MAX PLANCK INSTITUTE
FOR DYNAMICS OF COMPLEX
TECHNICAL SYSTEMS
MAGDEBURG

CSC
COMPUTATIONAL METHODS IN
SYSTEMS AND CONTROL THEORY

Low-Rank Methods for Linear Bayesian Inverse Problems

Peter Benner

GAMM Jahrestagung 2018
TU München, 19–23 March 2018

Joint work with Yue Qiu (MPI Magdeburg)
and Martin Stoll (TU Chemnitz)

Inverse Problems

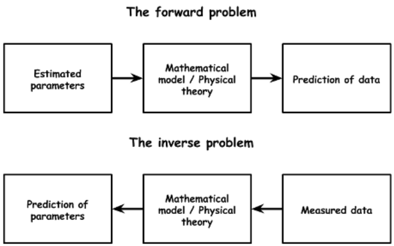

Many inverse problems governed by PDEs, we seek to infer unknown or uncertain varying parameter fields

- initial conditions
- coefficients
- boundary conditions ······

from limited and noisy observation.

Bayesian inference provides a systematic framework for incorporating uncertainties in

- observations
- forward models
- prior knowledge······

to quantify uncertainties.

Outline

Bayesian Framework

In Bayesian framework, view all parameters as random variables and define the parameter-to-observable map $g : \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}^m$

$$Y = g(U, E).$$

Here, Y, U, E are all random variables.

- $u \in \mathbb{R}^n$: model parameters to be recovered;
- $e \in \mathbb{R}^k$: error vector;
- $y \in \mathbb{R}^m$: observable variables;

are realizations of U, E, and Y, respectively.
Choose probability density functions (PDFs)

- $\pi_{\text{noise}} : \mathbb{R}^k \to \mathbb{R}$, modeling error and observation noise, etc.,
- $\pi_{\text{prior}} : \mathbb{R}^n \to \mathbb{R}$, prior information on parameters u,
- $\pi(y|u)$: describes relationship between observables y and parameters u.

Bayesian Inference

To get posterior PDF $\pi_{\text{post}} : \mathbb{R}^m \to \mathbb{R}$, we apply Bayes' theorem,

$$\pi_{\text{post}} := \pi(u|y_{\text{obs}}) = \frac{\pi_{\text{prior}}(u)\pi(y_{\text{obs}}|u)}{\pi(y_{\text{obs}})} \propto \pi_{\text{prior}}(u)\pi(y_{\text{obs}}|u).$$

Here we assume additive noise, so

$$Y = f(U) + E,$$

where $f : \mathbb{R}^n \to \mathbb{R}^m$ and e is additive noise that reflects both the modeling error and observation error.
Assume U and E are statistically independent, therefore,

$$\pi_{\text{post}} \propto \pi_{\text{prior}}(u)\pi_{\text{noise}}(y_{\text{obs}} - f(u)).$$

Linear Gaussian Case

Assuming that both probability density functions for u and e are Gaussian, we can rewrite the PDFs in the form

$$\pi_{\text{prior}}(u) \propto \exp\left(-\frac{1}{2}\left(u-\bar{u}_{\text{prior}}\right)^T \Gamma_{\text{prior}}^{-1}\left(u-\bar{u}_{\text{prior}}\right)\right),$$
$$\pi_{\text{noise}}(e) \propto \exp\left(-\frac{1}{2}\left(e-\bar{e}\right)^T \Gamma_{\text{noise}}^{-1}\left(e-\bar{e}\right)\right)$$

Bayes' theorem further gives

$$\begin{aligned}\pi_{\text{post}} &\propto \exp\left(-\frac{1}{2}\left(u-\bar{u}_{\text{prior}}\right)^T \Gamma_{\text{prior}}^{-1}\left(u-\bar{u}_{\text{prior}}\right)-\frac{1}{2}\left(e-\bar{e}\right)^T \Gamma_{\text{noise}}^{-1}\left(e-\bar{e}\right)\right)\\ &= \exp\left(-\frac{1}{2}\|u-\bar{u}_{\text{prior}}\|^2_{\Gamma_{\text{prior}}^{-1}}-\frac{1}{2}\|e-\bar{e}\|^2_{\Gamma_{\text{noise}}^{-1}}\right).\end{aligned}$$

Computations of Γ_{post}

For linear parameter-to-observable map, $y = f(u) = Au$, we have

$$\pi_{\text{post}} \propto \exp\left(-\frac{1}{2}\|u - \bar{u}_{\text{prior}}\|^2_{\Gamma^{-1}_{\text{prior}}} - \frac{1}{2}\|y_{\text{obs}} - Au - \bar{e}\|^2_{\Gamma^{-1}_{\text{noise}}}\right).$$

Then we get

$$\bar{u}_{\text{post}} = \operatorname{argmin}_u \underbrace{\left(\frac{1}{2}\|u - \bar{u}_{\text{prior}}\|^2_{\Gamma^{-1}_{\text{prior}}} + \frac{1}{2}\|y_{\text{obs}} - Au - \bar{e}\|^2_{\Gamma^{-1}_{\text{noise}}}\right)}_{\mathcal{J}(u)},$$

given by maximum a posterior (MAP) point.
The posterior covariance matrix Γ_{post} is the inverse of the Hessian of $\mathcal{J}(u)$

$$\Gamma_{\text{post}} = \left(A^T \Gamma^{-1}_{\text{noise}} A + \Gamma^{-1}_{\text{prior}}\right)^{-1}.$$

Since A stems from discretization of time-dependent PDE, A is large and dense, direct computation of Γ_{post} is impossible.

Approximation of Γ_{post}

Note that

$$\begin{aligned}\Gamma_{\text{post}} &= \left(A^T \Gamma_{\text{noise}}^{-1} A + \Gamma_{\text{prior}}^{-1}\right)^{-1} \\ &= \Gamma_{\text{prior}}^{1/2} (\underbrace{\Gamma_{\text{prior}}^{1/2} A^T \Gamma_{\text{noise}}^{-1} A \Gamma_{\text{prior}}^{1/2}}_{\tilde{\mathcal{H}}_{\text{mis}}} + I)^{-1} \Gamma_{\text{prior}}^{1/2}\end{aligned}$$

It can be shown that the so-called 'prior-preconditioned data misfit Hessian' $\tilde{\mathcal{H}}_{\text{mis}}$ has low numerical rank [Flath et al 2011] $\rightsquigarrow$ Lanczos method can be used to obtain low-rank approximation

$$\tilde{\mathcal{H}}_{\text{mis}} \approx V \Lambda V^T.$$

With Sherman-Morrison-Woodbury formula,

$$\Gamma_{\text{post}} \approx \Gamma_{\text{prior}} - \Gamma_{\text{prior}}^{1/2} V \tilde{\Lambda} V^T \Gamma_{\text{prior}}^{1/2}$$

where $\tilde{\Lambda} = \text{diag}(\frac{\lambda_i}{1+\lambda_i})$.

Computational Challenage

Applying the Lanczos method to $\tilde{\mathcal{H}}_{\text{mis}} = \Gamma_{\text{prior}}^{1/2} A^T \Gamma_{\text{noise}}^{-1} A \Gamma_{\text{prior}}^{1/2}$, at each iteration, we need to solve

- one forward PDE
- one adjoint PDE

For time-dependent PDE

$$\frac{\partial}{\partial t} y - \mathcal{L} y = f$$

computations of the solution $y = \begin{bmatrix} y_1 & y_2 & y_3 & \cdots & y_{n_t} \end{bmatrix}$ gives $\mathcal{O}(n_x n_t)$ complexity for both computations and storage.
Here

- n_x: number of variables in space
- n_t: number of time steps

$\rightsquigarrow$ high storage requirements!

Low-Rank Lanczos

To overcome storage problem

- we do not directly compute $y = \begin{bmatrix} y_1 & y_2 & y_3 & \cdots & y_{n_t} \end{bmatrix}$,
- but a low-rank approximation of $y \approx u_x v_t^T$ with

$$\operatorname{rank}(u_x) = \operatorname{rank}(v_t) = r, \qquad r \ll \{n_x, n_t\}$$

  by low-rank in time approach [Stoll/Breiten 2015].

Low-rank in time approach gives

- $\mathcal{O}(n_x + n_t)$ computational complexity
- $\mathcal{O}(n_x + n_t)$ memory consumption

Applying the low-rank in time approach to the Lanczos method yields low-rank Lanczos,

- reorthogonalization is necessary due to truncation (perturbation)!

Example

We assume uncertainty in the initial condition

$$\begin{aligned} \frac{\partial}{\partial t} y - \Delta y &= 0, && \Omega \times (0, T) \\ y &= u, && \Omega \times \{t = 0\} \\ y &= 0, && \partial\Omega_D \times (0, T) \\ \nabla y \cdot \mathbf{n} &= 0, && \partial\Omega_N \times (0, T) \end{aligned}$$

The objective function becomes

$$\min_u \left(\frac{\beta_{\text{noise}}}{2} \int_0^T \int_\Omega (y - y_{\text{obs}})^2 \, b(x, t) dx dt + \frac{\beta_{\text{prior}}}{2} \int_\Omega (u - \bar{u})^2 dx \right).$$

Here, y_{obs} can only be obtained in the vicinity of the sensor locations and $b(x, t)$ is the observation operator with

$$b(x, t) = \Sigma_j \delta(x - x_j).$$

Settings for uncertainty in input, boundary conditions, etc., are similar.

Tensor Formulation

Discretizing using

- finite elements in space, and
- implicit Euler in time

simultaneously gives $\mathcal{K}\mathbf{y} = \mathcal{C}u$ and

$$\min_{u}\left(\frac{1}{2}\left(\mathbf{y}-\mathbf{y}_{\mathrm{obs}}\right)^{T}\mathcal{B}^{T}\Gamma_{\mathrm{noise}}^{-1}\mathcal{B}(\mathbf{y}-\mathbf{y}_{\mathrm{obs}})+\frac{1}{2}(\mathbf{u}-\bar{\mathbf{u}})^{T}\Gamma_{\mathrm{prior}}^{-1}(\mathbf{u}-\bar{\mathbf{u}})\right),$$

where $\Gamma_{\mathrm{noise}} = \frac{1}{\beta_{\mathrm{noise}}} I_{n_t} \otimes M$, and $\Gamma_{\mathrm{prior}} = \frac{1}{\beta_{\mathrm{prior}}} M$.
Here $\mathcal{K} = C \otimes M + I_{n_t} \otimes (\tau K)$, $\mathcal{C} = e_1 \otimes (\tau M)$, M is the mass matrix, K is the stiffness matrix, e_1 is the canonical unit vector, $\tau = T/n_t$ and C is a bidiagonal matrix that comes from the implicit Euler method.
This gives

$$\tilde{\mathcal{H}}_{\mathrm{mis}} = \Gamma_{\mathrm{prior}}^{\frac{1}{2}}\mathcal{C}^{T}\mathcal{K}^{-T}\mathcal{B}^{T}\Gamma_{\mathrm{noise}}^{-1}\mathcal{B}\mathcal{K}^{-1}\mathcal{C}\Gamma_{\mathrm{prior}}^{\frac{1}{2}}.$$

Other settings for Γ_{prior} are also possible.

Preconditioning

Recall that

$$\tilde{\mathcal{H}}_{\text{mis}} = \Gamma_{\text{prior}}^{\frac{1}{2}} \mathcal{C}^T \mathcal{K}^{-T} \mathcal{B}^T \Gamma_{\text{noise}}^{-1} \mathcal{B} \mathcal{K}^{-1} \mathcal{C} \Gamma_{\text{prior}}^{\frac{1}{2}}.$$

Applying low-rank Lanczos to $\tilde{\mathcal{H}}_{\text{mis}}$ requires matrix-vector multiplications with the involved operators.

Take application of $\mathcal{K}^{-1}$ as an example:

$$(I_{n_t} \otimes L + C \otimes M)\, vec(X) = vec(F). \tag{1}$$

To solve (1), we use the alternative minimal energy (AMEn) approach and the tensor-train (TT) toolbox. At each AMEn iteration,

- left Galerkin projection

$$\left(I_n \otimes L + \hat{C} \otimes M\right) \hat{x} = \hat{b}, \tag{2}$$

- or right Galerkin projection

$$\left(I_n \otimes \tilde{L} + C \otimes \tilde{M}\right) \tilde{x} = \tilde{b}. \tag{3}$$

$P = \text{diag}(I_n) \otimes L + \text{diag}(\hat{C}) \otimes M$ for (2) and direct method for (3)

Implementation Details

Consider the uncertainty in the initial condition case. Two problems are used

- Heat equation
- Convection-diffusion equation

For both cases, we have data observed from 9 sensors which are uniformly located inside the domain.

Hardware

- Intel i5 CPU, 3.60GHZ
- 8GB RAM

Software

- MATLAB 2011b
- Tensor Train Toolbox (TT-Toolbox)
- tensor approximation tolerance: 10^{-8}
- PDE solver tolerance: 10^{-8}

Low-rank Lanczos vs 'eigs'

Heat equation, 32×32 mesh

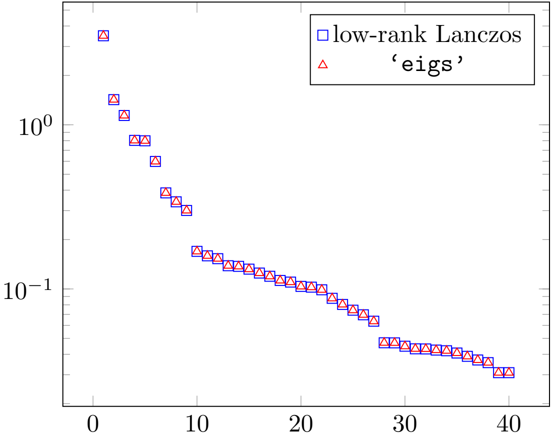

Figure: $\lambda(\tilde{\mathcal{H}}_{\text{mis}})$, $n_t = 60$

Maximum Rank vs n_t

64×64 mesh

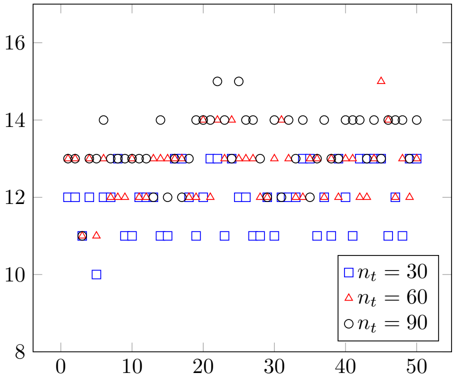

Figure: Heat equation

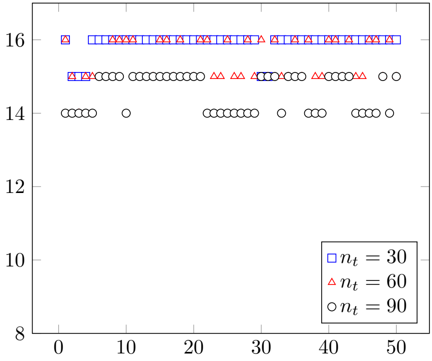

Figure: Convection-diffusion equation, $\nu = 10^{-2}$

Maximum Rank vs h and ν

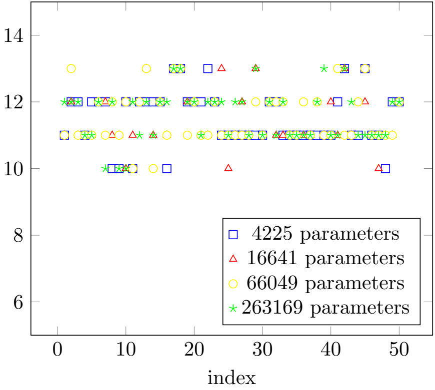

Figure: Heat equation, $n_t = 30$

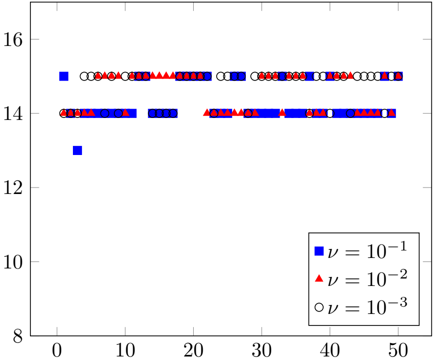

Figure: Convection-diffusion equation, 64×64 mesh, $n_t = 30$

Preconditioner Performance

Heat equation, 1st Lanczos iteration, 2nd AMEn (outer) iteration, GMRES for inner iteration

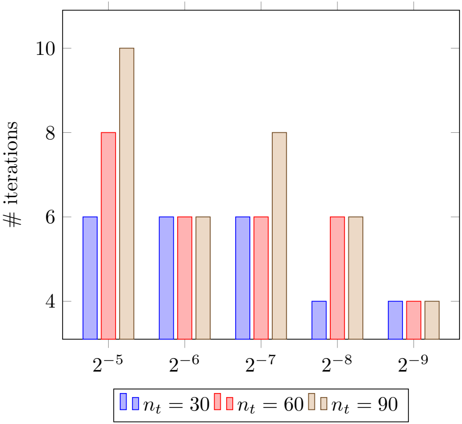

Preconditioner Performance

Convection-diffusion problem, 1st Lanczos iteration, 2nd AMEn (outer) iteration, GMRES for inner iteration

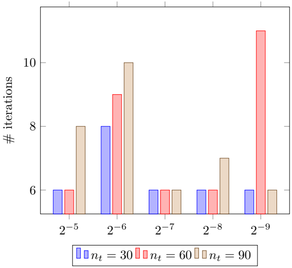

Figure: $\nu = 10^{-1}$

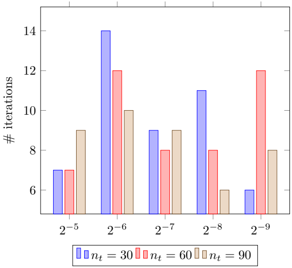

Figure: $\nu = 10^{-2}$

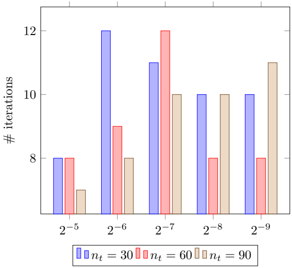

Figure: $\nu = 10^{-3}$

Approximated Γ_{post}

We show different truncation tolerance ε for the eigenvalues computations of low-rank Lanczos method. We plot the diagonal entries of Γ_{post} for a 64×64 mesh, $\beta_{\text{noise}} = 10^6 \beta_{\text{prior}}$.

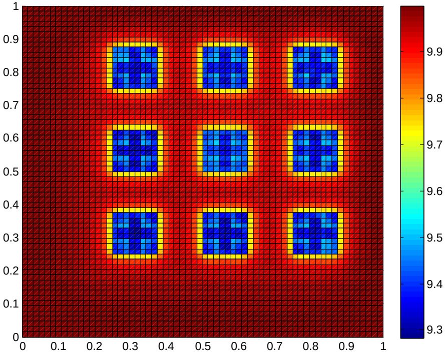

Figure: $\varepsilon = 10^0$, 72 Lanczos steps

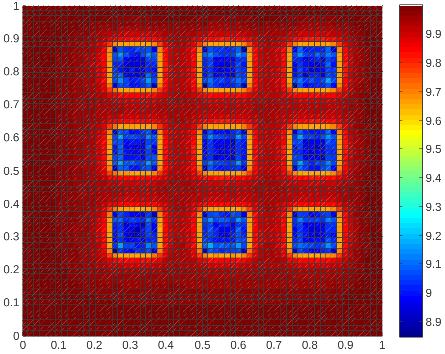

Figure: $\varepsilon = 10^{-2}$, 187 Lanczos steps

Approximated Γ_{post} (cont'd)

$\beta_{\text{noise}} = 10^8 \beta_{\text{prior}}$.

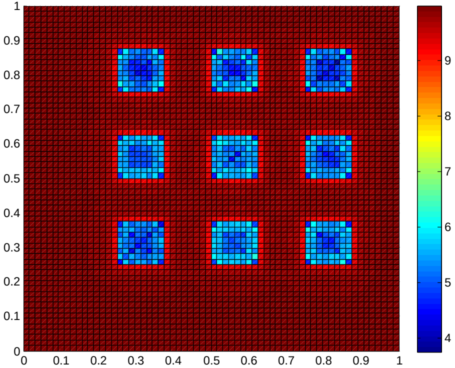

Figure: $\varepsilon = 10^0$, 347 Lanczos steps

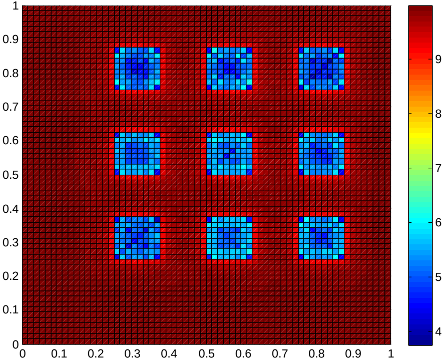

Figure: $\varepsilon = 10^{-2}$, 591 Lanczos steps

Conclusions & Remarks

- Low-rank method enables to solve even larger linear Bayesian inverse problems;
- Maximum rank remains bounded with the increase of n_x and n_t;
- Block diagonal preconditioner + AMEn solver give satisfactory performance;
- Extending to nonlinear inverse problem is challenging (non-Gaussian Posterior).

References

- A. M. Stuart. *Inverse problems: A Bayesian perspective*, Acta Numerica 19:451–559, 2010.
- H. P. Flath, L. C. Wilcox, V. Akelik, J. Hill, B. van Bloemen Waanders, and O. Ghattas. *Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations*, SIAM Journal on Scientific Computing 33(1):407–432, 2011.
- M. Stoll and T. Breiten. *A low-rank in time approach to PDE-constrained optimization*, SIAM Journal on Scientific Computing 37(1):B1–B29, 2015.
- P. Benner, Y. Qiu, and M. Stoll. *Low-rank computations of posterior covariance matrices in Bayesian inverse problems*, arXiv:1703.05638.
- TT-Toolbox, `https://github.com/oseledets/TT-Toolbox`.