



HAL
open science

Learning Controlled Stochastic Differential Equations

Luc Brogat-Motte, Riccardo Bonalli, Alessandro Rudi

► **To cite this version:**

Luc Brogat-Motte, Riccardo Bonalli, Alessandro Rudi. Learning Controlled Stochastic Differential Equations. 2024. hal-04765038

HAL Id: hal-04765038

<https://hal.science/hal-04765038v1>

Preprint submitted on 4 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning Controlled Stochastic Differential Equations

Luc Brogat-Motte*

Riccardo Bonalli*

Alessandro Rudi†

Abstract. Identification of nonlinear dynamical systems is crucial across various fields, facilitating tasks such as control, prediction, optimization, and fault detection. Many applications require methods capable of handling complex systems while providing strong learning guarantees for safe and reliable performance. However, existing approaches often focus on simplified scenarios, such as deterministic models, known diffusion, discrete systems, one-dimensional dynamics, or systems constrained by strong structural assumptions such as linearity. This work proposes a novel method for estimating both drift and diffusion coefficients of continuous, multidimensional, nonlinear controlled stochastic differential equations with non-uniform diffusion. We assume regularity of the coefficients within a Sobolev space, allowing for broad applicability to various dynamical systems in robotics, finance, climate modeling, and biology. Leveraging the Fokker-Planck equation, we split the estimation into two tasks: (a) estimating system dynamics for a finite set of controls, and (b) estimating coefficients that govern those dynamics. We provide strong theoretical guarantees, including finite-sample bounds for L^2 , L^∞ , and risk metrics, with learning rates adaptive to coefficients' regularity, similar to those in nonparametric least-squares regression literature. The practical effectiveness of our approach is demonstrated through extensive numerical experiments. Our method is available as an open-source Python library.

Keywords. Nonlinear dynamical systems identification, controlled stochastic differential equation, least-squares, statistical learning theory, kernel methods.

Mathematics Subject Classification. 68Q32, 60H10, 93B30.

1 Introduction

Modeling complex dynamical systems is pivotal across various fields, enabling tasks such as analysis, prediction, simulation, control, optimization, and fault detection [31, 45]. Deriving models from first principles—such as physical, electrical, mechanical, chemical, biological, or economic laws—requires extensive knowledge, which is often lacking in practice. In response, the literature has seen the emergence of data-driven modeling approaches since at least the 1970s, utilizing input-output data sets to identify the most suitable model within a hypothesis set of possible models [4].

Stochastic differential equations (SDEs) are a general mathematical tool for modeling dynamical systems subject to random fluctuations. Let $X(t)$ be a controlled n -dimensional stochastic process whose dynamics are governed by the controlled SDE

$$dX(t) = b(t, X(t), u(t)) dt + \sigma(t, X(t), u(t)) dW(t), \quad X(0) \sim p_0, \quad u \in \mathcal{H} \quad (1)$$

with coefficients $(b, \sigma^2) : [0, T] \times \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}^{n+n^2}$, where $W(t)$ is a standard Brownian motion, p_0 a probability density over \mathbb{R}^n , and $\mathcal{H} \subset \mathcal{F}([0, T], \mathbb{R}^d)$ a set of deterministic controls.

*Laboratoire des Signaux et Systèmes, Université Paris-Saclay, CNRS, CentraleSupélec, Bât. Bréguet, 3 Rue Joliot Curie, 91190 Gif-sur-Yvette, France. Emails: luc.brogat.motte@l2s.centralesupelec.fr, riccardo.bonalli@cnrs.fr.

†Inria, Ecole Normale Supérieure, CNRS, PSL Research University, Paris, France. Email: alessandro.rudi@inria.fr.

We consider the problem of estimating a controlled SDE from a data set of sample paths generated under various controls from \mathcal{H} . Namely, our goal is to estimate (b, σ^2) from a data set

$$(u_k, X_{u_k}(w_i^k, t_l))_{k \in [1, K], i \in [1, Q], l \in [1, M]}, \quad (2)$$

where X_{u_k} denotes the solution of Eq. (1) under the control $u_k : [0, T] \mapsto \mathbb{R}^d$. By treating controls as inputs and sample paths as outputs, we frame system identification as a supervised learning problem.

While controlled SDEs cover a wide range of scenarios, estimating them poses significant statistical and computational challenges. Consequently, the existing literature on dynamical systems (see Section 1.1) imposes various limitations to develop practical methods. These include deterministic assumptions (e.g., using ordinary differential equations, omitting diffusion), discrete assumptions (e.g., using difference equations), dimensional assumptions (e.g., one-dimensional systems), and structural assumptions on the coefficients (e.g., linearity or parametric models).

Contributions. This paper introduces a novel method addressing the controlled SDE estimation problem in its most general setting while enjoying strong theoretical guarantees. Our approach builds on recent advances in SDE estimation [11] and kernel representation [41, 44]. More precisely, we establish finite-sample learning bounds for our method, featuring statistical rates that are adaptive to standard regularity assumptions on the learning problem. Specifically, we prove L^2 learning rates between the probability density of the true SDE and that of the estimated coefficients, offering guarantees in expectation over controls, times, and positions. We then derive L^∞ learning rates, ensuring accuracy almost surely over times and controls—particularly relevant for optimal control. From these L^∞ rates, we establish CVaR learning rates, ensuring accuracy in tail estimation, crucial for risk-averse applications, such as risk-averse optimal control. For clarity of exposition, we focus on deterministic controls throughout the paper, with extensions to stochastic controls presented in a dedicated result. Furthermore, we provide implementations and experimental evaluations of the method proposed in [11] for estimating uncontrolled SDEs—as an insightful limit case of our approach where the control space reduces to a singleton—as well as of our method for controlled SDEs. Both implementations are available as open-source Python libraries on the GitHub repositories [lmotte/sde-learn](#) and [lmotte/controlled-sde-learn](#). Experimental results underscore the statistical and computational performance of the methods.

1.1 Related work

System identification. The foundation of system identification theory was established in the 1960s by [5, 29] and further strengthened in the 1990s by works such as [37, 61]. A wide variety of methods for identifying dynamical systems exists [32, 37]. These methods are generally categorized based on three key aspects: the model, the fitting method (which includes fitting criteria and optimization techniques), and—distinctly from standard supervised learning—the strategy for selecting inputs, referred to as experiment design. Given the extensive literature on system identification, we provide only a brief overview of the most prevalent methods in the field as they relate to our work and refer the reader to [31, 45] for detailed presentations of these methods.

- (Models). Dynamical systems modeling can be characterized by how states and controls interact over time (see, e.g., state-space, ARX, ARMAX models) and the form of these interactions. Common models include linear, polynomial, lookup tables, neural network (e.g., RNN, LSTM, CNN [3, 38, 68]), and fuzzy models. Nonlinear modeling, while more involved, addresses a broader spectrum of engineering challenges by enabling capturing more complex dynamics. More recently, models based on the Koopman operator have emerged [34, 48, 71].
- (Fitting methods). Models are typically fitted by optimizing a criterion that matches the outputs of the true and estimated systems across a selected set of inputs. Common methods include least-squares, Bayesian, and maximum likelihood estimations. For models that are linear in their parameters, optimization is often solved with closed-form solutions. For models that are nonlinear in their parameters, iterative optimization is commonly used.

- (Experiment design). Selecting inputs that generate the most informative data for accurate model estimation poses a significant challenge. Typically, this involves obtaining data sets that allow discrimination between any two functions within the hypothesis set. This leads for instance to concepts such as persistence of excitation, and maximization of Fisher information, as discussed in the literature [37]. Besides traditional experiment design strategies, there is a growing interest in using reinforcement learning and active learning techniques to dynamically select inputs in ways that maximize information gain for model identification [1, 66].

Regarding learning guarantees, substantial research exists on the identification of linear dynamical systems [20, 57, 60, 64]. In contrast, the identification of nonlinear dynamical systems has received much less attention [6, 15, 24, 40, 49, 58], and sample complexity in nonlinear system identification remains largely underexplored. Several studies, such as Bahmani and Romberg [6], Foster et al. [24], Oymak [49], Sattar and Oymak [58], investigate generalized linear dynamical systems of the form $x_{t+1} = \phi(Ax_t + Bu_t)$, where ϕ is a specified function. Mania et al. [40] analyze a parametric model defined as $x_{t+1} = A\phi(x_t, u_t) + w_t$. Their approach, however, is constrained by several assumptions, including a "warm start" (an informed initial model guess), the availability of a computational oracle bypassing trajectory planning intractability, and controllability (the ability to drive the system from any state to a target-aligned feature vector). While these results contribute valuable insights, they also highlight the necessity for a statistical theory designed specifically for nonparametric models, as developed in this work.

Online learning in system identification. Online learning refers to methods in which models are incrementally updated as new data becomes available, rather than being trained once on a static dataset. There is a substantial literature on online learning methods for system identification. These methods primarily focus on refining pre-existing nominal models, typically learned offline, by leveraging data gathered during system operation. The goal is to improve control performance by addressing uncertainties and adapting to environmental changes [28], rather than learning the system dynamics entirely from scratch [36, 40]. This is particularly important in many practical systems where the dynamics are only partially known, with time-varying or uncertain parameters (e.g., payloads or environmental conditions), necessitating adaptive methods that can handle these evolving conditions. A key assumption in these approaches is that the nominal model provides a reasonable approximation of the true system, ensuring safe data collection and refinement. In contrast, this work focuses on offline data collection and learning the system's dynamics from the ground up.

Uncontrolled SDE estimation. The literature on learning SDE coefficients primarily focuses on autonomous systems, which operate independently of any control inputs. There is a broad range of methods for learning SDE coefficients [47]. In terms of learning guarantees, most literature has focused on deriving guarantees for coefficient estimation by observing a single process up to time T at discrete intervals of size Δ , and then studying convergence rates as $T \rightarrow +\infty$ and $\Delta \rightarrow 0$ under ergodicity assumptions [2, 18, 19, 23, 26, 30, 33, 35, 46]. Such settings do not align with our needs in the control setting, where short time horizons are critical. Specifically, we aim to obtain arbitrarily accurate coefficients within a fixed time horizon as the number of observations increases. More recent studies consider the i.i.d. setting [11, 17, 48, 71], where data consists of i.i.d. sample paths with fixed horizons. Comte and Genon-Catalot [17] considers continuous observation of a one-dimensional process and provides risk bounds that hold in expectation. Bonalli and Rudi [11], Nüske et al. [48], Zhang and Zuazua [71] addresses a more realistic setting involving n -dimensional processes and discrete sampling observations, providing finite-sample bounds.

Controlled SDE estimation. Despite the extensive literature on system identification and SDE estimation, to our knowledge, the study by Nüske et al. [48] stands as the sole finite-data error analysis in the controlled setting. Their method is tailored for nonlinear control-affine SDEs, particularly under the structural assumption $b(t, x, u_\theta(t)) = b_0(x) + B_1(x)\theta$, where $b_0 : \mathbb{R}^n \rightarrow \mathbb{R}^n$, $B_1 : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ and the control space is parametrized by $\theta \in \mathbb{R}^m$ for $m \in \mathbb{N}^*$. Additionally, they assume that the diffusion σ does not depend on u and t . Importantly, their methodology does not provide explicit estimates of the drift and diffusion coefficients. In contrast, our approach provides explicit coefficient estimates, relying only on regularity assumptions and offering learning rates adaptive to coefficient regularity.

1.2 Paper organisation

The paper is organized as follows. In Section 2, we formulate the learning problem addressed in this work. In Section 3, we introduce our controlled SDEs estimation method. In Section 4, we provide the learning guarantees for this method. In Section 5, we present numerical experiments.

1.3 Paper notations

Let $q, r, s, d, n \in \mathbb{N}_+$. Let $\mathcal{F}(A, B)$ denote the set of all functions from the set A to the set B . Let $\mathcal{M}(A, B)$ denote the set of all measurable functions from the measurable space (A, \mathcal{A}) to (B, \mathcal{B}) , where \mathcal{A} and \mathcal{B} are the Borel σ -algebras on A and B , respectively. Let $W^q(A, B)$ denote the Sobolev Hilbert space from A to B whose weak derivatives are defined up to order q . In this work, A is either \mathbb{R}^r or the cylindrical domain $[0, T] \times \mathbb{R}^r$, and $B = \mathbb{R}^s$, ensuring well-definition. We denote the Loewner partial ordering by \preceq , such that for any two bounded linear operators A and B , $A \preceq B$ if and only if $B - A$ is positive. For any vectors u, v , $u \otimes v$ denotes the tensor product. Let $\mathcal{H} \subset \mathcal{F}([0, T], \mathbb{R}^d)$ denotes the space of possible controls, then for any sets A, C , $u \in \mathcal{H}$, and $f : \mathcal{H} \times A \rightarrow C$, we denote $f(u) \triangleq f(u(\cdot), \cdot)$. We denote $a \wedge b \triangleq \min(a, b)$, and $a \vee b \triangleq \max(a, b)$. Throughout this work, p_c denotes a probability measure on the control space \mathcal{H} . The following mixed norms are used

$$\|f\|_{L^2(p_c)L^2([0, T] \times \mathbb{R}^n)}^2 \triangleq \int_{u \in \mathcal{H}} \int_{t \in [0, T]} \int_{x \in \mathbb{R}^n} f(u, t, x)^2 p_c(u) dx dt du, \quad (3)$$

$$\|f\|_{L^\infty(p_c)L^\infty([0, T])L^2(\mathbb{R}^n)}^2 \triangleq \operatorname{ess\,sup}_{u \sim p_c} \sup_{t \in [0, T]} \int_{x \in \mathbb{R}^n} f(u, t, x)^2 dx, \quad (4)$$

where $\operatorname{ess\,sup}_{u \sim p_c}$ denotes the essential supremum taken p_c -almost surely over \mathcal{H} , ignoring sets of p_c -measure zero.

2 Problem setting

In this section, we introduce and discuss the learning problem addressed in this work.

Controlled SDE identification. Our goal is to estimate the coefficients $(b, \sigma^2) : [0, T] \times \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}^{n+n^2}$ of a n -dimensional and nonlinear controlled SDE given by

$$dX(t) = b(t, X(t), u(t)) dt + \sigma(t, X(t), u(t)) dW(t), \quad X(0) \sim p_0, \quad u \in \mathcal{H}, \quad (5)$$

where $T > 0$ is a fixed time horizon, $W(t)$ is a standard Brownian motion, p_0 is a probability density over \mathbb{R}^n , $\mathcal{H} \subset \mathcal{F}([0, T], \mathbb{R}^d)$ is a set of possible controls, thanks to a data set of controlled sample paths

$$(u_k, X_{u_k}(w_i^k, t_l))_{k \in [1, K], i \in [1, Q], l \in [1, M]}, \quad (6)$$

where X_{u_k} denotes the solution of Eq. (5) under control $u_k : [0, T] \rightarrow \mathbb{R}^d$.

Probability density associated with (b, σ^2) . For any coefficients $(b, \sigma^2) : [0, T] \times \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}^{n+n^2}$, we define the map

$$p_{b, \sigma} : \mathcal{H} \times [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}^+, \quad (7)$$

such that $p_{b, \sigma}(u, t, \cdot)$ represents the probability density of $X_u(t)$ governed by Eq. (5).

Learning problem. Given a hypothesis space $\mathcal{F} \subset \mathcal{M}([0, T] \times \mathbb{R}^n \times \mathbb{R}^d, \mathbb{R}^{n+n^2})$ for the SDE coefficients, the controlled SDE identification problem can be formulated as

$$\min_{(\hat{b}, \hat{\sigma}^2) \in \mathcal{F}} \|p_{\hat{b}, \hat{\sigma}} - p\|_{L^2(p_c)L^2([0, T] \times \mathbb{R}^n)}, \quad (8)$$

where $p \triangleq p_{b, \sigma}$ represents the probability density associated with the true controlled stochastic process.

The learning problem defined by Eq. (8) is both statistically and computationally challenging due to the complex mapping $\hat{b}, \hat{\sigma}^2 \mapsto p_{\hat{b}, \hat{\sigma}}$. In Section 3, we overcome this challenge, following the approach of recent work by Bonalli and Rudi [11], by leveraging the Fokker-Planck equation to cast this non-convex learning objective into a convex least-squares objective. This transformation allows for efficient resolution, with strong learning guarantees, harnessing the well-established literature on least-squares.

Experiment design. The choice of (\mathcal{H}, p_c) will determine the controls over which the estimated SDE is accurate (see Theorem 4.1). Therefore, it should align with the intended usage of the estimated SDE, and depend on the specific application at hand. We leave this choice to users, based on their specific experimental context. Additionally, the choice of the initial distribution p_0 can also be considered part of the experiment design. Similar to the choice of (\mathcal{H}, p_c) , it will determine the initial positions over which the estimated controlled SDE is accurate (see Theorem 4.1). We refer the curious reader to the related work section for a brief review of existing experiment design methods.

Remark 1 (Non-identifiability of SDE coefficients). *Estimating the true SDE coefficients (b, σ^2) from a set of sample paths, with L^2 norm guarantees such as $\|(\hat{b}, \hat{\sigma}^2) - (b, \sigma^2)\|_{L^2([0, T] \times \mathbb{R}^n \times \mathbb{R}^d)} \leq \epsilon(K, M, Q, \delta)$, with probability at least $1 - \delta$, given sample sizes $K, M, Q \in \mathbb{N}^*$, is generally infeasible without additional assumptions. This stems from the fact that identifying the true coefficients (b, σ^2) solely from sample paths constitutes an ill-posed problem, as distinct SDE coefficients can result in identical probability distributions of the sample paths due to unexplored regions of $[0, T] \times \mathbb{R}^n$. In contrast, recovering the probability density function $t, x \mapsto p_{b, \sigma}(t, x)$ from the sample paths is a well-posed problem. Then, from $p_{b, \sigma}$, one can identify coefficients $(\tilde{b}, \tilde{\sigma}^2)$ that, while potentially distinct from (b, σ^2) , yield the same probability density, i.e., such that $p_{\tilde{b}, \tilde{\sigma}} = p_{b, \sigma}$. See Example 1 for a prototypical case where different SDE coefficients lead to the same probability density. Investigating the conditions under which the distribution $p_{b, \sigma}$ uniquely determines (b, σ^2) is an interesting and complex question [13, 42, 51, 67], though it lies beyond the scope of this study. Finally, obtaining coefficients that accurately generate $p_{b, \sigma}$ is sufficient for many practical scenarios where the main goal is to accurately reproduce the true controlled dynamics. This justifies our learning problem proposed in Eq. (8).*

Example 1 (Prototypical non-identifiable SDE coefficients). *We provide an example where different coefficients yield the same probability density. Here, a change in the diffusion coefficient is offset by a corresponding adjustment in the drift coefficient. Consider the (uncontrolled) Ornstein-Uhlenbeck process $X(t)$ governed by the SDE*

$$dX(t) = \theta(\mu - X(t)) dt + \sigma dW(t), \quad \text{and} \quad X(0) \sim \mathcal{N}(m, s^2), \quad (9)$$

where $\theta > 0$, $\mu \in \mathbb{R}$, $\sigma > 0$, and $W(t)$ denotes standard Brownian motion. The process $X(t)$ follows a normal distribution, with its mean and variance evolving over time as

$$\mathbb{E}[X(t)] = \mu + (m - \mu)e^{-\theta t}, \quad \text{and} \quad \text{Var}(X(t)) = \frac{\sigma^2}{2\theta} (1 - e^{-2\theta t}) + s^2 e^{-2\theta t}. \quad (10)$$

If $m = \mu$ and $s^2 = \sigma^2/(2\theta)$, any pair $(\tilde{\theta}, \tilde{\sigma}) \neq (\theta, \sigma)$ where $\frac{\tilde{\sigma}^2}{2\tilde{\theta}} = \frac{\sigma^2}{2\theta}$ yield the same probability density, specifically that of a stationary normal distribution with mean m and variance s^2 .

3 Proposed method

In this section, we present our method for estimating controlled SDEs. We leverage the Fokker-Planck (FP) matching inequality to split the estimation into two more manageable problems: a) estimation of the

dynamics for a finite set of controls, and b) estimation of coefficients that could govern these dynamics. Section 3.1 introduces the FP matching inequality, while Section 3.2 details the estimator leveraging the decomposition principle.

3.1 Fokker-Planck matching inequality

We assume that the true SDE coefficients (b, σ^2) , the initial probability density p_0 , and the considered model \mathcal{F} meet a minimal smoothness assumption, characterized by a sufficiently high order of Sobolev regularity, essential for the validity of both the Strong FP equation and the FP matching inequality.

(A1) Smooth SDE. The initial density satisfies $p_0 \in W^3(\mathbb{R}^n, \mathbb{R})$, and there exists $c > 0$ such that

$$\forall (\hat{b}, \hat{\sigma}^2) \in \mathcal{F} \cup \{(b, \sigma^2)\}, u \in \mathcal{H}, \quad \|(\hat{b}, \hat{\sigma}^2)(\cdot, \cdot, u(\cdot))\|_{W^{4+\lfloor \frac{n}{2} \rfloor}} \leq c. \quad (11)$$

We assume a uniform ellipticity condition on σ , which is a standard requirement in SDE analysis to ensure that the diffusion coefficient σ remains uniformly non-degenerate.

(A2) Uniform ellipticity. There exists a constant $\kappa > 0$ such that, for all $(t, x) \in [0, T] \times \mathbb{R}^n$,

$$\sigma^2(t, x) \geq \kappa I_{\mathbb{R}^n \times n}. \quad (12)$$

Additionally, we assume the ability to sample from a probability density p_s over $[0, T] \times \mathbb{R}^n$ that covers the supports of the true coefficients and their derivatives.

(A3) Calibrated sampling. There exists a probability density p_s over $[0, T] \times \mathbb{R}^n$, and $a > 0$, such that

$$\forall (t, x) \in \text{supp}(p_s), \quad p_s(t, x) > a. \quad (13)$$

Moreover, the bounded support of p_s contains the support of (b, σ^2) and their derivatives. Namely,

$$\forall (b, \sigma^2) \in \mathcal{F}, v \in \mathbb{R}^d, \quad \text{supp}((b, \sigma^2 - \kappa I_{\mathbb{R}^n \times n})(\cdot, \cdot, v)) \subseteq D \triangleq \text{supp}(p_s), \quad (14)$$

with the same conditions applied to the derivatives of (b, σ^2) .

Under these assumptions, we derive the two following key lemmas for our method.

Lemma 3.1 (SFP equation). *Let $(\mathcal{L}^{(b, \sigma)(u)})^*$ denotes the dual of the Kolmogorov generator*

$$\left(\mathcal{L}^{(b, \sigma)(u)}\right)^* p \triangleq \frac{1}{2} \sum_{i, j=1}^n \frac{\partial^2 \sigma_{ij}^2(u) p}{\partial x_i \partial x_j} - \sum_{i=1}^n \frac{\partial b_i(u) p}{\partial x_i}. \quad (15)$$

Under Assumptions (A1) and (A2), the strong Fokker-Planck equation holds, namely,

$$\frac{\partial p}{\partial t}(u, t, x) = \left(\mathcal{L}^{(b, \sigma)(u)}\right)^* p(u, t, x), \quad \text{for almost every } (u, t, x) \in \mathcal{H} \times [0, T] \times \mathbb{R}^n. \quad (16)$$

Lemma 3.2 (FP matching inequality). *Under Assumptions (A1), (A2), and (A3), there exists $c > 0$ such that*

$$\forall (\hat{b}, \hat{\sigma}^2) \in \mathcal{F}, \quad \|p_{\hat{b}, \hat{\sigma}^2} - p\|_{L^2(p_c) L^2([0, T] \times \mathbb{R}^n)} \leq c \times FP(\hat{b}, \hat{\sigma}^2), \quad (17)$$

where $FP(\hat{b}, \hat{\sigma}^2)$ is defined as

$$FP(\hat{b}, \hat{\sigma}^2) \triangleq \mathbb{E}_{u \sim p_c, (t, x) \sim p_s} \left[\left(\frac{\partial p}{\partial t}(u, t, x) - \left(\mathcal{L}^{(\hat{b}, \hat{\sigma}^2)(u)}\right)^* p(u, t, x) \right)^2 \right]^{1/2}. \quad (\text{FP matching}) \quad (18)$$

Eq. (17) shows that our learning objective, as specified in Eq. (8), can be bounded by the convex alternative described in Eq. (18), referred to as Fokker-Planck (FP) matching. This provides theoretical support for addressing FP matching in order to solve our initial learning problem. Specifically, it allows to infer L^2 learning rates for $p_{\hat{b}, \hat{\sigma}^2}$ from FP matching learning rates for $(\hat{b}, \hat{\sigma}^2)$.

3.2 Estimator

Controlled SDE estimator. Consider a data set of controlled sample paths

$$(X_{u_k}(w_i^k, t_l), u_k)_{i \in \llbracket 1, Q \rrbracket, l \in \llbracket 1, M \rrbracket, k \in \llbracket 1, K \rrbracket}, \quad (19)$$

where the controls $(u_k)_{k=1}^K$ are i.i.d. from p_c , and the sample paths $(X_{u_k}(w_i^k, t_l))_{i \in \llbracket 1, Q \rrbracket, l \in \llbracket 1, M \rrbracket}$ are i.i.d. from $p(u_k, t_l, \cdot)$. For each control u_k , we denote \hat{p}_k as the estimator proposed by Bonalli and Rudi [11] for uncontrolled SDE density estimation, trained on the following data

$$(X_{u_k}(w_i^k, t_l))_{i \in \llbracket 1, Q \rrbracket, l \in \llbracket 1, M \rrbracket}. \quad (20)$$

Subsequently, from an i.i.d. set $(z_i)_{i=1}^N = (t_i, x_i)_{i=1}^N$ sampled from p_s , we define the estimator

$$(\hat{b}, \hat{\sigma}^2) \triangleq \arg \min_{(b, \sigma^2) \in \mathcal{F}} \frac{1}{KN} \sum_{k=1}^K \sum_{i=1}^N \left(\frac{\partial \hat{p}_k}{\partial t}(t_i, x_i) - (\mathcal{L}^{(b, \sigma)(u_k)})^* \hat{p}_k(t_i, x_i) \right)^2. \quad (21)$$

Defining the hypothesis space \mathcal{F} for (b, σ^2) with RKHSs. FP matching (Eq. (18)) constitutes a convex learning problem, specifically quadratic with respect to (b, σ^2) . This stems from the linearity of the dual Kolmogorov generator with respect to (b, σ^2) , which is a linear combination of the partial derivatives of (b, σ^2) . Consequently, RKHSs are well-suited for solving FP matching. In particular, they meet the key properties of maintaining the convexity that facilitates empirical risk minimization (Eq. (21)) while being universal approximators, leading both to good computational and statistical performance. Their ability for closed-form differentiation also significantly facilitates the computation of the Kolmogorov generator.

Uniformly elliptic diffusion σ^2 . Meaningful diffusion coefficient σ^2 should satisfy the uniform ellipticity condition of Assumption (A2). This requirement can be efficiently addressed as follows.

- (Soft shape constraint). One approach is to enforce this via pointwise constraints on a set of training points. Specifically, for $(z_i)_{i=1}^q \in ([0, T] \times \mathbb{R}^n \times \mathbb{R}^d)^q$ and $(x_j)_{j=1}^r \in (\mathbb{R}^n)^r$, the linear constraints $\forall (i, j) \in \llbracket 1, q \rrbracket \times \llbracket 1, r \rrbracket, x_j^T \sigma^2(z_i) x_j \geq \kappa$ are imposed. For uniform diffusion $\sigma^2 = \sigma_0^2 I_n$, this simplifies to $\forall i \in \llbracket 1, q \rrbracket, \sigma_0^2(z_i) \geq \kappa$.
- (Hard shape constraint). A recently introduced class of kernel models has been designed to represent PSD-valued functions Marteau-Ferey et al. [41], Muzellec et al. [44], while maintaining the aforementioned key properties of standard kernel models. These models utilize cone constraints on model parameters and have been successfully applied across various applications [9, 52, 56, 65]. With a bounded positive definite kernel k over $[0, T] \times \mathbb{R}^n \times \mathbb{R}^d$, the corresponding reproducing kernel Hilbert space (RKHS) \mathcal{G} has a feature map $\phi(t, x, v) \triangleq k((t, x, v), \cdot)$. The hypothesis space \mathcal{F} for (b, σ^2) is defined as $\mathcal{F} \triangleq \mathcal{F}_b \times \mathcal{F}_\sigma$. Here, $\mathcal{F}_b \triangleq \{b \mid \forall i \in \llbracket 1, n \rrbracket, b_i = \langle w_i, \phi(\cdot) \rangle_{\mathcal{G}}, w_i \in \mathcal{G}\}$, and $\mathcal{F}_\sigma \triangleq \{\sigma^2 + \kappa I_{\mathbb{R}^n \times \mathbb{R}^n} \mid \forall i, j \in \llbracket 1, n \rrbracket, \sigma_{ij}^2 = \langle \phi(\cdot), w_{ij} \phi(\cdot) \rangle_{\mathcal{G}}, (w_{ij})_{i,j=1}^n \in \mathcal{S}\}$, where \mathcal{S} is defined as $\{W \in \mathcal{G}^n \otimes \mathcal{G}^n \mid W^* = W, W \succeq 0\}$.

For a detailed discussion on the advantages of various PSD models, refer to Marteau-Ferey et al. [41], Muzellec et al. [44].

4 Learning guarantees

In Section 4.1, we provide L^2 -in-time-control-and-position learning rates for the proposed method. In Section 4.2, we provide refined L^2 learning rates under refined regularity assumptions on the learning problem. In Section 4.3, we provide L^∞ -in-time-and-control, L^2 -in-position learning rates. In Section 4.4, we provide CVaR learning rates. In Section 4.5, we illustrate our theoretical results by analyzing the special case of SDE coefficients in Sobolev spaces. In Section 4.6, we discuss the proposed method's adaptability to the regularity of the learning problem. Finally, in Section 4.7, we extend our approach to encompass a larger class of controls, expanding from deterministic to stochastic controls within closed-loop systems.

4.1 L^2 learning rates

To establish finite-sample bounds for our estimator, it is necessary to impose regularity assumptions on the learning problem, as underpinned by the No-Free-Lunch Theorem (see Devroye et al. [21]).

(A4) Attainability assumption. The true coefficients indeed belong to the chosen hypothesis space. Namely,

$$(b, \sigma^2) \in \mathcal{F}. \quad (22)$$

This is a standard assumption in the literature of regularized least-squares regression [14]. For concrete examples, refer to Section 4.5.

Theorem 4.1 (L^2 Learning rates). *For any $N, K \in \mathbb{N}^*$, let $(\hat{b}, \hat{\sigma}^2)$ be the proposed estimator, with hard shape-constrained PSD diffusion, trained with K controls $(u_k)_{k=1}^K$ i.i.d. from a probability measure p_c over \mathcal{H} , and N points $(t_i, x_i)_{i=1}^N$ i.i.d. from a probability measure p_s over $[0, T] \times \mathbb{R}^n$. Assuming Assumptions (A1), (A3), and (A4) hold true, then there exist constants $c_1, c_2 > 0$ that do not depend on N, K , or δ , such that for any $\delta \in (0, 1]$, defining $\varepsilon \triangleq \sup_{k \in \llbracket 1, K \rrbracket} L(\hat{p}_k, p(u_k))$ where L is a squared Sobolev norm defined in the appendix, and setting $\lambda = c_2(N^{-1} \log(N\delta^{-1}) + K^{-1} \log(K\delta^{-1}))$, if*

$$\varepsilon \leq \frac{1}{N} \log \frac{N}{\delta} + \frac{1}{K} \log \frac{K}{\delta}, \quad (23)$$

then with probability $1 - \delta$,

$$\|p_{\hat{b}, \hat{\sigma}} - p\|_{L^2(p_c)L^2([0, T] \times \mathbb{R}^n)} \leq c_1 \log \frac{2}{\delta} \left(\frac{\log \frac{N}{\delta}}{\sqrt{N}} + \frac{\log \frac{K}{\delta}}{\sqrt{K}} \right). \quad (24)$$

Sketch of the proof. The proof involves decomposing the error of the proposed estimator into three components:

1. Error stemming from the estimation of the probability densities $(\hat{p}_k)_{k=1}^K \approx (p(u_k))_{k=1}^K$,
2. Error stemming from the finite sampling approximation of times $[0, T]$ and positions \mathbb{R}^n ,
3. Error stemming from the finite sampling approximation of the space of controls \mathcal{H} .

Each error component is then bounded by representing all quantities as norms of linear operators, followed by appropriate decomposition and the application of Bernstein inequalities for sums of operators. \square

Theorem 4.1 demonstrates that, with sufficiently accurate density estimation (small enough ε), the estimator for the coefficients of the controlled SDE achieves the standard learning rate of kernel ridge regression with respect to K and N —without additional assumptions in the noiseless setting, where the output is unambiguous given the input.

Remark 2 (Dependency in K and N). *Note that despite utilizing KN data points $(u_k, t_i, x_i)_{i \in \llbracket 1, N \rrbracket, k \in \llbracket 1, K \rrbracket}$ in $\mathcal{H} \times [0, T] \times \mathbb{R}^n$, we do not achieve a convergence rate of $(KN)^{-1/2}$. Instead, the rate is $K^{-1/2} + N^{-1/2}$. This outcome is expected, as the points $(u_k, t_i, x_i)_{i,k}$ are not sampled independently. In particular, for a given finite number of observed controls $K \in \mathbb{N}^*$, the bound does not approach zero as N tends to infinity.*

Remark 3 (Dependency in the density estimation). *The proposed method employs the probability density estimation from Bonalli and Rudi [11]. However, our Theorem 4.1 remains valid regardless of the chosen density estimator. In particular, the learning accuracy of our method is independent of the specific choice of density estimator, given ϵ .*

Remark 4 (Control-dependent sampling of $[0, T] \times \mathbb{R}^n$). *For the sake of clarity, we present our method considering drawing the $(t_i, x_i)_{i=1}^N$ independently from the control u_k . Nevertheless, for each u_k , it might be advantageous to draw distinct data sets $(t_{k,i}, x_{k,i})_{i=1}^N$ i.i.d. from $p_s(\cdot | u_k)$, typically sampling with preference for the high-value regions of \hat{p}_k . In this case, our analysis and learning rates still apply, as we do not rely on this independence assumption in our proofs.*

Remark 5 (Soft shape Constraint). *When considering the soft shape constraint for the diffusion coefficient, the theoretical analysis involves examining kernel ridge regression under additional linear constraints, which is not expected to introduce significant difficulties. We anticipate that this approach could yield rates comparable to those obtained under hard shape constraints; nonetheless, rigorous investigation is necessary to confirm this. We leave this question for future investigation.*

4.2 Refined L^2 learning rates

Applying additional regularity conditions that finely measure the effective dimension of the learning problem generally enables the derivation of refined learning rates. In the context of kernel ridge regression, this is standardly achieved by measuring the regularity of the features and of the target, leading to learning rates that are adaptive to the strength of these regularities [14, 16]. This section focuses on such assumptions tailored to our specific least-squares problem: FP matching.

Kernel-model for $\frac{\partial p}{\partial t}$ induced from (b, σ^2) . Considering a RKHS modeling for the coefficients, as defined in Section 3, and based on Lemma 3.1, Lemma 3.2, and Assumption (A4), then there exists $w \in \mathcal{G}_n$ such that

$$\frac{\partial p}{\partial t}(u, t, x) = \langle w, \tilde{\phi}(u, t, x) \rangle_{\mathcal{G}_n}, \quad (25)$$

for almost every $(u, t, x) \in \mathcal{H} \times [0, T] \times \mathbb{R}^n$, where the feature map $\tilde{\phi} : \mathcal{H} \times [0, T] \times \mathbb{R}^n \rightarrow \mathcal{G}_n$ and feature space \mathcal{G}_n are defined in the appendix.

(A5) Regularity of $\frac{\partial p}{\partial t}$. Define $C \triangleq \mathbb{E}_{t,x,u}[(\tilde{\phi} \otimes \tilde{\phi})(u, t, x)]$. There exists $\alpha \in [0, 1]$ such that

$$\|wC^{-\alpha/2}\|_{\mathcal{G}_n} < \infty. \quad (26)$$

This condition is referred to as the source condition. Notably, Assumption (A5) with $\alpha = 0$ corresponds to Assumption (A4). It pertains to the regularity of $\frac{\partial p}{\partial t}$. For concrete examples, refer to Section 4.5.

(A6) Regularity of the features. Define $C(u) \triangleq \mathbb{E}_{t,x}[(\tilde{\phi} \otimes \tilde{\phi})(u, t, x)]$ and $C(x) \triangleq \mathbb{E}_u[(\tilde{\phi} \otimes \tilde{\phi})(u, t, x)]$.

(A6.1) There exists $r \in [0, 1]$ and $c > 0$ such that

$$\|C(u)^{-r/2} \tilde{\phi}(u, t, x)\|_{\mathcal{G}_n} < c \quad p_c p_s\text{-almost surely.} \quad (27)$$

(A6.2) There exists $s \in [0, 1]$ and $c > 0$ such that

$$\|C(x)^{-s/2}\tilde{\phi}(u, t, x)\|_{\mathcal{G}_n} < c \quad p_c p_s\text{-almost surely.} \quad (28)$$

Assumptions (A6), known as embedding properties [10, 22, 50], are always satisfied for $r = s = 0$ (as $\tilde{\phi}$ is bounded) and become stricter as $r, s \in [0, 1]$ increase. The embedding property relates to the effective dimension of the input distribution through the RKHS (see Remark 7 below). Furthermore, by quantifying the minimal alignment of the covariance with any features almost surely over the input space, it measures of how well almost all input points are represented within the least-squares objective, indicating typically the presence or absence of low-probability regions. For concrete examples, refer to Section 4.5.

Remark 6 (Decomposition of the embedding property). *Assumptions (A6) can be viewed as decomposing the embedding property on $[0, T] \times \mathbb{R}^n \times \mathcal{H}$ into two embedding properties on $[0, T] \times \mathbb{R}^n$ and \mathcal{H} . In particular, Assumptions (A6) imply $\|C^{-rs/2}\tilde{\phi}(u, t, x)\|_{\mathcal{G}_n} \leq c$ a.s. with $c > 0$. Indeed, using Jensen's operator inequality [27], we find $C(u) \preceq c\mathbb{E}_x[C(x)^s] \preceq cC^s$. Consequently, $\|C^{-rs/2}\tilde{\phi}(u, t, x)\|_{\mathcal{G}_n} \leq \|C^{-rs/2}C(u)^{r/2}\|_{\mathcal{G}_n}\|C(u)^{-r/2}\tilde{\phi}(u, t, x)\|_{\mathcal{G}_n} \leq c$, for $c > 0$.*

Remark 7 (Embedding property and capacity condition). *The embedding property is a finer assumption compared to the standard capacity condition. In particular, Assumptions (A6) imply the capacity condition $\text{Tr}(C^{1-rs}) = \mathbb{E}[\|C^{-rs/2}\tilde{\phi}\|_{\mathcal{G}_n}^2] \leq c$ for $c > 0$. While the capacity condition allows refining learning rates from $n^{-1/4}$ to $n^{-1/2}$ in the noisy least-squares setting, Assumptions (A6) allow refining learning rates from $n^{-1/2}$ to arbitrarily fast polynomial decay in the noiseless setting, as r and s transition from 0 to 1.*

We derive refined learning rates for the proposed estimator.

Theorem 4.2 (Refined L^2 learning rates). *Under the assumptions of Theorem 4.1, Assumption (A5), and Assumption (A6), then there exist constants $c_1, c_2 > 0$ that do not depend on N, K, δ , such that for any $\delta \in (0, 1]$, taking $\lambda = c_2 \left(\left(\frac{18c}{N} \log^2 \frac{N}{\delta} \right)^{\frac{1}{1-r}} + \left(\frac{18c}{K} \log^2 \frac{K}{\delta} \right)^{\frac{1}{1-s}} \right)$, if*

$$\varepsilon + N^{-1}\varepsilon_\infty \leq \left(\frac{\log \frac{N}{\delta}}{\sqrt{N}} \right)^{\frac{2+\alpha}{1-r}} + \left(\frac{\log \frac{K}{\delta}}{\sqrt{K}} \right)^{\frac{2+\alpha}{1-s}}, \quad (29)$$

then with probability $1 - \delta$,

$$\|p_{\hat{\mu}, \hat{\sigma}} - p\|_{L^2(p_c)L^2([0, T] \times \mathbb{R}^n)} \leq c_1 \log \frac{2}{\delta} \left(\frac{\log \frac{K \wedge N}{\delta}}{\sqrt{K \wedge N}} \right)^{\frac{\alpha+1}{1-r \wedge s}}. \quad (30)$$

where $\varepsilon_\infty \triangleq \sup_{k \in [1, K]} L_\infty(\hat{p}_k, p(u_k))$ with L_∞ a squared Sobolev norm defined in the appendix.

Theorem 4.2 shows that the rates of FP matching can improve at arbitrarily fast polynomial rates depending on the regularity of the learning problem. However, this improvement is constrained by the accuracy of \hat{p}_k . In particular, the error bound always exceeds the L^2 error of any \hat{p}_k . Furthermore, the lower of the two regularities measured by the parameters r and s , acts as a bottleneck, limiting the rate of convergence.

4.3 L^∞ learning rates

From an optimal control perspective, it is relevant to establish learning guarantees almost surely over the control space rather than in expectation over the control space. L^2 rates can be casted into L^∞ rates thanks to Assumption (A6).

Corollary 4.1 (L^∞ learning rates). *Under the same conditions as Theorem 4.2, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$, we have*

$$\|p_{\hat{b}, \hat{\sigma}} - p\|_{L^\infty(p_c) L^\infty([0, T]) L^2(\mathbb{R}^n)} \leq c \log \frac{2}{\delta} \left(\frac{\log \frac{K \wedge N}{\delta}}{\sqrt{K \wedge N}} \right)^{\frac{\alpha + r \wedge s}{1 - r \wedge s}}, \quad (31)$$

Corollary 4.1 provides accuracy guarantees almost surely over the set of possible controls \mathcal{H} , in contrast to Theorems 4.1 and 4.2, which offer guarantees only on average over \mathcal{H} . However, this stronger result comes at the cost of a reduced learning rate, $(\alpha + r \wedge s)/(1 - r \wedge s)$ with $r \wedge s \in [0, 1]$, instead of $(\alpha + 1)/(1 - r \wedge s)$, under the same assumptions.

4.4 CVaR learning rates

Conditional Value at Risk (CVaR). For any $\alpha \in [0, 1]$, the conditional value at risk of a random variable $X : \Omega \rightarrow \mathbb{R}$ is defined as

$$\text{CVaR}_\alpha(X) \triangleq \inf_{t \in \mathbb{R}} \{t + \alpha^{-1} \mathbb{E}[\max(0, X - t)]\}. \quad (32)$$

CVaR is a risk measure that can essentially be interpreted as the expected value conditional upon being within some percentage of the worst-case scenarios (i.e., high X).

Remark 8 (Risk averse optimal control). *CVaR has applications in risk-averse optimal control, where methods typically involve solving*

$$\inf_{u \in \mathcal{H}} \mathbb{E}[f(X_u(T)) + \lambda D(f(X_u(T))), \quad (33)$$

where D is a risk measure, f the loss, and $\lambda > 0$ a trade-off parameter. Compared to the approach of choosing $\lambda = 0$, it allows minimizing the dispersion of the loss around its mean. Various risk can be defined depending on the application. Typically, one can defined the risk measure as the variance. However, if one does not want to treats the excess over the mean equally as the shortfall, one will prefer the VaR or CVaR over the variance. The latter, enjoying better mathematical properties than the former, has become a standard tool in the litterature for managing risk [43, 59].

Lemma 4.3. *For any $f \in L^\infty(\mathbb{R}^n)$, $\alpha \in [0, 1]$, and two random variables $X_1, X_2 : \Omega \rightarrow \mathbb{R}^n$ with probability density functions p_1, p_2 , the following bound holds*

$$|\text{CVaR}_\alpha(f(X_1)) - \text{CVaR}_\alpha(f(X_2))| \leq \frac{\|f\|_\infty}{\alpha} \|p_1 - p_2\|_{L^1(\mathbb{R}^n)}. \quad (34)$$

Lemma 4.4. *For any $f \in L^2(\mathbb{R}^n)$, if there exists $\beta > 1$ such that $\|fx^\beta\|_{L^1(\mathbb{R}^n)} < +\infty$, then we have*

$$\|f\|_{L^1(\mathbb{R}^n)} \leq \|f\|_{L^2(\mathbb{R}^n)}^{\frac{\beta}{\beta+n/2}} (3 + \|fx^\beta\|_{L^1(\mathbb{R}^n)}). \quad (35)$$

Lemma 4.5. *Let $X_{\hat{b}, \hat{\sigma}}(u)$ be the solution to the SDE driven by the coefficients $(\hat{b}, \sigma^2) \in \mathcal{F}$ under control $u \in W^{4+\lfloor \frac{n}{2} \rfloor}([0, T], \mathbb{R}^d)$. Under Assumption (A1) and (A3), if $\mathbb{E}[\|X_0\|^\beta] < \infty$ for some $\beta > 2$, then there exists a constant $C > 0$ such that, for all $t \in [0, T]$,*

$$\mathbb{E}[\|X_{\hat{b}, \hat{\sigma}}(u)(t)\|^\beta] \leq C.$$

From Lemma 4.3, Lemma 4.4, and Lemma 4.5, we derive the following result.

Theorem 4.6. *Under Assumptions (A1), (A3), and (A4), and with $\mathbb{E}[\|X_0\|^\beta] < \infty$ for some $\beta > 2$, given $u \in W^{4+\lfloor \frac{\beta}{2} \rfloor}([0, T], \mathbb{R}^d)$, then for any $\alpha \in [0, 1]$, $f \in L^\infty(\mathbb{R}^n)$, $t \in [0, T]$, we have*

$$\left| CVaR_\alpha(f(X_{\hat{b}, \hat{\sigma}}(u, t))) - CVaR_\alpha(f(X_{b, \sigma}(u, t))) \right| \leq \frac{c}{\alpha} \|p_{\hat{b}, \hat{\sigma}}(u, t) - p(u, t)\|_{L^2(\mathbb{R}^n)}^{\frac{\beta}{\beta+n/2}}, \quad (36)$$

where $c > 0$ is a constant that does not depend on \hat{b} , $\hat{\sigma}$, or α .

Theorem 4.6, in conjunction with the L^∞ learning rates detailed in Corollary 4.1 enables the derivation of CVaR learning rates. Notably, this allows to derive accuracy guarantees for quantities expressed as

$$CVaR_\alpha[f(X_{(\hat{b}, \hat{\sigma})}(u, T))], \quad (37)$$

which are crucial in scenarios such as risk-averse optimal control, among other applications. This contrasts with Corollary 4.1, which allows to derive accuracy guarantees merely for quantities expressed as

$$\mathbb{E}[f(X_{(\hat{b}, \hat{\sigma})}(u, T))]. \quad (38)$$

4.5 Application to Sobolev spaces

In this section, we analyze the example of SDE coefficients in Sobolev spaces.

In this lemma, we examine the extent to which all previously discussed assumptions are satisfied when considering (b, σ^2) within a Sobolev space.

Lemma 4.7. *Let $q > \frac{1+n+m}{2} + 4$, and $\mathcal{H} = \{u_\theta | \theta \in \mathbb{R}^m\}$ with $\theta, t \mapsto u_\theta(t) \in W^q(\mathbb{R}^m \times [0, T])$. Assume σ^2 is uniformly elliptic, with $\forall (t, x) \in [0, T] \times \mathbb{R}^n$ $\sigma^2(t, x) \geq \kappa I_{\mathbb{R}^n \times \mathbb{R}^n}$ with $\kappa > 0$, and $(b, \sigma^2 - \kappa I_{\mathbb{R}^n \times \mathbb{R}^n}) \in W^q([0, T] \times \mathbb{R}^n \times \mathbb{R}^d, \mathbb{R}^{n+n^2})$. Assume further that $\theta, t, x \mapsto p(u_\theta, t, x) \in W^{q-2}(\mathbb{R}^m \times [0, T] \times \mathbb{R}^n, \mathbb{R})$. Additionally, let p_s, p_c be such that there exist constants $c_1, c_2, c_3, c_4 > 0$ ensuring $c_1 \leq p_s(t, x) \leq c_2$ holds p_s -almost surely and $c_3 \leq p_c(u) \leq c_4$ holds p_c -almost surely, and assume that the partial derivatives of $(b, \sigma^2 - \kappa I_{\mathbb{R}^n \times \mathbb{R}^n})$ are supported within the bounded set $D \triangleq \text{supp}(p_s)$. Then, using a kernel that induces the Sobolev RKHS $W^q(D \times \mathbb{R}^d, \mathbb{R}^{n+n^2})$, the Assumptions (A1), (A2), (A3), (A4) are met, Assumption (A5) is satisfied with $\alpha = 0$, Assumptions (A6) are satisfied with $r = 1 - \frac{m}{2(q-2)}$ and $s = 1 - \frac{n+1}{2(q-2)}$.*

Remark 9 (Regularity of p). *In Lemma 4.7, we operate under the assumption that the probability density function $u, t, x \mapsto p(u, t, x)$ associated with the stochastic differential equation is regular, specifically that it belongs to a Sobolev space of order $q - 2$ with respect to the parameter θ . This regularity assumption is typically valid when the coefficients $t, x, \theta \mapsto (b, \sigma^2)(t, x, u_\theta(t))$ of the SDE are themselves sufficiently regular in t, x, θ . Although a comprehensive proof of this regularity goes beyond the scope of this paper, the general validity of the assumption can be inferred from the regularity properties of the coefficients and controls. For further details on such regularity results, refer to the appropriate literature on stochastic calculus and Fokker-Planck equations (see, for example, Lunardi [39], specifically Section 8.3.1, and Friedman [25], Theorem 7).*

From Lemma 4.7, we deduce learning rates for our controlled SDE coefficients estimation method in the case of coefficients in Sobolev spaces.

Corollary 4.2 (Sobolev coefficients). *Under the assumptions of Lemma 4.7, the proposed controlled SDE estimation method verifies all learning rates provided in the previous sections for various risks. In particular, there exists a constant $c > 0$ that does not depend on N , K , or δ , such that, with sufficiently accurate density estimation (small enough ε_∞), with probability at least $1 - \delta$, we have*

$$\|p_{\hat{b}, \hat{\sigma}} - p\|_{L^\infty(p_c) L^\infty([0, T]) L^2(\mathbb{R}^n)} \leq c \log \frac{2}{\delta} \left(\frac{\log \frac{K \wedge N}{\delta}}{\sqrt{K \wedge N}} \right)^{\frac{2(q-2)}{m \vee (n+1)} - 1}. \quad (39)$$

Proof. Applying Lemma 4.7 and either Theorem 4.2 or Corollary 4.1 allows to conclude the proof. \square

The key insight from Corollary 4.2 is that although the proposed method’s learning rates suffer from the curse of dimensionality—with accuracy decreasing exponentially as the dimensions of the state and control spaces increase—this impact is mitigated by the smoothness of the SDE coefficients.

Remark 10 (Optimal learning rates). *Most existing literature focuses on the noisy regression setting. Nevertheless, in the noiseless setting, minimax rate in L^∞ -norm of $\mathcal{O}((\log^{1/2} K/\sqrt{K})^{\frac{2(q-2)}{m}})$ has been proven by Bauer et al. [7] when estimating $q-2$ -times continuously differentiable function on $[0, 1]^m$ with K i.i.d. data points. In comparison, our rates are $\mathcal{O}((\log K/\sqrt{K})^{\frac{2(q-2)}{m}(1-\frac{m}{2(q-2)})})$ slower than their rates, but they converge closely as q increases.*

We derive the following results as a consequence of Corollary 4.2.

Corollary 4.3 (Computational complexity). *Under the assumptions of Lemma 4.7, let $N \geq K$. To achieve a precision level $\eta > 0$ with respect to the L^∞ -in-time-and-control norm and L^2 -in-position norm, that is,*

$$\|p_{\hat{b}, \hat{\sigma}} - p\|_{L^\infty(p_c)L^\infty([0, T])L^2(\mathbb{R}^n)} \leq \eta, \quad (40)$$

the proposed controlled SDE estimation method requires a number of controls

$$K = \mathcal{O}\left(\eta^{-\frac{2(m \vee (n+1))}{2(q-2) - m \vee (n+1)}}\right), \quad (41)$$

with a corresponding computational cost for the Fokker-Planck matching step of

$$\mathcal{O}\left(\eta^{-\frac{12(m \vee (n+1))}{2(q-2) - m \vee (n+1)}}\right). \quad (42)$$

This result implies, for instance, that under sufficiently regular coefficients with $q \geq 2 + m \vee (n + 1)$, the proposed controlled SDE estimation method achieves accuracy η in the L^∞ -in-time-and-control and L^2 -in-position norms using $K = \mathcal{O}(\eta^{-2})$ controls and incurring a computational cost of $\mathcal{O}(\eta^{-12})$.

Remark 11 (Complexity analysis in M and Q). *In our analysis, we assume that density estimation \hat{p}_k for each $p(u_k)$ is sufficiently accurate, and therefore omits a detailed examination of the computational complexity in terms of parameters M and Q . While Bonalli and Rudi [11] address this complexity in the case of ε , an extension is required for ε_∞ . Here, we focus on the number of controls required to reach a desired accuracy $\eta > 0$, and the computational cost of the Fokker-Planck step, which dominates the computational cost in our numerical experiments of Section 5.*

4.6 Adaptive learning rates

In Section 4.5, we examined Sobolev coefficients. While analyzing stronger assumptions could potentially lead to faster learning rates, this is not the primary focus of our paper. For illustrative purposes, we provide two examples of stronger assumptions in Examples 2 and 3. However, a detailed analysis of these assumptions is beyond the scope of our study. The key result of this paper is the adaptability of the proposed method: it automatically benefits from the most favorable regularity assumptions, whether they pertain to Sobolev regularity or other types. Specifically, our method achieves the best learning rates based on the most advantageous parameters r and s in Assumptions (A6), without requiring prior knowledge of these optimal parameters.

Example 2 (Parametric probability density). *Consider the scenario where $p(u, t, \cdot)$ lies, for all $t, u \in [0, T] \times \mathcal{H}$, within a parametric family of dimension m . In this case, Fokker-Planck matching formulates as a least-squares problem over a space of dimension $m + 1 + n + d$, independent of the dimension of the control space \mathcal{H} . Specifically, by pushing forward $p_{c,p_s} : \mathcal{H} \times [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}^+$ through $t, u \mapsto u(t) \in \mathbb{R}^d$*

and $t, u \mapsto p(u, t, \cdot)$, the FP matching least-squares problem over $\mathcal{H} \times [0, T] \times \mathbb{R}^n$ can be reformulated as a least-squares problem over $\mathcal{P} \times \mathbb{R}^d \times [0, T] \times \mathbb{R}^n$, where \mathcal{P} is the set containing all distributions over \mathbb{R}^n reached given the control space and initial distribution. Specifically,

$$\mathcal{P} \triangleq \{x \in \mathbb{R}^n \mapsto p(u, t, x) \mid u \in \mathcal{H}, t \in [0, T]\}. \quad (43)$$

As an example, consider a controlled linear SDE. Given $A, C \in L^\infty([0, T], \mathbb{R}^{n^2})$ and $u, e \in L^2([0, T], \mathbb{R}^n)$, the following controlled SDE

$$dX(t) = (A(t)X(t) + u(t))dt + (C(t) + e(t))dW(t), \quad (44)$$

$$X(0) \sim \mathcal{N}(\mu_0, \Sigma_0), \quad (45)$$

where the explicit solution can be derived by means of the variation of constants formula, resulting in $X(t) \sim \mathcal{N}(\mu(t), \Sigma(t))$, with μ and Σ governed by linear ODEs. For more details, refer to Chapter 6.3 in [70]. In this setting, $p(u, t, \cdot)$ for any control $u \in \mathcal{H} \triangleq L^2([0, T], \mathbb{R}^n)$ belongs to a parametric family of dimension $m = n + n(n-1)/2$. Consequently, it is straightforward to show that Assumptions (A6) are satisfied with $r = 1 - \frac{d+m}{2(q-2)}$ and $s = 1 - \frac{n+1}{2(q-2)}$.

Example 3 (Identifiability). Consider a scenario where identifiability is assumed, characterized by constants $c_1, c_2 > 0$ and $\nu_1, \nu_2 > 0$ such that for any $(\hat{b}, \hat{\sigma}^2) \in \mathcal{F}$ we have, p_{cP_s} -almost surely,

$$\mathbb{E}_{t,x}[(\hat{b}, \hat{\sigma}^2)(t, x, u(t)) - (b, \sigma^2)(t, x, u(t))]^2 \leq c_1 \mathbb{E}_{t,x}[\|\mathcal{L}^{\hat{b}, \hat{\sigma}^2} p - \mathcal{L}^{b, \sigma} p\|^2]^{\nu_1}, \quad (46)$$

$$\mathbb{E}_u[(\hat{b}, \hat{\sigma}^2)(t, x, u(t)) - (b, \sigma^2)(t, x, u(t))]^2 \leq c_2 \mathbb{E}_u[\|\mathcal{L}^{\hat{b}, \hat{\sigma}^2} p - \mathcal{L}^{b, \sigma} p\|^2]^{\nu_2}. \quad (47)$$

Equations (46) and (47) ensure that Assumptions (A6) are satisfied with $r = 1 - \frac{\nu_2 d}{2(q-2)}$ and $s = 1 - \frac{\nu_1(n+1)}{2(q-2)}$, as they imply $\mathbb{E}_{t,x}[(\phi \otimes \phi)(t, x, u(t))] \leq c_1 \mathbb{E}_{t,x}[(\tilde{\phi} \otimes \tilde{\phi})(u, t, x)]^{\nu_1}$ and $\mathbb{E}_u[(\phi \otimes \phi)(t, x, u(t))] \leq c_2 \mathbb{E}_u[(\tilde{\phi} \otimes \tilde{\phi})(u, t, x)]^{\nu_2}$, and similar derivation than for Lemma 4.7 allows to conclude. This suggests that identifiability necessitates minimal mass conditions, such as avoiding regions of $[0, T] \times \mathbb{R}^n$ that, while having non-zero Lebesgue measure, are explored with zero probability. However, note that deriving conditions to ensure identifiability is generally intricate.

4.7 Extension to closed-loop control

For clarity in exposition, we considered position-independent (a.k.a. open-loop) controls in the previous sections. Nevertheless, our approach and results seamlessly extend to position-dependent controls $u(t, x)$, which is essential for control systems employing monitoring feedback. This extension involves updating the Fokker-Planck equation, specifically the Kolmogorov generator (or equivalently, $\tilde{\phi}$), to account for the dependencies on the system's state variables. From a practical standpoint, modifying $\tilde{\phi}$ is straightforward, although it requires knowledge of the controls' partial derivatives. From a theoretical standpoint, the required modifications are limited to Lemma 7.10 and are also straightforward. We expect the learning rates to remain the same, though the constants will be updated and will depend on the supremum bounds of the control's partial derivatives.

Lemma 4.8. When the controls $u \in \mathcal{H}$ are position-dependent, specifically $\mathcal{H} \subset \mathcal{F}([0, T] \times \mathbb{R}^n, \mathbb{R}^d)$, Assumption (A1) and (A4) yields

$$\left(\mathcal{L}^{(b, \sigma)(u)}\right)^* p(u, t, x) = \langle w, \tilde{\phi}(u, t, x) \rangle_{\mathcal{G}_n} \quad \text{for almost every } (u, t, x) \in \mathcal{H} \times [0, T] \times \mathbb{R}^n, \quad (48)$$

where

$$\tilde{\phi}(u, t, x) \triangleq \left((-\tilde{\phi}_i)_{i \in \llbracket 1, n \rrbracket} \mid (1/2 \tilde{\phi}_{ij})_{i, j \in \llbracket 1, n \rrbracket} \right) (u, t, x), \quad (49)$$

with, denoting $\phi^2 \triangleq \phi \otimes \phi$,

$$\tilde{\phi}_i(u, t, x) \triangleq \frac{\partial(\phi(u)p(u))}{\partial x_i}(t, x) \quad (50)$$

$$= \left(\phi(u) \frac{\partial p(u)}{\partial x_i} + p(u)(\phi(u))_i \right)(t, x), \quad (51)$$

$$\tilde{\phi}_{ij}(u, t, x) \triangleq \frac{\partial^2(\phi^2(u)p(u))}{\partial x_i \partial x_j}(t, x) \quad (52)$$

$$= \left(\phi^2(u) \frac{\partial^2 p(u)}{\partial x_i \partial x_j} + (\phi^2(u))_i \frac{\partial p(u)}{\partial x_j} + (\phi^2(u))_j \frac{\partial p(u)}{\partial x_i} + p(u)(\phi^2(u))_{ij} \right)(t, x), \quad (53)$$

where

$$(\phi(u))_i \triangleq \frac{\partial \phi(u)}{\partial x_i}(t, x), \quad (54)$$

$$(\phi^2(u))_{ij} \triangleq \frac{\partial^2 \phi^2(u)}{\partial x_i \partial x_j}(t, x), \quad (55)$$

and

$$\frac{\partial \phi(u)}{\partial x_i}(t, x) = \phi_{x_i}(t, x, u(t, x)) + \sum_{k=1}^d \frac{\partial u_k}{\partial x_i}(t, x) \phi_{u_k}(t, x, u(t, x)), \quad (56)$$

$$\begin{aligned} \frac{\partial^2 \phi^2(u)}{\partial x_i \partial x_j}(t, x) &= \phi_{x_i x_j}^2(t, x, u(t, x)) + \sum_{k=1}^d \frac{\partial u_k}{\partial x_i}(t, x) \phi_{x_i u_k}^2(t, x, u(t, x)) \\ &+ \sum_{k=1}^d \left(\frac{\partial^2 u_k}{\partial x_i \partial x_j}(t, x) \phi_{u_k}^2(t, x, u(t, x)) \right) + \sum_{k=1}^d \frac{\partial u_k}{\partial x_i}(t, x) \sum_{l=1}^d \frac{\partial u_l}{\partial x_j}(t, x) \phi_{u_l}^2(t, x, u(t, x)). \end{aligned} \quad (57)$$

Proof. The proof follows directly from applying the chain rule. Compared to position-independent controls, additional terms arise due to u 's dependency on position x , impacting the derivatives of ϕ . \square

5 Numerical experiments

In this section, we present numerical experiments on the problems of uncontrolled (Section 5.1) and controlled (Section 5.2) SDE estimation, illustrating the main behaviors of the evaluated methods.

5.1 Uncontrolled SDE estimation

Purpose of uncontrolled SDE experiments. We initiate our numerical study by implementing and evaluating the method proposed in Bonalli and Rudi [11] for uncontrolled SDE estimation. Note that the method proposed in Bonalli and Rudi [11] corresponds to our approach in the limit case where the control space is reduced to a singleton $\mathcal{H} = \{u\}$. Therefore, this study provides important insights for the more intricate scenario of controlled SDEs considered in Section 5.2.

Python open-source library. Our implementation is available as an open-source Python library on the GitHub repository [lmotte/sde-learn](https://github.com/lmotte/sde-learn). The library includes comprehensive documentation and example scripts demonstrating various use cases with light computational demands. In particular, our implementation supports the Nyström approximation [55], which aids in reducing the computational complexity of Fokker-Planck matching. Further details, including thorough derivations of all necessary formulas for implementing the proposed estimator and the computational complexity of each step, are provided in Section 8.

Experimental setup. We consider uniform diffusion and use the soft-shape constraint presented in Section 3 for the diffusion coefficient. All kernels are Gaussian kernels $k(x, y) = \exp(-\gamma\|x - y\|^2)$ with different parameters $\gamma > 0$. We select all hyperparameters using grid search and validation sets. The selection metrics are the log-likelihood for the probability density estimation step, i.e., $\hat{p} \mapsto \sum_{l=1}^{M_{val}} \sum_{i=1}^{Q_{val}} \log(\hat{p}(t_l, x_i))$, and the mean squared error for the Fokker-Planck matching step, i.e., $(\hat{b}, \hat{\sigma}^2) \mapsto \sum_{i=1}^N [\frac{\partial \hat{p}}{\partial t}(t_i, x_i) - (\mathcal{L}^{\hat{b}, \hat{\sigma}^2})^* p(t_i, x_i)]^2$. Further details can be found in the code repository.

Computational considerations. While the computational complexity of each step is provided in Section 8, we offer here a practical idea of the computational requirements by reporting the observed execution times from our experiments on the considered problem and data sets. These experiments were performed on a machine equipped with an Apple M3 Pro processor and 18 GB of RAM. For the 1D case, the probability density estimation step takes 0.015 seconds for training with 1000 sample paths and 100 time steps, and 0.11 seconds for prediction over a Fokker-Planck training set of size 2500. The Fokker-Planck matching step requires 6.0 seconds for training with 2500 data points and 3.7 seconds to generate 100 sample paths with 100 time steps. For the two 2D problems, probability density estimation training takes around 0.012 seconds for 3000 sample paths and 100 time steps, with prediction taking approximately 60 seconds over a Fokker-Planck training set of size 3000. The Fokker-Planck matching step requires about 10 seconds for training and 6 seconds to generate 100 sample paths with 100 time steps. These results indicate that, although the computational times are non-trivial, they remain manageable—ranging from a few seconds to just over a minute—demonstrating the practicality of the method for the problem sizes considered.

5.1.1 Linear scalar SDE

To carry out our first experiments, we consider Ornstein–Uhlenbeck processes.

Ornstein–Uhlenbeck process. The Ornstein-Uhlenbeck (OU) process is a simple example of stochastic process, which tends to revert to a mean over time under the influence of a mean-reverting term and a stochastic noise term. It is a useful tool, for example, for modeling phenomena such as price volatility. It is defined by the following SDE

$$dX(t) = \theta(\mu - X(t))dt + \sigma dW(t), \quad (58)$$

where μ is the mean to which the process reverts, θ is the mean-reverting coefficient giving the strength of reversion to μ , and σ is the amplitude of the stochastic noise. The probability density function of the OU process can be explicitly derived by solving the Fokker-Planck equation. Notably, if $X(0)$ is initially Gaussian-distributed with mean $\mu_0 \in \mathbb{R}$ and covariance $\sigma_0^2 > 0$, then $X(t)$ will follow a Gaussian distribution with the mean and covariance given by

$$\mu(t) \triangleq e^{-\theta t}(\mu_0 - \mu) + \mu, \quad \Sigma(t) \triangleq \left(\sigma_0^2 - \frac{\sigma^2}{2\theta}\right)e^{-2\theta t} + \frac{\sigma^2}{2\theta}. \quad (59)$$

Data sets. We consider an Ornstein-Uhlenbeck process with constant variance and a mean that increases from 0.5 to 2.5 from $t = 0$ to $t = 10$. Specifically, we set $\mu = 2.5, \theta = 0.5, \sigma^2 = \theta/4, \mu_0 = 0.5, \sigma_0^2 = \sigma^2/(2\theta)$, and $T = 10$. In Figure 3, we plot 100 sample paths generated from this OU process. For the probability density estimation steps, we draw training and validation sets with $Q/Q_{val} = 1000/100$ sample paths and $M/M_{val} = 100/100$ time steps, respectively. For the Fokker-Planck matching step, we draw a training set $(t_i, x_i)_{i=1}^N = \{t_i\}_{i=1}^{50} \times \{x_i\}_{i=1}^{50}$ by drawing times and positions uniformly within well-chosen intervals, and we draw a validation set with same size and distribution.

Step 1 (probability density estimation). In Figure 1, we plot the true and estimated probability densities, $p(t, x)$ and $\hat{p}(t, x)$, respectively. These densities are plotted on a uniformly spaced temporal grid, offset by 1/2 unit from the training time discretization, and a spatial grid with positions randomly

drawn from a uniform distribution within an interval. The estimated density reasonably approximates the true dynamics, as visually evident.

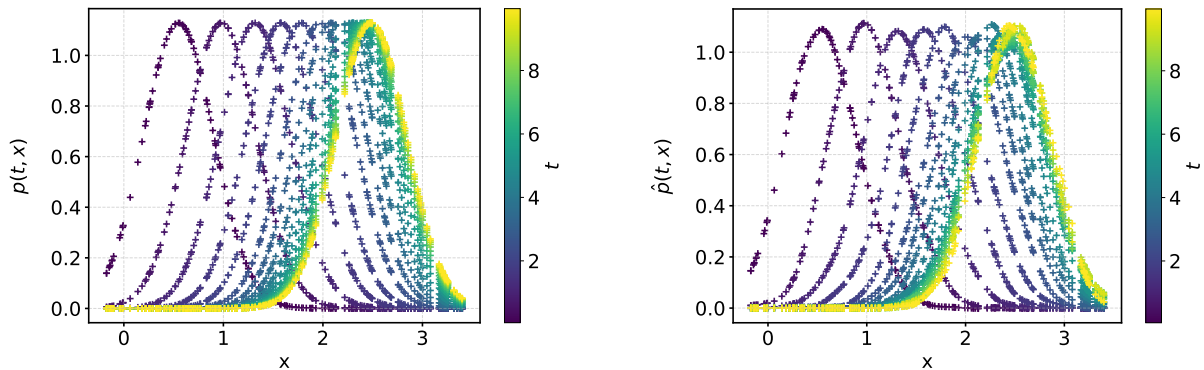


Figure 1: OU process. True and estimated probability densities, $p(t, x)$ (left) and $\hat{p}(t, x)$ (right), w.r.t x for several $t \in [0, 10]$.

Step 2 (coefficients estimation via FP matching). In Figure 2, we plot the estimated coefficients obtained via FP matching. Notice that these estimated coefficients differ significantly from the true coefficients. As discussed in Remark 1, measuring the accuracy of the estimated coefficients using the L^2 distance from the true coefficients is excluded due to the non-identifiability of the coefficients. Our learning guarantees do not ensure the recovery of the true coefficients but rather ensure that the estimated coefficients accurately replicate the true dynamics. Specifically, this involves obtaining an induced distribution $p_{\hat{b}, \hat{\sigma}}$ that closely approximates the true distribution $p_{b, \sigma}$ in terms of L^2 distance.

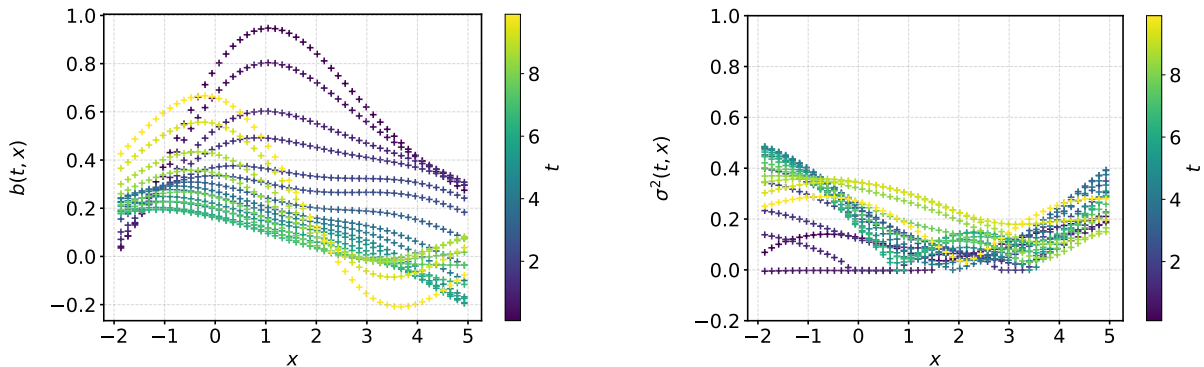


Figure 2: OU process. Estimated coefficients $\hat{b}(t, x)$ (left) and $\hat{\sigma}(t, x)^2$ (right) w.r.t. x for several $t \in [0, 10]$.

Recovering the true dynamic. We draw and plot 100 sample paths from the true and estimated coefficients in Figure 3 to assess the recovery of the true dynamics. Notice that the estimated coefficients lead to visually comparable probability distributions, with close means and variances over time. Interestingly, while the distributions are similar, individual paths appear quite different.




Figure 3: OU process. 100 samples from the SDEs associated with the true (*left*) and estimated (*right*) coefficients.

5.1.2 Nonlinear multivariate SDE

To carry out nonlinear multivariate SDE experiments, we consider two examples: Dubins process, and finite exponential sum process.

Dubins process. The Dubins process is defined by the following SDE

$$dX(t) = v \times (\cos(u(t)), \sin(u(t)))dt + \sigma dW(t) \in \mathbb{R}^2, \quad (60)$$

with $u(t) = \theta \sin(\frac{\pi t}{10})$, $v, \theta \in \mathbb{R}$, and $\sigma > 0$.




Figure 4: Dubins process. 100 samples from the true and estimated SDEs.

Finite exponential sum (FES) process. We introduce the process defined by the SDE associated with the following coefficients

$$b(t, x) = \sum_{i=1}^{n_b} b_i \exp(-\gamma \|x - x_i\| |t - t_i|), \quad \sigma(t, x) = \sum_{i=1}^{n_\sigma} s_i \exp(-\gamma \|x - x_i\| |t - t_i|), \quad (61)$$

for $(b_i)_i \in (\mathbb{R}^2)^{n_b}$, $(\sigma_i)_i \in (\mathbb{R}^2)^{n_\sigma}$, $\gamma > 0$, $n_b, n_\sigma \in \mathbb{N}^*$.




Figure 5: FES process. 100 samples from the true and estimated SDEs.

Data sets. For the Dubins process, we set the parameters as follows: $T = 10$, $n = 2$, $M = 100$, $Q = 3000$, $v = 2$, $\theta = 3$, and $\sigma = 0.3$. For the probability density estimation steps, the training and validation sets consist of $Q/Q_{val} = 3000/100$ sample paths and $M/M_{val} = 100/10$ time steps, respectively. For the Fokker-Planck matching step, training and a validation set are drawn by sampling the set of training sample paths' time-position pairs, with size $N = 3000/N_{val} = 1000$. For the finite exponential sum process, the parameters are $T = 3$, $n = 2$, $M = 100$, $Q = 3000$, $n_b = n_\sigma = 3$, $\gamma = 1$, and $(t_1, x_1) = (0, (0, 0))$, $(t_2, x_2) = (1, (1, 0))$, and $(t_3, x_3) = (3, (4, 6))$. For the probability density estimation steps, the training and validation sets consist of $Q/Q_{val} = 3000/100$ sample paths and $M/M_{val} = 100/100$ time steps, respectively. For the Fokker-Planck matching step, we use a training set $\{t_i\}_{i=1}^{30} \times \{x_i\}_{i=1}^{100}$ where the sets of times and positions are drawn uniformly in well-chosen interval and two-dimensional box, respectively. We use a validation set $\{t_i\}_{i=1}^{10} \times \{x_i\}_{i=1}^{100}$ where times and positions are uniform grids in well-chosen interval and two-dimensional box. For both processes, the initial condition is $X(0) \sim \mathcal{N}(0, 1/4I_{\mathbb{R}^2})$.

Recovering the true dynamic. We perform density estimation and Fokker-Planck (FP) matching for both processes, selecting the hyperparameters based on the validation sets as previously. We draw and plot 100 sample paths from both the true and estimated coefficients in Figure 4 and in Figure 5. Notably, the estimated coefficients lead to probability distributions that are visually comparable to those of the true dynamics, with closely matching means and variances over time.

Observations. Our experiments highlight key behaviors in the estimation of uncontrolled SDEs. Specifically, we observe the following phenomena.

- **Cumulative error.** The accuracy of the density $p_{\hat{b}, \hat{\sigma}}(t, x)$, associated with the estimated coefficients $(\hat{b}, \hat{\sigma}^2)$, decreases over time due to cumulative error. Even if the coefficients at a given time t_0 are accurate, earlier inaccuracies for $0 \leq t < t_0$ can lead to an inaccurate $p_{\hat{b}, \hat{\sigma}}(t_0, x)$.
- **Error amplification.** Paths that enter regions of low probability, whether in time or space, where the accuracy of the coefficients is poor, experience error amplification. This often results in path divergence or termination. Increasing the variance of $X(0)$ can facilitate a broader exploration of times and positions, thereby reducing this instability.

An illustrative analogy is a marble run: small deviations early on can lead to cumulative errors and amplified divergence, much like the behavior observed in the estimation of coefficients in SDEs.

5.2 Controlled SDE estimation

Python open-source library. Our implementation is available as an open-source Python library on the GitHub repository [lmotte/controlled-sde-learn](https://github.com/lmotte/controlled-sde-learn). The library includes comprehensive documentation and example scripts demonstrating various use cases with light computational demands. More details are provided in Section 9, including detailed derivations of all necessary formulas for implementing the proposed estimator.

Experimental setup. We consider uniform diffusion and the soft-shape constraint presented in Section 3 for the diffusion coefficient. All kernels are Gaussian kernels $k(x, y) = \exp(-\gamma\|x - y\|^2)$ with different parameters $\gamma > 0$. For the 1D SDE, we select all hyperparameters using grid search and validation sets. The selection metrics are the log-likelihood for the probability density estimation step, i.e., $\hat{p} \mapsto \sum_{l=1}^{M_{val}} \sum_{i=1}^{Q_{val}} \log(\hat{p}(t_l, x_i))$, and the mean squared error for the Fokker-Planck matching step, i.e., $(\hat{b}, \hat{\sigma}) \mapsto \sum_{i=1}^N \left[\frac{\partial p}{\partial t}(t_i, x_i) - (\mathcal{L}^{\hat{b}, \hat{\sigma}})^* p(t_i, x_i) \right]^2$. For the 2D SDE, to avoid excessive computation, we fix the hyperparameters for both the probability density estimation and the Fokker-Planck matching using previously selected hyperparameters for the estimation of uncontrolled Dubins process in Section 5.1.2. Further details can be found in the code repository.

Computational considerations. While the computational complexity of each step is provided in Section 9, we offer here a practical idea of the computational requirements based on the observed execution times from our experiments on the considered problem and data sets, using a machine equipped with an Apple M3 Pro processor and 18 GB of RAM. For the 1D case, the probability density estimation step takes 5.6 seconds for training with 1000 sample paths, 100 time steps, and 10 training controls, and 1.0 seconds for prediction over a Fokker-Planck training set of size 10000. The Fokker-Planck matching step requires 190 seconds for training with 10000 data points and 150 seconds to generate 100 sample paths with 100 time steps for 10 different controls. For the 2D problem, probability density estimation training takes around 0.032 seconds for 3000 sample paths, 100 time steps, and 20 training controls, with prediction taking approximately 160 seconds over a Fokker-Planck training set of size 10000. The Fokker-Planck matching step requires about 490 seconds for training and 140 seconds to generate 100 sample paths with 100 time steps for 5 different controls. These results show that, while the computational demands are higher for controlled SDEs, they remain manageable. As expected, the computational requirements increase as the sizes of the data sets are multiplied by the number of training controls.

5.2.1 Linear scalar SDE

To carry out our first experiments, we consider controlled Ornstein–Uhlenbeck processes.

Controlled Ornstein–Uhlenbeck process. Let $X(t)$ be a Ornstein–Uhlenbeck process with controlled mean, defined by the SDE

$$dX(t) = \theta(u(t) - X(t))dt + \sigma dW(t), \quad (62)$$

where $\theta \in \mathbb{R}$, $\sigma \in \mathbb{R}_+$, and $u : [0, T] \rightarrow \mathbb{R} \in \mathcal{H}$. Moreover, we define \mathcal{H} as the space of two-step functions from $[0, T]$ to \mathbb{R} , namely

$$u(t) = \begin{cases} u_0, & \text{if } x \in [0, t_1) \\ u_1, & \text{if } x \in [t_1, T] \end{cases} \quad \text{with } u_0, u_1 \in \mathbb{R}, t_1 \in [0, T]. \quad (63)$$

Data sets. We consider a controlled OU process, and set $\theta = 0.5$, $\sigma^2 = \theta/4$, $\mu_0 = 0.5$, $\sigma_0^2 = \sigma^2/(2\theta)$, $T = 10$, and $X(0) \sim \mathcal{N}(\mu_0, \sigma_0^2)$. In Figure 7, we plot 100 sample paths generated from this OU process for three different controls. We draw a data set of $K = 10$ controls, as shown in Figure 6, by drawing u_0, u_1, t_1 independently and uniformly in $[-2, 2]$, $[-2, 2]$, and $[3, 7]$, respectively. For the probability density estimation steps, we draw training and validation sets with $Q/Q_{val} = 1000/100$ sample paths

and $M/M_{val} = 100/100$ time steps, respectively. For the Fokker-Planck matching step, we draw a training set $(t_i, x_i)_{i=1}^N = \{t_i\}_{i=1}^{20} \times \{x_i\}_{i=1}^{50}$ by drawing times and positions uniformly within well-chosen intervals, and we draw a validation set with same size and distribution.




Figure 6: Training set of 10 i.i.d. piecewise-constant controls.

Recovering the true controlled dynamics. For $K_{te} = 3$ randomly drawn controls, we generate and plot 100 sample paths using both the true and estimated coefficients, as shown in Figure 7, to assess the recovery of the true controlled dynamics. The estimated coefficients produce probability distributions that are visually comparable to the true ones, with closely matching means and variances over time. It is important to note that the three controls used for evaluation are not part of the training data.

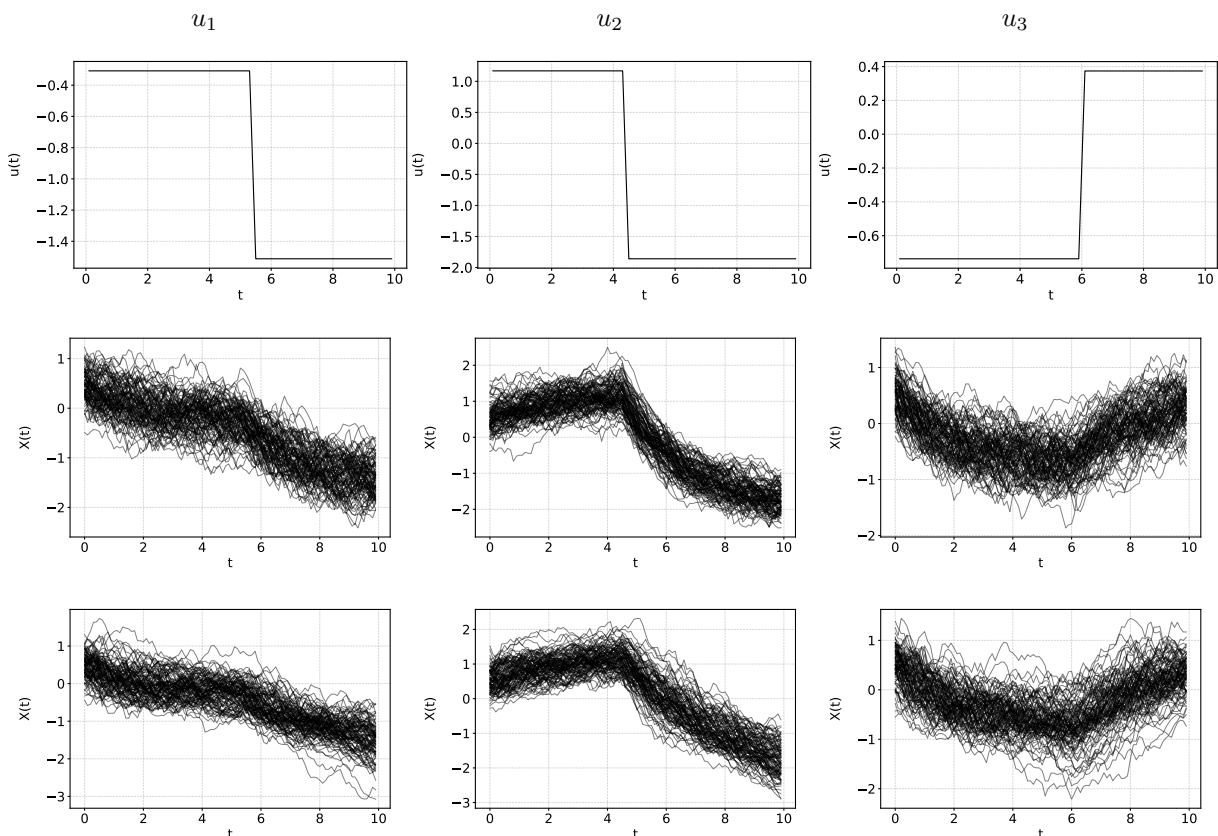


Figure 7: Controlled Ornstein-Uhlenbeck process. 100 samples from the true (*center*) and estimated (*bottom*) SDE coefficients for three random controls u_1, u_2, u_3 (*top*).

5.2.2 Nonlinear multivariate SDE

To conduct nonlinear multivariate SDE experiments, we consider controlled Dubins processes.

Controlled Dubins process. Let $X(t)$ be a Dubins process with a controlled angle, defined by the SDE

$$dX(t) = v \times (\cos(u(t)), \sin(u(t)))dt + \sigma dW(t) \in \mathbb{R}^2, \quad (64)$$

where $v \in \mathbb{R}$, $\sigma > 0$, and $u : [0, T] \rightarrow \mathbb{R} \in \mathcal{H}$. We define sinusoidal controls by specifying \mathcal{H} as

$$\mathcal{H} = \{u : t \mapsto a \sin(t\pi/10) \mid a \in \mathbb{R}\}. \quad (65)$$




Figure 8: Training set of 20 i.i.d. sinusoidal controls.

Data sets. We consider a controlled Dubins process and set $v = 2$, $\sigma = 0.3$, $\mu_0 = 0$, $\sigma_0 = 0.5$, $T = 10$, and $X(0) \sim \mathcal{N}(\mu_0, \sigma_0^2 I_{\mathbb{R}^2})$. In Figure 9, we plot 100 sample paths generated from this process for five different controls. We build a data set of $K = 20$ controls, as shown in Figure 8, by drawing a independently and uniformly from $[-1.2, 1.2]$. For the probability density estimation steps, we generate a training set with $Q = 3000$ sample paths and $M = 100$ time steps. For the Fokker-Planck matching step, we draw a training set $(u_k, t_i, x_i)_{k \in [1, 20], i \in [1, 500]}$ of size 10^4 . To avoid sampling regions where $p(u_k, t, x)$ is negligible, for each u_k , we draw a set of 500 pairs $(t_i, x_i)_{i=1}^{500}$ from a set of 5 sample paths with 100 time steps, generated using the same parameters as the controlled Dubins process but with an initial variance of $\sigma_0 = 2.5$.

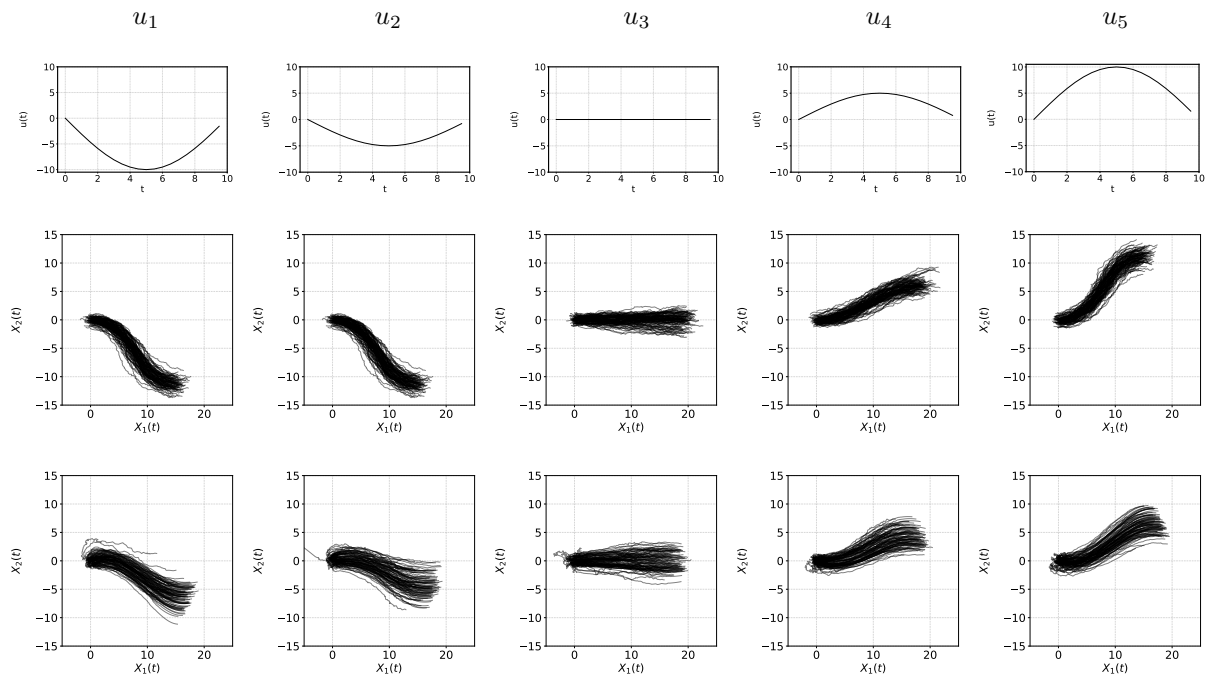


Figure 9: Controlled Dubins process. 100 samples from the true (*center*) and estimated (*bottom*) SDE coefficients for five controls (*top*) with $a = -1, -1/2, 0, 1/2, 1$ (from left to right).

Recovering the true controlled dynamics. For $K_{te} = 5$ controls with $a = -1, -1/2, 0, 1/2, 1$, spanning the set of possible controls \mathcal{H} with $a \in [-1, 1]$, we generate and plot 100 sample paths using both the true and estimated coefficients, as shown in Figure 9, to evaluate the recovery of the true controlled dynamics. The estimated coefficients yield probability distributions that are visually comparable to the true ones, with matching means and variances over time. Notably, the five controls used for evaluation are not included in the training data.

6 Conclusion

In this work, we address the problem of estimating continuous, multidimensional nonlinear controlled SDEs with non-uniform diffusion—a previously unaddressed challenge. We demonstrate how dynamical system identification can be approached through (a) density estimation of the dynamics for a finite set of controls, followed by (b) least-squares regression to estimate governing coefficients, using the Fokker-Planck matching inequality. This formulation enables us to derive strong theoretical guarantees by leveraging the rich and well-established literature on nonparametric least-squares regression. We believe this work lays a promising foundation for future research by providing a robust mathematical framework for identifying complex controlled dynamical systems, supporting further exploration of challenges such as experiment design, including online and safe exploration of the control space with theoretical guarantees. Although we illustrate our findings with prototype identification tasks, future research will focus on real-world applications, such as autonomous driving and space rendezvous.

7 Proofs

7.1 Notations

For the sake of readability, we employ the following notations.

- For any set \mathcal{X} , map $f : \mathcal{X} \times \mathbb{R}^d \rightarrow \mathbb{R}$, and $u : [0, T] \rightarrow \mathbb{R}^d$, we denote $f(u) \triangleq f(\cdot, u(\cdot))$.
- For any set \mathcal{X}, \mathcal{Y} , norm $\|\cdot\|_{\mathcal{Y}}$ over \mathcal{Y} , and bounded map $f : \mathcal{X} \rightarrow \mathcal{Y}$, we denote $\kappa_f \triangleq \sup_{x \in \mathcal{X}} \|f(x)\|_{\mathcal{Y}}$.
- For any map $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and any index $i \in \llbracket 1, n \rrbracket$, the first and second partial derivatives of f are denoted by $f_i \triangleq \frac{\partial f}{\partial x_i}$ and $f_{ij} \triangleq \frac{\partial^2 f}{\partial x_i \partial x_j}$, respectively. Alternatively, if context requires, these derivatives can be denoted as $f_{x_i} \triangleq \frac{\partial f}{\partial x_i}$ and $f_{x_i x_j} \triangleq \frac{\partial^2 f}{\partial x_i \partial x_j}$.

7.2 Organization of the proofs

The proofs are organized as follows.

1. **Proof of the Fokker-Planck matching inequality.** The FP matching inequality is proven in Section 7.3.
2. **Proof of L^2 learning rates.** We present necessary preliminary results in Section 7.4. The L^2 learning rates are then established in Section 7.5, based on four main lemmas detailed in Sections 7.6, 7.7, 7.8, and 7.9, respectively. These main lemmas are proven thanks to auxiliary lemmas, which are proven in Section 7.10. The auxiliary lemmas rely on concentration inequalities adapted to our needs, stated in Section 7.11.
3. **Proof of refined L^2 learning rates.** Refined L^2 learning rates are derived in Section 7.12 using a similar approach to the unrefined L^2 rates, but employing refined lemmas, which are proven in Section 7.13.
4. **Proof of L^∞ learning rates.** The proofs for the L^∞ learning rates are provided in Section 7.14.
5. **Proof of CVaR learning rates.** The proofs for deriving CVaR learning rates are detailed in Section 7.15.
6. **Proofs for Sobolev coefficients.** The embedding property of Fokker-Planck matching, when the coefficients belong to a Sobolev space, is derived in Section 7.16.

7.3 Proof of Lemma 3.2

In this section, we prove the FP matching inequality.

Lemma 7.1 (Fokker-Planck (FP) matching inequality for L^2 learning rates). *Let $(\mathcal{L}^{(b,\sigma)(u)})^*$ denotes the dual of the Kolmogorov generator*

$$(\mathcal{L}^{(b,\sigma)(u)})^* p \triangleq \frac{1}{2} \sum_{i,j=1}^n \frac{\partial^2 \sigma_{ij}^2(u) p}{\partial x_i \partial x_j} - \sum_{i=1}^n \frac{\partial b_i(u) p}{\partial x_i}. \quad (66)$$

Under Assumption (A1), the strong Fokker-Planck equation holds. Namely, we have

$$\frac{\partial p}{\partial t}(u, t, x) = (\mathcal{L}^{(b,\sigma)})^* p(u, t, x) \quad a.e. \quad (u, t, x) \in \mathcal{H} \times [0, T] \times \mathbb{R}^n. \quad (67)$$

Moreover, under Assumption (A3), there exists a constant $c > 0$ such that for any $(\hat{b}, \hat{\sigma}^2) \in \mathcal{F}$,

$$\|p_{\hat{b}, \hat{\sigma}} - p\|_{L^2(p_c)L^2([0, T] \times \mathbb{R}^n)} \leq c \times FP(\hat{b}, \hat{\sigma}^2), \quad (68)$$

where

$$FP(\hat{b}, \hat{\sigma}^2) \triangleq \mathbb{E}_{u \sim p_c, (t, x) \sim p_s} \left[\left(\frac{\partial p}{\partial t}(u, t, x) - \left(\mathcal{L}^{\hat{b}, \hat{\sigma}}(u) \right)^* p(u, t, x) \right)^2 \right]^{1/2}. \quad (FP \text{ matching}) \quad (69)$$

Proof. Define $\rho(u) \triangleq (p_{\hat{b}(u), \hat{\sigma}(u)} - p(u))$. For any $\phi \in C_c^\infty(\mathbb{R}^n, \mathbb{R})$, we have

$$\frac{\partial}{\partial t} \int_{\mathbb{R}^n} \phi(x) \rho(u)(t, x) dx = \frac{\partial}{\partial t} \int_{\mathbb{R}^n} \phi(x) p_{\hat{b}(u), \hat{\sigma}(u)}(t, x) dx - \frac{\partial}{\partial t} \int_{\mathbb{R}^n} \phi(x) p(u, t, x) dx \quad (70)$$

$$= \int_{\mathbb{R}^n} \mathcal{L}^{\hat{b}(u), \hat{\sigma}(u)} \phi(x) p_{\hat{b}(u), \hat{\sigma}(u)}(t, x) dx - \frac{\partial}{\partial t} \int_{\mathbb{R}^n} \phi(x) p(u, t, x) dx \quad (71)$$

$$= \int_{\mathbb{R}^n} \mathcal{L}^{\hat{b}(u), \hat{\sigma}(u)} \phi(x) \rho(u)(t, x) dx + \int_{\mathbb{R}^n} f(u, t, x) \phi(x) dx, \quad (72)$$

where $f(u, t, x) \triangleq \left(\mathcal{L}^{\hat{b}(u), \hat{\sigma}(u)} \right)^* p(u, t, x) - \frac{\partial p}{\partial t}(u, t, x)$.

Since $\rho(u)(0, \cdot) = 0$, applying Theorem 3.4 from Bonalli and Rudi [11] yields

$$\sup_{t \in [0, T]} \left\| p_{\hat{b}(u), \hat{\sigma}(u)}(t, \cdot) - p(u, t, \cdot) \right\|_{L^2(\mathbb{R}^n)}^2 \leq c(u) \left\| \left(\mathcal{L}^{\hat{b}(u), \hat{\sigma}(u)} \right)^* p(u, \cdot, \cdot) - \frac{\partial p}{\partial t}(u, \cdot, \cdot) \right\|_{L^2([0, T] \times \mathbb{R}^n)}^2, \quad (73)$$

where the constant $c(u) > 0$ depends continuously on $\|(\hat{b}, \hat{\sigma}^2)(u)\|_{W^{4+1\frac{1}{2}}}$.

Further, we derive

$$\left\| p_{\hat{b}, \hat{\sigma}} - p \right\|_{L^2(p_c)L^2([0, T] \times \mathbb{R}^n)}^2 \leq c \times \left\| \left(\mathcal{L}^{\hat{b}, \hat{\sigma}} \right)^* p - \frac{\partial p}{\partial t} \right\|_{L^2(p_c)L^2([0, T] \times \mathbb{R}^n)}^2. \quad (74)$$

where $c \triangleq T \sup_{u \in \mathcal{H}} c(u)$.

To conclude, we observe that under Assumptions (A1), (A2), and (A3), the following holds

$$\begin{aligned} & \mathbb{E}_u \left[\int_{t, x} \left(\frac{\partial p}{\partial t}(u, t, x) - \left(\mathcal{L}^{\hat{b}, \hat{\sigma}}(u) \right)^* p(u, t, x) \right)^2 dt dx \right] \\ & \leq \frac{1}{a} \mathbb{E}_{u, t, x} \left[\left(\frac{\partial p}{\partial t}(u, t, x) - \left(\mathcal{L}^{\hat{b}, \hat{\sigma}}(u) \right)^* p(u, t, x) \right)^2 \right], \end{aligned} \quad (75)$$

where we note that for almost every $(u, t, x) \in \mathcal{H} \times [0, T] \times \mathbb{R}^n$, $\frac{\partial p}{\partial t}(u, t, x) = \left(\mathcal{L}^{(b, \sigma)}(u) \right)^* p(u, t, x)$. Additionally $\left(\mathcal{L}^{(b, \sigma)}(u) \right)^* p(u, t, x) - \left(\mathcal{L}^{\hat{b}, \hat{\sigma}}(u) \right)^* p(u, t, x) = \left(\left(\mathcal{L}^{(b, \sigma - \kappa I)}(u) \right)^* - \left(\mathcal{L}^{\hat{b}, \hat{\sigma} - \kappa I}(u) \right)^* \right) p(u, t, x)$. Moreover, we have $\left(\mathcal{L}^{(b, \sigma - \kappa I)}(u) \right)^* p(u, t, x) = \left(\mathcal{L}^{\hat{b}, \hat{\sigma} - \kappa I}(u) \right)^* p(u, t, x) = 0$ for any $(t, x) \notin D$, with $(\hat{b}, \hat{\sigma}^2) \in \mathcal{F}$. \square

Lemma 7.2 (Fokker-Planck (FP) matching inequality for L^∞ learning rates). *Under Assumptions (A1) and (A3), there exists a constant $c > 0$ such that for any $(\hat{b}, \hat{\sigma}^2) \in \mathcal{F}$,*

$$\|p_{\hat{b}, \hat{\sigma}} - p\|_{L^\infty(p_c)L^\infty([0, T]L^2(\mathbb{R}^n))} \leq c \times FP^\infty(\hat{b}, \hat{\sigma}^2), \quad (76)$$

where

$$FP^\infty(\hat{b}, \hat{\sigma}^2) \triangleq \operatorname{ess\,sup}_{u \in \mathcal{H}} \mathbb{E}_{t, x \sim p_s} \left[\left(\frac{\partial p}{\partial t}(u, t, x) - \left(\mathcal{L}^{\hat{b}, \hat{\sigma}}(u) \right)^* p(u, t, x) \right)^2 \right]^{1/2}. \quad (77)$$

Proof. Similar proof as for the Fokker-Planck inequality for L^2 learning rates, noticing that

$$\sup_{u \in \mathcal{H}} \sup_{t \in [0, T]} \left\| p_{\hat{b}(u), \hat{\sigma}(u)}(t, \cdot) - p(u, t, \cdot) \right\|_{L^2(\mathbb{R}^n)}^2 \leq c \times \sup_{u \in \mathcal{H}} \left\| \left(\mathcal{L}^{\hat{b}, \hat{\sigma}}(u) \right)^* p(u) - \frac{\partial p}{\partial t}(u) \right\|_{L^2([0, T] \times \mathbb{R}^n)}^2 \quad (78)$$

where $c \triangleq T \sup_{u \in \mathcal{H}} c(u)$.

□

7.4 Useful preliminary results for proving FP matching learning rates

In this section, we state necessary results to prove FP matching learning rates for the estimator proposed in Section 3 with hard shape-constrained PSD diffusion.

FP matching as least-squares regression. The FP matching problem can be formulated as a least-squares regression problem, expressed as

$$\text{FP}^2(b, \sigma^2) = \mathbb{E}_{u, t, x} \left[\left(\frac{\partial p}{\partial t}(u, t, x) - \langle w, \tilde{\phi}(u, t, x) \rangle_{\mathcal{G}_n} \right)^2 \right], \quad (79)$$

where $\tilde{\phi}(u, t, x)$ is defined as

$$\tilde{\phi}(u, t, x) \triangleq \left((-\tilde{\phi}_i)_{i \in [1, n]} | (1/2\tilde{\phi}_{ij})_{i, j \in [1, n]} \right) (u, t, x) \in \mathcal{G}^n \triangleq \mathcal{G}^n \times (\mathcal{G} \otimes \mathcal{G})^{n^2}, \quad (80)$$

and where

$$\tilde{\phi}_i(u, t, x) \triangleq \left(\phi \frac{\partial p}{\partial x_i} + p \phi_i \right) (u, t, x) \in \mathcal{G}, \quad (81)$$

$$\tilde{\phi}_{ij}(u, t, x) \triangleq \left(\phi \otimes \phi \frac{\partial^2 p}{\partial x_i \partial x_j} + (\phi \otimes \phi)_i \frac{\partial p}{\partial x_j} + (\phi \otimes \phi)_j \frac{\partial p}{\partial x_i} + p(\phi \otimes \phi)_{ij} \right) (u, t, x) \in \mathcal{G} \otimes \mathcal{G}. \quad (82)$$

$\tilde{\phi} : \mathcal{Z} \rightarrow \mathcal{F}$ is well-defined from Lemma 4.34 in Steinwart and Christmann [62].

Empirical FP matching as ridge regression. By denoting $w = (w_i)_{i=1}^n | (w_{ij})_{i, j=1}^n \in \mathcal{G}_n$ and defining

$$\mathcal{S}_n \triangleq \mathcal{G}^n \times \mathcal{S} \subset \mathcal{G}_n, \quad (83)$$

regularized empirical FP matching (Eq. (21)) can be formulated as a ridge regression, expressed as

$$\hat{w}_S \triangleq \arg \min_{w \in \mathcal{S}_n} \frac{1}{KN} \sum_{k=1}^K \sum_{i=1}^N \left(\frac{\partial \hat{p}_k}{\partial t}(t_i, x_i) - \langle w, \hat{\phi}(t_i, x_i) \rangle_{\mathcal{G}_n} \right)^2 + \lambda \|w\|_{\mathcal{G}_n}^2, \quad (84)$$

where $\hat{\phi}(t, x)$ is defined as:

$$\hat{\phi}(t, x) \triangleq \left((-\hat{\phi}_i^k)_{i \in [1, n]} | (1/2\hat{\phi}_{ij}^k)_{i, j \in [1, n]} \right) (t, x, u_k(t)), \quad (85)$$

and where

$$\hat{\phi}_i^k(t, x, v) \triangleq \left(\phi \frac{\partial \hat{p}_k}{\partial x_i} + \hat{p}_k \phi_i \right) (t, x, v), \quad (86)$$

$$\hat{\phi}_{ij}^k(t, x, v) \triangleq \left((\phi \otimes \phi) \frac{\partial^2 \hat{p}_k}{\partial x_i \partial x_j} + (\phi \otimes \phi)_i \frac{\partial \hat{p}_k}{\partial x_j} + (\phi \otimes \phi)_j \frac{\partial \hat{p}_k}{\partial x_i} + \hat{p}_k(\phi \otimes \phi)_{ij} \right) (t, x, v). \quad (87)$$

Useful ridge estimators for Theorem 4.1 proof's decompositions. We define the ridge estimator \hat{w} obtained from Eq. (84) but without the cone constraint. Namely,

$$\hat{w} \triangleq \arg \min_{w \in \mathcal{G}_n} \frac{1}{KN} \sum_{k=1}^K \sum_{i=1}^N \left(\frac{\partial \hat{p}_k}{\partial t}(t_i, x_i) - \langle w, \hat{\phi}(t_i, x_i) \rangle_{\mathcal{G}_n} \right)^2 + \lambda \|w\|_{\mathcal{G}_n}^2. \quad (88)$$

Furthermore, we define the ridge estimator \tilde{w} whose discrepancy from the target w is due to the finite sampling of $\mathcal{H} \times [0, T] \times \mathbb{R}^n$ (excluding estimation errors due to $(\hat{p}_k)_{k=1}^K \approx (p(u_k))_{k=1}^K$). Namely,

$$\tilde{w} \triangleq \arg \min_{w \in \mathcal{G}_n} \frac{1}{KN} \sum_{k=1}^K \sum_{i=1}^N \left(\frac{\partial p}{\partial t}(u_k, t_i, x_i) - \langle w, \tilde{\phi}(u_k, t_i, x_i) \rangle_{\mathcal{G}_n} \right)^2 + \lambda \|w\|_{\mathcal{G}_n}^2. \quad (89)$$

Moreover, we define the ridge estimator w_K , whose discrepancy from the target w arises from the finite sampling of \mathcal{H} . Formally,

$$w_K \triangleq \arg \min_{w \in \mathcal{G}_n} \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\left(\frac{\partial p}{\partial t}(u_k, t, x) - \langle w, \tilde{\phi}(u_k, t, x) \rangle_{\mathcal{G}_n} \right)^2 \right] + \lambda \|w\|_{\mathcal{G}_n}^2. \quad (90)$$

This proposition characterizes the ridge estimators as products and inverses of covariance operators.

Proposition 7.3 (\hat{w}, \tilde{w}, w_K expressions). *For any operator A , define $A_\lambda \triangleq A + \lambda I$, where \otimes denotes the tensor product. \hat{w} yields the standard ridge regression solution*

$$\hat{w} = \hat{D} \hat{C}_\lambda^{-1}, \quad (91)$$

with $\hat{D} \triangleq \frac{1}{KN} \sum_{k=1}^K \sum_{i=1}^N \frac{\partial \hat{p}_k}{\partial t}(t_i, x_i) \hat{\phi}_k(t_i, x_i)$, and $\hat{C} \triangleq \frac{1}{KN} \sum_{k=1}^K \sum_{i=1}^N \hat{\phi}_k(t_i, x_i) \otimes \hat{\phi}_k(t_i, x_i)$.

Similarly, $\tilde{w} = \tilde{D} \tilde{C}_\lambda^{-1}$ with

$$\tilde{D} \triangleq \frac{1}{NK} \sum_{k=1}^K \sum_{i=1}^N \frac{\partial p}{\partial t}(u_k, t_i, x_i) \tilde{\phi}_k(t_i, x_i), \quad \text{and} \quad \tilde{C} \triangleq \frac{1}{NK} \sum_{k=1}^K \sum_{i=1}^N \tilde{\phi}_k(t_i, x_i) \otimes \tilde{\phi}_k(t_i, x_i). \quad (92)$$

We also have $w_K = D_K C_{K,\lambda}^{-1}$ with

$$D_K \triangleq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\frac{\partial p}{\partial t}(u_k, t, x) \tilde{\phi}(u_k, t, x) \right], \quad \text{and} \quad C_K \triangleq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\tilde{\phi}(u_k, t, x) \otimes \tilde{\phi}(u_k, t, x) \right]. \quad (93)$$

Proposition 7.4 (Attainability of $\frac{\partial p}{\partial t}$ from Assumptions (A1) and (A4)). *Assumption (A4) is equivalent to the existence of $W_b \in \mathbb{R}^n \otimes \mathcal{G}$ and $W_\sigma \in \mathbb{R}^{n^2} \otimes (\mathcal{G} \otimes \mathcal{G})$ such that*

$$(b, \sigma^2) = (W_b \phi(\cdot), W_\sigma(\phi(\cdot) \otimes \phi(\cdot))) \quad \text{with} \quad \|W_b\|_{\text{HS}} < +\infty, \|W_\sigma\|_{\text{HS}} < +\infty, \quad (94)$$

and such that the representation $[W_\sigma]_{\mathcal{G}^n \otimes \mathcal{G}^n}$ of W_σ in $\mathcal{G}^n \otimes \mathcal{G}^n$ satisfies

$$[W_\sigma]_{\mathcal{G}^n \otimes \mathcal{G}^n} \in \mathcal{S}. \quad (95)$$

Assumption (A1) further ensures the existence of $w \in \mathcal{S}_n$ such that

$$\frac{\partial p}{\partial t}(u, t, x) = \left((\mathcal{L}^{b,\sigma})^* p \right)(u, t, x) = \langle w, \tilde{\phi}(u, t, x) \rangle_{\mathcal{G}_n}. \quad (96)$$

Furthermore, the Fokker-Planck matching L^2 -risk can be written as follows (see Ciliberto et al. [16])

$$FP(\hat{b}, \hat{\sigma}^2) = \|(\hat{w}_S - w)C^{1/2}\|_{\mathcal{G}_n}, \quad (97)$$

with $C \triangleq \mathbb{E}_{u,t,x} [\tilde{\phi}(u, t, x) \otimes \tilde{\phi}(u, t, x)]$.

Remark 12 (Alternative assumptions to Assumptions (A6)). We denote $\tilde{\phi} \triangleq \tilde{\phi}(u, t, x)$. The three following assumptions can be found in the literature of least-squares regression as alternative conditions on the regularity of the feature space

- (1) There exists $r_1 \in [0, 1]$ and $c > 0$ such that

$$\text{Tr}(C^{1-r_1}) < c. \quad (98)$$

- (2) There exists $r_2 \in [0, 1]$ and $c > 0$ such that

$$\mathbb{E}[(\tilde{\phi}, C^{-r_2} \tilde{\phi})_{\mathcal{G}_n} \tilde{\phi} \otimes \tilde{\phi}] \preceq cC. \quad (99)$$

- (3) There exists $r_3 \in [0, 1]$ and $c > 0$ such that

$$\|C^{-r_3/2} \tilde{\phi}\|_{\mathcal{G}_n} < c \quad \text{almost surely.} \quad (100)$$

Assumptions (1) and (3) are known as the capacity condition and embedding property, respectively [22, 50]. Assumptions (1), (2), and (3) provide progressively finer conditions on the regularity of the features $\tilde{\phi}(u, t, x)$. Indeed, it is straightforward to show that Assumption (1) implies Assumption (2) with $r_2 = r_1$, and Assumption (2) implies Assumption (3) with $r_3 = r_2$ (see Remark 3 in Berthier et al. [10]). In the setting of noisy kernel ridge regression, Assumption (1) allows to refine learning rates from $n^{-1/4}$ ($r_1 = 0$) to $n^{-1/2}$ ($r_1 = 1$). In the setting of noiseless kernel ridge regression, Assumption (2) allows to refine learning rates from $n^{-1/2}$ ($r_2 = 0$) to n^{-1} ($r_2 = 1$), while Assumption (3) allows to refine them from n^{-1} ($r_3 = 0$) to arbitrarily fast polynomial decay as $r_3 \rightarrow +\infty$. For completeness, we present these three assumptions here; however, for clarity, in this work, we base our proofs on the more stringent Assumption (3) (instead of considering Assumptions (2) and (3)).

7.5 Proof of Theorem 4.1

Theorem 7.5 (L^2 Learning rates). For any $N, K \in \mathbb{N}^*$, let $(\hat{b}, \hat{\sigma}^2)$ be the proposed estimator, with hard shape-constrained PSD diffusion, trained with K controls $(u_k)_{k=1}^K$ i.i.d. from a probability measure p_c over \mathcal{H} , and N points $(t_i, x_i)_{i=1}^N$ i.i.d. from a probability measure p_s over $[0, T] \times \mathbb{R}^n$. Assuming Assumptions (A1), (A4), and (A3) hold true, then there exist constants $c_1, c_2 > 0$ that do not depend on N, K , or δ , such that for any $\delta \in (0, 1]$, defining $\varepsilon \triangleq \sup_{k \in \llbracket 1, K \rrbracket} L(\hat{p}_k, p(u_k))$ where L is a squared Sobolev norm defined as

$$L(p_1, p_2) \triangleq \int_0^T \left\| \frac{\partial p_1}{\partial t}(t, \cdot) - \frac{\partial p_2}{\partial t}(t, \cdot) \right\|_{L^2(\mathbb{R}^n)}^2 + \|p_1(t, \cdot) - p_2(t, \cdot)\|_{W^2(\mathbb{R}^n)}^2 dt, \quad (101)$$

and setting $\lambda = c_2(N^{-1} \log(N\delta^{-1}) + K^{-1} \log(K\delta^{-1}))$, if

$$\varepsilon \leq \frac{1}{N} \log \frac{N}{\delta} + \frac{1}{K} \log \frac{K}{\delta}, \quad (102)$$

then with probability $1 - \delta$,

$$\|p_{\hat{b}, \hat{\sigma}} - p\|_{L^2(p_c)L^2([0, T] \times \mathbb{R}^n)} \leq c_1 \log \frac{2}{\delta} \left(\frac{\log \frac{N}{\delta}}{\sqrt{N}} + \frac{\log \frac{K}{\delta}}{\sqrt{K}} \right). \quad (103)$$

Sketch of the proof. The proof involves decomposing the error of the proposed estimator into three components:

1. Error stemming from the estimation of the probability densities $(\hat{p}_k)_{k=1}^K \approx (p(u_k))_{k=1}^K$,
2. Error due to the finite sampling approximation of times $[0, T]$ and positions \mathbb{R}^n ,

3. Error due to the finite sampling approximation of the space of controls \mathcal{H} ,

and then bounding each error component.

Denoting $z = (t, x)$, $\tilde{\phi} = \tilde{\phi}(u, t, x)$, and $\hat{\mathbb{E}}[\cdot]$ as the empirical expectation over the training points, this corresponds to the following successive approximations.

$$\hat{C} \triangleq \hat{\mathbb{E}}_u \hat{\mathbb{E}}_z[\hat{\phi} \otimes \hat{\phi}] \quad \rightarrow \quad \tilde{C} \triangleq \hat{\mathbb{E}}_u \hat{\mathbb{E}}_z[\tilde{\phi} \otimes \tilde{\phi}] \quad \rightarrow \quad C_K \triangleq \hat{\mathbb{E}}_u \mathbb{E}_z[\tilde{\phi} \otimes \tilde{\phi}] \quad \rightarrow \quad C \triangleq \mathbb{E}_u \mathbb{E}_z[\tilde{\phi} \otimes \tilde{\phi}]. \quad (104)$$

□

Proof. The definitions of Section 7.4 allow for the following decomposition

$$\text{FP}(\hat{b}, \hat{\sigma}^2) = \|(\hat{w} - w)C^{1/2}\|_{\mathcal{G}_n} \leq \|(\hat{w} - \tilde{w})C^{1/2}\|_{\mathcal{G}_n} + \|(\tilde{w} - w_K)C^{1/2}\|_{\mathcal{G}_n} + \|(w_K - w)C^{1/2}\|_{\mathcal{G}_n}. \quad (105)$$

Using Lemmas 7.6, 7.7, 7.8, and 7.9, and Lemma 7.14, we determine that for any $\delta \in (0, 1]$, if $\lambda \geq c_2 \log \frac{2}{\delta}(N^{-1} + \varepsilon)$, $\lambda \geq \frac{9\kappa_{\tilde{\phi}}^2}{N} \log \frac{N}{\delta}$, and $\lambda \geq \frac{9\kappa_{\tilde{\phi}}^2}{K} \log \frac{K}{\delta}$, with probability $1 - \delta$,

$$\text{FP}(\hat{b}, \hat{\sigma}^2) \leq c_1 \left((N^{-1} + \varepsilon)\lambda^{-1/2} \log \frac{2}{\delta} + \frac{\log \frac{N}{\delta}}{\sqrt{N}} + \frac{\log \frac{K}{\delta}}{\sqrt{K}} + \lambda^{1/2} \right), \quad (106)$$

where constants $c_1, c_2 > 0$ are independent of N, K , and δ . Therefore, setting $\lambda = \frac{9\kappa_{\tilde{\phi}}^2}{N} \log^2 \frac{N}{\delta} + \frac{9\kappa_{\tilde{\phi}}^2}{K} \log^2 \frac{K}{\delta}$, we find

$$\text{FP}(\hat{b}, \hat{\sigma}^2) \leq c_1 \log \frac{2}{\delta} \left(\frac{\varepsilon}{\lambda^{1/2}} + \frac{\log \frac{N}{\delta}}{\sqrt{N}} + \frac{\log \frac{K}{\delta}}{\sqrt{K}} \right). \quad (107)$$

up to overloading $c_1 > 0$. Then, assuming $\varepsilon \leq \frac{1}{N} \log^2 \frac{N}{\delta} + \frac{1}{K} \log^2 \frac{K}{\delta}$, the condition $\lambda \geq c_2 \log \frac{2}{\delta}(N^{-1} + \varepsilon)$ holds true up to multiplying λ by a constant, and overloading the constant c_1 . Finally, Lemma 3.2 allows to conclude the proof. □

7.6 Proof of Lemma 7.6

This lemma shows that the error of \hat{w}_S can be bounded by that of the unconstrained ridge estimator \hat{w} . This stems from the fact that using the same empirical learning objective with a subset $\mathcal{S}_n \subset \mathcal{G}_n$, rather than the full set \mathcal{G}_n , cannot increase the error, provided that $w \in \mathcal{S}_n$.

Lemma 7.6. *Under Assumption (A4), the following inequality holds true.*

$$\|(\hat{w}_S - \hat{w})C^{1/2}\|_{\mathcal{G}_n} \leq \|(\hat{w} - w)C^{1/2}\|_{\mathcal{G}_n} \left(\|C_{\lambda}^{-1/2} \hat{C}_{\lambda}^{1/2}\|_{\infty} + \|\hat{C}_{\lambda}^{-1/2} C^{1/2}\|_{\infty} \right). \quad (108)$$

Proof. We have

$$\|(\hat{w}_S - w)C^{1/2}\|_{\mathcal{G}_n} \leq \|(\hat{w}_S - \hat{w})C^{1/2}\|_{\mathcal{G}_n} + \|(\hat{w} - w)C^{1/2}\|_{\mathcal{G}_n}. \quad (109)$$

Then,

$$\|(\hat{w}_S - \hat{w})C^{1/2}\|_{\mathcal{G}_n} \leq \|(\hat{w}_S - \hat{w})\hat{C}_{\lambda}^{1/2}\|_{\mathcal{G}_n} \|\hat{C}_{\lambda}^{-1/2} C^{1/2}\|_{\infty}. \quad (110)$$

Furthermore, it is straightforward to derive from the definition of \hat{w}_S that

$$\hat{w}_S = \arg \min_{\hat{w}_S \in \mathcal{S}_n} \|(\hat{w}_S - w)\hat{C}_{\lambda}^{-1/2} \hat{C}_{\lambda}^{1/2}\|_{\mathcal{G}_n}. \quad (111)$$

However, from Assumption (A4), we have $\hat{D} = w\hat{C}_\lambda^{-1}$. Hence,

$$\hat{w}_S = \arg \min_{\hat{w}_S \in \mathcal{S}_n} \|(\hat{w}_S - \hat{w})\hat{C}_\lambda^{1/2}\|_{\mathcal{G}_n}. \quad (112)$$

Since $w \in \mathcal{S}_n$, we have

$$\|(\hat{w}_S - \hat{w})\hat{C}_\lambda^{1/2}\|_{\mathcal{G}_n} \leq \|(w - \hat{w})\hat{C}_\lambda^{1/2}\|_{\mathcal{G}_n} \leq \|(w - \hat{w})C_\lambda^{1/2}\|_{\mathcal{G}_n} \|C_\lambda^{-1/2}\hat{C}_\lambda^{1/2}\|_\infty. \quad (113)$$

We conclude that

$$\|(\hat{w}_S - \hat{w})C^{1/2}\|_{\mathcal{G}_n} \leq \|(\hat{w} - w)C_\lambda^{1/2}\|_{\mathcal{G}_n} \left(\|C_\lambda^{-1/2}\hat{C}_\lambda^{1/2}\|_\infty + \|\hat{C}_\lambda^{-1/2}C^{1/2}\|_\infty \right). \quad (114)$$

□

7.7 Proof of Lemma 7.7

This lemma bounds the error stemming from the estimation of the probability densities $(\hat{p}_k)_{k=1}^K \approx (p(u_k))_{k=1}^K$.

Lemma 7.7 (Bound $\|(\hat{w} - \tilde{w})C^{1/2}\|_{\mathcal{G}_n}$). *There exist constants $c_1, c_2 > 0$ that do not depend on N, K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$*

$$\|(\hat{w} - \tilde{w})C^{1/2}\|_{\mathcal{G}_n} \leq c_1 \log \frac{2}{\delta} (N^{-1} + \varepsilon) \lambda^{-1/2}, \quad (115)$$

provided that $\lambda \geq c_2 \log \frac{2}{\delta} (N^{-1} + \varepsilon)$, $\lambda \geq \frac{18c}{N} \log \frac{N}{\delta}$ and $\lambda \geq \frac{18c}{K} \log \frac{K}{\delta}$.

Proof. We have

$$\|(\hat{w} - \tilde{w})C^{1/2}\|_{\mathcal{G}_n} = \|(\hat{D}\hat{C}_\lambda^{-1} - \tilde{D}\tilde{C}_\lambda^{-1})C^{1/2}\|_{\mathcal{G}_n} \quad (116)$$

$$\leq \|(\hat{D} - \tilde{D})\hat{C}_\lambda^{-1}C^{1/2}\|_{\mathcal{G}_n} + \|\tilde{D}(\hat{C}_\lambda^{-1} - \tilde{C}_\lambda^{-1})C^{1/2}\|_{\mathcal{G}_n}. \quad (117)$$

Moreover,

$$\|(\hat{D} - \tilde{D})\hat{C}_\lambda^{-1}C^{1/2}\|_{\mathcal{G}_n} \leq \|\hat{D} - \tilde{D}\|_{\mathcal{G}_n} \|\hat{C}_\lambda^{-1}C^{1/2}\|_\infty. \quad (118)$$

Furthermore, using $A^{-1} - B^{-1} = B^{-1}(B - A)A^{-1}$, we have

$$\|\tilde{D}(\hat{C}_\lambda^{-1} - \tilde{C}_\lambda^{-1})C^{1/2}\|_{\mathcal{G}_n} = \|\tilde{D}\tilde{C}_\lambda^{-1}(\tilde{C} - \hat{C})\hat{C}_\lambda^{-1}C^{1/2}\|_{\mathcal{G}_n}. \quad (119)$$

1. Bound $\|\hat{D} - \tilde{D}\|_{\mathcal{G}_n}$. From Lemma 7.11, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$

$$\|\hat{D} - \tilde{D}\|_{\mathcal{G}_n} \leq c \log \frac{2}{\delta} (N^{-1} + \varepsilon), \quad (120)$$

where $c > 0$ is a constant that does not depend on N, K, δ .

2. Bound $\|\hat{C}_\lambda^{-1}C^{1/2}\|_\infty$. We have

$$\|\hat{C}_\lambda^{-1}C^{1/2}\|_\infty \leq \lambda^{-1/2} \|\hat{C}_\lambda^{-1/2}\tilde{C}_\lambda^{1/2}\|_\infty \|\tilde{C}_\lambda^{-1/2}C^{1/2}\|_\infty. \quad (121)$$

By Lemmas 7.12 and 7.13, for any $\delta \in (0, 1]$, if $\lambda \geq c \log \frac{2}{\delta} (N^{-1} + \varepsilon)$, $\lambda \geq \frac{18c}{N} \log \frac{N}{\delta}$ and $\lambda \geq \frac{18c}{K} \log \frac{K}{\delta}$, then with probability at least $1 - \delta$,

$$\|\hat{C}_\lambda^{-1}C^{1/2}\|_\infty \leq 2\sqrt{2}\lambda^{-1/2}, \quad (122)$$

where $c > 0$ is a constant that does not depend on N, K, δ .

3. Bound $\|\tilde{D}\tilde{C}_\lambda^{-1}(\tilde{C} - \hat{C})\hat{C}_\lambda^{-1}C^{1/2}\|_{\mathcal{G}_n}$. From Assumption (A4), we have

$$\tilde{D} = w\tilde{C}. \quad (123)$$

Therefore,

$$\|\tilde{D}\tilde{C}_\lambda^{-1}(\tilde{C} - \hat{C})\hat{C}_\lambda^{-1}C^{1/2}\|_{\mathcal{G}_n} \leq \|w\|_{\mathcal{G}_n} \|\tilde{C}\tilde{C}_\lambda^{-1}(\tilde{C} - \hat{C})\hat{C}_\lambda^{-1}C^{1/2}\|_\infty \quad (124)$$

$$\leq \|w\|_{\mathcal{G}_n} \|\tilde{C} - \hat{C}\|_\infty \|\hat{C}_\lambda^{-1}C^{1/2}\|_\infty. \quad (125)$$

Then, as for bound 2., from Lemma 7.12 and Lemma 7.13, we have

$$\|\tilde{D}\tilde{C}_\lambda^{-1}(\tilde{C} - \hat{C})\hat{C}_\lambda^{-1}C^{1/2}\|_{\mathcal{G}_n} \leq c\lambda^{-1/2} \log \frac{2}{\delta} (N^{-1} + \varepsilon), \quad (126)$$

where $c > 0$ is a constant that does not depend on N, K, δ .

Conclusion. Combining all bounds, we conclude that there exist constants $c_1, c_2 > 0$ that do not depend on N, K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$

$$\|(\hat{w} - \tilde{w})C^{1/2}\|_{\mathcal{G}_n} \leq c_1 \log \frac{2}{\delta} (N^{-1} + \varepsilon) \lambda^{-1/2}, \quad (127)$$

if $\lambda \geq c_2 \log \frac{2}{\delta} (N^{-1} + \varepsilon)$, $\lambda \geq \frac{18c}{N} \log \frac{N}{\delta}$ and $\lambda \geq \frac{18c}{K} \log \frac{K}{\delta}$. \square

7.8 Proof of Lemma 7.8

This lemma bounds the error resulting from the finite sampling approximation of $[0, T] \times \mathbb{R}^n$.

Lemma 7.8 (Bound $\|(w_K - w)C^{1/2}\|_{\mathcal{G}_n}$). *Under Assumption (A4), there exist a constant $c > 0$ that does not depend on N, K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$*

$$\|(\tilde{w} - w_K)C^{1/2}\|_{\mathcal{G}_n} \leq c \frac{\log \frac{N}{\delta}}{\sqrt{N}}, \quad (128)$$

provided that $\lambda \geq \frac{9\kappa_\phi^2}{N} \log \frac{N}{\delta}$ and $\lambda \geq \frac{9\kappa_\phi^2}{K} \log \frac{K}{\delta}$.

Proof. From Assumption (A4), we have

$$\tilde{D} = w\tilde{C} \quad \text{and} \quad D_K = wC_K. \quad (129)$$

The error norm can be decomposed as follows

$$\|(\tilde{w} - w_K)C^{1/2}\|_{\mathcal{G}_n} = \|(\tilde{D}\tilde{C}_\lambda^{-1} - D_K C_{K,\lambda}^{-1})C^{1/2}\|_{\mathcal{G}_n} \quad (130)$$

$$\leq \|w\|_{\mathcal{G}_n} \|(\tilde{C}\tilde{C}_\lambda^{-1} - C_K C_{K,\lambda}^{-1})C^{1/2}\|_\infty. \quad (131)$$

Moreover, using $A^{-1} - B^{-1} = B^{-1}(B - A)A^{-1}$, we have

$$\|(\tilde{C}\tilde{C}_\lambda^{-1} - C_K C_{K,\lambda}^{-1})C^{1/2}\|_\infty = \lambda \|(\tilde{C}_\lambda^{-1} - C_{K,\lambda}^{-1})C^{1/2}\|_\infty \quad (132)$$

$$= \lambda \|\tilde{C}_\lambda^{-1}(C_K - \tilde{C})C_{K,\lambda}^{-1}C^{1/2}\|_\infty \quad (133)$$

$$\leq \lambda^{1/2} \|\tilde{C}_\lambda^{-1/2}C_{K,\lambda}^{1/2}\|_\infty \|C_{K,\lambda}^{-1/2}(\tilde{C} - C_K)C_{K,\lambda}^{-1/2}\|_\infty \|C_{K,\lambda}^{1/2}C^{1/2}\|_\infty. \quad (134)$$

1. Bound $\|C_{K,\lambda}^{-1/2}(\tilde{C} - C_K)C_{K,\lambda}^{-1/2}\|_\infty$. From Proposition 7.17, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$,

$$\|C_{K,\lambda}^{-1/2}(\tilde{C} - C_K)C_{K,\lambda}^{-1/2}\|_\infty \leq \frac{2\beta\kappa_\phi^2}{\lambda N} + \sqrt{\frac{2\kappa_\phi^2\beta}{3\lambda N}}, \quad (135)$$

with $\beta = \log \frac{4(\|C_K\|_\infty + \lambda)\text{Tr}(C_K)}{\delta\lambda\|C_K\|_\infty}$. Moreover,

$$\|C_K\|_\infty \leq \text{Tr}(C_K) = \frac{1}{K} \sum_{k=1}^K \mathbb{E}[\|\tilde{\phi}(u_k, t, x)\|_{\mathcal{G}_n}^2] \leq \kappa_\phi^2. \quad (136)$$

Hence, if $N^{-1} \leq \lambda \leq \|C_K\|_\infty$, it is straightforward to verify that $\beta \leq \log(4\kappa_\phi^2\lambda^{-1}\delta^{-1})$, and that there exists a constant $c > 0$ that does not depend on δ, N, K such that

$$\|C_{K,\lambda}^{-1/2}(\tilde{C} - C_K)C_{K,\lambda}^{-1/2}\|_\infty \leq c \frac{\log \frac{N}{\delta}}{\sqrt{\lambda N}}. \quad (137)$$

Moreover,

$$\|C_K\|_\infty \geq \|C\|_\infty - \|C_K - C\|_\infty, \quad (138)$$

and same proof as Proposition 7.17 gives $\|C_K - C\|_\infty \leq cK^{-1/2} \log(K\delta^{-1})$ for a constant $c > 0$ that does not depend on K, δ , if $K^{-1} \leq \lambda \leq \|C\|_\infty$, such that

$$\lambda \leq \|C\|_\infty - cK^{-1/2} \log(K\delta^{-1}), \quad (139)$$

ensures $\lambda \leq \|C_K\|_\infty$.

2. Bound $\|\tilde{C}_\lambda^{-1/2}C_{K,\lambda}^{1/2}\|_\infty\|C_{K,\lambda}^{1/2}C^{1/2}\|_\infty$. From Lemma 7.13, for any $\delta \in (0, 1]$, if $\lambda \geq \frac{9\kappa_\phi^2}{N} \log \frac{N}{\delta}$ and $\lambda \geq \frac{9\kappa_\phi^2}{K} \log \frac{K}{\delta}$, with probability at least $1 - \delta$,

$$\|\tilde{C}_\lambda^{-1/2}C_{K,\lambda}^{1/2}\|_\infty\|C_{K,\lambda}^{1/2}C^{1/2}\|_\infty \leq 2. \quad (140)$$

Conclusion. We conclude by combining all bounds that there exist a constant $c > 0$ that does not depend on N, K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$

$$\|(\tilde{w} - w_K)C^{1/2}\|_{\mathcal{G}_n} \leq c \frac{\log \frac{N}{\delta}}{\sqrt{N}}, \quad (141)$$

if $\lambda \geq \frac{9\kappa_\phi^2}{N} \log \frac{N}{\delta}$ and $\lambda \geq \frac{9\kappa_\phi^2}{K} \log \frac{K}{\delta}$. \square

7.9 Proof of Lemma 7.9

This lemma bounds the error stemming from the finite sampling approximation of \mathcal{H} , and the RKHS norm regularization.

Lemma 7.9 (Bound $\|(\tilde{w} - w)C^{1/2}\|_{\mathcal{G}_n}$). *Under assumption (A4), there exist a constant $c > 0$ that does not depend on K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$*

$$\|(w_K - w)C^{1/2}\|_{\mathcal{G}_n} \leq c \left(\frac{\log \frac{K}{\delta}}{\sqrt{K}} + \lambda^{1/2} \right), \quad (142)$$

if $\lambda \geq \frac{9\kappa_\phi^2}{K} \log \frac{K}{\delta}$.

Proof. We have

$$\|(\tilde{w} - w)C^{1/2}\|_{\mathcal{G}_n} \leq \|(\tilde{w} - w_\lambda)C^{1/2}\|_{\mathcal{G}_n} + \|(w_\lambda - w)C^{1/2}\|_{\mathcal{G}_n}, \quad (143)$$

defining $w_\lambda = DC_\lambda^{-1}$.

1. Bound $\|(w_K - w_\lambda)C^{1/2}\|_{\mathcal{G}_n}$ Employing the same reasoning as in Lemma 7.8, there exists a constant $c > 0$ that does not depend on K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$

$$\|(\tilde{w} - w_\lambda)C^{1/2}\|_{\mathcal{G}_n} \leq c \frac{\log \frac{K}{\delta}}{\sqrt{K}}, \quad (144)$$

if $\lambda \geq \frac{9\kappa_\tau^2}{K} \log \frac{K}{\delta}$.

2. Bound $\|(w_\lambda - w)C^{1/2}\|_{\mathcal{G}_n}$. From Assumption (A4), we have

$$D = wC. \quad (145)$$

Therefore,

$$\|(w_\lambda - w)C^{1/2}\|_{\mathcal{G}_n} = \|w(DCC_\lambda^{-1} - I)C^{1/2}\|_{\mathcal{G}_n} = \lambda \|wC_\lambda^{-1}C^{1/2}\|_{\mathcal{G}_n} \leq \lambda^{1/2} \|w\|_{\mathcal{G}_n}. \quad (146)$$

Conclusion. We conclude by combining all bounds, that there exist a constant $c > 0$ that does not depend on K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$

$$\|(w_K - w)C^{1/2}\|_{\mathcal{G}_n} \leq c \left(\frac{\log \frac{K}{\delta}}{\sqrt{K}} + \lambda^{1/2} \right), \quad (147)$$

if $\lambda \geq \frac{9\kappa_\tau^2}{K} \log \frac{K}{\delta}$.

□

7.10 Auxiliary lemmas

This section presents auxiliary lemmas used in the proofs of the main lemmas, along with their proofs.

This lemma bounds the L^2 discrepancy between each $\hat{\phi}_k(t, x)$ and $\tilde{\phi}(u_k, \cdot, \cdot)$ because of the use of \hat{p}_k instead of $p(u_k)$ in the expression of $\hat{\phi}_k(t, x)$.

Lemma 7.10 (Bound $\mathbb{E}_{t,x}[\|\tilde{\phi}(u_k, t, x) - \hat{\phi}_k(t, x)\|_{\mathcal{G}_n}^2]$). *For any $k \in \llbracket 1, K \rrbracket$, the following inequality holds*

$$\mathbb{E}_{t,x} \left[\left\| \tilde{\phi}(u_k, t, x) - \hat{\phi}_k(t, x) \right\|_{\mathcal{G}_n}^2 \right] \leq \kappa_{tot} L(\hat{p}_k, p(u_k)). \quad (148)$$

with $L(p_1, p_2) \triangleq \int_{t=0}^T \left(\left\| \frac{\partial p_1}{\partial t}(t, \cdot) - \frac{\partial p_2}{\partial t}(t, \cdot) \right\|_{L^2}^2 + \|p_1(t, \cdot) - p_2(t, \cdot)\|_{W^2}^2 \right) dt$,

and $\kappa_{tot} \triangleq \sum_{i,j=1}^n \left(\kappa_\phi^2 + \kappa_{\phi_i}^2 + \kappa_\phi^4 + \kappa_{\phi_i}^4 + \kappa_{\phi_j}^4 + \kappa_{\phi_{ij}}^4 \right)$.

Proof. We have

$$\begin{aligned} \|\tilde{\phi}(u_k, t, x) - \hat{\phi}_k(t, x)\|_{\mathcal{G}_n}^2 &= \sum_{i=1}^n \|\tilde{\phi}_i(u_k, t, x) - \hat{\phi}_i^k(t, x, u_k(t))\|_{\mathcal{G}_n}^2 \\ &\quad + \frac{1}{4} \sum_{i,j=1}^n \|\tilde{\phi}_{ij}(u_k, t, x) - \hat{\phi}_{ij}^k(t, x, u_k(t))\|_{\mathcal{G}_n}^2, \end{aligned} \quad (149)$$

and

$$\begin{aligned}
 \|\tilde{\phi}_i(u_k, t, x) - \hat{\phi}_i^k(t, x, u_k(t))\|_{\mathcal{G}_n}^2 &\leq 2\|\phi(t, x, u_k(t))\|_{\mathcal{G}_n}^2 \left| \frac{\partial p(u_k)}{\partial x_i}(t, x) - \frac{\partial \hat{p}_k(t, x)}{\partial x_i} \right|^2 \\
 &\quad + 2\|\phi_i(t, x, u_k(t))\|_{\mathcal{G}_n}^2 |p(u_k)(t, x) - \hat{p}_k(t, x)|^2 \\
 &\leq 2(\kappa_\phi^2 + \kappa_{\phi_i}^2) \left(\left| \frac{\partial p(u_k)}{\partial x_i}(t, x) - \frac{\partial \hat{p}_k(t, x)}{\partial x_i} \right|^2 \right. \\
 &\quad \left. + |p(u_k)(t, x) - \hat{p}_k(t, x)|^2 \right). \tag{150}
 \end{aligned}$$

Then, one can obtain a similar bound for $\|\tilde{\phi}_{ij}(u_k, t, x) - \hat{\phi}_{ij}^k(t, x, u_k(t))\|_{\mathcal{G}_n}^2$ involving a sum of the difference of the second order derivatives with respect to x . Using these two bounds, we obtain

$$\mathbb{E}_{t,x} \left[\left\| \tilde{\phi}(u_k, t, x) - \hat{\phi}_k(t, x) \right\|_{\mathcal{G}_n}^2 \right] \leq L(\hat{p}_k, p(u_k)) \times \sum_{i,j=1}^n \left(\kappa_\phi^2 + \kappa_{\phi_i}^2 + \kappa_\phi^4 + \kappa_{\phi_i}^4 + \kappa_{\phi_j}^4 + \kappa_{\phi_{ij}}^4 \right). \tag{151}$$

□

Lemma 7.10 allows us to derive the three following lemmas which bounds the two error terms appearing in the main lemmas' proofs, which stem from the estimation \hat{p}_k of $p(u_k)$.

Lemma 7.11 (Bound $\|\hat{D} - \tilde{D}\|_{\mathcal{G}_n}$). *For any $\delta \in (0, 1]$, with probability at least $1 - \delta$*

$$\|\hat{D} - \tilde{D}\|_{\mathcal{G}_n} \leq c \log \frac{2}{\delta} (N^{-1} + \varepsilon N^{-1/2} + \varepsilon), \tag{152}$$

where $c > 0$ is a constant that does not depend on N, K, δ .

Proof. We have

$$\hat{D} - \tilde{D} = \frac{1}{KN} \sum_{i=1}^N \sum_{k=1}^K \left(\frac{\partial \hat{p}_k}{\partial t}(t_i, x_i) \hat{\phi}_k(t_i, x_i) - \frac{\partial p}{\partial t}(u_k, t_i, x_i) \tilde{\phi}(u_k, t_i, x_i) \right). \tag{153}$$

Hence, defining

$$\hat{E}_1 \triangleq \frac{1}{KN} \sum_{i=1}^N \sum_{k=1}^K \left(\frac{\partial \hat{p}_k}{\partial t}(t_i, x_i) - \frac{\partial p}{\partial t}(u_k, t_i, x_i) \right) \hat{\phi}_k(t_i, x_i), \tag{154}$$

$$\hat{E}_2 \triangleq \frac{1}{KN} \sum_{i=1}^N \sum_{k=1}^K \frac{\partial p}{\partial t}(u_k, t_i, x_i) \left(\hat{\phi}_k(t_i, x_i) - \tilde{\phi}(u_k, t_i, x_i) \right), \tag{155}$$

$$E_1 \triangleq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\left(\frac{\partial \hat{p}_k}{\partial t}(t, x) - \frac{\partial p}{\partial t}(u_k, t, x) \right) \hat{\phi}_k(t, x) \right], \tag{156}$$

$$E_2 \triangleq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\frac{\partial p}{\partial t}(u_k, t, x) \left(\hat{\phi}_k(t, x) - \tilde{\phi}(u_k, t, x) \right) \right], \tag{157}$$

we have

$$\|\hat{D} - \tilde{D}\|_{\mathcal{G}_n} \leq \|\hat{E}_1 - E_1\|_{\mathcal{G}_n} + \|\hat{E}_2 - E_2\|_{\mathcal{G}_n} + \|E_1\|_{\mathcal{G}_n} + \|E_2\|_{\mathcal{G}_n}. \tag{158}$$

1. Bound $\|\hat{E}_1 - E_1\|_{\mathcal{G}_n}$. To apply Proposition 7.17, we define

$$\psi(t, x, k) \triangleq \left(\frac{\partial \hat{p}_k}{\partial t}(t, x) - \frac{\partial p}{\partial t}(u_k, t, x) \right), \tag{159}$$

$$\chi(t, x, k) \triangleq \hat{\phi}_k(t, x), \tag{160}$$

and

$$\kappa_\psi \triangleq \sup_{t,x,k} |\psi(t,x,k)| \leq \kappa_{\frac{\partial \hat{p}_k}{\partial t}} + \kappa_{\frac{\partial p}{\partial t}}, \quad (161)$$

$$\kappa_\chi \triangleq \sup_{t,x,k} \|\hat{\phi}_k(t,x)\|_{\mathcal{G}_n}, \quad (162)$$

and

$$\sigma^2 \triangleq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} [\|\psi(t,x,k) \otimes \chi(t,x,k)\|_{\text{HS}}^2] \quad (163)$$

$$= \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\left(\frac{\partial \hat{p}_k}{\partial t}(t,x) - \frac{\partial p}{\partial t}(u_k,t,x) \right)^2 \|\hat{\phi}_k(t,x)\|_{\mathcal{G}_n}^2 \right] \quad (164)$$

$$\leq \frac{\kappa_\chi^2}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\left(\frac{\partial \hat{p}_k}{\partial t}(t,x) - \frac{\partial p}{\partial t}(u_k,t,x) \right)^2 \right] \quad (165)$$

$$\leq \kappa_\chi^2 \sup_k L(\hat{p}_k, p(u_k)). \quad (166)$$

Then, denoting $\varepsilon \triangleq \sup_k L(\hat{p}_k, p(u_k))$, from Proposition 7.17, we have, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$,

$$\|\hat{E}_1 - E_1\|_{\mathcal{G}_n} \leq \frac{16(\kappa_{\frac{\partial \hat{p}_k}{\partial t}} + \kappa_{\frac{\partial p}{\partial t}})\kappa_{\hat{\phi}} \log \frac{2}{\delta}}{N} + \sqrt{\frac{2\varepsilon\kappa_{\hat{\phi}}^2 \log \frac{2}{\delta}}{N}}. \quad (167)$$

2. Bound $\|\hat{E}_2 - E_2\|_{\mathcal{G}_n}$. To apply Proposition 7.17, we define

$$\psi(t,x,k) \triangleq \frac{\partial p}{\partial t}(u_k,t,x), \quad (168)$$

$$\chi(t,x,k) \triangleq \left(\hat{\phi}_k(t,x) - \tilde{\phi}(u_k,t,x) \right), \quad (169)$$

and

$$\kappa_\psi \triangleq \sup_{t,x,k} |\psi(t,x,k)| \leq \kappa_{\frac{\partial p}{\partial t}}, \quad (170)$$

$$\kappa_\chi \triangleq \sup_{t,x,k} \|\hat{\phi}_k(t,x) - \tilde{\phi}(u_k,t,x)\|_{\mathcal{G}_n} \leq \kappa_{\hat{\phi}} + \kappa_{\tilde{\phi}}, \quad (171)$$

and, from Lemma 7.10, we obtain

$$\sigma^2 \triangleq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} [\|\psi(t,x,k) \otimes \chi(t,x,k)\|_{\text{HS}}^2] \quad (172)$$

$$= \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\left(\frac{\partial p}{\partial t}(u_k,t,x) \right)^2 \|\hat{\phi}_k(t,x) - \tilde{\phi}(u_k,t,x)\|_{\mathcal{G}_n}^2 \right] \quad (173)$$

$$\leq \frac{\kappa_{\frac{\partial p}{\partial t}}^2}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\|\hat{\phi}_k(t,x) - \tilde{\phi}(u_k,t,x)\|_{\mathcal{G}_n}^2 \right] \quad (174)$$

$$\leq \kappa_{\frac{\partial p}{\partial t}}^2 \kappa_{tot} \sup_k L(\hat{p}_k, p(u_k)). \quad (175)$$

Then, from Proposition 7.17, we have, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$,

$$\|\hat{E}_2 - E_2\|_{\mathcal{G}_n} \leq \frac{16\kappa_{\frac{\partial p}{\partial t}}(\kappa_{\hat{\phi}} + \kappa_{\tilde{\phi}}) \log \frac{2}{\delta}}{N} + \sqrt{\frac{2\varepsilon\kappa_{\frac{\partial p}{\partial t}}^2 \kappa_{tot} \log \frac{2}{\delta}}{N}}. \quad (176)$$

3. Bound $\|E_1\|_{\mathcal{G}_n}$. We have

$$\|E_1\|_{\mathcal{G}_n} \leq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\left\| \left(\frac{\partial \hat{p}_k}{\partial t}(t, x) - \frac{\partial p}{\partial t}(u_k, t, x) \right) \hat{\phi}_k(t, x) \right\|_{\mathcal{G}_n} \right] \quad (177)$$

$$\leq \kappa_{\hat{\phi}} \varepsilon. \quad (178)$$

4. Bound $\|E_2\|_{\mathcal{G}_n}$. We have

$$\|E_2\|_{\mathcal{G}_n} \leq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\left\| \frac{\partial p}{\partial t}(u_k, t, x) \left(\hat{\phi}_k(t, x) - \tilde{\phi}(u_k, t, x) \right) \right\|_{\mathcal{G}_n} \right] \quad (179)$$

$$\leq \kappa_{\frac{\partial p}{\partial t}} \kappa_{tot} \varepsilon, \quad (180)$$

where $c > 0$ □

Conclusion. We conclude by combining all bounds that there exists a constant $c > 0$ that does not depend on N, K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$

$$\|\hat{D} - \tilde{D}\|_{\mathcal{G}_n} \leq c \log \frac{2}{\delta} (N^{-1} + \varepsilon). \quad (181)$$

Lemma 7.12 (Bound $\|\hat{C} - \tilde{C}\|_{\mathcal{G}_n}$). *For any $\delta \in (0, 1]$, with probability at least $1 - \delta$,*

$$\|\hat{C} - \tilde{C}\|_{\mathcal{G}_n} \leq c \log \frac{2}{\delta} (N^{-1} + \varepsilon). \quad (182)$$

Moreover, if $\lambda \geq 2c \log \frac{2}{\delta} (N^{-1} + \varepsilon)$,

$$\|\hat{C}_\lambda^{-1/2} \tilde{C}^{1/2}\|_\infty \leq \sqrt{2}, \quad (183)$$

where $c > 0$ is a constant that does not depend on N, K, δ .

Proof. Defining the operator $B_n = \tilde{C}_\lambda^{-1/2} (\tilde{C} - \hat{C}) \tilde{C}_\lambda^{-1/2}$, if $\|B_n\|_\infty < 1$, we have

$$\|\hat{C}_\lambda^{-1/2} \tilde{C}^{1/2}\|_\infty = \|\tilde{C}^{1/2} \hat{C}_\lambda^{-1} \tilde{C}^{1/2}\|_\infty^{1/2} \quad (184)$$

$$\leq \|\tilde{C}_\lambda^{1/2} \hat{C}_\lambda^{-1} \tilde{C}_\lambda^{1/2}\|_\infty^{1/2} \quad (185)$$

$$= \|(I - B_n)^{-1}\|_\infty^{1/2} \quad (186)$$

$$\leq (1 - \|B_n\|_\infty)^{-1/2}. \quad (187)$$

Moreover, we have

$$\hat{C} - \tilde{C} = \frac{1}{KN} \sum_{i=1}^N \sum_{k=1}^K \left(\hat{\phi}_k(t_i, x_i) \otimes \hat{\phi}_k(t_i, x_i) - \tilde{\phi}(u_k, t_i, x_i) \otimes \tilde{\phi}(u_k, t_i, x_i) \right). \quad (188)$$

Hence, defining

$$\hat{E}_1 \triangleq \frac{1}{KN} \sum_{i=1}^N \sum_{k=1}^K \left(\hat{\phi}_k(t_i, x_i) - \tilde{\phi}(u_k, t_i, x_i) \right) \hat{\phi}_k(t_i, x_i), \quad (189)$$

$$\hat{E}_2 \triangleq \frac{1}{KN} \sum_{i=1}^N \sum_{k=1}^K \tilde{\phi}(u_k, t_i, x_i) \left(\hat{\phi}_k(t_i, x_i) - \tilde{\phi}(u_k, t_i, x_i) \right), \quad (190)$$

$$E_1 \triangleq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\left(\hat{\phi}_k(t, x) - \tilde{\phi}(u_k, t, x) \right) \hat{\phi}_k(t, x) \right], \quad (191)$$

$$E_2 \triangleq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\tilde{\phi}(u_k, t, x) \left(\hat{\phi}_k(t, x) - \tilde{\phi}(u_k, t, x) \right) \right], \quad (192)$$

we have

$$\|\hat{C} - \tilde{C}\|_{\mathcal{G}_n} \leq \|\hat{E}_1 - E_1\|_{\mathcal{G}_n} + \|\hat{E}_2 - E_2\|_{\mathcal{G}_n} + \|E_1\|_{\mathcal{G}_n} + \|E_2\|_{\mathcal{G}_n}. \quad (193)$$

Then, following the same proof as for Lemma 7.11 leads, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$

$$\|B_n\|_{\mathcal{G}_n} \leq c \log \frac{2}{\delta} (N^{-1} + \varepsilon) \lambda^{-1}, \quad (194)$$

where $c > 0$ is a constant that does not depend on N, K, δ .

Hence, if $\lambda \geq 2c \log \frac{2}{\delta} (N^{-1} + \varepsilon)$, we have

$$\|\hat{C}_\lambda^{-1/2} \tilde{C}^{1/2}\|_\infty \leq \sqrt{2}. \quad (195)$$

□

Lemma 7.13 (Bound $\|\tilde{C}_\lambda^{-1/2} C^{1/2}\|_\infty$). *For any $\delta \in (0, 1]$, if $\lambda \geq \frac{9\kappa_\phi^2}{N} \log \frac{N}{\delta}$ and $\lambda \geq \frac{9\kappa_\phi^2}{K} \log \frac{K}{\delta}$, with probability at least $1 - \delta$,*

$$\|\tilde{C}_\lambda^{-1/2} C^{1/2}\|_\infty \leq 2. \quad (196)$$

Proof. We define

$$C_K = \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{t,x} \left[\tilde{\phi}(u_k, t, x) \otimes \tilde{\phi}(u_k, t, x) \right]. \quad (197)$$

We have

$$\|\tilde{C}_\lambda^{-1/2} C^{1/2}\|_\infty \leq \|\tilde{C}_\lambda^{-1/2} C_{K,\lambda}^{1/2}\|_\infty \|C_{K,\lambda}^{-1/2} C^{1/2}\|_\infty. \quad (198)$$

Therefore, applying Lemma 7.18 two times gives us, for any $\delta \in (0, 1]$, if $\lambda \geq \frac{9\kappa_\phi^2}{N} \log \frac{N}{\delta}$ and $\lambda \geq \frac{9\kappa_\phi^2}{K} \log \frac{K}{\delta}$, then with probability at least $1 - \delta$,

$$\|\tilde{C}_\lambda^{-1/2} C^{1/2}\|_\infty \leq 2, \quad (199)$$

with $\kappa_{\tilde{\phi}} \triangleq \sup_{(u,t,x) \in \mathcal{H} \times [0,T] \times \mathbb{R}^n} \|\tilde{\phi}(u, t, x)\|$. □

The following lemma provide bounds for all supremum norms involves in our proofs.

Lemma 7.14 (Bound all κ s). *Recall that, for any set \mathcal{X} , any Hilbert space \mathcal{H} , and any map $f : \mathcal{X} \rightarrow \mathcal{Y}$, we note*

$$\kappa_f = \sup_{x \in \mathcal{X}} \|f(x)\|_{\mathcal{H}}. \quad (200)$$

The following list of constants κ s can be bounded by a constant that does not depend on K .

$$\kappa_{\frac{\partial p}{\partial t}}, (\kappa_{\frac{\partial p_k}{\partial t}})_{k=1}^K, \kappa_{\tilde{\phi}}, (\kappa_{\hat{\phi}_k})_{k=1}^K, \kappa_\phi, (\kappa_{\phi_i})_{i=1}^n, (\kappa_{\phi_{ij}})_{i,j=1}^n. \quad (201)$$

Proof. Let us start with

$$\begin{aligned} \kappa_{\tilde{\phi}} \triangleq \sup_{u,t,x} \|\tilde{\phi}(u, t, x)\|_{\mathcal{G}_n} &\leq \sum_i^n \left(\kappa_\phi \kappa_{\frac{\partial p}{\partial x_i}} + \kappa_{\phi_i} \kappa_p \right) \\ &\quad + \frac{1}{2} \sum_{i,j=1}^n \left(\kappa_\phi \kappa_{\frac{\partial^2 p}{\partial x_i \partial x_j}} + \kappa_{\phi_i} \kappa_{\frac{\partial p}{\partial x_j}} + \kappa_{\phi_j} \kappa_{\frac{\partial p}{\partial x_i}} + \kappa_{\phi_{ij}} \kappa_p \right). \end{aligned} \quad (202)$$

Same inequality holds for $\hat{\phi}_k$ but with \hat{p}_k instead of p .

We conclude by noticing that the following kappas are bounded in our assumptions (independently of k):

$$\kappa_\phi, \kappa_{\phi_i}, \kappa_{\phi_{ij}}, \kappa_p, \kappa_{\frac{\partial p}{\partial x_j}}, \kappa_{\frac{\partial^2 p}{\partial x_i \partial x_j}}, \kappa_{\frac{\partial p}{\partial t}}, \kappa_{\hat{p}_k}, \kappa_{\frac{\partial \hat{p}_k}{\partial x_j}}, \kappa_{\frac{\partial^2 \hat{p}_k}{\partial x_i \partial x_j}}, \kappa_{\frac{\partial \hat{p}_k}{\partial t}}. \quad \square$$

7.11 Concentration inequalities

In this section, we provide the concentration inequalities used in the lemmas' proofs.

The following inequality is essentially a restatement of Proposition 2 of Rudi and Rosasco [53].

Proposition 7.15 (Bernstein's inequality for sum of random vectors). *Let \mathcal{Z} be a Polish space, $M : \mathcal{Z} \rightarrow \mathcal{H}$ be bounded maps with values in a separable Hilbert space \mathcal{H} , such that $\kappa_M \triangleq \sup_{z \in \mathcal{Z}} \|M(z)\|_{\mathcal{H}}$. Let z_1, \dots, z_N be a sequence of independent and identically distributed random vectors. Then, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$,*

$$\left\| \frac{1}{N} \sum_{i=1}^N M(z_i) - \mathbb{E}[M(z)] \right\|_{\text{HS}} \leq \frac{16\kappa_M \log \frac{2}{\delta}}{N} + \sqrt{\frac{2\sigma^2 \log \frac{2}{\delta}}{N}}, \quad (203)$$

with $\sigma^2 = \mathbb{E}[\|M(z)\|_{\text{HS}}^2]$.

Proof. For $p \in \mathbb{N}^*$, we have

$$\mathbb{E}[\|M_i(z) - \mathbb{E}[M_i(z)]\|_{\text{HS}}^p] \leq \mathbb{E}_{z, z'}[\|M_i(z) - M_i(z')\|_{\text{HS}}^p] \quad (204)$$

$$\leq \mathbb{E}_{z, z'}[(\|M_i(z)\|_{\text{HS}} + \|M_i(z')\|_{\text{HS}})^p] \quad (205)$$

$$\leq 2^p \mathbb{E}_z[\|M_i(z)\|_{\text{HS}}^p] \quad (206)$$

$$\leq 2^p \kappa_M^{p-2} \mathbb{E}[\|M(z)\|_{\mathcal{H}}^2], \quad (207)$$

using $(a+b)^p \leq 2^{p-1}(a^p + b^p)$ for $a, b > 0, p \in \mathbb{N}^*$.

One can conclude that, for $p \geq 2$,

$$\mathbb{E}[\|M_i(z) - \mathbb{E}[M_i(z)]\|_{\text{HS}}^p] \leq \frac{p!}{2} \sigma^2 M^{p-2}, \quad (208)$$

with $\sigma^2 = \mathbb{E}[\|\psi(z) \otimes \phi(z)\|_{\text{HS}}^2]$, and $M = 8\kappa_M$.

Then, application of Proposition 2 of Rudi and Rosasco [53] gives us

$$\left\| \frac{1}{N} \sum_{i=1}^N M(z_i) - \mathbb{E}[M(z)] \right\|_{\text{HS}} \leq \frac{2M \log \frac{2}{\delta}}{n} + \sqrt{\frac{2\sigma^2 \log \frac{2}{\delta}}{n}}, \quad (209)$$

with probability at least $1 - \delta$. □

From Proposition 7.15, we deduce the following inequality.

Proposition 7.16 (Covariance estimation, $\|\cdot\|_{\text{HS}}$). *Let \mathcal{Z} be a Polish space, $\psi, \phi : \mathcal{Z} \rightarrow \mathcal{H}$ be bounded maps with values in a separable Hilbert space \mathcal{H} , such that $\kappa_\psi \triangleq \sup_{z \in \mathcal{Z}} \|\psi(z)\|, \kappa_\phi \triangleq \sup_{z \in \mathcal{Z}} \|\phi(z)\|$. Let z_1, \dots, z_N be a sequence of independent and identically distributed random vectors. Then, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$,*

$$\left\| \frac{1}{N} \sum_{i=1}^N \psi(z_i) \otimes \phi(z_i) - \mathbb{E}[\psi(z) \otimes \phi(z)] \right\|_{\text{HS}} \leq \frac{2M \log \frac{2}{\delta}}{N} + \sqrt{\frac{2\sigma^2 \log \frac{2}{\delta}}{N}}, \quad (210)$$

with $\sigma^2 = \mathbb{E}[\|\psi(z) \otimes \phi(z)\|_{\text{HS}}^2]$, and $M = 8\kappa_\psi \kappa_\phi$.

Proof. For $i \in \llbracket 1, N \rrbracket$, we consider the random Hilbert-Schmidt operators

$$M_i(z) = \psi(z_i) \otimes \phi(z_i), \quad (211)$$

with mean

$$\mathbb{E}[M_i(z)] = \mathbb{E}[\psi(z) \otimes \phi(z)], \quad (212)$$

and such that

$$\|M_i(z)\|_{\text{HS}} = \|\psi(z_i) \otimes \phi(z_i)\|_{\text{HS}} = \|\psi(z_i)\|_{\mathcal{H}} \|\phi(z_i)\|_{\mathcal{H}} \leq \kappa_\psi \kappa_\phi, \quad (213)$$

and also

$$\mathbb{E}_z[\|M_i(z)\|_{\text{HS}}^2] = \mathbb{E}[\|\psi(z) \otimes \phi(z)\|_{\text{HS}}^2]. \quad (214)$$

Then, application of Proposition 7.15 gives us the desired bound. \square

The following proposition adapts Proposition 7.16 for our purposes.

Proposition 7.17 (Mixed covariance estimation, $\|\cdot\|_{\text{HS}}$). *Let \mathcal{X}, \mathcal{Y} be Polish spaces, $\psi, \phi : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{H}$ be bounded maps with values in a separable Hilbert space \mathcal{H} , such that $\kappa_\psi \triangleq \sup_{(x,y) \in \mathcal{X} \times \mathcal{Y}} \|\psi(x,y)\|$, $\kappa_\phi \triangleq \sup_{(x,y) \in \mathcal{X} \times \mathcal{Y}} \|\phi(x,y)\|$. Let x_1, \dots, x_N be a sequence of independent and identically distributed random vectors. Then, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$,*

$$\left\| \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{y|x_i}[\psi(x_i, y) \otimes \phi(x_i, y)] - \mathbb{E}_{x,y}[\psi(x, y) \otimes \phi(x, y)] \right\|_{\text{HS}} \leq \frac{2M \log \frac{2}{\delta}}{N} + \sqrt{\frac{2\sigma^2 \log \frac{2}{\delta}}{N}}, \quad (215)$$

with $\sigma^2 = \mathbb{E}_{x,y}[\|\psi(x, y) \otimes \phi(x, y)\|_{\text{HS}}^2]$, and $M = 8\kappa_\psi \kappa_\phi$.

Proof. Same proof as Proposition 7.16 but defining $M_i = \mathbb{E}_y[\psi(x, y) \otimes \phi(x, y)]$ instead of $M_i = \psi(x) \otimes \phi(x)$. \square

The following proposition adapts Lemma 3.6 in Rudi et al. [54] to handle "mixed" covariance estimation.

Proposition 7.18 (Mixed covariance estimation, $\|\cdot\|_\infty$). *Let \mathcal{X}, \mathcal{Y} be Polish spaces, $\phi : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{H}$ a bounded maps with values in a separable Hilbert space \mathcal{H} such that $\kappa_\phi \triangleq \sup_{(x,y) \in \mathcal{X} \times \mathcal{Y}} \|\phi(x, y)\|$. Let x_1, \dots, x_N be a sequence of independent and identically distributed random vectors. Let*

$$C_N = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{y|x_i}[\phi(x_i, y) \otimes \phi(x_i, y)], \quad (216)$$

$$C = \mathbb{E}_{x,y}[\phi(x, y) \otimes \phi(x, y)], \quad (217)$$

Then, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$,

$$\|C_\lambda^{-1/2}(C - C_N)C_\lambda^{-1/2}\|_\infty \leq \frac{2\beta\kappa_\phi^2}{\lambda N} + \sqrt{\frac{2\kappa_\phi^2\beta}{3\lambda N}}, \quad (218)$$

with $\beta = \log \frac{4(\|C\|_\infty + \lambda) \text{Tr}(C)}{\delta \lambda \|C\|_\infty}$.

In particular, if $\lambda \geq \frac{18\kappa_\phi^2}{N} \log \frac{N}{\delta}$,

$$\|C_{N,\lambda}^{-1/2}C_\lambda^{1/2}\|_\infty \leq \sqrt{2}, \quad (219)$$

and

$$\|C_\lambda^{-1/2}C_{N,\lambda}^{1/2}\|_\infty \leq \sqrt{2}. \quad (220)$$

Proof. Same proof than Lemma 3.6 of Rudi et al. [54], but defining

$$Z_i \triangleq C_\lambda^{-1/2} \mathbb{E}_{y|x_i} [\phi(x_i, y) \otimes \phi(x_i, y)] C_\lambda^{-1/2}, \quad (221)$$

instead of $Z_i \triangleq C_\lambda^{-1/2} \phi(x_i) \otimes C_\lambda^{-1/2} \phi(x_i)$.

More precisely, defining $U_i \triangleq C_\lambda^{-1/2} \phi(x_i, y)$, we have $Z_i = \mathbb{E}_{y|x_i} [U_i \otimes U_i]$ instead of $Z_i = U_i \otimes U_i$.

Moreover, notice that

$$\|Z\|_\infty \leq \lambda^{-1} \|\mathbb{E}_{y|x_i} [\phi(x_i, y) \otimes \phi(x_i, y)]\|_\infty \quad (222)$$

$$\leq \frac{\kappa_\phi^2}{\lambda}, \quad (223)$$

and

$$T \triangleq \mathbb{E}_{x_i} [Z] = CC_\lambda^{-1}, \quad (224)$$

and also

$$\mathbb{E}[(Z - T)^2] = \mathbb{E}[Z^2] - T^2 \preceq \mathbb{E}[Z^2] \preceq \mathbb{E}_{x,y} [\|U\|_{\mathcal{H}}^2 U \otimes U] \preceq \kappa_\phi^2 \lambda^{-1} T, \quad (225)$$

using the inequality $\mathbb{E}(M^2) \preceq \mathbb{E}(M^2)$ for a random variable M with values in the space of bounded self-adjoint operators.

□

7.12 Proof of Theorem 4.2

The following theorem establishes refined learning rates for the proposed estimator, showing that for regular problems, the rates of FP matching can improve at arbitrarily fast polynomial rates. However, this gain is constrained by the accuracy of the \hat{p}_k . In particular, the error remains greater than the L^2 error of any \hat{p}_k .

Theorem 7.19 (Refined L^2 learning rates). *Under the assumptions of Theorem 4.1, Assumption (A5), and Assumption (A6), then there exist constants $c_1, c_2 > 0$ that do not depend on N, K, δ , such that for any $\delta \in (0, 1]$, defining $\varepsilon_\infty \triangleq \sup_{k \in \llbracket 1, K \rrbracket} L^\infty(\hat{p}_k, p(u_k))$ where L^∞ is a squared Sobolev norm defined as*

$$L^\infty(p_1, p_2) \triangleq \int_{t=0}^T \left\| \frac{\partial p_1}{\partial t}(t, \cdot) - \frac{\partial p_2}{\partial t}(t, \cdot) \right\|_{L^\infty(\mathbb{R}^n)}^2 + \|p_1(t, \cdot) - p_2(t, \cdot)\|_{W_\infty^2(\mathbb{R}^n)}^2 dt, \quad (226)$$

taking $\lambda = c_2 \left(\left(\frac{18c}{N} \log^2 \frac{N}{\delta} \right)^{\frac{1}{1-r}} + \left(\frac{18c}{K} \log^2 \frac{K}{\delta} \right)^{\frac{1}{1-s}} \right)$, if

$$\varepsilon + N^{-1} \varepsilon_\infty \leq \left(\frac{\log \frac{N}{\delta}}{\sqrt{N}} \right)^{\frac{2+\alpha}{1-r}} + \left(\frac{\log \frac{K}{\delta}}{\sqrt{K}} \right)^{\frac{2+\alpha}{1-s}}, \quad (227)$$

then with probability $1 - \delta$,

$$\|p_{\hat{p}, \sigma} - p\|_{L^2(p_c) L^2([0, T] \times \mathbb{R}^n)} \leq c_1 \log \frac{2}{\delta} \left(\frac{\log \frac{K \wedge N}{\delta}}{\sqrt{K} \wedge \sqrt{N}} \right)^{\frac{1+\alpha}{1-r \wedge s}}. \quad (228)$$

Proof. The definitions of Section 7.4 allow for the following decomposition

$$\|(\hat{w} - w)C^{1/2}\|_{\mathcal{G}_n} \leq \|(\hat{w} - \tilde{w})C^{1/2}\|_{\mathcal{G}_n} + \|(\tilde{w} - w_K)C^{1/2}\|_{\mathcal{G}_n} + \|(w_K - w)C^{1/2}\|_{\mathcal{G}_n}. \quad (229)$$

Using Lemma 7.6, refined Lemmas 7.21, 7.22, and 7.23, and Lemma 7.14, we determine that for any $\delta \in (0, 1]$, if $\lambda \geq c_2 \log \frac{2}{\delta} (N^{-1} \varepsilon_\infty + \varepsilon)$, $\lambda^{1-r} \geq \frac{18c}{N} \log \frac{N}{\delta}$, and $\lambda^{1-s} \geq \frac{18c}{K} \log \frac{K}{\delta}$, with probability $1 - \delta$,

$$\text{FP}(\hat{b}, \hat{\sigma}^2) \leq c_1 \left((\varepsilon_\infty N^{-1} + \varepsilon) \lambda^{-1/2} \log \frac{2}{\delta} + \frac{\log \frac{N}{\delta}}{\sqrt{N}} \lambda^{(\alpha+r)/2} + \frac{\log \frac{K}{\delta}}{\sqrt{K}} \lambda^{(\alpha+s)/2} + \lambda^{(1+\alpha)/2} \right), \quad (230)$$

where constants $c_1, c_2 > 0$ are independent of N, K , and δ . Therefore, setting $\lambda = \left(\frac{18c}{N} \log^2 \frac{N}{\delta} \right)^{\frac{1}{1-r}} + \left(\frac{18c}{K} \log^2 \frac{K}{\delta} \right)^{\frac{1}{1-s}}$, we find

$$\text{FP}(\hat{b}, \hat{\sigma}^2) \leq c_1 \log \frac{2}{\delta} \left(\frac{\varepsilon + N^{-1} \varepsilon_\infty}{\lambda^{1/2}} + \left(\frac{\log \frac{N}{\delta}}{\sqrt{N}} \right)^{\frac{1+\alpha}{1-r}} + \left(\frac{\log \frac{K}{\delta}}{\sqrt{K}} \right)^{\frac{1+\alpha}{1-s}} \right) \quad (231)$$

$$+ \left(\frac{\log \frac{N}{\delta}}{\sqrt{N}} \right) \left(\frac{\log \frac{K}{\delta}}{\sqrt{K}} \right)^{\frac{\alpha+r}{1-s}} + \left(\frac{\log \frac{K}{\delta}}{\sqrt{K}} \right) \left(\frac{\log \frac{N}{\delta}}{\sqrt{N}} \right)^{\frac{\alpha+s}{1-r}}. \quad (232)$$

up to overloading $c_1 > 0$. Assuming $\varepsilon + N^{-1} \varepsilon_\infty \leq \lambda^{1+\alpha/2}$, the condition $\lambda \geq c_2 \log \frac{2}{\delta} (N^{-1} \varepsilon_\infty + \varepsilon)$ holds true up to multiplying λ by a constant, and overloading the constants c_1 .

Moreover, notice that defining $\gamma = r \wedge s$, we have $\frac{\alpha+r}{1-s} \geq \frac{\alpha+\gamma}{1-\gamma}$, $\frac{\alpha+s}{1-r} \geq \frac{\alpha+\gamma}{1-\gamma}$, $\frac{\alpha+1}{1-r} \geq \frac{\alpha+1}{1-\gamma}$, and $\frac{\alpha+1}{1-s} \geq \frac{\alpha+1}{1-\gamma}$.

Then, defining $P = K \wedge N$, we have

$$\left(\frac{\log \frac{N}{\delta}}{\sqrt{N}} \right) \left(\frac{\log \frac{K}{\delta}}{\sqrt{K}} \right)^{\frac{\alpha+r}{1-s}} \leq \left(\frac{\log \frac{P}{\delta}}{\sqrt{M}} \right)^{\frac{1+\alpha}{1-\gamma}}, \quad (233)$$

$$\left(\frac{\log \frac{K}{\delta}}{\sqrt{K}} \right) \left(\frac{\log \frac{N}{\delta}}{\sqrt{N}} \right)^{\frac{\alpha+s}{1-r}} \leq \left(\frac{\log \frac{P}{\delta}}{\sqrt{M}} \right)^{\frac{1+\alpha}{1-\gamma}}. \quad (234)$$

Finally, Lemma 3.2 allows to conclude the proof. \square

7.13 Proof of refined lemmas

In this section, we present refined versions of the lemmas previously utilized to establish the non-refined L^2 learning rate. To maintain clarity in our presentation, we do not restate the assumptions for each lemma if they are unchanged from their original, non-refined versions. The main ingredient for these proofs involves utilizing Assumption (A5) and Assumption (A6), and performing derivations such as:

$$C_\lambda^{-1} C^a \preceq \lambda^{-(1-a)} I \quad (235)$$

to establish refined dependencies in λ . Furthermore, to avoid redundancy and streamline the presentation, we directly state the results without detailing the proofs when possible, focusing on clarifying the refinements made from their non-refined counterparts.

Proposition 7.20 (Refined Lemma 7.18). *Let*

$$\forall x \in \mathcal{X}, C(x) = \mathbb{E}_{y|x}[\phi(x, y) \otimes \phi(x, y)]. \quad (236)$$

Assume that there exist constants $r \geq 0$ and $c > 0$ such that

$$\|C^{-r/2} C(x)^{1/2}\|_\infty < c \quad \text{almost surely.} \quad (237)$$

Then, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$:

$$\|C_\lambda^{-1/2} (C - C_N) C_\lambda^{-1/2}\|_\infty \leq \frac{2\beta c}{3\lambda^{1-r} N} + \sqrt{\frac{2c\beta}{\lambda^{1-r} N}}, \quad (238)$$

where $\beta = \log \frac{4(\|C\|_\infty + \lambda) \text{Tr}(C)}{\delta \lambda \|C\|_\infty}$.

In particular, if $\lambda \geq \left(\frac{18c}{N} \log \frac{N}{\delta}\right)^{\frac{1}{1-r}}$,

$$\|C_{N,\lambda}^{-1/2} C_\lambda^{1/2}\|_\infty \leq \sqrt{2}, \quad (239)$$

and

$$\|C_\lambda^{-1/2} C_{N,\lambda}^{1/2}\|_\infty \leq \sqrt{2}. \quad (240)$$

Proof. We have

$$\|Z\|_\infty = \|C_\lambda^{-1/2} C(x)^{1/2}\|_\infty \quad (241)$$

$$\leq \|C_\lambda^{-1/2} C_\lambda^{r/2}\|_\infty^2 \|C_\lambda^{-r/2} C(x)^{1/2}\|_\infty^2 \quad (242)$$

$$\leq c\lambda^{-(1-r)}, \quad (243)$$

and

$$\mathbb{E}[(Z - T)^2] = \mathbb{E}[Z^2] - T^2 \preceq \mathbb{E}[Z^2] \preceq \mathbb{E}_{x,y}[\|U\|_{\mathcal{H}}^2 U \otimes U] \preceq c\lambda^{-(1-r)}T. \quad (244)$$

□

Lemma 7.21 (Refined Lemma 7.7). *There exist constants $c_1, c_2 > 0$ that do not depend on N, K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$*

$$\|(\hat{w} - \tilde{w})C^{1/2}\|_{\mathcal{G}_n} \leq c_1 \log \frac{2}{\delta} (N^{-1}\varepsilon_\infty + \varepsilon)\lambda^{-1/2}, \quad (245)$$

provided that $\lambda \geq c_2 \log \frac{2}{\delta} (N^{-1}\varepsilon_\infty + \varepsilon)$, $\lambda \geq \frac{18c}{N} \log \frac{N}{\delta}$ and $\lambda \geq \frac{18c}{K} \log \frac{K}{\delta}$.

Lemma 7.22 (Refined Lemma 7.8). *There exist a constant $c > 0$ that does not depend on N, K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$*

$$\|(\tilde{w} - w_K)C^{1/2}\|_{\mathcal{G}_n} \leq c \frac{\log \frac{N}{\delta}}{\sqrt{N}} \lambda^{(\alpha+r)/2}, \quad (246)$$

provided that $\lambda^{1-r} \geq \frac{18c}{N} \log \frac{N}{\delta}$ and $\lambda^{1-s} \geq \frac{18c}{K} \log \frac{K}{\delta}$.

Lemma 7.23 (Refined Lemma 7.9). *There exist a constant $c > 0$ that does not depend on K, δ , such that, for any $\delta \in (0, 1]$, with probability at least $1 - \delta$*

$$\|(w_K - w)C^{1/2}\|_{\mathcal{G}_n} \leq c \left(\frac{\log \frac{K}{\delta}}{\sqrt{K}} \lambda^{(\alpha+s)/2} + \lambda^{(1+\alpha)/2} \right), \quad (247)$$

if $\lambda^{1-s} \geq \frac{18c}{K} \log \frac{K}{\delta}$.

Lemma 7.24 (Refined Lemma 7.11). *For any $\delta \in (0, 1]$, with probability at least $1 - \delta$*

$$\|\hat{D} - \tilde{D}\|_{\mathcal{G}_n} \leq c \log \frac{2}{\delta} (\varepsilon_\infty N^{-1} + \varepsilon), \quad (248)$$

where $c > 0$ is a constant that does not depend on N, K, δ .

Lemma 7.25 (Refined Lemma 7.12). *For any $\delta \in (0, 1]$, with probability at least $1 - \delta$,*

$$\|\hat{C} - \tilde{C}\|_{\mathcal{G}_n} \leq c \log \frac{2}{\delta} (\varepsilon_\infty N^{-1} + \varepsilon). \quad (249)$$

Moreover, if $\lambda \geq 2c \log \frac{2}{\delta} (\varepsilon_\infty N^{-1} + \varepsilon N^{-1/2} + \varepsilon)$,

$$\|\hat{C}_\lambda^{-1/2} \tilde{C}^{1/2}\|_\infty \leq \sqrt{2}, \quad (250)$$

where $c > 0$ is a constant that does not depend on N, K, δ .

Lemma 7.26 (Refined Lemma 7.13). *For any $\delta \in (0, 1]$, if $\lambda^{1-r} \geq \frac{18c}{N} \log \frac{N}{\delta}$ and $\lambda^{1-s} \geq \frac{18c}{K} \log \frac{K}{\delta}$, with probability at least $1 - \delta$,*

$$\|\tilde{C}_\lambda^{-1/2} C_\lambda^{1/2}\|_\infty \leq 2. \quad (251)$$

7.14 Proof of L^∞ learning rates

Lemma 7.27. *Let $\gamma \triangleq r \wedge s$. Under the same assumptions as Theorem 4.2, for any $\alpha \in [0, 1]$, with probability at least $1 - \delta$, we have*

$$\|(\hat{w} - w)C^{\gamma/2}\|_{\mathcal{G}_n} \leq c \log \frac{2}{\delta} \left(\frac{\log \frac{K \wedge N}{\delta}}{\sqrt{K \wedge N}} \right)^{\frac{\alpha+\gamma}{1-\gamma}}, \quad (252)$$

where $c > 0$ is independent of N, K, δ .

Proof. We have

$$\|(\hat{w} - w)C^{\gamma/2}\|_{\mathcal{G}_n} \leq \|(\hat{w} - w)C_\lambda^{1/2}\|_{\mathcal{G}_n} \|C_\lambda^{1/2}C^{\alpha/2}\|_{\mathcal{G}_n} \quad (253)$$

$$\leq \left(\|(\hat{w} - w)C^{1/2}\|_{\mathcal{G}_n} + \lambda^{1/2} \|\hat{w} - w\|_{\mathcal{G}_n} \right) \lambda^{(\gamma-1)/2}. \quad (254)$$

Theorem 4.2 provides the bound

$$\|(\hat{w} - w)C^{1/2}\|_{\mathcal{G}_n} \leq \mathcal{B}(N, K, \delta) \triangleq c_1 \log \frac{2}{\delta} \left(\frac{\log \frac{K \wedge N}{\delta}}{\sqrt{K \wedge N}} \right)^{\frac{1+\alpha}{1-\gamma}}. \quad (255)$$

taking $\lambda = c_2 \left(\left(\frac{18c}{N} \log^2 \frac{N}{\delta} \right)^{\frac{1}{1-r}} + \left(\frac{18c}{K} \log^2 \frac{K}{\delta} \right)^{\frac{1}{1-s}} \right)$, where constants $c_1, c_2 > 0$ are independent of N, K , and δ .

The same proof as Theorem 4.2 leads, under the exact same assumptions, to the bound

$$\|\hat{w} - w\|_{\mathcal{G}_n} \leq \mathcal{B}(N, K, \delta) \times \lambda^{-1/2}. \quad (256)$$

Therefore,

$$\|(\hat{w} - w)C^{\alpha/2}\|_{\mathcal{G}_n} \leq 2\mathcal{B}(N, K, \delta) \lambda^{-(1-\gamma)/2} \quad (257)$$

$$\leq c \log \frac{2}{\delta} \left(\frac{\log \frac{K \wedge N}{\delta}}{\sqrt{K \wedge N}} \right)^{\frac{\gamma+\alpha}{1-\gamma}}, \quad (258)$$

for a constant $c > 0$.

□

Corollary 7.1 (L^∞ learning rates). *Under identical conditions to those in Theorem 4.2, then for any $\delta \in (0, 1]$, with probability at least $1 - \delta$*

$$\|p_{\hat{b}, \hat{\sigma}} - p\|_{L^\infty(p_c)L^2([0, T] \times \mathbb{R}^n)} \leq c \log \frac{2}{\delta} \left(\frac{\log \frac{K \wedge N}{\delta}}{\sqrt{K \wedge N}} \right)^{\frac{\alpha + r \wedge s}{1 - r \wedge s}}, \quad (259)$$

where $c > 0$ is a constant independent of N , K , or δ .

Proof. We have, defining $\gamma = r \wedge s$, almost everywhere,

$$\mathbb{E}_{t,x} \left[\left(\frac{\partial p}{\partial t}(u, t, x) - \left(\mathcal{L}^{(\hat{b}, \hat{\sigma})}(u) \right)^* p(u, t, x) \right)^2 \right]^{1/2} = c \|(\hat{w} - w)C(u)^{1/2}\|_{\mathcal{G}_n} \quad (260)$$

$$\leq c \|(\hat{w} - w)C^{\gamma/2}\|_{\mathcal{G}_n} \|C_\lambda^{-\gamma/2} C^{1/2}(u)\|_\infty \quad (261)$$

$$\leq cc_0 \|(\hat{w} - w)C^{\gamma/2}\|_{\mathcal{G}_n}. \quad (262)$$

Therefore, there exists $c > 0$, such that

$$\|p_{\hat{b}, \hat{\sigma}} - p\|_{L^\infty(p_c)L^\infty([0, T])L^2(\mathbb{R}^n)} \leq c \|(\hat{w} - w)C^{\gamma/2}\|_{\mathcal{G}_n}. \quad (263)$$

Then, applying Lemma 7.27 allows to conclude the proof. \square

7.15 Proof of CVaR learning rates

Value at Risk (VaR) For any $\alpha \in [0, 1]$, the value at risk $VaR_\alpha(X)$ of a random variable $X : \Omega \rightarrow \mathbb{R}$ is defined as the smallest of the $\alpha \times 100\%$ worst (greatest) values of the distribution of X , namely

$$VaR_\alpha(X) \triangleq \inf_{t \in \mathbb{R}} \{t : \mathbb{P}(X \leq t) \geq 1 - \alpha\}. \quad (264)$$

Conditional Value at Risk (CVaR). For any $\alpha \in [0, 1]$, the conditional value at risk of a variable $X : \Omega \rightarrow \mathbb{R}$ is defined as

$$CVaR_\alpha(X) \triangleq \inf_{t \in \mathbb{R}} \{t + \alpha^{-1} \mathbb{E}[\max(0, X - t)]\}. \quad (265)$$

When the cumulative distribution function of X is continuous at $VaR_\alpha(X)$, it holds that

$$CVaR_\alpha(X) = \mathbb{E}[X | X \geq VaR_\alpha(X)], \quad (266)$$

allowing to interpretate CVaR as the expected value conditional upon being within some percentage of the worst-case loss scenarios.

Lemma 7.28. *For any $f \in L^\infty(\mathbb{R}^n)$, $\alpha \in [0, 1]$, and two random variables $X_1, X_2 : \Omega \rightarrow \mathbb{R}^n$ with probability density functions p_1, p_2 , the following bound holds*

$$|CVaR_\alpha(f(X_1)) - CVaR_\alpha(f(X_2))| \leq \frac{\|f\|_\infty}{\alpha} \|p_1 - p_2\|_{L^1}. \quad (267)$$

Proof. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $\alpha \in [0, 1]$, we note $C_1, C_2 = CVaR_\alpha(f(X_1)), CVaR_\alpha(f(X_2))$, $V_1, V_2 = VaR_\alpha(f(X_1)), VaR_\alpha(f(X_2))$.

If $\max(V_1, V_2) = V_2$, for any $f : \mathbb{R}^n \rightarrow \mathbb{R}$, we have

$$\begin{aligned} |C_1 - C_2| &= \frac{1}{\alpha} \left| \int_{f(x) \geq V_1} f(x)p_1(x)dx - \int_{f(x) \geq V_2} f(x)p_2(x)dx \right| \\ &\leq \frac{1}{\alpha} \left| \int_{V_1 \leq f(x) \leq V_2} f(x)p_1(x)dx \right| + \frac{1}{\alpha} \left| \int_{f(x) \geq V_2} f(x)(p_1(x) - p_2(x))dx \right|. \end{aligned}$$

However,

$$\left| \int_{f(x) \geq V_2} f(x)(p_1(x) - p_2(x)) dx \right| \leq \int_{f(x) \geq V_2} |f(x)| |p_1(x) - p_2(x)| dx.$$

Moreover,

$$\begin{aligned} \left| \int_{V_1 \leq f(x) \leq V_2} f(x)p_1(x) dx \right| &\leq \|f\|_\infty \int_{V_1 \leq f(x) \leq V_2} p_1(x) dx \\ &= \|f\|_\infty \left| \int_{f(x) \leq V_2} p_1(x) dx - \int_{f(x) \leq V_1} p_1(x) dx \right| \\ &= \|f\|_\infty \left| \int_{f(x) \leq V_2} (p_1(x) - p_2(x)) dx + \int_{f(x) \leq V_2} p_2(x) dx - \int_{f(x) \leq V_1} p_1(x) dx \right| \\ &= \|f\|_\infty \int_{f(x) \leq V_2} |p_1(x) - p_2(x)| dx, \end{aligned}$$

as $\int_{f(x) \leq V_1} p_1(x) dx = \mathbb{P}(f(X_1) \leq V_1) = 1 - \alpha$ by definition of V_1 , and similarly $\int_{f(x) \leq V_2} p_2(x) dx = 1 - \alpha$.

Finally, we obtain $|C_1 - C_2| \leq \alpha^{-1} \|f\|_\infty \|p_1 - p_2\|_{L^1}$. If $\max(V_1, V_2) = V_1$, by symmetry, the same bound holds.

□

Lemma 7.29. *For any $f \in L^2(\mathbb{R}^n)$, if there exists $\beta > 1$ such that $\|fx^\beta\|_{L^1(\mathbb{R}^n)} < +\infty$, then we have*

$$\|f\|_{L^1(\mathbb{R}^n)} \leq \|f\|_{L^2(\mathbb{R}^n)}^{\frac{\beta}{\beta+n/2}} (3 + \|fx^\beta\|_{L^1(\mathbb{R}^n)}). \quad (268)$$

Proof. For $f \in L^2(\mathbb{R}^n)$, and $R > 0$, denoting $\mathcal{B} = B_R^{\mathbb{R}^n}(0)$, we have

$$\|f\|_{L^1(\mathbb{R}^n)} = \|f\|_{L^1(\mathcal{B})} + \|f\|_{L^1(\mathbb{R}^n \setminus \mathcal{B})}. \quad (269)$$

Moreover,

$$\|f\|_{L^1(\mathcal{B})} \leq \|f\|_{L^2(\mathcal{B})} 3R^{n/2} \quad (270)$$

$$\leq \|f\|_{L^2(\mathbb{R}^n)} 3R^{n/2}. \quad (271)$$

Furthermore, similar proof than for the standard Markov's inequality gives

$$\|f\|_{L^1(\mathbb{R}^n \setminus \mathcal{B})} \leq R^{-\beta} \|fx^\beta\|_{L^1(\mathbb{R}^n)}. \quad (272)$$

Hence,

$$\|f\|_{L^1(\mathbb{R}^n)} \leq \|f\|_{L^2(\mathbb{R}^n)} 3R^{n/2} + R^{-\beta} \|fx^\beta\|_{L^1(\mathbb{R}^n)}. \quad (273)$$

Now, taking $R = \|f\|_{L^2(\mathbb{R}^n)}^{-\frac{n/2}{\beta+n/2}}$, gives

$$\|f\|_{L^1(\mathbb{R}^n)} \leq \|f\|_{L^2(\mathbb{R}^n)}^{\frac{\beta}{\beta+n/2}} (3 + \|fx^\beta\|_{L^1(\mathbb{R}^n)}). \quad (274)$$

□

Lemma 7.30 (Bounded moments of SDE solutions). *Denote $X_{\hat{b}, \hat{\sigma}}(u)$ the solution to the SDE driven by $(\hat{b}, \hat{\sigma}^2) \in \mathcal{F}$ under control $u \in W^{4+\lfloor \frac{s}{2} \rfloor}([0, T], \mathbb{R}^d)$. Under Assumption (A1) and (A3), if $\mathbb{E}[\|X_0\|^\beta] < \infty$ for some $\beta > 2$, then there exists a constant $C > 0$ such that, for all $t \in [0, T]$,*

$$\mathbb{E}[\|X_{\hat{b}, \hat{\sigma}}(u)(t)\|^\beta] \leq C.$$

Proof. Applying Hölder's and Burkholder-Davis-Gundy inequalities, we can find a constant $C > 0$ such that, for all $t \in [0, T]$,

$$\begin{aligned} \mathbb{E}[\|X(t)\|^\beta] &\leq C \left(\mathbb{E}[\|X_0\|^\beta] + \mathbb{E} \left[\left\| \int_0^t b(s, X(s), u(s)) ds \right\|^\beta \right] + \mathbb{E} \left[\left\| \int_0^t \sigma(s, X(s), u(s)) dW_s \right\|^\beta \right] \right) \\ &\leq C \left(\mathbb{E}[\|X_0\|^\beta] + \mathbb{E} \left[\int_0^t \|b(s, X(s), u(s))\|^\beta ds \right] + \mathbb{E} \left[\left(\int_0^t \|\sigma(s, X(s), u(s))\|^2 ds \right)^{\frac{\beta}{2}} \right] \right). \end{aligned}$$

Using classical embedding results (see, for instance, Theorem 2.1 in Bonalli and Rudi [11]), results on composition for Sobolev spaces (see Theorem 9.1 in Bourdaud [12]), along with Assumptions (A1) and (A3), and the condition $u \in W^{4+\lfloor \frac{s}{2} \rfloor}([0, T], \mathbb{R}^d)$, we conclude that the drift $b(\cdot, \cdot, u(\cdot))$ and diffusion $\sigma(\cdot, \cdot, u(\cdot))$ are uniformly bounded. It follows that there exists a constant $C > 0$ such that

$$\mathbb{E}[\|X_{\hat{b}, \hat{\sigma}}(u)(t)\|^\beta] \leq C.$$

□

7.16 Proof of Lemma 4.7

We recall the following result.

Sobolev Spaces and regularity assumptions. Consider a least-squares regression problem

$$\min_{f \in \mathcal{F}} \mathbb{E}_{\rho_z} [(\hat{f}(x) - f^*(x))^2], \quad (275)$$

where f^* is s -times differentiable, and there exist constants $c_1, c_2 > 0$ such that $c_1 d\nu \leq d\rho_z \leq c_2 d\nu$ holds ρ_z -almost surely, with ν being the Lebesgue measure. Let $q > \frac{m}{2}$, and let $\mathcal{Z} \subset \mathbb{R}^m$ be a bounded subset of \mathbb{R}^m with Lipschitz boundary. The Matérn kernel

$$k_{q,m}(x, y) = \|x - y\|^{q - \frac{m}{2}} \mathcal{K}_{\frac{m}{2} - q}(\|x - y\|), \quad (276)$$

where $\mathcal{K}_{\frac{m}{2} - q}$ is the modified Bessel function of the second kind, induces the RKHS $\mathcal{F} = W^q(\mathcal{Z})$, where $W^q(\mathcal{Z})$ denotes the Sobolev space of smoothness q . The source condition is verified for $\alpha = \frac{s}{q} - 1$, and the embedding property is verified for any parameter $r > 1 - \frac{m}{2q}$. For more details, we refer the reader to Example 2 in Pillaud-Vivien et al. [50], Section 4 in Fischer and Steinwart [22], and Steinwart et al. [63], Wendland [69].

Regularity of Fokker-Planck matching. The previous result on Sobolev RKHS is not directly applicable to our setting. The chosen hypothesis spaces \mathcal{F} for the coefficients and \mathcal{H} for the controls induce the hypothesis space \mathcal{M} for the Fokker-Planck matching. More precisely,

$$\mathcal{M} \triangleq \{u, t, x \mapsto ((\mathcal{L}^{\hat{b}, \hat{\sigma}})^* p)(u, t, x) : \mathcal{H} \times [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R} \mid (\hat{b}, \hat{\sigma}^2) \in \mathcal{F}\}. \quad (277)$$

Nevertheless, the following lemma establishes the embedding property for Fokker-Planck matching when the coefficients belong to a Sobolev space, and the control spaces is smoothly parametrized.

Lemma 7.31. *Let $q > \frac{1+n+m}{2} + 4$, and $\mathcal{H} = \{u_\theta | \theta \in \mathbb{R}^m\}$ with $\theta, t \mapsto u_\theta(t) \in W^q(\mathbb{R}^m \times [0, T])$. Assume σ^2 is uniformly elliptic, with $\forall (t, x) \in [0, T] \times \mathbb{R}^n \sigma^2(t, x) \geq \kappa I_{\mathbb{R}^n \times n}$ with $\kappa > 0$, and $(b, \sigma^2 - \kappa I_{\mathbb{R}^n \times n}) \in W^q([0, T] \times \mathbb{R}^n \times \mathbb{R}^d, \mathbb{R}^{n+n^2})$. Assume further that $\theta, t, x \mapsto p(u_\theta, t, x) \in W^{q-2}(\mathbb{R}^m \times [0, T] \times \mathbb{R}^n, \mathbb{R})$. Additionally, let p_s, p_c be such that there exist constants $c_1, c_2, c_3, c_4 > 0$ ensuring $c_1 \leq p_s(t, x) \leq c_2$ holds p_s -almost surely and $c_3 \leq p_c(u) \leq c_4$ holds p_c -almost surely, and assume that the partial derivatives of $(b, \sigma^2 - \kappa I_{\mathbb{R}^n \times n})$ are supported within the bounded set $D \triangleq \text{supp}(p_s)$. Then, using a kernel that induces the Sobolev RKHS $W^q(D \times \mathbb{R}^d, \mathbb{R}^{n+n^2})$, the Assumptions (A1), (A2), (A3), (A4) are met, Assumption (A5) is satisfied with $\alpha = 0$, Assumptions (A6) are satisfied with $r = 1 - \frac{m}{2(q-2)}$ and $s = 1 - \frac{n+1}{2(q-2)}$.*

Proof. The induced space for Fokker-Planck matching

$$\mathcal{M} \triangleq \{\theta, t, x \mapsto ((\mathcal{L}^{\hat{b}, \hat{\sigma}})^* p)(u_\theta, t, x) : \mathbb{R}^m \times [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R} \mid (\hat{b}, \hat{\sigma}^2) \in \mathcal{F}\}, \quad (278)$$

is a subset of the RKHS with kernel

$$\tilde{k}((\theta, t, x), (\theta', t', x')) = \langle \tilde{\phi}(u_\theta, t, x), \tilde{\phi}(u_{\theta'}, t', x') \rangle_{\mathcal{F}}, \quad (279)$$

where $\tilde{\phi} : \mathcal{H} \times [0, T] \times \mathbb{R}^n \rightarrow \mathcal{F}$ is the feature map previously defined.

Define the bounded set $D_c = \text{supp}(p_c) \subset \mathbb{R}^m$. We have $\theta, t, x \mapsto (\hat{b}, \hat{\sigma}^2)(t, x, u_\theta) \in W^q(D_c \times D, \mathbb{R}^{n+n^2})$ for any $(\hat{b}, \hat{\sigma}^2) \in \mathcal{F}$. This holds because $\theta, t \mapsto u_\theta(t) \in W^q(D_c \times [0, T])$ and $\mathcal{F} \subset \{(\hat{b}, \hat{\sigma}^2 + \kappa I_{\mathbb{R}^n \times n}) \mid (\hat{b}, \hat{\sigma}^2) \in W_2^q(D \times \mathbb{R}^d, \mathbb{R}^{n+n^2})\}$, and using results on composition for Sobolev spaces (see Theorem 9.1 in Bourdaud [12]). Therefore, we have $\mathcal{M} \subset W_2^{q-2}(\mathbb{R}^m \times [0, T] \times \mathbb{R}^n, \mathbb{R})$ since $p \in W_2^{q-2}(\mathbb{R}^m \times [0, T] \times \mathbb{R}^n, \mathbb{R})$ (see Theorem 6.1 in Behzadan and Holst [8]).

Moreover, defining the bounded linear operator

$$A : \mathcal{F} \rightarrow W_2^{q-2}(\mathbb{R}^m \times [0, T] \times \mathbb{R}^n) \quad (280)$$

$$(b, \sigma^2) \mapsto (\mathcal{L}^{(b, \sigma)(u_\theta)})^* p(u_\theta), \quad (281)$$

for any $v \in \mathcal{F}$ and $\theta, t, x \in \mathbb{R}^m \times [0, T] \times \mathbb{R}^n$, we have

$$\langle v, \tilde{\phi}(u_\theta, t, x) \rangle_{\mathcal{F}} = (Av)(\theta, t, x) = \langle Av, \phi_{q-2}(\theta, t, x) \rangle_{\mathcal{M}} = \langle v, A^* \phi_{q-2}(\theta, t, x) \rangle_{\mathcal{F}}, \quad (282)$$

where $\phi_{q-2, n+1}(\theta, t, x) \triangleq k_{q-2, n+1}((\theta, t, x), \cdot)$. Hence, for any $\theta, t, x \in \mathbb{R}^m \times [0, T] \times \mathbb{R}^n$,

$$\tilde{\phi}(u_\theta, t, x) = A^* \phi_{q-2}(\theta, t, x). \quad (283)$$

Therefore, from the embedding property of Sobolev RKHSs, we have

$$(\tilde{\phi} \otimes \tilde{\phi})(u_\theta, t, x) = A^* ((\phi_{q-2} \otimes \phi_{q-2})(\theta, t, x)) A \quad (284)$$

$$\preceq A^* \mathbb{E}_{t, x} [(\phi_{q-2} \otimes \phi_{q-2})(\theta, t, x)]^{1 - \frac{n+1}{2(q-2)}} A \quad (285)$$

$$= \|A\|_{\infty}^{\frac{n+1}{q-2}} \mathbb{E}_{t, x} [(\tilde{\phi} \otimes \tilde{\phi})(u_\theta, t, x)]^{1 - \frac{n+1}{2(q-2)}} \quad (286)$$

$$= \|A\|_{\infty}^{\frac{n+1}{q-2}} C(u_\theta)^{1 - \frac{n+1}{2(q-2)}} \quad \text{a.s.} \quad (287)$$

This concludes the proof for obtaining Assumption (A6.1). Similar proof gives Assumption (A6.2). \square

8 Implementation details of uncontrolled SDE estimation

This section details the implementation of the uncontrolled SDE estimation method proposed in Bonalli and Rudi [11], available on GitHub ([lmotte/sde-learn](https://github.com/lmotte/sde-learn)) as an open-source Python library. We provide details on the computations involved, including vectorized versions for efficient computation with Python libraries such as NumPy, as well as the computational complexity of each step. In Section 8.1, we detail probability density estimation. In Section 8.2, Fokker-Planck matching is presented.

Notations. We use the following notations.

$$\mathbb{1} \triangleq (1, 1, \dots, 1) \in \mathbb{R}^N, \quad \hat{\rho}_i \triangleq \frac{\partial \hat{\rho}}{\partial x_i}, \quad \hat{\rho}_{ij} \triangleq \frac{\partial^2 \hat{\rho}(x)}{\partial x_i \partial x_j}, \quad (288)$$

and similar notations for the partial derivatives of \hat{g} and ρ .

8.1 Step 1: probability density estimation

8.1.1 Estimator closed-form

From Bonalli and Rudi [11], we have

$$\hat{\rho}(t, x) = k_t(t)^T K_t^{-1} \hat{g}(x), \quad (289)$$

with $k_t(t) = (k_t(t, t_l))_l \in \mathbb{R}^M$, $\hat{g}(x) = N^{-1} \mathbf{1}^T (\rho(x, X_{kl}))_k$, and $\rho(x, y) = \mu^n (2\pi)^{-n/2} \exp\left(-\frac{\mu^2}{2} \|x - y\|^2\right)$. We consider k_t defined as the Gaussian kernel with parameter ν : $k_t(t, t') \triangleq \exp(-\nu(t - t')^2)$.

8.1.2 Algorithm

Inputs. We are provided with a dataset of Q sample paths $(X^{tr}(w_i, t_l))_{i \in \llbracket 1, Q \rrbracket, l \in \llbracket 1, M \rrbracket}$ sampled from an unknown SDE, along with the hyper-parameters $\nu, \mu > 0$.

Step 1.1 ($\hat{\rho}$ fitting). We compute and store $K_t^{-1} = (k_t(t_k, t_l))_{k,l}^{-1} \in \mathbb{R}^{M \times M}$, $(T_l^{tr})_l \triangleq (t_l)_l \in \mathbb{R}^M$, and $(X_{kl})_{kl} \triangleq (X^{tr}(w_k, t_l))_{kl} \in \mathbb{R}^{Q \times M \times n}$.

The time and space complexities of this step are $\mathcal{O}(M^3)$ and $\mathcal{O}(M^2 + QMn)$, respectively.

Step 1.2 ($\hat{\rho}$ prediction). For any $T^{te} \in [0, T]^{M_{te}}$ and $X^{te} \in (\mathbb{R}^n)^{N_{te}}$, the evaluations for each $(t, x) \in T^{te} \times X^{te}$ can be computed as

$$(\hat{\rho}(t_i, x_j))_{i,j} = K_t^{te, tr} \times K_t^{-1} \times N^{-1} \mathbf{1}^T \times G^{te} \in \mathbb{R}^{M_{te} \times N_{te}}, \quad (290)$$

where

$$\begin{aligned} K_t^{te} &\triangleq (k_t(T_j^{te}, t_i))_{ji} \in \mathbb{R}^{M_{te} \times M}, \\ G^{te} &\triangleq (\rho(X_j^{te}, X_{il}))_{ilj} \in \mathbb{R}^{N \times M \times N_{te}}, \end{aligned}$$

and similar formulas hold for $\hat{\rho}_i$ (or $\hat{\rho}_{ij}$) by replacing ρ with ρ_i (or ρ_{ij}).

In particular, the fitting phase of step 2 requires the computation of the \hat{P}_{ij} and \hat{d} . Using a dataset $Z^{fp} = T^{fp} \times X^{fp}$, with $T^{fp} \in \mathbb{R}^{M_{fp}}$ and $X^{fp} \in \mathbb{R}^N$, and denoting $D_x \triangleq (X_j^{fp} - X_{il})_{ilj} \in \mathbb{R}^{N \times N \times M}$, the \hat{P}_{ij} can be computed recursively as follows

$$\begin{aligned} G^{fp} &= (\rho(X_{il}, X_j^{fp}))_{ilj}, \\ G_i^{fp} &= (\rho_i(X_{il}, X_j^{fp}))_{ilj} = -\mu^2 D_x e_i \odot G^{fp}, \\ G_{ij}^{fp} &= (\rho_{ij}(X_{il}, X_j^{fp}))_{ilj} = -\mu^2 (D_x e_j \odot G_i^{fp} + \delta_{ij} G^{fp}), \\ \hat{P} &= \text{vec}(K_t^{fp, tr} \times K_t^{-1} \times N^{-1} \mathbf{1}^T G^{fp}), \\ \hat{P}_i &= \text{vec}(K_t^{fp, tr} \times K_t^{-1} \times N^{-1} \mathbf{1}^T G_i^{fp}), \\ \hat{P}_{ij} &= \text{vec}(K_t^{fp, tr} \times K_t^{-1} \times N^{-1} \mathbf{1}^T G_{ij}^{fp}). \end{aligned}$$

Moreover, denoting $D_t \triangleq (T_j^{fp} - T_l)_{lj} \in \mathbb{R}^{M_{fp} \times M}$, we have

$$\hat{d} = \text{vec}((-2\nu D_t \odot K_t^{fp, tr}) \times K_t^{-1} \times N^{-1} \mathbf{1}^T G^{fp}).$$

The time and space complexities are $\mathcal{O}(QMN \max(M_{fp}, M))$ and $\mathcal{O}(QMN)$, respectively.

Outputs. We return \hat{P} , $(\hat{P}_i)_{i=1}^n$, $(\hat{P}_{ij})_{i,j=1}^n$ and \hat{d} .

8.2 Step 2: Fokker-Planck matching

8.2.1 Estimator closed-form

Optimization objective. We consider a uniform diffusion model where $\sigma^2(t, x) = \sigma_0(t, x)^2 I_{\mathbb{R}^n}$, and we enforce the positivity of σ^2 over a subset of the training points $(t_i, x_i)_{i \in I} \subset (t_i, x_i)_{i=1}^N$. Specifically, we solve the optimization problem:

$$\min_{\substack{(b, \sigma^2) \\ \forall i \in I, \sigma_0^2(t_i, x_i) \geq 0}} \frac{1}{N} \sum_{i=1}^N \left(\frac{\partial \hat{p}}{\partial t}(t_i, x_i) - (\mathcal{L}_t^{b, \sigma})^* \hat{p}(t_i, x_i) \right)^2 + \lambda \| (b, \sigma^2) \|_{\mathcal{F}}^2. \quad (291)$$

Solving the optimization problem. Given a p.d. kernel k over $[0, T] \times \mathbb{R}^n$, we consider the model

$$\begin{aligned} \forall i \in \llbracket 1, n \rrbracket, \quad b^i(t, x) &= \langle w_b^i, \phi(t, x) \rangle_{\mathcal{G}} \quad \text{with} \quad w_b^i \in \mathcal{G}, \\ \sigma^2(t, x) &= \langle w_\sigma, \phi(t, x) \rangle_{\mathcal{G}} I_{\mathbb{R}^n} \quad \text{with} \quad w_\sigma \in \mathcal{G}, \end{aligned}$$

where $\phi(t, x) \triangleq k((t, x), (\cdot, \cdot)) \in \mathcal{G}$. Therefore, defining $w = ((w_b^i)_{i=1}^n | w_\sigma) \in \mathcal{G}^{n+1}$, Eq. (291) is expressed as

$$\min_{w \in \mathcal{G}^{n+1}} \max_{\gamma \geq 0} \| (\hat{C} + \lambda I)^{1/2} w - (\hat{C} + \lambda I)^{-1/2} v \|_{\mathcal{G}^{n+1}}^2 - 2 \sum_{i \in I} \gamma_i \langle w_\sigma, \phi(t_i, x_i) \rangle_{\mathcal{G}}, \quad (292)$$

where $w = (w_b | w_\sigma) \in \mathcal{G}^{n+1}$, $\hat{C} = N^{-1} \sum_{i=1}^N (\tilde{\phi} \otimes \tilde{\phi})(t_i, x_i)$, $v = N^{-1} \sum_{i=1}^N (\frac{\partial \hat{p}}{\partial t} \tilde{\phi})(t_i, x_i)$, and

$$\tilde{\phi} = \left((-\tilde{\phi}_i)_i | 1/2 \sum_i \tilde{\phi}_{ii} \right) \in \mathcal{G}^{n+1}. \quad (293)$$

Note the change of feature map $\tilde{\phi}$ compared to the non-uniform diffusion model where

$$\tilde{\phi} = ((-\tilde{\phi}_i)_i | (1/2 \tilde{\phi}_{ij})_{ij}) \in \mathcal{G}^{n+n^2}. \quad (294)$$

This can be equivalently expressed as

$$\min_{w \in \mathcal{G}^{n+1}} \max_{\gamma \geq 0} \| w (\hat{C} + \lambda I)^{1/2} - (v + U^* \phi_\gamma) (\hat{C} + \lambda I)^{-1/2} \|_{\mathcal{G}^{n+1}}^2 - \| \hat{C}_\lambda^{-1/2} (v + U^* \phi_\gamma) \|_{\mathcal{G}^{n+1}}^2, \quad (295)$$

where $\phi_\gamma = \sum_{i \in I} \gamma_i \phi(t_i, x_i)$, $U = \sum_i e_i \otimes (0_{\mathcal{G}^n} | e_i)_i \in \mathcal{G} \otimes \mathcal{G}^{n+1}$, with $(e_i)_i$ being an orthonormal basis of \mathcal{G} , ensuring $w^\sigma = U w$.

Strong duality guarantees solvability for $w \in \mathcal{G}^{n+1}$ via

$$\hat{w}_{pc} \triangleq (\hat{C} + \lambda I)^{-1} (v + U^* \phi_\gamma), \quad (296)$$

and the optimal $\gamma \in \mathbb{R}^n$ is obtained by solving the dual problem

$$\max_{\gamma \geq 0} QP(\gamma), \quad (297)$$

with $QP(\gamma) \triangleq -\phi_\gamma^* U C_\lambda^{-1} U^* \phi_\gamma - 2 \phi_\gamma^* U C_\lambda^{-1} v$.

Closed-form formula for $\hat{b}_{pc}, \hat{\sigma}_{pc}^2$ and $QP(\gamma)$. We decompose

$$\hat{w}_{pc} = \hat{w}_{std} + \hat{w}_+, \quad (298)$$

where $\hat{w}_{std} \triangleq (\hat{C} + \lambda I)^{-1}v$ is the standard unconstrained ridge regression formula, and $\hat{w}_+ \triangleq (\hat{C} + \lambda I)^{-1}U^*\phi_\gamma$ is an adjustment stemming from the positivity constraint.

Additionally, by denoting $\tilde{K} = \sum_i \tilde{K}_i + \frac{1}{4} \sum_{i,j=1}^n \tilde{K}_{ii}^{jj}$, we deduce

$$\hat{w}_+ = N(\tilde{\Phi}^*\tilde{\Phi} + N\lambda I)^{-1}U^*\phi_\gamma \quad (299)$$

$$= \lambda^{-1}(U^*\phi_\gamma - \tilde{\Phi}^*(\tilde{\Phi}\tilde{\Phi}^* + N\lambda I)^{-1}\tilde{\Phi}U^*\phi_\gamma) \quad (300)$$

$$= \lambda^{-1}(U^*\phi_\gamma - \tilde{\Phi}^*\tilde{K}_{N\lambda}^{-1}\tilde{\Phi}U^*\phi_\gamma), \quad (301)$$

using the Woodbury identity for any operators A and B : $(AB + \lambda I)^{-1} = \lambda^{-1}(I - A(BA + \lambda I)^{-1}B)$. Hence, setting $V_i = \sum_j e_j \otimes (e_j | 0_{\mathcal{G}})$, we derive

$$\hat{b}_{pc}^i = \langle V_i \hat{w}_{pc}, \phi(t, x) \rangle_{\mathcal{G}} \quad (302)$$

$$= \underbrace{\langle V_i \hat{w}_{std}, \phi(t, x) \rangle_{\mathcal{G}}}_{\triangleq \hat{b}_{std}^i} + \underbrace{\langle V_i \hat{w}_+, \phi(t, x) \rangle_{\mathcal{G}}}_{\triangleq \hat{b}_+^i}, \quad (303)$$

Moreover, using $UV_i^* = 0$, we have

$$\hat{b}_+^i = \langle V_i \hat{w}_+, \phi(t, x) \rangle_{\mathcal{G}} \quad (304)$$

$$= \lambda^{-1} \left(\phi_\gamma^* UV_i^* \phi(t, x) - \phi_\gamma^* U \tilde{\Phi}^* \tilde{K}_{N\lambda}^{-1} \tilde{\Phi} V_i^* \phi(t, x) \right) \quad (305)$$

$$= -\lambda^{-1} \phi_\gamma^* U \tilde{\Phi}^* \tilde{K}_{N\lambda}^{-1} \tilde{\Phi} V_i^* \phi(t, x). \quad (306)$$

Similarly,

$$\hat{\sigma}_{pc}^2 = (\hat{\sigma}_{std}^2 + \hat{\sigma}_+^2) I_{\mathbb{R}^n}, \quad (307)$$

where $\hat{\sigma}_{std}^2$ is the standard unconstrained ridge regression estimator, and

$$\hat{\sigma}_+^2 = \lambda^{-1} \left(\phi_\gamma^* \phi(t, x) - \phi_\gamma^* U \tilde{\Phi}^* \tilde{K}_{N\lambda}^{-1} \tilde{\Phi} U^* \phi(t, x) \right). \quad (308)$$

Furthermore,

$$QP(\gamma) = -\phi_\gamma^* UC_\lambda^{-1} U^* \phi_\gamma - 2\phi_\gamma^* UC_\lambda^{-1} v \quad (309)$$

$$= -\gamma^T \Phi UC_\lambda^{-1} U^* \Phi^* \gamma - 2N^{-1} \gamma^T \Phi UC_\lambda^{-1} \tilde{\Phi}^* \hat{d}. \quad (310)$$

Therefore, we conclude

$$\hat{b}_{pc}^i(t, x) = \hat{b}_{std}^i(t, x) + \hat{b}_+^i(t, x), \quad (311)$$

$$\hat{b}_{std}^i(t, x) = \tilde{d}^T \tilde{K}_{N\lambda}^{-1} r_i^b(t, x), \quad (312)$$

$$\hat{b}_+^i(t, x) = -\lambda^{-1} \gamma^T (R_{fp,pc}^\sigma)^T \tilde{K}_{N\lambda}^{-1} r_i^b(t, x), \quad (313)$$

$$\hat{\sigma}_{pc}^2(t, x) = \hat{\sigma}_{std}^2(t, x) + \hat{\sigma}_+^2(t, x), \quad (314)$$

$$\hat{\sigma}_{std}^2(t, x) = \tilde{d}^T \tilde{K}_{N\lambda}^{-1} r^\sigma(t, x), \quad (315)$$

$$\hat{\sigma}_+^2(t, x) = \lambda^{-1} \gamma^T \left(k(t, x) - (R_{fp,pc}^\sigma)^T \tilde{K}_{N\lambda}^{-1} r^\sigma(t, x) \right) I_{\mathbb{R}^n}, \quad (316)$$

$$QP(\gamma) = -\lambda^{-1} \gamma^T (K - (R_{fp,pc}^\sigma)^T \tilde{K}_{N\lambda}^{-1} (R_{fp,pc}^\sigma)) \gamma - 2\gamma^T (R_{fp,pc}^\sigma)^T \tilde{K}_{N\lambda}^{-1} \hat{d}, \quad (317)$$

using the following notations

$$z, z' = (t, x), (t', x') \quad (318)$$

$$r_i^b(z, z') \triangleq -\langle \tilde{\phi}_i(z), \phi(z') \rangle_{\mathcal{G}}, \quad (319)$$

$$r_i^b(z) \triangleq (r_i^b(z_l, z))_l \in \mathbb{R}^N, \quad (320)$$

$$r_{ij}^\sigma(z, z') \triangleq \langle \tilde{\phi}_{ij}(z), \phi(z') \rangle_{\mathcal{G}}, \quad (321)$$

$$r^\sigma(z, z') \triangleq 1/2 \sum_i r_{ii}^\sigma(z, z'), \quad (322)$$

$$r^\sigma(z) \triangleq (r^\sigma(z_l, z))_l \in \mathbb{R}^N, \quad (323)$$

$$R_{fp,pc}^\sigma \triangleq 1/2 \sum_i \tilde{\Phi}_{ii} \Phi \in \mathbb{R}^{N \times n_{pc}}, \quad (324)$$

$$\hat{d} \triangleq \left(\frac{\partial \hat{p}}{\partial t}(z_l) \right)_l \in \mathbb{R}^N. \quad (325)$$

8.2.2 Gram matrices computations

The term $(\mathcal{L}_t^{b,\sigma})^* \hat{p}$ involves second derivatives of $\sigma \times \hat{p}$, leading to the feature maps $\tilde{\phi}_{ij}$, which are sums of four terms each. Consequently, the evaluation of the scalar product between two feature map values results in a sum of 16 terms.

\tilde{K} computation. Denoting \odot as the Hadamard product, and $A \odot BC \triangleq A \odot (BC)$, we have

$$\tilde{K} \triangleq (\langle \tilde{\phi}(z_p), \tilde{\phi}(z_q) \rangle_{\mathcal{G}^{n+1}})_{p,q=1}^N = \sum_{i=1}^n \tilde{K}_i + \frac{1}{4} \sum_{i,j=1}^n \tilde{K}_{ii}^{jj}, \quad (326)$$

with

$$\tilde{K}_i \triangleq (\langle \tilde{\phi}_i(z_p), \tilde{\phi}_i(z_q) \rangle_{\mathcal{G}^{n+1}})_{p,q=1}^N = K \odot \hat{P}_i \hat{P}_i^T + K_i \odot \hat{P} \hat{P}_i^T + K^i \odot \hat{P}_i \hat{P}^T + K_i^i \odot \hat{P} \hat{P}^T,$$

and

$$\begin{aligned} \tilde{K}_{kl}^{ij} \triangleq \left(\langle \tilde{\phi}_{kl}(z_p), \tilde{\phi}_{il}(z_q) \rangle_{\mathcal{G}^{n+1}} \right)_{p,q=1}^N &= K \odot \hat{P}_{kl} \hat{P}_{ij}^T + K_k \odot \hat{P}_l \hat{P}_{ij}^T + K_l \odot \hat{P}_k \hat{P}_{ij}^T + K_{kl} \odot \hat{P} \hat{P}_{ij}^T \\ &+ K^i \odot \hat{P}_{kl} \hat{P}_j^T + K_k^i \odot \hat{P}_l \hat{P}_j^T + K_l^i \odot \hat{P}_k \hat{P}_j^T + K_{kl}^i \odot \hat{P} \hat{P}_j^T \\ &+ K^j \odot \hat{P}_{kl} \hat{P}_i^T + K_k^j \odot \hat{P}_l \hat{P}_i^T + K_l^j \odot \hat{P}_k \hat{P}_i^T + K_{kl}^j \odot \hat{P} \hat{P}_i^T \\ &+ K^{ij} \odot \hat{P}_{kl} \hat{P}^T + K_k^{ij} \odot \hat{P}_l \hat{P}^T + K_l^{ij} \odot \hat{P}_k \hat{P}^T + K_{kl}^{ij} \odot \hat{P} \hat{P}^T, \end{aligned} \quad (327)$$

where $K_{ij}^{kl} = (k_{ij}^{kl}(z_p, z_q))_{p,q=1}^N$, and $\hat{P}_{ij} = (\hat{p}_{ij}(z_p))_{p=1}^N$.

Fast computation of the Gram matrices K_{ij}^{kl} with Gaussian kernel. If

$$k((t, x), (t', x')) \triangleq \exp(-\gamma \|(t, x), -(t', x')\|^2) \quad \text{for } \gamma > 0, \quad (328)$$

then, denoting $D_z \triangleq (Z_k - Z_l)_{k,l=1}^N$, the K_{pq}^{kl} can be computed recursively as follows. For $i, j, k, l \in \llbracket 1, n \rrbracket$,

$$\begin{aligned}
 K_i &= -2\gamma D_z e_i \odot K, \\
 K^i &= (-K_i), \\
 K_{ij} &= -2\gamma (D_z e_j \odot K_i + \delta_{ij} K), \\
 K^{ij} &= K_{ij}, \\
 K_i^j &= -K_{ij}, \\
 K_k^{ij} &= -2\gamma (D_z e_j \odot K_{ik} + \delta_{jk} K_i + \delta_{ij} K_k), \\
 K_{ij}^k &= -K_i^{jk}, \\
 K_{kl}^{ij} &= -2\gamma (-D_z e_j \odot K_{kl}^i + \delta_{jl} K_{ik} + \delta_{jk} K_{il} + \delta_{ij} K_{kl}).
 \end{aligned}$$

8.2.3 Algorithm

Inputs. We are provided with a dataset $(z_l)_{l=1}^N = ((x_l, t_l))_{l=1}^N$ i.i.d. from p_s , along with a p.d. kernel k over $[0, T] \times \mathbb{R}^n$, and the hyper-parameter $\lambda > 0$.

Step 2.1 ($(\hat{b}, \hat{\sigma}^2)$ fitting). From Sections 8.2.1 and 8.2.2, at fitting time, one needs to compute and store

$$(\tilde{K} + N\lambda I)^{-1} \in \mathbb{R}^{N \times N}. \quad (329)$$

The time and space complexities of this step are $\mathcal{O}(N^3)$ and $\mathcal{O}(N^2)$, respectively.

Step 2.2 ($(\hat{b}, \hat{\sigma}^2)$ predictions). Denoting $D_z^{te} = (Z_l - Z_l^{te})_{k,l} \in \mathbb{R}^{N \times N_{te}}$, we have

$$\begin{aligned}
 K_i^{fp,te} &= -2\gamma D_z^{te} e_i \odot K^{fp,te}, \\
 K_{ij}^{fp,te} &= -2\gamma (D_z^{te} e_j \odot K_i^{fp,te} + \delta_{ij} K^{fp,te}).
 \end{aligned}$$

Then, predictions can be computed from the formulas provided in Section 7.3.

The time and space complexities of this step are $\mathcal{O}(NN_{te})$ and $\mathcal{O}(NN_{te})$, respectively.

Outputs. We return the predicted SDE coefficients' values for the given test inputs.

8.2.4 FP matching with Nyström approximation

Time and space complexities of Fokker-Planck matching can be reduced from $\mathcal{O}(N^3)$ and $\mathcal{O}(N^2)$ to $\mathcal{O}(mN^2)$ and $\mathcal{O}(Nm)$ by using Nyström approximation with m anchors. More precisely, the Fokker-Planck matching objective with the Nyström approximation

$$w = \sum_{i=1}^m \alpha_i \tilde{\phi}(z_i), \quad (330)$$

writes

$$\min_{\alpha \in \mathbb{R}^m} \alpha^T \left(\tilde{K}_{nm}^T \tilde{K}_{nm} + n\lambda \tilde{K}_{mm} \right) \alpha - 2\alpha^T \left(\tilde{K}_{mn} Y_{tr} + R_{mn_{pc}}^\sigma \gamma \right). \quad (331)$$

It can be solved with

$$\alpha = \left(\tilde{K}_{nm}^T \tilde{K}_{nm} + n\lambda \tilde{K}_{mm} \right)^{-1} \left(\tilde{K}_{mn} Y_{tr} + R_{mn_{pc}}^\sigma \gamma \right). \quad (332)$$

We deduce the following formula for $\hat{b}_{ny,pc}$, $\hat{\sigma}_{ny,pc}^2$.

$$\hat{b}_{ny,pc}(z) = \alpha^T R_m^b(z), \quad (333)$$

$$\hat{\sigma}_{ny,pc}^2(z) = \alpha^T R_m^\sigma(z), \quad (334)$$

and the following formula for the dual problem

$$\min_{\gamma \geq 0} \gamma^T A \gamma + \gamma^T b, \quad (335)$$

with

$$A = (R_{mn_{pc}}^\sigma)^T \left(\tilde{K}_{nm}^T \tilde{K}_{nm} + n\lambda \tilde{K}_{mm} \right)^{-1} R_{mn_{pc}}^\sigma, \quad (336)$$

$$b = 2(R_{mn_{pc}}^\sigma)^T \left(\tilde{K}_{nm}^T \tilde{K}_{nm} + n\lambda \tilde{K}_{mm} \right)^{-1} \tilde{K}_{mn} \hat{d}. \quad (337)$$

9 Implementation details of controlled SDE estimation

This section details the implementation of the controlled SDE estimation method proposed in this work, available on GitHub ([lmotte/controlled-sde-learn](https://github.com/lmotte/controlled-sde-learn)) as an open-source Python library. We update the formulas and computational complexity of each step provided in Section 8 for uncontrolled SDEs to adapt them for controlled SDEs.

9.1 Step 1: probability density estimation

For controlled SDEs, the estimation of the probability density function involves repeating the algorithm from Step 1 of the uncontrolled case K times. Each iteration corresponding to a specific control setting. At the end of this step, we have computed and stored

$$(\hat{p}_k((t_i, x_j)))_{i,j,k} \in \mathbb{R}^{M_{fp} \times N_{fp} \times K}, \quad (338)$$

and also the partial derivatives evaluations.

The time and space complexities of this step are $\mathcal{O}(KM^3)$ and $\mathcal{O}(K(M^2 + QMn))$ for fitting, and $\mathcal{O}(KQMN \max(M_{fp}, M))$ and $\mathcal{O}(KQMN)$ for predictions.

9.2 Step 2: Fokker-Planck matching

Same algorithm but adding additional dimensions for the controls. More precisely, closed-form are updated with $z = (t, x, v) \in [0, T] \times \mathbb{R}^n \times \mathbb{R}^d$ instead of $z = (t, x) \in [0, T] \times \mathbb{R}^n$. We store

$$Z_{tr} = (t_i, x_i, u_k(t_i))_{i=1, \dots, N_{fp}, k=1, \dots, K}, \quad (339)$$

and compute and store

$$(\tilde{K} + NK\lambda I)^{-1} \in \mathbb{R}^{NK \times NK}. \quad (340)$$

The time and space complexities of this step are $\mathcal{O}(K^3N^3)$ and $\mathcal{O}(K^2N^2)$ for fitting, and $\mathcal{O}(KNN_{te})$ and $\mathcal{O}(KNN_{te})$ for predictions.

References

- [1] Ian Abraham and Todd D Murphey. Active learning of dynamics for data-driven control using koopman operators. *IEEE Transactions on Robotics*, 35(5):1071–1083, 2019.
- [2] Kweku Abraham. Nonparametric Bayesian posterior contraction rates for scalar diffusions with high-frequency data. *Bernoulli*, 25(4A):2696 – 2728, 2019. doi: 10.3150/18-BEJ1067. URL <https://doi.org/10.3150/18-BEJ1067>.
- [3] Carl Andersson, Antônio H Ribeiro, Koen Tiels, Niklas Wahlström, and Thomas B Schön. Deep convolutional networks in system identification. In *2019 IEEE 58th conference on decision and control (CDC)*, pages 3670–3676. IEEE, 2019.
- [4] Karl Johan Åström and Peter Eykhoff. System identification—a survey. *Automatica*, 7(2):123–162, 1971.
- [5] Karl-Johan Åström and Bohlin Torsten. Numerical identification of linear dynamic systems from normal operating records. *IFAC Proceedings Volumes*, 2(2):96–111, 1965.
- [6] Sohail Bahmani and Justin Romberg. Convex programming for estimation in nonlinear recurrent models. *Journal of Machine Learning Research*, 21(235):1–20, 2020.
- [7] Benedikt Bauer, Luc Devroye, Michael Kohler, Adam Krzyżak, and Harro Walk. Nonparametric estimation of a function from noiseless observations at random points. *Journal of Multivariate Analysis*, 160:93–104, 2017.
- [8] Ali Behzadan and Michael Holst. Multiplication in sobolev spaces, revisited. *Arkiv för Matematik*, 59(2):275–306, 2021.
- [9] Eloïse Berthier, Justin Carpentier, Alessandro Rudi, and Francis Bach. Infinite-dimensional sums-of-squares for optimal control. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 577–582. IEEE, 2022.
- [10] Raphaël Berthier, Francis Bach, and Pierre Gaillard. Tight nonparametric convergence rates for stochastic gradient descent under the noiseless linear model. *Advances in Neural Information Processing Systems*, 33:2576–2586, 2020.
- [11] Riccardo Bonalli and Alessandro Rudi. Non-parametric learning of stochastic differential equations with fast rates of convergence. *arXiv preprint arXiv:2305.15557*, 2023.
- [12] Gerard Bourdaud. An introduction to composition operators in sobolev spaces. *Eurasian Mathematical Journal*, 14(1):39–54, 2023.
- [13] Alexander P Browning, David J Warne, Kevin Burrage, Ruth E Baker, and Matthew J Simpson. Identifiability analysis for stochastic differential equation models in systems biology. *Journal of the Royal Society Interface*, 17(173):20200652, 2020.
- [14] Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7:331–368, 2007.
- [15] Yongxin Chen and Umesh Vaidya. Sample complexity for nonlinear stochastic dynamics. In *2019 American Control Conference (ACC)*, pages 3526–3531. IEEE, 2019.
- [16] Carlo Ciliberto, Lorenzo Rosasco, and Alessandro Rudi. A general framework for consistent structured prediction with implicit loss embeddings. *Journal of Machine Learning Research*, 21(98):1–67, 2020.
- [17] Fabienne Comte and Valentine Genon-Catalot. Nonparametric drift estimation for iid paths of stochastic differential equations. *The Annals of Statistics*, 48(6):3336–3365, 2020.
- [18] Fabienne Comte, Valentine Genon-Catalot, and Yves Rozenholc. Penalized nonparametric mean square estimation of the coefficients of diffusion processes. *Bernoulli*, 13(2):514 – 543, 2007. doi: 10.3150/07-BEJ5173. URL <https://doi.org/10.3150/07-BEJ5173>.

- [19] Arnak Dalalyan. Sharp adaptive estimation of the drift function for ergodic diffusions. *The Annals of Statistics*, 33(6):2507 – 2528, 2005. doi: 10.1214/009053605000000615. URL <https://doi.org/10.1214/009053605000000615>.
- [20] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20(4):633–679, 2020.
- [21] Luc Devroye, László Györfi, and Gábor Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.
- [22] Simon Fischer and Ingo Steinwart. Sobolev norm learning rates for regularized least-squares algorithms. *Journal of Machine Learning Research*, 21(205):1–38, 2020.
- [23] Danielle Florens-Zmirou. On estimating the diffusion coefficient from discrete observations. *Journal of applied probability*, 30(4):790–804, 1993.
- [24] Dylan Foster, Tuhin Sarkar, and Alexander Rakhlin. Learning nonlinear dynamical systems from a single trajectory. In *Learning for Dynamics and Control*, pages 851–861. PMLR, 2020.
- [25] Avner Friedman. *Partial differential equations of parabolic type*. Courier Dover Publications, 2008.
- [26] Valentine Genon-Catalot, Catherine Laredo, and Dominique Picard. Non-parametric estimation of the diffusion coefficient by wavelets methods. *Scandinavian Journal of Statistics*, pages 317–335, 1992.
- [27] Frank Hansen and Gert K Pedersen. Jensen’s operator inequality. *Bulletin of the London Mathematical Society*, 35(4):553–564, 2003.
- [28] Lukas Hewing, Kim P Wabersich, Marcel Menner, and Melanie N Zeilinger. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1):269–296, 2020.
- [29] BL Ho and Rudolf E Kálmán. Effective construction of linear state-variable models from input/output functions: Die konstruktion von linearen modeilen in der darstellung durch zustandsvariable aus den beziehungen für ein-und ausgangsgrößen. *at-Automatisierungstechnik*, 14(1-12):545–548, 1966.
- [30] Marc Hoffmann. Adaptive estimation in diffusion processes. *Stochastic processes and their Applications*, 79(1):135–163, 1999.
- [31] Rolf Isermann and Marco Münchhof. *Identification of dynamic systems: an introduction with applications*, volume 85. Springer, 2011.
- [32] Tohru Katayama et al. *Subspace methods for system identification*, volume 1. Springer, 2005.
- [33] Mathieu Kessler and Michael Sørensen. Estimating equations based on eigenfunctions for a discretely observed diffusion process. *Bernoulli*, pages 299–314, 1999.
- [34] Stefan Klus, Feliks Nüske, Sebastian Peitz, Jan-Hendrik Niemann, Cecilia Clementi, and Christof Schütte. Data-driven approximation of the koopman generator: Model reduction, system identification, and control. *Physica D: Nonlinear Phenomena*, 406:132416, 2020.
- [35] Yury A Kutoyants. *Statistical inference for ergodic diffusion processes*. Springer Science & Business Media, 2013.
- [36] Thomas Lew, Apoorva Sharma, James Harrison, Andrew Bylard, and Marco Pavone. Safe active dynamics learning and control: A sequential exploration–exploitation framework. *IEEE Transactions on Robotics*, 38(5):2888–2907, 2022.
- [37] L. Ljung. *System Identification: Theory for the User*. Prentice Hall information and system sciences series. Prentice Hall PTR, 1999. ISBN 9780136566953. URL <https://books.google.fr/books?id=nHFOqgAACAAJ>.

- [38] Lennart Ljung, Carl Andersson, Koen Tiels, and Thomas B Schön. Deep learning and system identification. *IFAC-PapersOnLine*, 53(2):1175–1181, 2020.
- [39] Alessandra Lunardi. *Analytic semigroups and optimal regularity in parabolic problems*. Springer Science & Business Media, 2012.
- [40] Horia Mania, Michael I Jordan, and Benjamin Recht. Active learning for nonlinear system identification with guarantees. *Journal of Machine Learning Research*, 23(32):1–30, 2022.
- [41] Ulysse Marteau-Ferey, Francis Bach, and Alessandro Rudi. Non-parametric models for non-negative functions. *Advances in neural information processing systems*, 33:12816–12826, 2020.
- [42] Hongyu Miao, Xiaohua Xia, Alan S Perelson, and Hulin Wu. On identifiability of nonlinear ode models and applications in viral dynamics. *SIAM review*, 53(1):3–39, 2011.
- [43] Christopher W Miller and Insoon Yang. Optimal control of conditional value-at-risk in continuous time. *SIAM Journal on Control and Optimization*, 55(2):856–884, 2017.
- [44] Boris Muzellec, Francis Bach, and Alessandro Rudi. Learning psd-valued functions using kernel sums-of-squares. *arXiv preprint arXiv:2111.11306*, 2021.
- [45] Oliver Nelles. *Nonlinear dynamic system identification*. Springer, 2020.
- [46] R Nickl and K Ray. Nonparametric statistical inference for drift vector fields of multi-dimensional diffusions. *Annals of Statistics*, 48:1383–1408, 2020. doi: 10.1214/19-AOS1851. URL <http://dx.doi.org/10.1214/19-AOS1851>.
- [47] Jan Nygaard Nielsen, Henrik Madsen, and Peter C Young. Parameter estimation in stochastic differential equations: an overview. *Annual Reviews in Control*, 24:83–94, 2000.
- [48] Feliks Nüske, Sebastian Peitz, Friedrich Philipp, Manuel Schaller, and Karl Worthmann. Finite-data error bounds for koopman-based prediction and control. *Journal of Nonlinear Science*, 33(1):14, 2023.
- [49] Samet Oymak. Stochastic gradient descent learns state equations with nonlinear activations. In *conference on Learning Theory*, pages 2551–2579. PMLR, 2019.
- [50] Loucas Pillaud-Vivien, Alessandro Rudi, and Francis Bach. Statistical optimality of stochastic gradient descent on hard learning problems through multiple passes. *Advances in Neural Information Processing Systems*, 31, 2018.
- [51] Xing Qiu, Tao Xu, Babak Soltanizadeh, and Hulin Wu. Identifiability analysis of linear ordinary differential equation systems with a single trajectory. *Applied Mathematics and Computation*, 430:127260, 2022.
- [52] Alessandro Rudi and Carlo Ciliberto. Psd representations for effective probability models. *Advances in Neural Information Processing Systems*, 34:19411–19422, 2021.
- [53] Alessandro Rudi and Lorenzo Rosasco. Generalization properties of learning with random features. *Advances in neural information processing systems*, 30, 2017.
- [54] Alessandro Rudi, Guillermo D Canas, and Lorenzo Rosasco. On the sample complexity of subspace learning. *Advances in Neural Information Processing Systems*, 26, 2013.
- [55] Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco. Less is more: Nyström computational regularization. *Advances in neural information processing systems*, 28, 2015.
- [56] Alessandro Rudi, Ulysse Marteau-Ferey, and Francis Bach. Finding global minima via kernel approximations. *Mathematical Programming*, pages 1–82, 2024.
- [57] Tuhin Sarkar and Alexander Rakhlin. Near optimal finite time identification of arbitrary linear dynamical systems. In *International Conference on Machine Learning*, pages 5610–5618. PMLR, 2019.
- [58] Yahya Sattar and Samet Oymak. Non-asymptotic and accurate learning of nonlinear dynamical systems. *Journal of Machine Learning Research*, 23(140):1–49, 2022.

- [59] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on stochastic programming: modeling and theory*. SIAM, 2021.
- [60] Max Simchowitz, Horia Mania, Stephen Tu, Michael I Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In *Conference On Learning Theory*, pages 439–473. PMLR, 2018.
- [61] T. Söderström and P. Stoica. *System identification*. Prentice-Hall, Inc., USA, 1988. ISBN 0138812365.
- [62] Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008.
- [63] Ingo Steinwart, Don R Hush, Clint Scovel, et al. Optimal rates for regularized least squares regression. In *COLT*, pages 79–93, 2009.
- [64] Stephen Tu, Roy Frostig, and Mahdi Soltanolkotabi. Learning from many trajectories. *arXiv preprint arXiv:2203.17193*, 2022.
- [65] Adrien Vacher, Boris Muzellec, Alessandro Rudi, Francis Bach, and Francois-Xavier Vialard. A dimension-free computational upper-bound for smooth optimal transport estimation. In *Conference on Learning Theory*, pages 4143–4173. PMLR, 2021.
- [66] Andrew Wagenmaker and Kevin Jamieson. Active learning for identification of linear dynamical systems. In *Conference on Learning Theory*, pages 3487–3582. PMLR, 2020.
- [67] Yuanyuan Wang, Xi Geng, Wei Huang, Biwei Huang, and Mingming Gong. Generator identification for linear sdes with additive and multiplicative noise. *Advances in Neural Information Processing Systems*, 36, 2024.
- [68] Manuel Watter, Jost Springenberg, Joschka Boedecker, and Martin Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. *Advances in neural information processing systems*, 28, 2015.
- [69] Holger Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.
- [70] Jiongmin Yong and Xun Yu Zhou. *Stochastic controls: Hamiltonian systems and HJB equations*, volume 43. Springer Science & Business Media, 2012.
- [71] Christophe Zhang and Enrique Zuazua. A quantitative analysis of koopman operator methods for system identification and predictions. *Comptes Rendus. Mécanique*, 351(S1):1–31, 2023.