

# EvSign: Sign Language Recognition and Translation with Streaming Events

Pengyu Zhang<sup>1,2\*</sup>, Hao Yin<sup>1</sup>, Zeren Wang<sup>1</sup>, Wenyue Chen<sup>1</sup>,  
Shengming Li<sup>1</sup>, Dong Wang<sup>1</sup>, Huchuan Lu<sup>1</sup>, and Xu Jia<sup>1†</sup>

<sup>1</sup> Dalian University of Technology

<sup>2</sup> National University of Singapore

**Abstract.** Sign language is one of the most effective communication tools for people with hearing difficulties. Most existing works focus on improving the performance of sign language tasks on RGB videos, which may suffer from degraded recording conditions, such as fast movement of hands with motion blur and textured signer’s appearance. The bio-inspired event camera, which asynchronously captures brightness change with high speed, could naturally perceive dynamic hand movements, providing rich manual clues for sign language tasks. In this work, we aim at exploring the potential of event camera in continuous sign language recognition (CSLR) and sign language translation (SLT). To promote the research, we first collect an event-based benchmark **EvSign** for those tasks with both gloss and spoken language annotations. EvSign dataset offers a substantial amount of high-quality event streams and an extensive vocabulary of glosses and words, thereby facilitating the development of sign language tasks. In addition, we propose an efficient transformer-based framework for event-based SLR and SLT tasks, which fully leverages the advantages of streaming events. The sparse backbone is employed to extract visual features from sparse events. Then, the temporal coherence is effectively utilized through the proposed local token fusion and gloss-aware temporal aggregation modules. Extensive experimental results are reported on both simulated (PHOENIX14T) and EvSign datasets. Our method performs favorably against existing state-of-the-art approaches with only **0.34%** computational cost (0.84G FLOPS per video) and **44.2%** network parameters. The project is available at <https://zhang-pengyu.github.io/EVSign>.

**Keywords:** Sign Language Recognition · Sign Language Translation · Event Camera

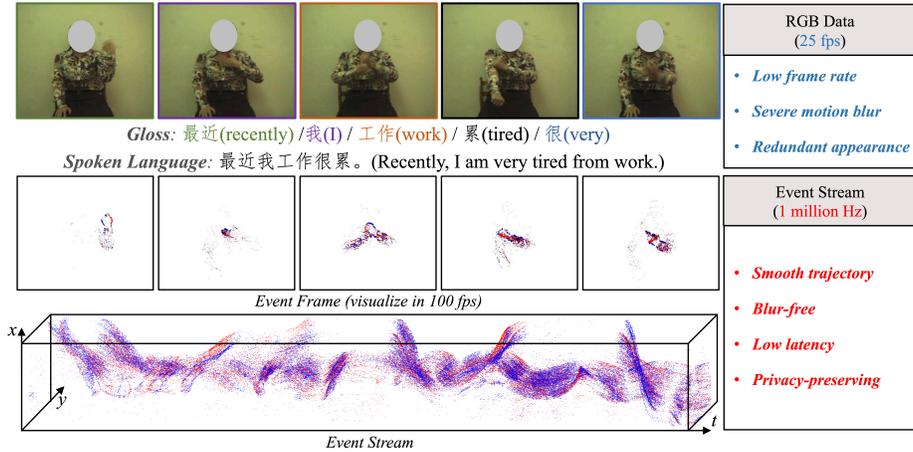
## 1 Introduction

As a main communication medium employed by the deaf community, sign language conveys multi-cue information by manual features (*e.g.*, hand shape, movement and pose) and non-manuals features, including facial expression, mouth

---

\* Work was done at Dalian University of Technology.

† Corresponding author: [jiayushenyang@gmail.com](mailto:jiayushenyang@gmail.com)



**Fig. 1:** Comparison between sign language recognition and translation with RGB and event data. We provide the first benchmark for event-based CSLR and SLT tasks, namely EvSign. Compared with RGB data, event stream can capture smooth movement within microsecond-level response, avoiding motion blur. Furthermore, the sparse event only stresses on the moving targets, such as hands and arms, which can be processed efficiently and protects personal privacy (facial information).

gesture and other body movements [2, 45]. According to the outputs, sign language based tasks can be mainly categorized into sign language recognition (SLR) and sign language translation (SLT). In SLR, the minimal lexical component, namely gloss, is predicted. SLR can be further categorized into isolated SLR (ISLR) and continuous SLR (CSLR). The purpose of SLT is to fully translate sign language into spoken language, which is often considered as a sequence-to-sequence learning problem. Recent works are based on videos captured by conventional frame-based sensors, which suffer from challenging scenarios, including severe motion blur and cluttering distractors. As shown in Fig. 1(a), the information will be degraded in extreme conditions, such as fast hand and arm movement, thereby leading to limited performances. To this end, existing works [8, 14, 20, 42, 64] emphasize temporal modeling, which leverages temporal cues in pixel- and frame-wise representations via 3D convolution and Long-Short Term Memory (LSTM) networks. Furthermore, STMC [64] and CorrNet [20] are designed to construct discriminative spatial representation, focusing on body trajectories.

Event camera, a biologically-inspired sensor, detects the variation of intensity along time. Rather than encoding visual appearance with still images, it generates sparse and asynchronous event stream with extremely high temporal resolution [3] (1M Hz *vs.* 120 Hz), high dynamic range and low latency, which is ideally suited for extracting motion cues [43, 61, 65]. As shown in Fig. 1(b), event camera could benefit sign language tasks from four perspectives. First, event data can capture richer motion information, thereby facilitating limb movements modeling effectively. Second, conventional images may be degraded due to the rapid movements, while event camera can record the sharp boundary

for further processing. Third, the event stream contains less redundant information, such as background and textured clothing, which can boost efficiency and avoid distractor interference. Fourth, from a privacy protection perspective, event cameras can avoid collecting static facial information. There have been several attempts [50, 53, 54, 56] to leverage event camera for sign language tasks. However, current event-based sign language datasets [53, 54, 56] only provide sign videos for ISLR. The limited vocabulary size and frame length cannot meet the requirements of real-world applications. Furthermore, the designed methods [50, 56] are based on the networks originally designed for frame sequences, such as AlexNet [33] and ResNet [16], which do not fully leverage the advantage of event data.

To unveil the power of event-based sign language tasks, we collect an event-based Chinese sign language benchmark, namely EvSign. To the best of our knowledge, it is the first dataset designed for event-based CSLR and SLT tasks. More than 6.7K high-resolution event streams are collected in EvSign, which is of comparable scale to the existing RGB-based SLR datasets. The large corpus with native expressions, precise gloss and spoken language annotations can promote the development of CSLR and SLT tasks.

Moreover, we present a transformer-based framework to make full use of the advantage of event data. First, a sparse backbone is employed to efficiently compute visual features on event data to obtain visual features efficiently. Then, temporal information is modeled via the proposed local token fusion and gloss-aware temporal aggregation modules, where the visual tokens are firstly combined to model local motion and reduce the computational cost. Subsequently, we aggregate the temporal information during the whole video into fused tokens hierarchically, which is then used for gloss and word prediction. Our method achieves very competitive performance on both synthetic and real datasets for both tasks with only 0.34% computational cost.

To sum up, our contributions can be concluded as three aspects:

- We propose the first benchmark for event-based CSLR and SLT tasks, which contains high-quality event streams, comprehensive corpus and precise gloss and spoken language annotation.
- We design a transformer-based algorithm for both tasks, which fully leverages the characteristics of event data.
- Experiments on both synthesized PHOENIX14T and EvSign datasets demonstrate that our method achieves favorable performance with only 0.34% FLOPS and 44.2% parameters against existing algorithms.

## 2 Related work

### 2.1 RGB-based sign language recognition

Sign language recognition can be categorized into two main directions: isolated SLR (ISLR) [17, 18] and continuous SLR (CSLR) [5, 15, 64, 67]. As for ISLR, word-level prediction is performed, while CSLR aims to predict a series of glosses

from longer-term videos, which has become the primary focus of research due to its closer alignment with real-world applications. To tackle CSLR, researchers mainly work on extracting discriminative spatial and temporal information. In the early stage, hand-crafted features [11, 45] are utilized to extract spatio-temporal representation. HMM-based algorithms [28, 30–32] are designed to predict gloss progressively. Recently, deep-based frameworks have been proposed, focusing on leveraging motion-aware information [18, 44, 49, 64]. CTC loss [13] is adopted to temporally align the classification results with unsegmented sequences, achieving end-to-end training. In addition, some works explore the use of other modalities to provide a complementary cue to visual signals, including skeleton [7, 10, 25, 48, 67] and depth [34, 46, 62], *etc.* In this paper, we exploit the effectiveness of a bio-inspired sensor, *i.e.*, event camera, for CSLR task.

## 2.2 RGB-based sign language translation

Compared to SLR, Sign Language Translation (SLT) aims to generate spoken language translations from sign language videos in a progressive manner. Camgoz *et al.* [4] first introduce SLT task and formalize it into Neural Machine Translation in an end-to-end setting, which extracts gloss features through a CNN model and SLT using a sequence-to-sequence model. Subsequent works [5, 6, 36, 60, 63] have focused primarily on how to better extract spatial and temporal features. Some studies have attempted gloss-free methods [39, 57, 59] to generate sentences without relying on gloss-level features, extensive experiments have shown that directly implementing an end-to-end Sign2Text model yields inferior results compared to using glosses as the intermediate supervision in the Sign2Gloss2Text model. Therefore, the current implementation of SLT task is predominantly based on the Sign2Gloss2Text approach. Recent works mainly focus on vision based SLT, while SLT with other modalities have not been fully exploited.

## 2.3 Event-based sign language recognition

Event camera captures the intensity variation of each pixel, recording the trajectory of fast-moving objects at high temporal resolution. Due to its property, it can provide sufficient temporal information, which is suitable for modeling object motion. A few attempts contribute to ISLR [50, 53, 54, 56]. Vasudevan *et al.* [53, 54] propose an event-based Spanish sign language dataset, namely SL-Animals-DVS, consisting of 1,102 event streams regarding animals. Two Spiking Neural Networks (SNN) [51, 58] are used for evaluation. Wang *et al.* [56] consider the event camera as a novel sensor in ISLR and collect an American sign language dataset, which contains 56 words. Shi *et al.* [50] design an event sampling strategy to select key event segments according to event distribution. The selected events are then fed to a CNN to obtain classification results. They also provide a synthetic dataset N-WLASL, where the event is collected by shooting an LCD monitor to record the videos from WLASL [35]. Above all, three main challenges limit the development of SLR. First, existing benchmarks are in small vocabulary, which cannot fully exploit the potential of sign language recognition.

**Table 1:** Summary of existing sign language recognition and translation benchmarks.

Dataset	Lang.	Gloss Vocab.	Text Vocab.	Num. Videos	Num. Signer	Continuous	SLT	Resolution	Source
WLASL [26]	ASL	1,000	–	25,513	222	×	×	Variable	Web
DEVISIGN [55]	CSL	2,000	–	24,000	8	×	×	598 × 448	Lab
PHOENIX-14 [29]	DGS	1,081	–	6,841	9	✓	×	210 × 260	TV
CCSL [23]	CSL	178	–	25,000	50	✓	×	1280 × 720	Lab
SIGNUM [1]	DGS	455	–	15,075	25	✓	×	776 × 578	Lab
PHOENIX-14T [4]	DGS	1,066	2,887	8,257	9	✓	✓	210 × 260	TV
CSL-Daily [63]	CSL	2,000	2,343	20,654	10	✓	✓	1920 × 1080	Lab
Youtube-ASL [52]	ASL	–	60,000	11,093	2,519	×	✓	Variable	Web
EvASL [56]	ASL	56	–	11,200	10	×	×	128 × 128	Lab
SL-Animals-DVS [54]	SSL	19	–	1,121	59	×	×	128 × 128	Lab
<b>EvSign(Ours)</b>	CSL	1,387	1,947	6,773	9	✓	✓	640 × 480	Lab

Second, the event data is collected using the out-of-date sensors with low spatial resolution in those datasets, leading to missing details in hand gesture and subtle movement. Third, all the datasets are designed for ISLR. It can solely be used for specific applications and cannot be generalized in real scenes. Therefore, it is crucial to collect a larger-scale dataset to promote the development of sign language tasks.

### 3 EvSign benchmark

#### 3.1 Benchmark Statistics

We use the DVXplorer-S-Duo camera from iniVation, which is a binocular camera capable of simultaneously capturing both event and RGB data. The spatial size of event stream is 640 × 480. We also record the RGB data with the size of 480 × 320 at 25 FPS for visualization and annotation.

To fit the practical usage, the corpus is sourced around daily life, such as shopping, education, medical care, travel and social communication, etc. The glosses are sampled from the Chinese national sign language dictionary [9] and CSL-Daily [63], and are then reorganized into a spoken sentence. To avoid differences in expression, we further provide glosses and sentences to signers for adjustment to suit the deaf community.

We recruit 9 professional volunteers from the deaf community, who are familiar with general sign language for data collection. We employ a two-step manner to avoid data ambiguity. When collecting sign data, the signers first watch a reference video and then start to perform the action. After recording, other three signers vote to determine whether the sign expression is precise and easy to understand. For each sample in the corpus, there are about three signers to perform the action. We separate the sign videos into training, development and test subsets, which contain 5,570, 553 and 650 clips, respectively. As shown in Table 1, the proposed dataset significantly surpasses existing datasets in vocabulary size, task scope, and data resolution, which provides a comprehensive corpus to exploit the power of event data in handling sign language tasks.

**Table 2:** Annotation statistics of EvSign dataset.

		Segments	Frame	Duration(h)	Vocab.	Avg. words	Tot. words	OOVs	Singletons
Sign Gloss	Train	5,570	606.7K	6.74	1,348	5.09	28,387	–	230
	Dev	553	75.1K	0.83	695	5.68	3,145	25	–
	Test	650	87.2K	0.97	723	5.48	3,562	18	–
Chinese	Train				1,825	7.76	43,276	–	407
	Dev	↑ same	↑ same	↑ same	880	8.58	4,746	79	–
	Test				912	8.23	5,350	53	–

### 3.2 Annotation

In EvSign, both sign gloss and spoken language annotations are provided. First, annotators identify all the glosses according to [9] in the RGB videos. We note that several signs may express the same meaning. Thus, the authors further revise the annotation to ensure that each sign language corresponds to a unique gloss annotation. Finally, the spoken language annotations are updated according to the gloss annotation. We employ tokenization method in HanLP<sup>1</sup> to separate a sentence into words. As shown in Table 2, EvSign provides 1.3K unique signs and 1.8K words, which cover various aspects of our daily life. Furthermore, more than 35K and 53K gloss and text annotations are totally labeled.

### 3.3 Evaluation Metrics

We provide two evaluation protocols for both SLR and SLT. As for SLR evaluation, we use Word Error Rate (WER) as the metric, which is widely used in sign language and speech recognition. WER measures the similarity of reference and hypothesis, which is based on the minimum number of operations required to convert the prediction into the reference sentence as:

$$\text{WER} = \frac{\#\text{sub} + \#\text{ins} + \#\text{del}}{\#\text{ref}} \quad (1)$$

where  $\#\text{sub}$ ,  $\#\text{ins}$  and  $\#\text{del}$  are the number of basic operations, including substitution, insertion and deletion.  $\#\text{ref}$  represents the number of words in the reference sentence. Lower WER indicates better performance. For SLT evaluation, we employ ROUGE [37] and BLEU [47] as evaluation metrics. Here, BLEU is calculated with n-grams from 1 to 4 and ROUGE-L [38] is used as our metric. The higher the ROUGE and BLEU scores, the better the performance.

## 4 Methodology

### 4.1 Overview

As shown in Fig. 2, the event stream  $\mathcal{E} = \{e_i\}_{i=1}^N$  is firstly split to  $P$  segments evenly and converted to a set of event representation  $\mathbf{E} = \{\mathbf{E}_1, \mathbf{E}_2, \dots, \mathbf{E}_P\} \in$

<sup>1</sup> <https://github.com/hankcs/HanLP>

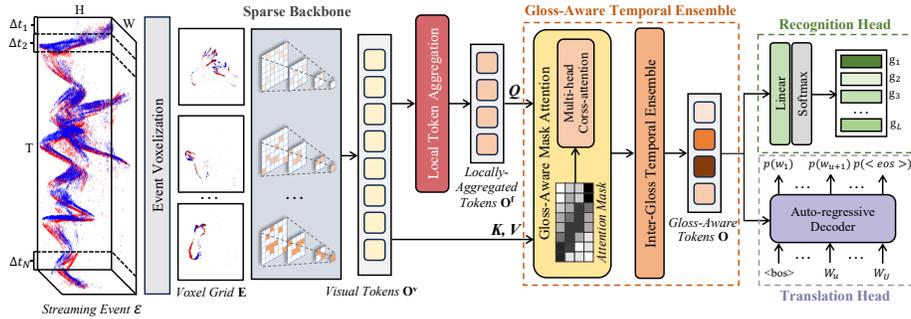


Fig. 2: Pipeline of the transformer-based framework for CSLR and SLT tasks.

$\mathbb{R}^{P \times B \times H \times W}$ , *i.e.*, voxel grid [66], where  $B$  denotes the bin size in voxel grid. Taken  $\mathbf{E}$  as input, the proposed network is to jointly predict all the glosses  $\mathcal{G} = \{g_1, g_2, \dots, g_Z\}$  and translate them into spoken language  $\mathcal{W} = \{w_1, w_2, \dots, w_U\}$  in a sequence-to-sequence manner.  $Z$  and  $U$  are the number of glosses and words in the sign language video. Our method contains five main parts, including sparse backbone (SConv), Local Token Fusion (LTF), Gloss-Aware Temporal Aggregation (GATA) and two task heads, *i.e.*, recognition head and translation head. First, SConv generates a set of visual tokens. Then, we employ LTF to fuse local motion from adjacent timestamps thereby reducing the number of tokens. Moreover, the temporal information is decoupled into intra-gloss and inter-gloss cues, which are learned hierarchically by GATA module. Specifically, we propose a gloss-aware mask attention to dynamically fuse the comprehensive motion information from visual tokens into the fused tokens. It measures the token’s similarity in time and feature spaces, which can be aware of various action lengths. Furthermore, the global coherence among tokens from different glosses is learned via inter-gloss temporal aggregation, thereby obtaining the gloss-aware tokens. Finally, those tokens are sent to recognition and translation heads to predict the probability of target gloss sequence and spoken language.

## 4.2 Overall framework

**Sparse backbone (SConv).** Due to the sparse property of event data, we build a sparse convolutional network [12, 40] with the architecture of ResNet18 [16]. The backbone can process the event representation to obtain the visual tokens  $\mathbf{O}^v = \{\mathbf{o}_1^v, \mathbf{o}_2^v, \dots, \mathbf{o}_P^v\} \in \mathbb{R}^{P \times C}$ . The sparse backbone is able to fully leverage the characteristics of data sparsity, thus significantly reducing the computational load. Compared to regular convolution layers, sparse backbone can also better maintain the sparsity at feature level, leading to a sharper boundary [12].

**Local Token Fusion (LTF).** Local motion integration is crucial for long-term temporal modeling, which can first build an effective representation for continuous actions within a short duration and reduce computational load [19]. To construct a powerful features for local motion, we introduce LTF to fuse neigh-

boring visual tokens, which contains two multi-head self-attention within local window (W-MSA) [41] and two max pooling (MaxPool) layers. A regular window partitioning scheme is adopted to all the visual tokens, where each window covers  $I$  tokens. The self-attention within non-overlapping window is calculated to aggregate local motion information. Then, we introduce a max pooling operation to reduce the number of tokens. We adopt another W-MSA and MaxPool with the same window size and ratio to capture a longer-term movement, thus obtaining the fused tokens  $\mathbf{O}^f \in \mathbb{R}^{L \times C}$ . The LTF can be formulated as,

$$\begin{aligned}\tilde{\mathbf{O}}^v &= \text{MaxPool}(\text{W-MSA}(\mathbf{O}^v) + \mathbf{O}^v) \\ \mathbf{O}^f &= \text{MaxPool}(\text{W-MSA}(\tilde{\mathbf{O}}^v) + \tilde{\mathbf{O}}^v)\end{aligned}\quad (2)$$

where the number of tokens  $L = \frac{P}{\gamma}$ , while  $\gamma$  represents the downsampling ratio. **Gloss-Aware Temporal Aggregation (GATA)**. Global temporal modeling is a key step to exploit the correspondence of continuous signs in long-term videos. Existing methods [18, 20, 42] learn the global motion cue by 1D temporal convolution and BiLSTM, which ignores the varying durations of different signs. Simply applying temporal modeling among the fixed frames will learn a non-optimal representation and involve redundant information from different glosses. In this work, we decouple the temporal information into intra-gloss and inter-gloss cues and model them hierarchically via the proposed GATA module, which consists of Gloss-Aware Mask Attention (GAMA) and Inter-Gloss Temporal Aggregation (IGTA). As for intra-gloss temporal aggregation, we aim to aggregate the gloss-level information from  $\mathbf{O}^v$  into  $\mathbf{O}^f$ . To achieve this, we propose GAMA by introducing a gloss-aware mask  $\mathbf{M} \in \mathbb{R}^{L \times P}$ ,

$$\text{GAMA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}, \mathbf{M}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}} \odot \mathbf{M}\right)\mathbf{V}\quad (3)$$

where,  $\mathbf{Q}$ ,  $\mathbf{K}$ ,  $\mathbf{V}$  are the query, key and value defined in cross-attention.  $d$  is the dimension of query or key. We claim that the tokens belonging to the same class tend to have highly-relevant representations. The mask  $\mathbf{M}$  can be considered as an attention weight, which measures the similarity between the fused and visual tokens. Thus, we first calculate the token similarity  $\rho \in \mathbb{R}^{L \times P}$ ,

$$\rho = \psi_f(\mathbf{O}^f)\psi_v(\mathbf{O}^v)^T\quad (4)$$

where,  $\psi_f()$  and  $\psi_v()$  are the linear embedding functions for the fused and visual tokens. Also, we add a distance constrain  $\delta \in \mathbb{R}^{L \times P}$  in time space to avoid computing attention between different glosses of the same category,

$$\delta_{i,j} = \mathcal{K}(t_i^f, t_j^v; \sigma)\quad (5)$$

where  $\mathcal{K}()$  is a Radial Basis Function (RBF) kernel with the parameter  $\sigma$ . Since the precise timestamp for each token is not accessible, we introduce pseudo timestamps  $t_i^f, t_j^v \in \mathbb{R}^1$  to represent the relative temporal position for fused and visual tokens, respectively. The pseudo timestamp for  $j$ -th visual token is set to

$t_j^y = j$ . For the fused tokens, we calculate the pseudo timestamp as the average of the pseudo timestamps of the fused tokens, defined as  $t_i^f = \sum_{x=i \times \gamma}^{x=(i+1) \times \gamma} x$ . Finally, the mask is defined as  $\mathbf{M} = \mathcal{N}(\rho) \odot \mathcal{N}(\delta)$ , where the  $\odot$  represents the element-wise product.  $\mathcal{N}$  is a zero-one normalization.

We extend the gloss-aware mask attention into multiple heads and all the attention heads share the same mask  $\mathbf{M}$ . In addition, a Feed Forward Network (FFN) module is utilized following GAMA to enhance the representation, which consists of two linear transformations with ReLU activation. Following SLT [5], we add a positional encoding (PE) process to the inputs. Thus, the intra-gloss temporal aggregation process can be summarized as follows,

$$\begin{aligned} \mathbf{Q} &= \mathbf{O}^f + \mathbf{P}_q; \quad \mathbf{K} = \mathbf{O}^v + \mathbf{P}_k; \quad \mathbf{V} = \mathbf{O}^v \\ \hat{\mathbf{O}}^f &= \mathbf{O}^f + \text{GAMA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) \\ \tilde{\mathbf{O}}^f &= \hat{\mathbf{O}}^f + \text{FFN}(\hat{\mathbf{O}}^f) \end{aligned} \quad (6)$$

where  $\mathbf{P}_q \in \mathbb{R}^{L \times C}$  and  $\mathbf{P}_k \in \mathbb{R}^{P \times C}$  are the temporal positional encodings generated by a predefined sine function.

After obtaining intra-gloss tokens  $\tilde{\mathbf{O}}^f$ , we apply inter-gloss temporal aggregation to model the global motion via a multi-head self-attention (MSA),

$$\begin{aligned} \mathbf{X} &= \tilde{\mathbf{O}}^f + \mathbf{P}_x \\ \mathbf{O} &= \tilde{\mathbf{O}}^f + \text{MSA}(\mathbf{X}) \end{aligned} \quad (7)$$

where the  $\mathbf{P}_x \in \mathbb{R}^{L \times C}$  is also the temporal positional encoding. After applying GATA, we obtain the gloss-aware tokens  $\mathbf{O} \in \mathbb{R}^{L \times C}$ , which learns the motion cues among all glosses comprehensively and is sent to the following task heads for predicting glosses and words.

**Task heads.** Following the existing methods, we employ a classifier  $\mathcal{F}_{\text{reg}}$  as Recognition Head (RH) to predict the logits  $\mathbf{G} = \{\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_L\} \in \mathbb{R}^{L \times Y}$ , where  $Y$  denotes the size of the sign language vocabulary with adding a ‘blank’ class.  $\mathcal{F}_{\text{reg}}$  consists of a fully-connected layer with a softmax activation. As for handling translation task, the translation head (TH) is to sequentially generate logits of spoken language sentences  $\mathbf{W} = \{p(w_1), p(w_2), \dots, p(w_U), p(\langle eos \rangle)\}$  conditioned by the gloss-aware tokens  $\mathbf{O}$ , which is an auto-regressive transformer decoder. We adopt the same translation head in [5]. Due to the page limitation, Please refer to [5] and the supplementary material for details.

## 5 Experiments

### 5.1 Datasets and evaluation protocol

**Dataset.** We conduct analysis on synthetic (PHOENIX14T) and real (EvSign) event-based sign language benchmarks. As an extension of PHOENIX14 [29], PHOENIX14T introduces German spoken language annotation and has become the primary benchmark for CSLR and SLT. It consists of 8,247 videos with

a vocabulary of 1,066 and 2,887 sign and words, respectively. The number of videos for training, development and test are 7,096, 519 and 642. Based on PHOENIX14T, we build a synthetic dataset for further research using an event simulator (V2E [22]), where the video frames are firstly interpolated to 350 frames per second by SuperSloMo [24] and used for event generation.

**Evaluation protocol.** We focus on both SLR and SLT tasks as follows:

- **SLR:** It only predicts the sign gloss from sign language videos without introducing any additional information. Specifically, we implement a variant (Ours<sub>S2G</sub><sup>2</sup>) by removing the translation branch and retrain our method for fair comparison. In this setting, we select four public-available methods including VAC [42], CorrNet [20], TLP [19], SEN [21] for comparison. All the methods are trained from scratch according to the original setting. We present WER as a metric for evaluation.
- **SLT:** We follow the Sign2 (Gloss+Text) protocol defined in [5], which is to jointly learn both recognition and translation branches in an end-to-end manner. We compare Ours<sub>S2GT</sub> with SLT [5]. Additionally, we equip the recent SLR algorithm with our translation head (CorrNet+TH, VAC+TH), which are trained under this SLT comparison. Both ROUGE-L and BLEU-X metrics are employed to quantitatively assess these methods.

## 5.2 Training Details

The feature dimension  $C$  is 1024. The downsample ratio  $\gamma$  in LTF module is 4. We adopt the widely-used CTC loss [13] for SLR supervision. Inspired by VAC [42], we set an additional recognition branch  $\mathcal{F}'_{\text{reg}}$  on the fused tokens  $\mathbf{O}^f$ . Two CTC losses ( $\mathcal{L}_{\text{inter}}$  and  $\mathcal{L}_{\text{final}}$ ) are applied to both intermediate and final outputs against the ground truth. The total loss  $\mathcal{L}_{\text{SLR}}$  for SLR protocol is,

$$\mathcal{L}_{\text{SLR}} = \lambda_{\text{inter}}\mathcal{L}_{\text{inter}} + \lambda_{\text{final}}\mathcal{L}_{\text{final}} \quad (8)$$

Under the SLT protocol, we add a cross-entropy loss  $\mathcal{L}_{\text{ce}}$  to supervise the output of translation head. Thus the overall loss of our method can be summarized as,

$$\mathcal{L}_{\text{SLT}} = \mathcal{L}_{\text{SLR}} + \lambda_{\text{ce}}\mathcal{L}_{\text{ce}} \quad (9)$$

The weights for those three losses  $\lambda_{\text{inter}}$ ,  $\lambda_{\text{final}}$  and  $\lambda_{\text{ce}}$  are all set to 1. The parameters  $\sigma$  in RBF is set to 16 in our experiment. The bin size of voxel grid  $B$  is 5. We adopt Adam [27] optimizer with cosine annealing strategy to adjust the learning rate. The initial learning rate, weight decay and batch size are set to  $3e^{-5}$ , 0.001 and 2, respectively. We train our method for 200 epochs to achieve convergence. To align the setting with existing methods, all the RGB frames are cropped to  $320 \times 320$  to remove the useless background and then resized to

<sup>2</sup> We use subscript to indicate which task is focused. S2G denotes the method is trained and evaluated solely on SLR task while S2GT represents the method is trained with the supervision of both gloss and spoken language and then used for SLT evaluation.

**Table 3:** Comparison results for SLR on PHOENIX14T and EvSign datasets. **Bold** and underline denotes the top-two ranking performance.

Method	Modal	PHOENIX14T		EvSign		FLOPS	Param (M)
		Dev (%)	Test (%)	Dev (%)	Test (%)		
VAC [42]	RGB	20.17	21.60	32.08	30.43	228.87G	31.64
TLP [19]	RGB	19.40	21.20	33.70	32.96	231.28G	59.69
SEN [21]	RGB	19.50	21.00	33.26	33.46	231.96G	34.70
CorrNet [20]	RGB	18.90	20.50	32.37	32.04	234.59G	32.04
VAC [42]	EV	24.99	24.77	30.84	30.71	<u>238.88G</u>	<u>31.65</u>
TLP [19]	EV	24.81	24.60	32.59	32.68	240.08G	59.69
SEN [21]	EV	24.63	<u>24.51</u>	33.34	32.71	242.00G	34.70
CorrNet [20]	EV	<u>24.57</u>	24.55	<u>29.98</u>	<u>29.95</u>	244.63G	32.05
<b>Ours<sub>S2G</sub></b>	EV	<b>23.89</b>	<b>24.03</b>	<b>29.19</b>	<b>28.69</b>	<b>0.84G</b>	<b>14.19</b>

$256 \times 256$ . As for event data, we firstly generate the voxel grid. Then, the voxel grid is cropped to  $480 \times 480$  and resized to  $256 \times 256$ . Other competitors are trained using their own settings. Note that we modify the input channel of the first convolutional layer to fit the event input. All the models are trained and tested on a single NVIDIA RTX 3090 GPU with 24G RAM.

### 5.3 Quantitative Results on Sign Language Recognition

Table 3 provides quantitative results on PHOENIX14T and EvSign datasets. Compared to existing methods that utilize RGB data, all the algorithms working with streaming event show lower WER consistently on EvSign dataset with real event, revealing the power of event stream in handling sign language recognition. However, the advantages of event data are not reflected in the results on the simulated dataset. The reason is that the video frames used for event synthesis are of poor quality with severe blur and limited frame rate. Compared to event-based competitors, our method shows superior performance on both synthetic PHOENIX14T and EvSign datasets. Notably, our method has significant advantage in both computational cost and number of parameters, due to the concise architecture and sparse data processing. We also evaluate the computational efficiency using FLOPS and number of parameters (Params) on EvSign dataset<sup>3</sup>. Compared to the most recent method (CorrNet [20]), Ours<sub>S2G</sub> achieves 0.79% and 1.26% improvement with respect to WER on development and test sets of EvSign with only 0.34% FLOPS and 44.2% parameters.

### 5.4 Quantitative Results on Sign Language Translation

Table 4 and 5 show the comparison results for SLT on PHOENIX14T and EvSign datasets.

<sup>3</sup> Since the FLOPS and Params fluctuate based on data sparsity and sequence lengths, we calculate their averages across videos in Dev and Test sets of EvSign.

**Table 4:** Comparison results for SLT on synthesized PHOENIX14T. (R: ROUGE-L, B-X: BLEU-X)

Method	Modal	PHOENIX-14T Dev					PHOENIX-14T Test				
		R	B-1	B-2	B-3	B-4	R	B-1	B-2	B-3	B-4
SLT [5]	RGB	41.35	40.06	29.73	23.24	18.86	38.69	38.26	28.36	22.06	18.06
VAC+TH [42]	RGB	39.38	39.22	29.47	23.44	19.40	38.57	39.87	29.53	23.11	19.00
CorrNet+TH [20]	RGB	39.82	40.05	30.34	24.04	19.86	40.26	41.23	31.42	24.86	20.48
SLT [5]	EV	<u>39.86</u>	38.99	29.01	23.54	18.41	39.21	39.84	29.25	23.09	19.23
VAC+TH [42]	EV	39.01	39.09	<u>29.64</u>	<b>23.74</b>	<u>19.64</u>	39.34	39.62	29.55	23.30	<u>19.92</u>
CorrNet+TH [20]	EV	39.52	<u>39.28</u>	29.60	23.41	19.24	<u>40.12</u>	<b>40.77</b>	<b>30.62</b>	<u>24.11</u>	19.83
<b>Ours<sub>S2GT</sub></b>	EV	<b>40.23</b>	<b>39.37</b>	<b>29.66</b>	<u>23.67</u>	<b>19.83</b>	<b>40.21</b>	<u>40.40</u>	<u>30.47</u>	<b>24.27</b>	<b>20.07</b>

**Table 5:** Comparison results for SLT on EvSign. (R: ROUGE-L, B-X: BLEU-X)

Method	Modal	EvSign Dev					EvSign Test					FLOPS	Param (M)
		R	B-1	B-2	B-3	B-4	R	B-1	B-2	B-3	B-4		
SLT [5]	RGB	39.75	39.64	23.75	15.80	10.86	40.05	39.84	23.54	15.60	10.63	242.71G	34.51
VAC+TH [42]	RGB	38.54	38.37	22.58	15.00	9.21	39.08	38.74	23.90	15.88	10.19	233.70G	45.26
CorrNet+TH [20]	RGB	38.05	38.74	22.41	13.98	9.04	39.41	39.45	23.74	15.68	10.57	239.46G	45.66
SLT [5]	EV	<u>39.92</u>	39.06	23.54	<u>15.89</u>	<b>11.21</b>	<u>41.54</u>	40.13	24.36	16.04	10.87	252.99G	<u>34.51</u>
VAC+TH [42]	EV	38.96	38.93	23.29	15.23	10.08	39.48	39.22	24.11	15.94	10.01	<u>243.84G</u>	45.26
CorrNet+TH [20]	EV	39.55	<u>39.58</u>	<u>24.09</u>	15.69	10.50	41.23	<u>40.85</u>	<u>25.34</u>	<u>16.95</u>	<u>11.83</u>	249.59G	45.66
<b>Ours<sub>S2GT</sub></b>	EV	<b>40.98</b>	<b>42.00</b>	<b>25.75</b>	<b>16.89</b>	<u>11.20</u>	<b>42.43</b>	<b>41.44</b>	<b>25.61</b>	<b>17.55</b>	<b>12.37</b>	<b>6.99G</b>	<b>28.06</b>

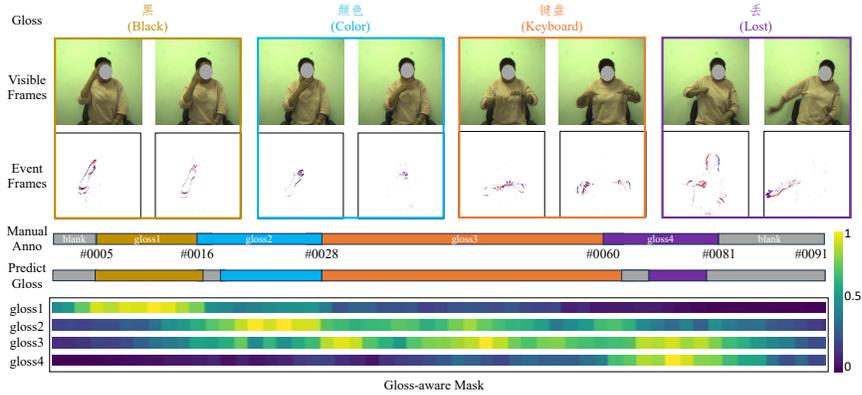
As shown in Table 5, our method achieves the best performance except BLEU-4 in the development set. On the other hand, the results on Phoenix14T (Table 4) cannot demonstrate the effectiveness of event camera. Compared with SLT, our method achieves 1.06% and 0.89% improvement in terms of ROUGE in development and test sets, respectively. Specifically, our method also exhibits significant advantages in terms of computational and parameter efficiency. Furthermore, methods with event streams are consistently better than those with RGB frames, which demonstrates the potential of event data in SLT task.

## 5.5 Further Analysis

**Ablation study.** As shown in Table 6, we conduct ablation analysis on both PHOENIX14T and EvSign. we introduce a modified VAC [42] as our baseline (B), with removing the visual alignment loss. Compared to the baseline with RGB input, the event-based baseline achieves better performance on real event, which reveals the potential of event in handling CSLR tasks. All of the proposed modules contribute positively to the recognition performance on both datasets. The sparse backbone (SConv) can significantly drop the computational load and parameters while maintaining recognition accuracy, which can fully leverage the sparsity of event data. The simple yet effective LTF module obtains 0.47% and 0.53% on the test set of PHOENIX14T and EvSign, respectively. Compared to the BiLSTM, the designed GATA module can learn the temporal cues more comprehensively, leading to 0.81% and 1.12% WER decrease on the test sets.

**Table 6:** Ablation analysis for SLR on PHOENIX14T and EvSign.

Methods	PHOENIX14T		EvSign		FLOPS	Params (M)
	Dev (%)	Test (%)	Dev (%)	Test (%)		
B(RGB)	20.38	21.74	32.00	30.29	228.87G	31.64
B(EV)	24.92(-4.54%)	25.19(-3.45%)	30.91(+1.09%)	30.08(+0.27%)	238.88G	31.65
B(EV)+SConv	24.68(+0.24%)	25.38(-0.19%)	30.21(+0.70%)	30.01(+0.07%)	1.22G	21.90
B(EV)+SConv+LTF	24.33(+0.35%)	24.91(+0.47%)	29.79(+0.42%)	29.48(+0.53%)	0.72G	13.13
B(EV)+SConv+GATA	24.04(+0.64%)	24.57(+0.81%)	29.41(+0.80%)	28.89(+1.12%)	1.69G	26.10
B(EV)+SConv+LTF+GATA	23.89(+0.15%)	24.03(+0.54%)	29.19(+0.22%)	28.69(+0.20%)	0.84G	14.19

**Fig. 3:** Visualization of the gloss-aware mask on EvSign dataset.

The final model obtains the best performance with regard to all the metrics. Compared to B(EV), the final model achieves 1.16% and 1.39% improvement on the test sets, which can serve as a strong baseline for further research.

**Visualization of gloss-aware mask.** As shown in Fig. 3, we provide a visualization results of SLR task and the gloss-aware mask in GATA module. With the guidance of GATA module, our method learns comprehensive motion cues, thus predicting all the glosses correctly. We visualize the gloss-aware mask between the tokens in corresponding gloss and visual tokens. It demonstrates that the gloss-aware mask can precisely provide an intra-gloss correlation without any supervision, achieving gloss-aware temporal token aggregation.

**Analysis of token aggregation strategy.** We compare several token aggregation strategies and the performance on EvSign dataset is shown in Table 7(a). We implement the ‘w/o aggregation’ by removing the selection module, where the visual tokens are directly sent to GATA modules to output the probability of gloss. The inferior performance indicates the necessity of token aggregation module. The useless information may lead to a delete or insert error, which significantly affects the recognition performance. We also compare the simplest selection strategies, which are denoted as MaxPooling and AvgPooling. Those methods can decrease the WER, while selection in a soft manner (AvgPooling) is better than the hard one (MaxPooling). This indicates the local fusion is necessary for learning discriminative tokens for temporal modeling. Both 1D-CNN designed in [42] and our method apply local aggregation before selection, leading

**Table 7:** Ablation analysis of the proposed modules on EvSign dataset.

(a) Ablation analysis of LTF.			(b) Ablation analysis of GATA.		
	EvSign(EV)			EvSign(EV)	
	Dev (%)	Test (%)		Dev (%)	Test (%)
w/o aggregation	40.95	40.34	w/o GAMA	30.38	31.14
MaxPooling	34.44	33.24	$\delta$ -only	29.89	29.83
AvgPooling	32.27	31.56	$\rho$ -only	<u>29.63</u>	29.74
1D-CNN [42]	<u>29.41</u>	<u>28.89</u>	GAMA-hard	29.74	<u>29.31</u>
LTF(Ours)	<b>29.19</b>	<b>28.69</b>	GAMA-Soft (Ours)	<b>29.19</b>	<b>28.69</b>

to favorable performances. Our method achieves 0.22% and 0.2% improvement than 1D-CNN, which shows the effectiveness of LTF module.

**Analysis of temporal aggregation module.** Temporal information module is the key component in sign language tasks. To this end, we compare various methods to demonstrate the capabilities of our method in temporal modeling. As shown in Table 7(b), we set our method without intra-gloss temporal aggregation (GAMA) as baseline. The method solely relies on inter-gloss temporal modeling via global self-attention.  $\rho$ -only and  $\delta$ -only GAMA denote that the mask  $\mathbf{M}$  in Eq. (3) is set to  $\rho$  and  $\delta$  in Eq. (4) and Eq. (5). Compared with the baseline,  $\rho$ -only and  $\delta$ -only GAMA achieve 1.4% and 1.31% performance gain in the test set. Furthermore, we compare the effectiveness of soft and hard mask in GAMA. GAMA-hard means that we binarize the learned  $\mathbf{M}$  with a threshold  $1e^{-3}$ . we find that the hard mask will lead to a performance decrease. Soft manners exhibit greater flexibility, resulting in improved aggregation performance.

## 6 Conclusion

In this paper, we unveil the power of events in sign language tasks. We first build a comprehensive dataset for event-based sign language recognition and translation. The dataset contains more than 6.7K sign videos captured by high-quality event camera, which covers most of daily topics. It can greatly promote the development of event-based vision and sign language related tasks. Furthermore, we propose a transformer-based framework for both tasks by fully exploiting the sparsity and high temporal resolution characteristics of events. The proposed gloss-aware temporal aggregation module could effectively model temporal information in a global-and-local manner and a gloss-aware representation is computed for SLR and SLT tasks. Our method shows favorable performance on the sign language datasets with synthetic and real events with only 0.34% FLOPS and 44.2% network parameters.

### Acknowledgment

The research was partially supported by the National Natural Science Foundation of China, (grants No. 62106036, U23B2010, 62206040, 62293540, 62293542), the Dalian Science and Technology Innovation Fund (grant No. 2023JJ11CG001).

## References

1. Agris, U.V., Kraiss, K.F.: Towards a video corpus for signer-independent continuous sign language recognition. In: International Gesture Workshop (2007)
2. Boyes-Braem, P., Sutton-Spence, R.: The Hands are the Head of the Mouth: The Mouth as Articulator in Sign Languages. Gallaudet University Press (2001)
3. Brandli, C., Berner, R., Yang, M., Liu, S., Delbrück, T.: A  $240 \times 180$  130 db 3  $\mu$ s latency global shutter spatiotemporal vision sensor. *IEEE J. Solid State Circuits* **49**(10), 2333–2341 (2014)
4. Camgöz, N.C., Hadfield, S., Koller, O., Ney, H., Bowden, R.: Neural sign language translation. In: CVPR (2018)
5. Camgoz, N.C., Koller, O., Hadfield, S., Bowden, R.: Sign language transformers: Joint end-to-end sign language recognition and translation. In: CVPR (2020)
6. Chen, Y., Wei, F., Sun, X., Wu, Z., Lin, S.: A simple multi-modality transfer learning baseline for sign language translation. In: CVPR (2022)
7. Chen, Y., Zuo, R., Wei, F., Wu, Y., Liu, S., Mak, B.: Two-stream network for sign language recognition and translation. In: NeurIPS (2022)
8. Cui, R., Liu, H., Zhang, C.: A deep neural framework for continuous sign language recognition by iterative training. *IEEE TMM* **21**(7), 1880–1891 (2019)
9. of the Deaf, C.A., for Sign Language, N.C., Research, B.: National Sign Language Dictionary. HuaXia Press (2019)
10. Duarte, A.C., Palaskar, S., Ventura, L., Ghadiyaram, D., DeHaan, K., Metze, F., Torres, J., Giró-i-Nieto, X.: How2sign: A large-scale multimodal dataset for continuous american sign language. In: CVPR (2021)
11. Freeman, W.T., Roth, M.: Orientation histograms for hand gesture recognition. In: FG (1995)
12. Graham, B., van der Maaten, L.: Submanifold sparse convolutional networks. *CoRR* **abs/1706.01307** (2017)
13. Graves, A., Fernández, S., Gomez, F.J., Schmidhuber, J.: Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In: ICML (2006)
14. Guo, L., Xue, W., Guo, Q., Liu, B., Zhang, K., Yuan, T., Chen, S.: Distilling cross-temporal contexts for continuous sign language recognition. In: CVPR (2023)
15. Hao, A., Min, Y., Chen, X.: Self-mutual distillation learning for continuous sign language recognition. In: ICCV (2021)
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
17. Hu, H., Zhao, W., Zhou, W., Wang, Y., Li, H.: Signbert: Pre-training of hand-model-aware representation for sign language recognition. In: ICCV (2021)
18. Hu, H., Zhou, W., Li, H.: Hand-model-aware sign language recognition. In: AAAI (2021)
19. Hu, L., Gao, L., Liu, Z., Feng, W.: Temporal lift pooling for continuous sign language recognition. In: ECCV (2022)
20. Hu, L., Gao, L., Liu, Z., Feng, W.: Continuous sign language recognition with correlation network. In: CVPR (2023)
21. Hu, L., Gao, L., Liu, Z., Feng, W.: Self-emphasizing network for continuous sign language recognition. In: AAAI (2023)
22. Hu, Y., Liu, S., Delbrück, T.: v2e: From video frames to realistic DVS events. In: CVPRW (2021)

23. Huang, J., Zhou, W., Zhang, Q., Li, H., Li, W.: Video-based sign language recognition without temporal segmentation. In: AAAI (2018)
24. Jiang, H., Sun, D., Jampani, V., Yang, M., Learned-Miller, E.G., Kautz, J.: Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation. In: CVPR (2018)
25. Jiao, P., Min, Y., Li, Y., Wang, X., Lei, L., Chen, X.: Cosign: Exploring co-occurrence signals in skeleton-based continuous sign language recognition. In: ICCV (2023)
26. Joze, H.R.V., Koller, O.: MS-ASL: A large-scale data set and benchmark for understanding american sign language. In: BMVC (2019)
27. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
28. Koller, O., Camgöz, N.C., Ney, H., Bowden, R.: Weakly supervised learning with multi-stream cnn-lstm-hmms to discover sequential parallelism in sign language videos **42**(9), 2306–2320 (2020)
29. Koller, O., Forster, J., Ney, H.: Continuous sign language recognition: Towards large vocabulary statistical recognition systems handling multiple signers. CVIU **141**, 108–125 (2015)
30. Koller, O., Ney, H., Bowden, R.: Deep hand: How to train a CNN on 1 million hand images when your data is continuous and weakly labelled. In: CVPR (2016)
31. Koller, O., Zargaran, S., Ney, H.: Re-sign: Re-aligned end-to-end sequence modelling with deep recurrent cnn-hmms. In: CVPR (2017)
32. Koller, O., Zargaran, S., Ney, H., Bowden, R.: Deep sign: Hybrid CNN-HMM for continuous sign language recognition. In: BMVC (2016)
33. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NeurIPS (2012)
34. Kumar, P., Gauba, H., Roy, P.P., Dogra, D.P.: A multimodal framework for sensor based sign language recognition. Neurocomputing **259**, 21–38 (2017)
35. Li, D., Opazo, C.R., Yu, X., Li, H.: Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In: WACV (2020)
36. Li, D., Xu, C., Yu, X., Zhang, K., Swift, B., Suominen, H., Li, H.: Tspnet: Hierarchical feature learning via temporal semantic pyramid for sign language translation. In: NeurIPS (2020)
37. Lin, C.Y.: ROUGE: A package for automatic evaluation of summaries. In: ACL (2004)
38. Lin, C., Och, F.J.: Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In: ACL (2004)
39. Lin, K., Wang, X., Zhu, L., Sun, K., Zhang, B., Yang, Y.: Gloss-free end-to-end sign language translation. In: ACL (2023)
40. Liu, B., Wang, M., Foroosh, H., Tappen, M.F., Pensky, M.: Sparse convolutional neural networks. In: CVPR (2015)
41. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV (2021)
42. Min, Y., Hao, A., Chai, X., Chen, X.: Visual alignment constraint for continuous sign language recognition. In: ICCV (2021)
43. Mitrokhin, A., Fermüller, C., Parameshwara, C., Aloimonos, Y.: Event-based moving object detection and tracking. In: IROS (2018)
44. Niu, Z., Mak, B.: Stochastic fine-grained labeling of multi-state sign glosses for continuous sign language recognition. In: ECCV (2020)
45. Ong, S.C.W., Ranganath, S.: Automatic sign language analysis: A survey and the future beyond lexical meaning. TPAMI **27**(6), 873–891 (2005)

46. Oszust, M., Krupski, J.: Isolated sign language recognition with depth cameras. In: KES (2021)
47. Papineni, K., Roukos, S., Ward, T., Zhu, W.: BLEU: a method for automatic evaluation of machine translation. In: ACL (2002)
48. Parelli, M., Papadimitriou, K., Potamianos, G., Pavlakos, G., Maragos, P.: Spatio-temporal graph convolutional networks for continuous sign language recognition. In: ICASSP (2022)
49. Pu, J., Zhou, W., Li, H.: Iterative alignment network for continuous sign language recognition. In: CVPR (2019)
50. Shi, Q., Ye, Z., Wang, J., Zhang, Y.: Qisampling: An effective sampling strategy for event-based sign language recognition. *Signal Process. Lett.* **30**, 768–772 (2023)
51. Shrestha, S.B., Orchard, G.: SLAYER: spike layer error reassignment in time. In: NeurIPS (2018)
52. Uthus, D., Tanzer, G., Georg, M.: Youtube-asl: A large-scale, open-domain american sign language-english parallel corpus. *CoRR* **abs/2306.15162** (2023)
53. Vasudevan, A., Negri, P., Ielsi, C.D., Linares-Barranco, B., Serrano-Gotarredona, T.: Sl-animals-dvs: event-driven sign language animals dataset. *Pattern Anal. Appl.* **25**(3), 505–520 (2022)
54. Vasudevan, A., Negri, P., Linares-Barranco, B., Serrano-Gotarredona, T.: Introduction and analysis of an event-based sign language dataset. In: FG (2020)
55. Wang, H., Chai, X., Hong, X., Zhao, G., Chen, X.: Isolated sign language recognition with grassmann covariance matrices. *ACM Trans. Access. Comput.* **8**(4), 14:1–14:21 (2016)
56. Wang, Y., Zhang, X., Wang, Y., Wang, H., Huang, C., Shen, Y.: Event-based american sign language recognition using dynamic vision sensor. In: WASA (2021)
57. Wong, R., Camgöz, N.C., Bowden, R.: Sign2gpt: Leveraging large language models for gloss-free sign language translation (2024)
58. Wu, Y., Deng, L., Li, G., Zhu, J., Shi, L.: Spatio-temporal backpropagation for training high-performance spiking neural networks. *CoRR* **abs/1706.02609** (2017)
59. Yin, A., Zhong, T., Tang, L., Jin, W., Jin, T., Zhao, Z.: Gloss attention for gloss-free sign language translation. In: CVPR (2023)
60. Zhang, B., Müller, M., Sennrich, R.: SLTUNET: A simple unified model for sign language translation. In: ICLR (2023)
61. Zhang, J., Yang, X., Fu, Y., Wei, X., Yin, B., Dong, B.: Object tracking by jointly exploiting frame and event domain. In: ICCV (2021)
62. Zheng, L., Liang, B.: Sign language recognition using depth images. In: ICARCV (2016)
63. Zhou, H., Zhou, W., Qi, W., Pu, J., Li, H.: Improving sign language translation with monolingual data by sign back-translation. In: CVPR (2021)
64. Zhou, H., Zhou, W., Zhou, Y., Li, H.: Spatial-temporal multi-cue network for continuous sign language recognition. *AAAI* (2020)
65. Zhu, A.Z., Yuan, L., Chaney, K., Daniilidis, K.: Ev-flownet: Self-supervised optical flow estimation for event-based cameras. In: *Robotics: Science and Systems* (2018)
66. Zhu, A.Z., Yuan, L., Chaney, K., Daniilidis, K.: Unsupervised event-based learning of optical flow, depth, and egomotion. In: CVPR (2019)
67. Zuo, R., Mak, B.:  $C^2$ slr: Consistency-enhanced continuous sign language recognition. In: CVPR (2022)