

GestSync: Determining who is speaking without a talking head

Sindhu B Hegde
sindhu@robots.ox.ac.uk

Andrew Zisserman
az@robots.ox.ac.uk

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

Abstract

In this paper we introduce a new synchronisation task, *Gesture-Sync*: determining if a person’s gestures are correlated with their speech or not. In comparison to Lip-Sync, *Gesture-Sync* is far more challenging as there is a far looser relationship between the voice and body movement than there is between voice and lip motion. We introduce a dual-encoder model for this task, and compare a number of input representations including RGB frames, keypoint images, and keypoint vectors, assessing their performance and advantages. We show that the model can be trained using self-supervised learning alone, and evaluate its performance on the LRS3 dataset. Finally, we demonstrate applications of *Gesture-Sync* for audio-visual synchronisation, and in determining who is the speaker in a crowd, without seeing their faces. The code, datasets and pre-trained models can be found at: <https://www.robots.ox.ac.uk/~vgg/research/gestsync>.



Figure 1: Who is speaking in these scenes? Our model, dubbed *GestSync* learns to identify whether a person’s gestures and speech are “in-sync”. The learned embeddings from our model are used to determine “who is speaking” in the crowd, without looking at their faces. Please refer to the demo video for examples.

1 Introduction

Imagine you are watching a movie scene where multiple people are in a conversation, but the camera view does not include their faces; how can you determine who is speaking? That it is possible at all is due to the motion of the speaker’s hands and arms, that may be correlated with the speech, as illustrated in Figure 1.

If the face is visible, then the synchronisation of the lip motion and speech provides a very strong signal for who is speaking. This *lip-sync* problem has been well explored

and solved in prior work [4, 11, 12, 15, 17, 18, 19]. It is a strong signal because speech and lip motion are dense in time, and causally related. However, there could be various scenarios where facial features, especially lips, may not be accessible or useful. Examples include: (i) detecting if a person is speaking from a long distance where the lips are too small to decipher their movements; (ii) people speaking into a microphone that occludes their lips; (iii) scenes in movies where there are often lots of pose and scene changes along with occlusion of face by other actors/objects; and (iv) ego-centric views of the wearer’s hands from a wearable device. In all these cases, relying on traditional lip synchronisation signals will not be applicable. Thus, our objective in this paper is to determine to what extent it is possible to answer the question ‘who is speaking?’ *from gestures alone*. This is very challenging because the correlation between speech and gestures is much looser than that between speech and lip motion.

McNeill [24] categorizes co-speech gestures into four types: *beats*, *iconics*, *metaphorics*, and *deictics*. Beat gestures are usually two phase hand movements (up/down, left/right etc) that mark speech prosody, reinforce speech intonation, or control speech flow. They do not carry semantic content [5] but are used to emphasize particular words or phrases, and match the rhythm of the speech. Iconics physically represent the semantic content of speech, e.g. an arch-like hand movement accompanying the word “bridge”. Metaphoric gestures represent abstract ideas that have no material instantiation or have metaphorical meaning – e.g. an upward gesture accompanying “the stock price went through the roof”. Deictics are pointing gestures that indicate points or directions of interest. Note, these spontaneous gestures are very different from the structured gestures used in sign language [28] or typical gesture recognition datasets, e.g. [30], and furthermore they are very speaker dependent – some speakers make little or no gestures when they speak.

We introduce a dual-encoder model for this task, where the visual and audio streams are ingested by different encoders and then compared. However, rather than directly using RGB frames, we first investigate the use of 2D pose keypoints, such as points marking out the lips, elbows or shoulders. There are three reasons for this: first, we can effortlessly ‘switch-off’ information by not including a subset of keypoints. So, for example, if no facial keypoints are included, then the model cannot make use of lip-sync cues at all for the task. Second, keypoints remove many ‘nuisance’ parameters for the task, such as lighting and clothing, and are a much more compressed representation than images, so the training and exploration of models is easier. Finally, keypoints can be computed at frame rate, so models for the task can run in real time. We also investigate keypoints for the task of lip-sync as this enables performance to be compared with previous work on lip-sync – since there is no prior work on gesture-sync that we know of to compare with.

For training, we follow the example of lip-sync, where models are trained using self-supervised learning by introducing temporal offsets, and are then used for the task of active-speaker detection. We demonstrate in this paper that *gesture-sync* models can also be trained by a similar method of self-supervision. However, it is necessary to introduce longer temporal offsets, as hand motions correlated with speech are slower than lip motions, and so a longer temporal scale must be used.

In summary, we make the following four contributions: (i) we introduce a new task of audio-visual synchronisation using gestures; (ii) we propose a model for this task, and show that it can be trained using self-supervised learning; (iii) we investigate different representations of the video data including keypoints, RGB frames, and keypoints rendered as images; finally, (iv), we show very promising quantitative results for this task on the LRS3 dataset [4], and outline a number of applications.

2 Related Work

Lip synchronisation: Before deep learning, lip-speech synchronisation was predicted using hand-crafted features, e.g. [27]. In the deep learning era, SyncNet, a dual encoder with one ConvNet for the the visual stream and another for the audio stream, was trained using self-supervised learning [10]. The model architecture and training were then improved over a series of papers [11, 15, 17, 19]. All of these required a face detector. More recent work [18] has demonstrated that with the introduction of attention mechanisms, synchronisation for talking heads could be learnt at the image level and also the talking head could be tracked [4].

Synchronisation from sparse signals: While lip movements are densely correlated with audio in time, other audio-visual associations are more sparse. [9, 16] explore audio-visual synchronisation for arbitrary videos in the wild. For example, in a clip of a dog, the audio and video signals only provide a synchronisation signal when the dog barks and not at other time. Our work explores a task with a similar nature. The gestures match with the speech quite clearly at certain time-steps, but could be quite ambiguous at others. Long-term context becomes crucial here in order to confidently predict an audio-visual match.

Keypoint representations: While early methods relied on handcrafted features, recent approaches leverage deep learning architectures like CNNs to estimate accurate human body keypoints [8, 14, 22, 23, 26]. Several works have considered using keypoints as inputs for various tasks. One of the popular ways is to construct graph neural networks (GNNs) using keypoint representations. Tasks such as action recognition [6], gesture recognition [2], and sign language segmentation and recognition [7, 12, 25, 29] have been attempted using GNNs that take in keypoints as inputs. These works have shown impressive results, indicating the effectiveness of using keypoint representations. In our work, we take a slightly different path: directly using keypoint representations as inputs to a Transformer based model.

Speech and gestures: [13] generates plausible hand/arm gesture sequences from a speech signal. A GAN based network is employed to perform cross-modal translation from speech to hand and arm motion. The paper focuses on generating gestures by training on each speaker separately. In this paper, we explore a different task of learning synchronisation between speech and gestures while also training on a large number of identities in the wild.

3 Learning to Synchronise Gestures and Speech

Figure 2 shows an overview of the proposed *GestSync* network. The network takes in a video sequence, denoted as $I_{vid} = \{v_1, v_2, \dots, v_m\}$, and a speech sequence, denoted as $I_{speech} = \{s_1, s_2, \dots, s_n\}$. The video and audio streams are processed independently to obtain two distinct embeddings: visual embedding E_v and speech embedding E_s . Our goal is to determine the correlation between gestures and speech patterns, which we approach as a synchronisation task. To accomplish this, we compute the similarity between the visual and speech embeddings. Below, we provide a detailed description of the modules involved, followed by the training procedure in Section 3.3, and implementation details in Section 3.4.

3.1 Video Representations

We consider multiple input video representations: (i) **RGB frames:** Raw RGB frames are used as input; (ii) **Keypoint images:** The extracted keypoints are rendered onto an image

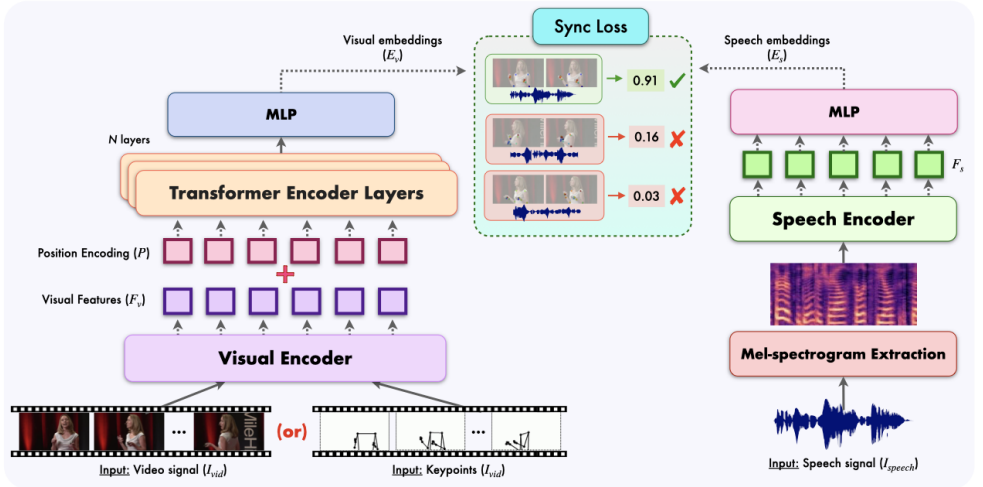


Figure 2: **The GestSync Network:** The model ingests one of the following video representations: (i) raw RGB frames, (ii) keypoint skeleton images, or (iii) keypoint vectors. The visual encoder processes the input signal and generates the visual features. These features are processed by the Transformer model to obtain the visual embeddings. For the speech input, the speech encoder uses Mel-spectrogram representation and generates the speech embeddings. The model is trained with a contrastive loss to have a high similarity between the visual and speech embeddings when synchronized, and low otherwise.

to create a skeleton graph using the keypoint (x, y) coordinates; (iii) **Keypoint vectors:** The keypoints are used directly as vectors. Figure 2 illustrates the RGB image representation and keypoint skeleton image representations.

Visual Encoder: We extract the visual features using a convolutional neural network (CNN). The input to the CNN is either the RGB frames or keypoint skeleton images or keypoint vectors. The aim (for both image and vector inputs) is to obtain the visual embeddings which we utilize further in computing the audio-visual similarity. For RGB frame and keypoint skeleton image inputs, we use 3D convolution layers, similar to previous audio-visual networks [10, 11], where the first layer has a temporal receptive field of 5 frames to capture the motion information. The network processes $(T \times 3 \times H \times W)$ -dimension frames (where T is the temporal window, and H and W are the frame height and width respectively), and generates visual features F_v as T vectors of dimension d . For keypoint vector inputs, we use 2D CNNs. The model takes in data of dimension $(T \times 2 \times N_{kp})$ (where N_{kp} is the number of keypoints), and these are treated as images of height \times width $(T \times N_{kp})$ with the (x, y) coordinate dimension acting as two channels. The output visual features F_v are T vectors of dimension $d \times N_{kp}$.

Transformer Encoder: The visual features F_v are encoded using a Transformer Encoder with N layers. The form of the Transformer depends on the input.

Transformer for image-based inputs: The Transformer attends to the temporal dimension in order to capture the temporal context. Sinusoidal positional encoding P for the frame number (time) are added to the visual features F_v . The output is T features of dimension d :

$$E_v = \text{TransformerEncoder}_{\text{img}}(F_v + P) \in \mathbb{R}^{T \times d} \quad (1)$$

Transformer for vector-based inputs: For keypoint vector based inputs, we design a Factorised attention transformer [6]. The goal here is to enable the Transformer to attend to the keypoints as well as to the temporal context. Learnable positional embeddings are used to label both keypoints and time, and are added to the visual features F_v . The Transformer has a stack of N Encoder layers, each comprising a feed-forward layer and two attention blocks: the first attention block attends across the time dimension and the second block attends across the keypoints. The Transformer outputs T features of dimension $N_{kp} \times d$:

$$E_v = \text{TransformerEncoder}_{\text{vec}}(F_v + P) \in \mathbb{R}^{T \times N_{kp} \times d} \quad (2)$$

Aggregation: We use a temporal aggregation layer which computes the mean for all the T -frames, followed by a linear layer. Thus, the output visual embedding E_v is a d -dimensional vector for image inputs, and N_{kp} d -dimensional vectors for keypoint vector inputs.

3.2 Speech Representation

For a speech segment I_{speech} , a mel-spectrogram representation is extracted using a window length of 25ms with a hop length of 10ms sampled at 16kHz. The mel-spectrograms (T' , 80) are given as input to the speech encoder which is similar to previous synchronisation networks [4]. The encoder is a stack of 2D convolutions with appropriate strides to match the visual time-steps T and outputs d dimension speech features F_s . We then use a linear layer to obtain speech embedding E_s .

3.3 Self-supervised Synchronisation Training

For each visual embedding E_v , we ask how similar is this embedding to the corresponding speech embedding E_s ? To answer this, we compute the cosine similarity between the visual and speech embeddings. Note that for the keypoint vector input, the visual embedding also has the spatial keypoint dimension, thus we consider the maximum over the spatial response as is done in AVObjects [4] to obtain the similarity.

The model is trained for the task of audio-visual synchronisation: the objective is to maximize the similarity between the visual input and speech input when they are synchronized, and to minimize the similarity between the shifted versions of the audio input. We sample K negative audio segments from the same video clip which are at least 1-second away from the positive window. This is illustrated in Figure 3. Thus, given the visual embedding E_v , synchronized speech embedding E_s , and K shifted speech embeddings $E_{s,1}, E_{s,2}, \dots, E_{s,K}$, we minimize the contrastive loss:

$$\mathcal{L} = -\log \frac{\exp(E_v \cdot E_s)}{\exp(E_v \cdot E_s) + \sum_{j=1}^K \exp(E_v \cdot E_{s,j})} \quad (3)$$

3.4 Implementation Details

Keypoint Extraction: To extract the keypoints, we utilize Mediapipe [23], which is a publicly available library built for human pose estimation. The major advantage of using Mediapipe over other libraries, such as OpenPose [8], is that it enables the keypoints to be




Figure 3: Illustration of self-supervised synchronisation training. The network is trained using contrastive loss by sampling K negative speech segments which are at least 1-second away from the positive segment.

extracted in real-time (on a CPU), thus reducing the computational burden. Mediapipe pose model generates 33 pose keypoints, out of which we discard the (sparse) face and lip keypoints and use the remaining 22 keypoints as inputs to the GestSync model. This ensures that the model learns to focus on the gestures rather than face or lip movements. We render the keypoint skeleton images using the 22 pose keypoint coordinates for the keypoint-image inputs, and use the normalized keypoint coordinates directly for the keypoint-vector inputs. For the task of lip synchronisation, we extract dense 128 face keypoints using the Mediapipe face model, and these form the input to the lip synchronisation models.

Hyperparameters: For training, videos are sampled at 25 FPS, at a resolution of 270×480 pixels. We extract 1-second audio and video segments ($T = 25$) which are considered as inputs. Our Transformer has 6 encoder layers, with 8 attention heads and generates 512-dimension embeddings. We train our network using the Adam optimizer [20] with a learning rate of $1e^{-4}$. The training is stopped once the validation loss does not improve for 5 epochs.

Sampling: The visual embeddings $E_{v,i}$ are computed for $T = 25$ frames, at a stride of 1 frame and the speech embeddings $E_{s,i}$ are computed for $T' = 100$ frames at a stride of 4 frames to obtain per-frame embeddings. The correlation between gestures and speech segments are not as fine-grained as that of lip synchronisation. Thus, for the task of gesture-sync: (i) we need longer context windows to enable the model learn well, and (ii) the negative speech segments need to be carefully sampled. It is possible that the gesture begins (ends) slightly before (after) the speech segment ends. To account for this, we ensure that the negatives are atleast 1-second away from the positive window.

Training procedure: For keypoint vector based inputs, the visual embeddings also have a spatial dimension (keypoint dimension). To make the training easier, we start by taking the mean of the embeddings and then switch to the maximum as explained in Section 3.3.

4 Experiments

We start by describing the datasets and metrics used for our tasks in Section 4.1. Since there are no existing works on our task of gesture synchronisation, we start our evaluation on the standard lip synchronisation task and compare our different representations against the state-of-the-art work in Section 4.2. We then move on to evaluate on the proposed gesture synchronisation task in Section 4.3.

4.1 Datasets and Metrics

Datasets: We conduct our experiments on the LRS3 dataset [2], which consists of TED videos downloaded from YouTube, totaling over 400 hours of video data. The dataset features a diverse set of speakers (over 9K) and a large vocabulary. We use the full-frame videos and resize them to (270×480) pixels. To ensure a fair evaluation, we create a new test split

from the LRS3 pre-train set, as the videos in the original LRS3 test set are short in duration. Our train and test sets do not have any speaker overlaps, and we have released our code, models, and data splits to encourage future research. For the task of lip-synchronisation, we report the results on LRS2 test set [9] to enable a direct comparison with the existing works.

Metrics: For evaluating lip synchronisation, we follow the same protocol as previous lip-sync methods [9, 10]. The video is shifted within ± 15 frames (31 negatives in total), and the synchronisation is considered correct if the predicted offset is within 1 video frame of the ground truth offset. We obtain per-frame scores by computing the cosine similarity of the visual and speech embeddings. The scores are averaged over F input frames (here $F = [5, 7, 9, 11, 13, 15]$). A random prediction would give an accuracy of 9.7% (3 out of 31 are correct). Lip movements rapidly change and are very fine-grained, meaning that there is a significant difference between consecutive frames. On the other hand, this is not the case for gestures. Thus, to evaluate gesture synchronisation, we shift the video by ± 50 frames with a gap of 5 frames between each negative (21 negatives in total) and average the per-frame score over F' input frames (here $F' = [25, 50, 75, 100]$). We deem the synchronisation to be correct if the predicted offset is within ± 10 frames of the ground truth. Note that since gestures and speech are coarsely correlated, we cannot expect the model to predict the finer-grained shifts (i.e. shifts less than ± 10 frames). Using this setup yields a chance accuracy of 23.8% (any of the closest 5 shifted versions out of the 21 is correct).

4.2 Lip Synchronisation

Comparison: We train our keypoint representation based models, both keypoint-image and keypoint-vector based methods for the task of lip synchronisation. We use the same settings as done in previous works, where the temporal window of 5 frames is considered for training and testing. We compare our models with three previous models: (i) SyncNet [10], (ii) Perfect Match [10], and (iii) AVObjects [9]. We start by re-implementing the existing RGB based AVObjects model, followed by our keypoint based models as explained in Section 3. Our keypoint based models use 128 face keypoints.

Results: Table 1 presents the lip synchronisation results for different number of frames used for averaging the scores on the LRS2 dataset [9]. Our RGB model’s scores validates the correctness of our experimental setup. We can observe that there is no significant difference between our keypoint based methods and the existing models. The fact that we are able to achieve an accuracy as high as $\approx 95\%$ using keypoint representations is remarkable and demonstrates the effectiveness of using the keypoint representations.

Table 1: Performance comparison of lip synchronisation accuracy (%) averaged over a given number of input frames on LRS2 dataset [9]. Our keypoint based methods achieve very similar results to the existing works, but with an advantage of being computationally efficient.

Method	5	7	9	11	13	15	Input type
SyncNet [10]	75.8	82.3	87.6	91.8	94.5	96.1	face crop
PerfectMatch [10]	88.1	93.8	96.4	97.9	98.7	99.1	face crop
AVObjects [9]	78.8	87.1	92.1	94.8	96.3	97.3	full image
RGB (AVObjects re-imp.)	80.2	88.1	92.0	95.5	96.1	97.5	full image
Ours: Keypoint-image	69.7	76.2	83.4	86.9	89.9	91.3	face points
Ours: Keypoint-vector	71.9	78.8	84.3	89.5	92.0	94.9	face points

4.3 Gesture Synchronisation

Baseline: As there are no existing works on our proposed task, we implement an RGB baseline model similar to AVObjects [4], but with a longer temporal window (25 frames) and an LSTM layer to aggregate the temporal information (Conv. + LSTM model). Note for this model and others that use RGB frames as input, the face is occluded by overlaying a mask so that the lips, and face/head movements cannot be used for the task.

Comparison: We present results on the Transformer network as explained in Section 3. We train the same models using keypoint-image representation. For the keypoint vector representation, we experiment with a baseline Conv. + LSTM model, followed by a Transformer model where we attend to the keypoint dimension, and finally show results on the Transformer with factorised attention. All our keypoint based models utilize 22 pose keypoints.

Results: Table 2 gives the results for the task of gesture synchronisation on the LRS3 dataset. The difference in scores between the lip synchronisation and gesture synchronisation validates that gesture-sync is a far more challenging problem. We can observe that for this task, the raw RGB based Transformer model performs the best. In all the three representations, the Transformer based models outperform the convolutional baselines. Unlike the task of lip-synchronisation, here there is a gap in performance between the keypoint based representations and the RGB representation. We speculate that one of the reasons for this difference could be that face or lips do not need any 3D information. However, this is not the case with gestures, as there could be occlusions present (for example, one arm in front of the other) and 3D motion (e.g. movements towards the camera). This information is completely lost when we consider 2D keypoint representations. Since the images naturally have this information available, their performance is superior when compared to the keypoints.

Table 2: Performance comparison of gesture synchronisation accuracy (%) averaged over a given number of input frames on the LRS3 dataset [4]. For the task of gesture-sync, RGB based model gives the best performance.

Input	Method	25	50	75	100
RGB	Conv. + LSTM baseline	49.8	59.3	68.2	75.7
	Ours (Image based Transformer)	53.7	66.1	72.6	77.5
Keypoint-image	Conv. + LSTM baseline	39.1	47.2	52.3	63.6
	Ours (Image based Transformer)	43.0	51.0	58.2	61.6
Keypoint-vector	Conv. + LSTM baseline	38.2	40.3	47.7	54.8
	Transformer (with temporal attn.)	40.6	43.5	49.9	56.5
	Ours (Transformer with Fact. attn.)	41.7	49.8	58.1	62.7

Visualization: Figure 4 depicts the visualization of the attention map from the keypoint-vector based model. The attention map indicates that the model is highly attentive to hands and arms amongst all the keypoints.

Computation Comparison: In Table 3, we compare the model parameters and the inference time for the best performing models in each of the input representations. For comparing the inference time, we report the time taken to load and generate embeddings for a 5-second input on a single NVIDIA V100 GPU. As we can see from the table, the keypoint-vector based model is very ef-

Table 3: Comparison of the model size (in million parameters) and the inference time (in seconds). The keypoint-vector based model has $\approx 1.4\times$ fewer parameters and is $\approx 6\times$ faster than the best performing RGB based Transformer model.

	RGB	Kp-img	Kp-vec
# params (M)↓	29.2	29.2	21.4
inf. time (sec)↓	12.12	10.81	2.20




Figure 4: Attention map visualization of the keypoint-vector Transformer model (with factorized attention). We can observe that the model focuses mostly on hands/arms to identify if the gestures are in-sync with speech. Note that the actual speaker frames are overlaid only for visualization.

ficient: it is $5.5\times$ and $4.9\times$ faster than raw RGB and keypoint skeleton based models respectively. Thus, in terms of compute, the keypoint-vector based model beats other models by a large margin. Note that we discard the time needed for keypoint extraction here for all the models, since Mediapipe runs in real-time on a CPU.

4.3.1 Bridging the gap between RGB and keypoint-vector representations

One of the major advantages of using the keypoint representations is the *computational efficiency*: keypoint-vector representations are $\approx 6\times$ faster than RGB representations. However, as observed in Table 2, this comes with a cost in the performance. Naturally a question arises is there a possibility to close or reduce the gap that exists between the RGB and keypoint representations? Below, we describe the techniques we experimented to achieve this feat.

In comparison to the RGB model which can *see* the entire frame – the arms of the person, not just their joints for example – the keypoint model has much less information. Thus, to provide more information, we gradually increase the number of keypoints between the joints, starting from 22 to extremely dense representations up to 142 keypoints. The results are presented in Table 4. We can observe that the 142 keypoints model achieved a significant boost in performance, reducing the gap between the RGB and keypoint representations from 14.4% to 7.9%. However, there is a trade-off in terms of computational complexity (calculated for a 5-second input on a single NVIDIA V100 GPU), as indicated in Table 4 (last column). Improving the performance of the keypoint-vector representation, while keeping the computational cost as low as possible is an interesting future direction to explore. Please refer to our Appendix for additional experiments and ablation studies.

Table 4: Performance comparison of gesture synchronisation accuracy (%) with different number of keypoints averaged over a given number of input frames (25, 50, 75, 100) on the LRS3 dataset [10]. The dense keypoint model with 142 keypoints gives the best performance among other keypoint models.

Input	# keypoints	25	50	75	100	inf. time (sec) ↓
Keypoint-vector	22	41.7	49.8	58.1	62.7	2.20
Keypoint-vector	48	44.7	52.7	60.2	67.0	4.32
Keypoint-vector	70	45.8	55.6	62.3	67.2	6.27
Keypoint-vector	142	46.1	57.5	65.6	68.9	12.10
RGB	-	53.7	66.1	72.6	77.5	12.12

5 Application: Who is speaking?

We investigate our model’s ability to identify the speaker when multiple individuals are present in the scene (as in Figure 1). The goal is to recognize the target speaker using our

trained model, without looking at their faces. There could be numerous scenarios where the face, or more precisely the lip movements of the speaker is corrupted, occluded, not covered by the camera angle or not present at all. In such cases, the traditional methods of spotting the speakers using lip-synchronisation fails, whereas our model can be used as an alternative.

Procedure and Evaluation: For this experiment, we create a pairwise test set by mixing the video of the target speaker with videos of other randomly selected speakers from the LRS3 dataset. Next, we calculate the similarity between the visual embedding of the target speaker and the speech embeddings of all the speakers to determine the likelihood of the speech belonging to the target speaker. Table 5 presents the model’s performance on the LRS3 dataset, where we report the accuracy of correctly identifying the target speaker among 1, 3, and 5 negative speakers. A random prediction for each of these settings would result in 50%, 25%, and 16% accuracy, respectively.

Results: We report the scores for the top-performing models (selected from Table 2) for each of the representations. The model’s performance is remarkably impressive, particularly in the case of 5 negative speakers, where our RGB-based model achieves an accuracy of 73.2%. This suggests a significant correlation between the speaker’s body movements and speech, which our model captures well. It also indicates the high quality of our learned visual and speech embeddings.

Table 5: Performance evaluation of identifying the target speaker in a crowd on the LRS3 dataset [1]. The accuracy of determining the positive speaker in the presence of: (i) 1-negative, (ii) 3-negative, and (iii) 5-negative speakers is computed.

Method	1-negative	3-negatives	5-negatives
Ours - RGB	87.0	77.6	73.2
Ours - Keypoint-image	82.4	70.6	62.6
Ours - Keypoint-vector	79.0	58.7	50.1

6 Summary, limitations, and future directions

In this work, we have presented a new task of gesture-speech synchronisation, which has not been explored before. We provided a detailed analysis on the different types of input representations, including raw RGB frames, keypoint skeleton images, and keypoint vectors. We demonstrated that it is indeed possible to identify “who is speaking” in a crowd using our model, by focusing only on the gestures without a need for the ‘talking head’ to be visible.

Compared to using lips to identify the speaker in a crowd, gestures suffer from the problem that some people do not gesture extensively. Also, a longer temporal context is required for gesture-sync compared to lip-sync.

There are several interesting directions to explore in the future, such as the correlation between gestures and language. In [1], initial results on predicting the native language of the speakers using gestures has been provided. Another direction is to explore if it is possible for some types of gestures (such as iconic or metaphoric) to ‘read’ the semantics of what is being said.

Acknowledgements. The authors would like to thank K R Prajwal, Tengda Han, Triantafyllos Afouras and Joon Son Chung for helpful discussions. This research is funded by EPSRC Programme Grant VisualAI EP/T028572/1, and a Royal Society Research Professorship RP/R1\191132.

References

- [1] Artem Abzaliev, Andrew Owens, and Rada Mihalcea. Towards understanding the relation between gestures and language. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5507–5520, 2022.
- [2] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. Lrs3-ted: a large-scale dataset for visual speech recognition. *arXiv preprint arXiv:1809.00496*, 2018.
- [3] Triantafyllos Afouras, Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Deep audio-visual speech recognition. *IEEE PAMI*, 2019.
- [4] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. Self-supervised learning of audio-visual objects from video. In *Proc. ECCV*, 2020.
- [5] Michael Andric and Steven L. Small. Gesture’s neural language. In *Front. Psychology*, 2012.
- [6] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, volume 2, page 4, 2021.
- [7] Hannah Bull, Michèle Gouiffès, and Annelies Braffort. Automatic segmentation of sign language into subtitle-units. In *ECCV Workshops*, 2020.
- [8] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [9] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. Audio-visual synchronisation in the wild. In *Proc. BMVC*, 2021.
- [10] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *Workshop on Multi-view Lip-reading, ACCV*, 2016.
- [11] Soo-Whan Chung, Joon Son Chung, and Hong-Goo Kang. Perfect match: Improved cross-modal embeddings for audio-visual synchronisation. In *Proc. ICASSP*, 2019.
- [12] Cleison Correia de Amorim, David Macêdo, and Cleber Zanchettin. Spatial-temporal graph convolutional networks for sign language recognition. In *Artificial Neural Networks and Machine Learning–ICANN*, pages 646–657. Springer, 2019.
- [13] Shiry Ginosar, Amir Bar, Gefen Kohavi, Caroline Chan, Andrew Owens, and Jitendra Malik. Learning individual styles of conversational gesture. In *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [14] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7297–7306, 2018.
- [15] Tavi Halperin, Ariel Ephrat, and Shmuel Peleg. Dynamic temporal alignment of speech to lips. In *Proc. ICASSP*, 2019.
- [16] Vladimir Iashin, Weidi Xie, Esa Rahtu, and Andrew Zisserman. Sparse in space and time: Audio-visual synchronisation with trainable selectors. In *Proc. BMVC*, 2022.

- [17] 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.
- [18] Naji Khosravan, Shervin Ardeshir, and Rohit Puri. On attention modules for audio-visual synchronization. In *Workshop on Sight and Sound, CVPR*, 2019.
- [19] 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.
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [21] Chuankun Li, Shuai Li, Yanbo Gao, Xiang Zhang, and Wanqing Li. A two-stream neural network for pose-based hand gesture recognition. *IEEE Transactions on Cognitive and Developmental Systems*, 14(4):1594–1603, 2021.
- [22] Marko Linna, Juho Kannala, and Esa Rahtu. Real-time human pose estimation from video with convolutional neural networks. *arXiv preprint arXiv:1609.07420*, 2016.
- [23] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*, 2019.
- [24] David McNeill. Hand and mind: What gestures reveal about thought. In *Bibliovault OAI Repository, the University of Chicago Press*, volume 27, 1994.
- [25] Lu Meng and Ronghui Li. An attention-enhanced multi-scale and dual sign language recognition network based on a graph convolution network. *Sensors*, 21(4):1120, 2021.
- [26] Tomas Pfister, James Charles, and Andrew Zisserman. Flowing convnets for human pose estimation in videos. In *Proc. ICCV*, 2015.
- [27] Malcolm Slaney and Michele Covell. Facesync: A linear operator for measuring synchronization of video facial images and audio tracks. *NeurIPS*, 2000.
- [28] Rachel Sutton-Spence and Bencie Woll. *The Linguistics of British Sign Language: An Introduction*. Cambridge University Press, 1999.
- [29] Manuel Vazquez-Enriquez, Jose L Alba-Castro, Laura Docío-Fernández, and Eduardo Rodriguez-Banga. Isolated sign language recognition with multi-scale spatial-temporal graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3462–3471, 2021.
- [30] Jun Wan, Stan Z. Li, Yibing Zhao, Shuai Zhou, Isabelle Guyon, and Sergio Escalera. ChaLearn looking at people RGB-D isolated and continuous datasets for gesture recognition. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 761–769, 2016.
- [31] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.