

Sparse in Space and Time: Audio-visual Synchronisation with Trainable Selectors

Vladimir Iashin¹
vladimir.iashin@tuni.fi

Weidi Xie^{2,3}
weidi@robots.ox.ac.uk

Esa Rahtu¹
esa.rahtu@tuni.fi

Andrew Zisserman³
az@robots.ox.ac.uk

¹ Computing Sciences
Tampere University
Tampere, Finland

² Coop. Medianet Innovation Center
Shanghai Jiao Tong University
Shanghai, China

³ Visual Geometry Group
Department of Engineering Science
University of Oxford
Oxford, UK

Abstract

The objective of this paper is audio-visual synchronisation of general videos ‘in the wild’. For such videos, the events that may be harnessed for synchronisation cues may be spatially small and may occur only infrequently during a many seconds-long video clip, *i.e.* the synchronisation signal is ‘sparse in space and time’. This contrasts with the case of synchronising videos of talking heads, where audio-visual correspondence is dense in both time and space.

We make four contributions: *(i)* in order to handle longer temporal sequences required for sparse synchronisation signals, we design a multi-modal transformer model that employs ‘selectors’ to distil the long audio and visual streams into small sequences that are then used to predict the temporal offset between streams. *(ii)* We identify artefacts that can arise from the compression codecs used for audio and video and can be used by audio-visual models in training to artificially solve the synchronisation task. *(iii)* We curate a dataset with only sparse in time and space synchronisation signals; and *(iv)* the effectiveness of the proposed model is shown on both dense and sparse datasets quantitatively and qualitatively. Project page: v-iashin.github.io/SparseSync

1 Introduction

Audio-visual synchronisation is the task of determining the temporal offset between the audio (sound) and visual (image) streams in a video. In recent literature, this task has been explored by exploiting strong correlations between the audio and visual streams, *e.g.* in human speech [0, 8, 10] and playing instruments [3, 23], to provide a training signal for deep neural networks. In such scenarios, effective signals for synchronisation can be discovered between the lip or body movements and audio at almost every second. Despite the tremendous success achieved by these methods, for the most part, existing models are still limited to specialised domains, and not directly applicable to general (non-face, non-music) classes.




Figure 1: Audio-visual synchronisation requires a model to relate changes in the visual and audio streams. Open-domain videos often have a small visual indication, *i.e.* sparse in space. Moreover, cues may be intermittent and scattered, *i.e.* sparse across time, *e.g.* a lion only roars once during a video clip. This differs from a tight face crop of a speaker where cues are dense in space and time.

Our goal in this paper is to develop the next-generation audio-visual (AV) synchroniser. Rather than focusing on a specialised domain, such as human speech, we explore architectures for AV synchronisation for videos of general thematic content, *e.g.* daily videos [9, 20, 21] and live sports [13]. A solution for this task would be extremely useful for a number of applications that improve a user’s viewing experience – in order to avoid or at least automatically detect AV synchronisation offsets. Applications such as video conferencing, television broadcasts, and video editing, are currently largely done by ‘off-line’ measurements or heavy manual processing [11, 26, 28].

However, upgrading the existing audio-visual synchronisation systems to general videos is not straightforward, due to the following challenges: (i), in general videos, the synchronisation signal is often *sparse and instantaneous in time*, (a lion roaring or a tennis volley), rather than *dense in time* (a recorded monologue); (ii), objects that emit sounds can vary in size or appear in the distance making their presence on the frame small or *sparse in space* (a ball being hit in tennis), whereas synchronisation of a talking-head video may rely on visual cues from the localised mouth region, *i.e.* *dense in space*; (iii), some sound sources do not have a useful visual signal for synchronisation, *e.g.* stationary sounds (a car engine or electric trimmer), ambient sounds (wind, water, crowds, or traffic), and off-screen distractors (commentary track or advertisements); (iv), video encoding algorithms compress unperceived redundancy of a signal, this, however, can introduce artefacts that may lead to a trivial solution when training for audio-visual synchronisation; lastly, (v) due to its challenging nature, a public benchmark to measure progress has not yet been established.

In this paper, (i) we introduce a novel multi-modal transformer architecture, **SparseSelector**, that can digest long videos with linear scaling complexity with respect to the number of input tokens, and predict the temporal offset between the audio and visual streams. We achieve this by using a set of learnable *queries* to select informative signals from the ‘sparse’ video events across a wide time span. (ii) We show that for specific common audio and visual coding standards, a model can detect compression artefacts during training. We present a few simple indicators to determine if a model has learnt using these artefacts, as well as suggest several ways to mitigate the problem. Specifically, for the RGB stream, we recommend avoiding the MPEG-4 Part 2 codec, as well as reducing the sampling rate for audio. (iii) Additionally, to measure the progress of audio-visual synchronisation on general thematic content, we curate a subset of VGGSound with ‘sparse’ audio-visual correspondence called VGGSound-Sparse. We validate the effectiveness of the new model with thorough experiments on the existing lip reading benchmark (LRS3) and natural videos from VGGSound-Sparse and demonstrate state-of-the-art performance.

2 Related Work

Audio-visual synchronisation. During the pre-deep-learning era, the audio-visual human face synchronisation models relied on manually crafted features and statistical models [16, 27]. With the advent of deep learning, [9] introduced a two-stream architecture that was trained in a self-supervised manner using a binary contrastive loss. Later improvements were brought by multi-way contrastive training [10], and Dynamic Time Warping [24] used by [12]. Khosravan *et al.* [20] demonstrated the benefits of spatio-temporal attention and Kim *et al.* [21] employed a cross-modal embedding matrix to predict the offset for synchronisation. The progress was followed by [13] who introduced an architecture called VocaLiST with three transformer decoders [29]: two that cross-attend individual modalities and a third that fuses the outputs of the first two. These methods achieve impressive performance but focus on human speech rather than open-domain videos.

Although audio-visual synchronisation of general classes is a novel task, a few promising attempts have been made. In particular, Casanovas *et al.* [6] studied a handful of different scenes captured from a set of cameras. More recently, Chen *et al.* [7] adapted the transformer architecture and used a subset of VGGSound [8] covering 160 classes. In contrast to prior work, we focus on more challenging classes that have ‘sparse’ rather than ‘dense’ synchronisation signals.

Video coding artefacts. Since the early work of Doersch *et al.* on self-supervision [17], it has been known that network training can find shortcuts. Similarly, shortcuts due to video editing and coding artefacts have been noted in Wei *et al.* [30] and Arandjelović *et al.* [3]. In particular, [30] tackled the arrow-of-time in videos and studied artificial cues caused by black regions on video frames. While [3] noticed a slight impact of MPEG-encoding on audio-visual correspondence training and attributed it to the way negative samples are picked with respect to the start time of a positive sample. In this work, we study the ways to easily spot that the data contains artificial signals, as well as provide a few recommendations on how to prevent leaking such artefacts into data.

3 SparseSelector: an Audio-visual Synchronisation Model

In this section, we describe our audio-visual synchronisation model, where the audio-visual correspondence may only be available at sparse events in the ‘in the wild’ videos. This requires the model to handle longer video clips so that there is a high probability that a synchronisation event will occur. To this end, we propose *SparseSelector*, a transformer-based architecture that enables the processing of long videos with linear complexity with respect to the duration of a video clip. It achieves this by ‘compressing’ the audio and visual input tokens into two small sets of learnable *selectors*. These selectors form an input to a transformer which predicts the temporal offset between the audio and visual streams.

Architecture overview. The overview of the model is shown in Fig. 2. Given an audio spectrogram $\mathcal{A} \in \mathbb{R}^{H_a \times W_a \times 1}$ (H_a , W_a are frequency and time dimensions) and a stack of RGB frames $\mathcal{V} \in \mathbb{R}^{T_v \times H_v \times W_v \times 3}$, the audio-visual synchronisation model outputs the offset Δ between audio (\mathcal{A}) and visual (\mathcal{V}) streams:

$$\Delta = \Phi_{\text{Sync}}\left(\Phi_{\text{A-Sel}}\left(\Phi_{\text{A-Feat}}(\mathcal{A})\right), \Phi_{\text{V-Sel}}\left(\Phi_{\text{V-Feat}}(\mathcal{V})\right)\right). \quad (1)$$

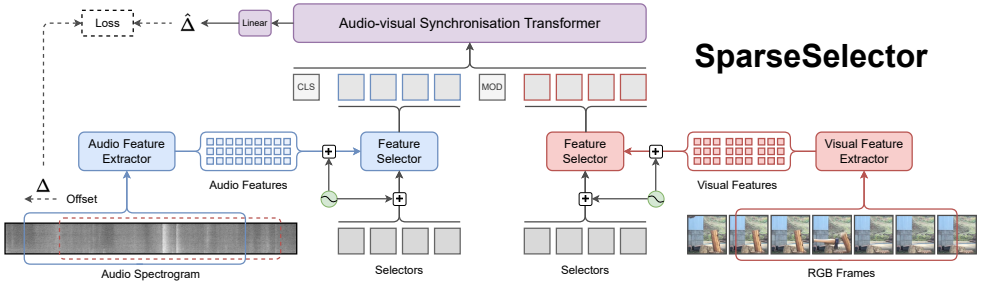


Figure 2: **An overview of SparseSelector.** The input is a spectrogram of the audio waveform and RGB frames from the video stream. These are passed through corresponding feature extractors, and the resulting features are refined with trainable selectors that ‘pick’ useful cues for synchronisation. As a result, the synchronisation transformer operates on substantially shorter sequences than the original input. The visual and audio selector queries are concatenated with classification (CLS) and separation tokens (MOD) as input to the transformer. Finally, the CLS token of the transformer output is used to predict the audio offset using a linear classification head. RGB frames are zoomed-in for visualisation purposes. The model is trained by off-setting the audio spectrogram. Dashed lines illustrate train-time behaviour.

First, audio and visual streams are independently encoded in feature extraction modules $\Phi_{A/V-Feat}$. Next, trainable *selectors* are passed to *Feature Selectors* ($\Phi_{A/V-Sel}$) along with the encoded features where they ‘summarise’ informative signals from the features that contain ‘sparse’ information across time and space. Finally, the selectors are used in *Synchronisation Transformer* (Φ_{Sync}) to predict the temporal offset Δ between audio and visual streams.

Feature encoding. Audio & visual inputs are encoded in spatio-temporal feature extractors:

$$a = \Phi_{A-Feat}(\mathcal{A}) \in \mathbb{R}^{h_a \times w_a \times d_a}, \quad v = \Phi_{V-Feat}(\mathcal{V}) \in \mathbb{R}^{t_v \times h_v \times w_v \times d_v}, \quad (2)$$

where t, h, w , and d denote time, height, width and channel dimensions, respectively. For the audio backbone, we use a variant of ResNet18 [18], which we pre-train on VGGSound [6] for sound classification. As for the visual backbone, we adopt S3D [61] pre-trained for action recognition on Kinetics 400 [19]. Although the setting allows employing any visual recognition network, we found that training a synchronisation model with a frame-wise feature extractor was significantly more difficult.

Feature Selectors. To utilize sparsely occurring synchronisation cues, the model should be able to handle longer input sequences. Moreover, accurate synchronisation requires a higher visual frame rate than other video understanding tasks (*e.g.* action recognition), which further increases the input size. For this reason, drawing on the idea of trainable queries [9, 17, 52], we propose to use a small number of trainable ‘selectors’ that learn to attend to the most useful modality features for synchronisation and, thus, reducing sequence length.

The architecture of the *Feature Selector* is similar to the transformer decoder [49]. Specifically, we start by flattening audio and visual features into sequences $a \in \mathbb{R}^{h_a w_a \times d_a}$ and $v \in \mathbb{R}^{t_v h_v w_v \times d_v}$. After adding trainable positional encoding (PE_*) for each dimension, trainable selectors and modality features are passed to the separate *Feature Selectors* as follows:

$$\hat{q}_a = \Phi_{A-Sel}(a + PE_a, q_a + PE_{q_a}), \quad \hat{q}_v = \Phi_{V-Sel}(v + PE_v, q_v + PE_{q_v}), \quad (3)$$

where $q_a, \hat{q}_a \in \mathbb{R}^{k_a \times d}$ and $q_v, \hat{q}_v \in \mathbb{R}^{k_v \times d}$ while $k_{a/v}$ are the numbers of selectors.

Note that the selectors provide a ‘short summary’ of the context features through the cross-attention mechanism while making the memory footprint more manageable. The reduced memory requirement is a consequence of (a) casting the complexity from quadratic to linear w.r.t. the input length, and (b) setting $k_v \ll t_v \cdot h_v \cdot w_v$ and $k_a \ll h_a \cdot w_a$. The number of selectors (k_v or k_a) can be conveniently tweaked according to the memory budget.

Audio-visual synchronisation transformer. To fuse the audio-visual cues from individual selectors, we adopt the standard transformer encoder layers to jointly process them, and to predict the offset, *i.e.* relative temporal shift between audio and visual streams, as follows:

$$\hat{\Delta} = \Phi_{\text{Sync}}([\text{CLS}; q_v; \text{MOD}; q_a]) \quad (4)$$

Here we concatenate the visual-audio selectors with two learnable special tokens, namely the classification token [CLS], and the modality token [MOD] that separates the two modalities. The offset prediction is obtained by applying a linear prediction head on the first token of the output sequence (omitted from Eq. (1) and (4) for clarity).

Training procedure. We assume that the majority of videos in the public datasets are synchronised to a good extent. With this assumption, we can artificially create temporal offsets between audio and visual streams from a video. We formulate the audio-visual synchronisation as a classification task onto a set of offsets from a pre-defined temporal grid space as $[-2.0, -1.8, \dots, 0.0, +0.2, \dots, +2.0]$ sec. The step size is motivated by the ± 0.2 sec human tolerance, where the ITU performed strictly controlled tests with expert viewers and found that the threshold for acceptability is -0.19sec to $+0.1\text{sec}$ [25]. To train the model, we employ the cross-entropy loss. For our experiments, we randomly trim a 5-sec segment out of 9 seconds such that both audio and visual streams are within the 9-second clip to make inputs of the same size and avoid padding that could hint if the input is off-sync.

4 Avoiding Temporal Artefacts

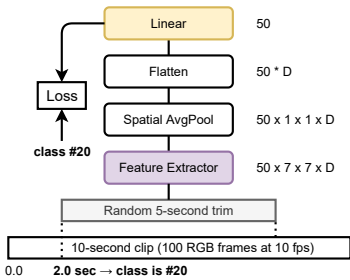
In this section, we describe our discovery of trivial solutions for training audio-visual synchronisation, that is, the model is able to exploit the video compression artefacts, to infer the time stamp for the specific video clip. Additionally, we also detail a suite of techniques that allows us to probe the artefacts and provide some practical suggestions to avoid them.

4.1 Identifying Temporal Artefact Leakage

We present two ways of identifying the temporal artefact leakage. In particular, training to predict the start time of a temporal crop (discussed next) and tracking metrics with temporal tolerance (discussed in the supplementary material).

Training to predict a video clip’s time stamp. A synchronisation model should rely solely on temporal positions of conceptual cues instead of, what we call, temporal artefacts. To check if data is polluted with artefacts, we suggest training a model to predict the start time of a random trim of an available video clip as shown in Tab. 1, left. Of course, it should not be possible to determine the start time of the trim in the original clip from the trim itself, and a network trained for this task should achieve only chance performance. However, for some audio and video codecs, the performance is far higher indicating artefact leakage.

The start-time classifier is a simple feature extractor (ResNet18). It is trained on three variations of the MJPEG-AoT dataset [60] obtained from the Vimeo streaming service: original ProRes videos, and ProRes videos transcoded into either MPEG-4 Pt. 2 (aka. mpeg4)



Codec	Acc@1	Acc@5
MPEG-4 Part 2 (mpeg4)	27.2	77.1
MPEG-4 Part 10 (H.264)	2.5	11.9
ProRes	2.7	13.4
AAC @ 44100Hz	86.7	100.0
AAC @ 22050Hz	23.0	74.3
AAC @ 16000Hz	6.3	19.3
Lossless @ 22050Hz	2.9	14.6

Table 1: **Commonly used coding standards may leak temporal artefacts – it is easy to test.** *Left:* a simple architecture trained to classify the start of a trim from a 10s clip to a pre-defined 0.1 sec-step grid (here for RGB). *Right:* Accuracy comparison for RGB and audio stream codecs predicting the start of an RGB or audio trim. Metrics are accuracy at 1 and 5 on 50 classes. Chance performance is 2 and 10 %. The higher accuracies indicate that an artefact is being used – see text for discussion.

or MPEG-4 Pt. 10 (aka. H.264). Note, frames in ProRes are compressed independently from others. If the visual stream of the video is encoded using `mpeg4`, the model trained to predict the start of the trim can do it significantly beyond a chance performance (Tab. 1, top-right).

Similarly, Advanced Audio Coding (AAC) might also leak temporal cues to the audio signal (Tab. 1, bottom-right). Since it is challenging to find a large set of videos with lossless audio compression, we used audio of randomly generated noise with a specified sampling rate and saved it to a disk losslessly (PCM) to obtain the performance with lossless compression. To obtain results on AAC, we transcoded these files to AAC with `ffmpeg`.

4.2 Preventing Temporal Artefact Leakage

Avoiding MPEG-4 Part 2 in favour of H.264. The algorithm that selects key-frames in MPEG-4 Part 2 (`mpeg4`) is less flexible than the one of H.264. In particular, `ffmpeg`, which is commonly used in practice, by default, encodes key frames every 12 frames. This means that each of the following 11 frames is merely a residual of the key frame and it is noticeable on the RGB stream (as we show in the supplementary). Such a temporal regularity can be picked up by a model and used to solve the task relying mostly on these artefacts. In contrast, each frame encoded by H.264 can reference up to 16 key-frames, which can be allocated more sparsely and their presence depends heavily on the scene rather than a rather strict interval as in MPEG-4 Pt. 2. This benefit is apparent when training a model to predict the start of a trim (see Tab. 1, top-right). A potential solution would be to avoid *inter-frame* codecs (`mpeg4` and H.264) in favor of an *intra-frame* codec (e.g. MJPEG, ProRes). However, this is a strong requirement for research datasets because it requires avoiding YouTube which stores videos compressed with inter-frame codecs (H.264 or VP9, according to view count).

Reducing audio sampling rate. There is a substantial difference in the model’s ability to predict the start of a trim depending on the sampling rate of the audio track (Fig. 1, bottom-right). While the reason behind the temporal artefacts in AAC is unknown, we recommend avoiding higher sampling rates. In our experiments, we rely on a 16kHz sampling rate as it provides a reasonable trade-off between audio quality and the start prediction performance. Ultimately one would want to have a dataset with lossless audio tracks yet, again, it is a strong requirement for a dataset as it is commonly used by YouTube.

5 Experiments

Dense in time dataset. The dataset is Lip Reading Sentences (LRS3) [10] which is obtained from TED talks for many speakers. We use two variations of the dataset. The first employs strict rectangular face crop coordinates that are extended to make a square (‘dense in time and space’). The second variation consists of full-frame videos without cropping (‘spatially sparse and temporally dense’). The raw videos are obtained from YouTube with RGB (25fps, H.264) and audio (16kHz, AAC) streams and referred to as ‘LRS3-H264’ and ‘LRS3-H264 (“No face crop”)’. We utilise the `pretrain` subset and split video ids into 8:1:1 parts for train, validation, and test sets. Only videos longer than 9 sec are used to unify it with the sparse dataset (discussed next). In total, we use $\sim 58k$ clips from $\sim 4.8k$ videos.

Sparse in time dataset. The dataset uses VGGSound [11] which consists of 10s clips collected from YouTube for 309 sound classes. A subset of ‘temporally sparse’ classes is selected using the following procedure: 5–15 videos are randomly picked from each of the 309 VGGSound classes, and manually annotated as to whether audio-visual cues are only sparsely available. After this procedure, 12 classes are selected ($\sim 4\%$) or 6.5k and 0.6k videos in the train and test sets, respectively (for class names see Fig. 3). Next, the second round of manual verification of a different subset of 20 videos from each class determines if it is feasible to align the sound based on the visual content. It is observed that $\sim 70\%$ of these video clips are synchronisable. We refer to this dataset as *VGGSound-Sparse*.

Baseline. Drawing on architectural details proposed in [10], we design a baseline as a transformer decoder that uses audio features as *queries* and visual features as context (*keys* and *values*) to predict the offset. The audio features are pooled across the spectrogram frequency dimension and trained from scratch. Apart from that, the feature extractors resemble ours.

Offset grid. We define the synchronisation task as classifying the offset on a 21-class grid ranging from -2 to $+2$ seconds with the 0.2-sec step size, as explained in Sec. 3. This can be regarded as a more challenging variant of the sync/off-sync task that prior work solves. We also experiment with a simpler setting with only 3 offset classes $[-1, 0, +1]$, that test if a model could predict if one track either lags, is in sync, or is ahead of the other one.

Metrics. Considering the human off-sync perception tolerance, in our experiments, we mainly report the Top-1 Accuracy with a ± 1 class of temporal tolerance (as described in Sec. 4.1). Note, that the training loss does not account for this tolerance. In the supplementary section, we additionally provide performance on accuracy without tolerance.

5.1 Results

Dense in time and space. Tab. 2 shows the comparison between the baseline and proposed architecture. As the task becomes more difficult (more sparse data or finer offset grid), we observe a larger gap between our model and the baseline. In particular, the baseline performs strongly on LRS3 (dense in time and space) in the setting with just three classes (98.4%). However, once the task gets more challenging, *i.e.* when training on finer offset shifts (21 class, 0.2 sec apart), the baseline performance deteriorates significantly ($\sim 89.8\%$). In contrast, the proposed model performs strongly even in the setting with finer offsets. This suggests that the proposed model is better suited for more challenging data tasks.

	Dense-Dense		Dense-Sparse		Sparse-Sparse	
	LRS3 (Face crop)		LRS3 (W/o face crop)		VGGSound-Sparse	
	3 cls	21 cls	3 cls	21 cls	3 cls*	21 cls
AVST _{dec}	98.4	89.8	95.8	83.1	52.2	29.3
Ours	96.4	95.6	95.5	96.9	60.3	44.3

Table 2: **The proposed model handles the increasing complexity of the setting and dataset better than the baseline while reaching a strong performance compared to the oracle.** ‘Dense-Dense’ refers to the face-cropped speech videos (LRS3), ‘Dense-Sparse’ for spatially-sparse LRS3 (‘No face crop’), ‘Sparse-Sparse’ is reported on VGGSound-Sparse which is sparse in time and space, *e.g.* lion roars once during a clip. The synchronisation performance is measured in two settings: the 3-class with $(-1, 0, +1)$ offsets given 5-sec clips, and 21 classes of offsets from -2.0 to $+2.0$ sec with 0.2-sec step size. The latter setting allows ± 1 temporal class tolerance (± 0.2 sec). *: oracle performance is 70 %.




Figure 3: **Performance per data and offset class on VGGSound-Sparse (test).**

Dense in time and sparse in space. A similar effect is observed on the dataset that is dense in time and sparse in space, *i.e.* LRS3 (No face crop). As both architectures drop their performance slightly on the 3-class setting after switching to a more difficult dataset, the drop is more significant for the baseline than for our model. Moreover, the baseline performance drops substantially in the 21-class setting ($>6\%$), while our model performs strong.

Sparse in time and space. Finally, the experiments on the VGGSound-Sparse reveal an even larger difference between the baseline method and our final model. For this experiment, we add additional data augmentation to mitigate overfitting. In particular, our model significantly outperforms the baseline showing the benefit of selectors on a more challenging dataset and setting. Ultimately, our model reaches 60 % in the 3-class setting which is close to the oracle performance ($\sim 70\%$: a human performance on 240 randomly picked videos), while achieving 44 % in the 21-class setting. We report performance per class in Fig. 3.

5.2 Ablation Study

In Tab. 3 we provide results for an ablation study. The results are reported on the VGGSound-Sparse dataset with 21 offset classes. More results are provided in the supplementary.

Feature selectors. The architecture with feature selectors outperforms the vanilla transformer showing the effectiveness of selectors in ‘compressing’ signals from the audio and visual features. Also, Fig. 4 shows a memory footprint comparison of the two, omitting the memory consumed by feature extractors that are the same. It is evident that memory demand grows rapidly with the input duration making it impossible to work with longer sequences and the transformer that inputs concatenated audio and visual streams.

Selectors	Sync Model pre-trained on LRS3	Pre-trained Feature Extractors	Unfrozen Feature Extractors	Accuracy ₂₁
✗	✓	✓	✓	40.1
✓	✗	✓	✓	12.1
✓	✓	✗	✓	29.6
✓	✓	✓	✗	33.5
✓	✓	✓	✓	44.3

Table 3: **Results of the ablation study.** The experiments are conducted on VGGSound-Sparse with 21 off-set classes. Metric is Accuracy with ± 1 class tolerance.




Figure 4: Working with longer sequences quickly becomes infeasible without selectors.

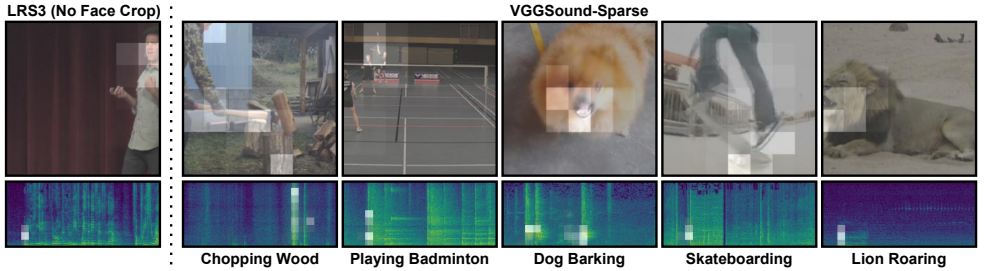


Figure 5: **Visual feature selectors focus on specific parts of the sparse signal that is useful for synchronisation.** Examples are from the hold-out set of LRS3 (‘No face crop’) and VGGSound-Sparse. Attention is captured from a selector to a visual or spectrogram feature token from a head within one of the layers. Attention values are min-max scaled.

Pre-training on dense signals. Pre-training a model on LRS3 (‘No face crop’) is an essential part of the training procedure on VGGSound-Sparse: 44.3 vs. 12.1 (near chance performance). For this reason, we also pre-train the baseline architecture (see Tab. 2).

Pre-trained feature extractors. The initialisation of audio and visual feature extractors with pre-trained weights has a strong positive effect on model performance. To initialise our feature extractors, we use weights of S3D pre-trained on Kinetics 400 for action recognition and ResNet18 pre-trained sound classification on VGGSound. The initialisation not only improves the final performance (43.3 vs. 33.5 %) but also significantly speeds up training. We attribute this improvement to the fact that such initialisation allows the model to ‘skip’ learning of the generic low-layer features and focus on training for synchronisation.

Frozen feature extractors during training. Allowing the gradients to reach raw data pixels is useful for audio-visual synchronisation as it makes the model sensitive to the smallest variations in the signal which is useful for synchronisation. In particular, having feature extractors to be trainable significantly boosts the performance from 34 to above 43 %, and the difference is even more pronounced on LRS3 (‘No face crop’) – see supplementary material.

Attention visualisation. Fig. 5 shows examples from LRS3-H.264 (‘No Face Crop’) and VGGSound-Sparse. Specifically, the attention exhibits spatial locality as the selectors learned to attend to the features extracted from the mouth region as expected from a model trained on a speech dataset. For a more challenging and diverse dataset, VGGSound-Sparse, the model highlights important parts of the visual and audio streams. In particular, the model accounts for the hit of the second badminton player who is far away in the background or attends to

Length (sec.)	VGGSound-Sparse	
	3 classes	21 classes
2	55.6	—
3	59.4	36.8
4	60.8	43.0
5	60.3	44.2
6	61.2	45.6
7	62.9	46.5

Table 4: Synchronisation accuracy improves with input length. We report results on two settings: with 3 offset classes ($-1, 0, +1$ sec) and 21 classes (± 2.0 sec grid with 0.2-sec step size). The results are reported on the test subset and accuracy is used as the metric. The accuracy for the 21-class setting is reported with ± 1 class tolerance. We use the same input lengths for pre-training, fine-tuning, and testing.

the axe during the chop, or the roaring mouth of the lion, yet these occur just once per video clip. Similarly, audio feature selectors point to specific parts of the spectrogram when the change occurs. More examples are provided in the supplementary material.

Input length. As the sparse synchronisation cues occur only occasionally within a video clip, processing shorter temporal crops decrease the chance of having sufficient cues for synchronisation, which, in turn, should decrease the performance. In Tab. 4, we show how performance varies with respect to the duration of input video clips. The results on the 3 and 21 offset classes illustrate the upward trend in model performance as the input duration extends. Note that the longer the inputs, the less unseen training data the model processes at each epoch. Specifically, a 10-second clip may be split into non-overlapping 3-second clips, which is not possible with clips longer than 5 seconds. Thus, this effect may undermine the current performance.

6 Conclusion

In this work, we study ‘in the wild’ videos that often have a synchronisation signal that is sparse in time. This requires a model to efficiently process longer input sequences as these synchronisation cues occur only rarely. To this end, we designed a transformer-based synchronisation model that has linear complexity with respect to the input length. This was made possible by using a small set of learnable ‘selectors’ that summarise long audio and visual features that are employed to solve the synchronisation task. To evaluate models in this challenging setup, we curate a dataset with only sparse events and train it on 5-second long clips. Finally, we discovered that compression artefacts caused by audio and video codecs might pose a threat to training for synchronisation, yet, as we show, these artefacts are easily identifiable and could be avoided to a certain extent.

Limitations. First, considering the complicated input-output relationship in the proposed model, it is challenging to determine which part of the input signal influences the output. Second, in this work, we considered signals that are ‘dense in time and space’, ‘dense in time but sparse in space’, and ‘sparse in time and space’. However, there is another interesting setting ‘sparse in time but dense in space’ yet it is not clear how to design such a dataset without making it ‘too artificial’. Third, despite showing strong performance on the proposed dataset, VGGSound-Sparse, there is still room for improvement.

Acknowledgements. Funding for this research was provided by the Academy of Finland projects 327910 and 324346, EPSRC Programme Grant VisualAI EP/T028572/1, and a Royal Society Research Professorship. We also acknowledge CSC – IT Center for Science, Finland, for computational resources.

References

- [1] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. LRS3-TED: a large-scale dataset for visual speech recognition. *arXiv preprint arXiv:1809.00496*, 2018.
- [2] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. Self-supervised learning of audio-visual objects from video. In *Proceedings of the European Conference on Computer Vision*, 2020.
- [3] Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *Proceedings of the European Conference on Computer Vision*, pages 435–451, 2018.
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision*, 2020.
- [5] Anna Llagostera Casanovas and Andrea Cavallaro. Audio-visual events for multi-camera synchronization. *Multimedia Tools and Applications*, 2015.
- [6] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. VGG-Sound: A large-scale audio-visual dataset. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, 2020.
- [7] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. Audio-visual synchronisation in the wild. In *Proceedings of the British Machine Vision Conference*, 2021.
- [8] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *Workshop on Multi-view Lip-reading, ACCV*, pages 251–263, 2016.
- [9] Joon Son Chung and Andrew Zisserman. Lip reading in the wild. In *Proceedings of the Asian Conference on Computer Vision*, pages 87–103, 2016.
- [10] Soo-Whan Chung, Joon Son Chung, and Hong-Goo Kang. Perfect match: Improved cross-modal embeddings for audio-visual synchronisation. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, 2019.
- [11] Vansh Dassani, Jon Bird, and Dave Cliff. Automated composition of picture-synched music soundtracks for movies. In *Proceedings of the European Conference on Computer Vision*, 2019.
- [12] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the International Conference on Computer Vision*, pages 1422–1430, 2015.
- [13] Joshua P. Ebeneze, Yongjun Wu, Hai Wei, Sriram Sethuraman, and Zongyi Liu. Detection of audio-video synchronization errors via event detection. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, 2021.
- [14] Tavi Halperin, Ariel Ephrat, and Shmuel Peleg. Dynamic temporal alignment of speech to lips. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, 2019.

- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [16] John Hershey and Javier Movellan. Audio vision: Using audio-visual synchrony to locate sounds. *Advances in Neural Information Processing Systems*, 1999.
- [17] Andrew Jaegle, Felix Gimeno, Andrew Brock, Andrew Zisserman, Oriol Vinyals, and Joao Carreira. Perceiver: General perception with iterative attention. *arXiv preprint arXiv:2103.03206*, 2021.
- [18] Venkatesh S Kadandale, Juan F Montesinos, and Gloria Haro. VocaLiST: An audio-visual synchronisation model for lips and voices. *arXiv preprint arXiv:2204.02090*, 2022.
- [19] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [20] Naji Khosravan, Shervin Ardeshtir, and Rohit Puri. On attention modules for audio-visual synchronization. In *Workshop on Sight and Sound, CVPR*, 2019.
- [21] You Jin Kim, Hee Soo Heo, Soo-Whan Chung, and Bong-Jin Lee. End-to-end lip synchronisation based on pattern classification. In *SLT Workshop*, 2021.
- [22] Bruno Korbar, Du Tran, and Lorenzo Torresani. Co-training of audio and video representations from self-supervised temporal synchronization. In *Advances in Neural Information Processing Systems*, 2018.
- [23] Andrew Owens and Alexei A. Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European Conference on Computer Vision*, 2018.
- [24] Lawrence Rabiner and Biing-Hwang Juang. *Fundamentals of speech recognition*. Prentice-Hall, Inc., 1993.
- [25] ITU Radiocommunication. Relative timing of sound and vision for broadcasting.
- [26] Prarthana Shrestha, Mauro Barbieri, Hans Weda, and Dragan Sekulovski. Synchronization of multiple camera videos using audio-visual features. *IEEE Transactions on Multimedia*, 2010. doi: 10.1109/TMM.2009.2036285.
- [27] Malcolm Slaney and Michele Covell. Facesync: A linear operator for measuring synchronization of video facial images and audio tracks. *Advances in Neural Information Processing Systems*, 2000.
- [28] Nicolas Staelens, Jonas De Meulenaere, Lizzy Bleumers, Glenn Van Wallendael, Jan De Cock, Koen Geeraert, Nick Vercammen, Wendy Van den Broeck, Brecht Vermeulen, Rik Van de Walle, et al. Assessing the importance of audio/video synchronization for simultaneous translation of video sequences. *Multimedia systems*, 2012.

- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.
- [30] Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8052–8060, 2018.
- [31] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proceedings of the European Conference on Computer Vision*, pages 305–321, 2018.
- [32] Mingze Xu, Yuanjun Xiong, Hao Chen, Xinyu Li, Wei Xia, Zhuowen Tu, and Stefano Soatto. Long short-term transformer for online action detection. In *Advances in Neural Information Processing Systems*, 2021.