

RESEARCH

Open Access



# Connecting image inpainting with denoising in the homogeneous diffusion setting

Daniel Gaa<sup>1\*</sup> , Vassillen Chizhov<sup>1</sup>, Pascal Peter<sup>1</sup>, Joachim Weickert<sup>1</sup> and Robin Dirk Adam<sup>1</sup>

\*Correspondence:

[gaa@mia.uni-saarland.de](mailto:gaa@mia.uni-saarland.de)

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

## Abstract

While local methods for image denoising and inpainting may use similar concepts, their connections have hardly been investigated so far. The goal of this work is to establish links between the two by focusing on the most foundational scenario on both sides – the homogeneous diffusion setting. To this end, we study a denoising by inpainting (DbI) framework. It averages multiple inpainting results from different noisy subsets. We derive equivalence results between DbI on shifted regular grids and homogeneous diffusion filtering in 1D via an explicit relation between the density and the diffusion time. We also provide an empirical extension to the 2D case. We present experiments that confirm our theory and suggest that it can also be generalized to diffusions with nonhomogeneous data or nonhomogeneous diffusivities. More generally, our work demonstrates that the hardly explored idea of data adaptivity deserves more attention – it can be as powerful as some popular models with operator adaptivity.

**Mathematics Subject Classification:** 65D18; 68U10; 94A08

**Keywords:** Diffusion; Denoising; Inpainting; Partial differential equations; Sampling

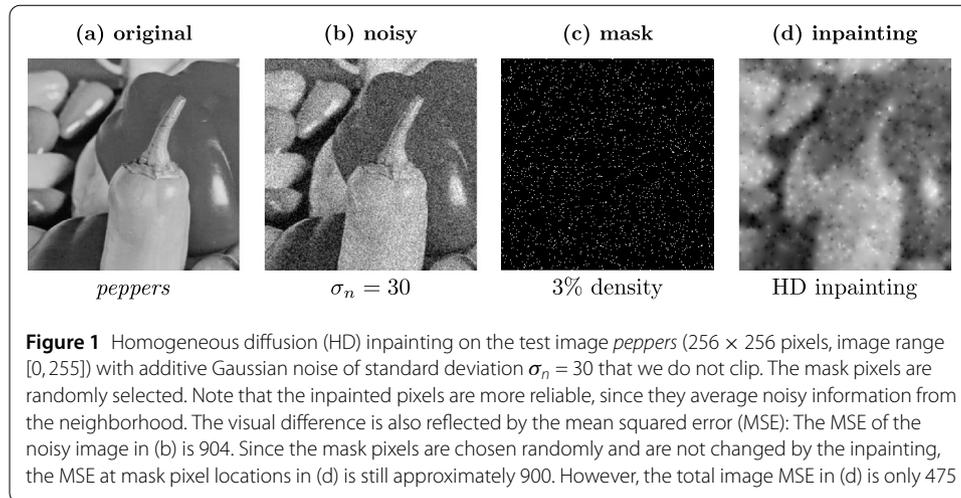
## 1 Introduction

Investigating connections between different fields in image analysis has often been rewarded with deep structural insights. Consider, for example, the link between variational image inpainting [1–5] and optic flow computation [6–8] via the concept of the *filling-in effect*. This effect is due to the smoothness term (regularizer) of the models, which inserts information at locations where the data term is absent or small in magnitude. The gradient flow for minimizing the variational energy functional leads to partial differential equations (PDEs) with a diffusion term.

While the filling-in effect has an obvious benefit for image inpainting, it can also lead to more powerful optic flow methods. It produces a dense flow field from the sparse information of the data term. Surprisingly, the parts of the flow field that are filled in by the diffusion-like regularization terms are usually those with the highest confidence [9].

Figure 1 shows a similar but hitherto hardly studied effect when performing *sparse inpainting* on noisy data. There the known data – the so-called *mask* – is a scattered set of pixels. The noisy mask pixels remain unchanged during the process, while the unknown areas in between are interpolated smoothly by averaging information from the noisy pix-

© The Author(s) 2025. **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/>.



els. We thus again have a scenario, where *the filled-in data are more reliable than the known data*. In the present manuscript we study how far this idea can lead us.

### 1.1 Our contribution

The goal of our work is to shed some light on the connections between PDE-based inpainting and denoising, two tasks which have coexisted for a long time, while their links have hardly been studied so far. We bridge this gap by a detailed investigation of the unconventional idea of denoising by inpainting. To facilitate a rigorous mathematical analysis, we focus on homogeneous diffusion. As will be explained below, it constitutes the most transparent and most foundational setting in both worlds.

The present paper builds upon our previous conference publication [10], in which the basic denoising by inpainting framework is established. This framework reconstructs a denoised version of an image by averaging the results of multiple inpaintings obtained from distinct masks. Furthermore, two concrete implementations of this framework are proposed in [10]: The first uses shifted regular masks and allows establishing a relation between denoising by inpainting and classical diffusion filtering in 1D, while the second uses probabilistic densification to adapt the masks to the image structures and enables an edge-preserving denoising behavior.

We extend the aforementioned results by a much broader study of the framework in [10], providing a fundamental understanding of the connections between PDE-based image inpainting and denoising. Since denoising methods can also be used as plug-and-play priors in algorithms for solving inverse problems [11–13], our relations between inpainting and denoising approaches may have an even broader application spectrum. Compared to [10], we introduce the following additional contributions:

- We show that the heuristically motivated DbI framework from [10] can be seen as a representative of a general probabilistic framework, for which we derive a sound theory. We argue that the denoising result obtained with such framework is an approximation of a minimum mean squared error (MMSE) estimate.
- We provide convergence estimates for the framework and propose a deterministic sampling approach to boost the convergence.

- We prove a general relation between the mask density of regular masks in the DbI framework and the diffusion time of homogeneous diffusion filtering in 1D. We also propose an empirical generalization of this result to 2D for uniform random masks.
- We integrate a step that optimizes the gray values at the selected mask pixels (tonal optimization) into the DbI framework. We investigate its effect on the MMSE estimate and perform experiments which confirm that tonal optimization can improve the denoising performance of DbI in practice.
- We show that the different spatial optimization approaches in the DbI framework correspond to specific posterior distributions. We compare two such strategies (that presented in [10] and a novel one) in terms of quality and provide the formulations for the respective probability distributions. Our experiments demonstrate that this data optimization leads to an edge-preserving denoising behavior.
- We replace homogeneous diffusion inpainting in the DbI framework by biharmonic inpainting and show that it is unable to improve denoising results. This confirms one of our key insights: The hitherto hardly practiced data optimization can be as powerful as widely used operator optimizations.

*Why homogeneous diffusion?* Our decision to focus on homogeneous diffusion is based on several reasons:

- For denoising and image simplification, one should keep in mind that homogeneous diffusion filtering is equivalent to Gaussian convolution. The Gaussian is the only convolution kernel that is separable and rotation invariant. The diffusion evolution generates a Gaussian scale-space representation [14–16], which is one of the most widely-used scale-spaces and forms the basis of highly successful interest point detectors such as SIFT [17] and its numerous variants.
- In inpainting applications, homogeneous diffusion is particularly popular in inpainting-based compression [18], where one stores only a sparse subset of all pixels and reconstructs the image in the decoding phase by inpainting. By optimizing the stored data, homogeneous diffusion can achieve surprisingly faithful reconstructions [19]. Moreover, its simplicity allows a detailed theoretical analysis [20], it frees the user from specifying parameters, and one can achieve real-time performance on current PC hardware even for large images [21].
- Last but not least, there exist already well-understood connections between diffusion processes for denoising and other approaches, such as variational regularization methods [22, 23] and wavelets [24, 25], but also deep neural network architectures [26, 27]. Thus, establishing also connections to inpainting ideas gives more comprehensive insights into various paradigms beyond diffusion-based denoising.

This discussion also implies that *it is not the goal of the present paper to design novel approaches that outperform the most recent state-of-the-art approaches for denoising or inpainting*. This is reserved for future research that may benefit from the foundational insights in the our manuscript.

## 1.2 Related work

Since we consider image inpainting as well as image denoising, we give an overview of some relevant methods from both fields and relate them to our work.

*PDE-based denoising and inpainting* We borrow several ideas from sparse PDE-based inpainting methods [18]. We mostly restrict ourselves to homogeneous diffusion inpainting [28], which can be implemented very efficiently [21, 29–33], and – in spite of its simplicity – can produce convincing results for suitably chosen data [20, 34–40]. Especially on piecewise constant images, such as cartoon images, depth maps or flow fields, homogeneous diffusion inpainting in conjunction with edge or segment information performs very well [28, 33, 41–45]. This even allows some of these methods [43, 44] to outperform HEVC [46] on such data. Nonlinear diffusion inpainting methods, e.g., edge-enhancing diffusion (EED) inpainting [18, 47], can improve reconstruction quality for sparse inpainting, enabling lossy image codecs [18, 48, 49] competitive to JPEG [50] and JPEG2000 [51]. On the other hand, such methods are more complex due to their nonlinearity. This complexity also carries over to the data optimization process. Higher-order inpainting operators can also be used for sparse inpainting [18, 36, 49, 52, 53], but can be more sensitive to noise. The quality of PDE-based sparse inpainting approaches strongly depends on the stored data, and in our denoising by inpainting framework we incorporate ideas from spatial optimization [20, 30, 33, 34, 36–40] and tonal optimization [30, 36, 38, 39, 54]. To interpret the filtering results of the denoising by inpainting framework, we compare to classical diffusion-based image denoising methods. Aside from the simple homogeneous diffusion [14], we also consider methods that adapt the diffusion operator to the given image, namely linear space-variant diffusion [55] and nonlinear diffusion [56]. We choose these methods because they are closest conceptually so we expect them to provide useful insights.

*Patch-based denoising and inpainting* Patch- or exemplar-based methods are another class of inpainting methods, and work especially well with textured data. The idea is to copy similar patches from known to unknown regions. Efros and Leung have proposed the first exemplar-based inpainting method [2], but many versions have been developed since then (e.g., [57–60]), including the method of Facciolo et al. for sparse inpainting [61]. Inpainting approaches combining PDE- and patch-based methods have also been presented [62, 63]. Inspired by the method of Efros and Leung [2], a patch-based denoising method called NL-means [64] has been proposed. It denoises an image based on a non-local weighted averaging of similar image patches. Other algorithms such as the famous BM3D algorithm [65] are also based on the filtering of image patches. These observations further substantiate the ties between denoising and inpainting. The NL-means method can even be interpreted as a case of a denoising by inpainting approach, although it does not use the inpainting ideas as directly as we do. Of course, a direct application of patch-based inpainting techniques would lead to the copying of erroneous noisy data, and not to a denoising effect.

*Sparse signal approximation* A popular approach in the field of image denoising relies on the idea that signals (and images) can be represented as a linear combination of a smaller number of basis signals – so-called atoms – that are selected from a dictionary [66]. Such a dictionary might for example consist of the basis vectors of a suitable transform, that makes the signal representation sparse (e.g., a wavelet transform [67] or a discrete cosine transform (DCT) [68]). The task is to then find those atoms, that best represent the given signal [69–71]. To fill in missing information in images, several authors also consider

sparse representations in some transform domain such as the DCT [72] or the shearlet domain [73]. This shows another bridge between the two tasks of denoising and inpainting. Hoffmann et al. [31] relate linear PDE-based inpainting methods to concepts from sparse signal approximation. They solve the inpainting problem with the help of discrete Green's functions [74, 75], which can be interpreted as atoms in a dictionary. This allows for a sparse representation of the inpainting solution. Kalmoun et al. [32] follow a similar approach by solving homogeneous diffusion inpainting with the charge simulation method [76, 77]. An application of homogeneous diffusion inpainting with Green's functions is the video codec by Andris et al. [29]. We justify certain design choices within the DbI framework with results from this field. Notably, homogeneous diffusion inpainting is based on the idea that the Laplacian of the reconstructed image is mostly sparse. On the other hand, the DbI framework combines multiple noisy sparse representations in order to get a denoised but nonsparse representation. The latter can be studied rigorously from a Bayesian denoising perspective, which is why we discuss this next.

*Bayesian denoising* The study of denoising has also been carried out from a probabilistic perspective. Here, the assumption is that some prior information regarding the noise distribution and/or the image distribution is available. This can be incorporated in a denoising framework through Bayes' rule, such that the final denoised result is conditioned on this information about the distributions. The latter provides a correspondence between classical denoising variational methods and specific Bayesian priors [78–80]. The standard approach is to employ statistical inference approaches, such as maximum likelihood (ML) estimation, maximum a posteriori (MAP) estimation, or minimum mean squared error (MMSE) estimation. Both the MAP and MMSE approach rely on a posteriori density, and as such they require a model of the distribution of considered classes of images. One of the first such models uses a Gibbs distribution for the prior [81]. Subsequently, a number of works have built upon this idea. The most relevant to our setting is that by Larsson and Selen [82], which studies MMSE estimation in the context of sparse vector representations. Our sparse inpaintings can be interpreted as such sparse vector representations. Moreover, in the current work we show that the averaging performed in [10] is, in fact, a Monte Carlo approach to approximate an MMSE estimate.

*Cross-validation* We also see the work of Craven and Wahba [83] on (generalized) cross-validation as conceptually related to parts of our work. Cross-validation can be used to optimize parameters in denoising models [82–84]. It removes data points from given noisy observations and judges the quality of a parameter selection in terms of the model's capability to reconstruct the data at these locations. Related ideas are also pursued in [85]. Probabilistic densification [42] and sparsification [39], two concepts from spatial optimization that we consider in our framework, also use the error of the inpainted reconstruction at left out locations – in our case also on noisy data. Yet, both applications differ, as the goal of the latter methods is to construct an inpainting mask and not to optimize model parameters.

*Neural denoising and inpainting* In recent years, many very powerful methods for inpainting and denoising have been proposed that rely on neural networks. They are, however, not a topic of our paper, since we aim at gaining structural insights into the connections between inpainting and denoising. Such results on classical approaches are still

relevant in the learning era [78]. They may serve as foundations for deep learning-based methods, and model- and learning-based approaches may be fused to obtain powerful and transparent algorithms. It is our hope that in the long run, our insights can also be beneficial to neural approaches.

### 1.3 Paper organization

In Sect. 2 we briefly introduce the basic idea behind diffusion filtering and its application to image denoising and image inpainting. In Sect. 3 we present the framework for denoising by inpainting from [10] and show that it can be interpreted as a Monte Carlo approach for approximating an MMSE estimate. We additionally provide convergence results, and suggest a method to boost the convergence by employing low-discrepancy sequences. In Sect. 4 we relate denoising by inpainting with nonadaptive masks to classical diffusion filtering. In Sect. 5 we present strategies for adaptively selecting the mask pixels in the DbI framework, which leads to space-variant denoising behavior. Our experiments and results are presented in Sect. 6, and we conclude the paper in Sect. 7.

## 2 Basics of diffusion filtering

In its original context of physics, diffusion is a process that equilibrates particle concentrations. When working with images, we interpret the gray values as particle concentrations and use diffusion processes as smoothing filters that balance gray value differences. To this end, we define the original grayscale image as a function  $f : \Omega \rightarrow \mathbb{R}$ , with  $\Omega \subset \mathbb{R}^2$  being a rectangular image domain. Similarly,  $u : \Omega \times [0, \infty) \rightarrow \mathbb{R}$  denotes the evolving, filtered image. Then the diffusion evolution is described by the following PDE:

$$\partial_t u(\mathbf{x}, t) = \operatorname{div}(g \nabla u(\mathbf{x}, t)) \quad \text{for } \mathbf{x} \in \Omega, t \in (0, \infty). \quad (1)$$

Here  $t$  denotes time,  $\nabla = (\partial_x, \partial_y)^\top$  is the spatial gradient,  $\operatorname{div}(\mathbf{v}) = \partial_x v_x + \partial_y v_y$  is the spatial divergence, and the scalar diffusivity  $g$  determines the local smoothing activity. We discuss different choices of  $g$  in Sect. 2.1. Note that  $g$  can be extended to a diffusion tensor to introduce anisotropy into the process [86], but since we do not consider such a case in this paper, we refrain from discussing it here. We equip the PDE with an initial condition at time  $t = 0$  and reflecting boundary conditions at the image boundary  $\partial\Omega$ :

$$u(\mathbf{x}, 0) = f(\mathbf{x}) \quad \text{for } \mathbf{x} \in \Omega, \quad (2)$$

$$\partial_n u(\mathbf{x}, t) = 0 \quad \text{for } \mathbf{x} \in \partial\Omega, t \in (0, \infty), \quad (3)$$

where  $\mathbf{n}$  is the outer normal vector at the image boundary. Solving this initial boundary value problem for  $u$  yields a family of filtered images  $\{u(\cdot, t) \mid t \geq 0\}$ .

### 2.1 Diffusion for image denoising

In image denoising the image  $f$  is a noisy version of the noise-free ground truth image  $f_r$ . In our case we assume zero-mean additive white Gaussian noise, i.e.,  $f = f_r + n$  with  $n \sim \mathcal{N}(0, \sigma_n^2)$ . Diffusion processes are good candidates for image denoising tasks thanks to their smoothing properties. Depending on the form of the diffusivity  $g$ , different processes are obtained.

### 2.1.1 Homogeneous diffusion

By setting  $g \equiv 1$ , (1) simplifies to  $\partial_t u = \Delta u$ , with  $\Delta u = \partial_{xx}u + \partial_{yy}u$  being the Laplacian operator. The resulting process is known as *homogeneous diffusion* [14]. Its analytical solution in the unbounded image domain  $\mathbb{R}^2$  is given by a convolution of the original image with a Gaussian kernel  $K_\sigma$  with standard deviation  $\sigma = \sqrt{2t}$ . The resulting images  $\{u(\cdot, t) \mid t \geq 0\}$  constitute the so-called Gaussian scale-space [14, 87]. Since  $g$  is selected to be constant, the smoothing strength is the same across the entire image. Therefore, not only the noise is reduced, but also semantically important image structures such as edges are smoothed.

### 2.1.2 Linear space-variant diffusion

To overcome the drawbacks of homogeneous diffusion, one can make the process space-variant by selecting a diffusivity function that varies depending on the structure of the *initial* image  $f$  [55]. This is called *linear space-variant diffusion*. If edges and other high-gradient features are to be preserved, the diffusivity should be decreasing with increasing gradient magnitude of the image, so that that the smoothing would be reduced at edges. An example for a suitable function is the Charbonnier diffusivity [88],

$$g(|\nabla f|^2) = \frac{1}{\sqrt{1 + \frac{|\nabla f|^2}{\lambda^2}}}, \tag{4}$$

where  $|\cdot|$  denotes the Euclidean norm. The contrast parameter  $\lambda > 0$  is used to distinguish locations where smoothing should be applied (for  $|\nabla f| \ll \lambda$ , we get  $g_\lambda \rightarrow 1$ ) and locations where it should be reduced (for  $|\nabla f| \gg \lambda$ , we obtain  $g_\lambda \rightarrow 0$ ).

### 2.1.3 Nonlinear diffusion

Alternatively, one can make the diffusivity function  $g$  dependent on the *evolving* image  $u$ . This allows updating the locations where smoothing is reduced during the evolution, by choosing them based on the image  $u$ , which becomes gradually smoother and less noisy. The resulting process  $\partial_t u = \operatorname{div}(g(|\nabla u|^2)\nabla u)$  is *nonlinear* [56]. The feedback mechanism throughout the evolution helps steering the process to achieve better results.

## 2.2 Diffusion for image inpainting

Diffusion processes can also be used to fill in missing information in images [28, 47, 89]. Particularly, they allow reconstructing an image from only a small number of pixels by propagating information from known to unknown areas [18]. The set of known pixels is called the *inpainting mask* and is denoted by  $K \subset \Omega$ . To recover the image, the information at the unknown locations is computed as the steady state ( $t \rightarrow \infty$ ) of a diffusion process, while the values at mask locations are preserved. The parabolic inpainting formulation is obtained by modifying (1) and (2) accordingly:

$$\partial_t u(\mathbf{x}, t) = \operatorname{div}(g\nabla u(\mathbf{x}, t)) \quad \text{for } \mathbf{x} \in \Omega \setminus K, t \in (0, \infty), \tag{5}$$

$$u(\mathbf{x}, t) = f(\mathbf{x}) \quad \text{for } \mathbf{x} \in K, t \in [0, \infty), \tag{6}$$

$$u(\mathbf{x}, 0) = 0 \quad \text{for } \mathbf{x} \in \Omega \setminus K, \tag{7}$$

$$\partial_n u(\mathbf{x}, t) = 0 \quad \text{for } \mathbf{x} \in \partial\Omega, t \in (0, \infty). \tag{8}$$

For  $g \equiv 1$ , (5) is the homogeneous diffusion PDE [14] and we talk about *homogeneous diffusion inpainting* (also called *harmonic inpainting*). We almost exclusively consider homogeneous diffusion inpainting in the remainder of this paper, so we set  $g \equiv 1$  in the following. Instead of computing the steady state of the parabolic diffusion equation, we may solve the corresponding boundary value problem:

$$-\Delta u(\mathbf{x}) = 0 \quad \text{for } \mathbf{x} \in \Omega \setminus K, \tag{9}$$

$$u(\mathbf{x}) = f(\mathbf{x}) \quad \text{for } \mathbf{x} \in K, \tag{10}$$

$$\partial_n u(\mathbf{x}) = 0 \quad \text{for } \mathbf{x} \in \partial\Omega. \tag{11}$$

The problem may be written equivalently using the variational formulation

$$\min_u \int_{\Omega} |\nabla u(\mathbf{x})|^2 d\mathbf{x}, \text{ such that } u(\mathbf{x}) = f(\mathbf{x}) \text{ for } \mathbf{x} \in K. \tag{12}$$

This suggests the interpretation that the inpainting is designed to penalize the gradient magnitude of the reconstruction, i.e., it inherently promotes smoothness. In order to simplify the discretization of the boundary value problem formulation, we introduce a mask indicator function  $c = \mathbb{1}_K$  (we use the term *mask* synonymously for the set  $K$  and the function  $c$ ), that takes the value 1 at points from  $K$  and 0 elsewhere. This allows us to combine (9) and (10) into a single equation

$$(c(\mathbf{x}) + (1 - c(\mathbf{x}))(-\Delta))u(\mathbf{x}) = c(\mathbf{x})f(\mathbf{x}) \quad \text{for } \mathbf{x} \in \Omega. \tag{13}$$

### 2.3 Discrete homogeneous diffusion inpainting

Since we are working with digital images, the above considerations need to be translated to the discrete setting. We therefore discretize the images on a regular pixel grid of size  $n_x \times n_y$ . Then we write them as vectors of length  $N = n_x n_y$  that are obtained by stacking the discrete images column-by-column, e.g.,  $\mathbf{f}, \mathbf{u} \in \mathbb{R}^N$ . Furthermore, let  $\mathbf{L} \in \mathbb{R}^{N \times N}$  denote the five-point stencil discretization matrix of the negated Laplacian  $(-\Delta)$  with reflecting boundary conditions  $\partial_n u(\mathbf{x}) = 0$  for  $\mathbf{x} \in \partial\Omega$ . Additionally, let  $\mathbf{C} = \text{diag}(\mathbf{c})$  be the diagonal matrix with the mask vector  $\mathbf{c} \in \{0, 1\}^N$  discretizing  $c$ , and let  $\mathbf{I}$  be the  $N \times N$  identity matrix. Then the discrete version of (13) can be formulated as the linear system of equations,

$$(\mathbf{C} + (\mathbf{I} - \mathbf{C})\mathbf{L}) \mathbf{u} = \mathbf{C}\mathbf{f}, \tag{14}$$

and the reconstruction can be written explicitly as

$$\mathbf{u} = \mathbf{r}(\mathbf{c}, \mathbf{f}) = (\mathbf{C} + (\mathbf{I} - \mathbf{C})\mathbf{L})^{-1} \mathbf{C}\mathbf{f}. \tag{15}$$

The inverse of the *inpainting matrix*  $\mathbf{M}_{\mathbf{c}} := \mathbf{C} + (\mathbf{I} - \mathbf{C})\mathbf{L}$  exists as long as  $\mathbf{C} \neq \mathbf{0}$  [33]. To deal with the case  $\mathbf{C} = \mathbf{0}$ , we define  $\mathbf{r}(\mathbf{0}, \mathbf{f}) := \frac{1}{N} \mathbf{1}^T \mathbf{f}$ , i.e., we take the average. If we want to approximate the image  $\mathbf{f}$  instead of interpolating it over  $\mathbf{C}$ , we can replace  $\mathbf{C}\mathbf{f}$  with  $\mathbf{C}\mathbf{g}$ , where

$$\mathbf{g} \in \underset{\mathbf{h}: \mathbf{h}|_{\bar{c}=\mathbf{0}}}{\text{argmin}} \|\mathbf{r}(\mathbf{c}, \mathbf{h}) - \mathbf{f}\|_2^2. \tag{16}$$

Here  $\mathbf{h}|_{\mathbf{c}}$  is the restriction of  $\mathbf{h}$  to  $\mathbf{c}$  and  $\mathbf{h}|_{\bar{\mathbf{c}}}$  is the restriction of  $\mathbf{h}$  to the complement  $\bar{\mathbf{c}} = \mathbf{1} - \mathbf{c}$ . The optimization is thus only over  $\mathbf{h}|_{\mathbf{c}}$  since the remainder of the values are irrelevant for the inpainting result, so we set them to zero. The least squares problem is known as the *tonal optimization* problem and we discuss its implications for the current work in Sect. 3.1.1. Additionally, we observe that the reconstruction is linear in  $\mathbf{g}$ . This motivates us to write it as a linear combination of basis vectors with weights given by  $\mathbf{g}|_{\mathbf{c}}$ . Let  $\mathbf{B}_{\mathbf{c}} := (\mathbf{M}_{\mathbf{c}}^{-1})|_{I \times \mathbf{c}}$  be the restriction of  $\mathbf{M}_{\mathbf{c}}^{-1}$  to the columns corresponding to nonzeros in  $\mathbf{c}$ , and we set  $m = \|\mathbf{c}\|_0$  to be the number of nonzeros in  $\mathbf{c}$ . By denoting the columns as  $\{\mathbf{b}_{\mathbf{c}}^k\}_{k=1}^m$ , i.e.,  $\mathbf{B}_{\mathbf{c}} = [\mathbf{b}_{\mathbf{c}}^1 \dots \mathbf{b}_{\mathbf{c}}^m]$ , we can write the reconstruction as

$$\mathbf{u} = \mathbf{r}(\mathbf{c}, \mathbf{g}) = \mathbf{M}_{\mathbf{c}}^{-1} \mathbf{C} \mathbf{g} = \mathbf{B}_{\mathbf{c}} \mathbf{g}|_{\mathbf{c}} = \sum_{k=1}^m (\mathbf{g}|_{\mathbf{c}})_k \mathbf{b}_{\mathbf{c}}^k. \tag{17}$$

We see that the columns of  $\mathbf{B}_{\mathbf{c}}$  are the basis vectors induced from  $\mathbf{r}$  and  $\mathbf{c}$ . They are also termed *inpainting echoes* [39, 90]. We note that inpainting with  $\mathbf{g}|_{\mathbf{c}} = \mathbf{f}|_{\mathbf{c}}$  constructs the interpolant over  $\mathbf{c}$  in the space  $\text{span}(\mathbf{B}_{\mathbf{c}}) \subseteq \mathbb{R}^N$ . Since the tonal optimization solution can be written as  $\mathbf{g}|_{\mathbf{c}} = (\mathbf{B}_{\mathbf{c}})^+ \mathbf{f}$ , where  $(\mathbf{B}_{\mathbf{c}})^+$  is the Moore–Penrose pseudoinverse, we note that  $\mathbf{r}(\mathbf{c}, \mathbf{g}) = \mathbf{B}_{\mathbf{c}} (\mathbf{B}_{\mathbf{c}})^+ \mathbf{f}$  is the orthogonal projection of  $\mathbf{f}$  on the subspace  $\text{span}(\mathbf{B}_{\mathbf{c}}) \subseteq \mathbb{R}^N$ , i.e., the best approximant of  $\mathbf{f}$  in this space.

### 3 Our denoising by inpainting framework

We now present the basic idea and the framework for denoising by inpainting proposed in our conference paper [10]. Since the framework inherently links inpainting and denoising, it is well suited to study connections between the two tasks. As previously mentioned, we use diffusion-based inpainting – specifically homogeneous diffusion inpainting – for image denoising, by only keeping a sparse subset of the noisy input data and by reconstructing the rest. Inpainting on noisy images differs from the classical setting and poses additional challenges. During the inpainting process, gray values at mask locations are not altered. As they might contain errors from the noise, these mask pixels are less trustworthy than inpainted pixels, which combine information from their surrounding mask pixels. While we want to exploit the filling-in effect in unknown areas, this observation implies that a single inpainted image cannot give satisfactory denoising results. Therefore, we compute multiple inpaintings with different masks and obtain the final result by averaging them. This ensures that none of the pixels remain unchanged (unless a pixel is contained in all masks). In the current work, we further mitigate the issue of noisy mask pixels by employing tonal optimization (see Sect. 3.1.1). If we denote the  $n$  different masks by  $\{\mathbf{c}^{\ell}\}_{\ell=1}^n$ , we can generate the inpaintings  $\{\mathbf{v}^{\ell}\}_{\ell=1}^n$  via

$$\mathbf{v}^{\ell} = \mathbf{r}(\mathbf{c}^{\ell}, \mathbf{f}) = (\mathbf{C}^{\ell} + (\mathbf{I} - \mathbf{C}^{\ell}) \mathbf{L})^{-1} \mathbf{C}^{\ell} \mathbf{f}. \tag{18}$$

We obtain the final denoising result  $\langle \mathbf{u} \rangle_n$  by averaging,

$$\langle \mathbf{u} \rangle_n = \frac{1}{n} \sum_{\ell=1}^n \mathbf{v}^{\ell} = \frac{1}{n} \sum_{\ell=1}^n \mathbf{r}(\mathbf{c}^{\ell}, \mathbf{f}). \tag{19}$$

As we fix the inpainting operator (for a discussion of denoising by biharmonic inpainting see Sect. 6.4), the only freedom in the framework lies in the selection of the different

masks. This is in contrast to the common strategy in denoising, where all available data is used and the operator is optimized instead. To study the effects of different data selection strategies, we will borrow several ideas from mask optimization for image compression. To obtain multiple different masks as our framework requires, we rely on some degree of randomness in the mask generation processes (see Sect. 5). Since we make use of stochastic strategies, we formalize and study DbI from a probabilistic point of view in the following subsection.

### 3.1 Probabilistic theory

As seen in (19), the denoised image is the result of averaging  $n$  inpaintings from  $n$  different masks, that are generated by some mask optimization process. In the following, we interpret this from a probabilistic point of view. This allows us to formalize the DbI framework from our conference paper [10] and provides us with tools to study and boost the convergence of our methods in Sects. 3.1.4 and 3.1.5, respectively. We take the masks  $\{\mathbf{c}^\ell\}_{\ell=1}^n$  to be independent and identically distributed samples from a predetermined distribution conditioned on  $\mathbf{f}$ , with a conditional probability mass function (PMF)  $p(\mathbf{c}|\mathbf{f})$ . Then the estimator  $\mathbf{u}$  converges to the following conditional expectation for  $n \rightarrow \infty$ :

$$\mathbb{E}[\langle \mathbf{u} \rangle_n | \mathbf{f}] = \mathbb{E} \left[ \frac{1}{n} \sum_{\ell=1}^n \mathbf{r}(\mathbf{c}^\ell, \mathbf{f}) \mid \mathbf{f} \right] = \frac{1}{n} \sum_{\ell=1}^n \mathbb{E}[\mathbf{r}(\mathbf{c}, \mathbf{f}) | \mathbf{f}] = \sum_{\mathbf{c} \in \{0,1\}^N} \mathbf{r}(\mathbf{c}, \mathbf{f}) p(\mathbf{c} | \mathbf{f}). \tag{20}$$

The second equality holds because the masks were assumed to be identically distributed, and thus  $\mathbb{E}[\mathbf{r}(\mathbf{c}^\ell, \mathbf{f}) | \mathbf{f}] = \mathbb{E}[\mathbf{r}(\mathbf{c}, \mathbf{f}) | \mathbf{f}]$  for any  $\mathbf{c}$  sampled with the same PMF  $p$ . The third equality follows from the definition of the conditional mathematical expectation. We note that from this probabilistic point of view, spatial adaptivity is provided through the design of the PMF  $p$ . The following proposition shows that the DbI result constitutes a minimum mean squared error (MMSE) estimate. This emphasizes its optimality under certain assumptions.

**Proposition 1** (DbI as an MMSE Estimate) *The expectation (20) of the DbI averaging (19) can be interpreted as an MMSE estimate under prior assumptions on the image and noise distributions, i.e., it solves the minimization problem*

$$\min_{\mathbf{u} \in \mathbb{R}^N} \mathbb{E}[\|\mathbf{u} - \mathbf{w}\|_2^2 | \mathbf{f}] = \min_{\mathbf{u} \in \mathbb{R}^N} \mathbb{E}[\|\mathbf{u} - \mathbf{r}(\mathbf{c}, \mathbf{f})\|_2^2 | \mathbf{f}]. \tag{21}$$

*Proof* We can rewrite the minimization problem (21) as

$$\min_{\mathbf{u} \in \mathbb{R}^N} \mathbb{E}[\|\mathbf{u} - \mathbf{r}(\mathbf{c}, \mathbf{f})\|_2^2 | \mathbf{f}] = \min_{\mathbf{u} \in \mathbb{R}^N} \sum_{\mathbf{c} \in \{0,1\}^N} \|\mathbf{u} - \mathbf{r}(\mathbf{c}, \mathbf{f})\|_2^2 p(\mathbf{c} | \mathbf{f}). \tag{22}$$

Taking the derivative with respect to  $\mathbf{u}$  and setting it to zero results in the MMSE estimate

$$\mathbf{u}^{\text{MMSE}} = \mathbb{E}[\mathbf{r}(\mathbf{c}, \mathbf{f}) | \mathbf{f}] = \sum_{\mathbf{c} \in \{0,1\}^N} \mathbf{r}(\mathbf{c}, \mathbf{f}) p(\mathbf{c} | \mathbf{f}). \tag{23}$$

By (20) this is the same as the expectation  $\mathbb{E}[\langle \mathbf{u} \rangle_n]$  of the DbI estimator  $\langle \mathbf{u} \rangle_n$ . □

The estimate  $\mathbf{u}^{\text{MMSE}}$  is close to  $\mathbf{f}_r$  (and  $\langle \mathbf{u} \rangle_n$  is close to  $\mathbf{f}_r$ ), whenever  $\mathbf{v} = \mathbf{r}(\mathbf{c}, \mathbf{f})$  with  $\mathbf{c} \sim p(\mathbf{c}|\mathbf{f})$  provides a good model for the distribution from which  $\mathbf{f}_r$  is assumed to originate. This formalization of DbI as an estimator for the MMSE estimate therefore provides an additional justification for the DbI framework as an image denoising approach.

### 3.1.1 MMSE and tonal optimization

The classical DbI formulation (19) from [10] employs an *interpolating* inpainting. It is natural to extend the framework to the best *approximating* inpainting, computing the denoised image  $\langle \mathbf{u} \rangle_n$  as

$$\langle \mathbf{u} \rangle_n = \frac{1}{n} \sum_{\ell=1}^n \mathbf{r}(\mathbf{c}^\ell, \mathbf{g}^\ell), \tag{24}$$

where the masks  $\{\mathbf{c}^\ell\}_{\ell=1}^n$  are selected as before, while  $\{\mathbf{g}^\ell\}_{\ell=1}^n$  are the solutions to the corresponding tonal optimization problems (16). Next we show that after relaxing assumptions on the gray values compared to Proposition 1, the MMSE estimate actually corresponds to DbI with an approximating inpainting instead of an interpolating one.

**Proposition 2** (DbI with Approximating Inpainting as an MMSE Estimate) *The DbI result based on a best approximating inpainting (24) can also be interpreted as an MMSE estimate, assuming that the gray values  $\mathbf{h}$  are now also a random variable conditioned on  $\mathbf{f}$ .*

*Proof* Firstly, we note that the minimization problem for the MMSE now differs, as the expectation has to be taken over the gray values  $\mathbf{h}$  as well:

$$\begin{aligned} \min_{\mathbf{u} \in \mathbb{R}^N} \mathbb{E}[\|\mathbf{u} - \mathbf{w}\|_2^2 | \mathbf{f}] &= \min_{\mathbf{u} \in \mathbb{R}^N} \mathbb{E}[\|\mathbf{u} - \mathbf{r}(\mathbf{c}, \mathbf{h})\|_2^2 | \mathbf{f}] \\ &= \min_{\mathbf{u} \in \mathbb{R}^N} \sum_{\mathbf{c} \in \{0,1\}^N} \mathbb{E}[\|\mathbf{u} - \mathbf{r}(\mathbf{c}, \mathbf{h})\|_2^2 | \mathbf{f}, \mathbf{c}] p(\mathbf{c} | \mathbf{f}) \\ &= \min_{\mathbf{u} \in \mathbb{R}^N} \sum_{\mathbf{c} \in \{0,1\}^N} \left( \int_{\mathbf{h} \in \mathbb{R}^N} \|\mathbf{u} - \mathbf{r}(\mathbf{c}, \mathbf{h})\|_2^2 p(\mathbf{h} | \mathbf{f}, \mathbf{c}) d\mathbf{h} \right) p(\mathbf{c} | \mathbf{f}). \end{aligned} \tag{25}$$

As before, differentiation with respect to  $\mathbf{u}$  yields the MMSE estimate

$$\mathbf{u}^{\text{MMSE}} = \mathbb{E}[\mathbf{r}(\mathbf{c}, \mathbf{h}) | \mathbf{f}] = \sum_{\mathbf{c} \in \{0,1\}^N} \mathbb{E}[\mathbf{r}(\mathbf{c}, \mathbf{h}) | \mathbf{f}, \mathbf{c}] p(\mathbf{c} | \mathbf{f}), \tag{26}$$

which is similar to (23), but now contains the expectation

$$\mathbb{E}[\mathbf{r}(\mathbf{c}, \mathbf{h}) | \mathbf{f}, \mathbf{c}] = \int_{\mathbf{h} \in \mathbb{R}^N} \mathbf{r}(\mathbf{c}, \mathbf{h}) p(\mathbf{h} | \mathbf{f}, \mathbf{c}) d\mathbf{h}. \tag{27}$$

To compute  $\mathbb{E}[\mathbf{r}(\mathbf{c}, \mathbf{h}) | \mathbf{f}, \mathbf{c}]$ , we need to know the a posteriori density  $p(\mathbf{h} | \mathbf{f}, \mathbf{c})$ . If we assume that the noise is normally distributed  $\mathbf{n} = (\mathbf{r}(\mathbf{c}, \mathbf{h}) - \mathbf{f}) \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$ , and that the gray values restricted to the mask  $\mathbf{h}|_c$  are normally distributed  $\mathbf{h}|_c \sim \mathcal{N}(\mathbf{0}, \sigma_{h|c}^2 \mathbf{I})$ , then the expectation can be calculated [82] as

$$\mathbb{E}[\mathbf{r}(\mathbf{c}, \mathbf{h}) | \mathbf{f}, \mathbf{c}] = \mathbf{B}_c \mathbb{E}[\mathbf{h}|_c | \mathbf{f}, \mathbf{c}] = \mathbf{B}_c \left( \frac{\sigma_n^2}{\sigma_{h|c}^2} \mathbf{I} + \mathbf{B}_c^T \mathbf{B}_c \right)^{-1} \mathbf{B}_c^T \mathbf{f}. \tag{28}$$

Since we do not know  $\sigma_{h|c}$  and because the assumption of the normality of the gray values may not be a very plausible one, we can dispense away with it by taking  $\sigma_{h|c} \rightarrow \infty$ , which results in a tonally optimized inpainting:

$$\lim_{\sigma_{h|c} \rightarrow \infty} \mathbb{E}[\mathbf{r}(\mathbf{c}, \mathbf{h})|\mathbf{f}, \mathbf{c}] = \mathbf{B}_c \lim_{\sigma_{h|c} \rightarrow \infty} \left( \frac{\sigma_n^2}{\sigma_{h|c}^2} \mathbf{I} + \mathbf{B}_c^T \mathbf{B}_c \right)^{-1} \mathbf{B}_c^T \mathbf{f} = \mathbf{B}_c (\mathbf{B}_c)^+ \mathbf{f}. \tag{29}$$

Using  $\mathbf{B}_c (\mathbf{B}_c)^+ \mathbf{f} = \mathbf{r}(\mathbf{c}, (\mathbf{B}_c)^+ \mathbf{f})$ , the new MMSE estimate differs with (23) only in that we have approximation instead of interpolation:

$$\mathbf{u}^{\text{MMSE}} = \sum_{\mathbf{c} \in \{0,1\}^N} \mathbb{E}[\mathbf{r}(\mathbf{c}, \mathbf{h})|\mathbf{f}, \mathbf{c}] p(\mathbf{c}|\mathbf{f}) = \sum_{\mathbf{c} \in \{0,1\}^N} \mathbf{r}(\mathbf{c}, (\mathbf{B}_c)^+ \mathbf{f}) p(\mathbf{c}|\mathbf{f}). \tag{30}$$

This corresponds exactly to the expectation of the approximating DbI formulation. □

We note that the above analysis did not require  $\mathbf{r}$  to be linear in  $\mathbf{f}$  except for the approximation of  $\mathbf{f}$ . Given a fixed  $\mathbf{c}$ , a natural extension to nonlinear operators could use nonlinear least-squares to compute something similar to  $\mathbf{B}_c^+ \mathbf{f}$ . By using the approximating formulation, we project the image onto the various subspaces induced by the inpainting operator  $\mathbf{r}$  and the mask  $\mathbf{c}$ . We will show in Sect. 6.3.2 that in practice, tonal optimization is able to improve quality and to reduce the variance of MMSE denoising, since it mitigates the error from the interpolation of noisy mask pixels and provides representations that are closer to  $\mathbf{f}$  in terms of MSE.

### 3.1.2 Interpreting tonal optimization as MAP estimate

Not directly related to the classical averaging formulation of DbI, but nevertheless interesting and a valuable extension, is the fact that spatial and tonal optimization for a single inpainting can also be framed as a maximum a posteriori (MAP) estimate. In MAP estimation, instead of minimizing the MSE, we want to find an inpainting  $\mathbf{w}$  that maximizes the posterior:

$$\operatorname{argmax}_{\mathbf{w}} p(\mathbf{w}|\mathbf{f}) = \operatorname{argmax}_{\mathbf{c}, \mathbf{h}} p(\mathbf{h}, \mathbf{c}|\mathbf{f}) = \operatorname{argmax}_{\mathbf{c}, \mathbf{h}} p(\mathbf{f}|\mathbf{h}, \mathbf{c}) p(\mathbf{h}|\mathbf{c}) p(\mathbf{c}). \tag{31}$$

We have assumed that  $\mathbf{w} = \mathbf{r}(\mathbf{c}, \mathbf{h})$  is an injection, so we have  $p(\mathbf{w}|\mathbf{f}) = p(\mathbf{r}(\mathbf{c}, \mathbf{h})|\mathbf{f}) = p(\mathbf{h}, \mathbf{c}|\mathbf{f})$ . In the noninjective case, one gets a set

$$p(\mathbf{w}|\mathbf{f}) = p(\mathbf{r}^{-1}(\mathbf{w})|\mathbf{f}) = p(\{\mathbf{h}, \mathbf{c} : \mathbf{w} = \mathbf{r}(\mathbf{c}, \mathbf{h})\}|\mathbf{f}), \tag{32}$$

which does not change the derivation meaningfully, except for introducing additional technical details. Thus, for the sake of clarity, we proceed with the injective case, but a similar argument holds in the general setting. The maximization problem (31) can be split into two optimization problems:

$$\max_{\mathbf{c}, \mathbf{h}} p(\mathbf{f}|\mathbf{h}, \mathbf{c}) p(\mathbf{h}|\mathbf{c}) p(\mathbf{c}) = \max_{\mathbf{c}} \left( \max_{\mathbf{h}} p(\mathbf{f}|\mathbf{h}, \mathbf{c}) p(\mathbf{h}|\mathbf{c}) \right) p(\mathbf{c}). \tag{33}$$

The inner one optimizes over the gray values  $\mathbf{h}$  given a mask  $\mathbf{c}$ , and the outer one optimizes over the masks  $\mathbf{c}$ . If we again assume that  $\mathbf{f} = \mathbf{r}(\mathbf{c}, \mathbf{h}) + \mathbf{n}$ , where  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$ , then the density  $p(\mathbf{f}|\mathbf{h}, \mathbf{c})$  is given by a Gaussian

$$p(\mathbf{f}|\mathbf{h}, \mathbf{c}) = \frac{1}{(2\pi\sigma_n^2)^{N/2}} \exp\left(-\frac{\|\mathbf{r}(\mathbf{c}, \mathbf{h}) - \mathbf{f}\|_2^2}{\sigma_n^2}\right). \tag{34}$$

Assuming also that the gray values are normally distributed, i.e.,  $\mathbf{h}|\mathbf{c} \sim \mathcal{N}(\mathbf{0}, \sigma_{h|\mathbf{c}}^2 \mathbf{I})$ , then the minimization problem with respect to  $\mathbf{h}$  is what we call *the regularized tonal optimization problem*:

$$\operatorname{argmax}_{\mathbf{h}|\mathbf{c}=\mathbf{0}} \exp\left(-\frac{\|\mathbf{r}(\mathbf{c}, \mathbf{h}) - \mathbf{f}\|_2^2}{\sigma_n^2} - \frac{\|\mathbf{h}|\mathbf{c}\|_2^2}{\sigma_{h|\mathbf{c}}^2}\right) = \operatorname{argmin}_{\mathbf{h}|\mathbf{c}=\mathbf{0}} \|\mathbf{B}_c \mathbf{h}|\mathbf{c} - \mathbf{f}\|_2^2 + \frac{\sigma_n^2}{\sigma_{h|\mathbf{c}}^2} \|\mathbf{h}|\mathbf{c}\|_2^2, \tag{35}$$

where the solution is the same as in (28), namely

$$\mathbf{h}|\mathbf{c}^* = \left(\frac{\sigma_n^2}{\sigma_{h|\mathbf{c}}^2} \mathbf{I} + \mathbf{B}_c^\top \mathbf{B}_c\right)^{-1} \mathbf{B}_c^\top \mathbf{f}. \tag{36}$$

Note that this can already be used for denoising with just a single inpainting with a mask  $\mathbf{c}$ , provided that we know the ratio of the variances of the noise and the gray values. The above expression suggests that we can then just apply a regularized tonal optimization to get the best MAP estimate. As before, we may take  $\sigma_{h|\mathbf{c}} \rightarrow \infty$  to get classical tonal optimization if desired. Of course, we also need to optimize with respect to the masks according to  $p(\mathbf{c})$ . In fact, if we take  $p(\mathbf{c}) = 0$  for  $\|\mathbf{c}\|_0 \neq m$ , and  $p(\mathbf{c})$  being equal for all  $\|\mathbf{c}\|_0 = m$ , then we get the spatial optimization problem with tonally optimized values

$$\min_{\|\mathbf{c}\|_0=m} \|\mathbf{r}(\mathbf{c}, \mathbf{h}|\mathbf{c}^*(\mathbf{f})) - \mathbf{f}\|_2^2. \tag{37}$$

If we take the interpolating case, we get the classical spatial optimization problem [39]

$$\min_{\|\mathbf{c}\|_0=m} \|\mathbf{r}(\mathbf{c}, \mathbf{f}) - \mathbf{f}\|_2^2. \tag{38}$$

The above further motivates using spatial optimization for denoising in both the interpolation and approximation cases; see Sect. 5.

### 3.1.3 Bayesian interpretation

In this subsection, we discuss how the above approaches fit in a general Bayesian perspective, which allows for meaningful interpretations of the occurring probabilities. This is valuable as MMSE and MAP estimates rely on a posterior  $p(\mathbf{w}|\mathbf{f})$ . Using Bayes' rule, this posterior can be rewritten as

$$p(\mathbf{w}|\mathbf{f}) = \frac{p(\mathbf{f}|\mathbf{w})p(\mathbf{w})}{p(\mathbf{f})} = \frac{p(\mathbf{f}|\mathbf{w})p(\mathbf{w})}{\int_{\mathbb{R}^N} p(\mathbf{f}|\mathbf{w})p(\mathbf{w}) d\mathbf{w}}, \tag{39}$$

where  $p(\mathbf{w})$  is the probability density function (PDF) for the distribution of images  $\mathbf{w}$  from which we assume  $\mathbf{f}$ , to originate. The likelihood  $p(\mathbf{f}|\mathbf{w})$  is the noise PDF, which in our case

is a Gaussian. The term  $p(\mathbf{f})$  is just a normalization constant that is irrelevant in practice, since it is not a function of  $\mathbf{w}$ . This shows that the task of finding a proper posterior distribution corresponds to introducing an appropriate prior  $p(\mathbf{w})$  under a given noise distribution  $p(\mathbf{f}|\mathbf{w})$ . This is known to be crucial for good denoising performance of Bayesian methods, and links our DbI framework to such approaches.

*Incorporating the inpainting operator* To introduce an inpainting operator  $\mathbf{r}$  into the above model, we make the assumption that any  $\mathbf{w}$  is synthesized as  $\mathbf{w} = \mathbf{r}(\mathbf{c}, \mathbf{h})$  for some mask  $\mathbf{c}$  and some gray values  $\mathbf{h}|\mathbf{c}$ . Since now the model depends on the masks we can rewrite the PDF as

$$p(\mathbf{w}|\mathbf{f}) = \sum_{\mathbf{c} \in \{0,1\}^N} p(\mathbf{w}|\mathbf{f}, \mathbf{c})p(\mathbf{c}|\mathbf{f}), \quad (40)$$

which is where the conditional mask PMF  $p(\mathbf{c}|\mathbf{f})$  comes into play – this is the other key ingredient for DbI along with the inpainting operator. We will see that this PMF allows us to introduce spatial adaptivity (Sect. 5.2, Fig. 9) for operators that are otherwise not spatially adaptive. Finally, we can also rewrite  $p(\mathbf{w}|\mathbf{f}, \mathbf{c})$  using Bayes' rule in order to relate the above formulation to (39):

$$p(\mathbf{w}|\mathbf{f}, \mathbf{c}) = \frac{p(\mathbf{f}|\mathbf{w}, \mathbf{c})p(\mathbf{w}|\mathbf{c})}{p(\mathbf{f}|\mathbf{c})} = \frac{p(\mathbf{f}|\mathbf{w}, \mathbf{c})p(\mathbf{w}|\mathbf{c})}{\int_{\mathbb{R}^N} p(\mathbf{f}|\mathbf{w}, \mathbf{c})p(\mathbf{w}|\mathbf{c}) d\mathbf{w}}. \quad (41)$$

This provides a similar interpretation, but now we have knowledge about the mask. As before  $p(\mathbf{f}|\mathbf{w}, \mathbf{c})$  models the noise, but now  $p(\mathbf{w}|\mathbf{c})$  models the distribution of the gray values defining  $\mathbf{w}$  given  $\mathbf{c}$ , i.e., the distribution of  $\mathbf{h}|\mathbf{c}$ . As before, the denominator is a normalization constant that is not practically relevant.

*The mask posterior* Bayes' rule allows us to explore further theoretical considerations about the involved mask probabilities. We can study the mask posterior  $p(\mathbf{c}|\mathbf{f})$  in more detail, using

$$p(\mathbf{c}|\mathbf{f}) = \frac{p(\mathbf{f}|\mathbf{c})p(\mathbf{c})}{p(\mathbf{f})}. \quad (42)$$

Now  $p(\mathbf{c})$  models the probability of the mask  $\mathbf{c}$  being generated (irrespective of  $\mathbf{f}$ ) and  $p(\mathbf{f}|\mathbf{c})$  models some measure of the noise and image content in relation to the mask. In practice, ideally the density  $\mathbf{1}^T \mathbb{E}[\mathbf{c}]/N$  should be chosen to be inversely proportional to the standard deviation of the noise. Similarly if we know that the noise distribution is space-variant, or if we suspect that features (e.g., edges) are present, we can choose the local density of  $\mathbf{c}$  to account for that: higher for more prominent edges, lower for higher noise variance. The weight of these choices are modeled by  $p(\mathbf{f}|\mathbf{c})$ . Selecting  $p(\mathbf{c})$  is less trivial, as it needs to match the mask distribution of natural images, i.e., the distribution of natural images from the perspective of the masks used in the inpainting operator. It is simpler to choose it based on the density, i.e.,  $p(\mathbf{c}) = p(\|\mathbf{c}\|_0/N)$ , which makes it blind to spatial variations, or to just choose it as a constant, if we have no data on it. Note that these considerations are meant to provide a different view on the mask posterior and an alternative strategy on

how to construct it. The adaptive mask selection methods that we consider in this work directly induce a mask posterior  $p(\mathbf{c}|\mathbf{f})$  and do not model  $p(\mathbf{f}|\mathbf{c})$  or  $p(\mathbf{c})$ . They are based on strategies from the noise-free case in image inpainting and we adapt and extend them to the noisy case. For all the approaches that we consider, we state their induced PMFs  $p(\mathbf{c}|\mathbf{f})$  (see Eq. (64), Proposition 5, Appendix B).

*On the importance of the inpainting operator* A crucial question is whether an inpainting operator  $\mathbf{r}$  is suitable for modeling natural images in a sparse and robust manner, such that noise can be attenuated by averaging multiple nearby representations of a noisy image from a lower-dimensional image manifold. For  $\mathbf{r}$  being homogeneous diffusion inpainting, we know that it has been used successfully for image compression of natural images with low to medium frequencies [33]. Moreover, we present new results in Sect. 4 that relate the MMSE estimate to homogeneous diffusion denoising. The large body of literature on sparse image approximation and compression should provide a reasonable selection of good inpainting operators  $\mathbf{r}$ . In the current work we also consider biharmonic inpainting (see Sect. 6.4).

*Interplay between the mask PMF and homogeneous diffusion* The basis vectors  $\mathbf{B}_c$  for homogeneous diffusion are generally low-frequent and smooth, with the local frequency depending on the local density of the mask points. For a constant PMF  $p(\mathbf{c}|\mathbf{f})$ , i.e., a homogeneous mask density, we get a process similar to isotropic homogeneous diffusion, and it is in fact approximately equivalent to it, as we demonstrate later in Sects. 4.2 and 4.3. As such it also shares its drawbacks, i.e., smoothing equally over image structures and noise. More sophisticated denoising methods such as space-variant diffusion allow for steering the smoothing away from image structures by relying on a guidance image, e.g., the gradient magnitude  $|\nabla u|$ . Similarly, we may use the PMF  $p(\mathbf{c}|\mathbf{f})$  to guide the denoising. One instance of a PMF that we consider is inspired by a result for mask selection in inpainting. Belhachmi et al. [20] have argued that the local density of an optimal inpainting mask  $\mathbf{c}$  should be proportional to the pixelwise magnitude of the Laplacian  $|\mathbf{L}\mathbf{f}|$ . In our setting this translates to constructing a PMF  $p$  such that  $\mathbb{E}[\mathbf{c}|\mathbf{f}] \sim |\mathbf{L}\mathbf{f}|$ ; see Sect. 5.2.

### 3.1.4 Convergence

A question which arises is how well the estimator  $\langle \mathbf{u} \rangle_n$  approximates the MMSE estimate  $\mathbf{u}^{\text{MMSE}} = \mathbb{E}[\langle \mathbf{u} \rangle_n | \mathbf{f}]$  as a function of the number of samples  $n$ . We consider this scaling behavior in the next proposition.

**Proposition 3** (Convergence of the DbI Estimator) *The root mean square error (RMSE)  $\sqrt{\text{MSE}(\langle \mathbf{u} \rangle_n, \mathbb{E}[\langle \mathbf{u} \rangle_n | \mathbf{f}])}$  between the estimator  $\langle \mathbf{u} \rangle_n$  and its expectation  $\mathbb{E}[\langle \mathbf{u} \rangle_n | \mathbf{f}]$  scales as  $O(n^{-1/2})$ , where  $n$  is the number of sampled masks.*

*Proof* We first recall that we can decompose the MSE between some estimator  $\hat{\theta}$  and some fixed parameter  $\theta$  into a variance and a bias part:

$$\begin{aligned} \text{MSE}(\hat{\theta}, \theta) &= \mathbb{E}[\|\hat{\theta} - \theta\|_2^2] \\ &= \mathbb{E}[\|\hat{\theta} - \mathbb{E}[\hat{\theta}]\|_2^2] + \|\mathbb{E}[\hat{\theta}] - \theta\|_2^2 \\ &= \mathbb{V}[\hat{\theta}] + \text{Bias}(\hat{\theta}, \theta)^2. \end{aligned} \tag{43}$$

If we consider the MSE between the estimator  $\langle \mathbf{u} \rangle_n$  and its expectation  $\mathbb{E}[\langle \mathbf{u} \rangle_n | \mathbf{f}]$ , the bias vanishes and we have  $\text{MSE}(\langle \mathbf{u} \rangle_n, \mathbb{E}[\langle \mathbf{u} \rangle_n | \mathbf{f}]) = \mathbb{V}[\langle \mathbf{u} \rangle_n | \mathbf{f}]$ . The variance  $\mathbb{V}[\langle \mathbf{u} \rangle_n | \mathbf{f}]$  is given by

$$\mathbb{V}[\langle \mathbf{u} \rangle_n | \mathbf{f}] = \mathbb{V} \left[ \frac{1}{n} \sum_{\ell=1}^n \mathbf{r}(\mathbf{c}^\ell, \mathbf{f}) \middle| \mathbf{f} \right] = \frac{1}{n^2} \sum_{\ell=1}^n \mathbb{V}[\mathbf{r}(\mathbf{c}, \mathbf{f}) | \mathbf{f}] = \frac{1}{n} \mathbb{V}[\mathbf{r}(\mathbf{c}, \mathbf{f}) | \mathbf{f}]. \tag{44}$$

The second equality holds because the masks are independent and identically distributed. For a finite variance  $\mathbb{V}[\mathbf{r}(\mathbf{c}, \mathbf{f}) | \mathbf{f}]$ , the root mean square error between the estimator and its expectation thus scales as  $O(n^{-1/2})$ .  $\square$

### 3.1.5 Acceleration by low-discrepancy sequences

When the masks are random variables, as noted in Sect. 3.1.4, we have a somewhat slow convergence of  $O(n^{-1/2})$ . Informally this means that to decrease the RMSE by a factor 4 we would need 16 times as many samples. The natural question arises whether we can do better by trading randomness for a more structured sampling strategy. The answer is positive, as in the context of integration (and our problem can be framed as such with respect to the counting measure), a prominent approach for speeding up convergence is the use of low-discrepancy sequences. These sequences fill up space more uniformly than random sequences. The uniformity is typically quantified using the (star) discrepancy of the sequence. Theoretically, the Koksma–Hlawka inequality [91] allows one to bound the numerical integration error, i.e.,  $\|\langle \mathbf{u} \rangle_n - \mathbb{E}[\langle \mathbf{u} \rangle_n | \mathbf{f}]\|_2$  in our case, by using the product of the discrepancy of the sequence and the variation of the integrand. In practice this usually translates to a convergence that can reach as high as  $O(n^{-1})$  which is much better than the  $O(n^{-1/2})$  convergence for the purely random case. Experimental results illustrating a boost to the convergence in the DbI setting are presented in Sect. 6.2.

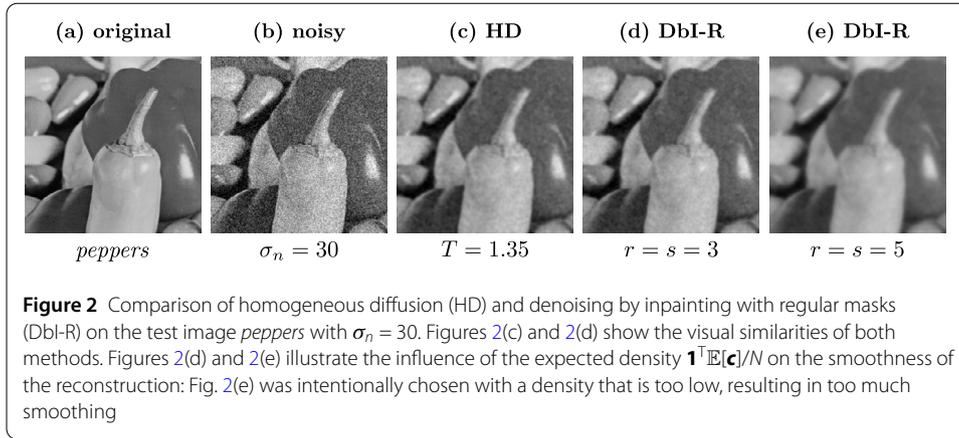
## 4 Linking denoising by inpainting to homogeneous diffusion

The simplest approaches for mask selection in the DbI framework are those, that are independent of the image that is to be filtered ( $p(\mathbf{c} | \mathbf{f}) \equiv p(\mathbf{c})$ ). We consider shifted regular masks as well as randomly selected masks. They are characterized by a spatially flat expectation  $\mathbb{E}[\mathbf{c}] = \text{const}$ . In the following, we briefly introduce regular masks, show how they can be used in the DbI framework and discuss the resulting filtering behavior. Then we derive relations between DbI with regular masks and homogeneous diffusion filtering in 1D. Afterwards, using random masks instead of regular masks, we empirically extend those results to the 2D setting.

### 4.1 Regular masks

Regular masks are created by generating a pattern with each  $r$ th pixel in the  $x$ - and each  $s$ th pixel in the  $y$ -direction being added to the mask. We can then shift such a mask in both directions to obtain multiple masks. If we assume an  $n_x \times n_y$  pixel grid, we can create such a regular mask via

$$c_{i,j} = \begin{cases} 1 & \text{if } i \bmod r = 0 \text{ and } j \bmod s = 0, \\ 0 & \text{else.} \end{cases} \tag{45}$$



We have  $r$  options of shifting this regular mask in the  $x$ -direction and  $s$  options in the  $y$ -direction, adding up to  $n = rs$  total possible configurations. Denoting by  $p \in \{0, \dots, r - 1\}$  and  $q \in \{0, \dots, s - 1\}$  the shift in the  $x$ - and  $y$ -direction, respectively, we can write the shifted masks as

$$c_{i,j}^{ps+q+1} = \begin{cases} 1 & \text{if } i \bmod r = p \text{ and } j \bmod s = q, \\ 0 & \text{else.} \end{cases} \tag{46}$$

Clearly, the created masks are independent of the image. Furthermore, the mask density is constant over the entire image, leading to the same smoothing strength at all locations, solely determined by the total mask density, i.e., by the spacing. If  $r = s$ , this smoothing is equally strong in the  $x$ - and  $y$ -direction. Visually one then observes a smoothing behavior that resembles the one of homogeneous diffusion filtering (see Figs. 2(c) and 2(d)). The influence of the mask density on the smoothing strength can be observed in Figs. 2(d) and 2(e).

The similarity between the methods can not only be observed visually, but also established theoretically. Next we provide a derivation in the 1D case for regular masks relating the diffusion time of homogeneous diffusion to the mask density in DbI.

### 4.2 Mathematical analysis in 1D

We consider a discrete 1D signal  $\mathbf{f}$  and regular inpainting masks with spacing  $r$  and shift  $p \in \{0, \dots, r - 1\}$ . It is known that in 1D, homogeneous diffusion inpainting and linear interpolation are equivalent. Thus, an inpainted pixel at position  $i$  can be described in terms of its two neighboring mask pixels. We denote the distance between the pixel  $i$  and its neighboring mask pixel on the left by  $\ell := |i - p| \bmod r$ , which implies that for mask pixels we have  $\ell = 0$ . Accordingly, the distance to the mask pixel on the right is given by  $r - \ell$ . The interpolated value at pixel  $i$  for mask  $\ell + 1$  is then

$$v_i^{\ell+1} = \frac{r - \ell}{r} f_{i-\ell} + \frac{\ell}{r} f_{i+r-\ell}. \tag{47}$$

To obtain the final result, the inpaintings from the  $r$  shifted masks are averaged. We get

$$\begin{aligned}
 u_i &= \frac{1}{r} \sum_{\ell=0}^{r-1} v_i^{\ell+1} = \frac{1}{r} \left( f_i + \sum_{\ell=1}^{r-1} \frac{r-\ell}{r} f_{i-\ell} + \frac{\ell}{r} f_{i+r-\ell} \right) \\
 &= \frac{1}{r^2} \left( r f_i + \sum_{\ell=1}^{r-1} \ell (f_{i-(r-\ell)} + f_{i+(r-\ell)}) \right),
 \end{aligned}
 \tag{48}$$

where the last line reveals the general form of the filter in dependence of the spacing  $r$ . The filter is given by a hat kernel with central weight  $1/r$  and width  $2r - 1$ . In Theorem 4 we demonstrate that this kernel can be seen as a consistent discretization of  $\partial_t u = \partial_{xx} u$ . Consequently, convolution with such a kernel approximates Gaussian smoothing, which explains the visual similarity of the results in Fig. 2. Since the spacing  $r$  determines the size of the smoothing kernel, we explicitly see the connection between the mask density and the smoothing strength. For the special case of  $r = 2$ , (48) yields

$$u_i = \frac{f_{i-1} + 2f_i + f_{i+1}}{4},
 \tag{49}$$

which is exactly a single step of an explicit scheme for homogeneous diffusion with step size  $T = \frac{1}{4}$  and initial signal  $\mathbf{f}$  (assuming grid size  $h = 1$ ). If we reformulate (48) in a way that resembles an explicit scheme for homogeneous diffusion, we can derive a general connection between the spacing  $r$  (and thus the density) of denoising by inpainting with regular masks and the time step size of such an explicit scheme, which we state in Theorem 4.

**Theorem 4** (Connection between Mask Density and Diffusion Time) *Given the  $r$  shifted regular inpainting masks in 1D, each of density  $d = 1/r$ , denoising by inpainting approximates explicit homogeneous diffusion at time*

$$\boxed{T = \frac{1 - d^2}{12d^2}}.
 \tag{50}$$

*Proof* In (48) we derived the general form of the filter corresponding to denoising by inpainting with regular masks of spacing  $r$  as

$$u_i = \frac{1}{r^2} \left( r f_i + \sum_{\ell=1}^{r-1} \ell (f_{i-(r-\ell)} + f_{i+(r-\ell)}) \right).
 \tag{51}$$

We can rewrite this in the following manner:

$$\begin{aligned}
 u_i &= \frac{1}{r^2} \left( r f_i + \sum_{\ell=1}^{r-1} \ell (f_{i-(r-\ell)} + f_{i+(r-\ell)}) \right) \\
 &= \frac{1}{r^2} \left( r^2 f_i - 2 \sum_{\ell=1}^{r-1} \ell f_i + \sum_{\ell=1}^{r-1} \ell (f_{i-(r-\ell)} + f_{i+(r-\ell)}) \right) \\
 &= f_i + \frac{1}{r^2} \sum_{\ell=1}^{r-1} \ell (f_{i-(r-\ell)} - 2f_i + f_{i+(r-\ell)}),
 \end{aligned}
 \tag{52}$$

where we have used that  $\sum_{\ell=1}^{r-1} \ell = \frac{(r-1)r}{2}$ . Then we may write

$$\begin{aligned} u_i - f_i &= \frac{1}{r^2} \sum_{\ell=1}^{r-1} \ell (f_{i-(r-\ell)} - 2f_i + f_{i+(r-\ell)}) \\ &= \frac{1}{r^2} \sum_{\ell=1}^{r-1} \ell(r-\ell)^2 \frac{f_{i-(r-\ell)} - 2f_i + f_{i+(r-\ell)}}{(r-\ell)^2} \\ &= \sum_{\ell=1}^{r-1} \frac{\ell(r-\ell)^2}{r^2} \frac{f_{i-(r-\ell)} - 2f_i + f_{i+(r-\ell)}}{(r-\ell)^2}. \end{aligned} \tag{53}$$

By approximating  $f_{i\pm(r-\ell)}$  via a Taylor expansion and using the sampling distance  $h$ , we can derive the time step size as

$$\begin{aligned} u_i - f_i &= \sum_{\ell=1}^{r-1} \frac{\ell(r-\ell)^2}{r^2} \frac{f_{i-(r-\ell)} - 2f_i + f_{i+(r-\ell)}}{(r-\ell)^2} \\ &= \sum_{\ell=1}^{r-1} \left( \frac{\ell(r-\ell)^2}{r^2} \right) \left( h^2 d_{xx}f|_i + \frac{h^4(r-\ell)^2}{12} d_{xxxx}f|_i + \mathcal{O}(h^6) \right) \\ &= h^2 \sum_{\ell=1}^{r-1} \left( \frac{\ell(r-\ell)^2}{r^2} \right) (d_{xx}f|_i + \mathcal{O}(h^2)) \\ &\approx h^2 \sum_{\ell=1}^{r-1} \left( \frac{\ell(r-\ell)^2}{r^2} \right) d_{xx}f|_i. \end{aligned} \tag{54}$$

We end up with an approximation of an explicit scheme with time step size

$$T = h^2 \sum_{\ell=1}^{r-1} \frac{\ell(r-\ell)^2}{r^2} = \frac{h^2(r^2-1)}{12}. \tag{55}$$

Using that the density is the inverse of the grid spacing and setting  $h = 1$ , we derive the final relation between  $T$  and the density  $d$ , given by

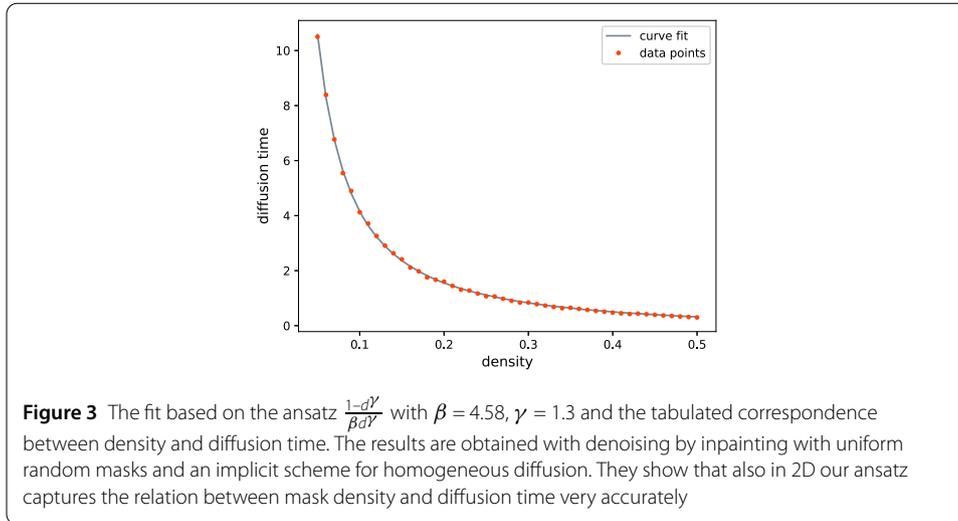
$$T = \frac{1-d^2}{12d^2}. \tag{56}$$

□

### 4.3 Empirical extension to 2D

To derive the relationship to the diffusion time in the 1D case we used the fact that the solution of the Laplace equation with Dirichlet boundaries is given by linear interpolation. That is, we know the closed form of the inpainting echoes in 1D. In 2D a closed form solution for those is not known, however, it may be computed numerically. Thus our goal is to establish a relationship between the diffusion time and the density numerically.

We take as a starting point the ansatz from the 1D case that the diffusion time  $T$  is given as  $\frac{1-d^2}{12d^2}$ , but generalize it to the form  $T \approx \frac{1-d^\gamma}{\beta d^\gamma}$ . Provided that this conjecture is correct, we only need to find the constants  $\beta$  and  $\gamma$ . Since regular masks only allow for a stepwise adaptation of the mask density, they are not well suited for generating a large number of



data points at different densities. Therefore, we use uniform random masks instead, which also have a spatially flat expectation, i.e.,  $\mathbb{E}[\mathbf{c}] = \text{const}$ .

First, we numerically tabulate the relationship between the density and the diffusion time. That is, given a density  $d$  we find the diffusion time  $T(d)$  which minimizes the difference between the filter matrices,

$$T(d) = \operatorname{argmin}_{T \geq 0} \|\mathbf{A}_{DbI}(d) - \mathbf{A}_{HD}(T)\|_F^2. \tag{57}$$

Here  $\|\cdot\|_F$  is the Frobenius norm, and the matrices are the DbI filter matrix resulting from a probability mass function for masks with expected density  $d$ , and the matrix modeling homogeneous diffusion at time  $T$  using an implicit Euler discretization:

$$\mathbf{A}_{DbI}(d) := \mathbb{E}[(\mathbf{C} + (\mathbf{I} - \mathbf{C})\mathbf{L})^{-1}\mathbf{C}], \quad \frac{1}{N}\mathbf{1}^\top \mathbb{E}[\mathbf{c}] = d, \tag{58}$$

$$\mathbf{A}_{HD}(T) := (\mathbf{I} + T\mathbf{L})^{-1}. \tag{59}$$

We estimate  $\mathbf{A}_{DbI}$  using 1024 sampled masks. Then, having the relationship  $d \mapsto T(d)$  we find that  $T(d) \approx \frac{1-d\gamma}{\beta d^\gamma}$  for  $\beta = 4.58$ ,  $\gamma = 1.3$ , which is illustrated in Fig. 3. Note the high quality of the data fit, which confirms the accuracy of the derived relation.

### 5 Spatial optimization for denoising by inpainting

As we have seen in Sect. 4, the use of nonadaptive masks restricts the DbI framework, as it entails a nonadaptive smoothing behavior. Furthermore, our results from Sect. 3.1 emphasize the importance of spatial optimization in the context of image denoising. In [10], an adaptive mask selection approach enables the framework to perform edge-preserving image filtering, although the simple homogeneous diffusion inpainting operator by itself is space-invariant. This approach thus implies a different paradigm for image denoising: *Instead of optimizing the denoising operator, one can optimize the data.* In this section, we will first present the strategy that was proposed in [10]. Then we propose an alternative, simpler approach that eventually gains its power by the application of tonal optimization.

### 5.1 Densification method

Two well-known mask selection strategies from image compression are probabilistic sparsification [39] and densification [42], which build the mask in an iterative way using a top-down and a bottom-up strategy, respectively.

In probabilistic sparsification, we start with a full mask and take away the least important pixels from a number of randomly selected candidates in each iteration. To identify those pixels, we temporarily exclude all candidates from the mask and compute an inpainting. Then the candidate locations with the highest local (i.e., pixelwise) reconstruction error are added back to the mask as they are assumed to be the most important, while the others remain permanently excluded. This process is repeated until the desired mask density is reached. In probabilistic densification, the initial mask is empty and again a number of candidate pixels are selected. Given an inpainting with the current mask (in the first step some pixels have to be chosen at random) we select and add those candidates to the mask that have the highest local reconstruction error.

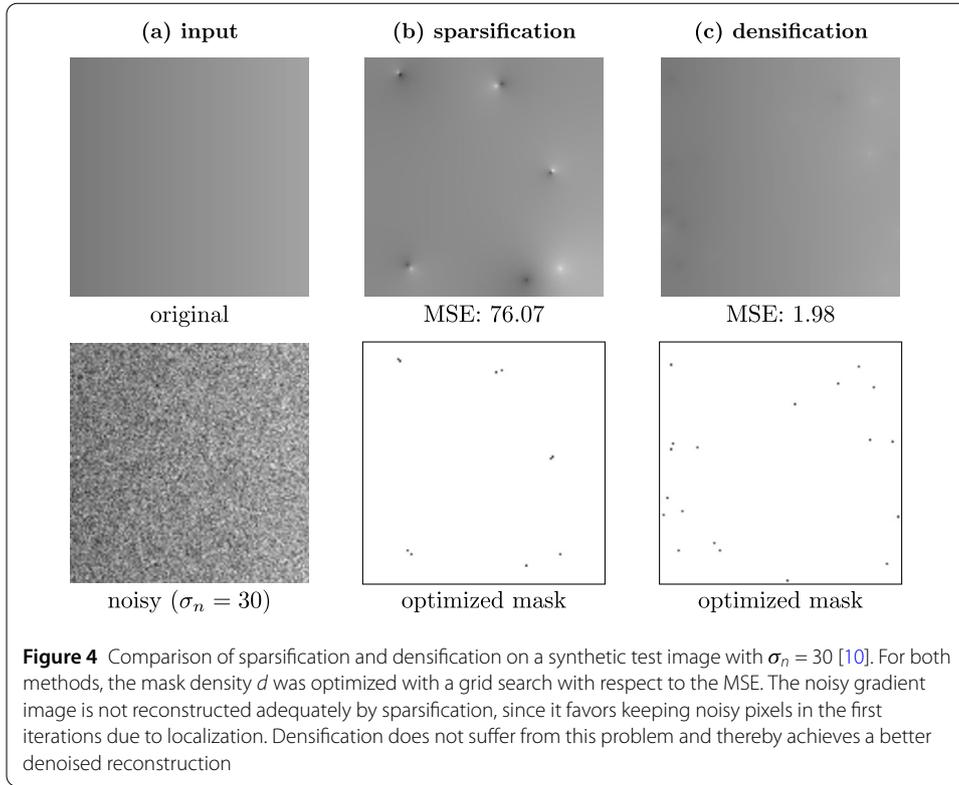
In the noisy setting, special care is required as the pixel selection based on the *local* reconstruction error is not reliable. The local error does not allow the algorithm to distinguish between noise and important image structures, such as edges. If a pixel contains strong noise, this creates a large local error because – just like edges – the noise cannot be reconstructed by the smooth inpainting. Introducing such a noisy pixel into the mask is not desirable. We cure this problem by judging the importance of a pixel based on its effect on the *global* reconstruction error. We do this by calculating a full inpainting for each candidate pixel. While this improves the quality of the selected mask, it drastically increases the run time.

Even though in the noise-free setting densification and sparsification yield results of comparable quality [38], this is different when handling noisy data. For sparsification, we initially have very dense masks. If we exclude candidate pixels from such masks, the reconstructions often only differ at the locations of these pixels. Therefore, sparsification tends to keep noisy pixels in the mask, even when a global reconstruction error is computed. This problem does not occur in probabilistic densification, as for a sparse mask, the candidate pixels have a global influence. The result of this effect is illustrated in Fig. 4. Here, densification is able to select appropriate pixels that lead to an almost perfect result while sparsification fails to reconstruct the image properly.

Thus, we opt for a *probabilistic densification* algorithm based on a *global* error computation, which is described in Algorithm 1 and has been proposed in [10]. An additional advantage of this probabilistic densification method is that it does not only select pixels at useful locations (e.g., close to edges), but also implicitly avoids picking pixels that are too noisy, as they would have a negative impact on the reconstruction quality. The method can be interpreted according to the probabilistic mask generation framework from Sect. 3, and we provide the implied mask probabilities in the following proposition.

**Proposition 5** (Mask Probabilities implied by the Densification Method) *A mask  $\mathbf{c}$  generated by probabilistic densification has the conditional probability density function*

$$p(\mathbf{c}|\mathbf{f}) = \sum_{\sigma \in S_m} p_{\sigma}(\mathbf{c}|\mathbf{f}), \quad (60)$$



where  $m = \|\mathbf{c}\|_0$  is the number of mask pixels, the sum is taken over the group  $S_m$  of permutations of the ordering of the  $m$  mask pixels, and  $p_\sigma(\mathbf{c}|\mathbf{f})$  denotes the probability that the  $m$  mask points were introduced in the order  $\sigma$ . The latter is the product of the probabilities of selecting one mask pixel at each step,

$$p_\sigma(\mathbf{c}|\mathbf{f}) = p_\sigma^m(\mathbf{c}|\mathbf{f}) \cdots p_\sigma^1(\mathbf{c}|\mathbf{f}). \tag{61}$$

The probability of picking the  $k$ th mask pixel (according to the permutation  $\sigma$ ) at step  $k$  has the following form:

$$p_\sigma^k(\mathbf{c}|\mathbf{f}) = \sum_{\beta=1}^{\alpha} \frac{1}{\beta} \frac{\binom{N_{eq}-1}{\beta-1} \binom{N_{gt}}{\alpha-\beta}}{\binom{N-k}{\alpha}}, \tag{62}$$

where  $\alpha$  is the number of candidates considered per step,  $N_{gt}$  is the number of nonmask pixels at step  $k$  that would have resulted in an inpainting with a higher MSE if they were chosen instead of the  $k$ th mask pixel in  $\sigma$ , and  $N_{eq}$  is the number of nonmask pixels that would have resulted in the same MSE.

*Proof* We present the proof of this result in Appendix A. □

### 5.2 Acceleration via the analytic results of Belhachmi et al

As the global error computation in the previous approach requires calculating an inpainting for each candidate pixel, the run time is substantial. Therefore, we propose another

---

**Algorithm 1** Mask densification with global error computation [10]

---

*Input:* Noisy image  $\mathbf{f} \in \mathbb{R}^N$ , number of candidates  $\alpha$ , desired final mask density  $d$ .

*Initialization:* Mask  $\mathbf{c} = 0$  is empty.

*Compute:*

**do**

1. Choose randomly a set  $A \subset \{k \in \{1, \dots, N\} \mid c_k = 0\}$  with  $\alpha$  candidates.

**for all**  $i \in A$  **do**

2. Set temporary mask  $\mathbf{m}^i$  such that  $\forall k \in \{1, \dots, \alpha\} \setminus \{i\} : m_k^i = c_k, m_i^i = 1$ .

3. Compute reconstruction  $\mathbf{u}^i$  from mask  $\mathbf{m}^i$  and image data  $\mathbf{f}$ .

**end for**

4. Set  $\mathbf{c} = \operatorname{argmin}_{\mathbf{m}^i, i \in A} \operatorname{MSE}(\mathbf{u}^i, \mathbf{f})$ . This adds one mask point to  $\mathbf{c}$ .

**while** pixel density of  $\mathbf{c}$  smaller than  $d$ .

*Output:* Mask  $\mathbf{c}$  of density  $d$ .

---

approach, with the goal of a faster mask generation process. We refer to this method as the *analytic method*. It is based on the results of Belhachmi et al. [20]. They have shown that the mask density for homogeneous diffusion inpainting should be proportional to the pointwise magnitude of the Laplacian  $|\mathbf{L}\mathbf{f}|$ . Additionally, they suggest using the Gaussian-smoothed version  $\mathbf{f}_\sigma := \mathbf{K}_\sigma * \mathbf{f}$  of  $\mathbf{f}$  even in the noise-free setting. Here  $\mathbf{K}_\sigma$  is a discrete approximation of a Gaussian with standard deviation  $\sigma$ . This step proves even more beneficial in our setting, since we are calculating the Laplacian of noisy data, and regularizing  $\mathbf{f}$  helps considerably for constructing a reasonable guidance image  $|\mathbf{L}\mathbf{f}_\sigma|$ .

As we require multiple different binary masks for our framework, we sample from  $|\mathbf{L}\mathbf{f}_\sigma|$  by using a simple and fast Poisson sampling. Given a density image  $\mathbf{d} \in [0, 1]^N$ , we can sample a mask according to it by generating a uniform random number  $v_i \sim U[0, 1]$  for each pixel  $i$  and then thresholding at  $d_i$ :

$$c_i = \begin{cases} 1 & \text{if } v_i \leq d_i, \\ 0 & \text{if } v_i > d_i. \end{cases} \tag{63}$$

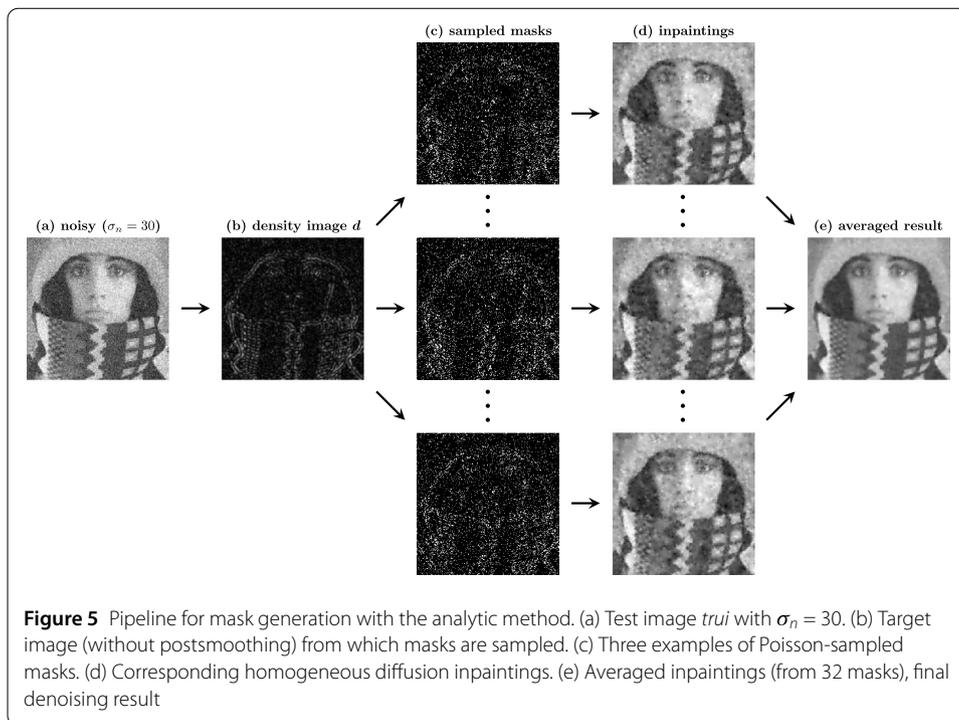
Then the probability mass function  $p_{\mathbf{d}}$  for sampling a mask  $\mathbf{c}$  given the density image  $\mathbf{d}$  is

$$p_{\mathbf{d}}(\mathbf{c}) = \frac{1}{P} \prod_{i=1}^N (d_i)^{c_i} (1 - d_i)^{1-c_i}, \quad P = \sum_{\mathbf{c} \in \{0,1\}^N} \prod_{i=1}^N (d_i)^{c_i} (1 - d_i)^{1-c_i}. \tag{64}$$

By construction the mask would have an expected density equal to the mean value of  $\mathbf{d}$ . In our approach we set the per pixel probabilities to

$$\mathbf{d} = \min \{C|\mathbf{L}\mathbf{f}_\sigma|, \mathbf{1}\}, \tag{65}$$

where the minima are taken pointwise, and  $C$  is a constant chosen such that the mean value of  $\mathbf{d}$  is equal to the desired mask density. Figure 5 shows the pipeline for mask generation with this method. One can observe in Fig. 5(b) that  $\mathbf{d}$  is strongly affected by the noise despite the presmoothing. This is because we calculate second-order derivatives that are even more sensitive to noise. When sampling from this image the mask is drawn towards



noisy pixels. To counteract this, we propose to perform an additional outer smoothing of the probability image  $\mathbf{d}$ , after the absolute value of the Laplacian is taken, thus modifying it to

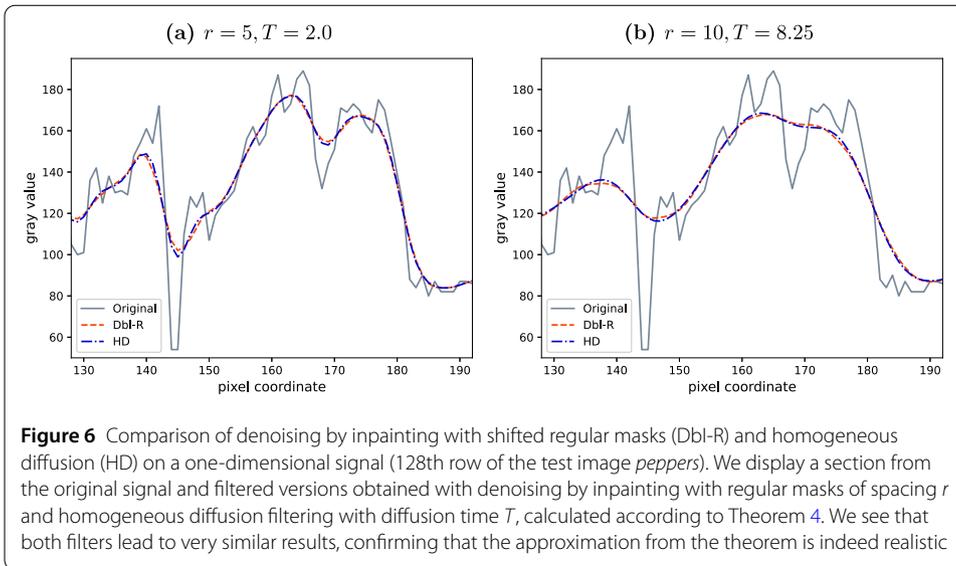
$$\mathbf{d} = \min \{ C (\mathbf{K}_\rho * |\mathbf{L}\mathbf{f}_\sigma|), \mathbf{1} \}, \tag{66}$$

with a postsMOOTHING parameter  $\rho$ . Our proposed selection strategy offers an instant generation of adaptive masks, in a sense that it does not require the calculation of any inpainting. Furthermore, it provides a transparent formulation of the mask PMF (see (64)) and as such exhibits a specifically simple interpretation in the context of our probabilistic framework in Sect. 3.1. On the other hand, contrary to probabilistic densification it does not have a mechanism to avoid noisy mask pixels. To obtain the best possible results, the presMOOTHING parameter  $\sigma$ , the postsMOOTHING parameter  $\rho$ , and the desired mask density have to be optimized depending on the image content and the noise level.

Note that Belhachmi et al. [20] apply Floyd–Steinberg dithering [92], which includes an error diffusion in the binarization process. This strategy can be equipped with a random component in order to generate multiple masks, which makes it an alternative to Poisson sampling for us. We have tested both methods and found that there is no advantage in using Floyd–Steinberg dithering. Thus, we opt for the simple Poisson sampling. Nonetheless, we give the mask probabilities for sampling with error diffusion methods in Appendix B.

### 6 Experiments

In this section, we present our experiments. They evaluate our theories and compare the different DbI strategies in practice. Firstly, we confirm the accuracy of the 1D relation that we derived for DbI with regular mask in Sect. 4.2. We also display the corresponding results in 2D. Next, we show that the theoretical convergence estimates from Sect. 3.1.4



also hold in practice and evaluate the gain through low-discrepancy-based sampling (see Sect. 3.1.5). Furthermore, we assess the spatial and tonal mask optimization approaches. To this end, we compare DbI to PDE-based methods of similar structural complexity. Aside from homogeneous diffusion, we choose linear space-variant diffusion and nonlinear diffusion as representatives of methods that are based on operator optimization. Lastly, we consider the denoising by biharmonic inpainting to further investigate the question of data optimization vs. operator optimization.

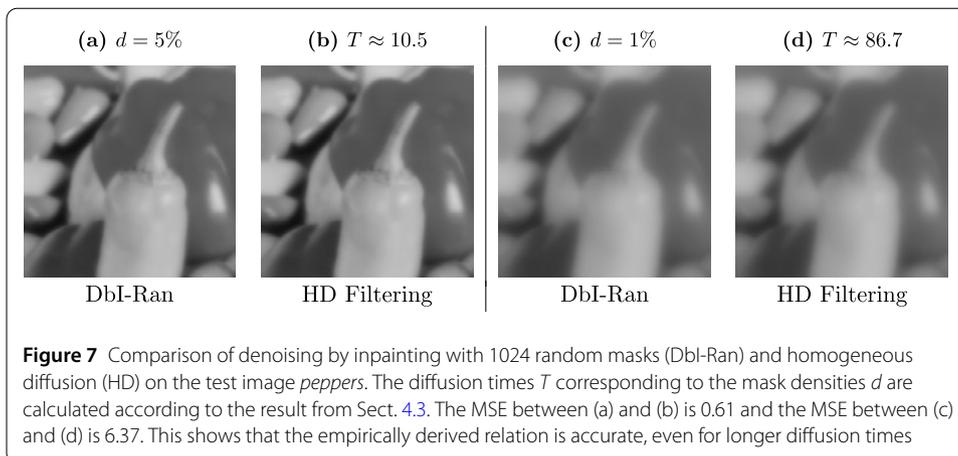
### 6.1 Relation between DbI and homogeneous diffusion

In Sect. 4.2 we derived a relation between the mask density  $d$  and the diffusion time  $T$ , given by  $T = (1 - d^2)/(12d^2)$ . To confirm that this relation allows for a good estimate of the diffusion time in practice, we perform an experiment on a 1D signal, which is generated by extracting the 128th row of the *peppers* test image. Homogeneous diffusion is implemented using explicit Euler and the spatial discretization from (49) with the number of iterations chosen such that the desired diffusion time  $T$  is reached. The result in Fig. 6 demonstrates that the diffusion time obtained via Theorem 4 is a good approximation.

In Sect. 4.3 we extended this relation to 2D, yielding  $T = (1 - d^\gamma)/(\beta d^\gamma)$  with  $\beta = 4.58$  and  $\gamma = 1.3$ . To confirm this, we now consider the 2D *peppers* test image. We perform denoising by inpainting with 1024 randomly selected masks, as well as homogeneous diffusion filtering with the diffusion time calculated according to the above relation and compare the results. The experiments in Fig. 7 visually and qualitatively confirm the accuracy of the relation in 2D.

### 6.2 Convergence

As we have shown in Sect. 3.1.4 the estimator converges to its expectation at a rate of  $O(n^{-1/2})$  with respect to the RMSE. In Sect. 3.1.5 we introduced the idea of using low-discrepancy sequences. Theoretically, they should lead to much faster convergence, thus here we test whether this also holds in practice. In the experiments we again use the  $256 \times 256$  test image *peppers*. We use two sampling strategies for the masks, whose sample means  $\mathbf{c} = \frac{1}{n} \sum_{\ell=1}^n \mathbf{c}^\ell$  converge to the same expectation  $\mathbb{E}[\mathbf{c}|\mathbf{f}]$ . As a representative of a

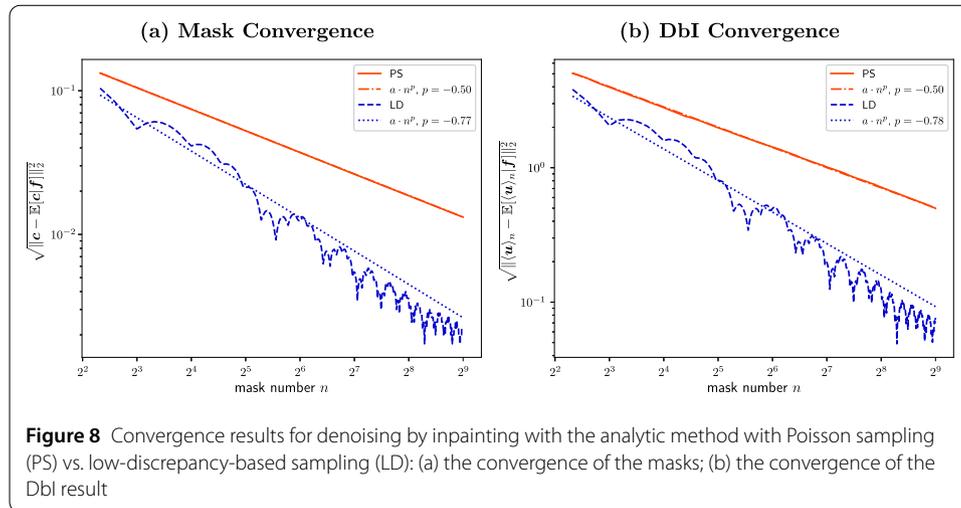


low-discrepancy sequence we use the R2 sequence [93] to create a sampling threshold in each pixel (see [93] for details). This leads to a more regular sampling pattern compared to using a purely random threshold. To make the experiment relevant to realistic scenarios, we use the analytic mask selection method from Sect. 5.2. We first test the mask convergence. To this end, we create  $2^{16} = 65,536$  masks via Poisson sampling and consider their average as converged to the expectation  $\mathbb{E}[\mathbf{c}|\mathbf{f}]$ . Then we sample masks with both sampling strategies and observe how the RMSE between sample mean and expectation evolves with  $n$ . Of course, we are more interested in the convergence of the DbI result  $\langle \mathbf{u} \rangle_n$ . Therefore, following a similar approach as for the masks, we create an individual “converged” DbI result for the two sampling methods, and again consider the RMSE between  $\langle \mathbf{u} \rangle_n$  and the respective reference images. Figure 8 shows that the simple Poisson sampling leads to a convergence rate of  $O(n^{-1/2})$  for the masks as well as for the DbI result, which is perfectly in line with the theory from Sect. 3.1.4. Through low-discrepancy sampling this rate approaches  $O(n^{-1})$ . By fitting a curve through the data, we get a convergence rate of  $O(n^{-0.77})$  for the masks and  $O(n^{-0.78})$  for the DbI result. The  $O(n^{-1})$  estimate is typically achieved for low dimensions, so the difference of our results can be explained by the high dimensionality of our sampling problem. The experiments confirm that the sampling strategy based on low-discrepancy sequences is indeed able to improve the convergence in practice.

### 6.3 Data optimization for denoising by inpainting

In the next step, we investigate the edge-preserving filtering behavior achieved by the use of adaptive masks. We first test the two spatial optimization methods and compare the results to classical diffusion models. We show that DbI can yield results comparable to certain space-variant diffusion methods. Then we discuss the effect of tonal optimization in the DbI setting. It should be noted that these experiments are meant to provide an illustration of the mask optimization strategies and not to achieve the best denoising quality. As we have shown, these strategies can be applied in a more general setting than DbI with homogeneous diffusion inpainting. They are valid for the general probabilistic framework from Sect. 3.1, and as such they also extend to more complex operators (including non-linear ones).

We perform experiments on the three standard test images *trui*, *peppers*, and *walter* with a resolution of  $256 \times 256$ , that are corrupted with additive Gaussian noise with standard deviations  $\sigma_n \in \{10, 20, 30\}$  that we do not clip. To ensure a fair comparison, we optimize



the mask density and if required the pre- and postsmoothing parameter for the denoising by inpainting methods with respect to the MSE to the original image. We do this individually for each image and for each noise level using a grid search. In practice, these parameters need to be adapted to the noise level and the image content. We create 32 masks with each of the mask selection methods, except for the regular masks where the number is determined by the spacing and thus by the density. For the proposed probabilistic densification algorithm we set the number of candidate pixels per iteration to 16.

### 6.3.1 Spatial optimization

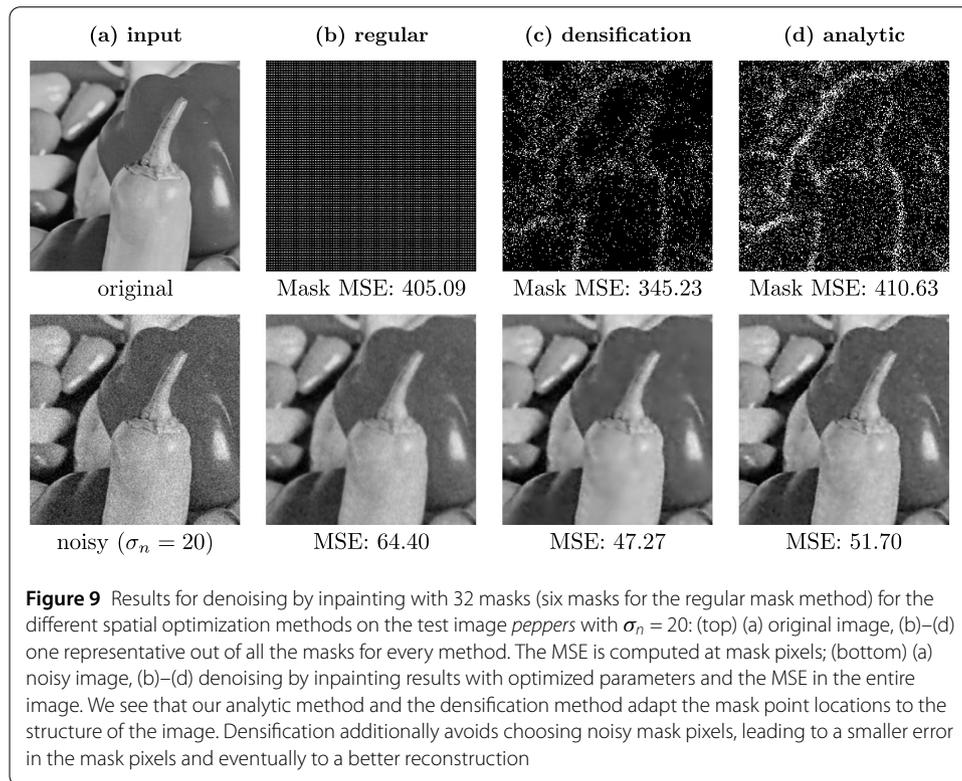
Firstly, we investigate the different spatial selection strategies proposed in Sect. 5 and compare the denoising results with the standard diffusion methods presented in Sect. 2.1. For the diffusion methods, which we discretize with an explicit scheme, we optimize the stopping time and if required the contrast parameter of the Charbonnier diffusivity [88].

As can be seen in Table 1, inpainting with regular masks leads to unsatisfying results, slightly worse than those obtained with homogeneous diffusion filtering. This is expected given the connections derived in Sect. 4.1. Note that the stopping time in homogeneous diffusion filtering can be tuned continuously, while the spacing of the regular mask can only be adapted in integer steps. The analytic method based on Poisson sampling of the smoothed Laplacian magnitude improves the results, especially at lower noise levels. Figure 9(c) shows how the mask pixels accumulate around important image structures, enabling an edge-preserving filtering behavior. The densification method is able to further improve those results. The reason for this improvement can be seen in Fig. 9(d). On top of selecting pixels at reasonable positions, the error in the mask is reduced drastically in comparison to the analytic method, because densification implicitly avoids noisy pixels. The adaptive mask selection strategies enable the denoising by inpainting method to produce results that are comparable to linear space-variant diffusion filtering. However, it cannot reach the quality of nonlinear diffusion. This is not surprising, as a feedback mechanism throughout the inpainting process is missing. Nonetheless, the results reveal that proper data optimization enables DbI to compete with methods that optimize the operator, if they are of comparable complexity.

Although qualitatively the densification approach is better than the analytic method, its required run time is orders of magnitude larger, and this only gets worse for images of

**Table 1** Results (MSE) for denoising by inpainting with regular masks, the densification method and the analytic method with 32 masks (fewer masks for the regular mask method). Comparison to classical diffusion-based denoising methods

		trui			peppers			walter		
noise level $\sigma_n$		10	20	30	10	20	30	10	20	30
Dbl	regular	27.30	57.29	86.46	35.31	64.40	91.79	22.63	50.13	79.16
	densification	<b>19.34</b>	<b>42.72</b>	<b>68.01</b>	<b>24.36</b>	<b>47.27</b>	<b>69.89</b>	<b>13.40</b>	<b>29.65</b>	<b>47.65</b>
	analytic	21.49	49.71	79.79	25.14	51.70	79.91	16.41	37.83	62.08
Diff	homogeneous	24.12	50.18	76.12	32.16	59.77	84.58	19.65	42.76	66.87
	lin. space-var.	17.89	42.62	69.57	24.03	47.47	72.67	13.31	32.30	55.37
	nonlinear	<b>16.21</b>	<b>34.99</b>	<b>54.66</b>	<b>22.63</b>	<b>40.48</b>	<b>57.54</b>	<b>11.89</b>	<b>25.31</b>	<b>39.49</b>



higher resolution. Due to the required number of inpaintings, the densification method takes about an hour to create a single mask with 10% density for our  $256 \times 256$  pixel test images. In contrast, the analytic and the regular approaches allow instant mask generation in approximately a millisecond. Thus, the analytic method yields a reasonable spatial mask pixel distribution in a very short time and clearly has potential, if the error in the mask pixels can be reduced. We show next that this can be achieved by complementing the mask selection strategies with tonal optimization.

### 6.3.2 Tonal optimization

As mentioned in Sect. 3.1.1, tonal optimization leads to an MMSE estimate that is approximating instead of interpolating. If one assumes that mask pixels are erroneous due to the noise, this is certainly a desirable behavior. We will evaluate its effect in the following. To this end, we apply tonal optimization to the masks obtained by each of our spatial opti-

**Table 2** Results (MSE) for denoising by inpainting with regular masks, the densification method and the analytic method with 32 masks (less masks for the regular mask method) including tonal optimization. Comparison to classical diffusion-based denoising methods

		trui			peppers			walter		
		10	20	30	10	20	30	10	20	30
Dbl	regular	22.32	48.77	76.06	32.65	60.92	87.05	16.36	39.54	64.45
	densification	18.46	41.56	67.72	24.42	47.28	70.21	12.35	28.13	45.92
	analytic	<b>17.24</b>	<b>39.49</b>	<b>63.17</b>	<b>23.68</b>	<b>46.43</b>	<b>68.55</b>	<b>12.08</b>	<b>27.66</b>	<b>45.36</b>
Diff	homogeneous	24.12	50.18	76.12	32.16	59.77	84.58	19.65	42.76	66.87
	lin. space-var.	17.89	42.62	69.57	24.03	47.47	72.67	13.31	32.30	55.37
	nonlinear	<b>16.21</b>	<b>34.99</b>	<b>54.66</b>	<b>22.63</b>	<b>40.48</b>	<b>57.54</b>	<b>11.89</b>	<b>25.31</b>	<b>39.49</b>

mization methods. We optimize the tonal values for each individual mask, before once again averaging the respective inpaintings to obtain the final denoised result.

The results in Table 2 reveal that the methods that do not consider the noise in the selection process get the greatest boost in performance. This confirms the conjecture that tonal optimization is able to mitigate the negative effect of noisy mask pixels selection. We also observe that tonal optimization decreases the error in the mask pixels for those methods. In Fig. 9, the MSE at mask locations decreases from 405.09 to 271.80 for the regular mask, and from 410.63 to 306.62 for the analytic method. For probabilistic densification, tonal optimization barely changes the final results, as well as the mask MSE (which even increases slightly from 345.23 to 356.12 in the example).

We see that tonal optimization enables the analytic method to produce results of quality comparable to those of the densification method, and of better quality than space-variant diffusion. Although the tonal optimization step takes some additional seconds, the analytic method is still orders of magnitude faster than the densification method. Figure 10 shows a selection of resulting images comparing the two adaptive mask selection methods with tonal optimization and linear space-variant diffusion, as the diffusion method that leads to the most similar results.

### 6.4 Denoising by biharmonic inpainting

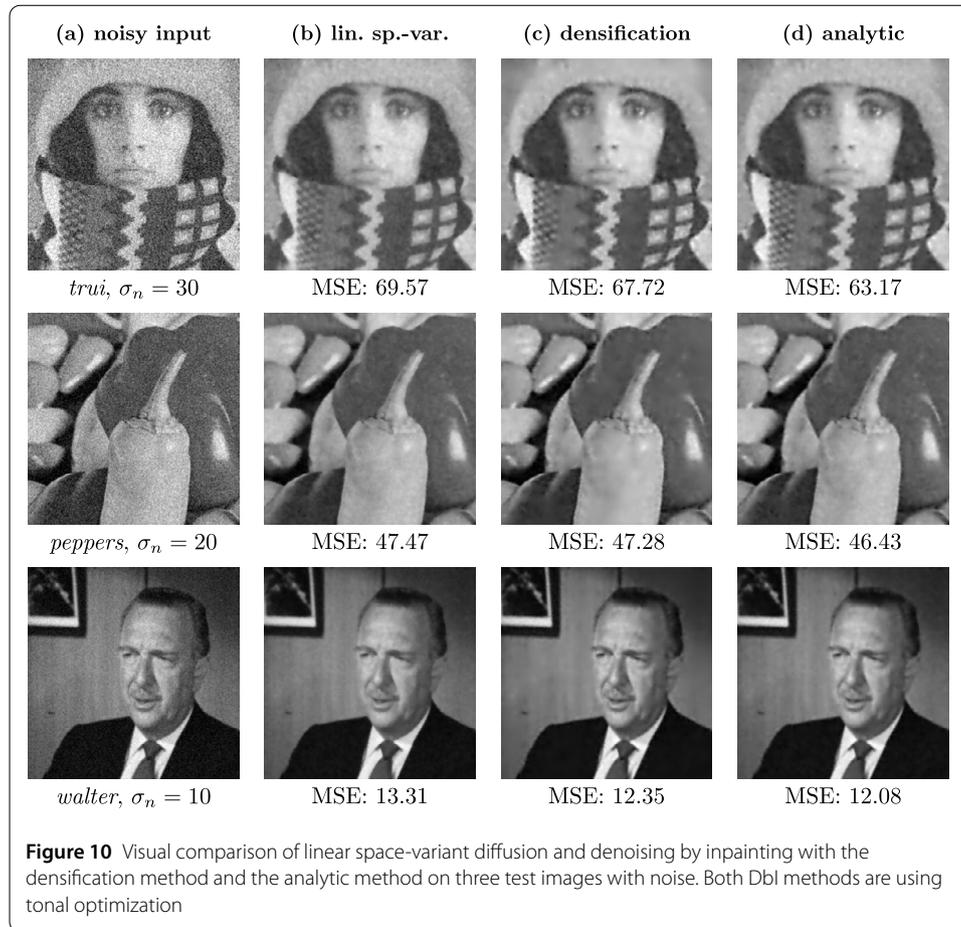
Our previous results reveal that optimizing the data instead of the operator constitutes an interesting alternative for image denoising. To further substantiate this idea, we now adapt the inpainting operator within the DbI framework. We consider biharmonic inpainting as a representative of a higher-order polyharmonic operator.

It has been shown that the biharmonic operator can have quality advantages over homogeneous diffusion (i.e., the harmonic operator) in classical sparse inpainting [18, 36, 49]. Biharmonic inpainting is given by the PDE

$$(c(\mathbf{x}) + (1 - c(\mathbf{x}))\Delta^2)u(\mathbf{x}) = c(\mathbf{x})f(\mathbf{x}) \quad \text{for } \mathbf{x} \in \Omega, \tag{67}$$

with  $\Delta^2 u = \partial_{xxxx}u + 2\partial_{xxyy}u + \partial_{yyyy}u$  and reflecting boundary conditions  $\partial_n u(\mathbf{x}) = 0$  and  $\partial_n \Delta u(\mathbf{x}) = 0$  for  $\mathbf{x} \in \partial\Omega$ . It can be derived from the following variational formulation (analogously to (12)):

$$\min_u \int_{\Omega} (\Delta u(\mathbf{x}))^2 d\mathbf{x}, \text{ such that } u(\mathbf{x}) = f(\mathbf{x}) \text{ for } \mathbf{x} \in K. \tag{68}$$



This shows that biharmonic inpainting penalizes second-order derivatives. Biharmonic inpainting does not suffer from the typical singularities at mask points that homogeneous diffusion inpainting produces. On the other hand it can produce over- and undershoots, since it does not guarantee a maximum–minimum principle. We evaluate the potential of biharmonic inpainting for denoising by comparing it to homogeneous diffusion inpainting. To ensure that the results reflect the quality of the operators, we first perform the experiment on fully random masks.

Our results in Table 3 show that biharmonic inpainting does lead to an improvement, and it is largest at low noise levels. This is to be expected, as the method is not as radical as homogeneous diffusion inpainting, since it penalizes second- instead of first-degree derivatives. However, already tonal optimization as a first data optimization step neutralizes this advantage and the two methods perform similarly. These results support our reasoning that data optimization plays a significant role for the denoising abilities of our framework, being more important than the use of more complex, higher-order models. Further experiments on spatially optimized masks (see Table 4) confirm our findings, and even shift the advantage towards homogeneous diffusion inpainting. When comparing to previous results from classical sparse image inpainting, one has to consider that the singularities, that homogeneous diffusion inpainting suffers from, are suppressed by the averaging in the Dbl framework. Thus, this disadvantage of homogeneous diffusion inpainting does not come into play in our scenario. Lastly, one should keep in mind that

**Table 3** Results (MSE) for denoising by inpainting with 32 random masks using homogeneous diffusion (HD) and biharmonic (BI) inpainting, without and with tonal optimization (TO)

noise level $\sigma_n$	trui			peppers			walter		
	10	20	30	10	20	30	10	20	30
HD, without TO	30.53	65.51	100.51	36.01	71.20	104.21	26.93	60.21	94.37
BI, without TO	<b>24.23</b>	<b>56.37</b>	<b>93.51</b>	<b>33.28</b>	<b>66.61</b>	<b>102.49</b>	<b>19.16</b>	<b>48.19</b>	<b>82.92</b>
HD, with TO	23.10	49.83	<b>76.57</b>	<b>31.98</b>	<b>59.75</b>	<b>85.28</b>	18.09	41.79	66.40
BI, with TO	<b>22.21</b>	<b>49.74</b>	77.25	33.27	61.84	87.92	<b>16.52</b>	<b>39.53</b>	<b>65.03</b>

**Table 4** Results (MSE) for denoising by inpainting with 32 masks obtained with the analytic method using homogeneous diffusion (HD) and biharmonic (BI) inpainting, without and with tonal optimization (TO)

noise level $\sigma_n$	trui			peppers			walter		
	10	20	30	10	20	30	10	20	30
HD, without TO	21.49	49.71	<b>79.79</b>	<b>25.14</b>	<b>51.70</b>	<b>79.91</b>	16.41	37.83	<b>62.08</b>
BI, without TO	<b>19.01</b>	<b>47.47</b>	82.39	25.83	55.42	90.28	<b>14.16</b>	<b>37.15</b>	68.25
HD, with TO	17.24	<b>39.49</b>	<b>63.17</b>	<b>23.68</b>	<b>46.43</b>	<b>68.55</b>	12.08	27.66	<b>45.36</b>
BI, with TO	<b>17.18</b>	40.45	66.13	25.35	49.27	72.68	<b>11.74</b>	<b>27.22</b>	45.70

biharmonic inpainting leads to a higher condition number of the inpainting matrix, and consequently each inpainting is numerically more burdensome and less efficient.

## 7 Conclusions

Our work is the first that links the tasks of PDE-based image inpainting and denoising in a systematic way, by providing an explicit connection between homogeneous diffusion inpainting and denoising through a relation between the diffusion time and the mask density. Our *denoising by inpainting (DbI)* framework achieves denoising by averaging inpainting results with different sparse masks of the same density. It constitutes a means to investigate the connections between PDE-based denoising and inpainting and allows us to evaluate the denoising potential of PDE-based inpainting methods. We have established a probabilistic theory with convergence estimates for the framework, and have extended it to a deterministic version by the use of low-discrepancy sequences. We have further shown that this framework computes an approximation to an MMSE estimate. For nonadaptive masks, we have linked the framework to classical diffusion via a one-to-one relationship between the mask density and the diffusion time. We have demonstrated that a simple operator can exhibit space-variant filtering behavior, when supplemented with adaptive data selection strategies. Experiments with a higher-order inpainting operator, which can be more powerful than homogeneous diffusion inpainting [18, 36, 49], have underlined the importance of choosing appropriate data over more complex operators. For data optimization specific to denoising by inpainting, we have presented two distinct, fundamental strategies. The densification method from our conference paper [10] aims at finding pixels that represent the data well. Thereby, it implicitly avoids the selection of noisy mask pixels during spatial optimization. On the contrary, we have proposed a new approach, where the selection of noisy pixels is tolerated in the spatial optimization but is compensated for by the tonal optimization.

Our work constitutes an unconventional, new viewpoint on image denoising. By using a simple inpainting operator but focusing on adequate data selection, we *shift the priority*

from optimizing the filter model to optimizing the considered data. Moreover, our densification strategy allows us to find the most trustworthy pixels in the data. This shows that simple filter operators such as homogeneous diffusion can give deep insights into data. Last but not least, we have seen that the filling-in effect is not only useful in variational optic flow models and in PDE-based inpainting, but also in denoising. This emphasizes its fundamental role in digital image analysis, which is in full agreement with classical results from biological vision [94].

While our focus in the present paper is on gaining fundamental insights into the potential of inpainting ideas for denoising, our future work will deal with various modifications to make these ideas also applicable to more recent denoising methods. To this end, we are going to consider more sophisticated inpainting operators [49] and data selection strategies [95], including neural ones [96], and the incorporation of more advanced types of data [97]. Such future work should also extend our theory to, e.g., space-variant and non-linear operators.

### Appendix A: Proof of Proposition 5

We derive the expression for the stated probability in Proposition 5 for step  $k + 1$  here. At the beginning of step  $k + 1$ , Algorithm 1 has already inserted  $k$  mask pixels yielding the mask  $\mathbf{c}^k$ . At the end of step  $k + 1$ , we want to have inserted a new mask pixel that is not in  $\mathbf{c}^k$ . Consequently, we can select a pixel only from the set  $\mathcal{I}^k$  of remaining empty mask pixel locations, with  $|\mathcal{I}^k| = N - k$ . The algorithm samples a set  $\mathcal{X}$  of  $\alpha$  distinct candidates from  $\mathcal{I}^k$  uniformly at random (there are  $C_\alpha^{N-k}$  different ways to do so):

$$\mathcal{X} = \{X_1, \dots, X_\alpha \in \mathcal{I}^k : X_i \neq X_j \text{ for } i \neq j\}. \tag{69}$$

Then one chooses the candidate  $X^* \in \mathcal{X}$  with lowest reconstruction error (with respect to the noisy image  $\mathbf{f}$ ):

$$X^* \in \mathcal{X}^* = \operatorname{argmin}_{X \in \mathcal{X}} E^k(X), \quad E^k(X) := \|\mathbf{r}(\mathbf{c}^k + \mathbf{e}_X, \mathbf{f}) - \mathbf{f}\|_2^2, \tag{70}$$

where  $\mathbf{e}_X \in \mathbb{R}^N$  is the zero vector modified with a one at the location corresponding to mask point  $X$ . The minimizer does not have to be unique; in fact the set of minimizers

$$\mathcal{X}^* = \{X_i \in \mathcal{X} : E^k(X_i) = \min_{X \in \mathcal{X}} E^k(X)\} \tag{71}$$

may have more than one element ( $|\mathcal{X}^*| > 1$ ) in which case we choose  $X^*$  uniformly at random from  $\mathcal{X}^*$  with probability  $\frac{1}{|\mathcal{X}^*|}$ . This completes step  $k + 1$ , now with a specific  $X^* = x^*$  and corresponding mask  $\mathbf{c}^{k+1} = \mathbf{c}^k + \mathbf{e}_{x^*}$ . If the desired number of mask points have been achieved the algorithm ends, otherwise one proceeds to step  $k + 2$  in the exact same manner.

After we have inserted mask pixel  $x^* \in \mathcal{I}^k$ , we want to be able to compute the probability  $\Pr(X^* = x^*)$  of this occurring. This is equal to the probability of  $x^*$  having been selected as a candidate,

$$\Pr(x^* \in \mathcal{X}) = C_{\alpha-1}^{N-k-1} / C_\alpha^{N-k} = \frac{\alpha}{N - k}, \tag{72}$$

multiplied by the probability  $\Pr(x^* \in \mathcal{X}^* | x^* \in \mathcal{X})$  that  $x^*$  ends up in  $\mathcal{X}^*$ , which is in turn multiplied by the probability  $\Pr(X^* = x^* | x^* \in \mathcal{X}^*) = \frac{1}{|\mathcal{X}^*|}$  of having picked  $x^*$  from  $\mathcal{X}^*$  uniformly at random. We thus have the following chain of conditional probabilities:

$$\begin{aligned} \Pr(X^* = x^*) &= \Pr(X^* = x^* | x^* \in \mathcal{X}^*) \Pr(x^* \in \mathcal{X}^*) \\ &= \Pr(X^* = x^* | x^* \in \mathcal{X}^*) \Pr(x^* \in \mathcal{X}^* | x^* \in \mathcal{X}) \Pr(x^* \in \mathcal{X}) \\ &= \frac{1}{|\mathcal{X}^*|} \Pr(x^* \in \mathcal{X}^* | x^* \in \mathcal{X}) \frac{\alpha}{N-k}. \end{aligned} \tag{73}$$

We can write the terms involving  $\mathcal{X}^*$  in the following manner:

$$\frac{1}{|\mathcal{X}^*|} \Pr(x^* \in \mathcal{X}^* | x^* \in \mathcal{X}) = \sum_{\beta=1}^{\alpha} \frac{1}{\beta} \Pr(x^* \in \mathcal{X}^* \wedge |\mathcal{X}^*| = \beta | x^* \in \mathcal{X}). \tag{74}$$

The probability on the right-hand side can be rewritten as requiring  $\beta$  of the candidates to have energy equal to  $E^k(x^*)$  and the remaining  $\alpha - \beta$  having a strictly larger energy:

$$\begin{aligned} &\Pr(x^* \in \mathcal{X}^* \wedge |\mathcal{X}^*| = \beta | x^* \in \mathcal{X}) \\ &= \Pr \left( \left( \bigwedge_{i=1}^{\beta} E^k(X_i) = E^k(x^*) \right) \wedge \left( \bigwedge_{j=\beta+1}^{\alpha} E^k(X_j) > E^k(x^*) \right) \middle| x^* \in \mathcal{X} \right). \end{aligned} \tag{75}$$

To compute the above probabilities, we would need to know the total number of pixels from  $\mathcal{I}^k$  with energy equal to  $E^k(x^*)$ ,

$$N_{eq} := |\{x \in \mathcal{I}^k : E^k(x) = E^k(x^*)\}|, \tag{76}$$

and the total number of pixels from  $\mathcal{I}^k$  having a strictly higher energy,

$$N_{gt} := |\{x \in \mathcal{I}^k : E^k(x) > E^k(x^*)\}|. \tag{77}$$

From the requirement  $|\mathcal{X}^*| = \beta$ , it follows that we need to choose  $\beta$  pixels that have energy equal to  $E^k(x^*)$ . However,  $x^* \in \mathcal{X}$  so  $E^k(X) = E^k(x^*)$  with probability 1 for at least one candidate  $X = x^*$ . Then  $\beta - 1$  elements  $X_i$  remain to be selected from  $N_{eq} - 1$  locations, the total number of possibilities being  $C_{\beta-1}^{N_{eq}-1}$ . Finally, the remaining  $\alpha - \beta$  candidates must be selected from  $N_{gt}$  locations, resulting in  $C_{\alpha-\beta}^{N_{gt}}$  options. Using this we can compute the probability

$$\frac{1}{|\mathcal{X}^*|} \Pr(x^* \in \mathcal{X}^* | x^* \in \mathcal{X}) = \sum_{\beta=1}^{\alpha} \frac{1}{\beta} \frac{C_{\beta-1}^{N_{eq}-1} C_{\alpha-\beta}^{N_{gt}}}{C_{\alpha-1}^{N-k-1}}. \tag{78}$$

Ultimately, we get the following probability for step  $k + 1$ :

$$\Pr(X^* = x^*) = \frac{\alpha}{N-k} \sum_{\beta=1}^{\alpha} \frac{1}{\beta} \frac{C_{\beta-1}^{N_{eq}-1} C_{\alpha-\beta}^{N_{gt}}}{C_{\alpha-1}^{N-k-1}} = \sum_{\beta=1}^{\alpha} \frac{1}{\beta} \frac{C_{\beta-1}^{N_{eq}-1} C_{\alpha-\beta}^{N_{gt}}}{C_{\alpha}^{N-k}}. \tag{79}$$

Through the probabilistic densification procedure, the exact same mask  $\mathbf{c}$ , with  $\|\mathbf{c}\|_0$  mask pixels, can be constructed in  $\|\mathbf{c}\|_0!$  different ways (the same set of mask pixels being introduced in all possible orders). That is, we get the probability mass function  $p_\sigma(\mathbf{c}|\mathbf{f})$  over masks that also retain the order of insertion of their mask pixels (e.g., we can modify  $\mathbf{c}$  by setting entries equal to one, to be equal to  $k -$  the step in which those were inserted). To get the usual probability mass function over binary masks we need to sum up the above probabilities over all  $\|\mathbf{c}\|_0!$  permutations of point insertion orders. The main issue for practicality is that  $N_{eq}$  and  $N_{gt}$  must be known, which would require evaluating all possible  $|\mathcal{I}^k| = N - k$  inpaintings for a single step. Nevertheless, Monte Carlo can be used to estimate the probabilities.

**Appendix B: Probability for error diffusion masks**

Error diffusion half-toning (e.g., Floyd–Steinberg dithering [92]) can be used to produce a binary mask  $\mathbf{c} \in \{0, 1\}^N$  from a continuous density image  $\mathbf{d} \in [0, 1]^N$ . The process involves iterating over the image pixels (e.g., in serpentine order), binarizing a single pixel at a given step, and then diffusing the error arising from the binarization to the set of currently nonvisited pixels. This results in a sequence of images  $\mathbf{d} = \mathbf{d}^1, \mathbf{d}^2, \dots, \mathbf{d}^{N+1} = \mathbf{c}$ . The binarization happens according to a thresholding step, which usually reads as

$$c_k = d_k^{k+1} = \begin{cases} 0 & \text{for } d_k^k < 0.5, \\ 1 & \text{for } d_k^k \geq 0.5. \end{cases} \tag{80}$$

Since we want to get multiple masks stochastically, we randomize the process by sampling a uniform random number  $v_k \in [0, 1]$  for pixel  $k$ , and then perform thresholding:

$$c_k = d_k^{k+1} = \begin{cases} 0 & \text{for } d_k^k < v_k, \\ 1 & \text{for } d_k^k \geq v_k. \end{cases} \tag{81}$$

Then the probability mass function for mask  $\mathbf{c}$  constructed from density image  $\mathbf{d}$  is

$$p_{\mathbf{d}}(\mathbf{c}) = \frac{1}{P} \prod_{k=1}^N (d_k^k(\mathbf{c}))^{c_k} (1 - d_k^k(\mathbf{c}))^{1-c_k}, \quad P = \sum_{\mathbf{c} \in \{0,1\}^N} \prod_{k=1}^N (d_k^k(\mathbf{c}))^{c_k} (1 - d_k^k(\mathbf{c}))^{1-c_k}. \tag{82}$$

In the above  $d_k^k(\mathbf{c})$  are assumed to be clamped to  $[0, 1]$ . Note that while this bears similarity to Poisson sampling (see (64)), the probability  $d_k^k(\mathbf{c})$  is conditioned on the probabilities in the  $k$  previous steps. Algorithmically, it is trivial to compute the numerator of the probability during the error diffusion process.

**Abbreviations**

Dbl, Denoising by Inpainting; DCT, Discrete Cosine Transform; EED, Edge-Enhancing Diffusion; MAP, Maximum A Posteriori; ML, Maximum Likelihood; MMSE, Minimum Mean Squared Error; MSE, Mean Squared Error; PDE, Partial Differential Equation; PDF, Probability Density Function; PMF, Probability Mass Function; RMSE, Root Mean Square Error.

**Acknowledgements**

Not applicable.

**Author contributions**

DG conducted experiments and provided theory. VC provided mathematical theory. DG and VC wrote the paper, with feedback from JW and PP. RDA, PP, and JW jointly contributed to the conference paper. The authors have read and approved the final manuscript.

### Funding information

Open Access funding enabled and organized by Projekt DEAL. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 741215, ERC Advanced Grant INCOVID).

### Data availability

The datasets used and/or analyzed during the current study are available from the corresponding author on reasonable request.

## Declarations

### Competing interests

The authors declare no competing interests.

Received: 6 September 2024 Accepted: 11 March 2025 Published online: 28 March 2025

### References

1. Bertalmio, M., Sapiro, G., Caselles, V., Ballester, C.: Image inpainting. In: Proc. SIGGRAPH 2000, New Orleans, pp. 417–424 (2000)
2. Efros, A.A., Leung, T.K.: Texture synthesis by non-parametric sampling. In: Proc. Seventh IEEE International Conference on Computer Vision, vol. 2, Corfu, Greece, pp. 1033–1038 (1999)
3. Guillemot, C., Le Meur, O.: Image inpainting: overview and recent advances. *IEEE Signal Process. Mag.* **31**(1), 127–144 (2014)
4. Masnou, S., Morel, J.-M.: Level lines based disocclusion. In: Proc. 1998 IEEE International Conference on Image Processing, vol. 3, Chicago, pp. 259–263 (1998)
5. Schönlieb, C.-B.: *Partial Differential Equation Methods for Image Inpainting*. Cambridge University Press, Cambridge (2015)
6. Horn, B., Schunck, B.: Determining optical flow. *Artif. Intell.* **17**, 185–203 (1981)
7. Nagel, H.-H., Enkelmann, W.: An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Trans. Pattern Anal. Mach. Intell.* **8**, 565–593 (1986)
8. Weickert, J., Schnörr, C.: A theoretical framework for convex regularizers in PDE-based computation of image motion. *Int. J. Comput. Vis.* **45**(3), 245–264 (2001)
9. Bruhn, A., Weickert, J.: A confidence measure for variational optic flow methods. In: Klette, R., Kozera, R., Noakes, L., Weickert, J. (eds.) *Geometric Properties from Incomplete Data*. Computational Imaging and Vision, vol. 31, pp. 283–297. Springer, Dordrecht (2006)
10. Adam, R.D., Peter, P., Weickert, J.: Denoising by inpainting. In: Lauze, F., Dong, Y., Dahl, A.B. (eds.) *Scale Space and Variational Methods in Computer Vision*. Lecture Notes in Computer Science, vol. 10302, pp. 121–132. Springer, Cham (2017)
11. Laumont, R., Bortoli, V.D., Almansa, A., Delon, J., Durmus, A., Pereyra, M.: Bayesian imaging using plug & play priors: when Langevin meets Tweedie. *SIAM J. Imaging Sci.* **15**(2), 701–737 (2022)
12. Romano, Y., Elad, M., Milanfar, P.: The little engine that could: regularization by denoising (RED). *SIAM J. Imaging Sci.* **10**(4), 1804–1844 (2017)
13. Venkatakrisnan, S.V., Bouman, C.A., Wohlberg, B.: Plug-and-play priors for model based reconstruction. In: Proc. 2013 IEEE Global Conference on Signal and Information Processing, pp. 945–948. IEEE Comput. Soc., Austin (2013)
14. Iijima, T.: Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bull. Electrotechn. Lab.* **26**, 368–388 (1962). In Japanese
15. Lindeberg, T.: *Scale-Space Theory in Computer Vision*. Kluwer, Boston (1994)
16. Sporring, J., Nielsen, M., Florack, L., Johansen, P. (eds.): *Gaussian Scale-Space Theory*. Computational Imaging and Vision, vol. 8. Kluwer, Dordrecht (1997)
17. Lowe, D.L.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.* **60**(2), 91–110 (2004)
18. Galić, I., Weickert, J., Welk, M., Bruhn, A., Belyaev, A., Seidel, H.-P.: Image compression with anisotropic diffusion. *J. Math. Imaging Vis.* **31**(2–3), 255–269 (2008)
19. Mainberger, M., Hoffmann, S., Weickert, J., Tang, C.H., Johannsen, D., Neumann, F., Doerr, B.: Optimising spatial and tonal data for homogeneous diffusion inpainting. 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. 26–37. Springer, Berlin (2012)
20. Belhachmi, Z., Bucur, D., Burgeth, B., Weickert, J.: How to choose interpolation data in images. *SIAM J. Appl. Math.* **70**(1), 333–352 (2009)
21. Kämper, N., Weickert, J.: Domain decomposition algorithms for real-time homogeneous diffusion inpainting in 4K. In: Proc. 2022 IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 1680–1684. IEEE Comput. Soc., Singapore (2022)
22. Nielsen, M., Florack, L., Deriche, R.: Regularization, scale-space and edge detection filters. *J. Math. Imaging Vis.* **7**, 291–307 (1997)
23. Scherzer, O., Weickert, J.: Relations between regularization and diffusion filtering. *J. Math. Imaging Vis.* **12**(1), 43–63 (2000)
24. Weickert, J., Steidl, G., Mrázek, P., Welk, M., Brox, T.: Diffusion filters and wavelets: what can they learn from each other? In: Paragios, N., Chen, Y., Faugeras, O. (eds.) *Handbook of Mathematical Models in Computer Vision*, pp. 3–16. Springer, New York (2006)
25. Dong, B., Jiang, Q., Shen, Z.: Image restoration: wavelet frame shrinkage, nonlinear evolution PDEs, and beyond. *Multiscale Model. Simul.* **15**(1), 606–660 (2017)

26. Ruthotto, L., Haber, E.: Deep neural networks motivated by partial differential equations. *J. Math. Imaging Vis.* **62**, 352–364 (2020)
27. Alt, T., Schrader, K., Augustin, M., Peter, P., Weickert, J.: Connections between numerical algorithms for PDEs and neural networks. *J. Math. Imaging Vis.* **65**, 185–208 (2023)
28. Carlsson, S.: Sketch based coding of grey level images. *Signal Process.* **15**(1), 57–83 (1988)
29. Andris, S., Peter, P., Mohideen, R.M.K., Weickert, J., Hoffmann, S.: Inpainting-based video compression in FullHD. 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. 425–436. Springer, Cham (2021)
30. Chizhov, V., Weickert, J.: Efficient data optimisation for harmonic inpainting with finite elements. In: Tsapatsoulis, N., Panayides, A., Theocharides, T., Lanitis, A., Pattichis, C., Vento, M. (eds.) *Computer Analysis of Images and Patterns. Lecture Notes in Computer Science*, vol. 13053, pp. 432–441. Springer, Cham (2021)
31. Hoffmann, S., Plonka, G., Weickert, J.: Discrete Green's functions for harmonic and biharmonic inpainting with sparse atoms. In: Tai, X.-C., Bae, E., Chan, T.F., Lysaker, M. (eds.) *Energy Minimization Methods in Computer Vision and Pattern Recognition. Lecture Notes in Computer Science*, vol. 8932, pp. 169–182. Springer, Berlin (2015)
32. Kalmoun, E.M., Nasser, M.M.S.: Harmonic image inpainting using the charge simulation method. *Pattern Anal. Appl.* **25**(4), 795–806 (2022)
33. Mainberger, M., Bruhn, A., Weickert, J., Forchhammer, S.: Edge-based compression of cartoon-like images with homogeneous diffusion. *Pattern Recognit.* **44**(9), 1859–1873 (2011)
34. Bonettini, S., Loris, I., Porta, F., Prato, M., Rebegoldi, S.: On the convergence of a linesearch based proximal-gradient method for nonconvex optimization. *Inverse Probl.* **33**(5), 055005 (2017)
35. Breuß, M., Hoeltgen, L., Radow, G.: Towards PDE-based video compression with optimal masks prolonged by optic flow. *J. Math. Imaging Vis.* **63**(2), 144–156 (2021)
36. Chen, Y., Ranftl, R., Pock, T.: A bi-level view of inpainting-based image compression. In: Proc. 19th Computer Vision Winter Workshop, Křtiny, Czech Republic, pp. 19–26 (2014)
37. Hoeltgen, L., Setzer, S., Weickert, J.: An optimal control approach to find sparse data for Laplace interpolation. In: Heyden, A., Kahl, F., Olsson, C., Oskarsson, M., Tai, X.-C. (eds.) *Energy Minimization Methods in Computer Vision and Pattern Recognition. Lecture Notes in Computer Science*, vol. 8081, pp. 151–164. Springer, Berlin (2013)
38. Hoffmann, S.: Competitive image compression with linear PDEs. PhD thesis, Department of Computer Science, Saarland University, Saarbrücken, Germany (December 2016)
39. Mainberger, M., Hoffmann, S., Weickert, J., Tang, C.H., Johannsen, D., Neumann, F., Doerr, B.: Optimising spatial and tonal data for homogeneous diffusion inpainting. In: Bruckstein, A., Haar Romeny, B., Bronstein, A., Bronstein, M. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 6667, pp. 26–37. Springer, Berlin (2011)
40. Ochs, P., Chen, Y., Brox, T., Pock, T.: iPiano: inertial proximal algorithm for nonconvex optimization. *SIAM J. Appl. Math.* **7**(2), 1388–1419 (2014)
41. Gautier, J., Le Meur, O., Guillemot, C.: Efficient depth map compression based on lossless edge coding and diffusion. In: Proc. 2012 Picture Coding Symposium, Kraków, Poland, pp. 81–84 (2012)
42. Hoffmann, S., Mainberger, M., Weickert, J., Puhl, M.: Compression of depth maps with segment-based homogeneous diffusion. 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. 319–330. Springer, Berlin (2013)
43. Jost, F., Peter, P., Weickert, J.: Compressing flow fields with edge-aware homogeneous diffusion inpainting. In: Proc. 2020 IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 2198–2202. IEEE Comput. Soc., Barcelona (2020)
44. Jost, F., Peter, P., Weickert, J.: Compressing piecewise smooth images with the Mumford-Shah cartoon model. In: Proc. 2020 European Signal Processing Conference, pp. 511–515. IEEE Comput. Soc., Amsterdam (2021)
45. Li, Y., Sjöström, M., Jennehag, U., Olsson, R.: A scalable coding approach for high quality depth image compression. In: Proc. 3DTV-Conference: The True Vision - Capture, Transmission and Display of 3D Video, Zurich, Switzerland, pp. 1–4 (2012)
46. Sullivan, G.J., Ohm, J.-R., Han, W.-J., Wiegand, T.: Overview of the high efficiency video coding (HEVC) standard. *IEEE Trans. Circuits Syst. Video Technol.* **22**(12), 1649–1668 (2012)
47. Weickert, J., Welk, M.: Tensor field interpolation with PDEs. In: Weickert, J., Hagen, H. (eds.) *Visualization and Processing of Tensor Fields*, pp. 315–325. Springer, Berlin (2006)
48. Jumakulyyev, I., Schultz, T.: Fourth-order anisotropic diffusion for inpainting and image compression. In: Özarslan, E., Schultz, T., Zhang, E., Fuster, A. (eds.) *Anisotropy Across Fields and Scales*, pp. 99–124. Springer, Cham (2021)
49. Schmaltz, C., Peter, P., Mainberger, M., Ebel, F., Weickert, J., Bruhn, A.: Understanding, optimising, and extending data compression with anisotropic diffusion. *Int. J. Comput. Vis.* **108**(3), 222–240 (2014)
50. Pennebaker, W.B., Mitchell, J.L.: *JPEG: Still Image Data Compression Standard*. Springer, New York (1992)
51. Taubman, D.S., Marcellin, M.W. (eds.): *JPEG 2000: Image Compression Fundamentals, Standards and Practice*. Kluwer, Boston (2002)
52. Burger, M., He, L., Schönlieb, C.-B.: Cahn–Hilliard inpainting and a generalization for grayvalue images. *SIAM J. Imaging Sci.* **2**(4), 1129–1167 (2009)
53. Tai, X.-C., Hahn, J., Chung, G.J.: A fast algorithm for Euler's elastica model using augmented Lagrangian method. *SIAM J. Imaging Sci.* **4**(1), 313–344 (2011)
54. Peter, P.: Fast inpainting-based compression: combining Shepard interpolation with joint inpainting and prediction. In: Proc. 2019 IEEE International Conference on Image Processing, pp. 3557–3561. IEEE Comput. Soc., Taipei (2019)
55. Fritsch, D.S.: A medial description of greyscale image structure by gradient-limited diffusion. In: Robb, R.A. (ed.) *Visualization in Biomedical Computing '92. Proceedings of SPIE*, vol. 1808, pp. 105–117. SPIE Press, Bellingham (1992)
56. Perona, P., Malik, J.: Scale space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**, 629–639 (1990)
57. Aguerrebere, C., Almansa, A., Delon, J., Gousseau, Y., Musé, P.: A Bayesian hyperprior approach for joint image denoising and interpolation, with an application to HDR imaging. *IEEE Trans. Comput. Imaging* **3**(4), 633–646 (2017)
58. Aujol, J.-F., Ladjal, S., Masnou, S.: Exemplar-based inpainting from a variational point of view. *SIAM J. Math. Anal.* **42**(3), 1246–1285 (2010)

59. Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: PatchMatch: a randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.* **28**(3), 1–11 (2009)
60. Criminisi, A., Pérez, P., Toyama, K.: Region filling and object removal by exemplar-based image inpainting. *IEEE Trans. Image Process.* **13**(9), 1200–1212 (2004)
61. Facciolo, G., Arias, P., Caselles, V., Sapiro, G.: Exemplar-based interpolation of sparsely sampled images. In: Cremers, D., Boykov, Y., Blake, A., Schmidt, F.R. (eds.) *Energy Minimization Methods in Computer Vision and Pattern Recognition*. Lecture Notes in Computer Science, vol. 5681, pp. 331–344. Springer, Berlin (2009)
62. Bertalmio, M., Vese, L., Sapiro, G., Osher, S.: Simultaneous structure and texture image inpainting. *IEEE Trans. Image Process.* **12**(8), 882–889 (2003)
63. Peter, P., Weickert, J.: Compressing images with diffusion- and exemplar-based inpainting. In: Aujol, J.-F., Nikolova, M., Papadakis, N. (eds.) *Scale-Space and Variational Methods in Computer Vision*. Lecture Notes in Computer Science, vol. 9087, pp. 154–165. Springer, Berlin (2015)
64. Buades, A., Coll, B., Morel, J.-M.: A review of image denoising algorithms, with a new one. *Multiscale Model. Simul.* **4**(2), 490–530 (2005)
65. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Trans. Image Process.* **16**(8), 2080–2095 (2007)
66. Elad, M.: *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer, New York (2010)
67. Mallat, S.: *A Wavelet Tour of Signal Processing*, 2nd edn. Academic Press, San Diego (1999)
68. Ahmed, N., Natarajan, T., Rao, K.R.: Discrete cosine transform. *IEEE Trans. Comput.* **C-23**(1), 90–93 (1974)
69. Chen, S.S., Donoho, D.L., Saunders, M.A.: Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.* **20**(1), 33–61 (1998)
70. Donoho, D.L., Johnstone, I.M.: Ideal spatial adaptation by wavelet shrinkage. *Biometrika* **81**(3), 425–455 (1994)
71. Mallat, S.G., Zhang, Z.: Matching pursuits with time-frequency dictionaries. *IEEE Trans. Signal Process.* **41**(12), 3397–3415 (1993)
72. Guleryuz, O.G.: Iterated denoising for image recovery. In: *Proc. 2002 Data Compression Conference*, pp. 3–12. IEEE Comput. Soc., Snowbird (2002)
73. King, E.J., Kutyniok, G., Lim, W.-Q.: Image inpainting: theoretical analysis and comparison of algorithms. In: Van De Ville, D., Goyal, V.K., Papadakis, M. (eds.) *Wavelets and Sparsity XV*. Proceedings of SPIE, vol. 8858. SPIE Press, Bellingham (2013)
74. Berger, J.M., Lasher, G.J.: The use of discrete Green's functions in the numerical solution of Poisson's equation. III. *J. Math.* **2**(4A), 593–607 (1958)
75. Chung, F., Yau, S.-T.: Discrete Green's functions. *J. Comb. Theory, Ser. A* **91**(1), 191–214 (2000)
76. Katsurada, M.: A mathematical study of the charge simulation method by use of peripheral conformal mappings. *Mem. Instit. Sci. Technol., Meiji Univ.* **37**, 195–212 (1999)
77. Kupradze, V.D., Aleksidze, M.A.: The method of functional equations for the approximate solution of certain boundary value problems. *USSR Comput. Math. Math. Phys.* **4**(4), 82–126 (1964)
78. Elad, M., Kwar, B., Vaksman, G.: Image denoising: the deep learning revolution and beyond – a survey paper. *SIAM J. Imaging Sci.* **16**(3), 1594–1654 (2023)
79. Hamza, A.B., Krim, H., Unal, G.B.: Unifying probabilistic and variational estimation. *IEEE Signal Process. Mag.* **19**(5), 37–47 (2002)
80. Lebrun, M., Colom, M., Buades, A., Morel, J.M.: Secrets of image denoising cuisine. *Acta Numer.* **21**(1), 475–576 (2012)
81. Geman, S., Geman, D.: Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **6**(6), 721–741 (1984)
82. Larsson, E.G., Selen, Y.: Linear regression with a sparse parameter vector. *IEEE Trans. Signal Process.* **55**(2), 451–460 (2007)
83. Craven, P., Wahba, G.: Smoothing noisy data with spline functions. *Numer. Math.* **31**(4), 377–403 (1978)
84. Weyrich, N., Warhola, G.T.: Wavelet shrinkage and generalized cross validation for image denoising. *IEEE Trans. Image Process.* **7**(1), 82–90 (1998)
85. Chan, R.H., Ho, C.-W., Nikolova, M.: Salt-and-pepper noise removal by median-type noise detectors and detail-preserving regularization. *IEEE Trans. Image Process.* **14**(10), 1479–1485 (2005)
86. Weickert, J.: *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart (1998)
87. Weickert, J., Ishikawa, S., Imiya, A.: Linear scale-space has first been proposed in Japan. *J. Math. Imaging Vis.* **10**(3), 237–252 (1999)
88. Charbonnier, P., Blanc-Féraud, L., Aubert, G., Barlaud, M.: Deterministic edge-preserving regularization in computed imaging. *IEEE Trans. Image Process.* **6**(2), 298–311 (1997)
89. Chan, T.F., Shen, J.: Non-texture inpainting by curvature-driven diffusions (CDD). *J. Vis. Commun. Image Represent.* **12**(4), 436–449 (2001)
90. Dam, E., Nielsen, M.: Exploring non-linear diffusion: the diffusion echo. In: Kerckhove, M. (ed.) *Scale-Space and Morphology in Computer Vision*. Lecture Notes in Computer Science, vol. 2106, pp. 264–272. Springer, Berlin (2001)
91. Kuipers, L., Niederreiter, H.: *Uniform Distribution of Sequences*. Dover, New York (2005)
92. Floyd, R.W., Steinberg, L.: An adaptive algorithm for spatial grey scale. *Proc. Soc. Inf. Disp.* **17**, 75–77 (1976)
93. Roberts, M.: the unreasonable effectiveness of quasirandom sequences (2018). <https://extremelarning.com/unreasonable-effectiveness-of-quasirandom-sequences/> Accessed 2023-03-06
94. Werner, H.: Studies on contour. *Am. J. Psychol.* **47**(1), 40–64 (1935)
95. Daropoulos, V., Augustin, M., Weickert, J.: Sparse inpainting with smoothed particle hydrodynamics. *SIAM J. Imaging Sci.* **14**(4), 1669–1705 (2021)
96. Peter, P., Schrader, K., Alt, T., Weickert, J.: Deep spatial and tonal data optimisation for homogeneous diffusion inpainting. *Pattern Anal. Appl.* **26**(4), 1585–1600 (2023)
97. Jost, F., Chizhov, V., Weickert, J.: Optimising different feature types for inpainting-based image representations. In: *Proc. 2023 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE Comput. Soc., Rhodes Island (2023)

**Publisher's note**

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

**Submit your manuscript to a SpringerOpen<sup>®</sup> journal and benefit from:**

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

---

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)

---