# Sequential Controlled Langevin Diffusions

**Junhua Chen**[*,†,1], **Lorenz Richter**[*,2,3], **Julius Berner**[*,†,4], **Denis Blessing**[*,5], **Gerhard Neumann**[5], **Anima Anandkumar**[6]

[1]University of Cambridge, [2]Zuse Institute Berlin, [3]dida Datenschmiede GmbH,[4]NVIDIA, [5]Karlsruhe Institute of Technology, [6]California Institute of Technology

[*] equal contribution

[†] Work partially done at the California Institute of Technology

# Problem Statement

We consider the task of sampling an unnormalized density

$$p_{\mathrm{target}} = \frac{\rho_{\mathrm{target}}}{Z} \quad \text{with} \quad Z := \int_{\mathbb{R}^d} \rho_{\mathrm{target}}(x)\mathrm{d}x,$$

where $\rho_{\mathrm{target}} \in C(\mathbb{R}^d, \mathbb{R}_{\geq 0})$ and its gradient can be evaluated pointwise, but the normalizing constant $Z$ is unknown.

Sampling from **unnormalized densities** is central in many applications (e.g. Bayesian inference, molecular dynamics).

# Prior Work & Motivation

Two popular approaches for Sampling are SMC and Diffusion Sampling:

- **Sequential Monte Carlo (SMC):** Combines resampling and MCMC.
  - Resampling helps focus computation on promising regions, Robust performance.
  - Slow sampling and requires tuning.
- **Diffusion-based samplers:** Learn a SDE to transport samples from a prior to the target.
  - Training helps samplers adapt to given densities, leading to fixed, short sampling times.
  - Training often time-consuming and unstable.

Our work **Sequential Controlled Langevin Diffusions** (SCLD) aims to **unify the two methods and combine their strengths** during training and sampling time.

**Key Idea 0 (Vargas et.al '24):** Consider the task of learning a *forward* transport $\vec{\mathbb{P}}^u$ from some tractable prior $\pi_0$ to target $\pi_1 = p_{\text{target}}$. We can do this by enforcing the time reversal property between the $\vec{\mathbb{P}}^u$ and some *backward process* $\overleftarrow{\mathbb{P}}^v$ with $\overleftarrow{\mathbb{P}}^v_{t=1} =^d \pi_1$ by minimizing (estimates of) $D(\vec{\mathbb{P}}^u, \overleftarrow{\mathbb{P}}^v)$ for a divergence $D$ between their induced path measures.

We can define the forward and backward processes $\vec{\mathbb{P}}^u, \overleftarrow{\mathbb{P}}^v$ as diffusion SDEs with drift:

$$\vec{\mathbb{P}}^u : \ \mathrm{d}X_t^u = u(X_t^u, t)\mathrm{d}t + \sigma(t)\vec{\mathrm{d}}W_t, \qquad X_0^u \sim \pi_0$$
$$\overleftarrow{\mathbb{P}}^v : \mathrm{d}Y_t^v = v(Y_t^v, t)\mathrm{d}t + \sigma(t)\overleftarrow{\mathrm{d}}W_t, \qquad Y_1^v \sim \pi_1$$

where $u$ and $v$ are parametrized functions, often constrained to give unique minima. The Radon-Nikodym derivative $\frac{\mathrm{d}\overleftarrow{\mathbb{P}}}{\mathrm{d}\vec{\mathbb{P}}}$ is tractable, and so divergences like the KL-divergence $\mathbb{E}_{X_{[0,1]}\sim\vec{\mathbb{P}}}[\log \frac{\mathrm{d}\overleftarrow{\mathbb{P}}}{\mathrm{d}\vec{\mathbb{P}}}]$ or the Log-Variance (LV) divergence $\mathrm{Var}[\log \frac{\mathrm{d}\overleftarrow{\mathbb{P}}}{\mathrm{d}\vec{\mathbb{P}}}]$ can be used.

# Method Overview: SCLD - Sampling

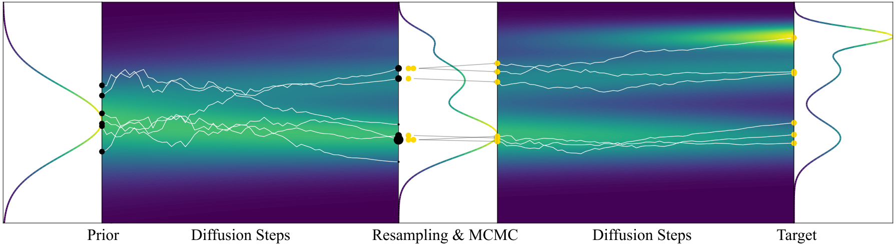


**Key Idea 1:** The trained Forward and Backward SDEs of a Diffusion Sampler can be readily used as the forward and backward kernels of an SMC sampler in conjunction with a probability annealing schedule $(\pi_t)_{0 \leq t \leq 1}$.

- SCLD uses controlled SDEs to evolve particles, and intersperses resampling and MCMC steps at arbitrary times. Can be seen as performing importance sampling on path space.

# Method Overview: SCLD - Training

**Key Idea 2:** Rather than learning a diffusion sampler separately first, the entire SCLD sampling setup can be trained end-to-end by minimizing a *sub-trajectory based divergence.*

For two collections of sub-trajectory path measures $(\vec{\mathbb{P}}_{[t_{n-1},t_n]})_{1\leq n\leq N}$ and $(\overleftarrow{\mathbb{P}}_{[t_{n-1},t_n]})_{1\leq n\leq N}$ (where $0 = t_0 \leq \ldots \leq t_N = 1$) we can define a *sub-trajectory* divergence by

$$\mathcal{L} = \sum_{n=1}^{N} D\left(\vec{\mathbb{P}}_{[t_{n-1},t_n]}, \overleftarrow{\mathbb{P}}_{[t_{n-1},t_n]}\right),$$

- We considered both $D = D_{KL}$ and $D = D_{LV}$, and found that $D_{LV}$ was both better in theory and in practice and use it going forward. Other divergences possible.
- We also adopt multiple algorithmic improvements such as *Replay Buffers* and *Learned Annealing Schedules*. There is room for further improvement in our design choices.

# Experimental Results - Quantitative

SCLD consistently delivers strong performance across our set of benchmark tasks.

Table: Comparison of different methods in terms of ELBOs, i.e., lower bounds on the log-normalization constant $\log Z$.

| ELBO (↑) | Seeds (26d) | Sonar (61d) | Credit (25d) | Brownian (32d) | LGCP (1600d) |
|---|---|---|---|---|---|
| **SMC** | $-74.63_{\pm 0.14}$ | $-111.50_{\pm 0.96}$ | $-589.82_{\pm 5.72}$ | $-2.21_{\pm 0.53}$ | $385.75_{\pm 7.65}$ |
| **SMC-ESS** | $-74.07_{\pm 0.60}$ | $-109.10_{\pm 0.17}$ | $-505.57_{\pm 0.18}$ | $0.49_{\pm 0.19}$ | $\mathbf{497.85_{\pm 0.11}}$ |
| **SMC-FC** | $-74.07_{\pm 0.02}$ | $-108.93_{\pm 0.02}$ | $-505.30_{\pm 0.02}$ | $-1.91_{\pm 0.04}$ | $-878.10_{\pm 2.20}$ |
| **CRAFT** | $-73.75_{\pm 0.02}$ | $-108.97_{\pm 0.16}$ | $-518.25_{\pm 0.52}$ | $0.90_{\pm 0.10}$ | $485.87_{\pm 0.37}$ |
| **DDS** | $-75.21_{\pm 0.21}$ | $-121.22_{\pm 5.99}$ | $-514.74_{\pm 1.22}$ | $0.56_{\pm 0.23}$ | NA |
| **PIS** | $-88.92_{\pm 2.05}$ | $-142.87_{\pm 3.29}$ | $-846.57_{\pm 2.42}$ | NA | $479.54_{\pm 0.40}$ |
| **CMCD-KL** | $-73.51_{\pm 0.01}$ | $-109.09_{\pm 0.01}$ | $-507.23_{\pm 6.40}$ | $0.86_{\pm 0.01}$ | $478.75_{\pm 0.34}$ |
| **CMCD-LV** | $-73.67_{\pm 0.01}$ | $-109.50_{\pm 0.03}$ | $-504.90_{\pm 0.02}$ | $0.54_{\pm 0.03}$ | $472.79_{\pm 0.44}$ |
| **SCLD (ours)** | $\mathbf{-73.45_{\pm 0.01}}$ | $\mathbf{-108.17_{\pm 0.25}}$ | $\mathbf{-504.46_{\pm 0.09}}$ | $\mathbf{1.00_{\pm 0.18}}$ | $486.77_{\pm 0.70}$ |

Table: Comparison of different methods in terms of Sinkhorn distances to groundtruth samples, when available

| Sinkhorn (↓) | Funnel (10d) | MW54 (5d) | Robot1 (10d) | Robot4 (10d) | GMM40 (50d) | MoS (50d) |
|---|---|---|---|---|---|---|
| **SMC** | $149.35_{\pm 4.73}$ | $20.71_{\pm 5.33}$ | $24.02_{\pm 1.06}$ | $24.08_{\pm 0.26}$ | $46370.34_{\pm 137.79}$ | $3297.28_{\pm 2184.54}$ |
| **SMC-ESS** | $\mathbf{117.48_{\pm 9.70}}$ | $1.11_{\pm 0.15}$ | $1.82_{\pm 0.50}$ | $2.11_{\pm 0.31}$ | $24240.68_{\pm 50.52}$ | $1477.04_{\pm 133.80}$ |
| **SMC-FC** | $211.43_{\pm 30.08}$ | $2.03_{\pm 0.17}$ | $0.37_{\pm 0.08}$ | $1.23_{\pm 0.02}$ | $39018.27_{\pm 159.32}$ | $3200.10_{\pm 95.35}$ |
| **CRAFT** | $133.42_{\pm 1.04}$ | $11.47_{\pm 0.90}$ | $2.92_{\pm 0.01}$ | $4.14_{\pm 0.50}$ | $28960.70_{\pm 354.89}$ | $1918.14_{\pm 108.22}$ |
| **DDS** | $142.89_{\pm 9.55}$ | $0.63_{\pm 0.24}$ | $11.44_{\pm 12.50}$ | $5.38_{\pm 2.44}$ | $5435.18_{\pm 172.20}$ | $2154.88_{\pm 3.86}$ |
| **PIS** | NA | $\mathbf{0.42_{\pm 0.01}}$ | $1.54_{\pm 0.72}$ | $2.02_{\pm 0.36}$ | $10405.75_{\pm 69.41}$ | $2113.17_{\pm 31.17}$ |
| **CMCD-KL** | $124.89_{\pm 8.95}$ | $0.57_{\pm 0.05}$ | $3.71_{\pm 1.00}$ | $2.62_{\pm 0.41}$ | $22132.28_{\pm 595.18}$ | $1848.89_{\pm 532.56}$ |
| **CMCD-LV** | $139.07_{\pm 9.35}$ | $0.51_{\pm 0.08}$ | $28.49_{\pm 0.07}$ | $27.00_{\pm 0.07}$ | $4258.57_{\pm 737.15}$ | $1945.71_{\pm 48.79}$ |
| **SCLD (ours)** | $134.23_{\pm 8.39}$ | $0.44_{\pm 0.06}$ | $\mathbf{0.31_{\pm 0.04}}$ | $\mathbf{0.40_{\pm 0.01}}$ | $\mathbf{3787.73_{\pm 249.75}}$ | $\mathbf{656.10_{\pm 88.97}}$ |

# Experimental Results - Qualitative

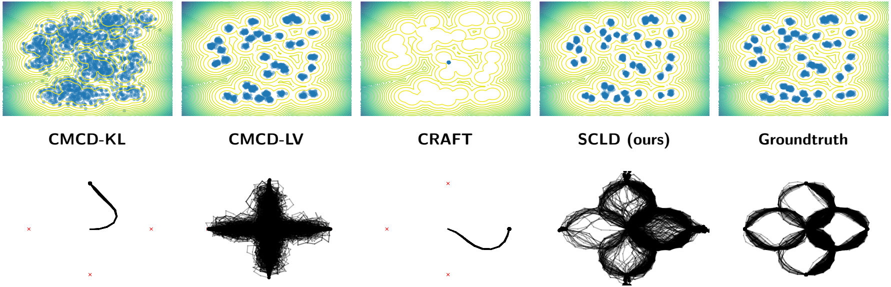


Figure: Samples from our considered methods and the groundtruth for the GMM40 (50d) (top) and Robot4 (10d) (bottom) tasks. Our SCLD method accurately finds all modes and avoids low probablity regions.

# Experimental Results - Efficiency

SCLD leverages the robustness of SMC (resampling/MCMC) and the flexibility of learned diffusions, facilitating speedup in convergence and up to a $10\times$ decrease in training time to attain performance of prior samplers.

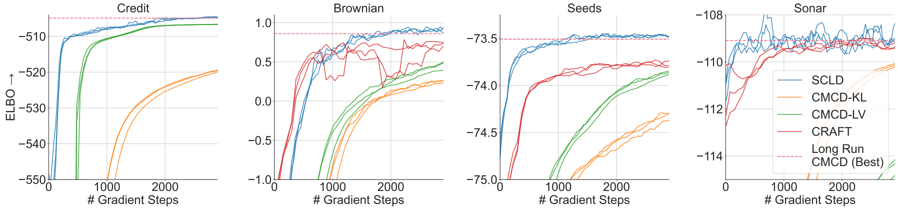


Figure: ELBOs during training for several tasks. We visualize the ELBO estimates attained by 4 methods as training progresses, running 3 seeds for each task. We mark the long run CMCD ELBOs (best out of KL and LV loss), corresponding to running for 40000 gradient steps.

# Conclusion & Future Work

**Our Contributions:**

- A framework for merging SMC and diffusion-based sampling.
- Resulting improvements in convergence speed and sample quality.

## Thank you for your attention!

**Paper:** arxiv.org/abs/2412.07081
**Code:** github.com/anonymous3141/SCLD/

junhua.chen.ios@gmail.com lorenz.richter@gmail.com mail@jberner.info
denis.blessing@kit.edu