

# The Importance of Facial Features in Vision-based Sign Language Recognition: Eyes, Mouth or Full Face?

Dinh Nam Pham

German Research Center for Artificial Intelligence (DFKI)  
Berlin, Germany  
dinh\_nam.pham@dfki.de

Eleftherios Avramidis

German Research Center for Artificial Intelligence (DFKI)  
Berlin, Germany  
eleftherios.avramidis@dfki.de

## Abstract

Non-manual facial features play a crucial role in sign language communication, yet their importance in automatic sign language recognition (ASLR) remains underexplored. While prior studies have shown that incorporating facial features can improve recognition, related work often relies on hand-crafted feature extraction and fails to go beyond the comparison of manual features versus the combination of manual and facial features. In this work, we systematically investigate the contribution of distinct facial regions—eyes, mouth, and full face—using two different deep learning models (a CNN-based model and a transformer-based model) trained on an SLR dataset of isolated signs with randomly selected classes. Through quantitative performance and qualitative saliency map evaluation, we reveal that the mouth is the most important non-manual facial feature, significantly improving accuracy. Our findings highlight the necessity of incorporating facial features in ASLR.

## CCS Concepts

• **Human-centered computing** → **Accessibility design and evaluation methods; Accessibility technologies**; • **Computing methodologies** → Interest point and salient region detections; Activity recognition and understanding; • **Applied computing** → Language translation.

## Keywords

Sign Language, Sign Language Recognition, Video Classification, Saliency Map, German Sign Language, Deep Learning, Computer Vision

### ACM Reference Format:

Dinh Nam Pham and Eleftherios Avramidis. 2025. The Importance of Facial Features in Vision-based Sign Language Recognition: Eyes, Mouth or Full Face?. In *ACM International Conference on Intelligent Virtual Agents (IVA Adjunct '25)*, September 16–19, 2025, Berlin, Germany. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3742886.3756718>

## 1 Introduction

Sign language (SL) serves as a vital means of communication for deaf and hard-of-hearing communities worldwide. However, similar to spoken languages, sign language is not universal. Different countries and regions have their own native sign languages, each with

unique grammatical and linguistic structures. Furthermore, there are significant communication barriers between sign language users and those unfamiliar with SL. In an effort to address this, automatic sign language recognition (ASLR) has gained growing attention in recent years, with the number of available datasets and studies increasing.

As a visual language, sign language conveys information through multiple channels, which can be broadly categorized into (a) manual and (b) non-manual features. Manual features are related to hands and include hand shapes, palm orientation, hand location, and hand movement. In contrast, non-manual features describe expressions which do not involve the hands, such as movements of the mouth, cheeks, head, shoulders, and eyebrows. Sign languages have complex grammars that include non-manual markers, making them essential for understanding SL [10]. Facial non-manual features are an essential part of SL grammar, not just visual cues. They serve core linguistic functions, with upper-face markers being particularly essential for syntactic and prosodic structures across different sign languages, while the mouth plays a crucial role for the lexical and morphological levels [22]. For example, eye gaze can mark agreement in relation to agreement verbs, while facial expressions like sucked-in or blown-out cheeks modify signs such as HOUSE to indicate size [10]. In German Sign Language (DGS), a head nod can mark past tense, and mood is often expressed through facial expressions. Additionally, some signs, like TRAURIG (sad) in DGS, require specific facial expressions to convey their meaning accurately [10]. Mouth actions can also help disambiguate signs with ambiguous manual features [25].

Despite this crucial role of non-manual features, only 6% of sign language recognition (SLR) results from 2015 to 2020 made use of the head, while 5% utilized mouth features, and 3% incorporated the eyes [12]. While recent advances in deep learning have significantly improved manual sign recognition, the extent to which facial features contribute to recognition performance remains underexplored. Given their importance in disambiguating signs and conveying grammatical information, a deeper understanding of their role in vision-based SLR is crucial. Unlike methods that rely on keypoints, skeleton tracking, or handcrafted features extracted from pose estimation models, our approach focuses on end-to-end deep learning-based SLR, where models learn directly from raw video input without performing explicit feature engineering. Additionally, we decided against approaches that depend on external modalities such as depth sensors, wearable gloves, or intrusive tracking devices, making our approach more signer-centric, practical and adaptable to real-world applications.

*IVA Adjunct '25, Berlin, Germany*

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM International Conference on Intelligent Virtual Agents (IVA Adjunct '25)*, September 16–19, 2025, Berlin, Germany, <https://doi.org/10.1145/3742886.3756718>.

In this study, we systematically investigate the contribution of different facial regions—eyes, mouth, and full face—to SLR performance. By analyzing classification accuracy and saliency maps on an isolated sign language dataset, we aim to determine which facial features are most informative and relevant for deep learning models in video-based SLR. Furthermore, we explore whether combining facial and manual features improves recognition performance, providing insights into the integration of non-manual signals in ASLR systems.

## 2 Related Work

While there have been works that specifically modeled the facial features as a parameter for SLR, the overwhelming majority rely on full-frame inputs [12, 27]. In such cases, models may implicitly learn to utilize facial cues, but the specific contribution of these features to overall performance remains understudied. For this, studies comparing SLR performance between using solely manual features and combining manual features with facial features could help to understand their significance better.

In [16], the authors combine hand gestures with facial expressions, which slightly outperforms the model relying only on hands. They extract 3D feature points from the hands and face using Leap Motion and Kinect sensors, then employ a hidden Markov model for each modality to recognize signs. The decisions from both modalities are combined using a Bayesian classifier. This approach achieves relatively small recognition rate gains, but its reliance on sensors limits its practicality for real-world scenarios and makes it less comparable to methods that use only RGB video data.

As one of the earliest studies to address the significance of facial features in ASLR, [31] reported improved recognition performance when combining facial and manual features for both isolated and continuous SLR. The authors extracted hand and facial features as coordinates, representing the face as a 16-dimensional vector and the hands as a 22-dimensional vector.

Meanwhile, the authors of [21] conducted a more fine-grained quantitative comparison. They evaluated not only manual features versus the combination of manual and facial features, but also examined specific subsets of facial features: 'Manual & only eyebrows,' 'Manual & only mouth,' 'Manual & eyebrows, mouth,' and 'Manual & face, eyebrows, mouth.' Using OpenPose [1], keypoints were extracted and classified with logistic regression. However, the dataset was explicitly designed to consist of signs that are similar in manual articulation and are distinguished by using non-manual features, leading to non-manual features achieving a mean accuracy of 77%, compared to 73.4% for manual features alone. The fact that the dataset solely consisted of cases where non-manual features play a decisive role by design, diminishes the effect of highlighting their importance in general.

Similarly, [25] addressed the function of the mouth to disambiguate such homonyms, based on a dataset of pairs with identical or similar manual signs, but different meanings. The authors demonstrated that the combination of mouth and manual features outperform the model using only manual features. Unlike the other works described before, this approach avoided hand-crafted feature extraction, using a deep learning model instead to classify the video streams end-to-end. Each video stream was processed by a

separate instance of the same model, with hand and mouth features combined in a late fusion fashion.

Although not primarily aimed to address the importance of facial features in SLR, a deep learning model was trained in [23] with full-frame inputs. The results interestingly demonstrated that omitting the face in the full-frame significantly decreased the accuracy whereas the utilization of face swapping as a data augmentation method lead to accuracy gains, further underscoring the impact of facial features in SLR.

In essence, previous studies have primarily evaluated the impact of combining manual and facial features. They often relied on hand-crafted features, though deep learning approaches were also used. While all works reported accuracy gains from incorporating facial features, the extent of improvement varied, particularly in datasets designed to emphasize non-manual features. Contrary to that, in this work, we construct an isolated SLR dataset with randomly selected classes, independent of their co-occurrence with non-manual articulation. We focus on deep learning models, which better reflect the currently commonly used methods [27], and conduct a detailed investigation into the contribution of the eyes, mouth, and full face. By quantitatively evaluating their integration with manual features and qualitatively analyzing saliency maps, we offer new insights into the importance of facial features in ASLR.

As this study does not include direct comparisons to other sign language recognition methods, our objective is not to achieve state-of-the-art accuracy, but to provide a focused and controlled analysis of the contribution of facial regions using modern vision models. Most prior works that aim to highlight the significance of non-manual features combine manual and facial features without isolating specific facial regions, or rely on hand-crafted features and sensor data—methods that are increasingly obsolete given the current dominance of end-to-end deep learning SLR models for RGB video analysis. In contrast, our approach uses state-of-the-art deep learning backbones under a consistent training setup, enabling a reproducible and architecture-agnostic evaluation of non-manual feature importance. We believe that these controlled conditions offer a complementary perspective, and our findings can inform the integration of facial cues into future multimodal SLR systems.

## 3 Dataset

In order to assess the importance of facial features in ASLR, a dataset is required to train and evaluate models. For this purpose, the Public DGS Corpus [13] was selected to create a dataset for isolated sign language recognition with glosses as class labels. The Public DGS Corpus is well-suited for this task due to its comprehensive annotations, diversity of speakers as well as the quality and quantity of video recordings. Data, including 550 hours of DGS signing, were collected from 330 signers across 12 different locations in Germany, ensuring a balanced representation of age, gender, and regional variation among participants [9]. Using the gloss annotations, we randomly selected 12 glosses, each occurring at least 500 times, to ensure a sufficient number of instances for robust model training. Using the timestamps associated with these annotations, we extracted the corresponding video clips from the recordings. Each instance in the dataset consists of a video clip displaying a signer performing the sign for a single gloss. In our dataset, the glosses




Figure 1: Class distribution in the dataset.

serve as the class labels, while the individual video clips represent the instances within each class.

The class distribution is illustrated in Figure 1. The class with the highest amount of instances contains 989 instances, while the gloss with the fewest instances includes 518 videos. Given that this imbalance ratio is less than 1:2 and considering that most studies on class imbalance focus on ratios starting from 1:4 [15], we consider the dataset to not be severely imbalanced. As a result, no additional measures were taken to adjust the class distribution.

In accordance with the annotation standard of the Public DGS corpus [14], the names of the labels begin with a gloss, i.e. a word of the spoken language that represents the core meaning of a specific sign. This is followed by a number that describes a lexical variant and a letter denoting a phonological variant. An optional asterisk indicates that the observed form of the sign differs slightly from its standard or citation form.

We split the dataset into training, validation and test sets using an 8:1:1 ratio while preserving the class distribution in each set. To address variations in video frame length, we normalized the videos by padding them with repetitions of the last frame until each video reached a uniform length of 32 frames. To evaluate the importance of different facial regions, we constructed four versions of the dataset by cropping each video to specific regions of interest (ROIs): the eyes, mouth, full face, and body. For cropping the mouth region, we adopted the implementation from [19]. The face region was extracted using the Face Recognition library<sup>1</sup>, while MediaPipe [18] was employed to detect the coordinates of the eyes, upper body, and hands for cropping. The upper body ROI was defined by determining the outermost coordinates from the detected upper body and hand landmarks, ensuring that both the body and hands were always fully visible. Finally, the videos of the ROIs were resized to 224x224 pixels since this is a common image size for pre-trained vision models.

## 4 Experiments

For the ASLR, we leveraged two distinct state-of-the-art architectures: Channel-Separated Convolutional Networks (CSN) [30] and Multiscale Vision Transformers (MVIT) [5]. More specifically, we used the CSN R101 and MVIT 32x3 model implementations from PyTorchVideo [4], which were pretrained on the Kinetics-400 [11] dataset under the same frame length and image size setting as the

<sup>1</sup>Adam Geitgey, *Face Recognition*, GitHub. [Online] Available: [https://github.com/ageitgey/face\\_recognition](https://github.com/ageitgey/face_recognition).


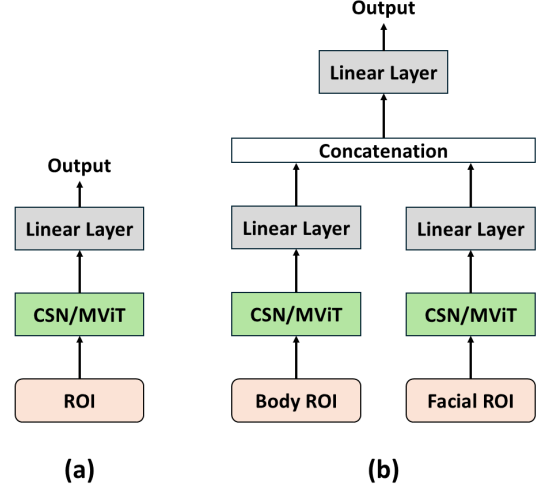


Figure 2: An example of the dataset [13] being cropped to the regions of interest: (a) eyes, (b) mouth, (c) face, and (d) body.

Figure 3: Model architectures for (a) single ROI and (b) two ROIs (in this example *body + facial*).

items in our dataset. The CSN architecture factorizes 3D convolutions, whereas MVIT integrates a multiscale feature hierarchy with the transformer architecture.

The reason we chose these models is that they differ fundamentally from each other in their design principles (convolutional vs. transformer-based), allowing us to assess whether observed trends generalize across architectures and gain complementary insights into the significance of different regions of interest. By comparing the performance and feature utilization of both architectures, we can identify consistent patterns in how facial features contribute to sign language recognition, leading to conclusions about the role of facial non-manual signals in vision-based models that can be generalized across the current state of the art. The use of the two models for SLR was also supported by the fact that they have demonstrated state-of-the-art performance on human action recognition, as they have outperformed successful and commonly used SLR architectures, such as I3D and R(2+1)D [6–8, 17, 23, 26], on the Kinetics400 benchmark [11].

We trained both CSN and MVIT on each of the four ROIs. For both architectures, the output layer was replaced with a linear layer with an output size of 12, corresponding to the number of class labels. To evaluate the impact of facial features on SLR with manual signs, we combined the body ROI with each facial ROI (eyes, mouth, and face). For each model  $M \in \{\text{CSN}, \text{MVIT}\}$  and each facial ROI  $R_F \in \{\text{Eyes}, \text{Mouth}, \text{Face}\}$ , we constructed a new model where

Region of Interest	CSN			MViT		
	Top-1 Acc.	Top-3 Acc.	F1-Score	Top-1 Acc.	Top-3 Acc.	F1-Score
Eyes	20.59 ± 2.97	44.82 ± 3.65	0.1853	22.41 ± 3.06	46.36 ± 3.66	0.2006
Mouth	48.18 ± 3.67	76.19 ± 3.12	0.4864	51.40 ± 3.67	77.31 ± 3.07	0.5067
Face	60.64 ± 3.58	82.07 ± 2.81	0.6148	59.66 ± 3.60	79.55 ± 2.96	0.6002
Body	80.53 ± 2.90	94.96 ± 1.60	0.8164	84.03 ± 2.69	93.42 ± 1.82	0.8496
Body + Eyes	82.77 ± 2.77	95.24 ± 1.56	0.8400	82.91 ± 2.76	94.54 ± 1.67	0.8485
Body + Mouth	<b>88.24</b> ± 2.36	97.20 ± 1.21	<b>0.8890</b>	86.42 ± 2.51	<b>97.76</b> ± 1.09	0.8735
Body + Face	87.26 ± 2.45	97.62 ± 1.12	0.8790	86.98 ± 2.47	96.22 ± 1.40	0.8767

**Table 1: Top-1 accuracy, top-3 accuracy and F1-Score for CSN and MViT across different regions of interest.**

one instance of  $M$  processes  $R_F$  and another processes the body ROI. The outputs of both instances were fused and propagated to a linear output layer. This late-fusion approach integrates facial non-manual features with manual signals, enabling observation to what extent they improve the performance. This architecture is illustrated in Figure 3.

For all experiments, the same hyperparameters and settings were employed. We used the cross-entropy loss, the Adam optimizer, a learning rate of  $10^{-5}$ , and a batch size of 3 due to limited memory capacity. Additionally, RandAugment [2] with a magnitude of 4 and  $N = 3$  was applied as a data augmentation technique to each video in the training set before it was fed into a model. The models were trained on the training set and validated on the validation set for 100 epochs. The weights yielding the highest top-1 accuracy on the validation set were used for final evaluation on the test set.

## 5 Results and Discussion

### 5.1 Classification Performance

We report the experimental results in Table 1, including accuracies with a confidence interval of 95% and F1-scores for both CSN and MViT. It is striking that CSN and MViT perform relatively similarly across all metrics and ROIs. Unsurprisingly, the eyes on their own achieved the lowest accuracies. However, with top-1 accuracies of 20.59% and 22.41%, it seems to contain useful information to some extent. The full face ROI strongly outperforms the other facial ROIs, which makes sense, since it contains both these facial ROIs and other features such as head pose and facial expressions as well. The body ROI, demonstrating the best performance across all ROIs, when using only one video stream, was also expected.

The accuracies obtained by combining eyes and body are statistically not significantly different from the models using only the body as the confidence intervals overlap. This suggests that while the eyes can provide useful features, their contribution becomes relatively insignificant when manual features are present—at least in this experimental setting. In contrast to that, the incorporation of the mouth as well as the face significantly improves the top-1 accuracy of CSN with the body ROI. However, for the MViT model, as its top-1 accuracy for the body ROI was already relatively high, a statistically significant difference cannot be observed for the "body + mouth" and "body + face" ROIs. Only MViT's top-3 accuracy was significantly improved by adding the mouth.

Interestingly, while the face outperformed the mouth ROI when using them on their own, the performance difference between them, when combined with the body, becomes insignificant. With overlapping confidence intervals across all accuracies and models, the results do not indicate a clear distinction between the two ROIs when the manual signals are present. This may indicate that *the mouth area is the most important facial feature*.

Furthermore, even though the addition of mouth and face significantly improves the top-1 accuracy of CSN compared to solely using the body, they fail to make a significant difference for the top-3 accuracy. This suggests that *the facial features can make a difference for vision-based models to better distinguish between top 3 candidates*. This aligns with the known functionality of non-manual signals, such as mouthing, to disambiguate between signs that share the same manual articulation but differ in meaning [25].

In an effort to explain why the inclusion of the mouth ROI increases the performance when modeled together with the body ROI, we use some linguistic observations. The signs for 6 of our classes (ABER, AUCH, SELBST, MUSS, NICHT, WENN) are based on the same handform (extended index finger) but with different orientation, position or movement. An assumption is that from a computer vision perspective, the respective video clips are similar (e.g. share identical frames when taken from different angles), so that the mouthings, which depict different vowels, help the models make the distinction.

### 5.2 Saliency Maps

In addition to the quantitative results, we generated saliency maps to visualize the most relevant regions for both models, providing further insights into the importance of facial features in a qualitative manner. For this, we computed vanilla gradient saliency maps [28] for videos from each class and applied SmoothGrad [29] with a sample size of 50 to reduce noise in the pixel attributions. Our analysis focused on two regions of interest: (1) the face alone and (2) the combination of the face and body. This approach allowed us to identify the most critical facial regions and examine how their relevance changes when contextual information from the body is included, which is of particular interest in this study. Examples are shown in Figure 4, which illustrates the first, fourth, seventh, tenth, and 13th frames of a single instance for a few classes. The top 0.5% attribution values were set to 1 while the remaining values were normalized relative to these outliers.



Figure 4: Saliency maps for some videos in the test set [13], visualising the models for the face as well as the "body + face" fusion.

Throughout all classes, the saliency maps of MViT appear to be much more fine-grained compared to those of CSN. This likely stems from the underlying architecture differences between the two models. MViT, with its hierarchical attention mechanism, is better equipped to capture localized and detailed spatial-temporal dependencies, resulting in finer feature attribution. Despite these different granularities, both models highlight the same crucial regions, making the observations more meaningful. For the face ROI, the pixels of the mouth area were consistently highlighted as the most relevant features, reinforcing the quantitative findings that the mouth is the most important facial feature. This aligns with the findings of eye-tracking studies where signers were observed to not only mainly focus on the face during sign language comprehension, but also primarily fixate on the mouth area while doing so [3, 20]. However, some distinctions exist in instances. For example, CSN sometimes highlights broader facial regions, including the cheeks, nose, and jawline, whereas MViT concentrates on finer details, particularly the lips. Moreover, in a noteworthy number of instances, the nose and eyes were significantly marked, speaking for the relevance of these features.

Furthermore, we often observed that the addition of the body leads to a diminished saliency of the face. This effect is particularly evident in the example of the MViT model for the class "DA" as shown in Figure 4 where the face region is nearly absent from the saliency map, shifting attention almost entirely to the hands. When it comes to the body ROI, we noticed that the saliency of the right hand was frequently higher than for the left hand, which aligns with the work of Fink et al. [6] who made the same observation on saliency maps for the I3D model on a French Belgian Sign Language dataset. We share their interpretation that the right hand is of bigger importance for models as most signers use their dominant hand to sign. This bias could lead to lower accuracy for left-handed signers and should be considered in future works by making SL dataset efforts or including horizontal flipping in data augmentation to mitigate this. Furthermore, in cases where the hands appear within the face ROI, such as in the "FRÜHER" class in Figure 4, the models strongly focused on the hands rather than the face. This suggests that even when trained primarily on facial features, the models implicitly recognize the relevance of manual features when they appear in the frame.

While our study mainly focused on facial features, it is important to acknowledge that non-facial non-manual features (e.g., head pose, shoulder movement, and torso orientation) may also play a significant role in ASLR. As our body ROI includes the upper body and hands, it primarily captures manual articulators (i.e., the hands and their movements) along with their spatial positioning relative to the signer's body. As such, signals like subtle shoulder shifts or head tilts are not explicitly disentangled in our analysis. Future work should explore these non-facial non-manual cues more explicitly to better understand their contribution and potential interaction with facial and manual features.

## 6 Conclusion

In this paper, we systematically investigated the role of facial features in vision-based sign language recognition using CSN and MViT models. Through experiments on an isolated SLR dataset, we

demonstrated that incorporating facial features can significantly improve SLR models, aligning with prior works and underscoring the need for future SLR models to integrate non-manual signals. Our approach focuses on state-of-the-art end-to-end deep learning methods as opposed to previous work using manual feature extraction such as keypoints. Our results suggest that the mouth is the most important facial feature when combined with manual features, as there were no statistically significant performance differences when manual features were combined with either the mouth or the full face, across all models and metrics. Additionally, it is shown that the facial features can aid the vision-based models to better distinguish between top candidates. The saliency map analysis further confirmed that the mouth is the most important facial feature, consistently highlighting the mouth as the most relevant facial region. Compared to that, the eyes contributed less significantly, although they appeared to provide some useful cues to some extent. Our findings emphasize the need of incorporating facial non-manual markers—especially mouth actions—into SLR models.

Future research should explore finer-grained non-manual features such as eye gaze, blinks, nose movement, head pose, tongue position, and cheek articulation, with further guidance from linguistic theory observations. The experiments could be expanded to include more models and perform training and evaluation on a number of different SLR datasets, including a larger number of glosses and continuous SLR, to provide even more robust and generalizable findings. Additionally, leveraging transfer learning from related tasks, such as automatic lip reading [24] and eye tracking, could advance the integration of non-manual features in SLR systems.

## Ethical Considerations

In our work, we present experiments on the German Sign Language (DGS), which should be seen and respected as the primary languages of the respective language community. While the overarching objective of this research is to promote equitable access to language technologies for sign language users, the predominance of hearing researchers in NLP entails the risk of developments that are not in accordance with the will of the respective communities, and therefore it is required that every research step takes them in constant consideration. To address this concern, we have included members of the Deaf/deaf and hard-of-hearing communities in our broader research line, as part of the research team, consultants and participants in user studies and workshops, and we have been in co-operation with related unions and communication centers. It should also be noted that our experiments are part of a broader series of research projects, and the results presented here should be by no means considered ready for production nor used as final products without the agreement of the communities. The use of datasets follows their respective licenses and restrictions, and every follow-up work should adhere to those.

## Acknowledgments

The research reported in this paper was supported by BMBF (German Federal Ministry of Education and Research) via the project SocialWear (grant no. 01IW20002).



## References

- [1] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2021. OpenPose: Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43, 1 (2021), 172–186. doi:10.1109/TPAMI.2019.2929257
- [2] Ekin Dogus Cubuk, Barret Zoph, Jon Shlens, and Quoc Le. 2020. RandAugment: Practical Automated Data Augmentation with a Reduced Search Space. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 18613–18624.
- [3] Karen Emmorey, Robin Thompson, and Rachael Colvin. 2009. Eye gaze during comprehension of American Sign Language by native and beginning signers. *J. Deaf Stud. Deaf Educ.* 14, 2 (2009), 237–243.
- [4] Haoqi Fan, Tullie Murrell, Heng Wang, Kalyan Vasudev Alwala, Yanghao Li, Yilei Li, Bo Xiong, Nikhila Ravi, Meng Li, Haichuan Yang, Jitendra Malik, Ross Girshick, Matt Feiszli, Aaron Adcock, Wan-Yen Lo, and Christoph Feichtenhofer. 2021. PyTorchVideo: A Deep Learning Library for Video Understanding. In *Proceedings of the 29th ACM International Conference on Multimedia (Virtual Event, China) (MM '21)*. Association for Computing Machinery, New York, NY, USA, 3783–3786. doi:10.1145/3474085.3478329
- [5] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. 2021. Multiscale Vision Transformers. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. 6804–6815. doi:10.1109/ICCV48922.2021.00675
- [6] Jérôme Fink, Benoît Frénay, Laurence Meurant, and Anthony Cleve. 2021. LSFB-CONT and LSFB-ISOL: Two New Datasets for Vision-Based Sign Language Recognition. In *2021 International Joint Conference on Neural Networks (IJCNN)*. 1–8. doi:10.1109/IJCNN52387.2021.9534336
- [7] Ye Gao, Ruixiang Hu, Tian Ma, Songyi Guo, Yizhou Yang, and Xinlei Zhou. 2022. Dynamic Sign Language Recognition Based on Improved R(2+1)D Algorithm. In *2022 7th International Conference on Image, Vision and Computing (ICIVC)*. 7–15. doi:10.1109/ICIVC55077.2022.9886615
- [8] Xiangzu Han, Fei Lu, Jianqin Yin, Guohui Tian, and Jun Liu. 2022. Sign Language Recognition Based on R(2+1)D With Spatial–Temporal–Channel Attention. *IEEE Transactions on Human-Machine Systems* 52, 4 (2022), 687–698. doi:10.1109/THMS.2022.3144000
- [9] Thomas Hanke, Marc Schulder, Reiner Konrad, and Elena Jahn. 2020. Extending the Public DGS Corpus in Size and Depth. In *Proceedings of the LREC2020 9th Workshop on the Representation and Processing of Sign Languages: Sign Language Resources in the Service of the Language Community, Technological Challenges and Application Perspectives*, Eleni Efthimiou, Stavroula-Evita Fotinea, Thomas Hanke, Julie A. Hochgesang, Jette Kristoffersen, and Johanna Mesch (Eds.). European Language Resources Association (ELRA), Marseille, France, 75–82. <https://www.sign-lang.uni-hamburg.de/lrec/pub/20016.pdf>
- [10] Annika Herrmann. 2014. *Modal and Focus Particles in Sign Languages: A Cross-Linguistic Study*. De Gruyter Mouton, Berlin, Boston. doi:10.1515/9781614511816
- [11] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. 2017. The Kinetics Human Action Video Dataset. *arXiv preprint arXiv:1705.06950* (2017).
- [12] Oscar Koller. 2020. Quantitative Survey of the State of the Art in Sign Language Recognition. *arXiv preprint arXiv:2008.09918* (2020).
- [13] Reiner Konrad, Thomas Hanke, Gabriele Langer, Dolly Blanck, Julian Bleicken, Ilona Hofmann, Olga Jeziorski, Lutz König, Susanne König, Rie Nishio, Anja Regen, Uta Salden, Sven Wagner, Satu Worbeck, Oliver Böse, Elena Jahn, and Marc Schulder. 2020. MEINE DGS – annotiert. Öffentliches Korpus der Deutschen Gebärdensprache, 3. Release / MY DGS – annotated. Public Corpus of German Sign Language, 3rd release. doi:10.25592/dgs.corpus-3.0
- [14] Reiner Konrad, Thomas Hanke, Gabriele Langer, Susanne König, Lutz König, Rie Nishio, and Anja Regen. 2022. *Öffentliches DGS-Korpus: Annotationskonventionen / Public DGS Corpus: Annotation Conventions* (4.1 ed.). Project Note AP03-2018-01. DGS-Korpus project, IDGS, Hamburg University, Hamburg, Germany. doi:10.25592/uhhfdm.822
- [15] Bartosz Krawczyk. 2016. Learning from imbalanced data: open challenges and future directions. *Progress in Artificial Intelligence* 5, 4 (April 2016), 221–232. doi:10.1007/s13748-016-0094-0
- [16] Pradeep Kumar, Partha Pratim Roy, and Debi Prosad Dogra. 2018. Independent Bayesian classifier combination based sign language recognition using facial expression. *Information Sciences* 428 (2018), 30–48. doi:10.1016/j.ins.2017.10.046
- [17] Dongxu Li, Cristian Rodriguez-Opazo, Xin Yu, and Hongdong Li. 2019. Word-level Deep Sign Language Recognition from Video: A New Large-scale Dataset and Methods Comparison. *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)* (2019), 1448–1458.
- [18] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, Wan-Teh Chang, Wei Hua, Manfred Georg, and Matthias Grundmann. 2019. MediaPipe: A Framework for Building Perception Pipelines. *arXiv preprint arXiv:1906.08172* (2019).
- [19] Pingchuan Ma, Stavros Petridis, and Maja Pantic. 2022. Visual speech recognition for multiple languages in the wild. *Nature Machine Intelligence* 4, 11 (Oct. 2022), 930–939. doi:10.1038/s42256-022-00550-z
- [20] Eliana Mastrantuono, David Saldaña, and Isabel R. Rodríguez-Ortiz. 2017. An Eye Tracking Study on the Perception and Comprehension of Unimodal and Bimodal Linguistic Inputs by Deaf Adolescents. *Frontiers in Psychology* Volume 8 - 2017 (2017). doi:10.3389/fpsyg.2017.01044
- [21] Medet Mukushev, Arman Sabyrov, Alfarabi Imashev, Kenessary Koishybay, Vadim Kimmelman, and Anara Sandygulova. 2020. Evaluation of Manual and Non-manual Components for Sign Language Recognition. In *12th International Conference on Language Resources and Evaluation (LREC 2020)*. European Language Resources Association (ELRA), Marseille, France, 6073–6078.
- [22] Nina-Kristin Pendzich. 2020. *Lexical Nonmanuals in German Sign Language: Empirical Studies and Theoretical Implications*. De Gruyter Mouton, Berlin, Boston. doi:10.1515/9783110671667
- [23] Marina Perea-Trigo, Enrique J. López-Ortiz, Luis M. Soria-Morillo, Juan A. Álvarez García, and J. J. Vegas-Olmos. 2024. Impact of face swapping and data augmentation on sign language recognition. *Universal Access in the Information Society* (July 2024). doi:10.1007/s10209-024-01133-y
- [24] Dinh Nam Pham and Eleftherios Avramidis. 2025. Transfer Learning from Visual Speech Recognition to Mouthing Recognition in German Sign Language. In *2025 19th IEEE International Conference on Automatic Face and Gesture Recognition (FG)*. 1–6.
- [25] Dinh Nam Pham, Vera Czehmann, and Eleftherios Avramidis. 2023. Disambiguating Signs: Deep Learning-based Gloss-level Classification for German Sign Language by Utilizing Mouth Actions. In *31st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN 2023, Bruges, Belgium, October 4-6, 2023*. doi:10.14428/ESANN/2023.ES2023-168
- [26] Noha Sarhan and Simone Frintrop. 2020. Transfer Learning For Videos: From Action Recognition To Sign Language Recognition. In *2020 IEEE International Conference on Image Processing (ICIP)*. 1811–1815. doi:10.1109/ICIP40778.2020.9191289
- [27] Noha Sarhan and Simone Frintrop. 2023. Unraveling a Decade: A Comprehensive Survey on Isolated Sign Language Recognition. In *2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. 3202–3211. doi:10.1109/ICCVW60793.2023.00345
- [28] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.).
- [29] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda B. Viégas, and Martin Wattenberg. 2017. SmoothGrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825* (2017).
- [30] Du Tran, Heng Wang, Matt Feiszli, and Lorenzo Torresani. 2019. Video Classification With Channel-Separated Convolutional Networks. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. 5551–5560. doi:10.1109/ICCV.2019.00565
- [31] Ulrich von Agris, Moritz Knorr, and Karl-Friedrich Kraiss. 2008. The significance of facial features for automatic sign language recognition. In *2008 8th IEEE International Conference on Automatic Face & Gesture Recognition*. 1–6. doi:10.1109/AFGR.2008.4813472