In J. Blanc-Talon, D. Helbert, W. Philips, D. Popescu, P. Scheunders (Eds.): Advanced Concepts for Intelligent Vision Systems. Lecture Notes in Computer Science, Vol. 11182, pp. 547–558, Springer, Cham, 2018. The final publication is available at link.springer.com.

Optimising Data for Exemplar-based Inpainting

Lena Karos, Pinak Bheed, Pascal Peter, and Joachim Weickert

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

Abstract. Optimisation of inpainting data plays an important role in inpainting-based codecs. For diffusion-based inpainting, it is well-known that a careful data selection has a substantial impact on the reconstruction quality. However, for exemplar-based inpainting, which is advantageous for highly textured images, no data optimisation strategies have been explored yet. In our paper, we propose the first data optimisation approach for exemplar-based inpainting. It densifies the known data iteratively: New data points are added by dithering the current error map. Afterwards, the data mask is further improved by nonlocal pixel exchanges. Experiments demonstrate that our method yields significant improvements for exemplar-based inpainting with sparse data.

Keywords: exemplar-based inpainting, texture, spatial data optimisation, error maps, dithering, densification.

1 Introduction

Inpainting fills in missing parts of an image from available data [11, 17]. Originating from the field of texture synthesis [7, 12, 15], the class of exemplar-based inpainting techniques developed during the last few decades. These methods are of a non-local nature in the sense that unknown image regions are reconstructed based on known data from different locations in the image domain [2, 6, 8]. During the inpainting process, information is exchanged between known and unknown image patches, which are usually disk- or rectangle-shaped regions. Moving away from the classical concept of completely known patches, Facciolo et al. [8] introduced a framework that acts upon patches that can contain unknown pixels. This allows the method to inpaint images with sparse known data, which makes it interesting for applications such as compression [19].

Most inpainting methods that are well-known to work with sparse image data rely on partial differential equations (PDEs) [4, 17, 22]. Numerous publications have shown that a careful selection of the known data is crucial for the reconstruction quality [3, 5, 10, 13, 16, 21]. Despite the success of these data optimisation techniques for PDE-based inpainting, this idea has not yet been studied for exemplar-based inpainting. The latter is far more challenging, because exemplar-based techniques are nonlocal and thus computationally much more expensive than PDE-based approaches. However, since they offer advantages for strongly textured images, it would be desirable to have suitable data optimisation strategies.

Our Contribution. In this paper, we introduce data optimisation to exemplar-based inpainting. To this end, we use the sparse inpainting method of Facciolo et al. [8]. Building upon the idea of iterative data selection [14, 16], we come up with a densification approach that picks the most important pixels in a bottom-up fashion. In order to adapt our technique to the specific setting of exemplar-based inpainting in an efficient manner, we propose the concept of error maps. They measure the spatial distribution of the reconstruction error. We apply a modified form of the Floyd-Steinberg dithering technique [9] to convert error maps into binary images that specify the location of important known data. An evaluation of our experimental results shows that our new approach is well-suited for exemplar-based inpainting.

Related Work. Working in a patch-based setting, exemplar-based methods pursue different strategies for the search of patches and consolidation of information for the purpose of inpainting. Criminisi et al. [6] proposed a direct extension of the idea by Efros and Leung [7], where the priority order for patches to be inpainted depends on the local structural information. The encouraging results of both methods led to a popularisation of exemplar-based inpainting. Contrary to these methods, the approach of Facciolo et al. [8] relies on the idea of partial patches: Regardless of whether a patch contains unknown pixels, information is exchanged between a pair of known and unknown pixels in the compared patches. The resulting ability to inpaint from sparse known data makes the approach of Facciolo et al. particularly useful for our purpose.

Compression applications based on image inpainting are a classic motivation for data optimisation, since they need to select a sparse set of representative image points. Galić et al. [10], Schmaltz et al. [21], and Peter et al. [18] validated the potential of diffusion-based codecs with subdivision strategies for sparsifying data. These techniques restrict known locations to an adaptive regular grid that is easy to store. Their results are competitive to the quasi-standards JPEG and JPEG 2000. In particular, a proof-of-concept by Peter and Weickert [19] combined diffusion-based inpainting with exemplar-based methods to address textured images. However, they did not optimise the known data w.r.t. the exemplar-based inpainting operator due to runtime issues. The impact of this omission remains unknown.

The success of early PDE-based compression codecs [10] also sparked research on optimising inpainting data without constrained locations. For homogeneous diffusion inpainting, Belhachmi et al. [3] showed that optimal known data should be chosen according to the Laplacian magnitude. For practical purposes, they propose a dithering of this magnitude. An iterative strategy by Mainberger et al. [16] called probabilistic sparsification gradually removes pixels that are least significant for the image reconstruction. Thus, it selects the known data in a top-down manner. In contrast, the probabilistic densification method by Hoffmann et al. [14] adds significant pixels iteratively. As a postprocessing step to the sparsification approach, Mainberger et al. [16] introduced another iterative method called nonlocal pixel exchange in which known and unknown pixels pairs are swapped if the error of reconstruction from such an exchange improves. While the previous approaches rely strictly on local per-pixel errors, the densification method by Adam et al. [1] computes the global error in every iteration.

Since the probabilistic data optimisation methods [1, 13, 16] are least restricted to a specific inpainting operator such as homogeneous diffusion inpainting, our densification approach builds upon some of their core ideas. However, our adaptation to the computationally expensive setting of exemplar-based inpainting leads to key differences in judging the importance of known data: We capture the nonlocal nature of exemplar-based inpainting through global interactions by dithering of error maps.

Concerning dithering, a wide variety of image halftoning techniques are available. The method Floyd and Steinberg [9] is popular due to its simplicity: Based on the concept of error diffusion, it achieves dithering by distributing pixel errors in their respective neighbourhood. For alternative and more sophisticated dithering approaches, we refer to the monograph of Ulichney [23] and more recent papers such as that by Schmaltz et al. [20].

Structure of the Paper. First the essentials of the sparse exemplar-based inpainting method of Facciolo et al. [8] are presented in Section 2. Section 3 reviews spatial data optimisation for inpainting-based image compression. We introduce our novel approach in Section 4. Experiments and their results are presented in Section 5. Finally, we conclude the paper with a summary and give an outlook on future work in Section 6.

2 Exemplar-based Inpainting of Sparse Data

Let us consider a greyscale image as a function $f : \Omega \rightarrow \mathbb{R}$, where Ω denotes some rectangular image domain in $\mathbb{R}^2$. Given the set of sparse known data $K \subset \Omega$, the so-called *inpainting mask*, the intent of inpainting is to find an image u that is a reconstruction of the image f in the inpainting domain $\Omega \setminus K$.

Exemplar-based inpainting methods exchange information between patches according to their similarity. In the case of sparse data, the idea of using complete patches fails. To this end, Facciolo et al. [8] introduced a similarity measure V for partial patches, that only relies on the known pixels within a patch. This similarity measure between two patches that are centred at $\boldsymbol{x}$ and $\boldsymbol{x'}$ is given by the weighted squared difference function

$$V(\boldsymbol{x}, \boldsymbol{x'}) = \int_D g(\boldsymbol{x}, \boldsymbol{x'}, \boldsymbol{y}) \; \left(u(\boldsymbol{x}+\boldsymbol{y}) - u(\boldsymbol{x'}+\boldsymbol{y})\right)^2 \, d\boldsymbol{y}. \tag{1}$$

Here D denotes the domain of a patch centred at the origin, usually taken as a disk or a rectangle. The vector $\boldsymbol{y}$ indicates the displacement w.r.t. to this centre. The weighting function g is defined as

$$g\left(\boldsymbol{x}, \boldsymbol{x'}, \boldsymbol{y}\right) = \frac{g_\sigma\left(\boldsymbol{y}\right)}{\rho\left(\boldsymbol{x}, \boldsymbol{x'}\right)} \left(\alpha \mathcal{X}_K\left(\boldsymbol{x}+\boldsymbol{y}\right) + \beta \mathcal{X}_K\left(\boldsymbol{x'}+\boldsymbol{y}\right)\right). \tag{2}$$

Its normalisation parameter $\rho\left(\boldsymbol{x}, \boldsymbol{x'}\right)$ allows for a similarity measure that does not depend on the amount of known data contained in the patches. The Gaussian function $g_\sigma(\boldsymbol{y})$ with a standard deviation σ provides a spatial weighting which depends on the distance from the patch centre. The various combinations of the constants α and β for the characteristic function $\mathcal{X}_K$ result in different schemes, called A, B and AB as described in the original paper [8].

Following this definition of similarity, inpainting corresponds to the minimisation of the *data coherence term*

$$F\left(u, w\right) = \int_\Omega \int_K w\left(\boldsymbol{x}, \boldsymbol{x'}\right) V\left(\boldsymbol{x}, \boldsymbol{x'}\right) \, d\boldsymbol{x'} d\boldsymbol{x}, \tag{3}$$

where $w\left(\boldsymbol{x}, \boldsymbol{x'}\right)$ is called the *non-local weight function*. Such a weighting allows for a pixel $\boldsymbol{x}$ to be reconstructed from several different patches in the image. As w is unknown, the framework introduces a second term, namely the *weight entropy term*

$$H\left(\boldsymbol{x}, w\right) = -\int_K w\left(\boldsymbol{x}, \boldsymbol{x'}\right) \log w\left(\boldsymbol{x}, \boldsymbol{x'}\right) \, d\boldsymbol{x'}. \tag{4}$$

This term is maximised following the *principle of maximum entropy* from information theory which states that the maximum entropy distribution is the least biased estimate under the constraint of known information. Together these two terms form the energy functional

$$E\left(u, w\right) = \frac{1}{h} F\left(u, w\right) - \int_\Omega H\left(\boldsymbol{x}, w\right) \, d\boldsymbol{x}, \tag{5}$$

where h is a constant used to balance the influence of the two terms in the energy functional. The following constrained optimisation problem is then solved to find the solution u^* to the inpainting problem:

$$\left(u^*, w^*\right) = \operatorname*{arg\,min}_{u,w} E\left(u, w\right), \tag{6}$$

$$s.t.: \quad \int_K w\left(\boldsymbol{x}, \boldsymbol{x'}\right) \, d\boldsymbol{x'} = 1 \qquad \forall\, \boldsymbol{x} \in \Omega, \tag{7}$$

$$u\left(\boldsymbol{x}\right) = f\left(\boldsymbol{x}\right) \qquad \forall\, \boldsymbol{x} \in K. \tag{8}$$

Implementation. The minimisation of this energy functional is done by an alternate coordinate descent of two steps: the weight update and a step that updates the image u. Moving slightly away from the variational framework, Facciolo et al. [8] proposed the scheme O: Here they compute the patch distance V differently for different steps. The weight update step uses the scheme A, whereas the image update step uses the scheme AB.

In our experiments, we use the reference implementation and the default parameter settings provided by Facciolo et al. [8] for the scheme O.

3 Spatial Data Optimisation

In the following, we discuss the core ideas of iterative data selection that are useful to understand our method: a bottom-up [16] and a top-down [14] approach for probabilistic optimisation, and an important postprocessing strategy [16]. In the remainder of this paper we need discrete concepts. We use the symbol $\boldsymbol{\Omega}$ for the discretised image domain, and $\boldsymbol{K}$ for its set of mask pixels. The notion $\boldsymbol{f} = (f_{i,j})$ describes a discrete version of the image f, where the index (i,j) denotes the pixel. We represent the set $\boldsymbol{K}$ also by its binary mask function $\boldsymbol{c} = (c_{i,j})$ with $c_{i,j} = 1$ for $(i,j) \in \boldsymbol{K}$ and $c_{i,j} = 0$ elsewhere.

Probabilistic Sparsification by Mainberger et al. [16] is an iterative greedy algorithm that follows the natural idea of removing unimportant pixels. It starts with a full inpainting mask and removes a set of pixels in every iteration. After randomly removing a fraction p of all pixels, the so-called *candidate set*, an inpainting yields per-pixel errors for each candidate. Then a fraction q of the candidate pixels with the least error is permanently removed. This corresponds to removing $p \cdot q \cdot |\boldsymbol{K}|$ pixels that are easily reconstructed in every iteration, where $|\boldsymbol{K}|$ represents the total number of known pixels in the inpainting mask. The algorithm terminates when the desired density of pixels is achieved.

Probabilistic Densification. In contrast to probabilistic sparsification, the probabilistic densification method of Hoffmann et al. [14] starts with an empty mask. In every iteration, the algorithm first computes the inpainted image and then randomly selects a candidate set containing a fraction p of all non-mask pixels. A fraction q of pixels from the candidate set, which have the largest local error, are permanently added to the mask. This is repeated until the required density is reached. In our approach, we modify probabilistic densification.

Nonlocal Pixel Exchange. One key shortcoming of sparsification is that when a pixel is removed it never gets added back into the mask. As the algorithm relies on selecting pixels randomly for the candidate set, there exists a possibility of removing significant pixels. Therefore, Mainberger et al. [16] introduced the nonlocal pixel exchange to compensate for this disadvantage and to avoid getting trapped in a suboptimal local minimum. This iterative method exchanges a set of randomly selected known pixels with a set of unknown pixels that exhibit a large local error. The exchange is accepted if the global mean squared error (MSE) obtained by inpainting the new mask is reduced. Otherwise, the exchange is discarded and a new one is made. The inherent nature of the algorithm guarantees that any change in the mask yields a better mask. We use nonlocal pixel exchange as a postprocessing step to our method.

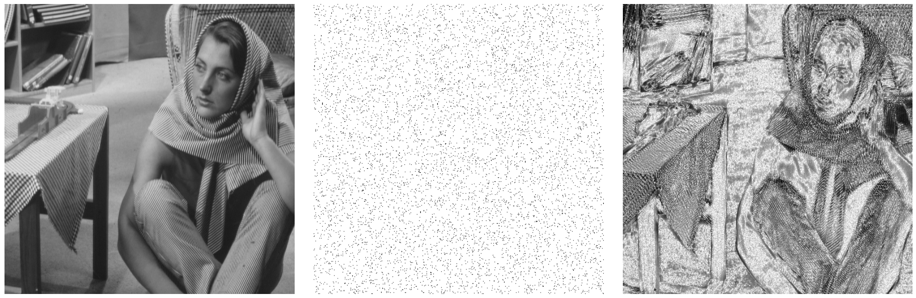

Fig. 1. Error map of Barbara image. *Left:* Original image. *Middle:* Random mask with a density of 3%. *Right:* Error map computed after exemplar-based inpainting. *Note:* For better visualisation, the error map is logarithmically rescaled such that the maximum error is denoted by black and the minimum error by white.

4 Our Approach: Densification by Dithering

While the probabilistic methods from Section 3 can be combined with any inpainting operator, they have a specific weakness with regard to exemplar-based inpainting: The error computation in classic sparsification and densification methods is highly localised due to the use of per-pixel errors, while patch-based inpainting is nonlocal. For PDE-based methods, Adam et al. [1] provide evidence that this local error computation can have a negative impact on the results of probabilistic optimisation. Instead, they propose to compute the global error associated with each candidate pixel. However, such a strategy is not feasible for exemplar-based inpainting because the non-locality of these methods leads to significantly larger runtimes than for PDE-based methods. In order to find a remedy for this issue, we consider the concepts of error maps and dithering.

Error Maps. As the two previously mentioned densification methods are iterative, they require to add mask pixels given a set of already known data in every step. In such a situation, a natural way to identify regions of importance for data selection is through error maps. Since we are interested in minimising the mean squared error of the inpainted image, we compute our error map $\boldsymbol{e} = (e_{i,j})$ pixelwise as the squared difference between the original image $\boldsymbol{f} = (f_{i,j})$ and the inpainted image $\boldsymbol{u} = (u_{i,j})$.

The error map depends directly on the inpainting mask and its density, because it is computed from the inpainting result. In particular, it gives a global overview of the spatial error distribution within the entire image domain. Figure 1 shows an example. As we shall see next, the error map becomes particularly useful when we combine it with dithering concepts.

Modified Floyd-Steinberg Dithering. The established probabilistic sparsification and densification methods suffer from two drawbacks: 1. They select only a small subset of the error map as candidate pixels. 2. They only consider local per-pixel errors, which might result from noise or small-scale features, such that choosing them would lead to a bad global approximation.

Interestingly, binary dithering can solve both problems: 1. Dithering replaces a random sampling by a deterministic, global strategy that takes into account all error values. 2. Classical binary dithering methods such as Floyd-Steinberg dithering place a data point at a certain location if the accumulated values from a neighbourhood are sufficiently large. Thus, they overcome the drawback of purely local per-pixel errors. This motivates us to perform binary dithering on error maps. To this end, we regard the error map as a greyscale image, such that its dithering provides us with the inpainting mask.

Of the many dithering methods, the classical Floyd-Steinberg dithering [9] is popular due to its simplicity. In order to adapt it to our setting, we have to take into account that the locations of mask pixels from previous iterations are already set and should remain unaltered. The dithering should only create additional mask points at different locations.

The original Floyd-Steinberg algorithm sequentially scans the image from left to right and top to bottom. Each pixel is binarised to 0 or 255, and the binarisation error is distributed to its unvisited neighbours. The stencil weights used for distribution sum up to 1, thereby preserving the average grey value. This allows us to determine how many new mask pixels are added by rescaling the average grey value of the error map accordingly (see Belhachmi et al. [3]).

In order to preserve existing mask points, we always binarise them to 0. The accumulated error at these pixels is then transferred over to the next pixel in the scan order. Overall, this corresponds to skipping known pixels for the purpose of dithering. An additional modification is the use of serpentine scan order i.e., left to right for odd numbered rows, right to left for even numbered rows and top to bottom, since this improves the results [23]. See Algorithm 1 for details.

Densification by Dithering. Since error maps depend on the mask, adding new mask pixels will change the error map. Therefore, it has to be recomputed, which requires an inpainting. As exemplar-based inpaintings are computationally expensive, we insert multiple points at the same time and recompute the error map afterwards. In practise, we partition the desired number of mask pixels into n equally sized batches. This leads to the following iterative *densification by dithering (DbD)* approach (Algorithm 2). In each iteration, DbD performs three straightforward steps:

1. Inpaint with the current inpainting mask.
2. Compute the error map.
3. Apply modified Floyd-Steinberg dithering and update the mask.

We repeat the process until we reach the required density d. Note that in the first iteration, the inpainting mask is still empty and therefore, no error map can be computed. Instead of relying on error map dithering, we choose the initial set of known pixels randomly.

Nonlocal Pixel Exchange. Since DbD is a greedy algorithm that may get trapped in a bad local optimum, we can expect to improve the reconstruction quality by postprocessing the mask with a nonlocal pixel exchange (NLPE). For a random mask pixel, we consider a set of 10 randomly chosen non-mask pixels

Algorithm 1: Modified Floyd-Steinberg Dithering of the Error Map

Input: rescaled error map $\boldsymbol{e} : \boldsymbol{\Omega} \to [0, 255]$,
binary inpainting mask $\boldsymbol{c} : \boldsymbol{\Omega} \to \{0, 1\}$

Output: modified inpainting mask $\boldsymbol{c}$

for *all pixels (i,j) scanned in serpentine order* **do**
- *{Binarise and set new mask pixel}*
- $old \leftarrow e_{i,j}$
- **if** $c_{i,j} = 0$ *and* $e_{i,j} > 127.5$ **then**
  - $e_{i,j} \leftarrow 255.0$
  - $c_{i,j} \leftarrow 1$
- **end**
- **else**
  - $e_{i,j} \leftarrow 0$
- **end**
- *{Compute binarisation error}*
- $\epsilon \leftarrow old - e_{i,j}$
- *{Diffuse binarisation error to unvisited neighbours}*
- **if** *scanning from left to right* **then**
  - $e_{i+1,j} \leftarrow e_{i+1,j} + \frac{7}{16}\epsilon$
  - $e_{i-1,j+1} \leftarrow e_{i-1,j+1} + \frac{3}{16}\epsilon$
  - $e_{i,j+1} \leftarrow e_{i,j+1} + \frac{5}{16}\epsilon$
  - $e_{i+1,j+1} \leftarrow e_{i+1,j+1} + \frac{1}{16}\epsilon$
- **end**
- **else** *scanning from right to left*
  - $e_{i-1,j} \leftarrow e_{i-1,j} + \frac{7}{16}\epsilon$
  - $e_{i+1,j+1} \leftarrow e_{i+1,j+1} + \frac{3}{16}\epsilon$
  - $e_{i,j+1} \leftarrow e_{i,j+1} + \frac{5}{16}\epsilon$
  - $e_{i-1,j+1} \leftarrow e_{i-1,j+1} + \frac{1}{16}\epsilon$
- **end**

end

as candidates. As an acceleration strategy, we use the local error to identify the most promising candidate. Following Adam et al. [1] we then compute the global MSE that would result from a swap of the mask pixel with this non-mask pixel. To speed up the process further, we compute eight NLPE steps in parallel on the different threads of a multicore CPU and take the one with the smallest error. This pixel exchange attempt is accepted if it leads to an MSE improvement. It appears natural to measure the amount of exchange attempts in terms of *cycles*: One cycle denotes as many exchange attempts as mask pixels $|\boldsymbol{K}|$. This ensures that on average each mask pixel is given a chance to be exchanged. As we consider eight parallel NLPE steps, this counts as eight exchange attempts.

Algorithm 2: Densification by Dithering

Input: original image $\boldsymbol{f} : \boldsymbol{\Omega} \to [0, 255]$,
desired pixel density d,
number of iterations n

Output: final inpainting mask $\boldsymbol{c} : \boldsymbol{\Omega} \to \{0, 1\}$

Choose a random inpainting mask $\boldsymbol{c}$ with pixel density $\frac{d}{n}$.
for $i = 1$ *to* $n - 1$ **do**
  Reconstruct with exemplar-based inpainting: $\boldsymbol{u} := inpaint(\boldsymbol{c}, \boldsymbol{f})$.
  Compute the error map $\boldsymbol{e} = (e_{i,j})$ via $e_{i,j} \leftarrow |f_{i,j} - u_{i,j}|^2$.
  Rescale the error map to $\frac{255 \cdot d}{mean(\boldsymbol{e}) \cdot n} \cdot \boldsymbol{e}$.
  Apply modified Floyd-Steinberg dithering on the error map $\boldsymbol{e}$.
  Update the current inpainting mask $\boldsymbol{c}$ by adding the dithered mask.
end

Fig. 2. Original images. *From left to right: Synthetic* (256×256), *Trui* (256×256), *Barbara* (512×512), and *Baboon* (256×256).

5 Experiments

In this section, we evaluate the performance of DbD for the four test images shown in Figure 2: *Synthetic* and *Barbara* with a mask pixel density of 5%, and *Trui* and *Baboon* with 10% density. For each of these test settings, we perform exemplar-based inpainting with random masks, DbD-optimised ones, and DbD masks with NLPE. For DbD, we present results for $n = 10$ iterations, which was found empirically to give the best quality for modified Floyd-Steinberg dithering. NLPE is run for four cycles. The largest improvements by NLPE are achieved in the first cycle, while allowing for more cycles reduces the error only incrementally.

Figure 3 displays the reconstructed images and their MSE. It shows that both visually and quantitatively, DbD leads to substantially better inpainting results than random masks, and NLPE gives significant further quality improvements. The corresponding optimised masks are depicted in Figure 4. For the *Synthetic* image, we observe that our algorithm correctly identifies that the pixels near texture boundaries are important to the exemplar-based inpainting operator. Thus, in contrast to random masks, many incorrect inpaintings at texture boundaries are avoided. For the *Trui* and the *Barbara* images, the recovery of highly textured regions is more faithful, since our data optimisation leads to higher pixel

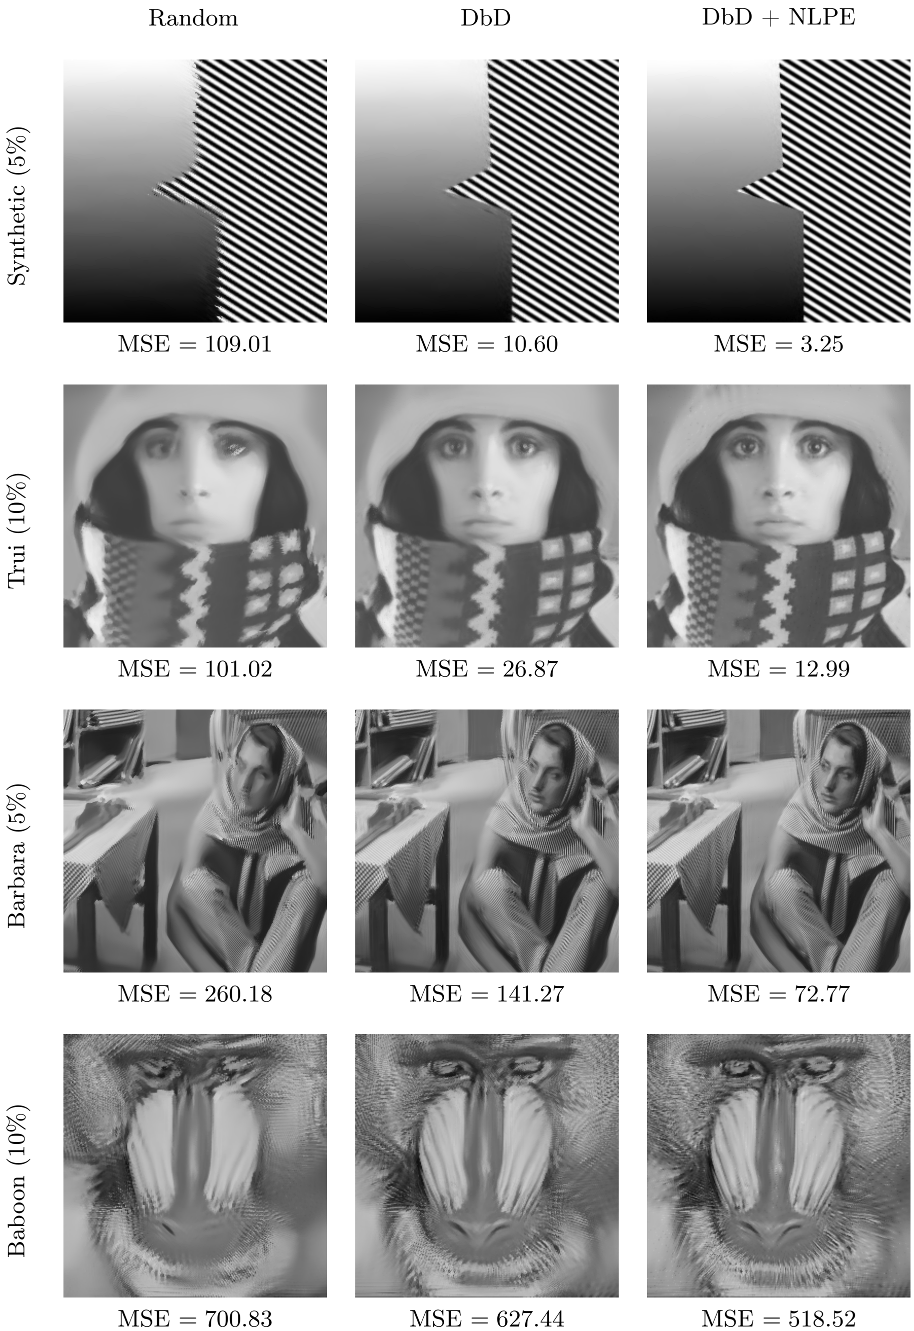

Fig. 3. Comparison of different data selection approaches. The densification by dithering (DbD) algorithm is run for $n = 10$ iterations, and we perform 4 cycles of nonlocal pixel exchange (NLPE).

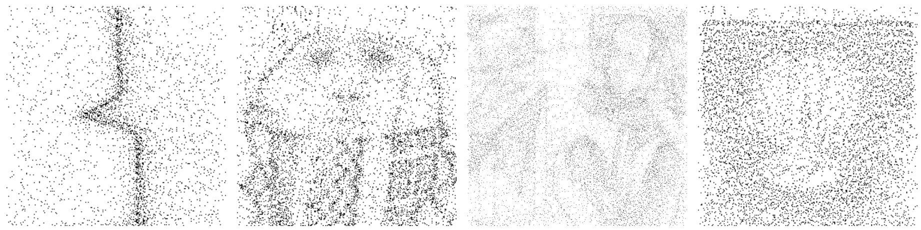

Fig. 4. Optimised inpainting masks. *From left to right:* Masks after DbD + NLPE optimisation for *Synthetic*, *Trui*, *Barbara*, and *Baboon*. The black pixels represent the known mask points.

densities in these areas. Small-scale details such as the semantically important eye regions are not oversmoothed. The *Baboon* image requires a high, almost equal mask density everywhere, such that the MSE is fairly high at 10% density, and improvements by DbD and NLPE are more moderate.

6 Conclusions and Outlook

We have presented the first paper that optimises data for exemplar-based inpainting. We achieve this with a novel approach called densification by dithering. It relies on iterative dithering of error maps for selection of important pixels. Further improvements are obtained with nonlocal pixel exchange. Our method shows that even for images which are heavily dominated by textures, a smart selection of the inpainting mask improves the quality substantially.

In our ongoing work, we explore the concept of error maps further and study more sophisticated dithering methods. Since we did not analyse the entropy of the resulting sparse data, the potential for compression applications is still unknown. Therefore, in the future we will also address this aspect and embed our method into a full image compression framework.

Acknowledgement. Part of our research has been funded by the ERC Advanced Grant INCOVID. This is gratefully acknowledged.

References

1. Adam, R.D., Peter, P., Weickert, J.: Denoising by inpainting. In: Lauze, F., Dong, Y., Dahl, A.B. (eds.) Scale Space and Variational Methods in Computer Vision, Lecture Notes in Computer Science, vol. 10302, pp. 121–132. Springer, Berlin (2017)
2. Arias, P., Facciolo, G., Caselles, V., Sapiro, G.: A variational framework for exemplar-based image inpainting. International Journal of Computer Vision 93(3), 319–347 (Jul 2011)
3. 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)

4. Bertalmio, M., Sapiro, G., Caselles, V., Ballester, C.: Image inpainting. In: Proc. SIGGRAPH '00. pp. 417–424. ACM Press, New Orleans, LI (Jul 2000)
5. Chen, Y., Ranftl, R., Pock, T.: A bi-level view of inpainting-based image compression. In: Kúkelová, Z., Heller, J. (eds.) Proc. Computer Vision Winter Workshop. Křtiny, Czech Republic (Feb 2014)
6. Criminisi, A., Pérez, P., Toyama, K.: Region filling and object removal by exemplar-based image inpainting. IEEE Transactions on Image Processing 13(9), 1200–1212 (Sep 2004)
7. Efros, A.A., Leung, T.K.: Texture synthesis by non-parametric sampling. In: Proc. IEEE International Conference on Computer Vision. vol. 2, pp. 1033–1038. IEEE, Kerkyra, Greece (Sep 1999)
8. Facciolo, G., Arias, P., Caselles, V., Sapiro, G.: Exemplar-based interpolation of sparsely sampled images. In: Cremers, D., Boykov, Y., Blake, A., Schmidt, F.R. (eds.) Energy Minimization Methods in Computer Vision and Pattern Recognition, Lecture Notes in Computer Science, vol. 5681, pp. 331–344. Springer, Berlin (2009)
9. Floyd, R.W., Steinberg, L.: An adaptive algorithm for spatial greyscale. Society of Information Display 17(2), 75–77 (1976)
10. 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)
11. Guillemot, C., Meur, O.L.: Image inpainting: Overview and recent advances. IEEE Signal Processing Magazine 31(1), 127–144 (Jan 2014)
12. Heeger, D.J., Bergen, J.R.: Pyramid-based texture analysis / synthesis. In: Proc. SIGGRAPH '95. pp. 229–238. ACM Press, Los Angeles, CA (Jul 1995)
13. Hoeltgen, L., Setzer, S., Weickert, J.: An optimal control approach to find sparse data for Laplace interpolation. In: Heyden, A., Kahl, F., Olsson, C., Oskarsson, M., Tai, X.C. (eds.) Energy Minimization Methods in Computer Vision and Pattern Recognition, Lecture Notes in Computer Science, vol. 8081, pp. 151–164. Springer, Berlin (2013)
14. Hoffmann, S., Mainberger, M., Weickert, J., Puhl, M.: Compression of depth maps with segment-based homogeneous diffusion. In: Kuijper, A., Bredies, K., Pock, T., Bischof, H. (eds.) Scale Space and Variational Methods in Computer Vision, Lecture Notes in Computer Science, vol. 7893, pp. 319–330. Springer, Berlin (2013)
15. Liang, L., Liu, C., Xu, Y.Q., Guo, B., Shum, H.Y.: Real-time texture synthesis by patch-based sampling. ACM Transactions on Graphics 20(3), 127–150 (Jul 2001)
16. Mainberger, M., Hoffmann, S., Weickert, J., Tang, C.H., Johannsen, D., Neumann, F., Doerr, B.: Optimising spatial and tonal data for homogeneous diffusion inpainting. In: Bruckstein, A.M., ter Haar Romeny, B.N., Bronstein, A.M., Bronstein, M.M. (eds.) Scale Space and Variational Methods in Computer Vision, Lecture Notes in Computer Science, vol. 6667, pp. 26–37. Springer, Berlin (2011)
17. Masnou, S., Morel, J.M.: Level lines based disocclusion. In: Proc. 1998 IEEE International Conference on Image Processing. vol. 3, pp. 259–263. Chicago, IL (Oct 1998)
18. Peter, P., Kaufhold, L., Weickert, J.: Turning diffusion-based image colorization into efficient color compression. IEEE Transactions on Image Processing 26(2), 860–869 (Feb 2017)
19. Peter, P., Weickert, J.: Compressing images with diffusion- and exemplar-based inpainting. In: Aujol, J.F., Nikolova, M., Papadakis, N. (eds.) Scale Space and Variational Methods in Computer Vision, Lecture Notes in Computer Science, vol. 9087, pp. 154–165. Springer, Berlin (2015)

20. Schmaltz, C., Gwosdek, P., Bruhn, A., Weickert, J.: Electrostatic halftoning. Computer Graphics Forum 29(8), 2313–2327 (Dec 2010)
21. 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)
22. Schönlieb, C.B.: Partial Differential Equation Methods for Image Inpainting. Cambridge University Press, Cambridge (2015)
23. Ulichney, R.: Digital Halftoning. The MIT Press, London (1987)