

# End-to-End Multi-View Lipreading

Stavros Petridis<sup>1</sup>  
stavros.petridis04@imperial.ac.uk

Yujiang Wang<sup>1</sup>  
yujiang.wang14@imperial.ac.uk

Zuwei Li<sup>1</sup>  
zuwei.li15@imperial.ac.uk

Maja Pantic<sup>1,2</sup>  
m.pantic@imperial.ac.uk

<sup>1</sup> iBUG Group  
Dept. Computing  
Imperial College London  
London, UK

<sup>2</sup> EEMCS  
University Of Twente  
Enschede, The Netherlands

---

## Abstract

Non-frontal lip views contain useful information which can be used to enhance the performance of frontal view lipreading. However, the vast majority of recent lipreading works, including the deep learning approaches which significantly outperform traditional approaches, have focused on frontal mouth images. As a consequence, research on joint learning of visual features and speech classification from multiple views is limited. In this work, we present an end-to-end multi-view lipreading system based on Bidirectional Long-Short Memory (BLSTM) networks. To the best of our knowledge, this is the first model which simultaneously learns to extract features directly from the pixels and performs visual speech classification from multiple views and also achieves state-of-the-art performance. The model consists of multiple identical streams, one for each view, which extract features directly from different poses of mouth images. The temporal dynamics in each stream/view are modelled by a BLSTM and the fusion of multiple streams/views takes place via another BLSTM. An absolute average improvement of 3% and 3.8% over the frontal view performance is reported on the OuluVS2 database when the best two (frontal and profile) and three views (frontal, profile, 45°) are combined, respectively. The best three-view model results in a 10.5% absolute improvement over the current multi-view state-of-the-art performance on OuluVS2, without using external databases for training, achieving a maximum classification accuracy of 96.9%.

## 1 Introduction

Lipreading, also known as visual speech recognition, is the process of recognising speech by observing only the lip movements without having access to the audio signal. Several approaches have been presented [10, 23, 28, 33] which extract features from a mouth region of interest (ROI) and attempt to model their dynamics in order to recognise speech. Such a system has the potential to enhance acoustic speech recognition in noisy environments since the visual signal is not corrupted by acoustic noise and can also enable the use of silent interfaces.

Recently, several deep learning approaches [0, 8, 24, 25, 26, 30] have been presented which replace the traditional feature extraction process and automatically extract features

from the pixels. There are also a few end-to-end approaches which attempt to jointly learn the extracted features and perform visual speech classification [6, 9, 27, 31]. This has given rise to a second generation of lipreading systems based on deep learning which significantly outperform the traditional approaches.

The vast majority of previous works has focused on frontal view lipreading. This is in contrast to evidence in the literature that human lip-readers prefer non-frontal views [6, 18], probably due to lip protrusion and lip rounding being more pronounced. It is therefore reasonable to assume that non-frontal lip views contain useful information which can be used to enhance the performance of frontal view lipreading or in cases where frontal view of the mouth ROI is not available. This is especially true in realistic in-the-wild scenarios where the face is rarely frontal. However, research on multi-view lipreading has been very limited and mainly restricted to evaluating non-frontal views independently with few works combining up to two views.

Such a system would be useful in meeting rooms where multiple cameras can record the participants simultaneously. A car environment is another potential application where multiple cameras can be easily installed and multiple views are easily available. Finally, some new smartphones have dual frontal cameras so that is another potential scenario although the views will not be too far apart.

In this work, we present an end-to-end model which jointly learns to extract features directly from the pixels and performs visual speech classification from multiple views. To the best of our knowledge, this is the first end-to-end model which performs multi-view lipreading and also achieves state-of-the-art performance. The proposed model consists of multiple identical streams, one per view, which extract features directly from the raw images. Each stream consists of an encoder which compresses the high dimensional input image to a low dimensional representation. The encoding layers in each stream are followed by a BLSTM which models the temporal dynamics. Finally, the information of the different streams/views is fused via a BLSTM which also provides a label for each input frame.

The second contribution of this work is an extensive comparison of all possible combinations up to five views. We evaluate the proposed model on the OuluVS2 database [1], which to the best of our knowledge is the only publicly available database containing 5 different views between  $0^\circ$  and  $90^\circ$ . We first perform an evaluation of each view independently and we see that the frontal and profile views are the best single views. The combination of these two views results in the best 2-view model which leads to a 3% absolute increase in classification accuracy over the frontal view. The addition of the  $45^\circ$  view results in a further small improvement of 0.8% over the best 2-view model. This means that multi-view lipreading can indeed enhance the performance of frontal view lipreading. However, the addition of more views (4 or 5) does not lead to any further improvements. We also show that non-frontal combinations like  $30^\circ$  or  $45^\circ$  and  $90^\circ$  outperform the frontal view which means that such combinations can be successfully used when frontal lip views are not available. Finally, in all single-view or multi-view scenarios the proposed model improves the state-of-the-art performance and achieves a maximum classification accuracy of 96.9% which is the highest performance reported on the OuluVS2 database.

## 2 Related Work

Previous works can be grouped into three categories: 1) evaluation of different views, 2) pose-invariant lipreading and 3) multi-view lipreading. The first group contains works which

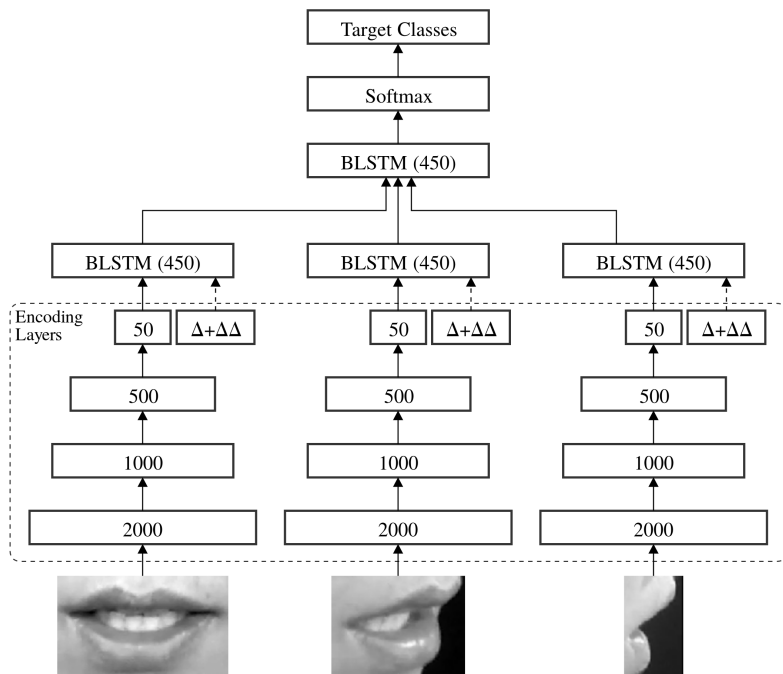


Figure 1: Overview of the end-to-end visual speech recognition system. One stream per view is used for feature extraction directly from the raw images. Each stream consists of an encoder which compresses the high dimensional input image to a low dimensional representation. The  $\Delta$  and  $\Delta\Delta$  features are also computed and appended to the bottleneck layer. The encoding layers in each stream are followed by a BLSTM which models the temporal dynamics. A BLSTM is used to fuse the information from all streams and provides a label for each input frame.

evaluate non-frontal views individually and compare their performance with the frontal view. It is not possible to draw a conclusion as to which is the most informative view since conflicting results are presented in the literature. Early works considered only frontal and profile views. Lucey and Potamianos [21] report a moderate performance degradation when profile view is used compared to frontal view. On the other hand, Kumar et al. [16] show results where the profile view outperforms the frontal view. Recently, few studies considering more views have been published. Lan et al. [17] evaluated 5 different views,  $0^\circ$ ,  $30^\circ$ ,  $45^\circ$ ,  $60^\circ$  and  $90^\circ$ , for lipreading and showed that their system performed best at  $30^\circ$ . The same 5 views have been evaluated on the OuluVS2 database and the results are still conflicting. The frontal view was found to be the best by Saitoh et al. [29], the profile view by Lee et al. [19], the  $30^\circ$  view by Zimmermann et al. [32] and the  $60^\circ$  view was found to be the best performing in [8]. The results presented by Saitoh et al. provide some insight as to why this may happen. Three different convolutional neural networks (CNNs), GoogLeNet, AlexNet and Network in Network, were trained on OuluVS2 using data augmentation. Each model led to different performance across the views. This is probably an indication that the best view depends on the model and maybe even the features used. This could explain the different conclusions

reached by different studies.

Pose-invariant lipreading works belong to the second group. The main goal of such works is to reduce the impact of different poses as it is known that the performance decreases in mismatched train/test conditions, i.e., the classifier is trained and tested on different poses. There are two main approaches in the literature for pose-invariant lipreading. The first one trains classifiers using data from all available views in order to build a generic classifier [22]. The second approach applies a mapping to transform features from non-frontal views to the frontal view. Lucey et al. [23] apply a linear mapping to transform profile view features to frontal view features. This approach has been extended to mapping other views like  $30^\circ$ ,  $45^\circ$  and  $60^\circ$  to the frontal view [24] or to the  $30^\circ$  view [25]. This approach makes it possible to collect a large amount of data on the optimal view and train a robust model. However, as the number of features to be generated by the linear mapping increases the performance is degraded [26].

To the best of our knowledge, there are only three works which have attempted multi-view lipreading, i.e., using multiple views of the lips simultaneously. Lucey and Potamianos [27] concatenated discrete cosine transform features extracted from frontal and profile views and fed them to a Hidden Markov Model (HMM). The performance of the 2-view system outperformed the frontal view model. Lee et al. [19] experimented with different CNN architectures which takes as input multiple lip views. The CNNs were combined with LSTMs and trained end-to-end on the OuluVS2 database. They just tested the combination of all 5 views which turned out to be worse than the frontal and profile views, possibly due to lack of enough training data. Finally, Zimmermann et al. [28] used principal component analysis (PCA) networks together with LSTMs and HMMs in order to combine multiple views from the OuluVS2 database. The only combination which outperformed the frontal view is frontal +  $30^\circ$ . Similarly to [19], the combination of all views led to worse performance than most individual views.

### 3 End-to-End Multi-view Lipreading

The proposed deep learning system for multi-view lipreading is shown in Fig. 1. It consists of three identical streams which extract features directly from the raw input images. Each stream corresponds to one view and consists of two parts: an encoder and a BLSTM. The encoder follows a bottleneck architecture in order to compress the high dimensional input image to a low dimensional representation at the bottleneck layer. The same architecture as in [24] is used, with 3 fully connected hidden layers of sizes 2000, 1000 and 500, respectively, followed by a linear bottleneck layer. The rectified linear unit is used as the activation function for the hidden layers. The  $\Delta$  (first derivatives) and  $\Delta\Delta$  (second derivatives) [29] features are also computed, based on the bottleneck features, and they are appended to the bottleneck layer. In this way, during training we force the encoding layers to learn compact representations which are discriminative for the task at hand but also produce discriminative  $\Delta$  and  $\Delta\Delta$  features. This is in contrast to the traditional approaches which pre-compute the  $\Delta$  and  $\Delta\Delta$  features at the input level and as a consequence there is no control over their discriminative power.

The second part is a BLSTM layer added on top of the encoding layers in order to model the temporal dynamics associated with each view. The BLSTM outputs of each stream are concatenated and fed to another BLSTM in order to fuse the information from all streams and model the temporal dynamics associated with all views. The output layer is a softmax

Table 1: Size of mouth ROIs in pixels for each view.

Views	0°	30°	45°	60°	90°
Height/Width	29/50	29/44	29/43	35/44	44/30

layer which provides a label for each input frame. The entire system is trained end-to-end which enables the joint learning of features and classifier. In other words, the encoding layers learn to extract features from raw images which are useful for classification using BLSTMs.

### 3.1 Single Stream Training

**Initialisation:** First, each stream/view is trained independently. The encoding layers are pre-trained in a greedy layer-wise manner using Restricted Boltzmann Machines (RBMs) [13]. Since the input (pixels) is real-valued and the hidden layers are either rectified linear or linear (bottleneck layer) four Gaussian RBMs [13] are used. Each RBM is trained for 20 epochs with a mini-batch size of 100 and L2 regularisation coefficient of 0.0002 using contrastive divergence. The learning rate is fixed to 0.001 as suggested in [13] when the visible/hidden units are linear.

**End-to-End Training:** Once the encoder has been pretrained then the BLSTM is added on top and its weights are initialised using glorot initialisation [12]. The Adam training algorithm is used for end-to-end training with a mini-batch size of 10 utterances. The default learning rate of 0.001 led to unstable training so it was reduced to 0.0003. Early stopping with a delay of 5 epochs was also used in order to avoid overfitting and gradient clipping was applied to the LSTM layers.

### 3.2 Multi-Stream Training

**Initialisation:** Once the single streams have been trained then they are used for initialising the corresponding streams in the multi-stream architecture. Then another BLSTM is added on top of all streams in order to fuse the single stream outputs. Its weights are initialised using glorot initialisation.

**End-to-End Training:** Finally, the entire network is trained jointly using Adam [13], 0.0001, to fine-tune the entire network. Early stopping and gradient clipping were also applied similarly to single stream training.

## 4 Experiments

The database used in this study is the OuluVS2 [4] which to the best of our knowledge is the only publicly available database with 5 lip views between 0° and 90°. It contains 52 speakers saying 10 utterances, 3 times each, so in total there are 156 examples per utterance. The utterances are the following: “Excuse me”, “Goodbye”, “Hello”, “How are you”, “Nice to meet you”, “See you”, “I am sorry”, “Thank you”, “Have a good time”, “You are welcome”. The mouth ROIs are provided and they are downscaled as shown in Table 1 in order to keep the aspect ratio of the original videos constant.

We first partition the data into training, validation and test sets. The protocol suggested by the creators of the OuluVS2 database is used [29] where 40 subjects are used for training

Table 2: Mean (standard deviation) single-view classification accuracy of 10 runs on the OuluVS2 database. The end-to-end model is trained using the established evaluation protocol where 40 subjects are used for training and validation and 12 for testing. † denotes that the difference with the frontal view mean classification accuracy is statistically significant.

Views	0°	30°	45°	60°	90°
End-to-end Encoder + BLSTM	91.8 (1.1)	87.3 <sup>†</sup> (1.6)	88.8 <sup>†</sup> (1.4)	86.4 <sup>†</sup> (0.6)	91.2 (1.3)

Table 3: Comparison of the maximum single-view classification accuracy over 10 runs with previous works on OuluVS2 database. All works follow the established evaluation protocol where 40 subjects are used for training and validation and 12 for testing. \* In cross-view training, the model is first trained with data from all views and then fine-tuned with data from the corresponding view. \*\* These models are pretrained on the BBC dataset [10], which is a large database, and then fine-tuned on OuluVS2. DA: Data Augmentation, LVM: Latent Variable Models

Views	0°	30°	45°	60°	90°
CNN + DA [29]	85.6	82.5	82.5	83.3	80.3
End-to-end CNN + LSTM [19]	81.1	80.0	76.9	69.2	82.2
CNN + LSTM, Cross-view Training * [19]	82.8	81.1	85.0	83.6	86.4
PCA Network + LSTM + GMM-HMM [32]	74.1	76.8	68.7	63.7	63.1
Baseline: Raw Pixels + LVM [9]	73.0	75.0	76.0	75.0	70.0
End-to-end Encoder + BLSTM (Ours)	94.7	89.7	90.6	87.5	93.1
CNN pretrained on BBC dataset + DA** [10]	93.2	-	-	-	-
CNN pretrained on BBC dataset + DA + LSTM** [8]	94.1	-	-	-	-

and validation and 12 for testing. We randomly divided the 40 subjects into 35 and 5 subjects for training and validation purposes, respectively. This means that there are 1050 training utterances, 150 validation utterances and 360 test utterances.

The target classes are a one-hot encoding for the 10 utterances. Each frame is labelled based on the label of the utterance and the end-to-end model is trained with these labels. The model provides a label for each frame and the majority label over each utterance is used in order to label the entire sequence.

Since all experiments are subject independent we first need to reduce the impact of subject dependent characteristics. This is done by subtracting the mean image, computed over the entire utterance, from each frame. The next step is the normalisation of data. As recommended in [13] the data should be z-normalised, i.e. the mean and standard deviation should be equal to 0 and 1 respectively, before training an RBM with linear input units. Hence, each image is z-normalised before pre-training the encoding layers.

Due to randomness in initialisation, every time a deep network is trained the results vary. In order to present a more objective evaluation we run each experiment 10 times and we report the mean and standard deviation of classification accuracy on the utterance level.

Table 4: Mean (standard deviation) 2-view classification accuracy over 10 runs of the proposed end-to end model on the OuluVS2 database. Due to lack of space we present results for the 6 best 2-view combinations. † denotes that the difference with the frontal view mean classification accuracy is statistically significant.

Views	0°+ 30°	0°+ 45°	0°+ 60°	0°+ 90°	30°+ 90°	45°+ 90°
End-to-end						
Encoder + BLSTM	91.7 (0.8)	93.6† (0.7)	92.0 (1.0)	94.8† (0.7)	93.8† (0.8)	93.6† (0.7)

Table 5: Comparison of the maximum 2-view classification accuracy over 10 runs with previous works on the OuluVS2 database. All approaches are trained using the established evaluation protocol where 40 subjects are used for training and validation and 12 for testing. Due to lack of space we present results for the 6 best 2-view combinations.

Views	0°+ 30°	0°+ 45°	0°+ 60°	0°+ 90°	30°+ 90°	45°+ 90°
PCA Network						
+ LSTM + HMM [32]	82.9	73.9	73.1	72.7	-	-
End-to-end						
Encoder + BLSTM	93.3	95.0	93.1	96.7	95	94.4

The proposed model was developed in Theano [10] using the Lasagne [11] library. The code and best models for the best view combinations are publicly available<sup>1</sup>.

## 4.1 Results on OuluVS2

### 4.1.1 Single-View Results

In the single view scenario, we train and test models on data recorded from a single view. This means that only a single stream is used from the model shown in Fig. 1. Table 2 shows the mean classification accuracy and standard deviation of the 10 models trained for each view. The best performance is achieved by the frontal and profile views followed by the 45°, 30° and 60° views.

In almost all previous works just a single accuracy value is provided, which is most likely the maximum performance achieved, with no standard deviation. In order to facilitate a fair comparison, we also provide the maximum performance achieved over the 10 runs in Table 3 together with previous results. It is obvious that even our mean performance is consistently higher for all views than previous works, which do not use external data for training. When it comes to maximum performance the proposed end-to-end architecture sets the new state-of-the-art for all views. It is worth pointing out, that the proposed system outperforms even the CNN models [7], [8] trained using external data. Both models are pre-trained on the BBC dataset [9], which is a large dataset, and fine-tuned on OuluVS2.

### 4.1.2 Multi-view Results

In the multi-view scenario, we train and test models on data recorded from multiple views. Table 4 shows the mean classification accuracy and standard deviation over 10 runs for com-

<sup>1</sup><https://ibug.doc.ic.ac.uk/resources/EndToEndLipreading/>

Table 6: 3-view classification accuracy over 10 runs on the OuluVS2 database of the proposed end-to-end model. The mean (standard deviation) and the maximum accuracy over 10 runs are presented. Due to lack of space we present results for the 4 best 3-view combinations. † denotes that the difference with the frontal view mean classification accuracy is statistically significant.

Views	$0^\circ + 45^\circ + 90^\circ$	$0^\circ + 30^\circ + 90^\circ$	$30^\circ + 45^\circ + 90^\circ$	$0^\circ + 60^\circ + 90^\circ$
Mean (st. dev.)	95.6 <sup>†</sup> (0.5)	95.2 <sup>†</sup> (0.5)	94.8 <sup>†</sup> (0.8)	94.8 <sup>†</sup> (1.0)
Max	96.9	96.1	95.8	96.4

binations of two views. Due to lack of space we present only the six best combinations and the other four are shown in the supplementary material. The combination of frontal and profile views is the best 2-view combination resulting in a 3% absolute increase in classification accuracy over the frontal view. The only other combinations that outperform the frontal one are the following:  $0^\circ + 45^\circ$ ,  $45^\circ + 90^\circ$ ,  $30^\circ + 90^\circ$ . It is beneficial to combine a frontal ( $0^\circ$ ) or partially frontal ( $30^\circ$  or even  $45^\circ$ ) view, which contains information mostly about the lips appearance, with a side ( $90^\circ$ ) or partially side ( $45^\circ$ ) which contains information mostly about lip protrusion. It is also evident that combining neighbouring views is not useful since there is not enough complementary information. Finally, it is worth pointing out that although a  $60^\circ$  view is a partially side view its combination with any other view is not beneficial. This is because it is the worst performing view as can be seen in Table 2.

Table 5 shows the maximum performance over the 10 runs for the 2-view combinations. To the best of our knowledge only Zimmerman et al. [34] present results on 2-view lipreading on OuluVS2. It is clear that the proposed end-to-end model significantly outperforms [34] and sets the new state-of-the-art maximum performance for all combinations of 2 views.

Table 6 shows the results of the four best 3-view combinations, the rest can be found in the supplementary material. All four combinations outperform the frontal view performance up to 3.8%, however only  $0^\circ + 45^\circ + 90^\circ$  results in a statistically significant improvement of 0.8% over the best 2-view combination (frontal + profile). Results for the combinations of 4 views and all views can be found in the supplementary material. All combinations outperform the frontal view up to 3.8%, however none of them results in a statistically significant improvement over the best 3-view combination. This reveals that 2-view lipreading leads to the biggest improvement over frontal view lipreading, and the addition of a third view offers a further small improvement but addition of more views does not really offer any performance benefit.

We should also point out that the maximum performance achieved of 96.9% is the highest classification accuracy reported on the short phrases part of OuluVS2, outperforming the previous state-of-the-art accuracy of 86.4% achieved by [19] (see Table 3) when no external data are used. It also outperforms the state-of-the-art performance of 94.1%, when external data are used for training, reported in [8] (see Table 5).

Fig. 2 shows the classification accuracy per subject for the best single-view, 2-view and 3-view combinations, respectively. We observe that for all subjects except 34 the combinations of frontal and profile views and frontal +  $45^\circ$  + profile results in similar or better performance than the frontal view. It is clear that there is not a very large performance deviation across the different test subjects. All of them except 6 achieve a classification accuracy over 90% with 5 of them (8, 9, 26, 30, 44) achieving 100%.






Figure 2: Mean classification accuracy and standard deviation per subject over 10 runs for the best performing views, frontal (blue), frontal + profile (red), frontal + 45° + profile (gray).




Figure 3: Confusion matrices for the best performing single-view, 2-view and 3-view models. The numbers correspond to the phrases in section 4. There are 12 subjects on the test set so there are 36 utterances per phrase.

Fig. 3 shows the confusion matrices for the best single-view, 2-view and 3-view combinations, respectively. It is clear that the number of confusions is reduced when multiple views are considered. The most common confusion pair in the single view is between “Hello” (3rd phrase) and “Thank you” (8th phrase) which is consistent with confusions presented in [19, 27]. This is followed by confusions between “See you” (6th phrase) and “Hello”, “Goodbye” (2nd phrase) and “You are welcome” (10th phrase), “You are welcome” and “How are you” (4th phrase).

Finally, we should also mention that we experimented with CNNs for the encoding layers but this led to worse performance than the proposed system. Chung and Zisserman [9] report that it was not possible to train a CNN on OuluVS2 without the use of external data. Similarly, Saitoh et al. [29] report that they were able to train CNNs on OuluVs2 only after data augmentation was used. This is likely due to the rather small training set. We also experimented with data augmentation which improved the performance but did not exceed the performance of the proposed system.

## 5 Conclusion

In this work, we present an end-to-end multiview lipreading system which jointly learns to extract features directly from the pixels and performs classification using BLSTM networks.

Results on the OuluVS2 demonstrate that the combination of multiple views is indeed beneficial for lipreading. The combination of the frontal and profile views leads to significant improvement over the frontal view. Further addition of the 45° view leads to a further small increase in performance. However, addition of more views does not lead to any further improvement. We have also demonstrated that combinations of non-frontal views, like 30° or 45° and 90°, can outperform the frontal view which is useful in cases where frontal lip views are not available. The proposed model achieves state-of-the-art performance on the OuluVS2 without using external data for training or even data augmentation. However, we should stress that the provided mouth ROIs are well cropped and this might not be the case when automatic tools for mouth ROI detection are used. It is well known that the accuracy of automatic detectors might degrade with non-frontal views so it would be interesting to investigate the impact of automatic mouth ROI cropping on multi-view lipreading. Finally, the model can be easily extended to multiple streams so we are planning to add an audio stream in order to evaluate its performance on audiovisual multi-view speech recognition.

## 6 Acknowledgements

This work has been funded by the European Community Horizon 2020 under grant agreement no. 645094 (SEWA).

## References

- [1] <http://deeplearning.net/software/theano/index.html/>.
- [2] <https://lasagne.readthedocs.io/en/latest/>.
- [3] <http://www.ee.oulu.fi/research/imag/OuluVS2/preliminary.html>.
- [4] I. Anina, Z. Zhou, G. Zhao, and M. Pietikäinen. Ouluvs2: A multi-view audiovisual database for non-rigid mouth motion analysis. In *IEEE FG*, pages 1–5, 2015.
- [5] Y. M. Assael, B. Shillingford, S. Whiteson, and N. de Freitas. Lipnet: Sentence-level lipreading. *arXiv preprint arXiv:1611.01599*, 2016.
- [6] S. L. Bauman and G. Hambrecht. Analysis of view angle used in speechreading training of sentences. *American journal of audiology*, 4(3):67–70, 1995.
- [7] J. S. Chung and A. Zisserman. Lip reading in the wild. In *Asian Conference on Computer Vision*, pages 87–103. Springer, 2016.
- [8] J. S. Chung and A. Zisserman. Out of time: automated lip sync in the wild. In *Workshop on Multiview Lipreading, Asian Conference on Computer Vision*, pages 251–263. Springer, 2016.
- [9] J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman. Lip reading sentences in the wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

- [10] S. Dupont and J. Luetin. Audio-visual speech modeling for continuous speech recognition. *IEEE Trans. on Multimedia*, 2(3):141–151, 2000.
- [11] V. Estellers and J. P. Thiran. Multipose audio-visual speech recognition. In *European Signal Processing Conference*, pages 1065–1069, 2011.
- [12] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Aistats*, volume 9, pages 249–256, 2010.
- [13] G. Hinton. A practical guide to training restricted boltzmann machines. In *Neural Networks: Tricks of the Trade*, pages 599–619. Springer, 2012.
- [14] G. Hinton and R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- [15] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] K. Kumar, T. Chen, and R. M. Stern. Profile view lip reading. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 4, pages 429–432, 2007.
- [17] Y. Lan, B. J. Theobald, and R. Harvey. View independent computer lip-reading. In *IEEE International Conference on Multimedia and Expo*, pages 432–437, 2012.
- [18] Y. Lan, B. J. Theobald, and R. Harvey. View independent computer lip-reading. In *2012 IEEE International Conference on Multimedia and Expo*, pages 432–437, 2012.
- [19] D. Lee, J. Lee, and K. E. Kim. Multi-view automatic lip-reading using neural network. In *Workshop on Multi-view Lip-reading Challenges*, pages 290–302. Asian Conference on Computer Vision, 2016.
- [20] P. Lucey and G. Potamianos. Lipreading using profile versus frontal views. In *IEEE Workshop on Multimedia Signal Processing*, pages 24–28, 2006.
- [21] P. Lucey, G. Potamianos, and S. Sridharan. An extended pose-invariant lipreading system. In *International Workshop on Auditory-Visual Speech Processing*, 2007.
- [22] P. Lucey, S. Sridharan, and D. B. Dean. Continuous pose-invariant lipreading. In *Interspeech*, pages 2679–2682, 2008.
- [23] I. Matthews, T. F. Cootes, A. Bangham, S. Cox, and R. Harvey. Extraction of visual features for lipreading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(2):198–213, 2002.
- [24] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y Ng. Multimodal deep learning. In *Proc. of ICML*, pages 689–696, 2011.
- [25] H. Ninomiya, N. Kitaoka, S. Tamura, Y. Iribe, and K. Takeda. Integration of deep bottleneck features for audio-visual speech recognition. In *Conf. of the International Speech Communication Association*, 2015.

- [26] S. Petridis and M. Pantic. Deep complementary bottleneck features for visual speech recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 2304–2308. IEEE, 2016.
- [27] S. Petridis, Z. Li, and M. Pantic. End-to-end visual speech recognition with lstms. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 2592–2596. IEEE, 2017.
- [28] G. Potamianos, C. Neti, G. Gravier, A. Garg, and A. W. Senior. Recent advances in the automatic recognition of audiovisual speech. *Proceedings of the IEEE*, 91(9):1306–1326, 2003.
- [29] T. Saitoh, Z. Zhou, G. Zhao, and Matti Pietikäinen. Concatenated frame image based CNN for visual speech recognition. In *Workshop on Multi-view Lip-reading Challenges*, pages 277–289. Asian Conference on Computer Vision, 2016.
- [30] C. Sui, R. Togneri, and M. Bennamoun. Extracting deep bottleneck features for visual speech recognition. In *IEEE ICASSP*, pages 1518–1522, 2015.
- [31] M. Wand, J. Koutn, and J. Schmidhuber. Lipreading with long short-term memory. In *IEEE ICASSP*, pages 6115–6119, 2016.
- [32] S. Young, G. Evermann, M. Gales, T. Hain, D. Kershaw, X. Liu, G. Moore, J. Odell, D. Ollason, D. Povey, et al. The HTK book. 3:175, 2002.
- [33] G. Zhao, M. Barnard, and M. Pietikainen. Lipreading with local spatiotemporal descriptors. *IEEE Transactions on Multimedia*, 11(7):1254–1265, 2009.
- [34] M. Zimmermann, M. M. Ghazi, H. K. Ekenel, and J. P. Thiran. Visual speech recognition using PCA networks and LSTMs in a tandem GMM-HMM system. In *Asian Conference on Computer Vision*, pages 264–276. Springer, 2016.