



Generalised Diffusion Probabilistic Scale-Spaces

Pascal Peter¹

Received: 15 September 2023 / Accepted: 30 May 2024 / Published online: 17 June 2024
© The Author(s) 2024

Abstract

Diffusion probabilistic models excel at sampling new images from learned distributions. Originally motivated by drift-diffusion concepts from physics, they apply image perturbations such as noise and blur in a forward process that results in a tractable probability distribution. A corresponding learned reverse process generates images and can be conditioned on side information, which leads to a wide variety of practical applications. Most of the research focus currently lies on practice-oriented extensions. In contrast, the theoretical background remains largely unexplored, in particular the relations to drift-diffusion. In order to shed light on these connections to classical image filtering, we propose a generalised scale-space theory for diffusion probabilistic models. Moreover, we show conceptual and empirical connections to diffusion and osmosis filters.

Keywords Diffusion probabilistic models · Scale-spaces · Drift-diffusion · Osmosis

1 Introduction

Diffusion probabilistic models [1] have recently risen to the state-of-the-art in image generation, surpassing generative adversarial networks [2] in popularity. In addition to significant research activity, the availability of pre-trained latent diffusion networks [3] has also brought diffusion models to widespread public attention [4]. Practical applications are numerous, including the generation of convincing, high-fidelity images from text prompts or partial image data.

Initial diffusion probabilistic models [1, 5–10] relied on a forward drift-diffusion process that gradually perturbs input images with noise and can be reversed by deep learning. Recently, it has been shown that the concrete mechanism that gradually destroys information in the forward process has a significant impact on the image generation by the reverse process. Alternative proposed image degradations include blur [11], combinations of noise and blur [12–14], or image masking [12].

So far, diffusion probabilistic research was mostly of practical nature. Some theoretical contributions established connections to other fields such as score-matching [7–10],

variational autoencoders [6], or normalising flows [15]. Diffusion probabilistic models have been initially motivated [1] by drift-diffusion, a well-known process in physics. However, its connections to other physics-inspired methods remain mostly unexplored. Closely related concepts have a long tradition in model-based visual computing, such as osmosis filtering proposed by Weickert et al. [16]. In addition, there is a wide variety of diffusion-based scale-spaces [17–20]. Conceptually, these scale-spaces embed given images into a family of simplified versions. This resembles the gradual removal of image features in the forward process of diffusion probabilistic models.

Despite this multitude of connections, there is a distinct lack of systematic analysis of diffusion probabilistic models from a scale-space perspective. This is particularly surprising due to the impact of the forward process on the generative performance [13, 14]. It indicates that a deeper understanding of the information reduction could also lead to further practical improvements in the future.

1.1 Our Contribution

With our previous conference publication [21], we made first steps to bridge this gap between the scale-space and deep learning communities. To this end, we introduced first generalised scale-space concepts for diffusion probabilistic models. In this work, we further explore the theoretical

✉ Pascal Peter
peter@mia.uni-saarland.de

¹ Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science, Saarland University, Campus E1.7, 66041 Saarbrücken, Saarland, Germany

background of this successful paradigm in deep learning. In contrast to traditional scale-spaces, we consider the evolution of probability distributions instead of images. Despite this departure from conventional families of images, we can show scale-space properties in the sense of Alvarez et al. [17]. These include architectural properties, invariances, and entropy-based measures of simplification.

In addition to our previous findings [21], our novel contributions include

- A generalisation of our scale-space theory for diffusion probabilistic models (DPMs) which includes both variance-preserving and variance-exploding approaches,
- Generalised scale-space properties for the reverse process of DPM,
- A scale-space theory for inverse heat dissipation [13] and blurring diffusion [14],
- And a significantly extended theoretical and empirical comparison of three diffusion probabilistic models to homogeneous diffusion [18] and osmosis filtering [16].

2 Related Work

Besides diffusion probabilistic models themselves, two additional research areas are relevant for our own work. Since we adopt a scale-space perspective, classical scale-space research acts as the foundation for our generalised theory. Furthermore, we discuss connections to osmosis filters, which have a tradition in model-based visual computing.

2.1 Diffusion Probabilistic Models

Large parts of our scale-space theory are based on the work of Sohl-Dickstein et al. [1], which pioneered diffusion probabilistic models (DPMs). The latent diffusion model by Rombach et al. [3] was integral for the gain of popularity of this approach. The public availability of code and trained models sparked many practical applications and caused a shift [4] away from generative adversarial networks [2]. Applications range from image generation [1, 3–5] over image inpainting [22, 23], super-resolution [24], segmentation [25], and deblurring [26] to the generation of different types of media, including video [24] and audio [27].

Following the same principles as early DPMs, most approaches rely on adding noise in the forward process. Since coarse-to-fine strategies were shown to improve DPMs [24, 28], they inspired Lee et al. [29] to use blurring with Gaussian convolution instead. Rissanen et al. [13] leveraged the equivalence of Gaussian blur to homogeneous diffusion [18] to establish inverse heat dissipation models. They only add small amounts of observation noise, while the model of Hoogeboom and Salimans [14] generalises this process and

allows more substantial contributions of noise. Due to the close relations to both classical diffusion and osmosis filtering, we address both of these models in detail in Sect. 5.3. Image masking was recently also proposed as an alternative degradation in the forward process by Daras et al. [12].

We also rely on a more general version of the original DPM model that was proposed by Kingma et al. [6]. In addition, they introduced the notion of variational diffusion models and showed connections to variational autoencoders. Interestingly, early results of Vincent [30] established connections between denoising autoencoders and score-matching. Later publications showed more direct relations of score-based approaches to DPM [5]. Song and Ermon [7] related diffusion models to score-based Langevin dynamics with later follow-up results [8–10] that also include continuous time processes.

Theoretical contributions that are related to our own work are rare. Recently, Hagemann et al. [15] connected diffusion probabilistic models to a multitude of different concepts under the common framework of normalising flows. This includes relations to osmosis filtering, but not from a scale-space perspective. Franzese et al. [31] have investigated functional diffusion processes as a time-continuous generalisation of classical DPMs. They also allow both noise and blur as image degradations. In contrast, our models are time- and space-discrete.

2.2 Scale-Spaces

The second major field of research that forms the foundation of our contribution is scale-space theory. Scale-spaces have a long tradition in visual computing. Most of them rely on partial differential [17–20, 32] or pseudo-differential equations [33, 34], but they have also been considered for wavelets [35], sparse inpainting-based image representations [36], or hierarchical quantisation operators [37]. Such classical scale-spaces describe the evolution of an input image over multiple scales, which gradually simplifies the image. Since they obey a hierarchical structure and provide guarantees for simplification, they allow to analyse image features that are specific to individual scales. This makes them useful for tasks such as corner detection [38], modern invariant feature descriptors [39, 40], or motion estimation [41–43].

General principles for classical scale-spaces are vital for our contributions. They form the foundation for our generalised scale-space theory for DPM. We establish architectural, invariance, and information reduction properties in the sense of Alvarez et al. [17] for this new setting. In Sect. 5, we also mention where we drew inspiration from this contribution and other sources [18, 20] in more detail.

There are many different classes of scale-spaces, originating from the early work by Iijima [18], which was later popularised by Witkin [44]. They proposed a scale-space that

can be interpreted as evolutions according to homogeneous diffusion. These initial Gaussian scale-spaces [17, 18, 44–47] have been generalised with pseudodifferential operators [33, 34, 48, 49] or nonlinear diffusion equations [19, 20]. Moreover, a comprehensive theory for shape analysis exists in the form of morphological scale-spaces [17, 50–54]. Wavelet shrinkage as a form of blurring [35] and sparse image representations [36, 37] have been considered from a scale-space perspective as well. Among this wide variety of different options, for us, the original Gaussian scale-space is still the most relevant. It is closely related to the blurring diffusion processes we consider in Sect. 5.3.

Our novel class of stochastic scale-spaces considers families of probability distributions instead of sequences of images. Conceptually similar approaches are rare. The Ph.D. thesis of Majer [55] proposes a stochastic concept, which also considers drift-diffusion. However, it is not related to deep learning and simplifies images in a different way. Instead of adding noise or blur, it shuffles image pixels. Similarly, Koenderink and Van Doorn [56] proposed “locally orderless images”, a local pixel shuffling as an alternative to blur. Other probabilistic scale-space concepts are only broadly related. There have been theoretical considerations of connections between diffusion filters and the statistics of natural images [57] and practical applications in stem cell differentiation [58].

In parallel to our conference publication [21], Zach et al. [59] have used homogeneous diffusion scale-spaces on probability densities. However, they learn image priors via denoising score matching with practical applications to image denoising. We on the other hand focus on the scale-space theory of generative diffusion probabilistic models.

2.3 Osmosis Filtering

In visual computing, osmosis filtering is a successful class of filters that has been introduced by Weickert et al. [16] and generalises diffusion filtering [20]. Even though it creates deterministic image evolutions, it is connected to statistical physics. Namely, it is closely related to the Fokker-Planck equation [60] and by extension also to Langevin formulations and the Beltrami flow [61]. This suggests that there could also be connections to diffusion probabilistic models.

Since such connections to drift-diffusion also apply to diffusion probabilistic models, we investigate connections between these approaches in Sect. 6.1. There, we also discuss the continuous theory for osmosis filters as it was originally proposed by Weickert et al. [16] and later extended by Schmidt [62]. Vogel et al. [63] introduced both the corresponding discrete theory and a fast implicit solver which we use for our experiments.

Osmosis filters are well suited to integrate conflicting information from multiple images, which makes them an

excellent tool for image editing [16, 63, 64]. Additionally, they have been successfully used for shadow removal [16, 64, 65], the fusion of spectral images [66, 67], and image blending [68]. There are also applications for osmosis that do not deal with images. Notably, Hagenburg et al. [69] used osmosis to enhance numerical methods and considered a Markov chain formulation. While we deal with Markov processes in this paper, our interpretation of osmosis and the context in which we use it is significantly different.

There are also conceptually similar methods in visual computing that are also connected to drift-diffusion and predate osmosis. Namely, Hagenburg et al. [70] proposed a lattice Boltzmann model for dithering. Other broadly related filters are the directed diffusion models of Illner and Neunzert [71] and the covariant derivative approach of Georgiev [72].

3 Organisation of the Paper

We introduce the basic ideas of diffusion probabilistic models in Sect. 4, including Markov formulations for the forward and reverse processes. Based on these foundations, we propose generalised scale-space properties for three classes of probabilistic forward diffusion processes in Sect. 5 and briefly address reverse processes as well. As a link to classical scale-spaces and deterministic image filters, we investigate relations of diffusion probabilistic models to homogeneous diffusion and osmosis filtering in Sect. 6. We conclude with a discussion and an outlook in Sect. 7.

4 Diffusion Probabilistic Models

Diffusion probabilistic models [1] are generative approaches which have the goal to create new samples from a desired distribution. This distribution is unknown except for a set of given representatives. For image processing purposes, this training data typically consists of a set of images $f_1, \dots, f_{n_t} \in \mathbb{R}^{n_x n_y n_c}$ with n_c colour channels of size $n_x \times n_y$ and $n = n_x n_y n_c$ pixels. From a stochastic point of view, these images are realisations of a random variable F with an unknown probability density function $p(F)$. DPMs aim to sample from this *target distribution*.

4.1 The Forward Process

In a first step, the so-called *forward process*, diffusion probabilistic models map from the target distribution to a simpler distribution. While there are many alternatives, the standard normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ is a typical choice. Here, $\mathbf{I} \in \mathbb{R}^{n \times n}$ denotes the unit matrix. Thus, the forward process takes training images as an input and maps them to samples of

multivariate Gaussian noise. These noise samples act as seeds for the *reverse process*. Like other generative models such as generative adversarial networks [2], it maps from samples of this simple distribution back to the approximate target distribution. These diffusion probabilistic models (DPMs) create image structures from pure noise.

In Sect. 5.3, we also present alternatives to purely noise-based forward processes. Typically, the forward process is straightforward both conceptually and in its implementation. For practical tasks, a major challenge is the estimation of the corresponding reverse process, which is implemented with deep learning. Our focus lies mostly on a theoretical analysis of the forward process from a scale-space perspective. We also discuss the reverse process in Sect. 5.2, but due to its approximate nature, theoretical results are less comprehensive. The design of the forward process also has significant impact on the performance of the generative model [13, 14]. Thus, it also constitutes a more attractive direction for a scale-space focused investigation: Understanding the nature of existing forward processes might allow to carry over useful properties from existing classical scale-spaces.

Therefore, in Sect. 5, we show that a wide variety of existing diffusion probabilistic models fulfil generalised scale-space properties. This new class of scale-spaces differs significantly from classical approaches, since it does not consider the evolution of images, but of probability distributions instead. First, we need to establish a mathematical definition of the probabilistic forward process on a time-dependent random variable $U(t)$.

At time $t = 0$, this random variable has the initial distribution $p(F)$. For subsequent times $t_1 < t_2 < \dots < t_m$ a *trajectory* is defined as a sequence of temporal realisations $\mathbf{u}_1, \dots, \mathbf{u}_m$ of $U(t)$. It represents one possible evolution according to random additions of noise in each time step and is visualised in Fig. 1. Importantly, each image \mathbf{u}_i in a trajectory only depends on \mathbf{u}_{i-1} . This implies that the corresponding conditional transition probabilities fulfil the Markov property

$$p(\mathbf{u}_i | \mathbf{u}_{i-1}, \dots, \mathbf{u}_0) = p(\mathbf{u}_i | \mathbf{u}_{i-1}). \tag{1}$$

Here, we consider the probability of observing \mathbf{u}_i as a realisation of $U(t)$ at time t_i given $U(t_{i-1}) = \mathbf{u}_{i-1}$. Thus, the stochastic forward evolution is a Markov process [73] and we can write the probability density of the trajectory in terms of the transition probabilities from (1) and the initial distribution $p(\mathbf{u}_0) = p(F)$:

$$p(\mathbf{u}_0, \dots, \mathbf{u}_m) = p(\mathbf{u}_0) \prod_{i=1}^m p(\mathbf{u}_i | \mathbf{u}_{i-1}). \tag{2}$$

This property is also integral to establishing central architectural properties of our generalised scale-space in Sect. 5. In contrast to our earlier conference publication [21], we consider a more general transition probability than the original model of Sohl-Dickstein et al. [1]. Relying on the model of Kingma et al. [6], we use Gaussian distributions of the type

$$p(\mathbf{u}_i | \mathbf{u}_{i-1}) = \mathcal{N}(\alpha_i \mathbf{u}_{i-1}, \beta_i^2 \mathbf{I}). \tag{3}$$

Since $\mathbf{I} \in \mathbb{R}^{n \times n}$ denotes the unit matrix, the covariance matrix of this multivariate Gaussian is diagonal. Thus, for every pixel j , we consider independent, identically distributed Gaussian noise with mean $\alpha_i u_{i-1,j}$ and standard deviation β_i . Overall, the forward process has the free parameters $\alpha_i > 0$ and $\beta_i \in (0, 1)$. In practice, these parameters can be learned or chosen by a user. Often α_i is also defined as a function of β_i , which we discuss in more detail in Sect. 5.

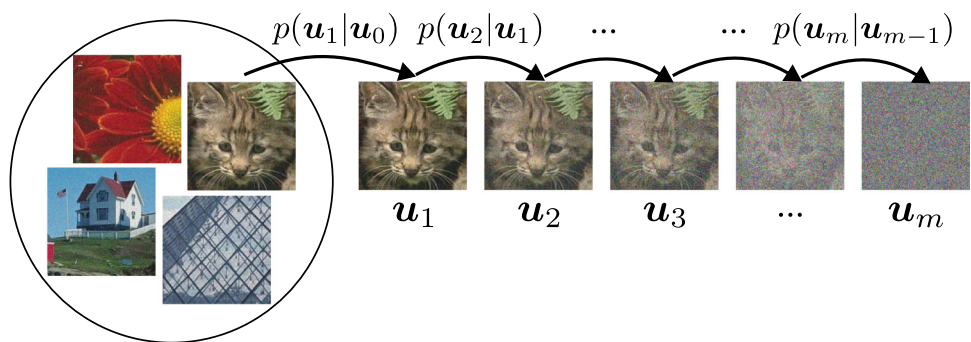
4.2 The Reverse Process

Sohl-Dickstein et al. [1] motivate the reverse process by a partial differential equation (PDE) that is associated to the forward process. In particular, they rely on the results of Feller [74]. These require the existence of the stochastic moments

$$m_k(\mathbf{u}_t, t) = \lim_{h \rightarrow 0} \frac{1}{h} \int p(\mathbf{u}_{t+h}, \mathbf{u}_t) (\mathbf{u}_{t+h} - \mathbf{u}_t)^k d\mathbf{u}_{t+h} \tag{4}$$

with $k \in \{1, 2\}$. Under this assumption, the probability density of the Markov process from Eq. (2) is a solution of the

Fig. 1 Forward DPM Trajectory. Starting from each sample of the initial distribution $p(\mathbf{u}_0)$, infinitely many trajectories exist. In each step of the trajectory, noise is added according to the transition probability $p(\mathbf{u}_i | \mathbf{u}_{i-1})$



partial differential equation

$$\frac{\partial}{\partial t} p = \frac{1}{2} \frac{\partial^2}{\partial \mathbf{u}_t \partial \mathbf{u}_t} (m_2(\mathbf{u}_t, t) p(\mathbf{u}_\tau, \mathbf{u}_t)) + \frac{\partial}{\partial \mathbf{u}_t} (m_1(\mathbf{u}_t, t) p(\mathbf{u}_\tau, \mathbf{u}_t)) . \tag{5}$$

Here, $p(\mathbf{u}_\tau, \mathbf{u}_t)$ denotes the probability density for a transition from \mathbf{u}_τ to \mathbf{u}_t with $\tau < t$. In Sect. 6.1, we use the fact that Eq. (5) is a drift-diffusion equation to discuss connections to osmosis filtering.

For practical purposes, it is important that Feller has proven that a solution of Eq. (5) also solves the backward equation

$$\frac{\partial}{\partial \tau} p = \frac{1}{2} m_2(\mathbf{u}_\tau, \tau) \frac{\partial^2}{\partial \mathbf{u}_\tau \partial \mathbf{u}_\tau} p + m_1(\mathbf{u}_\tau, \tau) \frac{\partial}{\partial \mathbf{u}_\tau} p . \tag{6}$$

Here, the backward perspective is obtained due to the exchange of roles of the earlier time τ with the later time t . Sohl-Dickstein et al. [1] exploit the close similarity of the backward equation to the forward equation. It implies that the reverse process from the normal distribution to the target distribution also has Gaussian transition probabilities. However, the mean and standard deviation are unknown and are estimated with a neural network instead. In particular, the training minimises the cross entropy to the target distribution $p(F)$. We discuss the reverse process in more detail in Sect. 5.2.

The capabilities of diffusion probabilistic models go beyond merely using the reverse process to sample from the target distribution. Additionally, it is possible to condition this distribution with side information such as partial image information or textual descriptions of the image content. This is useful for restoring missing image parts with inpainting [1, 3] or for text-to-image models [3]. However, our main focus are theoretical properties of multiple different forward processes. Since the estimation of the parameters for the reverse process is not relevant for our contributions, we refer to [3, 5, 6, 9] for more details.

5 Generalised Diffusion Probabilistic Scale-Spaces

In our previous conference publication [21], we introduced scale-space properties for the original forward diffusion probabilistic model (DPM) of Sohl-Dickstein et al. [1]. We generalise these results in Sect. 5.1 to a wider variety of noise schedules. Moreover, we introduce a generalised scale-space theory for the corresponding backward direction in Sect. 5.2. Finally, we address the recent inverse heat dissipation [13] and blurring diffusion models [14] in Sect. 5.3.

Before we discuss scale-space properties, we need to establish some preliminaries that allow us to rewrite transition probabilities in a useful way. The transition probabilities from Eq. (3) allow us to express the random variable at time t_i in terms of the random variable at time t_{i-1} according to

$$\mathbf{U}_i = \alpha_i \mathbf{U}_{i-1} + \beta_i \mathbf{G} . \tag{7}$$

Here, \mathbf{G} denotes Gaussian noise from the standard normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. This generalises the model of Sohl-Dickstein et al. [1] who use $\alpha_i = \sqrt{1 - \beta_i^2}$. Kingma et al. [6] also investigate variance-exploding diffusion [7] with $\alpha_i^2 = 1$. We discuss both types of models in Sect. 5.1.

Proposition 1 (Transition Probability from the Initial Distribution) *We can directly transition from \mathbf{U}_0 to \mathbf{U}_i by*

$$\mathbf{U}_i = \mathbf{U}_0 \prod_{\ell=1}^i \alpha_\ell + \left(\sum_{k=1}^{i-1} \beta_k \prod_{\ell=k+1}^i \alpha_\ell + \beta_i \right) \mathbf{G} \tag{8}$$

Proof For $i = 1$, the statement is fulfilled according to

$$\mathbf{U}_1 = \alpha_1 \mathbf{U}_0 + \beta_1 \mathbf{G} . \tag{9}$$

We prove the statement by induction. Applying the hypothesis for the step from i to $i + 1$, we obtain

$$\mathbf{U}_{i+1} = \alpha_{i+1} \mathbf{U}_i + \beta_{i+1} \mathbf{G} \tag{10}$$

$$= \alpha_{i+1} \left(\mathbf{U}_0 \prod_{\ell=1}^i \alpha_\ell + \left(\sum_{k=1}^{i-1} \beta_k \prod_{\ell=k+1}^i \alpha_\ell + \beta_i \right) \mathbf{G} \right) + \beta_{i+1} \mathbf{G} \tag{11}$$

$$= \mathbf{U}_0 \prod_{\ell=1}^{i+1} \alpha_\ell + \left(\sum_{k=1}^i \beta_k \prod_{\ell=k+1}^{i+1} \alpha_\ell + \beta_{i+1} \right) \mathbf{G} . \tag{12}$$

□

Thereby, we have established the transition probability from time 0 to time t_i as

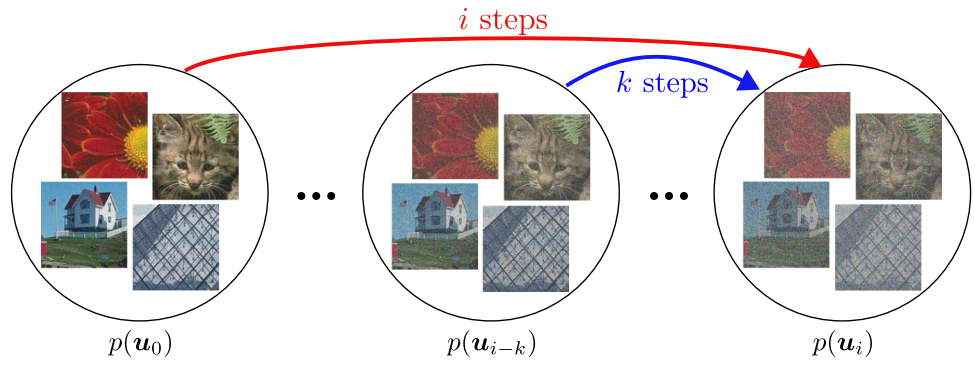
$$p(\mathbf{u}_i | \mathbf{u}_0) = \mathcal{N}(\lambda_i \mathbf{u}_0, \gamma_i^2 \mathbf{I}) \tag{13}$$

where the mean λ_i and standard deviation γ_i of this multivariate Gaussian distribution are

$$\lambda_i := \prod_{\ell=1}^i \alpha_\ell , \quad \gamma_i := \sum_{k=1}^{i-1} \beta_k \prod_{\ell=k+1}^i \alpha_\ell + \beta_i . \tag{14}$$

An interesting special case of the proposition above arises for the parameter choice $\alpha_i = \sqrt{1 - \beta_i^2}$ of the variance-preserving case [1]. Ho et al. [5] have shown that under this

Fig. 2 Semigroup Property for Forward DPM. Due to the Markov property, each intermediate scale i can be reached either from the training distribution $p(\mathbf{u}_0)$ in i steps, or from $p(\mathbf{u}_{i-k})$ in k steps. Note that this property does not apply to individual images as in classical scale-spaces. Instead, it refers to probability distributions which are visualised by samples from four different trajectories



condition, the transition probability becomes

$$p(\mathbf{u}_i | \mathbf{u}_0) = \mathcal{N} \left(\sqrt{\prod_{j=1}^i (1 - \beta_j^2)} \mathbf{u}_0, \mathbf{I} - \prod_{j=1}^i (1 - \beta_j^2) \mathbf{I} \right). \tag{15}$$

These insights are helpful for establishing a generalised scale-space theory for diffusion probabilistic models.

5.1 Generalised-Scale-Space Properties for Forward DPM

In the following, we propose central architectural properties for a generalised DPM scale-space and also discuss invariances. To this end, we consider the sequence of the marginal distributions of the random variable $\mathbf{U}(t)$. These can be obtained by integrating over all possible paths from the starting distribution to scale i according to

$$p(\mathbf{u}_i) = \int p(\mathbf{u}_0, \dots, \mathbf{u}_i) d\mathbf{u}_0 \cdots d\mathbf{u}_{i-1}. \tag{16}$$

Thus at each scale i , we consider the marginal distribution of $\mathbf{U}(t_i)$. Individual images are samples from the distributions at a given scale.

Property 1: Initial State

By definition, the initial distribution for the Markov process is the distribution $p(\mathbf{F})$ of the training database. Thus, it also defines the initial state $p(\mathbf{u}_0)$ of the scale-space.

Property 2: Semigroup Property

One central architectural property of scale-spaces is the ability to recursively construct a coarse scale from finer scales, i.e. the path from the initial state can be split into intermediate scales. This concept has been already established by Iijima [18] in the pioneering works on Gaussian scale-spaces. Intuitively, diffusion probabilistic models fulfil this property

since they are Markov processes. The property is visualised in Fig. 2.

Proposition 2 (Semigroup Property) *The distribution $p(\mathbf{u}_i)$ at scale i can be reached equivalently in i steps from $p(\mathbf{u}_0)$ or in ℓ steps from $p(\mathbf{u}_{i-\ell})$.*

Proof The probability density of the forward trajectory is defined in a recursive way in Eq. (2). Thus, we have to show that this property also carries over to the marginal distributions of the scale-space. We can reach $p(\mathbf{u}_i)$ either directly from \mathbf{u}_0 or from an intermediate scale ℓ by using the definition of the joint probability density of the Markov process:

$$p(\mathbf{u}_i) = \int p(\mathbf{u}_0) \prod_{j=1}^i p(\mathbf{u}_j | \mathbf{u}_{j-1}) d\mathbf{u}_0 \cdots d\mathbf{u}_{i-1} \tag{17}$$

$$= \int p(\mathbf{u}_{i-\ell}) \prod_{j=i-\ell+1}^i p(\mathbf{u}_j | \mathbf{u}_{j-1}) d\mathbf{u}_{i-\ell} \cdots d\mathbf{u}_{i-1}. \tag{18}$$

□

Property 3: Lyapunov Sequences

In classical scale-spaces (e.g. with diffusion), Lyapunov sequences quantify the change in the evolving image with increasing scale parameter. They constitute a measure of image simplification [20] in terms of monotonic functions. In practice, they often represent the information content of an image at a given scale. Here, we define a Lyapunov sequence on the evolving probability density instead.

To this end, we consider the conditional entropy of the random variable \mathbf{U}_i at time t_i given the random variable \mathbf{U}_0 . It constitutes a measure for the gradual removal of the image information from the initial distribution $p(\mathbf{u}_0)$.

Proposition 3 (Increasing Conditional Entropy) *The conditional entropy*

$$H_p(U_i|U_0) = - \iint p(\mathbf{u}_i, \mathbf{u}_0) \ln p(\mathbf{u}_i|\mathbf{u}_0) \, d\mathbf{u}_0 \, d\mathbf{u}_i. \tag{19}$$

increases with i under the assumption $\beta_j \in (0, 1)$ for all j with $\beta_{j+1} \geq (1 - \alpha_{j+1})\gamma_j$ with γ_j as defined in Eq. (14).

Proof We can reduce the problem of showing that the conditional entropy is monotonically increasing to a statement on the differential entropy of $p(\mathbf{u}_i|\mathbf{u}_0)$ since

$$H_p(U_i|U_0) = \int p(\mathbf{u}_0) \underbrace{\left(- \int p(\mathbf{u}_i|\mathbf{u}_0) \ln p(\mathbf{u}_i|\mathbf{u}_0) \, d\mathbf{u}_i \right)}_{=H_p(\mathbf{W}_i)} \, d\mathbf{u}_0 \geq H_p(U_{i-1}|U_0) \tag{20}$$

$$\Leftrightarrow H_p(\mathbf{W}_i) \geq H_p(\mathbf{W}_{i-1}). \tag{21}$$

According to Eq. (13), \mathbf{W}_i is from $\mathcal{N}(\lambda_i \mathbf{u}_0, \gamma_i^2 \mathbf{I})$. Therefore, the entropy of \mathbf{W}_i only depends on the covariance matrix $\gamma_i \mathbf{I}$ and yields

$$H_p(\mathbf{W}_i) = \frac{1}{2} \ln \left((2\pi e \gamma_i^2)^n \right). \tag{22}$$

Thus, the entropy is increasing if $\gamma_{i+1} \geq \gamma_i$. Furthermore, due to Eq. (14) we have

$$\begin{aligned} \gamma_{i+1} &= \sum_{k=1}^i \beta_k \prod_{\ell=k+1}^{i+1} \alpha_\ell + \beta_{i+1} = \alpha_{i+1} \left(\sum_{k=1}^{i-1} \beta_k \prod_{\ell=k+1}^i \alpha_\ell + \beta_i \right) \\ &+ \beta_{i+1} = \alpha_{i+1} \gamma_i + \beta_{i+1}. \end{aligned} \tag{23}$$

Since $\beta_i > 0$ and $\gamma_i > 0$, we require

$$\alpha_{i+1} \gamma_i + \beta_{i+1} \geq \gamma_i \Leftrightarrow \beta_{i+1} \geq (1 - \alpha_{i+1}) \gamma_i. \tag{24}$$

Again, we can also consider $\alpha_i = \sqrt{1 - \sigma_i}$, $\beta_i = \sqrt{\sigma_i}$ as in [1]. This gives us more concrete expressions for γ_i according to Eq. (15). With this we obtain

$$\gamma_{i+1} = \sqrt{1 - \prod_{j=1}^{i+1} (1 - \beta_j^2)} = \sqrt{1 - (1 - \beta_{i+1}^2) \prod_{j=1}^i (1 - \beta_j^2)} > \gamma_i. \tag{25}$$

Since $\beta_{i+1} \in (0, 1)$, this holds without further conditions with the noise schedule of Sohl-Dickstein et al. [1]. \square

Property 4: Permutation Invariance

The 1-D drift-diffusion process acts independently on each image pixel. Therefore, the spatial configuration of the pixels

does not matter for the process. In the following we provide formal arguments for a permutation invariance of all distributions created by the trajectories of the drift-diffusion.

Let $P(\mathbf{f})$ denote a permutation function that arbitrarily shuffles the position of the pixels in the image \mathbf{f} from the initial database. In particular, such permutations also include cyclic translations as well as rotations by 90° increments.

Proposition 4 (Permutation Invariant Trajectories) *Let \mathbf{u}_0 denote an image from the initial distribution and $\mathbf{v}_0 := P(\mathbf{u}_0)$ its permutation. Then, any trajectory $\mathbf{v}_0, \dots, \mathbf{v}_m$ obtained from the process in Eq. (7) is given by $\mathbf{v}_i = P(\mathbf{u}_i)$ for a trajectory $\mathbf{u}_0, \dots, \mathbf{u}_m$ starting with the original image \mathbf{u}_0 .*

Proof Consider the transition of \mathbf{v}_{i-1} to \mathbf{v}_i via a realisation \mathbf{g}_i of the random variable \mathbf{G} in Eq. (7). Then $\tilde{\mathbf{g}}_i := P^{-1}(\mathbf{g}_i)$ is also from the distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and we can obtain \mathbf{u}_i from \mathbf{u}_{i-1} by using this permuted transition noise.

Since $\mathbf{v}_0 = P(\mathbf{u}_0)$ holds by definition, we can inductively show the claim by considering

$$\mathbf{v}_i = \alpha_i \mathbf{v}_{i-1} + \beta_i \mathbf{g}_i = \alpha_i P(\mathbf{u}_{i-1}) + \beta_i P(\tilde{\mathbf{g}}_i) = P(\mathbf{u}_i). \tag{26}$$

\square

Any permutation P is a bijection. Thus every trajectory from a permuted image corresponds exactly to one trajectory starting from the original image. This directly implies permutation invariance of the corresponding distributions.

Corollary 5 (Permutation Invariant Distributions) *Let \mathbf{u}_0 denote an image from the initial distribution and $\mathbf{v}_0 := P(\mathbf{u}_0)$ its permutation. Then, for any image \mathbf{v} from $p(\mathbf{v}_i)$ there exists exactly one image \mathbf{u} from $p(\mathbf{u}_i)$ such that $\mathbf{v} = P(\mathbf{u})$.*

Property 5: Steady State

The steady-state distribution for $i \rightarrow \infty$ is a multivariate Gaussian distribution $\mathcal{N}(\mathbf{0}, \beta^2 \mathbf{I})$ with mean $\mathbf{0}$ and a covariance matrix $\beta^2 \mathbf{I}$ with $\beta < 2$. Convergence to a noise distribution is a cornerstone of diffusion probabilistic models and thus well known. For the sake of completeness, we provide formal arguments for this property. Moreover, we verify that for the noise schedule of Sohl-Dickstein et al. [1], we obtain $\beta = 1$. Thus, we verify that their steady state is the normal distribution.

Proposition 6 (Convergences to a Normal Distribution) *Let α_i and β_i be bounded from above by $a, b \in (0, 1)$, i.e. $\alpha_i \in (0, a]$ and $\beta_i \in (0, b]$. Moreover, let the assumptions of Proposition 3 be fulfilled. Then, for $i \rightarrow \infty$, the forward process from Eq. (2) converges to a normal distribution $\mathcal{N}(\mathbf{0}, \gamma^2 \mathbf{I})$ with $\gamma \leq \frac{b}{1-a}$.*

Proof According to Eqs. (8) and (14), we have

$$U_i = \lambda_i U_0 + \gamma_i \mathbf{G} \quad (27)$$

with \mathbf{G} from $\mathcal{N}(\mathbf{0}, \mathbf{I})$. For $i \rightarrow \infty$, we can immediately conclude $\lambda_i \rightarrow 0$ for the mean of the steady-state distribution since it is the product of i factors $\alpha_\ell < 1$.

Under the assumptions of Proposition 3, we have already shown that γ_i is increasing. Now let us consider the boundedness of γ_i , starting with definition Eq. (14) and using the assumptions $0 < \alpha_i \leq a$ and $0 < \beta_i \leq b$:

$$\begin{aligned} \gamma_i &= \sum_{k=1}^{i-1} \beta_k \prod_{\ell=k+1}^i \alpha_\ell + \beta_i \leq \sum_{k=1}^{i-1} b \prod_{\ell=k+1}^i a + b \quad (28) \\ &= b \left(\sum_{k=1}^{i-1} \prod_{\ell=k+1}^i a + 1 \right) = b \sum_{k=0}^{i-1} a^k \xrightarrow{i \rightarrow \infty} \frac{b}{1-a}. \quad (29) \end{aligned}$$

Overall, this shows that for $i \rightarrow \infty$, every trajectory converges to $\gamma \mathbf{G}$ with a $\gamma \leq \frac{b}{1-a}$ and \mathbf{G} from $\mathcal{N}(\mathbf{0}, \mathbf{I})$. \square

We have only specified an upper bound for the variance γ^2 of the steady-state distribution so far. The original DPM model [1] with $\alpha_i = \sqrt{1 - \beta_i^2}$ does not only act as an example that verifies reasonable parameter choices are possible under the assumptions of our proposition. Additionally, we can also explicitly infer that $\gamma = 1$ in this special case. Due to

$$\gamma_i = \sqrt{1 - \prod_{j=1}^i (1 - \beta_j^2)} \quad (30)$$

and $0 < 1 - \beta_j^2 < 1$, we obtain $\gamma_i \xrightarrow{i \rightarrow \infty} 1$. Thus, the original DPM model converges to the standard normal distribution.

Note that the variance-exploding model is not covered by the steady-state criterion without additional assumptions. Due to the parameter choice $\alpha_\ell = 1$, the sequence γ_i is given by

$$\gamma_i = \sum_{k=1}^i \beta_k. \quad (31)$$

As the name of the model suggests, the variance is thus not necessarily bounded for $i \rightarrow \infty$. For special choices of β_i , e.g. $\beta_i = \beta_0^{i-1}$ with $\beta_0 = 1$, we can however still get convergence to a normal distribution with a fixed variance. In the case of the example it would be $(1 - \beta_0)^{-1}$. Also note that due to $\alpha_\ell = 1$ in Eq. (14), $\lambda_i = 1$ for all i . Thus, the mean remains constant in this case.

The noisy steady state marks a clear difference to traditional scale-spaces. For instance, diffusion scale-spaces on images [20] converge to a flat steady state instead. However, the new class of diffusion probabilistic scale-spaces still underlies the same core concept: It removes information from the initial state recursively and hierarchically, leading to a state of minimal information w.r.t. the initial distribution.

5.2 Generalised Scale-Space Properties for Reverse DPM

Sohl-Dickstein et al. [1] argue via results of Feller [74] that for infinitesimal β_t , the distribution of forward and reverse trajectories becomes identical. However, these results are also tied in an inverse proportional way to the length of the trajectory. For arbitrary small β_t , the number of steps goes to infinity. For such a case of identical distributions, our results for the forward process would carry over to the reverse process: Initial and steady state are swapped and our Lyapunov sequences are decreasing instead of increasing. The remainder of the properties carry over verbatim.

In practice, however, the time-discrete reverse process used for DPMs is an approximation. Neural networks estimate the parameters for this reverse process. In the following, we comment on properties that can be established under these conditions.

To this end, we consider the reverse process of Sohl-Dickstein et al. [1] and denote its distributions with q . It takes the normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ as a starting distribution $q(\mathbf{u}_M)$. Transitions in the reverse direction from t_i to t_{i-1} fulfil

$$q(\mathbf{u}_{i-1} | \mathbf{u}_i) = \mathcal{N}(\boldsymbol{\mu}(\mathbf{u}_i, i), \boldsymbol{\Sigma}(\mathbf{u}_i, i)). \quad (32)$$

While these learned distributions are still Gaussian, they are significantly more complex than in the forward process. Both the learned mean and variance do not reduce to common scalars for all pixels. Furthermore, they depend on the current time step and on the image \mathbf{u}_i itself. Therefore, we can establish less properties for the reverse process than before, but central ones still carry over.

Property 1: Normal Distribution as Initial State

By definition, the distribution $q(\mathbf{u}_M)$ at time t_M is given by $\mathcal{N}(\mathbf{0}, \mathbf{I})$.

Property 2: Semigroup Property

The distribution $q(\mathbf{u}_i)$ at time t_i , $i < M$ can be reached in $M - i$ steps from $p(\mathbf{u}_M)$ or in $M - \ell - i$ steps from $p(\mathbf{u}_{M-\ell})$ with $M - \ell > i$. Since the Markov property is fulfilled, the

proof for the semigroup property is analogous to the forward case in Sect. 5.1.

Property 3: Lyapunov Sequence

As a by-product of their derivation of conditional bounds for the reverse process, Sohl-Dickstein et al. [1] have already concluded that both the entropy $H_q(\mathbf{u}_i)$ and the conditional entropy $H_q(\mathbf{u}_0|\mathbf{u}_i)$ are decreasing for the backward direction, i.e.

$$H_q(\mathbf{u}_{i-1}) \leq H_q(\mathbf{u}_i).$$

This is plausible, since the evolution starts with noise, a state of maximum entropy, and sequentially introduces more structure to it.

Property 4: Steady State

DPM reverse processes have the goal to enable sampling from the unknown initial distribution $p(\mathbf{u}_0)$. Convergence to this distribution is only guaranteed for the ideal case with identical distributions p and q . However, even if this is not strictly fulfilled, the parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are chosen such that they maximise a lower bound of the log likelihood

$$\int q(\mathbf{u}_0) \log p(\mathbf{u}_0) d\mathbf{u}_0. \tag{33}$$

In this sense, the reverse process approximates the distribution of the training data at time $t = 0$.

Property 4 from Sect. 5.1, the permutation invariance, does in general not apply to the reverse process. Since the parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ depend on the previous steps \mathbf{u}_i of the trajectory, the configuration of the pixels matters. Invariances will only be present if the network that estimates the parameters enforces them in its architecture.

5.3 Generalised Scale-Space Properties for Blurring Diffusion

Inverse heat dissipation [13] and blurring diffusion [14] models do not solely rely on adding noise in order to destroy features of the original image. Instead, they combine it with a deterministic homogeneous diffusion filter [18] to gradually blur this image. Such diffusion filters are well known as the origin of scale-space theory [18] and also constitute a special case of the osmosis filters we consider in more detail in Sect. 6.1. First, we discuss equivalent formulations of blurring diffusion in the spatial and transform domain [13, 14]. These allow us to transfer our scale-space results for DPM to this new setting.

A discrete linear diffusion operator can be interpreted either as the discretisation of a continuous time evolution described by a partial differential equation (see also Sect. 6.1) or as a Gaussian convolution. In the following, we consider only greyscale images with $N = n_x n_y$ pixels. Colour images can be processed by filtering each channel separately. Rissanen et al. [13] use an operator $\mathbf{A}_i = \exp(t_i \boldsymbol{\Delta})$, where $\boldsymbol{\Delta} \in \mathbb{R}^{N \times N}$ is a discretisation of the Laplacian $\Delta u = \partial_{xx} u + \partial_{yy} u$ with reflecting boundary conditions. By adding Gaussian noise, they turn the deterministic diffusion evolution into a probabilistic process given by

$$\mathbf{u}_i = \mathbf{A}_i \mathbf{u}_0 + \boldsymbol{\epsilon}_i, \quad \boldsymbol{\epsilon}_i \in \mathcal{N}(\mathbf{0}, \beta_i^2 \mathbf{I}) \tag{34}$$

In particular, they make use of a change of basis. To this end, let $\mathbf{V} \in \mathbb{R}^{N \times N}$ denote the basis transform operator of the orthogonal discrete cosine transform (DCT). Furthermore, we use the notation $\tilde{\mathbf{u}} = \mathbf{V} \mathbf{u}$ to denote the DCT representation of a spatial variable \mathbf{u} . Then, the diffusion operator $\mathbf{A}_t = \mathbf{V}^T \mathbf{B}_t \mathbf{V}$ reduces to a diagonal matrix $\mathbf{B}_t = \text{diag}(\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_N)$ in the DCT domain. The entries of \mathbf{B}_t result from the eigendecomposition of the Laplacian. Let the vector index $j \in \{1, \dots, N\}$ correspond to the position $(k, \ell) \in \{0, \dots, n_x - 1\} \times \{0, \dots, n_y - 1\}$ in the two-dimensional frequency domain. Then the entries of \mathbf{B}_t are given by

$$\alpha_j = \exp\left(-t\pi^2 \left(\frac{k^2}{n_x} + \frac{\ell^2}{n_y}\right)\right). \tag{35}$$

Note that for the frequency $(k, \ell)^T = (0, 0)^T$, $\alpha_0 = 1$. Hence, the average grey value of the image is preserved. Moreover, for all $j \neq 0$, we have $\alpha_j < 1$ for $t > 0$. Additionally, since \mathbf{V} is orthogonal and $\mathbf{V}^T \mathbf{V} = \mathbf{I}$, it also preserves Gaussian noise: For a sample $\boldsymbol{\epsilon}$ from the normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$, the transform $\tilde{\boldsymbol{\epsilon}}$ is from the same distribution. These properties are important for some of our scale-space considerations later on.

Rissanen et al. [13] proposed the DCT representation of their process for a fast implementation. However, Hoogeboom and Salimans [14] used it to derive a more general formulation in the DCT domain. Since \mathbf{B}_t is diagonal, a step from time t_i to time t_{i+1} can be considered for individual scalar frequencies j :

$$\tilde{u}_{i+1,j} = \alpha_{i,j} \tilde{u}_{i,j} + \beta_i \epsilon. \tag{36}$$

Here, ϵ is from $\mathcal{N}(0, 1)$. In contrast to Rissanen et al. [13], they do not limit the noise variance σ_t to minimal observation noise, but also allow to choose a noise schedule as in the DPM from Sect. 5.1. This formulation is particularly useful for us since it constitutes another 1-D drift-diffusion process as in Eq. (7). It allows us to transfer some of our previous

findings to the new setting. However, note that compared to the model in Eq. (7) we operate on frequencies instead of images as realisations of the random variable. Moreover, each frequency has its own individual parameters $\alpha_{i,j}$. In addition, β_i could be made frequency specific. However, in practice, Hoogeboom and Salimans [14] choose the same β_i for all frequencies. This also entails that the Markov criterion for this process in DCT space is given by

$$p(\tilde{\mathbf{u}}_{i+1}|\tilde{\mathbf{u}}_i) = \mathcal{N}(\tilde{\mathbf{u}}_i|\mathbf{B}_i\tilde{\mathbf{u}}_{i-1}, \beta_i^2\mathbf{I}). \tag{37}$$

In the following we consider a scale-space that is defined by the marginal distributions of the random variable $\mathbf{U}_t := \mathbf{V}^\top \tilde{\mathbf{U}}_t$, i.e. the backtransform of the trajectories in DCT space. Note that every trajectory in the frequency domain has exactly one corresponding trajectory in the spatial domain. Therefore, we can argue equivalently in the DCT domain or the spatial domain, depending on what is more convenient. As in Sect. 5.1, the process starts with the distribution of the training data or, respectively, the distribution of its discrete cosine transform.

Property 1: Initial State

By definition, the distribution $p(\mathbf{u}_0)$ at time $t_0 = 0$ is given by the distribution $p(\mathbf{F})$ of the training data.

Property 2: Semigroup Property

The distribution $p(\mathbf{u}_i)$ at scale i can be reached equivalently in i steps from $p(\mathbf{u}_0)$ or in ℓ steps from $p(\mathbf{u}_{i-\ell})$. The Markov property is fulfilled in the DCT domain. Thus, the proof for the semigroup property is analogous to the DPM model in Sect. 5.1. For each intermediate scale, we can switch back to the spatial domain by multiplication with \mathbf{V}^\top .

Property 3: Lyapunov Sequence

To establish this information reduction property, we require an analogous statement to Eq. (8). By using a similar induction proof for each frequency, we obtain

$$\tilde{\mathbf{U}}_i = \mathbf{M}_i\tilde{\mathbf{U}}_0 + \boldsymbol{\Sigma}_i\mathbf{G}. \tag{38}$$

Here, \mathbf{G} is from $\mathcal{N}(\mathbf{0}, \mathbf{I})$, $\mathbf{M}_i = \text{diag}(\lambda_{i,0}, \dots, \lambda_{i,N})$ with

$$\lambda_{i,j} = \prod_{\ell=1}^i \alpha_{\ell,j}, \tag{39}$$

and $\boldsymbol{\Sigma}_u = \text{diag}(\gamma_{i,0}, \dots, \gamma_{i,N})$ with

$$\gamma_{i,j} = \left(\sum_{k=1}^{i-1} \beta_i \prod_{\ell=k+1}^i \alpha_{\ell,j} + \beta_i \right). \tag{40}$$

This is a frequency-dependent analogue statement to the direct transition from time t_0 to time t_i in the DPM setting in Eq. (14).

Proposition 7 (Increasing Conditional Entropy) *The conditional entropy*

$$H_p(\tilde{\mathbf{U}}_i|\tilde{\mathbf{U}}_0) = - \int \int p(\tilde{\mathbf{u}}_i, \tilde{\mathbf{u}}_0) \ln p(\tilde{\mathbf{u}}_i|\tilde{\mathbf{u}}_0) d\tilde{\mathbf{u}}_0 d\tilde{\mathbf{u}}_i \tag{41}$$

increases with i under the assumption that for all frequencies j and $\beta_i \in (0, 1)$ we have $\beta_{i+1} \geq (1 - \alpha_{i+1,j})\gamma_{i,j}$.

Proof As in Sect. 5.1, the statement is equivalent to showing that the entropy $H_p(\mathbf{W}_i)$ of the distribution $p(\tilde{\mathbf{u}}_i|\tilde{\mathbf{u}}_0)$ is increasing. Thus, we need to show that $H_p(\mathbf{W}_i) \geq H_p(\mathbf{W}_{i-1})$. The probability distribution of $p(\tilde{\mathbf{u}}_i|\tilde{\mathbf{u}}_0)$ can be inferred from Eq. (13), but it is more complex than in the DPM case due to the frequency dependent parameters.

Fortunately, the entropy of the multivariate Gaussian distribution $\mathcal{N}(\mathbf{M}_i\tilde{\mathbf{u}}_0, \boldsymbol{\Sigma}_i)$ only depends on its covariance matrix $\boldsymbol{\Sigma}_i$ and is given by

$$H_p(\mathbf{W}_i) = \frac{1}{2} \ln \left((2\pi e)^n \det \boldsymbol{\Sigma}_i \right) = \frac{1}{2} \ln \left((2\pi e)^n \prod_{j=1}^N \gamma_{i,j} \right) \tag{42}$$

If $\gamma_{i+1,j} \geq \gamma_{i,j}$ holds for all frequencies j , the entropy is increasing. For a fixed frequency j , we can transfer the previous result from Eq. (23) to the scalar setting:

$$\gamma_{i+1,j} = \alpha_{i+1,j}\gamma_{i,j} + \beta_{i+1}. \tag{43}$$

Since $\beta_i > 0$ and $\gamma_i > 0$, we require

$$\alpha_{i+1,j}\gamma_{i,j} + \beta_{i+1} \geq \gamma_{i,j} \iff \beta_{i+1} \geq (1 - \alpha_{i+1,j})\gamma_{i,j} \tag{44}$$

This is a direct extension of our previous result in Sect. 5.1 to the frequency setting. \square

Note that there are also previous results for deterministic diffusion filters that use entropy as a Lyapunov sequence [20]. However, there the entropy is defined on the pixel values of the image instead of an evolving probability distribution. Individual trajectories rely on a deterministic diffusion filter. However, due to the added noise, the entropy statements of classical scale-spaces do not transfer to the trajectories of blurring diffusion.

Property 3: Preservation of the Average Grey Value

The DPM scale-space from Sect. 5.1 ensures convergence to Gaussian noise with zero mean, independently of the initial image. This requires the mean of the image to change. The behaviour of blurring diffusion scale-spaces is significantly different.

Proposition 8 (Preservation of the Average Grey Value) *Let \mathbf{u}_0 denote an image from the initial distribution $p(\mathbf{f})$. Then, all images in the trajectory $\mathbf{u}_0, \dots, \mathbf{u}_m$ have the same average grey value.*

This statement directly follows from an observation on the spatial version of the process. In every step, we add Gaussian noise with mean zero to a diffusion filtered image. Since diffusion filtering preserves the average grey value [20], this also holds for blurring diffusion. In colour images, the preservation of the average colour value applies for each channel.

Property 4: Rotation Invariance

In contrast to the forward DPM scale-space, blurring diffusion takes into account the neighbourhood configuration in the spatial domain due to the blurring of the homogeneous diffusion operator. Thus, permutation invariance does not apply to blurring diffusion scale-spaces.

However, in a space-continuous setting, the diffusion operator is rotationally invariant. The same applies to Gaussian noise samples: Under rotation, they remain samples from the same noise distribution. In the fully discrete setting, this rotation invariance is typically partially lost since only 90° rotations align perfectly with the pixel grid. However, this depends on the concrete implementation of the process. From this observation, we can directly deduce the following statement.

Proposition 9 (Rotation Invariant Trajectories) *Let \mathbf{u}_0 denote an image from the initial distribution and $\mathbf{v}_0 := R(\mathbf{u}_0)$ a rotation by a multiple of 90°. Then, any trajectory $\mathbf{v}_0, \dots, \mathbf{v}_m$ obtained from the process in Eq. (7) is given by $\mathbf{v}_i = R(\mathbf{u}_i)$ for a trajectory $\mathbf{u}_0, \dots, \mathbf{u}_m$ starting with the original image \mathbf{u}_0 .*

Property 5: Steady State

Due to the preservation of the average grey value in Property 3, a blurring diffusion scale-space cannot converge to a noise distribution with zero mean unless the initial image already had a zero mean. Images typically have nonnegative pixel values (e.g. a range of [0, 255] or [0, 1]). Thus, a zero mean would be only possible for a flat image or after a transformation to a range that is symmetric to zero (e.g. [-1, 1]). We do

not make any such assumption. However, for the sake of simplicity, we consider greyscale images in the following. For colour images, the same statements apply for each channel.

Proposition 10 (Convergence to a Mixture of Normal Distributions) *Under the assumptions of Proposition 7, $\alpha_{i,j} \in (0, a_j]$ with $0 < a_j < 1$, $\beta_i \in (0, b]$ with $0 < b < 1$, and $i \rightarrow \infty$, the forward process from Eq. (2) converges to a mixture of normal distributions $\mathcal{N}(\mu_k \mathbf{I}, \Sigma)$ with $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_N)$ and $\sigma_i \leq \frac{b}{1-a_j}$. The value μ_k with $k \in \{0, \dots, n_f\}$ assumes all possible average grey levels from the training database.*

Proof Consider a trajectory starting from an arbitrary image \mathbf{u}_0 from the training database with mean μ . As in the proof for Proposition 7, the findings from Eq. (8) and Eq. (14) for DPM carry over to our setting in the scalar case for each frequency j :

$$\tilde{u}_{i,j} = \lambda_{i,j} \tilde{u}_{0,j} + \gamma_{i,j} \epsilon \quad (45)$$

with ϵ from $\mathcal{N}(0, 1)$. For the convergence of $\lambda_{i,j}$, we have to consider the product of all frequency specific $\alpha_{i,j}$ according to Eq. (39). Here we have the special case $\alpha_{i,0} = 1$ for the lowest frequency, i.e. $\lambda_{i,0} = 1$ for all i . This is consistent with our previous findings in Proposition 8: The lowest frequency represents the average grey level, which remains unchanged.

For $j > 0$, we have $\alpha_{i,j} < 1$ and thus $\lambda_{i,j} \rightarrow 0$ for $i \rightarrow \infty$. Thus, for all other frequencies the contribution of \mathbf{u}_i vanishes and only its mean μ is preserved. The convergence of $\gamma_{i,j}$ to $\sigma_i \leq \frac{b}{1-a_j}$ is analogous to the proof of Proposition 6. This determines the standard deviation of the noise for each frequency. \square

Overall, blurring diffusion constitutes a scale-space that resembles the DPM scale-space in its architectural properties. The key difference lies in the incorporation of 2-D neighbourhood relationships between pixel values in the spatial domain. In the DCT domain, this translates to an individual set of process parameters for each frequency.

6 Diffusion Probabilistic Models and Osmosis

All diffusion probabilistic models considered in Sect. 5 have in common that they are connected to drift-diffusion processes. In image processing with partial differential equations (PDEs), osmosis filters proposed by Weickert et al. [16] are successfully applied to image editing and restoration tasks [16, 63–65]. Since they share their origin with diffusion probabilistic models, namely the Fokker-Planck equation [60], we discuss connections in the following. First, we briefly review the PDE formulation of osmosis filtering. Afterwards, we

discuss common properties as well as differences between both classes of models. Finally, we compare all four models experimentally.

6.1 Continuous Osmosis Filtering

Unlike the first part of the manuscript, we consider a grey value image as a function $f : \Omega \rightarrow \mathbb{R}_+$ defined on the image domain $\Omega \subset \mathbb{R}^2$. It maps each coordinate $\mathbf{x} \in \Omega$ to a positive grey value. In the following we limit our description to grey value images for the sake of simplicity. The filter can be extended to colour or arbitrary other multi-channel images (e.g. hyperspectral) by applying it channel-wise.

As in diffusion probabilistic models, osmosis filters consider the evolution of an image $u : \Omega \times [0, \infty) \rightarrow \mathbb{R}_+$ over time. However, here the evolution is entirely deterministic. Its initial state is given by a starting image f , i.e. for all \mathbf{x} we have $u(\mathbf{x}, 0) = f(\mathbf{x})$. In addition, the second major factor that influences the evolution is the *drift vector field* $\mathbf{d} : \Omega \rightarrow \mathbb{R}^2$ that can be chosen independently of the initial image and is typically used for filter design.

Given these two degrees of freedom, the image evolution fulfils the PDE [16]

$$\partial_t u = \Delta u - \mathbf{div}(\mathbf{d}u) \quad \text{on } \Omega \times (0, T]. \quad (46)$$

At the boundaries of Ω , reflecting boundary conditions avoid any exchange of information with areas outside of the image domain. There is a direct connection of this model to the inverse heat dissipation [13] and blurring diffusion [14]: All of these processes build on homogeneous diffusion [18], which is a special case of osmosis with $\mathbf{d} = \mathbf{0}$.

However, a non-trivial drift vector field enables to describe evolutions that do not merely smooth an image. The Laplacian Δu , which corresponds to the diffusion part of the PDE in Eq. (46), represents a symmetric exchange of information between neighbouring pixels. This symmetry can be broken by the drift component $-\mathbf{div}(\mathbf{d}u) = -\partial_x(d_1u) - \partial_y(d_2u)$. This is also vital for our own goal of relating osmosis image evolutions to diffusion probabilistic models and will allow us to introduce stochastic elements without any need to modify the PDE above.

6.2 Relating Osmosis and Diffusion Probabilistic Models

As noted in the previous section, for $\mathbf{d} = \mathbf{0}$, osmosis is directly connected to blurring diffusion. For the original DPM model, the connection is less obvious, but equally close. To this end, it is instructive to consider a 1-D osmosis process. Its evolution is described by the PDE

$$\partial_t u = \partial_{xx}u - \partial_x(\mathbf{d}u) \quad \text{on } \Omega \subset \mathbb{R} \times (0, T]. \quad (47)$$

A comparison with the PDE formulation of the DPM forward process in Eq. (5) reveals a significant structural similarity. Both equations describe an evolution w.r.t. the time t : 1-D osmosis considers an evolving image u , while the DPM evolution is defined on a probability density p . The derivative ∂_{u_i} w.r.t. positions of individual particles in the diffusion probabilistic model corresponds to the spatial derivative ∂_x in Eq. (47). The moments correspond to diffusivity and drift factors for the osmosis PDE and thus we can interpret 1-D osmosis as the deterministic counterpart to the probabilistic drift-diffusion process of DPM.

While both models are derived from the same physical principle, a 1-D osmosis filter would typically act on a 1-D signal instead of an image. In this 1-D setting, it also considers neighbourhood relations and exchanges information between neighbouring signal elements. In contrast, the 1-D drift-diffusion of DPM acts on individual pixels of a 2-D image which correspond to positions of individual particles. Thus, 1-D osmosis creates evolving 1-D signals, while a DPM trajectory consists of 2-D images.

For a meaningful comparison between the properties of osmosis filtering and the three diffusion probabilistic models, we require 2-D osmosis as in Eq. (46). Moreover, we need to add a stochastic component. Ideally, every trajectory of this new probabilistic osmosis process should inherit the theoretical properties of the deterministic model. Inverse heat dissipation [13] can be seen as a similar approach in that it provides a stochastic counter part to homogeneous diffusion [18]. However, there the addition of noise marks a clear departure from the PDE model and also implies that properties of the deterministic model do not carry over to the trajectory.

Instead of adding noise, we design the drift vector field \mathbf{d} in such a way that the osmosis evolution naturally converges to a noise sample. Weickert et al. [16] have shown that osmosis preserves the average grey value. Thus, for a given image f from the training data, we choose a noise image ϵ , where at every location $x \in \Omega$, $\epsilon(x)$ is from $\mathcal{N}(\mu_f, \sigma^2)$. Here, μ_f is the average grey value of f . Note that osmosis requires positive images. In practice, we clip the noise sample to $[10^{-6}, 1]$. While this truncates the tails of the Gaussian distribution, it has little impact for small standard deviation σ .

Now we can use another result of Weickert et al. [16] and consider the *compatible* case for osmosis. Our noise sample ϵ acts as guidance image for the osmosis process. The corresponding canonical drift vector field is defined by

$$\mathbf{d} = \frac{\nabla \epsilon}{\epsilon} \quad (48)$$

with the spatial gradient ∇ . Since $\mu_f = \mu_\epsilon$, the osmosis process converges to

$$\frac{\mu_f}{\mu_\epsilon} \epsilon = \epsilon. \tag{49}$$

Considering intermediate results of the evolution at times t_i yields a trajectory u_0, \dots, u_m of the probabilistic osmosis process with $u_0 = f$. Even though the intermediate marginal probabilities or the transition probabilities are not known, the process starts with the target distribution $p(F)$ and converges to an approximation to a normal distribution. In the following, we can empirically compare trajectories of this osmosis process to trajectories of diffusion probabilistic models.

6.3 Comparing Osmosis Filters and Diffusion Probabilistic Models

For concrete experiments, we require a discrete implementation of the continuous osmosis model from Sect. 6.1. As Vogel et al. [63], we use a stabilised BiCGSTAB solver [75]. For the noise guidance images, we use a standard deviation of $\sigma = 0.1$. All experiments are conducted on the Berkeley segmentation data set *BSDS500* [76].

We compare to the three models from Sect. 5. According to Eq. (7), we implement forward DPM [1] by successively adding Gaussian noise with the standard parameter choice $\alpha_i = \sqrt{1 - \beta_i^2}$ and $\beta_i = 0.1$. Inverse heat dissipation [13] and blurring diffusion [14] can be implemented in many equivalent ways. We based our implementation on the reference code of Rissanen et al. [13], which implements the Laplacian in the DCT domain according to Eq. (35). For inverse heat dissipation, we use the standard parameter $\beta_i = 0.01$. For blurring diffusion, we choose $\beta_i = 0.1$ such that it coincides with the guidance noise of our probabilistic osmosis. Furthermore, we also include homogeneous diffusion [18] without any added noise in our comparisons.

6.3.1 Visual Comparison

Figure 3 reveals visual similarities and differences between model trajectories. All five models successively remove features of the initial training sample: They are drowned out by noise, by blur, or a combination of both. For the probabilistic models, we have shown that this information reduction is quantified by entropy-based Lyapunov sequences. A similar statement holds for images in an osmosis trajectory. As shown by Schmidt [62], the relative entropy of u w.r.t. the noise sample ϵ is increasing:

$$L(t) := - \int_{\Omega} u(x, t) \ln \left(\frac{u(x, t)}{w(x)} \right) dx. \tag{50}$$

This also reflects the transition from the initial image f to the steady state ϵ . Similar statements on an unconditioned entropy apply to homogeneous diffusion [20].

Notably, only DPM does not preserve the average colour value of the initial image and converges to noise with mean zero. For visualisation purposes, the images of the DPM trajectory have therefore been affinely mapped to $[0, 1]$. DPM is also the only process that does not take neighbourhood relations between pixels into account. Therefore, edge features, such as the stripe pattern in Fig. 3a, remain sharp until they are completely overcome by noise.

This observation directly results from the 1-D drift-diffusion: DPM models the microscopic aspect of Brownian motion with colour values as particle positions. All other models consider the macroscopic aspect of drift-diffusion instead which considers colour values as particle concentrations per pixel cell. Consequentially, all other models apply 2-D blur in the image domain. Due to the very small amount of observation noise added by Rissanen et al. [13], the trajectory of heat dissipation is visually very similar to homogeneous diffusion.

Similarly, osmosis and blurring diffusion lead to visually almost identical trajectories. They mainly differ in the way how noise is added: Blurring diffusion uses explicit addition, while osmosis transitions to noise due to the drift vector field. In the following, we verify these observations quantitatively.

6.3.2 Variance Comparison

On the entire *BSDS500* database, we evaluate the evolution of the image variance over time in Fig. 4a. As expected, the pairs homogeneous diffusion/heat dissipation and osmosis/blurring diffusion exhibit very similar evolutions of the variance. On the way to the flat steady state, homogeneous diffusion and heat dissipation approach zero variance. Osmosis and blurring diffusion converge to a noise variance defined by the input parameters while DPM very slowly converges to the standard distribution. These observations coincide with the expectations from our theoretical results.

6.3.3 FID Comparison

Additionally, we can judge the similarity of intermediate distributions in the scale-space with the Fréchet-Inception distance (FID) [77]. It is widely used to judge the quality of generative models in terms of the approximation quality towards the target distribution. We use the implementation clean-fid [78] that avoids discretisation artefacts due to sampling and quantisation. Note that we measure the FID of probabilistic osmosis distributions relative to results of the other four models. Thus, a low FID indicates how closely each filter approximates osmosis.

Figure 4b also confirms our previous hypothesis: Heat dissipation and blurring diffusion consistently differ most from osmosis results since they rely mostly on blur and not on noise. DPM comes close to osmosis in its noisy steady state,

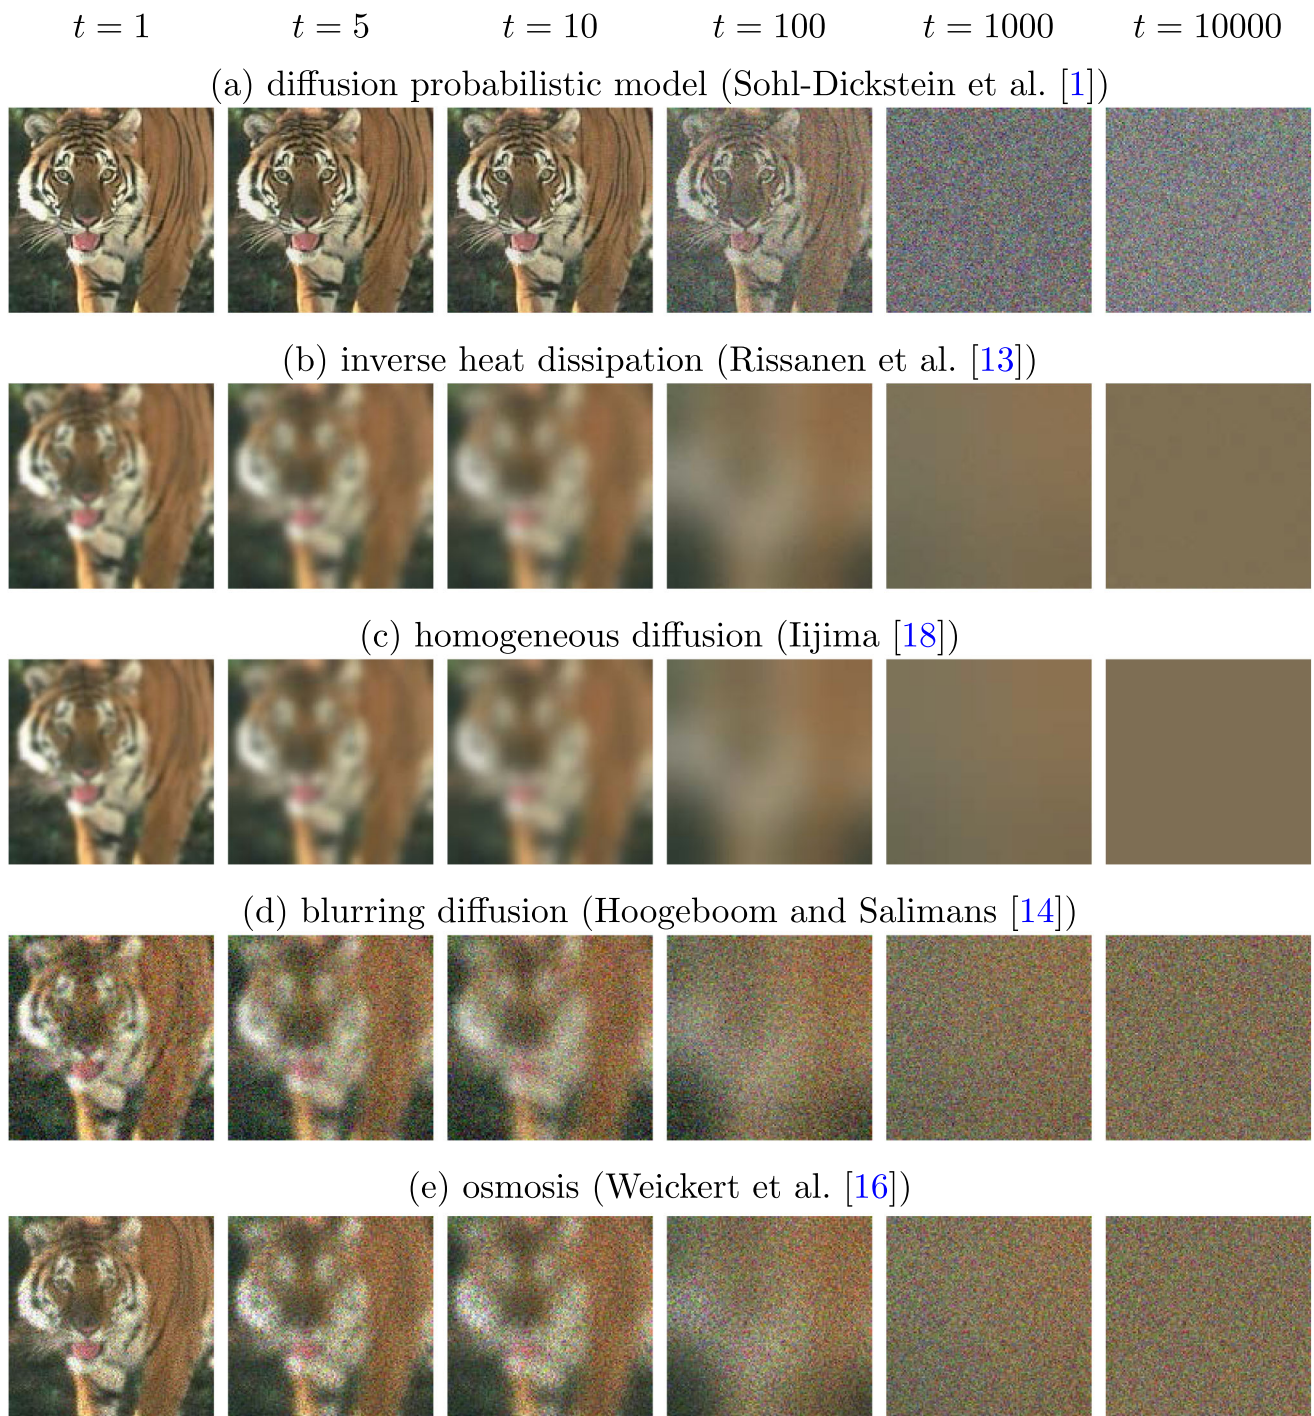


Fig. 3 Visual comparison of trajectories. The diffusion probabilistic model (a) behaves distinctively different compared to the other approaches since it does not perform blurring in the image domain.

Due to the minimal amounts of added noise, inverse heat dissipation (b) closely resembles homogeneous diffusion (c). With a suitable noise schedule, blurring diffusion (d) closely resembles osmosis filtering (e)

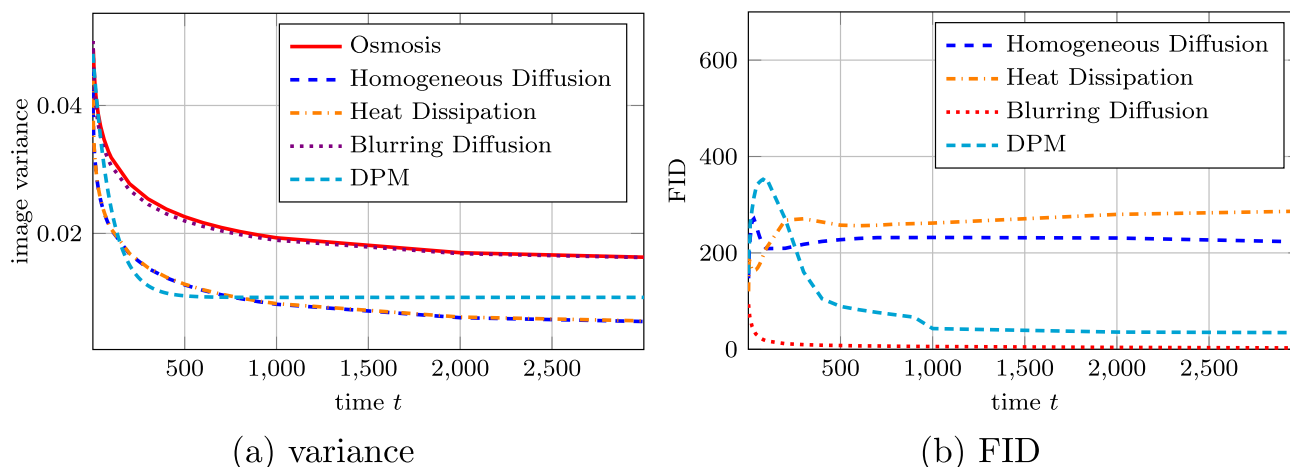


Fig. 4 Quantitative Comparison of Diffusion Probabilistic Models and Model-based Filters. Both the variance evolution over time in (a) and the FID w.r.t. the osmosis distributions in (b) suggest that DPM differs

but differs significantly in the initial evolution due the lack of 2-D smoothing. Blurring diffusion approximates osmosis surprisingly closely over the whole evolution.

Hoogeboom and Salimans [14] have found that heat dissipation improves the overall quality of generative models compared to DPM. Blurring diffusion yields even better results. Given our findings, we can interpret these observations from a scale-space perspective: The integration of 2-D neighbourhood relationships in the scale-space evolution is important for good diffusion probabilistic models. However, also the addition of sufficient amounts of stochastic perturbations is vital. Overall, recent advances can be interpreted as an increasingly accurate approximation of osmosis filtering, with an approximation to diffusion scale-spaces as an intermediate model. Using such an approximation instead of directly applying 2-D osmosis is convenient due to the more straightforward relation to the reverse process.

7 Conclusions and Outlook

Inspired by diffusion probabilistic filters, we have proposed the first class of generalised stochastic scale-spaces that describe evolutions of probability distributions instead of images. While the setting differs significantly from classical scale-spaces, central properties such as gradual, quantifiable simplification and causality still apply. These results suggest that in general, sequential generative models from deep learning are closely connected to scale-space theory. Therefore, we hope that in the future, the scale-space community will benefit from the discovery of new scale-spaces that might also be used in different contexts. In particular, existing generative models are mostly focused on the steady states as

significantly from the classical diffusion and osmosis filters. Heat dissipation approximates diffusion, while blurring diffusion approximates osmosis

the practically relevant output. The intermediate results of the associated scale-spaces could however also be useful in future applications.

On the flip side, trajectories of recent diffusion probabilistic models approximate well-known classical scale-space evolutions. This suggests that in the opposite direction, the deep learning community can potentially benefit from existing knowledge about scale-spaces by incorporating them into deep learning approaches.

Acknowledgements I thank Vassilen Chizhov, Michael Ertel, Kristina Schaefer, Karl Schrader, and Joachim Weickert for fruitful discussions and advice. We gratefully acknowledge the stimulating research environment of the GRK 2853/1 “Neuroexplicit Models of Language, Vision, and Action”, funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under project number 471607914.

Author Contributions As the single author, P.P. wrote the manuscript and conducted all research described in it.

Funding Open Access funding enabled and organized by Projekt DEAL.

Availability of Data and Materials The data sets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Declarations

Conflict of interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence,

unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: Bach, F., Blei, D. (eds.) Proc. 32nd International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 37, pp. 2256–2265. Lille, France (2015)
- Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial nets. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) Proc. 28th International Conference on Neural Information Processing Systems. Advances in Neural Information Processing Systems, vol. 27, pp. 2672–2680. Montréal, Canada (2014)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proc. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, pp. 10684–10695 (2022)
- Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis. In: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P.S., Vaughan, J.W. (eds.) Proc. 35th International Conference on Neural Information Processing Systems. Advances in Neural Information Processing Systems, vol. 34, pp. 8780–8794 (2021)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Advances in Neural Information Processing Systems, vol. 33, pp. 6840–6851. NeurIPS Foundation, San Diego, CA (2020)
- Kingma, D., Salimans, T., Poole, B., Ho, J.: Variational diffusion models. *Adv. Neural. Inf. Process. Syst.* **34**, 21696–21707 (2021)
- Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution **32** (2019)
- Song, Y., Ermon, S.: Improved techniques for training score-based generative models. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) Proc. 34th International Conference on Neural Information Processing Systems. Advances in Neural Information Processing Systems, vol. 33, pp. 12438–12448 (2020)
- Song, Y., Durkan, C., Murray, I., Ermon, S.: Maximum likelihood training of score-based diffusion models. *Adv. Neural. Inf. Process. Syst.* **34**, 1415–1428 (2021)
- Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: Proc. 2021 International Conference on Learning Representations, Kigali, Rwanda (2021)
- Bansal, A., Borgnia, E., Chu, H.-M., Li, J.S., Kazemi, H., Huang, F., Goldblum, M., Geiping, J., Goldstein, T.: Cold diffusion: Inverting arbitrary image transforms without noise. *arXiv preprint [arXiv:2208.09392](https://arxiv.org/abs/2208.09392)* [cs.CV] (2022)
- Daras, G., Delbraccio, M., Talebi, H., Dimakis, A., Milanfar, P.: Soft diffusion: Score matching with general corruptions. *Trans. Mach. Learn. Res.* (2023). issn. 2835-8856
- Rissanen, S., Heinonen, M., Solin, A.: Generative modelling with inverse heat dissipation. In: Proc. 11th International Conference on Learning Representations, Kigali, Rwanda (2023)
- Hoogeboom, E., Salimans, T.: Blurring diffusion models. *arXiv preprint [arXiv:2209.05557](https://arxiv.org/abs/2209.05557)* (2022)
- Hagemann, P.L., Hertrich, J., Steidl, G.: Generalized Normalizing Flows Via Markov Chains. *Elements in Non-local Data Interactions: Foundations and Applications*. Cambridge University Press, Cambridge (2023)
- Weickert, J., Hagenburg, K., Breuß, M., Vogel, O.: Linear osmosis models for visual computing. In: Heyden, A., Kahl, F., Olsson, C., Oskarsson, M., Tai, X.-C. (eds.) Energy Minimisation Methods in Computer Vision and Pattern Recognition. Lecture Notes in Computer Science, vol. 8081, pp. 26–39. Springer, Berlin (2013)
- Alvarez, L., Guichard, F., Lions, P.-L., Morel, J.-M.: Axioms and fundamental equations in image processing. *Arch. Ration. Mech. Anal.* **123**, 199–257 (1993)
- Iijima, T.: Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bull. Electrotech. Lab.* **26**, 368–388 (1962). (In Japanese)
- Scherzer, O., Weickert, J.: Relations between regularization and diffusion filtering. *J. Math. Imaging Vis.* **12**(1), 43–63 (2000)
- Weickert, J.: Anisotropic Diffusion in Image Processing. Teubner, Stuttgart (1998)
- Peter, P.: Generalised scale-space properties for probabilistic diffusion models. In: Calatroni, L., Donatelli, M., Morigi, S., Prato, M., Santaviesaria, M. (eds.) Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science, vol. 14009, pp. 601–613. Springer, Cham (2023)
- Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., Van Gool, L.: Repaint: Inpainting using denoising diffusion probabilistic models. In: Proc. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, pp. 11461–11471 (2022)
- Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: Proc. 2022 ACM SIGGRAPH Conference, Vancouver, Canada (2022). articleno. 15
- Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.* **23**(1), 2249–2281 (2022)
- Wolleb, J., Sandkühler, R., Bieder, F., Valmaggia, P., Cattin, P.C.: Diffusion models for implicit image segmentation ensembles. In: Konukoglu, E., Menze, B., Venkataraman, A., Baumgartner, C., Dou, Q., Albarqouni, S. (eds.) Proc. 5th International Conference on Medical Imaging with Deep Learning, Proceedings of Machine Learning Research, vol. 172, pp. 1336–1348. Zurich, Switzerland (2022)
- Ren, M., Delbraccio, M., Talebi, H., Gerig, G., Milanfar, P.: Image deblurring with domain generalizable diffusion models (2023)
- Kong, Z., Ping, W., Huang, J., Zhao, K., Catanzaro, B.: DiffWave: A versatile diffusion model for audio synthesis. In: Proc. 9th International Conference on Learning Representations, Vienna, Austria (2021)
- Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: Meila, M., Zhang, T. (eds.) Proc. 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 8162–8171. Honolulu, HI (2021)
- Lee, S., Chung, H., Kim, J., Ye, J.C.: Progressive deblurring of diffusion models for coarse-to-fine image synthesis. In: Proc. NeurIPS 2022 Workshop on Score-Based Methods (2022)
- Vincent, P.: A connection between score matching and denoising autoencoders. *Neural Comput.* **23**(7), 1661–1674 (2011)
- Franzese, G., Rossi, S., Rossi, D., Heinonen, M., Filippone, M., Michiardi, P.: Continuous-time functional diffusion processes (2023)
- Lindeberg, T.: Generalized Gaussian scale-space axiomatics comprising linear scale-space, affine scale-space and spatio-temporal scale-space. *J. Math. Imaging Vis.* **40**, 36–81 (2011)
- Duits, R., Florack, L., de Graaf, J., ter Haar Romeny, B.: On the axioms of scale space theory. *J. Math. Imaging Vis.* **20**, 267–298 (2004)

34. Schmidt, M., Weickert, J.: Morphological counterparts of linear shift-invariant scale-spaces. *J. Math. Imaging Vis.* **56**(2), 352–366 (2016)
35. Chambolle, A., Lucier, B.L.: Interpreting translationally-invariant wavelet shrinkage as a new image smoothing scale space. *IEEE Trans. Image Process.* **10**(7), 993–1000 (2001)
36. Cárdenas, M., Peter, P., Weickert, J.: Sparsification scale-spaces. In: Lellmann, J., Burger, M., Modersitzki, J. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 11603, pp. 303–314. Springer, Cham (2019)
37. Peter, P.: Quantisation scale-spaces. In: Elmoataz, A., Fadili, J., Quéau, Y., Rabin, J., Simon, L. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 12679, pp. 15–26. Springer, Cham (2021)
38. Alvarez, L., Morales, F.: Affine morphological multiscale analysis of corners and multiple junctions. *Int. J. Comput. Vision* **25**, 95–107 (1994)
39. Alcantarilla, P.F., Bartoli, A., Davison, A.J.: KAZE features. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Saito, Y., Schmid, C. (eds.) *Computer Vision - ECCV 2012. Lecture Notes in Computer Science*, vol. 7574, pp. 214–227. Springer, Berlin (2012)
40. Lowe, D.L.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision* **60**(2), 91–110 (2004)
41. Alvarez, L., Weickert, J., Sánchez, J.: Reliable estimation of dense optical flow fields with large displacements. *Int. J. Comput. Vision* **39**(1), 41–56 (2000)
42. Demetz, O., Weickert, J., Bruhn, A., Zimmer, H.: Optic flow scale-space. In: Bruckstein, A.M., Haar Romeny, B., Bronstein, A.M., Bronstein, M.M. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 6667, pp. 713–724. Springer, Berlin (2012)
43. Agustsson, E., Minnen, D., Johnston, N., Ballé, J., Hwang, S.J., Toderici, G.: Scale-space flow for end-to-end optimized video compression. In: *Proc. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Seattle, WA, pp. 8503–8512 (2020)
44. Witkin, A.P.: Scale-space filtering. In: *Proc. Eighth International Joint Conference on Artificial Intelligence*, vol. 2. Karlsruhe, West Germany, pp. 945–951 (1983)
45. Lindeberg, T.: *Scale-Space Theory in Computer Vision*. Kluwer, Boston (1994)
46. Sporring, J., Nielsen, M., Florack, L., Johansen, P. (eds.): *Gaussian Scale-Space Theory. Computational Imaging and Vision*, vol. 8. Kluwer, Dordrecht (1997)
47. Florack, L.: *Image Structure. Computational Imaging and Vision*, vol. 10. Springer, Dordrecht (2013)
48. Felsberg, M., Sommer, G.: Scale-adaptive filtering derived from the Laplace equation. In: Radig, B., Florczyk, S. (eds.) *Pattern Recognition. Lecture Notes in Computer Science*, vol. 2032, pp. 95–106. Springer, Berlin (2001)
49. Burgeth, B., Didas, S., Weickert, J.: Relativistic scale-spaces. In: Kimmel, R., Sochen, N., Weickert, J. (eds.) *Scale Space and PDE Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 3459, pp. 1–12. Springer, Berlin (2005)
50. Brockett, R.W., Maragos, P.: Evolution equations for continuous-scale morphology. In: *Proc. 1992 IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 3. San Francisco, CA, pp. 125–128 (1992)
51. van den Boomgaard, R., Smeulders, A.: The morphological structure of images: The differential equations of morphological scale-space. *IEEE Trans. Pattern Anal. Mach. Intell.* **16**, 1101–1113 (1994)
52. Caselles, V., Sbert, C.: What is the best causal scale space for three-dimensional images? *SIAM J. Appl. Math.* **56**(4), 1199–1246 (1996)
53. Kimia, B.B., Siddiqi, K.: Geometric heat equation and non-linear diffusion of shapes and images. *Comput. Vis. Image Underst.* **64**, 305–322 (1996)
54. Sapiro, G., Tannenbaum, A.: Affine invariant scale-space. *Int. J. Comput. Vision* **11**, 25–44 (1993)
55. Majer, P.: A statistical approach to feature detection and scale selection in images. PhD thesis, Department of Mathematics, University of Göttingen, Göttingen, Germany (2000)
56. Koenderink, J.J., Van Doorn, A.J.: The structure of locally orderless images. *Int. J. Comput. Vision* **31**(2), 159–168 (1999)
57. Pedersen, K.S.: Properties of Brownian image models in scale-space. In: Griffin, L.D., Lillholm, M. (eds.) *Scale-Space Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 2695, pp. 281–296. Springer, Berlin (2003)
58. Huckemann, S., Kim, K.-R., Munk, A., Rehfeldt, F., Sommerfeld, M., Weickert, J., Wollnik, C.: The circular sizer, inferred persistence of shape parameters and application to early stem cell differentiation. *Bernoulli* **22**(4), 2113–2142 (2016)
59. Zach, M., Pock, T., Kobler, E., Chambolle, A.: Explicit diffusion of Gaussian mixture model based image priors. In: Calatroni, L., Donatelli, M., Morigi, S., Prato, M., Santavesaria, M. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, pp. 3–15. Springer, Cham (2023)
60. Risken, H.: *The Fokker-Planck Equation*. Springer, New York (1984)
61. Sochen, N.A.: Stochastic processes in vision: From Langevin to Beltrami. In: *Proc. Eighth International Conference on Computer Vision*, vol. 1. Vancouver, Canada, pp. 288–293 (2001)
62. Schmidt, M.: Linear scale-spaces in image processing: Drift-diffusion and connections to mathematical morphology. PhD thesis, Department of Mathematics, Saarland University, Saarbrücken, Germany (2018)
63. Vogel, O., Hagenburg, K., Weickert, J., Setzer, S.: A fully discrete theory for linear osmosis filtering. In: Kuijper, A., Bredies, K., Pock, T., Bischof, H. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 7893, pp. 368–379. Springer, Berlin (2013)
64. d’Autume, M., Morel, J.-M., Meinhardt-Llopis, E.: A flexible solution to the osmosis equation for seamless cloning and shadow removal. In: *Proc 2018 IEEE International Conference on Image Processing*, Athens, Greece, pp. 2147–2151 (2018)
65. Parisotto, S., Calatroni, L., Caliari, M., Schönlieb, C.-B., Weickert, J.: Anisotropic osmosis filtering for shadow removal in images. *Inverse Prob.* **35**(5), 054001 (2019)
66. Parisotto, S., Calatroni, L., Daffara, C.: Digital cultural heritage imaging via osmosis filtering. In: Mansouri, A., El Moataz, A., Nouboud, F., Mammass, D. (eds.) *Image and Signal Processing. Lecture Notes in Computer Science*, vol. 10884, pp. 407–415. Springer, Cham (2018)
67. Parisotto, S., Calatroni, L., Bugeau, A., Papadakis, N., Schönlieb, C.-B.: Variational osmosis for non-linear image fusion. *IEEE Trans. Image Process.* **29**, 5507–5516 (2020)
68. Bungert, P., Peter, P., Weickert, J.: Image blending with osmosis. In: Calatroni, L., Donatelli, M., Morigi, S., Prato, M., Santavesaria, M. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 14009, pp. 652–664. Springer, Cham (2023)
69. Hagenburg, K., Breuß, M., Weickert, J., Vogel, O.: Novel schemes for hyperbolic PDEs using osmosis filters from visual computing. In: Bruckstein, A.M., ter Haar Romeny, B., Bronstein, A.M., Bronstein, M.M. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 6667, pp. 532–543. Springer, Berlin (2012)
70. Hagenburg, K., Breuß, M., Vogel, O., Weickert, J., Welk, M.: A lattice Boltzmann model for rotationally invariant dithering. In: Bebis, G., Boyle, R., Parvin, B., Koracin, D., Kuno, Y., Wang, J., Pajarola,

- R., Lindstrom, P., Hinkenjann, A., Encarnação, M.L., Silva, C.T., Coming, D. (eds.) *Advances in Visual Computing. Lecture Notes in Computer Science*, vol. 5876, pp. 949–959. Springer, Berlin (2009)
71. Illner, R., Neunzert, H.: Relative entropy maximization and directed diffusion equations. *Math. Methods Appl. Sci.* **16**, 545–554 (1993)
 72. Georgiev, T.: Covariant derivatives and vision. In: Bischof, H., Leonardis, A., Pinz, A. (eds.) *Computer Vision - ECCV 2006, Part IV. Lecture Notes in Computer Science*, vol. 3954, pp. 56–69. Springer, Berlin (2006)
 73. Gardiner, C.W.: *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences. Springer Series in Synergetics*, vol. 13. Springer, Berlin (1985)
 74. Feller, W.: On the theory of stochastic processes, with particular reference to applications. In: *First Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley, CA, pp. 403–432 (1949)
 75. Meister, A.: *Numerik Linearer Gleichungssysteme*, 5th edn. Vieweg, Braunschweig (2015)
 76. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(5), 898–916 (2011)
 77. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. *Adv. Neural Inf. Process. Syst.* **30**, 6629–6640 (2017)
 78. Parmar, G., Zhang, R., Zhu, J.: On aliased resizing and surprising subtleties in GAN evaluation. In: *Proc. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, New Orleans, LA, pp. 10684–10695 (2022)



Pascal Peter is a member of the Mathematical Image Analysis Group and a tenured lecturer at Saarland University (Saarbrücken, Germany). At the same institution, he graduated in computer science, mathematics, and pedagogy for high school teachers in 2011 and became a postdoctoral researcher after obtaining his Ph.D. degree in computer science in 2016. His research activities focus on the intersection of computer science and mathematics. His current research interests include image inpainting, compression, and the mathematical foundations of deep learning.

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