

Sign-IDD: Iconicity Disentangled Diffusion for Sign Language Production

Shengeng Tang, Jiayi He, Dan Guo, Yanyan Wei*, Feng Li, Richang Hong*

School of Computer Science and Information Engineering, Hefei University of Technology
 {tangsg, guodan, weiy, fengli, hongrc}@hfut.edu.cn, hejy4396@mail.hfut.edu.cn

Abstract

Sign Language Production (SLP) aims to generate semantically consistent sign videos from textual statements, where the conversion from textual glosses to sign poses (G2P) is a crucial step. Existing G2P methods typically treat sign poses as discrete three-dimensional coordinates and directly fit them, which overlooks the relative positional relationships among joints. To this end, we provide a new perspective, constraining joint associations and gesture details by modeling the limb bones to improve the accuracy and naturalness of the generated poses. In this work, we propose a pioneering iconicity disentangled diffusion framework, termed Sign-IDD, specifically designed for SLP. Sign-IDD incorporates a novel Iconicity Disentanglement (ID) module to bridge the gap between relative positions among joints. The ID module disentangles the conventional 3D joint representation into a 4D bone representation, comprising the 3D spatial direction vector and 1D spatial distance vector between adjacent joints. Additionally, an Attribute Controllable Diffusion (ACD) module is introduced to further constrain joint associations, in which the attribute separation layer aims to separate the bone direction and length attributes, and the attribute control layer is designed to guide the pose generation by leveraging the above attributes. The ACD module utilizes the gloss embeddings as semantic conditions and finally generates sign poses from noise embeddings. Extensive experiments on PHOENIX14T and USTC-CSL datasets validate the effectiveness of our method.

Code — <https://github.com/NaVi-start/Sign-IDD>

Introduction

Sign Language Production (SLP) plays a crucial role in bridging the communication gap between the deaf and the general population, promoting inclusion and accessibility. This task is technically closely related to areas such as visual understanding (Wei et al. 2024; Li et al. 2024b; Guo et al. 2024; Li et al. 2024a) and cross-media reasoning (Song et al. 2024; Wu, Hong, and Tang 2024; Song et al. 2023). Given a textual description, SLP aims to transform it into the corresponding sequence of continuous signs automatically. These sequences can manifest as sign language

*Corresponding authors.

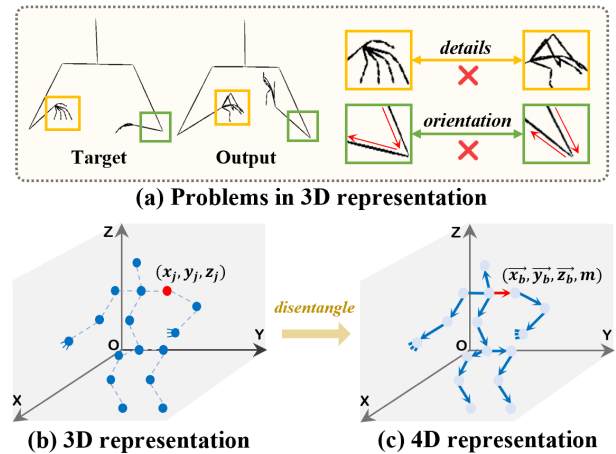


Figure 1: (a) Problems caused by using only traditional 3D representations for SLP. (b) Traditional 3D joint coordinate representation (Saunders, Camgoz, and Bowden 2020b; Tang et al. 2022b). (c) Proposed 4D bone representation. We take the neck joint as the root joint and define the parent-child joints for each bone along the skeletons. The Euclidean distance and direction vectors between parent-child joints are used to determine the length and orientation of the bones.

poses (Saunders, Camgoz, and Bowden 2020a,b) or sign language videos (Saunders, Camgoz, and Bowden 2022a). Currently, direct sign language video generation from spoken sentences remains a challenge due to the huge gap between sign vision and linguistics. Previous works usually translate spoken language into gloss¹ sequence (T2G) first and then generate sign pose video (G2P) based on gloss sequence (Saunders, Camgoz, and Bowden 2020a,b). Finally, the produced gesture poses are selectively used to generate realistic gesture videos (Saunders, Camgoz, and Bowden 2022a). Since T2G can be well addressed by Neural Machine Translation (NMT, language-to-language) based (Othman and Jemni 2011) and rule-based approaches (Moryossef et al. 2021), G2P remains the key procedure for this task at this stage and is the focus of this work.

¹Glosses refer to minimal lexical items that match the meaning of signs in linguistics.

Depending on the decoding strategy, current G2P methods are typically classified as either autoregressive (Saunders, Camgoz, and Bowden 2020a,b; Tang et al. 2022b) or non-autoregressive (Huang et al. 2021; Xie et al. 2024; Tang et al. 2024). These efforts have promoted the development of SLP tasks, especially G2P-DDM (Xie et al. 2024) and GCDM (Tang et al. 2024) as representative diffusion-based solutions, which improve the accuracy of generating sign poses. However, existing methods typically treat sign poses as discrete three-dimensional coordinates and only focus on the regression prediction of joint coordinates. These solutions overlook the modeling of relative positional relationships among joints, which hinders the effective generation of sign poses. Our goal is to generate clear and accurate gestures, especially in terms of poses that affect the semantic expression of sign language. In the example of Figure 1(a), most of the joints in the output are already close to the target distribution in spatial position. However, their relative positions are quite different from the facts, especially in the finger details that are susceptible to deviation (yellow box). In addition, the orientation of the generated limbs also shows unexpected deviations (green box).

A better constraint of bone details and limb orientation should be established in the generated pose using relative positions among the joints. Although some methods can deal with skeletal constraints by employing graph models (Saunders, Camgoz, and Bowden 2022b), they do not consider the inherent skeletal supervision and lead to higher computational costs. Fortunately, in the field of pose estimation, some emerging works have considered and verified the positive impact of constraining human bone length and orientation directly on prediction performance (Cai et al. 2024). Inspired by these works, we introduce an iconicity disentanglement strategy to enhance the relative position association (including orientations and distances) among joints.

To this end, we propose a novel iconicity disentanglement diffusion framework, Sign-IDD, which aims to improve the expressiveness and accuracy of sign language gestures by enhancing the spatial association among joints. Sign-IDD incorporates a novel Iconicity Disentanglement (ID) module, to improve the perception of relative positions among joints. As shown in Figure 1 (b) and (c), two adjacent joints along each bone are regarded as parent-child joints. The ID module is designed to disentangle the conventional 3D joint representation into a 4D bone representation, comprising the 3D spatial direction vector and 1D spatial distance vector from parent to child joint. As shown in Figure 2, we further construct a diffusion-based SLP framework, where gloss embeddings are used as semantic conditions to guide sign language gesture generation. In this framework, the Attribute Controllable Diffusion (ACD) module is another core component, which further strengthens the association learning of joints in the generated poses. The ACD module first incorporates gloss conditions into the pose embeddings to achieve semantic guidance. The attribute separation layer aims to separate the bone direction and length attributes, and the attribute control layer is designed to optimize the sign generation under the guidance of the above attributes. Our main contributions are summarized as follows:

- We innovatively introduce the concept of iconicity disentanglement, a novel strategy that transcends traditional joint coordinate regression fitting. Unlike most previous works that only adopt 3D joint coordinate representation, we utilize the disentangled 4D bone representation to further constrain the relative positions of joints, thereby ensuring the accuracy of sign pose details.
- We further propose a novel diffusion-based gloss-to-pose SLP approach, containing an attribute controllable diffusion module controlling the orientation and length of bones, capable of generating accurate and robust 3D sign poses according to textual glosses. The introduced constraint of bones (*i.e.*, \mathcal{L}_{bone}) further improves the quality of generated sign poses.
- Exhaustive experiments on the PHOENIX14T and USTC-CSL datasets show that Sign-IDD significantly enhances pose accuracy, skeletal coherence, and linguistic fidelity, which outperforms SOTA methods.

Related Work

Sign Language Production (SLP)

Sign language research is a classic and hot topic in artificial intelligence. Early works focus on Sign Language Recognition (SLR) (Guo et al. 2021; Cui, Liu, and Zhang 2017; Koller 2020; Guo et al. 2017; Wang, Chai, and Chen 2019) and Sign Language Translation (SLT) (Tang et al. 2022a; Camgoz et al. 2018, 2020; Orbay and Akarun 2020; Guo, Tang, and Wang 2019). Recently, increasing attention has been paid to Sign Language Production (SLP).

Early SLP works translate sentences into sign representations using synthetic animation techniques (Mazumder et al. 2021; McDonald et al. 2016; Segouat 2009). These methods rely on rule-based lookups to pre-capture phrases, resulting in high collection costs and limited to predefined phrases. The development of deep models has sparked extensive research in SLP. Stoll *et al.* (Stoll et al. 2020) adopt a multi-step process (*i.e.*, text-to-gloss, gloss-to-pose, and pose-to-video) to generate sign video from the text. Saunders *et al.* (Saunders, Camgoz, and Bowden 2020b) propose the first end-to-end SLP model to generate sign poses in an autoregressive manner. To improve the quality of generation, Saunders *et al.* (Saunders, Camgoz, and Bowden 2020a) also introduces a multi-channel model with an adversarial discriminator. Mixture Density (Saunders, Camgoz, and Bowden 2021a) combines transformers and mixture density networks to model multi-modal continuous sequences. FS-NET (Saunders, Camgoz, and Bowden 2022a) alleviates the error accumulation and the "mean sign pose" problem in the above autoregressive models. GEN-OBT (Tang et al. 2022b) utilizes online reverse translation to enhance constraints of semantics. SignDiff (Fang et al. 2023), G2P-DDM (Xie et al. 2024), and GCDM (Tang et al. 2024) are both diffusion-based solutions that generate coordinate representations of sign poses from Gaussian noises. However, these methods treat sign poses as discrete 3D coordinates, causing detail confusion and limb misorientation in generated poses. In contrast, our model employs an iconicity disentanglement strategy, to better assist pose generation.

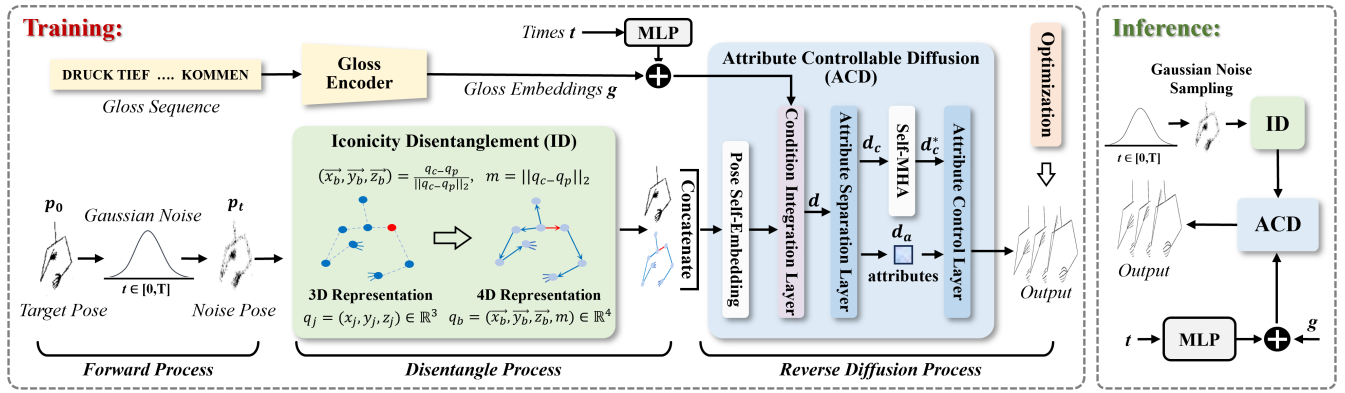


Figure 2: Overview of our framework. Given a gloss sequence, Sign-IDD generates a coherent sign pose video guided by gloss semantics. During training, we initially create Noise Pose p_t by adding Gaussian noise to Target Pose p_0 for t steps. Next, the 4D representation is derived from the 3D joint coordinates p_t through Iconicity Disentanglement (ID). Then, the combination of the 3D and 4D representations is fed into the Attribute Controllable Diffusion (ACD) module, where gloss embeddings are integrated as a semantic condition. The attribute separation and control layers aim to separate skeletal attributes and guide pose generation. The final poses are optimized by applying joint and bone constraints. During inference, the initial p_T is randomly sampled from Gaussian noise, with the disentanglement and reverse diffusion processes mirroring those used in training.

Diffusion Models

Early generation works focus on Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) and Variational AutoEncoders (VAEs) (Kingma and Welling 2014; Makhzani et al. 2015). Recently, diffusion models (Sohl-Dickstein et al. 2015) have emerged as a novel approach garnering increasing attention, which have demonstrated remarkable achievements in various domains, including image generation (Ho, Jain, and Abbeel 2020; Shang et al. 2024), text generation (Li et al. 2022), speech synthesis (Huang et al. 2022), and video generation (Ho et al. 2022).

The application of diffusion models in SLP is relatively rare and is still in its infancy. Baltatzis *et al.* propose a diffusion-based model for generating motion sequences from textual transcriptions (Baltatzis et al. 2024). G2P-DDM (Xie et al. 2024) proposes the Pose-VQVAE framework, which combines VAEs and vector quantization to transform the continuous pose space generation into a discrete sequence generation problem. GCDM (Tang et al. 2024) designs a gloss-driven conditional diffusion model and introduces a multi-hypothesis strategy to optimize sign pose generation. Different from existing diffusion-based methods, our Sign-IDD incorporates an attribute controllable diffusion module, to constrain joint associations by leveraging bone orientation and length attributes to guide the pose generation. This unique design enables the proposed solution to produce more accurate sign poses.

Preliminaries

Gloss to Pose Production

Gloss to Pose production (G2P) is a crucial step of the SLP task. Given a gloss sequence $\mathcal{G} = \{g_n | n = 1, 2, \dots, N\}$ with N glosses, G2P aims to transform it into a semantically consistent sign pose video $\mathcal{P} = \{p_s | s = 1, 2, \dots, S\}$ with S frames. The goal of G2P is to learn a mapping func-

tion $F(\mathcal{P}|\mathcal{G})$ that represents the probability distribution of generating a pose video \mathcal{P} based on the given gloss sequence \mathcal{G} . The progressive generation process can be formalized as:

$$F(\mathcal{P}|\mathcal{G}) = \prod_{s=1}^S F(p_s | p_{<s}, \mathcal{G}). \quad (1)$$

3D to 4D Representation

In existing works (Saunders, Camgoz, and Bowden 2020b; Tang et al. 2022b; Xie et al. 2024), sign poses are typically represented using a set of discrete 3D joint coordinates, which is widely adopted due to simplicity and computational efficiency. Each pose p in pose sequence \mathcal{P} is defined as a collection of joint points:

$$p = \left\{ q_j^{joint} = (x_j, y_j, z_j) \in \mathbb{R}^3 \mid j = 1, \dots, J \right\}, \quad (2)$$

where x_j , y_j and z_j represent the Cartesian coordinates of the j -th joint in 3D space, and J is the total number of joints.

However, these methods only consider the absolute positions of joints in 3D space and overlook the inherent constraints of human skeletal structure. To address these limitations, we introduce 4D bone representation, where each pose p is expressed as a collection of bones in 4D form:

$$p = \left\{ q_b^{bone} = (\vec{x}_b, \vec{y}_b, \vec{z}_b, m) \in \mathbb{R}^4 \mid b = 1, \dots, B \right\}, \quad (3)$$

where B denotes the number of bones.

Methodology

Given a gloss sequence, our goal is to generate a sign pose video with consistent semantics, as illustrated in Figure 2. We first derive 4D bone representation from the original 3D coordinates through iconicity disentanglement (Sec. **Iconicity Disentanglement of Pose**). This 4D representation, combined with 3D coordinates, is then fed into the ACD module to enable skeletal attribute controllable pose generation

The **condition integration layer** aims to introduce gloss semantics to guide pose feature embedding, which is implemented based on MHA. The process of obtaining pose features d with gloss semantics is expressed as:

$$d = MHA(Q, K, V)|_{Q=\tilde{p}_t, K=V=g}. \quad (10)$$

For the fused pose features d that already contain gloss semantics, we expect to separate the skeletal attributes (such as bone orientation and length) from them as supervisory cues to control the sign pose generation. Thus, we design an **attribute separation layer** that reprojects the pose feature d into the 7D space and separates the coordinate features $d_c \in \mathbb{R}^3$ and the attribute features $d_a \in \mathbb{R}^4$. This process is the inverse of Equation 6.

Next, the 3D coordinate features d_c are passed through an independent MHA layer, yielding updated coordinate features d_c^* . Then, an MHA-based **attribute control layer** integrates the attribute features d_a into d_c^* to refine and optimize skeletal details in sign poses. This process is formalized as:

$$d_p = MHA(Q, K, V)|_{Q=d_c^*, K=V=d_a}. \quad (11)$$

Finally, we obtain the 3D pose hypothesis \tilde{p}_0 from the noisy pose \tilde{p}_t through ACD. This constitutes the input of \mathcal{D} for the ensuing time step, expressed as

$$p_{t-1} = \sqrt{\bar{a}_{t-1}} \cdot \tilde{p}_0 + \sqrt{1 - \bar{a}_{t-1} - \sigma_t^2} \cdot \epsilon_t + \sigma_t \epsilon, \quad (12)$$

where t and $t-1$ are the current time step and the next time step, respectively, and the initial $t = T$. $\epsilon \sim \mathcal{N}(0, I)$ is a standard Gaussian noise independent of p_0 and

$$\begin{cases} \epsilon_t = (p_t - \sqrt{\bar{a}_t} \cdot \tilde{p}_0) / \sqrt{1 - \bar{a}_t}, \\ \sigma_t = \sqrt{(1 - \bar{a}_{t-1}) / (1 - \bar{a}_t)} \cdot \sqrt{1 - \bar{a}_t / \bar{a}_{t-1}}, \end{cases} \quad (13)$$

where ϵ_t is the noise at timestep t (derived from Equation 5) and σ_t controls the stochastic during the diffusion process.

In this stage, p_{t-1} is utilized as input of \mathcal{D} to regenerate and update p_0 , which is repeated I times. Initiated at T , the timestep for each iteration is computed as $T = T - (1 - i/I)$, where $i \in [0, I]$. The adjustable parameter i controls the diversity and quality of the generated results.

Pose Generation and Optimization

Pose Generation In this part, we describe the process of generating the final poses \hat{p}_0 from the output feature \tilde{p}_0 of the ACD module. In practice, we employ a Multi-Layer Perception (MLP) to reproject \tilde{p}_0 to \hat{p}_0 , represented as:

$$\hat{p}_s = MLP(LayerNorm(\tilde{p}_0)). \quad (14)$$

Here, we optimize the sign pose generation process by constraining the joint coordinates (*i.e.*, joint constraint) and the bone orientation (*i.e.*, bone constraint).

Joint Constraint Following (Huang et al. 2021; Saunders, Camgoz, and Bowden 2020b, 2021b; Viegas et al. 2023), we adopt a joint loss to constraint the accuracy of the joint positions in poses, ensuring precise matching with the ground truth. The joint constraint \mathcal{L}_{joint} is defined as follows:

$$\mathcal{L}_{joint} = \frac{1}{S} \sum_{s=1}^S |p_s - \hat{p}_s|, \quad (15)$$

where p_0^s and \hat{p}_0^s represent the ground-truth and generated 3D pose at the s -th frame, respectively.

Bone Constraint To better depict complex motion details during training, we introduce \mathcal{L}_{bone} to improve the accuracy of bone orientations in the generated poses:

$$\mathcal{L}_{bone} = \frac{1}{S} \sum_{s=1}^S (q_s^* - \hat{q}_s^*)^2, \quad (16)$$

where q_s^* and \hat{q}_s^* represent the bone orientations derived from p_0^s and \hat{p}_0^s according to Equation 4. The overall objective is:

$$\mathcal{L} = \mathcal{L}_{joint} + \lambda \mathcal{L}_{bone}. \quad (17)$$

Experiments

Experimental Settings

Datasets We evaluate the proposed method on two benchmarks: PHOENIX14T (Camgoz et al. 2018) and USTC-CSL (Huang et al. 2018). PHOENIX14T consists of 8,257 instances featuring 2,887 unique German words and 1,066 signs, known for its complexity. USTC-CSL encompasses 100 Chinese sign language sentences performed by 50 signers and is divided into 4,000 training instances and 1,000 testing instances (Guo et al. 2018).

Evaluation Metrics Following the existing works (Huang et al. 2021; Saunders, Camgoz, and Bowden 2020b, 2021b; Viegas et al. 2023), a SLT model named NSLT (Camgoz et al. 2018) is employed to back-translate sign poses into textual glosses and compare them with references for calculating metrics such as BLEU, ROUGE, and WER. In addition, we also report the FID, Mean Per Joint Position Error (MPJPE), and Mean Per Joint Angle Error (MPJAE) to directly measure the quality of generated poses.

Implementation Details Since PHOENIX14T lacks pose labels, we use OpenPose (Cao et al. 2017) to extract 2D joint coordinates and convert them to 3D using a skeletal correction model (Zelinka and Kanis 2020) as target poses. The Transformer-based Gloss Encoder is built with 2 layers, 4 heads, and an embedding size of 512. In addition, we set the timesteps t of the diffusion model to 1000 and the number of inferences i to 5. During training, we use the Adam optimizer (Kingma and Ba 2015) and a learning rate of 1×10^{-3} . Experiments are conducted using PyTorch on NVIDIA GeForce RTX 2080 Ti GPUs.

Comparison with State-of-the-Arts

PHOENIX14T Table 1 provides comparison results of the proposed Sign-IDD with other SOTA methods on PHOENIX14T. As shown in this table, Sign-IDD significantly outperforms other non-diffusion-based approaches, achieving 25.40% and 24.80% BLEU-1 on the DEV and TEST sets, respectively. Even compared with the best-performing non-diffusion-based method, our method achieves significant performance improvements on ROUGE and WER, *e.g.*, Sign-IDD surpasses GEN-OBT (Tang et al. 2022b) by margins of 3.09% and 5.12% on the TEST set. Considering the strong advantages of the diffusion model itself in content generation, we specifically compare our solution with several diffusion-based methods. It is noticeable

Methods	DEV							TEST						
	B1↑	B4↑	ROUGE↑	WER↓	FID↓	MPJPE↓	MPJAE↓	B1↑	B4↑	ROUGE↑	WER↓	FID↓	MPJPE↓	MPJAE↓
Ground Truth	29.77	12.13	29.60	74.17	0.00	0.00	0.00	29.76	11.93	28.98	71.94	0.00	0.00	0.00
NDBM														
PT-base [†] _{ECCV2020}	9.53	0.72	8.61	98.53	2.90	41.92	33.74	9.47	0.59	8.88	98.36	3.22	51.35	33.17
PT-GN [†] _{ECCV2020}	12.51	3.88	11.87	96.85	2.98	40.63	28.25	13.35	4.31	13.17	96.50	3.33	50.8	28.81
NAT-AT _{MM2021}	–	–	–	–	–	–	–	14.26	5.53	18.72	88.15	–	–	–
NAT-EA _{MM2021}	–	–	–	–	–	–	–	15.12	6.66	19.43	82.01	–	–	–
DET [*] _{SEM2023}	17.25	5.32	17.85	–	–	–	–	17.18	5.76	17.64	–	–	–	–
GEN-OBT _{MM2022}	24.92	8.68	25.21	82.36	2.54	41.47	26.64	23.08	8.01	23.49	81.78	2.97	52.9	27.53
DBM														
D3DP-sign [†] _{ICCV2023}	17.20	5.01	17.94	91.51	2.38	39.42	25.73	16.51	5.25	17.55	91.83	2.63	47.65	25.92
G2P-DDM _{AAAI2024}	–	–	–	–	–	–	–	16.11	7.50	–	77.26	–	–	–
GCDM _{TOMM2024}	22.88	7.64	23.35	82.81	–	–	–	22.03	7.91	23.20	81.94	–	–	–
Sign-IDD (Ours)	25.40	8.93	27.60	77.72	2.22	39.11	25.34	24.80	9.08	26.58	76.66	2.46	47.19	25.37

Table 1: Performance comparison on PHOENIX14T. ‘†’ indicates the model is tested by us under a fair setting. **NDBM**: Non-Diffusion Based Methods; **DBM**: Diffusion Based Methods.

Methods	B1↑	WER↓	FID↓	MPJPE↓	MPJAE↓
Ground Truth	69.10	47.38	0.00	0.00	0.00
PT-base _{ECCV2020}	22.32	87.64	0.49	175.14	21.93
PT-GN _{ECCV2020}	24.42	84.01	0.46	103.44	18.97
GEN-OBT _{MM2022}	38.31	70.50	0.41	78.98	12.86
D3DP-sign _{ICCV2023}	59.37	53.62	0.34	79.27	11.08
Sign-IDD (Ours)	65.26	50.15	0.31	72.20	10.92

Table 2: Performance comparison on USTC-CSL.

that our method achieves higher BLEU than the most recent diffusion-based SLP method, *e.g.*, Sign-IDD is 8.69% and 1.58% higher than G2P-DDM (Xie et al. 2024) on BLEU-1 and BLEU-4 metrics. Compared with diffusion-based methods using multiple hypothesis strategies during inference, *e.g.*, D3DP-sign (Shan et al. 2023) and GCDM (Tang et al. 2024), our method still achieves superior performance.

USTC-CSL Table 2 shows comparison results on a challenging Chinese sign language benchmark USTC-CSL. Note that no existing work has reported the SLP performance on USTC-CSL, so we test several typical solutions under a fair setting. Sign-IDD achieves the best performance on all back-translation metrics, *e.g.*, 65.26% and 50.15% on BLEU-1 and WER, which indicates that Sign-IDD can maintain high semantic accuracy during sign generation. In addition, our performance on direct metrics outperforms previous non-diffusion-based solutions and is especially significantly better than diffusion-based D3DP-sign (Shan et al. 2023) by 7.07% on MPJPE. This further demonstrates the superiority of the poses generated by Sign-IDD in terms of the accuracy of joint positions and bone orientations.

Visualization Results In Figure 4, we visualize sign pose examples generated by the proposed Sign-IDD and other methods, *i.e.*, PT-GN (Saunders, Camgoz, and Bowden 2020b) and GEN-OBT (Tang et al. 2022b). In the top example, the sign poses generated by Sign-IDD are noticeably superior to PT-GN, particularly in hand details (blue box), and demonstrate more accurate limb movements compared to GEN-OBT (red box). Furthermore, even in cases where the ground truth provides inaccurate pose labels due

Methods	DEV			TEST		
	B1↑	B4↑	WER↓	B1↑	B4↑	WER↓
Base	21.38	7.06	85.29	21.50	7.11	84.74
Base+ID	24.97	8.51	78.52	23.46	8.21	77.53
Base+ACD	24.33	8.40	82.04	23.42	8.16	79.92
Base+ID+ACD	25.40	8.93	77.72	24.80	9.08	76.66

Table 3: Ablation results of modules. **ACD**: Attribute Controllable Diffusion, **ID**: Iconicity Disentanglement.

Methods		DEV			TEST		
		B1↑	B4↑	WER↓	B1↑	B4↑	WER↓
λ	0.01	21.34	6.93	85.53	20.89	6.97	84.43
	0.1	25.40	8.93	77.72	24.80	9.08	76.66
	1	20.85	6.85	88.66	19.95	6.46	88.38
t	500	24.11	8.07	81.56	23.63	8.43	79.90
	1000	25.40	8.93	77.72	24.80	9.08	76.66
	1500	23.09	7.96	83.88	22.77	7.53	83.99
i	1	24.05	8.38	82.39	23.60	8.42	81.73
	5	25.40	8.93	77.72	24.80	9.08	76.66
	10	20.85	6.85	88.66	19.95	6.46	88.38

Table 4: Ablation results of parameters on PHOENIX14T.

to motion blur (yellow box), Sign-IDD consistently generates clear and precise results. The bottom example shows that Sign-IDD outperforms existing methods in generating both upper and lower limb movements, further highlighting the advantages of our approach in accurately generating joint positions and bone orientations during sign production.

Ablation Study

In this subsection, we present ablation results to verify the effectiveness of the Sign-IDD. All results are evaluated on PHOENIX14T, while USTC-CSL is not used for ablations.

Retain Stronger Sign Semantics Table 3 shows the ablation results of ID and ACD modules. We set a diffusion-based SLP model as a baseline, denoted as Base, which only uses 3D joint coordinates and replaces the proposed ACD with a denoiser that removes attribute separation/control layers. Base+ID refers to incorporating ID into Base (*i.e.*, introducing iconicity disentanglement strategy), which improves

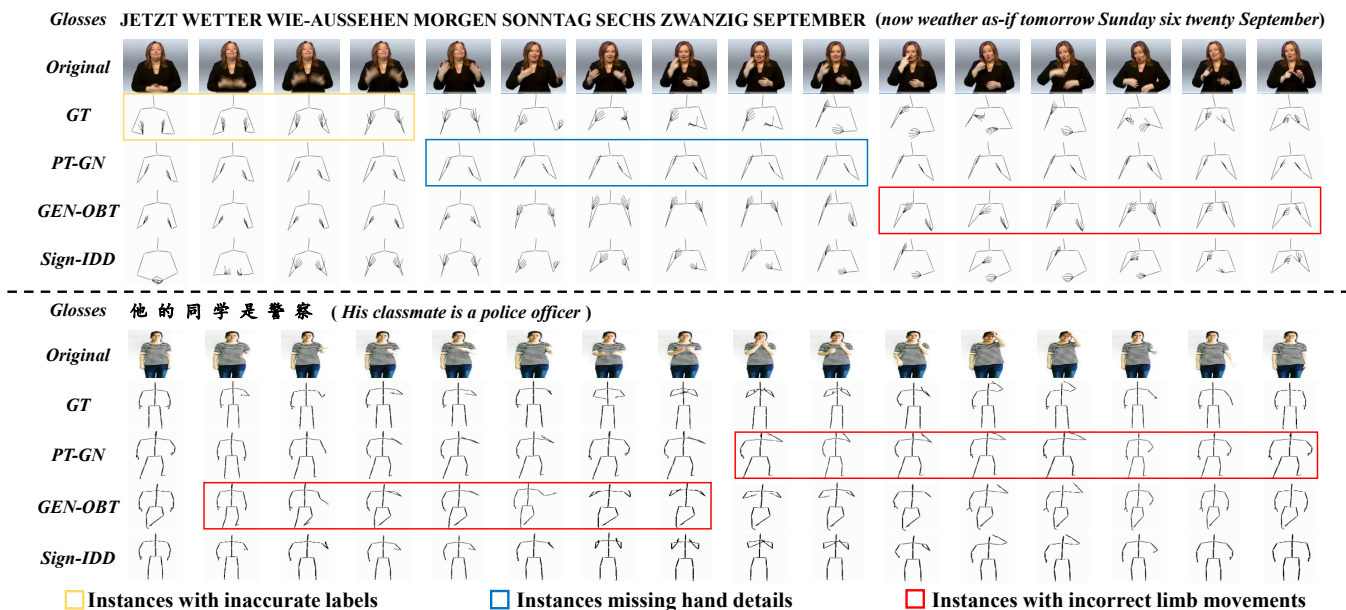


Figure 4: Visualization examples of generated poses on PHOENIX14T (top) and USTC-CSL (bottom). We compare Sign-IDD with PT-GN and GEN-OBT. Gloss annotations, original video frames, and ground-truth poses are attached for clear evaluation.

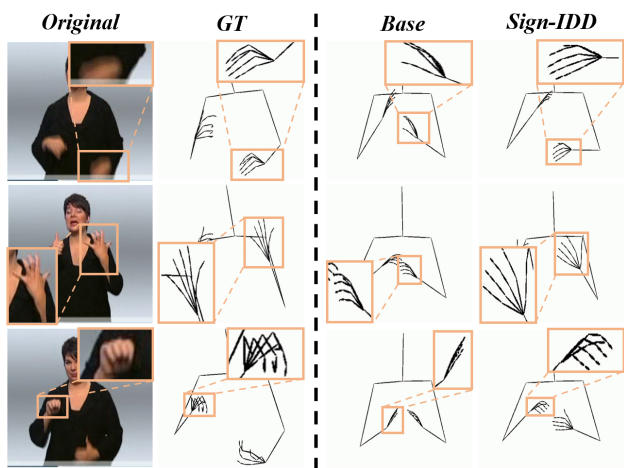


Figure 5: Visualization results of Sign-IDD and Base, which is a diffusion-based baseline without ID and ACD modules.

BLEU-1/BLEU-4 by 1.96%/1.10% on TEST. In addition, introducing ACD with an attribute control mechanism based on Base also improves performance, and the WER reaches 82.04%/79.92% (compared to 85.29%/84.74%). When both ID and ACD modules are adopted, our method achieves the best performance on all back-translation metrics, which indicates that poses generated by Sign-IDD retain richer and more accurate sign semantics.

Generate Clearer Local Details In Figure 5, the visualization further shows the direct impact of the proposed ID and ACD modules on the generated sign poses. In the top and bottom examples, the Base generates overlapping

fingers, obscuring important hand movements. In contrast, Sign-IDD generates clearer local details and the posture of fingers is also closer to the truth.

Capture More Accurate Limb Orientation In the middle example of Figure 5, although Base correctly captures the relative positions of the hand joints, the finger bone orientations differ significantly from the ground truth. In contrast, Sign-IDD accurately aligns both joint positions and bone orientations, with all fingers pointing in the correct direction, closely matching the target pose. Additionally, the top example further demonstrates that incorporating iconicity disentanglement strategy and attribute controllable diffusion enhances the orientation accuracy of generated limbs.

Analysis of Parameters Table 4 presents the ablation results for Sign-IDD, focusing on the weight λ of the \mathcal{L}_{bone} , the time step t in diffusion, and the sampling step i during inference. The model achieves optimal performance with $\lambda = 0.1$, $t = 1000$, and $i = 5$.

Conclusions

This work proposes to enhance the accuracy and naturalness of generated sign poses by modeling both positions and associations of joints. We introduce a novel iconicity disentanglement strategy that transforms 3D joint representations into 4D bone representations, covering 3D orientation and 1D length. We further design an attribute controllable diffusion module to separate skeletal attributes and utilize them to guide pose generation. Extensive experiments on two benchmarks validate the effectiveness of our approach.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grants No. U23B2031, 61932009, U20A20183, 62272144, 62302141, 62331003), the Anhui Provincial Natural Science Foundation, China (Grant No. 2408085QF191), the Major Project of Anhui Province (Grant No. 202423k09020001), and the Fundamental Research Funds for the Central Universities (Grants No. JZ2024HGTA0178, JZ2024HGTA0255).

References

- Baltatzis, V.; Potamias, R. A.; Ververas, E.; Sun, G.; Deng, J.; and Zafeiriou, S. 2024. Neural Sign Actors: A Diffusion Model for 3D Sign Language Production from Text. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1985–1995.
- Cai, Q.; Hu, X.; Hou, S.; Yao, L.; and Huang, Y. 2024. Disentangled Diffusion-Based 3D Human Pose Estimation with Hierarchical Spatial and Temporal Denoiser. In *AAAI Conference on Artificial Intelligence*, 882–890.
- Camgoz, N. C.; Hadfield, S.; Koller, O.; Ney, H.; and Bowden, R. 2018. Neural Sign Language Translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7784–7793.
- Camgoz, N. C.; Koller, O.; Hadfield, S.; and Bowden, R. 2020. Sign Language Transformers: Joint End-to-End Sign Language Recognition and Translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10023–10033.
- Cao, Z.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2017. Realtime Multi-Person 2d Pose Estimation Using Part Affinity Fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7291–7299.
- Cui, R.; Liu, H.; and Zhang, C. 2017. Recurrent Convolutional Neural Networks for Continuous Sign Language Recognition by Staged Optimization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7361–7369.
- Fang, S.; Sui, C.; Zhang, X.; and Tian, Y. 2023. SignDiff: Learning Diffusion Models for American Sign Language Production. *arXiv Preprint ArXiv:2308.16082*, arXiv–2308.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. *Advances in Neural Information Processing Systems*, 2672–2680.
- Guo, D.; Tang, S.; Hong, R.; and Wang, M. 2021. Sign Language Recognition. *Multimedia for Accessible Human Computer Interfaces*, 23–59.
- Guo, D.; Tang, S.; and Wang, M. 2019. Connectionist Temporal Modeling of Video and Language: A Joint Model for Translation and Sign Labeling. In *International Joint Conference on Artificial Intelligence*, 751–757.
- Guo, D.; Zhou, W.; Li, H.; and Wang, M. 2017. Online Early-Late Fusion Based on Adaptive HMM for Sign Language Recognition. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 1–18.
- Guo, D.; Zhou, W.; Li, H.; and Wang, M. 2018. Hierarchical LSTM for Sign Language Translation. In *AAAI Conference on Artificial Intelligence*, 6845–6852.
- Guo, M.; He, J.; Tang, S.; Wang, Z.; and Cheng, L. 2024. Shaping a Stabilized Video by Mitigating Unintended Changes for Concept-Augmented Video Editing. *arXiv preprint arXiv:2410.12526*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. *Advances in Neural Information Processing Systems*, 6840–6851.
- Ho, J.; Salimans, T.; Gritsenko, A.; Chan, W.; Norouzi, M.; and Fleet, D. J. 2022. Video Diffusion Models. *Advances in Neural Information Processing Systems*, 8633–8646.
- Huang, J.; Zhou, W.; Zhang, Q.; Li, H.; and Li, W. 2018. Video-Based Sign Language Recognition without Temporal Segmentation. In *AAAI Conference on Artificial Intelligence*, 2257–2264.
- Huang, R.; Lam, M.; Wang, J.; Su, D.; Yu, D.; Ren, Y.; and Zhao, Z. 2022. FastDiff: A Fast Conditional Diffusion Model for High-Quality Speech Synthesis. In *International Joint Conference on Artificial Intelligence*, 4157–4163.
- Huang, W.; Pan, W.; Zhao, Z.; and Tian, Q. 2021. Towards Fast and High-Quality Sign Language Production. In *ACM International Conference on Multimedia*, 3172–3181.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 1–15.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. *Stat.*, 1.
- Koller, O. 2020. Quantitative Survey of the State of the Art in Sign Language Recognition. *arXiv Preprint ArXiv:2008.09918*, arXiv–2008.
- Li, F.; Cong, R.; Wu, J.; Bai, H.; Wang, M.; and Zhao, Y. 2024a. SRConvNet: A Transformer-Style ConvNet for Lightweight Image Super-Resolution. *International Journal of Computer Vision*, 1–17.
- Li, F.; Wu, Y.; Li, A.; Bai, H.; Cong, R.; and Zhao, Y. 2024b. Enhanced Video Super-Resolution Network towards Compressed Data. *ACM Transactions on Multimedia Computing, Communications and Applications*, 1–21.
- Li, X.; Thickstun, J.; Gulrajani, I.; Liang, P. S.; and Hashimoto, T. B. 2022. Diffusion-LM Improves Controllable Text Generation. *Advances in Neural Information Processing Systems*, 4328–4343.
- Makhzani, A.; Shlens, J.; Jaitly, N.; Goodfellow, I.; and Frey, B. 2015. Adversarial Autoencoders. *arXiv e-prints*, arXiv–1511.
- Mazumder, S.; Mukhopadhyay, R.; Namboodiri, V. P.; and Jawahar, C. 2021. Translating Sign Language Videos to Talking Faces. In *Indian Conference on Computer Vision, Graphics and Image Processing*, 1–10.
- McDonald, J.; Wolfe, R.; Schnepf, J.; Hochgesang, J.; Jamrozik, D. G.; Stumbo, M.; Berke, L.; Bialek, M.; and Thomas, F. 2016. An Automated Technique for Real-Time Production of Lifelike Animations of American Sign Language. *Universal Access in the Information Society*, 551–566.

- Moryossef, A.; Yin, K.; Neubig, G.; and Goldberg, Y. 2021. Data Augmentation for Sign Language Gloss Translation. In *International Workshop on Automatic Translation for Signed and Spoken Languages*, 1–11.
- Nielsen, A. K.; and Dingemans, M. 2021. Iconicity in Word Learning and Beyond: A Critical Review. *Language and Speech*, 52–72.
- Orbay, A.; and Akarun, L. 2020. Neural Sign Language Translation by Learning Tokenization. In *IEEE International Conference on Automatic Face and Gesture Recognition*, 222–228.
- Othman, A.; and Jemni, M. 2011. Statistical Sign Language Machine Translation: from English Written Text to American Sign Language Gloss. *International Journal of Computer Science Issues*, 65.
- Saunders, B.; Camgoz, N. C.; and Bowden, R. 2020a. Adversarial Training for Multi-Channel Sign Language Production. In *British Machine Vision Conference*, 1–15.
- Saunders, B.; Camgoz, N. C.; and Bowden, R. 2020b. Progressive Transformers for End-to-End Sign Language Production. In *European Conference on Computer Vision*, 687–705.
- Saunders, B.; Camgoz, N. C.; and Bowden, R. 2021a. Continuous 3d Multi-Channel Sign Language Production via Progressive Transformers and Mixture Density Networks. *International journal of computer vision*, 2113–2135.
- Saunders, B.; Camgoz, N. C.; and Bowden, R. 2021b. Mixed Signals: Sign Language Production via A Mixture of Motion Primitives. In *IEEE/CVF International Conference on Computer Vision*, 1919–1929.
- Saunders, B.; Camgoz, N. C.; and Bowden, R. 2022a. Signing at Scale: Learning to Co-Articulate Signs for Large-Scale Photo-Realistic Sign Language Production. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5141–5151.
- Saunders, B.; Camgoz, N. C.; and Bowden, R. 2022b. Skeletal Graph Self-Attention: Embedding a Skeleton Inductive Bias into Sign Language Production. In *International Workshop on Sign Language Translation and Avatar Technology*, 95–102.
- Segouat, J. 2009. A Study of Sign Language Coarticulation. *ACM Sigaccess Accessibility and Computing*, 31–38.
- Shan, W.; Liu, Z.; Zhang, X.; Wang, Z.; Han, K.; Wang, S.; Ma, S.; and Gao, W. 2023. Diffusion-Based 3D Human Pose Estimation with Multi-Hypothesis Aggregation. In *IEEE/CVF International Conference on Computer Vision*, 14761–14771.
- Shang, S.; Shan, Z.; Liu, G.; Wang, L.; Wang, X.; Zhang, Z.; and Zhang, J. 2024. Resdiff: Combining Cnn and Diffusion Model for Image Super-Resolution. In *AAAI Conference on Artificial Intelligence*, 8975–8983.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep Unsupervised Learning Using Nonequilibrium Thermodynamics. In *International Conference on Machine Learning*, 2256–2265.
- Song, P.; Guo, D.; Yang, X.; Tang, S.; and Wang, M. 2024. Emotional Video Captioning With Vision-Based Emotion Interpretation Network. *IEEE Transactions on Image Processing*, 1122–1135.
- Song, P.; Guo, D.; Yang, X.; Tang, S.; Yang, E.; and Wang, M. 2023. Emotion-Prior Awareness Network for Emotional Video Captioning. In *ACM International Conference on Multimedia*, 589–600.
- Stoll, S.; Camgoz, N. C.; Hadfield, S.; and Bowden, R. 2020. Text2Sign: Towards Sign Language Production Using Neural Machine Translation and Generative Adversarial Networks. *International Journal of Computer Vision*, 891–908.
- Tang, S.; Guo, D.; Hong, R.; and Wang, M. 2022a. Graph-Based Multimodal Sequential Embedding for Sign Language Translation. *IEEE Transactions on Multimedia*, 4433–4445.
- Tang, S.; Hong, R.; Guo, D.; and Wang, M. 2022b. Gloss Semantic-Enhanced Network with Online Back-Translation for Sign Language Production. In *ACM International Conference on Multimedia*, 5630–5638.
- Tang, S.; Xue, F.; Wu, J.; Wang, S.; and Hong, R. 2024. Gloss-Driven Conditional Diffusion Models for Sign Language Production. *ACM Transactions on Multimedia Computing, Communications, and Applications*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All You Need. In *Advances in Neural Information Processing Systems*, 6000–6010.
- Viegas, C.; Inan, M.; Quandt, L.; and Alikhani, M. 2023. Including Facial Expressions in Contextual Embeddings for Sign Language Generation. In *Joint Conference on Lexical and Computational Semantics*, 1–10.
- Wang, H.; Chai, X.; and Chen, X. 2019. A Novel Sign Language Recognition Framework Using Hierarchical Grassmann Covariance Matrix. *IEEE Transactions on Multimedia*, 2806–2814.
- Wei, Y.; Zhang, Y.; Li, K.; Wang, F.; Tang, S.; and Zhang, Z. 2024. Leveraging Vision-Language Prompts for Real-World Image Restoration and Enhancement. *Computer Vision and Image Understanding*, 104222.
- Wu, J.; Hong, R.; and Tang, S. 2024. Intermediary-Generated Bridge Network for RGB-D Cross-Modal Re-Identification. *ACM Transactions on Intelligent Systems and Technology*, 1–25.
- Xie, P.; Zhang, Q.; Taiying, P.; Tang, H.; Du, Y.; and Li, Z. 2024. G2P-DDM: Generating Sign Pose Sequence from Gloss Sequence with Discrete Diffusion Model. In *AAAI Conference on Artificial Intelligence*, 6234–6242.
- Zelinka, J.; and Kanis, J. 2020. Neural Sign Language Synthesis: Words Are Our Glosses. In *IEEE/CVF Winter Conference on Applications of Computer Vision*, 3395–3403.