

Discrete Geometric Homogenisation and Inverse Homogenisation of an Elliptic Operator

Thesis by
Roger David Donaldson

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2008
(Defended May 2, 2008)

© 2008

Roger David Donaldson

All Rights Reserved

Abstract

We show how to parameterise a homogenised conductivity in \mathbb{R}^2 by a scalar function $s(x)$, despite the fact that the conductivity parameter in the related up-scaled elliptic operator is typically tensor valued. Ellipticity of the operator is equivalent to strict convexity of $s(x)$, and with consideration to mesh connectivity, this equivalence extends to discrete parameterisations over triangulated domains. We apply the parameterisation in three contexts: (i) sampling $s(x)$ produces a family of stiffness matrices representing the elliptic operator over a hierarchy of scales; (ii) the curvature of $s(x)$ directs the construction of meshes well-adapted to the anisotropy of the operator, improving the conditioning of the stiffness matrix and interpolation properties of the mesh; and (iii) using electric impedance tomography to reconstruct $s(x)$ recovers the up-scaled conductivity, which while anisotropic, is unique. Extensions of the parameterisation to \mathbb{R}^3 are introduced.

Preface

A common undertaking in science is to discover what properties emerge from systems as we observe them over a range of scales. The interaction between genetics and environmental change gives us evolution just as gravitational interactions between small astronomical bodies builds solar systems. This thesis is concerned with the properties that emerge when we view heterogeneous materials on a variety of scales.

It is an amazing pattern in science that we can do experiments and form theories at particular scales as if these scales were isolated. Knowing the makeup of the nuclei in a fluid phase does not help us form the Navier-Stokes equations, nor does modelling the stream of radiation from the Sun help us build a theory of planetary motion. This pattern is probably related to the universality of the stationary action principle. At a great many scales, following the flow of energy in a system provides good predictions of experiment, and energy by and large does not seem to leap between disparate scales.

This said, it also seems that interesting science happens when this assumption of scale separation is violated. Looking at black holes questions the consistency between theories of gravitation and quantum mechanics, and understanding how molecules form bonds with each other ought to explain why cracks form easily in some materials while other materials can be greatly deformed before any catastrophic failure. Even water, a phase composed of a single type of molecule, transfers energy between fine eddies and its average flow in the relatively mild circumstances under which it flows turbulently.

The general approach in engineering and computation for accounting for the interactions between scales is to either argue that their affect is insignificant, or by encapsulating the small scales in a model. For example, this thesis studies conductivity, a material parameter relating the amount of flux, either of mass or energy, to an applied potential. A conductivity model is typically an encapsulation of a microscopic flow, such as electrons through a wire, or molecular vibrations through an atomic lattice. While conductivity is often constant

in homogeneous media, we can further encapsulate it by averaging it over compositions of materials, heterogeneous at fine scales. This further encapsulation is *homogenisation*, and the first part of this thesis considers which features of conductivity should be preserved by homogenisation.

The second part of this thesis concerns a question dual to building a consistent multi-scale theory. The dual problem is a measurement problem, and it asks what we can determine about a system given that the scale at which we perform our measurements is macroscopic relative to the scale of the underlying physics. Again, we study conductivity, and based on flux and potential measurements at the boundary of a body, we ask what we can determine about its conductivity throughout its interior.

Many people have in some way contributed to this work, either by their early influence in my life, or by the more recent contributions of hot meals and moral support. I will thank the great many of you in person over the months and years ahead, and shall confine my thanks here to the earliest and the latest players. Our lives tend to begin with our parents, and mine is no different: the influence of my parents, Judith Ann and Robert W. Donaldson, can be seen directly in this thesis for their conviction that I be educated both in math and in writing. Equally important, I thank them for the freedom they offered me to work on my own projects from a young age. This freedom has been mirrored by a great many of my educators, including my current advisors Hooman Owhadi and Mathieu Desbrun. The idea to work on electric impedance tomography, the starting point of this work, was one that I brought to them, and they never lacked for interest in its development.

Roger Donaldson

Pasadena, California

April 2008

Table of Contents

Abstract	iii
Preface	iv
Table of Contents	vi
List of Tables	viii
List of Figures	ix
1 Introduction: Homogenising Conductivity	1
1.1 A multi-scale elliptic operator	2
1.2 Homogenising scalar conductivities	7
1.3 Anisotropic conductivities in \mathbb{R}^2	12
1.4 A simple example: the laminated domain	15
2 The Forward Problem	20
2.1 Discretizing the conductivity	21
2.2 Multi-scale elliptic operators	30
2.2.1 Homogenised operator convergence	30
2.2.2 Operator convergence: numerical experiment	35
2.3 Well-adapted triangulations	44
2.3.1 Mesh connectivity by weighted Delaunay triangulations	46
2.3.2 Optimising vertex locations	50
2.3.3 Numerical results	54
3 The Inverse Problem	61
3.1 What aspect of conductivity can we measure?	63

3.2	Theory	66
3.3	Reconstructing $s(x)$	71
3.3.1	Least-squares method	72
3.3.2	Regularisation	75
3.3.3	Boundary values for s	79
3.3.4	Enforcing convexity	81
3.3.5	Numerical results	81
3.3.6	Evaluating algorithms for divergence-free EIT	86
3.4	Unsuccessful reconstruction algorithms	90
3.4.1	Resistor network interpretation	90
3.4.2	Down-scaling the conductivity	93
3.4.3	Calderón iteration	98
3.4.4	Variational constraints	99
4	Extensions	103
4.1	Parameterising conductivities in \mathbb{R}^3	104
4.2	Discrete parameterisations of conductivities	106
	Bibliography	109
	Appendices	
A	Software credits	117
B	Notation	119

List of Tables

1.1	Linear constitutive laws for potential-driven flux	8
2.1	Convergence of an approximated operator	41
2.2	Improvement of interpolation by mesh adaptation	57
2.3	Improvement of matrix conditioning by mesh adaptation	58
2.4	Comparison of eigenvalues of the stiffness matrix	59
2.5	Leading constants in quality bounds	60

List of Figures

1.1	Homogenising the conductivity of a laminated domain	3
1.2	Relationships between the different representations for conductivity	14
1.3	The harmonic coordinate $F_2(x, y)$ for the horizontally laminated domain	16
2.1	Notation for computing q_{ij}	23
2.2	Notation for confirming divergence-free property	25
2.3	Positivity of Q and convexity of interpolated s	29
2.4	Edge flips altering convexity	30
2.5	Laminated domain used in computations	38
2.6	Fine-scale $s(x)$ for the laminated domain	39
2.7	Coarse-scale $s(x)$ for the laminated domain	39
2.8	Comparison of homogenisation solutions at coarse and fine scales	40
2.9	Convergence of an approximated operator	41
2.10	Reference solution detail	42
2.11	Solution error showing fine-scale error	43
2.12	Comparison of an isotropic and anisotropic mesh	45
2.13	Convex hull construction of a Delaunay triangulation	46
2.14	Geometry of a weighted Delaunay triangulation	48
2.15	L_2 interpolation quality of adapted meshes	56
2.16	h_1 interpolation quality of adapted meshes	56
2.17	Matrix conditioning of adapted meshes	57
2.18	Growth in the number of vertices with boundary resolution	58
3.1	Voronoi region associated to a vertex	77
3.2	Stencil for approximating Q	78
3.3	Ghost vertices and triangles at the mesh boundary	79

3.4	Counting the edges and vertices of a triangle mesh	81
3.5	A sample isotropic conductivity for testing reconstruction	82
3.6	Reconstruction of an isotropic conductivity	84
3.7	A sample conductivity, anisotropic at the coarse scale	84
3.8	Anisotropic reconstruction I	85
3.9	Anisotropic reconstruction II	86
3.10	Output of the harmonic coordinate iteration	95
3.11	Instability of harmonic coordinate iteration	96
3.12	Variational constraints for EIT	101

Chapter 1

Introduction: Homogenising Conductivity

Computer simulation and qualitative understanding of a real-world material often requires only coarse-scale solutions despite the material's fine-scale variability. Solution fields tend to vary at scales similar to those of the material's properties, so if we desire only coarse-scale solutions, we need only coarse-scale representations of the materials. The trouble is that material properties tend to appear as parameters in the partial differential equations (PDE) modelling the system. Averaging material properties is thus often the same problem as coarsening the scale at which the differential operators act. Compressing a differential operator with the intent of averaging the effects of its rapidly varying parameters is called homogenisation.

In this thesis, we study the homogenisation of conductivity. Conductivity is the capacity of a material to conduct flow under an applied potential. The flow may be electric current through a wire or hydrogen molecules through a fuel cell or reindeer through a thick forest, but the point is that this capacity for flow may change rapidly from one region to another. When homogenised, this fine-scale variation can have a strange effect: the direction of the flow may not be the same as the gradient of the potential. This directional dependence means that conductivity is not a scalar-valued function of space, but is tensor valued. The magic here is that if we homogenise the conductivity, symmetries emerge that reduce its dimensionality. In \mathbb{R}^2 , we can parameterise the homogenised conductivity again by a scalar function, and with this new parameterisation, taking averages of conductivity is as easy as coarsely-sampling the scalar parameterisation, just as we would do to a solution field.

1.1 A multi-scale elliptic operator

We are concerned with the variable coefficient linear elliptic problem

$$\begin{aligned} -\operatorname{div}(Q\nabla u) &= f, & x \in \Omega, \\ u &= 0, & x \in \partial\Omega. \end{aligned} \tag{1.1}$$

$\Omega \subset \mathbb{R}^d$ is a simply connected closed and bounded domain with boundary $\partial\Omega$, f is a source term, u is a potential field, and Q is a tensor representing a material property such as conductivity, or heat or chemical diffusivity. The computational problem is to resolve a potential field u at a coarse scale quickly and with accuracy, despite the fact that conductivity Q may vary at fine spatial scales. Hence, we are looking for a reduction in the dimensionality of solutions of problem (1.1) without undue compromise in accuracy. The main result of homogenisation is that while simple spatial averages of a given field u give good approximations to its fine-scale variation, parameterising the differential operator by naïve averages of Q does not provide coarse-scale solutions close to the fine-scale solutions.

To illustrate the care that must be exercised in averaging material properties, consider an observer measuring the conductivity of the material in Figure 1.1. The material is laminated. Although locally isotropic, it is composed of alternating layers having large and small conductivities. A very small observer, one small enough to reside on a single layer, sees no preferred direction in the material. Unless the small observer is lucky enough to be sitting on the boundary between two layers, the observer will see only black or white material in his immediate surroundings. With no preferred direction, the small observer deduces that the material is isotropic.

A large observer sees exactly the opposite. Seeing many boundaries between black and white laminates, the large observer knows that the conductivity parallel to the laminates will be larger than the conductivity perpendicular to them, even if the conductivity is constant and isotropic within each layer. At a coarse scale, an anisotropic conductivity is the best model for the material, even though the conductivity is scalar at its finest scale. Taking simple averages of the fine-scale conductivities does not predict this anisotropy.

There are at least three ways to take averages of conductivities, and hence of elliptic




Figure 1.1: A scalar conductivity alternating at a fine scale acts as an anisotropic conductivity at a coarse scale. The lengths of the axes represent the effective homogenised conductivities in each principal direction.

operators: the representative volume element (RVE) technique, asymptotic homogenisation, and metric-based up-scaling. RVE takes a small sample of the domain, larger than the scale of the material's inhomogeneities, and takes measurements of the potential gradient ∇u , and the resulting flux $\mathbf{j} = Q\nabla u$ averaged over the entire sample. The RVE experimenter is ignorant of any variations in the material smaller than the size of the sample. While physically intuitive, this method assumes a particular sample scale, and does not indicate how to correctly change scale, composing larger-scale samples of the material from smaller-scale ones. Mathematically, there is no limit taken, and so there is no notion of the RVE sample converging to a particular value. In fact, since the scale at which averaging occurs is fixed and finite (not infinitesimal), some authors do not classify this as a homogenisation method. The RVE method appears in particular in the elasticity literature, see [47], for example, and the references therein.

Asymptotic homogenisation, perhaps the most widely known approach in the mathematics community, assumes that the material has periodic structure, the fine-scale period ϵ being much smaller than the diameter of the domain. Asymptotic homogenisation expands the exact solution as an asymptotic series, identifying the material's *cell problem*, a PDE parameterised by the exact fine-scale conductivity, but which is solved only over a single fine-scale period in the domain. The conductivity parameterising the homogenised operator is given in terms of solutions to the cell problem for $\epsilon \rightarrow 0$. Although this method does take advantage of periodic structure to reduce the size of the calculations required, one of its shortcomings is requiring a separation between the macro-scale (the size of the domain

itself) and the micro-scale (the period of the structure). The calculation is described in detail in [48, Chapter 1], for example.

The work in this thesis is based on the third method, metric-based up-scaling, a name coined by Owhadi and Zhang in their papers [65, 66, 88] which followed earlier work of Babuška et al [20]. Like the other methods, metric-based up-scaling reduces the effective dimensionality of the numerical PDE that must be solved, but does so without requiring a separation of scale. Without separation of scale, and without periodicity or other patterns in the material, full knowledge of the fine-scale conductivity must be accounted for, and the method begins by solving d high-dimensional problems, where the domain $\Omega \subset \mathbb{R}^d$. The solutions to these d problems are a change of variables $F : \Omega \rightarrow \Omega$, and it is in this new coordinate system that we perform up-scaling. The new coordinates F account for the fine-scale variability in solutions to the original problem, such that solutions in the new coordinate system are remarkably smooth.

Numerically, computing these so-called *harmonic* coordinates F requires discretizing the problem at a scale smaller than the finest scale of the conductivity. Thus, this method offers no advantage over simply solving the problem at a fine resolution unless the operator is to be re-used. This happens, for example, when solutions are required for several source terms f , or when the f is replaced by a time-dependent term, such as a time-derivative of u , as is the case in the diffusion equation. We study this method here because its independence on scale — barring the resolution of the d initial solves — is precisely what allows us to use *linear* operations to rescale the homogenised operator to the desired resolution of the solution fields, despite the fact that the underlying conductivity may have a continuum of scales.

Briefly, a conductivity obtained by metric-based up-scaling is a tensor field with three properties: it is *positive-definite*, *symmetric* and its columns are *divergence-free*. The first two conditions are conventional, and the third is a consequence of metric-based up-scaling. These three constraints reduce the dimensionality of the conductivity, and for dimension $d = 2$, the up-scaled conductivity may be represented as scalar convex function, $s : \mathbb{R}^2 \rightarrow \mathbb{R}$: roughly speaking, the curvatures of $s(x)$ represent the directions and strength of the

anisotropy of the conductivity. In \mathbb{R}^2 , we write

$$Q = F_*\sigma = \begin{pmatrix} \partial_{yy}s & -\partial_{xy}s \\ -\partial_{xy}s & \partial_{xx}s \end{pmatrix}, \quad (1.2)$$

Q is the push-forward of σ by the harmonic coordinates. Since solutions u pushed forward by F are smooth, Q can be up-scaled while preserving the global effects inherited by fine-scale features of σ , such as those of a laminated domain. The bulk of this thesis, where we consider conductivities in \mathbb{R}^2 , regards computations consequential to this parameterisation. Development of this parameterisation and the recent work by Owhadi and Zhang will be introduced in further detail in Section 1.2.

Since conductivity is a parameter specifying the elliptic differential operator, up-scaling conductivity is equivalent to up-scaling an operator. The original treatment of metric-based up-scaling specifies a discretised operator from the original conductivity as a non-linear operation. An important development by the use of our parameterisation $s(x)$ is that up-scaling — selecting a particular scale for computation — can now be accomplished as the linear operation of a function approximation, such as interpolation. Once $s(x)$ is computed on the finest scale, its coarser scale approximations give a family of consistent operators that can be quickly computed.

In this thesis, we identify applications of the homogenised operator to both forward and inverse elliptic problems. The forward problem is the basic problem (1.1), where we seek a potential field u for given a conductivity Q and source f . Of course, variations on (1.1) with different boundary conditions and source terms are common — the focus here is on the operator $\operatorname{div}(Q\nabla\cdot)$. After homogenisation, the operator in the forward problem is discretized and represented as a scalar function $s(x)$ at the finest scale. We then re-sample $s(x)$ to form coarse scale representations of the operator. Discretization of $s(x)$ in the form (1.2) requires some care. We shall show that while Q is represented pointwise by second derivatives of $s(x)$, second derivatives of $s(x)$ can be understood in the sense of distributions, so we need only at least a piecewise linear interpolation of $s(x)$ to represent discretizations of operator $\operatorname{div}(Q\nabla\cdot)$.

This representation for $s(x)$ advantages forward solution methods in at least two ways. First, multi-scale methods for solving elliptic PDE require a hierarchy of discrete operators.

For example, the multi-grid method works by resolving solutions at increasingly fine scales, and hence requires a hierarchy of discretizations of the operator. For unstructured meshes, methods for computing this hierarchy for variable-coefficient problems are not known. Second, we can use our parameterisation to better condition a forward solver's linear systems. A source of ill-conditioning in these solves is a mismatch between the anisotropy of the problem and the mesh over which solutions are interpolated. We use our scalar function $s(x)$ to produce meshes that are well-adapted to the anisotropy of Q . The idea is related to the construction of Delaunay triangulations formed by projecting points in the plane onto a regular paraboloid. We devote Chapter 2 to the discretization of $s(x)$, and to these two applications of our parameterisation to the forward problem.

The inverse problem is a measurement problem, the goal being to reconstruct conductivity Q throughout the domain from the boundary Neumann-to-Dirichlet map. The Neumann-to-Dirichlet map

$$\Lambda_\sigma : H^{-\frac{1}{2}}(\partial\Omega) \rightarrow H^{\frac{1}{2}}(\partial\Omega), \quad (1.3)$$

$$\Lambda_\sigma(f) = g \quad (1.4)$$

is a positive-definite boundary operator that gives boundary trace $g = u|_{\partial\Omega}$ of the solution to

$$\begin{aligned} -\operatorname{div}(\sigma\nabla u) &= 0, & x \in \Omega, \\ \frac{\partial}{\partial n}(\sigma u) &= f, & x \in \partial\Omega, \\ \int_{\partial\Omega} f &= 0. \end{aligned} \quad (1.5)$$

$H^{-\frac{1}{2}}$ and $H^{\frac{1}{2}}$ are the appropriate spaces for the boundary traces of flux and potential when the solution $u \in H^1(\Omega)$. In practise, measuring the Neumann-to-Dirichlet map amounts to applying fluxes to the boundary $\partial\Omega$ and measuring the resulting solution values on $\partial\Omega$, then using this data to deduce the conductivity throughout the entire domain. This problem appears first in the literature in the 1980 paper by Calderón [29]. This reconstruction is ill-posed in the sense that large perturbations to Q incur only small perturbations to the boundary data, particularly when Q is perturbed far from the boundary, see [39], for example. Building the symmetry, positivity, and divergence-free properties directly into the parameterisation for Q limits the space over which we search for solutions, regularising the

reconstruction. Because the problem is exponentially ill-posed — the modulus of continuity for the map from boundary data to the conductivity grows exponentially with increasing resolution — coarsening the scale at which we resolve Q is also regularising, but only if the up-scaled reconstruction is consistent with the fine-scale, unresolved conductivity. Chapter 3 applies the discretization of Q as scalar function $s(x)$ to this inverse problem.

Save occasional one-dimensional examples, all of the work in the first three chapters considers conductivities in dimension two. Extending the results to dimensions three and higher is not trivial. Not only do we require more than a single scalar parameter to represent the symmetric and divergence-free up-scaled conductivity, but in even simple examples it is possible for the up-scaled conductivity to lose ellipticity. The added freedom in \mathbb{R}^3 allows the formation of interlocking current loops, which up-scaled, seem to give negative conductivities in some directions. Chapter 4 considers these difficulties, and, based on our discretization and work with the inverse problem in \mathbb{R}^2 , suggests directions of future research.

Homogenisation by metric-based up-scaling is at the heart of the multi-scale representation for conductivity. Even if a conductivity is isotropic (scalar) at the finest scale, it can appear anisotropic at coarser scales. Indeed, it is possible to build meta-materials with fine-scale structure manifesting advantages at the coarse-scale. Two examples are the design of laminated materials having increased strength [14], and electromagnetic materials exhibiting cloaking behaviour [43, 74]. This is a third application of homogenisation, that of material design. We do not treat this application here, but refer the reader to [14] for example. The remainder of this chapter introduces metric-based up-scaling, representation (1.2) of homogenised conductivity by second derivatives of a scalar function $s(x)$, and uses metric-based up-scaling to compute the homogenisation for the simple example of a laminated domain.

1.2 Homogenising scalar conductivities

Material properties may be characterised as isotropic — independent of orientation — or anisotropic, the opposite. An observer standing on an isotropic material has no information from which to deduce directionality, while an observer standing in an anisotropic material does see a preferred direction. The anisotropy of problem (1.1) appears in the tensor-valued

	conductivity, Q	flow density, \mathbf{j}	potential, u
Ohm's Law	electric conductivity	current density	voltage
Fick's Law	diffusivity	substance flow	concentration
Fourier's Law	thermal conductivity	heat flow	temperature
Darcy's Law	diffusivity	substance flow	pressure

Table 1.1: Linear constitutive laws for potential-driven flux. In each case, a flux, \mathbf{j} is proportional to the gradient of a potential ∇u with constant of proportionality Q . Q , the conductivity, need not be scalar, indicating that \mathbf{j} and ∇u are not necessarily parallel.

conductivity parameter Q of the elliptic differential operator.

To see how Q represents anisotropy, consider the physical interpretation of problem (1.1). Setting \mathbf{j} the flux, $-\operatorname{div} \mathbf{j} = f$ states that the net flux leaving a point equals the externally-provided source flux f . u is a potential, and $\mathbf{j} = Q\nabla u$ is a linear constitutive law. This law is Ohm's law for electromagnetism, Fick's law for chemical diffusion, Fourier's law for heat conduction, and Darcy's law for porous-media flow (see Table 1.1). Only when Q is a scalar multiple of the identity do the flux and potential gradient have the same ratio irrespective of direction. That is, only scalar conductivities are isotropic. What is interesting is that, as we have already described for the special case of the laminated domain, the anisotropy of the conductivity depends on the scale at which we observe the material.

While not limited to treating laminated materials, asymptotic homogenisation requires a separation between the coarse scale where we observe anisotropy, and a fine scale, where material properties are nearly constant. In our example, the fine scale is given by the pitch ϵ of the laminate in Figure 1.1. Metric-based up-scaling works differently. Instead of assuming the period of the finest scale features from the outset, metric-based up-scaling transforms the fine-scale conductivity by a change of variables. The transformed conductivity varies on the same scales as the original conductivity, but its arithmetic spatial averages faithfully parameterise the up-scaled operator.

Consider the basic problem (1.1), where the conductivity varies over fine scales:

$$\begin{aligned} -\operatorname{div}(\sigma\nabla u) &= f, & x \in \Omega, \\ u &= 0, & x \in \partial\Omega. \end{aligned} \tag{1.6}$$

$\sigma(x)$ represents a conductivity as a positive-definite symmetric tensor, where σ may vary

over a continuum of scales from very coarse to very fine. The homogenisation begins by solving for the d σ -harmonic coordinates, $F : \Omega \rightarrow \Omega$:

$$\begin{aligned} -\operatorname{div}(\sigma \nabla F) &= 0, & x \in \Omega, \\ F &= x, & x \in \partial\Omega. \end{aligned} \tag{1.7}$$

These problems can be in practise quite expensive to solve, as F must be resolved on a scale comparable to the finest scale of σ . However, once we have the σ -harmonic coordinates, the reward is that if we compute the push-forward of σ by transformation F , denoting

$$Q(x) = \frac{\nabla F^T \sigma \nabla F}{|\det \nabla F|} \circ F^{-1}(x), \tag{1.8}$$

then solutions to problem (1.1) will be close to solutions to problem (1.6), even if u is approximated in a (coarse) low-dimensional space. Here we are using the convention $(\nabla F)_{ij} = \partial_{x_i} F_j$ for row i , column j of matrix ∇F . Provided our conductivity σ is scalar at the very finest scale, we can compute coarse-scale approximations, as long as we are willing to compute with the tensor-valued conductivities predicted by the laminate example. Note that $F : \Omega \rightarrow \Omega$, is 1-1 onto.

This choice, Q , for the homogenised σ that allows us to approximate solutions in a low-dimensional space is the main result of Owhadi and Zhang [65, 88]. Details are provided in that work, but we summarise the argument here. In their paper, the authors prove that under coercivity and boundedness conditions on σ , if u solves (1.6), then $\hat{u} = u \circ F \in W^{2,p}$, $p > 2$, the Sobolev space of functions having second-order weak derivatives in L^p . In \mathbb{R}^2 , the application of a Sobolev embedding theorem shows that \hat{u} is surprisingly smooth: $u \circ F \in C^{1,\alpha}(\Omega)$, the Hölder space of functions having α -Hölder continuous first derivatives. In particular,

$$\|(\nabla F)^{-1} \nabla u\|_{C^\alpha} = \|\nabla \hat{u}\|_{C^\alpha} \leq C \|f\|_\infty, \tag{1.9}$$

for constants $\alpha, C > 0$. Both C and the Hölder parameter α depend on Ω , on bounds on σ and on the eigenvalues of $\nabla F^T \sigma \nabla F$. In addition, C depends on d . The solution $u \in H^1$, and will vary at small scales if σ does. Hence, we may interpret this bound to mean that the harmonic coordinates F contain all of the fine-scale variation of u , and subsequently provide a form of compensation, giving us the solution's smooth counterpart, \hat{u} . This

bound holds irrespective of the scales at which σ varies, implying that $\hat{u} = u \circ F^{-1}$ may be well-approximated in a (low-dimensional) space, for example, the space of piecewise linear polynomials.

In this thesis we choose \hat{u}, \hat{v} from $\mathcal{P}_1^h = \text{span}(\varphi_i)$, the space of piecewise linear polynomials of scale h interpolating functions at vertices i of a triangulation of Ω . The φ_i are the so-called *hat functions*. In such case, solutions to (1.6) may be approximated by linear combinations of functions $\psi_i = \varphi_i \circ F$. Setting $u^h \in \text{span}(\psi_i)$, where h represents the scale of the triangulation, [65] gives

$$\|u - u^h\|_{H^1} \leq Ch^\alpha \|f\|_{L^\infty}. \quad (1.10)$$

This justifies using Q , the push-forward of σ under transformation F , as the tensor parameterising the homogenised elliptic operator.

One way to see why bound (1.10) holds in \mathbb{R}^2 is to note that interpolating solution $u \in C^{1,\alpha}$ by a piecewise linear function $v^h \in \mathcal{P}_1^h$ gives an approximation error

$$\|u - v^h\|_{H^1} \leq Ch^\alpha. \quad (1.11)$$

The regularity of u determines α and the magnitude of the second derivatives of u determine C . Next, it is shown in [65] that the Q -norm,

$$\|v\|_Q = \int_{\Omega} \nabla v^T Q \nabla v \quad (1.12)$$

is equivalent to the H^1 -norm. Now we let u^h be the finite element solution to our PDE in \mathcal{P}_1^h . Céa's lemma says that

$$\|u - u^h\|_Q \leq \inf_{v \in \mathcal{P}_1^h} \|u - v\|_Q. \quad (1.13)$$

Our interpolation $v^h \in \mathcal{P}_1^h$, and the Q -norm and H^1 -norm are equivalent, so the finite element solution satisfies (1.10). The explicit dependence on $\|f\|_{L^\infty}$ makes sense because C in (1.11) depends on the magnitude of the weak second derivatives of solution u .

From definition (1.8), we can see that when homogenised from a scalar conductivity, Q is indeed positive-definite, symmetric, and divergence-free. Positivity follows directly from

the positivity of σ , the physical requirement that a non-zero potential gradient imparts a non-zero flux approximately parallel to the gradient. Symmetry follows from the left and right applications of ∇F . Symmetry is consistent with the condition that Q be used to compute the energy of a potential field. For example, in electromagnetism, the power liberated by a potential field is the inner product of flux and the potential gradient,

$$E(u) = \int_{\Omega} \nabla u^T \mathbf{j} = \int_{\Omega} \nabla u^T Q \nabla u. \quad (1.14)$$

This is the Dirichlet norm for the operator, and forms an inner product given a symmetric positive-definite Q .

While the positivity and symmetry properties are standard in representing physical conductivities, the divergence-free property is special to metric-based up-scaling. Q divergence-free means $\operatorname{div}(Q\mathbf{e}) = 0$ for all constant vectors $\mathbf{e} \in \mathbb{R}^d$. Choosing an arbitrary constant vector \mathbf{e} and test function $w \in H_0^1(\Omega)$, we integrate by parts:

$$\int_{\Omega} w \operatorname{div}(Q\mathbf{e}) = - \int_{\Omega} \nabla w^T \frac{\nabla F^T \sigma \nabla F}{|\det \nabla F|} \circ F^{-1} \mathbf{e} \quad (1.15)$$

$$= - \int_{\Omega} \nabla(w \circ F)^T (\sigma \nabla F) \mathbf{e} \quad (1.16)$$

$$= \mathbf{e}^T \int_{\Omega} w \circ F \operatorname{div}(\sigma \nabla F) = 0. \quad (1.17)$$

So Q is at least divergence-free in a weak sense. While we construct our parameterisation for Q from the standpoint that columns of Q be pointwise divergence-free, we only require the weak condition in our numerical implementations. The divergence-free property constrains the up-scaled conductivity to admit affine potential functions u as solutions to $\operatorname{div}(Q\nabla u) = 0$. On the coarsest scale, this is exactly the kind of average physical behaviour we expect: applying a large potential to one side of a domain and a small potential to the other, the coarsest average response is linear decay across the domain. Metric-based up-scaling gives a conductivity that can be appropriately averaged over *any* scale, beginning with the coarsest, so it is appropriate that our homogenisation method should support the coarsest-scale response.

1.3 Anisotropic conductivities in \mathbb{R}^2

Having recalled the metric-based homogenisation of σ , we will now apply the three constraints of positivity, symmetry, and zero divergence to parameterise Q in \mathbb{R}^2 . Saving mention of higher-dimensional problems for the final chapter, we will constrain our calculations to \mathbb{R}^2 . This choice is not only for clarity of the presentation and simplicity in the numerics: the added freedom afforded in three and higher dimensions can give singular homogenised conductivities, as pointed out in [65, Figure 1.6]. Further discussion of this issue appears in Chapter 4.

We begin construction of our parameterisation by satisfying the divergence-free property, checking that $\operatorname{div}(Q\mathbf{e}^{(i)}) = 0$ for $e_j^{(i)} = \delta_{ij}$. In \mathbb{R}^2 , this amounts to

$$\partial_x Q_{11} + \partial_y Q_{12} = 0, \quad (1.18)$$

$$\partial_x Q_{21} + \partial_y Q_{22} = 0. \quad (1.19)$$

This is reminiscent of a harmonicity condition, where we might choose scalar functions h and k such that

$$Q = \begin{pmatrix} \partial_y h & -\partial_y k \\ -\partial_x h & \partial_x k \end{pmatrix}. \quad (1.20)$$

Observing that Q is symmetric, we further impose $\partial_y k = \partial_x h$. Now we define scalar function $s(x, y)$ such that $h = \partial_y s, k = \partial_x s$, and we have our stated parameterisation, equation (1.2):

$$Q = \begin{pmatrix} \partial_{yy} s & -\partial_{xy} s \\ -\partial_{xy} s & \partial_{xx} s \end{pmatrix}. \quad (1.2)$$

So, given a sufficiently smooth function $s(x, y)$, its second derivatives define a symmetric and divergence-free tensor Q . Furthermore, as a consequence of the Poincaré lemma, if Ω is simply-connected, every irrotational vector field is also a conservative field. That is, we can always compute h, k from a given Q , and then compute s from h, k . For example, denoting vector field $\omega = (-Q_{12}, Q_{11})$, the divergence-free condition gives $\operatorname{curl} \omega = 0$. With Ω simply-connected, we have that there exists scalar function α such that $\omega = \nabla \alpha$, and identifying $\alpha = h$, we have that $\partial_x h = -Q_{12}$, and $\partial_y h = Q_{11}$. Justification for the existence of k and s is similar. h, k and s are not unique. h, k are determined only up to addition of

an arbitrary constant, and s up to the addition of an arbitrary affine function.

Finally, we determine which $s(x, y)$ represents positive-definite Q . In \mathbb{R}^2 , Q is positive-definite if and only if its determinant and trace are both positive. In terms of $s(x, y)$, we require

$$\det Q = \partial_{xx}s\partial_{yy}s - (\partial_{xy}s)^2 \quad (1.21)$$

$$\operatorname{tr} Q = \partial_{xx}s + \partial_{yy}s > 0. \quad (1.22)$$

Furthermore, note that setting H_s the Hessian of $s(x)$, and the 90-degree rotation matrix

$$R = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad (1.23)$$

we have that $H_s = RQR^T$. Hence, Q has the same eigenvalues as the Hessian of $s(x)$, while the directions of the principal curvatures of $s(x)$ are rotated 90 degrees with respect to the eigenvectors of Q . Hence, Q positive-definite is equivalent to H_s positive-definite, and Q positive-definite is equivalent to $s(x)$ being a strictly convex function.

We have thus far been very much treating Q as a tensor field rather than as a parameter for the bilinear operator. Rather, in Chapter 2 when we discretize this operator in the basic problem (1.1), and in Chapter 3, where we discuss what aspect of material conductivities can actually be measured, we consider Q only as an integrated quantity. In particular, if $\operatorname{span}(\varphi_i)$ form a basis for approximating solutions \hat{u} , we are concerned with

$$q_{ij} = - \int_{\Omega} \nabla \varphi_i^T Q \nabla \varphi_j \quad (1.24)$$

$$= - \int_{\Omega} \nabla \psi_i^T \sigma \nabla \psi_j, \quad (1.25)$$

where, again, $\psi_i = \varphi_i \circ F$. This is the *stiffness matrix* familiar from the finite element literature. The negative sign is a convention chosen such that $q_{ij} > 0, i \neq j$ for most q_{ij} . Referring back to equation (1.17), it is because Q appears in integrals for the stiffness matrix that it is sufficient to consider the divergence-free property only in the weak sense. Furthermore, when we compute the relationship between q_{ij} and piecewise linear interpolants s_i of $s(x)$, we can take advantage of this integral form to consider second derivatives of our

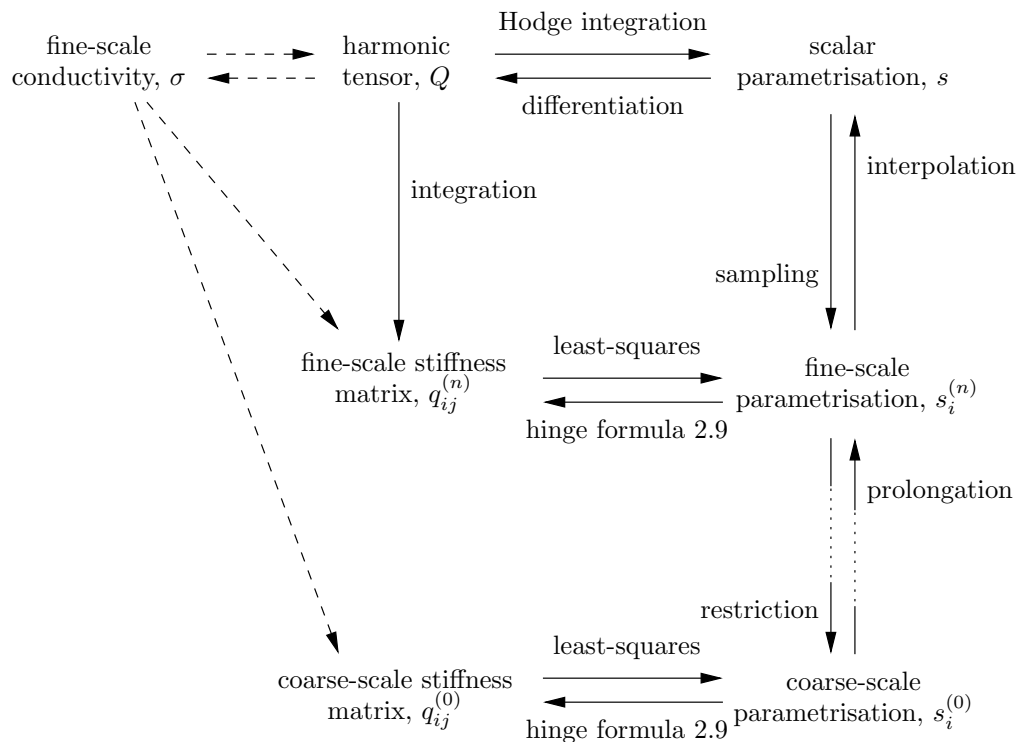


Figure 1.2: Relationships between the different representations for conductivity. Dashed arrows represent non-linear homogenisation through harmonic coordinates, and solid arrows represent linear operations. We apply prolongation and restriction over a sequence of discrete scales as represented by the dotted arrows. Refer to the text for details.

parameterisation $s(x)$ in the sense of distributions.

Figure 1.2 summarises the relationships between the quantities representing conductivity in \mathbb{R}^2 . The top line is the sequence of representations for conductivities given in the continuum. Computation of Q from σ is summarised here and presented in detail by Owhadi and Zhang [65]. Computation of σ from Q is new in this thesis and is presented as Theorem 3.1 in Chapter 3. The relationship between σ and Q is non-linear, but the benefit is that once Q is known, subsequent up-scaling is linear.

Existence of $s(x)$ given Q is given above, consequential to the Hodge decomposition, and computation of $s(x)$ from Q can be accomplished by a series of PDE solves for h, k , then s . In practise, we compute the fine-scale interpolants $s_i^{(n)}$ by a least-square solve from the stiffness matrix elements $q_{ij}^{(n)}$. We up-scale $s(x)$ from its fine-scale discretisation by prolongation of the function $s(x)$ to coarser scales, then subsequent computation of $q_{ij}^{(k)}$ from these coarse-scale interpolants. An alternative is to compute a family of stiffness

matrices, each directly from integration of the original conductivity σ against appropriate test functions. Our parameterisation avoids the need for this non-linear approach.

The appearance of Q exclusively in the elliptic operator implies that $s(x)$ need only be in H^2 , rather than C^2 . Furthermore, as a consequence of the smoothness of solutions found in harmonic coordinates, arithmetic averages of Q give appropriate approximations. We never need second derivatives of $s(x)$ itself, but only integrals of these second derivatives. We take advantage of this weaker regularity condition, and approximate $s(x)$ by piecewise-linear functions in our discretizations. Our discretization and up-scaling procedure will be the subject of Chapter 2.

Finally, we refer to the upward-pointing arrows on the right-hand side of Figure 1.2. These are inverse operations, whereby we are producing fine-scale representations of conductivity from coarse-scale ones. Without additional information, this inverse procedure is not well-defined, and such down-scaling typically makes assumptions of smoothness, either of $s(x)$ or of Q itself. Previous work in inverse problems down-scales σ directly. However, this procedure is typically applied only to scalar conductivities, and we know that in general, conductivities are not effectively scalar at all scales. Chapter 3 studies the inverse problem in more detail, and before considering the discretization of $s(x)$, we further illustrate the up-scaling problem by closing this chapter with an example.

1.4 A simple example: the laminated domain

To get a sense of the relationship between the second derivatives, or curvatures of $s(x)$, the homogenised conductivity Q , and the fine-scale conductivity σ , consider again the conductivity of the laminated domain of Figure 1.1. We will explicitly compute $Q(x)$ for this conductivity, and then take its spatial arithmetic average. Although we arrive at the same homogenised conductivity computed by asymptotic homogenisation, we see this equivalence only after the averaging step. The point is that while Q may be safely averaged, it still varies at the same fine scale as σ .

Let the laminated domain be the unit square; $\Omega = [0, 1] \times [0, 1]$. Partition the square into subsets A and B , where connected components of A and B alternate in the y -direction to form the lamination pattern of Figure 1.1. Let each lamination, that is, each connected component of A, B , have height $\epsilon/2$. Let σ_A, σ_B be the constant scalar conductivities of

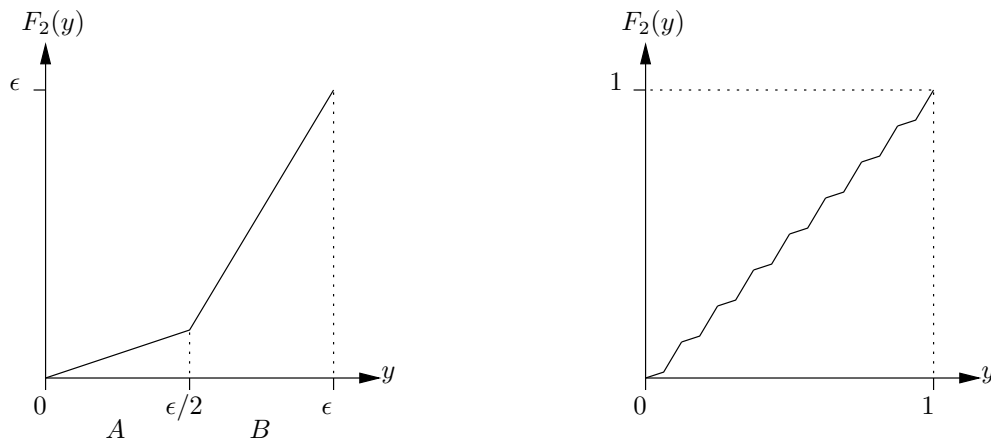


Figure 1.3: The harmonic coordinate $F_2(x, y)$ for domain laminated by horizontal layers where $\sigma_A > \sigma_B$. The figure on the left shows the solution over an ϵ -cell, and the figure on the right shows how this harmonic coordinate captures the fine-scale variation in σ over the entire domain.

A, B , respectively. Given these constant conductivities, we compute the harmonic coordinates $F = (F_1(x, y), F_2(x, y))$ to be

$$F_1(x, y) = x, \quad (1.26)$$

$$F_2(x, y) = \begin{cases} \frac{2\sigma_B}{\sigma_A + \sigma_B}y + b_i, & y \in A, \\ \frac{2\sigma_A}{\sigma_A + \sigma_B}y + c_i, & y \in B. \end{cases}$$

b_i, c_i are constants such that F_2 is continuous at lamination boundaries. The calculation for F_1 is obvious; F_2 is computed by observing that since σ_A, σ_B are constant, F_2 is piecewise linear. To determine the gradient of F_2 , consider a single pair of laminations, height ϵ . Over each of these pairs, a *cell* in the asymptotic homogenisation terminology, F_2 has a net slope in the y -direction of 1, as its net slope over the entire y -extent is 1 also, and each cell is identical. Solving equation (1.7) over the cell gives the gradients in (1.26). Figure 1.3 shows $F_2(y)$ over a single cell and over the entire domain.

Denote \hat{Q} the value of Q prior to composition with F^{-1} , $\hat{Q} = \frac{\nabla F^T \sigma \nabla F}{|\det \nabla F|}$. σ piecewise

constant gives \hat{Q} constant over each of A, B :

$$\hat{Q} = \begin{cases} Q_A = \begin{pmatrix} \frac{\sigma_A}{2\sigma_B}(\sigma_A + \sigma_B) & 0 \\ 0 & \frac{2\sigma_A\sigma_B}{\sigma_A + \sigma_B} \end{pmatrix}, & x \in A, \\ Q_B = \begin{pmatrix} \frac{\sigma_B}{2\sigma_A}(\sigma_A + \sigma_B) & 0 \\ 0 & \frac{2\sigma_A\sigma_B}{\sigma_A + \sigma_B} \end{pmatrix}, & x \in B. \end{cases} \quad (1.27)$$

Since $\hat{Q}(x) = Q(F(x))$, we have

$$Q = \begin{cases} Q_A, & F(x) \in A \\ Q_B, & F(x) \in B, \end{cases} \quad (1.28)$$

or

$$Q = \begin{cases} Q_A, & x \in F(A) \\ Q_B, & x \in F(B). \end{cases} \quad (1.29)$$

Q is piecewise constant, and varies at the same scale as σ , so it appears that this calculation has done nothing but make things more complicated by turning our scalar σ into a matrix. However, we can now affect the homogenisation by approximating Q by its arithmetic integral average, $\langle Q \rangle$. This is justified, for example, by a finite element treatment of the problem. If φ_i, φ_j are functions in the piecewise linear basis having overlapping support, then to solve problems using Q , we will compute stiffness matrix elements

$$q_{ij} = - \int_{\Omega} \nabla \varphi_i^T Q \nabla \varphi_j. \quad (1.30)$$

However, on individual triangles, $\nabla \varphi_i, \nabla \varphi_j$ are constant, so these matrix elements are exactly given by integrals of Q over triangles T_j , namely, the quantities

$$- \int_{T_j} \nabla \varphi_i^T \langle Q \rangle_{T_j} \nabla \varphi_j, \quad \langle Q \rangle_{T_j} = \frac{1}{|T|} \int_T Q. \quad (1.31)$$

This choice to take averages over triangles with an appeal to the piecewise linear basis is arbitrary: it is only at this step where we set the scale to which we up-scale the conductivity.

To compute $\langle Q \rangle$, we note that (1.29) implies that $Q = Q_A$ over fraction $\lambda_A = \frac{\sigma_B}{\sigma_A + \sigma_B}$ of the domain, and similarly for Q_B . Since Q is piecewise constant, the integral amounts to a weighted sum:

$$\langle Q \rangle = \lambda_A Q_A + \lambda_B Q_B \quad (1.32)$$

$$= \begin{pmatrix} \frac{\sigma_A + \sigma_B}{2} & 0 \\ 0 & \frac{2\sigma_A\sigma_B}{\sigma_A + \sigma_B} \end{pmatrix}. \quad (1.33)$$

This weighted sum holds provided we integrate over an integral number of pairs of laminates: the smallest domain over which we can average has height ϵ , exactly the size of the single cell identified in asymptotic homogenisation, and this $\langle Q \rangle$ is that computed by asymptotic homogenisation for small ϵ .

The homogenised conductivity is a good representation of the average conductivity of the domain. Despite the fact that the fine-scale conductivity is isotropic, the homogenised, or coarse-scale, conductivity is the arithmetic mean, $\bar{\sigma}$, of σ_A, σ_B parallel to the laminations, and their harmonic mean, $\underline{\sigma}$, perpendicular to the laminations¹. Note that a simple volume average of σ would give an isotropic conductivity of $\bar{\sigma}$, and does not predict the anisotropy of the averaged media.

Given Q , we would now like to calculate its corresponding parameterisation $s(x, y)$. This is non-trivial, even in this simple case, precisely because σ is itself not constant. σ , hence Q , has fine-scale variations, so the curvatures of $s(x)$ will also vary at a fine scale. However, we are seeking coarse-scale solutions to the basic problem (1.6), and would thus prefer to use the homogenisation $\langle Q \rangle$. When averaged over sufficiently large sub-domains, $\langle Q \rangle$ is constant over the domain, and we claim that up to the addition of an arbitrary affine function,

$$s(x, y) \approx \frac{1}{2} \langle Q \rangle_{22} x^2 - \langle Q \rangle_{12} xy + \frac{1}{2} \langle Q \rangle_{11} y^2, \quad (1.34)$$

$$= \frac{1}{2} \underline{\sigma} x^2 + \frac{1}{2} \bar{\sigma} y^2. \quad (1.35)$$

¹These results may be familiar as the formulae for conductances in serial and parallel: the effective conductance of conductors in parallel is their arithmetic sum, and of conductors in serial is their harmonic sum.

This approximation to $s(x, y)$ represents a constant conductivity, which is trivially symmetric and divergence free, and provided $\sigma_A, \sigma_B > 0$, likewise is $s(x, y)$ convex, and is $\langle Q \rangle$ positive definite. In the next chapter, we demonstrate that when we seek solutions to the basic problem (1.1) at sufficiently coarse scales — at scales greater than ϵ in this case — averaging $s(x, y)$ is equivalent to averaging Q and hence homogenising σ .

Chapter 2

The Forward Problem

Although this thesis project began with a study of the inverse problem, it is more natural in this presentation to begin our closer look at the basic problem

$$\begin{aligned} -\operatorname{div}(Q\nabla u) &= f, & x \in \Omega, \\ u &= 0, & x \in \partial\Omega, \end{aligned} \tag{1.1}$$

by considering the forward problem of determining the potential field u for a given conductivity σ . In solving the homogenised forward problem, we need to construct a discrete approximation of $s(x)$, and in so doing, we will see what to expect when we attempt to reconstruct $s(x)$ in the inverse problem. Inspired by recent work in discrete differential geometry, see for example [34, 35], we choose to calculate fields and homogenise conductivities using piecewise linear finite element spaces over simplicial (triangle) meshes, working nearly exclusively in finite spaces.

There are several reasons to discretize the problem over simplicial meshes supporting piecewise linear functions. With this choice, we use data structures that treat single elements at a time, where each element has information only about which elements are its immediate neighbours. Unlike regular meshes, where we often rely on each element's position in a global coordinate system, relying only on neighbour information reflects the locality of the differential problem we are discretizing. From a more practical point of view, the small support of the basis functions gives us simple formulae relating problem variables, and hence gives linear operators that can be represented as sparse matrices. The sparse relationships also makes interpreting our formulae straightforward — for example, we develop an interpretation of positivity of Q via convexity of the discretized $s(x)$ in Section 2.1. Finally, we can take advantage of the wealth of existing open software available for building, manipulating, and viewing triangle meshes. Appendix A lists the software we use in our

calculations.

In the first section of this chapter we discretize the conductivity through a piecewise linear interpolation of $s(x)$. This will complete the connection between the fine-scale σ and the multi-scale stiffness matrices for the problem anticipated by Figure 1.2 and already begun with the continuum construction in Section 1.3. The remaining two sections of this chapter apply the discretized $s(x)$ to the construction of meshes well-adapted to the anisotropy of the up-scaled conductivity, then the construction of multi-scale approximations to $\operatorname{div}(Q\nabla\cdot)$. These applications show connections to weighted Delaunay triangulations and multi-grid solution methods, respectively.

2.1 Discretizing the conductivity

We solve the problem using a Galerkin method, seeking weak solutions to the basic problem (1.1) in a finite subspace of $H_0^1(\Omega)$. We use the notation $H_0^1(\Omega)$ to be the space of functions having L_2 -integrable weak first derivatives taking value zero on $\partial\Omega$. We will determine the relationship between the homogenised stiffness matrix and the piecewise linear interpolation of $s(x)$.

To specify the Galerkin method, fix \mathcal{T}^h a triangle mesh with nodes \mathcal{N}^h . Set h the largest circumdiameter over all triangles in the mesh. h is used throughout to denote the linear scale of our discretizations. Single indices i, j, k, l denote vertices, pairs of indices ij, jk denote edges, and triples ijk, ijl denote triangles. Vertices have location (x_i, y_i) or just x_i when the dimension of the mesh is understood. The mesh is oriented: ij is the edge beginning at vertex i and ending at vertex j ; ijk is the triangle traversed from vertex i to j to k . When ijk appear in counterclockwise order around the triangle, the triangle has positive area, otherwise, the area is negative. Further discussion of mesh orientation is in [34, 63].

Set $V^h \subset H_0^1(\Omega)$ where we choose $V^h = \mathcal{P}_1^h$ to be the set of piecewise linear polynomials at scale, linear on triangles of scale h . Basis functions φ_i are often referred to as *hat functions*, where i indexes the vertex where $\varphi_i = 1$. Multiplying equation (1.1) by test

function $v^h \in V^h$, and integrating by parts, we seek the solution u_h satisfying

$$\int_{\Omega} (\nabla v^h)^T Q \nabla u^h = \int_{\Omega} v^h f, \quad \forall v^h \in V^h. \quad (2.1)$$

Setting $u_i = u^h(x_i)$, this amounts to solving the linear system

$$\int_{\Omega} \nabla \varphi_i^T Q \nabla u^h = \int_{\Omega} \varphi_i f \quad (2.2)$$

$$\sum_{j \in \mathcal{N}_h} u_j \int_{\Omega} \nabla \varphi_i^T Q \nabla \varphi_j = \int_{\Omega} \varphi_i f. \quad (2.3)$$

Recall that provided Q is the metric up-scaled σ , the accurate approximation of the true solution $\hat{u} \in C^{1,\alpha}$ in this basis is justified. What this equation implies is that the only information we need about the conductivity is the elements

$$q_{ij} = - \int_{\Omega} \nabla \varphi_i^T Q \nabla \varphi_j. \quad (2.4)$$

These quantities form the *stiffness matrix* familiar in the finite element literature. As mentioned in Chapter 1, the negative sign, while arbitrary, makes most $q_{ij} > 0$.

There is redundancy in the stiffness matrix. First, it is symmetric in i, j , inheriting the symmetry of tensor Q . Second, since piecewise linear functions exactly interpolate constants, $q_{ii} = - \sum_{j \in \mathcal{N}(i)} q_{ij}$. $\mathcal{N}(i)$ is the set of vertices sharing an edge with vertex i . Thus, we are solving the system

$$\sum_{j \in \mathcal{N}(i)} q_{ij} (u_i - u_j) = f_i, \quad (2.5)$$

where $f_i = \int_{\Omega} \varphi_i f$. This equation relies only on the connectivity of the triangulation, and we can conveniently associate the q_{ij} with edges, and the u_i, f_i with nodes of the mesh.

The only information we have about conductivity is contained in the stiffness matrix. That is, it is the integrated, not pointwise, values of Q that matter to finding approximate discrete solutions to the basic problem. This indicates why up-scaling σ to give Q is important: construction of the stiffness matrix takes a weighted arithmetic average of the conductivity which we know can give inaccurate solutions when σ has fine-scale variations. Instead of parameterising q_{ij} in terms of Q , however, we next show how to build the stiffness




Figure 2.1: Notation for computing q_{ij} over a hinge from s_i, s_j, s_k , and s_l . The supports of φ_i and φ_j are shown, with the support of the integrand of (2.4) shaded.

matrix directly from a piecewise linear approximation of $s(x)$.

Referring to Figure 2.1, we will recover q_{ij} as a linear combination of s_i, s_j, s_k , and s_l . Just as for scalar $u(x)$, we set $s_i = s(x_i)$. Coefficients of the relationship between the s_i and the stiffness matrix depend only on the geometry of the triangle mesh. The collection of the four vertices and two triangles surrounding an interior edge is often referred to as a *hinge*.

Rewrite (2.4) in the rotated coordinate system of Figure 2.1:

$$q_{ij} = - \int_{\Omega} (\nabla \varphi_i)^T \begin{pmatrix} \partial_{\eta\eta} s & -\partial_{\xi\eta} s \\ -\partial_{\xi\eta} s & \partial_{\xi\xi} s \end{pmatrix} \nabla \varphi_j. \quad (2.6)$$

A change of variables confirms that integral (2.6) is invariant under rotation and translation. We abuse notation in that the second derivatives are understood here as distributional derivatives: we are about to interpolate $s(x)$ by piecewise linear functions, which do not have pointwise second derivatives everywhere.

We have chosen coordinate system ξ - η such that edge ij is parallel to the η -axis. We are concerned with the values of $s(x)$ interpolated at i, j, k , and l , as these are associated to only the corresponding hat basis functions sharing support with those at i and j . The second derivatives of the φ are non-zero only on edges, and due to the support of the

gradients of the φ , contributions of the second derivatives at edges ik , jk , il , and jl are also zero. Finally, the $\partial_{\xi\eta}\varphi$ and $\partial_{\eta\eta}\varphi$ are zero along ij , so the only contributions of $s(x)$ to the integral (2.4) are its second derivatives with respect to ξ along edge ij .

The contributions of four integrals remain, and by symmetry, we have only two integrals to compute. Noting that the singularities in the first and second derivatives are not coincident, from direct computation of the gradients of the basis functions and integration by parts we have

$$\int_{ijk \cup ijl} \partial_{\eta}\varphi_i \partial_{\xi\xi}\varphi_i \partial_{\eta}\varphi_j = \frac{1}{|ij|^2} (\cot \theta_{ijk} + \cot \theta_{ijl}), \quad (2.7)$$

$$\int_{ijk \cup ijl} \partial_{\eta}\varphi_i \partial_{\xi\xi}\varphi_k \partial_{\eta}\varphi_j = -\frac{1}{2A_{ijk}}. \quad (2.8)$$

The only contribution to these integrals is in the neighbourhood of edge ij . A_{ijk} is the unsigned area of triangle ijk and, for example, θ_{ijk} is the interior angle of triangle ijk at vertex j . Combining these results, we have that the elements of the stiffness matrix are

$$\begin{aligned} q_{ij} = & -\frac{1}{|ij|^2} (\cot \theta_{ijk} + \cot \theta_{ijl}) s_i - \frac{1}{|ij|^2} (\cot \theta_{jik} + \cot \theta_{jil}) s_j \\ & + \frac{1}{2A_{ijk}} s_k + \frac{1}{2A_{ijl}} s_l. \end{aligned} \quad (2.9)$$

Diagonal elements q_{ii} of the stiffness matrix are computed by $q_{ii} = -\sum_{j \in \mathcal{N}(i)} q_{ij}$.

Equation (2.9) is valid only for interior edges. Because of our choice to interpolate $s(x)$ by piecewise linear functions, we have concentrated all of the curvature of $s(x)$ on the edges of the mesh, and we need a complete hinge, an edge with two incident triangles, in order to approximate this curvature. Without values for $s(x)$ outside of Ω and hence exterior to the mesh, we do not have a complete hinge on boundary edges. This will become important where we solve the inverse problem in Chapter 3, but in our model forward problem our homogeneous boundary conditions make irrelevant the values of q_{ij} on boundary edges.

Equation (2.9) represents the discretized homogenised elliptic operator, and preserves several properties of the continuous operator exactly. We next confirm by direct calculation that the discretization preserves the divergence-free property, that the operator is unchanged under the addition of arbitrary affine functions to $s(x)$, and that our formula is equivalent to the well-known cotangent formula when Q is isotropic and constant.




Figure 2.2: One-ring surrounding an interior vertex i , with the notation for confirming that the divergence-free property is exactly preserved by our discretization of the elliptic operator.

First, the discretized operator exactly preserves the divergence-free property. In the discrete setting, we check that for all choices of $s(x_i, y_i)$,

$$\sum_{j \in \mathcal{N}(i)} q_{ij} (\mathbf{x}_i - \mathbf{x}_j) = 0, \quad (2.10)$$

where $\mathbf{x}_i = (x_i, y_i)$ are the coordinates of the mesh vertices. (2.10) is the discretization of the divergence-free condition in the continuum,

$$\operatorname{div}(Q\mathbf{e}^{(k)}) = \operatorname{div}(Q\nabla x^{(k)}) = 0, \quad k = 1 \dots d. \quad (2.11)$$

The $\mathbf{e}^{(p)}$ form the usual orthonormal basis, which for the convenience of considering the q_{ij} directly, we represent as the gradients of identity map, $\nabla x = \operatorname{Id}$. Referring to Figure 2.2, choose interior vertex i to be at the origin, and take value $s = (x_i, y_i) = 0$. Write the locations of the surrounding vertices in complex notation, such that $s(x_j, y_j) = s(a_j e^{I\alpha_j}) = s_j$. Here, $I = \sqrt{-1}$. Let the area of the triangle counterclockwise from edge ij be A_j , and

let index addition be modulo n , where n is the valence of vertex i . Then, directly compute

$$\sum_{j \in \mathcal{N}(i)} q_{ij} (\mathbf{x}_i - \mathbf{x}_j) = \sum_{j \in \mathcal{N}(i)} q_{ij} a_j e^{I\theta_j}, \quad (2.12)$$

$$= - \sum_{j \in \mathcal{N}(i)} a_j e^{I\theta_j} \left[\frac{s_j}{a_j^2} (\cot(\alpha_{j+1} - \alpha_j) + \cot(\alpha_j - \alpha_{j-1})) - \frac{s_{j+1}}{2A_j} - \frac{s_{j-1}}{2A_{j-1}} \right], \quad (2.13)$$

$$= - \sum_{j \in \mathcal{N}(i)} s_j e^{I\theta_j} \left[\frac{s_j}{a_j} (\cot(\alpha_{j+1} - \alpha_j) + \cot(\alpha_j - \alpha_{j-1})) - \frac{a_{j-1} e^{I(\alpha_{j-1} - \alpha_j)}}{2A_{j-1}} - \frac{a_{j+1} e^{I(\alpha_{j+1} - \alpha_j)}}{2A_j} \right]. \quad (2.14)$$

We observe that $2A_j = a_j a_{j+1} \sin(\alpha_{j+1} - \alpha_j)$ and compute the bracketed term in the sum:

$$\frac{1}{a_j} \left[\cot(\alpha_{j+1} - \alpha_j) + \cot(\alpha_j - \alpha_{j-1}) - \frac{\cos(\alpha_j - \alpha_{j-1})}{\sin(\alpha_j - \alpha_{j-1})} - \frac{\cos(\alpha_{j+1} - \alpha_j)}{\sin(\alpha_{j+1} - \alpha_j)} + I \left(\frac{\sin(\alpha_j - \alpha_{j-1})}{\sin(\alpha_j - \alpha_{j-1})} - \frac{\sin(\alpha_{j+1} - \alpha_j)}{\sin(\alpha_{j+1} - \alpha_j)} \right) \right] = 0. \quad (2.15)$$

Hence, for any s_j and any a_j, α_j coordinates around vertex i , the elliptic operator parameterised by $s(x)$ is exactly divergence-free.

This result is encouraging, as divergence-free operators in the continuum by definition admit the d solutions $u^{(k)} = x^{(k)}, k = 1 \dots d$ to the Dirichlet problems

$$\begin{aligned} \operatorname{div}(Q \nabla u^{(k)}) &= 0, & x \in \Omega, \\ u^{(k)} &= x^{(k)}, & x \in \partial\Omega, \end{aligned} \quad (2.16)$$

and the piecewise linear functions which we use to interpolate the solutions u , exactly interpolate first-order polynomials. Céa's lemma says that with conditions on Q , if $u^h \in V^h$ solves the weak problem (2.1), then this solution is optimal in V^h , in the sense that if $u \in H_0^1$ is the exact solution, then

$$\|u - u^h\| \leq C \inf_{v \in V^h} \|u - v\|. \quad (2.17)$$

However, this statement is only true when the integrations in (2.1) are exact. By interpolating $s(x)$, we have no such guarantee of exactness. Fortunately, we now find that all solutions interpolated by $V^h = \mathcal{P}_1^h$ are exact solutions. The row-sum condition $q_{ii} = - \sum_{j \in \mathcal{N}(i)} q_{ij}$ guarantees that the approximated stiffness matrix admits constant fields as solutions, and

now the construction of q_{ij} from samples s_i guarantees that the approximated stiffness matrix admits exact linear solutions also. Preservation of constant solutions in the null-space of the discretized elliptic operator appears to be important for stability of discretized problems, as is observed in [21]. Similarly, preservation of the divergence-free property — exact admission of affine solutions to the null-space — is important in discretizing elliptic operators.

Next, we compute the stiffness matrix for the case where $s(x, y)$, rather than $u(x, y)$, is an affine function. First, suppose $s(x, y) = \text{constant}$. The q_{ij} are linear in the s_i , so choose $s(x, y) = 1$. In this case, (2.9) becomes

$$q_{ij} = -\frac{1}{\|ij\|^2} (\cot \theta_{ijk} + \cot \theta_{ijl}) - \frac{1}{\|ij\|^2} (\cot \theta_{jik} + \cot \theta_{jil}) + \frac{1}{2A_{ijk}} + \frac{1}{2A_{ijl}} \quad (2.18)$$

$$= \frac{1}{2A_{ijk}} \left(\frac{ij^T ki}{ij^T ij} + \frac{ij^T jk}{ij^T ij} + \frac{ij^T ij}{ij^T ij} \right) + \frac{1}{2A_{ijl}} \left(\frac{ij^T li}{ij^T ij} + \frac{ij^T jl}{ij^T ij} + \frac{ij^T ij}{ij^T ij} \right) \quad (2.19)$$

$$= 0, \quad (2.20)$$

since the sum of edges around a triangle, $ij + jk + ki = 0$, for example. Choosing $s(x, y) = ax + by$ for constants a, b similarly gives $q_{ij} = 0$. We don't show this calculation in detail, but the idea is as follows: we translate the mesh such that the origin is at vertex i , an arbitrary choice, since translation of a linear function is equivalent to the addition of a constant. Then we compute that the contributions of linear functions parallel and perpendicular to edge ij are both zero. Combining this with $ax + by$ in the local rotated frame of each hinge shows that $q_{ij} = 0$ for arbitrary a, b .

The fact that expression (2.9) gives $q_{ij} = 0$ exactly for affine functions is consistent with our interpolation for s . Piecewise linear functions exactly interpolate first-order polynomials, and the addition of functions $ax + by + c$ to $s(x)$ contribute exactly zero to Q in the continuum.

Our final special case is where the conductivity is constant and isotropic. Set $Q = \text{Id}$. In this case, we have by inspection,

$$Q = \text{Id} \Rightarrow s(x, y) = \frac{1}{2}x^2 + \frac{1}{2}y^2 + a(x, y), \quad (2.21)$$

where $a(x, y)$ is an arbitrary affine function. The addition of arbitrary $a(x, y)$ is equivalent to translating $s(x, y)$ around the plane, so given hinge $ijkl$, we choose vertex i to be at the origin. Thus, $s(x_p, y_p) = \frac{1}{2}\|ip\|^2 = \frac{1}{2}ip^T ip$ for $p = j, k, l$. Substituting $\cot \theta_{ijk} = -\frac{ij^T jk^T}{2A_{ijk}}$, for vertices i, j, k , equation (2.9) gives

$$q_{ij} = \frac{1}{2} \left(\frac{ij^T ki}{2A_{ijk}} + \frac{ij^T li}{2A_{ijl}} \right) + \frac{ik^T ik}{4A_{ijk}} + \frac{il^T il}{4A_{ijl}}, \quad (2.22)$$

$$= \frac{ik^T(ik - ij)}{4A_{ijk}} + \frac{il^T(il - ij)}{4A_{ijl}}, \quad (2.23)$$

$$= \frac{1}{2} (\cot \theta_{ikj} + \cot \theta_{ilj}). \quad (2.24)$$

Again, all triangle areas are unsigned. This result is precisely the famous cotangent formula for the stiffness matrix of the discretized Laplacian. Not only is this the result obtained by direct discretization of solutions by piecewise linear functions, but also appears as the discrete Laplacian computed, for example, by discrete exterior calculus: see [34] and its references. That a piecewise linear interpolation of a quadratic $s(x)$ should give an exact result for the stiffness matrix is not immediately obvious, and we return to this point in Section 2.2 when we discuss convergence of the operator under mesh refinement.

We next consider the convexity of the piecewise linear surface interpolated by the s_i to confirm the positivity of Q , just as we have established the relationship between positivity and convexity in the continuum. We claim that when the s_i form a convex interpolation of a convex function, the linear operator represented in (2.5) is positive-definite. Using formula (2.9), we next show that when the dihedral angle of a hinge defined by the interpolants s_i is less than 180 degrees if and only if $q_{ij} > 0$. Furthermore, when all edges have convex dihedral angles, the resulting surface is convex.

To see that $q_{ij} > 0$ is equivalent to a convex interior dihedral angle $\theta_d < \pi$, we use the invariance of the stiffness matrix under the addition of an arbitrary affine function to $s(x)$ to consider the special case $s_i = s_j = s_l = 0$. In this case, over a single hinge,

$$q_{ij} = \frac{s_k}{2A_{ijk}}. \quad (2.25)$$

Since we are using the unsigned area, $q_{ij} > 0 \Leftrightarrow s_k > 0 \Leftrightarrow \theta_d < \pi$ (see Figure 2.3).

To relate this result to the positivity of the stiffness matrix, and the positivity of the




Figure 2.3: The connection between convexity of the surface interpolating $s(x)$ and the positivity of Q can be seen by a calculation of the dihedral angles of each hinge.

elliptic operator it represents, we know that as a consequence of the Gershgorin circle theorem that if all off-diagonal entries of the stiffness matrix are negative, and the stiffness matrix has row-sum zero, then the matrix is at least positive semi-definite. Recalling that we have chosen to consider q_{ij} the negative of the stiffness matrix, the result applies here. However, while $q_{ij} > 0$ is sufficient to guarantee that the stiffness matrix is positive semi-definite, we know from other work in the finite element literature that this element-wise positivity is not a necessary condition [77]. That is, it is possible to have a few $q_{ij} < 0$ and still have a positive-definite stiffness matrix. The consequence, however, is that the matrix is ill-conditioned

From our geometric form for Q , instead of requiring the piecewise linear surface interpolating $s(x)$ be convex, that is, that all edges have $q_{ij} > 0$, we only require that the s_i interpolate a convex surface. For example, consider the quadrilateral in Figure 2.4: its corners lie on the surface $s(x)$. A triangulation of this quadrilateral could be completed by connecting either diagonal, but unless the four corners are co-planar, only one choice gives a convex interpolation. Although single edge flips cannot always make a non-convex interpolation into a convex one, if the underlying surface $s(x)$ is convex, there is always a choice of connectivity that gives a convex interpolating triangulation. This triangulation can be constructed by taking the convex hull of the interpolating vertices. Indeed, this construction motivates the adapted triangulations we consider in Section 2.3. Following the arguments in the previous paragraph, the convex choice also gives the best-conditioned stiffness matrix, and we discuss this further in the next section.




Figure 2.4: Edge flips can replace non-convex edges, where $\theta_d > \pi \Rightarrow q_{ab} < 0$, with convex edges without changing the interpolated values s_i . For the given hinge, the diagonal giving a negative edge is on the left; a positive edge is on the right.

2.2 Multi-scale elliptic operators

It's tidy to have a sparse representation for divergence-free conductivities, but in order to be able to use the discretization of $s(x)$ in computation, we need to verify its convergence. That is, we need to verify that the operator which the discrete $s(x)$ represents is convergent, and, in turn, that the solutions produced in the finite subspaces of H^1 converge to something reasonable. In particular, we would like to use coarse samples of $s(x)$ to produce a hierarchy of operators. We may wish to use the coarsest operators as a preconditioner for finer-scale solves: such a hierarchy is employed in multi-grid solvers [28]. However, before considering the preconditioning application, we must check that the sequence of operators give convergent solutions. If Q is represented exactly, then we already know from results in metric-based up-scaling that a sequence of ever-finer finite element solutions converges to the homogenisation solution. By sampling $s(x)$, however, we are not expressing Q exactly, and we must check that error due to this discretization does not overwhelm error due to the finite element method alone.

2.2.1 Homogenised operator convergence

Set Q the anisotropic conductivity, computed on the continuum, and Q^h the conductivity at scale h . Q^h is computed from $s^h(x)$, an approximation of $s(x)$ over a triangulated domain.

In our case, we have chosen $s^h(x)$ to be a piecewise linear interpolating polynomial, where small h represents a fine scale interpolation; specifically, we set h to be the diameter of the largest circumcircle over all triangles in the mesh. Since $s^h(x)$ is piecewise linear, its second derivatives exist only in the sense of distributions, so we instead prefer to work with the bilinear operators

$$s(x) \Rightarrow Q(x) \Rightarrow \mathcal{B}(u, v) = \int_{\Omega} \nabla u^T Q \nabla v, \quad (2.26)$$

$$s^h(x) \Rightarrow Q^h(x) \Rightarrow \mathcal{B}^h(u, v) = \int_{\Omega} \nabla u^T Q^h \nabla v. \quad (2.27)$$

The order of the argument is as follows: We assume \mathcal{B}^h to converge to \mathcal{B} , showing that \mathcal{B}^h inherits the continuous and coercivity properties of \mathcal{B} . We use coercivity to show that solutions computed using \mathcal{B}^h converge to those of \mathcal{B} . Finally, we confirm that for our specific approximation of $s(x)$, we have convergence of \mathcal{B}^h to \mathcal{B} .

Assume that we control the error incurred by discretizing the operator, whereupon there exists a function $\tilde{C}(h) > 0$ such that

$$|\mathcal{B}^h(u, v) - \mathcal{B}(u, v)| \leq \tilde{C}(h) \|u\|_{H^1} \|v\|_{H^1}, \quad u, v \in V^h. \quad (2.28)$$

$V^h \subset H_0^1$ is where we approximate solutions to the homogenised problem. We choose V^h such that $\forall v^h \in V^h, \exists v \in H_0^1$ such that $\|v^h - v\|_{H^1} \leq Ch \|v\|_{H^1}$. The space of piecewise linear polynomials is one such example of V^h with this property. With such control, we find that \mathcal{B}^h can be continuous and coercive if \mathcal{B} is:

$$\begin{aligned} |\mathcal{B}^h(u, v)| &\leq |\mathcal{B}(u, v) - \mathcal{B}^h(u, v)| + |\mathcal{B}(u, v)| \\ &\leq C(\tilde{C}(h) + 1) \|u\|_{H^1} \|v\|_{H^1}, \end{aligned} \quad (2.29)$$

$$\begin{aligned} |\mathcal{B}^h(u, u)| &\geq |\mathcal{B}(u, u)| - |\mathcal{B}(u, u) - \mathcal{B}^h(u, u)| \\ &\geq C(C_0 - \tilde{C}(h)) \|u\|_{H^1}^2, \end{aligned} \quad (2.30)$$

where here and throughout the following $C > 0$ denotes a constant independent of h . While continuity of \mathcal{B}^h is guaranteed irrespective of $\tilde{C}(h)$, preservation of coercivity depends on

the size of $\tilde{C}(h)$ with respect to C_0 , a constant proportional to the smallest eigenvalue,

$$\lambda_{\min} = \min_{\|u\|_{H^1} \leq 1} \frac{\mathcal{B}(u, u)}{\|u\|_{H^1}}. \quad (2.31)$$

In our special case of approximation of s by piecewise linear polynomials, Lemma 2.2 will show that $\tilde{C}(h)$ is a constant such that $C(C_0 - \tilde{C}) > 0$ for all h .

Next, we repose the basic problem (1.1) and two approximations in weak form. In each case, set $L(v) = \int_{\Omega} f v$. All functions u, u^h, \tilde{u}^h, v are zero on $\partial\Omega$:

$$u \in H_0^1 \quad \text{solves} \quad \mathcal{B}(u, v) = L(v), \quad \forall v \in H_0^1, \quad (2.32)$$

$$u^h \in V^h \quad \text{solves} \quad \mathcal{B}(u^h, v) = L(v), \quad \forall v \in V^h, \quad (2.33)$$

$$\tilde{u}^h \in V^h \quad \text{solves} \quad \mathcal{B}^h(u^h, v) = L(v), \quad \forall v \in V^h. \quad (2.34)$$

We assume the Q in problem (2.32) to be a conductivity up-scaled by the metric, and as such we note $u \in H_0^1$, following [65]. With these definitions, and the preservation of coercivity and continuity under approximation, we have our main result for this section:

Theorem 2.1. *Let $u \in H_0^1$ and $\tilde{u}^h \in V^h$ solve problems (2.32) and (2.34), respectively. Furthermore, suppose operators \mathcal{B}^h converge to \mathcal{B} in the sense of equation (2.28) at rate $\tilde{C}(h) = C_0 h^p$ for some $p > 0$. Then there exists a constant C independent of h such that*

$$\|u - \tilde{u}^h\|_{H^1} \leq C h^\gamma, \quad \gamma = \min(p, \alpha). \quad (2.35)$$

Even though we have approximated the operator, we don't lose an order of convergence in approximating both the operator and the solution. We will see this behaviour in our numerical example, where the irregularity of the homogenisation solution, rather than the approximation of $s(x)$, dominates the error approximating the solution.

Proof of Theorem 2.1. First, we examine the error due to the approximation of the operator.

For $v \in V^h$,

$$|\mathcal{B}^h(\tilde{u}^h - u^h, v)| = \left| \mathcal{B}^h(\tilde{u}^h, v) - \mathcal{B}(u^h, v) + \left(\mathcal{B}(u^h, v) - \mathcal{B}^h(u^h, v) \right) \right| \quad (2.36)$$

$$= |L(v) - L(v) + \mathcal{B}(u^h, v) - \mathcal{B}^h(u^h, v)| \quad (2.37)$$

$$\leq Ch^p \|u^h\|_{H^1} \|v\|_{H^1} \quad (2.38)$$

We have specified the $\mathcal{O}(h^p)$ convergence of our approximated operator. Choose $v = \tilde{u}^h - u^h$, and apply coercivity:

$$|\mathcal{B}^h(\tilde{u}^h - u^h, \tilde{u}^h - u^h)| \leq Ch \|u^h\|_{H^1} \|\tilde{u}^h - u^h\|_{H^1} \quad (2.39)$$

$$\|\tilde{u}^h - u^h\|_{H^1} \leq Ch^p \|u^h\|_{H^1}. \quad (2.40)$$

Finally, u solves (2.32), and since the operator is constructed from a metric up-scaled conductivity, we have

$$\|u - \tilde{u}^h\|_{H^1} \leq \|u^h - \tilde{u}^h\|_{H^1} + \|u - u^h\|_{H^1} \quad (2.41)$$

$$\leq C \left(h^p \|u^h\|_{H^1} + h^\alpha \|f\|_{L^\infty} \right) \quad (2.42)$$

$$\leq Ch^\gamma \|f\|_{L^\infty}, \quad \gamma = \min(p, \alpha). \quad (2.43)$$

We have applied bound (1.10) from [65] to control the approximation of u by u^h . \square

To complete the argument, we check the convergence of \mathcal{B}^h to \mathcal{B} for our particular approximation of $s(x)$ by piecewise linear interpolating polynomials.

Lemma 2.2. *Set $s(x) \in H^2$ such that*

$$\mathcal{B}(u, v) = \int_{\Omega} \nabla u \begin{pmatrix} \partial_{yy}s & -\partial_{xy}s \\ -\partial_{xy}s & \partial_{xx}s \end{pmatrix} \nabla v, \quad (2.44)$$

and interpolate $s(x)$ by piecewise linear $s^h(x)$ over a triangulated domain. Then,

$$|\mathcal{B}^h(u, v) - \mathcal{B}(u, v)| \leq Ch^2 \|u\|_{H^1} \|v\|_{H^1}, \quad u, v \in V^h, \quad (2.45)$$

where all triangles have circumcircles of diameter less than or equal to h , and $V^h \subset H_0^1$ is

the space where we approximate solutions to the continuous problem.

Proof. Recalling the calculation of q_{ij} from the interpolated $s_i = s(x_i)$, since $s^h(x)$ is defined only in the sense of distributions, we find that by integration by parts, instead of comparing second derivatives of s, s^h over triangle faces, we compare their first derivatives over edges. Setting (φ_i) piecewise linear hat functions interpolating solutions at vertices, consider

$$q_{ij} = - \int_{\Omega} \nabla \varphi_i^T Q \nabla \varphi_j, \quad (2.46)$$

$$q_{ij}^h = - \int_{\Omega} \nabla \varphi_i^T Q^h \nabla \varphi_j. \quad (2.47)$$

Using the $\xi - \eta$ coordinates in Figure 2.1, considering $q_{ij} - q_{ij}^h$ leads, for example, to

$$\begin{aligned} & (\partial_{\eta} \varphi_i \partial_{\eta} \varphi_j) \int_{\Omega} \operatorname{div}((\partial_{\xi} s - \partial_{\xi} s^h, 0)) \\ &= (\partial_{\eta} \varphi_i \partial_{\eta} \varphi_j) \int_{\partial \Omega} (\partial_{\xi} s - \partial_{\xi} s^h, 0)^T \hat{n} \end{aligned} \quad (2.48)$$

$$\leq Ch^2, \quad (2.49)$$

where we have applied Taylor's theorem and multiplied by the length of the boundary integral in the last line. Calculations for the other contributions of the s_i to q_{ij}^h give similar estimates, so

$$|q_{ij} - q_{ij}^h| \leq Ch^2. \quad (2.50)$$

Now we estimate the operator error. When V^h is the space of piecewise linear polynomials, (φ_i) forms a basis, and setting $u_i = u(x_i), v_i = v(x_i)$, expressions (2.46) and (2.47) give

$$|\mathcal{B}^h(u, v) - \mathcal{B}(u, v)| = \left| \sum_{i \sim j} (q_{ij} - q_{ij}^h)(u_i - u_j)(v_i - v_j) \right| \quad (2.51)$$

$$\leq C \max_{i \sim j} |q_{ij} - q_{ij}^h| \left| \sum_{i \sim j} \gamma_{ij} (u_i - u_j)(v_i - v_j) \right| \quad (2.52)$$

$$= C |q_{ij} - q_{ij}^h|_{\infty} \int_{\Omega} \nabla u^T \nabla v \leq Ch^2 \|\nabla u\|_{L_2} \|\nabla v\|_{L_2} \quad (2.53)$$

$$\leq Ch^2 \|u\|_{H^1} \|v\|_{H^1}. \quad (2.54)$$

We use the notation $i \sim j$ to denote vertices i, j sharing edge ij of the triangle mesh. The γ_{ij} are the cotangents representing the geometry of the triangles such that the integral is exact for $u, v \in V^h$ — in \mathbb{R}^2 , the γ_{ij} are independent of h . Finally, the Poincaré inequality implies the equivalence of the h_0^1 semi-norm and the H_0^1 norm. \square

2.2.2 Operator convergence: numerical experiment

In this section, we construct $s(x)$ numerically, and use it to compute solutions to a test problem over a laminated domain. As previously alluded, we construct \mathcal{B}^h by computing $s(x)$ on a fine scale, and then resampling $s(x)$ at the desired scale h . Although $s(x)$ is convex at the finest scale, we have already noted that a piecewise linear interpolation of $s(x)$ is only convex up to edge flips. Since $s(x)$ and its interpolants belong not to linear spaces, but only to the cone of convex functions, numerical calculations for $s(x)$ are delicate. All of our numerical experiments are performed using C and C++ code compiled with gcc version 3.4.2 on a Linux platform. The hardware is a desktop machine with a dual-Pentium 3.0 GHz processor and 1.5 GB RAM.

Construction of a coarse-scale operator is in three steps. First, we compute the stiffness matrix q_{ij} from σ at a fine scale. Generally, this scale is chosen smaller than the finest scale of σ . Second, we invert equation (2.9), giving $s_i = s(x_i)$. Finally, we resample our fine-scale $s(x)$, and apply equation (2.9) to give the stiffness matrix for the downscaled operator.

The first step, computing q_{ij} from σ is the most expensive. In this step, we compute the fine-scale stiffness matrix

$$q_{ij} = \int_{\Omega} \nabla \varphi_i Q \nabla \varphi_j = \int_{\Omega} \nabla \psi_i \sigma \nabla \psi_j, \quad (2.55)$$

where we recall that the notation $\psi_i = \varphi_i \circ F$ from Chapter 1 implies that solutions produced using this stiffness matrix are the homogenised solutions, $\hat{u} = u \circ F^{-1} \in C^{1,\alpha}$. Numerically, there are two choices: 1) we could compute q_{ij} from Q , the first equality of (2.55); or 2) we could compute q_{ij} directly from σ , the second equality of (2.55). In order that the result is accurate, both choices require computing F on a scale finer than that of the triangulation for q_{ij} . We cannot avoid this fine-scale calculation since F contains all variation in σ . Numerically, we find that the second choice is easier to implement, since it does not require

explicit calculation of F^{-1} from F . This calculation can be done on a desktop machine, but only if care is taken to compute contributions to q_{ij} one element at a time. In our experiments, we choose the mesh for representing q_{ij} at a scale 7–10 times smaller than the smallest feature in σ , and choose the mesh for solving for F to 7–10 times smaller than the scale for q_{ij} . This is about the limit of what we can hope to study given the memory and speed constraints of our desktop machine.

To compute the s_i from the q_{ij} , we solve equations (2.9) by least-squares. That is, we determine

$$\operatorname{argmin} \|As - \mathbf{q}\|_{l_2}^2 + \nu \|\mathbf{s}\|_{l_2}^2. \quad (2.56)$$

Here, \mathbf{s} is the vector of s_i , and \mathbf{q} is the given vector of q_{ij} . The least-squares solve is regularised by some $\nu \ll 1$, since $s(x)$ is only unique up to the addition of affine functions, but with only mild regularisation — $\nu/\|\mathbf{s}\| \approx 10^{-4}$ in practise. Our calculations show that A is rank-deficient, having three zero singular values corresponding to the three degrees of freedom of affine functions. By this solve, for optimal \mathbf{s}^* , we also have a fine-scale stiffness matrix $\mathbf{q}^* = A\mathbf{s}^*$ that is divergence-free. As a final step, we compute the convex hull of the surface defined by (x_i, y_i, s_i) . The reason we do this is so that any interpolation we make of the fine scale surface is itself convex.

Building coarser scale stiffness matrices from a non-convex approximation of $s(x)$ has two pitfalls: 1) the matrix may not be positive-definite; and 2) the matrix may be ill-conditioned. The first problem arises due to numerical error in computing F , and then the fine-scale q_{ij} , and finally the s_i by least-squares. Due to this error, we have produced an $s(x)$ that gives an indefinite operator, even though the original σ is positive everywhere. Since a convex discrete $s(x)$ gives positive-definite stiffness matrix, taking the convex hull of the s_i can avoid this problem without changing the approximation for the s_i significantly. In practise, when we build F and q_{ij} from sufficiently fine meshes, less than 0.1% of points are culled when taking the convex hull.

Even if all points are on the convex hull, indicating that the stiffness matrix is positive-definite, the surface may have some edges with negative dihedral angles, giving the stiffness matrix positive off-diagonals, ill-conditioning the matrix. In this case, taking the convex hull of the s_i does not cull any points, but merely changes the connectivity of the mesh. The simplest such change is the repair of a single negative edge by an edge flip as we saw in

Figure 2.4. We should not in general expect the surface defined by the s_i to form a convex interpolation of $s(x)$ without taking the convex hull — the only way this could happen is if we were to anticipate the shape of $s(x)$ prior to triangulating the domain. More will be said of triangulations for given σ in the next section.

The final step in computing our coarse-scale operator is to resample $s(x)$ at a coarse scale. Again, we take the convex hull of the resampled $s(x)$ before applying equation (2.9) to compute a well-conditioned coarse-scale stiffness matrix.

As a final note on the coarse operator, we comment that there is a way to compute the s_i from the matrix Q . Recalling our construction of $s(x)$ through the intermediate integrals of h and k , we have

$$Q = \begin{pmatrix} Q_{11} & Q_{12} \\ Q_{12} & Q_{22} \end{pmatrix} = \begin{pmatrix} \partial_y h & -\partial_y k \\ -\partial_x h & \partial_x k \end{pmatrix} = \begin{pmatrix} \partial_{yy} s & -\partial_{xy} s \\ -\partial_{xy} s & \partial_{xx} s \end{pmatrix}. \quad (2.57)$$

Given Q , this gives equations

$$\Delta h = \partial_y Q_{11} - \partial_x Q_{12}, \quad (2.58)$$

$$\Delta k = -\partial_y Q_{12} + \partial_x Q_{22}, \quad (2.59)$$

$$\Delta s = \partial_y k + \partial_x h, \quad (2.60)$$

and we can numerically solve for h, k and then s . We can only solve for $s(x)$ weakly since even at the finest scale we approximate $s(x)$ by piecewise linear functions, but this is an explicit construction. Furthermore, by applying Green's theorem to the equations for h and k , we only need the entries of Q piecewise constant on triangles. Unfortunately, in practise, this method fails to produce an $s(x)$ that is remotely convex: barely 30% of the (x_i, y_i, s_i) produced in this manner lie on their convex hull. It may be that we are not discretizing the three problems in a consistent manner, or that Q does not average over triangles as well as it does as a stiffness matrix. Comparing the success of our two approaches, we are content to use the former in our codes, however, returning to this method may be fruitful in a study of the problem in \mathbb{R}^3 .

Our numerical experiment in computing $s(x)$ begins with the laminated scalar conductivity of Figure 2.5. From this laminated domain, we compute the fine-scale $s(x)$, projected

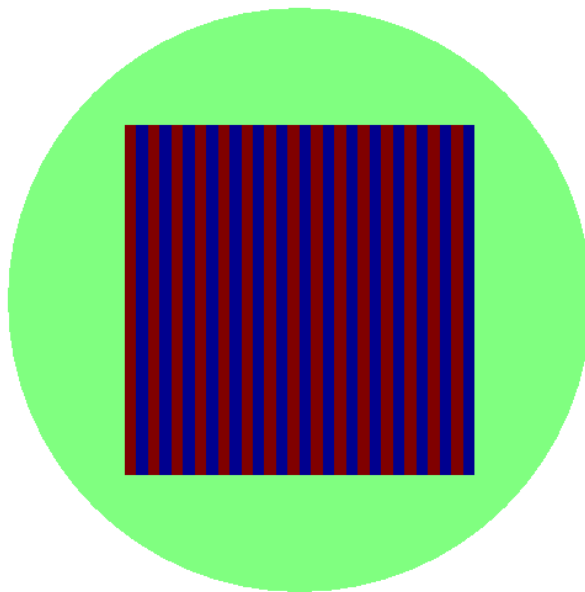


Figure 2.5: Laminated domain used in numerical experiments. Alternating blue and red strips have conductivity 0.05 and 1.95, respectively, and the surrounding green area has conductivity 1. The circle has radius 1, and the strips each have width 0.04.

in Figure 2.6. Note that unlike asymptotic homogenisation, the fine-scale structure of the conductivity is visible at the finest scale of $s(x)$. However, the coarse-scale anisotropy is also apparent: the curvature is larger parallel to the x -axis, consistent with a larger conductivity in the y -direction. (Recall that the up-scaled conductivity Q is the *rotated* Hessian of $s(x)$.) Figure 2.7 shows the homogenising effect of down-scaling $s(x)$: the fine detail reflecting the laminations in σ are not visible, but the anisotropy represented by the anisotropic curvature remains.

Having a method for computing the linear interpolant $s^h(x)$ over many scales h , we turn now to a test of convergence of solutions to $\mathcal{B}^h(u^h, v) = L(v), v \in V^h$ as $h \rightarrow 0$. We choose the test problem

$$\begin{aligned} -\operatorname{div}(Q\nabla u) &= f & x \in \Omega, \\ u &= 0, & x \in \partial\Omega, \end{aligned} \tag{2.61}$$

$$\begin{aligned} f &= \exp[-10((x - 0.4)^2 + (y - 0.4)^2)] \\ &\quad - \exp[-10((x + 0.4)^2 + (y + 0.4)^2)]. \end{aligned} \tag{2.62}$$




Figure 2.6: Two views of the fine-scale $s(x)$ surface for the laminated domain. The left-hand view reflects the fine-scale pattern in σ , and the right-hand view highlights the coarse-scale anisotropy in the curvature.




Figure 2.7: Two views of the $s(x)$ surface resampled and linearly interpolated at a coarse scale. While the left-hand view no longer shows the fine-scale pattern in σ , the right-hand view still shows the overall anisotropy of the up-scaled operator.




Figure 2.8: A comparison of homogenisation solutions at coarse and fine scales. The figure on the left shows a fine-scale solution; note in particular the undulations at the peak due to the laminations in σ . The coarse-scale solution on the right is much smoother.

Note that Q is the homogenised σ , and we are computing the homogenisation solution u rather than solution to the original problem with conductivity σ . The reason for this is that the solution to the unhomogenised problem is not smooth in general, and up-scaling the non-smooth solution is unjustified. However, the solution to the homogenised problem is $C^{1,\alpha}$, and its coarse-scale averages do approximate the fine-scale homogenised solution faithfully.

For our test problem, we do not have an exact solution and can only compute a reference solution u_{ref} , over our finest scale for q_{ij} . In this case, our reference solution is over a polygonal domain with 512 boundary points on the unit circle; with a uniform mesh this corresponds to $h \approx 0.012$. This reference solution and a coarse-scale solution on a mesh having $h \approx 0.070$ are shown in Figure 2.8. Convergence of the sequence of solutions for h decreasing to the finest numerical scale is shown by Table 2.1 and Figure 2.9.

Note that we observe convergence in H_0^1 although the convergence is only $\mathcal{O}(h^{0.43})$. This is suggestive of Hölder continuity of the first derivative with exponent $\alpha \approx 0.43$. Figure 2.10 shows a rendered close-up of the reference solution, on which we see a lack of smoothness at the lamination boundaries, a nice depiction of a solution that is C^1 , but almost certainly

V_B	N	h	$\ e\ _{L_2}$	$\ e\ _{L^\infty}$	$ e _{h^1}$	$\ \nabla e\ _{L^\infty}$
32	85	3.13×10^{-2}	3.57×10^{-3}	7.59×10^{-3}	3.63×10^{-2}	1.44×10^{-3}
45	190	2.22×10^{-2}	1.70×10^{-3}	4.62×10^{-3}	3.01×10^{-2}	1.32×10^{-3}
64	389	1.56×10^{-2}	1.02×10^{-3}	3.16×10^{-3}	2.81×10^{-2}	1.24×10^{-3}
90	872	1.11×10^{-2}	6.36×10^{-4}	1.97×10^{-3}	2.50×10^{-2}	1.11×10^{-3}
128	1631	7.81×10^{-3}	4.21×10^{-4}	1.62×10^{-3}	2.12×10^{-2}	1.12×10^{-3}
180	2616	5.56×10^{-3}	3.08×10^{-4}	1.31×10^{-3}	1.88×10^{-2}	8.58×10^{-4}
256	5172	3.91×10^{-3}	2.03×10^{-4}	7.38×10^{-4}	1.50×10^{-2}	7.43×10^{-4}
360	8099	2.78×10^{-3}	1.73×10^{-4}	8.74×10^{-4}	1.29×10^{-2}	1.05×10^{-3}

Table 2.1: Convergence of an approximated operator. Test cases are indexed by the V_B , the number of boundary points. N is the number of vertices in the mesh, h is the mesh size, and error $e = \tilde{u}^h - u_{\text{ref}}$. u_{ref} is the reference solution on a circle having 512 boundary points. This data is plotted in Figure 2.9.




Figure 2.9: Plot of the convergence of solutions to the reference solution as given in Table 2.1. The observed convergence rates are approximately: $\|e\|_{L_2} \sim \mathcal{O}(h^{1.2})$, $\|e\|_{L^\infty} \sim \mathcal{O}(h^{1.0})$, $|e|_{h^1} \sim \mathcal{O}(h^{0.43})$, and $\|\nabla e\|_{L^\infty} \sim \mathcal{O}(h^{0.23})$.




Figure 2.10: Detail of the reference solution near the peak at $(x, y) = (0.4, 0.4)$. This rendering suggests that although flipping edges could make the solution appear smoother — a problem with representing any solution on a piecewise linear mesh — there are irregularities in the solution at the lamination boundaries which affect convergence more so than the approximation of $s(x)$.

not C^2 .

Finally for this experiment, we look at the direct difference between the reference solution and a coarse-scale solution. The error pattern in Figure 2.11 shows how up-scaling the operator and coarsening the mesh limits the ability to resolve the fine-scale structure of the solution. Figure 2.11 also shows how the laminations are distorted by the harmonic change of coordinates F , just as we saw in the example calculation in Chapter 1.

A nice application of this down-scaling would be to multi-grid methods, where a hierarchy of operators is needed to provide rapidly convergent coarse-scale corrections to the slowly convergent fine-scale problem [28]. Following the multi-grid literature, set A^h the stiffness matrix at scale h , P_{2h}^h the prolongation operator, which takes coarse-scale solutions to approximate fine-scale solutions, and R_h^{2h} the restriction operator, which decimates fine-scale solutions to coarse-scale approximations. In such case, two conditions are generally




Figure 2.11: Difference between $V_B = 256$ and $V_B = 512$, showing the fine-scale structure not captured by the coarser mesh. The distortion of the laminations by F is also apparent.

demanded to ensure consistency between scales, lending stability to the method as a whole:

$$A^{2h} = R_h^{2h} A^h P_{2h}^h, \quad (2.63)$$

$$R_h^{2h} = c(P_{2h}^h)^T, \quad c \text{ constant.} \quad (2.64)$$

The question is, by changing the scale of the operator (the stiffness matrix) as we have done by resampling $s(x)$, do we produce a hierarchy of stiffness matrices satisfying the first of these conditions? It seems that convergence of the operators should be sufficient, but the multi-grid condition is much more specific, applying in particular to hierarchical meshes, where our family of ever-finer meshes need not bear any relation to each other. For the moment, we do not have an answer to this question.

2.3 Well-adapted triangulations

In addition to being able to resolve the elliptic operator on a variety of scales, it is typically useful to examine the quality of the discretized operator. In particular, this section focuses on adapting the triangle mesh to the anisotropy of the up-scaled operator. If Q is highly anisotropic and a regular mesh is used to discretize the domain, the resulting stiffness matrix can be ill-conditioned, and the interpolation of the solution can be adversely affected. Fortunately, while the problem's physical parameters may be given to us as σ , we often have a choice about how we build our computational mesh, and good mesh design can give us well-conditioned stiffness matrices and good interpolation properties. A good introduction to this problem is given by Shewchuk in his review of finite element quality [77]. One method for constructing anisotropic meshes is given in [57], and assumes that the anisotropy of the problem is given as a map from the problem domain to an isotropic space. Our method works directly with the conductivity to produce meshes well-adapted to the anisotropy of Q .

Figure 2.12 compares an isotropic to an anisotropic mesh. The isotropic mesh is suitable for the basic problem where $Q = \text{Id}$, and the anisotropic mesh for the case where the conductivity is greater in the y -direction than in the x -direction. This case corresponds to $Q_{11} \ll Q_{22} \Rightarrow \partial_{xx}s \gg \partial_{yy}s$. The mesh anisotropy improves both interpolation of the solution and conditioning of the stiffness matrix.

We can see the effect on interpolation from a physical argument: potentials change rapidly when conductivities are low, thus demanding increased resolution in this case, in the x -direction. Alternatively, recalling that Q is divergence-free, we have from the strong statement of the original problem,

$$-\text{div}(Q\nabla u) = -\partial_{xx}u\partial_{yy}s + 2\partial_{xy}u\partial_{xy}s - \partial_{yy}u\partial_{xx}s = f. \quad (2.65)$$

If this expression is approximately zero the curvature of u is large in directions where the curvature of $s(x)$ is small. Recalling that Q is the 90-degree rotated Hessian of $s(x)$ recovers the intuition that u should be better resolved orthogonal to the principal curvature of $s(x)$.

The improvement in conditioning can be seen by recalling our goal of making off-diagonal elements of the stiffness matrix as negative as possible. Since we build the stiffness matrix




Figure 2.12: Comparison of an isotropic and an anisotropic mesh. The figure on the left shows the lack of directional bias expected for a mesh suitable for the isotropic problem, while the figure on the right is suitable for the case where the conductivity is greater in the y -direction than in the x -direction.

by interpolating the surface $s(x)$, this means building a mesh that best interpolates $s(x)$ has its finest resolution in the direction where the second derivative of $s(x)$ is greatest. Looking again at expression (2.65), if $f \approx 0$, and, for illustration, the principal curvatures of $s(x)$ are parallel to the axes, so $\partial_{xy}s = 0$, $|\partial_{xx}s| \gg |\partial_{yy}s|$ is consistent with $|\partial_{xx}u| \gg |\partial_{yy}u|$. A mesh that is suitable for interpolating $s(x)$ and providing a well-conditioned operator is also a suitable candidate for interpolating the solution.

That the optimal mesh for interpolation coincides with that for well-conditioned stiffness matrices is not always the case. An example in [77] shows how the choice of boundary conditions can void this correspondence. In our case, the choice of source function f can give this problem, but predicting the best interpolating mesh from source data departs from our concern with the elliptic operator itself.

We propose a method to build meshes directly from our parametrisation $s(x)$ of the anisotropic scalable conductivity. Two ingredients contribute to quality meshes: the connectivity of the mesh vertices, and the placement of those vertices. The final algorithm for mesh generation iterates over two steps: 1) Select the optimal connectivity for the given vertex locations; 2) Select the optimal vertex locations for the given connectivity. By ex-




Figure 2.13: One-dimensional illustration of the convex hull construction of the Delaunay triangulation. Points x_a, x_b are given points, projected onto the surface $s(x)$. The slope of the interpolation $s^h(x)$ equals the slope of $s(x)$ at x_c , analogous to the circumcentre of a triangle in a two-dimensional construction.

Examining how these issues are addressed in isotropic problems, we will apply those same principals to use $s(x)$ in our anisotropic mesh generation.

2.3.1 Mesh connectivity by weighted Delaunay triangulations

For both interpolation and matrix conditioning, a reliable connectivity criterion is that the triangulation be Delaunay. This means that the circumcircle of each triangle in the mesh contains no other vertex of the mesh. (There are several equivalent definitions of a Delaunay triangulation, but this one is directly useful for our purposes.) Given a set of points, we can construct its Delaunay triangulation geometrically. We project points given in the xy -plane up to the paraboloid $z = \frac{1}{2}x^2 + \frac{1}{2}y^2$. We then take the convex hull of these points which defines a polyhedron having triangular faces. These triangles, when projected back to the xy -plane, give the Delaunay triangulation of the original points. O'Rourke gives an illustrated discussion of this construction [63, pp 182–190]. Figure 2.13 shows a one-dimensional illustration of this construction.

The convex hull construction uses exactly the points that interpolate $s(x, y)$ corresponding to the unit constant isotropic conductivity, that is, the discretization of the Laplacian operator. Furthermore, in taking the convex hull of these points, all edges of the interpolated paraboloid are convex by construction, and all off-diagonal entries of the stiffness matrix for the Laplacian on this mesh, given by the cotangent formula, are negative. This explains why quality triangulations should at least be Delaunay. Referring back to Figure 2.4, it is possible for the convex paraboloid $z = \frac{1}{2}x^2 + \frac{1}{2}y^2$ to have a non-convex interpolation

if the connectivity of the mesh is non-Delaunay. With non-convex edges, a non-Delaunay triangulation gives a discrete Laplacian with positive off-diagonal entries. Even though the matrix as a whole is positive-definite, having just a few positive off-diagonal entries gives the stiffness matrix a few large eigenvalues, increases the spectral condition number, and ill-conditions the matrix [77, pp 31–32, 43]. The main point here is that Delaunay triangulations give stiffness matrices such that

$$-\int_{\Omega} \nabla \varphi_i^T \nabla \varphi_j > 0, \quad i \sim j, \quad (2.66)$$

which in turn give the best-conditioned stiffness matrix for the given vertices [77].

In the case of anisotropic media, we seek the same property, that

$$q_{ij} = -\int_{\Omega} \nabla \varphi_i^T Q \nabla \varphi_j > 0, \quad i \sim j. \quad (2.67)$$

Hence, instead of creating Delaunay meshes, we build meshes that are Q -Delaunay by way of our $s(x)$ parameterisation. That is, instead of projecting our set of points up to a paraboloid, we project them to the $s(x)$ surface, take the convex hull of these points, and use the projection of the resulting polyhedron to the xy -plane as our Q -Delaunay triangulation. Since we have taken the convex hull of the points over the $s(x)$ surface, each hinge has a convex dihedral angle and each off-diagonal element of the stiffness matrix will be negative. Building a triangulation by way of the convex hull isn't expensive — there is an $\mathcal{O}(N^2)$ algorithm for N vertices — but there is a way that we can build our Q -Delaunay triangulation as a weighted Delaunay triangulation, and hence take advantage of faster and less memory-intensive algorithms.

A weighted Delaunay triangulation assigns to each vertex a scalar quantity, w_i , the *weight* of the vertex. In a standard Delaunay triangulation, connecting the circumcentres of triangles determines the dual *Voronoi* diagram for the triangulation. If the triangulation is Delaunay, each polyhedron of the Voronoi diagram contains exactly one vertex of the triangulation. In a weighted Delaunay triangulation, the same is true, but the weights act to move the weighted Voronoi centre of each triangle away from its circumcentre, and the connectivity of the original points changes accordingly.

Figure 2.14 shows the geometry of a weighted Delaunay triangle. We interpret the




Figure 2.14: The figure on the left shows the geometry of a triangle in a weighted Delaunay triangulation, showing the shift of the dual edges away from the midpoint of the primal edges, and the subsequent shift of the Voronoi centre of the triangle. The right-hand figure shows an entire hinge, identifying the primal and dual edge lengths at the hinge.

weights w_i as the square of the radii of circles centred at the triangle vertices. (There is also an interpretation for the case where $w_i < 0$; see [68].) The associated dual Voronoi edges lie not at the midpoints of the edges, as they would in a standard Delaunay triangulation, but are the lines equidistant from the circles at each primal edge's vertices. The intersection of the weighted Voronoi edges is the weighted centre of the triangle. Hence, a weighted Delaunay triangulation having equal weights on all vertices is a standard Delaunay triangulation. Our goal is to compute weights for the triangulation such that all $q_{ij} > 0$, just as the regular Delaunay triangulation provides for the isotropic case.

Glickenstein [42] examines the discrete Dirichlet energy on a weighted Delaunay triangulation. As is standard, he defines the discrete Dirichlet energy

$$E(u) = \frac{1}{2} \sum_{i \sim j} \frac{|ij^*|}{|ij|} (u_i - u_j)^2, \quad (2.68)$$

where the edge-wise values $\frac{|ij^*|}{|ij|}$ are the ratios of the dual to primal edge lengths at each hinge: see Figure 2.14. More generally, the Dirichlet energy associated to conductivity Q is given by

$$E(u) = \frac{1}{2} \int_{\Omega} \nabla u^T Q \nabla u, \quad (2.69)$$

and interpolating u in the piecewise linear basis, we have the discrete energy (2.68) provided

we can make the identification

$$q_{ij} = \frac{|ij^*|}{|ij|}. \quad (2.70)$$

If we choose the dual edges according to their positions dictated by the weights of the triangulation, $\frac{|ij^*|}{|ij|} > 0$ by definition. The magic is that for given s_i , we can always find those weights such that the identification (2.70) holds, and hence so $q_{ij} > 0$ for all edges of the mesh.

The calculation by Glickenstein shows that at hinge $ijkl$ of a weighted triangulation,

$$\begin{aligned} \frac{|ij^*|}{|ij|} &= \frac{1}{2} (\cot \theta_{ikj} + \cot \theta_{ilj}) \\ &+ \frac{1}{2|ij|^2} (\cot \theta_{ijk} + \cot \theta_{ijl}) w_i + \frac{1}{2|ij|^2} (\cot \theta_{jik} + \cot \theta_{jil}) w_j \\ &- \frac{1}{4A_{ijk}} w_k - \frac{1}{4A_{ijl}} w_l. \end{aligned} \quad (2.71)$$

We identify the first two terms with the edge ratio for the Laplacian, where $Q = \text{Id}$. The remaining terms are half of our value for q_{ij} , where the w_i replace the s_i . This is to say that when u is a piecewise linear interpolation,

$$E(u) = \frac{1}{2} \int_{\Omega} \nabla u^T \left(\text{Id} - \frac{1}{2} Q_w \right) \nabla u, \quad (2.72)$$

where the subscript w defines

$$Q_w = \begin{pmatrix} \partial_{yy} w & -\partial_{xy} w \\ -\partial_{xy} w & \partial_{xx} w \end{pmatrix}, \quad (2.73)$$

just as our conductivity Q is built from derivatives of $s(x)$. Hence, choosing w so that $Q = \text{Id} - \frac{1}{2} Q_w$ will give us $q_{ij} > 0$. Since we use second derivatives of $s(x)$ to compute Q , this implies

$$s(x, y) = \frac{1}{2} x^2 + \frac{1}{2} y^2 - \frac{1}{2} w(x, y), \quad (2.74)$$

so choosing weights

$$w_i = x_i^2 + y_i^2 - s_i \quad (2.75)$$

gives a triangulation on which $q_{ij} > 0$ on each edge. For a given set of points, this is equivalent to the construction over the convex hull of $s(x)$, and gives the best-conditioned

stiffness matrix we could hope for.

A remarkable fact about the convex hull construction is that in the isotropic case, the (x, y) -projected coordinates of the circumcentre of each Delaunay triangle are at the point where the slope of the paraboloid equals the slope of the triangle interpolating the paraboloid (again, see Figure 2.13 [63]). This fact, we believe, is the ultimate reason that equation (2.9) exactly gives the cotangent formula for $s(x, y) = \frac{1}{2}x^2 + \frac{1}{2}y^2$, even though $s^h(x)$, the interpolation of $s(x)$, is only piecewise linear.

2.3.2 Optimising vertex locations

Choosing appropriate triangulations alone improves the conditioning of stiffness matrices, but by exercising our freedom in choosing vertex locations, we can further improve upon the conditioning and interpolation properties of our discretization. Again, there is a way to use $s(x)$ directly to this end, and the method of Alliez et al for producing quality isotropic meshes suggests the correct extension to the anisotropic case [15].

In the isotropic case [15] seeks the arrangement of vertices (both connectivity and location) minimising the energy

$$E = \|p - p^h\|_{L_1}, \quad (2.76)$$

where $p(x, y) = \frac{1}{2}x^2 + \frac{1}{2}y^2$ and $p^h(x, y)$ is its piecewise linear interpolating polynomial. In one dimension, the relationship between p and p^h is exactly the same as that depicted between $s(x)$ and s^h in Figure 2.13. The one-dimensional picture is slightly misleading in that while it does illustrate the freedom to move the interpolation knots, it does not show the freedom in choosing the connectivity that we have in higher dimensions. (In one dimension, we can show that equispaced knots minimise (2.76).) We replace p, p^h with s, s^h , and so seek the connectivity and vertex positions that minimise

$$E = \|s - s^h\|_{L_1}. \quad (2.77)$$

As in [15], we propose a two step iterative process. First, for a given set of points, we compute its optimal connectivity as its Q -weighted Delaunay triangulation. Second, we adjust the points towards the minimum E for that triangulation. We have already discussed the Delaunay step, and the remainder of this section develops the adjustment step, tailored

to our given Q .

Since $s^h \geq s$, the energy

$$E = \|s - s^h\|_{L_1} = \int_{\Omega} s^h - s \quad (2.78)$$

$$= \sum_i s(x_i) \int_{\Omega} \varphi_i(x) - \int_{\Omega} s(x) \quad (2.79)$$

$$= \frac{1}{3} \sum |\Omega_i| s(x_i) - \int_{\Omega} s(x). \quad (2.80)$$

We have used the linear basis \mathcal{P}_1^h spanned by the hat functions (φ_i) for s^h , and identified $|\Omega_i|$ as the area of all triangles incident to vertex i . Minimising the energy for this connectivity means setting $\nabla_{x_i} E = 0$ for all vertex locations x_i . The integral of $s(x)$ is constant, so we have

$$\nabla_{x_i} E = \frac{|\Omega_i|}{3} \nabla s|_{x=x_i} + \frac{1}{3} \sum_{T_j \in \Omega_i} \left(\nabla_{x_i} |T_j| \left[\sum_{\substack{k \in T_j \\ k \neq i}} s(x_k) \right] \right), \quad (2.81)$$

where in addition to observing how E changes with respect to $s(x)$, we also need to include the contribution due to changing the area $|T_j|$ of each triangle incident to vertex i . $k \in T_j$ refers to the set of vertices of triangle T_j . A straightforward computation shows that for positively-oriented triangle ijk , and corners with coordinates $\{(x_i, y_i), (x_j, y_j), (x_k, y_k)\}$,

$$\begin{aligned} \partial_{x_i} |T_{ijk}| &= \frac{1}{2} (y_j - y_k), \\ \partial_{x_i} |T_{ijk}| &= \frac{1}{2} (x_k - x_j), \end{aligned} \quad (2.82)$$

which is only dependent on vertex positions.

At the optimal vertex positions, $\nabla_{x_i} E = 0$. We examine two cases, first where the conductivity is constant, and second where it varies over the domain. In both cases, we use (2.81) to compute a greedy update rule for the vertex positions.

In the case $Q = \text{constant}$, we have

$$s(x, y) = \frac{Q_{22}}{2} x^2 - Q_{12} xy + \frac{Q_{11}}{2} y^2, \quad (2.83)$$

$$\nabla s = \begin{pmatrix} Q_{22}x - Q_{12}y \\ -Q_{12}x + Q_{11}y \end{pmatrix} = H_s x, \quad (2.84)$$

where H_s is the Hessian of $s(x)$. This means that within each one-ring, vertices have optimal position x_i^* satisfying the 2×2 linear system

$$H_s x_i^* = -\frac{1}{|\Omega_i|} \sum_{T_j \in \Omega_i} \left(\nabla_{x_i} |T_j| \left[\sum_{\substack{k \in T_j \\ k \neq i}} s(x_k) \right] \right). \quad (2.85)$$

Following [15], we use this expression to build an update rule by shifting all vertices by the current vertex position x_i .

In the constant case, the interpolation error is unaffected if we shift $s(x)$ by constant value x_i . To see this, compute

$$s(x - x_i) = \frac{Q_{22}}{2}(x - x_i)^2 - Q_{12}(x - x_i)(y - y_i) + \frac{Q_{11}}{2}(y - y_i)^2 \quad (2.86)$$

$$= \frac{Q_{22}}{2}x^2 - Q_{12}xy + \frac{Q_{11}}{2}y^2 + ex + fy + g \quad (2.87)$$

$$= s(x, y) + ex + fy + g \quad (2.88)$$

for constants e, f, g . Since $s^h(x), s^h(x - x_i)$ are piecewise linear, and interpolate $s(x), s(x - x_i)$, respectively, we have that

$$s^h(x - x_i) = s^h(x) + ex + fy + g \quad (2.89)$$

holds exactly for the same e, f, g , so

$$\int_{\Omega} s^h(x - x_i) - s(x - x_i) = \int_{\Omega} s^h(x) - s(x). \quad (2.90)$$

Hence, updating vertex position x_i to position x_i^* means solving

$$H_s x_i^* = H_s x_i - \frac{1}{|\Omega_i|} \sum_{T_j \in \Omega_i} \left(\nabla_{x_i} |T_j| \left[\sum_{k \in T_j} s(x_k - x_i) \right] \right), \quad (2.91)$$

where we have $s(0) = 0$. Setting $s(x) = \frac{1}{2}x^2 + \frac{1}{2}y^2$ recovers the update rule in [15] for isotropic meshes. At each update step, we iterate over the entire mesh, computing $dx_i =$

$x_i^* - x_i$. We then adjust each vertex,

$$x_i^{(n+1)} = x_i^{(n)} + \alpha dx_i^{(n)}, \quad (2.92)$$

where n is the iteration number, and α is a relaxation parameter. In practise, even for constant Q , we find that for particularly anisotropic conductivities, choosing $\alpha \approx \frac{1}{2}$ stabilises the iteration. The need for some relaxation is not surprising, since we are using a greedy algorithm to adjust the vertex points, and on anisotropic meshes where triangles get long and thin it doesn't take much to move a vertex far from its original one-ring.

When updating x_i in the case of non-constant Q , we observe that the interpolation error is not invariant under translation, and we cannot use (2.91). Instead, we seek x_i^* satisfying

$$H_s(x_i^*) = -\frac{1}{|\Omega_i|} \sum_{T_j \in \Omega_i} \left(\nabla_{x_i} |T_j| \left[\sum_{\substack{k \in T_j \\ k \neq i}} s(x_k) \right] \right). \quad (2.93)$$

The system for x_i^* is no longer a linear system, since the Hessian of $s(x)$ does not vary linearly in x . The solution is to linearise the Hessian around the current vertex location x_i and to further relax the vertex movement such that in any given timestep, no vertex leaves its Voronoi area. This slows convergence, but prevents inversion of triangles in the mesh. The inversion of triangles isn't necessarily destabilising, as the weighted Delaunay step returns the points to the convex hull of $s(x)$, but in practise we see that if vertices move too far from their original positions, stability is adversely affected.

Thus far, we have implied that we begin with an arrangement of points, and then perform weighted Delaunay and vertex adjustment steps until convergence. This method works, but convergence can be slow as the connectivity can become trapped far from its optimal value with a complete complement of points. Instead, we begin with a full set of boundary points plus a similar number of interior points randomly distributed throughout the domain. With each iteration, we add new points to the domain in the neighbourhood of badly formed or excessively large triangles. The precise locations where these points are introduced is flexible, as the vertex adjustment step tends to move them towards their optimal positions immediately. Triangles smaller than a preset-set minimum area are not

refined, as allowing this lets the algorithm refine the mesh indefinitely. Point addition is not sensitive to the choice of quality measure, so long as an anisotropic quality measure is used. [77, pp 53–54] lists several quality measures, of which we use

$$\text{quality} = \frac{\tilde{A}_T}{\tilde{l}_1^2 + \tilde{l}_2^2 + \tilde{l}_3^2}. \quad (2.94)$$

A_T is triangle area, and l_1, l_2, l_3 are the lengths of its three sides. The tilde indicates that the measurements are made on the triangle pulled back to isotropic space. We use measure (2.94), as it is scale independent and bounds the condition number of the triangle stiffness matrix.

Original seed boundary points are not moved. Moreover, any vertices sent outside the domain by the vertex adjustment step are projected back to the domain boundary, where they are allowed to move either back into the domain or along the boundary as the iteration proceeds.

Over a small number of iterations (12 is typical) the algorithm converges, whereupon the vertex adjustment step moves the vertices very little, the weighted Delaunay triangulation doesn't change the connectivity in more than a few places, and no new points are added after vertex adjustment. In the numerical experiments, after these first few steps, we see minimal improvement to the conditioning of the stiffness matrix or to the interpolation properties of the mesh. Algorithm 2.1 summarises the complete adaptive algorithm.

Algorithm 2.1 Build Q -adapted mesh

Read the input boundary points
 Select initial interior points
 Construct Q -weighted Delaunay triangulation
repeat
 Adjust vertices
 Add new points
 Construct Q -weighted Delaunay triangulation
until converged

2.3.3 Numerical results

In order to clearly see the behaviour of our iteration, we perform our numerical tests for the simple case that Q is anisotropic but constant over the domain. Our test domain is again

the unit disk, with

$$Q = \begin{pmatrix} 0.1 & 0 \\ 0 & 10 \end{pmatrix} \quad (2.95)$$

corresponding to $s(x, y) = \frac{1}{2}10x^2 + \frac{1}{2}0.1y^2$. This is the conductivity used to produce the adapted mesh already seen in Figure 2.12.

Figures 2.15, 2.16 and 2.17 plot comparisons of the interpolation quality in the L_2 -norm, the interpolation quality in the h_1 -semi-norm, and the conditioning of the stiffness matrices. Interpolation quality is judged by how well the mesh interpolates $s(x, y)$. That is, if s^h is the piecewise linear interpolant of $s(x)$, then we measure

$$\|e_s\|_{L_2} = \|s^h - s\|_{L_2}, \quad (2.96)$$

$$|e_s|_{h^1} = \|\nabla s^h - \nabla s\|_{L_2}. \quad (2.97)$$

As already mentioned, this is an optimistic measure of interpolation quality. Without further knowledge of the problem, however, $s(x)$ is a reasonable representative of the problem's anisotropy. See Table 2.2 for the interpolation error. We also measure the spectral condition number κ_2 of the stiffness matrices for Q consequential to the meshes, as the ratio of largest to smallest eigenvalues of a matrix is typically associated with the difficulty of solving its resulting linear systems. See Table 2.3 for the condition number data. In all cases, we index the test meshes by the number of boundary vertices, V_B . For fixed V_B , the isotropic and adapted anisotropic meshes tend to give comparable interpolation quality, while the isotropic meshes tend to require far more vertices than do the adapted meshes. See Figure 2.18 for this comparison. The isotropic meshes are generated by Shewchuk's Triangle program [75].

We measure both interpolation and conditioning quality with respect to the number of vertices, N , of the mesh, rather than with respect to h , the circumradius of the largest triangle in the mesh. Although $N \sim h^{-2}$ for both isotropic and anisotropic meshes, using circumradius gives the appearance that the long, thin triangles of the anisotropic meshes will appear much larger than the regular triangles of the isotropic meshes. One way to standardise the comparison is to perform measurements of the anisotropic triangles in isotropic space, deforming the anisotropic triangles by a function representing the anisotropy of Q .

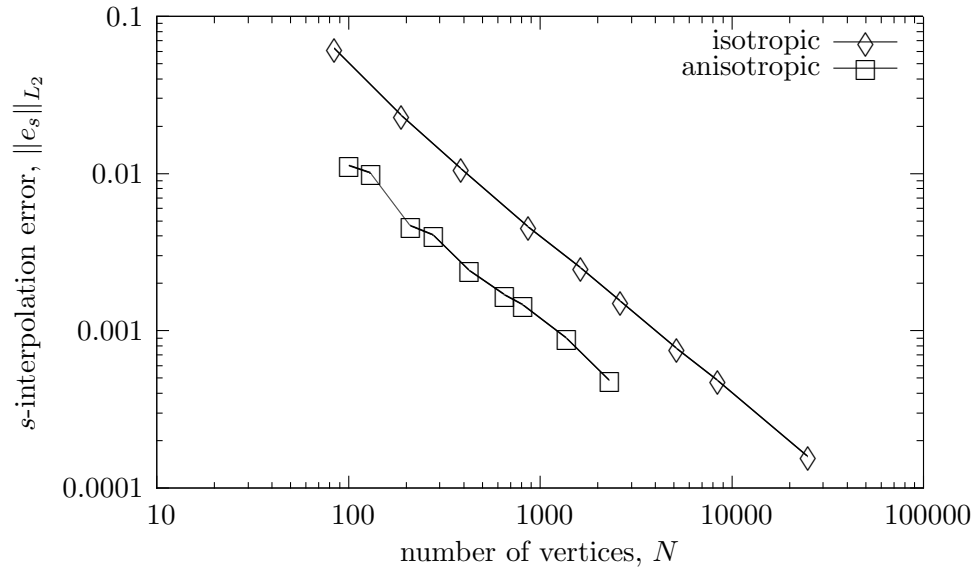


Figure 2.15: Interpolation quality of adapted meshes measured by the L_2 -norm error in a linear interpolation of $s(x, y)$. Error diminishes as $\mathcal{O}(N^{-1})$ in both cases, but is offset by a factor of about 4 in the adapted anisotropic meshes.




Figure 2.16: Interpolation quality of adapted meshes measured by the h^1 -semi-norm error in a linear interpolation of $s(x, y)$. Error diminishes as $\mathcal{O}(N^{-\frac{1}{2}})$ in both cases, but is offset by a factor of about 3 in the adapted anisotropic meshes.




Figure 2.17: Matrix conditioning quality of adapted meshes measured by the spectral condition number κ_2 of the stiffness matrix. The condition number grows as $\mathcal{O}(N)$ in both cases, but is offset by a factor of about 5 in the adapted anisotropic meshes.

V_B	isotropic			anisotropic		
	N	$\ e_s\ _{L_2}$	$ e_s _{h^1}$	N	$\ e_s\ _{L_2}$	$ e_s _{h^1}$
32	85	6.27×10^{-2}	1.25	101	1.13×10^{-2}	3.43×10^{-1}
45	190	2.34×10^{-2}	7.42×10^{-1}	131	1.01×10^{-2}	3.11×10^{-1}
64	389	1.07×10^{-2}	5.04×10^{-1}	212	4.68×10^{-3}	2.16×10^{-1}
90	873	4.59×10^{-3}	3.28×10^{-1}	280	4.04×10^{-3}	1.94×10^{-1}
128	1631	2.52×10^{-3}	2.43×10^{-1}	428	2.43×10^{-3}	1.52×10^{-1}
180	2633	1.55×10^{-3}	1.91×10^{-1}	656	1.69×10^{-3}	1.21×10^{-1}
256	5201	7.72×10^{-4}	1.35×10^{-1}	817	1.46×10^{-3}	1.13×10^{-1}
360	8453	4.81×10^{-4}	1.07×10^{-1}	1380	8.97×10^{-4}	8.89×10^{-2}
512	25019	1.58×10^{-4}	6.11×10^{-2}	2312	4.83×10^{-4}	6.56×10^{-2}

Table 2.2: Improvement of interpolation by meshes adapted to an anisotropic conductivity. The test meshes are indexed by their number of boundary vertices, V_B . Other values are: N , the number of vertices; $\|e_s\|_{L_2}$, the L_2 -norm error in interpolating $s(x)$; and $|e_s|_{h^1}$, the h^1 -semi-norm error in interpolating $s(x)$.

V_B	isotropic		anisotropic	
	N	κ_2	N	κ_2
32	85	40	101	6.0
45	190	99	131	6.7
64	389	220	212	15
90	873	534	280	18
128	1631	1037	428	26
180	2633	1721	656	35
256	5201	3551	817	46
360	8453	5974	1380	71
512	25019	17849	2312	132

Table 2.3: Improvement of matrix conditioning by meshes adapted to an anisotropic conductivity. The test meshes are indexed by their number of boundary vertices, V_B . Other values are: N , the number of vertices; and κ_2 the spectral condition number of the stiffness matrix.

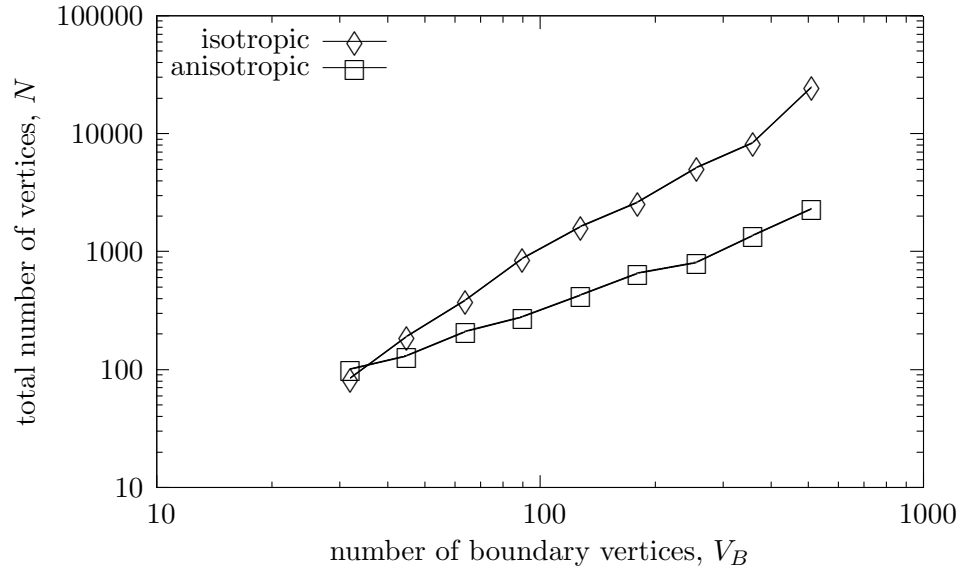


Figure 2.18: Comparison of the growth of the total number of vertices with respect to the number of boundary vertices. This is particularly interesting for the problem of computing boundary flux data from boundary value data, where the size of the problem is determined by the boundary resolution. For the isotropic mesh, $N \sim V_B^{1.9}$, and for the anisotropic mesh, $N \sim V_B^{1.1}$.

V_B	isotropic			anisotropic		
	N	λ_{\min}	λ_{\max}	N	λ_{\min}	λ_{\max}
32	85	1.19	48.1	101	1.16	7.01
45	190	5.20×10^{-1}	51.6	131	0.964	6.46
64	389	2.40×10^{-1}	52.9	212	0.504	7.46
90	873	9.92×10^{-2}	53.0	280	0.411	7.37
128	1631	5.40×10^{-2}	56.0	428	0.310	8.17
180	2633	3.30×10^{-2}	56.8	656	0.214	7.53
256	5201	1.67×10^{-2}	59.3	817	0.179	8.24
360	8453	9.91×10^{-3}	59.2	1380	0.124	8.82
512	25019	3.44×10^{-3}	61.4	2312	0.0586	7.71

Table 2.4: Comparison of the eigenvalues of the stiffness matrices for both isotropic and anisotropic meshes. The maximal eigenvalues λ_{\max} are nearly constant, while the smallest eigenvalues λ_{\min} are $\mathcal{O}(h^2)$.

This is indeed a standard method in building anisotropic meshes [57, 77]. We choose to measure quality with respect to N because of its independence on anisotropy, and for its relevance in practical computational constraints: the memory footprint of data is as important as processor loading in modern computing.

The matrix condition numbers are determined by the smallest eigenvalues, which shrink in proportion with h^2 , the area of the triangles, and the largest eigenvalues, which are determined by triangle quality. Table 2.4 shows that the largest eigenvalues are relatively unaffected by the mesh size in both cases, and that even by adapting the mesh to the conductivity, we cannot escape the $\mathcal{O}(h^2)$ shrinkage in the smallest eigenvalues. However, the smallest triangles in the anisotropic mesh need not have areas as small as those of the isotropic mesh in order to achieve the same interpolation accuracy.

More generally, quality measures of both the anisotropic and isotropic meshes converge at the same rate. Recalling that $N \sim h^{-2}$, the estimates

$$\|s^h - s\|_{L_2} \leq C_{L_2} N^{-1} \quad (2.98)$$

$$|s^h - s|_{h^1} \leq C_{h^1} N^{-\frac{1}{2}} \quad (2.99)$$

$$\kappa_2 \leq C_{\kappa_2} N \quad (2.100)$$

hold for all meshes, irrespective of their isotropy. The gain in adaptively meshing the domain to suit Q comes as the size of the constants $C_{L_2}, C_{h^1}, C_{\kappa_2}$ are each smaller for well-

	isotropic	anisotropic	iso/aniso
κ_2	0.385	0.0723	5.33
$\ e_s\ _{L_2}$	4.69	1.13	4.15
$ e_s _{h^1}$	10.5	3.66	2.87

Table 2.5: Comparison of leading constants in quality bounds for isotropic and anisotropic meshes.

adapted meshes. The ratios of these constants given in Table 2.5 shows that, for example, to achieve comparable L_2 -accuracy on an unadapted mesh requires an approximately twenty-fold increase in the condition number over a well-adapted mesh.

Finally, mesh adaptation is only useful if its cost does not overwhelm that of solving the larger or badly conditioned problem on the isotropic mesh. Fortunately, the cost is reasonable: for N vertices, the weighted Delaunay triangulation is approximately $\mathcal{O}(N \log N)$, the vertex adjustment is $\mathcal{O}(N)$, and the number of iterations required is, in practise, constant. Approximately 12 iterations construct the anisotropic meshes studied in this example.

Chapter 3

The Inverse Problem

Perhaps the most widely applicable definition of an “inverse problem” is one that deduces that parameters of a PDE from information about its solutions. Examples include the measurement of the inhomogeneous reflectivity of sedimentary rock, the absorptivity of the human body by X-rays, the relaxation time of hydrogen nuclei in a magnetic field, and the electrical conductivity of inhomogeneous materials. These examples represent, respectively, reflection seismology for oil exploration, X-ray tomography for finding bone breaks, nuclear magnetic resonance (NMR) for identifying chemical composition, and electric impedance tomography (EIT) for identifying abnormal tissue. In each of these examples, we are typically inferring what material we are seeing (oil, bone, chemical, tumour) from the uniqueness of that material’s properties in the context in which it is measured.

With the possible exception of NMR, each of the above examples is an imaging example, whereupon the result of the inverse problem is an image representing how the parameter of the PDE changes with respect to location. Our inverse problem is the EIT problem. Our model is that at each point in space, the current density, \mathbf{j} , obeys Ohm’s law with respect to electric potential u :

$$\mathbf{j}(x) = \sigma(x)\nabla u(x), \quad x \in \Omega \subset \mathbb{R}^d. \quad (3.1)$$

We have explicitly identified the spatial dependence of the conductivity σ over the domain Ω , and our image will represent the size of the material property conductivity at each point in space. Measuring conductivity is useful in at least two applications: 1) in medical imaging, where abnormal tissue (cysts, tumours) has a different conductivity than normal surrounding tissue; and 2) in soil remediation, where oil-contaminated soil has a lower conductivity than the surrounding water-moistened soil.

EIT measures $\sigma(x)$ by taking measurements of voltages consequential to applied currents

at the boundary of the domain. Specifically, we define the Neumann-to-Dirichlet map,

$$\Lambda_\sigma : H^{-\frac{1}{2}}(\partial\Omega) \rightarrow H^{\frac{1}{2}}(\partial\Omega) \quad (3.2)$$

$$\Lambda_\sigma(f) = g, \quad (3.3)$$

where for given σ , $g = u|_{\partial\Omega}$ is the boundary trace of the solution to

$$\begin{aligned} -\operatorname{div}(\sigma \nabla u) &= 0, & x \in \Omega, \\ \sigma \frac{\partial u}{\partial n} &= f & x \in \partial\Omega, \\ \int_{\partial\Omega} f &= 0. \end{aligned} \quad (3.4)$$

The EIT problem is:

Given knowledge of the Neumann-to-Dirichlet map Λ_σ , deduce the conductivity

$$\sigma(x), x \in \Omega.$$

Practically, this means that from coincident measurements of currents and potentials on $\partial\Omega$, we attempt to recover σ throughout the interior of the domain. We note that deducing σ from the Dirichlet-to-Neumann, denoted $(\Lambda_\sigma)^{-1}$, is an equivalent problem. Λ_σ has an inverse provided we restrict ourselves to boundary sources $\int_{\partial\Omega} f = 0$, and solutions $\int_{\partial\Omega} u = 0$. We choose to refer to Λ_σ because its smoothing action is inherited by numerical reconstruction methods.

The EIT problem is first identified in the mathematics literature in a 1980 paper by Calderón [29], although the technique had been known in geophysics since the 1930s. Since the 1980 publication, a large community has formed around this problem. From work of Uhlmann, Sylvester, Kohn, Vogelius, Isakov and more recently, Alessandrini and Vessella, we know that complete knowledge of Λ_σ uniquely determines $\sigma \in L^\infty(\Omega)$, where $\Omega \subset \mathbb{R}^d$, $d \geq 2$ [13, 50, 54, 82]. Specific examples show that not only is this problem exponentially ill-posed, but that this extreme sensitivity is realised in simple examples [11, 12]. There has also been numerical work, studying reconstruction methods [24, 25, 26, 36, 41, 52], and the sensitivity of this problem to error in boundary measurements [37, 39]. Recently, there has also been work recovering σ from partial boundary data and from cases where the boundary location is imperfectly known [51, 55, 56]. Finally, in addition to the technical reports of

EIT in medical applications [40, 49, 58, 60, 89] and in geophysical applications [38], EIT has been the subject of several reviews [23, 30, 83].

What can we add to such a well-studied, though ill-posed problem? As we will show in Section 3.2, it is difficult to recover σ at a fine resolution. Hence, one approach to regularising this problem is to only coarsely resolve the domain, in accordance with the resolution and accuracy of the boundary data. However, homogenisation tells us that even if the conductivity is scalar at a fine scale, and hence we know that there is a 1-1 correspondence between Λ_σ and σ , if we resolve σ at a coarse scale, we can expect that we will measure an effective anisotropic conductivity, inconsistent with our assumption of isotropy. Hence, our approach is to recover an anisotropic, divergence-free conductivity using the parameterisation $s(x)$. Given our interest in exploring the application of $s(x)$ to this inverse problem, everything we do in this chapter, save a one-dimensional example in Section 3.1, is in \mathbb{R}^2 .

Before discussing the theory of this inverse homogenisation in Section 3.2, we motivate the choice to seek divergence-free conductivities by example in Section 3.1. Section 3.3 gives a reconstruction algorithm with some numerical results. Finally, as parameterisation of conductivity by $s(x)$ is new, there are a host of reconstruction methods using metric-based up-scaling that are not successful. Instead of polluting the inverse problems literature with these failed attempts, we confine their report to Section 3.4 of this thesis.

3.1 What aspect of conductivity can we measure?

An inverse problem seeks to measure a physical quantity². Ohm's law, equation (3.1) defines conductivity pointwise in correspondence with pointwise values of the potential gradient and current density. However, in practise, we cannot measure the gradient of a potential at a point, nor the density of a current flux: instead, scientists and engineers resort to measuring integrated quantities, such as the difference in potential between two points, and the current flux which is an integral of current density over an area. For practical purposes, conductivity will turn out to be an integral average of the pointwise σ appearing in Ohm's Law. The idea of reporting integral quantities is not specific to inverse problems: it is central to discrete differential geometry, an idea presented in [34, 73].

²We might argue that every scientific measurement provides data for an inverse problem.

Starting simply, in one dimension, what can we measure about conductivity? By correspondence to our problem in \mathbb{R}^2 , we are given the one-dimensional rod $0 \leq x \leq 1$ with conductivity $\sigma(x)$. We measure the potentials $u(0), u(1)$, and the boundary currents $\sigma \frac{du}{dx} \Big|_{x=0,1}$.

The potential obeys the conservation law

$$-\frac{d}{dx} \left(\sigma \frac{du}{dx} \right) = 0, \quad 0 < x < 1, \quad (3.5)$$

so we immediately note that consistent with our physical intuition, the current $\sigma \frac{du}{dx} = f =$ constant throughout the rod. Furthermore, since only the first derivative of u appears in the conservation law, we can add a constant to $u(0), u(1)$ without changing the solution, and only the difference $u(1) - u(0)$ is the relevant measurement of potential. Finally, the equation is linear: In this context, the only unique data we can collect about this rod is the potential difference $g = u(1) - u(0)$ when the current is equal to f , and when, for example, we double f , g will double also. Hence, the Neumann-to-Dirichlet map in one dimension is

$$\Lambda_\sigma(f) = ag \quad (3.6)$$

for some constant a . This value a is parameterised by σ and is the only information about σ that we can hope to obtain from studying Λ_σ .

It is helpful for what comes next to know what that constant a actually is. Multiplying (3.5) by u and integrating by parts, we have

$$\int_0^1 \sigma \left(\frac{du}{dx} \right)^2 = u \sigma \frac{du}{dx} \Big|_0^1 \quad (3.7)$$

$$\left(\sigma \frac{du}{dx} \right)^2 \int_0^1 \frac{1}{\sigma} = fg \quad (3.8)$$

$$f/g = \left(\int_0^1 \frac{1}{\sigma} \right)^{-1} = \sigma. \quad (3.9)$$

We have used the fact that the current, $f = \sigma \frac{du}{dx}$, is constant throughout the domain. The interpretation is that we identify the harmonic mean of σ as the effective conductivity of the domain. In one dimension, the harmonic mean is the only quantity we can measure about conductivity based only on the Neumann-to-Dirichlet map³.

³This is really a consequence of the coincidence of the gradient and divergence in one dimension: since

From an experimentalist's point of view, even if we are able to probe the material at its interior points, with only a finite number of measurements, all we can deduce is the harmonic mean of σ between our sample points — we can know integral values, but not pointwise values. The harmonic mean is the same conductivity we compute when we homogenise by the RVE method, the asymptotic homogenisation, and, most importantly, by metric-based up-scaling. The message is that the averaged conductivities we measure using the Neumann-to-Dirichlet map are also the homogenised conductivities. In fact, the only experiment we can do to determine conductivity is measure a Neumann-to-Dirichlet map. And, this is not just true in one dimension.

When the conductivity is homogenised by metric-based up-scaling, the homogenised conductivity admits piecewise linear solutions. In one dimension, this means that the homogenisation solution has constant gradient, so since the current $Q \frac{du}{dx}$ is constant, the homogenised conductivity is constant also. However, Q homogenised from σ is almost never constant in more than one dimension, but Q 's admission of solutions with constant gradient can tell us precisely what we can measure about homogenisation of conductivity using the Neumann-to-Dirichlet map.

Consider our domain to be the triangle $T = ijk$. As usual, (φ_i) form the piecewise linear basis over the triangle. To each test potential $g_i = \varphi_i|_{\partial T}$, we associate a test current $f_i|_{\partial T}$. In keeping with our restriction to integrated quantities, the information available to us are the f_i integrated over edges of the triangle. For example, for test potential i and edge ij , we have

$$f_i^{(ij)} = \int_{ij} f_i \quad (3.10)$$

$$= \int_{ij} (Q \nabla \varphi_i)^T d\hat{n} \quad (3.11)$$

$$= \int_{\partial T} (\varphi_i + \varphi_j) (Q \nabla \varphi_i)^T d\hat{n} \quad (3.12)$$

$$= \int_T \nabla(\varphi_i + \varphi_j)^T Q \nabla \varphi_i \quad (3.13)$$

$$= - \int_T \nabla \varphi_k^T Q \nabla \varphi_i = q_{ik}. \quad (3.14)$$

We have used the divergence theorem and the fact that $\nabla \varphi_i + \nabla \varphi_j + \nabla \varphi_k = 0$ on a triangle.

$\operatorname{div}(\sigma \nabla u) = \nabla(\sigma \nabla u)$, $x \in \mathbb{R}$, we know $\sigma \nabla u = \text{constant}$.

Similar calculations show that the three independent elements q_{ij} of the stiffness matrix, namely

$$q_{ij} = - \int_T \nabla \varphi_i^T Q \nabla \varphi_j, \quad (3.15)$$

together completely determine the relationship between these test potentials and integrated currents. In some sense this is obvious: the stiffness matrix for an element determines the relationship between the solution u and the sources $\int_T f \varphi_j$ over the element. However, we use the piecewise linear basis functions as test potentials because we know that they are exact solutions to the homogenised equation using Q at every interior point of the triangle.

Of course, the uniqueness results in the inverse problems literature say we can do far more in dimensions two and greater, stating that given Λ_σ , we can uniquely specify the corresponding isotropic conductivity everywhere in the domain [13, 50, 54, 82]. However, even in the broader context of an entire triangulated domain, deducing the elements of the stiffness matrix, the (q_{ij}) for the piecewise linear basis, is the correct thing to do, and we should interpret these edge-wise values as stiffness matrix elements for the homogenised conductivity. Further, since Q is divergence-free, we are justified in parameterising its stiffness matrix by $s(x)$, just as we did for the forward problem.

3.2 Theory

Our contribution is to the reconstruction not of σ , but to its homogenisation Q obtained through metric-based up-scaling. Every other practical reconstruction method parametrises σ as a scalar function. This is for good reason: uniqueness of σ with respect to Λ_σ is not only known for scalar σ , but we also know that with no further constraints, if σ is anisotropic, then it is only known up to a push-forward by some diffeomorphism on Ω [19, 45, 61, 81]. However, the main result of this section is that we can identify two unique conductivities consistent with the Neumann-to-Dirichlet map: the isotropic conductivity, and the divergence-free conductivity we obtain from metric-based up-scaling.

One way to see the non-uniqueness in recovering anisotropic conductivities is to look at the Neumann-to-Dirichlet map for anisotropic conductivity σ under a diffeomorphism $\Phi : \Omega \rightarrow \Omega$, where Φ is the identity in a neighbourhood of $\partial\Omega$. Under this mapping, we

consider the push-forward of the conductivity σ and a solution u :

$$\tilde{\sigma} = \Phi_*\sigma = \frac{\nabla\Phi^T\sigma\nabla\Phi}{|\det\nabla\Phi|} \circ \Phi^{-1}, \quad (3.16)$$

$$\tilde{u} = \Phi_*u = u \circ \Phi^{-1}. \quad (3.17)$$

With this mapping by Φ , for boundary source f , $\int_{\partial\Omega} f = 0$, the two problems

$$\begin{aligned} -\operatorname{div}(\sigma\nabla u) &= 0, & x \in \Omega, \\ \sigma \frac{\partial u}{\partial n} &= f, & x \in \partial\Omega, \end{aligned} \quad (3.18)$$

and

$$\begin{aligned} -\operatorname{div}(\tilde{\sigma}\nabla\tilde{u}) &= 0, & x \in \Omega, \\ \tilde{\sigma} \frac{\partial\tilde{u}}{\partial n} &= f, & x \in \partial\Omega, \end{aligned} \quad (3.19)$$

give equal boundary potentials $u = \tilde{u} = g$. The calculation is similar to that used to produce Q from σ in Chapter 1. When Φ is not the identity mapping at least one of $\sigma, \tilde{\sigma}$ is anisotropic, and we conclude that anisotropic conductivities cannot be in 1-1 correspondence with Neumann-to-Dirichlet maps.

Anisotropy presents a difficulty to this inverse problem. On the one hand, we know that the problem is exponentially ill-posed. That is, while we have Lipschitz continuity between σ and Λ_σ , for two conductivities σ_1, σ_2 in a finite subspace of L^∞ ,

$$\|\sigma_1 - \sigma_2\|_{L^\infty(\Omega)} \leq C \|\Lambda_{\sigma_1} - \Lambda_{\sigma_2}\|_{\mathcal{L}(H^{-\frac{1}{2}}(\partial\Omega), H^{\frac{1}{2}}(\partial\Omega))}. \quad (3.20)$$

($\|\cdot\|_{\mathcal{L}(H^{-\frac{1}{2}}(\partial\Omega), H^{\frac{1}{2}}(\partial\Omega))}$ is the natural operator norm for Λ_σ .) The constant C grows exponentially with the number of elements, N , we use to represent the conductivity [13]:

$$C(N) \geq A \exp\left(BN^{\frac{1}{d-1}}\right). \quad (3.21)$$

d is the dimension and A, B are absolute constants. This bound is confirmed by example in a note by Luca Rondi [71]. Hence, a way to regularise the problem is to demand reconstruction over fewer elements, reducing N to match the desired resolution of σ .

On the other hand, homogenisation reminds us that unless we sample an isotropic

conductivity at a sufficiently fine scale, we do not necessarily observe its samples to be isotropic. Practical reconstruction methods assume that σ is isotropic at the coarse scale of reconstruction, ignoring the effects of homogenisation. Indeed, this is the only example of how EIT appears to be robust: while otherwise ill-posed, numerical reconstruction does not seem to suffer in practise from the effect of homogenisation. While simulated test cases in the literature appear to lack any fine-scale texture required to produce an anisotropic homogenisation effect, a better explanation is provided by studying anisotropic EIT.

It is true that the Neumann-to-Dirichlet map defines an equivalence class of conductivities, where elements of each class are push-forwards of each other under diffeomorphisms. Following the exposition of [19], which has the strongest result we know of for uniqueness of anisotropic conductivities over bounded domains $\Omega \subset \mathbb{R}^2$, we consider the class of matrix functions $\sigma = [\sigma_{ij}]$ such that

$$\sigma \in L^\infty(\Omega; \mathbb{R}^{2 \times 2}), \quad \sigma^T = \sigma, \quad \lambda_{\min}(\sigma) > 0, \quad (3.22)$$

where λ_{\min} is the minimal eigenvalue of σ . Further, we identify the set

$$\Sigma(\Omega) = \{\sigma \in L^\infty(\Omega; \mathbb{R}^{2 \times 2}) \mid \lambda_{\max}(\sigma) < \infty\}. \quad (3.23)$$

It is known by example that if σ has an infinite eigenvalue, even its equivalence class cannot be recovered. The unbounded eigenvalue case gives an example of electric cloaking [44, 45].

The main result of [19] is that if $\Omega \subset \mathbb{R}^2$ is a simply connected bounded domain and $\sigma \in L^\infty(\Omega; \mathbb{R}^{2 \times 2})$ satisfies conditions (3.22), then the Neumann-to-Dirichlet map Λ_σ determines the equivalence class of conductivities

$$E_\sigma = \{\sigma_1 \in \Sigma(\Omega) \mid \sigma_1 = \Phi_*\sigma, \quad (3.24)$$

$$\Phi : \Omega \rightarrow \Omega \text{ is an } H^1\text{-diffeomorphism and } \Phi|_{\partial\Omega} = x\},$$

$\Phi_*\sigma$ denotes the push-forward of σ by diffeomorphism Φ . Q is a member of this equivalence class, since $Q = \Phi_*\sigma$ when Φ are the σ -harmonic coordinates F . This result is the anisotropic extension of the uniqueness result in [13] which requires $\sigma \in L^\infty(\Omega)$ to be scalar for its unique reconstruction from Λ_σ .

While we have uniqueness only up to the equivalence class (3.24), the following theorem

identifies two important elements of this class.

Theorem 3.1. *Let fixed $\sigma \in L^\infty(\Omega; \mathbb{R}^{2 \times 2})$ which satisfies conditions (3.22) identify the equivalence class E_σ . Then in dimension $d = 2$:*

- (i) *There exists a unique $Q \in E_\sigma$ such that Q is divergence-free; and*
- (ii) *There exists a unique $\gamma \in E_\sigma$ such that γ is scalar.*

A proof of the theorem follows a few comments.

Anisotropy requires that we be specific about which conductivity we are trying to recover. We are fortunate that for a given Neumann-to-Dirichlet map, we will always be able to find a scalar conductivity consistent with our measurements — it is the unique scalar conductivity in the equivalence class of conductivities defined by Λ_σ . This is what nearly all reconstruction algorithms do, and this explains why the algorithms rarely encounter difficulty due to anisotropy from up-scaling. A scalar conductivity consistent with Λ_σ always exists.

However, that scalar conductivity is only L^∞ , and may vary at scales finer than the resolution of the mesh used in the reconstruction. In this case, seeking the homogenisation of the original conductivity may be more reasonable. The inverse problem asks “What can we measure?” Since the ill-posedness of this problem restricts the scale at which we can make measurements, the element of the equivalence class associated to the Neumann-to-Dirichlet map that we can actually measure is the unique divergence-free homogenised conductivity. There is no one “true” conductivity for the domain, but the homogenised conductivity is a unique conductivity that is well-behaved with respect to change of scale.

We close this section with a proof of the theorem:

Proof of Theorem 3.1. For claim (i), recall that the homogenisation of σ does not require σ to be isotropic, so Q exists by construction of the σ -harmonic coordinates F .

Conversely, suppose $Q = \Phi_*\sigma$ is divergence-free. Then the development of equation (1.17) in Section 1.2 shows that

$$\int_{\Omega} w \operatorname{div}(\sigma \nabla \Phi) = 0 \tag{3.25}$$

for all test functions $w \in H_0^1$. But since $\Phi = x|_{\partial\Omega}$, Φ uniquely solves the equation for the σ -harmonic coordinates, $\Phi = F \in H_0^1$, and $Q = F_*\sigma$ uniquely when Q is divergence-free.

For claim (ii), we will show that scalar conductivity γ is the push-forward of σ by coordinates G which uniquely solve the equations

$$\begin{aligned} \operatorname{div} \left(\frac{\sigma}{\sqrt{\det \sigma}} \nabla G \right) &= 0, & x \in \Omega, \\ G &= x, & x \in \partial\Omega. \end{aligned} \quad (3.26)$$

In this case, we will show that $\gamma \operatorname{Id} = \left(\sqrt{\det \sigma} \circ G^{-1} \right) \operatorname{Id} = G_*\sigma$.

Let $G = (G_1, G_2)$ satisfy (3.26). Then

$$\operatorname{div} \left(\frac{\sigma}{\sqrt{\det \sigma}} \nabla G_2 \right) = 0 \quad (3.27)$$

which implies that we can find a function P_1 such that

$$R \nabla P_1 = \frac{\sigma}{\sqrt{\det \sigma}} \nabla G_2, \quad P_1 = x|_{\partial\Omega} \quad (3.28)$$

We can enforce the boundary condition on P_1 provided the domain is closed, bounded and simply connected.

Observe that for $A \in \mathbb{R}^{2 \times 2}$,

$$(A^{-1})^T = \frac{1}{\det A} R A R^T, \quad R = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}. \quad (3.29)$$

This implies that we can rearrange (3.28),

$$\sqrt{\det \sigma} R^T \sigma^{-1} R \nabla P_1 = R^T \nabla G_2 \quad (3.30)$$

$$\frac{\sigma}{\sqrt{\det \sigma}} \nabla P_1 = R^T \nabla G_2 \quad (3.31)$$

implying

$$\operatorname{div} \left(\frac{\sigma}{\sqrt{\det \sigma}} \nabla P_1 \right) = 0. \quad (3.32)$$

By the uniqueness of solutions to (3.26), we have that $P_1 = G_1$. Together, we have

$$\frac{\sigma}{\sqrt{\det \sigma}} \nabla G_2 = R \nabla G_1, \quad (3.33)$$

$$\frac{\sigma}{\sqrt{\det \sigma}} \nabla G_1 = R^T \nabla G_2, \quad (3.34)$$

or, more compactly,

$$\frac{\sigma}{\sqrt{\det \sigma}} \nabla G R = R \nabla G. \quad (3.35)$$

A second application of identity (3.29), this time to ∇G gives

$$\frac{\sigma}{\sqrt{\det \sigma}} \nabla G = R \nabla G R^T, \quad (3.36)$$

$$\frac{\nabla G^T \sigma \nabla G}{|\det \nabla G|} = \sqrt{\det \sigma}. \quad (3.37)$$

Modulo composition by G^{-1} , we identify the left-hand side as the push-forward of σ by G , and setting $\gamma = \sqrt{\det \sigma} \circ G^{-1}$, we have that $G_* \sigma = \gamma \text{Id}$. \square

Finally, we note that the authors in [19, 80, 81] use quasi-conformal changes of coordinates in \mathbb{R}^2 to show that any anisotropic conductivity can be changed to an isotropic one, a construction central to their observation of their observation of the equivalence class E_σ , and a result we discover via metric-based up-scaling. Both proofs, however, rely on arguments in \mathbb{R}^2 . It is not known, for example, if the construction of coordinates in (3.26) is valid in \mathbb{R}^3 , and a similar result in dimensions $d \geq 3$ is future work for the time-being.

3.3 Reconstructing $s(x)$

From the equivalence class of conductivities determined by the Neumann-to-Dirichlet map, we have identified two conductivities from this class that we might naturally choose to reconstruct. The first is the scalar conductivity, and the second is the divergence-free conductivity. As mentioned in the introduction to this chapter, the reconstruction of the scalar conductivity has been well studied. However, the reconstruction of divergence-free conductivities, interpreted as homogenised conductivities, is new in this thesis, and a reconstruction algorithm is the topic of this section.

We tried a great variety of reconstruction algorithms, and we present the most successful choice here. Section 3.4 enumerates our less-successful reconstruction attempts. Our algorithm parameterises the stiffness matrix of divergence-free conductivity Q using the piecewise linear interpolation of $s(x)$. We pose the problem on a triangulated domain as a least-squares minimisation program with a Tikhonoff regularisation, much as is introduced for EIT in [36]. This method solves for a divergence-free stiffness matrix for Q . Finally, we use the values for interpolated $s(x)$ to estimate its curvatures, hence corresponding values for Q on triangles of the mesh.

3.3.1 Least-squares method

A least squares problem for $s(x)$ is

$$\begin{aligned} & \underset{s \in \mathcal{P}_1^h}{\text{minimise}} && \frac{1}{2} \sum_{i=1}^M \|\Lambda_s f^{(i)} - g^{(i)}\|_{L^2(\partial\Omega)}^2 + \nu \|\nabla Q_s\|, \\ & \text{subject to} && s(x) \text{ convex} \end{aligned} \tag{3.38}$$

We have identified the dependence of the Neumann-to-Dirichlet map Λ_s and the conductivity Q_s on the parameterisation $s(x)$. We have approximated the operator norm $\mathcal{L}(H^{-\frac{1}{2}}(\partial\Omega), H^{\frac{1}{2}}(\partial\Omega))$ by measuring M current-voltage pairs $(f^{(i)}, g^{(i)})$, $i = 1 \dots M$ over the boundary of the domain. That is, for $u^{(i)}$ solving

$$\begin{aligned} -\operatorname{div}(Q_s \nabla u^{(i)}) &= 0, & x \in \Omega, \\ \frac{\partial}{\partial n}(Q_s u^{(i)}) &= f^{(i)}, & x \in \partial\Omega, \end{aligned} \tag{3.39}$$

we have $g^{(i)} = u^{(i)}|_{\partial\Omega}$. $\nu \ll 1$ is a regularisation parameter, and $\|\nabla Q_s\|$ measures the smoothness of Q_s . We discuss specific choices for this smoothness norm shortly. We refer to the first and second terms of the objective in program (3.38) as the boundary-error and regularisation terms, respectively. By arguments of Chapter 1, the constraint, that $s(x)$ is a convex function, ensures that Q_s is positive-definite. Problem (3.38) is not a convex program. This is due to the nonlinear relationship between conductivity and the Neumann-to-Dirichlet map. However, with the small amount of regularisation, we can use a convex optimisation algorithm to solve this problem.

In order to use a convex optimisation algorithm to solve the least-squares problem, we need to compute the gradient of the boundary-error term with respect to $s(x)$. We do so by adapting the adjoint method outlined in [36], which we repeat here and frame in the context of the piecewise linear discretization of $s(x)$. To compute the gradient, we consider the effect of a variation δs of $s(x)$ on the solution u :

$$s(x) \Rightarrow s(x) + \delta s(x), \quad (3.40)$$

$$Q \Rightarrow Q + \delta Q, \quad (3.41)$$

$$u(x) \Rightarrow u(x) + \delta u(x) + \delta^2 u(x) + \dots \quad (3.42)$$

Q is linear in $s(x)$, so we have furthermore identified the first variation δQ in Q . $(s, \delta s)$ and $(Q, \delta Q)$ are used interchangeably in the following. Suppressing the test mode index i for clarity, the variation of an element of the boundary-error term with respect to s is

$$D \left(\frac{1}{2} \|\Lambda_s f - g\|_{L^2(\partial\Omega)}^2 \right) (\delta s) = (\Lambda_s f - g, D\Lambda_s(\delta s))_{\partial\Omega}, \quad (3.43)$$

where $(\cdot, \cdot)_{\partial\Omega}$ denotes the $L^2(\partial\Omega)$ inner product, and $D\Lambda_s(\delta s)$ is the variation of the Neumann-to-Dirichlet map with respect to s . Our task is to determine this variation.

Writing the Neumann-to-Dirichlet problem (3.39) in terms of these expansions and grouping terms by their order gives the 0th-order problem,

$$\begin{aligned} -\operatorname{div}(Q\nabla u) &= 0, & x \in \Omega, \\ \frac{\partial}{\partial n}(Qu) &= f, & x \in \partial\Omega, \end{aligned} \quad (3.44)$$

and the 1st-order problem,

$$\begin{aligned} -\operatorname{div}(Q\nabla\delta u) &= \operatorname{div}(\delta Q\nabla u), & x \in \Omega, \\ \frac{\partial}{\partial n}(Q\delta u) &= -\frac{\partial}{\partial n}(\delta Qu), & x \in \partial\Omega. \end{aligned} \quad (3.45)$$

Next, we identify the derivative of the Neumann-to-Dirichlet map

$$D\Lambda_s(\delta s) = \delta u|_{\partial\Omega}, \quad (3.46)$$

and define the L^2 -adjoint derivative with respect to a test function w :

$$\int_{\partial\Omega} w D\Lambda_s(\delta s) = \int_{\Omega} \delta s D^* \Lambda_s(w). \quad (3.47)$$

Finally, we define a third problem, the problem adjoint to the 0th-order problem:

$$\begin{aligned} -\operatorname{div}(Q\nabla u^*) &= 0, & x \in \Omega, \\ \frac{\partial}{\partial n}(Qu^*) &= w, & x \in \partial\Omega, \end{aligned} \quad (3.48)$$

where w is again a test function. We multiply the 1st-order problem by u^* and integrate by parts, applying the boundary conditions of each of the three problems:

$$\int_{\Omega} u^* \operatorname{div}(Q\nabla \delta u) = - \int_{\Omega} u^* \operatorname{div}(\delta Q \nabla u), \quad (3.49)$$

$$\begin{aligned} \Rightarrow \int_{\Omega} (\nabla u^*)^T Q \nabla \delta u + \int_{\Omega} (\nabla u^*)^T Q \nabla u &= \\ \int_{\partial\Omega} u^* (\delta Q \nabla \delta u)^T d\hat{n} + \int_{\partial\Omega} u^* (\delta Q \nabla u)^T d\hat{n}, & \end{aligned} \quad (3.50)$$

$$\Rightarrow - \int_{\Omega} \delta u \operatorname{div}(Q\nabla u^*) + \int_{\partial\Omega} (\delta u) w = - \int_{\Omega} (\nabla u^*)^T \delta Q \nabla u, \quad (3.51)$$

$$\Rightarrow \int_{\partial\Omega} w D\Lambda_s(\delta s) = - \int_{\Omega} (\nabla u^*)^T \delta Q \nabla u, \quad (3.52)$$

and we identify the adjoint derivative of Λ_σ as

$$(\delta s, D^* \Lambda_\sigma(w))_\Omega = - \int_{\Omega} \nabla u^T \delta Q \nabla u^*, \quad (3.53)$$

where u, u^* solve the primal and adjoint problems, respectively, and $(\cdot, \cdot)_\Omega$ is the $L^2(\Omega)$ inner product. In terms of our original calculation of the variation of the boundary error, we have

$$D \left(\frac{1}{2} \|\Lambda_s f - g\|_{L^2}^2 \right) (\delta s) = - \int_{\Omega} \nabla u^T \delta Q \nabla u^*, \quad (3.54)$$

where u^* solves the adjoint problem with $w = \Lambda_s f - g$, representing the error in our guess for the Neumann-to-Dirichlet map for current-voltage pair (f, g) . When this error is zero for all modes $i = 1 \dots M$, the variation of the boundary-error term is zero. The role of the adjoint derivative framework is to interpret variability on the boundary — the boundary error — to variability in the interior — the error in conductivity.

In the discrete setting, the right-hand side of equation (3.54) is a sum over edges, weighted by the variations δq_{ij} in the stiffness matrix elements:

$$-\sum_{i \sim j} \delta q_{ij} (u_i - u_j)(u_i^* - u_j^*). \quad (3.55)$$

Applying formula (2.9) for the q_{ij} in terms of the interpolants s_i , over hinge $ijkl$,

$$\begin{aligned} \delta q_{ij} = & -\frac{1}{|ij|^2} (\cot \theta_{ijk} + \cot \theta_{ijl}) \delta s_i - \frac{1}{|ij|^2} (\cot \theta_{jik} + \cot \theta_{jil}) \delta s_j \\ & + \frac{1}{2A_{ijk}} \delta s_k + \frac{1}{2A_{ijl}} \delta s_l. \end{aligned} \quad (3.56)$$

This completes the linear relationship between the variation of the boundary-error term and the variations δs_i , from which we compute the gradient of the boundary error. Since the adjoint method is constructed using integration by parts, developing the discrete adjoint derivative starting with the discrete finite element form of the primal and adjoint equations, and using summation by parts produces exactly the result we have here.

Calculating the gradient of the boundary-error term for a particular $s(x)$ means solving two PDE, the 0th-order primal problem and the adjoint problem. These problems have the same stiffness matrix, and so need to be assembled only once. This adjoint method can be extended to compute higher-order variations of the boundary-error term. In particular, in order to minimise the objective, we need to compute its Hessian. However, we find that this problem is sufficiently well-behaved that we can use finite-difference approximations to the Hessian in our codes.

3.3.2 Regularisation

The idea of the regularisation term is to provide a criterion for choosing between several qualitatively different choices for the conductivity that might reduce the boundary error, given that the problem is ill-posed. Our regularisation selects for conductivities that are in some way smooth at the scale of our mesh, and is called Tikhonhoff regularisation for where it appears in the minimisation problem. The parameter ν weighs the importance of matching the data, given in the boundary-error term, against the importance of having a smooth solution. The regularisation criteria we have found appropriate for this problem

control the smoothness of the trace of Q . In particular, we have investigated regularisation using $\|\text{tr } Q\|_{\text{TV}}$ and $\|\nabla \text{tr } Q\|_{L^2}$, that is, the total variation norm and the $L^2(\Omega)$ -norm of the gradient, respectively.

There are two reasons we choose to regularise the problem using the smoothness of $\text{tr } Q$. First, $\text{tr } Q$ is the sum of the eigenvalues of Q and hence is rotation invariant. Another rotation-invariant choice is $\det Q$, the product of the eigenvalues. However, our second reason for choosing $\text{tr } Q$ is that it is linear in $s(x)$, and so in our interpolants s_i , while the determinant is not.

In the continuous setting, Q is constructed from weak second derivatives of $s(x)$. However, our representation for $s(x)$ is piecewise linear, and we must justify how we can use $s(x)$ to represent Q . This question is important, since a meaningful report of the conductivity is Q itself, not $s(x)$. From an implementation standpoint, we need to approximate Q itself from the interpolants s_i in order to perform our $\text{tr } Q$ regularisation.

A specific calculation justifies using only piecewise linear $s(x)$ to approximate Q . Set V_i the Voronoi region associated to vertex i , and ∂V_i its boundary. See Figure 3.1. Now we compute using the divergence theorem

$$\int_{V_i} \text{tr } Q = \int_{V_i} \partial_{xx}s + \partial_{yy}s \quad (3.57)$$

$$= \int_{V_i} \Delta s \quad (3.58)$$

$$= \int_{\partial V_i} \nabla s^T d\hat{n} \quad (3.59)$$

$$= \frac{1}{2} \sum_{j \in \mathcal{N}(i)} (\cot \alpha_{ij} + \cot \alpha_{ji}) (s_i - s_j) \quad (3.60)$$

The last line assumes a piecewise constant discretisation of $s(x)$, which gives ∇s constant on each triangle. The angles α_{ij}, α_{ji} are the angles opposite edge ij on its hinge. This is the classic cotangent formula for approximating the Laplacian of piecewise linear functions.

At this point, we have only the integral of $\text{tr } Q$ over the Voronoi area, rather than a pointwise value. However, because of the smoothness of solutions associated to Q , we know that arithmetic averages of Q scale stably. Furthermore, we can also compute averages of




Figure 3.1: The Voronoi region associated to the centre vertex is shaded. The boundary of the Voronoi region is comprised of the edges dual to the edges incident to the centre vertex.

each of the second partial derivatives of $s(x)$. For example,

$$\langle \partial_{xx}s \rangle = \frac{1}{|V_i|} \int_{V_i} \partial_{xx}s = \frac{1}{|V_i|} \int_{\partial V_i} \begin{pmatrix} \partial_x s \\ 0 \end{pmatrix}^T d\hat{n} \quad (3.61)$$

$\partial_x s$ is constant on each triangle, so this quantity is well-defined for piecewise linear approximations of $s(x)$. This example illustrates that since we can be satisfied with arithmetic averages of Q , piecewise linear $s(x)$ are sufficient to reconstruct of conductivity. However, in our calculations, we choose a slightly different approximation to integral averages of Q .

Consistent with these arguments are other stencils for measuring integrals of curvature presented in the literature [46, 59, 70, 72]. We choose to approximate Q on triangles using the stencil of Figure 3.2, whose adequate approximation properties are discussed in [70]. This choice gives a smaller stencil for approximating $\nabla \text{tr } Q$, and allows us to directly read the conductivity as second derivatives of an interpolating polynomial for the purposes of reporting results.

The six points of the stencil are sufficient to define a quadratic interpolant in the neighbourhood of the triangle. This gives

$$s_{ijk}(x, y) = ax^2 + bxy + cy^2 + dx + ey + f \quad (3.62)$$

for constants a, b, c, d, e, f linear in the six interpolants s_i in the stencil. Differentiating twice gives a constant approximation

$$Q_{ijk} = \begin{pmatrix} 2a & b \\ b & 2c \end{pmatrix}. \quad (3.63)$$




Figure 3.2: Stencil for approximating Q on triangle ijk from the interpolants s_i at nearby vertices.

Since a, c are linear in the s_i , so too is $\text{tr } Q_{ijk} = 2(a + c)$. $\det Q_{ijk}$ is only quadratic in the s_i , and so it isn't onerous to calculate its derivatives with respect to the s_i . In fact, since $\sqrt{\det Q} = \sigma \circ F$, for the transformation F from scalar σ to Q , we might prefer to control $\det Q$ directly. However, $\|\nabla \det Q\|$ is not convex, and we see in the simulation results that the image of $\det Q$ closely matches that of $\text{tr } Q$ in practise.

On the triangulation, we compute the regularisation norm using, for example,

$$\|\text{tr } Q\|_{TV} = \sum_{i \sim j} |ij| |\text{tr } Q_{ij} - \text{tr } Q_{ji}|. \quad (3.64)$$

The subscripts ij and ji denote values on either side of edge ij . This formula for total variation of piecewise constant functions over triangular meshes follows from [53]. This is a convex function of the s_i , since $\text{tr } Q_{ij} - \text{tr } Q_{ji}$ is linear in the s_i , and the absolute value function is convex. Other choices are the approximated norms

$$\|\text{tr } Q\|_{L^1} = \sum_{i \sim j} |ij^*| |ij| |\text{tr } Q_{ij} - \text{tr } Q_{ji}|, \quad (3.65)$$

$$\|\text{tr } Q\|_{L^2}^2 = \sum_{i \sim j} |ij^*| |ij| (\text{tr } Q_{ij} - \text{tr } Q_{ji})^2, \quad (3.66)$$

$$(3.67)$$

but we find that the approximation to the total variation norm produces the best results. At each boundary edge, where Q is defined only on one side of the edge, we assume a value




Figure 3.3: Placement of ghost vertices and triangles at the boundary of the mesh for specifying values of the stiffness matrix on boundary edges.

for Q on an imaginary triangle outside the domain. This is similar to the assumption of σ near the boundary of the domain, as is done in most other EIT algorithms. However, in contrast to these other algorithms, we use this assumed boundary value for the conductivity only to facilitate regularisation, and not to otherwise parameterise the conductivity. Still, at boundaries, we must be careful about how we parameterise $s(x)$, and that is the topic of the next section.

3.3.3 Boundary values for s

In Chapter 2, we developed the relationship between interpolants s_i and the stiffness matrix entries q_{ij} for q_{ij} at a hinge of the mesh. This presents a problem at the boundary of the mesh, on edges that border only one triangle. Indeed, by interpolating $s(x)$ as a piecewise linear function, we have concentrated all of the curvature of $s(x)$, and hence the contribution of the approximated Q to the stiffness matrix, to the edges on each hinge. We avoided the boundary problem in Chapter 2 by considering only problems where $u = 0|_{\partial\Omega}$, but we do not have homogeneous boundary data in the inverse problem. The solution is to provide “ghost” vertices and triangles on the exterior of the mesh.

Figure 3.3 gives an example of how ghost triangles and vertices are added at the boundary of a mesh. $s(x)$ is interpolated at the ghost vertices, providing sufficient degrees of freedom to specify the stiffness matrix on boundary edges, and these interpolants become additional variables in the optimisation. The locations of the ghost vertices are somewhat arbitrary — at each boundary edge, we choose to reflect the vertex inside the domain across the edge to determine its ghost vertex location. This is one choice that is a compromise between setting the ghost vertices on the point at infinity, and setting the ghost vertices

on the boundary edge itself. In the first instance, the area of the associated ghost triangle goes to infinity, its coefficient in the q_{ij} formula, the reciprocal of the triangle area, is zero, and $s(x)$ interpolated at the ghost vertex has no contribution to the stiffness matrix. In the second instance, the area of the ghost triangle goes to zero, and the stiffness matrix is extremely sensitive to the value of the ghost vertex interpolant. Keeping the area and the angles of the ghost triangle comparable to the dimensions of the other triangle of the hinge seems a reasonable choice.

We are in effect expanding the domain in order to compute the stiffness matrix at the boundary, despite the fact that the finite element method does not otherwise require that we do so. We see one justification for adding degrees of freedom to $s(x)$ at ghost vertices by comparing the degrees of freedom in the stiffness matrix to the degrees of freedom given by the interpolants of $s(x)$. Suppose we have a mesh with E edges and V vertices, of which V_I are interior vertices, and V_B are boundary vertices. There are E elements q_{ij} in the stiffness matrix, given the symmetry of the matrix, and the fact that its row sums are zero. Furthermore, since Q is divergence-free, we have the conditions

$$\sum_{j \in \mathcal{N}(i)} q_{ij} (\mathbf{x}_i - \mathbf{x}_j) = 0, \quad (2.10)$$

of which there are two at each interior vertex in dimension 2. Hence, the stiffness matrix has $E - 2V_I$ degrees of freedom.

Now, suppose we count the number of edges in a triangle mesh with respect to the number of its vertices. Building the mesh up from a single triangle, each time we add a vertex, we also add two edges. (This is known as adding an “ear” to the mesh.) This is true except when we add an edge that encloses a vertex in a one-ring, creating an interior vertex, see Figure 3.4. This means that the total number of edges is

$$E = 3 + 2(V - 3) + V_I. \quad (3.68)$$

Rearranging this relationship, and noting $V = V_I + V_B$ gives

$$E - 2V_I = V + V_B - 3. \quad (3.69)$$




Figure 3.4: Computing the number of edges in a mesh with respect to its number of vertices and interior vertices. The left figure shows how adding a vertex to a boundary edge adds two new edges, while the right figure shows that creating an interior vertex requires adding an edge to complete a one-ring.

The left-hand side is the number of degrees of freedom in the stiffness matrix, and the right-hand side is the number of degrees of freedom in the piecewise linear interpolation of $s(x)$. There are $V + V_B$ knots for $s(x)$ if we include the ghost vertices, less three degrees of freedom since $s(x)$ is itself unique only up to the addition of affine functions.

3.3.4 Enforcing convexity

As noted in Chapter 2, convexity of the piecewise linear surface interpolating $s(x)$ is equivalent to $q_{ij} > 0$. This constraint is linear in our solution variables s_i , and so seems straightforward to enforce in the optimisation problem. However, as we showed by the example of edge-flips in Figure 2.4, even if $s(x)$ is convex, its interpolation may not be so. All the same, even without adjusting the mesh during the optimisation for the s_i , requiring $q_{ij} > 0$ does not in practise present a difficulty to our nonlinear optimising code. One way we might see such difficulty is where the solution has $q_{ij} = 0$ for some edge, indicating that the optimal solution is to have no curvature at that edge. Likely, this indicates an edge that should be flipped in the mesh. However, the effect of having a zero-valued edge in the stiffness matrix seems to be sufficiently local that it does not qualitatively affect our solves. Moreover, flipping edges during the optimisation seems to destabilise the solve.

3.3.5 Numerical results

As a test of the use of $s(x)$ as a suitable parameterisation for the inverse problem, we test the reconstruction using the boundary least-squares error minimisation (3.38). The algorithm we use is an interior point method provided in the C++ code IpOpt [2]. Details




Figure 3.5: A sample isotropic conductivity for testing reconstruction. The image on the left is σ , while the image on the right is $\sqrt{\det Q} = \sigma \circ F^{-1}$. The dark blue background has conductivity 1.0, the red circle has conductivity 10.0 and the yellow bar has conductivity 5.0. In this case, all of the features shrink in harmonic coordinates.

of the algorithm appear in [84, 85, 86]. Despite the fact that the boundary-error term is not convex in the conductivity, and so is not convex in the piecewise linear parameterisation for $s(x)$, this minimisation does give us sufficiently good results that we can assess the role of σ -harmonic coordinates in this inverse problem.

Our first test is the reconstruction of a conductivity pattern that is piecewise constant at a coarse scale. This is similar to the test patterns used in previous work, such as in [25, 26, 39, 52], and helps us understand what the Neumann-to-Dirichlet map tells us about divergence-free conductivities. Figure 3.5 shows the original conductivity pattern, a few large geometric shapes embedded in a constant background conductivity. The first image in Figure 3.5 is the conductivity that we would expect to recover from boundary data under the assumption that the conductivity is isotropic. We compare this to the second image in Figure 3.5, which shows $\sqrt{\det Q} = \sigma \circ F^{-1}$ for Q the push-forward of σ by its harmonic coordinates. Although this picture does not capture the anisotropy of our divergence-free conductivity, it does indicate a difficulty in trying to capture up-scaled conductivities — in this case, the area of the up-scaled features is significantly smaller than the original features, and it is difficult to detect small features with EIT, especially far from the domain boundary.

To reconstruct the conductivity from boundary data, we solve for the Dirichlet data

given Neumann data sinusoidal over the boundary of the circle.

$$f^{(i)} = \begin{cases} \cos(2\pi \lfloor i \rfloor \theta), & i \text{ even,} \\ \sin(2\pi \lfloor i \rfloor \theta), & i \text{ odd.} \end{cases} \quad (3.70)$$

We parameterise the circle boundary by the polar angle θ . $\lfloor x \rfloor$ is the largest integer not larger than x . This gives us the $(f^{(i)}, g^{(i)})$ pairs. Note that the decision to use sinusoidal boundary excitations is arbitrary. As pointed out in [39], the best boundary excitations can be obtained by diagonalising the boundary map, and this can be done in practise with only boundary data. However, arguments in [29, 82] which use geometric optics solutions to show unicity of isotropic EIT reconstruction indicate that sinusoidal excitations are always nearly optimal.

Images of our reconstruction for the conductivity in Figure 3.5 are given in Figure 3.6. These reconstructions use $M = 18$ test modes and we choose regularisation parameter $\nu = 3 \times 10^{-3}$ with the $\|\text{tr } Q\|_{TV}$ regularisation. By design of the regularisation, the reconstruction tends to favour piecewise constant conductivities, and like other reconstruction methods, blurs the original σ . The circular blob, although having a large conductivity, is difficult to resolve; as we have already noted, it appears quite small in harmonic coordinates. When the scale of the conductivity pattern is large, it may indeed be best to search directly for σ instead of its divergence-free up-scaled value.

The situation is different when the conductivity pattern varies at fine scales. Consider the laminated pattern in Figure 3.7, the same example conductivity we use in Chapter 2. Again the first image in Figure 3.7 gives the original conductivity, while the second image gives $\sqrt{\det Q} = \sigma \circ F^{-1}$, a representation of the up-scaled conductivity. The fine-scale structure is difficult to make out in harmonic coordinates, but the homogenised pattern, which is anisotropic, is detectable by EIT.

Figures 3.8 and 3.9 show the results of the reconstruction. The first set of figures, Figure 3.8 shows the anisotropy captured by the reconstruction. Comparing the rendering of $s(x)$ to that in Figure 2.6, page 39, we see that the overall anisotropy is captured — conductivity is greater in the y -direction than in the x -direction — but the dynamic range of the reconstruction is muted. The approximate ratio of the y -conductivity to the x -conductivity is approximately 8:1 in the original, but only 2:1 in the reconstruction. However, we obtain




Figure 3.6: Reconstruction of the isotropic conductivity in Figure 3.5. The left-hand figure shows $\text{tr } Q$, while the right-hand figure shows $\sqrt{\det Q}$. The reconstruction blurs the original σ , similar to other methods in the literature, but does not underestimate the dynamic range of the large rectangle.




Figure 3.7: A sample conductivity that, while isotropic at the fine scale, is anisotropic at the coarse scale. This is the same conductivity of Figure 2.5 which we considered in the forward problem. The right-hand image gives $\sqrt{\det Q} = \sigma \circ F^{-1}$, showing how in harmonic coordinates, the fine-scale structure of the medium is difficult to see pointwise. Fortunately, we can still recover the coarse-scale anisotropy.

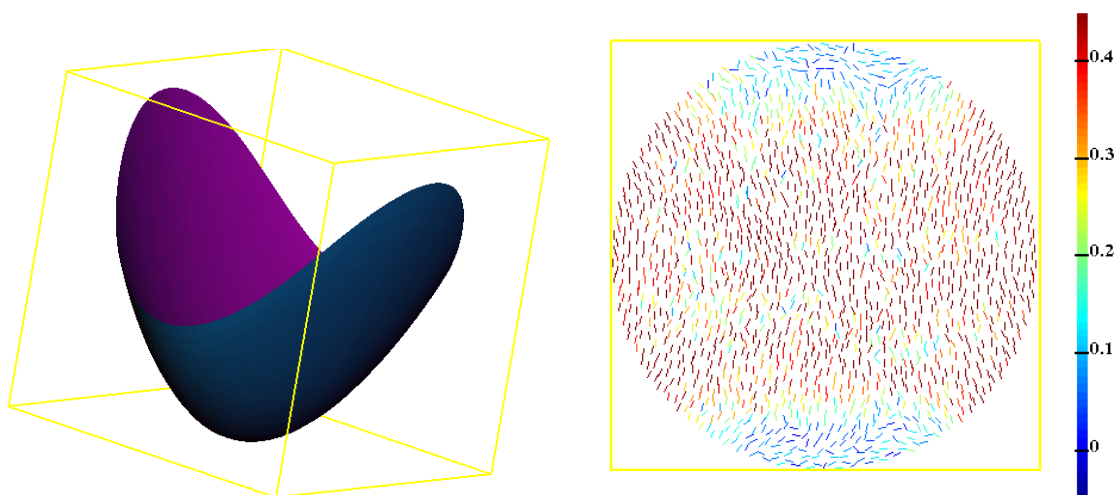


Figure 3.8: Anisotropic reconstruction showing the parameterisation $s(x)$ we recover using EIT (left image), and the pattern of anisotropy we see by rendering the orientation of the maximal eigenvalue of Q (right image). The colour-bar in the right image indicates the strength of the anisotropy as $|\lambda_{\max} - \lambda_{\min}|/\text{tr } Q$.

this result over a large range of regularisation weights, $\nu = 1 \times 10^{-6}$ to 1×10^{-3} , and over a large range of initial guesses for the conductivity, suggesting that it is indeed the boundary data, and not any particular bias in our reconstruction method, that indicates anisotropy. Without the parameterisation $s(x)$ for divergence-free conductivities, we could not identify this particular anisotropic conductivity from the equivalence class defined by our Neumann-to-Dirichlet map.

For comparison with the other reconstruction, we show the trace and determinant of Q in Figure 3.9. Included in these figures are the ghost boundary ears we add to give complete hinges on the boundary edges. The plot of the trace is smoother (in the sense of total variation) than that of the determinant, which is not surprising, since we have controlled the trace directly in our regularisation. The fact that the trace and determinant patterns are nearly coincident suggests that regularising the trace of Q is a sufficient substitute for regularising over the determinant, despite the fact that $\det Q = \sigma^2 \circ F^{-1}$, and relates directly to the original scalar conductivity.




Figure 3.9: The determinant and trace of the reconstructed divergence-free conductivity Q . That $\det Q$ and $\text{tr} Q$ are nearly coincident even for this anisotropic case suggests that regularising one of the two regularises the other.

3.3.6 Evaluating algorithms for divergence-free EIT

Reconstructing conductivities by minimisation of program 3.38 is admittedly not competitive with existing methods for tomography of piecewise-constant scalar conductivities. However, determining divergence-free conductivities is certainly less-well studied than the many reconstruction techniques for isotropic conductivities, and here, we have succeeded in simply introducing divergence-free conductivities as new reconstruction candidates. These candidates are well-suited to representing conductivities having fine-scale variability, while having low-dimensional representations. While we cannot claim discovery of an optimal reconstruction method, we can specify criteria appropriate for evaluating future reconstruction methods which seek divergence-free conductivities.

We propose the following test: An experimenter begins with a homogeneous simply-connected material, carves out regions of the material, and replaces the removed material with inclusions having different conductivities than the original. The experimenter knows where these inclusions are, and asks the mathematician to find them using EIT. If the mathematician's reconstruction method is perfect, he will recover the locations and conductivities of all of the experimenter's inclusions. Succeeding in this test many times will convince the experimenter that the mathematician's EIT method is sound.

Whether simulated or tested by an experimental apparatus, this is the very test performed on isotropic EIT. Owing to the ill-posedness of EIT, experimenters and mathematicians alike are typically satisfied with a reconstruction method that produces the locations of all inclusions, even if the dynamic range of the reconstructed conductivities is attenuated. For divergence-free EIT, we need a similar check. The difficulty is that the divergence-free conductivity is related to the original conductivity by the harmonic coordinates, which are solutions to an elliptic problem parameterised by the original isotropic conductivity solved over the entire domain. This means that the divergence-free conductivity at a point depends on the isotropic conductivity everywhere. Fortunately, we have a relationship between the isotropic and divergence-free conductivities that holds locally, and we can in some sense evaluate the possibility of falsely detecting or missing inclusions based on a divergence-free reconstruction.

Suppose that the experimenter constructs test domains from materials isotropic at the finest scale. In this case, we know that the divergence-free conductivity Q satisfies

$$\sigma = \sqrt{\det Q} \circ F, \quad (3.71)$$

where $F : \Omega \rightarrow \Omega$ are the harmonic coordinates for the original conductivity σ . We remind the reader that F satisfies an elliptic boundary value problem, just as does $G = F^{-1}$,

$$\begin{aligned} \operatorname{div} \left(\frac{Q}{\sqrt{\det Q}} \nabla G \right) &= 0, & x \in \Omega, \\ G &= x, & x \in \partial\Omega, \end{aligned} \quad (3.72)$$

consequential to Theorem 3.1. In \mathbb{R}^2 , $F, F^{-1} \in H^1(\Omega)$, and H^1 embeds in C^α [10]. Therefore, connected paths remain connected when transformed by F , and likewise by F^{-1} . Also, in \mathbb{R}^2 , $\det \nabla F > 0$ strictly, so simple paths remain simple when transformed by F , and again likewise for F^{-1} . This means that isotropic inclusions are in 1-1 correspondence with inclusions in $\det Q$. If a test domain has piecewise constant isotropic inclusions, then perfect divergence-free reconstruction will detect those inclusions.

Having detected the inclusions, one final task remains for the mathematician: determination of their location. Relationship (3.71) indicates that the location of inclusions will be distorted by the harmonic coordinates. Fortunately, (3.72) determines $G = F^{-1}$, and

given a perfect reconstruction of divergence-free Q , we can exactly determine σ . Testing the success of divergence-free EIT with respect to an experimental collection of isotropic inclusions is meaningful. Note that these arguments are restricted to \mathbb{R}^2 . As we discuss further in Chapter 4, in dimensions $d \geq 3$, simple paths can become crossed under F , and we have no proof of formula (3.72).

This test of divergence-free EIT falls short in three ways: first, it does not treat conductivities that are anisotropic at the finest scale; second, it makes no comment on the stability of divergence-free EIT; and third, it does not take advantage of the divergence-free representation of conductivity to directly represent real fine-scale conductivities, without transformation by F^{-1} . Recovering conductivities which are anisotropic at the finest scale cannot be done uniquely: we cannot hope to recover the experimenter's inclusions. The best we can hope to do is to identify elements of the equivalence class determined by the Neumann-to-Dirichlet map. The two natural choices are the isotropic conductivity and the divergence-free conductivity. Since the divergence-free conductivity has good homogenisation properties, we expect this to be a good choice, since well-resolved divergence-free conductivities produce solutions close to their poorly-resolved approximations. This is different than suggesting the divergence-free conductivity is close to the experimenter's conductivity, which may not be directly homogenisable.

The second issue is stability: can two similar Neumann-to-Dirichlet maps be consistent with substantially different divergence-free conductivities? It appears that the answer is no, or at least, the difference can be no worse than that of an isotropic reconstruction. Owing to relation (3.71), and the fact that $\sigma > 0$, we have that the maximal eigenvalue of Q satisfies $|\lambda_{\max}| \in L^\infty$. This furthermore means that $\|Q\|_* \in L^\infty$ in common matrix norms of Q , where $\|\cdot\|_*$ represents one of the induced p -norms, the Frobenius norm, or the trace norm. The converse also holds, so in these common matrix norms, $\sigma \in L^\infty$ and $\|Q\|_* \in L^\infty$ are topologically equivalent, and Alessandrini's stability result in [13] will hold for divergence-free conductivities. That is, we approximate Q in a finite subspace of the set of functions where $\|Q\|_*$ is bounded, so

$$Q = \sum_{i=1}^N \xi_i(x) Q_i \tag{3.73}$$

for basis functions $\xi_i(x)$. Then, for Q_1, Q_2 in this subspace,

$$\| |Q_1 - Q_2|_* |_\infty \leq C(N) \|\Lambda_{Q_1} - \Lambda_{Q_2}\|_{\mathcal{L}(H^{-\frac{1}{2}}(\partial\Omega), H^{\frac{1}{2}}(\partial\Omega))}, \quad (3.74)$$

where $C(N)$ depends at worst exponentially on N , the size of the basis used to approximate Q . (Recall that $\|\cdot\|_{\mathcal{L}(H^{-\frac{1}{2}}(\partial\Omega), H^{\frac{1}{2}}(\partial\Omega))}$ is an operator norm for NtD .)

Fortunately, bound (3.74) may not be as sharp as it is known to be for isotropic conductivities. Or, it may simply be that owing to the good homogenisation qualities of Q , N can be smaller than it could be for direct reconstructions of isotropic conductivities. In any case, it appears that reconstructing divergence-free conductivities is not a detriment to stability, even when it comes to evaluating sensitivity to missing or falsely detected inclusions. Whether reconstructing divergence-free conductivities can improve stability is not yet known.

Finally, we will take best advantage of divergence-free reconstructions when we have methods for directly interpreting which fine-scale conductivities are implied by their divergence-free counterparts. This means building an associative library. A divergence-free conductivity, for example, that is strongly anisotropic is likely to represent a conductivity laminated on the finest scale. The difficulty with interpreting divergence-free conductivities, as already stated, is their non-locality with respect to the fine scale. However, the relationship $\sigma = \sqrt{\det Q} \circ F$ suggests that there may exist approximate local relationships, up to composition with the harmonic coordinates. In practise, we furthermore find that $F(x) \approx x$, so interpreting divergence-free conductivities may not be as difficult as the initial non-locality suggests. That is, many conductivities appear to be nearly divergence-free, a fact that may allow us to interpret the divergence-free reconstruction as one close to the true conductivity, even in the case where that conductivity may be anisotropic, and EIT reconstruction is known to be non-unique. In any case, further research in associating divergence-free to fine-scale isotropic conductivities will benefit not only EIT, but also the forward problem, which is currently limited by the need to compute the harmonic coordinates at a fine scale. An associative library could provide an appropriate shortcut to this procedure.

3.4 Unsuccessful reconstruction algorithms

This final section on EIT in two dimensions itemises the several unsuccessful reconstruction algorithms we have attempted over the last several years. In order, we discuss reconstruction of the conductivity as a resistor network, the down-scaling of the conductivity by inverting the principles of metric-based up-scaling, an adaptation of Calderón’s original algorithm to the $s(x)$ parameterisation, and the incorporation of so-called variational constraints, already successful in the EIT problem for scalar conductivities. While unsuccessful as reconstruction methods, each case illustrates some further understanding of conductivity and the inverse problem. Our difficulties show the importance in resolving the conductivity in the cone of $\sigma > 0$, Q positive-definite, or $s(x)$ convex, that the conductivity not be otherwise too constrained, and that the reconstruction be regularised in some way.

3.4.1 Resistor network interpretation

In the case of linear finite elements, we have shown that it is natural to report elements of the stiffness matrix as values on edges of a triangulation. In turn, we could use this representation to pose the inverse problem as an inverse conductance network problem, where the stiffness matrix value on each edge represents the conductance of the edge. This is a topological interpretation of the problem, and gives further insight into what really is measurable about conductivity.

In the case of electric conductance, the physical interpretation of this model is that if $u_i - u_j$ is the potential difference between endpoints of edge ij , then $J_{ij} = q_{ij}(u_i - u_j)$ is the electric current flowing through the edge from vertex i to vertex j . The conservation law at each vertex is exactly the same,

$$\sum_{j \in \mathcal{N}(i)} J_{ij} = \sum_{j \in \mathcal{N}(i)} q_{ij}(u_i - u_j) = f_i, \quad (3.75)$$

where f_i is the current injected into the network at vertex i . On interior vertices, $f_i = 0$, and on boundary vertices, f_i is an integral of the test current we use to measure the Neumann-to-Dirichlet map for the network. There is no buildup of charge in the network: indeed, the row-sum condition on the stiffness matrix implies that for solvability, $\sum_i f_i = 0$.

We emphasise that there is no manifold in this statement of the inverse resistor network

problem: there is only a network of edges and vertices. However, recalling Glickenstein's development [42], summarised by us in Section 2.3, q_{ij} is the ratio between the length of a dual edge to the length of its edge in a weighted Delaunay triangulation. This is reminiscent of the expression for the conductance of a one-dimensional bar, namely

$$\text{Conductance} = \frac{1}{\text{Length}} \times \text{Cross-section Area} \times \text{Conductivity}, \quad (3.76)$$

where we interpret Cross-section Area as the dual edge length, and Length as the edge length. If Conductivity, an isotropic material parameter in this setting, were constant, this suggests that conductances q_{ij} are determined not by any material property, but rather by the metric geometry of the mesh. However, as a model for the continuum, referring to Theorem 3.1, while we can always find an isotropic σ in the equivalence class determined by Λ_σ , that unique isotropic conductivity may not be constant. Hence, conductivity is not simply a discrete metric over the triangle mesh, at least not as a model of a metric in the continuum.

While our idea of conductivity as a metric may not be consistent with the continuum view, this affects only our interpretation of the stiffness matrix, and there appears to be no difficulty in treating the problem as an inverse resistor network problem. This problem has been studied by Curtis and Morrow to the extent that these authors provide a direct (though ill-conditioned) method for solving the inverse resistor problem on specific networks [31, 32, 33]. Indeed, in the discrete setting, there appears to be plenty of boundary data from which to deduce the edge values in the interior. With V_B boundary vertices, the Neumann-to-Dirichlet map has

$$|\Lambda_\sigma| = \frac{V_B(V_B - 1)}{2} \quad (3.77)$$

degrees of freedom. This accounts for the fact that this linear map is symmetric with row-sum zero. Recalling our calculation from Section 3.3, the number of edges E in a triangulation having V_I interior vertices is

$$E = 3V_I + 2V_B - 3. \quad (3.78)$$

Unless our mesh choice is pathological, $V_I = \mathcal{O}(V_B^2)$. However, we find in practise that

the constant for this growth relation is less than one-half, and the number of edges we are solving for does not exceed the cardinality of the boundary data.

Such is the case for unstructured triangular networks: Curtis and Morrow go to lengths to construct patterned networks that not only meet cardinality criteria, but also can be solved directly from boundary data. On these special networks, conductances are determined one at a time, with known values used in the computation of succeeding conductances in their iteration. It happens that error builds exponentially as the solve proceeds, but with exact arithmetic, their method is correct. The direct method is sensitive to error in boundary data, but still, recent work uses the circular planar networks of Curtis and Morrow, combined with a parameterisation of conductivity using the finite-volume method to compute isotropic conductivities on circular domains [25].

There are few examples of triangular networks meeting the exactness criteria of Curtis and Morrow, so we attempt to solve the inverse network problem on unstructured networks, minimising the boundary least-squares functional, our optimisation (3.38) of Section 3.3 without the regularisation term. As we are interested in resolving the homogenisation divergence-free conductivity, we search the space of interpolants s_i rather than the space of edge conductivities, q_{ij} . Owing to the geometric information in the relationship between q_{ij} and s_i , this parameterisation makes the link between the topological connectivity of the network and its geometry as a triangular mesh embedded in Ω . Were we to solve the inverse network problem for the q_{ij} without the divergence-free constraint, the locations of the network vertices would be arbitrary. Choosing the vertex locations independently of solving for the q_{ij} is equivalent to choosing a particular element of the equivalence class determined by the Neumann-to-Dirichlet map. Optimising for the divergence-free element avoids this ambiguity. Parameterising the problem by the s_i returns us to the more successful method already presented, where we furthermore find that some regularisation is needed.

Recognising that the boundary least-squares functional is not convex, we attempt solving this problem using simulated annealing, a stochastic global optimisation method for non-convex problems. Our problem is constrained to have positive conductances, $q_{ij} > 0$, and we are fortunate for the development of constrained simulated annealing algorithms. Simulated annealing is introduced in the classic text of Nocedal and Wright [62]; see [87] for an example use of Lagrange multipliers in implementing an inequality-constrained constrained

algorithm.

Using non-convex optimisation is no minor undertaking, as we found methods for constrained simulated annealing to be sensitive to our choice of mesh. Simulated annealing works by randomly adjusting the problem variables, occasionally taking steps that increase the value of the cost function. The idea is to prevent the solve from settling in local minima. The frequency of these uphill steps decreases as the solve proceeds, with the preference against uphill steps decreasing at a problem-specific so-called cooling rate. A large cooling rate tends to trap the solution in a local minimum, while a small cooling rate lets the solve run for so long that it never seems to converge. In our case, the optimal cooling rate is sensitive to the choice of mesh geometry, even for a constant mesh size, so in practise, we choose a small cooling rate. The algorithm then takes so long to converge that it can only be applied in practise to coarse meshes.

The trouble with simulated annealing may be due to a sensitivity to finding edge values near zero. As in the regularised convex optimisation, this is an indication that the mesh connectivity requires a non-convex interpolation, even though the underlying $s(x)$ is convex. One way to correct this is to allow the algorithm to flip edges of the triangulation throughout the solve: unfortunately, this destabilises the annealing solver in practise.

Solving the problem directly as an inverse network problem does not give a unique solution unless the network is one of a small class of networks, or unless the edge conductances are divergence-free. Imposing the divergence-free constraint recovers the same boundary-error term in the objective that we have in the continuum interpretation, and, save the regularisation term, there is no difference between the two interpretations, save the lack of regularisation in the network model. In neither case does a numerical solver converge without some form of regularisation.

3.4.2 Down-scaling the conductivity

In our early attempts at incorporating metric-based up-scaling into EIT, we would first solve for a coarse stiffness matrix q_{ij} , and then down-scale the conductivity to produce an estimate of the conductivity on the fine scale. The first of these attempts at downscaling is an iteration to determine isotropic σ and the σ -harmonic coordinates F . Recognising a numeric instability in down-scaling and determining the F -harmonic coordinates simultaneously,

our second attempt first down-scales Q before attempting to compute isotropic σ . Neither method is particularly successful, but both failures show us that confining our admissible conductivities to those in the equivalence class determined by the Neumann-to-Dirichlet map is indeed restrictive.

The iteration for the σ -harmonic coordinates is in two steps. First, for given q_{ij} on a coarse mesh, we determine the σ on a fine mesh that is consistent with the given q_{ij} and has optimal smoothness. We experiment with several smoothness metrics, of which the total-variation of σ seems to give the best result. Hence, this first step is to solve

$$\begin{aligned} & \underset{\sigma \in \mathcal{P}_0^h}{\text{minimise}} && \|\sigma^{(n)}\|_{\text{TV}} \\ \text{subject to} &&& q_{ij} = - \int_{\Omega} \nabla(\varphi_i \circ F^{(n)})^T \sigma^{(n)} \nabla(\varphi_j \circ F^{(n)}), \\ &&& \sigma^{(n)} > 0, \end{aligned} \tag{3.79}$$

where we use superscripts n to denote the iteration number. σ is piecewise constant over triangles. This first step requires that we assume the harmonic coordinates F . Once we have σ , the second step is to update the harmonic coordinates by solving

$$\begin{aligned} -\operatorname{div}(\sigma^{(n)} \nabla F^{(n+1)}) &= 0 && x \in \Omega, \\ F^{(n+1)} &= x && x \in \partial\Omega. \end{aligned} \tag{3.80}$$

We initialise the iteration by choosing $F^{(0)} = x$. The trouble is that this iteration does not converge, although after only a single iteration, this iteration does produce a reasonable result. See Figure 3.10 for an example. Figure 3.11 is a plot of several convergence metrics as an iteration proceeds. We see the best images in the first few iterations, where the constraints that the given q_{ij} , our input data, match the stiffness matrix for our current guess for Q , are most badly violated.

Suspecting that the non-convergence of the iteration is due to the simultaneous down-scaling of σ and the computation of the harmonic coordinates, as an alternative, we attempted to first resolve the divergence-free conductivity Q at the fine scale, then compute




Figure 3.10: Output of the harmonic coordinate iteration. The figure on the top left is the coarse mesh produced by simulated annealing, the input to the harmonic coordinate iteration. Left to right, top to bottom, the remaining three images show the progression of the iteration at 1, 10, and 20 steps, showing its instability. The true conductivity is that of Figure 3.5.

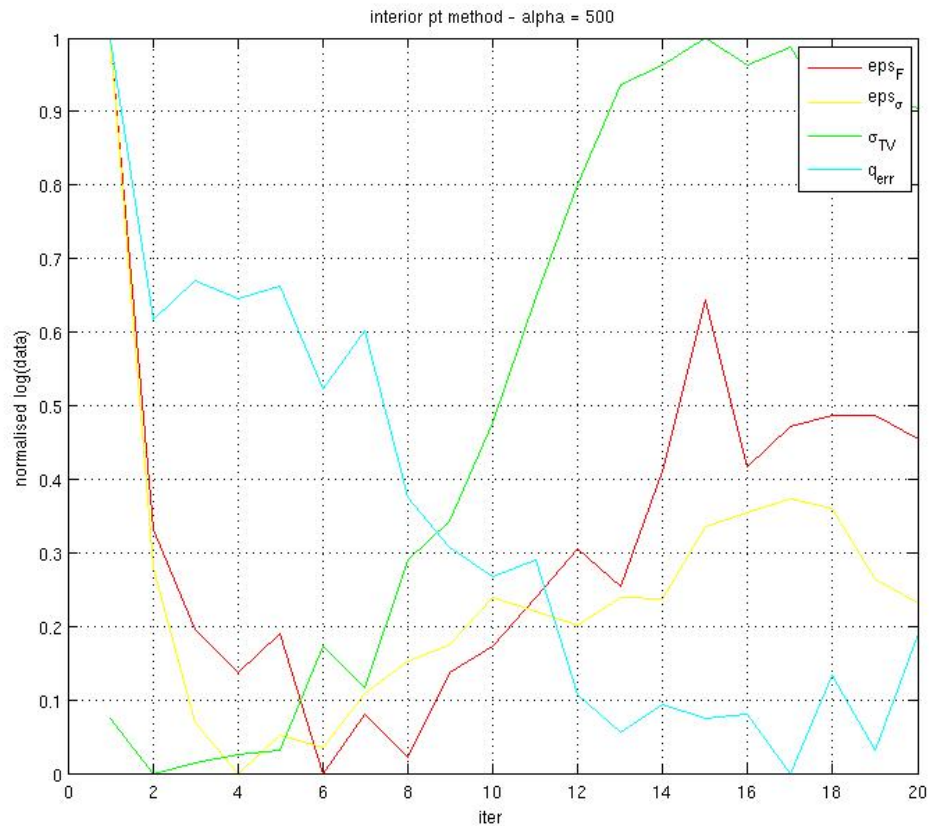


Figure 3.11: Convergence metrics of the harmonic coordinate iteration. eps_F and eps_σ are the change in F and σ between iterations, σ_{TV} is the total variation smoothness of σ , and q_{err} is the violation of the q_{ij} constraints in equation (3.79) following the solve for F . We see the best images in the first few steps, where the q_{ij} constraints are most violated.

σ and F . This means solving

$$\begin{aligned} & \underset{Q \in \mathcal{P}_0^h}{\text{minimise}} && \|\nabla Q\|_* \\ & \text{subject to} && q_{ij} = - \int_{\Omega} \nabla \varphi_i^T Q \nabla \varphi_j, \\ & && Q \text{ positive-definite and divergence-free.} \end{aligned} \tag{3.81}$$

This is a semidefinite convex program, provided the smoothness measure $\|\nabla \cdot\|_*$ for Q is convex. (For example, $\|\nabla Q\|_*$ could be the total variation of Q .) From a practical perspective, the code to implement this program is quite complicated, and the solver tends not to converge. The semidefinite solver has difficulty finding any feasible Q , let alone one that minimises any smoothness criterion.

Finally, we have attempted solving directly for $G = F^{-1}$ using our coarse-scale Q , and equations (3.26).

$$\begin{aligned} \operatorname{div} \left(\frac{Q}{|\det Q|^{1/2}} \nabla G \right) &= 0, & x \in \Omega, \\ G &= x, & x \in \partial\Omega. \end{aligned} \tag{3.26}$$

Although the harmonic coordinates F vary on a fine scale, the hope is that their coarse approximation will be sufficiently close to the true solution that we can initialise the F - σ iteration and rely on more local convergence of the iteration. Unfortunately, we observe that solutions to (3.26) on the coarse scale fail to satisfy the consistency condition implied by (3.35), which relates G_1, G_2 through Q , indicating that the G we recover is not terribly close to the true solution. Moreover, even if we solve for Q on a fine scale, using the regularised boundary-error optimisation (3.38), the consistency condition G is sufficiently well-violated that the solutions to (3.26) are nonsensical. A critical step in solving equations (3.26) is computing a value for $\det Q$ to build a stiffness matrix. The values of this determinant can be approximated from $s(x)$ using the piecewise-quadratic approximation for $s(x)$ given, as we recall, by the stencil in Figure 3.2. We suspect that it is the inaccuracy of these values which spoils the solve for $G(x)$.

The boundary data defines an equivalence class of conductivities, and we use this boundary data to produce a coarse-scale stiffness matrix. The lack of regularisation notwithstanding, all of these failures suggest that the coarse-scale stiffness matrix is too far from any

fine-scale conductivity, either scalar or divergence-free, that we can step toward the fine scale using a local iteration or convex optimisation. By starting with a coarse scale stiffness matrix, we are, in effect, searching an equivalence class of conductivities, and within this equivalence class, there are no good guesses we can make for fine-scale conductivities sufficiently close to a true conductivity that we can take advantage of any supposed local convergence properties.

3.4.3 Calderón iteration

In his seminal paper on EIT, Calderón showed that the Neumann-to-Dirichlet map is locally in 1-1 correspondence with an isotropic conductivity [29]. Moreover, the proof is constructive, providing a linearised reconstruction algorithm: as long as the initial guess for the conductivity is not too far from the actual conductivity, Calderón's algorithm converges. The method is so successful, that until the quite recent development of the D-bar method, the Calderón algorithm was used in EIT medical devices [40, 49]. One of our attempts at reconstruction is to mimic Calderón's method, in our case recovering divergence-free conductivities by our $s(x)$ parameterisation. The adapted algorithm fails: unless the initial guess is extremely close to the true solution, the algorithm tends to lose positivity in Q , computing non-convex $s(x)$. However, the idea is so simple, and the Calderón algorithm so central to practical EIT, that we briefly present the idea here.

Calderón's insight is that if a supposed conductivity \tilde{Q} is close to the true conductivity Q , then the solution \tilde{u} in the interior of the domain is close to the true solution u . That is, if we solve

$$\begin{aligned} -\operatorname{div}(\tilde{Q}\nabla\tilde{u}) &= 0, & x \in \Omega, \\ \tilde{u} &= g, & x \in \partial\Omega, \end{aligned} \tag{3.82}$$

then $\tilde{u} \approx u|_{\Omega}$ when $\|\tilde{Q} - Q\|$ is small. This observation motivates computing

$$\begin{aligned} \int_{\Omega} \nabla u^{(1)} Q \nabla u^{(2)} - \int_{\Omega} \nabla u^{(2)} \tilde{Q} \nabla \tilde{u}^{(1)} \\ = \int_{\partial\Omega} u^{(1)} (Q \nabla u^{(2)})^T d\hat{n} - \int_{\partial\Omega} u^{(2)} (\tilde{Q} \nabla \tilde{u}^{(1)})^T d\hat{n} \end{aligned} \tag{3.83}$$

$$= \int_{\partial\Omega} g^{(1)} f^{(2)} - \int_{\partial\Omega} g^{(2)} \tilde{f}^{(1)}. \tag{3.84}$$

The pairs $(f^{(i)}, g^{(i)})$ are Neumann-Dirichlet solution pairs obtained as boundary measurements of the Neumann-to-Dirichlet map, and $\tilde{f} = (\tilde{Q}\nabla\tilde{u})^T\hat{n}$ in the solution to the approximate problem (3.82). Setting $\delta Q = Q - \tilde{Q}$, we make the approximation $u \approx \tilde{u}$, and we have an linear equation for updating Q based only on boundary data:

$$\int_{\Omega} \nabla u^{(1)} \delta Q \nabla u^{(2)} = \int_{\partial\Omega} g^{(1)} f^{(2)} - g^{(2)} \tilde{f}^{(1)}. \quad (3.85)$$

(Compare this to the similar update rule, equation (3.54) on page 74 for the nonlinear optimisation, obtained by the primal-adjoint method.) Calderón uses complex geometric optics solutions to make all of this rigorous for sufficiently small δQ .

In the discrete setting, we have the linear system

$$\sum_{i \sim j} \delta q_{ij} (u_i^{(a)} - u_j^{(a)}) (u_i^{(b)} - u_j^{(b)}) = h^{(ab)}, \quad (3.86)$$

where the $h^{(ab)}$ are a quadrature of the boundary data for solution pairs a, b . Using our solutions u_i from our supposed q_{ij} , this gives us one equation for each solution pair a, b . As we have commented in the context of the resistor network interpretation, most reasonable triangular meshes have enough boundary nodes that the cardinality of the discrete Neumann-to-Dirichlet is greater than the number of edges, and this linear system, though ill-posed, is in practise overdetermined.

Unfortunately, although we regularise equations (3.86), using regularised least-squares or a truncated singular value decomposition, this method tends to produce search directions δq that give negative edge conductivities. Edge flipping is again destabilising. All of this suggests this linearised update rule is inappropriate for finding divergence-free conductivities. We may also be significantly violating the approximation $u = \tilde{u}$: Calderón assumes that the boundary test potentials contain only low-order Fourier modes. By interpolating solutions as piecewise linear functions, we may be violating this assumption too grossly.

3.4.4 Variational constraints

The EIT reconstruction algorithms most robust in the presence of noise is perhaps the work of Borcea [26], which adds so-called variational constraints to the least-squares boundary-error minimisation following earlier work of Berryman and Kohn [22]. The addition of these

constraints tends to better localise high-conductivity inclusions and reduce the reconstruction's sensitivity to noise in the boundary data.

The constraints rely on the fact that the energy lost to the conductance of the domain equals the energy expended at the boundary, which is the inner product of the voltage and the current density over $\partial\Omega$. We understand this relationship as integration by parts of the Dirichlet energy:

$$\int_{\Omega} \nabla u^T Q \nabla u = \int_{\partial\Omega} u (Q \nabla u)^T d\hat{n} = \int_{\partial\Omega} fg. \quad (3.87)$$

Recognising the monotonicity between the Dirichlet energy and the conductivity Q is the starting point for the variational constraints, whereupon we restrict Q to a set such that its Dirichlet energy either exceeds or is exceeded by the known boundary energy. Parameterising Q by $s(x)$ puts these constraints in our framework of divergence-free conductivities. Consistent with our other findings, however, combining the divergence-free constraint with the variational constraints makes it difficult for our nonlinear convex optimiser to find a feasible conductivity, let alone one that minimises the boundary energy.

Variational constraints identify two sets of conductivities, the Dirichlet feasible set and the Thomson feasible set. A conductivity σ^* is said to be Dirichlet feasible for boundary potential g if

$$(g, (\Lambda_{\sigma})^{-1}g)_{\partial\Omega} = \min_{u=g|_{\partial\Omega}} \int_{\Omega} \nabla u^T \sigma \nabla u \geq (g, (\Lambda_{\sigma^*})^{-1}g)_{\partial\Omega}, \quad (3.88)$$

and is Thomson feasible for boundary current f if

$$(f, \Lambda_{\sigma}f)_{\partial\Omega} = \min_{\frac{\partial}{\partial n}(\sigma u)=f|_{\partial\Omega}} \int_{\Omega} \nabla u^T \sigma \nabla u \geq (f, \Lambda_{\sigma^*}f)_{\partial\Omega}. \quad (3.89)$$

At a solution, where $\Lambda_{\sigma^*} = \Lambda_{\sigma}$, equality holds for both the Dirichlet and Thomson inequalities, and it can be shown that solutions to the inverse problem lie on the intersection of the boundaries of both sets. The important observation in [22, 26] is that the interior of the Dirichlet infeasible set lies entirely in the Thomson feasible set, and the interior of the Thomson infeasible set lies entirely in the Dirichlet feasible set. This means that there are relatively few conductivities which can be both Thomson and Dirichlet feasible. See Figure 3.12.

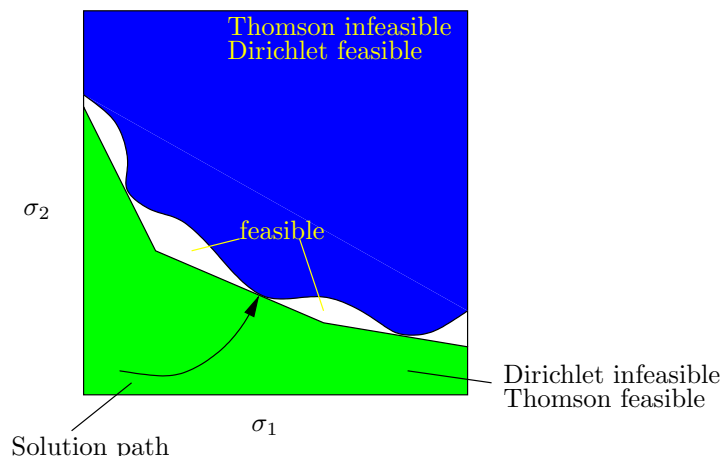


Figure 3.12: Feasibility constraints for EIT for a conductivity σ parameterised by only two conductivities, σ_1, σ_2 . The Thomson feasible region (green \cup white) is non-convex, while the Dirichlet feasible region (blue \cup white) is convex. The solution lies at the intersection between the two feasible sets, and the solution path is entirely within the Thomson feasible region. Similar images appear in [22, 26].

As the boundary-error term is minimised, the conductivity approaches the intersection of the boundaries of both the Thomson and Dirichlet infeasible sets. There are at least two ways to use Dirichlet or Thomson feasibility to improve the boundary least-squares error optimisation. The first is to restrict the solution to the set of Thomson infeasible conductivities, whereupon the convex minimisation algorithm approaches the boundary of the Thomson feasible set. The second is to restrict the solution to the set of Dirichlet infeasible conductivities. This second choice is preferred by Borcea for numerical reasons, as it takes advantage of the convexity of the Dirichlet feasible set. (See again Figure 3.12.) The complete constrained nonlinear program for an isotropic conductivity σ is

$$\begin{aligned}
 & \text{minimise} && \sum_{i=1}^M \|\Lambda_\sigma f^{(i)} - g^{(i)}\| + \nu \|\nabla \sigma\|, \\
 & \text{subject to} && (g^{(i)}, (\Lambda_\sigma)^{-1} g^{(i)})_{\partial\Omega} \leq (g^{(i)}, f^{(i)})_{\partial\Omega}, \quad i = 1 \dots M \\
 & && \sigma > 0.
 \end{aligned} \tag{3.90}$$

Borcea uses a convex optimisation algorithm to solve this problem, and does so successfully, despite the fact that the problem is non-convex. Neither the Dirichlet constraint nor the boundary-error term in the objective are convex (although the positivity constraint on σ is).

This non-convexity may be the source of our difficulty in applying this method to determine divergence-free anisotropic conductivities. Our divergence-free constraint, in addition to the variational constraint, may be beyond the flexibility of our convex optimisation algorithm. That said, Borcea does point out several other equivalent programs that use variational constraints, and we have yet to explore all of the possibilities numerically.

Chapter 4

Extensions

Further work in discretizing homogenised conductivities obtained by metric-based up-scaling could include extensions to other systems with linear constitutive laws, such as linearised elasticity, or addressing the issue of optimal design by homogenisation. However, instead of focusing on the applications, we propose in this final chapter the extension of discretized up-scaled conductivities to higher dimensions, where homogenisation is not so straightforward.

The initial motivation for this extension comes from recent work of Stern et al, which provides a family of numerical schemes for the free space Maxwell's equations. These schemes preserve energy and momentum even for coarse discretizations of time and space [79]. In the presence of variable electric permittivity and magnetic permeability, however, coarse spatial discretizations do not faithfully represent materials having fine-scale structure, an inference from our understanding of conductivity up-scaling. A homogenisation method for Maxwell's equations based on integral averages of the fields has already been proposed [78], although the method is much like representative volume averaging (RVE) in that it not only requires a separation of coarse and fine scales, but also dictates what the fine scale should be, without regard for convergence. Fortunately, a formalism already exists for applying metric-based up-scaling to the spatial components of hyperbolic PDE [64], a formalism that in two spatial dimensions gives the same parameterisation $s(x)$ and its discrete analogue we have developed throughout this thesis.

The natural discretization of Maxwell's equations is in four-dimensional space-time. Just as we have seen for conductivity, the material parameters permittivity and permeability define a metric on the discretized solution mesh. Looking even further ahead, studying the effect of varying this metric may provide a stepping stone for extending the variational methods of Stern et al to numerical relativity. However, remaining in context of this work on homogenisation, we will close by discussing some aspects of metric-based up-scaling in three dimensions, both over the continuum, and in the discrete setting.

4.1 Parameterising conductivities in \mathbb{R}^3

Suppose we extend our parameterisation $s(x)$ of up-scaled conductivities to three dimensions. We will still have a conductivity $Q = F_*\sigma$ that is divergence-free, and we can take advantage of its redundancy to parameterise Q . In two dimensions, Q is symmetric, giving three degrees of freedom, from which we subtract the two divergence-free constraints, leaving Q in a one-parameter family, our functions $s(x)$. In three dimensions, Q symmetric gives six degrees of freedom, from which we subtract the three divergence-free constraints, leaving Q in a three parameter family of functions. Naming these three functions q, r, s , one way to parameterise Q is

$$Q = \begin{pmatrix} \partial_{yy}q + \partial_{zz}s & -\partial_{xy}q & -\partial_{zx}s \\ -\partial_{xy}q & \partial_{zz}r + \partial_{xx}q & -\partial_{yz}r \\ -\partial_{zx}s & -\partial_{yz}r & \partial_{xx}s + \partial_{yy}r \end{pmatrix}. \quad (4.1)$$

In dimension d , there are $d(d-1)/2$ such parameters, one associated to each second derivative $\partial_{x_i x_j}, i < j$. Setting, say, $r, s = 0$ in the three-dimensional case, indicating zero conductivity in the z -direction, gives the divergence-free conductivity in two dimensions.

A constraint determining the positivity of Q is not clear, however. In two dimensions, we identify convexity of $s(x, y)$ with the positivity of Q . In three dimensions, writing

$$Q_{xy} = \begin{pmatrix} \partial_{yy}q & -\partial_{xy}q & 0 \\ -\partial_{xy}q & \partial_{xx}q & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (4.2)$$

and likewise for Q_{yz} associated to r and Q_{zx} associated to s , we have $Q = Q_{xy} + Q_{yz} + Q_{zx}$. We see that convexity of q, r, s in the xy, yz and zx planes, is sufficient for positivity of Q . However, this is not a necessary condition, as we see by example: choose

$$Q = \alpha \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} + \beta \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} + \gamma \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (4.3)$$

and set $\alpha = \beta = 1, \gamma = -1/2$. Q_{xy}, Q_{yz} are positive-definite, but even though Q_{zx} is

negative-definite, Q is still positive-definite. Moreover, when the up-scaled Q is constant, the harmonic coordinates F are just $F(x) = x$, and Q is the same as the original conductivity: we can't blame this result on metric-based up-scaling. We even see in this example that Q uniquely determines q, r, s (up to affine functions, which don't affect their convexity), so we can't change our choices of q, r, s such that they are all convex. Failure to represent Q in this simple case by convex functions suggests that convexity of the functions generating Q is not the correct way to determine its positivity.

More difficult still, with the added freedom in three dimensions, there arises the possibility that when up-scaled, a conductivity operator can lose ellipticity. With the added freedom afforded to the flux field in \mathbb{R}^3 , it's possible at the fine scale for the potential field to lose monotonicity for boundary conditions $u = x|_{\partial\Omega}$. At the coarse-scale, the geometry of the fine-scale conductivity giving rise to this inversion is lost, and is replaced by up-scaled conductivity that is not everywhere positive-definite. Briane et al give an example of a periodic lattice of fine-scale interlocking rings which has this property in asymptotic homogenisation, a difficulty likewise seen in metric-based up-scaling [27, 65]⁴.

If a homogenised conductivity does not necessarily form an elliptic operator, this also raises questions for the inverse problem. Very little is known about anisotropy in general in the inverse problem for $d \geq 3$, let alone a result identifying members of an equivalence class determined by the Neumann-to-Dirichlet map. Here is where a study of what we can actually measure could contribute to understanding homogenisation. If all we can measure about a piece of conductive material is its Neumann-to-Dirichlet map, then we are blind to its internal structure. Yet, if a unique divergence-free conductivity that is consistent with the boundary map can be found, then metric-based up-scaling is in some way producing a sensical up-scaled conductivity. Despite the local loss of ellipticity, it seems that we should still be able to solve for homogenisation solutions using the up-scaled conductivity since the up-scaled domain has a well-defined Neumann-to-Dirichlet map.

The situation is analogous to the formation of shocks in nonlinear hyperbolic conservation laws such as Burgers' equation. Solving Burgers' equation by the method of characteristics gives solutions that are well-defined in the characteristic variables of the problem, but when projected back to the original space-time coordinates give multivalued functions.

⁴Fortunately, Briane et al also show that this loss of ellipticity cannot happen in any dimension if the material is merely laminated.

Our interpretation is that this gives shocks (discontinuities) in the solutions, although the conservation law, an integral law, implied by the original PDE is preserved. Likewise, conductivity may be positive-definite in its original coordinates, but when pushed forward by its harmonic coordinates (a nonlinear operation) it loses positivity, but it does not lose its conservation properties, and we can still observe a sensible Neumann-to-Dirichlet map. It may be simply that we require the Neumann-to-Dirichlet map to be positive-definite, rather than requiring Q itself to be pointwise positive-definite. Perhaps this will be a fruitful perspective for studying non-ellipticity due to homogenisation.

4.2 Discrete parameterisations of conductivities

An answer for how to restrict Q in three dimensions so that it maintains its conservation properties may be discovered from the discrete view of the functions q, r, s which generate the up-scaled conductivity. One way to expand our view of divergence-free conductivity is to frame our parameterisations for q, r, s and for the resulting stiffness matrix elements in terms of a discrete differential complex. Such complexes are introduced in [34] for discrete exterior calculus, and in [16, 17] where they are used for constructing mixed finite element families having favourable stability properties.

In the continuum, differential forms and their exterior derivatives together form a complex. Setting $\Lambda^{(k)}(\Omega)$ the set of k -forms over manifold Ω , and $d^{(k)}$ the exterior derivative on elements of $\Lambda^{(k)}(\Omega)$, we form the complex where the dimension⁵ of Ω is n ,

$$\mathbb{R} \hookrightarrow \Lambda^{(0)} \xrightarrow{d^{(0)}} \Lambda^{(1)} \xrightarrow{d^{(1)}} \dots \xrightarrow{d^{(n-1)}} \Lambda^{(n)} \longrightarrow 0. \quad (4.4)$$

This is the de Rham complex. In the continuum, it satisfies $d^{(k+1)} \circ d^{(k)} = 0$, a consequence of Stokes' theorem. Moreover, provided the domain Ω is "reasonable" (Ω simply connected is sufficient), the dual of this statement, the Poincaré lemma, holds: If $\omega \in \Lambda^{(k+1)}$, and $d\omega = 0$, then there is $\alpha \in \Lambda^{(k)}$ such that $\omega = d\alpha$. If forms are discretized properly, then a discrete complex analogous to (4.4) has these same consistencies.

We discretize the de Rham complex by noting that forms are quantities that can be

⁵We must switch notation here: d is so commonly used for the exterior derivative that we cannot also use it for the dimension of the space.

integrated over regions of sub-manifolds. Following [34], discrete exterior calculus recognises that these regions are quantised when the manifold is a simplicial complex (such as a triangulation), so discrete differential forms are defined only on elements of the simplicial complex. For example, 1-forms are defined on edges and 2-forms are defined on faces, corresponding to their integration over paths and areas. Most importantly, there is no interpolation of forms between simplicial elements in this context. (This perspective is specific to [34] – see [16] for an example of a framework where forms are interpolated.) Avoiding interpolation contributes to the coordinate-free representation of forms as values on vertices, edges, faces, and so forth.

Discrete exterior calculus discretizes the exterior derivative, completing a discrete construction of the complex (4.4). In this setting, $d^{(k)}$ is a linear operation from k -forms to $(k+1)$ -forms which can be represented as a matrix multiplication. To give an idea of the form of this matrix, if α is a 0-form with values α_i, α_j on vertices of oriented edge ij , then ω_{ij} , the value of $\omega = d\alpha$ on edge ij , is $\alpha_j - \alpha_i$. Similar definitions exist for higher-order exterior derivatives, with the design that every discrete exterior derivative satisfy a discrete Stokes' theorem. The important consequence of the discrete Stokes' theorem is that by their definition, every discrete exterior derivative satisfies $d^{(k+1)} \circ d^{(k)} = 0$. This means that even in the discrete setting, (4.4) is a complex. Further work in [35] shows that a discrete Poincaré lemma also holds exactly in this framework.

Looking back at conductivities in \mathbb{R}^2 , we see that we have formed a differential complex, though it is not the de Rham complex. We have

$$\mathbb{R} \hookrightarrow S \xrightarrow{A} Q \xrightarrow{B} V \longrightarrow 0. \quad (4.5)$$

Working immediately with discretized quantities, A is the linear operation that computes the stiffness matrix $q_{ij} \in Q$ from the parameterisation $s_i \in S$. B is the linear operation that computes the divergences $\mathbf{v}_i \in V$ of the stiffness matrix at interior vertices by equation (2.10):

$$\mathbf{v}_i = \sum_{j \in \mathcal{N}(i)} q_{ij} (\mathbf{x}_i - \mathbf{x}_j). \quad (4.6)$$

By construction, we have $B \circ A = 0$, and so (4.5) is a complex. We know in the continuum that this complex supports an analogue to the Poincaré lemma: if a conductivity is

divergence free, then we can parameterise it by some $s(x)$, unique up to the addition of affine functions. Numerical experiments suggest that this is true in the discrete setting, but we do not yet have a proof. Arnold et al identify the continuous analogue of (4.5) as the elasticity complex in their application of this complex to produce stable mixed finite element methods for two-dimensional elasticity [17, 18]. They use higher-order elements than our parameterisation of $s(x)$, and it would be interesting to know if our low-order parameterisations give similarly stable methods.

It remains to be seen if divergence-free stiffness matrices can be incorporated into higher-order complexes, either for the purpose of furthering higher-dimensional linearised elasticity, or representing homogenised conductivities. Arnold et al claim that no such framework in higher dimensions yet exists [17]; a good starting point may be to establish the relationship between piecewise linear interpolants for q, r, s and the conductivity stiffness matrix. For example, just as we developed a coordinate-free relationship between the s_i and q_{ij} in \mathbb{R}^2 , we should be able to do the same in higher dimensions, removing the rectangular bias of our Q_{xy}, Q_{yz}, Q_{zx} .

For homogenisation, the benefit of extending our parameterisation to higher dimensions could be remarkable, as we would reap the numerical rewards of operator compression to a variety of scales, increasingly important in higher dimensions. We may be able to form a discrete perspective to decide which divergence-free conductivities give admissible up-scaled operators based on geometric properties of q, r, s . Appropriate discretization of up-scaled conductivities could also benefit the inverse problem in higher dimensions, where our understanding of anisotropic conductivities, largely limited to two dimensions, could be expanded by the question of what we are really measuring about conductivity by the response of a material at its boundary.

Bibliography

- [1] *ConfigFile*, <http://www-personal.umich.edu/~wagnerr/ConfigFile.html>.
- [2] *IpOpt, Interior Point Optimizer, version 3.2.4*, <https://projects.coin-or.org/Ipopt>.
- [3] *OpenGL*, <http://www.opengl.org>.
- [4] *Qt, open source edition, version 3.3*, <http://doc.trolltech.com/3.3/>.
- [5] BLAS, *Basic Linear Algebra Subroutines*, <http://www.netlib.org/blas/>.
- [6] CGAL, *Computational Geometry Algorithms Library, version 3.3.1*, <http://www.cgal.org>.
- [7] GNU *Linear Programming Kit*, <http://www.gnu.org/software/glpk/>.
- [8] GSL, *Gnu Scientific Library*, <http://www.gnu.org/software/gsl/>.
- [9] LSQR: *Sparse equations and least squares*, <http://www.stanford.edu/group/SOL/software/lsqr.html>.
- [10] Robert A. Adams and John J. F. Fournier, *Sobolev spaces*, (2003).
- [11] Giovanni Alessandrini, *Stable determination of conductivity by boundary measurements*, *Applied Analysis* **27** (1988), 153–172.
- [12] ———, *Examples of instability in inverse boundary-value problems*, *Inverse Problems* **11** (1997), no. 5, L17–L24.
- [13] Giovanni Alessandrini and Sergio Vessella, *Lipschitz stability for the inverse conductivity problem*, *Advances in Applied Mathematics* **35** (2005), 207–241.
- [14] Grégoire Allaire, *Shape optimization by the homogenization method*, Springer, New York, 2002.

-
- [15] Pierre Alliez, David Cohen-Steiner, Mariette Yvinec, and Mathieu Desbrun, *Variational tetrahedral meshing*, ACM Transactions on Graphics (SIGGRAPH '05) **24** (2005), no. 3, 617–625, <http://www.geometry.caltech.edu/pubs/ACYD05.pdf>.
- [16] Douglas N. Arnold, Richard S. Falk, and Ragnar Winther, *Compatible spatial discretizations*, The IMA Volumes in Mathematics and its Applications, vol. 142, Springer, New York, 2006, pp. 23–46.
- [17] ———, *Compatible spatial discretizations*, The IMA Volumes in Mathematics and its Applications, vol. 142, Springer, New York, 2006, pp. 47–67.
- [18] Douglas N. Arnold and Ragnar Winther, *Mixed finite elements for elasticity*, Numerische Mathematik **92** (2002), no. 3, 401–419.
- [19] Kari Astala, Matti Lassas, and Lassi Paivarinta, *Calderon's inverse problem for anisotropic conductivity in the plane*, Communications in Partial Differential Equations **30** (2005), no. 2, 207–224, DOI <http://arxiv.org/abs/math/0401410v1>.
- [20] Ivo Babuška, Gabriel Caloz, and John E. Osborn, *Special finite element methods for a class of second order elliptic problems with rough coefficients*, SIAM Journal of Numerical Analysis **31** (1994), no. 4, 945–981.
- [21] Ivo Babuška, Uday Banerjee, John E. Osborn, and Qiaoluan Li, *Quadrature for meshless methods*, Tech. report, Institute for Computational Engineering and Sciences, 2007, <http://www.ices.utexas.edu/research/reports/2007/0704.pdf>.
- [22] James G. Berryman and Robert V. Kohn, *Variational constraints for electrical-impedance tomography*, Physical Review Letters **65** (1990), no. 3, 325–328.
- [23] Liliana Borcea, *Electrical impedance tomography*, Inverse Problems **18** (2002), R99–R136.
- [24] Liliana Borcea, J. G. Berryman, and G. C. Papanicolau, *High contrast impedance tomography*, Inverse Problems **12** (1996), 935–958.
- [25] Liliana Borcea, Vladimir Druskin, and Fernando Guevara Vasquez, *Electrical impedance tomography with resistor networks*, (2007), submitted to Inverse Problems, http://www.caam.rice.edu/~borcea/Publications/EIT_NETS.pdf.

-
- [26] Liliana Borcea, Genetha Anne Gray, and Yin Zhang, *Variationally constrained numerical solution of electrical impedance tomography*, *Inverse Problems* **19** (2003), 1159–1184.
- [27] Marc Briane, Graeme W. Milton, and Vincenzo Nesi, *Change of sign for the corrector’s determinant for homogenization in three-dimensional conductivity*, *Archive for Rational Mechanics and Analysis* **173** (2004), no. 1, 133–150.
- [28] William L. Briggs, Van Emden Henson, and Steve F. McCormick, *A multigrid tutorial*, 2nd ed., SIAM, Philadelphia, 2000.
- [29] A. P. Calderón, *On an inverse boundary value problem*, *Seminar on Numerical Analysis and its Applications to Continuum Physics (Rio de Janeiro)*, Soc. Brasileira de Mathematica, 1980.
- [30] M. Cheney, D. Isaacson, and J. C. Newell, *Electrical impedance tomography*, *SIAM Review* **41** (1999), no. 1, 85–101.
- [31] Edward B. Curtis, D. Ingerman, and James A. Morrow, *Circular planar graphs and resistor networks*, *Linear Algebra and its Applications* **283** (1998), 115–150.
- [32] Edward B. Curtis and James A. Morrow, *Determining the resistors in a network*, *SIAM Journal on Applied Mathematics* **50** (1990), no. 3, 931–941.
- [33] ———, *Inverse problems for electrical networks*, World Scientific, 2000.
- [34] Mathieu Desbrun, Eva Kanso, and Yiyong Tong, *Discrete differential forms for computational modeling*, *SIGGRAPH’05: ACM SIGGRAPH’05 Courses (New York)*, ACM Press, 2005, DOI <http://doi.acm.org/10.1145/1198555.1198666>.
- [35] Mathieu Desbrun, Melvin Leok, and Jerrold E. Marsden, *Discrete Poincaré lemma*, *Applied Numerical Mathematics* **53** (2005), 231–248.
- [36] David C. Dobson, *Convergence of a reconstruction method for the inverse conductivity problem*, *SIAM Journal on Applied Mathematics* **52** (1992), no. 2, 442–458.
- [37] ———, *Estimates on resolution and stabilization for the linearized inverse conductivity problem*, *Inverse Problems* **8** (1992), 71–81.

-
- [38] David C. Dobson and Peter G. Kaup, *Monitoring underground flows with electromagnetic methods*, *Inverse Problems* **15** (1999), 597–613.
- [39] David C. Dobson and Fadil Santosa, *Resolution and stability analysis of an inverse problem in electrical impedance tomography: Dependence on the input current patterns*, *SIAM Journal on Applied Mathematics* **54** (1994), no. 6, 1542–1560.
- [40] Peter M Edic, Gary J Saulnier, Jonathan C Newell, and David Isaacson, *A real-time electrical impedance tomograph*, *IEEE Transactions on Biomedical Engineering* **42** (1995), no. 9, 849–859.
- [41] D. G. Gisser, D. Isaacson, and J. C. Newell, *Electric current computed tomography and eigenvalues*, *SIAM Journal on Applied Mathematics* **50** (1990), no. 6, 1623–1634.
- [42] David Glickenstein, *A monotonicity property for weighted Delaunay triangulations*, *Discrete and Computational Geometry* **38** (2007), no. 4, 651–664,
<http://math.arizona.edu/~glickenstein/monotonic.pdf>.
- [43] Allan Greenleaf, Yaroslav Kurylev, Matti Lassas, and Gunther Uhlmann, *Full-wave invisibility of active devices at all frequencies*, *Communications in Mathematical Physics* **275** (2007), no. 3.
- [44] Allan Greenleaf, Matti Lassas, and Gunther Uhlmann, *Anisotropic conductivities that cannot be detected by eit*, *Physiological Measurement* **24** (2003), 412–420.
- [45] ———, *On nonuniqueness for Calderón’s inverse problem*, *Mathematical research letters* **10** (2003), no. 5, 685, DOI <http://arxiv.org/abs/math/0302258v2>.
- [46] Eitan Grinspun, Yotam Gingold, Jason Reisman, and Denis Zorin, *Computing discrete shape operators on general meshes*, *Computer Graphics Forum* **25** (2006), no. 3, 547–556.
- [47] Z. Hashin, *Analysis of composite materials, a survey*, *Journal of Applied Mechanics* **50** (1983), 481–505.
- [48] Ulrich Hornung (ed.), *Homogenization and porous media*, Springer, New York, 1997.

-
- [49] D. Isaacson, J. L. Mueller, J. C. Newell, and S. Siltanen, *Imaging cardiac activity by the D -bar method for electrical impedance tomography*, *Physiological Measurement* **27** (2006), S43–S50.
- [50] Victor Isakov, *Uniqueness and stability in multi-dimensional inverse problems*, *Inverse Problems* **9** (1993), 579–621.
- [51] C Kenig, J Sjöstrand, and Gunther Uhlmann, *The Calderón problem with partial data*, *Annals of Mathematics* **165** (2007), no. 2, 567–591, DOI <http://arxiv.org/abs/math.AP/0405486v3>.
- [52] K. Knudsen, J. L. Mueller, and SI Siltanen, *Numerical solution method for the d -bar-equation in the plane*, *Journal of Computational Physics* **198** (2004), 500–517.
- [53] Roger Koenker and Ivan Mizera, *Penalized triograms: Total variation regularization for bivariate smoothing*, *Journal of the Royal Statistical Society. Series B. Statistical Methodology* **66** (2004), no. 1, 145–163.
- [54] R. V. Kohn and M. Vogelius, *Determining conductivity by boundary measurements*, *Communications on Pure and Applied Mathematics* **37** (1984), 113–123.
- [55] Ville Kolehmainen, Matti Lassas, and Petri Ola, *The inverse conductivity problem with an imperfectly known boundary*, *SIAM Journal on Applied Mathematics* **66** (2005), no. 2, 365–383, DOI <http://arxiv.org/pdf/math/0408237v1>.
- [56] ———, *The inverse conductivity problem with an imperfectly known boundary in three dimensions*, 2006, DOI <http://arxiv.org/pdf/math/0606640v1>.
- [57] François Labelle and Jonathan Richard Shewchuk, *Anisotropic voronoi diagrams and guaranteed-quality anisotropic mesh generation*, *Proceedings of the Nineteenth Annual Symposium on Computational Geometry (San Diego, CA)*, Association for Computing Machinery, 2003, <http://www.cs.berkeley.edu/~jrs/papers/aniso.pdf>, pp. 191–200.
- [58] W. R. B. Lionheart, *EIT reconstruction algorithms: pitfalls, challenges and recent developments*, *Physiological Measurement* **25** (2004), no. 1, 125–142.

-
- [59] Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H. Barr, *Discrete differential-geometry operators for triangulated 2-manifolds*, Visualization and mathematics III (Berlin), Springer-Verlag, 2003, pp. 35–57.
- [60] Jennifer L. Mueller, David Isaacson, and Jonathan C. Newell, *A reconstruction algorithm for electrical impedance tomography data collected on rectangular electrode arrays*, IEEE Transactions on Biomedical Engineering **46** (1999), no. 1, 1379–1386.
- [61] Adrian I. Nachman, *Global uniqueness for a two-dimensional inverse boundary value problem*, Annals of Mathematics **142** (1995), 71–96.
- [62] Jorge Nocedal and Stephen J Wright, *Numerical optimization*, Springer, New York, 2000.
- [63] Joseph O’Rourke, *Computational geometry in c*, 2nd ed., Cambridge University Press, New York, 1998.
- [64] Houman Owhadi and Lei Zhang, *Homogenization of the acoustic wave equation with a continuum of scales*, (2006), Submitted, DOI <http://arxiv.org/abs/math/0604380v2>.
- [65] ———, *Metric-based upscaling*, Communications on Pure and Applied Mathematics **60** (2006), no. 5, 675–723, DOI <http://arxiv.org/abs/math/0505223v5>.
- [66] ———, *Homogenization of parabolic equations with a continuum of space and time scales*, SIAM Journal on Numerical Analysis **46** (2007), no. 1, 1–36, DOI <http://arxiv.org/abs/math/0512504v1>.
- [67] Christopher C. Paige and Michael A. Saunders, *Lsqr: An algorithm for sparse linear equations and sparse least squares*, ACM Transactions on Mathematical Software **8** (1982), no. 1, 43–71.
- [68] D. Pedoe, *Geometry, a comprehensive course*, 2nd ed., Dover, New York, 1988.
- [69] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery, *Numerical recipes in c++*, 2nd ed., Cambridge University Press, New York, 2002.

-
- [70] Jason Reisman, Eitan Grinspun, and Denis Zorin, *A note on the triangle-centered quadratic interpolation discretization of the shape operator*, Tech. report, New York University, 2007.
- [71] Luca Rondi, *A remark on a paper by alessandrini and vessella*, *Advances in Applied Mathematics* **36** (2006), 67–69.
- [72] Szymon Rusinkiewicz, *Estimating curvatures and their derivatives on triangle meshes*, Proceedings, 2nd international symposium on 3D data processing, visualisation and transmission, 2004, pp. 486–493.
- [73] Peter Schröder, *What can we measure?*, 2005, Course notes, <http://www.multires.caltech.edu/pubs/GeoMeasure.pdf>.
- [74] D. Schurig, J. J. Mock, B. J. Justice, S. A. Cummer, J. B. Pendry, A. F. Starr, and D. R. Smith, *Metamaterial electromagnetic cloak at microwave frequencies*, *Science* **314** (2006), no. 5801, 977–980.
- [75] Jonathan Richard Shewchuk, *Triangle: A two-dimensional quality mesh generator and delaunay triangulator*, <http://www.cs.cmu.edu/~quake/triangle.html>.
- [76] ———, *An introduction to the conjugate gradient method without the agonizing pain*, 1994, <http://www.cs.cmu.edu/quake-papers/painless-conjugate-gradient.pdf>.
- [77] ———, *What is a good linear element? Interpolation, conditioning, and quality measures*, Proceedings of the 11th International Meshing Roundtable (Albuquerque, NM), Sandia National Laboratories, 2002, <http://www.cs.berkeley.edu/~jrs/papers/elemj.pdf>, pp. 115–126.
- [78] David R. Smith and John B. Pendry, *Homogenization of metamaterials by field averaging*, *Journal of the Optical Society of America B* **23** (2006), no. 3, http://www.ee.duke.edu/~drsmith/pubs_smith_group/smith_josab_2006.pdf.
- [79] Ari Stern, Yiyong Tong, Mathieu Desbrun, and Jerrold E. Marsden, *Computational electromagnetism with variational integrators and discrete differential forms*, (2008), Submitted, DOI <http://arxiv.org/pdf/0707.4470v2>.

-
- [80] Z. Sun and G. Uhlmann, *Anisotropic inverse problems in two dimensions*, *Inverse Problems* **16** (2003), 681–699.
- [81] John Sylvester, *An anisotropic inverse boundary value problem*, *Communications on Pure and Applied Mathematics* **43** (1990), 201–232.
- [82] John Sylvester and Gunther Uhlmann, *A global uniqueness theorem for an inverse boundary value problem*, *Annals of Mathematics* **125** (1987), 153–169.
- [83] Gunther Uhlmann, *Developments in inverse problems since Calderón’s foundational paper*, *Harmonic analysis and partial differential equations* (Chicago, IL), University of Chicago Press, 1999, pp. 295–345.
- [84] A. Wächter and L. T. Biegler, *Line search filter methods for nonlinear programming: Local convergence*, *SIAM Journal on Optimization* **16** (2005), no. 1, 32–48.
- [85] ———, *Line search filter methods for nonlinear programming: Motivation and global convergence*, *SIAM Journal on Optimization* **16** (2005), no. 1, 1–31.
- [86] ———, *On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming*, *Mathematical Programming* **106** (2006), no. 1, 25–57.
- [87] Benjamin W Wah and Tao Wang, *Simulated annealing with asymptotic convergence for nonlinear constrained global optimization*, *Proceedings of the 5th International Conference on Principles and Practice of Constraint Programming* (London), *Lecture Notes In Computer Science*, vol. 1713, Springer, 1999, pp. 461–475.
- [88] Lei Zhang, *Metric based upscaling for partial differential equations with a continuum of scales*, Ph.D. thesis, California Institute of Technology, 2007.
- [89] Sharon Zlochiver, M Michal Radai, Shimon Abboud, Moshe Rosenfeld, Xiu-Zhen Dong, Rui-Gang Liu, Fu-Sheng You, Hai-Yan Xiang, and Xue-Tao Shi, *Induced current electrical impedance tomography system: experimental results and numerical simulations*, *Physiological Measurement* **25** (2004), no. 1, 239–255.

Appendix A

Software credits

All of the software to produce the computational results for this work is written in C++ and compiled on a Linux machine using gcc version 3.4.6. Much assistance in the form of computational libraries is provided by the computing community. Although credit to the authors of software is provided in the bibliography, the following lists the packages used to avoid the reader having to sift through the bibliography for software citations. Wherever possible, we cite the version number used in our codes — a change in the version currently available to the reader should indicate packages that are continuing to be developed.

CGAL-3.3.1 <http://www.cgal.org> — Computational Geometry Algorithms Library, a collection of geometric algorithms and data structures, used in nearly all code for its mesh data structures and interpolation algorithms; also has built-in variable-precision numeric types, particularly useful for computing signed triangle areas [6]

Qt-3.3 <http://doc.trolltech.com/3.3/> — Qt, a GUI toolkit, used in nearly all code as a user interface to the OpenGL library [4]

GSL-1.10 <http://www.gnu.org/software/gsl/> — Gnu Scientific Library, a numerical library for C and C++ programmers [8]

Triangle-1.6 <http://www.cs.cmu.edu/~quake/triangle.html> — Triangle, a 2D quality mesh generator, used to generate triangle meshes [75]

IpOpt-3.2.4 <https://projects.coin-or.org/Ipopt/> — Interior point optimisation library for C++, used to solve the inverse problem optimisation program [2]

Glpk-4.8 <http://www.gnu.org/software/glpk/> — Gnu Linear Programming Kit, used to solve the inverse problem when posed as an L1-minimisation [7]

-
- lsqr** <http://www.stanford.edu/group/SOL/software/lsqr.html> — A conjugate gradient type method for solving sparse least-squares problems, used to compute (s_i) from (q_{ij}) [9, 67]
- ConfigFile** <http://www-personal.umich.edu/~wagnerr/ConfigFile.html> — Configuration file reader for C++, used everywhere for setting solver parameters at runtime [1]
- OpenGL-1.2** <http://www.opengl.org/> — Open Graphics Library, an industry standard for high-performance 3D rendering, used here with the OpenGL Utility Library (GLU) to render the surface $s(x)$; particularly useful for checking by eye that $s(x)$ is convex [3]
- cblas** <http://www.netlib.org/blas/> — Basic Linear Algebra Subroutines in C, low-level libraries used by several other packages, as well as to optimise my own conjugate gradient solver [5]
- Numerical Recipes** <http://www.nr.com/> — Numerical Recipes in C++, used for its SVD code [69]
- CGeom** <http://maven.smith.edu/~orourke/books/ftp.html> — Computational Geometry in C, code from a book by Joseph O'Rourke; though the book offers a great deal more, only the $\mathcal{O}(n^2)$ convex hull code is used here [63]
- CG** Although by itself not software, all of the linear solves for this project are done using the diagonally preconditioned conjugate gradient method. This is not the fastest method, but for positive-definite systems, it is one of the simplest, particularly given the help found in Shewchuk's explanation [76]

Appendix B

Notation

$\lfloor x \rfloor$	Largest integer not larger than x
∇F	When $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$, element $(\nabla F)_{ij} = \partial_{x_i} F_j$
$\lambda_{\min}, \lambda_{\max}$	Extremal eigenvalues of a matrix
Ω	A domain in \mathbb{R}^d , closed and bounded unless otherwise indicated
$\partial\Omega$	Boundary of domain Ω
$\partial_{x_i x_j} f$	The second partial derivative of function f with respect to coordinates x_i, x_j
A^T	Transpose of matrix A
A_{ij}	The element in row i , column j of matrix A
d	The dimension; the domain $\Omega \subset \mathbb{R}^d$
h	The scale of a triangulation of Ω
H_s	Hessian of scalar function $s(x)$
I	The imaginary number $\sqrt{-1}$
\mathbf{j}	Flux due to a potential
$\delta v, \delta^2 v$	First, second variations of v
φ_i	Piecewise linear hat function at node i
ψ_i	The pull-back of the piecewise linear hat function, $\varphi_i \circ F$
σ	A conductivity, typically having fine-scales
F	Harmonic coordinate map

F_*u	Push-forward of function u by coordinate change F ; u may be scalar- or tensor-valued
Q	A conductivity, often the homogenisation from metric-based up-scaling
q_{ij}	Stiffness matrix value on edge ij . When plural, the entire stiffness matrix is implied
$s(x), s(x, y)$	Scalar function parameterising a conductivity
s_i	Value of the parameterisation $s(x)$ interpolated at vertex i . When plural, the entire set of interpolants is implied
u	A scalar potential field
$(\cdot, \cdot)_\Omega$	L^2 inner product over the domain
$(\cdot, \cdot)_{\partial\Omega}$	L^2 inner product over the domain boundary
$\ \cdot\ $	Unspecified norm
i, j, k	Mesh vertices
$i \sim j$	Vertices i, j sharing an edge of a triangulation
ij	Oriented mesh edge from vertex i to vertex j
ijk	Oriented mesh triangle with vertices i, j, k
$x_i, (x_i, y_i)$	Location of mesh vertex i
$C^{1,\alpha}$	Space of functions with C^α first derivatives
C^α	Space of Hölder continuous functions with parameter α
$H^1(\Omega)$	First Sobolev space, $W^{1,2}(\Omega)$. The domain Ω is often implied
$H_0^1(\Omega)$	Set of functions in $H^1(\Omega)$ taking the value zero on $\partial\Omega$
V^h	Discrete function space approximating solutions at scale h
$W^{m,p}(\Omega)$	Sobolev space of functions having support in domain Ω , with m th-order (and lower) weak derivatives in L^p . The domain is often implied

\mathcal{P}_0	Space of piecewise constant functions
\mathcal{P}_0^h	Space of functions piecewise constant over scale h
\mathcal{P}_1	Space of piecewise linear polynomials
\mathcal{P}_1^h	Subspace of \mathcal{P}_1 piecewise linear over scale h
$\mathcal{B}(u, v)$	The bilinear operator associated to the elliptic problem
$\mathcal{B}^h(u, v)$	The bilinear operator associated to the elliptic problem, approximated at scale h
Λ_s	Neumann-to-Dirichlet map parameterised by $s(x)$
Λ_σ	Neumann-to-Dirichlet map for conductivity σ
E_σ	Equivalence class of conductivities consistent with the Neumann-to-Dirichlet map determined by σ
$f^{(i)}, g^{(i)}$	Neumann and Dirichlet boundary data, respectively. Index i identifies $f^{(i)}$ with $g^{(i)}$ by $\Lambda_\sigma(f^{(i)}) = g^{(i)}$