An adaptive algorithm based on spectral likelihood expansion for efficient Bayesian calibration

Author(s):
Wagner, Paul-Remo; Lataniotis, Christos; Marelli, Stefano; Sudret, Bruno

Publication date:
2019-06-24

Permanent link:
https://doi.org/10.3929/ethz-b-000350795

An adaptive algorithm based on spectral likelihood expansion for efficient Bayesian calibration

P.-R. Wagner, C. Lataniotis, S. Marelli, B. Sudret

Chair of Risk, Safety and Uncertainty Quantification - ETH Zürich

UNCECOMP, Crete, Greece, 24.06.2019

Motivation

Typical problem:

- Consider a computational model $\mathcal{M}$ that is to be calibrated using data $\mathcal{Y}$
- The framework of Bayesian model calibration allows the computation of the distribution of the input parameters $\boldsymbol{X}$ conditioned on the data (posterior)

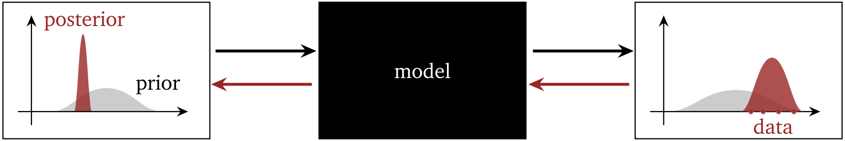

In practice the computation of posterior quantities is challenging

Outline

1. Bayesian model calibration
2. Stochastic spectral embedding
3. Application: Heat transfer problem
4. Conclusions

Framework

Given parameters $\boldsymbol{X} \sim \pi(\boldsymbol{x})$ and measurements $\mathcal{Y}$, the Bayesian inverse problem reads:

$$\pi(\boldsymbol{x}|\mathcal{Y}) = \frac{\mathcal{L}(\boldsymbol{x};\mathcal{Y})\pi(\boldsymbol{x})}{Z} \quad \text{where} \quad Z = \int_{\mathcal{D}_{\boldsymbol{X}}} \mathcal{L}(\boldsymbol{x};\mathcal{Y})\pi(\boldsymbol{x})\mathrm{d}\boldsymbol{x}$$

with:

- $\mathcal{L} : \mathcal{D}_{\boldsymbol{X}} \to \mathbb{R}^+$: likelihood function (measure of how well the model fits the data)
- $\pi(\boldsymbol{x}|\mathcal{Y})$: posterior density function

Quantities of Interest (QoI)

Often one is interested in expectations of QoI following calibration: $h(\boldsymbol{x}) : \mathcal{D}_{\boldsymbol{X}} \to \mathbb{R}$:

$$\mathbb{E}\left[h(\boldsymbol{X})|\mathcal{Y}\right] = \int_{\mathcal{D}_{\boldsymbol{X}}} h(\boldsymbol{x})\pi(\boldsymbol{x}|\mathcal{Y})\mathrm{d}\boldsymbol{x}$$

e.g. posterior moments.

Outline

1. Bayesian model calibration
2. Stochastic spectral embedding
3. Application: Heat transfer problem
4. Conclusions

Spectral likelihood expansions (SLE)

Principle

Nagel & Sudret, J. Comput. Phys. (2016)

Assuming $\boldsymbol{X} \sim \pi(\boldsymbol{x}) = \prod_{i=1}^{M} \pi(x_i)$ (independent), SLE approximates $\mathcal{L}(\boldsymbol{X})$ (for fixed observations $\mathcal{Y}$) with a finite sum of orthonormal polynomials:

$$\mathcal{L}(\boldsymbol{X}) = \sum_{\boldsymbol{\alpha} \in \mathbb{N}^M} a_{\boldsymbol{\alpha}} \Psi_{\boldsymbol{\alpha}}(\boldsymbol{X}) \approx \sum_{\boldsymbol{\alpha} \in \mathcal{A}} a_{\boldsymbol{\alpha}} \Psi_{\boldsymbol{\alpha}}(\boldsymbol{X})$$

where $\Psi_{\boldsymbol{\alpha}}(\boldsymbol{x}) \stackrel{\text{def}}{=} \prod_{i=1}^{M} \psi_{\alpha_i}^{(i)}(x_i)$ and $\int_{\mathcal{D}_{X_i}} \psi_j^{(i)}(x_i) \psi_k^{(i)}(x_i) \pi(x_i) \mathrm{d}x_i = \delta_{jk}$

Practical implementation

- Sparse polynomial chaos expansions of the likelihood function
- Using least-angle regression

Blatman & Sudret, J. Comput. Phys. (2011)

All calculations are carried out with UQLab

Spectral likelihood expansions (SLE)

Analytical posterior

After the computation of the coefficients $a_{\boldsymbol{\alpha}}$, the full posterior distribution or QoI can be computed analytically:

Post-processing $a_{\boldsymbol{\alpha}}$

Nagel & Sudret, J. Comput. Phys. (2016)

$$Z = \mathbb{E}\left[\mathcal{L}(\boldsymbol{X})\right] \approx a_{\boldsymbol{0}}$$

$$\pi(\boldsymbol{x}|\mathcal{Y}) \approx \pi(\boldsymbol{x}) \sum_{\boldsymbol{\alpha}\in\mathcal{A}} \frac{a_{\boldsymbol{\alpha}}}{a_{\boldsymbol{0}}} \Psi_{\boldsymbol{\alpha}}(\boldsymbol{x})$$

$$\mathbb{E}\left[h(\boldsymbol{X})|\mathcal{Y}\right] \approx \frac{1}{a_{\boldsymbol{0}}} \sum_{\boldsymbol{\alpha}\in\mathcal{A}} a_{\boldsymbol{\alpha}} b_{\boldsymbol{\alpha}} \quad \text{after} \quad h(\boldsymbol{X}) \approx \sum_{\boldsymbol{\alpha}\in\mathcal{A}} b_{\boldsymbol{\alpha}} \Psi_{\boldsymbol{\alpha}}(\boldsymbol{X})$$

Features

- No MCMC simulation, analytical posterior moments, covariance, etc.
- However, likelihood functions typically do not have a sparse representation

Stochastic spectral embedding (SSE)

Proposed approach

Sequentially approximate likelihood with sum of low-degree PCEs $f^k_{\mathrm{PC}}(\boldsymbol{X}) = \sum_{\boldsymbol{\alpha}\in\mathcal{A}^k} a^k_{\boldsymbol{\alpha}} \Psi^k_{\boldsymbol{\alpha}}(\boldsymbol{X})$ on multiple levels and subdomains $\mathcal{D}^k_{\boldsymbol{X}}$ using the residual of the previous PCEs

$$\mathcal{L}(\boldsymbol{X}) \approx \mathcal{L}^{\mathrm{SSE}}(\boldsymbol{X}) \stackrel{\mathrm{def}}{=} \sum_{k\in\mathcal{K}} f^k_{\mathrm{PC}}(\boldsymbol{X}) \mathbf{1}_{\mathcal{D}^k_{\boldsymbol{X}}}(\boldsymbol{X})$$

How to sequentially build embedded PCEs?

Construction of SSE - 1D example

Algorithm

Initialize, repeat:

1. Select $\mathcal{D}_{\boldsymbol{X}}^{k}$ with maximum error (among terminal domains $\mathcal{T}$)
2. Optional: Enrich experimental design
3. Split $\mathcal{D}_{\boldsymbol{X}}^{k}$
4. Construct PCEs using the residual of $\mathcal{L}$

Construction of SSE - higher dimensions

Algorithm

Initialize, repeat:

1. Select $\mathcal{D}_{\boldsymbol{X}}^{k}$ with maximum error (among terminal domains $\mathcal{T}$)
2. Optional: Enrich experimental design
3. Split $\mathcal{D}_{\boldsymbol{X}}^{k}$ into $2M$ candidate subdomains
4. Keep subdomain pair with minimum error
5. Construct PCEs using the residual of $\mathcal{L}$

Post-processing of SSE

Analytical posterior

After the construction of an SSE approximation of $\mathcal{L}$, the full posterior distribution or QoI can be computed analytically:

Post-processing $a^k_{\boldsymbol{\alpha}}$

$$Z = \mathbb{E}\left[\mathcal{L}(\boldsymbol{X})\right] \approx \sum_{k \in \mathcal{K}} c^k a^k_{\boldsymbol{0}}, \quad \text{where} \quad c^k = \int_{\mathcal{D}^k_{\boldsymbol{X}}} \pi(\boldsymbol{x}) \mathrm{d}\boldsymbol{x}$$

$$\pi(\boldsymbol{x}|\mathcal{Y}) \approx \frac{\pi(\boldsymbol{x})}{Z} \sum_{k \in \mathcal{K}} f^k_{\mathrm{PC}}(\boldsymbol{x}) \mathbf{1}_{\mathcal{D}^k_{\boldsymbol{X}}}(\boldsymbol{x})$$

$$\mathbb{E}\left[h(\boldsymbol{X})|\mathcal{Y}\right] \approx \frac{1}{Z} \sum_{k \in \mathcal{K}} c^k \cdot \sum_{\boldsymbol{\alpha} \in \mathcal{A}^k} a^k_{\boldsymbol{\alpha}} b^k_{\boldsymbol{\alpha}} \quad \text{after} \quad h(\boldsymbol{x}) \approx \sum_{\boldsymbol{\alpha} \in \mathcal{A}^k} b^k_{\boldsymbol{\alpha}} \Psi^k_{\boldsymbol{\alpha}}(\boldsymbol{x})$$

Outline

Heat transfer problem

- Temperature measurements at 20 locations $\mathcal{Y} = \{T_1, \dots, T_N\}$
- Computational forward model solves the steady-state heat equation (FE-method):

$$\nabla(\kappa \nabla T) = 0$$

- Likelihood with $\boldsymbol{\kappa} \stackrel{\text{def}}{=} (\kappa_1, \dots, \kappa_6)$:

$$\mathcal{L}(\boldsymbol{\kappa}; \mathcal{Y}) = \prod_{i=1}^{N} \mathcal{N}(T_i | \mathcal{M}(\boldsymbol{\kappa}), \sigma^2)$$

- Prior distributions:

$$\pi(\boldsymbol{\kappa}) = \prod_{i=1}^{M} \mathcal{LN}(\mu = 30, \sigma = 6 \text{ W/mK})$$

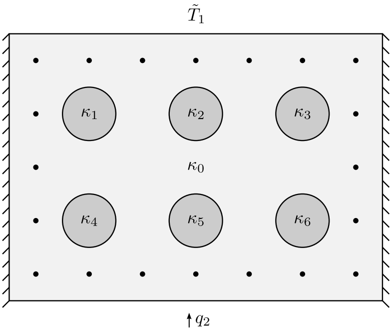

Heat transfer problem

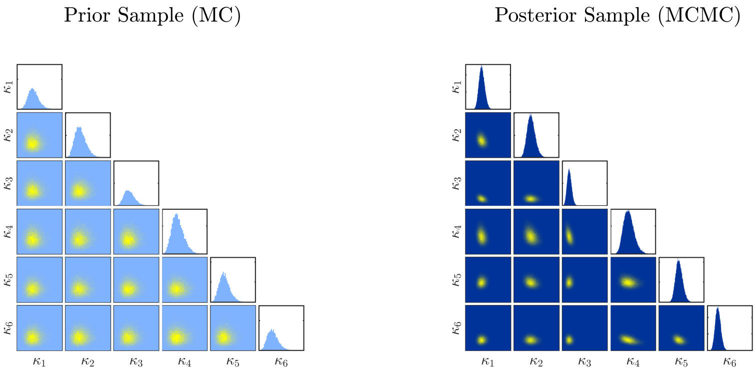

The reference solution is obtained by MCMC (AIES, 10^5 $\mathcal{L}$ evaluations)...

Goodman & Weare, Comm. Appl. Math. Comp. Sci. (2010)

Heat transfer problem

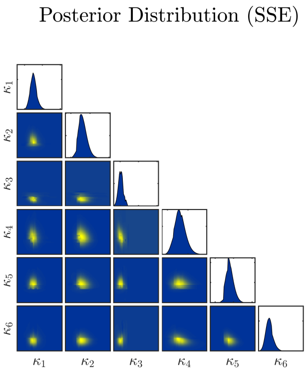

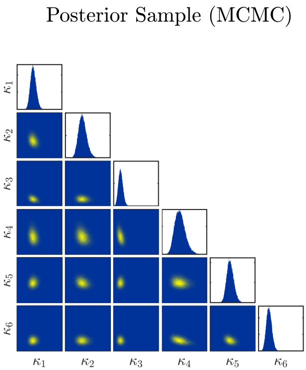

. . . and compared to the SSE solution (10^4 $\mathcal{L}$ evaluations)

Posterior moments and correlations

SSE (10^4 $\mathcal{L}$ evaluations) vs. MCMC (10^5 $\mathcal{L}$ evaluations)

Moments

	κ_1	κ_2	κ_3	κ_4	κ_5	κ_6
$\mathbb{E}[\cdot\|\mathcal{Y}]$ (W/mK)	30.1(29.8)	32.5(32.3)	20.7(20.7)	32.4(32.4)	36(36.4)	26.4(26.2)
$\mathrm{Var}[\cdot\|\mathcal{Y}]$ ($\mathrm{W}^2/\mathrm{mK}^2$)	12.5(10.5)	17.4(17.9)	6.51(6.12)	27.9(26.8)	13.6(14.7)	12.8(9.02)

Correlations

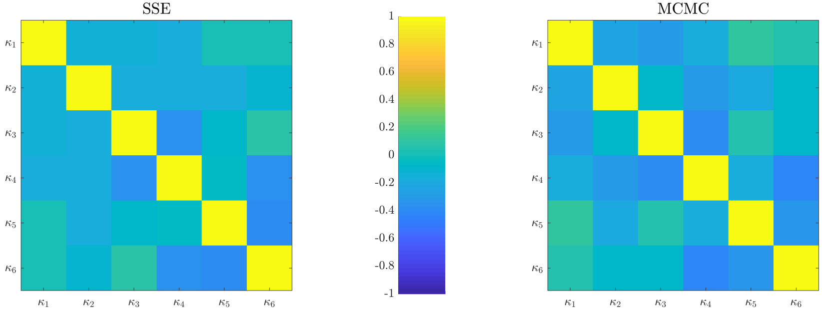

Convergence of SSE

Track convergence of Z estimate

Estimate $Z \overset{\text{def}}{=} \mathbb{E}\left[\mathcal{L}(\boldsymbol{X})\right]$ at step t by:

$$Z \approx \tilde{Z}_t \overset{\text{def}}{=} \mathbb{E}\left[\mathcal{L}^{\mathrm{SSE},t}(\boldsymbol{X})\right] = \sum_{k\in\mathcal{K}_t} c^k a_{\mathbf{0}}^k$$

- $\mathcal{K}_t$: Domains at step t
- $c^k \overset{\text{def}}{=} \int_{\mathcal{D}_{\boldsymbol{X}}^k} \pi(\boldsymbol{x})\mathrm{d}\boldsymbol{x}$: prior mass in the $k-$th domain
- $a_{\mathbf{0}}^k$: coefficient of $\Psi_{\mathbf{0}}^k$

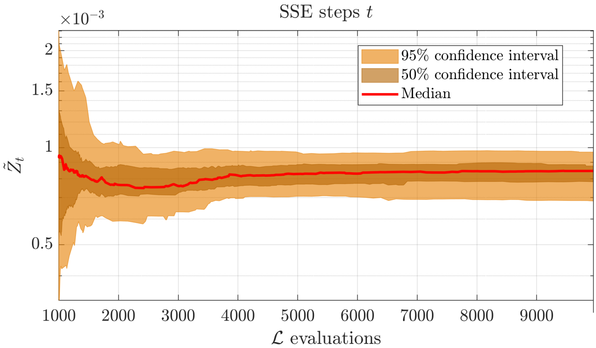

Convergence of SSE

Track error of posterior expectation

$$\varepsilon_{\kappa_i,t} \overset{\text{def}}{=} \frac{\left(\mu_{\kappa_i} - \tilde{\mu}_{\kappa_i,t}\right)^2}{\mathrm{Var}\left[\kappa_i\right]}$$

- μ_{κ_i}: MCMC reference
- $\tilde{\mu}_{\kappa_i,t}$: estimate at step t

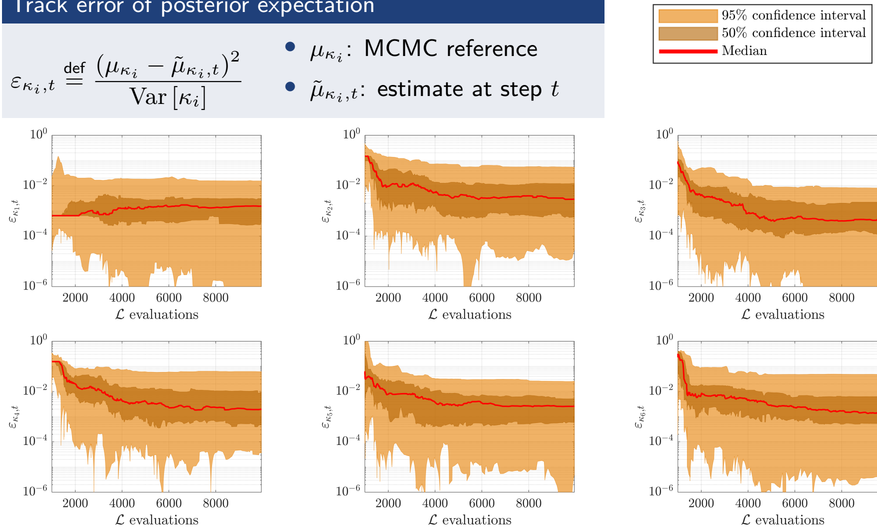

Outline

Conclusion

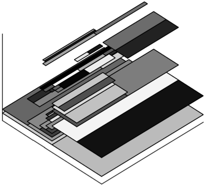

- Bayesian inversion is a powerful tool for model calibration
- Spectral likelihood expansion aims at avoiding any MCMC sampling, by expanding the likelihood function onto a polynomial chaos
- To avoid a single, dense, high-degree expansion, stochastic spectral embedding is built as a sum of local polynomials, with adaptive domain refinement and error estimation.
- First investigations show a good performance, at a cost one order of magnitude smaller than MCMC

Questions ?

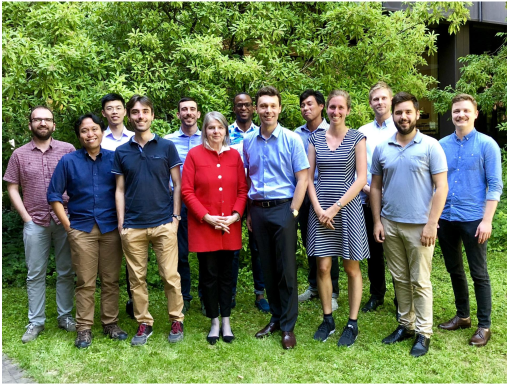

Chair of Risk, Safety & Uncertainty Quantification

www.rsuq.ethz.ch

The Uncertainty Quantification Software

www.uqlab.com

UQLab

Thank you very much for your attention !