

Content4All Open Research Sign Language Translation Datasets

Necati Cihan Camgöz¹ Ben Saunders¹ Guillaume Rochette¹

Marco Giovanelli² Giacomo Inches² Robin Nachtrab-Ribback³ Richard Bowden¹

¹CVSSP, University of Surrey, Guildford, UK ²Fincons Group, Switzerland ³SWISSTXT, Switzerland

Abstract—Computational sign language research lacks the large-scale datasets that enables the creation of useful real-life applications. To date, most research has been limited to prototype systems on small domains of discourse, e.g. weather forecasts. To address this issue and to push the field forward, we release six datasets comprised of 190 hours of footage on the larger domain of news. From this, 20 hours of footage have been annotated by Deaf experts and interpreters and is made publicly available for research purposes. In this paper, we share the dataset collection process and tools developed to enable the alignment of sign language video and subtitles, as well as baseline translation results to underpin future research.

I. INTRODUCTION

Sign languages are natural visual languages [36]. They are the main medium of communication of the Deaf and have rich grammatical rules with unique syntax and vocabulary. Signers utilize multiple complementary articulators to convey information via spatio-temporal constructs [3]. Machine understanding of sign languages is a challenging task and an active research field [4], that has seen recent progress with the availability of benchmark recognition [18], [31], [30] and translation [6], [1] datasets.

As with any machine translation task, sign language translation [9] and production [34] require large-scale corpora to build models that generalize effectively. However, large-scale datasets with high quality annotations, on par with their spoken language counterparts, are not available for computational sign language research. This not only limits the advancement of the field, but also gives the research and Deaf communities a false sense of technological readiness by reporting promising results on small datasets with limited domains of discourse.

To address these issues, in this work we introduce 190 hours of public datasets, an order of magnitude larger than currently available resources, aimed at supporting sign language translation research [6]. In collaboration with broadcast partners, we collected sign language interpretation footage on the large domain of news. We adapted an annotation tool [15] which was used to manually align spoken language subtitles to sign language videos¹. We extracted human body pose features from all the videos and share them as part of the dataset to reduce redundant compute. Additionally, we trained state-of-the-art transformer based translation models [9] using the pose features and report baseline results. The datasets are publicly available² for research purposes

¹<https://gitlab.com/content4all-public/data-annotation>

²<https://www.cvssp.org/data/c4a-news-corpus>



Fig. 1. Anonymized video frame samples from SWISSTXT (top) and VRT (bottom) broadcast footage

upon agreeing to a user license (CC 4.0 BY-NC-SA³).

The rest of this paper is structured as follows: In Section II we give a brief (non-exhaustive) overview of existing datasets that have been used for computational sign language research. We introduce the Content4All base datasets in Section III. We then give details of the subset selection and annotation process, and introduce the manually aligned datasets in Section IV. We report baseline translation performance in Section V to underpin future research. Finally, we conclude this paper and discuss possible future directions in Section VI.

II. RELATED WORK

Compared to their spoken counterparts, computational sign language research is an under-resourced field, especially with regards to data availability [4]. Until recently, most computational sign language research was focused on isolated sign recognition [7], [12], [22], working on datasets collected in controlled environments with a limited vocabulary [16], [8], [11]. This resulted in most research being incomparable whilst also robbing the field of competitive progress [24].

³<https://creativecommons.org/licenses/by-nc-sa/4.0/>

TABLE I DETAILED STATISTICS OF THE RELEASED THREE BASE (RAW) DATASETS.

	SWISSTXT-RAW-NEWS	SWISSTXT-RAW-WEATHER	VRT-RAW
Sign Interpretations	DSGS	DSGS	VGT
Subtitle Language	Swiss-German	Swiss-German	Flemish
Number of Videos	183	181	120
Duration (hours:minutes)	76:05	12:09	100:21
Video Resolution@FPS	1280x720@50fps	1280x720@50fps	1280x720@25fps
Video Codec/Container	H264/mp4	H264/mp4	H264/mp4

The available large-scale linguistic resources, such as BSLCP [35] and mein-DGS [20], have been considered as alternate resources for sign language recognition [14]. However, they were not collected with vision-based sign language research in mind. They lack the necessary inter- and intra- signer variance on shared content, which current approaches rely on to build generalizable statistical models.

Relatively recently, the field has shifted focus to continuous recognition of co-articulated sign sequences [26]. To support this move, researchers have collected datasets based on sentence repetition, where signers were asked to repeat preconstructed sign phrases [31], [1], [30]. Although these datasets are useful for their target application, such as sign language assessment [37], [2], [42], they have a limited vocabulary and phrase variance, thus making them less useful for building large-scale sign language translation models.

Sign interpretations from broadcast footage have also been considered as alternate sources of data [13], [29]. Given the variety and quantity of the signed content, they are an ideal resource for training sign language translation models that generalize to large domains of discourse. A good example of broadcast footage utilization for computational sign language research is the Phoenix corpora [17], [18]. Comprised of sign language interpretations from broadcast weather forecast footage, the Phoenix datasets have been widely used for continuous sign language recognition research [43], [25], [5]. With the release of the aligned spoken language translations [6], Phoenix2014T has become the benchmark dataset for sign language translation. Although researchers have developed successful models which achieve promising results on Phoenix2014T [9], [10], [41], the application of such models has been limited to the weather domain, due to the limited domain of discourse of the dataset. Furthermore, compared to contemporary broadcast technology (1280x720@50fps), the dataset is of poor visual quality. It has low image resolution (210x260) and suffer from motion blur and interlacing artifacts, due to being recorded between the period of 2009-2013.

To address the shortcomings of the available public datasets, in this work we release six new sign language translation datasets of broadcast sign language interpretations. We share an order of magnitude larger amount of data compared to the benchmark Phoenix2014T dataset [6] (190 hours vs 10 hours, respectively) in two different sign languages, namely Swiss German Sign Language - Deutschschweizer Gebärdensprache (DSGS) and Flemish Sign Language - Vlaamse Gebarentaal (VGT). We provide both sign language

interpretations in the large domain of news broadcasts and the corresponding spoken language subtitles for each program, as well as manual spoken language text-sign video alignments for three of the datasets.

III. CONTENT4ALL BASE (RAW) DATASETS

In this section we introduce the three base datasets that have been created using the broadcast sign language interpretation footage. Broadcast partners SWISSTXT and VRT recorded and provided news footage from the periods March-August 2020 and March-July 2020, respectively. All of the video footage contains picture-in-picture sign language interpretations in DSGS for SWISSTXT and VGT for VRT.

In total, SWISSTXT and VRT provided 183 and 120 days of news footage, yielding 152 hours 40 minutes and 100 hours 21 minutes of video data, respectively. Both sources consist of mp4 files with a resolution of 1280x720. Both broadcasters provided time-coded subtitles in their respective spoken language, which we also make publicly available.

VRT videos contain uninterrupted sign language interpretations (VRT-RAW). However, SWISSTXT broadcast footage is interlaced with advertisements, which are not signed. We have manually annotated and removed the advertisement regions, creating two new datasets, namely SWISSTXT-RAW-NEWS and SWISSTXT-RAW-WEATHER. The detailed statistics of all the base (RAW) datasets can be seen in Table I.

To encourage the use of the datasets and to reduce redundant compute, we share extracted 2D and 3D human pose estimation features for all the video frames. We first ran the BODY-135 variant of the OpenPose library [21] and obtained 2D human pose information. The obtained 2D body skeletons are then uplifted using a deep learning based 2D-to-3D uplifting method proposed in [32]. Both 2D and 3D pose information are stored in json files, with one file per frame. We compress all json files corresponding to the same video into a single tar.xz file, for the sake of storage ease and efficiency. We also share a sample Python script to access them efficiently.

All of the footage has been anonymized to protect the privacy of the individuals present in the released datasets. We utilize the facial landmark predictions of OpenPose to locate each individual's face and blur the corresponding regions using a Gaussian kernel. Anonymized video frame samples can be seen in Figure 1. Interested parties are advised to contact the broadcast partners for the non-anonymized versions of the video files.

TABLE II DETAILED STATISTICS OF THE THREE MANUALLY ALIGNED DATASETS IN CONTRAST TO PHOENIX2014T.

	Phoenix2014T	SWISSTXT-WEATHER	SWISSTXT-NEWS	VRT-NEWS
#sequences	8257	811	6031	7174
#frames	947,756	156,891	1,693,391	810,738
Duration (hour:mins:secs)	10:31:50	00:52:18	09:24:28	09:00:30
Total Words (TW)	99,081	6,675	72,892	79,833
Vocabulary Size (VS)	2,287	1,248	10,561	6,875
NS (/VS)	1,077 (0.37)	701 (0.56)	5,969 (0.57)	3,405 (0.50)
< 3 (/VS)	1,427 (0.49)	874 (0.70)	7,561 (0.72)	4423 (0.64)
< 5 (/VS)	1,758 (0.60)	999 (0.80)	8,779 (0.83)	5334 (0.78)
<10 (/VS)	2,083 (0.72)	1,124 (0.90)	9,661 (0.91)	6009 (0.87)

In summary, the curated base (RAW) datasets include:

- Anonymized news footage with picture-in-picture sign language interpretations.
- Corresponding spoken language subtitles.
- Human pose estimation features, namely 2D OpenPose [21] and 3D skeleton [32] outputs.

IV. CONTENT4ALL MANUALLY ALIGNED DATASETS

Automatic sign language translation has so far only been successfully realized for a small domain of discourse, namely weather forecasts using the Phoenix2014T dataset [6]. However, the same techniques fail to learn successful translation models on larger domains of discourse using similar amounts of data [14].

To move the field beyond the weather domain, whilst also considering the aforementioned limitations of learning-based translation techniques, we selected 10 hour subsets from both the SWISSTXT-RAW-NEWS and VRT-RAW datasets that best resembles the Phoenix2014T dataset statistics. We manually aligned spoken language text and sign language video pairs in these subsets with the help of Deaf experts and sign language interpreters. These two new datasets, SWISSTXT-NEWS and VRT-NEWS, contain over 13,000 sign language interpretation and spoken language translation pairs. Whilst building up the annotation process, we also annotated a subset of the SWISSTXT-RAW-WEATHER dataset, which we also make publicly available for research purposes (SWISSTXT-WEATHER).

The detailed statistics of all three manually annotated datasets in comparison to Phoenix2014T can be seen in Table II. We calculate the total number of words, vocabulary size and number of singletons (words that only occur once). We also share the number of unique words that occur less than three (<3), five (<5) and ten (<10) times. To enable a comparison between the different datasets, we also report metrics proportional to the vocabulary size (/VS).

Similar to the base (RAW) datasets (Section III), we also provide anonymized sign language videos and 2D & 3D human body pose estimates for these annotated datasets. In addition, we provide a json file for each dataset, containing the spoken language sentence-sign language interpretation alignments. Each element in the json files has a unique ID that corresponds to a spoken language sentence-sign phrase pair. Alongside the spoken language translations, the beginning

and end timing information in *minutes:seconds:milliseconds* and in frames are provided, which indicates the location of the sign video-sentence alignment in the corresponding video. Within each element, we also store the relative path of corresponding videos and the 2D/3D skeletal files.

In the rest of this section we give further details on the subset selection and manual alignment processes.

A. Subset Selection

We first parsed the subtitles and manually cleaned the text to eliminate special characters, i.e. replaced periods (“.”) which are not used to terminate sentences with a special token. We then tokenized subtitle sentences into words, and calculated word statistics. To determine the size of the domain of discourse, we used spoken language word statistics as a proxy.

To select 10 hour subsets for manual subtitle alignment, we implemented a beam search algorithm. In order to constrain the domain of discourse, we capitalised on the COVID focused news during the capture period and used the first 22 minutes of each video. Considering some parts of the videos are eliminated during the annotation process, we selected 30 videos for each dataset, thus yielding a total of 11 hours of footage for each dataset. We used a beam size of 50 and use vocabulary size as the key metric to select a combination of videos that have similar statistics to the Phoenix2014T corpus. The selected subset was then annotated using the annotation tool described in the following section.

B. Manual Alignment

Current sign language translation techniques require parallel corpora, which provide spoken language translations that correspond to segmented sign language videos. The curation of such a dataset is a challenging task, as spoken and sign languages do not share a one-to-one word mapping and have a non-monotonic relationship. We utilized subtitles obtained from broadcast footage and aligned these subtitles to sign language interpretations.

Even though there are sign language video annotation tools already available (e.g. ELAN [39]), they mainly focus on data entry for linguistic purposes, such as the annotation of glosses. In our specific use case, we need to manually adjust weakly pre-aligned sentences (based on subtitles) with the corresponding sign language interpreter video in order to speed up the annotation work.

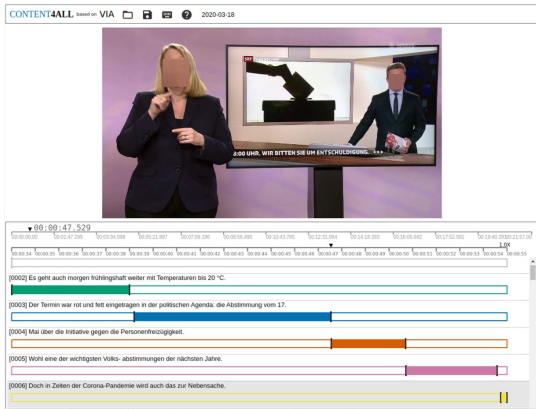


Fig. 2. Sign language interpretation - subtitle alignment tool user interface.

We used the VIA annotation tool [15] developed by the University of Oxford, and adapted it to facilitate the alignment of spoken language subtitles to sign language videos. We converted the subtitle files into VIA project file format. The project files were then sent to the Deaf experts and interpreters via email with a link to the annotation platform. The annotators were asked to align spoken language text to sign language videos when applicable. An equivalent of three person-months was needed to annotate 11 hours of video in each sign language. The user interface of the modified tool used by the annotators can be seen in Figure 2.

V. BASELINE EXPERIMENTS

To underpin future research on the released datasets, we train sign language translation models and report baseline results on SWISSTXT-NEWS and VRT-NEWS.

Protocols: To make future research on our datasets comparable, we set training and testing protocols. We divide each dataset into three partitions, namely training, development and test set. Detailed partition statistics for SWISSTXT-NEWS and VRT-NEWS can be seen in Table III.

TABLE III STATISTICS OF THE EVALUATION PROTOCOLS.

	SWISSTXT-NEWS			VRT-NEWS		
	Train	Dev	Test	Train	Dev	Test
#segments	5,126	442	463	6,098	542	534
#frames	1,445,804	122,666	130,952	698,077	60,887	58,948
Total Word Count	62,302	5,382	5,774	67,973	6,060	5,858
Vocabulary Size	9,623	1,976	2,093	6,325	1,612	1,584
#Out-of-Vocab	-	455	500	-	304	257
#Singletons	5,535	-	-	3,190	-	-

To measure the translation performance, we use Recall-Oriented Understudy for Gisting Evaluation (ROUGE) [27] and Bilingual Evaluation Understudy (BLEU) [28] scores, which are the most common metrics used for machine and sign language translation [9], [10], [33].

Experimental Setup and Results: We employ a state-of-the-art transformer [38] based sign language translation approach [10]. We represent video frames with the extracted 2D human pose features and use them as inputs to our models. We construct our transformers using three layers, with 512 hidden units, 1024 feed forward size and 8 heads

per layer. We use Xavier initialization [19] and train all our networks from scratch. We also utilize dropout with a drop rate of 0.3 to help our models generalize.

We use Adam [23] optimizer with a batch size of 32 and a learning rate of 10^{-3} . We employ a plateau learning rate scheduling algorithm which tracks the development set performance. We evaluate our network every 100 iterations. If the development set performance does not improve for 8 evaluation cycles, we decrease the learning rate by a factor of 0.7. This continues until the learning rate drops below 10^{-5} .

During inference, we utilize a beam search decoding with a length penalty [40]. We determine the best performing beam width and length penalty α values based on development set performance. We then use these values for reporting results on the test set.

TABLE IV BASELINE RESULTS ON OUR DATASETS.

	Dev		Test	
	ROUGE	BLEU-4	ROUGE	BLEU-4
SWISSTXT-NEWS	15.72	0.46	15.34	0.41
VRT-NEWS	17.63	0.45	17.77	0.36

We achieve similar baseline translation performances on both datasets (See Table IV). These results not only display the challenge of moving to larger domains of discourse, but also encourage the development of more specialized sign language translation approaches to realize real-life computational sign language applications.

VI. CONCLUSIONS

In this paper we introduced six new open research datasets aimed at automatic sign language translation research. Sign language interpretation footage was captured by the broadcast partners SWISSTXT and VRT. Raw footage was anonymized and processed to extract 2D and 3D human body pose information. From the roughly 190 hours of processed data, three base (RAW) datasets were released, namely 1) SWISSTXT-RAW-NEWS, 2) SWISSTXT-RAW-WEATHER and 3) VRT-RAW. Each dataset contains sign language interpretations, corresponding spoken language subtitles, and extracted 2D/3D human pose information.

A subset from each base dataset was selected and manually annotated to align spoken language subtitles and sign language interpretations. The subset selection was done to resemble the benchmark Phoenix 2014T dataset. Our aim is for these three new annotated public datasets, namely 4) SWISSTXT-NEWS, 5) SWISSTXT-WEATHER and 6) VRT-NEWS to become benchmarks and underpin future research as the field moves closer to translation and production on larger domains of discourse.

ACKNOWLEDGEMENTS

We would like to thank Gregg Young (VRT), Harry Witzthum and Swiss Federation of the Deaf for their contributions to the dataset creation process. This work received funding from the European Union’s Horizon2020 research and innovation programme under grant agreement no. 762021 ‘Content4All’. This work reflects only the authors view and the Commission is not responsible for any use that may be made of the information it contains.

REFERENCES

- [1] N. Adaloglou, T. Chatzis, I. Papastratis, A. Stergioulas, et al. A Comprehensive Study on Deep Learning-based Methods for Sign Language Recognition. *IEEE MultiMedia*, 2020.
- [2] O. Aran, I. Ari, L. Akarun, B. Sankur, A. Benoit, A. Caplier, P. Campr, A. H. Carrillo, and F.-X. Fanard. SignTutor: An Interactive System for Sign Language Tutoring. *IEEE MultiMedia*, 16(1), 2009.
- [3] P. Boyes-Braem and R. Sutton-Spence. *The Hands are the Head of the Mouth: The Mouth as Articulator in Sign Languages*. Gallaudet University Press, 2001.
- [4] D. Bragg, O. Koller, M. Bellard, L. Berke, P. Boudreault, et al. Sign Language Recognition, Generation, and Translation: An Interdisciplinary Perspective. In *Proceedings of the International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)*, 2019.
- [5] N. C. Camgoz, S. Hadfield, O. Koller, and R. Bowden. SubUNets: End-to-end Hand Shape and Continuous Sign Language Recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [6] N. C. Camgoz, S. Hadfield, O. Koller, H. Ney, and R. Bowden. Neural Sign Language Translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [7] N. C. Camgoz, A. A. Kindiroglu, and L. Akarun. Sign Language Recognition for Assisting the Deaf in Hospitals. In *Proceedings of the International Workshop on Human Behavior Understanding (HBU)*, 2016.
- [8] N. C. Camgoz, A. A. Kindiroglu, S. Karabuklu, M. Kelepir, A. S. Ozsoy, and L. Akarun. BosphorusSign: A Turkish Sign Language Recognition Corpus in Health and Finance Domains. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*, 2016.
- [9] N. C. Camgoz, O. Koller, S. Hadfield, and R. Bowden. Multi-channel Transformers for Multi-articulatory Sign Language Translation. In *16th European Conference on Computer Vision Workshops (ECCVW)*, 2020.
- [10] N. C. Camgoz, O. Koller, S. Hadfield, and R. Bowden. Sign Language Transformers: Joint End-to-end Sign Language Recognition and Translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [11] X. Chai, H. Wang, and X. Chen. The DEVISIGN Large Vocabulary of Chinese Sign Language Database and Baseline Evaluations. *Technical Report VIPL-TR-14-SLR-001*, 2014.
- [12] H. Cooper and R. Bowden. Sign Language Recognition using Boosted Volumetric Features. In *Proceedings of the IAPR International Conference on Machine Vision Applications (MVA)*, 2007.
- [13] H. Cooper and R. Bowden. Learning Signs from Subtitles: A Weakly Supervised Approach to Sign Language Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [14] K. Cormier, N. Fox, B. Woll, A. Zisserman, N. Camgoz, and R. Bowden. ExTOL: Automatic recognition of British Sign Language using the BSL Corpus. In *Proceedings of the Sign Language Translation and Avatar Technology (SLTAT)*, 2019.
- [15] A. Dutta and A. Zisserman. The VIA Annotation Software for Images, Audio and Video. In *Proceedings of the 27th ACM International Conference on Multimedia*, 2019.
- [16] S. Ebling, N. C. Camgoz, P. B. Braem, K. Tissi, S. Sidler-Miserez, S. Stoll, S. Hadfield, T. Haug, R. Bowden, S. Tornay, M. Razavi, and M. Magimai-Doss. SMILE Swiss German Sign Language Dataset. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*, 2018.
- [17] J. Forster, C. Schmidt, T. Hoyoux, O. Koller, U. Zelle, J. H. Piater, and H. Ney. RWTH-PHOENIX-Weather: A Large Vocabulary Sign Language Recognition and Translation Corpus. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*, 2012.
- [18] J. Forster, C. Schmidt, O. Koller, M. Bellgardt, and H. Ney. Extensions of the Sign Language Recognition and Translation Corpus RWTH-PHOENIX-Weather. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*, 2014.
- [19] X. Glorot and Y. Bengio. Understanding the Difficulty of Training Deep Feedforward Neural Networks. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2010.
- [20] T. Hanke, L. König, S. Wagner, and S. Matthes. DGS Corpus & Dicta-Sign: The Hamburg Studio Setup. In *Proceedings of the Representation and Processing of Sign Languages: Corpora and Sign Language Technologies*, 2010.
- [21] G. Hidalgo, Y. Raaj, H. Idrees, D. Xiang, H. Joo, T. Simon, and Y. Sheikh. Single-Network Whole-Body Pose Estimation. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [22] H. R. V. Joze and O. Koller. MS-ASL: A Large-Scale Data Set and Benchmark for Understanding American Sign Language. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2019.
- [23] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.
- [24] O. Koller. Quantitative Survey of the State-of-the-art in Sign Language Recognition. *arXiv:2008.09918*, 2020.
- [25] O. Koller, N. C. Camgoz, R. Bowden, and H. Ney. Weakly Supervised Learning with Multi-Stream CNN-LSTM-HMMs to Discover Sequential Parallelism in Sign Language Videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2019.
- [26] O. Koller, J. Forster, and H. Ney. Continuous Sign Language Recognition: Towards Large Vocabulary Statistical Recognition Systems Handling Multiple Signers. *Computer Vision and Image Understanding (CVIU)*, 141, 2015.
- [27] C.-Y. Lin. ROUGE: A Package for Automatic Evaluation of Summaries. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics, Text Summarization Branches Out Workshop*, 2004.
- [28] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. BLEU: A Method for Automatic Evaluation of Machine Translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2002.
- [29] T. Pfister, J. Charles, M. Everingham, and A. Zisserman. Automatic and Efficient Long Term Arm and Hand Tracking for Continuous Sign Language TV Broadcasts. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2012.
- [30] G. Potamianos, K. Papadimitriou, E. Efthimiou, S.-E. Fotinea, G. Sarpountzaki, and P. Maragos. SL-ReDu: Greek sign language recognition for educational applications. Project description and early results. In *Proceedings of the 13th ACM International Conference on Pervasive Technologies Related to Assistive Environments*, 2020.
- [31] J. Pu, W. Zhou, and H. Li. Iterative Alignment Network for Continuous Sign Language Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [32] G. Rochette, C. Russell, and R. Bowden. Weakly-Supervised 3D Pose Estimation from a Single Image using Multi-View Consistency. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2019.
- [33] B. Saunders, N. C. Camgoz, and R. Bowden. Adversarial Training for Multi-Channel Sign Language Production. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020.
- [34] B. Saunders, N. C. Camgoz, and R. Bowden. Progressive Transformers for End-to-End Sign Language Production. In *16th European Conference on Computer Vision (ECCV)*, 2020.
- [35] A. Schembri, J. Fenlon, R. Rentelis, S. Reynolds, and K. Cormier. Building the British Sign Language Corpus. *Language Documentation & Conservation (LD&C)*, 7, 2013.
- [36] R. Sutton-Spence and B. Woll. *The Linguistics of British Sign Language: An Introduction*. Cambridge University Press, 1999.
- [37] S. Tornay, M. Razavi, N. C. Camgoz, R. Bowden, and M. Magimai-Doss. HMM-based Approaches to Model Multichannel Information in Sign Language Inspired from Articulatory Features-based Speech Processing. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2019.
- [38] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is All You Need. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [39] P. Wittenburg, H. Brugman, A. Russel, A. Klassmann, and H. Sloetjes. ELAN: A Professional Framework for Multimodality Research. In *Proceedings of the 5th International Conference on Language Resources and Evaluation (LREC)*, 2006.
- [40] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, et al. Google's Neural Machine Translation System: Bridging the Gap Between Human and Machine Translation. *arXiv:1609.08144*, 2016.
- [41] K. Yin and J. Read. Better Sign Language Translation with STMC-Transformer. In *Proceedings of the 28th International Conference on Computational Linguistics*, 2020.
- [42] Z. Zafrulla, H. Brashear, T. Starner, H. Hamilton, and P. Presti. American Sign Language Recognition with the Kinect. In *Proceedings of the ACM International Conference on Multimodal Interfaces*, 2011.
- [43] H. Zhou, W. Zhou, Y. Zhou, and H. Li. Spatial-Temporal Multi-Cue Network for Continuous Sign Language Recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.