



# Image embedding for denoising generative models

Andrea Asperti<sup>1</sup> · Davide Evangelista<sup>2</sup> · Samuele Marro<sup>1</sup> · Fabio Merizzi<sup>1</sup>

© The Author(s) 2023

## Abstract

Denoising Diffusion models are gaining increasing popularity in the field of generative modeling for several reasons, including the simple and stable training, the excellent generative quality, and the solid probabilistic foundation. In this article, we address the problem of *embedding* an image into the latent space of Denoising Diffusion Models, that is finding a suitable “noisy” image whose denoising results in the original image. We particularly focus on Denoising Diffusion Implicit Models due to the deterministic nature of their reverse diffusion process. As a side result of our investigation, we gain a deeper insight into the structure of the latent space of diffusion models, opening interesting perspectives on its exploration, the definition of semantic trajectories, and the manipulation/conditioning of encodings for editing purposes. A particularly interesting property highlighted by our research, which is also characteristic of this class of generative models, is the independence of the latent representation from the networks implementing the reverse diffusion process. In other words, a common seed passed to different networks (each trained on the same dataset), eventually results in identical images.

**Keywords** Denoising diffusion models · Generative models · Embedding · Latent space · Representation learning

---

✉ Andrea Asperti  
andrea.asperti@unibo.it

Davide Evangelista  
davide.evangelista5@unibo.it

Samuele Marro  
samuele.marro@unibo.it

Fabio Merizzi  
fabio.merizzi@unibo.it

<sup>1</sup> Department of Informatics: Science and Engineering (DISI), University of Bologna, Mura Anteo Zamboni 7, Bologna 40126, Italy

<sup>2</sup> Department of Mathematics, University of Bologna, Piazza di Porta San Donato 1, Bologna 40126, Italy


## 1 Introduction

Denosing Diffusion Models (DDM) (Ho et al. 2020) are rapidly imposing as the new state-of-the-art technology in the field of deep generative modeling, challenging the role held so far by Generative Adversarial Networks (Dhariwal and Nichol 2021). The impressive text-to-image generation capability shown by models like DALL·E 2 (Ramesh et al. 2022) and Imagen Saharia et al. (2022), recently extended to videos in Ho et al. (2022), clearly proved the qualities of this technique, comprising excellent image synthesis quality, good sampling diversity, high sensibility and easiness of conditioning, stability of training and good scalability.


In very rough terms, a diffusion model trains a single network to denoise images with a parametric amount of noise, and generates images by iteratively denoising pure random noise. This latter process is traditionally called *reverse diffusion* since it is meant to “invert” the *direct diffusion* process consisting in adding noise. In the important case of Implicit Diffusion models (Song et al. 2021), reverse diffusion is deterministic, but obviously not injective: many noisy images can be denoised to a single common result. Let us call  $emb(x)$  (embedding of  $x$ ) the set of points whose reverse diffusion generate  $x$ . The problems we are interested in are investigating the shape of  $emb(x)$  (e.g. is it a connected, convex space?), finding a “canonical” element in it (i.e. a sort of center of gravity) and, in case such a canonical element exists, finding an efficient way to compute it. This would allow us to embed an arbitrary image into the “latent” space of a diffusion model, providing functionality similar to GAN-recoders (see Sect. 2), or to encoders in the case of Variational AutoEncoders (Kingma and Welling 2019; Asperti et al. 2021).

Since reverse diffusion is the “inversion” of the diffusion process, it might be natural to expect  $emb(x)$  to be composed by noisy versions of  $x$ , and that the canonical element we are looking for could be  $x$  itself. This is not the case: indeed,  $x$  does not seem to belong to  $emb(x)$ . Figure 1 details some examples of the output obtained by using the image itself as input to the reverse diffusion process.

Since the input signal is clearly too strong, we may be tempted to reduce it using a multiplicative factor equal to the minimum signal rate used to train the denoising network (0.02 in our case), or a similarly low value. Examples of results are shown in Fig. 2. Although some macroscopic aspects of the original image like orientation and illumination are roughly preserved, most of the information is not embedded in these seeds: scaling does not result in a reasonable embedding. We also attempted to inject some additional noise into the initial seed, hoping to obtain a more entropic signal that



**Fig. 1** Examples of faces obtained using the image itself as input to the reverse diffusion process. The input signal is clearly far too strong



**Fig. 2** Examples of faces obtained using a weak version of the image itself as input signal. The first row shows the original image, while the second shows its weak version, which has been scaled by the minimum signal rate used to train the denoising network (0.02). This weaker image constitutes the initial seed. In the following four rows, we see the reconstructions obtained through reverse diffusion from the initial seed and progressively stronger versions of it, varying the signal rate between 0.02 and 0.08

is similar to the typical input of the reverse diffusion process, but this merely resulted in a less deterministic output.


Therefore, the embedding problem is both far from trivial and very interesting. Understanding the embedding would give us a better grasp of the reverse diffusion process, as well as a deeper, semantic insight into the structure of its latent space.

Our approaches to the embedding problem are discussed in Sect. 5. Overall, we find that we can obtain pretty good results by directly training a Neural Network to compute a kind of “canonical” seed (see Fig. 3).


The reconstruction quality is very high, with an MSE of around 0.0015 in the case of CelebA (Liu et al. 2015). More detailed values are provided in Sect. 5.2.

A typical application of the embedding process consists in transforming a signal into an element of the data manifold sufficiently close to it (the same principle behind denoising autoencoders). An amusing utilization is for the reification of artistic portraits, as exemplified in Fig. 4.

Another interesting possibility is that of making sketchy modifications to an image (a portrait, or a photograph) and delegating to the embedder-generator pair the burden of integrating them in the original picture in a satisfactory way (see Fig. 5).



**Fig. 3** CelebA: examples of face embeddings (second row) and reconstructions (third row). No cherry-picking



**Fig. 4** Reification of portraits. The portrait is first embedded into the latent space, and then pulled back into the data manifold



**Fig. 5** Scribbles over David's Napoleon

## 1.1 Structure of the work

The article is structured as follows. In Sect. 2, we discuss related works, mostly focusing on the embedding problem for Generative Adversarial Networks. Section 3 is devoted to formally introducing the notion of Denoising Diffusion Models, in addition to the deterministic variant of Denoising Diffusion Implicit Models we are particularly interested in. In the same section, we also discuss an intuitive interpretation of denoising diffusion models in terms of a “gravitational analogy” (Sect. 3.3), which drove many of our investigations and plays an important role in understanding the structure of datapoint embeddings. A major consequence of this interpretation, which to the best of our knowledge has never been pointed out before, is the *invariance of the latent space with respect to different models*: a given seed, passed as input to different models, always produces the same image. In Sect. 4, we provide architectural details about our implementation of the Denoising Diffusion model. Our methodology to address the embedding problem is discussed in Sect. 5. Two main approaches have been considered, one based on a gradient descent technique, which allows us to synthesize large clouds of different seeds in the embedding space of specific data points (Sect. 5.2), and another one based on training a neural network to compute a single “canonical” seed for the given image: essentially, a sort of encoder. Conclusions and future works are discussed in Sect. 6.

**Code.** The source code of the experiments described in this work is freely available at the GitHub repository <https://github.com/asperti/Embedding-in-Diffusion-Models>, along with links to weights for pre-trained models.

## 2 Related works

The embedding problem has been extensively investigated in the case of Generative Adversarial Networks (GANs) (Xia et al. 2022). Similarly to Denoising Generative Models, GANs lack a direct encoding process of the original input sample into the latent space.

Several approaches to inversion have been investigated (Perarnau et al. 2016; Bau et al. 2019; Daras et al. 2020; Anirudh et al. 2020), mostly with the purpose of editing. The most common approaches are based on synthesis of the latent encoding via gradient descent techniques (Creswell and Bharath 2019), or by training a suitable neural network to produce encodings able to reconstruct the original input with sufficient approximation. While the former technique generally tends to achieve better reconstruction errors, the latter has faster inference times and can take advantage of the fact that, since a GAN produces an infinite stream of training data, over-fitting is much less likely. Hybrid methods combining both techniques have also been explored (Zhu et al. 2016, 2020).

Recent works have mostly focused on the inversion of the popular StyleGAN and its successors (Karras et al. 2019, 2020, 2021), building on previous work with a variety of inversion structures and minimization objectives, or aiming to generalize/transfer to arbitrary datasets (Abdal et al. 2019; Collins et al. 2020; Abdal et al. 2020; Poirier-Ginter et al. 2022; Alaluf et al. 2022).

As we already mentioned, the typical application of the embedding is for exploration of the latent space, either for disentanglement purposes or in view of editing; the two issues are in fact tightly intertwined, since knowledge about semantically meaningful directions (e.g. color, pose, shape) can be exploited to tweak an image with the desired features. For instance,

InterFaceGAN (Shen et al. 2022) uses regression techniques to find a hyperplane in the latent space whose normal vector allows for a gradual modification of the feature. Further work based on this idea searches for these directions as an iterative or an optimization problem (Li et al. 2021) and also extends it to controllable walks in the latent space (Li et al. 2021). In the same vein, Kwon et al. (2022) studies the feature space of the U-Net bottleneck of the diffusion model, finding that it can be used as an alternative latent space with highly semantic directions.

Another important application of embeddings is for the comparison of the latent space of different generative models (Asperti and Tonelli 2022): having the possibility to embed the same image in different spaces allows us to create a supervised dataset suitable to learn direct mappings from one space to another.

In the realm of diffusion models, much work has been done on the refinement of the reverse diffusion process (Nichol and Dhariwal 2021; Choi et al. 2021; Dhariwal and Nichol 2021), but relatively little attention has been so far devoted to its inversion. DALL-E 2 Ramesh et al. (2022) relies on a form of image embedding, but this is a pre-trained contrastive model not learnt as the inversion of the generative task. An additional difference with respect to our work is that we are also interested in investigating and understanding the *structure* of the embedding cloud for each image since it could highlight the organization of the latent space and the sampling process.

Finally, in the context of text-conditioned generative models, interesting attempts to invert not just the image but a user-provided concept have been investigated in Gal et al. (2022). The concept is represented as a new pseudo-word in the model's vocabulary, which can be then used as part of a prompt (e.g. "a flower in the style of  $S_*$ ", where  $S_*$  refers to an image). The mapping is achieved by optimizing the conditioning vector in order to minimize the reconstruction error (similarly to the technique described in Sect. 5.1). A similar approach is used in Dong et al. (2022).

### 3 Denoising diffusion models

In this section, we provide a general overview of diffusion models from a mathematical perspective. All results in Sect. 3.1 and Sect. 3.2 are known in the literature; in Sect. 3.3 we propose an original interpretation of the reverse diffusion process in terms of a gravitational collapse of the latent space over the data manifold.

#### 3.1 Diffusion and reverse diffusion

Consider a distribution  $q(x_0)$  generating the data. Generative models aim to find a parameter vector  $\theta$  such that the distribution  $p_\theta(x_0)$ , parameterized by a neural network, approximates  $q(x_0)$  as accurately as possible. In Denoising Diffusion Probabilistic Models (DDPM) (Ho et al. 2020), the generative distribution  $p_\theta(x_0)$  is assumed to have the form

$$p_\theta(x_0) = \int p_\theta(x_{0:T}) dx_{1:T} \quad (1)$$

for a given time range horizon  $T > 0$ , where

$$p_\theta(x_{0:T}) = p_\theta(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t) \tag{2}$$

with  $p_\theta(x_T) = \mathcal{N}(x_T|0;I)$  and  $p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}|\mu_\theta(x_t, \alpha_t); \sigma_t^2 \cdot I)$ . Similarly, the diffusion model  $q(x_{0:T})$  is considered to be a Markov chain of the form

$$q(x_t|x_{t-1}) = \mathcal{N}\left(x_t \left| \sqrt{\frac{\alpha_t}{\alpha_{t-1}}}x_{t-1}; \left(1 - \frac{\alpha_t}{\alpha_{t-1}}\right) \cdot I \right.\right) \tag{3}$$

with  $\{\alpha_t\}_{t \in [0,T]}$  being a decreasing sequence in the interval  $[0, 1]$ . The parameters of the generative model  $p_\theta(x_0)$  are then trained to fit  $q(x_0)$  by minimizing the negative Evidence Lower BOund (ELBO) loss, defined as

$$\mathcal{L}(\theta) = -\mathbb{E}_{q(x_{0:T})}[\log p_\theta(x_{0:T}) - \log q(x_{1:T})]. \tag{4}$$

The ELBO loss can be rewritten in a computable form by noticing that, as a consequence of Bayes' Theorem,  $q(x_{t-1}|x_t, x_0) = \mathcal{N}(x_{t-1}|\tilde{\mu}(x_t, x_0); \sigma_q^2 \cdot I)$ . Consequently,

$$\mathcal{L}(\theta) = \sum_{t=1}^T \gamma_t \mathbb{E}_{q(x_t|x_0)} \left[ \|\mu_\theta(x_t, \alpha_t) - \tilde{\mu}(x_t, x_0)\|_2^2 \right] \tag{5}$$

which can be interpreted as the weighted mean squared error between the reconstructed image from  $p_\theta(x_t|x_0)$  and the true image obtained by the reverse diffusion process  $q(x_{t-1}|x_t, x_0)$  for each time  $t$ .

In Song et al. (2021), the authors considered a non-Markovian diffusion process

$$q_\sigma(x_{1:T}|x_0) = q_\sigma(x_T|x_0) \prod_{t=2}^T q_\sigma(x_{t-1}|x_t, x_0) \tag{6}$$

where  $q_\sigma(x_T|x_0) = \mathcal{N}(x_T|\sqrt{\alpha_T}x_0, (1 - \alpha_T) \cdot I)$ , and

$$q_\sigma(x_{t-1}|x_t, x_0) = \mathcal{N}\left(x_{t-1} \left| \mu_{\sigma_t}(x_0, \alpha_{t-1}); \sigma_t^2 \cdot I \right.\right) \tag{7}$$

with

$$\mu_{\sigma_t}(x_0, \alpha_{t-1}) = \sqrt{\alpha_{t-1}}x_0 + \sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \frac{x_t - \sqrt{\alpha_t}x_0}{\sqrt{1 - \alpha_t}}. \tag{8}$$

This construction implies that the forward process is no longer Markovian, since it depends both on the starting point  $x_0$  and on  $x_{t-1}$ . Moreover, Song et al. (2021) proved that, with this choice of  $q_\sigma(x_{1:T}|x_0)$ , the marginal distribution  $q_\sigma(x_t|x_0) = \mathcal{N}(x_t|\sqrt{\alpha_t}x_0; (1 - \alpha_t) \cdot I)$ , recovers the same marginals as in DDPM, which implies that  $x_t$  can be diffused from  $x_0$  and  $\alpha_t$  by generating a realization of normally distributed noise  $\epsilon_t \sim \mathcal{N}(\epsilon_t|0;I)$  and defining

$$x_t = \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon_t. \tag{9}$$

Note that when in Equation (7)  $\sigma_t = 0$ , the reverse diffusion  $q_\sigma(x_{t-1}|x_t, x_0)$  becomes deterministic. With such a choice of  $\sigma_t$ , the resulting model is named Denoising Diffusion Implicit Models (DDIM) by the authors in Song et al. (2021). Interestingly, in DDIM, the parameters of the generative model  $p_\theta(x_{t-1}|x_t)$  can be simply optimized by training a neural

network  $\epsilon_{\theta}^{(t)}(x_t, \alpha_t)$  to map a given  $x_t$  to an estimate of the noise  $\epsilon_t$  added to  $x_0$  to construct  $x_t$  as in (9). Consequently,  $p_{\theta}(x_{t-1}|x_t)$  becomes a  $\delta_{f_{\theta}^{(t)}}$ , where

$$f_{\theta}^{(t)}(x_t, \alpha_t) = \frac{x_t - \sqrt{1 - \alpha_t} \epsilon_{\theta}^{(t)}(x_t, \alpha_t)}{\sqrt{\alpha_t}}. \quad (10)$$

Intuitively, the network in (10) is just a denoiser that takes as input the noisy image  $x_t$  and the variance of the noise  $\alpha_t$  and returns an estimate of the denoised solution  $x_0$ . In DDIM, one can generate new data by first considering random Gaussian noise  $x_T \sim p_{\theta}(x_T)$  with  $\alpha_T = 1$ . Then,  $x_T$  is processed by  $f_{\theta}^{(T)}(x_T, \alpha_T)$  to generate an estimation of  $x_0$ , which is then corrupted again by the reverse diffusion  $q(x_{T-1}|x_T, f_{\theta}^{(T)}(x_T, \alpha_T))$ . This process is repeated until a new datum  $x_0$  is generated by  $f_{\theta}^{(1)}(x_1, \alpha_1)$ .

The sampling procedure of DDIM generates a trajectory  $\{x_T, x_{T-1}, \dots, x_0\}$  in the image space. In Song et al. (2020); Khruikov and Oseledets (2022) the authors found that the (stochastic) mapping from  $x_T$  to  $x_0$  in DDPM follows a score-based stochastic differential equation (SDE), where the dynamic is governed by terms related to the gradient of the ground-truth probability distribution from which the true data is generated. The sampling procedure for DDIM can be obtained by discretizing the deterministic *probability flow* (Song et al. 2020) associated with this dynamics. Consequently, training a DDIM model leads to an approximation of the score function of the ground-truth distribution.

### 3.2 The diffusion schedule

An important aspect in implementing diffusion models is the choice of the diffusion noise  $\{\alpha_t\}_{t=1}^T$ , defining the mean and the variance of  $q(x_t|x_0)$ . In Ho et al. (2020), the authors showed that the diffusion process  $q(x_t|x_0)$  converges to a normal distribution if and only if  $\alpha_T \approx 0$ . Moreover, to improve the generation quality,  $\alpha_t$  has to be chosen such that it slowly decays to 0.

The specific choice for the sequence  $\alpha_t$  defines the so-called *diffusion schedule*.

In Ho et al. (2020), the authors proposed to use linear or quadratic schedules. This choice was criticized in Kingma et al. (2021); Nichol and Dhariwal (2021) since it exhibits a too steep decrease during the first time steps, causing difficulties during generation for the neural network model. To remedy this situation, alternative scheduling functions with a gentler decrease have been proposed in the literature, such as the *cosine* or *continuous cosine* schedule. The behavior of all these functions is compared in Fig. 6.



The quantity of noise added by each schedule is also represented in Fig. 7, where a single image is injected with increasing noise according to the given schedule. It is not hard to see that the cosine and the continuous cosine schedules exhibit a more uniform transaction between the original image and the pure noise.

### 3.3 The gravitational analogy

Similarly to other generative models, developing an intuition of the actual behavior of diffusion models (and of the mapping from a latent encoding to its visible outcome) can be challenging. In this section, we propose a simple gravitational analogy that we found extremely useful to get an intuitive grasp of these models, and which suggested us some interesting conjectures about the actual shape of the embedding clouds for each object.



**Fig. 6** Comparison of different schedules. Generation is better if noise variance does not change too abruptly, so cosine and continuous cosine schedules usually work better than the linear or quadratic ones



**Fig. 7** Increasing noise added by the different scheduling: in order, from top to bottom, linear, quadratic, cosine, and continuous cosine schedules. For each row, from left to right, the time  $t$  is increased linearly from 0 to  $T$ . The corresponding  $\alpha_t$  for each schedule and for any  $t$  is reported above each image

Simply stated, the idea is the following. You should think of the datapoints as corps with a gravitational attraction. Regions of the space where the data manifold has high probability are equivalent to regions with high density. The denoising model essentially learns the gravitational map induced over the full space: any single point of the space gets mapped to the point where it would naturally “land” if subject to the “attraction” of the data manifold.

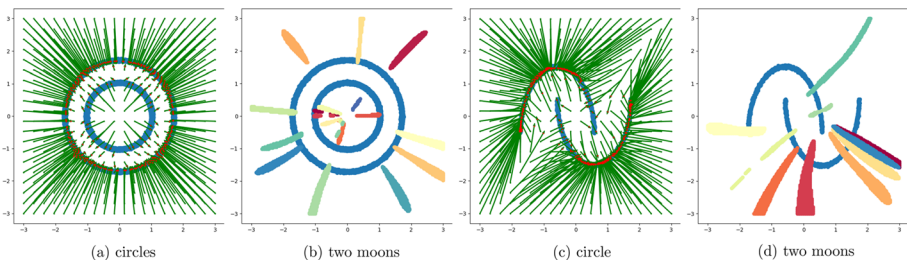

In more explicit terms, any point  $z$  of the space can be seen as a noisy version of any point  $x$  in the dataset. The “attraction” exerted by  $x$  on  $z$  (i.e. the loss) is directly proportional to their distance, usually an absolute or quadratic error.

However, the probability to train the network to reconstruct  $x$  from  $z$  has a Gaussian distribution  $\mathcal{N}(z|x;\sigma_z^2 \cdot I)$ , with  $\sigma_z^2$  depending on the denoising step. Hence, the *weighted attraction* exerted by  $x$  on  $z$  at each step is

$$\mathcal{N}(z|x;\sigma_z^2 \cdot I) \cdot \|x - z\|_1 \tag{11}$$

To get a grasp of the phenomenon, in Fig. 8 we compare the gravitational low for a corp  $x$  with the *weighted attraction* reported in Equation (11), under the assumption that the

**Fig. 8** Gravitational analogy. The orange line is the usual gravitational low for a corp with radius 1 and constant density. The blue line is a weighted loss  $\mathcal{N}(z|x;1) \cdot \|x - z\|_1$ . The two lines have been rescaled to have an equal integral




**Fig. 9** Gravitational map and embeddings for the “circles” (a, b) and “two moons” (c, d) datasets. Data-points are in blue. We consider a dense grid of seeds in the latent space, depicted in green. To visualize the maps (a) and (c) we draw an arrow pointing from each seed to the corresponding point generated by reverse diffusion (in red). To visualize embeddings (b) and (d) we consider a set of elements in the datasets, and for each element  $x$  we consider all points in the grid generating a sample  $\hat{x}$  sufficiently close to  $x$

variance  $\sigma$  has to be compared with the radius of the corp (with constant density, for simplicity).

According to the gravitational analogy, the embedding space  $emb(x)$  of each datapoint  $x$  should essentially coincide with the set of points in the space corresponding to trajectories ending in  $x$ . We can study this hypothesis on synthetic datasets. In Fig. 9 we show the gravitational map for the well-known “circle” (a) and “two moons” datasets (b); examples of embeddings are given in figures (c) and (d).

From the pictures, it is clear in which way the model “fills the space”, that is associating to each datapoint  $x$  all “trajectories” landing in  $x$ . The trajectories are almost straight and oriented along directions orthogonal to the data manifold. We believe that this behavior can be formally understood by exploiting the dynamics of the trajectories introduced in Song et al. (2020), as mentioned in Sect. 3. We aim to deeply investigate those aspects in a future work.

The most striking consequence of the “gravitational” interpretation is, however, the independence of the latent encoding from the neural network or its training: the gravitational map only depends on the data manifold and it is unique, so distinct networks or different trainings of the same network, if successful, should eventually end up with the



**Fig. 10** Uniqueness of the generative model. Different diffusion models generate essentially identical images when fed with the same seed. The two models in the picture are different versions of the U-Net: the first one has three downsampling blocks with channels [32, 64, 128], and the second one has four downsampling blocks with channels [48, 96, 192, 384]. The training sets are CIFAR10 (top), MNIST (middle), and Oxford Flowers 102 (bottom). Seeds have been randomly generated

same results. This seems miraculous: if we pick a random seed in an almost immense space, and pass it as input to two diffusion (deterministic) models for the same dataset, they should generate essentially *identical* images.

We experimentally verified and confirmed the previous property on a large number of variants of generative diffusion models and different datasets (see Fig. 10 for some results relative to CIFAR10, MNIST and Oxford Flowers). In particular, we tested different variants of the U-Net, with different numbers of downsampling blocks, different channel dimensions, and different layers in each block. We also optionally added different kinds of attention layers, in the traditional spatial form, or acting on channels like in squeeze-and-excitation layers (Hu et al. 2020) or in the recent NAFNet (Chen et al. 2022).

Provided generative models produce acceptable samples, the average quadratic distance between images generated by different generators on a same latent seed is always very small: typically, two-three orders of magnitude smaller than the average quadratic distance between random samples.

The fact that the same encoding works for different models seems to be peculiar to this kind of generative models. In Asperti and Tonelli (2022), it was observed that we can essentially pass from a latent space to another of different generative models with a simple *linear map*: however, an identity or even a permutation of latent variables does not usually suffice.<sup>1</sup>

While the uniqueness of the latent space is, in our opinion, a major discovery, it is not the main focus of this article, and we plan to conduct a more exhaustive and principled investigation of this property in future works.

## 4 Denoising architecture

The pseudocodes explaining training and sampling for diffusion models are respectively given in Algorithms 1 and 2 below.

---

### Algorithm 1 Training

---

```

1: repeat
2:    $x_0 \sim q(x_0)$ 
3:    $t \sim \text{Uniform}(1, \dots, T)$ 
4:    $\epsilon \sim \mathcal{N}(0; I)$ 
5:   Take gradient descent step on
      $\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}x_0 + \sqrt{1-\alpha_t}\epsilon, \alpha_t)\|^2$ 
6: until converged

```

---



---

### Algorithm 2 Sampling

---

```

1:  $x_T \sim \mathcal{N}(0, I)$ 
2: for  $t = T, \dots, 1$  do
3:    $\epsilon = \epsilon_\theta(x_t, \alpha_t)$ 
4:    $\tilde{x}_0 = \frac{1}{\sqrt{\alpha_t}}(x_t - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}}\epsilon)$ 
5:    $x_{t-1} = \sqrt{\alpha_{t-1}}\tilde{x}_0 + \sqrt{1-\alpha_{t-1}}\epsilon$ 
6: end for

```

---

As explained in Sect. 3, they are iterative algorithms; the only trainable component is the denoising network  $\epsilon_\theta(x_t, \alpha_t)$ , which takes as input a noisy image  $x_t$  and a noise variance  $\alpha_t$ , and tries to estimate the noise present in the image. This model is trained as a traditional denoising network, taking a sample  $x_0$  from the dataset, corrupting it

<sup>1</sup> It remains to be checked if imposing a spatial structure to the latent space of GANs and VAEs is enough to induce uniqueness in that case too. We plan to investigate this issue in a forthcoming work.

with the expected amount of random noise, and trying to identify the noise in the noisy image.

As a denoising network, it is quite standard to consider a conditional variant of the U-Net. This is a very popular network architecture originally proposed for semantic segmentation (Ronneberger et al. 2015) and subsequently applied to a variety of image manipulation tasks. In general, the network is structured with a downsample sequence of layers followed by an upsample sequence, with skip connections added between the layers of the same size.

To improve the sensibility of the network to the noise variance,  $\alpha$ , is taken as input, which is then embedded using an ad-hoc sinusoidal transformation by splitting the value in a set of frequencies, in a way similar to positional encodings in Transformers (Vaswani et al. 2017). The embedded noise variance is then vectorized and concatenated to the noisy images along the channel axes before being passed to the U-Net. This can be done for each convolution blocks separately, or just at the starting layer; we adopted the latter solution due to its simplicity and the fact that it does not seem to entail any loss in performance.

Having worked with a variety of datasets, we used slightly different implementations of the previously described model. The U-Net is usually parameterized by specifying the number of downsampling blocks, and the number of channels for each block; the upsampling structure is symmetric. The spatial dimension does not need to be specified, since it is inferred from the input. Therefore, the whole structure of a U-Net is essentially encoded in a single list such as [32, 64, 96, 128] jointly expressing the number of downsampling blocks (4, in this case), and the respective number of channels (usually increasing as we decrease the spatial dimension).

For our experiments, we have mainly worked with two basic architectures, mostly adopting [32, 64, 96, 128] for simple datasets such as MNIST of Fashion MNIST, and using more complex structures such as [48, 96, 192, 384] for CIFAR10 or CelebA. We also used different U-Net variants to extensively test the independence of the latent encoding discussed in Sect. 3.3.


## 5 Embedding

We experimented with several different approaches for the embedding task. The most effective ones have been the direct synthesis through gradient descent, and the training of ad-hoc neural networks. Both techniques have interesting aspects worth discussing.

The gradient descent technique is intrinsically non-deterministic, producing a variegated set of “noisy” versions of a given image  $x$ , all able to reconstruct  $x$  via reverse diffusion. The investigation of this set allows us to draw interesting conclusions on the shape of  $emb(x)$ .

Gradient descent is, however, pretty slow. A direct network can be trained to compute a single element inside  $emb(x)$ . Interestingly enough, this single element seems to be very close to the *average* of all noisy versions of  $x$  synthesized by the previous technique, suggesting evidence of its “canonical” nature.

The two techniques will be detailed in the following subsections.



**Fig. 11** Examples of seeds in the latent space. The image on the left is the original. On the right, we see 5 different seeds and their corresponding generations through the reverse diffusion process

## 5.1 Gradient descent synthesis

In Sect. 3.3, we computed the shape of embeddings for a few synthetic datasets by defining a dense grid of points in the latent space and looking for their final mapping through the reverse denoising process. Unfortunately, the number of points composing the grid grows exponentially in the number of features, and the technique does not scale to more complex datasets.

A viable alternative is the gradient descent approach, where we synthesize inputs starting from random noise, using the distance from a given target image as the objective function. In particular, given a sample  $x_0 \in \mathbb{R}^n$ , we propose solving the minimization problem


$$\min_{x_T \in \mathbb{R}^d} \frac{1}{2} \|f_\theta(x_T, \{\alpha_t\}_{t \in [0, T]}) - x_0\|_2^2 \quad (12)$$

where  $f_\theta(x_T, \{\alpha_t\}_{t \in [0, T]})$  models the sampling process described above with schedule  $\{\alpha_t\}_{t \in [0, T]}$ . Due to the non-convex nature of (12), the obtained solution strongly depends on the starting guess that initializes the optimization algorithm. Thus, by repeating the procedure above with different starting guesses  $x_T^0 \sim \mathcal{N}(0, I)$ , we were able to obtain multiple samples from  $emb(x_0)$ .


Generation usually requires several thousand steps, but it can be done in parallel on a batch of inputs. This allows us to compute, within a reasonable time, a sufficiently large number of samples in  $emb(x)$  for any given  $x$  (Fig. 11). Having a full cloud of data, we can use standard techniques like PCA to investigate its shape, as well as to study how the image changes when moving along the components of the cloud (see Sect. 5.1.1). For PCA investigations we need an embedding cloud with a dimension larger than the dimension of the latent space. We typically worked with clouds of 2k points for MNIST and Fashion MNIST and 4k points for CIFAR10.

A first interesting observation is that the average Euclidean distance among samples in  $emb(x)$  is typically very high, around 0.9: they are not concentrated in a small portion of the latent space. However, they seem to occupy a convex region. In Fig. 12 we show images obtained by reverse diffusion from 100 random *linear combinations* of seeds belonging to the embedding of the image on the left: all of them result in very similar reconstructions of the starting image.

Due to the convexity of the space, its mean is also comprised in it. In Fig. 13 we see the reconstructions obtained by considering as seed the average of a progressive number of seeds. The resulting images stabilize soon, although the result is slightly more blurry compared to using a single seed. The seeds on the borders of  $emb(x)$  seem to provide slightly



**Fig. 12** Linear combination of seeds. Given the original image (1 and 3) we compute by gradient descent a large cloud of seeds (4K) in its embedding. Then, we compute 100 *internal points*, as a 1-sum random linear combination of the given seeds. Images 2 and 4 contain the results of these linear combinations. All generated images are similar between each other and are very close to the original image. Therefore, all internal points seem to belong to the embedding




**Fig. 13** Progressive averaging in CIFAR10 and CelebA. The first row shows seeds computed as the mean of a progressive number of seeds in  $emb(x)$ , in a linear progression between 1 and 16; the second row shows their respective output through the reverse denoising process. The output is very similar. Additionally, observe that the original image becomes identifiable in the seeds, even averaging a relatively small number of a samples

better reconstructions than internal points (which makes the quest for a “canonical”, high-quality seed even more challenging).

### 5.1.1 PCA decomposition

Principal Component Analysis allows us to fit an ellipsoid over the cloud of datapoints, providing a major tool for investigating the actual shape of embeddings. According to the “gravitational” intuition exposed in Sect. 3.3,  $emb(x)$  should be elongated along directions orthogonal to the data manifold: moving along those directions should not sensibly



**Fig. 14** Fashion MNIST: movements along directions with minimal eigenvalues. The two groups of images refer to different components: starting from a mean seed generating the image in the middle, we move along a given component by the indicated positive or negative factor of the normalized eigenvector. Observe the progressive change in intensity and shape

influence generation, which should instead be highly affected by movements along minor components. Moreover, since the data manifold is likely oriented along a relatively small number of directions (due to the low dimensionality of the manifold), we expect that most PCA components in each cloud will be orthogonal to the manifold, and have relatively high eigenvalues.

For instance, in the case of the clouds of seeds for CIFAR10, eigenvalues along all 3072 components typically span between 0.0001 and 4. We observe significant modifications only moving along the minor components of the clouds: in fact, they provide the shortest way to leave the embedding space of a given point. However, as soon as we leave the embedding space of  $x$  we should enter the embedding space of some “adjacent” point  $x'$ . In other words, the minor components should define directions *inside* the data manifold, and possibly have a “semantical” (likely entangled) interpretation (Fig. 14).

## 5.2 Embedding networks

The second approach consists in training a neural network to directly compute a sort of “canonical” embedding for each image of the data manifold. The network takes as input an image  $x$  and produces a seed  $z_x \in emb(x)$ ; the loss function used to train the network is simply the distance between  $x$  and the result  $\hat{x}$  of the denoising process starting from  $z_x$ .

We tested several different networks; metrics relative to the most significant architectures are reported in Table 1.

A visual comparison of the behavior of the different networks is given in Fig. 15, relative to CIFAR10. More examples on CelebA are given below.


We started our investigation with a very simple network: a single convolution with a  $5 \times 5$  kernel. The reason for this choice is that, according to the discussion we made in the introduction and the visualization of the mean element of the embedding clouds of Fig. 13, we expected the latent encoding to be similar to a suitably rescaled version of the source image. The results on a simple dataset like MNIST confirmed this hypothesis,



**Table 1** Comparing the Mean Square Error (MSE) through the embedding-reconstruction process using different embedding networks; the MSE standard deviation is below the last reported decimal


Network	Params	MSE				
		MNIST	Fashion MNIST	CIFAR10	Oxford Flowers	CeleBA
layers: 1 conv. 5×5	78	.00704	.0152	.0303	.0372	.0189
layers: 3 conv. 5×5 channels: 16-16-out	7233	.00271	.00523	.0090	.0194	.0101
layers: 3 conv. 5×5 channels: 64-64-out	105,729	.00206	.00454	.0061	.0153	.00829
layers: 2conv.5×5 3conv.3×3 channels	859,009	.00121	.00172	.0038	.00882	.00396
128-128-128-128-out U-Net	9,577,683	.000361	.000890	.0012	.00248	.00147

The number of parameters refers to the instance of the network for the CeleBA dataset




**Fig. 15** Visual comparison with different Embedding Networks. We consider a set of test images from CIFAR10 (first row) and compute the embedding with one of the Embedding Networks of Table 1. We then use the embeddings to generate the corresponding images (remaining rows)

but on more complex ones like CIFAR10 it does not seem to be the case, as exemplified in Fig. 15. We then progressively improved the model's architecture by augmenting their depth and channel dimensions, with the latter being typically the most effective way to improve their performance. In the end, the best results were obtained with a U-Net architecture that is practically identical to the denoising network. Many additional experiments have been performed, comprising autoencoders, residual networks,



**Fig. 16** Embedding examples for the CelebA dataset. The first row contains the original examples, the second the synthesized latent seed, and the third the reconstructed image. Reconstruction is very good, with just a slight blurriness



**Fig. 17** Embedding examples on generated images. In this case, we start from images created by the generator (first row) and re-embed them inside the latent space (second row) using the Embedding Network. In the third row, we show the reconstruction, which is almost perfect. This could be either explained by a deficiency of the generator, or just by the fact that generated images are “simpler”, and hence can be more easily embedded than real ones

inception modules, and variants with different padding modalities or regularizations. However, they did not prove to be particularly effective and were thus dropped from our discussion.

In Fig. 16, we show some examples of embeddings and relative reconstructions in the case of the CelebA dataset.

The quality of the reconstruction is definitely high, with just a slight blurriness. There are two possible justifications for the tiny inaccuracy of this result: it could either be a fault of the generator, which is unable to create the requested images (as it is frequently the case with Generative Adversarial Networks (Asperti and Tonelli 2022)), or it could be a fault of the Embedding Network, which is unable to compute the correct seed.

To better investigate the issue, we performed two experiments. First, we restricted the reconstruction to images produced by the generator: in this case, if the Embedding network works well, it should be able to reconstruct almost perfect images. Secondly, we tried to improve the seeds computed by the Embedding Network through gradient descent, looking for better candidates.


We report the result of the first experiment in Fig. 17.

While the reconstruction is qualitatively accurate, we can also confirm the effectiveness in a more analytical way. In Table 2 we compare the mean squared error of the

**Table 2** Reconstruction error

Source Images	MSE
Dataset	0.00147
Generated	0.00074

In the first case, images are taken from the CelebA dataset: in the second case, they have been generated through the reverse diffusion process. The mean squared error (MSE) was computed over 1000 examples. Both experiments achieve a small reconstruction error, although the second one is even smaller



**Fig. 18** Gradient descent fine-tuning. The seeds obtained through the embedding network (second row) are refined through gradient descent (fourth row). The respective resulting reconstructions are depicted in rows 3 and 5. The improvement is almost imperceptible


reconstruction starting from original CelebA images versus generated data: the latter is sensibly smaller.

The fact that embedding works better for generated images is, however, not conclusive: it could either be explained by a deficiency of the generator, unable to generate all images in the CelebA dataset, or just by the fact that generated images are “simpler” than real ones (observe the well-known patinated look, which is typical of most generative models) and hence more easily embeddable.

Even the results of the second experiment are not easily deciphered. From a visual point of view, refining the embedding through gradient descent is not producing remarkable results, as exemplified in Fig. 18.

However, numerically, we see an improvement from an MSE of 0.00147 to an MSE of 0.00058, which seems to suggest some margin of improvement for the embedding network.

In conclusion, both the generator and the embedder can likely still be improved. However, a really interesting research direction seems to be the possibility to modify the latent representation to improve the realism of the resulting image, even if possibly not in the direction of the original. Therefore, a basic embedder, even if not fully accurate, could still provide the starting point for very interesting manipulations.



**Fig. 19** Interpolation between samples of the CelebA dataset. In the first row, we have the linear interpolation between the source and the target. In the second row, the linear interpolation between the latent embedding of the source, and the latent embedding of the target. In the third row, the reconstructed images

### 5.3 Latent space interpolation

A typical application of the embedding network is for the investigation of semantical properties of the latent space, starting from real samples and their attributes. As a preliminary step in this direction, in this section we provide examples of latent-space interpolations: the crucial additional ability added by the embedder is in the choice of the starting and ending point, that can be the embeddings of *real* data samples: this allows us to produce smooth interpolations between any pair of images in the dataset.

In Fig. 19 we show an example relative to the CelebA dataset. The linear interpolation in the visible space between a source and a target sample, depicted in the first row, does not produce satisfactory results: the superposition of the two images is clearly visible, introducing annoying artifacts. A better result can be achieved by first embedding both source and target into the latent space, and then moving along their (linear) interpolation (second row). The images generated from the interpolated latent points provide a smooth transition from the source to the target, as shown in the third row of Fig. 19. In this case, the “artifacts” of the latent representations are automatically corrected by the generator, trained to produce realistic faces.

## 6 Conclusions

In this article we addressed the problem of embedding data into the latent space of Deterministic Diffusion models, providing functionality similar to the encoder in a Variational Autoencoder, or the so-called *recoder* for Generative Adversarial Networks. The main source of complexity when inverting a diffusion model is the non-injective nature of the generator: for each sample  $x$ , there exists a cloud of elements  $z$  able to generate  $x$ . We call this set the embedding of  $x$ , denoted as  $emb(x)$ . We performed a deep investigation of the typical shape of  $emb(x)$ , which suggests that embeddings are usually orthogonal to the dataset. These studies point to a sort of gravitational interpretation of the reverse diffusion process, according to which the space is progressively collapsing over the data manifold. In this perspective,  $emb(x)$  is just the set of all trajectories in the space ending in  $x$ . We tested

our interpretation on both low- and high-dimensional datasets, highlighting a quite amazing result: the latent space of a DDIM generator does not significantly depend on the specific generative model, but just on the data manifold. In other words, passing the same seed as input to different DDIMs will result in almost identical outputs. In order to compute embeddings, we considered both gradient descent approaches, as well as the definition and training of specific Embedding Networks. We showed that, among all the architectures we tested, a U-Net obtained the best results, achieving a high-quality reconstruction from both a quantitative and qualitative point of view.

Embedding networks have a lot of interesting applications, largely exemplified in the introduction. More generally, the simplicity and ease of use of Embedding Networks open a wide range of fascinating perspectives about the exploration of semantic trajectories in the latent space, the disentanglement of the different aspects of variations, and the possibility of data editing. We thus hope that our results, by expanding the current understanding of generative models, can guide future research efforts.

**Funding** Open access funding provided by Alma Mater Studiorum - Università di Bologna within the CRUI-CARE Agreement. Research partially supported by the Future AI Research (FAIR) project of the National Recovery and Resilience Plan (NRRP), Mission 4 Component 2 Investment 1.3 funded from the European Union - NextGenerationEU.

## Declarations

**Conflict of interest** On behalf of all authors, the corresponding author states that there is no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Abdal R, Qin Y, Wonka P (2019) Image2stylegan: How to embed images into the stylegan latent space? In: 2019 IEEE/CVF international conference on computer vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019, IEEE, pp 4431–4440. <https://doi.org/10.1109/ICCV.2019.00453>
- Abdal R, Qin Y, Wonka P (2020) Image2stylegan++: How to edit the embedded images? In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 8296–8305
- Alaluf Y, Tov O, Mokady R, Gal R, Bermano A (2022) Hyperstyle: Stylegan inversion with hypernetworks for real image editing. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 18511–18521
- Anirudh R, Thiagarajan JJ, Kailkhura B, Bremer PT (2020) Mimicgan: Robust projection onto image manifolds with corruption mimicking. *Int J Comput Vis* 128(10):2459–2477
- Asperti A, Evangelista D, Piccolomini EL (2021) A survey on variational autoencoders from a green AI perspective. *SN Comput Sci* 2(4):301. <https://doi.org/10.1007/s42979-021-00702-9>
- Asperti A, Tonelli V (2022) Comparing the latent space of generative models. *Neural Comput Appl*. <https://doi.org/10.1007/s00521-022-07890-2>
- Bau D, Strobel H, Peebles WS, Wulff J, Zhou B, Zhu J, Torralba A (2019) Semantic photo manipulation with a generative image prior. *ACM Trans Graph* 38(4):59–15911

- Chen L, Chu X, Zhang X, Sun J (2022) Simple baselines for image restoration. In: Computer Vision—ECCV 2022: 17th European conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII, Springer, pp 17–33
- Choi J, Kim S, Jeong Y, Gwon Y, Yoon S (2021) ILVR: conditioning method for denoising diffusion probabilistic models. In: 2021 IEEE/CVF International conference on computer vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021, pp 14347–14356. <https://doi.org/10.1109/ICCV48922.2021.01410>
- Collins E, Bala R, Price B, Susstrunk S (2020) Editing in style: Uncovering the local semantics of gans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp 5771–5780
- Creswell A, Bharath AA (2019) Inverting the generator of a generative adversarial network. *IEEE Trans Neural Networks Learn Syst* 30(7):1967–1974
- Daras G, Odena A, Zhang H, Dimakis AG (2020) Your local gan: designing two dimensional local attention mechanisms for generative models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 14531–14539
- Dhariwal P, Nichol AQ (2021) Diffusion models beat gans on image synthesis. In: Ranzato M, Beygelzimer A, Dauphin YN, Liang P, Vaughan JW (eds) Advances in neural information processing systems 34: annual conference on neural information processing systems 2021, NeurIPS 2021, December 6–14, 2021, Virtual, pp 8780–8794 . <https://proceedings.neurips.cc/paper/2021/hash/49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html>
- Dong Z, Wei P, Lin L (2022) Dreamartist: Towards controllable one-shot text-to-image generation via contrastive prompt-tuning. [arxiv:2211.11337](https://arxiv.org/abs/2211.11337)
- Gal R, Alaluf Y, Atzmon Y, Patashnik O, Bermano AH, Chechik G, Cohen-Or D (2022) An image is worth one word: Personalizing text-to-image generation using textual inversion. [arxiv:2208.01618](https://arxiv.org/abs/2208.01618)
- Ho J, Jain A, Abbeel P (2020) Denoising diffusion probabilistic models. In: Larochelle H, Ranzato M, Hadsell R, Balcan M, Lin H (eds) Advances in neural information processing systems 33: annual conference on neural information processing systems 2020, NeurIPS 2020, December 6–12, 2020, Virtual. <https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html>
- Ho J, Salimans T, Gritsenko A, Chan W, Norouzi M, Fleet DJ (2022) Video diffusion models. [arXiv:2204.03458](https://arxiv.org/abs/2204.03458)
- Hu J, Shen L, Albanie S, Sun G, Wu E (2020) Squeeze-and-excitation networks. *IEEE Trans Pattern Anal Mach Intell* 42(8):2011–2023. <https://doi.org/10.1109/TPAMI.2019.2913372>
- Karras T, Aittala M, Laine S, Härkönen E, Hellsten J, Lehtinen J, Aila T (2021) Alias-free generative adversarial networks. *Adv Neural Inf Process Syst* 34:852–863
- Karras T, Laine S, Aila T (2019) A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 4401–4410
- Karras T, Laine S, Aittala M, Hellsten J, Lehtinen J, Aila T (2020) Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 8110–8119
- Khrulkov V, Oseledets I (2022) Understanding ddpm latent codes through optimal transport. [arxiv:2202.07477](https://arxiv.org/abs/2202.07477)
- Kingma D, Salimans T, Poole B, Ho J (2021) Variational diffusion models. *Adv Neural Inf Process Syst* 34:21696–21707
- Kingma DP, Welling M (2019) An introduction to variational autoencoders. *Found Trends Mach Learn* 12(4):307–392. <https://doi.org/10.1561/22000000056>
- Kwon M, Jeong J, Uh Y (2022) Diffusion models already have a semantic latent space. *CoRR* [arxiv:2210.10960](https://arxiv.org/abs/2210.10960), <https://doi.org/10.48550/arXiv.2210.10960>
- Li Z, Tao R, Wang J, Li F, Niu H, Yue M, Li B (2021) Interpreting the latent space of gans via measuring decoupling. *IEEE Trans Artif Intell* 2(1):58–70
- Li G, Liu Y, Wei X, Zhang Y, Wu S, Xu Y, Wong HS (2021) Discovering density-preserving latent space walks in gans for semantic image transformations. In: Proceedings of the 29th ACM international conference on multimedia, pp 1562–1570
- Liu Z, Luo P, Wang X, Tang X (2015) Deep learning face attributes in the wild. In: Proceedings of International Conference on Computer Vision (ICCV), pp 3730–3738
- Nichol AQ, Dhariwal P (2021) Improved denoising diffusion probabilistic models. In: International conference on machine learning, PMLR, pp 8162–8171
- Perarnau G, Van De Weijer J, Raducanu B, Álvarez JM (2016) Invertible conditional gans for image editing. *arXiv preprint* [arXiv:1611.06355](https://arxiv.org/abs/1611.06355)

- Poirier-Ginter Y, Lessard A, Smith R, Lalonde JF (2022) Overparameterization improves stylegan inversion. [arxiv:2205.06304](https://arxiv.org/abs/2205.06304)
- Ramesh A, Dhariwal P, Nichol A, Chu C, Chen M (2022) Hierarchical text-conditional image generation with CLIP latents. [arXiv:2204.06125](https://arxiv.org/abs/2204.06125)
- Ronneberger O, Fischer P, Brox T (2015) U-net: Convolutional networks for biomedical image segmentation. In: International conference on medical image computing and computer-assisted intervention, Springer pp 234–241
- Saharia C, Chan W, Saxena S, Li L, Whang J, Denton E, Ghasemipour SKS, Ayan BK, Mahdavi SS, Lopes RG, Salimans T, Ho J, Fleet DJ, Norouzi M (2022) Photorealistic text-to-image diffusion models with deep language understanding. [CoRR arxiv: 2205.11487](https://arxiv.org/abs/2205.11487), <https://doi.org/10.48550/arXiv.2205.11487>
- Shen Y, Yang C, Tang X, Zhou B (2022) Interfacegan: interpreting the disentangled face representation learned by gans. *IEEE Trans Pattern Anal Mach Intell* 44(4):2004–2018. <https://doi.org/10.1109/TPAMI.2020.3034267>
- Song Y, Sohl-Dickstein J, Kingma DP, Kumar A, Ermon S, Poole B (2020) Score-based generative modeling through stochastic differential equations. [arxiv:2011.13456](https://arxiv.org/abs/2011.13456)
- Song J, Meng C, Ermon S (2021) Denoising diffusion implicit models. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. [OpenReview.net, ??? . https://openreview.net/forum?id=St1giarCHLP](https://openreview.net/forum?id=St1giarCHLP)
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L, Polosukhin I (2017) Attention is all you need. In: advances in neural information processing systems 30: Annual conference on neural information processing systems 2017, December 4-9, 2017, Long Beach, CA, USA, pp 5998–6008. <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>
- Xia W, Zhang Y, Yang Y, Xue JH, Zhou B, Yang MH (2022) Gan inversion: a survey. In: *IEEE transactions on pattern analysis and machine intelligence*
- Zhu J, Krähenbühl P, Shechtman E, Efros AA (2016) Generative visual manipulation on the natural image manifold. In: *Computer Vision - ECCV 2016 - 14th European conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V. Lecture notes in computer science*, vol 9909. Springer, pp 597–613. [https://doi.org/10.1007/978-3-319-46454-1\\_36](https://doi.org/10.1007/978-3-319-46454-1_36)
- Zhu J, Shen Y, Zhao D, Zhou B (2020) In-domain GAN inversion for real image editing. In: *Computer Vision - ECCV 2020 - 16th European conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XVII. Lecture notes in computer science*, vol 12362. Springer, pp 592–608. [https://doi.org/10.1007/978-3-030-58520-4\\_35](https://doi.org/10.1007/978-3-030-58520-4_35)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.