Variational inference via Gaussian interacting particles
in the Bures-Wasserstein geometry

Giacomo Borghi (Heriot–Watt Edinburgh)
José Antonio Carrillo (Oxford)

ICML 2026

Problem:

$$\mu^\star \in \underset{\mu\in\Pi}{\text{argmin}}\, \mathcal{E}(\mu) = \text{KL}(\mu, \rho^{\text{targ}}) \qquad \text{with} \quad \rho^{\text{targ}} \propto e^{-V} \tag{VI}$$

$$\Pi = \mathcal{N}^d := \{\mu = \mathcal{N}(m, \Sigma) \mid m \in \mathbb{R}^d,\, \Sigma \in \text{Sym}_d^+\}$$

State of the art: Wasserstein Gradient Flows restricted to $\mathcal{N}^d$
[Alvarez-Melis et al., 2022; Lambert et al., 2022; Diao et al.,2023]
Hellinger–Kantorovich GFs [Liero et al., 2025b];
Laplace method [Tierney & Kadane, 1986]

Issue: The problem might be non-convex if ρ^{targ} non-longcave

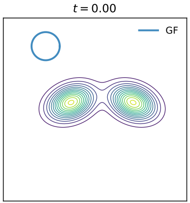

Gaussian Consensus Based Optimization (GaussCBO)

High-level idea of CBO [Pinnau et al. '17]

- Candidate solutions as interacting particles
- Deterministic drift towards **consensus point**
- Stochastic **exploration** of search space

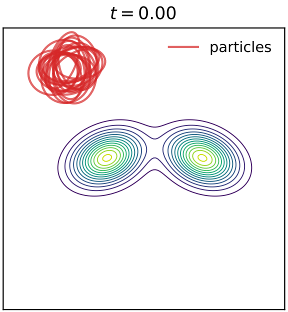

Gaussian version

✓ Gaussian particles $\mu^i = \mathcal{N}(m^i, \Sigma^i)$

✓ **consensus point** $\to$ **barycenter**
via Bures-Wasserstein (BW) Riemannian geometry

? Stochasticity $\to$ BW manifold is NOT geodesically complete

Linearized Bures–Wasserstein (LBW) space $\to$ linearize around $\mu^0 = \mathcal{N}(m^0, \Sigma^0)$

Tangent space at $(m^0, \Sigma^0) \in \mathbb{R}^d \times \mathrm{Sym}_d^{++}$:

$$(m, T) \in \mathbb{R}^d \times \mathrm{Sym}_d \cong \mathbb{R}^{d+d(d+1)/2} \qquad (\mathrm{Sym}_d^{++} \to \mathrm{Sym}_d \checkmark)$$

The particle dynamics in LBW

LBW particles: $(m^i, \Sigma^i) \quad \longrightarrow \quad (m^i, T^i)$ with $\Sigma^i = \exp^{\mathsf{BW}}_{\Sigma^0}(T^i)$

Riemannian exponential: $\exp^{\mathsf{BW}}_{\Sigma^0}(T) = (I+T)\Sigma^0(I+T)$

LBW **weighted barycenter** ($\approx$ best particle)

$$\overline{m}^{\alpha}[\rho^N] = \sum_{i=1}^{N} \omega^i m^i \quad \overline{T}^{\alpha}[\rho^N] = \sum_{i=1}^{N} w^i T^i \qquad \text{with} \quad w^i \propto e^{-\alpha \mathcal{E}(\mu^i)}, \quad \mu^i = \mathcal{N}(m^i, \Sigma^i)$$

$$\begin{cases} \mathrm{d}m_t^i &= \lambda(\overline{m}^{\alpha}[\rho_t^N] - m_t^i)\mathrm{d}t + \sigma(\overline{m}^{\alpha}[\rho_t^N] - m_t^i) \odot \mathrm{d}B_t^{i,m} \qquad i \in [N] \\ \mathrm{d}T_t^i &= \lambda(\overline{T}^{\alpha}[\rho_t^N] - T_t^i)\mathrm{d}t + \sigma(\overline{T}^{\alpha}[\rho_t^N] - T_t^i) \odot \mathrm{d}B_t^{i,T} \qquad i \in [N] \end{cases}$$

$\alpha \gg 1 \quad \lambda > 0$ drift parameter, $\quad \sigma > 0$ exploration parameter

$\rightarrow$ Well-posedness & convergence analysis via **mean-field approximation**

Validation tests

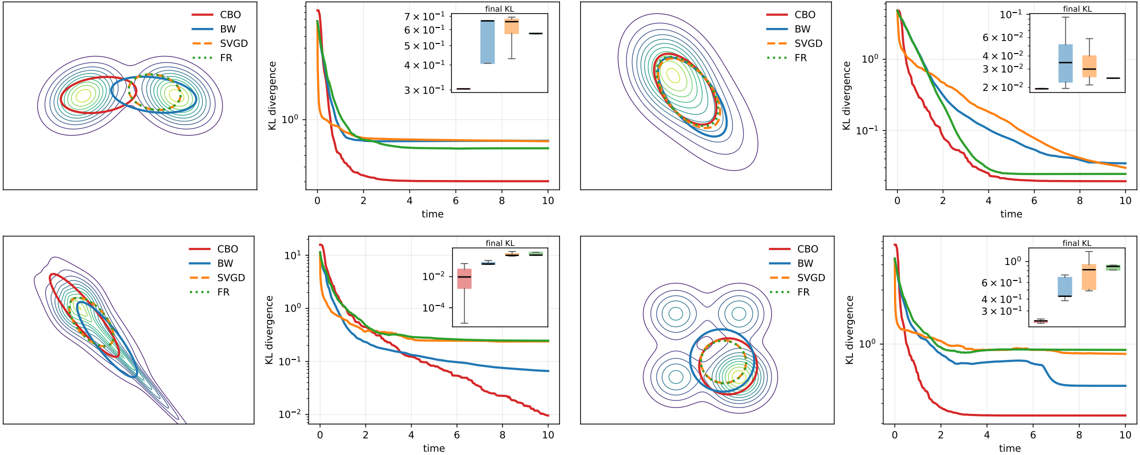

<u>Baselines:</u> Gaussian Wass GF, Gaussian SVGD, Gaussian Fisher–Rao

$\lambda = 1, \sigma = 5, N = 20, \alpha = 10^4$. Reference measure for LBW is $\mathcal{N}(0, I)$

$\rightarrow$ higher dimensional tests, and extension to Gaussian Mixture Models in the paper