

# Hand1000: Generating Realistic Hands from Text with Only 1,000 Images

Haozhuo Zhang<sup>1</sup>, Bin Zhu<sup>2\*</sup>, Yu Cao<sup>2</sup>, Yanbin Hao<sup>3</sup>

<sup>1</sup>Peking University

<sup>2</sup>Singapore Management University

<sup>3</sup>University of Science and Technology of China

2100013132@stu.pku.edu.cn, binzhu@smu.edu.sg, yu.cao.2022@msc.smu.edu.sg, haoyanbin@hotmail.com

## Abstract

Text-to-image generation models have achieved remarkable advancements in recent years, aiming to produce realistic images from textual descriptions. However, these models often struggle with generating anatomically accurate representations of human hands. The resulting images frequently exhibit issues such as incorrect numbers of fingers, unnatural twisting or interlacing of fingers, or blurred and indistinct hands. These issues stem from the inherent complexity of hand structures and the difficulty in aligning textual descriptions with precise visual depictions of hands. To address these challenges, we propose a novel approach named **Hand1000** that enables the generation of realistic hand images with target gesture using only 1,000 training samples. The training of Hand1000 is divided into three stages with the first stage aiming to enhance the model’s understanding of hand anatomy by using a pre-trained hand gesture recognition model to extract gesture representation. The second stage further optimizes text embedding by incorporating the extracted hand gesture representation, to improve alignment between the textual descriptions and the generated hand images. The third stage utilizes the optimized embedding to fine-tune the Stable Diffusion model to generate realistic hand images. In addition, we construct the first publicly available dataset specifically designed for text-to-hand image generation. Based on the existing hand gesture recognition dataset, we adopt advanced image captioning models and LLaMA3 to generate high-quality textual descriptions enriched with detailed gesture information. Extensive experiments demonstrate that Hand1000 significantly outperforms existing models in producing anatomically correct hand images while faithfully representing other details in the text, such as faces, clothing and colors.

## Project Page —

<https://haozhuo-zhang.github.io/Hand1000-project-page/>

## Introduction

In recent years, text-to-image generation models have achieved significant progress, such as Stable Diffusion (Rombach et al. 2022; Esser et al. 2024), Multi-modal Large Language Model (MM-LLM) (Koh, Fried, and Salakhutdinov 2024), and DiT (Peebles and Xie 2023).

\*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aai.org). All rights reserved.



Figure 1: Comparison of hand image generation results between Stable Diffusion and our Hand1000. Given the same text prompt, Stable Diffusion produces deformed and chaotic hands. In contrast, our proposed Hand1000 manages to generate anatomically correct and realistic hands while preserving details such as character, clothing, and colors.

However, these models often struggle when generating images involving human bodies, particularly hands. The generated images frequently exhibit distorted or misaligned body parts that do not conform to normal physical characteristics. As shown in the top of Figure 1, this problem is especially pronounced with hands, which often appear blurry, with fingers twisted or interlaced, and an incorrect number of fingers. Despite capturing the color and texture of human hands, these generated images clearly display noticeable inaccuracies.

The underlying causes of this phenomenon can be attributed to two primary factors: (a) The inherent complexity of the human hand’s structure, along with the occlusions between fingers, means that the Stable Diffusion model lacks the necessary prior knowledge of the precise physical structure of human hands, resulting in its inability to generate anatomically correct hands. (b) When the text prompt includes information related to human hands (e.g., “thumbs up” and “phone call gesture”), the model struggles to as-

sociate this with corresponding hand images, indicating its inability to accurately comprehend text containing hand-related information and thus failing to produce realistic hand visual appearance. Recent studies have explored methods to generate anatomically correct hands using diffusion models, such as (Narasimhaswamy et al. 2024) and (Lu et al. 2023), yet they fail to fully address the fundamental issues outlined in the previous paragraph. Furthermore, these approaches require extensive training datasets, often comprising hundreds of thousands of images, which significantly reduces training efficiency and limits their practical applicability. In contrast, recent advancements in text-based image editing techniques (Kawar et al. 2023; Lin et al. 2024; Bar-Tal et al. 2022) have demonstrated the ability to achieve background alterations or modifications in facial expressions and actions with only a few images.

Inspired by text-based image editing works, we propose a novel method named Hand1000 to generate anatomically correct hands using only 1,000 images for each target hand gesture in three stages. To address the issue (a) from the previous paragraph, our proposed method enhances hand prior knowledge about correct hand anatomy by leveraging a pre-trained hand gesture recognition model (i.e., Mediapipe hands (Zhang et al. 2020)) to extract gesture representation in the first stage. To tackle issue (b), in the second stage, we optimize the text embedding for each image by integrating the hand gesture representation in the first stage to ensure alignment between text and hand image. This optimized embedding is then used to fine-tune the stable diffusion model to generate realistic hand images in the third stage. In addition, we construct a dataset specifically for text-to-hand-image generation by leveraging hand gesture recognition datasets (i.e. HaGRID dataset (Kapitanov et al. 2024)). Image captioning models (e.g. BLIP2 (Li et al. 2023b), PaliGemma (Beyer et al. 2024), and VitGpt2 (Mishra et al. 2024)) are first used to produce a textual description, and LLaMA3 model (Touvron et al. 2023a) is employed to enrich the textual description with hand gesture information. After training on 1,000 images in our dataset, our Hand1000 significantly outperforms the Stable Diffusion model in generating correct hands for the same textual prompts, which is shown at the bottom of Figure 1.

In summary, the contributions of this paper are as follows:

- We propose a new method, Hand1000, that empowers the Stable Diffusion model with the ability to generate realistic hands using only 1,000 images. To the best of our knowledge, Hand1000 is the first text-to-image model capable of achieving accurate hand generation with such a small-scale dataset.
- By leveraging image captioning models and the LLaMA3, we construct the first publicly available dataset specifically designed for generating hand images from textual descriptions. This dataset aims to facilitate and advance research within the community on this task.
- We showcase the power of enhancing hand features and text optimization to generate anatomically correct and realistic hand images from text. Hand1000 achieves significant improvements compared with existing methods.

## Related Work

### Text-to-Image Generation

Text-based image generation has long been a prominent research area, with classical methods including Generative Adversarial Networks (GANs) and their variants (Kang et al. 2023; Tao et al. 2022; Xu et al. 2018; Zhu and Ngo 2020; Ye et al. 2021; Zhu et al. 2019b,a; Wang et al. 2024; Li et al. 2023a; Ma et al. 2023), as well as stable diffusion models and their derivatives (Sauer et al. 2023; Saharia et al. 2022b; Dhariwal and Nichol 2021; Ho, Jain, and Abbeel 2020; Ho et al. 2022; Sohl-Dickstein et al. 2015; Yang et al. 2024; Ye et al. 2024; Yu et al. 2024). These methods leverage the capabilities of CLIP (Radford et al. 2021) and Transformer (Vaswani et al. 2017), empowering them to generate images from textual descriptions. In recent years, numerous related studies have emerged, achieving notable success in the field such as GigaGAN (Kang et al. 2023), Imagen (Saharia et al. 2022b), and SDXL Turbo (Sauer et al. 2023). With the advent of ChatGPT (Brown et al. 2020; Achiam et al. 2023) and the development of large language models (LLM) (Chowdhery et al. 2023; Touvron et al. 2023b; Team et al. 2023), several multimodal models for text-to-image generation have also emerged, such as DALL-E (Ramesh et al. 2021). Despite the notable advancements these methods have achieved in image generation, they frequently encounter significant issues when producing images that include human hands, often resulting in severe distortions or anomalies in hand representations.

### Realistic Hand Generation

Text-based diffusion models are effective for generating various images but struggle with rendering human hands. Recent works address this challenge in two ways. Some studies use image inpainting, like HandRefiner (Lu et al. 2023), which identifies and crops hand regions, then refines them using ControlNet (Zhang, Rao, and Agrawala 2023). However, this approach depends on another model for initial image generation and fails with severely distorted hands. Alternatively, methods like HandDiffuser (Narasimhaswamy et al. 2024) directly generate hand images by reconstructing text prompts into 3D hand parameters, though the training process is complex and expensive. Additionally, the above methods require an enormous training dataset, in the order of hundreds of thousands of samples, to achieve noticeable results. In contrast, our **Hand1000** method fine-tunes on just **1,000** images to generate realistic hands without requiring explicit hand pose modeling, offering a more efficient and flexible approach. Note that we cannot use HandDiffuser (Narasimhaswamy et al. 2024) as a baseline since its source code is not publicly available.

### Text-based Image Editing

Text-based image editing has seen rapid advancements in recent years, allowing for the modification of an image’s background, style, subject size, and movements through direct textual input. These technologies, often derived from improvements in GANs (Abdal et al. 2021; Härkönen et al. 2020; Lang et al. 2021; Patashnik et al. 2021; Shen et al.

2020; Shen and Zhou 2021; Xia et al. 2023) or Stable Diffusion models (Ruiz et al. 2023; Kawar et al. 2023; Nichol and Dhariwal 2021; Rombach et al. 2022; Saharia et al. 2022a,c; Song, Meng, and Ermon 2020; Song and Ermon 2019, 2020), typically require minimal input images to achieve the desired effects. For instance, Imagic (Kawar et al. 2023) firstly fine-tunes text embeddings, then fine-tunes the diffusion model, and finally performs interpolation. This method focuses on text embeddings, and achieves remarkable results. Considering the challenge of generating realistic hand images, it seems plausible to borrow from the aforementioned image editing concepts. However, image editing algorithms cannot be directly applied to generate realistic hand images because they fundamentally rely on the diffusion model’s inherent capabilities, which lack prior knowledge of realistic hands. Besides, these methods usually take several minutes to modify a single picture, making it impossible to generate a large amount of images in a short time. To address these problems, we propose **Hand1000**, a method that fine-tunes the model using image editing principles. This approach not only imparts the model with knowledge of realistic hands but also benefits from the characteristics of image editing algorithms, requiring fewer training images and faster training speed.

## Method

This section begins by introducing the preliminaries of the text-to-image generation model, followed by a detailed explanation of the Hand1000 method. As depicted in Figure 2, the training process of Hand1000 is divided into three stages. In the first stage, 1,000 images from the training set are processed by a hand gesture recognition model (i.e. Mediapipe hands (Zhang et al. 2020)) to obtain mean hand gesture feature. The second stage involves sequentially feeding text-image pairs for training, where text embedding is obtained using CLIP (Radford et al. 2021) from text description. The text embedding is then concatenated with the mean hand gesture feature from the first stage, which is mapped to obtain the fused embedding. The fused embedding is linearly combined with text embedding to produce double-fused embedding, which is optimized by using the reconstruction loss with a frozen stable diffusion model to produce optimized embedding. The third stage leverages the frozen optimized embedding to further fine-tune the stable diffusion model to generate realistic hand images.

### Preliminaries

Text-to-image generation aims to generate an image that accurately represents a given textual description. The process of a standard text-to-image diffusion model (Rombach et al. 2022) is as follows: Given an initial noise image  $\epsilon \sim \mathcal{N}(0, I)$ , where the noise follows a Gaussian distribution, the CLIP text encoder (Radford et al. 2021)  $\Gamma$  encodes the *text* into a text embedding  $e = \Gamma(\textit{text})$ . Using  $e$  and the time step  $t$  as conditions, the diffusion model generates a latent representation  $\mathbf{z}_t := \alpha_t \mathbf{x} + \sigma_t \epsilon$  based on the initial noise image  $\epsilon$ . Subsequently, an image is generated from this latent representation. The entire process is guided by the following

mean squared error loss:

$$\mathbb{E}_{\mathbf{x}, \mathbf{e}, \epsilon, t} \left[ w_t \|\hat{\mathbf{x}}_\theta(\alpha_t \mathbf{x} + \sigma_t \epsilon, \mathbf{e}) - \mathbf{x}\|_2^2 \right], \quad (1)$$

where  $\mathbf{x}$  represents the real image, and  $\mathbf{e}$  denotes the text embedding. The terms  $\alpha_t$ ,  $\sigma_t$ , and  $w_t$  are parameters that control the noise scheduling and sample quality, and they are functions of the diffusion process time  $t \sim U([0, 1])$ .

### Stage I: Hand Gesture Feature Extraction

Assume we have a set of images depicting a target hand gesture, such as “phone call”. To extract hand features associated with this gesture, we feed the images into a gesture recognition model (i.e. Mediapipe hands (Zhang et al. 2020)) to obtain features from the final layer of the network. Subsequently, these features are averaged to obtain a **Mean Hand Gesture Feature** representation of the gesture, which is used for training in the following stages. Simultaneously, the Mean Hand Gesture Feature is also preserved and utilized during the inference phase.

### Stage II: Text Embedding Optimization

In the second stage of training, the main objective is to integrate text embedding with hand gesture features to facilitate the diffusion model to generate realistic hand images. Given a textual description as input, CLIP text encoder is employed to obtain the **Text Embedding**. Subsequently, the text embedding is concatenated with the **Mean Hand Gesture Feature** obtained from the first stage. To fuse the text and hand gesture features, a fully connected (FC) layer is employed to map the concatenated embedding to the same dimension of text embedding, resulting in the **Fused Embedding**, which smoothly incorporates the gesture’s general characteristics. Notably, the FC layer is kept frozen to reduce computational costs, as the embedding will undergo further optimization in Stage II. A linear fusion of the fused embedding with the original text embeddings is subsequently performed to obtain the **Double Fused Embedding**:

$$e_d = \lambda \times e_t + (1 - \lambda) \times e_f, \quad (2)$$

where  $\lambda$  is the hyperparameter,  $e_d$  is Double Fused Embedding,  $e_t$  is Text Embedding, and  $e_f$  is Fused Embedding. This linear fusion reinforces the text embedding’s unique features, allowing it to leverage the gesture feature without overpowering its own nuances.

Inspired by image editing techniques (Kawar et al. 2023), we proceed to optimize the Double Fused Embedding to better align with the real hand images. Specifically, the stable diffusion model is kept frozen while the Double Fused Embedding is fine-tuned using the loss function defined in Equation 1, resulting in what we term **Optimized Embedding**. The Optimized Embedding, enriched and refined with hand gesture information, establishes a stronger alignment with the corresponding hand images, making it ideal for the subsequent hand image generation in Stage III.

### Stage III: Stable Diffusion Fine-tuning

In the final stage of training, the Optimized Embedding obtained in Stage II is used as a condition to fine-tune the

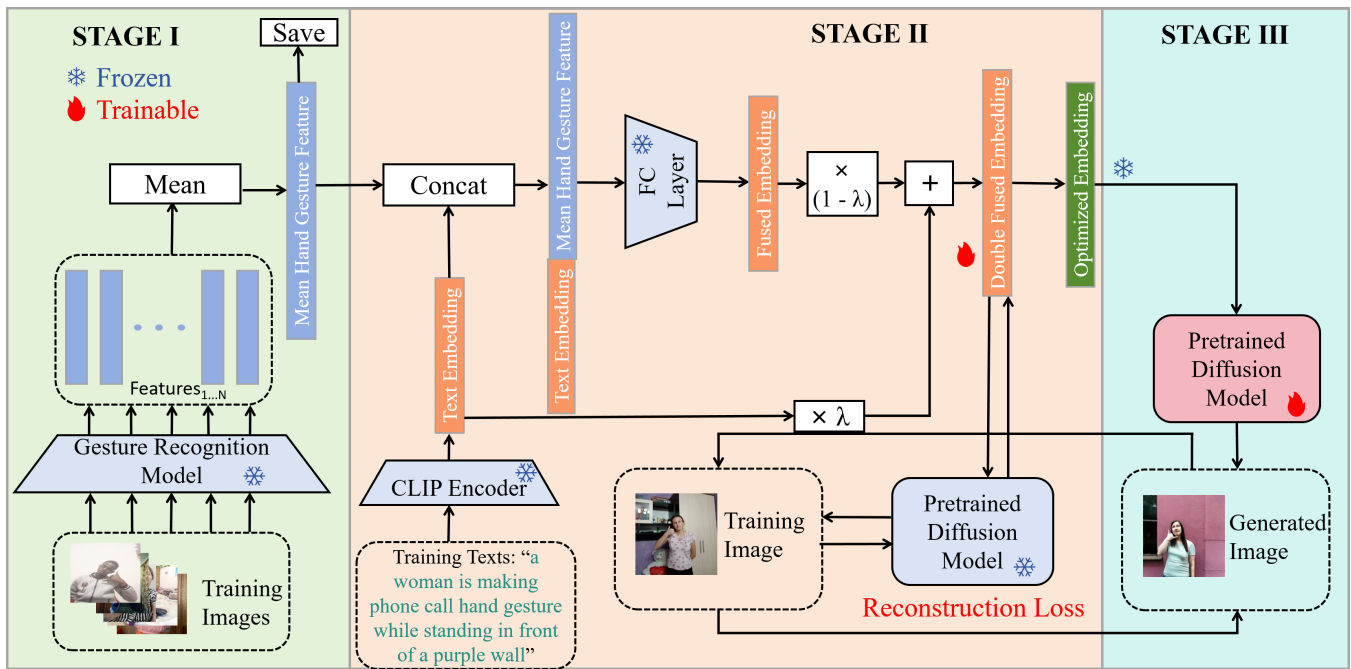


Figure 2: The proposed Hand1000 is designed with a three-stage training process. In Stage I, the primary objective is to compute mean hand gesture feature from images. Stage II builds on this by concatenating the mean hand gesture feature obtained in Stage I with the corresponding text embeddings. These concatenated features are then mapped into a fused embedding, which is further enhanced by linearly fusing it with the original text embedding, resulting in a double-fused embedding. This embedding is optimized using a reconstruction loss through a frozen Stable Diffusion model, ensuring that the final embedding is well-optimized. Stage III involves fine-tuning the Stable Diffusion model for image generation, leveraging the frozen optimized embedding obtained from Stage II.

Stable Diffusion model with the loss function specified in Equation 1. The fine-tuning procedure follows the standard text-to-image diffusion model training, with the only modification being the replacement of the usual text embedding with our Optimized Embedding from Stage II. To preserve the hand gesture information integrated in Stage II, the Optimized Embedding is kept frozen during this stage, in line with image editing works (Kawar et al. 2023).

### Inference

As illustrated in Figure 3, during the inference phase, a textual description containing a hand gesture is used as input. Similar to the training stage, the CLIP text encoder is adopted to obtain the Text Embedding. The Mean Hand Gesture Feature, stored from during the Stage I in the training phase, is concatenated with the Text Embedding. This concatenated feature is then reduced in dimension using a fully connected layer to obtain the Fused Embedding. The Text Embedding and the Fused Embedding are linearly fused to produce the Double Fused Embedding:

$$e_d = \mu \times e_t + (1 - \mu) \times e_f, \quad (3)$$

where  $\mu$  is the hyperparameter,  $e_d$  is Double Fused Embedding,  $e_t$  is Text Embedding, and  $e_f$  is Fused Embedding. The Double Fused Embedding is then used as a condition to input into the trained Stable Diffusion model in Stage III, en-

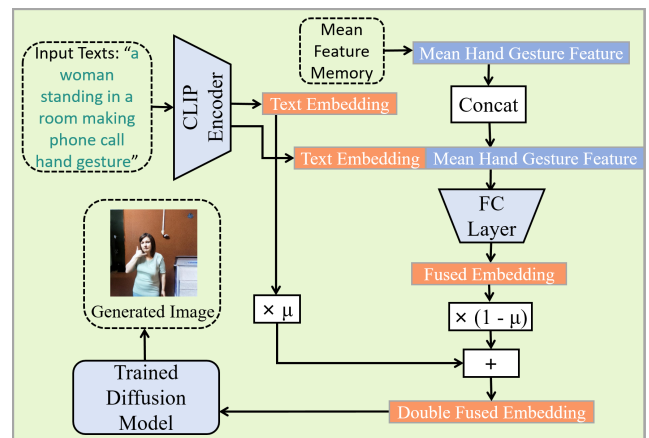


Figure 3: Overview of inference phase. First, text embedding is computed given a textual description as input. Next, the stored mean hand gesture feature is concatenated and fused with text embedding to obtain double-fused embedding. Finally, the double-fused embedding is fed into the trained diffusion model in Stage III to generate hand images.

abling the generation of images that incorporate the specified gesture and maintain a realistic hand appearance.




Figure 4: The dataset construction begins with generating a textual description using an image captioning model (e.g., BLIP2) from image. The textual description, along with gesture labels, is then fed into the LLaMA3 model (Touvron et al. 2023a) to produce a text description enriched with gesture label information.

## Experiments

### Dataset Construction

Due to the lack of publicly available datasets for text-to-hand image generation, we construct a new dataset, as shown in Figure 4. Our dataset is based on HaGRID (Kapitanov et al. 2024), a diverse collection of hand gestures with clear images, varied individuals, clothing, and colors, making it ideal for training and evaluation. We use six hand gestures from HaGRID—“phone call,” “four,” “like,” “mute,” “ok,” and “palm”—with 1,000 images per gesture for training and testing respectively. These gestures cover a range of poses, from high finger separation (e.g., “palm”) to tightly closed fingers (e.g., “mute”) and complex poses (e.g., “ok”), ensuring a comprehensive evaluation of the model.

However, HaGRID (Kapitanov et al. 2024) contains only gesture labels and lacks detailed textual descriptions, making it unsuitable for direct model training. To address this, we firstly use image captioning models to generate descriptions for each image. Secondly, we input gesture labels and textual descriptions into the LLaMA3 model (Touvron et al. 2023a) to enrich the original text with rich hand gesture information. We employ three different image captioning models—BLIP2 (Li et al. 2023b), PaliGemma (Beyer et al. 2024), and VitGpt2 (Mishra et al. 2024)—to generate the textual descriptions. To evaluate the textual descriptions, we use CLIP’s (Radford et al. 2021) text and image encoders to calculate the similarity between descriptions and images, as well as between texts. As shown in Table 1, postprocessing with LLaMA3 (Touvron et al. 2023a) improves text-image similarity, while textual similarity does not increase significantly, demonstrating the effectiveness of this approach. According to Table 1, BLIP2 (Li et al. 2023b) shows the best performance, which is thus chosen as the image captioning model for generating textual descriptions for our dataset.

### Evaluation Metrics

we employ five evaluate metrics to assess the quality of the generated images: Frechet Inception Distance (FID), Kernel Inception Distance (KID), FID-H, KID-H, and HAND-CONF (Heusel et al. 2017; Parmar, Zhang, and Zhu 2022). FID-H involves using an open-source hand gesture recognition model (i.e. Mediapipe hands (Zhang et al. 2020)) to segment the hand regions in the images, and then computing the FID on these segmented patches; KID-H is calcu-



Figure 5: Comparison of images in hand gesture of four fingers up generated by stable diffusion, fine-tuned stable diffusion, stable diffusion enhanced with Imagic (Kawar et al. 2023), fine-tuned stable diffusion enhanced with Imagic (Kawar et al. 2023), stable diffusion enhanced with HandRefiner (Lu et al. 2023) and our Hand1000.

lated similarly. The first four metrics measure the distance between the feature distributions of generated images and real images, with smaller distances indicating that the generated images are closer to real ones and thus of higher quality. HAND-CONF is computed by running a pre-trained hand detector (i.e. Mediapipe hands (Zhang et al. 2020)) and calculating the confidence scores of the detections. Higher confidence scores suggest that the detector is more certain about a region being a hand, indicating better quality of hand generation.

### Implementation Details

For each image in the training set, the first stage of training involves 10 epochs with a learning rate of 1e-3. The sec-

Dataset Processing Method	Image-Caption Consistency $\uparrow$	Caption-Caption Similarity $\downarrow$
PaliGemma	0.237	0.835
VitGpt2	0.254	0.643
<b>BLIP2</b>	<b>0.288</b>	<b>0.600</b>
PaliGemma(post-processed)	0.252	0.715
VitGpt2(post-processed)	0.273	<b>0.617</b>
<b>BLIP2(post-processed)</b>	<b>0.298</b>	0.636

Table 1: Dataset construction quality comparison using different image captioning models and LLaMA3 for postprocessing.

Method	FID $\downarrow$	KID $\downarrow$	FID-H $\downarrow$	KID-H $\downarrow$	HAND-CONF $\uparrow$
Stable Diffusion	114.768	0.079	109.168	0.094	0.895
Stable Diffusion(Fine-tuned)	<b>76.120</b>	<b>0.040</b>	123.254	0.080	0.913
HandRefiner	97.565	0.068	88.997	0.056	0.919
<b>Hand1000</b>	86.169	0.062	<b>78.960</b>	<b>0.053</b>	<b>0.948</b>

Table 2: Performance comparison with Stable Diffusion, fine-tuned Stable Diffusion and HandRefiner.  $\downarrow$  indicates that a lower value is better for that metric, while  $\uparrow$  signifies that a higher value is better.

ond stage consists of 20 epochs with a learning rate of  $1e-6$ . Besides, We use Mediapipe hands (Zhang et al. 2020) as the hand gesture recognition model while the text encoder is the CLIP (Radford et al. 2021) text encoder. The sampling scheme is DDIM (Song, Meng, and Ermon 2020) and the optimizer is Adam (Kingma 2014). We used SDv1.4 for our experiments, which is similar to recent works like HandDiffuser (Narasimhaswamy et al. 2024), HandRefiner (Lu et al. 2023), and Imagic (Kawar et al. 2023).

### Quantitative Performance Comparison

As shown in Table 2, our Hand1000 model demonstrates a significant performance improvement over the Stable Diffusion (Rombach et al. 2022) in terms of all the metrics. Specifically, Hand1000 achieves a 28.599 reduction in FID, a 0.017 reduction in KID, a 30.208 reduction in FID-H, and a 0.041 reduction in KID-H, while its HAND-CONF score is elevated by 0.053. To further validate the effectiveness of our method, we also perform comparison with fine-tuned stable diffusion using the same dataset. Although the fine-tuned Stable Diffusion model demonstrates slightly superior performance in terms of FID and KID compared to our method, several studies (Kynkäänniemi et al. 2019; Borji 2019) have suggested that these two metrics are not entirely reliable. Nevertheless, the quality of hand image generation is significantly inferior, as indicated by FID-H, KID-H, and HAND-CONF. Compared to the fine-tuned stable diffusion model, our method achieves a 44.294 reduction in FID-H, a 0.027 reduction in KID-H, and a 0.035 increase in HAND-CONF, underscoring the effectiveness of our approach. Furthermore, we also conduct experiments on the open-source hand image refinement method, HandRefiner (Lu et al. 2023), and the results demonstrate that our method outperforms HandRefiner (Lu et al. 2023) across all metrics, with an 11.396 reduction in FID, a 0.006 reduction in KID, a 10.037 reduction in FID-H, a 0.003 reduction in KID-H, and a 0.029 increase in HAND-CONF.

### Qualitative Example Comparison

By providing identical text prompts, we conduct a comparative analysis of the results with Stable Diffusion, finetuned Stable Diffusion, using Imagic (Kawar et al. 2023) to edit the results of Stable Diffusion and fine-tuned Stable Diffusion as well as HandRefiner. As illustrated in Figure 5, our method not only generates realistic hand images but also effectively renders other components such as characters and clothing, demonstrating superior performance compared to all the other methods. Furthermore, though Imagic is adopted to edit the image generated by Stable Diffusion, the result led to subtle alterations, such as the lightening of the dress color in the first set of images in the bottom row. In other words, Imagic fails to recover realistic hands and bodies from the generated images, leaving them in a chaotic, deformed, and distorted state. Additional experimental results can be found in the supplementary materials. Besides, HandRefiner fails to interpret gesture information accurately, despite some improvement in hand deformation, and shows no effect when only partial hands are present, as in the rightmost example.

### Ablation Studies

To validate the effectiveness of our design choices, we present a series of ablation experiments, including the ablation of various main components of the method and the ablation of several hyperparameters.

#### Effect of Embedding Optimization and Feature Fusion in Stage II

The ablation results for the feature fusion and embedding optimization in Stage II are shown in Table 3. Since FID-H and KID-H metrics are calculated after extracting the hand region, they provide a more accurate measure of the model’s hand generation capability. With the addition of embedding optimization, FID-H is decreased by 25.003 and KID-H by 0.035. Incorporating Feature Fusion, which integrates features extracted by the hand gesture recognition model with text embeddings, leads to a reduction of 20.319

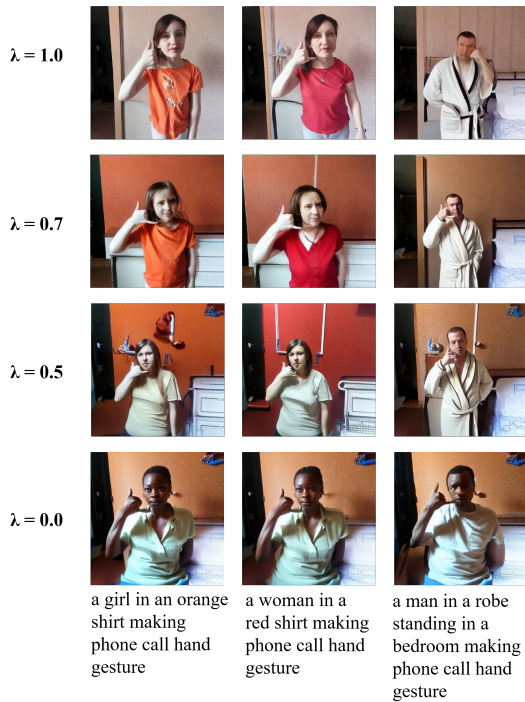


Figure 6: The effect of different values of  $\lambda$  is evaluated. For the “phone call” gesture, the optimal value is found to be  $\lambda = 0.7$ . When  $\lambda$  is too large, the hand regions appear abnormal, whereas when  $\lambda$  is too small, the generated images tend to become homogenized.

in FID-H and 0.048 in KID-H. When both techniques are used together, FID-H is decreased by 29.636 and KID-H by 0.051, demonstrating the superiority of our designs.

**The Impact of Hyperparameter  $\lambda$**   $\lambda$ , as defined in Equation 2, controls the proportion of the Text Embedding in the linear fusion step. A larger  $\lambda$  indicates a higher proportion of the Text Embedding. The optimal value of  $\lambda$  varies across different hand gestures and datasets. Accordingly, the parameter  $\mu$  in Equation 3 must be adjusted during the inference phase to reflect changes in  $\lambda$ . Figure 6 illustrates the effects of varying  $\lambda$  on the gesture “phone call” when  $\mu$  is optimally tuned. It can be observed that when  $\lambda$  is too low, the generated hand regions are accurate and well-formed, but other features, such as the person and clothing, tend to become homogenized. This is because a low  $\lambda$  reduces the diversity of the Double Fused Embedding. Conversely, when  $\lambda$  is too high, the diversity of the images is maintained, but the generated hand regions may exhibit distortions, as the hand feature fusion is not fully utilized. For the “phone call” gesture, an optimal value of  $\lambda = 0.7$  was found, which produces normal hand regions while maintaining diversity and realism in other parts of the image.

**Size of Training Set** To determine the optimal training set size for our method, we conduct ablation experiments using FID-H as the evaluation metric. As shown in Figure 7, the experiments reveal that when the training set size is below

Emb Optimization	Feature Fusion	FID-H↓	KID-H↓
×	×	108.596	0.104
×	✓	88.277	0.056
✓	×	83.593	0.069
✓	✓	<b>78.960</b>	<b>0.053</b>

Table 3: Ablation Study of Embedding Optimization and Feature Fusion in Stage II.

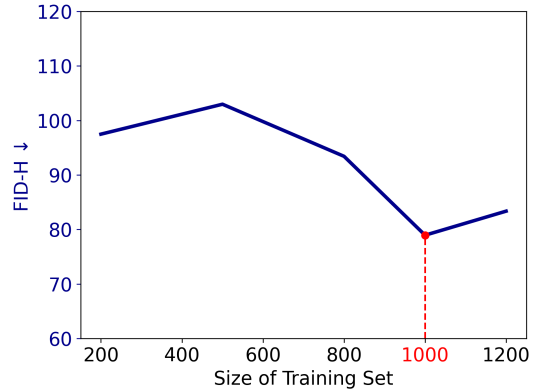


Figure 7: Ablation for the size of the training set.

1,000, the FID-H is relatively high. The FID-H begins to decrease at a training set size of 500, reaching its minimum value at **1,000**. At a training set size of 1,200, the FID-H shows a slight increase, but the magnitude of this increase is small, which may indicate overfitting of the embedding, a known challenge in image-editing methods and worth further study. Therefore, we select 1,000 as the optimal training set size, and ultimately decided to use only 1,000 images for all experiments to ensure both efficiency and performance.

## Conclusion

We have presented a novel method based on the diffusion model to generate realistic hand images from text using only 1,000 images. Through quantitative and qualitative experiments, we demonstrate the superiority and effectiveness of our approach. In particular, we show that it is significant to incorporate hand gesture information and optimize text embeddings for accurate hand image generation. As text-to-image generation continues to evolve, the methods and insights presented in this paper pave the way for more accurate and reliable visual representations, particularly in the challenging domain of human hands. While encouraging, this work is limited by the lack of exploration into simultaneous target gesture generation and hand-object interaction, primarily due to the unavailability of datasets, which will be our future work.

We would also like to gratefully acknowledge Bitdeer.AI for providing the cloud services and computing resources that were essential to this work.

## References

- Abdal, R.; Zhu, P.; Mitra, N. J.; and Wonka, P. 2021. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ToG*, 40(3): 1–21.
- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Bar-Tal, O.; Ofri-Amar, D.; Fridman, R.; Kasten, Y.; and Dekel, T. 2022. Text2live: Text-driven layered image and video editing. In *ECCV*, 707–723. Springer.
- Beyer, L.; Steiner, A.; Pinto, A. S.; Kolesnikov, A.; Wang, X.; Salz, D.; Neumann, M.; Alabdulmohsin, I.; Tschannen, M.; Bugliarello, E.; et al. 2024. PaliGemma: A versatile 3B VLM for transfer. *arXiv preprint arXiv:2407.07726*.
- Borji, A. 2019. Pros and cons of GAN evaluation measures. *Computer vision and image understanding*, 179: 41–65.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *NeurIPS*, 33.
- Chowdhery, A.; Narang, S.; Devlin, J.; Bosma, M.; Mishra, G.; Roberts, A.; Barham, P.; Chung, H. W.; Sutton, C.; Gehrmann, S.; et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240): 1–113.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *NeurIPS*, 34.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *ICML*.
- Härkönen, E.; Hertzmann, A.; Lehtinen, J.; and Paris, S. 2020. Ganspace: Discovering interpretable gan controls. *NeurIPS*, 33.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *NeurIPS*, 33.
- Ho, J.; Saharia, C.; Chan, W.; Fleet, D. J.; Norouzi, M.; and Salimans, T. 2022. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47): 1–33.
- Kang, M.; Zhu, J.-Y.; Zhang, R.; Park, J.; Shechtman, E.; Paris, S.; and Park, T. 2023. Scaling up gans for text-to-image synthesis. In *CVPR*, 10124–10134.
- Kapitanov, A.; Kvanchiani, K.; Nagaev, A.; Kraynov, R.; and Makhliarchuk, A. 2024. HaGRID–HAnd Gesture Recognition Image Dataset. In *WACV*, 4572–4581.
- Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2023. Imagic: Text-based real image editing with diffusion models. In *CVPR*, 6007–6017.
- Kingma, D. 2014. Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Koh, J. Y.; Fried, D.; and Salakhutdinov, R. R. 2024. Generating images with multimodal language models. *NeurIPS*, 36.
- Kynkäänniemi, T.; Karras, T.; Laine, S.; Lehtinen, J.; and Aila, T. 2019. Improved precision and recall metric for assessing generative models. *NeurIPS*, 32.
- Lang, O.; Gandelsman, Y.; Yarom, M.; Wald, Y.; Elidan, G.; Hassidim, A.; Freeman, W. T.; Isola, P.; Globerson, A.; Irani, M.; et al. 2021. Explaining in style: Training a gan to explain a classifier in stylespace. In *ICCV*, 693–702.
- Li, B.; Ma, T.; Zhang, P.; Hua, M.; Liu, W.; He, Q.; and Yi, Z. 2023a. ReGANIE: rectifying GAN inversion errors for accurate real image editing. In *AAAI*, volume 37, 1269–1277.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023b. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*. PMLR.
- Lin, Q.; Zhang, J.; Ong, Y. S.; and Zhang, M. 2024. Make Me Happier: Evoking Emotions Through Image Diffusion Models. *arXiv preprint arXiv:2403.08255*.
- Lu, W.; Xu, Y.; Zhang, J.; Wang, C.; and Tao, D. 2023. Handrefiner: Refining malformed hands in generated images by diffusion-based conditional inpainting. *arXiv preprint arXiv:2311.17957*.
- Ma, T.; Li, B.; Liu, W.; Hua, M.; Dong, J.; and Tan, T. 2023. CF2T-GAN: cross-domain feature fusion transformer for exemplar-based image translation. In *AAAI*, volume 37, 1887–1895.
- Mishra, S.; Seth, S.; Jain, S.; Pant, V.; Parikh, J.; Jain, R.; and Islam, S. M. 2024. Image Caption Generation using Vision Transformer and GPT Architecture. In *InCACCT*, 1–6. IEEE.
- Narasimhaswamy, S.; Bhattacharya, U.; Chen, X.; Dasgupta, I.; Mitra, S.; and Hoai, M. 2024. Handdiffuser: Text-to-image generation with realistic hand appearances. In *CVPR*, 2468–2479.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *ICML*. PMLR.
- Parmar, G.; Zhang, R.; and Zhu, J.-Y. 2022. On aliased resizing and surprising subtleties in gan evaluation. In *CVPR*, 11410–11420.
- Patashnik, O.; Wu, Z.; Shechtman, E.; Cohen-Or, D.; and Lischinski, D. 2021. Styleclip: Text-driven manipulation of stylegan imagery. In *ICCV*, 2085–2094.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *ICCV*, 4195–4205.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*. PMLR.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *ICML*. PMLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*, 10684–10695.

- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 22500–22510.
- Saharia, C.; Chan, W.; Chang, H.; Lee, C.; Ho, J.; Salimans, T.; Fleet, D.; and Norouzi, M. 2022a. Palette: Image-to-image diffusion models. In *SIGGRAPH*, 1–10.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022b. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*, 35.
- Saharia, C.; Ho, J.; Chan, W.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2022c. Image super-resolution via iterative refinement. *TPAMI*, 45(4): 4713–4726.
- Sauer, A.; Lorenz, D.; Blattmann, A.; and Rombach, R. 2023. Adversarial diffusion distillation. *arXiv preprint arXiv:2311.17042*.
- Shen, Y.; Gu, J.; Tang, X.; and Zhou, B. 2020. Interpreting the latent space of gans for semantic face editing. In *CVPR*, 9243–9252.
- Shen, Y.; and Zhou, B. 2021. Closed-form factorization of latent semantics in gans. In *CVPR*, 1532–1540.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*. PMLR.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *NeurIPS*, 32.
- Song, Y.; and Ermon, S. 2020. Improved techniques for training score-based generative models. *NeurIPS*, 33.
- Tao, M.; Tang, H.; Wu, F.; Jing, X.-Y.; Bao, B.-K.; and Xu, C. 2022. Df-gan: A simple and effective baseline for text-to-image synthesis. In *CVPR*, 16515–16525.
- Team, G.; Anil, R.; Borgeaud, S.; Wu, Y.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023a. Open and efficient foundation language models. *Preprint at arXiv. <https://doi.org/10.48550/arXiv.2302>*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023b. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *NeurIPS*, 30.
- Wang, Y.; Pang, M.; Chen, S.; and Rao, H. 2024. Consistency-gan: Training gans with consistency model. In *AAAI*, volume 38, 15743–15751.
- Xia, M.; Shu, Y.; Wang, Y.; Lai, Y.-K.; Li, Q.; Wan, P.; Wang, Z.; and Liu, Y.-J. 2023. FEditNet: few-shot editing of latent semantics in GAN spaces. In *AAAI*, volume 37, 2919–2927.
- Xu, T.; Zhang, P.; Huang, Q.; Zhang, H.; Gan, Z.; Huang, X.; and He, X. 2018. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *CVPR*, 1316–1324.
- Yang, Z.; Peng, D.; Kong, Y.; Zhang, Y.; Yao, C.; and Jin, L. 2024. Fontdiffuser: One-shot font generation via denoising diffusion with multi-scale content aggregation and style contrastive learning. In *AAAI*, volume 38, 6603–6611.
- Ye, H.; Yang, X.; Takac, M.; Sunderraman, R.; and Ji, S. 2021. Improving text-to-image synthesis using contrastive learning. *arXiv preprint arXiv:2107.02423*.
- Ye, Y.; Xu, K.; Huang, Y.; Yi, R.; and Cai, Z. 2024. DiffusionEdge: Diffusion Probabilistic Model for Crisp Edge Detection. In *AAAI*, volume 38, 6675–6683.
- Yu, Z.; Li, H.; Fu, F.; Miao, X.; and Cui, B. 2024. Accelerating Text-to-Image Editing via Cache-Enabled Sparse Diffusion Inference. In *AAAI*, volume 38, 16605–16613.
- Zhang, F.; Bazarevsky, V.; Vakunov, A.; Tkachenka, A.; Sung, G.; Chang, C.-L.; and Grundmann, M. 2020. Mediapipe hands: On-device real-time hand tracking. *arXiv preprint arXiv:2006.10214*.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *ICCV*, 3836–3847.
- Zhu, B.; and Ngo, C.-W. 2020. CookGAN: Causality based text-to-image synthesis. In *CVPR*, 5519–5527.
- Zhu, B.; Ngo, C.-W.; Chen, J.; and Hao, Y. 2019a. R2gan: Cross-modal recipe retrieval with generative adversarial network. In *CVPR*, 11477–11486.
- Zhu, M.; Pan, P.; Chen, W.; and Yang, Y. 2019b. Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis. In *CVPR*, 5802–5810.