Geometry-Aware Score-Repellent Monte Carlo

Aligning History-Driven MCMC Repulsion with Fisher Geometry

Jie Hu[1] · **Lingyun Chen**[2] · **Do Young Eun**[2]

[1]Oakland University · [2]NC State University · NeurIPS 2026

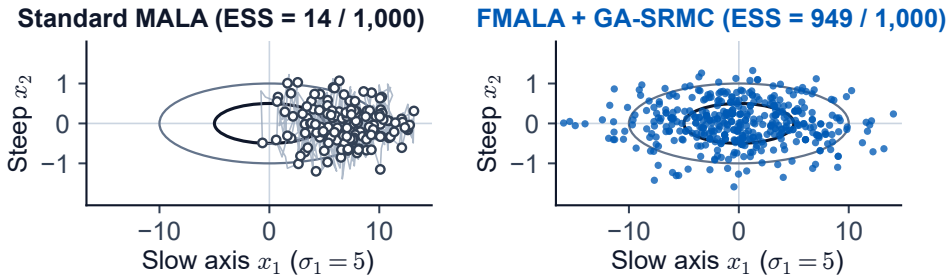

Why Isotropic SRMC[1] Fails Under Anisotropy

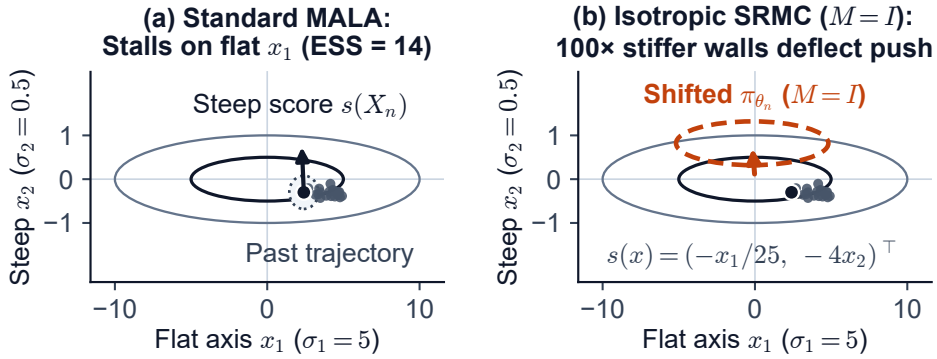

Scalar α dilemma: small α stalls on flat x_1; large α overshoots $100\times$ stiffer x_2.

[1] Hu, J., Chen, L., Kim, G., Choi, J., Han, B. & Eun, D.Y. Score-Repellent Monte Carlo: Toward Efficient Non-Markovian Sampler with Constant Memory in General State Spaces. ICML (**spotlight**), 2026.

Geometry-Aware Surrogate Target

$$\pi_{\theta_n}^{(M)}(x) \;\propto\; \pi(x)\,\exp\!\big(-\alpha\,\theta_n^{\top} M\, s(x)\big) \qquad (M = I_d\text{: SRMC} \;\rightarrow\; M \propto S^{-1}\text{: GA-SRMC})$$

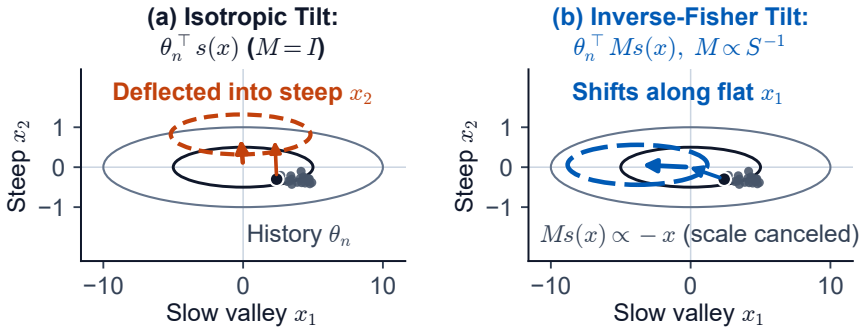

- Linear $M \propto S^{-1}$ cancels curvature mismatch and recovers π as $\theta_n \to 0$.
- Scale-free $\cos(\theta, s)$ or $s/\|s\|$ never switches off as $\theta_n \to 0$, biasing π.

Complementarity: Proposal (FisherMALA[2]) vs. Target

$$Y = X_n + \epsilon A_n \nabla \log \pi_{\theta_n}^{(M_n)}(X_n) + \sqrt{2\epsilon}\, A_n^{1/2}\, \xi_n \quad (A_n\text{: proposal step} \cdot M_n\text{: repellent target})$$

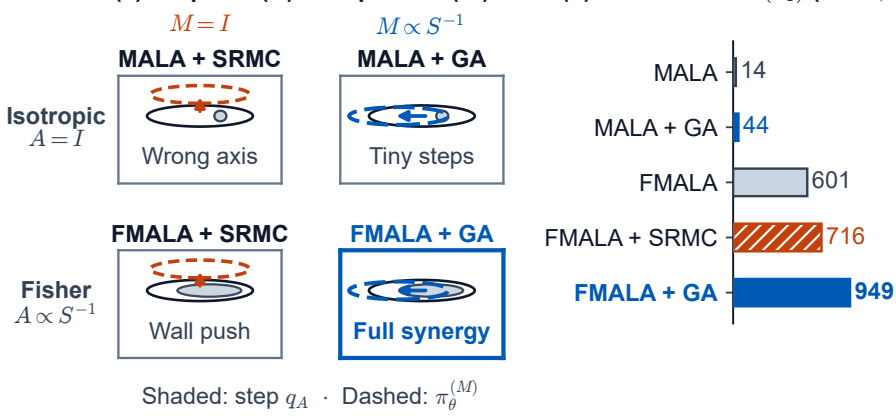

[2] Titsias, Michalis. Optimal preconditioning and fisher adaptive langevin sampling. NeurIPS, 2023.

Shared-Fisher Online Wrapper

$$M_{\text{opt}} = \frac{\operatorname{tr}(S)}{d} S^{-1}, \qquad S = \operatorname{Cov}_\pi\big(s(X)\big) \quad \text{(online: } \widehat{S}_{n,\delta} = \widehat{S}_n + \delta I_d\text{)}$$

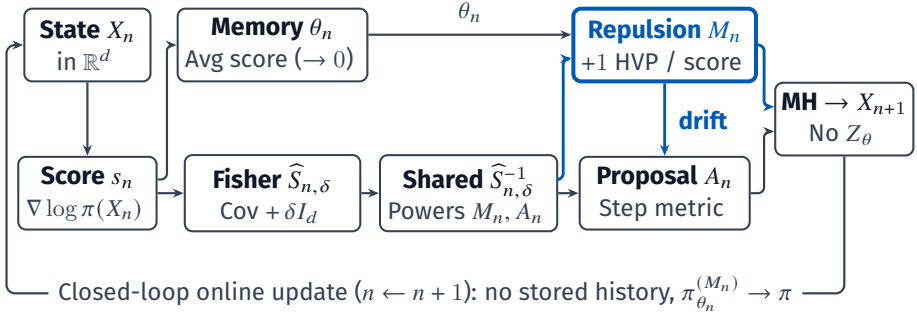

0 extra factorizations over FMALA; $+1$ HVP* per score eval (2/MALA step); exact MH.

*HVP: Hessian–vector product (via automatic differentiation, without forming the full Hessian).

Theoretical Optimality and Online Convergence

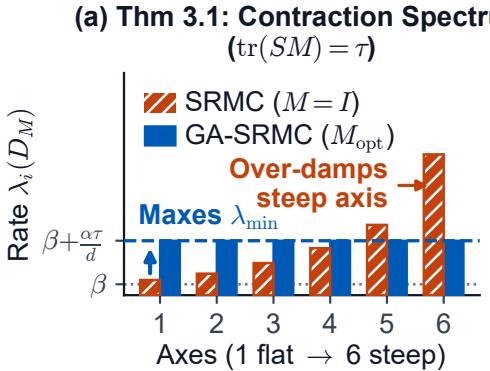

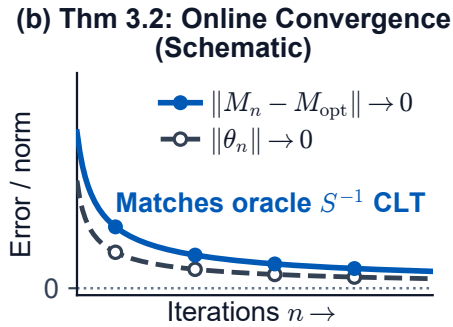

Equalizes contraction (**Thm 3.1**) and matches oracle S^{-1} CLT online (**Thm 3.2**).

Empirical Mixing and Steep-Axis Stability

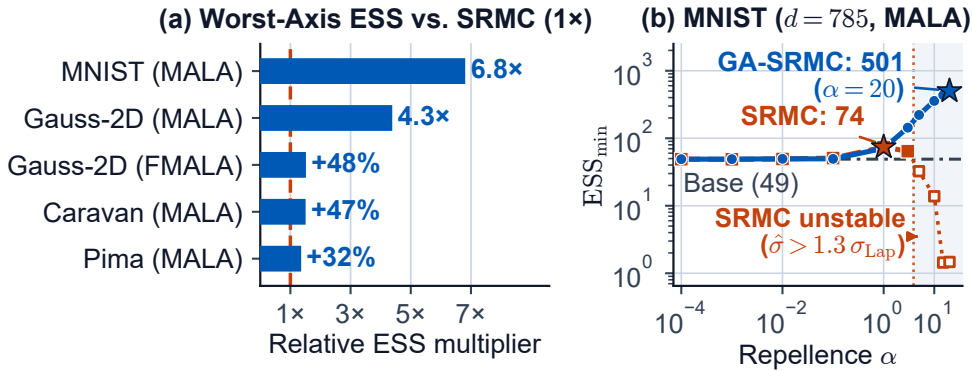

8 targets (best stable α): wins/ties **14/16** pairs; stable to $\alpha = 20$ on MNIST.

Posterior-Mean Accuracy Across All 16 Pairs

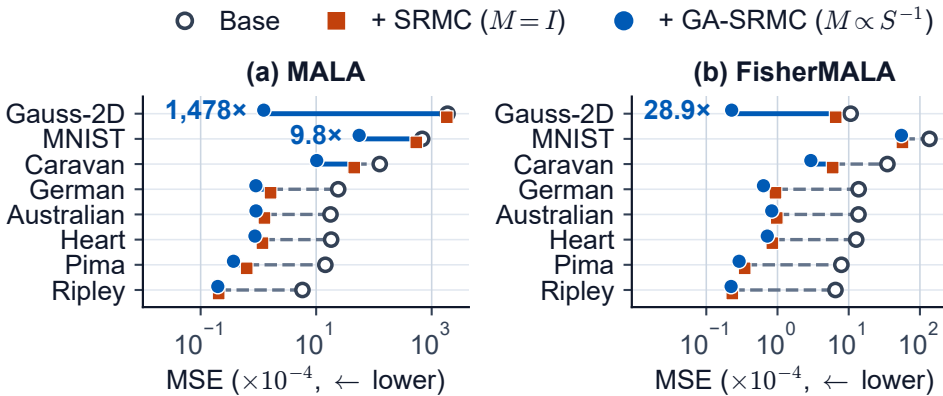

16/16 lower MSE: full $O(d^2)$ if coupled, diagonal $O(d)$ for axis-aligned targets.

Summary

1. Geometry-Aware Score Repulsion

Weighting score alignment $\theta_n^\top M\, s(x)$ by inverse Fisher $M \propto S^{-1}$ directs chains along slow flat valleys without steep-wall overshoot, recovering exact π as $\theta_n \to 0$.

2. Shared Online Fisher Wrapper

One online inverse-Fisher factor $\widehat{S}_{n,\delta}^{-1}$ powers both proposal preconditioning (A_n) and target repulsion (M_n) with zero extra factorizations over FisherMALA and oracle CLT efficiency.

3. Empirical Gains and Extensions

Boosts worst-axis ESS up to $6.8\times$ over SRMC and lowers posterior-mean MSE across all $16/16$ target–sampler pairs ($d \le 785$); future work: HMC, Riemannian HMC, and NUTS.

Contact: jiehu@oakland.edu · lchen47@ncsu.edu · dyeun@ncsu.edu