

Design of Gradient Flows for Sampling

Yifan Chen

Courant Institute, New York University

YMSC CAM seminar, Nov 2023


The Paper

[Chen, Huang, Huang, Reich, Stuart 2023]

Sampling via Gradient Flows in the Space of Probability Measures




Daniel Huang
Peking University



Jiaoyang Huang
University of
Pennsylvania



Sebastian Reich
University of
Potsdam



Andrew Stuart
Caltech

Link: <https://arxiv.org/abs/2310.03597>.

Context

The sampling problem

Goal: draw (approximate) samples from

$$\rho^*(\theta) \propto \exp(-V(\theta))$$

Context

The sampling problem

Goal: draw (approximate) samples from

$$\rho^*(\theta) \propto \exp(-V(\theta))$$

*Set-up: $V(\theta)$ **available**, rather than samples in generative modeling*

Context

The sampling problem


Goal: draw (approximate) samples from

$$\rho^*(\theta) \propto \exp(-V(\theta))$$

Set-up: $V(\theta)$ **available**, rather than samples in generative modeling

Many applications in

- Uncertainty quantification
- Bayes inverse problems
- Filtering
- Active learning
- ...



Dynamics for sampling

Idea: construct a **dynamics of** ρ_t that gradually converges to

$$\rho^*(\theta) \propto \exp(-V(\theta))$$

Note: for simplicity we consider continuous-time

Dynamics for sampling

Idea: construct a **dynamics of** ρ_t that gradually converges to

$$\rho^*(\theta) \propto \exp(-V(\theta))$$

Note: for simplicity we consider continuous-time

- **Finite time dynamics** $\rho_1 = \rho^*$, from a given ρ_0 (e.g. prior)
 - Sequential Monte Carlo, e.g., $\rho_t \propto \exp(-tV(\theta))$, ...

Dynamics for sampling

Idea: construct a **dynamics of ρ_t** that gradually converges to

$$\rho^*(\theta) \propto \exp(-V(\theta))$$

Note: for simplicity we consider continuous-time

- **Finite time dynamics** $\rho_1 = \rho^*$, from a given ρ_0 (e.g. prior)
 - Sequential Monte Carlo, e.g., $\rho_t \propto \exp(-tV(\theta))$, ...
- **Infinite time dynamics** $\rho_\infty = \rho^*$, from arbitrary ρ_0
 - MCMC, Langevin's dynamics, ...

Dynamics through Gradient Flows (GFs)

Gradient flow dynamics for sampling

Idea: construct a **gradient flow dynamics of ρ_t** that converges to

$$\rho^*(\theta) \propto \exp(-V(\theta))$$

Namely, dynamics comes from gradient based optimization methods

- Langevin's dynamics and Wasserstein GFs
[Jordan, Kinderlehrer, Otto 1998], ...
- Stein variational GD and Stein variational GFs
[Liu, Wang 2016], [Liu 2017], ...
- Interaction between optimization and sampling
[Wibisono 2018], ...
- A recent review paper
[Trillos, Hosseini, Sanz-Alonso 2023]
- ...

Gradient Flows

Ingredients in gradient flows

Formally: (\mathcal{P} is the space of probability densities)

- **An energy functional** to minimize

$$\mathcal{E} : \mathcal{P} \rightarrow \mathbb{R}$$

- **A metric** for descent direction

$$g_\rho : T_\rho \mathcal{P} \times T_\rho \mathcal{P} \rightarrow \mathbb{R}, \quad g_\rho(\sigma_1, \sigma_2) = \langle M(\rho)\sigma_1, \sigma_2 \rangle_{L^2}$$

$$\implies \text{Flow: } \frac{\partial \rho_t}{\partial t} = -\nabla_g \mathcal{E}(\rho_t) = -M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}$$

- $T_\rho \mathcal{P}$ (tangent space) is the space of measures integrated to 0
- $\frac{\delta \mathcal{E}}{\delta \rho}$ is the first variation of \mathcal{E} at ρ
- $M(\rho_t)^{-1}$ can be understood as a **preconditioner**

Sampling through Numerical Approximation of GFs

Gradient flow equation

$$\frac{\partial \rho_t}{\partial t} = -M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}$$

Numerical approximations of GFs then lead to sampling methods

- Particle methods such as mean field SDEs

$$d\theta_t = f(\theta_t; \rho_t, \rho^*)dt + h(\theta_t; \rho_t, \rho^*)dW_t$$

Sampling through Numerical Approximation of GFs

Gradient flow equation

$$\frac{\partial \rho_t}{\partial t} = -M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}$$

Numerical approximations of GFs then lead to sampling methods

- Particle methods such as mean field SDEs

$$d\theta_t = f(\theta_t; \rho_t, \rho^*)dt + h(\theta_t; \rho_t, \rho^*)dW_t$$

Example: Langevin's dynamics

- $\mathcal{E}(\rho) = \text{KL}[\rho \parallel \rho^*]$ and $M(\rho)^{-1} = -\nabla \cdot (\rho \nabla \cdot)$
- The PDE: $\frac{\partial \rho_t}{\partial t} + \nabla \cdot (\rho_t \nabla_{\theta} \log \rho^*) = \Delta \rho_t$
- The SDE: $d\theta_t = \nabla_{\theta} \log \rho^*(\theta_t)dt + \sqrt{2}dW_t$

The Focus of this Talk

The question:

Any guiding principles for designing \mathcal{E} and $M(\rho)$?

The Focus of this Talk

The question:

Any guiding principles for designing \mathcal{E} and $M(\rho)$?

Some desired properties of gradient flows for sampling:

- **Numerical approximation is tractable**
 - e.g. Langevin dynamics $d\theta_t = \nabla_{\theta} \log \rho^*(\theta_t) dt + \sqrt{2} dW_t$ does not need the normalization constants of ρ^*
- **Fast convergence of the flow**
 - e.g. Langevin dynamics converges fast for well conditioned, single mode distributions, but may struggle for others

Design of Gradient Flows for Sampling

- 1 On Choosing Energy Functionals: KL is Special
- 2 On Choosing Metrics: Fisher-Rao is Special
- 3 On Numerical Approximations by Gaussians and Mixtures

On Choosing the Energy Functionals

Recap: Gradient flow equation

$$\frac{\partial \rho_t}{\partial t} = -M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}$$

- Most popular choice of $\mathcal{E}(\rho)$: Kullback–Leibler divergence

$$\mathcal{E}(\rho; \rho^*) = \text{KL}[\rho \parallel \rho^*] = \int \rho \log \left(\frac{\rho}{\rho^*} \right) d\theta$$

On Choosing the Energy Functionals

Recap: Gradient flow equation

$$\frac{\partial \rho_t}{\partial t} = -M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}$$

- Most popular choice of $\mathcal{E}(\rho)$: Kullback–Leibler divergence

$$\mathcal{E}(\rho; \rho^\star) = \text{KL}[\rho \parallel \rho^\star] = \int \rho \log \left(\frac{\rho}{\rho^\star} \right) d\theta$$

- Property: $\mathcal{E}(\rho; c\rho^\star) = \mathcal{E}(\rho; \rho^\star) - \log c$ for any $c \in \mathbb{R}_+$

On Choosing the Energy Functionals

Recap: Gradient flow equation

$$\frac{\partial \rho_t}{\partial t} = -M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}$$

- Most popular choice of $\mathcal{E}(\rho)$: Kullback–Leibler divergence

$$\mathcal{E}(\rho; \rho^\star) = \text{KL}[\rho \parallel \rho^\star] = \int \rho \log \left(\frac{\rho}{\rho^\star} \right) d\theta$$

- Property: $\mathcal{E}(\rho; c\rho^\star) = \mathcal{E}(\rho; \rho^\star) - \log c$ for any $c \in \mathbb{R}_+$
 \Rightarrow first variation $\frac{\delta \mathcal{E}}{\delta \rho}$ is independent of c

On Choosing the Energy Functionals

Recap: Gradient flow equation

$$\frac{\partial \rho_t}{\partial t} = -M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}$$

- Most popular choice of $\mathcal{E}(\rho)$: Kullback–Leibler divergence

$$\mathcal{E}(\rho; \rho^\star) = \text{KL}[\rho \parallel \rho^\star] = \int \rho \log \left(\frac{\rho}{\rho^\star} \right) d\theta$$

- Property: $\mathcal{E}(\rho; c\rho^\star) = \mathcal{E}(\rho; \rho^\star) - \log c$ for any $c \in \mathbb{R}_+$
 - \Rightarrow first variation $\frac{\delta \mathcal{E}}{\delta \rho}$ is independent of c
 - \Rightarrow the gradient flow equation is independent of c

On Choosing the Energy Functionals

Recap: Gradient flow equation

$$\frac{\partial \rho_t}{\partial t} = -M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}$$

- Most popular choice of $\mathcal{E}(\rho)$: Kullback–Leibler divergence

$$\mathcal{E}(\rho; \rho^*) = \text{KL}[\rho \parallel \rho^*] = \int \rho \log \left(\frac{\rho}{\rho^*} \right) d\theta$$

- Property: $\mathcal{E}(\rho; c\rho^*) = \mathcal{E}(\rho; \rho^*) - \log c$ for any $c \in \mathbb{R}_+$
 - \Rightarrow first variation $\frac{\delta \mathcal{E}}{\delta \rho}$ is independent of c
 - \Rightarrow the gradient flow equation is independent of c

Implication: no need to worry about **normalization consts** of ρ^*

The question

Any other choices of \mathcal{E} that have such invariance property?

The question

Any other choices of \mathcal{E} that have such invariance property?

The answer is **NO** among a large class of \mathcal{E}

KL Divergence is Special

Theorem [Chen, Huang, Huang, Reich, Stuart 2023]

Among all **f -divergence** with continuously differentiable f , KL divergence is the only one, up to scaling, whose first variation is **invariant to the normalization const of ρ^***

- f -divergence: for $f(0) = 1$ and f convex

$$D_f[\rho||\rho^*] = \int \rho^* f\left(\frac{\rho}{\rho^*}\right) d\theta$$

- Kullback–Leibler divergence: $f(x) = x \log x$
- χ^2 divergence: $f(x) = (x - 1)^2$
- Hellinger distance: $f(x) = (\sqrt{x} - 1)^2$
- ...

KL Divergence is Special

Theorem [Chen, Huang, Huang, Reich, Stuart 2023]

Among all **f -divergence** with continuously differentiable f , KL divergence is the only one, up to scaling, whose first variation is **invariant to the normalization const of ρ^***

- f -divergence: for $f(0) = 1$ and f convex

$$D_f[\rho \parallel \rho^*] = \int \rho^* f\left(\frac{\rho}{\rho^*}\right) d\theta$$

- Kullback–Leibler divergence: $f(x) = x \log x$
- χ^2 divergence: $f(x) = (x - 1)^2$
- Hellinger distance: $f(x) = (\sqrt{x} - 1)^2$
- ...

Use KL divergence from now on

Design of Gradient Flows for Sampling

- 1 On Choosing Energy Functionals: KL is Special
- 2 On Choosing Metrics: Fisher-Rao is Special
- 3 On Numerical Approximations by Gaussians and Mixtures

Two Metrics

Wasserstein metric [Jordan, Kinderlehrer, Otto 1998]

$$\text{Metric: } M(\rho)^{-1}\psi = -\nabla \cdot (\rho \nabla \psi)$$

$$\text{Flow: } \frac{\partial \rho_t}{\partial t} = -\nabla_{\theta} \cdot (\rho_t \nabla_{\theta} \log \rho^*) + \nabla \cdot (\nabla \rho_t)$$

$$\text{SDEs: } d\theta_t = \nabla_{\theta} \log \rho^* dt + \sqrt{2}dW_t$$

- Optimal transport [Villani 2003, 2008]

Fisher-Rao metric [Rao 1945]

$$\text{Metric: } M(\rho)^{-1}\psi = \rho(\psi - \mathbb{E}_{\rho}[\psi])$$

$$\text{Flow: } \frac{\partial \rho_t}{\partial t} = \rho_t(\log \rho^* - \log \rho_t) - \rho_t \mathbb{E}_{\rho_t}[\log \rho^* - \log \rho_t]$$

- Information geometry [Amari 2016], [Ay, Jost, Lê, Schwachhöfer, 2017]

Convergence Property of Wasserstein Gradient Flow

Theorem [Markowich, Villani 2000]

Assume $\exists \lambda > 0$ such that

$$-D^2 \log \rho^*(\cdot) \succeq \lambda I$$

Then, for all $t \geq 0$,

$$\text{KL}[\rho_t \|\rho^*] \leq \text{KL}[\rho_0 \|\rho^*] e^{-2\lambda t}$$

Rate of exponential convergence depends on ρ^*

Two Metrics

Wasserstein metric [Jordan, Kinderlehrer, Otto 1998]

$$\text{Metric: } M(\rho)^{-1}\psi = -\nabla \cdot (\rho \nabla \psi)$$

$$\text{Flow: } \frac{\partial \rho_t}{\partial t} = -\nabla_{\theta} \cdot (\rho_t \nabla_{\theta} \log \rho^*) + \nabla \cdot (\nabla \rho_t)$$

$$\text{SDEs: } d\theta_t = \nabla_{\theta} \log \rho^* dt + \sqrt{2}dW_t$$

- Optimal transport [Villani 2003, 2008]

Fisher-Rao metric [Rao 1945]

$$\text{Metric: } M(\rho)^{-1}\psi = \rho(\psi - \mathbb{E}_{\rho}[\psi])$$

$$\text{Flow: } \frac{\partial \rho_t}{\partial t} = \rho_t(\log \rho^* - \log \rho_t) - \rho_t \mathbb{E}_{\rho_t}[\log \rho^* - \log \rho_t]$$

- Information geometry [Amari 2016], [Ay, Jost, Lê, Schwachhöfer, 2017]

A Closer Look at Fisher-Rao

Fisher-Rao gradient flow

$$\frac{\partial \rho_t}{\partial t} = \rho_t (\log \rho^* - \log \rho_t) - \rho_t \mathbb{E}_{\rho_t} [\log \rho^* - \log \rho_t]$$

Apply transformation of any **diffeomorphism** $\varphi : \mathbb{R}^{d_\theta} \rightarrow \mathbb{R}^{d_\theta}$

- $\tilde{\rho}_t = \varphi \# \rho_t$ is the transformed distribution at time t
- $\tilde{\rho}^* = \varphi \# \rho^*$ is the transformed target distribution

Then, the form of the flow equation remains **invariant**

$$\frac{\partial \tilde{\rho}_t}{\partial t} = \tilde{\rho}_t (\log \tilde{\rho}^* - \log \tilde{\rho}_t) - \tilde{\rho}_t \mathbb{E}_{\tilde{\rho}_t} [\log \tilde{\rho}^* - \log \tilde{\rho}_t]$$

Why Care About Invariance?

Implication of invariance

Convergence rates of the gradient flow are **the same** for
general ρ^* and **Gaussian** ρ^*

Why Care About Invariance?

Implication of invariance

Convergence rates of the gradient flow are **the same** for general ρ^* and **Gaussian** ρ^*

- Assume there exists a diffeomorphism φ such that

$$\tilde{\rho}^* = \varphi \# \rho^* = \text{Gaussian}$$

- Recall the property of the KL divergence

$$\text{KL}[\rho_t \|\rho^*] = \text{KL}[\varphi \# \rho_t \|\varphi \# \rho^*] = \text{KL}[\tilde{\rho}_t \|\tilde{\rho}^*]$$

Thus, a general ρ^* problem \sim a simpler **Gaussian** ρ^* problem

Convergence of Fisher-Rao gradient flows

[Chen, Huang, Huang, Reich, Stuart 2023]

Let ρ_t satisfy the Fisher-Rao gradient flow. Assume

- there exist constants $K, B > 0$ such that ρ_0 satisfies

$$e^{-K(1+|\theta|^2)} \leq \rho_0(\theta)/\rho^*(\theta) \leq e^{K(1+|\theta|^2)}$$

- the second moments of ρ_0, ρ^* are both bounded by B

Then, for any $t \geq \log((1+B)K)$,

$$\text{KL}[\rho_t \|\rho^*] \leq (2 + B + eB)Ke^{-t}$$

See also: [Lu, Slepčev, Wang 2022], [Domingo-Enrich, Pooladian 2023]

Convergence of Fisher-Rao gradient flows

[Chen, Huang, Huang, Reich, Stuart 2023]

Let ρ_t satisfy the Fisher-Rao gradient flow. Assume

- there exist constants $K, B > 0$ such that ρ_0 satisfies

$$e^{-K(1+|\theta|^2)} \leq \rho_0(\theta)/\rho^*(\theta) \leq e^{K(1+|\theta|^2)}$$

- the second moments of ρ_0, ρ^* are both bounded by B

Then, for any $t \geq \log((1+B)K)$,

$$\text{KL}[\rho_t \|\rho^*] \leq (2 + B + eB)Ke^{-t}$$

See also: [Lu, Slepčev, Wang 2022], [Domingo-Enrich, Pooladian 2023]

“Unconditional” uniform exponential convergence

- In sharp contrast to Wasserstein gradient flows whose convergence rates depend on ρ^*

Does This Mean Fisher-Rao is All You Need?

Numerical approximations of Fisher-Rao GFs can be tricky

Particle methods (i.e. Diracs ansatz)

- Birth-death dynamics [Lu, Lu, Nolen 2019], [Lu, Slepčev, Wang 2022]
- Ensemble MCMC [Lindsey, Weare, Zhang 2021]

Need ways to move the support of the particles to explore the space and choices of smoothing kernels

Does This Mean Fisher-Rao is All You Need?

Numerical approximations of Fisher-Rao GFs can be tricky

Particle methods (i.e. Diracs ansatz)

- Birth-death dynamics [Lu, Lu, Nolen 2019], [Lu, Slepčev, Wang 2022]
- Ensemble MCMC [Lindsey, Weare, Zhang 2021]

Need ways to move the support of the particles to explore the space and choices of smoothing kernels

Before going that far, let us first ask a **basic question**

The question:

Any other choices of metric having such invariance property?

Does This Mean Fisher-Rao is All You Need?

Numerical approximations of Fisher-Rao GFs can be tricky

Particle methods (i.e. Diracs ansatz)

- Birth-death dynamics [Lu, Lu, Nolen 2019], [Lu, Slepčev, Wang 2022]
- Ensemble MCMC [Lindsey, Weare, Zhang 2021]

Need ways to move the support of the particles to explore the space and choices of smoothing kernels

Before going that far, let us first ask a **basic question**

The question:

Any other choices of metric having such invariance property?

The answer is again, **NO**

Fisher-Rao Metric is Special

Unique property of Fisher-Rao metric

[Cencov 2000], [Ay, Jost, L e, Schwachh ofer 2015], [Bauer, Bruveris, Michor 2016]

The Fisher-Rao metric is the **only Riemannian metric on smooth positive densities** (up to scaling) that is invariant under any diffeomorphism of the parameter space.

*No other alternatives if we ask for **diffeomorphism invariance!***

Fisher-Rao Metric is Special

Unique property of Fisher-Rao metric

[Cencov 2000], [Ay, Jost, Lê, Schwachhöfer 2015], [Bauer, Bruveris, Michor 2016]

The Fisher-Rao metric is the **only Riemannian metric on smooth positive densities** (up to scaling) that is invariant under any diffeomorphism of the parameter space.

*No other alternatives if we ask for **diffeomorphism invariance!***

But can ask for a relaxed **affine invariance**

- Affine invariant MCMC [Goodman, Weare 2010]
- Preconditioned Langevin and Kalman-Wasserstein GFs
[Reich Cotter 2015], [Leimkuhler, Matthews, Weare 2018], [Garbuno-Inigo, Hoffmann, Li, Stuart 2020]
- Other affine invariant gradient flow examples in the paper
 - e.g., affine invariant Stein gradient flow

Design of Gradient Flows for Sampling

- 1 On Choosing Energy Functionals: KL is Special
- 2 On Choosing Metrics: Fisher-Rao is Special
- 3 On Numerical Approximations by Gaussians and Mixtures**

Numerical Approximation of the Fisher-Rao Gradient Flow

Particle methods (i.e. Diracs ansatz)

- Birth-death dynamics [Lu, Lu, Nolen 2019], [Lu, Slepčev, Wang 2022]
- Ensemble MCMC [Lindsey, Weare, Zhang 2021]

Need ways to move the support of the particles to explore the space and choices of smoothing kernels

Our focus: **parametric approximation** (full support ansatz)

- Gaussian approximations
- Gaussian mixtures for multimodal problems

Gaussian Approximation by Moment Closures

The general procedures:

- Consider any dynamics in the **density space**

$$\frac{\partial \rho_t(\theta)}{\partial t} = \sigma_t(\theta, \rho_t)$$

- Write down the dynamics of the **mean and covariance**

$$\begin{aligned}\frac{dm_t}{dt} &= \int \sigma_t(\theta, \rho_t) \theta d\theta \\ \frac{dC_t}{dt} &= \int \sigma_t(\theta, \rho_t) (\theta - m_t)(\theta - m_t)^T d\theta\end{aligned}$$

- Closure: replace ρ_t in the above RHS by $\rho_{a_t} = \mathcal{N}(m_t, C_t)$
Notation: $a_t = (m_t, C_t)$

References: Moment closure in variational Kalman filtering [Särkkä, 2007], and in Wasserstein gradient flow [Lambert, Chewi, Bach, Bonnabel, Rigollet 2022]

Gaussian Approximation by Moment Closures

Gaussian approximate Fisher-Rao gradient flow

$$\begin{aligned}\frac{dm_t}{dt} &= C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \log \rho^*], \\ \frac{dC_t}{dt} &= C_t + C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho^*] C_t\end{aligned}$$

- Derived using Stein's lemma
- Equivalent to **natural gradient** flow [Amari 1998] for

Gaussian variational inference: $\min_{m, C} \text{KL}[\mathcal{N}(m, C) \parallel \rho^*]$

Key: Fisher information matrix is used for preconditioning

$$\frac{d}{dt}(m_t, C_t) = -\text{FI}(m_t, C_t)^{-1} \nabla_{m_t, C_t} \text{KL}$$

Convergence Guarantee

Gaussian target

If $\rho^\star = \mathcal{N}(m_\star, C_\star)$, and $C_0 = \lambda_0 I$, $\lambda_0 > 0$, then

$$\|m_t - m_\star\|_2 = \mathcal{O}(e^{-t}), \quad \|C_t - C_\star\|_2 = \mathcal{O}(e^{-t})$$

- Same story, due to **invariance** property

Convergence Guarantee

Logconcave target [Chen, Huang, Huang, Reich, Stuart 2023]

Assume

- $\alpha I \preceq -\nabla_{\theta} \nabla_{\theta} \log \rho^* \preceq \beta I$
- $\lambda_{0,\min} I \preceq C_0 \preceq \lambda_{0,\max} I$

Then

$$\text{KL}[\rho_{a_t} \|\rho^*] - \text{KL}[\rho_{a_*} \|\rho^*] \leq e^{-Kt} (\text{KL}[\rho_{a_0} \|\rho^*] - \text{KL}[\rho_{a_*} \|\rho^*])$$

where

- $a_t = (m_t, C_t), \rho_{a_t} = \mathcal{N}(m_t, C_t)$
- $a_* = \operatorname{argmin}_a \text{KL}[\rho_a \|\rho^*]$
- $K = \alpha \min\{1/\beta, \lambda_{0,\min}\}$

- Inspired by the proof for the Wasserstein gradient flow in Gaussian variational inference for logconcave target

[Lambert, Chewi, Bach, Bonnabel, Rigollet 2022]

Local Convergence Rates: Linearized Analysis

Theorem [Chen, Huang, Huang, Reich, Stuart 2023]

Assume $\alpha I \preceq -\nabla_{\theta} \nabla_{\theta} \log \rho^* \preceq \beta I$. For $N_{\theta} = 1$, let $\lambda_{\star, \max} < 0$ denote the largest eigenvalue of the linearized Jacobian matrix of the flow around a_{\star} . Then we have

$$-\lambda_{\star, \max} \geq \frac{1}{\left(7 + \frac{4}{\sqrt{\pi}}\right) \left(1 + \log\left(\frac{\beta}{\alpha}\right)\right)}$$

Moreover, the bound is sharp: it is possible to construct a sequence of triplets ρ_n^* , α_n and β_n , where $\lim_{n \rightarrow \infty} \frac{\beta_n}{\alpha_n} = \infty$, such that, if we let $\lambda_{\star, \max, n}$ denote the corresponding largest eigenvalues of the linearized Jacobian matrix for the n -th triple, then, it holds that

$$-\lambda_{\star, \max, n} = \mathcal{O}\left(1 / \log \frac{\beta_n}{\alpha_n}\right)$$

Convergence rates only depend on $\log(\text{condition number})$

Numerical Examples

- **2D Convex Potential:** $\theta = (\theta^{(1)}, \theta^{(2)})$

$$V(\theta) = \frac{(\sqrt{\lambda}\theta^{(1)} - \theta^{(2)})^2}{20} + \frac{(\theta^{(2)})^4}{20} \quad \text{with } \lambda = 0.01, 0.1, 1$$

- **Method:** Gaussian approximation of Fisher-Rao GF, Wasserstein GF and vallina GF
- **Configuration:** we initialize the Gaussian at

$$\mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}\right)$$

We integrate the mean and covariance dynamics to $t = 15$

Numerical Examples

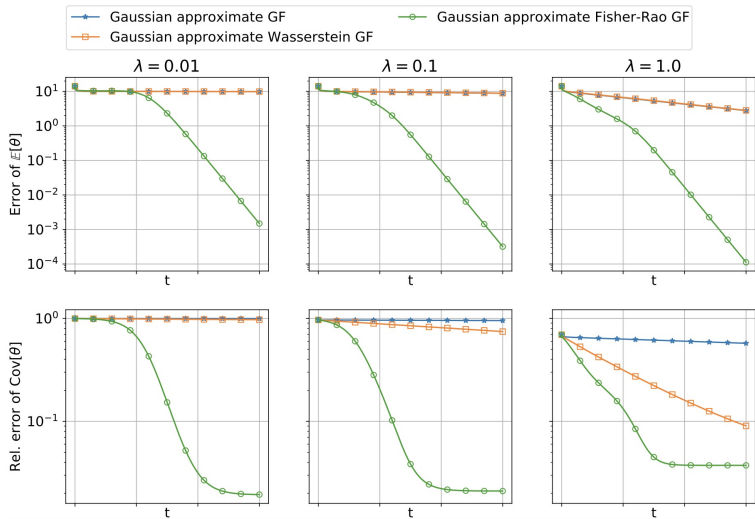


Figure: x axis is from $t = 0$ to 15. Convergence rate of Gaussian approximate Fisher-Rao gradient flows not influenced by values of λ

Ongoing Work: Gaussian Mixtures + Kalman Methods

Consider the 2d Darcy flow problem ($\theta \in \mathbb{R}^{128}$)

$$\begin{aligned} -\nabla \cdot (a(x, \theta) \nabla p(x)) &= f(x) = 1000 \sin(4\pi x_{(2)}) & x \in D \\ p(x) &= 0 & x \in \partial D \end{aligned}$$




Figure: The reference pressure field p and observations $\frac{p([x_{(1)}, x_{(2)}]^T) + p([1-x_{(1)}, x_{(2)}]^T)}{2}$ at 28 equidistant points (solid black dots) and their mirroring points (empty black dots).




Figure: The truth log permeability field $a(x; \theta_{ref})$, and log permeability fields obtained by 3-mode GMKI (Gaussian Mixture Kalman Inversion)




Figure: The truth expansion parameters $\theta_{(i)}$ (black crosses), and mean estimations of $\theta_{(i)}$ for each modes (circles) and the associated marginal distributions obtained GMKI at the **30-th iteration**.

Summary

Gradient flows for sampling

- **Energy functional:** KL divergence is special
 - Invariance to normalization consts
- **Metric:** Fisher-Rao metric is special
 - Invariance to any diffeomorphism of the parameter space
⇒ unconditional uniform exponential convergence
 - Relaxed to affine invariance and many constructions
- **Gaussian approximation**
 - Moment closures = natural gradient in Gaussian VI
 - Convergence guarantee for Gaussian and logconcave targets
- **Further directions**
 - Gaussian mixture approximations for multimodal targets
 - Derivative free approximations via Kalman's methodology
 - Other approximations to sample curved distributions

Thank You!

[Chen, Huang, Huang, Reich, Stuart 2023]

Sampling via Gradient Flows in the Space of Probability Measures

Link: <https://arxiv.org/abs/2310.03597>.