In F. Lauze, Y. Dong, A. B. Dahl (Eds.): Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science, Vol. 10302, pp. 121–132, Springer, Cham, 2017. The final publication is available at link.springer.com.

Denoising by Inpainting

Robin Dirk Adam, Pascal Peter, and Joachim Weickert

Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science, Campus E1.7, Saarland University, 66041 Saarbrücken, Germany.
{adam, peter, weickert}@mia.uni-saarland.de

Abstract. The filling-in effect of diffusion processes has been successfully used in many image analysis applications. Examples include image reconstructions in inpainting-based compression or dense optic flow computations. As an interesting side effect of diffusion-based inpainting, the interpolated data are smooth, even if the known image data are noisy: Inpainting averages information from noisy sources. Since this effect has not been investigated for denoising purposes so far, we propose a general framework for denoising by inpainting. It averages multiple inpainting results from different selections of known data. We evaluate two concrete implementations of this framework: The first one specifies known data on a shifted regular grid, while the second one employs probabilistic densification to optimise the known pixel locations w.r.t. the inpainting quality. For homogeneous diffusion inpainting, we demonstrate that our regular grid method approximates the quality of its corresponding diffusion filter. The densification algorithm with homogeneous diffusion inpainting, however, shows edge-preserving behaviour. It resembles space-variant diffusion and offers better reconstructions than homogeneous diffusion filters.

Keywords: diffusion, denoising, inpainting, densification

1 Introduction

Image inpainting is the task of reconstructing missing image parts from available known data [2, 10, 17, 25]. Diffusion filters have been proven to be capable of recovering images from very sparse pixel sets in high quality [7, 13, 16], which is particularly useful in the context of image compression [12, 24]. This filling-in effect has also been used successfully for more than three decades in variational models for optic flow computation such as [14, 18]. Here, dense displacement vector fields are created by inpainting at locations where no flow can be measured and the data term vanishes. Surprisingly, these reconstructed parts of the flow fields are often more reliable than the measured flow vectors, since the diffusion-based inpainting solution averages information from many noisy data in the neighbourhood [3]. One can also observe a similar effect in diffusion-based compression applications: For low compression rates, the compressed image can look smoother and visually more pleasing than the original.

Denoising is another classic image processing task that can be solved by diffusion. From the simple original homogeneous diffusion filter [15], a plethora of fairly sophisticated approaches has evolved; see e.g. [19, 27]. Although nowadays non-local denoising methods such as BM3D [9] are very popular, modern diffusion-reaction models that rely on learning yield competitive results [6]. Both the classic and the more recent diffusion-based denoising methods have in common that they approach the task by directly applying smoothing to the image. To our best knowledge, the potential denoising capabilities of diffusion-based *inpainting*, however, have not been investigated so far.

Our Contribution. In order to close this gap, we propose a general framework for *denoising by inpainting*: In order to denoise an image, we average several inpainting results that use different selections of the noisy original pixels as known data. Moreover, we introduce two different implementations of this framework that both rely on linear homogeneous diffusion inpainting, but differ w.r.t. the selection strategy for the known image points. Our investigations show that inpainting from specified points at shifted, non-overlapping regular grid positions approximates the quality of linear homogeneous diffusion. We also propose a more sophisticated probabilistic strategy inspired by the sparsification approach of Mainberger et al. [16]: It adapts the locations of known data to the image structure. Our evaluation reveals that this method possesses edge-preserving properties similar to space-variant diffusion, while using a space-invariant differential operator.

Related Work. Our work makes extensive use of diffusion filters. In particular, our implementations rely on linear homogeneous diffusion, which goes back to Iijima [15]. We also consider nonlinear isotropic diffusion, which was first introduced by Perona and Malik [19]. In contrast to linear diffusion, this filter adapts to the local image structure in order to preserve edges. After these classic models, many more have been proposed (e.g. [27]), but a full review would be beyond the scope of this work. Since we are mainly interested in gaining insight into new applications for inpainting and do not aim to produce state-of-the-art denoising results, we focus solely on linear homogeneous and nonlinear isotropic diffusion filters.

Interestingly, diffusion filtering can be related to many other types of denoising methods. For instance, Scherzer and Weickert [23] have shown connections between variational methods such as Tikhonov [26] or TV regularisation [22] and fully implicit time discretisations of diffusion filters. Furthermore, a large variety of diffusion filters for denoising can be interpreted as Bayesian denoising models; see e.g. [20]. Thereby, they can be seen as special cases of probabilistic approaches such as the field-of-experts model [21]. Also, relations to wavelet shrinkage have been established [28].

There are also other classes of denoising filters that we cannot discuss in detail in this work. In particular, non-local methods like BM3D [9] and its successors can be regarded as a sophisticated extension of the NL-means filter of Buades et al. [4]: Such approaches search for similar image patches and average over those to remove noise. NL-means can also be seen as a specific non-local representative of a denoising by inpainting strategy, since Buades et al. have been inspired by the exemplar-based inpainting method of Efros and Leung [10].

With respect to image inpainting, parts of our paper rely on spatial optimisation techniques. It has been shown that in cases where a sparse image representation can be chosen from the fully available original image, the selection of known data has a significant impact on reconstruction quality [7, 13, 16]. In particular, we focus on the probabilistic sparsification method by Mainberger et al. [16] that iteratively removes image points which are easy to reconstruct. In a broad conceptual sense, the sparsification and densification strategies that we consider in this paper are related to the generalised cross-validation methods by Craven and Wahba [8], since here also approximation accuracy under removal of data is considered. However, generalised cross-validation is usually used to determine model parameters, for instance for denoising based on wavelet shrinkage [29]. In our application, we iteratively remove known data to obtain sparse image representations.

Organisation of the Paper. Since the concepts of diffusion-based denoising and inpainting are integral to our work, we review them in Section 2. In Section 3 we propose our general framework for denoising by inpainting. With an approach based on regular masks in Section 4 and a densification scheme in Section 5, we also present two concrete implementations of this framework. Finally, we evaluate the new denoising methods in Section 6 and conclude our work with a discussion and an outlook in Section 7.

2 Diffusion-based Denoising and Inpainting

2.1 Diffusion-based Denoising

Our goal is to apply a diffusion filter to a noisy image $f : \Omega \to \mathbb{R}$ that maps the rectangular image domain $\Omega \subset \mathbb{R}^2$ to the grey value range $\mathbb{R}$. To this end we start with the initial condition $u(x, y, 0) = f(x, y)$ and compute filtered versions $\{u(x, y, t) \,|\, (x, y) \in \Omega,\, t \geq 0\}$ of $f(x, y)$ with diffusion time t as solutions of the following initial boundary value problem:

$$\partial_t u = \operatorname{div}(g\, \boldsymbol{\nabla} u) \quad \text{on} \quad \Omega \times (0, \infty), \tag{1}$$

$$u(x, y, 0) = f(x, y) \quad \text{on} \quad \Omega, \tag{2}$$

$$\partial_{\boldsymbol{n}} u = 0 \quad \text{on} \quad \partial\Omega \times (0, \infty). \tag{3}$$

Here we use the outer normal vector $\boldsymbol{n}$ to the image boundary $\partial\Omega$ and the corresponding directional derivative $\partial_{\boldsymbol{n}}$ to specify reflecting boundary conditions. By $\boldsymbol{\nabla} = (\partial_x, \partial_y)^\top$ we denote the spatial nabla operator, and div is its corresponding divergence operator. The diffusion time t embeds the filtered images u into a scale-space: Increasing the diffusion time simplifies the image.

The scalar-valued diffusivity $g : [0, \infty) \to (0, \infty)$ is a positive function of the local image structure. The magnitude of g determines how much smoothing the diffusion filter performs at a given l ocation in the image. Three scenarios are relevant for our paper: homogeneous, linear space-variant, and nonlinear diffusion.

Homogeneous diffusion [15] uses the constant diffusivity $g = 1$. This leads to the linear diffusion equation (also known as heat equation) $\partial_t u = \Delta u$ with the Laplacian $\Delta = \partial_{xx} + \partial_{yy}$. This evolution generates the well-known Gaussian scale-space [15]. It is simple to implement and uses no additional parameters apart from t. Since homogeneous diffusion smoothes equally at all locations of the image, it is *space-invariant*. However, as it cannot distinguish between noise and edges, it also blurs semantically important image edges.

Linear space-variant diffusion [11] avoids this drawback by adapting the diffusive evolution to the *initial* image f. This can be achieved by choosing $g = g(|\boldsymbol{\nabla} f|^2)$ with a decreasing diffusivity function that becomes small at edges where $|\boldsymbol{\nabla} f|$ is large. An example is the Charbonnier diffusivity [5]

$$g_{\mathrm{C}}(s^2) := \left(1 + \frac{s^2}{\lambda^2}\right)^{-1/2}. \tag{4}$$

Note that locations where $|\boldsymbol{\nabla} f| \gg \lambda$ are regarded as edges where the diffusivity is close to 0, while we have full diffusion in regions with $|\boldsymbol{\nabla} f| \ll \lambda$. Therefore, $\lambda > 0$ acts as a contrast parameter.

Nonlinear diffusion [19] goes one step further and chooses the diffusivity as a function of the *evolving* image $u(., t)$. Using $g = g(|\boldsymbol{\nabla} u|^2)$ instead of $g = g(|\boldsymbol{\nabla} f|^2)$ introduces a nonlinear feedback into the evolution. Since the evolving image becomes gradually smoother, one often obtains better denoising results than for linear space-variant diffusion.

To keep everything simple and focus on structural insights, we do not consider more advanced diffusion processes that use a diffusion tensor [27].

2.2 Diffusion-based Inpainting

With some small modifications, the diffusion filters from the previous section can be applied to image inpainting problems. Let the original image f only be known on the *inpainting mask* $K \subset \Omega$. Our goal is to reconstruct the missing data in the *inpainting domain* $\Omega \setminus K$. We achieve this by computing the steady state ($t \to \infty$) of the image evolution of $u(x, y, t)$ that is described by

$$\partial_t u = \operatorname{div}(g\, \boldsymbol{\nabla} u) \quad \text{on} \quad \Omega \setminus K \times (0, \infty), \tag{5}$$

$$u(x, y, t) = f(x, y) \quad \text{on} \quad K \times [0, \infty), \tag{6}$$

$$\partial_{\boldsymbol{n}} u = 0 \quad \text{on} \quad \partial\Omega \times (0, \infty). \tag{7}$$

In contrast to Eq. (1)–(3), the diffusion PDE is only applied to the inpainting domain $\Omega \setminus K$, while Dirichlet boundary conditions fix the known data on K. This leads to a non-flat steady state. Equivalently, we can formulate the inpainting problem with the elliptic PDE

$$(1 - c(\boldsymbol{x}))Lu - c(\boldsymbol{x})(u - f) = 0 \tag{8}$$

where $Lu := \operatorname{div}(g\,\boldsymbol{\nabla} u)$, and c is a binary confidence function that is 1 on K, and 0 on $\Omega \setminus K$. Note that on $\partial\Omega$, the reflecting boundary conditions still apply. In the following sections we favour this more compact notation, and we use the term *inpainting mask* for both the set K and its associated indicator function c.

3 A General Framework for Denoising by Inpainting

In our new denoising approach we want to exploit that inpainting reconstructions are smooth, even if we apply the diffusion operator to noisy known data. Thus, given a confidence mask c as in the previous section, we expect our pixels in the inpainting domain $\Omega \setminus K$ to be more reliable than our known noisy data K.

Obviously this has the undesired effect that all noisy pixels which belong to our inpainting mask are not affected by the filter at all. To obtain a denoised image u in *all* pixels, we average n reconstructions $(v^\ell)_{\ell=0}^{n-1}$ that are computed with the same differential operator L, but with *different* inpainting masks $(c^\ell)_{\ell=0}^{n-1}$. This leads to the following general formulation:

$$(1 - c^\ell(\boldsymbol{x}))Lv^\ell - c^\ell(\boldsymbol{x})(v^\ell - f) = 0\,, \qquad \ell \in \{0, ..., n-1\}\,, \tag{9}$$

$$u = \frac{1}{n}\sum_{\ell=0}^{n-1} v^\ell\,. \tag{10}$$

Compared to standard diffusion-based denoising, we do not have to choose a diffusion time any more, since all reconstructions correspond to steady states of the n inpaintings in Eq. (9). Instead, the mask density d (the percentage of known data points) is the free parameter of our model that determines the amount of smoothing: Decreasing the mask density leads to more smoothing.

In the following, we introduce two different strategies to choose the mask locations corresponding to this density parameter. In both cases, we choose $L = \Delta$ as our differential operator, i.e. we use homogeneous diffusion inpainting.

4 Denoising with Regular Masks

For our first inpainting-based denoising model, we choose the masks as shifted versions of a regular grid. Let the spacing between known pixels in the grid be given by r in x-direction and s in y-direction. Then we have $n := r \cdot s$ ways to shift this grid in a non-overlapping way. For a discrete image with resolution $M \times N$ and grid size h, we define the space-discrete masks $\boldsymbol{c}^\ell$ with $\ell \in \{0, ..., n-1\}$ by

$$c_{i,j}^{ps+q} = c^{ps+q}(ih, jh) := \begin{cases} 1 & i = p \bmod r \;\text{ and }\; j = q \bmod s\,, \\ 0 & \text{else}\,. \end{cases} \tag{11}$$

original **noisy** **homog. diffusion** **regular inpainting**

$\sigma = 30$ MSE: 75.37 MSE: 85.54

Fig. 1. Experiment: Denoising with Inpainting on Regular Masks versus Homogeneous Diffusion. For this test on *trui* with Gaussian noise ($\sigma = 30$), optimal diffusion time and grid spacing were chosen for each method respectively. We compare both methods w.r.t. the mean squared error (MSE).

Here, $p \in \{0, r-1\}$ and $q \in \{0, ..., s-1\}$ are the admissible grid offsets. Each pixel in the image domain Ω is covered by exactly one mask. For our denoising model, this means that at each location the confidence in the known data is equal: We always average $n-1$ inpainting results and the original pixel.

Experimentally, we determine that this scheme with regular masks is indeed capable of denoising, but performs slightly worse than homogeneous diffusion filtering. For a typical result see Fig. 1. However, both quantitatively and visually, the results of our regular mask inpainting approach with homogeneous diffusion appear to approximate homogeneous diffusion filtering. In the following, we justify this behaviour with considerations in the 1-D setting.

1-D Analysis. Let us consider inpainting with 1-D homogenous diffusion and regular masks with spacing n. For a pixel position i and a mask shifted by $p \in \{0, ..., n-1\}$, we define the offset $\ell = |i-p| \bmod n$ relative to i. This implies that for the mask c^ℓ with $\ell = 0$, the location i is known. For general choices $\ell \in \{0, ..., n-1\}$, the known points that are closest to i are $i-\ell$ on the left and $i+n-\ell$ on the right. Since in 1-D, inpainting with homogeneous diffusion is equivalent to linear interpolation between adjacent known points, we obtain the reconstruction $\boldsymbol{v}^\ell$ at location i as

$$v_i^\ell = \frac{n-\ell}{n} f_{i-\ell} + \frac{\ell}{n} f_{i+n-\ell} \,. \tag{12}$$

Now we average our inpainting solutions $(\boldsymbol{v}^\ell)_{\ell=0}^{n-1}$ to end up with the denoised image $\boldsymbol{u}$. This yields

$$u_i = \frac{1}{n} \sum_{\ell=0}^{n-1} v_i^\ell = \frac{1}{n} \sum_{\ell=0}^{n-1} \Big(\frac{n-\ell}{n} f_{i-\ell} + \frac{\ell}{n} f_{i+n-\ell} \Big) \tag{13}$$

$$= \frac{1}{n^2} \Big(n \cdot f_i + \sum_{\ell=1}^{n-1} \ell \cdot (f_{i-n+\ell} + f_{i+n-\ell}) \Big) \,. \tag{14}$$

For the highest non-trivial regular mask density, which comes down to storing every second pixel ($n = 2$), we obtain

$$u_i = \frac{f_{i-1} + 2f_i + f_{i+1}}{4} \quad \Longleftrightarrow \quad \frac{u_i - f_i}{\tau} = \frac{f_{i+1} - 2f_i + f_{i-1}}{h^2} \tag{15}$$

with $\tau := \frac{4}{h^2}$. Since the right equation is an explicit finite difference step of $\partial_t u = \partial_{xx} u$ with initial value f and time step size τ, this is equivalent to applying a homogeneous diffusion filter. For larger choices of n, Eq. 14 corresponds to convolving the image with a symmetric sampling of a hat function that has $2n + 1$ non-zero samples. This visually resembles Gaussian convolution.

5 Denoising with Adaptive Masks

In order to improve our denoising results compared to the non-adaptive masks from the previous section, we want to rely on *spatial mask optimisation* [7, 13, 16], a successful concept in PDE-based compression. Optimising the location of the known data can improve each individual reconstruction $\boldsymbol{v}^\ell$ and thereby also the average $\boldsymbol{u}$. For homogeneous diffusion inpainting, the theory of Belhachmi et al. [1] recommends to choose locations left and right of image edges. However, in images with large amounts of noise, edge detection is by no means an easy task. Moreover, we require multiple different masks for our general denoising by inpainting framework from Section 3.

Among the wide variety of different approaches for spatial optimisation, the probabilistic approach by Mainberger et al. [16] seems to be the most promising for our purpose: It does not rely on edge detection and contains a random component that we can use to generate different adaptive masks.

Sparsification. The original probabilistic sparsification starts with a mask that contains all image points and successively reduces the amount of known pixels until it reaches a target density d. In each iteration, it removes a fixed percentage α of known data. After inpainting with the resulting smaller mask, it adds a percentage β of the removed pixels with the highest reconstruction error back to the mask. Thus, out of $\alpha\%$ candidates, we remove the $\beta\%$ pixels that can be reconstructed best.

Unfortunately, applying sparsification directly to our denoising problem with homogeneous diffusion yields unsatisfactory results due to its local error computation: It considers the deviation of each candidate pixel from the corresponding image point in the noisy input image and preserves those candidates with the largest deviation. However, a large local difference can not only result from fine scale detail that should be preserved. Since the original data are noisy, the sparsification algorithm preserves noise that deviates from the smooth reconstruction. One solution to avoid this problem is to consider the impact of removing a single pixel on the overall reconstruction: If the noise has zero mean, computing the global error between inpainting solution and noisy image should give a better estimate to the error w.r.t. the unperturbed original. However, even with this change, sparsification selects noise pixels (see Fig. 2(b)).

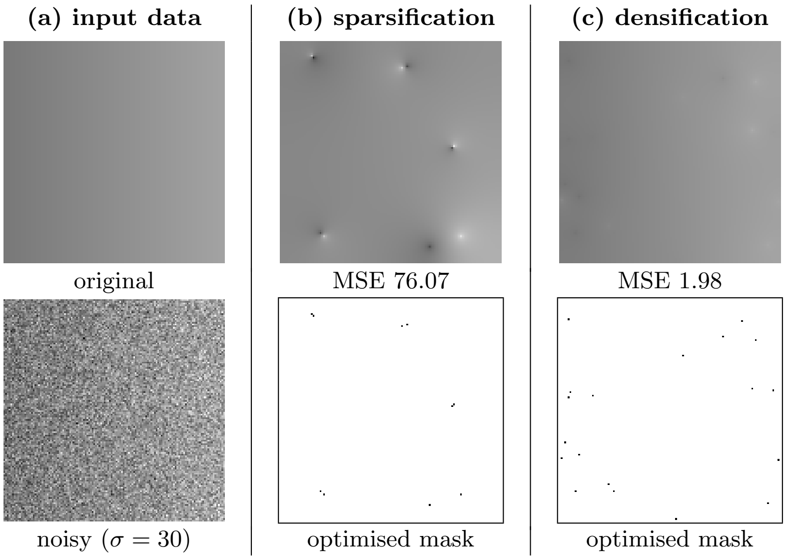

Fig. 2. Experiment: Densification versus Sparsification. For both methods, the mask density d was optimised with a grid search w.r.t. the MSE. The noisy gradient image is not reconstructed adequately by sparsification, since it prefers to keep noisy pixels in the first iterations due to localisation. Densification does not suffer from this problem and thereby achieves a more accurate inpainting.

The reason for this behaviour is a second source of locality: The influence of a pixel on the reconstruction result is determined by the mask density of its surroundings. This is illustrated by two extreme cases: In a mask consisting of a single pixel, its influence is truly global. It determines the average grey value of the flat steady state. In contrast, a pixel surrounded entirely by known data does not influence the inpainting at all. Since we start with a dense pixel mask in sparsification, each pixel initially has a very small influence which gradually increases the more points are removed. This leads to the preference of noisy data.

Densification. In order to remove this second source of locality, we instead propose a densification approach in Algorithm 1. We start with an empty mask and consider α randomly selected candidates that do not belong to the mask. We then only add the single pixel that improves the overall reconstruction error w.r.t. the noisy image the most.

6 Experiments

In the following we evaluate the performance of our two approaches for denoising by inpainting from Sections 4 and 5. We add Gaussian noise to the test images *trui*, *peppers*, and *lena* to compare our methods to diffusion filters. In order to reveal the full potential of each algorithm, we select the respective parameters

Input: Noisy image $\boldsymbol{f} \in \mathbb{R}^{MN}$, number α of candidates, desired final mask density d.
Initialisation: Mask $\boldsymbol{c} = 0$ is empty.
Compute:
do
1. Choose randomly a set $A \subset \{k \in \{1, ..., MN\} \,|\, c_k = 0\}$ with α candidates.
   for all $i \in A$ **do**
   2. Set temporary mask $\boldsymbol{m}^i$ such that $\forall k \in \{1, ..., \alpha\} \setminus \{i\} : m_k^i = c_k$, $m_i^i = 1$.
   3. Compute reconstruction $\boldsymbol{u}^i$ from mask $\boldsymbol{m}^i$ and image data $\boldsymbol{f}$.
   end for
4. Set $\boldsymbol{c} = \operatorname{argmin}_{\boldsymbol{m}^i} \operatorname{MSE}(\boldsymbol{u}^i, \boldsymbol{f})$. This adds one mask point to $\boldsymbol{c}$.

while pixel density of $\boldsymbol{c}$ smaller than d.
Output: Mask $\boldsymbol{c}$ of density d.

Algorithm 1: Mask densification with global error computation.

such that the mean squared error (MSE) w.r.t. the ground truth is minimised. This includes the stopping time of all diffusion processes, the contrast parameter λ in the diffusivity (4) for the linear space-variant and nonlinear diffusion models, as well as the mask density for the inpainting approaches. For this optimisation, we use a straightforward grid search. For all experiments, we have fixed the number of different masks in our densification approach to $n = 128$.

The results in Fig. 1 and Tab. 1 confirm that our inpainting approach with regular masks approximates the quality of homogeneous diffusion filtering. It is slightly worse than its diffusion counterpart.

Our densification method, however, proves to be consistently better than denoising with homogeneous diffusion. Note that it does not only offer a better quantitative performance: Due to the preservation of edges, the results are also visually more pleasing (see Fig. 3).

Suprisingly, the densification method is even superior to linear space-variant diffusion filtering in 8 out of 9 cases considered in Tab. 1. In order to understand this behaviour, we should remember that the densification method achieves adaptivity by searching for the most useful pixels as inpainting data. Typically these are those pixels which are less degraded by Gaussian noise than their neighbours. Linear space-variant diffusion lacks such a mechanism to identify the most reliable pixels: All edge pixels with the same gradient magnitude are assigned the same diffusivity, regardless of their individual reliability. This explains the slightly weaker performance of linear space-variant diffusion. At the same time, our model is simpler: Since it uses homogeneous diffusion, there is no need to choose a diffusivity model (e.g. Eq. 4) or the parameter λ.

Finally, comparing the densification method based on homogeneous diffusion inpainting with a nonlinear diffusion filter shows its limitations, in particular for high noise levels. This is an unfair comparison: Since the densification approach lacks a nonlinear feedback mechanism, it is not suprising that its performance is dominated by nonlinear diffusion filtering.

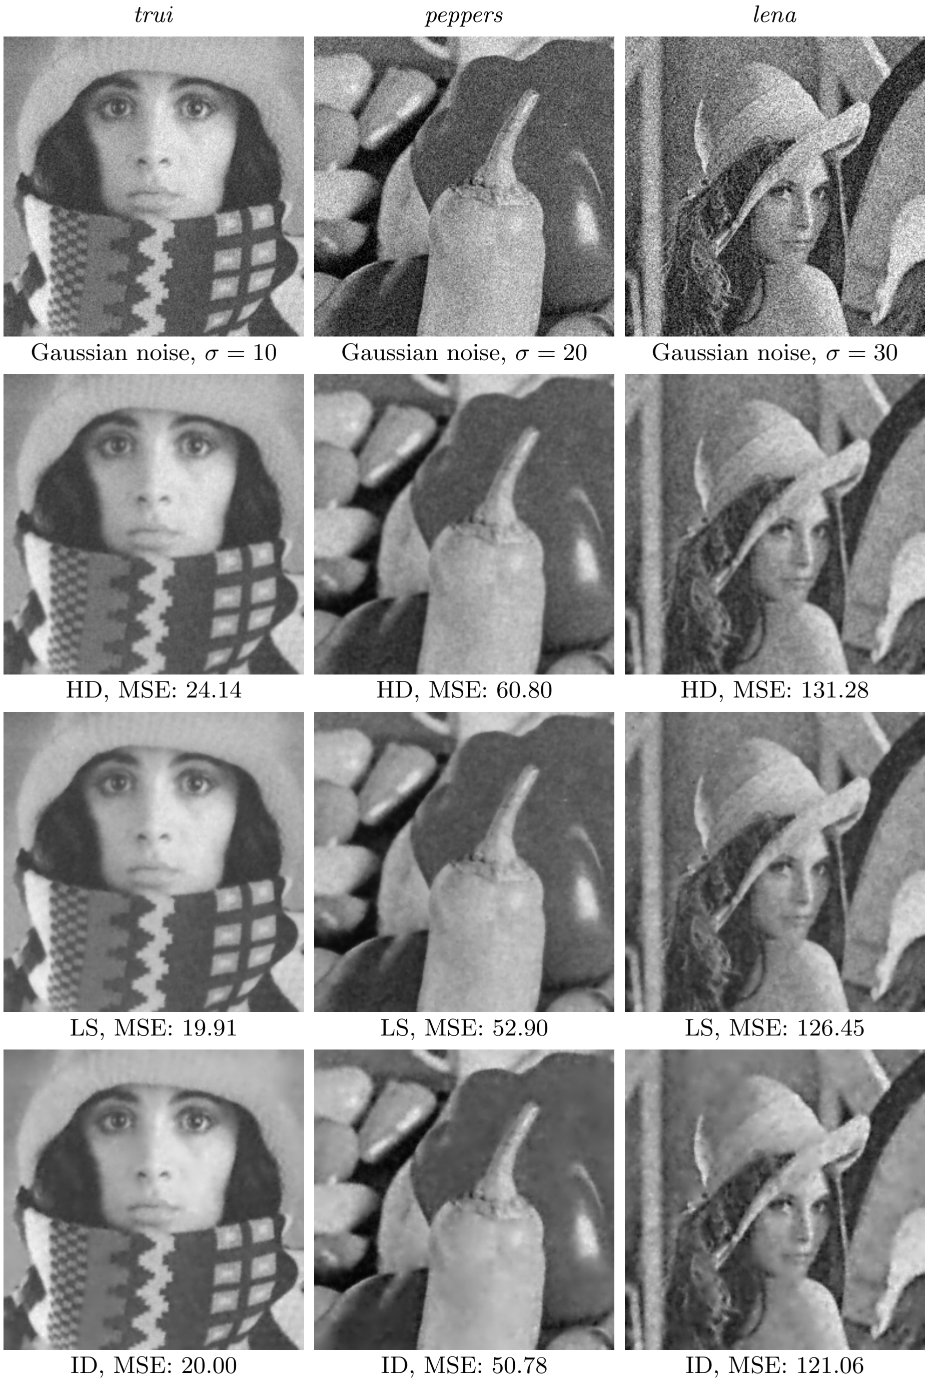

Fig. 3. Comparison of homogeneous diffusion (HD), linear space-variant diffusion (LS), and denoising by inpainting with densification (ID)

test image	*trui*			*peppers*			*lena*		
noise scale σ	10	20	30	10	20	30	10	20	30
inpainting with reg. masks	27.25	56.81	85.54	35.69	65.94	97.34	44.56	91.79	134.79
inpainting by densification	20.00	44.61	73.18	25.04	50.78	75.10	31.43	76.55	121.06
homogeneous diffusion	24.14	49.73	75.37	32.32	60.80	89.14	43.04	89.58	131.28
linear space-var. diffusion	19.91	46.25	73.51	25.24	52.90	82.98	32.13	77.04	126.45
nonlinear diffusion	16.43	35.12	55.03	22.34	41.17	62.79	28.16	64.28	99.84

Table 1. Denoising Results. We compare our two denoising by inpainting strategies (that employ homogeneous diffusion inpainting) with three diffusion filters.

7 Conclusions

Our work is the first that explicitly demonstrates the denoising capabilities of PDE-based inpainting methods. In particular, implementing our general framework with adaptive inpainting masks introduces space-variant behaviour to purely homogeneous processes. The resulting densification strategy based on homogeneous diffusion inpainting does not only outperform homogeneous diffusion filtering, but even linear space-variant diffusion filters. This shows a fundamental principle for denoising that has been widely ignored: *Adaptivity in the filter model can be replaced by adaptivity of data selection.* Exploring this encouraging road further by gaining more theoretical insights is part of our ongoing work.

References

1. Belhachmi, Z., Bucur, D., Burgeth, B., Weickert, J.: How to choose interpolation data in images. SIAM Journal on Applied Mathematics 70(1), 333–352 (Jun 2009)
2. Bertalmío, M., Sapiro, G., Caselles, V., Ballester, C.: Image inpainting. In: Proc. SIGGRAPH 2000. pp. 417–424. New Orleans, LI (Jul 2000)
3. 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)
4. Buades, A., Coll, B., Morel, J.M.: A review of image denoising algorithms, with a new one. Multiscale Modeling and Simulation 4(2), 490–530 (2005)
5. Charbonnier, P., Blanc-Féraud, L., Aubert, G., Barlaud, M.: Deterministic edge-preserving regularization in computed imaging. IEEE Transactions on Image Processing 6(2), 298–311 (1997)
6. Chen, Y., Yu, W., Pock, T.: On learning optimized reaction diffusion processes for effective image restoration. In: Proc. 2015 IEEE Conference on Computer Vision and Pattern Recognition. pp. 5261–5269. Boston, MA (Jun 2015)
7. Chen, Y., Ranftl, R., Pock, T.: A bi-level view of inpainting-based image compression. In: Proc. 19th Computer Vision Winter Workshop. pp. 19–26. Křtiny, Czech Republic (Feb 2014)
8. Craven, P., Wahba, G.: Smoothing noisy data with spline functions. Numerische Mathematik 31(4), 377–403 (Jan 1978)

9. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-D transform-domain collaborative filtering. IEEE Transactions on Image Processing 16(8), 2080–2095 (Jul 2007)
10. Efros, A.A., Leung, T.K.: Texture synthesis by non-parametric sampling. In: Proc. Seventh IEEE International Conference on Computer Vision. vol. 2, pp. 1033–1038. Corfu, Greece (Sep 1999)
11. 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)
12. Galić, I., Weickert, J., Welk, M., Bruhn, A., Belyaev, A., Seidel, H.P.: Image compression with anisotropic diffusion. Journal of Mathematical Imaging and Vision 31(2–3), 255–269 (Jul 2008)
13. 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)
14. Horn, B., Schunck, B.: Determining optical flow. Artificial Intelligence 17, 185–203 (Aug 1981)
15. Iijima, T.: Basic theory on normalization of pattern (in case of typical one-dimensional pattern). Bulletin of the Electrotechnical Laboratory 26, 368–388 (Jan 1962), in Japanese
16. 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., ter 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)
17. Masnou, S., Morel, J.M.: Level lines based disocclusion. In: Proc. 1998 IEEE International Conference on Image Processing. vol. 3, pp. 259–263. Chicago, IL (Oct 1998)
18. Nagel, H.H., Enkelmann, W.: An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. IEEE Transactions on Pattern Analysis and Machine Intelligence 8, 565–593 (Sep 1986)
19. Perona, P., Malik, J.: Scale space and edge detection using anisotropic diffusion. IEEE Transactions on Pattern Analysis and Machine Intelligence 12, 629–639 (Jul 1990)
20. Peter, P., Weickert, J., Munk, A., Krivobokova, T., Li, H.: Justifying tensor-driven diffusion from structure-adaptive statistics of natural images. 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. 263–277. Springer, Berlin (2015)
21. Roth, S., Black, M.J.: Fields of experts. International Journal of Computer Vision 82(2), 205–229 (Apr 2009)
22. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. Physica D 60(1), 259–268 (Nov 1992)
23. Scherzer, O., Weickert, J.: Relations between regularization and diffusion filtering. Journal of Mathematical Imaging and Vision 12(1), 43–63 (Feb 2000)
24. Schmaltz, C., Peter, P., Mainberger, M., Ebel, F., Weickert, J., Bruhn, A.: Understanding, optimising, and extending data compression with anisotropic diffusion. International Journal of Computer Vision 108(3), 222–240 (Jul 2014)

25. Schönlieb, C.B.: Partial Differential Equation Methods for Image Inpainting. Cambridge University Press, Cambridge (2015)
26. Tikhonov, A.N., Arsenin, V.Y.: Solutions of Ill–Posed Problems. Wiley, Washington, DC (1977)
27. Weickert, J.: Anisotropic Diffusion in Image Processing. Teubner, Stuttgart (1998)
28. 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)
29. Weyrich, N., Warhola, G.T.: Wavelet shrinkage and generalized cross validation for image denoising. IEEE Transactions on Image Processing 7(1), 82–90 (Aug 1998)