In J. Lellmann, M. Burger, J. Modersitzki (Eds.): Scale Space and Variational Methods. Lecture Notes in Computer Science, Vol. 11603, pp. 303–314, Springer, Cham, 2019. The final publication is available at link.springer.com.

Sparsification Scale-Spaces

Marcelo Cárdenas, Pascal Peter, and Joachim Weickert

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

Abstract. We introduce a novel scale-space concept that is inspired by inpainting-based lossy image compression and the recent denoising by inpainting method of Adam et al. (2017). In the discrete setting, the main idea behind these so-called sparsification scale-spaces is as follows: Starting with the original image, one subsequently removes a pixel until a single pixel is left. In each removal step the missing data are interpolated with an inpainting method based on a partial differential equation. We demonstrate that under fairly mild assumptions on the inpainting operator this general concept indeed satisfies crucial scale-space properties such as gradual image simplification, a discrete semigroup property or invariances. Moreover, our experiments show that it can be tailored towards specific needs by selecting the inpainting operator and the pixel sparsification strategy in an appropriate way. This may lead either to uncommitted scale-spaces or to highly committed, image-adapted ones.

Keywords: scale-space, sparsification, inpainting, diffusion, image compression

1 Introduction

Lossy image compression methods and scale-spaces share similar philosophies in a number of aspects: The former ones are based on information reduction, while the latter ones aim at image simplification. Both the amount of information reduction as well as the degree of image simplification can be steered by a free parameter: the quality parameter which influences the compression rate, and the scale parameter. Moreover, both concepts reveal naturally certain denoising properties: For instance, wavelet shrinkage can be used both for compression as well as for denoising, and nonlinear diffusion methods create scale-spaces that offer structure-preserving denoising.

In view of these similarities, it is surprising that attempts of both communities to fertilise each other are still fairly limited. One of the reasons lies in the fact that both paradigms traditionally rely on different techniques: Typical lossy compression methods are based on orthogonal or unitary transforms such as the discrete cosine or the discrete wavelet transforms [13, 23], whereas scale-spaces usually involve partial differential equations (PDEs) [2, 9, 19, 25] or pseudodifferential evolutions [6, 21].

One of the few notable attempts to connect both worlds goes back to Chambolle and Lucier [4], who interpreted iterated shift-invariant wavelet shrinkage as a nonlinear smoothing scale-space. Although these authors emphasised that it is not given by a PDE, this statement may be questioned in view of several equivalence results between iterated wavelet shrinkage and nonlinear diffusion processes [26].

Inpainting-based compression [7] constitutes an alternative to classical transform-based codecs: One stores only a carefully selected subset of all pixels and reconstructs the missing image structures by inpainting which usually involves PDEs. The less pixels are stored, the more information is discarded. This suggests that also such approaches reveal scale-space properties. This claim is also supported by a recent paper by Adam et al. [1], who use closely related inpainting ideas and data sparsification concepts to design novel denoising methods.

Our Goals. In the present paper we use these ideas to introduce a novel class of scale-space processes which we name sparsification scale-spaces. For the sake of simplicity, we restrict ourselves to the discrete setting. The basic idea of constructing such a scale-space is as follows: One starts with the original image and gradually removes all pixels until only a single pixel remains. At those locations where no information is kept, we interpolate the missing image structure with some PDE-based inpainting. While this framework is strongly inspired by inpainting-based image compression which uses similar sparsification strategies [11], several questions have to be answered in this context, e.g.

– Can one prove that this scale-space representation creates indeed a simplifying transformation? Does it also satisfy other typical scale-space properties such as a semigroup structure or invariances?
– How powerful and flexible is this family of scale-spaces? Can it be made uncommited or can one adapt it in order to reward important structures with a larger lifetime over scales? Is it strongly dependent on the inpainting operator?

Our paper will give answers to all these questions.

Structure of the Paper. Our publication is organised as follows. In Section 2 we formalise the concept of discrete sparsification scale-spaces and establish a number of theoretical properties which show that our processes under consideration indeed qualify as scale-spaces under fairly mild assumptions on the inpainting operator. Section 3 is devoted to additional aspects of specific interest, and it presents several experiments which illustrate the flexibility of the sparsification scale-space framework. The paper is concluded with a summary and an outlook in Section 4.

2 Theoretical Results

In this section we provide a formalisation of discrete sparsification scale-spaces, and we show that they satisfy all essential properties of a space- and time-discrete scale-space.

2.1 Formalisation of Discrete Sparsification Scale-Spaces

To define our sparsification scale-space in a discrete setting, we have to select its two components first: a sparsification strategy, and an inpainting method that involves a scale-space operator. Thus, we will first elaborate on these two concepts and then present our formal definition.

Sparsification Strategy. We represent a discrete greyscale image with N pixels by a vector $\boldsymbol{f} \in \mathbb{R}^N$. Let its corresponding index set be denoted by $J = \{1, ..., N\}$. We define a *known data sequence* $(K^\ell)_{\ell=0}^{N-1}$ of N nested subsets of J by imposing two conditions:

- The first set K^0 satisfies $K^0 = J$.
- For $\ell = 1,...,N-1$, every set K^ℓ is a subset of $K^{\ell-1}$ and contains $N-\ell$ elements.

Thus, K^ℓ is created by removing a single index m_ℓ from $K^{\ell-1}$. Different strategies for pixel removal will be discussed in Section 3.

Instead of specifying a known data sequence $(K^\ell)_{\ell=0}^{N-1}$, we can equivalently specify its *sparsification path* $\boldsymbol{m} = (m_1, ..., m_{N\text{-}1})$. Obviously the sparsification path offers a very compact representation of the known data sequence, which is useful when we want to store this information.

Another equivalent representation of the known data sequence is the *mask sequence*. It consists of a sequence $(\boldsymbol{C}^\ell)_{\ell=0}^{N-1}$ of diagonal matrices $\boldsymbol{C}^\ell \in \mathbb{R}^{N\times N}$ with the following property: The diagonal elements $c_{i,i}^\ell$ are 1 if $i \in K^\ell$, and 0 elsewhere.

Scale-Space Induced Inpainting. Apart from these sparsification concepts, we also need a discrete inpainting method. To this end, we consider a linear operator $\boldsymbol{A} \in \mathbb{R}^{N\times N}$ or a nonlinear operator $\boldsymbol{A}(\boldsymbol{u}) : \mathbb{R}^N \to \mathbb{R}^{N\times N}$ that generates a space-discrete scale-space evolution $\{\boldsymbol{u}(t)\,|\,t \geq 0\}$ via

$$\boldsymbol{u}(0) = \boldsymbol{f}, \tag{1}$$

$$\frac{d\boldsymbol{u}}{dt} = \boldsymbol{A}(\boldsymbol{u})\,\boldsymbol{u}. \tag{2}$$

Here we have chosen the nonlinear formulation, which includes the linear one. We assume that this scale-space satisfies classical requirements that can be classified into three categories [2]:

- architectural properties (e.g. the semigroup property),

– simplification properties (such as energy minimisation properties, the existence of a Lyapunov functional, or causality in terms of a maximum–minimum principle),
– invariances (e.g. w.r.t. additive shifts in the grey values).

This space-discrete scale-space can be seen as a dynamical system. Typically it arises from a space discretisation of a continuous scale-space which evolves an image under a partial differential equation (PDE), e.g. a diffusion equation [25]. Moreover, the structure of $\boldsymbol{A}$ also involves discretisations of boundary conditions, such as reflecting (homogeneous Neumann) boundary conditions.

To use the operator $\boldsymbol{A}(\boldsymbol{u})$ for inpainting purposes, let us assume that we are given some image $\boldsymbol{f} \in \mathbb{R}^N$, and that its grey values are only reliable in a subset K of its pixel set J. Then one can inpaint the missing information in $J \setminus K$ by solving the following problem for $\boldsymbol{u}$:

$$\boldsymbol{C}(\boldsymbol{u}-\boldsymbol{f}) \;-\; (\boldsymbol{I}-\boldsymbol{C})\;\boldsymbol{A}(\boldsymbol{u})\,\boldsymbol{u} \;=\; \boldsymbol{0} \tag{3}$$

where $\boldsymbol{I} \in \mathbb{R}^{N\times N}$ is the identity matrix, and the diagonal matrix $\boldsymbol{C}$ denotes the inpainting mask associated to the known pixel set K. Thus, our restoration $\boldsymbol{u}$ is identical to $\boldsymbol{f}$ on K, and it is inpainted with the operator $\boldsymbol{A}(\boldsymbol{u})$ in $J \setminus K$. This operator uses the data $\boldsymbol{f}$ in K as additional Dirichlet boundary data.

One may wonder if it is natural to base an inpainting process on a scale-space operator. Actually this makes a lot of sense: It is not difficult to see that basically all PDE-based inpainting operators enjoy scale-space properties, although some representatives such as higher-order operators may not have been used in this context. For more details on PDE-based inpainting we refer to [22].

Sparsification Scale-Spaces. Having a sparsification strategy and an inpainting process at our disposal, it is not difficult to define a sparsification scale-space. Let $\boldsymbol{f} = (f_i)_{i=1}^N \in \mathbb{R}^N$ be our image, $(K^\ell)_{\ell=0}^{N-1}$ a known data sequence, $(\boldsymbol{C}^\ell)_{\ell=0}^{N-1}$ its associated mask sequence, and let $\boldsymbol{A}(\boldsymbol{u})$ denote a scale-space induced inpainting operator. Then we define a (space- and time-discrete) *sparsification scale-space* $(\boldsymbol{u}^\ell)_{\ell=0}^{N-1}$ of $\boldsymbol{f}$ as the solution set of the following sequence of inpainting problems:

$$\boldsymbol{C}^\ell\big(\boldsymbol{u}^\ell-\boldsymbol{f}\big) \;-\; \big(\boldsymbol{I}-\boldsymbol{C}^\ell\big)\;\boldsymbol{A}\left(\boldsymbol{u}^\ell\right)\boldsymbol{u}^\ell \;=\; \boldsymbol{0} \qquad (\ell = 0, ..., N-1). \tag{4}$$

If we interpret this equation as a space-discrete elliptic PDE, our time-discrete scale-space solves a sequence of elliptic problems. In this sense, it has some structural similarities with scale-spaces created by the iterative regularisation methods of Radmoser et al. [17]. The central remaining question, however, is the following: In which sense does a sparsification scale-space satisfy the typical scale-space properties? It shall be answered next.

2.2 Scale-Space Properties

The good news after the somewhat tedious formalisation in the previous subsection is that it makes a subsequent analysis of the scale-space properties of sparsification scale-spaces surprisingly simple and intuitive. Let us now verify six important properties.

Property 1 (Original Image as Initial State).

For $\ell = 0$, Equation (4) comes down to

$$\boldsymbol{C}^0\big(\boldsymbol{u}^0 - \boldsymbol{f}\big) \;-\; \big(\boldsymbol{I} - \boldsymbol{C}^0\big)\;\boldsymbol{A}\left(\boldsymbol{u}^0\right)\boldsymbol{u}^0 \;=\; \boldsymbol{0}\,, \tag{5}$$

where $\boldsymbol{C}^0$ denotes the mask associated to K^0. However, K^0 is defined as the full index set J, which implies that $\boldsymbol{C}^0$ is identical to the identity matrix $\boldsymbol{I}$. In this case, Equation (5) simplifies to $\boldsymbol{u}^0 - \boldsymbol{f} = \boldsymbol{0}$.

Property 2 (Semigroup Property).

The semigroup property states that one can construct a scale-space in a cascadic manner. This property follows directly from the fact that $\boldsymbol{u}^\ell$ depends only on $\boldsymbol{f}$ and the nested known data sequence $(K^\ell)_{\ell=0}^{N-1}$. As long as one keeps $\boldsymbol{f}$ and this sequence, the sparsification scale-space has an obvious semigroup structure: A result $\boldsymbol{u}^{\ell+n}$ can be obtained in $\ell+n$ steps by starting from $\boldsymbol{u}^0 = \boldsymbol{f}$, or in n steps by starting from $\boldsymbol{u}^\ell$.

Property 3 (Maximum–Minimum Principle).

Traditionally, maximum–minimum principles are related to causality in the sense of Hummel [8]. They state that during the scale-space evolution, the greyscale range of the gradually simplified images lies in the range of the original image. For sparsification scale-spaces this property is always satisfied whenever the scale-space induced inpainting operator $\boldsymbol{A}(\boldsymbol{u})$ fulfils a maximum-minimum principle: Then the inpainting solution $\boldsymbol{u}^\ell = (u_i^\ell)_{i=1}^N$ satisfies

$$\min_{j\in J} f_j \;\le\; \min_{j\in K^\ell} f_j \;\le\; u_i^\ell \;\le\; \max_{j\in K^\ell} f_j \;\le\; \max_{j\in J} f_j \qquad \text{for all } i \in J, \tag{6}$$

since the known data sequence $(K^\ell)_{\ell=0}^{N-1}$ is nested and starts with $K^0 = J$.

Property 4 (Lyapunov Sequences).

Often the scale-space evolution (2) can be derived as the gradient descent evolution of some energy function $E(\boldsymbol{u})$. For instance, in a linear setting with operator $\boldsymbol{A}$, one minimises the quadratic form

$$E(\boldsymbol{u}) = \tfrac{1}{2}\,\boldsymbol{u}^\top \boldsymbol{A}\boldsymbol{u}\,. \tag{7}$$

In our inpainting setting, we have additional interpolation constraints on the set K^ℓ, which obviously results in a larger energy. By increasing ℓ, we reduce the nested constraint set K^ℓ, which means that we also reduce the energy gradually. This discussion shows that energy minimisation properties of the scale-space operator $\boldsymbol{A}(\boldsymbol{u})$ allow us to interpret the energy as a Lyapunov sequence for the sparsification scale-space, i.e.

$$E(\boldsymbol{u}^\ell) \le E(\boldsymbol{u}^{\ell-1}) \qquad \text{for } \ell = 1,\ldots,N-1. \tag{8}$$

Lyapunov sequences provide important criteria for the simplification properties of scale-spaces [25]. Interestingly, they may even exist when no energy function is known: For instance, it is sufficient that the operator $\boldsymbol{A}(\boldsymbol{u})$ satisfies a maximum–minimum principle. Then the nested known data sequence $(K^\ell)_{\ell=0}^{N-1}$ implies that the total image contrast

$$\Phi(\boldsymbol{u}^\ell) \ := \ \max_{i \in J} u_i^\ell - \min_{i \in J} u_i^\ell \ = \ \max_{i \in K^\ell} f_i - \min_{i \in K^\ell} f_i \tag{9}$$

decreases in ℓ. Thus, $(\Phi(\boldsymbol{u}^\ell))_{\ell=0}^{N-1}$ can be regarded as a Lyapunov sequence. Recently such sequences based on the total image contrast have been introduced by Welk et al. [27] to establish scale-space properties for FAB diffusion.

Property 5 (Invariances).

Any invariance of the scale-space induced inpainting operator $\boldsymbol{A}(\boldsymbol{u})$ is directly inherited to its sparsification scale-space.

Property 6 (Convergence to a Flat Steady-State).

A discrete sparsification scale-space consists of N images $\{\boldsymbol{u}^0, ..., \boldsymbol{u}^{N-1}\}$. The final image $\boldsymbol{u}^{N-1}$ is always obtained by inpainting with a single mask pixel. In this case, a scale-space induced inpainting operator $\boldsymbol{A}(\boldsymbol{u})$ that implements reflecting boundary conditions will create a flat image.

This completes our discussion of the essential properties that are satisfied by sparsification scale-spaces. We observe that all properties are either directly inherited from the scale-space induced inpainting operator or follow from the nested structure of the known data sequence.

3 Specific Aspects and Experiments

Our theoretical framework from the previous section allows to define sparsification scale-spaces for any PDE-based scale-space operator that can be expressed by the minimisation of an energy. Therefore, many well-known filters are viable for this task. The simplest one is the harmonic operator (also known as homogeneous diffusion) [9], but our model also covers higher-order linear operators such as biharmonic inpainting [5]. Moreover, a variety of nonlinear filters are also applicable. For instance, all operators investigated by Peter et al. [16] fit into the framework. This includes the nonlinear isotropic Perona–Malik model [14] as well as anisotropic filters, such as the approach of Tschumperlé and Deriche [24] or the method of Roussos and Maragos [18].

The experiments in Fig. 1 illustrate the influence of different inpainting methods on the behaviour of the corresponding sparsification scale-space. With the same randomly sparsified masks, both scale-space simplify the image in a similar fashion. In both cases, we start with the original image at 100 % mask density. The density acts as the scale parameter: With decreasing number of mask pixels, the reconstructions gradually loose detail until only a single pixel remains and we reach a constant steady-state. Note that this implies that for an image with

$m \times n$ pixels, we reach a finite extinction time after $m \cdot n - 1$ discrete time steps. Visually, the inpainting results at intermediate time steps differ significantly due to the individual properties of the operators: Harmonic inpainting creates visible singularities at the locations of known data. They are caused by the logarithmic singularity of the Green's function of the 2D Laplacian. Biharmonic inpainting yields much smoother results, since its Green's function is continuously differentiable. For more information on Green's functions we refer the reader to [12].

In addition to the influence of the inpainting approach, sparsification scale-spaces depend significantly on the order in which mask points are removed. The framework from Section 2 only requires that a single pixel is removed in each discrete time step, but it does not prescribe which one. For a given inpainting operator, this order of removal, the sparsification path, uniquely defines the scale-space.

On one hand, sparsification strategies can be uncommitted, such as the random removal of one mask point per step in Fig. 1. It does not depend on the original image in any way. By randomising the removal, we ensure that the lifetime expectation of each pixel is identical. Thus, the notion of uncommittedness must be understood in a statistical sense. On the other hand, our framework also allows adaptive sparsification that takes into account the image structure. Fig. 2 shows results of a global adaptive sparsification strategy that has been inspired by the denoising by inpainting approach of Adam et al. [1]. In each step, we remove the single mask point that yields the smallest mean squared error (MSE) of the reconstruction. To speed up this algorithm, we consider 200 randomly selected mask points as candidates for removal in every iteration instead of an exhaustive search. Both the homogeneous and biharmonic inpainting yield considerably different results compared to the uncommitted sparsification path of Fig. 2. The adaptive sparsification assigns different life times to different pixels according to their importance for the reconstruction. This preserves important features of the original image over a longer period of time. Although the inpainting operators are still linear, this adaptive behaviour resembles nonlinear scale-spaces such as the Perona-Malik model, where image edges selected by a contrast parameter have a longer survival time in scale-space [14]. However, it should be noted that the adaptive sparsification is parameter-free.

Apart from denoising by inpainting, adaptive inpainting scale-spaces are of significant importance for image compression. Stochastic sparsification approaches similar to the one in our experiment have been proposed to identify known data to be stored in inpainting-based codecs [11, 15]. In particular, there are variants of sparsification [15] that allow to store the corresponding masks efficiently as a binary tree. This concept is the foundation of the most successful inpainting-based codecs [7, 20]. Thus, sparsification scale-spaces have a direct impact on practical applications in the field of image compression.

The distinction between uncommitted and committed scale-spaces sheds some light on an interesting additional aspect: Traditional uncommitted scale-spaces such as the Gaussian scale-space try to remove information as quickly and completely as possible. Sometimes this criterion is even stated explicitly, e.g. in Iijima's principle of maximum loss of figure impression [10]. On the other hand, a committed scale-space that aims at representations which are useful for image compression tries to achieve the exact opposite: It simplifies an image while aiming at a maximal preservation of the figure impression. It is nice and useful that sparsification scale-spaces allow to accommodate both philosophies.

4 Conclusions and Future Work

We have introduced and analysed a new scale-space concept: sparsification scale-spaces. They simplify a digital image via a gradual removal of pixels and inpaint the discarded image structures. By construction, they offer a great flexibility which lies in the freedom how their two ingredients are chosen:

- The sparsification path can be uncommitted or adaptive. In the latter case one follows a specific task-driven strategy, which allows e.g. structure-preserving image simplification.
- For the inpainting operator, any PDE-based scale-space operator is permissible that can be derived from energy minimisation or satisfies a maximum–minimum principle.

Our theoretical analysis has shown that sparsification scale-spaces satisfy all reasonable assumptions on a space- and time-discrete scale-space. They are either inherited from the scale-space properties of the underlying inpainting operator or they follow from the nested structure of the gradually sparsified inpainting data.

Interestingly, sparsification scale-spaces differ from most other scale-spaces in a number of aspects:

- The space discretisation of the image directly implies a natural time discretisation: It is determined by the number of pixels. We are not aware of any other scale-space which shares this property.
- After finitely many time steps, adaptive scale-spaces reach their finite extinction time. Apart from a few exceptions such as the total variation (TV) flow [3] and evolutions based on mean curvature motion or its affine invariant variant [2], this behaviour is not often met in the scale-space literature.
- Sparsification scale-spaces are perfectly suited for compression applications: If we disregard specific coding aspects such as quantisation, an adaptive sparsification scale-space can contain the entire family of compressed versions of the original image. This natural link between scale-space ideas and compression concepts is new and has been made possible by the paradigm of inpainting-based image compression. It may lay the foundations of a more fruitful exchange of ideas between both fields, and it may also inspire novel denoising methods.

Fig. 1: **Randomly Sparsified Scale-Spaces**: A sparsification with purely random selection of the removed pixels provides the same mask for harmonic and biharmonic inpainting. These uncommitted scale-spaces differ only due to the inpainting operator.

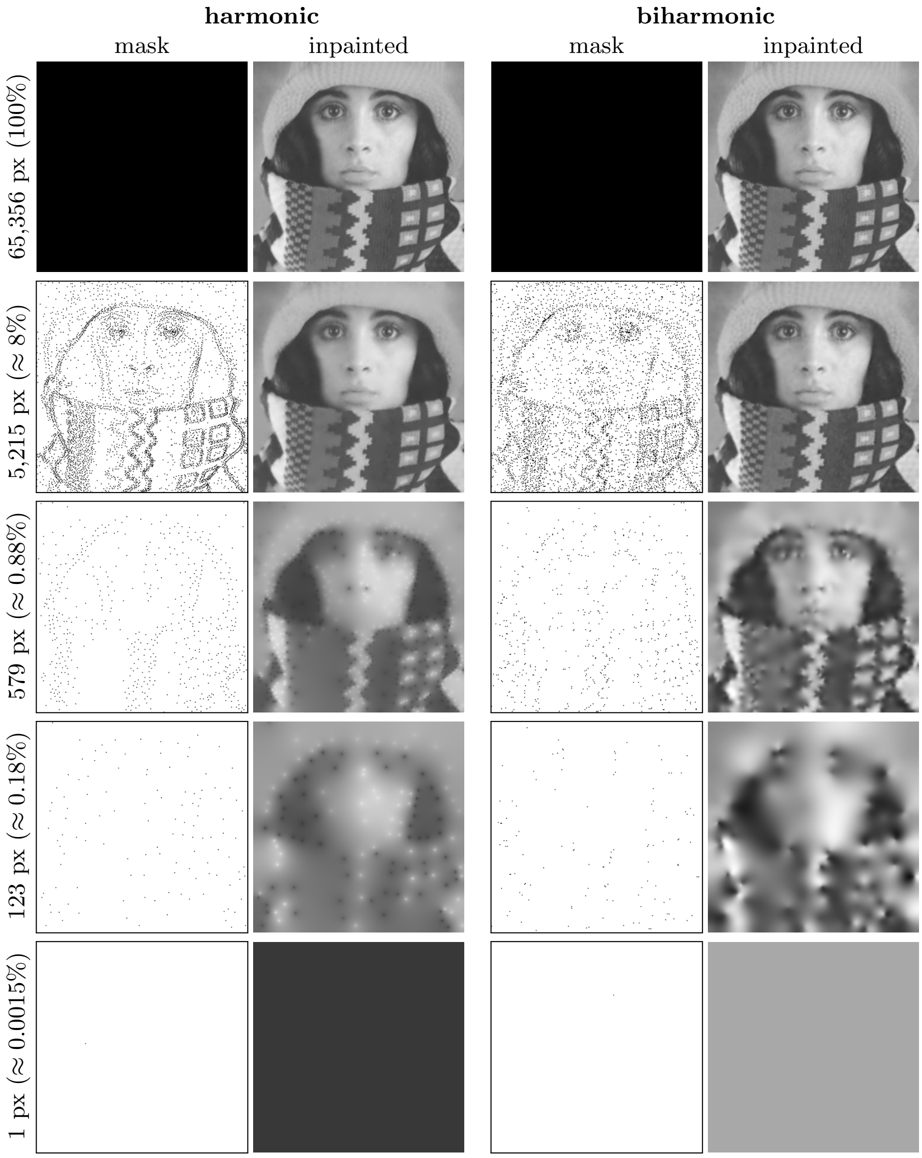

Fig. 2: **Adaptively Sparsified Scale-Spaces**: An image-adaptive stochastic sparsification leads to different masks for harmonic and biharmonic inpainting. Important image structures such as edges are better preserved than for an uncommitted sparsification scale-spaces.

In our future work, we are going to analyse and evaluate a broader spectrum of representatives within the large family of sparsification scale-spaces, covering e.g. also nonlinear methods and approaches that are not derived from a variational formulation.

It should be noted that we have restricted ourselves to discrete scale-spaces so far. Interestingly, the entire framework can also be extended to continuous sparsification scale-spaces. However, since their theory involves more technical challenges, it will be treated in a journal paper.

Acknowledgements. This project 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).

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, Cham (2017)
2. Alvarez, L., Guichard, F., Lions, P.L., Morel, J.M.: Axioms and fundamental equations in image processing. Archive for Rational Mechanics and Analysis **123**, 199–257 (1993)
3. Andreu–Vaillo, F., Caselles, V., Mazon, J.M.: Parabolic Quasilinaer Equations Minimizing Linear Growth Functionals, Progress in Mathematics, vol. 223. Birkhäuser, Basel (2004)
4. Chambolle, A., Lucier, B.L.: Interpreting translationally-invariant wavelet shrinkage as a new image smoothing scale space. IEEE Transactions on Image Processing **10**(7), 993–1000 (2001)
5. Duchon, J.: Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. RAIRO Analyse Numérique **10**, 5–12 (1976)
6. Duits, R., Florack, L., de Graaf, J., ter Haar Romeny, B.: On the axioms of scale space theory. Journal of Mathematical Imaging and Vision **20**, 267–298 (2004)
7. Galić, I., Weickert, J., Welk, M., Bruhn, A., Belyaev, A., Seidel, H.P.: Image compression with anisotropic diffusion. Journal of Mathematical Imaging and Vision **31**(2–3), 255–269 (Jul 2008)
8. Hummel, R.A.: Representations based on zero-crossings in scale space. In: Proc. 1986 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 204–209. IEEE Computer Society Press, Miami Beach, FL (Jun 1986)
9. Iijima, T.: Basic theory on normalization of pattern (in case of typical one-dimensional pattern). Bulletin of the Electrotechnical Laboratory **26**, 368–388 (1962), in Japanese
10. Iijima, T.: Basic equation of figure and observational transformation. Systems, Computers, Controls **2**(4), 70–77 (1971), in English
11. Mainberger, M., Hoffmann, S., Weickert, J., Tang, C.H., Johannsen, D., Neumann, F., Doerr, B.: Optimising spatial and tonal data for homogeneous diffusion inpainting. In: Bruckstein, A.M., ter Haar Romeny, B., 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 (2012)
12. Melnikov, Y.A., Melnikov, M.Y.: Green's Functions: Construction and Applications. De Gruyter, Berlin (2012)
13. Pennebaker, W.B., Mitchell, J.L.: JPEG: Still Image Data Compression Standard. Springer, New York (1992)
14. Perona, P., Malik, J.: Scale space and edge detection using anisotropic diffusion. IEEE Transactions on Pattern Analysis and Machine Intelligence **12**, 629–639 (1990)
15. Peter, P., Hoffmann, S., Nedwed, F., Hoeltgen, L., Weickert, J.: Evaluating the true potential of diffusion-based inpainting in a compression context. Signal Processing: Image Communication **46**, 40–53 (Aug 2016)
16. Peter, P., Weickert, J., Munk, A., Krivobokova, T., Li, H.: Justifying tensor-driven diffusion from structure-adaptive statistics of natural images. In: Tai, X.C., Bae, E., Chan, T.F., Lysaker, M. (eds.) Energy Minimization Methods in Computer Vision and Pattern Recognition, Lecture Notes in Computer Science, vol. 8932, pp. 263–277. Springer, Berlin (2015)
17. Radmoser, E., Scherzer, O., Weickert, J.: Scale-space properties of nonstationary iterative regularization methods. Journal of Visual Communication and Image Representation **11**(2), 96–114 (Jun 2000)
18. Roussos, A., Maragos, P.: Tensor-based image diffusions derived from generalizations of the total variation and Beltrami functionals. In: Proc. 17th IEEE International Conference on Image Processing. pp. 4141–4144. Hong Kong (Sep 2010)
19. Scherzer, O., Weickert, J.: Relations between regularization and diffusion filtering. Journal of Mathematical Imaging and Vision **12**(1), 43–63 (Feb 2000)
20. 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)
21. Schmidt, M., Weickert, J.: Morphological counterparts of linear shift-invariant scale-spaces. Journal of Mathematical Imaging and Vision **56**(2), 352–366 (Oct 2016)
22. Schönlieb, C.B.: Partial Differential Equation Methods for Image Inpainting. Cambridge University Press, New York (2015)
23. Taubman, D.S., Marcellin, M.W. (eds.): JPEG 2000: Image Compression Fundamentals, Standards and Practice. Kluwer, Boston (2002)
24. Tschumperlé, D., Deriche, R.: Vector-valued image regularization with PDEs: A common framework for different applications. IEEE Transactions on Pattern Analysis and Machine Intelligence **27**(4), 506–516 (Apr 2005)
25. Weickert, J.: Anisotropic Diffusion in Image Processing. Teubner, Stuttgart (1998)
26. Weickert, J., Steidl, G., Mrázek, P., Welk, M., Brox, T.: Diffusion filters and wavelets: What can they learn from each other? In: Paragios, N., Chen, Y., Faugeras, O. (eds.) Handbook of Mathematical Models in Computer Vision, pp. 3–16. Springer, New York (2006)
27. Welk, M., Weickert, J., Gilboa, G.: A discrete theory and efficient algorithms for forward-and-backward diffusion filtering. Journal of Mathematical Imaging and Vision **60**(9), 1399–1426 (Nov 2018)