

DexAvatar: 3D Sign Language Reconstruction with Hand and Body Pose Priors

Kaustubh Kundu¹, Hrishav Bakul Barua^{1,2}, Lucy Robertson-Bell¹, Zhixi Cai¹, Kalin Stefanov¹
¹Monash University
²TCS Research

{kaustubh.kundu, hrishav.barua, lucy.robertson-bell, zhixi.cai, kalin.stefanov}@monash.edu



Figure 1. **DexAvatar** recovers bio-mechanically accurate 3D hand and body poses from monocular sign language videos.

Abstract

The trend in sign language generation is centered around data-driven generative methods that require vast amounts of precise 2D and 3D human pose data to achieve an acceptable generation quality. However, currently, most sign language datasets are video-based and limited to automatically reconstructed 2D human poses (i.e., keypoints) and lack accurate 3D information. Furthermore, existing state-of-the-art for automatic 3D human pose estimation from sign language videos is prone to self-occlusion, noise, and motion blur effects, resulting in poor reconstruction quality. In response to this, we introduce DexAvatar, a novel framework to reconstruct bio-mechanically accurate fine-grained hand articulations and body movements from in-the-wild monocular sign language videos, guided by learned 3D hand and body priors. DexAvatar achieves strong performance in the SGNify motion capture dataset, the only benchmark available for this task, reaching an improvement of 35.11% in the estimation of body and hand poses compared to the state-of-the-art. The official website of this work is: <https://github.com/kaustesseract/DexAvatar>.

1. Introduction

Sign languages are the primary mode of communication for approximately 466 million Deaf or hard-of-hearing individuals worldwide [41]. They are spatio-temporal languages that utilize the visual-gestural modality to convey meaning through manual hand articulations in combination with non-manual elements like the face and body. Similarly to spoken languages, sign languages follow linguistic rules [5], but lack standardized written forms e.g., American Sign Language is not a visual representation of spoken English.

Existing sign language datasets [15, 24, 27, 33, 48, 54] provide videos and 2D keypoints and have contributed significantly to sign language generation [4, 51, 52]. However, 2D representations cannot capture the full spatial complexity of signing. Depth information and hand-body contact are critical for conveying meaning, yet different 3D hand configurations can project to identical 2D keypoints. There is therefore a clear need for datasets with accurate 3D information to support realistic spatial sign language generation. Recent advances in whole-body parametric models, such as SMPL-X [44], enable expressive whole-body mesh recovery that jointly estimates 3D body and hand pose and facial expressions from videos.

Despite significant progress, 3D hand pose estimation remains challenging for sign language videos. Hand movements in sign language are more complex and intricate compared to everyday scenarios. The main challenges arise from rapid and complex hand articulations [7], frequent hand-to-hand and hand-to-body interactions that cause self-occlusions [42], and motion blur from fast signing motions. These factors hinder accurate hand pose reconstruction across many video frames, and existing whole-body and hand-only mesh recovery methods struggle to estimate reliable poses from sign language videos.

Current approaches to whole-body mesh recovery fall into two categories: regression-based and optimization-based methods. Regression methods predict mesh parameters directly and are efficient [9, 35, 36, 55], but are usually trained on general-purpose data and often fail to capture the hand articulations specific to signing. Optimization methods fit parametric models using priors and multi-term objectives and are more computationally intensive, but can incorporate sign language aware constraints to yield accurate and stable hand poses under self-occlusion. We therefore introduce **DexAvatar**, an optimization framework that incorporates signing-based hand and body priors to reconstruct 3D signing avatars from monocular videos.

A key limitation of existing work [6, 44] is that priors are trained on general-purpose datasets that fail to capture the distinctive characteristics of signed communication. In sign languages, meaning relies heavily on precise hand shape and orientation, yet accurate 3D hand data are not publicly available. To address this, we collected a motion capture sign language dataset to capture fine finger articulations and train our proposed hand prior, SignHPoser. The proposed body prior, SignBPoser, is trained with a subset of the 3D data published in [65], reconstructed from the How2Sign [14] dataset. The pretrained body and hand priors can be utilized in any regression- or optimization-based approach for whole-body mesh recovery from sign language videos. Our main contributions are:

- We introduce two sign language-aware pose priors, SignHPoser for hands and SignBPoser for body, trained to learn compact latent spaces that preserve phonologically meaningful variations and discourage anatomically implausible configurations.
- We integrate the priors into DexAvatar, an optimization pipeline that reconstructs 3D signing avatars from monocular videos. DexAvatar employs the priors as differentiable regularizers, together with temporal consistency and contact-aware terms, to stabilize estimation under self-occlusion and noisy 2D keypoints in upper-body-only and one-handed signing videos.
- Extensive experiments show that DexAvatar consistently achieves lower reconstruction errors compared to strong baselines for whole-body and hand-only mesh recovery.

2. Related Work

A foundational element of our ability to interact with others is the recognition of body poses. Similarly, for intelligent machines, the accurate understanding of body poses is essential for interacting with humans.

2D Human Pose Estimation. Research on 2D keypoints estimation has evolved from early tree- and random forest-based models [50, 57, 58] to recent deep learning methods [13, 22, 28, 61]. Recently, 3D joint estimation methods [19, 37, 56, 69] have shown promise in overcoming the depth ambiguity inherent to 2D approaches. However, these approaches predict a sparse set of skeletal keypoints, limiting expressivity.

3D Human Pose Reconstruction. Expressive 3D body models such as SMPL-X [44], MANO [47], and FLAME [34] have enabled research on estimating 3D body meshes. For holistic 3D mesh recovery, prior work leverages multi-task learning [35], hybrid priors [55], and attention mechanisms [9, 36]. For hand-specific recovery, methods incorporate auxiliary components for hand localization and bounding box refinement [46, 67], or integrate optimization-based refinement and temporal filtering [12, 45] to improve stability and accuracy. Although these methods offer strong potential for expressive sign language reconstruction, they are generic and fail to handle long self-occlusions, enforce realistic hand-hand and hand-body contact, and capture fine finger articulations in sign language videos.

3D Avatar Reconstruction for Sign Language. SG-Nify [20] introduced one of the first dedicated pipelines for whole-body mesh recovery from sign language videos, adding linguistic priors that constrain 3D hand pose to resolve ambiguities. Built atop SMPLify-X [44], it estimates SMPL-X [44] parameters from images. However, reliance on pseudo-ground truth from off-the-shelf 2D keypoint detectors such as ViTPose [62] and MediaPipe [39] limits accuracy. The method also struggles with severe self-occlusions, intricate hand-hand and hand-body interactions, motion blur, and frequent cropping—all inherent challenges in sign language reconstruction.

OSX [36] proposed the UBody pipeline for downstream upper body reconstruction tasks, including sign language, gesture, and emotion generation. While OSX demonstrates strong performance on general-purpose datasets, it often misses nuanced articulations and produces unrealistic hand poses in sign language videos.

Neural Sign Actors [6] and SignAvatars [65] used OSX for initialization. The former used OSX to initialize pose and shape for an optimization-based pipeline and added a mesh prior based on a principal component analysis-based space derived from the AMASS [40] and ARCTIC [16] datasets. The latter fused OSX with ACR [64] and PARE [32] and used a pseudo-ground truth 2D key-

points [39, 61] within a SMPLify-X [44] framework.

EVA [23] leveraged multiple sources as pseudo-ground truth, including initial estimates of camera parameters, SMPL-X parameters [9], 2D keypoints [45], and 3D hand parameters [9] from off-the-shelf tools [63]. Despite their effectiveness in denoising and generating plausible motions, existing pipelines suffer from critical limitations for sign language reconstruction. They rely on general-purpose pose detectors and body priors not tailored to sign language articulations, leading to domain shift, depth ambiguity, and over-regularization. Moreover, these methods are trained on everyday human motion and lack exposure to the linguistic and cultural nuances of sign languages.

3. Method

We introduce DexAvatar, a method for reconstructing 3D whole-body pose and mesh from monocular sign language videos. In sign languages, meaning is conveyed through the coordinated use of hand gestures, facial expressions, and upper body movements within 3D space. To represent these modalities jointly, we employ the SMPL-X [44] parametric model. DexAvatar, similar to SMPLify-X [44], is an optimization-based method that uses pseudo-ground truth from off-the-shelf tools, including initial camera parameters, SMPL-X parameters, 2D keypoints, and 3D hand parameters [9, 28, 45]. Since off-the-shelf tools are not dedicated to sign language and may fail to capture fine-grained hand articulations and upper-body expressions, we introduce specialized priors: SignHPoser for hands and SignBPoser for body.

3.1. Preliminaries

SMPLify-X is a widely used optimization method based on SMPL-X that reconstructs 3D body and hand poses and facial expressions from monocular images (see Section 1 in the Supplementary Material). The reconstruction is achieved by minimizing the following objective function with respect to β , ψ , and θ :

$$\mathcal{L} = \mathcal{L}_{\text{joint}} + \lambda_{\zeta} \mathcal{L}_{\zeta} + \lambda_{\text{pen}} \mathcal{L}_{\text{pen}}. \quad (1)$$

$\mathcal{L}_{\text{joint}}$, defined in Eq. (2), minimizes the error between the detected 2D keypoints \mathcal{K} and the corresponding 3D keypoints $P(\mathcal{D})$ projected onto the image plane via $P(\cdot)$. \mathcal{J} represents the set of 3D joints, and $P(\cdot)$ projects a 3D joint $\mathcal{D}_i \in \mathbb{R}^3$ from the world coordinate system to the 2D image coordinates. The term \mathcal{K}_i represents the corresponding 2D keypoint, which may come from pseudo or ground truth annotations. ω_i denotes the confidence of \mathcal{K}_i , γ_i is a predefined weight for joint \mathcal{D}_i , and ψ represents a robust Geman-McClure loss function [18] to prevent the disturbance from noisy supervision signals. The objective is given as:

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

\mathcal{L}_{ζ} , defined in Eq. (3), is a zero-mean Gaussian prior on the VPoser learnable body pose embedding $\zeta \in \mathbb{R}^d$, where the negative log-prior yields an ℓ_2 penalty on each latent dimension. Here, $i \in \{1, \dots, d\}$ indexes the latent dimensions and σ_i^2 denotes the variance of the i -th dimension under the diagonal Gaussian prior. The loss term is added to the total objective with weight λ_{ζ} .

$$\mathcal{L}_{\zeta} = \sum_{i=1}^d \frac{\zeta_i^2}{\sigma_i^2}, \quad \text{with } \sigma_i^2 = 1 \forall i. \quad (3)$$

\mathcal{L}_{pen} , defined in Eq. (4), prevents self-collision by first detecting colliding face pairs \mathcal{C} with a bounding volume hierarchy and then penalizing the bi-directional intrusion depth for each pair (f_s, f_t) . For any face f , $\mathcal{V}(f)$ is its vertex set and $\Psi_f(v)$ is the conic signed distance to f (negative inside), so $\max(0, -\Psi_f(v))$ measures how far a vertex penetrates. The loss sums these depths for vertices of f_s in the f_t field and vice versa, optionally normalizing by $|\mathcal{C}|$, and is added to the total objective with weight λ_{pen} .

$$\mathcal{L}_{\text{pen}} = \frac{1}{|\mathcal{C}|} \sum_{(f_s, f_t) \in \mathcal{C}} \left[\sum_{v \in \mathcal{V}(f_s)} (\max(0, -\Psi_{f_t}(v)))^2 + \sum_{v \in \mathcal{V}(f_t)} (\max(0, -\Psi_{f_s}(v)))^2 \right]. \quad (4)$$

3.2. Data Preprocessing

3.2.1. Body Data

We use the 3D body data provided by SignAvatars [65], derived from the How2Sign [14] dataset. Since these data constitute pseudo-ground truth, they may contain residual noise and bias in the recovered motions. Training SignBPoser on such data without additional safeguards can degrade performance. Therefore, we filter out implausible poses using established bio-mechanical constraints of human body joints [21, 31], as shown in Fig. 3. Since major movements in sign languages employ a subset of body joints, that is, shoulders, elbows/forearms, and wrists, we focus on those in our preprocessing. For these joints, we define plausible ranges of motion constrained by physiological degrees of freedom and signer space [8, 60]—a torso-centric 3D region where signs are produced and perceived (see Section 2 of Supplementary Material for details). Specifically, we remove frames where joint angles fall outside the defined ranges of motion, such as overly elevated, excessively retracted, or fully outstretched arms that are inconsistent with real-life signing.

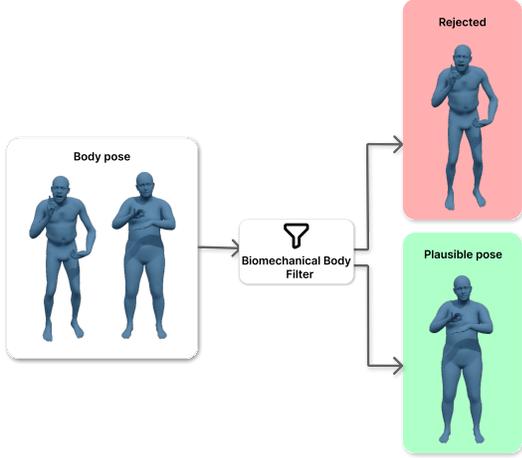


Figure 3. **Bio-mechanical body filter.** For body data from [14], we enforce joint range of motion and signer space constraints on shoulders, elbows/forearm, and wrists. Frames that violate these envelopes are rejected, and only plausible body poses are retained for training.

the reference mesh M , measured as a per-vertex squared ℓ_2 error,

$$\mathcal{L}_{\text{mesh}} = \|M - \hat{M}\|_2^2. \quad (8)$$

$\mathcal{L}_{\text{orth}}$, defined in Eq. (9), constrains rotations to be valid by enforcing orthogonality and unit determinant, where $\hat{R} \in \mathbb{R}^{3 \times 3}$ denotes the output rotation matrix for each joint and I is the 3×3 identity,

$$\mathcal{L}_{\text{orth}} = \|\hat{R}\hat{R}^\top - I\|_2^2. \quad (9)$$

\mathcal{L}_{reg} , defined in Eq. (10), discourages overfitting by penalizing the ℓ_2 norm of the trainable parameters ϕ ,

$$\mathcal{L}_{\text{reg}} = \|\phi\|_2^2. \quad (10)$$

$\mathcal{L}_{\text{biomech}}$, defined in Eq. (11), enforces per-joint anatomical limits. For joint j , let $\bar{\theta}_{j,\min}$ and $\bar{\theta}_{j,\max}$ denote the lower and upper angular bounds. The penalty is zero when θ_j lies within $[\bar{\theta}_{j,\min}, \bar{\theta}_{j,\max}]$ and grows quadratically outside this range. \max is taken element-wise. We apply this to $J = 6$ body joints for SignBPoser and $J = 15$ hand joints for SignHPoser,

$$\mathcal{L}_{\text{biomech}} = \sum_{j=1}^J \left\| \max\{\theta_j - \bar{\theta}_{j,\max}, \bar{\theta}_{j,\min} - \theta_j, \mathbf{0}\} \right\|_2^2. \quad (11)$$

3.4. Optimization

Initializing with reliable SMPL-X estimates improves avatar reconstruction from videos. We therefore initialize

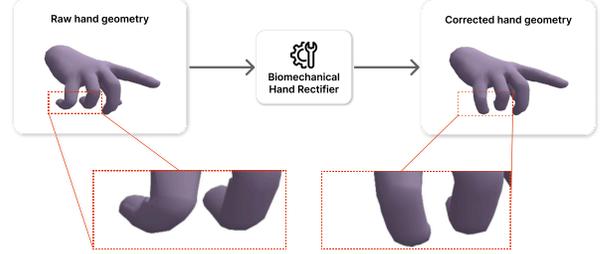


Figure 4. **Bio-mechanical hand rectifier.** For raw mocap hand data, we correct implausible joint configurations by enforcing per-joint limits on bending, splaying, and twisting (15 hand joints). The rectifier outputs corrected hand geometry for training.

DexAvatar optimization with off-the-shelf SMPL-X parameters, 2D keypoints, and camera estimates [9, 28, 63], as shown in Fig. 2. However, these detectors often struggle with sign language data, requiring additional improvements. Hence, the DexAvatar optimization objective follows Eq. (1) by retaining the joint loss and interpenetration term from SMPLify-X, while replacing the generic VPoser prior with SignBPoser for the body and adding the SignHPoser prior for the hands. In addition, we introduce temporal consistency and bio-mechanical penalties for the hand and body, tailored to signing dynamics. We also restrict lower-body joints and deactivate the non-dominant arm for one-handed signs to improve stability. The optimization minimizes the following objective:

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{\text{joint}} + \lambda_1 \mathcal{L}_{\text{bprior}} + \lambda_2 \mathcal{L}_{\text{hprior}} + \lambda_3 \mathcal{L}_{\text{pen}} \\ & + \lambda_4 \mathcal{L}_{\text{temp}} + \lambda_5 \mathcal{L}_{\text{bbiomech}} + \lambda_6 \mathcal{L}_{\text{hbiomech}}. \end{aligned} \quad (12)$$

$\mathcal{L}_{\text{joint}}$ adopts the formulation in Eq. (2), leveraging 2D keypoints from Sapiens [28] for the body and HaMeR [45] for the hands. Since signing typically involves minimal to no lower-body motion, we set $\omega_i = 0$ in Eq. (2) for all lower body joints to exclude them from optimization. Furthermore, we utilize the sign classifier from [20] to distinguish between one-handed and two-handed signs. In the case of one-handed signing, our Hand Decision Maker (see Fig. 2) disables optimization of the non-dominant arm (shoulder, elbow, and wrist) and non-dominant hand by assigning $\omega_i = 0$ in Eq. (2). This strategy prevents spurious updates and ensures the optimization focuses only on the active parts of the body where refinement is necessary.

$\mathcal{L}_{\text{bprior}}$ (for the body), defined in Eq. (13), integrates SignBPoser to regularize infeasible body poses. Since SignBPoser is trained on sign language data, it provides compact mapping from low-dimensional ζ to axis-angle representation of the body pose θ_b . To better optimize the low-dimensional embedding, the estimated body pose $\hat{\theta}_b$ from SMPLerX [9] is treated as a supervisory signal, in coordination with the added regularization term $\bar{\zeta}$ similar to

Eq. (3) with a weighting of $\lambda_{\bar{\zeta}}$. The body prior loss term is formulated as follows,

$$\mathcal{L}_{bprior} = \psi(\theta_b - \hat{\theta}_b) + \lambda_{\bar{\zeta}} \mathcal{L}_{\bar{\zeta}}. \quad (13)$$

\mathcal{L}_{hprior} (for the hand), defined in Eq. (14), integrates SignHPoser to regularize infeasible hand poses. We follow the same formulation as Eq. (13). We treat the left and right hand pose $\hat{\theta}_h$ from HaMeR [45] as a supervisory signal along with ϵ , which is the low-dimensional representation for each hand (see Eq. (3)).

$$\mathcal{L}_{hprior} = \psi(\theta_h - \hat{\theta}_h) + \lambda_{\epsilon^l} \mathcal{L}_{\epsilon^l} + \lambda_{\epsilon^r} \mathcal{L}_{\epsilon^r}. \quad (14)$$

\mathcal{L}_{temp} ensures temporal consistency across frames, and incorporates the pose parameters from the previous frame, denoted as θ_b^{pre} . We enforce smooth transitions between frames by penalizing discrepancies in the motion of corresponding joints.

$$\mathcal{L}_{temp} = \psi(\theta_b - \theta_b^{pre}). \quad (15)$$

$\mathcal{L}_{bbiomech}$ and $\mathcal{L}_{hbiomech}$ incorporate bio-mechanical loss constraints for the body [8, 60] and the hands [10] as additional supervision during optimization. The associated loss term is unchanged and follows Eq. (11).

4. Experiments

Datasets. We evaluated DexAvatar on the motion capture dataset used in SGNify [20]. The dataset consists of 57 German signs. Following the standard evaluation protocol for this dataset, we evaluate on central portions of each sign from the raw videos and compute the quantitative results on only these central frames (in total, 2,872 frames).

Evaluation Metrics. Following prior work SGNify [20] and Neural Sign Actors [6], we evaluated DexAvatar using mean vertex-to-vertex error (TR-V2V), restricting computation to vertices above the pelvis. We report results for three regions: Upper Body (excluding the face), Left Hand, and Right Hand. To assess SignBPoser and SignHPoser independently, we compute MPJPE and MPVPE on the recovered joints and meshes.

Implementation Details. DexAvatar is implemented in PyTorch [25] and optimized using LBFGS [43]. All experiments are performed on NVIDIA-RTX 4090 with 24 GB GPU memory and 64 GB CPU memory. SignBPoser and SignHPoser are built using an encoder-decoder VAE, each with 3 linear layers with an embedding size of 512. During training, we use the Adam optimizer [29] and the learning rate of $1e-3$.

Table 1. **Quantitative results.** We compare DexAvatar with the current state-of-the-art methods on TR-V2V error (mm). We report results in three regions, *i.e.*, UBody(-F): Upper Body excluding the face, LHand: Left Hand, and RHand: Right Hand. EVA* denotes our modification of EVA [23] to accommodate one-handed signs.

Method	UBody (-F)↓	LHand↓	RHand↓
FrankMoCap [49]	78.07	20.47	19.62
PIXIE [17]	60.11	25.02	22.42
PyMAF-X [66]	68.61	21.46	19.19
SMPLify-SL [20]	56.07	22.23	18.83
SGNify [20]	55.63	19.22	17.50
OSX [36]	47.32	18.34	18.12
Neural Sign Actors [6]	46.42	16.17	15.23
EVA*	40.38	13.73	13.68
DexAvatar (Ours)	30.13	13.53	13.08

Table 2. We conduct an ablation study to understand the effectiveness of SignBPoser within DexAvatar. We conduct experiments with three different variants. BP_u: trained on unfiltered data, BP_f: trained on bio-mechanically filtered data, BP_{f+bio}: trained on filtered data and with body bio-mechanical loss. We evaluate over four vertex subsets: FBody: Full Body (10,475 vertices), UBody: Upper Body (vertices above the pelvis), UBody (-H): Upper Body without head (above-pelvis without the head).

Method	FBody↓	UBody↓	UBody (-H)↓	UBody (-F)↓
BP _u	43.18	29.95	44.72	34.06
BP _f	42.32	26.78	41.35	30.28
BP _{f+bio}	42.38	26.93	41.88	30.44

5. Results

5.1. Quantitative Results

Table 1 reports a comparison of DexAvatar against state-of-the-art baselines on the SGNify dataset. DexAvatar achieves the best performance across regions, outperforming FrankMocap [49], PyMAF-X [66], PIXIE [17], SMPLify-SL [20], SGNify [20], OSX [36], and Neural Sign Actors [6]. As the dataset contains both one- and two-handed signing, we modify EVA [23] to accommodate one-handed signs, and denote it as EVA*. DexAvatar surpasses Neural Sign Actors on the left and right hands by 16.32% and 14.11%, respectively, and yields a substantial **35.11%** improvement on the upper body.

Effectiveness of SignBPoser. Table 2 summarizes ablation results of the body prior, SignBPoser, within DexAvatar. We perform a hyperparameter search (see section 4 of Supplementary Material for details) for three variants and report results based on the best configurations: BP_u trained on unfiltered data, BP_f trained on bio-mechanically

Table 3. Ablation study of SignHPoser within DexAvatar.

Method	UBody (-F)↓	LHand↓	RHand↓
HP _u	31.34	14.19	13.92
HP _f	30.17	13.55	13.06
HP _{f+bio}	30.13	13.53	13.08

filtered data, and BP_{f+bio} trained on filtered data with bio-mechanical loss. The first two rows of Table 2 show the results for BP_u and BP_f. BP_f consistently outperforms BP_u across all metrics, with relative error reductions of **2.0%** (FBody), **10.6%** (UBody), **7.5%** (UBody (-H)), and **11.1%** (UBody (-F)). This establishes the efficacy of our data preprocessing using bio-mechanical constraints. Using a prior trained on the filtered data with bio-mechanical loss in BP_{f+bio} yields a slight degradation (+0.14% FBody, +0.56% UBody, +1.28% UBody (-H), +0.53% UBody (-F), suggesting mild over-regularization during optimization. Finally, adding bio-mechanical loss during optimization (excluded from Table 2) while retaining BP_f produces the best results on all subsets, with relative error reductions of **0.17%** (FBody), **0.37%** (UBody), **0.05%** (UBody (-H)), and **0.33%** (UBody (-F)) compared to BP_f.

Effectiveness of SignHPoser. We evaluate the hand prior SignHPoser using the best performing configuration for the body prior, that is BP_f and body-based bio-mechanical loss during optimization. Table S2 summarizes ablation results. We perform a hyperparameter search (see section 4 of Supplementary Material for details) for three hand prior variants and report results based on the best configurations for HP_u trained on uncorrected hand data, HP_f trained on the bio-mechanically corrected data, and HP_{f+bio} trained on the corrected data with hand bio-mechanical loss. The first two rows of Table S2 show the results for HP_u and HP_f variants. It can be observed that HP_f outperforms HP_u on all metrics, with relative error reductions of **3.7%** on Upper Body, **4.5%** on Left Hand, and **6.2%** on Right Hand. This demonstrates the importance of the correction process using bio-mechanical constraints on the hand data. Adding a bio-mechanical regularizer to the filtered prior in HP_{f+bio} increases accuracy compared to HP_f, yielding slight improvements on Upper Body (0.13%) and Left Hand (0.15%). However, Right Hand performance degrades slightly (0.2%). This indicates that introducing bio-mechanical constraints provides useful physical regularization in the fitting process. Please see section 5 of Supplementary Material for ablation of SignHPoser with Vposer.

5.2. Qualitative Results

In Fig. 5, we present qualitative results of DexAvatar on the SGNify [20] motion capture dataset, compared against existing baselines. Our approach consistently reconstructs

bio-mechanically accurate 3D sign language avatars from monocular videos. DexAvatar preserves fine-grained hand articulations and finger orientations that are often missing or distorted in baseline reconstructions. For the sign *Sonne*, methods such as PIXIE [17], OSX [36], and EVA* generate avatars that deviate noticeably from ground truth, whereas DexAvatar maintains close alignment. The sign *BesuchenEinmischen* highlights a complex sequence involving simultaneous hand-body interactions. Competing methods including PIXIE [17], PyMAF-X [66], SGNify [20], OSX [36], and EVA* [23] frequently produce misaligned wrists or unnatural limb orientations, while DexAvatar preserves structural consistency guided by our learned hand and body priors. For the sign *Muell*, PyMAF-X and EVA* perform particularly poorly, and other methods struggle with wrist and finger orientations; in contrast, DexAvatar most closely approximates ground truth. Overall, these results demonstrate that DexAvatar produces stable and coherent reconstructions, effectively capturing both global body pose and subtle finger dynamics. This produces avatars that are anatomically plausible and well-aligned with the video. We also inspected the evaluation dataset and found some implausible hand poses in the ground truth (see section 6 of Supplementary Material for details). Please see section 7 of Supplementary Material for additional evaluations on motion blur, noisy, and self-occlusion cases in input frames.

6. Conclusion

We introduced DexAvatar, a method that generates 3D avatars of humans involved in sign language communication from monocular videos. Quantitative and qualitative results demonstrate consistent improvements over prior work, driven by our sign language-aware priors and an improved optimization pipeline that produces realistic reconstructions. These findings highlight the importance of domain-specific priors for sign language reconstruction. Future work will scale the training data and further strengthen the priors to cover a broader range of signers and signing styles.

Acknowledgements

This research was partly supported by the Australian Government through the Australian Research Council’s Discovery Early Career Researcher Award (project DE230100049). The views expressed herein are those of the authors and are not necessarily those of the Australian Government or Australian Research Council. We also acknowledge Monash University and National Computational Infrastructure for providing High Performance Computing infrastructure used in this research.

Figure 5. **Qualitative results.** We compare 3D holistic human mesh reconstruction methods on SGNify [20] evaluation dataset. DexAvatar produces significantly better SL reconstructions with plausible body and hand poses.



Supplementary Material

S1. Background Knowledge

SMPL-X [44] is an advanced parametric human body model, an extension of the original SMPL [38], integrating hand articulations through MANO [47] and facial expressions through FLAME [34]. This enables comprehensive full-body representations that include hand and face dynamics. SMPL-X is defined by a mapping function $M(\boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\psi}) : \mathbb{R}^{|\boldsymbol{\theta}| \times |\boldsymbol{\beta}| \times |\boldsymbol{\psi}|} \rightarrow \mathbb{R}^{3N}$, parameterized by the pose $\boldsymbol{\theta} \in \mathbb{R}^{3(K+1)}$, where K is the number of body joints in addition to a joint for global rotation. $\boldsymbol{\beta}$ represents shape coefficients, and $\boldsymbol{\psi}$ are the facial expression coefficients.

The model uses vertex-based linear blend skinning with $N = 10,475$ vertices and $K = 54$ joints, including joints for hands, neck, jaw, and eyeballs. The formulation of SMPL-X is defined as follows:

$$M(\boldsymbol{\beta}, \boldsymbol{\theta}, \boldsymbol{\psi}) = W(T_P(\boldsymbol{\beta}, \boldsymbol{\theta}, \boldsymbol{\psi}), J(\boldsymbol{\beta}), \boldsymbol{\theta}, \mathcal{W}), \quad (16)$$

where,

$$T_P(\boldsymbol{\beta}, \boldsymbol{\theta}, \boldsymbol{\psi}) = \bar{T} + B_S(\boldsymbol{\beta}; \mathcal{S}) + B_E(\boldsymbol{\psi}; \mathcal{E}) + B_P(\boldsymbol{\theta}; \mathcal{P}). \quad (17)$$

$B_P(\cdot)$, $B_S(\cdot)$, and $B_E(\cdot)$ in Eq. (17) represent the pose, shape, and expression-dependent corrective blend functions. \mathcal{S} , \mathcal{E} , and \mathcal{P} represent the orthonormal principal components of vertex displacements of shape, pose, and expression blend shape variations. These functions apply vertex displacements to the canonical template mesh \bar{T} based on the pose parameters $\boldsymbol{\theta}$, shape parameters $\boldsymbol{\beta}$, and expression parameters $\boldsymbol{\psi}$. In particular, $B_P(\boldsymbol{\theta})$ and $B_S(\boldsymbol{\beta})$ capture non-linear deformations specific to pose and shape variations. After these corrections, the deformed mesh is processed using linear blend skinning, denoted as \mathcal{W} , which rotates the vertices around the joints $J(\boldsymbol{\beta})$ according to the skeletal kinematics. The final mesh is smoothed using a pre-defined set of blend weights W , resulting in the articulated 3D human body mesh.

S2. Range of Motion and Signer Space

This section complements section 3.2.1 of the main paper. We constrain the upper limb using the physiological degrees of freedom (DOFs) [21, 31] of the joints most active in sign language. The shoulder has three DOFs, *i.e.*, flexion/extension, abduction/adduction, and internal/external rotation. The elbow-forearm complex has two DOFs, *i.e.*, humeroulnar flexion/extension and radioulnar pronation/supination. For the wrist, we adopt a three DOF formulation covering forearm pronation/supination, wrist flexion/extension, and radial/ulnar deviation.

Clinical range of motion (ROM) [21, 31] for these DOFs is typically reported as *unsigned* magnitudes in anatomical planes (*e.g.*, wrist flexion 90° and extension 90°). To use these values, we convert them to *signed* bounds in an anatomical Euler [59] convention. For each DOF, we align a local axis with the corresponding motion, adopt the right-hand rule for signs, and verify the orientation on the rig.

We express ROMs as a single signed interval in the aligned Euler convention. For bilateral joints, we mirror the sign across the sagittal plane (including wrist flexion/extension), matching the left/right labeling in Fig. S1. This deterministic normalization yields SMPL-X [44] compatible signed bounds from clinical ROM with a brief visual sanity check.

Having normalized clinical ROM to SMPL-X compatible signed Euler bounds, we further constrain the shoulder using the notion of signer space [8, 60], a torso-centric 3D region where signs are typically produced (see Fig. S3). We model this as a compact volume anchored to the torso, bounded laterally near the shoulders, vertically from the lower chest to the forehead, and in depth slightly in front of the chest.

We restrict shoulder motion to be consistent with a torso-anchored, ground-parallel signer space envelope (see Fig. S3). Accordingly, we cap horizontal abduction by disallowing motion behind the torso (see Fig. S2a) while permitting substantial horizontal adduction to support cross-body movements (see Fig. S2b). This yields a simple deterministic rule that filters out poses with shoulder angles outside these bounds.

S3. Motion Capture Data Acquisition

This section complements section 3.2.2 of the main paper. We export MANUS [1] motion from Unity as FBX and process it in Blender. The FBX hierarchy in Fig. S4 comprises a top-level Hands node that controls global translation, child nodes `ManusHand_L` and `ManusHand_R`, per-hand `SK_Hand` meshes, and a terminal `root` that contains the armature.

Rotations in mocap reside on the finger and thumb bones of the armature, while the parent nodes position the gloves about the wrist pivot. After testing several approaches, the most reliable process was to set the armature to a flat “rest” pose, then delete keyframes on the parent containers so that only the finger animations remained. The armature was then aligned as closely as possible with the SMPL-X finger bones as shown in Fig. S5. Differences in finger length did not affect retargeting substantially, but misaligned knuckles caused distortion and stretching in the SMPL-X rig, so careful attention was paid to joint spacing. Some discrepancies remained due to restrictions on altering the SMPL-X

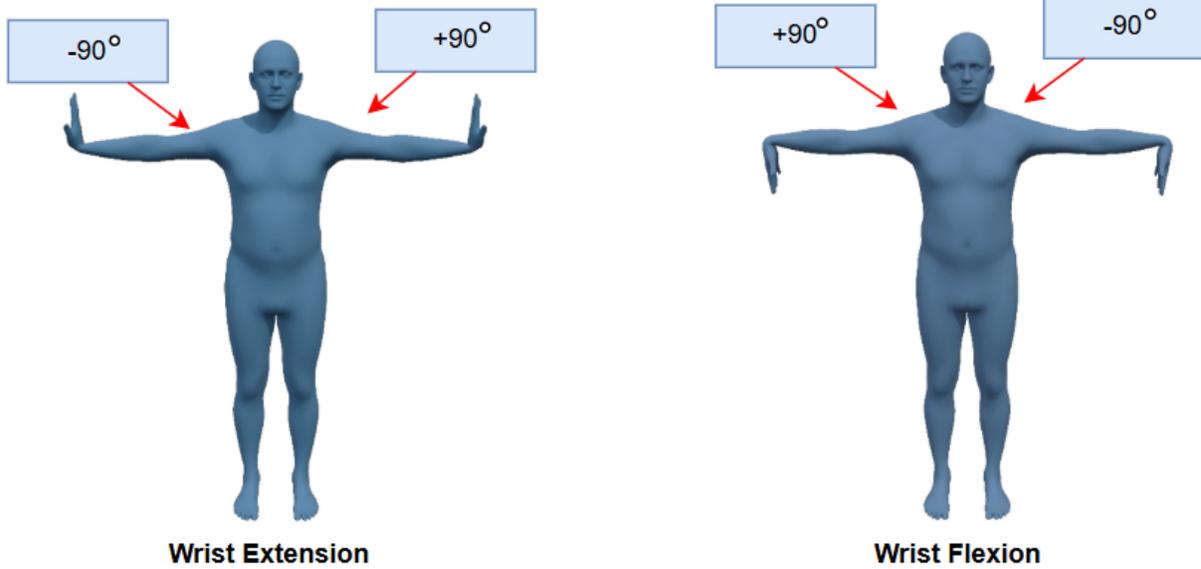


Figure S1. We show extreme poses at $\pm 90^\circ$ for wrist extension/flexion sign convention with left and right sign mirroring consistent with SMPL-X euler angle setup.

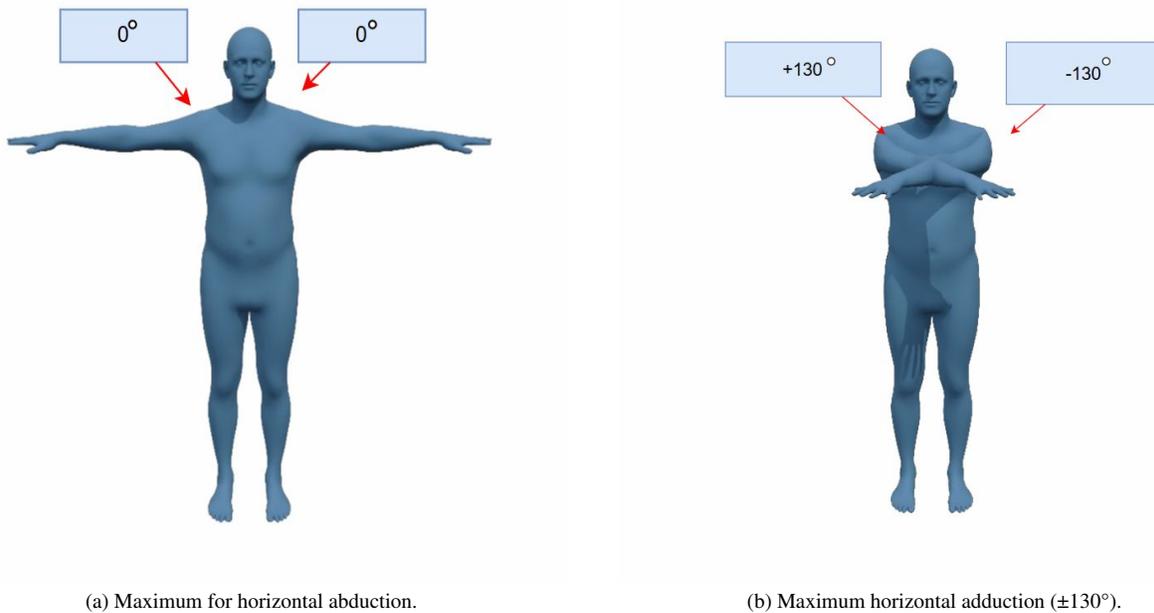


Figure S2. Shoulder horizontal ad/abduction within signer space in SMPL-X compatible euler angle setup. Shoulder motion is constrained within a torso-anchored, ground-parallel signer space, disallowing horizontal abduction behind the torso while permitting substantial horizontal adduction for cross-body movements.

T-pose.

Since the hands were on separate armatures, they were retargeted one at a time as shown in Fig. S6. We dupli-

cated the SMPL-X rig, retargeted the right hand first, then retargeted the left hand on the duplicate. The left-hand keyframes from the duplicate were copied back to the origi-

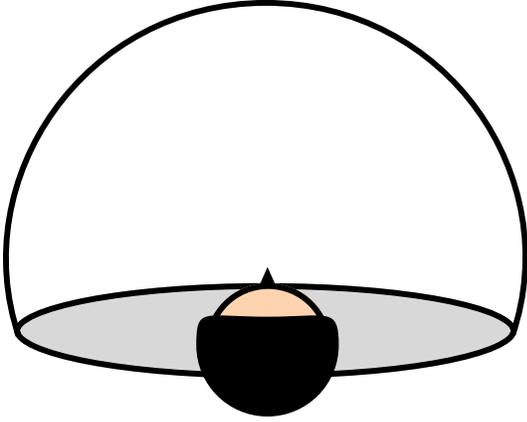


Figure S3. Bird's eye view of signer space envelope showing a torso-anchored 3D workspace. The figure has been adapted from [8].

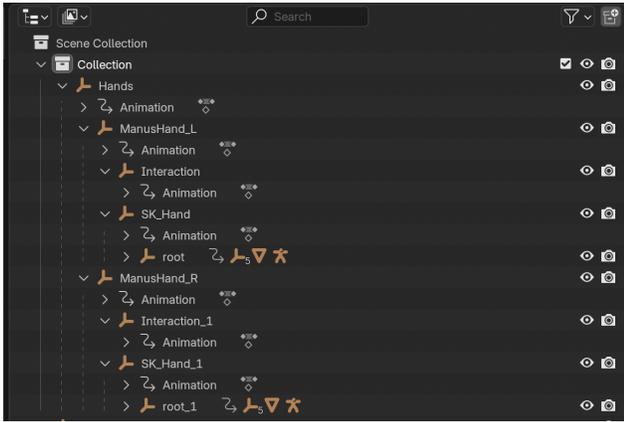


Figure S4. FBX hierarchy from the MANUS export. Hands controls global placement. ManusHand_L and ManusHand_R parent the SK_Hand meshes and a per-hand root armature with finger and thumb mocap. Animation and Interaction store non-finger transforms that we remove for retargeting.

nal SMPL-X rig, consolidating both hands' motion. At this stage, the finger animations were successfully transferred to SMPL-X. A fresh copy of the .fbx file was imported, and only the parent "Hands" keyframes were removed. This preserved the movement of the hands in space, allowing them to be positioned naturally in front of the SMPL-X body.

To enable the arms to move with the MANUS gloves, an inverse kinematics (IK) setup was added. Each arm was given an IK handle constrained to a duplicate wrist bone, which was then linked to ManusHand_L and ManusHand_R. This enabled the arms to follow the hand movements using a single bone, rather than requiring manual adjustment of the forearm and upper arm. Wrist rotations could not be transferred due to incompatibilities be-

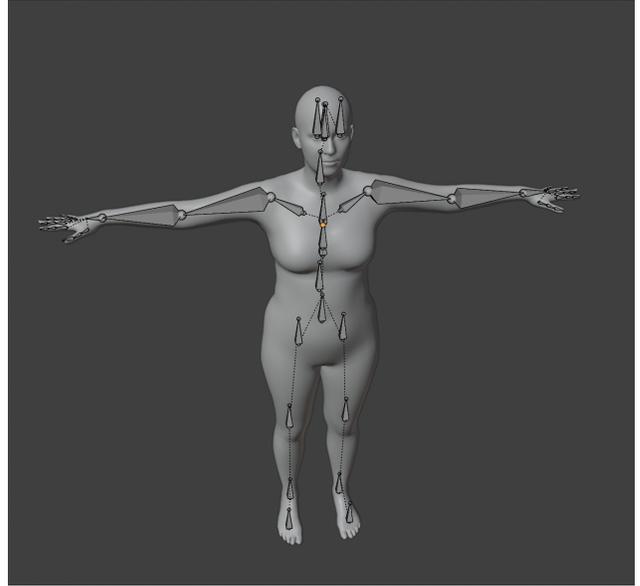


Figure S5. SMPL-X in the rest pose with visible armature.

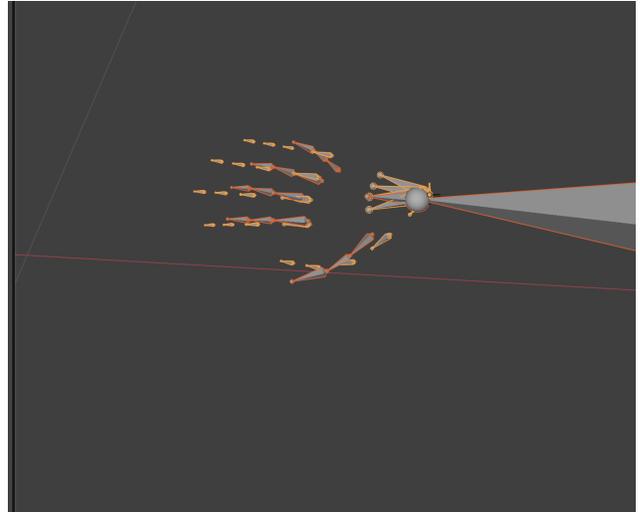


Figure S6. Separate hand armatures during retargeting. We retarget each hand one at a time. We duplicate the SMPL-X rig, retarget the right hand first, then retarget the left hand on the duplicate. We copy the left-hand keyframes from the duplicate back to the original SMPL-X rig, consolidating both hands' motion. Finger animations are then transferred to SMPL-X.

tween the SMPL-X T-pose and the bone roll of the MANUS rig, which meant that constraints could not replicate these rotations accurately. Finally, we baked the animations into the SMPL-X armature, replacing the constraints with explicit per-frame rotation keyframes. The temporary constraints and auxiliary IK bones were removed, leaving a clean animation that could be extracted and used.

S4. Analysis of Prior Parameters

This section complements section 5.1 of the main paper. We study how data filtering and lightweight bio-mechanical constraints affect body and hand pose estimation. Tables S1 and S2 report hyperparameter sweeps for SignBPoser and SignHPoser under matched architectures. Each setting varies only the training data correction and the presence of a bio-mechanical loss, while we select the best hyperparameter on Evaluation (DEV) and TEST data. We evaluate with MPJPE and MPVPE on both splits, and we summarize the main trends below.

Table S1 reports hyperparameter tuning for three settings of SignBPoser. BP_u uses the unfiltered data, BP_f uses the bio-mechanically filtered data, and BP_{f+bio} adds a body bio-mechanical loss on top of the filtered data. We select the best hyperparameter for each setting on the DEV and TEST sets.

In the unfiltered setting BP_u latent 31 performs worst. Increasing to 32 reduces error on DEV by 21% on MPJPE and 17% on MPVPE, and on TEST by about 18% and 16%. A further increase to 33 brings additional reductions of about 2% and 11% on DEV, and about 5% and 14% on TEST. Latent 33 is therefore the most reliable choice in this setting.

On the other hand, in the bio-mechanical filtered setting BP_f , latent 32 has the highest error. Switching to 31 reduces error on DEV by about 1% MPJPE and 6% MPVPE, and on TEST by about 3% and 2%. Increasing to 33 brings smaller additional gains of roughly 2% on both metrics on DEV and about 2% on TEST. The effect of latent size is therefore mild in this setting.

Finally for filtered-plus-constraint setting BP_{f+bio} , the extremes 0.5 and 2.5 give slightly higher errors. Setting the weight to 1.5 reduces MPJPE by about 1–3% and MPVPE by about 1% on both DEV and TEST. The configurations are close to each other, which indicates stable behavior with the constraint.

In Table S2 we compare three settings for SignHPoser. HP_u is trained on the uncorrected data, HP_f is trained on a bio-mechanically corrected data, and HP_{f+bio} keeps the corrected poses and adds a lightweight hand bio-mechanical loss during training. The architecture matches SignBPoser, and the only differences are data correction and the presence of the hand constraint.

In the HP_u setting, moving from latent 22 to 23 lowers error by about 5% on MPJPE and about 5% on MPVPE on TEST, with similar gains on DEV. Increasing to 24 gives an additional reduction of roughly 2% on MPJPE across splits, and about 2% on MPVPE on TEST and about 5% on DEV. Latent 24 is therefore the most reliable within this setting.

For HP_f , both latent 24 and 22 are worse than 23. From 24 to 23 the error drops by about 8% on MPJPE and about 8% on MPVPE on both DEV and TEST. From 22 to 23, the drop is smaller on MPJPE at about 3%, but larger on

MPVPE at about 12–14% across splits. Latent 23 is the preferred choice.

Finally for HP_{f+bio} , weight 2.5 performs worst. Reducing it to 1.5 lowers MPJPE by about 9–10% and MPVPE by about 16–18% on both DEV and TEST. Compared with 0.5, the 1.5 setting also improves by about 2–3% on MPJPE and about 8% on MPVPE. The hand constraint is most effective at 1.5.

From the above results, we can conclude that training for SignBPoser and SignHPoser remains stable under a similar architecture. Simple choices like latent size and lightweight bio-mechanical constraints guide accuracy without instability across the DEV and TEST sets.

S5. Ablation of SignHPoser with Vposer

This section complements section 5.1 of the main paper. We evaluate the hand prior SignHPoser using VPoser. Table S3 summarizes ablation results. The first two rows of Table S3 show the results for HP_u and HP_f variants. It can be observed that HP_f outperforms HP_u on all metrics, with relative error reductions of **1.2%** on Upper Body, **1.3%** on Left Hand, and **3.2%** on Right Hand. This demonstrates the importance of the correction process using bio-mechanical constraints on the hand data. Adding a bio-mechanical regularizer to the filtered prior in HP_{f+bio} increases accuracy compared to HP_f , yielding slight improvements on Upper Body (0.05%), Left Hand (0.15%), and Right Hand (1.7%). This indicates that introducing bio-mechanical constraints provides useful physical regularization in the fitting process.

S6. Limitations of the SGNify Ground Truth

This section complements section 5.2 of the main paper. We evaluate using TR-V2V against the SGNify ground truth, which contains occasional implausible hand configurations. DexAvatar shows consistent improvements for both hands and the upper body. The margin remains modest because TR-V2V penalizes distance to the ground truth mesh. When the reference encodes anatomically inconsistent finger postures or knuckle spacing, moving toward a plausible pose does not always reduce vertex distance. The ground truth often contains collapsed fingers and irregular knuckle spacing, which explains why plausibility corrections may not yield large vertex reductions. Qualitative comparisons in Fig. S7 show cleaner finger alignment. For certain signs like BESUCHENEINMISCHEN and FRECH, DexAvatar generates more plausible hand poses compared to the ground truth. These effects arise as SignHPoser optimizes toward anatomically plausible hand poses learned from our mocap data.

Table S1. **Hyper-parameter tuning of SignBPoser.** We denote SignBPoser trained on BP_u: unfiltered body data, BP_f: filtered body data, BP_{f+bio}: filtered body data with body bio-mechanical loss. We evaluate using Mean Per Joint Position Error (MPJPE) and Mean Per Vertex Position Error (MPVPE), on the recovered joints and meshes.

Variant	Parameters				DEV		TEST	
	KL	Neuron	Latent	Biomech constant	MPJPE↓	MPVPE↓	MPJPE↓	MPVPE↓
BP _u	0.001	512	33	X	5.87	3.73	5.69	3.62
	0.001	512	32	X	5.99	4.17	5.98	4.21
	0.001	512	31	X	7.56	5.05	7.28	5.00
BP _f	0.001	512	33	X	7.21	4.33	7.04	4.14
	0.001	512	32	X	7.45	4.68	7.43	4.32
	0.001	512	31	X	7.37	4.41	7.17	4.24
BP _{f+bio}	0.001	512	33	0.5	7.42	4.43	7.21	4.32
	0.001	512	33	1.5	7.30	4.39	7.10	4.25
	0.001	512	33	2.5	7.37	4.42	7.29	4.29

Table S2. **Hyper-parameter tuning of SignHPoser.** We denote SignHPoser trained on HP_u: unfiltered hand data, HP_f: filtered hand data, HP_{f+bio}: filtered hand data with hand bio-mechanical loss.

Variant	Parameters				DEV		TEST	
	KL	Neuron	Latent	Biomech constant	MPJPE↓	MPVPE↓	MPJPE↓	MPVPE↓
HP _u	0.0001	512	24	X	0.56	0.55	0.56	0.54
	0.0001	512	23	X	0.57	0.58	0.57	0.55
	0.0001	512	22	X	0.59	0.58	0.60	0.58
HP _f	0.0001	512	24	X	0.40	0.38	0.40	0.38
	0.0001	512	23	X	0.37	0.35	0.37	0.35
	0.0001	512	22	X	0.38	0.40	0.38	0.40
HP _{f+bio}	0.0001	512	23	0.5	0.40	0.41	0.40	0.41
	0.0001	512	23	1.5	0.39	0.38	0.39	0.38
	0.0001	512	23	2.5	0.43	0.45	0.43	0.45

Table S3. Ablation study of SignHPoser with Vposer within DexAvatar.

Method	UBody (-F)↓	LHand↓	RHand↓
HP _u	37.25	13.56	14.53
HP _f	36.79	13.39	14.06
HP _{f+bio}	36.77	13.37	13.82

S7. Additional qualitative evaluations with SOTA method

This section complements section 5.2 of the main paper. In Fig. S8, S9, and S10 we present qualitative results of DexAvatar on the MM-WLAuslan [53] SL dataset, compared against SGNify [20] and EVA* under challenging scenarios such as motion blur, self-occlusion, and noise.

S7.1. Qualitative evaluation on motion blur images

Fig. S8 presents three examples under motion blur. In Example 1, EVA* shows overspread fingers with a un-

natural bending and uneven gaps, while SGNify is an incorrect wrist configuration although bio-mechanically correct, whereas DexAvatar maintains a compact rounded configuration with evenly spaced fingertips that matches the blurred target contact. Building on this, in Example 2, EVA* produces distorted fingers with incorrect spacing and no bio-mechanical stability, SGNify is again an incorrect detection although bio-mechanically correct, while DexAvatar preserves clear fingers with realistic curl and fingertip positions that support the intended overlap with a clean stable contact. Continuing this pattern, in Example 3, EVA* compresses the fingers into a tight bundle so separation is lost and local interpenetration appears along the contact between the index finger and thumb, and also shows a body misalignment relative to the image, SGNify exhibits penetration of fingers with palm that are inconsistent with the image, whereas DexAvatar forms a coherent rounded cluster with visible fingertip order and contact that follows the blurred evidence without fusion and is correct. Overall, DexAvatar maintains accurate and plausible hand contacts under blur.

S7.2. Qualitative evaluation on self-occluded images

Fig. S9 presents three examples under self-occlusion. In Example 1, EVA* overspreads the fingers with interpenetration of right hand with the left and inconsistent fingertip spacing, SGNify is bio-mechanically reasonable but misplaces the hand contact, whereas DexAvatar maintains a compact closed configuration with fingertips aligned to the intended contact. Similarly, in Example 2, EVA* keeps both index fingers unnaturally extended while the others curl, SGNify reconstruction consists of overlapping of hands with a wrong wrist orientation for the left hand, while DexAvatar preserves clear fingers with realistic curl and a clean stable contact. Finally, in Example 3, EVA* shows right hand finger postures beyond plausibility, SGNify is plausible but infers the overlap order incorrectly since the right hand should occlude the left rather than the reverse, whereas DexAvatar provides the correct estimation with plausible hands throughout. Overall, DexAvatar maintains accurate and plausible hand contacts under occlusion.

S7.3. Qualitative evaluation on images with gaussian noise

We add Gaussian noise to the input frames and compare EVA*, SGNify [20], and DexAvatar in Fig. S10. In Example 1, EVA* reconstructs excessive spacing between the pinky and ring fingers which is bio-mechanically implausible, while SGNify yields a bio-mechanically reasonable hand that nevertheless does not match the target configuration in the image. In contrast, DexAvatar recovers the intended finger arrangement and maintains bio-mechanical constraints across both hands. Moving to Example 2, EVA* degrades to an implausible finger arrangement and SGNify fails to produce any mesh due to missing keypoints under noise, whereas DexAvatar still reconstructs a complete and accurate pose with plausible finger angles and stable contact. Finally, in Example 3, EVA* shows a left-thumb posture beyond plausibility together with distorted right-hand fingers and SGNify again produces no mesh because no keypoints are detected, while DexAvatar returns an anatomically plausible reconstruction with unbroken fingers and consistent bilateral alignment. Overall, DexAvatar remains stable under noisy frames and preserves both accuracy and bio-mechanical plausibility.

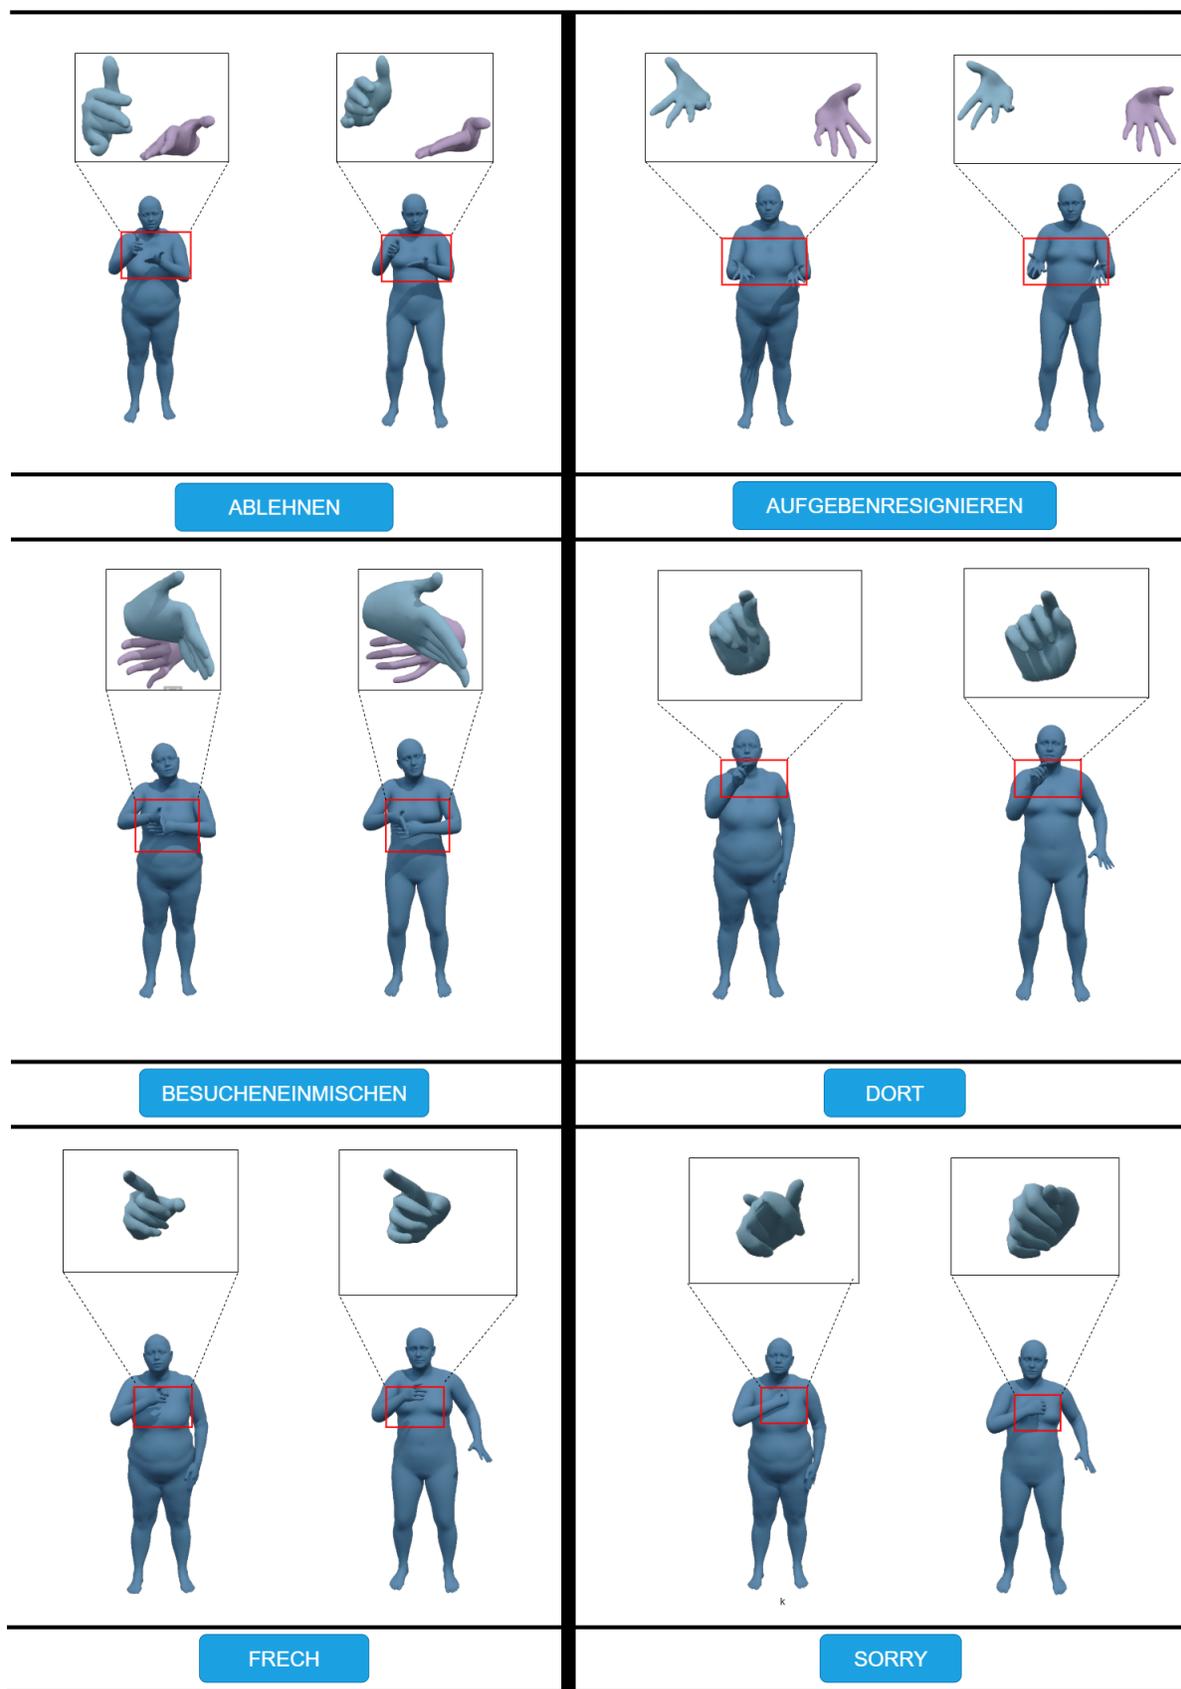


Figure S7. **Examples of independent signs.** Each panel shows two signers, the left signer is the ground truth from the SGNiFy evaluation set and the right signer is DexAvatar generated by our fitting optimization. The SGNiFy [20] ground truth often contains low-quality hand shapes and placements, while DexAvatar produces more plausible poses. The improvement comes from a SignHPrior trained on our mocap dataset.

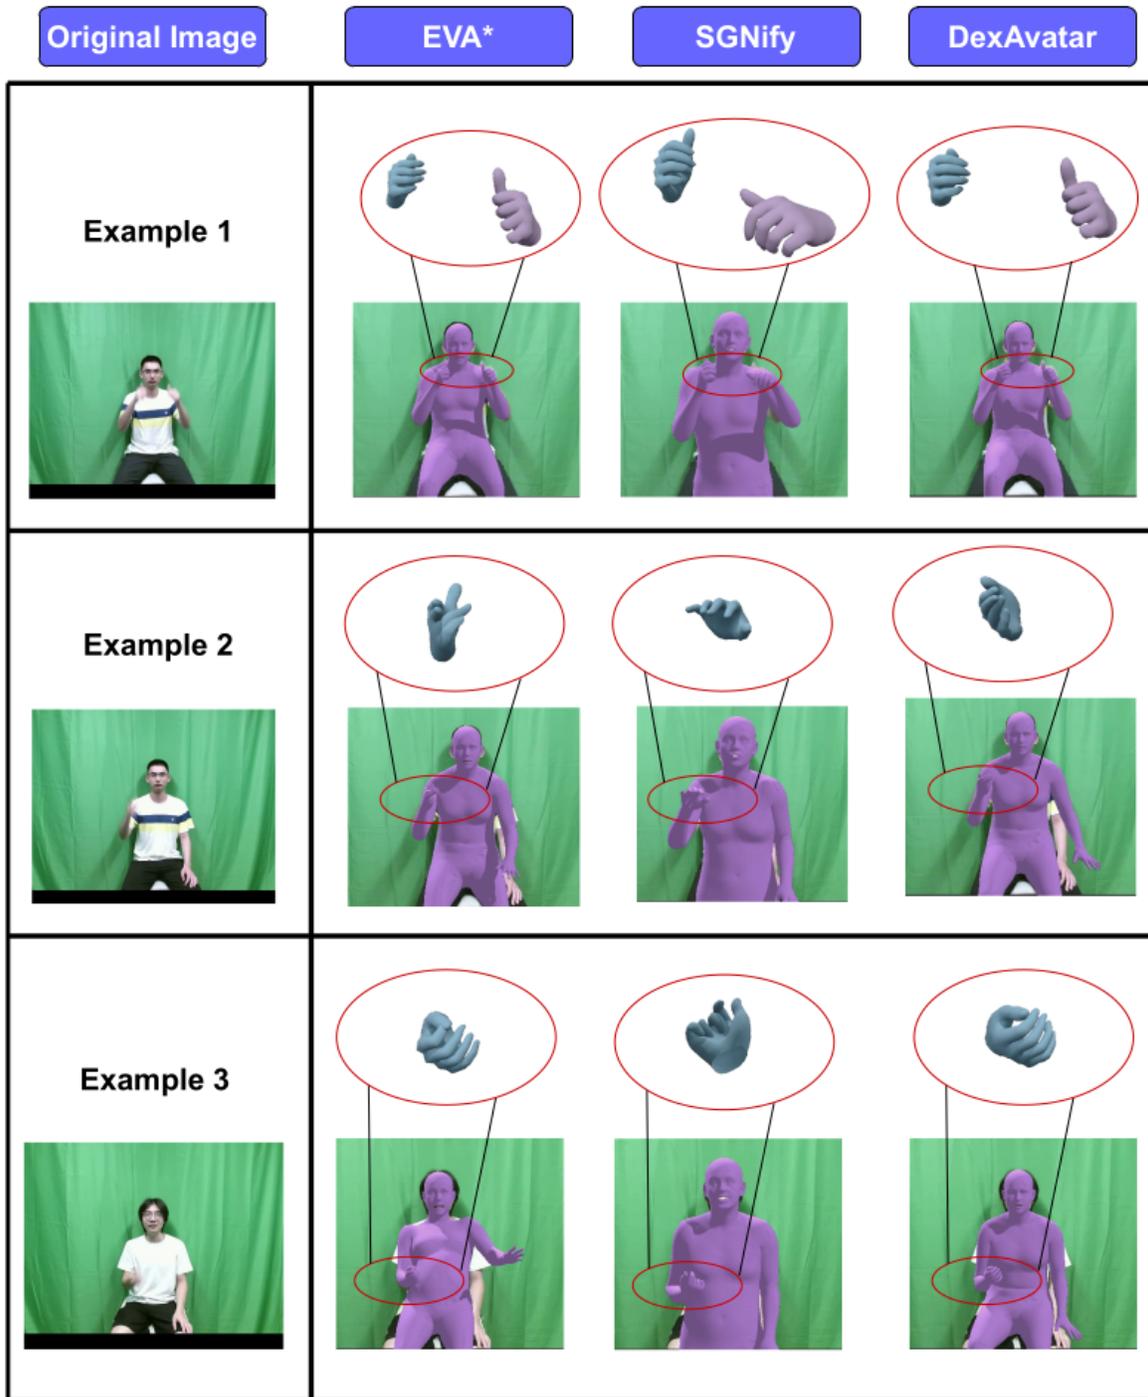


Figure S8. **Qualitative evaluation under motion blur.** We compare EVA*, SGNify [20], and DexAvatar. DexAvatar preserves compact rounded finger configurations and clean contact, while EVA* overspreads or distorts the fingers and SGNify yields incorrect detections or misrepresented contact, with additional body misalignment appearing in harder cases.

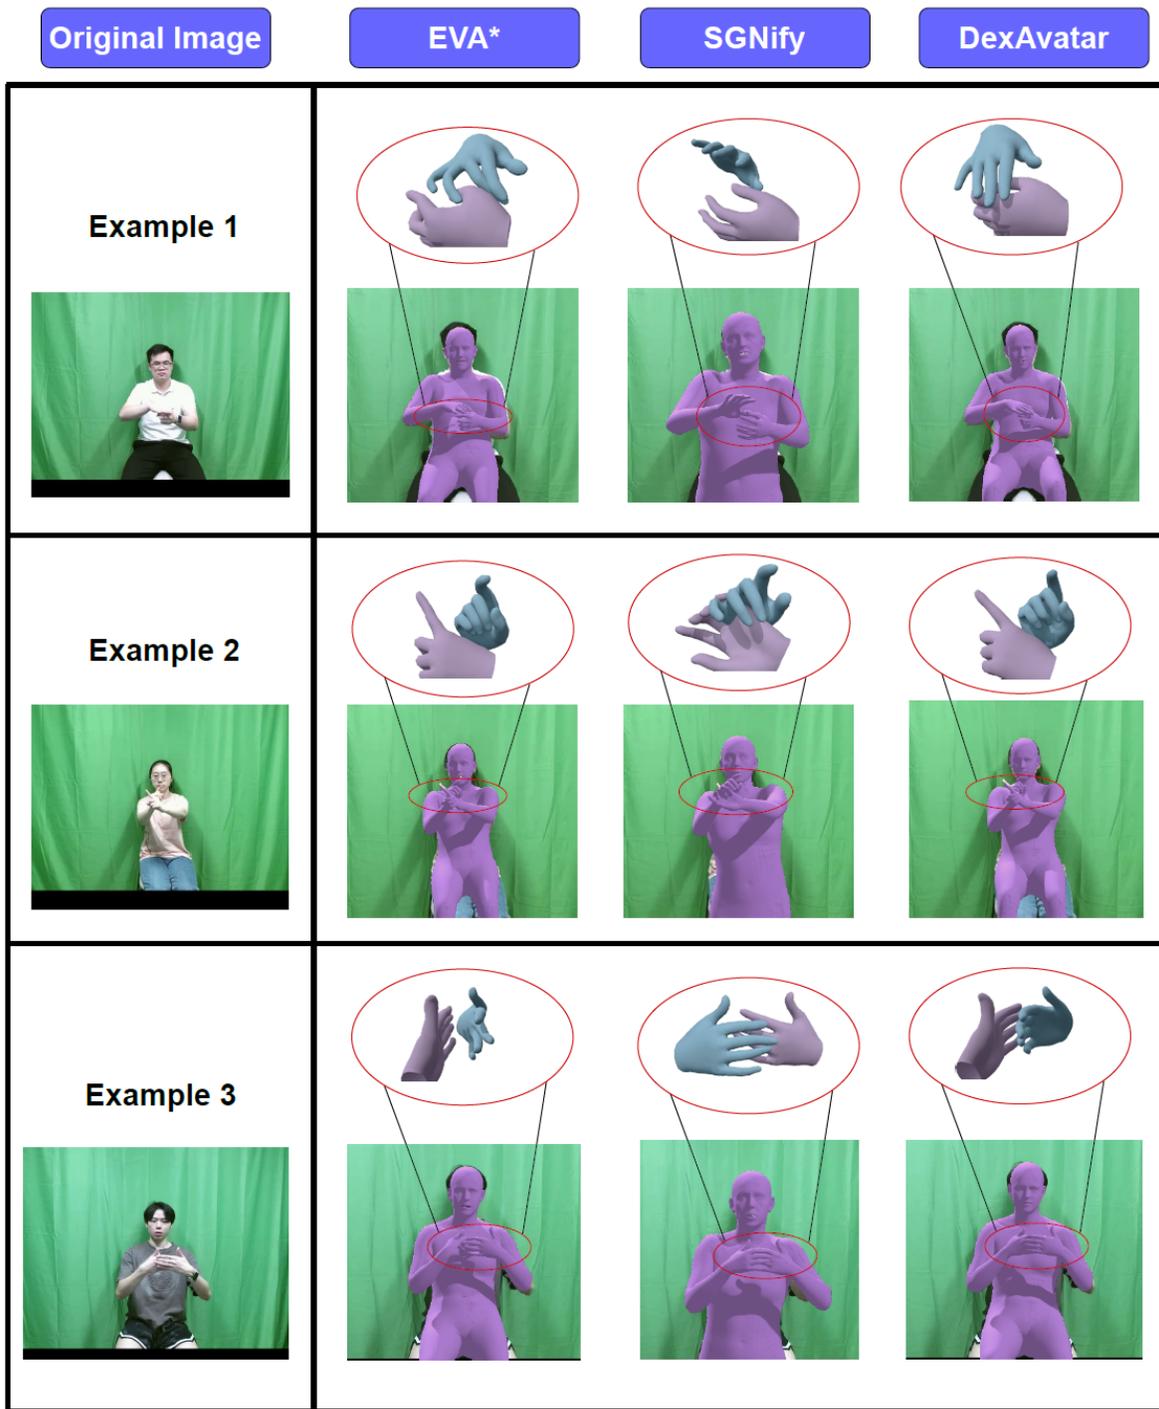


Figure S9. **Qualitative evaluation under self-occlusion.** We compare EVA*, SGNify [20], and DexAvatar. DexAvatar maintains compact finger configurations and correct overlap, while EVA* overspreads or distorts the fingers and SGNify misplaces contact or infers the overlap order incorrectly.

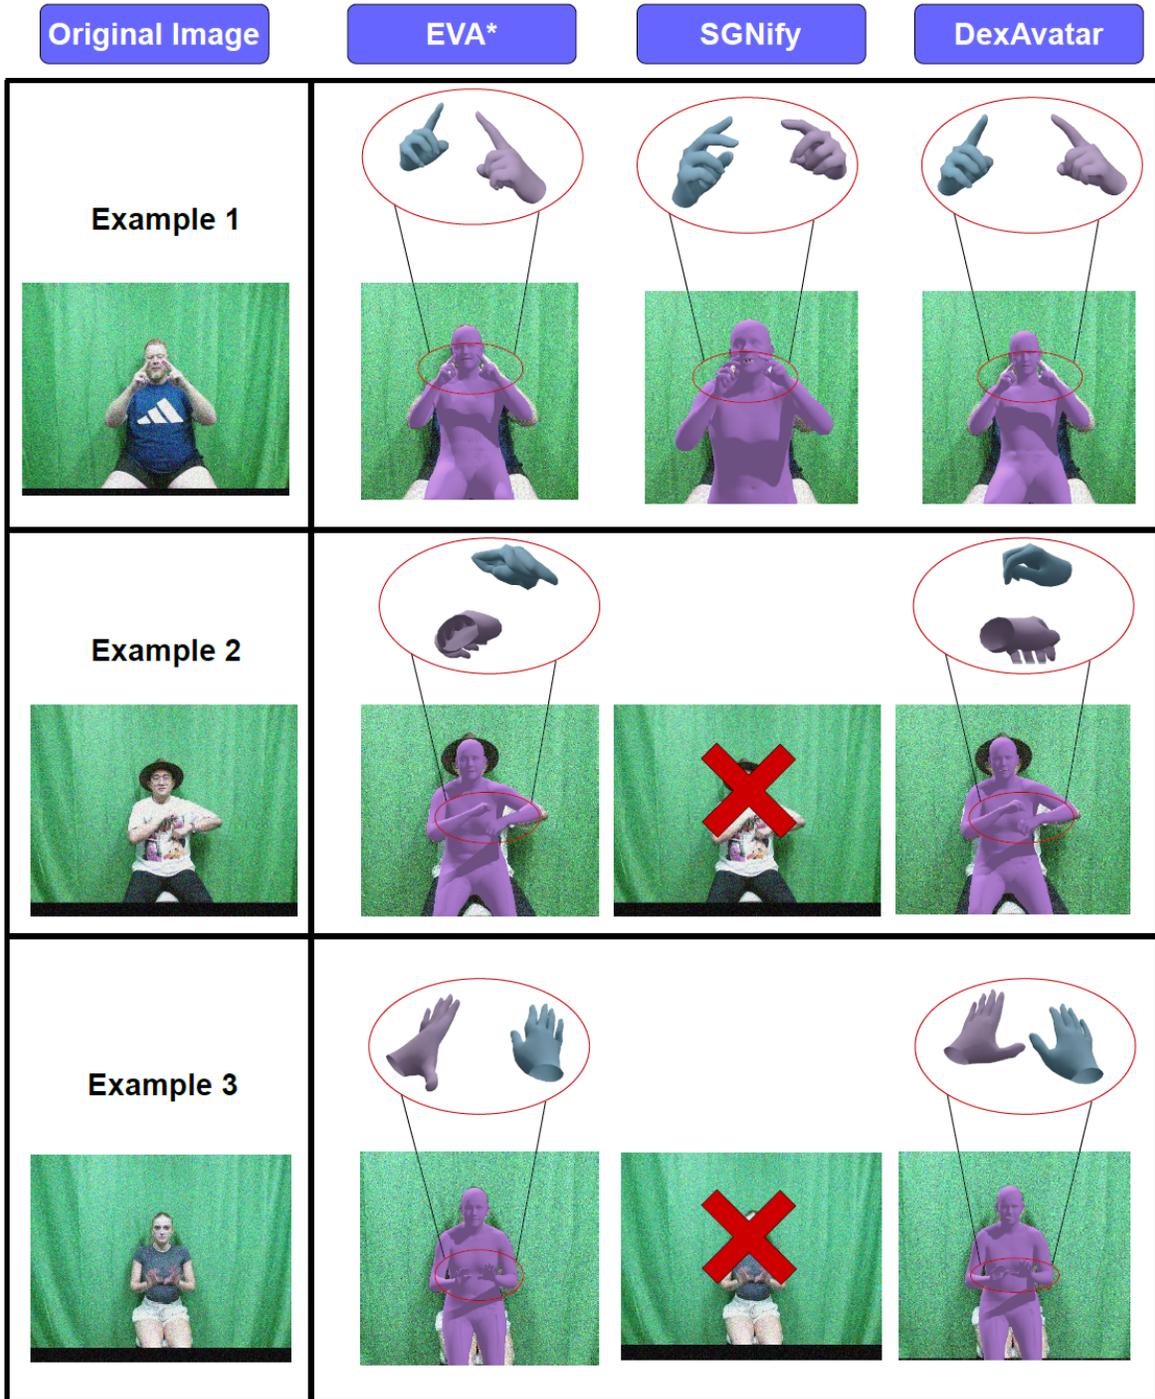


Figure S10. **Qualitative evaluation under gaussian noise.** We add Gaussian noise to input frames and compare EVA*, SGNify [20], and DexAvatar. DexAvatar preserves plausible finger shape and clean contact, while EVA* exhibits implausible spacing and distortions, and SGNify fails to produce a mesh in harder cases due to missing keypoints.

References

- [1] Metagloves pro, 2025. 4, 1
- [2] Devices, 2025. 4
- [3] Ijaz Akhter and Michael J Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1446–1455, 2015. 4
- [4] Rotem Shalev Arkushin, Amit Moryossef, and Ohad Fried. Ham2pose: Animating sign language notation into pose sequences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21046–21056, 2023. 1
- [5] DF Armstrong. *Gesture and the Nature of Language*. Cambridge University Press, 1995. 1
- [6] Vasileios Baltatzis, Rolandos Alexandros Potamias, Evangelos Ververas, Guanxiong Sun, Jiankang Deng, and Stefanos Zafeiriou. Neural sign actors: A diffusion model for 3d sign language production from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1985–1995, 2024. 2, 6
- [7] Sara Bilal, Rini Akmeliawati, Momoh Jimoh El Salami, and Amir A Shafie. Vision-based hand posture detection and recognition for sign language—a study. In *2011 4th International Conference on Mechatronics (ICOM)*, pages 1–6. IEEE, 2011. 2
- [8] Chiara Branchini, Lara Mantovan, et al. *A Grammar of Italian Sign Language (LIS)*. Edizioni Ca’Foscari, 2020. 3, 4, 6, 1
- [9] Zhongang Cai, Wanqi Yin, Ailing Zeng, Chen Wei, Qingping Sun, Wang Yanjun, Hui En Pang, Haiyi Mei, Mingyuan Zhang, Lei Zhang, et al. Smpler-x: Scaling up expressive human pose and shape estimation. *Advances in Neural Information Processing Systems*, 36:11454–11468, 2023. 2, 3, 4, 5
- [10] Fai Chen Chen, Silvia Appendino, Alessandro Battezzato, Alain Favetto, Mehdi Mousavi, and Francesco Pescarmona. Constraint study for a hand exoskeleton: human hand kinematics and dynamics. *Journal of Robotics*, 2013(1):910961, 2013. 4, 6
- [11] CMU Graphics Lab. Carnegie mellon university graphics lab motion capture database, 2003. Accessed: YYYY-MM-DD. 4
- [12] Haoye Dong, Aviral Chharia, Wenbo Gou, Francisco Vicente Carrasco, and Fernando D De la Torre. Hamba: Single-view 3d hand reconstruction with graph-guided bi-scanning mamba. *Advances in Neural Information Processing Systems*, 37:2127–2160, 2024. 2
- [13] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [14] Amanda Duarte, Shruti Palaskar, Lucas Ventura, Deepti Ghadiyaram, Kenneth DeHaan, Florian Metze, Jordi Torres, and Xavier Giro-i Nieto. How2sign: a large-scale multi-modal dataset for continuous american sign language. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2735–2744, 2021. 2, 3, 5
- [15] Sergio Escalera, Xavier Baró, Jordi Gonzalez, Miguel A Bautista, Meysam Madadi, Miguel Reyes, Víctor Ponce-López, Hugo J Escalante, Jamie Shotton, and Isabelle Guyon. Chalearn looking at people challenge 2014: Dataset and results. In *Computer Vision-ECCV 2014 Workshops: Zurich, Switzerland, September 6-7 and 12, 2014, Proceedings, Part I 13*, pages 459–473. Springer, 2015. 1
- [16] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J Black, and Otmar Hilliges. Arctic: A dataset for dexterous bimanual hand-object manipulation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12943–12954, 2023. 2
- [17] Yao Feng, Vasileios Choutas, Timo Bolkart, Dimitrios Tzionas, and Michael J Black. Collaborative regression of expressive bodies using moderation. In *2021 International Conference on 3D Vision (3DV)*, pages 792–804. IEEE, 2021. 6, 7
- [18] JA Fessler. Statistical methods for tomographic image reconstruction. *University of Michigan, GE CRD*, 2000. 3
- [19] Lin Geng Foo, Tianjiao Li, Hossein Rahmani, Qihong Ke, and Jun Liu. Unified pose sequence modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13019–13030, 2023. 2
- [20] Maria-Paola Forte, Peter Kulits, Chun-Hao P Huang, Vasileios Choutas, Dimitrios Tzionas, Katherine J Kuchenbecker, and Michael J Black. Reconstructing signing avatars from video using linguistic priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12791–12801, 2023. 2, 5, 6, 7, 8, 9, 10
- [21] Joseph Hamill and Kathleen M Knutzen. *Biomechanical basis of human movement*. Lippincott Williams & Wilkins, 2006. 3, 1
- [22] Yanbin Hao, Zi-Niu Liu, Hao Zhang, Bin Zhu, Jingjing Chen, Yu-Gang Jiang, and Chong-Wah Ngo. Person-level action recognition in complex events via tsd-tsm networks. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4699–4702, 2020. 2
- [23] Hezhen Hu, Zhiwen Fan, Tianhao Wu, Yihan Xi, Seoyoung Lee, Georgios Pavlakos, Zhangyang Wang, et al. Expressive gaussian human avatars from monocular rgb video. *Advances in Neural Information Processing Systems*, 37:5646–5660, 2024. 3, 6, 7
- [24] Jie Huang, Wengang Zhou, Houqiang Li, and Weiping Li. Attention-based 3d-cnns for large-vocabulary sign language recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, 29(9):2822–2832, 2018. 1
- [25] Sagar Imambi, Kolla Bhanu Prakash, and GR Kanagachidambaresan. Pytorch. In *Programming with TensorFlow: solution for edge computing applications*, pages 87–104. Springer, 2021. 6
- [26] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. 4
- [27] Tomasz Kapuscinski, Mariusz Oszust, Marian Wysocki, and Dawid Warchol. Recognition of hand gestures observed by

- depth cameras. *International Journal of Advanced Robotic Systems*, 12(4):36, 2015. 1
- [28] Rawal Khirodkar, Timur Bagautdinov, Julieta Martinez, Su Zhaoen, Austin James, Peter Selednik, Stuart Anderson, and Shunsuke Saito. Sapiens: Foundation for human vision models. *arXiv preprint arXiv:2408.12569*, 2024. 2, 3, 4, 5
- [29] Diederik Kinga, Jimmy Ba Adam, et al. A method for stochastic optimization. In *International conference on learning representations (ICLR)*. California, 2015. 6
- [30] Diederik P Kingma, Max Welling, et al. Auto-encoding variational bayes, 2013. 4
- [31] Duane Knudson. *Fundamentals of biomechanics*. Springer, 2007. 3, 1
- [32] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11127–11137, 2021. 2
- [33] Dongxu Li, Cristian Rodriguez, Xin Yu, and Hongdong Li. Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1459–1469, 2020. 1
- [34] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017. 2, 1
- [35] Yiheng Li, Ruibing Hou, Hong Chang, Shiguang Shan, and Xilin Chen. Unipose: A unified multimodal framework for human pose comprehension, generation and editing. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 27805–27815, 2025. 2
- [36] Jing Lin, Ailing Zeng, Haoqian Wang, Lei Zhang, and Yu Li. One-stage 3d whole-body mesh recovery with component aware transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21159–21168, 2023. 2, 6, 7
- [37] Jiajie Liu, Mengyuan Liu, Hong Liu, and Wenhao Li. Tcformer: Learning temporal correlation with implicit pose proxy for 3d human pose estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5478–5486, 2025. 2
- [38] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. pages 851–866, 2023. 1
- [39] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Ubowaja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*, 2019. 2, 3
- [40] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. 2
- [41] website main. World Health Organization, 2019. 1
- [42] Gyeongsik Moon, Shoou-I Yu, He Wen, Takaaki Shiratori, and Kyoung Mu Lee. Interhand2. 6m: A dataset and baseline for 3d interacting hand pose estimation from a single rgb image. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 548–564. Springer, 2020. 2
- [43] Philipp Moritz, Robert Nishihara, and Michael Jordan. A linearly-convergent stochastic l-bfgs algorithm. In *Artificial intelligence and statistics*, pages 249–258. PMLR, 2016. 6
- [44] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10975–10985, 2019. 1, 2, 3, 4
- [45] Georgios Pavlakos, Dandan Shan, Ilija Radosavovic, Angjoo Kanazawa, David Fouhey, and Jitendra Malik. Reconstructing hands in 3d with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9826–9836, 2024. 2, 3, 4, 5, 6
- [46] Rolandos Alexandros Potamias, Jinglei Zhang, Jiankang Deng, and Stefanos Zafeiriou. Wilor: End-to-end 3d hand localization and reconstruction in-the-wild. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 12242–12254, 2025. 2
- [47] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610*, 2022. 2, 1
- [48] Franco Ronchetti, Facundo Manuel Quiroga, César Estrebow, Laura Lanzarini, and Alejandro Rosete. Lsa64: an argentinian sign language dataset. *arXiv preprint arXiv:2310.17429*, 2023. 1
- [49] Yu Rong, Takaaki Shiratori, and Hanbyul Joo. Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1749–1759, 2021. 6
- [50] Benjamin Sapp, Alexander Toshev, and Ben Taskar. Cascaded models for articulated pose estimation. In *Computer Vision—ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part II 11*, pages 406–420. Springer, 2010. 2
- [51] Ben Saunders, Necati Cihan Camgoz, and Richard Bowden. Progressive transformers for end-to-end sign language production. In *European Conference on Computer Vision*, pages 687–705. Springer, 2020. 1
- [52] Ben Saunders, Necati Cihan Camgoz, and Richard Bowden. Mixed signals: Sign language production via a mixture of motion primitives. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1919–1929, 2021. 1
- [53] Xin Shen, Heming Du, Hongwei Sheng, Shuyun Wang, Hui Chen, Huiqiang Chen, Zhuojie Wu, Xiaobiao Du, Jiaying Ying, Ruihan Lu, et al. Mm-wlauslan: Multi-view multi-modal word-level australian sign language recognition dataset. *Advances in Neural Information Processing Systems*, 37:69700–69715, 2024. 5
- [54] Ozge Mercanoglu Sincan and Hacer Yalim Keles. Autsl: A large scale multi-modal turkish sign language dataset and baseline methods. *IEEE access*, 8:181340–181355, 2020. 1

- [55] Qingping Sun, Yanjun Wang, Ailing Zeng, Wanqi Yin, Chen Wei, Wenjia Wang, Haiyi Mei, Chi-Sing Leung, Ziwei Liu, Lei Yang, et al. Aios: All-in-one-stage expressive human pose and shape estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1834–1843, 2024. 2
- [56] Zhenhua Tang, Zhaofan Qiu, Yanbin Hao, Richang Hong, and Ting Yao. 3d human pose estimation with spatio-temporal criss-cross attention. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4790–4799, 2023. 2
- [57] Fang Wang and Yi Li. Beyond physical connections: Tree models in human pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 596–603, 2013. 2
- [58] Yang Wang and Greg Mori. Multiple tree models for occlusion and spatial constraints in human pose estimation. In *Computer Vision—ECCV 2008: 10th European Conference on Computer Vision, Marseille, France, October 12–18, 2008, Proceedings, Part III 10*, pages 710–724. Springer, 2008. 2
- [59] Wikipedia contributors. Euler angles. *Wikipedia, The Free Encyclopedia*. Accessed 2025-09-20. 1
- [60] Sherman Wilcox and Rocío Martínez. The conceptualization of space: Places in signed language discourse. *Frontiers in Psychology*, 11:1406, 2020. 3, 4, 6, 1
- [61] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vitpose: Simple vision transformer baselines for human pose estimation. *Advances in Neural Information Processing Systems*, 35:38571–38584, 2022. 2, 3
- [62] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vitpose: Simple vision transformer baselines for human pose estimation. *Advances in neural information processing systems*, 35:38571–38584, 2022. 2
- [63] Zhendong Yang, Ailing Zeng, Chun Yuan, and Yu Li. Effective whole-body pose estimation with two-stages distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4210–4220, 2023. 3, 5
- [64] Zhengdi Yu, Shaoli Huang, Chen Fang, Toby P Breckon, and Jue Wang. Acr: Attention collaboration-based regressor for arbitrary two-hand reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12955–12964, 2023. 2
- [65] Zhengdi Yu, Shaoli Huang, Yongkang Cheng, and Tolga Birdal. Signavatars: A large-scale 3d sign language holistic motion dataset and benchmark. In *European Conference on Computer Vision*, pages 1–19. Springer, 2024. 2, 3
- [66] Hongwen Zhang, Yating Tian, Yuxiang Zhang, Mengcheng Li, Liang An, Zhenan Sun, and Yebin Liu. Pymaf-x: Towards well-aligned full-body model regression from monocular images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10):12287–12303, 2023. 6, 7
- [67] Jinglei Zhang, Jiankang Deng, Chao Ma, and Rolandos Alexandros Potamias. Hawor: World-space hand motion reconstruction from egocentric videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1805–1815, 2025. 2
- [68] Yufei Zhang, Jeffrey O Kephart, and Qiang Ji. Weakly-supervised 3d hand reconstruction with knowledge prior and uncertainty guidance. In *European Conference on Computer Vision*, pages 106–125. Springer, 2024. 4
- [69] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. Motionbert: A unified perspective on learning human motion representations. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15085–15099, 2023. 2