

End-to-End Diffusion Latent Optimization Improves Classifier Guidance

Bram Wallace
Salesforce AI

b.wallace@salesforce.com

Akash Gokul
Salesforce AI

agokul@salesforce.com

Stefano Ermon
Stanford University

ermon@cs.stanford.edu

Nikhil Naik
Salesforce AI

nnaik@salesforce.com

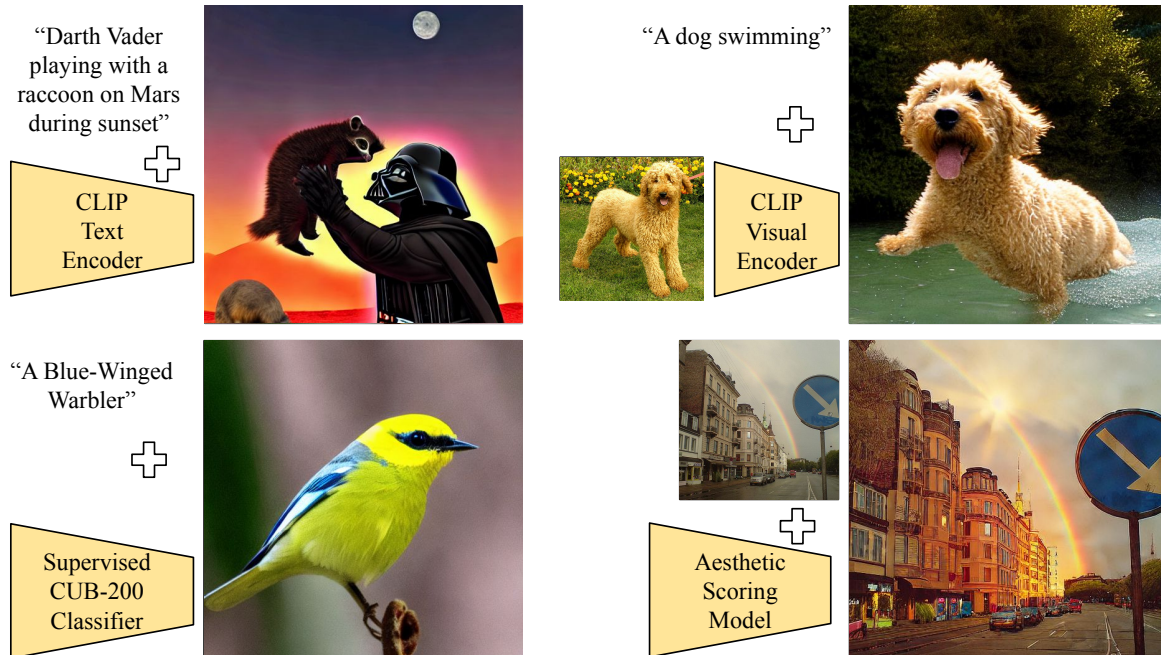


Figure 1: We propose DOODL - a process that directly optimizes diffusion latents w.r.t. a model-based loss on the final generation. Our method improves on vanilla classifier guidance in all tested settings and we demonstrate capabilities novel to this class of methods such as vocabulary expansion, entity personalization, and perceived aesthetic value improvement.

Abstract

Classifier guidance—using the gradients of an image classifier to steer the generations of a diffusion model—has the potential to dramatically expand the creative control over image generation and editing. However, currently classifier guidance requires either training new noise-aware models to obtain accurate gradients or using a one-step denoising approximation of the final generation, which leads to misaligned gradients and sub-optimal control. We highlight this approximation’s shortcomings and propose a novel guidance method: *Direct Optimization of Diffusion Latents (DOODL)*, which enables plug-and-play guidance by optimizing diffusion latents w.r.t. the gradients of a pre-trained classifier on the true generated pixels, using an invertible diffusion process to achieve memory-efficient

backpropagation. Showcasing the potential of more precise guidance, DOODL outperforms one-step classifier guidance on computational and human evaluation metrics across different forms of guidance: using CLIP guidance to improve generations of complex prompts from DrawBench, using fine-grained visual classifiers to expand the vocabulary of Stable Diffusion, enabling image-conditioned generation with a CLIP visual encoder, and improving image aesthetics using an aesthetic scoring network.

1. Introduction

Text-conditioned denoising diffusion models (DDMs), such as Latent/Stable Diffusion, Imagen and numerous others, have been widely adopted for synthesizing realistic im-

ages given an input text prompt [35, 36, 38, 34, 2]. The *conditioning* setup requires DDMs to be trained using paired data of images and the conditioning modality—image-caption pairs in the case of text conditioning. Once trained, the DDM can be steered to generate images using the conditioning modality. However, the *conditioning* paradigm constricts the image generation capabilities of a trained DDM model; a text-conditioned diffusion model cannot readily utilize other modalities such as depth maps or image classifiers or audio as conditioning signals. While DDMs conditioned on multiple modalities have been trained [30, 36] a “plug-and-play” approach that allows for a pretrained DDM to be guided by any external function that determines whether some generation criterion is satisfied is desirable.

In principle, *classifier guidance* [43, 8] enables this capability in DDMs. Classifier guidance, named so since it was first demonstrated using pretrained image classification models, combines the score estimate of the diffusion model with the gradient of the image classifier to steer the generation process to produce images that correspond to a particular class. Any differentiable loss function can be used for classifier guidance. In addition to class-conditional generation, it has been shown to improve compositionality using cross-modal guidance [24]. There are two existing ways of incorporating classifier guidance. In the first approach, we train a noise-aware classifier that can be used to compute an accurate gradient w.r.t. an intermediate generation step [43, 8]. This approach requires re-training the classifier model, which can be computationally expensive/infeasible due to lack of access to training data. In the second approach, at a time step t , we denoise the image with a single application of the DDM and compute the gradient using this approximately denoised image [24]. This one-step approximation is necessitated by the prohibitive memory requirements of computing a gradient w.r.t. the latents through the entire diffusion process containing many steps. Since this approach only obtains gradients using an one-step denoising approximation of the final generation, the approximately denoised images are often misaligned with the final generations which classifier guidance is aiming to modify (Figure 2), leading to sub-optimal guidance signal.

To enable flexible and exact model guidance, without noise-aware classifiers or approximations, we propose Direct Optimization Of Diffusion Latents (DOODL). DOODL optimizes the initial diffusion noise vectors w.r.t. a model-based loss on images generated from the full-chain diffusion process. We leverage EDICT, a recently developed drop-in *discretely invertible* diffusion algorithm [46], which admits backpropagation with constant memory cost w.r.t the number of diffusion steps, to compute classifier gradients on the pixels of the final generation w.r.t. the original noise vectors. This enables efficient iterative optimization of diffusion latents w.r.t. any differentiable loss on the image pixels




Figure 2: Stable Diffusion with prompt “A kangaroo in a field” using 50 DDIM steps. vs. a single step. The one-step approximation is inaccurate for high noise levels, even though it is relied upon by classifier guidance. DOODL calculates loss gradients with respect to fully formed realistic images as opposed to approximations.

and accurate calculation of gradients for classifier guidance.

We demonstrate the efficacy of DOODL through a diverse set of guidance signals (Figure 1) using quantitative and human evaluation studies. First, we show that CLIP classifier guidance using DOODL improves generation of images guided by text prompts from the DrawBench [38] dataset, which test compositionality and the ability to guide using unusual captions. Second, we show the ability to expand the vocabulary of a pretrained Stable Diffusion model using fine-grained visual classifiers; a capability that one-step classifier guidance does not have. Third, we demonstrate that DOODL can be used for personalized entity generation (e.g. “A dog in sunglasses”), *with zero retraining of any new network*—a first to our knowledge. Finally, we utilize DOODL to perform a novel task; increasing the perceived aesthetic quality of generated/real images. We hope that DOODL can enable and inspire diverse plug-and-play capabilities for pretrained diffusion models.

2. Related Work

2.1. Text-to-Image Diffusion Models

Text-to-image diffusion models [17, 43, 41] such as GLIDE [28], DALLE-2 [34], Imagen [38, 16], Latent Diffusion [35, 36], and eDiffi [2] have recently emerged at the forefront of image generation, based on methods in non-equilibrium thermodynamics [39]. Classifier guidance [43, 8], and its adoptions [28, 24], use the gradients of pretrained classifier models to guide such generations. Instead of sequential denoising, [43] traverses at a fixed noise level before each denoising step. Concurrent work [3] modifies classifier guidance to refine the gradient prediction at each noise level before continuing. DiffuseIT [23] is another work that uses pretrained recognition models in conjunction with DDMs. Specifically, it performs either text- or image-guided image translation, leveraging a DINO ViT [5] model to enforce structural similarity of a DDM generation to a source image via a contrastive loss and a CLIP model to steer the image semantics towards the desired tar-

Method Type	Requirements		
	Training	Data	Pretrained “Classifier”
<i>Learning-Based</i> [12, 37, 48]	✓	✓	✗
<i>Guidance-Based</i> (DOODL, Clf. Guidance)	✗	✗	✓

Table 1: *Learning-Based* methods require data and training while *Guidance-Based* methods require pretrained recognition networks (trained on non-noisy data in our setting).

get. DiffuseIT calculates updates at incremental noise levels as opposed to DOODL which optimizes x_t based off of x_0 . Table 1 shows the requirements for learned methods such as ControlNet [48] vs. these guidance-based methods. The former require *data* and *training*, while the latter substitute this need for *pretrained recognition models*. Directly optimizing latent variables of other generative models such as GANs or VAEs [11, 32, 29] w.r.t. a pixel-based loss has been shown to generate in-distribution images that fit targeted criteria. To the best of our knowledge, these techniques have not been applied to diffusion models before DOODL.

At the intersection of diffusion models and invertible neural networks is a recently proposed approach called EDICT [46] which algorithmically reformulates the denoising diffusion process to be invertible. This prior work focuses solely on the applications to image editing and does not consider properties of Invertible Neural Networks (INNs) or similar processes. Methods such as DDIM are theoretically invertible in the limit of discretization, but this limit cannot be achieved in practice [46].

2.2. Invertible Neural Networks (INNs)

While neural networks tend to be non-dimensionality-preserving functions, there has been prior work on constructing reversible architectures. A predominant class of such INNs are normalizing flow models [9, 10, 22]. A modified version of the “coupling layers” in normalizing flow architectures are incorporated into the EDICT [46] algorithm employed in this work. [4] proposes an architecture guaranteed to be invertible via well-conditioned inverse problems instead of a closed-form solution. The memory savings of such architectures has been used in long-sequence recurrent neural networks [26] and to study inverse problems [1].

3. Background

3.1. Invertible Neural Networks w.r.t Memory

When using gradient descent for optimizing neural networks, with network parameters $\Xi = \{\xi_p\}_{p=1}^{p=P}$, network input x , network output $y = f(x)$, and loss function c ,

the derivative $\frac{dc(y)}{d\xi}$ is calculated and gradient descent performed to minimize $\mathbb{E}_{Data}c(y) = \mathbb{E}_{Data}c(f(x))$. Here f is implicitly conditioned on Ξ . Consider f as the composition of n functions (layers) $f^n \circ f^{n-1} \circ \dots \circ f^1$. To optimize ξ , a parameter of the i^{th} layer f^i the derivative $\frac{dc(y)}{d\xi}$ is calculated. Denote $f^k \circ f^{k-1} \circ \dots \circ f^j = F_j^k$. Since $y = F_1^n(x)$ the derivative w.r.t. ξ can be calculated using the chain rule:

$$\begin{aligned} \frac{dc(y)}{d\xi} &= \frac{dc(F_1^n(x))}{d\xi} \\ &= \frac{dc(F_1^n(x))}{dF_1^n(x)} \cdot \frac{dF_1^n(x)}{dF_1^{n-1}(x)} \dots \frac{dF_1^i(x)}{dF_1^{i-1}(x)} \cdot \frac{dF_1^{i-1}(x)}{dx} \end{aligned} \quad (1)$$

In the general case, calculating $\frac{dc(y)}{d\xi}$ requires storing all intermediate activations, a key bottleneck to backpropagation. Network sharding [33] across processors reduces the per-processor hardware memory requirement but the total remains the same. Gradient checkpointing [13, 49] decreases memory cost, but increases runtime linearly w.r.t. saved memory. INNs can recover intermediate states/inputs from the output, reducing memory costs by avoiding activation caching. If every f^j is invertible in Equation (1), the denominator terms can be reconstructed during the backwards pass; such methods have been leveraged to train large INNs much faster than non-invertible equivalents(Section 2.2).

3.2. Denoising Diffusion Models (DDMs)

Image DDMS are trained to predict the noise ϵ added to an image x [40, 17, 8, 18, 42]. Noise levels are discretized into a set $\mathcal{T} = \{0, 1, \dots, T\}$ that index a noising schedule $\{\alpha_t\}_{t=0}^T$, $\alpha_T = 0$, $\alpha_0 = 1$. $t \in \mathcal{T}$ are randomly sampled during training and paired with data $x^{(i)}$ (images or autoencoded representations) to generate noisy samples

$$x_t^{(i)} = \sqrt{\alpha_t}x^{(i)} + \sqrt{1 - \alpha_t}\epsilon \quad (3)$$

where $\epsilon \sim N(0, I)$. DDMs conditioned on the timestep t and auxiliary information (e.g, image caption) C are trained to approximate the added noise $DDM(x_t^{(i)}, t, C) \approx \epsilon$. At generation time, an $x_T \sim N(0, 1)$ is drawn and the DDM is iteratively applied to hallucinate a real image from the noise. Following the DDIM [40] sampling model, the final generation x_0 is equal to the composition of S denoising functions: applications of Θ conditioned on C and varying timesteps t . Denoting $\Theta(x, t, C)$ as $\Theta_{(t,C)}(x)$

$$x_0 = [\Theta_{(0,C)} \circ \Theta_{(1,C)} \circ \dots \circ \Theta_{(T,C)}](x_T) \quad (4)$$

3.2.1 Classifier Guidance

In addition to C , other guidance signals can steer generations to objectives. The foremost example, “classifier

guidance” incorporates gradients of a loss (c_{clf} , from a classifier network Φ) on estimated pixels into the noise prediction [43, 8]. From a theoretical perspective, this typically is the gradient of the log-conditional probability $\nabla \log p_{\Phi}(y|x_t)$. There are two primary ways of incorporating classifier guidance:

1: A *noise-aware* classifier is trained for direct use on intermediate (noisy) x_t , with $\nabla_{x_t} c_{clf}(x_t)$ incorporated into the denoising prediction [28]. Training noise-aware models is effective but often infeasible due to computational expense and data availability (e.g. medical images). This results in there being very few publicly available noise-aware models.

2: x_0 is approximated with a single application of $\Theta_{(t,C)}$ as in the training objective [24]. Here, the incorporated gradient is $\nabla_{x_t} c_{clf}(x_0^*)$, where x_0^* is a 1-step approximate solve by substituting ϵ for $\Theta_{(t,C)}$ in Equation (3). While a standard (noise-unaware) model can be used, the gradients are calculated w.r.t. an *approximation* of x_0 (Figure 2) and as such may be misaligned with the true derivative of $\frac{dc_{clf}(x_t)}{dx_0}$.

3.2.2 Exact Inversion of the Diffusion Process

Recently, EDICT [46], an exactly invertible variant of the discrete (time-stepped) diffusion process was proposed. EDICT operates on a latent pair (x_t, y_t) instead of a single variable. Initially $x_T = y_T \sim N(0, I)$, followed by iteratively denoising using the reverse diffusion process:

$$\begin{aligned} x_t^{inter} &= a_t \cdot x_t + b_t \cdot \Theta_{(t,C)}(y_t) \\ y_t^{inter} &= a_t \cdot y_t + b_t \cdot \Theta_{(t,C)}(x_t^{inter}) \\ x_{t-1} &= p \cdot x_t^{inter} + (1-p) \cdot y_t^{inter} \\ y_{t-1} &= p \cdot y_t^{inter} + (1-p) \cdot x_{t-1} \end{aligned} \quad (5)$$

where (a_t, b_t) are time-dependent coefficients, and $p \in [0, 1]$ is a mixing parameter to mitigate latent drift. Intuitively, this process first updates the x and y sequences based on the current state of the counterpart, and then invertibly “averages” them together. The above equations admit linear solves to invert them, defining the inverse process:

$$\begin{aligned} y_{t+1}^{inter} &= (y_t - (1-p) \cdot x_t) / p \\ x_{t+1}^{inter} &= (x_t - (1-p) \cdot y_{t+1}^{inter}) / p \\ y_{t+1} &= (y_{t+1}^{inter} - b_{t+1} \cdot \Theta_{(t+1,C)}(x_{t+1}^{inter})) / a_{t+1} \\ x_{t+1} &= (x_{t+1}^{inter} - b_{t+1} \cdot \Theta_{(t+1,C)}(y_{t+1})) / a_{t+1} \end{aligned} \quad (6)$$

We employ this construction in DOODL, and use Equation (6) to encode images x_0 to latents x_T in Section 5.3.

4. Direct Optimization of Diffusion Latents

We aim to develop a method that overcomes the shortcomings of classifier guidance discussed in Sec-

tion 3.2.1. Concretely, a method that (1) does not require re-training/finetuning an existing pretrained classifier model, (2) computes gradients w.r.t. the true output instead of a one-step approximation, and (3) incorporates the guidance in a semantically meaningful way (as opposed to an adversarial-style perturbation). We emphasize the last point in particular. As shown in the literature on adversarial attacks [44, 47, 14], gradients w.r.t. pixels can satisfy a classifier loss while not perceptually changing the content of an image. This is as opposed to techniques such as latent optimization in GANs, where the regularization provided by the decoder means that optimization happens in a space where perturbations typically result in perceptually meaningful changes that satisfy the desired objective. In this work, we aim to directly optimize *diffusion latents*, a first in the literature to our knowledge.

As shown in Equation (4), it is mathematically trivial to optimize x_T for a desired outcome on x_0 . There is a closed form expression for $\frac{dx_0}{dx_T}$ as in Equation (1). However, due to activation caching, the naive memory cost is linear in the number of DDIM sampling steps due to the T applications of Θ . With a typical value of $S = 50$, this memory cost nears a terabyte for state-of-the-art diffusion models which is impractical for most uses. Gradient checkpointing trades memory for computational complexity, resulting in the computational complexity of each backwards pass increasing by a factor of S if memory costs are held constant.

We draw inspiration from INNs (Section 3.1) to optimize x_T w.r.t. criteria on x_0 in feasible runtime. Using invertible $\Theta_{(i,C)}$ in Equation (4), intermediate states can be reconstructed during the backwards pass using only a constant number of applications of Θ w.r.t. T , circumventing the prohibitive memory cost without sacrificing runtime.

We turn to the recently developed EDICT [46] as an invertible reverse diffusion process that admits constant-memory implementation of optimization of x_t . Given conditioning C , differentiable model-based cost function c and a latent draw $x_T^{(0)}$, the EDICT generative process is performed (50 steps, $p = 0.93$, StableDiffusion v1.4) yielding initial output $f(x_T^{(0)}) = x_0^{(0)}$, which is then used to calculate a loss $c(x_0^{(0)})$ and corresponding gradient $\nabla_{x_t} c(f(x_T^{(0)}))$. This gradient can then be used to perform a step of gradient descent optimization on $x_T^{(0)}$.

We modify vanilla gradient descent in several key ways to achieve realistic images that satisfy the guidance criteria. After each optimization step, the EDICT “fully noised” latent pair $x_T^{(j)}$ and $y_T^{(j)}$ (from Equations (5) and (6)) are averaged together and renormalized to the original norm of the initial draw $x_T^{(0)}$. The averaging prevents latent drift which degrades quality, as noted in [46]. Normalizing to the original norm keeps the latent on the “gaussian shell” and in-distribution for our diffusion model.

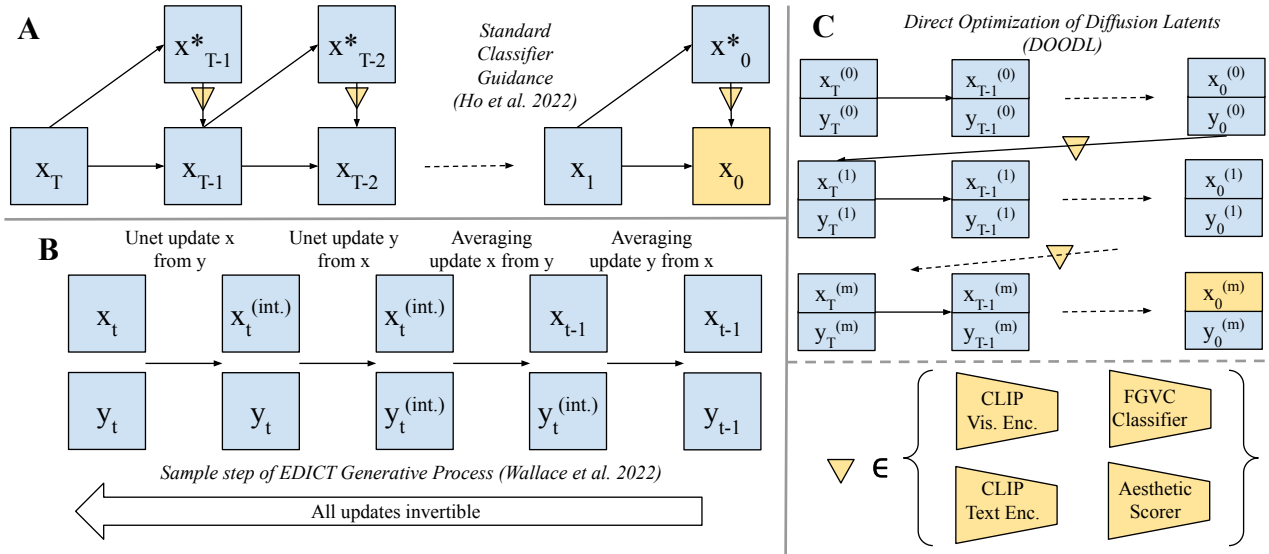


Figure 3: Method diagrams. **A** Standard classifier guidance: at each timestep, t , a one-step denoising approximation of x_0 is computed and the loss is calculated w.r.t the pixels of this generation. The gradient of this loss is incorporated into the subsequent diffusion step. **B** EDICT [46], an invertible variant of the diffusion process which admits backpropagation through the entire chain with no additional memory cost. **C** DOODL, our proposed method. We leverage EDICT and demonstrate that the gradients of model losses computed w.r.t. the final generation can be used to optimize the fully noised x_T directly. ∇ indicates a gradient calculation from a differentiable model-based loss with networks employed in this work displayed.

Method	Vanilla T2I	Clf. Guidance	DOODL
U-Net Calls	T	T	$2mT$
Scorer Calls	0	T	m
Runtime (s)	6	33	$45 \cdot m$

Table 2: Functional calls for guiding a diffusion process of T steps with m optimization steps. A factor of 2 for classifier-free guidance is omitted due to parallelization. Metrics for A100 GPU with $T = 50$ shown. Note backwards passes take relatively much longer than the forward counterparts. Typically $m \in [10, 50]$.

We also perform multi-crop data augmentation on the generated (x_0, y_0) , sampling 16 crops per image (details in Supplementary). Momentum with $\eta = 0.9$ is employed. We do not find Nesterov momentum to be useful. Finally, to increase stability and realism of outputted images, we perform element-wise clipping of g at magnitude 10^{-3} and perturb x_T by $\mathcal{N}(0, 10^{-4} \cdot I)$ at each update. Our algorithm is formally laid out in Algorithm 1. Figure 3 provides an overview of our method alongside classifier guidance and Table 2 compares computational complexity as well as wall-clock time on an A100 GPU.

We note that DOODL can be theoretically grounded by considering it as Langevin dynamics. In the general form of Langevin dynamics, the movement of a particle in a space with energy potential E (particle moves high to low poten-

tial) is described by the update rule:

$$x_{t+1} = x_t - \nabla E \Delta t + \epsilon \Delta t \quad (7)$$

where ϵ is a normal noise Brownian motion term. The update rule of DOODL,

$$g_i = 0.9 \cdot g_{i-1} + \text{Clip}(-\lambda \cdot \nabla_{x_T} L) \quad (8)$$

$$x_T^{(i+1)} = x_T + g_i + \mathcal{N}(0, 10^{-4} \cdot I) \quad (9)$$

is sum of the previous position, velocity (with momentum, induced by energy potential L), and drift term.

5. Applications and Results

Previous work has used dataset- or webly-supervised classifiers such as CLIP [32, 28]. We consider these as well as previously unexplored settings. The optimization learning rate, λ , is a key hyperparameter and we note its values throughout. The focus of our work is on how the gradients from *non-noise-aware* classifiers can be better incorporated into the diffusion process. As such, we do not consider noise-aware classifier methods [28]. For larger images and other example results, please refer to the Supplementary.

5.1. Reinforcing Text Guidance

Classifier guidance using a CLIP text-image similarity loss has been used to reinforce the text conditioning C pro-

Algorithm 1 DOODL

Input: λ (Learning rate), C (Model conditioning), c (Cost function), T (# Diffusion steps), m (# Optimization steps), Θ (Diffusion Model)
Initialize $x_T^{(0)} = \mathcal{N}(0, I)$.
 $g_{-1} = 0$
for $i = 0$ **to** m **do**
 $x_T, y_T = x_T^{(i)}$
for $t = T - 1$ **to** 0 **do**
 $x_t^{inter} = a_t \cdot x_t + b_t \cdot \Theta_{(t,C)}(y_t)$
 $y_t^{inter} = a_t \cdot y_t + b_t \cdot \Theta_{(t,C)}(x_t^{inter})$
 $x_{t-1} = p \cdot x_t^{inter} + (1 - p) \cdot y_t^{inter}$
 $y_{t-1} = p \cdot y_t^{inter} + (1 - p) \cdot x_{t-1}$
end for
 $L = 0.5 \cdot (c(MultiCrop(x_0)) + c(MultiCrop(y_0)))$
 $g_i = -\lambda \cdot \nabla_{x_T} L$ **Linear in T b/c of EDICT**
 $g_i = Clip(g_i, -10^{-3}, 10^{-3})$
 $g_i = 0.9 \cdot g_{i-1} + g_i$
 $x_T^{(i+1)} = x_T + g_i + \mathcal{N}(0, 10^{-4} \cdot I)$
end for
yield $x_T^{(m)}$

vided as input to a text-to-image diffusion model [28]. Previous work [24] shows that classifier guidance with gradients calculated on 1-step approximations of x_0 enhance the ability to generate complex prompts. This setting serves as an initial validation of our proposed method.

We employ a spherical distance loss from the literature:

$$d(x, y) = 2\lambda \sqrt{\sin^{-1}\left(\frac{\|x - y\|}{2}\right)} \quad (10)$$

computing a distance between the embedding of C by the CLIP text encoder, $CLIP_{Text}(C)$, and the embedding of $x_0^{(i)}$ by the image encoder, $CLIP_{Image}(x_0^{(i)})$. $\lambda \in [10, 100]$ is effective for the baseline one-step guidance, we use $\lambda = 30$. For DOODL, we find $\lambda = 0.1$ to be preferable. We attribute the lower value to the gradient magnitudes increasing when computed through repeated applications of the network, as well as the incorporation of momentum.

We test on the DrawBench benchmark [38], a collection containing 11 categories of text prompts designed to probe different aspects of text-to-image generation. In particular, we evaluate on the *DALLE* and *Reddit* categories of DrawBench, which focus on compositionality and highly implausible scenes, being especially suited for testing fidelity to conditioning text. Image-text alignment is measured both automatically (CLIP score) and with human evaluation in Figure 5. Experiments are performed across 9 seeds for all method-prompt pairs. We use a LAION-trained CLIP model [20] for classifier guidance and an OpenAI-trained



Figure 4: Generations using DrawBench prompts[38] for DOODL and baselines from the same seed. Note the prompting style of the *DALLE* category uses multiple sentences to reinforce a single concept.

Per-Class % FID Change (Lower Better)			
Dataset	CUB	Dogs	Aircraft
DOODL	-3.2%	-5.3%	-0.4%
Baseline	+6.1%	+0.27%	+13.7%

Table 3: Quantitative results for rare vocabulary generation. Per-class FID is measured between generated samples and the validation set. Change relative to the original SD generations is shown. DOODL achieves a more similar set of images in all instances, while the baseline fails in all.

CLIP model [32] for CLIP score evaluation. While there is an overlap in the training sets of these models; we find them sufficiently independent to provide automated validation. Generated samples are shown in Figure 4.

DOODL relatively improves on baseline classifier guidance’s CLIP score by 3.1% and 2.6% on *DALLE* and *Reddit* respectively. For human study, we define prompt alignment success rate (PASR) as the percent of generations rated as humans to be aligned with the prompt in head-to-head comparisons (see Figure 5 caption). The PASR of DOODL is highest, notably increasing on the complex unusual prompts of *Reddit* by 11.3% as compared to vanilla Stable Diffusion, with a 6% increase over one-step classifier guidance.




Figure 5: DrawBench results. **Left:** CLIP Scores. **Right:** Human evaluation. 5 labelers are given a pairwise comparison of DOODL to another method for the same prompt-seed inputs and asked which represents the prompt better, with options for equal success or failure. We display the fraction of time for each method that it is classified by the majority of labelers as “Better” or “Both Achieve”. Value for DOODL averaged across two comparison runs.

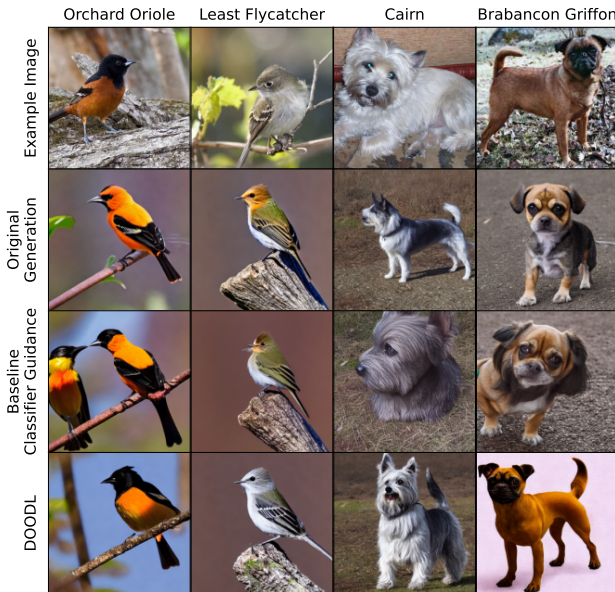


Figure 6: Qualitative generations for fine-grained classes with rare vocabulary. *Example Image* is an exemplar of the class from the CUB [45]/Dogs [21] dataset. Though the optimization process of DOODL starts from pixels not matching the target class (*Original Generation*) due to the rare vocabulary failure of StableDiffusion, it is able to produce an image of the targeted species using the gradients of a supervised classifier trained solely on the specialized dataset.

5.2. Vocabulary Expansion

Fine-grained visual classification (FGVC) models aim to classify datasets with subtle variations between classes, such as the Caltech-UCSD Birds (CUB) [45], Stanford Dogs (Dogs) [21], and FGVC-Aircraft (FGVC-A) [27] datasets. We seek to expand and refine the vocabulary of a pretrained DDM using DOODL such that it can generate

a specific class learnt by an FGVC classifier. For classifier guidance, we use binary cross-entropy (BCE) loss from supervised models trained on FGVC datasets. Given a model m trained on a dataset with classes $\{class_i\}_{i=1}^{n_c}$ we generate instances of $class_j$ via loss $c(x_0^{(i)}) = \lambda BCE(m(x_0^{(i)}), j)$, $(\lambda_{DOODL}, \lambda_{Baseline}) = (0.05, 5)$.

We evaluate DOODL’s ability to expand Stable Diffusion’s vocabulary using FGVC classifiers trained, respectively, on CUB, Dogs, and FGVC-A using WS-DAN [19]. The majority of concepts present in these datasets are rare or non-existent in the training data and as such, the pretrained DDM cannot generate accurate images for them (Figure 6-second row). We measure the FID [15] between a set of generated images (4 seeds) and the validation set of the FGVC dataset being studied. Despite overlap between ImageNet and the targeted datasets, FID has historically still been considered a useful generative evaluation metric [6]. In Table 3 we see that DOODL reduces FID compared to original Stable Diffusion generation on all datasets (average decrease of $\sim 3\%$), while classifier guidance does not improve the original generation FID on any. This indicates that DOODL is able to better incorporate the gradient signal from the classifiers as compared to the one-step approximation. Indeed, example images (Figure 6) show that DOODL can produce images with requisite fine-grained features that identify the object category in question (e.g., a bird species or a dog breed), such as colors, textures, and shapes. DOODL obtains the least relative improvement on Aircrafts, perhaps due to the class differences being largely subtle structural changes with very few textural cues that can act as guiding signals for the optimization process.

5.3. Aesthetic Improvement

Image generation research has emphasized the ability to generate aesthetically pleasing images. Notably, the publicly available Stable Diffusion model is trained on the LAION-5B dataset pre-filtered using an “Aesthetics Predictor” model. The “Aesthetics Predictor” model a is a linear head trained on top of CLIP visual embeddings to predict a scalar value in the range of [1,10], indicating perceived aesthetic quality (trained on [31]). This can easily be incorporated into DOODL’s framework with the cost function $c_{aes}(x_0^{(i)}) = \lambda \cdot |a(x_0^{(i)}) - A|$, $\lambda_{DOODL} = 1$ where $A = 10$ for aesthetic maximization and 1 for minimization.

While we find improvement signal for incorporating c_{aes} into text-to-image generations, we place the results in the Supplementary and instead focus on *editing* an existing image to *increase* its perceptual appeal, a novel task with a practical application for users who may wish to improve the aesthetic quality of their photographs. We employ EDICT to unconditionally invert an existing image to a latent pair (x_T, y_T) and then optimize the latents w.r.t. c_{aes} . Qualitative results of incorporating this cost function



Figure 7: We use EDICT to invert a real-world image to latents which are then optimized using DOODL, targeting an “aesthetic score” of either 1 (top row) or 10 (bottom).

are shown in Figure 7 where in-the-wild images are edited to have lower or higher aesthetic quality. We quantify DOODL’s editing ability using human evaluation in Figure 8, seeing that DOODL produces satisfactory edits significantly more often than the baseline classifier guidance. For the baselines in the latter, we find that inverting the images to $t = T$ results in generation content largely decorrelated from the original image, and invert to $t = \frac{T}{2}$ which we find to produce satisfactory edits.

5.4. Visual Personalization Guidance

Visual personalization, making a diffusion model generate images that contain a highly specific entity or concept based off of a limited exemplar set of images, is an area that has garnered much research attention. Most methods, such as Dreambooth [37] or Textual Inversion [12] learn or finetune components such as text embeddings or subsets of the diffusion model parameters in order to generate visually personalized images. While highly effective, the training cost and subsequent model storage are drawbacks. We propose a novel on-the-fly personalization paradigm using




Figure 8: Quantitative evaluation of aesthetic quality editing. 5 labelers are asked whether the edited image both retains content and improves perceptual appeal. Displayed are the success rates for DOODL and the baseline method (two guidance scale settings) for 3 increasingly stringent approval thresholds. Accounting for labeling noise, we plot the “success rate” of a $Binomial(5, 0.5)$ distribution. Evaluation is done on 96 random images from COCO [25].

DOODL to employ a pretrained recognition model in lieu of any additional generative model training or tuning.

We utilize the distance from Equation (10) taken between the image embedding of a conditioning image, $CLIP_{Image}(C_{Image})$, and that of the current generation $CLIP_{Image}(x_0^{(i)})$. We find ensembling ViT/B-32, ViT/L-14, and ViT/g-14 CLIP models with loss weights of 0.5, 0.25, 0.25 respectively achieves good performance. Input text conditioning to the diffusion model is as standard. A similar regime of λ_{DOODL} performs here as in Section 5.1, $\lambda_{DOODL} = 3$ for displayed results. This setting is unusually difficult for DOODL, requiring more (5 \times) optimization iterations than other reported experiments (additionally, we find a much higher $\lambda_{Baseline}$ needed).

We visualize generated samples in Figure 9, placing dogs pictured in the top row (C_{Image}) into various contexts. Prompts generically refer to “a dog”, with all identifying information coming from C_{Image} . DOODL maintains faithfulness to both the text prompt and visual conditioning; a first for non-learned methods to our knowledge.

We perform a human evaluation, sampling four target dog images from Imagenet [7] and using four prompt templates ($A\ dog\ \{at\ the\ Acropolis/swimming/in\ a\ bucket/wearing\ sunglasses\}$) across 6 random seeds, resulting in 96 generated images per method. For each image, three labelers are asked whether the dog in the generated image appears to be the same dog as the original with the context matching the caption. A majority of labelers classify the original generation and baseline as successful just 2% and 1% of the time respectively as opposed to 14% for DOODL, 7 \times more often than either baseline. Further analysis is given in the Supplementary.

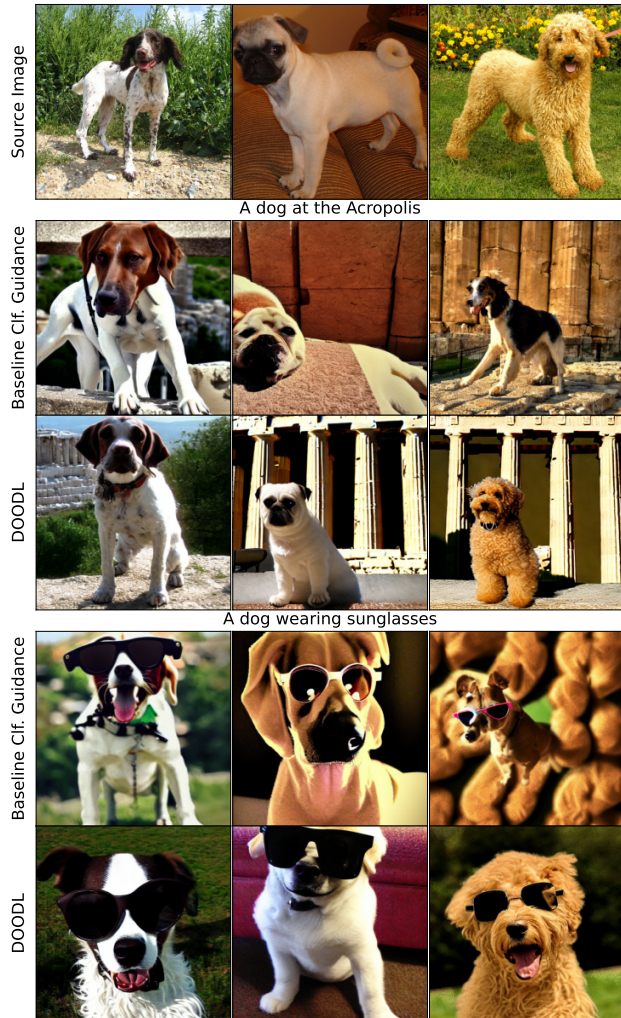


Figure 9: Personalization generations from a single seed. Baseline Clf. Guidance does not depict the pictured dog without destabilizing. DOODL produces realistic images containing both the desired text and image conditioning.

6. Conclusion & Future Work

In this work, we demonstrated that Direct Optimization Of Diffusion Latents (DOODL) offers an exciting new way to incorporate the knowledge of pretrained recognition networks in the generative process of diffusion models. We expect future work to both expand the types of guidances incorporated as well as sophisticate and accelerate the optimization process of DOODL to incorporate the process into applications that require higher compute efficiency.

References

[1] Lynton Ardizzone, Jakob Kruse, Sebastian J. Wirkert, Daniel Rahner, Eric W. Pellegrini, Ralf S. Klessen, Lena Maier-

Hein, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. *CoRR*, abs/1808.04730, 2018. 3

[2] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. 2

[3] Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models, 2023. 2

[4] Jens Behrmann, Will Grathwohl, Ricky TQ Chen, David Duvenaud, and Jörn-Henrik Jacobsen. Invertible residual networks. *International Conference on Machine Learning*, pages 573–582, 2019. 3

[5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 2

[6] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *Advances in neural information processing systems*, 29, 2016. 7

[7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 8

[8] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 2, 3, 4

[9] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014. 3

[10] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016. 3

[11] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis, 2020. 3

[12] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 3, 8

[13] Audrunas Gruslys, Remi Munos, Ivo Danihelka, Marc Lanctot, and Alex Graves. Memory-efficient backpropagation through time. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. 3

[14] Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens van der Maaten. Countering adversarial images using input transformations. In *International Conference on Learning Representations*, 2018. 4

[15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a

- two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 7
- [16] 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. 2
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2, 3
- [18] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 3
- [19] Tao Hu, Honggang Qi, Qingming Huang, and Yan Lu. See better before looking closer: Weakly supervised data augmentation network for fine-grained visual classification. *arXiv preprint arXiv:1901.09891*, 2019. 7
- [20] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hananeh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip, July 2021. If you use this software, please cite it as below. 6
- [21] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Fei-Fei Li. Novel dataset for fine-grained image categorization: Stanford dogs. In *Proc. CVPR workshop on fine-grained visual categorization (FGVC)*, volume 2. Cite-seer, 2011. 7
- [22] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018. 3
- [23] Gihyun Kwon and Jong Chul Ye. Diffusion-based image translation using disentangled style and content representation. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [24] Wei Li, Xue Xu, Xinyan Xiao, Jiachen Liu, Hu Yang, Guohao Li, Zhanpeng Wang, Zhifan Feng, Qiaoqiao She, Yajuan Lyu, et al. Upainting: Unified text-to-image diffusion generation with cross-modal guidance. *arXiv preprint arXiv:2210.16031*, 2022. 2, 4, 6
- [25] 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. 8
- [26] Matthew MacKay, Paul Vicol, Jimmy Ba, and Roger B Grosse. Reversible recurrent neural networks. *Advances in Neural Information Processing Systems*, 31, 2018. 3
- [27] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 7
- [28] 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, 4, 5, 6
- [29] Pascal Notin, José Miguel Hernández-Lobato, and Yarin Gal. Improving black-box optimization in vae latent space using decoder uncertainty, 2021. 3
- [30] Justin Pinkney. Image variation diffusion. <https://github.com/LambdaLabsML/Lambda-diffusers>, 2023. 2
- [31] John David Pressman, Katherine Crowson, and Simulacra Captions Contributors. Simulacra aesthetic captions. Technical Report Version 1.0, Stability AI, 2022. url <https://github.com/JD-P/simulacra-aesthetic-captions>. 7
- [32] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 3, 5, 6
- [33] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. ArXiv, May 2020. 3
- [34] 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. 2
- [35] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *CVPR*, pages 10684–10695, 2022. 2
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. Stablediffusion-2.0. <https://github.com/Stability-AI/stablediffusion>, 2023. 2
- [37] 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, 8
- [38] 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. 2, 6
- [39] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. *CoRR*, abs/1503.03585, 2015. 2
- [40] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3
- [41] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34:1415–1428, 2021. 2
- [42] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019. 3
- [43] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 2, 4

- [44] Jiawei Su, Danilo Vasconcellos Vargas, and Kouichi Sakurai. One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation*, 23(5):828–841, 2019. [4](#)
- [45] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. [7](#)
- [46] Bram Wallace, Akash Gokul, and Nikhil Naik. Edict: Exact diffusion inversion via coupled transformations. *arXiv preprint arXiv:2211.12446*, 2022. [2](#), [3](#), [4](#), [5](#)
- [47] Chaowei Xiao, Bo Li, Jun yan Zhu, Warren He, Mingyan Liu, and Dawn Song. Generating adversarial examples with adversarial networks. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 3905–3911. International Joint Conferences on Artificial Intelligence Organization, 7 2018. [4](#)
- [48] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. [3](#)
- [49] Geoffrey Zweig. Exact alpha-beta computation in logarithmic space with application to map word graph construction. In *Proceedings of ICSLP*, January 2000. [3](#)