1
Introduction

This book is concerned with digital image processing techniques that use partial differential equations (PDEs) for the task of image inpainting. *Image inpainting* is an artistic term for virtual image restoration or image interpolation whereby missing or occluded parts of images are filled in based on information provided by the intact parts of the image. Computer graphic designers, artists and photographers have long used manual inpainting to digitally restore damaged paintings or manipulate photographs. Today, mathematicians apply powerful methods based on PDEs to automate this task. They operate in much the same way that trained restorers do: they propagate information from the structure around a hole into the hole to fill it in.

Virtual image restoration is an important challenge in our modern computerised society. From the reconstruction of crucial information in satellite images of the Earth to the renovation of digital photographs and ancient artwork, virtual image restoration is ubiquitous. The example in Figure 1.1 is entitled *Mathematical Analysis Can Make You Fly*, and it should give you a first impression of the idea of image inpainting with PDEs. The PDE model used for this example is called TV-H^{-1} *inpainting* and will be discussed in great detail in Section 5.3.

1.1. Digital Image Restoration in Modern Society

Digital images are one of the main sources of information today. The vast number of images and videos that exist in digital form 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.

Figure 1.1. 2013 EPSRC science photo competition winner: *Mathematical Analysis Can Make You Fly.* How is this possible? How can Joana – the woman in the picture and a master's 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 PDE that can be used for digital image inpainting. Inpainting is the process by which specified parts of an image are filled in based on the remaining part of the image. In this 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, the enhancement of medical images and the renovation of digital photographs and artwork to special effects in images and videos. As in this photograph, image inpainting is ubiquitious. (Image courtesy of Joana Grah, Kostas Papafitsoros and the author. Winner of the 2013 EPSRC Science Photo Competition in the 'People' category.)

In today's society, we encounter digital images on a daily basis: from everyday life, where analogue cameras have long been replaced by digital ones, to their professional use in medicine, earth sciences, the arts and security applications. In these contexts, we often have to deal with the processing of images, for example, 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. Medical imaging tools such as magnetic resonance imaging (MRI) usually produce noisy or incomplete image data. Satellite images of our Earth are often poorly resolved and blurred. Furthermore, digital image restoration is used in art preservation and restoration, where digital photographs are taken of historical artwork and are digitally restored and stored. As such, they can serve as templates for restorers and can be kept in a database for preservation. In sum, digital image restoration provides an effective tool to recover or complete lost image information. Keywords in this context are *image de-noising*, *image de-blurring*, *image decomposition*, *image inpainting* and *image synthesising*.

Considering the vast number of image restoration applications and problems in this area that have not been completely and satisfactorily resolved yet, it is not surprising that this is a very active and broad field of research. From engineers to computer scientists and mathematicians, a large group of people has been and is still working in this area.

1.2. What is a Digital Image?

To appreciate the following theory and the image inpainting applications, we first need to understand what a digital image is. Roughly speaking, a *digital image* is obtained from an analogue image (representing the continuous world) by sampling and quantisation. Similar to our eye, a digital camera superimposes a regular grid on an analogue image and assigns a value to each grid element. For example, in the case of a digital photograph, each grid element stores the mean brightness in the recorded field encoded in the photon counts. Thus, a digital image can be considered to be a sample of the continuous world.

In the terminology of digital images, the sampling points are grid elements called *pixels* (from *pict*ure *el*ement*s*). The image content is then described by the grey values or colour values prescribed in each pixel. The grey values are usually scalar values ranging between 0 (black) and 255 (white). Colour values are represented by vectors, most commonly (r,g,b), where each channel r,g and b represents the red, green and blue component of the colour, respectively, ranging again from 0 to 255. The mathematical representation of a digital image is a so-called image function u defined (for now) on a two-dimensional (in general, rectangular) grid. Indeed, in some applications, images are three-dimensional (e.g., videos and medical imaging) or even four-dimensional objects involving three spatial dimensions and time. However, in this book we concentrate on the two-dimensional case. Figure 1.2 illustrates the connection between a digital image and its image function for the case of a grey value image.

Typical sizes of digital images range from $2,000 \times 2,000$ pixels in images taken with a simple digital camera to $10,000 \times 10,000$ pixels in images taken

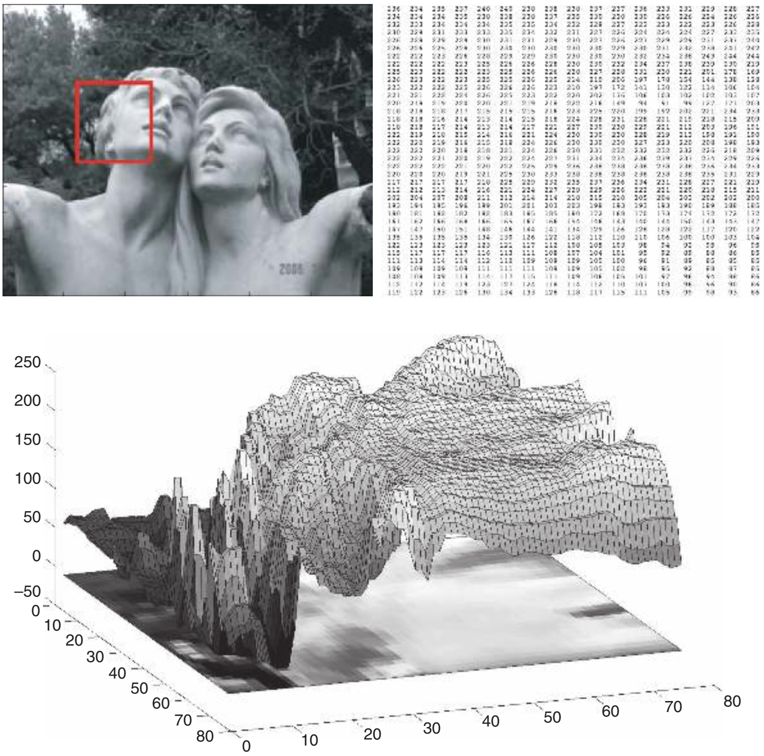

Figure 1.2. Digital image represented as an image function. In the top row on the left, a grey value photograph is shown; on the right, the grey values for a small detail of the digital photograph are displayed in matrix form. In the bottom row, the image function within a small selection of the digital photograph is shown where the grey value $u(x,y)$ is plotted as the height over the (x,y)-plane.

with high-resolution cameras used by professional photographers. The sizes of images in medical imaging applications depend on the task at hand. Positron emission tomography (PET), for example, produces three-dimensional (3D) image data, where a full-length body scan has a typical size of $200 \times 200 \times 500$ pixels.

Since the image function is a mathematical object, we can treat it as such and apply mathematical operations to it. These mathematical operations are summarised by the term *image processing techniques* and range from statistical methods and morphological operations to solving a PDE for the image function (cf. Section 1.1). We are especially interested in the latter, that is, PDE methods for image inpainting.

Remark 1.2.1 Digital Images and Continuous Models Note that although digital images are discrete finite-dimensional objects, most image processing methods discussed in this book are modelled and analysed in infinite-dimensional function space rather than $\mathbb{R}^d$. This reflects our aim of finding and reconstructing from a finite-dimensional sample (the digital image) an image from the physical, continuous world (the analogue image). On a microscopic level, a line in a digital image is only a collection of pixels with common characteristics that allow us to identify it as a continuous line on a macroscopic level where pixels are invisible to us. Refer to Chapter 3 for more discussion on the perception of geometrical objects from point clouds. Moreover, performing the modelling and analysis of image reconstruction in function space allows us to specify finer properties of a desired reconstruction. For instance, it is easy to talk about discontinuities for a function that lives in the continuum, whereas the definition of discontinuities for a finite-dimensional function is unclear. This is similar to analogous situations in many other areas of mathematics, for example, in kinetic theory, where particles are replaced by densities, or in statistics, where finite-dimensional analysis is replaced by infinite-dimensional stochastic analysis.

1.3. Image Inpainting

An important task in image processing is the process of filling in missing parts of a damaged image based on the information obtained from the intact part of the image. It is a type of interpolation called *inpainting*.

Let g represent some given image defined on an image domain Ω. The problem is to reconstruct the original image u in the (damaged) domain $D \subset \Omega$, called the *inpainting domain* or *hole/gap* (Figure 1.3).

The term *inpainting* was invented by art restoration workers [EM76, Wal85, WG04] and first appeared in the framework of digital image restoration in the work of Bertalmio et al. [BSCB00]. Therein the authors designed a discrete inpainting model motivated by a PDE which intends to imitate the restoration work of museum artists. Their method will be explained in more detail in Section 6.1.

To give you a first ‘gusto’ for real-world problems where digital inpainting might be useful, let us evoke one example from the recent history of art restoration. In August 2012, Cecilia Giménez, an eighty-year-old amateur artist from a small village near Zaragoza, Spain, gained fame by an attempt to restore a wall painting in a local church. She produced the now-famous painting dubbed ‘Ecce Mono’ (‘Behold the Monkey’) when aiming to restore

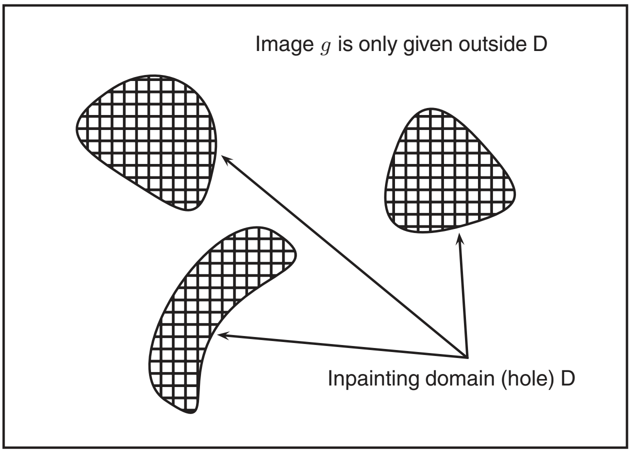

Figure 1.3. The inpainting task.

Figure 1.4. *Ecce Homo* (left) and *Ecce Mono* (right).

the wall painting *Ecce Homo* (*Behold the Man*) by Spanish painter Elías García Martínez (Figure 1.4).

Although Cecilia's restoration clearly follows a conceptual approach, the result seems to us rather suboptimal. This shows the complexity of the image interpolation problem and the challenge of defining and formalising rules which produce a visually appealing restoration result. To see what virtual image restoration methods make of *Ecce Homo*, continue reading (see Figure 2.4) or see [Sch14].

Applications Applications of digital image inpainting are numerous. From automatic scratch removal in old photographs and films [BSCB00, SC02, KMFR95b] to digital restoration of ancient paintings for conservation purposes [BFM$^+$08], text erasing such as the removal of dates, subtitles or publicity from a photograph [BSCB00, BBC$^+$01, SC02, CS01b], special effects such as object removal [BSCB00, CS01b], disocclusion [NMS93, MM98, Mas98, Mas02], spatial/temporal zooming and super-resolution [BBC$^+$01, SC02, Mal00, MG01a, TYJW01], error concealment [WZ98], lossy perceptual image coding [SC02], removal of the laser dazzling effect [CCBT03] and sinogram inpainting in X-ray imaging [GZY$^+$06], to name only a few.

The Beginnings of Digital Image Inpainting The history of digital image inpainting has its beginning in the works of engineers and computer scientists. Their methods were based on statistical and algorithmic approaches in the context of image interpolation [AKR97, KMFR95a, KMFR95b], image replacement [IP97, WL00], error concealment [JCL94, KS93] and image coding [Cas96, LFL96, RF95]. In [KMFR95b], for example, the authors present a method for video restoration. Their algorithm uses intact information from earlier and later frames to restore the current frame and is therefore not applicable to still images. In interpolation approaches for 'perceptually motivated' image coding [Cas96, LFL96, RF95], the underlying image model is based on the concept of 'raw primal sketch' [Mar82]. More precisely, this method assumes that the image consists of mainly homogeneous regions separated by discontinuities, that is, edges. The coded information then just consists of the geometrical structure of the discontinuities and the amplitudes at the edges. Some of these early coding techniques already used PDEs for this task (see, e.g., [Car88, Cas96, CT94]).

Initiated by pioneering works [NMS93, MM98, CMS98a, BSCB00 and SC02], the mathematics community got involved in image restoration using PDEs and variational methods for this task. Their approach and some of their methods shall be honoured in this book.

2

Overview of Mathematical Inpainting Methods

Digital inpainting methods are being designed with the desire for an *automated* and *visually convincing* interpolation of images. In this chapter we give an overview of approaches and trends in digital image inpainting and provide a preview of our discussion in Chapters 4 through 7. Before we start with this, let us raise our consciousness about the challenges and hurdles we might face in the design of inpainting problems.

The first immediate issue of image inpainting is, of course, that we do not know the truth but can only guess. We can make an educated guess, but still it will never be more than a guess. This is so because once something is lost, it is lost, and without additional knowledge (based on the context, e.g., historical facts), the problem of recovering this loss is an ambiguous one. Just look at Figure 2.1, and I ask you: is it a black stripe behind a grey stripe or a grey stripe behind a black stripe? Thus, the challenge of image inpainting is that *the answer to the problem might not be unique*. We will discuss this and strategies to make 'good' guesses based on the way our perception works in Chapter 3.

When inspecting different inpainting methods in the course of this book, you should be aware of the fact that *mathematical inpainting methods are designed for inpainting the image completely automatically*, that is, without intervention (supervision) by the user. Hence, the art of designing efficient and qualitatively high inpainting methods is really the skill of modelling the mechanisms that influence what the human brain can usually do in an instant. At present, we are still far away from a fair competition with the human brain. Digital inpainting methods are currently not (will never be?) as smart as our brain. In particular, *no all-round inpainting model exists* that can solve a variety of inpainting problems with sufficient quality. One of the main shortcomings of inpainting methods is their inability to realistically reconstruct both structure and texture simultaneously (see Section 2.2).

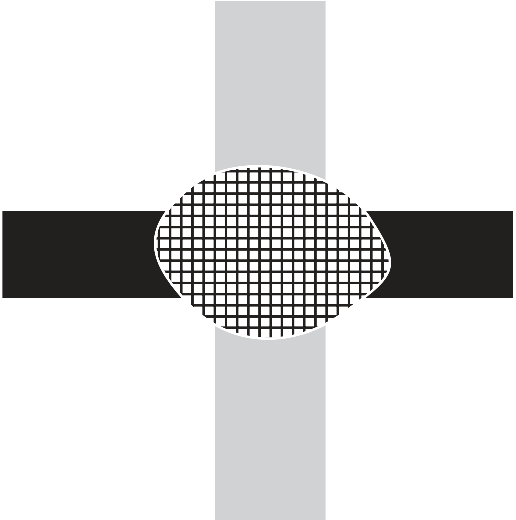

Figure 2.1. Non-uniqueness of image inpainting. The image got lost inside the hatched domain. Should the reconstructed image show a black stripe behind a grey stripe or a grey stripe behind a black stripe?

Finally, let us also emphasise that the *difficulty of an image inpainting problem increases radically with the size of the damage* (the inpainting domain) in the image. The inpainting of large, connected gaps is more challenging than the inpainting of a collection of small holes – even if both their total areas cover the same amount of pixels. Intuitively, this is so because the larger the inpainting domain, the larger the distance across which image contents have to be interpolated and the less the known image information suggested by the intact image weighs for inpainting the inner parts of the inpainting domain. Lower-order inpainting methods – as we will discuss in Chapter 4 – may be sufficient for the inpainting of small holes because the interpolation model can act more locally there. For large gaps, more sophisticated inpainting models must be consulted, in particular, non-linear partial differential equations (PDEs) of higher order, as discussed in Chapter 5, and transport inpainting, as discussed in Chapter 6, as well as exemplar-based inpainting methods. One could loosely say that the sophistication of an interpolation method must increase with the size of the holes. Figure 2.2 shows two examples of harmonic inpainting, that is, interpolation from the grey values from the boundary of the inpainting domain ∂D by harmonic extension, as discussed in Section 4.2. This is one of the most basic inpainting methods. It constitutes a second-order linear elliptic PDE whose properties are

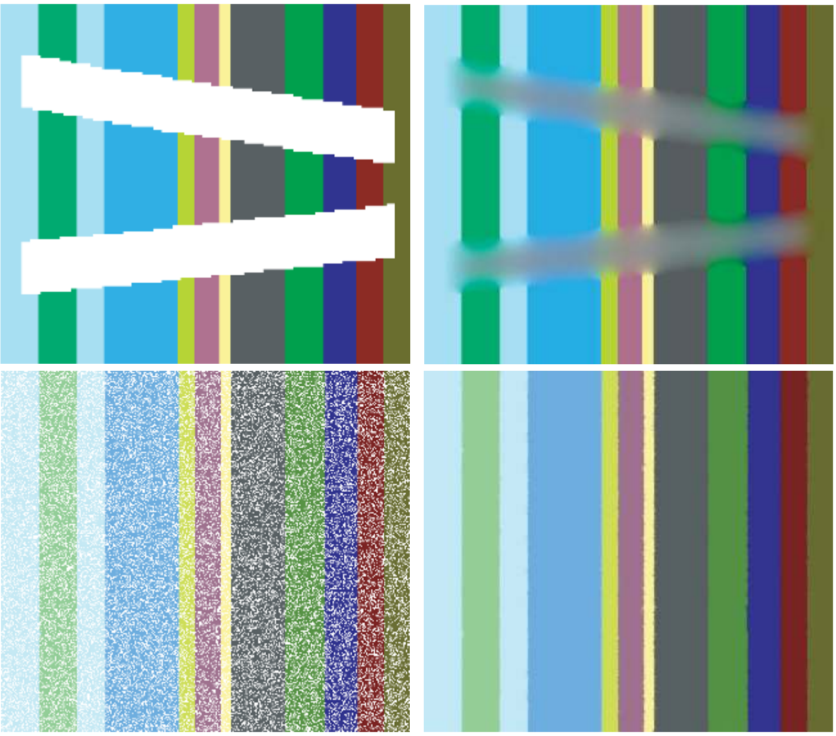

Figure 2.2. Large versus small gap image restoration using harmonic inpainting (see Section 4.2). On the left, the damaged images with inpainting domain marked in white. On the right, the corresponding restored images with harmonic inpainting. The percentage of damaged pixels is 30 per cent in both examples.

very well understood [Eva98] and for whose numerical solution very efficient solvers exist, for instance, iterative methods, spectral solvers and many more (cf. [Ise09]). This method performs very well for the inpainting of many small points of damage. If the inpainting domain, however, is one large occlusion, it fails to return a visually suggestive answer. Note that the percentage of damaged pixels in both examples is exactly the same. In the presence of large gaps, more sophisticated inpainting methods, such as inpainting by coherence transport (see Section 6.2), have to be consulted (see Figure 2.3).

2.1. Variational and PDE Methods

A *variational* approach for image inpainting is a method that computes the inpainted image as a minimiser of an objective functional. The latter is the