

HyperLoRA: Parameter-Efficient Adaptive Generation for Portrait Synthesis

Mengtian Li* Jinshu Chen* Wanquan Feng*† Bingchuan Li Fei Dai
Songtao Zhao Qian He
Intelligent Creation, ByteDance



Figure 1. We propose HyperLoRA, a parameter-efficient adaptive method for portrait synthesis. Given an input face image, HyperLoRA generates personalized LoRA weights without online fine-tuning. Due to the natural interpolability of LoRA, it is easy to support multiple inputs by simple averaging. Leveraging the generated LoRA, we can create personalized portraits with high photorealism and fidelity.

Abstract

Personalized portrait synthesis, essential in domains like social entertainment, has recently made significant progress. Person-wise fine-tuning based methods, such as LoRA and DreamBooth, can produce photorealistic outputs but need training on individual samples, consuming time and resources and posing an unstable risk. Adapter based techniques such as IP-Adapter freeze the foundational model parameters and employ a plug-in architecture to enable zero-shot inference, but they often exhibit a lack of naturalness and authenticity, which are not to be overlooked in portrait synthesis tasks. In this paper, we introduce a parameter-efficient adaptive generation method,

namely **HyperLoRA**, that uses an adaptive plug-in network to generate LoRA weights, merging the superior performance of LoRA with the zero-shot capability of adapter scheme. Through our carefully designed network structure and training strategy, we achieve zero-shot personalized portrait generation (supporting both single and multiple image inputs) with high photorealism, fidelity, and editability.

1. Introduction

The personalized portrait synthesis task [16, 47, 48] aims to generate images sharing the same identity with what is contained in the given single or multiple portrait images. By feeding the identity information into specific generative model, the task enables the users to create diverse, high-quality and creative personalized portraits, thus hold-

* Equal contribution † Corresponding author

ing substantial importance across various domains, such as entertainment and social media.

One key challenge of the personalized portrait synthesis task is the **fidelity**: the expected generated images should retain the given identity to the greatest extent and avoid irrelevant outputs. Another challenge lies in **editability**: while preserving the given identity, the portrait synthesis model should maintain the abilities of flexibly editing the other contents of the generated images, for example, the composition of the image, the positioning and the clothing of the character. Additionally, the inference **speed** is crucial, as fast inference is friendly to not only the computational costs but also the user experience.

Considering the challenges mentioned above, most of the current solutions have different shortcomings. Existing personalized image synthesis methods can generally be divided into tuning-based and tuning-free approaches. Tuning-based methods, such as LoRA (Low-Rank Adaptation) [20] and DreamBooth [38], update model parameters based on specific user-provided data, which can accurately comprehend and preserve the input identity information. Theoretically, these methods are mainly designed for few-shot cases, and customize different models for each groups of provided data with different identities. Thus, disadvantages like high computing costs, unstable training, and lengthy training processes make such methods **bloated** when dealing with various provided identities.

Tuning-free methods, such as IP-Adapter [48], typically utilize an additional plug-in structure to encode user identity information. By feeding such identity-related tokens into the frozen foundational model, these methods make influences on the inference process, and ultimately yielding results that retain the given identities. Such additional adapter structure is trained with large portrait datasets in advance, and frozen when inference. Therefore, these methods get rid of the online training, and are convenient for the actual uses. However, for these methods, the interaction between encoded identity information and the foundational model relies on **a group of fixed extra cross-attention modules**, which might degrade the quality of generated images, leading to unnatural and non-photorealistic facial features, and fail to recover fine-grained details. In summary, the previous methods fail to simultaneously strike a balance between the fidelity, the editability, and the speed.

In this paper, we introduce a parameter-efficient adaptive generation method, namely **HyperLoRA**, which achieves great identity fidelity, high editability and fast inference speed. HyperLoRA uses an adaptive plug-in network to encode identity images, which follows the design of tuning-free methods like IP-Adapter. However, **unlike such methods which utilize the encoded tokens for the attention computation, we directly generate the LoRA weights with the plug-in network**. This strategy not only maintains the zero-

shot capability of the adapter strategy, but also inherits the representation and high-quality generation effect of the LoRA strategy which adjusts the global parameters without real fine-tuning process.

Considering the substantial number of parameters in LoRA, to reduce the training difficulty, **we initially project the LoRA parameters onto a linear LoRA parameter space**, rather than directly training in the original parameter space. In this way, the generated LoRA weights can be represented as a linear combination of a trainable basis. **We employ a perceiver resampler (as in IP-Adapter [48]) to convert the input identity images into combination coefficients for the basis**. To encode the input images, **we utilize a CLIP ViT [36] and an AntelopeV2 encoder [22] to extract pixel-level and identity-sensitive features**. These feature extractors mentioned above **stay frozen during all the training process**. To avoid overfitting, we **decompose** LoRA at the parameter level into an ID-LoRA part and a Base-LoRA part, thereby the ID and ID-irrelevant information are fitted into different parts correspondingly. The whole training pipeline runs in a multi-stage way: **only the Base-LoRA part is trained in the first stage as a warm-up start; and in the second stage, we train the ID-LoRA with CLIP ViT features; in the final stage, we only finetune the attention modules related to ID embeddings**. Besides the case where only a single image is given, our proposed method can also handle cases that there are multiple inputs by simply interpolating the predicted LoRA weights. Our contributions can be summarized as follows:

- We propose HyperLoRA, a parameter-efficient adaptive method for the personalized portrait synthesis task, **merging the superior performance of the LoRA strategy with the zero-shot capability of the adapter strategy**.
- We **construct a low-dimension linear LoRA space**, which reduces the training difficulty of the plug-in network while ensuring the identity preservation.
- We design a strategy that **decomposes LoRA weights at the parameter level into a base part and an ID part**, thereby effectively decoupling the identity information restored in the given images and significantly reducing overfitting.
- We conduct various experiments to demonstrate the capacities of our framework. Experimental results show that our method achieves both high fidelity and editability.

2. Related Work

2.1. Text-to-Image Synthesis

Text-to-image (T2I) synthesis has been rapidly developed for several years. The task aims to generate high-quality images which are highly correlated with given texts or instructions. Some early works [23, 45, 50] employ transformer [46] or GAN [11] architectures to solve this task, while the performance of the solutions is not that satisfy-

ing owing to the limited capacities of their pipelines. Recently, rapid progress has been made by the methods adopting diffusion pipelines [6, 18, 19, 25, 43, 44]. With a step-wise noising-denoising design, the generated results of diffusion pipelines have far higher qualities than other self-supervised methods, especially on the datasets with huge intra-domain variance [5, 41]. In this paper, we follow a diffusion pipeline [37] to build up our pipeline to ensure the qualities and the varieties of the outputs.

2.2. Tuning-based Personalized Portrait Synthesis

Tuning-based methods finetune models based on user-provided data to comprehend user identity information. Dreambooth [38] and its derivative versions [7, 33, 39] treat the entire T2I model trainable, and utilize certain special captions to represent the given identities. LoRA [20] and its variants [8, 14, 15, 29, 49, 52] represent another category of tuning-based methods. Unlike Dreambooth, LoRA greatly enhances the flexibility of the fine-tuning scheme by creating several trainable low-rank matrices within the T2I model along with much fewer parameters than the whole model. Nevertheless, a common issue with all the aforementioned methods is that they lack the ability to handle zero-shot cases. Each fine-tuned model matches only a single customized task, thus the fine-tuning incurs a high burden in terms of time, computational resources, and storage resources. In this work, our proposed HyperLoRA does not rely on online training of user-provided data, thereby circumventing the resource consumption problem brought about by the training process.

2.3. Tuning-free Personalized Portrait Synthesis

Tuning-free methods, such as [28, 31, 32, 47, 48], pre-design a plug-in structure containing few parameters that can be inserted into the T2I base model. After pre-training and freezing such parameters on a large-scale dataset, tuning-free methods utilize such insertable structure to extract the identity information features from the given inputs, and inject them into the base model during inference. The online tuning is not necessary for these methods, thereby they can save resource costs and time. Concurrently, the design of the detachable parameters is also rather flexible. However, these methods merely learn the mapping relationship between the input image and the encoded tokens, which are utilized by a group of additional cross-attention modules, and the outputs are simply blended with the original ones. Our proposed HyperLoRA learns the mapping relationship between the input image and the network parameters, which is more essential. The implicit information contained in the network parameters is far more abundant and friendly to the base model than that in tokens. Compared with the methods mentioned above, our method has better adaptability, higher nonlinearity, and better fidelity.

3. Method

In this section, we introduce the details of our method, overview of which is illustrated in Fig. 2. In Sec. 3.1, we introduce the low-dimensional linear LoRA space. In Sec. 3.2, we introduce the network structure. In Sec. 3.3, we introduce the designed multi-stage training process.

3.1. Low-Dimensional Linear LoRA Space

The number of parameters in LoRA is typically large, which makes it difficult to generate the LoRA weights directly. Recognizing the interpolative nature of LoRA, we are inspired to exploit its linearity, and map the LoRA parameters onto a low-dimensional linear space. Note that our discussion below and in the rest of the paper is based on the implementation from [40], which is the most widely acknowledged way. Each layer of LoRA includes 8 matrices, which consist of the down and up matrices for query, key, value, and output. We project each matrix onto a K -dimensional basis. Mark the number of LoRA layers as N , then the overall degrees of freedom for the entire LoRA can be calculated as $N \times 8K$, which is significantly smaller than the origin parameter number in most cases. In our paper, we use $K = 128$, and the number of parameters after projection becomes approximately only 1.2% (0.14M vs. 11.6M when the LoRA rank is 8) of the original one. To decouple the given specific identity information from the irrelevant features, we design to split a LoRA matrix into a base part and an ID part. We denote the basis of the ID part as $\{\mathbf{M}_{id}^k\}_{k=1}^K$, the combination coefficients as $\{\alpha_k\}_{k=1}^K$. Further, we denote the ID part as:

$$\mathbf{M}_{id} = \sum_{k=1}^K \alpha_k \cdot \mathbf{M}_{id}^k. \quad (1)$$

Similarly, the base part, denoted as \mathbf{M}_{base} , follows the same format as \mathbf{M}_{id} :

$$\mathbf{M}_{base} = \sum_{k=1}^K \beta_k \cdot \mathbf{M}_{base}^k. \quad (2)$$

Although the mathematical format of Base-LoRA is similar to ID-LoRA, it possesses a different significance. Specifically, we configure \mathbf{M}_{base} to contain identity-irrelevant contents across all the training samples, such as the background and clothing, essentially capturing information from the input image that should not be included in \mathbf{M}_{id} . Explicitly modeling Base-LoRA can help ID-LoRA learn more essential identity-related information and ignore the irrelevant rest. In summary, we represent a HyperLoRA parameter matrix \mathbf{M} as:

$$\mathbf{M} = \mathbf{M}_{base} + \mathbf{M}_{id} = \sum_{k=1}^K \beta_k \cdot \mathbf{M}_{base}^k + \sum_{k=1}^K \alpha_k \cdot \mathbf{M}_{id}^k. \quad (3)$$

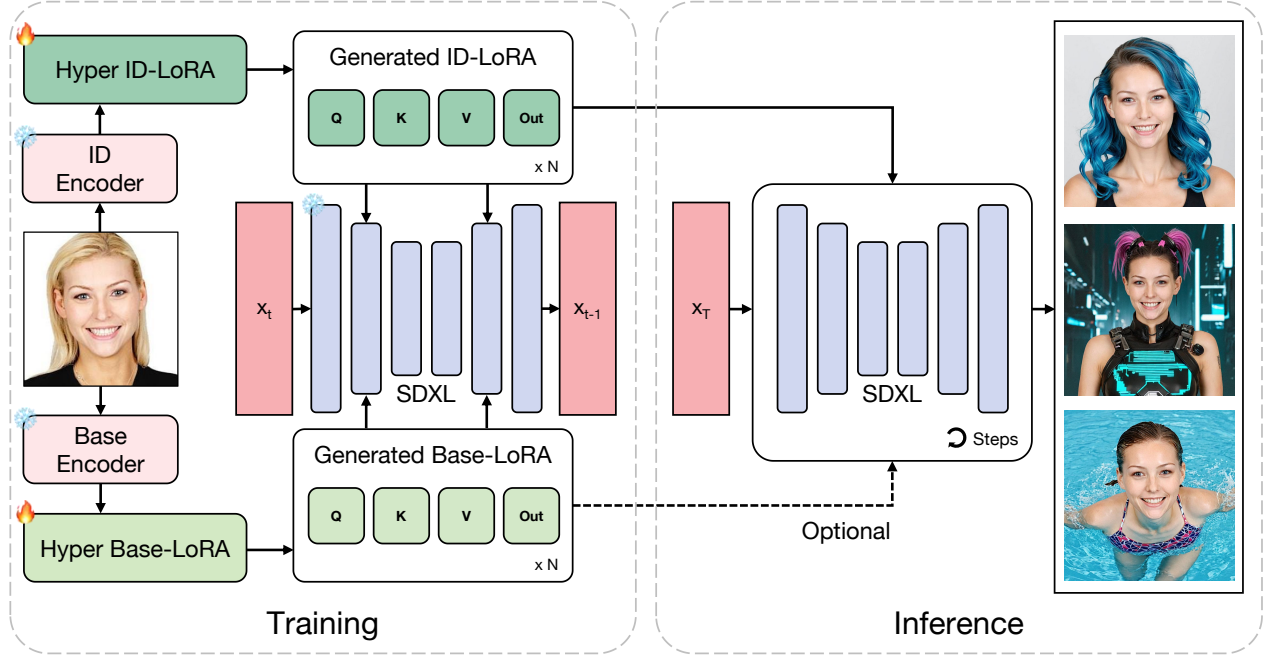


Figure 2. Overview of HyperLoRA. We explicitly decompose the HyperLoRA into a Hyper ID-LoRA and a Hyper Base-LoRA. The former is designed to learn ID information while the latter is expected to fit others, e.g. background and clothing. Such a design helps to prevent irreverent features leaking to ID-LoRA. During the training, we fix the weights of the pretrained SDXL base model and encoders, allowing only HyperLoRA modules updated by Backpropagation. At the inference stage, the Hyper ID-LoRA integrated into SDXL generates personalized images while the Hyper Base-LoRA is optional.

It is noteworthy that although the degrees of freedom for the LoRA parameters are significantly compressed, they still retain the capacity to represent identity information, as Fig. 3 shows. With our designed projection method and the setting of $K = 128$, the generated results in the illustration of overfitting a face dataset suggest that, it is still capable of reconstructing the identity of the reference image for the compressed LoRA module without obvious degradation.

3.2. Network Structure

In this section, we introduce our network structure, which is shown in Fig. 4. The whole network can be divided into several main components: the image encoder, the preceiver resampler and the LoRA combination.

Image Encoder. The input image is passed into two image encoders, including a CLIP [36] ViT and an AntelopeV2 [22] encoder. The features extracted by the former are dense and contain strong structural information, while the features extracted by the latter are more abstract identity information. To explicitly distinguish the training processes of the base part and the id part, we deal with the input image of the CLIP in two different ways: for the training of M_{base} , we first crop the face region of an original image, and inside the cropped image, we blur the facial features and contour, preventing the base part reconstructing the specific face; for

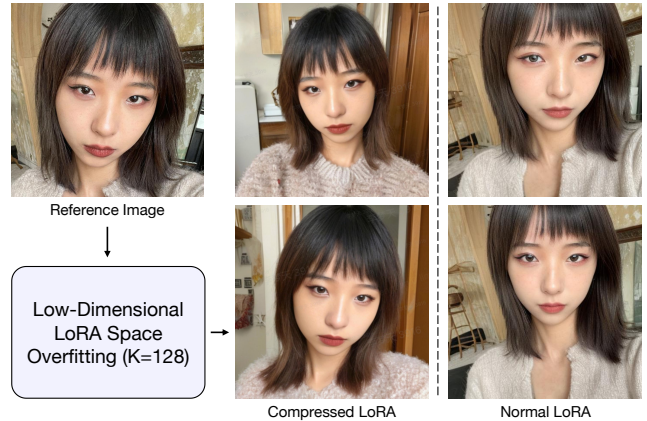


Figure 3. The identity reconstruction ability on the low-dimensional linear LoRA space. In this example, we project LoRA parameters onto 128-dim basis and train on a face dataset (about 400K samples). Compared to normal LoRA, our compressed LoRA can also maintain the identity of the reference image well.

the training of M_{id} , we directly feed the cropped face image into CLIP ViT. Meanwhile, it also serves as the input of AntelopeV2 encoder to extract ID embedding. We utilize a simple ID Projector, consisting of a Linear layer followed by a Reshape and a LayerNorm [2], to transform the ID em-

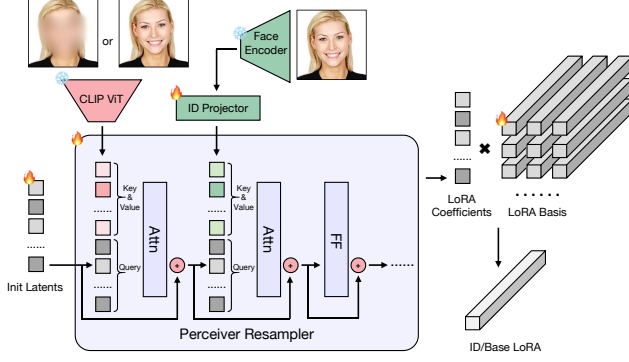


Figure 4. Network structure of HyperLoRA. We apply a perceiver resampler to convert the image features into a group of LoRA coefficients, thereby generating the whole LoRA by multiplied with LoRA basis. Two independent perceiver resamplers are instantiated for Hyper Base and ID LoRAs. Note that the second attention block (interacting with green tokens from ID Projector) is absent in Hyper Base-LoRA.

bedding into some ID tokens. The details of the approaches mentioned above can be found in Fig. 4.

Perceiver Resampler. The perceiver resampler [1, 48] employs a transformer backbone and incorporates image tokens into the computation through an attention mechanism. Currently, we apply a 4-layer perceiver resampler, each layer of which consists of two attention modules interacting with the tokens from CLIP ViT and ID Projector. Note that we skip the second attention module when training M_{base} . We obtain the output of the perceiver resampler, the combination coefficients $\{\alpha_k\}_{k=1}^K$ and $\{\beta_k\}_{k=1}^K$, which would be used to compute the output LoRA weights subsequently.

LoRA Combination. Both the Base-LoRA basis $\{M_{base}^k\}_{k=1}^K$ and the ID-LoRA basis $\{M_{id}^k\}_{k=1}^K$ are trainable variables. They are optimized during the training phase and fixed as constants during inference. Once the combination coefficients $\{\alpha_k\}_{k=1}^K$ and $\{\beta_k\}_{k=1}^K$ are predicted, the network can compute the entire LoRA by Eq. 3. It is noteworthy that the combination way can be adjusted flexibly to achieve various effects, for example, if the editability is a priority, the weight of M_{base} in Eq. 3 can be reduced to drop the information unrelated to the input identity.

3.3. Training

In this section, we introduce our multi-stage training process. Our approach begins with training the Base-LoRA as an initial warm-up. Subsequently, we expand the training to ID-LoRA with CLIP ViT. Finally, we fine-tune the model to enhance its identity preserving capability by using AntelopeV2 face encoder. For all stages, we employ the same loss term as in DDPM [18].

Base-LoRA Warm-up. Training the ID-LoRA directly from an open domain base model is quite challenging, as

it tends to inadvertently incorporate extraneous details such as composition, background, and attire. Therefore, as described in Sec.3.1, we train non-facial information into the Base-LoRA, which serves as a warm-up in our training process. In this stage, we blur the facial information of the original input image before passing it through the CLIP ViT, and bypass the AntelopeV2 face encoder.

ID-LoRA with CLIP. After the Base-LoRA warm-up, we add ID-LoRA into the training process. In this stage, we add the AntelopeV2 branch besides the existed CLIP branch, so that the abstract identity information can be efficiently extracted. We observe that training with CLIP features converges quickly but tends to overfit structural details, whereas training with ID embedding converges more slowly. Consequently, in this stage we stop training with CLIP features as long as it reaches the approximate convergence. To stabilize the training state in the early iterations, we initialize the output layers of the attention modules that are fused with ID embeddings to zero, allowing only the CLIP features to guide the early generation. To enable ID-LoRA to effectively represent the identity in the presence of trigger words (e.g., 3 rare words *fcsks fxxhks fhyks* in this paper), while ensuring that the results remain unaffected in their absence, each iteration of this stage is set to one of the following situations: 1). The given prompt contains trigger words, both Base-LoRA and ID-LoRA are enabled, and the reconstruction target is the target image; 2). The given prompt does not contain trigger words, only ID-LoRA is enabled, and the resulting noise prediction should align with that of the original foundational model; 3). The given prompt does not contain trigger words, both Base-LoRA and ID-LoRA are enabled, and the resulting noise prediction should be consistent with that when only Base-LoRA is enabled. This approach minimizes disruption to the base model parameters and further concentrates the parameters on those related to identity information.

ID-Embedding Fine-tune. After the previous training stages achieve rough convergence, we fix the weights of the CLIP-related attention modules, and continue to fine-tune the attention modules related to ID embedding (the second attention in Fig. 4). The training loss of this stage is exactly the same as in the second stage.

4. Experiments

4.1. Implementation Details

To train our HyperLoRA model, we collect a subset from LAION-2B [42] which comprises 4.4 million portrait images. In detail, InsightFace [12] is leveraged to detect faces and we filter out those non-single portrait images. Meanwhile, we generate a caption for each image in our dataset using BLIP-2 [27] to enable Text-to-Image training. During the training, we randomly crop images around the face for




Figure 5. Adapter and LoRA show different tolerances to CFG. The first row: images generated by InstantID where higher CFG leads to oversaturation. The second row: HyperLoRA always yield reasonable portrait images for CFG ranged from 3 to 7.



Figure 6. Inference with Base-LoRA improves face fidelity but limits the layout of generated images more similar to the cropped region for training.

data augmentation and focus more on the face region. Considering memory consumption, we currently set the LoRA rank of ID-LoRA and Base-LoRA to 8 and 4 respectively, which costs about 60G memory when batch size is 1. We further apply the 8-step gradient accumulation to simulate a larger batch size. We train the model based on SDXL-Base-1.0 [35] with AdamW [30] optimizer and learning rate 0.00001. In this 3-stage training process, we optimize the Base-LoRA, ID-LoRA (CLIP) and ID-LoRA (ID embedding) for 20K, 15K and 55K iterations in turn, and set the probability of three situations to 0.9, 0.05 and 0.05 for ID-LoRA training. We implement our method with PyTorch and Diffusers. The total training process is performed on 16 NVIDIA A100 GPUs for about 10 days.

4.2. Settings

We build workflows in ComfyUI [4] to evaluate our HyperLoRA and three SOTA ID preservation methods, i.e. InstantID [47], IP-Adapter [48] (specifically, IP-Adapter-FaceID-Portrait) and PuLID [13]. To ensure a fair comparison, we apply the same seed, prompt and sampler across all methods. Note that we additionally concatenate the lora trigger words as a prefix to the prompt when testing HyperLoRA, and set different Classifier-Free Guidance [17] (CFG) scales

to generate reasonable images. As shown in Fig. 5, Adapter-based methods usually require a lower CFG scale while HyperLoRA allows a wide range of CFG scales. Although these methods are trained on SDXL-Base-1.0, they show sufficient generalization to other SDXL models. Thus, we use LEOSAM [26], an open-source SDXL model for high quality portraits, as the base model for evaluations, which typically improve image aesthetics. In quantitative comparison, we utilize a recent open-source test set, Unsplash-50 [9], which covers diverse genders, ages and races.

Table 1. Quantitative comparison.

	Fidelity		Editability	
	CLIP-I \uparrow	ID Sim. \uparrow	CLIP-I \uparrow	CLIP-T \uparrow
IP-Adapter	0.764	0.566	0.725	0.244
InstantID	0.734	0.681	0.688	0.237
PuLID	0.771	0.613	0.805	0.259
Arc2Face	0.786	0.643	-	-
Ours (Full)	0.853	0.678	0.710	0.243
Ours (ID)	0.831	0.625	0.748	0.252

4.3. Comparisons

Quantitative Comparison. We evaluate fidelity and editability as the quantitative metrics for comparison across these methods. For fidelity, we leverage the cosine of CLIP image features (CLIP-I) to measure face fidelity, and the cosine of ID embeddings (ID Sim.) to measure face ID similarity. We generate four images with prompt headshot of a girl/boy, high definition, 4k for each sample in our test set using each method, and compute the average of them. For editability, we utilize CLIP-I to measure the similarity of images generated with and without ID plug-in networks, and the cosine of CLIP image and text features (CLIP-T) to measure image-text similarity. We collect 10 prompts which cover hair, accessory, clothing and background, and evaluate all combinations of prompts and ID images. The specific CLIP model we used is CLIP-L/14 which is different from the text encoder in SDXL, and CurricularFace [21] (not used as ID encoder in all methods) is adopted to extract ID embeddings. We further incorporate Arc2Face [34], a human face base model, into comparison, which is trained using a fixed caption and a cropped region of faces, thus restricting its editability. Ours (Full) in Table. 1 means inference with both Base-LoRA (weight is set to 0.4) and ID-LoRA, while Ours (ID) represents only using ID-LoRA. As shown in Table. 1 and Fig. 6, Base-LoRA helps to improve fidelity but limits editability. Although our design decouples the image features into different LoRAs, it's hard to avoid leaking mutually. Thus, we can adjust the weight of Base-LoRA to adapt to different application scenarios. Our HyperLoRA (Full and ID) achieve the best and second-best face fidelity while InstantID shows superiority



Figure 7. Qualitative comparison. From top to bottom, we generate images with the scene prompts: **white shirt**, **wolf ears**, **wavy hair and brown sweater**, and **play with snowman**.

in face ID similarity but lower face fidelity. Both these metrics should be considered together to evaluate fidelity, since the face ID similarity is more abstract and face fidelity reflects more details. This is evidenced in the following visual results in Fig. 7. In the aspect of editability, PuLID possesses a clear advantage, the key contribution of which is a well-designed loss function to reduce intrusion into the base model, thereby improving editability. Despite HyperLoRA (ID) inferior to PuLID, it surpasses other methods, meanwhile has good fidelity. It is noteworthy that there is a trade-off between fidelity and editability, and our HyperLoRA reaches a proper balance.

Qualitative Comparison. Fig. 7 shows the visual comparisons across IP-Adapter, InstantID, PuLID and our HyperLoRA (only ID-LoRA). The skin of portraits generated by IP-Adapter and InstantID has apparent AI-generated texture, which is a little oversaturation and far from photorealism. It is a common shortcoming of Adapter-based meth-

ods. PuLID improves this problem by weakening the intrusion to base model, outperforming IP-Adapter and InstantID but still suffering from blurring and lack of details. In contrast, LoRA directly modifies the base model weights instead of introducing extra attention modules, usually generating highly detailed and photorealistic images. Our HyperLoRA naturally inherits this advantage from LoRA. Furthermore, the manner in which LoRA injects ID features exhibits higher nonlinearity. Thus, our method achieves better face fidelity, capturing fine-grained details, such as pupil color. The layout of generated images of our HyperLoRA and PuLID is more similar compared with other methods, which is because of the higher CLIP-I and CLIP-T of editability. In comparison, IP-Adapter and InstantID have inferior editability, resulting in unreasonable layout (the 4th row of IP-Adapter) and failure to follow the prompt (the 3rd row of InstantID).



Figure 8. Impact of Hyper Base-LoRA. The first row: the model trained with Hyper Base-LoRA generates correct clothing and background. The second row: the model trained without it takes many irrelevant features into the generated images.



Figure 9. Comparison between different encoders. The first row: our CLIP + ID embeddings training strategy successfully generates glasses and mask. The second row: training with only CLIP fails to follow the prompts.

4.4. Ablation Study

No Hyper Base-LoRA. In Fig. 8, we demonstrate the effectiveness of Hyper Base-LoRA. As aforementioned, our Hyper Base-LoRA helps to decouple ID information from input images. This is evidenced in the second row of Fig. 8, where the model trained without Hyper Base-LoRA fails to properly reflect the prompts dress, spacesuit and beach. As expected, the background and clothing are more editable if trained with Hyper Base-LoRA, since it prevents the ID-irrelevant features leaking to ID-LoRA.

Only CLIP Encoder. Our multi-stage training strategy is crucial for successfully training ID-LoRA. Especially, we



Figure 10. Multiple input images contribute to a higher ID similarity and a more realistic and aesthetic synthesized portrait.

leverage CLIP ViT features only during the early iterations of training and then switch to ID embeddings to learn more details. Fig. 9 shows that the problem arises when only adopting CLIP as encoder, where we see that the facial elements are challenging to modify. In contrast, predominantly learning facial features with ID embeddings relaxes the structural constraints on the face.

Multiple Input Images. Thanks to the natural interpolability of LoRA and our low-dimensional linear LoRA space (as discussed in Sec. 3.1), it is straightforward to support portrait synthesis with multiple input images by averaging the LoRA coefficients generated from each single image. Fig. 10 illustrates that multiple ID images help HyperLoRA to identify more stable and pure ID features, thereby improving both ID fidelity and image quality.

5. Conclusion

We have proposed HyperLoRA, the first zero-shot personalized portrait generation method based on LoRA, which is trained in an end-to-end manner. With the advantage of LoRA, our method modifies the base model weights to inject ID information rather than introducing new cross attention as adapter-based methods do, exhibiting higher nonlinearity, thereby yielding more detailed and realistic portrait images. To make HyperLoRA training feasible, we represent the whole LoRA parameters with a linear combination of a group of LoRA basis, which reduces the number of HyperLoRA parameters magnificently. We further decompose the LoRA into two parts explicitly. With our well-designed training strategy, the ID part tends to learn facial features while other aspects, like background and clothing, are fitted into the base part, thereby reducing the risk of overfitting and improving editability. It's recognized that only a rank 8 HyperLoRA is allowed to train when considering the memory limitation of current GPU. Nevertheless, rank 8 is still significantly smaller than what is frequently used in normal LoRA. An advanced LoRA parameters compression method, such as LoRA autoencoder, may enable larger rank HyperLoRA training, unlocking more powerful capability.

HyperLoRA: Parameter-Efficient Adaptive Generation for Portrait Synthesis

Supplementary Material

A. Appendix

A.1. Inference Speed

We evaluate the inference speed of IP-Adapter [48], InstantID [47], PuLID [13] and our HyperLoRA on a single NVIDIA V100 GPU. For each method, the same 512×512 image is adopted as the input of ID plug-in network, and we generate a 1024×1024 personalized image using the DPM-2M Karras sampler with 16 denoising steps. We repeat the same inference process four times and record the average inference time in Table. 2. The complete inference process consists of two parts: one is the preprocess and another is the real inference of base model. In the preprocess stage, we crop the input image, obtaining ID embedding and predicting the ID tokens (for HyperLoRA, generating the LoRA weights and merging into base model). The sequence length of LoRA coefficients is significantly larger than that of ID tokens. Thus, HyperLoRA is usually more time-consuming at this stage. However, HyperLoRA shows remarkable superiority in the second stage, since it doesn't introduce extra attention modules compared with adapter-based methods. Once merged into the base model, our HyperLoRA maintains the same inference performance as the original model.

Table 2. Inference speed of different methods (milliseconds).

Method	IP-Adapter	InstantID	PuLID	HyperLoRA
Preprocess	2996	758	236	1143
Inference	6148	8037	6616	4327

A.2. Interpolation

Benefit from the natural interpolability of LoRA, our HyperLoRA not only supports multiple input images easily but also enables smooth ID interpolation. Fig. 12 illustrates that the personalized images generated by the mixture of two Hyper ID-LoRAs can interpolate naturally from one ID to another. Furthermore, we find that the space of our generated LoRAs has similar properties as the $\mathcal{W}+$ space of StyleGAN [24]. Specifically, the slider LoRA [10] can be generated by HyperLoRA given only a pair of images, as shown in Fig. 11. The image pair consists of a normal face image and an edited one (modifying the attributes of face, e.g., age or eye size). By feeding this pair into HyperLoRA, we obtain two LoRAs, and then subtract the LoRA weights of the original image from the LoRA weights of the edited image. Fig. 11 demonstrates the difference weights between two LoRAs perform a similar behavior as a slider LoRA, possessing the ability to edit local attributes.

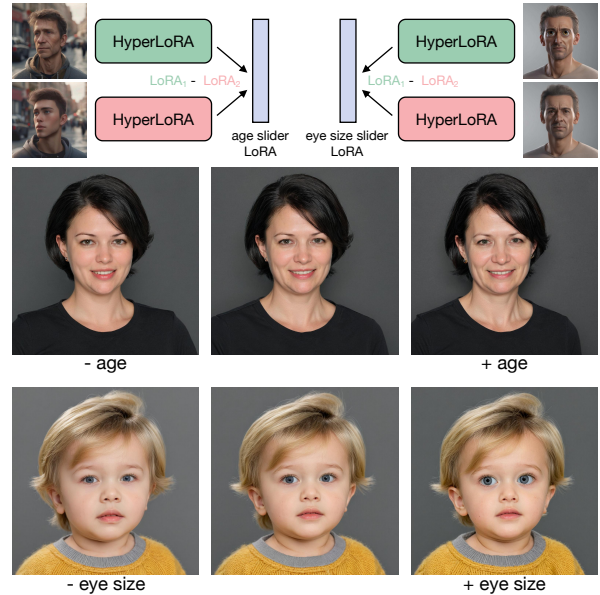


Figure 11. Using the slider LoRA generated by HyperLoRA from paired images, we are allowed to edit the local attributes (e.g. age or eye size) of synthesized portraits.

A.3. More Visual Results

Inference with ControlNet. ControlNet [51] is a practical network to control the diffusion models with additional condition images (edge, pose, depth and more). It helps us to generate satisfied images more easily, getting rid of spending effort trying different seeds, prompts and other parameters. Thus, it is crucial for an ID plug-in to be compatible with ControlNet. As shown in Fig. 13, our HyperLoRA can also generate portrait images with high photorealism and fidelity, when equipped with different ControlNets.

More Text-to-image Results. Fig. 14 and Fig. 15 present more portrait images synthesized by our HyperLoRA across various IDs and scenes, demonstrating its generalization.

A.4. Discussion on Fidelity and Editability

HyperLoRA relies on a large dataset with various facial features for training the LoRA basis, constructing a LoRA space covering sufficient face IDs, so that an unseen face can be reconstructed by the trained LoRA basis. Currently, our dataset comprises only 4.4 million images (in contrast, InstantID is trained with 60 million data). It is crucial to expand the dataset, aiming to further improve the fidelity. In the aspect of editability, it is a feasible solution to introduce the alignment loss of PuLID into HyperLoRA training, weakening the intrusion into the base model.



Figure 12. ID interpolation with HyperLoRA.




Figure 13. Portrait images generated by HyperLoRA and ControlNet.




Figure 14. More Text-to-image results with HyperLoRA.



Figure 15. More Text-to-image results with HyperLoRA.

A.5. Details of Data Filtering

Our data filtering pipeline consists of four steps: **1)** We eliminate the images with a resolution lower than 768×768 ; **2)** LAION Aesthetics Predictor [3] is utilized to assess the aesthetics score of images, and only those with a score higher than 5.5 are retained; **3)** We detect faces from images, and only keep those that contain a single face and the ratio of face region to the entire image is greater than 0.03; **4)** We further filter out those images where plural personal pronouns appear in the corresponding captions.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022. [5](#)
- [2] Jimmy Lei Ba. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [4](#)
- [3] Romain Beaumont and Christoph Schuhmann. Laion-aesthetics predictor. [12](#)
- [4] comfyanonymous. Comfyui: The most powerful and modular diffusion model gui and backend. [6](#)
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [3](#)
- [6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [3](#)
- [7] Zhongyi Fan, Zixin Yin, Gang Li, Yibing Zhan, and Heliang Zheng. Dreambooth++: Boosting subject-driven generation via region-level references packing. In *ACM Multimedia 2024*, 2024. [3](#)
- [8] Rinon Gal, Moab Arar, Yuval Atzmon, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Encoder-based domain tuning for fast personalization of text-to-image models. *ACM Transactions on Graphics (TOG)*, 42(4):1–13, 2023. [3](#)
- [9] Rinon Gal, Or Lichter, Elad Richardson, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Lcm-lookahead for encoder-based text-to-image personalization. *arXiv preprint arXiv:2404.03620*, 2024. [6](#)
- [10] Rohit Gandikota, Joanna Materzyńska, Tingrui Zhou, Antonio Torralba, and David Bau. Concept sliders: Lora adaptors for precise control in diffusion models. In *European Conference on Computer Vision*, pages 172–188. Springer, 2024. [9](#)
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. [2](#)
- [12] Jia Guo and Jiankang Deng. Insightface: 2d and 3d face analysis project. [5](#)
- [13] Zinan Guo, Yanze Wu, Zhuowei Chen, Lang Chen, and Qian He. Pulid: Pure and lightning id customization via contrastive alignment. *Advances in neural information processing systems*, 2024. [6](#), [9](#)
- [14] Ligong Han, Yinxiao Li, Han Zhang, Peyman Milanfar, Dimitris Metaxas, and Feng Yang. Svdif: Compact parameter space for diffusion fine-tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7323–7334, 2023. [3](#)
- [15] Soufiane Hayou, Nikhil Ghosh, and Bin Yu. Lora+: Efficient low rank adaptation of large models. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, 2024. [3](#)
- [16] Zecheng He, Bo Sun, Felix Juefei-Xu, Haoyu Ma, Ankit Ramchandani, Vincent Cheung, Siddharth Shah, Anmol

- Kalia, Harihar Subramanyam, Alireza Zareian, et al. Imagine yourself: Tuning-free personalized image generation. *arXiv preprint arXiv:2409.13346*, 2024. 1
- [17] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 6
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 3, 5
- [19] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 23(47):1–33, 2022. 3
- [20] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*, 2022. 2, 3
- [21] Yuge Huang, Yuhan Wang, Ying Tai, Xiaoming Liu, Pengcheng Shen, Shaoxin Li, Jilin Li, and Feiyue Huang. Curricularface: adaptive curriculum learning loss for deep face recognition. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5901–5910, 2020. 6
- [22] Insightface. Antelopev2, 2023. 2, 4
- [23] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up gans for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10124–10134, 2023. 2
- [24] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 9
- [25] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in neural information processing systems*, 34:21696–21707, 2021. 3
- [26] LEOSAM. Leosam’s helloworld xl, 2023. 6
- [27] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. 5
- [28] Zhen Li, Mingdeng Cao, Xintao Wang, Zhongang Qi, Ming-Ming Cheng, and Ying Shan. Photomaker: Customizing realistic human photos via stacked id embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8640–8650, 2024. 3
- [29] Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, 2024. 3
- [30] I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [31] Jian Ma, Junhao Liang, Chen Chen, and Haonan Lu. Subject-diffusion: Open domain personalized text-to-image generation without test-time fine-tuning. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–12, 2024. 3
- [32] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4296–4304, 2024. 3
- [33] Lianyu Pang, Jian Yin, Baoquan Zhao, Feize Wu, Fu Lee Wang, Qing Li, and Xudong Mao. Attndreambooth: Towards text-aligned personalized text-to-image generation. *arXiv preprint arXiv:2406.05000*, 2024. 3
- [34] Foivos Paraperas Papantoniou, Alexandros Lattas, Stylianos Moschoglou, Jiankang Deng, Bernhard Kainz, and Stefanos Zafeiriou. Arc2face: A foundation model for id-consistent human faces. In *European Conference on Computer Vision*, pages 241–261. Springer, 2024. 6
- [35] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 6
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. 2021. 2, 4
- [37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 3
- [38] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. 2022. 2, 3
- [39] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Wei Wei, Tingbo Hou, Yael Pritch, Neal Wadhwa, Michael Rubinstein, and Kfir Aberman. Hyperdreambooth: Hypernetworks for fast personalization of text-to-image models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6527–6536, 2024. 3
- [40] Simo Ryu. Low-rank adaptation for fast text-to-image diffusion fine-tuning. 3
- [41] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 3
- [42] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 5
- [43] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using

- nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 3
- [44] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 3
- [45] Ming Tao, Hao Tang, Fei Wu, Xiao-Yuan Jing, Bing-Kun Bao, and Changsheng Xu. Df-gan: A simple and effective baseline for text-to-image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16515–16525, 2022. 2
- [46] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 2
- [47] Qixun Wang, Xu Bai, Haofan Wang, Zekui Qin, and Anthony Chen. Instantid: Zero-shot identity-preserving generation in seconds. *arXiv preprint arXiv:2401.07519*, 2024. 1, 3, 6, 9
- [48] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. 2023. 1, 2, 3, 5, 6, 9
- [49] Shih-Ying Yeh, Yu-Guan Hsieh, Zhidong Gao, Bernard BW Yang, Giyeong Oh, and Yanmin Gong. Navigating text-to-image customization: From lycoris fine-tuning to model evaluation. In *The Twelfth International Conference on Learning Representations*, 2023. 3
- [50] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2(3):5, 2022. 2
- [51] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 9
- [52] Qingru Zhang, Minshuo Chen, Alexander Bukharin, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. Adaptive budget allocation for parameter-efficient fine-tuning. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*, 2023. 3