Uncertainty Quantification in Bayesian Inversion

Jonas Latz
Technische Universität München
Mathematics Faculty
Chair of Numerical Mathematics (M2)
November 17, 2016

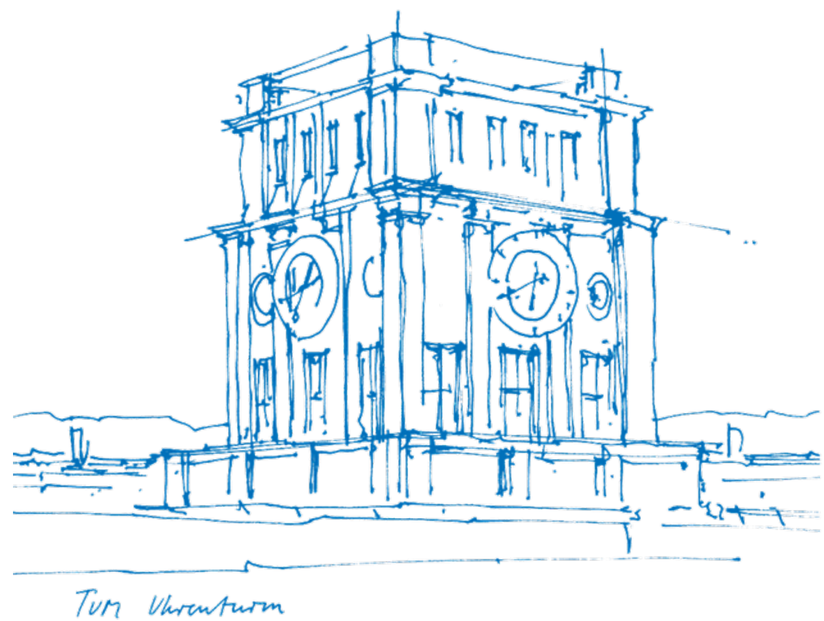

Outline

- Brief introduction to Bayesian Statistics and Bayesian Inverse Problems
- Andrew M. Stuart - Uncertainty Quantification in Bayesian Inversion.

Bayesian Statistics

Setting:

- Observation: data point $y \in \mathbb{R}^n$,
- Parameterized distribution of y: $\mathbb{P}(y \in \cdot | u)$, given by a Lebesgue-density

$$\frac{\mathrm{d}\mathbb{P}(y \in \cdot | u)}{\mathrm{d}\lambda^n}(x) =: L(x|u),$$

- Parameter $u \in X$.

Task: Identify the parameter u based on the observed data point y.

Bayesian Statistics

Setting:

- Observation: data point $y \in \mathbb{R}^n$,
- Parameterized distribution of y: $\mathbb{P}(y \in \cdot | u)$, given by a Lebesgue-density

$$\frac{\mathrm{d}\mathbb{P}(y \in \cdot | u)}{\mathrm{d}\lambda^n}(x) =: L(x|u),$$

- Parameter $u \in X$.

Task: Identify the parameter u based on the observed data point y.

Classical Inference: Maximum likelihood estimate: $\widehat{u} :\in \mathrm{argmax}_{u \in X}[\log L(y|u)]$

Bayesian Statistics

Setting:

- Observation: data point $y \in \mathbb{R}^n$,
- Parameterized distribution of y: $\mathbb{P}(y \in \cdot | u)$, given by a Lebesgue-density

$$\frac{\mathrm{d}\mathbb{P}(y \in \cdot | u)}{\mathrm{d}\lambda^n}(x) =: L(x|u),$$

- Parameter $u \in X$.

Task: Identify the parameter u based on the observed data point y.

Classical Inference: Maximum likelihood estimate: $\widehat{u} :\in \operatorname{argmax}_{u \in X}[\log L(y|u)]$

Bayesian Inference: The estimate $\widehat{u}$ is now modelled as a random variable, reflecting knowledge about u.

Bayesian Statistics

Setting:

- Observation: data point $y \in \mathbb{R}^n$,
- Parameterized distribution of y: $\mathbb{P}(y \in \cdot | u)$, given by a Lebesgue-density

$$\frac{\mathrm{d}\mathbb{P}(y \in \cdot | u)}{\mathrm{d}\lambda^n}(x) =: L(x|u),$$

- Parameter $u \in X$.

Task: Identify the parameter u based on the observed data point y.

Classical Inference: Maximum likelihood estimate: $\widehat{u} :\in \operatorname{argmax}_{u \in X} [\log L(y|u)]$

Bayesian Inference: The estimate $\widehat{u}$ is now modelled as a random variable, reflecting knowledge about u.

Prior: $\mu_{\text{prior}} := \mathbb{P}(u \in \cdot)$ (knowledge/distribution before seeing the data)

Bayesian Statistics

Setting:

- Observation: data point $y \in \mathbb{R}^n$,
- Parameterized distribution of y: $\mathbb{P}(y \in \cdot | u)$, given by a Lebesgue-density

$$\frac{\mathrm{d}\mathbb{P}(y \in \cdot | u)}{\mathrm{d}\lambda^n}(x) =: L(x|u),$$

- Parameter $u \in X$.

Task: Identify the parameter u based on the observed data point y.

Classical Inference: Maximum likelihood estimate: $\widehat{u} :\in \mathrm{argmax}_{u \in X}[\log L(y|u)]$

Bayesian Inference: The estimate $\widehat{u}$ is now modelled as a random variable, reflecting knowledge about u.

Prior: $\mu_{\text{prior}} := \mathbb{P}(u \in \cdot)$ (knowledge/distribution before seeing the data)

Posterior: $\mu_{\text{post}} := \mathbb{P}(u \in \cdot | y)$ (knowledge/distribution after seeing the data)

Bayesian Statistics: How to derive the posterior?

Under further assumptions, the posterior can be derived using Bayes' formula:

$$\frac{\mathrm{d}\mu_{\text{post}}}{\mathrm{d}\mu_{\text{prior}}}(u) = \frac{L(y|u)}{\int L(y|u)\mathrm{d}\mu_{\text{prior}}(u)} =: \frac{L(y|u)}{Z(y)}.$$

Assume $u \in X = \mathbb{R}^k$ and μ_{prior} has a pdf f_{prior}, then:

$$f_{\text{post}}(u) = \frac{L(y|u)}{\int L(y|u)\mathrm{d}\mu_{\text{prior}}} f_{\text{prior}}(u) = \frac{L(y|u)}{Z(y)} f_{\text{prior}}(u)$$

.

Bayesian Statistics: How to derive the posterior?

Analytical: Only possible given specific prior + likelihood pairs. (*Conjugate priors*)

Bayesian Statistics: How to derive the posterior?

Analytical: Only possible given specific prior + likelihood pairs. (*Conjugate priors*)

Computational: Produce (weighted) samples of the posterior to approximate it empirically.

- Importance Sampling (requires to estimate $Z(y)$)
- Markov Chain Monte Carlo (MCMC; does not yield independent samples of μ_{post})
- Sequential Monte Carlo (SMC; efficient, even if posterior is multimodal or concentrated; requires to estimate $Z(y)$)

Inverse Problem

- $\mathscr{G} : X \to Y$ is the forward response operator,
- $\eta \sim \mathrm{N}(0, \Gamma)$ is noise,
- $u^{\text{truth}} \in X$ is the true model parameter,
- $y \in Y$ is (noisy) observed data of the model, i.e. given by $y := \mathscr{G}(u^{\text{truth}}) + \eta$.

Identify the parameter u^{truth}, based on the data y.

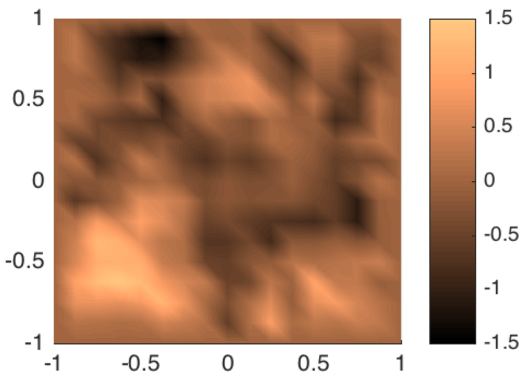

(a) True Parameter

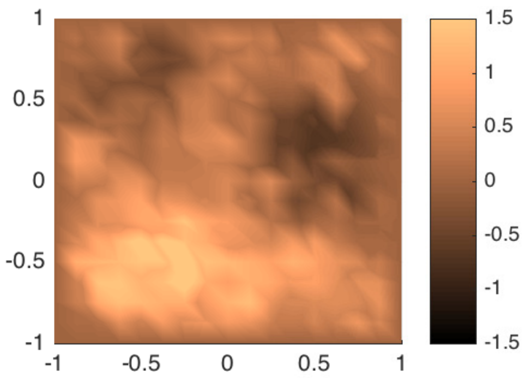

(b) Estimation

Figure: log-Permeability of an Oil Reservoir

Bayesian Inverse Problem

Let $u \sim \mu_0 := \mu_{\text{prior}}$ and (u, η) independent. Then,

$$\mathscr{G}(u) + \eta = y \Leftrightarrow \eta = y - \mathscr{G}(u),$$

and therefore,

$$y - \mathscr{G}(u) \sim \mathrm{N}(0, \Gamma).$$

The likelihood is then given by

$$L(y|u) := \phi_{0,\Gamma}(y - \mathscr{G}(u)) := \exp(-\tfrac{1}{2}\|\Gamma^{-\frac{1}{2}}(y - \mathscr{G}(u))\|_2^2)$$

The posterior $\mu^y := \mu_{\text{post}}$ is then given by Bayes' formula:

$$\mu^y = \frac{1}{Z(y)} \exp(-\tfrac{1}{2}\|\Gamma^{-\frac{1}{2}}(y - \mathscr{G}(\cdot))\|_2^2)\mu_0.$$

References

[1] Dashti, M. and Stuart, A.M. (2015) - *The Bayesian Approach to Inverse Problems*, Handbook of Uncertainty Quantification, Springer.

[2] Iglesias, M.A., Law, K.J.H. and Stuart, A.M. (2013) - *Ensemble Kalman methods for inverse problems*, IOPscience

[3] Latz, J. (2016) - *Bayes Linear Methods for Inverse Problems*, Master's Thesis University of Warwick.

[4] K.J.H. Law, Stuart, A.M. and Zygalakis, K.C. - *Data Assimilation: A Mathematical Introduction*, Springer.

[5] Schillings, C. and Stuart, A.M. (2016) - *Analysis of the ensemble Kalman filter for inverse problems*, preprint.

[6] Stuart, A.M. (2010) - *Inverse problems: a Bayesian perspective*, Acta Numerica 19.

[7] Sullivan, T.J. (2015) - *Introduction to Uncertainty Quantification*, Springer.