

# CiCo: Domain-Aware Sign Language Retrieval via Cross-Lingual Contrastive Learning

Yiting Cheng<sup>1</sup>, Fangyun Wei<sup>2\*</sup>, Jianmin Bao<sup>2</sup>, Dong Chen<sup>2</sup>, Wenqiang Zhang<sup>1\*</sup>

<sup>1</sup>School of Computer Science, Fudan University, <sup>2</sup>Microsoft Research Asia

{ytcheng18, wqzhang}@fudan.edu.cn, {fawe, jianbao, doch}@microsoft.com

## Abstract

This work focuses on sign language retrieval—a recently proposed task for sign language understanding. Sign language retrieval consists of two sub-tasks: text-to-sign-video (T2V) retrieval and sign-video-to-text (V2T) retrieval. Different from traditional video-text retrieval, sign language videos, not only contain visual signals but also carry abundant semantic meanings by themselves due to the fact that sign languages are also natural languages. Considering this character, we formulate sign language retrieval as a cross-lingual retrieval problem as well as a video-text retrieval task. Concretely, we take into account the linguistic properties of both sign languages and natural languages, and simultaneously identify the fine-grained cross-lingual (i.e., sign-to-word) mappings while contrasting the texts and the sign videos in a joint embedding space. This process is termed as cross-lingual contrastive learning. Another challenge is raised by the data scarcity issue—sign language datasets are orders of magnitude smaller in scale than that of speech recognition. We alleviate this issue by adopting a domain-agnostic sign encoder pre-trained on large-scale sign videos into the target domain via pseudo-labeling. Our framework, termed as domain-aware sign language retrieval via Cross-lingual Contrastive Learning or CiCo for short, outperforms the pioneering method by large margins on various datasets, e.g., +22.4 T2V and +28.0 V2T R@1 improvements on How2Sign dataset, and +13.7 T2V and +17.1 V2T R@1 improvements on PHOENIX-2014T dataset. Code and models are available at: <https://github.com/FangyunWei/SLRT>.

## 1. Introduction

Sign languages are the primary means of communication used by people who are deaf or hard of hearing. Sign language understanding [1, 10, 12–15, 18, 32, 33, 61, 73] is significant for overcoming the communication barrier between

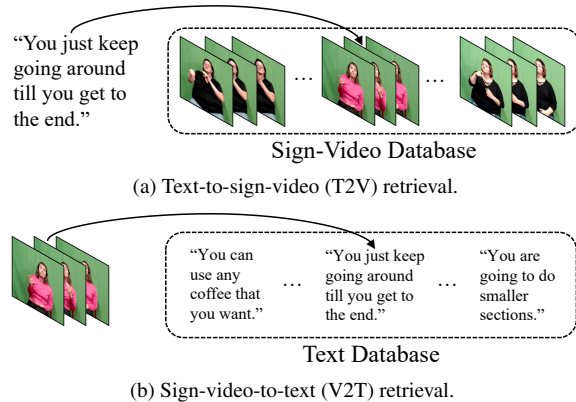


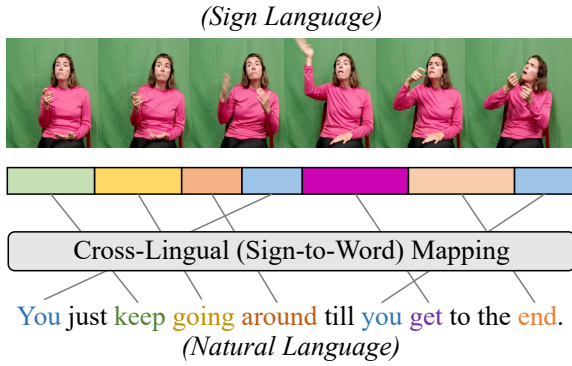
Figure 1. Illustration of: (a) T2V retrieval; (b) V2T retrieval.

the hard-of-hearing and non-signers. Sign language recognition and translation (SLRT) has been extensively studied, with the goal of recognizing the *arbitrary* semantic meanings conveyed by sign languages. However, the lack of available data significantly limits the capability of SLRT. In this paper, we focus on developing a framework for a recently proposed sign language retrieval task [18]. Unlike SLRT, sign language retrieval focuses on retrieving the meanings that signers express from a *closed-set*, which can significantly reduce error rates in realistic deployment.

Sign language retrieval is both similar to and distinct from the traditional video-text retrieval. On the one hand, like video-text retrieval, sign language retrieval is also composed of two sub-tasks, i.e., text-to-sign-video (T2V) retrieval and sign-video-to-text (V2T) retrieval. Given a free-form written query and a large collection of sign language videos, the objective of T2V is to find the video that best matches the written query (Figure 1a). In contrast, the goal of V2T is to identify the most relevant text description given a query of sign language video (Figure 1b).

On the other hand, different from the video-text retrieval, sign languages, like most natural languages, have their own grammars and linguistic properties. Therefore, sign language videos not only contain visual signals, but also carry

\*Corresponding author.



(a) While contrasting the sign videos and the texts in a joint embedding space, we simultaneously identify the fine-grained cross-lingual (sign-to-word) mappings of sign languages and natural languages via the proposed cross-lingual contrastive learning. Existing datasets do not annotate the sign-to-word mappings.



(b) We show four instances of the sign “Book” in How2Sign [19] dataset, which are identified by our approach. Please refer to the supplementary material for more examples.

Figure 2. Illustration of: (a) cross-lingual (sign-to-word) mapping; (b) sign-to-word mappings identified by our CiCo.

semantics (*i.e.*, sign<sup>1</sup>-to-word mappings between sign languages and natural languages) by themselves, which differentiates them from the general videos that merely contain visual information. Considering the linguistic characteristics of sign languages, we formulate sign language retrieval as a cross-lingual retrieval [6, 34, 59] problem in addition to a video-text retrieval [5, 24, 42–44, 57, 68] task.

Sign language retrieval is extremely challenging due to the following reasons: (1) Sign languages are completely separate and distinct from natural languages since they have unique linguistic rules, word formation, and word order. The transcription between sign languages and natural languages is complicated, for instance, the word order is typically not preserved between sign languages and natural languages. It is necessary to automatically identify the sign-to-word mapping from the cross-lingual retrieval perspective; (2) In contrast to the text-video retrieval datasets [45, 48] which contain millions of training samples, sign language datasets are orders of magnitude smaller in scale—for example, there are only 30K video-text pairs in How2Sign [19] training set; (3) Sign languages convey information through the handshape, facial expression, and body movement, which requires models to distinguish fine-

<sup>1</sup>We use sign to denote lexical item within a sign language vocabulary.

grained gestures and actions; (4) Sign language videos typically contain hundreds of frames. It is necessary to build efficient algorithms to lower the training cost and fit the long videos as well as the intermediate representations into limited GPU memory.

In this work, we concentrate on resolving the challenges listed above:

- We consider the linguistic rules (*e.g.*, word order) of both sign languages and natural languages. We formulate sign language retrieval as a cross-lingual retrieval task as well as a video-text retrieval problem. While contrasting the sign videos and the texts in a joint embedding space as achieved in most vision-language pre-training frameworks [5, 44, 57], we simultaneously identify the fine-grained cross-lingual (sign-to-word) mappings between two types of languages via our proposed cross-lingual contrastive learning as shown in Figure 2.
- Data scarcity typically brings in the over-fitting issue. To alleviate this issue, we adopt transfer learning and adapt a recently released domain-agnostic sign encoder [61] pre-trained on large-scale sign-videos to the target domain. Although this encoder is capable of distinguishing the fine-grained signs, direct transferring may be sub-optimal due to the unavoidable domain gap between the pre-training dataset and sign language retrieval datasets. To tackle this problem, we further fine-tune a domain-aware sign encoder on pseudo-labeled data from target datasets. The final sign encoder is composed of the well-optimized domain-aware sign encoder and the powerful domain-agnostic sign encoder.
- In order to effectively model long videos, we decouple our framework into two disjoint parts: (1) a sign encoder which adopts a sliding window on sign-videos to pre-extract their vision features; (2) a cross-lingual contrastive learning module which encodes the extracted vision features and their corresponding texts in a joint embedding space.

Our framework, called domain-aware sign language retrieval via **Cross-lingual Contrastive learning** or CiCo for short, outperforms the pioneer SPOT-ALIGN [18] by large margins on various datasets, achieving 56.6 (+22.4) T2V and 51.6 (+28.0) V2T R@1 accuracy (improvement) on How2Sign [19] dataset, and 69.5 (+13.7) T2V and 70.2 (+17.1) V2T R@1 accuracy (improvement) on PHOENIX-2014T [8] dataset. With its simplicity and strong performance, we hope our approach can serve as a solid baseline for future research.

## 2. Related Work

**Sign Language Understanding.** Sign language understanding aims at interpreting the semantic information con-

veyed within sign videos. Researchers have explored such capability on various tasks including sign language recognition [13, 20, 21, 30, 38, 58, 73], sign spotting [1, 18, 47, 61], sign language translation [8, 9, 12, 37, 72] and our focused sign language retrieval [18].

One of the fundamental tasks of sign language understanding is sign language recognition (SLR), which aims to transcribe a sign video into a gloss sequence. Previous works on SLR focus on designing carefully engineered features [20, 21, 58] or modeling temporal dependencies [55, 56]. Recently, the success of 3D convolutional neural networks in action related tasks [41, 60, 63] is transferred to SLR. In particular, the I3D [11] architecture has proven to be effective for this task [1, 30, 36, 38, 61]. In this work, we also adopt this network architecture in our sign encoder.

Sign spotting is a particular variant of sign language recognition. It aims to localize all instances of a given sign within an untrimmed video. Recent works tackle this task with auxiliary cue of subtitles, introducing automatic annotation systems by using mouthing [1], dictionaries [47] and attention maps of Transformer [61]. SPOT-ALIGN [18] extends existing spotting methods [1, 47] with an iterative training schema, which alternates between repeated sign spotting and model fine-tuning. In this work, pseudo-labeling is served as our sign spotting approach to localize isolated signs in untrimmed videos from target sign language retrieval datasets. Compared with above efforts, our approach only employs a pre-trained sign encoder without utilizing additional auxiliary cue, which is proved to be simple yet efficient.

Early works of sign language retrieval primarily investigate query-by-example searching [4, 70], which queries individual instances with given sign examples. Our work focuses on free-form textual retrieval—a recently introduced task by SPOT-ALIGN [18]. It symbolizes the real-world scenario of searching sign language videos with natural languages. The pioneer SPOT-ALIGN [18] purely formulates sign language retrieval as a video-text retrieval task, where the cross-modal alignment is modeled upon overall global embeddings of sign videos and texts. However, the linguistic properties of sign languages are ignored. In contrast, we formulate sign language retrieval as a joint task of text-video retrieval and cross-lingual retrieval.

**Vision-Language Models.** Learning general-purpose representations for visual and textual modalities is a long-standing topic [7, 35]. The idea has been investigated decades ago [49]. Recently, the prominent success of CLIP [51], ALIGN [28] and ALBEF [39] has demonstrated the capability of learning joint cross-modal representations with large-scale web data using simple image-text contrastive learning. Similar idea has also been explored in video-text retrieval, which is the closest area of our work.

The common practice in image-text retrieval [23, 25, 29, 51] and video-text retrieval [5, 24, 42–44, 57, 68] is to encode the images/videos and texts into the overall representations. It is reasonable since the visual signals of typical image-text and video-text retrieval datasets mainly describe the certain objects or describable events. In contrast, sign videos, as the carriers of sign languages, convey abundant semantics by themselves. There exist fine-grained mappings between sign videos and natural languages. We find that identifying such cross-lingual (sign-to-word) mappings significantly boosts the performance of sign language retrieval.


**Cross-Lingual Information Retrieval.** In this work, we also formulate sign-language retrieval as a cross-lingual task. In natural language processing community, cross-lingual information retrieval (CLIR) [34, 59, 65] refers to the task of retrieving documents between different languages. One of the most common approaches of CLIR is to learn sentence-level embedding alignment by mapping pre-acquired monolingual embeddings [46, 50, 69] of different languages into a shared space. Recently, researchers have exploited fine-grained word-level mappings with self-training [2, 3, 59]. They find that mappings can be learned by initiating with a seed dictionary and alternating between alignment modeling and dictionary mining. Our work also exploits fine-grained word-level mappings, called sign-to-word mappings, in the context of sign language retrieval. Identifying sign-to-word mappings is challenging since sign languages are expressed in visual modality. The proposed cross-lingual contrastive learning tackles this challenge by exploiting the fine-grained cross-modal interactions.

### 3. Methodology

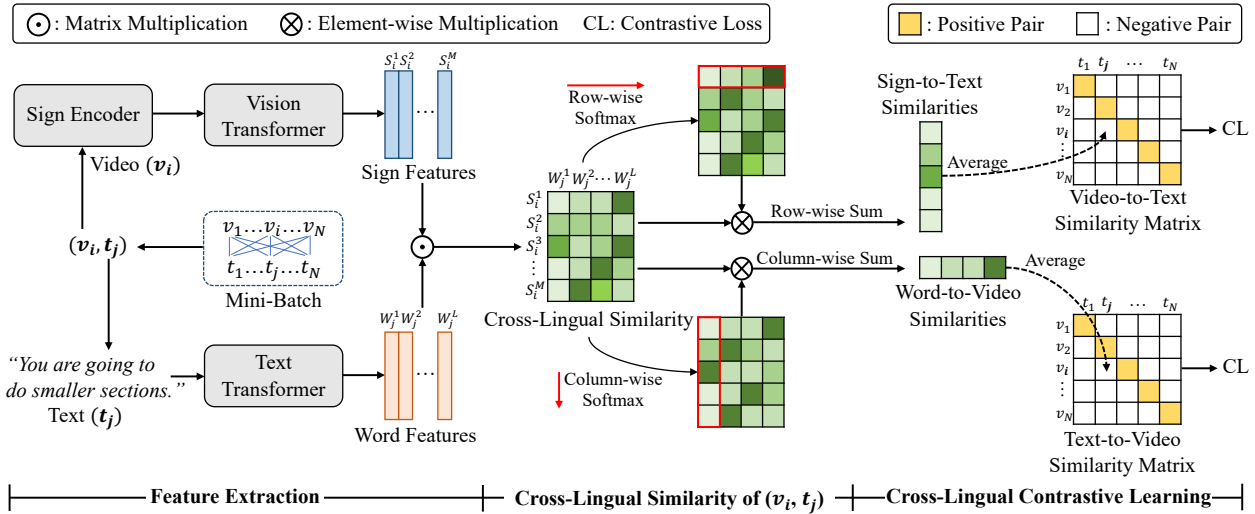
In this section, we first formulate the task of sign language retrieval in Section 3.1. Next, we introduce our framework which is composed of two parts: 1) a sign encoder which extracts discriminative and domain-aligned features of sign videos (Section 3.2); 2) a cross-lingual contrastive learning framework, which contrasts sign-video-text pairs while concurrently identifying the fine-grained sign-to-word mappings (Section 3.3). At last, we explore text augmentations in Section 3.4.

#### 3.1. Task Formulation

Let  $\mathcal{V}$  and  $\mathcal{T}$  denote a set of sign videos and their corresponding texts (transcriptions), respectively. Sign language retrieval consists of two tasks namely text-to-sign-video retrieval (T2V), and sign-video-to-text retrieval (V2T), respectively. The objective of T2V is to find the sign video  $v \in \mathcal{V}$  whose signing content best matches the text query. In contrast, the reverse task V2T requires the model to identify the most relevant text (transcription)  $t \in \mathcal{T}$  given a query of sign video. We resolve sign language retrieval by learning a joint embedding space of sign videos and texts.



(a) Illustration of sign encoder.



(b) Illustration of cross-lingual contrastive learning.

Figure 3. Overview of our framework. (a) Sign encoder is composed of a powerful domain-agnostic sign encoder pre-trained on large-scale sign videos, and a domain-aware sign encoder fine-tuned on pseudo-labeled data from target datasets. We adopt a sliding window manner to extract a discriminative and domain-aligned feature per clip. (b) Cross-lingual contrastive learning takes  $N$  sign-video-text pairs as inputs and contrasts paired data in a shared embedding space while implicitly identifying the fine-grained sign-to-word mappings during training.

### 3.2. Sign Encoder

**Process Sign Videos with Sliding Window.** Sign videos from sign language retrieval datasets typically contain hundreds of frames. To efficiently train our model and lower the usage of GPU memory, given a sign video, we adopt a sliding window manner with stride of 1 and window size of 16 to produce  $M$  temporally overlapped clips. Next, we separately feed each clip into a sign encoder to extract its feature. The final sign-video feature is yielded by stacking features coming out of  $M$  clips along temporal dimension. A powerful sign encoder is crucial.

**Overview of Sign Encoder.** Recent advances in sign spotting [47, 61] greatly facilitate the collection of large-scale

sign language datasets, enabling powerful representation learning abilities of convolutional neural networks on the sign classification task. Previous methods [16, 18] have demonstrated the feasibility of transferring a sign encoder pre-trained on large-scale sign-spotting data into downstream tasks. We follow this practice and use an I3D network pre-trained on BSL-1K [61], a sign classification dataset collected via sign spotting, as our primary sign encoder. Due to its favorable transfer performance, we term this model as a domain-agnostic sign encoder. Nevertheless, the domain gap between BSL-1K and sign language retrieval datasets is non-negligible. To tackle this problem, we further fine-tune a domain-aware sign encoder, which has an identical architecture to the domain-agnostic sign en-



coder, on target datasets through pseudo-labeling. The final sign encoder is composed of the well-optimized domain-aware sign encoder and the powerful domain-agnostic sign encoder, as illustrated in Figure 3a.

**Pseudo-Labeling on Target Datasets.** Now we describe the details of pseudo-labeling. Given a sign video from a target dataset, we adopt a sliding window with the stride of 1 and the window size of 16 to generate a set of temporally overlapped clips. For each clip, we first utilize the pre-trained domain-agnostic sign encoder to produce its prediction. Then we binarize the prediction with a pre-defined threshold  $\lambda$  to generate the corresponding pseudo label. The invalid samples, whose maximum score is lower than  $\lambda$ , are filtered. We repeat the above process for all sign videos and eventually build a pseudo-labeled set. Our domain-aware sign encoder, which is initialized by the domain-agnostic sign encoder, is fine-tuned on the pseudo-labeled set via a standard cross-entropy loss for classification training.

**Feature Extraction with Sign Encoder.** So far, we acquire a domain-aware sign encoder approximately aligned in target domain. Nevertheless, its capability is restricted by the unavoidable noises in pseudo-labels and the limited amount of pseudo-labeled samples. Recall that we already have a powerful domain-agnostic sign encoder pre-trained on large-scale dataset in hand, inspiring us to make use of both domain-agnostic sign encoder  $h_\xi(\cdot)$  and domain-aware sign encoder  $h_\theta(\cdot)$  to extract discriminative and domain-aligned features. Our final sign encoder  $H(\cdot)$ , as shown in Figure 3a, is a weighted combination of  $h_\xi(\cdot)$  and  $h_\theta(\cdot)$  with a trade-off hyper-parameter  $\alpha$ . As described above,  $H(\cdot)$  encodes sign videos in a sliding window manner. For simplicity, we use  $H(v)$  to denote feature extraction on sign video  $v$ , which is formulated as:

$$H(v) = \alpha h_\xi(v) + (1 - \alpha)h_\theta(v). \quad (1)$$

### 3.3. Cross-Lingual Contrastive Learning

The objective of cross-lingual contrastive learning (CLCL) is to learn a joint embedding space of sign videos and texts while concurrently identifying the fine-grained sign-to-word mappings during training. An overview is shown in Figure 3b. CLCL takes a mini-batch  $\{(v_n, t_n)\}_{n=1}^N$  containing  $N$  sign-video-text pairs as input, and contrasts paired data in a shared embedding space for sign language retrieval.

**Sign Features and Word Features.** Given a sign video  $v \in \{v_n\}_{n=1}^N$ , we first adopt our sign encoder  $H(\cdot)$  described in Section 3.2 to extract its intermediate feature. Note  $H(\cdot)$  encodes sign videos in a sliding window manner, and thus there are no interactions among different clips. To facilitate information exchange, we further append a 12-layer Transformer [62]  $F(\cdot)$  onto  $H(\cdot)$  to extract sign features  $\mathbf{S}$  of sign video  $v$ , which is formulated

as  $\mathbf{S} = F(H(v)) \in \mathbb{R}^{M \times D}$ , where  $M$  denotes the number of clips, and  $D$  is the hidden dimension. Given a text  $t \in \{t_n\}_{n=1}^N$ , we convert  $t$  into a lower-cased byte pair encoding (BPE) representation [53], which is subsequently fed into another 12-layer Transformer  $G(\cdot)$  to generate the word features  $\mathbf{W} = G(t) \in \mathbb{R}^{L \times D}$ , where  $L$  represents word number.

Since CLIP [51] shows excellent transfer capability in various downstream tasks [17, 22, 27, 54, 66, 67], we initialize  $F(\cdot)$  and  $G(\cdot)$  with CLIP’s image encoder (ViT-B) and text encoder, respectively, to ease the learning. Though CLIP’s vision encoder takes image patches as inputs, we experimentally find that it generalizes well in our scenario where input data is in a different modality.

**Cross-Lingual Similarity.** There exist inherent sign-to-word mappings between sign languages and natural languages. To incorporate this prior knowledge into learning, we introduce cross-lingual similarity—an indicator to identify sign-to-word mappings between  $i$ -th sign video  $v_i$  and  $j$ -th text  $t_j$ . Concretely, given sign features  $\mathbf{S}_i \in \mathbb{R}^{M \times D}$  of  $v_i$ , and word features  $\mathbf{W}_j \in \mathbb{R}^{L \times D}$  of  $t_j$ , we calculate a cross-lingual similarity matrix  $\mathbf{E}_{(i,j)} = \mathbf{S}_i \cdot \mathbf{W}_j^T \in \mathbb{R}^{M \times L}$ . Each element in  $\mathbf{E}_{(i,j)}$  represents the similarity of a sign clip in  $v_i$  and a word in  $t_j$ .

**Cross-Lingual Contrastive Learning.** Directly apply supervisions on token-wise similarity matrix  $\mathbf{E}_{(i,j)}$  is infeasible due to the absence of fine-grained sign-to-word annotations. Inspired by the recent progress of vision-language contrastive learning [23, 25, 29, 39, 40, 51], we turn to contrast the global representations of sign videos and texts. The underlying idea is to calculate a global similarity  $z$  of  $v_i$  and  $t_j$  based on  $\mathbf{E}_{(i,j)} \in \mathbb{R}^{M \times L}$ .

**Sign-Video-to-Text Contrast.** We first introduce sign-video-to-text contrast as shown in Figure 3b. To be specific, we utilize a Softmax operation to each row of  $\mathbf{E}_{(i,j)}$ , and multiply the resulting matrix with  $\mathbf{E}_{(i,j)}$  to generate a sign-to-word similarity matrix  $\mathbf{E}'_{(i,j)} \in \mathbb{R}^{M \times L}$ , where each row represents the re-weighted similarities between a sign clip in  $v_i$  and all words in  $t_j$ . After that, we adopt a row-wise addition operation on  $\mathbf{E}'_{(i,j)}$  to yield the sign-to-text similarity vector  $\mathbf{e}_{(i,j)} \in \mathbb{R}^M$ , where each element denotes a similarity of a sign clip in  $v_i$  and whole text  $t_j$ . At last, we average all elements in  $\mathbf{e}_{(i,j)}$  to produce the global similarity  $z$  of sign video  $v_i$  and text  $t_j$ .

In the same way, we can calculate the similarities for both positive pairs  $\{(v_i, t_i)\}_{i=1}^N$  and negative pairs  $\{(v_i, t_j)\}_{i=1, j=1, i \neq j}^N$  in a mini-batch, yielding a video-to-text similarity matrix  $\mathbf{Z}_{V2T} \in \mathbb{R}^{N \times N}$ , where  $\mathbf{Z}_{V2T}^{(i,j)}$  denotes the global similarity of  $v_i$  and  $t_j$ . Following CLIP [51], we adopt InfoNCE loss [26] to pull the embeddings of matched image-text pairs together while pushing those of non-matched pairs apart, which is formulated as

follows:

$$\begin{aligned} \mathcal{L}_{V2T} = & -\frac{1}{2N} \sum_{i=1}^N \log \frac{\exp(\mathbf{Z}_{V2T}^{(i,j)}/\tau)}{\sum_{j=1}^N \exp(\mathbf{Z}_{V2T}^{(i,j)}/\tau)} \\ & -\frac{1}{2N} \sum_{j=1}^N \log \frac{\exp(\mathbf{Z}_{V2T}^{(i,j)}/\tau)}{\sum_{i=1}^N \exp(\mathbf{Z}_{V2T}^{(i,j)}/\tau)}, \end{aligned} \quad (2)$$

where  $\tau$  is a trainable temperature parameter.

**Text-to-Sign-Video Contrast.** Up to now, we have introduced sign-video-to-text contrast. A symmetrical version, termed text-to-sign-video contrast, shares the similar spirit as shown in Figure 3b. The implementation of text-to-sign-video contrast is extremely simple: we replace the row-wise operations (*i.e.*, Softmax and addition) in sign-video-to-text contrast with column-wise ones and keep the remaining processes unchanged. We use  $\mathcal{L}_{T2V}$  to denote the loss function of text-to-sign-video contrast. In our implementation, we reuse the loss defined in Eq 2 but substitute the input with the text-to-video similarity matrix  $\mathbf{Z}_{T2V}$ .

**Loss Function.** The overall loss for cross-lingual contrastive learning is a weighted sum of  $\mathcal{L}_{V2T}$  and  $\mathcal{L}_{T2V}$  with a trade-off hyper-parameter  $\beta$ :

$$\mathcal{L} = \beta \mathcal{L}_{V2T} + (1 - \beta) \mathcal{L}_{T2V}. \quad (3)$$

### 3.4. Text Augmentation

Considering that the datasets of sign language retrieval are typically small-scale, we explore text augmentations to improve the generalization of our approach. EDA [64] introduces three simple yet efficient data augmentations in text classification task: random delete randomly removes words in a sentence; synonym replacement randomly selects words from a sentence that are not stop words and replaces them with synonyms; random swap randomly chooses two words in a sentence and swaps their positions. The first two augmentations have been proven effective in text classification task. However, we experimentally find that our focused sign language retrieval is sensitive to random delete and synonym replacement augmentations. To guarantee that the augmented texts preserve the original semantic meanings, we only adopt the random swap augmentation in our approach. We suppose there are two reasons: 1) the word order of sign languages and natural languages are constitutionally distinct, and reordering does not affect semantic meanings; 2) the proposed cross-lingual contrastive learning is insensitive to word order.

## 4. Experiment

### 4.1. Datasets and Implementation Details

**Datasets.** We primarily focus on *How2Sign* [19] dataset. Our model is also evaluated on *PHOENIX-2014T* [8] and

*CSL-Daily* [71], which are primarily used for sign language recognition and translation in previous works.

*How2Sign* is a large-scale continuous American sign language (ASL) dataset consisting of a parallel corpus of about 80 hours of sign videos with subtitle annotations. It covers a wide range of instructional videos corresponding to various categories. There are 31164, 1740 and 2356 sign-video-text pairs in training, validation and test sets, respectively. Following SPOT-ALIGN [18], we remove the invalid pairs where the subtitle alignment is detected to exceed the video duration, remaining 31085, 1739 and 2348 available pairs in training, validation and test sets, respectively. The resolution of sign videos is 1280×720, we crop the human bodies of signers with Faster R-CNN [52] to generate valid videos.

*PHOENIX-2014T* is a German sign language (Deutsche Gebärdensprache, DGS) dataset collected in the domain of weather forecast from TV broadcast, consisting of 7096, 519 and 642 video-text pairs in training, validation and test sets, respectively.

*CSL-Daily* is a recently released Chinese sign language (CSL) dataset. The topic of CSL-Daily revolves around people’s daily lives, including 18401, 1077, 1176 parallel samples in training, validation and test set, respectively.

**Evaluation Metric.** Following previous works [18, 43, 51, 57], retrieval performance is evaluated by recall at rank K (R@K, higher is better) and median rank (MedR, lower is better). We evaluate our approach on both text-to-sign-video (T2V) retrieval and sign-video-to-text (V2T) retrieval tasks. We report R@1, R@5, and R@10 in all experiments, and additionally report MedR when comparing with state-of-the-art approaches.

**Implementation Details.** The sign encoder takes videos of resolution of 256×256 as input. The domain-agnostic sign encoder is a I3D [11] network pre-trained on BSL-1K [61]. In pseudo label generation, we set the threshold  $\lambda$  to 0.6 to filter samples with low-confidence. Non-maximum suppression (NMS) with a temporal window of 24 frames is utilized to remove the duplicates among the pseudo-labeled samples. A collection of approximate 64K pseudo-labeled samples covering a vocabulary of 1220 words is eventually generated. The domain-aware sign encoder is initialized with the domain-agnostic one and fine-tuned with a learning rate of  $1 \times 10^{-2}$  and batch size of 4 for 15 epochs. In the training of cross-lingual contrastive learning, the vision transformer and text transformer are initialized by the image encoder and text encoder in CLIP (ViT-B/32) [51]. The maximum length of sign clip features and text features are set to 64 and 32, respectively. The model is fine-tuned with Adam optimizer [31] with batch size of 512. The initial learning rate is set to  $1 \times 10^{-5}$ , which is decreased with a cosine schedule following the CLIP [51]. We set  $\alpha = 0.8$  in Eq. 1 and  $\beta = 0.5$  in Eq.3. The languages of *PHOENIX-2014T* and *CSL-Daily* are German and Chinese

Model	T2V				V2T			
	R@1↑	R@5↑	R@10↑	MedR↓	R@1↑	R@5↑	R@10↑	MedR↓
SA-SR [18]	18.9	32.1	36.5	62.0	11.6	27.4	32.5	69.0
SA-CM [18]	24.3	40.7	46.5	16.0	17.9	40.1	46.9	14.0
SA-COMB [18]	34.2	48.0	52.6	8.0	23.6	47.0	53.0	7.5
<b>Ours</b>	<b>56.6</b>	<b>69.9</b>	<b>74.7</b>	<b>1.0</b>	<b>51.6</b>	<b>64.8</b>	<b>70.1</b>	<b>1.0</b>

Table 1. Comparison with the different variants of the pioneer SPOT-ALIGN (SA) [18] on How2Sign [19] dataset.

Model	T2V				V2T			
	R@1↑	R@5↑	R@10↑	MedR↓	R@1↑	R@5↑	R@10↑	MedR↓
Translation [10]	30.2	53.1	63.4	4.5	28.8	52.0	60.8	56.1
SA-CM [18]	48.6	76.5	84.6	2.0	50.3	78.4	84.4	<b>1.0</b>
SA-COMB [18]	55.8	79.6	87.2	<b>1.0</b>	53.1	79.4	86.1	<b>1.0</b>
<b>Ours</b>	<b>69.5</b>	<b>86.6</b>	<b>92.1</b>	<b>1.0</b>	<b>70.2</b>	<b>88.0</b>	<b>92.8</b>	<b>1.0</b>

Table 2. Comparison with the different variants of the pioneer SPOT-ALIGN (SA) [18] on PHOENIX2014T [8] dataset.

Model	T2V				V2T			
	R@1↑	R@5↑	R@10↑	MedR↓	R@1↑	R@5↑	R@10↑	MedR↓
Ours	75.3	88.2	91.9	1.0	74.7	89.4	92.2	1.0

Table 3. We additionally provide a baseline for CSL-Daily [71] dataset.

respectively. Since CLIP is trained on English corpus, to reuse CLIP’s text encoder, we utilize Google translation to translate the texts of these two datasets into English.

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

We compare our method with different variants of the pioneer, called SPOT-ALIGN [18], on How2Sign and PHOENIX-2014T. We also provide the results on CSL-Daily as a baseline.

Table 1 and Table 2 show the comparisons between our approach and SPOT-ALIGN [18] on How2Sign and PHOENIX-2014T, respectively. SPOT-ALIGN builds the final combination (COMB) model by integrating its primary cross-modal (CM) model with an auxiliary retrieval model (sign recognition (SR) model for How2Sign and Translation [10] for PHOENIX-2014T). Our method outperforms the COMB model, which achieves best results in SPOT-ALIGN, by large margins, achieving +22.4 T2V and +28.0 V2T R@1 improvements on How2Sign, +13.7 T2V and +17.1 V2T R@1 improvements on PHOENIX-2014T. It is worth mentioning that the SPOT-ALIGN conducts three rounds of sign spotting [1, 47] and encoder training. In contrast, we simplify the training of sign encoder and only perform a single round of training on pseudo-labeled data. We also provide a baseline on CSL-Daily dataset as shown in Table 3, demonstrating that our model can be generalized to various sign languages.

Method	T2V			V2T		
	R@1	R@5	R@10	R@1	R@5	R@10
Baseline	31.5	49.6	57.9	26.4	44.4	52.8
+SE	33.3	51.9	59.8	30.9	48.8	56.4
+SE+CLCL	54.0	67.1	71.9	50.0	63.2	68.1
+SE+CLCL+TA	<b>56.6</b>	<b>69.9</b>	<b>74.7</b>	<b>51.6</b>	<b>64.8</b>	<b>70.1</b>

Table 4. Ablation study of each proposed component. SE: sign encoder; CLCL: cross-lingual contrastive learning; TA: text augmentation.

Encoder	T2V			V2T		
	R@1	R@5	R@10	R@1	R@5	R@10
Single-Ag [61]	53.1	68.0	73.4	47.3	62.9	67.0
Single-Aw	54.1	67.5	73.1	49.1	61.8	67.4
Fusion-Average	54.7	68.7	73.8	49.6	63.7	69.0
<b>Fusion-Weighted Sum</b>	<b>56.6</b>	<b>69.9</b>	<b>74.7</b>	<b>51.6</b>	<b>64.8</b>	<b>70.1</b>

Table 5. Results of domain-agnostic sign encoder (Single-Ag) and domain-aware sign encoder (Single-Aw). We also study different ways to integrate the features extracted by Single-Ag and Single-Aw, including average and weighted sum.

## 4.3. Ablation Study

We conduct all ablation studies on the most challenging How2Sign dataset.

**The Effectiveness of Each Proposed Component.** Table 4 shows the ablation of each component. We first build a baseline where the sign encoder is the domain-agnostic one and contrastive learning is trained with the standard contrastive loss [51]. Then we add the proposed sign encoder (SE), cross-lingual contrastive learning (CLCL) and text augmentation (TA) to the baseline step by step. SE encodes domain-relevant and discriminative features, yielding +1.8 T2V and +4.5 V2T R@1 gains. The proposed CLCL significantly boosts the performance by +20.7 T2V and +19.1 V2T R@1 improvements, demonstrating that identifying fine-grained sign-to-word mappings is essential for sign language retrieval. The introduction of TA further promotes the retrieval task, achieving 56.6 R@1 on V2T and 51.6 R@1 on T2V, respectively.

**Various Sign Encoders.** As described in Section 3.2, our sign encoder is composed of a domain-agnostic sign encoder (Single-Ag) [61] and a domain-aware sign encoder (Single-Aw). We first report the results of each individual encoder in Table 5. Next, we study the different ways to integrate the features extracted by these two encoders, including average and weighted sum as defined in Eq. 1. The results are also shown in Table 5. We experimentally find that the weighted sum strategy yields best results.

**Variants of Cross-Lingual Contrastive Learning.** As described in Section 3.3 and illustrated in Figure 3b, through the use of a combination of softmax, multiplication, sum, and average operations, we convert the fine-grained cross-lingual (sign-to-word) similarity to the coarse-grained sign-video-to-text similarity to enable contrastive learning. We

Strategy	T2V			V2T		
	R@1	R@5	R@10	R@1	R@5	R@10
Mean	33.1	52.5	59.7	29.8	47.8	55.3
Max	42.2	59.7	66.0	38.5	55.1	62.0
<b>Softmax</b>	<b>56.6</b>	<b>69.9</b>	<b>74.7</b>	<b>51.6</b>	<b>64.8</b>	<b>70.1</b>

Table 6. Study on different strategies to identify the fine-grained sign-to-word mappings in cross-lingual contrastive learning.

Text Augmentation	T2V			V2T		
	R@1	R@5	R@10	R@1	R@5	R@10
None	54.0	67.1	71.9	50.0	63.2	68.1
RD	52.7	66.3	71.6	47.6	62.1	67.4
SR	53.9	68.7	73.0	47.5	61.2	65.9
<b>RS</b>	<b>56.6</b>	<b>69.9</b>	<b>74.7</b>	<b>51.6</b>	<b>64.8</b>	<b>70.1</b>

Table 7. Ablation study on different text augmentations. RD: random delete; SR: synonym replacement; RS: random swap.

refer to this process as ‘‘Softmax’’. Here we study two variants termed ‘‘Mean’’ and ‘‘Max’’. The ‘‘Mean’’ and ‘‘Max’’ strategies simply replace the combination of softmax, multiplication and sum operations with a simple mean operation and a max operation, respectively. The results are shown in Table 6. We observe that the ‘‘Mean’’ strategy performs worst since it merely evaluates the overall similarity of a text and a sign video and ignores the fine-grained sign-to-word mappings during training. The ‘‘Max’’ strategy identifies hard sign-to-word mappings, *i.e.*, one sign is associated with the most similar word, and vice versa. Nevertheless, since we do not have the ground-truth of sign-to-word mappings, it is challenging for models to identify the confident one-to-one mappings in a weakly supervised manner, as discussed in previous works on multiple instance learning [47]. In contrast, our default ‘‘Softmax’’ strategy localizes the soft sign-to-word mappings and achieves the best results.

**Study of Text Augmentations.** In Table 7, we study three text augmentations described in Section 3.4: random delete (RD), synonym replacement (SR) and random swap (RS). We observe a slight performance drop when introducing RD and SR into training. Sign language retrieval is not only a video-text retrieval task, but also a cross-lingual retrieval challenge, deletion and replacement may break the intrinsic sign-to-word mappings. In contrast, RS augmentation preserves the semantics of texts and we observe a +2.6 T2V and a +1.6 V2T R@1 improvements over the counterpart without any text augmentations.

**The Effects of CLIP Initialization.** In our framework, the vision and text Transformer are initialized with the CLIP’s vision encoder and text encoder. Note that the input modality of our vision Transformer and that of CLIP’s vision encoder are different. Ours takes a feature sequence as input while the input of CLIP’s image encoder is a set of image patches. In Table 8, we compare a randomly initialized baseline with the one initialized by CLIP. We find that CLIP-initialization significantly improves the perfor-

Initialization	T2V			V2T		
	R@1	R@5	R@10	R@1	R@5	R@10
Random	45.2	60.3	67.5	40.6	54.5	59.7
<b>CLIP</b>	<b>56.6</b>	<b>69.9</b>	<b>74.7</b>	<b>51.6</b>	<b>64.8</b>	<b>70.1</b>

Table 8. Comparison between CLIP initialization and random initialization.

Frozen layers	T2V			V2T		
	R@1	R@5	R@10	R@1	R@5	R@10
<b>None</b>	<b>56.6</b>	<b>69.9</b>	<b>74.7</b>	<b>51.6</b>	<b>64.8</b>	<b>70.1</b>
1	54.4	68.2	73.4	49.7	63.5	68.8
2	53.8	68.6	73.3	48.5	62.8	67.9
6	51.7	66.6	71.1	46.2	60.9	66.6
12	49.5	65.2	71.4	44.7	59.7	65.2

Table 9. Study of freezing different layers of text Transformer.

mance, yielding +11.4 T2V and +11.0 V2T R@1 improvements though the input originates from a completely distinct modality.

**Study on Text Transformer.** The text Transformer in our model is initialized with the CLIP’s text encoder. Considering the generalization capability of the pre-trained CLIP, we attempt to freeze the shallow layers of our text transformer. As shown in Table 9, the performance gradually decreases as the number of the frozen layers increases. CLIP is trained on large-scale image-text pairs, however, the domain gap between image-text data and sign-video-text data is non-negligible. Though CLIP shows promising transfer capacity, it is optimal to fine-tune the whole model on target datasets.

## 5. Conclusion

In this paper, we propose a novel framework named **Cross-lingual Contrastive learning (CiCo)** for recently introduced sign language retrieval task. We formulate sign language retrieval as a cross-lingual retrieval task as well as a video-text retrieval problem. CiCo models the fine-grained cross-lingual mappings between sign videos and texts via the proposed cross-lingual contrastive learning. We also introduce a sign encoder, which is composed of a domain-agnostic encoder and a domain-aware one, to extract discriminative and domain-aligned features. Our CiCo outperforms the pioneer SPOT-ALIGN by large margins on How2Sign and PHOENIX-2014T benchmarks. We also provide a baseline on CSL-Daily. We hope our approach could serve as a solid baseline for future research.

**Acknowledgement.** This work was partially supported by National Natural Science Foundation of China (No. 62072112), National Key R&D Program of China (No. 2020AAA0108301), and Scientific and Technological Innovation Action Plan of Shanghai Science and Technology Committee (No. 22511101502, 22511102202).



## References

- [1] Samuel Albanie, Gül Varol, Liliane Momeni, Triantafyllos Afouras, Joon Son Chung, Neil Fox, and Andrew Zisserman. Bsl-1k: Scaling up co-articulated sign language recognition using mouthing cues. In *European Conference on Computer Vision*, pages 35–53. Springer, 2020. 1, 3, 7
- [2] Mikel Artetxe, Gorka Labaka, and Eneko Agirre. Learning bilingual word embeddings with (almost) no bilingual data. In *Meeting of the Association for Computational Linguistics*, pages 451–462, 2017. 3
- [3] Mikel Artetxe, Gorka Labaka, and Eneko Agirre. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. *arXiv preprint arXiv:1805.06297*, 2018. 3
- [4] Vassilis Athitsos, Carol Neidle, Stan Sclaroff, Joan Nash, Alexandra Stefan, Ashwin Thangali, Haijing Wang, and Quan Yuan. Large lexicon project: American sign language video corpus and sign language indexing/retrieval algorithms. In *Workshop on the Representation and Processing of Sign Languages: Corpora and Sign Language Technologies*, volume 2, pages 11–14, 2010. 3
- [5] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *IEEE/CVF International Conference on Computer Vision*, pages 1728–1738, 2021. 2, 3
- [6] Lisa Ballesteros and Bruce Croft. Dictionary methods for cross-lingual information retrieval. In *International Conference on Database and Expert Systems Applications*, pages 791–801. Springer, 1996. 2
- [7] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013. 3
- [8] Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, Hermann Ney, and Richard Bowden. Neural sign language translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7784–7793, 2018. 2, 3, 6, 7
- [9] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. Multi-channel transformers for multi-articulatory sign language translation. In *European Conference on Computer Vision*, pages 301–319. Springer, 2020. 3
- [10] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. Sign language transformers: Joint end-to-end sign language recognition and translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10023–10033, 2020. 1, 7
- [11] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 3, 6
- [12] Yutong Chen, Fangyun Wei, Xiao Sun, Zhirong Wu, and Stephen Lin. A simple multi-modality transfer learning baseline for sign language translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5120–5130, 2022. 1, 3
- [13] Yutong Chen, Ronglai Zuo, Fangyun Wei, Yu Wu, Shujie LIU, and Brian Mak. Two-stream network for sign language recognition and translation. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 1, 3
- [14] Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, and Richard Bowden. Subunets: End-to-end hand shape and continuous sign language recognition. In *IEEE/CVF International Conference on Computer Vision*, pages 3056–3065, 2017. 1
- [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. 1
- [16] Tonni Das Jui, Gissella Maria Bejarano, and Pablo Rivas. A machine learning-based segmentation approach for measuring similarity between sign languages. In *sign-lang@LREC 2022*, pages 94–101. European Language Resources Association (ELRA), 2022. 4
- [17] Yu Du, Fangyun Wei, Zihe Zhang, Miaoqing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14084–14093, 2022. 5
- [18] Amanda Duarte, Samuel Albanie, Xavier Giró-i Nieto, and Gül Varol. Sign language video retrieval with free-form textual queries. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14094–14104, 2022. 1, 2, 3, 4, 6, 7
- [19] Amanda Duarte, Shruti Palaskar, Lucas Ventura, Deepti Ghadiyaram, Kenneth DeHaan, Florian Metze, Jordi Torres, and Xavier Giro-i Nieto. How2Sign: A Large-scale Multimodal Dataset for Continuous American Sign Language. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 2, 6, 7
- [20] Ali Farhadi, David Forsyth, and Ryan White. Transfer learning in sign language. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. 3
- [21] Holger Fillbrandt, Suat Akyol, and K-F Kraiss. Extraction of 3d hand shape and posture from image sequences for sign language recognition. In *2003 IEEE International SOI Conference*, pages 181–186. IEEE, 2003. 3
- [22] Stanislav Fort, Jie Ren, and Balaji Lakshminarayanan. Exploring the limits of out-of-distribution detection. *Advances in Neural Information Processing Systems*, 34:7068–7081, 2021. 5
- [23] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems*, 26, 2013. 3, 5
- [24] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *European Conference on Computer Vision*, pages 214–229. Springer, 2020. 2, 3
- [25] Yunchao Gong, Qifa Ke, Michael Isard, and Svetlana Lazebnik. A multi-view embedding space for modeling internet images, tags, and their semantics. *International Journal of Computer Vision*, 106(2):210–233, 2014. 3, 5

- [26] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *International Conference on Artificial Intelligence and Statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010. 5
- [27] Tony Huang, Jack Chu, and Fangyun Wei. Unsupervised prompt learning for vision-language models. *arXiv preprint arXiv:2204.03649*, 2022. 5
- [28] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. 3
- [29] Armand Joulin, Laurens van der Maaten, Allan Jabri, and Nicolas Vasilache. Learning visual features from large weakly supervised data. In *European Conference on Computer Vision*, pages 67–84. Springer, 2016. 3, 5
- [30] Hamid Reza Vaezi Joze and Oscar Koller. Ms-asl: A large-scale data set and benchmark for understanding american sign language. *arXiv preprint arXiv:1812.01053*, 2018. 3
- [31] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. 6
- [32] 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*, 42(9):2306–2320, 2019. 1
- [33] 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, 2015. 1
- [34] Victor Lavrenko, Martin Choquette, and W Bruce Croft. Cross-lingual relevance models. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 175–182, 2002. 2, 3
- [35] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. 3
- [36] 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 *IEEE/CVF winter conference on applications of computer vision*, pages 1459–1469, 2020. 3
- [37] Dongxu Li, Chenchen Xu, Xin Yu, Kaihao Zhang, Benjamin Swift, Hanna Suominen, and Hongdong Li. Tspnet: Hierarchical feature learning via temporal semantic pyramid for sign language translation. *Advances in Neural Information Processing Systems*, 33:12034–12045, 2020. 3
- [38] Dongxu Li, Xin Yu, Chenchen Xu, Lars Petersson, and Hongdong Li. Transferring cross-domain knowledge for video sign language recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6205–6214, 2020. 3
- [39] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in Neural Information Processing Systems*, 34:9694–9705, 2021. 3, 5
- [40] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208*, 2021. 5
- [41] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *IEEE/CVF International Conference on Computer Vision*, pages 7083–7093, 2019. 3
- [42] Song Liu, Haoqi Fan, Shengsheng Qian, Yiru Chen, Wenkui Ding, and Zhongyuan Wang. Hit: Hierarchical transformer with momentum contrast for video-text retrieval. In *IEEE/CVF International Conference on Computer Vision*, pages 11915–11925, 2021. 2, 3
- [43] Yang Liu, Samuel Albanie, Arsha Nagrani, and Andrew Zisserman. Use what you have: Video retrieval using representations from collaborative experts. *arXiv preprint arXiv:1907.13487*, 2019. 2, 3, 6
- [44] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022. 2, 3
- [45] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *IEEE/CVF International Conference on Computer Vision*, pages 2630–2640, 2019. 2
- [46] Tomas Mikolov, Quoc V Le, and Ilya Sutskever. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*, 2013. 3
- [47] Liliane Momeni, Gül Varol, Samuel Albanie, Triantafyllos Afouras, and Andrew Zisserman. Watch, read and lookup: learning to spot signs from multiple supervisors. In *Asian Conference on Computer Vision*, 2020. 3, 4, 7, 8
- [48] Mathew Monfort, SouYoung Jin, Alexander Liu, David Harwath, Rogerio Feris, James Glass, and Aude Oliva. Spoken moments: Learning joint audio-visual representations from video descriptions. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14871–14881, 2021. 2
- [49] Yasuhide Mori, Hironobu Takahashi, and Ryuichi Oka. Image-to-word transformation based on dividing and vector quantizing images with words. In *International Workshop on Multimedia Intelligent Storage and Retrieval Management*, pages 1–9. Citeseer, 1999. 3
- [50] Aitor Ormazabal, Mikel Artetxe, Gorka Labaka, Aitor Soroa, and Eneko Agirre. Analyzing the limitations of cross-lingual word embedding mappings. *arXiv preprint arXiv:1906.05407*, 2019. 3
- [51] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language super-

- vision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 3, 5, 6, 7
- [52] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Conference on Neural Information Processing Systems*, 28, 2015. 6
- [53] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015. 5
- [54] Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. How much can clip benefit vision-and-language tasks? *arXiv preprint arXiv:2107.06383*, 2021. 5
- [55] Thad Starner, Joshua Weaver, and Alex Pentland. Real-time american sign language recognition using desk and wearable computer based video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12):1371–1375, 1998. 3
- [56] Thad E Starner. Visual recognition of american sign language using hidden markov models. Technical report, Massachusetts inst of tech Cambridge dept of brain and cognitive sciences, 1995. 3
- [57] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *IEEE/CVF International Conference on Computer Vision*, pages 7464–7473, 2019. 2, 3, 6
- [58] Rachel Sutton-Spence. Mouthings and simultaneity in british sign language. *Amsterdam Studies in the Theory and History of Linguistic Science Series*, 281:147, 2007. 3
- [59] Chau Tran, Yuqing Tang, Xian Li, and Jiatao Gu. Cross-lingual retrieval for iterative self-supervised training. *Advances in Neural Information Processing Systems*, 33:2207–2219, 2020. 2, 3
- [60] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018. 3
- [61] Gul Varol, Liliane Momeni, Samuel Albanie, Triantafyllos Afouras, and Andrew Zisserman. Read and attend: Temporal localisation in sign language videos. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16857–16866, 2021. 1, 2, 3, 4, 6, 7
- [62] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. 2017. 5
- [63] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018. 3
- [64] Jason Wei and Kai Zou. EDA: Easy data augmentation techniques for boosting performance on text classification tasks. In *2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6383–6389, Hong Kong, China, Nov. 2019. Association for Computational Linguistics. 6
- [65] Jinxi Xu, Ralph Weischedel, and Chanh Nguyen. Evaluating a probabilistic model for cross-lingual information retrieval. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 105–110, 2001. 3
- [66] Mengde Xu, Zheng Zhang, Fangyun Wei, Han Hu, and Xiang Bai. Side adapter network for open-vocabulary semantic segmentation. *arXiv preprint arXiv:2302.12242*, 2023. 5
- [67] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In *European Conference on Computer Vision*, pages 736–753. Springer, 2022. 5
- [68] Youngjae Yu, Jongseok Kim, and Gunhee Kim. A joint sequence fusion model for video question answering and retrieval. In *European Conference on Computer Vision*, pages 471–487, 2018. 2, 3
- [69] Mozhi Zhang, Keyulu Xu, Ken-ichi Kawarabayashi, Stefanie Jegelka, and Jordan Boyd-Graber. Are girls neko or shōjo? cross-lingual alignment of non-isomorphic embeddings with iterative normalization. *arXiv preprint arXiv:1906.01622*, 2019. 3
- [70] Shilin Zhang and Bo Zhang. Using revised string edit distance to sign language video retrieval. In *International Conference on Computational Intelligence and Natural Computing*, volume 1, pages 45–49. IEEE, 2010. 3
- [71] Hao Zhou, Wengang Zhou, Weizhen Qi, Junfu Pu, and Houqiang Li. Improving sign language translation with monolingual data by sign back-translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1316–1325, 2021. 6, 7
- [72] Hao Zhou, Wengang Zhou, Yun Zhou, and Houqiang Li. Spatial-temporal multi-cue network for sign language recognition and translation. *IEEE Transactions on Multimedia*, 24:768–779, 2021. 3
- [73] Ronglai Zuo, Fangyun Wei, and Brian Mak. Natural language-assisted sign language recognition. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 1, 3