

Optimising Spatial and Tonal Data for Homogeneous Diffusion Inpainting

Markus Mainberger¹, Sebastian Hoffmann¹, Joachim Weickert¹, Ching Hoo Tang²,
Daniel Johannsen³, Frank Neumann⁴, and Benjamin Doerr²

¹ Mathematical Image Analysis Group,
Faculty of Mathematics and Computer Science, Campus E1.1
Saarland University, 66041 Saarbrücken, Germany
{mainberger, hoffmann, weickert}@mia.uni-saarland.de

² Department 1: Algorithms and Complexity
Max Planck Institute for Informatics, Campus E1.4
66123 Saarbrücken, Germany
{doerr, chtang}@mpi-inf.mpg.de

³ School of Mathematical Sciences,
Tel Aviv University, Ramat Aviv, Tel Aviv 69978, Israel
johannse@tau.ac.il

⁴ School of Computer Science, Innova21 Building
University of Adelaide, Adelaide, SA 5005, Australia
frank@cs.adelaide.edu.au

Abstract. Finding optimal inpainting data plays a key role in the field of image compression with partial differential equations (PDEs). In this paper, we optimise the spatial as well as the tonal data such that an image can be reconstructed with minimised error by means of discrete homogeneous diffusion inpainting. To optimise the spatial distribution of the inpainting data, we apply a probabilistic data sparsification followed by a nonlocal pixel exchange. Afterwards we optimise the grey values in these inpainting points in an exact way using a least squares approach. The resulting method allows almost perfect reconstructions with only 5% of all pixels. This demonstrates that a thorough data optimisation can compensate for most deficiencies of a suboptimal PDE interpolant.

Keywords: image compression, partial differential equations (PDEs), inpainting, optimisation, homogeneous diffusion

1 Introduction

Research on PDE-based data compression suffers from poverty, but enjoys liberty [1, 2, 8, 18]: Unlike in pure inpainting research [14, 3], one has an extremely tight pixel budget for reconstructing some given image. However, one is free to choose where and how one spends this budget.

Let us explain the problem of PDE-based image compression in more detail. The basic idea is to reconstruct some given image by inpainting from a sparse set of pixels with a suitable partial differential equation (PDE). There is an evident tradeoff between




Fig. 1. Reconstruction of the test image *trui* using only 5% of all pixels and homogeneous diffusion inpainting. (a) Original image. (b) Unoptimised data (randomly selected from original image). (c) Optimised tonal and spatial data.

the number of pixels to store and the achievable reconstruction quality. Even if the number of pixels and the PDE are already specified, we still have many degrees of freedom: On one hand we can place the pixels wherever we want. On the other hand we can freely choose the grey value (or colour value) in each selected pixel.

The goal of the present paper is to optimise this spatial and tonal data selection. In order to show the real potential behind this approach, we choose an extremely simple PDE that has a bad reputation for inpainting tasks: We interpolate with the steady state of a homogeneous diffusion process, i.e. we solve the Laplace equation. Figure 1 illustrates the huge potential that one can exploit with this optimisation. Even with homogeneous diffusion and a pixel density of only 5%, astonishing results can be achieved. One should note that we did not optimise our algorithm with respect to its runtime, as we regard it as a proof-of-concept only. Thus, our methods can require several hours to days to process typical images. However, we are confident that this runtime can be significantly reduced, and are going to address this issue in our ongoing research.

Organisation of the paper. Section 2 gives a brief introduction to homogeneous diffusion inpainting. In Section 3 we present two approaches that are applied sequentially to optimise the pixel locations: a probabilistic sparsification method, followed by a nonlocal pixel exchange. Afterwards, in Section 4, we show how the results can be improved further by an exact optimisation of the tonal data. Finally, we summarise and conclude our paper in Section 5.

Related work. The most similar work to our paper is a recent publication by Belhachmi et al. [2] where a continuous analysis on spatially optimal data selection for homogeneous diffusion interpolation is presented. Their framework is based on the theory of shape optimisation and suggests to choose a pixel density that is an increasing function of the modulus of the image Laplacian. In order to make this result applicable to the practically relevant discrete setting, dithering techniques must be applied that can introduce additional errors. In the experiments we compare our results with the ones from [2]. It should be mentioned that in [2] no tonal optimisation is performed.

There is a long tradition to restore image data by homogeneous diffusion inpainting from edges [5, 7, 9, 13, 20] or specific feature points in Gaussian scale space [10–12]. Although such features can be perceptually relevant, one cannot expect that they are optimal w.r.t. some error norm.

In order to come up with data-adaptive point distributions, some publications use subdivision strategies in connection with anisotropic diffusion [8, 18]. They offer the advantage that the resulting tree structures allow an inexpensive coding of the selected pixels, but they severely constrain the set of admissible point distributions. For a more sensitive interpolant such as homogeneous diffusion, this restriction is too prohibitive.

The holographic image representation presented in [4] maps the image into a sequence of sample pixels, such that any partition of this sequence allows for a reconstruction of the whole image with similar quality. This requires the samples in each portion to be equally optimal. On the contrary, our goal is to reduce the image to only one set of optimal samples.

From the Green function of the Laplace operator it follows that homogeneous diffusion inpainting involves radial basis functions. These functions are popular for scattered data interpolation, and some of them have also been used for inpainting corrupted images [6, 19]. However, such problems usually do not allow to optimise the location and the grey values of the inpainting data set.

2 Image Inpainting with Homogeneous Diffusion

Continuous formulation. Let $f(\mathbf{x})$ be a continuous grey value image, where $\mathbf{x} = (x, y)^\top$ denotes the location within a rectangular image domain $\Omega \subset \mathbb{R} \times \mathbb{R}$. Furthermore, let $\Omega_K \subset \Omega$ be a subset of the image domain, denoting *known* data. A reconstruction $u(\mathbf{x})$ by means of homogeneous diffusion inpainting can be obtained by keeping known data and using them as Dirichlet boundary conditions, while solving the Laplace equation on the set of *unknown* data $\Omega \setminus \Omega_K$:

$$\begin{aligned} u(\mathbf{x}) &= f(\mathbf{x}) & \text{for } \mathbf{x} \in \Omega_K, \\ \Delta u(\mathbf{x}) &= 0 & \text{for } \mathbf{x} \in \Omega \setminus \Omega_K, \end{aligned} \quad (1)$$

with homogeneous (reflecting) Neumann boundary conditions across the image boundary $\partial\Omega$. These two equations can be combined to a single equation

$$c(\mathbf{x})(u(\mathbf{x}) - f(\mathbf{x})) - (1 - c(\mathbf{x}))\Delta u(\mathbf{x}) = 0, \quad (2)$$

by using a confidence function $c(\mathbf{x})$ which specifies whether a point is known or not:

$$c(\mathbf{x}) = \begin{cases} 1 & \text{for } \mathbf{x} \in \Omega_K, \\ 0 & \text{for } \mathbf{x} \in \Omega \setminus \Omega_K. \end{cases} \quad (3)$$

Discrete formulation. To apply the homogeneous inpainting process to a digital image, we need a discrete formulation of Equation 2. The discrete version of a continuous image f is represented as a one-dimensional vector $\mathbf{f} = (f_1, \dots, f_N)^\top = (f_i)_{i \in J}$,

where $J = \{1, \dots, N\}$ denotes the set of all pixel indices. Analogously, \mathbf{u} describes the solution vector and \mathbf{c} the binary pixel mask that indicates whether a pixel is known or not. The set K contains the pixel indices i of known pixels, i.e. for which $c_i = 1$. The Laplacian Δu is discretised by means of finite differences [15]. Then the discrete formulation reads

$$\mathbf{C}(\mathbf{u} - \mathbf{f}) - (\mathbf{I} - \mathbf{C})\mathbf{A}\mathbf{u} = \mathbf{0}, \quad (4)$$

where \mathbf{I} is the identity matrix, $\mathbf{C} := \text{diag}(\mathbf{c})$ is a diagonal matrix having the components of \mathbf{c} as diagonal entries, and \mathbf{A} is a symmetric $N \times N$ matrix, describing the discrete Laplace operator Δ with homogeneous Neumann boundary conditions. Its entries are given by

$$a_{i,j} = \begin{cases} \frac{1}{h_\ell^2} & (j \in \mathcal{N}_\ell(i)) \\ - \sum_{\ell \in \{x,y\}} \sum_{j \in \mathcal{N}_\ell(i)} \frac{1}{h_\ell^2} & (j = i) \\ 0 & (\text{else}), \end{cases} \quad (5)$$

where $\mathcal{N}_\ell(i)$ are the neighbours of pixel i in ℓ -direction.

Reformulating Equation 4 yields a linear system of equations:

$$\underbrace{(\mathbf{C} - (\mathbf{I} - \mathbf{C})\mathbf{A})}_{=: \mathbf{M}} \mathbf{u} = \mathbf{C}\mathbf{f}. \quad (6)$$

This linear system of equations has a unique solution and can be solved efficiently by using bidirectional multigrid methods [13].

3 Optimising Spatial Data

Now that we know how an image can be reconstructed by means of homogeneous diffusion inpainting, let us optimise the spatial data. This means we are looking for a pixel mask that selects for example only 5% of all pixels and that minimises the reconstruction error.

The good news on the pixel selection is that it is a discrete problem and thus it is finite and a global optimum exists. The bad news is that selecting the best 5% pixels of a 256×256 image offers already $\binom{65536}{3277} \approx 1.72 \cdot 10^{5648}$ possible solutions.

We overcome this problem by introducing two optimisation approaches. The first one is the *probabilistic sparsification*, which step by step removes pixels until the desired amount of pixels is left. Since this method can be trapped in local minima, we apply in a second step a method which we call *nonlocal pixel exchange*. It takes the mask that was created by the probabilistic sparsification and tries to improve the result by globally exchanging mask pixels with non-mask pixels.

3.1 Probabilistic Sparsification

Given an fixed discrete image \mathbf{f} , let $r(\mathbf{c}, \mathbf{f})$ be a function which computes the solution \mathbf{u} of the discrete homogeneous inpainting process described by Equation 6, depending

Input:	Original image \mathbf{f} , fraction p of mask pixels used for candidate set, fraction q of candidate pixels that are finally removed, desired pixel density d .
Output:	Pixel mask \mathbf{c} , s.t. $\sum_{i \in J} c_i = d \cdot J $.
Initialisation:	$\mathbf{c} := (1, \dots, 1)^\top$, thus $K = J$.

While $|K| > d \cdot |J|$ do

1. Choose randomly $p \cdot |K|$ pixel indices from K into a candidate set T .
2. For all $i \in T$ reassign $c_i := 0$.
3. Compute $\mathbf{u} := r(\mathbf{c}, \mathbf{f})$.
4. For all $i \in T$ compute the local error $e_i = |u_i - f_i|$.
5. For all i of the $(1 - q) \cdot |T|$ largest values of $\{e_i | i \in T\}$, reassign $c_i := 1$.
6. Update K and clear T .

Fig. 2. Probabilistic sparsification

on a mask \mathbf{c} . Our goal is to obtain a pixel mask \mathbf{c} , marking only a predefined fraction d of all pixels J such that the *mean squared error (MSE)*

$$\text{MSE}(\mathbf{u}) = \frac{1}{|J|} \sum_{i \in J} (f_i - u_i)^2 \quad (7)$$

is minimised. To obtain a suitable, approximatively optimal pixel mask, we suggest a method that we call *probabilistic sparsification*.

In each iteration, we first randomly remove a fraction of mask pixels, inpaint, compute the error in each removed pixel, and put a subset of the removed pixels with largest error back into the mask again. Thus, pixels which are supposed to be least significant are step by step removed until the desired fraction of pixels remains. The algorithm is given in detail in Figure 2.

Note that our algorithm removes $p \cdot q \cdot |K|$ pixels in each step. Thus, in the k -th iteration, there are $(1 - pq)^k \cdot |J|$ mask pixels left, since K is initially J . In total, we need $\log_{(1-pq)} d$ many steps to obtain the desired fraction d . Hence, the larger p and q are chosen, the faster the algorithm converges. On the other hand, it is then more likely that significant pixels are removed. Since we aim for an optimal pixel mask, we suggest to choose small values, such as $p = 0.02$ and $q = 0.02$.

3.2 Nonlocal Pixel Exchange

The previously presented method has the disadvantage that once a pixel is removed from the mask, it will never be put back again. Moreover, by selecting pixels randomly, it might be possible that we also remove some significant pixels. To this end, we add a post-optimisation step, called *nonlocal pixel exchange*.

In each iteration, we choose randomly a fixed amount of non-mask pixels into a candidate set. A subset of those which exhibit the largest inpainting error, are exchanged with randomly chosen mask pixels. If the inpainting error for the new mask does not decrease, we reset the mask to its previous configuration. Thus, we allow mask pixels to

Input:	Original image \mathbf{f} , (pre-optimised) pixel mask \mathbf{c} , size m of candidate set and number n of mask pixels exchanged per iteration.
Output:	Post-optimised pixel mask \mathbf{c} .
Initialisation:	$\mathbf{u} := r(\mathbf{c}, \mathbf{f})$ and $\mathbf{c}^{\text{new}} := \mathbf{c}$.
Repeat:	
1.	Choose randomly $m \leq K $ pixel indices from $J \setminus K$ into a candidate set T and compute for all $i \in T$ the local error $e_i = u_i - f_i $.
2.	Choose randomly $n \leq T $ pixel indices i from K and reassign $c_i^{\text{new}} := 0$.
3.	For all i of the n largest values of $\{e_i i \in T\}$, reassign $c_i^{\text{new}} := 1$.
4.	Compute $\mathbf{u}^{\text{new}} := r(\mathbf{c}^{\text{new}}, \mathbf{f})$.
5.	If $\text{MSE}(\mathbf{u}) > \text{MSE}(\mathbf{u}^{\text{new}})$ $\mathbf{u} := \mathbf{u}^{\text{new}}$ and $\mathbf{c} := \mathbf{c}^{\text{new}}$. Update K .
else	Reset $\mathbf{c}^{\text{new}} := \mathbf{c}$.
6.	Clear T .

Fig. 3. Nonlocal pixel exchange.

move globally as long as the reconstruction result improves. The details of the algorithm are given in Figure 3.

For our experiments, we exchange only one pixel per iteration ($n = 1$) and keep the candidate set small by choosing $m = 10$. Interestingly, the results cannot be improved by larger candidate sets. Restricting it to this size adds some moderate amount of randomness such that we are not trapped in the next local minimum.

3.3 Results

Let us now evaluate the capabilities of the probabilistic sparsification and the nonlocal pixel exchange. To this end, we consider the test image *trui*, which is depicted in Figure 1(a). We apply the probabilistic sparsification to select only 5% of all pixels. For comparison, we choose the same amount of pixels randomly. In addition, we compare our method with an inpainting mask which relies on the analytic approach of Belhachmi et al. [2]: We first compute the Laplace magnitude $|\Delta f_\sigma|$ of the Gaussian presmoothed original image, using a standard deviation σ . Then we rescale the obtained data and apply electrostatic halftoning [17] such that we obtain a dithered version which contains only 5% of all pixels. We decided to favour the electrostatic halftoning over simpler dithering approaches, since it has proven to be the state-of-the-art method for discretising a continuous distribution function. In the following, we say ‘‘analytic approach’’ when referring to this method and choose the standard deviation σ of the Gaussian presmoothing such that the MSE is minimal.

The resulting masks as well as the corresponding reconstruction results are depicted in Figure 6(a)–(c) and (e)–(g). As expected, the random mask gives poor quality reconstructions. Comparing the reconstructed images of the analytic approach and the




Fig. 4. Convergence behaviour of the nonlocal pixel exchange ($m = 10$, $n = 1$, 500,000 iterations) applied to the mask obtained by the probabilistic sparsification (see Figure 6(c)) with 5% of all pixels.

probabilistic sparsification, we observe that the latter has a lower reconstruction error. This shows that we cannot immediately deduce an optimal pixel set from the optimal continuous theory.

If we now additionally apply the nonlocal pixel exchange to the mask that was obtained by the probabilistic sparsification, we also get a visually more pleasant result (see Figure 6(d) and (h)). The MSE is decreased to 23.21, which is much better than the MSE of the analytic approach that is 49.47.

The plot in Figure 4 shows that the nonlocal pixel exchange achieves the most significant improvement during the first 50,000 iterations. On the other hand, it illustrates that after 500,000 iterations the real optimum is still not reached, even though we are probably rather close to it.

4 Optimising Tonal Data

So far we explained how to obtain approximatively optimal positions for a predefined amount of pixels. This is considered as spatial data optimisation. However, it is also possible to optimise the data with respect to the tonal data (i.e. the co-domain).

For inpainting, we usually use the original grey values of the input image. Now we allow arbitrary grey values and thus accept to introduce some error at the positions of mask pixels in favour of a lower overall reconstruction error.

4.1 Grey Value Optimisation

Let us start by stressing that the homogeneous inpainting function $r(\mathbf{c}, \mathbf{f})$ is a linear function with respect to the grey values \mathbf{f} . This allows us to formulate a least squares approach, with which we can compute the optimal grey values for a given mask exactly.

For a given mask \mathbf{c} and given data \mathbf{f} the solution $\mathbf{u} = r(\mathbf{c}, \mathbf{f})$ of the discrete homogeneous inpainting process (6) is given by

$$r(\mathbf{c}, \mathbf{f}) := \mathbf{M}^{-1} \mathbf{C} \mathbf{f} . \quad (8)$$

Since \mathbf{M} only depends on \mathbf{c} it follows directly that r is a linear function in \mathbf{f} .

Least squares approximation. Our goal is to find grey values \mathbf{g} such that $\text{MSE}(r(\mathbf{c}, \mathbf{g}))$ becomes minimal for a fixed mask \mathbf{c} . To this end, we suggest the following minimisation approach:

$$\operatorname{argmin}_{\boldsymbol{\alpha}} \|\mathbf{f} - r(\mathbf{c}, \mathbf{f} + \boldsymbol{\alpha})\|^2 , \quad (9)$$

such that $\mathbf{g} = \mathbf{f} + \boldsymbol{\alpha}$. Let \mathbf{e}_i denote the vector with a 1 in the i -th coordinate and zeros elsewhere. Then we call $r(\mathbf{c}, \mathbf{e}_i)$ the *inpainting echo* of the i -th pixel. By linearity and $\boldsymbol{\alpha} = \sum_{i \in J} \alpha_i \mathbf{e}_i$ it follows that

$$r(\mathbf{c}, \mathbf{f} + \boldsymbol{\alpha}) = r(\mathbf{c}, \mathbf{f}) + r(\mathbf{c}, \boldsymbol{\alpha}) = r(\mathbf{c}, \mathbf{f}) + \sum_{i \in J} \alpha_i r(\mathbf{c}, \mathbf{e}_i). \quad (10)$$

Since the $r(\mathbf{c}, \mathbf{e}_i)$ is $\mathbf{0}$ if $c_i = 0$ (i.e. $i \in J \setminus K$), we get

$$r(\mathbf{c}, \mathbf{f} + \boldsymbol{\alpha}) = r(\mathbf{c}, \mathbf{f}) + \sum_{i \in K} \alpha_i r(\mathbf{c}, \mathbf{e}_i). \quad (11)$$

For our minimisation problem (9), this means that α_i can be chosen arbitrarily if $i \in J \setminus K$. Thus, for the sake of simplicity, we set $\alpha_i = 0$ for $i \in J \setminus K$. The remaining α_i with $i \in K$ can be obtained by considering the least squares problem:

$$\operatorname{argmin}_{\boldsymbol{\alpha}_K} \|\mathbf{U}\boldsymbol{\alpha}_K - \mathbf{b}\|^2, \quad (12)$$

where $\mathbf{b} = r(\mathbf{c}, \mathbf{f}) - \mathbf{f}$ is a vector of size $|J|$, $\boldsymbol{\alpha}_K = (\alpha_i)_{i \in K}$ is a vector of size $|K|$, and \mathbf{U} is a $|J| \times |K|$ matrix which contains the vectors $r(\mathbf{c}, \mathbf{e}_i)$, $i \in K$ as columns.

Its solution is given by solving the normal equations:

$$\mathbf{U}^\top \mathbf{U} \boldsymbol{\alpha}_K = \mathbf{U}^\top \mathbf{b}. \quad (13)$$

Let us first prove that the matrix $\mathbf{U}^\top \mathbf{U}$ is invertible: Since \mathbf{U} contains the vectors $r(\mathbf{c}, \mathbf{e}_i)$, $i \in K$ as columns, it is sufficient to show that the vectors $r(\mathbf{c}, \mathbf{e}_i)$ with $i \in K$ are linearly independent. It holds that

$$r(\mathbf{c}, \mathbf{e}_i) = \mathbf{M}^{-1} \mathbf{C} \mathbf{e}_i \stackrel{i \in K}{=} \mathbf{M}^{-1} \mathbf{e}_i. \quad (14)$$

Hence, $r(\mathbf{c}, \mathbf{e}_i)$ is the i -th column of \mathbf{M}^{-1} and since \mathbf{M}^{-1} exists [13], the vectors $r(\mathbf{c}, \mathbf{e}_i)$ have to be linearly independent. Thus, $\mathbf{U}^\top \mathbf{U}$ is invertible.

Iterative approach. The linear system given by Equation 13 can be solved exactly by using standard methods such as an LU-decomposition. Since this is rather slow, we suggest the following iterative solver.

Let us for a moment consider the simplified optimisation problem, where a vector \mathbf{g} and the inpainting result $\mathbf{u} = r(\mathbf{c}, \mathbf{g})$ are initially given. We want to optimise only the i -th grey value and keep the remaining grey values fixed:

$$\operatorname{argmin}_{\alpha} \|\mathbf{f} - r(\mathbf{c}, \mathbf{g} + \alpha \mathbf{e}_i)\|^2. \quad (15)$$

Then the solution is

$$\alpha = \frac{r(\mathbf{c}, \mathbf{e}_i)^\top (\mathbf{f} - \mathbf{u})}{r(\mathbf{c}, \mathbf{e}_i)^\top r(\mathbf{c}, \mathbf{e}_i)}. \quad (16)$$

The optimised grey value can be computed as $g_i := g_i + \alpha$. Moreover, provided we have precomputed all inpainting echos $r(\mathbf{c}, \mathbf{e}_i)$, $i \in K$, we can not only efficiently compute

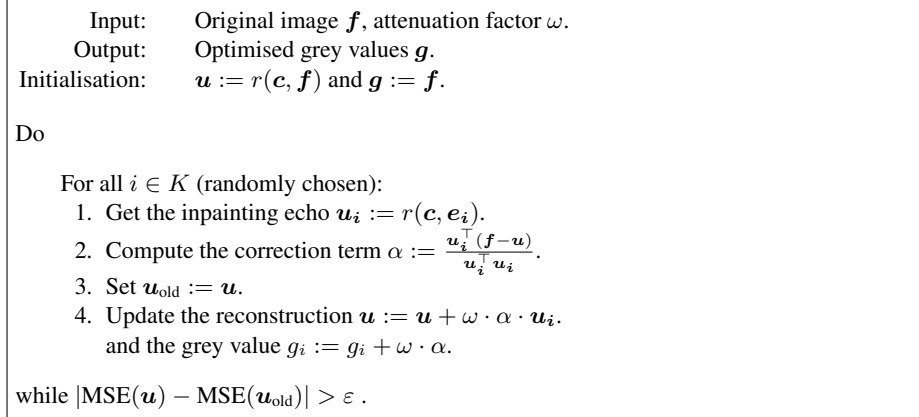


Fig. 5. Grey value optimisation.

α , but also the inpainting result for the updated image \mathbf{g} . To this end, we exploit again the linearity of r :

$$r(\mathbf{c}, \mathbf{g}) := r(\mathbf{c}, \mathbf{g} + \alpha \mathbf{e}_i) = r(\mathbf{c}, \mathbf{g}) + \alpha \cdot r(\mathbf{c}, \mathbf{e}_i) = \mathbf{u} + \alpha \cdot r(\mathbf{c}, \mathbf{e}_i). \quad (17)$$

If we apply this optimisation for each $i \in K$ iteratively, and update the grey values in each step directly, we obtain an algorithm which corresponds to the Gauss-Seidel method [16] for the previously presented linear system of equations (see Equation 13).

Optimising one grey value at a time means that this grey value might be shifted extremely in order to reduce the inpainting error in its neighbourhood. However, there could be mask points nearby which are not optimised yet and a combined optimisation would lead to a smaller shift for each of them. Thus, to prevent over- and undershoots we suggest to introduce an attenuation factor ω . This can be seen as a variant of the so-called *successive over-relaxation method (SOR)* [16] with under-relaxation instead of over-relaxation. Figure 5 summarises our iterative algorithm.

We terminate our algorithm when the qualitative improvement from one to the next iteration step decreases to a value smaller than $\varepsilon = 0.001$. Moreover, note that we choose the indices $i \in K$ randomly in each run. This allows more stable results, since we do not rely on a specific pixel ordering for each run and thus the approximation error is better distributed over the whole image.

4.2 Results

We apply the presented grey value optimisation to the test image *trui* and the mask obtained by our spatial optimisation method (see Figure 6(d)). However, we actually can use the grey value optimisation to optimise the grey values for any fixed mask. Thus, for the sake of comparison, we also consider the random mask, the mask obtained by the analytic approach, and the mask created with the probabilistic sparsification (see Figure 6(a),(b) and (c)). The reconstruction results are depicted in the last row of Figure 6.

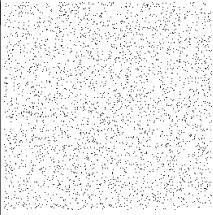
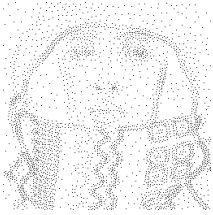
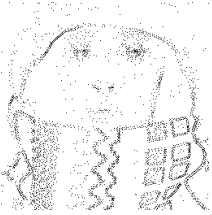





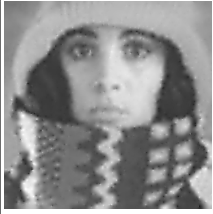



	randomly selected	analytic approach	probabilistic sparsification	+ nonlocal pixel exchange
mask				
MSE	189.90	49.47	41.08	23.21
reconstruction				
MSE	106.17	31.62	20.68	17.17
reconstruction with optimal tonal data				

Fig. 6. Evaluation of different inpainting data using 5% of all pixels. **Top row:** Different masks obtained by (a) random selection, (b) analytic approach ($\sigma = 1.44$), (c) probabilistic sparsification ($p = 0.02$, $q = 0.02$, $d = 0.05$), (d) nonlocal pixel exchange ($m = 10$, $n = 1$, 500,000 iterations) applied to (c). **Middle row:** (e-h) Reconstructions with homogeneous diffusion inpainting using the masks (a-d). **Bottom row:** (i-l) As middle row but using optimal tonal data.

For all examples the MSE has decreased. However, we observe that the worse the spatial data are selected, the larger the improvement that can be achieved by the grey value optimisation. The explanation for this behaviour is simple. Both our spatial optimisation method as well as the analytic approach select the pixels depending on the grey values of the original image. Thus, the spatial data are optimised by incorporating these tonal data. If we choose random spatial data, it is more likely that we can compensate bad locations by adapting the grey values.

Besides this observation, the smallest MSE of only 17.17 is obtained by the combination of the presented spatial and tonal optimisation methods. To further evaluate this combined approach, we apply it to two other test images. The results are depicted in Figure 7. Moreover, Table 1 gives a comparison with the analytic approach and the results obtained with a random pixel mask. In all cases, we obtain by far the best




Fig. 7. Reconstruction results with 5% of all pixels, spatially and tonally optimised, for the test images *walter* and *peppers256*.

Table 1. Comparison of the reconstruction error (MSE) with 5% of all pixels for different test images and different inpainting data.

	unoptimised (randomly selected)	analytic approach ($ \Delta f_\sigma $ dithered)	spatially and tonally optimised
<i>trui</i>	198.90	49.47 ($\sigma = 1.44$)	17.17
<i>walter</i>	183.37	24.59 ($\sigma = 1.37$)	8.14
<i>peppers256</i>	179.22	49.71 ($\sigma = 1.15$)	19.38

reconstruction results with our new approach. This confirms that it is a suitable method for the selection of optimal inpainting data.

5 Conclusion

While many researchers have tried to find highly sophisticated PDEs for inpainting problems with given data, we have investigated the opposite way: finding optimal data for a given PDE. We have shown that even for the simplest inpainting PDE, namely homogeneous diffusion, one can obtain reconstructions of astonishing quality using only 5% of all pixels. However, this requires to optimise the data carefully in the domain and the co-domain.

Since we are able to reduce the amount of data needed for high quality reconstructions drastically, our ongoing research addresses the problem how these data can be encoded efficiently. This includes appropriate adaptations of the grey value optimisation to quantised data. Moreover, we are interested in applying our optimisation framework also to nonlinear inpainting methods. As a result, we might obtain similar qualitative reconstructions with even less data, allowing further cuts in our pixel budget.

Acknowledgements. Our research is partly funded by the Deutsche Forschungsgemeinschaft (DFG) through a Gottfried Wilhelm Leibniz Prize. This is gratefully acknowledged. We also thank Pascal Gwosdek and Christian Schmalz for providing the electrostatic halftoning images for us.

References

1. Bae, E., Weickert, J.: Partial differential equations for interpolation and compression of surfaces. In: Daehlen, M., Floater, M., Lyche, T., Merrien, J.L., Mørken, K., Schumaker, L.L. (eds.) *Mathematical Methods for Curves and Surfaces*, Lecture Notes in Computer Science, vol. 5862, pp. 1–14. Springer, Berlin (2010)
2. Belhachmi, Z., Bucur, D., Burgeth, B., Weickert, J.: How to choose interpolation data in images. *SIAM Journal on Applied Mathematics* 70(1), 333–352 (2009)
3. Bertalmío, M., Sapiro, G., Caselles, V., Ballester, C.: Image inpainting. In: *Proc. SIGGRAPH 2000*. pp. 417–424. New Orleans, LA (Jul 2000)
4. Bruckstein, A.M., Holt, R.J., Netravali, A.N.: Holographic representations of images. *IEEE Transactions on Image Processing* 7(11), 1583–1597 (Nov 1998)
5. Carlsson, S.: Sketch based coding of grey level images. *Signal Processing* 15, 57–83 (1988)
6. Di Blasi, G., Francomano, E., Tortorici, A., Toscano, E.: A smoothed particle image reconstruction method. *Calcolo* (in press)
7. Elder, J.H.: Are edges incomplete? *International Journal of Computer Vision* 34(2/3), 97–122 (1999)
8. 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)
9. Hummel, R., Moniot, R.: Reconstructions from zero-crossings in scale space. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 37, 2111–2130 (1989)
10. Johansen, P., Skelboe, S., Grue, K., Andersen, J.D.: Representing signals by their toppoints in scale space. In: *Proc. Eighth International Conference on Pattern Recognition*. pp. 215–217. Paris, France (Oct 1986)
11. Kanters, F.M.W., Lillholm, M., Duits, R., Jansen, B.J.P., Platel, B., Florack, L.M.J., ter Haar Romeny, B.M.: On image reconstruction from multiscale top points. In: Kimmel, R., Sochen, N., Weickert, J. (eds.) *Scale Space and PDE Methods in Computer Vision*, Lecture Notes in Computer Science, vol. 3459, pp. 431–439. Springer, Berlin (2005)
12. Lillholm, M., Nielsen, M., Griffin, L.D.: Feature-based image analysis. *International Journal of Computer Vision* 52(2/3), 73–95 (2003)
13. Mainberger, M., Bruhn, A., Weickert, J., Forchhammer, S.: Edge-based image compression of cartoon-like images with homogeneous diffusion. *Pattern Recognition* (in press)
14. 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)
15. Morton, K.W., Mayers, L.M.: *Numerical Solution of Partial Differential Equations*. Cambridge University Press, Cambridge, UK, second edn. (2005)
16. Saad, Y.: *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia, second edn. (2003)
17. Schmaltz, C., Gwosdek, P., Bruhn, A., Weickert, J.: Electrostatic halftoning. *Computer Graphics Forum* 29(8), 2313–2327 (Dec 2010)
18. Schmaltz, C., Weickert, J., Bruhn, A.: Beating the quality of JPEG 2000 with anisotropic diffusion. In: Denzler, J., Notni, G., Süße, H. (eds.) *Pattern Recognition*, Lecture Notes in Computer Science, vol. 5748, pp. 452–461. Springer, Berlin (2009)
19. Uhlir, K., Skala, V.: Reconstruction of damaged images using radial basis functions. In: *Proc. 13th European Signal Processing Conference (EUSIPCO)*. pp. 160–163. Antalya, Turkey (Sep 2005)
20. Zeevi, Y., Rotem, D.: Image reconstruction from zero-crossings. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 34, 1269–1277 (1986)