

In J. Blanc-Talon, C. Distant, W. Philips, D. Popescu, P. Scheunders (Eds.):  
Advanced Concepts for Intelligent Vision Systems.  
Lecture Notes in Computer Science, Vol. 10016, 1-13, Springer, Cham, 2016.  
The final publication is available at [link.springer.com](http://link.springer.com)

## Gradients versus Grey Values for Sparse Image Reconstruction and Inpainting-Based Compression

Markus Schneider<sup>1</sup>, Pascal Peter<sup>1</sup>, Sebastian Hoffmann<sup>1</sup>, Joachim Weickert<sup>1</sup>,  
and Enric Meinhardt-Llopis<sup>2</sup>

<sup>1</sup> Mathematical Image Analysis Group  
Faculty of Mathematics and Computer Science, Campus E1.7  
Saarland University, 66041 Saarbrücken, Germany  
`{schneider,peter,hoffmann,weickert}@mia.uni-saarland.de`

<sup>2</sup> École Normale de Supérieure de Cachan  
61, Avenue du Président Wilson, 94235 Cachan, France  
`enric.meinhardt@cmla.ens-cachan.fr`

**Abstract.** Interpolation methods that rely on partial differential equations can reconstruct images with high quality from a few prescribed pixels. A whole class of compression codecs exploits this concept to store images in terms of a sparse grey value representation. Recently, Brinkmann et al. (2015) have suggested an alternative approach: They propose to store gradient data instead of grey values. However, this idea has not been evaluated and its potential remains unknown. In our paper, we compare gradient and grey value data for homogeneous diffusion inpainting w.r.t. two different aspects: First, we evaluate the reconstruction quality, given a comparable amount of data of both kinds. Second, we assess how well these sparse representations can be stored in compression applications. To this end, we establish a framework for optimising and encoding the known data. It allows a fair comparison of both the grey value and the gradient approach. Our evaluation shows that gradient-based reconstructions avoid visually distracting singularities involved in the reconstructions from grey values, thus improving the visual fidelity. Surprisingly, this advantage does not carry over to compression due to the high sensitivity to quantisation.

**Keywords:** partial differential equations (PDEs), Laplace interpolation, Poisson equation, inpainting, image compression, derivatives.

# 1 Introduction

Interpolation methods based on partial differential equations (PDEs) have been successfully used for image restoration [2, 4, 12]: So-called inpainting approaches propagate known data into missing or damaged image areas. If the known data can be freely selected from the original image, PDE-based inpainting even allows reconstructions with high quality from much sparser grey value data [8, 11, 16]. PDE-based compression codecs such as [4, 7, 17] use this fact and only store the locations and grey values of a few pixels.

Many approaches rely on homogeneous diffusion inpainting due to its simplicity and the availability of efficient solvers [4, 8, 11, 14, 16]. However, this simple differential equation has a drawback: The inpainting solution can be expressed as a superposition of Green’s functions which have singularities [13]. This behaviour often leads to unpleasant artifacts that spoil the visual perception of the reconstructions. Several existing codecs [7, 17] avoid these artifacts by using anisotropic PDEs at the price of a higher complexity and computational cost.

However, using different known data than grey values might be an alternative to circumvent the drawbacks of homogeneous diffusion inpainting: Brinkmann et al. [3] have suggested that gradient data could be useful for sparse image representations and compression. So far, this idea has not been implemented, and the potential of gradient data for PDE-based is yet to be explored.

*Our Contribution.* We fill this gap with a comparative evaluation of grey value and gradient data. First, we investigate the advantages and drawbacks of both kinds of data for sparse image representations with homogeneous diffusion inpainting. To enable a fair comparison, we embed both approaches into a common framework that provides probabilistic algorithms to optimise the known data. In particular, we examine if gradient data can avoid reconstruction artifacts that are common for homogeneous grey value inpainting. Secondly, we analyse the impact of widely used compression techniques such as quantisation and entropy coding on grey values and derivatives. With our experiments on well-known test images, we evaluate their overall compression quality, and also compare to JPEG.

*Outline.* Section 2 recaps PDE-based reconstructions from grey values and shows an image reconstruction approach from sparse gradients. The following sections also review existing concepts based on grey values, and continue with their adaption to a gradient-based setting. The data optimisation is covered in Section 3. We combine these optimisation strategies with compression techniques to form a complete codec in Section 4. We evaluate the grey value-based methods and their gradient-based counterparts in Section 5, firstly for image reconstructions from sparse data, and secondly in a compression context as proposed by Brinkmann et al. [3]. Finally, we conclude with a summary in Section 6.

## 2 PDE-Based Reconstructions from Sparse Data

### 2.1 Reconstructions from Grey Values

The goal of PDE-based inpainting is to fill in missing areas in an image  $f : \Omega \rightarrow \mathbb{R}$  with image domain  $\Omega \subset \mathbb{R}^2$ . A binary image  $c : \Omega \rightarrow \{0, 1\}$ , the so-called *mask*, characterises each pixel  $\mathbf{x}$  of  $f$  as known ( $c(\mathbf{x}) = 1$ ) or unknown ( $c(\mathbf{x}) = 0$ ). A homogeneous diffusion inpainting propagates the known data equally in all directions. The steady state solves the inpainting equation

$$c(\mathbf{x}) \cdot (u(\mathbf{x}) - f(\mathbf{x})) - (1 - c(\mathbf{x})) \cdot \Delta u(\mathbf{x}) = 0 \quad \forall \mathbf{x} \in \Omega \quad (1)$$

for  $u$  under reflecting boundary conditions. For the unknown pixels, the Laplace equation  $\Delta u = 0$  imposes a smoothness constraint. In addition, the known data stays fixed and determines the inpainting result in terms of Dirichlet boundary conditions. A discretisation of (1) leads to a linear system of equations that we solve efficiently with multigrid methods [10].

### 2.2 Reconstructions from Gradients

Let us now describe how to reconstruct an image from a few known gradients.

Given two masks  $c_x$  and  $c_y$  representing the locations of the stored  $x$ - and  $y$ -derivatives, we denote the values for the  $x$ - and  $y$ -derivatives at these positions by  $p$  and  $q$ . Because one of the two derivatives might have much more structure than the other in a certain image region or in the whole image, it makes sense to use separate masks. This allows a direct adaption of the masks to the local image structure.

The first reconstruction step performs a componentwise homogeneous diffusion inpainting of the derivatives: We apply an inpainting as given in (1) once to inpaint the  $x$ -derivative on a mask  $c_x$ , and once to inpaint the  $y$ -derivative on  $c_y$ . Together, both yield an approximation  $\mathbf{v}$  of the original gradient field  $\nabla f$ .

The second step is a numerical integration, that eliminates the singularities the inpainting suffers from. In general, the dense vector field  $\mathbf{v}$  will be non-integrable, because there is not necessarily an image  $u$  whose gradient field is  $\mathbf{v}$ . Therefore, we integrate  $\mathbf{v}$  numerically by minimising the Poisson functional

$$E(u) = \int_{\Omega} |\nabla u - \mathbf{v}|^2 d\mathbf{x} . \quad (2)$$

This process searches for a differentiable image reconstruction  $u$  whose gradient field is closest to the given vector field  $\mathbf{v}$  in terms of the squared Euclidean distance. Its minimiser  $u$  has to satisfy the Poisson equation  $\Delta u = \operatorname{div} \mathbf{v}$  with reflecting boundary conditions and is determined up to an additive constant. Requiring  $u$  to have the same mean  $\mu$  as the original image, the minimiser becomes unique. Thus, we additionally need the original mean  $\mu$  in the reconstruction step.

Altogether, the reconstruction procedure consists of the following three steps:

1. Inpaint the  $x$ -derivative by solving  $c_x(v_1 - p) - (1 - c_x)\Delta v_1 = 0$  for  $v_1$ .
2. Inpaint the  $y$ -derivative by solving  $c_y(v_2 - q) - (1 - c_y)\Delta v_2 = 0$  for  $v_2$ .
3. Minimise (2), under the constraint that  $u$  has mean  $\mu$ .

*Efficient Algorithmic Realisation.* We discretise the first derivatives by forward differences and the Laplacian by its standard discretisation. Consequently, the first derivatives are given on a shifted grid. Discretising the divergence by backward differences finally ensures a discretisation with second order of consistency.

We compute both inpaintings in parallel, as they are independent of each other. To solve the Poisson equation, we apply the discrete Fourier transform and solve the corresponding equation in the Fourier domain [6]. Hereby, the Fourier coefficient  $\mathcal{F}[u]_{0,0}$  in the origin is set to the scaled mean  $\mu$ . In summary, this yields an efficient solver, if the Fast Fourier Transform is used.

### 3 Data Optimisation for PDE-Based Reconstructions

To find a compact image representation we optimise both the *locations* of the stored grey values / gradients (*spatial optimisation*), and the grey / gradient *values* themselves at these positions (*tonal optimisation*). We consider both aspects for grey values in Section 3.1 and for gradients in Section 3.2.

#### 3.1 Optimisation Algorithms for Grey Values

*Spatial Optimisation.* We first aim at finding an optimal image representation without caring about how efficiently it can be stored later on. We call the resulting masks *exact masks*, since they allow to place known data freely and with pixel accuracy. Belhachmi et al. [1] proved that one should choose the mask pixels as an increasing function of the Laplacian magnitude. An optimal control approach of Hoeltgen et al. [8] minimises the trade-off between reconstruction error and sparsity of the known data. Another approach was introduced by Mainberger et al. [11]: Starting with a full mask, they successively remove a certain fraction of randomly chosen mask pixels which have the smallest error. They call this process probabilistic sparsification. Since a pixel is never put back into the mask, the algorithm typically runs into a local minimum. Therefore, they post-optimize the resulting mask by a nonlocal pixel exchange allowing randomised swaps between mask and non-mask pixels. This is the approach we later adapt for gradients.

*Tonal Optimisation.* For the optimisation of the grey values, we assume the mask to be fixed. We now switch to a discrete setting by reordering all the pixels of an image  $u$  row-wise into a vector  $\mathbf{u}$ . Given the grey values at a few locations, collected in a sparse vector  $\mathbf{g}$ , one can express the inpainting in matrix-vector notation as  $\mathbf{u} = \mathbf{M}\mathbf{g}$  [10]. It computes the inpainting solution  $\mathbf{u}$  from the grey values  $\mathbf{g}$ . The matrix  $\mathbf{M}$  implicitly contains the mask.

To find the best grey values to store we minimise the reconstruction error

$$\operatorname{argmin}_{\mathbf{g}} |\mathbf{M}\mathbf{g} - \mathbf{f}|^2 . \quad (3)$$

This least squares problem leads to a linear system of equations. We use an algorithm introduced by Mainberger et al. [11], which is based on so-called inpainting echoes. An inpainting echo is the reconstruction  $\mathbf{M}\mathbf{e}_i$  on the  $i$ -th unit vector  $\mathbf{e}_i$ . This allows to express the final reconstruction  $\mathbf{u}$  as a linear combination of the inpainting echoes. The flexibility of this approach allows an adaption to the gradient-based setting.

### 3.2 Optimisation Algorithms for Gradients

*Spatial Optimisation.* Since we aim at a fair comparison, we optimise both the grey values and the gradients using the probabilistic approach of Mainberger et al. [11]. We apply their algorithm on the  $x$ -derivative and on the  $y$ -derivative of the original image to generate two masks containing the positions for  $x$ - and  $y$ -derivatives. Our modified version of their nonlocal pixel exchange minimises the mean squared error in the final reconstruction  $u$  after the integration step, and the pixel exchange step allows the mask pixels to swap from one mask into the other. We thereby automatically adapt the number of mask pixels of  $\partial_x$ - and  $\partial_y$ -mask to the proportion of structures in horizontal and vertical direction in the original image. This modified nonlocal pixel exchange converges towards the global optimum. This allows us to assess the true potential of gradient data for image reconstructions.

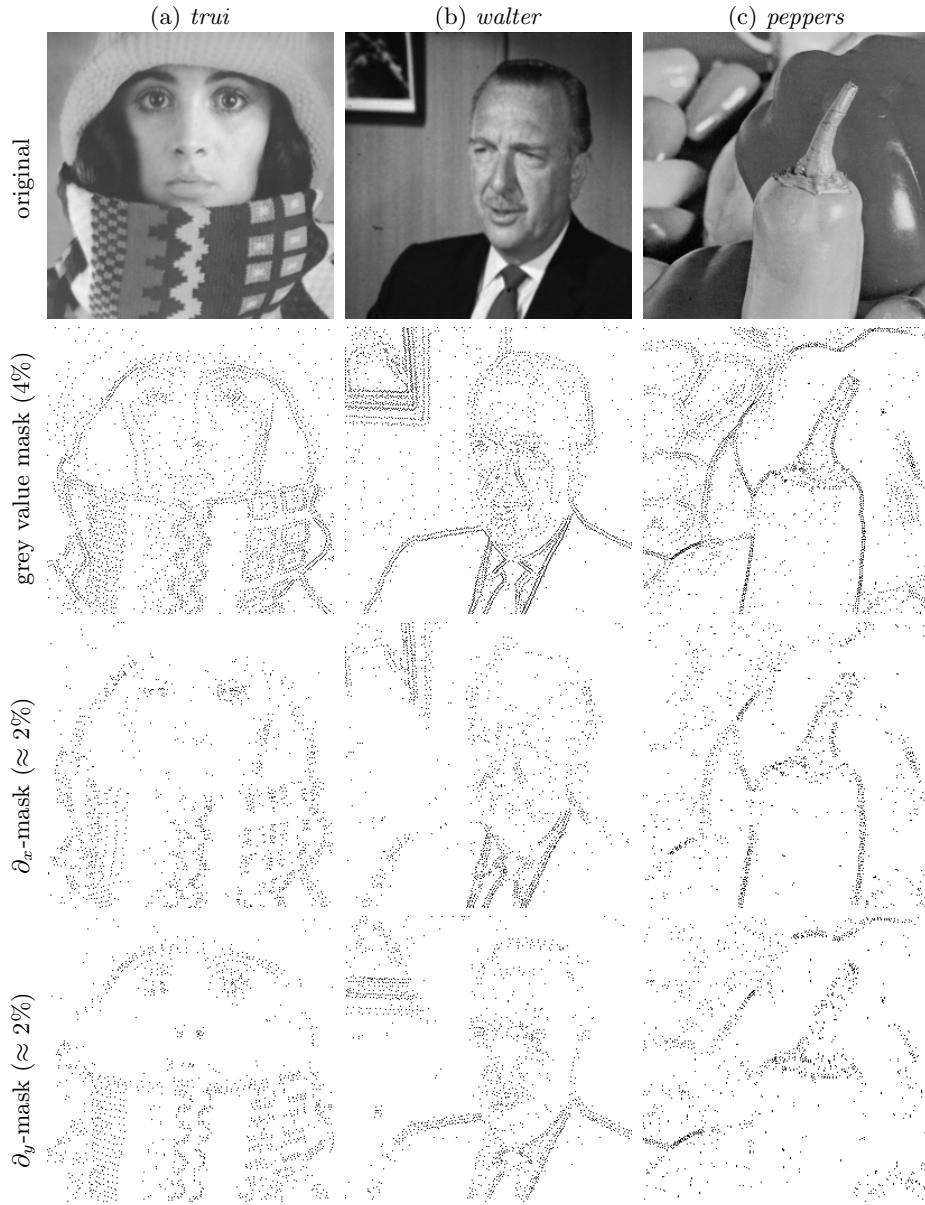
In Figure 1, we show three pairs of exact masks for the standard test images *trui*, *walter* and *peppers*. As expected, we observe that the  $\partial_x$ - and the  $\partial_y$ -mask prefer locations with structures in horizontal and vertical direction, respectively. The resulting  $\partial_x$ - and  $\partial_y$ -masks roughly have the same density, as the structures in our test images are more or less balanced between both derivative directions.

*Tonal Optimisation.* Given a mask pair  $(\mathbf{c}_x, \mathbf{c}_y)$  and the sparse stored gradients  $\mathbf{d} = (\mathbf{p}, \mathbf{q})$  at these locations, we can express the reconstruction procedure in matrix-vector notation as  $\mathbf{u} = \mathbf{L}\mathbf{d} + \boldsymbol{\mu}$ . It computes the reconstructed image  $\mathbf{u}$  from the stored gradients  $\mathbf{d}$ , which includes the inpaintings and the integration.  $\mathbf{L}$  implicitly contains the masks  $\mathbf{c}_x$  and  $\mathbf{c}_y$ . We define  $\mathbf{L}$  to return an image with zero mean. The vector  $\boldsymbol{\mu}$  represents the mean correction by an additive shift.

The minimisation problem to find the best gradients values to store reads:

$$\operatorname{argmin}_{\mathbf{d}} |(\mathbf{L}\mathbf{d} + \boldsymbol{\mu}) - \mathbf{f}|^2 \quad (4)$$

To solve this problem, we extend the idea of inpainting echoes [11] to our setting. For this purpose, we define the *reconstruction echo*  $\mathbf{r}_i := \mathbf{L}\mathbf{e}_i$  for the  $i$ -th stored data point as the overall reconstruction with zero mean on the  $i$ -th unit vector  $\mathbf{e}_i$ . Thereby, the reconstruction can be expressed as a linear combination of the



**Fig. 1.** Exact masks. **First row:** Original test images *trui*, *walter* and zoom into *peppers*. **Second row:** Masks for 4% grey values from the probabilistic approach by Mainberger et al. [11]. **Third and fourth row:**  $\partial_x$ - and  $\partial_y$ -mask from the gradient-based approach for 4% derivatives in total. The density of each derivative mask can slightly differ from 2% since we allow a flexible adaption of this proportion. Each mask pair of  $\partial_x$ - and  $\partial_y$ -mask, however, contains exactly 4% of all pixels.

reconstruction echoes followed by the mean correction:

$$\mathbf{u} = \mathbf{L}\mathbf{d} + \boldsymbol{\mu} = \mathbf{L} \sum_{i \in K} d_i \mathbf{e}_i + \boldsymbol{\mu} = \sum_{i \in K} d_i (\mathbf{L}\mathbf{e}_i) + \boldsymbol{\mu} = \sum_{i \in K} d_i \mathbf{r}_i + \boldsymbol{\mu} \quad (5)$$

Hereby,  $K$  denotes the set of all mask pixels of the  $\partial_x$ - and the  $\partial_y$ -mask. That is, if a stored gradients value  $d_i$  is changed by some value  $\alpha$ , (5) directly specifies how the reconstruction changes, namely by  $\alpha \mathbf{r}_i$ . This connection enables the adaption of the iterative algorithm of Mainberger et al. [11].

## 4 Encoding Framework

### 4.1 Encoding of Grey Value Data

The previously mentioned methods for exact masks [5, 8, 11] yield sparse image representations in the spatial domain, which can also be useful for image compression. However, this optimal data might be expensive to store, even if we use efficient coding techniques such as a block coding scheme to encode exact masks [19]. In the following, we discuss an alternative approach that incorporates a trade-off between reconstruction quality and memory cost.

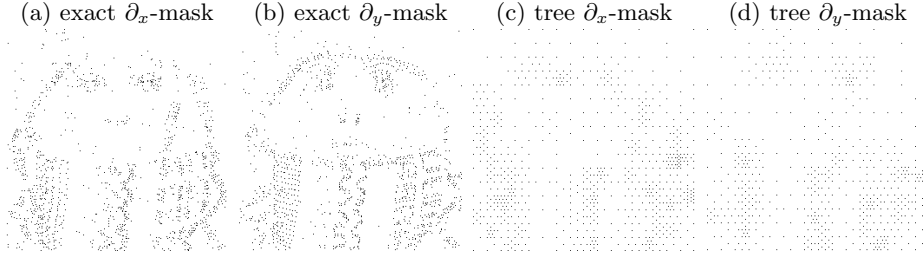
*Subdivision Trees.* Galić et al. [7] restricted the allowed positions to triangular subdivisions, which can be efficiently encoded by a binary tree. Thereby, nodes efficiently encode mask positions. Such a tree describes a subdivision of the image into regions. In areas which are difficult to reconstruct, it allows a local refinement of the mask. This concept was improved by Schmaltz et al. [17], who used rectangular subdivisions, and was adapted to a probabilistic approach by Peter et al. [16]. They transfer the probabilistic approach for exact masks to binary trees, where tree nodes play the role of pixels. This is the approach that we choose and adapt to gradients.

*Quantisation and Encoding.* To store the grey values, we have to quantise them. An equidistant quantisation, where successive quantisation levels have the same distance, is sufficient. For both exact and tree-based masks, we write the data into a single file and apply the lossless entropy coder PAQ of Mahoney [9], which is a context mixing scheme and well-suited for heterogeneous data.

### 4.2 Encoding of Gradient Data

Just as for grey values, we adapt the tree-based approach to the gradient setting. We would like to generate one tree for each gradient component.

*Subdivision Trees.* In order to adapt the stochastic tree densification of Peter et al. [16], we consider the node sets of both trees together. This automatically tailors the proportion between the size of  $\partial_x$ - and  $\partial_y$ -tree to the image: For images with more structures in  $x$ -direction, the algorithm will put more mask pixels into



**Fig. 2.** Exact mask pair and tree-based mask pair for the test image *trui* (shown in Figure 3) with approximately 2%  $x$ -derivatives and 2%  $y$ -derivatives.

the  $\partial_x$ -mask than into the  $\partial_y$ -mask. We post-optimize the resulting pair of trees by an adapted version of the nonlocal node exchange [16]. Similar to the nonlocal pixel exchange for exact masks [11], it prevents the densification from getting trapped in local minima. Besides the fact that we build all the node sets over both trees, another modification is that we consider the final reconstruction error after the integration step whenever the algorithm has to decide if a node swap is kept or not. Just as the nonlocal pixel exchange, this global nonlocal node exchange converges to the global optimum, now with respect to the restriction to the rectangular subdivisions. Figure 2 (c) and (d) show an example for our tree-based masks in comparison to a corresponding exact mask pair for *trui*.

*Quantisation of the Gradient Data.* The optimised gradients  $\mathbf{d}$  in (4) are not constrained to integer values, but given in floating point precision. However, in order to store them efficiently, they must be quantised. We use the same equidistant quantisation for  $x$ - and  $y$ -derivatives. The remaining freedom is to choose the number of quantised values  $q$  and their range, yielding two quantisation parameters.

We optimise the quantisation parameters in a direct search during the tonal optimisation and store both in the encoded image file. Within each iteration, we optimise the gradient values w.r.t. the current quantisation.

## 5 Evaluation of Gradients versus Grey Values

### 5.1 PDE-Based Image Reconstruction

For a fair comparison of gradient-based reconstructions and homogeneous grey value inpaintings, we use the optimisation strategies from Section 3 to select the same amount of data points for each method. Figure 3 shows three exemplary reconstructions from exact masks. We evaluate the reconstruction quality not only in terms of the mean squared error (MSE), but also in terms of the *mean squared Sobolev error* (MSSE)

$$\text{MSSE}(u, f) = \frac{1}{|\Omega|} \int_{\Omega} ((u - f)^2 + (u_x - f_x)^2 + (u_y - f_y)^2) \, d\mathbf{x} , \quad (6)$$



which rewards a faithful approximation of the gradient. Pure gradient-based reconstructions often achieve a better MSE than homogeneous inpainting on grey values. In terms of MSSE, the gradient-based reconstructions even yield a better error on all the three test images. Moreover, a closer look at the zoomed sections reveals that they indeed avoid the unpleasant singularities (see zooms in Figure 3). This improves the perceived quality considerably. For instance, around the textured scarf in *trui*, the improvement becomes obvious.

## 5.2 PDE-Based Image Compression



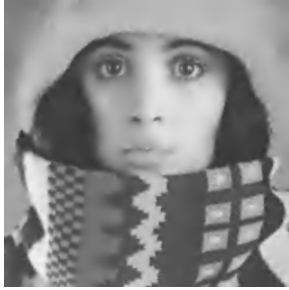















Comparing the reconstruction quality of grey value-based compared to gradient-based methods on the same amount of stored data points neglects how efficiently this data can be stored. This section evaluates the proposal of Brinkmann et al. [3] to use gradient data for compression. Thereby, we focus on three questions:

- Should one store  $x$ - and  $y$ -derivatives at different positions, or can we perform better with joint positions, saving the memory costs of a second mask?
- Does it pay off to use exact masks or should we restrict the allowed positions to the structure given by efficiently storable subdivision trees [16, 17]?
- How can we quantise the gradient data efficiently?

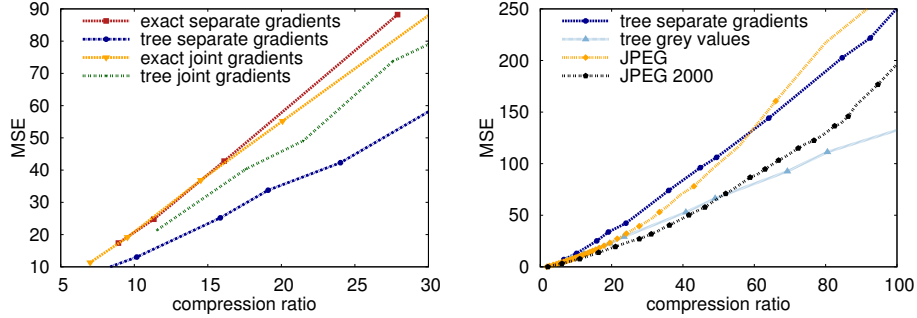
To answer the first two questions, we compare joint masks to an individual selection of known data for each derivative, and do this for exact and for tree-based masks. Finally, we compare the corresponding compression capabilities. In the second part of this section, we have a closer look at the quantisation effects.

*Mask Generation.* Figure 4 (a) shows the compression results for four variants of the gradient-based approach: For exact or for tree-based masks, we consider one joint mask or separated masks, respectively. Generating two exact masks and encoding them performs the worst, because storing two exact masks is too expensive: The proportion between the storage costs for two encoded masks and the derivative values typically is about 4:1 for exact masks, but about 1:8 for tree-based masks. One joint mask ( $c_x = c_y$ ) does not allow to get a significant improvement in the overall performance. For tree-based masks, two separate tree-based masks perform significantly better: It pays off to allow different positions for  $x$ - and  $y$ -derivatives, since these masks are fairly inexpensive to store.

*Quantisation Effects.* Figure 4 (b) compares the best gradient-based method to existing codecs. Surprisingly, the grey value inpainting now performs better than the gradient-based approach, even though its reconstruction quality was worse before. The main reason for this behaviour is the high sensitivity of gradient data to quantisation: In regions with few mask points a small error in known derivative values propagates into its neighbourhood by the inpainting and is amplified further by the integration step. Table 1 illustrates this behaviour by showing the effect of a coarser quantisation. For grey value inpainting, we see that restricting the number of quantised values does not significantly degrade the reconstruction.

(a) original	(b) grey values	(c) gradients
		
		
<i>trui</i>	MSE: 27.24, MSSE: 90.39	MSE: <b>25.09</b> , MSSE: <b>61.42</b>
		
		
<i>walter</i>	MSE: 12.45, MSSE: 39.80	MSE: <b>11.55</b> , MSSE: <b>26.93</b>
		
		
<i>peppers</i>	MSE: <b>25.10</b> , MSSE: 108.66	MSE: 28.07, MSSE: <b>104.64</b>

**Fig. 3. First column:** Test image. **Second column:** Reconstruction on exact masks with homogeneous grey value inpainting on 4% grey values. **Third column:** Our gradient-based result on 2%  $x$ - and 2%  $y$ -derivatives. The corresponding exact masks are shown in Figure 1. Zoomed sections below. The best results are highlighted in bold.



**Fig. 4.** Compression results. **(a) Left:** Comparison of two separate masks ( $c_x$ ,  $c_y$ ) versus one joint mask ( $c_x = c_y$ ), and exact versus tree-based masks. **(b) Right:** Comparison of the best gradient-based method (two separate tree-based masks) against existing codecs for *trui*, including JPEG [15] and JPEG 2000 [18].

**Table 1.** Comparison of the MSE in dependence on the quantisation for test image *trui* with 4% data points in total. **First row:** The optimal data in float precision subsequently is quantised to different quantisations. We allow  $q$  different values in the range  $[0, 255]$  for the grey values, and  $q$  different values in the range  $[-255, 255]$  for the gradient data. **Second row:** After optimisation w.r.t. the discrete quantisation.

	float precision	$q = 256$	$q = 128$	$q = 64$
grey values	27.24	27.26	27.31	27.50
optimised for discrete levels	-	27.25	27.29	27.42
gradients	25.09	56.69	80.84	209.36
optimised for discrete levels	-	26.11	27.50	37.17

For gradients, however, already small restrictions in the number of quantised values introduce large errors. Optimising the gradient values w.r.t. the discrete quantisation levels attenuates these effects again. Thus, the optimisation tailored to the discrete quantisation levels is crucial for gradient-based compression.

The gradient-based codec is capable of beating JPEG for high compression ratios. However, the limitations imposed by lossy compression steps like quantisation seem too severe to achieve better results with purely gradient-based techniques. This shows that the proposal of Brinkmann et al. [3] has severe drawbacks for image compression, even though it allows for a compact representation of an image in the gradient domain and avoids artifacts in the reconstruction.

## 6 Conclusions

We have established an evaluation framework which allows to explore the potential of gradient-based image reconstructions compared to grey value-based methods. Derivatives allow for a sparse image representation in gradient domain: Considering the same amount of known data points, gradients typically

provide a better quality in both a quantitative and a perceptual sense. In contrast to grey value-based methods, they avoid unpleasant artifacts, have a higher smoothness, and preserve the average grey value.

However, if one employs this approach for image compression as proposed by Brinkmann et al. [3], it turns out that one has to pay a high price for these benefits: Compared to grey value-based methods, the data optimisation becomes more challenging. Furthermore, the gradient data reveals a much higher sensitivity to its quantisation, which is the main reason why a pure gradient-based model in the sense of Brinkmann et al. [3] is not promising for compression. In our future work, we concentrate on models that *combine* gradient data and grey value data in order to unify the advantages of both.

## 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. Bertalmio, M., Sapiro, G., Caselles, V., Ballester, C.: Image inpainting. In: *Proc. SIGGRAPH 2000*. pp. 417–424. New Orleans, LI (Jul 2000)
3. Brinkmann, E.M., Burger, M., Grah, J.: Regularization with sparse vector fields: From image compression to TV-type reconstruction. 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. 191–202. Springer, Berlin (2015)
4. Carlsson, S.: Sketch based coding of grey level images. *Signal Processing* 15(1), 57–83 (Jul 1988)
5. 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)
6. Frankot, R.T., Chellappa, R.: A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10(4), 439–451 (Jul 1988)
7. 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)
8. 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)
9. Mahoney, M.: Adaptive weighing of context models for lossless data compression. *Tech. Rep. CS-2005-16*, Florida Institute of Technology, Melbourne, FL (Dec 2005)
10. Mainberger, M., Bruhn, A., Weickert, J., Forchhammer, S.: Edge-based compression of cartoon-like images with homogeneous diffusion. *Pattern Recognition* 44(9), 1859–1873 (Sep 2011)
11. 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.) *Proc. Third International Conference on Scale Space and Variational Methods in Computer Vision, Lecture Notes in Computer Science*, vol. 6667, pp. 26–37. Springer, Berlin (Jun 2011)

12. 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)
13. Melnikov, Y.A., Melnikov, M.Y.: Green’s Functions: Construction and Applications. De Gruyter, Berlin (2012)
14. Ochs, P., Chen, Y., Brox, T., Pock, T.: iPiano: Inertial proximal algorithm for nonconvex optimization. *SIAM Journal on Applied Mathematics* 7(2), 1388–1419 (Jun 2014)
15. Pennebaker, W.B., Mitchell, J.L.: JPEG: Still Image Data Compression Standard. Springer, New York (1992)
16. Peter, P., Hoffmann, S., Nedwed, F., Hoeltgen, L., Weickert, J.: Evaluating the true potential of diffusion-based inpainting in a compression context. *Signal Processing: Image Communication* 46, 40–53 (Aug 2016)
17. 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)
18. Taubman, D.S., Marcellin, M.W. (eds.): JPEG 2000: Image Compression Fundamentals, Standards and Practice. Kluwer, Boston (2002)
19. Zeng, G., Ahmed, N.: A block coding technique for encoding sparse binary patterns. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 37(5), 778–780 (May 1989)