Electrostatics-based Particle Sampling & Approx. Inference

Yongchao Huang (yongchao.huang@abdn.ac.uk)

Frontiers in Probabilistic Inference @ ICLR 2025

1495 UNIVERSITY OF ABERDEEN

Summary

We developed a **physics-based** sampling and variational inference method, ***EParVI*** [1], which exploits principles of *electrostatics* and *Newton mechanics* to simulate an *interacting particle system* (IPS). There are two sets of particles in the system: freely-moving negative charges and positive charges at fixed grids, with magnitudes proportional to the (un-normalised) target distribution. Particles interact with each other via *attraction* and *repulsion* induced by the resulting electric fields described by Poisson's equation. The IPS evolves towards a steady-state where the distribution of negative charges conforms to the target distribution.
This physics-inspired method offers **deterministic**, **gradient-free** sampling and inference, achieving comparable performance as other particle-based and MCMC methods in benchmark tasks of inferring low-dimensional, multi-modal densities, Bayesian logistic regression and dynamical system identification.

Background

Probabilistic approaches [2,3] are used in many learning and decision-making tasks, e.g. Bayesian inference and generative modelling (e.g. diffusion models), and many of them involve estimating the distribution of some quantities of interest. When exact inference is infeasible, approximate inference methods such as MCMC and variational inference become relevant. *Particle-based variational inference* (**ParVI**), such as SVGD [4], SMC [5] and EVI [6], have been developed which iteratively transport a set of particles to match the target distribution, leveraging gradient information for efficient inference. In this work, we propose an electrostatics-based particle sampling method inspired by principles from physics, hopefully facilitating more effective exploration of complex distributions.

Methodology

EParVI is inspired from *electrostatic halftoning* [7] - originally used for for image dithering: particles evolve to reach an equilibrium state that forms an accurate halftone representation of the original image. In EParVI, a free-moving negative charge j is repelled by other negative charges and attracted by fixed positive charges. The combined electrostatics force acting on this negative charge is:

$$\mathbf{F}_j^t = \Big[\sum_{i=1, i\neq j}^{M^{neg}} \frac{q_i q_j \Gamma(\frac{d}{2})}{2\pi^{d/2}\varepsilon_0} \frac{\mathbf{x}_j - \mathbf{x}_i}{\|\mathbf{x}_j - \mathbf{x}_i\|_2^{d/2}} - \sum_{i'=1, i'\neq j}^{M^{pos}} \frac{q p(\mathbf{x}_{i'}) q_j \Gamma(\frac{d}{2})}{2\pi^{d/2}\varepsilon_0} \frac{\mathbf{x}_j - \mathbf{x}_{i'}}{\|\mathbf{x}_j - \mathbf{x}_{i'}\|_2^{d/2}} \Big]$$

Under this force, the negative charge updates its position as per Newton's law e.g. $\mathbf{x}_j^{t+1} = \mathbf{x}_j^t + \tau \mathbf{F}_j^t$. At transient or equilibrium time, the configuration of negative charges approximates the target density.

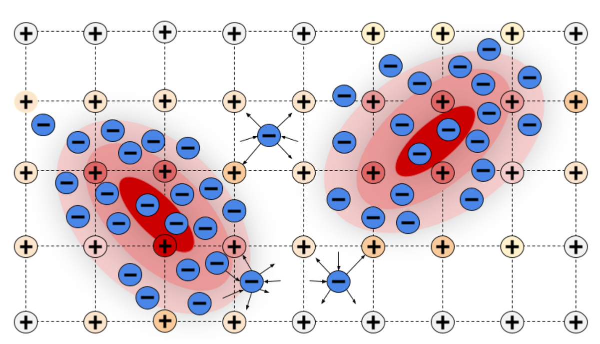

Fig: illustrative diagram of particle distributions. Red darkness indicates the density value of two-Gaussian mixture, arrows denote example forces. Positive charges are fixed at the grid points with magnitude proportional to density value, negative charges can move freely.

Experiments

- Toy densities inference

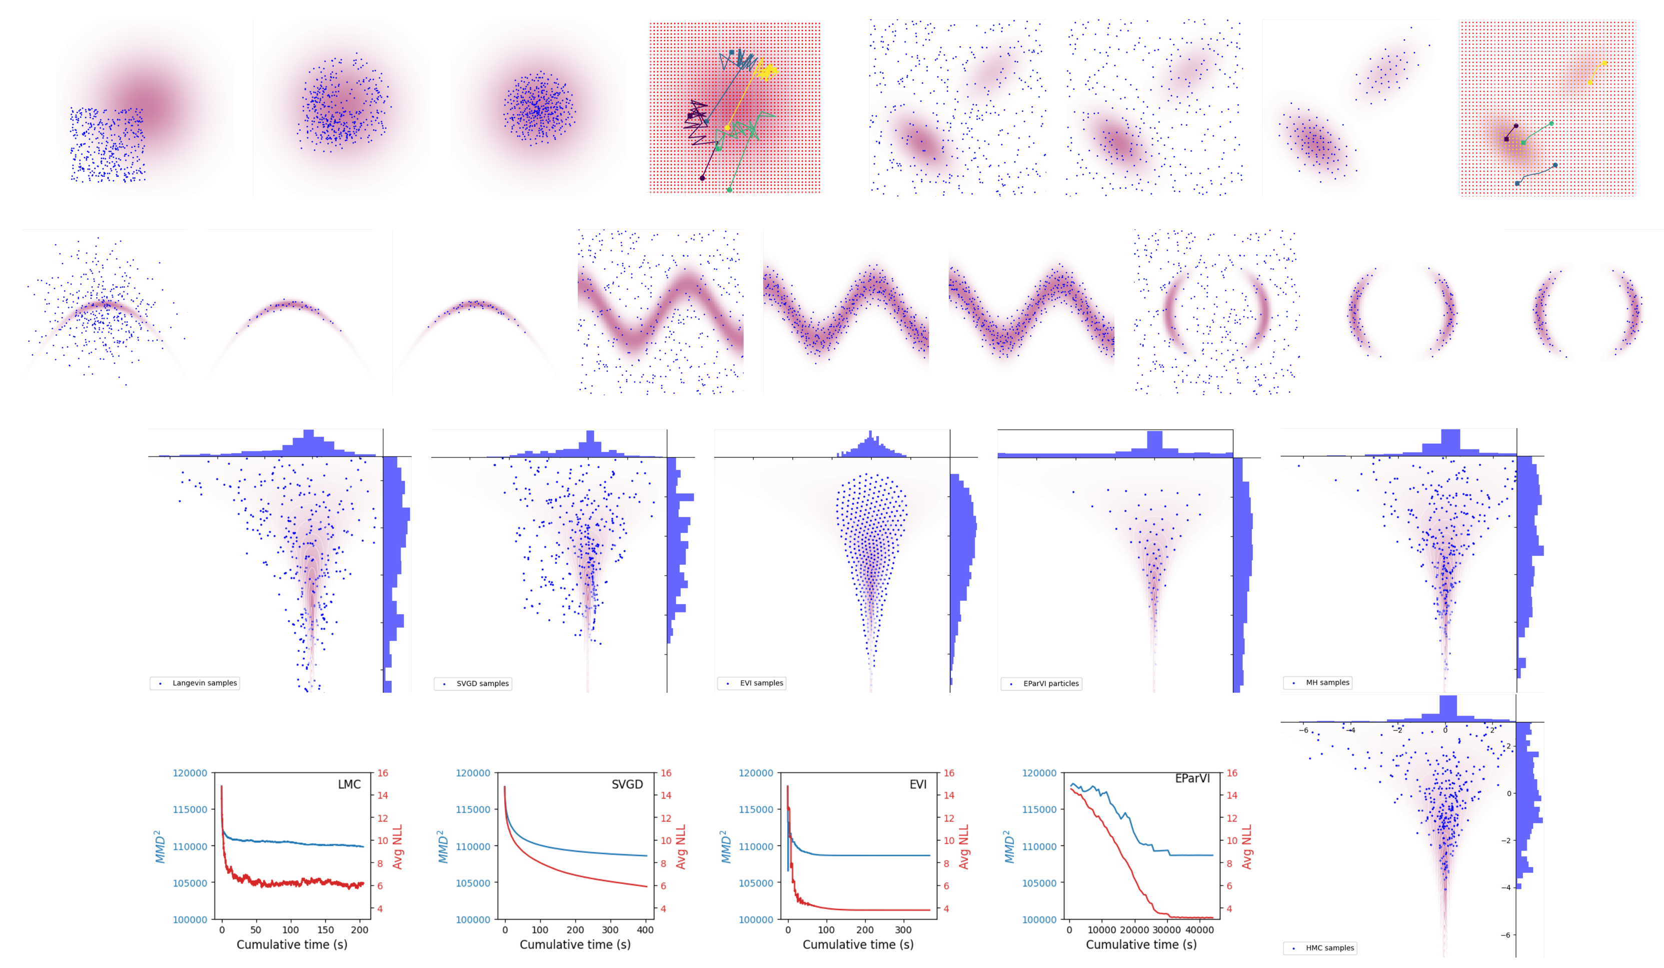

- Bayesian logistic regression: *Iris* classification

$$p(\boldsymbol{\omega}|y, \mathbf{X}) \propto \prod_{i=1}^{N} \left[\frac{1}{1+\exp(-\boldsymbol{\omega}^T \mathbf{x}_i)}\right]^{y_i} \times \left[1 - \frac{1}{1+\exp(-\boldsymbol{\omega}^T \mathbf{x}_i)}\right]^{1-y_i} \times \exp\left(-\frac{1}{2\alpha}\boldsymbol{\omega}^T\boldsymbol{\omega}\right)$$

Method	$\hat{\omega}_1$	$\hat{\omega}_2$	$\hat{\omega}_3$	$\hat{\omega}_4$	NLL/MMD^2	Runtime
MLE-LR	-0.64	1.89	-1.75	-1.67	NA	< 1s
HMC	-0.69	1.98	-1.84	-1.86	NA	< 30s
SVGD	-0.76	1.87	-1.74	-1.67	12.83 / 2.41	< 100s
EVI	-0.78	1.81	-1.63	-1.56	11.39 / 10.14	< 100s
EParVI	-0.64	1.62	-1.50	-1.42	11.32 / 23.25	$> 10^3$s

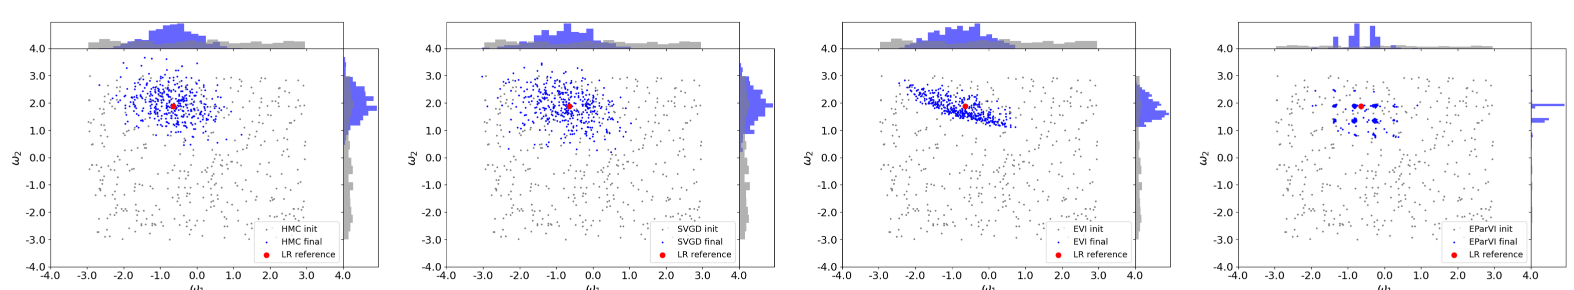

- Dynamical system inference

Conclusion & Future Work

The physics-based ParVI method, EParVI, features *gradient-free sampling with deterministic dynamics*, *simplicity*, *non-parametric flexibility*, *expressiveness* and *good accuracy*.
It also suffers from curse of dimensionality due to its *mesh-based* nature. Future work includes exploring theoretical guarantees of convergence from an energy optimisation perspective, and improving its computational efficiency.

References

[1] Yongchao Huang (2025). *Electrostatics-based Particle Sampling*, https://arxiv.org/abs/2406.20044.
[2] Zoubin Ghahramani (2015). *Probabilistic machine learning and artificial intelligence*, Nature 521, 452–459.
[3] Michael I. Jordan et al. (1999). *An Introduction to Variational Methods for Graphical Models*, Machine Learning 37, 183-233.
[4] Qiang Liu and Dilin Wang (2016). *Stein variational gradient descent: A general purpose Bayesian inference algorithm*, NeurIPS proceeding, 2378–2386.
[5] Arnaud Doucet et al. (2001). *Sequential Monte Carlo Methods in Practice*. Springer, New York.
[6] Yiwei Wang et al. (2021). *Particle-based energetic variational inference*. Statistics and Computing, 31(34).
[7] Christian Schmaltz et al. (2010). *Electrostatic halftoning*. Computer Graphics Forum, 29(8), 2313–2327.