

AnnoTheia: A Semi-Automatic Annotation Toolkit for Audio-Visual Speech Technologies

José-M. Acosta-Triana¹, David Gimeno-Gómez², Carlos-D. Martínez-Hinarejos²

¹ ValgrAI - Valencian Graduate School and Research Network of Artificial Intelligence
Camino de Vera, s/n, 3Q Building, 46022, València, Spain

² Pattern Recognition and Human Language Technologies Research Center
Universitat Politècnica de València, Camino de Vera, s/n, 46022, València, Spain
joactr@inf.upv.es, dagigo1@dsic.upv.es, cmartine@dsic.upv.es

Abstract

More than 7,000 known languages are spoken around the world. However, due to the lack of annotated resources, only a small fraction of them are currently covered by speech technologies. Albeit self-supervised speech representations, recent massive speech corpora collections, as well as the organization of challenges, have alleviated this inequality, most studies are mainly benchmarked on English. This situation is aggravated when tasks involving both acoustic and visual speech modalities are addressed. In order to promote research on low-resource languages for audio-visual speech technologies, we present AnnoTheia, a semi-automatic annotation toolkit that detects when a person speaks on the scene and the corresponding transcription. In addition, to show the complete process of preparing AnnoTheia for a language of interest, we also describe the adaptation of a pre-trained model for active speaker detection to Spanish, using a database not initially conceived for this type of task. The AnnoTheia toolkit, tutorials, and pre-trained models are available at <https://github.com/joactr/AnnoTheia/>.

Keywords: data annotation, speech technologies, audio-visual databases, active speaker detection

1. Introduction

Speech technologies aim to develop tools capable of enabling human-machine interaction via different types of communication. Multiple research areas are included in this vast field, e.g., speech recognition (Baeviski et al., 2020; Radford et al., 2023; Shi et al., 2022; Ma et al., 2021); speech synthesis (Shen et al., 2018; Li et al., 2019b; Liu et al., 2023); speech translation (Jia et al., 2019; Lee et al., 2022; Huang et al., 2023); sign language recognition (Koller et al., 2015; Cihan Camgöz et al., 2020; Albanie et al., 2021); silent speech interfaces (Denby et al., 2010; González-López et al., 2020); speech-based disease detection (Ablimit et al., 2022; Milling et al., 2022); speech enhancement and separation (Pascual et al., 2017; Yen et al., 2023; Huang et al., 2022); speaker diarization (Fujita et al., 2019; Bredin and Laurent, 2021); spoken dialogue systems (Jokinen and McTear, 2022). All these studies reveal the importance that speech technologies and their applications can represent in our daily lives.

One of the most widely studied tasks in the field is automatic speech recognition. However, although more than 7,000 known languages are spoken around the world (Eberhard et al., 2023), only a small fraction of them are currently covered by these speech technologies (Pratap et al., 2023). Taking into account that many of these languages are considered endangered (Bromham et al., 2022), the lack of support from current tech-

nologies may contribute negatively to this situation. For this reason, to avoid any type of discrimination or exclusion, different works have promoted scaling speech technologies to cover the greatest possible number of languages. Massive multi-lingual speech corpora collections (Ardila et al., 2020; Salesky et al., 2021; Anwar et al., 2023; Pratap et al., 2023), large multi-lingual speech recognition models (Radford et al., 2023; Pratap et al., 2023; Zhang et al., 2023; Chen et al., 2023), self-supervised speech representations (Baeviski et al., 2020; Ravanelli et al., 2020; Babu et al., 2022; Shi et al., 2022), as well as the organization of challenges, such as the ML-SUPERB (Shi et al., 2023), have been some of the most common approaches to alleviate this inequality. Nevertheless, despite all these efforts, most studies are still benchmarked on English (Pratap et al., 2023).

In addition, the described approaches mainly focused on auditory-based speech recognition, contrasting with the recent trend towards developing audio-visual speech technologies (Ma et al., 2021; Shi et al., 2022; Liu et al., 2023). Tasks such as automatic lipreading (Afouras et al., 2018; Ma et al., 2022; Prajwal et al., 2022), audio-visual speech recognition (Ma et al., 2021; Shi et al., 2022; Burchi and Timofte, 2023), or audio-visual speech synthesis (Liu et al., 2023) have received increasing interest in the last decades. However, the scarcity of data for low-resource languages is exacerbated when this type of tasks is addressed due to the difficulty of collecting data where audio and video

	English	Persian	Spanish	French	Portuguese	Italian	Chinese	Russian	German	Greek	Arabic
MuAViC ¹	438.9 [†]	-	181.0	179.0	156.0	104.0	-	52.0	13.0	29.0	19.0
CMU-MOSEAS ²	-	-	16.3	15.6	16.0	-	-	-	18.6	-	-
LRS2-BBC ³	224.5	-	-	-	-	-	-	-	-	-	-
LRS3-TED ⁴	438.9	-	-	-	-	-	-	-	-	-	-
CMLR ⁵	-	-	-	-	-	-	86.5	-	-	-	-
Arman-AV ⁶	-	220.0	-	-	-	-	-	-	-	-	-
LIP-RTVE ⁷	-	-	13.0	-	-	-	-	-	-	-	-
Total of Hours	663.4	220.0	210.3	194.6	172.0	104.0	86.5	52.0	31.6	29.0	19.0

¹(Anwar et al., 2023) ²(Zadeh et al., 2020) ³(Afouras et al., 2018) ⁴(Afouras et al., 2018a) ⁵(Zhao et al., 2020) ⁶(Peymanfard et al., 2023) ⁷(Gimeno-Gómez and Martínez-Hinarejos, 2022) [†] The MuAViC database was sourced from TED and TEDx talks: thus, it includes directly the LRS3-TED corpus.

Table 1: Language coverage in terms of hours by databases focused on unconstrained continuous audio-visual speech recognition. Corpora oriented to word-level classification, collected in controlled recording studios, or not providing transcriptions were not considered.

are aligned. Table 1 reflects the limited number of languages covered by the audio-visual corpora that have been publicly released for unconstrained audio-visual speech recognition. Corpora oriented to word-level classification (Chung and Zisserman, 2017; Yang et al., 2019; Egorov et al., 2021) or collected in controlled recording studios (Harte and Gillen, 2015; Fernandez-Lopez et al., 2017) were not considered, since they do not represent the actual nature of speech¹.

A special case is the AV-Speech (Ephrat et al., 2018), a large-scale database offering over 4k hours of data for multiple languages, which was also not considered because it does not provide transcriptions. It should be noted that this corpus could be used for pre-training multilingual encoders in a self-supervised manner (Shi et al., 2022). However, the current lack of audio-visual language resources would impede a more effective use of this technology (Anwar et al., 2023). Despite all these efforts in the audio-visual domain, there is no comparison to large-scale audio-based speech databases that cover more than a thousand languages with thousands of hours of data (Pratap et al., 2023; Zhang et al., 2023). Besides, the inequality in terms of the number of hours among the different languages reflected in Table 1 and the fact that most studies based on audio-visual technologies are mainly benchmarked on English, highlights the need to promote and encourage the collection and annotation of new audio-visual databases to cover the largest possible number of languages.

For this reason, numerous works have described multi-step pipelines to collect and annotate their own audio-visual databases (Afouras et al., 2018,a; Peymanfard et al., 2023; Ephrat et al., 2018). However, while building these pipelines

¹A more comprehensive review of existing audio-visual databases can be found in (Fernandez-Lopez and Sukno, 2018).

is not trivial, to the best of our knowledge, no toolkit has been publicly released for the research community. In addition, one of the most important components of these pipelines is the Active Speaker Detection (ASD) module (Tao et al., 2021; Liao et al., 2023; Min et al., 2022), which allows us to identify the person on the scene who is talking. Nevertheless, these ASD models are also heavily benchmarked on English (Tao et al., 2021; Liao et al., 2023; Min et al., 2022).

Motivated by this fact, we present AnnoTheia, a semi-automatic toolkit that detects when a person speaks on the scene and the corresponding transcription. One of the most notable aspects of the proposed toolkit is the flexibility to replace a module with another of our preference or, where appropriate, adapted to our language of interest. Therefore, to show the complete process of preparing AnnoTheia for a language of interest, we also describe in this paper the adaptation of a pre-trained ASD model to Spanish, using a database not initially conceived for this type of task.

These were the main reasons that motivated our work, whose key contributions are:

- We present AnnoTheia, a semi-automatic annotation toolkit to encourage the largest possible language coverage for audio-visual speech technologies, highlighting its intuitive user interface and its flexibility to replace the different modules that composed it.
- We described how to adapt a pre-trained ASD model to a language of interest, even in those occasions where there are no databases originally conceived for this type of task.
- We consequently present RTVE-ASD, a new corpus for ASD in Spanish collected from TV newscast programs, where multiple challenging in-the-wild situations can be found.



Figure 1: AnnoTheia toolkit user interface. **A**: Video display of the scene candidate to be a new sample of the future database. An overlying green bounding box highlights the active speaker detected by the toolkit. **B**: Keyword legend to control the video display. **C**: Transcription automatically generated by the toolkit. It can be edited by the annotator. **D**: Buttons to allow the annotator to accept or discard the candidate scene sample. **E**: Navigation buttons through candidate scenes. It can be useful to correct possible annotation mistakes.

2. The AnnoTheia Toolkit

We present AnnoTheia, the first open-source toolkit for the semi-automatic annotation of audiovisual speech resources, promoting the further development of speech technologies for low-resource languages. In this section, we describe the multiple modules that compose the toolkit, as well as the design of the user interface. However, we highlight the flexibility with which the toolkit was provided when replacing one of these modules with another of our preference or, as described in Section 3, adapted to our language of interest.

Scene Bounding Detection. PySceneDetection² is a toolkit designed to detect bounding between scenes in videos. Specifically, it computes a score representing the difference in content between the current and previous frame in the HSV color space for each video frame. In this way, we were able to accelerate the processing of large video clips by checking if at least one person was present on the scene simply by checking the first frame that composed each detected scene. Hence, using the face detector described below, we discarded those scenes where no people were present.

Face Detection. Using Dual Shot Face Detector (Li et al., 2019a), we were able to accurately extract faces from the scenes previously selected. However, more than one person could appear

on the scene at the same time, hindering the subsequent Active Speaker Detection. Therefore, we implemented a frame-by-frame face-distance comparison algorithm, where each face was matched to a specific individual based on proximity, allowing for the precise assignment of faces to their respective individuals throughout the video sequence. The detected bounding boxes were saved not only because we needed the cropped faces for the Active Speaker Detection module described below, but also because this information might be useful once the future database is compiled. An alternative would be using the state-of-the-art RetinaFace face detector (Deng et al., 2020).

Face Alignment. On multiple occasions, it can also be useful to consider a specific region of the face, as it is the case with the visual speech recognition task, also known as automatic lipreading (Afouras et al., 2018; Ma et al., 2022; Shi et al., 2022). For this reason, we also introduced a face alignment module based on the Face Alignment Network architecture proposed by Bulat and Tzimiropoulos (2017). Hence, our toolkit will also provide 68 2-dimensional facial landmarks for each frame of the detected scene. An alternative would be using the face landmark detector provided by the Google’s MediaPipe framework³.

Active Speaker Detection. Once the different

²<https://github.com/Breakthrough/PySceneDetect>

³https://developers.google.com/mediapipe/solutions/vision/face_landmarker

people appearing on the scene were detected, we used the TalkNet-ASD model (Tao et al., 2021) to identify who of them was really speaking, since, on certain occasions, we may encounter dubbing or voiceovers. Besides, we studied different strategies when processing large video clips, concluding that a non-overlapping window-slicing method offered the best trade-off between quality and speed. A more detailed description of the TalkNet-ASD model architecture can be found in Subsection 3.1. An alternative would be using the state-of-the-art Light-ASD model (Liao et al., 2023).

Post-Process. The ASD module previously described is not exempt from errors. Therefore, we applied a smoothing average strategy for stabilizing the complete scene detection process. This method entailed calculating the mean score provided by the ASD model for a defined sliding window of frames, assigning this averaged score to the central frame within the window. The window then shifts by one frame at a time, repeating the process across the entire sequence. This approach effectively mitigates potential fluctuations in classification scores by incorporating information from neighboring frames, leading to a more consistent and stable classification.

Scene Trimming. We only considered as candidate scene samples those that encompassed frames in which the ASD module detected that a person was actively talking, trimming the scenes with a margin of a small number of frames both at the beginning and at the end of the scene.

Automatic Transcription. Whisper is a multilingual speech recognition system covering around 100 different languages (Radford et al., 2023). By applying this technology to the raw waveform of our trimmed scenes, we were able to obtain word-level aligned transcriptions, thus facilitating the annotator’s task. Although the model automatically detects the spoken language, it can also be manually declared by the user. An alternative would be using the state-of-the-art MMS models (Pratap et al., 2023)⁴.

User Interface. In order to facilitate the annotator’s task, an intuitive user interface has been designed to accelerate the data collection process. Details of the design can be found in Figure 1.

3. Fine-Tuning an ASD Model

This section describes the entire process we followed to prepare AnnoTheia for a language of interest by fine-tuning a pre-trained ASD model even in those situations where there are no databases originally conceived for this type of task. In our case, we adapted the TalkNet-ASD model (Tao

et al., 2021) pre-trained for the AVA-ActiveSpeaker benchmark (Roth et al., 2020) to the Spanish language using the LIP-RTVE database (Gimeno-Gómez and Martínez-Hinarejos, 2022). Hence, in this section, we described the model architecture, how to construct a dataset to estimate an ASD model, and the training setup, as well as a discussion of the findings we observed in our results.

3.1. The TalkNet-ASD Model

The TalkNet-ASD model (Tao et al., 2021) was explicitly designed for the ASD task, where the model aims to identify who of the people appearing in a video clip is actually speaking and at what time. It required two simultaneous inputs: raw grayscale frames of the cropped person’s face, providing the visual cues, and pre-processed audio features in the form of 13 Mel Frequency Cepstral Coefficients (MFCC) (Gales and Young, 2008), representing the audio content. It was able to achieve state-of-the-art results on the AVA-ActiveSpeaker benchmark (Roth et al., 2020). As reflected in Figure 2, the TalkNet-ASD architecture is composed of three main modules.

Window Sampling. Due to the unfeasibility, in computational terms, of processing large videos in their entirety, TalkNet-ASD worked at the window sample level. In the original work, the authors studied multiple window sizes, covering contexts of different numbers of seconds.

Feature Representation Frontend. Once the context windows are sampled, the visual and acoustic cues described above are each processed by an independent encoder.

- **Visual Temporal Encoder.** This encoder was aimed to capture long-term facial expression dynamics. It comprises a visual frontend for spatial information and a visual temporal block for temporal patterns. The frontend consists of a 3D convolutional layer followed by a 2D ResNet-18 (He et al., 2016), providing a sequence of frame-based embeddings. Subsequently, a block mainly based on depth-wise separable convolutional layers (Guo et al., 2018) with residual connections is used to model temporal relationships of the previous visual representation.
- **Audio Temporal Encoder.** By using a 2D ResNet-34 (He et al., 2016) with a squeeze-and-excitation module (Hu et al., 2018), this encoder was able to learn audio representations from temporal dynamics. Given the MFCCs described above, the encoder produces a sequence of audio embeddings that, thanks to the use of dilated convolutions when defining the ResNet-34, match the length of the visual embedding sequence.

⁴<https://github.com/facebookresearch/fairseq/tree/main/examples/mms>

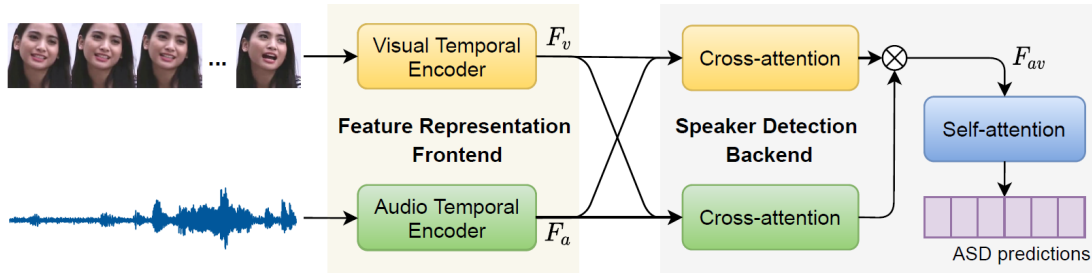


Figure 2: The overall architecture of the TalkNet-ASD model (Tao et al., 2021).

Speaker Detection Backend. This module is composed of two network blocks mainly based on attention mechanisms (Vaswani et al., 2017). First, a cross-attention network is employed to dynamically capture audio-visual interactions by aligning the audio and visual embeddings previously obtained by the feature frontend. Once a joint audio-visual representation is computed, a self-attention network focuses on temporal relationships, distinguishing whether the person is speaking or not across the video frames.

Post-Process Another important aspect to highlight is the use of an optimum classification threshold. Because the main purpose of the proposed toolkit is to provide the annotator with candidate scene samples to create a new database, we were interested in losing, within reason, as few candidate scenes as possible due to false negatives that the model could predict, delegating the final decision to the annotator’s criteria. Therefore, we computed the optimum threshold⁵ to find the best trade-off between the true-positive and false-positive rates on a validation set.

3.2. The ASD-RTVE corpus

This section describes how, from a database that was not originally conceived to address this type of tasks, we were able to construct a dataset to estimate and evaluate an ASD model.

We present ASD-RTVE, a new corpus for ASD in Spanish. It was constructed using the LIP-RTVE database (Gimeno-Gómez and Martínez-Hinarejos, 2022), where multiple challenging in-the-wild situations can be found, as it was collected from TV newscast programs. Due to the fact that LIP-RTVE was conceived to address audio-visual speech recognition, all their samples include a person talking during the entire video clip. Therefore, in order to estimate a model focused on the ASD task, we defined different types of samples, intending to cover all the possible situations we might find when applying our model in a realistic application scenario.

⁵<https://machinelearningmastery.com/threshold-moving-for-imbalanced-classification/>

Positive Samples. We consider a positive sample when audio and video correspond to each other, i.e., the voice is aligned with the talking face.

Negative Samples. With the intention of providing the model with a certain robustness against all types of undesired situations that might occur in a realistic application scenario, we defined three different types of negative samples:

- **Temporal Mismatch.** The audio and video cues correspond to each other. However, the audio was randomly shifted to simulate a temporal mismatch, as long as the overlap between both modalities did not exceed 50%. This type of sample is aimed to make the model discard those scenes where audio and video cues are not aligned.
- **Partial Speaker Mismatch.** In this case, although the voice and the face belong to the same person, the audio and video cues do not correspond to each other. This type of sample is aimed to ensure that the model does not learn to simply associate a voice with its corresponding face.
- **Complete Speaker Mismatch.** The voice and the face belong to different people. This type of sample was aimed to provide the model examples of situations we can find in a real scenario, such as dubbing or voiceovers.

Specifically, we split the corpus into training, validation, and test sets, ensuring a balance between all types of samples previously described. More details about the data sets defined for the ASD-RTVE database can be found in Table 2.

Dataset	No. of Speakers			No. of Utterances
	Female	Male	Total	
Training	19	10	29	100k
Validation	65	86	151	30k
Test	76	67	143	30k
Total	160	163	323	160k

Table 2: Details of the ASD-RTVE corpus.

It should be noted that, if we do not have a database from which to generate data to train an

ASD model for our language of interest, we can always use public sources to compile it, e.g., a set of vlogs where a single person is usually talking in front of the camera during the entire video clip can be used to obtain raw audio-visual data, and then, by following, the procedure described above, the ASD-oriented dataset could be created. In the worst-case scenario, an alternative approach would be to use an ASD module trained for English, for instance, and then set its decision threshold to 0.0. We are aware that this may entail more effort from the annotator’s point of view, but we still believe the toolkit can speed up the process.

3.3. Implementation Details

Experiments were conducted on a GeForce RTX 4090 GPU with 24GB memory. The design of our ASD model was based on the official implementation⁶ of the TalkNet-ASD model.

Training. In all our experiments, we used the Adam optimizer (Kingma and Ba, 2014) with a learning rate of 0.0001 decaying through a linear scheduler across 9 epochs, with a batch size of 32 samples. Because we fine-tuned the TalkNet-ASD model pre-trained for the AVA-ActiveSpeaker database (Roth et al., 2020), its original architecture was not modified, referring readers to Tao et al. (2021) to a more detailed description. Nonetheless, for the rest of the hyper-parameters, multiple settings were explored until we found the optimum we described above.

Evaluation. Similar to other works in the literature (Tao et al., 2021; Liao et al., 2023; Min et al., 2022), we reported our results in terms of different evaluation metrics, namely the accuracy, the mean Average Precision (mAP), and the Area Under the Curve (AUC).

3.4. Results & Discussion

Similar to Tao et al. (2021), we studied how the size of the context window was affecting our system performance. Table 3 illustrates a clear positive correlation on the validations set. As the number of context seconds increased, the ASD model showed a consistent upward trend, suggesting that incorporating more contextual information leads to more accurate and reliable predictions regarding speech activity detection.

Furthermore, we evaluated the ASD models estimated for each window size in terms of inference time. However, no significant differences were found between the studied context windows. Therefore, for AnnoTheia, we used the ASD model trained with 51-frame windows, which was able to achieve, on the ASD-RTVE test set, results around 95.4±0.2 accuracy, 91.6 % mAP, and 99.3 % AUC.

⁶<https://github.com/TaoRuijie/TalkNet-ASD>

Context Window		Accuracy	mAP (%)	AUC (%)
no. of frames	no. of seconds			
5	0.20	70.2±0.5	60.0	75.3
9	0.36	79.2±0.5	72.0	84.2
13	0.52	83.5±0.4	77.7	88.0
17	0.68	85.7±0.4	78.4	91.2
21	0.84	89.9±0.3	85.2	94.8
25	1.00	91.8±0.3	87.5	95.8
35	1.40	93.1±0.3	90.7	98.5
43	1.72	97.0±0.2	94.8	99.3
51	2.04	97.3±0.2	95.8	99.5

Table 3: Results for the ASD-RTVE validation set depending on the context window size.

4. Conclusions & Future Work

In this work, we presented the AnnoTheia toolkit to facilitate the task of annotating audio-visual resources for speech technologies, thus promoting further research on low-resource languages in the field. In addition, because some of its modules are language-dependent, we showed the complete process of preparing AnnoTheia for a language of interest, even in those occasions where there are no databases originally conceived for that purpose. Regarding our future work, we are considering introducing further refinements to improve the toolkit both from a performance and usability perspective. One important contribution would be integrating a module capable of ensuring speaker-independent partitions, e.g., via face, body, and voice clustering (Brown et al., 2021; Deng et al., 2019). Furthermore, we plan to conduct a comprehensive experiment to properly assess the effectiveness and advantages of using the AnnoTheia toolkit versus a completely manual procedure.

5. Limitations

Our toolkit’s main limitation is the need to adapt the ASD module to the language of interest, which could hinder the experience for users unfamiliar with training machine learning-based systems. Another aspect to consider is the possibility that the pre-trained models composing the toolkit can be biased. Due to the lack of demographically diverse data, certain social groups could be under-represented because of their age, gender, and cultural background, among other factors.

6. Ethical Considerations

The toolkit we present can be used to collect and annotate audio-visual databases from social media where a person’s face and voice are recorded. Both information cues could be considered as biometrics (Nagrani et al., 2017), raising privacy concerns regarding a person’s identity. Therefore, any data collection must protect privacy rights, taking necessary measures and always asking all the people involved for permission beforehand.

7. Acknowledgements

The work of José-M. Acosta-Triana was in the framework of the Valencian Graduate School and Research Network for Artificial Intelligence (ValgrAI) funded by Generalitat Valenciana. The work of David Gimeno-Gómez and Carlos-D. Martínez-Hinarejos was partially supported by Grant CIACIF/2021/295 funded by Generalitat Valenciana and by Grant PID2021-124719OB-I00 under project LLEER (PID2021-124719OB-100) funded by MCIN/AEI/10.13039/501100011033/ and by ERDF, EU A way of making Europe.

8. Bibliographical References

- A. Ablimit, C. Botelho, A. Abad, T. Schultz, and I. Trancoso. 2022. [Exploring dementia detection from speech: Cross corpus analysis](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6472–6476.
- T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman. 2018. Deep audio-visual speech recognition. *Transactions on PAMI*, 44(12):8717–8727.
- S. Albanie, G. Varol, L. Momeni, H. Bull, T. Afouras, H. Chowdhury, N. Fox, B. Woll, R. Cooper, A. McParland, and A. Zisserman. 2021. BOBSL: BBC-Oxford British Sign Language Dataset.
- Mohamed Anwar, Bowen Shi, Vedanuj Goswami, Wei-Ning Hsu, Juan Pino, and Changhan Wang. 2023. [MuAViC: A Multilingual Audio-Visual Corpus for Robust Speech Recognition and Robust Speech-to-Text Translation](#). In *Proc. INTER-SPEECH*, pages 4064–4068.
- A. Babu, C. Wang, A. Tjandra, K. Lakhotia, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino, A. Baevski, A. Conneau, and M. Auli. 2022. [XLS-R: Self-supervised Cross-lingual Speech Representation Learning at Scale](#). In *Proc. Interspeech*, pages 2278–2282.
- A. Baevski, Y. Zhou, A. Mohamed, and M. Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460.
- H. Bredin and A. Laurent. 2021. [End-To-End Speaker Segmentation for Overlap-Aware Re-segmentation](#). In *Proc. Interspeech*, pages 3111–3115.
- L. Bromham, R. Dinnage, H. Skirgård, A. Ritchie, M. Cardillo, F. Meakins, S. Greenhill, and X. Hua. 2022. Global predictors of language endangerment and the future of linguistic diversity. *Nature ecology & evolution*, 6(2):163–173.
- Andrew Brown, Vicky Kalogeiton, and Andrew Zisserman. 2021. Face, body, voice: Video person-clustering with multiple modalities. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 3184–3194.
- A. Bulat and G. Tzimiropoulos. 2017. [How far are we from solving the 2d & 3d face alignment problem? \(and a dataset of 230,000 3d facial landmarks\)](#). In *ICCV*, pages 1021–1030.
- M. Burchi and R. Timofte. 2023. Audio-visual efficient conformer for robust speech recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2258–2267.
- W. Chen, B. Yan, J. Shi, Y. Peng, S. Maiti, and S. Watanabe. 2023. [Improving massively multilingual asr with auxiliary ctc objectives](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1–5.
- Necati Cihan Camgöz, Oscar Koller, Simon Hadfield, and Richard Bowden. 2020. [Sign language transformers: Joint end-to-end sign language recognition and translation](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10020–10030.
- B. Denby, T. Schultz, K. Honda, T. Hueber, J. M. Gilbert, and J. S. Brumberg. 2010. Silent speech interfaces. *Speech Communication*, 52(4):270–287.
- J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou. 2020. Retinaface: Single-shot multi-level face localisation in the wild. In *CVPR*, pages 5202–5211.
- Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. 2019. [Arcface: Additive angular margin loss for deep face recognition](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4685–4694.
- David M. Eberhard, Gary F. Simons, and Charles D. Fennig. 2023. Ethnologue: languages of the world. *Online version: <http://www.ethnologue.com>*.
- Y. Fujita, N. Kanda, S. Horiguchi, K. Nagamatsu, and S. Watanabe. 2019. [End-to-End Neural Speaker Diarization with Permutation-Free Objectives](#). In *Proc. Interspeech*, pages 4300–4304.

- M. Gales and S. Young. 2008. *The application of hidden Markov models in speech recognition*. Now Publishers Inc.
- J. A. González-López, A. Gómez-Alanís, J. M. Martín Doñas, J. L. Pérez-Córdoba, and A. M. Gómez. 2020. Silent speech interfaces for speech restoration: A review. *IEEE Access*, 8:177995–178021.
- Jianbo Guo, Yuxi Li, Weiyao Lin, Yurong Chen, and Jianguo Li. 2018. Network decoupling: From regular to depthwise separable convolutions. In *British Machine Vision Conference*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Jie Hu, Li Shen, and Gang Sun. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141.
- R. Huang, H. Liu, X. Cheng, Y. Ren, L. Li, Z. Ye, J. He, L. Zhang, J. Liu, X. Yin, and Z. Zhao. 2023. AV-TranSpeech: Audio-visual robust speech-to-speech translation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8590–8604.
- Zili Huang, Shinji Watanabe, Shu-wen Yang, Paola García, and Sanjeev Khudanpur. 2022. Investigating self-supervised learning for speech enhancement and separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6837–6841.
- Y. Jia, R. Weiss, F. Biadsy, W. Macherey, M. Johnson, Z. Chen, and Y. Wu. 2019. Direct Speech-to-Speech Translation with a Sequence-to-Sequence Model. In *Proc. Interspeech*, pages 1123–1127.
- Kristina Jokinen and Michael McTear. 2022. *Spoken dialogue systems*. Springer Nature.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- O. Koller, J. Forster, and H. Ney. 2015. Continuous sign language recognition: Towards large vocabulary statistical recognition systems handling multiple signers. *Computer Vision and Image Understanding*, 141:108–125.
- A. Lee, P. Chen, C. Wang, J. Gu, S. Popuri, X. Ma, A. Polyak, Y. Adi, Q. He, Y. Tang, J. Pino, and W. Hsu. 2022. Direct speech-to-speech translation with discrete units. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pages 3327–3339.
- J. Li, Y. Wang, C. Wang, Y. Tai, J. Qian, J. Yang, C. Wang, J. Li, and F. Huang. 2019a. Dsfd: Dual shot face detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5055–5064.
- N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu. 2019b. Neural speech synthesis with transformer network. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, page 6706–6713.
- J. Liao, H. Duan, K. Feng, W. Zhao, Y. Yang, and L. Chen. 2023. A light weight model for active speaker detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22932–22941.
- X. Liu, E. Lakomkin, K. Vougioukas, P. Ma, H. Chen, R. Xie, M. Doulaty, N. Moritz, J. Kolar, S. Petridis, M. Pantic, and C. Fuegen. 2023. Synthvsr: Scaling up visual speech recognition with synthetic supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18806–18815.
- P. Ma, S. Petridis, and M. Pantic. 2021. End-to-end audio-visual speech recognition with conformers. In *ICASSP*, pages 7613–7617.
- P. Ma, S. Petridis, and M. Pantic. 2022. Visual speech recognition for multiple languages in the wild. *Nature Machine Intelligence*, 4(11):930–939.
- Manuel Milling, Florian B Pokorny, Katrin D Bartl-Pokorny, and Björn W Schuller. 2022. Is speech the new blood? recent progress in ai-based disease detection from audio in a nutshell. *Frontiers in Digital Health*, 4.
- K. Min, S. Roy, S. Tripathi, T. Guha, and S. Majumdar. 2022. Learning long-term spatial-temporal graphs for active speaker detection. In *ECCV*, pages 371–387, Cham. Springer Nature Switzerland.
- A. Nagrani, J. S. Chung, and A. Zisserman. 2017. *VoxCeleb: A Large-Scale Speaker Identification*

- [Dataset](#). In *Proc. Interspeech*, pages 2616–2620.
- S. Pascual, A. Bonafonte, and J. Serrà. 2017. SEGAN: Speech Enhancement Generative Adversarial Network. In *Proc. Interspeech*, pages 3642–3646.
- KR Prajwal, T. Afouras, and A. Zisserman. 2022. Sub-word level lip reading with visual attention. In *CVPR*, pages 5162–5172.
- V. Pratap, A. Tjandra, B. Shi, P. Tomasello, A. Babu, S. Kundu, A. Elkahky, Z. Ni, A. Vyas, M. Fazel-Zarandi, A. Baevski, Y. Adi, Xiaohui Z., W. Hsu, A. Conneau, and Auli M. 2023. Scaling speech technology to 1,000+ languages. *arXiv preprint arXiv:2305.13516*.
- A. Radford, J. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*, pages 28492–28518. PMLR.
- M. Ravanelli, J. Zhong, S. Pascual, P. Swietojanski, J. Monteiro, J. Trmal, and Y. Bengio. 2020. [Multi-task self-supervised learning for robust speech recognition](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6989–6993.
- J. Shen, R. Pang, R. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, R. Saurous, Y. Agiomvrgiannakis, and Y. Wu. 2018. Natural tts synthesis by conditioning wavenet on mel spectrogram predictions. In *2018 IEEE international conference on acoustics, speech and signal processing*, pages 4779–4783. IEEE.
- Bowen Shi, Wei-Ning Hsu, Kushal Lakhota, and Abdelrahman Mohamed. 2022. Learning audio-visual speech representation by masked multimodal cluster prediction. *arXiv preprint arXiv:2201.02184*.
- J. Shi, D. Berrebbi, W. Chen, H. Chung, E. Hu, W. Huang, X. Chang, S. Li, A. Mohamed, H. Lee, and S. Watanabe. 2023. MI-superb: Multilingual speech universal performance benchmark. *arXiv preprint arXiv:2305.10615*.
- R. Tao, Z. Pan, R. Das, X. Qian, M. Shou, and h. Li. 2021. [Is someone speaking? exploring long-term temporal features for audio-visual active speaker detection](#). In *Proceedings of the 29th ACM International Conference on Multimedia*, MM '21, page 3927–3935. Association for Computing Machinery.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N Gomez, L. Kaiser, and I. Polosukhin. 2017. Attention is all you need. *NeurIPS*, 30:6000–6010.
- H. Yen, F. Germain, G. Wichern, and J. Roux. 2023. Cold diffusion for speech enhancement. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1–5.
- Y. Zhang, W. Han, J. Qin, Y. Wang, A. Bapna, Z. Chen, N. Chen, B. Li, V. Axelrod, G. Wang, Z. Mengh, K. Hu, A. Rosenberg, R. Prabhavalkar, D. Park, P. Haghani, J. Riesa, G. Perng, H. Soltau, T. Strohmaier, B. Ramabhadran, T. Sainath, P. Moreno, C. Chlu, J. Schalkwyk, F. Beaufays, and Y. Wu. 2023. Google usm: Scaling automatic speech recognition beyond 100 languages. *arXiv preprint arXiv:2303.01037*.

9. Language Resource References

- T. Afouras, J. Chung, and A. Zisserman. 2018a. Lrs3-ted: a large-scale dataset for visual speech recognition. *arXiv preprint arXiv:1809.00496*.
- T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman. 2018b. Deep audio-visual speech recognition. *Transactions on PAMI*, 44(12):8717–8727.
- Mohamed Anwar, Bowen Shi, Vedanuj Goswami, Wei-Ning Hsu, Juan Pino, and Changhan Wang. 2023. [MuAViC: A Multilingual Audio-Visual Corpus for Robust Speech Recognition and Robust Speech-to-Text Translation](#). In *Proc. Interspeech*, pages 4064–4068.
- R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber. 2020. Common voice: A massively-multilingual speech corpus. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4218–4222.
- J. Chung and A. Zisserman. 2017. Lip reading in the wild. In *13th Asian Conference on Computer Vision*, pages 87–103. Springer.
- E. Egorov, V. Kostyumov, M. Konyk, and S. Kolesnikov. 2021. LRWR: large-scale benchmark for lip reading in Russian language. *arXiv preprint arXiv:2109.06692*.
- A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. Freeman, and M. Rubinstein. 2018. [Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation](#). *ACM Trans. Graph.*, 37(4):1–11.

- A. Fernandez-Lopez, O. Martinez, and F. M. Sukno. 2017. [Towards estimating the upper bound of visual-speech recognition: The visual lip-reading feasibility database](#). In *12th FG*, pages 208–215. IEEE.
- A. Fernandez-Lopez and F. Sukno. 2018. Survey on automatic lip-reading in the era of deep learning. *Image and Vision Computing*, 78:53–72.
- D. Gimeno-Gómez and C.-D. Martínez-Hinarejos. 2022. LIP-RTVE: An Audiovisual Database for Continuous Spanish in the Wild. In *LREC*, pages 2750–2758. ELRA.
- N. Harte and E. Gillen. 2015. [TCD-TIMIT: An audio-visual corpus of continuous speech](#). *IEEE Transactions on Multimedia*, 17(5):603–615.
- J. Peymanfard, S. Heydarian, A. Lashini, H. Zeinali, M. Mohammadi, and N. Mozayani. 2023. A multi-purpose audio-visual corpus for multi-modal persian speech recognition: the arman-av dataset. *arXiv preprint arXiv:2301.10180*.
- V. Pratap, A. Tjandra, B. Shi, P. Tomasello, A. Babu, S. Kundu, A. Elkahky, Z. Ni, A. Vyas, M. Fazel-Zarandi, A. Baevski, Y. Adi, Xiaohui Z., W. Hsu, A. Conneau, and Auli M. 2023. Scaling speech technology to 1,000+ languages. *arXiv preprint arXiv:2305.13516*.
- Joseph Roth, Sourish Chaudhuri, Ondrej Klejch, Radhika Marvin, Andrew Gallagher, Liat Kaver, Sharadh Ramaswamy, Arkadiusz Stopczynski, Cordelia Schmid, Zhonghua Xi, and Caroline Pantofaru. 2020. AVA Active Speaker: An Audio-Visual Dataset for Active Speaker Detection. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4492–4496. IEEE.
- E. Salesky, M. Wiesner, J. Bremerman, R. Cattoni, M. Negri, M. Turchi, D. Oard, and M. Post. 2021. [The Multilingual TEDx Corpus for Speech Recognition and Translation](#). In *Proc. Interspeech*, pages 3655–3659.
- S. Yang, Y. Zhang, D. Feng, M. Yang, C. Wang, J. Xiao, K. Long, S. Shan, and X. Chen. 2019. [LRW-1000: A Naturally-Distributed Large-Scale Benchmark for Lip Reading in the Wild](#). In *14th IEEE International Conference on Automatic Face & Gesture Recognition*, pages 1–8.
- A. B. Zadeh, Y. Cao, S. Hessner, P. P. Liang, S. Poria, and L.-P. Morency. 2020. CMU-MOSEAS: A multimodal language dataset for spanish, portuguese, german and french. In *EMNLP*, pages 1801–1812.
- Y. Zhao, R. Xu, and M. Song. 2020. [A cascade sequence-to-sequence model for chinese mandarin lip reading](#). In *Proceedings of the ACM Multimedia Asia*, pages 1–6. Association for Computing Machinery.