Localization for high dimensional data assimilation and MCMC

Xin T Tong

National University of Singapore
Joint work with Matthias Morzfeld and Youssef Marzouk

Potsdam 20/07/2018

Content

- Ensemble Kalman filter (EnKF)
- High dimensional challenges for EnKF.
- Sparse/Localized scenario:
  - EnKF with domain localization,
  - with a stable localized structure
  - reaches its proclaimed performance,
  - if the ensemble size $K > C_L \log d$ for a constant C_L.
- Localization for inverse problems.
- Gibbs sampler on Gaussian distributions.
- l-MwG for inverse problems.

Signal-observation system

$$\text{Signal:}\quad X_{n+1} = A_n X_n + B_n + \xi_{n+1}, \quad \xi_{n+1} \sim \mathcal{N}(0, \Sigma_n)$$
$$\text{Observation:}\quad Y_{n+1} = H X_{n+1} + \zeta_{n+1}, \quad \zeta_{n+1} \sim \mathcal{N}(0, \sigma_o^2 I_q)$$

Goal: estimate X_n based on $Y_1, \ldots, Y_n$

Weather forecast

- Signal: $X_{n+1} = A_n X_n + B_n + \xi_{n+1}$,
  Observation: $Y_{n+1} = H X_{n+1} + \zeta_{n+1}$.
- Weather forecast:
  Signal: atmosphere and ocean, “follows” a PDE.
  Obs: weather station, satellite, sensors.....
- Main challenge: high dimension, $d \sim 10^6 - 10^8$.

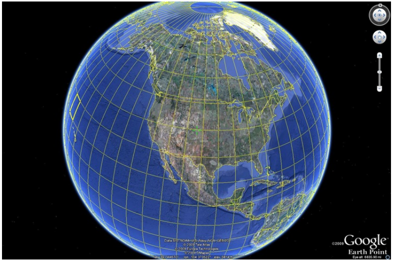

Kalman filter

- Use Gaussian: $X_n|_{Y_{1\dots n}} \sim \mathcal{N}(m_n, R_n)$
- Forecast step: $\hat{m}_{n+1} = A_n m_n + B_n, \widehat{R}_{n+1} = A_n R_n A_n^T + \Sigma_n$.
- Assimilation step: apply the Kalman update rule

$$m_{n+1} = \hat{m}_{n+1} + \mathcal{G}(\widehat{R}_{n+1})(Y_{n+1} - H\hat{m}_{n+1}), \quad R_{n+1} = \mathcal{K}(\widehat{R}_{n+1})$$

$$\mathcal{G}(C) = CH^T(\sigma_o^2 I_q + HCH^T)^{-1}, \quad \mathcal{K}(C) = C - \mathcal{G}(C)HC$$

- Complexity: $O(d^3)$.

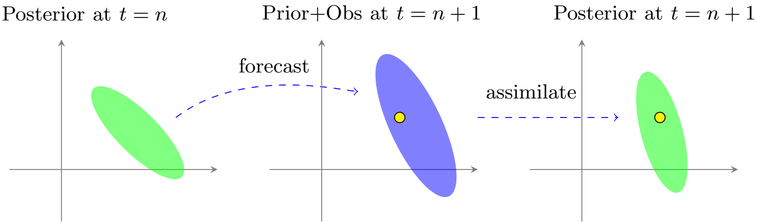

Sampling+Gaussian

- Monte Carlo: use samples to represent a distribution:

$$X^{(1)}, \dots, X^{(K)} \sim p, \quad \frac{1}{K}\sum_{k=1}^{K} \delta_{X^{(k)}} \approx p.$$

- Ensemble $\{X_n^{(k)}\}_{k=1}^{K}$ to represent $\mathcal{N}(\overline{X}_n, C_n)$

$$\overline{X}_n = \frac{\sum X_n^{(k)}}{K}, \quad C_n = \frac{1}{K-1}\sum_k (X_n^{(k)} - \overline{X}_n) \otimes (X_n^{(k)} - \overline{X}_n).$$

Ensemble Kalman filter

- Forecast step

$$\widehat{X}^{(k)}_{n+1} = A_n X^{(k)}_n + B_n + \xi^{(k)}_{n+1}, \quad \widehat{C}_{n+1} = \frac{\widehat{S}_{n+1}\widehat{S}^T_{n+1}}{K-1}.$$

- Assimilation step

$$X^{(k)}_{n+1} = \widehat{X}^{(k)}_{n+1} + \mathcal{G}(\widehat{C}_{n+1})(Y_{n+1} - H\widehat{X}^{(k)}_{n+1} + \zeta^{(k)}_{n+1}).$$

Gain matrix: $\mathcal{G}(C) = CH^T(\sigma_o^2 I_q + HCH^T)^{-1}$.

- Complexity $O(K^2 d)$.

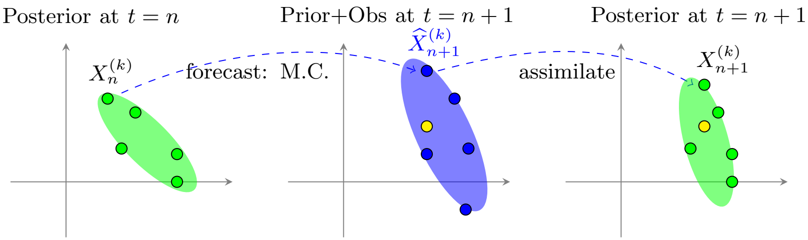

Application and theory of EnKF

Application:

- Successful weather forecast and oil reservoir management.
- Recently been applied to deep neural networks.
- $K = 50$ ensembles can forecast $d = 10^6$ dimensional systems.
- Extreme savings: $10^{10} = dK^2 \ll d^3 = 10^{18}$.

Theoretical Literature

- Focused on showing ensemble version $(\overline{X}_n, C_n) \to (m_n, R_n)$
- Require $K \to \infty$ (Mandel, Cobb, Beezley 11)
- Fixed d sufficiently large K, $|A| < 1$ (Del Moral, Tugaut 16)
- Perturbation interpretation (Bishop, Del Moral, Pathiraja 17)
- Fixed K, well definedness $\mathbb{E}|X_n^{(k)}|^2 < \infty$ (Law, Kelly, Stuart, 14)
- Fixed K, boundedness $\sup_n \mathbb{E}|X_n^{(k)}|^2 < \infty$ (Tong, Majda, Kelly 15)
- Continuous version (de Wilijes, Reich, Stannat 17)

Gap: dependence or independence of K on d.

Application and theory of EnKF

Application:

- Successful weather forecast and oil reservoir management.
- Recently been applied to deep neural networks.
- $K = 50$ ensembles can forecast $d = 10^6$ dimensional systems.
- Extreme savings: $10^{10} = dK^2 \ll d^3 = 10^{18}$.

Theoretical Literature

- Focused on showing ensemble version $(\overline{X}_n, C_n) \to (m_n, R_n)$
- Require $K \to \infty$ (Mandel, Cobb, Beezley 11)
- Fixed d sufficiently large K, $|A| < 1$ (Del Moral, Tugaut 16)
- Perturbation interpretation (Bishop, Del Moral, Pathiraja 17)
- Fixed K, well definedness $\mathbb{E}|X_n^{(k)}|^2 < \infty$ (Law, Kelly, Stuart, 14)
- Fixed K, boundedness $\sup_n \mathbb{E}|X_n^{(k)}|^2 < \infty$ (Tong, Majda, Kelly 15)
- Continuous version (de Wilijes, Reich, Stannat 17)

Gap: dependence or independence of K on d.

Application:

- Successful weather forecast and oil reservoir management.
- Recently been applied to deep neural networks.
- $K = 50$ ensembles can forecast $d = 10^6$ dimensional systems.
- Extreme savings: $10^{10} = dK^2 \ll d^3 = 10^{18}$.

Theoretical Literature

- Focused on showing ensemble version $(\overline{X}_n, C_n) \to (m_n, R_n)$
- Require $K \to \infty$ (Mandel, Cobb, Beezley 11)
- Fixed d sufficiently large K, $|A| < 1$ (Del Moral, Tugaut 16)
- Perturbation interpretation (Bishop, Del Moral, Pathiraja 17)
- Fixed K, well definedness $\mathbb{E}|X_n^{(k)}|^2 < \infty$ (Law, Kelly, Stuart, 14)
- Fixed K, boundedness $\sup_n \mathbb{E}|X_n^{(k)}|^2 < \infty$ (Tong, Majda, Kelly 15)
- Continuous version (de Wilijes, Reich, Stannat 17)

Gap: dependence or independence of K on d.

Challenges from high dimension

Ensemble size K to represent uncertainty of dimension d:

- Spurious correlation in high dimension.
  Suppose $X_n^{(k)} \sim \mathcal{N}(0, I_d)$ i.i.d, by Bai-Yin's law

$$\|C_n - I_d\| \approx \sqrt{d/K} \quad \text{with large probability}$$

- Rank deficiency: $C_n = \frac{\sum_{k=1}^{K}(X_n^{(k)} - \overline{X}_n)(X_n^{(k)} - \overline{X}_n)^T}{K-1}$
  Has rank$(C_n) \leq K - 1$, see as $\begin{bmatrix} C_n & 0 \\ 0 & 0 \end{bmatrix} \begin{matrix} \} \text{ K-1} \\ \} \text{ d-K+1} \end{matrix}$

Heuristic answers

We need conditions! Answers from practitioners

- Low effective dimension.
- Localized covariance structure.

As comparison: for high dimensional numerical problems,

- Low rank structure
- Sparse structure

can be exploited for efficient computation.

Local interaction

- High dimension often comes from dense grids.
- Interaction often is local: PDE discritization:

$$\partial_x x(t) \Rightarrow \frac{1}{2h}(x_{i+1}(t) - x_{i-1}(t)).$$

- Example: Lorenz 96 model

$$\dot{x}_i(t) = (x_{i+1} - x_{i-2})x_{i-1} - x_i dt + F, \quad i = 1, \cdots, d$$

- Information travels along interaction, and is dissipated.

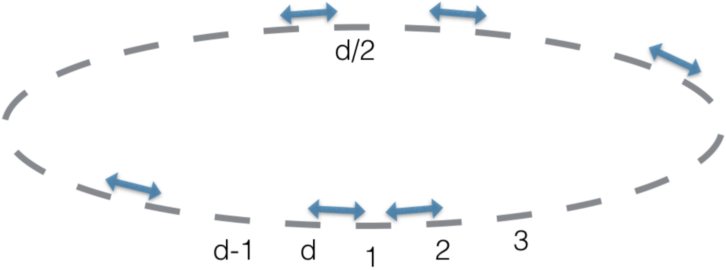

Sparsity: local covariance

- Correlation depends on information propagation.
- Correlation decays quickly with the distance.
- Covariance is localized with a structure Φ, e.g. $\Phi(x) = \rho^x$

$$[\widehat{C}_n]_{i,j} \propto \Phi(|i-j|)$$

$\Phi(x) \in [0,1]$ is decreasing. Distance can be general.

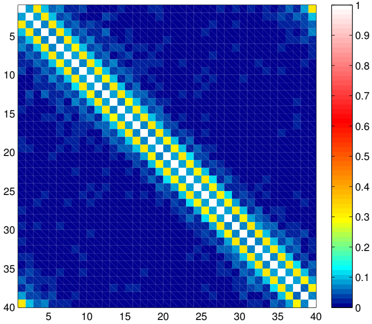

Correlation of Lorenz 96

- Spurious correlation may exist for far away terms.
- Localization: simply ignore far away correlations.
- Implementation: Schur product with a mask

$$[\widehat{C}_n \circ \mathbf{D}_L]_{i,j} = [\widehat{C}_n]_{i,j} \cdot [\mathbf{D}_L]_{i,j}$$

  Use $\widehat{C}_n \circ \mathbf{D}_L$ to describe uncertainty
- $[\mathbf{D}_L]_{i,j} = \phi(|i-j|)$, with a radius L.
  Gaspari-Cohn matrix: $\phi(x) = \exp(-4x^2/L^2)\mathbf{1}_{|i-j|\leq L}$.
  Cutoff/Branding matrix: $\phi(x) = \mathbf{1}_{|i-j|\leq L}$.
- Also resolves rank deficiency, e.g.

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} \circ \begin{bmatrix} 1 & 0.2 & 0 \\ 0.2 & 1 & 0.2 \\ 0 & 0.2 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0.2 & 0 \\ 0.2 & 1 & 0.2 \\ 0 & 0.2 & 1 \end{bmatrix}.$$

Domain localization

Two types LEnKF: Domain localization and covariance tempering.

Domain localization with radius l:
Assume H is a partial observation matrix
Use information in $\mathcal{I}_i = \{j : |i-j| \le l\}$ to update component i

$$C^i = P_{\mathcal{I}_i} C P_{\mathcal{I}_i}, \quad \mathcal{G}^i(C) = C^i H^T(\sigma_o^2 I_q + HC^i H^T)^{-1}$$

$$\mathcal{G}^L(C) = \sum e_i e_i^T \mathcal{G}^i(C)$$

$$X_n^{(k)} = \widehat{X}_n^{(k)} + \mathcal{G}^L(\widehat{C}_n)(Y_n - H\widehat{X}_n^{(k)} + \zeta_n^{(k)}).$$

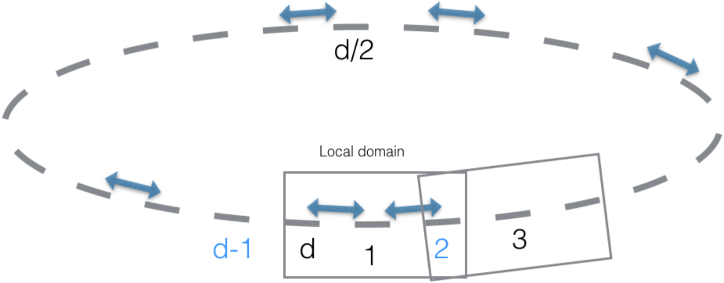

Advantage with localization

Intuitively, ignoring the long distance covariance terms, reduces the sampling difficulty, and necessary sampling size.

Theorem (Bickel, Levina 08)

If $X^{(1)}, \dots, X^{(K)} \sim \mathcal{N}(0, \Sigma)$, *denote* $C = \frac{1}{K}\sum_{k=1}^{K} X^{(k)} \otimes X^{(k)}$.
$\|\mathbf{D}_L\|_1 = \max_i \sum_j |\mathbf{D}_L|_{i,j}$. *There is a constant* c*, and for any* $t > 0$

$$\mathbb{P}(\|C \circ \mathbf{D}_L - \Sigma \circ \mathbf{D}_L\| > \|\mathbf{D}_L\|_1 t) \leq 8\exp(2\log d - cK\min\{t, t^2\})$$

This indicates that $K \propto \|\mathbf{D}_L\|_1^2 \log d$ *is the necessary sample size.*

$\|\mathbf{D}_L\|$ is independent of d, e.g, the cut-off/branding matrix,
$[\mathbf{D}^L_{cut}]_{i,j} = \mathbf{1}_{|i-j|\leq L}$, $\|\mathbf{D}^L_{cut}\|_1 \approx 2L$.

Advantage with localization

Intuitively, ignoring the long distance covariance terms, reduces the sampling difficulty, and necessary sampling size.

Theorem (Bickel, Levina 08)

If $X^{(1)}, \dots, X^{(K)} \sim \mathcal{N}(0, \Sigma)$, *denote* $C = \frac{1}{K}\sum_{k=1}^{K} X^{(k)} \otimes X^{(k)}$. $\|\mathbf{D}_L\|_1 = \max_i \sum_j |\mathbf{D}_L|_{i,j}$. *There is a constant* c, *and for any* $t > 0$

$$\mathbb{P}(\|C \circ \mathbf{D}_L - \Sigma \circ \mathbf{D}_L\| > \|\mathbf{D}_L\|_1 t) \leq 8 \exp(2 \log d - cK \min\{t, t^2\})$$

This indicates that $K \propto \|\mathbf{D}_L\|_1^2 \log d$ *is the necessary sample size.*

$\|\mathbf{D}_L\|$ is independent of d, e.g, the cut-off/branding matrix, $[\mathbf{D}^L_{cut}]_{i,j} = \mathbf{1}_{|i-j|\leq L}$, $\|\mathbf{D}^L_{cut}\|_1 \approx 2L$.

Main result: localized EnKF

Theorem (T. 18)

Suppose the system coefficients have bandwidth l, and the LEnKF ensemble covariance admits a stable localized structure, then for any $\delta > 0$, LEnKf reaches its proclaimed performance with high probability $1 - O(\delta)$:

$$1 - \frac{1}{T}\sum_{t=1}^{T}\mathbb{P}(\mathbb{E}_S \hat{e}_n \otimes \hat{e}_n \preceq (1+\delta)(\widehat{C}_n \circ \mathbf{D}^{4l}_{cut} + \rho I_d)) \leq \frac{1}{T}D_0 + D_1\delta,$$

if the sample size $K > D_{l,\delta} \log d$.

$\mathbb{E}_S$ conditioned on the information of the sampling noise realization.

Main result: localized EnKF

Theorem (T. 18)

Suppose the system coefficients have bandwidth l, and the LEnKF ensemble covariance admits a stable localized structure, then for any $\delta > 0$, LEnKf reaches its proclaimed performance with high probability $1 - O(\delta)$:

$$1 - \frac{1}{T}\sum_{t=1}^{T} \mathbb{P}(\mathbb{E}_S \hat{e}_n \otimes \hat{e}_n \preceq (1+\delta)(\widehat{C}_n \circ \mathbf{D}_{cut}^{4l} + \rho I_d)) \leq \frac{1}{T} D_0 + D_1 \delta,$$

if the sample size $K > D_{l,\delta} \log d$.

$\mathbb{E}_S$ conditioned on the information of the sampling noise realization.

Proclaimed performance

Proclaimed/estimated performance

- EnKF estimates X_n by $\overline{X}_n = \frac{1}{K}\sum X_n^{(k)}$.
- Error $e_n = \overline{X}_n - X_n$. Covariance : $\mathbb{E}e_n e_n^T = \mathbb{E}e_n \otimes e_n$.
- EnKF estimates its performance by ensemble covariance C_n^ρ.
- Can it captures the error covariance?

$$\mathbb{E}C_n^\rho \succeq \mathbb{E}e_n \otimes e_n$$

$$1 - \frac{1}{T}\sum_{t=1}^{T}\mathbb{P}(\mathbb{E}_S \hat{e}_n \otimes \hat{e}_n \preceq (1+\delta)(\widehat{C}_n \circ \mathbf{D}_{cut}^{4l} + \rho I_d)) \leq \frac{1}{T}D_0 + D_1\delta,$$

Proclaimed performance

Proclaimed/estimated performance

- EnKF estimates X_n by $\overline{X}_n = \frac{1}{K}\sum X_n^{(k)}$.
- Error $e_n = \overline{X}_n - X_n$. Covariance : $\mathbb{E}e_n e_n^T = \mathbb{E}e_n \otimes e_n$.
- EnKF estimates its performance by ensemble covariance C_n^ρ.
- Can it captures the error covariance?

$$\mathbb{E}C_n^\rho \succeq \mathbb{E}e_n \otimes e_n$$

$$1 - \frac{1}{T}\sum_{t=1}^{T} \mathbb{P}(\mathbb{E}_S \hat{e}_n \otimes \hat{e}_n \preceq (1+\delta)(\widehat{C}_n \circ \mathbf{D}_{cut}^{4l} + \rho I_d)) \leq \frac{1}{T}D_0 + D_1\delta,$$

Stable localized structure

- Intuitively, we need some conditions on the covariance structure.
- Stable localized structure: with local structure function Φ, e.g. $\Phi(x) = \lambda^x$,

$$[\widehat{C}_n]_{i,j} \leq M_n \Phi(|i-j|), \quad \sum_{n=1}^{T} \mathbb{E} M_n \leq T M_*.$$

  M_n describes how localized the sample covariance matrix is.
- Why is this necessary?

LEnKF inconsistency

An intrinsic bias/inconsistency in LEnKF.

- Localization creates a bias.
- Target covariance by Bayes formula

$$(I-\mathcal{G}^L(\widehat{C}_n)H)[\widehat{C}_n\circ \mathbf{D}_L](I-\mathcal{G}^L(\widehat{C}_n)H)^T+\sigma_o^2\mathcal{G}^L(\mathcal{G}^L)^T.$$

- LEnKF implementation

$$X_n^{(k)}=\widehat{X}_n^{(k)}+\mathcal{G}^L(\widehat{C}_n)(Y_n-H\widehat{X}_n^{(k)}+\zeta_n^{(k)})$$

- Average ensemble covariance

$$C_n\circ \mathbf{D}_L=[(I-\mathcal{G}^L(\widehat{C}_n)H)\widehat{C}_n(I-\mathcal{G}^L(\widehat{C}_n)H)^T+\sigma_o^2\mathcal{G}^L(\mathcal{G}^L)^T]\circ \mathbf{D}_L.$$

- Difference: commuting the localization and Kalman update.
- Previously investigated numerically by Nerger 2015, the inconsistency can lead to error growth.

When is the inconsistency small?

- localization is applied, covariance is assumed localized.
- Given localized structure Φ, find M_n so that

$$[\widehat{C}_n]_{i,j} \leq M_n \Phi(|i-j|).$$

- Interestingly, when $\mathbf{D}_L$ is $\mathbf{D}_{4l}^{cut}$, the

$$\text{Localization inconsistency} \leq C M_n \Phi(2l).$$

If $2l$ is large, $\Phi(x) = \lambda^x$, this difference can be controlled.
- Localized covariance leads to small localization inconsistency.
- Therefore, we need M_n to be a stable sequence,

$$\sum_{n=1}^{T} \mathbb{E} M_n \leq T M_*.$$

When is the inconsistency small?

- localization is applied, covariance is assumed localized.
- Given localized structure Φ, find M_n so that

$$[\widehat{C}_n]_{i,j} \leq M_n \Phi(|i-j|).$$

- Interestingly, when $\mathbf{D}_L$ is $\mathbf{D}_{4l}^{cut}$, the

$$\text{Localization inconsistency} \leq C M_n \Phi(2l).$$

  If $2l$ is large, $\Phi(x) = \lambda^x$, this difference can be controlled.
- Localized covariance leads to small localization inconsistency.
- Therefore, we need M_n to be a stable sequence,

$$\sum_{n=1}^{T} \mathbb{E} M_n \leq T M_*.$$

When is covariance localized?

Practical perspective

- Simply assumed.
- Numerically checked.

Theoretical perspective: does covariance localize for any stochastic system?

- Linear system: covariance can be computed.
- Nonlinear: difficult, e.g. Lorenz 96.
- LEnKF: difficult since assimilation is nonlinear.
- Under strong conditions:
  - Weak local interaction, strong dissipation.
  - Sparse observation for simplicity.
- Also scales with the noise strength.

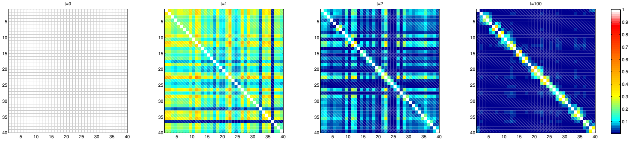

When is covariance localized?

Practical perspective

- Simply assumed.
- Numerically checked.

Theoretical perspective: does covariance localize for any stochastic system?

- Linear system: covariance can be computed.
- Nonlinear: difficult, e.g. Lorenz 96.
- LEnKF: difficult since assimilation is nonlinear.
- Under strong conditions:
  - Weak local interaction, strong dissipation.
  - Sparse observation for simplicity.
- Also scales with the noise strength.

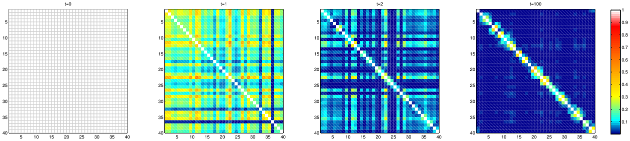

Theorem

Suppose the following, then a stable localized structure with $\Phi(x) = \lambda_A^x$

1) *The system noise is diagonal and the observations are sparse* $\Sigma_n = \sigma_\xi^2 I_d, \quad \mathbf{d}(o_i, o_j) > 2l, \ \forall i \neq j.$
2) *There is a* $\lambda_A < r^{-1}$, $\max_i \left\{\sum_{k=1}^d |[A_n]_{i,k}|\lambda_A^{-\mathbf{d}(i,k)}\right\} \leq \lambda_A.$
3) *There are constants such that* $\psi_{\lambda_A}(M_*, \delta_*) \leq M_*$

$$0 < \delta_* < \min\{0.25, \tfrac{1}{2}(\lambda_A^{-1} - r)\}, \quad M_* \geq \frac{(r + 2\delta_*)\sigma_\xi^2}{1 - \lambda_A},$$

$$\psi_{\lambda_A}(M, \delta) = (r+\delta) \max\left\{\lambda_A M \left(1 + \sigma_o^{-2} M\right)^2 + \lambda_A \sigma_o^{-2} M^2, \lambda_A^2 M + \sigma_\xi^2\right\}.$$

4) *Denote* $n_* = 2L + \lceil \frac{\log 4\delta_*^{-1}}{\log \lambda_A^{-1}} \rceil$. *The sample size* K *exceeds*

$$K > \max\left\{-\frac{1}{c\delta_*^2 \lambda_A^{2L}} \log(16 d^2 n_* \delta_*^{-2}), \Gamma(2r\delta_*^{-1}, d)\right\}.$$

A real example

A stochastically forced dissipative advection equation:

$$\frac{\partial u(x,t)}{\partial t} = c\frac{\partial u(x,t)}{\partial x} - \nu u(x,t) + \mu \frac{\partial^2 u(x,t)}{\partial x^2} + \sigma_x \dot{W}(x,t).$$

Discretization

$$X_{n+1,i} = a_- X_{n,i-1} + a_0 X_{n,i} + a_+ X_{n,i+1} + \sigma_x \sqrt{\Delta t} W_{n+1,i}, \quad i = 1, \dots, d;$$
$$a_- = \tfrac{\mu \Delta t}{h^2} - \tfrac{c\Delta t}{2h}, \quad a_0 = 1 - \tfrac{2\mu\Delta t}{h^2} - \nu \Delta t, \quad a_+ = \tfrac{\mu \Delta t}{h^2} + \tfrac{c\Delta t}{2h}.$$

Observe $Y_{n,k} = X_{n,p(k-1)+1} + \sigma_o B_{n,k}$.

Stable regime

Strong damping+weak advection

$$h = 1, \quad \Delta t = 0.1, \quad p = 5, \quad \nu = 5, \quad c = 0.1, \quad \mu = 0.1, \quad \sigma_x = \sigma_o = 1.$$

Direct verification of the conditions is possible.

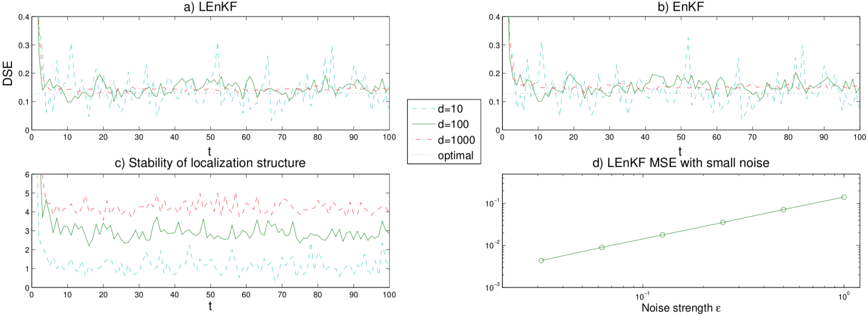

Strong advection regime

Weak damping+strong advection

$$h = 0.2, \quad \Delta t = 0.1, \quad p = 5, \quad \nu = 0.1, \quad c = 2, \quad \mu = 0.1, \quad \sigma_x = \sigma_o = 1.$$

Direct verification of the conditions is not possible.

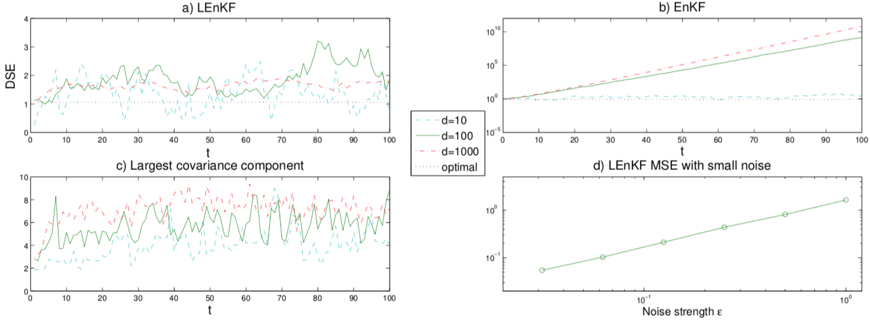

Extension of localization

- Localization has made EnKF very effective for high dimensional DA problems.
- Various generalization to particle filters.
- Often relies on Gaspari Cohn matrices.
- Makes non-Gaussian application difficult.
- Non-ad hoc ways generalize localization to PF?
- Can we apply localization to other UQ problem?

Bayesian inverse problem

- Suppose $\mathbf{x} \sim p_0 = \mathcal{N}(\mathbf{m}, \mathbf{C})$, we observe

$$\mathbf{y} = h(\mathbf{x}) + \mathbf{v}, \quad \mathbf{v} \sim \mathcal{N}(0, R).$$

Try to recover the value and uncertainty of $\mathbf{x}$.

- Possible applications:
  - $\mathbf{x}$ is the real image, h defocus map.
  - $\mathbf{x}$ initial condition, h forward map of a PDE.
  - $\mathbf{x}$ model parameters, h gives model outcome.

Often $\mathbf{x}$ is high dimension.

- Bayesian approach: try to sample the posterior

$$p(\mathbf{x}|\mathbf{y}) \propto p_0(\mathbf{x}) p_l(\mathbf{y}|\mathbf{x}).$$

$p_l(\mathbf{y}|\mathbf{x}) = \mathcal{N}(h(\mathbf{x}), R)$.

Bayesian inverse problem

- Suppose $\mathbf{x} \sim p_0 = \mathcal{N}(\mathbf{m}, \mathbf{C})$, we observe

$$\mathbf{y} = h(\mathbf{x}) + \mathbf{v}, \quad \mathbf{v} \sim \mathcal{N}(0, R).$$

Try to recover the value and uncertainty of $\mathbf{x}$.

- Possible applications:
  - $\mathbf{x}$ is the real image, h defocus map.
  - $\mathbf{x}$ initial condition, h forward map of a PDE.
  - $\mathbf{x}$ model parameters, h gives model outcome.

Often $\mathbf{x}$ is high dimension.

- Bayesian approach: try to sample the posterior

$$p(\mathbf{x}|\mathbf{y}) \propto p_0(\mathbf{x}) p_l(\mathbf{y}|\mathbf{x}).$$

$p_l(\mathbf{y}|\mathbf{x}) = \mathcal{N}(h(\mathbf{x}), R)$.

Bayesian inverse problem

- Suppose $\mathbf{x} \sim p_0 = \mathcal{N}(\mathbf{m}, \mathbf{C})$, we observe

$$\mathbf{y} = h(\mathbf{x}) + \mathbf{v}, \quad \mathbf{v} \sim \mathcal{N}(0, R).$$

  Try to recover the value and uncertainty of $\mathbf{x}$.

- Possible applications:
  - $\mathbf{x}$ is the real image, h defocus map.
  - $\mathbf{x}$ initial condition, h forward map of a PDE.
  - $\mathbf{x}$ model parameters, h gives model outcome.

  Often $\mathbf{x}$ is high dimension.

- Bayesian approach: try to sample the posterior

$$p(\mathbf{x}|\mathbf{y}) \propto p_0(\mathbf{x}) p_l(\mathbf{y}|\mathbf{x}).$$

  $p_l(\mathbf{y}|\mathbf{x}) = \mathcal{N}(h(\mathbf{x}), R)$.

MCMC sampler

- Given a target distribution $p(\mathbf{x})$, generate a sequence of samples

$$\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \cdots, \mathbf{x}^{(N)}$$

  Use sample statistics to approximate population ones.

- Standard MCMC steps
  - Generate proposals $\mathbf{x}' \sim q(\mathbf{x}^{(k)}, \mathbf{x}')$
  - Accept with prob.
    $\alpha(\mathbf{x}, \mathbf{x}') = \min\{1, p(\mathbf{x}')q(\mathbf{x}', \mathbf{x}^{(k)})/q(\mathbf{x}^{(k)}, \mathbf{x}')p(\mathbf{x})\}$
- Popular choices of proposals $\xi_k \sim \mathcal{N}(0, I_d)$.
  - RWM: $\mathbf{x}' = \mathbf{x}_k + \sigma\xi_k$
  - MALA: $\mathbf{x}' = \mathbf{x}_k + \frac{\sigma^2}{2}\nabla \log p(\mathbf{x}_k) + \sigma\xi_k$.
  - pCN: $\Delta\mathbf{x}'_{k+1} = \sqrt{1-\beta^2}\Delta\mathbf{x}_k + \beta\,\xi_k$.

  Also emcee and Hamiltonian MCMC.
  σ, β are tuning parameters.

Sample isotropic Gaussian $p = \mathcal{N}(0, I_d)$.
Measurement of efficiency: integrated auto-correlation time (IACT)
Measure how many iterations to get an "uncorrelated" sample.

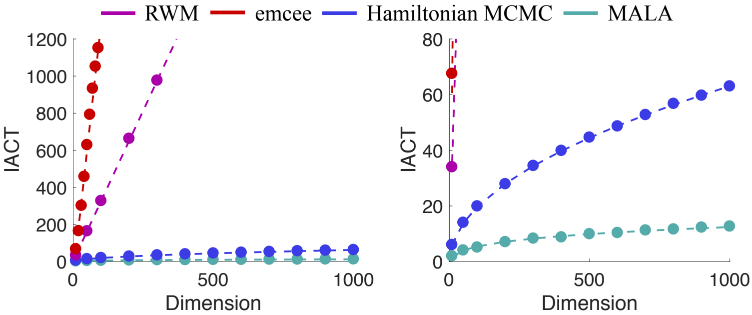

All increases with dimension

Where is the problem?

- Let's look at RWM, assume $\mathbf{x}_k = \mathbf{0}$.
  - Propose $\mathbf{x}' = \sigma\xi_k$
  - Accept with probability $\exp(-\frac{1}{2}\sigma^2\|\xi_k\|^2) \sim \exp(-\frac{1}{2}\sigma^2 d)$.
- If we keep $\sigma = 1$, "never" accept if $d > 20$.
- If we want acceptance at a constant rate, $\sigma = d^{-\frac{1}{2}}$.
  But then $\mathbf{x}' = \mathbf{x}_k + \sigma\xi_k$ is highly correlated with $\mathbf{x}_k$.
- Similar for MALA, $\sigma = d^{-1/3}$. Hamiltonian MCMC, $\sigma = d^{-1/4}$.
- Is it possible to break this curse of dimensionality?
- Is high dimensionality an issue in other related fields?

Where is the problem?

- Let's look at RWM, assume $\mathbf{x}_k = \mathbf{0}$.
  - Propose $\mathbf{x}' = \sigma\xi_k$
  - Accept with probability $\exp(-\frac{1}{2}\sigma^2\|\xi_k\|^2) \sim \exp(-\frac{1}{2}\sigma^2 d)$.
- If we keep $\sigma = 1$, "never" accept if $d > 20$.
- If we want acceptance at a constant rate, $\sigma = d^{-\frac{1}{2}}$.
  But then $\mathbf{x}' = \mathbf{x}_k + \sigma\xi_k$ is highly correlated with $\mathbf{x}_k$.
- Similar for MALA, $\sigma = d^{-1/3}$. Hamiltonian MCMC, $\sigma = d^{-1/4}$.
- Is it possible to break this curse of dimensionality?
- Is high dimensionality an issue in other related fields?

Where is the problem?

- Let's look at RWM, assume $\mathbf{x}_k = \mathbf{0}$.
  - Propose $\mathbf{x}' = \sigma\xi_k$
  - Accept with probability $\exp(-\frac{1}{2}\sigma^2\|\xi_k\|^2) \sim \exp(-\frac{1}{2}\sigma^2 d)$.
- If we keep $\sigma = 1$, "never" accept if $d > 20$.
- If we want acceptance at a constant rate, $\sigma = d^{-\frac{1}{2}}$.
  But then $\mathbf{x}' = \mathbf{x}_k + \sigma\xi_k$ is highly correlated with $\mathbf{x}_k$.

- Similar for MALA, $\sigma = d^{-1/3}$. Hamiltonian MCMC, $\sigma = d^{-1/4}$.
- Is it possible to break this curse of dimensionality?
- Is high dimensionality an issue in other related fields?

From EnKF to MCMC?

- Localization for EnKF:
  - Update only in small local blocks.
  - Works when covariance have local structure.

  How to apply localization to MCMC?
- How to update $\mathbf{x}^i$ component by component?
- Gibbs sampling implements this idea exactly!
  - Write $\mathbf{x}^i = [\mathbf{x}_1^i, \mathbf{x}_2^i, \cdots, \mathbf{x}_m^i]$.
  - $\mathbf{x}_k^i$ can be of dimension q, then $d = qm$.
  - Generate $\mathbf{x}_1^{i+1} \sim p(\mathbf{x}_1 | \mathbf{x}_2^i, \mathbf{x}_3^i \cdots, \mathbf{x}_m^i)$.
  - Generate $\mathbf{x}_2^{i+1} \sim p(\mathbf{x}_2 | \mathbf{x}_1^{i+1}, \mathbf{x}_3^i \cdots, \mathbf{x}_m^i)$.
  - $\cdots$
  - Generate $\mathbf{x}_m^{i+1} \sim p(\mathbf{x}_m | \mathbf{x}_1^{i+1}, \mathbf{x}_2^{i+1} \cdots, \mathbf{x}_{m-1}^{i+1})$.

From EnKF to MCMC?

- Localization for EnKF:
  - Update only in small local blocks.
  - Works when covariance have local structure.

  How to apply localization to MCMC?
- How to update $\mathbf{x}^i$ component by component?

- Gibbs sampling implements this idea exactly!
  - Write $\mathbf{x}^i = [\mathbf{x}^i_1, \mathbf{x}^i_2, \cdots, \mathbf{x}^i_m]$.
  - $\mathbf{x}^i_k$ can be of dimension q, then $d = qm$.
  - Generate $\mathbf{x}^{i+1}_1 \sim p(\mathbf{x}_1 | \mathbf{x}^i_2, \mathbf{x}^i_3 \cdots, \mathbf{x}^i_m)$.
  - Generate $\mathbf{x}^{i+1}_2 \sim p(\mathbf{x}_2 | \mathbf{x}^{i+1}_1, \mathbf{x}^i_3 \cdots, \mathbf{x}^i_m)$.
  - $\cdots$
  - Generate $\mathbf{x}^{i+1}_m \sim p(\mathbf{x}_m | \mathbf{x}^{i+1}_1, \mathbf{x}^{i+1}_2 \cdots, \mathbf{x}^{i+1}_{m-1})$.

How efficient is Gibbs?

First just test with $p = \mathcal{N}(0, I_n)$

- Generate $\mathbf{x}_1^{i+1} \sim p(\mathbf{x}_1|\mathbf{x}_2^i, \mathbf{x}_3^i \cdots, \mathbf{x}_m^i) = \mathcal{N}(0, I_q)$.
- Generate $\mathbf{x}_2^{i+1} \sim p(\mathbf{x}_2|\mathbf{x}_1^{i+1}, \mathbf{x}_3^i \cdots, \mathbf{x}_m^i) = \mathcal{N}(0, I_q)$.
- $\cdots$
- Generate $\mathbf{x}_m^{i+1} \sim p(\mathbf{x}_m|\mathbf{x}_1^{i+1}, \mathbf{x}_2^{i+1} \cdots, \mathbf{x}_{m-1}^{i+1}) = \mathcal{N}(0, I_q)$.

Gibbs naturally exploits the component independence.
It works efficiently against the dimension.
How about component with sparse/local independence?

How efficient is Gibbs?

First just test with $p = \mathcal{N}(0, I_n)$

- Generate $\mathbf{x}_1^{i+1} \sim p(\mathbf{x}_1 | \mathbf{x}_2^i, \mathbf{x}_3^i \cdots, \mathbf{x}_m^i) = \mathcal{N}(0, I_q)$.
- Generate $\mathbf{x}_2^{i+1} \sim p(\mathbf{x}_2 | \mathbf{x}_1^{i+1}, \mathbf{x}_3^i \cdots, \mathbf{x}_m^i) = \mathcal{N}(0, I_q)$.
- $\cdots$
- Generate $\mathbf{x}_m^{i+1} \sim p(\mathbf{x}_m | \mathbf{x}_1^{i+1}, \mathbf{x}_2^{i+1} \cdots, \mathbf{x}_{m-1}^{i+1}) = \mathcal{N}(0, I_q)$.

Gibbs naturally exploits the component independence.
It works efficiently against the dimension.
How about component with sparse/local independence?

Block tridiagonal matrix

- Local covariance matrix $\mathbf{C}$:
  $[\mathbf{C}]_{i,j}$ decays to zero quickly when $|i-j|$ becomes large.
- Localized covariance matrix $\mathbf{C}$:
  $[\mathbf{C}]_{i,j} = 0$ when $|i-j| > L$. $\mathbf{C}$ has a bandwidth $2L$.
- We will see "local" is a perturbation of "localized"
- We can choose $q = L$ in $\mathbf{x}^i = [\mathbf{x}^i_1, \mathbf{x}^i_2, \cdots, \mathbf{x}^i_m]$,
  Then $\mathbf{C}$ is block tridiagonal.

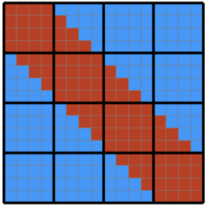

Localized covariance

Theorem (Morzfeld, T., Marzouk)

Apply Gibbs sampler with block-size q to $p = \mathcal{N}(\mathbf{m}, \mathbf{C})$. Suppose $\mathbf{C}$ is q-block-tridiagonal. Then the distribution of $\mathbf{x}^k$ converges to p geometrically fast in all coordinates, and we can couple $\mathbf{x}^k$ and a sample $\mathbf{z} \sim \mathcal{N}(\mathbf{m}, \mathbf{C})$ such that

$$\mathbb{E}\|\mathbf{C}^{-1/2}(\mathbf{x}^k - \mathbf{z})\|^2 \leq \beta^k d(1 + \|\mathbf{C}^{-1/2}(\mathbf{x}^0 - \mathbf{m})\|^2),$$

where

$$\beta \leq \frac{2(1-\mathcal{C}^{-1})^2\mathcal{C}^4}{1+2(1-\mathcal{C}^{-1})^2\mathcal{C}^4},$$

with $\mathcal{C}$ being the condition number of $\mathbf{C}$.

Localized covariance+mild condition$\Rightarrow$ dimension free covergence.

Localized precision

Theorem (Morzfeld, T., Marzouk)

Apply Gibbs sampler with block-size q to $p = \mathcal{N}(\mathbf{m}, \mathbf{C})$. Suppose $\Sigma = \mathbf{C}^{-1}$ is q-block-tridiagonal. Then the distribution of $\mathbf{x}^k$ converges to p geometrically fast in all coordinates, and we can couple $\mathbf{x}^k$ and a sample $\mathbf{z} \sim \mathcal{N}(\mathbf{m}, \mathbf{C})$ such that

$$\mathbb{E}\|\mathbf{C}^{-1/2}(\mathbf{x}^k - \mathbf{z})\|^2 \leq \beta^k d(1 + \|\mathbf{C}^{-1/2}(\mathbf{x}^0 - \mathbf{m})\|^2),$$

where

$$\beta \leq \frac{\mathcal{C}(1-\mathcal{C}^{-1})^2}{1+\mathcal{C}(1-\mathcal{C}^{-1})^2},$$

with $\mathcal{C}$ being the condition number of $\mathbf{C}$.

Localized precision+mild condition$\Rightarrow$ dimension free covergence.

Covariance v.s. Precision

Why both localized covariance and precision?

- A lemma in Bickle & Lindner 2012.
- Localized covariance+mild condition $\Rightarrow$ local precision.
- Localized precision+mild condition $\Rightarrow$ local covariance.
- We will see "local" is a perturbation of "localized"

For computation of Gibbs sampler, localized precision is superior:

$$\mathbf{x}_j^{k+1} \sim \mathcal{N}\left(\mathbf{m}_j - \sum_{i<j} \mathbf{\Omega}_{j,j}^{-1}\mathbf{\Omega}_{j,i}(\mathbf{x}_i^{k+1} - \mathbf{m}_i) - \sum_{i>j} \mathbf{\Omega}_{j,j}^{-1}\mathbf{\Omega}_{j,i}(\mathbf{x}_i^{k} - \mathbf{m}_i), \mathbf{\Omega}_{j,j}^{-1}\right).$$

When Ω is sparse, meaning fast computation.

Why both localized covariance and precision?

- A lemma in Bickle & Lindner 2012.
- Localized covariance+mild condition ⇒ local precision.
- Localized precision+mild condition ⇒ local covariance.
- We will see "local" is a perturbation of "localized"

For computation of Gibbs sampler, localized precision is superior:

$$\mathbf{x}_j^{k+1} \sim \mathcal{N}\left(\mathbf{m}_j - \sum_{i<j} \mathbf{\Omega}_{j,j}^{-1}\mathbf{\Omega}_{j,i}(\mathbf{x}_i^{k+1} - \mathbf{m}_i) - \sum_{i>j} \mathbf{\Omega}_{j,j}^{-1}\mathbf{\Omega}_{j,i}(\mathbf{x}_i^{k} - \mathbf{m}_i), \mathbf{\Omega}_{j,j}^{-1}\right).$$

When Ω is sparse, meaning fast computation.

Generalization

- Gibbs works for Gaussian sampling, with localized covariance or precision.
- How about Bayesian inverse problem?

$$\mathbf{y} = h(\mathbf{x}) + \mathbf{v}, \quad \mathbf{v} \sim \mathcal{N}(0, R), \mathbf{x} \sim p_0 = \mathcal{N}(\mathbf{m}, \mathbf{C}).$$

  If h is linear, p is also Gaussian, Gibbs is directly applicable.
- What to do when $\mathbf{C}$ e.t.c. are not localized but local?

Metropolis within Gibbs

- Add in Metropolis steps to incorporate information
  - Generate $\mathbf{x}_1' \sim p_0(\mathbf{x}_1|\mathbf{x}_2^i, \mathbf{x}_3^i \cdots, \mathbf{x}_m^i)$
  - Accept as $\mathbf{x}_1^{i+1}$ with $\alpha_1(\mathbf{x}_1^i, \mathbf{x}_1', \mathbf{x}_{2:m}^i)$

$$\alpha_1(\mathbf{x}_1^i, \mathbf{x}_1', \mathbf{x}_{2:m}^i) = \min\left\{1, \frac{\exp(-\frac{1}{2}\|\mathbf{y} - h(\mathbf{x}')\|_R^2)}{\exp(-\frac{1}{2}\|\mathbf{y} - h(\mathbf{x}^i)\|_R^2)}\right\},$$

  where $\mathbf{x}' = (\mathbf{x}_1', \mathbf{x}_{2:m}^i)$.
  - Repeat for all $2, ..., m$ blocks
- When h has a dimension free Lipschitz constant, $\|\mathbf{y} - h(\mathbf{x}')\|_R^2 - \|\mathbf{y} - h(\mathbf{x}^i)\|_R^2$ is independent of d.
- Dimension independent acceptance rate.
- Should have fast convergence, though proof is unclear.

NUS National University of Singapore

Often $\boldsymbol{\Omega}$ and h are local

- $[\boldsymbol{\Omega}]_{i,j}$ decays to zero quickly when $|i-j|$ increases.
- $[h(\mathbf{x})]_j$ depends significantly only over a few $\mathbf{x}_i$.

Fast sparse computation is possible with *localized* parameters

- $[\boldsymbol{\Omega}]_{i,j}$ decays to zero quickly when $|i-j|$ increases.
- $[h(\mathbf{x})]_j$ depends significantly only over a few $\mathbf{x}_i$.

Localization: truncate the near zero terms, $\Omega \to \Omega^L$, $h \to h^L$.

We call MwG with localization as l-MwG.

Theorem (Morzfeld, T., Marzouk 2018)

The perturbation to the inverse problem is of order

$$\max\left\{\|\Omega-\Omega^L\|_1, \sqrt{\|(H-H^L)(H-H^L)^T\|_1}\right\}.$$

$\|A\|_1 = \max_{1\le j\le d}\sum_{i=1}^d |A_{i,j}|$.

Comparison with function space MCMC

Function space MCMC:

- Discretization refines, domain const.
- Number of obs. const.
- Effective dimension const.
- Low-rank priors.
- Low-rank prior to posterior update.

Solved by dimension reduction.

MCMC for local problems:

- Domain size increases, discretization is const.
- Number of obs. increases.
- Effective dimension increases.
- High-rank, sparse priors.
- High-rank prior to posterior update.

Solved by localization.

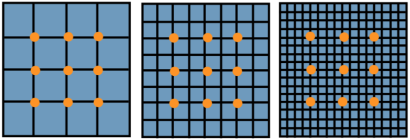

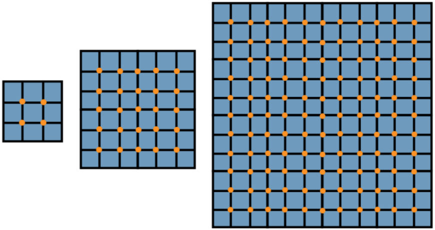

Example I: image deblurring

- Truth $\mathbf{x} \sim \mathcal{N}(0, \delta^{-1} L^{-2})$, L is Laplacian.
- Defocus obs: $\mathbf{y} = A\mathbf{x} + \eta, \eta \sim \mathcal{N}(0, \lambda^{-1}\mathbf{I})$.
- Dimension is large $O(10^4)$.

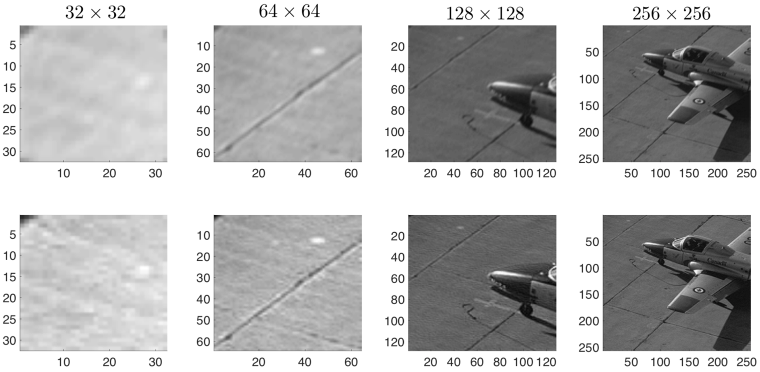

Example I: image deblurring

$\Omega = \lambda A^T A + \delta L x$

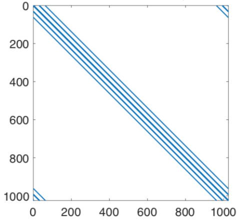

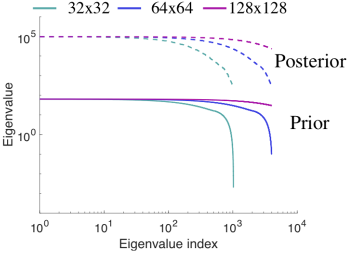

Precision is sparse. Effective dimension is large.

Image size	32 x 32	64 x 64	128 x 128	256 x 256
Dimension	1,024	4,096	16,348	16,536
Eff. Dimension	$4.8 \cdot 10^8$	$7.4 \cdot 10^9$	$1.2 \cdot 10^{11}$	-
IACT (Gibbs)	2.92	2.97	1.74	1.11
Blocksize (Gibbs)	16	16	32	64

Example II: Lorenz 96 inverse

- Truth $\mathbf{x}_0 \sim p_0$, p_0 is Gaussian Climatology.
- $\Psi_t : \mathbf{x}_0 \mapsto \mathbf{x}_t : d\mathbf{x}_i = (\mathbf{x}_{i+1} - \mathbf{x}_{i-2})\mathbf{x}_{i-1} - \mathbf{x}_i + 8$
- Observe every other $\mathbf{x}_t$, $\mathbf{y} = H(\Psi_t(\mathbf{x}_0)) + \xi$.

	MALA	pCN	l-MwG-B2	l-MwG-B4	l-MwG-B8
$n = 40$	686	1051	55	60	266
$n = 400$	3,153	3,257	43	81	257

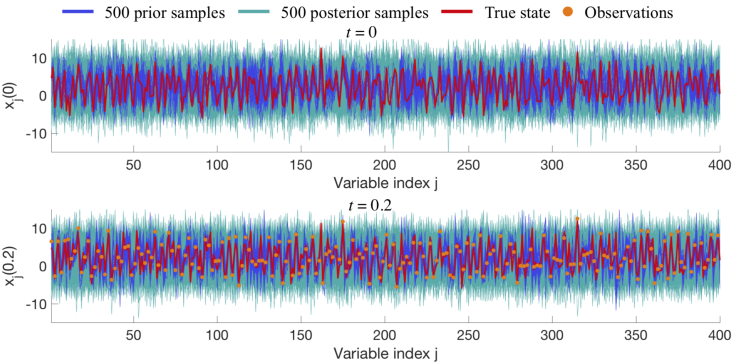

Summary

- Most MCMC suffers from high dimensionality due to degenerate acceptance.
- Localization technique in EnKF significantly reduces sampling complexity.
- Gibbs sampler has dimension free convergence sampling local Gaussian dist.
- Local proposals help MCMC has dimension free acceptance.
- Different setting comparing with functional space MCMC.
- Successful applications with image deblurring and Lorenz inverse problem.

Reference

- Localization for MCMC: sampling high-dimensional posterior distributions with local structure. arXiv:1710.07747
- Performance analysis of local ensemble Kalman filter. to appear on J. Nonlinear Science.

Links and slides can be found at www.math.nus.edu.sg/∼mattxin.

Thank you!