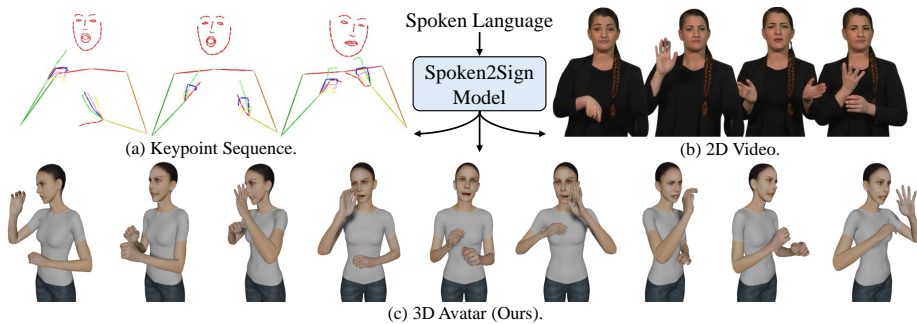


# A Simple Baseline for Spoken Language to Sign Language Translation with 3D Avatars

Ronglai Zuo<sup>1\*</sup>, Fangyun Wei<sup>2\*†</sup>, Zenggui Chen<sup>2</sup>, Brian Mak<sup>1</sup>, Jialong Yang<sup>2</sup>,  
and Xin Tong<sup>2</sup>

<sup>1</sup>The Hong Kong University of Science and Technology    <sup>2</sup>Microsoft Research Asia  
{rzuo,mak}@cse.ust.hk    {fawe,v-zenchen,jiaoyan,xtong}@microsoft.com



**Fig. 1:** Prior works have presented Spoken2Sign translation results through either (a) keypoint sequences [56] or (b) 2D videos [58]. In contrast, we utilize a (c) 3D avatar to display the translation results, enabling the visualization of results from any viewpoint.

**Abstract.** The objective of this paper is to develop a functional system for translating spoken languages into sign languages, referred to as Spoken2Sign translation. The Spoken2Sign task is orthogonal and complementary to traditional sign language to spoken language (Sign2Spoken) translation. To enable Spoken2Sign translation, we present a simple baseline consisting of three steps: 1) creating a gloss-video dictionary using existing Sign2Spoken benchmarks; 2) estimating a 3D sign for each sign video in the dictionary; 3) training a Spoken2Sign model, which is composed of a Text2Gloss translator, a sign connector, and a rendering module, with the aid of the yielded gloss-3D sign dictionary. The translation results are then displayed through a sign avatar. As far as we know, we are the first to present the Spoken2Sign task in an output format of 3D signs. In addition to its capability of Spoken2Sign translation, we also demonstrate that two by-products of our approach—3D keypoint augmentation and multi-view understanding—can assist in keypoint-based sign language understanding. Code and models are available at <https://github.com/FangyunWei/SLRT>.

\* Equal contribution.

† Corresponding author.

## 1 Introduction

Sign languages are the primary means of communication for the deaf. Numerous previous works [8, 9, 39, 71, 72, 79] have focused on sign language translation, with the goal of translating sign languages into spoken languages (Sign2Spoken). However, this paper shifts the focus to the reverse process: translating spoken languages into sign languages (Spoken2Sign) to further bridge the communication gap between the deaf and the hearing.

A majority of prior works [26–28, 53, 54, 57, 62, 68] on Spoken2Sign translation (aka sign language production) focused on expressing translation outcomes through keypoints (Figure 1a). However, the keypoint representations often pose interpretable challenges for signers [50]. With the evolution of generative models [2, 76], several studies [16, 58, 61] have employed these keypoints to animate signer images, subsequently creating a sign video (Figure 1b). However, the 2D video format is prone to blurriness and visual distortions. In this work, we introduce an innovative method for Spoken2Sign translation by utilizing a 3D avatar to represent the translation results (Figure 1c). In contrast to earlier attempts that utilized generative models [51, 59], our method prioritizes enhancing understandability and incorporates a 3D human pose prior with a special emphasis on signing poses, allowing for multi-view representations of the translation results.

We begin with the basic concepts of sign languages:

- *Isolated sign*. An isolated sign is a single gesture comprising handshapes, along with movements of the body and hands. Sometimes, facial expressions also play a role in conveying information.
- *Continuous sign*. It is a sequence of several isolated signs along with co-articulations.
- *Co-articulation*. It refers to the movement of the hands and body between two adjacent signs in a continuous sign.
- *Gloss*. A gloss refers to the label assigned to a specific isolated sign. It is usually represented as a word or phrase.
- *Gloss sequence*. A label sequence for a continuous sign. In our system, it serves as an intermediate representation.
- *Text*. It refers to the translation of a continuous sign. Generally, the *gloss sequence* does not equate with the *text*, as the linguistic rules, *e.g.*, word order, between a sign language and its corresponding spoken language can be different.

As shown in Figure 2, our Spoken2Sign translation baseline includes three steps: 1) dictionary construction; 2) 3D sign estimation for each entry in the dictionary; and 3) Spoken2Sign translation in a retrieve-then-connect paradigm.

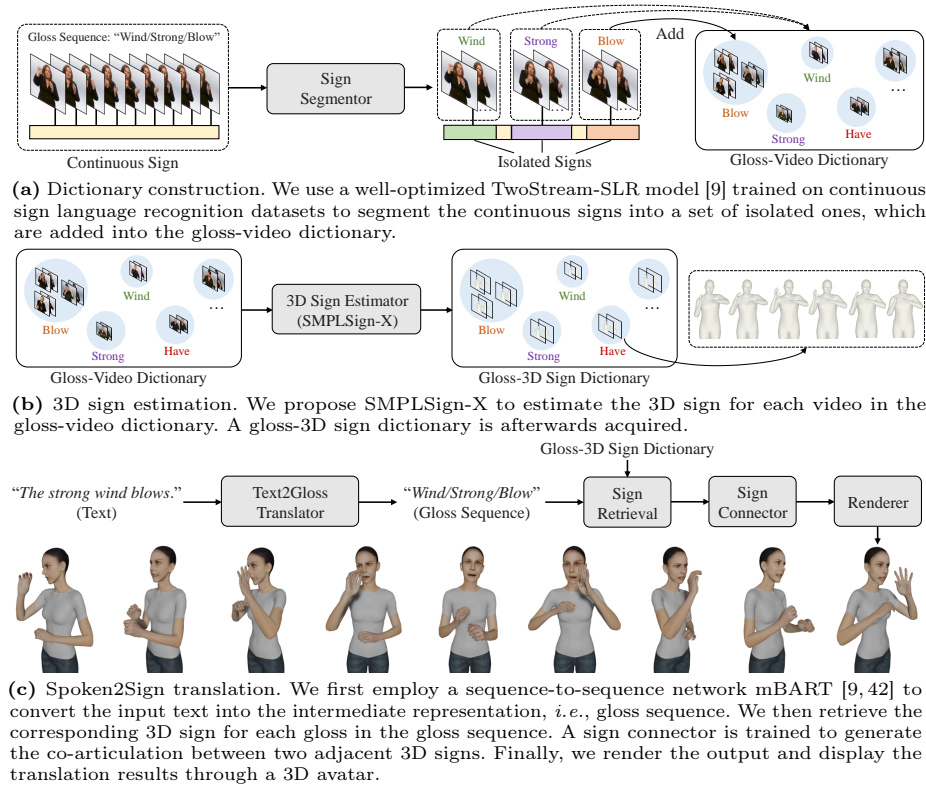
**Dictionary Construction.** As shown in Figure 2a, we initially create a gloss-video dictionary comprising  $M$  glosses, each of which may contain multiple isolated sign videos that express the same meaning. Existing sign language recognition (SLR) datasets are typically divided into: 1) datasets for isolated SLR (ISLR), *e.g.*, WLASL [38] and MSASL [32]; 2) datasets for continuous SLR (CSLR), *e.g.*, Phoenix-2014T [4] and CSL-Daily [80]. The objective of this work is to develop a

Spoken2Sign system, that translates the input text into a sequence of 3D signs, using glosses as intermediate representations. Although existing ISLR datasets are inherent sign language dictionaries, they lack parallel text-gloss sequence data, posing a challenge for the essential text-to-gloss sequence translation. An alternative is to leverage the CSLR datasets; nevertheless, these datasets do not include a dictionary. One solution is utilizing an external dictionary, as demonstrated in FS-Net [58]. Another approach employs a CSLR model trained with a connectionist temporal classification (CTC) loss [20], effectively segmenting a continuous sign language video into its constituent signs [66, 85]. The isolated signs produced by this latter method are more suitable for subsequent sign connection since they exclude meaningless motions, such as raising and lowering hands, which may be present in the beginnings and endings of signs in the external dictionary. In this paper, we adopt the state-of-the-art CTC-based CSLR model, TwoStream-SLR [9], as the sign segmentor to construct a dictionary.

**3D Sign Estimation.** We currently have a gloss-video dictionary at our disposal. Our next step converts each sign video within this dictionary into a 3D sign. This transformation is essential to mimicking co-articulations more naturally and stitching two signs together in 3D space during Spoken2Sign translation. Inspired by recent advances in 3D whole-body parametric models and monocular 3D whole-body reconstruction [31, 49, 69], we introduce a sign-language-specific version of SMPLify-X [49]: SMPLSign-X. This approach is dedicated to estimating 3D signs from monocular sign videos. The vanilla SMPLify-X estimates a holistic and expressive 3D human representation by optimizing shape, pose, and expression parameters. To enhance its application to 3D sign estimation, we incorporate a series of improvements including promoting temporal consistency and involving a more effective 3D pose prior tailored for sign language avatars. Following these improvements, we construct a gloss-3D sign dictionary. The entire process is depicted in Figure 2b.

**Spoken2Sign Translation.** Figure 2c illustrates the core idea behind the proposed Spoken2Sign translation. Glosses serve as a link between sign languages and spoken languages. Therefore, training a text-to-gloss-sequence (Text2Gloss) model is necessary to enable Spoken2Sign translation. Following [7, 9], we adopt mBART [42] as our Text2Gloss model due to its promising sequence-to-sequence translation capability. For each gloss in the gloss sequence, we retrieve its corresponding 3D sign from the gloss-3D sign dictionary. A significant challenge arises in seamlessly stitching two adjacent signs, as it is crucial to imitate the co-articulations between them to produce visually appealing translation results. Prior works have either overlooked this challenge [60, 75] or have addressed it through fixed-length interpolation over 2D keypoints [58]. In this paper, we develop a lightweight sign connector that can dynamically predict the duration of each co-articulation, allowing it to seamlessly stitch adjacent signs in 3D space. Our experiments reveal that this simple connector plays a key role in enhancing the overall performance of the Spoken2Sign system. The final translation results are obtained by rendering the output of the sign connector with a sign avatar.

In summary, the contributions of this work are:



**Fig. 2:** Overview of our methodology. It consists of (a) dictionary construction; (b) 3D sign estimation; (c) Spoken2Sign translation. Sign videos are from Phoenix-2014T [4], a German sign language benchmark.

- We present a simple yet effective baseline for Spoken2Sign translation. To the best of our knowledge, this is the first work of developing a practical Spoken2Sign system that utilizes a 3D avatar to display the translation results.
- Considering the unique characteristics of sign languages, we propose SMPLSign-X, a novel method for 3D sign estimation. Our method aims to develop comprehensive 3D sign language dictionaries for well-established benchmarks including Phoenix-2014T and CSL-Daily. We will release these dictionaries to facilitate future research.
- Our Spoken2Sign system sets a new state-of-the-art performance in back-translation evaluation.
- Besides demonstrating the capability of Spoken2Sign translation, we also show that two by-products of our approach significantly enhance keypoint-based models for sign language understanding.

## 2 Related Work

**Sign Language Understanding.** There are several research directions in sign language understanding, such as sign language recognition (SLR), sign language translation (SLT), sign spotting [1, 45, 46, 64], and sign language retrieval [11, 14]. The objective of SLR is to transcribe a sign video into its constituent glosses. It can be categorized into isolated SLR (ISLR) [22, 23, 29, 32, 38, 40, 77, 84] and continuous SLR (CSLR) [8–10, 21, 24, 44, 47, 78, 81–83]. The former is a classification task aimed at predicting the gloss of an isolated sign, while the latter focuses on recognizing a sequence of signs in a video and generating the corresponding gloss sequence. SLT takes a step further—it aims at translating sign languages into spoken languages. Existing works [4–6, 8, 9, 18, 39, 73, 74] attempt to formulate the SLT task as a neural machine translation problem, using a visual encoder and a language model. Inspired by the recent success of transferring a pre-trained language model to SLT [9], we adopt mBART [42] as our Text2Gloss translator.

**Spoken2Sign Translation.** In contrast to sign language understanding, this area remains under-explored. Most previous works [25–28, 53, 54, 57, 62, 68] focus on expressing translation results through keypoints, which are often difficult for signers to understand [50]. With the advances of generative models, several studies [15, 16, 58, 61] leverage these keypoints to animate sign videos. However, these approaches often encounter challenges such as blurriness and visual distortions. Some papers [35, 43] proposed the use of avatars to display sign languages, but they were too early to benefit from the advanced parametric 3D human models [49]. In contrast, our approach incorporates a 3D signer pose prior, enhancing both the understandability and the temporal consistency of the animations.

FS-Net [58] represents a notable work in the field. Our approach diverges from FS-Net in four key aspects: 1) we are the pioneers in utilizing 3D avatars as outcomes in Spoken2Sign translation; 2) unlike FS-Net, which relies on an external dictionary, our dictionary is built by segmenting continuous sign videos into isolated ones, which are more suitable for the subsequent sign connection; 3) our lightweight sign connector is designed to flexibly predict the lengths of co-articulations and stitch signs in 3D space, avoiding unrealistic fixed-length interpolation over 2D keypoints; 4) we demonstrate that two by-products of 3D sign representations can significantly aid in the understanding of sign languages.

**3D Whole-Body Estimation.** Recent monocular 3D whole-body estimation approaches adopt parametric models such as Frank [31] and SMPL-X [49] for 3D body representation. These approaches can be divided into fitting-based and regression-based methods. The former utilizes optimization algorithms to fit the parametric model to the given 2D observations, such as body keypoints and segmentation masks [36, 49, 67, 69], while the latter directly predicts model parameters without iterative optimization [3, 33, 41, 70], although they often rely on precise SMPL-X annotations [3]. Parametric models have also been applied in sign language research [17, 22, 23, 37]. However, none of these studies have developed a complete Spoken2Sign pipeline. In this work, we tailor the widely-used SMPL-X parametric model by incorporating sign-language-specific priors for monocular 3D sign estimation and introduce a functional Spoken2Sign system.

### 3 Methodology

This section details our methodology, which comprises three stages: dictionary construction (Section 3.1), 3D sign estimation (Section 3.2), and Spoken2Sign translation (Section 3.3). Additionally, we discuss two by-products derived from 3D signs in Section 3.4. An overview is depicted in Figure 2.

#### 3.1 Dictionary Construction

Our Spoken2Sign system is built upon the availability of a sign dictionary, which consists of gloss-video pairs. Nevertheless, existing datasets [4, 80] do not provide such dictionaries. Inspired by the capability of a well-trained CSLR model to segment a continuous sign video into a collection of isolated signs by finding the optimal alignment path through the CTC forced alignment algorithm [13, 20, 52, 66], we adopt the state-of-the-art CSLR model, TwoStream-SLR [9], as the sign segmentor. This model is used to build a sign dictionary containing the segmented isolated signs for each dataset, as shown in Figure 2a. Subsequently, we acquire a sign dictionary containing  $M$  glosses. We also generate a set of co-articulations (which are transitions between two adjacent signs, corresponding to the blank class in the CTC loss [20]) to train our sign connector. More details can be found in the supplementary materials.

#### 3.2 3D Sign Estimation

We introduce the process of estimating the 3D representation for each isolated sign from a monocular video in the dictionary, as shown in Figure 2b.

**Preliminaries of SMPLify-X and SMPL-X.** SMPLify-X [49] is a widely used method to estimate a 3D representation of human body pose, hand pose, and facial expression from a single monocular image. The yielded 3D human representation is termed as a SMPL-X model [49], which is defined by a series of learnable parameters. These parameters include global orientation  $\zeta \in \mathbb{R}^3$ , body shape  $\beta \in \mathbb{R}^{10}$ , facial expression  $\psi \in \mathbb{R}^{10}$ , and body pose  $\theta \in \mathbb{R}^{3N}$ , where  $N = 54$  denotes the number of joints. Note that pose parameters represent the relative axis-angle rotations to the parent joints defined in a kinematics map. Using these parameters, the SMPL-X model could produce a mesh comprising 10,475 vertices and a set of 118 3D joints  $\mathcal{D} \in \mathbb{R}^{118 \times 3}$ . For simplicity, we use  $\mathcal{D}_i \in \mathbb{R}^3$  ( $1 \leq i \leq 118$ ) to denote the  $i$ -th joint in  $\mathcal{D}$ . Note that only 54 out of the 118 joints are associated with the pose parameters  $\theta$ .

To fit the SMPL-X model to a single monocular image, we first use HRNet [65] pre-trained on COCO-Wholebody [30], to estimate the image’s 2D keypoints  $\mathcal{K} \in \mathbb{R}^{118 \times 2}$ . We employ only a subset of the COCO-Wholebody keypoints for aligning the joints defined by SMPL-X. Subsequently, SMPLify-X [49] seeks to minimize the following objective function by optimizing  $\zeta$ ,  $\beta$ ,  $\psi$  and  $\theta$ :

$$\mathcal{L} = \mathcal{L}_{joint} + \mathcal{L}_{prior} + \mathcal{L}_{penetration}, \quad (1)$$

where  $\mathcal{L}_{joint}$  is the major loss function that minimizes the distance between the 2D keypoints  $\mathcal{K}$  and the projected keypoints  $P(\mathcal{D})$ ;  $\mathcal{L}_{prior}$  denotes a combination of losses that incorporate prior knowledge of hand pose, facial pose, body shape, and facial expressions, and penalize extreme body states;  $\mathcal{L}_{penetration}$  is a regularization term designed to prevent the SMPL-X model from penetrations and self-collisions. In SMPLify-X,  $\mathcal{L}_{joint}$  is formulated as:

$$\mathcal{L}_{joint} = \frac{1}{|\mathcal{J}|} \sum_{i \in \mathcal{J}} \gamma_i \omega_i \ell_r(P(\mathcal{D}_i) - \mathcal{K}_i), \quad (2)$$

where  $\mathcal{J}$  represents the set of 118 3D joints;  $P(\cdot)$  is a function that projects each 3D joint  $\mathcal{D}_i \in \mathbb{R}^3$  from the world coordinate to image coordinate;  $\mathcal{K}_i$  is the associated 2D keypoint (pseudo ground truth) of  $\mathcal{D}_i$ ;  $\omega_i$  is the confidence (yielded by HRNet) of  $\mathcal{K}_i$ ;  $\gamma_i$  is the pre-defined weight of joint  $\mathcal{D}_i$ ;  $\ell_r$  represents a robust Geman-McClure loss function [19]. More details can be found in [49].

**SMPLSign-X.** We enhance SMPLify-X by adapting its input from a single monocular *image* to a sign *video*. Additionally, we consider the unique properties of sign languages to further improve estimation quality. The new estimator is named SMPLSign-X. Our improvements are based on three observations: 1) optimization targets are absent for joints that do not appear in sign videos (*e.g.*, the lower body and a dropping hand); 2) the upper body of a signer remains upright during signing; 3) independently fitting each frame to the SMPL-X model results in temporal inconsistencies and visually unsatisfactory outcomes. To address these issues, we define the objective function for SMPLSign-X as:

$$\mathcal{L} = \mathcal{L}_{joint} + \mathcal{L}_{prior} + \mathcal{L}_{penetration} + \lambda_1 \mathcal{L}_{unseen} + \lambda_2 \mathcal{L}_{upright} + \lambda_3 \mathcal{L}_{smooth}. \quad (3)$$

$\mathcal{L}_{unseen}$  (Eq. 4) is a regularization term that pushes the unseen keypoints to approach those of the rest pose. We use confidence scores predicted by the pre-trained HRNet to identify the unseen keypoints. Concretely, for a 2D keypoint  $\mathcal{K}_i$  and its confidence score  $\omega_i$ , we regard  $\mathcal{K}_i$  as an unseen keypoint if  $\omega_i < \lambda$ , where  $\lambda$  is a pre-defined threshold, and  $\lambda = 0.65$  by default. SMPLify-X provides joint mappings between the 2D keypoints  $\mathcal{K}$  and the pose parameters  $\theta$ . Therefore, we can easily identify the set of unseen joints ( $\mathcal{J}_{unseen}$ ) in the  $\theta$  space. We use  $\hat{\theta}$  to denote the pose parameters of the rest pose, which are frozen during training.

$$\mathcal{L}_{unseen} = \sum_{i \in \mathcal{J}_{unseen}} \ell_r(\theta_i - \hat{\theta}_i). \quad (4)$$

$\mathcal{L}_{upright}$  (Eq. 5) denotes a regularization term for encouraging an upright posture in the upper body. To accomplish this, we define a set of keypoints  $\mathcal{J}_{upright}$ , including the neck and pelvis keypoints within the context of  $\theta$  space. Our goal is to preserve depth consistency across all keypoints in  $\mathcal{J}_{upright}$ , as specified by the loss in Eq. 5, where  $d$  represents depth.

$$\mathcal{L}_{upright} = \sum_{i, j \in \mathcal{J}_{upright}} \ell_r(d_i - d_j). \quad (5)$$

Finally, given the pose parameters  $\theta_i^{pre}$  of the previous frame, we use Eq. 6 to preserve the temporal consistency for the current frame.  $\mathcal{J}_\theta$  is the set containing all joints in the  $\theta$  space.  $\gamma_i$  denotes the weight of the  $i$ -th joint. These joint weights are pre-defined by SMPLify-X [49].

$$\mathcal{L}_{smooth} = \sum_{i \in \mathcal{J}_\theta} \gamma_i \ell_r(\theta_i - \theta_i^{pre}). \quad (6)$$

We estimate the 3D representation for each sign video from the gloss-video dictionary using the objective function in Eq. 3, on a frame-by-frame basis. Subsequently, we construct a gloss-3D sign dictionary.


### 3.3 Spoken2Sign Translation

As shown in Figure 2c, our Spoken2Sign translation pipeline primarily consists of three components: 1) a Text2Gloss translator that translates the input text into a gloss sequence; 2) a sign connector, which stitches two adjacent 3D signs together; and 3) a rendering module for producing the final animated sign avatar.

**Text2Gloss Translator.** Inspired by the recent success of sign language translation (SLT) [9], we adopt mBART [42] as our Text2Gloss translator. Existing SLT benchmarks [4, 80] typically provide annotations of text-gloss-sequence pairs. In the vein of traditional SLT models, which train a Gloss2Text network using gloss sequences as inputs and texts as outputs, we train a Text2Gloss translator by simply reversing the inputs and outputs. Our Text2Gloss translator achieves a promising BLEU-4 score of 29.27/31.88 on the dev set of Phoenix-2014T/CSL-Daily. More details are available in the supplementary materials.

**Sign Retrieval.** For each gloss predicted by the Text2Gloss translator, we retrieve its corresponding 3D sign from the gloss-3D sign dictionary. Since a single gloss may be associated with multiple 3D signs, we develop a retrieval strategy to identify the optimal one. Specifically, we train an ISLR model [84] on all instances in the dictionary. During retrieval, we feed candidate signs into the ISLR model and select the sign with the highest confidence for the gloss query.

**Sign Connector.** As shown in Figure 3, we present a sign connector to predict the duration of co-articulation between two adjacent 3D signs when they are



**Fig. 3:** Illustration of the sign connector. The objective is to predict the duration of the co-articulation between two adjacent 3D signs,  $S_{k-1}$  and  $S_k$ , followed by generating the co-articulation through interpolation in the 3D joint space.



stitched together. As described in Section 3.1, we generate a set of co-articulations, each of which is denoted as a triplet  $(\mathcal{D}_{\mathcal{J}_{SC}}^{pre}, L, \mathcal{D}_{\mathcal{J}_{SC}}^{next})$ , where  $L$  is the duration of the co-articulation,  $\mathcal{J}_{SC}$  denotes the set of 3D joints used in our sign connector, and  $\mathcal{D}_{\mathcal{J}_{SC}}^{pre}$  and  $\mathcal{D}_{\mathcal{J}_{SC}}^{next}$  represent the 3D joints of the preceding and succeeding signs of the co-articulation, respectively. In our implementation,  $\mathcal{J}_{SC}$  includes the joints of hands, wrists, and elbows. Our sign connector is a 4-layer MLP taking the concatenation of  $\mathcal{D}_{\mathcal{J}_{SC}}^{pre}$ ,  $\mathcal{D}_{\mathcal{J}_{SC}}^{next}$ , and their Euclidean coordinate distance  $\mathcal{D}_{\mathcal{J}_{SC}}^{pre} - \mathcal{D}_{\mathcal{J}_{SC}}^{next}$  as inputs. We minimize the loss between the prediction yielded by the MLP and the target  $L$  using the L1 loss function.

Once the sign connector is optimized, given a 3D sign sequence  $\{S_1, S_2, \dots, S_K\}$ , it can predict the duration  $\hat{L}_k$  of the co-articulation between two adjacent 3D signs  $S_{k-1}$  and  $S_k$  ( $2 \leq k \leq K$ ). Subsequently, we interpolate  $\hat{L}_k$  frames between the last frame of  $S_{k-1}$  and the first frame of  $S_k$  in the 3D joint space to simulate the co-articulation  $C_k^{k-1}$ . Finally, a 3D sign sequence  $\{S_1, C_2^1, S_2, \dots, S_{K-1}, C_K^{K-1}, S_K\}$  is yielded for rendering.

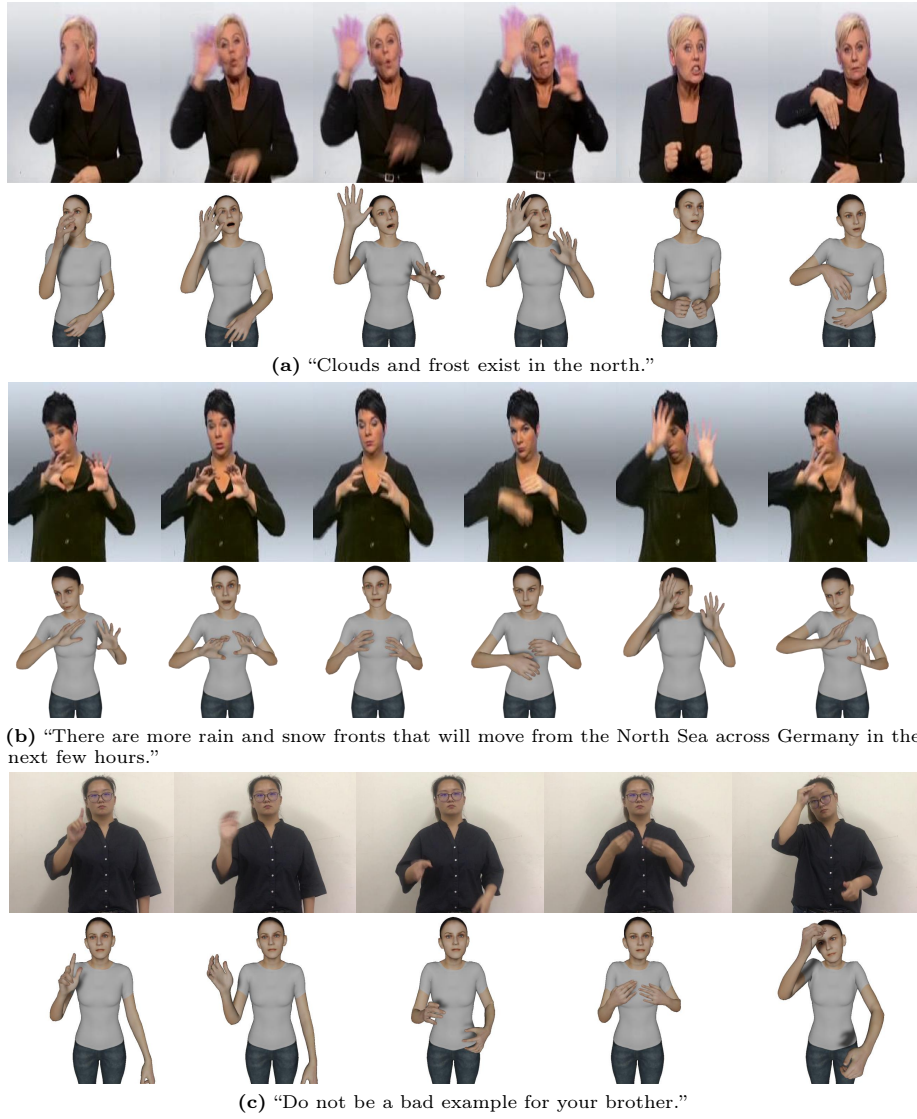
**Rendering Module.** We render the 3D sign sequence frame by frame using the Blender toolkit [12]. The pose and facial expression parameters of the SMPL-X model are utilized to drive the avatar.

### 3.4 By-Products of 3D Signs

The generated 3D signs implicitly integrate the human pose prior within the SMPLSign-X. Below, we discuss two by-products of these 3D signs, which significantly enhance keypoint-based models for sign language understanding.

**3D Keypoint Augmentation.** In contrast to existing Spoken2Sign translation works [53, 54, 56–58], which only generate frontal-view keypoints/videos, our approach outputs a sequence of 3D signs. The 3D nature of these signs allows the development of 3D keypoint augmentation to enhance keypoint-based models for sign language understanding. Specifically, during each training iteration, we sample an angle,  $\delta$ , from a pre-defined range of  $[-\Delta, \Delta]$ , where  $\Delta = 20^\circ$  by default. We then rotate the original 3D keypoints by  $\delta$  through modifying the global orientation parameters  $\zeta$ . Finally, the projected 2D keypoints are fed into the sign language understanding models as inputs.

**Understanding from Multiple Views.** Current datasets for sign language understanding commonly include 2D videos of signs captured from a frontal perspective. This limitation leads to the development of models that are only capable of interpreting sign languages from this single viewpoint. The introduction of 3D signs naturally enables sign language understanding from multiple views. Drawing inspiration from the TwoStream Network [9], which simultaneously processes 2D sign videos and their corresponding keypoint sequences, we modify the network to process two keypoint sequences: one sequence represents the original frontal view (a projection of the 3D keypoints without rotation), while the other sequence represents the side view (a projection of the 3D keypoints resulted from a  $60^\circ$  rotation).



**Fig. 4:** Qualitative results on P-2014T [4] (a and b) and CSL [80] (c). In each sub-figure, we display the text in the caption, and show the ground truth sign video and our translation result in the first row and second row, respectively.

## 4 Experiments

**Datasets and Evaluation Metrics.** Due to the absence of SMPL-X annotations in existing sign language datasets, back-translation is adopted as the primary metric [56–58] to evaluate the Spoken2Sign system on Phoenix-2014T

(P-2014T) [4] and CSL-Daily (CSL) [80]. We utilize SingleStream-Keypoint [9] as the back translator by default due to its superior keypoint-based Sign2Spoken performance and the availability of its code. P-2014T is a German sign language dataset with a vocabulary size of 1,066 for glosses and 2,887 for German text and there are 7,096/519/642 samples in its training/dev/test set. CSL is a large-scale Chinese sign language dataset, consisting of 18,401/1,077/1,176 samples in the training/dev/test set. Its vocabulary includes 2,000 glosses and 2,343 Chinese words. Following [56–58], we report BLEU-4 and ROUGE-L scores. In the P-2014T and CSL datasets, signers are recorded while standing stationary in front of the camera, maintaining an almost fixed pose. To verify the effectiveness of 3D keypoint augmentation, we conduct classification experiments on two extra datasets, WLASL [38] and MSASL [32], which feature more pronounced variations in signer poses. We report per-instance/class top-1/5 accuracy [84].

**Implementation Details.** Following SMPLify-X [49], we optimize Eq. 3 in multiple stages with the Limited-memory BFGS optimizer (L-BFGS) [48]. By default, we set  $\lambda_1 = 3e5$ ,  $\lambda_2 = 7e5$ , and  $\lambda_3 = 1e3$ , in Eq. 3. Each frame takes 300 epochs to fit the SMPL-X model [49]. To keep better temporal consistency, we pre-estimate the unified shape parameters  $\beta$  and global orientation  $\zeta$  of SMPL-X and extrinsic camera parameters for the input sign video. When training the sign connector, we set the learning rate as  $1e-5$  and use the Adam optimizer [34]. We also filter out outlying samples with too long co-articulations for stable training. For a fair comparison with existing works [53, 54, 56–58], the inputs of all models are the re-projected 2D keypoints.

#### 4.1 Qualitative Evaluation

The objective of this work is to translate spoken languages into sign languages. The translation results are displayed through a 3D avatar. To verify the effectiveness of our system, we show several qualitative results on P-2014T and CSL in Figure 4. It can be seen that the translation results are aesthetically comparable to the ground truth sign videos. Refer to the supplementary materials for more qualitative results.

#### 4.2 Quantitative Evaluation

We employ the widely adopted back-translation metric [58] to assess the effectiveness of our 3D sign estimator. The estimated 3D signs consist of meshes and 3D keypoints. For a fair comparison with previous works, we project these 3D keypoints back into 2D space, adopting a frontal view. Then, the back translator [9] is trained using these re-projected 2D keypoint sequences to translate sign languages into spoken languages.

**Comparison with State-of-the-Art Works.** Table 1 presents a comparative analysis of our approach against other state-of-the-art works, focusing on back-translation. The prevalent approach in existing works [16, 53, 54, 56–58] uses generated keypoint sequences to represent sign languages. Additionally, some studies [16, 58] incorporate a GAN [55] or a Diffusion model [76] to further generate

**Table 1:** Comparison of back-translation performance with existing Spoken2Sign translation works on the P-2014T benchmark. \*: using SignDiff’s [16] back translator.

Method	Dev		Test	
	BLEU-4↑	ROUGE↑	BLEU-4↑	ROUGE↑
Progressive Transformer [54]	11.82	33.18	10.51	32.46
Adversarial Training [53]	12.65	33.68	10.81	32.74
Mixture Density Networks [56]	11.54	33.40	11.68	33.19
Mixture of Motion Primitives [57]	14.03	37.76	13.30	36.77
FS-Net [58]	16.92	35.74	21.10	42.57
SignDiff [16]	18.26	39.62	22.15	46.82
Ours*	22.28	47.48	22.57	48.29
Ours	<b>24.16</b>	<b>49.12</b>	<b>25.46</b>	<b>49.68</b>

synthetic 2D videos from these keypoints. In contrast, our novel Spoken2Sign system produces a 3D sign sequences, visualized through an avatar. In line with the common practice, where back-translation is utilized for evaluation, our approach attains a BLEU-4 score of 24.16 on the P-2014T dev set. It is worth noting that we also re-implement the back translator [63] used in the previous best work, SignDiff [16], for a fairer comparison. Although this translator is weaker than our default SingleStream-Keypoint, our approach can still outperform SignDiff on the P-2014T dev set by more than 4 points in BLEU-4 score.

**3D Sign Estimator.** In Section 3.2, we present SMPLSign-X, an enhancement of the vanilla SMPLify-X that effectively estimates 3D signs from monocular sign videos. To demonstrate the superior performance of our 3D sign estimator, we benchmark it against three state-of-the-art whole-body estimation methods: SMPLify-X [49], SMPLer-X [3], and OSX [41], with the back-translation results detailed in Table 2. Our approach outperforms all three methods, regardless they are optimization-based (SMPLify-X) or regression-based (SMPLer-X and OSX). Furthermore, our method achieves the lowest 2D keypoint loss and the highest temporal consistency, indicating the precision of the estimated signs and the smoothness of the Spoken2Sign translation results. We attribute these improvements to the incorporation of prior knowledge about sign languages. Additionally, under the same experimental setting, we introduce a strong baseline that generates 2D keypoints for each frame using a pre-trained HRNet [65], contrasting with our method and other baselines that re-project the estimated 3D keypoints back into 2D space. We observe that our approach outperforms this baseline on P-2014T. This superiority can be attributed to the consideration of temporal consistency.

**User Study with Deaf Participants.** The primary goal of our Spoken2Sign system is to narrow the communication gap between the deaf and the hearing. Consequently, conducting a user study with deaf participants is essential. We invite four Chinese signers to evaluate our Spoken2Sign results on three aspects: naturalness (checking for awkward poses), smoothness (observing for noticeable

**Table 2:** Comparison with other state-of-the-art 3D sign estimation methods using back-translation. We re-implement SMPLify-X [49], SMPLer-X [3] and OSX [41]. ‘‘HRNet’’ denotes a strong baseline that directly estimates 2D keypoints (pseudo ground truth) from raw videos. \*: regression-based methods trained on large-scale 3D datasets with SMPL-X annotations. 2D KL: Euclidean distance between re-projected 2D keypoints and pseudo ground truth. TC: temporal consistency measured by calculating the average cosine similarity across all consecutive rendered frames in pixel space.

Dataset	Method	Dev				Test			
		BLEU-4↑	ROUGE↑	2D KL↓	TC↑	BLEU-4↑	ROUGE↑	2D KL↓	TC↑
P-2014T	HRNet [65]	22.94	48.81	0.00	0.961	24.95	49.13	0.00	0.962
	SMPLify-X [49]	19.21	44.28	31.56	0.945	19.69	43.65	31.06	0.943
	SMPLer-X* [3]	19.49	44.95	29.50	0.966	20.01	44.93	29.72	0.965
	OSX* [41]	22.31	47.71	26.87	0.969	23.00	47.23	27.91	0.966
	Ours	<b>24.16</b>	<b>49.12</b>	<b>22.09</b>	<b>0.982</b>	<b>25.46</b>	<b>49.68</b>	<b>23.48</b>	<b>0.981</b>
CSL	HRNet [65]	22.14	51.02	0.00	0.969	21.29	50.97	0.00	0.970
	SMPLify-X [49]	19.31	46.46	30.69	0.947	18.96	46.71	29.66	0.949
	SMPLer-X* [3]	20.76	48.93	29.59	0.967	20.86	49.29	28.05	0.968
	OSX* [41]	20.44	49.00	29.75	0.965	20.29	49.60	27.40	0.965
	Ours	<b>21.66</b>	<b>49.69</b>	<b>26.51</b>	<b>0.980</b>	<b>21.44</b>	<b>49.80</b>	<b>24.40</b>	<b>0.981</b>

**Table 3:** User study with deaf participants. Nat.: naturalness; Smo.: smoothness; Sim.: similarity to the raw video. Score range: 1-5.

Dataset	Method	Nat.↑	Smo.↑	Sim.↑
P-2014T	SMPLify-X [49]	1.52	1.98	2.41
	Ours	<b>3.58</b>	<b>4.04</b>	<b>3.94</b>
CSL	SMPLify-X [49]	1.27	1.75	1.69
	Ours	<b>3.78</b>	<b>4.14</b>	<b>3.78</b>

**Table 4:** Ablation study for the proposed loss functions,  $\mathcal{L}_{unseen}$ ,  $\mathcal{L}_{upright}$ , and  $\mathcal{L}_{smooth}$ , on P-2014T using back-translation.

Method	Dev		Test	
	BLEU-4↑	ROUGE↑	BLEU-4↑	ROUGE↑
w/o $\mathcal{L}_{unseen}$	22.57	47.25	23.70	47.62
w/o $\mathcal{L}_{upright}$	23.10	47.93	23.99	48.00
w/o $\mathcal{L}_{smooth}$	23.64	48.50	24.36	48.35
w/ all (default)	<b>24.16</b>	<b>49.12</b>	<b>25.46</b>	<b>49.68</b>

shaking between frames), and similarity to raw videos<sup>1</sup>. We provide 100 randomly selected videos from each dataset to the participants, who then give a rating ranging from 1 to 5 for each aspect. The average ratings across all videos and participants are reported in Table 3. It is evident that our approach significantly outperforms the baseline, SMPLify-X [49], in all aspects.

**Loss Functions.** In Section 3.2, we introduce three loss functions designed to facilitate 3D sign estimation:  $\mathcal{L}_{unseen}$ , which draws the unseen keypoints closer to those of the rest pose;  $\mathcal{L}_{upright}$ , aimed at encouraging an upright posture in the upper body; and  $\mathcal{L}_{smooth}$ , which preserves temporal consistency to generate visually appealing results. As shown in Table 4, omitting any of these loss functions leads to degraded performance.

<sup>1</sup> Though the participants are not native in German sign language, their familiarity with general sign language rules enables them to evaluate the similarity to sign videos for P-2014T.

**Table 5:** Ablation study on 3D keypoint (Kp.) augmentation (Aug.).

Dataset	3D Kp.	Per-instance		Per-class	
	Aug.	Top-1↑	Top-5↑	Top-1↑	Top-5↑
WLASL		46.20	78.81	43.72	77.55
	✓	<b>47.66</b>	<b>79.71</b>	<b>45.10</b>	<b>78.16</b>
MSASL		51.81	73.78	48.52	71.76
	✓	<b>53.10</b>	<b>75.06</b>	<b>49.71</b>	<b>72.71</b>

**Table 6:** Ablation study on multi-view Sign2Spoken translation.

Dataset	View		Dev		Test	
	Frontal	Side	BLEU-4↑	ROUGE↑	BLEU-4↑	ROUGE↑
P-2014T		✓	24.16	49.12	25.46	49.68
	✓	✓	22.60	47.31	23.49	47.31
CSL		✓	21.66	49.69	21.44	49.80
	✓	✓	<b>23.14</b>	<b>52.13</b>	<b>22.22</b>	<b>51.47</b>

### 4.3 Effectiveness of the By-Products

Our approach is able to estimate 3D signs from monocular sign videos. In Section 3.4, we introduce two by-products, 3D keypoint augmentation and multi-view understanding, which have the potential to improve keypoint-based models for sign language understanding.

**3D Keypoint Augmentation.** 3D signs can be rotated at any angle, motivating us to develop 3D keypoint augmentation. This is achieved by randomly rotating the input 3D sign by a small angle before projecting it into 2D space. To validate its effectiveness, we turn to two challenging ISLR datasets, WLASL and MSASL, where the variation of signer poses is more dramatic than that in P-2014T and CSL. We use NLA-SLR-Keypoint-64 [84] as the base model. As shown in Table 5, 3D keypoint augmentation consistently improves model performance across all metrics on the dev sets of both datasets, with almost no extra training cost.

**Multi-View Spoken2Sign Translation.** Another by-product is the use of side-view keypoints to enhance the performance of the model trained on frontal-view keypoints. Table 6 shows the effectiveness of multi-view Spoken2Sign translation, with the SingleStream-Keypoint [9] as the base model.

## 5 Conclusion

This paper focuses on Spoken2Sign translation, a reverse process to traditional Sign2Spoken translation, aimed at narrowing the communication gap between deaf and hearing individuals. In contrast to prior works that produce translation results in 2D space, our innovative method generates 3D signs using the proposed techniques such as SMPLSign-X and sign connector. The translation results are displayed through an avatar. Our method involves three main steps: 1) building a sign language dictionary; 2) estimating the 3D representation for each sign in this dictionary; and 3) executing Spoken2Sign translation and rendering an avatar. Additionally, we introduce two by-products, 3D keypoint augmentation and multi-view understanding, which significantly promote the keypoint-based models. Extensive experiments demonstrate the effectiveness of our approach.

## References

1. Albanie, S., Varol, G., Momeni, L., Afouras, T., Chung, J.S., Fox, N., Zisserman, A.: BSL-1K: Scaling up co-articulated sign language recognition using mouthing cues. In: ECCV. pp. 35–53 (2020)
2. Bao, J., Chen, D., Wen, F., Li, H., Hua, G.: Cvae-gan: fine-grained image generation through asymmetric training. In: ICCV. pp. 2745–2754 (2017)
3. Cai, Z., Yin, W., Zeng, A., Wei, C., Sun, Q., Wang, Y., Pang, H.E., Mei, H., Zhang, M., Zhang, L., et al.: Smpler-x: Scaling up expressive human pose and shape estimation. In: NeurIPS (2023)
4. Camgoz, N.C., Hadfield, S., Koller, O., Ney, H., Bowden, R.: Neural sign language translation. In: CVPR. pp. 7784–7793 (2018)
5. Camgoz, N.C., Koller, O., Hadfield, S., Bowden, R.: Multi-channel transformers for multi-articulatory sign language translation. In: ECCV. pp. 301–319. Springer (2020)
6. Camgöz, N.C., Koller, O., Hadfield, S., Bowden, R.: Sign language transformers: Joint end-to-end sign language recognition and translation. In: CVPR. pp. 10020–10030 (2020)
7. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML. pp. 1597–1607. PMLR (2020)
8. Chen, Y., Wei, F., Sun, X., Wu, Z., Lin, S.: A simple multi-modality transfer learning baseline for sign language translation. In: CVPR. pp. 5120–5130 (2022)
9. Chen, Y., Zuo, R., Wei, F., Wu, Y., Liu, S., Mak, B.: Two-stream network for sign language recognition and translation. In: NeurIPS (2022)
10. Cheng, K.L., Yang, Z., Chen, Q., Tai, Y.: Fully convolutional networks for continuous sign language recognition. In: ECCV. vol. 12369, pp. 697–714 (2020)
11. Cheng, Y., Wei, F., Bao, J., Chen, D., Zhang, W.: Cico: Domain-aware sign language retrieval via cross-lingual contrastive learning. In: CVPR (2023)
12. Community, B.O.: Blender - a 3D modelling and rendering package. Blender Foundation, Stichting Blender Foundation, Amsterdam (2018), <http://www.blender.org>
13. Cui, R., Liu, H., Zhang, C.: A deep neural framework for continuous sign language recognition by iterative training. IEEE TMM **PP**, 1–1 (07 2019)
14. Duarte, A., Albanie, S., Giró-i Nieto, X., Varol, G.: Sign language video retrieval with free-form textual queries. In: CVPR. pp. 14094–14104 (2022)
15. Duarte, A., Palaskar, S., Ventura, L., Ghadiyaram, D., DeHaan, K., Metze, F., Torres, J., Giro-i Nieto, X.: How2sign: a large-scale multimodal dataset for continuous american sign language. In: CVPR. pp. 2735–2744 (2021)
16. Fang, S., Sui, C., Zhang, X., Tian, Y.: Signdiff: Learning diffusion models for american sign language production (2023)
17. Forte, M.P., Kulits, P., Huang, C.H.P., Choutas, V., Tzionas, D., Kuchenbecker, K.J., Black, M.J.: Reconstructing signing avatars from video using linguistic priors. In: CVPR. pp. 12791–12801 (2023)
18. Gan, S., Yin, Y., Jiang, Z., Xia, K., Xie, L., Lu, S.: Contrastive learning for sign language recognition and translation. In: IJCAI. pp. 763–772 (2023)
19. Geman, S.: Statistical methods for tomographic image reconstruction. Bulletin of International Statistical Institute **4**, 5–21 (1987)
20. Graves, A., Fernández, S., Gomez, F., Schmidhuber, J.: Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In: ICML. pp. 369–376 (2006)

21. Hao, A., Min, Y., Chen, X.: Self-mutual distillation learning for continuous sign language recognition. In: ICCV. pp. 11303–11312 (October 2021)
22. Hu, H., Zhao, W., Zhou, W., Li, H.: Signbert+: Hand-model-aware self-supervised pre-training for sign language understanding. IEEE TPAMI (2023)
23. Hu, H., Zhou, W., Li, H.: Hand-model-aware sign language recognition. In: AAAI. vol. 35, pp. 1558–1566 (2021)
24. Hu, L., Gao, L., Feng, W., et al.: Self-emphasizing network for continuous sign language recognition. In: AAAI (2023)
25. Hu, L., Gao, L., Liu, Z., Feng, W.: Continuous sign language recognition with correlation network. In: CVPR (2023)
26. Huang, W., Pan, W., Zhao, Z., Tian, Q.: Towards fast and high-quality sign language production. In: MM. pp. 3172–3181 (2021)
27. Huang, W., Zhao, Z., He, J., Zhang, M.: Dualsign: Semi-supervised sign language production with balanced multi-modal multi-task dual transformation. In: MM. pp. 5486–5495 (2022)
28. Hwang, E., Kim, J.H., Park, J.C.: Non-autoregressive sign language production with gaussian space. In: BMVC (2021)
29. Jiang, S., Sun, B., Wang, L., Bai, Y., Li, K., Fu, Y.: Skeleton aware multi-modal sign language recognition. In: CVPRW. pp. 3413–3423 (2021)
30. Jin, S., Xu, L., Xu, J., Wang, C., Liu, W., Qian, C., Ouyang, W., Luo, P.: Whole-body human pose estimation in the wild. In: ECCV. pp. 196–214 (2020)
31. Joo, H., Simon, T., Sheikh, Y.: Total capture: A 3d deformation model for tracking faces, hands, and bodies. In: CVPR. pp. 8320–8329 (2018)
32. Joze, H.R.V., Koller, O.: MS-ASL: A large-scale data set and benchmark for understanding American sign language. In: BMVC (2019)
33. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: CVPR. pp. 7122–7131 (2018)
34. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
35. Kipp, M., Heloir, A., Nguyen, Q.: Sign language avatars: Animation and comprehensibility. In: Intelligent Virtual Agents: 10th International Conference, IVA 2011, Reykjavik, Iceland, September 15–17, 2011. Proceedings 11. pp. 113–126. Springer (2011)
36. Lassner, C., Romero, J., Kiefel, M., Bogo, F., Black, M.J., Gehler, P.V.: Unite the people: Closing the loop between 3d and 2d human representations. In: CVPR. pp. 6050–6059 (2017)
37. Lee, T., Oh, Y., Lee, K.M.: Human part-wise 3d motion context learning for sign language recognition. In: ICCV. pp. 20740–20750 (2023)
38. Li, D., Rodriguez, C., Yu, X., Li, H.: Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In: WACV. pp. 1459–1469 (2020)
39. Li, D., Xu, C., Yu, X., Zhang, K., Swift, B., Suominen, H., Li, H.: Tspnet: Hierarchical feature learning via temporal semantic pyramid for sign language translation. In: NeurIPS. vol. 33, pp. 12034–12045 (2020)
40. Li, D., Yu, X., Xu, C., Petersson, L., Li, H.: Transferring cross-domain knowledge for video sign language recognition. In: CVPR. pp. 6205–6214 (2020)
41. Lin, J., Zeng, A., Wang, H., Zhang, L., Li, Y.: One-stage 3d whole-body mesh recovery with component aware transformer. In: CVPR. pp. 21159–21168 (2023)
42. Liu, Y., Gu, J., Goyal, N., Li, X., Edunov, S., Ghazvininejad, M., Lewis, M., Zettlemoyer, L.: Multilingual denoising pre-training for neural machine translation. *TACL* **8**, 726–742 (2020)



43. McDonald, J., Wolfe, R., Schnepf, J., Hochgesang, J., Jamrozik, D.G., Stumbo, M., Berke, L., Bialek, M., Thomas, F.: An automated technique for real-time production of lifelike animations of american sign language. *Universal Access in the Information Society* **15**, 551–566 (2016)
44. Min, Y., Hao, A., Chai, X., Chen, X.: Visual alignment constraint for continuous sign language recognition. In: *ICCV*. pp. 11542–11551 (October 2021)
45. Momeni, L., Bull, H., Prajwal, K., Albanie, S., Varol, G., Zisserman, A.: Automatic dense annotation of large-vocabulary sign language videos. In: *ECCV*. pp. 671–690 (2022)
46. Momeni, L., Varol, G., Albanie, S., Afouras, T., Zisserman, A.: Watch, read and lookup: learning to spot signs from multiple supervisors. In: *ACCV* (2020)
47. Niu, Z., Zuo, R., Mak, B., Wei, F.: A hong kong sign language corpus collected from sign-interpreted tv news. In: *LREC-COLING*. pp. 636–646 (2024)
48. Nocedal, J., Wright, S.J.: Nonlinear equations. *Numerical Optimization* pp. 270–302 (2006)
49. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: *CVPR*. pp. 10975–10985 (2019)
50. Rastgoo, R., Kiani, K., Escalera, S., Sabokrou, M.: Sign language production: A review. In: *CVPRW*. pp. 3451–3461 (2021)
51. Razavi, A., Van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with vq-vae-2. *NeurIPS* **32** (2019)
52. Sak, H., Senior, A., Rao, K., Irsoy, O., Graves, A., Beaufays, F., Schalkwyk, J.: Learning acoustic frame labeling for speech recognition with recurrent neural networks. In: *ICASSP*. pp. 4280–4284 (2015)
53. Saunders, B., Camgöz, N.C., Bowden, R.: Adversarial training for multi-channel sign language production. In: *BMVC* (2020)
54. Saunders, B., Camgoz, N.C., Bowden, R.: Progressive transformers for end-to-end sign language production. In: *ECCV*. pp. 687–705 (2020)
55. Saunders, B., Camgoz, N.C., Bowden, R.: Anonymsign: Novel human appearance synthesis for sign language video anonymisation. In: *FG 2021*. pp. 1–8 (2021)
56. Saunders, B., Camgoz, N.C., Bowden, R.: Continuous 3d multi-channel sign language production via progressive transformers and mixture density networks. *IJCV* **129**(7), 2113–2135 (2021)
57. Saunders, B., Camgoz, N.C., Bowden, R.: Mixed signals: Sign language production via a mixture of motion primitives. In: *ICCV*. pp. 1919–1929 (2021)
58. Saunders, B., Camgoz, N.C., Bowden, R.: Signing at scale: Learning to co-articulate signs for large-scale photo-realistic sign language production. In: *CVPR*. pp. 5141–5151 (2022)
59. Shmelkov, K., Schmid, C., Alahari, K.: How good is my gan? In: *ECCV*. pp. 213–229 (2018)
60. Stoll, S., Camgöz, N.C., Hadfield, S., Bowden, R.: Sign language production using neural machine translation and generative adversarial networks. In: *BMVC* (2018)
61. Stoll, S., Camgoz, N.C., Hadfield, S., Bowden, R.: Text2sign: towards sign language production using neural machine translation and generative adversarial networks. *IJCV* **128**(4), 891–908 (2020)
62. Tang, S., Hong, R., Guo, D., Wang, M.: Gloss semantic-enhanced network with online back-translation for sign language production. In: *MM*. pp. 5630–5638 (2022)
63. Tarrés, L., Gállego, G.I., Duarte, A., Torres, J., Giró-i Nieto, X.: Sign language translation from instructional videos. In: *CVPRW*. pp. 5624–5634 (2023)

64. Varol, G., Momeni, L., Albanie, S., Afouras, T., Zisserman, A.: Read and attend: Temporal localisation in sign language videos. In: CVPR. pp. 16857–16866 (2021)
65. Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., et al.: Deep high-resolution representation learning for visual recognition. *IEEE TPAMI* **43**(10), 3349–3364 (2020)
66. Wei, F., Chen, Y.: Improving continuous sign language recognition with cross-lingual signs. In: ICCV. pp. 23612–23621 (October 2023)
67. Xiang, D., Joo, H., Sheikh, Y.: Monocular total capture: Posing face, body, and hands in the wild. In: CVPR. pp. 10965–10974 (2019)
68. Xie, P., Zhang, Q., Li, Z., Tang, H., Du, Y., Hu, X.: Vector quantized diffusion model with codeunet for text-to-sign pose sequences generation (2023)
69. Xu, H., Bazavan, E.G., Zanfir, A., Freeman, W.T., Sukthankar, R., Sminchisescu, C.: Ghum & ghuml: Generative 3d human shape and articulated pose models. In: CVPR. pp. 6184–6193 (2020)
70. Xu, Y., Zhu, S.C., Tung, T.: Denserac: Joint 3d pose and shape estimation by dense render-and-compare. In: ICCV. pp. 7760–7770 (2019)
71. Yao, H., Zhou, W., Feng, H., Hu, H., Zhou, H., Li, H.: Sign language translation with iterative prototype. In: ICCV. pp. 15592–15601 (October 2023)
72. Yin, A., Zhao, Z., Liu, J., Jin, W., Zhang, M., Zeng, X., He, X.: Simulslt: End-to-end simultaneous sign language translation. In: MM. pp. 4118–4127 (2021)
73. Yin, K., Read, J.: Better sign language translation with stmc-transformer. *COLING* (2020)
74. Yu, P., Zhang, L., Fu, B., Chen, Y.: Efficient sign language translation with a curriculum-based non-autoregressive decoder. In: IJCAI. pp. 5260–5268 (2023)
75. Zelinka, J., Kanis, J.: Neural sign language synthesis: Words are our glosses. In: WACV. pp. 3395–3403 (2020)
76. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: ICCV. pp. 3836–3847 (2023)
77. Zhao, W., Hu, H., Zhou, W., Shi, J., Li, H.: BEST: BERT pre-training for sign language recognition with coupling tokenization. In: AAAI (2023)
78. Zheng, J., Wang, Y., Tan, C., Li, S., Wang, G., Xia, J., Chen, Y., Li, S.Z.: CVT-SLR: Contrastive visual-textual transformation for sign language recognition with variational alignment. In: CVPR (2023)
79. Zhou, B., Chen, Z., Clapés, A., Wan, J., Liang, Y., Escalera, S., Lei, Z., Zhang, D.: Gloss-free sign language translation: Improving from visual-language pretraining. In: ICCV. pp. 20871–20881 (October 2023)
80. Zhou, H., Zhou, W., Qi, W., Pu, J., Li, H.: Improving sign language translation with monolingual data by sign back-translation. In: CVPR. pp. 1316–1325 (2021)
81. Zhou, H., Zhou, W., Zhou, Y., Li, H.: Spatial-temporal multi-cue network for continuous sign language recognition. In: AAAI. pp. 13009–13016 (2020)
82. Zuo, R., Mak, B.: C2SLR: Consistency-enhanced continuous sign language recognition. In: CVPR. pp. 5131–5140 (2022)
83. Zuo, R., Mak, B.: Improving continuous sign language recognition with consistency constraints and signer removal. *ACM TOMM* **20**(6), 1–25 (2024)
84. Zuo, R., Wei, F., Mak, B.: Natural language-assisted sign language recognition. In: CVPR (2023)
85. Zuo, R., Wei, F., Mak, B.: Towards online sign language recognition and translation. arXiv preprint arXiv:2401.05336 (2024)