In J.-F. Aujol, M. Nikolova, N. Papadakis (Eds.):
Scale Space and Variational Methods in Computer Vision.
Lecture Notes in Computer Science, Vol. 9087, 154-165, Springer, Berlin, 2015.
The final publication is available at link.springer.com

Compressing Images with Diffusion- and Exemplar-based Inpainting

Pascal Peter and Joachim Weickert

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

Abstract. Diffusion-based image compression methods can surpass state-of-the-art transform coders like JPEG 2000 for cartoon-like images. However, they are not well-suited for highly textured image content. Recently, advances in exemplar-based inpainting have made it possible to reconstruct images with non-local methods from sparse known data. In our work we compare the performance of such exemplar-based and diffusion-based inpainting algorithms, dependent on the type of image content. We use our insights to construct a hybrid compression codec that combines the strengths of both approaches. Experiments demonstrate that our novel method offers significant advantages over state-of-the-art diffusion-based methods on textured image data and can compete with transform coders.

Keywords: exemplar-based inpainting, diffusion-based inpainting, image compression, texture

1 Introduction

From the initial compression approach with diffusion proposed by Galić et al. [9], a whole class of diffusion-based codecs has evolved during the last years. The R-EED codec by Schmaltz et al. [23] has demonstrated that diffusion coders can beat JPEG [18] and JPEG 2000 [25] on greyscale images. Peter and Weickert have shown in [20] that this is also possible on colour data.

The key element to the success of these methods is the ability of edge-enhancing anisotropic diffusion (EED) [26] to reconstruct images from sparse pixel data. Unfortunately, EED is not well-suited for reproducing fine-scale textures from small amounts of data. This implies that the performance of diffusion-based algorithms degrades with increasing amount of texture content in the original images. In such cases, transform-based coders still provide superior results.

In 1999, Efros and Leung [7] pioneered exemplar-based inpainting for the purpose of extending images and filling in missing or corrupted image parts. More recently, Facciolo et al. [8] have proposed exemplar-based inpainting methods that are suited for sparse known data. In our work we explore the potential of this sparse exemplar-based inpainting for image compression.

Our contribution. We assess the suitability of exemplar-based inpainting for image compression and compare it to diffusion-based inpainting. Following the results of this analysis, we construct a novel compression codec that combines the strengths of both approaches while minimising the effect of their drawbacks. In our experiments on well-known test images we demonstrate that this hybrid inpainting approach can beat established diffusion-based methods and is competitive to transform codecs also for images with rich texture.

Related work. Regarding diffusion-based compression, we rely on the anisotropic approach by Schmalz et al. [22, 23]. While our novel codec is the only diffusion-based method that deals specifically with textured images, there are PDE-based coders dedicated to other classes of data such as cartoon images [16], 3-D data [19], or depth maps [10, 12, 13].

The exemplar-based inpainting on sparse images by Facciolo et al. [8] that we use in our paper is related to a long line of classic patch-based approaches, starting with texture synthesis methods like the influental work of Efros and Leung [7]. Since a full review of the field is beyond the scope of our paper, we focus on selected publications that are related to our own work and refer to Arias et al. [2] for an in-depth review.

During the last decade, the concept of combining structure adaptive inpainting with exemplar-based ideas has been explored in several different directions. The approach of Bertalmío et al. [3] comes closest to our method since it also employs an explicit decomposition into a cartoon and a texture image. The cartoon reconstruction relies on an inpainting process that propagates information along isophotes. Patch-based inpainting restores missing parts of the texture image. However, in contrast to our paper, their decomposition is additive. This doubles the amount of original data, which is disadvantageous for compression.

Many patch-based approaches incorporate the image structure as additional guidance information. Sun et al. [24] perform patch-based texture reconstruction along manually specified curves, while Criminisi et al. [6] prioritise the reconstruction of missing image points in such a way that existing image structures are continued. In the work of Cao et al. [4], level lines extracted from a simplified version of the image are the guidance feature for exemplar-based inpainting. A different approach is pursued by Arias et al. [2] who include gradient information in a variational model for exemplar-based inpainting.

All aforementioned publications focus on image inpainting. In regards to actual compression, there are two related approaches that modify existing transform-based coders with exemplar-based inpainting. Rane et al. [21] propose a scheme that removes selected JPEG blocks and reconstructs them either with the method of Efros and Leung [7] or structure inpainting. The method of Liu et al. [14] focuses on removing visual redundancy in transform coders like H.264 or JPEG without minimising the pixel-wise error. To this end, the image is decomposed into edge-regions that are reconstructed with a combination of structure propagation and exemplar-based inpainting and texture regions that are synthesised with purely patch-based methods. Moreover, there are distantly related methods from the area of compressed sensing. These dictionary approaches (e.g. [1, 11]) use databases of image prototypes or patches for reconstruction instead of relying on partially known data like the exemplar-based methods. Aharon et al. [1] also explicitly propose compression as one application of their approach.

Organisation of the paper. First we explain the concepts of the two inpainting techniques that we combine in our paper: Section 2 covers diffusion-based inpainting, while exemplar-based inpainting is reviewed in Section 3. In Section 4 we assess the strengths and weaknesses of both approaches with respect to image compression. From these conclusions we motivate a novel hybrid compression scheme in Section 5 and analyse its performance in Section 6. We conclude our paper with a summary and outlook on future work in Section 7.

2 Diffusion-based Inpainting

Let us consider an image $f : \Omega \to \mathbb{R}$ that maps a rectangular image domain Ω to the corresponding grey values. In image compression, diffusion-based inpainting is used to reconstruct an image from a small amount of known data. Known data is only provided on the subset $K \subset \Omega$, the so-called *inpainting mask*.

The role of the diffusion process is to propagate the known information to the *inpainting domain* $\Omega \setminus K$. Thereby, the missing parts are filled in. This process of data propagation follows the partial differential equation (PDE)

$$\partial_t u = \operatorname{div}(\boldsymbol{D}\boldsymbol{\nabla} u) \quad \text{ on } \Omega \setminus K, \tag{1}$$

with reflecting boundary conditions on $\partial\Omega$. Note that the known data on K imposes Dirichlet boundary conditions on the PDE. This implies that the diffusion process converges to a nontrivial steady-state for $t \to \infty$ which yields the reconstruction of the image. Experiments show that the reconstruction does not depend on the initialisation in the inpainting domain $\Omega \setminus K$.

The most important part of the diffusion equation (1) is the diffusion tensor $\boldsymbol{D} \in \mathbb{R}^{2\times 2}$. It guides the diffusion process in terms of its eigenvalues λ_1 and λ_2 that specify the amount of diffusion in the direction of the corresponding eigenvectors $\boldsymbol{v}_1$ and $\boldsymbol{v}_2$. Thus, the choice of $\boldsymbol{D}$ is essential for the quality of the reconstruction.

For the task of image compression, Schmaltz et al. [22] have shown that edge-enhancing anisotropic diffusion (EED) [26] is particularly well-suited. EED uses an anisotropic, structure-adaptive diffusion tensor of the form

$$\boldsymbol{D} := \lambda_1(\boldsymbol{\nabla} u_\sigma)\boldsymbol{v}_1\boldsymbol{v}_1^\top + \lambda_2\boldsymbol{v}_2\boldsymbol{v}_2^\top, \tag{2}$$

$$\boldsymbol{v}_1 \parallel \boldsymbol{\nabla} u_\sigma, \quad \lambda_1(\boldsymbol{\nabla} u_\sigma) := g(|\boldsymbol{\nabla} u_\sigma|^2), \tag{3}$$

$$\boldsymbol{v}_2 \perp \boldsymbol{\nabla} u_\sigma, \quad \lambda_2 := 1, \tag{4}$$

where $u_\sigma := K_\sigma * u$ denotes a convolution of the evolving image u with a Gaussian K_σ of standard deviation σ. The tensor design in Eq. (2) implies that diffusion *across* edges is inhibited by the Charbonnier diffusivity [5]

$$g(s^2) := \frac{1}{\sqrt{1 + s^2/\lambda^2}} \tag{5}$$

with some contrast parameter $\lambda > 0$. Full diffusion *along* edges is achieved with a constant second eigenvalue $\lambda_2 := 1$. Optimising both the positions of the known data, i.e. the inpainting mask K, as well as the contrast parameter λ can improve the reconstruction quality.

Experiments in [22] show that the Gaussian convolution with K_σ in Eq. (2) plays an important role for the application of EED in image compression. It propagates structural information into the neighbourhood of each pixel and thereby allows reconstruction of edges from a very sparse inpainting mask. In EED-based compression codecs, the set K therefore usually contains scattered, isolated known pixels.

3 Exemplar-based Inpainting

The so-called *non-local inpainting* (NLI) approach of Facciolo et al. [8] follows the core idea of all patch-based methods: Missing information is filled in by exchanging information between image patches. However, in contrast to other algorithms from the field, it allows inpainting from sparse data. Therefore, we can compare it to the inpainting capabilities of EED from Section 2. For the sake of comprehensibility, we only discuss a special case of the flexible NLI framework: We have chosen algorithm AB with patch-wise non-local means from [8] for the specific task at hand.

Let us consider the same inpainting problem as in the previous section, namely finding a reconstruction u of the missing data on $\Omega \setminus K$ from the sparse known data on $K \subset \Omega$. In essence, NLI reconstructs u by minimising a patch similarity function V between pairs of image patches. It forces unknown pixel values in one patch to be similar to known values in the other patch. To this end, consider two disk-shaped patches centred in image points $\boldsymbol{x}$ and $\boldsymbol{x}'$, respectively. The similarity function V is defined as the weighted squared difference

$$V(\boldsymbol{x}, \boldsymbol{x}') = \int_D g(\boldsymbol{x}, \boldsymbol{x}', \boldsymbol{y}) \Big(u(\boldsymbol{x} + \boldsymbol{y}) - u(\boldsymbol{x}' + \boldsymbol{y}) \Big)^2 \, d\boldsymbol{y}. \tag{6}$$

Here, D is a disk around the origin, and $\boldsymbol{y}$ a coordinate relative to the respective patch centre. A common practice in patch-based methods is to rescale the individual pixel differences $u(\boldsymbol{x} + \boldsymbol{y}) - u(\boldsymbol{x}' + \boldsymbol{y})$ with Gaussian weights that reflect descending importance with the distance to the patch centres. However, in NLI, the weights g additionally account for the fact that, given a sparse inpainting mask, both patches can contain similar amounts of known data. Thus, a mutual exchange of information can be beneficial. The weights are defined as

$$g(\boldsymbol{x}, \boldsymbol{x}', \boldsymbol{y}) = \frac{K_\sigma(\boldsymbol{y})}{\rho(\boldsymbol{x}, \boldsymbol{x}')} (\chi_K(\boldsymbol{x} + \boldsymbol{y}) + \chi_K(\boldsymbol{x}' + \boldsymbol{y})), \tag{7}$$

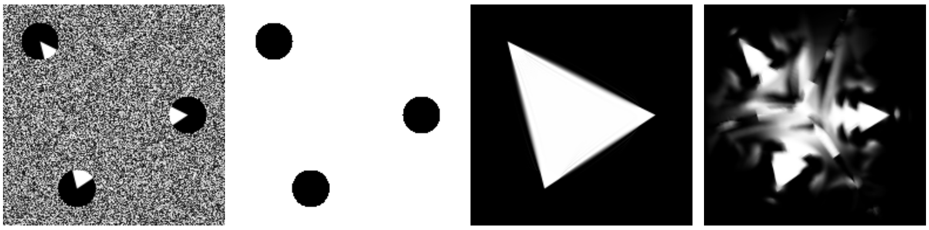

(a) initialisation (b) mask (c) EED (d) exemplar-based

Fig. 1. Experiment: Structure Propagation Known data is only given at black locations in the mask. EED ($\lambda = 0.01$, $\sigma = 4$) continues the edge structures into the inpainting domain and reconstructs an almost perfect triangle. The exemplar-based method ($h = 100$) propagates structure only locally and creates copies of structure.

where $\rho(\boldsymbol{x}, \boldsymbol{x}')$ is a normalisation term that ensures $\int_D g(\boldsymbol{x}, \boldsymbol{x}', \boldsymbol{y}) d\boldsymbol{y} = 1$ and K_σ is a Gaussian with standard deviation σ. The characteristic function χ_K of the set of known data indicates where similarities between patches should be enforced. Let $\boldsymbol{x}_1 := \boldsymbol{x} + \boldsymbol{y}$ denote an image point from the first patch and $\boldsymbol{x}_2 := \boldsymbol{x}' + \boldsymbol{y}$ the corresponding point in the second patch. If both points are unknown, i.e. not contained in K, g becomes 0 and thus, no information exchange takes place. If at least one of the two points $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ is known, we have $g > 0$ and thus V enforces similarity between those two pixels.

The second important ingredient of NLI is the decision, for which pairs of patches the similarity function V should be minimised. To this end, Facciolo et al. introduce a patch similarity weight function w and minimise the energy

$$E(u, w) = \frac{1}{h} \int_\Omega \int_K w(\boldsymbol{x}, \boldsymbol{x}') V(\boldsymbol{x}, \boldsymbol{x}') \, d\boldsymbol{x} \, d\boldsymbol{x}' - \int_\Omega H(\boldsymbol{x}, w) \, d\boldsymbol{x}, \tag{8}$$

$$\text{s.t.} \quad \int_K w(\boldsymbol{x}, \boldsymbol{x}') \, d\boldsymbol{x}' = 1. \tag{9}$$

Optimal weights w minimise the weighted total patch error according to V while maximising the entropy

$$H(\boldsymbol{x}, w) = -\int_K w(\boldsymbol{x}, \boldsymbol{x}') \log w(\boldsymbol{x}, \boldsymbol{x}') \, d\boldsymbol{x}'. \tag{10}$$

For a given u, the patch similarity weights w impose a Gaussian-like weighting of the patch differences $V(\boldsymbol{x}, \boldsymbol{x}')$:

$$w(\boldsymbol{x}, \boldsymbol{x}') = \exp\Big(-\frac{1}{h} V(\boldsymbol{x}, \boldsymbol{x}') \Big). \tag{11}$$

Thus, the parameter $h \in \mathbb{R}$ from Eq. (8) steers the standard deviation of the Gaussian weights w. In practice, the reconstruction u is found by alternating minimisation of u and w.

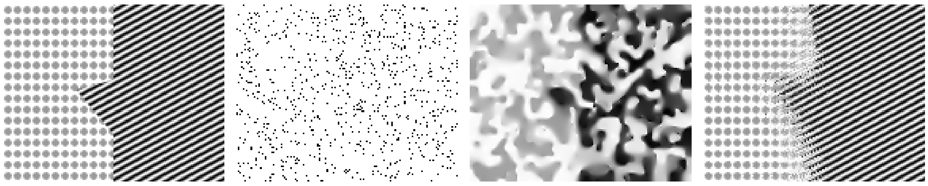

(a) original (b) mask (c) EED (d) exemplar-based

Fig. 2. Experiment: Texture. Known data from the original image is only given at black locations in the mask. EED with $\sigma = 0.8$ and $\lambda = 0.9$ completely fails to reconstruct the texture and creates a coarser pattern instead. Exemplar-based inpainting ($h = 150$) reconstructs regular texture very well, but has problems at the interface between textures.

4 Strengths and Weaknesses of the Inpainting Techniques

In the following sections we assess the advantages and drawbacks of diffusion and exemplar-based inpainting in the context of image compression. To this end, we demonstrate specific properties with simple synthetic examples and discuss their implications for practical purposes.

Let us first consider the capabilities of both algorithms in respect to structure propagation. To this end, we consider a variation of the well-known Kanizsa triangle that was used in [22] to demonstrate the capabilities of EED. For a human observer, the known data in Fig. 1(a) suggests that three corners of a triangle are given here, since human perception tends to continue sharp edges.

EED is able to preserve sharp edges and propagates image structure due to the locally adaptive diffusion tensor. Therefore, with adequate parameter choices, it is possible to match the expected reconstruction very well (see Fig. 1(c)). In contrast, exemplar-based inpainting continues structures only in close vicinity to the known data (Fig. 1(d)). In regions of the inpainting domain where known data is far away, structures are copied and multiplied.

For image compression, this behaviour implies that EED is well-suited to reconstruct coarse-scale image features from sparse known data, if they consist of mostly homogeneous areas that are separated by high contrast edges. Exemplar-based inpainting, however, tends to create visually distracting artefacts in such a setting. The reconstructions of the test image *barbara* in Fig. 3(a) and (b) illustrate the practical effects well. For example in the face region, exemplar-based inpainting repeats vertical structures of the hood in the cheeks, while EED produces a much more convincing reconstruction. Similar effects can be observed throughout the whole image.

In a second synthetic experiment we consider the reconstruction of textured areas. Fig. 2 displays the test image *interface* from [8], a representative for another extreme type of image content, namely repetitive texture. Here, EED completely fails to reconstruct the texture in a satisfying way. If an isolated region like e.g. a grey dot in the left hand side of the image is not represented by several known pixels that encode both its grey value and its shape, EED has

(a) EED (b) NLI (c) block decomposition

Fig. 3. Reconstruction of the image *barbara* with EED (a) and NLI (b) from the same known data. The reconstructions correspond to intermediate results from steps 1 and 2 of the hybrid algorithm for a compression rate of $\approx 18:1$. The block decomposition (c) with $b = 48$ indicates where EED (black) and NLI (white) yield better results.

no chance to recreate it. In contrast, exemplar-based inpainting benefits from its tendency to copy structure and create regular patterns. Its reconstruction in Fig. 2(d) is fairly close to the original, except for the sharp boundary between both repetitive patterns.

In practical compression applications, EED struggles with repetitive small-scale structure even when a lot of known data is given (e.g. Barbara's trousers in Fig. 3(a)). Therefore, compression algorithms that purely use EED for reconstruction have to store textures almost verbatim to achieve a good reconstruction. NLI produces a visually much more pleasing texture inpainting (see Fig. 3(b)) that is also close to the ground truth in regard to quantitative error measures. Therefore, a sparse inpainting mask in combination with NLI inpainting can potentially be used for compression.

5 A Hybrid Compression Algorithm for Textured Images

Block Decomposition. The core idea of our hybrid algorithm is to combine the strengths of diffusion- and exemplar-based inpainting by decomposing the image into EED and NLI blocks. From a common set of known data, EED blocks are reconstructed with diffusion inpainting and NLI blocks with exemplar-based inpainting. Shared known data for both methods offers two distinct advantages: *storage efficiency* and *direct decomposition*. Since the inpainting mask is only stored once, the only overhead generated by employing two different inpainting methods is the block decomposition and the respective model parameters. In addition, no a priori method for texture/cartoon decomposition is needed. Blocks can be directly classified as EED and NLI blocks by comparing the corresponding reconstructions to the original file, which is (in contrast to the inpainting case) available during compression.

Point Selection and Storage. In the previous sections we have only discussed diffusion and patch-based ideas in an inpainting context, where the inpainting mask $K \subset \Omega$ is already known. For compression, in addition to the right reconstruction method, also a good inpainting mask must be chosen and stored efficiently.

To this end, we employ a rectangular subdivision technique that proves to be successful in the R-EED codec [23, 22]. It limits the choice of the known data positions K to a rectangular adaptive grid that can be represented by a binary decision tree. We iteratively refine this grid by adding known pixels in image regions, where the local reconstruction error with EED inpainting exceeds a given threshold. This strategy yields a mask that is optimised for a good diffusion-based reconstruction and efficient storage in form of a binary tree. As soon as the mask is known, we apply a so-called brightness optimisation step: Introducing errors to the small amount of known pixel values can improve the reconstruction quality in the large inpainting domain $\Omega \setminus K$. For more details, we refer to [22]. The optimised grey values are finally quantised and stored with the entropy coder PAQ [15].

Our goal to create a hybrid algorithm requires some modifications to this point selection strategy. Since the subdivision grid adapts to the reconstruction abilities of EED, the point density in an R-EED inpainting mask is low in homogeneous regions, medium near coarse scale edges and very high in textured areas. Since our goal is to reconstruct homogeneous areas and sharp edges with EED and textures with exemplar-based inpainting, this point distribution is not ideal. In particular, textured areas are over-represented in the inpainting mask at the cost of more coarsely quantised grey values. Therefore, we limit the depth N of the binary tree and by that also the minimum grid size of the adaptive inpainting mask.

Avoiding Block Artefacts. For compression algorithms that use block decomposition steps, there is always the danger of visually very distracting discontinuities at block boundaries. In order to keep such effects to a minimum and simultaneously improve the overall reconstruction quality, we propose a modified diffusion-reconstruction in the decompression step. In addition to the known data on K, we also consider the reconstructed NLI blocks as Dirichlet boundary data for the final EED reconstruction. This ensures smoother transitions between NLI and EED blocks and can even improve the EED block reconstructions due to the good approximation of additional known data.

Compression Algorithm. The complete compression pipeline for our hybrid scheme consists of five steps.

1. **Depth-Limited Subdivision**: Perform rectangular subdivision with a maximum tree depth N to avoid oversampling in highly textured areas. Create a preliminary diffusion reconstruction of the whole image with EED.
2. **Exemplar-Based Inpainting**: Reconstruct the image with NLI and the inpainting mask acquired in the previous step. In order to provide a good prior for structure propagation in non-texture areas to the exemplar-based

(a) *bridge* (b) *barbara* (c) error for *barbara*

Fig. 4. (a)+(b) Original images *bridge* (No. 22090 in the Berkeley database [17]) and *barbara*. **(c)** Error comparison of *barbara* at different compression rates.

method, we initialise the inpainting domain $\Omega \setminus K$ with the diffusion reconstruction from Step 1.

3. **Block Decomposition**: Compute a block decomposition: If the mean square error (MSE) of the diffusion reconstruction is lower than the MSE of the exemplar-based inpainting in a given block, consider it to be an EED block, otherwise mark it as an NLI block. Optimising the number b of blocks in x- and y-direction can improve the overall compression quality.
4. **Encoding**: Store the known data in a modified R-EED file format (see [22]) with the number of blocks in the file header. Encode the block decomposition row-wise as a sequence of binary flags for each block (1: EED block, 0: texture block). Here, the context mixing method PAQ [15] yields the best results by encoding the binary tree, block decomposition and grey values jointly.

Decompression Algorithm. Decompression comes down to three straightforward inpainting steps, since the compressed file provides all parameters, known data and the cartoon/texture-decomposition.

1. **Diffusion-Based Reconstruction**: Extract the inpainting mask and R-EED parameters from the compressed file and reconstruct all missing data on $\Omega \setminus K$ with EED inpainting.
2. **Exemplar-Based Reconstruction**: Initialise $\Omega \setminus K$ with the result from step 1 and perform an NLI reconstruction of the inpainting domain.
3. **Final Inpainting**: Reconstruct the EED blocks with EED inpainting. Use the data of the inpainting mask Ω as well as the reconstructed NLI blocks as known data to improve the final reconstruction.

6 Experiments

In the following we evaluate the performance of our hybrid approach in comparison to the R-EED codec and the transform-based coders JPEG and JPEG 2000. For the first step in our hybrid algorithm and the results of R-EED we use the same reference implementation with PAQ for entropy coding. We optimise all model parameters as described in [23]. In our R-EED experiments, we allow larger tree depths N than in the depth-restricted hybrid step wherever this improves the result. For the implementation of NLI in our hybrid scheme we use the publicly available reference implementation of Facciolo et al. [8]. In particular, we apply the fast approximation to the variational NLI scheme AB that is referred to as algorithm O in [8]. All JPEG and JPEG 2000 images were created with the converter from ImageMagick 6.8.3-6 2013-03-04 Q16.

The results on the test image *barbara* in Fig. 4(c) and Fig. 5 demonstrate that on images with significant amounts of regular texture content, the hybrid scheme offers a significant quality gain over R-EED. R-EED stores some of the texture almost completely as known data (e.g. in the trouser region) and fails to reconstruct other parts completely (e.g. parts of the hood and tablecloth). Depending on the compression rate, the hybrid algorithm improves the mean square error (MSE) by more then 50% and is visually much more compelling. It also surpasses JPEG quantitatively by a large margin and does not suffer from similarly obvious artefacts. For compression rates larger than 25:1, the hybrid algorithm is quantitatively on par with JPEG 2000. While is not able to beat JPEG 2000, yet, it is the first time that a diffusion-based algorithm achieves comparable results on images with such a high amount of texture.

On images with irregular texture, e.g. *bridge* from Fig. 4(a), the quality gain of the hybrid algorithm over R-EED is less significant, but can still reach around 10% depending on the compression rate. While the image quality is quantitatively worse than JPEG, it is subjectively better due to the absence of block artefacts, especially in zoom-ins.

7 Conclusion

With the first combination of diffusion-based and exemplar-based inpainting from sparse data, textured images can be compressed efficiently. Our hybrid algorithm uses the full spectrum of inpainting ideas for compression while keeping the resulting overhead small. This is an important step towards closing the gap between the widely accepted general purpose encoders JPEG and JPEG 2000 and diffusion-based methods that have thus far shown their advantages primarily in more specialised applications like depth-map encoding.

For future work it would be particularly interesting to investigate how far the quality could be further improved if exemplar-based inpainting is treated equally with EED instead of using it as a post processing step. In particular, a lot of additional potential lies in optimising the choice of known data for a good trade-off between quality in EED and NLI reconstructions instead.

References

1. Aharon, M., Elad, M., Bruckstein, A.: K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. IEEE Transactions on Signal Processing 54(11), 4311–4322 (2006)

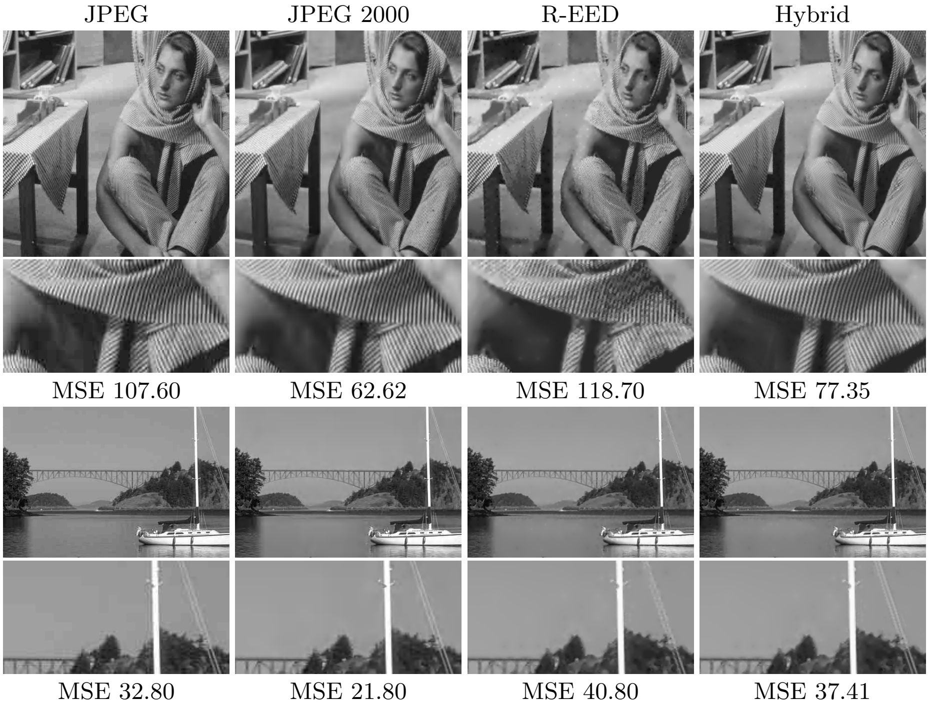

Fig. 5. Top rows: Results for *barbara* with ratio $\approx 18:1$. The hybrid algorithm uses the upper tree limit $N = 16$, block parameter $b = 48$ and the NLI parameter $h = 100$. **Bottom rows:** Results for *bridge* with ratio $\approx 19:1$. The hybrid algorithm uses the upper tree limit $N = 16$, the block size $b = 62$ and the NLI parameter $h = 25$.

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 (2011)
3. Bertalmio, M., Vese, L., Sapiro, G., Osher, S.: Simultaneous structure and texture image inpainting. IEEE Transactions on Image Processing 12(8), 882–889 (2003)
4. Cao, F., Gousseau, Y., Masnou, S., Pérez, P.: Geometrically guided exemplar-based inpainting. SIAM Journal on Imaging Sciences 4(4), 1143–1179 (2011)
5. Charbonnier, P., Blanc-Féraud, L., Aubert, G., Barlaud, M.: Deterministic edge-preserving regularization in computed imaging. IEEE Transactions on Image Processing 6(2), 298–311 (1997)
6. 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 (2004)
7. Efros, A.A., Leung, T.K.: Texture synthesis by non-parametric sampling. In: Proc. Seventh IEEE International Conference on Computer Vision. vol. 2, pp. 1033–1038. Corfu (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.

(eds.) Energy Minimization Methods in Computer Vision and Pattern Recognition, Lecture Notes in Computer Science, vol. 5681, pp. 331–344. Springer, Berlin (2009)

9. Galić, I., Weickert, J., Welk, M., Bruhn, A., Belyaev, A., Seidel, H.P.: 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, pp. 37–48. Springer, Berlin (2005)
10. Gautier, J., Meur, O.L., Guillemot, C.: Efficient depth map compression based on lossless edge coding and diffusion. In: Proc. 29th Picture Coding Symposium. pp. 81–84. Krakow, Poland (May 2012)
11. Hays, J., Efros, A.A.: Scene completion using millions of photographs. ACM Transactions on Graphics 26(3), 4 (2007)
12. 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)
13. Li, Y., Sjostrom, M., Jennehag, U., Olsson, R.: A scalable coding approach for high quality depth image compression. In: 3DTV-Conference: The True Vision - Capture, Transmission and Display of 3D Video. pp. 1–4. Zurich, Switzerland (Oct 2012)
14. Liu, D., Sun, X., Wu, F., Li, S., Zhang, Y.Q.: Image compression with edge-based inpainting. IEEE Transactions on Circuits, Systems and Video Technology 17(10), 1273–1286 (2007)
15. Mahoney, M.: Adaptive weighing of context models for lossless data compression. Tech. Rep. CS-2005-16, Florida Institute of Technology, Melbourne, Florida (Dec 2005)
16. Mainberger, M., Bruhn, A., Weickert, J., Forchhammer, S.: Edge-based compression of cartoon-like images with homogeneous diffusion. Pattern Recognition 44(9), 1859–1873 (2011)
17. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: Proc. Eigth International Conference on Computer Vision. pp. 416–423. Vancouver, Canada (Jul 2001)
18. Pennebaker, W.B., Mitchell, J.L.: JPEG: Still Image Data Compression Standard. Springer, New York (1992)
19. Peter, P.: Three-dimensional data compression with anisotropic diffusion. In: Weickert, J., Hein, M., Schiele, B. (eds.) Pattern Recognition, Lecture Notes in Computer Science, vol. 8142, pp. 231–236. Springer, Berlin (2013)
20. Peter, P., Weickert, J.: Colour image compression with anisotropic diffusion. In: Proc. 21st IEEE International Conference on Image Processing. Paris, France (Oct 2014), in press
21. Rane, S.D., Sapiro, G., Bertalmio, M.: Structure and texture fillingin of missing image blocks in wireless transmission and compression applications. IEEE Transactions on Image Processing 12(3), 296–302 (2003)
22. 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)
23. Schmaltz, C., Weickert, J., Bruhn, A.: Beating the quality of JPEG 2000 with anisotropic diffusion. In: Denzler, J., Notni, G., Süße, H. (eds.) Pattern Recognition, Lecture Notes in Computer Science, vol. 5748, pp. 452–461. Springer, Berlin (2009)

24. Sun, J., Yuan, L., Jia, J., Shum, H.Y.: Image completion with structure propagation. ACM Transactions on Graphics 24(3), 861–868 (2005)
25. Taubman, D.S., Marcellin, M.W. (eds.): JPEG 2000: Image Compression Fundamentals, Standards and Practice. Kluwer, Boston (2002)
26. Weickert, J.: Theoretical foundations of anisotropic diffusion in image processing. Computing Supplement 11, 221–236 (1996)