

Sign Language Translation with Iterative Prototype

Huijie Yao¹ Wengang Zhou^{1,2,*} Hao Feng¹ Hezhen Hu¹ Hao Zhou¹ Houqiang Li^{1,2,*}

¹ CAS Key Laboratory of Technology in GIPAS, EEIS Department, University of Science and Technology of China

² Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

{yaohuijie,haof,alexhu,zhouh156}@mail.ustc.edu.cn, {zhwg,lihq}@ustc.edu.cn

Abstract

This paper presents *IP-SLT*, a simple yet effective framework for sign language translation (SLT). Our *IP-SLT* adopts a recurrent structure and enhances the semantic representation (prototype) of the input sign language video via an iterative refinement manner. Our idea mimics the behavior of human reading, where a sentence can be digested repeatedly, till reaching accurate understanding. Technically, *IP-SLT* consists of feature extraction, prototype initialization, and iterative prototype refinement. The initialization module generates the initial prototype based on the visual feature extracted by the feature extraction module. Then, the iterative refinement module leverages the cross-attention mechanism to polish the previous prototype by aggregating it with the original video feature. Through repeated refinement, the prototype finally converges to a more stable and accurate state, leading to a fluent and appropriate translation. In addition, to leverage the sequential dependence of prototypes, we further propose an iterative distillation loss to compress the knowledge of the final iteration into previous ones. As the autoregressive decoding process is executed only once in inference, our *IP-SLT* is ready to improve various SLT systems with acceptable overhead. Extensive experiments are conducted on public benchmarks to demonstrate the effectiveness of the *IP-SLT*.

1. Introduction

Sign language translation (SLT) aims to automatically generate spoken language translations based on sign language videos, which holds both social significance and academic value. On the one hand, a high-quality SLT system can greatly facilitate communication between deaf-mute and hearing individuals [4, 11, 12, 41]. On the other hand, SLT as an interdisciplinary research topic necessitates a comprehensive understanding of computer vision [31, 13, 36, 19] and natural language processing [43, 40], given its

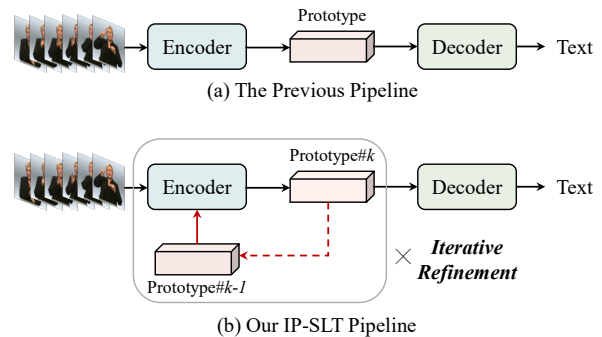


Figure 1. Illustration of the pipeline of the previous works and our *IP-SLT*. (a) The previous studies rely on a one-pass forward process to generate the final translation. (b) To mitigate the vision-text gap, we introduce the iterative refinement module into the original SLT system. The refinement module updates the current prototype conditioned on the sign language video, which can be run iteratively to obtain a better representation of the semantic meaning of the sign language video.

involvement with vision and text modalities. As a result, SLT has emerged as a vital research topic, garnering increasing attention [5, 49, 31, 2, 21, 20, 42].

SLT is a challenging task, which faces a tough domain gap between the input video and output text, as well as a limited dataset scale due to costly data collection and annotation [49, 7, 44, 21]. Since SLT is typically viewed as a sequence-to-sequence mapping problem, the existing SLT systems [23, 5, 7] commonly adopt the one-pass forward pipeline based on encoder-decoder architecture [40, 30] (as shown in Fig. 1 (a)). In such a framework, the encoder transforms the sign video into its semantic representation (prototype), which is then fed into the decoder to obtain the final translation. However, due to the inherent gap between vision and text, it may be hard to conduct such mapping within the vanilla one-pass architecture.

In this study, we present *IP-SLT* with the iterative prototype to boost sign language translation (as shown in Fig. 1 (b)), which is inspired by the human reading process. During this process, we note that repeatedly digging into the

*Corresponding authors: Wengang Zhou and Houqiang Li

source materials is necessary for accurate understanding. Similarly, when we are trying to translate a sign language video into a sentence, we commonly do not directly write it down. Instead, we would recall and go back to the original sign video to check our answers. To implement the above idea, our IP-SLT adopts a recurrent structure that enhances the semantic representation (prototype) of the input sign language video via an iterative refinement process. IP-SLT generally contains three main components, including feature extraction, prototype initialization, and iterative prototype refinement. Given a sign video to be translated, we first extract its visual representation, which is then used for generating an initial raw prototype. Subsequently, we iteratively leverage the attention mechanism [43] to update the prototype toward the semantic meaning of the sign video. At each iteration, we refine the previous prototype by aggregating it with the original visual representation. In this way, the network repeatedly digs the semantic context from the sign video to polish the prototype. Through iterative refinement, the prototype finally converges to a stable and accurate state, producing a high-quality translation.

In addition, our IP-SLT introduces a novel design discussed next. Firstly, to leverage the sequential dependence between different iterations, we further propose the iterative distillation loss which allows the previous prototypes to obtain supervision from the final one. Since the final prototype converges to a more stable and accurate state, it is possible for IP-SLT to achieve better performance. Secondly, during training, all predicted prototypes are transformed into their corresponding translations to provide guidance for each iteration. Our inference process is neat since only the final prototype is used for the autoregressive decoding process. Thirdly, our IP-SLT can easily work with different visual backbones. Through end-to-end optimization, our IP-SLT achieves significant performance improvements over the baselines.

In summary, our contributions are three-fold:

- We propose IP-SLT, a novel framework to ameliorate sign language translation, which iteratively refines the prototypes by aggregating the previous translation progress and the original visual representation.
- We propose an iterative distillation loss to enhance the basic supervision, by leveraging the sequential dependence between the outputs at each iteration.
- We conduct extensive experiments to validate the proposed method, and show encouraging improved results on the two prevalent benchmarks, *i.e.*, CSL-Daily [49] and PHOENIX-2014T [5].

2. Related Work

In this section, we briefly review the related works, *i.e.*, sign language translation and iterative refinement methods.

Sign Language Translation. Camgoz *et al.* [5] pioneer the neural SLT task and publish the neural dataset PHOENIX-2014T and regard the SLT as a sequence-to-sequence problem. They implement the neural SLT system using the encoder-decoder paradigm [3]. This paradigm is adopted by subsequent studies which focus on addressing the challenges of data scarcity and domain gap. Considering the lack of frame-level annotation in sign language datasets, Li *et al.* [27] design the temporal semantic pyramid structure to obtain more discriminative features. Camgoz *et al.* [7] explore the mutual benefits of SLT and continuous sign language recognition through joint optimization. Zhou *et al.* [49] leverage gloss annotation to transform the monolingual texts into pseudo-videos. According to the characteristic of sign language, several works [46, 6, 50] propose multi-channel SLT systems which explicitly extract and align the key parts of sign language expression. Jin *et al.* [23] leverage the additional prior knowledge to obtain high-quality translations. Chen *et al.* [9] propose a transfer learning baseline for SLT by leveraging external resources from related tasks. Chen *et al.* [10] further combine the raw videos and the keypoint sequences to achieve better semantic understanding with auxiliary supervision.

In contrast, our proposed approach employs an iterative refinement process that utilizes the previous prototype as an additional clue to enhance the accuracy of the mapping between sign videos and their translations.

Iterative Refinement Methods. The idea of iterative refinement is applied to various computer vision tasks, such as image generation [1, 18, 37], instance segmentation [34, 28, 29, 47], image rectification [15], *etc.*, which shows promising performance improvements. CARN [1] maintains training stability in super-resolution tasks and improves the quality of output images. Ling *et al.* [28] regard the object instance segmentation task as a regression task and proposes the Curve-GCN to iteratively predict the locations of all vertices. DeepSanke [34] uses the deep network to iteratively enclose the object boundary based on an initial contour. The previous studies [48, 35, 14, 36, 31] in continuous sign language recognition task follow an iterative training scheme to enhance the discriminative power of feature extraction modules which use the convolutional neural network and their variants [8, 22, 39]. They leverage alignment proposals given by the connectionist temporal classification (CTC) [17] decoding as supervision at frame-wise granularity, which cannot be directly applied to the SLT.

Different from the aforementioned methods, we explore how to reduce the vision-text gap by proposing an iterative refinement module to the existing SLT system. To reduce the complexity, we design the refinement module in a shared-weight manner. Moreover, we put forward the iterative distillation loss to leverage the sequential dependence between different iterations.

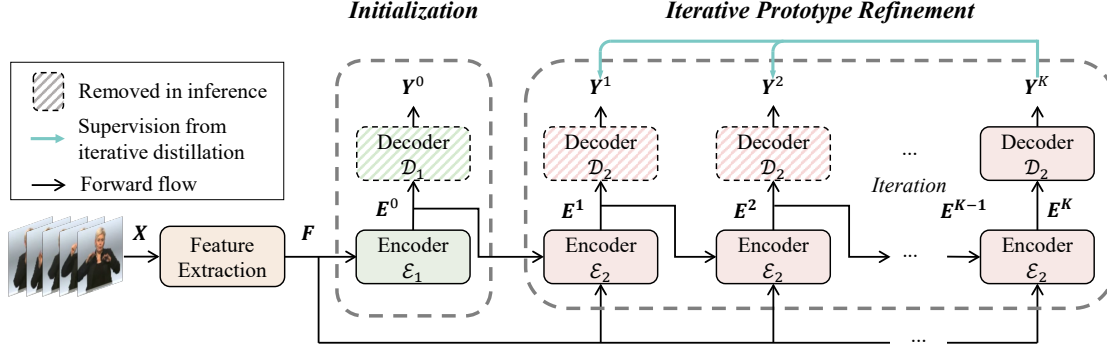


Figure 2. An overview of the proposed IP-SLT framework. Given a sign video \mathbf{X} , the feature extraction module is responsible for embedding the input into visual representation \mathbf{F} . The initialization module (the encoder \mathcal{E}_1 and decoder \mathcal{D}_1) generates the initial prototype \mathbf{E}^0 and the raw translation \mathbf{Y}^0 . The refinement module (the encoder \mathcal{E}_2 and decoder \mathcal{D}_2) first takes the initial prototype \mathbf{E}^0 as input and generates a prototype for the current step by fusing it with the original visual representation \mathbf{F} . Through K times refinement, the prototype sequence $\mathbf{E} = \{\mathbf{E}^0, \mathbf{E}^1, \dots, \mathbf{E}^K\}$ and corresponding translation sequence $\mathbf{Y} = \{\mathbf{Y}^0, \mathbf{Y}^1, \dots, \mathbf{Y}^K\}$ are obtained. In light of the fact that the decoding part of IP-SLT consists of $K + 1$ branches based on the iteration order, we introduce the iterative distillation loss to improve the underlying supervision. It should be noted that the parts enclosed in dashed boxes can be removed in inference.

3. Methodology

In this section, we first introduce the overall architecture of our IP-SLT, and then separately elaborate individual components. Finally, we propose the design of the training objective and inference strategy for the IP-SLT.

3.1. Framework Overview

The primary objective of the SLT system is to acquire knowledge about the mapping $f : \mathcal{X} \mapsto \mathcal{Y}$, where \mathcal{X} and \mathcal{Y} denote the collections of N sign language videos and spoken language sentences associated with vocabulary \mathcal{V} , respectively. Most SLT systems adopt the encoder-decoder architecture [40], where the input $\mathbf{X} \in \mathcal{X}$ is first encoded to derive a high-level context representation. It is then passed to the decoder to generate the output $\mathbf{Y} \in \mathcal{Y}$. The encoder and decoder can be specialized using different types of neural networks, such as GRU [3], CNN [16], and Transformer [43]. Considering the performance of existing SLT systems, we adopt the Transformer as well.

With the goal of narrowing the domain gap between vision and text, we augment the original translation process with an iterative refinement step. Fig. 2 provides an overview of the proposed IP-SLT model, which consists of three stages, namely feature extraction, prototype initialization, and iterative prototype refinement. As with previous approaches, the initialization and iterative refinement module adopt the encoder-decoder architecture. Given the sign language video $\mathbf{X} = \{\mathbf{x}_t\}_{t=1}^{T_x}$ with T_x frames, the feature extraction module embeds it into the spatial-temporal feature $\mathbf{F} = \{\mathbf{f}_t\}_{t=1}^{T_f}$. Next, the encoder \mathcal{E}_1 and decoder \mathcal{D}_1 are employed in the initialization module to derive the initial prototype \mathbf{E}^0 and initial translation \mathbf{Y}^0 from the visual feature \mathbf{F} .

Subsequently, the refinement module iteratively refines the previous prototype and generates the final translation $\mathbf{Y}^K = \{\mathbf{y}_t^K\}_{t=1}^{T_{y,K}}$ with $T_{y,K}$ words after total K iterations. At the k -th iteration, the encoder \mathcal{E}_2 estimates the prototype \mathbf{E}^k for the current step by augmenting the original visual feature \mathbf{F} with the previous prototype \mathbf{E}^{k-1} . Next, the decoder \mathcal{D}_2 predicts the corresponding translation \mathbf{Y}^k . Finally, the translation sequence $\mathbf{Y} = \{\mathbf{Y}^0, \mathbf{Y}^1, \dots, \mathbf{Y}^K\}$ is obtained according to the prototype sequence $\mathbf{E} = \{\mathbf{E}^0, \mathbf{E}^1, \dots, \mathbf{E}^K\}$. For the IP-SLT optimization, by dividing the decoding part into $K + 1$ branches, we add iterative distillation supervision from the final translation to the middle translations. Since the refinement process takes place in the encoder \mathcal{E}_2 , in inference, the proposed IP-SLT can generate the translation directly from the K -th prototype which causes acceptable overhead.

3.2. Feature Extraction

The feature extraction module embeds a series of video frames $\mathbf{X} \in \mathbb{R}^{T_x \times H \times W \times 3}$ with width W and height H into its visual feature $\mathbf{F} \in \mathbb{R}^{T_f \times C}$ with the dimension of feature C . Since its goal is to extract a distinguishable representation for SLT, we can draw on the visual backbone used in CSLR [14, 26, 25, 13] to extract the valid representation. Generally, with sliding window size w and stride size s , the sign video is split into $T_f = \lceil \frac{T_x}{s} \rceil$ clips. By passing sign videos through it, the spatial-temporal embeddings $\mathbf{F} = \{\mathbf{f}_t\}_{t=1}^{T_f}$ are extracted as:

$$\{\mathbf{f}_t\}_{t=1}^{T_f} = \text{Extractor}(\{\mathbf{x}_t\}_{t=1}^{T_x}). \quad (1)$$

3.3. Prototype Initialization

After a visual feature $\mathbf{F} \in \mathbb{R}^{T_f \times C}$ is extracted by the feature extraction module, it is first fed into the initialization

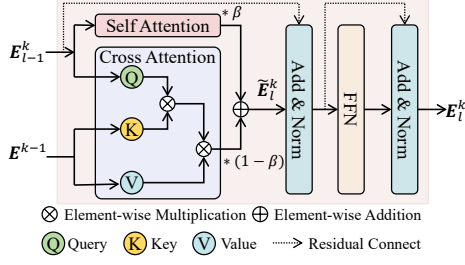


Figure 3. An illustration of the prototype update process in the l -th layer of the encoder \mathcal{E}_2 at the k -th iteration.

module. The initialization module consists of an encoder \mathcal{E}_1 and a decoder \mathcal{D}_1 .

The visual representation \mathbf{F} is first fed into the encoder \mathcal{E}_1 , and encoded into T_f raw states $\mathbf{E}^0 = \{e_t^0\}_{t=1}^{T_f} \in \mathbb{R}^{T_f \times C}$. Then the decoder \mathcal{D}_1 reads the prototype \mathbf{E}^0 and produces the initial translation $\mathbf{Y}^0 = \{\mathbf{y}_t^0\}_{t=1}^{T_{y,0}}$ according to the predicted logits $\mathbf{U}^0 = \{\mathbf{u}_t^0\}_{t=1}^{T_{y,0}}$. Specifically, it predicts the conditional probability of the translation sequence, which is formulated as:

$$p^{0-th}(\mathbf{Y}^0|\mathbf{X}) = \prod_{t=1}^{T_{y,0}} p(\mathbf{y}_t^0|\mathbf{F}, \mathbf{y}_{0:t-1}^0), \quad (2)$$

where $\mathbf{y}_{0:t-1}^0 = \{\mathbf{y}_0^0, \mathbf{y}_1^0, \dots, \mathbf{y}_{t-1}^0\}$ denotes the previous output sub-sequence at the t -th step. The initial token \mathbf{y}_0^0 represents the beginning of a sentence. The predicted probability of each token in the translation is computed as:

$$\begin{aligned} p(\mathbf{y}_t^0|\mathbf{F}, \mathbf{y}_{1:t-1}^0) &= \text{softmax}(\mathbf{u}_t^0)_{\mathbf{y}_t^0} \\ &= \text{softmax}(\mathbf{h}_t^0 \cdot \mathbf{W})_{\mathbf{y}_t^0}, \end{aligned} \quad (3)$$

where $\mathbf{h}_t^0 \in \mathbb{R}^C$ represents the output of the final layer at the t -th step, and $\mathbf{W} \in \mathbb{R}^{C \times |\mathcal{V}|}$ denotes a linear mapping to projects the hidden state \mathbf{h}_t^0 into the predicted logits over the target vocabulary \mathcal{V} . The probability is calculated by applying the $\text{softmax}(\cdot)$ function to the logits. Notably, our goal is to obtain a more accurate prototype for SLT, thus, the decoder \mathcal{D}_1 is only used in the training process to provide guidance for the initialization module.

3.4. Iterative Prototype Refinement

Once the initial prototype \mathbf{E}^0 is obtained through the initialization module, we feed it together with the original visual representation \mathbf{F} into the iterative refinement module. The module maintains a single prototype which is iteratively refined. In this way, the coarse semantic feature finally converges to a stable state where the prototype best fits the sign language semantics. The refinement module is divided into two sub-processes, *i.e.*, iterative prototype aggregation in the encoder \mathcal{E}_2 and translation generation in the decoder \mathcal{D}_2 .

Iterative prototype aggregation. To utilize the prototype as a reference, each layer of the encoder \mathcal{E}_2 attends over the maintained semantic features through the attention mechanism. As shown in Fig. 3, we illustrate the prototype aggregation process in the l -th layer of the encoder \mathcal{E}_2 at the k -th iteration. The input of the refinement module consists of the original visual feature \mathbf{F} and the $(k-1)$ -th prototype \mathbf{E}^{k-1} . The attention mechanism $\text{attn}(\mathbf{q}, \mathbf{K}, \mathbf{V})$ is originally used in Transformer [43], which is formulated as:

$$\begin{aligned} \text{attn}(\mathbf{q}, \mathbf{K}, \mathbf{V}) &= \sum_{i=1}^{|\mathcal{V}|} \alpha_i \mathbf{W}_v \mathbf{v}_i, \\ \alpha_i &= \text{softmax}((\mathbf{W}_q \mathbf{q})^T (\mathbf{W}_k \mathbf{k}_i)), \end{aligned} \quad (4)$$

where \mathbf{W}_q , \mathbf{W}_k and \mathbf{W}_v are learnable parameters. To make better use of the $(k-1)$ -th prototype, we consider the $(k-1)$ -th prototype $\mathbf{E}^{k-1} = \{e_t^{k-1}\}_{t=1}^{T_f}$ as key-value pair, and the original visual feature $\mathbf{F} = \{\mathbf{f}_t\}_{t=1}^{T_f}$ as the query. In this way, we inject the rich semantic information into the new prototype $\mathbf{E}^k = \{e_t^k\}_{t=1}^{T_f}$.

The encoder \mathcal{E}_2 is composed of L_e identical layers, where \mathbf{E}_l^k denotes the output of the l -th layer at the k -th iteration. Similar to [51], the hidden state is computed as:

$$\begin{aligned} \tilde{\mathbf{E}}_l^k &= \beta \cdot \text{attn}_s(\mathbf{E}_{l-1}^k, \mathbf{E}_{l-1}^k, \mathbf{E}_{l-1}^k) \\ &\quad + (1 - \beta) \cdot \text{attn}_c(\mathbf{E}_{l-1}^k, \mathbf{E}^{k-1}, \mathbf{E}^{k-1}), \end{aligned} \quad (5)$$

where attn_s and attn_c denote the self-attention sub-layer and cross-attention sub-layer used in Transformer [43], respectively. For the first layer of the encoder \mathcal{E}_2 , \mathbf{E}_{l-1}^k is equal to the visual feature \mathbf{F} . The output of the encoder \mathcal{E}_2 is $\mathbf{E}_{L_e}^k$ (*i.e.*, \mathbf{E}^k). $\beta \in [0, 1]$ is a hyperparameter that weights the importance of previous prototypes during training and inference. To further fuse and refine the prototype, it is linked with a fully connected sub-layer $\text{FFN}(\cdot)$ using the residual connection. The output of the l -th layer at the k -th iteration is formulated as:

$$\mathbf{E}_l^k = \text{FFN}(\text{LN}(\tilde{\mathbf{E}}_l^k + \mathbf{E}_{l-1}^k)) + \text{LN}(\tilde{\mathbf{E}}_l^k + \mathbf{E}_{l-1}^k), \quad (6)$$

where $\text{LN}(\cdot)$ is the layer normalization operation.

Translation generation. The decoder \mathcal{D}_2 iteratively takes the prototype \mathbf{E}^k as input and generates the corresponding translation \mathbf{Y}^k . Take the k -th iteration as an example. The decoder \mathcal{D}_2 predicts the conditional probability of translation $\mathbf{Y}^k = \{\mathbf{y}_t^k\}_{t=1}^{T_{y,k}}$ based on the predicted logits $\mathbf{U}^k = \{\mathbf{u}_t^k\}_{t=1}^{T_{y,k}}$, which is computed as:

$$\begin{aligned} p^{k-th}(\mathbf{Y}^k|\mathbf{X}) &= \prod_{t=1}^{T_{y,k}} p(\mathbf{y}_t^k|\mathbf{F}, \mathbf{E}^{k-1}, \mathbf{y}_{0:t-1}^k), \\ p(\mathbf{y}_t^k|\mathbf{F}, \mathbf{E}^{k-1}, \mathbf{y}_{0:t-1}^k) &= \text{softmax}(\mathbf{u}_t^k)_{\mathbf{y}_t^k}, \end{aligned} \quad (7)$$

where \mathbf{u}_t^k is the predicted logits of the decoder \mathcal{D}_2 at the t -

th step. The decoder \mathcal{D}_2 iteratively takes the prototype given by the encoder \mathcal{E}_2 as input to generate the final translation \mathbf{Y}^K . With total K iterations, the outputs of encoder \mathcal{E}_1 and \mathcal{E}_2 compose the prototype sequence $\mathbf{E} = \{\mathbf{E}^k\}_{k=0}^K$, where $\mathbf{E}^k = \{e_t^k\}_{t=1}^{T_f}$ is the semantic feature at the k -th iteration. Accordingly, the decoder \mathcal{D}_1 and \mathcal{D}_2 generate $K + 1$ translations $\mathbf{Y} = \{\mathbf{Y}^k\}_{k=0}^K$, where $\mathbf{Y}^k = \{\mathbf{y}_t^k\}_{t=1}^{T_{y,k}}$ is the translation at the k -th iteration during training. Note that after K iterations, the decoder \mathcal{D}_2 obtains the converged prototype \mathbf{E}^K and generates the translation only once in inference.

3.5. Training Objective

We introduce two kinds of losses in the training period of the IP-SLT system. Firstly, the cross entropy loss is adopted to supervise the final generated sentence. Secondly, we put forward an iterative distillation loss. As the decoder \mathcal{D}_1 and \mathcal{D}_2 generate translation sequence $\mathbf{Y} = \{\mathbf{Y}^0, \mathbf{Y}^1, \dots, \mathbf{Y}^K\}$, we naturally divide the sequence \mathbf{Y} into an initial prediction, $K - 1$ intermediate predictions, and the final prediction according to the order of iteration. Conceptually, the $K - 1$ intermediate predictions are regarded as the student model and distill knowledge from the final prediction which is regarded as the teacher model.

Cross entropy loss. As mentioned above, the IP-SLT generates the translation sequence based on the conditional probability provided by the decoder \mathcal{D}_1 and \mathcal{D}_2 . The cross-entropy loss [43] is computed with the ground truth from N training samples and the outputs of the decoder \mathcal{D}_1 and \mathcal{D}_2 . Its training objective in the proposed approach is to maximize the log-likelihood which is equal to minimizing the cross entropy loss formulated as:

$$L_{CE,k} = -\log p^{k-th}(\hat{\mathbf{Y}}|\mathbf{X}), \quad (8)$$

where $\hat{\mathbf{Y}}$ denotes translation annotation. We apply this at the output of the initialization module and the K -th output of the refinement module.

Iterative distillation loss. Since the KL (Kullback-Leibler) divergence loss can affect the teacher’s networks to the student’s networks, we compute it between the $K - 1$ intermediate predictions and the final prediction. As the final translation is based on the previous prototypes, we are able to get better translations by approximating more vital characterization capabilities of middle prototypes. The iterative distillation loss (IDL) is formulated as:

$$L_{IDL} = \sum_{k=1}^{K-1} KL(\mathbf{U}^k, \mathbf{U}^K), \quad (9)$$

where \mathbf{U}^k and \mathbf{U}^K are the predicted output of the decoder \mathcal{D}_2 at the k -th and K -th iteration, respectively. By computing the Kullback-Leibler divergence between the \mathbf{U}^k and

\mathbf{U}^K , the IP-SLT is encouraged to approximate the performance of the final iteration. We apply this at the $K - 1$ shallow outputs of the refinement module.

Overall, the loss function of our IP-SLT is formulated as:

$$L = L_{CE,0} + L_{CE,K} + \lambda \cdot L_{IDL}. \quad (10)$$

Under the guidance of CE loss, we first train the initial generation module until convergence as a warm start, and then apply the loss function as Equ. (10) for optimizing the IP-SLT system in an end-to-end manner.

3.6. Inference Strategy

Since the refinement process only involves in the encoder \mathcal{E}_2 , the initialization and iterative refinement modules behave differently during training and inference. The decoder \mathcal{D}_1 of the initialization module is leveraged to provide guidance for the encoder \mathcal{E}_1 during the training process, while it is not required in inference. Similarly, the autoregressive decoder \mathcal{D}_2 of the iterative refinement module generates translations for each iteration during training, while it just decodes the final translation once in inference.

In inference, given a sign language video to be translated, the feature extraction module first converts it into a visual feature. The encoder \mathcal{E}_1 of the initialization module transforms it into the raw prototype. Then, the encoder \mathcal{E}_2 of the refinement module further iteratively refines it by fusing it with the original visual feature. Finally, the decoder \mathcal{D}_2 generates the translation based on the final prototype.

4. Experiments

4.1. Experimental Setup

Dataset. We evaluate our approach on two public sign language translation datasets, *i.e.*, PHOENIX14T [5] and CSL-Daily [49]. Both datasets provide gloss-level and spoken-sentence-level annotations. The PHOENIX14T dataset [5] is the first large-scale neural SLT dataset created with 9 German sign language interpreters. The dataset is split into a training set (7, 096), a development set (519), and a test set (642). The CSL-Daily dataset [49] is a Chinese SLT dataset containing 20, 654 annotated sign language videos. We follow the previous experimental setting [49] and split it into the training, development, and test set.

Evaluation metrics. Following the work [49], we quantitatively assess the quality of translations according to the BLEU- N [32] and ROUGE [38]. The BLEU- N (N ranges from 1 to 4) cares more about the accuracy of the predicted translation while the ROUGE cares more about the consistency of sentences. For both evaluation metrics, a higher value indicates a better performance.

Training settings. We implement our approach on Pytorch [33]. The encoder and decoder of the initialization

Methods	LLM	Dev					Test				
		ROUGE	BLEU-1	BLEU-2	BLEU-3	BLEU-4	ROUGE	BLEU-1	BLEU-2	BLEU-3	BLEU-4
Joint-SLRT [7]	✗	-	47.26	34.40	27.05	22.38	-	46.61	33.73	26.19	21.32
PET [23]	✗	-	-	-	-	-	49.97	49.54	37.19	29.30	24.02
BN-TIN-Transf. [49]	✗	46.87	46.90	33.98	26.49	21.78	46.98	47.57	34.64	26.78	21.68
STMC [50]	✗	48.24	47.60	36.43	29.18	24.09	46.65	46.98	36.09	28.70	23.65
IP-SLT	✗	54.43	54.10	41.56	33.66	28.22	53.72	54.25	41.51	33.45	27.97
MMTLB [9]	✓	53.10	53.95	41.12	33.14	27.61	52.65	53.97	41.75	33.84	28.39
TwoStream-SLT [10]	✓	54.08	54.32	41.99	34.15	28.66	53.48	54.90	42.43	34.46	28.95

Table 1. Performance comparison of IP-SLT with methods for SLT on PHOENIX-2014T. ‘LLM’ denotes adopting pre-trained large language models.

Methods	LLM	Dev					Test				
		ROUGE	BLEU-1	BLEU-2	BLEU-3	BLEU-4	ROUGE	BLEU-1	BLEU-2	BLEU-3	BLEU-4
SL-Luong [5]	✗	34.28	34.22	19.72	12.24	7.96	34.54	34.16	19.57	11.84	7.56
Joint-SLRT [7]	✗	37.06	37.47	24.67	16.86	11.88	36.74	37.38	24.36	16.55	11.79
BN-TIN-Transf. [49]	✗	37.29	40.66	26.56	18.06	12.73	37.67	40.74	26.96	18.48	13.19
IP-SLT	✗	44.33	45.26	31.77	22.87	16.74	44.09	44.85	31.50	22.66	16.72
MMTLB [9]	✓	53.38	53.81	40.84	31.29	24.42	53.25	53.31	40.41	30.87	23.92
TwoStream-SLT [10]	✓	55.10	55.21	42.31	32.71	25.76	55.72	55.44	42.59	32.87	25.79

Table 2. Performance comparison of IP-SLT with methods for SLT on CSL-Daily. ‘LLM’ denotes adopting pre-trained large language models.

and refinement module consist of 3 layers, respectively. The dimension of the feed-forward network is 2048. The visual feature of the STMC [50] is 1024-dimension while the visual feature of BN-TIN-Transf. [49] and VAC [31] are 512-dimension. To alleviate over-fitting, we set dropout and attention head to 0.1 and 8, respectively. The training optimizer is Adam [24]. During training, the learning rate is fixed to 5×10^{-5} . To ensure the features provided by the previous prototype and original sign video are fully utilized, we apply the drop-net [51] during training, which effects Equ. (5). During training, for any layer in the encoder \mathcal{E}_2 , with probability β , the hidden state \tilde{E}_i^k in Equ. (5) is the output of the self-attention sub-layer attn_s ; with probability $1 - \beta$, it is the output of the cross-attention sub-layer attn_c . In inference, the hidden state \tilde{E}_i^k is computed as Equ. (5).

Inference details. In inference, we use the beam search strategy [45] to improve the decoding accuracy. For the PHOENIX-2014T dataset and the CSL-Daily dataset, we set the beam search width and the length penalty to 3 and 1.0, respectively. To reduce the computational complexity, our IP-SLT just decodes once in inference for each input sign language video.

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

We compare the proposed IP-SLT with the previous SLT systems on two public benchmarks, *i.e.*, PHOENIX14T [5] and CSL-Daily [49], and the performance of our IP-SLT is shown in Tab. 1 and Tab. 2, respectively. For PHOENIX14T and CSL-Daily dataset, we adopt the STMC [50] and BN-

TIN-Transf. [49] as the baseline, respectively. Our IP-SLT follows the sign-to-text (S2T) paradigm, which directly transforms the sign language video into translation. Note MMTLB [9] and TwoStream-SLT [10] adopt pre-trained large-scale language models that leverage more model parameters and extra resources than IP-SLT.

By combining all proposed components together, our IP-SLT achieves substantial improvements against the baseline. The IP-SLT achieves 28.22 and 16.74 BLEU-4 scores on the DEV set of PHOENIX14T and CSL-Daily, respectively. The quantitative results demonstrate that our IP-SLT achieves promising performance improvements. Our IP-SLT delivers promising performance gains on DEV and test sets by leveraging the iterative refinement process. The results prove the advantage of aggregating the previous translation progress and the original visual representation, which distinguishes our IP-SLT from previous SLT systems.

4.3. Ablation Studies

In this section, we put forward several ablation studies on the DEV set of PHOENIX-2014T. Unless otherwise specified, we adopt the STMC [50] as the baseline for the following experiments.

Impact of network architecture. The main difference between our proposed method and the existing work is to leverage the previous information as an additional clue to enhance the current prototype. To evaluate the effectiveness of each component, we gradually add the refinement module and the iterative distillation loss to the baseline SLT sys-

Setting	ROUGE	BLEU-1	BLEU-2	BLEU-3	BLEU-4
Baseline	50.10	50.21	37.12	29.41	24.31
+Refinement	51.22	51.04	38.40	30.61	25.39
+IDL	54.43	54.10	41.56	33.66	28.22
6-6 Layers	50.49	50.58	37.57	29.64	24.53

Table 3. Effect of our proposed components. ‘Refinement’ denotes applying the refinement process. ‘IDL’ denotes applying the iterative distillation loss. ‘6-6 layers’ denotes enlarging the encoder and decoder of the baseline system from 3 to 6 layers.

Model	RM	ROUGE	BLEU-1	BLEU-2	BLEU-3	BLEU-4
STMC	w/o	50.10	50.21	37.12	29.41	24.31
	w/	54.43	54.10	41.56	33.66	28.22
BN-TIN-Transf.	w/o	47.41	47.99	34.94	27.33	22.35
	w/	52.06	52.06	39.01	31.08	25.69
VAC-Transf.	w/o	49.48	50.01	37.00	29.12	23.91
	w/	53.68	53.60	41.28	33.47	28.07

Table 4. Generalization of IP-SLT. ‘RM’ denotes leveraging the refinement process. ‘w/’ and ‘w/o’ denote the baseline SLT system with and without a refinement process, respectively.

λ	ROUGE	BLEU-1	BLEU-2	BLEU-3	BLEU-4
0	51.30	50.32	38.01	30.29	25.14
5	53.07	52.31	40.08	32.29	27.02
10	53.71	53.51	40.92	32.94	27.47
15	54.43	54.10	41.56	33.66	28.22
20	54.42	54.16	41.51	33.44	27.87

Table 5. The weight λ of iterative distillation loss to CE loss.

tem. Directly applying the refinement process to the baseline delivers a performance gain of 1.08 BLEU-4. We further apply the iterative distillation loss to improve the performance. The results suggest that adding distillation supervision can be helpful with a gain of 2.83 BLEU-4. Besides, to keep the number of parameters unchanged, we enlarge the depth from 3 to 6 layers and evaluate the performance. Naively enlarging the model scale slightly improves the performance (+0.22 BLEU-4). The results are shown in Tab. 3.

Generalization of the IP-SLT. We conduct three sets of experiments by changing the visual backbone to evaluate the generalization of the proposed IP-SLT approach in Tab. 4. Specifically, the BN-TIN-Transf. [49] uses a basic CNN network to get the dense representation of sign video. The STMC [50] extracts and aligns the key parts of sign language expression to achieve better performance. The VAC [31] proposes two auxiliary supervision methods to enhance the feature extraction module. VAC-Transf. replaces the feature extractor in BN-TIN-Transf. [49] with the visual backbone of VAC [31]. Applying our proposed IP-SLT methods to the BN-TIN-Transf., VAC-Transf. and STMC, we achieve 25.69, 28.07, and 28.22 BLEU-4 scores

β	ROUGE	BLEU-1	BLEU-2	BLEU-3	BLEU-4
0.0	50.10	50.21	37.12	29.41	24.31
0.2	54.35	53.81	41.40	33.52	28.02
0.4	54.39	54.02	41.24	33.19	27.70
0.5	54.43	54.10	41.56	33.66	28.22
0.6	53.96	53.63	40.94	33.10	27.72
0.8	54.31	53.39	40.92	32.98	27.53
-	50.97	50.67	38.24	30.52	25.41

Table 6. The weight β of the original visual feature to the previous prototype. ‘0.0’ denotes the assessment of the baseline. ‘-’ denotes setting β as 0.5 without using the drop-net [51].

Setting	ROUGE	BLEU-1	BLEU-2	BLEU-3	BLEU-4
None	50.10	50.21	37.12	29.41	24.31
Con-input	54.37	53.83	41.10	32.97	27.39
Con-feature	53.97	52.76	40.56	32.81	27.47
Add-feature	54.43	54.10	41.56	33.66	28.22

Table 7. Effect of the different refinement methods. ‘Con-input’ denotes directly concatenating the original visual feature and previous prototype as input. ‘Con-feature’ denotes concatenating the original feature and the previous feature given by the cross-attention mechanism in each layer. On top of that, ‘Add-feature’ denotes changing it to an addition operation.

on the DEV set, surpassing the baselines by 3.34, 4.16, 3.91, respectively. Using a high-quality visual backbone delivers further quality gains.

Impact of λ . In our experiments, the weight λ of iterative distillation loss is set to 15. It is a hyper-parameter that is designed to balance the effect of cross-entropy loss and the iterative distillation loss. We conduct experiments by varying the weight λ . Tab. 5 shows that our IP-SLT achieves the best performance when the weight λ is set to 15.

Impact of β . In the above experiments, the weight β is fixed to 0.5. The weight β represents the importance of the previous prototype compared with the original visual feature. As a hyper-parameter of our proposed methods, the weight β is examined with a set of different values in Tab. 6. When the weight is 0.5, the performance is the highest one. This indicates that the previous prototype to SLT is as important as the original visual feature. Besides, to fully use the previous prototype, the drop-net is required.

Impact of refinement method. We also examine a set of refinement methods for IP-SLT considering the fusion mechanism as a key part of our proposed method in Tab. 7. Directly concatenating the original visual feature and previous prototype in the time dimension improves the performance of SLT from 24.31 to 27.39. We further add the cross-attention mechanism to explicitly leverage the useful information from the previous prototype. And then, we concatenate the representation from the original visual feature and the representation from the previous prototype in fea-

K	ROUGE	BLEU-1	BLEU-2	BLEU-3	BLEU-4
0	50.10	50.21	37.12	29.41	24.31
1	51.22	51.04	38.40	30.61	25.39
2	53.73	53.39	40.76	32.98	27.66
3	54.43	54.10	41.56	33.66	28.22
4	54.63	53.91	41.40	33.53	28.01

Table 8. Effect of the iteration number K in the iterative refinement module.

Model	SW	I-P(M)	T-P(M)	FLOPs(B)	ROUGE	BLEU-4
STMC	-	92.6	92.6	28.7	50.10	24.31
	w/	128.7	172.6	32.6	54.43	28.22
	w/o	238.6	357.9	32.6	54.81	28.48
BN-TIN-Transf.	-	28.3	28.3	9.6	47.41	22.35
	w/	37.7	53.4	10.6	52.06	25.69
	w/o	66.1	110.0	10.6	52.24	25.96
VAC-Transf.	-	28.3	28.3	9.6	49.48	23.91
	w/	37.7	53.4	10.6	53.68	28.07
	w/o	66.1	110.0	10.6	54.82	28.27

Table 9. Comparing the baseline and non-shared weight prototype refinement method with IP-SLT. ‘SW’ denotes the sharing weight across all iterations. ‘-’ denotes the baseline SLT system. ‘w/’ and ‘w/o’ denote the refinement module in different iterations with and without shared parameters, respectively. ‘I-P’ and ‘T-P’ denote the parameter amount calculated in inference and during training, respectively.

ture dimension, which delivers performance gains of 3.16 BLEU-4. On top of that, changing the concatenating process to an element-wise addition operation achieves a further quality gain of 3.91.

Impact of iteration number K . The iteration number K is an important hyper-parameter and is fixed to 3 in the previous experiments. We conduct experiments with different iteration numbers to explore the effect of iteration number K . Tab. 8 shows that there is a culmination in the performance at iteration 3. Before this culmination, the performance improvement of our proposed method increase fast. When the iteration number is bigger than 3, the performance of our method is slightly weakened (-0.21 BLEU-4).

Impact of sharing weights between different iterations.

In Tab. 9, we examine the storage and translation quality of our proposed IP-SLT with different baselines *i.e.*, STMC, BN-TIN-Transf, and VAC-Transf. In order to ensure the efficiency of the model, our refinement module shares weight across all iterations. Since the feature extraction module of each group is identical, the parameters of the feature extraction module are not included in the calculation of the storage of different models. As in inference, several parts can be removed from IP-SLT without changing the performance, we report the number of parameters in both training and inference, respectively. The results indicate that our iterative process causes acceptable overhead with re-

Type	Text
GT	ich wünsche ihnen einen schönen abend und machen sie es gut.
Baseline	jetzt wünsche ich ihnen noch einen schönen abend.
Our	ihnen einen schönen abend und machen sie es gut.
GT	in der neuen woche wird es milder aber es bleibt wechselhaft.
Baseline	dann wird es wieder milder.
Our	in der neuen woche wird es dann wieder milder.

Table 10. Qualitative evaluation. ‘GT’ denotes the spoken language translation annotation. ‘Baseline’ and ‘Our’ denote the translation result of baseline and our IP-SLT, respectively.

markable performance improvements. We further conduct experiments in which the parameters of each iteration are independent. The results demonstrate that the proposed parameter-shared method achieves near performance gain (28.22 v.s 28.48 BLEU-4, 25.69 v.s 25.96 BLEU-4, and 28.07 v.s 28.27 BLEU-4) with non-shared one but leveraging much less parameter number.

Computation comparison with baseline. The FLOPs is a key factor of computation efficiency. We conduct experiments to compare the computation of the proposed IP-SLT method with different baselines. Similarly, we exclude the FLOPs of the feature extraction module and report relevant computational costs in inference. Tab. 9 shows that leveraging the iterative refinement process can cause acceptable computation costs while achieving promising performance.

Case Study. To provide a more intuitionistic view of our proposed method, we list some translation samples of the proposed IP-SLT method in Tab. 10. We observe that based on the previous prototype and original visual representation, the IP-SLT can generate more accurate and fluent sentences.

5. Conclusion

In this work, we propose a new framework IP-SLT which introduces the iterative refinement into a conventional SLT system. With the goal to polish the semantic representation by leveraging the previous results, we present IP-SLT to support the iterative refinement process. The proposed method is differentiable and optimized in an end-to-end manner to achieve its best performance. On top of it, we put forward the iterative distillation loss to further improve the translation quality by leveraging the sequential dependence between the outputs of each iteration. In inference, the autoregressive decoding process is required once to generate the translation based on the final prototype, applying IP-SLT does not significantly affect efficiency. The experimental results demonstrate the effectiveness of the IP-SLT.

Acknowledgments: This work was supported by NSFC under Contract U20A20183 and 62021001. It was also supported by the GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC, and the Supercomputing Center of the USTC.

References

- [1] Namhyuk Ahn, Byungkon Kang, and Kyung-Ah Sohn. Image super-resolution via progressive cascading residual network. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018. 2
- [2] 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 *Computer Vision European Conference*, 2020. 1
- [3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014. 2, 3
- [4] Danielle Bragg, Oscar Koller, Mary Bellard, Larwan Berke, Patrick Boudreault, Annelies Braffort, Naomi Caselli, Matt Huenerfauth, Hernisa Kacorri, Tessa Verhoef, et al. Sign language recognition, generation, and translation: An interdisciplinary perspective. In *International ACM SIGACCESS Conference on Computers and Accessibility*, 2019. 1
- [5] Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, Hermann Ney, and Richard Bowden. Neural sign language translation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2, 5, 6
- [6] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. Multi-channel transformers for multi-articulatory sign language translation. In *Computer Vision–ECCV 2020 Workshops*, 2020. 2
- [7] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. Sign language transformers: Joint end-to-end sign language recognition and translation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 1, 2, 6
- [8] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [9] Yutong Chen, Fangyun Wei, Xiao Sun, Zhirong Wu, and Stephen Lin. A simple multi-modality transfer learning baseline for sign language translation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 2, 6
- [10] Yutong Chen, Ronglai Zuo, Fangyun Wei, Yu Wu, Shujie Liu, and Brian Mak. Two-stream network for sign language recognition and translation. *Advances in Neural Information Processing Systems*, 2022. 2, 6
- [11] Helen Cooper, Brian Holt, and Richard Bowden. Sign language recognition. *Visual Analysis of Humans: Looking at People*, 2011. 1
- [12] Kearsy Cormier, Neil Fox, Bencie Woll, Andrew Zisserman, Necati Cihan Camgöz, and Richard Bowden. ExTOL: Automatic recognition of british sign language using the bsl corpus. In *Workshop on Sign Language Translation and Avatar Technology*, 2019. 1
- [13] 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*, 2017. 1, 3
- [14] Runpeng Cui, Hu Liu, and Changshui Zhang. A deep neural framework for continuous sign language recognition by iterative training. *IEEE Transactions on Multimedia*, 2019. 2, 3
- [15] Hao Feng, Wengang Zhou, Jiajun Deng, Qi Tian, and Houqiang Li. DocScanner: robust document image rectification with progressive learning. *arXiv preprint arXiv:2110.14968*, 2021. 2
- [16] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. In *International Conference on Machine Learning*, 2017. 3
- [17] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *International Conference on Machine Learning*, 2006. 2
- [18] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Rezende, and Daan Wierstra. DRAW: A recurrent neural network for image generation. In *International Conference on Machine Learning*, 2015. 2
- [19] Jun He, Richang Hong, Xueliang Liu, Mingliang Xu, and Qianru Sun. Revisiting local descriptor for improved few-shot classification. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 2022. 1
- [20] Hezhen Hu, Weichao Zhao, Wengang Zhou, and Houqiang Li. Signbert+: Hand-model-aware self-supervised pre-training for sign language understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 1
- [21] Jie Huang, Wengang Zhou, Qilin Zhang, Houqiang Li, and Weiping Li. Video-based sign language recognition without temporal segmentation. In *AAAI Conference on Artificial Intelligence*, 2018. 1
- [22] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3D convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2012. 2
- [23] Tao Jin, Zhou Zhao, Meng Zhang, and Xingshan Zeng. Prior knowledge and memory enriched transformer for sign language translation. In *Findings of the Association for Computational Linguistics*, 2022. 1, 2, 6
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [25] 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. 3
- [26] Oscar Koller, Hermann Ney, and Richard Bowden. Deep hand: How to train a CNN on 1 million hand images when your data is continuous and weakly labelled. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 3
- [27] 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*, 2020. 2
- [28] Huan Ling, Jun Gao, Amlan Kar, Wenzheng Chen, and Sanja Fidler. Fast interactive object annotation with Curve-GCN. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [29] Zichen Liu, Jun Hao Liew, Xiangyu Chen, and Jiashi Feng. DANCE: A deep attentive contour model for efficient instance segmentation. In *IEEE Winter Conference on Applications of Computer Vision*, 2021. 2
- [30] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*, 2015. 1
- [31] Yuecong Min, Aiming Hao, Xiujuan Chai, and Xilin Chen. Visual alignment constraint for continuous sign language recognition. In *IEEE International Conference on Computer Vision*, 2021. 1, 2, 6, 7
- [32] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: a method for automatic evaluation of machine translation. In *Annual Meeting of the Association for Computational Linguistics*, 2002. 5
- [33] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 5
- [34] Sida Peng, Wen Jiang, Huaijin Pi, Xiuli Li, Hujun Bao, and Xiaowei Zhou. Deep snake for real-time instance segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [35] Junfu Pu, Wengang Zhou, and Houqiang Li. Dilated convolutional network with iterative optimization for continuous sign language recognition. In *International Joint Conference on Artificial Intelligence*, 2018. 2
- [36] Junfu Pu, Wengang Zhou, and Houqiang Li. Iterative alignment network for continuous sign language recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 1, 2
- [37] Dongwei Ren, Wangmeng Zuo, Qinghua Hu, Pengfei Zhu, and Deyu Meng. Progressive Image Deraining Networks: A better and simpler baseline. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [38] Lin CY ROUGE. A package for automatic evaluation of summaries. In *Workshop on Text Summarization of ACL*, 2004. 5
- [39] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *Advances in Neural Information Processing Systems*, 2014. 2
- [40] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *Advances in Neural Information Processing Systems*, 2014. 1, 3
- [41] Rachel Sutton-Spence and Bencie Woll. *The linguistics of British Sign Language: an introduction*. Cambridge University Press, 1999. 1
- [42] Shengeng Tang, Dan Guo, Richang Hong, and Meng Wang. Graph-based multimodal sequential embedding for sign language translation. *IEEE Transactions on Multimedia*, 2021. 1
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 1, 2, 3, 4, 5
- [44] Hanjie Wang, Xiujuan Chai, Xiaopeng Hong, Guoying Zhao, and Xilin Chen. Isolated sign language recognition with grassmann covariance matrices. *ACM Transactions on Accessible Computing*, 2016. 1
- [45] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016. 6
- [46] Kayo Yin and Jesse Read. Better sign language translation with stmc-transformer. *arXiv preprint arXiv:2004.00588*, 2020. 2
- [47] Tao Zhang, Shiqing Wei, and Shunping Ji. E2EC: An end-to-end contour-based method for high-quality high-speed instance segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [48] Hao Zhou, Wengang Zhou, and Houqiang Li. Dynamic pseudo label decoding for continuous sign language recognition. In *IEEE International conference on Multimedia and Expo*, 2019. 2
- [49] Hao Zhou, Wengang Zhou, Weizhen Qi, Junfu Pu, and Houqiang Li. Improving sign language translation with monolingual data by sign back-translation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 1, 2, 5, 6, 7
- [50] 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. 2, 6, 7
- [51] Jinhua Zhu, Yingce Xia, Lijun Wu, Di He, Tao Qin, Wengang Zhou, Houqiang Li, and Tie-Yan Liu. Incorporating bert into neural machine translation. *arXiv preprint arXiv:2002.06823*, 2020. 4, 6, 7