Contents

Preface vii

1 Non-linear statistical inverse problems 1
- 1.1 Model examples 2
  - 1.1.1 Boundary measurements in tomography 2
  - 1.1.2 Parameter identification for elliptic PDEs 3
  - 1.1.3 Data assimilation 6
- 1.2 Bayesian regression 7
  - 1.2.1 The forward map $\mathcal{G}$ 7
  - 1.2.2 A random design regression model with normal errors 7
  - 1.2.3 The Bayesian prior and posterior distribution 9
  - 1.2.4 Posterior computation by MCMC 10
- 1.3 The frequentist perspective 14
  - 1.3.1 Information distances for random design regression 14
  - 1.3.2 A first posterior contraction theorem 17
- 1.4 Notes 20
  - 1.4.1 Exercises 20
  - 1.4.2 Remarks and comments 20

2 Global stability and posterior consistency 23
- 2.1 Analytical hypotheses and PDEs 24
  - 2.1.1 Forward regularity conditions for $\mathcal{G}$ 24
  - 2.1.2 Injectivity and stability estimates 27
- 2.2 Regularisation with Gaussian process priors 31
- 2.3 Convergence of posterior measure and mean 35
- 2.4 Notes 38
  - 2.4.1 Exercises 38
  - 2.4.2 Remarks and comments 39

3 Information operators and curvature 41
- 3.1 Information geometry 41
  - 3.1.1 The LAN expansion 42
  - 3.1.2 Cramer–Rao bounds and inverse information 45
- 3.2 Gradient stability and concentration 49
  - 3.2.1 Convexity of $-\ell_N$ and the gradient of $\mathcal{G}$ 49
  - 3.2.2 A concentration result for the empirical Hessian 51

3.3 Information operators for elliptic PDEs 58
3.3.1 Schrödinger equation 58
3.3.2 Diffusion equation 59
3.3.3 Injectivity and local identifiability 61
3.4 Notes 63
3.4.1 Exercises 63
3.4.2 Remarks and comments 64

4 Bernstein–von Mises theorems 67
4.1 Gaussian asymptotics for cylindrical laws 68
4.1.1 Asymptotic normality of linear functionals of the posterior . . 68
4.1.2 Asymptotic distribution of the posterior mean 77
4.1.3 Applications to uncertainty quantification 80
4.2 Solving information equations in PDE models 81
4.2.1 A Bernstein–von Mises theorem for the Schrödinger equation 82
4.2.2 Impossibility of the BvM-phenomenon for Darcy's problem . 84
4.3 Notes 88
4.3.1 Exercises 88
4.3.2 Remarks and comments 89

5 Posteriors are probably log-concave 91
5.1 Wasserstein approximation of the posterior 92
5.1.1 Construction of a log-concave surrogate posterior 94
5.1.2 The log-concave approximation theorem 95
5.2 Computational complexity of MCMC in high dimensions 101
5.2.1 Gradient methods for approximately log-concave posteriors . . 101
5.2.2 On failure of 'cold start' MCMC in high dimensions 106
5.3 Application to PDE models 113
5.3.1 Darcy's problem on the eigen-spaces of the Laplacian 114
5.3.2 Polynomial time computation of the posterior mean 119
5.4 Notes 121
5.4.1 Exercises 121
5.4.2 Remarks and comments 121

A Analytical background 125
A.1 Sobolev and related spaces 125
A.2 Elliptic second-order differential operators 126
A.3 Orthonormal discretisation of L^2 and metric entropy 128
A.4 Feynman–Kac formulæ 132
A.5 Elliptic regularity estimates 133

B Further auxiliary results 139
B.1 Results from Gaussian process theory 139
B.2 A concentration inequality for empirical processes 143
B.3 Mixing time bounds for Langevin diffusions 145
B.4 A characterisation of vanishing efficient information 147

References 151

Index 159