

Signing at Scale: Learning to Co-Articulate Signs for Large-Scale Photo-Realistic Sign Language Production

Ben Saunders, Necati Cihan Camgoz, Richard Bowden
University of Surrey

{b.saunders, n.camgoz, r.bowden}@surrey.ac.uk




Figure 1. **Photo-Realistic Sign Language Production:** Given a spoken language sentence from an unconstrained domain of discourse (a), an initial translation is conducted to a gloss sequence (b). FS-NET next produces a co-articulated continuous skeleton pose sequence from dictionary signs (c), which SIGNGAN generates into a photo-realistic sign language video in a given style (d).

Abstract

Sign languages are visual languages, with vocabularies as rich as their spoken language counterparts. However, current deep-learning based Sign Language Production (SLP) models produce under-articulated skeleton pose sequences from constrained vocabularies and this limits applicability. To be understandable and accepted by the deaf, an automatic SLP system must be able to generate co-articulated photo-realistic signing sequences for large domains of discourse.

In this work, we tackle large-scale SLP by learning to co-articulate between dictionary signs, a method capable of producing smooth signing while scaling to unconstrained domains of discourse. To learn sign co-articulation, we propose a novel Frame Selection Network (FS-NET) that improves the temporal alignment of interpolated dictionary signs to continuous signing sequences. Additionally, we propose SIGNGAN, a pose-conditioned human synthesis model that produces photo-realistic sign language videos direct from skeleton pose. We propose a novel keypoint-

based loss function which improves the quality of synthesized hand images.

We evaluate our SLP model on the large-scale *meineDGS* (mDGS) corpus, conducting extensive user evaluation showing our FS-NET approach improves co-articulation of interpolated dictionary signs. Additionally, we show that SIGNGAN significantly outperforms all baseline methods for quantitative metrics, human perceptual studies and native deaf signer comprehension.

1. Introduction

Sign languages are rich visual languages with large lexical vocabularies [52] and intricate co-articulated movements of both manual (hands and body) and non-manual (facial) features. Sign Language Production (SLP), the automatic translation from spoken language sentences to sign language sequences, must be able to produce photo-realistic continuous signing for large domains of discourse to be useful to the deaf communities.

Prior deep-learning approaches to SLP have either pro-

duced concatenated isolated sequences that disregard the natural co-articulation between signs [53, 68] or continuous sequences end-to-end [22, 44, 46, 68] which suffer from under-articulation [43]. Furthermore, these methods have struggled to generalise beyond the limited domain of weather [14].

In this paper, we propose an SLP method to produce photo-realistic continuous sign language videos direct from unconstrained spoken language sequences. Firstly, we translate from spoken language to gloss¹ sequences. We next learn the temporal co-articulation between gloss-based dictionary signs, modelling the temporal prosody of sign language [3].

To model sign co-articulation, we propose a novel Frame Selection Network (FS-NET) that learns the optimal subset of frames that best represents a continuous signing sequence (Fig. 2 middle). We build a transformer encoder with cross-attention [58] to predict a temporal alignment path supervised by Dynamic Time Warping (DTW).

The resulting skeleton pose sequences are subsequently used to condition a video-to-video synthesis model capable of generating photo-realistic sign language videos, named SIGNGAN (Fig. 2 right). Due to the natural presence of motion blur in sign language datasets from fast moving hands [15], a classical application of a hand discriminator leads to an increase in blurred hand generation. To avoid this, we propose a novel keypoint-based loss that significantly improves the quality of hand image synthesis in our photo-realistic signer generation module. To enable training on diverse sign language datasets, we propose a method for controllable video generation that models a multi-modal distribution of sign language videos in different styles.

Our deep-learning based SLP model is able to generalise to large domains of discourse, as it is trivial to increase vocabulary with a few examples of this new sign in a continuous signing context. We conduct extensive deaf user evaluation on a translation protocol of mDGS [20], showing that FS-NET improves the natural signing motion of interpolated dictionary sequences and is overwhelmingly preferred to baseline SLP methods [47]. Additionally, we achieve state-of-the-art back translation performance on RWTH-PHOENIX-Weather-2014T (PHOENIX14T) with a 43% improvement over baselines, highlighting the understandable nature of our approach.

Furthermore, we evaluate SIGNGAN using the high quality Content4All (C4A) dataset [5], outperforming state-of-the-art synthesis methods [6, 54, 62, 63] for quantitative evaluation and human perception studies. Finally, we conduct a further deaf user evaluation to show that SIGNGAN is more understandable than the skeletal sequences previously used to represent sign [44].

¹Glosses are a written representation of sign that follow sign language ordering and grammar, defined as minimal lexical items [52].

The contributions of this paper can be summarised as:

- The first SLP model to produce large-scale sign language sequences from an unconstrained domain of discourse to a level understandable by a native deaf signer
- A novel Frame Selection Network, FS-NET, that learns to co-articulate between dictionary signs via a monotonic alignment to continuous sequences
- A method to generate photo-realistic continuous sign language videos, SIGNGAN, with a novel hand keypoint loss that improves the hand synthesis quality
- Extensive user evaluation of our proposed approach, showing preference of our proposed method, alongside state-of-the-art back translation results

2. Related Work

Sign Language Production The initial focus of computational sign language technology was Sign Language Recognition (SLR) [11, 18, 28] with few works tackling unconstrained SLR [8, 25, 29]. More recently, focus has shifted to Sign Language Translation (SLT) [4, 7, 27].

Sign Language Production (SLP), the translation from spoken to sign language, has been historically tackled using animated avatars [9, 26, 36] with rules-based co-articulation that does not generalise to unseen sequences [49].

Initial deep learning-based SLP methods concatenated isolated signs with no regards for natural co-articulation [53, 68]. Recently, continuous SLP methods have directly regressed sequences of multiple signs [22, 44, 46–48], but exhibit under-articulated signing motion due to regression to the mean. To overcome under-articulation, we avoid generating pose directly and learn the optimal temporal alignment between dictionary and continuous sign sequences.

In addition, prior work has represented sign language as skeleton pose sequences, which have been shown to reduce the deaf comprehension compared to a photo-realistic production [59]. Previous works have attempted photo-realistic signer generation [10, 45, 54], but of low-resolution isolated signs. In this work, we produce high-resolution photo-realistic continuous sign language videos directly from spoken language input, from unrestricted domains of discourse.

Pose-Conditioned Human Synthesis Generative Adversarial Networks (GANs) [17] have achieved impressive results in image [23, 41, 63, 70] and, more recently, video generation tasks [35, 57, 60–62]. Specific to pose-conditioned human synthesis, there has been concurrent research focusing on the generation of whole body [1, 34, 37, 50, 55, 71], face [12, 31, 67] and hand [33, 56, 66] images.

However, there has been no research into accurate hand generation in the context of full body synthesis, with current

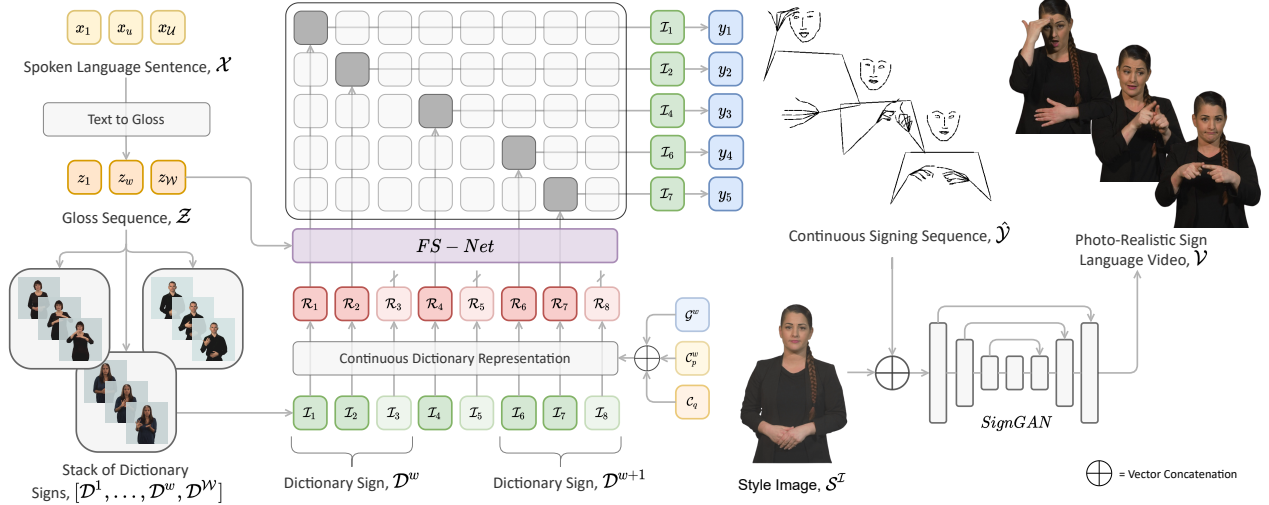


Figure 2. Overview of our proposed large-scale SLP method. An initial Text to Gloss (left) animates an interpolated dictionary sequence, \mathcal{I} , with a Frame Selection Network (FS-NET), learning the temporal alignment (middle) to a continuous signing sequence, \mathcal{Y} . Finally, SIGNGAN generates a photo-realistic sign language video, \mathcal{V} , (right) from the continuous skeleton pose and a given style image, $\mathcal{S}^{\mathcal{I}}$.

methods failing to generate high-quality hand images [59]. Due to the hands being high fidelity objects, they are often overlooked in model optimisation. Chan *et al.* introduced FaceGAN for high resolution face generation [6], but no similar work has been proposed for the more challenging task of hand synthesis in the context of sign language, where hand to hand interaction is ubiquitous. In this work, we propose a keypoint-based loss to enhance hand synthesis.

The task of human motion transfer, transferring motion from source to target videos via keypoint extraction, is relevant to our task [6, 65, 69]. However, there has been limited research into the generation of novel poses, which we produce from a given spoken language sentence. Additionally, works have attempted to produce unseen appearances in a few-shot manner [61, 67], but continue to produce only a single style at inference.

Sign Language Co-Articulation Sign language co-articulation can be defined as “the articulatory influence of one phonetic element on another across more than one intervening element” [19] and is an important distinction between isolated and natural continuous signing [39].

Co-articulation involves both the motion and duration of signs, with a particular focus on the transition between signs [39]. The boundaries of a sign are also modified depending on the context, with continuous signing typically produced faster than their isolated counterparts [49]. In this work, we model temporal co-articulation by learning the optimal alignment between isolated signs and continuous signing sequences, predicting the duration, boundary and transition of each sign in context.

3. Large-Scale Photo-Realistic SLP

The true aim of a large-scale SLP model is to translate a spoken language sequence from an unconstrained domain of discourse, $\mathcal{X} = (x_1, \dots, x_U)$ with U words, to a continuous photo-realistic sign language video, $\mathcal{V}^{\mathcal{S}} = (v_1, \dots, v_{\mathcal{T}})$ with \mathcal{T} frame. This is a challenging task due to the large vocabulary of unconstrained signing and the intricate spatial nature of sign, with a requirement for temporal co-articulation indicative of natural continuous signing.

We approach this problem as a multi-stage sequence-to-sequence task. Firstly, spoken language is translated to sign gloss, $\mathcal{Z} = (z_1, \dots, z_W)$, as an intermediate representation (Sec. 3.1). Next, our FS-NET model co-articulates between gloss-based dictionary signs to produce a full continuous signing sequence, $\mathcal{Y} = (y_1, \dots, y_{\mathcal{T}})$ (Sec. 3.2). Finally, given \mathcal{Y} and a style image, $\mathcal{S}^{\mathcal{I}}$, our video-to-video signer generation module generates a photo-realistic sign language video, $\mathcal{Z}^{\mathcal{S}}$ (Sec. 3.3). An overview of our approach can be seen in Fig. 2. In the remainder of this section, we shall describe each component of our approach in detail.

3.1. Text to Gloss

Given a spoken language sequence, \mathcal{X} , we first translate to a sign language grammar and order, represented by a gloss sequence, $\mathcal{Z} = (z_1, \dots, z_W)$ with W glosses (Fig. 2 left). We formulate this as a sequence-to-sequence problem, due to the non-monotonic relationship between the two sequences of different lengths. We use an encoder-decoder transformer [58] to perform this translation, formalised as:

$$f_t = E_{T2G}(x_t | x_{1:\mathcal{T}}) \quad (1)$$

$$g_{w+1} = D_{T2G}(g_w | g_{1:w-1}, f_{1:\mathcal{T}}) \quad (2)$$

where f_t and g_w are the encoded source and target tokens respectively and g_0 is the encoding of the special $\langle \text{bos} \rangle$ token. The output gloss tokens can be computed as $z_w = \text{argmax}_i(g_w)$ until the special $\langle \text{eos} \rangle$ token is predicted.

3.2. Gloss to Pose

Next, motivated by the monotonic relationship between glosses and signs, we produce a continuous signing pose sequence, $\hat{\mathcal{Y}} = (y_1, \dots, y_{\mathcal{T}})$ with \mathcal{T} frames, from the translated gloss sequence, \mathcal{Z} , using a learnt co-articulation of dictionary signs. We first encode the gloss sequence using a transformer encoder with self-attention, as:

$$h_w = E_{G2S}(z_w | z_{1:w}) \quad (3)$$

where h_w is the encoded gloss token for step w . We next collect a dictionary sample, \mathcal{D}^w , for every sign present in the gloss vocabulary. By definition, dictionary signs contain accurate and articulated sign content. Furthermore, it is trivial to expand to larger domains of discourse, simply collecting dictionary examples of the expanded vocabulary.

Interpolated Dictionary Representation Given the translated gloss sequence, \mathcal{Z} , we create a stack of ordered dictionary signs, $[\mathcal{D}^1, \dots, \mathcal{D}^w, \mathcal{D}^{\mathcal{W}}]$ (Bottom left of Fig. 2). As in previous works [46], we represent each dictionary sign as a sequence of skeleton pose, $\mathcal{D}^w = (s_1^w, \dots, s_{\mathcal{P}^w}^w)$ with \mathcal{P}^w frames. We first convert the stack of dictionary signs into a continuous sequence by linearly interpolating between neighbouring signs for a predefined fixed \mathcal{N}_{LI} frames. The final interpolated dictionary sequence, $\mathcal{I} = (\mathcal{I}_1, \dots, \mathcal{I}_{\mathcal{Q}})$ with \mathcal{Q} frames, is the combination of skeleton pose and the respective linear interpolation.

We next build a continuous dictionary sequence representation to be used as input to FS-NET. Alongside the skeleton pose of \mathcal{I} , we learn a gloss embedding, \mathcal{G}^w , unique to each gloss in the vocabulary, with a separate shared embedding for all interpolation frames, \mathcal{G}^{LI} . Additionally, we use a counter embedding proposed by Saunders *et al.* [44], expanded to both a specific counter, \mathcal{C}_p^w , relating to the progression of each dictionary sign and a global counter, \mathcal{C}_q , relating to the progress of the full sequence, \mathcal{I} . The final continuous dictionary representation, $\mathcal{R} = (\mathcal{R}_1, \dots, \mathcal{R}_{\mathcal{Q}})$ with \mathcal{Q} frames, is constructed by concatenating the corresponding skeleton, gloss and counter embeddings per frame, as:

$$\mathcal{R}_q = [s_p^w, \mathcal{G}^w, \mathcal{C}_p^w, \mathcal{C}_q] \quad (4)$$

where frame q represents a time step p frames into gloss w .

Frame Selection Network To co-articulate between dictionary signs, we propose a Frame Selection Network (FS-NET) that learns to predict the temporal alignment to a continuous signing sequence, $\mathcal{Y} = (y_0, \dots, y_{\mathcal{T}})$ with \mathcal{T} frames (Fig. 2 middle). We note that this is a monotonic sequence-to-sequence task, due to the matching order of signing and the different sequence lengths ($\mathcal{Q} \neq \mathcal{T}$).

Formally, FS-NET predicts a discrete sparse monotonic temporal alignment path, $\hat{\mathcal{A}} \in \mathbb{R}^{\mathcal{Q} \times \mathcal{Q}}$:

$$\hat{\mathcal{A}} = \text{FS-NET}(\mathcal{R}, h_{1:\mathcal{W}}) \quad (5)$$

where $\hat{\mathcal{A}}$ contains binary decisions representing either frame selection or skipping. Fig. 2 shows an example alignment that skips the production of frames 3, 5 and 8 in the output sequence, removing redundant frames to create a smoother co-articulated continuous sequence. We build FS-NET as a transformer encoder [58] with an additional cross-attention to the encoded gloss sequence. To produce the co-articulated continuous signing pose sequence, $\hat{\mathcal{Y}}$, a matrix multiplication can be applied between \mathcal{I} and $\hat{\mathcal{A}}$, as:

$$\hat{\mathcal{Y}} = \mathcal{I} \times \hat{\mathcal{A}} \quad (6)$$

This enables the mapping between varied length sequences, with the end of sequence prediction determined as the alignment selection of the final dictionary frame.

Dynamic Time Warping Supervision In practice, directly predicting the 2D alignment, $\hat{\mathcal{A}}$, provides weak gradients due to the sparse nature of the alignment. We therefore propose to train FS-NET using a Dynamic Time Warping (DTW) supervision signal [2] designed to learn the optimal monotonic temporal alignment. We pre-compute the DTW path, $\mathcal{A}^* = \text{DTW}(\mathcal{Q}, \mathcal{T})$, between the interpolated dictionary sequence, \mathcal{I} , and the target continuous sequence, \mathcal{Y} . Due to the intractability of 2D alignment path prediction, we collapse the alignment down to a 1D sequence during training, $\hat{\mathcal{A}} \in \mathbb{R}^{\mathcal{Q}} = \text{argmax}_q(\hat{\mathcal{A}})$. This enables a temporal mask prediction over \mathcal{I} , selecting which frames of the interpolated dictionary sequence to animate in turn to create a continuous sequence.

We argue that for the majority of sequences (88% for mDGS), $\mathcal{Q} \gg \mathcal{T}$, due to the faster tempo of continuous sign [39]. We thus assume that no frames are added during temporal alignment, only removed. To train FS-NET, we compute a cross entropy loss \mathcal{L}_{CE} between the predicted 1D temporal alignment, $\hat{\mathcal{A}} \in \mathbb{R}^{\mathcal{Q}}$, and the ground truth DTW alignment, $\mathcal{A}^* \in \mathbb{R}^{\mathcal{Q} \times 1}$, as:

$$\mathcal{L}_{CE}(\hat{\mathcal{A}}, \mathcal{A}^*) = -\frac{1}{\mathcal{Q}} \sum_{q=1}^{\mathcal{Q}} \mathcal{A}_q^* \cdot \log(\hat{\mathcal{A}}_q) \quad (7)$$

The final continuous sign pose sequence, $\hat{\mathcal{Y}} = (y_1, \dots, y_{\mathcal{T}})$, is produced as shown in Eq. 6.

3.3. Pose to Video

To generate a photo-realistic sign language video, \mathcal{V}^S , conditioned on the produced sign pose sequence, $\hat{\mathcal{Y}}$, we propose a method for video-to-video signer generation, SIGNGAN (Fig. 2 right). Taking inspiration from [6], in the conditional GAN setup, a generator network, G , competes in a min-max game against a multi-scale discriminator, $D = (D_1, D_2, D_3)$. The goal of G is to synthesise images of similar quality to ground-truth images, in order to fool D . Conversely, the aim of D is to discern the “fake” images from the “real” images. For our purposes, G synthesises images of a signer, v^S given a human pose, y_t , and a style image, S^I .

Following [23], we introduce skip connections to the architecture of G in a U -Net structure [42] between each down-sampling layer i and up-sampling layer $n-i$, where n is the total number of up-sampling layers. Skip connections propagate pose information across the networks, enabling the generation of fine-grained details. Specifically, we add skip connections between each down-sampling layer i and up-sampling layer $n-i$, where n is the total number of up-sampling layers.

Controllable Video Generation To enable training on diverse sign language datasets, we use a style-controllable video generation approach [45]. A style image, S^I , is provided to condition synthesis alongside the pose sequence, as seen in Figure 2. SIGNGAN learns to associate the given style, S , with the person-specific aspects of the corresponding target image, v_t^S , such as the clothing or face, but disentangle the signer-invariant skeleton pose.

Controllable generation allows SIGNGAN to make use of the variability in signer appearance in the data. A multi-modal distribution of sign language videos in different styles, \mathcal{V}^S , can be produced, where $S \in \{1, N_S\}$ represents the styles seen during training².

Hand Keypoint Loss Previous pose-conditioned human synthesis methods have failed to generate realistic and accurate hand images [59]. To enhance the quality of hand synthesis, we introduce a novel loss that operates in the keypoint space, as shown in Figure 3. A pre-trained 2D hand pose estimator [16], H , is used to extract hand keypoints, k_H , from cropped hand regions (*i.e.* a 60x60 patch centered around the middle knuckle), v_H , as $k_H = H(v_H)$. We avoid operating in the image space due to the existence of blurry hand images in the dataset, whereas the extracted keypoints are invariant to motion-blur. A hand keypoint discriminator, D_H , then attempts to discern between the

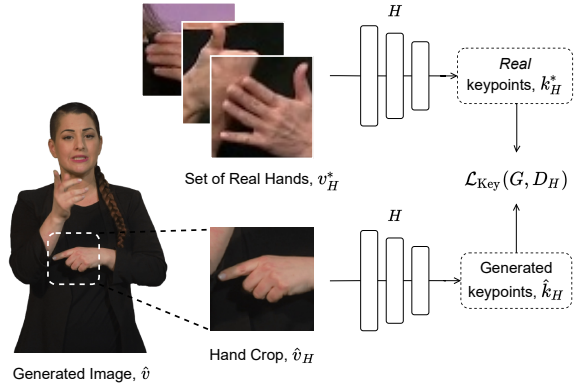


Figure 3. Hand keypoint loss overview. A keypoint discriminator, D_H , compares keypoints extracted from generated hands, \hat{k}_H , and real hands, k_H^* .

“real” keypoints, $k_H^* = H(v_H)$, and the “fake” keypoints, $\hat{k}_H = H(G(y_H))$, leading to the objective:

$$\mathcal{L}_{\text{KEY}}(G, D_H) = \mathbb{E}_{y_H, z_H} [\log D_H(k_H^*)] + \mathbb{E}_{y_H} [\log(1 - D_H(\hat{k}_H))] \quad (8)$$

Full Objective In standard image-to-image translation frameworks [23,63], G is trained using a combination of adversarial and perceptual losses. We update the multi-scale adversarial loss, $\mathcal{L}_{GAN}(G, D)$, to reflect the controllable generation with a joint conditioning on sign pose, y_t , and style image, S^I :

$$\mathcal{L}_{GAN}(G, D) = \sum_{i=1}^k \mathbb{E}_{y_t, z_t} [\log D_i(z_t | y_t, S^I)] + \mathbb{E}_{y_t} [\log(1 - D_i(G(y_t, S^I) | y_t, S^I))] \quad (9)$$

where $k = 3$ reflects the multi-scale discriminator. The adversarial loss is supplemented with two feature-matching losses; $\mathcal{L}_{FM}(G, D)$, the discriminator feature-matching loss presented in pix2pixHD [63], and $\mathcal{L}_{VGG}(G, D)$, the perceptual reconstruction loss [24] which compares pre-trained VGGNet [51] features at multiple layers of the network. Our full SIGNGAN objective, \mathcal{L}_{Total} , is a weighted sum of these, alongside our proposed hand keypoint loss (Eq. 8), as:

$$\mathcal{L}_{Total} = \min_G \left(\max_{D_i} \sum_{i=1}^k \mathcal{L}_{GAN}(G, D_i) \right) + \lambda_{FM} \sum_{i=1}^k \mathcal{L}_{FM}(G, D_i) + \lambda_{VGG} \mathcal{L}_{VGG}(G(y_t, I^S), z_t) + \lambda_{KEY} \mathcal{L}_{KEY}(G, D_H) \quad (10)$$

where $k = 3$ and $\lambda_{FM}, \lambda_{VGG}, \lambda_{KEY}$ weight the contributions of each loss.

²For qualitative examples (e.g. in Fig 4), we share a single signer appearance, as we have consent from this signer to use their appearance for publication purposes.




Figure 4. Example photo-realistic frames with skeleton pose produced using FS-NET and photo-realistic video generated using SIGNGAN.

4. Experiments

In this section, we evaluate our large-scale photo-realistic SLP approach. We outline our experimental setup then perform quantitative, qualitative & user evaluation.

4.1. Experimental Setup

To train our large-scale SLP approach, we set a new translation protocol on the Meine DGS (mDGS) corpus³ [20], a large German Sign Language - Deutsche Gebärdensprache (DGS) linguistic resource capturing free-form signing from 330 deaf participants, with a vocabulary of 10,042 glosses. To adapt the corpus for translation, we segment the free-flowing discourse into 40,230 segments of German sentences, sign gloss translations and sign language videos. We pre-process the mDGS gloss annotations [30] and create two protocols, with either gloss variants included (mDGS-V) or removed (mDGS). We publicly release these translation protocols⁴ to facilitate the future growth in large-scale SLP and SLT research, with further details provided in the appendix. A license must be obtained from the University of Hamburg to use mDGS for computational research.

For additional experiments, we use the benchmark PHOENIX14T dataset [4] from the constrained weather broadcast domain, with setup and skeletal pose configuration as in [44]. We collect exhaustive dictionary examples of every DGS sign present in mDGS and PHOENIX14T, trimmed to remove the sign onset and offset. For samples without expressive mouthings, we insert the facial features present in a example of the respective gloss from the continuous training set. For photo-realistic video generation, we use the C4A dataset [5] due to its high video quality and diverse interpreter appearance. We use a heat-map representation as pose condition, with each skeletal limb plotted on a separate feature channel.

We build our Text to Gloss models with 2 layers, 4 heads and hidden size of 128 and our FS-NET with 2 layers, 4 heads and 64 hidden size. We set the interpolation frames, \mathcal{N}_{LI} , to 5 and the learning rate to 10^{-3} . Our code is based on JoeyNMT [32], and implemented using PyTorch [40].

³With permission from the University of Hamburg.

⁴<https://github.com/BenSaunders27/meineDGS-Translation-Protocols>

4.2. Quantitative Evaluation

4.2.1 Text to Gloss

We first evaluate our Text to Gloss translation described in Sec. 3.1. Table 1 shows a performance of 21.93 BLEU-4 on PHOENIX14T, outperforming [44] (20.23) but falling short of [38] (23.17) who use larger training data. Translation performance is considerably lower on both meineDGS-Variants (mDGS-V) and mDGS due to the larger domain, showing that further research is required to scale the task to larger vocabularies.

Dataset:	DEV SET		TEST SET	
	BLEU-4	ROUGE	BLEU-4	ROUGE
mDGS-V	1.96	24.51	1.16	25.34
mDGS	3.17	32.93	3.08	32.52
PHOENIX14T	21.93	57.25	20.08	56.63

Table 1. Text to Gloss results on mDGS and PHOENIX14T.

4.2.2 Gloss to Pose

Back Translation Back translation has developed as the state-of-the-art SLP evaluation metric [44]. We train an SLT model [7] on PHOENIX14T with skeleton pose sequences generated using our FS-NET approach. Table 2 shows considerable performance gains (43%) compared to baseline methods on the Gloss to Pose task [43, 44, 46, 47]. This highlights the increased comprehension provided by FS-NET compared to baseline end-to-end regression methods, with an ability to overcome the poor quality of the PHOENIX14T dataset. Furthermore, it can be seen that interpolated dictionary sequences (Comparable to that of

Approach:	DEV SET		TEST SET	
	BLEU-4	ROUGE	BLEU-4	ROUGE
Progressive Transformers [44]	11.93	34.01	10.43	32.02
Adversarial Training [43]	13.16	36.75	12.16	34.19
Mixture Density Networks [46]	13.14	39.06	11.94	35.19
Mixture of Motion Primitives [47]	13.32	37.58	12.67	35.61
Interpolated Dictionary Sequence	16.28	38.11	16.27	36.95
FS-NET (Ours)	19.14	40.94	18.78	40.60

Table 2. Back translation results on the PHOENIX14T dataset for the *Gloss to Pose* task.

Approach:	DEV SET		TEST SET	
	BLEU-4	ROUGE	BLEU-4	ROUGE
Progressive Transformers [44]	11.82	33.18	10.51	32.46
Adversarial Training [43]	12.65	33.68	10.81	32.74
Mixture Density Networks [46]	11.54	33.40	11.68	33.19
Mixture of Motion Primitives [47]	14.03	37.76	13.30	36.77
FS-NET (Ours)	16.92	35.74	21.10	42.57

Table 3. Back translation results on the PHOENIX14T dataset for the *Text to Pose* task.

[53]) achieve worse back translation results, highlighting the effect of FS-NET co-articulation for comprehension.

Additionally, Table 3 shows further state-of-the-art results on the full pipeline of text to pose, with an initial text to gloss translation and a subsequent sign animation. This highlights the effectiveness of the full spoken language to photo-realistic video pipeline required for true SLP.

Sign User Evaluation We next perform extensive user evaluation with native signers, animating our skeleton pose outputs using SIGNGAN. All baselines are also generated by SIGNGAN, to alleviate visual differences in comparison. In total, 10 participants completed our sign user evaluation, of which all were fluent signers and 20% were deaf. We provide all the generated user evaluation videos in the supplementary materials.

We first compare the comprehension of FS-NET compared to the state-of-the-art deep learning based SLP method [47]. We show participants pairs of generated videos from the same sequence, asking to select which video was more understandable. Table 4 shows how our productions were unanimously preferred to baselines for both mDGS and PHOENIX14T. This overwhelming result highlights both the increased comprehension of FS-NET, alongside the inability of previous methods to scale to unconstrained domains of discourse.

We next evaluate how understandable our large-scale sign productions are in isolation. We show each participant a produced video alongside a list of 10 signs, of which 5 are signed in the video, and ask them to select which signs they believe are being signed. For FS-NET productions, an average of 4.8 signs were recognised for each video. This shows our productions are easily understandable by native signers, an essential result for accurate large-scale SLP.

Our final user evaluation evaluates how co-articulated

Dataset	FS-NET	Baseline [47]	Equal
mDGS	95%	0%	5%
PHOENIX14T	95%	0%	5%

Table 4. Comprehension user evaluation results, showing the percentage of participants who chose productions from FS-NET or a baseline [47] to be more understandable, or equal.

	FS-NET	Isolated	Equal
Non-Trimmed	100%	0%	0%
Trimmed	40%	47%	13%

Table 5. Co-articulation user evaluation results, showing the percentage of participants who believed the video with the smoothest transitions was from FS-NET, isolated concatenation or equal.

our FS-NET productions are. We show participants two videos of the same sequence, one isolated dictionary sequence and one co-articulated continuous video generated by FS-NET, and ask them to select which they believed had the most smooth transitions between signs. We first evaluate dictionary signs without trimming the sign onset and offset, with Table 5 showing our productions were unanimously preferred. Moving to trimmed dictionary signs, FS-NET productions were preferred 40% of the time, with 13% equal preference. This highlights the effectiveness of FS-NET at improving co-articulation between dictionary signs and temporally aligning to continuous signing sequences.

4.2.3 Pose to Video

Finally, we evaluate our photo-realistic sign language video approach, SIGNGAN. We compare the performance of SIGNGAN against state-of-the-art image-to-image and video-to-video translation methods [6, 54, 62, 63], conditioned on skeletal pose images. We measure the quality of synthesized images using the following metrics; 1) SSIM: Structural Similarity [64] over the full image. 2) Hand SSIM: SSIM metric over a crop of each hand. 3) Hand Pose: Absolute distance between 2D hand keypoints of the produced and ground truth hand images, using a pre-trained hand pose estimation model [16]. 4) FID: Fréchet Inception Distance [21] over the full image.

Baseline Comparison We first compare SIGNGAN to baseline methods for photo-realistic generation given a sequence of ground truth poses as input. Table 6 shows results on the C4A data, with SIGNGAN outperforming all baselines particularly for the Hand SSIM and FID. We believe this is due to the improved quality of synthesized hand images by using the proposed hand keypoint loss.

	SSIM \uparrow	Hand SSIM \uparrow	Hand Pose \downarrow	FID \downarrow
EDN [6]	0.737	0.553	23.09	41.54
vid2vid [62]	0.750	0.570	22.51	56.17
Pix2PixHD [63]	0.737	0.553	23.06	42.57
Stoll <i>et al.</i> [54]	0.727	0.533	23.17	64.01
SIGNGAN (Ours)	0.759	0.605	22.05	27.75

Table 6. Baseline model comparison results of photo-realistic sign language video generation.

	SSIM \uparrow	Hand SSIM \uparrow	Hand Pose \downarrow	FID \downarrow
Baseline	0.743	0.582	22.87	39.33
Hand Discriminator	0.738	0.565	22.81	39.22
Hand Keypoint Loss	0.759	0.605	22.05	27.75

Table 7. Ablation study results of SIGNGAN

Ablation Study We perform an ablation study of SIGNGAN, with results in Table 7. As suggested in Sec. 3.3, the hand discriminator performs poorly for both SSIM and hand SSIM, due to the generation of blurred hands. However, our proposed hand keypoint loss increases model performance considerably and particularly for Hand SSIM, emphasizing the importance of an adversarial loss invariant to blurring.

Perceptual Study We perform an additional perceptual study of our photo-realistic generation, showing participants pairs of 10 second videos generated by SIGNGAN and a corresponding baseline method. Participants were asked to select which video was more visually realistic, with a separate focus on the body and hands. In total, 46 participant completed the study, of which 28% were signers, each viewing 2 randomly selected videos from each baseline. Table 8 shows the percentage of participants who preferred the outputs of SIGNGAN to the baseline method. It can be seen that SIGNGAN outputs were unanimously preferred for both body (96.2% average) and hand (95.6% average) synthesis. Vid2vid [62] was the strongest contender, with our productions preferred only 85% of the time.

Deaf User Evaluation Our final user evaluation compares the comprehension of photo-realistic videos against the previously-used skeletal pose representation [44]. We provided 5 30-second videos of ground-truth skeletal sequences and corresponding photo-realistic videos to deaf participants, asking them to rate each video out of 5 for understandability. Synthesised videos were rated higher for comprehension, at 3.9 compared to 3.2 for skeletal sequences. This suggests that photo-realistic content is more understandable to a deaf signer than a skeleton sequence.

4.3. Qualitative Evaluation

We show example generated photo-realistic frames in Fig. 4, highlighting the production quality. We provide further qualitative evaluation in supplementary materials.

	Body	Hand
EDN [6]	100%	97.8%
vid2vid [62]	85.9%	84.8%
Pix2PixHD [63]	98.9%	100%
Stoll <i>et al.</i> [54]	100%	100%

Table 8. Perceptual study results, showing the percentage of participants who preferred SIGNGAN to the baseline model.

5. Potential Negative Societal Impact

We acknowledge the potential use of SLP technology to remove the reliance on human interpreters. However, we see this work as enabling a larger provision of signed content, especially where interpretation doesn’t exist [13]. We also recognise the potential harm if this technology produced incorrect sign language content, particularly in emergency settings. Although this paper significantly advances the SLP field, we would like to state that SLP technology is still under development and should not yet be relied upon.

6. Conclusion

Large-scale photo-realistic SLP is important to provide high quality signing content to deaf communities. In this paper, we proposed the first SLP method to achieve both large-scale signing and photo-realistic video generation. We proposed FS-NET, which learns to co-articulate between dictionary signs by modelling the optimal temporal alignment to continuous sequences. Furthermore, we proposed SIGNGAN to produce photo-realistic sign language videos. We proposed a novel keypoint-based loss function that improves the quality of hand synthesis, operating in the keypoint space to avoid issues caused by motion blur.

We showed how our approach can scale to unconstrained domains of discourse and be understood by native signers, with considerable state-of-the-art PHOENIX14T back translation performance. Additionally, we performed extensive user evaluation showing our approach increases the realism of interpolated dictionary signs, can be understood by native deaf signers and is overwhelmingly preferred to baseline methods. Finally, we showed that SIGNGAN outperforms all baseline methods for quantitative metrics, human evaluation and native deaf signer comprehension.

Our approach is limited by the current performance of text to gloss translation for large-scale domains. Available gloss annotations are limited, making sign language translation tasks a low-resource machine translation task [38]. Improvements on both architectures and datasets are required to compete with spoken language Neural Machine Translation (NMT) methods. For future work, we plan to tackle spatial co-articulation between dictionary signs.

Acknowledgments

This project was supported by the EPSRC project ExTOL (EP/R03298X/1), SNSF project SMILE-II (CR-SII5 193686) and EU project EASIER (ICT-57-2020-101016982). This work reflects only the authors view and the Commission is not responsible for any use that may be made of the information it contains. We thank Thomas Hanke and University of Hamburg for use of the mDGS data. We thank SWISSTXT for the use of interpreter appearance data.

References

- [1] Guha Balakrishnan, Amy Zhao, Adrian V Dalca, Fredo Durand, and John Guttag. Synthesizing Images of Humans in Unseen Poses. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [2] Donald J Berndt and James Clifford. Using Dynamic Time Warping to Find Patterns in Time Series. In *AAAI-94 Workshop on Knowledge Discovery in Databases*, 1994. 4
- [3] Diane Brentari, Joshua Falk, Anastasia Giannakidou, Annika Herrmann, Elisabeth Volk, and Markus Steinbach. Production and Comprehension of Prosodic Markers in Sign Language Imperatives. *Frontiers in Psychology*, 2018. 2
- [4] Necati Cihan Camgöz, Simon Hadfield, Oscar Koller, Hermann Ney, and Richard Bowden. Neural Sign Language Translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 6
- [5] Necati Cihan Camgöz, Ben Saunders, Guillaume Rochette, Marco Giovanelli, Giacomo Inches, Robin Nachtrab-Ribback, and Richard Bowden. Content4All Open Research Sign Language Translation Datasets. In *IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, 2021. 2, 6
- [6] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody Dance Now. In *Proceedings of the IEEE International Conference on Computer Vision (CVPR)*, 2019. 2, 3, 5, 7, 8
- [7] Necati Cihan Camgöz, Oscar Koller, Simon Hadfield, and Richard Bowden. Sign Language Transformers: Joint End-to-end Sign Language Recognition and Translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 6
- [8] Helen Cooper and Richard Bowden. Large Lexicon Detection of Sign Language. In *International Workshop on Human-Computer Interaction*, 2007. 2
- [9] Stephen Cox, Michael Lincoln, Judy Tryggvason, Melanie Nakisa, Mark Wells, Marcus Tutt, and Sanja Abbott. TESSA, a System to Aid Communication with Deaf People. In *Proceedings of the ACM International Conference on Assistive Technologies*, 2002. 2
- [10] Runpeng Cui, Zhong Cao, Weishen Pan, Changshui Zhang, and Jianqiang Wang. Deep Gesture Video Generation With Learning on Regions of Interest. *IEEE Transactions on Multimedia*, 2019. 2
- [11] Runpeng Cui, Hu Liu, and Changshui Zhang. Recurrent Convolutional Neural Networks for Continuous Sign Language Recognition by Staged Optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [12] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and Controllable Face Image Generation via 3D Imitative-Contrastive Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [13] Jules Dickinson. *Sign Language Interpreting in the Workplace*. Gallaudet University Press, 2017. 8
- [14] Jens Forster, Christoph Schmidt, Thomas Hoyoux, Oscar Koller, Uwe Zelle, Justus H Piater, and Hermann Ney. RWTH-PHOENIX-Weather: A Large Vocabulary Sign Language Recognition and Translation Corpus. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*, 2012. 2
- [15] Jens Forster, Christoph Schmidt, Oscar Koller, Martin Bellgardt, and Hermann Ney. Extensions of the Sign Language Recognition and Translation Corpus RWTH-PHOENIX-Weather. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*, 2014. 2
- [16] Lihao Ge, Zhou Ren, Yuncheng Li, Zehao Xue, Yingying Wang, Jianfei Cai, and Junsong Yuan. 3D Hand Shape and Pose Estimation from a Single RGB Image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 5, 7
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, 2014. 2
- [18] Kirsti Grobel and Marcell Assan. Isolated Sign Language Recognition using Hidden Markov Models. In *IEEE International Conference on Systems, Man, and Cybernetics*, 1997. 2
- [19] Michael Andrew Grosvald. *Long-Distance Coarticulation: A Production and Perception Study of English and American Sign Language*. University of California, Davis, 2009. 3
- [20] Thomas Hanke, Lutz König, Sven Wagner, and Silke Matthes. DGS Corpus & Dicta-Sign: The Hamburg Studio Setup. In *4th Workshop on the Representation and Processing of Sign Languages: Corpora and Sign Language Technologies (CSLT 2010)*, Valletta, Malta, 2010. 2, 6
- [21] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, 2017. 7
- [22] Wencan Huang, Wenwen Pan, Zhou Zhao, and Qi Tian. Towards Fast and High-Quality Sign Language Production. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2021. 2
- [23] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-Image Translation with Conditional Adversarial Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 5
- [24] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual Losses for Real-Time Style Transfer and Super-Resolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 5
- [25] Timor Kadir, Richard Bowden, Eng-Jon Ong, and Andrew Zisserman. Minimal Training, Large Lexicon, Unconstrained Sign Language Recognition. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2004. 2
- [26] Kostas Karpouzis, George Caridakis, S-E Fotinea, and Eleni Efthimiou. Educational Resources and Implementation of a Greek Sign Language Synthesis Architecture. *Computers & Education (CAEO)*, 2007. 2

- [27] Sang-Ki Ko, Chang Jo Kim, Hyedong Jung, and Choongsang Cho. Neural Sign Language Translation based on Human Keypoint Estimation. *Applied Sciences*, 2019. 2
- [28] Oscar Koller. Quantitative Survey of the State of the Art in Sign Language Recognition. *arXiv preprint arXiv:2008.09918*, 2020. 2
- [29] Oscar Koller, Jens Forster, and Hermann Ney. Continuous Sign Language Recognition: Towards Large Vocabulary Statistical Recognition Systems Handling Multiple Signers. *Computer Vision and Image Understanding (CVIU)*, 2015. 2
- [30] Reiner Konrad, Thomas Hanke, Gabriele Langer, Susanne König, Lutz König, Rie Nishio, and Anja Regen. Public DGS Corpus: Annotation Conventions. Technical report, Project Note AP03-2018-01, DGS-Korpus project, IDGS, Hamburg University, 2018. 6
- [31] Marek Kowalski, Stephan J Garbin, Virginia Estellers, Tadas Baltrušaitis, Matthew Johnson, and Jamie Shotton. CONFIG: Controllable Neural Face Image Generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [32] Julia Kreutzer, Joost Bastings, and Stefan Riezler. Joey NMT: A Minimalist NMT Toolkit for Novices. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2019. 6
- [33] Yahui Liu, Marco De Nadai, Gloria Zen, Nicu Sebe, and Bruno Lepri. Gesture-to-Gesture Translation in the Wild via Category-Independent Conditional Maps. In *Proceedings of the 27th ACM International Conference on Multimedia*, 2019. 2
- [34] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose Guided Person Image Generation. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. 2
- [35] Arun Mallya, Ting-Chun Wang, Karan Sapra, and Ming-Yu Liu. World-Consistent Video-to-Video Synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [36] John McDonald, Rosalee Wolfe, Jerry Schnepf, Julie Hochgesang, Diana Gorman Jamrozik, Marie Stumbo, Larian Berke, Melissa Bialek, and Farah Thomas. Automated Technique for Real-Time Production of Lifelike Animations of American Sign Language. *Universal Access in the Information Society (UAIS)*, 2016. 2
- [37] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable Person Image Synthesis with Attribute-Decomposed GAN. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [38] Amit Moryossef, Kayo Yin, Graham Neubig, and Yoav Goldberg. Data Augmentation for Sign Language Gloss Translation. In *Proceedings of the Biennial Machine Translation Summit, International Workshop on Automatic Translation for Signed and Spoken Languages*, 2021. 6, 8
- [39] Lucie Naert, Caroline Larboulette, and Sylvie Gibet. Coarticulation Analysis for Sign Language Synthesis. In *International Conference on Universal Access in Human-Computer Interaction*, 2017. 3, 4
- [40] 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. In *NIPS Autodiff Workshop*, 2017. 6
- [41] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv preprint arXiv:1511.06434*, 2015. 2
- [42] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional Networks for Biomedical Image Segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MIC-CAI)*, 2015. 5
- [43] Ben Saunders, Necati Cihan Camgöz, and Richard Bowden. Adversarial Training for Multi-Channel Sign Language Production. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020. 2, 6, 7
- [44] Ben Saunders, Necati Cihan Camgöz, and Richard Bowden. Progressive Transformers for End-to-End Sign Language Production. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2, 4, 6, 7, 8
- [45] Ben Saunders, Necati Cihan Camgöz, and Richard Bowden. AnonySign: Novel Human Appearance Synthesis for Sign Language Video Anonymisation. *arXiv preprint arXiv:2107.10685*, 2021. 2, 5
- [46] Ben Saunders, Necati Cihan Camgöz, and Richard Bowden. Continuous 3D Multi-Channel Sign Language Production via Progressive Transformers and Mixture Density Networks. 2021. 2, 4, 6, 7
- [47] Ben Saunders, Necati Cihan Camgöz, and Richard Bowden. Mixed SIGNALs: Sign Language Production via a Mixture of Motion Primitives. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 2, 6, 7
- [48] Ben Saunders, Necati Cihan Camgöz, and Richard Bowden. Skeletal Graph Self-Attention: Embedding a Skeletal Inductive Bias into Sign Language Production. *arXiv preprint arXiv:2112.05277*, 2021. 2
- [49] Jérémie Segouat. A Study of Sign Language Coarticulation. *ACM SIGACCESS Accessibility and Computing*, 2009. 2, 3
- [50] Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuilière, and Nicu Sebe. Deformable GANs for Pose-Based Human Image Generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [51] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5
- [52] William C Stokoe. Sign Language Structure. *Annual Review of Anthropology*, 1980. 1, 2
- [53] Stephanie Stoll, Necati Cihan Camgöz, Simon Hadfield, and Richard Bowden. Sign Language Production using Neural Machine Translation and Generative Adversarial Networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2018. 2, 7
- [54] Stephanie Stoll, Necati Cihan Camgöz, Simon Hadfield, and Richard Bowden. Text2Sign: Towards Sign Language Production using Neural Machine Translation and Generative

- Adversarial Networks. *International Journal of Computer Vision (IJCV)*, 2020. 2, 7, 8
- [55] Hao Tang, Song Bai, Li Zhang, Philip HS Torr, and Nicu Sebe. XingGAN for Person Image Generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [56] Hao Tang, Wei Wang, Dan Xu, Yan Yan, and Nicu Sebe. GestureGAN for Hand Gesture-to-Gesture Translation in the wild. In *Proceedings of the 26th ACM International Conference on Multimedia*, 2018. 2
- [57] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. MoCoGAN: Decomposing Motion and Content for Video Generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [58] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. 2, 3, 4
- [59] Lucas Ventura, Amanda Duarte, and Xavier Giró-i Nieto. Can Everybody Sign Now? Exploring Sign Language Video Generation from 2D Poses. In *ECCV Sign Language Recognition, Translation and Production Workshop*, 2020. 2, 3, 5
- [60] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating Videos with Scene Dynamics. In *Advances in Neural Information Processing Systems (NIPS)*, 2016. 2
- [61] Ting-Chun Wang, Ming-Yu Liu, Andrew Tao, Guilin Liu, Jan Kautz, and Bryan Catanzaro. Few-shot Video-to-Video Synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2, 3
- [62] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-Video Synthesis. In *Advances in Neural Information Processing Systems (NIPS)*, 2018. 2, 7, 8
- [63] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 5, 7, 8
- [64] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing*, 2004. 7
- [65] Dongxu Wei, Xiaowei Xu, Haibin Shen, and Kejie Huang. GAC-GAN: A General Method for Appearance-Controllable Human Video Motion Transfer. *IEEE Transactions on Multimedia*, 2020. 3
- [66] Zhenyu Wu, Duc Hoang, Shih-Yao Lin, Yusheng Xie, Liangjian Chen, Yen-Yu Lin, Zhangyang Wang, and Wei Fan. MM-Hand: 3D-Aware Multi-Modal Guided Hand Generative Network for 3D Hand Pose Synthesis. *arXiv preprint arXiv:2010.01158*, 2020. 2
- [67] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-Shot Adversarial Learning of Realistic Neural Talking Head Models. In *Proceedings of the IEEE International Conference on Computer Vision (CVPR)*, 2019. 2, 3
- [68] Jan Zelinka and Jakub Kanis. Neural Sign Language Synthesis: Words Are Our Glosses. In *The IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2020. 2
- [69] Yipin Zhou, Zhaowen Wang, Chen Fang, Trung Bui, and Tamara Berg. Dance Dance Generation: Motion Transfer for Internet Videos. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2019. 3
- [70] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [71] Zhen Zhu, Tengting Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. Progressive Pose Attention Transfer for Person Image Generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2