuhak Kim, and Bing Liu. Continual learning of language models. In *International Conference on Learning Representations (ICLR)*, 2023.
- [31] Sanghwan Kim, Lorenzo Noci, Antonio Orvieto, and Thomas Hofmann. Achieving a better stability-plasticity trade-off via auxiliary networks in continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11930–11939, 2023.
- [32] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017.

- [33] Saksham Singh Kushwaha. Analyzing the effect of equal-angle spatial discretization on sound event localization and detection. In *Proceedings of the 7th Detection and Classification of Acoustic Scenes and Events 2022 Workshop (DCASE2022)*, 2022.
- [34] Saksham Singh Kushwaha and Magdalena Fuentes. A multimodal prototypical approach for unsupervised sound classification. *arXiv preprint arXiv:2306.12300*, 2023.
- [35] Jaewoo Lee, Jaehong Yoon, Wonjae Kim, Yunji Kim, and Sung Ju Hwang. Lifelong audio-video masked autoencoder with forget-robust localized alignments. *arXiv preprint arXiv:2310.08204*, 2023.
- [36] Kibok Lee, Kimin Lee, Jinwoo Shin, and Honglak Lee. Overcoming catastrophic forgetting with unlabeled data in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 312–321, 2019.
- [37] Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *International Conference on Machine Learning*, pages 3925–3934. PMLR, 2019.
- [38] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- [39] Yan-Shuo Liang and Wu-Jun Li. Adaptive plasticity improvement for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7816–7825, 2023.
- [40] Zilin Luo, Yaoyao Liu, Bernt Schiele, and Qianru Sun. Class-incremental exemplar compression for class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11371–11380, 2023.
- [41] Haoxin Ma, Jiangyan Yi, Jianhua Tao, Ye Bai, Zhengkun Tian, and Chenglong Wang. Continual learning for fake audio detection. *arXiv preprint arXiv:2104.07286*, 2021.
- [42] Divyam Madaan, Jaehong Yoon, Yuanchun Li, Yunxin Liu, and Sung Ju Hwang. Representational continuity for unsupervised continual learning. In *International Conference on Learning Representations (ICLR)*, 2022.
- [43] Fei Mi, Liangwei Chen, Mengjie Zhao, Minlie Huang, and Boi Faltings. Continual learning for natural language generation in task-oriented dialog systems. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, volume EMNLP 2020, pages 3461–3474, 2020.
- [44] Shentong Mo, Weiguo Pian, and Yapeng Tian. Class-incremental grouping network for continual audio-visual learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7788–7798, 2023.
- [45] Shentong Mo and Yapeng Tian. Audio-visual grouping network for sound localization from mixtures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10565–10574, 2023.
- [46] Xing Nie, Shixiong Xu, Xiyang Liu, Gaofeng Meng, Chunlei Huo, and Shiming Xiang. Bilateral memory consolidation for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16026–16035, 2023.
- [47] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *International conference on machine learning*, pages 2642–2651. PMLR, 2017.
- [48] Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jahnichen, and Moin Nabi. Learning to remember: A synaptic plasticity driven framework for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11321–11329, 2019.
- [49] 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.

- [50] Jaeyoo Park, Minsoo Kang, and Bohyung Han. Class-incremental learning for action recognition in videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13698–13707, 2021.
- [51] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [52] Quang Pham, Chenghao Liu, and Steven Hoi. Dualnet: Continual learning, fast and slow. *Advances in Neural Information Processing Systems*, 34:16131–16144, 2021.
- [53] Weiguo Pian, Shentong Mo, Yunhui Guo, and Yapeng Tian. Audio-visual class-incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7799–7811, 2023.
- [54] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 524–540. Springer, 2020.
- [55] Senthil Purushwalkam, Pedro Morgado, and Abhinav Gupta. The challenges of continuous self-supervised learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI*, pages 702–721. Springer, 2022.
- [56] 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.
- [57] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017.
- [58] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [59] Thomas Scialom, Tuhin Chakrabarty, and Smaranda Muresan. Fine-tuned language models are continual learners. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6107–6122, 2022.
- [60] Charles Spence. Audiovisual multisensory integration. *Acoustical science and technology*, 28(2):61–70, 2007.
- [61] Tejas Srinivasan, Ting-Yun Chang, Leticia Leonor Pinto Alva, Georgios Chochlakis, Mohammad Rostami, and Jesse Thomason. CLiMB: A continual learning benchmark for vision-and-language tasks. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022.
- [62] Barry E Stein and Terrence R Stanford. Multisensory integration: current issues from the perspective of the single neuron. *Nature reviews neuroscience*, 9(4):255–266, 2008.
- [63] Yiyang Su, Ali Vosoughi, Shijian Deng, Yapeng Tian, and Chenliang Xu. Separating invisible sounds toward universal audiovisual scene-aware sound separation. *arXiv preprint arXiv:2310.11713*, 2023.
- [64] Elyse S Sussman. Integration and segregation in auditory scene analysis. *The Journal of the Acoustical Society of America*, 117(3):1285–1298, 2005.
- [65] Reuben Tan, Arijit Ray, Andrea Burns, Bryan A Plummer, Justin Salamon, Oriol Nieto, Bryan Russell, and Kate Saenko. Language-guided audio-visual source separation via trimodal consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10575–10584, 2023.

- [66] Yu-Ming Tang, Yi-Xing Peng, and Wei-Shi Zheng. Learning to imagine: Diversify memory for incremental learning using unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9549–9558, 2022.
- [67] Yapeng Tian, Di Hu, and Chenliang Xu. Cyclic co-learning of sounding object visual grounding and sound separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2745–2754, 2021.
- [68] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. Audio-visual event localization in unconstrained videos. In *Proceedings of the European conference on computer vision (ECCV)*, pages 247–263, 2018.
- [69] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35:10078–10093, 2022.
- [70] Efthymios Tzinis, Scott Wisdom, Aren Jansen, Shawn Hershey, Tal Remez, Daniel PW Ellis, and John R Hershey. Into the wild with audioscope: Unsupervised audio-visual separation of on-screen sounds. *arXiv preprint arXiv:2011.01143*, 2020.
- [71] Efthymios Tzinis, Scott Wisdom, Tal Remez, and John R Hershey. Audioscopev2: Audio-visual attention architectures for calibrated open-domain on-screen sound separation. In *European Conference on Computer Vision*, pages 368–385. Springer, 2022.
- [72] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. FOSTER: feature boosting and compression for class-incremental learning. In *In Proceedings of the European Conference on Computer Vision (ECCV)*, pages 398–414, 2022.
- [73] Yu Wang, Nicholas J Bryan, Mark Cartwright, Juan Pablo Bello, and Justin Salamon. Few-shot continual learning for audio classification. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 321–325. IEEE, 2021.
- [74] Zhepei Wang, Cem Subakan, Xilin Jiang, Junkai Wu, Efthymios Tzinis, Mirco Ravanelli, and Paris Smaragdis. Learning representations for new sound classes with continual self-supervised learning. *IEEE Signal Processing Letters*, 29:2607–2611, 2022.
- [75] Zhepei Wang, Cem Subakan, Efthymios Tzinis, Paris Smaragdis, and Laurent Charlin. Continual learning of new sound classes using generative replay. In *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 308–312. IEEE, 2019.
- [76] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 139–149, 2022.
- [77] Adam Weisser. *Complex acoustic environments: Concepts, methods, and auditory perception*. PhD thesis, PhD thesis). Macquarie University, Sydney, Australia. doi: 1959.14/1266534, 2018.
- [78] Jia-Wen Xiao, Chang-Bin Zhang, Jiekang Feng, Xialei Liu, Joost van de Weijer, and Ming-Ming Cheng. Endpoints weight fusion for class incremental semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7204–7213, 2023.
- [79] Xudong Xu, Bo Dai, and Dahua Lin. Recursive visual sound separation using minus-plus net. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 882–891, 2019.
- [80] Shipeng Yan, Lanqing Hong, Hang Xu, Jianhua Han, Tinne Tuytelaars, Zhenguo Li, and Xuming He. Generative negative text replay for continual vision-language pretraining. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVI*, pages 22–38. Springer, 2022.
- [81] Yuxin Ye, Wenming Yang, and Yapeng Tian. Lavss: Location-guided audio-visual spatial audio separation. *arXiv preprint arXiv:2310.20446*, 2023.

- [82] Xiaohui Zhang, Jiangyan Yi, Jianhua Tao, Chenglong Wang, and Chu Yuan Zhang. Do you remember? overcoming catastrophic forgetting for fake audio detection. In *International Conference on Machine Learning*, pages 41819–41831. PMLR, 2023.
- [83] Hang Zhao, Chuang Gan, Wei-Chiu Ma, and Antonio Torralba. The sound of motions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1735–1744, 2019.
- [84] Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. The sound of pixels. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [85] Da-Wei Zhou, Qi-Wei Wang, Zhi-Hong Qi, Han-Jia Ye, De-Chuan Zhan, and Ziwei Liu. Deep class-incremental learning: A survey. *arXiv preprint arXiv:2302.03648*, 2023.
- [86] Da-Wei Zhou, Yuanhan Zhang, Jingyi Ning, Han-Jia Ye, De-Chuan Zhan, and Ziwei Liu. Learning without forgetting for vision-language models. *arXiv preprint arXiv:2305.19270*, 2023.
- [87] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *European Conference on Computer Vision*, pages 350–368. Springer, 2022.
- [88] Lingyu Zhu and Esa Rahtu. Visually guided sound source separation and localization using self-supervised motion representations. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1289–1299, 2022.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction clearly reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitations of the work in Sec. 4.4.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our paper does not contain theoretical results or proof.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We include an implementation paragraph section in Sec. 4.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Our source code has been released. Datasets used in our experiments are public dataset.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We include an implementation paragraph section in Sec. 4.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Because of the limitation of computation resource, we didn't report the error bars.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The paper provides sufficient information of the computer resources needed to reproduce the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Our paper does not violate any respect of the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The paper discusses the potential societal impacts. We include it in Sec. 5.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The creators or original owners of assets (e.g., code, data, models), used in the paper, are properly credited and are the license and terms of use explicitly mentioned and properly respected.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets are introduced in the paper

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.