

# Pose-Based Sign Language Spotting via an End-to-End Encoder Architecture

Samuel Ebimobowei Johnny<sup>1</sup>    Blessed Guda<sup>1,2</sup>

Emmanuel Enejo Aaron<sup>1</sup>    Assane Gueye<sup>1,2</sup>

{sjohnny, blessedg, eaaron, assaneg}@andrew.cmu.edu

<sup>1</sup>Carnegie Mellon University Africa, Kigali, Rwanda

<sup>2</sup>Carnegie Mellon University, Pittsburgh, USA

## Abstract

Automatic Sign Language Recognition (ASLR) has emerged as a vital field for bridging the gap between deaf and hearing communities. However, the problem of sign-to-sign retrieval or detecting a specific sign within a sequence of continuous signs remains largely unexplored. We define this novel task as Sign Language Spotting. In this paper, we present a first step toward sign language retrieval by addressing the challenge of detecting the presence or absence of a query sign video within a sentence-level gloss or sign video. Unlike conventional approaches that rely on intermediate gloss recognition or text-based matching, we propose an end-to-end model that directly operates on pose keypoints extracted from sign videos. Our architecture employs an encoder-only backbone with a binary classification head to determine whether the query sign appears within the target sequence. By focusing on pose representations instead of raw RGB frames, our method significantly reduces computational cost and mitigates visual noise. We evaluate our approach on the Word Presence Prediction dataset from the WSLP 2025 shared task, achieving 61.88% accuracy and 60.00% F1-score. These results demonstrate the effectiveness of our pose-based framework for Sign Language Spotting, establishing a strong foundation for future research in automatic sign language retrieval and verification. Code is available at [this repository](#).

## 1 Introduction

Sign language, which globally consists of more than 300 different sign languages (United Nations, 2023), was developed to address the need for effective communication for the deaf and hearing-impaired population (Tunga et al., 2021). Each sign language comprises a complex combination of hand gestures, facial expressions, and body movements that collectively encode the semantics and grammatical structures of spoken languages (Tang et al., 2025; Rastgoo et al., 2024). However, there

is still a challenge and a communication gap between the deaf and hearing community (Das et al., 2024), (Venugopalan and Reghunadhan, 2021). Previous works have focused on sign language translation(SLT) where researchers have attempted to translate sign language either as RGB or poses to either text(that is word word-level semantically meaningful) (Yin and Read, 2020; Kan et al., 2022) or glosses (Zhou et al., 2023; Low, 2025).

Sign language recognition(SLR) could be isolated and continuous SLR. Isolated sign language (ISLR) (Kumari and Anand, 2024; Baihan et al., 2024; Ren et al., 2025) translation involves word-level focuses on recognizing individual signs in isolation, treating each sign as an independent classification problem. In contrast, Continuous Sign Language Recognition (CSLR) (Wang et al., 2025; Jian He et al., 2025; Zheng et al., 2023; Low, 2025) involves sentence-level SL, which addresses a more challenging task of translating continuous signing sequences into semantically correct sentences or gloss annotations, requiring models to handle temporal dependencies, co-articulation effects, and variable-length sequences.

While significant progress has been made in recognition and translation, the ability to search, retrieve, or verify specific signs within continuous signing videos remains underexplored. This capability- known as sign spotting - is critical for applications such as SL retrieval, dictionary lookup, and educational tools. This requires robust sign spotting capabilities, that is, the ability to locate and identify specific signs within continuous signing videos. Traditional approaches to this problem have relied on text-based intermediate representations.

For word spotting for CSLR, researchers have attempted to spot words using Large Language Models (LLMs). (Walsh et al., 2023) proposed using LLMs such as BERT and Word2Vec to leverage alignment to improve isolated signs from continu-

ous signs. Their approach solves text gloss mapping using LLMs; their model provides an effective method, which was evaluated on MeieneDGS (Konrad et al., 2020) and BOBSL(Albanie et al., 2021).

Recent work on sign spotting addresses the challenge of SLT by decomposing it into modular stages. Spotter+GPT (Jian He et al., 2025) proposes an approach to eliminate the need for SLT-specific end-to-end training, significantly reducing computational costs. Their approach extracts I3D motion and ResNeXt-101 handshape features, matches them to a sign dictionary using DTW and cosine similarity, and passes spotted signs, which are the top-k glosses, to GPT for sentence generation. While Spotter+GPT demonstrates the effectiveness of modular SLT, our work addresses a fundamentally different task: word presence verification.

In this paper, we introduce a novel end-to-end video-to-video sign spotting framework that eliminates the need for textual or gloss-based intermediates. Given a query sign video and a sentence-level sign video, our model determines whether the query sign is present within the sentence. We adopt an encoder-only architecture with a binary classification head, operating directly on pose keypoints rather than RGB frames. This design reduces computational complexity and suppresses visual noise while maintaining discriminative spatial-temporal information. We evaluate our approach on the Word Presence Prediction dataset from the WSLP 2025 shared task<sup>1</sup>. To the best of our knowledge, this work represents the first study to address sign language spotting purely through video-to-video matching, establishing a foundation for future research in automatic sign language retrieval, verification, and search.

## 2 Methodology

We propose a video-to-video sign spotting architecture that jointly models visual–semantic alignment and binary word presence prediction. The framework learns robust cross-modal representations that generalize across signers and sentence contexts. Our approach consists of three main components: pose extraction, feature encoding, and presence prediction, as illustrated in Figure 1.

<sup>1</sup><https://exploration-lab.github.io/WSLP/task/>

### 2.1 Pose Extraction

We use MediaPipe (Lugaresi et al., 2019) to convert RGB video sequences to pose-based representations, allowing for a more generalized, efficient, and resilient architecture. For each frame, MediaPipe estimates the pose keypoints of the signer in the video. Following (Johnny et al., 2025), we extract holistic pose features containing 42 hand keypoints (21 per hand), 8 body keypoints, and 19 facial landmarks. As suggested by (Johnny et al., 2025), we used only the hand and body features in this study.

### 2.2 Problem formulation

Given a sentence sequence  $\mathbf{X}_s \in \mathbb{R}^{T_s \times F}$  and a query sequence  $\mathbf{X}_q \in \mathbb{R}^{T_q \times F}$ , where  $T_s$  and  $T_q$  denote the temporal lengths of the sentence and query sequences respectively, and  $F$  represents the dimensionality of the pose features, the objective is to determine whether the query sign appears within the sentence.

Let  $f(\mathbf{X}_s, \mathbf{X}_q; \theta)$  be a parameterized model, where  $\theta$  denotes the set of learnable parameters. The model outputs a probability score  $\hat{y} = f(\mathbf{X}_s, \mathbf{X}_q; \theta)$ , representing the likelihood that the query sign occurs in the given sentence. The binary prediction is made as: The training objective is to optimize the model parameters  $\theta$  by minimizing a loss function  $\mathcal{L}(\theta)$  over the training data:

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\theta).$$

### 2.3 Pose CNN Encoder

Each pose frame is represented as a vector  $\mathbf{x}_t \in \mathbb{R}^{100}$ , corresponding to 50 keypoints with 2D coordinates  $(x, y)$ . To preserve the spatial topology of the human skeleton, each vector is reshaped into a 2D array of size  $\mathbb{R}^{50 \times 2}$ , considering only the hand and body keypoints as described in Section 2.1.

A 2D CNN is applied independently to each frame to extract local spatial dependencies among keypoints. Specifically, each pose frame passes through three Conv2D blocks, each consisting of a convolutional layer, Batch Normalization, and ReLU activation. These blocks progressively capture hierarchical geometric patterns while maintaining spatial coherence across keypoints.

Following the convolutional layers, an **adaptive average pooling** layer reduces the spatial dimensions to a fixed-size representation, which is then

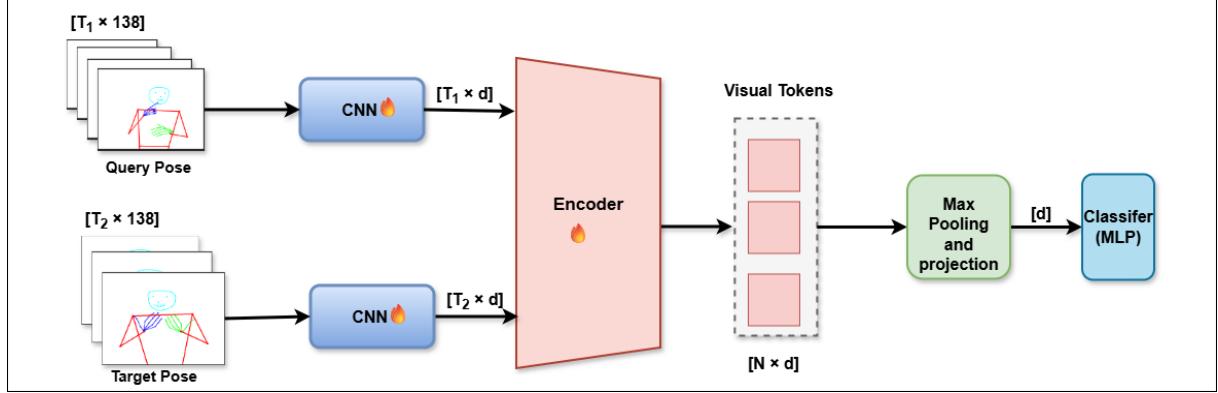


Figure 1: Architecture overview. Pose sequences are encoded using 2D CNNs and then processed by a Transformer encoder, which produces visual tokens. The [CLS] token is max-pooled to predict query presence using binary cross-entropy loss ( $\mathcal{L}_{BCE}$ ).

linearly projected to a feature vector of dimension  $d = 128$ .

This process yields a sequence of per-frame embeddings:

$$\mathbf{H} = \{\mathbf{h}_t\}_{t=1}^T \in \mathbb{R}^{T \times 128}, \quad (1)$$

where  $T$  denotes the total frames in the video. Each embedding  $\mathbf{h}_t$  encodes the spatial structure of the signer’s body and hand poses at time step  $t$ . The resulting feature sequence is passed to the transformer encoder for temporal modeling.

## 2.4 Visual Transformer Encoder model

To enable temporal dependencies and cross-sequence interactions between the query and sentence embeddings, we adopt a BERT-style sequence modeling approach. Specifically, a [CLS] token is prepended for global sequence-level classification, while a [SEP] token is inserted to explicitly separate the *query pose tokens* from the *candidate pose tokens*. This design enables the model to attend across the boundary between the two sequences, allowing direct interaction between corresponding temporal segments.

Learnable positional encodings and token-type embeddings are incorporated to preserve temporal order and to distinguish between query and candidate sequences. The transformer encoder then processes the concatenated sequence using multi-head self-attention, where the **attention scores between query and candidate pose tokens** serve as a key mechanism for measuring their semantic and spatial correspondence. These cross-sequence attention patterns help the model identify whether visual and structural similarities exist between the

query sign and any segment of the candidate video, thereby assisting the sign spotting task.

### 2.4.1 Classification Loss

Since the expected outcome is binary(present or absent), our model employs the binary cross-entropy loss(BCE) to penalize incorrect and overconfident predictions. We extract the [CLS] token representation via max pooling and project it to an MLP classifier to generate the corresponding logits  $\hat{y} \in \mathbb{R}$ . The final prediction is obtained by applying the sigmoid( $\sigma$ ) to the logits. The binary cross-entropy loss is computed as:

$$\mathcal{L}_{BCE} = -\frac{1}{B} \sum_{i=1}^B [y_i \log p_i + (1 - y_i) \log(1 - p_i)] \quad (2)$$

where  $p = \sigma(\hat{y}_i)$ ,  $\sigma$  is the sigmoid function,  $y_i \in \{0, 1\}$  is the ground truth label, and  $B$  is the batch size.

## 3 Experiments

### 3.1 Dataset and Evaluation Metrics

For this experiment, we evaluate our model with *Word Presence Dataset*<sup>2</sup>, an ASL sign spotting dataset designed to determine if a query sign appears within a sentence sequence.

The dataset comprises 25,432 sentence-query pairs constructed from 7,857 unique sentence sequences and 1,410 unique query sequences. The dataset is balanced, with equal distribution of positive (query present) and negative (query absent) samples. We employ an 80:20 train-validation split.

<sup>2</sup><https://huggingface.co/datasets/Exploration-Lab/WSLP-AAACL-2025>

Table 1: Performance comparison on Word Presence Prediction (Test Set).

Method	Acc.	F1	Prec.	Rec.
Ours (1D CNN)	60.95	<b>59.62</b>	62.70	61.01
<b>Ours (2D CNN)</b>	<b>61.66</b>	58.42	<b>67.16</b>	<b>61.74</b>

The test set contains 1,266 unique sentence sequences and 555 unique query sequences, ensuring minimal overlap with the training distribution.

During Evaluation, we use standard classification evaluation metrics, i.e., Accuracy, Precision, Recall, and F1-score.

## 4 Implementation Details

We train the model end-to-end with an initial learning rate of 0.0005 since commonly used values (e.g., 0.001 or 0.01) resulted in suboptimal convergence, using the AdamW optimizer and a temperature of 0.07 for contrastive losses. A dropout of 0.02 was applied to prevent overfitting. Training is carried out for 50 epochs with a patience of 5 if no future improvements. This was done using a single NVIDIA L40S GPU.

To ensure our model focuses on important features, we skipped all early and late frames with no finger movement. During training, we applied different data augmentation techniques such as sequence masking, scaling, jittering, and Gaussian noise to ensure robustness.

### 4.1 Evaluation on Test Set

Table 1 presents our results on the test set. Given that this is a novel task introduced in the WSLP 2025 shared task, with no prior work to the best of our knowledge, hence no baseline to compare against. Our 2D-CNN approach achieves 61.66% accuracy, outperforming linear(1D-CNN) projection. Notably, 2D-CNN significantly improves precision, indicating fewer false positives, though F1 slightly decreases due to the precision-recall trade-off.

### 4.2 Ablation study and analysis

To evaluate the robustness of our model, we conduct some ablations using different training choices with a concentration on Accuracy and F1 Scores.

#### 4.2.1 Effect of different loss function

As shown in Table 2 demonstrate that using only contrastive loss underperforms when compared to

Table 2: Ablation study on validation set(val set).

Configuration	Acc.	F1	Prec.	Rec.
<i>Loss Functions</i>				
<b>BCE only (ours)</b>	<b>63.04</b>	<b>70.36</b>	59.19	86.71
Contrast only	57.20	69.27	54.39	<b>95.38</b>
BCE + Contrast	61.39	64.13	<b>60.49</b>	68.25
<i>Pose Encoding</i>				
1D Conv	53.65	<b>67.58</b>	52.29	95.53
<b>2D Conv (ours)</b>	<b>61.39</b>	64.13	<b>60.49</b>	<b>68.25</b>

using BCE, indicating that contrastive supervision is not satisfactory enough for this task. While combining both losses with contrastive weight  $\lambda = 0.5$  achieves 61.39% accuracy, the result is still below BCE-only performance. The contrastive objective may interfere with classification if the weight is not carefully tuned; using either higher or lower weights results in lower performance, and the embedding space learned through mean pooling may be less discriminative than the [CLS] token representation for this verification task.

#### 4.2.2 Effect of Pose Encoding

Table 2 demonstrates that 2D-CNN outperforms other methods in encoding postures. 1D-CNN captures temporal patterns but treats keypoints as a sequence without using their geometric correlations. In contrast, 2D-CNN preserves spatial structure by reshaping each frame as an  $n \times 2$  grid, where  $n$  is the number of keypoints, allowing the network to learn spatial patterns such as hand configurations and body postures.

## 5 Conclusion

In the work, we present the first video-to-video word presence verification in sign language, where both the sentence and query are video sequences. Our approach proposes using pose sequence in, combining 2D CNN encoding with a Transformer temporal model, achieving 61.66% accuracy on the *word presence dataset*.

To the best of our knowledge at the time of this research, no prior work has been done in video-to-video sentence-to-query word spotting. Ablation studies and analysis show that 2D spatial encoding of poses and BCE loss are critical design choices. Our work establishes a strong baseline for this task and demonstrates the effectiveness of pose-based representations for SL understanding.

## References

Samuel Albanie, Gü̈l Varol, Liliane Momeni, Hannah Bull, Triantafyllos Afouras, Himel Chowdhury, Neil Fox, Bencie Woll, Rob Cooper, Andrew McParland, and Andrew Zisserman. 2021. BOBSL: Bbc-oxford british sign language dataset. <https://www.robots.ox.ac.uk/~vgg/data/bobs1>.

Abdullah Baihan, Ahmed I. Alutaibi, Mohammed Al-shehri, and Sunil Kumar Sharma. 2024. Sign language recognition using modified deep learning network and hybrid optimization: a hybrid optimizer (ho) based optimized cnnsa-lstm approach. *Scientific Reports*, 14:26111.

Subhadeep Das, Subrata Kr. Biswas, and Bidyut Purkayastha. 2024. Occlusion robust sign language recognition system for indian sign language using cnn and pose features. *Multimedia Tools and Applications*, 83(36):84141–84160.

Low Jian He, Ozge Mercanoglu Sincan, and Richard Bowden. 2025. Sign spotting disambiguation using large language models. In *Adjunct Proceedings of the 25th ACM International Conference on Intelligent Virtual Agents*, IVA Adjunct '25, New York, NY, USA. Association for Computing Machinery.

Samuel Ebimobowei Johnny, Blessed Guda, Andrew Blayama Stephen, and Assane Gueye. 2025. Autosign: Direct pose-to-text translation for continuous sign language recognition. *arXiv preprint arXiv:2507.19840*.

Jichao Kan, Kun Hu, Markus Hagenbuchner, Ah Chung Tsoi, Mohammed Bennamoun, and Zhiyong Wang. 2022. Sign language translation with hierarchical spatio-temporal graph neural network. In *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2131–2140.

Reiner Konrad, Thomas Hanke, Gabriele Langer, Dolly Blanck, Julian Bleicken, Ilona Hofmann, Olga Jeziorski, Lutz König, Susanne König, Rie Nishio, Anja Regen, Uta Salden, Sven Wagner, Satu Worseck, Oliver Böse, Elena Jahn, and Marc Schulder. 2020. Meine dgs – annotiert. öffentliches korpus der deutschen gebärdensprache, 3. release / my dgs – annotated. public corpus of german sign language, 3rd release.

Diksha Kumari and Radhey Shyam Anand. 2024. Isolated video-based sign language recognition using a hybrid cnn-lstm framework based on attention mechanism. *Electronics*, 13(7).

Jian He Low. 2025. Sage: Segment-aware gloss-free encoding for token-efficient sign language translation. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV 2025)*. Institute of Electrical and Electronics Engineers (IEEE).

Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, and 1 others. 2019. Mediapipe: A framework for building perception pipelines. In *arXiv preprint arXiv:1906.08172*.

Rahimeh Rastgoo, Kourosh Kiani, and Sergio Escalera. 2024. Word separation in continuous sign language using isolated signs and post-processing. *Expert Systems with Applications*, 249:123695.

Yazhou Ren, Hongkai Li, Yuhao Li, Jingyu Pu, Xiaorong Pu, Siyuan Jing, Peng Jin, and Lifang He. 2025. Multi-modal isolated sign language recognition based on self-paced learning. *Expert Systems with Applications*, 291:128340.

Siliang Tang, Fan Xue, Jing Wu, Shanghang Wang, and Richang Hong. 2025. Gloss-driven conditional diffusion models for sign language production. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 21(4):105:1–105:17.

Abhishek Tunga, Sravanthi V. Nuthalapati, and Juan Wachs. 2021. Pose-based sign language recognition using gcn and bert. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision Workshops (WACVW)*, pages 31–40. IEEE.

United Nations. 2023. International day of sign languages. <https://www.un.org/en/observances/sign-languages-day>. According to the United Nations, there are more than 300 different sign languages used around the world.

Adithya Venugopalan and Rajesh Reghunadhan. 2021. Applying deep neural networks for the automatic recognition of sign language words: A communication aid to deaf agriculturists. *Expert Syst. Appl.*, 185(C).

Harry Walsh, Ozge Mercanoglu Sincan, Ben Saunders, and Richard Bowden. 2023. Gloss alignment using word embeddings. *CoRR*, abs/2308.04248.

Zhen Wang, Dongyuan Li, Renhe Jiang, and Manabu Okumura. 2025. Continuous sign language recognition with multi-scale spatial-temporal feature enhancement. *IEEE Access*, 13:5491–5506.

Kayo Yin and Jesse Read. 2020. Better sign language translation with STMC-transformer. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5975–5989, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Jiabin Zheng, Yile Wang, Cheng Tan, Siyuan Li, Ge Wang, Jun Xia, Yidong Chen, and Stan Z. Li. 2023. CVT-SLR: Contrastive Visual-Textual Transformation for Sign Language Recognition with Variational Alignment. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23141–23150, Los Alamitos, CA, USA. IEEE Computer Society.

Benjia Zhou, Zhigang Chen, Albert Clapés, Jun Wan, Yanyan Liang, Sergio Escalera, Zhen Lei, and Du Zhang. 2023. Gloss-free sign language translation: Improving from visual-language pretraining. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20871–20881.