

EDICT: Exact Diffusion Inversion via Coupled Transformations

Bram Wallace
Salesforce Research
b.wallace@salesforce.com

Akash Gokul
Salesforce Research
agokul@salesforce.com

Nikhil Naik
Salesforce Research
nnaik@salesforce.com

Abstract

Finding an initial noise vector that produces an input image when fed into the diffusion process (known as inversion) is an important problem in denoising diffusion models (DDMs), with applications for real image editing. The standard approach for real image editing with inversion uses denoising diffusion implicit models (DDIMs [29]) to deterministically noise the image to the intermediate state along the path that the denoising would follow given the original conditioning. However, DDIM inversion for real images is unstable as it relies on local linearization assumptions, which result in the propagation of errors, leading to incorrect image reconstruction and loss of content. To alleviate these problems, we propose Exact Diffusion Inversion via Coupled Transformations (EDICT), an inversion method that draws inspiration from affine coupling layers. EDICT enables mathematically exact inversion of real and model-generated images by maintaining two coupled noise vectors which are used to invert each other in an alternating fashion. Using Stable Diffusion [25], a state-of-the-art latent diffusion model, we demonstrate that EDICT successfully reconstructs real images with high fidelity. On complex image datasets like MS-COCO, EDICT reconstruction significantly outperforms DDIM, improving the mean square error of reconstruction by a factor of two. Using noise vectors inverted from real images, EDICT enables a wide range of image edits—from local and global semantic edits to image stylization—while maintaining fidelity to the original image structure. EDICT requires no model training/finetuning, prompt tuning, or extra data and can be combined with any pretrained DDM.

1. Introduction

Using the iterative denoising diffusion principle, denoising diffusion models (DDMs) trained with web-scale data can generate highly realistic images conditioned on input text, layouts, and scene graphs [24, 25, 27]. After image generation, the next important application of DDMs being explored by the research community is that of im-

age editing. Models such as DALL-E-2 [24] and Stable Diffusion [25] can perform inpainting, allowing users to edit images through manual annotation. Methods such as SDEdit [20] have demonstrated that both synthetic and real images can be edited using stroke or composite guidance via DDMs. However, the goal of holistic image editing tools that can edit any real/artificial image using purely text is still a field of active research.

The generative process of DDMs starts with an initial noise vector (x_T) and performs iterative denoising (typically with a guidance signal e.g. in the form of text-conditional denoising), ending with a realistic image sample (x_0). Reversing this generative process is key in solving this image editing problem for one family of approaches. Formally, this problem is known as “inversion” i.e., finding the initial noise vector that produces the input image when passed through the diffusion process.

A naïve approach for inversion is to add Gaussian noise to the input image and perform a predefined number of diffusion steps, which typically results in significant distortions [7]. A more robust method is adapting Denoising Diffusion Implicit Models (DDIMs) [29]. Unlike the commonly used Denoising Diffusion Probabilistic Models (DDPMs) [9], the generative process in DDIMs is defined in a non-Markovian manner, which results in a deterministic denoising process. DDIM can also be used for inversion, deterministically noising an image to obtain the initial noise vector ($x_0 \rightarrow x_T$).

DDIM inversion has been used for editing real images through text methods such as DDIBs [30] and Prompt-to-Prompt (P2P) image editing [7]. After DDIM inversion, P2P edits the original image by running the generative process from the noise vector and injecting conditioning information from a new text prompt through the cross-attention layers in the diffusion model, thus generating an edited image that maintains faithfulness to the original content while incorporating the edit. However, as noted in the original P2P work [7], the DDIM inversion is unstable in many cases—encoding from x_0 to x_T and back often results in inexact reconstructions of the original image as in Fig. 2. These distortions limit the ability to perform significant ma-

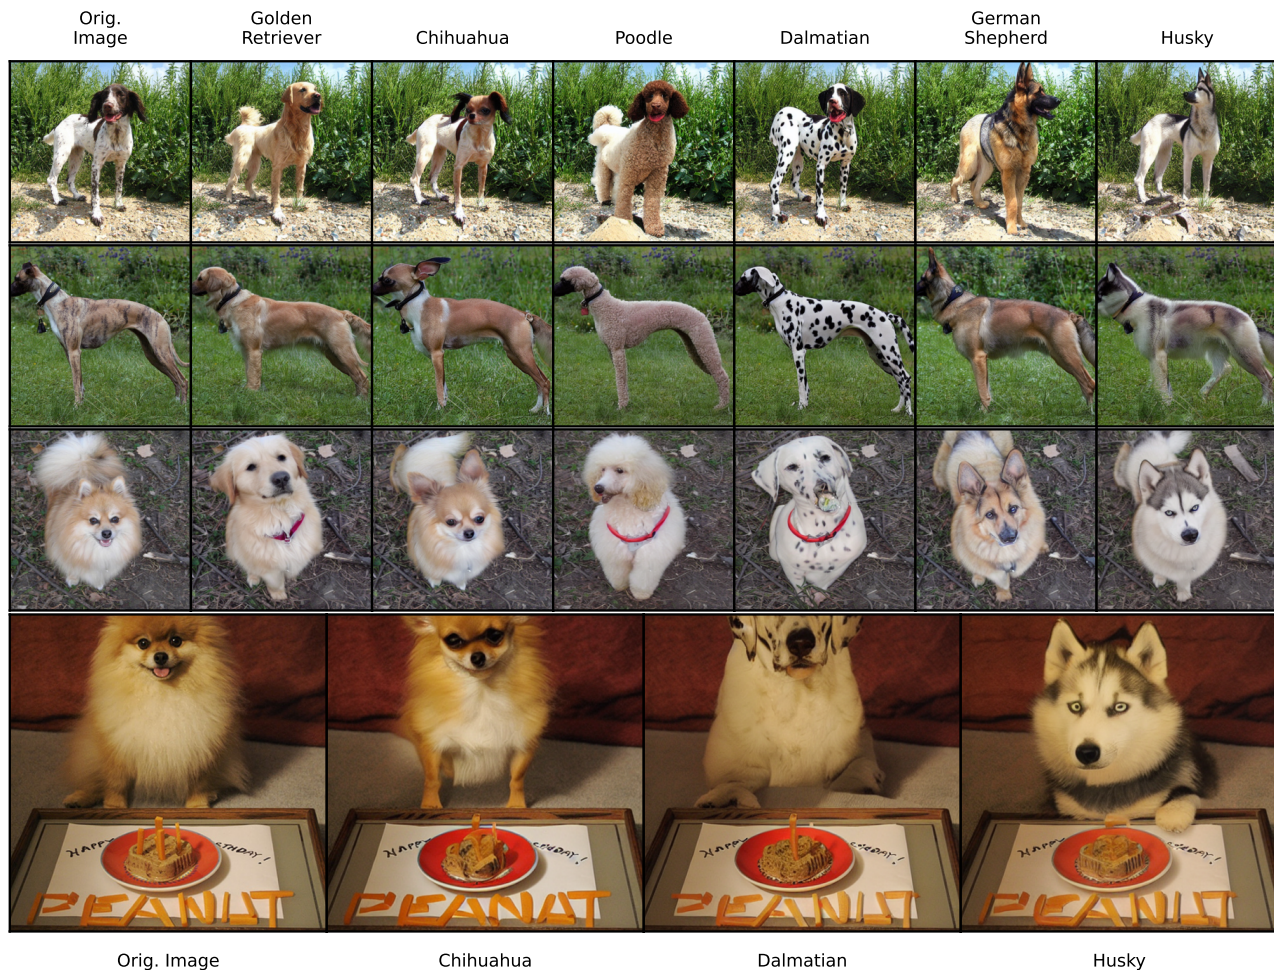


Figure 1. EDICT enables complex real image edits, such as editing dog breeds. We highlight the fine-grain text preservation in the bottom row of examples, with the message remaining even as the dog undergoes dramatic transformations. More examples, including baseline comparisons for all image-breed pairs, are included in the Supplementary. All original images from the ImageNet 2012 validation set.

nipulations through text as increase in the corruption is correlated with the strength of the conditioning.

To improve the inversion ability of DDMs and enable robust real image editing, we diagnose the problems in DDIM inversion, and offer a solution: Exact Diffusion Inversion via Coupled Transformations (EDICT). EDICT is a re-formulation of the DDIM process inspired by coupling layers in normalizing flow models [4, 5, 12] that allows for mathematically exact inversion. By maintaining two coupled noise vectors in the diffusion process, EDICT enables recovery of the original noise vector in the case of model-generated images; and for real imagery, initial noise vectors that are guaranteed to map to the original image when the EDICT generative process is run. While EDICT doubles the computation time of the diffusion process, it can be combined with any pretrained DDM model and does

not require any computationally-expensive model finetuning, prompt tuning, or multiple images.

For the standard generative process, EDICT approximates DDIM well, resulting in nearly identical generations given equal initial conditions. For real images, EDICT can recover a noise vector which yields an exact reconstruction when used as input to the generative process. Experiments with the COCO dataset [16] show that EDICT can recover complex image features such as detailed textures, thin objects, subtle reflections, faces, and text, while DDIM fails to do so consistently. Finally, using the initial noise vectors derived from a real image with EDICT, we can sample from a DDM and perform complex edits or transformations to real images using textual guidance. We show editing capabilities including local and global modifications of objects and background and object transformations (Fig. 1).




Figure 2. While both unconditional and conditional DDIM [29] often fail to accurately reconstruct real images, leading to loss of global image structure and/or finer details, EDICT is able to almost perfectly reconstruct even complex scenes. Examples from ImageNet and COCO with 50 DDIM steps. Captions used only in DDIM Conditional reconstruction with a guidance scale of 3.

2. Related Work

Diffusion Models and Normalizing Flows: Denoising diffusion models (DDMs), drawing on nonequilibrium thermodynamics [28], have emerged at the forefront of image generation. Models such as GLIDE [22], DALLE-2 [24], Imagen (Video) [8, 27], Latent/Stable Diffusion [25], and eDiffi [1] all utilize concepts borrowed from thermodynamics to hallucinate an image from pure noise by training on intermediately noised images. While a commonly used sampling process in DDMs is the stochastic Denoising Diffusion Probabilistic Models (DDPMs) method [9], a deterministic sampling method was introduced in Denoising Diffusion Implicit Models (DDIM) [29]. Multi-step or higher-order methods that parallel DDIM have also been proposed [17, 19], including methods that aim to reducing the computational time for generation [13, 18]. Another class of generative models relevant to our work are normalizing flow models [4, 5, 12]. In these models, an invertible mapping is learned between a latent gaussian distribution and image space. The methods of invertibility, specifically coupling layers, are used as an inspiration for our method. Invertible neural networks have been studied in areas outside of normalizing flows as well. Neural ODEs [2] have many parallels with the diffusion process and can be inverted using a variety of ODE solvers.

Editing in Diffusion Models: The seminal work in apply-

ing DDMs to image editing is SDEdit [20] where coarse layouts are used to guide the generative process by noising the layout to resemble an intermediately noised image. Prompt-to-Prompt [7] combines query-key pairs from one prompt with values from another in the attention layers of a DDM to enable prompt-guided image editing from intermediate latents obtained by sampler inversion [29, 31]. DiffEdit [3] edits real/synthetic images using automatically generated masks for regions of an input image that should be edited given a text query. Kwon et al. [14] introduce style and structure losses to guide the sampling process to enable text-guided image translation. CycleDiffusion [32] uses a deterministic DPM encoder to enable zero-shot image-to-image translation. Concurrent work [21] also stabilizes text-conditioned inversion, optimizing the null embedding \emptyset for consistency. While highly effective, the algorithmic exactness is not guaranteed and per-case optimization is required. Another set of methods [6, 11, 26] finetune the model with the target image and/or learn a new conditioning prompt to enable indirect image editing via sampling. EDICT, our proposed approach, does not require any specialized training/finetuning, and can be paired with any pretrained DDM.

3. Background

3.1. Denoising Diffusion Models

DDMs are trained on a simple denoising objective. A set of timesteps index a monotonic strictly increasing noising schedule $\{\alpha_t\}_{t=0}^T$, $\alpha_T = 0$, $\alpha_0 = 1$. Images (or autoencoded latents) $x \in X$ are noised with draws $\epsilon \sim N(0, I)$ according to the noising schedule following the formula

$$x_t = \sqrt{\alpha_t}x + \sqrt{1 - \alpha_t}\epsilon \quad (1)$$

The time-aware DDM Θ is trained on the objective $MSE(\Theta(x_t, t, C), \epsilon)$ to predict the noise added to the original image where C is a conditioning signal (typically in the form of a text embedding) with some probability of random assignment to the null conditioning \emptyset . To generate a novel image from a gaussian draw $\epsilon_T \sim N(0, I)$, partial denoising is applied at each t . The most common sampling scheme is that of DDIM [29] where intermediate steps are calculated as

$$x_{t-1} = \sqrt{\alpha_{t-1}} \frac{x_t - \sqrt{1 - \alpha_t} \Theta(x_t, t, C)}{\sqrt{\alpha_t}} + \sqrt{1 - \alpha_{t-1}} \Theta(x_t, t, C) \quad (2)$$

In practice, for text-to-image models to hallucinate from random noise an x_0 that matches conditioning C to desired levels, the model has to be biased more heavily towards generations aligned with C . To do so, a pseudo-gradient $G \cdot (\Theta(x_t, t, C) - \Theta(x_t, t, \emptyset))$ is added to the unconditional prediction $\Theta(x_t, t, \emptyset)$ to up-weight the effect of conditioning, where G is a weighting parameter, Substituting

$\Phi(x_t, t, C, G) = \Theta(x_t, t, \emptyset) + G \cdot (\Theta(x_t, t, C) - \Theta(x_t, t, \emptyset))$ into the prior equation for the Θ term, we simplify the notation $\Phi(x_t, t, C, G) \rightarrow \epsilon(x_t, t)$ and rewrite the previous equation as $x_{t-1} = a_t x_t + b_t \epsilon(x_t, t)$ where

$$a_t = \sqrt{\alpha_{t-1}/\alpha_t} \quad (3)$$

$$b_t = -\sqrt{\alpha_{t-1}(1-\alpha_t)/\alpha_t} + \sqrt{1-\alpha_{t-1}} \quad (4)$$

3.2. Denoising Diffusion Implicit Model (DDIM)

As noted in DDIM [29], the above denoising process is approximately invertible; that is x_t is approximately recoverable from x_{t-1}

$$x_t = \frac{x_{t-1} - b_t \epsilon(x_t, t)}{a_t} \approx \frac{x_{t-1} - b_t \epsilon(x_{t-1}, t)}{a_t} \quad (5)$$

where the approximation is a linearization assumption that $\epsilon(x_t, t) \approx \epsilon(x_{t-1}, t)$ (necessary due to the discrete nature of both computation and the underlying noise schedule). This corresponds with reversing the Euler integration which is a first-order ODE solver. More sophisticated solvers such as multi-step Euler [17] have been shown to stabilize the generative process, and correspondingly the deterministic inversion process, with fewer time steps. However, such methods are also approximations where the inversion accuracy ultimately relies on the strength of the linearization assumption and the reconstruction is not exactly equal. This assumption is largely accurate for unconditional DDIM generation, but the pseudo-gradient of classifier-free guidance $G \cdot (\Theta(x_t, t, C) - \Theta(x_t, t, \emptyset))$ is inconsistent across time steps as shown in the Supplementary.

While unconditional reconstructions have relatively insignificant errors (Fig. 2), *conditional* reconstructions are extremely distorted when noised to high levels. This phenomenon was noted in [7], where the guidance scale must be heavily downweighted in order for inversions on real-world images to be stable, thus limiting the strength of edits. Obtaining an x_t from x_0 allows for the generative process to be run with novel conditioning. In SDEdit [20], this process is done stochastically to obtain broad sample diversity, at the cost of controllability and faithfulness to the original image contents. In contrast, the inverse DDIM process produces a unique x_t from a single x_0 in a deterministic manner, yielding only one sample but enabling higher strength edits while preserving finer-grain structure and content.

3.3. Affine Coupling Layers

Affine Coupling Layers (ACL) are invertible neural network layers introduced in [4, 5] and used in other normalizing flow models such as Glow [12]. The layer input z , is split into two equal-dimensional halves z_a and z_b . A modified version of z_a is then calculated, according to:

$$z'_a = \Psi(z_b)z_a + \psi(z_b) \quad (6)$$

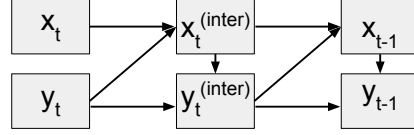


Figure 3. Information flow of EDICT. Denoising process (sampling) follows the arrows forwards, sampler inversion backwards. From (x_t, y_t) , (x_{t-1}, y_{t-1}) can be calculated while holding a single member of each sequence in memory at a time. All steps are invertible, so (x_t, y_t) can be exactly recovered from (x_{t-1}, y_{t-1}) , as opposed to current methods which compute an approximation.

where Ψ and ψ are neural networks. The layer output z' is the concatenation of z'_a and z_b in accordance with the original splitting function. ACL can parameterize complex functions and z can be exactly recovered given z' :

$$z_a = (z'_a - \psi(z_b))/\Psi(z_b) \quad (7)$$

Noting the similarity of equations 6 and 7 to the simplified form of Eq. (2), we parallel this construction in our method (described next) where two separate quantities are tracked and alternately modified by transformations that are affine with respect to the original modified quantity and a non-linear transformation of its counterpart.

4. Exact Diffusion Inversion via Coupled Transformations (EDICT)

4.1. Making an Invertible Diffusion Process

As summarized in Sec. 3.3, affine coupling layers track two quantities which can then be used to invert each other. These two quantities are partitions of a latent representation with a network specifically designed to operate in a fitting alternating manner. Without training of a new DDM, this method can not be applied out-of-the-box to the forward diffusion process. We consider the simplified form of the forward step equation from Sec. 3.1 below

$$x_{t-1} := a_t x_t + b_t \epsilon(x_t, t) \quad (8)$$

If the noise prediction term, $\epsilon(x_t, t) = \varepsilon$ was independent of x_t , this would be an affine function in both x_t , and ε . Paralleling Eq. (6), by creating a new variable $y_t = x_t$ the stepping equation fits the desired form. Consider performing this computation, so we have the variables x_t , $y_t = x_t$, $x_{t-1} = a_t x_t + b_t \epsilon(y_t, t)$. x_t can be recovered exactly from x_{t-1} in the non-trivial form:

$$x_t = (x_{t-1} - b_t \cdot \epsilon(y_t, t))/a_t \quad (9)$$

$y_t = x_t$ is trivial and will not be true in the general case. Now consider the initialization of the reverse (denoising) diffusion process, where $x_T \sim \mathcal{N}(0, 1)$, we similarly initialize $y_T = x_T$. Following the above process, we define the update rule



Figure 4. Images generated from the same text prompt and random seed. **EDICT x/y (No mixing)** employs coupling layers but not intermediate mixing layers, resulting in distorted and inconsistent images as x and y diverge. **EDICT** is our full method with the sequences explicitly contracted together, resulting in identical images that match DDIM (**Baseline**) in quality and composition. Top row prompt: *A white horse galloping through a forest* Bottom row prompt: *A couple of glasses are sitting on a table*

$$\begin{aligned} x_{t-1} &= a_t x_t + b_t \cdot \epsilon(y_t, t) \\ y_{t-1} &= a_t y_t + b_t \cdot \epsilon(x_{t-1}, t) \end{aligned} \quad (10)$$

Note that the noise prediction term in the second line is a function of the *other* sequence value at the next timestep. Only one member of each sequence ($x_t, y_{t(\pm 1)}$) must be held in memory at any given time. The sequences can be recovered exactly according to

$$\begin{aligned} y_t &= (y_{t-1} - b_t \cdot \epsilon(x_{t-1}, t)) / a_t \\ x_t &= (x_{t-1} - b_t \cdot \epsilon(y_t, t)) / a_t \end{aligned} \quad (11)$$

As illustrated in Fig. 3, the entire sequence can be reconstructed from any two adjacent x_i and y_i . In sum, our method re-uses the linearization assumption of Euler DDIM inversion, that $\epsilon(x_t, t) \approx \epsilon(x_{t-1}, t)$, but crucially does not rely on it for invertibility, guaranteeing recovery up to machine precision. We note that the functional form of our approach bears similarities to the ODE solver Heun’s method where derivative values at initial predictions are used to refine predictions, but due to the need for invertibility we cannot exploit this for more accurate/faster sampling.

4.2. Stabilization

While our method assures invertibility by design, realism and faithfulness to the original diffusion process are not automatic. When naively applied for a typical, low number of DDIM steps (e.g., $T = 50$), the sequences x_t and y_t can diverge (Fig. 4). This is a result of the strong linearization assumption not holding in practice. To alleviate this problem, we introduce intermediate *mixing layers* after each diffusion step computing weighted averages of the form

$$x' = px + (1 - p)y, \quad 0 \leq p \leq 1 \quad (12)$$

which are invertible affine transformations. Note that this averaging layer becomes a *dilating* layer during deterministic noising; the inversion being

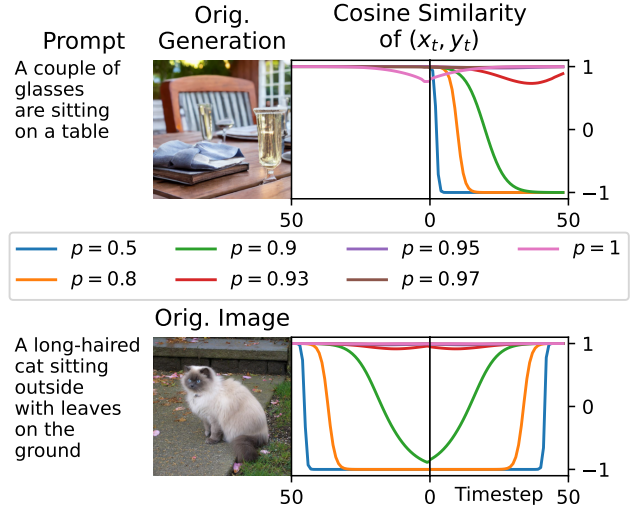


Figure 5. We visualize the effect of the mixing coefficient p . Top row: an EDICT-generated image with $p = 0.93$, for each listed p a generation is computed from $T = 50$ to $T = 0$ followed by reconstructing initial latents ($T = 50$). Cosine similarity between (x_t, y_t) is computed at each step. $p = 0.97$ suffices for generative convergence, but the *inverse* process diverges when p is too small. Repeated dilation can exponentially exaggerate small floating-point differences, annulling the theoretical guarantee. These trends hold in the bottom row for a non-generated image put through the EDICT inversion/generation process. While all reconstructions ultimately are aligned, a sufficiently large p is needed for the latents to maintain alignment throughout.

$$x = \frac{x' - (1 - p)y}{p} \quad (13)$$

A high (near 1) value of p results in the averaging layer not being strong enough to prevent divergence of the x and y series during denoising, while a low value of p results in a numerically unstable exponential dilation in the backwards pass (Fig. 5). Note that in both cases that our generative process remains exactly mathematically invertible with no further assumptions, but there is a degradation in utility of the results. Typically we employ $p = 0.93$, with values in the interval $[0.9, 0.97]$ generally being effective for 50 steps.

4.3. Complete Summary of the Method

We dub our presented process EDICT: **Exact Diffusion Inversion via Coupled Transformations**. In sum, EDICT uses a combination of “coupling” and “averaging/dilating” steps for exact inversion of the diffusion process. Given x_t and y_t , we calculate the denoising process by

$$\begin{aligned} x_t^{inter} &= a_t \cdot x_t + b_t \cdot \epsilon(y_t, t) \\ y_t^{inter} &= a_t \cdot y_t + b_t \cdot \epsilon(x_t^{inter}, t) \\ x_{t-1} &= p \cdot x_t^{inter} + (1 - p) \cdot y_t^{inter} \\ y_{t-1} &= p \cdot y_t^{inter} + (1 - p) \cdot x_{t-1} \end{aligned} \quad (14)$$

Method	LDM AE	EDICT (UC)	EDICT (C)	DDIM (UC)	DDIM (C)
50 Steps	0.015	0.015	0.015	0.030	0.420
100 Steps	0.015	0.015	0.015	0.027	0.471
200 Steps	0.015	0.015	0.015	0.023	0.497

Table 1. Mean-square reconstruction error for COCO-val using the first listed prompt as conditioning with $G = 7$. The latent diffusion model autoencoder (LDM AE) is the lower bound on reconstruction error. Using half precision increases 50-step EDICT (C) MSE by 6%. More step values are in the Supplementary.

and the deterministic noising inversion process by:

$$\begin{aligned}
 y_{t+1}^{inter} &= (y_t - (1 - p) \cdot x_t) / p \\
 x_{t+1}^{inter} &= (x_t - (1 - p) \cdot y_{t+1}^{inter}) / p \\
 y_{t+1} &= (y_{t+1}^{inter} - b_{t+1} \cdot \epsilon(x_{t+1}^{inter}, t + 1)) / a_{t+1} \\
 x_{t+1} &= (x_{t+1}^{inter} - b_{t+1} \cdot \epsilon(y_{t+1}, t + 1)) / a_{t+1}
 \end{aligned} \tag{15}$$

Recall that the conditioning C is implicitly included in the ϵ terms. In practice, we alternate the order in which the x and y series are calculated at each step in order to symmetrize the process with respect to both sequences. We cast all operations to double floating point precision (from the native half precision) to mitigate roundoff floating point errors.

4.4. Image Editing

While EDICT can theoretically operate on either pixel-based or latent-diffusion models we present the latter case in this work. Given an image I , we edit the semantic contents to match text conditioning C_{target} . We describe the current content by text C_{base} in a parallel manner to C_{target} . We compute an autoencoder latent $x_0 = VAE_{enc}(I)$, initializing $y_0 = x_0$. We run the deterministic noising process of Eq. (15) on (x_0, y_0) using text conditioning C_{base} for $s \cdot S$ steps, where S is the number of global timesteps and s is a chosen editing strength. This yields partially “noised” latents (x_t, y_t) which are not necessarily equal and, in practice, tend to diverge by a small amount due to linearization error and the dilation of the mixing layers. These intermediate representations are then used as input to Eq. (14) using text condition C_{target} , and identical step parameters, (s, S) . The resulting image outputs $(VAE_{dec}(x_0^{edit}), VAE_{dec}(y_0^{edit}))$ are empirically nearly identical as seen in Fig. 4, this is by design of the method and in particular, the mixing layers. For all methods we find that a guidance scale of 3 performs well (as opposed to the standard 7.5 for generation).

5. Experiments

We now describe results using the Stable Diffusion 1.4 latent diffusion model [25].

5.1. Image Reconstruction

We demonstrate the exact invertibility of EDICT using the MS-COCO-2017 validation set ($n = 5,000$), which contains both simple object-centric images and complex scene images [16]. Given an image-caption pair, inverted latents are calculated and used to reconstruct the image. Mean-square error is calculated on pixels normalized to $[-1, 1]$ and averaged across all images in the dataset. This process is performed both with and without the text as conditioning (C vs. UC). For COCO, we use the first listed prompt as conditioning. The LDM autoencoder reconstruction error serves as a lower bound. EDICT maintains complete latent recovery in all examples for both 50 and 200 steps, with error 50-75% that of DDIM (UC) (Tab. 1). DDIM (C) is unstable for inversions (also noted in [7]), which results in an error an order of magnitude greater than any other, with failure to reconstruct as shown in Fig. 2.

5.2. Image Editing

We show EDICT’s ability to perform complex editing tasks on real images in Fig. 6. In the first row, a diverse set of objects are added to a lake scene, demonstrating **object addition**. In the *giraffe* and *car* examples, we see interaction between introduced and original objects. Additionally, both the *giraffe* and *castle* examples capture reflections of the added objects, while the pattern of the water is maintained. Throughout all edits, details such as the cloud patterns and patches of tree color are preserved. The second row shows **object-preserving global changes**. A chair is placed into a variety of settings while keeping near-perfect identity and detail, even when occlusions are generated (*grass* and *snow*). In all examples the chair keeps a realistic footing, despite ground changes.

In the third row, we demonstrate that EDICT is able to perform **object deformations**, a challenging class of edits for DDM methods, as many broad compositional components are determined very early in the generation [7]. The previously most successful method for these types of edits, Imagic [11] requires model finetuning. EDICT makes the sculpture assume a broad set of poses, spatially changing the semantic map in a refined way. Novel views, such as “A statue from behind” are able to be plausibly rendered. Across these large-scale edits, fine-grained details such as the face and dress of the statue—as well as the foliage and path—are preserved. In the fourth row, we show that EDICT is able to perform **global style changes** while maintaining layout and details where appropriate. The layout is nearly identical across images (note the preserved cloud pattern), but can be changed when needed (e.g., the lack of trees in the *Mars* example). Specific art-styles are capable of being generated, including challenging concepts such as cubism.

In Fig. 1, EDICT makes **semantic entity edits** to and from a variety of dog breeds. It proves adept at holding the

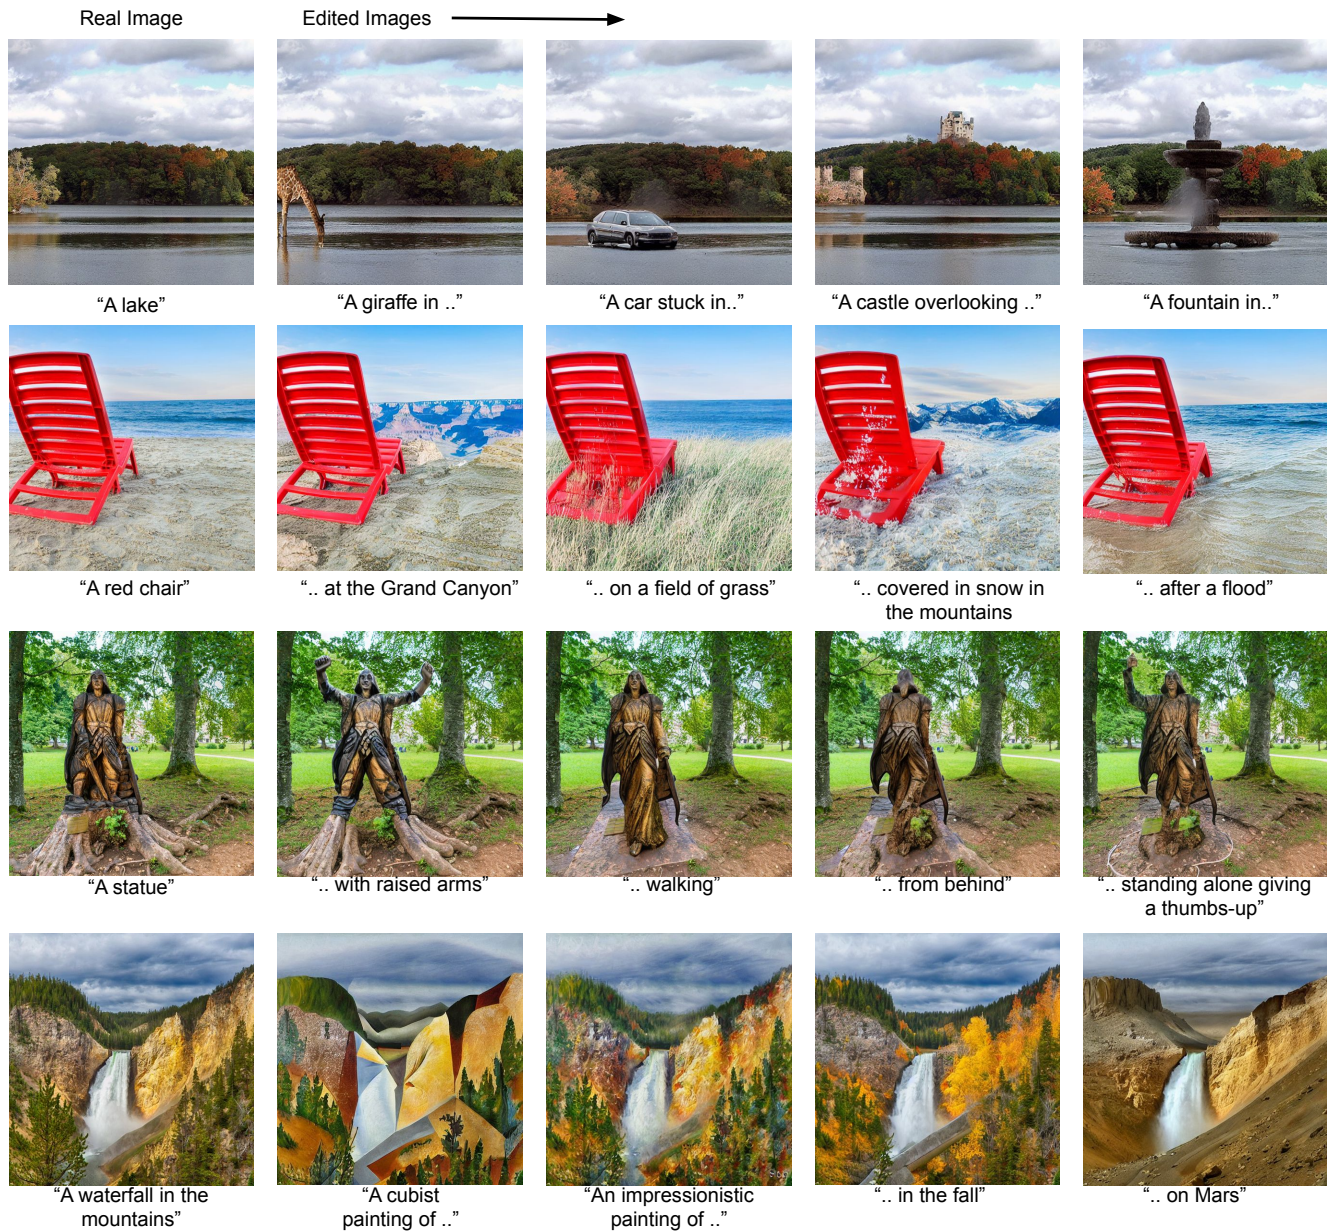


Figure 6. Examples of the strength and varied editing ability of EDICT. From top to bottom: object additions, global context changes with object preservation, deformations while preserving identity, and global style changes. The bidirectional reversibility of EDICT’s process allows for large-scale changes to the image while maintaining auxiliary details.

original subject pose, including in the third row where the dog is viewed in a very atypical position. The realism in the *Chihuahua* examples is particularly interesting, due to the relatively small size of the breed. We again highlight the preservation of details such as the background foliage or ground in the upper three rows. The fourth row is of specific interest as small text is preserved across examples (a typical failure case for DDIM).

Baseline Comparison: In Fig. 7, we demonstrate EDICT’s superior performance to other DDM sampling-based methods for image editing: conditional and unconditional DDIM inversion, prompt-to-prompt image editing, and SDEdit (as a stochastic baseline). All methods are run with 50 steps (we do not observe improvement of baselines when the number of steps are increased, see Supplementary). Since methods that require model finetuning or prompt tuning [6, 11, 26] can be combined with EDICT, we view them



Figure 7. EDICT compared to baseline editing methods. EDICT demonstrates superior preservation of unedited components while still performing the edit. Further comparison between EDICT and baselines (including concurrent work ([14, 32]) is in the Supplementary.

as complementary to, rather than competitive with, EDICT, and do not perform a direct comparison with them. We quantitatively compare visual metrics of edits in Fig. 8.

6. Discussion

Limitations and Future Work: EDICT is deterministic, only outputs one generation per image-prompt pair (vs. SDEdit). Also, the computational time is approximately twice that of baseline DDIM. As with all editing methods, performance can vary across inputs in a hard-to-predict manner and sometimes requires careful prompt selection. For future work, we note that the induced latent space constructed by the inversion process admits operations such as latent interpolation [10] which has not been widely applied to real images. As in prompt tuning [15], formalizing the process of prompt selection could further improve EDICT. Adding a controllable degree of randomness to EDICT could yield multiple candidate generations per input.

Ethics: Like other image generation and editing models, EDICT will produce images that may reflect the socioeconomic biases of the training data or images that could be considered inappropriate. Image editing methods can also be utilized for malicious purposes, including harassment and misinformation spread. Practitioners utilizing EDICT or its methodologies in a production setting should consider these limitations. An in-depth discussion on the ethics of image generation can be found in Imagen [27].

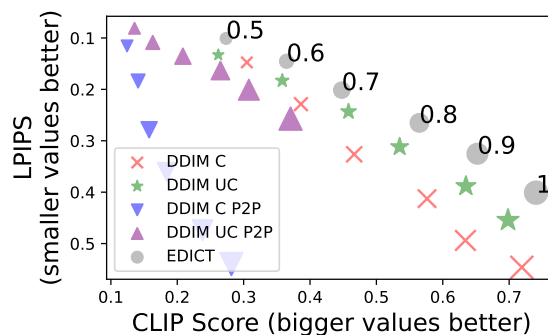


Figure 8. We quantitatively benchmark all methods on image editing using images from five ImageNet mammal classes (*African Elephant*, *Ram*, *Egyptian Cat*, *Brown Bear*, and *Norfolk Terrier*). Four experiments are performed, one swapping the pictured animal’s species to each of the other classes (20 species editing pairs in total), two contextual changes (*A [animal] in the snow* and *A [animal] in a parking lot*), and one stylistic (*An impressionistic painting of a [animal]*). The LPIPS [33] (visual similarity) score to the original image and a CLIP score [23] (semantic similarity) to a set of five related text queries is computed per edit. Metrics are averaged across images and experiments. Details and human evaluation are provided in the Supplementary. We sweep across the fraction of inverted steps (overlaid in black on top of gray EDICT bubbles), observing that EDICT forms a Pareto frontier.

References

- [1] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. [3](#)
- [2] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018. [3](#)
- [3] 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](#)
- [4] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014. [2](#), [3](#), [4](#)
- [5] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016. [2](#), [3](#), [4](#)
- [6] 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. [3](#), [7](#)
- [7] 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](#), [3](#), [4](#), [6](#)
- [8] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. [3](#)
- [9] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. [1](#), [3](#)
- [10] Andrej Karpathy. Stable diffusion latent interpolation. <https://gist.github.com/karpathy/00103b0037c5aaea32fe1da1af553355>. Accessed Nov. 10 2022. [8](#)
- [11] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. *arXiv preprint arXiv:2210.09276*, 2022. [3](#), [6](#), [7](#)
- [12] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018. [2](#), [3](#), [4](#)
- [13] Zhifeng Kong and Wei Ping. On fast sampling of diffusion probabilistic models. *CoRR*, abs/2106.00132, 2021. [3](#)
- [14] Gihyun Kwon and Jong Chul Ye. Diffusion-based image translation using disentangled style and content representation. *arXiv preprint arXiv:2209.15264*, 2022. [3](#), [8](#)
- [15] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *Empirical Methods in Natural Language Processing*, pages 3045–3059, Nov. 2021. [8](#)
- [16] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. *European conference on computer vision*, pages 740–755, 2014. [2](#), [6](#)
- [17] Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. *arXiv preprint arXiv:2202.09778*, 2022. [3](#), [4](#)
- [18] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. [3](#)
- [19] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *arXiv preprint arXiv:2206.00927*, 2022. [3](#)
- [20] 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. [1](#), [3](#), [4](#)
- [21] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. *arXiv preprint arXiv:2211.09794*, 2022. [3](#)
- [22] 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. [3](#)
- [23] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [8](#)
- [24] 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](#), [3](#)
- [25] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *CVPR*, pages 10684–10695, 2022. [1](#), [3](#), [6](#)
- [26] 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. [3](#), [7](#)
- [27] 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, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. [1](#), [3](#), [8](#)
- [28] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. *CoRR*, abs/1503.03585, 2015. [3](#)

- [29] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 1, 3, 4
- [30] Xuan Su, Jiaming Song, Chenlin Meng, and Stefano Ermon. Dual diffusion implicit bridges for image-to-image translation. *arXiv preprint arXiv:2203.08382*, 2022. 1
- [31] Github User trygvebw. Sampler inversion gist. <https://gist.github.com/trygvebw/c71334dd127d537a15e9d59790f7f5e1>. Accessed Nov. 10 2022. 3
- [32] 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, 8
- [33] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 8