

ANONYSIGN: Novel Human Appearance Synthesis for Sign Language Video Anonymisation

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



Fig. 1: Novel Appearance Human Synthesis examples generated by ANONYSIGN

Abstract—The visual anonymisation of sign language data is an essential task to address privacy concerns raised by large-scale dataset collection. Previous anonymisation techniques have either significantly affected sign comprehension or required manual, labour-intensive work.

In this paper, we formally introduce the task of Sign Language Video Anonymisation (SLVA) as an automatic method to anonymise the visual appearance of a sign language video whilst retaining the meaning of the original sign language sequence. To tackle SLVA, we propose ANONYSIGN, a novel automatic approach for visual anonymisation of sign language data. We first extract pose information from the source video to remove the original signer appearance. We next generate a photo-realistic sign language video of a novel appearance from the pose sequence, using image-to-image translation methods in a conditional variational autoencoder framework. An approximate posterior style distribution is learnt, which can be sampled from to synthesise novel human appearances. In addition, we propose a novel *style loss* that ensures style consistency in the anonymised sign language videos.

We evaluate ANONYSIGN for the SLVA task with extensive quantitative and qualitative experiments highlighting both realism and anonymity of our novel human appearance synthesis. In addition, we formalise an anonymity perceptual study as an evaluation criteria for the SLVA task and showcase that video anonymisation using ANONYSIGN retains the original sign language content.

I. INTRODUCTION

Sign languages are rich visual languages, and the main mode of communication for Deaf communities. As with the computational study of spoken languages, sign language research requires large-scale datasets to effectively tackle the

tasks of Sign Language Translation (SLT) and Sign Language Production (SLP). However, due to privacy concerns regarding the visual nature of sign language datasets, there is a hesitancy to contribute to data collection [3].

In this work, we formally introduce the task of Sign Language Video Anonymisation (SLVA) as an automatic method to anonymise the visual appearance of a sign language video whilst retaining the original sign language content. Full anonymisation requires two aspects of the signer to be unidentifiable: visual appearance and signing style. In this work, we focus on the former and leave the latter task of signing style anonymisation for future work.

Visual anonymisation of sign language datasets has been shown to increase the willingness of data participation due to the removal of signer-specific appearance [3]. However, previous approaches to sign language video anonymisation, such as pixelation [31], blackening [2] and greyscale filtering [3], have fallen woefully short of the mark. These methods significantly affect sign comprehension or worse, in the case of animoji filtering [3], deride the signers appearance.

To tackle SLVA, we propose ANONYSIGN, a novel automatic method that achieves visual anonymisation of sign language data whilst retaining sign language comprehension. We build ANONYSIGN as a conditional variational generative adversarial network (cVAE-GAN) [1] with a combination of Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs). To remove the original signers appearance, we first extract a skeleton pose sequence from the source sign language video via pose estimation [8].

During training, we use a style encoder to encode style

images into a style code that represents signer-specific appearance features such as facial attributes and body shape. Implicitly, ANONYSIGN learns an approximate posterior style distribution from the space of all pose-independent style images, in the framework of a conditional VAE. We sample from this distribution at inference time to generate multiple novel human appearances not seen during training. To ensure both style consistency and the disentanglement of pose and style, we propose a novel *style loss* that minimises the distance between style distributions extracted from the style image and generated image, respectively.

Finally, we reconstruct a photo-realistic human appearance in the given pose by using a generator with multiple Adaptive Instance Normalisation (AdaIN) layers [18]. AdaIN learns the affine transform parameters relating to the desired style of the sampled style code. This enables ANONYSIGN to puppeteer the novel signer appearance with the original sign language pose sequence, resulting in a visual anonymisation of the original video that retains all sign language content.

We evaluate ANONYSIGN for the SLVA task using a large-scale isolated sign language dataset [13]. We conduct extensive quantitative and qualitative evaluation that showcase the effectiveness of ANONYSIGN for generating novel signer appearances. In addition, we formalise an anonymity perceptual study as an evaluation criteria for the SLVA task, drawing parallels to super-recogniser tests [12]. Finally, we show that our SLVA approach retains the original sign language content using a downstream SLT task.

The **contributions** of this work can be summarised as:

- The formal introduction of the automatic Sign Language Video Anonymisation (SLVA) task
- ANONYSIGN, a novel automatic SLVA method that uses a conditional variational autoencoder framework to generate novel human appearances not seen during training
- A novel style loss that enhances style consistency in generated sign language videos

The rest of this paper is organised as follows: In Section II, we review the literature in sign language anonymisation and human synthesis. In Section III, we outline the proposed ANONYSIGN approach. Section IV presents quantitative and qualitative evaluation, whilst Section V concludes.

II. RELATED WORK

A. Sign Language Production

Combining the domains of vision and language, computational sign language research has been prominent for the last 30 years [41]. Previous research has focused on isolated sign recognition [29], Continuous Sign Language Recognition (CSLR) [4] and SLT [5], [6]. Sign Language Production (SLP), the translation from spoken language to sign language, has been traditionally tackled using graphical avatars [27], [14] with a recent growth of deep learning-based approaches [33], [35], [36], [40], [45].

Photo-realistic SLP has been proposed to overcome the difficulty of skeleton pose comprehension [40], [11]. Saun-

ders *et al.* proposed SIGNGAN [34] to generate photo-realistic continuous sign video sequences, overcoming the previous lack of detail in hand synthesis. In this work, we overcome the privacy issues of signer puppeteering with the synthesis of novel human appearances. In addition to anonymisation, the ability to sample multiple novel appearances further enables diversity in signer generation, which has been highlighted as important by Deaf focus groups [24].

B. Sign Language Video Anonymisation

The field of Neural Machine Translation (NMT) has seen incredible recent progress with the availability of large-scale datasets [7], [10], [26]. However, there has been a lack of large-scale computational sign language datasets collected. Due to its visual medium, the collection of large-scale sign language corpora requires the storing of easily identifiable video data. Bragg *et al.* suggests this introduces privacy concerns over data misuse, that significantly impacts the collection of sign language datasets [3].

There has been little previous research into sign language video anonymisation. Video effects of blackening [2] and pixelation [31] have been used to conceal sensitive information, which is unfeasible for extension to full data corpora. Bragg *et al.* [3] suggest the use of greyscale or animoji filtering for video anonymisation, but neither provides full signer anonymity and both significantly impact sign comprehension. Focus groups with Deaf participants have suggested the use of actors or avatars to reproduce the data [38], but this requires labour-intensive work [20] and often results in non-realistic production [24].

Previously, the idea of complete sign language video anonymisation has been deemed impossible [20], [30], due to the required visibility of both face and hands. Disputing this claim, we formalise the task of Sign Language Video Anonymisation (SLVA) to promote future research and propose a novel automatic visual anonymisation method.

C. Pose Conditioned Human Synthesis

Generative Adversarial Networks (GANs) [16] have achieved impressive results in pose-conditioned human synthesis, for the concurrent tasks of pose transfer [9] and controllable body generation [28], [32]. Latent space encoding has provided disentanglement of shape and appearance [28], and body part sampling [15], [32]. Concurrently, Adaptive Instance Normalisation (AdaIN) [18] has been used to decode style into a generated image, for controllable image synthesis [22] and multimodal human synthesis [28]. We expand pose conditioned human synthesis to the generation of a novel signer appearance not seen during training.

D. Variational Autoencoders

Variational Autoencoders (VAEs) are generative models that enable a sampling from the prior at inference [23]. They have been extended to the conditional framework [39] and combined with GANs [1] to achieve high-quality multimodal synthesis of human pose and appearance [15], [32]. Our method builds on conditional VAEs to synthesise a signer with a novel appearance via latent code sampling.

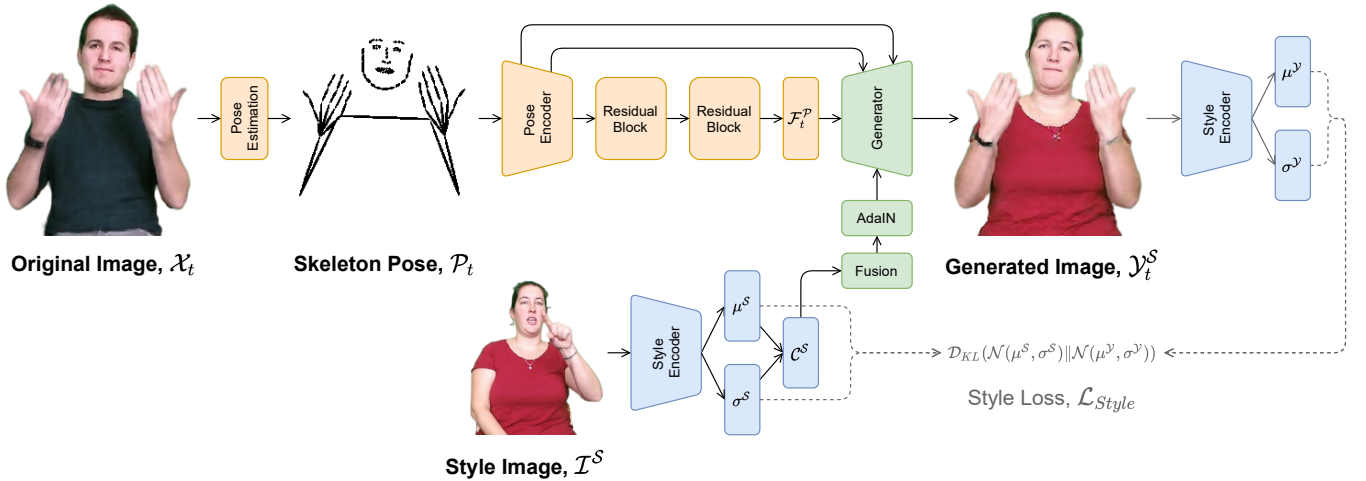


Fig. 3: An overview of the ANONYSIGN network for Sign Language Video Anonymisation (SLVA). At training, skeleton pose, \mathcal{P}_t , is extracted from an original image, \mathcal{X}_t using a pre-trained *Pose Estimator*. The skeleton pose is then encoded by a *Pose Encoder* to a set of pose features, \mathcal{F}_t^P . A Style Image, \mathcal{I}^S , is encoded by a *Style Encoder* to a style distribution, $\mathcal{N}(\mu^S, \sigma^S)$, where a *Style Code*, \mathcal{C}^S , is sampled using the re-parameterisation trick. A photo-realistic signer image, \mathcal{Y}^S , of pose \mathcal{P} and style \mathcal{S} is generated by a *Generator* with multiple AdaIN layers to embed the style from \mathcal{C}^S . To disentangle style and pose, out *Style Loss*, \mathcal{L}_{Style} , applies a KL divergence between the encoded style distributions of \mathcal{I}^S , $\mathcal{N}(\mu^S, \sigma^S)$, and \mathcal{Y}^S , $\mathcal{N}(\mu^Y, \sigma^Y)$.

III. METHODOLOGY

Given a source sign language video, $\mathcal{X} = (\mathcal{X}_1, \dots, \mathcal{X}_T)$ with T frames, our goal is to generate an anonymised target video with the same sign language content, $\mathcal{Y} = (\mathcal{Y}_1^S, \dots, \mathcal{Y}_T^S)$, where \mathcal{S} represents a novel appearance not seen in the training data. We approach this problem as an image-to-image translation task with a conditional variational autoencoder framework. Our proposed ANONYSIGN architecture is outlined in Fig. 3.

During training, we first extract pose, \mathcal{P}_t from the t^{th} frame using a pose estimation model and embed it in a high-dimensional feature space, \mathcal{F}_t^P , using a *Pose Encoder*. We next learn an approximate posterior style distribution from a given style image, \mathcal{I}^S , by using a *Style Encoder* to encode a latent style code, \mathcal{C}^S . Finally, a photo-realistic reconstruction of the source pose in the given style is generated, \mathcal{Y}_t^S , using a generator with AdaIN layers [18] that embed the style code.

To enable sampling of a novel anonymous style at inference time, we build ANONYSIGN using a conditional Variational Autoencoder (VAE) framework [23], [39]. As formulated in latent vector models, we assume that the style of the generated image is solely dependent on the latent variable, \mathcal{C}^S . The goal is then to learn an approximate posterior distribution, $q(\mathcal{C}^S | \mathcal{I}^S)$, from the training style images. A style code, \mathcal{C}^S , can be sampled from this learnt distribution at inference to synthesise a novel signer appearance. A key assumption of our model is that the latent variable \mathcal{C}^S depends *only* on the signer style and is independent of pose. We enforce this condition by proposing a novel *Style Loss* that prompts a similarity between the style distributions encoded from the style image and the generated image, which contain different poses.

A. Pose Encoder (PE)

To remove the original signers appearance, we first use a pre-trained pose estimation model [8] that extracts skeleton pose, \mathcal{P}_t , from each source frame, \mathcal{X}_t , as seen on the left of Fig. 3. We represent both manual and non-manual features¹ of the signer as joint parameters of the upper body, hands and facial landmarks. Alternatively, ANONYSIGN can be used as the photo-realistic output of a SLP model that translates a spoken language sentence (text) into a sign pose sequence, as in [34].

We next encode the extracted pose, \mathcal{P}_t , into a higher-dimensional feature space, \mathcal{F}_t^P , using a *Pose Encoder*, $\text{PE}(\cdot)$:

$$\mathcal{F}_t^P = \text{PE}(\mathcal{P}_t), \quad (1)$$

where \mathcal{F}_P is the extracted pose features.

B. Style Encoder (SE)

During training, we extract style information from a given style image, \mathcal{I}^S , using a *Style Encoder*, $\text{SE}(\cdot)$, as seen at the bottom of Fig. 3. We represent the appearance of the signer as a latent style code, \mathcal{C}^S , of dimension \mathcal{D}_C . We define style as the person-specific appearance aspects such as body shape, face, hair or clothing. The pose information is signer-invariant and is thus learnt independently from style. During training, we randomly select \mathcal{I}^S from the training data that contains identical signer style but different pose to the source frame, \mathcal{X}_t^S .

As for a VAE [23], the style encoder learns an approximate posterior style distribution, $P(\mathcal{C}^S | \mathcal{I}^S)$, from the space of

¹Manual features are the hand shape and motion, whereas non-manuals are the mouthings, facial expressions, gait etc.




Fig. 4: Novel human appearances, \hat{Y}_t^S , can be generated from ANONYSIGN by sampling a style code from the prior distribution, $\mathcal{C}^S \sim \mathcal{N}(0, I)$. Varying \mathcal{P}_t whilst keeping \mathcal{C}^S consistent enables the anonymisation of a full sign language video.

pose-independent style images, \mathcal{I}^S . We assume the style code distribution to be Gaussian, $p(\mathcal{C}^S | \mathcal{I}^S) \equiv \mathcal{N}(\mu, \sigma)$ and thus represent style with parameters $\mu_S, \sigma_S \in \mathbb{R}^{D_c}$. Using the re-parameterisation trick, we can sample the style code, \mathcal{C}^S , as:

$$\mathcal{C}^S \sim \mathcal{N}(\mu^S, \sigma^S) = \text{SE}(\mathcal{I}^S) \quad (2)$$

Style Loss: To ensure the style code is disentangled from pose information, we propose a style consistency loss. As seen on the right of Fig. 3, a KL divergence is measured between the style distribution encoded from the style image, $\mathcal{N}(\mu^S, \sigma^S)$, and the generated image, $\mathcal{N}(\mu^Y, \sigma^Y)$, as:

$$\mathcal{L}_{Style} = \mathcal{D}_{KL}(\mathcal{N}(\mu^S, \sigma^S) \| \mathcal{N}(\mu^Y, \sigma^Y)), \quad (3)$$

where $\mathcal{D}_{KL}(p \| q)$ is the Kullback-Leibler divergence [25] between the probability distributions $p(x)$ and $q(x)$. This provides extra supervision to SE whilst also enforcing style consistency in the generated image.

C. Photo-Realistic Image Generator (G)

Given the encoded pose features, \mathcal{F}_P , and style code, \mathcal{C}^S , a generator model, $G(\cdot)$, decodes a photo-realistic image, \mathcal{Y}_t^S , with the corresponding pose and style:

$$\mathcal{Y}_t^S = G(\mathcal{F}_t^P, \mathcal{C}^S) \quad (4)$$

To embed the style from \mathcal{C}^S into the generated image, we use multiple Adaptive Instance Normalisation (AdaIN) layers [18]. AdaIN layers first apply an instance normalisation [43], followed by a further learnt feature style normalisation. This encourages the features to have similar statistics to the desired style using the learnt affine transform parameters scale, μ_A , and shift, σ_A , as:

$$\text{ADAIn}(x) = \sigma_A \cdot \frac{x - \mu(x)}{\sigma(x)} + \mu_A, \quad (5)$$

where $\mu(x)$ and $\sigma(x)$ are calculated from the pre-normalised feature channel. We apply an additional fusion module to extract the affine parameters from the style code [28]. Additionally, we utilise multiple skip connections from the encoded pose layers, to enable a direct pose conditioning.

D. Training Details

We use the re-parameterisation trick [23] to enable a differentiable pipeline, jointly training PE, SE and G. Expanding from the standard image-to-image framework [21], [44], we train G with the objective \mathcal{L}_{Total} , a combination of adversarial, perceptual and prior losses balanced by their respective weights, λ_* .

Adversarial Loss: To enable the generation of realistic images of both pose and style, we apply a multi-scale adversarial loss [44]. Following [47], we use two discriminators: A pose discriminator, \mathcal{D}_P , that prompts the pose alignment of the generated image, \mathcal{Y}_t^S , with the target pose, \mathcal{P}_t ; and a style discriminator, \mathcal{D}_S , that ensures a style similarity between the generated image, \mathcal{Y}_t^S and the style image, \mathcal{I}^S . The full adversarial loss is formulated as $\mathcal{L}_{Adv} = \lambda_{DP} \mathcal{L}_{DP}(G, \mathcal{D}_P) + \lambda_{DS} \mathcal{L}_{DS}(G, \mathcal{D}_S)$.

Hand Keypoint Loss: To ensure high-quality hand synthesis for sign language comprehension, we apply the hand keypoint loss proposed in [34]. A hand keypoint discriminator, \mathcal{D}_H , is applied that operates in the keypoint space to overcome motion blur. Formally, $\mathcal{L}_{Hand} = \lambda_{DH} \mathcal{L}_{DH}(G, \mathcal{D}_H)$.

Perceptual Loss: We include two feature-matching perceptual losses: $\mathcal{L}_{FM}(G, \mathcal{D}_P)$, the discriminator feature-matching loss presented in pix2pixHD [44]; and $\mathcal{L}_{VGG}(G, \mathcal{D}_P)$, the perceptual reconstruction loss that compares pretrained VGGNet [37] features at multiple layers of the network. The full perceptual loss is formulated as $\mathcal{L}_{Perc} = \lambda_{FM} \mathcal{L}_{FM}(G, \mathcal{D}_P) + \lambda_{VGG} \mathcal{L}_{VGG}(G, \mathcal{D}_P)$.

Prior Loss: To ease the sampling at inference, we encourage the style encoding, $\text{SE}(\mathcal{I}^S)$ to be close to a standard Gaussian distribution, where the prior distribution on the style code, \mathcal{C}^S , is assumed to be $\mathcal{N}(0, I)$. The prior loss is formulated as $\mathcal{L}_{Prior} = \lambda_{KL} \mathcal{D}_{KL}(\text{SE}(\mathcal{I}^S) \| \mathcal{N}(0, I))$.

Our full objective, \mathcal{L}_{Total} , is a weighted sum of these losses, alongside our proposed style loss (Eq. 3):

$$\mathcal{L}_{Total} = \mathcal{L}_{Adv} + \mathcal{L}_{Hand} + \mathcal{L}_{Perc} + \mathcal{L}_{Prior} + \lambda_{Style} \mathcal{L}_{Style} \quad (6)$$

With the re-parameterisation trick, we train ANONYSIGN end-to-end by optimising the parameters of PE, SE and G. The final objective, \mathcal{L}_{Total} , is minimised with respect to PE, SE and G, whilst maximised against \mathcal{D}_P , \mathcal{D}_S and \mathcal{D}_H .

E. Inference: Sampling Novel Appearances

During inference, we *sample* a style code from the prior distribution, $C^S \sim \mathcal{N}(0, I)$, to synthesise a novel signer appearance in a given pose, as seen in Fig. 4. A style can also be extracted from the approximate posterior distribution, as in Fig. 3, by encoding a given style image, I^S , using the style encoder, *i.e.*, $C^S = \mu_S$, where $\mu_S, \sigma_S = SE(\mathcal{I}^S)$. A novel human appearance can be synthesised in this sampled style, \hat{Y}_t^S . By keeping C^S fixed and varying the pose sequence, $\mathcal{P}_{1:T}$, a full-length anonymised sign language video sequence can be generated that retains the sign language content of the original video, $\hat{Y}_{1:T}^S$.

IV. EXPERIMENTS

In this section, we evaluate the performance of ANONYSIGN for the SLVA task. We first outline our experimental setup then perform quantitative, perceptual and qualitative evaluation.

A. Experimental Setup

Dataset: We train our anonymous human synthesis model using the SMILE isolated sign language dataset [13]. As outlined in previous work, we use a heat-map representation as pose condition and perform background segmentation [34]. We also add a random horizontal shifting of up to 20 pixels to provide further data augmentation.

Implementation Details: We train ANONYSIGN for the output resolution of 256×256 . PE() is comprised of four down-sampling and six residual blocks. SE() contains six down-sampling convolutional layers with a latent style code dimension, \mathcal{D}_C , of 64. We omit normalisation layers from SE() to retain the original feature mean and variance that contains important style information [19]. G() consists of four up-sampling blocks, each with a skip connection from the features of PE() and a respective AdaIN layer [18]. The loss weights in Eq. 6 are set empirically to $\lambda_{DP} = 1.0$, $\lambda_{DS} = 1.0$, $\lambda_{DH} = 1.0$, $\lambda_{FM} = 1.0$, $\lambda_{VGG} = 1.0$, $\lambda_{KL} = 0.01$, $\lambda_{Style} = 0.01$. All parts of our network are initialized using Xavier method and trained with an Adam optimizer using default parameters and a learning rate of 2×10^{-3} .

Baselines: We compare ANONYSIGN to the following baseline methods: 1) **VUNet** [15] performs disentanglement between shape and appearance and provides conditional human synthesis. Appearance can be sampled from the stochastic latent representation with a shape preservation. 2) **Pix2PixHD + Noise** [44] is a stochastic extension to pix2pixHD, with random noise added to the input alongside pose. We tile sampled noise, $z \in \mathbb{R}^{\mathcal{D}_C}$, across the \mathcal{D}_P pose condition to create a $\mathcal{D}_P + \mathcal{D}_C$ channel input vector and train pix2pixHD for a stochastic human synthesis. We note a similar baseline was used in [32]. Additionally, we ablate ANONYSIGN with the following conditions: 3) **Deterministic Autoencoder (DAE)** is a deterministic version of ANONYSIGN, where C^S is a single vector output of SE(.) without any KL constraint. At inference, we sample the style code from a unit Gaussian distribution, $C^S \sim \mathcal{N}(0, I)$. 4) **Decoder Style Concatenation** removes the AdaIN layers

and provides style information via a direct concatenation between C^S and \mathcal{F}_t^P before the photo-realistic image generation module. To enable concatenation, we reshape C^S into a multi-channel tensor, with a dimension equal to \mathcal{F}_t^P .

Evaluation Metrics: We evaluate our model and baselines for both diversity and realism. We randomly select 100 poses from the validation set and generate 50 styles for each pose. We evaluate with the following metrics: 1) **Pairwise LPIPS distance:** We compute the Learned Perceptual Image Patch Similarity (LPIPS) [46] between generated samples corresponding to each pose, and take the mean distance over all poses. A higher LPIPS value defines a larger style diversity in the generated images. 2) **FID:** We compute the Fréchet Inception Distance (FID) [17] between the generated samples and a subset of the training data. FID captures how similar the distribution of generated samples is to the ground truth data, with a lower score representing higher realism.

B. Quantitative Evaluation

Baseline Comparisons: We first evaluate ANONYSIGN for the synthesis of novel human appearances. Table I shows that ANONYSIGN significantly outperforms all baselines in terms of both diversity and realism. VUNet [15] struggles with the large appearance diversity in the dataset, with the low LPIPS performance showing it fails to generate diverse appearances. Pix2PixHD+Noise [21] generates more realistic images as indicated by the higher FID score, but the lack of control in the stochastic noise input during training results in a generation of only appearances seen during training.

Additionally, ablation results show the importance of the proposed ANONYSIGN model. A removal of the conditional variational autoencoder framework in the DAE setup results in a style distribution not centered around a prior. Even though realistic appearances can be generated via sampling, they are often of a consistent appearance lacking in diversity as shown in the weaker LPIPS score. The use of decoder concatenation instead of AdaIN layers is not powerful enough to encode style in the generated image, as shown by the low LPIPS performance.

Realism Perceptual Study: Our first perceptual study evaluates the realism of our novel human appearance synthesis. We show each participant 10 pairs of images, each containing a random shuffle of a real appearance and a sampled novel appearance. The participant is asked to select

TABLE I: Baseline model comparison for novel human appearance synthesis. We use LPIPS [46] distance between generated appearances to measure diversity and FID [17] against a subset of the training data to measure realism.

	LPIPS (\uparrow)	FID (\downarrow)
VUNet [15]	6×10^{-4}	157.3
Pix2PixHD + Noise [21]	0.001	117.34
DAE	0.169	81.34
Decoder Concat.	2×10^{-5}	103.4
ANONYSIGN (Ours)	0.243	49.48



Fig. 5: Qualitative baseline comparisons, highlighting the diversity and realism of our novel human appearance synthesis.

which appearance they believe is real, on a Likert scale from 1 to 5, where 1 is ‘left image is definitely real’, 2 is ‘left image is probably real’, 3 is ‘both left and right look similarly as real’ etc. We generate both sets of images using the photo-realistic generator of ANONYSIGN, either from an extracted (real) or sampled (novel) style code, to remove other factors that may affect the decision. We note similar perceptual evaluation has been used in [21]. In total, 24 participants completed the study. Our novel appearances were given a score of 3.36 on average, a majority choice of ‘both left and right look similarly as real’. Discretising the Likert scale down to a preference for each appearance, 23% of participants believed our novel appearances were real (Likert choice 1 or 2) and 32% thought they were equally as real as the original appearances (Likert choice 3). This highlights the realism of our novel human appearance synthesis for the SLVA task.

Anonymity Perceptual Study: Determining whether a novel appearance resembles the original data is an inherently

TABLE II: Anonymity perceptual study results, showing the percentage of time participants chose a novel appearance to be most similar to a *base appearance* (\mathcal{I}^{S1} or \mathcal{I}^{S2}), a *random appearance* or *none*.

	Participant Selection		
	Base	Random	None
Novel Appearances	10.5%	57.5%	32.0%

TABLE III: Downstream SLT results on the RWTH-PHOENIX-Weather-2014T (PHOENIX14T) dataset.

Data:	DEV SET		TEST SET	
	BLEU-4	ROUGE	BLEU-4	ROUGE
Original	13.75	32.81	13.52	32.14
Anonymised	13.58	32.12	13.09	30.24

subjective task. Taking inspiration from super-recogniser tests [12], we formalise an anonymity perceptual study for the SLVA task that uses a photo lineup. We generate a novel appearance by blending between the style codes of two *base appearances*, \mathcal{I}^{S1} and \mathcal{I}^{S2} , thus generating an appearance with a mixed style, \mathcal{Y}^{mix} . Formally, we define the mixing result between style codes \mathcal{C}^{S1} and \mathcal{C}^{S2} as:

$$\mathcal{C}^{mix} = \beta \mathcal{C}^{S1} + (1 - \beta) \mathcal{C}^{S2} \quad (7)$$

where $\beta \in (0, 1)$ and β decreases from 1 to 0 in uniform steps. To evaluate the anonymity of this novel appearance, we show the participant a selection of 8 original appearances found in the training dataset. Two of these appearances are the *base* style images used to generate the novel appearance, \mathcal{I}^{S1} and \mathcal{I}^{S2} , and the other 6 are appearances randomly selected from the training dataset. We ask the participant to choose which original appearance looks most like the novel appearance, with an option for none of the appearances.

We suggest that the novel appearance is anonymous if it bears a closer resemblance to a different original appearance than \mathcal{I}^{S1} or \mathcal{I}^{S2} . Table II provides results of our anonymity perceptual study, showing that participants chose the correct *base* appearance only 10.5% of the time. 57.5% of the time the participant chose a random original appearance, with a 32% choice of none. This highlights that ANONYSIGN can generate convincing novel human appearances that do not resemble the original appearances it was sampled from, suggesting strong anonymity performance for the SLVA task.

Downstream Translation Task: To evaluate the retention of sign language content when performing video anonymisation, we perform a downstream SLT task using both the original and anonymised data. We use the challenging PHOENIX14T dataset released by Camgoz *et al.* [5] and report BLEU and ROUGE scores. We implement the state-of-the-art SLT model [6] as our backbone and represent the photo-realistic videos using EfficientNet-B7 features [42].



Fig. 6: Style code traversal between two original appearances generates novel human appearances using ANONYSIGN

Table III shows that sign language data anonymised by ANONYSIGN achieves comparable SLT BLEU-4 scores to the original data for both the development (13.75 vs. 13.58) and test (13.52 vs. 13.09) sets. This suggests that visual anonymisation using ANONYSIGN retains the original sign language content.

C. Qualitative Evaluation

Baseline Comparisons: Baseline comparisons are shown in Fig. 5. It can be seen that ANONYSIGN generates realistic and diverse novel appearances more consistently than baselines. VUNet [15] struggles to replicate the high-quality and diverse nature of the SMILE dataset, resulting in a blended generation of blurry appearances. Pix2PixHD+Noise [21] generates realistic samples, but experiences mode collapse with a replication of only appearances seen during training.

DAE generates high-quality outputs with little diversity. Due to its deterministic training, the learnt style manifold has an arbitrary unknown prior distribution. This makes the model difficult to sample from, resulting in a consistent appearance synthesis. Decoder style concatenation generates a consistent appearance regardless of the style code, showing that the configuration does not correctly encode style.

Style Code Traversal: Using a trained ANONYSIGN model, we can traverse along the style code manifold between two individuals, shown in Eq. 7 as a blend between style codes C^{S1} to C^{S2} . Two style code traversal examples are shown in Fig. 6. It can be seen that interpolating between known appearances provides a way to create realistic novel human appearances with a blended appearance.

Style Consistency: Due to the proposed *style loss*, the sample of a novel appearance results in a consistent style over multiple frames and poses. Fig. 7 shows an example of style consistency across a pose sequence.



Fig. 7: Style consistency over a pose sequence is achieved with our proposed *style loss*

V. CONCLUSION

In this paper we formally introduced the task of Sign Language Video Anonymisation (SLVA), an essential step towards addressing the privacy concerns risen whilst collecting large-scale sign language datasets. Furthermore, we proposed ANONYSIGN, a novel automatic method for the visual anonymisation of sign language data with a retention of the original sign language content. Using skeletal pose extracted from the source video, ANONYSIGN generates an anonymised photo-realistic video using a conditional variational autoencoder-based approach.

To enable the generation of novel signer appearances not seen during training, we learnt an approximate posterior style distribution over input signer appearances. At inference, we sampled from the style distribution to decode a photo-realistic anonymous signer appearance. In addition, we proposed a novel *style loss* that ensures a style consistency in the generated anonymous video.

As future work, we aim to learn a structured latent space to ease customisation of our novel human appearance synthesis. In addition, we wish to tackle the anonymisation of person-specific signing styles that could be used to identify the original signer.

VI. ACKNOWLEDGEMENTS

This project was supported by the EPSRC project ExTOL (EP/R03298X/1), the SNSF project SMILE2 (CR-SII5_193686) and the 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 would also like to thank NVIDIA Corporation for their GPU grant.

REFERENCES

- [1] J. Bao, D. Chen, F. Wen, H. Li, and G. Hua. CVAE-GAN: Fine-Grained Image Generation through Asymmetric Training. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [2] J. Bleicken, T. Hanke, U. Salden, and S. Wagner. Using a Language Technology Infrastructure for German in order to Anonymize German Sign Language Corpus Data. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*, 2016.
- [3] D. Bragg, O. Koller, N. Caselli, and W. Thies. Exploring Collection of Sign Language Datasets: Privacy, Participation, and Model Performance. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, 2020.
- [4] N. C. Camgoz, S. Hadfield, O. Koller, and R. Bowden. SubUNets: End-to-end Hand Shape and Continuous Sign Language Recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [5] N. C. Camgoz, S. Hadfield, O. Koller, H. Ney, and R. Bowden. Neural Sign Language Translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [6] N. C. Camgoz, O. Koller, S. Hadfield, and R. 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.
- [7] N. C. Camgoz, B. Saunders, G. Rochette, M. Giovanelli, G. Inches, R. Nachtrab-Ribback, and R. Bowden. Content4All Open Research Sign Language Translation Datasets. *arXiv:2105.02351*, 2021.
- [8] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh. OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [9] C. Chan, S. Ginosar, T. Zhou, and A. A. Efros. Everybody Dance Now. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [10] C. Chelba, T. Mikolov, M. Schuster, Q. Ge, T. Brants, P. Koehn, and T. Robinson. One Billion Word Benchmark for Measuring Progress in Statistical Language Modeling. In *Annual Conference of the International Speech Communication Association*, 2014.
- [11] R. Cui, Z. Cao, W. Pan, C. Zhang, and J. Wang. Deep Gesture Video Generation With Learning on Regions of Interest. *IEEE Transactions on Multimedia*, 2019.
- [12] J. D. Dunn, S. Summersby, A. Towler, J. P. Davis, and D. White. UNSW Face Test: A Screening Tool for Super-Recognizers. *PloS one*, 2020.
- [13] S. Ebling, N. C. Camgöz, P. B. Braem, K. Tissi, S. Sidler-Miserez, S. Stoll, S. Hadfield, T. Haug, R. Bowden, S. Tornay, et al. Smile swiss german sign language dataset. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*, 2018.
- [14] S. Ebling and M. Huenerfauth. Bridging the Gap between Sign Language Machine Translation and Sign Language Animation using Sequence Classification. In *Proceedings of SLPAT 2015: 6th Workshop on Speech and Language Processing for Assistive Technologies*, 2015.
- [15] P. Esser, E. Sutter, and B. Ommer. A Variational U-Net for Conditional Appearance and Shape Generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative Adversarial Nets. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [17] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. 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.
- [18] X. Huang and S. Belongie. Arbitrary Style Transfer in Real-Time With Adaptive Instance Normalization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [19] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz. Multimodal Unsupervised Image-to-Image Translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [20] A. Isard. Approaches to the Anonymisation of Sign Language Corpora. In *Proceedings of the LREC Workshop on the Representation and Processing of Sign Languages*, 2020.
- [21] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-Image Translation with Conditional Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [22] T. Karras, S. Laine, and T. Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [23] D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. 2014.
- [24] M. Kipp, Q. Nguyen, A. Heloir, and S. Matthes. Assessing the Deaf User Perspective on Sign Language Avatars. In *ACM SIGACCESS Conference on Computers and Accessibility (ASSETS)*, 2011.
- [25] S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 1951.
- [26] M. Marcus, B. Santorini, and M. A. Marcinkiewicz. Building a Large Annotated Corpus of English: The Penn Treebank. 1993.
- [27] J. McDonald, R. Wolfe, J. Schnepp, J. Hochgesang, D. G. Jamrozik, M. Stumbo, L. Berke, M. Bialek, and F. Thomas. Automated Technique for Real-Time Production of Lifelike Animations of American Sign Language. *Universal Access in the Information Society*, 2016.
- [28] Y. Men, Y. Mao, Y. Jiang, W.-Y. Ma, and Z. Lian. Controllable Person Image Synthesis with Attribute-Decomposed GAN. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [29] O. Özdemir, N. C. Camgöz, and L. Akarun. Isolated Sign Language Recognition using Improved Dense Trajectories. In *Proceedings of the Signal Processing and Communication Application Conference (SIU)*, 2016.
- [30] J. Quer and M. Steinbach. Handling Sign Language Data: The Impact of Modality. *Frontiers in Psychology*, 2019.
- [31] A. Rudge. Analysing British Sign Language through the Lens of Systemic Functional Linguistics. 2018.
- [32] K. Sarkar, L. Liu, V. Golyanik, and C. Theobalt. HumanGAN: A Generative Model of Humans Images. *arXiv preprint arXiv:2103.06902*, 2021.
- [33] B. Saunders, N. C. Camgoz, and R. Bowden. Adversarial Training for Multi-Channel Sign Language Production. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020.
- [34] B. Saunders, N. C. Camgoz, and R. Bowden. Everybody Sign Now: Translating Spoken Language to Photo Realistic Sign Language Video. *arXiv preprint arXiv:2011.09846*, 2020.
- [35] B. Saunders, N. C. Camgoz, and R. Bowden. Progressive Transformers for End-to-End Sign Language Production. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [36] B. Saunders, N. C. Camgoz, and R. Bowden. Continuous 3D Multi-Channel Sign Language Production via Progressive Transformers and Mixture Density Networks. In *International Journal of Computer Vision (IJCV)*, 2021.
- [37] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv:1409.1556*, 2014.
- [38] J. L. Singleton, G. Jones, and S. Hanumantha. Toward Ethical Research Practice With Deaf Participants. *Journal of Empirical Research on Human Research Ethics*, 2014.
- [39] K. Sohn, H. Lee, and X. Yan. Learning Structured Output Representation using Deep Conditional Generative Models. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [40] S. Stoll, N. C. Camgoz, S. Hadfield, and R. Bowden. Text2Sign: Towards Sign Language Production using Neural Machine Translation and Generative Adversarial Networks. *International Journal of Computer Vision (IJCV)*, 2020.
- [41] S. Tamura and S. Kawasaki. Recognition of Sign Language Motion Images. *Pattern Recognition*, 1988.
- [42] M. Tan and Q. Le. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In *International Conference on Machine Learning (ICML)*, 2019.
- [43] D. Ulyanov, A. Vedaldi, and V. Lempitsky. Improved Texture Networks: Maximizing Quality and Diversity in Feed-forward Stylization and Texture Synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [44] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. 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.
- [45] J. Zelinka and J. Kanis. Neural Sign Language Synthesis: Words Are Our Glosses. In *The IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2020.
- [46] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [47] Z. Zhu, T. Huang, B. Shi, M. Yu, B. Wang, and X. Bai. Progressive Pose Attention Transfer for Person Image Generation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.