

Null-text Inversion for Editing Real Images using Guided Diffusion Models

Ron Mokady^{*1,2}, Amir Hertz^{*1,2}, Kfir Aberman¹, Yael Pritch¹, and Daniel Cohen-Or^{†1,2}

¹Google Research

²The Blavatnik School of Computer Science, Tel Aviv University

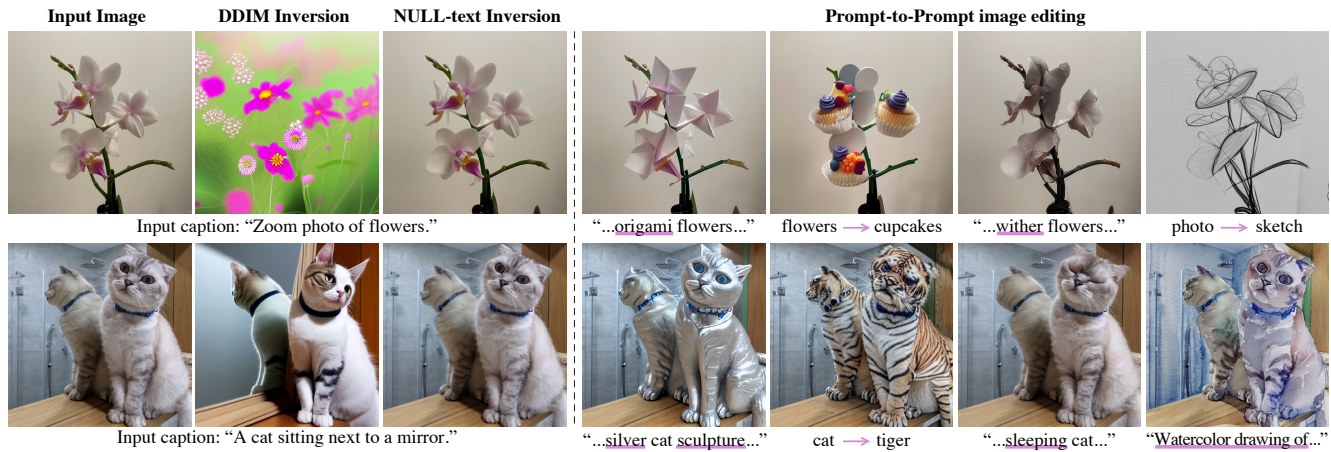


Figure 1. **Null-text inversion for real image editing.** Our method takes as input a real image (leftmost column) and an associated caption. The image is inverted with a DDIM diffusion model to yield a diffusion trajectory (second column to the left). Once inverted, we use the initial trajectory as a pivot for null-text optimization that accurately reconstructs the input image (third column to the left). Then, we can edit the inverted image by modifying only the input caption using the editing technique of Prompt-to-Prompt [18].

Abstract

Recent large-scale text-guided diffusion models provide powerful image generation capabilities. Currently, a massive effort is given to enable the modification of these images using text only as means to offer intuitive and versatile editing tools. To edit a real image using these state-of-the-art tools, one must first invert the image with a meaningful text prompt into the pretrained model’s domain. In this paper, we introduce an accurate inversion technique and thus facilitate an intuitive text-based modification of the image. Our proposed inversion consists of two key novel components: (i) Pivotal inversion for diffusion models. While current methods aim at mapping random noise samples to a single input image, we use a single pivotal noise vector for each timestamp and optimize around it. We demonstrate that a direct DDIM inversion is inadequate on its own, but does provide a rather good anchor for our optimization. (ii) Null-text optimization, where we only modify the unconditional textual embedding that is used for classifier-free guidance, rather than the input text embedding. This allows for keeping both the model weights and the conditional embedding intact and hence enables applying prompt-based editing while avoiding the cumbersome tuning of the model’s weights. Our null-text inversion, based on the publicly available Stable Diffusion model, is extensively evaluated on a variety of images and various

prompt editing, showing high-fidelity editing of real images.

1. Introduction

The progress in image synthesis using text-guided diffusion models has attracted much attention due to their exceptional realism and diversity. Large-scale models [29, 32, 34] have ignited the imagination of multitudes of users, enabling image generation with unprecedented creative freedom. Naturally, this has initiated ongoing research efforts, investigating how to harness these powerful models for image editing. Most recently, intuitive text-based editing was demonstrated over synthesized images, allowing the user to easily manipulate an image using text only [18].

However, text-guided editing of a real image with these state-of-the-art tools requires *inverting* the given image and textual prompt. That is, finding an initial noise vector that produces the input image when fed with the prompt into the diffusion process while preserving the editing capabilities of the model. The inversion process has recently drawn considerable attention for GANs [7, 44], but has not yet been fully addressed for text-guided diffusion models. Although an effective DDIM inversion [15, 37] scheme was suggested for unconditional diffusion models, it is found lacking for text-guided diffusion models when classifier-free guidance [20], which is necessary for meaningful editing, is applied.

^{*}Equal contribution. Performed this work while working at Google.

[†]Performed this work while working at Google.

In this paper, we introduce an effective inversion scheme, achieving near-perfect reconstruction, while retaining the rich text-guided editing capabilities of the original model (see Fig. 1). Our approach is built upon the analysis of two key aspects of guided diffusion models: classifier-free guidance and DDIM inversion.

In the widely used classifier-free guidance, in each diffusion step, the prediction is performed twice: once unconditionally and once with the text condition. These predictions are then extrapolated to amplify the effect of the text guidance. While all works concentrate on the conditional prediction, we recognize the substantial effect induced by the unconditional part. Hence, we optimize the embedding used in the unconditional part in order to invert the input image and prompt. We refer to it as *null-text optimization*, as we replace the embedding of the empty text string with our optimized embedding.

DDIM Inversion consists of performing DDIM sampling in reverse order. Although a slight error is introduced in each step, this works well in the unconditional case. However, in practice, it breaks for text-guided synthesis, since classifier-free guidance magnifies its accumulated error. We observe that it can still offer a promising starting point for the inversion. Inspired by GAN literature, we use the sequence of noised latent codes, obtained from an initial DDIM inversion, as pivot [31]. We then perform our optimization around this pivot to yield an improved and more accurate inversion. We refer to this highly efficient optimization as *Diffusion Pivotal Inversion*, which stands in contrast to existing works that aim to map all possible noise vectors to a single image.

To the best of our knowledge, our approach is the first to enable the text editing technique of Prompt-to-Prompt [18] on real images. Moreover, unlike recent approaches [21, 41], we do not tune the model weights, thus avoiding damaging the prior of the trained model and duplicating the entire model for each image. Throughout comprehensive ablation study and comparisons, we demonstrate the contribution of our key components to achieving a high-fidelity reconstruction of the given real image, while allowing meaningful and intuitive editing abilities. Our code, built upon the publicly available Stable Diffusion model, will be released.

2. Related Work

Large-scale diffusion models, such as Imagen [34], DALL-E 2 [29], and Stable Diffusion [32], have recently raised the bar for the task of generating images conditioned on plain text, known as text-to-image synthesis. Exploiting the powerful architecture of diffusion models [19, 19, 32, 36–38], these models can generate practically any image by simply feeding a corresponding text, and so have changed the landscape of artistic applications.

However, synthesizing very specific or personal objects which are not widespread in the training data has been chal-




Figure 2. **Real image editing using our method.** We first apply a single NULL-text inversion over the real input image, achieving high-fidelity reconstruction. Then, various Prompt-to-Prompt text-based editing operations are applied. As can be seen, our inversion scheme provides high fidelity while retaining high editability. See additional examples in the supplementary materials.

lenging. This requires an *inversion* process that given input images would enable regenerating the depicted object using a text-guided diffusion model. Inversion has been studied extensively for GANs [7, 12, 24, 44–46], ranging from latent-based optimization [1, 2] and encoders [30, 40] to feature space encoders [42] and fine-tuning of the model [3, 31]. Motivated by this, Gal et al. [17] suggest a textual inversion scheme for diffusion models that enables regenerating a user-provided concept out of 3 – 5 images. Concurrently, Ruiz et al. [33] tackled the same task with model-tuning. However, these works struggle to edit a given real image while accurately reproducing the unedited parts.

Naturally, recent works have attempted to adapt text-guided diffusion models to the fundamental challenge of single-image editing, aiming to exploit their rich and diverse semantic knowledge. DALL-E 2 [29] demonstrates editing capabilities by interpolating between the CLIP image embeddings and a normalised difference of the CLIP text embeddings produced from two text descriptions. Meng et al. [25] add noise to the input image and then perform a text-guided denoising process from a predefined step. Yet, they struggle to accurately preserve the input image details. To overcome this, several works [4, 5, 27] assume that the user provides a mask to restrict the region in

which the changes are applied, achieving both meaningful editing and background preservation.

However, requiring that users provide a precise mask is burdensome. Furthermore, masking the image content removes important information, which is mostly ignored in the inpainting process. While some text-only editing approaches are bound to global editing [14, 22, 23], Bartal et al. [6] propose a text-based localized editing technique without using any mask. Their technique allows high-quality texture editing, but not modifying complex structures, since only CLIP [28] is employed as guidance instead of a generative diffusion model. Hertz et al. [18] suggest an intuitive editing technique, called Prompt-to-Prompt, of manipulating local or global details by modifying only the text prompt when using text-guided diffusion models. By injecting internal cross-attention maps, they preserve the spatial layout and geometry which enable the regeneration of an image while modifying it through prompt editing. Still, without an inversion technique, their approach is limited to synthesized images. Sheynin et al. [35] suggest training the model for local editing without the inversion requirement, but their expressiveness and quality are inferior compared to current large-scale diffusion models.

Concurrent to our work, DiffEdit [11] uses DDIM inversion for image editing, but avoids the emergent distortion by automatically producing a mask that allows background preservation. Wu et al. [43] and Chandramouli et al. [8] employ a deterministic forward diffusion and the desired manipulation is achieved by providing the target text to condition the reverse diffusion process.

Also concurrent, Imagic [21] and UniTune [41] have demonstrated impressive editing results using the powerful Imagen model [34]. Yet, they both require the restrictive fine-tuning of the model. Moreover, Imagic requires a new tuning for each editing, while UniTune involves a parameter search for each image. Our method enables us to apply the text-only intuitive editing of Prompt-to-Prompt [18] on real images. We do not require any fine-tuning and provide highly-quality local and global modifications using the publicly available Stable Diffusion [32] model.

3. Method

Let \mathcal{I} be a real image. Our goal is to edit \mathcal{I} , using only text guidance, to get an edited image \mathcal{I}^* . We use the setting defined by Prompt-to-Prompt [18], where the editing is guided by source prompt \mathcal{P} and edited prompt \mathcal{P}^* . This requires the user to provide a source prompt. Yet, we found that automatically producing the source prompt using an off-the-shelf captioning model [26] works well (see Sec. 4). For example, see Fig. 2, given an image and a source prompt "A baby wearing...", we replace the baby with a robot by providing the edited prompt "A robot wearing...".

Such editing operations first require inverting \mathcal{I} to the model's output domain. Namely, the main challenge is faithfully reconstructing \mathcal{I} by feeding the source prompt \mathcal{P}

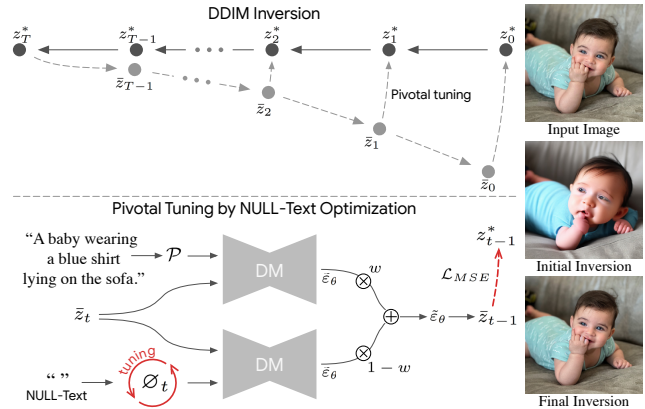


Figure 3. **NULL-text Inversion overview.** *Top: pivotal inversion.* We first apply an initial DDIM inversion on the input image which estimates a diffusion trajectory $\{z_t^*\}_0^T$. Starting the diffusion process from the last latent z_T^* results in unsatisfying reconstruction as the latent codes become farther away from the original trajectory. We use the initial trajectory as a pivot for our optimization which brings the diffusion backward trajectory $\{\bar{z}_t\}_1^T$ closer to the original image encoding z_0^* . *Bottom: NULL-text optimization for timestamp t .* Recall that classifier-free guidance consists of performing the prediction ϵ_θ twice – using text condition embedding and unconditionally using NULL-text embedding \emptyset (bottom-left). Then, these are extrapolated with guidance scale w (middle). We optimize only the unconditional embeddings \emptyset_t by employing a reconstruction MSE loss (in red) between the predicted latent code z_{t-1} to the pivot z_{t-1}^* .

to the model, while still retaining the intuitive text-based editing abilities.

Our approach is based on two main observations. First, DDIM inversion produces unsatisfying reconstruction when classifier-free guidance is applied, but provides a good starting point for the optimization, enabling us to efficiently achieve high-fidelity inversion. Second, optimizing the unconditional null embedding, which is used in classifier-free guidance, allows an accurate reconstruction while avoiding the tuning of the model and the conditional embedding. Thereby preserving the desired editing capabilities.

Next, we provide a short background, followed by a detailed description of our approach in Sec. 3.2 and Sec. 3.3. A general overview is provided in Fig. 3.

3.1. Background and Preliminaries

Text-guided diffusion models aim to map a random noise vector z_t and textual condition \mathcal{P} to an output image z_0 , which corresponds to the given conditioning prompt. In order to perform sequential denoising, the network ϵ_θ is trained to predict artificial noise, following the objective:

$$\min_{\theta} E_{z_0, \epsilon \sim N(0, I), t \sim \text{Uniform}(1, T)} \|\epsilon - \epsilon_\theta(z_t, t, \mathcal{C})\|_2^2. \quad (1)$$

Note that $\mathcal{C} = \psi(\mathcal{P})$ is the embedding of the text condition and z_t is a noised sample, where noise is added to the sampled data z_0 according to timestamp t . At inference, given a




Figure 4. **Ablation Study.** *Top: we compare the performance of our full algorithm (green line) to different variations, evaluating the reconstruction quality by measuring the PSNR score as a function of number optimization iterations and running time in minutes. Bottom: we visually show the inversion results after 200 iterations of our full algorithm (on right) compared to other baselines. Results for all iterations are shown in the supplementary materials.*

noise vector z_T , The noise is gradually removed by sequentially predicting it using our trained network for T steps.

Since we aim to accurately reconstruct a given real image, we employ the deterministic DDIM sampling [37]:

$$z_{t-1} = \sqrt{\frac{\alpha_{t-1}}{\alpha_t}} z_t + \sqrt{\alpha_{t-1}} \left(\sqrt{\frac{1}{\alpha_{t-1}} - 1} - \sqrt{\frac{1}{\alpha_t} - 1} \right) \cdot \varepsilon_\theta(z_t)$$

For the definition of α_t and additional details, please refer to the supplementary materials. Diffusion models often operate in the image pixel space where z_0 is a sample of a real image. In our case, we use the popular and publicly available Stable Diffusion model [32] where the diffusion forward process is applied on a latent image encoding $z_0 = E(x_0)$ and an image decoder is employed at the end of the diffusion backward process $x_0 = D(z_0)$.

Classifier-free guidance. One of the key challenges in text-guided generation is the amplification of the effect induced by the conditioned text. To this end, Ho et al. [20] have presented the classifier-free guidance technique, where the prediction is also performed unconditionally, which is then extrapolated with the conditioned prediction. More formally, let $\emptyset = \psi(\text{" "})$ be the *embedding* of a null text and let w be the guidance scale parameter, then the classifier-free guidance prediction is defined by:

$$\tilde{\varepsilon}_\theta(z_t, t, \mathcal{C}, \emptyset) = w \cdot \varepsilon_\theta(z_t, t, \mathcal{C}) + (1 - w) \cdot \varepsilon_\theta(z_t, t, \emptyset).$$

E.g., $w = 7.5$ is the default parameter for Stable Diffusion.

DDIM inversion. A simple inversion technique was suggested for the DDIM sampling [15, 37], based on the assumption that the ODE process can be reversed in the limit of small steps:

$$z_{t+1} = \sqrt{\frac{\alpha_{t+1}}{\alpha_t}} z_t + \sqrt{\alpha_{t+1}} \left(\sqrt{\frac{1}{\alpha_{t+1}} - 1} - \sqrt{\frac{1}{\alpha_t} - 1} \right) \cdot \varepsilon_\theta(z_t)$$

In other words, the diffusion process is performed in the reverse direction, that is $z_0 \rightarrow z_T$ instead of $z_T \rightarrow z_0$, where z_0 is set to be the encoding of the given real image.

3.2. Pivotal Inversion

Recent inversion works use random noise vectors for each iteration of their optimization, aiming at mapping every noise vector to a single image. We observe that this is inefficient as inference requires only a single noise vector. Instead, inspired by GAN literature [31], we seek to perform a more "local" optimization, ideally using only a single noise vector. In particular, we aim to perform our optimization around a pivotal noise vector which is a good approximation and thus allows a more efficient inversion.

We start by studying the DDIM inversion. In practice, a slight error is incorporated in every step. For unconditional diffusion models, the accumulated error is negligible and the DDIM inversion succeeds. However, recall that meaningful editing using the Stable Diffusion model [32] requires applying classifier-free guidance with a large guidance scale $w > 1$. We observe that such a guidance scale amplifies the accumulated error. Therefore, performing the DDIM inversion procedure with classifier-free guidance results not only in visual artifacts, but the obtained noise vector might be out of the Gaussian distribution. The latter decreases the editability, i.e., the ability to edit using the particular noise vector.

We do recognize that using DDIM inversion with guidance scale $w = 1$ provides a rough approximation of the original image which is highly editable but far from accurate. More specifically, the reversed DDIM produces a T steps trajectory between the image encoding z_0 to a Gaussian noise vector z_T^* . Again, a large guidance scale is essential for editing. Hence, we focus on feeding z_T^* to the diffu-



Figure 5. **Fine control editing using attention re-weighting.** We can use attention re-weighting to further control the level of dryness over the field or create a denser zebra pattern over the couch.

sion process with classifier-free guidance ($w > 1$). This results in high editability but inaccurate reconstruction, since the intermediate latent codes deviate from the trajectory, as illustrated in Fig. 3.

Motivated by the high editability, we refer to this initial DDIM inversion with $w = 1$ as our *pivot* trajectory and perform our optimization around it with the standard guidance scale, $w > 1$. That is, our optimization maximizes the similarity to the original image while maintaining our ability to perform meaningful editing. In practice, we execute a separate optimization for each timestamp t in the order of the diffusion process $t = T \rightarrow t = 1$ with the objective of getting close as possible to the initial trajectory z_T^*, \dots, z_0^* :

$$\min \|z_{t-1}^* - z_{t-1}\|_2^2, \quad (2)$$

where z_{t-1} is the intermediate result of the optimization. Since our pivotal DDIM inversion provides a rather good starting point, this optimization is highly efficient compared to using random noise vectors, as demonstrated in Sec. 4.

Note that for every $t < T$, the optimization should start from the endpoint of the previous step ($t + 1$) optimization, otherwise our optimized trajectory would not hold at inference. Therefore, after the optimization of step t , we compute the current noisy latent \bar{z}_t , which is then used in the optimization of the next step to ensure our new trajectory would end near z_0 (see Eq. (3) for more details).

3.3. Null-text optimization

To successfully invert real images into the model’s domain, recent works optimize the textual encoding [17], the network’s weights [33, 41], or both [21]. Fine-tuning the model’s weight for each image involves duplicating the entire model which is highly inefficient in terms of memory consumption. Moreover, unless fine-tuning is applied for each and every edit, it necessarily hurts the learned prior of the model and therefore the semantics of the edits. Direct optimization of the textual embedding results in a non-interpretable representation since the optimized tokens does not necessarily match pre-existing words. Therefore, an intuitive prompt-to-prompt edit becomes more challenging.

Instead, we exploit the key feature of the classifier-free guidance — the result is highly affected by the unconditional prediction. Therefore, we replace the default null-text embedding with an optimized one, referred to as *null-text optimization*. Namely, for each input image, we optimize only the unconditional embedding \emptyset , initialized with the null-text embedding. The model and the conditional textual embedding are kept unchanged.

This results in high-quality reconstruction while still allowing intuitive editing with Prompt-to-Prompt [18] by simply using the optimized unconditional embedding. Moreover, after a single inversion process, the same unconditional embedding can be used for multiple editing operations over the input image. Since null-text optimization is naturally less expressive than fine-tuning the entire model, it requires the more efficient pivotal inversion scheme.

We refer to optimizing a single unconditional embedding \emptyset as a *Global* null-text optimization. During our experiments, as shown in Fig. 4, we have observed that optimizing a different “null embedding” \emptyset_t for each timestamp t significantly improves the reconstruction quality while this is well suited for our pivotal inversion. And so, we use per-timestamp unconditional embeddings $\{\emptyset_t\}_{t=1}^T$, and initialize \emptyset_t with the embedding of the previous step \emptyset_{t+1} .

Putting the two components together, our full algorithm is presented in algorithm 1. The DDIM inversion with $w = 1$ outputs a sequence of noisy latent codes z_T^*, \dots, z_0^* where $z_0^* = z_0$. We initialize $\bar{z}_T = z_T^*$, and perform the following optimization with the default guidance scale $w = 7.5$ for the timestamps $t = T, \dots, 1$, each for N iterations:

$$\min_{\emptyset_t} \|z_{t-1}^* - z_{t-1}(\bar{z}_t, \emptyset_t, \mathcal{C})\|_2^2. \quad (3)$$

For simplicity, $z_{t-1}(\bar{z}_t, \emptyset_t, \mathcal{C})$ denotes applying DDIM sampling step using \bar{z}_t , the unconditional embedding \emptyset_t , and the conditional embedding \mathcal{C} . At the end of each step, we update

$$\bar{z}_{t-1} = z_{t-1}(\bar{z}_t, \emptyset_t, \mathcal{C}).$$

We find that early stopping reduces time consumption, resulting in ~ 1 minute using a single A100 GPU.

Finally, we can edit the real input image by using the noise $\bar{z}_T = z_T^*$ and the optimized unconditional embeddings $\{\emptyset_t\}_{t=1}^T$. Please refer to the supplementary materials for additional implementation details.

4. Ablation Study

In this section, we validate the contribution of our main components, thoroughly analyzing the effectiveness of our design choices by conducting an ablation study. We focus on the fidelity to the input image which is an essential evaluation for image editing. In Sec. 5 we demonstrate that our method performs high-quality and meaningful manipulations.

Experimental setting. Evaluation is provided in Fig. 4. We have used a subset of 100 images and captions pairs, randomly selected from the COCO [9] validation set. We

Algorithm 1: Null-text inversion

- 1 **Input:** A source prompt embedding $\mathcal{C} = \psi(\mathcal{P})$ and input image \mathcal{I} .
- 2 **Output:** Noise vector z_T and optimized embeddings $\{\varnothing_t\}_{t=1}^T$.

- 3 Set guidance scale $w = 1$;
- 4 Compute the intermediate results z_T^*, \dots, z_0^* using DDIM inversion over \mathcal{I} ;
- 5 Set guidance scale $w = 7.5$;
- 6 Initialize $\bar{z}_T \leftarrow z_T^*$, $\varnothing_T \leftarrow \psi(\text{""})$;
- 7 **for** $t = T, T - 1, \dots, 1$ **do**
- 8 **for** $j = 0, \dots, N - 1$ **do**
- 9 $\varnothing_t \leftarrow \varnothing_t - \eta \nabla_{\varnothing} \|z_{t-1}^* - z_{t-1}(\bar{z}_t, \varnothing_t, \mathcal{C})\|_2^2$;
- 10 **end**
- 11 Set $\bar{z}_{t-1} \leftarrow z_{t-1}(\bar{z}_t, \varnothing_t, \mathcal{C})$, $\varnothing_{t-1} \leftarrow \varnothing_t$;
- 12 **end**
- 13 **Return** $\bar{z}_T, \{\varnothing_t\}_{t=1}^T$

then applied our approach on each image-caption pair using the default Stable Diffusion hyper-parameters for an increasing number of iterations per diffusion step, $N = 1, \dots, 20$ (see algorithm 1). The reconstruction quality was measured in terms of mean PSNR. We now turn to analyze different variations of our algorithm.

DDIM inversion. We mark the DDIM inversion as a lower bound for our algorithm, as it is the starting point of our optimization, producing unsatisfying reconstruction when classifier-free guidance is applied (see Sec. 3.2).

VQAE. For an upper bound, we consider the reconstruction using the VQ auto-encoder [16], denoted VQAE, which is used by the Stable Diffusion model. Although the latent code or the VQ-decoder can be further optimized according to the input image, this is out of our scope, since it would be only applicable to this specific model [32] while we aim for a general algorithm. Therefore, our optimization treats its encoding z_0 as ground truth, as the obtained reconstruction is quite accurate in most cases.

Our method. As can be seen in Fig. 4, our method converges to a near-optimal reconstruction with respect to the VQAE upper bound after a total number of 500 iterations ($N = 10$) and even after 250 iterations (~ 1 minute on an A100 GPU) we achieve high-quality inversion.

Random Pivot. We validate the importance of the DDIM initialization by replacing the DDIM-based trajectory with a *single random trajectory* of latent codes, sharing the same starting point z_0 — the input image encoding. In other words, we randomly sample a single Gaussian noise $\sim N(0, I)$ for each image and use it to noise the corresponding encoding z_0 from $t = 1$ to $t = T$ using the diffusion scheduler. As presented in Fig. 4, the DDIM initial-



Figure 6. **Comparison.** Text2LIVE [6] excels at replacing textures locally but struggles to perform more structured editing, such as replacing a kid with a tiger. VQGAN+CLIP [13] obtains inferior realism. SDEdit [25] fails to faithfully reconstruct the original image, resulting in identity drift when humans are involved. Our method achieves realistic editing of both textures and structured objects while retaining high fidelity to the original image.

ization is crucial for fast convergence, since the initial error becomes significantly larger when the pivot is random.

Robustness to different input captions. Since we require an input caption, it is only natural to ask whether our method is highly sensitive to the chosen caption. We take this to the extreme by sampling a *random caption* from the dataset for each image. Even with unaligned captions, the optimization converges to an optimal reconstitution with respect to the VQAE. Therefore, we conclude that our inversion is robust to the input caption. Clearly, choosing a random caption is undesired for text-based editing. But, providing any reasonable and editable prompt would work, including using an off-the-shelf captioning model [26, 39]. This is illustrated in the supplementary materials. We invert an image using multiple captions, demonstrating that the edited parts should be included in the source caption in order to produce semantic attention maps for editing. For example, to edit the print on the shirt, the source caption should include a "shirt with a drawing" or a similar phrase.

Global null-text embedding. We refer to optimizing a single embedding \varnothing for all timestamps as a *Global* embedding. As can be seen, such optimization struggles to converge, since it is less expressive than our final approach, which uses embedding per-timestamp $\{\varnothing_t\}_{t=1}^T$.

Textual inversion. We compare our method to *textual inversion*, similar to the proposed method by Gal et al. [17].

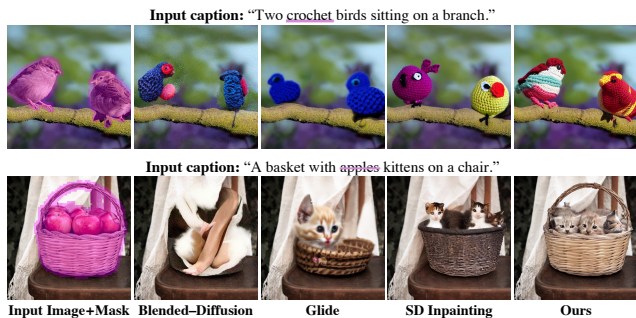


Figure 7. **Comparison to mask-based methods**. As can be seen, mask-based methods do not require inversion as the region outside the mask is kept. However, unlike our approach, such methods often struggle to preserve details that are found inside the masked region. For example, basket size is not preserved.

We optimize the textual embedding $\mathcal{C} = \psi(\mathcal{P})$ using random noise samples instead of pivotal inversion. That is, we randomly sample a different Gaussian noise for each optimization step and obtain z_t by adding it to z_0 according to the diffusion scheduler. Intuitively, this objective aims to map all noise vectors to a single image, in contrast to our pivotal tuning inversion which focuses on a single trajectory of noisy vectors. The optimization objective is then defined:

$$\min_{\mathcal{C}} E_{z_0, \varepsilon \sim N(0, I), t} \|\varepsilon - \varepsilon_{\theta}(z_t, t, \mathcal{C})\|_2^2. \quad (4)$$

Note that Gal et al. [17] have attempted to regenerate a specific object rather than achieve an accurate inversion. As presented in Fig. 4, the convergence is much slower than ours and results in poor reconstruction quality.

Textual inversion with a pivot. We also consider performing textual inversion around a pivot, i.e., similar to our pivotal inversion but optimizing the conditioned embedding. This results in a comparable reconstruction to ours, but with poor editability. We observe that in this case, visual information is often undesirably encoded in single or few tokens, weakening intuitive text-based editing. A visual example is provided in the supplementary materials.

Null-text optimization without pivotal inversion. We observe that optimizing the unconditional null-text embedding using random noise vectors, instead of pivotal inversion as described in previous paragraphs, completely breaks the null-text optimization. The results are inferior even to the DDIM inversion baseline as presented in the supplementary materials. We hypothesize that null-text optimization is less expressive than model-tuning and thus depends on the efficient pivotal inversion, as it struggles to map all noise vectors to a single image.

5. Results

Real image editing is presented in Figs. 1, 2 and 5, showing our method not only reaches remarkable reconstruction quality but also retains high editability. In particular, we use the intuitive approach of Prompt-to-Prompt [18] and demonstrate that the editing capabilities which previously

Table 1. **User study results**. The participants were asked to select the best editing result in terms of fidelity to both the input image and the textual edit instruction.

VQGAN+CLIP	Text2Live	SDEDIT	Ours
3.8%	16.6%	14.5%	65.1%

were constrained to synthesized images are now applied to real images using our inversion technique.

As can be seen in Fig. 2, our method effectively modifies both textures (“floral shirt”) and structured objects (“baby” to “robot”). Since we support the local editing of Prompt-to-Prompt and achieve high-fidelity reconstruction, the original identity is well preserved, even in the challenging case of a baby face. Fig. 2 also illustrates that our method requires only a single inversion process to perform multiple editing operations. Using a single inversion procedure, we can modify hair color, glasses, expression, background, and lighting and even replace objects or put on a joker make-up (bottom rows). Using Prompt-to-Prompt, we also can attenuate or amplify the effect of a specific word over real images, as appeared in Fig. 5. For additional examples, please refer to the supplementary materials.

Visual results for our high-fidelity reconstruction are presented in Fig. 1 and ?? and the supplementary materials, supporting our quantitative measures in Sec. 4. Evaluating additional editing technique and limitations are presented in the supplementary materials.

5.1. Comparisons

Our method aims attention at intuitive editing using only text, and so we compare our results to other text-only editing methods: (1) VQGAN+CLIP [13], (2) Text2Live [6], and (3) SDEdit [25]. We evaluated these on the images used by Bar-Tal et al. [6] and photos that include more structured objects, such as humans and animals, which we gathered from the internet. In total, we use 100 samples of images, input captions, and edited captions. We also compare our method to the mask-based methods of (4) Glide [27], (5) Blended-Diffusion [5], and (6) Stable Diffusion Inpaint [32]. The latter fine-tunes the diffusion model using an inpainting objective, allowing simple editing by inpainting a masked region using a target prompt. Lastly, we consider the concurrent work of (7) Imagic [21], which employs model-tuning per editing operation and has been designed for the Imagen model [34]. We refrain from comparing to the concurrent works of Unitune [41] and DiffEdit [11] as there are no available implementations.

Qualitative Comparison. As presented in Fig. 6, VQGAN+CLIP [13] mostly produces unrealistic results. Text2LIVE [6] handles texture modification well but fails to manipulate more structured objects, e.g., placing a boat (3rd row). Both struggle due to the use of CLIP [28] which lacks a generative ability. In SDEdit [25], the noisy image is fed to an intermediate step in the diffusion process, and therefore, it struggles to faithfully reconstruct the original details. This results in severe artifacts when fine details are

involved, such as human faces. For instance, identity drifts in the top row, and the background is not well preserved in the 2nd row. Contrarily, our method successfully preserves the original details, while allowing a wide range of realistic and meaningful editing, from simple textures to replacing well-structured objects.

Fig. 7 presents a comparison to mask-based methods, showing these struggles to preserve details that are found inside the masked region. This is due to the masking procedure that removes important structural information, and therefore, some capabilities are out of the inpainting reach.

A comparison to Imagic [21], which operates in a different setting – requiring model-tuning for each editing operation, is provided in the supplementary materials. We first employ the unofficial Imagic implementation for Stable Diffusion and present the results for different values of the interpolation parameter $\alpha = 0.6, 0.7, 0.8, 0.9$. This parameter is used to interpolate between the target text embedding and the optimized one [21]. In addition, the Imagic authors applied their method using the Imagen model over the same images, using the following parameters $\alpha = 0.93, 0.86, 1.08$. As can be seen, Imagic produces highly meaningful editing, especially when the Imagen model is involved. However, Imagic struggles to preserve the original details, such as the identity of the baby (1st row) or cups in the background (2nd row). Furthermore, we observe that Imagic is quite sensitive to the interpolation parameter α , as a high value reduces the fidelity to the image and a low value reduces the fidelity to the text guidance, while a single value cannot be applied to all examples. Lastly, Imagic takes a longer inference time, as shown in supplementary materials.

Quantitative Comparison. Since ground truth is not available for text-based editing of real images, quantitative evaluation remains an open challenge. Similar to [6, 18], we present a user study in Tab. 1. 50 participants have rated a total of 48 images for each baseline. The participants were recruited using *Prolific* (prolific.co). We presented side-by-side images produced by: VQGAN+CLIP, Text2LIVE, SDEdit, and our method (in random order). We focus on methods that share a similar setting to ours – no model tuning and mask requirement. The participants were asked to choose the method that better applies the requested edit while preserving most of the original details. A print screen is provided in the supplementary materials. As shown in Tab. 1, most participants favored our method. Quantitative comparison to Imagic is presented in the supplementary materials, using the unofficial Stable Diffusion implementation. According to these measures, our method achieves better scores for LPIPS perceptual distance, indicating a better fidelity to the input image.

We provide an additional evaluation in Tab. 2 using the ImageNet dataset according to the protocol presented by Couairon et al. [10], which includes a set of 1000 test im-

Table 2. Quantitative evaluation on ImageNet.

	SDEdit [25]	Text2Live [6]	FlexIT [10]	Ours
LPIPS ↓	31.5	20.0	24.7	29.1
Acc.% ↑	36.5	47.5	51.3	61.1

ages and their respective editing operations. This evaluation is well-defined and can be easily reproduced since the evaluation code is public [10]. Note that these editing operations consist of replacing the main object with another and therefore, the editing quality is measured by classification accuracy. We compare our approach to FlexIT [10], Text2Live [6], and SDEdit [25]. As presented in Tab. 2, our method achieves significantly better accuracy which indicates more meaningful editing. Although other methods achieve better similarity to the original image (LPIPS), we observe that most objects capture large portions of the image, and therefore, high similarity actually indicates less meaningful editing.

6. Conclusions

We have presented an approach to invert real images with corresponding captions into the latent space of a text-guided diffusion model while maintaining its powerful editing capabilities. Our two-step approach first uses DDIM inversion to compute a sequence of noisy codes, which roughly approximate the original image (with the given caption), then uses this sequence as a fixed pivot to optimize the input null-text embedding. Its fine optimization compensates for the inevitable reconstruction error caused by the classifier-free guidance component. Once the image-caption pair is accurately embedded in the output domain of the model, prompt-to-prompt editing can be instantly applied at inference time. By introducing two new technical concepts to text-guided diffusion models – pivotal inversion and null-text optimization, we were able to bridge the gap between reconstruction and editability. Our approach offers a surprisingly simple and compact means to reconstruct an arbitrary image, avoiding the computationally intensive model-tuning. We believe that null-text inversion paves the way for real-world use case scenarios for intuitive, text-based, image editing.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4432–4441, 2019. 2
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8296–8305, 2020. 2
- [3] Yuval Alaluf, Omer Tov, Ron Mokady, Rinon Gal, and Amit H. Bermano. Hyperstyle: Stylegan inversion with hypernetworks for real image editing, 2021. 2
- [4] Omri Avrahami, Ohad Fried, and Dani Lischinski. Blended latent diffusion. *arXiv preprint arXiv:2206.02779*, 2022. 2

- [5] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18208–18218, 2022. 2, 7
- [6] Omer Bar-Tal, Dolev Ofri-Amar, Rafail Fridman, Yoni Kasten, and Tali Dekel. Text2live: Text-driven layered image and video editing. *arXiv preprint arXiv:2204.02491*, 2022. 3, 6, 7, 8
- [7] Amit H Bermano, Rinon Gal, Yuval Alaluf, Ron Mokady, Yotam Nitzan, Omer Tov, Oren Patashnik, and Daniel Cohen-Or. State-of-the-art in the architecture, methods and applications of stylegan. In *Computer Graphics Forum*, volume 41, pages 591–611. Wiley Online Library, 2022. 1, 2
- [8] Paramanand Chandramouli and Kanchana Vaishnavi Gandikota. Ldedit: Towards generalized text guided image manipulation via latent diffusion models. *arXiv preprint arXiv:2210.02249*, 2022. 3
- [9] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 5
- [10] Guillaume Couairon, Asya Grechka, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Flexit: Towards flexible semantic image translation. In *CVPR 2022*. 8
- [11] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*, 2022. 3, 7
- [12] Antonia Creswell and Anil Anthony Bharath. Inverting the generator of a generative adversarial network. *IEEE transactions on neural networks and learning systems*, 30(7):1967–1974, 2018. 2
- [13] Katherine Crowson. Vqgan + clip, 2021. <https://colab.research.google.com/drive/1L8oL-vLjXVcRzCFbPwOoMkPKJ8-aYdPN>. 6, 7
- [14] Katherine Crowson, Stella Biderman, Daniel Kornis, Dashiell Stander, Eric Hallahan, Louis Castricato, and Edward Raff. Vqgan-clip: Open domain image generation and editing with natural language guidance. *arXiv preprint arXiv:2204.08583*, 2022. 3
- [15] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 1, 4
- [16] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis, 2020. 6
- [17] 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, 5, 6, 7
- [18] 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. 1, 2, 3, 5, 7, 8
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2
- [20] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. 1, 4
- [21] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Hui-Tang Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. *ArXiv*, abs/2210.09276, 2022. 2, 3, 5, 7, 8
- [22] 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. 3
- [23] Gihyun Kwon and Jong Chul Ye. Clipstyler: Image style transfer with a single text condition. *arXiv preprint arXiv:2112.00374*, 2021. 3
- [24] Zachary C Lipton and Subarna Tripathi. Precise recovery of latent vectors from generative adversarial networks. *arXiv preprint arXiv:1702.04782*, 2017. 2
- [25] Chenlin Meng, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 2, 6, 7, 8
- [26] Ron Mokady, Amir Hertz, and Amit H Bermano. Clip-cap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*, 2021. 3, 6
- [27] 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. 2, 7
- [28] 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. *arXiv preprint arXiv:2103.00020*, 2021. 3, 7
- [29] 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, 2
- [30] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. *arXiv preprint arXiv:2008.00951*, 2020. 2
- [31] Daniel Roich, Ron Mokady, Amit H. Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on Graphics (TOG)*, 2022. 2, 4
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 1, 2, 3, 4, 6, 7
- [33] 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. 2, 5
- [34] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. 1, 2, 3, 7

- [35] Shelly Sheynin, Oron Ashual, Adam Polyak, Uriel Singer, Oran Gafni, Eliya Nachmani, and Yaniv Taigman. Knn-diffusion: Image generation via large-scale retrieval. *arXiv preprint arXiv:2204.02849*, 2022. [3](#)
- [36] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. [2](#)
- [37] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2020. [1](#), [2](#), [4](#)
- [38] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#)
- [39] Matteo Stefanini, Marcella Cornia, Lorenzo Baraldi, Silvia Cascianelli, Giuseppe Fiameni, and Rita Cucchiara. From show to tell: a survey on deep learning-based image captioning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [6](#)
- [40] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *arXiv preprint arXiv:2102.02766*, 2021. [2](#)
- [41] Dani Valevski, Matan Kalman, Yossi Matias, and Yaniv Leviathan. Unitune: Text-driven image editing by fine tuning an image generation model on a single image. *arXiv preprint arXiv:2210.09477*, 2022. [2](#), [3](#), [5](#), [7](#)
- [42] Tengfei Wang, Yong Zhang, Yanbo Fan, Jue Wang, and Qifeng Chen. High-fidelity gan inversion for image attribute editing. *ArXiv*, abs/2109.06590, 2021. [2](#)
- [43] Chen Henry Wu and Fernando De la Torre. Unifying diffusion models’ latent space, with applications to cyclediffusion and guidance. *arXiv preprint arXiv:2210.05559*, 2022. [3](#)
- [44] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. Gan inversion: A survey, 2021. [1](#), [2](#)
- [45] Raymond A. Yeh, Chen Chen, Teck Yian Lim, Alexander G. Schwing, Mark Hasegawa-Johnson, and Minh N. Do. Semantic image inpainting with deep generative models, 2017. [2](#)
- [46] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *European conference on computer vision*, pages 597–613. Springer, 2016. [2](#)