

# Continuous Sign Language Recognition using Multimodal Input and Handshape-aware Boundary Detection

Mingyu Zhao<sup>1</sup>, Zhanfu Yang<sup>1</sup>, Yang Zhou<sup>1</sup>, Zhaoyang Xia<sup>1</sup>, Can Jin<sup>1</sup>,  
Xiaoxiao He<sup>1</sup>, Shuhang Lin<sup>1</sup>, Carol Neidle<sup>2</sup>, Dimitris N. Metaxas<sup>1</sup>

<sup>1</sup>Rutgers University, <sup>2</sup>Boston University

<sup>1</sup>110 Frelinghuysen Road, Piscataway, NJ 08854, USA

<sup>2</sup>Boston University Linguistics, 111 Cummington Mall, Boston, MA 02215, USA

{zhao.mingyu, zhanfu.yang, eta.yang, zx149, can.jin}@rutgers.edu

{xiaoxiao.he, shuhang.lin}@rutgers.edu, dnm@cs.rutgers.edu, carol@bu.edu

## Abstract

This paper employs a multimodal approach for continuous sign recognition by first using ML for detecting the start and end frames of signs in videos of American Sign Language (ASL) sentences, and then by recognizing the segmented signs. For improved robustness, we use 3D skeletal features extracted from sign language videos to take into account the convergence of sign properties and their dynamics that tend to cluster at sign boundaries. Another focus of this paper is the incorporation of information from 3D hand configuration for boundary detection. To detect handshapes normally expected at the beginning and end of signs, we pretrain a handshape classifier for detection of 87 linguistically defined canonical handshape categories using a dataset that we created by integrating and normalizing several existing datasets. A multimodal fusion module is then used to unify the pretrained sign video segmentation framework and handshape classification models. Finally, the estimated boundaries are used for sign recognition, where the recognition model is trained on a large database containing both citation-form isolated signs and signs pre-segmented (based on manual annotations) from continuous signing—as such signs often differ a bit in certain respects. We evaluate our method on the ASLLRP corpus and demonstrate significant improvements over previous work.

**Keywords:** sign language recognition, boundary detection, handshape, multimodal fusion, skeletal features

## 1. Introduction

Signed languages are fully developed visual-gestural languages used by Deaf and hard-of-hearing communities worldwide. Unlike spoken languages, they convey meaning through hand shapes, orientations, movements, facial expressions, head gestures, and body posture (Brentari, 2019). American Sign Language (ASL) is the primary language used by more than 500,000 people in the United States (ASL-deafined, 2017), and it is used in other parts of North America, as well. It is also ranked as the third most studied non-native language studied in US colleges (Lusin et al., 2023). Despite advances in automatic speech and natural language processing, Sign Language Recognition (SLR) remains a complex problem for 2 main reasons: (1) the high-dimensional spatio-temporal dynamics of gestures, and (2) a general lack of large-scale annotated datasets (Huang et al., 2018). Much research to date has focused on Isolated Sign Language Recognition (ISLR), where each input clip contains exactly one sign. However, in real-world communication, signed languages are produced as continuous streams of signing, not isolated signs.

Compared with ISLR, Continuous Sign Language Recognition (CSLR) is more complex, in part because of significant differences in sign production in continuous signing as compared with

citation forms (Neidle, 2023). Among many complicating factors in continuous signing are the potential independence of signing on the 2 hands, and the tendency for sign production to be influenced by the properties of adjacent segments. Note that in the current research, all of the signers are right-handed, and we are only attending to signs that include production by the dominant (i.e., right) hand, i.e., 1-handed signs produced on the dominant hand and 2-handed signs. We are not considering 1-handed signs produced on the non-dominant hand.

A major challenge in CSLR lies in the difficulty of locating precise temporal boundaries for each sign segment, i.e., identifying the start and end frames corresponding to each sign (Neidle, 2023). Another major difficulty arises from the fact that natural signing includes things that cannot just be looked up in a “dictionary” of isolated signs (even if they could be perfectly segmented)—including potentially: false starts, things like classifiers and gestures, not all of which belong to any finite classes of elements, as well as fingerspelled signs, potentially corresponding to any word or proper noun (name) from the spoken language (often with letters missing or transposed). This is why any ultimate strategy for recognition of signs from continuous signing needs to be able to distinguish among types of signs, as different recognition strategies need to be applied to different sign types (Yanovich et al., 2016). Note, however, that the research reported here does not

do this, so we have submitted our segmented signs for sign recognition only in cases where we have at least 6 examples of those signs in our training dataset. Thus, we are not currently carrying out end-to-end sign recognition, although that is the long-term goal. The sign recognition component of this research is intended to demonstrate the usefulness of segmentation for sign recognition.

Much of the prior work relies on relatively small datasets (Zhang and Duh, 2023; Zhang et al., 2026), which achieve better results compared to the use of larger datasets for experimentation (Gonzalez and Collet, 2011; Moryossef et al., 2023, e.g.). Other work (Rastgoo et al., 2024, e.g.) attempts to synthesize pseudo-continuous sign language datasets by concatenating isolated sign videos. However, relatively few of these strategies have been applied to real large-scale continuous sign language datasets for boundary detection (Moryossef et al., 2023). Most approaches to sign boundary detection follow a similar 2-stage pipeline. They typically first extract visual features from raw videos using tools such as OpenPose (skeletons) or 3D convolutional neural networks, and then feed these features into deep learning models such as LSTMs or Transformers for temporal modeling and boundary prediction (Rastgoo et al., 2024; Renz et al., 2021; Moryossef et al., 2023). These methods have shown some success, but they generally neglect 2 important perspectives: (1) multimodal insights from patterns involving handshape, motion, velocity, acceleration, facial expressions, upper body position, and head gestures, known to correlate with sign transitions, which have the potential to provide strong signals for boundary detection; and (2) the use of more specialized skeleton-based approaches for pose modeling, as some of methods rely on general-purpose deep models and overlook structured methods that may be better suited to use skeletal data.

Building on these observations, we propose a multimodal boundary detection framework, incorporating a pretrained handshape classification model with a sign boundary detection model. We use the estimated boundaries for sign recognition in a well-trained sign recognition model.

There is no standard written representation for signs. Researchers often use English-based ID glossing to identify signs, as we also do. However, a major challenge is ensuring consistency in gloss labeling, i.e., a 1-1 correspondence between signs and glosses, and many datasets fail to meet this requirement (Neidle et al., 2022a, e.g.). One major advantage of the datasets we are using is that we ensure consistency in gloss representations across all the datasets we use.

Our contributions are summarized as follows:

- We design a segmentation module based on

a spatio-temporal convolutional network, incorporating velocity and acceleration features.

- We propose a novel method to enhance boundary detection by incorporating handshape information. A multimodal fusion module is introduced, which integrates features from a pretrained handshape classifier into the segmentation stream using a cross-attention mechanism with a gating mechanism.
- We present a new CSLR pipeline combining a segmentation model with a sign language recognition model. Experiments demonstrate the effectiveness of our segmentation framework.

## 2. Related Work

### 2.1. Isolated Sign Language Recognition

Recent approaches to Isolated Sign Language Recognition (ISLR) have leveraged both visual and skeletal modalities to recognize signs from trimmed video clips. Early methods mostly relied on handcrafted features and Hidden Markov Models (HMMs) (Starnier and Pentland, 1998). With the rise of deep learning, convolutional neural networks (CNNs) (Pigou et al., 2015) and hierarchical attention networks (HANs) (Huang et al., 2018) became widely used to model spatial and temporal dependencies in isolated signs.

More recent approaches adopt BiLSTM (Camgöz et al., 2017), Transformer-based architectures (Rastgoo et al., 2024), and graph-based models (e.g., ST-GCN (Zhou et al., 2024)) to capture spatio-temporal dynamics. Some work also explores multimodal fusion (Zhou et al., 2020), combining RGB, depth, optical flow, and skeletal data to improve recognition accuracy. In particular, skeletal features extracted from pose estimation models offer a compact and robust representation of sign gestures, leading to improved performance in data-limited scenarios (Li et al., 2020).

Despite impressive progress, ISLR methods typically assume well-segmented input clips, limiting applicability to real-world CSL scenarios. This motivates the development of methods that can extend to continuous sign recognition.

### 2.2. Continuous Sign Language Recognition

Continuous Sign Language Recognition (CSLR) aims to segment and recognize signs from natural sign language sentences. This has important potential applications to real-world scenarios, as it could provide useful input to facilitate such things as live interpretation or video captioning.

Early approaches relied more on alignment-based strategies, which offered a good way to model temporal sequences. Vogler and Metaxas

(1999; 2001) proposed parallel HMM frameworks that modeled the simultaneous movements of multiple articulators (e.g., hands and facial expressions), enabling recognition of continuous signing without requiring manual segmentation. Such work laid the foundation for structured temporal modeling in CSLR. Deep learning methods introduced alignment-free techniques such as Connectionist Temporal Classification (CTC) (Graves et al., 2006), which allowed sequence-level supervision without the need for frame-level annotations.

Later research introduced iterative alignment schemes using RNNs or TCNs to refine temporal segmentation. More recently, to address the intrinsic ambiguity in gloss-to-frame alignment, several studies propose incorporating linguistic priors and contextual constraints. For example, Zhang et al. (2023) introduce a gloss-conditioned decoder integrating prior gloss context into the decoding process. Guo et al. (2023) enhance temporal modeling via cross-temporal knowledge distillation. Researchers have also introduced gloss-level embeddings as guidance for visual feature learning (Guo et al., 2024; Zhang et al., 2025).

### 2.3. Sign Boundary Detection

For CSLR, it is essential to identify the precise temporal boundaries of individual signs within continuous video streams. Early approaches treated boundary detection as a frame-wise classification problem or relied on dynamic programming to align gloss-level annotations with video frames (Gonzalez and Collet, 2011). Still, these methods overlook the movement dynamics of continuous signing and sign boundaries.

Subsequent methods leveraged temporal convolutional networks, such as TCNs (Renz et al., 2021), or recurrent neural networks, such as LSTM models (Jui et al., 2022), to better capture dynamic dependencies in signing sequences for estimation of sign boundaries. Moryossef et al. (2023) later encoded pose features, which they used as input into an LSTM; they applied a greedy segmentation decoder to identify sign segments.

Recently, Transformer-based architectures have been proposed for sign boundary detection. Varol et al. (2021) use transformers for boundary detection in CSL. Bull et al. (2021) utilize BERT (Devlin et al., 2019) for boundary detection. Rastgoo et al. (2024) propose a Transformer-based framework for boundary detection, but simplify the problem by using artificially constructed CSL sentences rather than natural sign language sentences. Rastgoo et al. (2025) employ a GCN-based model to extract skeletal features that are then embedded into a Transformer-based encoder for boundary detection. These transformer-based approaches exhibit strong performance in modeling multiframe dependencies,

but still fail to integrate richer information, such as motion patterns and acceleration, which could provide additional signals for more accurate boundary detection and recognition.

## 3. Method

### 3.1. Problem Formulation

In this work, we formulate the sign language segmentation task as a frame-wise "BIO" annotation problem over continuous sign language video sequences. That is: given a sequence of video frames  $X = \{x_1, x_2, \dots, x_N\}$ , the goal is to predict a corresponding label sequence  $L = \{l_1, l_2, \dots, l_N\} \in \{B, I, O\}^N$ , with the label "B" indicating the boundary frame of a sign segment, "I" indicating that the frame lies within a sign, specifically between the start frame and the end frame, and "O" indicating a frame that is outside any sign segment. This BIO annotation is better suited for modeling the continuous sign language segmentation task than traditional IO tagging, as sign segments in continuous sign language video often follow one another without any O frames in between, which results in these two adjacent segments being merged and counted as one segment (Moryossef et al., 2023).

### 3.2. Spatio-Temporal Convolutional Segmentation

We define  $S(x)$  as the skeleton extracted model (Fang et al., 2023). From a video frame sequence  $X = \{x_1, x_2, \dots, x_N\}$ , we extract the corresponding frame-wise 27 skeleton joint representations  $Y = \{Y_1, Y_2, \dots, Y_N\}$ , where  $Y_i = S(x_i)$ . To meet input requirements of the STGCN-based model, we pad all extracted skeleton sequences along the temporal dimension to a fixed length of  $T$  frames. Ground-truth labels are extended to include 4 classes: B, I, O, and P (padding), corresponding to label indices 3, 2, 1, and 0, respectively.

The extracted skeleton sequences  $Y = \{Y_1, Y_2, \dots, Y_T\}$  are then fed into a Spatio-Temporal Convolutional Segmentation Network. As shown in Figure 1, the front part of the network is constructed based on the ST-GCN framework proposed by Han et al. (2020), which effectively models the spatio-temporal dynamics of human motion. Specifically, we stack 10 ST-GCN layers to progressively capture high-level spatio-temporal patterns across the skeleton sequence.

The output features from the final ST-GCN layer are then passed into a segmentation module designed for frame-wise prediction. To aggregate information across joints, we first apply a spatial pooling operation over the 27 joints at each time step. The pooled features are then processed through 2 temporal convolutional layers, followed by a linear

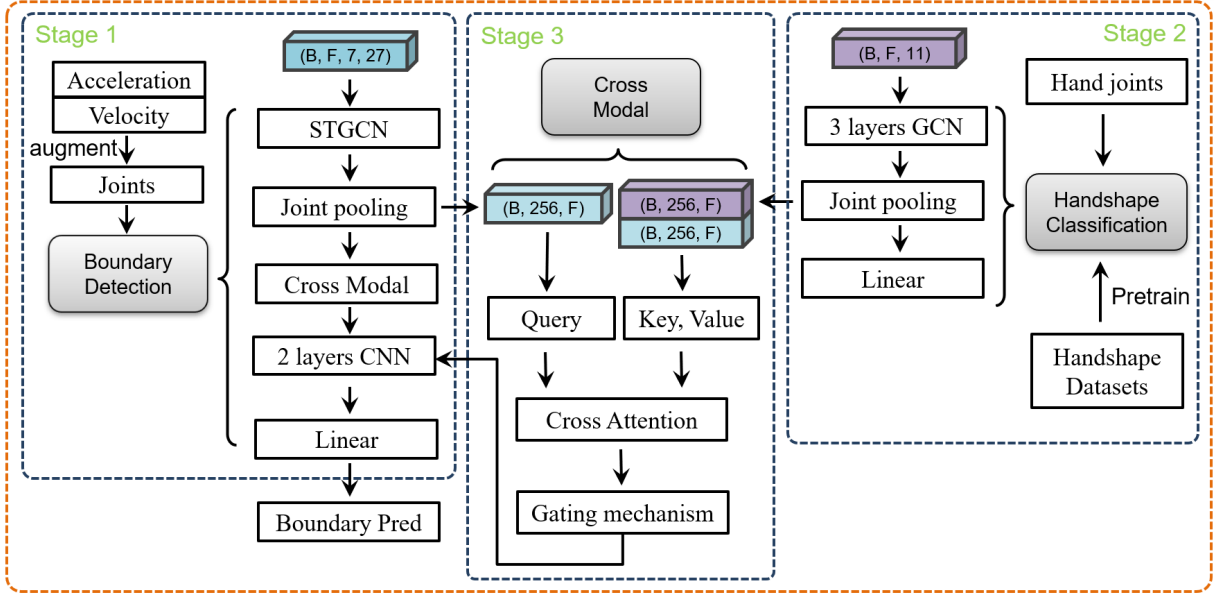


Figure 1: Overview of our proposed multimodal boundary detection. The left branch is the segmentation network based on ST-GCN and CNN modules, which processes all skeletal joints augmented with their velocity and acceleration. We pretrain this branch without the Cross-Modal layer in stage 1. The right branch shows the handshape classification module, which is pretrained on curated handshape datasets using a 3-layer GCN and later used to extract frame-wise handshape features. We pretrain this branch in stage 2. In the middle, a cross-modal attention module with a gating mechanism fuses the handshape and segmentation features to enhance temporal boundary prediction. We fine-tune the right branch with the cross-modal attention module in stage 3. The final segmentation output is further used to extract sign clips and then used for sign recognition.

projection layer that produces frame-level predictions over the segmentation labels. This design allows the model to first learn rich spatial and temporal dependencies using the ST-GCN backbone and then map the encoded motion features to fine-grained frame-wise segmentation outputs.

We adopt a combined loss function consisting of a frame-wise prediction loss and a boundary loss.

The frame-wise prediction loss is defined as a weighted cross-entropy loss. Specifically, we assign a higher class weight to boundary frames (class 3) and a lower weight to background frames (class 0), in order to emphasize the importance of accurate boundary detection and reduce the influence of dominant background frames. The loss is computed as:

$$\mathcal{L}_{\text{frame}} = - \sum_{t=1}^T w_{y_t} \cdot \log(\hat{y}_t[y_t]),$$

where  $\hat{y}_t \in \mathbb{R}^C$  is the predicted class probability at time  $t$ ,  $y_t \in \{0, 1, 2, 3\}$  is the ground-truth label, and  $w_{y_t}$  is the class-specific weight.

The boundary loss is defined as the absolute difference between the number of predicted boundaries and ground-truth boundaries, averaged over the utterance:

$$\mathcal{L}_{\text{boundary}} = \frac{1}{U} \sum_{u=1}^U |N_{\text{pred}}^{(u)} - N_{\text{true}}^{(u)}|,$$

where  $N_{\text{pred}}^{(u)}$  and  $N_{\text{true}}^{(u)}$  represent the number of predicted and true boundaries in the  $u$ -th utterance, and  $U$  is the number of all utterances.

The final loss  $\mathcal{L}_{\text{seg}}$  is a weighted combination of the 2 components:

$$\mathcal{L}_{\text{seg}} = \mathcal{L}_{\text{frame}} + \lambda \cdot \mathcal{L}_{\text{boundary}},$$

where  $\lambda$  is a hyperparameter controlling the contribution of the boundary loss.

### 3.3. Handshape Classification

We use the set of 87 handshapes from the NCSLGR (National Center for Sign Language and Gesture Resources, 2012) that can be linguistically significant. We formulate the handshape recognition task as a frame-wise classification problem over hand regions in sign video sequences. Given a video frame sequence  $X = \{x_1, x_2, \dots, x_N\}$ , we extract the corresponding frame-wise hand joint representations  $X_H = \{x_1^h, x_2^h, \dots, x_N^h\}$ . We use the established set of 87 canonical handshape categories, indexed from 0 to 86. Rather than assigning an absolute label to each frame, the goal is to estimate the degree of similarity between the observed handshape and these canonical categories. Thus, we want to find a handshape classifier denoted as  $H(x_H) \in \{0, 1, \dots, 86\}^N$ . However, in practice, some frames exhibit blurred or

occluded hand configurations. To address this, we manually sample a subset of clear frames to construct a high-quality training dataset denoted as  $X_H^S = \{x_{h_1^s}, x_{h_2^s}, \dots, x_{h_k^s}\}$ , where  $h_1^s, h_2^s, \dots, h_k^s$  are selected  $k$  frames. The handshape classifier  $H(x^H)$  is then pretrained on this dataset and integrated into our multimodal segmentation framework. See Section 4.1.1 for details of the handshape dataset and its construction.

For the handshape recognition model, since we operate directly on skeletal data, we adopt a straightforward 3-layer GCN architecture. This design is intuitive and intentionally kept shallow to mitigate overfitting, while still being expressive enough to capture the discriminative skeletal patterns of different handshapes, as shown in Figure 1. We use standard cross-entropy loss for training.

It is worth noting that our model is not intended to assign a definitive handshape label to every frame in a video (nor would it be realistic to attempt to do so). This is because almost all frames contain hand configurations that deviate in some way from the canonical handshapes. These canonical handshapes are idealized abstractions, whereas in practice, handshapes often fall somewhere in between. Instead, the purpose of our classifier is to extract handshape features by estimating the degree of similarity between the observed handshape and handshapes in the canonical set.

### 3.4. Multimodal Module

The goal of the multimodal module is to incorporate the additional per-frame information from the handshape classification module into the segmentation network. To do this, we introduce a unified cross-attention-based fusion module into the segmentation network; see Figure 1.

#### 3.4.1. Self-Attention

For context, we begin by reviewing the self-attention mechanism (Vaswani et al., 2017). Given an input sequence  $X \in \mathbb{R}^{T \times d}$ , the attention is computed by projecting it into query, key, and value spaces via learned matrices  $W_Q, W_K, W_V$ , and computing:

$$\text{Attn}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d}}\right)V, \quad (1)$$

where  $Q = W_Q X$ ,  $K = W_K X$ , and  $V = W_V X$ ,  $\sqrt{d}$  is the scaling factor computed from  $Q, K$  and  $V$ .

#### 3.4.2. Cross-Attention Module

We treat the segmentation features  $X^{\text{seg}} \in \mathbb{R}^{T \times d}$  as the source of the query, and concatenate the keys and values from both the segmentation and handshape streams to form the joint attention pool.

Let  $X^{\text{hand}} \in \mathbb{R}^{T \times d}$  be the pooled handshape features. We compute:

$$Q_c = W_Q X^{\text{seg}}, \\ K_c = \begin{bmatrix} W_K X^{\text{seg}} \\ W_K X^{\text{hand}} \end{bmatrix}, \quad V_c = \begin{bmatrix} W_V X^{\text{seg}} \\ W_V X^{\text{hand}} \end{bmatrix}. \quad (2)$$

The resulting cross-attention is then used to update the segmentation representation as follows:

$$\hat{X}^{\text{seg}} = X^{\text{seg}} + \sigma(g) \cdot \text{softmax}\left(\frac{Q_c K_c^\top}{\sqrt{d_c}}\right)V_c. \quad (3)$$

Here,  $\sigma(g)$  is a learnable scalar function with parameter  $g$  that modulates the fused output, and  $\sqrt{d_c}$  is a scaling factor (Vaswani et al., 2017).

This formulation enables the model to leverage both internal temporal dynamics and relevant external cues from the handshape stream, while the gating mechanism ensures that cross-modal information is selectively integrated in a learnable and data-dependent way.

### 3.5. Sign Recognition Using Boundaries from Segmentation

The predicted boundaries are used to segment the video into sign clips, among which only those segments that have been successfully associated with a ground truth sign and for which we have at least 6 examples in our training dataset are passed to a sign recognition model. Specifically, we integrate into our framework a sign classification model (Zhou et al., 2024), which is used to predict the gloss label for each segment correctly identified by our boundary detection model. Given the output of the boundary detection model, we first identify the corresponding segments from the input sequence. These segments are then passed to the sign language classification model for recognition, laying the groundwork for a future continuous sign language recognition pipeline (Zhou et al., 2024). We use this to show how good boundary segmentation is for sign recognition.

## 4. Experiments

### 4.1. Dataset Construction

#### 4.1.1. Handshape Recognition Dataset

We integrated 3 different ASLLRP ASL datasets (Neidle et al., 2022b)—two collections with isolated signs, ASLLVD (Neidle and Metaxas, 2024a; Neidle et al., 2012) and DSP (Neidle and Metaxas, 2024b), and one collection with continuous signing, ASLLRP Sentences (Neidle and Metaxas, 2025a) (henceforth ASLLRP-S)—to construct a unified dataset for handshape recognition. Note that the ASLLVD and ASLLRP-S datasets are shared on the Web from <https://dai.cs.rutgers.edu/dai/s/signbank>. Statistical information about

these datasets is available from <https://dai.cs.rutgers.edu/dai/s/runningstats>.

We also used a set of videos collected specifically to illustrate the 87 handshapes in motion from different camera angles and made available through the Boston University NCSLGR (National Center for Sign Language and Gesture Resources, 2012). For each handshape category, the NCSLGR videos provide synchronized multi-view video recordings for the 87 linguistically defined handshapes. We sampled at least 20 frames per handshape from different camera angles, enabling the model to learn cross-view representations of each handshape. In addition, we augmented the dataset by sampling 30 instances per handshape from the 3 other sign language datasets. For rare handshapes that do not meet the 30-sample threshold, we collected as many representative samples as possible. Through this process, we constructed a comprehensive handshape dataset that captures rich visual diversity across categories and views.

#### 4.1.2. Sign Recognition Dataset

For continuous sign data, we adopt the method used in Zhou et al. (2024). We constructed a unified sign classification dataset by integrating video clips from 6 different sources. We use 4 collections with citation-form ASL signs—WLASL (Li et al., 2020) (with modified glossing, as explained below), RIT (Neidle and Metaxas, 2024c), ASLLVD (Neidle and Metaxas, 2024a; Neidle et al., 2012), and DSP (Neidle and Metaxas, 2024b)—as well as 2 datasets in which signs had been extracted from ASL sentences based on manual annotations carried out using SignStream® (see <http://www.bu.edu/asllrp/SignStream/3>)—ASLLRP-S (Neidle and Metaxas, 2025a) and DSP Sentences (Neidle and Metaxas, 2025b) (hereafter DSP\_S).

As previously mentioned, it is critically important that the datasets for this research all make use of consistent glossing conventions. To ensure this, we have modified some of the gloss ID labels for the WLASL dataset to conform to the conventions used for the other datasets (Neidle et al., 2022a; Dafnis et al., 2022; Neidle and Ballard, 2022b,a). Note also that we are relying on “class label” ID glosses to represent signs, which group together minor variants in sign production (Neidle and Opoku, 2022; Neidle and Ballard, 2022a).

We exclude from the sentence datasets signs designated as “hidden”. This label identified false starts and signs that otherwise seriously deviate from their canonical articulation. For further details, see <https://dai.cs.rutgers.edu/asllvd/signbank/Additional-information-6-29-25.pdf>. We used the hidden signs for purposes of segmentation, but excluded them from the sign recognition research. We further restrict the

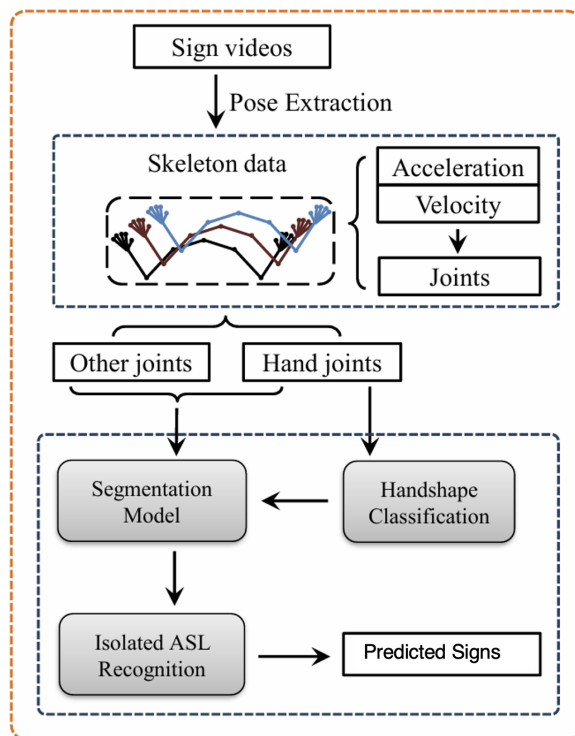


Figure 2: Overview of the experimental procedure. We begin by extracting 2D skeletal keypoints from sign language videos using a pose estimation tool. We use velocity and acceleration information to augment the joint feature for each frame. These are then separated into two branches: hand joints are fed into a pretrained handshape classification model, while all joints are used as input to a segmentation model. The outputs of both branches are fused to produce improved segmentation boundaries. For sign recognition, the segmented sign clips are further passed into a sign recognition model trained on both citation-form signs and signs segmented from continuous signing, to produce final sign predictions.

recognition research to signs that were correctly segmented and for which our training set includes at least 6 examples.

#### 4.1.3. Boundary Detection Dataset

For boundary detection, we use all of the pre-segmented clips from the ASLLRP-S (Neidle and Metaxas, 2025a) dataset, including the “hidden signs,” as described in the previous section. Training for segmentation was performed on the basis of the annotations for the start and end frames of each sign.

## 4.2. Experiment Setup

### 4.2.1. Data Preprocessing

For our constructed datasets, we first extract skeletal features, including keypoints of the torso, arms,

and hands, from RGB frames using the AlphaPose model (Fang et al., 2023). We also incorporate velocity and acceleration information derived from the keypoint sequences to enrich the temporal dynamics. For both the boundary detection and continuous sign language recognition tasks, we employ a skeletal graph consisting of 27 nodes. In contrast, for the pretrained handshape recognition model, we utilize a skeletal graph constructed solely from hand joints.

In addition, we apply normalization and data augmentation techniques inspired by Jiang et al. (2024), Boháček and Hruz (2022), Zhou et al. (2022), and Jin et al. (2025). Specifically, we normalize the skeleton keypoints to the range  $[-1, 1]$  to enhance cross-dataset transferability. To further improve model robustness, we also apply random horizontal flipping as a data augmentation strategy during training.

#### 4.2.2. Experiment Details

To ensure a fair comparison with the baseline (Renz et al., 2021), we adopt the same 4:1 train-test split ratio and randomly partition the ASLLRP-S dataset. As shown in Figure 2, our overall pipeline consists of skeletal feature extraction, a pretrained handshape classification model, and segmentation and recognition modules. We fine-tune the model with LoRA (Jin et al., 2025) for parameter-efficient adaptation. For the boundary detection task, the model is trained only up to the segmentation module to predict temporal boundaries. For the sign language recognition task, the recognition results of each correctly predicted segment are used to further optimize the boundary detection module. All models are trained on 8 NVIDIA RTX 8000 GPUs.

#### 4.2.3. Evaluation Metrics

We define different evaluation metrics for different tasks. For boundary detection, to facilitate comparison with prior related research, we adopt mF1B (mean F1 score of boundary) and mF1S (mean F1 score of segment) proposed by Renz et al. (2021). The mF1B metric measures the mean F1 score for boundary detection, where a predicted boundary is considered correct if its distance to a ground-truth boundary is less than a predefined threshold. The mF1S metric measures the mean F1 score for sign segmentation, where a predicted sign segment is considered correct if its intersection-over-union (IoU) with a ground-truth segment exceeds a predefined threshold. We adopt the same threshold settings as the baseline paper (Renz et al., 2021), to ensure fair comparison.

We also propose a more intuitive evaluation metric. Specifically, we introduce a tolerance-based matching criterion to evaluate the alignment be-

tween predicted and ground-truth sign segments, which we will refer to as Boundary Tolerance. A predicted segment is considered correct if both its start frame and end frame are within a predefined tolerance of the ground-truth start and end frames, respectively. The tolerance is adaptive to the duration of the sign:  $\pm 2$  frames for signs with duration of 1-5 frames,  $\pm 3$  frames for signs with 6-10 frames,  $\pm 4$  frames for signs with 11-23 frames, and  $\pm 5$  frames for longer signs. This adaptive metric ensures that both boundaries are accurately captured within a reasonable margin. Allowing for some divergence takes account of the fact that neighboring frames may be very similar in appearance; there may also be a little variability in manual annotations for the same reason. Thus, this tolerance allows for reasonable evaluation.

For the CSLR task, we use our tolerance-based matching criterion, Boundary Tolerance, to determine whether we have achieved a successful segmentation. We consider a predicted segment to be aligned with a sign in the ground truth if both the start and end frames of the two signs are within the tolerance range. We then compute the top-1 recognition accuracy across all predicted sign segments that were successfully matched based on this criterion.

### 4.3. Evaluation Benchmark

We separately evaluate our model on 3 tasks: boundary detection, sign recognition from continuous signing, and handshape classification.

#### 4.3.1. Segmentation Comparison

We compare our segmentation results with Renz et al. (2021), although our experiments are carried out on the ASLLRP-S dataset, which is more complex and challenging than the PHOENIX14 dataset (Forster et al., 2014). Specifically, ASLLRP-S contains approximately 2,000 distinct sign classes and fewer than 20,000 annotated sign clips, whereas PHOENIX14 includes around 1,200 sign classes with approximately 23,000 annotated sign clips.

Despite the fact that our dataset is significantly more complex, we adopt the same training-testing split ratio used in prior work and directly compare segmentation performance. As shown in the comparison with other methods in Table 1, Schembri et al. (2013) report a segmentation performance of 61.12 (mF1B) on the BSL-1K test dataset. For the PHOENIX14 dataset, Renz et al. (2021) achieve an mF1B of 71.50.

In contrast, our model achieves substantially stronger segmentation performance on the more challenging ASLLRP-S dataset, with an mF1B of 79.40. Given that ASLLRP-S presents greater complexity than PHOENIX14 in terms of vocabulary

Method	Dataset	mF1B	mF1S
<i>RF<sub>GEO</sub></i>	BSLCORPUS	68.68	41.71
<i>I3D,TCN</i>	PHOENIX14	71.50	52.78
<i>Ours MHB</i>	ASLLRP-S	<b>79.40</b>	<b>58.26</b>

Table 1: Comparison with other methods. We compare to the geometric features proposed by Farag and Brock (2019) on BSLCORPUS and the I3D + MS-TCN method proposed by Renz et al. (2021) on Phoenix14. Our method significantly outperforms the others on the different metrics.

size and class distribution, these results demonstrate the robustness and effectiveness of our proposed segmentation approach.

#### 4.3.2. Sign Recognition

In the sign recognition task, we adopt a different setting than was used for the standalone boundary detection task, as described earlier. To the best of our knowledge, this is the first work to evaluate sign recognition on continuous sign language datasets using segments derived from predicted boundaries. Therefore, during evaluation we assess performance from two perspectives: boundary detection and sign recognition.

For boundary detection, we adopt our Boundary Tolerance matching criterion. Specifically, we report 2 metrics to evaluate segmentation performance: (i) the proportion of ground-truth signs that have predicted segments that match under our Boundary Tolerance criterion; and (ii) the proportion of predicted segments that match a ground-truth sign under our Boundary Tolerance criterion. This provides a clear assessment of segmentation accuracy, as shown in Table 2.

Segment type	Match	Total	Prop
Ground-truth signs	3783	6595	57.4%
Predicted sign segmtns	3783	6477	58.5%

Table 2: Segmentation accuracy evaluated based on our tolerance-based matching criterion. For each segment type, **Match** denotes the number of sign segments that have corresponding sign segments under our matching criterion, **Total** is the total number of segments in the ground-truth or predicted sign segments, and **Prop** is the ratio of matched to total (First row is the Recall; Second Row is the Prediction).

We evaluate the improvement in sign recognition based on the predicted segmentation boundaries by conducting experiments under varying thresholds  $k$  for the minimum number of training samples per sign class. Evaluation is restricted to predicted sign segments that can be aligned to ground-

truth sign segments under our Boundary Tolerance criterion. For varying values of  $k$ , we report the top-1 sign recognition accuracy in Table 3, which shows that top-1 recognition accuracy on boundary-matched segments increases consistently as the minimum sample threshold  $k$  rises, ranging from 80.2% to 83.3%. These results demonstrate that our predicted segmentation boundaries provide reliable alignment with ground truth, and that recognition performance is strong and improves with stricter data thresholds, confirming the robustness of our approach for downstream sign recognition tasks.

Threshold $k$	Top-1 Accuracy
6	80.23%
10	80.86%
15	81.67%
20	82.17%
30	<b>83.30%</b>

Table 3: Sign recognition results on the ASLLRP-S dataset under different thresholds  $k$ , where  $k$  denotes the minimum number of training samples required per class. Top-1 Accuracy is the recognition accuracy based on the Top-1 prediction for each segment.

#### 4.4. Ablation Study

We also conduct an ablation study; see Table 4. When both the handshape classification module and the multimodal fusion module are removed, leaving only the spatio-temporal convolutional segmentation model, we observe a small but significant performance drop. The model still outperforms existing methods, demonstrating the strength of our base segmentation framework.

Method	mF1B	mF1S
<i>I3D,TCN</i>	71.50	52.78
<i>MSTCN, HaMeR</i>	76.22	50.18
<i>Ours w/o handshape module</i>	77.29	56.98
<i>Ours with handshape module</i>	<b>79.40</b>	<b>58.26</b>

Table 4: Ablation study for the boundary detection task. We compare to the I3D + MS-TCN method proposed by Renz et al. (2021) and the MS-TCN + HaMeR method proposed by He et al. (2025). The 2nd row shows our method without the handshape module, and the 3rd presents the full model. Best performance is achieved when the handshape module is incorporated.

## 5. Conclusion

Here we address the task of boundary detection for CSLR and propose a novel pipeline for continuous sign language recognition. We construct a

merged ASL handshape classification dataset and evaluate our method on the ASLLRP-S dataset, measuring boundary detection performance and the accuracy of recognition of successfully segmented signs. The results demonstrate that our approach significantly outperforms existing baselines in boundary detection, as shown in table 4. The ablation studies and other experiments confirm that accurate boundary modeling contributes to the overall recognition performance.

## 6. Acknowledgments

We are grateful to the many people who have helped with the collection, linguistic annotation, and sharing of the ASL data upon which we have relied for this research. In particular, we are indebted to the many ASL signers who have contributed to our database; to Matt Huenerfauth and his team for data collection at RIT; to the many who have helped with linguistic annotations; and to our invaluable Web designer, Augustine Opoku. Thanks also to anonymous reviewers for their suggestions. This work was supported in part by NSF grants 2212302, 2212301, and 2212303, but any opinions, findings, and conclusions expressed here are those of the authors and do not necessarily reflect the views of the National Science Foundation.

## 7. Bibliographical References

- ASL-deafined. 2017. Deaf and Hard of Hearing Population in the United States (expanded overview). <https://blog.asldeafined.com/2025/04/deaf-and-hard-of-hearing-population-in-the-united-states-expanded-overview/>.
- Matyáš Boháček and Marek Hruží. 2022. [Sign pose-based transformer for word-level sign language recognition](#). In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 182–191.
- Diane Brentari. 2019. *Sign Language Phonology*. University of Chicago Press.
- Hannah Bull, Triantafyllos Afouras, Gül Varol, Samuel Albanie, Liliane Momeni, and Andrew Zisserman. 2021. [Aligning subtitles in sign language videos](#). In *Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11552–11561.
- Necati Cihan Camgöz, Simon Hadfield, Oscar Koller, and Richard Bowden. 2017. [Subunets: End-to-end hand shape and continuous sign language recognition](#). In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 3075–3084.
- Konstantinos M. Dafnis, Evgenia Chroni, Carol Neidle, and Dimitri Metaxas. 2022. [Isolated sign recognition using ASL datasets with consistent text-based gloss labeling and curriculum learning](#). In *Proceedings of the 7th International Workshop on Sign Language Translation and Avatar Technology: The Junction of the Visual and the Textual: Challenges and Perspectives*, pages 13–20, Marseille, France. European Language Resources Association.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hao-Shu Fang, Jiefeng Li, Hongyang Tang, Chao Xu, Haoyi Zhu, Yuliang Xiu, Yong-Lu Li, and Cewu Lu. 2023. [Alphapose: Whole-body regional multi-person pose estimation and tracking in real-time](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6):7157–7173.
- Iva Farag and Heike Brock. 2019. [Learning motion disfluencies for automatic sign language segmentation](#). In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7360–7364.
- Jens Forster, Christoph Schmidt, Oscar Koller, Martin Bellgardt, and Hermann Ney. 2014. [Extensions of the sign language recognition and translation corpus RWTH-PHOENIX-weather](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 1911–1916, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Matilde Gonzalez and Christophe Collet. 2011. [Sign segmentation using dynamics and hand configuration for semi-automatic annotation of sign language corpora](#). In *International Gesture Workshop*, pages 204–215. Springer.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd International Conference on Machine Learning*, pages 369–376.

- Leming Guo, Wanli Xue, Qing Guo, Bo Liu, Kaihua Zhang, Tiantian Yuan, and Shengyong Chen. 2023. Distilling cross-temporal contexts for continuous sign language recognition. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 10771–10780.
- Leming Guo, Wanli Xue, Bo Liu, Kaihua Zhang, Tiantian Yuan, and Dimitris Metaxas. 2024. Gloss prior guided visual feature learning for continuous sign language recognition. *IEEE Transactions on Image Processing*, 33:3486–3495.
- Haoyu Han, Mengdi Zhang, Min Hou, Fuzheng Zhang, Zhongyuan Wang, Enhong Chen, Hongwei Wang, Jianhui Ma, and Qi Liu. 2020. Stgcn: A spatial-temporal aware graph learning method for poi recommendation. In *2020 IEEE International Conference on Data Mining (ICDM)*, pages 1052–1057.
- Low Jian He, Harry Walsh, Ozge Mercanoglu Sincan, and Richard Bowden. 2025. Hands-on: Segmenting individual signs from continuous sequences. In *2025 IEEE 19th International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–5. IEEE.
- Jiani Huang, Wengang Zhou, Houqiang Zhang, and Weiping Li. 2018. Video-based sign language recognition without temporal segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, pages 2257–2264. Article No. 275.
- Zifan Jiang, Gerard Sant, Amit Moryossef, Mathias Müller, Rico Sennrich, and Sarah Ebling. 2024. SignCLIP: Connecting text and sign language by contrastive learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9171–9193, Miami, Florida, USA. Association for Computational Linguistics.
- Can Jin, Ying Li, Mingyu Zhao, Shiyu Zhao, Zhen-ting Wang, Xiaoxiao He, Ligong Han, Tong Che, and Dimitris N. Metaxas. 2025. Lor-VP: Low-rank visual prompting for efficient vision model adaptation. In *The Thirteenth International Conference on Learning Representations*, pages 9359–9376.
- Tonni Das Jui, Gissella Bejarano, and Pablo Rivas. 2022. A machine learning-based segmentation approach for measuring similarity between sign languages. In *Proceedings of the LREC2022 10th Workshop on the Representation and Processing of Sign Languages: Multilingual Sign Language Resources*, pages 94–101.
- Dongxu Li, Cristian Rodriguez, Xin Yu, and Hongdong Li. 2020. Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In *Proceedings of the IEEE/CVF winter conference on Applications of Computer Vision*, pages 1459–1469.
- Natalia Lusin, Terri Peterson, Christine Sulewski, and Rizwana Zafer. 2023. Enrollments in languages other than English in US institutions of higher education, fall 2021 - Modern Language Association of America.
- Amit Moryossef, Zifan Jiang, Mathias Müller, Sarah Ebling, and Yoav Goldberg. 2023. Linguistically motivated sign language segmentation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12703–12724, Singapore. Association for Computational Linguistics.
- Boston University (BU) National Center for Sign Language and Gesture Resources. 2012. NC-SLGR Handshapes. Videos available from <https://www.bu.edu/asllrp/cslgr/pages/nicslgr-handshapes.html>.
- Carol Neidle. 2023. Challenges for linguistically-driven computer-based sign recognition from continuous signing for American Sign Language. *arXiv preprint arXiv:2311.00762*.
- Carol Neidle and Carey Ballard. 2022a. Revised Gloss Labels for Signs from the WLASL Dataset - Boston University. <http://www.bu.edu/asllrp/wlasl-alt-glosses.pdf>.
- Carol Neidle and Carey Ballard. 2022b. Why alternative gloss labels will increase the value of the WLASL dataset - ASLLRP Project Report no. 21, Boston University. Project reports available from <https://www.bu.edu/asllrp/reports.html>.
- Carol Neidle and Dimitris Metaxas. 2024a. American Sign Language Linguistic Research Project (ASLLRP) Sign Bank: Boston University American Sign Language Lexicon Video Dataset (ASLLVD) version 8 - June. Videos and annotations available from <https://dai.cs.rutgers.edu/dai/s/signbank>.
- Carol Neidle and Dimitris Metaxas. 2024b. American Sign Language Linguistic Research Project (ASLLRP) Sign Bank: DawnSignPress (DSP) ASL Sign Collection, version 6. Videos provided by DawnSignPress.
- Carol Neidle and Dimitris Metaxas. 2024c. American Sign Language Linguistic Research Project (ASLLRP) Sign Bank: Rochester Institute of Technology (RIT) collection, version 6. Videos

- and annotations available from <https://dai.cs.rutgers.edu/dai/s/signbank>.
- Carol Neidle and Dimitris Metaxas. 2025a. American Sign Language Linguistic Research Project (ASLLRP) Sign Bank: ASLLRP Continuous Signing Corpora - version 3 - June 2025. Videos and annotations available from <https://dai.cs.rutgers.edu/dai/s/signbank>.
- Carol Neidle and Dimitris Metaxas. 2025b. American Sign Language Linguistic Research Project (ASLLRP) Sign Bank: DawnSignPress (DSP) Sentences Collection, version 4. Videos provided by DawnSignPress.
- Carol Neidle and Augustine Opoku. 2022. [Addendum to: Documentation for Download of ASLLRP Sign Bank Citation-form Sign Datasets - ASLLRP Project Report no. 20, Boston University](#). Project reports available from <https://www.bu.edu/asllrp/reports.html>.
- Carol Neidle, Augustine Opoku, Carey Ballard, Konstantinos M Dafnis, Evgenia Chroni, and Dimitris Metaxas. 2022a. [Resources for computer-based sign recognition from video, and the criticality of consistency of gloss labeling across multiple large ASL video corpora](#). In *Proceedings of the LREC2022 10th Workshop on the Representation and Processing of Sign Languages: Multilingual Sign Language Resources, Marseille, France, 25 June 2022*, pages 165–172.
- Carol Neidle, Augustine Opoku, and Dimitris Metaxas. 2022b. ASL Video Corpora & a Sign Bank: Resources available through the American Sign Language Linguistic Research Project (ASLLRP). *arXiv preprint arXiv:2201.07899*.
- Carol Neidle, Ashwin Thangali, and Stan Sclaroff. 2012. [Challenges in development of the American Sign Language Lexicon Video Dataset \(ASLLVD\) corpus](#). In *Proceedings of the LREC2012 5th Workshop on the Representation and Processing of Sign Languages: Interactions between Corpus and Lexicon*, pages 143–150, Istanbul, Turkey. European Language Resources Association (ELRA).
- Lionel Pigou, Sander Dieleman, Pieter-Jan Kindermans, and Benjamin Schrauwen. 2015. Sign language recognition using convolutional neural networks. In *Computer Vision – ECCV 2014 Workshops*, pages 572–578. Springer.
- Razieh Rastgoo, Kouros Kiani, and Sergio Escalera. 2024. [A transformer model for boundary detection in continuous sign language](#). *Multimedia Tools and Applications*, 83(42):89931–89948.
- Razieh Rastgoo, Kouros Kiani, and Sergio Escalera. 2025. A non-anatomical graph structure for boundary detection in continuous sign language. *Scientific Reports*, 15(1):25683.
- Katrin Renz, Nicolaj C Stache, Samuel Albanie, and Gül Varol. 2021. [Sign language segmentation with temporal convolutional networks](#). In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2135–2139.
- Adam Schembri, Jordan Fenlon, Ramas Rentelis, Sally Reynolds, and Kearsy Cormier. 2013. [Building the British Sign Language corpus](#). *Language Documentation & Conservation*, 7:136–154.
- Thad Starner and Alex Pentland. 1998. [Real-time American Sign Language recognition using desk and wearable computer-based video](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12):1371–1375.
- 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 *NIPS’17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6000–6010.
- Christian Vogler and Dimitris Metaxas. 1999. [Parallel Hidden Markov Models for American Sign Language Recognition](#). In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 116–122.
- Christian Vogler and Dimitris Metaxas. 2001. [A framework for recognizing the simultaneous aspects of American Sign Language](#). *Computer Vision and Image Understanding*, 81(3):358–384.
- Polina Yanovich, Carol Neidle, and Dimitris Metaxas. 2016. Detection of major ASL sign types in continuous signing for ASL recognition. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3067–3073.
- Huaiwen Zhang, Zihang Guo, Yang Yang, Xin Liu, and De Hu. 2023. [C2st: Cross-modal contextualized sequence transduction for continuous sign language recognition](#). In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21053–21062.
- Sen Zhang, Xiaoxiao He, Di Liu, Zhaoyang Xia, Mingyu Zhao, Chaowei Tan, Vivian Li, Bo Liu, Dimitris N. Metaxas, and Mubbasir Kapadia. 2026. [Large sign language models: Toward 3D American Sign Language translation](#). In *Proceedings of the IEEE/CVF Winter Conference on*

*Applications of Computer Vision (WACV)*, pages 3109–3119.

Xuan Zhang and Kevin Duh. 2023. [Handshape-aware sign language recognition: Extended datasets and exploration of handshape-inclusive methods](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2993–3002.

Yan Zhang, Wanli Xue, Yuxi Zhou, Tiantian Yuan, and Shengyong Chen. 2025. [Core: multi-link graph attention network with inter-regional collaboration for continuous sign language recognition](#). *Pattern Recognition*, 167:111716.

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.

Yang Zhou, Zhaoyang Xia, Yuxiao Chen, Carol Neidle, and Dimitris N. Metaxas. 2024. [A multi-modal spatio-temporal GCN model with enhancements for isolated sign recognition](#). In *Proceedings of the LREC-COLING 2024 11th Workshop on the Representation and Processing of Sign Languages: Evaluation of Sign Language Resources*, pages 408–419, Torino, Italia. ELRA and ICCL.