Aalto University
Department of
Computer Science

Non-Parametric Optimization for Scalable Learning in Stochastic Decision Problems

Mohsen Amidzadeh, Lauri Viitasaari, Mario Di Francesco

ICML 2026
Forty-Third International Conference on Machine Learning

Overview/Motivation:

Our work addresses *time-varying stochastic optimization* (TV-SO), finds *non-parametric* optimality conditions via Malliavin calculus, and devises a *scalable* neural algorithm, competitive with the baselines from *both* complexity and performance perspectives. Non-parametric framework enable us to provide a learning mechanism *insensitive to the parameterization dimension*, the challenge which most of the baselines such as adjoint sensitivity models and path-wise differentiation methods (PDMs) struggle with.

Key Feature

A scalable algorithm *parallel* to PDMs and adjoint sensitivity models for stochastic optimization under distribution drift.

Application

Our work has applications in machine learning settings where parameters / variables undergo *distribution drift* as well as in *non-stationary environments*. Evaluations on *least-square recovery* and *logistic regression* problems show the effectiveness of our method against baselines, achieving a *10-fold* smaller optimality distance than the stochastic gradient based baseline, and having *extremely less runtime* than the conventional neural-based baseline (PDM).

Problem Statement

Solve the stochastic counterpart of the time-varying optimization (TV-O) problem: (TV-O) : $\min_{(\boldsymbol{x}_t)_{t\in[0,T]}} f(\boldsymbol{x}_t, t)\,,\ \forall t \in [0,T].$

Specifically, solving time-varying stochastic optimization (TV-SO) problem: $\min_{X=(\boldsymbol{x}_t)_{t\in[0,T]}} \mathbb{E}_{(\Xi,\mathbf{X})\sim\mathbb{P}\otimes\mathbb{Q}}\{f(\boldsymbol{x}_t, \boldsymbol{\xi}_t, t)\},\ \forall t \in [0,T],$

where (i) $\Xi = (\boldsymbol{\xi_t} \in \mathbb{R}^m)_{t\in[0,T]}$ Malliavin differentiable *known process* under law $\mathbb{P}$, (ii) $\mathbf{X} = (\boldsymbol{x}_t \in \mathbb{R}^n)_{t\in[0,T]}$ Malliavin differentiable *decision process* under law $\mathbb{Q}$, (iii) $f(\cdot,\cdot,\cdot)$: $\mathbb{R}^m \times \mathbb{R}^m \times [0,T] \to \mathbb{R}$ a square-integrable function representing a time-varying measure.

Theory

Theorem: Consider problem TV-SO under Assumptions (i)–(iii), with $\mathbf{X} = (\boldsymbol{x}_t)_{t\in[0,T]}$ and $\Xi = (\boldsymbol{\xi}_t)_{t\in[0,T]}$ being $\mathcal{F}$-adapted, Malliavin-differentiable, and square-integrable processes. Then TV-SO is optimized at a given instant t if the following necessary conditions hold:

$$\mathbb{E}\{\nabla_x^2 f(\boldsymbol{x}_t, \boldsymbol{\xi}_t, t)\, D_s\boldsymbol{x}_t + \boldsymbol{\nabla}_{x\xi} f(\boldsymbol{x}_t, \boldsymbol{\xi}_t, t)\, D_s\boldsymbol{\xi}_t\} = \mathbf{0}, \quad \forall s \le t$$

$$\mathbb{E}\{\nabla_x f(\boldsymbol{x}_t, \boldsymbol{\xi}_t, t)\} = \mathbf{0}$$

where $D_s\boldsymbol{x}_t$ (respectively, $D_s\boldsymbol{\xi}_t$) is the Malliavin derivative of $\boldsymbol{x}_t$ (respectively, $\boldsymbol{\xi}_t$) with respect to Wiener process $\boldsymbol{w}_t$.

Proposition: Consider the Malliavin derivative $D_s\boldsymbol{x}_t$ for the SDE $\mathrm{d}\boldsymbol{x}_t = \boldsymbol{a}(\boldsymbol{x}_t, t)dt + \sum_{l=1}^d \boldsymbol{b}_l(\boldsymbol{x}_t, t)dw_t^l$. Then, $D_s\boldsymbol{x}_t = \boldsymbol{\Gamma}_{s,t}\boldsymbol{B}(\boldsymbol{x}_s, s)$ for $s \le t$, where $\boldsymbol{B}(\boldsymbol{x}_s, s) = [\boldsymbol{b}_1, \dots, \boldsymbol{b}_d](\boldsymbol{x}_s, s)$ *and s-to-t* stochastic flow $\boldsymbol{\Gamma}_{s,t} \in \mathbb{R}^{n\times n}$ represents a linear SDE, adapted to the filtration $\mathcal{F}_s$, given by

$$d\boldsymbol{\Gamma}_{s,t} = \left(\sum_{l=1}^{d} \mathcal{J}_x\, \boldsymbol{b}_l(\boldsymbol{x}_t, t)dw_t^l + \mathcal{J}_x\boldsymbol{a}(\boldsymbol{x}_t, t)dt\right)\boldsymbol{\Gamma}_{s,t}, \qquad \boldsymbol{\Gamma}_{s,s} = \mathbf{I}, \quad \mathcal{J}_x[\cdot] \to \text{Jacobian operator}$$

Comparison of Optimality Conditions	**TV-O** (literature)**:** $\nabla_x^2 f \frac{dx_t}{dt} + \frac{\partial}{\partial t}\nabla_x f = \mathbf{0}$	**TV-SO** (our problem)**:** $\mathbb{E}\{\nabla_x^2 f\, D_s\boldsymbol{x}_t + \nabla_\xi \nabla_x f\, D_s\xi_t\} = \mathbf{0}$

A Stochastic Path Follower (SPF) algorithm to solve TV-SO

(i) **Representation:** represent decision process $\boldsymbol{x}_t$ using a neural SDE

$$\mathrm{d}\boldsymbol{x}_t^{\boldsymbol{\theta}} = \boldsymbol{a}\big(\boldsymbol{x}_t^{\boldsymbol{\theta}}, t; \boldsymbol{\theta}\big)dt + \sum_{l=1}^d \boldsymbol{b}_l\big(\boldsymbol{x}_t^{\boldsymbol{\theta}}, t; \boldsymbol{\theta}\big)dw_t^l$$

(ii) **Forward Pass**$^{(1)}$**:** simulate the process $\boldsymbol{x}_t^{\boldsymbol{\theta}}$ using any SDE-solver, given NN parameters $\boldsymbol{\theta}$

(iii) **Forward Pass**$^{(2)}$**:** simulate the Malliavin derivatives, $(D_s\boldsymbol{x}_t^{\boldsymbol{\theta}}, D_s\boldsymbol{\xi}_t)$, using any SDE-solver and based on the *s-to-t* stochastic flow $\boldsymbol{\Gamma}_{s,t}$

(iv) **Loss Evaluation:** include both the optimality conditions by the following energy-functional criterion as the learning loss:

$$\mathcal{L}(\boldsymbol{\theta}; \boldsymbol{x}, D\boldsymbol{x}) = \int_0^T \left(\|\mathbb{E}\,\nabla_x f\|^2 + \int_0^t \left\|\mathbb{E}\{\nabla_x^2 f\, D_s\boldsymbol{x}_t^{\boldsymbol{\theta}} + \nabla_{x\xi} f\, D_s\boldsymbol{\xi}_t\}\right\|^2 ds\right) dt$$

(v) **Update:** update the parameters of the neural SDE using a convergence-guaranteed Adam-type algorithm.

Algorithm 1: Stochastic Path Follower (SPF)

Goal: Solving stochastic optimization problem TV-SO .
Input: Neural drift and diffusion functions $\mathbf{a}(\cdot, t, \boldsymbol{\theta}) : \mathbb{R}^{n\times p} \times [0,T] \times \Theta \to \mathbb{R}^{n\times p}$, $\mathbf{B}(\cdot, t, \boldsymbol{\theta}) : \mathbb{R}^{n\times p} \times [0,T] \times \Theta \to \mathbb{R}^{n\times p\times d}$, initial process $\mathbf{x}_0 \in \mathbb{R}^{n\times p}$, number of sample paths P, number of episodes N_{eps}.
Output: Diffusion process $\mathbf{X}$ with optimal solution $\boldsymbol{\theta}$ of TV-SO.
for $E = 1$ **to** N_{eps} **do**
  // Simulate the neural SDE to obtain sample paths:
  $(\mathbf{x}_t, \mathbf{w}_t) \leftarrow \texttt{SDEsolve}(\mathbf{a}, \mathbf{B}, \mathbf{x}_0, \boldsymbol{\theta})$;
  // Compute Malliavin derivative $D_s\mathbf{x}_t$ by simulating the SDE
  // of the s-to-t stochastic flow $\boldsymbol{\Gamma}_{s,t}$
  $D_s\mathbf{x}_t \leftarrow \texttt{SDEsolve}(\mathcal{J}_\mathbf{x}\mathbf{a}, \mathcal{J}_\mathbf{x}\mathbf{B}, \boldsymbol{\Gamma}_{s,s})$;
  // Evaluate the energy-functional loss $\mathcal{L}(\theta; \mathbf{x}, D\mathbf{x})$
  $\mathcal{L}(\boldsymbol{\theta}) \leftarrow \mathcal{L}(\theta; \mathbf{x}, D\mathbf{x})$;
  // Update parameters of neural SDE:
  $\boldsymbol{\theta} \leftarrow \text{AdamOptimizer}(\boldsymbol{\theta}, \mathcal{L}(\boldsymbol{\theta}))$;
end

Experiments

Approaches	
Stochastic Path Follower SPF (ours)	Scalable neural algorithm; needs $n + n^2$ SDEs for solution (notably independent of parametrization)
Path-wise Differentiation Method (PDM) [Tze19]	Neural baseline; needs $n + n\phi$ SDEs for solution with ϕ parameterization dimension
Proximal Stochastic Gradient (PSG) [Cut23]	Neural-free gradient-based approach

Least-squares recovery with distribution drift:
$\min_{\boldsymbol{X}} \mathbb{E}\frac{1}{2}\|\boldsymbol{A}\boldsymbol{x}_t - \boldsymbol{n}_t\|^2, \quad \boldsymbol{A} \in \mathbb{R}^{n\times d}, \boldsymbol{n}_t \in \mathbb{R}^n, \boldsymbol{x}_t \in \mathbb{R}^d$
$\boldsymbol{A}$: measurement matrix, $\boldsymbol{n}_t \sim \mathcal{N}(\boldsymbol{A}\boldsymbol{x}_t^*, \boldsymbol{C})$: observed parameter, $\boldsymbol{C} = \sigma^2\boldsymbol{I}$ $\boldsymbol{x}_t^*$ follows an Ornstein–Uhlenbeck (OU) process

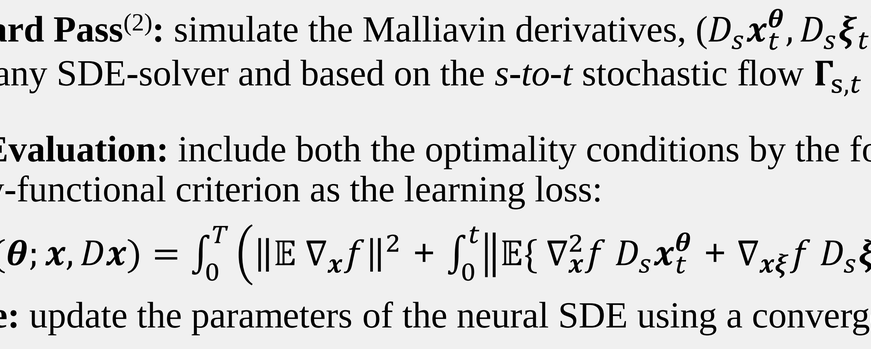

Logistic regression with distribution drift:
$\min_{\boldsymbol{X}} \mathbb{E}\frac{1}{n}\left(\sum_{i=1}^n \log\left(1 + \exp\left(\boldsymbol{a}_i^\top \boldsymbol{x}_t\right)\right) - \boldsymbol{b}_t^\top \boldsymbol{A}\boldsymbol{x}_t\right), \quad \boldsymbol{A} = \{\boldsymbol{a}_i\}_{i=1}^n \in \mathbb{R}^{n\times d}, \boldsymbol{b}_t \in \mathbb{R}^n, \boldsymbol{x}_t \in \mathbb{R}^d$
$\boldsymbol{A}$: measurement matrix, $\boldsymbol{b}_t = 1/(1 + e^{-\boldsymbol{y}_t})$: soft labels with $\boldsymbol{y}_t$ following an OU process

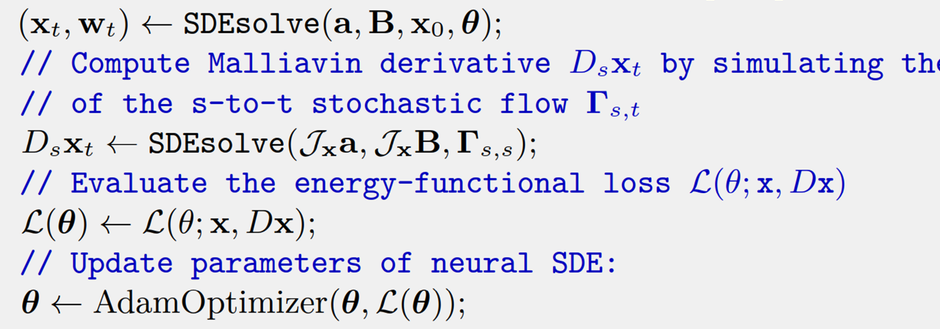

High-dimensional evaluation on least-squares recovery (n = 100, d = 50)

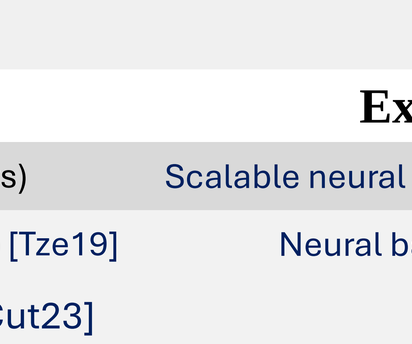

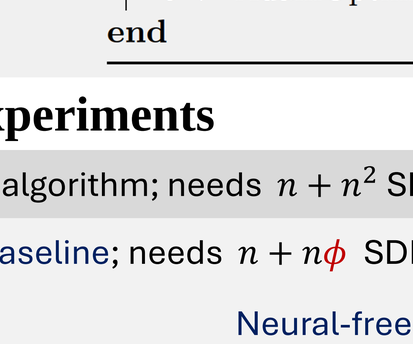

Quick Access

References: Cutler, J., Drusvyatskiy, D., and Harchaoui, Z. Stochastic optimization under distributional drift. Journal of Machine Learning Research, 24(147):1–56, 2023.
Tzen, B. and Raginsky, M. Neural stochastic differential equations: Deep latent gaussian models in the diffusion limit, arXiv:1905.09883, 2019.