

Optimal control for estimation in partially observed elliptic and hypoelliptic linear stochastic differential equations

Quentin Clairon · Adeline Samson

Received: date / Accepted: date

Abstract Multi-dimensional Stochastic Differential Equations (SDEs) are a powerful tool to describe dynamics of phenomena that change over time. We focus on the parametric estimation of such SDEs based on partial observations when only a one-dimensional component of the system is observable. We consider two families of SDE, the elliptic family with a full-rank diffusion coefficient and the hypoelliptic family with a degenerate diffusion coefficient. Estimation for the second class is much more difficult and only few estimation methods have been proposed. Here, we adopt the framework of the optimal control theory to derive a contrast (or cost function) based on the best control sequence mimicking the (unobserved) Brownian motion. We propose a full data-driven approach to estimate the drift and of diffusion coefficient parameters. Numerical simulations made on different examples (Harmonic Oscillator, FitzHugh-Nagumo, Lotka-Volterra) reveal our method produces good pointwise estimate for an acceptable computational price with, interestingly, no performance drop for hypoelliptic systems.

Keywords Stochastic Differential Equations · Hypoellipticity · Estimation · Optimal Control Theory · Linear-Quadratic Theory · Pontryagin maximum principle

1 Introduction

We focus on statistical inference for d -dimensional stochastic dynamical systems modeled by a stochastic differential equation (SDE). We are interested in partial observations: only the first one-dimensional coordinate, denoted V_t , of the system

Q. Clairon
Newcastle University, School of Mathematics and Statistics, Newcastle Upon Tyne, United Kingdom
E-mail: Quentin.Clairon@newcastle.ac.uk

A. Samson
Univ. Grenoble Alpes, CNRS, Grenoble INP*, LJK, 38000 Grenoble, France
* Institute of Engineering Univ. Grenoble Alpes.
E-mail: adeline.leclercq-samson@univ-grenoble-alpes.fr

is observed, the other $(d-1)$ -dimensional coordinates, denoted U_t , are unobserved. The system is written:

$$\begin{aligned} dV_t &= a_1(V_t, U_t, t; \theta)dt + \sigma_1 dW_{1t} \\ dU_t &= a_2(V_t, U_t, t; \theta)dt + \sigma_2 dW_{2t} \end{aligned} \quad (1)$$

where a_1 and a_2 are the two drift functions that depend both on V_t and U_t , θ are the drift parameters, $(W_{1t})_t$ and $(W_{2t})_t$ are two independent Brownian processes and σ_1, σ_2 are the two diffusion coefficients. Here σ_1 is scalar and σ_2 is a $d-1$ dimensional vector and the multiplication $\sigma_2 dW_{2t}$ has to be understood componentwise in the last equation.

Let us denote $B_\sigma = \begin{pmatrix} \sigma_1 & 0_{1,d-1} \\ 0_{d-1,1} & \text{Diag}(\sigma_2) \end{pmatrix}$ the diffusion coefficient. We consider two classes of models (1). The first class, called *elliptic*, corresponds to an SDE with a full non degenerate diffusion coefficient i.e. the matrix $B_\sigma B_\sigma^T$ is full rank, X^T being the transposed matrix of X . The second class, called *hypoelliptic*, corresponds to an SDE with degenerate stochastic noise: the diffusion coefficient $B_\sigma B_\sigma^T$ is not invertible, for example when $\sigma_1 = 0$, but the noise is still propagated to V_t through the influence of U_t . Assumptions on the drift and diffusion coefficient structures to ensure hypoellipticity are stated in the next section.

These two specificities, partial observations and hypoelliptic properties, are of increasing importance in many applications, some examples are given below. Let us first remark that these two specificities are not of the same nature. The first is linked to the type of observations. In many examples, the system is complex and is modeled by a multi-dimensional system, while the experimentalists are only able to measure, often at discrete times, a one-dimensional signal. This increases the difficulty of estimating parameters of model (1). The second specificity is not linked to any experimental constraint, but is a mathematical way of describing the intrinsic noise of process (V_t, U_t) . It might be difficult for the modeler to know in advance if the system is elliptic rather than hypoelliptic. Unfortunately, estimation methods are often strongly different depending on the nature of the noise (see more details below), and may fail down when applied to the ‘wrong’ class of models. This is an advantage of our method which is the same in both cases.

Let us now give some examples of applications. Partially observed SDEs have been used in pharmacokinetics for modeling the concentration of a drug in the body, either in a elliptic or a hypoelliptic version (Ditlevsen et al., 2005; Cuenod et al., 2011). In systems biology, the stochastic version of the famous deterministic Lotka-Volterra model (see Lotka, 1925) which can be found in Mao et al. (2002); Meeds and Welling (2015); Graham and Storkey (2017) describes the interaction between two species, predator and prey, through a two-dimensional elliptic system. It is often possible to observe only one of the two species, leading to partial observations. In neurosciences, several stochastic systems have been proposed to model the dynamic of one single neuron. The first equation V_t corresponds to the dynamics of the membrane potential of the neuron and U_t to a recovery variable, or a synaptic conductance, that can not be measured. We can cite the synaptic-conductance based models (Paninski et al., 2010, 2012; Ditlevsen and Greenwood, 2013) or the FitzHughNagumo model (Gerstner and Kistler, 2002). These models have been proposed with stochastic noise on the synaptic conductance dynamic (U_t) only, leading to hypoelliptic SDEs (see e.g. Paninski et al., 2012, and references therein), or on both coordinates leading to elliptic SDE (Ditlevsen and

Greenwood, 2013). A last class of models is the stochastic hypoelliptic Damping Hamiltonian system where the first coordinate represents the position of a particle and the second its velocity ($d = 2$). It is natural that the (Brownian) noise appears only in the velocity coordinate, the position being defined as the deterministic infinitesimal integral of the speed. The position can be measured with precision, but the speed is not directly available. In these models, the first equation of the system reduces to $dV_t = U_t dt$.

Let us now review the estimation methods. Estimation of elliptic SDE has been widely studied. In the complete observations cases (both V_t and U_t observed), we can cite among others Bibby and Sorensen (1995); Pedersen (1995); Kessler (1997); Aït-Sahalia (2008); Durham and Gallant (2002); Sørensen (2004); Beskos et al. (2006); Jensen et al. (2012); Van der Meulen and Schauer (2016). Partial observations have been considered with several approaches. The unobserved coordinates are treated as missing data and are imputed, see for examples Elerian et al. (2001); Eraker (2001); Golightly and Wilkinson (2006, 2008); Ditlevsen and Samson (2014). Some methods propose to approximate the transition density by the Euler-Maruyama scheme and consider a Monte-Carlo approximation to impute and filter the unobserved coordinates. Therefore, they are computationally intensive. We will show that the methodology we develop is less demanding in terms of computation time than these previous methods, even if its computational cost is not negligible.

Let us now explain why the estimation of hypoelliptic systems is more difficult. Imagine that the complete observations of (V_t, U_t) are available in continuous time. Estimating θ is naturally performed through the Girsanov formula, that gives directly the likelihood (Lipster and Shiryaev, 2001), the diffusion coefficient being estimated by the quadratic variation for example. Girsanov formula requires the existence of a solution w.r.t. $u(z)$ to the system of algebraic equations $B_\sigma u(z) = a(z) - \alpha(z)$ where $\alpha(z)$ is the drift of the dominating measure (see section 7.6.4. in Lipster and Shiryaev, 2001). This solution does not exist for hypoelliptic systems. The same problem occurs for most of estimation methods developed for elliptic systems that would typically require $B_\sigma B_\sigma^T$ to be invertible. They can thus not be applied to hypoelliptic systems.

There are only few references for hypoelliptic SDEs. The stochastic Damping Hamiltonian system (with $dV_t = U_t dt$) has been the most studied. In the parametric framework, Gloter (2006); Samson and Thieullen (2012); Leon et al. (2019) propose Euler contrasts with a correction of the bias due to partial observations. Pokern et al. (2009) propose a Gibbs sampling in a bayesian approach and emphasize the bias induced by the hypoellipticity of the system when using the Euler pseudo-likelihood in the calculation of the posterior distribution. Bierkens et al. (2018) propos a new simulation method fo hypoelliptic conditional diffusions, which could be used in a Bayesian framework. Cattiaux et al. (2014b,a); Comte et al. (2017) deal with the non-parametric framework. For hypoelliptic SDEs that are more general than the stochastic Damping Hamiltonian system, we are only aware of the works of Ditlevsen and Samson (2017); Melnykova (2019). Their approach is based on a discretization scheme of order 1.5, a particle filter to approximate the unobserved coordinate and a stochastic maximization of a statistical contrast. The main drawbacks are the computational time induced by the particle filter and the fastidious mathematical calculations to exhibit the sufficient statistics of the likelihood for the considered SDE model, that are then

stochastically approximated. There is thus a need to develop new approaches, and especially estimation methods that can be applied to both elliptic and hypoelliptic systems.

In this paper, we propose an alternative which is computationally demanding but can be applied to a larger class of models, including elliptic and hypoelliptic systems and that does not require additional mathematical calculations. This strategy is based on optimal control theory. Let us recall the main question the optimal control theory aims to address. For a given dynamical system in a given initial state, which control do we have to apply to steer it to a desired behavior in an optimal way? This problem is formulated as a constrained optimization problem where the optimality is defined through the introduction of a cost function and the proposed dynamical model belongs to the constraints. Under this theory a large variety of theoretical and numerical tools have been developed to find the so-called optimal control which minimizes the proposed cost function. Recently, this theory has been advantageously used for statistical purpose. We can cite the pioneer work of Martin et al. (2001) for non-parametric estimation of B-splines. Parametric approaches have been proposed more recently by Brunel and Clairon (2015); Clairon and Brunel (2018, 2017); Iolov et al. (2017). One way of using the optimal control theory is to write an estimation criterion, typically a likelihood, as a tracking problem. Tracking problems are specific optimal control problems: the aim is to find the optimal control such that one coordinate of the system is the closest possible to a target trajectory, here the observations, on a given observation interval. The optimal control problem can be solved by the Pontryagin maximum principle (Pontryagin et al., 1962; Trelat, 2005; Sontag, 1998).

This idea has already been successfully developed by Brunel and Clairon (2015) to estimate the parameters of ordinary differential equations. It proves to be numerically efficient and stable, especially when the problem is ill-conditioned. Moreover, the consistency and the rate of convergence of the corresponding estimator have been proven (Clairon and Brunel, 2017, 2018). However, their method, based on the deterministic theory of optimal control, is restricted to the case of deterministic systems. In this paper, we extend this idea to SDE systems. To do so, we resort to the framework of the discrete optimal control theory. Theoretical and numerical results have been fully developed for linear models, this is the discrete linear-quadratic (LQ) theory, a particular case of the Pontryagin maximum principle. Indeed, this theory ensures the existence, uniqueness and gives the closed form of the solution of the control problem defining our estimation criterion. The main advantage of this theory is that it applies without any hypothesis on the diffusion coefficient B_σ . Especially, it applies also to the hypoelliptic case. This is the approach that we use for linear SDEs. When estimating parameters of a nonlinear SDE, we propose to rewrite the estimation criterion to enter this theory as well. Unfortunately, the consistency and convergence results presented for ODE models in Brunel and Clairon (2015) do not apply here and we were not able to prove a theoretical result for our estimator. Nevertheless, our method gives good estimates on the different tested models during the numerical simulations and encounter no additional difficulties for hypoelliptic models comparing to the elliptic ones. The computational time of the method is reasonable though not small.

Our estimation procedure follows three steps. First we define a criterion to estimate the drift parameter θ alone, σ being considered known. This criterion is minimized through the linear-quadratic theory. This method introduces a weight-

ing parameter w which needs to be selected. The second step consists in constructing an external criterion based on moments of the Brownian process allowing to data-select the weight w . The third step is the estimation of σ by profiling the functional used in step 2. The final estimation method is fully data-driven and there is no tuning parameter. It can be viewed as a plug-and-play methodology that can be applied to both elliptic and hypoelliptic systems.

The paper is organised as follows. Section 2 presents the models and the objectives of the estimation problem. Section 3 introduces the optimal control theory approach. Section 4 an adaptive selection of the weight w and Section 5 gives the complete procedure. Section 6 illustrates the approach on simulated examples. The paper ends with some discussions (Section 7).

2 Models and objectives

The estimation procedure that we propose, based on the discrete optimal control theory, is fully developed for linear models. We thus introduce a linear SDE. In section 3, we nonetheless present an adaptation of the estimation method to a certain kind of non-linear SDEs.

2.1 Elliptic and hypoelliptic stochastic differential equations

We consider a d -dimensional state variable $Z_t \in \mathbb{R}^d$, $d \geq 2$, defined for t in a time interval $[0, T]$. We distinguish in the following the first observed state variable denoted $V_t \in \mathbb{R}$ from the last $d - 1$ other unobserved variables denoted $U_t \in \mathbb{R}^{d-1}$. The dynamic of $(Z_t = (V_t, U_t))_{t \geq 0}$ is described by the following stochastic dynamical system:

$$\begin{pmatrix} dV_t \\ dU_t \end{pmatrix} = A_\theta(t) \begin{pmatrix} V_t \\ U_t \end{pmatrix} dt + B_\sigma dW_t \quad (2)$$

with initial conditions $Z_0 = (V_0, U_0)$ possibly unknown and where W_t is a m dimensional Brownian motion. The drift is assumed linear with respect to $Z_t = (V_t, U_t)$. The $d \times d$ -matrix $A_\theta(t)$ depends on the unknown parameter vector θ and may be a function of time t . The $d \times m$ -matrix B_σ is called the diffusion coefficient and depends on an unknown parameter vector σ . We consider two cases:

- **Elliptic SDE:** $m = d$ and the $d \times d$ matrix B_σ is not singular:

$$\det(B_\sigma B_\sigma^T) > 0 \quad (3)$$

- **Hypoelliptic SDE:** $m < d$ and the matrix B_σ is singular:

$$B_\sigma = \begin{pmatrix} 0_m \\ b_\sigma \end{pmatrix}, \quad (4)$$

with 0_m the m -dimensional row vector of zeros and b_σ a $(d - 1) \times m$ -matrix such that

$$\det(b_\sigma b_\sigma^T) > 0.$$

A noticeable feature in the hypoelliptic SDE is that the equation ruling V_t does not contain a stochastic part. The matrix b_σ models the way the stochastic disturbance acts on the unobserved variables U_t of the system and indirectly on V_t through U_t . In the hypoelliptic case, we consider the following assumption:

- (H1) Let $A_{1,j}(t)$ denote the j th element of the first column of $A_\theta(t)$ and $B_{,j}$ the j th column of B_σ . For any t , there exists at least one $j = 2, \dots, d$ such that

$$A_{1,j}(t)B_{,j} \neq 0.$$

Under assumption (H1), the noise is propagated also to the first coordinate V_t . This property that the noise generates the entire space \mathbb{R}^d is a characteristic of the hypoelliptic property see e.g. Mattingly et al. (2002). Note that the standard assumptions to ensure the hypoellipticity are different from (H1) see e.g. Samson and Thieullen (2012), but it reduces to (H1) for a linear system.

2.2 Objectives and issues

The state variables $Z_t = (V_t, U_t)$ are split in two because we observe only the first one-dimensional state variable V_t . We denote $t \mapsto Y(t)$, the realization of V_t . We assume that Y is discretely observed on the interval $[0, T]$ at times $0 = t_0 < \dots < t_n = T$ without measurement error and denote $Y = (Y_0, \dots, Y_n)$ these observations. We thus have $Y_i = CZ_{t_i}$, with C a $1 \times d$ -matrix $C = (1, 0_{d-1})^T$.

The aim of the paper is to estimate the unknown parameters θ and the initial condition Z_0 (if unknown) of model (2) in the elliptic and hypoelliptic cases using the discrete observations (Y_0, \dots, Y_n) . We will distinguish two cases: 1/ σ is known and only θ is estimated; 2/ both θ and σ are estimated.

Before introducing our approach, let us explain why the estimation problem is difficult. To estimate θ and σ of an elliptic, partially observed SDE, one can start with the discretization of the diffusion (2). The Euler–Maruyama discretization scheme at time (t_1, \dots, t_n) is defined as follows for $i = 0, \dots, n-1$ and $Z_i = (V_i, U_i)$:

$$Z_{i+1} = Z_i + \Delta_i A_\theta(t_i) Z_i + B_\sigma \eta_i =: \mathbf{A}_\theta(t_i) Z_i + B_\sigma \eta_i \quad (5)$$

where Z_i, U_i, V_i stand for $Z_{t_i}, U_{t_i}, V_{t_i}$, $\Delta_i = t_{i+1} - t_i$, $\mathbf{A}_\theta(t_i) = I_d + \Delta_i A_\theta(t_i)$ and the η_i are independent variables distributed as $\mathcal{N}(0, \Delta_i)$. When the system is elliptic, that is when $B_\sigma B_\sigma^T$ is invertible, minus twice the log-likelihood of this discretized process, assuming both V_t and U_t discretely observed, is then

$$L_{Euler}((V_i, U_i)_{i=0, \dots, n}, \theta, \sigma) = \sum_{i=0}^{n-1} (Z_{i+1} - \mathbf{A}_\theta(t_i) Z_i)^T (B_\sigma B_\sigma^T)^{-1} (Z_{i+1} - \mathbf{A}_\theta(t_i) Z_i) + n \log(\det(B_\sigma B_\sigma^T)). \quad (6)$$

For partial observations, the log-likelihood (6) has to be integrated with respect to (U_0, \dots, U_n) . The estimator is thus defined as

$$\arg \min_{\theta, \sigma} \int L_{Euler}((V_i, U_i)_{i=0, \dots, n}, \theta, \sigma) d(U_0, \dots, U_n).$$

This integral can be viewed as a filtering problem: the unobserved trajectory (U_0, \dots, U_n) is filtered with respect to observations (V_0, \dots, V_n) . Kalman filter can be applied to linear SDEs. For non-linear SDEs, the extended Kalman filter can be used, or a particle albeit at the price of increased computing time. For hypoelliptic SDE, criterion (6) can not even be computed, $B_\sigma B_\sigma^T$ being not invertible.

Therefore, there is a need of alternatives. In this paper, we take advantage of the optimal control theory to filter the unobserved coordinate or more precisely to find a surrogate value for the (unobserved) realization of the Brownian motion, under the form of a control sequence that drives the trajectory. This sequence allows to define an estimation criterion that applies also when the stochastic noise is degenerate. The estimation procedure is presented in Sections 3-5.

3 Estimation of θ via optimal control theory

In this section, we assume σ known. First we write the statistical problem into an optimal control one. The procedure is the same whether B_σ is singular or not. The optimal control problem is then solved for linear models by using the linear quadratic theory (Section 3.2). We end this section by extending the estimation procedure to a particular but meaningful subset of nonlinear SDEs (Section 3.3).

3.1 Optimal control problem and definition of $\hat{\theta}$

The control problem introduces a *control sequence* such that data (Y_0, \dots, Y_n) are close to a solution of the discretized model (5), close meaning of the order of the Euler scheme error Δ (Bally and Talay, 1996). In the context of SDEs, the natural control sequence is the increment of the Brownian motion. Let us introduce $u_i = W_{i+1} - W_i$, $i = 0, \dots, n-1$, which will play the role of the control value at time t_i . Note that $u_i \in \mathbb{R}^m$, $i = 0, \dots, n-1$ with a dimension m that can be different from d , thus allowing the noise to be degenerate.

The discretized model (5) can be reformulated under the form of a *discrete controlled system*:

$$\begin{pmatrix} V_{i+1} \\ U_{i+1} \end{pmatrix} = \mathbf{A}_\theta(t_i) \begin{pmatrix} V_i \\ U_i \end{pmatrix} + B_\sigma u_i, \quad (V_{t_0}, U_{t_0}) = Z_0. \quad (7)$$

We denote u the vector of discrete values taken by the control: $u = (u_0, \dots, u_{n-1})$; $Z_{i,\theta,\sigma,u} = (V_{i,\theta,\sigma,u}, U_{i,\theta,\sigma,u})$, the solution of (7) corresponding to the given θ , σ and u .

When model (2) is true, there exists one realization of the Brownian motion such that data (Y_0, \dots, Y_n) are a sample of model (2). Control sequence u can be viewed as this specific realization of the Brownian motion. The objective is to infer sequence u by extracting knowledge from data Y . We define a cost function $C(u, Y; \theta, \sigma)$ such that the optimum in u corresponds to this realization of the Brownian motion.

A natural cost function is the conditional distribution (also called posterior distribution) $P(u|Y; \theta, \sigma)$ of u knowing data Y . We might select filtering methods

that consist in computing this conditional distribution but that are computationally demanding. An alternative is to compute the maximum of the conditional distribution $P(u|Y; \theta, \sigma)$, that can also be called Maximum A Posteriori (MAP) by analogy with the Bayesian settings. Maximizing $P(u|Y; \theta, \sigma)$ could be numerically difficult, the function being likely not concave. The optimization problem can be stabilized by regularizing it using a weight.

Let us introduce the MAP and then the weighted cost function. For a fixed value of the parameters (θ, σ) , the MAP would be defined as follows:

$$\begin{aligned} \hat{u}_{MAP} &= \arg \max_u P(u|Y; \theta, \sigma) \\ &= \arg \max_u \left(\frac{P(Y|u; \theta, \sigma)P(u; \theta, \sigma)}{P(Y; \theta, \sigma)} \right) \end{aligned}$$

where $P(Y|u; \theta, \sigma)$ is the density of the data given the u , $P(u; \theta, \sigma)$ is the density of the Brownian motion and the likelihood $P(Y; \theta, \sigma)$ does not depend on u . Let us now detail the two first terms. Conditioning on u means that $V_{i, \theta, \sigma, u}$ is known. Thus conditional density $P(Y|u; \theta, \sigma)$ is the distribution of the difference between Y_i and $V_{i, \theta, \sigma, u}$, which is the distribution of the error of discretization of the SDE. This error has a variance of order the time step Δ (Bally and Talay, 1996). Second term $P(u; \theta, \sigma)$ is the density of a discretized Brownian motion, which has a variance Δ . Therefore, the MAP reduces to

$$\begin{aligned} \hat{u}_{MAP} &= \arg \max_u \left(- \sum_{i=0}^n \frac{(V_{i, \theta, \sigma, u} - Y_i)^2}{\Delta_i} - \sum_{i=0}^{n-1} \frac{u_i^T u_i}{\Delta_i} \right) \\ &= \arg \min_u C(u, Y; \theta, \sigma) \end{aligned} \quad (8)$$

with $C(u, Y; \theta, \sigma, Z_0) = \sum_{i=0}^n \frac{(V_{i, \theta, \sigma, u} - Y_i)^2}{\Delta_i} + \sum_{i=0}^{n-1} \frac{u_i^T u_i}{\Delta_i}$.

Optimisation of the MAP (8) reveals to be intractable in practice. By replacing the scaling $1/\Delta_i$ of the first term by a weight $w > 0$, we obtain a tractable optimisation problem which falls into the linear quadratic theory framework, where the new cost function is:

$$C_w(u, Y; \theta, \sigma, Z_0) = w (V_{n, \theta, \sigma, u} - Y_n)^2 + \sum_{i=0}^{n-1} \left(w (V_{i, \theta, \sigma, u} - Y_i)^2 + \frac{u_i^T u_i}{\Delta_i} \right).$$

In the last equation, we exhibit the last observation Y_n which plays a specific role in control problems. The weight w has to be chosen by the user. In Section 4 a procedure is proposed to select w adaptively from data Y . For a given weight w , the *best control* is the sequence u that minimizes $C_w(u, Y; \theta, \sigma, Z_0)$ under the constraint of model (7). It is defined as the solution of the following optimal control problem:

$$\begin{aligned} &\text{Minimize in } u: C_w(u, Y; \theta, \sigma, Z_0) = \\ &\quad w (V_{n, \theta, \sigma, u} - Y_n)^2 + \sum_{i=0}^{n-1} \left(w (V_{i, \theta, \sigma, u} - Y_i)^2 + \frac{1}{\Delta_i} u_i^T u_i \right) \\ &\text{Subject to: } \begin{cases} \begin{pmatrix} V_{i+1, \theta, \sigma, u} \\ U_{i+1, \theta, \sigma, u} \end{pmatrix} = \mathbf{A}_\theta(t_i) \begin{pmatrix} V_{i, \theta, \sigma, u} \\ U_{i, \theta, \sigma, u} \end{pmatrix} + B_\sigma u_i \\ (U_{0, \theta, \sigma, u} \ V_{0, \theta, \sigma, u}) = Z_0. \end{cases} \end{aligned} \quad (9)$$

This is called the *optimal control* and denoted $\bar{u}_{\theta,\sigma}$. For a fixed value σ and given the discrete observations (Y_0, \dots, Y_n) , we then define the estimator of θ as:

$$\hat{\theta}_w(\sigma) = \arg \min_{\theta} S_w(Y; \theta, \sigma) \quad (10)$$

where

$$S_w(Y; \theta, \sigma) := \min_u C_w(u, Y; \theta, \sigma, Z_0) = C_w(\bar{u}_{\theta,\sigma}, Y; \theta, \sigma, Z_0)$$

is the profiled cost C_w over the set of possible sequences u , or if the initial conditions Z_0 is unknown:

$$S_w(Y; \theta, \sigma) := \min_{Z_0} \{ \min_u C_w(u, Y; \theta, \sigma, Z_0) \} = \min_{Z_0} \{ C_w(\bar{u}_{\theta,\sigma}, Y; \theta, \sigma, Z_0) \}$$

the profiled cost C_w over u then over the possible initial condition values.

The computations of S_w and $\bar{u}_{\theta,\sigma}$ require solving optimal control problem (9). We have to prove the existence and uniqueness of the solution of problem (9), and that this unique solution is numerically computable. These two results are given by the linear quadratic theory Trelat (2005); Sontag (1998) which is exposed below. Interestingly, this theory also ensures profiling on Z_0 is not computationally demanding.

Remark 1 If, instead of solving problem (9), we set $u = 0$, we end up with an estimator somewhat similar to the Trajectory Fitting Estimator (TFE) (Kutoyants (1991); Dietz (2001)). However, the addition of the penalization term $\frac{1}{\Delta_i} u_i^T u_i$ in C_w and its optimization over u allows us to be able to deal with the partially observed framework for both elliptic and hypoelliptic models which is not the case for TFE.

3.2 Linear quadratic theory

Linear quadratic theory is derived from the Pontryagin maximum principle (Pontryagin et al., 1962) for linear models. For a given (θ, σ) , it ensures the existence and uniqueness of $\bar{u}_{\theta,\sigma}$, the solution of (9), and that $\bar{u}_{\theta,\sigma}$ and $S_w(Y; \theta, \sigma)$ can be computed via a backward finite difference equation, called the Riccati equation.

Let us be more precise. For a given (θ, σ) , and a given weight w , let us denote the following matrices, for $i = 0, \dots, n-1$:

$$A_{i,\theta} = \begin{pmatrix} \mathbf{A}_{\theta}(t_i) & 0_d^T \\ 0_d & 1 \end{pmatrix} \in \mathbb{R}^{d+1} \times \mathbb{R}^{d+1}, \quad Q_i = \begin{pmatrix} C^T C & -C^T Y_i \\ -C Y_i & Y_i^2 \end{pmatrix} w \in \mathbb{R}^{d+1} \times \mathbb{R}^{d+1},$$

and

$$B_{\sigma}^0 = \begin{pmatrix} B_{\sigma} \\ 0_m \end{pmatrix} \in \mathbb{R}^{d+1} \times \mathbb{R}^m.$$

Set $R_{n,\theta,\sigma} = Q_n$ and let $R_{i,\theta,\sigma}$ be the sequence of positive symmetric matrices, solution of the discrete backward Riccati equation, for $i = n-1, \dots, 0$:

$$R_{i,\theta,\sigma} = Q_i + A_{i,\theta}^T R_{i+1,\theta,\sigma} \left(I_{d+1} - B_{\sigma}^0 \left(\frac{1}{\Delta_i} I_m + B_{\sigma}^{0T} R_{i+1,\theta,\sigma} B_{\sigma}^0 \right)^{-1} B_{\sigma}^{0T} R_{i+1,\theta,\sigma} \right) A_{i,\theta}. \quad (11)$$

We can establish the following theorem.

Theorem 1 *The optimal control problem (9) has a unique solution $\bar{u}_{\theta,\sigma} = (\bar{u}_{0,\theta,\sigma}, \dots, \bar{u}_{n-1,\theta,\sigma})$, the value $\bar{u}_{i,\theta,\sigma} \in \mathbb{R}^m$ at time t_i given by:*

$$\bar{u}_{i,\theta,\sigma} = - \left(\frac{1}{\Delta_i} I_m + B_\sigma^{0T} R_{i+1,\theta,\sigma} B_\sigma^0 \right)^{-1} B_\sigma^{0T} R_{i+1,\theta,\sigma} A_{i,\theta} \begin{pmatrix} V_{i,\theta,\sigma,\bar{u}_{\theta,\sigma}} \\ U_{i,\theta,\sigma,\bar{u}_{\theta,\sigma}} \\ 1 \end{pmatrix}. \quad (12)$$

Moreover, the minimum value of C_w is equal to:

$$\min_{u \in \mathbb{R}^{m \times (n-1)}} C_w(u, Y; \theta, \sigma, Z_0) = (Z_0^T, 1) R_{0,\theta,\sigma} (Z_0^T, 1)^T.$$

When Z_0 is known,

$$S_w(Y; \theta, \sigma) = (Z_0^T, 1) R_{0,\theta,\sigma} (Z_0^T, 1)^T.$$

When Z_0 is unknown,

$$S_w(Y; \theta, \sigma) = -h_{\theta,\sigma}^T E_{\theta,\sigma}^{-1} h_{\theta,\sigma} + \alpha_{\theta,\sigma}$$

where $h_{\theta,\sigma}$, $E_{\theta,\sigma}$ and $\alpha_{\theta,\sigma}$ are the components of the symmetric matrix $R_{0,\theta,\sigma}$ i.e.

$$R_{0,\theta,\sigma} = \begin{pmatrix} E_{\theta,\sigma} & h_{\theta,\sigma} \\ h_{\theta,\sigma}^T & \alpha_{\theta,\sigma} \end{pmatrix}.$$

The fundamental theorem used to derive these expressions for S_w is recalled in appendix. Theorem 1 provides a closed form expression for S_w that depends on matrix $R_{0,\theta,\sigma}$ obtained by solving Riccati equation (11).

Let us now highlight a noticeable advantage of our approach. Theorem 1 holds for elliptic ($m = d$) SDEs as well as for hypoelliptic ($m < d$) SDEs. Indeed C_w is strictly convex with respect to the sequence u due to the term $u_i^T u_i$, no matter the dimension of the noise. Applications of linear quadratic theory does not require the matrix $B_\sigma B_\sigma^T$ to be full-conditioned. It can thus be applied to both elliptic and hypoelliptic cases without any additional calculations.

Note that the cost function can be easily profiled on initial condition Z_0 . This is due to the quadratic nature of $\min_{u \in \mathbb{R}^{m \times (n-1)}} C_w(u, Y; \theta, \sigma, Z_0)$ with respect to Z_0 . Thus, the construction of S_w for unknown initial conditions is similar to the deterministic Kalman Filter state estimator derivation (see Sontag (1998)).

The computation of the cost function C_w is in the framework of the discrete LQ theory. This ensures the well posedness nature of the optimization problem (uniqueness, existence and continuity of the solution w.r.t problem parameters) defining the cost S_w as well as an efficient method to compute it no matter the elliptic nature of the SDE model.

From equation (11), we can easily establish by induction that for every i , $R_{i,\theta,\sigma}$ has the same degree of regularity with respect to model parameters (θ, σ) than A_θ and B_σ . This in turn applies as well to S_w . So, if we assume θ belongs to a bounded subset Θ and $\theta \mapsto A_\theta$ is continuous on Θ , $\hat{\theta}$ is defined as the minimizer of a continuous functional S_w on a bounded subset which ensures its existence. This also allows the use of classic optimization methods during the numerical analysis in Section 6.

3.3 Extension to some nonlinear models

To extend the method, we resort to the pseudo-linear formalism presented for example in Cimen and Banks (2004a,b). A nonlinear SDE can always be written through a pseudo-linear representation:

$$\begin{pmatrix} dV_t \\ dU_t \end{pmatrix} = A_\theta(V_t, U_t, t) \begin{pmatrix} V_t \\ U_t \end{pmatrix} dt + B_\sigma dW_t \quad (13)$$

even if this representation is non unique. Here, we focus on the case where a pseudo-linear matrix A_θ which only depends on the observed coordinate V_t can be found i.e. $A_\theta(V_t, U_t, t) = A_\theta(V_t, t)$. In this case the Euler–Maruyama discretization scheme becomes

$$Z_{i+1} = Z_i + \Delta_i A_\theta(Z_i, t_i) Z_i + B_\sigma u_i = \mathbf{A}_\theta(t_i) Z_i + B_\sigma u_i \quad (14)$$

with $\mathbf{A}_\theta(t_i) := I_d + \Delta_i A_\theta(Y_i, t_i)$ and we are able to use Theorem 1 to compute the profiled cost S_w . The next step is the selection of the weight w involved in the definitions of the optimal control problem (9) and the estimator $\hat{\theta}_w(\sigma)$. This is explained in the next section.

4 Adaptive w selection

The choice of the weight w is critical to ensure the numerical stability of the procedure. We propose to select it adaptively by minimizing an external functional criterion $G(Y; \sigma, w, \theta)$. We propose two possible choices for G :

1. $G^{(1)}$, a criteria quantifying the consistency between the estimated optimal control $\bar{u}_{\theta, \sigma}$ and a moment of the quadratic growth of Brownian motion increments,
2. $G^{(2)}$, a contrast function based on a state variable estimator obtained as a by-product of $\bar{u}_{\theta, \sigma}$ estimation.

For the two functionals $G^{(1)}$ and $G^{(2)}$, the selected weight w and the final estimator of θ are computed through the following steps.

Definition 1 (Estimator $\hat{\theta}(\sigma)$, σ known) Assume σ known. The estimation of θ is defined as follows:

1. For any w , compute the estimator $\hat{\theta}_w(\sigma) = \arg \min_{\theta} S_w(Y; \theta, \sigma)$.
2. Minimize the functional $G = G^{(1)}$ or $G = G^{(2)}$ w.r.t. w , using the plugged value $\hat{\theta}_w(\sigma)$, and define

$$\hat{w} := \arg \min_w G(Y; \sigma, w, \hat{\theta}_w(\sigma)) \quad (15)$$

3. Define the final estimator of θ as

$$\hat{\theta}(\sigma) = \hat{\theta}_{\hat{w}}(\sigma) \quad (16)$$

The two functionals are presented in the next sections.

4.1 w selection via a quadratic growth moment condition

We present the first functional $G^{(1)}$ derived as a constraint of the distribution of the control. The control $\bar{u}_{\theta,\sigma}$ mimics the increments of the Brownian motion. Thus, it is natural to impose that the values $\frac{\bar{u}_{i,\theta,\sigma}}{\sqrt{\Delta_i}}$ are independent and distributed with $\mathcal{N}(0, I_m)$. We can then derive constraints related to this distribution. Let us denote $\bar{u}_{i,j,\theta,\sigma}$ the j -th component of the m dimensional Brownian motion at time t_i , for $j \in \llbracket 1, m \rrbracket$. The law of large numbers implies that, almost surely,

$$\frac{1}{n} \sum_{i=0}^{n-1} \frac{\bar{u}_{i,j,\theta,\sigma}^2}{\Delta_i} \longrightarrow 1.$$

We thus propose the following functional to select w :

$$G^{(1)}(Y; \sigma, w, \theta) = \sum_{j=1}^m \left(\sum_{i=0}^{n-1} \frac{\bar{u}_{i,j,\theta,\sigma}^2}{\Delta_i} - n \right)^2. \quad (17)$$

Here, the dependency of $G^{(1)}(\cdot)$ with respect to w is made through the optimal control sequence $\bar{u}_{\theta,\sigma}$. This functional gives importance to the control sequence property.

4.2 w selection via a data fidelity criterion

The second functional is driven by the idea of giving importance to the data fidelity term, that is to minimize the distance between model solution $Z_{i+1,\theta,\sigma,u}$ and data Y_i .

We can show by recurrence that for $i = 0, \dots, n-1$ and any integer $0 \leq l \leq i-1$:

$$Z_{i+1,\theta,\sigma,u} = \left(\prod_{j=0}^l \mathbf{A}_{\theta}(t_{i-j}) \right) Z_{i-l,\theta,\sigma,u} + \sum_{k=0}^l \left(\prod_{j=0}^{k-1} \mathbf{A}_{\theta}(t_{i-j}) \right) B_{\sigma} u_{i-k}$$

where $\prod_{j=0}^{k-1} \mathbf{A}_{\theta}(t_{i-j}) = 1$ when $k = 0$. Thus, by multiplying by C , we can explicitly link the observations Y and the control:

$$Y_{i+1} - C \left(\prod_{j=0}^l \mathbf{A}_{\theta}(t_{i-j}) \right) Z_{i-l,\theta,\sigma,u} = C \sum_{k=0}^l \sqrt{\Delta_{i-k}} \left(\prod_{j=0}^{k-1} \mathbf{A}_{\theta}(t_{i-j}) \right) B_{\sigma} \frac{u_{i-k}}{\sqrt{\Delta_{i-k}}}.$$

The functional $G^{(2)}$ is based on the distribution of $Y_{i+1} - C \prod_{j=0}^l \mathbf{A}_{\theta}(t_{i-j}) Z_{i-l,\theta,\sigma,u}$. To ease the reading, let us introduce the m -vector

$$\Gamma_{\theta,\sigma}(k, t_i) = C \sqrt{\Delta_{i-k}} \left(\prod_{s=0}^{k-1} \mathbf{A}_{\theta}(t_{i-s}) \right) B_{\sigma}$$

and the sequence of scalars, for $i = 0, \dots, n-1$:

$$X_{i,\theta,\sigma,u} = Y_{i+1} - C \prod_{j=0}^l \mathbf{A}_{\theta}(t_{i-j}) Z_{i-l,\theta,\sigma,u}. \quad (18)$$

We can derive the law followed by $X_{i,\theta,\sigma,u} = \sum_{k=0}^l \Gamma_{\theta,\sigma}(k, t_i) \frac{u_{i-k}}{\sqrt{\Delta_{i-k}}}$:

$$X_{i,\theta,\sigma,u} \sim N(0, \gamma_{l,\theta,\sigma}^2(t_i)) \quad \text{with} \quad \gamma_{l,\theta,\sigma}^2(t_i) = \sum_{k=0}^l \Gamma_{\theta,\sigma}(k, t_i) \Gamma_{\theta,\sigma}(k, t_i)^T.$$

Let us choose l such that $\gamma_{l,\theta,\sigma}^2(t_i)$ is non-zero. Let us explain the intuition why choosing a sufficiently large index l yields $\gamma_{l,\theta,\sigma}^2(t_i) > 0$. The idea is to give time to the stochastic elements u_i to diffuse. In other words, after “enough time” (characterized by l), the elements u_i are able to perturb the observations Y_{i+l+1} and σ becomes univocally identified.

Now set $X_{i,\theta,\sigma,u}^l := \frac{X_{i,\theta,\sigma,u}}{\gamma_{l,\theta,\sigma}(t_i)}$, then $\mathbb{E} \left[(X_{i,\theta,\sigma,u}^l)^2 \right] = 1$. Note that $X_{i,\theta,\sigma,u}^l$ depends on (u_{i-l}, \dots, u_i) and $X_{i+l+1,\theta,\sigma,u}^l$ depends on $(u_{i+1}, \dots, u_{i+l+1})$ thus they are independent. For the i.i.d sequence of $\left\{ X_{(l+1)i+1,\theta,\sigma,\bar{u}_{\theta,\sigma}}^l \right\}_{0 \leq i \leq \frac{n-1}{l+1}}$ of size $L := \left\lceil \frac{n+1}{l+1} \right\rceil$, the law of large number implies that, almost surely,

$$\frac{1}{L} \sum_{i=0}^{L-1} \left(X_{(l+1)i+1,\theta,\sigma,u}^l \right)^2 \longrightarrow 1.$$

We can now define the functional $G^{(2)}$ as:

$$G^{(2)}(Y; \sigma, w, \theta) = \left(\sum_{i=0}^{L-1} \left(X_{(l+1)i+1,\theta,\sigma,\bar{u}_{\theta,\sigma}}^l \right)^2 - L \right)^2. \quad (19)$$

In practice, we choose the smallest value l which ensures $\gamma_{l,\theta,\sigma}^2(t_i) > 0$ in order to use the largest sequence $\left\{ X_{(l+1)i+1,\theta,\sigma,\bar{u}_{\theta,\sigma}}^l \right\}_{0 \leq i \leq \frac{n-1}{l+1}}$. For elliptic SDEs, it is sufficient to take $l = 0$. Indeed $\Gamma_{\theta,\sigma}(0, t_i) = C\sqrt{\Delta_i}B_\sigma \neq 0_m$ since B_σ is of full rank. For hypoelliptic SDEs, we have $\Gamma_{\theta,\sigma}(0, t_i) = C\sqrt{\Delta_i}B_\sigma = 0_m$. With $l = 1$, we have $\gamma_{1,\theta,\sigma}^2(t_i) = \Gamma_{\theta,\sigma}(1, t_i) \Gamma_{\theta,\sigma}(1, t_i)^T = \Delta_{i-1} C \mathbf{A}_\theta(t_i) B_\sigma B_\sigma^T \mathbf{A}_\theta(t_i)^T C^T$ and assumption (H1) implies

$$\gamma_{1,\theta,\sigma}^2(t_i) > 0$$

The computations of $\gamma_{0,\theta,\sigma}^2(t_i)$ and $\gamma_{1,\theta,\sigma}^2(t_i)$ are detailed in Section 6 for some examples of SDEs.

5 Estimation of σ^2

We now consider the estimation of σ^2 . The quadratic variation of the stochastic process divided by T provides an estimator of σ^2 . It can not be applied with partial observations because only the first coordinate is observed. We therefore propose to minimize an estimation criterion. It turns out that the two functionals $G^{(1)}$ and $G^{(2)}$ presented in Section 4 can be used to estimate σ^2 . It might seem counter-intuitive that the same criterion could be used to estimate a weight parameter and a diffusion coefficient. However, let us recall that the two functionals have been constructed as constraints of the optimal control problem. These constraints are

not specifically linked to the weight parameter but to the model itself. They can thus be judiciously used to estimate also the diffusion coefficient B_σ .

Our proposal is a nested procedure that provides the final estimation of θ , σ^2 and the data-driven selection of the weight w .

Definition 2 (Estimators $\widehat{\sigma^2}$ and $\widehat{\theta}$) The estimation of (θ, σ^2) is defined as follows:

1. For any w and σ^2 , compute $\widehat{\theta}_w(\sigma^2) = \arg \min_{\theta} S_w(Y; \theta, \sigma)$.
2. Minimize the functional $G = G^{(1)}$ or $G = G^{(2)}$ w.r.t. σ^2 , using the plugged value $\widehat{\theta}_w(\sigma^2)$, and define

$$\widehat{\sigma^2}(w) := \arg \min_{\sigma^2} G(Y; \sigma, w, \widehat{\theta}_w(\sigma^2)). \quad (20)$$

3. Minimize the functional $G = G^{(1)}$ or $G = G^{(2)}$ with respect to w , using the plugged value $\widehat{\theta}_w(\widehat{\sigma^2}(w))$, and define

$$\widehat{w} := \arg \min_w \left\{ G(Y; \widehat{\sigma^2}(w), w, \widehat{\theta}_w(\widehat{\sigma^2}(w))) \right\}. \quad (21)$$

4. Define the final estimators of θ and σ^2 as

$$\begin{aligned} \widehat{\sigma^2} &:= \widehat{\sigma^2}(\widehat{w}) \\ \widehat{\theta} &:= \widehat{\theta}_{\widehat{w}}(\widehat{\sigma^2}). \end{aligned} \quad (22)$$

6 Simulation study

A simulation study is conducted on elliptic (Section 6.1) and hypoelliptic (Section 6.2) SDEs.

We examine the estimation procedures given in definition 1 when σ is known and definition 2 when σ is estimated. When σ is known, the two estimators denoted $\widehat{\theta}^{(1)}(\sigma)$ ($\widehat{\theta}^{(2)}(\sigma)$, respectively) defined by (16) and the adapted weight w selected by the functional $G = G^{(1)}$ ($G = G^{(2)}$, respectively) are computed. When σ is also estimated, the two estimators denoted $\widehat{\theta}^{(1)}$ and $\widehat{\sigma}^{(1)}$ ($\widehat{\theta}^{(2)}$ and $\widehat{\sigma}^{(2)}$, respectively) from definition 2 using $G = G^{(1)}$ ($G = G^{(2)}$, respectively) are computed.

For the sake of comparison, we also estimate parameters by maximum likelihood approximation via particle filter methods. For elliptic models, we rely on the iterative filtering algorithm proposed by Ionides et al. (2006) implemented in the R package 'pomp' (King et al., 2016). This method embeds a parameter perturbation step. It is thought to be computationally faster than the direct approach of plugging a likelihood estimate from a particle filter directly into a stochastic optimization algorithm (Ionides et al., 2011). The algorithm is tuned as suggested in their paper. The corresponding estimators are denoted with the upper symbol I . For the hypoelliptic case, we are only aware of the method developed by Ditlevsen and Samson (2017) based on the 1.5 order discretization scheme of the SDE and which uses a particle filter coupled to a stochastic approximation EM algorithm. The corresponding estimators are denoted with the upper symbol DS .

For each model, 100 trajectories are simulated with a discretization scheme with $n_s = 10^5$ equidistant points from which we subsample $n = 10^3$ equispaced

points for parameter estimation purpose. For the elliptic case, we use an Euler-Maruyama scheme and for hypoelliptic models, a 1.5 order scheme to propagate the noise into the observed state variable Kloeden and Platen (1992).

The mean and standard error (SE) of the estimators are reported as well as the mean computational time. Our method was coded with Matlab, the required time was computed by using the function 'cputime'. Optimization required for inner and outer criteria minimization were both conducted using the function 'fminsearch' (a derivative free simplex search method). The filtering based methods was implemented in R. The simulations have been run on a personal Dell machine using one core Intel(R) Core (TM) i7-4790S CPU @ 3.20Ghz.

Three models are used: the Lotka-Volterra model, the FitzHugh-Nagumo model and the Harmonic Oscillator (HO). The two firsts are non linear and HO is linear. For all cases, we profile on initial conditions.

6.1 Simulation study for elliptic systems

Two models are studied in their elliptic version, the Lotka-Volterra system and the FitzHugh Nagumo model. The Iterated Filtering requires the setting of four hyperparameters, the number of iterations M , here $M = 200$, the number of particles P required by the embedded sequential Monte-Carlo algorithm estimating the likelihood, here $P = 1500$, the cooling factor a , here $a = 0.8$ and the perturbation scales σ_p which is the standard deviation of the perturbed random walk, here we set $\sigma_p = 0.2$. In contrast, we recall our method only involves one weighting parameter w for which an adaptive selection method has been proposed. For each model, we precise an interval $[w_{min}, w_{max}]$ where the value \hat{w} will be searched by using a simple dichotomy method.

6.1.1 Elliptic Lotka-Volterra process

Stochastic versions of Lotka-Volterra model (see Lotka, 1925) have been proposed (Meeds and Welling, 2015; Graham and Storkey, 2017; Mao et al., 2002). It is a predator-prey model that describes the dynamics of two interacting populations. The preys are assumed to have an unlimited food supply. Let V_t and U_t denote the number of predators and preys at time t , respectively. The dynamics of V_t and U_t are described as:

$$\begin{aligned} dV_t &= (-\theta_1 V_t + \theta_2 V_t U_t) dt + \sigma dW_{1,t} \\ dU_t &= (\theta_3 U_t - \theta_4 V_t U_t) dt + \sigma dW_{2,t} \\ (V_0, U_0) &= (0.5, 1) \end{aligned} \quad (23)$$

where θ_1 is the death rate of the predator, θ_2 is the growth rate of the predator population, θ_3 is the exponential growth of the prey, θ_4 is the rate of predation upon the prey which reflects the interaction between the two species, $(W_{1,t})_{t \geq 0}$ and $(W_{2,t})_{t \geq 0}$ are two independent Brownian motions, and σ is the diffusion coefficient assumed to be the same for both coordinates. The system is thus elliptic.

We simulate the trajectories on the interval $[0, 10]$. Parameters are set to $(\theta_1, \theta_2, \theta_3, \theta_4) = (0.5, 0.5, 0.5, 0.5)$ and $\sigma = 0.2$. An example of trajectory for the model (23) is presented in Figure 1.




Fig. 1 Lotka-Volterra model: simulated dynamics along time of the number of predators (green plain line) and the number of preys (blue dotted line).

The objective is to estimate these parameters from discrete observations of the predator population only. For the sake of parametric identifiability, θ_2 and θ_3 are considered known. The system (23) is non linear due to the interaction term between predators and preys. Thus, we need to use the nonlinear extension approach exposed in Section 3.3. Matrix

$$A_\theta(V_t, t) = \begin{pmatrix} -\theta_1 & \theta_2 V_t \\ 0 & \theta_3 - \theta_4 V_t \end{pmatrix} \quad (24)$$

is used in equation (13) to proceed to parametric estimation as in the linear case. We have $B_\sigma = \begin{pmatrix} \sigma & 0 \\ 0 & \sigma \end{pmatrix}$. Note that to ensure that $G^{(2)}$ is well defined, it is sufficient to take $l = 0$, as $\gamma_{0,\theta,\sigma}^2(t_i) = \Delta_i \sigma^2 > 0$. We set $[w_{min}, w_{max}] = [10, 10^3]$.

Results are presented in Table 1. The two weight selection procedures with σ known give similar results. Interestingly, they give less biased estimators than the one proposed in Ionides et al. (2006) and at a very reasonable computational cost. When σ is also estimated, $G^{(1)}$ and $G^{(2)}$ minimization give different estimators for the drift and different computational time. However, they both give two accurate estimators of σ comparing to $\hat{\sigma}^I$. In particular, $\hat{\theta}^{(2)}$ and $\hat{\sigma}^{(2)}$ are less biased estimators than $\hat{\theta}^I$ and $\hat{\sigma}^I$ and are obtained much more quickly. One can argue that other hyperparameter values would have lead $\hat{\sigma}^I$ to be more accurate. To investigate this, we plot in figure 2 the graphs of mean square error with respect to varying cooling factor a and perturbation scales σ_p , the other hyperparameters being set to the value given in the previous subsection. The chosen value for a does not have a huge impact on estimation and the selected σ_p gives accurate estimation comparing to the other values. Overall the selected hyperparameters (a, σ_p) do not undermine the accuracy of the method.

6.1.2 Elliptic FitzHugh-Nagumo

The FitzHugh-Nagumo model (FHN) describes the dynamic of an excitable neuron, modeling the characteristic spikes of the neuron (FitzHugh, 1961; Nagumo




Fig. 2 Lotka-Volterra model: Mean square for varying a (left) and σ_p (right).

Table 1 Estimation results for the elliptic Lotka Volterra model. The values are Means (and standard errors) from 100 simulations. Mean individual CPU times are given for each estimation procedure.

true value	θ_1 0.5	θ_4 0.5	σ 0.2	CPU time
σ known, estimation procedure given in Definition 1				
$\hat{\theta}^{(1)}$	0.53 (0.18)	0.52 (0.13)	–	2min30s
$\hat{\theta}^{(2)}$	0.54 (0.19)	0.52 (0.13)	–	1min20s
$\hat{\theta}^I$	0.34 (0.08)	0.49 (0.02)	–	2h15min
σ unknown, estimation procedure given in Definition 2				
$\hat{\theta}^{(1)}$ and $\hat{\sigma}^{(1)}$	0.59 (0.20)	0.50 (0.10)	0.19 (0.01)	3h30min
$\hat{\theta}^{(2)}$ and $\hat{\sigma}^{(2)}$	0.49 (0.19)	0.54 (0.11)	0.20 (0.02)	44min20s
$\hat{\theta}^I$ and $\hat{\sigma}^I$	0.36 (0.15)	0.65 (0.08)	1.18 (0.14)	2h20min

et al., 1962). It is defined by a system of two differential equations which model the membrane potential of the neuron and a recovery channel mimicking the opening/closing of ion channels. More formally, let V_t and U_t denote the membrane potential at time t and the value of the recovery channel at time t , respectively. The stochastic FHN model is defined as:

$$\begin{aligned}
 dV_t &= \frac{1}{\epsilon} (V_t - V_t^3 - U_t) dt + \sigma_1 dW_{1,t} \\
 dU_t &= (\gamma V_t - U_t + \beta) dt + \sigma_2 dW_{2,t} \\
 (V_0, U_0) &= (0, 0)
 \end{aligned} \tag{25}$$

where ϵ is a time scale parameter, γ and β are kinetic parameters, σ_1 and σ_2 are the two diffusion coefficients, $(W_{1,t})_t$ and $(W_{2,t})_t$ are two independent Brownian motions. Only the membrane potential V_t can be measured experimentally.

Trajectories are simulated on the interval $[0, 20]$. Parameters are set to $\epsilon = 0.1$, $\gamma = 1.5$, $\beta = 0.8$, $\sigma_1 = 0.1$, and $\sigma_2 = 0.3$. An example of trajectory for the model (25) is presented on Figure 3. The objective is to estimate $\theta = (\epsilon, \gamma, \beta)$ and $\sigma = (\sigma_1, \sigma_2)$ from the discrete observations of the first coordinate (Y_1, \dots, Y_n) .

The FHN model (25) includes a non-linear term V^3 in the first equation and a constant term β in the second equation. To apply our estimation procedure, we consider a pseudo-linear representation of FHN by introducing a new variable R_t




Fig. 3 Elliptic FitzHugh-Nagumo model: example of a simulation with the dynamics along time of the membrane potential (left) and the recovery variable (right)

assumed to be constant and equal to 1. This can be done by increasing the number of coordinates by one and stating $dR_t = 0$ and $R_0 = 1$:

$$A_\theta(V_t, t) = \begin{pmatrix} \frac{1}{\epsilon}(1 - V_t^2) & -\frac{1}{\epsilon} & 0 \\ \gamma & -1 & \beta \\ 0 & 0 & 0 \end{pmatrix} \quad (26)$$

As explained in section 3.3, the pseudo linear representation is not necessarily unique. Another choice for $A_\theta(V_t, t)$ could have been:

$$\begin{pmatrix} \frac{1}{\epsilon} & -\frac{1}{\epsilon} & -\frac{1}{\epsilon}V_t^3 \\ \gamma & -1 & \beta \\ 0 & 0 & 0 \end{pmatrix}.$$

In this case, we would have taken Y^3 to replace V^3 instead of replacing $V - V^3$ by $(1 - Y^2)V$ in the original SDE (25). But this reveals to be numerically unstable due to the dramatic propagation of the discretization error committed by using Y^3 instead of V^3 . The estimation procedures are then applied with matrix (26)

with $d = 3$, $C = (1 \ 0 \ 0)$ and $B_\sigma = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \\ 0 & 0 \end{pmatrix}$. To estimate $\hat{\sigma}^{(2)}$ defined by (19),

note that again the index $l = 0$ is sufficient as $\gamma_{0,\theta,\sigma}^2(t_i) = \Delta_i \sigma_1^2 > 0$. Here, we set $[w_{min}, w_{max}] = [50, 100]$.

Results are presented in Table 2. When σ is known the situation is similar to the Lotka-Volterra example. This is not the case when σ is also estimated, $G^{(2)}$ minimization gives a more accurate estimator of σ_2 than $G^{(1)}$ minimization. It has to be noted σ_2 is the diffusion of the unobserved state variable. Interestingly, the estimator given by Ionides et al. (2006) encounters difficulties for estimating the diffusion σ , in particular σ_2 . Only our approach gives an accurate estimator of the diffusion at a reasonable computational price.

Table 2 Estimation results for the elliptic FitzHugh-Nagumo model. The values are Means (and standard errors) from 100 simulations. Mean individual CPU times are given for each estimation procedure.

true value	ϵ	γ	β	σ_1	σ_2	CPU time
σ known, estimation procedure given in Definition 1						
$\hat{\theta}^{(1)}$	0.09 (0.01)	1.28 (0.18)	0.66 (0.13)	–	–	2min20s
$\hat{\theta}^{(2)}$	0.09 (0.01)	1.28 (0.19)	0.66 (0.14)	–	–	4min30s
$\hat{\theta}^I$	0.08 (0.01)	1.68 (0.24)	1.16 (0.35)	–	–	2h20min
σ unknown, estimation procedure given in Definition 2						
$\hat{\theta}^{(1)}$ and $\hat{\sigma}^{(1)}$	0.20 (0.74)	1.37 (1.10)	0.71 (0.84)	0.06 (0.02)	0.02 (0.01)	4h5min
$\hat{\theta}^{(2)}$ and $\hat{\sigma}^{(2)}$	0.09 (0.01)	1.28 (0.19)	0.67 (0.13)	0.12 (0.02)	0.37 (0.02)	1h45min
$\hat{\theta}^I$ and $\hat{\sigma}^I$	0.1 (0.01)	1.54 (0.14)	0.85 (0.12)	0.04 (0.02)	1.50 (0.13)	2h25min

6.2 Simulation study for hypoelliptic systems

We now compare the estimation procedures on two hypoelliptic systems, the Harmonic Oscillator and a hypoelliptic FitzHugh-Nagumo neuronal model. The approach developed in Ditlevsen and Samson (2017) is based on a Stochastic approximation of the EM algorithm which requires to set a number of iterations, a number of particles and a decreasing sequence of positive numbers. For both tested models, we choose the value given by Ditlevsen and Samson (2017) for these hyperparameters. Moreover, this approach requires the formal computation of model specific sufficient statistics which can be mathematically cumbersome. In contrast, the method we propose here does not require any modification comparing to the elliptic case and we select w as before.

6.2.1 Harmonic oscillator

The Harmonic Oscillator is a mechanistic model describing oscillations governed by a white noise. It is described by a system of two equations, denoted V_t and U_t , the noise entering only in the second equation (Pokern et al., 2009). The model is defined as follows:

$$\begin{cases} dV_t = U_t dt \\ dU_t = (-DV_t - \delta U_t) dt + \sigma dW_t \\ (V_0, U_0) = (0, 0) \end{cases} \quad (27)$$

with $D, \delta, \sigma > 0$ and $(W_t)_t$ a Brownian motion.

Trajectories are simulated on the interval $[0, 20]$. Parameters are set to $D = 4$, $\delta = 0.5$ and $\sigma = 0.5$. The objective is to estimate $\theta = (D, \delta)$ and σ from the discrete observations of the first coordinate (Y_1, \dots, Y_n) . We set $[w_{min}, w_{max}] = [10^{10}, 10^{15}]$. Again, let us give some details on the procedure for obtaining $\hat{\sigma}^{(2)}$. We have $B_\sigma = \begin{pmatrix} 0 \\ \sigma \end{pmatrix}$ and $A_\theta = \begin{pmatrix} 0 & 1 \\ -D & -\delta \end{pmatrix}$. So $\Gamma_{\theta, \sigma}(0, t_i) = 0$, $\Gamma_{\theta, \sigma}(1, t_i) = \Delta_i^{\frac{3}{2}} \sigma$. We can deduce that $\gamma_{1, \theta}^2(t_i) = \Delta_i^3 \sigma^2 \neq 0$. This order of variance for V_t is the true one, as explained in Samson and Thieullen (2012). Thus our procedure automatically propagates the noise from the second coordinate to the first one, yielding to an invertible covariance matrix of the process. In that sense, it is close to what is proposed by Ditlevsen and Samson (2017).

Table 3 Estimation results for the hypoelliptic Harmonic Oscillator model. The values are Means (and standard errors) from 100 simulations. Mean individual CPU times are given for each estimation procedure.

	D	δ	σ	CPU time
true value	4	0.5	0.5	
σ known, estimation procedure given in Definition 1				
$\hat{\theta}(\sigma)^{(1)}$	4.18 (0.10)	0.48 (0.02)	–	2min
$\hat{\theta}(\sigma)^{(2)}$	4.13 (0.21)	0.49 (0.04)	–	1min55s
σ unknown, estimation procedure given in Definition 2				
$\hat{\theta}^{(1)}$ and $\hat{\sigma}^{(1)}$	4.08 (0.49)	0.49 (0.22)	0.41 (0.01)	2h45min
$\hat{\theta}^{(2)}$ and $\hat{\sigma}^{(2)}$	4.07 (0.50)	0.48 (0.22)	0.41 (0.01)	1h50min
$\hat{\theta}^{DS}$ and $\hat{\sigma}^{DS}$	4.08 (0.50)	0.66 (0.27)	0.50 (0.01)	3h00min

Results are presented in Table 3. For σ known, both weight selection methods give accurate estimation at a reasonable computational cost. When σ is estimated, our approach gives good results and is faster than the Ditlevsen and Samson (2017) procedure.

6.2.2 Hypoelliptic FitzHugh-Nagumo

We consider the FHN model without noise on the first coordinate, as studied by Ditlevsen and Samson (2017). The model is thus defined by:

$$\begin{aligned} dV_t &= \frac{1}{\epsilon} (V_t - V_t^3 - U_t) dt \\ dU_t &= (\gamma V_t - U_t + \beta) dt + \sigma dW_t \\ (V_0, U_0) &= (0, 0) \end{aligned} \quad (28)$$

and we refer to Section 6.1.2 for the description of the variables and parameters.

Trajectories are simulated on the interval $[0, 20]$. Parameters are set to $\epsilon = 0.1$, $\gamma = 1.5$, $\beta = 0.8$, and $\sigma = 0.3$. As in the elliptic case, the objective is to estimate $\theta = (\epsilon, \gamma, \beta)$ and σ from the discrete observations of the first coordinate. We set $[w_{min}, w_{max}] = [50, 500]$. Model (28) is non-linear and time-inhomogenous. As done in the elliptic case, we choose the pseudo linear matrix (26).

For $\hat{\sigma}^{(2)}$ defined by (19), we have $C = (1 \ 0 \ 0)$, $B_\sigma = \begin{pmatrix} 0 \\ \sigma \\ 0 \end{pmatrix}$ and $A_\theta(V_t, t)$

defined by (26), so $\Gamma_{\theta, \sigma}(0, t_i) = 0$ and $\Gamma_{\theta, \sigma}(1, t_i) = -\Delta_i^{\frac{3}{2}} \frac{\sigma}{\epsilon}$. We can deduce from that $\gamma_{1, \theta}^2(t_i) = \Delta_i^3 \frac{\sigma^2}{\epsilon^2} \neq 0$. This corresponds to the first term of the exact variance of V_t , as proved by Ditlevsen and Samson (2017).

Results are presented in Table 4. Our method shows a larger bias than Ditlevsen and Samson (2017) for parameters γ and β but very good results for ϵ which is known to be difficult to estimate.

The most striking and somewhat counter-intuitive result with our method is that the computational time in the hypoelliptic case, especially when σ is estimated is smaller than for the elliptic case. Indeed, for a fixed σ , the complexity of the optimization problem linked to the drift parameter estimation θ does not depend on the singularity of B_σ . It only relies on the dimension of θ and the dimension of the control space. So the smaller the dimension of σ , the quicker the minimizer of

Table 4 Estimation results for the hypoelliptic FitzHugh-Nagumo model. The values are Means (and standard errors) from 100 simulations. Mean individual CPU times are given for each estimation procedure.

true value	ϵ 0.1	γ 1.5	β 0.8	σ 0.3	CPU time
σ known, estimation procedure given in Definition 1					
$\hat{\theta}(\sigma)^{(1)}$	0.09 (0.01)	1.26 (0.16)	0.65 (0.10)	–	2min30s
$\hat{\theta}(\sigma)^{(2)}$	0.09 (0.01)	1.26 (0.16)	0.65 (0.10)	–	2min50s
σ unknown, estimation procedure given in Definition 2					
$\hat{\theta}^{(1)}$ and $\hat{\sigma}^{(1)}$	0.09 (0.01)	1.27 (0.15)	0.65 (0.10)	0.08 (0.01)	1h35min
$\hat{\theta}^{(2)}$ and $\hat{\sigma}^{(2)}$	0.09 (0.01)	1.27 (0.14)	0.65 (0.10)	0.31 (0.06)	45min
$\hat{\theta}^{DS}$ and $\hat{\sigma}^{DS}$	0.10 (0.01)	1.59 (0.16)	0.87 (0.13)	0.31 (0.02)	6h00min

S_w is found. Having the dimension of σ smaller in the hypoelliptic example than in the elliptic one explains the observed difference in the computational time.

7 Conclusion and discussion

In this work, we propose a new method based on control theory to estimate parameters in SDEs. Its main feature is to propose a unified framework for both elliptic and hypoelliptic models by using a criterion focusing on estimating the Brownian motion realization given the observation rather than solely on the observation. By doing so, we manage to construct a criterion S_w well defined no matter the structure of B_σ with only one hyperparameter to be set. Another advantage of the method is the reasonable computational time. The use of the discrete LQ theory allows us to avoid the use of MCMC and/or other computationally costly stochastic approximation.

However, because of the nested structure of our estimation procedure, we can see that the computational time drastically increases when σ has to be estimated. In particular it is very sensitive to the dimension of σ . This can be a limitation for high dimensional systems. Interestingly, in contrast to classic statistical approaches, the hypoelliptic nature of a system is an advantage for us because the dimension of σ is smaller than the one present in its elliptic counterpart. A perspective could be to investigate a criterion which allows us to simultaneously estimate θ and σ instead of using a nested procedure. However, so far we have not managed to find a criterion which fits in the framework of the discrete LQ theory. This leads us to our second limitation, our method is currently restricted to linear models or non-linear ones with a specific structure (i.e models only nonlinear w.r.t the observed state variable). A main challenge would be to extend our method to general nonlinear SDEs.

Acknowledgements

Quentin Clairon was supported by EPSRC award EP/M015637/1.

Adeline Samson has been partially supported by the LabEx PERSYVAL-Lab (ANR-11-LABX-0025-01) funded by the French program Investissements d'avenir.

Conflict of interest

On behalf of all authors, the corresponding author states that there is no conflict of interest.

References

- Aït-Sahalia, Y. (2008) Closed-form likelihood expansions for multivariate diffusions. *Ann. Statist.*, **36**, 906–937.
- Bally, V. and Talay, D. (1996) The law of the euler scheme for stochastic differential equations i. convergence rate of the distribution function. *Monte Carlo Methods and Applications*, 93–128.
- Bertsekas, D. (2005) *Dynamic Programming and Optimal Control*. Athena Scientific.
- Beskos, A., Papaspiliopoulos, O., Roberts, G. O. and Fearnhead, P. (2006) Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **68**, 333–382.
- Bibby, B. and Sorensen, M. (1995) Martingale estimating functions for discretely observed diffusion processes. *Bernoulli*, **1**, 17–39.
- Bierkens, J., Van der Meulen, F. and Schauer, M. (2018) Simulation of elliptic and hypo-elliptic conditional diffusions. *submitted*.
- Brunel, N. J.-B. and Clairon, Q. (2015) A tracking approach to parameter estimation in linear ordinary differential equations. *Electronic Journal of Statistics*, **9**, 2903–2949.
- Cattiaux, P., León, J. and Prieur, C. (2014a) Estimation for stochastic damping hamiltonian systems under partial observation. i. invariant density. *Stochastic Process. Appl.*, **124**, 1236–1260.
- (2014b) Estimation for stochastic damping hamiltonian systems under partial observation. ii. drift term. *ALEA*, **11**, 359–384.
- Cimen, T. and Banks, S. (2004a) Global optimal feedback control for general nonlinear systems with nonquadratic performance criteria. *Systems and Control Letters*, **53**, 327–346.
- (2004b) Nonlinear optimal tracking control with application to super-tankers for autopilot design. *Automatica*, **40**, 1845–1863.
- Clairon, Q. and Brunel, N. (2017) Optimal control and additive perturbations help in estimating ill-posed and uncertain dynamical systems. *JASA*.
- (2018) Tracking for parameter and state estimation in possibly misspecified partially observed linear ordinary differential equations. *JSPI*.
- Comte, F., Prieur, C. and Samson, A. (2017) Adaptive estimation for stochastic damping hamiltonian systems under partial observation. *Stochastic processes and their applications*.
- Cuenod, C., Favetto, B., Genon-Catalot, V., Rozenholc, Y. and Samson, A. (2011) Parameter estimation and change-point detection from dynamic contrast enhanced MRI data using stochastic differential equations. *Mathematical Biosciences*, **233**, 68–76.
- Dietz, H. (2001) Asymptotic behaviour of trajectory fitting estimators for certain non-ergodic sde. *Statistical Inference for Stochastic Processes*, **4**, 249–258.

- Ditlevsen, S. and Greenwood, P. (2013) The morris-lecar neuron model embeds a leaky integrate-and-fire model. *Journal of Mathematical Biology*, **67**, 239–259.
- Ditlevsen, S. and Samson, A. (2014) Estimation in the partially observed stochastic morris-lecar neuronal model with particle filter and stochastic approximation methods. *Annals of Applied Statistics*, **2**, 674–702.
- (2017) Hypocoelliptic diffusions: discretization, filtering and inference from complete and partial information. *arXiv:1707.04235v1*, 1–33.
- Ditlevsen, S., Yip, K. and Holstein-Rathlou, N. (2005) Parameter estimation in a stochastic model of the tubuloglomerular feedback mechanism in a rat nephron. *Mathematical Biosciences*, **194**, 49–69.
- Durham, G. B. and Gallant, A. R. (2002) Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *J. Bus. Econom. Statist.*, **20**, 297–338.
- Elerian, O., Chib, S. and Shephard, N. (2001) Likelihood inference for discretely observed nonlinear diffusions. *Econometrica*, **69**, 959–993.
- Eraker, B. (2001) MCMC analysis of diffusion models with application to finance. *J. Bus. Econom. Statist.*, **19**, 177–191.
- FitzHugh, R. (1961) Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, **6**, 445–466.
- Gerstner, W. and Kistler, W. (2002) *Spiking Neuron Models*. Cambridge University Press.
- Gloter, A. (2006) Parameter estimation for a discretely observed integrated diffusion process. *Scand. J. Statist.*, **33**, 83–104.
- Golightly, A. and Wilkinson, D. J. (2006) Bayesian sequential inference for nonlinear multivariate diffusions. *Stat. Comput.*, **16**, 323–338.
- (2008) Bayesian inference for nonlinear multivariate diffusion models observed with error. *Comput Stat & Data Analysis*, **52**, 1674–1693.
- Graham, M. and Storkey, A. (2017) Asymptotically exact inference in differentiable generative models. *arXiv:1605.07826*, 14.
- Iolov, A., Ditlevsen, S. and Longtin, A. (2017) Optimal design for estimation in diffusion processes from first hitting times. *SIAM J. Uncertainty Quantification*, **5**, 88–110.
- Ionides, E. L., Bhadra, A., Atchade, Y. and King, A. A. (2011) Iterated filtering. *Annals of Statistics*, **39**, 1776–1802.
- Ionides, E. L., Breto, C. and King, A. A. (2006) Inference for nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, **103**, 18438–18443.
- Jensen, A. C., Ditlevsen, S., Kessler, M. and Papaspiliopoulos, O. (2012) Markov chain Monte Carlo approach to parameter estimation in the FitzHugh-Nagumo model. *Physical Review E*, **86**, 041114.
- Kessler, M. (1997) Estimation of an ergodic diffusion from discrete observations. *Scand. J. Statist.*, **24**, 211–229.
- King, A. A., Nguyen, D. and Ionides, E. L. (2016) Statistical inference for partially observed markov processes via the r package pomp. *Journal of Statistical Software*, **69**.
- Kloeden, P. and Platen, E. (1992) *Numerical Solution of Stochastic Differential Equations*. Springer-Verlag Berlin.
- Kutoyants, Y. (1991) Minimum distance parameter estimation for diffusion type observation. *Comptes rendus de l Academie des sciences*, **312**, 637.

- Leon, J., Rodriguez, L. and Ruggiero, R. (2019) Consistency of a likelihood estimator for stochastic damping hamiltonian systems. totally observed data. *ESAIM PS*, **23**, 1–36.
- Lipster, R. and Shiryaev, A. (2001) *Statistics of random processes I : general theory*. Springer.
- Lotka, A. (1925) *Elements of Mathematical Biology*. Dover.
- Mao, X., Marion, G. and Renshaw, E. (2002) Environmental brownian noise suppresses explosions in populations dynamics. *Stochastic Processes and Their Applications*, **97**, 95–110.
- Martin, C., Sun, S. and Egerstedt, M. (2001) Optimal control, statistics and path planning. *Mathematical and Computer Modelling*, **33**, 237–253.
- Mattingly, J., Stuart, A. M. and Higham, D. (2002) Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stochastic Process. Appl.*, **101**, 185–232.
- Meeds, E. and Welling, M. (2015) Optimization monte carlo: Efficient and embarrassingly parallel likelihood-free inference. In *Advances in Neural Information Processing Systems*.
- Melnykova, A. (2019) Parametric inference for multi-dimensional hypo-elliptic diffusion with full observations. *Stat. Inf. Stoc. Proc.*
- Van der Meulen, F. and Schauer, M. (2016) Bayesian estimation of discretely observed multi-dimensional diffusion processes using guided proposals. *submitted*.
- Nagumo, J., Animoto, S. and Yoshizawa, S. (1962) An active pulse transmission line simulating nerve axon. *Proc. Inst. Radio Eng.*, **50**, 2061–2070.
- Paninski, L., Ahmadian, Y., Ferreira, D. G., Koyama, S., Rad, K. R., Vidne, M., Vogelstein, J. and Wu, W. (2010) A new look at state-space models for neural data. *Journal of Computational Neuroscience*, **29**, 107–126.
- Paninski, L., Vidne, M., DePasquale, B. and Ferreira, D. (2012) Inferring synaptic inputs given a noisy voltage trace via sequential Monte Carlo methods. *Journal of Computational Neuroscience*, **33**, 1–19.
- Pedersen, A. (1995) A new approach to maximum likelihood estimation for stochastic differential equations based on discrete observations. *Scand. J. Statist.*, **22**, 55–71.
- Pokern, Y., Stuart, A. and Wiberg, P. (2009) Parameter estimation for partially observed hypoelliptic diffusions. *J. Roy. Stat. Soc. B*, **71**, 49–73.
- Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V. and Mischenko, E. F. (1962) *The Mathematical Theory of Optimal Processes*. Wiley-Interscience.
- Samson, A. and Thieullen, M. (2012) A contrast estimator for completely or partially observed hypoelliptic diffusion. *Stochastic Processes and their Applications*, **122**, 2521–2552.
- Sontag, E. (1998) *Mathematical Control Theory: Deterministic finite-dimensional systems*. Springer-Verlag (New-York).
- Sørensen, H. (2004) Parametric inference for diffusion processes observed at discrete points in time: a survey. *Int Stat Rev*, **72**, 337–354.
- Trelat, E. (2005) *Contrôle optimal : Théories et applications*. Vuibert.

A Discrete LQ theory: Main theorem

In section 3, the theorem which gives us the profiled cost value, as well as the corresponding control sequence, is a particular case of a more general theorem. This theorem, fundamental in LQ theory, ensures the existence, uniqueness and gives a closed form for the global minimizer of a cost under the form:

$$J(u) = x_N^T S x_N + \sum_{i=0}^{N-1} x_i^T Q_i x_i + u_i^T R_i u_i \quad (29)$$

where $u = \{u_0, \dots, u_{N-1}\}$ and the state variable sequence $x = \{x_0, \dots, x_N\}$ are linked by the finite difference equation:

$$x_{k+1} = A_k x_k + B_k u_k. \quad (30)$$

The derivation of the next theorem by the optimality principle can be found in Bertsekas (2005).

Theorem 2 *Let us assume that S is positive semi-definite, for all $i \in \llbracket 0, N-1 \rrbracket$, Q_i is positive semi-definite and R_i is positive definite. Then the cost (29) reaches its global minimum for the control sequence u^* given by:*

$$u_k^* = -[R_k + B_k^T P_{k+1} B_k]^{-1} B_k^T P_{k+1} A_k x_k$$

and the minimal cost value is equal to:

$$J(u^*) = x_0^T P_0 x_0$$

where P_k is given by the discrete time Riccati difference equation:

$$\begin{aligned} P_k &= A_k^T P_{k+1} A_k + Q_k - A_k^T P_{k+1} B_k [R_k + B_k^T P_{k+1} B_k]^{-1} B_k^T P_{k+1} A_k \\ P_N &= S. \end{aligned}$$