

Diffusion Models in Vision: A Survey

Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, *Member, IEEE*, and Mubarak Shah, *Fellow, IEEE*

Abstract—Denoising diffusion models represent a recent emerging topic in computer vision, demonstrating remarkable results in the area of generative modeling. A diffusion model is a deep generative model that is based on two stages, a forward diffusion stage and a reverse diffusion stage. In the forward diffusion stage, the input data is gradually perturbed over several steps by adding Gaussian noise. In the reverse stage, a model is tasked at recovering the original input data by learning to gradually reverse the diffusion process, step by step. Diffusion models are widely appreciated for the quality and diversity of the generated samples, despite their known computational burdens, *i.e.* low speeds due to the high number of steps involved during sampling. In this survey, we provide a comprehensive review of articles on denoising diffusion models applied in vision, comprising both theoretical and practical contributions in the field. First, we identify and present three generic diffusion modeling frameworks, which are based on denoising diffusion probabilistic models, noise conditioned score networks, and stochastic differential equations. We further discuss the relations between diffusion models and other deep generative models, including variational auto-encoders, generative adversarial networks, energy-based models, autoregressive models and normalizing flows. Then, we introduce a multi-perspective categorization of diffusion models applied in computer vision. Finally, we illustrate the current limitations of diffusion models and envision some interesting directions for future research.

Index Terms—diffusion models, denoising diffusion models, noise conditioned score networks, score-based models, image generation, deep generative modeling.

1 INTRODUCTION

DIFFUSION models [1]–[11] form a category of deep generative models which has recently become one of the hottest topics in computer vision (see Figure 1), showcasing impressive generative capabilities, ranging from the high level of details to the diversity of the generated examples. We can even go as far as stating that these generative models raised the bar to a new level in the area of generative modeling, particularly referring to models such as Imagen [12] and Latent Diffusion Models (LDMs) [10]. This statement is confirmed by the image samples illustrated in Figure 2, which are generated by Stable Diffusion, a version of LDMs [10] that generates images based on text prompts. The generated images exhibit very few artifacts and are very well aligned with the text prompts. Notably, the prompts are purposely chosen to represent unrealistic scenarios (never seen at training time), thus demonstrating the high generalization capacity of diffusion models.

To date, diffusion models have been applied to a wide variety of generative modeling tasks, such as image generation [1]–[7], [10], [11], [13]–[23], image super-resolution [10], [18], [24]–[27], image inpainting [1], [3], [4], [10], [24], [26], [28]–[30], image editing [31]–[33], image-to-image translation [32], [34]–[38], among others. Moreover, the latent representation learned by diffusion models was also found to be useful in discriminative tasks, *e.g.* image segmentation




Fig. 1. The rough number of papers on diffusion models per year.

[39]–[42], classification [43] and anomaly detection [44]–[46]. This confirms the broad applicability of denoising diffusion models, indicating that further applications are yet to be discovered. Additionally, the ability to learn strong latent representations creates a connection to representation learning [47], [48], a comprehensive domain that studies ways to learn powerful data representations, covering multiple approaches ranging from the design of novel neural architectures [49]–[52] to the development of learning strategies [53]–[58].

According to the graph shown in Figure 1, the number of papers on diffusion models is growing at a very fast pace. To outline the past and current achievements of this rapidly developing topic, we present a comprehensive review of articles on denoising diffusion models in computer vision. More precisely, we survey articles that fall in the category of generative models defined below. *Diffusion models* represent a category of deep generative models that are based on (i) a forward diffusion stage, in which the input data is gradually perturbed over several steps by adding Gaussian noise, and (ii) a reverse (backward) diffusion stage, in which a

- F.A. Croitoru, V. Hondru and R.T. Ionescu are with the Department of Computer Science, University of Bucharest, Bucharest, Romania. F.A. Croitoru and V. Hondru have contributed equally. R.T. Ionescu is the corresponding author.
E-mail: raducu.ionescu@gmail.com
- M. Shah is with the Center for Research in Computer Vision (CRCV), Department of Computer Science, University of Central Florida, Orlando, FL, 32816.

Manuscript received April 19, 2022; revised August 26, 2022.




Fig. 2. Images generated by Stable Diffusion [10] based on various text prompts, via the <https://beta.dreamstudio.ai/dream> platform.

generative model is tasked at recovering the original input data from the diffused (noisy) data by learning to gradually reverse the diffusion process, step by step.

We underline that there are at least three sub-categories of diffusion models that comply with the above definition. The first sub-category comprises denoising diffusion probabilistic models (DDPMs) [1], [2], which are inspired by the non-equilibrium thermodynamics theory. DDPMs are latent variable models that employ latent variables to estimate the probability distribution. From this point of view, DDPMs can be viewed as a special kind of variational auto-encoders (VAEs) [50], where the forward diffusion stage corresponds to the encoding process inside VAE, while the reverse diffusion stage corresponds to the decoding process. The second sub-category is represented by noise conditioned score networks (NCSNs) [3], which are based on training a shared neural network via score matching to estimate the score function (defined as the gradient of the log density) of the perturbed data distribution at different noise levels. Stochastic differential equations (SDEs) [4] represent an alternative way to model diffusion, forming the third sub-category of diffusion models. Modeling diffusion via forward and reverse SDEs leads to efficient generation strategies as well as strong theoretical results [59]. This latter formulation (based on SDEs) can be viewed as a generalization over DDPMs and NCSNs.

We identify several defining design choices and synthesize them into three generic diffusion modeling frameworks corresponding to the three sub-categories introduced above. To put the generic diffusion modeling framework into context, we further discuss the relations between diffusion models and other deep generative models. More specifi-

cally, we describe the relations to variational auto-encoders (VAEs) [50], generative adversarial networks (GANs) [52], energy-based models (EBMs) [60], [61], autoregressive models [62] and normalizing flows [63], [64]. Then, we introduce a multi-perspective categorization of diffusion models applied in computer vision, classifying the existing models based on several criteria, such as the underlying framework, the target task, or the denoising condition. Finally, we illustrate the current limitations of diffusion models and envision some interesting directions for future research. For example, perhaps one of the most problematic limitations is the poor time efficiency during inference, which is caused by a very high number of evaluation steps, *e.g.* thousands, to generate a sample [2]. Naturally, overcoming this limitation without compromising the quality of the generated samples represents an important direction for future research.

In summary, our contribution is twofold:

- Since many contributions based on diffusion models have recently emerged in vision, we provide a comprehensive and timely literature review of denoising diffusion models applied in computer vision, aiming to provide a fast understanding of the generic diffusion modeling framework to our readers.
- We devise a multi-perspective categorization of diffusion models, aiming to help other researchers working on diffusion models applied to a specific domain in quickly finding relevant related works in the respective domain.

2 GENERIC FRAMEWORK

Diffusion models are a class of probabilistic generative models that learn to reverse a process that gradually degrades




Fig. 3. A generic framework composing three alternative formulations of diffusion models based on: stochastic differential equations (SDEs), denoising diffusion probabilistic models (DDPMs) and noise conditioned score networks (NCSNs). In general, a diffusion model consists of two processes. The first one, called the forward process, transforms data into noise, while the second one is a generative process that reverses the effect of the forward process. This latter process learns to transform the noise back into data. We illustrate these processes for all three formulations. The forward SDE shows that a change over time in x is modeled by a function f plus a stochastic component $\partial \omega \sim \mathcal{N}(0, \partial t)$ scaled by $\sigma(t)$. We underline that different choices of f and σ will lead to different diffusion processes. This is why the SDE formulation is a generalization of the other two. The reverse (generative) SDE shows how to change x in order to recover the data from pure noise. We keep the random component and modify the deterministic one using the gradients of the log probability $\nabla_x \log p_t(x)$, so that x moves to regions where the data density $p(x)$ is high. DDPMs sample the data points during the forward process from a normal distribution $\mathcal{N}(x_t; \sqrt{1 - \beta_t} \cdot x_{t-1}, \beta_t \cdot \mathbf{I})$, where $\beta_t \ll 1$. This iterative sampling slowly destroys information in data, and replaces it with Gaussian noise. The sampling is illustrated via the reparametrization trick (see details in Section 2.1). The reverse process of DDPM also performs iterative sampling from a normal distribution, but the mean $\mu_\theta(x_t, t)$ of the distribution is derived by subtracting the noise, estimated by a neural network, from the image at the previous step x_t . The variance is equal to the one used in the forward process. The initial image going into the reverse process contains only Gaussian noise. The forward process of NCSN simply adds normal noise to the image at the previous step. This can also be seen as sampling from a normal distribution $\mathcal{N}(x_t; x_{t-1}, (\sigma_t^2 - \sigma_{t-1}^2) \cdot \mathbf{I})$, with the mean being the image at the previous step. The reverse process of NCSN is based on an algorithm described in Section 2.2. Best viewed in color.

the training data structure. Thus, the training procedure involves two phases: the forward diffusion process and the backward denoising process.

The former phase consists of multiple steps in which low-level noise is added to each input image, where the scale of the noise varies at each step. The training data is progressively destroyed until it results in pure Gaussian noise.

The latter phase is represented by reversing the forward diffusion process. The same iterative procedure is employed, but backwards: the noise is sequentially removed, and hence, the original image is recreated. Therefore, at inference time, images are generated by gradually reconstructing them starting from random white noise. The noise subtracted at each time step is estimated via a neural network, typically based on a U-Net architecture [65], allowing the preservation of dimensions.

In the following three subsections, we present three formulations of diffusion models, namely denoising diffusion probabilistic models, noise conditioned score networks, and the approach based on stochastic differential equations that generalizes over the first two methods. For each formulation, we describe the process of adding noise to the data, the method which learns to reverse this process, and how new samples are generated at inference time. In Figure 3, all

three formulations are illustrated as a generic framework. We dedicate the last subsection to discussing connections to other deep generative models.

2.1 Denoising Diffusion Probabilistic Models (DDPMs)

Forward process. DDPMs [1], [2] slowly corrupt the training data using Gaussian noise. Let $p(x_0)$ be the data density, where the index 0 denotes the fact that the data is uncorrupted (original). Given an uncorrupted training sample $x_0 \sim p(x_0)$, the noised versions $x_1, x_2, \dots, x_T$ are obtained according to the following Markovian process:

$$p(x_t | x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t} \cdot x_{t-1}, \beta_t \cdot \mathbf{I}), \forall t \in \{1, \dots, T\}, \quad (1)$$

where T is the number of diffusion steps, $\beta_1, \dots, \beta_T \in [0, 1]$ are hyperparameters representing the variance schedule across diffusion steps, $\mathbf{I}$ is the identity matrix having the same dimensions as the input image x_0 , and $\mathcal{N}(x; \mu, \sigma)$ represents the normal distribution of mean μ and covariance σ that produces x . An important property of this recursive formulation is that it also allows the direct sampling of x_t , when t is drawn from a uniform distribution, i.e. $\forall t \sim \mathcal{U}(\{1, \dots, T\})$:

$$p(x_t | x_0) = \mathcal{N}(x_t; \sqrt{\hat{\beta}_t} \cdot x_0, (1 - \hat{\beta}_t) \cdot \mathbf{I}), \quad (2)$$

where $\hat{\beta}_t = \prod_{i=1}^t \alpha_i$ and $\alpha_t = 1 - \beta_t$. Essentially, Eq. (2) shows that we can sample any noisy version x_t via a single step, if we have the original image x_0 and fix a variance schedule β_t .

The sampling from $p(x_t|x_0)$ is performed via a reparametrization trick. In general, to standardize a sample x of a normal distribution $x \sim \mathcal{N}(\mu, \sigma^2 \cdot \mathbf{I})$, we subtract the mean μ and divide by the standard deviation σ , resulting in a sample $z = \frac{x-\mu}{\sigma}$ of the standard normal distribution $z \sim \mathcal{N}(0, \mathbf{I})$. The reparametrization trick does the inverse of this operation, starting with z and yielding the sample x by multiplying z with the standard deviation σ and adding the mean μ . If we translate this process to our case, then x_t is sampled from $p(x_t|x_0)$ as follows:

$$x_t = \sqrt{\hat{\beta}_t} \cdot x_0 + \sqrt{(1 - \hat{\beta}_t)} \cdot z_t, \quad (3)$$

where $z_t \sim \mathcal{N}(0, \mathbf{I})$.

Properties of β_t . If the variance schedule $(\beta_t)_{t=1}^T$ is chosen such that $\hat{\beta}_T \rightarrow 0$, then, according to Eq. (2), the distribution of x_T should be well approximated by the standard Gaussian distribution $\pi(x_T) = \mathcal{N}(0, \mathbf{I})$. Moreover, if each $(\beta_t)_{t=1}^T \ll 1$, then the reverse steps $p(x_{t-1}|x_t)$ have the same functional form as the forward process $p(x_t|x_{t-1})$ [1], [66]. Intuitively, the last statement is true when x_t is created with a very small step, as it becomes more likely that x_{t-1} comes from a region close to where x_t is observed, which allows us to model this region with a Gaussian distribution. To conform to the aforementioned properties, Ho *et al.* [2] choose $(\beta_t)_{t=1}^T$ to be linearly increasing constants between $\beta_1 = 10^{-4}$ and $\beta_T = 2 \cdot 10^{-2}$, where $T = 1000$.

Reverse process. By leveraging the above properties, we can generate new samples from $p(x_0)$ if we start from a sample $x_T \sim \mathcal{N}(0, \mathbf{I})$ and follow the reverse steps $p(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu(x_t, t), \Sigma(x_t, t))$. To approximate these steps, we can train a neural network $p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t))$ that receives as input the noisy image x_t and the embedding at time step t , and learns to predict the mean $\mu_\theta(x_t, t)$ and the covariance $\Sigma_\theta(x_t, t)$.

In an ideal scenario, we would train the neural network with a maximum likelihood objective such that the probability assigned by the model $p_\theta(x_0)$ to each training example x_0 is as large as possible. However, $p_\theta(x_0)$ is intractable because we have to marginalize over all the possible reverse trajectories to compute it. The solution to this problem [1], [2] is to minimize a variational lower-bound of the negative log-likelihood instead, which has the following formulation:

$$\begin{aligned} \mathcal{L}_{vlb} = & -\log p_\theta(x_0|x_1) + \text{KL}(p(x_T|x_0) \parallel \pi(x_T)) \\ & + \sum_{t>1} \text{KL}(p(x_{t-1}|x_t, x_0) \parallel p_\theta(x_{t-1}|x_t)), \end{aligned} \quad (4)$$

where KL denotes the Kullback-Leibler divergence between two probability distributions. The full derivation of this objective is presented in Appendix A. Upon analyzing each component, we can see that the second term can be removed because it does not depend on θ . The last term shows that the neural network is trained such that, at each time step t , $p_\theta(x_{t-1}|x_t)$ is as close as possible to the true posterior of the forward process when conditioned on the original image. Moreover, it can be proven that the posterior $p(x_{t-1}|x_t, x_0)$ is a Gaussian distribution, implying closed-form expressions for the KL divergences.

Algorithm 1 DDPM sampling method

Input:

T – the number of diffusion steps.

$\sigma_1, \dots, \sigma_T$ – the standard deviations for the reverse transitions.

Output:

x_0 – the sampled image.

Computation:

```

1:  $x_T \sim \mathcal{N}(0, \mathbf{I})$ 
2: for  $t = T, \dots, 1$  do
3:   if  $t > 1$  then
4:      $z \sim \mathcal{N}(0, \mathbf{I})$ 
5:   else
6:      $z = 0$ 
7:    $\mu_\theta = \frac{1}{\sqrt{\alpha_t}} \cdot \left( x_t - \frac{1-\alpha_t}{\sqrt{1-\hat{\beta}_t}} \cdot z_\theta(x_t, t) \right)$ 
8:    $x_{t-1} = \mu_\theta + \sigma_t \cdot z$ 
```

Ho *et al.* [2] propose to fix the covariance $\Sigma_\theta(x_t, t)$ to a constant value and rewrite the mean $\mu_\theta(x_t, t)$ as a function of noise, as follows:

$$\mu_\theta = \frac{1}{\sqrt{\alpha_t}} \cdot \left(x_t - \frac{1-\alpha_t}{\sqrt{1-\hat{\beta}_t}} \cdot z_\theta(x_t, t) \right). \quad (5)$$

These simplifications (more details in Appendix B) unlocked a new formulation of the objective $\mathcal{L}_{vlb}$, which measures, for a random time step t of the forward process, the distance between the real noise z_t and the noise estimation $z_\theta(x_t, t)$ of the model:

$$\mathcal{L}_{simple} = \mathbb{E}_{t \sim [1, T]} \mathbb{E}_{x_0 \sim p(x_0)} \mathbb{E}_{z_t \sim \mathcal{N}(0, \mathbf{I})} \|z_t - z_\theta(x_t, t)\|^2, \quad (6)$$

where $\mathbb{E}$ is the expected value, and $z_\theta(x_t, t)$ is the network predicting the noise in x_t . We underline that x_t is sampled via Eq. (3), where we use a random image x_0 from the training set.

The generative process is still defined by $p_\theta(x_{t-1}|x_t)$, but the neural network does not predict the mean and the covariance directly. Instead, it is trained to predict the noise from the image, and the mean is determined according to Eq. (5), while the covariance is fixed to a constant. Algorithm 1 formalizes the whole generative procedure.

2.2 Noise Conditioned Score Networks (NCSNs)

The score function of some data density $p(x)$ is defined as the gradient of the log density with respect to the input, $\nabla_x \log p(x)$. The directions given by these gradients are used by the Langevin dynamics algorithm [3] to move from a random sample (x_0) towards samples (x_N) in regions with high density. *Langevin dynamics* is an iterative method inspired from physics that can be used for data sampling. In physics, this method is used to determine the trajectory of a particle in a molecular system that allows interactions between the particle and the other molecules. The trajectory of the particle is influenced by a drag force of the system and by a random force motivated by the fast interactions between the molecules. In our case, we can think of the gradient of the log density as a force that drags a random sample through the data space into regions with high data density $p(x)$. The other term ω_i accounts, in physics, for

the random force, but for us, it is useful to escape local minima. Lastly, the value of γ weighs the impact of both forces, because it represents the friction coefficient of the environment where the particle resides. From the sampling point of view, γ controls the magnitude of the updates. In summary, the iterative updates of the Langevin dynamics are the following:

$$x_i = x_{i-1} + \frac{\gamma}{2} \nabla_x \log p(x) + \sqrt{\gamma} \cdot \omega_i, \quad (7)$$

where $i \in \{1, \dots, N\}$, γ controls the magnitude of the update in the direction of the score, x_0 is sampled from a prior distribution, the noise $\omega_i \sim \mathcal{N}(0, \mathbf{I})$ addresses the issue of getting stuck in local minima, and the method is applied recursively for $N \rightarrow \infty$ steps. Therefore, a generative model can employ the above method to sample from $p(x)$ after estimating the score with a neural network $s_\theta(x) \approx \nabla_x \log p(x)$. This network can be trained via score matching, a method that requires the optimization of the following objective:

$$\mathcal{L}_{sm} = \mathbb{E}_{x \sim p(x)} \|s_\theta(x) - \nabla_x \log p(x)\|_2^2. \quad (8)$$

In practice, it is impossible to minimize this objective directly, because $\nabla_x \log p(x)$ is unknown. However, there are other methods such as denoising score matching [67] and sliced score matching [68] that overcome this problem.

Although the described approach can be used for data generation, Song *et al.* [3] emphasize several issues when applying this method on real data. Most of the problems are linked with the manifold hypothesis. For example, the score estimation $s_\theta(x)$ is inconsistent when the data resides on a low-dimensional manifold and, among other implications, this could cause the Langevin dynamics to never converge to the high-density regions. In the same work [3], the authors demonstrate that these problems can be addressed by perturbing the data with Gaussian noise at different scales. Furthermore, they propose to learn score estimates for the resulting noisy distributions via a single noise conditioned score network (NCSN). Regarding the sampling, they adapt the strategy in Eq. (7) and use the score estimates associated with each noise scale.

Formally, given a sequence of Gaussian noise scales $\sigma_1 < \sigma_2 < \dots < \sigma_T$ such that $p_{\sigma_1}(x) \approx p(x_0)$ and $p_{\sigma_T}(x) \approx \mathcal{N}(0, \mathbf{I})$, we can train an NCSN $s_\theta(x, \sigma_t)$ with denoising score matching so that $s_\theta(x, \sigma_t) \approx \nabla_x \log(p_{\sigma_t}(x))$, $\forall t \in \{1, \dots, T\}$. We can derive $\nabla_x \log(p_{\sigma_t}(x))$ as follows:

$$\nabla_x \log p_{\sigma_t}(x_t|x) = -\frac{x_t - x}{\sigma_t^2}, \quad (9)$$

given that:

$$\begin{aligned} p_{\sigma_t}(x_t|x) &= \mathcal{N}(x_t; x, \sigma_t^2 \cdot \mathbf{I}) \\ &= \frac{1}{\sigma_t \cdot \sqrt{2\pi}} \cdot \exp\left(-\frac{1}{2} \cdot \left(\frac{x_t - x}{\sigma_t}\right)^2\right), \end{aligned} \quad (10)$$

where x_t is a noised version of x , and $\exp$ is the exponential function. Consequently, generalizing Eq. (8) for all $(\sigma_t)_{t=1}^T$ and replacing the gradient with the form in Eq. (9) leads to training $s_\theta(x_t, \sigma_t)$ by minimizing the following objective, $\forall t \in \{1, \dots, T\}$:

$$\mathcal{L}_{dsm} = \frac{1}{T} \sum_{t=1}^T \lambda(\sigma_t) \mathbb{E}_{p(x)} \mathbb{E}_{x_t \sim p_{\sigma_t}(x_t|x)} \left\| s_\theta(x_t, \sigma_t) + \frac{x_t - x}{\sigma_t^2} \right\|_2^2, \quad (11)$$

Algorithm 2 Annealed Langevin dynamics

Input:

$\sigma_1, \dots, \sigma_T$ – a sequence of Gaussian noise scales.

N – the number of Langevin dynamics iterations.

$\gamma_1, \dots, \gamma_T$ – the update magnitudes for each noise scale.

Output:

x_0^0 – the sampled image.

Computation:

```

1:  $x_T^0 \sim \mathcal{N}(0, \mathbf{I})$ 
2: for  $t = T, \dots, 1$  do
3:   for  $i = 1, \dots, N$  do
4:      $\omega \sim \mathcal{N}(0, \mathbf{I})$ 
5:      $x_t^i = x_t^{i-1} + \frac{\gamma_t}{2} \cdot s_\theta(x_t^{i-1}, \sigma_t) + \sqrt{\gamma_t} \cdot \omega$ 
6:    $x_{t-1}^0 = x_t^N$ 
```

where $\lambda(\sigma_t)$ is a weighting function. After training, the neural network $s_\theta(x_t, \sigma_t)$ will return estimates of the scores $\nabla_{x_t} \log(p_{\sigma_t}(x_t))$, having as input the noisy image x_t and the corresponding time step t .

At inference time, Song *et al.* [3] introduce the annealed Langevin dynamics, formally described in Algorithm 2. Their method starts with white noise and applies Eq. (7) for a fixed number of iterations. The required gradient (score) is given by the trained neural network conditioned on the time step T . The process continues for the following time steps, propagating the output of one step as input to the next. The final sample is the output returned for $t = 0$.

2.3 Stochastic Differential Equations (SDEs)

Similar to the previous two methods, the approach presented in [4] gradually transforms the data distribution $p(x_0)$ into noise. However, it generalizes over the previous two methods because, in its case, the diffusion process is considered to be continuous, thus becoming the solution of a stochastic differential equation (SDE). As shown in [69], the reverse process of this diffusion can be modeled with a reverse-time SDE which requires the score function of the density at each time step. Therefore, the generative model of Song *et al.* [4] employs a neural network to estimate the score functions, and generates samples from $p(x_0)$ by employing numerical SDE solvers. As in the case of NCSNs, the neural network receives the perturbed data and the time step as input, and produces an estimation of the score function.

The SDE of the forward diffusion process $(x_t)_{t=0}^T$, $t \in [0, T]$ has the following form:

$$\frac{\partial x}{\partial t} = f(x, t) + \sigma(t) \cdot \omega_t \iff \partial x = f(x, t) \cdot \partial t + \sigma(t) \cdot \partial \omega, \quad (12)$$

where ω_t is Gaussian noise, f is a function of x and t that computes the drift coefficient, and σ is a time-dependent function that computes the diffusion coefficient. In order to have a diffusion process as a solution for this SDE, the drift coefficient should be designed such that it gradually nullifies the data x_0 , while the diffusion coefficient controls how much Gaussian noise is added. The associated reverse-time SDE [69] is defined as follows:

$$\partial x = [f(x, t) - \sigma(t)^2 \cdot \nabla_x \log p_t(x)] \cdot \partial t + \sigma(t) \cdot \partial \hat{\omega}, \quad (13)$$

Algorithm 3 Euler-Maruyama sampling method**Input:** $\Delta t < 0$ – a negative step close to 0. f – a function of x and t that computes the drift coefficient. σ – a time-dependent function that computes the diffusion coefficient. $\nabla_x \log p_t(x)$ – the (approximated) score function. T – the final time step of the forward SDE.**Output:** x – the sampled image.**Computation:**

```

1:  $t = T$ 
2: while  $t > 0$  do
3:    $\Delta x = [f(x, t) - \sigma(t)^2 \cdot \nabla_x \log p_t(x)] \cdot \Delta t + \sigma(t) \cdot \Delta \hat{w}$ 
4:    $x = x + \Delta x$ 
5:    $t = t + \Delta t$ 

```

where $\hat{w}$ represents the Brownian motion when the time is reversed, from T to 0. The reverse-time SDE shows that, if we start with pure noise, we can recover the data by removing the drift responsible for data destruction. The removal is performed by subtracting $\sigma(t)^2 \cdot \nabla_x \log p_t(x)$.

We can train the neural network $s_\theta(x, t) \approx \nabla_x \log p_t(x)$ by optimizing the same objective as in Eq. (11), but adapted for the continuous case, as follows:

$$\begin{aligned} \mathcal{L}_{dsm}^* = \\ = \mathbb{E}_t \left[\lambda(t) \mathbb{E}_{p(x_0)} \mathbb{E}_{p_t(x_t|x_0)} \|s_\theta(x_t, t) - \nabla_{x_t} \log p_t(x_t|x_0)\|_2^2 \right], \end{aligned} \quad (14)$$

where λ is a weighting function, and $t \sim \mathcal{U}([0, T])$. We underline that, when the drift coefficient f is affine, $p_t(x_t|x_0)$ is a Gaussian distribution. When f does not conform to this property, we cannot use denoising score matching, but we can fallback to sliced score matching [68].

The sampling for this approach can be performed with any numerical method applied on the SDE defined in Eq. (13). In practice, the solvers do not work with the continuous formulation. For example, the Euler-Maruyama method fixes a tiny negative step Δt and executes Algorithm 3 until the initial time step $t = T$ becomes $t = 0$. At step 3, the Brownian motion is given by $\Delta \hat{w} = \sqrt{|\Delta t|} \cdot z$, where $z \sim \mathcal{N}(0, \mathbf{I})$.

Song *et al.* [4] present several contributions in terms of sampling techniques. They introduce the Predictor-Corrector sampler which generates better examples. This algorithm first employs a numerical method to sample from the reverse-time SDE, and then uses a score-based method as a corrector, for example the annealed Langevin dynamics described in the previous subsection. Furthermore, they show that ordinary differential equations (ODEs) can also be used to model the reverse process. Hence, another sampling strategy unlocked by the SDE interpretation is based on numerical methods applied to ODEs. The main advantage of this latter strategy is its efficiency.

2.4 Relation to Other Generative Models

We discuss below the connections between diffusion models and other types of generative models. We start with

likelihood-based methods and finish with generative adversarial networks.

Diffusion models have more aspects in common with VAEs [50]. For instance, in both cases, the data is mapped to a latent space and the generative process learns to transform the latent representations into data. Moreover, in both situations, the objective function can be derived as a lower-bound of the data likelihood. Nevertheless, there are essential differences between the two approaches and, further, we will mention some of them. The latent representation of a VAE contains compressed information about the original image, while diffusion models destroy the data entirely after the last step of the forward process. The latent representations of diffusion models have the same dimensions as the original data, while VAEs work better when the dimensions are reduced. Ultimately, the mapping to the latent space of a VAE is trainable, which is not true for the forward process of diffusion models because, as stated before, the latent is obtained by gradually adding Gaussian noise to the original image. The aforementioned similarities and differences can be the key for future developments of the two methods. For example, there already exists some work that builds more efficient diffusion models by applying them on the latent space of a VAE [17], [19].

Autoregressive models [62], [70] represent images as sequences of pixels. Their generative process produces new samples by generating an image pixel by pixel, conditioned on the previously generated pixels. This approach implies a unidirectional bias that clearly represents a limitation of this class of generative models. Esser *et al.* [28] see diffusion and autoregressive models as complementary and solve the above issue. Their method learns to reverse a multinomial diffusion process via a Markov chain where each transition is implemented as an autoregressive model. The global information provided to the autoregressive model is given by the previous step of the Markov chain.

Normalizing flows [63], [64] are a class of generative models that transform a simple Gaussian distribution into a complex data distribution. The transformation is done via a set of invertible functions which have an easy-to-compute Jacobian determinant. These conditions translate in practice into architectural restrictions. An important feature of this type of model is that the likelihood is tractable. Hence, the objective for training is the negative log-likelihood. When comparing with diffusion models, the two types of models have in common the mapping of the data distribution to Gaussian noise. However, the similarities between the two methods end here, because normalizing flows perform the mapping in a deterministic fashion by learning an invertible and differentiable function. These properties imply, in contrast to diffusion models, additional constraints on the network architecture, and a learnable forward process. A method which connects these two generative algorithms is DiffFlow. Introduced in [71], DiffFlow extends both diffusion models and normalizing flows such that the reverse and forward processes are both trainable and stochastic.

Energy-based models (EBMs) [60], [61], [72], [73] focus on providing estimates of unnormalized versions of density functions, called energy functions. Thanks to this property and in contrast to the previous likelihood-based methods, this type of model can be represented with any regression

neural network. However, due to this flexibility, the training of EBMs is difficult. One popular training strategy used in practice is score matching [72], [73]. Regarding the sampling, among other strategies, there is the Markov Chain Monte Carlo (MCMC) method, which is based on the score function. Therefore, the formulation from Subsection 2.2 of diffusion models can be considered to be a particular case of the energy-based framework, precisely the case when the training and sampling only require the score function.

GANs [52] were considered by many as state-of-the-art generative models in terms of the quality of the generated samples, before the recent rise of diffusion models [5]. GANs are also known as being difficult to train due to their adversarial objective [74], and often suffer from mode collapse. In contrast, diffusion models have a stable training process and provide more diversity because they are likelihood-based. Despite these advantages, diffusion models are still inefficient when compared to GANs, requiring multiple network evaluations during inference. A key aspect for comparison between GANs and diffusion models is their latent space. While GANs have a low-dimensional latent space, diffusion models preserve the original size of the images. Furthermore, the latent space of diffusion models is usually modeled as a random Gaussian distribution, being similar to VAEs. In terms of semantic properties, it was discovered that the latent space of GANs contains subspaces associated with visual attributes [75]. Thanks to this property, the attributes can be manipulated with changes in the latent space [75], [76]. In contrast, when such transformations are desired for diffusion models, the preferred procedure is the guidance technique [5], [77], which does not exploit any semantic property of the latent space. However, Song *et al.* [4] demonstrate that the latent space of diffusion models has a well-defined structure, illustrating that interpolations in this space lead to interpolations in the image space. In summary, from the semantic perspective, the latent space of diffusion models has been explored much less than in the case of GANs, but this may be one of the future research directions to be followed by the community.

3 A CATEGORIZATION OF DIFFUSION MODELS

We categorize diffusion models into a multi-perspective taxonomy considering different criteria of separation. Perhaps the most important criteria to separate the models are defined by (i) the task they are applied to, and (ii) the input signals they require. Furthermore, as there are multiple approaches in formulating a diffusion model, (iii) the underlying framework is another key factor for classifying diffusion models. Finally, the (iv) data sets used during training and evaluation are also of high importance, because they provide the means to compare different models on the same task. Our categorization of diffusion models according to the criteria enumerated above is presented in Table 1.

In the remainder of this section, we present several contributions on diffusion models, choosing the target task as the primary criterion to separate the methods. We opted for this classification criterion as it is fairly well-balanced and representative for research on diffusion models, facilitating a quick grasping of related works by readers working on specific tasks. Although the main task is usually related to

image generation, a considerable amount of work has been conducted to match and even surpass the performance of GANs on other topics, such as super-resolution, inpainting, image editing, image-to-image translation or segmentation.

3.1 Unconditional Image Generation

The diffusion models presented below are used to generate samples in an unconditional setting. Such models do not require supervision signals, being completely unsupervised. We consider this as the most basic and generic setting for image generation.

3.1.1 Denoising Diffusion Probabilistic Models

The work of Sohl-Dickstein *et al.* [1] formalizes diffusion models as described in Section 2.1. The proposed neural network is based on a convolutional architecture containing multi-scale convolution.

Austin *et al.* [78] extend the approach of Sohl-Dickstein *et al.* [1] to discrete diffusion models, studying different choices for the transition matrices used in the forward process. Their results are competitive with previous continuous diffusion models for the image generation task.

Ho *et al.* [2] extend the work presented in [1], proposing to learn the reverse process by estimating the noise in the image at each step. This change leads to an objective that resembles the denoising score matching applied in [3]. To predict the noise in an image, the authors use the Pixel-CNN++ architecture, which was introduced in [70].

On top of the work proposed by Ho *et al.* [2], Nichol *et al.* [6] introduce several improvements, observing that the linear noise schedule is suboptimal for low resolution. They propose a new option that avoids a fast information destruction towards the end of the forward process. Further, they show that it is required to learn the variance in order to improve the performance of diffusion models in terms of log-likelihood. This last change allows faster sampling, somewhere around 50 steps being required.

Song *et al.* [7] replace the Markov forward process used in [2] with a non-Markovian one. The generative process changes such that the model first predicts the normal sample, and then, it is used to estimate the next step in the chain. The change leads to a faster sampling procedure with a small impact on the quality of the generated samples. The resulting framework is known as the denoising diffusion implicit model (DDIM).

The work of Sinha *et al.* [16] presents the diffusion-decoding model with contrastive representations (D2C), a generative method which trains a diffusion model on latent representations produced by an encoder. The framework, which is based on the DDPM architecture presented in [2], produces images by mapping the latent representations to images.

In [94], the authors present a method to estimate the noise parameters given the current input at inference time. Their change improves the Fréchet Inception Distance (FID), while requiring less steps. The authors employ VGG-11 to estimate the noise parameters, and DDPM [2] to generate images.

The work of Nachmani *et al.* [93] suggests replacing the Gaussian noise distributions of the diffusion process with

TABLE 1

Our multi-perspective categorization of diffusion models applied in computer vision. To classify existing models, we consider three criteria: the task, the denoising condition, and the underlying approach (architecture). Additionally, we list the data sets on which the surveyed models are applied. We use the following abbreviations in the architecture column: D3PM (Discrete Denoising Diffusion Probabilistic Models), DSB (Diffusion Schrödinger Bridge), BDDM (Bilateral Denoising Diffusion Models), PNDM (Pseudo Numerical Methods for Diffusion Models), ADM (Ablated Diffusion Model), D2C (Diffusion-Decoding Models with Contrastive Representations), CCDF (Come-Closer-Diffuse-Faster), VQ-DDM (Vector Quantised Discrete Diffusion Model), BF-CNN (Bias-Free CNN), FDM (Flexible Diffusion Model), RVD (Residual Video Diffusion), RaMViD (Random Mask Video Diffusion).

Paper	Task	Denoising Condition	Architecture	Data Sets
Austin <i>et al.</i> [78]	image generation	unconditional	D3PM	CIFAR-10
Bao <i>et al.</i> [20]	image generation	unconditional	DDIM, Improved DDPM	CelebA, ImageNet, LSUN Bedroom, CIFAR-10
Benny <i>et al.</i> [79]	image generation	unconditional	DDPM, DDIM	CIFAR-10, ImageNet, CelebA
Bond-Taylor <i>et al.</i> [80]	image generation	unconditional	DDPM	LSUN Bedroom, LSUN Church, FFHQ
Choi <i>et al.</i> [81]	image generation	unconditional	DDPM	FFHQ, AFHQ-Dog, CUB, MetFaces
De <i>et al.</i> [82]	image generation	unconditional	DSB	MNIST, CelebA
Deasy <i>et al.</i> [83]	image generation	unconditional	NCSN	MNIST, Fashion-MNIST, CIFAR-10, CelebA
Deja <i>et al.</i> [84]	image generation	unconditional	Improved DDPM	Fashion-MNIST, CIFAR-10, CelebA
Dockhorn <i>et al.</i> [21]	image generation	unconditional	NCSN++, DDPM++	CIFAR-10
Ho <i>et al.</i> [2]	image generation	unconditional	DDPM	CIFAR-10, CelebA-HQ, LSUN
Huang <i>et al.</i> [59]	image generation	unconditional	DDPM	CIFAR-10, MNIST
Jing <i>et al.</i> [30]	image generation	unconditional	NCSN++, DDPM++	CIFAR-10, CelebA-256-HQ, LSUN Church
Jolicœur <i>et al.</i> [85]	image generation	unconditional	NCSN	CIFAR-10, LSUN Church, Stacked-MNIST
Jolicœur <i>et al.</i> [86]	image generation	unconditional	DDPM++, NCSN++	CIFAR-10, LSUN Church, FFHQ
Kim <i>et al.</i> [87]	image generation	unconditional	NCSN++, DDPM++	CIFAR-10, CelebA, MNIST
Kingma <i>et al.</i> [88]	image generation	unconditional	DDPM	CIFAR-10, ImageNet
Kong <i>et al.</i> [89]	image generation	unconditional	DDIM, DDPM	LSUN Bedroom, CelebA, CIFAR-10
Lam <i>et al.</i> [90]	image generation	unconditional	BDDM	CIFAR-10, CelebA
Liu <i>et al.</i> [91]	image generation	unconditional	PNDM	CIFAR-10, CelebA
Ma <i>et al.</i> [92]	image generation	unconditional	NCSN, NCSN++	CIFAR-10, CelebA, LSUN Bedroom, LSUN Church, FFHQ
Nachmani <i>et al.</i> [93]	image generation	unconditional	DDIM, DDPM	CelebA, LSUN Church
Nichol <i>et al.</i> [6]	image generation	unconditional	DDPM	CIFAR-10, ImageNet
Pandey <i>et al.</i> [19]	image generation	unconditional	DDPM	CelebA-HQ, CIFAR-10
San <i>et al.</i> [94]	image generation	unconditional	DDPM	CelebA, LSUN Bedroom, LSUN Church
Sehwag <i>et al.</i> [95]	image generation	unconditional	ADM	CIFAR-10, ImageNet
Sohl-Dickstein <i>et al.</i> [1]	image generation	unconditional	DDPM	MNIST, CIFAR-10, Dead Leaf Images
Song <i>et al.</i> [13]	image generation	unconditional	NCSN	FFHQ, CelebA, LSUN Bedroom, LSUN Tower, LSUN Church Outdoor
Song <i>et al.</i> [15]	image generation	unconditional	DDPM++	CIFAR-10, ImageNet 32×32
Song <i>et al.</i> [7]	image generation	unconditional	DDIM	CIFAR-10, CelebA, LSUN
Vahdat <i>et al.</i> [17]	image generation	unconditional	NCSN++	CIFAR-10, CelebA-HQ, MNIST
Wang <i>et al.</i> [96]	image generation	unconditional	DDIM	CIFAR-10, CelebA
Wang <i>et al.</i> [97]	image generation	unconditional	StyleGAN2, ProjectedGAN	CIFAR-10, STL-10, LSUN Bedroom, LSUN Church, AFHQ, FFHQ
Watson <i>et al.</i> [98]	image generation	unconditional	DDPM	CIFAR-10, ImageNet
Watson <i>et al.</i> [8]	image generation	unconditional	Improved DDPM	CIFAR-10, ImageNet 64×64
Xiao <i>et al.</i> [99]	image generation	unconditional	NCSN++	CIFAR-10
Zhang <i>et al.</i> [71]	image generation	unconditional	DDPM	CIFAR-10, MNIST
Zheng <i>et al.</i> [100]	image generation	unconditional	DDPM	CIFAR-10, CelebA, CelebA-HQ, LSUN Bedroom, LSUN Church
Bordes <i>et al.</i> [101]	conditional image generation	conditioned on latent representations	Improved DDPM	ImageNet
Campbell <i>et al.</i> [102]	conditional image generation	unconditional, conditioned on sound	DDPM	CIFAR-10, Lakh Pianoroll
Chao <i>et al.</i> [103]	conditional image generation	conditioned on class	Score SDE, Improved DDPM	CIFAR-10, CIFAR-100
Dhariwal <i>et al.</i> [5]	conditional image generation	unconditional, classifier guidance	ADM	LSUN Bedroom, LSUN Horse, LSUN Cat

Ho <i>et al.</i> [104]	conditional image generation	conditioned on label	DDPM	LSUN, ImageNet
Ho <i>et al.</i> [77]	conditional image generation	unconditional, classifier-free guidance	ADM	ImageNet 64×64, ImageNet 128×128
Karras <i>et al.</i> [105]	conditional image generation	unconditional, conditioned on class	DDPM++, NCSN++, DDPM, DDIM	CIFAR-10, ImageNet 64×64
Liu <i>et al.</i> [106]	conditional image generation	conditioned on text, image, style guidance	DDPM	FFHQ, LSUN Cat, LSUN Horse, LSUN Bedroom
Liu <i>et al.</i> [22]	conditional image generation	conditioned on text, 2D positions, relational descriptions between items, human facial attributes	Improved DDPM	CLEVR, Relational CLEVR, FFHQ
Lu <i>et al.</i> [107]	conditional image generation	unconditional, conditioned on class	DDIM	CIFAR-10, CelebA, ImageNet, LSUN Bedroom
Salimans <i>et al.</i> [108]	conditional image generation	unconditional, conditioned on class	DDIM	CIFAR-10, ImageNet, LSUN
Singh <i>et al.</i> [109]	conditional image generation	conditioned on noise	DDIM	ImageNet
Sinha <i>et al.</i> [16]	conditional image generation	unconditional, conditioned on label	D2C	CIFAR-10, CIFAR-100, fMoW, CelebA-64, CelebA-HQ-256, FFHQ-256
Ho <i>et al.</i> [34]	image-to-image translation	conditioned on image	Improved DDPM	ctest10k, places10k
Li <i>et al.</i> [37]	image-to-image translation	conditioned on image	DDPM	Face2Comic, Edges2Shoes, Edges2Handbags
Sasaki <i>et al.</i> [110]	image-to-image translation	conditioned on image	DDPM	CMP Facades, KAIST Multi-spectral Pedestrian
Wang <i>et al.</i> [36]	image-to-image translation	conditioned on image	DDIM	ADE20K, COCO-Stuff, DIODE
Wolleb <i>et al.</i> [38]	image-to-image translation	conditioned on image	Improved DDPM	BRATS
Zhao <i>et al.</i> [35]	image-to-image translation	conditioned on image	DDPM	CelebA-HQ, AFHQ
Gu <i>et al.</i> [111]	text-to-image generation	conditioned on text	VQ-Diffusion	CUB-200, Oxford 102 Flowers, MS-COCO
Jiang <i>et al.</i> [23]	text-to-image generation	conditioned on text	Transformer-based encoder-decoder	DeepFashion-MultiModal
Ramesh <i>et al.</i> [112]	text-to-image generation	conditioned on text	ADM	MS-COCO, AVA
Rombach <i>et al.</i> [11]	text-to-image generation	conditioned on text	LDM	OpenImages, WikiArt, LAION-2B-en, ArtBench
Saharia <i>et al.</i> [12]	text-to-image generation	conditioned on text	Imagen	MS-COCO, DrawBench
Shi <i>et al.</i> [9]	text-to-image generation	unconditional, conditioned on text	Improved DDPM	Conceptual Captions, MS-COCO
Zhang <i>et al.</i> [113]	text-to-image generation	unconditional, conditioned on text	DDIM	CIFAR-10, CelebA, ImageNet
Daniels <i>et al.</i> [25]	super-resolution	conditioned on image	NCSN	CIFAR-10, CelebA
Saharia <i>et al.</i> [18]	super-resolution	conditioned on image	DDPM++	FFHQ, CelebA-HQ, ImageNet-1K
Avrahami <i>et al.</i> [114]	image editing	conditioned on image and mask	DDPM, ADM	ImageNet, CUB, LSUN Bedroom, MS-COCO
Avrahami <i>et al.</i> [31]	region image editing	text guidance	DDPM	PaintByWord
Meng <i>et al.</i> [33]	image editing	conditioned on image	Score SDE, DDPM, Improved DDPM	LSUN, CelebA-HQ
Lugmayr <i>et al.</i> [29]	inpainting	unconditional	DDPM	CelebA-HQ, ImageNet
Nichol <i>et al.</i> [14]	inpainting	conditioned on image, text guidance	ADM	MS-COCO
Amit <i>et al.</i> [42]	image segmentation	conditioned on image	Improved DDPM	Cityscapes, Vaihingen, MoNuSeg
Baranchuk <i>et al.</i> [39]	image segmentation	conditioned on image	Improved DDPM	LSUN, FFHQ-256, ADE-Bedroom-30, CelebA-19
Batzolis <i>et al.</i> [24]	multi-task (inpainting, super-resolution, edge-to-image)	conditioned on image	DDPM	CelebA, Edges2Shoes
Batzolis <i>et al.</i> [115]	multi-task (image generation, super-resolution, inpainting, image-to-image translation)	unconditional	DDIM	ImageNet, CelebA-HQ, CelebA, Edges2Shoes
Blattmann <i>et al.</i> [116]	multi-task (image generation)	unconditional, conditioned on text, class	LDM	ImageNet
Choi <i>et al.</i> [32]	multi-task (image generation, image-to-image translation, image editing)	conditioned on image	DDPM	FFHQ, MetFaces
Chung <i>et al.</i> [26]	multi-task (inpainting, super-resolution, MRI reconstruction)	conditioned on image	CCDF	FFHQ, AFHQ, fastMRI knee
Esser <i>et al.</i> [28]	multi-task (image generation, inpainting)	unconditional, conditioned on class, image and text	ImageBART	ImageNet, Conceptual Captions, FFHQ, LSUN
Gao <i>et al.</i> [117]	multi-task (image generation, inpainting)	unconditional, conditioned on image	DDPM	CIFAR-10, LSUN, CelebA
Graikos <i>et al.</i> [40]	multi-task (image generation, image segmentation)	conditioned on class	DDIM	FFHQ-256, CelebA

Hu <i>et al.</i> [118]	multi-task (image generation, inpainting)	unconditional, conditioned on image	VQ-DDM	CelebA-HQ, LSUN Church
Khrulkov <i>et al.</i> [119]	multi-task (image generation, image-to-image translation)	conditioned on class	Improved DDPM	AFHQ, FFHQ, MetFaces, ImageNet
Kim <i>et al.</i> [120]	multi-task (image translation, multi-attribute transfer)	conditioned on image, portrait, stroke	DDIM	ImageNet, CelebA-HQ, AFHQ-Dog, LSUN Bedroom, Church
Luo <i>et al.</i> [121]	multi-task (point cloud generation, auto-encoding, unsupervised representation learning)	conditioned on shape latent	DDPM	ShapeNet
Lyu <i>et al.</i> [122]	multi-task (image generation, image editing)	unconditional, conditioned on class	DDPM	CIFAR-10, CelebA, ImageNet, LSUN Bedroom, LSUN Cat
Preechakul <i>et al.</i> [123]	multi-task (latent interpolation, attribute manipulation)	conditioned on latent representation	DDIM	CelebA-HQ
Rombach <i>et al.</i> [10]	multi-task (super-resolution, image generation, inpainting)	unconditional, conditioned on image	VQ-DDM	ImageNet, CelebA-HQ, FFHQ, LSUN
Shi <i>et al.</i> [124]	multi-task (super-resolution, inpainting)	conditioned on image	Improved DDPM	MNIST, CelebA
Song <i>et al.</i> [3]	multi-task (image generation, inpainting)	unconditional, conditioned on image	NCSN	MNIST, CIFAR-10, CelebA
Kadkhodaie <i>et al.</i> [125]	multi-task (Spatial super-resolution, Deblurring, Compressive sensing, Inpainting, Random missing pixels)	conditioned on linear measurements	BF-CNN	MNIST, Set5, Set68, Set14
Song <i>et al.</i> [4]	multi-task (image generation, inpainting, colorization)	unconditional, conditioned on image, class	NCSN++, DDPM++	CelebA-HQ, CIFAR-10, LSUN
Hu <i>et al.</i> [126]	medical image-to-image translation	conditioned on image	DDPM	ONH
Chung <i>et al.</i> [127]	medical image generation	conditioned on measurements	NCSN++	fastMRI knee
Özbey <i>et al.</i> [128]	medical image generation	conditioned on image	Improved DDPM	IXI, Gold Atlas - Male Pelvis
Song <i>et al.</i> [129]	medical image generation	conditioned on measurements	NCSN++	LIDC, LDCT Image and Projection, BRATS
Wolleb <i>et al.</i> [41]	medical image segmentation	conditioned on image	Improved DDPM	BRATS
Sanchez <i>et al.</i> [130]	medical image segmentation and anomaly detection	conditioned on image and binary variable	ADM	BRATS
Pinaya <i>et al.</i> [44]	medical image segmentation and anomaly detection	conditioned on image	DDPM	MedNIST, UK Biobank Images, WMH, BRATS, MSLUB
Wolleb <i>et al.</i> [45]	medical image anomaly detection	conditioned on image	DDIM	CheXpert, BRATS
Wyatt <i>et al.</i> [46]	medical image anomaly detection	conditioned on image	ADM	NFBS, 22 MRI scans
Harvey <i>et al.</i> [131]	video generation	conditioned on frames	FDM	GQN-Mazes, MineRL Navigate, CARLA Town01
Ho <i>et al.</i> [132]	video generation	unconditional, conditioned on text	DDPM	101 Human Actions
Yang <i>et al.</i> [133]	video generation	conditioned on video representation	RVD	BAIR, KTH Actions, Simulation, Cityscapes
Höppe <i>et al.</i> [134]	video generation and infilling	conditioned on frames	RaMViD	BAIR, Kinetics-600, UCF-101
Giannone <i>et al.</i> [135]	few-shot image generation	conditioned on image	Improved DDPM	CIFAR-FS, mini-ImageNet, CelebA
Jeanneret <i>et al.</i> [136]	counterfactual explanations	unconditional	DDPM	CelebA
Sanchez <i>et al.</i> [137]	counterfactual estimates	conditional	ADM	MNIST, ImageNet
Kawar <i>et al.</i> [27]	image restoration	conditioned on image	DDIM	FFHQ, ImageNet
Özdenizci <i>et al.</i> [138]	image restoration	conditioned on image	DDPM	Snow100K, Outdoor-Rain, RainDrop
Kim <i>et al.</i> [139]	image registration	conditioned on image	DDPM	Radboud Faces, OASIS-3
Nie <i>et al.</i> [140]	adversarial purification	conditioned on image	Score SDE, Improved DDPM, DDIM	CIFAR-10, ImageNet, CelebA-HQ
Wang <i>et al.</i> [141]	semantic image generation	conditioned on semantic map	DDPM	Cityscapes, ADE20K, CelebAMask-HQ
Zhou <i>et al.</i> [142]	shape generation and completion	unconditional, conditional shape completion	DDPM	ShapeNet, PartNet
Zimmermann <i>et al.</i> [43]	classification	conditioned on label	DDPM++	CIFAR-10

two other distributions, a mixture of two Gaussians and the Gamma distribution. The results show better FID values and faster convergence thanks to the Gamma distribution that has higher modeling capacity.

Lam *et al.* [90] learn the noise scheduling for sampling. The noise schedule for training remains linear as before. After training the score network, they assume it to be close to the optimal value in order to use it for noise schedule training. The inference is composed of two steps. First, the

schedule is determined by fixing two initial hyperparameters. The second step is the usual reverse process with the determined schedule.

Bond-Taylor *et al.* [80] present a two-stage process, where they apply vector quantization to images to obtain discrete representations, and use a transformer [143] to reverse a discrete diffusion process, where the elements are randomly masked at each step. The sampling process is faster because the diffusion is applied to a highly compressed representa-

tion, which allows fewer denoising steps (50-256).

Watson *et al.* [98] propose a dynamic programming algorithm that finds the optimal inference schedule, having a time complexity of $\mathcal{O}(T)$, where T is the number of steps. They conduct their image generation experiments on CIFAR-10 and ImageNet, using the DDPM architecture.

In a different work, Watson *et al.* [8] begin by presenting how a reparametrization trick can be integrated within the backward process of diffusion models in order to optimize a family of fast samplers. Using the Kernel Inception Distance as loss function, they show how optimization can be done using stochastic gradient descent. Next, they propose a special parametrized family of samplers, which, using the same process as before, can achieve competitive results with fewer sampling steps. Using FID and Inception Score (IS) as metrics, the method seems to outperform some diffusion model baselines.

Similar to Bond-Taylor *et al.* [80] and Watson *et al.* [8], [98], Xiao *et al.* [99] try to improve the sampling speed, while also maintaining the quality, coverage and diversity of the samples. Their approach is to integrate a GAN in the denoising process to discriminate between real samples (forward process) and fake ones (denoised samples from the generator), with the objective of minimizing the softened reverse KL divergence [144]. However, the model is modified by directly generating a clean (fully denoised) sample and conditioning the fake example on it. Using the NCSN++ architecture with adaptive group normalization layers for the GAN generator, they achieve similar FID values in both image synthesis and stroke-based image generation, at sampling rates of about 20 to 2000 times faster than other diffusion models.

Kingma *et al.* [88] introduce a class of diffusion models that obtains state-of-the-art likelihoods on image density estimation. They add Fourier features to the input of the network to predict the noise, and investigate if the observed improvement is specific to this class of models. Their results confirm the hypothesis, *i.e.* previous state-of-the-art models did not benefit from this change. As a theoretical contribution, they show that the diffusion loss is impacted by the signal-to-noise ratio function only through its extremes.

Following the work presented in [117], Bao *et al.* [20] propose an inference framework that does not require training using non-Markovian diffusion processes. By first deriving an analytical estimate of the optimal mean and variance with respect to a score function, and using a pretrained scored-based model to obtain score values, they show better results, while being 20 to 40 times more time-efficient. The score is approximated by Monte Carlo sampling. However, the score is clipped within some precomputed bounds in order to diminish any bias of the pretrained DDPM model.

Zheng *et al.* [100] suggest truncating the process at an arbitrary step, and propose a method to inverse the diffusion from this distribution by relaxing the constraint of having Gaussian random noise as the final output of the forward diffusion. To solve the issue of starting the reverse process from a non-tractable distribution, an implicit generative distribution is used to match the distribution of the diffused data. The proxy distribution is fit either through a GAN or a conditional transport. We note that the generator utilizes the same U-Net model as the sampler of the diffusion model,

thus not adding extra parameters to be trained.

Deja *et al.* [84] begin by analyzing the backward process of a diffusion model and postulate that it is formed of two models, a generator and a denoiser. Thus, they propose to explicitly split the process into two components: the denoiser via an auto-encoder, and the generator via a diffusion model. Both models use the same U-Net architecture.

Wang *et al.* [97] start from the idea presented by Arjovsky *et al.* [145] and Sønderby *et al.* [146] to augment the input data of the discriminator by adding noise. This is achieved in [97] by injecting noise from a Gaussian mixture distribution composed of weighted diffused samples from the clean image at various time steps. The noise injection mechanism is applied to both real and fake images. The experiments are conducted on a wide range of data sets covering multiple resolutions and high diversity.

3.1.2 Score-Based Generative Models

Starting from a previous work [3], Song *et al.* [13] present several improvements which are based on theoretical and empirical analyses. They address both training and sampling phases. For training, the authors show new strategies to choose the noise scales and how to incorporate the noise conditioning into NCSNs [3]. For sampling, they propose to apply an exponential moving average to the parameters and select the hyperparameters for the Langevin dynamics such that the step size verifies a certain equation. The proposed changes unlock the application of NCSNs on high-resolution images.

Jolicœur-Martineau *et al.* [85] introduce an adversarial objective along with denoising score matching to train score-based models. Furthermore, they propose a new sampling procedure called Consistent Annealed Sampling and prove that it is more stable than the annealed Langevin method. Their image generation experiments show that the new objective returns higher quality examples without an impact on diversity. The suggested changes are tested on the architectures proposed in [2], [3], [13].

Song *et al.* [15] improve the likelihood of score-based diffusion models. They achieve this through a new weighting function for the combination of the score matching losses. For their image generation experiments, they use the DDPM++ architecture introduced in [4].

In [82], the authors present a score-based generative model as an implementation of Iterative Proportional Fitting (IPF), a technique used to solve the Schrödinger bridge problem. This novel approach is tested on image generation, as well as data set interpolation, which is possible because the prior can be any distribution.

Vahdat *et al.* [17] train diffusion models on latent representations. They use a VAE to encode to and decode from the latent space. This work achieves up to 56 times faster sampling. For the image generation experiments, the authors employ the NCSN++ architecture introduced in [4].

3.1.3 Stochastic Differential Equations

DiffFlow is introduced in [71] as a new generative modeling approach that combines normalizing flows and diffusion probabilistic models. From the perspective of diffusion models, the method has a sampling procedure that is up to 20 times more efficient, thanks to a learnable forward process

which skips unneeded noise regions. The authors perform experiments using the same architecture as in [2].

Jolicœur-Martineau *et al.* [86] introduce a new SDE solver that is between $2\times$ and $5\times$ faster than Euler-Maruyama and does not affect the quality of the generated images. The solver is evaluated in a set of image generation experiments with pretrained models from [3].

Wang *et al.* [96] present a new deep generative model based on Schrödinger bridge. This is a two-stage method, where the first stage learns a smoothed version of the target distribution, and the second stage derives the actual target.

Focusing on score-based models, Dockhorn *et al.* [21] utilize a critically-damped Langevin diffusion process by adding another variable (velocity) to the data, which is the only source of noise in the process. Given the new diffusion space, the resulting score function is demonstrated to be easier to learn. The authors extend their work by developing a more suitable score objective called hybrid score matching, as well as a sampling method, by solving the SDE through integration. The authors adapt the NCSN++ and DDPM++ architectures to accept both data and velocity, being evaluated on unconditional image generation and outperforming similar score-based diffusion models.

Motivated by the limitations of high-dimensional score-based diffusion models due to the Gaussian noise distribution, Deasy *et al.* [83] extend denoising score matching to generalize to the normal noising distribution. By adding a heavier tailed distribution, their experiments on several data sets show promising results, as the generative performance improves in certain cases (depending on the shape of the distribution). An important scenario in which the method excels is on data sets with unbalanced classes.

Jing *et al.* [30] try to shorten the duration of the sampling process of diffusion models by reducing the space onto which diffusion is realized, *i.e.* the larger the time step in the diffusion process, the smaller the subspace. The data is projected onto a finite set of subspaces, at specific times, each being associated with a score model. This results in reduced computational costs, while the performance is increased. The work is limited to natural image synthesis. Evaluating the method in unconditional image generation, the authors achieve similar or better performance compared with state-of-the-art models, while having a lower inference time. The method is demonstrated to work for the inpainting task as well.

Kim *et al.* [87] propose to change the diffusion process into a non-linear one. This is achieved by using a trainable normalizing flow model which encodes the image in the latent space, where it can now be linearly diffused to the noise distribution. A similar logic is then applied to the denoising process. This method is applied on the NCSN++ and DDPM++ frameworks, while the normalizing flow model is based on ResNet.

Ma *et al.* [92] aim to make the backward diffusion process more time-efficient, while maintaining the synthesis performance. Within the family of score-based diffusion models, they begin to analyze the reverse diffusion in the frequency domain, subsequently applying a space-frequency filter to the sampling process, which aims to integrate information about the target distribution into the initial noise sampling. The authors conduct experiments with NCSN [3] and

NCSN++ [4], where the proposed method clearly shows speed improvements in image synthesis (by up to 20 times less sampling steps), while keeping the same satisfactory generation quality for both low and high-resolution images.

3.2 Conditional Image Generation

We next showcase diffusion models that are applied to conditional image synthesis. The condition is commonly based on various source signals, in most cases some class labels being used. Some methods perform both unconditional and conditional generation, which are also discussed here.

3.2.1 Denoising Diffusion Probabilistic Models

Dhariwal *et al.* [5] introduce few architectural changes to improve the FID of diffusion models. They also propose classifier guidance, a strategy which uses the gradients of a classifier to guide the diffusion during sampling. They conduct both unconditional and conditional image generation experiments.

Bordes *et al.* [101] examine representations resulting from self-supervised tasks by visualizing and comparing them to the original image. They also compare representations generated from different sources. Thus, a diffusion model is used to generate samples that are conditioned on these representations. The authors implement several modifications to the U-Net architecture presented by Dhariwal *et al.* [5], such as adding conditional batch normalization layers, and mapping the vector representation through a fully connected layer.

The method presented in [95] allows diffusion models to produce images from low-density regions of the data manifold. They use two new losses to guide the reverse process. The first loss guides the diffusion towards low-density regions, while the second enforces the diffusion to stay on the manifold. Moreover, they demonstrate that their diffusion model does not memorize the examples from the low-density neighborhoods, generating novel images. The authors employ an architecture similar to that of Dhariwal *et al.* [5].

Kong *et al.* [89] define a bijection between the continuous diffusion steps and the noise levels. With the defined bijection, they are able to construct an approximate diffusion process which requires less steps. The method is tested using the previous DDIM [7] and DDPM [2] architectures on image generation.

Pandey *et al.* [19] build a generator-refiner framework, where the generator is a VAE and the refiner is a DDPM conditioned by the output of the VAE. The latent space of the VAE can be used to control the content of the generated image because the DDPM only adds the details. After training the framework, the resulting DDPM is able to generalize to different noise types. More specifically, if the reverse process is not conditioned on the VAE's output, but on different noise types, the DDPM is able to reconstruct the initial image.

Ho *et al.* [104] introduce Cascaded Diffusion Models (CDM), an approach for generating high-resolution images conditioned on ImageNet classes. Their framework contains multiple diffusion models, where the first model from the pipeline generates low-resolution images conditioned on

the image class. The subsequent models are responsible for generating images of increasingly higher resolutions. These models are conditioned on both the class and the low-resolution image.

Benny *et al.* [79] study the advantages and disadvantages of predicting the image instead of the noise during the reverse process. They conclude that some of the discovered problems could be addressed by interpolating the two types of output. They modify previous architectures to return both the noise and the image, as well as a value that controls the importance of the noise when performing the interpolation. The strategy is evaluated on top of the DDPM and DDIM architectures.

Choi *et al.* [81] investigate the impact of the noise levels on the visual concepts learned by diffusion models. They modify the conventional weighting scheme of the objective function to a new one that enforces diffusion models to learn rich visual concepts. The method groups the noise levels into three categories (coarse, content and clean-up) according to the signal-to-noise ratio, *i.e.* small SNR is coarse, medium SNR is content, large SNR is clean-up. The weighting function assigns lower weights to the last group.

Singh *et al.* [109] propose a novel method for conditional image generation. Instead of conditioning the signal throughout the sampling process, they present a method to condition the noise signal (from where the sampling starts). Using Inverting Gradients [147], the noise is injected with information about localization and orientation of the conditioned class, while maintaining the same random Gaussian distribution.

Describing the resembling functionality of diffusion models and energy-based models, and leveraging the compositional structure of the latter models, Liu *et al.* [22] propose to combine multiple diffusion models for conditional image synthesis. In the reverse process, the composition of multiple diffusion models, each associated with a different condition, can be achieved either through conjunction or negation.

3.2.2 Score-Based Generative Models

The works of Song *et al.* [4] and Dhariwal *et al.* [5] on scored-based conditional diffusion models based on classifier guidance inspired Chao *et al.* [103] to develop a new training objective which reduces the potential discrepancy between the score model and the true score. The loss of the classifier is modified into a scaled cross-entropy added to a modified score matching loss.

3.2.3 Stochastic Differential Equations

Ho *et al.* [77] introduce a guidance method that does not require a classifier. It just needs one conditional diffusion model and one unconditional version, but they use the same model to learn both cases. The unconditional model is trained with the class identifier being equal to 0. The idea is based on the implicit classifier derived from the Bayes rule.

Liu *et al.* [91] investigate the usage of conventional numerical methods to solve the ODE formulation of the reverse process. They find that these methods return lower quality samples compared with the previous approaches. Therefore, they introduce pseudo-numerical methods for

diffusion models. Their idea splits the numerical methods into two parts, the gradient part and the transfer part. The transfer part (standard methods have a linear transfer part) is replaced such that the result is as close as possible to the target manifold. As a last step, they show how this change solves the problems discovered when using conventional approaches.

Tachibana *et al.* [148] address the slow sampling problem of DDPMs. They propose to decrease the number of sampling steps by increasing the order (from one to two) of the stochastic differential equation solver (denoising part). While preserving the network architecture and score matching function, they adopt the Itô-Taylor expansion scheme for the sampler, as well as substitute some derivative terms in order to simplify the calculation. They reduce the number of backward steps while retaining the performance. On top of these, another contribution is the new noise schedule.

Karras *et al.* [105] try to separate diffusion scored-based models into individual components that are independent of each other. This separation allows modifying a single part without affecting the other units, thus facilitating the improvement of diffusion models. Using this framework, the authors first present a sampling process that uses Heun's method as the ODE solver, which reduces the neural function evaluations while maintaining the FID score. They further show that a stochastic sampling process brings great performance benefits. The second contribution is related to training the score-based model by preconditioning the neural network on its input and the corresponding targets, as well as using image augmentation.

Within the context of both unconditional and class-conditional image generation, Salimans *et al.* [108] propose a technique for reducing the number of sampling steps. They distill the knowledge of a trained teacher model, represented by a deterministic DDIM, into a student model that has the same architecture, but halving the number of sampling steps. In other words, the target of the student is to take two consecutive steps of the teacher. Furthermore, this process can be repeated until the desired number of sampling steps is reached, while maintaining the same image synthesis quality. Finally, three versions of the model and two loss functions are explored in order to facilitate the distillation process and reduce the number of sampling steps (from 8192 to 4).

Campbell *et al.* [102] demonstrate a continuous-time formulation of denoising diffusion models that is capable of working with discrete data. The work models the forward continuous-time Markov chain diffusion process via a transition rate matrix, and the backward denoising process via a parametric approximation of the inverse transition rate matrix. Further contributions are related to the training objective, the matrix construction, and an optimized sampler.

The interpretation of diffusion models as ODEs proposed by Song *et al.* [4] is reformulated by Lu *et al.* [107] in a form that can be solved using an exponential integrator. Other contributions of Lu *et al.* [107] are an ODE solver that approximates the integral term of the new formulation using Taylor expansion (first order to third order), and an algorithm that adapts the time step schedule, being 4 to 16 times faster.

3.3 Image-to-Image Translation

Saharia *et al.* [34] propose a framework for image-to-image translation using diffusion models, focusing on four tasks: colorization, inpainting, uncropping and JPEG restoration. The proposed framework is the same across all four tasks, meaning that it does not suffer custom changes for each task. The authors begin by comparing L_1 and L_2 losses, suggesting that L_2 is preferred, as it leads to a higher sample diversity. Finally, they reconfirm the importance of self-attention layers in conditional image synthesis.

To translate an unpaired set of images, Sasaki *et al.* [110] propose a method involving two jointly trained diffusion models. During the reverse denoising process, at every step, each model is also conditioned on the other's intermediate sample. Furthermore, the loss function of the diffusion models is regularized using the cycle-consistency loss [149].

The aim of Zhao *et al.* [35] is to improve current image-to-image translation score-based diffusion models by utilizing data from a source domain with an equal significance. An energy-based function trained on both source and target domains is employed in order to guide the SDE solver. This leads to generating images that preserve the domain-agnostic features, while translating characteristics specific to the source domain to the target domain. The energy function is based on two feature extractors, each specific to a domain.

Leveraging the power of pretraining, Wang *et al.* [36] employ the GLIDE model [14] and train it to obtain a rich semantic latent space. Starting from the pretrained version and replacing the head to adapt to any conditional input, the model is fine-tuned on some specific image generation downstream tasks. This is done in two steps, where the first step is to freeze the decoder and train only the new encoder, and the second step is to train them simultaneously. Finally, the authors employ adversarial training and normalize the classifier-free guidance to enhance generation quality.

Li *et al.* [37] introduce a diffusion model for image-to-image translation that is based on Brownian bridges, as well as GANs. The proposed process begins by encoding the image with a VQ-GAN [150]. Within the resulting quantized latent space, the diffusion process, formulated as a Brownian bridge, maps between the latent representations of the source domain and target domain. Finally, another VQ-GAN decodes the quantized vectors in order to synthesize the image in the new domain. The two GAN models are independently trained on their specific domains.

Continuing their previous work proposed in [45], Wolleb *et al.* [38] extend the diffusion model by replacing the classifier with another model specific to the task. Thus, at every step of the sampling process, the gradient of the task-specific network is infused. The method is demonstrated with a regressor (based on an encoder) or with a segmentation model (using the U-Net architecture), whereas the diffusion model is based on existing frameworks [2], [6]. This setting has the advantage of eliminating the need to retrain the whole diffusion model, except for the task-specific model.

3.4 Text-to-Image Synthesis

Perhaps the most impressive results of diffusion models are attained on text-to-image synthesis, where the capability of combining unrelated concepts, such as objects, shapes and

textures, to generate unusual examples comes to light. To confirm this statement, we used Stable Diffusion [10] to generate images based on various text prompts, and the results are shown in Figure 2.

Imagen is introduced in [12] as an approach for text-to-image synthesis. It consists of one encoder for the text sequence and a cascade of diffusion models for generating high-resolution images. These models are also conditioned on the text embeddings returned by the encoder. Moreover, the authors introduce a new set of captions (DrawBench) for text-to-image evaluations. Regarding the architecture, the authors develop Efficient U-Net to improve efficiency, and apply this architecture in their text-to-image generation experiments.

Gu *et al.* [111] introduce the VQ-Diffusion model, a method for text-to-image synthesis that does not have the unidirectional bias of previous approaches. With its masking mechanism, the proposed method avoids the accumulation of errors during inference. The model has two stages, where the first stage is based on a VQ-VAE that learns to represent an image via discrete tokens, and the second stage is a discrete diffusion model that operates on the discrete latent space of the VQ-VAE. The training of the diffusion model is conditioned on caption embeddings. Inspired from masked language modeling, some tokens are replaced with a *[mask]* token.

Avrahami *et al.* [31] present a text-conditional diffusion model conditioned on CLIP [151] image and text embeddings. This is a two-stage approach, where the first stage generates the image embedding, and the second stage (decoder) produces the final image conditioned on the image embedding and the text caption. To generate image embeddings, the authors use a diffusion model in the latent space. They perform a subjective human assessment to evaluate their generative results.

Addressing the slow sampling inconvenience of diffusion models, Zhang *et al.* [113] focus their work on a new discretization scheme that reduces the error and allows a greater step size, *i.e.* a lower number of sampling steps. By using high-order polynomial extrapolations in the score function and an Exponential Integrator for solving the reverse SDE, the number of network evaluations is drastically reduced, while maintaining the generation capabilities.

Shi *et al.* [9] combine a VQ-VAE [152] and a diffusion model to generate images. Starting from the VQ-VAE, the encoding functionality is preserved, while the decoder is replaced by a diffusion model. The authors use the U-Net architecture from [6], injecting the image tokens into the middle block.

Building on top of the work presented in [116], Rombach *et al.* [11] introduce a modification to create artistic images using the same procedure: extract the k -nearest neighbors in the CLIP [151] latent space of an image from a database, then generate a new image by guiding the reverse denoising process with these embeddings. As the CLIP latent space is shared by text and images, the diffusion can be guided by text prompts as well. However, at inference time, the database is replaced with another one that contains artistic images. Thus, the model generates images within the style of the new database.

Jiang *et al.* [23] present a framework to generate images

of full-body humans with rich clothing representation given three inputs: a human pose, a text description of the clothes' shape, and another text of the clothing texture. The first stage of the method encodes the former text prompt into an embedding vector and infuses it into the module (encoder-decoder based) that generates a map of forms. In the second stage, a diffusion-based transformer samples an embedded representation of the latter text prompt from multiple multi-level codebooks (each specific to a texture), a mechanism suggested in VQ-VAE [152]. Initially, the codebook indices at coarser levels are sampled, and then, using a feed-forward network, the finer-level indices are predicted. The text is encoded using Sentence-BERT [153].

3.5 Image Super-Resolution

Saharia *et al.* [18] apply diffusion models to super-resolution. Their reverse process learns to generate high quality images conditioned on low-resolution versions. This work employs the architectures presented in [2], [6] and the following data sets: CelebA-HQ, FFHQ and ImageNet.

Daniels *et al.* [25] use score-based models to sample from the Sinkhorn coupling of two distributions. Their method models the dual variables with neural networks, then solves the problem of optimal transport. After training the neural networks, the sampling can be performed via Langevin dynamics and a score-based model. They run experiments on image super-resolution using a U-Net architecture.

3.6 Image Editing

Meng *et al.* [33] utilize diffusion models in various guided image generation tasks, *i.e.* stroke painting or stroke-based editing and image composition. Starting from an image that contains some form of guidance, its properties (such as shapes and colors) are preserved, while the deformations are smoothed out by progressively adding noise (forward process of the diffusion model). Then, the result is denoised (reverse process) to create a realistic image according to the guidance. Images are synthesized with a generic diffusion model by solving the reverse SDE, without requiring any custom data set or modifications for training.

One of the first approaches for editing specific regions of images based on natural language descriptions is introduced in [31]. The regions to be modified are specified by the user via a mask. The method relies on CLIP guidance to generate an image according to the text input, but the authors observe that combining the output with the original image at the end does not produce globally coherent images. Hence, they modify the denoising process to fix the issue. More precisely, after each step, the authors apply the mask on the latent image, while also adding the noisy version of the original image.

Extending the work presented in [10], Avrahami *et al.* [114] apply latent diffusion models for editing images locally, using text. A VAE encodes the image and the adaptive-to-time mask (region to edit) into the latent space where the diffusion process occurs. Each sample is iteratively denoised, while being guided by the text within the region of interest. However, inspired by Blended Diffusion [31], the image is combined with the masked region in the latent space that is noised at the current time step. Finally, the

sample is decoded with the VAE to generate the new image. The method demonstrates superior performance while being comparably faster.

3.7 Image Inpainting

Nichol *et al.* [14] train a diffusion model conditioned on text descriptions and also study the effectiveness of classifier-free and CLIP-based guidance. They obtain better results with the first option. Moreover, they fine-tune the model for image inpainting, unlocking image modifications based on text input.

Lugmay *et al.* [29] present an inpainting method agnostic to the mask form. They use an unconditional diffusion model for this, but modify its reverse process. They produce the image at step $t - 1$ by sampling the known region from the masked image, and the unknown region by applying denoising to the image obtained at step t . With this procedure, the authors observe that the unknown region has the right structure, while also being semantically incorrect. Further, they solve the issue by repeating the proposed step for a number of times and, at each iteration, they replace the previous image from step t with a new sample obtained from the denoised version generated at step $t - 1$.

3.8 Image Segmentation

Baranchuk *et al.* [39] demonstrate how diffusion models can be used in semantic segmentation. Taking the feature maps (middle blocks) at different scales from the decoder of the U-Net (used in the denoising process) and concatenating them (upsampling the feature maps in order to have the same dimensions), they can be used to classify each pixel by further attaching an ensemble of multi-layer perceptrons. The authors show that these feature maps, extracted at later steps in the denoising process, contain rich representations. The experiments show that segmentation based on diffusion models outperforms most baselines.

Amit *et al.* [42] propose the use of diffusion probabilistic models for image segmentation through extending the architecture of the U-Net encoder. The input image and the current estimated image are passed through two different encoders and combined together through summation. The result is then supplied to the encoder-decoder of the U-Net. Due to the stochastic noise infused at every time step, multiple samples for a single input image are generated and used to compute the mean segmentation map. The U-Net architecture is based on a previous work [6], while the input image generator is built with Residual Dense Blocks [154]. The denoised sample generator is a simple 2D convolutional layer.

3.9 Multi-Task Approaches

A series of diffusion models have been applied to multiple tasks, demonstrating a good generalization capacity across tasks. We discuss such contributions below.

Song *et al.* [3] present the noise conditional score network (NCSN), an approach which estimates the score function at different noise scales. For sampling, they introduce an annealed version of Langevin dynamics and use it to report results in image generation and inpainting. The NCSN

architecture is mainly based on the work presented in [155], with small changes such as replacing batch normalization with instance normalization.

Kadkhodaie *et al.* [125] train a neural network to restore images corrupted with Gaussian noise, generated using random standard deviations that are restricted to a particular range. After training, the difference between the output of the neural network and the noisy image received as input is proportional with the gradient of the log-density of the noisy data. This property is based on previous work done in [156]. For image generation, the authors use the mentioned difference as gradient (score) estimation and sample from the implicit data prior of the network by employing an iterative method similar to the annealed Langevin dynamics from [3]. However, the two sampling methods have some dissimilarities, for example the noise injected in the iterative updates follow distinct strategies. In [125], the injected noise is adapted according to the network's estimate, while in [3], it is fixed. Moreover, the gradient estimates in [3] are learned by score matching, while Kadkhodaie *et al.* [125] rely on the previously mentioned property to compute the gradients. The contribution of Kadkhodaie *et al.* [125] develops even further by adapting the algorithm to linear inverse problems, such as deblurring and super-resolution.

The SDE formulation of diffusion models introduced in [4] generalizes over several previous methods [1]–[3]. Song *et al.* [4] present the forward and reverse diffusion processes as solutions of SDEs. This technique unlocks new sampling methods, such as the Predictor-Corrector sampler, or the deterministic sampler based on ODEs. The authors carry out experiments on image generation, inpainting and colorization.

Batzolis *et al.* [115] introduce a new forward process in diffusion models, called non-uniform diffusion. This is determined by each pixel being diffused with a different SDE. Multiple networks are employed in this process, each corresponding to a different diffusion scale. The paper further demonstrates a novel conditional sampler that interpolates between two denoising score-based sampling methods. The model, whose architecture is based on [2] and [4], is evaluated on unconditional synthesis, super-resolution, inpainting and edge-to-image translation.

Esser *et al.* [28] propose ImageBART, a generative model which learns to revert a multinomial diffusion process on compact image representations. A transformer is used to model the reverse steps autoregressively, where the encoder's representation is obtained using the output at the previous step. ImageBART is evaluated on unconditional, class-conditional and text-conditional image generation, as well as local editing.

Gao *et al.* [117] introduce diffusion recovery likelihood, a new training procedure for energy-based models. They learn a sequence of energy-based models for the marginal distributions of the diffusion process. Thus, instead of approximating the reverse process with normal distributions, they derive the conditional distributions from the marginal energy-based models. The authors run experiments on both image generation and inpainting.

Batzolis *et al.* [24] analyze the previous score-based diffusion models on conditional image generation. Moreover, they present a new method for conditional image generation

called conditional multi-speed diffusive estimator (CMDE). This method is based on the observation that diffusing the target image and the condition image at the same rate, might be suboptimal. Therefore, they propose to diffuse the two images, which have the same drift and different diffusion rates, with an SDE. The approach is evaluated on inpainting, super-resolution and edge-to-image synthesis.

Liu *et al.* [106] introduce a framework which allows text, content and style guidance from a reference image. The core idea is to use the direction that maximizes the similarity between the representations learned for image and text. The image and text embeddings are produced by the CLIP model [151]. To address the need of training CLIP on noisy images, the authors present a self-supervised procedure that does not require text captions. The procedure uses pairs of normal and noised images to maximize the similarity between positive pairs and minimize it for negative ones (contrastive objective).

Choi *et al.* [32] propose a novel method, which does not require further training, for conditional image synthesis using unconditional diffusion models. Given a reference image, *i.e.* the condition, each sample is drawn closer to it by eliminating the low frequency content and replacing it with content from the reference image. The low pass filter is represented by a downsampling operation, which is followed by an upsampling filter of the same factor. The authors show how this method can be applied on various image-to-image translation tasks, *e.g.* paint-to-image, and editing with scribbles.

Hu *et al.* [118] propose to apply diffusion models on discrete representations given by a discrete VAE. They evaluate the idea in image generation and inpainting experiments, considering the CelebA-HQ and LSUN Church data sets.

Rombach *et al.* [10] introduce latent diffusion models, where the forward and reverse processes happen on the latent space learned by an auto-encoder. They also include cross-attention in the architecture, which brings further improvements on conditional image synthesis. The method is tested on super-resolution, image generation and inpainting.

The method introduced by Preechakul *et al.* [123] contains a semantic encoder that learns a descriptive latent space. The output of this encoder is used to condition an instance of DDIM. The proposed method allows DDPMs to perform well on tasks such as interpolation or attribute manipulation.

Chung *et al.* [26] introduce an algorithm for sampling, which reduces the number of steps required for the conditional case. Compared to the standard case, where the reverse process starts from Gaussian noise, their approach first executes one forward step to obtain an intermediary noised image and resumes the sampling from this point on. The approach is tested on inpainting, super-resolution, and magnetic resonance imaging (MRI) reconstruction.

In [120], the authors fine-tune a pretrained DDIM to generate images according to a text description. They propose a local directional CLIP loss that basically enforces the direction between the generated image and the original image to be as close as possible to the direction between the reference (original domain) and target text (target domain).

The tasks considered in the evaluation are image translation between unseen domains, and multi-attribute transfer.

Starting from the formulation of diffusion models as SDEs of Meng *et al.* [33], Khurlov *et al.* [119] investigate the latent space and the resulting encoder maps. As per Monge formulation, it is shown that these encoder maps are the optimal transport maps, but this is demonstrated only for multivariate normal distributions. The authors further support this with numerical experiments, as well as practical experiments, using the model implementation of Dhariwal *et al.* [5].

Shi *et al.* [124] start by observing how an unconditional score-based diffusion model can be formulated as a Schrödinger bridge, which can be solved using a modified version of Iterative Proportional Fitting. The previous method is reformulated to accept a condition, thus making conditional synthesis possible. Further adjustments are made to the iterative algorithm in order to optimize the time required to converge. The method is first validated with synthetic data from Kovachki *et al.* [157], showing improved capabilities in estimating the ground truth. The authors also conduct experiments on super-resolution, inpainting, and biochemical oxygen demand, the latter task being inspired by Marzouk *et al.* [158].

Inspired by the Retrieval Transformer [159], Blattmann *et al.* [116] propose a new method for training diffusion models. First, a set of similar images is fetched from a database using a nearest neighbor algorithm. The images are further encoded by an encoder with fixed parameters and projected into the CLIP [151] feature space. Finally, the reverse process of the diffusion model is conditioned on this latent space. The method can be further extended to use other conditional signals, *e.g.* text, by simply enhancing the latent space with the encoded representation of the signal.

Lyu *et al.* [122] introduce a new technique to reduce the number of sampling steps of diffusion models, boosting the performance at the same time. The idea is to stop the diffusion process at an earlier step. As the sampling cannot start from a random Gaussian noise, a GAN or VAE model is used to encode the last diffused image into a Gaussian latent space. The result is then decoded into an image which can be diffused into the starting point of the backward process.

The aim of Graikos *et al.* [40] is to separate diffusion models into two independent parts, a prior (the base part) and a constraint (the condition). This enables models to be applied on various tasks without further training. Changing the equation of DDPMs from [2] leads to independently training the model and using it in a conditional setting, given that the constraint becomes differentiable. The authors conduct experiments on conditional image synthesis and image segmentation.

3.10 Medical Image Generation and Translation

Wolleb *et al.* [41] introduce a method based on diffusion models for image segmentation within the context of brain tumor segmentation. The training consists of diffusing the segmentation map, then denoising it to obtain the original image. During the backward process, the brain MR image is concatenated into the intermediate denoising steps in order to be passed through the U-Net model, thus conditioning

the denoising process on it. Furthermore, for each input, the authors propose to generate multiple samples, which will be different due to stochasticity. Thus, the ensemble can generate a mean segmentation map and its variance (associated with the uncertainty of the map).

Song *et al.* [129] introduce a method for score-based models that is able to solve inverse problems in medical imaging, *i.e.* reconstructing images from measurements. First, an unconditional score model is trained. Then, a stochastic process of the measurement is derived, which can be used to infuse conditional information into the model via a proximal optimization step. Finally, the matrix that maps the signal to the measurement is decomposed to allow sampling in closed-form. The authors carry out multiple experiments on different medical image types, including computed tomography (CT), low-dose CT and MRI.

Within the area of medical imaging, but focusing on reconstructing the images from accelerated MRI scans, Chung *et al.* [127] propose to solve the inverse problem using a score-based diffusion model. A score model is pretrained only on magnitude images in an unconditional setting. Then, a variance exploding SDE solver [4] is employed in the sampling process. By adopting a Predictor-Corrector algorithm [4] interleaved with a data consistency mapping, the split image (real and imaginary parts) is fed through, enabling conditioning the model on the measurement. Furthermore, the authors present an extension of the method which enables conditioning on multiple coil-varying measurements.

Özbey *et al.* [128] propose a diffusion model with adversarial inference. In order to increase each diffusion step, and thus make fewer steps, inspired by [99], the authors employ a GAN model in the reverse process to estimate the denoised image at every step. Using a similar method as [149], they introduce a cycle-consistent architecture to allow training on unpaired data sets.

The aim of Hu *et al.* [126] is to remove the speckle noise in optical coherence tomography (OCT) b-scans. The first stage is represented by a method called self-fusion, as described in [160], where additional b-scans close to the given 2D slice of the input OCT volume are selected. The second stage consists of a diffusion model whose starting point is the weighted average of the original b-scan and its neighbors. Thus, the noise can be removed by sampling a clean scan.

3.11 Anomaly Detection in Medical Images

Auto-encoders are widely used for anomaly detection [161]. Since diffusion models can be seen as a particular type of VAEs, it seems natural to employ diffusion models for the same tasks as VAEs. So far, diffusion models have shown promising results in detecting anomalies in medical images, as further discussed below.

Wyatt *et al.* [46] train a DDPM on healthy medical images. The anomalies are detected at inference time by subtracting the generated image from the original image. The work also proves that using simplex noise instead of Gaussian noise yields better results for this type of task.

Wolleb *et al.* [45] propose a weakly-supervised method based on diffusion models for anomaly detection in medical images. Given two unpaired images, one healthy and one

with lesions, the former is diffused by the model. Then, the denoising process is guided by the gradient of a binary classifier in order to generate the healthy image. Finally, the sampled healthy image and the one containing lesions are subtracted to obtain the anomaly map.

Pinaya *et al.* [44] propose a diffusion-based method for detecting anomalies in brain scans, as well as segmenting those regions. The images are encoded by a VQ-VAE [152], and the quantized latent representation is obtained from a codebook. The diffusion model operates in this latent space. Averaging the intermediate samples from median steps of the backward process and then applying a precomputed threshold map, a binary mask implying the anomaly location is created. Starting the backward process from the middle, the binary mask is used to denoise the anomalous regions, while maintaining the rest. Finally, the sample at the final step is decoded, resulting in a healthy image. The segmentation map of the anomaly is obtained by subtracting the input image and the synthesized image.

Sanchez *et al.* [130] follow the same principle for detecting and segmenting anomalies in medical images: a diffusion model generates healthy samples which are then subtracted from the original images. The input image is diffused using the model, reversing the denoising equation and nullifying the condition, and then the backward conditioned process is applied. Utilizing a classifier-free model, the guidance is achieved through an attention mechanism integrated in the U-Net. The training utilizes both healthy and unhealthy examples.

3.12 Video Generation

The recent progress towards making diffusion models more efficient has enabled their application in the video domain. We next present works applying diffusion models to video generation.

Ho *et al.* [132] introduce diffusion models to the task of video generation. When compared to the 2D case, the changes are applied only to the architecture. The authors adopt the 3D U-Net from [162], presenting results in unconditional and text conditional video generation. Longer videos are generated in an autoregressive manner, where the latter video chunks are conditioned on the previous ones.

Yang *et al.* [133] generate videos frame by frame, using diffusion models. The reverse process is entirely conditioned on a context vector provided by a convolutional recurrent neural network. The authors perform an ablation study to decide if predicting the residual of the next frame returns better results than the case of predicting the actual frame. The conclusion is that the former option works better.

Höppe *et al.* [134] present random mask video diffusion (RaMViD), a method which can be used for video generation and infilling. The main contribution of their work is a novel strategy for training, which randomly splits the frames into masked and unmasked frames. The unmasked frames are used to condition the diffusion, while the masked ones are diffused by the forward process.

The work of Harvey *et al.* [131] introduces flexible diffusion models, a type of diffusion model that can be used with multiple sampling schemes for long video generation. As in [134], the authors train a diffusion model by randomly

choosing the frames used in the diffusion and those used for conditioning the process. After training the model, they investigate the effectiveness of multiple sampling schemes, concluding that the sampling choice depends on the data set.

3.13 Other Tasks

There are some pioneering works applying diffusion models to new tasks, which have been scarcely explored via diffusion modeling. We gather and discuss such contributions below.

Luo *et al.* [121] apply diffusion models on 3D point cloud generation, auto-encoding, and unsupervised representation learning. They derive an objective function from the variational lower bound of the likelihood of point clouds conditioned on a shape latent. The experiments are conducted using PointNet [163] as the underlying architecture.

Zhou *et al.* [142] introduce Point-Voxel Diffusion (PVD), a novel method for shape generation which applies diffusion on point-voxel representations. The approach addresses the tasks of shape generation and completion on the ShapeNet and PartNet data sets.

Zimmermann *et al.* [43] show a strategy to apply score-based models for classification. They add the image label as a conditioning variable to the score function and, thanks to the ODE formulation, the conditional likelihood can be computed at inference time. Thus, the prediction is the label with the maximum likelihood. Further, they study the impact of this type of classifier on out-of-distribution scenarios considering common image corruptions and adversarial perturbations.

Kim *et al.* [139] propose to solve the image registration task using diffusion models. This is achieved via two networks, a diffusion network, as per [2], and a deformation network that is based on U-Net, as described in [164]. Given two images (one static, one moving), the role of the former network is to assess the deformation between the two images, and feed the result to the latter network, which predicts the deformation fields, enabling sample generation. This method also has the ability to synthesize the deformations through the whole transition. The authors carried out experiments for different tasks, one on 2D facial expressions and one on 3D brain images. The results confirm that the model is capable of producing qualitative and accurate registration fields.

Jeanneret *et al.* [136] apply diffusion models for counterfactual explanations. The method starts from a noised query image, and generates a sample with an unconditional DDPM. With the generated sample, the gradients required for guidance are computed. Then, one step of the reversed guided process is applied. The output is further used in the next reverse steps.

Sanchez *et al.* [137] adapt the work of Dhariwal *et al.* [5] for counterfactual image generation. As in [5], the denoising process is guided by classifier gradients to generate samples from the desired counterfactual class. The key contribution is the algorithm used to retrieve the latent representation of the original image, from where the denoising process starts. Their algorithm inverts the deterministic sampling procedure of [7] and maps each original image to a unique latent representation.

Nie *et al.* [140] demonstrate how a diffusion model can be used as a defensive mechanism for adversarial attacks. Given an adversarial image, it gets diffused up until an optimally computed time step. The result is then reversed by the model, producing a purified sample at the end. To optimize the computations of solving the reverse-time SDE, the adjoint sensitivity method of Li *et al.* [165] is used for the gradient score calculations.

In the context of few-shot learning, an image generator based on diffusion models is proposed by Giannone *et al.* [135]. Given a small set of images that condition the synthesis, a visual transformer encodes these, and the resulting context representation is integrated (via two different techniques) into the U-Net model employed in the denoising process.

Wang *et al.* [141] present a framework based on diffusion models for semantic image synthesis. Leveraging the U-Net architecture of diffusion models, the input noise is supplied to the encoder, while the semantic label map is passed to the decoder using multi-layer spatially-adaptive normalization operators [166]. To further improve the sampling quality and the condition on the semantic label map, an empty map is also supplied to the sampling method to generate the unconditional noise. Then, the final noise uses both estimates.

Concerning the task of restoring images negatively affected by various weather conditions (*e.g.* snow, rain), Özdenizci *et al.* [138] demonstrate how diffusion models can be used. They condition the denoising process on the degraded image by concatenating it channel-wise to the denoised sample, at every time step. In order to deal with varying image sizes, at every step, the sample is divided into overlapping patches, passed in parallel through the model, and combined back by averaging the overlapping pixels. The employed diffusion model is based on the U-Net architecture, as presented in [2], [4], but modified to accept two concatenated images as input.

Formulating the task of image restoration as a linear inverse problem, Kawar *et al.* [27] propose the use of diffusion models. Inspired by Kawar *et al.* [167], the linear degradation matrix is decomposed using singular value decomposition, such that both the input and the output can be mapped onto the spectral space of the matrix where the diffusion process is carried out. Leveraging the pre-trained diffusion models from [2] and [5], the evaluation is conducted on various tasks: super-resolution, deblurring, colorization and inpainting.

3.14 Theoretical Contributions

Huang *et al.* [59] demonstrate how the method proposed by Song *et al.* [4] is linked with maximizing a lower bound on the marginal likelihood of the reverse SDE. Moreover, they verify their theoretical contribution with image generation experiments on CIFAR-10 and MNIST.

4 CLOSING REMARKS AND FUTURE DIRECTIONS

In this paper, we reviewed the advancements made by the research community in developing and applying diffusion models to various computer vision tasks. We identified

three primary formulations of diffusion modeling based on: DDPMs, NCSNs, and SDEs. Each formulation obtains remarkable results in image generation, surpassing GANs while increasing the diversity of the generated samples. The outstanding results of diffusion models are achieved while the research is still in its early phase. Although we observed that the main focus is on conditional and unconditional image generation, there are still many tasks to be explored and further improvements to be realized.

Limitations. The most significant disadvantage of diffusion models remains the need to perform multiple steps at inference time to generate only one sample. Despite the important amount of research conducted in this direction, GANs are still faster at producing images. Other issues of diffusion models can be linked to the commonly used strategy to employ CLIP embeddings for text-to-image generation. For example, Ramesh *et al.* [112] highlight that their model struggles to generate readable text in an image and motivate the behavior by stating that CLIP embeddings do not contain information about spelling. Therefore, when such embeddings are used for conditioning the denoising process, the model can inherit this kind of issues.

Future directions. To reduce the uncertainty level, diffusion models generally avoid taking large steps during sampling. Indeed, taking small steps ensures the data sample generated at each step is explained by the learned Gaussian distribution. A similar behavior is observed when applying gradient descent to optimize neural networks. Indeed, taking a large step in the negative direction of the gradient, *i.e.* using a very large learning rate, can lead to updating the model to a region with high uncertainty, having no control over the loss value. In future work, transferring update rules borrowed from efficient optimizers to diffusion models could perhaps lead to a more efficient sampling (generation) process.

Aside from the current tendency of researching more efficient diffusion models, future work can study diffusion models applied in other computer vision tasks, such as image dehazing, video anomaly detection, or visual question answering. Even if we found some works studying anomaly detection in medical images [44]–[46], this task could also be explored in other domains, such as video surveillance or industrial inspection.

An interesting research direction is to assess the quality and utility of the representation space learned by diffusion models in discriminative tasks. This could be carried out in at least two distinct ways. In a direct way, by learning some discriminative model on top of the latent representations provided by a denoising model, to address some classification or regression task. In an indirect way, by augmenting training sets with realistic samples generated by diffusion models. The latter direction might be more suitable for tasks such as object detection, where inpainting diffusion models could do a good job at blending in new objects in images.

Another future work direction is to employ conditional diffusion models to simulate possible futures in video. The generated videos could further be given as input to reinforcement learning models.

Recent diffusion models [132] have shown impressive text-to-video synthesis capabilities compared to the previous state of the art, significantly reducing the number of

artifacts and reaching an unprecedented generative performance. However, we believe this direction requires more attention in future work, as the generated videos are rather short. Hence, modeling long-term temporal relations and interactions between objects remains an open challenge.

In future, the research on diffusion models can also be expanded towards learning multi-purpose models that solve multiple tasks at once. Creating a diffusion model to generate multiple types of outputs, while being conditioned on various types of data, *e.g.* text, class labels or images, might take us closer to understanding the necessary steps towards developing artificial general intelligence (AGI).

ACKNOWLEDGMENTS

This work was supported by a grant of the Romanian Ministry of Education and Research, CNCS - UEFISCDI, project no. PN-III-P2-2.1-PED-2021-0195, contract no. 690/2022, within PNCDI III.

REFERENCES

- [1] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using non-equilibrium thermodynamics," in *Proceedings of ICML*, pp. 2256–2265, 2015.
- [2] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Proceedings of NeurIPS*, vol. 33, pp. 6840–6851, 2020.
- [3] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," in *Proceedings of NeurIPS*, vol. 32, pp. 11918–11930, 2019.
- [4] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-Based Generative Modeling through Stochastic Differential Equations," in *Proceedings of ICLR*, 2021.
- [5] P. Dhariwal and A. Nichol, "Diffusion models beat GANs on image synthesis," in *Proceedings of NeurIPS*, vol. 34, pp. 8780–8794, 2021.
- [6] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," in *Proceedings of ICML*, pp. 8162–8171, 2021.
- [7] J. Song, C. Meng, and S. Ermon, "Denoising Diffusion Implicit Models," in *Proceedings of ICLR*, 2021.
- [8] D. Watson, W. Chan, J. Ho, and M. Norouzi, "Learning fast samplers for diffusion models by differentiating through sample quality," in *Proceedings of ICLR*, 2021.
- [9] J. Shi, C. Wu, J. Liang, X. Liu, and N. Duan, "DiVAE: Photorealistic Images Synthesis with Denoising Diffusion Decoder," *arXiv preprint arXiv:2206.00386*, 2022.
- [10] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-Resolution Image Synthesis with Latent Diffusion Models," in *Proceedings of CVPR*, pp. 10684–10695, 2022.
- [11] R. Rombach, A. Blattmann, and B. Ommer, "Text-Guided Synthesis of Artistic Images with Retrieval-Augmented Diffusion Models," *arXiv preprint arXiv:2207.13038*, 2022.
- [12] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, B. K. Ayan, S. S. Mahdavi, R. G. Lopes, *et al.*, "Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding," *arXiv preprint arXiv:2205.11487*, 2022.
- [13] Y. Song and S. Ermon, "Improved techniques for training score-based generative models," in *Proceedings of NeurIPS*, vol. 33, pp. 12438–12448, 2020.
- [14] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen, "GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models," in *Proceedings of ICML*, pp. 16784–16804, 2021.
- [15] Y. Song, C. Durkan, I. Murray, and S. Ermon, "Maximum likelihood training of score-based diffusion models," in *Proceedings of NeurIPS*, vol. 34, pp. 1415–1428, 2021.
- [16] A. Sinha, J. Song, C. Meng, and S. Ermon, "D2C: Diffusion-decoding models for few-shot conditional generation," in *Proceedings of NeurIPS*, vol. 34, pp. 12533–12548, 2021.
- [17] A. Vahdat, K. Kreis, and J. Kautz, "Score-based generative modeling in latent space," in *Proceedings of NeurIPS*, vol. 34, pp. 11287–11302, 2021.
- [18] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi, "Image super-resolution via iterative refinement," *arXiv preprint arXiv:2104.07636*, 2021.
- [19] K. Pandey, A. Mukherjee, P. Rai, and A. Kumar, "VAEs meet diffusion models: Efficient and high-fidelity generation," in *Proceedings of NeurIPS Workshop on DGMs and Applications*, 2021.
- [20] F. Bao, C. Li, J. Zhu, and B. Zhang, "Analytic-DPM: an Analytic Estimate of the Optimal Reverse Variance in Diffusion Probabilistic Models," in *Proceedings of ICLR*, 2022.
- [21] T. Dockhorn, A. Vahdat, and K. Kreis, "Score-based generative modeling with critically-damped Langevin diffusion," in *Proceedings of ICLR*, 2022.
- [22] N. Liu, S. Li, Y. Du, A. Torralba, and J. B. Tenenbaum, "Compositional Visual Generation with Composable Diffusion Models," in *Proceedings of ECCV*, 2022.
- [23] Y. Jiang, S. Yang, H. Qiu, W. Wu, C. C. Loy, and Z. Liu, "Text2Human: Text-Driven Controllable Human Image Generation," *ACM Transactions on Graphics*, vol. 41, no. 4, pp. 1–11, 2022.
- [24] G. Batzolis, J. Stanczuk, C.-B. Schönlieb, and C. Etmann, "Conditional image generation with score-based diffusion models," *arXiv preprint arXiv:2111.13606*, 2021.
- [25] M. Daniels, T. Maunu, and P. Hand, "Score-based generative neural networks for large-scale optimal transport," in *Proceedings of NeurIPS*, pp. 12955–12965, 2021.
- [26] H. Chung, B. Sim, and J. C. Ye, "Come-Closer-Diffuse-Faster: Accelerating Conditional Diffusion Models for Inverse Problems through Stochastic Contraction," in *Proceedings of CVPR*, pp. 12413–12422, 2022.
- [27] B. Kavar, M. Elad, S. Ermon, and J. Song, "Denoising diffusion restoration models," in *Proceedings of DGM4HSD*, 2022.
- [28] P. Esser, R. Rombach, A. Blattmann, and B. Ommer, "ImageBART: Bidirectional Context with Multinomial Diffusion for Autoregressive Image Synthesis," in *Proceedings of NeurIPS*, vol. 34, pp. 3518–3532, 2021.
- [29] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. Van Gool, "RePaint: Inpainting using Denoising Diffusion Probabilistic Models," in *Proceedings of CVPR*, pp. 11461–11471, 2022.
- [30] B. Jing, G. Corso, R. Berlinghieri, and T. Jaakkola, "Subspace diffusion generative models," *arXiv preprint arXiv:2205.01490*, 2022.
- [31] O. Avrahami, D. Lischinski, and O. Fried, "Blended diffusion for text-driven editing of natural images," in *Proceedings of CVPR*, pp. 18208–18218, 2022.
- [32] J. Choi, S. Kim, Y. Jeong, Y. Gwon, and S. Yoon, "ILVR: Conditioning Method for Denoising Diffusion Probabilistic Models," in *Proceedings of ICCV*, pp. 14347–14356, 2021.
- [33] C. Meng, Y. Song, J. Song, J. Wu, J.-Y. Zhu, and S. Ermon, "SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations," in *Proceedings of ICLR*, 2021.
- [34] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi, "Palette: Image-to-image diffusion models," in *Proceedings of SIGGRAPH*, pp. 1–10, 2022.
- [35] M. Zhao, F. Bao, C. Li, and J. Zhu, "EGSDE: Unpaired Image-to-Image Translation via Energy-Guided Stochastic Differential Equations," *arXiv preprint arXiv:2207.06635*, 2022.
- [36] T. Wang, T. Zhang, B. Zhang, H. Ouyang, D. Chen, Q. Chen, and F. Wen, "Pretraining is All You Need for Image-to-Image Translation," *arXiv preprint arXiv:2205.12952*, 2022.
- [37] B. Li, K. Xue, B. Liu, and Y.-K. Lai, "VQBB: Image-to-image Translation with Vector Quantized Brownian Bridge," *arXiv preprint arXiv:2205.07680*, 2022.
- [38] J. Wolleb, R. Sandkühler, F. Bieder, and P. C. Cattin, "The Swiss Army Knife for Image-to-Image Translation: Multi-Task Diffusion Models," *arXiv preprint arXiv:2204.02641*, 2022.
- [39] D. Baranchuk, I. Rubachev, A. Voynov, V. Khurlov, and A. Babenko, "Label-Efficient Semantic Segmentation with Diffusion Models," in *Proceedings of ICLR*, 2022.
- [40] A. Graikos, N. Malkin, N. Jojic, and D. Samaras, "Diffusion models as plug-and-play priors," *arXiv preprint arXiv:2206.09012*, 2022.
- [41] J. Wolleb, R. Sandkühler, F. Bieder, P. Valmaggia, and P. C. Cattin, "Diffusion Models for Implicit Image Segmentation Ensembles," in *Proceedings of MIDL*, 2022.
- [42] T. Amit, E. Nachmani, T. Shaharabany, and L. Wolf, "SegDiff: Image Segmentation with Diffusion Probabilistic Models," *arXiv preprint arXiv:2112.00390*, 2021.

- [43] R. S. Zimmermann, L. Schott, Y. Song, B. A. Dunn, and D. A. Klindt, "Score-based generative classifiers," in *Proceedings of NeurIPS Workshop on DGMs and Applications*, 2021.
- [44] W. H. Pinaya, M. S. Graham, R. Gray, P. F. Da Costa, P.-D. Tudosiu, P. Wright, Y. H. Mah, A. D. MacKinnon, J. T. Teo, R. Jager, et al., "Fast Unsupervised Brain Anomaly Detection and Segmentation with Diffusion Models," *arXiv preprint arXiv:2206.03461*, 2022.
- [45] J. Wolleb, F. Bieder, R. Sandkühler, and P. C. Cattin, "Diffusion Models for Medical Anomaly Detection," *arXiv preprint arXiv:2203.04306*, 2022.
- [46] J. Wyatt, A. Leach, S. M. Schmon, and C. G. Willcocks, "AnoD-DPM: Anomaly Detection With Denoising Diffusion Probabilistic Models Using Simplex Noise," in *Proceedings of CVPRW*, pp. 650–656, 2022.
- [47] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.
- [48] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [49] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [50] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in *Proceedings of ICLR*, 2014.
- [51] I. Higgins, L. Matthey, A. Pal, C. P. Burgess, X. Glorot, M. M. Botvinick, S. Mohamed, and A. Lerchner, "beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework," in *Proceedings of ICLR*, 2017.
- [52] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proceedings of NIPS*, pp. 2672–2680, 2014.
- [53] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," in *Proceedings of NeurIPS*, vol. 33, pp. 9912–9924, 2020.
- [54] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *Proceedings of ICML*, vol. 119, pp. 1597–1607, 2020.
- [55] F.-A. Croitoru, D.-N. Grigore, and R. T. Ionescu, "Discriminability-enforcing loss to improve representation learning," in *Proceedings of CVPRW*, pp. 2598–2602, 2022.
- [56] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [57] S. Laine and T. Aila, "Temporal ensembling for semi-supervised learning," in *Proceedings of ICLR*, 2017.
- [58] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Proceedings of NIPS*, vol. 30, pp. 1195–1204, 2017.
- [59] C.-W. Huang, J. H. Lim, and A. C. Courville, "A variational perspective on diffusion-based generative models and score matching," in *Proceedings of NeurIPS*, vol. 34, pp. 22863–22876, 2021.
- [60] Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, and F. J. Huang, "A tutorial on energy-based learning," in *Predicting Structured Data*, MIT Press, 2006.
- [61] J. Ngiam, Z. Chen, P. W. Koh, and A. Y. Ng, "Learning Deep Energy Models," in *Proceedings of ICML*, pp. 1105–1112, 2011.
- [62] A. van den Oord, N. Kalchbrenner, L. Espeholt, K. Kavukcuoglu, O. Vinyals, and A. Graves, "Conditional Image Generation with PixelCNN Decoders," in *Proceedings of NeurIPS*, vol. 29, pp. 4797–4805, 2016.
- [63] L. Dinh, D. Krueger, and Y. Bengio, "NICE: Non-linear Independent Components Estimation," in *Proceedings of ICLR*, 2015.
- [64] L. Dinh, J. Sohl-Dickstein, and S. Bengio, "Density estimation using Real NVP," in *Proceedings of ICLR*, 2017.
- [65] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Proceedings of MICCAI*, pp. 234–241, 2015.
- [66] W. Feller, "On the Theory of Stochastic Processes, with Particular Reference to Applications," in *First Berkeley Symposium on Mathematical Statistics and Probability*, pp. 403–432, 1949.
- [67] P. Vincent, "A Connection Between Score Matching and Denoising Autoencoders," *Neural Computation*, vol. 23, pp. 1661–1674, 2011.
- [68] Y. Song, S. Garg, J. Shi, and S. Ermon, "Sliced Score Matching: A Scalable Approach to Density and Score Estimation," in *Proceedings of UAI*, p. 204, 2019.
- [69] B. D. Anderson, "Reverse-time diffusion equation models," *Stochastic Processes and their Applications*, vol. 12, no. 3, pp. 313–326, 1982.
- [70] T. Salimans, A. Karpathy, X. Chen, and D. P. Kingma, "PixelCNN++: Improving the PixelCNN with Discretized Logistic Mixture Likelihood and Other Modifications," in *Proceedings of ICLR*, 2017.
- [71] Q. Zhang and Y. Chen, "Diffusion normalizing flow," in *Proceedings of NeurIPS*, vol. 34, pp. 16280–16291, 2021.
- [72] K. Swersky, M. Ranzato, D. Buchman, B. M. Marlin, and N. Freitas, "On Autoencoders and Score Matching for Energy Based Models," in *Proceedings of ICML*, pp. 1201–1208, 2011.
- [73] F. Bao, K. Xu, C. Li, L. Hong, J. Zhu, and B. Zhang, "Variational (gradient) estimate of the score function in energy-based latent variable models," in *Proceedings of ICML*, pp. 651–661, 2021.
- [74] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved Techniques for Training GANs," in *Proceedings of NeurIPS*, pp. 2234–2242, 2016.
- [75] Y. Shen, J. Gu, X. Tang, and B. Zhou, "Interpreting the Latent Space of GANs for Semantic Face Editing," in *Proceedings of CVPR*, pp. 9240–9249, 2020.
- [76] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in *Proceedings of ICLR*, 2016.
- [77] J. Ho and T. Salimans, "Classifier-Free Diffusion Guidance," in *Proceedings of NeurIPS Workshop on DGMs and Applications*, 2021.
- [78] J. Austin, D. D. Johnson, J. Ho, D. Tarlow, and R. van den Berg, "Structured denoising diffusion models in discrete state-spaces," in *Proceedings of NeurIPS*, vol. 34, pp. 17981–17993, 2021.
- [79] Y. Benny and L. Wolf, "Dynamic Dual-Output Diffusion Models," in *Proceedings of CVPR*, pp. 11482–11491, 2022.
- [80] S. Bond-Taylor, P. Hesse, H. Sasaki, T. P. Breckon, and C. G. Willcocks, "Unleashing Transformers: Parallel Token Prediction with Discrete Absorbing Diffusion for Fast High-Resolution Image Generation from Vector-Quantized Codes," in *Proceedings of ECCV*, 2022.
- [81] J. Choi, J. Lee, C. Shin, S. Kim, H. Kim, and S. Yoon, "Perception Prioritized Training of Diffusion Models," in *Proceedings of CVPR*, pp. 11472–11481, 2022.
- [82] V. De Bortoli, J. Thornton, J. Heng, and A. Doucet, "Diffusion Schrödinger bridge with applications to score-based generative modeling," in *Proceedings of NeurIPS*, vol. 34, pp. 17695–17709, 2021.
- [83] J. Deasy, N. Simidjievski, and P. Liò, "Heavy-tailed denoising score matching," *arXiv preprint arXiv:2112.09788*, 2021.
- [84] K. Deja, A. Kuzina, T. Trzcinski, and J. M. Tomczak, "On Analyzing Generative and Denoising Capabilities of Diffusion-Based Deep Generative Models," *arXiv preprint arXiv:2206.00070*, 2022.
- [85] A. Jolicœur-Martineau, R. Piché-Taillefer, I. Mitliagkas, and R. T. des Combes, "Adversarial score matching and improved sampling for image generation," in *Proceedings of ICLR*, 2021.
- [86] A. Jolicœur-Martineau, K. Li, R. Piché-Taillefer, T. Kachman, and I. Mitliagkas, "Gotta go fast when generating data with score-based models," *arXiv preprint arXiv:2105.14080*, 2021.
- [87] D. Kim, B. Na, S. J. Kwon, D. Lee, W. Kang, and I.-C. Moon, "Maximum Likelihood Training of Implicit Nonlinear Diffusion Models," *arXiv preprint arXiv:2205.13699*, 2022.
- [88] D. Kingma, T. Salimans, B. Poole, and J. Ho, "Variational diffusion models," in *Proceedings of NeurIPS*, vol. 34, pp. 21696–21707, 2021.
- [89] Z. Kong and W. Ping, "On Fast Sampling of Diffusion Probabilistic Models," in *Proceedings of INNF+*, 2021.
- [90] M. W. Lam, J. Wang, R. Huang, D. Su, and D. Yu, "Bilateral denoising diffusion models," *arXiv preprint arXiv:2108.11514*, 2021.
- [91] L. Liu, Y. Ren, Z. Lin, and Z. Zhao, "Pseudo Numerical Methods for Diffusion Models on Manifolds," in *Proceedings of ICLR*, 2022.
- [92] H. Ma, L. Zhang, X. Zhu, J. Zhang, and J. Feng, "Accelerating Score-Based Generative Models for High-Resolution Image Synthesis," *arXiv preprint arXiv:2206.04029*, 2022.
- [93] E. Nachmani, R. S. Roman, and L. Wolf, "Non-Gaussian denoising diffusion models," *arXiv preprint arXiv:2106.07582*, 2021.

- [94] R. San-Roman, E. Nachmani, and L. Wolf, "Noise estimation for generative diffusion models," *arXiv preprint arXiv:2104.02600*, 2021.
- [95] V. Sehwag, C. Hazirbas, A. Gordo, F. Ozgenel, and C. Canton, "Generating High Fidelity Data from Low-Density Regions Using Diffusion Models," in *Proceedings of CVPR*, pp. 11492–11501, 2022.
- [96] G. Wang, Y. Jiao, Q. Xu, Y. Wang, and C. Yang, "Deep generative learning via Schrödinger bridge," in *Proceedings of ICML*, pp. 10794–10804, 2021.
- [97] Z. Wang, H. Zheng, P. He, W. Chen, and M. Zhou, "Diffusion-GAN: Training GANs with Diffusion," *arXiv preprint arXiv:2206.02262*, 2022.
- [98] D. Watson, J. Ho, M. Norouzi, and W. Chan, "Learning to efficiently sample from diffusion probabilistic models," *arXiv preprint arXiv:2106.03802*, 2021.
- [99] Z. Xiao, K. Kreis, and A. Vahdat, "Tackling the generative learning trilemma with denoising diffusion GANs," in *Proceedings of ICLR*, 2022.
- [100] H. Zheng, P. He, W. Chen, and M. Zhou, "Truncated diffusion probabilistic models," *arXiv preprint arXiv:2202.09671*, 2022.
- [101] F. Bordes, R. Balestriero, and P. Vincent, "High fidelity visualization of what your self-supervised representation knows about," *Transactions on Machine Learning Research*, 2022.
- [102] A. Campbell, J. Benton, V. De Bortoli, T. Rainforth, G. Deligiannidis, and A. Doucet, "A Continuous Time Framework for Discrete Denoising Models," *arXiv preprint arXiv:2205.14987*, 2022.
- [103] C.-H. Chao, W.-F. Sun, B.-W. Cheng, Y.-C. Lo, C.-C. Chang, Y.-L. Liu, Y.-L. Chang, C.-P. Chen, and C.-Y. Lee, "Denoising Likelihood Score Matching for Conditional Score-Based Data Generation," in *Proceedings of ICLR*, 2022.
- [104] J. Ho, C. Saharia, W. Chan, D. J. Fleet, M. Norouzi, and T. Salimans, "Cascaded Diffusion Models for High Fidelity Image Generation," *Journal of Machine Learning Research*, vol. 23, no. 47, pp. 1–33, 2022.
- [105] T. Karras, M. Aittala, T. Aila, and S. Laine, "Elucidating the Design Space of Diffusion-Based Generative Models," *arXiv preprint arXiv:2206.00364*, 2022.
- [106] X. Liu, D. H. Park, S. Azadi, G. Zhang, A. Chopikyan, Y. Hu, H. Shi, A. Rohrbach, and T. Darrell, "More control for free! Image synthesis with semantic diffusion guidance," *arXiv preprint arXiv:2112.05744*, 2021.
- [107] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, "DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps," *arXiv preprint arXiv:2206.00927*, 2022.
- [108] T. Salimans and J. Ho, "Progressive distillation for fast sampling of diffusion models," in *Proceedings of ICLR*, 2022.
- [109] V. Singh, S. Jandial, A. Chopra, S. Ramesh, B. Krishnamurthy, and V. N. Balasubramanian, "On Conditioning the Input Noise for Controlled Image Generation with Diffusion Models," *arXiv preprint arXiv:2205.03859*, 2022.
- [110] H. Sasaki, C. G. Willcocks, and T. P. Breckon, "UNIT-DDPM: Unpaired Image Translation with Denoising Diffusion Probabilistic Models," *arXiv preprint arXiv:2104.05358*, 2021.
- [111] S. Gu, D. Chen, J. Bao, F. Wen, B. Zhang, D. Chen, L. Yuan, and B. Guo, "Vector quantized diffusion model for text-to-image synthesis," in *Proceedings of CVPR*, pp. 10696–10706, 2022.
- [112] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with CLIP latents," *arXiv preprint arXiv:2204.06125*, 2022.
- [113] Q. Zhang and Y. Chen, "Fast Sampling of Diffusion Models with Exponential Integrator," *arXiv preprint arXiv:2204.13902*, 2022.
- [114] O. Avrahami, O. Fried, and D. Lischinski, "Blended latent diffusion," *arXiv preprint arXiv:2206.02779*, 2022.
- [115] G. Batzolis, J. Stanczuk, C.-B. Schönlieb, and C. Etmann, "Non-Uniform Diffusion Models," *arXiv preprint arXiv:2207.09786*, 2022.
- [116] A. Blattmann, R. Rombach, K. Oktay, and B. Ommer, "Retrieval-Augmented Diffusion Models," *arXiv preprint arXiv:2204.11824*, 2022.
- [117] R. Gao, Y. Song, B. Poole, Y. N. Wu, and D. P. Kingma, "Learning Energy-Based Models by Diffusion Recovery Likelihood," in *Proceedings of ICLR*, 2021.
- [118] M. Hu, Y. Wang, T.-J. Cham, J. Yang, and P. N. Suganthan, "Global Context with Discrete Diffusion in Vector Quantised Modelling for Image Generation," in *Proceedings of CVPR*, pp. 11502–11511, 2022.
- [119] V. Khruikov and I. Oseledets, "Understanding DDPM Latent Codes Through Optimal Transport," *arXiv preprint arXiv:2202.07477*, 2022.
- [120] G. Kim, T. Kwon, and J. C. Ye, "DiffusionCLIP: Text-Guided Diffusion Models for Robust Image Manipulation," in *Proceedings of CVPR*, pp. 2426–2435, 2022.
- [121] S. Luo and W. Hu, "Diffusion probabilistic models for 3D point cloud generation," in *Proceedings of CVPR*, pp. 2837–2845, 2021.
- [122] Z. Lyu, X. Xu, C. Yang, D. Lin, and B. Dai, "Accelerating Diffusion Models via Early Stop of the Diffusion Process," *arXiv preprint arXiv:2205.12524*, 2022.
- [123] K. Preechakul, N. Chatthee, S. Wizadwongsa, and S. Suwanajakorn, "Diffusion Autoencoders: Toward a Meaningful and Decodable Representation," in *Proceedings of CVPR*, pp. 10619–10629, 2022.
- [124] Y. Shi, V. De Bortoli, G. Deligiannidis, and A. Doucet, "Conditional Simulation Using Diffusion Schrödinger Bridges," in *Proceedings of UAI*, 2022.
- [125] Z. Kadhodaie and E. P. Simoncelli, "Stochastic solutions for linear inverse problems using the prior implicit in a denoiser," in *Proceedings of NeurIPS*, vol. 34, pp. 13242–13254, 2021.
- [126] D. Hu, Y. K. Tao, and I. Oguz, "Unsupervised denoising of retinal OCT with diffusion probabilistic model," in *Proceedings of SPIE Medical Imaging*, vol. 12032, pp. 25–34, 2022.
- [127] H. Chung and J. C. Ye, "Score-based diffusion models for accelerated MRI," *Medical Image Analysis*, vol. 80, p. 102479, 2022.
- [128] M. Özbey, S. U. Dar, H. A. Bedel, O. Dalmaz, Ş. Öztürk, A. Güngör, and T. Çukur, "Unsupervised Medical Image Translation with Adversarial Diffusion Models," *arXiv preprint arXiv:2207.08208*, 2022.
- [129] Y. Song, L. Shen, L. Xing, and S. Ermon, "Solving inverse problems in medical imaging with score-based generative models," in *Proceedings of ICLR*, 2022.
- [130] P. Sanchez, A. Kascenas, X. Liu, A. Q. O'Neil, and S. A. Tsafaris, "What is healthy? generative counterfactual diffusion for lesion localization," in *Proceedings of DGM4MICCAI*, 2022.
- [131] W. Harvey, S. Naderiparizi, V. Masrani, C. Weilbach, and F. Wood, "Flexible Diffusion Modeling of Long Videos," *arXiv preprint arXiv:2205.11495*, 2022.
- [132] J. Ho, T. Salimans, A. A. Gritsenko, W. Chan, M. Norouzi, and D. J. Fleet, "Video diffusion models," in *Proceedings of DGM4HSD*, 2022.
- [133] R. Yang, P. Srivastava, and S. Mandt, "Diffusion Probabilistic Modeling for Video Generation," *arXiv preprint arXiv:2203.09481*, 2022.
- [134] T. Höppe, A. Mehrjou, S. Bauer, D. Nielsen, and A. Dittadi, "Diffusion Models for Video Prediction and Infilling," *arXiv preprint arXiv:2206.07696*, 2022.
- [135] G. Giannone, D. Nielsen, and O. Winther, "Few-Shot Diffusion Models," *arXiv preprint arXiv:2205.15463*, 2022.
- [136] G. Jeanneret, L. Simon, and F. Jurie, "Diffusion Models for Counterfactual Explanations," *arXiv preprint arXiv:2203.15636*, 2022.
- [137] P. Sanchez and S. A. Tsafaris, "Diffusion causal models for counterfactual estimation," in *Proceedings of CLeaR*, vol. 140, pp. 1–21, 2022.
- [138] O. Özdenizci and R. Legenstein, "Restoring Vision in Adverse Weather Conditions with Patch-Based Denoising Diffusion Models," *arXiv preprint arXiv:2207.14626*, 2022.
- [139] B. Kim, I. Han, and J. C. Ye, "DiffuseMorph: Unsupervised Deformable Image Registration Along Continuous Trajectory Using Diffusion Models," *arXiv preprint arXiv:2112.05149*, 2021.
- [140] W. Nie, B. Guo, Y. Huang, C. Xiao, A. Vahdat, and A. Anandkumar, "Diffusion models for adversarial purification," in *Proceedings of ICML*, 2022.
- [141] W. Wang, J. Bao, W. Zhou, D. Chen, D. Chen, L. Yuan, and H. Li, "Semantic image synthesis via diffusion models," *arXiv preprint arXiv:2207.00050*, 2022.
- [142] L. Zhou, Y. Du, and J. Wu, "3D shape generation and completion through point-voxel diffusion," in *Proceedings of ICCV*, pp. 5826–5835, 2021.
- [143] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proceedings of NIPS*, vol. 30, pp. 6000–6010, 2017.
- [144] M. Shannon, B. Poole, S. Mariooryad, T. Bagby, E. Battenberg, D. Kao, D. Stanton, and R. Skerry-Ryan, "Non-saturating GAN training as divergence minimization," *arXiv preprint arXiv:2010.08029*, 2020.

- [145] M. Arjovsky and L. Bottou, "Towards principled methods for training generative adversarial networks," in *Proceedings of ICLR*, 2017.
- [146] C. K. Sønderby, J. Caballero, L. Theis, W. Shi, and F. Huszár, "Amortised map inference for image super-resolution," in *Proceedings of ICLR*, 2017.
- [147] J. Geiping, H. Bauermeister, H. Dröge, and M. Moeller, "Inverting gradients – How easy is it to break privacy in federated learning?," in *Proceedings of NeurIPS*, vol. 33, pp. 16937–16947, 2020.
- [148] H. Tachibana, M. Go, M. Inahara, Y. Katayama, and Y. Watanabe, "Itô-Taylor Sampling Scheme for Denoising Diffusion Probabilistic Models using Ideal Derivatives," *arXiv preprint arXiv:2112.13339*, 2021.
- [149] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of ICCV*, pp. 2223–2232, 2017.
- [150] P. Esser, R. Rombach, and B. Ommer, "Taming transformers for high-resolution image synthesis," in *Proceedings of CVPR*, pp. 12873–12883, 2021.
- [151] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," in *Proceedings of ICML*, vol. 139, pp. 8748–8763, 2021.
- [152] A. Van Den Oord, O. Vinyals, and K. Kavukcuoglu, "Neural discrete representation learning," *Proceedings of NIPS*, vol. 30, pp. 6309–6318, 2017.
- [153] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," in *Proceedings of EMNLP*, pp. 3982–3992, 2019.
- [154] X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, and C. Change Loy, "ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks," in *Proceedings of ECCVW*, pp. 63–79, 2018.
- [155] G. Lin, A. Milan, C. Shen, and I. Reid, "RefineNet: Multi-Path Refinement Networks for High-Resolution Semantic Segmentation," in *Proceeding of CVPR*, pp. 5168–5177, 2017.
- [156] K. Miyasawa, "An empirical Bayes estimator of the mean of a normal population," *Bulletin of the International Statistical Institute*, vol. 38, pp. 181–188, 1961.
- [157] N. Kovachki, R. Baptista, B. Hosseini, and Y. Marzouk, "Conditional sampling with monotone GANs," *arXiv preprint arXiv:2006.06755*, 2020.
- [158] Y. Marzouk, T. Moselhy, M. Parno, and A. Spantini, "Sampling via Measure Transport: An Introduction," in *Handbook of Uncertainty Quantification*, (Cham), pp. 1–41, Springer, 2016.
- [159] S. Borgeaud, A. Mensch, J. Hoffmann, T. Cai, E. Rutherford, K. Millican, G. B. Van Den Driessche, J.-B. Lespiau, B. Damoc, A. Clark, D. De Las Casas, A. Guy, J. Menick, R. Ring, T. Hennigan, S. Huang, L. Maggiore, C. Jones, A. Cassirer, A. Brock, M. Paganini, G. Irving, O. Vinyals, S. Osindero, K. Simonyan, J. Rae, E. Elsen, and L. Sifre, "Improving Language Models by Retrieving from Trillions of Tokens," in *Proceedings of ICML*, vol. 162, pp. 2206–2240, 2022.
- [160] I. Oguz, J. D. Malone, Y. Atay, and Y. K. Tao, "Self-fusion for OCT noise reduction," in *Proceedings of SPIE Medical Imaging*, vol. 11313, p. 113130C, SPIE, 2020.
- [161] R. T. Ionescu, F. S. Khan, M.-I. Georgescu, and L. Shao, "Object-Centric Auto-Encoders and Dummy Anomalies for Abnormal Event Detection in Video," in *Proceedings of CVPR*, pp. 7842–7851, 2019.
- [162] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation," in *Proceedings of MICCAI*, pp. 424–432, 2016.
- [163] R. Q. Charles, H. Su, M. Kaichun, and L. J. Guibas, "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation," in *Proceedings of CVPR*, pp. 77–85, 2017.
- [164] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, "An unsupervised learning model for deformable medical image registration," in *Proceedings of CVPR*, pp. 9252–9260, 2018.
- [165] X. Li, T.-K. L. Wong, R. T. Chen, and D. Duvenaud, "Scalable gradients for stochastic differential equations," in *Proceedings of AISTATS*, pp. 3870–3882, 2020.
- [166] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic image synthesis with spatially-adaptive normalization," in *Proceedings of CVPR*, pp. 2337–2346, 2019.
- [167] B. Kavar, G. Vaksman, and M. Elad, "SNIPS: Solving noisy inverse problems stochastically," in *Proceedings of NeurIPS*, vol. 34, pp. 21757–21769, 2021.



learning.

Florinel-Alin Croitoru is a Ph.D. student at the University of Bucharest, Romania. He obtained his bachelor's degree from the Faculty of Mathematics and Computer Science of the University of Bucharest in 2019. In 2021, he obtained his masters degree in Artificial Intelligence with a thesis on action spotting in football videos. His domains of interest include machine learning, computer vision and deep



Vlad Hondru is a Ph.D. student at the University of Bucharest, Romania. He obtained his bachelor's degree from the University of Manchester in Mechatronic Engineering, then he graduated from Imperial College London, studying towards an MSc in Computing Science, with a Visual Computing and Robotics specialization, focusing on Artificial Intelligence. He did a year-long placement at Rolls-Royce as a software engineer, as well as undertaking a summer internship within the Robotics Group of the University of Manchester. He currently works as a machine learning engineer, developing NLP products.



Radu Ionescu is professor at the University of Bucharest, Romania. He completed his Ph.D. at the University of Bucharest in 2013, receiving the 2014 Award for Outstanding Doctoral Research from the Romanian Ad Astra Association. His research interests include machine learning, computer vision, image processing, computational linguistics and medical imaging. He published over 100 articles at international venues (including CVPR, NeurIPS, ICCV, ACL, EMNLP, NAACL, TPAMI, IJCV, CVIU), and a research monograph with Springer. Radu received the "Caianiello Best Young Paper Award" at ICIAP 2013. Radu also received the 2017 "Young Researchers in Science and Engineering" Prize for young Romanian researchers and the "Danubius Young Scientist Award 2018 for Romania".



Mubarak Shah is the UCF Trustee chair professor and the founding director of the Center for Research in Computer Vision at the University of Central Florida (UCF). He is a fellow of the NAI, IEEE, AAAS, IAPR and SPIE. He is an editor of an international book series on video computing, was editor-in-chief of Machine Vision and Applications and an associate editor of ACM Computing Surveys and IEEE TPAMI. His research interests include video surveillance, visual tracking, human activity recognition, visual analysis of crowded scenes, video registration, UAV video analysis, among others. He has served as an ACM distinguished speaker and IEEE distinguished visitor speaker. He is a recipient of ACM SIGMM Technical Achievement award; IEEE Outstanding Engineering Educator Award; Harris Corporation Engineering Achievement Award; an honorable mention for the ICCV 2005 "Where Am I?" Challenge Problem; 2013 NGA Best Research Poster Presentation; 2nd place in Grand Challenge at the ACM Multimedia 2013 conference; and runner up for the best paper award in ACM Multimedia Conference in 2005 and 2010. At UCF, he has received the Pegasus Professor Award, University Distinguished Research Award, Faculty Excellence in Mentoring Doctoral Students, Scholarship of Teaching and Learning Award, Teaching Incentive Program Award, Research Incentive Award.

APPENDIX A

VARIATIONAL BOUND.

We emphasize that the derivation presented below is also shown in [1], [2]. The variational bound of the log density of the data can be derived as in the case of VAEs [50], where the latent variables are the noisy images $x_{1:T}$ and the observed variable is the original image x_0 . We start by writing the log likelihood of the data $\log p_\theta(x_0)$ as the log marginal of the joint probability $p_\theta(x_{0:T})$:

$$\begin{aligned} \log p_\theta(x_0) &= \log \int p_\theta(x_{0:T}) \partial x_{1:T} \\ &= \log \int p_\theta(x_{0:T}) \cdot \frac{p(x_{1:T}|x_0)}{p(x_{1:T}|x_0)} \partial x_{1:T} \\ &= \log \int p(x_{1:T}) \cdot \frac{p_\theta(x_{0:T})}{p(x_{1:T}|x_0)} \partial x_{1:T} \\ &= \log \mathbb{E}_{x_{1:T} \sim p(x_{1:T}|x_0)} \left[\frac{p_\theta(x_{0:T})}{p(x_{1:T}|x_0)} \right]. \end{aligned} \quad (15)$$

Jensen's inequality states that, given a random variable Y and a convex function f , the following is true:

$$f(\mathbb{E}[Y]) \leq \mathbb{E}[f(Y)]. \quad (16)$$

If we apply Eq. (16) to Eq. (15) and change the inequality sign because the log function is concave, then we obtain the following:

$$\begin{aligned} \log p_\theta(x_0) &\geq \mathbb{E}_{x_{1:T} \sim p(x_{1:T}|x_0)} \left[\log \frac{p_\theta(x_{0:T})}{p(x_{1:T}|x_0)} \right] \cdot (-1) \\ -\log p_\theta(x_0) &\leq \mathbb{E}_{x_{1:T} \sim p(x_{1:T}|x_0)} \left[\log \frac{p(x_{1:T}|x_0)}{p_\theta(x_{0:T})} \right]. \end{aligned} \quad (17)$$

Eq. (17) shows that we can minimize the right-hand side of the inequality, instead of minimizing the expected negative log likelihood of the data for our generative model. We focus further on this term such that, at the end, we will derive the objective from Eq. (4).

By definition, the forward and reverse processes are Markovian. Based on this, we can rewrite the probabilities from Eq. (17) as follows:

$$\begin{aligned} p(x_{1:T}|x_0) &= p(x_T|x_{1:T-1}, x_0) \cdot p(x_{1:T-1}|x_0) \\ &= p(x_T|x_{T-1}) \cdot p(x_{T-1}|x_{1:T-2}, x_0) \cdot p(x_{1:T-2}|x_0) \\ &= \dots = \prod_{t=1}^T p(x_t|x_{t-1}), \\ p_\theta(x_{0:T}) &= p_\theta(x_0|x_{1:T}) \cdot p_\theta(x_{1:T}) \\ &= p_\theta(x_0|x_1) \cdot p_\theta(x_1|x_{2:T}) \cdot p_\theta(x_{2:T}) \\ &= \dots = p_\theta(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t). \end{aligned} \quad (18)$$

We replace the probabilities in the right-hand side of Eq. (17) with the products from Eq. (18) and apply the log property that transforms the products into sums:

$$\begin{aligned} \mathbb{E}_p \left[\log \frac{p(x_{1:T}|x_0)}{p_\theta(x_{0:T})} \right] &= \\ &= \mathbb{E}_{x_{1:T} \sim p(x_{1:T}|x_0)} \left[-\log p_\theta(x_T) + \sum_{t=1}^T \log \frac{p(x_t|x_{t-1})}{p_\theta(x_{t-1}|x_t)} \right]. \end{aligned} \quad (19)$$

The term $p(x_t|x_{t-1})$ can be transformed using the Bayes rule into $\frac{p(x_{t-1}|x_t) \cdot p(x_t)}{p(x_{t-1})}$, but the true posterior of the forward process $p(x_{t-1}|x_t)$ is intractable. However, if we additionally condition on the initial image x_0 , the posterior becomes tractable. Moreover, we know that $p(x_t|x_{t-1}, x_0) = p(x_t|x_{t-1})$ is true because the forward process is Markovian. Hence, if we apply the Bayes rule on $p(x_t|x_{t-1}, x_0)$, we will obtain the additional conditioning on the true posterior:

$$p(x_t|x_{t-1}, x_0) = \frac{p(x_{t-1}|x_t, x_0) \cdot p(x_t|x_0)}{p(x_{t-1}|x_0)}. \quad (20)$$

If we apply the derivation from Eq. (20) to Eq. (19) for all $t \geq 2$, the result is as follows:

$$\begin{aligned} \mathcal{L}_{vlb} &= \mathbb{E}_p \left[-\log p_\theta(x_T) + \sum_{t=1}^T \log \frac{p(x_t|x_{t-1})}{p_\theta(x_{t-1}|x_t)} \right] \\ &= \mathbb{E}_p[-\log p_\theta(x_T)] + \mathbb{E}_p \left[\sum_{t=2}^T \log \frac{p(x_{t-1}|x_t, x_0)}{p_\theta(x_{t-1}|x_t)} \right] \\ &\quad + \mathbb{E}_p \left[\sum_{t=2}^T \log \frac{p(x_t|x_0)}{p(x_{t-1}|x_0)} + \log \frac{p(x_1|x_0)}{p_\theta(x_0|x_1)} \right]. \end{aligned} \quad (21)$$

Observing that the terms of the second sum cancel out, $\mathcal{L}_{vlb}$ becomes:

$$\begin{aligned} \mathcal{L}_{vlb} &= \mathbb{E}_p[-\log p_\theta(x_T)] + \mathbb{E}_p \left[\sum_{t=2}^T \log \frac{p(x_{t-1}|x_t, x_0)}{p_\theta(x_{t-1}|x_t)} \right] \\ &\quad + \mathbb{E}_p \left[\log \frac{p(x_T|x_0)}{p(x_1|x_0)} + \log \frac{p(x_1|x_0)}{p_\theta(x_0|x_1)} \right]. \end{aligned} \quad (22)$$

Finally, if we rearrange the terms and transform the log rates into Kulback-Leibler divergences, the result is the formulation from Eq. (4):

$$\begin{aligned} \mathcal{L}_{vlb} &= \mathbb{E}_p[-\log p_\theta(x_0|x_1)] + \mathbb{E}_p \left[\sum_{t=2}^T \log \frac{p(x_{t-1}|x_t, x_0)}{p_\theta(x_{t-1}|x_t)} \right] \\ &\quad + \mathbb{E}_p \left[\log \frac{p(x_T|x_0)}{p_\theta(x_T)} \right] \\ &= -\log p_\theta(x_0|x_1) + KL(p(x_T|x_0) \| p_\theta(x_T)) \\ &\quad + \sum_{t=2}^T KL(p(x_{t-1}|x_t, x_0) \| p_\theta(x_{t-1}|x_t)). \end{aligned} \quad (23)$$

APPENDIX B

NOISE ESTIMATION.

In this section, we focus on the simplifications suggested by Ho *et al.* [2] and the adjustments needed to reach the simple objective from Eq. (6).

The first simplification is to avoid training the covariance of $p_\theta(x_{t-1}|x_t)$, fixing it beforehand to be $\sigma_t^2 \cdot \mathbf{I}$ instead. In practice, Ho *et al.* [2] propose to use $\sigma_t^2 = \beta_t$. This change impacts the Kulback-Leibler term of $\mathcal{L}_{vlb}$, because if the covariance is not trainable, then the divergence can be rewritten as being the distance between the means of the distributions plus some constant that does not depend on θ :

$$\begin{aligned} \mathcal{L}_{kl} &= KL(p(x_{t-1}|x_t, x_0) \| p_\theta(x_{t-1}|x_t)) \\ &= \frac{1}{2 \cdot \sigma_t^2} \cdot \|\tilde{\mu}(x_t, x_0) - \mu_\theta(x_t, t)\|^2 + C, \end{aligned} \quad (24)$$

where $\tilde{\mu}(x_t, x_0)$ is the mean of $p(x_{t-1}|x_t, x_0)$, $\mu_\theta(x_t, t)$ is the mean of $p_\theta(x_{t-1}|x_t)$ and C is a constant. We underline that at this point, the output of the neural network is $\mu_\theta(x_t, t)$.

The next change is based on the observation that the mean $\tilde{\mu}(x_t, x_0)$ can be expressed as a function of x_t and z_t , as follows:

$$\tilde{\mu}(x_t, x_0) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \hat{\beta}_t}} \cdot z_t \right). \quad (25)$$

This implies, according to Eq. (24), that $\mu_\theta(x_t, t)$ has to approximate this expression. However, x_t is the input of the model. Therefore, Ho *et al.* [2] propose to reparametrize $\mu_\theta(x_t, t)$ in the same way:

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \hat{\beta}_t}} \cdot z_\theta(x_t, t) \right), \quad (26)$$

where $z_\theta(x_t, t)$ is now the output of the neural network, namely an estimation for the noise z_t , given the noisy image x_t .

If we replace the means in $\mathcal{L}_{kl}$ with the parametrizations from Eq. (25) and Eq. (26), the result is the following:

$$\mathcal{L}_{kl} = \frac{\beta_t^2}{2\sigma_t^2 \alpha_t (1 - \hat{\beta}_t)} \|z_t - z_\theta(x_t, t)\|^2. \quad (27)$$

This term is essentially a time-weighted distance between the true noise of the image x_t and the estimation of the network. Ho *et al.* [2] simplify this term even further and discard the weights $\frac{\beta_t^2}{2\sigma_t^2 \alpha_t (1 - \hat{\beta}_t)}$, yielding a form that also covers the first term of $\mathcal{L}_{vllb}$. Therefore, with this latter changes in place, the final objective becomes the simplified version from Eq. (6):

$$\mathcal{L}_{simple} = \mathbb{E}_{t \sim [1, T]} \mathbb{E}_{x_0 \sim p(x_0)} \mathbb{E}_{z_t \sim \mathcal{N}(0, \mathbf{I})} \|z_t - z_\theta(x_t, t)\|^2. \quad (28)$$