

Improving Continuous Sign Language Recognition with Cross-Lingual Signs

Fangyun Wei
Microsoft Research Asia
fawe@microsoft.com

Yutong Chen
Microsoft Research Asia
chenytjudy@gmail.com

Abstract

This work dedicates to continuous sign language recognition (CSLR), which is a weakly supervised task dealing with the recognition of continuous signs from videos, without any prior knowledge about the temporal boundaries between consecutive signs. Data scarcity heavily impedes the progress of CSLR. Existing approaches typically train CSLR models on a monolingual corpus, which is orders of magnitude smaller than that of speech recognition. In this work, we explore the feasibility of utilizing multilingual sign language corpora to facilitate monolingual CSLR. Our work is built upon the observation of cross-lingual signs, which originate from different sign languages but have similar visual signals (e.g., hand shape and motion). The underlying idea of our approach is to identify the cross-lingual signs in one sign language and properly leverage them as auxiliary training data to improve the recognition capability of another. To achieve the goal, we first build two sign language dictionaries containing isolated signs that appear in two datasets. Then we identify the sign-to-sign mappings between two sign languages via a well-optimized isolated sign language recognition model. At last, we train a CSLR model on the combination of the target data with original labels and the auxiliary data with mapped labels. Experimentally, our approach achieves state-of-the-art performance on two widely-used CSLR datasets: *Phoenix-2014* and *Phoenix-2014T*.

1. Introduction

Sign languages are visual-spatial signals for communication among deaf communities. These languages are primarily expressed through hand shape but are also greatly aided by the movement of the body, head, mouth, and eyes. Sign language recognition, which aims to establish communication between hearing people and deaf people, can be roughly categorized into two sorts: isolated sign language recognition (ISLR) [49, 26, 19, 27] and continuous sign language recognition (CSLR) [24, 8, 22, 47, 46, 10]. ISLR is a supervised classification task—it requires models to recognize and classify isolated signs from videos. In contrast,



DGS: SONNE (Sun)




CSL: 灯 (Light)

(a) Cross-lingual signs usually have distinct meanings.



DGS: KALT (Cold)



CSL: 冷 (Cold)

(b) Cross-lingual signs occasionally convey the same meaning.

Figure 1: Cross-lingual signs are those that originate from different sign languages but have similar visual signals (e.g. hand shape and motion). We show two examples identified by our approach from a German sign language (DGS) dataset and a Chinese sign language (CSL) dataset.

CSLR is a weakly supervised task dedicated to the recognition of continuous signs from videos, without any prior knowledge about the temporal boundaries between consecutive signs. The objective of this work is to develop a CSLR framework, with the assistance of an ISLR model and multilingual corpus.

In contrast to the promising achievements in automatic speech recognition [28, 2, 35, 41], the lack of large-scale training data heavily impedes the progress of CSLR. In general, training a satisfying speech recognition model requires thousands of hours of training data [20, 34]. However, existing sign language datasets [23, 6, 44] are orders of magnitude smaller, containing only fewer than 20




Figure 2: Illustration of our approach. The objective is to train a continuous sign language recognition (CSLR) model in DGS with the assistance of a CSL dataset. We first build two sign dictionaries containing isolated signs in DGS and CSL respectively. Then we train an isolated sign language recognition (ISLR) model to identify the CSL-to-DGS mapping for each isolated sign in the CSL dictionary. Finally, the CSLR-DGS model is trained on a combination of the DGS dataset with original labels and the CSL dataset with *mapped* labels through the CTC loss [16].

hours of paralleled samples. Previous approaches [24, 8, 22, 47, 46, 10] typically train CSLR models on monolingual corpora such as German sign language (Deutsche Gebärdensprache, DGS) datasets (*e.g.* Phoenix-2014 [23] and Phoenix-2014T [6]), and Chinese sign language (CSL) datasets (*e.g.* CSL-Daily [44]). Unfortunately, the limited training data immensely restricts the recognition capacity. One possible technique to alleviate the data scarcity issue is semi-supervised learning, which requires a large volume of unlabeled data in addition to a collection of labeled data. However, existing CSLR datasets are collected in specific domains, *e.g.*, Phoenix-2014 [23] and Phoenix-2014T [6] are concentrated in the domain of weather forecast. Collecting domain-relevant data for CSLR becomes challenging and unpractical, impeding semi-supervised learning.

Nevertheless, since sign languages are visual languages, it is possible for them to employ the same sign to express either the same or different meanings. If the hypothesis is true, it will be feasible to utilize multilingual sign language corpora to enrich training data. Fortunately, we conduct experiments to confirm the existence of these signs, which are referred to as cross-lingual signs in our work. Cross-lingual signs are those that originate from different sign languages but have similar visual signals (*e.g.*, hand shape and motion). Since most sign languages are mutually unintelligible, cross-lingual signs typically have distinct meanings in different sign languages (Figure 1a). Interestingly, we find that they might convey the same meanings occasionally (Figure 1b). The discovery of cross-lingual signs inspires us to identify these signs in one sign language and properly leverage them as auxiliary training data to improve the recognition performance of the other.

Consider a scenario where we are going to train a CSLR model in DGS given a primary DGS dataset and an auxiliary dataset in another sign language, *e.g.*, CSL. In general, the size of the auxiliary dataset should exceed that of the primary one. The underlying idea behind our approach is to find the CSL-to-DGS mapping for each isolated sign in the CSL dataset. Figure 2 illustrates our so-

lution, which contains three steps: (1) build two sign language dictionaries containing isolated signs that appear in CSL and DGS datasets, respectively; (2) identify the CSL-to-DGS mapping for each isolated sign in CSL dictionary according to the cross-lingual sign similarity calculated by a well-optimized isolated sign language recognition model; (3) train a CSLR model in DGS on the combination of DGS data with original labels and CSL data with mapped labels through the well-established CTC loss [16]. It is worth mentioning that it is non-trivial to directly build a sign language dictionary from a CSLR dataset due to the absence of sign boundary annotations. To tackle this problem, we adopt a pre-trained CSLR model to split the continuous signs into the isolated ones for dictionary construction.

The contributions of this work can be summarized as:

- We are the first to utilize a multilingual sign language corpus to facilitate monolingual CSLR based on the finding of cross-lingual signs.
- We present a comprehensive solution for seamlessly incorporating an auxiliary dataset—though in another sign language—into the training of CSLR.
- Our approach achieves state-of-the-art performance on two widely-used CSLR datasets: Phoenix-2014 [23] and Phoenix-2014T [6].

2. Related Works

Sign Language Recognition can be categorized into isolated sign language recognition (ISLR) and continuous sign language recognition (CSLR). ISLR aims to classify an isolated sign video into a single sign class [49, 26, 19, 27, 1]. CSLR aims to word-by-word transcribe a co-articulated sign video into a sign sequence [24, 8, 10, 46, 14, 29]. Similar to autonomous speech recognition (ASR) [43, 38, 28], CSLR is a weakly supervised sequence-to-sequence task without temporal boundary annotation available and hence is typically trained using CTC loss [16]. To address the scarcity of parallel CSLR annotation, Pu et al. estimate

the temporal boundaries and stitch video clips to generate pseudo parallel data for data augmentation [36]. Some methods iteratively optimize the model with a sequence-level CTC loss and an auxiliary frame-level classification loss [37, 14]. Other works exploit multi-stream networks with multi-modal inputs or labels [45, 10, 22, 7], or recognize the meanings that signers express through sign language retrieval [12]. Our CSLR method alleviates data scarcity by leveraging multilingual training data and integrating ISLR models into the CSLR framework to find cross-lingual signs.

Cross-lingual Transfer is well explored in speech-to-text recognition [40, 25, 21, 3] and spoken language understanding [18, 30, 13]. This line of research commonly finds that training a model on multilingual corpora can yield superior performance on low-resource languages compared to training it on a monolingual corpus. In the context of sign language understanding, there are also a few publications investigating cross-lingual transfer [32, 39, 42]. For sign language translation, MLSLT [42] proposes a multilingual translation network with language-specific parameters that outperforms the monolingual baseline. For ISLR, Gokul et al. combine labels from different sign languages based on their word meanings [32]. Despite bringing improvements, this approach overlooks the fact that the same words can be expressed differently across different sign languages. For CSLR, Tornay et al. train a hand movement model using a different sign language resource before optimizing the classifier using the target sign language data in their KL-HMM framework [39]. However, their cross-lingual model falls short of the monolingual reference. Our CSLR method utilizes both target and auxiliary sign language data via a cross-lingual sign mapping and shows improvements over the monolingual baseline.

3. Method

Objective. The goal of continuous sign language recognition (CSLR) is to recognize a sequence of signs from videos. CSLR is a weakly supervised task since the temporal boundaries between consecutive signs are unprovided. In this work, we attempt to utilize multilingual sign language corpus to facilitate monolingual CSLR based on the finding of cross-lingual signs. Concretely, given a *primary* CSLR dataset \mathcal{D}_P in *target* sign language, and an *auxiliary* CSLR dataset \mathcal{D}_A in *source* sign language, the objective of this work is to train a CSLR model using both \mathcal{D}_P and \mathcal{D}_A to improve the recognition performance on \mathcal{D}_P . In general, the size of \mathcal{D}_A should exceed that of \mathcal{D}_P .

Overview. Figure 3 visualizes the training process of our approach, which consists of three steps: 1) build two sign dictionaries \mathcal{C}_P and \mathcal{C}_A containing isolated signs which appear in \mathcal{D}_P and \mathcal{D}_A , respectively (Section 3.1); 2) identify

source-to-target mapping for each isolated sign in \mathcal{C}_A (Section 3.2); 3) train a CSLR model in target sign language using \mathcal{D}_P with original labels and \mathcal{D}_A with mapped labels (Section 3.3).

3.1. Dictionary Construction

A sign dictionary \mathcal{C} is a set of isolated signs, each of which has a number of video instances associated with it. Existing CSLR datasets do not provide the corresponding dictionaries. To circumvent this issue, we propose to construct sign dictionaries on CSLR datasets. However, it is not feasible to simply split continuous signs into isolated ones without effort, as sign boundaries are not readily available in CSLR datasets. Fortunately, a well-optimized CSLR model is able to predict the boundaries of signs in addition to its main capability of recognizing a sequence of signs, which inspires us to employ a pre-trained CSLR model to automatically split the continuous signs into isolated ones to construct dictionaries, as shown in Figure 3a.

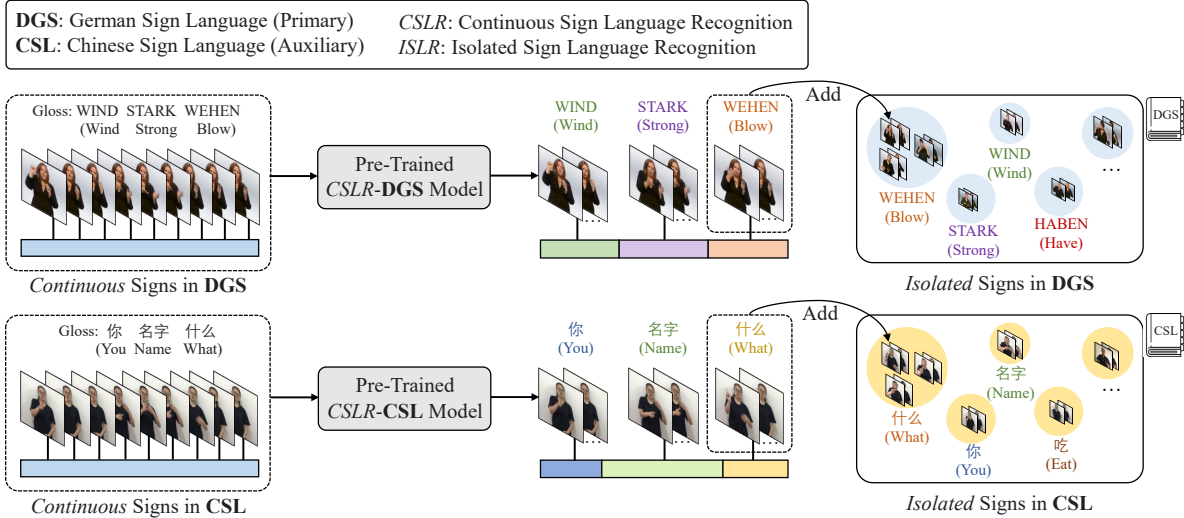
Concretely, given a CSLR dataset \mathcal{D} (\mathcal{D}_P or \mathcal{D}_A), we first train a CSLR model on it. This model can be any CSLR model [45, 47, 9, 36, 10] which is able to produce frame-wise predictions. In this work, we adopt the TwoStream-SLR model [10] due to its superior performance. Once the CSLR model is well optimized, given a training video $v \in \mathcal{D}$ containing T frames and its associated ground-truth sequence label $s = (s_1, \dots, s_N)$ containing N consecutive signs, we use the pre-trained CSLR model to compute the most probable alignment path $\pi^* = (\pi_1, \dots, \pi_T)$, $\pi_t \in \{s_i\}_{i=1}^N$ where π_t predicts which sign is being performed at the t -th frame.

The computation of π^* can be efficiently implemented via the dynamic time warping (DTW) algorithm [4]. The details are described in the supplementary material. With π^* indicating frame-wise prediction, we can split v into N non-overlapped clips where each clip is associated with an isolated sign appeared in s . Then we add each clip to the collection of its associated sign in dictionary \mathcal{C} . We repeat the above process for each video in \mathcal{D} . Finally, we obtain a sign dictionary \mathcal{C} with an alphabet \mathcal{S} . The alphabet \mathcal{S} is the set of all signs appearing in \mathcal{D} . In the constructed dictionary \mathcal{C} , each sign $s \in \mathcal{S}$ has a collection of isolated sign videos.


We apply two CSLR models trained on \mathcal{D}_P and \mathcal{D}_A to construct their corresponding sign dictionaries \mathcal{C}_P with alphabet \mathcal{S}_P , and \mathcal{C}_A with alphabet \mathcal{S}_A , respectively.

3.2. Cross-lingual Sign Mapping


Now we introduce the process of identifying source-to-target mapping for each isolated sign in \mathcal{S}_A . To measure sign-to-sign similarities, we adopt a simple yet effective approach by training a multilingual isolated sign classifier using the two dictionaries \mathcal{C}_A and \mathcal{C}_P . The unified classifier encodes signs from both dictionaries into a shared embed-




(a) Dictionary construction for DGS and CSL datasets.



(b) Train a multilingual ISLR model.



(c) Identify CSL-to-DGS sign mappings through cross-lingual similarities.



(d) Train a CSLR model in DGS on the combination of the DGS dataset with raw labels and the CSL dataset with mapped labels.

Figure 3: Overview of our method. Given a primary dataset in DGS and an auxiliary dataset in CSL, we train a CSLR model in DGS with the assistance of an auxiliary CSL dataset. Our method consists of four steps as illustrated in (a)-(d).

ding space where visually similar signs (cross-lingual signs) are closer to each other than dissimilar ones. By doing so, we can employ the classifier to find visually similar signs across the two sign languages and establish a cross-lingual sign mapping.

3.2.1 Multilingual ISLR

Isolated sign language recognition (ISLR) [19, 26, 27] aims to classify isolated signs. Unlike prior work which trains ISLR models on monolingual datasets, we train a unified classifier on the combination of \mathcal{C}_P and \mathcal{C}_A for multilingual

prediction. Figure 3b shows our ISLR model, which consists of a shared vision encoder and two language-aware classification layers. The vision encoder encodes input videos from \mathcal{C}_P and \mathcal{C}_A into a shared embedding space. The two separate classification layers then project embeddings into two probability distributions over \mathcal{S}_P and \mathcal{S}_A respectively. We train the multilingual ISLR model by minimizing the sum of two cross-entropy losses on \mathcal{C}_P and \mathcal{C}_A . Besides, as both isolated training sets are heavily long-tailed, we filter out signs with low frequency to alleviate the class

imbalance issue. We find that once optimized, the shared vision encoder can automatically align the embedding spaces of \mathcal{C}_P and \mathcal{C}_A —it draws close sign videos that are visually similar regardless of the data source.

3.2.2 Cross-lingual Sign Mapping

By using the trained multilingual ISLR model which encodes isolated sign videos in a multilingual embedding space, we are now able to identify the cross-lingual sign mappings between the signs in \mathcal{S}_A and those in \mathcal{S}_P . Formally, given a sign $s \in \mathcal{S}_A$ belonging to the auxiliary language, we intend to find its mapped sign $\mathcal{M}_{A \rightarrow P}(s) \in \mathcal{S}_P$ belonging to the primary sign language. We propose two strategies for identifying such mappings, utilizing our multilingual ISLR model, as described below.

Cross-lingual Prediction. This is our default mapping strategy as shown in Figure 3c. Recall that our ISLR model is composed of a shared vision encoder to extract vision features, and two classification layers to classify sign videos from \mathcal{C}_A and \mathcal{C}_P respectively. For the purpose of computing cross-lingual similarities, given a sign video from \mathcal{C}_A , we feed it through the vision encoder and then the classification layer of \mathcal{C}_P to predict its probability distribution over \mathcal{S}_P . In other words, we can generate a cross-lingual probability for each video in \mathcal{C}_A by switching the classification layer to that of \mathcal{C}_P . We denote the cross-lingual probability of $v \in \mathcal{C}_A$ as $p_{A \rightarrow P}(v) \in \mathbb{R}^{|\mathcal{S}_P|}$, which predicts the distribution of video v over the alphabet \mathcal{S}_P . The highest activation in $p_{A \rightarrow P}(v)$ indicates the most similar sign in \mathcal{S}_P to the video v . To map a sign $s \in \mathcal{S}_A$ to \mathcal{S}_P , we average the cross-lingual predictions of its all associated video instances:

$$\begin{aligned} \mathcal{M}_{A \rightarrow P}^{\text{cls}}(s) &= \operatorname{argmax}(p_{A \rightarrow P}(s)), \\ p_{A \rightarrow P}(s) &= \frac{1}{|\mathcal{C}_A(s)|} \sum_{v \in \mathcal{C}_A(s)} p_{A \rightarrow P}(v), \quad s \in \mathcal{S}_A, \end{aligned} \quad (1)$$

where $\mathcal{C}_A(s)$ denotes the set of videos in \mathcal{C}_A which are associated with the sign $s \in \mathcal{S}_A$. This type of mapping is referred to as a *class-level mapping*. Note that all video instances belonging to the same sign share identical mapping. In addition, we also explore a variant, *instance-level mapping*, which identifies mapping for each individual video instance. We use $\mathcal{M}_{A \rightarrow P}^{\text{ins}}(v)$ to denote the instance-level mapping of a video $v \in \mathcal{C}_A$, which is formulated as:

$$\mathcal{M}_{A \rightarrow P}^{\text{ins}}(v) = \operatorname{argmax}(p_{A \rightarrow P}(v)), \quad v \in \mathcal{C}_A. \quad (2)$$

Dot-product of Weight Matrices. Another way to compute cross-lingual similarity is to utilize the weights of two well-optimized classification layers, which we denote as $\mathbf{W}_A \in \mathbb{R}^{|\mathcal{S}_A| \times d}$ and $\mathbf{W}_P \in \mathbb{R}^{|\mathcal{S}_P| \times d}$, respectively. The i -th row of the weight matrix $\mathbf{W} \in \mathbb{R}^{|\mathcal{B}| \times d}$ can be interpreted as the learned prototype of the i -th sign. Therefore, we can calculate pair-wise similarities between \mathcal{S}_A and

\mathcal{S}_P by dot-producting their weight matrices as $\mathbf{W}_A \mathbf{W}_P^T \in \mathbb{R}^{|\mathcal{S}_A| \times |\mathcal{S}_P|}$, where the i -th row denotes the similarities between the i -th sign in \mathcal{S}_A and all signs in \mathcal{S}_P . Then we normalize the row vectors via a softmax operation to obtain the mapping from each auxiliary sign to the primary alphabet, which is formulated as:

$$\begin{aligned} \mathcal{M}_{A \rightarrow P}^{\text{weight}}(s_i) &= \operatorname{argmax}(p_{A \rightarrow P}(s_i)), \\ p_{A \rightarrow P}(s_i) &= \operatorname{Softmax}((\mathbf{W}_A \mathbf{W}_P^T)[i, :]), \quad s_i \in \mathcal{S}_A. \end{aligned} \quad (3)$$

Cross-lingual Mapping. Now for each sign $s \in \mathcal{S}_A$, we can compute its mapped sign $\mathcal{M}_{A \rightarrow P}(s) \in \mathcal{S}_P$ by either Eq. 1 or Eq. 3. In this case, all video instances belonging to the same sign share identical mapping. For the variant defined by Eq. 2, we simply map each sign video $v \in \mathcal{C}_A$ at the instance level. This means that video instances belonging to the same sign from the source dataset can be assigned to distinct signs in the target dataset. We visualize some examples in Figure 1 and the supplementary materials.

3.3. Training CSLR on Multilingual Corpus

Now we could map the labels of sign videos in the auxiliary CSLR dataset \mathcal{D}_A from the source sign language to the target sign language. Specifically, for each continuous sign video $v \in \mathcal{D}_A$ with raw labels $\mathbf{s} = (s_1, \dots, s_N)$, we replace each s_i with its mapped label as described in Section 3.2. We use $\mathcal{D}_{A \rightarrow P}$ to denote the yielded dataset. Then a CSLR model is trained on the combination of \mathcal{D}_P and $\mathcal{D}_{A \rightarrow P}$ as shown in Figure 3d. Following previous CSLR methods [10, 47, 45], we use CTC loss [16] to train a TwoStreamSLR network [10]. Given a video v and its corresponding sequence label \mathbf{s} , the CTC loss is formulated as:

$$\mathcal{L}(v, \mathbf{s}) = -\ln \sum_{\pi \in \Pi(\mathbf{s})} p(\pi|v). \quad (4)$$

where π is an alignment path and $\Pi(\mathbf{s})$ is the set of all possible alignments for \mathbf{s} . The conditional probability $p(\pi|v)$ is a product of cross-entropy losses over all time steps.

Our CSLR model is trained by minimizing

$$\mathcal{L}^{\text{CSLR}} = \mathbb{E}_{(v, \mathbf{s}) \sim [\mathcal{D}_P \cup \mathcal{D}_{A \rightarrow P}]_{1:\alpha}} \mathcal{L}(v, \mathbf{s}), \quad (5)$$

where $[\cdot]_{1:\alpha}$ denotes that the training samples in a mini-batch are randomly sampled from \mathcal{D}_P and $\mathcal{D}_{A \rightarrow P}$ with a ratio of $1 : \alpha$. We set $\alpha < 1$ to ensure the primary dataset \mathcal{D}_P dominates the training.

4. Experiments

4.1. Implementation Details

Datasets. Our experiments involve two German sign language (DGS) datasets namely Phoenix-2014 [23] and Phoenix-2014T [6], and a Chinese sign language (CSL)

dataset named CSL-Daily [44]. The Phoenix-2014 dataset is split into train/dev/test with 5672/540/629 videos respectively and contains a total of 1231 signs. The Phoenix-2014T dataset has a split of 7096/519/642 videos and contains 1085 signs. The two DGS datasets share 958 signs in common. The CSL-Daily dataset has a split of 18401/1077/1176 videos and includes 2000 unique signs.

Settings. Since the two DGS datasets are smaller in scale and vocabulary size than the CSL dataset, we choose Phoenix-2014/Phoenix-2014T as the primary dataset \mathcal{D}_P and CSL-Daily as the auxiliary dataset \mathcal{D}_A in our experiments. This ensures that \mathcal{D}_A provides sufficient cross-lingual signs to complement the training source of \mathcal{D}_P . Additionally, we can also use one of the two DGS datasets as the auxiliary dataset for the other. Due to the significant overlap in their vocabularies, we can directly merge their vocabularies without using the cross-lingual mappings described in Section 3.2. In summary, we verify our methodology under four settings (auxiliary dataset \rightarrow primary dataset): 1) CSL-Daily \rightarrow Phoenix-2014T; 2) CSL-Daily \rightarrow Phoenix-2014; 3) CSL-Daily+Phoenix-2014 \rightarrow Phoenix-2014T; 4) CSL-Daily+Phoenix-2014T \rightarrow Phoenix-2014.

Model Architecture. Our CSLR network follows the architecture of TwoStream-SLR [10], which adopts a dual encoder to model RGB videos and keypoint sequences. We conduct ablation studies on SingleStream-SLR with only RGB inputs due to its computational efficiency. More details about the architecture can be found in [10]. We also train separate monolingual TwoStream-SLRs to segment the co-articulated videos for dictionary construction (Section 3.1). Our ISLR model adopts the same network architecture except that an average pooling layer and a classification layer are appended on top of the network.

Training. For the ISLR model, we train it for 100 epochs with a batch size of 32 and a learning rate of $1e-4$. During training, we filter out the tail classes with a frequency threshold of 8 for Phoenix-2014T/Phoenix-2014 and 20 for CSL-Daily. During inference, we forward all samples to compute their cross-lingual predictions (Section 3.2.2). For the CSLR model, we follow the training scheme of TwoStream-SLR [10]—training it for 40 epochs with a batch size of 8 and a learning rate of $1e-3$. We set α as 0.2 in Eq 5. We show more details in the supplementary.

Evaluation. We evaluate our CSLR model on the primary dev/test set using CTC decoding [16] with beam width set to 5. Following prior works [24, 45, 8, 10], we use Word Error Rate (WER) as the evaluation metric, which is a normalized edit distance between the prediction and the reference [16]. Lower WER indicates higher recognition performance. We run each experiment three times with different random seeds and report the score of the best checkpoint.

Method	Dev	Test
CNN-LSTM-HMMs [22]	22.1	24.1
SFL [33]	25.1	26.1
FCN [11]	23.3	25.1
Joint-SLRT [8]	24.6	24.5
SignBT [44]	22.7	23.9
MMTLB [9]	21.9	22.5
SMKD [17]	20.8	22.4
STMC-R [45]	19.6	21.0
C ² SLR [47]	20.2	20.4
TwoStream-SLR [10]	<u>17.7</u>	<u>19.3</u>
CSL-Daily \rightarrow Phoenix-2014T	17.3	18.6
CSL-Daily+Phoenix-2014 \rightarrow Phoenix-2014T	16.9	18.5

Table 1: Comparison with previous work on **Phoenix-2014T** with WER as the evaluation metric. We underline the best results in previous work and bold the best results achieved by our methods.

4.2. Comparison with State-of-the-art Methods

Phoenix-2014T. We evaluate our method on the Phoenix-2014 benchmark and compare with previous work in Table 1. TwoStream-SLR [10] achieves the leading performance among previous works by using a dual encoder to model both RGB and keypoint inputs. We implement our cross-lingual method using their TwoStream architecture and utilize CSL-Daily as an auxiliary dataset to improve the performance of Phoenix-2014T (CSL-Daily \rightarrow Phoenix-2014T). We show that incorporating CSL-Daily improves the performance of its monolingual counterpart (TwoStream-SLR) by 0.4/0.7 WER on the dev/test sets. Further, we employ Phoenix-2014 as another auxiliary dataset besides CSL-Daily (CSL-Daily+Phoenix-2014 \rightarrow Phoenix-2014T). Adding Phoenix-2014 can further improve the performance of Phoenix-2014T, achieving a state-of-the-art with WER of 16.9/18.5 on the dev/test sets.

Phoenix-2014. Similar to the experimental setting on Phoenix-2014T, we evaluate our method on the Phoenix-2014 benchmark and compare its performance with previous work in Table 2. The results on Phoenix-2014 are in line with those on Phoenix-2014T. Leveraging CSL-Daily as an auxiliary (CSL-Daily \rightarrow Phoenix-2014) outperforms monolingual TwoStream-SLR [10] by 0.4/0.3 WER. Adding both Phoenix-2014T and CSL-Daily (CSL-Daily+Phoenix-2014T \rightarrow Phoenix-2014) reduces the WER to 15.7/16.7.

4.3. Analysis and Ablation Study

We first report the performance of our ISLR model. Next, we show the effectiveness of our proposed cross-lingual mapping by comparing it with a monolingual baseline and a multi-task training method. Then we compare the results of using different sign mappings mentioned in Section 3.2. Besides, we study the influence of the number of mapped auxiliary signs and the sampling ratio. To save

Method	Dev	Test
SubUNets [5]	40.8	40.7
IAN [37]	37.1	36.7
ReSign [24]	27.1	26.8
CNN-LSTM-HMMs [22]	26.0	26.0
SFL [33]	24.9	25.3
DNF [15]	23.8	24.4
FCN [11]	23.7	23.9
DNF [15]	23.1	22.9
VAC [31]	21.2	22.3
LCSA [48]	21.4	21.9
CMA [36]	21.3	21.9
SMKD [17]	20.8	21.0
STMC-R [45]	21.1	20.7
C ² SLR [47]	20.5	20.4
TwoStream-SLR [10]	18.4	18.8
CSL-Daily→Phoenix-2014	18.0	18.5
CSL-Daily+Phoenix-2014T→Phoenix-2014	15.7	16.7

Table 2: Comparison with previous work on **Phoenix-2014** with WER as the evaluation metric.

Dataset	Dev				Test			
	Per-instance top-1	Per-instance top-5	Per-class top-1	Per-class top-5	Per-instance top-1	Per-instance top-5	Per-class top-1	Per-class top-5
Phoenix-2014T	81.4	94.3	59.6	83.2	81.7	94.3	60.9	82.4
CSL-Daily	73.6	84.5	73.0	90.4	73.3	85.1	70.5	90.5

Table 3: Performance of our ISLR (isolated sign language recognition) model on Phoenix-2014T and CSL-Daily.

Method	SingleStream		TwoStream	
	Dev	Test	Dev	Test
Phoenix-2014T	21.1	22.4	17.7	19.3
CSL-Daily, Phoenix-2014T	20.7	21.9	17.5	19.1
CSL-Daily→Phoenix-2014T	20.6	21.3	17.3	18.6

Table 4: Using two types of network architectures (SingleStream and TwoStream) proposed by [10], we compare our cross-lingual method with monolingual baseline and multi-task learning approach on Phoenix-2014T with WER as the evaluation metric.

computational costs, we mainly conduct our ablation studies on CSL-Daily→Phoenix-2014T using SingleStream network with RGB inputs. Last, we show the advantage of our method in a data-scarcity scenario when using DGS as the auxiliary dataset and CSL as the primary dataset.

4.3.1 Isolated Sign Language Recognition

We employ a well-trained CSLR model to automatically partition the continuous signs into the isolated signs (Section 3.1), which are then used to train an ISLR model (Section 3.2). Although there is no ground truth for the segmentation, we can evaluate the ISLR model on the unseen dev

Level	Dev	Test	Mapping	Dev	Test
Instance	20.8	21.5	DP	20.7	21.9
Class	20.6	21.3	CLP	20.6	21.3

(a) Instance level vs class level.

(b) DP: Dot-product, CLP: Cross-lingual prediction.

Table 5: Comparison between different cross-lingual mapping strategies by using SingleStream network under CSL-Daily→Phoenix-2014T setting.

and test samples as a proxy. We report the performance of our multilingual ISLR model on Phoenix-2014T and CSL-Daily in Table 3 with per-instance/class top-1/5 accuracy as evaluation metrics. For both datasets, the per-instance top-1 accuracy exceeds 70%, implying that our automatic segmentation is able to spot video clips of the same sign word to build the dictionary, which provides a precondition for our subsequent cross-lingual CSLR training. The per-class accuracy for both datasets is much lower than the per-instance accuracy due to the imbalanced vocabulary distribution, particularly for Phoenix-2014T.

4.3.2 Effectiveness of Cross-lingual Sign Mapping

To demonstrate the advantage of the proposed cross-lingual training methodology, we compare our CSL-Daily→Phoenix-2014T model with two baselines:

- *Phoenix-2014T only.* We train a monolingual CSLR model with the identical network architecture using only Phoenix-2014T.
- *Multi-task learning on CSL-Daily and Phoenix-2014T.* We directly incorporate the original CSL-Daily dataset into the training of Phoenix-2014T without the sign mapping operation and append another CSL-Daily classification layer on top of the CSLR model, forming a multi-task learning framework.

We compare these two baselines with our method using both SingleStream and TwoStream models proposed by [10], with other configurations held the same. As shown in Table 4, multi-task learning (CSL-Daily, Phoenix-2014T) surpasses the monolingual baseline. Our multilingual method (CSL-Daily→Phoenix-2014T) further reduces the WER for both architectures. This suggests that cross-lingual signs identified by our approach indeed enrich the training sources and hence facilitate the recognition capability of the primary dataset.

4.3.3 Sign Mapping Strategies

We describe different sign mapping strategies in Section 3.2, including 1) *class-level* versus *instance-level*; 2) *cross-lingual prediction* versus *dot-product of weight matrices*. Now we experiment with these variants to show which strategy performs the best. We report their results in Table 5.




Figure 4: Effects of reducing the number of mapped signs by varying the threshold on the CSL-Daily→ Phoenix-2014T setting. The bars show the ratio of mapped signs and the solid lines show the WER on Phoenix-2014T dev and test set (lower is better).

First, Table 5a shows that when using cross-lingual prediction, the class-level mapping is superior to the instance-level mapping. This may be because averaging similarity scores over all instances belonging to the same sign could mitigate the noise raised by our imprecise segmentation of isolated sign videos. Second, Table 5b shows that “cross-lingual prediction” outperforms “dot-product of weight matrices” for the class-level mapping. Therefore, we adopt the class-level mapping produced by the cross-lingual prediction as our default sign mapping strategy.

4.3.4 Similarity Threshold

As mentioned in Section 3.2, for a sign mapping $\mathcal{M}_{A \rightarrow P}(s)$, we can use its maximum probability value, *i.e.* the confidence of the cross-lingual prediction, as a proxy to estimate how similar the pair of cross-lingual signs are. The higher the confidence is, the more similar the pair of signs are. This suggests we may need to avoid mapping some auxiliary signs of low confidence. To study this, we set a confidence threshold for cross-lingual mapping. If the mapping confidence is higher than the threshold, we map the auxiliary sign to its cross-lingual sign. Otherwise, we preserve its original label for CSLR training. We reduce the number of mapped signs by raising the threshold from 0 to 0.5 to see how the result changes accordingly, and plot the results in Figure 4. Note that a threshold of 1 preserves all auxiliary signs and becomes the multi-task learning baseline, *i.e.* “CSL-Daily, Phoenix-2014T” in Table 4. We find that a majority of cross-lingual mappings have low confidence scores. However, our cross-lingual training method is insensitive to the change of the confidence threshold — any threshold between 0 and 0.3 outperforms the multi-task learning baseline on the test set. Hence, we map all the auxiliary signs and set the threshold to zero by default.

Ratio (α)	0	0.1	0.2	0.3	0.4	0.6	0.8
Dev	21.1	20.8	20.6	21.0	20.4	21.1	21.3
Test	22.4	21.6	21.3	21.3	21.9	21.7	22.1

Table 6: Ablation on different sampling ratios between the auxiliary dataset and the primary dataset. The experiment is conducted on CSL-Daily→ Phoenix-2014T.

Auxiliary	20% CSL		40% CSL		60% CSL	
	Dev	Test	Dev	Test	Dev	Test
-	45.2	45.5	36.1	35.9	32.7	32.8
Phoenix-2014T	44.0	43.2	35.5	35.7	31.9	31.6
Phoenix-2014	44.0	44.5	35.3	35.2	31.9	32.0

Table 7: We compare the baseline without using auxiliary data, with Phoenix-2014/2014T→20/40/60% CSL-Daily.

4.3.5 Ratio of Auxiliary Dataset

Here we study the effect of the sampling ratio between \mathcal{D}_A and \mathcal{D}_P , *i.e.* the coefficient α in Eq 5. We train our CSLR models with various α and compare their results in Table 6. Almost all models involving cross-lingual training surpass the monolingual baseline ($\alpha = 0$).

4.3.6 Alleviating data-scarcity

While CSL→ DGS improves over training on DGS alone, we do not observe improvements in experiments of DGS→ CSL. We suspect this is because the whole CSL dataset is larger in scale than the two DGS datasets and benefits little from auxiliary DGS samples. To prove the effectiveness of our method particularly when the primary dataset is scarce, we sample a subset of CSL-Daily as the primary dataset, use a DGS dataset as the auxiliary dataset and conduct the DGS→ CSL experiments. As shown in Table 7, leveraging either Phoenix-2014T or Phoenix-2014 has an advantage over training on the CSL subset alone, bringing the largest gain when only 20% of CSL is available.

5. Conclusion

We present a novel approach to improve the monolingual performance of continuous sign language recognition (CSLR) by leveraging multilingual corpora and identifying visually similar signs across different sign languages, known as cross-lingual signs. Our method begins by constructing isolated sign dictionaries from the CSLR datasets. Next, we train a multilingual isolated sign recognition classifier on the two dictionaries and use it to identify the cross-lingual sign-to-sign mappings. Finally, we train our CSLR model on both the primary dataset and the remapped auxiliary dataset. By addressing the data scarcity issue, our approach achieves state-of-the-art results on two CSLR benchmarks. We are the first to demonstrate the effectiveness of cross-lingual transfer in CSLR and hope that our work will offer valuable insights for future research.

References

- [1] Samuel Albanie, Gül Varol, Liliane Momeni, Hannah Bull, Triantafyllos Afouras, Himel Chowdhury, Neil Fox, Bencie Woll, Rob Cooper, Andrew McParland, and Andrew Zisserman. BOBSL: BBC-Oxford British Sign Language Dataset. 2021.
- [2] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al. Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning*, pages 173–182. PMLR, 2016.
- [3] Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber. Common voice: A massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670*, 2019.
- [4] Donald J Berndt and James Clifford. Using dynamic time warping to find patterns in time series. In *KDD workshop*, volume 10, pages 359–370. Seattle, WA, USA., 1994.
- [5] Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, and Richard Bowden. Subunets: End-to-end hand shape and continuous sign language recognition. In *2017 IEEE International Conference on Computer Vision*, 2017.
- [6] Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, Hermann Ney, and Richard Bowden. Neural sign language translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [7] Necati Cihan Camgöz, Oscar Koller, Simon Hadfield, and Richard Bowden. Multi-channel transformers for multi-articulatory sign language translation. In *European Conference on Computer Vision 2020 Workshops*, 2020.
- [8] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. Sign language transformers: Joint end-to-end sign language recognition and translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020.
- [9] Yutong Chen, Fangyun Wei, Xiao Sun, Zhirong Wu, and Stephen Lin. A simple multi-modality transfer learning baseline for sign language translation. *arXiv preprint arXiv:2203.04287*, 2022.
- [10] Yutong Chen, Ronglai Zuo, Fangyun Wei, Yu Wu, Shujie Liu, and Brian Mak. Two-stream network for sign language recognition and translation, 2022.
- [11] Ka Leong Cheng, Zhaoyang Yang, Qifeng Chen, and Yu-Wing Tai. Fully convolutional networks for continuous sign language recognition. In *European Conference on Computer Vision*. Springer International Publishing, 2020.
- [12] Yiting Cheng, Fangyun Wei, Jianmin Bao, Dong Chen, and Wenqiang Zhang. Cico: Domain-aware sign language retrieval via cross-lingual contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19016–19026, 2023.
- [13] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*, 2019.
- [14] Runpeng Cui, Hu Liu, and Changshui Zhang. Recurrent convolutional neural networks for continuous sign language recognition by staged optimization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [15] Runpeng Cui, Hu Liu, and Changshui Zhang. A deep neural framework for continuous sign language recognition by iterative training. *IEEE Transactions on Multimedia*, 21(7):1880–1891, 2019.
- [16] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376, 2006.
- [17] Aiming Hao, Yuecong Min, and Xilin Chen. Self-mutual distillation learning for continuous sign language recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11303–11312, October 2021.
- [18] Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351, 2017.
- [19] Hamid Reza Vaezi Joze and Oscar Koller. MS-ASL: A large-scale data set and benchmark for understanding american sign language. In *30th British Machine Vision Conference, BMVC*, 2019.
- [20] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P. E. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux. Libri-light: A benchmark for asr with limited or no supervision. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7669–7673, 2020. <https://github.com/facebookresearch/libri-light>.
- [21] Anjali Kannan, Arindrima Datta, Tara N Sainath, Eugene Weinstein, Bhuvana Ramabhadran, Yonghui Wu, Ankur Bapna, Zhifeng Chen, and Seungji Lee. Large-scale multilingual speech recognition with a streaming end-to-end model. *arXiv preprint arXiv:1909.05330*, 2019.
- [22] Oscar Koller, Necati Cihan Camgoz, Hermann Ney, and Richard Bowden. 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*, 2019.
- [23] Oscar Koller, Jens Forster, and Hermann Ney. Continuous sign language recognition: Towards large vocabulary statistical recognition systems handling multiple signers. *Computer Vision and Image Understanding*, 141:108–125, Dec. 2015.
- [24] Oscar Koller, Sepehr Zargaran, and Hermann Ney. Re-sign: Re-aligned end-to-end sequence modelling with deep recurrent CNN-HMMs. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [25] Bo Li, Yu Zhang, Tara Sainath, Yonghui Wu, and William Chan. Bytes are all you need: End-to-end multilingual

- speech recognition and synthesis with bytes. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5621–5625. IEEE, 2019.
- [26] Dongxu Li, Cristian Rodriguez, Xin Yu, and Hongdong Li. Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In *The IEEE Winter Conference on Applications of Computer Vision*, 2020.
- [27] Dongxu Li, Xin Yu, Chenchen Xu, Lars Petersson, and Hongdong Li. Transferring cross-domain knowledge for video sign language recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [28] Jinyu Li et al. Recent advances in end-to-end automatic speech recognition. *APSIPA Transactions on Signal and Information Processing*, 11(1), 2022.
- [29] Zekang Liu Lianyu Hu, Liqing Gao and Wei Feng. Temporal lift pooling for continuous sign language recognition. In *European conference on computer vision*. Springer, 2022.
- [30] Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 2020.
- [31] Yuecong Min, Aiming Hao, Xiujuan Chai, and Xilin Chen. Visual alignment constraint for continuous sign language recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11542–11551, October 2021.
- [32] Gokul NC, Manideep Ladi, Sumit Negi, Prem Selvaraj, Pratyush Kumar, and Mitesh M Khapra. Addressing resource scarcity across sign languages with multilingual pretraining and unified-vocabulary datasets. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022.
- [33] Zhe Niu and Brian Mak. Stochastic fine-grained labeling of multi-state sign glosses for continuous sign language recognition. In *European Conference on Computer Vision*, pages 172–186. Springer, 2020.
- [34] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5206–5210. IEEE, 2015.
- [35] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. SpecAugment: A simple data augmentation method for automatic speech recognition. *arXiv preprint arXiv:1904.08779*, 2019.
- [36] Junfu Pu, Wengang Zhou, Hezhen Hu, and Houqiang Li. *Boosting Continuous Sign Language Recognition via Cross Modality Augmentation*, page 1497–1505. Association for Computing Machinery, New York, NY, USA, 2020.
- [37] Junfu Pu, Wengang Zhou, and Houqiang Li. Iterative alignment network for continuous sign language recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2019.
- [38] Lawrence Rabiner and Biing-Hwang Juang. *Fundamentals of speech recognition*. Prentice-Hall, Inc., 1993.
- [39] Sandrine Tornay, Marzieh Razavi, and Mathew Magimai-Doss. Towards multilingual sign language recognition. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6309–6313, 2020.
- [40] Shubham Toshniwal, Tara N Sainath, Ron J Weiss, Bo Li, Pedro Moreno, Eugene Weinstein, and Kanishka Rao. Multilingual speech recognition with a single end-to-end model. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 4904–4908. IEEE, 2018.
- [41] Shinji Watanabe, Takaaki Hori, Shigeaki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, et al. Espnet: End-to-end speech processing toolkit. *arXiv preprint arXiv:1804.00015*, 2018.
- [42] Aoxiong Yin, Zhou Zhao, Weike Jin, Meng Zhang, Xingshan Zeng, and Xiaofei He. Mslt: Towards multilingual sign language translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5109–5119, June 2022.
- [43] Dong Yu and Li Deng. *Automatic speech recognition*, volume 1. Springer, 2016.
- [44] Hao Zhou, Wengang Zhou, Weizhen Qi, Junfu Pu, and Houqiang Li. Improving sign language translation with monolingual data by sign back-translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [45] Hao Zhou, Wengang Zhou, Yun Zhou, and Houqiang Li. Spatial-temporal multi-cue network for sign language recognition and translation. *IEEE Transactions on Multimedia*, 2021.
- [46] Hao Zhou, Wengang Zhou, Yun Zhou, and Houqiang Li. Spatial-temporal multi-cue network for sign language recognition and translation. *IEEE Transactions on Multimedia*, 2022.
- [47] Ronglai Zuo and Brian Mak. C2SLR: Consistency-enhanced continuous sign language recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5131–5140, 2022.
- [48] Ronglai Zuo and Brian Mak. Local context-aware self-attention for continuous sign language recognition. In *Proc. Interspeech 2022*, pages 4810–4814, 2022.
- [49] Ronglai Zuo, Fangyun Wei, and Brian Mak. Natural language-assisted sign language recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14890–14900, 2023.