arXiv:2204.02090v2 [cs.CV] 30 Jun 2022

VocaLiST: An Audio-Visual Synchronisation Model for Lips and Voices

Venkatesh S. Kadandale, Juan F. Montesinos, Gloria Haro

Dept. of Information and Communication Technologies, Universitat Pompeu Fabra, Spain

{venkatesh.kadandale, juanfelipe.montesinos, gloria.haro}@upf.edu

Abstract

In this paper, we address the problem of lip-voice synchronisation in videos containing human face and voice. Our approach is based on determining if the lips motion and the voice in a video are synchronised or not, depending on their audio-visual correspondence score. We propose an audio-visual cross-modal transformer-based model that outperforms several baseline models in the audio-visual synchronisation task on the standard lip-reading speech benchmark dataset LRS2. While the existing methods focus mainly on lip synchronisation in speech videos, we also consider the special case of the singing voice. The singing voice is a more challenging use case for synchronisation due to sustained vowel sounds. We also investigate the relevance of lip synchronisation models trained on speech datasets in the context of singing voice. Finally, we use the frozen visual features learned by our lip synchronisation model in the singing voice separation task to outperform a baseline audio-visual model which was trained end-to-end. The demos, source code, and the pre-trained models are available on https://ipcv.github.io/VocaLiST/

Index Terms: audio-visual, speech, singing voice, synchronisation, source separation, self-supervision, cross-modal

1. Introduction

Online media is brimming with user-generated videos involving human voice in the form of speech and singing voice. A large number of these videos suffer from misalignment between the audio and visual streams arising due to the lack of appropriate care during video editing. As a result of this misalignment, the video viewers notice that the lip motion is not perfectly synchronised with the voice in the audio. Lip-voice synchronisation in such videos could be corrected by compensating the offset between the audio and visual modalities, which is rendering the video out-of-sync. A common method to determine this offset is by training a model to tell lip-synchronised audio-visual (AV) pairs from the non-synchronised pairs and then choose the alignment between the audio and visual signals in the non-synchronised pairs that maximise the synchronisation score. We design and train a new model that can discriminate between synchronised and unsynchronised AV pairs by learning their mutual AV correspondence in a self-supervised fashion.

A system capable of detecting lip synchronisation issues in a video has many applications. Such a system can be deployed to automatically fix the synchronisation error in online user-generated videos that suffer from this problem. It can also be used to automatically evaluate the performance of automatic AV dubbing of movies from one language to another. The features learned by the model trained to detect synchronisation have already been shown to be useful for other AV applications like visual speech recognition [1], cross-modal retrieval [2], same language dubbing [3], dubbing [4], source separation [5, 6, 7, 8] and source localisation [5, 6].

We make several contributions through this paper. We propose a novel AV transformer-based model for detecting lip-voice synchronisation that estimates the extent of synchronisation between the lips motion and the voice in a given voice video. Our lip synchronisation detector outperforms state-of-the-art models on the speech benchmark dataset Lip Reading Sentences 2 (LRS2) [9]. Besides, we also train and test our model on the Acappella [10] dataset which contains videos of solo singing performances. We are the first to explore the relevance of lip-sync detectors trained in speech videos to detect lip-sync in the singing voice. We use the learned visual features in the lip synchronisation detection task to outperform a singing voice separation baseline model to showcase the practical utility value of our work. The source code and the pre-trained weights of our model will be provided for reproducibility.

The rest of the paper is organised as follows. We discuss the related work in the next section. In Sec. 3, we delineate our network architecture. Sec. 4 is dedicated to the experimentation, results, and discussion. Finally, we end the paper with a conclusion section and pointers for future work.

2. Related Work

The seminal work *Audio vision* [11] used AV synchronisation, measured as the mutual information between non-learned audio and visual features, to localise the active speaker in a video. More recently, with the advent of deep learning techniques, different models have addressed the synchronisation of AV signals in a self-supervised way, by automatically creating positive and negative AV sync/out-of-sync pairs. While some works focus on synchronisation in general sounds [6, 12, 13], others are specialised in speech signals [1, 3, 4, 14, 15]. The common trend is to extract features from each modality with an audio and a visual stream and then measure similarity/distance between the two embeddings using a sliding window to infer the offset of synchronisation. The first deep learning-based model for AV synchronisation in speech [14] uses a contrastive loss with a positive and a negative pair. The use of N negative pairs in a multi-way cross-entropy loss with softmax function [1] further improved the performance of the same model. Unlike these works, [15] directly train the model to determine the offset in the AV pairs.

Transformers have emerged as powerful deep learning architectures capable of capturing long-range dependencies in time series. Lately, transformers have been explored for several AV tasks such as source separation [16, 17], source localisation [18] and speech recognition [19], including synchronisation [13]. Our work in this paper is closest to Audio-Visual Synchronisation with Transformers (AVST) [13]. At the high-level, both the models share the same outline as shown in the left part of Fig. 1. However, the overall architecture of our model is different with regard to the choice of audio and visual encoders and the design of the synchronisation block. The details of our model architecture are clearly outlined in Sec. 3.1. Besides,

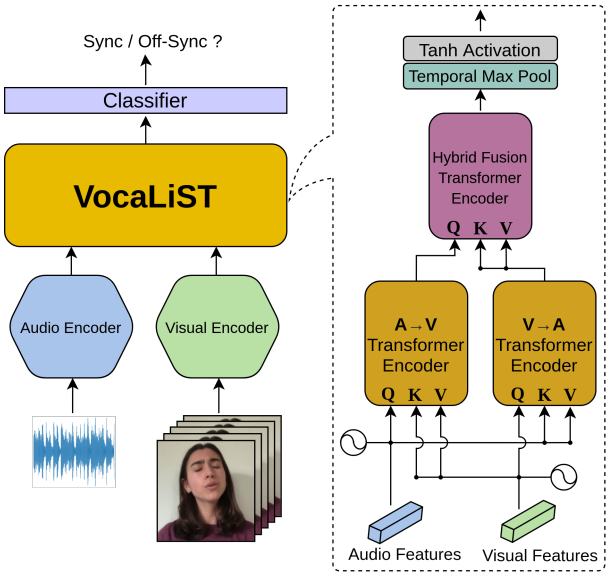

Figure 1: *Architecture of our lip synchronisation detector.*

[13] use the InfoNCE [20] loss for optimising their model, we use binary cross entropy loss. The main focus in [13] is AV synchronisation in general audio classes, while we focus on lip-synchronisation in speech and singing voice videos only. Finally, unlike in [13], we demonstrate a real-world application of the learned features of our synchronisation model.

3. Method

3.1. Architecture

The architecture of our model is shown in Fig. 1. We use a transformer-based classification model which uses the audio and visual features estimated by the audio and visual encoders to classify if a given face-voice pair is synchronised or not.

3.1.1. Audio Encoder

Our audio encoder (see Fig. 2) is very similar to the one used in the lip-sync expert discriminator [4]. The audio encoder is a stack of 2D convolutional layers with residual skip-connections that operate on the mel-spectrogram input of dimensions $1 \times 80 \times t_a$. The mel-spectrograms are obtained using 80 mel filterbanks with a hop size of 200 and window size of 800. The audios have a 16kHz sampling rate. The audio features are of the dimension $512 \times t_a$. The audio encoder conserves the temporal resolution of the input.

3.1.2. Visual Encoder

The visual encoder ingests a sequence of RGB images cropped around the mouth having dimensions $3 \times 48 \times 96 \times t_v$. Its architecture is inspired by the visual encoder of the lip-sync expert discriminator. Unlike in the latter, we apply 3D convolutions and conserve the temporal resolution in the feature maps. The output visual features are of dimension $512 \times t_v$. The conservation of temporal resolution in both the audio and visual features is helpful for learning the synchronisation patterns between the two modalities spread across the temporal dimension when we feed them into the synchronisation module. The visual frames are sampled from videos of 25 fps.

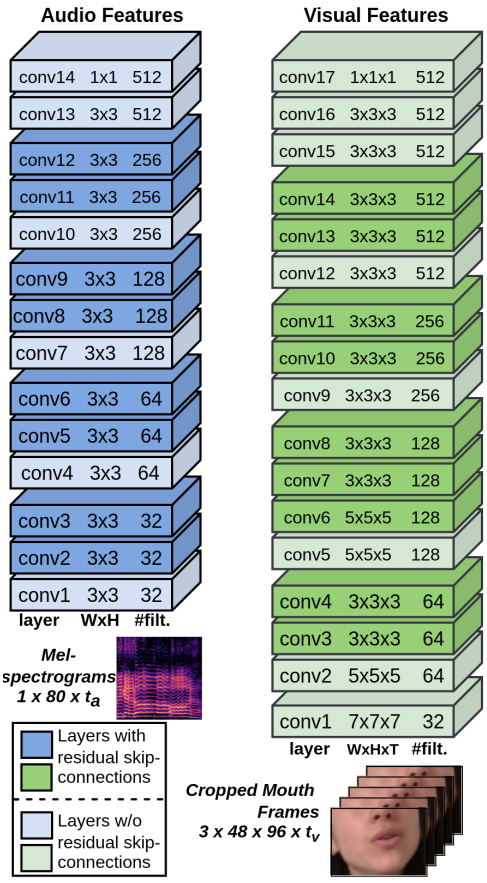

Figure 2: *Our audio encoder (left) and visual encoder (right).*

3.1.3. Synchronisation block

We design a powerful cross-modal AV transformer that can use the AV representations learned in its cross-modal attention modules to deduce the inherent AV correspondence in a synchronised voice and lips motion pair. We refer to our transformer model as VocaLiST, the **Voca**l **Li**p **S**ync **T**ransformer. Its design is inspired by the cross- modal transformer from [21]. The cross-modal attention blocks track correlations between signals across modalities.

The synchronisation block contains three cross-modal transformer encoders, each made up of 4 layers, 8 attention heads, and the hidden unit dimension of 512. The A→V unit takes in audio features as the query and the visual features as the key and values. The roles of these audio and visual features are swapped in the V→A unit. The output of the A→V unit forms the query to the hybrid fusion transformer unit, while its key and values are sourced from the output of the V→A unit. We max-pool the output of this hybrid fusion unit along the temporal dimension and pass it through *tanh* activation. Fig. 1 shows the architecture of VocaLiST.

Finally, there is a fully-connected layer acting as a classifier which outputs a score indicating if the voice and lips motion is synchronised or not. The whole architecture can handle any length of audio and visual inputs.

3.2. Training Setup

We train our model for estimating the AV correspondence score for a given AV pair in an end-to-end manner. The positive examples correspond to the synchronised pairs in which the audio corresponds to the visual. The negative examples are obtained by introducing random temporal misalignment between the in-sync AV pairs. This allows us to follow a self-supervised training pipeline. During the training, the positive and the negative examples are sampled equally. The AV correspondence score needs to be maximised for the positive examples and minimised for the negative examples. The binary cross-entropy loss criterion is used to optimise the model parameters during the training. Unless stated otherwise, we train our model with an AV input corresponding to a sequence length of 5 visual frames (0.2s) sampled at 25 fps. We use this training setup in 4.1 and 4.2.

4. Experiments

Here, we investigate lip synchronisation in two different settings: speech and singing voice. We also show the usefulness of the lip synchronisation models for a practical AV application.

4.1. Lip Synchronisation in Speech

Dataset. For the task of lip synchronisation in speech videos, we consider the LRS2 [9] dataset. The baseline methods SyncNet [14], Perfect Match (PM) [1] and Audio-Visual Synchronisation with Transformers (AVST) [13] also use the same dataset. This allows us to directly compare our method against these baselines. As in the baseline methods, we train our model on the 'pretrain' subset of LRS2 and evaluate on the 'test' subset.

Evaluation protocol. For a fair comparison, we mimic the evaluation protocol followed by the baseline methods. Given a sequence of cropped mouth frames and the mel-spectrograms, our model is tasked with predicting the correct synchronisation in the given pair of inputs. Human observers cannot tell unsynchronised speech videos from the synchronised ones when the temporal offset is in the ± 1 frame range. Hence, for a synchronised pair of speech audio and lips motion inputs, the synchronisation is estimated to be correct if the predicted offset between the pair is within $\pm$ 1 frame range. As in the baseline models, we estimate this offset by finding the index of maximum AV correspondence between the set of 5-frame visual features and all the audio feature sets (each temporally matching the length of 5 visual frames) lying in the ± 15 frame range.

Though we mainly train our model on AV input pairs corresponding to the length of 5 visual frames, the model can be tested on inputs of larger lengths. This is particularly useful when a set of 5 visual frames is non-informative (e.g. silence in speech). To evaluate on inputs of a larger context window, the AV correspondence score is obtained for each possible 5-frame set within the evaluation window (with a temporal stride of 1 frame) and then averaged before determining the offset. This averaging has been shown to improve the accuracy of the baseline models. On the other hand, [13] showed that training models for specific context window sizes (other than 5) leads to better performance than the previous averaging strategy in the same context window.

Results and discussion. Table 1 shows a direct comparison of the lip synchronisation accuracy of our model against the baseline models on the LRS2 test set. The accuracy is computed for multiple context windows ranging from sizes of 5 to 15. The accuracy improves across all the models as we increase the context window size while the rate of improvement reduces.

Table 1: *Accuracy of lip synchronisation models in LRS2*

Models	# params	Clip Length in frames (seconds)					
		5 (0.2s)	7 (0.28s)	9 (0.36s)	11 (0.44s)	13 (0.52s)	15 (0.6s)
SyncNet [14]	13.6M	75.8	82.3	87.6	91.8	94.5	96.1
PM [1]	13.6M	88.1	93.8	96.4	97.9	98.7	99.1
AVST [13]	42.4M	92.0	95.5	97.7	98.8	99.3	99.6
VocaLiST	80.1M	**92.8**	**96.7**	**98.4**	**99.3**	**99.6**	**99.8**

(M = million)

For each of the context window sizes, our model outperforms all the listed baseline models on the LRS2 test set. This improved performance can be attributed to the powerful cross-modal transformer blocks in our model in combination with our 3D-CNN based visual encoder. We tried reducing the complexity of our model by eliminating up to 2 cross-modal transformer blocks and also by replacing our visual encoder with the 18-layered Mixed Convolution network [22]. Both these attempts resulted in poorer performance.

4.2. Lip Synchronisation in Singing Voice

Dataset. We use Acappella [10], the only publicly available AV singing voice dataset. It consists of around 46-hours of solo-singing videos spanning four language categories. We test on the unseen-unheard test subset of the dataset which contains 83 song performances evenly distributed across each of the language categories and the genders. All the videos in the Acappella dataset have been manually curated from YouTube with care taken to ensure that the samples appear synchronised to human visual observation. On the other hand, in LRS2, an automatic method was used to better synchronise the AV signal pairs.

Models. We consider two models for the singing voice lip synchronisation: a baseline model and our VocaLiST. Instead of selecting SyncNet [14] as the baseline model, we choose the Lip Sync Expert Discriminator [4]. We refer to this new baseline model as SyncNet*. SyncNet* improves upon its predecessor SyncNet in the following ways: Unlike in SyncNet, i) SyncNet* operates on the RGB images, ii) the model is significantly deeper than the former with residual skip connections, and iii) the model is optimised using a cosine-similarity distance metric in combination with binary cross-entropy loss. We train both our VocaLiST and SyncNet* on the Acappella dataset's training set in an end-to-end manner.

Evaluation protocol. Unlike in speech videos, it is difficult for humans to notice out-of-sync in videos belonging to the general sound categories [13] when the offset between the modalities is less than 5. Singing voice also falls in such general sound category. The presence of sustained vowel sounds in singing voice, makes it difficult to notice synchronisation errors for the offsets less than 5. Thus, we consider synchronisation to be correctly estimated if the model predicts the maximum AV correspondence within the ± 5 frame range with respect to the ground truth. Unlike in [13], we do not decode the videos in lower frame rate for evaluation in singing voice.

Table 2: *Accuracy of lip synchronisation in Acappella dataset*

Models	Var	Trained on	Clip Length in frames (seconds)				
			5 (0.2s)	10 (0.4s)	15 (0.6s)	20 (0.8s)	25 (1s)
SyncNet*	N	Acappella	57.7	63.9	69.9	75.1	78.7
SyncNet*	Y	Acappella	57.7	65.9	—	—	73.6
VocaLiST	N	LRS2	56.7	65.1	72.2	77.2	81.2
VocaLiST	N	Acappella	58.8	65.4	71.6	76.5	80.5
VocaLiST	Y	Acappella	**58.8**	**66.4**	—	—	**85.2**

Results and discussion. We show the results in Table 2. Firstly,

we test the models trained on the speech dataset LRS2 directly on the singing voice samples. Among the baseline models discussed in Table 1, only the SyncNet model is publicly available. The Var column indicates if the results on larger context windows are obtained by the model trained with the exact length of the context window as the ones used for testing. Several observations arise from the results. First, the synchronisation accuracy in singing voice is lower than in the case of speech, even with a larger offset tolerance, showing that singing voice synchronisation is a harder task. Also, the increase of the context window size supposes a larger improvement, compared to the speech case. The best results are achieved with a dedicated network trained for 1s-length window. It can be noticed how our model trained for speech synchronisation generalises quite well for the singing voice case with large enough context window sizes. We hypothesise that with larger contexts it is more probable to find portions of singing voice excerpt with non-sustained vowels, or consonants, producing audio-visual cues more similar to the ones found in speech. The LRS2 'pretrain' subset that we use for training spans around 195 hours, while the training set of Acappella totals up to less than 37 hours of videos. Hence, compared to LRS2, Acappella is a small dataset and, therefore, the models trained on LRS2 tend to perform better in Acappella than the other way round.

4.3. Singing Voice Separation

We would like to demonstrate the effectiveness of the features learned by the visual encoder of our synchronisation network by using them in the singing voice separation task.

Dataset. We use the Acappella dataset [10], as in Sec. 4.2.

Training. We use the same training pipeline that is used for training the model Y-Net-mr in [10]. Y-Net-mr is a U-Net [23] conditioned by visual features extracted from cropped mouth frames using an 18-layer mixed convolution network [22]. Y-Net-mr estimates the complex masks corresponding to a voice present in an audio mixture when given with a spectrogram of an audio mixture and the temporal sequence of the cropped mouth frames corresponding to the target voice as input. During training, half of the input audio mixtures contain one voice mixed with a musical accompaniment, while the other half contains an additional voice besides the accompaniment.

Models. The main baseline model here is the Y-Net-mr from [10] which has been trained end-to-end for singing voice separation. It is a state-of-the-art audio-visual singing voice separation model among the models that operate directly on the cropped mouth frames. To investigate the contribution of the features learned by our lip synchronisation model, we propose Y-Net-mr-V. Y-Net-mr-V is nothing but a Y-Net-mr with its visual encoder replaced with that of our lip synchronisation model. In one setting, we train the Y-Net-mr-V for singing voice separation by loading the pre-trained weights of the visual encoder which was trained as a part of VocaLiST for singing voice lip synchronisation and keep them frozen throughout the training. In another setting, we train the entire model Y-Net-mr-V in an end-to-end manner without using the pre-trained weights learned from the lip synchronisation task. Finally, we also consider Y-Net-mr-S*, which is a Y-Net-mr with its visual encoder replaced by that of the SyncNet* [4]. For Y-Net-mr-S*, we only consider the setting where the pre-trained visual encoder weights corresponding to the singing voice lip synchronisation task are loaded from the SyncNet* and frozen as we train the rest of Y-Net-mr-S* for singing voice separation.

Evaluation. We evaluate the singing voice separation performance by computing the source separation metrics [24] Source-to-Distortion Ratio (SDR) and Source-to-Interference Ratio (SIR) on the estimated target voices in the unseen-unheard test subset of the Acappella dataset in the two-lead voice setup. In the two-lead voice setup, the music mixtures are artificially generated by randomly mixing two solo singing voices from the Acappella dataset in addition to another accompaniment audio. The higher these metrics are, the better the performance.

Table 3: *Performance metrics for Singing Voice Separation*

Architecture	Method	Source Separation Metrics	
		SDR	SIR
Y-Net-mr [10]	E2E	5.03	15.80
Y-Net-mr-V	E2E	1.14	11.72
Y-Net-mr-S*	PT - SyncNet*	5.44	16.17
Y-Net-mr-V	PT - VocaLiST	**6.32**	**17.08**

Results and Discussion. Table 3 shows the performance of the models in singing voice separation. Our model Y-Net-mr-V outperforms Y-Net-mr when we use the pre-trained weights from the lip synchronisation task for the visual encoder. It is challenging to train a source separation network that operates directly with video frames when the dataset is small, which is the case of Acappella. There is a tendency of over-fitting in such small datasets. Under such circumstances, knowledge transfer from the audio-visual synchronisation could guide the source separation despite the dataset size limitations. To further highlight the importance of this knowledge transfer, we also train the Y-Net-mr-V end to end without using the pre-trained synchronisation visual features. In this case, the model does not generalise well to the test-unseen subset despite having 38M trainable parameters solely in the visual encoder. Note that even Y-Net-mr-S* outperforms Y-Net-mr here. But since SyncNet* didn't perform as good as the VocaLisT in lip synchronisation task (see Table 2), as we expected, Y-Net-mr-S* did not outperform Y-Net-mr-V trained with knowledge transfer.

5. Conclusions

This paper presents VocaLiST, a transformer-based model for voice-lip synchronisation. The model has been analysed both in speech and singing voice, producing state-of-the-art results. We have shown that it learns powerful visual features that are useful for solving the problem of singing voice separation in a mixture with more than one voice. It could perform even better for larger input contexts if we have dedicated models trained for each specific length of the input, as shown in Table 2 and also in [13].

6. Acknowledgements

We would like to thank Honglie Chen (VGG, Oxford) and Soo-Whan Chung (Naver Corporation) for the insightful discussions. We acknowledge support by MICINN/FEDER UE project PID2021-127643NB-I00; H2020-MSCA-RISE-2017 project 777826 NoMADS. V. S. K. has received support through "la Caixa" Foundation (ID 100010434), fellowship code: LCF/BQ/DI18/11660064 and the Marie SkłodowskaCurie grant agreement No. 713673. J. F. M. acknowledges support by FPI scholarship PRE2018-083920.We gratefully acknowledge NVIDIA Corporation for the donation of GPUs used for the experiments.

7. References

[1] S.-W. Chung, J. S. Chung, and H.-G. Kang, "Perfect match: Improved cross-modal embeddings for audio-visual synchronisation," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 3965–3969.

[2] S. W. Chung, J. S. Chung, and H.-G. Kang, "Perfect match: Self-supervised embeddings for cross-modal retrieval," *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 3, pp. 568–576, 2020.

[3] Y. Shalev and L. Wolf, "End to end lip synchronization with a temporal autoencoder," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2020, pp. 341–350.

[4] K. Prajwal, R. Mukhopadhyay, V. P. Namboodiri, and C. Jawahar, "A lip sync expert is all you need for speech to lip generation in the wild," in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 484–492.

[5] T. Afouras, A. Owens, J. S. Chung, and A. Zisserman, "Self-supervised learning of audio-visual objects from video," in *European Conference on Computer Vision*. Springer, 2020, pp. 208–224.

[6] A. Owens and A. A. Efros, "Audio-visual scene analysis with self-supervised multisensory features," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 631–648.

[7] L. Zhu and E. 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*, 2022, pp. 1289–1299.

[8] Z. Pan, R. Tao, C. Xu, and H. Li, "Selective listening by synchronizing speech with lips," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2022.

[9] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, "Deep audio-visual speech recognition," *IEEE transactions on pattern analysis and machine intelligence*, 2018.

[10] J. F. Montesinos, V. S. Kadandale, and G. Haro, "A cappella: Audio-visual singing voice separation," in *32nd British Machine Vision Conference, BMVC*, 2021.

[11] J. Hershey and J. Movellan, "Audio vision: Using audio-visual synchrony to locate sounds," *Advances in neural information processing systems*, vol. 12, 1999.

[12] B. Korbar, D. Tran, and L. Torresani, "Cooperative learning of audio and video models from self-supervised synchronization," *Advances in Neural Information Processing Systems*, vol. 31, 2018.

[13] H. Chen, W. Xie, T. Afouras, A. Nagrani, A. Vedaldi, and A. Zisserman, "Audio-visual synchronisation in the wild," *arXiv preprint arXiv:2112.04432*, 2021.

[14] J. S. Chung and A. Zisserman, "Out of time: automated lip sync in the wild," in *Asian conference on computer vision*. Springer, 2016, pp. 251–263.

[15] Y. J. Kim, H. S. Heo, S.-W. Chung, and B.-J. Lee, "End-to-end lip synchronisation based on pattern classification," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 598–605.

[16] T.-D. Truong, C. N. Duong, H. A. Pham, B. Raj, N. Le, K. Luu *et al.*, "The right to talk: An audio-visual transformer approach," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 1105–1114.

[17] J. F. Montesinos, V. S. Kadandale, and G. Haro, "Vovit: Low latency graph-based audio-visual voice separation transformer," *arXiv preprint arXiv:2203.04099*, 2022.

[18] Y.-B. Lin and Y.-C. F. Wang, "Audiovisual transformer with instance attention for audio-visual event localization," in *Proceedings of the Asian Conference on Computer Vision*, 2020.

[19] P. Ma, S. Petridis, and M. Pantic, "End-to-end audio-visual speech recognition with conformers," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7613–7617.

[20] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.

[21] Y.-H. H. Tsai, S. Bai, P. P. Liang, J. Z. Kolter, L.-P. Morency, and R. Salakhutdinov, "Multimodal transformer for unaligned multimodal language sequences," in *Proceedings of the conference. Association for Computational Linguistics. Meeting*, vol. 2019. NIH Public Access, 2019, p. 6558.

[22] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri, "A closer look at spatiotemporal convolutions for action recognition," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018, pp. 6450–6459.

[23] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.

[24] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE trans. on audio, speech, and language process.*, vol. 14, no. 4, pp. 1462–1469, 2006.