

# An ensemble filter for heavy-tailed $t$ -distributions

**Mathieu Le Provost**<sup>1</sup>, Ricardo Baptista<sup>2</sup>,  
Youssef Marzouk<sup>3</sup> and Jeff D. Eldredge<sup>4</sup>

<sup>1</sup>Long Island University, <sup>2</sup>California Institute of Technology,  
<sup>3</sup>Massachusetts Institute of Technology, <sup>4</sup>University of California, Los Angeles

18th EnKF workshop

May 4, 2023

LIU

Caltech

MIT

UCLA

# Motivations

**Departure from Gaussian tails** is a common feature of geophysical inference problems due to the nonlinear dynamical and observation processes and the uncertainty from the physical sensors.

Many filters like the EnKF assume **at least** that the **tails of the forecast distribution are Gaussian** and not suited for heavy-tailed distributions.

**Objective:** How can do consistent inference in heavy-tailed filtering problems?

## Problem setting

We consider a generic state-space model:

The evolution of the state  $(\mathbf{X}_t)_{t \geq 0}$  is fully described by the initial distribution  $\pi_{\mathbf{X}_0}$  and the dynamical model:

$$\mathbf{X}_t = f(\mathbf{X}_{t-1}) + \mathbf{W}_t$$

We collect observations  $(\mathbf{Y}_t)_{t \geq 0}$  at every time step according to the observation model:

$$\mathbf{Y}_t = h(\mathbf{X}_t) + \mathcal{E}_t$$

**Objective:** Sequentially estimate **the filtering density**  $\pi_{t,j,t} := \pi_{\mathbf{X}_t,j} \mathbf{Y}_{1:t}=y_{1:t}$


# Generic ensemble filtering algorithm

Ensemble filtering methods propagate a set of  $M$  particles  $f_X^{(1)}, \dots, X^{(M)}$  to form an empirical approximation for the filtering density  $\pi_{t|j_t}$ .

1. Forecast step: Forecast dist. at time  $t-1$   $\pi_{t-1|j_{t-1}}$  ! Forecast dist.  $\pi_{t|j_{t-1}}$   
We obtain samples  $f_X^{(1)}, \dots, X^{(M)}$   $\pi_{t|j_{t-1}}$
2. Analysis step: Forecast dist.  $\pi_{t|j_{t-1}}$  ! Filtering dist. at time  $t$   $\pi_{t|j_t}$   
We obtain samples  $f_X^{(1)}, \dots, X^{(M)}$   $\pi_{t|j_t}$

Ensemble filtering algorithms share the same forecast step but differ in the analysis step.

## A “transformative” view of the analysis step



Analysis step: application of the **analysis map**  $T$ : Prior  $\pi_{t|t-1}$  → Posterior  $\pi_{t|t}$


The analysis map of the Kalman filter  $T_{KF}$ :

$$T_{KF}(y, x) = x - \Sigma_{X,Y} \Sigma_Y^{-1} (y - y^*) = x - K(y - y^*)$$

The ensemble Kalman filter (EnKF) [Evensen, 1994] constructs an estimate  $\hat{K} \in \mathbb{R}^{n \times d}$  from limited samples  $f^1, \dots, f^M$  of the forecast distribution.

# Transport map between two probability measures

- Seek a transport map  $S$  that pushes forward  $\pi$  to  $\eta$ , i.e.  $S_{\#}\pi = \eta$ .
- Generate cheap and independent samples  $x \sim \pi \Rightarrow S(x) \sim \eta$ .



## Looking for a good map [Marzouk et al., 2016]

Consider the **Knothe-Rosenblatt (KR) rearrangement**  $S$  s.t.  $S_{\#}\pi = \eta$

$$S(z) = S(z_1, z_2, \dots, z_m) = \begin{bmatrix} S^1(z_1) \\ S^2(z_1, z_2) \\ \vdots \\ S^m(z_1, z_2, \dots, z_m) \end{bmatrix}.$$


The KR has many nice features for Bayesian inference, e.g. **easily invertible** and  $\det r S(x)$  is **simple to evaluate** [Marzouk et al., 2016, Baptista et al., 2020].

The 1D map  $\xi \mapsto S^k(x_1, x_2, \dots, x_{k-1}, \xi)$  characterizes the **marginal conditional**  $\pi_{X_k | X_{1:k-1} = x_{1:k-1}}(\xi)$ .


### Gaussian case

Consider  $\mathbf{X} \sim \pi_{\mathbf{X}} = N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  and let  $LL^T = \boldsymbol{\Sigma}^{-1}$  be the Cholesky factorization of  $\boldsymbol{\Sigma}^{-1}$ . Then  $S(x) = L(x - \boldsymbol{\mu})$  is the KR that pushes forward  $\pi_{\mathbf{X}}$  to  $\eta = N(\mathbf{0}_n, I_n)$ .

# Construction of the analysis map [Spantini et al., 2022]



$T(y, x)$








# Construction of the analysis map [Spantini et al., 2022]

Consider the KR rearrangement  $S$  s.t.  $S_{\#}\pi_{\mathbf{Y},\mathbf{X}} = \eta$

$$S(y, x) = \begin{bmatrix} S^{\mathbf{y}}(y) \\ S^{\mathbf{x}}(y, x) \end{bmatrix},$$




# Construction of the analysis map [Spantini et al., 2022]

Consider the KR rearrangement  $S$  s.t.  $S_{\#}\pi_{\mathbf{Y},\mathbf{X}} = \eta$

$$S(y, x) = \begin{bmatrix} S^{\mathbf{y}}(y) \\ S^{\mathbf{x}}(y, x) \end{bmatrix},$$

The map  $\xi \mapsto S^{\mathbf{x}}(y^*, \xi)$  pushes forward  $\pi_{\mathbf{X}|\mathbf{Y}}(j y^*)$  to  $\eta$




# Construction of the analysis map [Spantini et al., 2022]

Consider the KR rearrangement  $S$  s.t.  $S_{\#}\pi_{\mathbf{Y},\mathbf{X}} = \eta$

$$S(y, x) = \begin{bmatrix} S^{\mathbf{y}}(y) \\ S^{\mathbf{x}}(y, x) \end{bmatrix},$$

The map  $\xi \mapsto S^{\mathbf{x}}(y^*, \xi)$  pushes forward  $\pi_{\mathbf{X}|\mathbf{Y}}(\cdot|y^*)$  to  $\eta$

$S^{\mathbf{x}}(\mathbf{Y}, \mathbf{X}) \quad \eta$




# Construction of the analysis map [Spantini et al., 2022]

Consider the KR rearrangement  $S$  s.t.  $S_{\#}\pi_{\mathbf{Y},\mathbf{X}} = \eta$

$$S(y, x) = \begin{bmatrix} S^{\mathbf{y}}(y) \\ S^{\mathbf{x}}(y, x) \end{bmatrix},$$

The map  $\xi \mapsto S^{\mathbf{x}}(y^*, \xi)$  pushes forward  $\pi_{\mathbf{X}|\mathbf{Y}}(\cdot|y^*)$  to  $\eta$

$$S^{\mathbf{x}}(\mathbf{Y}, \mathbf{X}) \quad \eta$$



# Derivation of the analysis map of the Kalman filter [Spantini et al., 2022]

$$T(y, x) = S^{\mathcal{X}}(y^*, \cdot)^{-1} \circ S^{\mathcal{X}}(y, x)$$

We recover the analysis map of the Kalman filter when

1.  $S^{\mathcal{X}}$  is linear
2. The reference density is the standard normal distribution  $\eta = N(\mathbf{0}_{d+n}, I_{d+n})$

Ensemble filters differ in the choice of

the reference density

the class of functions to represent  $S^{\mathcal{X}}$

the estimation of  $S^{\mathcal{X}}$  from samples

# Limitations of light-tailed filter for heavy-tailed distributions

Many ensemble filters like the EnKF assume **at least** that the **tails of the forecast distribution are Gaussian**.

These filters don't provide consistent inference for heavy-tailed filtering problems.

**Our contribution:** Introduce a new ensemble filter called **ensemble robust filter (EnRF)** based on the following assumptions:

We restrict  $S^{\mathcal{X}}$  to be linear

We choose a reference distribution whose **tail-heaviness can be adapted to the data**.

## $t$ -distributions

$t$ -distributions are a family of distributions parameterized by a mean  $\mu_{\mathbf{x}} \in \mathbb{R}^n$ , a scale matrix  $C_{\mathbf{x}} \in \mathbb{R}^{n \times n}$ , and a degree of freedom  $\nu_{\mathbf{x}} \in [1, \infty[$ .

The degree of freedom  $\nu_{\mathbf{x}}$  characterizes the tail-heaviness:

For  $\nu_{\mathbf{x}} = 1$ , we recover the Cauchy distribution


For  $\nu_{\mathbf{x}} = \infty$ , we recover the Gaussian distribution

Relation between the parameters and the moments of a  $t$ -distribution:

$$\mathbb{E}_{\pi_{\mathbf{x}}} [X] = \mu_{\mathbf{x}}, \text{ for } \nu_{\mathbf{x}} > 1$$

$$\mathbb{E}_{\pi_{\mathbf{x}}} [(X - \mu_{\mathbf{x}})(X - \mu_{\mathbf{x}})^{\top}] = \frac{\nu_{\mathbf{x}}}{\nu_{\mathbf{x}} - 2} C_{\mathbf{x}} \text{ for } \nu_{\mathbf{x}} > 2.$$


## Probability density function of $t$ -distributions



PDF of the univariate standard  $t$ -distribution  $St(0, 1, \nu)$  for  $\nu = 2, 5, 10, 20, 100$  and the univariate standard Gaussian distribution  $N(0, 1) = St(0, 1, \infty)$  (black).



## Quantiles of the $t$ -distributions



Evolution of  $\alpha$ -quantile  $q_\alpha$  for  $\alpha = 1\%, 2\%, 5\%, 10\%$  with the degree of freedom  $\nu$  of the univariate standard  $t$ -distribution. Dashed lines corresponds to the  $\alpha$ -quantiles for the standard Gaussian distribution  $N(0, 1) = St(0, 1, 1)$ .

## Useful properties of $t$ -distributions:

1.  $t$ -distributions are closed under affine transformations:

$$\text{If } \mathbf{X} \sim St(\boldsymbol{\mu}_X, C_X, \nu_X), \text{ then } \mathbf{Z} = A\mathbf{X} + b \sim St\left(A\boldsymbol{\mu}_X + b, AC_XA^T, \nu_X\right)$$

2. Conditional and marginal distributions are known in closed form.

$$\text{Consider } \begin{bmatrix} \mathbf{Y} \\ \mathbf{X} \end{bmatrix} \sim St\left(\begin{bmatrix} \boldsymbol{\mu}_X \\ \boldsymbol{\mu}_Y \end{bmatrix}, \begin{bmatrix} C_Y & C_{X,Y}^T \\ C_{X,Y} & C_X \end{bmatrix}, \nu\right)$$

**Marginal:**  $\mathbf{X} \sim St(\boldsymbol{\mu}_X, C_X, \nu_X)$

**Conditional:**  $\pi_{\mathbf{X}|\mathbf{Y}=y} \sim St\left(\boldsymbol{\mu}_{\mathbf{X}|\mathbf{Y}=y}, C_{\mathbf{X}|\mathbf{Y}=y}, \nu_{\mathbf{X}|\mathbf{Y}=y}\right)$  with

$$\begin{aligned} \boldsymbol{\mu}_{\mathbf{X}|\mathbf{Y}=y} &= \boldsymbol{\mu}_X + C_{X,Y}C_Y^{-1}(y - \boldsymbol{\mu}_Y) \\ C_{\mathbf{X}|\mathbf{Y}=y} &= \underbrace{\frac{\nu + (y - \boldsymbol{\mu}_Y)^T C_Y^{-1}(y - \boldsymbol{\mu}_Y)}{\nu + d}}_{\alpha_Y(y) > 0} \underbrace{\left( C_X - C_{X,Y}C_Y^{-1}C_{X,Y}^T \right)}_{\text{Schur complement } C_{X|Y}} \end{aligned}$$

$$\nu_{\mathbf{X}|\mathbf{Y}=y} = \nu + d$$

## Analysis map $T_\nu$ for $t$ -distributions

Let  $S_\nu$  the KR that pushes forward the joint  $t$ -distribution  $\pi_{\mathbf{Y}, \mathbf{X}}$  with dof  $\nu$  to a “judicious”  $t$ -distribution  $\eta_\nu$  with same dof, i.e.  $S_{\nu\#}\pi_{\mathbf{Y}, \mathbf{X}} = \eta_\nu$ .

**Key:**  $S_\nu$  can be computed in closed form (new result).

We obtain the analysis map  $T_\nu : \mathbb{R}^d \times \mathbb{R}^n \rightarrow \mathbb{R}^n$  by partial inversion of  $S_\nu^{\mathcal{X}}$ :

$$\begin{aligned} T_\nu(y, x) &= S_\nu^{\mathcal{X}}(y^*, \cdot)^{-1} S_\nu^{\mathcal{X}}(y, x) \\ &= \boldsymbol{\mu}_{\mathbf{X}} + C_{\mathbf{X}, \mathbf{Y}} C_{\mathbf{Y}}^{-1}(y^* - \boldsymbol{\mu}_{\mathbf{Y}}) + \sqrt{\frac{\alpha_{\mathbf{Y}}(y^*)}{\alpha_{\mathbf{Y}}(y)}} [(x - \boldsymbol{\mu}_{\mathbf{X}}) - C_{\mathbf{X}, \mathbf{Y}} C_{\mathbf{Y}}^{-1}(y - \boldsymbol{\mu}_{\mathbf{Y}})] \end{aligned}$$

Note that  $C_{\mathbf{X}, \mathbf{Y}} C_{\mathbf{Y}}^{-1} = \boldsymbol{\Sigma}_{\mathbf{X}, \mathbf{Y}} \boldsymbol{\Sigma}_{\mathbf{Y}}^{-1}$ .

## Connection with the Kalman filter

Let's perform an asymptotic expansion of  $T_\nu$  for large  $\nu$ :

$$T_\nu = \underbrace{T_1}_{T_{KF}} + \text{higher order terms in } O\left(\frac{1}{\nu}\right)$$

The zeroth order term  $T_1$  is exactly the analysis map of the Kalman filter  $T_{KF}$ .

The higher order terms correct the analysis map  $T_{KF}$  to account for the finite degree of freedom.

### Takeaway

The analysis map  $T_\nu$  generalizes the analysis map of the Kalman filter for  $t$ -distributions with finite degree of freedom!

# Sensitivity to outlying synthetic observations

**Context:** Observation models often rely on simplified physics and suffer from mis-specifications of the observation operator and observation noise.

Let  $(y^{(j)}, x^{(j)})$  be a joint forecast sample with an outlying synthetic observation  $y^{(j)}$  generated by the likelihood model  $\pi_{\mathbf{Y}}|_{\mathbf{X}=x^{(j)}}$ , such that  $\delta_{\mathbf{Y}}(y^{(j)}) = (y^{(j)} - \mu_{\mathbf{Y}})^{\top} C_{\mathbf{Y}}^{-1} (y^{(j)} - \mu_{\mathbf{Y}}) \gg 1$ . Then, the analysis map  $T_{\nu}$  reduces to

$$T_{\nu}(y, x) = \mu_{\mathbf{X}} + C_{\mathbf{X}, \mathbf{Y}} C_{\mathbf{Y}}^{-1} (y - \mu_{\mathbf{Y}}).$$

## Takeaways:

The outlying observation  $y^{(j)}$  vanishes from the analysis map.

The prior sample  $x^{(j)}$  is mapped to the posterior mean/median/mode.

## What is the pushforward of $\pi_{\mathbf{X}}$ by the maps $T_{\nu}$ and $T_{\text{KF}}$ ?

$T_{\nu}$  leads to exact inference:  $T_{\nu\#}\pi_{\mathbf{X}} = \pi_{\mathbf{X}jy^*} = \text{St}(\boldsymbol{\mu}_{\mathbf{X}jy^*}, C_{\mathbf{X}jy^*}, \nu_{\mathbf{X}jy^*})$ :

$$\begin{aligned}\boldsymbol{\mu}_{\mathbf{X}jy^*} &= \boldsymbol{\mu}_{\mathbf{X}} + C_{\mathbf{X},\mathbf{Y}}C_{\mathbf{Y}}^{-1}(y^* \quad \boldsymbol{\mu}_{\mathbf{Y}}), \\ C_{\mathbf{X}jy^*} &= \alpha_{\mathbf{Y}}(y^*)C_{\mathbf{X}|\mathbf{Y}}, \\ \nu_{\mathbf{X}jy^*} &= \nu + d,\end{aligned}\tag{1}$$

We interpret  $\alpha_{\mathbf{Y}}(y^*)$  as an **adaptive** and **data-dependent** multiplicative inflation.

$T_{\text{KF}\#}\pi_{\mathbf{X}} = \text{St}(\boldsymbol{\mu}_{T_{\text{KF}\#}\pi_{\mathbf{X}}}, C_{T_{\text{KF}\#}\pi_{\mathbf{X}}}, \nu_{T_{\text{KF}\#}\pi_{\mathbf{X}}})$  with

$$\begin{aligned}\boldsymbol{\mu}_{T_{\text{KF}\#}\pi_{\mathbf{X}}} &= \boldsymbol{\mu}_{\mathbf{X}} + C_{\mathbf{X},\mathbf{Y}}C_{\mathbf{Y}}^{-1}(y^* \quad \boldsymbol{\mu}_{\mathbf{Y}}), \\ C_{T_{\text{KF}\#}\pi_{\mathbf{X}}} &= \mathbf{1} \quad C_{\mathbf{X}|\mathbf{Y}}, \\ \nu_{T_{\text{KF}\#}\pi_{\mathbf{X}}} &= \nu + \mathbf{0}.\end{aligned}\tag{2}$$

The Kalman filter is only consistent to estimate the mean value of  $t$ -distributions.

# Challenges of estimating heavy-tailed distributions from samples

The classical sample mean and covariance estimators are derived from a maximum likelihood approach for Gaussian distributions.

These light-tailed estimators are very sensitive to outliers and introduce additional variance.

We use a “regularized” expectation-maximization algorithm (EMq) to estimate the heavy-tailed joint forecast distribution  $\pi_{\mathbf{Y}, \mathbf{X}}$  from  $f(y^{(i)}, x^{(i)})g$  [Doğru et al., 2018].

We call **ensemble robust filter (EnRF)** the ensemble filter that estimates  $T_\nu$  with the EMq from the joint forecast samples  $f(y^{(i)}, x^{(i)})g$ .

# Empirical performance of light-tailed and heavy-tailed estimators


Consider  $\begin{bmatrix} \mathbf{Y} \\ \mathbf{X} \end{bmatrix} \sim St \left( \begin{bmatrix} \mu_{\mathbf{X}} \\ \mu_{\mathbf{Y}} \end{bmatrix}, \begin{bmatrix} C_{\mathbf{Y}} & C_{\mathbf{X},\mathbf{Y}}^> \\ C_{\mathbf{X},\mathbf{Y}} & C_{\mathbf{X}} \end{bmatrix}, \nu \right)$  with  $\nu = 2.1$ ,  $\mathbf{X} \in \mathbb{R}^{10}$  and  $\mathbf{Y} \in \mathbb{R}^5$ .

Experiment:


1. We generate  $M$  samples  $f(y^{(i)}, x^{(i)})g \sim \pi_{\mathbf{Y},\mathbf{X}}$ .
2. We apply the analysis map of the EnKF and the EnRF to assimilate a realization  $y^* \sim \pi_{\mathbf{Y}}$ .
3. We compute the sample mean and covariance for the two posterior ensembles.



# Posterior mean estimates




**Left panel:** Evolution of the error RMSE  $\|\hat{\mu}_{filter} - \mu_{X|y^*}\| / \bar{n}$  with the ensemble size  $M$ .




**Right panel:** Empirical distribution of the RMSE over 1000 realizations of  $y^*$  using  $M = 600$  samples.

Both filters are consistent,  
but the sEnKF has a much slower convergence rate than the EnRF.

# Posterior covariance estimates



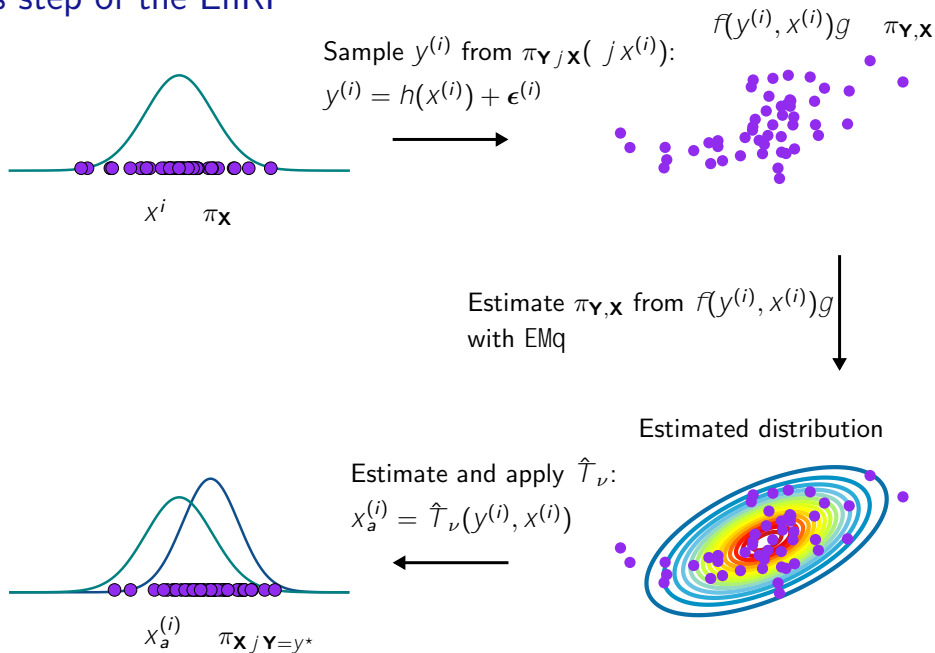
**Left panel:** Evolution of the error  $\|\hat{\Sigma}_{filter} - \Sigma_{x_j y^*}\|_F / \bar{n}$  with the ensemble size  $M$ .



**Right panel:** Empirical distribution of the error  $\|\hat{\Sigma}_{filter} - \Sigma_{x_j y^*}\|_F / \bar{n}$  over 1000 realizations of  $y^*$  using  $M = 600$ .

The sEnKF is not consistent to estimate covariances of  $t$ -distributions.

# Analysis step of the EnRF



## Estimate the tail-heaviness on the fly

We can estimate the dof at each assimilation cycle with the EMq


The EnRF can adapt its analysis map to the tail heaviness of the data!

Limitation: Computational cost of the EMq when the dof is unknown.

**Idea:** Present 3 variants of the EnRF that differ in the frequency and the samples used to estimate the dof


1. The RefreshEnRF: We maintain a buffer of past filtering samples, and estimate the dof with this larger ensemble every  $\Delta t_{refresh}$ .
2. The FixedEnRF: The dof is estimated once from a free-run of the state-space model and fixed for the entire assimilation ( $\Delta t_{refresh} = 1$ )
3. The AdaptEnRF: The dof is estimated at each assimilation cycle ( $\Delta t_{refresh} = \Delta t_{obs}$ )

# Lorenz-63



Empirical degree of freedom  $\hat{\nu}(\mathbf{y}_t, \mathbf{x}_t)$  with  $M = 150, 1000$  samples of the joint forecast density  $f(\mathbf{y}_t^{(i)}, \mathbf{x}_t^{(i)})g_{\pi(\mathbf{y}_t, \mathbf{x}_t) | \mathbf{Y}_{1:t-1}}$  for the Lorenz-63 problem with  $t$ -distributed observation noise.

## RMSE results for Lorenz-63



■ sEnKF   ■ FixedEnRF   ■ RefreshEnRF   ■ AdaptEnRF

Evolution of the RMSE with the ensemble size  $M$  for the Lorenz-63 model with  $t$ -distributed observation noise with  $\nu = 3.0$ .

We optimally tune the multiplicative inflation of the sEnKF.

25% reduction of the RMSE with the EnRF!

**The different EnRFs don't require tuning: Plug and Play!**

# Conclusion and Outlook

## Summary:


We introduce the EnRF that generalizes the EnKF to heavy-tailed  $t$ -distributions

The EnRF adapts its prior-to-posterior to the tail-heaviness of the data

⇒) **Adaptive** and **data-dependent** multiplicative **inflation**

The EnRF requires **no tuning**: **Plug and Play!**

**Software:** The EnRF will be soon available in `TransportBasedInference.jl`



**Acknowledgements:** Support from National Science Foundation is gratefully acknowledged.

**Contact:** mathieuleprovost1@gmail.com

# References I

-  Baptista, R., Zahm, O., and Marzouk, Y. (2020).  
An adaptive transport framework for joint and conditional density estimation.  
*arXiv preprint arXiv:2009.10303*.
-  Dođru, F. Z., Bulut, Y. M., and Arslan, O. (2018).  
Doubly reweighted estimators for the parameters of the multivariate t-distribution.  
*Communications in Statistics-Theory and Methods*, 47(19):4751–4771.
-  Evensen, G. (1994).  
Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics.  
*Journal of Geophysical Research: Oceans*, 99(C5):10143–10162.
-  Marzouk, Y., Moselhy, T., Parno, M., and Spantini, A. (2016).  
Sampling via measure transport: An introduction.  
*Handbook of Uncertainty Quantification*, pages 1–41.
-  Spantini, A., Baptista, R., and Marzouk, Y. (2022).  
Coupling Techniques for Nonlinear Ensemble Filtering.  
*SIAM Review*, 64(4):921–953.  
Publisher: Society for Industrial and Applied Mathematics.