

Unsupervised Region-Based Image Editing of Denoising Diffusion Models

Zixiang Li^{1,2}, Yue Song⁴, Renshuai Tao^{1,2}, Xiaohong Jia⁵, Yao Zhao^{1,2,3*}, Wei Wang^{1,2*}

¹Institute of Information Science, Beijing Jiaotong University

²Visual Intelligence +X International Cooperation Joint Laboratory of MOE

³Pengcheng Laboratory, Shenzhen, China

⁴University of Trento, Italy

⁵Lanzhou Jiaotong University

{zixiangli, yzhao}@bjtu.edu.cn

Abstract

Although diffusion models have achieved remarkable success in the field of image generation, their latent space remains under-explored. Current methods for identifying semantics within latent space often rely on external supervision, such as textual information and segmentation masks. In this paper, we propose a method to identify semantic attributes in the latent space of pre-trained diffusion models without any further training. By projecting the Jacobian of the targeted semantic region into a low-dimensional subspace which is orthogonal to the non-masked regions, our approach facilitates precise semantic discovery and control over local masked areas, eliminating the need for annotations. We conducted extensive experiments across multiple datasets and various architectures of diffusion models, achieving state-of-the-art performance. In particular, for some specific face attributes, the performance of our proposed method even surpasses that of supervised approaches, demonstrating its superior ability in editing local image properties.

Introduction

Diffusion models have achieved remarkable success in the field of image generation, with notable examples including the Denoising Diffusion Probability Model (DDPM) (Ho, Jain, and Abbeel 2020), the Denoising Diffusion Implicit Model (DDIM) (Song, Meng, and Ermon 2020), and score-based generative models (Song et al. 2020). Diffusion models have set new benchmarks in various applications such as image generation, video generation, image editing, and inverse problems. When deploying diffusion models for downstream tasks like image editing, additional supervision such as textual descriptions or sketches is typically required, and this often involves the integration of supplementary modules or fine-tuning pre-trained models (e.g., Lora (Hu et al. 2021), DiffusionClip (Kim, Kwon, and Ye 2022)). These techniques are meticulously engineered to minimize undesired alterations to the noise space, thereby avoiding irregular changes or artifacts. The necessity for such additional supervision or training underscores the current limitations in the inherent information available within diffusion models. Nevertheless, as pre-trained diffusion models have

*Corresponding author




Figure 1: **Motivation of our proposed method.** In this figure, mouth editing is used as an example. Direct editing often results in significant changes to areas beyond the mouth. By using Jacobian matrix projection, we can suppress these unwanted changes, allowing for more precise editing.

demonstrated unexpected versatility across various domains, it becomes imperative to explore and leverage their intrinsic capabilities to fully realize their potential. Therefore, it is more favored to leverage the pre-trained models for unsupervised image editing.

Diffusion models have achieved impressive results in guided generation tasks. The most outstanding one is the use of Contrastive Language-Image Pretraining (CLIP) (Radford et al. 2021) to convert the user input text prompt into the condition of the diffusion model input, thereby generating high-quality images. GLIDE (Nichol et al. 2021), DALL-E 2 (Ramesh et al. 2022), and Stable Diffusion (Rombach et al. 2022) are some of the best models currently. These models use CLIP to align the text space with the latent space of the diffusion model, so as to achieve text-guided image generation. However, there is still no comprehensive exploration of the properties of the latent space of the diffusion model itself. Current research on the latent space of diffusion models remains a bit under-explored.

In the denoising process of the diffusion model, the image space is generally represented as $\mathbf{x}_{\{1:T\}}$, where $\{1:T\}$ represents the time step of denoising. In the denoising process of the diffusion model, the denoised image space is generally represented as \mathbf{x} . All vector dimensions in $\mathbf{x}_{\{1:T\}}$ are the same as the initial input \mathbf{x}_0 and the output $\epsilon_{\{1:T\}}$ after the denoising network. In this article, we abbreviate the pixel space of the diffusion model as \mathbf{x} -space. The \mathbf{x} -space within diffusion models is generally considered to be devoid of seman-

tics. To address this, Diffusion autoencoders (Preechakul et al. 2022) introduces a semantic encoder that maps x -space into semantically rich vectors. Subsequently, Kwon (Kwon, Jeong, and Uh 2022) proposed that in diffusion models employing the U-Net architecture, the intermediate skip connections, termed h -space, can be utilized to control image semantics. Subsequent investigations (Park et al. 2023) have aimed to establish a linkage between h -space and x -space. Despite these efforts, the supervisory information can only be applied within the x -space, and these approaches strive to connect the x -space with the h -space. However, these supervised methodologies fail to fully exploit the potential abilities of identifying inherent semantics in h -space. Subsequent research by Boundary Diffusion (Zhu et al. 2024) examined the concurrent alterations in x -space and h -space to facilitate semantic editing. However, their approach relied on simple classification methods to identify vectors controlling semantics, resulting in uncontrollable semantic changes. Semdiff (Haas et al. 2024) introduced significant advancements by incorporating masks to locate the local Jacobian matrix. This innovation considerably enhanced the efficiency of identifying local semantics, and as an unsupervised method, it greatly enriched the discovered semantics. Nonetheless, similar to preceding methods, Semdiff focuses solely on local semantic changes and fails to maintain structural details outside the targeted edited area.

In our work, we introduce a novel algorithm named Region-Based Editing (RBE) to identify semantics based on regions while preserving the rest details of the image. Figure 1 shows the core idea of our method. Our method is a fully unsupervised algorithm that neither requires fine-tuning of pre-trained models nor the addition of auxiliary modules. The primary objective is to modify only the pixels within the region of interest, leaving the exterior pixels unchanged. This approach ensures post-editing, where only the targeted region transforms while the surrounding areas remain unaffected. Our pixel-based optimization method concerns decomposing the Jacobian matrix of the noise prediction model, and the goal is to maximize the pixel changes within the region of interest while minimizing changes outside this region. We only need to add a rough bounding box to the area to be edited, apply this box to the noise prediction network, and calculate the Jacobian matrix corresponding to the inside and outside of the box for the modified noise prediction network. Finally, we use the vectors of the area to be edited as a set of bases and use projection to suppress the influence of the vectors represented by the Jacobian matrix outside the area. Our method incorporates information outside the region of interest during local semantic editing within diffusion models. This process requires only a coarse mask, eliminating the need for a fine segmentation network. Moreover, due to the universality and computational efficiency of our framework, it can be applied to all pre-trained diffusion models utilizing the U-Net architecture. Additionally, we investigate the relationship between the h -space of diffusion models and semantic discovery and editing, deepening our understanding of the model’s latent space. Extensive experiments and comprehensive evaluations validate the effectiveness of our proposed method.

Background

Image Editing in Diffusion Model. In the field of image editing, a large number of works based on diffusion models have emerged in recent years. These works highlight the potential and versatility of diffusion models in improving image editing performance. The image editing of diffusion models has multiple categories, such as text guidance, reference image guidance, semantic segmentation map guidance, and mask guidance, etc. These methods include both earlier traditional and current multimodal conditional methods (Huang et al. 2024).

Text-based methods are the most commonly used. GLIDE (Nichol et al. 2021) is the first work to use text to directly control image generation. Unlike image generation tasks, image editing focuses on changing the appearance, content or structure of existing images. DiffusionClip (Kim, Kwon, and Ye 2022) is an early method that uses CLIP for diffusion models. The success of CLIP has inspired many new explorations (Liu et al. 2024; Tan et al. 2024). For image editing based on diffusion models guided by other conditions, controlnet (Zhang, Rao, and Agrawala 2023) is a guidance that adds spatial condition control. Zhang et al. (2023) further explores control methods in video generation. Mask-based image guidance is a prevalent method that primarily relies on segmentation networks or manual definitions to apply masks. DiffEdit (Couairon et al. 2022) simplifies the process of semantic editing by automatically generating masks that isolate specific areas for modification.

Another method is to operate on the latent space of the diffusion models. Their advantage is that the ability to edit images can be achieved without a large amount of paired training data, which depends on the pre-trained diffusion model. DragDiffusion (Mou et al. 2023) is one of the representatives of the methods, which can directly guide editing through image features. It observes that there is rich semantic information in the features of U-Net, and this information can be used to construct energy functions to guide editing.

Although many methods of image editing have good results in changing content, color, texture, etc., when dealing with complex structures, such as fingers, eyes, corners of the mouth and other details, artifacts are often generated. It is still a challenge to improve the artifact generation during image editing of diffusion models.

Latent Space Disentanglement. Extensive research has been conducted on the latent space disentanglement of generative models. β -VAE (Higgins et al. 2017) introduces a hyperparameter β to the KL divergence term, enhancing the gap between prior and posterior distributions to achieve effective disentanglement. Additional unsupervised methods based on statistical features, such as β -TC-VAE (Kim and Mnih 2018), DIP-VAE (Kumar, Sattigeri, and Balakrishnan 2017), and Guided-VAE (Ding et al. 2020), have also demonstrated significant success. Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) have a well-defined latent space to disentanglement.

Some recent efforts have proposed different approaches to improve the disentanglement of GANs from various perspective (Karras, Laine, and Aila 2019; Karras et al. 2020; Ling et al. 2021; Wang et al. 2022). Furthermore, Re-

SeFa (Huang et al. 2023) addresses region-based semantic discovery as a dual optimization problem, achieving semantic disentanglement through a properly defined generalized Rayleigh quotient. Song et al. (2024) use Flow Factorized Representation Learning to achieve more effective structured representation.

Unlike VAES and GANs, the x-space in the diffusion model is considered to be an implicit representation lacking semantics. Diffusion autoencoders (Preechakul et al. 2022) explores the possibility of representation learning of x-space through a semantic encoder. They achieve meaningful encoding and decodability of the input image. Kwon, Jeong, and Uh (2022) further investigate the relationship between the x-space and h-space. Their approaches leverage CLIP to facilitate semantic disentanglement. Zhu et al. (2024) and Park et al. (2023) explore the relationship between x-space and h-space, achieving a certain degree of disentanglement. Haas et al. (2024) identify semantic directions within the h-space. Dalva and Yanardag (2024) use contrastive learning to achieve unsupervised semantic direction discovery. Wu and Zheng (2024) combine content with mask to discover visual concepts. Hu et al. (2024) use transformer-based flow matching to achieve latent space editing.

To sum up, modifying local attributes often results in unintended global changes, as current methods focus narrowly on regional alterations, overlooking broader harmony. Moreover, they typically learn only a limited set of concepts. Achieving diverse regional attribute discovery and precise local editing remains a significant challenge.

Methodology

In this section, we explain the reasons why previous classifier-guided methods alter local attributes while simultaneously affecting global attributes. Subsequently, we introduce our region-based semantic discovery method and a technique for controlling the generation of time-step edits.

Preliminary: Denoising Diffusion Models

Denoising Diffusion Probability Models (DDPM) is a generative model inspired by the diffusion phenomenon in thermodynamic systems. The original diffusion models do not have a latent space but recently Kwon(Kwon, Jeong, and Uh 2022) identified a potential semantic latent space at the bottleneck level of U-Net, referred to as h-space. They reformulated the inverse process of the diffusion model as:

$$x_{t-1} = \sqrt{\alpha_{t-1}}\mathbf{P}_t(\epsilon_t^\theta(x_t)) + \mathbf{D}_t(\epsilon_t^\theta(x_t)) + \sigma_t \mathbf{z}_t. \quad (1)$$

where $\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, ϵ_t^θ is a neural network to predict noise from \mathbf{x}_t . At the same time, $\mathbf{P}_t(\epsilon_t^\theta(x_t))$ and $\mathbf{D}_t(\epsilon_t^\theta(x_t))$ are expressed as:

$$\mathbf{P}_t(\epsilon_t^\theta(x_t)) = \frac{\mathbf{x}_t - \sqrt{1 - \alpha_t}\epsilon_t^\theta(x_t)}{\sqrt{\alpha_t}}. \quad (2)$$

and

$$\mathbf{D}_t(\epsilon_t^\theta(x_t)) = \sqrt{1 - \alpha_{t-1} - \sigma_t^2}\epsilon_t^\theta(x_t). \quad (3)$$

$\sigma_t = \eta\sqrt{(1 - \alpha_{t-1})/(1 - \alpha_t)}\sqrt{1 - \alpha_t/\alpha_{t-1}}$. The range of η is from 0 to 1, where 0 corresponds to DDIM and 1

corresponds to DDPM. Its value represents the degree of randomness. During the image editing process of Asypr, the semantic vector $\Delta\mathbf{h}_t$ is only injected into \mathbf{P}_t without changing \mathbf{D}_t . They proved that the shift caused by $\Delta\mathbf{h}_t$ is offset by the shift \mathbf{D}_t caused by $\Delta\mathbf{h}_t$. Therefore, applying $\Delta\mathbf{h}_t$ on \mathbf{P}_t and \mathbf{D}_t leads to the same output as the original output.

Region-Based Semantic Discovery

We argue that h-space has local Euclidean properties and consistency in different timesteps. This kind of space allows us to propose approaches analogous to the latent disentanglement methods in VAEs and GANs. Therefore, we associate semantic discovery with the matrix of the generative network, thereby achieving semantic discovery that does not require supervision.

Unsupervised Editing through Jacobian. In generative models with explicit latent spaces, a semantic direction can usually maximize the output variations, which means a relation between image space and latent space. In diffusion models, if we assume that h-space is a latent space, then we also hope to find the connection between image space and h-space. Given a time step t and the corresponding \mathbf{x} , our goal is to find a function that establishes a connection between the denoising network $\mathbf{P}_t(\epsilon_t^\theta(x_t))$ and h-space. To achieve this goal, we define the following function:

$$f(h_t) := \epsilon_t^\theta(x_t, h_t). \quad (4)$$

where \mathbf{x}_t is determined image in specific timestep. The process of semantic editing can be described as:

$$x^{\text{edit}} = f(h_t + \alpha\Delta h_t). \quad (5)$$

Inspired by previous work, we relate the direction of maximum semantic change to the Jacobian matrix of the function. The Jacobian matrix is a bridge that plays an important role in attribute editing. The Jacobian matrix is defined as $J_t = \frac{\partial f}{\partial h_t}$. Descending along the direction of the Jacobian matrix will result in the most significant change in the function, potentially identifying a semantic direction. To analyze this Jacobian matrix, we can perform Singular Value Decomposition (SVD):

$$J_t = U_t \Sigma_t V_t^T. \quad (6)$$

The corresponding right singular vectors V_t constitute the orthogonal vector set of the image perturbations in h-space. Finally, our editing method becomes:

$$f(h_t + \alpha\Delta h_t) \approx f(h_t) + \alpha V_t. \quad (7)$$

However, directly computing the Jacobian matrix J_t remains computationally expensive. Fortunately, solving the vector product $J_t^T J_t$ using power iteration is significantly more feasible. The eigenvectors of $J_t^T J_t$ can be accurately approximated to the right singular vectors of J_t . Next, we demonstrate why J_t can represent the perturbation of h-space and is equivalent to the direction of the maximum change of the image. The known image prediction network \mathbf{P}_t is expressed as follows:

$$\mathbf{P}_t(\epsilon_t^\theta(x_t)) = \frac{\mathbf{x}_t - \sqrt{1 - \alpha_t}\epsilon_t^\theta(x_t)}{\sqrt{\alpha_t}}. \quad (8)$$

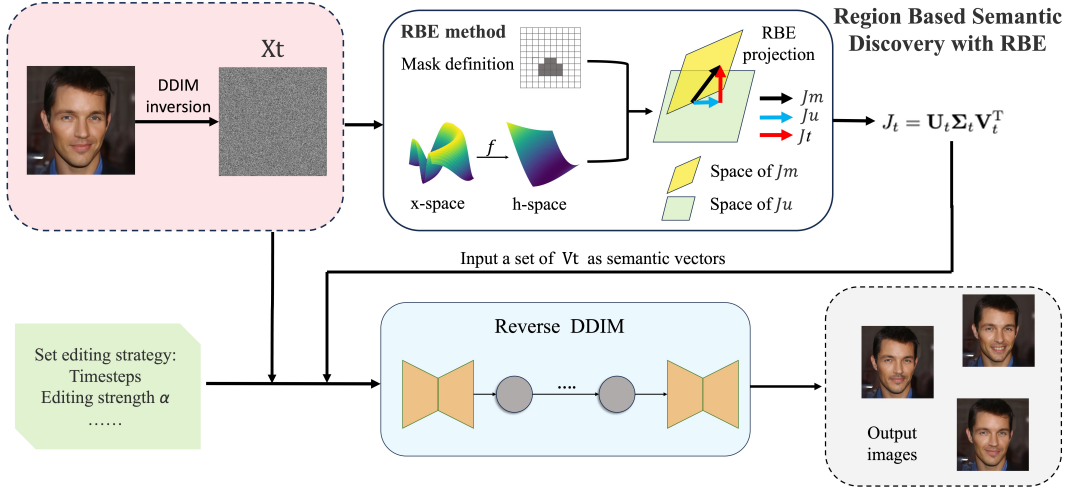


Figure 2: **Overview of our semantic discovery method and editing method.** Firstly, we define mask M and function f , which can be found in our method section. We also use DDIM inversion to precompute x_t and h_t for later use. We use the power iteration method to calculate the Jacobian matrix J_t , and use J_t to calculate the required V_t . Finally, we set the modified timesteps, edit intensity and other parameters, and use DDIM to generate images. All algorithms and specific experimental settings can be found in the appendix.

Under the premise that x_t is fixed, if \mathbf{P}_t is partial derivative with respect to h_t , the result is:

$$\frac{\partial}{\partial \mathbf{h}_t} \mathbf{P}_t(\mathbf{x}_t, \mathbf{h}_t) = -\frac{\sqrt{1-\alpha_t}}{\sqrt{\alpha_t}} \frac{\partial}{\partial \mathbf{h}_t} \epsilon_t^\theta(\mathbf{x}_t, \mathbf{h}_t). \quad (9)$$

which further equals to:

$$\frac{\partial}{\partial \mathbf{h}_t} \mathbf{P}_t(\mathbf{x}_t, \mathbf{h}_t) = -\frac{\sqrt{1-\alpha_t}}{\sqrt{\alpha_t}} J_t. \quad (10)$$

The above derivations show that J_t represents the direction of the maximum value of image change.

Local Mask Definition. To specify the editing area and the non-editing area, we employed two approaches. The first method involves utilizing a pre-trained segmentation network to delineate multiple regions. Subsequently, we select the area of interest as the editing region and apply a mask to the non-editing area, using the Hadamard product to binarize the selected area. The second method consists of directly using a rectangular bounding box to define the area for editing, followed by the application of a mask to the non-editing region in a manner analogous to the previous approach. Compared to the first method, the second is more straightforward. Additionally, for the same region of interest, the area selected using this method is generally larger than that identified by the segmentation network. Overall, the second method tends to facilitate richer changes. Since we expect to find the semantics of the mask area, that is, to find the direction of change in the mask area, we add the mask to the noise prediction network $\epsilon_t(\mathbf{x}_t, \mathbf{h}_t)$:

$$\tilde{\epsilon}_t(\mathbf{x}_t, \mathbf{h}_t) = \epsilon_t^\theta(\mathbf{x}_t, \mathbf{h}_t) \odot \text{mask}. \quad (11)$$

$$J_t^{\text{masked}} = \frac{\partial}{\partial \mathbf{h}_t} \tilde{\epsilon}_t(\mathbf{x}_t, \mathbf{h}_t). \quad (12)$$

Suppressing Non-masked Semantic Variations. Recall that our goal is to maintain the global attributes while allowing for changes in local attributes. The aforementioned algorithm identifies the vector in h-space that corresponds to the maximum change in local attributes. This approach focuses solely on local modifications, but does not address the potential for global alterations. We have previously proved that the change in x-space is equivalent to the noise prediction network $\epsilon_t^\theta(\mathbf{x}_t, \mathbf{h}_t)$. Therefore, our optimization goal can be expressed as:

$$\arg_{\mathbf{h}_t} \max \left(\frac{\partial}{\partial \mathbf{h}_t} \tilde{\epsilon}_t(\mathbf{x}_t, \mathbf{h}_t) \right) \min \left(\frac{\partial}{\partial \mathbf{h}_t} \epsilon_t^\theta(\mathbf{x}_t, \mathbf{h}_t) \right). \quad (13)$$

Here $\epsilon_t^\theta(\mathbf{x}_t, \mathbf{h}_t)$ refers to the noise prediction network in the unmasked region.

As previously discussed, changes in h-space affect the extent of modifications in \mathbf{x} . If a vector h_1 in h-space exhibits the largest change within the masked area but also significantly affects the non-masked area, we can find an alternative vector h_2 in h-space that induces the largest change specifically in the non-masked area. To address this, we need to consider the optimization objective in terms of Equation 13.

The Jacobian matrix J describes how a change in the h-space vector influences changes in \mathbf{x} . Therefore, J_m captures the sensitivity of the masked area, and J_u captures the sensitivity of the non-masked area. By leveraging these Jacobian matrices, the orthogonal projection of h_2 onto the non-masked area direction helps refine h_1 such that it optimizes the changes in the masked area without causing unintended modifications elsewhere. Therefore, we perform the following orthogonal Jacobian projection to achieve the optimization goal:

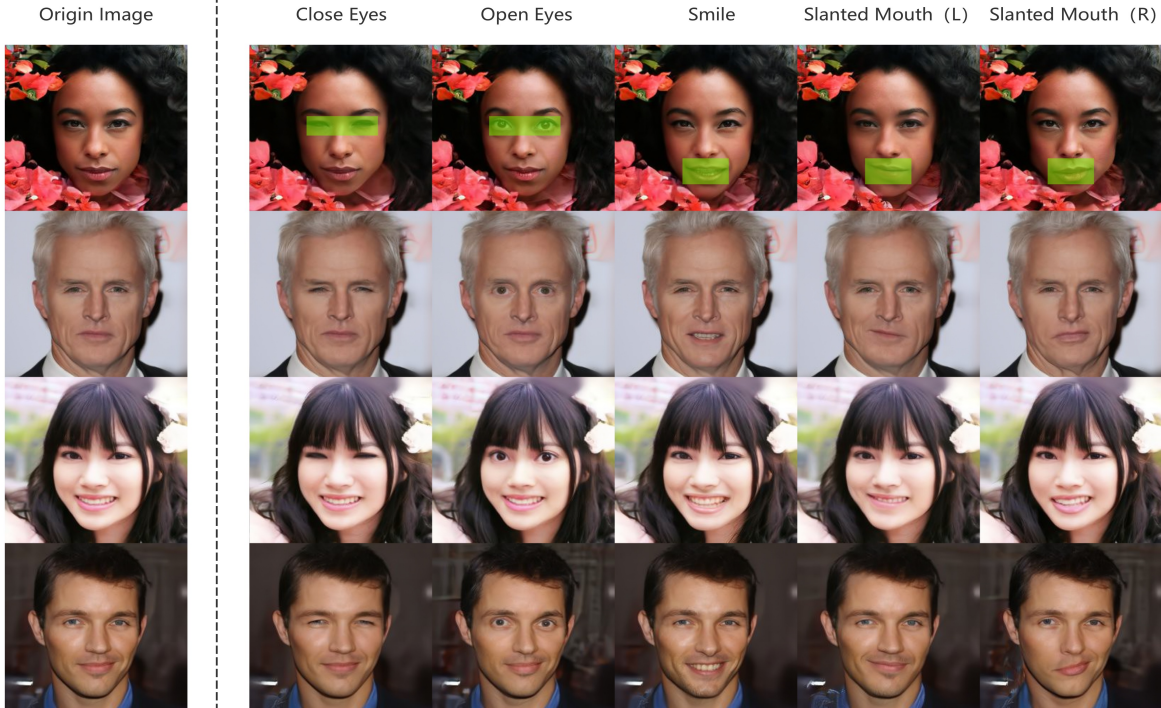


Figure 3: **Qualitative results of our method.** We experimented with the pre-trained DDPM model on the CelebA-HQ dataset with a resolution of 256*256. The leftmost image is the original image. The green box corresponds to the area where we use the mask. The area of the image mask in each column is the same. Please note that we only use the mask during training and do not need to add it during testing. In our experiment, the mask of the same area may find different attribute results. For example, for the mouth, it can be a smile or a slanted mouth in one direction.

$$J = J_m - \frac{J_m \cdot J_u}{J_u \cdot J_u} \cdot J_u. \quad (14)$$

Where J_m and J_u stand for the Jacobians of the masked and unmasked areas, respectively.

Summary

With the above analysis and discussion, we determine the space h-space where we are looking for semantics. We also define a new function f determined by the noise prediction network e_t^θ , the Jacobian matrix and its calculation method, and give a mask adding method. Solve the optimization target Equation 13 by orthogonal projection (Equation 14), so as to effectively obtain the semantic vectors of the target region.

Experiments

Since our method focuses on unsupervised semantic discovery and local image editing of diffusion models, there are not many baselines. We therefore mainly compare our method with supervised semantic discovery methods such as Asyryp, Boundary Diffusion which requires attribute labels as priors, and Semdiff which mainly searches for semantic attributes globally.

methods	FID↓	ID↑	MSE($\times 10^{-4}$)↓	LPIPS↓
close-eyes	11.53	0.8507	5.86	0.0346
open-eyes	11.50	0.8418	6.06	0.0312
smile	14.22	0.8627	5.74	0.0289
Slanted Mouth1	10.08	0.9064	4.51	0.0238
Slanted Mouth2	10.30	0.9052	4.62	0.0231

Table 1: **Quantitative performance** of our method under different attributes editing conditions

Experimental Setup

We conduct extensive experiments on multiple datasets and diffusion models with different structures to demonstrate the effectiveness of our approach. We conduct experiments on the datasets CelebA-HQ (Liu et al. 2015), LSUN-church (Yu et al. 2015), LSUN-bedroom (Yu et al. 2015), and the diffusion model architectures DDPM and iDDPM (Nichol and Dhariwal 2021). For the comparison methods Asyryp and Boundary Diffusion, we use the checkpoints provided by the official. For the method Semdiff, since they did not provide official checkpoints, we trained it exactly as they did and found comparable properties, and then performed a fair comparison with our method. We mainly use Fréchet In-

ception Distance(FID) (Heusel et al. 2017), Mean Squared Error (MSE) (Mathieu, Couprie, and LeCun 2015), Identity loss (ID) (Deng et al. 2019) and LPIPS (Zhang et al. 2018) as objective indicators. We use FID mainly to reflect the quality of image generation. For MSE, we not only measure the global MSE, which represents the overall change of the image but also measure the MSE of the masked area and the non-masked area, which represents the local change and Changes other than local areas. We use ArcFace net to measure the face identity similarity (ID) metric. This metric describes how similar the identity of the person after the modification is to the original image, reflecting the quality of the editing. The overall structure of our experiment is as follows. We randomly select 500 images on CelebA-HQ for testing. In the first part, we show the semantic directions found by our method in the CelebA-HQ dataset. In the second part, we compare our method with three baseline methods. In the third part, we compare the editing results of our method with Semdiff in local and global regions. In the fourth part, We summarize the experimental results and analyze the advantages and disadvantages. All our experiments can be performed on a single RTX 3090 GPU.



Figure 4: More semantic editing results

Main Results of Attribute Editing

In this subsection, we present the overall performance of our approach. Our specific experimental process is as follows: First, add a mask to a specific area on the image according to the Equation 11. Then we compute the Jacobian matrix of the mask and non-mask areas through Equation 12. After that, we use Equation 14 to get semantic vectors J . Finally, We apply the Jacobian matrix in timesteps of reverse process to obtain the results.

We mainly select the eyes region and the mouth region to find semantics. In Figure 3, we show the semantics found by our method in the eye and mouth regions. In the eye semantics, there are eyes open and eyes closed; in the mouth attributes, there are smiles, and slanting to the left and right. Interestingly, we also found some semantics for specific groups of people, such as editing the beard area for men, and changing the hair of people with long curly hair to straight hair. We show the above in Figure 4 and put more semantic editing results of other regions in the Appendix.

We further test the quantitative results shown in Table 1.




Figure 5: **Qualitative results.** Our method has the best results on the overall structure and details.

Our results show that while achieving good editing results, the image quality indicators FID, MSE, LPIPS, and the identity change ID of the person have achieved extremely high levels.

methods	FID↓	ID↑	MSE($\times 10^{-4}$)↓	LPIPS↓
Asyrrp	54.04	0.4892	16.60	0.2172
BoundaryDiffusion	53.04	0.5671	15.14	0.1765
Semdiff	<u>21.77</u>	<u>0.7224</u>	5.23	<u>0.0424</u>
ours	14.22	0.8627	<u>5.74</u>	0.0289

Table 2: **Quantitative comparison** between different local editing approaches and ours. We compare the editing effects of the attribute 'Smile'. **Bold** and underline represent the best and second-best performance, respectively.

Comparison against Other Baselines

In this section, we compare our method with the state-of-the-art methods both qualitatively and quantitatively. We compare with Asyrrp, Boundary diffusion, and Semdiff, which are current methods for semantic discovery and editing using diffusion models. Although our method finds local semantics, the smile attribute exists in the semantics of the mouth. And for fair comparison, their method has an open source smile attribute checkpoint, so we used the smile attribute that is unfavorable to ours for comparison. But we are surprised to find that even so, our unsupervised method is better than the supervised method, and has a great improvement effect than the unsupervised semdiff.

In Figure 5, we show the qualitative results. As can be seen from Figure 5, Asyrrp will bring about a great change in the image structure, and for images with complex character

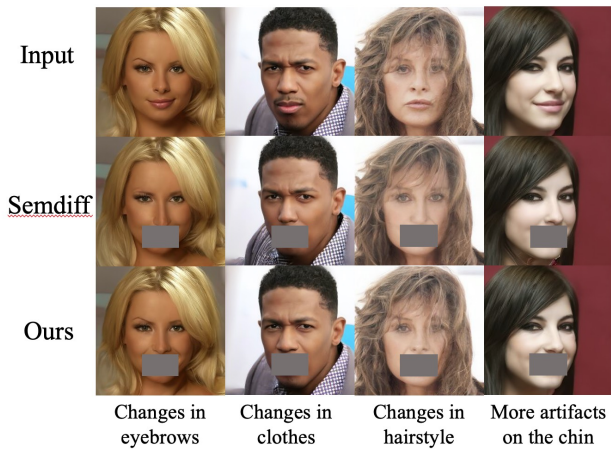


Figure 6: **Exemplary comparison of Semdiff and our RBE.** We compare the remaining parts after editing the mask of the region. Our method has a clear advantage in image preservation outside the region.

postures, it will bring more serious distortion results. Boundary Diffusion is better than Asyrp, but it still has a more serious impact on character identity, background and character details. The effect of Semdiff is intuitively good, but because this method only considers local semantic attributes, it does not consider the influence of other global parts. Therefore, the details of the characters, such as the noses in the first and third images, and the beards in the second male image, have more serious changes, which is not desirable. The last row is the result of our method. Our method not only ensures the invariance of the overall image (character identity, background, etc.), but also ensures the invariance of character details (clothes, nose, beard, etc.).

Table 2 presents a quantitative evaluation of baselines and our method using metrics such as FID, ID, MSE, and LPIPS. Our method delivers best results across most metrics, particularly excelling in FID and ID, which assess image quality and person identity, respectively. MSE reflects image quality on the one hand and editing intensity on the other. Asyrp and Boundary Diffusion exhibit significantly higher MSE values due to larger structural changes. Although our method and Semdiff have slightly higher in MSE values, it does not necessarily indicate inferiority, as other metrics should also be considered for a comprehensive assessment.

Changes of Local and Non-local Regions

To demonstrate the superiority of our method in local region editing, we compared the degree of change in local and non-local regions with Semdiff. We compared the MSE inside and outside the mask region when editing the eyes and mouth. Compared with Semdiff, when the overall MSE scale is close, our method is more inclined to edit the inner area of the mask, and the changes to the outer area of the mask are less than Semdiff. The qualitative results in Figure 7 also prove this point: when editing the masked region, Semdiff will cause changes in the rest of the image, while our method

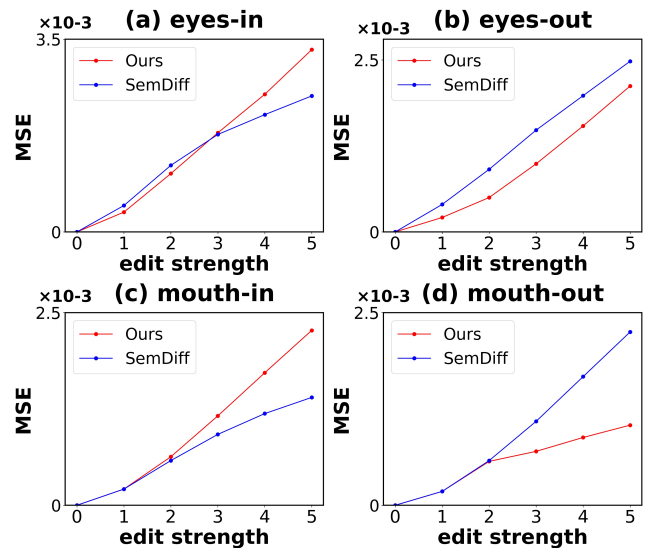


Figure 7: **Quantitative results of pixel change in mouth editing and eyes editing.** Our method is shown in red and Semdiff is shown in blue. "In" and "out" refer to the region of interest and its surroundings, respectively. For effective local editing, a higher change within the region of interest ("in") and a lower change in the surrounding area ("out") are expected.

maintains the non-masked region well.

Discussion

We have demonstrated the state-of-the-art performance of our method in local attribute editing. More experimental results, including other attribute edits and different datasets, are available in the Appendix. First, since our method is aimed at local editing, when editing global attributes, the mask area needs to be expanded to a large enough area. Under this premise, our method will not be significantly better than previous research on global attribute editing. Moreover, when the mask area is very small (such as an eye, a tooth, etc.), our method cannot find significant semantic directions. Our future research will be directed towards extending our method to more fine-grained areas and global attributes.

Conclusion

To explore the latent potential within diffusion models, we propose a novel method utilizing the Jacobian projection technique for precise semantic editing of pre-trained models without additional training. Our approach allows for the identification and manipulation of semantic attributes in local regions by projecting the Jacobian of the masked area into a low-dimensional subspace orthogonal to the non-masked regions. This technique significantly enhances control over local image attributes while preserving overall image harmony. Furthermore, our method outperforms existing supervised approaches, especially in specific tasks such as facial attribute editing, demonstrating its substantial potential to advance unsupervised image editing techniques.

Acknowledgments

This work was supported in part by the Fundamental Research Funds for the Central Universities (No.2022JBZY043), Fundamental Research Funds for the Central Universities (No.2022XKRC015), National NSF of China (No.62120106009, No.62372033) and Beijing Natural Science Foundation (L242021).

References

- Couairon, G.; Verbeek, J.; Schwenk, H.; and Cord, M. 2022. Diffedit: Diffusion-based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*.
- Dalva, Y.; and Yanardag, P. 2024. Noiseclr: A contrastive learning approach for unsupervised discovery of interpretable directions in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24209–24218.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4690–4699.
- Ding, Z.; Xu, Y.; Xu, W.; Parmar, G.; Yang, Y.; Welling, M.; and Tu, Z. 2020. Guided variational autoencoder for disentanglement learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7920–7929.
- Goodfellow, I. J.; et al. 2014. Generative Adversarial Nets. In *NIPS*.
- Haas, R.; Huberman-Spiegelglas, I.; Mulayoff, R.; Graßhof, S.; Brandt, S. S.; and Michaeli, T. 2024. Discovering interpretable directions in the semantic latent space of diffusion models. In *2024 IEEE 18th International Conference on Automatic Face and Gesture Recognition (FG)*, 1–9. IEEE.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Higgins, I.; Matthey, L.; Pal, A.; Burgess, C. P.; Glorot, X.; Botvinick, M. M.; Mohamed, S.; and Lerchner, A. 2017. beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR (Poster)*, 3.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Hu, V. T.; Zhang, W.; Tang, M.; Mettes, P.; Zhao, D.; and Snoek, C. 2024. Latent space editing in transformer-based flow matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2247–2255.
- Huang, N.; Tang, F.; Dong, W.; Lee, T.-Y.; and Xu, C. 2023. Region-aware diffusion for zero-shot text-driven image editing. *arXiv preprint arXiv:2302.11797*.
- Huang, Y.; Huang, J.; Liu, Y.; Yan, M.; Lv, J.; Liu, J.; Xiong, W.; Zhang, H.; Chen, S.; and Cao, L. 2024. Diffusion model-based image editing: A survey. *arXiv preprint arXiv:2402.17525*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8110–8119.
- Kim, G.; Kwon, T.; and Ye, J. C. 2022. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2426–2435.
- Kim, H.; and Mnih, A. 2018. Disentangling by factorising. In *International conference on machine learning*, 2649–2658. PMLR.
- Kumar, A.; Sattigeri, P.; and Balakrishnan, A. 2017. Variational inference of disentangled latent concepts from unlabeled observations. *arXiv preprint arXiv:1711.00848*.
- Kwon, M.; Jeong, J.; and Uh, Y. 2022. Diffusion models already have a semantic latent space. *arXiv preprint arXiv:2210.10960*.
- Ling, H.; Kreis, K.; Li, D.; Kim, S. W.; Torralba, A.; and Fidler, S. 2021. Editgan: High-precision semantic image editing. *Advances in Neural Information Processing Systems*, 34: 16331–16345.
- Liu, H.; Tan, Z.; Tan, C.; Wei, Y.; Wang, J.; and Zhao, Y. 2024. Forgery-aware adaptive transformer for generalizable synthetic image detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10770–10780.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Mathieu, M.; Couprie, C.; and LeCun, Y. 2015. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*.
- Mou, C.; Wang, X.; Song, J.; Shan, Y.; and Zhang, J. 2023. Dragondiffusion: Enabling drag-style manipulation on diffusion models. *arXiv preprint arXiv:2307.02421*.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, 8162–8171. PMLR.
- Park, Y.-H.; Kwon, M.; Choi, J.; Jo, J.; and Uh, Y. 2023. Understanding the latent space of diffusion models through the lens of riemannian geometry. *Advances in Neural Information Processing Systems*, 36: 24129–24142.

- Preechakul, K.; Chatthee, N.; Wizadwongsa, S.; and Suwanajakorn, S. 2022. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10619–10629.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2): 3.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, Y.; Keller, A.; Sebe, N.; and Welling, M. 2024. Flow factorized representation learning. *Advances in Neural Information Processing Systems*, 36.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Tan, C.; Tao, R.; Liu, H.; Gu, G.; Wu, B.; Zhao, Y.; and Wei, Y. 2024. C2P-CLIP: Injecting Category Common Prompt in CLIP to Enhance Generalization in Deepfake Detection. *arXiv preprint arXiv:2408.09647*.
- Wang, T.; Zhang, Y.; Fan, Y.; Wang, J.; and Chen, Q. 2022. High-fidelity gan inversion for image attribute editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11379–11388.
- Wu, A.; and Zheng, W.-S. 2024. Factorized Diffusion Autoencoder for Unsupervised Disentangled Representation Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 5930–5939.
- Yu, F.; Seff, A.; Zhang, Y.; Song, S.; Funkhouser, T.; and Xiao, J. 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3836–3847.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhang, Y.; Wei, Y.; Jiang, D.; Zhang, X.; Zuo, W.; and Tian, Q. 2023. Controlvideo: Training-free controllable text-to-video generation. *arXiv preprint arXiv:2305.13077*.
- Zhu, Y.; Wu, Y.; Deng, Z.; Russakovsky, O.; and Yan, Y. 2024. Boundary guided learning-free semantic control with diffusion models. *Advances in Neural Information Processing Systems*, 36.