

A Multimodal German Dataset for Automatic Lip Reading Systems and Transfer Learning

Gerald Schwiebert, Cornelius Weber, Leyuan Qu, Henrique Siqueira, Stefan Wermter

Knowledge Technology, Department of Informatics, University of Hamburg

Vogt-Koelln-Str. 30, 22527 Hamburg

g.schwiebert@mailbox.org

{weber, qu, siqueira, wermter}@informatik.uni-hamburg.de

Abstract

Large datasets as required for deep learning of lip reading do not exist in many languages. In this paper we present the dataset GLips (German Lips) consisting of 250,000 publicly available videos of the faces of speakers of the Hessian Parliament, which was processed for word-level lip reading using an automatic pipeline. The format is similar to that of the English language LRW (Lip Reading in the Wild) dataset, with each video encoding one word of interest in a context of 1.16 seconds duration, which yields compatibility for studying transfer learning between both datasets. By training a deep neural network, we investigate whether lip reading has language-independent features, so that datasets of different languages can be used to improve lip reading models. We demonstrate learning from scratch and show that transfer learning from LRW to GLips and vice versa improves learning speed and performance, in particular for the validation set.

Keywords: Audio-visual, Dataset, Lip reading, Automatic Speech Recognition, Deep Learning, Transfer Learning, Computer Vision

1. Introduction

Lip reading is the ability of drawing conclusions about what is being said by visually observing a speaker’s lips. In practice, however, it can rarely be considered on its own in human speech, since a variety of additional information is usually available in communication which, in combination with lip reading, can increase the information content of the incoming communication (Klucharev et al., 2003). These can consist of, for example, audio information, context, heuristics, gestures, facial expressions, or prior knowledge about what is being said. Also, special capabilities of the human brain and auditory sense allow us to apply filters that increase the focus on the desired communication, such as the well-known cocktail party effect (Cherry, 1953). Thus, when phonemes, syllables, words, or entire sentences have an information deficit from the sender or receiver related to the communication pathway, lip reading can be one among many ways to compensate for this deficit. Ambient noise, hearing loss, slurred pronunciation, unfamiliar words, distance, or soundproof barriers are examples of communication problems between sender and receiver. Most people automatically look at the speaker’s lips when intelligibility suffers, so they are all lip readers with varying degrees of skill (Woodhouse et al., 2009).

There is a variety of possible applications for lip reading, such as disability support, sports communication in the press, voice control in noisy environments, additional accuracy for ASR systems, law enforcement and preservation of evidence. In recent years, enormous progress has been made in a number of technical areas, the effects of which enable an efficient tech-

nical evaluation of lip reading (Paszke et al., 2019; Chetlur et al., 2014). High-quality cameras nowadays have high-resolution sensors whose light sensitivity, dynamic range and noise suppression in post-processing enable clear, sharp and detailed images. The memory capacity and computing power of graphic cards have increased to such an extent that complex computations are also possible on consumer PCs (Lemley et al., 2017). New deep learning models require large-scale labelled training data, which is not available for many languages yet. Our motivation is twofold:

1. to contribute German data¹ for the generation of corpus-based lip reading models, and
2. to evaluate this dataset by training a deep neural network and transfer learning to and from the English LRW dataset.

2. Background

Datasets for lip reading are not yet as common as those for speech recognition. In this section, we describe the details of the “Lip Reading in the Wild” (LRW) dataset created by Chung and Zisserman (2016), which is a popular benchmark and high-quality dataset for Automatic Lip Reading (ALR) and thematically related tasks such as Automatic Speech Recognition (ASR) in general. Other published datasets for lip reading are the Chinese LRW-1000 (Yang et al., 2019) and the Romanian (Jitaru et al., 2020). We will use the English LRW for comparison and transfer learning. In the second part

¹GLips is available at the following link: <https://www.inf.uni-hamburg.de/en/inst/ab/wtm/research/corpora.html>

of this section, we briefly address the particular situation of data protection and copyright in Germany and consider the practical obstacles that arise when creating large datasets.

2.1. LRW – Lip Reading in the Wild

The LRW dataset consists of short videos of people’s faces uttering a defined word (Chung and Zisserman, 2016). Care was taken to ensure that the speakers’ lips are clearly visible.

LRW is composed of MPEG4 videos, each 1.16s long and recorded at 25 frames per second (fps). The 500 classes of words, each with approximately 900-1100 instances, spoken by hundreds of different speakers were cut from BBC archives of news broadcasts, talk shows and interviews in 256×256 pixels format and center-focused on the speaker via face detection. In each video clip, the word labelled as a directory name consisting of 5-12 letters is pronounced in a time-centred manner. Here, 50-word instances each are divided between validation and test directories, while the rest is reserved for training purposes. In addition, for each word there is a metadata file in .txt format containing BBC internal data such as disk reference, channel and program start. As an externally usable entry, the duration of the pronunciation of the respective word is noted in seconds. In our evaluation (cf. Section 5), we cropped the videos of the LRW dataset to 96×96 pixels focusing on the lips of the speakers to train our lip reading models.

2.2. Copyright

Copyright law (UrhG) in Germany and its related ancillary copyrights deal with the creation of works and the rights and powers of their creators. Videos that can be accessed and downloaded from publicly accessible platforms are in principle subject to §1 UrhG² as a work as soon as they have a certain so-called creative level, i.e., they are enhanced, for example, by creative editing. Choosing such videos as a source for creating a dataset would involve a great deal of communication effort, as the permission of the author would have to be obtained in writing for each individual work. Videos from webcams or surveillance cameras without further significant creative editing generally lack this level of creativity, which is why they are particularly suitable as a data source for dataset creation, as long as the rules of the DSGVO³ are followed. Furthermore, in creating GLips, we comply with two special exceptions embedded into the German copyright law. First, we pursue a legitimate scientific interest for helping to enhance the support for the hearing impaired through the creation of our dataset and second, the politicians shown

²§1 UrhG-Allgemeines-dejure.org: <https://dejure.org/gesetze/UrhG/1.html>

³Datenschutz-Grundverordnung (DSGVO) - [dejure.org: https://dejure.org/gesetze/DSGVO](https://dejure.org/gesetze/DSGVO)

are persons of public interest recorded in a publically available parliament-recorded environment.

2.3. General Data Protection Regulation in Germany (DSGVO)

When people are recorded on film, these videos are subject to the DSGVO, which has been valid in the European Union since 25.05.2018. The DSGVO requirements cover, among other things, in Art. 5: (1) legality, (2) public interest for a specific purpose, (3) data minimization, (4) correctness, (5) storage time limit, and (6) integrity and confidentiality, whereby special regulations regarding items 2 and 5 apply for scientific research purposes, which e.g., allow a longer storage period. This binding framework increases the bureaucratic, legal and possibly also personnel effort for a data protection officer, which further reduces the availability of datasets for machine learning purposes. But since biometric features are usually not mutable and we cannot yet fully estimate how many derivations from biometric data are possible in the future (Faundez-Zanuy, 2005), we as dataset creators have a special responsibility with regard to copyright, data protection and licensing.

3. GLips Dataset Creation

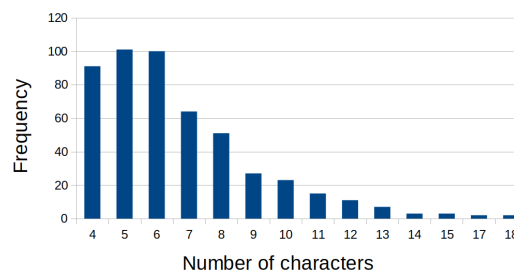


Figure 1: Word length distribution in GLips

This section gives an overview of the creation procedure of the dataset German Lips (GLips). Despite the focus of this paper and the name on the ALR domain, GLips should be applicable in the whole scientific ASR domain as versatile as possible, because, as explained in Section 2, the legally compliant creation of large video datasets in the German-speaking area is connected with some hurdles. Therefore, the creation of GLips is oriented towards LRW in order to ensure a high compatibility for methods such as transfer learning and experiments on the topic of Language Independence and to advance scientific knowledge in these areas. Furthermore, by choosing video material based on naturally spoken language in a natural environment, we decided to use this approach for ASR systems, as it produces more robust results for real-world applications than artificially generated datasets with as little noise as possible (Burton et al., 2018).

GLips consists of 250,000 H264-compressed MPEG-4 videos of speakers’ faces from parliamentary sessions

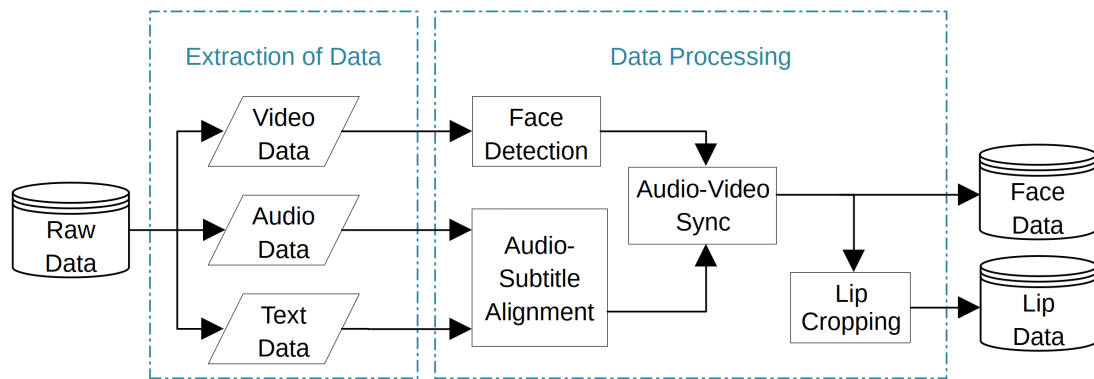


Figure 2: Pipeline for the generation of training data

of the Hessian Parliament, which are divided into 500 different words of 500 instances each. The word length distribution is shown in Fig.1. As with LRW, each video is 1.16s long at a frame rate of 25fps. The audio track was stored separately in an MPEG AAC audio file (.m4a). For each video there is an additional metadata textfile with the fields:

- Spoken word,
- Start time of utterance in seconds,
- End time of utterance in seconds,
- Duration of utterance in seconds,
- Corresponding numerical filename in the database.

Start- and end-time of utterance refers to the complete original video and not to the occurrence of the word in the clip.

3.1. Acquisition

With the permission of the Hessian Parliament, we used over 1000 videos and their respective subtitles. The Hessian parliament has published a superset of these videos also on its YouTube channel⁴. The subtitles are available as a separate text file and include manually created subtitles with time intervals. Similar to LRW subtitle editing (Chung and Zisserman, 2016), this leads to the issue that not all subtitles are verbatim, as in rare cases the content but not the exact spoken words have been reproduced in the subtitle, which means that despite checks, there are likely to be some words in the dataset that do not match the lip profile of the speaker. In order to create GLips, we also need the exact time of pronunciation and the duration of the utterance for each selected word. However, the subtitle files only contain one interval for each of several words. The solution to this problem via alignment using the WebMAUS service is discussed in section 3.3.

⁴YouTube - Hessischer Landtag: <https://www.youtube.com/c/HessischerLandtagOnline>

3.2. Multimodal Processing Pipeline

The technical creation of the multimodal dataset GLips is thematically divided into the two areas of extraction and processing of data. In Fig. 2, the entire pipeline is shown schematically from the existence of the raw data to the creation of training data suitable for machine learning, of which GLips represents a subset.

Since the original audio, video and subtitle data are already available in a separate form, the technical part of the data extraction is limited to the acquisition of all data and the cleaning of the text data from meta information so that only the spoken words are available as input for the next step. The more complex part of the data processing is described in more detail in the following two sections and is divided into the two subsections audio subtitle alignment using WebMAUS and face detection. The audio and video files are synchronized in the last step, but are stored in separate files for the sake of more diverse processing options.



Figure 3: Example of GLips cropped to 96×96 pixels

The output of this pipeline consists of structured, processed, and augmented data suitable as potential training data for various areas of machine learning. The TextGrid files with their phonetic information, which are no longer needed by GLips, longer excerpts from aligned videos for sentence-based lip reading approaches, or video clips with several people to test attention mechanisms, are just a few ideas for research. For the transfer learning in Section 4, we use a modified GLips dataset that was reduced in size from 256×256 pixels to 96×96 pixels (see Fig. 3) by additionally cropping the videos to focus on lip reading learning and to ensure better computability on consumer hardware.

3.3. Subtitle Alignment using WebMAUS

The Munich Automatic Segmentation System (MAUS) is a software for manifold speech data processing developed by Schiel et al. (2011) which is also available as a web service called WebMAUS⁵ (Kisler et al., 2012; Ide, 2017). For us, the modules G2P, Chunker and MAUS from this software package are of particular interest as an automatic pipeline via RESTful API to be able to create GLips. From our previously cleaned subtitle text files, a phonological transcript is created using G2P in BAS score format⁶, which is an open format for describing segmental information. A chunker is also needed to keep the audio and text information more easily computable, since some audio files come from hours of parliamentary sessions that can push the WebMAUS server to its capacity limits when sent in aggregate over several days. From these smaller segments, the aligner from the WebMAUS service can use the corresponding audio file to temporally match the text with the audio file in order to generate a TextGrid file with matching phonetic and word segments. There are three levels of analyzable information in this. As can be seen in Fig. 5, the levels ORT-MAU (orthographic information), KAN-MAU (canonical-phonemic word representation) and MAU (aligned transcription) including the corresponding time axis are available. With this information about an exact time period of the pronunciation of each word spoken in the audio file, it is possible to extract the temporally related video clip from the original video.

3.4. Face Detection

The face detection was implemented based on the two Python libraries OpenCV⁷ with Nvidia-CUDA⁸ support for more efficient performance and

⁵BAS web service interface: <https://clarin.phonetik.uni-muenchen.de/BASWebServices/interface>

⁶Phonetik BAS: <https://www.phonetik.uni-muenchen.de/Bas/BasFormatsdeu.html>

⁷OpenCV: <https://opencv.org/>

⁸CUDA Toolkit: <https://developer.nvidia.com/cuda-toolkit>



Figure 4: Example full screen view of the raw video data from the Hessian Parliament

Video Properties	LRW	GLips
Total number	~500,000	250,000
Number of different speakers	hundreds	~100
Format	MPEG4	MPEG4
Resolution in pixels	256×256	256×256
Length	1,16s	1,16s
Framerate	25fps	25fps
Word classes	500	500
Instances	~1000	500
Word length in letters	5-12	4-18
Metadata file	yes	yes
Separate audio file	no	yes
TextGrid file	no	yes
Equipment level	professional	webcam
Lighting	professional	indoor standard

Table 1: Comparison of LRW and GLips

face_recognition⁹. As shown in Fig. 4, the external conditions for speaker detection are almost perfectly suited to run face detection only on a section of the video due to the fixed podium and almost constant position of the webcam, which both reduces the processing time of all videos and makes attention mechanisms less important for speaker detection. Very few complications occur in the videos that these processing conditions do not satisfy, s.a. the speaker moving too vividly or being very tall which could cause the face detection to confuse the speakers face with the person sitting in the elevated position behind him.

4. Comparison of GLips and LRW

Both GLips and LRW contain large quantities of videos of speakers in frontal view perspectives to ensure a clear view on their lips. As seen in Table 1, special

⁹GitHub - ageitgey face_recognition: https://github.com/ageitgey/face_recognition




Figure 5: The spoken word “Landtags” including context visualized from a TextGrid file using Praat software (Boersma, 2001)

attention was paid to the aspect of compatibility between the two datasets. The camera equipment of the BBC is clearly of higher quality than the webcam of the Hessian Parliament so that despite nominally the same video resolution, there is a difference in quality between the video datasets due to dissimilar dynamic range of the camera sensor, possibly existing camera-internal post-processing, as well as more elaborately calculated and manufactured lenses. In addition, external factors such as shorter camera distance (object distance), the partially existing professional and intelligibility-oriented speech training of the news presenters and the more professional lighting in the BBC dataset provide a clearer, higher-contrast and sharper image of the lip movements, so that it can be expected that LRW-trained models for lip reading will have a higher performance in terms of word recognition than will be the case with GLips. The quality of the audio recordings, which were integrated into the .mp4 format in LRW and are available separately as .m4a in GLips, is less deviant due to the use of high-quality microphones in the Hessian Parliament. However, for training our lip reading models we will only use the visual information. Furthermore, the number of speakers in LRW is several hundred, which is significantly higher than in GLips, which is estimated to be around 100.

5. Model Evaluation




Figure 6: X3D Architecture

There are several recurrent (e.g. (Stafylakis and Tzimiropoulos, 2017)) and feedforward (e.g. (Xinshuo Weng, 2019), (Martínez et al., 2020)) models for lip reading on word level. We chose the X3D convolutional neural network model by Feichtenhofer (2020) since it is efficient for video classification in terms of accuracy and computational cost and well designed for the processing of spatiotemporal features. The model is

depicted in Figure 6, has 4D tensors as its main layers, with one dimension each for the temporal dimension (T), height and width ($H \times W$) of spatial dimensions and number of channels (C). It is an expanding image processing architecture that uses channel-wise convolutions as building blocks. Synchronized stochastic gradient descent (SGD) was performed of parallel workers following the linear scaling rule for learning rate and minibatch size to reduce training time (Goyal et al., 2018).

We used the official model implementation¹⁰ that is included in Pytorch Lightning-Flash¹¹ and for video processing we use the PyTorchVideo (Fan et al., 2021) library. The model was implemented and tested on a single NVIDIA Geforce RTX2080Ti.

5.1. Experiments with GLips and LRW

To evaluate whether the word recognition rate of the lip reading models can be improved by transfer learning, we conducted two experiments. To keep the transfer learning computations manageable, we create two subsets of each of the LRW and GLips datasets, which we call LRW_{15} and $GLips_{15}$, and which consist of only 15 randomly selected words of 500 instances each instead of all 500 words. Two further subsets named $GLips_{15-small}$ and $LRW_{15-small}$ consist only of a total of 95 word instances of the same 15 words as the former subsets. We cropped the videos to 96×96 pixels around the lip region to increase the performance in computation and to improve the focus of learning on the lips.

In Experiment 1, we use the largest versions of the $GLips_{15}$ and LRW_{15} datasets among themselves for transfer learning. We test here whether lip reading abilities can be passed among models trained on the same size datasets in a language-independent manner by comparing the word recognition rate of the transfer-learned models with those learned from scratch.

In Experiment 2, we test whether the transfer learning benefits are greatest, as theoretically expected, where

¹⁰GitHub - X3D implementation: https://github.com/facebookresearch/SlowFast/blob/main/MODEL_ZOO.md

¹¹GitHub - PyTorch Lightning Flash: <https://github.com/PyTorchLightning/lightning-flash>

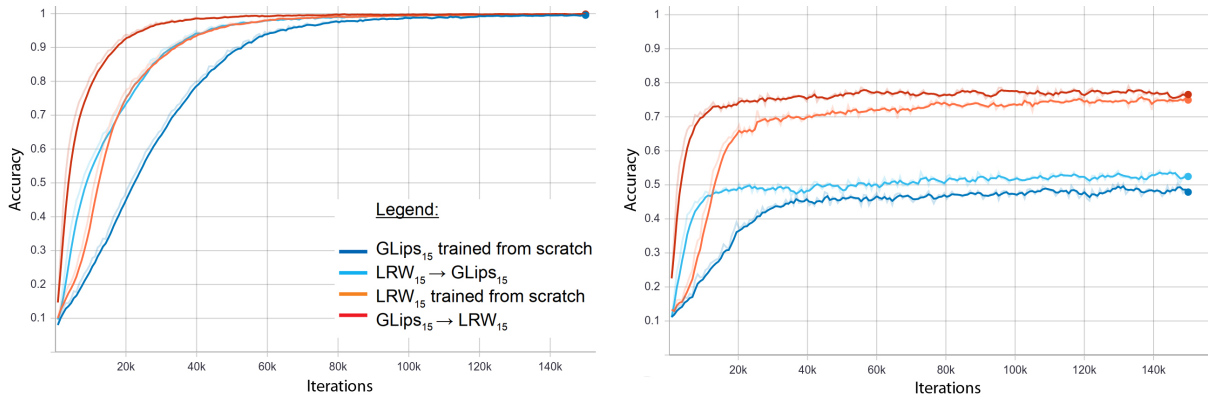


Figure 7: Results of training (left) and validation (right) accuracy per iteration for Experiment 1: transfer learning between same-size datasets

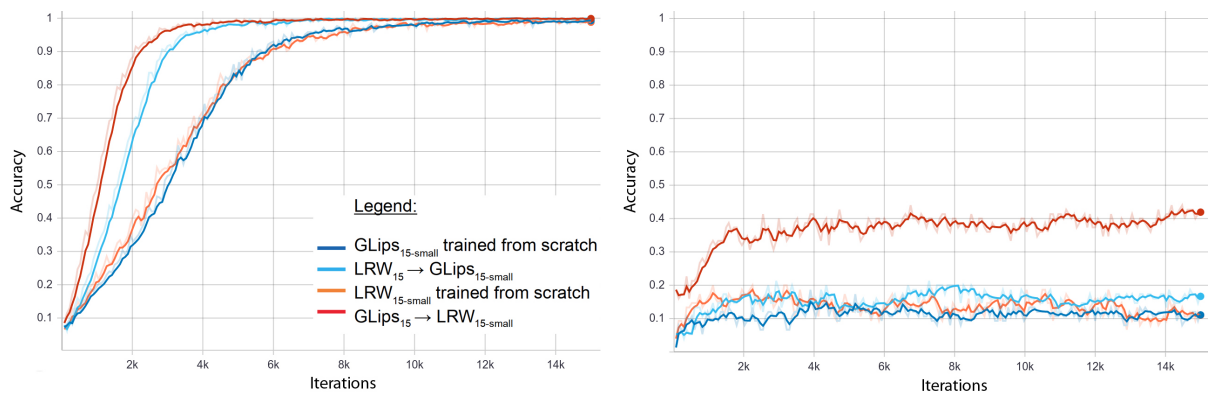


Figure 8: Results of training (left) and validation (right) accuracy per iteration for Experiment 2: transfer learning from a large dataset to a small dataset

models learned on large datasets transfer to models learned on smaller datasets by learning from $LRW_{15} \rightarrow GLips_{15-small}$ and also from $GLips_{15} \rightarrow LRW_{15-small}$, and also by comparing the results to models learned from scratch.

5.2. Experimental Results

The smoothed curves of the results were plotted in TensorBoard¹². As seen in Fig. 7, in Experiment 1 the validation accuracies of $GLips_{15}$ and LRW_{15} trained from scratch are lower than those of the transfer-learned models $GLips_{15} \rightarrow LRW_{15}$ and $LRW_{15} \rightarrow GLips_{15}$. Also both transfer-learned curves rise steeper in the beginning, which means that the networks learn faster and maintain their higher level across all epochs, manifesting their learning advantage. Additionally, $GLips_{15} \rightarrow LRW_{15}$ has a higher starting point, which accelerates the learning rate even more.

The results of Experiment 2 in Fig. 8 also clearly demonstrate the advantages of transfer learning, but look different in detail. $LRW_{15} \rightarrow GLips_{15-small}$ as well as $GLips_{15} \rightarrow LRW_{15-small}$ achieve an advantage over the respective networks trained from scratch. In both experiments the average validation accuracies of the LRW-networks reach a higher score than the GLips-networks. This is particularly pronounced for $GLips_{15} \rightarrow LRW_{15-small}$, which appears has the largest advantage in the validation experiment. It is surprising that GLips as source of transfer learning helps learning LRW more (dark red curve) than vice versa, since GLips as the source of transfer learning has lower-quality videos. An explanation could be that given the low number of data points in this experiment, GLips with its noisier visual features did not allow the model to overfit, while on the other hand, the LRW-pretrained model (light blue) might have overfit to distinct features in LRW.

¹²TensorBoard: <https://www.tensorflow.org/tensorboard>

6. Discussion

The successful transfer learning between the two languages indicates that there are features in both datasets regarding lip reading that are language-independent and thus can be transferred to another language.

Due to the better performance of the LRW-trained networks, we hypothesize that the difference between the models trained on GLips and LRW lies in the quality of the data. In GLips, in comparison, overall learning is subjected to more noise in addition to the features important for the complex task of lip reading. However, the evaluation of the curves shows clear advantages of transfer learning compared to learning from scratch, both in the overall performance and in the speed of learning. The quality of a dataset is of utmost importance for the performance of a model using it. Transfer learning in neural networks amplifies the effects of the quality differences between datasets. In both experiments, we can observe the advantage of the higher quality dataset of LRW for transfer learning, as anticipated in Section 4. In particular, for models trained from small datasets, the benefit of transfer learning compared to learning from scratch becomes substantial.

The benefit of transfer learning for lip reading has recently been shown between English, Romanian and Chinese language (Jitaru et al., 2021). Their work showcases transfer learning from multiple source languages. Here one needs to take into account that some languages have larger and higher-quality datasets, benefiting transfer learning more than other languages.

Speaker independence according to Bear and Taylor (2017) is only possible in GLips if the dataset would be split into training, validation and test speakers according to different speakers. We did not implement a speaker-independent split in GLips in favor of the larger word set. If splitting GLips in a speaker-independent fashion, the number of words in the training set would reduce to less than 500 for some words. As a technical implementation, it would be possible to perform a face recognition on the original videos and to link the different identities with the spoken word instances. With this extended information dimension, disjoint test sets can be created with respect to the training and validation sets. Data augmentation like translation or altering the RGB-values of the videos were not used in the training for performance reasons. However, such techniques may be useful on more powerful architectures to reduce the risk of overfitting (Krizhevsky et al., 2017) and to increase generalization performance.

7. Conclusion

With GLips we have created the first large-scale German lip reading dataset, which is compatible to the large English LRW dataset. Using an X3D deep neural network, we demonstrated successful transfer learning between the two languages in both directions. This new dataset can support various applications in areas of dis-

ability support, communication in noisy environments, boosting of existing ASR systems, etc., progressing state-of-the-art assistive technologies. Revisiting the publicly available source of the dataset, further applications would become possible, such as learning automatic speech recognition, and extended TextGrid information will allow to create a dataset for sentence-level recognition from the original videos.

8. Acknowledgements

We gratefully acknowledge partial support from the German Research Foundation (DFG) under the project Crossmodal Learning (CML, Grant TRR 169).

9. Bibliographical References

- Bear, H. L. and Taylor, S. (2017). Visual speech recognition: aligning terminologies for better understanding. *arXiv preprint arXiv:1710.01292*.
- Boersma, P. (2001). Praat, a system for doing phonetics by computer. *Glott. Int.*, 5(9):341–345.
- Burton, J., Frank, D., Saleh, M., Navab, N., and Bear, H. L. (2018). The speaker-independent lipreading play-off; a survey of lipreading machines. In *2018 IEEE International Conference on Image Processing, Applications and Systems (IPAS)*, pages 125–130. IEEE.
- Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *The Journal of the Acoustical Society of America*, 25(5):975–979.
- Chetlur, S., Woolley, C., Vandermersch, P., Cohen, J., Tran, J., Catanzaro, B., and Shelhamer, E. (2014). cuDNN: Efficient primitives for deep learning. *arXiv preprint arXiv:1410.0759*.
- Chung, J. S. and Zisserman, A. (2016). Lip reading in the wild. In *Asian Conference on Computer Vision*, pages 87–103. Springer.
- Fan, H., Murrell, T., Wang, H., Alwala, K. V., Li, Y., Li, Y., Xiong, B., Ravi, N., Li, M., Yang, H., Malik, J., Girshick, R., Feiszli, M., Adcock, A., Lo, W.-Y., and Feichtenhofer, C. (2021). PyTorchVideo: A deep learning library for video understanding. In *Proceedings of the 29th ACM International Conference on Multimedia*. <https://pytorchvideo.org/>.
- Faundez-Zanuy, M. (2005). Privacy issues on biometric systems. *IEEE Aerospace and Electronic Systems Magazine*, 20(2):13–15.
- Feichtenhofer, C. (2020). X3D: Expanding architectures for efficient video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 203–213.
- Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., and He, K. (2018). Accurate, large minibatch SGD: Training ImageNet in 1 hour.
- Ide, N. (2017). *Handbook of Linguistic Annotation*. Springer Netherlands Imprint Springer, Dordrecht.

- Jitaru, A. C., Abdulamit, Ş., and Ionescu, B. (2020). LRRo: a lip reading data set for the under-resourced romanian language. In *Proceedings of the 11th ACM Multimedia Systems Conference*, pages 267–272.
- Jitaru, A., Stefan, L.-D., and Ionescu, B. (2021). Toward language-independent lip reading: A transfer learning approach. In *2021 International Symposium on Signals, Circuits and Systems (ISSCS)*.
- Kisler, T., Schiel, F., and Sloetjes, H. (2012). Signal processing via web services: the use case WebMAUS. In *Digital Humanities Conference 2012*.
- Klucharev, V., Möttönen, R., and Sams, M. (2003). Electrophysiological indicators of phonetic and non-phonetic multisensory interactions during audiovisual speech perception. *Cognitive Brain Research*, 18(1):65–75.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2017). ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90.
- Lemley, J., Bazrafkan, S., and Corcoran, P. (2017). Deep learning for consumer devices and services: Pushing the limits for machine learning, artificial intelligence, and computer vision. *IEEE Consumer Electronics Magazine*, 6(2):48–56.
- Martínez, B., Ma, P., Petridis, S., and Pantic, M. (2020). Lipreading using temporal convolutional networks. *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6319–6323.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32:8026–8037.
- Schiel, F., Draxler, C., and Harrington, J. (2011). Phonemic segmentation and labelling using the MAUS technique. In *Workshop New Tools and Methods for Very-Large-Scale Phonetics Research*.
- Stafylakis, T. and Tzimiropoulos, G. (2017). Combining Residual Networks with LSTMs for Lipreading. In *Proc. Interspeech 2017*, pages 3652–3656.
- Woodhouse, L., Hickson, L., and Dodd, B. (2009). Review of visual speech perception by hearing and hearing-impaired people: clinical implications. *International Journal of Language & Communication Disorders*, 44(3):253–270.
- Xinshuo Weng, K. K. (2019). Learning spatio-temporal features with two-stream deep 3D CNNs for lipreading. In *Proceedings of British Machine Vision Conference (BMVC '19)*, September.
- Yang, S., Zhang, Y., Feng, D., Yang, M., Wang, C., Xiao, J., Long, K., Shan, S., and Chen, X. (2019). LRW-1000: A naturally-distributed large-scale benchmark for lip reading in the wild. In *2019 14th IEEE International Conference on Automatic Face Gesture Recognition (FG 2019)*, pages 1–8.

10. Language Resource References

- Gerald Schwiebert, Cornelius Weber, Leyuan Qu, Henrique Siqueira, Stefan Wermter. (2022). *GLips - German Lipreading Dataset*. University of Hamburg.