

Accelerating Text-to-Image Editing via Cache-Enabled Sparse Diffusion Inference

Zihao Yu¹, Haoyang Li¹, Fangcheng Fu¹, Xupeng Miao², Bin Cui^{1,3}

¹ School of Computer Science & Key Lab of High Confidence Software Technologies (MOE), Peking University

² Carnegie Mellon University

³ Institute of Computational Social Science, Peking University (Qingdao), China

{zihao.yu, ccchengff, bin.cui}@pku.edu.cn, lihaoyang@stu.pku.edu.cn, xupeng@cmu.edu

Abstract

Due to the recent success of diffusion models, text-to-image generation is becoming increasingly popular and achieves a wide range of applications. Among them, text-to-image editing, or continuous text-to-image generation, attracts lots of attention and can potentially improve the quality of generated images. It's common to see that users may want to slightly edit the generated image by making minor modifications to their input textual descriptions for several rounds of diffusion inference. However, such an image editing process suffers from the low inference efficiency of many existing diffusion models even using GPU accelerators.

To solve this problem, we introduce Fast Image Semantically Edit (FISEdit), a cached-enabled sparse diffusion model inference engine for efficient text-to-image editing. The key intuition behind our approach is to utilize the semantic mapping between the minor modifications on the input text and the affected regions on the output image. For each text editing step, FISEdit can 1) automatically identify the affected image regions and 2) utilize the cached unchanged regions' feature map to accelerate the inference process. For the former, we measure the differences between cached and ad hoc feature maps given the modified textual description, extract the region with significant differences, and capture the affected region by masks. For the latter, we develop an efficient sparse diffusion inference engine that only computes the feature maps for the affected region while reusing the cached statistics for the rest of the image. Finally, extensive empirical results show that FISEdit can be 3.4× and 4.4× faster than existing methods on NVIDIA TITAN RTX and A100 GPUs respectively, and even generates more satisfactory images.

Introduction

The recent months have witnessed the astonishing improvement of diffusion models. Text-to-image generation is one of the most popular applications of diffusion models, which involves generating realistic images based on textual descriptions. While diffusion models have been shown to generate high-quality images with good coverage of the data distribution, they require substantial computational power to generate high-quality images. Thus, it is an emerging research area to make the generation process of diffusion models more effective and efficient.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Our target scenario. In real-world cases, after an image has been generated, it is common that the user may be unsatisfied with some specific regions of the generated image and desire to make subsequent **semantic edits**. Figure 1 demonstrates a concrete example. The user may wish to interactively edit the image by making minor modifications to the textual description. These modifications involve adding specifications (e.g., adding a “river” or the “fireworks”) or replacing words/phrases (e.g., replacing the “river” with “snow mountain”). And in each round of editing, the user expects to make changes to only a small, specific region (affected region) and keep the rest (non-affected regions) untouched.

Challenges. However, we find that the current implementations fall short in supporting such a semantic editing task effectively and efficiently. First, due to the diversity nature of diffusion models, existing methods tend to make significant changes to the image. Although there are applications that allow users to provide masks to control the in-painting regions, it would lead to poor user experience. Second, even though only a small region is expected to be modified, existing works typically follow the vanilla inference process of diffusion models, which generates the entire image from scratch, without taking the spatial property into account. Undoubtedly, this is a waste of time and computational resources since the generation of non-affected regions is pointless.

In essence, to support the semantic editing task effectively and efficiently, there involves two technical issues: (1) *how to automatically and accurately detect the affected region to be revised given the original image and the modified description* and (2) *how to focus on the target region during the generation while skipping the non-affected regions?*

Summary of contributions. In this paper, we develop a brand new framework for the semantic editing task. The technical contributions of this work are summarized as follows.

- We introduce Fast Image Semantically Edit (FISEdit), a framework for minor image editing with semantic textual changes. FISEdit is a direct application to the inference process of pre-trained diffusion models and does not need to train new models from scratch. To support the interactive semantic editing task, FISEdit caches the feature maps of the latest generation process and addresses the

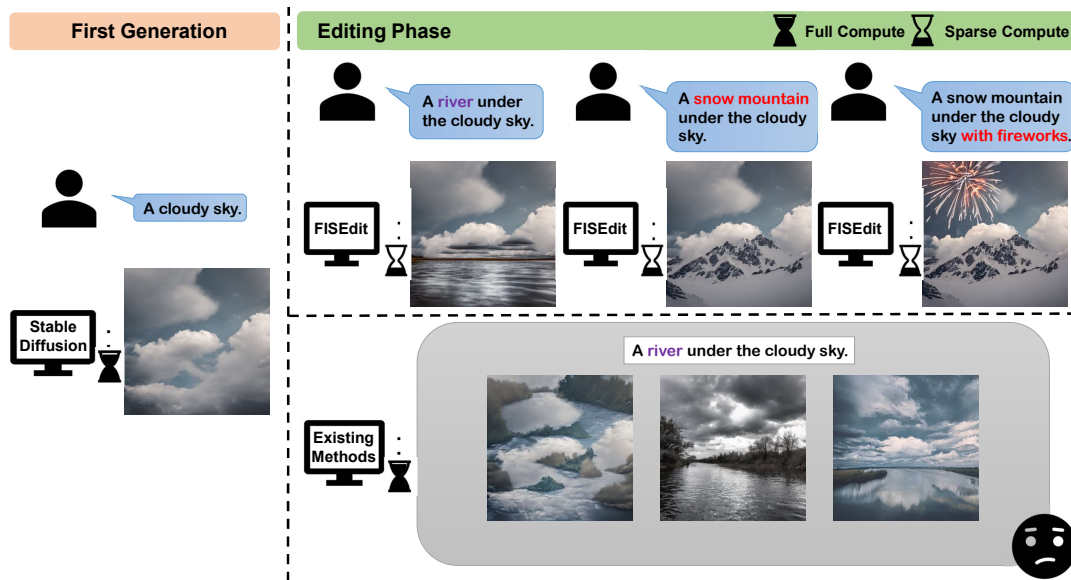


Figure 1: A real example of the user’s interaction with FISEdit and existing methods.

two technical issues respectively.

- We present a mask generation algorithm to capture the specific region for editing. To be specific, we quantify the distance between the cached and ad hoc latents given the modified textual description. Subsequently, we extract the region displaying substantial differences.
- Based on the detected region and the cached feature maps, we propose a sparse updating approach to minimize unnecessary computational burden. More precisely, we develop an end-to-end sparse diffusion inference mechanism that selectively computes feature maps solely for the affected region and re-uses the cached statistics for the non-affected regions.
- Compared with existing text-based image editing works, FISEdit can generate competitive high-quality images that are consistent to the semantic edit. Furthermore, FISEdit outperforms existing works in terms of generation efficiency by $4.4\times$ on NVIDIA TITAN RTX and $3.4\times$ on NVIDIA A100.

Related Works

Recent advances in diffusion models have enabled state-of-the-art image synthesis (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2021; Lu et al. 2022a; Balaji et al. 2022; Ramesh et al. 2022; Saharia et al. 2022; Huang et al. 2022; Bianchi et al. 2022; Ramasinghe et al. 2021), including image editing and text-to-image generation. Compared to other generative models such as GANs (Karras, Laine, and Aila 2019; Goodfellow et al. 2020) and VAEs (Kingma and Welling 2014), diffusion models have been shown to generate high-quality images with good coverage of the data distribution. However, they necessitate a substantial amount of computational power for both training and high-quality image generation, which makes them expensive and challenging to deploy in production systems.

Image Editing Methods

The user’s requirement for image modification can be satisfied by **image editing models** (Zhu et al. 2016; Abdal, Qin, and Wonka 2019, 2020; Zhu et al. 2020; Choi et al. 2021; Kim, Kwon, and Ye 2022; Patashnik et al. 2021; Rombach et al. 2022; Brooks, Holynski, and Efros 2022; Wei et al. 2023) and **controllable image generation methods** (Meng et al. 2022b; Park et al. 2019; Nichol et al. 2022; Hertz et al. 2022; Couairon et al. 2023; Orgad, Kawar, and Belinkov 2023; Parmar et al. 2023; Feng et al. 2023).

Image editing models can accurately modify images, but they require to train new diffusion models, which is computationally expensive and time-consuming. Furthermore, most of the existing models are difficult to control the affected region without a user-drawn mask. We only compare our method with InstructPix2Pix (Brooks, Holynski, and Efros 2022) and stable diffusion in-painting models, since they have shown the best editing effect yet.

Controllable image generation methods are able to edit image using existing pre-trained weights, but they always require equal or even more computations to detect affected regions and meet textual change. For example, DIFFEdit (Couairon et al. 2023) need extra reverse diffusion steps to calculate difference mask, and pix2pix-zero (Parmar et al. 2023) fails to edit high-resolution images due to its overuse of memory.

Diffusion Model Acceleration

As diffusion model training and inference can be computationally expensive and time-consuming, especially when dealing with high-resolution images, various methods have been explored to accelerate the training and inference of diffusion models. Firstly, many recent works (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2021; Lu et al.

2022a,b; Watson et al. 2021; Meng et al. 2022a; Kong and Ping 2021; Xiao, Kreis, and Vahdat 2022; Salimans and Ho 2022) have successfully reduced the number of iteration steps required for diffusion model inference while maintaining or improving the quality of the generated samples. Moreover, some deep learning frameworks such as TensorFlow (Abadi et al. 2016) and PyTorch (Paszke et al. 2019) speed up diffusion models by optimizing GPU kernels and paralleling computation streams. All these research directions are orthogonal to ours and can be integrated with FISEdit.

Another way to accelerate diffusion models is reducing computation by exploiting the sparsity of the input data or the model parameters. There are several techniques (Dong et al. 2017; Liu et al. 2018; Ren et al. 2018; Liu et al. 2019; Child et al. 2019; Bolya et al. 2023; Han et al. 2023; Bi et al. 2023; Zhang et al. 2023; Yi et al. 2022) available for implementing sparse computation. However, directly applying existing techniques such as tiled sparse convolution (Ren et al. 2018) fails to achieve considerable speedup and image quality in our scenario.

In particular, recent works have successfully applied sparse computation to generative models (Bolya and Hoffman 2023; Li et al. 2022). SIGE (Li et al. 2022) is a general-purpose sparse inference framework for generative models, which integrates sparse convolution techniques into image editing tasks and makes sufficient use of data sparsity. However, SIGE (Li et al. 2022) faces limitations when applied to Stable Diffusion (Rombach et al. 2022) models due to its significant additional memory overhead requirement and inability to support textual changes. Therefore, we will not include SIGE as a baseline in the experimental part. In comparison to SIGE, our proposed framework supports more efficient sparse kernels and addresses the limitations encountered by SIGE, including batched input restrictions, intermediate activation management and fine-grained mask generation designed to detect areas affected by semantic changes.

Method

Figure 2 shows the overview structure of FISEdit. This method mainly derives from our observation that by sparsely calculating the crucial region of the feature maps in U-Net according to a fine-grained mask, we can effectively maintain the original image structure as well as get close to the new prompt semantics. Meanwhile, the sparsity of the computation could also benefit a lot to the model efficiency. Below, we first introduce a method to detect affected areas and generate difference mask for the subsequent sparse computation. Then we present our improved implementation of sparse computation algorithms, along with a non-trivial image editing pipeline.

Fine-Grained Mask Generation

Difference masks are used in image editing tasks to index changed regions, but existing models have sub-optimal performance when the user cannot draw a precise mask. To address this, we propose a mask generation method, which automatically identifies the image areas that require substantial change according to the modification of textual description.

Target Area Capture. To detect the affected areas in feature maps, we observe the feature maps in each step of the Stable Diffusion (Rombach et al. 2022) inference process and compare them using two similar textual prompts. Unfortunately, even though we have controlled the random factors, the changed regions remain large and unpredictable. Drawing inspiration from Prompt-to-Prompt (Hertz et al. 2022), these irregular spatial changes can be controlled by enhancing cross-attention (Lin et al. 2022) module. Specifically, we cache and preserve part of the cross-attention map corresponding to the unchanged text prompts in the calls with new description, and then the changed regions can be structured and easily located through few denoise steps. As demonstrated in Figure 3, the difference caused by textual change can be observed clearly in step 5 to 10 and convert to a mask after finding the proper threshold ϵ via the following formula:

$$\begin{aligned} diff &:= \text{Normalize} \left(\sum_{t=t_1}^{t_2} |X_t - Y_t| \right), \quad t_1, t_2 \in [1, 10] \\ A &:= \{i | diff_{pixel_i} < \epsilon\}, B := \{j | diff_{pixel_j} \geq \epsilon\} \\ N_1 &:= \text{Card}(A), N_2 := \text{Card}(B) \\ \arg \max_{\epsilon \in [0, 1]} OTSU(\epsilon) &:= \frac{N_1 N_2}{(N_1 + N_2)^2} \left(\frac{1}{N_1} \sum_{i \in A} diff_{pixel_i} \right. \\ &\quad \left. - \frac{1}{N_2} \sum_{j \in B} diff_{pixel_j} \right)^2 \end{aligned}$$

Consequently, we no longer need extra diffusion steps to calculate the mask as our previous work did (e.g., DIFFEdit (Couairon et al. 2023)).

Sparsity Analysis. We evaluate our method by generating mask for the human-written dataset used by Instruct-Pix2Pix (Brooks, Holynski, and Efros 2022), where each sample in the dataset contains an original and manually modified caption for an image. We find that the average edited size for these samples is less than 30%, reflecting an astonishing fact that most image regions can be reused and numerous calculations can be potentially saved. However, when the feature map is compressed to low-resolution, sparsity decreases correspondingly. Consequently, we adopt the sparse computation for high-resolution inputs only, and forgo those with low-resolution.

Sparse Computation

In order to accelerate diffusion models and control the semantic changes, we focus on the U-Net (Ronneberger, Fischer, and Brox 2015) modules since they are the most time-consuming and influential components in the image generation process. The U-Net architecture is composed of several attention and ResNet (He et al. 2016) modules. Based on the observation in the previous section, it is practicable to perform sparse calculations on a small proportion of the whole image in the above modules. Below, sparse optimization on these modules will be discussed respectively.

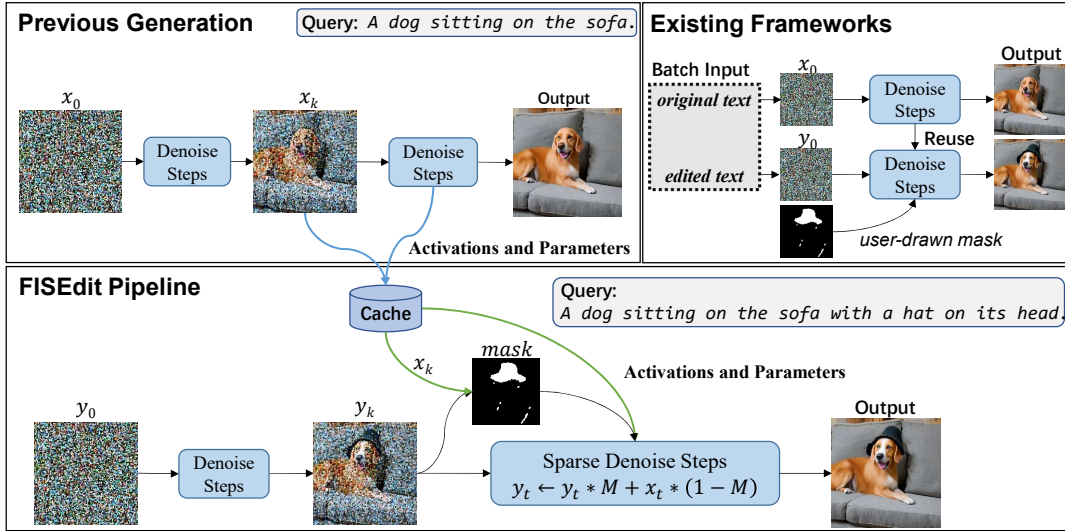


Figure 2: Overview structure of FISEdit. When a query arrives, our system first executes k denoise steps, and then generates a difference mask according to the output latents of k steps. In the remaining denoise steps, the pre-computed results (activations and parameters) of each layers in U-Net will be reused according to the mask and the feature maps will be computed sparsely. Compared to existing frameworks which leverage the batched inputs, we collect and cache the results of previous generation to avoid redundant computation, and use the mask to control as well as accelerate the new T2I generation process at U-Net level.

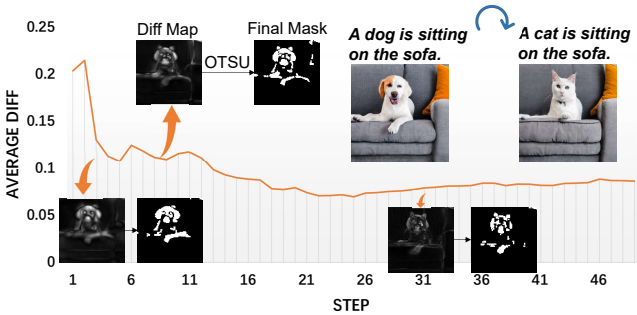


Figure 3: Variation of latent difference with iteration steps.

Adaptive Pixel-Wise Sparse Convolution. Tiled sparse convolution algorithm was proposed in SBNets (Ren et al. 2018) to handle sparse data within ResNet modules. However, as we mentioned in Section , sparsity will decrease when processing low-resolution feature maps, and this is particularly obvious in tiled sparse convolution layers. Therefore, directly applying existing technique fails to achieve considerable speedup and image quality in our scenario.

Given that the efficiency of this algorithm is largely contingent on the distribution of difference masks and the size of split blocks, the constant block-gather strategy used in SBNets doesn’t seem plausible. Therefore, it is essential to select a compatible block size with the given mask. We introduce Adaptive Pixel-wise Sparse Convolution (APSC) to ensure tiled sparse convolution achieves desired effect when dealing with capricious mask. We denote (h, w) , $N_{h,w}^m$ as the block’s spatial size and the number of blocks gathers us-

ing given mask m respectively, and assume the kernel size is (h_k, w_k) . We search for the optimal division strategy heuristically via the following approximation:

$$\arg \min_{h,w \in \{2,4,8,16,32\}} f(h, w) := (h-h_k+1) * (w-w_k+1) * N_{h,w}^m$$

All the variables mentioned above can be determined during the preprocessing stage, ensuring that APSC does not introduce any additional computational costs to model inference.

Approximate Normalization. Since some normalization layers such as group normalization (Wu and He 2018) entail calculating the mean and variance of the feature map in spatial dimension, it is infeasible to use sparse feature map in these layers. Inspired by batch normalization (Ioffe and Szegedy 2015), which uses the mean and variance during training for model inference, it is plausible to use approximate methods in normalization layer. With the assumption that a tiny textual modification will not affect the distribution of the whole image, we are able to cache the mean and variance computed in the previous image generation and directly reuse them to facilitate the sparse inference process. Therefore, without the reliance on the statistics of the input feature map, sparse computation is available in all normalization layers.

Approximate Attention. Recent diffusion models have incorporated self-attention (Vaswani et al. 2017) and cross-attention (Lin et al. 2022) modules to enhance the image quality and the correlation between image and text in text-to-image tasks, resulting in computational bottlenecks. Readers might suspect whether it is possible to utilize the sparsity of edited areas in attention layers, since the whole feature

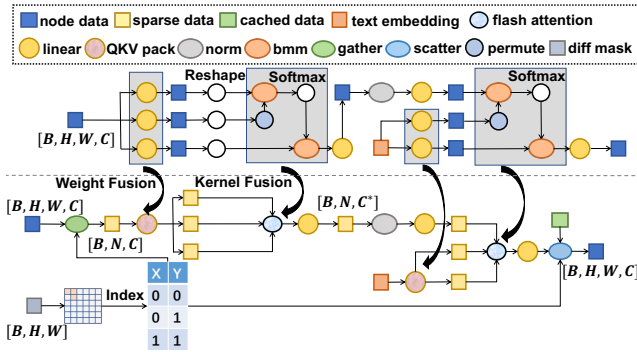


Figure 4: Overall improvement of the attention module.

map is required to be used in self-attention layers in order to learn the long-term relevance among pixels. Intuitively, if we only calculate the results using non-zero regions in difference mask, the output will be inevitably deviated from the accurate result. Fortunately, such precision errors can be controlled to a tolerable range through some accuracy compensation measures. Specifically, we only apply sparse computation in the self-attention layers with a high-resolution input feature map. In these layers, non-zero regions account for only a small proportion of the entire feature map, and their influence on the unaffected area is consequently negligible, particularly after the SoftMax layer.

Figure 4 shows the overview of our implementation of attention modules. Masked regions are collected before being input into the subsequent self-attention and cross-attention layers, utilizing sparse data exclusively within the attention layers to avoid unnecessary data movement.

Cache-Based Editing Pipeline

In order to support sparse computation, we have to store activation and parameters in the prior inference procedure, and load them in the following image generations. However, for the diffusion models, it’s imperative to store and load activations for all iteration steps (e.g., 50 for DDIM (Song, Meng, and Ermon 2021) scheduler and 10 for DPM-Solver++ (Lu et al. 2022b)), implying that there will be a significant amount of extra memory storage. For instance, when we edit an image with a 768×768 resolution, the memory storage for each iteration step is up to 4.2G, culminating in a total memory storage of up to 210 GB for 50 steps. Obviously, additional devices (e.g., CPU, HDD) are required for storage, and dealing with the movement of such a substantial amount of data during inference also presents a formidable challenge.

As depicted in Figure 2, existing image semantic editing methods operate under an inefficient working mode, which relies on batch input of the text before and after modification, in order to avoid the storage and loading of intermediate activation. Unfortunately, this mechanism involves a significant amount of repetitious computation. In fact, data movement and kernel computation are asynchronous on GPU, so the data transfer time can be overlapped by kernel computation. However, as our method dramatically reduces

the computational time, data transmission will inevitably become a new bottleneck. To resolve this impasse, we introduce a cache-based editing pipeline to support asynchronous data movement and distributed storage. Concretely, our system automatically moves stored data according to the memory usage of CPU and GPU, making the data load and storage completely asynchronous. Besides, we reuse the memory of tensors with same shapes during the inference so that memory usage can be lowered.

Experiments

We implement our system based on the HuggingFace’s diffusers¹, which is a generic framework for training and inference of diffusion models. We clone this project and integrate it with our self-developed sparse inference engine Hetu² (Miao et al. 2022c,a,b).

Baselines. We compare our method with the following baselines, and all the implementations for them can be found in HuggingFace’s diffusers. All the baselines and our method use the pre-trained weights and configurations of the Stable Diffusion v2 model³, and more details about our evaluation configurations can be found in our repository⁴.

- *Vanilla stable diffusion text-to-image pipeline (SDTP)*. A simple attempt to edit image is to fix the internal randomness (e.g., the initial latent) and regenerate using the edited text prompt.
- *Stable diffusion in-painting pipeline (SDIP)*. A variant of stable diffusion models, aiming at painting from a draft and text description. User-drawn masks are required in this model, and we share the difference mask generated by our method with it in order to make a fair comparison.
- *Prompt-to-Prompt* (Hertz et al. 2022). A text-driven image editing method based on cross-attention control.
- *SDEdit* (Meng et al. 2022b). An image editing method based on a diffusion model generative prior, which synthesizes realistic images by iteratively denoising through a stochastic differential equation.
- *InstructPix2Pix* (Brooks, Holynski, and Efros 2022). A model that edits images according to user-written instructions instead of descriptions.
- *DIFFEdit* (Couairon et al. 2023). An image semantic editing method with automatic mask generation algorithm.
- *PPAP* (Tumanyan et al. 2023). A semantic editing framework based on self-attention optimizations.
- *Pix2Pix-Zero* (Parmar et al. 2023). An image-to-image translation method that can preserve the content of the original image without manual prompting.

Datasets. We select LAION-Aesthetics (Schuhmann et al. 2022) as the evaluation dataset, and follow the preprocessing method of InstructPix2Pix. The processed dataset⁵ consists

¹<https://github.com/huggingface/diffusers>

²<https://github.com/PKU-DAIR/Hetu/>

³<https://huggingface.co/stabilityai/stable-diffusion-2>

⁴<https://github.com/Hankpipi/diffusers-hetu>

⁵<http://instruct-pix2pix.eecs.berkeley.edu/>

of 454,445 examples, each example contains a text and an edited one.

Metrics. Following prior works (Brooks, Holynski, and Efros 2022; Meng et al. 2022b), we use four metrics to evaluate the quality of image editing:

- *Image-Image similarity.* Cosine similarity of the original and edited CLIP (Radford et al. 2021) image embeddings, which reflects how much the edited image agrees with the input image.
- *Directional CLIP similarity* (Gal et al. 2022). It represents how much the change in text agrees with the change in the images.
- *Fréchet Inception Distance (FID)* (Heusel et al. 2017). A metric compares the distribution of generated images with the distribution of a set of real images. Due to the absence of standard images in the dataset we used, we regard the images generated by vanilla stable diffusion text-to-image model (prompt with edited text) as the standard image set.
- *CLIP Accuracy* (Hessel et al. 2021). The percentage of instances where the edited text has a higher similarity to the target image, as measured by CLIP (Radford et al. 2021), than to the original source image, indicating how much the target image agrees with the edited text.

Main Results

Image Quality. Our main target in image quality is to achieve equal effect as baselines since we have reduced the computational complexity. We show the quantitative results in Figure 5 and the qualitative results in Figure 6. We find that image-image similarity is competing with CLIP accuracy and directional CLIP similarity, which means that increasing the degree to which the output images correspond to a desired text will reduce their consistency with the input image. In order to show the tradeoff between two metrics, we have varied representative parameters for the methods. It can be inferred from these metrics that our method can better locate and modify changing areas as well as reflect the meaning of the text.

Model Efficiency. We present efficiency comparison results in Table 3. We observe that all baselines struggle to generate or edit images efficiently since they cannot leverage the sparsity to avoid unnecessary computations. Worse, InstructPix2Pix (Brooks, Holynski, and Efros 2022) even encounter both computation and memory bottlenecks. In contrast, our method can accurately edit the details of image, and also reduces the computation of diffusion models by up to $4.9\times$ when the image size is 768×768 and the edit size is 5%. Eventually, we accelerate text-to-image inference by up to $4.4\times$ on NVIDIA TITAN RTX and $3.4\times$ on NVIDIA A100 when the edit size is 5%.

Ablation Study

Table 1 shows the results of ablation studies conducted on TITAN RTX, validating the effectiveness of each improvement, and table 2 outlines the impact of the cache system on data storage and transmission.

Size	Optimizations			Efficiency		
	Op	Conv w/o APSC	APSC	Attn	Value	Ratio
(512,512)					5.42it/s	1.0 \times
	✓				6.80it/s	1.2 \times
	✓	✓			6.86it/s	1.2 \times
	✓		✓		7.2it/s	1.3 \times
	✓		✓	✓	10.17it/s	1.9\times
(768,768)					1.83it/s	1.0 \times
	✓				2.73it/s	1.5 \times
	✓	✓			3.17it/s	1.7 \times
	✓		✓		3.61it/s	2.0 \times
	✓		✓	✓	8.13it/s	4.4\times

Table 1: Ablation study of each optimization. Op: CUDA kernel improvements (e.g., Flash Attention (Dao et al. 2022)). Conv: Combining tiled sparse convolution with or without APSC. Attn: Using approximate attention.

Device	Edit size	Compute	Cached data		Transfer	
			w/o	FISEdit	w/o	FISEdit
TITAN RTX	30%	247ms		987		108
	15%	186ms	4192	1192	387	114
	5%	128ms		1348		121
A100 PCIe	30%	108ms		987		67
	15%	84ms	4192	1192	232	72
	5%	72ms		1348		76

Table 2: Comparison of computation time, data transfer time (ms) and cached data size (MB) with and without FISEdit.

Conclusion

In real-world text-to-image model applications, users often desire minor edits to generated images based on textual descriptions. However, existing methods fail to exploit the redundant computation between previous and current diffusion model calls. To sufficiently reuse prior results, we have solved the three challenges: difference detection, sparse computation and pre-computed data management. The experimental results indicate that our method can accurately edit the details of images while achieves a considerable speedup.

Limitations. Our method has a poor performance when editing a low-resolution images (e.g., 256×256). This is because, the image changes caused by semantic modification do not exhibit adequate sparsity in low-resolution images. However, GANs hold a competitive position in low-resolution image generation, and have less computational overhead compared to diffusion models, meaning that large image generation is the main challenge for real-world use.

Prospects. With our method, it is possible to design a cache system for a real-world text-to-image service. Specifically, when a user query that contains a text is received, we can search for the most similar text-image pairs stored in the cache and process this query sparsely and incrementally, thereby achieving better system throughput.

Method	Property			Efficiency					
	Without Training	Partial Edit	Edit Size	MACs		TITAN RTX		A100 PCIe 40GB	
				Value	Ratio	Value	Ratio	Value	Ratio
SDTP	✓	×							
SDIP	×	✓	-	1901G	1.0×	1.83it/s	1.0×	3.86it/s	1.0×
Prompt-to-Prompt	✓	×							
pix2pix-zero	✓	×	-	-	-	OOM	-	OOM	-
PPAP	✓	×							
InstructPix2Pix	×	×	-	4165G	0.5×	OOM	-	2.29it/s	0.6×
SDEdit	✓	×	-	1521G	1.2×	2.28it/s	1.2×	4.82it/s	1.2×
DIFFEdit	✓	✓	-	3041G	0.6×	1.15it/s	0.6×	2.42it/s	0.6×
			30%	835G	2.3×	4.06it/s	2.2×	9.57it/s	2.5×
FISEdit	✓	✓	15%	485G	3.9×	5.64it/s	3.1×	11.6it/s	3.0×
			5%	386G	4.9×	8.14it/s	4.4×	13.1it/s	3.4×

Table 3: Property and efficiency evaluation of each method. We use the Multiply-Accumulate Operations (MACs) of U-Net to measure computational cost and the number of U-Net calls completed per second as a measure of speed. Due to the fact that the inference speed of the baseline models will not be affected by the size of modifications, we don't discuss edit size for them.




Figure 5: Quantitative results of the images generated by baselines and our method, demonstrating the trade-off between consistency with the input image and consistency with the semantic change. We vary InstructPix2Pix’s image guidance scale between [1.0, 2.5], SDEdit’s strength between [0.5, 0.75], DIFFEdit’s strength between [0.5, 1.0], and edited size between [0.25, 0.75] for SDIP and our method.

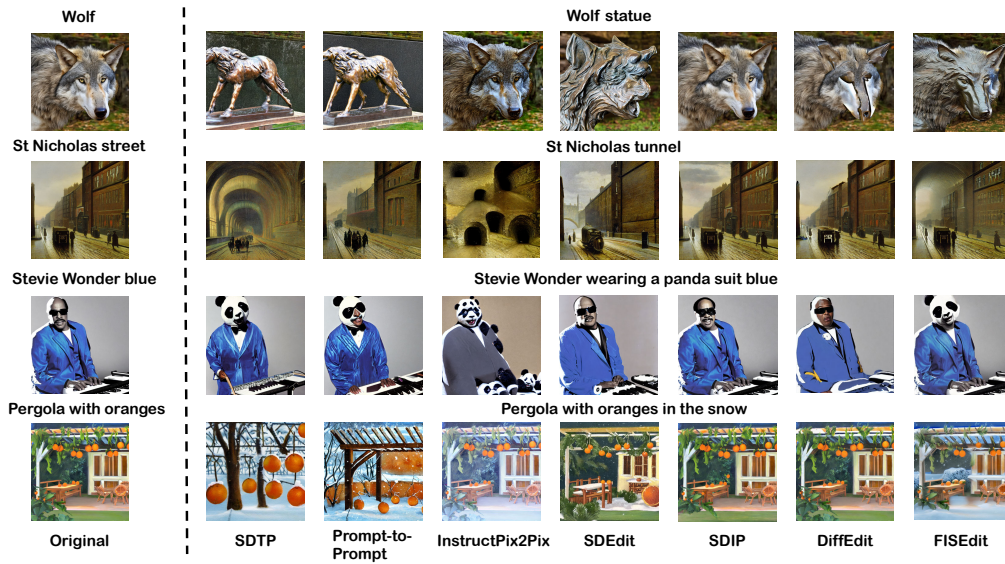


Figure 6: Qualitative results of the (768, 768) images generated by baselines and our method.

Acknowledgements

This work is supported by National Natural Science Foundation of China (U23B2048 and U22B2037), China National Postdoctoral Program for Innovative Talents (BX20230012), and PKU-Tencent joint research Lab. Bin Cui and Xupeng Miao are the co-corresponding authors.

References

- Abadi, M.; Barham, P.; Chen, J.; Chen, Z.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Irving, G.; Isard, M.; et al. 2016. Tensorflow: a system for large-scale machine learning. In *OSDI*.
- Abdal, R.; Qin, Y.; and Wonka, P. 2019. Image2StyleGAN: How to Embed Images Into the StyleGAN Latent Space? In *ICCV*.
- Abdal, R.; Qin, Y.; and Wonka, P. 2020. Image2StyleGAN++: How to Edit the Embedded Images? In *CVPR*.
- Balaji, Y.; Nah, S.; Huang, X.; Vahdat, A.; Song, J.; Kreis, K.; Aittala, M.; Aila, T.; Laine, S.; Catanzaro, B.; et al. 2022. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *ArXiv preprint*.
- Bi, H.; Song, Y.; Yin, Y.; Zhang, J.; Hong, W.; and Wu, Y. 2023. A novel sparse SAR unambiguous imaging method based on mixed-norm optimization. *Science China Information Sciences*, 66(11): 1–2.
- Bianchi, F.; Kalluri, P.; Durmus, E.; Ladhak, F.; Cheng, M.; Nozza, D.; Hashimoto, T.; Jurafsky, D.; Zou, J.; and Caliskan, A. 2022. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. *ArXiv preprint*.
- Bolya, D.; Fu, C.-Y.; Dai, X.; Zhang, P.; Feichtenhofer, C.; and Hoffman, J. 2023. Token Merging: Your ViT But Faster. *ICLR*.
- Bolya, D.; and Hoffman, J. 2023. Token Merging for Fast Stable Diffusion. *ArXiv preprint*.
- Brooks, T.; Holynski, A.; and Efros, A. A. 2022. Instructpix2pix: Learning to follow image editing instructions. *ArXiv preprint*.
- Child, R.; Gray, S.; Radford, A.; and Sutskever, I. 2019. Generating long sequences with sparse transformers. *ArXiv preprint*.
- Choi, J.; Kim, S.; Jeong, Y.; Gwon, Y.; and Yoon, S. 2021. ILVR: Conditioning Method for Denoising Diffusion Probabilistic Models. In *ICCV*.
- Couairon, G.; Verbeek, J.; Schwenk, H.; and Cord, M. 2023. Diffedit: Diffusion-based semantic image editing with mask guidance. *ICLR*.
- Dao, T.; Fu, D.; Ermon, S.; Rudra, A.; and Ré, C. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *NeurIPS*.
- Dong, X.; Huang, J.; Yang, Y.; and Yan, S. 2017. More is Less: A More Complicated Network with Less Inference Complexity. In *CVPR*.
- Feng, Y.; Chen, B.; Tang, M.; Qi, Y.; Bai, M.; Wang, C.; Wang, H.; Zhan, X.; Zhang, J.; Liu, J.; et al. 2023. Near-threshold-voltage operation in flash-based high-precision computing-in-memory to implement Poisson image editing. *Science China Information Sciences*, 66(12): 222402.
- Gal, R.; Patashnik, O.; Maron, H.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022. StyleGAN-NADA: CLIP-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*.
- Han, P.; Zhou, S.; Yu, J.; Xu, Z.; Chen, L.; and Shang, S. 2023. Personalized Re-ranking for Recommendation with Mask Pretraining. *Data Science and Engineering*, 8(4): 357–367.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *CVPR*.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *ArXiv preprint*.
- Hessel, J.; Holtzman, A.; Forbes, M.; Le Bras, R.; and Choi, Y. 2021. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. In *EMNLP*.
- 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. In *NeurIPS*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *NeurIPS*.
- Huang, R.; Lam, M. W.; Wang, J.; Su, D.; Yu, D.; Ren, Y.; and Zhao, Z. 2022. Fastdiff: A fast conditional diffusion model for high-quality speech synthesis. *ArXiv preprint*.
- Ioffe, S.; and Szegedy, C. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A Style-Based Generator Architecture for Generative Adversarial Networks. In *CVPR*.
- Kim, G.; Kwon, T.; and Ye, J. C. 2022. DiffusionCLIP: Text-Guided Diffusion Models for Robust Image Manipulation. In *CVPR*.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *ICLR*.
- Kong, Z.; and Ping, W. 2021. On fast sampling of diffusion probabilistic models. *ArXiv preprint*.
- Li, M.; Lin, J.; Meng, C.; Ermon, S.; Han, S.; and Zhu, J.-Y. 2022. Efficient spatially sparse inference for conditional gans and diffusion models. *NeurIPS*.
- Lin, H.; Cheng, X.; Wu, X.; and Shen, D. 2022. Cat: Cross attention in vision transformer. In *IEEE ICME*.
- Liu, L.; Deng, L.; Hu, X.; Zhu, M.; Li, G.; Ding, Y.; and Xie, Y. 2019. Dynamic Sparse Graph for Efficient Deep Learning. In *ICLR*.

- Liu, X.; Pool, J.; Han, S.; and Dally, W. J. 2018. Efficient Sparse-Winograd Convolutional Neural Networks. In *ICLR*.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022a. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *ArXiv preprint*.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022b. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *ArXiv preprint*.
- Meng, C.; Gao, R.; Kingma, D. P.; Ermon, S.; Ho, J.; and Salimans, T. 2022a. On distillation of guided diffusion models. *ArXiv preprint*.
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.; and Ermon, S. 2022b. SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations. In *ICLR*.
- Miao, X.; Nie, X.; Zhang, H.; Zhao, T.; and Cui, B. 2022a. Hetu: A highly efficient automatic parallel distributed deep learning system. *Sci. China Inf. Sci.*
- Miao, X.; Shi, Y.; Zhang, H.; Zhang, X.; Nie, X.; Yang, Z.; and Cui, B. 2022b. HET-GMP: A Graph-based System Approach to Scaling Large Embedding Model Training. In *Proceedings of SIGMOD Conference*, 470–480. ACM.
- Miao, X.; Zhang, H.; Shi, Y.; Nie, X.; Yang, Z.; Tao, Y.; and Cui, B. 2022c. HET: Scaling out Huge Embedding Model Training via Cache-enabled Distributed Framework. *Proc. VLDB Endow.*, 15(2): 312–320.
- Nichol, A. Q.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2022. GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models. In *ICML*.
- Orgad, H.; Kawar, B.; and Belinkov, Y. 2023. Editing Implicit Assumptions in Text-to-Image Diffusion Models. *ArXiv preprint*.
- Park, T.; Liu, M.; Wang, T.; and Zhu, J. 2019. Semantic Image Synthesis With Spatially-Adaptive Normalization. In *CVPR*.
- Parmar, G.; Singh, K. K.; Zhang, R.; Li, Y.; Lu, J.; and Zhu, J.-Y. 2023. Zero-shot image-to-image translation. *ArXiv preprint*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*.
- Patashnik, O.; Wu, Z.; Shechtman, E.; Cohen-Or, D.; and Lischinski, D. 2021. StyleCLIP: Text-Driven Manipulation of StyleGAN Imagery. In *ICCV*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*.
- Ramasinghe, S.; Ranasinghe, K. N.; Khan, S. H.; Barnes, N.; and Gould, S. 2021. Conditional Generative Modeling via Learning the Latent Space. In *ICLR*.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *ArXiv preprint*.
- Ren, M.; Pokrovsky, A.; Yang, B.; and Urtasun, R. 2018. SBNet: Sparse Blocks Network for Fast Inference. In *CVPR*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*.
- Salimans, T.; and Ho, J. 2022. Progressive Distillation for Fast Sampling of Diffusion Models. In *ICLR*.
- Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *ArXiv preprint*.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *ICLR*.
- Tumanyan, N.; Geyer, M.; Bagon, S.; and Dekel, T. 2023. Plug-and-play diffusion features for text-driven image-to-image translation. In *CVPR*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *NeurIPS*.
- Watson, D.; Ho, J.; Norouzi, M.; and Chan, W. 2021. Learning to efficiently sample from diffusion probabilistic models. *ArXiv preprint*.
- Wei, K.-K.; Luo, W.-Q.; Tan, S.-Q.; and Huang, J.-W. 2023. CTNet: A Convolutional Transformer Network for Color Image Steganalysis. *Journal of Computer Science and Technology*.
- Wu, Y.; and He, K. 2018. Group normalization. In *ECCV*.
- Xiao, Z.; Kreis, K.; and Vahdat, A. 2022. Tackling the Generative Learning Trilemma with Denoising Diffusion GANs. In *ICLR*.
- Yi, P.; Li, J.; Choi, B.; Bhowmick, S. S.; and Xu, J. 2022. FLAG: towards graph query autocompletion for large graphs. *Data Science and Engineering*, 7(2): 175–191.
- Zhang, C.; Wang, H.-Z.; Liu, H.-W.; and Chen, Y.-L. 2023. Fine-Tuning Channel-Pruned Deep Model via Knowledge Distillation. *Journal of Computer Science and Technology*.
- Zhu, J.-Y.; Krähenbühl, P.; Shechtman, E.; and Efros, A. A. 2016. Generative visual manipulation on the natural image manifold. In *ECCV*.
- Zhu, P.; Abdal, R.; Qin, Y.; and Wonka, P. 2020. SEAN: Image Synthesis With Semantic Region-Adaptive Normalization. In *CVPR*.