

# Sign Language Translation with Sentence Embedding Supervision

Yasser Hamidullah and Josef van Genabith and Cristina España-Bonet

{yasser.hamidullah, Josef.van\_Genabith, cristinae}@dfki.de

German Research Center for Artificial Intelligence (DFKI GmbH)

Saarland Informatics Campus, Saarbrücken, Germany

## Abstract

State-of-the-art sign language translation (SLT) systems facilitate the learning process through gloss annotations, either in an end2end manner or by involving an intermediate step. Unfortunately, gloss labelled sign language data is usually not available at scale and, when available, gloss annotations widely differ from dataset to dataset. We present a novel approach using sentence embeddings of the target sentences at training time that take the role of glosses. The new kind of supervision does not need any manual annotation but it is learned on raw textual data. As our approach easily facilitates multilinguality, we evaluate it on datasets covering German (PHOENIX-2014T) and American (How2Sign) sign languages and experiment with mono- and multilingual sentence embeddings and translation systems. Our approach significantly outperforms other gloss-free approaches, setting the new state-of-the-art for data sets where glosses are not available and when no additional SLT datasets are used for pretraining, diminishing the gap between gloss-free and gloss-dependent systems.

## 1 Introduction

Sign Language Translation (SLT) aims at generating text from sign language videos. There are several approaches to SLT reported in the literature, with *sign2text* and *sign2gloss2text* the most widely used. While *sign2text* directly translates video into text with or without the help of glosses (Camgöz et al., 2018), *sign2gloss2text* passes through an intermediate gloss step before translation into spoken language text (Ormel et al., 2010). That is, *sign2gloss2text* breaks down the problem into two independent sub-problems using *glosses* as a pivot language. A gloss is a textual label associated with a sign, and, although human signers do not in general use them, performance in automatic SLT has long been upper bounded by the gloss supervision and their use as an intermediate representation

(Camgöz et al., 2018). The advantage of translation without glosses is that collecting data is much easier. Even though translation results are better for approaches that use glosses as intermediate representation (Chen et al., 2022a,b), this comes at the cost of annotating all the video data with glosses which is a time consuming manual task. For many data sets glosses are simply not available. On the plus side, with gloss supervision-based SLT architectures, one can take full advantage of the maturity of text2text machine translation between glosses and spoken language text.

In this work, we present a novel approach *sign2(sem+text)*, a model that gets rid of glosses and adds supervision through sentence embeddings, SEM, pretrained on raw text and finetuned for sign language. Our experimental results demonstrate the strength of the novel approach on both standard small datasets with gloss annotation and larger datasets without. In the latter case, we achieve state-of-the-art results for the American Sign Language (ASL) dataset How2Sign when no additional SLT datasets are used<sup>1</sup> improving over Tarrés et al. (2023) by 4 BLEU points. For German Sign Language (DGS), our new approach achieves translation quality scores between the previous best gloss-free system (Zhou et al., 2023) and the current state-of-the-art using glosses (Chen et al., 2022b) on the PHOENIX-2014T dataset. Our code and models are publicly available.<sup>2</sup>

## 2 Related Work

Camgöz et al. (2018) proposed three formalisations considering SLT as a seq2seq problem that converts a sequence of signs into a sequence of words: (i) *sign2text*, a model that encodes video frames using pretrained 2D CNNs as spatial features and then

<sup>1</sup>Uthus et al. (2023) and Rust et al. (2024) obtain better results by using the YouTube-ASL dataset for pretraining.

<sup>2</sup><https://github.com/yhamidullah/sem-slt>

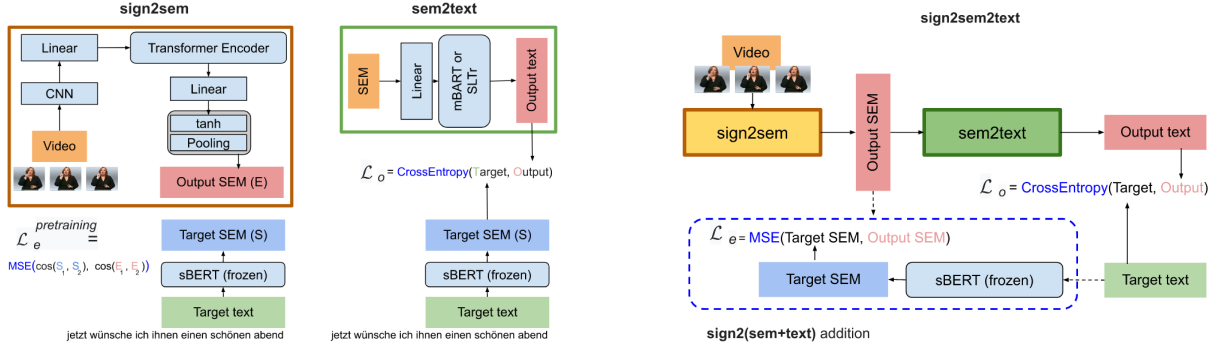


Figure 1: sign2sem and sem2text independent modules for the SLT task (left plot). End2end architectures: pipeline system sign2sem2text and multitask system sign2(sem+text) (right plot).

uses an RNN to generate the text; (ii) gloss2text, a model learning the translation between a sequence of textual glosses and fluent spoken language text; and (iii) a sign2gloss2text model that adds an extra intermediate gloss layer between the video and output text levels of a sign2text architecture to provide additional gloss supervision using a CTC loss.

In follow-up work, Camgöz et al. (2020) proposed an architecture for joint learning continuous sign language recognition (CSLR) and SLT which uses the same input as Camgöz et al. (2018); Zhou et al. (2021), that is, pretrained visual features, but a transformer (Vaswani et al., 2017) for text generation. Camgöz et al. (2020) conjectured that gloss2text results with ground-truth glosses provide an upper bound for SLT. Supporting this assumption, their translation quality on PHOENIX-2014T as measured by BLEU (Papineni et al., 2002) achieved 24.5 on gloss2text and 21.8 on sign2text.

Yin and Read (2020) used a different visual representation with a multi-cue network (Zhou et al., 2020) to encode videos. Cues included face, hands and pose besides the full frame. With a BLEU score of 24.0 they improved over sign2gloss2text Camgöz et al. (2020) and concluded that their visual representation was better than the spatial frame embeddings used by the Camgöz et al. (2020).

Chen et al. (2022a,b) used both pretraining of a network based on S3Ds (Xie et al., 2018) on action recognition for CSLR (sign2gloss) and pretraining of a textual transformer (gloss2text) with mBART-25 (Liu et al., 2020). Both types of pretraining are progressively adapted to the domain of the task by adding data closer to the domain. An additional mapping network between the vision and language parts allows Chen et al. (2022a) to build an end2end sign2text model relying on inter-

nal gloss supervision. To the best of our knowledge, Chen et al. (2022b) is the current state of the art for both sign2text (BLEU=28.95) and sign2gloss2text (BLEU=26.71), all on the PHOENIX-2014T data set.

Over the last few years, several gloss-free models have emerged (Li et al., 2020; Zhao et al., 2022; Yin et al., 2023). Zhou et al. (2023) obtains the current state-of-the-art in this category by utilising visual-language pretraining following CLIP (Radford et al., 2021). On the datasets (Camgöz et al. (2018); Zhou et al. (2021)) where the two approaches can be compared, translation quality diminishes by up to 7 BLEU points when the glosses are not used (Yin et al., 2023; Zhou et al., 2023). Tarrés et al. (2023) uses the How2Sign dataset (Duarte et al. (2021)) (where no gloss information is available) with I3D (Carreira and Zisserman (2017)) features for video representations and a Transformer. Uthus et al. (2023) introduces a new dataset, YouTube-ASL, 10 times larger than the previous one (Duarte et al. (2021)), and uses 2D pose estimation and pretraining to improve on Tarrés et al. (2023) best results on How2Sign (BLEU=8.09 vs BLEU=12.4). Simultaneously with our work, Rust et al. (2024) pretrains a self-supervised and privacy-aware visual model on YouTube-ASL to achieve the new state-of-the-art performance on How2Sign (BLEU=15.5).

### 3 SEM-based Architectures

In our work we build two systems that revolve around textual sentence embeddings, SEM, as depicted in Figure 1. The figure presents two independent modules sign2sem and sem2text (left plots) that we later combine in sign2sem2text and sign2(sem+text) in an end2end setting (right plot).

- **sign2sem Module** This module predicts an intermediate SEM vector. Given a set of frames (video) features, sign2sem produces a vector representing the sentence signed in the video using a transformer encoder.

**Pretraining** the visual feature sentence embedding model on text. We follow Reimers and Gurevych (2019) and train a Siamese network with twin subnetworks 1 and 2. We compute the loss as the minimum squared error (MSE):

$$\mathcal{L}_e = \frac{1}{N} \sum_{i=1}^N (\cos(S_{1,i}, S_{2,i}) - \cos(E_{1,i}, E_{2,i}))^2$$

where  $N$  is the batch size, and  $S$  and  $E$  contain the target text SEM vectors and the predicted output SEM vectors respectively. In our experiments, the target SEM vector is given by sBERT (Reimers and Gurevych, 2019) here and in our models below.

- **sem2text Module** This module is responsible for the text reconstruction from sentence embeddings SEM. It produces the text translation of the video features encoded in a given SEM vector. The core sem2text model is a transformer model; we compare encoder–decoder and only decoder systems for the task:

- **Encoder–decoder (SLTr)**: this version uses the sign language transformer (SLTr) architecture as in Camgöz et al. (2020). We use a transformer base with a linear projection from the SEM vector input instead of the usual word embedding layer.

- **Decoder only with pretrained mBART**: this version uses a pre-trained mBART-25 decoder and a linear layer to project the SEM vectors into the mBART model dimensions.

**Pretraining** We train both transformers (SLTr from scratch and the already pretrained mBART-25) with Wikipedia data and then finetune them on the SL datasets. We compute the translation output loss as the cross-entropy:

$$\mathcal{L}_o = \text{CE}(T, O) = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M (T_{ij} \cdot \log(O_{ij}))$$

where  $N$  is the batch size,  $M$  the vocabulary size,  $T$  is the target text and  $O$  is the output text.

After pretraining each component (sign2sem and sem2text), we combine them together for end2end training. We explore two approaches: an approach that only uses the output loss  $\mathcal{L}_o$ , **sign2sem2text**, and an approach that integrates an additional supervision loss  $\mathcal{L}_e$ , **sign2(sem+text)**.

- **sign2sem2text** is a simple pipeline combination of sign2sem and sem2text where the output SEM of the first module is used as input by the second module to obtain the final text prediction. The two pretrained modules (with both variants **SLTr** and **mBART**) are put together and trained in an end2end manner without any intermediate supervision. This formalisation is the sentence embedding equivalent to the sign2gloss2text approach.

- **sign2(sem+text)** performs translation using the same components as sign2sem2text. However, it uses the sign2sem SEM output as additional intermediate supervision using MSE loss computed against the target text SEM in a multitask learning approach. Both,  $\mathcal{L}_e$  (sentence embedding) and  $\mathcal{L}_o$  (output text), are used jointly to train the model.

For **SLTr**, we take the SEM before the tanh and pooling (see Figure 1 (left–middle)), and project it into the SLTr model dimension. The supervision is applied after the SLTr encoder. For **mBART**, the supervision happens right before the mBART.

Our architectures can be trained both monolingually and multilingually simply by using multilingual embeddings and merging multilingual training data.

## 4 Experimental Settings

We use two diverse (language and domain) **datasets** for our experiments:

**RWTH-PHOENIX-2014T (Camgöz et al., 2018)** 11 hours of weather forecast videos from 9 signers. Signers use German Sign Language and both transcriptions and glosses are available.

**How2Sign (Duarte et al., 2021)** 80 hours of instructional videos with speech and transcriptions and their corresponding American Sign Language videos (glosses unavailable) from 11 signers.

Detailed statistics for each dataset are provided in Appendix A. We **preprocess** the textual part of the datasets in a way that allows us to compare to the results obtained by Camgöz et al. (2018). We tokenise and lowercase the input for both training and evaluation. We apply BPE (Sennrich et al., 2016) with a vocabulary size of 1500 for Phoenix-2014T and 5000 for How2Sign. When pretraining sem2text SLTr, we use a shared (en–de) vocabulary size of 32000. In cases where we use pretrained models, we keep the tokenisation of the model.

		PHOENIX-2014T (DGS)				How2Sign (ASL)			
		BLEU <sub>val</sub>	BLEU	chrF	BLEURT	BLEU <sub>val</sub>	BLEU	chrF	BLEURT
SLTr	sign2sem2text - mono	14.22	13.4±1.4	33.5±1.5	0.379±0.016	6.69	5.7±0.4	21.2±0.4	0.382±0.005
	sign2sem2text - multi	13.05	12.7±1.3	32.3±1.3	0.343±0.014	6.48	6.4±0.4	22.0±0.5	0.403±0.006
	sign2(sem+text) - mono	19.10	18.8±1.7	40.1±1.5	0.437±0.016	10.41	9.5±0.5	27.4±0.5	0.445±0.006
	sign2(sem+text) - multi	17.03	16.6±1.6	37.9±1.5	0.412±0.016	7.85	7.8±0.4	25.4±0.5	0.430±0.006
mBART	sign2sem2text - mono	16.67	17.3±1.6	38.2±1.5	0.434±0.016	9.32	9.8±0.5	<b>31.2±0.5</b>	0.477±0.006
	sign2sem2text - multi	16.91	16.5±1.6	37.3±1.5	0.425±0.016	9.11	9.6±0.5	<b>31.2±0.5</b>	0.475±0.006
	sign2(sem+text) - mono	24.07	<b>24.2±1.9</b>	<b>46.3±1.6</b>	<b>0.483±0.017</b>	12.20	<b>11.7±0.5</b>	<b>32.0±0.5</b>	<b>0.487±0.006</b>
	sign2(sem+text) - multi	24.12	<b>24.1±1.9</b>	<b>46.1±1.6</b>	<b>0.481±0.017</b>	12.34	<b>12.0±0.5</b>	<b>31.8±0.5</b>	<b>0.483±0.006</b>

Table 1: Translation performance of our models on validation (val) and test. Best models at 95% confidence level are highlighted. Previous state-of-the-art for gloss-free systems is BLEU=21.44 for PHOENIX (Zhou et al., 2023) and 8.03 for How2Sign (Tarrés et al., 2023). Chen et al. (2022b) achieves 28.95 on PHOENIX with their gloss-assisted system sign2text and 26.71 with sign2gloss2text. Rust et al. (2024) achieves 15.5 on How2Sign pretraining with YouTube-ASL.

For video files, we extract frames using ffmpeg. We normalise the images, and resize them to 224x224. In this step, we initially obtain frame features from a pretrained model (Tan and Le, 2019), which does not contain gloss information. We then apply pooling to remove the spatial dimensions, followed by batch normalisation with ReLU, following the approach outlined by Camgöz et al. (2020). This generic approach facilitates the combination of datasets in the multilingual setting.

We use two multilingual **pretrained models** that cover both German and English, sBERT (Reimers and Gurevych, 2019)<sup>3</sup> for sentence embeddings and mBART (Liu et al., 2020)<sup>4</sup> as a language model. For further pretraining we use 26 million sentences per language from the English and German Wikipedia dumps extracted with Wikitailor (España-Bonet et al., 2023).

Following Müller et al. (2022) and Müller et al. (2023), we **evaluate** the models using three common automatic metrics in machine translation: BLEU (Papineni et al., 2002), chrF (Popović, 2015) and BLEURT (Sellam et al., 2020). Specifics can be found in Appendix C. In all cases, we estimate 95% confidence intervals (CI) via bootstrap resampling (Koehn, 2004) with 1000 samples.

## 5 Results and Discussion

Table 1 presents the results for our models and variants. Two major trends are observed: (i) massive pretraining of the sem2text module (mBART vs SLTr) significantly improves the results, confirming the observations by Chen et al. (2022a)

and (ii) the multitask approach sign2(sem+text) is better than the pipeline approach sign2sem2text. These findings hold for all three evaluation metrics at 95% confidence level.

Potentially beneficial effects of multilinguality are less evident. Monolingual and multilingual approaches are not distinguishable within the 95% CIs, possibly due to large differences in the domain of the datasets preventing effective transfer between languages.

Our best system, sign2(sem+text) with the pretrained text decoder, achieves state-of-the-art results on How2Sign when no additional SLT dataset is used for pretraining, improving from 8 to 12 BLEU points over Tarrés et al. (2023). For PHOENIX-2014T, we surpass all previous gloss-free approaches (24 vs 21 BLEU), but we are still below the best approach that uses glosses (Chen et al., 2022b) (24 vs 29 BLEU).

**Reconstruction quality: sem2text.** In our approach, sentence embeddings take the role of manually produced glosses in previous work. Our sem2text translation module defines the upper-bound results for the full system as gloss2text did in previous work. Our best sign2(sem+text) models with mBART produce a reconstruction score of BLEU 38.0±2.4/23.3±1.0, chrF 57.5±1.9/43.9±0.9 and BLEURT 0.588±0.019/0.571±0.008 for PHOENIX-2014T/How2Sign (see Table 2). Where the comparison with glosses is available (PHOENIX), we improve over gloss2text by up to 10 BLEU points. We hypothesise that a sentence is better represented by its embedding than by a string of glosses and this explains why the translation

<sup>3</sup>We use all-MiniLM-L12-v2 model with 384 dimensions.

<sup>4</sup>We use mBART-25 1024 dimensions.

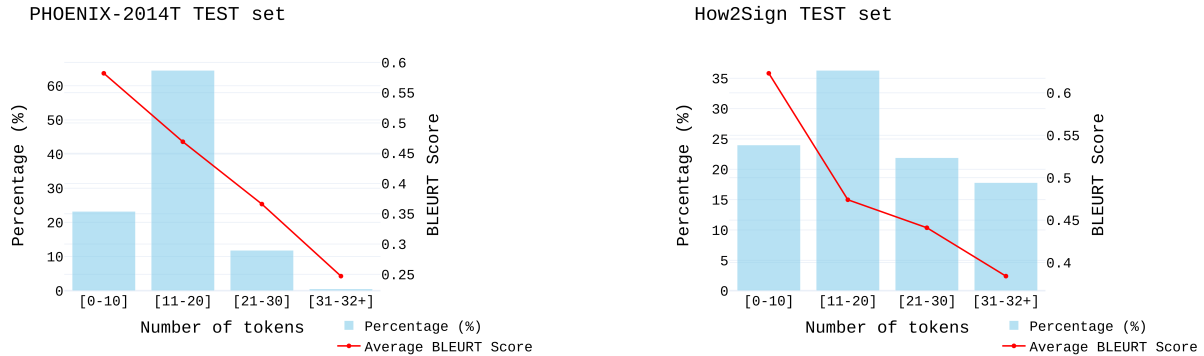


Figure 2: Average BLEURT score on different token length intervals on PHOENIX-2014T and How2Sign test.

	PHOENIX-2014T (DGS)			
	BLEU <sub>val</sub>	BLEU	chrF	BLEURT
Camgöz (2018)	20.16	19.26	–	–
Chen (2022a)	27.61	28.39	–	–
SLTr - mono	31.53	30.1±2.0	52.2±1.7	0.526±0.018
SLTr - multi	29.20	31.3±2.1	52.9±1.7	0.530±0.018
mBART - mono	37.11	<b>38.0±2.4</b>	<b>57.5±1.9</b>	<b>0.588±0.019</b>
mBART - multi	36.91	<b>37.5±2.3</b>	<b>57.4±1.8</b>	<b>0.584±0.018</b>

	How2Sign (ASL)			
	BLEU <sub>val</sub>	BLEU	chrF	BLEURT
SLTr - mono	13.24	14.6±0.6	34.3±0.6	0.489±0.006
SLTr - multi	16.17	16.4±0.7	36.5±0.7	0.529±0.007
mBART - mono	23.04	<b>22.8±1.0</b>	<b>43.3±0.9</b>	<b>0.577±0.008</b>
mBART - multi	24.60	<b>23.3±1.0</b>	<b>43.9±0.9</b>	<b>0.571±0.008</b>

Table 2: Reconstruction quality for the sem2text subtask of our models and gloss2text state-of-the-art on validation (val) and test. Best models at 95% confidence level are highlighted.

quality for sem2text is higher than for gloss2text. If these components (sem2text and gloss2text) are the upper-bound to the end2end sign2text translation, SEM-based systems are potentially at an advantage. These results, together with the fact that SEM models can be applied to raw data without annotations, highlight the promising future prospects of, especially, sign2(sem+text).

**SEM-based vs gloss-based SLT.** For comparison purposes, we integrate SEM supervision in a state-of-the-art gloss-based SLT system, Signjoey (Camgöz et al., 2020), by replacing their gloss supervision by SEM supervision. We perform no pretraining and train the two systems under the same conditions. We observe that convergence with SEM is faster and requires less than half of the iterations to finish (5k vs 12k) using the same setting and resources. The detailed training evolution

is shown in Appendix D.

**Translation quality vs output length.** Figure 2 shows the token length distribution of PHOENIX-2014T and How2Sign along with the average BLEURT score on each interval. The equivalent plots for chrF and BLEU are in Figures 4 and 5 in Appendix E respectively. In the PHOENIX test set, almost 90% of the sentences contain 20 tokens or less, while the number decreases to 60% for How2Sign. The 10-20 token range is the one with the best scores. While the drop in performance in translation quality for long sentences is smaller in How2Sign, the difference in the distribution affects the global quality.

## 6 Conclusions

We present a new approach to sign language translation using automatically computed sentence embeddings instead of manual gloss labels as intermediate representation with (sign2(sem+text)) and without (sign2sem2text) SEM supervision. We outperform the state-of-the-art of gloss-free SLT when no additional SLT datasets are used for pretraining, closing the gap to gloss-based SLT.

According to the upper-bound set by sem2text translation quality, there is still room for improvement for the end2end SEM-based SLT models. In this work, we limited ourselves to existing visual feature extractors, in the future we plan to train a SEM-based visual feature extractor on SL datasets in order to get closer to our sem2text upper-bound and match gloss-based performance.

## Limitations

Our SL datasets cover American English and German. Sentence embeddings for these languages are good quality as lots of textual data is available

for pre-training. It remains to be studied how the quality of the embeddings affects the final translation quality. This is important for low-resourced languages, i.e. languages with limited amounts of monolingual text data but, to the best of our knowledge, no public sign language data set exists for them.

## Acknowledgements

This work has been funded by by BMBF (German Federal Ministry of Education and Research) via the project SocialWear (grant no. 01IW20002).

## References

- Necati Cihan Camgöz, Simon Hadfield, Oscar Koller, Hermann Ney, and Richard Bowden. 2018. [Neural sign language translation](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7784–7793.
- 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 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10020–10030. Computer Vision Foundation / IEEE.
- J. Carreira and A. Zisserman. 2017. [Quo vadis, action recognition? a new model and the kinetics dataset](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4733, Los Alamitos, CA, USA. IEEE Computer Society.
- Yutong Chen, Fangyun Wei, Xiao Sun, Zhirong Wu, and Stephen Lin. 2022a. [A simple multi-modality transfer learning baseline for sign language translation](#). In *CVPR*, pages 5110–5120. IEEE.
- Yutong Chen, Ronglai Zuo, Fangyun Wei, Yu Wu, Shujie Liu, and Brian Mak. 2022b. [Two-stream network for sign language recognition and translation](#). *Advances in Neural Information Processing Systems*, 35:17043–17056.
- Amanda Duarte, Shruti Palaskar, Lucas Ventura, Deepti Ghadiyaram, Kenneth DeHaan, Florian Metze, Jordi Torres, and Xavier Giro-i Nieto. 2021. [How2Sign: A Large-Scale Multimodal Dataset for Continuous American Sign Language](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2735–2744.
- Cristina España-Bonet, Alberto Barrón-Cedeño, and Lluís Màrquez. 2023. [Tailoring and Evaluating the Wikipedia for in-Domain Comparable Corpora Extraction](#). *Knowledge and Information Systems*, 65(3):1365–1397.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André F. T. Martins. 2022. [Results of wmt22 metrics shared task: Stop using bleu – neural metrics are better and more robust](#). In *Proceedings of the Seventh Conference on Machine Translation*, pages 46–68, Abu Dhabi. Association for Computational Linguistics.
- Tom Kocmi, Christian Federmann, Roman Grundkiewicz, Marcin Junczys-Dowmunt, Hitokazu Matsushita, and Arul Menezes. 2021. [To ship or not to ship: An extensive evaluation of automatic metrics for machine translation](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 478–494, Online. Association for Computational Linguistics.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Dongxu Li, Chenchen Xu, Xin Yu, Kaihao Zhang, Ben Swift, Hanna Suominen, and Hongdong Li. 2020. [TSPNet: Hierarchical Feature Learning via Temporal Semantic Pyramid for Sign Language Translation](#). In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20*, Red Hook, NY, USA. Curran Associates Inc.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Mathias Müller, Malihe Alikhani, Eleftherios Avramidis, Richard Bowden, Annelies Braffort, Necati Cihan Camgöz, Sarah Ebling, Cristina España-Bonet, Anne Göhring, Roman Grundkiewicz, Mert Inan, Zifan Jiang, Oscar Koller, Amit Moryossef, Annette Rios, Dimitar Shterionov, Sandra Sidler-Miserez, Katja Tissi, and Davy Van Landuyt. 2023. [Findings of the second WMT shared task on sign language translation \(WMT-SLT23\)](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 68–94, Singapore. Association for Computational Linguistics.
- Mathias Müller, Sarah Ebling, Eleftherios Avramidis, Alessia Battisti, Michèle Berger, Richard Bowden, Annelies Braffort, Necati Cihan Camgöz, Cristina España-Bonet, Roman Grundkiewicz, Zifan Jiang, Oscar Koller, Amit Moryossef, Regula Perrollaz, Sabine Reinhard, Annette Rios, Dimitar Shterionov, Sandra Sidler-Miserez, Katja Tissi, and Davy Van Landuyt. 2022. [Findings of the First WMT Shared Task on Sign Language Translation \(WMT-SLT22\)](#). In *Proceedings of the Seventh Conference on Machine Translation*, pages 744–772, Abu Dhabi. Association for Computational Linguistics.
- Ellen Ormel, Onno Crasborn, Els van der Kooij, Lianne van Dijken, Ellen Yassine Nauta, Jens Forster, and

- Daniel Stein. 2010. [Glossing a multi-purpose sign language corpus](#). In *Proceedings of the LREC2010 4th Workshop on the Representation and Processing of Sign Languages: Corpora and Sign Language Technologies*, pages 186–191, Valletta, Malta. European Language Resources Association (ELRA).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Phillip Rust, Bowen Shi, Skyler Wang, Necati Cihan Camgöz, and Jean Maillard. 2024. [Towards privacy-aware sign language translation at scale](#).
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Mingxing Tan and Quoc Le. 2019. [EfficientNet: Rethinking model scaling for convolutional neural networks](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114. PMLR.
- Laia Tarrés, Gerard I. Gállego, Amanda Duarte, Jordi Torres, and Xavier Giró-i Nieto. 2023. Sign language translation from instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 5625–5635.
- Dave Uthus, Garrett Tanzer, and Manfred Georg. 2023. [YouTube-ASL: A Large-Scale, Open-Domain American Sign Language-English Parallel Corpus](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 29029–29047. Curran Associates, Inc.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. 2018. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV (15)*, volume 11219 of *Lecture Notes in Computer Science*, pages 318–335. Springer.
- A. Yin, T. Zhong, L. Tang, W. Jin, T. Jin, and Z. Zhao. 2023. [Gloss attention for gloss-free sign language translation](#). In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2551–2562, Los Alamitos, CA, USA. IEEE Computer Society.
- Kayo Yin and Jesse Read. 2020. [Better sign language translation with STMC-transformer](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5975–5989, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jian Zhao, Weizhen Qi, Wengang Zhou, Nan Duan, Ming Zhou, and Houqiang Li. 2022. [Conditional sentence generation and cross-modal reranking for sign language translation](#). *IEEE Transactions on Multimedia*, 24:2662–2672.
- Benjia Zhou, Zhigang Chen, Albert Clapés, Jun Wan, Yanyan Liang, Sergio Escalera, Zhen Lei, and Du Zhang. 2023. Gloss-free sign language translation: Improving from visual-language pretraining. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20871–20881.

Hao Zhou, Wen gang Zhou, Weizhen Qi, Junfu Pu, and Houqiang Li. 2021. [Improving sign language translation with monolingual data by sign back-translation](#). *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1316–1325.

Hao Zhou, Wengang Zhou, Yun Zhou, and Houqiang Li. 2020. Spatial-temporal multi-cue network for continuous sign language recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13009–13016.



## A Datasets Statistics

Table 3 summarises the statistics for the corpora used in the experiments.

	Phoenix-2014T	How2Sign
Src. Lang.	German	Am. English
Tgt. Lang.	DGS	ASL
Hours	11	80
Signers	9	11
Sentences	7000	35191
Val. Size	540	1741
Test Size	629	2322

Table 3: Statistics of the corpora used in the experiments. Source (Src.Lang.) and target (Tgt.Lang.) refer to the direction in which the corpora were created; all our experiments involve sign2text.

## B Infrastructure and Network Hyperparameters

We implement our SLT framework using PyTorch, and libraries from sBERT (Reimers and Gurevych, 2019) and Huggingface (Wolf et al., 2019). Our code is publicly available at Github.<sup>5</sup>

Tables 4, 5 and 6 show the hyperparameters and training times for the sign2sem and sem2text with SLTr and mBART transformers respectively. We run our experiments using 8 A100-80GB GPUs. For sign2sem2text and sign2(sem+text), each experiment runs for 72 hours and the configurations are inherited from the standalone modules sign2sem and sem2text.

Parameter	Value
model	all-MiniLM-L12-v2
batch_size_per_device	16
learning_rate	1e-5
input_projection_dim	1024
scheduler	warmuplinear
<b>Training time</b>	72 hours (5 GPU)

Table 4: Hyperparameters for the sign2sem module, we use the defaults of sBERT trainer for the rest.

## C Automatic Evaluation

Following Müller et al. (2022) and Müller et al. (2023), we evaluate the models using three common automatic metrics in machine translation: BLEU (Papineni et al., 2002), chrF (Popović, 2015) and BLEURT (Sellam et al., 2020). Notice that even though other semantic metrics based on embeddings might correlate better with human judge-

<sup>5</sup><https://github.com/yhamidullah/sem-slt>

Parameter	Value
num_encoder_layers	3
num_decoder_layers	3
d_model	512
ff_size	2048
input_projection_dim	1024
batch_size_per_device_train	32
batch_size_per_device_val	32
learning_rate	1e-5
lr_scheduler	reduceLROnPlateau
freeze_word_embeddings	True
<b>Training time</b>	1 hour (1GPU)

Table 5: Hyperparameters for the sem2text module with SLTr transformer, the rest are inherited from Camgöz et al. (2020).

Parameter	Value
input_projection_dim	1024
batch_size_per_device_train	4
batch_size_per_device_val	4
learning_rate	1e-5
fp16	True
freeze_word_embeddings	True
<b>Training time</b>	156 hours (8 GPUs)

Table 6: Hyperparameters for the sem2text module with mBART decoder, the rest are inherited from the Huggingface trainer default values.

ments (Kocmi et al., 2021; Freitag et al., 2022), they cannot be used for sign language translation because the source is video and not text. We use sacreBLEU (Post, 2018) for BLEU<sup>6</sup> and chrF<sup>7</sup> and the python library for BLEURT.<sup>8</sup>

Previous work starting with Camgöz et al. (2018) does mainly report only BLEU scores, but they do not specify the BLEU variant used or the signature in sacreBLEU. Therefore, comparisons among systems might not be strictly fair.

## D Gloss-based vs SEM-based Systems' Training Performance

Figure 3 shows the training evolution for a simple SLT system with no additional supervision (top), additional gloss supervision (middle) and SEM supervision (bottom) implemented in the Signjoey framework (Camgöz et al., 2020) and trained on PHOENIX-2014T. We use the best hyperparameters in Camgöz et al. (2020) and add our SEM supervision as a replacement of their recognition loss.

The three plots in Figure 3 include a red line at

<sup>6</sup>BLEU|nrefs:1|bs:1000|seed:16|case:mixed|eff:no|tok:13a|smooth:exp|version: 2.4.0

<sup>7</sup>chrF2|nrefs:1|bs:1000|seed:16|case:mixed|eff:yes|nc:6|nw:0|space:no|version: 2.4.0

<sup>8</sup>BLEURT v0.0.2 using checkpoint BLEURT-20.



Figure 3: Validation BLEU on PHOENIX without supervision (top plot), with gloss supervision (middle plot) and with SEM supervision (bottom plot).

translation quality BLEU=20 for reference. The first thing to notice is that both supervision methods reach the red line, but the one lacking any additional supervision lays behind. Second, we observe that the system with the additional SEM supervision reaches BLEU=20 earlier than the system with glosses: the gloss system needs 12k to finish and only 5k iterations are needed in the case of SEM. In both cases, we use early stopping with BLEU patience 7. Finally, notice that the gloss and SEM systems achieve the same translation quality but one does not need any data annotation with SEM.

## E Ablation Study on Sentence Length

Following the analysis of Section 5, we include the translation quality scores BLEU and chrF per sentence length.

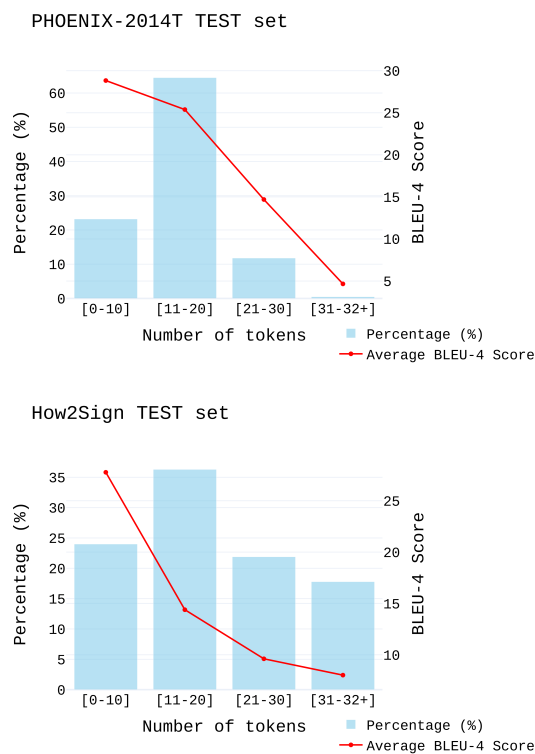


Figure 4: Variation of the average BLEU score on different token length intervals on PHOENIX-2014T (top) and How2Sign (bottom) test sets.




Figure 5: Variation of the average chrF score on different token length intervals on PHOENIX-2014T (top) and How2Sign (bottom) test sets.