Understanding MAP estimators for high dimensional non-linear filtering based on concentration inequalities

Daniel Paulin

University of Oxford, Department of Statistics

Joint work with Alexandros Beskos, Dan Crisan and Ajay Jasra.

Cargèse, May 18, 2018

Outline

Introduction

Theory

Idea of proofs

Flow-dependent 4D-Var

References

High dimensional filtering - Introduction

- ▷ We study the filtering problem for partially observed high dimensional deterministic dynamical systems.
- ▷ Such models are widely used in weather forecasting and engineering.
- ▷ Consider an ODE of the form $\frac{d\boldsymbol{u}}{dt} = F(\boldsymbol{u})$, where $\boldsymbol{u} : \mathbb{R}^+ \to \mathbb{R}^d$.
- ▷ This is usually discretisation of a PDE.
- ▷ In the case of non-linear F, such ODEs often exhibit chaotic behaviour.

▷ We assume F to be quadratic, which leads to ODEs of the form

$$\frac{d\boldsymbol{u}}{dt} = F(\boldsymbol{u}) = -\boldsymbol{A}\boldsymbol{u} - \boldsymbol{B}(\boldsymbol{u}, \boldsymbol{u}) + \boldsymbol{f}, \tag{1}$$

▷ where $\boldsymbol{u} : \mathbb{R}^+ \to \mathbb{R}^d$ is a dynamical system,

▷ $\boldsymbol{A}$ is a linear operator ($d \times d$ matrix)

▷ $\boldsymbol{B}$ is a bilinear form ($d \times d \times d$ array) causing nonlinearity,

▷ $\boldsymbol{f} \in \mathbb{R}^d$ is a constant vector, the so-called *forcing term*.

▷ We make the following assumption.

Assumption 1 (Trapping ball assumption)

Assume that for some $R > 0$*, we have* $\langle F(\boldsymbol{v}), \boldsymbol{v} \rangle < 0$ *for every* $\boldsymbol{v} \in \mathbb{R}^d$ *with* $\|\boldsymbol{v}\| = R$*.*

▷ Let $\mathcal{B}_R := \{\boldsymbol{v} \in \mathbb{R}^d : \|\boldsymbol{v}\| \leq R\}$. Under Assumption 1, for every initial point $\boldsymbol{v} \in \mathcal{B}_R$, a unique solution of (1), denoted by $\boldsymbol{v}(t)$, exists for every $t \geq 0$. In particular, $\boldsymbol{v}(0) = \boldsymbol{v}$, and $\|\boldsymbol{v}(t)\| \leq R$.

- ▷ The observations happen at times $t_i = ih$, for $0 \leq i \leq k$, $h > 0$.
- ▷ h is the assimilation *time step*
- ▷ $T = t_k$ is the size of the *observation window*.
- ▷ We assume that the noisy observations are of dimension d_o,

$$\boldsymbol{Y}_i := \boldsymbol{H}\boldsymbol{u}(t_i) + \boldsymbol{Z}_i,$$

- ▷ where $\boldsymbol{H} : \mathbb{R}^d \to \mathbb{R}^{d_o}$ is a linear observation operator, and
- ▷ $\boldsymbol{Z}_i$ are i.i.d. random vectors with distribution $N(0, \sigma_Z^2 \boldsymbol{I}_{d_o})$.

▷ **Problem**: how to estimate $\boldsymbol{u}(0)$ (smoothing) or $\boldsymbol{u}(T)$ (filtering) given observations $\boldsymbol{Y}_{0:k}$.

▷ The partial observations and the non-linearity makes this hard.

▷ A theoretical solution is to take some prior $\boldsymbol{u} \sim q$, and set

$$\overline{\boldsymbol{u}}^{\mathrm{sm}} := \mathbb{E}(\boldsymbol{u}(0)|\boldsymbol{Y}_{0:k}), \quad \overline{\boldsymbol{u}}^{\mathrm{fi}} := \mathbb{E}(\boldsymbol{u}(T)|\boldsymbol{Y}_{0:k}).$$

▷ These are optimal in MSE, i.e. for any $\hat{\boldsymbol{u}}^{\mathrm{sm}}(\boldsymbol{Y}_{0:k})$, $\hat{\boldsymbol{u}}^{\mathrm{fi}}(\boldsymbol{Y}_{0:k})$,

$$\mathbb{E}(\|\overline{\boldsymbol{u}}^{\mathrm{sm}} - \boldsymbol{u}(0)\|^2) \le \mathbb{E}(\|\hat{\boldsymbol{u}}^{\mathrm{sm}}(\boldsymbol{Y}_{0:k}) - \boldsymbol{u}(0)\|^2),$$

$$\mathbb{E}(\|\overline{\boldsymbol{u}}^{\mathrm{fi}} - \boldsymbol{u}(T)\|^2) \le \mathbb{E}(\|\hat{\boldsymbol{u}}^{\mathrm{fi}}(\boldsymbol{Y}_{0:k}) - \boldsymbol{u}(T)\|^2).$$

▷ However, computing them in high dimensions is difficult.

Filtering methods

▷ In the case of linear dynamics, and Gaussian prior, the filtering distribution is Gaussian, so we can use the Kalman filter

▷ There are several variants of the Kalman filter for non-linear dynamics: Extended Kalman filter, Ensemble Kalman filter, etc.

▷ There are also variational methods such as 3D-Var and 4D-Var.

▷ Their consistency for non-linear dynamics is unknown in general.

▷ [Sanz-Alonso and Stuart, 2015] has shown consistency of the 3D-Var filter for some non-linear ODEs, strong assumptions on d_o.

The 4D-Var method

- ▷ The 4D-Var method (see [Talagrand and Courtier, 1987]) consists of finding the MAP for the smoother, and propagating it forward to the filter
- ▷ **Key idea:** the gradient of the log-likelihood can be computed at $O(d)$ cost by adjoint equation, total cost of finding MAP is $O(d)$
- ▷ This allows weather models with dimension $d = 10^9$ and higher
- ▷ It is the most frequently used data assimilation method in weather forecasting centres. First implemented in ECMWF in 1997.

$\triangleright$ In practice, the conditioning of the Hessian can be bad due to the partiality of the observations

$\triangleright$ This makes gradient descent inefficient

$\triangleright$ Newton and Gauss-Newton methods with linear solvers based on Preconditioned Conjugate Gradient (PCG) work much better

$\triangleright$ **Key point:** although the Hessian $\mathcal{H}$ cannot be stored, matrix-vector products $\mathcal{H}\boldsymbol{v}$ can be computed with $O(d)$ cost

Assumption on the observations

Assumption 2 (Recoverability from partial observations)

Suppose that $\|\boldsymbol{u}\| < R$*, and there is an index* $j \in \mathbb{N}$ *such that the system of equations in* $\boldsymbol{v}$ *defined as*

$$\boldsymbol{H} \left. \frac{d^i \boldsymbol{u}(t)}{dt^i} \right|_{t=0} = \boldsymbol{H} \left. \frac{d^i \boldsymbol{v}(t)}{dt^i} \right|_{t=0} \quad \text{for every } 0 \le i \le j$$

has a unique solution $\boldsymbol{v} := \boldsymbol{u}$ *in* $\mathcal{B}_R$*, and*

$$\mathrm{span}\left\{ \nabla_{\boldsymbol{u}} \left(\boldsymbol{H} \left. \frac{d^i \boldsymbol{u}(t)}{dt^i} \right|_{t=0} \right)_l : 0 \le i \le j, 1 \le l \le d_o \right\} = \mathbb{R}^d.$$

▷ Assumption 2 implies in particular that there is a constant $c(\boldsymbol{u}, T) > 0$ depending on $\boldsymbol{u}$ but independent of h and σ_Z such that for every $\boldsymbol{v} \in \mathcal{B}_R$,

$$\sum_{l=0}^{k} \|\boldsymbol{H}\boldsymbol{v}(t_l) - \boldsymbol{H}\boldsymbol{u}(t_l)\|^2 \geq \frac{c(\boldsymbol{u}, T)}{h} \|\boldsymbol{v} - \boldsymbol{u}\|^2.$$

▷ [Paulin et al., 2018] has shown several consistency results for the 4D-Var under Assumptions 1 and 2.

Gaussian approximation

Theorem 1 (Gaussian approximation of smoother)

For quadratic dynamics, under some mild assumptions on $\boldsymbol{u}$, q *and* h, *there are some constants* $C_{\mathrm{TV}}^{(1)}$, $C_{\mathrm{TV}}^{(2)}$, $C_{\mathrm{W}}^{(1)}$, $C_{\mathrm{W}}^{(2)}$ *depending on* $\boldsymbol{u}$ *and* T *and a multivariate normal distribution* $\mu_{\mathcal{G}}^{\mathrm{sm}}(\cdot|\boldsymbol{Y}_{0:k})$ *such that*

$$\mathbb{P}\Big[d_{\mathrm{TV}}(\mu^{\mathrm{sm}}(\cdot|\boldsymbol{Y}_{0:k}),\mu_{\mathcal{G}}^{\mathrm{sm}}(\cdot|\boldsymbol{Y}_{0:k}))\le (C_{\mathrm{TV}}^{(1)}+C_{\mathrm{TV}}^{(2)}\log^2(1/\varepsilon))\sigma_Z\sqrt{h}$$
$$\& \, d_{\mathrm{W}}(\mu^{\mathrm{sm}}(\cdot|\boldsymbol{Y}_{0:k}),\mu_{\mathcal{G}}^{\mathrm{sm}}(\cdot|\boldsymbol{Y}_{0:k}))\le (C_{\mathrm{W}}^{(1)}+C_{\mathrm{W}}^{(2)}\log^2(1/\varepsilon))\sigma_Z^2 h\Big|\boldsymbol{u}\Big]$$
$$\ge 1-\varepsilon \text{ for every } \varepsilon>0.$$

Theorem 2 (Gaussian approximation of filter)

For quadratic dynamics, under some mild assumptions on $\boldsymbol{u}$*,* q *and* h*, there are some constants* $D_{\mathrm{TV}}^{(1)}$*,* $D_{\mathrm{TV}}^{(2)}$*,* $D_{\mathrm{W}}^{(1)}$*,* $D_{\mathrm{W}}^{(2)}$ *depending on* $\boldsymbol{u}$ *and* T *and a multivariate normal distribution* $\mu_{\mathcal{G}}^{\mathrm{fi}}(\cdot|\boldsymbol{Y}_{0:k})$ *such that*

$$\mathbb{P}\Big[d_{\mathrm{TV}}(\mu^{\mathrm{fi}}(\cdot|\boldsymbol{Y}_{0:k}), \mu_{\mathcal{G}}^{\mathrm{fi}}(\cdot|\boldsymbol{Y}_{0:k})) \leq (D_{\mathrm{TV}}^{(1)} + D_{\mathrm{TV}}^{(2)} \log^2(1/\varepsilon))\sigma_Z\sqrt{h}$$

$$\& \, d_{\mathrm{W}}(\mu^{\mathrm{fi}}(\cdot|\boldsymbol{Y}_{0:k}), \mu_{\mathcal{G}}^{\mathrm{fi}}(\cdot|\boldsymbol{Y}_{0:k})) \leq (D_{\mathrm{W}}^{(1)} + D_{\mathrm{W}}^{(2)} \log^2(1/\varepsilon))\sigma_Z^2 h \Big| \boldsymbol{u}\Big]$$

$$\geq 1 - \varepsilon \text{ for every } \varepsilon > 0.$$

Asymptotic optimality of MAP estimators

Theorem 3 (Asymptotic optimality of MAP for smoother)

For quadratic dynamics, under some mild assumptions on $\boldsymbol{u}$, q and h, there are constants $S_{\max}^{\mathrm{sm}} > 0$, $C_{\mathrm{MAP}}^{\mathrm{sm}}$ depending on $\boldsymbol{u}$ and T such that for $\sigma_Z\sqrt{h} \le S_{\max}^{\mathrm{sm}}$, we have

$$\frac{\mathbb{E}\left[\|\hat{\boldsymbol{u}}_{\mathrm{MAP}}^{\mathrm{sm}} - \boldsymbol{u}\|^2 | \boldsymbol{u}\right]}{\mathbb{E}\left[\|\overline{\boldsymbol{u}}^{\mathrm{sm}} - \boldsymbol{u}\|^2 | \boldsymbol{u}\right]} \le 1 + C_{\mathrm{MAP}}^{\mathrm{sm}} \cdot \sigma_Z\sqrt{h}.$$

Theorem 4 (Asymptotic optimality of MAP for filter)

Let $\hat{\boldsymbol{u}}^{\mathrm{fi}} := \hat{\boldsymbol{u}}^{\mathrm{sm}}_{\mathrm{MAP}}(T)$. *For quadratic dynamics, under some mild assumptions on* $\boldsymbol{u}$, q *and* h, *there are constants* $S^{\mathrm{fi}}_{\max} > 0$, $C^{\mathrm{fi}}_{\mathrm{MAP}}$ *depending on* $\boldsymbol{u}$ *and* T *such that for* $\sigma_Z\sqrt{h} \le S^{\mathrm{fi}}_{\max}$, *we have*

$$\frac{\mathbb{E}\left[\|\hat{\boldsymbol{u}}^{\mathrm{fi}} - \boldsymbol{u}(T)\|^2 | \boldsymbol{u}\right]}{\mathbb{E}\left[\|\overline{\boldsymbol{u}}^{\mathrm{fi}} - \boldsymbol{u}(T)\|^2 | \boldsymbol{u}\right]} \le 1 + C^{\mathrm{fi}}_{\mathrm{MAP}} \cdot \sigma_Z\sqrt{h}.$$

Newton's method

▷ The negative log-likelihood (excluding the normalising constant) is of the form

$$J(\boldsymbol{v}) := -\log(q(\boldsymbol{v})) + \frac{1}{2\sigma_Z^2}\sum_{i=0}^{k}\|\boldsymbol{Y}_i - \boldsymbol{H}\boldsymbol{v}(t_i)\|^2. \qquad (2)$$

▷ Newton's method is defined recursively as

$$\boldsymbol{x}_{i+1} := \boldsymbol{x}_i - (\nabla^2 J(\boldsymbol{x}_i))^{-1} \cdot \nabla J(\boldsymbol{x}_i) \text{ for } i \in \mathbb{N}. \qquad (3)$$

Theorem 5 (Convergence of Newton's method to MAP)

Under some mild assumptions of $\boldsymbol{u}$, q, h, *for every* $0 < \varepsilon \leq 1$, *there are finite constants* $S^{\mathrm{sm}}_{\max}(\boldsymbol{u}, T, \varepsilon)$, $N^{\mathrm{sm}}(\boldsymbol{u}, T)$ *and* $D^{\mathrm{sm}}_{\max}(\boldsymbol{u}, T) \in (0, N^{\mathrm{sm}}(\boldsymbol{u}, T)]$ *such that for* $\sigma_Z \sqrt{h} \leq S^{\mathrm{sm}}_{\max}(\boldsymbol{u}, T, \varepsilon)$, *if the initial point* $\boldsymbol{x}_0 \in \mathcal{B}_R$ *satisfies* $\|\boldsymbol{x}_0 - \boldsymbol{u}\| < D^{\mathrm{sm}}_{\max}(\boldsymbol{u}, T)$, *then*

$$\mathbb{P}\Bigg(\boldsymbol{x}_i \text{ is well defined and } \|\boldsymbol{x}_i - \hat{\boldsymbol{u}}^{\mathrm{sm}}_{\mathrm{MAP}}\| \leq N^{\mathrm{sm}}(\boldsymbol{u}, T)\left(\frac{\|\boldsymbol{x}_0 - \boldsymbol{u}\|}{N^{\mathrm{sm}}(\boldsymbol{u}, T)}\right)^{2^i} \text{ for every } i \in \mathbb{N} \,\Bigg|\, \boldsymbol{u}\Bigg) \geq 1 - \varepsilon. \tag{4}$$

Initial estimator

- ▷ For Newton's method to work, we need to have an initial estimator $\boldsymbol{x}_0$ satisfying that $\|\boldsymbol{x}_0 - \boldsymbol{u}\| < D^{\mathrm{sm}}_{\max}(\boldsymbol{u}, T)$
- ▷ Suppose that $\hat{\Phi}^{(0)}, \dots, \hat{\Phi}^{(j)}$ are estimators of $\boldsymbol{H}\boldsymbol{u}, \dots, \boldsymbol{H}\left.\frac{d^j \boldsymbol{u}(t)}{dt^j}\right|_{t=0}$ based on $\boldsymbol{Y}_{0:k}$.
- ▷ If there is a function $G : (\mathbb{R}^{d_o})^{j+1} \to \mathbb{R}^d$ independent of $\boldsymbol{u}$ such that $G\left(\boldsymbol{H}\boldsymbol{u}, \dots, \boldsymbol{H}\left.\frac{d^j \boldsymbol{u}(t)}{dt^j}\right|_{t=0}\right) = \boldsymbol{u}$, then we can use $G(\hat{\Phi}^{(0)}, \dots, \hat{\Phi}^{(j)})$ as an initial estimator.

Lorenz 96' model

▷ The Lorenz 96' model is a d dimensional chaotic ODE of the form

$$\frac{d}{dt}u_i = -u_{i-1}u_{i-2} + u_{i-1}u_{i+1} - u_i + f. \tag{5}$$

▷ The indices are understood modulo d, and usually $f := 8$.

▷ This is of the form (1), and we also have $\langle \boldsymbol{B}(\boldsymbol{v}, \boldsymbol{v}), \boldsymbol{v} \rangle = 0$ for every $\boldsymbol{v} \in \mathbb{R}^d$.

▷ Therefore the trapping ball assumption holds for $R > |f| \cdot \sqrt{d}$.

Choice of F

- We observe either
  1. coordinates $1, 2, 3,\ 7, 8, 9,\ \ldots$.
  2. coordinates $1, 2, 3$.
- By rearrangement of the ODE, we have

$$u_i = \left(\frac{du_{i-1}}{dt} - f + u_{i-1} + u_{i-2}u_{i-3} \right) \Big/ u_{i-2}, \quad \text{and}$$
$$u_i = \left(f - \frac{du_{i+2}}{dt} - u_{i+2} + u_{i+1}u_{i+3} \right) \Big/ u_{i+1}.$$

- The un-observed coordinates are expressed in terms of the observed coordinates and derivatives of order 1 or $\lceil (d-3)/3 \rceil$.

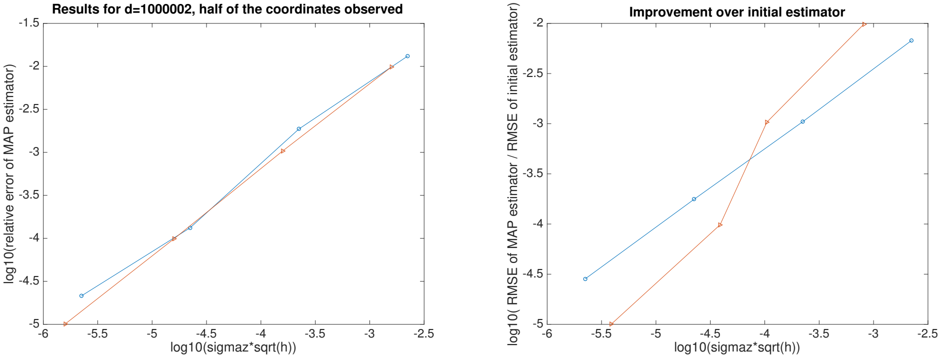

Figure: Dependence of RMSE of estimator on σ_Z and h for $d = 1000002$

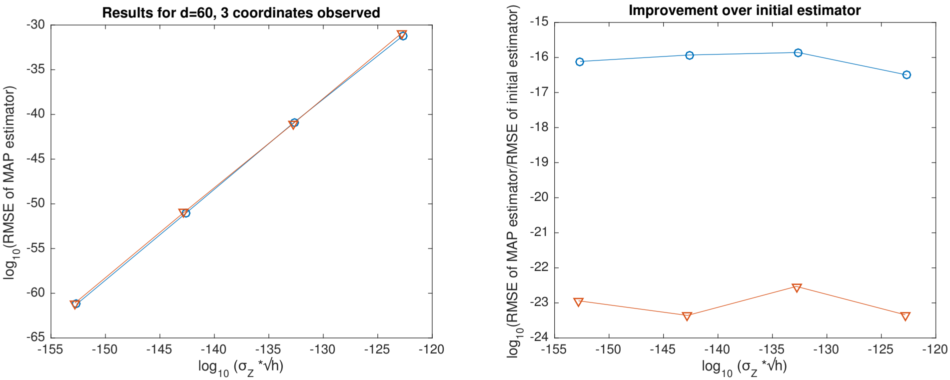

Figure: Dependence of RMSE of estimator on σ_Z and h for $d = 60$, 3 coordinates observed

Idea of proofs

- Denote $\Phi_t(\boldsymbol{v}) := \boldsymbol{H}\boldsymbol{v}(t)$.
- The prior-free negative log-likelihoods of the smoother and its Gaussian approximation are

$$I^{\mathrm{sm}}(\boldsymbol{v}) := \sum_{i=0}^{k} \left(\|\Phi_{t_i}(\boldsymbol{v}) - \Phi_{t_i}(\boldsymbol{u})\|^2 + 2\left\langle \Phi_{t_i}(\boldsymbol{v}) - \Phi_{t_i}(\boldsymbol{u}), \boldsymbol{Z}_i \right\rangle \right), \text{ and}$$

$$I^{\mathrm{sm}}_{\mathcal{G}}(\boldsymbol{v}) := (\boldsymbol{v} - \boldsymbol{u})^T \boldsymbol{A}_k (\boldsymbol{v} - \boldsymbol{u}) + 2\left\langle \boldsymbol{v} - \boldsymbol{u}, \boldsymbol{B}_k \right\rangle, \text{ where}$$

$$\boldsymbol{A}_k := \sum_{i=0}^{k} \left((\boldsymbol{J}\Phi_{t_i}(\boldsymbol{u}))^T \boldsymbol{J}\Phi_{t_i}(\boldsymbol{u}) + \boldsymbol{J}^2\Phi_{t_i}(\boldsymbol{u})[\cdot, \cdot, \boldsymbol{Z}_i] \right),$$

$$\boldsymbol{B}_k := \sum_{i=0}^{k} (\boldsymbol{J}\Phi_{t_i}(\boldsymbol{u}))^T \cdot \boldsymbol{Z}_i.$$

▷ Using matrix concentration inequalities ([Tropp, 2015]), it follows that $\boldsymbol{A}_k$ is pos. def. with high probability when $\sigma_Z\sqrt{h}$ is small.

▷ Using concentration inequalities for empirical processes, one can show the following type of results.

Proposition 1 (Bound on the difference $|l^{\mathrm{sm}}(\boldsymbol{v}) - l^{\mathrm{sm}}_{\mathcal{G}}(\boldsymbol{v})|$)

Under mild assumptions on $\boldsymbol{u}$ and h, for any $0 < \varepsilon \leq 1$, $\sigma_Z > 0$,

$$\mathbb{P}\left(|l^{\mathrm{sm}}(\boldsymbol{v}) - l^{\mathrm{sm}}_{\mathcal{G}}(\boldsymbol{v})| \leq \|\boldsymbol{v} - \boldsymbol{u}\|^3 \cdot \frac{C_2(\boldsymbol{u}, T) + C_3(\boldsymbol{u}, T, \varepsilon)\sigma_Z\sqrt{h}}{h} \right.$$

$$\left. \text{for every } \boldsymbol{v} \in \mathcal{B}_R \middle| \boldsymbol{u} \right) \geq 1 - \varepsilon.$$

▷ The next lemma is used in the proof of Proposition 1. It is based of Corollary 13.2 and Theorem 5.8 of [Boucheron et al., 2013].

Lemma 1

For every $l \in \mathbb{N}$*, define the sets*

$$\mathcal{T}_l := \{(r, \boldsymbol{s}_1, \ldots, \boldsymbol{s}_l) \in [0, 2R] \times \mathcal{B}_1^l : \boldsymbol{u} + r\boldsymbol{s}_1 \in \mathcal{B}_R\}, \quad \overline{\mathcal{T}}_l := \mathcal{B}_R \times \mathcal{B}_1^l.$$

For any two points $(r, \boldsymbol{s}_1, \ldots, \boldsymbol{s}_l), (r, \boldsymbol{s}_1', \ldots, \boldsymbol{s}_l') \in \mathcal{T}_l$*, let*

$$d_l((r, \boldsymbol{s}_1, \ldots, \boldsymbol{s}_l), (r, \boldsymbol{s}_1', \ldots, \boldsymbol{s}_l')) := \frac{|r - r'|}{2R} + \sum_{i=0}^{l} \|\boldsymbol{s}_i - \boldsymbol{s}_i'\|. \quad (6)$$

▷ Let $\boldsymbol{Z}_0, \ldots, \boldsymbol{Z}_k$ be i.i.d. d_o dimensional standard normal random vectors,

▷ Let $\varphi_0, \ldots, \varphi_k : \mathcal{T}_l \to \mathbb{R}^{d_o}$ be functions that are L-Lipschitz with respect to the distance d_l on $\mathcal{T}_l$, and satisfy that $\|\varphi_i(r, \boldsymbol{s}_1, \ldots, \boldsymbol{s}_l)\|_\infty \leq M$ for any $0 \leq i \leq k$.

▷ Then $W_l := \sup_{(r, \boldsymbol{s}_1, \ldots, \boldsymbol{s}_l) \in \mathcal{T}_l} \sum_{i=0}^{k} \langle \varphi_i(r, \boldsymbol{s}_1, \ldots, \boldsymbol{s}_l), \boldsymbol{Z}_i \rangle$ satisfies that for any $0 < \varepsilon \leq 1$,

$$\mathbb{P}(W_l \geq C^{(l)}(\boldsymbol{u}, k, \varepsilon)) \leq \varepsilon \text{ for} \tag{7}$$

$$C^{(l)}(\boldsymbol{u}, k, \varepsilon) := 10(l+1)L\sqrt{(k+1)(ld+1)d_o} + \sqrt{2(k+1)Md_o \log\left(\frac{1}{\varepsilon}\right)}.$$

From the definitions, we have

$$I^{\mathrm{sm}}(\boldsymbol{v}) = \sum_{i=0}^{k} \|\Phi_{t_i}(\boldsymbol{v}) - \Phi_{t_i}(\boldsymbol{u})\|^2 + 2\sum_{i=0}^{k} \langle \Phi_{t_i}(\boldsymbol{v}) - \Phi_{t_i}(\boldsymbol{u}), \boldsymbol{Z}_i \rangle \text{ , and}$$

$$I_{\mathcal{G}}^{\mathrm{sm}}(\boldsymbol{v}) = \sum_{i=0}^{k} \|J\Phi_{t_i}(\boldsymbol{u}) \cdot (\boldsymbol{v} - \boldsymbol{u})\|^2 + \sum_{i=0}^{k} J^2\Phi_{t_i}(\boldsymbol{u})[\boldsymbol{v} - \boldsymbol{u}, \boldsymbol{v} - \boldsymbol{u}, \boldsymbol{Z}_i]$$

$$+ 2\sum_{i=0}^{k} \langle J\Phi_{t_i}(\boldsymbol{u}) \cdot (\boldsymbol{v} - \boldsymbol{u}), \boldsymbol{Z}_i \rangle , \tag{8}$$

$$|I^{\mathrm{sm}}(\boldsymbol{v}) - I_{\mathcal{G}}^{\mathrm{sm}}(\boldsymbol{v})| \leq \sum_{i=0}^{k} \left| \|\Phi_{t_i}(\boldsymbol{v}) - \Phi_{t_i}(\boldsymbol{u})\|^2 - \|J\Phi_{t_i}(\boldsymbol{u}) \cdot (\boldsymbol{v} - \boldsymbol{u})\|^2 \right|$$

$$+ 2\left| \sum_{i=0}^{k} \left\langle \Phi_{t_i}(\boldsymbol{v}) - \Phi_{t_i}(\boldsymbol{u}) - J\Phi_{t_i}(\boldsymbol{u}) \cdot (\boldsymbol{v} - \boldsymbol{u}) - \frac{1}{2} J^2\Phi_{t_i}(\boldsymbol{u})[\boldsymbol{v} - \boldsymbol{u}, \boldsymbol{v} - \boldsymbol{u}, \cdot], \boldsymbol{Z}_i \right\rangle \right| . \tag{9}$$

▷ The first term in the right hand side of the above inequality can be upper bounded by $\frac{C_1(\boldsymbol{u},T)}{h}\|\boldsymbol{v}-\boldsymbol{u}\|^3$ for some constant $C_1(\boldsymbol{u},T)$.

▷ For the second term, for $(r,\boldsymbol{s}) \in \mathcal{T}_1$, $r>0$, let

$$\varphi_i(r,\boldsymbol{s}) := \left(\Phi_{t_i}(\boldsymbol{u}+r\boldsymbol{s}) - \Phi_{t_i}(\boldsymbol{u}) - \boldsymbol{J}\Phi_{t_i}(\boldsymbol{u})\cdot \boldsymbol{s}r - \frac{1}{2}\boldsymbol{J}^2\Phi_{t_i}(\boldsymbol{u})[r\boldsymbol{s}, r\boldsymbol{s}, \cdot]\right)/r^3.$$

For $r=0$, this can be continuously extended as

$$\varphi_i(0,\boldsymbol{s}) := \lim_{r\to 0}\varphi_i(r,\boldsymbol{s}) = \frac{1}{6}\boldsymbol{J}^3\Phi_{t_i}(\boldsymbol{u})[\boldsymbol{s},\boldsymbol{s},\boldsymbol{s},\cdot].$$

▷ We define $W_1 := \sup_{(r,\boldsymbol{s})\in\mathcal{T}_1}\sum_{i=0}^{k}\langle\varphi_i(r,\boldsymbol{s}), \boldsymbol{Z}_i\rangle$, and $W_1' := \sup_{(r,\boldsymbol{s})\in\mathcal{T}_1}\sum_{i=0}^{k}\langle-\varphi_i(r,\boldsymbol{s}), \boldsymbol{Z}_i\rangle$.

▷ The second term in (9) is bounded by $2\|\boldsymbol{v} - \boldsymbol{u}\|^3 \max(W_1, W_1')$.

▷ The Lipschitz coefficient of φ_i can be bounded via the partial derivatives, and the claim of Proposition 1 now follows by Lemma 1.

Proposition 2

Suppose that $\Omega \subset \mathbb{R}^d$ *is an open set, and* $g : \Omega \to \mathbb{R}$ *is a 3 times continuously differentiable function satisfying that*

1. g *has a local minimum at a point* $\boldsymbol{x}^* \in \Omega$*,*
2. *there exists a radius* $r^* > 0$ *and constants* $C_H > 0, L_H < \infty$ *such that* $B(\boldsymbol{x}^*, r^*) \subset \Omega$*,* $\nabla^2 g(\boldsymbol{x}) \succeq C_H \cdot \boldsymbol{I}_d$ *for every* $\boldsymbol{x} \in B(\boldsymbol{x}^*, r^*)$*, and* $\nabla^2 g(\boldsymbol{x})$ *is* L_H*-Lipschitz on* $B(\boldsymbol{x}^*, r^*)$*.*

Suppose that $\|\boldsymbol{x}_0 - \boldsymbol{x}^*\| < \min\left(r^*, 2\frac{C_H}{L_H}\right)$. *Then*

$$\boldsymbol{x}_{i+1} := \boldsymbol{x}_i - (\nabla^2 g(\boldsymbol{x}_i))^{-1} \cdot \nabla g(\boldsymbol{x}_i) \textit{ always stay in } B(\boldsymbol{x}^*, r^*), \textit{ and}$$

$$\|\boldsymbol{x}_i - \boldsymbol{x}^*\| \leq \frac{2C_H}{L_H} \cdot \left(\frac{L_H}{2C_H}\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|\right)^{2^i} \textit{ for every } i \in \mathbb{N}.$$

▷ The constants C_H (Hessian lower bound) and L_H (Hessian Lipschitz constant) can be bounded for the log-likelihood of the smoothing distribution using concentration inequalities.

Flow-dependent 4D-Var

- ▷ Due to the chaotic nature of the systems, likelihood is multimodal if T is too large
- ▷ Thus T has to be kept sufficiently short, and previous windows are taken into account by the prior (background) distribution
- ▷ In [Paulin et al., 2017], we have proposed a flow-dependent Gaussian background distributions by propagating forward the current Gaussian approximations via the dynamics from the previous b windows for some $b \geq 1$.

▷ If $Z \sim N(m, P^{-1})$, then for a continuously differentiable function φ, $\varphi(Z)$ is approximately distributed as

$$N\left(\varphi(m), \left[\left((\boldsymbol{J}\varphi(m))^{-1}\right)^T P(\boldsymbol{J}\varphi(m))^{-1}\right]^{-1}\right).$$

▷ For a parameter $b \geq 1$, we first set $P_{-b} := P_{fix}$

▷ Then set $P_{-b+1} = (J_{-b}^{-1})^T(P_{-b} + D_{-b})J_{-b}^{-1}, \dots$

▷ $P_{-k+1} = (J_{-k}^{-1})^T(P_{-k} + D_{-k})J_{-k}^{-1}, \dots$

▷ $P_0 = (J_{-1}^{-1})^T(P_{-1} + D_{-1})J_{-1}^{-1}$ is the flow-dependent precision matrix for the current interval.

The following figure illustrates the definition of the prior in a flow-dependent way:

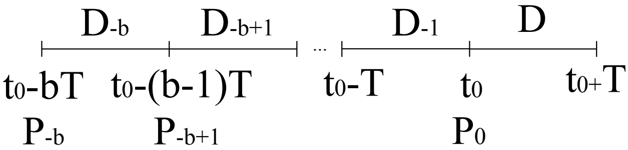

Figure: Definition of the prior precision matrices in a flow-dependent way. D_{-k} corresponds to the Hessian of the negative log-likelihood terms from the data in the interval $[t_0 - kT, t_0 - (k-1)T)$.

▷ The matrix-vector products $\mathcal{H}\boldsymbol{v}$ are at most b times more expensive to compute than for fixed background covariances, still $O(d)$ cost.

Simulations

▷ Consider the shallow-water equations, [Salmon, 2015],

$$\frac{\partial u}{\partial t} = \left(-\frac{\partial u}{\partial y} + f\right) v - \frac{\partial}{\partial x}\left(\frac{1}{2}u^2 + gh\right) + \nu \nabla^2 u - c_b u; \tag{10}$$

$$\frac{\partial v}{\partial t} = -\left(\frac{\partial v}{\partial x} + f\right) u - \frac{\partial}{\partial y}\left(\frac{1}{2}v^2 + gh\right) + \nu \nabla^2 v - c_b v; \tag{11}$$

$$\frac{\partial h}{\partial t} = -\frac{\partial}{\partial x}((h+o)u) - \frac{\partial}{\partial y}((h+o)v). \tag{12}$$

▷ Here, u and v are the velocities in the x and y directions, and h is the the height of the wave, o is the depth of the ocean.

- ▷ The shallow water equations are applied in tsunami modelling.
- ▷ [Saito et al., 2011] estimate the initial distribution of the tsunami waves after the 2011 Japan earthquake.
- ▷ They use data from 17 locations in the ocean, where the wave heights were observed continuously in time.
- ▷ We have used these estimates as our initial condition for the heights, and set the initial velocities to zero (as they are unknown).
- ▷ Using publicly available bathymetry data for o, and the above described initial condition, we have run a simulation of 40 minutes for our model.

▷ We have tested the efficiency of the data assimilation methods also on this simulated dataset, considering a time interval from 10 to 40 minutes.

▷ Thus the initial condition corresponds to the value of the model after 10 minutes.

▷ The following figures show the evolution of the waves according to our model, and the results of the data assimilation experiments.

▷ We have assumed that the heights are observed everywhere, and the velocities are only observed at 49 points.

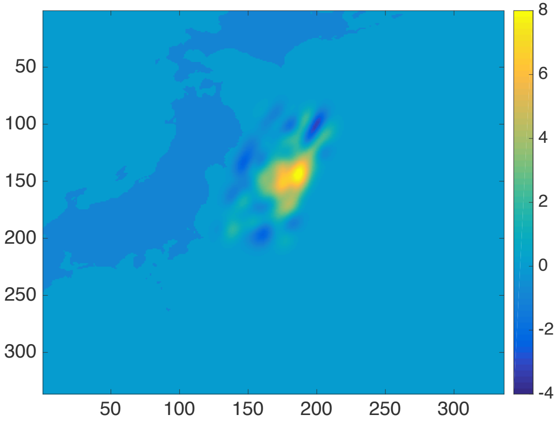

Figure: The height of the tsunami waves (in meters) at 0 mins (grid size $n = 336$).

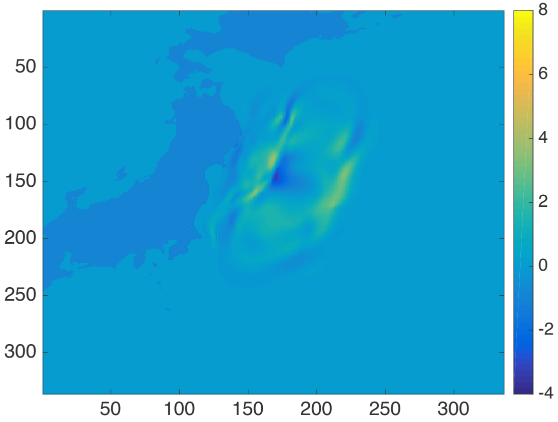

Figure: Evolution of the height of the tsunami waves at 10 mins (grid size $n = 336$).

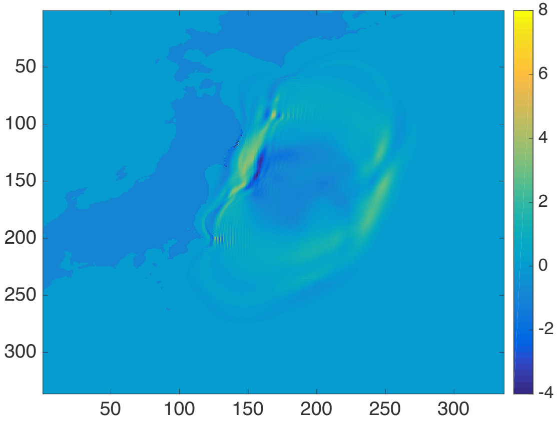

Figure: The height of the tsunami waves at 20 mins (grid size $n = 336$).

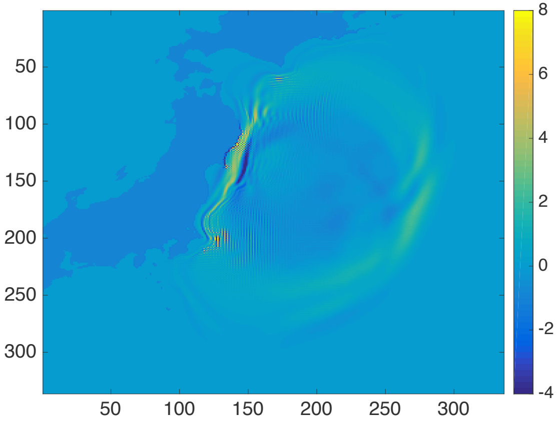

Figure: The height of the tsunami waves at 30 mins (grid size $n = 336$).

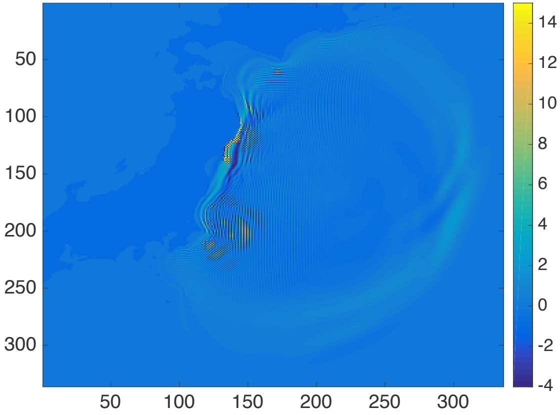

Figure: The height of the tsunami waves at 40 mins (grid size $n = 336$).

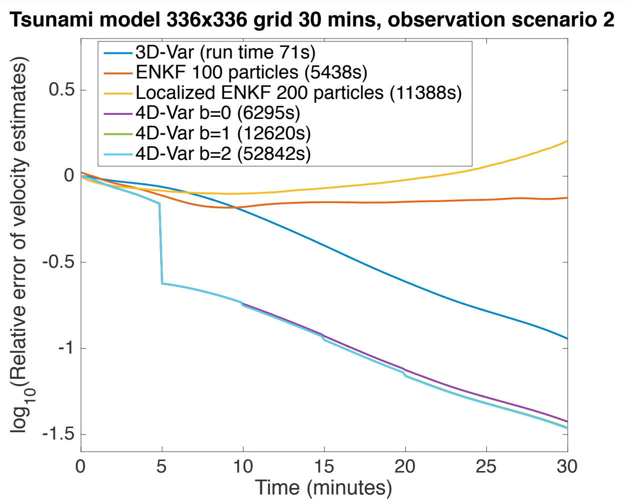

Figure: Relative error of estimates of velocities for tsunami data, all methods. Setting: $n = 336$, $k = 30$, $T = 5\text{mins}$, $\sigma_Z = 10^{-2}$, $\Delta = 2\text{km}$.

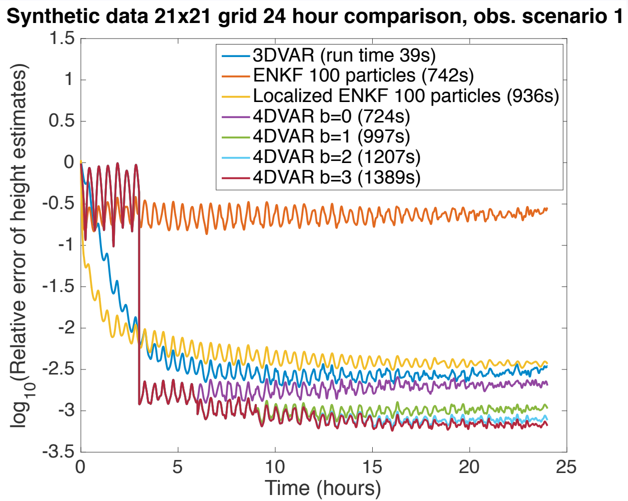

Figure: Relative errors in the case of synthetic data for all methods. Setting: $n = 21$, $k = 1080$, $T = 3\text{h}$, $\sigma_Z = 10^{-2}$, $\Delta = 10\text{km}$.

Conclusion

▷ By starting Newton's method at an appropriate initial point (based on derivatives), we can find the MAP with high probability.

▷ Flow-dependent prior distributions can improve the performance.

▷ This method is competitive with state-of-the art data assimilation techniques for the shallow-water equations.

▷ Performs better than ENKF and localised ENKF at the same computational cost when the background and forecast error covariances are non-localised due to longer assimilation windows.

Open problems

- ▷ Consistency for the flow-dependent case
- ▷ Generalise results to infinite dimensional nonparameteric setting.

References I

Boucheron, S., Lugosi, G., and Massart, P. (2013).

Concentration inequalities.

Oxford University Press, Oxford.

A nonasymptotic theory of independence, With a foreword by Michel Ledoux.

Paulin, D., Jasra, A., Beskos, A., and Crisan, D. (2017).

A 4D-Var method with flow-dependent background covariances for the shallow-water equations.

arXiv preprint arXiv:1710.11529.

References II

Paulin, D., Jasra, A., Crisan, D., and Beskos, A. (2018).

Optimization based methods for partially observed chaotic systems.

Foundations of Computational Mathematics.

Saito, T., Ito, Y., Inazu, D., and Hino, R. (2011).

Tsunami source of the 2011 Tohoku-Oki earthquake, Japan: Inversion analysis based on dispersive tsunami simulations.

Geophysical Research Letters, 38(7).

References III

Salmon, R. (2015).

Introduction to Ocean Waves.

Scripps Institution of Oceanography, University of California, San Diego.

Available at http://pordlabs.ucsd.edu/rsalmon/111.textbook.pdf.

Sanz-Alonso, D. and Stuart, A. M. (2015).

Long-time asymptotics of the filtering distribution for partially observed chaotic dynamical systems.

SIAM/ASA J. Uncertain. Quantif., 3(1):1200–1220.

References IV

Talagrand, O. and Courtier, P. (1987).

Variational assimilation of meteorological observations with the adjoint vorticity equation. i: Theory.

Quarterly Journal of the Royal Meteorological Society, 113(478):1311–1328.

Tropp, J. A. (2015).

An introduction to matrix concentration inequalities.

Foundations and Trends in Machine Learning, 8(1-2):1–230.

THANK YOU!