## Inference via low-dimensional couplings

Youssef Marzouk joint work with Alessio Spantini and Daniele Bigoni

Department of Aeronautics and Astronautics

Center for Computational Engineering Statistics and Data Science Center
Massachusetts Institute of Technology
http://uqgroup.mit.edu

Support from DOE Office of Advanced Scientific Computing Research
1 May 2017

## Bayesian inference in large-scale models

## Observations y <br> Parameters x



$$
\pi_{\mathrm{pos}}(x):=\underbrace{\pi(x \mid y) \propto \pi(y \mid x) \pi_{\mathrm{pr}}(x)}_{\text {Bayes' rule }}
$$

- Need to characterize the posterior distribution (density $\pi_{\text {pos }}$ )
- This is a challenging task since:
- $x \in \mathbb{R}^{n}$ is typically high-dimensional (e.g., a discretized function)
- $\pi_{\text {pos }}$ is non-Gaussian
- evaluations of $\pi_{\text {pos }}$ may be expensive
- $\pi_{\text {pos }}$ can be evaluated up to a normalizing constant


## Sequential Bayesian inference



- State estimation (e.g., filtering and smoothing) or joint state and parameter estimation, in a Bayesian setting
- Need recursive, online algorithms for characterizing the posterior


## Computational challenges

- Extract information from the posterior (means, covariances, event probabilities, predictions) by evaluating posterior expectations:

$$
\mathbb{E}_{\pi_{\mathrm{pos}}}[h(x)]=\int h(x) \pi_{\mathrm{pos}}(x) d x
$$

- Key strategies for making this computationally tractable
- Approximations of the forward model, e.g., polynomial approximations, local interpolants, reduced order models, multi-fidelity approaches
- Efficient and structure-exploiting sampling (integration) schemes


## Deterministic coupling of probability measures



## Core idea

- Choose $\pi_{\text {ref }}$ (e.g., Gaussian). Set $\pi_{\text {tar }}:=\pi_{\text {pos }}$.
- Seek a transport map $T: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ such that $T_{\sharp} \pi_{\text {ref }}=\pi_{\text {tar }}$


## Deterministic coupling of probability measures



## Core idea

- Choose $\pi_{\text {ref }}$ (e.g., Gaussian). Set $\pi_{\text {tar }}:=\pi_{\text {pos }}$.
- Seek a transport map $T: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ such that $T_{\sharp} \pi_{\text {ref }}=\pi_{\text {tar }}$
- Useful outcomes...
- Independent and unweighted samples from the target
- "Precondition" other sampling or quadrature schemes


## Various types of transport

- Optimal transport:

$$
\begin{aligned}
T^{\mathrm{opt}}= & \arg \min _{T} \int_{\mathbb{R}^{n}} c(x, T(x)) \mathrm{d} \pi_{\mathrm{ref}}(x) \\
& \text { s.t. } T_{\sharp} \pi_{\mathrm{ref}}=\pi_{\mathrm{tar}}
\end{aligned}
$$

- Monge (1781) problem; many nice properties, but numerically challenging in general continuous cases...
- Knothe-Rosenblatt rearrangement:

$$
T(x)=\left[\begin{array}{l}
T^{1}\left(x_{1}\right) \\
T^{2}\left(x_{1}, x_{2}\right) \\
\vdots \\
T^{n}\left(x_{1}, x_{2}, \ldots, x_{n}\right)
\end{array}\right]
$$

- Exists and is unique (up to ordering) under mild conditions
- Jacobian determinant easy to evaluate
- Monotonicity is essentially one-dimensional: $\partial_{x_{k}} T^{k}>0$


## Computation of transports

Variational characterization of the direct map $T$ [Moselhy \& M 2012]:

$$
\min _{T \in \mathcal{T}_{\triangle}} \mathcal{D}_{K L}\left(T_{\sharp} \pi_{\text {ref }} \| \pi_{\mathrm{tar}}\right)
$$

- $\mathcal{T}_{\Delta}$ is the set of monotone lower triangular maps
- Contains the Knothe-Rosenblatt rearrangement
- Expectation is with respect to reference measure
- Compute via, e.g., Monte Carlo, QMC, quadrature
- Use evaluations of $\pi_{\text {tar }}$ (and its gradients) directly; avoid MCMC or importance sampling altogether!
- Parameterize $k$-th component map $T^{k}(x)$ with coefficients $\mathbf{f}_{k} \in \mathbb{R}^{p_{k}}$
- Example: monotone parameterization, $\partial_{x_{k}} T^{k}>0$ :

$$
T^{k}\left(x_{1}, \ldots, x_{k}\right)=a_{k}\left(x_{1}, \ldots, x_{k-1}\right)+\int_{0}^{x_{k}} \exp \left(b_{k}\left(x_{1}, \ldots, x_{k-1}, w\right)\right) d w
$$

## Simple example

$$
\min _{\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}} \mathbb{E}_{\pi_{\mathrm{ref}}}\left[-\log \pi_{\mathrm{tar}} \circ T-\sum_{k} \log \partial_{x_{k}} T^{k}\right]
$$

- Parameterized map $T\left(\mathbf{x} ; \mathbf{f}_{1}, \ldots, \mathbf{f}_{n}\right)$
- Optimize over $\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}$
- Use gradient-based optimization (here, BFGS)
- Approximate $\mathbb{E}_{\pi_{\text {ref }}}[g] \approx \sum_{i} w_{i} g\left(\mathbf{x}_{i}\right)$
- The posterior is in the tail of the reference!



## Simple example

$$
\min _{\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}} \mathbb{E}_{\pi_{\mathrm{ref}}}\left[-\log \pi_{\mathrm{tar}} \circ T-\sum_{k} \log \partial_{x_{k}} T^{k}\right]
$$

- Parameterized map $T\left(\mathbf{x} ; \mathbf{f}_{1}, \ldots, \mathbf{f}_{n}\right)$
- Optimize over $\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}$
- Use gradient-based optimization (here, BFGS)
- Approximate $\mathbb{E}_{\pi_{\text {ref }}}[g] \approx \sum_{i} w_{i} g\left(\mathbf{x}_{i}\right)$
- The posterior is in the tail of the reference!



## Simple example

$$
\min _{f_{1}, \ldots, f_{n}} \mathbb{E}_{\pi_{\text {ref }}}\left[-\log \pi_{\mathrm{tar}} \circ T-\sum_{k} \log \partial_{x_{k}} T^{k}\right]
$$

- Parameterized map $T\left(\mathbf{x} ; \mathbf{f}_{1}, \ldots, \mathbf{f}_{n}\right)$
- Optimize over $\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}$
- Use gradient-based optimization (here, BFGS)
- Approximate $\mathbb{E}_{\pi_{\text {ref }}}[g] \approx \sum_{i} w_{i} g\left(\mathbf{x}_{i}\right)$
- The posterior is in the tail of the reference!



## Simple example

$$
\min _{\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}} \mathbb{E}_{\pi_{\mathrm{ref}}}\left[-\log \pi_{\mathrm{tar}} \circ T-\sum_{k} \log \partial_{x_{k}} T^{k}\right]
$$

- Parameterized map $T\left(\mathbf{x} ; \mathbf{f}_{1}, \ldots, \mathbf{f}_{n}\right)$
- Optimize over $\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}$
- Use gradient-based optimization (here, BFGS)
- Approximate $\mathbb{E}_{\pi_{\text {ref }}}[g] \approx \sum_{i} w_{i} g\left(\mathbf{x}_{i}\right)$
- The posterior is in the tail of the reference!



## Simple example

$$
\min _{f_{1}, \ldots, f_{n}} \mathbb{E}_{\pi_{\text {ref }}}\left[-\log \pi_{\mathrm{tar}} \circ T-\sum_{k} \log \partial_{x_{k}} T^{k}\right]
$$

- Parameterized map $T\left(\mathbf{x} ; \mathbf{f}_{1}, \ldots, \mathbf{f}_{n}\right)$
- Optimize over $\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}$
- Use gradient-based optimization (here, BFGS)
- Approximate $\mathbb{E}_{\pi_{\text {ref }}}[g] \approx \sum_{i} w_{i} g\left(\mathbf{x}_{i}\right)$
- The posterior is in the tail of the reference!



## Simple example

Other possible transports:

- Stein variational gradient descent [Liu \& Wang 2016]
- Normalizing flows [Rezende \& Mohamed 2015]
- Particle flows [Heng et al. 2015; Doucet, Daum...]
- Approximations of the optimal transport [Tabak 2013-16]



## Potential advantages

$$
\min _{\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}} \mathbb{E}_{\pi_{\text {ref }}}\left[-\log \pi_{\operatorname{tar}} \circ T-\sum_{k}^{n} \log \partial_{x_{k}} T^{k}\right]
$$

- Move samples; don't just reweigh them
- Use optimization to enhance integration


## Potential advantages

$$
\min _{\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}} \mathbb{E}_{\pi_{\text {ref }}}\left[-\log \pi_{\mathrm{tar}} \circ T-\sum_{k}^{n} \log \partial_{x_{k}} T^{k}\right]
$$

- Move samples; don't just reweigh them
- Use optimization to enhance integration
- Independent, unweighted, and cheap samples from the target (or close to it): $x_{i} \sim \pi_{\text {ref }} \Rightarrow T\left(x_{i}\right) \sim \pi_{\text {tar }}$
- Clear convergence criterion, even with unnormalized target density:

$$
\mathcal{D}_{K L}\left(T_{\sharp} \pi_{\mathrm{ref}} \| \pi_{\mathrm{tar}}\right) \approx \frac{1}{2} \mathbb{V a r}_{\pi_{\mathrm{ref}}}\left[\log \pi_{\mathrm{ref}}-\log T_{\sharp}^{-1} \bar{\pi}_{\mathrm{tar}}\right]
$$

- Key steps are embarrassingly parallel


## Potential advantages

$$
\min _{\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}} \mathbb{E}_{\pi_{\text {ref }}}\left[-\log \pi_{\operatorname{tar}} \circ T-\sum_{k}^{n} \log \partial_{x_{k}} T^{k}\right]
$$

- Move samples; don't just reweigh them
- Use optimization to enhance integration
- Independent, unweighted, and cheap samples from the target (or close to it): $x_{i} \sim \pi_{\text {ref }} \Rightarrow T\left(x_{i}\right) \sim \pi_{\text {tar }}$
- Clear convergence criterion, even with unnormalized target density:

$$
\mathcal{D}_{K L}\left(T_{\sharp} \pi_{\mathrm{ref}} \| \pi_{\mathrm{tar}}\right) \approx \frac{1}{2} \mathbb{V a r}_{\pi_{\text {ref }}}\left[\log \pi_{\mathrm{ref}}-\log T_{\sharp}^{-1} \bar{\pi}_{\mathrm{tar}}\right]
$$

- Key steps are embarrassingly parallel
- Yet we exchange a high-dimensional sampling task for a high-dimensional optimization problem
- Major bottleneck: representation of the map, e.g., cardinality of the map basis $\mathbf{f}_{1}, \ldots, \mathbf{f}_{n}$


## Low-dimensional structure

- How to make the construction/representation of high-dimensional transports tractable?
- Key idea: exploit Markov structure of the posterior
- Leads to various low-dimensional properties of transport maps:
(1) Decomposability
(2) Sparsity
(3) Low-rank/near-identity structure
- Property \#1 above will yield new online algorithms for Bayesian filtering, smoothing, and joint parameter/state estimation


## Markov networks

- Let $Z_{1}, \ldots, Z_{n}$ be random variables with joint density $\pi>0$


$$
(i, j) \notin \mathcal{E} \quad \text { iff } \quad Z_{i} \Perp Z_{j} \mid \mathbf{Z}_{\mathcal{V} \backslash\{i, j\}}
$$

- $\mathcal{G}$ encodes conditional independence (I-map for $\pi$ )
- Theorem: Define $\mathcal{G}$ s.t. $(i, j) \notin \mathcal{E}$ if and only if $\partial_{x_{i}, x_{j}} \log \pi=0$ The resulting $\mathcal{G}$ is the unique minimal $l$-map for $\pi$
- Choice of the probabilistic model $\Longrightarrow$ graphical structure


## A motivating example



- Fix an independent reference density $\eta=\prod_{j} \eta_{X_{j}}$ (left)
- Seek a transport map $T: \mathbb{R}^{6} \rightarrow \mathbb{R}^{6}$ from $\eta$ to $\pi$ (right)
- Is there a low-dimensional T?
- Yes, but we need two ingredients!
(1) Pullback density $T^{\sharp} \pi$ : if $\mathbf{Z} \sim \pi$, then $T^{-1}(\mathbf{Z}) \sim T^{\sharp} \pi$
(2) Graph decomposition
- Remark: if $T$ were the exact transport, we would have $T^{\sharp} \pi=\eta$


## Graph decomposition



## Definition

A triple $(A, S, B)$ of disjoint nonempty subsets of the vertex set $\mathcal{V}$ forms a decomposition of $\mathcal{G}$ if the following hold
(1) $\mathcal{V}=A \cup S \cup B$
(2) $S$ separates $A$ from $B$ in $\mathcal{G}$

## Step 1: build a local map



- For a given decomposition $(A, S, B)$, consider $\mathfrak{M}_{1}: \mathbb{R}^{3} \rightarrow \mathbb{R}^{3}$ s.t.
(1) $\mathfrak{M}_{1}\left(\mathbf{x}_{A}, \mathbf{x}_{S}\right)=\left[\begin{array}{l}A_{1}\left(\mathbf{x}_{S}, \mathbf{x}_{A}\right) \\ B_{1}\left(\mathbf{x}_{S}\right)\end{array}\right]$ pushes forward marginal $\eta_{\mathbf{X}_{\text {SUA }}}$ to $\pi_{\mathbf{x}_{\text {SUA }}}$
(2) Embed $\mathfrak{M}_{1}$ in $T_{1}\left(\mathbf{x}_{A}, \mathbf{x}_{S}, \mathbf{x}_{B}\right)=\left[\begin{array}{l}A_{1}\left(\mathbf{x}_{S}, \mathbf{x}_{A}\right) \\ B_{1}\left(\mathbf{x}_{S}\right) \\ \mathbf{x}_{B}\end{array}\right], T_{1}: \mathbb{R}^{6} \rightarrow \mathbb{R}^{6}$
- What can we say about the pullback density $T_{1}^{\sharp} \pi$ ?


## Local graph sparsification



$$
T=T_{1}
$$

- Figure: Markov structure of the pullback of $\pi$ through $T$
- Just remove any edge incident to any node in $A$
- $T_{1}$ is essentially a 3-D map
- Pulling back $\pi$ through $T_{1}$ makes $\mathbf{Z}_{A}$ independent of $\mathbf{Z}_{S \cup B}$ !


## Do it recursively!



$$
T=T_{1}
$$

- Figure: Markov structure of the pullback of $\pi$ through $T$
- Recursion at step $k$
(1) Consider a new decomposition $(A, S, B)$
(2) Compute transport $T_{k}$
(3) Pull back through $T_{k}$


## Step k: new decomposition and local map



$$
T=T_{1}
$$

- Figure: Markov structure of the pullback of $\pi$ through $T$
- Recursion at step $k$
(1) Consider a new decomposition $(A, S, B)$
(2) Compute transport $T_{k}$
(3) Pull back through $T_{k}$


## Step k: local graph sparsification



- Figure: Markov structure of the pullback of $\pi$ through $T$
- $T_{2}$ is essentially a 4-D map
- Each time we pull back by a new map we remove edges
- Intuition: Continue the recursion until no edges are left. . .


## And so on. . .



$$
T=T_{1} \circ T_{2}
$$

- Figure: Markov structure of the pullback of $\pi$ through $T$
- $T_{2}$ is essentially a 4-D map
- Each time we pullback by a new map we remove edges
- Intuition. Continue the recursion until no edges are left...


## Decomposable maps



- Figure: Markov structure of the pullback of $\pi$ through $T$
- Decomposability of $\mathcal{G} \Rightarrow$ existence of decomposable couplings
- Anisotropic triangular structure of $\left(T_{i}\right)$ is essential
- Idea: inference decomposed into smaller steps (no need for marginals!)
- In fact, we can make this more general...


## Decomposition theorem

## Theorem [Decomposition of transports]

Let $\mathcal{G}$ be an I-map for $\pi$ and let $\eta=\prod_{j} \eta_{X_{j}}$ be a reference density. If $(A, S, B)$ is a decomposition of $\mathcal{G}$, then
(1) $\exists$ a transport map:

$$
T=T_{1} \circ T_{2}
$$

- $T_{1}$ is a monotone triangular transport s.t. $\eta \xrightarrow{T_{1}} \pi_{X_{\text {AUS }}} \cdot\left(\prod_{j \in B} \eta_{X_{j}}\right)$
- $T_{1}$ is the identity map along components in $B: T_{1}^{k}(\mathbf{x})=x_{k}$ for $k \in B$
- $T_{2}$ is any transport s.t. $\eta \xrightarrow{T_{2}} T_{1}^{\sharp} \pi$
(2) $\mathbf{X}_{A}$ is independent of $\mathbf{X}_{S \cup B}$ w.r.t. the pullback density $T_{1}^{\sharp} \pi$
- $T_{2}$ is the identity along components in $A: T_{2}^{k}(\mathbf{x})=x_{k}$ for $k \in A$
- Strategy: recursively apply theorem to further decompose $T_{2}$


## Applications to Bayesian filtering/smoothing

- Nonlinear non-Gaussian state-space model: $\pi_{\mathbf{Z}_{k} \mid \mathbf{Z}_{k-1}}, \pi_{\mathbf{Y}_{k} \mid \mathbf{Z}_{k}}$

- Ideally, interested in recursively updating the full Bayesian solution: $\pi_{\mathbf{Z}_{0: k} \mid \mathbf{Y}_{0: k}} \rightarrow \pi_{\mathbf{Z}_{0: k+1} \mid \mathbf{Y}_{0: k+1}}$ (more difficult)
- Or focus on approximating the filtering distribution:
$\pi_{\mathbf{Z}_{k} \mid \mathbf{Y}_{0: k}} \rightarrow \pi_{\mathbf{Z}_{k+1} \mid \mathbf{Y}_{0: k+1}}$ (marginals of the full Bayesian solution)

Apply the decomposition theorem to $\pi_{\mathbf{Z}_{0}, \ldots, \mathbf{Z}_{k} \mid \mathbf{Y}_{0}, \ldots, \mathbf{Y}_{k}}$ (just a tree!)

## Coupling with an independent process



- Let $\mathbf{X}_{0}, \mathbf{X}_{1}, \ldots$ be an independent process with marginals $\left(\eta_{\mathbf{X}_{k}}\right)_{k}$
- Seek a coupling between $\mathbf{X}_{0}, \ldots, \mathbf{X}_{N}$ and $\mathbf{Z}_{0}, \ldots, \mathbf{Z}_{N} \mid \mathbf{Y}_{0}, \ldots, \mathbf{Y}_{N}$
- Ideally, we would like a low-dimensional decomposable coupling!
- Let's see. . .

First step: compute a 2-D map


- Compute $\mathfrak{M}_{0}: \mathbb{R}^{2 n} \rightarrow \mathbb{R}^{2 n}$ s.t.

$$
\mathfrak{M}_{0}\left(\mathbf{x}_{0}, \mathbf{x}_{1}\right)=\left[\begin{array}{l}
A_{0}\left(\mathbf{x}_{0}, \mathbf{x}_{1}\right) \\
B_{0}\left(\mathbf{x}_{1}\right)
\end{array}\right]
$$

- Reference: $\eta_{\mathrm{X}_{0}} \eta_{\mathrm{X}_{1}}$
- Target: $\pi_{\mathbf{Z}_{0}} \pi_{\mathbf{Z}_{1} \mid \mathbf{Z}_{0}} \pi_{\mathbf{Y}_{0} \mid \mathbf{Z}_{0}} \pi_{\mathbf{Y}_{1} \mid \mathbf{Z}_{1}}$
- $\operatorname{dim}\left(\mathfrak{M}_{0}\right) \simeq 2 \times \operatorname{dim}\left(\mathbf{Z}_{0}\right)$

$$
T_{0}(\mathbf{x})=\left[\begin{array}{l}
A_{0}\left(\mathbf{x}_{0}, \mathbf{x}_{1}\right) \\
B_{0}\left(\mathbf{x}_{1}\right) \\
\mathbf{x}_{2} \\
\mathbf{x}_{3} \\
\mathbf{x}_{4} \\
\mathbf{x}_{5} \\
\vdots \\
\mathbf{x}_{N}
\end{array}\right]
$$

## Second step: compute a 2-D map



- Compute $\mathfrak{M}_{1}: \mathbb{R}^{2 n} \rightarrow \mathbb{R}^{2 n}$ s.t.

$$
\mathfrak{M}_{1}\left(\mathbf{x}_{1}, \mathbf{x}_{2}\right)=\left[\begin{array}{l}
A_{1}\left(\mathbf{x}_{1}, \mathbf{x}_{2}\right) \\
B_{1}\left(\mathbf{x}_{2}\right)
\end{array}\right]
$$

$$
T_{1}(\mathbf{x})=\left[\begin{array}{l}
\mathbf{x}_{0} \\
A_{1}\left(\mathbf{x}_{1}, \mathbf{x}_{2}\right) \\
B_{1}\left(\mathbf{x}_{2}\right) \\
\mathbf{x}_{3} \\
\mathbf{x}_{4} \\
\mathbf{x}_{5} \\
\vdots \\
\mathbf{x}_{N}
\end{array}\right]
$$

## Proceed recursively forward in time



- Compute $\mathfrak{M}_{2}: \mathbb{R}^{2 n} \rightarrow \mathbb{R}^{2 n}$ s.t.

$$
\mathfrak{M}_{2}\left(\mathbf{x}_{2}, \mathbf{x}_{3}\right)=\left[\begin{array}{l}
A_{2}\left(\mathbf{x}_{2}, \mathbf{x}_{3}\right) \\
B_{2}\left(\mathbf{x}_{3}\right)
\end{array}\right]
$$

- Reference: $\eta_{\mathrm{X}_{2}} \eta_{\mathrm{X}_{3}}$
- Target: $\eta_{\mathbf{X}_{2}} \pi_{\mathbf{Y}_{3} \mid \mathbf{Z}_{3}} \pi_{\mathbf{Z}_{3} \mid \mathbf{Z}_{2}}\left(\cdot \mid B_{1}(\cdot)\right)$
- Uses only one component of $\mathfrak{M}_{1}$

$$
T_{2}(\mathbf{x})=\left[\begin{array}{l}
\mathbf{x}_{0} \\
\mathbf{x}_{1} \\
A_{2}\left(\mathbf{x}_{2}, \mathbf{x}_{3}\right) \\
B_{2}\left(\mathbf{x}_{3}\right) \\
\mathbf{x}_{4} \\
\mathbf{x}_{5} \\
\vdots \\
\mathbf{x}_{N}
\end{array}\right]
$$

## A decomposition theorem for chains



Theorem.
(1) $\left(B_{k}\right)_{\sharp} \eta_{\mathbf{X}_{k+1}}=\pi_{\mathbf{Z}_{k+1}} \mid \mathbf{Y}_{0: k+1}$
(2) $\left(\mathfrak{M}_{k}\right)_{\sharp} \eta_{\mathrm{X}_{k: k+1}} \simeq \pi_{\mathrm{Z}_{k}, \mathrm{Z}_{k+1}} \mid \mathrm{Y}_{0: k+1}$
(3) $\left(T_{1} \circ \cdots \circ T_{k}\right)_{\sharp} \eta_{\mathbf{x}_{0: k+1}}=\pi_{\mathbf{Z}_{0: k+1} \mid} \mid \mathbf{Y}_{0: k+1}$
(lag-1 smoothing)
(filtering)
(full Bayesian solution)

## A nested decomposable coupling!

$-\mathfrak{T}_{k}=T_{0} \circ T_{1} \circ \cdots \circ T_{k}$ characterizes the full joint dist $\pi_{Z_{0: k+1} \mid} \mid Y_{0: k+1}$

$$
\mathfrak{T}_{k}(\mathbf{x})=\underbrace{\left[\begin{array}{l}
A_{0}\left(\mathbf{x}_{0}, \mathbf{x}_{1}\right) \\
B_{0}\left(\mathbf{x}_{1}\right) \\
\mathbf{x}_{2} \\
\mathbf{x}_{3} \\
\mathbf{x}_{4} \\
\mathbf{x}_{5} \\
\vdots \\
\mathbf{x}_{N}
\end{array}\right]}_{T_{0}} \circ \underbrace{\left[\begin{array}{l}
\mathbf{x}_{0} \\
A_{1}\left(\mathbf{x}_{1}, \mathbf{x}_{2}\right) \\
B_{1}\left(\mathbf{x}_{2}\right) \\
\mathbf{x}_{3} \\
\mathbf{x}_{4} \\
\mathbf{x}_{5} \\
\vdots \\
\mathbf{x}_{N}
\end{array}\right]}_{T_{1}} \circ \underbrace{\left[\begin{array}{l}
\mathbf{x}_{0} \\
\mathbf{x}_{1} \\
A_{2}\left(\mathbf{x}_{2}, \mathbf{x}_{3}\right) \\
B_{2}\left(\mathbf{x}_{3}\right) \\
\mathbf{x}_{4} \\
\mathbf{x}_{5} \\
\vdots \\
\mathbf{x}_{N}
\end{array}\right]}_{T_{2}}
$$

- $\mathfrak{T}_{k}$ is dense and high-dimensional but decomposable!
- Trivial to go from $\mathfrak{T}_{k}$ to $\mathfrak{T}_{k+1}$ : just append a new map $T_{k+1}$
- No need to recompute $T_{0}, \ldots, T_{k}$ (nested transports)


## A single-pass algorithm for online estimation

## - Algorithm:

(1) Compute the maps $\mathfrak{M}_{0}, \mathfrak{M}_{1}, \ldots$, each of dimension $2 \times \operatorname{dim}\left(\mathbf{Z}_{0}\right)$
(2) Embed each $\mathfrak{M}_{j}$ into an identity map to form $T_{j}$
(3) Evaluate $T_{0} \circ \cdots \circ T_{k}$ for the full Bayesian solution

- Remarks:
- A single pass on the state-space model
- Maps $\mathfrak{M}_{0}, \mathfrak{M}_{1}, \ldots$ need not be recomputed given new data
- Constant effort per assimilated observation (online estimation)
- Variational algorithm: no particles and no particle degeneracy!
- Of course, we still need to compute each $\mathfrak{M}_{j}$ (many options)
- In spirit, a non-Gaussian generalization of the RTS smoother

Full Bayesian solution $\simeq$ lag-1 smoothing (using couplings)

## Joint parameter/state estimation

- Can be generalized to sequential joint parameter/state estimation

- $\left(T_{0} \circ \cdots \circ T_{k}\right)_{\sharp} \eta_{\Theta} \eta_{\mathbf{X}_{0: k+1}}=\pi_{\Theta, \mathrm{Z}_{0: k+1} \mid \mathrm{Y}_{0: k+1}}$ (full Bayesian solution)
- But now $\operatorname{dim}\left(\mathfrak{M}_{j}\right)=2 \times \operatorname{dim}\left(\mathbf{Z}_{j}\right)+\operatorname{dim}(\Theta)$
- Remarks:
- Online algorithm (unlike, e.g., particle marginal Metropolis Hastings)
- No artificial dynamic for the static parameters
- No a priori fixed-lag smoothing approximation


## Numerical example: stochastic volatility model

- Stochastic volatility model: Latent log-volatilities take the form of an $\operatorname{AR}(1)$ process for $t=1, \ldots, N$ :

$$
Z_{t+1}=\mu+\phi\left(Z_{t}-\mu\right)+\eta_{t}, \quad \eta_{t} \sim \mathcal{N}(0,1), \quad Z_{1} \sim \mathcal{N}\left(0,1 / 1-\phi^{2}\right)
$$

- Observe the mean return for holding an asset at time $t$

$$
Y_{t}=\varepsilon_{t} \exp \left(0.5 Z_{t}\right), \quad \varepsilon_{t} \sim \mathcal{N}(0,1), \quad t=1, \ldots, N
$$

- Markov structure for $\pi \sim \mu, \phi, \mathbf{Z}_{1: N} \mid \mathbf{Y}_{1: N}$ is given by:

- Joint state/parameter estimation problem


## Stochastic volatility model with hyperparameters

- Build the decomposition recursively

$$
T=\mathbf{I d}
$$



- Figure: Markov structure for the pullback of $\pi$ through $T$
- Start with the identity map


## Stochastic volatility model with hyperparameters

- Build the decomposition recursively

$$
T=\mathbf{I d}
$$



- Figure: Markov structure for the pullback of $\pi$ through $T$
- Find a good first decomposition of $\mathcal{G}$


## Stochastic volatility model with hyperparameters

- Build the decomposition recursively

$$
T=T_{1}
$$



- Figure: Markov structure for the pullback of $\pi$ through $T$
- Compute an (essentially) 4-D $T_{1}$ and pull back $\pi$
- Underlying approximation of $\mu, \phi, \mathbf{Z}_{1} \mid \mathbf{Y}_{1}$


## Stochastic volatility model with hyperparameters

- Build the decomposition recursively

$$
T=T_{1}
$$



- Figure: Markov structure for the pullback of $\pi$ through $T$
- Find a new decomposition
- Underlying approximation of $\mu, \phi, \mathbf{Z}_{1} \mid \mathbf{Y}_{1}$


## Stochastic volatility model with hyperparameters

- Build the decomposition recursively

$$
T=T_{1} \circ T_{2}
$$



- Figure: Markov structure for the pullback of $\pi$ through $T$
- Compute an (essentially) 4-D $T_{2}$ and pull back $\pi$
- Underlying approximation of $\mu, \phi, \mathbf{Z}_{1: 2} \mid \mathbf{Y}_{1: 2}$


## Stochastic volatility model with hyperparameters

- Build the decomposition recursively

$$
T=T_{1} \circ T_{2}
$$



- Figure: Markov structure for the pullback of $\pi$ through $T$
- Continue the recursion until no edges are left...
- Underlying approximation of $\mu, \phi, \mathbf{Z}_{1: 2} \mid \mathbf{Y}_{1: 2}$


## Stochastic volatility model with hyperparameters

- Build the decomposition recursively

$$
T=T_{1} \circ T_{2} \circ T_{3}
$$



- Figure: Markov structure for the pullback of $\pi$ through $T$
- Continue the recursion until no edges are left...
- Underlying approximation of $\mu, \phi, \mathbf{Z}_{1: 3} \mid \mathbf{Y}_{1: 3}$


## Stochastic volatility model with hyperparameters

- Build the decomposition recursively

$$
T=T_{1} \circ T_{2} \circ T_{3} \circ \cdots \circ T_{N-2}
$$



- Figure: Markov structure for the pullback of $\pi$ through $T$
- Continue the recursion until no edges are left...
- Underlying approximation of $\mu, \phi, \mathbf{Z}_{1: N-1} \mid \mathbf{Y}_{1: N-1}$


## Stochastic volatility model with hyperparameters

- Build the decomposition recursively

$$
T=T_{1} \circ T_{2} \circ T_{3} \circ \cdots \circ T_{N-2} \circ T_{N-1}
$$



- Figure: Markov structure for the pullback of $\pi$ through $T$
- Each map $T_{k}$ is essentially 4-D regardless of $N$
- Underlying approximation of $\mu, \phi, \mathbf{Z}_{1: N} \mid \mathbf{Y}_{1: N}$


## Stochastic volatility example (102-dim)



- Joint parameter/state inference problem solved with a single forward pass [filtering]


## Stochastic volatility example (102-dim)



- Joint parameter/state inference problem solved with a single forward pass, by composing low-dimensional transports [smoothing]


## Stochastic volatility example



- Online parameter estimation: marginals of hyperparameters $\mu, \phi$, conditioning on successively more observations $\mathbf{y}_{0: k}$


## Stochastic volatility example




- Marginals of hyperparameters $\mu, \phi$ : transport maps (solid), MCMC with ESS $=10^{5}$ (dashed)


## Stochastic volatility example



- Quantiles of smoothing marginals of the state $\mathbf{Z}_{0: N}$ (red) compared to MCMC (black)


## Stochastic volatility example

- If $\eta \sim \mathcal{N}(0, \mathbf{I})$ and $T_{\sharp} \eta=\pi$, then $T^{\sharp} \pi$ should be Gaussian!

- Figure: 2-D random conditionals of the pullback density $T^{\sharp} \pi$
- Variance diagnostic $\approx 8.05 \times 10^{-2}$


## Dual property: sparsity

## Theorem [Sparsity of triangular transports]

If $\mathcal{G}$ is an I-map for $\pi_{\text {pos }}$, then we can determine tight lower bounds on the sparsity patterns of:

- Direct transport $T_{\sharp} \pi_{\text {ref }}=\pi_{\text {pos }}$
- Inverse transport $S_{\sharp} \pi_{\text {pos }}=\pi_{\text {ref }}$
only by performing operations on the graph $\mathcal{G}$ (no need to evaluate $\pi_{\text {pos }}$ ).
- Example: Sparsity of inverse transport $S_{\sharp} \pi_{\text {pos }}=\pi_{\text {ref }}$

- Result: enforce sparsity structure in the approximation space $\mathcal{S}_{\triangle}$, e.g., $\min _{S \in \mathcal{S}_{\triangle}} \mathcal{D}_{K L}\left(\pi_{\text {ref }} \| S_{\sharp} \pi_{\text {pos }}\right)$


## Too many cycles. . .



- For certain graphs, sparsity/decomposability do not imply decoupling between the nominal dimension of the problem and the dimension of each transport $T_{i}$ (or the sparsity of $S$ )
- Here, $\mathcal{G}$ is an $n \times n$ grid graph
- $T^{S \cup A}$ acts on $2 n$ dimensions at each stage
- Nonetheless, the notion of composition of transports has still potential. . .


## Beyond the Markov properties of $\pi$

- Key idea: seek low-rank structure and near-identity maps
- Example: fix target $\pi$ to be the posterior density of a Bayesian inference problem,

$$
\pi(\mathbf{z}):=\pi_{\mathrm{pos}}(\mathbf{z}) \propto \pi_{\mathrm{Y} \mid \mathrm{Z}}(\mathbf{y} \mid \mathbf{z}) \pi_{\mathrm{Z}}(\mathbf{z})
$$

- Let $T_{\mathrm{pr}}$ push forward the reference $\eta$ to the prior $\pi_{\mathbf{Z}}$ (prior map)

$$
\widehat{\pi}_{\mathrm{pos}}(\mathbf{z}):=T_{\mathrm{pr}}^{\sharp} \pi_{\mathrm{pos}}(\mathbf{z}) \propto \pi_{\mathrm{Y} \mid \mathrm{Z}}\left(\mathbf{y} \mid T_{\mathrm{pr}}(\mathbf{z})\right) \eta(\mathbf{z})
$$

## Theorem [Graph decoupling]

If $\eta=\prod_{i} \eta_{X_{i}}$ and

$$
\operatorname{rank} \mathbb{E}_{\eta}[\nabla \log R \otimes \nabla \log R]=k, \quad R=\widehat{\pi}_{\mathrm{pos}} / \eta=\pi_{\mathrm{Y} \mid \mathrm{Z}} \circ T_{\mathrm{pr}}
$$

then there exists a rotation $Q$ such that:

$$
Q^{\sharp} \widehat{\pi}_{\mathrm{pos}}(\mathbf{z})=g\left(z_{1}, \ldots, z_{k}\right) \prod_{i>k}^{n} \eta_{X_{i}}\left(z_{i}\right)
$$

## Changing the Markov structure. . .

- The pullback has a different Markov structure:

$$
Q^{\sharp} \widehat{\pi}_{\mathrm{pos}}(\mathbf{z})=g\left(z_{1}, \ldots, z_{k}\right) \prod_{i>k}^{n} \eta_{X_{i}}\left(z_{i}\right)
$$



G

$\mathcal{G}$ Pullback

- Corollary: There exists a transport $T_{\sharp} \eta=Q^{\sharp} \widehat{\pi}_{\text {pos }}$ of the form $T(\mathbf{x})=\left[g\left(\mathbf{x}_{1: k}\right), x_{k+1}, \ldots, x_{n}\right]$, where $g: \mathbb{R}^{k} \rightarrow \mathbb{R}^{k}$.
- The composition $T_{\mathrm{pr}} \circ Q \circ T$ pushes forward $\eta$ to $\pi_{\text {pos }}$
- Why low rank structure? For example, few data-informed directions.


## Log-Gaussian Cox process



- 4096-D GMRF prior, $\mathbf{Z} \sim \mathcal{N}(\mu, \Gamma), \Gamma^{-1}$ specified through $\triangle+\kappa^{2}$ Id
- 30 sparse observations at locations $i \in \mathcal{I}, \mathbf{Y}_{i} \mid \mathbf{Z}_{i} \sim \operatorname{Pois}\left(\exp \mathbf{Z}_{i}\right)$
- Posterior density $\mathbf{Z} \mid \mathbf{Y} \sim \pi_{\text {pos }}$ is:

$$
\pi_{\mathrm{pos}}(\mathbf{z}) \propto \prod_{i \in \mathcal{I}} \exp \left[-\exp \left(z_{i}\right)+z_{i} \cdot y_{i}\right] \exp \left[-\frac{1}{2}(\mathbf{z}-\boldsymbol{\mu})^{\top} \Gamma^{-1}(\mathbf{z}-\boldsymbol{\mu})\right]
$$

- What is an independence map $\mathcal{G}$ for $\pi_{\text {pos }}$ ?


## Log-Gaussian Cox process



- 4096-D GMRF prior, $\mathbf{Z} \sim \mathcal{N}(\mu, \Gamma), \Gamma^{-1}$ specified through $\triangle+\kappa^{2}$ Id
- 30 sparse observations at locations $i \in \mathcal{I}, \mathbf{Y}_{i} \mid \mathbf{Z}_{i} \sim \operatorname{Pois}\left(\exp \mathbf{Z}_{i}\right)$
- Posterior density $\mathbf{Z} \mid \mathbf{Y} \sim \pi_{\text {pos }}$ is:

$$
\pi_{\mathrm{pos}}(\mathbf{z}) \propto \prod_{i \in \mathcal{I}} \exp \left[-\exp \left(z_{i}\right)+z_{i} \cdot y_{i}\right] \exp \left[-\frac{1}{2}(\mathbf{z}-\boldsymbol{\mu})^{\top} \Gamma^{-1}(\mathbf{z}-\boldsymbol{\mu})\right]
$$

- What is an independence map $\mathcal{G}$ for $\pi_{\text {pos }}$ ? A $64 \times 64$ grid.


## Log-Gaussian Cox process

- Fix $\pi_{\text {ref }} \sim \mathcal{N}(0, \mathbf{I})$ and let $T_{\text {pr }}$ push forward $\pi_{\text {ref }}$ to $\pi_{\mathrm{pr}}$ (prior map)
- Consider the pullback $\widehat{\pi}_{\text {pos }}=T_{\text {pr }}^{\sharp} \pi_{\text {pos }}$ and find that

$$
\text { rank } \mathbb{E}_{\pi_{\text {ref }}}[\nabla \log R \otimes \nabla \log R]=30 \ll n, \quad R=\widehat{\pi}_{\text {pos }} / \pi_{\text {ref }}
$$

- Deflate the problem and compute a transport map in 30 dimensions
- Change from prior to posterior concentrated in a low-dimensional subspace (LIS Cui, Law, M 2014; AS Constantine 2015)

truth

posterior sample

posterior mean


## Log-Gaussian Cox process



- (left) $\mathbb{E}[\mathbf{Z} \mid \mathbf{y}]$, (right) $\operatorname{Var}[\mathbf{Z} \mid \mathbf{y}]$. (top) transport; (bottom) MCMC
- Excellent match with reference MCMC solution, on a problem of $n=4096$ dimensions


## Conclusions

- Bayesian inference through the variational construction of deterministic couplings
- Computation of transport maps in high dimensions, leveraging the Markov structure of the posterior:
(1) Decomposability of direct transports
- New online algorithms for Bayesian filtering, smoothing, and parameter estimation
(2) Sparsity of triangular transports
(3) Near-identity transports
- Much ongoing work...
- Adaptive parameterizations of monotone maps
- Nonparametric transports and gradient flows
- Preconditioning sparse quadrature and QMC schemes
- Approximately sparse Markov structures


## References

- A. Spantini, D. Bigoni, Y. Marzouk. "Inference via low-dimensional couplings." arXiv:1703.06131 (main reference for this talk)
- Y. Marzouk, T. Moselhy, M. Parno, A. Spantini, "An introduction to sampling via measure transport." Handbook of Uncertainty Quantification, R. Ghanem, D. Higdon, H. Owhadi, eds. Springer (2016). arXiv:1602.05023.
- A. Spantini, D. Bigoni, Y. Marzouk. "Variational inference via decomposable transports: algorithms for Bayesian filtering and smoothing" and "Adaptive construction of measure transports for Bayesian inference." NIPS workshop on Advances in Approximate Bayesian Inference (2016).
- M. Parno, T. Moselhy, Y. Marzouk, "A multiscale strategy for Bayesian inference using transport maps." SIAM JUQ, 4: 1160-1190 (2016).
- M. Parno, Y. Marzouk, "Transport map accelerated Markov chain Monte Carlo." arXiv:1412.5492.
- T. Moselhy, Y. Marzouk, "Bayesian inference with optimal maps." J. Comp. Phys., 231: 7815-7850 (2012).
- Python code just released at http://transportmaps.mit.edu

