



# Deep spatial and tonal data optimisation for homogeneous diffusion inpainting

Pascal Peter<sup>1</sup> · Karl Schrader<sup>1</sup> · Tobias Alt<sup>1</sup> · Joachim Weickert<sup>1</sup>

Received: 9 September 2022 / Accepted: 21 March 2023  
© The Author(s) 2023

## Abstract

Diffusion-based inpainting can reconstruct missing image areas with high quality from sparse data, provided that their location and their values are well optimised. This is particularly useful for applications such as image compression, where the original image is known. Selecting the known data constitutes a challenging optimisation problem, that has so far been only investigated with model-based approaches. So far, these methods require a choice between either high quality or high speed since qualitatively convincing algorithms rely on many time-consuming inpaintings. We propose the first neural network architecture that allows fast optimisation of pixel positions and pixel values for homogeneous diffusion inpainting. During training, we combine two optimisation networks with a neural network-based surrogate solver for diffusion inpainting. This novel concept allows us to perform backpropagation based on inpainting results that approximate the solution of the inpainting equation. Without the need for a single inpainting during test time, our deep optimisation accelerates data selection by more than four orders of magnitude compared to common model-based approaches. This provides real-time performance with high quality results.

**Keywords** Image inpainting · Diffusion · Partial differential equations · Data optimisation · Deep learning

## 1 Introduction

The classical inpainting problem [1–4] deals with input images that have been partially corrupted and aims at reconstructing these missing areas. However, inpainting can be also useful when the whole image is known. For inpainting-based image compression [5–18], the encoder stores only a small percentage of known data from which the decoder restores the discarded remainder of the image with inpainting. Some approaches [5–7, 10–12] such as the pioneering work of Carlsson [5] limit the choice of known data to

edge locations. Following the diffusion-based codec of Galić et al. [8, 19], many later approaches [14–16, 20] rely on careful optimisation of the placement of known data in the image domain without the restriction to semantic image features. Inpainting with partial differential equations (PDEs) [5] has been able to outperform state-of-the-art codecs: Already simple homogeneous diffusion [21] can compress depth-maps or flow fields better than HEVC [22] with suitably selected known data [20, 23, 24]. The problem of choosing the right positions of mask pixels, the so-called inpainting mask, is also vital for other applications such as denoising [25] or adaptive sampling [26]. In addition to this *spatial optimisation*, compression also benefits from *tonal optimisation*: The values of the known pixels can be adjusted to optimise the reconstruction quality as well.

However, even for a simple inpainting operator, spatial and tonal optimisation constitute challenging problems. This sparked a plethora of non-neural approaches [14–16, 20, 24, 27–41]. We systematically review those in Sect. 1.2. Among these methods, most require many inpaintings per iteration, which tend to be computationally expensive or rely on sophisticated implementations for acceleration. For instance, probabilistic methods for spatial optimisation [34,

---

✉ Pascal Peter  
peter@mia.uni-saarland.de  
Karl Schrader  
schrader@mia.uni-saarland.de  
Tobias Alt  
alt@mia.uni-saarland.de  
Joachim Weickert  
weickert@mia.uni-saarland.de

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

[37] yield high quality masks, but come with a high computational cost. Theoretical optimality results are rare, but have been derived from shape optimisation [27] for homogeneous diffusion inpainting. This allows near instantaneous spatial optimisation without the need for a single inpainting. However, so far, existing discrete implementations of this concept do not realise the full potential of the theoretical results from the continuous setting.

With our deep data optimisation for homogeneous diffusion inpainting, we aim for the best of both worlds. We train neural networks that can optimise both mask positions and values without the need for a single inpainting. This allows real-time performance while maintaining competitive quality for our data selection. During training, we leverage new hybrid concepts that combine model-based inpainting with deep learning.

## 1.1 Our contribution

We propose a deep learning framework for inpainting with homogeneous diffusion. It is the first neural network approach that allows tonal optimisation in addition to the selection of spatial positions. In order to merge model-based inpainting relying on PDEs with learning ideas, we propose the concept of a *surrogate solver*: During training, a neural network efficiently and accurately emulates diffusion-based inpainting while allowing for straightforward backpropagation. This concept enables us to train spatial optimisation networks that generate inpainting masks and tonal networks that output optimised known pixel values for a given mask.

Our contributions extend our previous conference contribution on deep mask learning [42] in four distinct ways:

1. In addition to a learning approach for spatial optimisation, we also propose the first tonal optimisation network.
2. We improve the network architecture and investigate the impact of individual components.
3. We discuss and evaluate options to directly generate binary masks during training [43]. With an ablation study we show that our surrogate solver is robust under non-binary mask optimisation.
4. By extending our approach to colour data and performing experiments on significantly larger databases we provide further evidence for the practical relevance of deep data optimisation.

Overall, this constitutes the first comprehensive deep learning framework for data optimisation targeted at diffusion-based inpainting. The resulting mask network marries the quality of probabilistic approaches [37] with the computational efficiency of instantaneous spatial optimisation [27]. Similarly, our neural tonal optimisation consistently offers

real-time performance at good quality. Compared to model-based approaches, its speed is independent of the amount of known data in the inpainting mask. In addition, our networks do not require any parameter tuning after training, making them attractive for practical applications.

## 1.2 Related work

In the following, we discuss prior work for spatial and tonal optimisation, as well as related deep learning approaches.

### 1.2.1 Spatial optimisation

Finding good positions for sparse known pixels constitutes a challenging optimisation problem that has sparked significant research activities. In the following, we mostly focus on approaches for diffusion-based inpainting, but there are also more broadly related works, for instance the free knot problem for spline interpolation. For instance, Schütze and Schwetlick [44] have proposed a data selection algorithm for the 2-D setting which can also be applied to images. Model-based methods for diffusion inpainting can be organised in four categories.

1. *Analytic Approaches*. From the theory of shape optimisation, Belhachmi et al. [27] derived optimality statements in the continuous setting. In practice, these can be approximated by dithering the Laplacian magnitude of the input image. This approach does not require inpainting to find the mask pixels and is therefore very fast. However, the dithering yields only an imperfect approximation with limited quality [34, 37].
2. *Nonsmooth Optimisation Strategies*. Combining concepts from optimal control with sophisticated strategies such as primal-dual solvers, multiple works [28, 29, 33, 39, 40] leverage nonsmooth optimisation for the selection of mask positions. These produce results of high quality, but are difficult to adapt to different inpainting operators. Moreover, they do not allow to specify the target amount of mask points a priori. For applications in compression, the non-binary masks need to be binarised, which reduces quality and requires tonal optimisation [35] for good results.
3. *Sparsification Methods*. Mainberger et al. [37] have proposed *probabilistic sparsification (PS)* to tackle the combinatorial complexity of spatial optimisation. They start with a full mask and iteratively remove candidate pixels. Among those candidates the algorithm discards a fraction of pixels with the smallest inpainting error permanently, while returning the remainder to the mask. This process is repeated until the desired percentage of mask points, the target density, is achieved. Besides good quality, this approach can easily be adapted to

many different inpainting operators, including inpainting with PDEs [34, 37] or interpolation on triangulations [32]. This flexibility and quality comes at the cost of many inpainting operations. Nonetheless, sparsification is popular and widely used due to its advantages and its simplicity.

4. *Densification Approaches.* For applications such as compression or denoising, low densities are required. In such cases it can make sense to start with an empty mask and fill it successively instead of using sparsification. Such strategies [30, 31, 36] share the benefits of simplicity, good quality, and broad applicability with sparsification. They have been successfully used for diffusion-based [30, 31] and exemplar-based [36] inpainting operators. For compression, densification also has been applied to data structures such as subdivision trees instead of individual pixels [8, 14, 16, 45]. However, all of these strategies still require a significant amount of inpaintings to obtain masks of sufficient quality.

The approaches of Categories 3 and 4 are greedy strategies that can become stuck in local minima. To address this problem, a relocation strategy, the so-called *nonlocal pixel exchange* (NLPE) [37], has been proposed as a post-processing. It is a probabilistic method that iteratively swaps point locations randomly with heuristics for candidate selection based on the inpainting error. While it can yield significant additional improvements, it also requires more inpaintings and tends to converge slowly. Similar strategies have also been used for interpolation on triangulations [38].

Note that the approach from Category 1 is the only one to require no inpaintings or complex solvers. Unfortunately, this near instantaneous spatial optimisation yields clearly worse results in terms of quality than the methods from Categories 2–4. With our deep learning framework, we aim at achieving the best of both worlds: Fast spatial optimisation without the need for any inpaintings while producing results of a quality comparable to Categories 2–4.

### 1.2.2 Tonal optimisation

So far, we have discussed methods that focus on finding optimal positions at which the original image data is kept. However, in a data optimisation scenario, we are not confined to selecting the location, but can also alter the value of mask pixels. This tonal optimisation introduces errors at mask pixels if those lead to a more accurate reconstruction in larger missing areas. Also for tonal optimisation, one can distinguish several categories:

1. *Least Squares Approaches.* For spatially fixed mask pixels, tonal optimisation leads to a least squares problem. The resulting linear system of equations is given

by the normal equations. It has as many unknowns as mask pixels. The system matrix is a quadratic, dense matrix that is symmetric and positive definite [37]. To solve it numerically, various algorithms can be applied. Direct methods include Cholesky, LU, and QR factorisations, while conjugate gradients and the LSQR algorithm constitute suitable iterative approaches [46]. Other iterative methods that have been used for tonal optimisation are the L-BFGS algorithm [29] and a gradient descent with cyclically varying step sizes [34]. All of these approaches suffer from the fact that they require to store the full matrix, which can become prohibitive for masks with too many pixels. A potential remedy of this memory restriction consists of subsequently computing a so-called inpainting echo in a mask pixel [37]. It describes the influence of the mask pixel on the final inpainting result and can be used to adjust the grey or colour value accordingly. Doing this in random order for all mask pixels can be interpreted as a randomised Gauss–Seidel or SOR iteration step. If one does not store all inpainting echoes but computes them again in each iteration step, one achieves low memory requirements at the expense of a long runtime. Discrete Green’s functions offer another way to decompose the inpainting problem into pixel-wise contributions [47]. From this dictionary, the inpainting result can be assembled with simple linear superposition. Hoffmann [48] have used this property to derive an alternative least squares formulation for tonal optimisation which can be solved efficiently with a Cholesky solver. While its solution is equivalent to the direct least squares approach, it benefits from speed-ups for low amounts of mask pixels which are represented by only a few entries from the Green’s function dictionary. A recent alternative goes back to Chizhov and Weickert [30]. It uses nested conjugate gradient approaches within a finite element framework and it is both efficient w.r.t. memory and runtime.

2. *Nonsmooth Optimisation Methods.* Hoeltgen and Weickert [35] have shown that thresholded non-binary spatial mask optimisation [28, 29, 34, 39, 40] is equivalent to a combined selection of binary masks and a tonal optimisation. Thus, the previously discussed nonsmooth strategies also indirectly perform tonal optimisation. However, this is inherently coupled to a spatial optimisation with the advantages and drawbacks described in the previous section.
3. *Localisation Approaches.* Since the influence of a single mask pixel mainly affects its local neighbourhood, tonal optimisation can be sped up by localisation. Strategies exist for localised operators such as Shepard interpolation with truncated Gaussians [41], 1-D linear interpolation [49], or smoothed particle hydrodynamics [31].

Other approaches limit the influence artificially by subdivision trees [15] or segmentation [20, 24].

4. *Quantisation-based Strategies.* All compression codecs rely on quantisation, the coarse discretisation of the colour domain. It can be beneficial to directly take quantisation into account during tonal optimisation instead of applying it in postprocessing. Thereby, one replaces the continuous optimisation problem by a discrete one. To this end, Schmaltz et al. [16] proposed a simple strategy that visits pixels in random order and changes their values if increasing or decreasing the quantisation level yields a better results. Peter et al. [14] instead augment the Gauss-Seidel strategy with echoes [37] with a projection to the quantised grey levels. For interpolation on triangulations, Marwood et al. [38] use a stochastic approach that randomly assigns different quantisation levels in combination with spatial optimisation.

In addition to tonal optimisation itself, there are also related strategies. Galić et al. [8] proposed an early predecessor that modified tonal values to avoid singularities in PDE-based inpainting. To avoid visually unpleasant singularities at mask pixels, Schmaltz et al. [16] use interpolation swapping: After the initial inpainting, they remove disks around the known data and use the more reliable reconstruction for a second inpainting.

The tonal Category 1 is restricted to linear diffusion operators, including homogeneous diffusion. Category 2 marks the indirect tonal optimisation performed by nonsmooth spatial methods and Categories 3 and 4 are mainly relevant for practical applications in compression. We aim at providing a neural network alternative to Category 1 methods for homogeneous diffusion inpainting. As for spatial inpainting, our goal is to propose a deep optimisation approach that offers high speed at good quality.

### 1.2.3 Relations to deep learning approaches

To our best knowledge, deep learning approaches for sparse data optimisation are still very rare and so far, only spatial optimisation has been covered at all. Dai et al. [26] have proposed a deep learning method for adaptive sampling that trains an inpainting and an optimisation network separately. Joint training for spatial optimisation and inpainting with Wasserstein GANs was introduced by Peter [43]. Both approaches differ significantly from the current one, since they aim at learning both a spatial optimisation CNN and the inpainting operator. In contrast, we optimise known data for model-based diffusion inpainting with a surrogate solver for homogeneous diffusion inpainting. Moreover, our deep data selection is the first to consider both spatial and tonal optimisation.

In addition, a plethora of deep inpainting methods exist (e.g. [50–56]). A full review is beyond the scope of this paper, because these approaches do not consider any form of data optimisation. Since the selection of known data is decisive for the quality of inpainting-based compression, the current lack of research in this direction is the primary reason why deep inpainting has not played a role in this area, yet.

## 1.3 Organisation of the paper

After a brief review of diffusion-inpainting and model-based optimisation in Sect. 2, we propose our deep mask learning approach in Sect. 3. Section 4 provides an ablation study and a comparison to model-based data optimisation. We conclude our paper with a summary and outlook on future work in Sect. 5.

## 2 Diffusion-based inpainting and data optimisation

Consider a grey value image  $f : \Omega \rightarrow \mathbb{R}$  that is only known on the inpainting mask, a subset  $K \subset \Omega$  of the rectangular image domain  $\Omega \subset \mathbb{R}^2$ . Diffusion-based inpainting [5, 58] reconstructs the missing areas  $\Omega \setminus K$  by propagating the information of the fixed known pixels from  $K$  over the diffusion time  $t$ . The inpainted image is the steady state  $t \rightarrow \infty$  of this evolution. Figure 1 illustrates such a propagation over time. For our inpainting purposes, we are only interested in the steady state and not the intermediate steps of the evolution.

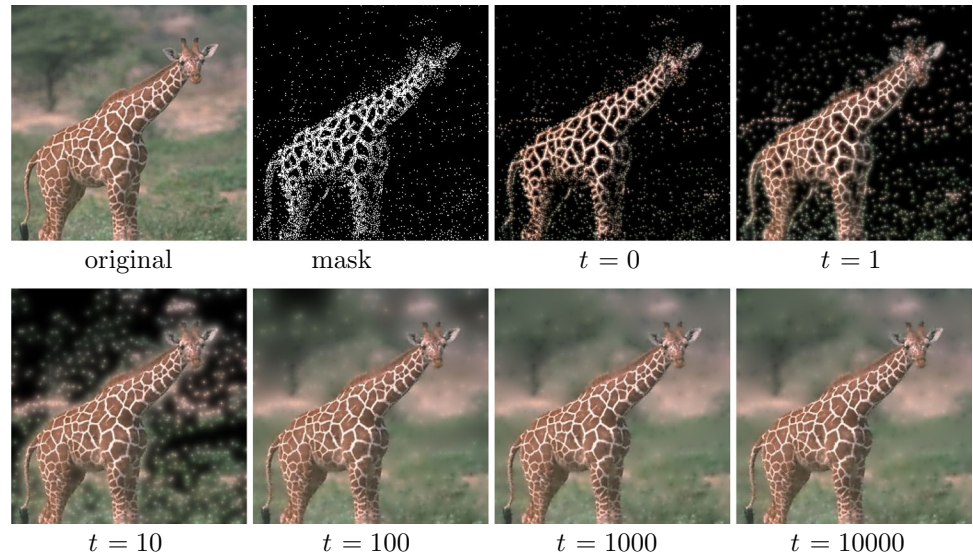
There are sophisticated anisotropic diffusion approaches [8, 15, 16, 58, 59] that adapt the amount of propagation in different directions to the image structure and can achieve results of very good quality even if the dataset  $K$  is not highly optimised. However, in the following, we consider simple homogeneous diffusion [21] for inpainting. It is parameter-free and can achieve surprisingly high quality for a well-optimised dataset. In this case, the inpainted image  $u$  fulfils the inpainting equation

$$(1 - c)\Delta u - c(u - f) = 0, \quad (1)$$

which arises as the steady state if one inpaints with the homogeneous diffusion equation  $\partial_t u = \Delta u$ . Here,  $\Delta u = \partial_{xx}u + \partial_{yy}u$  denotes the Laplacian and  $c$  is a binary confidence function with  $c(\mathbf{x}) = 1$  for known data in  $K$  and  $c(\mathbf{x}) = 0$  otherwise. At the image boundaries  $\partial\Omega$  we impose reflecting boundary conditions. Note that it is also possible to use non-binary confidence values [35], which we will do in Sect. 3.2.1. Since homogeneous diffusion is a linear operator, colour inpainting is implemented by channel-wise processing.



**Fig. 1** Image evolution of homogeneous diffusion inpainting. This figure shows a reconstruction of image 130,014 from the BSDS500 database [57] cropped to size  $256 \times 256$ . White mask pixels indicate a total of 10% known data. At time  $t = 0$ , we assign the original pixel values to known areas and initialise the unknown regions with zero (black). For  $t \rightarrow \infty$ , diffusion propagates the known values and yields the inpainted image as the steady state



In practice, we implement this method on a discrete input image  $f \in \mathbb{R}^{n_x, n_y}$  with resolution  $n_x \times n_y$ . Discretising Eq. 1 with finite differences leads to a linear system of equations. Then, reconstructing the image  $u \in \mathbb{R}^{n_x, n_y}$  is achieved with a suitable numerical solver.

The discrete problem of mask optimisation for homogeneous diffusion inpainting consists in finding the binary mask  $c \in \{0, 1\}^{n_x, n_y}$  with a user-specified target density  $d$  such that  $\|c\|_1 / (n_x n_y) = d$  where  $\|\cdot\|_1$  denotes the 1-norm. This density can be seen as a budget that specifies the percentage of image pixels that should be contained in the final mask.

For comparisons, we consider the *analytic approach* of Belhachmi et al. [27]. It is based on results from the theory of shape optimisation that demonstrate that mask pixels should be placed at locations of large Laplace magnitude. In the discrete setting, they use a Floyd–Steinberg dithering [60] of the Laplace magnitude. This leads to an imperfect, but very fast approximation of the theoretical optimum. This algorithm is a representative for simple approaches that do not require any inpaintings to determine the optimised mask.

As a prototype for better performing mask optimisation algorithms, we consider the widely used *probabilistic sparsification* of Mainberger et al. [37]. It yields better results than the analytic approach by taking the discrete nature into account directly and greedily removing pixels that are not important for the reconstruction. It starts with a full inpainting mask. In each iteration, it removes a fraction  $p$  of candidate pixels from the mask. After an inpainting with the new mask, it analyses the local inpainting error: Candidate pixels which have a high local inpainting error are hard to reconstruct and should thus not be removed. Therefore, the algorithm adds back the fraction  $q$  of candidates with the largest errors. The iterations are repeated until the target density  $d$  is reached.

Further improvements can be achieved with the *nonlocal pixel exchange* [37]. It is designed to escape from potential local minima by moving a set of  $p$  candidate locations from the inpainting mask to locations in the unknown image areas. If this positional exchange improves the overall inpainting, it is maintained, otherwise it is reverted. While this guarantees that mask quality cannot deteriorate, each step requires an inpainting and therefore, convergence tends to be slow.

In Fig. 2, a comparison of the three aforementioned spatial optimisation techniques with a uniformly random mask highlights their significant impact. Carefully optimised known data are integral for good inpainting results.

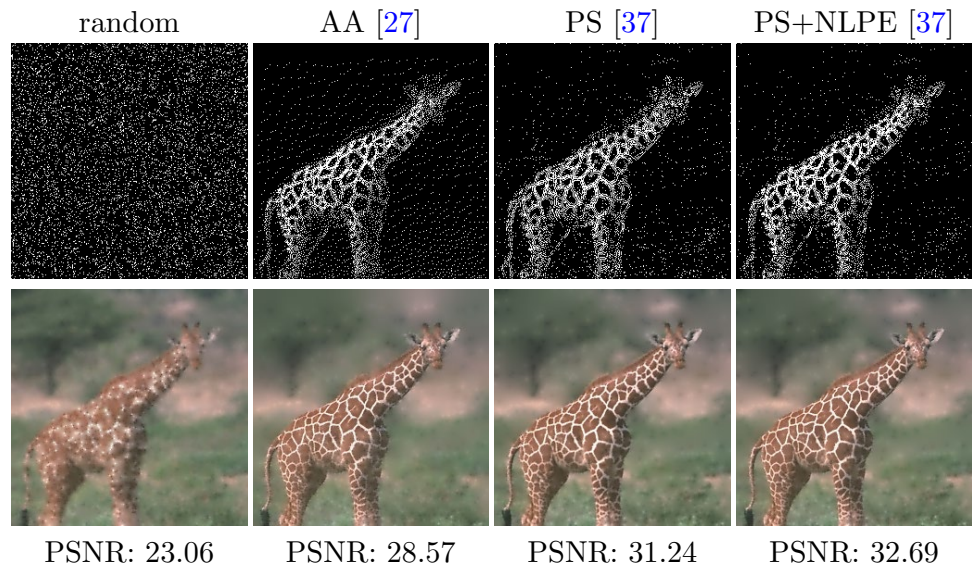
Since we consider homogeneous diffusion and do not require quantisation, we use a least squares approach for tonal optimisation. Due to the similar quality of the tonal methods from Sect. 1.2, we choose the Green’s formulation by Hoffmann et al. [48] equipped with a Cholesky solver. It offers good quality at fairly low computational cost, in particular for very sparse masks.

In the following sections we introduce a deep learning approach that does not require inpaintings during spatial or tonal optimisation and approximates the quality of probabilistic methods and model-based tonal optimisation.

### 3 Spatial and tonal optimisation with surrogate inpainting

In this section, we describe the three types of networks that act as the building blocks for our neural data optimisation framework. The centrepiece required for our different pipelines is the *surrogate inpainting network*. It approximates inpainting with homogeneous diffusion by minimising the residual of the inpainting equation. We only use it during

**Fig. 2** Spatial optimisation techniques. For reference, we consider a uniformly random mask with 10% known data and the corresponding reconstruction of image 130,014 from Fig. 1 with homogeneous diffusion inpainting. The analytic approach [27] (AA) already yields a significant improvement over the random mask. Probabilistic sparsification (PS) and non-local pixel exchange (NLPE) [37] refine the background as well as the fur patterns of the giraffe. A gain of more than 9dB PSNR illustrates the vital importance of spatial optimisation for homogeneous diffusion inpainting



training. Its sole purpose is to act as a fast approximate solver for the inpainting problem that is still differentiable and allows backpropagation.

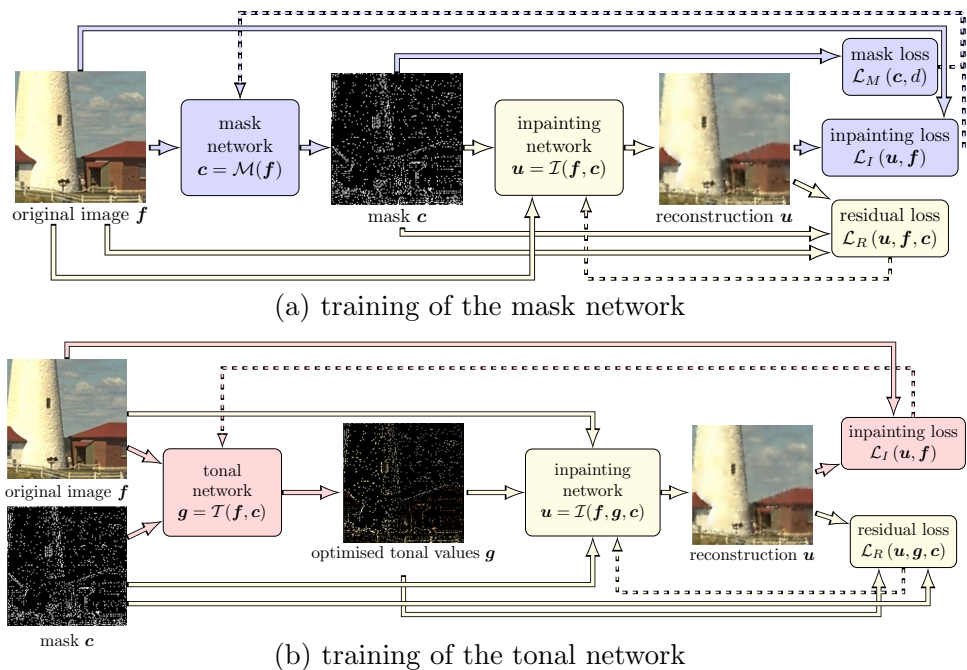
For the data optimisation, we consider a *mask network* for spatial optimisation and a *tonal network* for optimisation of the pixel values. Each of them is trained together with a separate surrogate inpainting network. Both data optimisation networks minimise the inpainting error w.r.t. the reconstruction by the respective surrogate solver.

In addition, the mask network requires a separate loss to approximate the intended mask density  $d$ . The macro architecture of our spatial approach with can be found in Fig. 3a.

For the tonal setting in Fig. 3b, we have a similar overall setup. However, here the binary masks are already part of the training dataset. In practice, we use our tonal network to generate these inputs, but also other sources such as model-based spatial optimisation approaches or even randomly generated masks could be used instead. Note that here, the optimised mask values are fed into the surrogate solver instead of the original ones.

All three types of networks use a similar U-net structure [61] that we discuss in more detail in Sect. 3.4. In the following sections on the individual networks, we only discuss deviations from this standard U-net architecture.

**Fig. 3** Structure of our deep tonal optimisation framework. The surrogate inpainting network and its associated residual loss are marked in yellow, the mask network and loss in blue, and the tonal network in red. Forward passes between the networks are indicated by solid arrows, while dashed arrows represent backpropagation



Deploying our networks for practical applications comes down to first applying the mask network to the input image. The resulting mask is then optionally fed into the tonal network together with the original. This yields the complete known data for homogeneous diffusion inpainting. The surrogate solver is never used in an evaluation scenario. Instead, we use model-based inpainting.

### 3.1 The surrogate inpainting network

To train our mask and tonal networks, we require backpropagation from inpainting results. For instance, this could be achieved by translating a classical discrete implementation of a diffusion process into a neural network [62], which results in a sequence of ResNet [63] blocks. However, this might require very deep networks to reach the steady state of the diffusion process since the number of ResNet blocks is tied to the diffusion time in such a scenario. Instead, we propose an alternative that approximates inpainting results more efficiently by also having access to the ground truth.

The *surrogate inpainting network*  $\mathcal{I}$  takes known data specified in terms of the locations in a binary or non-binary mask  $\mathbf{c}$  and pixel values  $\mathbf{g}$  as an input. Note that these known values do not necessarily need to coincide with the corresponding data in the original  $\mathbf{f}$ . In addition, it has access to the full known image  $\mathbf{f}$ . This network will only be used during training, and for evaluation, a model-based solver is responsible for the inpainting. Therefore, having access to the unknown pixels in  $\Omega \setminus K$  eases the networks task and does not compromise the validity of data optimisation in any way.

The reconstruction  $\mathbf{u} = \mathcal{I}(\mathbf{f}, \mathbf{g}, \mathbf{c})$  should solve the discrete inpainting equation

$$(\mathbf{I} - \mathbf{C})\mathbf{A}\mathbf{u} - \mathbf{C}(\mathbf{u} - \mathbf{g}) = \mathbf{0}, \tag{2}$$

which is a discretised version of Eq. (1). The finite difference discretisation of the Laplacian is represented by the matrix  $\mathbf{A} \in \mathbb{R}^{n_x n_y \times n_x n_y}$  and  $\mathbf{C} \in [0, 1]^{n_x n_y \times n_x n_y}$  is a diagonal matrix containing the mask entries.

Since the network aims at simulating a numerical solver for Eq. (2), we follow the ideas of Alt et al. [62] and define a corresponding *residual loss*

$$\mathcal{L}_R(\mathbf{u}, \mathbf{g}, \mathbf{c}) = \frac{1}{n_x n_y} \|(\mathbf{I} - \mathbf{C})\mathbf{A}\mathbf{u} - \mathbf{C}(\mathbf{u} - \mathbf{g})\|_2^2. \tag{3}$$

Here  $\|\cdot\|_2$  denotes the Euclidean norm. Note that the inpainting network is explicitly *not* trained to minimise any reconstruction loss w.r.t. the original  $\mathbf{f}$ . The residual loss only makes sure that the networks produces a good approximation of homogeneous diffusion inpainting given the mask  $\mathbf{c}$  and the pixel values  $\mathbf{g}$ . It follows similar principles as deep energy approaches [64]. This ensures that the surrogate

solver’s access to the full original image does not skew the data optimisation.

### 3.2 The mask network

Given the original image  $\mathbf{f}$ , our *mask network*  $\mathcal{M}$  outputs positional data in terms of the mask  $\mathbf{c} = \mathcal{M}(\mathbf{f})$  with a density  $d$ .

#### 3.2.1 Non-binary mask networks

In the approach from our conference paper [42], our network outputs non-binary masks with values in  $[0, 1]$ . Our goal is that the choice of  $\mathbf{c}$  should yield the best possible inpainting. Therefore, our network is equipped with an *inpainting loss* that measures the deviation of the reconstruction  $\mathbf{u}$  from the original  $\mathbf{f}$  in terms of

$$\mathcal{L}_I(\mathbf{u}, \mathbf{f}) = \frac{1}{n_x n_y} \|\mathbf{u} - \mathbf{f}\|_2^2. \tag{4}$$

While this loss establishes a connection between mask positions and reconstruction quality, it does not address the density. To this end, we apply a sigmoid activation at the last layer of our mask U-net, which limits the non-binary mask outputs to  $[0, 1]$ . If the preliminary mask  $\hat{\mathbf{c}}$  exceeds the target density  $d$ , we rescale it according to

$$\mathbf{c} = \frac{d\hat{\mathbf{c}}}{\frac{\|\hat{\mathbf{c}}\|_1}{n_x n_y} + \epsilon}. \tag{5}$$

With  $\epsilon = 10^{-5}$  we avoid rounding issues for very low estimated mask densities and potential division by zero.

During training, our network passes on the non-binary confidence values. Values close to 1 indicate that the mask network sees this position as highly important, and a value close to 0 marks unimportant positions. For practical applications, however, we still require binary masks. These can be extracted with a simple postprocessing: Interpreting the confidence values as a probability, we perform a weighted coin flip for each confidence value.

Our experiments show that this non-binary mask optimisation creates a challenging energy landscape. During the training process, the mask network can get stuck in local minima that assign equal confidence to every mask pixel. Combined with the coin flip, this can lead to a uniform random mask. As a remedy, we propose an additional *mask loss*  $\mathcal{L}_M$  that acts as a regulariser by penalising the inverse variance

$$\mathcal{L}_M(\mathbf{c}) = \alpha (\sigma_c^2 + \epsilon)^{-1} \tag{6}$$

As in Eq. (5),  $\epsilon$  avoids division by zero. The regularisation parameter  $\alpha$  balances the influence of the mask loss with the inpainting loss. Not only does this discourage flat masks



with equal confidence in every pixel, but it also encourages confidence values close to 0 and 1. This yields the additional benefit of a closer approximation of binary masks during training.

### 3.2.2 Binary mask networks

Recently, strategies for deep data optimisation of neural network-based inpainting have been proposed that also allow direct output of binary masks [43]. This constitutes a challenge since the binarisation of real input values is a non-differentiable operation. However, end-to-end approaches that also learn the inpainting benefit from this binarisation, since the training of the inpainting network tends to be biased by a non-binary mask input. This leads to worse results during deployment of the inpainting network.

For our own strategy, we investigate two different alternatives for direct binarisation and evaluate their performance in Sect. 4.2.

**Strategy 1: Quantisation.** First, we directly adopt the strategy of Peter [43]: We interpret binarisation of  $x \in \mathbb{R}$  by hard rounding  $x \mapsto \lfloor c + 0.5 \rfloor$  as very coarse quantisation. Theis et al. [65] have shown that simply approximating the derivative by 1 yields very good results among more sophisticated alternatives.

For this strategy, the variance-based regularisation from Eq. (6) is not necessary. However, the enforcement of the target density via rescaling from the non-binary approach also does not work in this case. Therefore, we define the mask loss directly as the deviation from the target density  $d$  according to

$$\mathcal{L}_M(\mathbf{c}) = \left| \frac{\|\mathbf{c}\|_1}{n_x n_y} - d \right|. \tag{7}$$

Since the mask contains only binary values, the 1-norm  $\|\cdot\|_1$  yields the number of mask points and thus the mask loss measures the deviation from the target density  $d$ . While the non-binary strategy does not require a density loss, we found in our experiments that it can have a stabilising effect on training if added to the regulariser loss from Eq. (6).

**Strategy 2: Coin Flip.** Instead of quantisation, we can also modify our non-binary approach to output binary masks. We keep the regularisation mask loss and rescaling from Sect. 3.2.1, yielding a non-binary confidence mask. However, during training, we directly add the coin flip binarisation. This can be seen as an alternative quantisation approach instead of the rounding operation in Strategy 1. We apply the same synthetic gradient as in the first binary mask approach.

In Sect. 4.2 we evaluate the binary and non-binary alternatives for mask generation in an ablation study.

### 3.3 The tonal network

Finally, our tonal network takes both the original image  $f$  and a mask  $\mathbf{c}$  as an input. The mask can either originate from the mask network or an external source.

Fortunately, we do not require binarisation layers, since the input masks are already binary. Furthermore, the mask density is already fixed. Therefore, the tonal network uses the U-net described in Sect. 3.4 without further need for modifications. It feeds the optimised pixel values  $\mathbf{g} = \mathcal{T}(f, \mathbf{c})$  into the inpainting loss from Eq. (4).

The residual network is trained with the residual loss w.r.t. the optimised known data  $\mathcal{L}_R(\mathbf{u}, \mathbf{g}, \mathbf{c})$  as well. While this works well, we have found in our experiments that the training of the surrogate solver can be stabilised by also minimising the residual  $\mathcal{L}_R(\mathbf{u}, \mathbf{f}, \mathbf{c})$  w.r.t. the original known data. This provides a fixed reference point for the residual solver, since in contrast to  $\mathbf{g}$ , the known data from  $\mathbf{f}$  is not influenced by the training progress of the tonal network. This prevents the training of the residual solver from getting trapped in local minima.

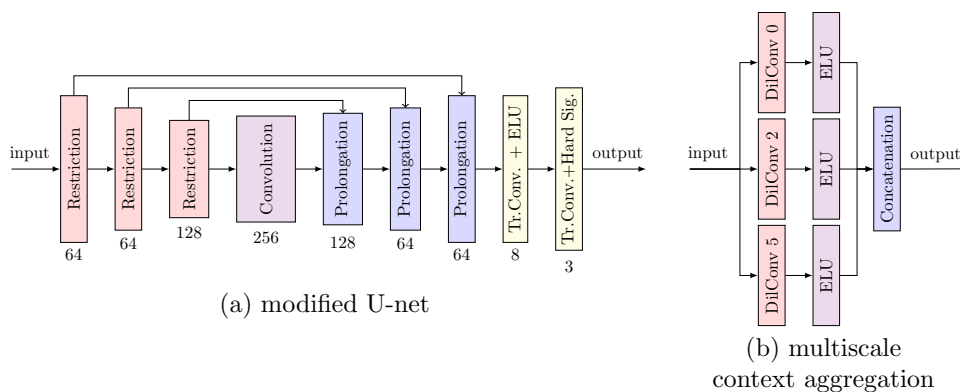
### 3.4 Network architecture

For all three networks, we use a U-net [61] architecture, since U-nets implement the core principles of multigrid solvers for PDE-based inpainting [62]. This makes them a perfect fit for the surrogate solver. U-nets and multigrid have in common that they operate on multiple scales, first restricting the image in multiple stages down to the coarsest scale and then prolongating it again to the finest scale. We follow this general structure in Fig. 4a.

However, in contrast to our conference paper [42], we also rely on modifications to the standard U-net approach that were first used for inpainting by Vařata et al. [66]. They replace traditional convolutional layers by multiple parallel dilated convolutions with dilation factors 0, 2, and 5 followed by ELU activations. As shown in Fig. 4b, the results are concatenated to a joint output. This so-called multiscale context aggregation was originally designed by Yu and Koltun [67] to increase the receptive field for segmentation. We discuss its benefits for our application in Sect. 4.2.1 with an ablation study.

For restriction, we also use context aggregation [67] with  $5 \times 5$  dilated convolutions followed by a  $2 \times 2$  max pooling. The corresponding prolongation uses the same structure, but with  $5 \times 5$  transposed convolutions and  $2 \times 2$  upsampling. Two context aggregation blocks without any upsampling or max pooling perform postprocessing on the coarsest scale. The final hard sigmoid activation limits the results to the original image range  $[0, 1]$ . Only in the case of our binary mask networks, this is followed by a quantisation or coin flip binarisation layer. As commonly the case in multiscale





**Fig. 4** Modified U-Net Architecture. **a** The original inputs are subsampled by max pooling three times, pass through the bottleneck and are then prolonged by upsampling again to the finest scale. Exchange of information between scales is implemented by skip connections. Two post processing blocks with transposed convolutions

architectures, the number of channels increases for coarser scales. It ranges from 64 to 256 (see Fig. 4a for details), which is half of the channel bandwidth used by Vařata et al. [66]. In Sect. 4.2, we have verified that such smaller networks suffice for our task.

## 4 Experimental evaluation

After an overview of the technical details of our evaluation in Sect. 4.1, we justify our design decisions for the networks with an ablation study in Sect. 4.2. We compare with model-based approaches for spatial optimisation in Sect. 4.3 and with tonal optimisation methods in Sect. 4.4. In both cases, we assess reconstruction quality and speed.

### 4.1 Experimental setup

Unless stated otherwise, all networks rely on the modified U-net architecture from Sect. 3.4 with  $\approx 2.9$  million parameters per network.

All of our networks have been trained on an Intel Xeon E5-2689 v4 CPU (2 cores), together with an Nvidia Pascal P100 16 GB GPU. For training, we use a subset of 100,000 images randomly sampled from ImageNet [68] by Dai et al. [26] and the corresponding validation dataset containing 1, 000 images. We use centre crops to reduce the size of the images, thus speeding up the training process. For model selection we crop to  $64 \times 64$ , while the remainder of the experiments are performed on size  $128 \times 128$ . All networks were with the Adam optimiser [69] and a learning rate of  $5 \cdot 10^{-5}$ . We used 50 epochs for the spatial experiments, and 100 for tonal experiments. For evaluation, we used an AMD Ryzen 7 5800X CPU equipped with an Nvidia RTX

conclude the pipeline. The numbers below each context aggregation block indicate the number of channels. **b** In a context aggregation block, three parallel dilated convolution increase the receptive field of the filter. All results are concatenated together

3090 24GB GPU. We performed model selection based on the lowest achieved inpainting error on the validation set. Our test set is based on all 500 images of the BSDS500 database [57]. These were centre cropped to size  $128 \times 128$  in order to fit the size of the training data. The cropping also speeds up the model-based competitors and thus allows us to compare with them on a larger variety of images. We measure qualitative results with the peak signal-to-noise ratio (PSNR).

We compare with three spatial optimisation methods. The analytic approach by Belhachmi et al. [27] (AA) acts as a representative of very fast spatial optimisation. It is implemented with Floyd–Steinberg dithering [60] of the Laplace magnitude. Probabilistic sparsification (PS) in combination with a non-local pixel exchange (NLPE) provides qualitative benchmarks. These methods have been implemented with a conjugate gradient solver, ensuring convergence up to a relative residual of  $10^{-6}$  for the diffusion inpainting. NLPE is run for 5 so-called cycles, each consisting of  $\|c\|_1$  iterations.

### 4.2 Ablation study

In the following, we first evaluate different architectures and design principles, to select the best among those for the comparison with model-based approaches.

#### 4.2.1 Network architecture

Compared to the standard U-net architecture in our conference publication [42], the modified U-net from Sect. 3.4 benefits from the context aggregation and more sophisticated postprocessing layers after upsampling to the finest scale. In [42], we used sequential  $3 \times 3$  convolutions on each scale. Therefore, propagation of information

over larger distances works mainly via downsampling to coarse scales and upsampling. On each individual scale, the receptive field of the simple convolutions is relatively small. In contrast, the context aggregation allows our network to perceive larger regions of the image on each individual scale. Our evaluation in Table 1(a) contrasts these modifications with the standard U-net using a similar total amount of weights. The modifications yield up to 2.3 dB improvement w.r.t. PSNR, especially on challenging very sparse masks.

We also evaluated other modifications to the U-net structure such as gated convolutions, but the context aggregation yielded the best combination of good qualitative performance and stability during training.

In Table 1(b), we compare the full size U-net proposed by Vařata et al. [66] with our leaner version from Sect. 3.4 on  $128 \times 128$  color images. The large U-net uses twice the amount of channels in relation to Fig. 4a in all but the last two postprocessing layers. This results in  $\approx 11.5$  million parameters, compared to our significantly lower  $\approx 2.9$  million. The larger network does not yield a qualitative advantage over a wide range of densities. The PSNR results of the large network only deviate marginally from those of the small masknet in Table 1(b). However, it increases training times from 43 min to 93 min per epoch. Therefore, we use our lean nets instead.

**Table 1** Ablation experiments

Density	PSNR (dB)		
	1%	5%	10%
<i>(a) Binary vs. non-binary</i>			
Non-binary [42]	19.40	25.45	28.34
Our non-binary	<b>21.72</b>	<b>25.92</b>	<b>29.06</b>
Coinflip	18.61	24.99	22.58
Binary	20.08	24.00	26.04
<i>(b) Small vs. large masknet</i>			
Small	21.10	25.48	28.58
Large	21.04	25.64	28.63

Bold represents the highest PSNR values and thus highlight the best model

All experiments have been conducted on  $64 \times 64$  grey value centre crops from BSDS500 [57] with mask densities 1%, 5%, and 10%. **a** The non-binary masks outperform both binary options qualitatively over the full range of mask densities. In addition, our modified network architecture and training methodology outperforms the earlier non-binary mask network [42]. The coinflip variant showed instabilities during training for high densities and did thus not yield satisfying results for 10%. **b** Reducing the number of channels in the modified U-net by a factor 2 does not deteriorate the quality

### 4.2.2 Non-binary vs. binary masks

In Sect. 3.2, we have proposed three possible output options for our mask networks: non-binary masks, binary masks based on quantisation, and binary masks produced by a coin flip. For full deep learning based approaches [43], the binarisation during training is a key component of their architecture.

Surprisingly, our ablation study in Table 1(a) paints a different picture: The non-binary mask network clearly outperforms both binary options. This results from a key difference in our method compared to full deep learning approaches. Using a non-binary mask while simultaneously training an inpainting network introduces a bias. This deteriorates inpainting quality during testing [26]. However, our surrogate solver is only deployed during training and is not coupled directly to an inpainting loss. It merely approximates diffusion-based inpainting. During testing, we use a model-based implementation of homogeneous diffusion inpainting.

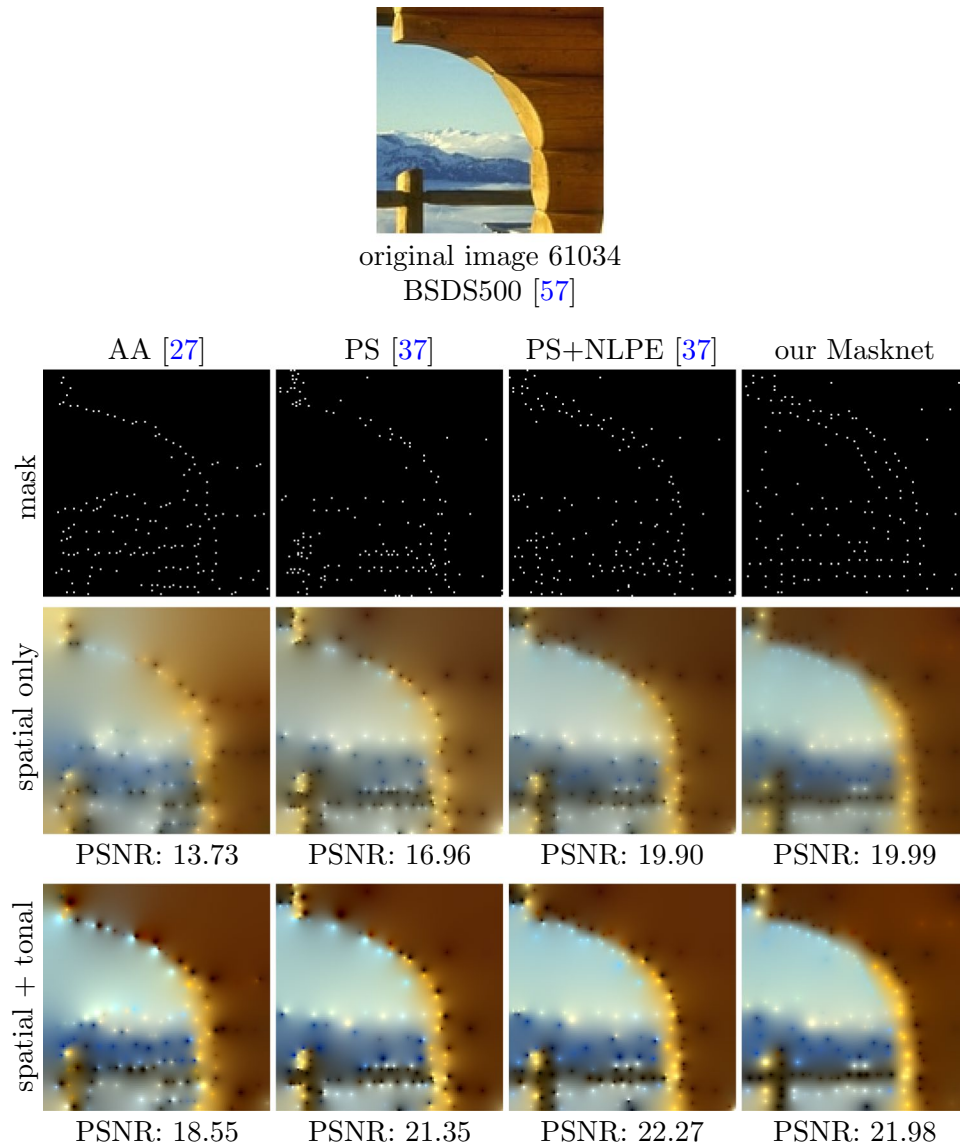
Therefore, we benefit from a non-binary mask network that does not rely on synthetic gradients for binarisation layers. Consequentially, we use the non-binary variant for our comparisons with model-based data optimisation.

### 4.3 Spatial optimisation

In our conference paper [42], we have shown that our approach yields similar results as probabilistic sparsification [37] on a small dataset of five greyscale images. Here we extend the evaluation of our improved networks to the significantly larger greyscale BSDS500 database in Fig. 7a and double the range of evaluated mask densities to 20%. Our mask network not only consistently outperforms both the analytic approach [27] (AA) and probabilistic sparsification [37] (PS), but very closely approximates the quality of PS + NLPE.

The same ranking also applies in the case of the full colour version of BSDS500 in Fig. 7b. Thus, our mask network rivals the best model-based approach in the comparison. Visually, it yields similar results as the probabilistic methods in Figs. 5 and 6. Especially for low densities, there is a large quality gap between the analytical approach and all other competitors.

Even though our mask net offers a similar quality as PS + NLPE, it requires significantly less computational time since it does not rely on any inpaintings during inference. On the CPU, it accelerates mask computation by up to a factor 3500 and even up to a factor 140.000 on the GPU in Fig. 9a. Only the analytic approach is faster with  $\approx 0.2$  ms. However, there the speed comes at the cost of a



**Fig. 5** Visual Comparison for 1% Mask Density. Visually, probabilistic sparsification (PS) in combination with non-local pixel exchange (NLPE) [37], and our network-based approach significantly outperform the analytic approach (AA) [27]. For such sparse masks, the visual impact of tonal optimisation is apparent

significantly diminished quality. Without compromising on quality, our mask net is also real-time capable with 1.4 ms on GPU and 55 ms on CPU.

Thus, our mask network reaches our goal of providing an easy to use, parameter-free spatial optimisation which approximates the quality of stochastic methods at a computational cost close to the instantaneous analytic approach.

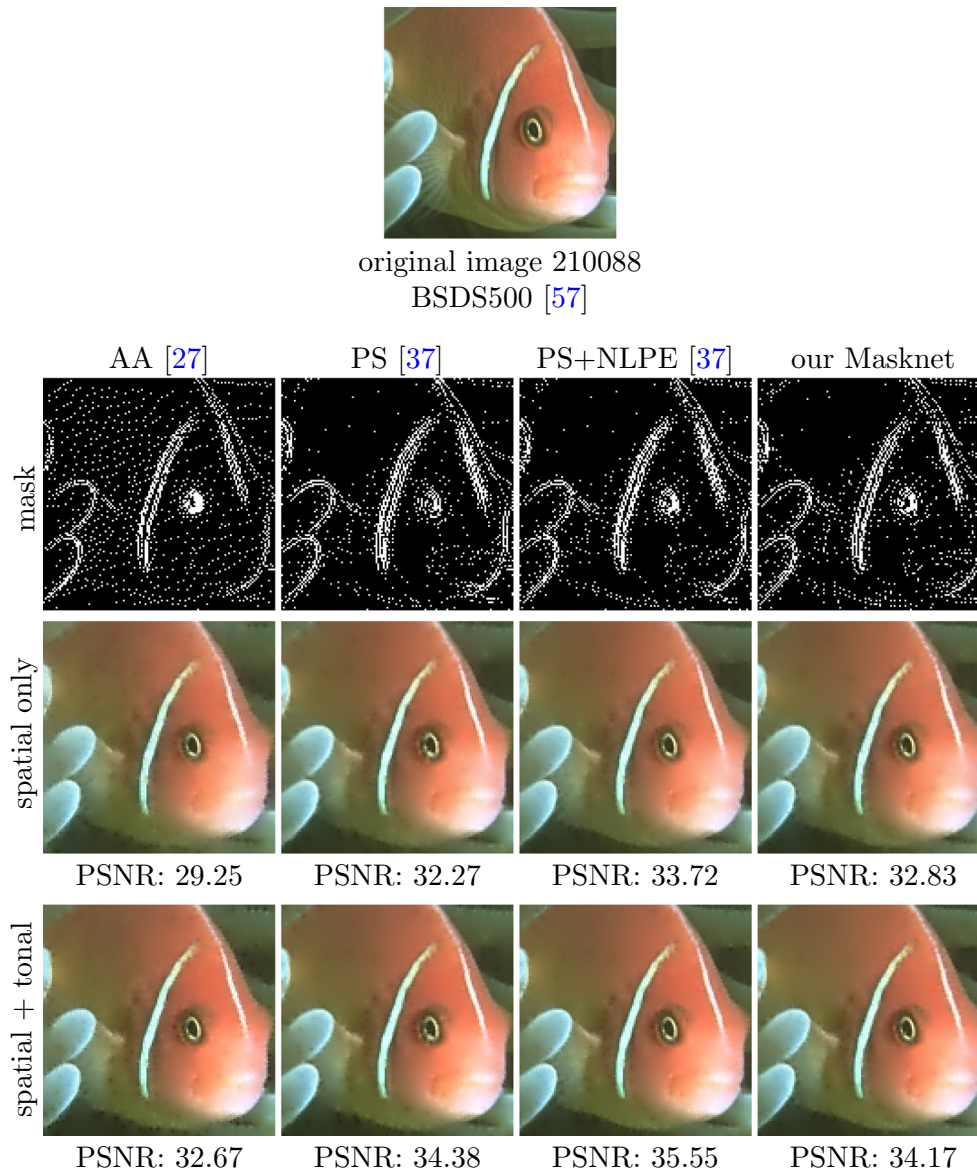
#### 4.4 Tonal optimisation

In Fig. 8a, we compare our tonal network with the Green’s function approach of Hoffmann [48] on the masks obtained

from our mask network. Especially for sparse known data, our deep tonal optimisation reaches a similar quality as the model-based approach. Only above 15%, the improvements over the unoptimised data from the mask net decline.

Our results in Fig. 5 show that our network approach also remains competitive to PS + NLPE when adding tonal optimisation. Here we apply the tonal network for our own deep learning method and the Green’s function optimiser for all model-based competitors.

As for spatial optimisation, our tonal network offers a viable alternative for time critical applications. Figure 9b shows that the computational cost of the Green’s function approach grows significantly with the number of mask



**Fig. 6** Visual Comparison for 10% Mask Density. Also at high density, our mask network yields results that are comparable to the probabilistic approaches PS and PS + NLPE [37]. The visual gap towards the analytic approach (AA) [27] is smaller, but still noticeable

values that need to be optimised. In contrast, the computational time of the tonal network is independent of the mask density. For densities larger than 5%, speed-ups by multiple orders of magnitude can be achieved with our mask net.

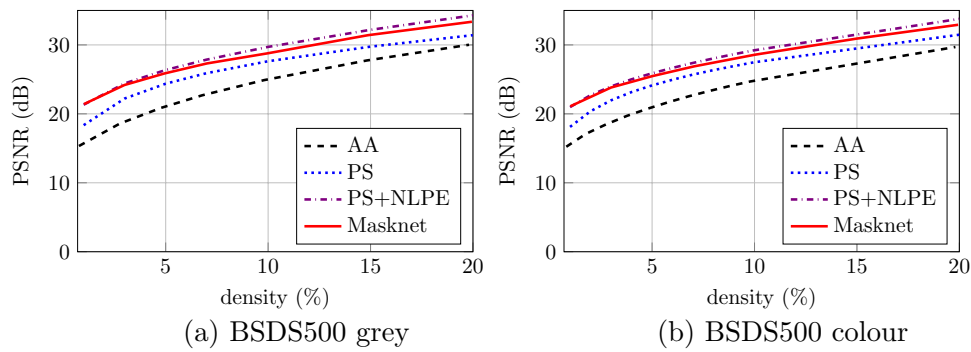
Thus, a combination of our spatial and tonal networks is a viable option for real-time applications that does not require to sacrifice quality for speed.

## 5 Conclusions

Our data optimisation approach merges classical inpainting with partial differential equations and deep learning with a surrogate solver. This allows us to select both position and values of known data for homogeneous diffusion inpainting that minimise the reconstruction error.

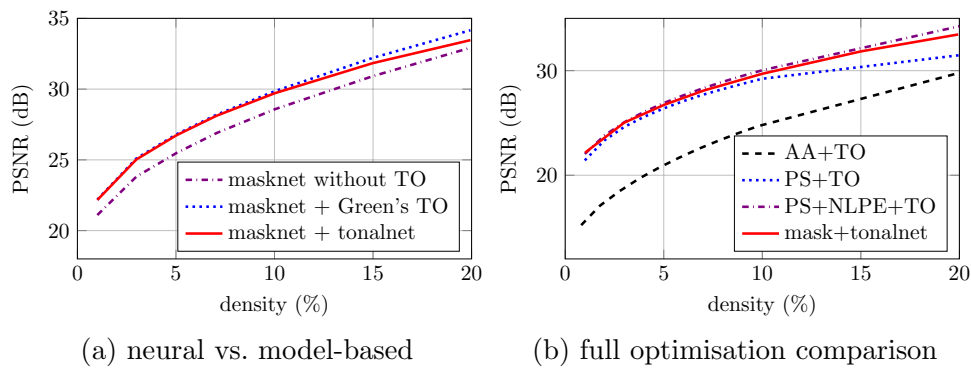
With this new strategy for sparse data optimisation we obtain real-time results in hitherto unprecedented quality.





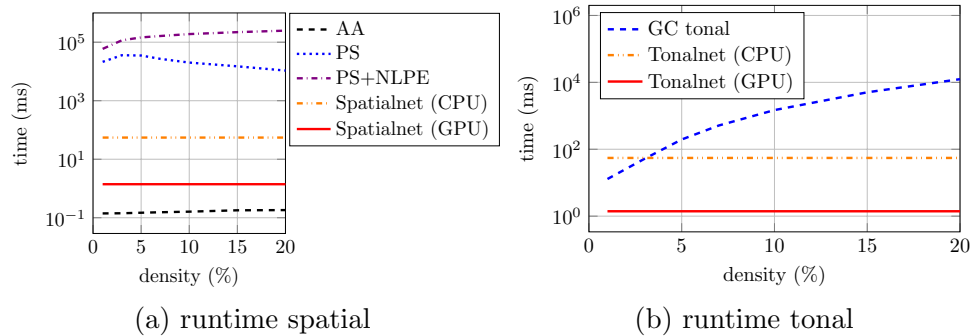
**Fig. 7** Spatial Optimisation. **a** On grey level images, our network consistently outperforms the analytic spatial optimisation (AA) [27] and probabilistic sparsification (PS) [37]. Especially for lower densities, Masknet results rival the quality of PS with non-local pixel

exchange (NLPE) [37] as postprocessing. **b** For colour images, our mask network also closely approximates the quality of PS + NLPE for the whole range from 1 to 20%



**Fig. 8** Tonal Optimisation. **a** We compare our tonal network with the Green's function approach [48] on masks from our mask network. Our tonal optimisation (TO) reaches a comparable quality up to 15% known data. **b** In a comparison that combines the spatially optimised

mask with tonal optimisation, our full network approach yields competitive results to PS + NLPE combined with tonal optimisation. All model-based approaches use the Green's function approach [48] for tonal optimisation



**Fig. 9** Runtime Comparison with Logarithmic Time Axis. **a** The analytic approach (AA) [27] is the fastest method, followed closely by our Masknet on the GPU and CPU. These three methods all have a constant speed independently of mask density. The speed of PS and PS + NLPE [37] is density dependent. Overall, our methods are consistently faster by several orders of magnitude compared to proba-

bilistic approaches. **b** The situation for tonal optimisation is comparable. The Green's function-based solver [48] becomes increasingly slower with rising mask density. Only for very low densities its speed is comparable to our tonal network on the CPU. The networks have constant runtime independent of density and are faster by 1–5 orders of magnitude

They yield reconstructions that rival the results of probabilistic sparsification with postprocessing by non-local pixel exchange and tonal optimisation. Simultaneously, they are reaching the near instantaneous speed of the qualitatively inferior analytic approach. This improvement of computational time by multiple orders of magnitude at comparable quality demonstrates the high potential of a fusion between model- and learning-based principles. We see this as a milestone on our way to bring the best of both worlds together in the area of inpainting and data optimisation.

In the future, we plan to incorporate our framework into image compression codecs. Time-consuming spatial and tonal optimisation still present a bottleneck in this area. This holds true especially for practical applications with high demand for computational efficiency, such as video coding. While real-time decoding is already possible with diffusion [70–72], the data selection during encoding will benefit from our deep optimisation.

**Acknowledgements** This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 741215, ERC Advanced Grant INCOVID).

**Funding** Open Access funding enabled and organized by Projekt DEAL.

**Data Availability** The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

## Declarations

**Conflict of interests** The authors have no competing interests to declare that are relevant to the content of this article.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Masnou S, Morel J-M (1998) Level lines based disocclusion. In: Proceedings of the 1998 IEEE international conference on image processing. Chicago, IL, vol 3, pp 259–263
- Efros AA, Leung T (1999) Texture synthesis by non-parametric sampling. In: Proceedings of the seventh international conference on computer vision. Kerkyra, Greece, vol 2, pp 1033–1038
- Bertalmío M, Sapiro G, Caselles V, Ballester C (2000) Image inpainting. In: Proceedings of the SIGGRAPH 2000, New Orleans, LI, pp 417–424
- Guillemot C, Le Meur O (2014) Image inpainting: overview and recent advances. *IEEE Signal Process Mag* 31(1):127–144
- Carlsson S (1988) Sketch based coding of grey level images. *Signal Process* 15:57–83
- Acar T, Gökmen M (1994) Image coding using weak membrane model of images. In: Katsaggelos AK (ed) Visual communications and image processing'94, vol 2308. Proceedings of SPIE. SPIE Press, Bellingham, pp 1221–1230
- Desai UY, Mizuki MM, Masaki I, Horn BKP (1996) Edge and mean based image compression. Technical Report 1584 (A.I. Memo), Artificial Intelligence Lab., Massachusetts Institute of Technology, Cambridge, MA, Nov 1996
- Galić I, Weickert J, Welk M, Bruhn A, Belyaev A, Seidel H-P (2008) Image compression with anisotropic diffusion. *J Math Imaging Vis* 31(2–3):255–269
- Wu Y, Zhang H, Sun Y, Guo H (2009) Two image compression schemes based on image inpainting. In: Proceedings of the 2009 international joint conference on computational sciences and optimization, Sanya, China, pp 816–820
- Bastani V, Helfroush M, Kasiri K (2010) Image compression based on spatial redundancy removal and image inpainting. *J Zhejiang Univ Sci C (Comput Electron)* 11(2):92–100
- Zhao C, Du M (2011) Image compression based on PDEs. In: Proceedings of the 2011 international conference of computer science and network technology, Harbin, China, pp 1768–1771
- Gautier J, Le Meur O, Guillemot C (2012) Efficient depth map compression based on lossless edge coding and diffusion. In: Proceedings of the 2012 picture coding symposium, Kraków, Poland, pp 81–84
- Li Y, Sjöström M, Jennehag U, Olsson R (2012) A scalable coding approach for high quality depth image compression. In: Proceedings of the 3DTV-conference: the true vision—capture, transmission and display of 3D Video, Zurich, Switzerland
- Peter P, Hoffmann S, Nedwed F, Hoeltgen L, Weickert J (2016) Evaluating the true potential of diffusion-based inpainting in a compression context. *Signal Process: Image Commun* 46:40–53
- Peter P, Kaufhold L, Weickert J (2017) Turning diffusion-based image colorization into efficient color compression. *IEEE Trans Image Process* 26(2):860–869
- Schmaltz C, Peter P, Mainberger M, Ebel F, Weickert J, Bruhn A (2014) Understanding, optimising, and extending data compression with anisotropic diffusion. *Int J Comput Vis* 108(3):222–240
- Breuß M, Hoeltgen L, Radow G (2021) Towards PDE-based video compression with optimal masks prolonged by optic flow. *J Math Imaging Vis* 63(2):144–156
- Jumakulyyev I, Schultz T (2021) Lossless PDE-based compression of 3D medical images. In: Elmoataz A, Fadili J, Quéau Y, Rabin J, Simon L (eds) Scale space and variational methods in computer vision. Lecture notes in computer science, vol 12679. Springer, Cham, pp 450–462
- Galić I, Weickert J, Welk M, Bruhn A, Belyaev A, Seidel H-P (2005) Towards PDE-based image compression. In: Paragios N, Faugeras O, Chan T, Schnörr C (eds) Variational, geometric and level-set methods in computer vision. Lecture notes in computer science, vol 3752. Springer, Berlin, pp 37–48
- Hoffmann S, Mainberger M, Weickert J, Pohl M (2013) Compression of depth maps with segment-based homogeneous diffusion. In: Kuijper A, Bredies K, Pock T, Bischof H (eds) Scale space and variational methods in computer vision. Lecture notes in computer science, vol 7893. Springer, Berlin, pp 319–330

21. Iijima T (1962) Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bull Electrotech Labor* 26:368–388 (In Japanese)
22. Sullivan GJ, Ohm JR, Han WJ, Wiegand T (2012) Overview of the high efficiency video coding (HEVC) standard. *IEEE Trans Circuits Syst Video Technol* 22(12):1649–1668
23. Jost F, Peter P, Weickert J (2020) Compressing flow fields with edge-aware homogeneous diffusion inpainting. In: *Proceedings of the 2020 international conference on acoustics, speech, and signal processing*, Barcelona, Spain, pp 2198–2202
24. Jost F, Peter P, Weickert J (2021) Compressing piecewise smooth images with the Mumford–Shah cartoon model. In: *Proceedings of the 28th European signal processing conference*, Amsterdam, Netherlands, pp 511–515
25. Adam RD, Peter P, Weickert J (2017) Denoising by inpainting. In: Lauze F, Dong Y, Dahl AB (eds) *Scale space and variational methods in computer vision. Lecture notes in computer science*, vol 10302. Springer, Cham, pp 121–132
26. Dai Q, Chopp H, Pouyet E, Cossairt O, Walton M, Katsaggelos AK (2019) Adaptive image sampling using deep learning and its application on X-ray fluorescence image reconstruction. *IEEE Trans Multimedia* 22(10):2564–2578
27. Belhachmi Z, Bucur D, Burgeth B, Weickert J (2009) How to choose interpolation data in images. *SIAM J Appl Math* 70(1):333–352
28. Bonettini S, Loris I, Porta F, Prato M, Rebegoldi S (2017) On the convergence of a linesearch based proximal-gradient method for nonconvex optimization. *Inverse Prob* 33(5):055005
29. Chen Y, Ranftl R, Pock T (2014) A bi-level view of inpainting-based image compression. In: Kúkelová Z, Heller J (eds) *Proceedings of the 19th computer vision winter workshop*, Křtiny, Czech Republic
30. Chizhov V, Weickert J (2021) Efficient data optimisation for harmonic inpainting with finite elements. In: Tsapatsoulis N, Panayides A, Theocharides T, Lanitis A, Pattichis CS, Vento M (eds) *Computer analysis of images and patterns. Part 2. Lecture notes in computer science*, vol 13053. Springer, Cham, pp 432–441
31. Daropoulos V, Augustin M, Weickert J (2021) Sparse inpainting with smoothed particle hydrodynamics. *SIAM J Appl Math* 14(4):1669–1704
32. Demaret L, Dyn N, Iske A (2006) Image compression by linear splines over adaptive triangulations. *Signal Process* 86(7):1604–1616
33. Hoeltgen L, Setzer S, Weickert J (2013) 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 minimisation methods in computer vision and pattern recognition. Lecture notes in computer science*, vol 8081. Springer, Berlin, pp 151–164
34. Hoeltgen L, Mainberger M, Hoffmann S, Weickert J, Tang CH, Setzer S, Johannsen D, Neumann F, Doerr B (2017) Optimising spatial and tonal data for PDE-based inpainting. In: Bergounioux M, Peyré G, Schnörr C, Caillaud J-P, Haberkorn T (eds) *Variational methods in imaging and geometric control. Radon series on computational and applied mathematics*, vol 18. De Gruyter, Berlin, pp 35–83
35. Hoeltgen L, Weickert J (2015) Why does non-binary mask optimisation work for diffusion-based image compression? In: Tai X-C, Bae E, Chan TF, Leung SY, Lysaker M (eds) *Energy minimisation methods in computer vision and pattern recognition. Lecture notes in computer science*, vol 8932. Springer, Berlin, pp 85–98
36. Karos L, Bheed P, Peter P, Weickert J (2018) Optimising data for exemplar-based inpainting. In: Blanc-Talon J, Helbert D, Philips W, Popescu D, Scheunders P (eds) *Advanced concepts for intelligent vision systems. Lecture notes in computer science*, vol 11182. Springer, Cham, pp 547–558
37. Mainberger M, Hoffmann S, Weickert J, Tang CH, Johannsen D, Neumann F, Doerr B (2012) Optimising spatial and tonal data for homogeneous diffusion inpainting. In: Bruckstein AM, ter Haar Romeny B, Bronstein AM, Bronstein MM (eds) *Scale space and variational methods in computer vision. Lecture notes in computer science*, vol 6667. Springer, Berlin, pp 26–37
38. Marwood D, Massimino P, Covell M, Baluja S (2018) Representing images in 200 bytes: compression via triangulation. In: *Proceedings of the 2018 IEEE international conference on image processing*, Athens, Greece, pp 405–409
39. Nahme R (2015) Inertial proximal algorithms in diffusion-based image compression. Master's thesis, Department of Mathematics, University of Göttingen, Germany
40. Ochs P, Chen Y, Brox T, Pock T (2014) iPiano: inertial proximal algorithm for nonconvex optimization. *SIAM J Imag Sci* 7(2):1388–1419
41. Peter P (2019) Fast inpainting-based compression: combining Shepard interpolation with joint inpainting and prediction. In: *Proceedings of the 26th IEEE international conference on image processing*, Taipei, Taiwan, pp 3557–3561
42. Alt T, Peter P, Weickert J (2022) Learning sparse masks for diffusion-based image inpainting. In: Pinho AJ, Georgieva P, Teixeira LF, Sánchez JA (eds) *Pattern recognition and image analysis. Lecture notes in computer science*, vol 13256. Springer, Cham, pp 528–539
43. Peter P (2022) A Wasserstein GAN for joint learning of inpainting and its spatial optimisation. [arXiv:2202.05623](https://arxiv.org/abs/2202.05623) [eess.IV]
44. Schütze T, Schwetlick H (2003) Bivariate free knot splines. *BIT Numer Math* 43(1):153–178
45. Distasi R, Nappi M, Vitulano S (1997) Image compression by B-tree triangular coding. *IEEE Trans Commun* 45(9):1095–1100
46. Björck Å (1996) *Numerical methods for least squares problems*. SIAM, Philadelphia
47. Hoffmann S, Plonka G, Weickert J (2015) Discrete Green's functions for harmonic and biharmonic inpainting with sparse atoms. In: Tai X-C, Bae E, Chan TF, Lysaker M (eds) *Energy minimization methods in computer vision and pattern recognition. Lecture notes in computer science*, vol 8932. Springer, Berlin, pp 169–182
48. Hoffmann S (2017) *Competitive image compression with linear PDEs*. PhD thesis, Department of Computer Science, Saarland University, Saarbrücken, Germany
49. Peter P, Contelly J, Weickert J (2019) Compressing audio signals with inpainting-based sparsification. In: Lellmann J, Burger M, Modersitzki J (eds) *Scale space and variational methods in computer vision. Lecture notes in computer science*, vol 11603. Springer, Cham, pp 92–103
50. Liu H, Jiang B, Xiao Y, Yang C (2017) Coherent semantic attention for image inpainting. In: *Proceedings of the 2019 IEEE/CVF international conference on computer vision*, Seoul, Korea, pp 4170–4179
51. Pathak D, Krähenbühl P, Donahue J, Darrell T, Efros AA (2016) Context encoders: feature learning by inpainting. In: *Proceedings of the 2016 IEEE conference on computer vision and pattern recognition*, Las Vegas, NV, pp 2536–2544
52. Xie J, Xu L, Chen E (2012) Image denoising and inpainting with deep neural networks. In: Bartlett PL, Pereira FCN, Burges CJC, Bottou L, Weinberger KQ (eds) *Proceedings of the 26th international conference on neural information processing systems. Advances in neural information processing systems*, vol 25. Lake Tahoe, NV, pp 350–358
53. Yang C, Lu X, Lin Z, Shechtman E, Wang O, Li H (2017) High-resolution image inpainting using multi-scale neural patch synthesis. In: *Proceedings of the 2017 IEEE conference on computer vision and pattern recognition*, Honolulu, HI, pp 6721–6729

54. Yu J, Lin Z, Yang J, Shen X, Lu X, Huang TS (2018) Generative image inpainting with contextual attention. In: Proceedings of the 2018 IEEE conference on computer vision and pattern recognition, Salt Lake City, UT, pp 5505–5514
55. Wang W, Zhang J, Niu L, Ling H, Yang X, Zhang L (2021) Parallel multi-resolution fusion network for image inpainting. In: Proceedings of the 2021 IEEE/CVF international conference on computer vision, pp 14559–14568
56. Wang N, Zhang Y, Zhang L (2021) Dynamic selection network for image inpainting. *IEEE Trans Image Process* 30:1784–1798
57. Arbelaez P, Maire M, Fowlkes C, Malik J (2011) Contour detection and hierarchical image segmentation. *IEEE Trans Pattern Anal Mach Intell* 33(5):898–916
58. Weickert J, Welk M (2006) Tensor field interpolation with PDEs. In: Weickert J, Hagen H (eds) Visualization and processing of tensor fields. Springer, Berlin, pp 315–325
59. Jumakulyyev I, Schultz T (2021) Fourth-order anisotropic diffusion for inpainting and image compression. In: Özarlan E, Schultz T, Zhang E, Fuster A (eds) Anisotropy across fields and scales. Mathematics and visualization. Springer, Cham, pp 99–124
60. Floyd RW, Steinberg L (1976) An adaptive algorithm for spatial grey scale. *Proc Soc Inf Disp* 17:75–77
61. Ronneberger O, Fischer P, Brox T (2015) U-net: convolutional networks for biomedical image segmentation. In: Navab N, Hornegger J, Wells W, Frangi A (eds) Medical image computing and computer-assisted intervention—MICCAI 2015. Lecture Notes in Computer Science, vol 9351. Springer, Cham, pp 234–241
62. Alt T, Schrader K, Augustin M, Peter P, Weickert J (2022) Connections between numerical algorithms for PDEs and neural networks. *J Math Imaging Vis* 65:185–208
63. He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: Proceedings of the 2016 IEEE conference on computer vision and pattern recognition, Las Vegas, NV, pp 770–778
64. Golts A, Freedman D, Elad M (2021) Deep energy: task driven training of deep neural networks. *IEEE J Sel Top Signal Process* 15(2):324–338
65. Theis L, Shi W, Cunningham A, Huszár F (2017) Lossy image compression with compressive autoencoders. In: Proceedings of the 5th international conference on learning representations, Toulouse, France
66. Vařata D, Halama T, Friedjungová M (2021) Image inpainting using Wasserstein generative adversarial imputation network. In: Farkař I, Masulli P, Otte S, Wermter S (eds) Artificial neural networks and machine learning—ICANN 2021. Lecture notes in computer science, vol 12892. Springer, Cham, pp 575–586
67. Yu F, Koltun V (2016) Multi-scale context aggregation by dilated convolutions. In: Proceedings of the 4th international conference on learning representations, San Juan, Puerto Rico
68. Deng J, Dong W, Socher R, Li L-J, Li K, Fei-Fei L (2009) Imagenet: a large-scale hierarchical image database. In: Proceedings of the 2009 IEEE computer society conference on computer vision and pattern recognition, Miami, FL, pp 248–255
69. Kingma DP, Ba J (2015) Adam: a method for stochastic optimization. In: Proceedings of the 3rd international conference on learning representations, San Diego, CA
70. Köstler H, Stürmer M, Freundl C, Rude U (2007) PDE based video compression in real time. Technical Report 07-11, Lehrstuhl für Informatik 10, Univ. Erlangen–Nürnberg, Germany
71. Peter P, Schmaltz C, Mach N, Mainberger M, Weickert J (2015) Beyond pure quality: progressive mode, region of interest coding and real time video decoding in PDE-based image compression. *J Vis Commun Image Represent* 31:256–265
72. Andris S, Peter P, Mohideen RMK, Weickert J, Hoffmann S (2021) Inpainting-based video compression in FullHD. In: Elmoataz A, Fadili J, Quéau Y, Rabin J, Simon L (eds) Scale space and variational methods in computer vision. Lecture notes in computer science, vol 12679. Springer, Cham, pp 425–436

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