Sample complexity of Schrödinger potential estimation

Nikita Puchkin [1] Denis Suchkov [1,2] Iurii Pustovalov [1]
Yuri Sapronov [1,3] Alexey Naumov [1] Denis Belomestny [1,4]

[1] HSE University [2] Skoltech [3] MIPT [4] Duisburg-Essen University

Problem statement

Base process:

$$X_0^0 \sim \mathcal{N}(0, I_d), \quad \mathrm{d}X_t^0 = b\left(m - X_t^0\right)\mathrm{d}t + \Sigma^{1/2}\,\mathrm{d}W_t, \quad 0 \leqslant t \leqslant T,$$

where $b > 0$, $m \in \mathbb{R}^d$, and Σ is a positive-semidefinite symmetric matrix.

Controlled process with a Schrodinger log-potential ψ:

$$\mathrm{d}X_t^\psi = b\left(m - X_t^\psi\right)\mathrm{d}t + \Sigma\nabla\log\left(\mathcal{T}_{T-t}[e^\psi](X_t^\psi)\right)\mathrm{d}t + \Sigma^{1/2}\,\mathrm{d}W_t, \quad X_0^\psi \sim \mathcal{N}(0, I_d),$$

where $\mathcal{T}$ is the Ornstein-Uhlenbeck operator:

$$\mathcal{T}_t[g](x) = \mathbb{E}_{\xi\sim\mathcal{N}(m_t(x),\Sigma_t)}\, g(\xi), \quad m_t(x) = (1 - e^{-bt})m + e^{-bt}x,\ \Sigma_t = \frac{1 - e^{-2bt}}{2b}\,\Sigma.$$

Optimality property. It is known that X_t^ψ, $0 \leqslant t \leqslant T$, transforms $\mathcal{N}(0, I_d)$ into

$$\rho_T^\psi(y) = (2\pi)^{-d/2}\int\limits_{\mathbb{R}^d} \frac{\mathsf{q}(y \,|\, x) e^{\psi(y)}}{\mathcal{T}_T[e^\psi](x)}\, e^{-\|x\|^2/2}\,\mathrm{d}x,$$

where $\mathsf{q}(y \,|\, x)$ is the density of $\mathcal{N}\left(m_T(x), \Sigma_T\right)$, **with minimal efforts** (see [1, Th. 3.2]).

Log-potential learning. Given i.i.d. samples $Y_1, \dots, Y_n \sim \rho_T^*$ and a reference class of log-potentials Ψ, we consider $\widehat\psi$ minimizing the empirical entropy:

$$\widehat\psi \in \operatorname*{argmin}_{\psi\in\Psi}\left\{-\frac{1}{n}\sum_{i=1}^n \log\rho_T^\psi(Y_i)\right\}.$$

Goal. Let $\widehat\rho_T$ be the corresponding density of $X_T^{\widehat\psi}$. Provide a bound on $\mathsf{KL}(\rho_T^*, \widehat\rho_T)$.

Main contribution

We provide a theoretical high-probability upper bound on $\mathsf{KL}(\rho_T^*, \widehat\rho_T)$ under very mild assumptions on ρ_T^* and Ψ. The established rate of convergence is potentially much faster than $\mathcal{O}(1/\sqrt{n})$ and can be of order $\mathcal{O}\left((\log n)^2/n\right)$.

Assumptions

Bounded sub-Gaussian density. The density ρ_T^* satisfies $\|\rho_T^*\|_{L_\infty} = \rho_{\max} < +\infty$. Moreover, ρ_T^* has sub-Gaussian tails:

$$\mathbb{E}_{Y\sim\rho_T^*}\, e^{u^\top Y} \leqslant e^{\mathrm{v}^2\|u\|^2/2} \quad \text{for any } u \in \mathbb{R}^d.$$

Sub-quadratic log-potentials. There exist non-negative constants Λ and M such that

$$-\Lambda\left\|\Sigma^{-1/2}(x - m)\right\|^2 - M \leqslant \psi(x) \leqslant M \quad \text{for all } x \in \mathbb{R}^d \text{ and } \psi \in \Psi.$$

Moreover, for any $\psi \in \Psi$ it holds that $\mathcal{T}_\infty\psi = \mathbb{E}\psi(X_\infty^0) = 0$.

Smooth parametrization. The class Ψ has a form $\Psi = \left\{\psi_\theta : \theta \in \Theta \subseteq [-R, R]^D\right\}$. Furthermore, there exists $L \geqslant 0$ such that

$$|\psi_\theta(x) - \psi_{\theta'}(x)| \leqslant L\left(1 + \|x\|^2\right)\|\theta - \theta'\|_\infty \quad \text{for all } \theta, \theta' \in \Theta \text{ and all } x \in \mathbb{R}^d.$$

Main theorem

Grant the aforementioned assumptions. Assume that T is sufficiently large in a sense that

$$bT \geqslant (5 + \log d) \vee \log\left(160b\,(\mathrm{v}^2 \vee 1)\left\|\Sigma^{-1}\right\|\right).$$

Then, for any $\delta \in (0, 1/2)$, with probability at least $1 - 2\delta$, it holds that

$$\mathsf{KL}(\rho_T^*, \widehat\rho_T) - \inf_{\psi\in\Psi}\mathsf{KL}(\rho_T^*, \rho_T^\psi) \lesssim \sqrt{\Upsilon(n,\delta)\inf_{\psi\in\Psi}\mathsf{KL}(\rho_T^*, \rho_T^\psi)} + \Upsilon(n,\delta),$$

where

$$\Upsilon(n,\delta) = (\Lambda d + M + d)\left(d + \log\frac{RLn}{\delta} + (M \vee \log\Lambda)\sqrt{d}e^{-bT}\right)\frac{D\log n}{n}.$$

The hidden constant behind $\lesssim$ depends on Σ, m, b, and v only.

Practical aspects

Log-density approximation:

$$\log\rho_T^\psi(y) = \psi(y) - \frac{1}{\mathcal{K}(T)}\log\mathcal{T}_\infty[e^\psi] + \mathcal{O}(M\sqrt{d}e^{-bT}),$$

where

$$\log\mathcal{K}(t) = 2e^2\sqrt{d}\arcsin(e^{-bt}) - 5e^2\sqrt{d}\log\left(1 - e^{-2bt}\right) = \mathcal{O}(\sqrt{d}e^{-bt}).$$

Approximate minimizer. Instead of $\widehat\psi$, we can consider

$$\widetilde\psi \in \operatorname*{argmin}_{\psi\in\Psi}\left\{-\frac{1}{n}\sum_{i=1}^n \psi(Y_i) + \frac{1}{\mathcal{K}(T)}\log\mathcal{T}_\infty[e^\psi]\right\}.$$

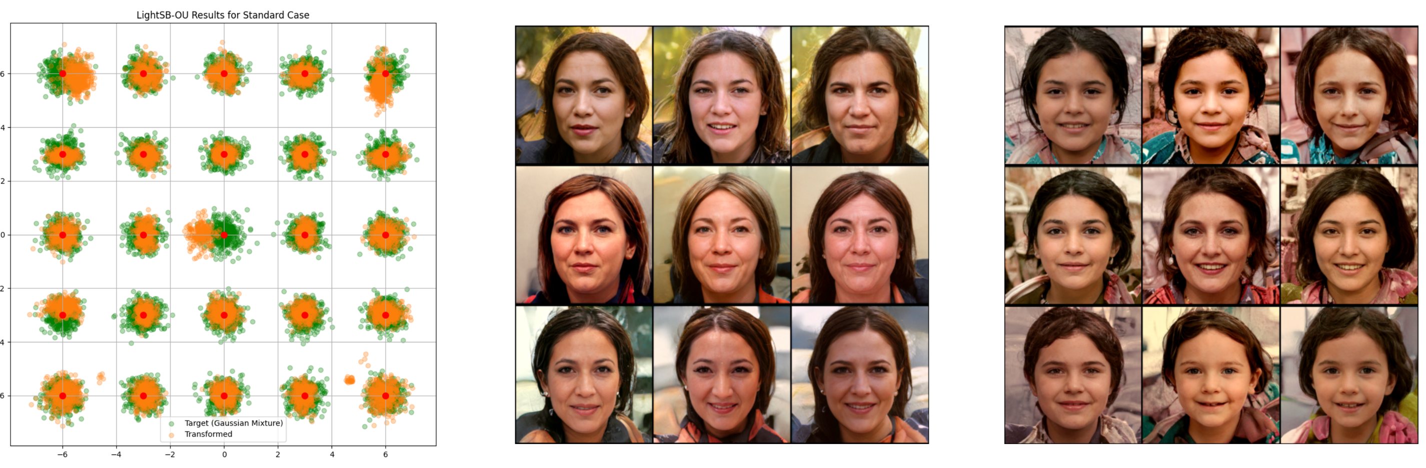

Figure: examples of sampling Gaussian mixture with 25 components (left) and human faces (center and right). In all the cases, $\Sigma = \varepsilon I_d$ and $\exp\left\{\widehat\psi(x) - b\varepsilon^{-1}\|x\|^2/(1 - e^{-bT})\right\}$ is a Gaussian mixture.

References

[1] P. Dai Pra. A stochastic control approach to reciprocal diffusion processes. *Applied Mathematics and Optimization*, 23(1):313–329, 1991.