23 May 2017

Mathematics can make you fly?

Dr Carola-Bibiane Schönlieb

Figure 1: Mathematics can make you fly? Image courtesy of Joana Grah, Kostas Papafitsoros and Carola-Bibiane Schönlieb.

How is this possible? How can Joana – the woman in the picture and a master student in mathematics – fly? Does she have supernatural powers? The clue to the solution of this mystery can be seen on the blackboard. It is a partial differential equation that can be used for digital image inpainting. Inpainting denotes the process whereby specified parts of images are filled in, based on the remaining part of the image. In the present example we solve this equation numerically and are able to remove the stool on which Joana was sitting originally. She appears to fly! While this may seem like gadgetry image inpainting has wide ranging practical applications: from the restoration of satellite images, enhancement of medical images, the renovation of digital photographs and artwork, to special effects in images and videos as in the present photograph. Image inpainting is ubiquitous. The picture in Figure 1 won the 2014 EPSRC Science Photo Award in the Category 'People'.

Digital image restoration - what and where? In modern society we encounter digital images in a lot of different situations: from everyday life, where analogue cameras have long been replaced by digital ones, to their professional use in medicine, earth sciences, arts, and security applications. The plain vastness of images and videos that exist in our digital system nowadays makes their unaided processing and interpretation by humans impossible. Automatic storage management, processing and analysis algorithms are needed to be able to retrieve only the essence of what the visual world has up its sleeve. Moreover, certain acquisition devices – such as magnetic resonance tomography or remote sensing of the atmosphere – do not immediately provide us with the kind of information relevant to our needs. Mathematical inversion algorithms are needed to extract this information from the physical and statistical laws that relate the measurements with the image.

The organization and processing of digital images is known under the name of image processing or computer vision.

We often have to deal with the processing of images, e.g., the restoration of images corrupted by noise, blur, or intentional scratching. The idea behind image processing is to provide methods that improve the quality of these images by postprocessing them. For an introduction to digital image processing we refer to [3, 9].Virtual image restoration or image interpolation (also referred to as “inpainting”) denotes the methodology whereby missing parts of damaged images are filled in, based on the information obtained from the intact part of the image and a-priori assumptions made on the missing image structures, cf. Figure 6. For an introduction to image inpainting see [17]. Virtual image restoration is an important challenge in our modern computerized society: From the reconstruction of crucial information in satellite images of our earth to the renovation of digital photographs and ancient artwork, virtual image restoration is ubiquitous. Considering this huge – but by no means complete – amount of image processing applications and the fact that there are still problems in this area which have not been completely and satisfactorily solved, it is not surprising that this is a very active and broad field of research. From mathematicians, to engineers and computer scientists, a large group of people have been and are still working in this area.

A digital image - a mathematical object? In order to appreciate the following theory and the image processing applications, we first need to understand what a digital image really is. Roughly speaking a digital image is obtained from an analogue image (representing the continuous world) by sampling and quantization. Basically this means that the digital camera superimposes a regular grid on an analogue image and assigns a value, e.g., the mean brightness in this field, to each grid element. In the terminology of digital images these grid elements are called pixels. The image content is then described by grey values or colour values prescribed in each pixel. The grey values are scalar values ranging between 0 (black) and 255 (white). The colour values are vector values, e.g., (r, g, b), where each channel r, g and b represents the red, green, and blue component of the colour and ranges, as the grey values, from 0 to 255.

The mathematical representation of a digital image is a so-called image function u defined on a two dimensional (in general rectangular) image domain, the grid. Indeed, in some applications, images are three dimensional (e.g. videos, 3D medical imaging) or even four dimensional (involving three spatial dimensions and time) objects, but for simplicity we focus on the two dimensional case for the following conceptual presentation. The image function is either scalar valued in the case of a grey value image, or vector valued in the case of a colour image. Here the function value $u(x, y)$ denotes the grey value, i.e., colourvalue, of the image in the pixel (x, y) of the image domain. Figure 2 visualizes the connection between the digital image and its image function.

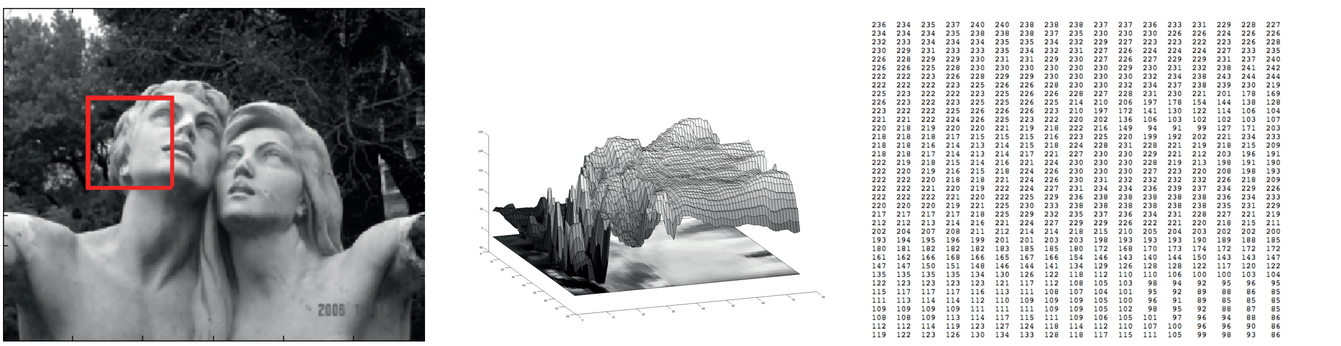

Figure 2: Digital image versus image function: On the very left a grey value photograph; in the middle the image function within a small selection of the digital photograph is shown where the grey value $u(x, y)$ is plotted as the height over the (x, y) - plane; on the very right the grey values for a small detail of the digital photograph are displayed in matrix form.

Typical sizes of digital images range from 2000×2000 pixels in images taken with a simple digital camera, to 10000×10000 pixels in images taken with high-resolution cameras used by professional photographers. The size of images in medical imaging applications depends on the task at hand. PET for example produces three dimensional image data, where a full-length body scan has a typical size of $175 \times 175 \times 500$ pixels.

Now, since the image function is a mathematical object we can treat it as such and apply mathematical operations to it. These mathematical operations are summarized by the term *image processing techniques*, and range from statistical methods, morphological operations, to solving a partial differential equation for the image function. We are especially interested in the last, i.e., PDE- and variational methods used in virtual restoration.

We have introduced a *digital* image as a sampled and quantised version of an *analog* (also called physical or real) image. The higher the resolution of a digital image, the closer it is to the analog image in the real-world. While digital image processing is indeed concerned with digital images the methods used are often motivated from considerations in the continuum, that is methods are formulated for the analog image. In this course we take up this mathematically more challenging and analytically more beautiful position, and let our image u be a continuous object defined on a rectangular domain $\Omega = (a, b) \times (c, d)$. Within this framework, there are many possibilities how images can be modelled, compare [9, Chapter 3].

From local to global image features - what is important for virtual image restoration? An important task in image processing is the process of filling in missing parts of damaged or occluded images based on the information obtained from the intact parts in the image. It is essentially a type of interpolation and we will refer to it as virtual image restoration or inpainting[1].

Let f represent some given image defined on an image domain Ω. Loosely speaking, the problem is to reconstruct the original image u in the (damaged) domain $D \subset \Omega$, called inpainting domain or a hole/gap (cf. Figure 6). Virtual image restoration methods can be roughly divided into two groups: 1) local inpainting, and 2) global inpainting methods. The main difference between these two classes lies in the type of image information used from the intact part of the image, as well as the different kind of inpainting processes with which this information is propagated into the missing domain.

A method is local if the information that to fill in D is only taken from of the boundary ∂D (or a small neighbourhood of D). In a local inpainting method the restored image u can be formalised as a solution of either a variational problem or a partial differential equation (PDE). The easiest example is harmonic inpainting, where the restored image u solves so-called Laplace equation

$$\begin{cases} \Delta u = 0 & \text{in } D \\ u = f & \text{on } \partial D. \end{cases}, \tag{1}$$

where $\Delta u = \frac{\partial^2 u(x,y,t)}{\partial x^2} + \frac{\partial^2 u(x,y,t)}{\partial y^2}$ the so-called Laplacian of u. In other words, we solve the Laplace equation inside the hole D with grey values from f from the surrounding of the hole ∂D (in mathematics these are called boundary conditions). As such u can be seen as the harmonic extension of f from ∂D into D. This can also be interpreted as the stationary (time t goes to infinity) state of the heat equation. Start with $u(x, y, t = 0) = f(x, y)$ a given image and evolve it along

$$\underbrace{\frac{\partial u(x,y,t)}{\partial t}}_{\text{rate of change of image greyvalues}} = \underbrace{\frac{\partial^2 u(x,y,t)}{\partial x^2} + \frac{\partial^2 u(x,y,t)}{\partial y^2}}_{\text{local averaging of greyvalues}}$$

Compare Figure 3. For image inpainting we solve the heat equation inside the black holes in the sculpture photograph in Figure 4.

Of course, any image structures such as image edges are not preserved by the harmonic extension (rather diffused into D). More sophisticated local inpainting methods have been proposed in the community during the last fifteen years that are able to propagate geometric image information such as object edges, their orientation and curvature. These approaches are mainly based on extensions of (1) to non-smooth variational problems and nonlinear (and often higher-order) PDEs, respectively. For different types of local inpainting methods we

[1] Various terminologies are used for image interpolation depending on application at hand. It is called also inpainting in the context of virtual image restoration [5] and disocclusion for the recovery of occluded objects [16]

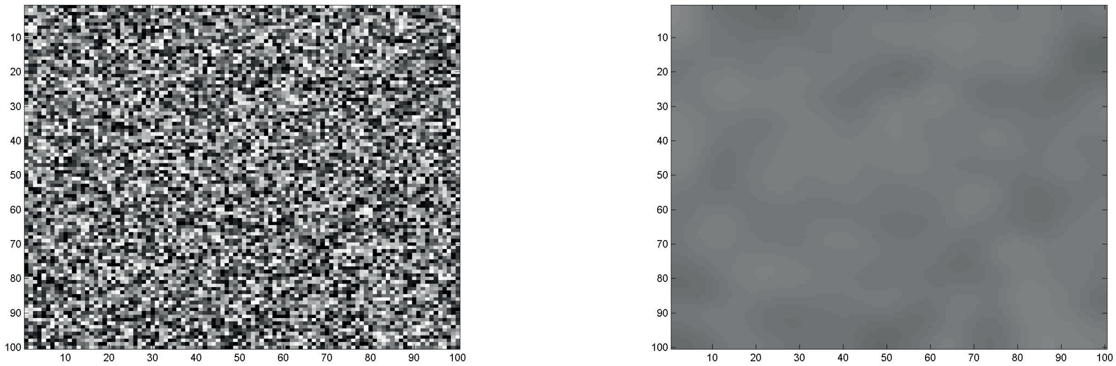

(a) $u(x, y, t = 0)$ (b) u(x,y,t=10)

Figure 3: An image $u(x, y, t)$ evolved along the heat equation.

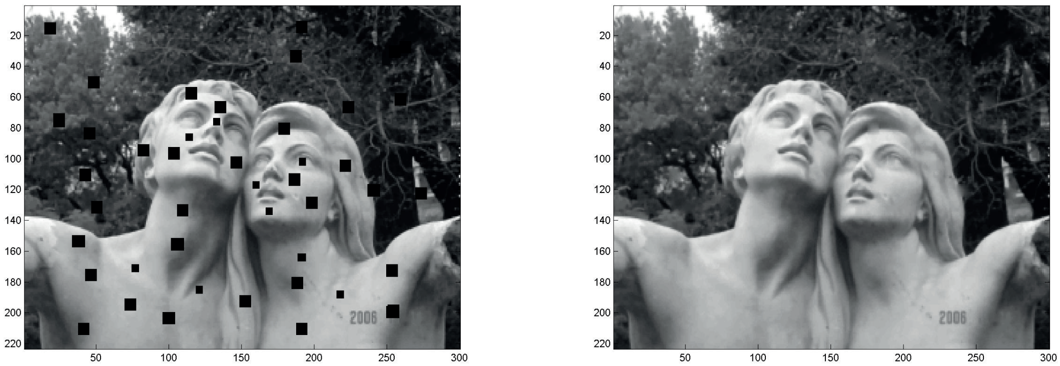

(a) $u(x, y, t = 0)$ (b) u(x,y,t=30)

Figure 4: The heat equation applied to image inpainting of the black holes in the left image.

refer the reader for instance to [5, 6, 15] (inpainting via transport, cf. Figure 5), [18, 7] (TV inpainting), [10] (curvature driven diffusion inpainting), [12] (Mumford- Shah based inpainting) and [16, 19] (Euler's elastica inpainting).

What all of these methods have in common is that they can only reproduce local image features (encoded in the value of the image function and its derivatives in a pixel) but are not able to pick up image patterns or texture which are non-local image features. Global (or non–local) inpainting methods[2] take into account all the information from the known part of the image, usually weighted by its distance and similarity (measured in a certain way) to a neighbourhood of the point that is to be filled in. Such methods usually work on image patches rather than on single image pixels. They are mathematically formalised as non-local variational problems and engineering–type discrete algorithms. This class of methods is very powerful, allowing to fill in structures and textures almost equally well. However, they still have some disadvantages. One major one is the high computational cost involved in their solution. For some of these methods another disadvantage is their dependence on the initial guess for the restored image in D. Local methods are sometimes more desirable, especially when the inpainting domain is relatively small. If D is large a local method can serve as a good initialisation for the global inpainting method. For more discussion on global methods the reader is referred to [8, 11, 1, 2].

[2]Global inpainting methods are also referred to as exemplar based inpainting, patch-based inpainting, copy-and-paste approaches, or texture-synthesis.

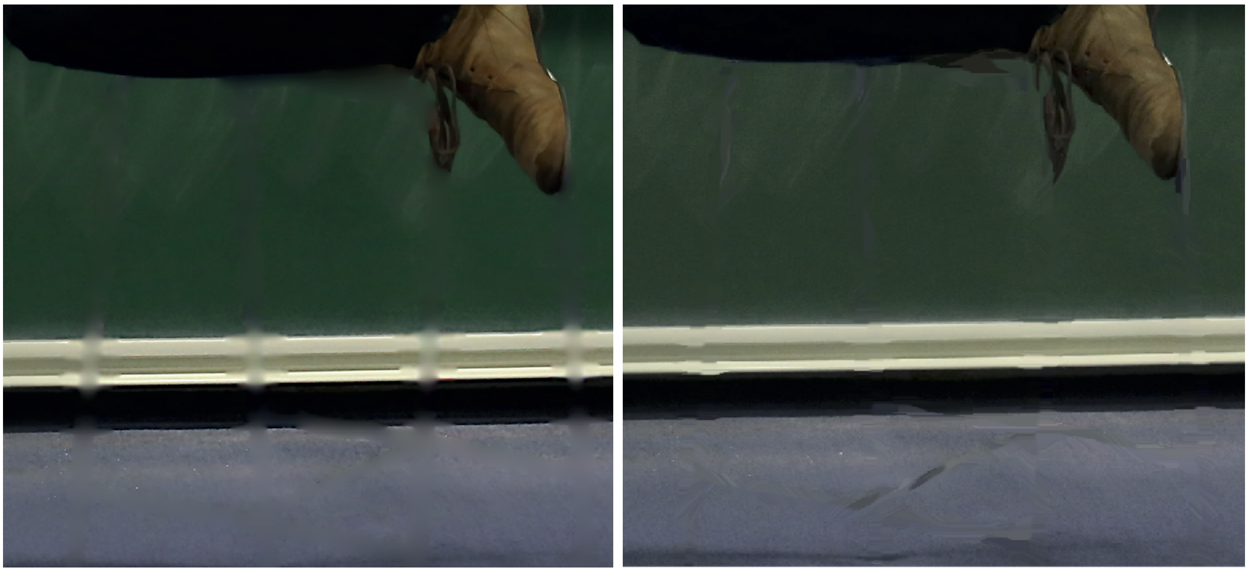

(a) Diffusion inpainting (b) Transport inpainting

Figure 5: Diffusion versus transport inpainting for the example in Figure 1.

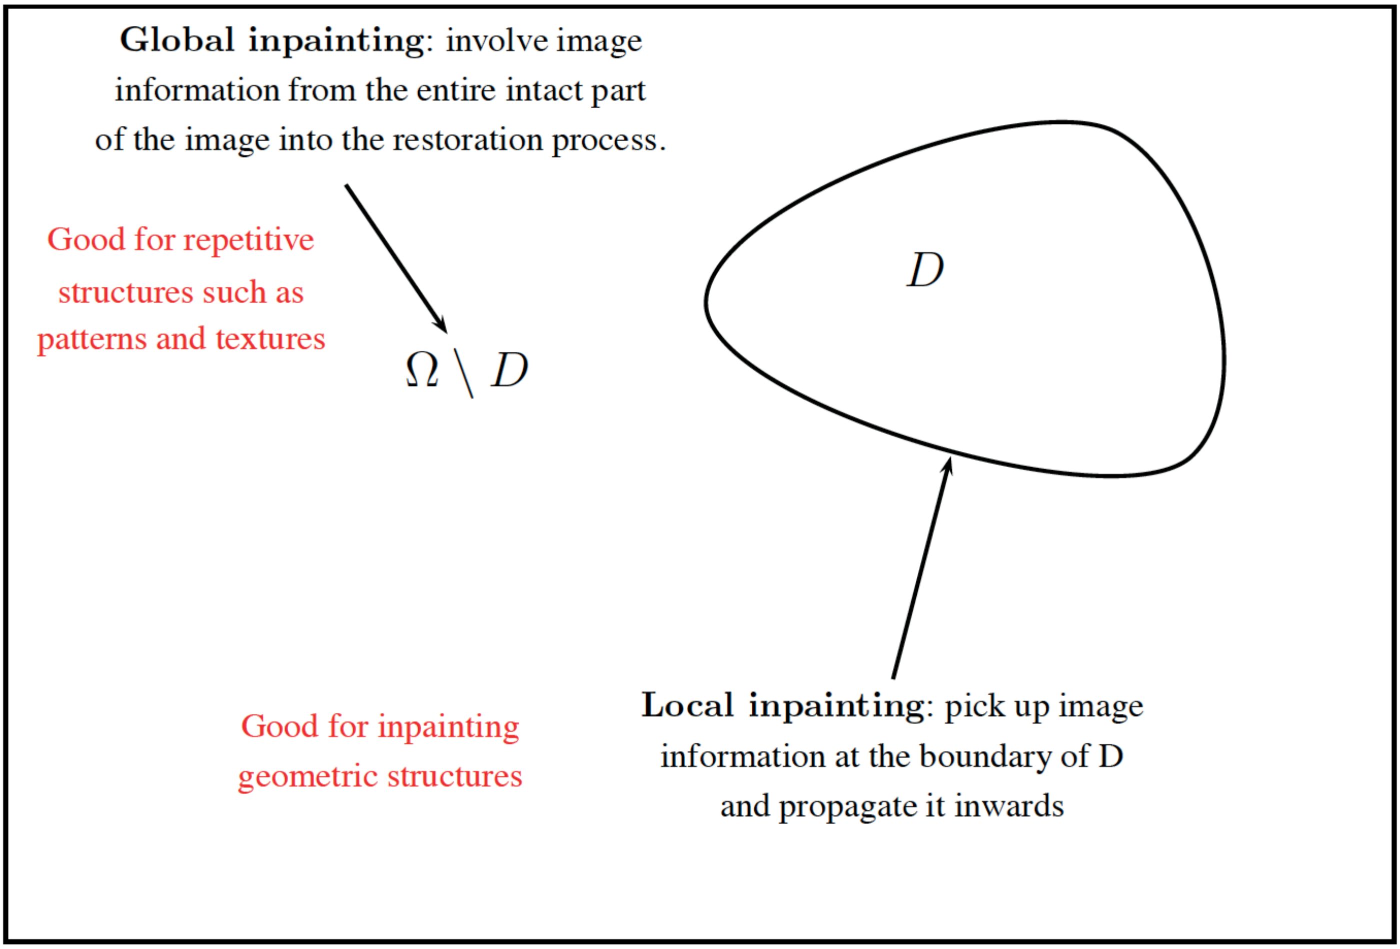

Figure 6: Virtual image restoration: based on the intact image information f inside $\Omega \setminus D$ one seeks for the inpainted image u that extends f into the inpainting domain D. The difference between local and global inpainting lies in its conceptually different method of recovering u from f.

Mathematical algorithms versus an amateur's attempt In August 2012 Cecilia Giménez, an eighty year old amateur artist from a small village near Zaragoza (Spain) gained fame by an attempt to restorate a wall painting in a local church. She produced the by now famous painting dubbed "Ecce Mono" (Behold the Monkey) when aiming to restore the wall painting "Ecce Homo" (Behold the Man) by the spanish painter Elías García Martínez, compare Figure 7.

Figure 7: "Ecce Homo" (left) and "Ecce Mono" (right).

Lets see what virtual image restoration methods make of this. In Figure 8 a local and a global inpainting result for the head of the Jesus figure are shown. For the local inpainting we used higher-order total variation inpainting [7] and for the global inpainting method a variational exemplar-based method with the L^1–norm as similarity measure between image patches [2]. Local inpainting is doing pretty well in recovering the main structures in the painting but smoothing out small-scale features and texture. Being initialised with the local inpainting result the global inpainting method performs reasonably well. Let the reader decide which restoration is more realistic: Cecilia's "Ecce mono" in Figure 7 or the mathematically formalised inpainting in Figure 8.

What is to be learned by this? Mathematical concepts such as nonlinear PDEs and variational calculus offer a beautiful and rich framework for formalising and solving real world problems in imaging. Of course, virtual image restoration can not (yet or never?) replace the human expertise. In fact, virtual image restoration algorithms have been very much influenced by the experience and guidelines from art restorers, aiming to frame what art restorers do in mathematical equations [5]. And this is not a one-way street but rather constitutes a fruitful two-way traffic. Virtual image restoration is able to produce digital templates for damaged art pieces[4] as well as offer a formalism for explaining what art restorers do.

Moreover: What I conceiled in this article is that these equations and variational problems are beautiful mathematical objects and their analysis and numerical treatment are challenging and very interesting tasks.

[3]Photo courtesy of Rob Hocking using the algorithm described in [1, 2]

[4]See `http://www.spiegel.de/international/zeitgeist/bombed-fresco-using-math-to-piece-together-a-lost-treasure-a-792781 html` for an example where mathematics helped the real restoration.

(a) Mask for restoration

(b) Initialisation of the restoration algorithm with random colours

(c) Restored image with local inpainting

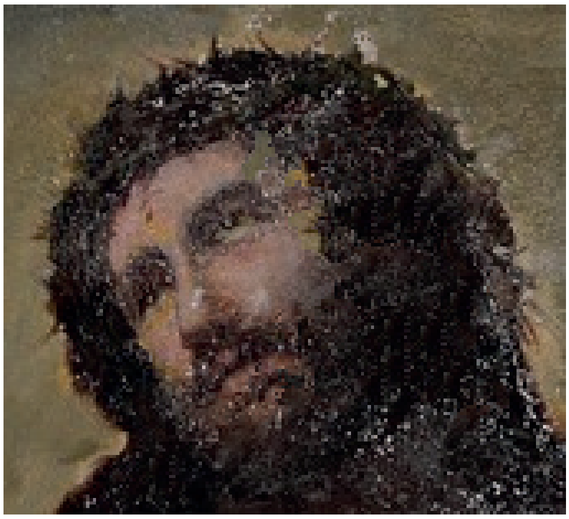

(d) Restored image with global inpainting[3]

Figure 8: Mathematical image restoration of "Ecce homo".

Some more applications PDEs as the one that can help Joana fly have applications also in other areas of sciences, e.g., medical imaging, forensics, geosciences, cf. Figures 9–11 for some examples. For more discussion on the role of mathematics in the big data revolution see [13].

References

[1] Pablo Arias, Vicent Caselles, and Guillermo Sapiro. A variational framework for non-local image inpainting. In *Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 345–358. Springer, 2009.

[2] Pablo Arias, Gabriele Facciolo, Vicent Caselles, and Guillermo Sapiro. A variational framework for exemplar-based image inpainting. *International journal of computer vision*, 93(3):319–347, 2011.

[3] G. Aubert and P. Kornprobst. *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*, volume 147. Springer Science + Business Media, 2006.

Figure 9: Closing gaps in data: enhancement of low quality fingerprints. On the left the raw fingerprint and on the right its enhanced version using an anisotropic diffusion equation, cf. [20].

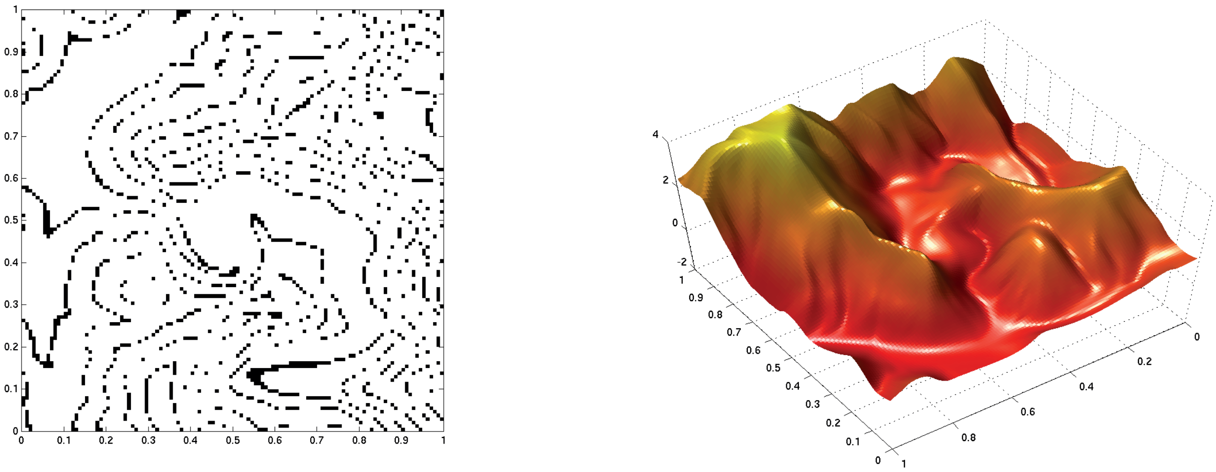

Figure 10: Reconstruction of digital elevation maps from sparse height data. Input is height information in the points given on the left. Output is a high resolution surface on the right. This has been interpolated with a nonlinear and anisotropic PDE, cf. [14].

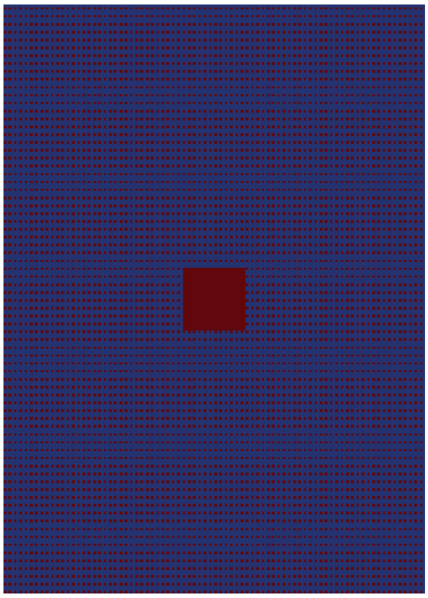

(a) Fourier sampling

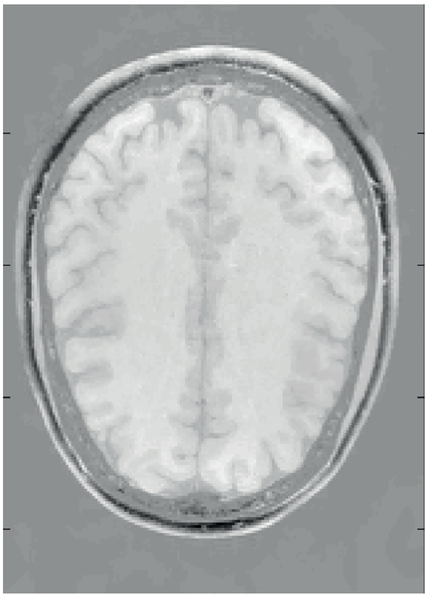

(b) Sparse reconstruction

Figure 11: Sparse reconstruction in parallel MRI. Measured datum $f = (f_1, \ldots, f_N)$ is a vector of Fourier samples measured by N coils. Measurements f and image u are related by the equation $f_j = (\mathcal{F}u \cdot c_j)_{|\Lambda}$, where c_j are so-called coil sensitivities (which are in general also unknown) and Λ is the sampling in Fourier space shown in the picture on the left. The image u shown on the right has been computed by a nonlinear interpolation and inversion approach using the so-called total variation, cf. [4] for more information. Thanks to Thorsten Hohage and the MPI for biophysical chemistry in Göttingen for the data and help with its processing.

[4] Martin Benning, Lynn Gladden, Daniel Holland, Carola-Bibiane Schönlieb, and Tuomo Valkonen. Phase reconstruction from velocity-encoded mri measurements–a survey of sparsity-promoting variational approaches. *Journal of Magnetic Resonance*, 238:26–43, 2014.

[5] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 417–424. ACM Press/Addison-Wesley Publishing Co., 2000.

[6] Folkmar Bornemann and Tom März. Fast image inpainting based on coherence transport. *Journal of Mathematical Imaging and Vision*, 28(3):259–278, 2007.

[7] Martin Burger, Lin He, and Carola-Bibiane Schönlieb. Cahn-hilliard inpainting and a generalization for grayvalue images. *SIAM Journal on Imaging Sciences*, 2(4):1129–1167, 2009.

[8] Frédéric Cao, Yann Gousseau, Simon Masnou, Patrick Pérez, et al. Geometrically guided exemplar-based inpainting. 2009.

[9] Tony Chan and Jianhong Shen. *Image processing and analysis: variational, PDE, wavelet, and stochastic methods*. Siam, 2005.

[10] Tony F Chan and Jianhong Shen. Nontexture inpainting by curvature-driven diffusions. *Journal of Visual Communication and Image Representation*, 12(4):436–449, 2001.

[11] Antonio Criminisi, Patrick Pérez, and Kentaro Toyama. Region filling and object removal by exemplar-based image inpainting. *Image Processing, IEEE Transactions on*, 13(9):1200–1212, 2004.

[12] Selim Esedoglu and Jianhong Shen. Digital inpainting based on the mumford–shah–euler image model. *European Journal of Applied Mathematics*, 13(04):353–370, 2002.

[13] Arieh Iserles and Carola-Bibiane Schönlieb. Mathematics of information. *Mathematics Today, IMA@50 special issue*, 50(1):53–58, 2013.

[14] Jan Lellmann, Jean-Michel Morel, and Carola-Bibiane Schönlieb. Anisotropic third-order regularization for sparse digital elevation models. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 161–173. Springer, 2013.

[15] Thomas März. Image inpainting based on coherence transport with adapted distance functions. *SIAM Journal on Imaging Sciences*, 4(4):981–1000, 2011.

[16] Simon Masnou and J-M Morel. Level lines based disocclusion. In *Image Processing, 1998. ICIP 98. Proceedings. 1998 International Conference on*, pages 259–263. IEEE, 1998.

[17] Carola-Bibiane Schönlieb. *Partial Differential Equation Methods for Image Inpainting*, volume 29. Cambridge University Press, 2015.

[18] Jianhong Shen and Tony F Chan. Mathematical models for local nontexture inpaintings. *SIAM Journal on Applied Mathematics*, 62(3):1019–1043, 2002.

[19] Jianhong Shen, Sung Ha Kang, and Tony F Chan. Euler's elastica and curvature-based inpainting. *SIAM Journal on Applied Mathematics*, 63(2):564–592, 2003.

[20] Duy Hoang Thai and Carsten Gottschlich. Global variational method for fingerprint segmentation by three-part decomposition. *IET Biometrics*, 5(2):120–130, 2016.