

# DreamSalon: A Staged Diffusion Framework for Preserving Identity-Context in Editable Face Generation

Haonan Lin<sup>1\*</sup>

<sup>1</sup>School of Computer Science and Technology, MOEKLINNS Lab, Xi’an Jiaotong University  
 linhaonan@stu.xjtu.edu.cn

## Abstract

While large-scale pre-trained text-to-image models can synthesize diverse and high-quality human-centered images, novel challenges arise with a nuanced task of “identity fine editing” – precisely modifying specific features of a subject while maintaining its inherent identity and context. Existing personalization methods either require time-consuming optimization or learning additional encoders, adept in “identity re-contextualization”. However, they often struggle with detailed and sensitive tasks like human face editing. To address these challenges, we introduce **DreamSalon**, a noise-guided, staged-editing framework, uniquely focusing on detailed image manipulations and identity-context preservation. By discerning editing and boosting stages via the frequency and gradient of predicted noises, DreamSalon first performs detailed manipulations on specific features in the editing stage, guided by high-frequency information, and then employs stochastic denoising in the boosting stage to improve image quality. For more precise editing, DreamSalon semantically mixes source and target textual prompts, guided by differences in their embedding covariances, to direct the model’s focus on specific manipulation areas. Our experiments demonstrate DreamSalon’s ability to efficiently and faithfully edit fine details on human faces, outperforming existing methods both qualitatively and quantitatively.

## 1. Introduction

Recent work on text-to-image (T2I) generation models has brought significant capabilities in the creation of visual content using textual prompts [9, 29, 33, 38]. A critical challenge known as “identity re-contextualization” (i.e., *ID preservation*, Fig. 1, left) [23, 36], requires modifying a subject’s broad context while preserving its identity. However, a more nuanced task, “identity fine editing” (i.e., *ID-context*



Figure 1. Unlike “identity re-contextualization” (Dreambooth [36]), “identity fine editing” precisely manipulates details while preserving identity and context (DreamSalon).

*preservation*, Fig. 1, right), which demands precise manipulations to a subject’s specific features (like lip color or specific accessories), remains underexplored. This task is crucial for applications such as professional photo editing or social networks, where subtle and accurate changes are required without modifying the subject’s original identity and context [52]. Despite progress in realistically editing object details [17, 47, 53], current models still struggle with detail-oriented human face editing. This difficulty arises due to the complexity of human facial features [8, 43], the high sensitivity of humans to recognize even minor facial changes [41], and technical challenges in maintaining photorealism during editing [39]. These factors make face editing more challenging than editing other objects or scenes.

In “identity fine editing”, a key aspect is to preserve the subject’s inherent identity and source context. We argue that the source textual prompt is essential to achieve this goal, particularly for reconstructing and safeguarding the conceptual essence of the image. However, current methods, which focus on editing images using target textual prompts, often struggle to preserve both identity and context [6, 28, 36]. This imbalance between source and target prompts can lead to the loss of parts that do not need to be edited, underscoring the need for a more nuanced approach that better guides the diffusion process for “identity fine editing”. Another key aspect of “identity fine editing” is the ability to make precise manipulations. GAN-based methods propose to disentangle

\*This work was done with Mengmeng Wang, Yan Chen, Wenbin An, Yuzhe Yao, Guang Dai, Qianying Wang, Liu yong, and Jingdong Wang, during the internship at SGIT AI Lab, State Grid Corporation of China. Due to omission in submission, other authors cannot be added in this paper.

attributes of human faces [11, 14, 30], showing precise editing effects. However, these methods rely on delicate control of hyperparameters or training extra networks to provide guidance. Furthermore, current diffusion-based T2I methods focus on attention mechanisms and rely on manual selections of stages for controllable editing [3, 15, 46]. They often overlook the role of noises in the diffusion process, which carries crucial information about varying distributions and interpretation of latent codes [16]. We posit that predicted noises in the sampling process [42] are informative and help develop a method to intuitively discern the appropriate stages for precise, disentangled editing.

DreamSalon introduces a staged and noise-guided editing framework that adeptly balances image manipulations and *ID-context preservation*. Initially, DreamSalon discerns aggressive-editing and quality-boosting stages using the frequency and gradient of predicted noises. In the editing stage, it uses this frequency information to guide a weighted mix of source and target text embeddings, primarily focusing on target prompts for detailed editing. Then in the boosting stage, its focus is shifted to the source prompt to maintain the original image’s identity and context. As the process moves into the boosting stage, it employs stochastic denoising to enhance the quality of details. Staged editing ensures balance in both specific manipulations and *ID-context preservation*. Additionally, in the editing stage, DreamSalon employs semantic mixing of source and target prompts, leveraging the differences in their embedding covariances for “identity fine editing.” This approach enables precise, context-sensitive adjustments by capturing token variances and relationships, ensuring targeted and nuanced image manipulations. Experiments demonstrate that DreamSalon can edit specific details on human faces efficiently and faithfully. Our contributions are as follows:

1. **Challenges:** We identify challenges of an underexplored T2I task, “*identity fine editing*”, which requires manipulations of specific features and *ID-context preservation*.
2. **Techniques:** We present a staged, noise-guided editing framework, **DreamSalon**. It leverages high-frequency information for detailed image manipulations in the editing stage and stochastic denoising for image quality improvement in the boosting stage. Furthermore, its adaptive and semantic mixing of source and target prompts balances the editability and *ID-context preservation*.
3. **Superiority:** Experiments show the superiority of DreamSalon in editing facial details, surpassing existing methods both qualitatively and quantitatively.

## 2. Related Work

### 2.1. Text-to-Image Generation

Text-conditioned image generation has seen significant advancements recently. Previously, GAN-based [7, 25, 35, 44]

and VAE-based models [5, 10] gain lots of interest at creating high-quality and diverse images. However, these models often struggle to accurately reflect user descriptions and demand substantial optimization time. In comparison, diffusion models stand out for their exceptional semantic understanding and capability to generate varied, photorealistic images directly from textual prompts, offering a distinct advantage in controllability and image quality [2, 21, 32, 51].

### 2.2. Personalized Image Synthesis for Face Identity

Recent works in personalization have shown significant promise in the realm of generating customized concepts [12, 23, 36, 37], but their extensive optimization requirements limit broader applications. Meanwhile, other works have turned to train extra encoders as a means to efficiently synthesize personalized images [6, 27, 40, 49]. Unlike methods that primarily concentrate on preserving identity while altering context significantly, our work is centered on “*identity fine editing*”, aiming to make precise manipulations on specific features while preserving both identity and context. A concurrent work [50] utilizes a weighted-mixing strategy to disentangle the editing targets. However, this method fails to capture the semantic relation between prompts during mixing and overlooks the requirements of adaptive editing in the diffusion process. In contrast, our approach distinguishes itself by utilizing the inherent capabilities of diffusion models, rather than training extra encoders. Moreover, it allows for detailed image manipulations without extensive optimization, leading to more efficient personalization.

## 3. Methods

Our framework focuses on fast, fine editing and *ID-context preservation* (Fig. 2). It begins by generating personalized weights for each identity. Once these weights are loaded into a pre-trained text-to-image model (Sec. 3.1), a noise-guided, staged-editing approach is employed for detailed editing and *ID-context preservation* (Sec. 3.2). Additionally, more precise editing is achieved by semantically mixing source and target prompts (Sec. 3.3).

### 3.1. Preliminary

#### 3.1.1 Denoising Diffusion Implicit Model (DDIM)

DDIM [42] redefines DDPM [16] as  $q_\sigma(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$  a non-Markovian process. DDIM’s reverse process reads:

$$\mathbf{x}_{t-1} = \underbrace{\sqrt{\alpha_{t-1}} \left( \frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \epsilon_t^\theta(\mathbf{x}_t)}{\sqrt{\alpha_t}} \right)}_{\text{predicted } \mathbf{x}_0(\mathbf{P}_t(\epsilon_t^\theta(\mathbf{x}_t)))} + \underbrace{\sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \epsilon_t^\theta(\mathbf{x}_t)}_{\text{direction pointing to } \mathbf{x}_t(\mathbf{D}_t(\epsilon_t^\theta(\mathbf{x}_t)))} + \underbrace{\sigma_t \epsilon_t}_{\text{random noise}}, \quad (1)$$

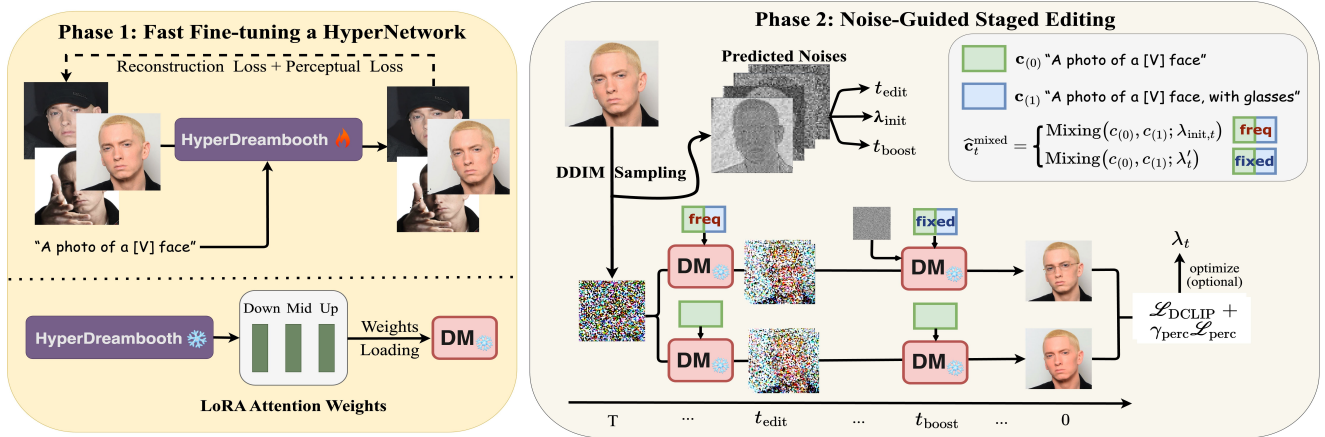


Figure 2. **DreamSalon pipeline.** Phase 1: fast fine-tuning a hypernetwork per identity, obtaining personalization weights for the Latent Diffusion Model. Phase 2: noise-guided staged editing, where the aggressive-editing stage (before  $t_{edit}$ ) and the quality-boosting stage (after  $t_{boost}$ ) are discerned via predicted noises.

where  $t$  is the timestep,  $\alpha_t$  and  $\sigma_t$  are the variance schedule,  $\epsilon_t^\theta$  is the noise predictor, and  $\epsilon_t \sim \mathcal{N}(0, \mathbf{I})$  is standard Gaussian noise. The denoising process in Eq. 1 is stochastic when  $\eta = 1$ , and is deterministic when  $\eta = 0$ .

### 3.1.2 Personalized Weights Generation

The Latent Diffusion Model (LDM) [34] first transforms an input image  $\mathbf{x}$  into a lower-resolution latent space  $\mathbf{z}$  via a Variational Auto-Encoder (VAE) [22]. Then a text-conditioned diffusion model [31] is trained to generate the latent code of the target image from embeddings of text input  $\mathbf{c}$ . In this work, a fast personalized method, HyperDreambooth [37], is used to initialize the attention weights of the LDM. The optimization of these weights is guided by:

$$\mathcal{L}_{\text{HyperDreambooth}} = \alpha \mathbb{E}_{\epsilon, \mathbf{z}, \mathbf{c}} [\|\epsilon - \epsilon^\theta(\mathbf{z}_t, \mathbf{c})\|_2^2 + \|\hat{\theta} - \theta\|_2^2], \quad (2)$$

where  $\hat{\theta}$  are the pre-optimized weight parameters of the personalized model for image  $\mathbf{x}$ , and  $\alpha$  is the hyperparameter that controls the relative weight for two loss terms. This process resembles Dreambooth but is significantly faster (6x) and more storage-efficient (10x). By using 2~4 images of an identity, we obtain personalized weights (Fig. 2, left). These weights are then loaded to the LDM, for the generation of various customized images for that identity (Fig. 2, right).

## 3.2. Staged Editing

In image editing using diffusion models, an adaptive approach to the denoising process is critical, due to the varied distributions of intermediate latent codes [24]. We advocate for aggressive edits in the early stage, followed by quality boosting in the later stage. However, identifying the optimal timing for each stage remains under-researched. This section explores how to perform adaptive, staged editing.

### 3.2.1 Boosting Stage with Stochastic Denoising

To boost the images' quality, we follow the insights of previous work [20] to employ stochastic denoising instead of the deterministic one (Eq. 1). However, indiscriminate noise addition throughout the diffusion process may introduce error accumulation, which modifies content significantly, resulting in undesired editing. We posit that the gradient of predicted noise plays an important role in identifying the appropriate timing ( $t_{boost}$ ) to add noise. Since when latents' gradients are small, their value changes become more consistent, indicating a lower risk of major changes to the content. This stage (e.g., 25% quantile of all gradients, timesteps 20 ( $t_{boost}$ )~0 in Fig. 3) is ideal for applying Gaussian noise:

$$\mathbf{z}_{t-1} = \sqrt{\alpha_{t-1}} \mathbf{P}(\epsilon_t^\theta(\mathbf{z}_t, \mathbf{c})) + \mathbf{D}(\epsilon_t^\theta(\mathbf{z}_t, \mathbf{c})) + \mathbf{I}_{[t < t_{boost}]}(\sigma_t \epsilon_t),$$

where  $\mathbf{I}$  denotes an indicator function for noise addition.

### 3.2.2 Editing Stage with Frequency Guidance

DreamSalon discerns the editing stage for manipulations on specific and detailed features, by identifying the high-frequency predicted noises (e.g., 75% quantile of all frequencies, timesteps 50~30 ( $t_{edit}$ ) in Fig. 3). As higher frequencies correlate strongly with the finer aspects of an image and signify rapid intensity changes [13], intermediates with high frequency are ideal for editing without impacting the overall structure, offering efficient ways to alter specific details. Conversely, low-frequency components, indicative of broader, uniform areas, are more influential in modifying the overall appearance rather than specific details. To achieve adaptive control during the diffusion process, we start with the mixing of source and target prompts. Since our goal is to edit the specific details while preserving identity and context, amplifying the contribution of the target prompt during

the editing stage aids the manipulation of specific details a lot. This mixing is controlled by a weighting strategy:

$$\mathbf{c}_t^{\text{mixed}} = (1 - \lambda_t)\mathbf{c}_{(0)} + \lambda_t\mathbf{c}_{(1)}, \quad (3)$$

where  $\mathbf{c}_{(0)}$  and  $\mathbf{c}_{(1)}$  represent the CLIP embeddings of the source and target prompts, respectively. The weighting factor  $\lambda_t$  at each time step  $t$ , is the only (optional) parameter optimized in our framework.  $\mathbf{c}_t^{\text{mixed}}$  are mixed text embeddings, which balances the impact of the source and target prompts on the generation of edited images.

For “*identity fine editing*”, our intuition is using the source prompt to help guide the diffusion model to maintain the identity and context, while incorporating the detailed edits suggested by the target prompt. Thus, in the editing stage, DreamSalon puts greater weights on target prompts for aggressive editing (*i.e.*,  $\lambda_t$  should be larger). And after that, it focuses on source prompts to retain the core conceptual content of the original image (*i.e.*,  $\lambda_t$  should be set as relatively smaller). However, the initialization of the weight factors  $\lambda_t$  is intricate. Inappropriate initialization can degrade the editing performance and increase optimization time significantly. Rather than manually setting these factors, we propose to use the inherent attributes of intermediates  $\mathbf{x}_t = \psi(\mathbf{z}_t)$  from the DDIM sampling process, where  $\psi$  denotes the pretrained decoder. Recall that  $\mathbf{x}_t$  with high frequency is mostly present at the early stage, indicating that more aggressive editing is required at these timesteps. Therefore, the initialized weight factors and the mixed text embeddings can be represented as:

$$\lambda_{\text{init}} = \begin{cases} \text{Normalize}(\text{FFT}(\psi(\mathbf{z}_t))) & \text{if } t \leq t_{\text{edit}}, \\ \lambda' & \text{otherwise,} \end{cases} \quad (4a)$$

$$\mathbf{c}_t^{\text{mixed}} = (1 - \lambda_t^{\text{init}})\mathbf{c}_{(0)} + \lambda_t^{\text{init}}\mathbf{c}_{(1)}, \quad (4b)$$

where FFT represents the Fast Fourier Transform,  $\lambda_{\text{init}}$  is a vector composed of all  $\lambda_t$ , and  $\lambda'$  is a hyperparameter used after the editing stage. We set  $\lambda'$  as 0.2 in our experiments, to leverage the source prompt to improve the reconstruction and preserve identity and context better.

### 3.3. Covariance Guidance for Detailed Editing

For sensitive tasks like facial attribute editing, directly mixing text embeddings using a weighted sum (Eq. 4b) can not effectively capture the expressiveness of source and target prompts. This is because a weighted sum uniformly treats each token, failing to account for its specific roles and context. To address this, we propose a covariance guidance method. Since the covariance matrix can assess how variations in one token of the embedding relate to another (Eq. 5a), it helps develop a more context-aware integration of prompts [48]. In particular, by comparing the covariance matrices of source ( $\text{Cov}_{\mathbf{c}_{(0)}}$ ) and target ( $\text{Cov}_{\mathbf{c}_{(1)}}$ ) embeddings,

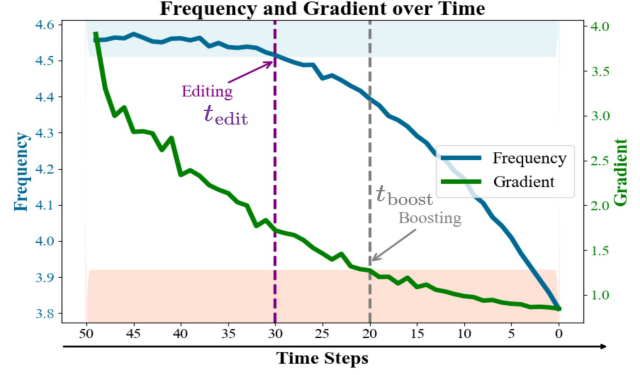


Figure 3. Stage discernment based on frequency and gradient of predicted noises. The editing stage is determined by high-frequency predicted noises (75% quantile), followed by the boosting stage where gradients are relatively smaller (25% quantile). More details about the frequency and gradient of predicted noises are available in the *Suppl.*

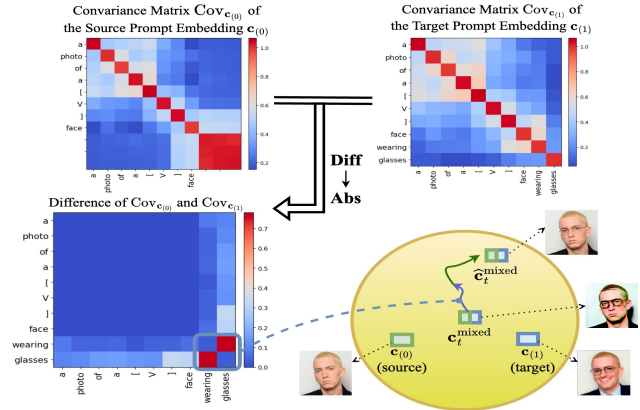


Figure 4. Covariance analysis in prompt embeddings: differences in covariance matrices for source and target prompt embeddings, guide the semantic mixing of prompts for precise attribute editing in generated images.

we identify tokens that significantly contribute to desired changes from the difference between  $\text{Cov}_{\mathbf{c}_{(0)}}$  and  $\text{Cov}_{\mathbf{c}_{(1)}}$ , as presented in Fig. 4. Tokens with significant differences are marked in red, indicating key tokens where the target prompt diverges from the source, thus requiring more attention in the editing process. To utilize this information to guide the diffusion process, we present  $\text{CovDiff}$  as Eq. 5b, which acts as a metric for token-specific guidance. It emphasizes tokens in the target prompt that bring desired manipulations, enabling precise editing control.

$$\text{Cov}_{\mathbf{c}_{(0)}} = \mathbf{c}_0\mathbf{c}_0^T / (n_{(0)} - 1) \quad (5a)$$

$$\text{CovDiff} = \text{Normalize}(\max_{i \text{ or } j} |\text{Cov}_{\mathbf{c}_{(1)}} - \text{Cov}_{\mathbf{c}_{(0)}}|) \quad (5b)$$

In Eq. 5,  $n_{(0)}$  is the dimension of position tokens and

$\max_i$  or  $j$  represents maximization via the x-axis or y-axis since the covariance matrix is symmetric. CovDiff consists of values in the range (0,1), as many as there are tokens. For more detailed image editing, tokens with higher values in CovDiff should be prioritized in the mixing of text embeddings. We offer token-level, semantic control for improved editability by integrating CovDiff with Eq. 4b as follows:

$$\hat{\mathbf{c}}_t^{\text{mixed}} = \begin{cases} \text{CovDiff} \odot ((1 - \lambda_t^{\text{init}})\mathbf{c}_{(0)} + \lambda_t^{\text{init}}\mathbf{c}_{(1)}) & \text{if } t \leq t_{\text{edit}}, \\ (1 - \lambda_t^{\text{init}})\mathbf{c}_{(0)} + \lambda_t^{\text{init}}\mathbf{c}_{(1)} & \text{otherwise,} \end{cases} \quad (6)$$

where  $\odot$  denotes the Hamilton multiplication. By multiplying CovDiff with the adaptive weighted sum of source and target embeddings, we can modulate the editing intensity of each token to control the contribution of each token in the mixed text embedding. This allows the model to perform more precise manipulations on target features based on semantic guidance.

### 3.4. Overall Loss

Instead of using the initialized weight factors  $\lambda_{\text{all}}$ , we provide an option to optimize them by combining directional CLIP loss and perceptual loss. With CLIP’s image encoder  $\mathcal{E}_I$  and text encoder  $\mathcal{E}_T$  [31], directional loss with cosine distance achieves homogeneous editing without mode collapse [11]:

$$\mathcal{L}_{\text{DCLIP}}(\mathbf{x}_{\text{edit}}, y_{\text{target}}; \mathbf{x}_{\text{source}}, y_{\text{source}}) = 1 - \frac{\Delta \mathbf{I} \cdot \Delta \mathbf{T}}{\|\Delta \mathbf{I}\| \|\Delta \mathbf{T}\|}, \quad (7)$$

where  $\Delta \mathbf{T} = \mathcal{E}_T(y_{\text{target}}) - \mathcal{E}_T(y_{\text{source}})$  and  $\Delta \mathbf{I} = \mathcal{E}_I(\mathbf{x}_{\text{edit}}) - \mathcal{E}_I(\mathbf{x}_{\text{source}})$  with source and edited image  $(\mathbf{x}_{\text{source}}, \mathbf{x}_{\text{edit}})$ , source and target prompts  $(y_{\text{source}}, y_{\text{target}})$ .

For edits that require identity preservation, we use the perceptual loss [18] defined in Eq. 8a, to prevent drastic changes in semantic content. Our total loss function is a weighted combination of these losses:

$$\mathcal{L}_{\text{perc}}(\mathbf{x}_0^s, \mathbf{x}_0^t) = \|\phi(\mathbf{x}_0^s) - \phi(\mathbf{x}_0^t)\|_1, \quad (8a)$$

$$\mathcal{L}_{\text{total}}(w) = \mathcal{L}_{\text{DCLIP}}(w) + \gamma_{\text{perc}} \mathcal{L}_{\text{perc}}(w), \quad (8b)$$

where  $\phi(\cdot)$  denotes a perceptual network that encodes a given image, and  $\gamma_{\text{perc}}$  is a weighting hyperparameter.

## 4. Experiments

### 4.1. Experimental Settings

**Datasets.** We first conduct experiments on the CelebA-HQ dataset [19]. Following existing works [23, 36, 49], we use 35 identities for image editing. Each identity is edited with 20 different prompts, and we randomly generated five edited images per identity-prompt combination, totaling 3,500 images. Furthermore, acknowledging the performance on unaligned face images, we construct a **Face-oriented Fine**

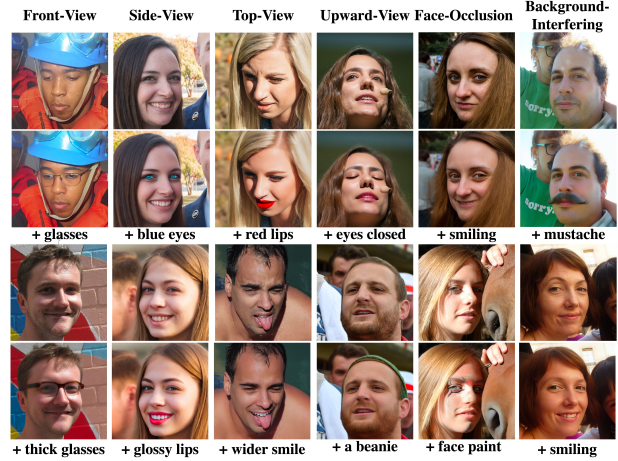


Figure 5. The FFE-Bench for fine-grained face editing across different views and challenging conditions, with DreamSalon’s edits like attribute additions and expression changes.

**Editing Benchmark** dataset from FFHQ-unaligned (*FFE-Bench*, Fig. 5), which features 600 face images with six types: four views (front, side, top, upward) and two conditions (face-occlusion, background-interfering). Each image is annotated with 20 source/target prompts, editing instructions, and an editing mask for metric calculations.

**Evaluation Metrics.** Following Dreambooth [36], we evaluate our method with three metrics: *CLIP-I*, *CLIP-T*, and *DINO-I*, which assess visual similarity, text-image alignment, and identity uniqueness, respectively. For *CLIP-I*, we calculate the CLIP visual similarity between the source and the generated images. For *CLIP-T*, we calculate the CLIP text-image similarity between the generated images and the text prompts given. For *DINO-I*, we calculate cosine similarity between the ViT-B/16 DINO [4] embeddings of source and generated images. Moreover, we adopt fine-tuning and editing time as a metric to evaluate the efficiency.

**Implementation Details.** We choose SD 1.5 as our base T2I model. During obtaining the weights using HyperDreambooth [37], the learning rate is set to 5e-5, the embedding regularization weight is set to 1e-4. The weights are generated with 2~4 images for each identity. During inference, we use the PLMs sampler [26] with 50 timesteps, and the scale of classifier-free guidance is 5. For adaptive editing, we use  $\lambda' = 0.2$  after the editing stage. Optional optimization is performed for 3 iterations with a learning rate of 5e-2.

### 4.2. Experimental Results

We present both qualitative and quantitative results to highlight our method’s superiority in “*identity fine editing*” over current SOTA diffusion-based models. Our comparisons include both fine-tuning methods [36, 37] and fine-tuning free methods [1, 15, 45–47, 49, 50], in which some methods

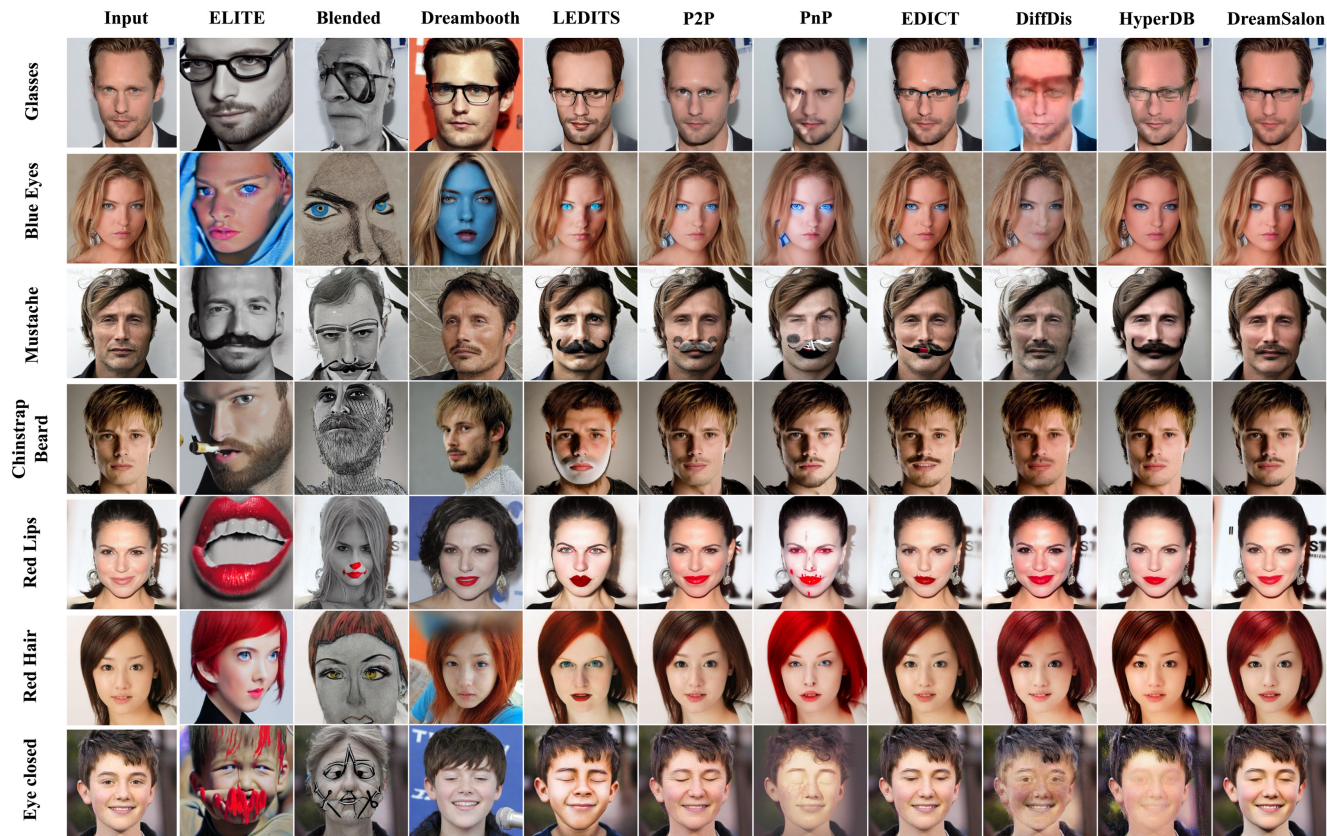


Figure 6. Comparative analysis showcases DreamSalon’s precision in aligning images with text descriptions while maintaining identity and context, surpassing other leading methods.

require additional guidance, such as masks [1, 49].

#### 4.2.1 Qualitative Results

Our method, DreamSalon, precisely edits image details according to text prompts, demonstrating versatility in various edits like adding accessories, altering facial hair, or changing eye color, as shown in Fig. 5, 6. A critical aspect of our method’s performance is its ability to preserve both identity and context in the images, avoiding altering prompt-unrelated image aspects, and ensuring detail-specific changes. In comparison, ELITE and Blended Latent Diffusion often change both identity and context, Dreambooth alters context, and LEDITS maintains context but changes identity. Other methods either introduce undesired changes or fail to make prompt-specific edits. For instance, PnP and P2P inaccurately edit earring colors when prompted for eye color changes, showing less control (2nd row); HyperDreambooth and EDICTS struggle with hair color changes (6th row). DreamSalon’s edits maintain realism, blending seamlessly with original image characteristics like lighting and texture, in contrast to the stylistic, lighting, and texture changes seen in Blended Latent Diffusion and

LEDITS. More qualitative results are in the *Suppl.*

#### 4.2.2 Quantitative Results

As presented in Tab. 1, DreamSalon excels in our performance comparison, achieving the highest CLIP-I score, indicating superior *ID-context preservation* on face images during editing. This means that both identity and context in edited images closely match their originals. DreamSalon also leads in CLIP-T scores (0.247), reflecting its accuracy in mirroring text prompts in image edits, demonstrating a robust understanding and application of textual instructions. Additionally, it tops in DINO-I scores, maintaining conceptual similarity with target images. Regarding time efficiency (Tab. 3), DreamSalon, assisted by HyperDreambooth, fine-tunes 6x faster and requires 10x smaller storage than Dreambooth. After fine-tuning, DreamSalon’s editing time is slightly longer than Dreambooth and HyperDreambooth, but its editing performance significantly outdoes them. Among fine-tuning-free methods, ELITE is the fastest, while DreamSalon matches this editing speed. Other methods like P2P, PnP, and LEDITS take longer to edit than DreamSalon. When making multiple different ed-

Table 1. Quantitative comparisons on CelebA-HQ and our FFE-Bench. Our method outperforms SOTA T2I methods in terms of face similarity, text-alignment, and conceptual similarity.

Method	CLIP-I $\uparrow$	CLIP-T $\uparrow$	DINO-I $\uparrow$
<b>CelebA-HQ</b>			
DB [36]	0.705	0.210	0.150
LEDITS [45]	0.606	0.242	0.940
P2P [15]	0.731	0.229	0.937
PnP [46]	0.621	0.245	0.928
EDICT [47]	0.812	0.212	0.948
DiffDis [50]	0.779	0.205	0.881
HyperDB [37]	0.675	0.224	0.620
<b>DreamSalon</b>	<b>0.837</b>	<b>0.247</b>	<b>0.958</b>
<b>FFE-Bench</b>			
PnP	0.603	0.238	0.924
EDICT	0.808	0.204	0.898
HyperDB	0.639	0.215	0.598
<b>DreamSalon</b>	<b>0.815</b>	<b>0.242</b>	<b>0.932</b>

Table 2. Ablation study on personalized weight generation, staged editing and covariance-guided prompts mixing.

Method	CLIP-I $\uparrow$	CLIP-T $\uparrow$	DINO-I $\uparrow$
w/o HyperDB	0.778	0.224	0.847
w/o Staged	0.734	0.228	0.792
w/o Frequency	0.750	0.233	0.835
w/o Boosting	0.804	0.241	0.904
w/o CovDiff	0.802	0.234	0.877
w/o Opt $\lambda_t$	0.812	0.239	0.939
<b>DreamSalon</b>	<b>0.837</b>	<b>0.247</b>	<b>0.958</b>




Figure 7. Qualitative comparisons with various components omitted, culminating in the full DreamSalon method which integrates all features for optimal editing outcomes.

its per identity, the time cost of fine-tuning is negligible.

### 4.3. Ablation Studies

Ablation studies, as seen in Fig. 7 and Tab. 2, demonstrate the impact of each method on the performance of face editing. “w/o HyperDB” indicates without using HyperDream-

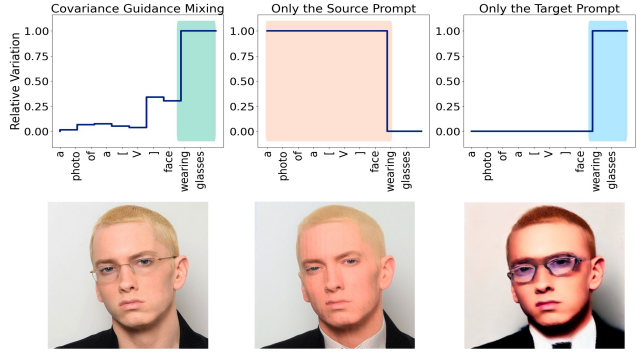


Figure 8. Impact of covariance guidance on prompt mixing: the effect of mixing source and target prompts with covariance guidance versus using only source or only target prompt, and their respective influence on the edited image outcome.

booth to fine-tune, showing reduced identity preservation. “w/o Frequency” indicates using a prefixed value for  $\lambda_t$  of all timesteps, which leads to a loss of specific details like mouth and collar, indicating the need for a dynamic value for  $\lambda_t$  at different timesteps. “w/o Boosting” denotes the use of deterministic denoising instead of stochastic one during the boosting stage, resulting in the loss of finer details such as eyes. “w/o Staged” condition is the combination of “w/o Frequency” and “w/o Boosting”. “w/o CovDiff” alters overall texture and saturation due to the absence of semantic, detailed guidance in mixing text embeddings. Finally, “w/o Opt  $\lambda_t$ ” introduces some artifacts. DreamSalon outperforms all ablated versions in all metrics, affirming the efficacy of its entire method.

## 4.4. Experimental Analysis

### 4.4.1 Different Covariance Guidance

We provide further analysis of the role of covariance guidance (Eq. 5b, 6) in Fig. 8. As stated before, tokens with high values in CovDiff signal its importance for precise edits, and low CovDiff indicates less impact, necessitating balance to avoid overemphasis and maintain overall context. The first column represents the CovDiff used in our methods, considering the difference between the covariance matrices of the source and target embeddings. Tokens that present a larger variance in the target text prompt, as indicated by a larger CovDiff, guide the model to focus its editing efforts, introducing new details (adding glasses), while preserving identity-context. The second column employs an extreme condition where the tokens unique in the target prompt are ignored, the editing fails to correspond to the text instruction in the target prompt. The third column considers another extreme condition, where tokens in the source prompt are ignored, even though editing corresponding to the text instruction in the target prompt is performed, there is a failure to preserve identities and context. By focusing on tokens

Table 3. Comparison of fine-tuning time and editing time with fine-tuning free and fine-tuning based methods.

Methods	Fine-tuning Free Methods						Fine-tuning based Methods		
	Blended	P2P	PnP	LEDITS	EDICT	ELITE	Dreambooth	HyperDB	DreamSalon
Editing Time (s)	62	97	65	48	42	28	14	14	26
Fine-tuning Time (s)	N/A						640	106	106

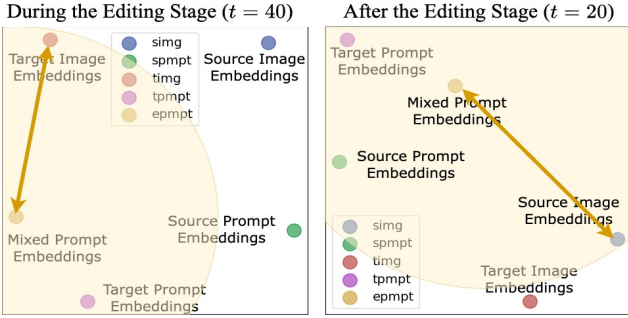


Figure 9. The embeddings of images and prompts during and after the editing stage, highlighting the shift from a focus on target image embeddings to integration with source image embeddings for precise and context-aware image manipulation.

that bring notable changes in the editing stage, DreamSalon ensures precise editing on minor details, enabling the creation of images that are not only visually appealing but also contextually consistent with both source and target prompts.

#### 4.4.2 Mixed Text Embeddings with Image Embeddings

To better understand how the adaptive mixing of prompt embeddings contributes to the precise editing, we assess the distance between mixed text embeddings and source/target image embeddings during and after the editing stage, as depicted in Fig. 4. During the editing stage (left), the mixed prompt embeddings closely align with the target image embeddings, indicating heightened editability for detailed modifications. On the contrary, after the editing stage (right), the mixed prompts exhibit closer proximity to the source image embeddings, signifying a greater focus on reconstructing aspects from the source image. This emphasis on source image reconstruction contributes to identity and context preservation after the editing stage.

#### 4.4.3 Discernment of Different Stages

As shown in Fig. 10, our investigation into discerning the editing and boosting stages involves selecting different quantiles for these phases. The first column adheres to our method’s quantile settings, resulting in precise editing while preserving both identity and context. The second column demonstrates a broader range for quality boosting, leading to significant identity or context changes due to excessive

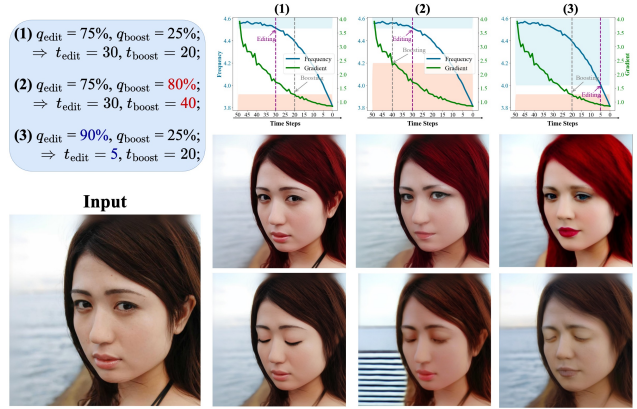


Figure 10. Selections of editing and boosting stages with different quantiles, influencing editing intensities and noise-boosting percentages. The first row depicts varying stage durations, while the second and third rows demonstrate the application of these stages in altering hair color and eye state, respectively.

noise addition. The third column presents a range for more aggressive editing, where ID-preservation is compromised as a result of insufficient integration with the source prompt, which is essential for maintaining the original essence.

## 5. Conclusions

In summary, DreamSalon offers a framework in “*identity fine editing*” for text-to-image models, adeptly manipulating specific features while safeguarding the subject’s identity and context. It outperforms recent work with a noise-guided, staged-editing framework that precisely manipulates image details through adaptive editing and semantic prompt mixing. Our experiments showcase DreamSalon’s exceptional performance in precise and efficient human face editing, marking its advance over existing approaches.

**Acknowledgement** This work was supported by National Science and Technology Major Project (2022ZD0117102), National Natural Science Foundation of China (62293551, 62377038, 62177038, 62277042). Project of China Knowledge Centre for Engineering Science and Technology, Project of Chinese academy of engineering “The Online and Offline Mixed Educational Service System for ‘The Belt and Road’ Training in MOOC China”. “LENOVO-XJTU” Intelligent Industry Joint Laboratory Project.



## References

- [1] Omri Avrahami, Ohad Fried, and Dani Lischinski. Blended latent diffusion. *ACM Transactions on Graphics (TOG)*, 42(4):1–11, 2023. 5, 6
- [2] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. 2
- [3] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiao-hu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. *arXiv preprint arXiv:2304.08465*, 2023. 2
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 5
- [5] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation via masked generative transformers. *arXiv preprint arXiv:2301.00704*, 2023. 2
- [6] Zhuowei Chen, Shancheng Fang, Wei Liu, Qian He, Mengqi Huang, Yongdong Zhang, and Zhendong Mao. Dreamidentity: Improved editability for efficient face-identity preserved image generation. *arXiv preprint arXiv:2307.00300*, 2023. 1, 2
- [7] Jun Cheng, Fuxiang Wu, Yanling Tian, Lei Wang, and Dapeng Tao. Rifegan: Rich feature generation for text-to-image synthesis from prior knowledge. In *CVPR*, pages 10911–10920, 2020. 2
- [8] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, pages 4685–4694, 2018. 1
- [9] Ziyi Dong, Pengxu Wei, and Liang Lin. Dreamartist: Towards controllable one-shot text-to-image generation via contrastive prompt-tuning. *arXiv preprint arXiv:2211.11337*, 2022. 1
- [10] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 2
- [11] Rinon Gal, Or Patashnik, Haggai Maron, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *arXiv preprint arXiv:2108.00946*, 2021. 2, 5
- [12] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 2
- [13] Rafael C Gonzales and Paul Wintz. *Digital image processing*. Addison-Wesley Longman Publishing Co., Inc., 1987. 3
- [14] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. *Advances in neural information processing systems*, 33:9841–9850, 2020. 2
- [15] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 2, 5, 7
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2
- [17] Wenjing Huang, Shikui Tu, and Lei Xu. Pfb-diff: Progressive feature blending diffusion for text-driven image editing. *arXiv preprint arXiv:2306.16894*, 2023. 1
- [18] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 694–711. Springer, 2016. 5
- [19] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 5
- [20] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35:26565–26577, 2022. 3
- [21] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2426–2435, 2022. 2
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 3
- [23] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multiconcept customization of text-to-image diffusion. *arXiv preprint arXiv:2212.04488*, 2022. 1, 2, 5
- [24] Mingi Kwon, Jaeseok Jeong, and Youngjung Uh. Diffusion models already have a semantic latent space. *arXiv preprint arXiv:2210.10960*, 2022. 3
- [25] Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip H.S. Torr. Controllable text-to-image generation. *Neural Information Processing Systems, Neural Information Processing Systems*, 2019. 2
- [26] Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. *arXiv preprint arXiv:2202.09778*, 2022. 5
- [27] Yiyang Ma, Huan Yang, Wenjing Wang, Jianlong Fu, and Jiaying Liu. Unified multi-modal latent diffusion for joint subject and text conditional image generation. *arXiv preprint arXiv:2303.09319*, 2023. 2
- [28] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6038–6047, 2023. 1
- [29] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation

- and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. [1](#)
- [30] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2085–2094, 2021. [2](#)
- [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [3](#), [5](#)
- [32] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. [2](#)
- [33] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. [1](#)
- [34] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. [3](#)
- [35] Shulan Ruan, Yong Zhang, Kun Zhang, Yanbo Fan, Fan Tang, Qi Liu, and Enhong Chen. Dae-gan: Dynamic aspect-aware gan for text-to-image synthesis. In *ICCV*, pages 13960–13969, 2021. [2](#)
- [36] 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. *arXiv preprint arXiv:2208.12242*, 2022. [1](#), [2](#), [5](#), [7](#)
- [37] 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. *arXiv preprint arXiv:2307.06949*, 2023. [2](#), [3](#), [5](#), [7](#)
- [38] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. [1](#)
- [39] Shunsuke Saito, Lingyu Wei, Liwen Hu, Koki Nagano, and Hao Li. Photorealistic facial texture inference using deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5144–5153, 2017. [1](#)
- [40] Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. Instantbooth: Personalized text-to-image generation without test-time finetuning. *arXiv preprint arXiv:2304.03411*, 2023. [2](#)
- [41] Pawan Sinha, Benjamin Balas, Yuri Ostrovsky, and Richard Russell. Face recognition by humans: Nineteen results all computer vision researchers should know about. *Proceedings of the IEEE*, 94(11):1948–1962, 2006. [1](#)
- [42] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. [2](#)
- [43] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1708, 2014. [1](#)
- [44] Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Xiao-Yuan Jing, Fei Wu, and Bingkun Bao. Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis. *arXiv preprint arXiv:2008.05865*, 2020. [2](#)
- [45] Linoy Tsaban and Apolinário Passos. Ledits: Real image editing with ddpm inversion and semantic guidance. *arXiv preprint arXiv:2307.00522*, 2023. [5](#), [7](#)
- [46] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1921–1930, 2023. [2](#), [7](#)
- [47] Bram Wallace, Akash Gokul, and Nikhil Naik. Edict: Exact diffusion inversion via coupled transformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22532–22541, 2023. [1](#), [5](#), [7](#)
- [48] Jonathan J Webster and Chunyu Kit. Tokenization as the initial phase in nlp. In *COLING 1992 volume 4: The 14th international conference on computational linguistics*, 1992. [4](#)
- [49] Yuxiang Wei, Yabo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. *arXiv preprint arXiv:2302.13848*, 2023. [2](#), [5](#), [6](#)
- [50] Qiucheng Wu, Yujian Liu, Handong Zhao, Ajinkya Kale, Trung Bui, Tong Yu, Zhe Lin, Yang Zhang, and Shiyu Chang. Uncovering the disentanglement capability in text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1900–1910, 2023. [2](#), [5](#), [7](#)
- [51] 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*, 2022. [2](#)
- [52] Fangneng Zhan, Yingchen Yu, Rongliang Wu, Jiahui Zhang, Shijian Lu, Lingjie Liu, Adam Kortylewski, Christian Theobalt, and Eric Xing. Multimodal image synthesis and editing: A survey. *arXiv preprint arXiv:2112.13592*, 2021. [1](#)
- [53] Zhixing Zhang, Ligong Han, Arnab Ghosh, Dimitris N Metaxas, and Jian Ren. Sine: Single image editing with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6027–6037, 2023. [1](#)