

KD-MSLRT: Lightweight Sign Language Recognition Model Based on MediaPipe and 3D to 1D Knowledge Distillation

Yulong Li^{1*}, Bolin Ren^{1*}, Ke Hu², Changyuan Liu¹, Zhengyong Jiang¹, Kang Dang^{1†}, Jionglong Su^{1†}

¹School of AI and Advanced Computing, XJTLU Entrepreneur College (Taicang), Xi'an Jiaotong-Liverpool University, Suzhou, 215123, China

²School of Humanities and Social Sciences, Xi'an Jiaotong-Liverpool University, Suzhou, 215123, China
{Kang.Dang, Jionglong.Su}@xjtlu.edu.cn

Abstract

Artificial intelligence has achieved notable results in sign language recognition and translation. However, relatively few efforts have been made to significantly improve the quality of life for the 72 million hearing-impaired people worldwide. Sign language translation models, relying on video inputs, involves with large parameter sizes, making it time-consuming and computationally intensive to be deployed. This directly contributes to the scarcity of human-centered technology in this field. Additionally, the lack of datasets in sign language translation hampers research progress in this area. To address these, we first propose a cross-modal multi-knowledge distillation technique from 3D to 1D and a novel end-to-end pre-training text correction framework. Compared to other pre-trained models, our framework achieves significant advancements in correcting text output errors. Our model achieves a decrease in Word Error Rate (WER) of at least 1.4% on PHOENIX14 and PHOENIX14T datasets compared to the state-of-the-art CorrNet. Additionally, the TensorFlow Lite (TFLite) quantized model size is reduced to 12.93 MB, making it the smallest, fastest, and most accurate model to date. We have also collected and released extensive Chinese sign language datasets, and developed a specialized training vocabulary. To address the lack of research on data augmentation for landmark data, we have designed comparative experiments on various augmentation methods. Moreover, we performed a simulated deployment and prediction of our model on Intel platform CPUs and assessed the feasibility of deploying the model on other platforms.

Introduction

72 million hearing-impaired individuals worldwide rely on sign language as their primary mode of communication, as estimated by the World Federation of the Deaf (Secretariat 2009). However, the lack of widespread sign language understanding creates significant barriers to their full participation in society. While the rapid advancements in Natural Language Processing (NLP) and Computer Vision (CV) have made video-based sign language recognition (SLR) more feasible (Koller, Zargaran, and Ney 2017), existing sign language recognition (SLR) models often fall short in

*These authors contributed equally.

†Corresponding authors.

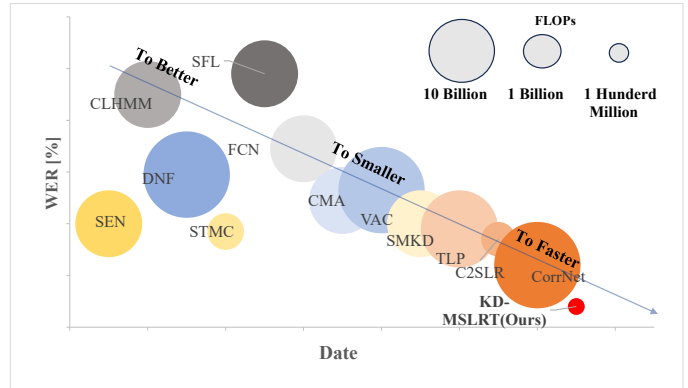


Figure 1: Publication time versus model performance. Larger circles indicate higher FLOPs, reflecting greater computational complexity.

addressing the practical needs of the hearing-impaired community.

Current SLR approaches, primarily based on complex video analysis, face three major limitations. First, these models are computationally expensive (Camgöz et al. 2016), hindering their deployment on resource-constrained devices like mobile phones, crucial for real-time communication accessibility. Second, despite significant progress in American Sign Language (ASL) recognition, research on other sign languages, particularly Chinese Sign Language (CSL), remains limited due to a lack of publicly available datasets (Selvaraj et al. 2021). Third, even state-of-the-art models are also prone to output errors, requiring additional mechanisms for grammatical and contextual correction to ensure accurate and understandable translations.

Our research addresses these challenges by introducing KD-MSLRT, a novel **K**nowledge **D**istillation-Based **M**ediaPipe **S**ign **L**anguage **R**ecognition model with **T**ext **C**orrection. Our approach prioritizes both accuracy and efficiency, aiming for practical deployment in real-world scenarios. Instead of relying on computationally intensive video processing (Cihan Camgoz et al. 2017), KD-MSLRT leverages MediaPipe (Grishchenko and Bazarevsky 2020), a powerful framework for real-time human pose estimation. By extracting skeletal landmarks from sign language videos, we simplify the input data from 3D to 1D, significantly reducing computational burden without compromising accuracy (Cao et al. 2017a). This allows for efficient processing

on even resource-constrained devices (Chen et al. 2022; Hu et al. 2025), making real-time sign language translation accessible to a wider audience. Figure 1 shows the publication time and performance of various models. Compared to other recently released baseline models, KD-MSLRT leads in both Word Error Rate (WER) and FLOPs. This can be attributed to three key innovations to enhance KD-MSLRT’s performance:

- **3D to 1D Knowledge Distillation.** This technique enables our lightweight model to learn from a more complex and accurate video-based model (Cheng et al. 2020a), achieving comparable performance with significantly reduced resource requirements.
- **Novel Temporal and Spatial Data Augmentation Algorithms.** Specifically designed for landmark-based sign language data, these algorithms significantly improve the model’s ability to generalize and recognize diverse signing styles.
- **Specialized Text Correction Network.** A novel text correction model trained on a self-supervised approach ensures grammatically correct and contextually relevant translations, addressing the common issue of errors in language model outputs.

Recognizing the scarcity of CSL resources, we also introduce a new large-scale dataset of 8,976 samples, focusing on long and complex sentences commonly found in news broadcasts. This dataset will be a valuable resource for researchers and developers, fostering further advancements in CSL recognition.

Our contributions are as follows:

- We developed KD-MSLRT, a highly accurate and efficient sign language recognition model that achieves a decrease in WER of at least 1.4% on PHOENIX14 and PHOENIX14T compared to state-of-the-art (SOTA) models. With a TensorFlow Lite (TFLite) quantized model size of just 12.93 MB, it can be deployed by enterprises at minimal cost.
- To the best of our knowledge, we are the first to propose temporal and spatial data augmentation algorithms specifically for landmark-format sign language data, which significantly enhances the effectiveness of landmark-based sign language recognition models.
- We collect and release the largest Chinese sign language recognition dataset in the CSLR field, containing 8,976 Chinese sign language samples with corresponding labels of complete news sentences.
- We performed a simulated deployment and prediction of our model on Intel platform CPUs and assessed the feasibility of deploying the model on other platforms.

Related Work

Continuous Sign Language Recognition. Continuous Sign Language Recognition (CSLR) is the recognition of consecutive sign language videos as an annotated sequence of sign languages after manual annotation (Sincan and Kelles 2020; Yin and Read 2020). Early methods in CSLR (Fels and Hinton 1993; Graves et al. 2006) always use manually

crafted features or systems based on Hidden Markov Models (Lewis et al. 2019) to perform temporal modeling and gradually translate sentences. HMM-based systems first use feature extractors to capture visual features and then use HMMs to perform long-term temporal modeling (Li et al. 2020a). With the introduction of Connectionist Temporal Classification (CTC) loss (Graves et al. 2006), significant improvements were achieved in the field of speech recognition. It can also be applied in CSLR (Li et al. 2020c), i.e., by sequentially aligning the target sentence with the input frames. The basis for using CTC loss in CSLR lies in extracting semantic information from sign language videos with long-term temporal dependencies. This requires the feature extractor to have a subtle structural design and to ensure that the spatial information is accurately extracted from the spatial information to the spatial features that contain the semantic information (Wang et al. 2022; Zhao et al. 2024). Therefore, Correlation Network is proposed (Albanie et al. 2020), which introduces Correlation operation as a spatial feature extractor based on a ResNet network, and employs 1D CNN and bidirectional LSTM (BiLSTM) for short-term and long-term temporal modeling, respectively, and employs classifiers to predict sentences.

With the proposal of models such as MediaPipe for extracting pose keypoints, C2SLR (Zuo and Mak 2022) uses pre-extracted pose keypoints as supervision to guide the model to focus explicitly on hand and face regions. However, as the model structure becomes more complex, the number of parameters and required computational resources has grown exponentially, leading to challenges in industrial applications. To address these, we propose a lightweight and fast sign language recognition model based on MediaPipe.

Application of MediaPipe to CSLR. OpenPose (Cao et al. 2017b) and MediaPipe Holistic (Grishchenko and Bazarevsky 2020) models, which are fast and highly accurate, have achieved remarkable results in the field of pose recognition. They can quickly and accurately extract the airspace information of key nodes such as hands, face, and pose in the video. Hence, OpenPose and MediaPipe are applied in ISLR modeling (Camgoz et al. 2020; Liu et al. 2021). Models using human body boundary markers as data require fewer parameters than video-based models, and the model architecture is simpler without the need for complex model structure design (Hu et al. 2023b). The rapid development of the field of sign language recognition (SLR) and interpreting cannot be promoted without commercialized applications. Therefore, a SLR model based on MediaPipe is more suitable for commercial deployment, and MediaPipe supports fast deployment on mobile devices such as Android and iOS.

However, SLR models based on landmarks are far inferior performance-wise to other models that are based on video (Li et al. 2020b). The primary reason for this performance gap is that landmark data only includes spatial information about hand and body movements, whereas video-based sign language often includes more contextual information, such as facial expressions, body posture, and temporal dynamics, which are crucial for accurate interpretation. Relying solely on the limited information provided by landmarks may not

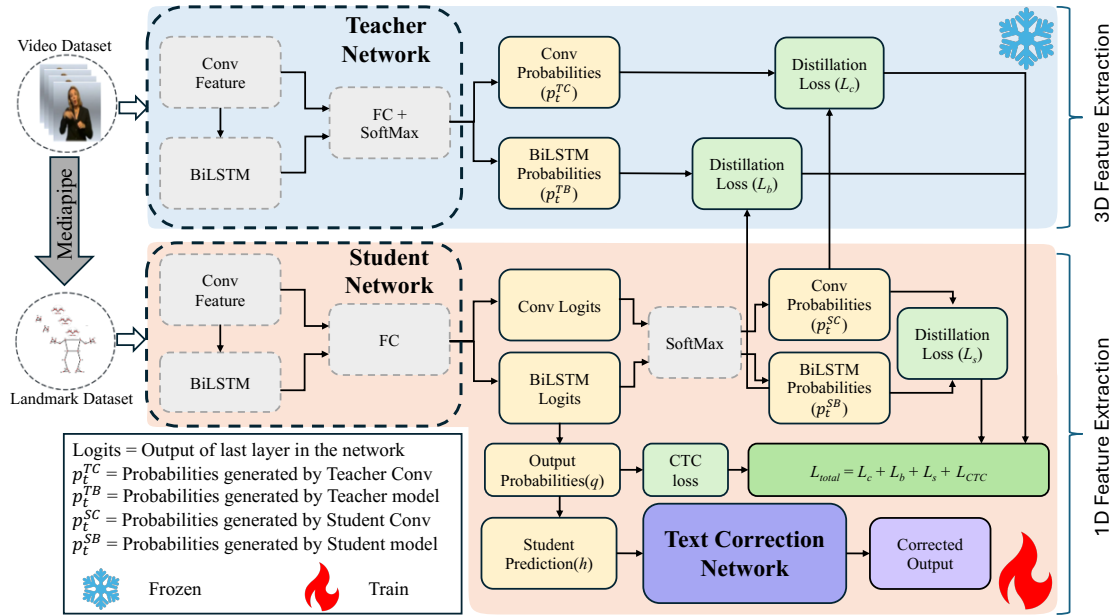


Figure 2: An overview of our KD-MSLRT. We adopt the SOTA video-based sign language model CorrNet (Hu et al. 2023a) as the teacher network. We conduct knowledge distillation separately on the probabilities obtained from the convolutional layers and bidirectional LSTM. The probabilities from the bidirectional LSTM are used to predict sentences, followed by further refinement of the predicted sentences using a text correction network trained on a large corpus.

fully capture the entire linguistic and semantic context conveyed by sign language. To address this, we propose a 3D to 1D knowledge distillation technique and a self-supervised learning text correction model.

Text correction network. Recent years have seen significant advancements in the research on text correction within the field of NLP (Khurana et al. 2023; Kang et al. 2020). The introduction of sequence-to-sequence (Seq2Seq) models (Sutskever, Vinyals, and Le 2014) has led to substantial advancements in this field, particularly in the development of text correction networks. Seq2Seq-based approaches, with their ability to manage complex sequence alignment and transformation, have established a robust framework for addressing textual errors (Yuan et al. 2018). Furthermore, Transformer-based models have significantly advanced the SOTA in text correction (Vaswani et al. 2017). BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al. 2018) and its derivatives have been extensively utilized for pre-training text correction systems, leveraging their capacity to capture rich contextual semantics (Li et al. 2024). However, models based on these two architectures can encounter issues when dealing with random errors, as a single error in the input sequence may lead to a cascade of errors in the output sequence. This error propagation effect is particularly pronounced when processing longer sentences or more complex structures. Additionally, these models primarily rely on patterns and rules learned during pre-training (Tang et al. 2024; Hu et al. 2024). They may fail to recognize the nature of random errors, potentially misinterpreting them as signals. To address these challenges,

we propose a novel end-to-end self-supervised pre-training network specifically designed to correct random errors in the output text of language translation and recognition models.

Methodology

This research introduces the **Knowledge Distillation-Based MediaPipe Sign Language Recognition model with Text Correction Network (KD-MSLRT)**, a lightweight and efficient sign language recognition model designed for resource-constrained environments. KD-MSLRT leverages a novel combination of 3D to 1D knowledge distillation, a MediaPipe-based architecture, landmark data augmentation, and a novel text correction network to address the challenges of accurate and efficient sign language recognition.

Overall Network Design and Knowledge Distillation

KD-MSLRT comprises three primary modules, as illustrated in Figure 2: a robust teacher network, a lightweight student network, and a novel text correction network.

The objective of both the teacher and student networks in this research is to translate the input video or extracted landmark data $X = \{x_t\}_{t=1}^T$ into a series of glosses $Y = \{y_i\}_{i=1}^N$ representing a sentence, where N denotes the length of the glosses label, T represents the time frame. Finally, the classifier uses CTC loss to predict the probability of the target gloss sequence $p(Y|X)$.

The teacher network utilizes CorrNet, a state-of-the-art model for video-based Continuous Sign Language Recognition (CSLR), to provide robust knowledge for the stu-

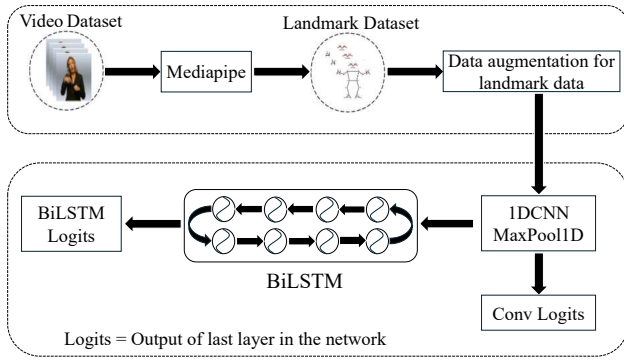


Figure 3: The lightweight sign language recognition model MSLR proposed based on MediaPipe in this research.

dent network. CorrNet operates on individual frames of a sign language video, effectively extracting features and capturing temporal dependencies without explicitly modeling frame continuity. Its key components include a Feature Extraction Layer, a 1D CNN for capturing spatial relationships within frames, a BiLSTM for modeling temporal dependencies between frames, and a Classification Layer for predicting the probability of a gloss sequence. Furthermore, CorrNet incorporates a Correlation Module for identifying body trajectories and an Identification Module for focusing on information-rich regions.

The student network, named the **MediaPipe Sign Language Recognition model (MSLR)**, is designed for lightweight and efficient sign language recognition. MSLR operates on landmark data extracted from videos using MediaPipe, significantly reducing the computational burden while preserving essential information for sign recognition. As shown in Figure 3, MSLR consists of a Convolutional Feature Extraction Layer (four layers of 1D CNN) and a BiLSTM Layer for capturing temporal dependencies, mirroring the teacher network’s structure. Its input is the human landmark data extracted by MediaPipe. The original video data $X_{video} = \{x_t\}_{t=1}^T \in \mathbb{R}^{T \times C \times H \times W}$ is processed by MediaPipe to obtain $X_{landmark} = \{x_t\}_{t=1}^T \in \mathbb{R}^{T \times C \times F}$, where F is the selected key nodes. In this research, only 72 keypoints, including hand and partial facial landmarks, are used as input to the student network, so $F = 72$.

KD-MSLRT employs a 3D to 1D knowledge distillation strategy, transferring knowledge from the teacher network’s convolutional and BiLSTM layers to the student network. This process encourages the student network to learn discriminative features and temporal dependencies similar to the teacher network, enhancing its recognition capabilities. Additionally, both networks share a final classification layer and undergo self-knowledge distillation, further refining the learned feature representations.

The knowledge distillation process is implemented by minimizing a combined loss function that considers the distillation losses from the convolutional layer, BiLSTM layer, and self-knowledge distillation, as defined in Equations 1-5.

$$\mathcal{L}_c = \alpha \sum_{t=1}^T (p_t^{TC} (\log p_t^{TC} - p_t^{SC})), \quad (1)$$

$$\mathcal{L}_b = \alpha \sum_{t=1}^T (p_t^{TB} (\log p_t^{TB} - p_t^{SB})), \quad (2)$$

$$\mathcal{L}_s = \alpha \sum_{t=1}^T (p_t^{SB} (\log p_t^{SB} - p_t^{SC})), \quad (3)$$

$$\mathcal{L}_{CTC}(X, Y) = -\log p(Y | X), \quad (4)$$

$$\mathcal{L}_{total} = \mathcal{L}_c + \mathcal{L}_b + \mathcal{L}_s + \mathcal{L}_{CTC}, \quad (5)$$

where $\mathcal{L}_c, \mathcal{L}_b, \mathcal{L}_s$ are the distillation loss of the convolutional layer, BiLSTM and self-knowledge distillation, respectively, T is the time frame corresponding to the input, α is the temperature parameter, and $p_t^{TC}, p_t^{TB}, p_t^{SC}, p_t^{SB}$ are the probability outputs of different modules as defined in Figure 2.

Landmark Data Augmentation

To enhance the robustness and generalization ability of the student network, KD-MSLRT introduces landmark data augmentation techniques. These techniques augment the training data by introducing variations in landmark positions and sequences, simulating real-world scenarios and improving the network’s ability to generalize to unseen data.

To enhance robustness and generalization, three landmark data augmentation techniques are employed: spatial rotation, which simulates variations in camera viewpoint and signer orientation by rotating landmark coordinates; random translation, which accounts for signer positioning variations within the video frame by shifting the entire landmark sequence; and frame fusion, which generates new frames by combining information from consecutive frames using either average fusion (averaging landmark coordinates) or weighted fusion (combining landmarks with varying degrees of influence).

These augmentation techniques, aligned with the video augmentation techniques used for the teacher network, ensure consistency in the training process and contribute to the student network’s robustness and accuracy.

Text Correction Network

The final component of KD-MSLRT is the text correction network, designed to refine the output from the student network and address potential recognition errors. This network identifies and corrects common errors such as misidentifications, omissions, insertions, and disordered sequences in the recognized gloss sequences, ensuring a more accurate and grammatically correct final translation.

As shown in Figure 4, the text correction network consists of two lightweight Transformer models with a high-encoding, low-decoding structure, enabling efficient real-time performance on resource-constrained devices. The first Transformer Module focuses on restoring the original meaning of the recognized text, preserving the overall intent of the sentence, while the second Transformer Module performs fine-grained text corrections, addressing remaining errors.

A novel self-supervised pre-training approach is employed to train the text correction network effectively. This approach generates training data by introducing artificial errors into correct text, simulating the types of errors commonly made by the student network. This allows the network to learn to identify and rectify similar errors in the

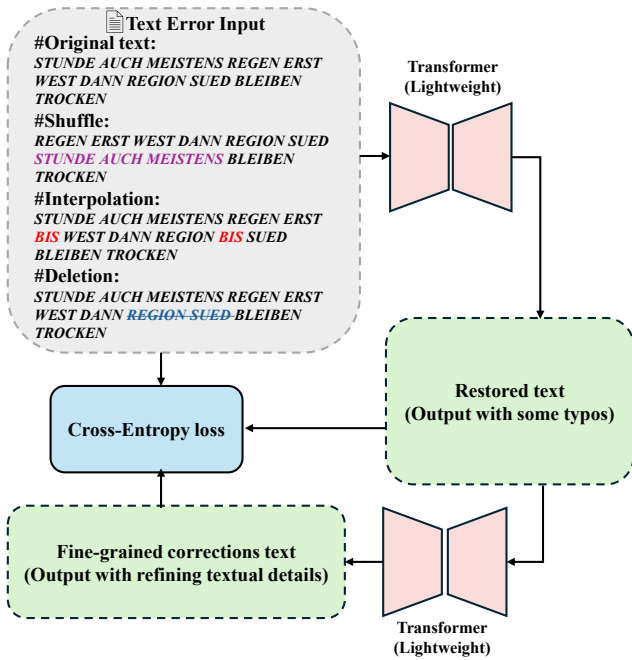


Figure 4: The self-supervised training text correction model proposed based on the error types of the SLR model outputs in this research.




Figure 5: The Chinese long sentence sign language dataset collected for this release consists entirely of news content, including a rich array of professional terms.

actual recognized text, significantly enhancing the accuracy and readability of the final sign language translations.

Our self-supervised pre-training process utilizes both original and noise-augmented text versions. Noise augmentation includes techniques like shuffling, random deletion, and insertion of sequences. This combined approach offers several benefits:

- Preserves Meaning: The original text helps the model learn full word relationships and sentence meaning.
- Expands Training Data: Combining original and modified texts exposes the model to more word combinations, improving its ability to handle new sentences.
- Prevents Overfitting: Including original text ensures the model learns actual semantics, not just patterns from pre-processing.
- Enhances Word Understanding: Direct input of original

text helps the model capture word usage in different contexts.

This approach ultimately enhances the model’s understanding of sentence structure and its ability to interpret relationships between sentences. It also increases the model’s robustness, making it more resilient to noisy or incomplete input. By learning to balance local word relationships with the overall sentence structure, the model can effectively handle complex outputs and correct errors in challenging scenarios, particularly when dealing with noisy data.

Chinese Standard Sign Language Dataset

This research addresses the scarcity of publicly available Chinese Standard Sign Language data, particularly for recognizing long sentences, by introducing a new dataset collected from the official Chinese TV channel program Common Concern, as illustrated in Figure 5. This dataset comprises 8,976 sign language data samples with a resolution of 480*720, making it the largest dataset for recognizing long sentences in Chinese sign language to date.

The dataset’s content focuses on international news, with sign language translations performed by professionally trained sign language interpreters, ensuring data accuracy. Notably, the dataset reflects real-world application scenarios by including instances where some nouns may have missing sign language counterparts, while their corresponding labels remain complete news sentences. Both the original video data and the landmark data processed by MediaPipe will be publicly released, contributing to the advancement of Chinese sign language recognition research.

Experiments

Experiments Setup

Dataset To ensure a fair comparison, the model and data augmentation algorithms proposed in this research will be compared with SOTA methods on the PHOENIX14 (Koller, Forster, and Ney 2015) and PHOENIX14T (Camgoz et al. 2018b) datasets.

PHOENIX14 is recorded from a German weather forecast broadcast with nine actors before a clean back-ground with a resolution of 210 × 260. It contains 6841 sentences with a vocabulary of 1295 signs, divided into 5672 training samples, 540 development (Dev) samples and 629 testing (Test) samples. PHOENIX14-T is a dataset that comprises 8,247 sentences with a vocabulary of 1,085 signs, divided into 7,096 training instances, 519 development instances, and 642 testing instances.

Training details We trained our model for 80 epochs with an initial learning rate of 0.0001, using a weight ratio of 1:25 for CTC loss and KD loss. We used the Adam optimizer with a weight decay of 0.0001 and employed a linear learning rate decay. For the input frames of the teacher network, we cropped them to 224*224 pixels and applied horizontal flipping on 50%, rotation on 30%, and temporal scaling on 20% of the frames. For the landmarks of the student network, we applied horizontal flipping on 50%, rotation on 30%, and frame fusion on 20% of the frames. Model training and evaluation were conducted on an Nvidia RTX 3080

Methods	PHOENIX14		PHOENIX14T		Mean CPU (milliseconds)	FLOPs	Input Shape
	Dev (%)	Test (%)	Dev (%)	Test (%)			
	WER	WER	WER	WER			
SFL (Niu and Mak 2020)	26.2	26.8	25.1	26.1	-	-	-
FCN (Cheng et al. 2020b)	23.7	23.9	23.3	25.1	-	-	-
CMA (Pu et al. 2020)	21.3	21.9	-	-	-	-	-
VAC (Min et al. 2021)	21.2	22.3	-	-	-	5×10^{11}	-
SMKD (Hao, Min, and Chen 2021)	20.8	21	20.8	22.4	-	-	-
TLP (Hu et al. 2022)	19.7	20.8	19.4	21.2	-	4×10^{11}	-
SEN (Huang et al. 2018)	19.5	21	19.3	20.7	-	-	-
CNN+LSTM+HMM (Camgoz et al. 2018a)	26	26	22.1	24.1	-	-	-
DNF* (Koller et al. 2019)	23.1	22.9	-	-	-	5×10^{11}	-
STMC* (Cui, Liu, and Zhang 2019)	21.1	20.7	19.6	21.0	-	9×10^{10}	-
C2SLR* (Zuo and Mak 2022)	20.5	20.4	20.2	20.4	-	8×10^{10}	-
CorrNet (Hu et al. 2023a)	18.8	19.4	18.9	20.5	5484.94	5×10^{11}	[T,3,224,244]
Landmark base							
MSLR	39.6	41.2	40.2	41.1	28.02	4×10^8	[T,276]
KD-MSLR	28.7	30.7	29.6	31.2	28.02	4×10^8	[T,276]
KD-MSLRT	17.4	17.8	16.9	18.4	42.64	1.9×10^9	[T,276]

Table 1: Comparison with the SOTA video-based SLR model on the PHOENIX 14 and PHOENIX14T datasets. * indicates that both video and landmark data types are used. Results better than the SOTAs are in bold.

graphics card, while CPU speed tests were performed on an Intel 13th generation i7 processor.

Evaluation Metric In the CSLR field, the Word Error Rate (WER) serves as a common evaluation metric. It categorizes errors in predicted sentences into substitutions, insertions, and deletions, then calculates the minimum sum of these errors as the evaluation standard as:

$$WER = \frac{\#substitutions + \#insertions + \#deletions}{\#glosses\ in\ reference}, \quad (6)$$

Note: The smaller the WER, the better the performance.

Comparison with SOTA Methods:

Table 1 gives the comparison between our KD-MSLRT model and other SOTA models. Currently, SOTA models on the PHOENIX14 and PHOENIX14T datasets are all based on video data format. Therefore, these models have larger parameters and computational requirements, resulting in relatively longer inference times. In contrast, our proposed model is based on landmark data extracted using MediaPipe, which compresses the data size by nearly 700 times. Our KD-MSLRT model reduces FLOPs by 260 times compared to SOTA models, significantly improving inference speed and runtime by 130 times. The TensorFlow Lite (TFLite) quantized model size is reduced to 12.93 MB, achieving millisecond-level inference speed on CPUs. Additionally, the KD-MSLRT model achieves a WER of 17.4% (17.8%), a 1.4% (1.6%) improvement on Dev (Test) set of PHOENIX14, and 16.9% (18.4%), a 2% (2.1%) improvement on PHOENIX14T. Although the FLOPs of KD-MSLRT increases significantly after applying the text correction model, the input of the text correction model consists of only 10-30 characters, resulting in minimal computational overhead, with slight increase of 14 milliseconds in the runtime.

Industrial Deployment

We conducted a simulated deployment on mainstream development boards within the Intel i5 platform to validate the feasibility and cost-effectiveness of deploying our model in industrial environments. Table 2 presents the performance of our model on the PHOENIX14 and PHOENIX14T datasets at FP32 and INT8 precision levels, while Table 3 details the frames per second (FPS) achieved by our model during inference. The experimental results demonstrate that at INT8 precision, the model achieves a Word Error Rate (WER) comparable to that of FP32, with significantly lower computational costs. We further evaluated the model’s inference speed, and the results indicate that on the Intel i5 platform, the model achieves high FPS across different precision levels, meeting the real-time requirements of industrial applications. The efficient inference at INT8 precision suggests that deployment on edge computing devices is feasible and can potentially reduce operational costs. This makes our model highly cost-effective for industrial deployment and suitable for large-scale application.

Datasets	FP32 (WER)	INT8 (WER)	Diff	Orisize	Onnsize
PHOENIX14	17.4%/17.8%	17.5%/17.9%	0.5%/0.4%	22.34MB	12.93MB
PHOENIX14T	16.9%/18.4%	16.8%/18.5%	-0.1%/0.1%	23.84MB	13.81MB

Table 2: Model performance comparison between FP32 and INT8 quantization.

Datasets	FP32 (frame/s)	INT8 (frame/s)	Ratio (INT8/FP32)
PHOENIX14	11678	31369	2.69
PHOENIX14T	10892	30154	2.74

Table 3: FPS achieved by the model during inference.

Ablation Experiments

We present the results of ablation experiments on the development (Dev) and testing (Test) sets of the PHOENIX14 and PHOENIX14T.

Exploring the effects of data augmentation, text correction (TC), and knowledge distillation (KD) on MSLR performance. Table 4 gives that the WER of using only the text correction network is not as effective as when knowledge distillation and text correction network are used together. This is because the text correction model adjusts sentences based on the most dominant meaning expressed in the text. When the model’s recognized text has a WER of over 40%, the sentence may contain multiple conflicting meanings. If the number of words representing the incorrect meaning exceeds those representing the correct meaning, the text correction model may revise the text to reflect the incorrect meaning. Therefore, the text correction network performs better when sentences contain a sufficient amount of accurate information.

Data Augmentation	KD	TC	Dev (%)	Test (%)
×	×	×	43.8/42.7	45.2/45.6
✓	×	×	39.6/38.2	41.2/40.7
✓	×	✓	30.2/31.4	31.7/31.9
×	✓	×	28.7/28.4	30.7/29.3
×	×	✓	31.4/30.9	32.7/32.1
×	✓	✓	18.2/17.9	18.7/19.1
✓	✓	✓	17.4/16.9	17.8/18.4

Table 4: Ablations for the performance improvement of MSLR by different modules. The best results are in bold.

Exploring the effects of single Transformer, dual Transformer, and data preprocessing on KD-MSLRT performance. Table 5 gives the results of ablation experiments conducted on different modules of KD-MSLRT. The use of a dual Transformer (tf) significantly improves the model’s WER. Additionally, performing data preprocessing on the raw input text achieved further improvements of 0.9% (1%) and 1.2% (0.7%) on the Dev (Test) set of PHOENIX14 and PHOENIX14T, respectively.

Single tf	Dual tf	Pre-process	Dev (%)	Test (%)
✓	×	×	22.1/21.3	23.2/22.4
✓	✓	×	18.3/18.1	18.8/19.1
✓	×	✓	20.1/19.3	20.6/19.9
✓	✓	✓	17.4/16.9	17.8/18.4

Table 5: Ablations for the performance improvement of KD-MSLRT by different modules. The best results are in bold.

Research on the impact of 1D CNN convolutional kernel size on landmark data processing performance. In Table 6, the results indicate that in the MSLR model, large convolutional kernels perform poorly when processing landmark data. The optimal kernel size is $K = 5$, achieving WER of 43.8% (45.2%) and 42.6% (44.8%) on the Dev (Test) set of two datasets, respectively. Although using only the basic MSLR model yields significant advantages in terms of inference speed and parameter size, there is a considerable gap in performance compared to models based on video. This is because a significant amount of spatial information are lost during the process of converting video data to landmark.

Configurations	Dev (%)	Test (%)
$K = 3$	45.7/44.9	46.8/47.1
$K = 5$	43.8/42.6	45.2/44.8
$K = 7$	47.9/46.3	48.7/47.3
$K = 9$	49.1/48.2	50.9/50.3

Table 6: Ablations for the number of 1DCNN convolutional kernels on PHOENIX14 and PHOENIX14T datasets.

Exploring the effect of knowledge distillation on MSLR improvement. Knowledge distillation improves model performance by transferring knowledge from a complex teacher model to a simpler student model, enhancing accuracy and efficiency. Using only convolutional layer knowledge distillation (KD_{conv}) in the Table 7 resulted in MSLR improving by 9.2% (9.1%) and 8.8% (8.4%) on Dev (Test) set of two datasets respectively. Using only knowledge distillation of BiLSTM (KD_{bi}) led to MSLR improving by 13.6% (13.8%) and 11.9% (14.7%) on Dev (Test) set. However, the improvement from using only self distillation (KD_{self}) was not significant, with only a 4.3% (4%) and 2.6% (5.2%) improvement on Dev (Test) set. It shows that using a 3D model as the teacher network for distilling knowledge into a 1D model significant improvements. The combined use of all three resulted in a 15.1% (15.1%) and 13.6% (16.2%) improvement on Dev (Test) set.

KD_{conv}	KD_{bi}	KD_{self}	Dev (%)	Test (%)
×	×	×	43.8/42.7	45.2/45.6
✓	×	×	34.6/33.9	36.1/37.2
✓	×	✓	32.1/31.4	34.3/34.9
×	✓	×	30.2/30.8	31.4/30.9
×	×	✓	39.5/40.1	41.2/40.4
×	✓	✓	29.3/29.6	30.7/31.2
✓	✓	✓	28.7/29.1	30.1/29.4

Table 7: Performance of Base MSLR with different KD configurations. The best results are in bold.

Conclusion

In this research, we propose a lightweight SLR model based on mediapipe and 3D to 1D knowledge distillation. It addresses the issues of high resource consumption, long runtime, and high deployment costs associated with video-based SLR models. Furthermore, to address the gap in augmentation methods for landmark data, we propose augmentation methods specifically tailored for landmark data based on techniques used for video and image augmentation. Additionally, to address the current lack of long-sentence datasets in the sign language pattern recognition field, we release a dataset containing 8,976 long-sentence sign language samples. We will make all aspects of this work publicly available, including datasets, in hope that these efforts will provide valuable insights for future research.

References

- Albanie, S.; Varol, G.; Momeni, L.; Afouras, T.; Chung, J. S.; Fox, N.; and Zisserman, A. 2020. BSL-1K: Scaling up co-articulated sign language recognition using mouthing cues. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, 35–53. Springer.
- Camgoz, N. C.; Hadfield, S.; Koller, O.; Ney, H.; and Bowden, R. 2018a. Neural sign language translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7784–7793.
- Camgoz, N. C.; Hadfield, S.; Koller, O.; Ney, H.; and Bowden, R. 2018b. Neural sign language translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7784–7793.
- Camgöz, N. C.; Kindiroğlu, A. A.; Karabüklü, S.; Keleş, M.; Özsoy, A. S.; and Akarun, L. 2016. BosphorusSign: A Turkish sign language recognition corpus in health and finance domains. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, 1383–1388.
- Camgoz, N. C.; Koller, O.; Hadfield, S.; and Bowden, R. 2020. Multi-channel transformers for multi-articulatory sign language translation. In *Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, 301–319. Springer.
- Cao, Z.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2017a. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7291–7299.
- Cao, Z.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2017b. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7291–7299.
- Chen, Y.; Wei, F.; Sun, X.; Wu, Z.; and Lin, S. 2022. A simple multi-modality transfer learning baseline for sign language translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5120–5130.
- Cheng, K. L.; Yang, Z.; Chen, Q.; and Tai, Y.-W. 2020a. Fully convolutional networks for continuous sign language recognition. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16*, 697–714. Springer.
- Cheng, K. L.; Yang, Z.; Chen, Q.; and Tai, Y.-W. 2020b. Fully convolutional networks for continuous sign language recognition. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16*, 697–714. Springer.
- Cihan Camgoz, N.; Hadfield, S.; Koller, O.; and Bowden, R. 2017. Subnets: End-to-end hand shape and continuous sign language recognition. In *Proceedings of the IEEE international conference on computer vision*, 3056–3065.
- Cui, R.; Liu, H.; and Zhang, C. 2019. A deep neural framework for continuous sign language recognition by iterative training. *IEEE Transactions on Multimedia*, 21(7): 1880–1891.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Fels, S. S.; and Hinton, G. E. 1993. Glove-talk: A neural network interface between a data-glove and a speech synthesizer. *IEEE transactions on Neural Networks*, 4(1): 2–8.
- Graves, A.; Fernández, S.; Gomez, F.; and Schmidhuber, J. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, 369–376.
- Grishchenko, I.; and Bazarevsky, V. 2020. Mediapipe holistic—simultaneous face, hand and pose prediction, on device. *Google AI Blog. Dec*.
- Hao, A.; Min, Y.; and Chen, X. 2021. Self-mutual distillation learning for continuous sign language recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11303–11312.
- Hu, L.; Gao, L.; Liu, Z.; and Feng, W. 2022. Temporal lift pooling for continuous sign language recognition. In *European conference on computer vision*, 511–527. Springer.
- Hu, L.; Gao, L.; Liu, Z.; and Feng, W. 2023a. Continuous sign language recognition with correlation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2529–2539.
- Hu, L.; Gao, L.; Liu, Z.; and Feng, W. 2023b. Continuous sign language recognition with correlation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2529–2539.
- Hu, M.; Xia, P.; Wang, L.; Yan, S.; Tang, F.; Xu, Z.; Luo, Y.; Song, K.; Leitner, J.; Cheng, X.; et al. 2025. Ophnet: A large-scale video benchmark for ophthalmic surgical workflow understanding. In *ECCV*.
- Hu, M.; Yuan, K.; Shen, Y.; Tang, F.; Xu, X.; Zhou, L.; Li, W.; Chen, Y.; Xu, Z.; Peng, Z.; et al. 2024. OphCLIP: Hierarchical Retrieval-Augmented Learning for Ophthalmic Surgical Video-Language Pretraining. *arXiv preprint arXiv:2411.15421*.
- Huang, J.; Zhou, W.; Zhang, Q.; Li, H.; and Li, W. 2018. Video-based sign language recognition without temporal segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Kang, Y.; Cai, Z.; Tan, C.-W.; Huang, Q.; and Liu, H. 2020. Natural language processing (NLP) in management research: A literature review. *Journal of Management Analytics*, 7(2): 139–172.
- Khurana, D.; Koli, A.; Khatter, K.; and Singh, S. 2023. Natural language processing: state of the art, current trends and challenges. *Multimedia tools and applications*, 82(3): 3713–3744.
- Koller, O.; Camgoz, N. C.; Ney, H.; and Bowden, R. 2019. Weakly supervised learning with multi-stream CNN-LSTM-HMMs to discover sequential parallelism in sign language videos. *IEEE transactions on pattern analysis and machine intelligence*, 42(9): 2306–2320.

- Koller, O.; Forster, J.; and Ney, H. 2015. Continuous sign language recognition: Towards large vocabulary statistical recognition systems handling multiple signers. *Computer Vision and Image Understanding*, 141: 108–125.
- Koller, O.; Zargaran, S.; and Ney, H. 2017. Re-sign: Re-aligned end-to-end sequence modelling with deep recurrent CNN-HMMs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4297–4305.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Li, D.; Rodriguez, C.; Yu, X.; and Li, H. 2020a. 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*, 1459–1469.
- Li, D.; Rodriguez, C.; Yu, X.; and Li, H. 2020b. 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*, 1459–1469.
- Li, D.; Xu, C.; Yu, X.; Zhang, K.; Swift, B.; Suominen, H.; and Li, H. 2020c. Tspnet: Hierarchical feature learning via temporal semantic pyramid for sign language translation. *Advances in Neural Information Processing Systems*, 33: 12034–12045.
- Li, J.; Tang, T.; Zhao, W. X.; Nie, J.-Y.; and Wen, J.-R. 2024. Pre-trained language models for text generation: A survey. *ACM Computing Surveys*, 56(9): 1–39.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Min, Y.; Hao, A.; Chai, X.; and Chen, X. 2021. Visual alignment constraint for continuous sign language recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11542–11551.
- Niu, Z.; and Mak, B. 2020. Stochastic fine-grained labeling of multi-state sign glosses for continuous sign language recognition. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, 172–186. Springer.
- Pu, J.; Zhou, W.; Hu, H.; and Li, H. 2020. Boosting continuous sign language recognition via cross modality augmentation. In *Proceedings of the 28th ACM international conference on multimedia*, 1497–1505.
- Secretariat, G. 2009. Deaf People and Human Rights.
- Selvaraj, P.; Nc, G.; Kumar, P.; and Khapra, M. 2021. Openhands: Making sign language recognition accessible with pose-based pretrained models across languages. *arXiv preprint arXiv:2110.05877*.
- Sincan, O. M.; and Keles, H. Y. 2020. Autsl: A large scale multi-modal turkish sign language dataset and baseline methods. *IEEE access*, 8: 181340–181355.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.
- Tang, F.; Xu, Z.; Qu, Z.; Feng, W.; Jiang, X.; and Ge, Z. 2024. Hunting Attributes: Context Prototype-Aware Learning for Weakly Supervised Semantic Segmentation. In *CVPR*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, J.; Huang, Q.; Tang, F.; Meng, J.; Su, J.; and Song, S. 2022. Stepwise feature fusion: Local guides global. In *MICCAI*. Springer.
- Yin, K.; and Read, J. 2020. Better sign language translation with STMC-transformer. *arXiv preprint arXiv:2004.00588*.
- Yuan, X.; Wang, T.; Meng, R.; Thaker, K.; Brusilovsky, P.; He, D.; and Trischler, A. 2018. One size does not fit all: Generating and evaluating variable number of keyphrases. *arXiv preprint arXiv:1810.05241*.
- Zhao, X.; Tang, F.; Wang, X.; and Xiao, J. 2024. Sfc: Shared feature calibration in weakly supervised semantic segmentation. In *AAAI*.
- Zuo, R.; and Mak, B. 2022. C2slr: Consistency-enhanced continuous sign language recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5131–5140.