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

Quentin Clairon[1], Adeline Samson[2]

[1] *Newcastle University, School of Mathematics and Statistics, Newcastle Upon Tyne, United Kingdom*
[2] *Univ. Grenoble Alpes, CNRS, Grenoble INP*, LJK, 38000 Grenoble, France*
* *Institute of Engineering Univ. Grenoble Alpes*

October 15, 2017

Abstract

Multi-dimensional Stochastic Differential Equations (SDEs) are a powerful tool to describe dynamics of several fields (pharmacokinetic, neurosciences, ecology, etc). The estimation of the parameters of these systems has been widely studied. We focus in this paper in the case of partial observations, only a one-dimensional observation is available. We consider two families of SDE, the elliptic family with a full-rank diffusion coefficient and the hypoelliptic family with a degenerate diffusion coefficient. The estimation for the second class is much more difficult and only few references have proposed estimation strategies in that case. Here, we adopt the framework of the optimal control theory to derive an estimation 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 parameters of the drift and of the diffusion coefficient. The estimation reveals to be very stable in a simulation study conducted on different examples (Harmonic Oscillator, FitzHugh-Nagumo, Lotka-Volterra).

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

1 Introduction

We focus on the statistical inference for d-dimensional stochastic dynamical systems modeled by a stochastic differential equation (SDE). We are interested in the case of partial observations: only the first one-dimensional coordinate, denoted V_t, of the system is observed while the other $(d-1)$-dimensional coordinates, denoted U_t, are unobserved. The system is written as follows:

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

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 motions and σ_1, σ_2 are the two diffusion coefficients.

We consider two classes of models (1). The first class, called *elliptic*, corresponds to an SDE with a full non degenerate diffusion coefficient. This means that denoting $B_\sigma = (\sigma_1, \sigma_2)$ the diffusion coefficient, 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 a degenerate stochastic noise: the diffusion coefficient $B_\sigma B_\sigma^T$ is not invertible. For example when $\sigma_1 = 0$.

These two specificities, partial observations and hypoelliptic/elliptic properties, are of increasing importance in many applications, and we give some examples below. But before describing the examples, let us remark that these two specificities are not of the same nature. The first one 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. As will be recalled later when reviewing the literature, this increases the difficulty of estimating the 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 the process (V_t, U_t). The two systems, elliptic and hypoelliptic, may not have the same interpretation depending on the applications (see below examples in neurosciences). It might nevertheless be difficult for the modeler to know in advance if the system is more likely 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 on the 'wrong' class of models. This is an advantage of our method which is the same for elliptic and hypoelliptic SDE.

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, Donnet and Samson, 2013]. In system biology, the famous stochastic Lotka-Volterra model [Lotka, 1925, Meeds and Welling, 2015, Graham and Storkey, 2017, Mao et al., 2002] 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 [Pospischil et al., 2009, Paninski et al., 2010, 2012, Ditlevsen and Greenwood, 2013, Ditlevsen and Samson, 2017] 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 dynamical system reduces to $dV_t = U_t dt$.

Let us now review the main estimation methods that have been proposed in the literature to estimate the drift parameter θ and the diffusion coefficient B_σ.

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], Ditlevsen and Samson [2014], van der Meulen and Schauer [2016a]. The case of partial observations has also been considered with several approaches. The unobserved coordinates are treated as missing data and are imputed, see for examples Elerian et al. [2001], Bjørn [2001], Golightly and Wilkinson [2006, 2008], Ditlevsen and Samson [2014], van der Meulen and Schauer [2016b]. Most 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 time of computation.

Let us now explain why the estimation of hypoelliptic systems is more difficult. Let us imagine that the complete observations of (V_t, U_t) are available in continuous time. Estimating θ would be naturally performed through the Girsanov formula, that gives directly the likelihood [Lipster and Shiryaev, 2001]. However, the Girsanov formula requires the matrix $B_\sigma B_\sigma^T$ to be invertible. Because of the singularity of B_σ in the hypoelliptic case, this inverse does not exist and the likelihood is not properly defined. The same problem occurs for the estimation methods developed for elliptic systems that have been cited above, as they also generally require $B_\sigma B_\sigma^T$ to be invertible. They can thus not be applied to hypoelliptic systems. There are thus 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] 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, but do not correct the bias. In the non-parametric framework, Cattiaux et al. [2014a,b], Comte et al. [2017] consider the estimation of the drift, the diffusion coefficient and the invariant density using kernel estimators. For hypoelliptic SDEs that are more general than the stochastic Damping Hamiltonian system, we are only aware of the work of Ditlevsen and Samson [2017]. Their approach is based on a discretization scheme of order 1.5, a particle filter to approximate the unobserved coordinate and a maximization of a statistical contrast by stochastic approximation. The main drawbacks are the computational time induced by the particle filter and the fastidious calculations to exhibit the sufficient statistics of the likelihood that are then stochastically approximated. There is thus a need to develop new approaches.

In this paper, we propose a strategy based on the 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 input, or control, do we have to apply to it in order to steer it to a desired behavior in an optimal way ? This problem is formulated as an optimization problem under constraint where the optimality is defined through the introduction of a cost function and the proposed dynamical model belongs to the contraints. Under this theory a large variety of theoretical and numerical tools have been developed to solve this kind of optimization problems, that is 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 among others 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 [2016, 2017], Iolov et al. [2017], Zhang et al. [2017]. One way of using the optimal control theory is to rewrite an estimation criteria, typically a likelihood or a posterior distribution as a tracking problem. Tracking problems are a specific class of optimal control theory problems: the aim is to find the optimal control which leads a coordinate of the system the closest possible to a target trajectory, here the observations, on a given observation interval. A way of solving an optimal control problem is the Pontryagin maximum principle [Pontryagin et al., 1962, Trelat, 2005, Sontag, 1998] which allows to calculate this optimal control as well as the corresponding system response.

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, 2016, 2017]. However, their method, based on the deterministic theory of optimal control, is restricted to the case of deterministic system. In this paper, we adapt this idea to the problem of SDE estimation. To do so, we resort to the framework of the discrete optimal control theory. Theoretical and numerical results have been fully developed for linear model, 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 criteria. 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 SDE. When estimating parameters of a nonlinear SDE, we propose to rewrite the estimation criteria 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, it reveals to be efficient in practice.

Our estimation procedure follows three steps. First we define a criteria to estimate the drift parameter θ alone, σ being considered known. This criteria is minimized with the help of the linear-quadratic theory. This method introduces a weighting parameter w which needs to be selected. The second step consists in constructing an external criteria 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 paper is organized as follows. Section 2 presents the elliptic and hypoelliptic models. Section 3 introduces the estimation of θ when σ is known. In section 4, the method used to select the weighting parameter w is presented. Section 5 explains the estimation of σ. Section 6 illustrates the procedure on four elliptic or hypoelliptic models: an elliptic Lotka-Volterra and FitzHugh-Nagumo model, and then an hypoelliptic Harmonic Oscillator and FitzHugh-Nagumo model.

2 Models and objectives

As explained in the introduction, 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 and the corresponding estimation criteria. In the simulation section, we will show an adaptation of the method for a non-linear SDE.

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 a known initial condition (V_0, U_0) 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 $A_{1j}(t)$ denote the jth element of the first column of $A_\theta(t)$ and $B_{,j}$ the jth column of B_σ. For any t, there exists at least one $j = 2, \ldots, d$ 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 $t \longmapsto Y(t)$, the realization of V_t. We assume that Y is discretely observed on the interval $[0, T]$ at times $0 = t_0 < \ldots < t_n = T$ without measurement error and denote $(Y_0, \ldots, Y_n)$ these observations. We thus have $Y_i = CZ_{t_i}$, with C a $1 \times d$-matrix $C = (1, 0_{d-1})^T$ and 0_{d-1} is the row vector of zeros of size $d-1$.

The aim of the paper is to estimate the unknown parameters of model (2) in the elliptic and hypoelliptic cases using the discrete observations $(Y_0, \ldots, Y_n)$. 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 based on the optimal control theory, let us explain why the estimation problem is difficult. When one wants to estimate θ and σ of an elliptic, partially observed SDE, the standard approach starts with the discretization of the diffusion (2). The Euler–Maruyama discretization scheme at time $(t_{1,}, \ldots t_n)$ is defined as follows for $i = 0, \ldots, n-1$ and $Z_i = (V_i, U_i)$:

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

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

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

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

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

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

Therefore, there is a need to propose 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 criteria, that involves only the non-degenerate part of the stochastic noise. This criteria is thus suitable for elliptic as well as for hypoelliptic SDEs. The estimation procedure is presented in Sections 3-5.

3 Estimation of θ via optimal control theory

In this section, we assume σ known and fixed at its true value. We expose our estimation procedure: its main principle is to turn the statistical problem into an optimal control one. The procedure is the same whether B_σ is singular or not. We start this section by deriving the general optimal control problem from which our estimator is defined. The optimal control problem is then solved by using a result coming from the linear quadratic theory (Section 3.2).

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

The principle of the control problem is to introduce a *control sequence* such that the 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 $\triangle$ [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*:

$$\begin{array}{l} \begin{pmatrix} V_{i+1} \\ U_{i+1} \end{pmatrix} = \mathbf{A}_\theta(t_i) \begin{pmatrix} V_i \\ U_i \end{pmatrix} + B_\sigma u_i \\ (V_{t_0}, U_{t_0}) = (V_0, U_0) \end{array} \tag{7}$$

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 the data $(Y_0, \dots, Y_n)$ are a sample of model (2). The control sequence u can be viewed as this specific realization of the Brownian motion. The objective is to infer this sequence u by extracting knowledge from the data. For that purpose, a cost function $C(u, Y; \theta, \sigma)$ is defined such that the optimum in u corresponds to this realization of the Brownian motion. We say that the sequence u is *designed* through the cost function $C(u, Y; \theta, \sigma)$.

A natural cost function is the conditional posterior probability $P(u|Y; \theta, \sigma)$ of u knowing the data Y. Filtering consists in computing the conditional expectation of this posterior distribution. An alternative is to compute the maximum a posteriori (MAP) by maximizing $P(u|Y; \theta, \sigma)$ with respect to u. That is the way we consider to compute the control in this paper. A direct computation of the MAP reveals to be numerically unstable or too strict for SDEs. Therefore, we introduce a weight to stabilize the optimization problem. Let us introduce the basic 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 $P(Y;\theta,\sigma)$, the likelihood, does not depend on u. Let us now detail the two terms that depend on u. The conditional density $P(Y|u;\theta,\sigma)$ is described by the error induced by the Euler discretisation of the SDE. It is well known that this error has a variance of order $\triangle$ the time step [Bally and Talay, 1996]. The second term $P(u;\theta,\sigma)$ is the density of a discretized Brownian motion, which has a variance $\triangle$. Therefore, one can formulate the MAP as

$$\hat{u}_{MAP} = \arg\max_{u}\left(-\sum_{i=0}^{n}\frac{\left(V_{i,\theta,\sigma,u}-Y_i\right)^2}{\triangle_i}-\sum_{i=0}^{n-1}\frac{u_i^T u_i}{\triangle_i}\right) = \arg\min_{u} C(u,Y;\theta,\sigma) \quad (8)$$

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

The optimisation of the MAP (8) reveals to be too strict and intractable in practice. Therefore, we relax the constraints by introducing a weight between the two terms. To ease the writing of the cost function, we replace the scaling $\frac{1}{\triangle_i}$ of the first term by a weight $w > 0$. The cost function of the control is finally defined as

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

The weight w has to be chosen by the user. In Section 4 a procedure is proposed to select w adaptively from the data. To enter the optimal control theory, we exhibit the last observation Y_n which plays a specific role in a control problem. The cost function can thus be rewritten:

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

For a given weight w, the *best control* is the sequence u that minimizes $C_w(u,Y;\theta,\sigma)$ under the constraint of model (7). It is defined as the solution of the following optimal control problem:

$$\begin{array}{ll}\text{Minimize in } u: & C_w(u,Y;\theta,\sigma) = w\left(V_{n,\theta,\sigma,u}-Y_n\right)^2 + \sum_{i=0}^{n-1}\left(w\left(V_{i,\theta,\sigma,u}-Y_i\right)^2 + \frac{1}{\triangle_i}u_i^T u_i\right) \\ \text{Subject to:} & \left\{\begin{array}{l}\begin{pmatrix}V_{i+1,\theta,\sigma,u}\\ U_{i+1,\theta,\sigma,u}\end{pmatrix} = \mathbf{A}_\theta(t_i)\begin{pmatrix}V_{i,\theta,\sigma,u}\\ U_{i,\theta,\sigma,u}\end{pmatrix} + B_\sigma u_i \\ \left(U_{0,\theta,\sigma,u}\, V_{0,\theta,\sigma,u}\right) = \left(V_0,\, U_0\right).\end{array}\right.\end{array} \quad (9)$$

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

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

where

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

is the profiled cost C_w over the set of possible sequences of the Brownian motion increments.

The computations of $\overline{u}_{\theta,\sigma}$ and $S_w(Y;\theta,\sigma)$ require to solve the optimal control problem (9). We thus have to prove the existence and uniqueness of the solution of problem (9), and that this unique solution is numerically computable to obtain $\overline{u}_{\theta,\sigma}$ and $S_w(Y;\theta,\sigma)$. These two results are given by the linear quadratic theory [Trelat, 2005, Sontag, 1998] and are exposed below.

3.2 Linear quadratic theory

The linear quadratic theory is derived from the Pontryagin maximum principle [Pontryagin et al., 1962] in the restricted framework of linear models. For a given (θ, σ), the linear quadratic theory ensures, for the problem (9):

- the existence and uniqueness of $\overline{u}_{\theta,\sigma}$, the solution of (9),
- that $\overline{u}_{\theta,\sigma}$ and $S_w(Y;\theta,\sigma)$ can be computed by the introduction of 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 $i = 0, \ldots, n-1$:

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

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 positive solution of the discrete backward Riccati equation, for $i = n-1, \ldots, 0$:

$$R_{i,\theta,\sigma} = Q_i + A_{i,\theta}^T \left(R_{i+1,\theta,\sigma} - R_{i+1,\theta,\sigma} B_\sigma^0 \left(\frac{1}{\triangle_i} I_m + B_\sigma^{0\,T} R_{i+1,\theta,\sigma} B_\sigma^0 \right)^{-1} B_\sigma^{0\,T} R_{i+1,\theta,\sigma} \right) A_{i,\theta} \tag{11}$$

We can establish the following theorem.

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

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

Moreover, the minimum value of C_w is equal to:

$$\begin{aligned} S_w(Y;\theta,\sigma) &= \min_{u \in \mathbb{R}^{m \times (n-1)}} C_w(u, Y; \theta, \sigma) = C_w(\overline{u}_{\theta,\sigma}, Y; \theta, \sigma) \\ &= (V_0, U_0^T, 1)\, R_{0,\theta,\sigma}\, (V_0, U_0^T, 1)^T \end{aligned} \tag{13}$$

The proof is given in Appendix.

Theorem 1 provides a closed expression for S_w that depends on the matrix $R_{0,\theta,\sigma}$ which is easily obtained by solving the Riccati equation (11) and on (V_0, U_0) the initial condition which is assumed to be known. The estimator of θ can thus be rewritten as

$$\widehat{\theta}_w(\sigma) = \arg\min_\theta \left\{ (V_0, U_0^T, 1)\, R_{0,\theta,\sigma}\, (V_0, U_0^T, 1)^T \right\}$$

Note that using the expression of $\overline{u}_{\theta,\sigma}$ given by (12) in the discretized model (7), the system dynamics under the optimal control is given by:

$$\begin{cases} Z_{i+1,\theta,\sigma,\overline{u}_{\theta,\sigma}} = \left(\mathbf{A}_\theta(t_i) - B_\sigma (\frac{1}{\triangle_i} I_m + B_\sigma^{0\,T} R_{i+1} B_\sigma^0)^{-1} B_\sigma^{0\,T} R_{i+1} A_i \right) Z_{i,\theta,\sigma,\overline{u}_{\theta,\sigma}} \\ Z_{0,\theta,\sigma,\overline{u}_{\theta,\sigma}} = (V_0, U_0)\,. \end{cases} \tag{14}$$

Let us now enlighten the noticeable advantage of our approach. Theorem 1 holds for elliptic ($m = d$) SDEs as well as for hypoelliptic ($m < d$) SDEs. The estimation criteria does not involve the inverse of the matrix $B_\sigma B_\sigma^T$, contrary to the standard estimation approaches for SDEs. Moreover, the linear quadratic theory itself does not require the matrix $B_\sigma B_\sigma^T$ to be full-conditioned. This can be explained by our choice of the cost function $C_\omega(u, Y; \theta, \sigma)$, which is strictly convex with respect to the sequence u due to the term $u_i^T u_i$. The cost function is thus automatically adapted to the dimension of the noise. Once again, it is noticeable that the solution to the optimal control problem is the same for the two classes of SDEs and that no specific calculations are required for the hypoelliptic case.

The relaxation made on the original cost C coupled with the Euler-Maruyama discretization scheme allows us to make our estimation problem fits in the framework of the discrete LQ theory. This ensures in turn 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. However, it comes with the price that the theoretical methods used for deriving consistency and convergence rate in Brunel and Clairon [2015] cannot be applied anymore. In Brunel and Clairon [2015], the control problem was a continuous one, the original cost C was not approximated and the original ODE model was not discretized, only perturbed by the addition of a control function u. Thus, the source of error was due to the noisy nature of observation and thus theoretical results have been derived by classic methods for M-estimator (see van der Vaart [1998] for example). Nevertheless, our procedure reveals to be efficient in practice, see Section 6.

The next step is the selection of the weight w involved in the definitions of the optimal control problem (9) and the estimator $\widehat{\theta}_w(\sigma)$. This is explained in the next section.

4 Adaptive w selection

The choice of the weight w is critical to ensure the stability of the procedure. We propose to select it adaptively by minimizing an external functional criteria $G(Y; \sigma, w, \theta)$ that depends on w, the data and the parameters θ and σ. We propose hereafter two possible choices for G: $G^{(1)}$ uses a property of the average quadratic growth of a Brownian process and gives importance to the term of adequacy of the control sequence u; $G^{(2)}$ relies on a new contrast that gives more importance to the data fidelity term.

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

1. For any weight parameter w, compute the estimator $\widehat{\theta}_w(\sigma) = \arg\min_\theta S_w(Y; \theta, \sigma)$.
2. Minimize the functional $G = G^{(1)}$ or $G = G^{(2)}$ with respect to w, using the plugged value $\widehat{\theta}_w(\sigma)$, and define
$$\widehat{w} := \arg\min_w G(Y; \sigma, w, \widehat{\theta}_w(\sigma)) \tag{15}$$
3. Define the final estimator of θ as
$$\widehat{\theta}(\sigma) = \widehat{\theta}_{\widehat{w}}(\sigma) \tag{16}$$

The two functionals are presented in the next sections.

4.1 w selection via a quadratic growth moment condition

We now present the first functional $G^{(1)}$. It comes from a constraint linked to the distribution of the control. The control $\overline{u}_{\theta,\sigma}$ mimics the increments of the Brownian motion. Thus, it is natural to impose that the values $\frac{\overline{u}_{i,\theta,\sigma}}{\sqrt{\triangle_i}}$ are independent and distributed with $N(0,\, I_m)$. We can then derive constraints related to this distribution. For example, let us denote $\overline{u}_{i,j,\theta,\sigma}$ the $j-$th component of the m dimensional Brownian motion at time t_i, for $j \in [\![1,\, m]\!]$. The law of large number implies that, almost surely,

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

We thus propose the following functional to select w:

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

Here, the dependency of $G^{(1)}(Y;\sigma,w,\theta)$ w.r.t w is made through the optimal control sequence $\overline{u}_{\theta,\sigma}$. This functional gives importance to the control sequence property.

4.2 w selection via a data fidelity criteria

The second functional is driven by the idea of giving importance to the data fidelity term. We can show by recurrence that for $i = 0, \ldots, n-1$ and any integer $0 \leq l \leq i-1$:

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

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

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

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

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

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

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

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

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

Let us choose l such that $\gamma_{l,\theta,\sigma}^2(t_i)$ is non 0. 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 $X_{i,\theta,\sigma,u}^l := \frac{X_{i,\theta,\sigma,u}}{\gamma_{l,\theta}(t_i)}$. Then $\mathbb{E}\left[\left(X_{i,\theta,\sigma,u}^l\right)^2\right] = 1$. Note that $X_{i,\theta,\sigma,u}^l$ depends on $(u_{i-l},\ldots,\, u_i)$ and $X_{i+l+1,\theta,\sigma,u}^l$ depends on $(u_{i+1},\ldots,\, u_{i+l+1})$ thus they are independent. For the i.i.d sequence of $\left\{X_{(l+1)i+1,\theta,\sigma,\overline{u}_{\theta,\sigma}}^l\right\}_{0\le i\le \frac{n-1}{l+1}}$ of size $L := \left[\frac{n+l}{l+1}\right]$, 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:

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

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 $\left\{X_{(l+1)i+1,\theta,\sigma,\overline{u}_{\theta,\sigma}}^l\right\}_{0\le i\le \frac{n-1}{l+1}}$.

- For elliptic SDEs, it is sufficient to take $l = 0$. Indeed $\Gamma_{\theta,\sigma}(0,t_i) = C\sqrt{\triangle_i}B_\sigma \neq 0_m$ since B_σ is of full rank.
- For hypoelliptic SDEs, we have $\Gamma_{\theta,\sigma}(0,t_i) = C\sqrt{\triangle_i}B_\sigma = 0_m$. With $l = 1$, we have $\gamma_{1,\theta,\sigma}^2(t_i) = \Gamma_{\theta,\sigma}(1,t_i)\,\Gamma_{\theta,\sigma}(1,t_i)^T = \triangle_{i-1}C\mathbf{A}_\theta(t_i)B_\sigma B_\sigma^T\mathbf{A}_\theta(t_i)^T C^T$ 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

So far, the vector σ of the parameters involved in the diffusion coefficient B_σ has been considered known. We now present its estimation. The standard estimator of a diffusion coefficient is the quadratic variation of the stochastic process. It can however not be applied in the context of partial observations because only the first coordinate is observed. We therefore propose to minimize an estimation criteria. It turns out that the two functionals $G^{(1)}$ and $G^{(2)}$ presented in Section 4 can be used to estimate σ. It might seem counter-intuitive that the same criteria 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 used judiciously to estimate also the diffusion coefficient B_σ.

Let us now present the procedure. It is a nested procedure that provides the final estimation of θ, σ and the data-driven selection of the weight w.

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

1. For any weight parameter w and diffusion coefficient σ, compute the estimator $\widehat{\theta}_w(\sigma) = \arg\min_\theta S_w(Y; \theta, \sigma)$.

2. Minimize the functional $G = G^{(1)}$ or $G = G^{(2)}$ with respect to σ, using the plugged value $\widehat{\theta}_w(\sigma)$, and define
$$\widehat{\sigma}(w) := \arg\min_\sigma G(Y; \sigma, w, \widehat{\theta}_w(\sigma)) \tag{20}$$

3. Minimize the functional $G = G^{(1)}$ or $G = G^{(2)}$ with respect to w, using the plugged value $\widehat{\theta}_w(\sigma)$, and define
$$\widehat{w} := \arg\min_w \left\{ G(Y; \widehat{\sigma}(w), w, \widehat{\theta}_w(\widehat{\sigma}(w))) \right\} \tag{21}$$

4. Define the final estimators of θ and σ as
$$\begin{aligned} \widehat{\sigma} &:= \widehat{\sigma}(\widehat{w}) \\ \widehat{\theta} &:= \widehat{\theta}_{\widehat{w}}(\widehat{\sigma}) \end{aligned} \tag{22}$$

The two functions $G^{(1)}$ and $G^{(2)}$ can be used in this nested procedure.

6 Simulation study

We want to evaluate the two estimation procedures given in definition 1 when σ is known and definition 2 when σ is also estimated. A simulation study is conducted on several SDEs, elliptic (Section 6.1) and hypoelliptic (Section 6.2). For each model, a hundred trajectories with $n = 1\ 000$ equidistant points are simulated. Then the mean and standard error (SE) of the estimators are computed. The mean required CPU time (computed with the MATLAB function cputime) is also reported.

When σ is known, two estimators are computed using the nested procedure given in definition 1. We denote $\widehat{\theta}^{(1)}(\sigma)$ ($\widehat{\theta}^{(2)}(\sigma)$, respectively) the estimator defined by (16) with the weight w adaptively selected by choosing the functional $G = G^{(1)}$ ($G = G^{(2)}$, respectively). When σ is also estimated, two estimators are also computed. We denote $\widehat{\theta}^{(1)}$ and $\hat{\sigma}^{(1)}$ ($\widehat{\theta}^{(2)}$ and $\hat{\sigma}^{(2)}$, respectively) the estimators given in definition 2 using $G = G^{(1)}$ ($G = G^{(2)}$, respectively).

Three models are used: the Lokta-Volterra model, the FitzHugh-Nagumo model and the Harmonic Oscillator (HO). The two firsts are non linear and HO is linear. As explained in the introduction, the procedures that we propose apply to linear SDE only. This allows to take advantage of the explicit solution to the optimal control problem thanks to the linear-quadratic theory. For the non-linear SDE, we thus propose strategies to bypass this limit and adapt our estimators to these non-linear models.

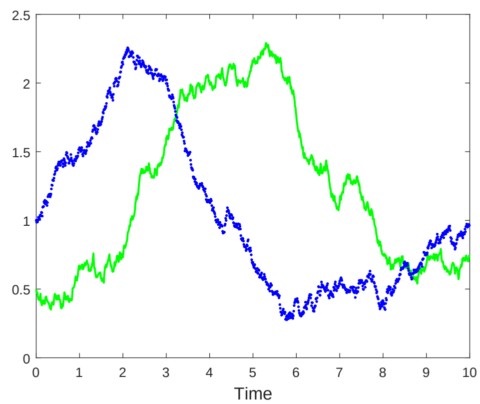

Figure 1: Lotka-Volterra model: example of a simulation with the dynamics along time of the number of predators (green plain line) and the number of preys (blue dotted line).

6.1 Simulation study for elliptic systems

Two models are studied in their elliptic version, the Lotka-Volterra system and the FitzHugh Nagumo model.

6.1.1 Elliptic Lotka-Volterra process

Stochastic versions of the Lotka-Volterra model have been proposed, for examples by Lotka [1925], Meeds and Welling [2015], Graham and Storkey [2017], Mao et al. [2002, 2003]. Originally, it is a predator-prey model that describes the dynamics of biological systems with two populations, named predator and prey, which interact together. The preys are assumed to have an unlimited food supply. The perturbed model is defined through a system of two differential equations disturbed by a stochastic noise. 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 is described as:

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

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.

To simulate the trajectories, a Euler-Maruyama discretization scheme is used with $\triangle = 0.01$. The 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.

The objective is to estimate these parameters from discrete observations of the predator population only. For the sake of parametric identifiability, θ_2 is considered as known.

The system (23) is non linear due to the interaction term between predators and preys. The estimation procedure proposed in Section 3 can thus not be applied directly. To deal with the nonlinear interaction term $V_t U_t$, we take advantage of the discrete observations Y of the first coordinate. The estimation procedure only uses the discretized version of the continuous model. Instead of a discretized version of model (23), we consider a discrete model where the interaction term $V_i U_i$ is replaced by $Y_i \times U_i$.

$$\begin{array}{l} V_{t_{i+1}} = V_{t_i} + \Delta_i \left(-\theta_1 V_{t_i} + \theta_2 Y_i U_{t_i}\right) + \sigma \left(W_{1,t_{i+1}} - W_{1,t_i}\right) \\ U_{t_{i+1}} = U_{t_i} + \Delta_i (\theta_3 U_{t_i} - \theta_4 Y_i U_{t_i}) + \sigma \left(W_{2,t_{i+1}} - W_{2,t_i}\right) \\ (V_0,\, U_0) = (0.5,\, 1)\,. \end{array} \tag{24}$$

The estimation procedures are then applied to the model (24). 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 $\gamma^2_{0,\theta,\sigma}(t_i) = \triangle_i \sigma^2 > 0$.

Results are presented in Table 1. The four procedures give overall good results. The two procedures with σ known are very similar, in terms of mean, SE of the estimators and computation time. When σ is also estimated, the second method performs better, especially to estimate the diffusion parameter σ. This could be explained by the fact that the $G^{(2)}$-based procedure relies more on the observations $\{Y_i\}$ than $G^{(1)}$. As the diffusion coefficient is the same for both coordinates, the estimation of σ is more accurate with $G^{(2)}$. Note that the computation time of the procedures estimating the σ's is about hundred times larger for $G^{(1)}$ and twenty times larger for $G^{(2)}$.

True values	θ_1 0.5	θ_3 0.5	θ_4 0.5	σ 0.2	CPU time (s)
σ known, estimation procedure given in Definition 1					
$\widehat{\theta}(\sigma)^{(1)}$	0.64 (0.34)	0.58 (0.15)	0.62 (0.40)	–	158.3
$\widehat{\theta}(\sigma)^{(2)}$	0.65 (0.35)	0.59 (0.17)	0.64 (0.42)	–	158.5
σ unknown, estimation procedure given in Definition 2					
$\widehat{\theta}^{(1)}$ and $\hat{\sigma}^{(1)}$	0.68 (0.35)	0.57 (0.15)	0.65 (0.45)	0.16 (0.01)	11405
$\widehat{\theta}^{(2)}$ and $\hat{\sigma}^{(2)}$	0.68 (0.35)	0.59 (0.16)	0.65 (0.44)	0.19 (0.01)	2438

Table 1: Estimation results for the Lotka Volterra model. From a hundred simulated trajectories, mean and standard error (in brackets) of the estimators given in Definition 1 (σ known) and Definition 2 (σ unknown) obtained with the two functionals $G^{(1)}$ and $G^{(2)}$. Mean individual CPU times are given for each estimation procedure.

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 et al., 1962]. It is defined by a system of two differential equations, that we consider perturbed by a random noise. The two coordinates 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

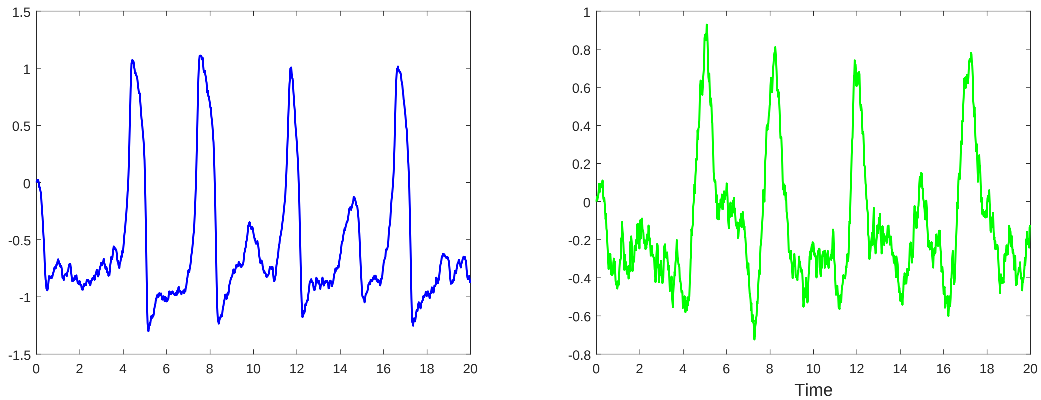

Figure 2: Elliptic FitzHugh-Nagumo model: example of a simulation with the dynamics along time of the membrane potential (left) and the recovery variable (right)

model is defined as:

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

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.

The trajectories are simulated with a Euler-Maruyama scheme with $\triangle = 0.02$. The 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 2. The time scale ϵ is difficult to estimate and is considered as known (and set to 0.1) in the following. The objective is to estimate $\theta = (\gamma, \beta, \sigma_1, \sigma_2)$ from the discrete observations of the first coordinate $(Y_1, \ldots, Y_n)$.

The estimation procedures require a homogeneous linear model. As done for the Lotka-Volterra model, we will propose a linear model close to (25) and apply the estimation procedures to that linearized model. In FHN model, there is a non-linear term (V^3) in the first equation and a non-homogeneous term (β) in the second equation. Both have to be replaced. The trajectory V can be replaced by Y, the observed realization of V. As we will see, several choices can be done to replace the cubic term. For the non-homogeneous term, a strategy is to introduce a new variable R_t assumed to be constant equal to 1. This can be done by increasing the number of coordinates by one and stating $dR_t = 0$.

A first attempt to propose a linear model consists in replacing V^3 by Y^3 in a discretized version of (28). But this reveals to be very unstable, due to the dramatic propagation of the discretization error committed by using Y^3 instead of V^3. An alternative is replacing $V - V^3$ by $(1 - Y^2)V$. The non-homogeneous term is replaced by a new variable Z_t. We then consider the following discretized linear system:

$$\begin{array}{l} V_{t_{i+1}} = V_{t_i} + \triangle_i \frac{1}{\epsilon}\left(\left(1 - Y_i^2\right) V_{t_i} - U_{t_i}\right) + \sigma_1(W_{1,t_{i+1}} - W_{1,t_i}) \\ U_{t_{i+1}} = U_{t_i} + \triangle_i \left(\gamma V_{t_i} - U_{t_i} + \beta R_i\right) + \sigma_2(W_{2,t_{i+1}} - W_{2,t_i}) \\ R_{t_{i+1}} = 1 \\ (V_0,\, U_0,\, R_0) = (0,\, 0,\, 1) \end{array} \tag{26}$$

The estimation procedures are then applied to the model (26) with $d = 3$, $C = \begin{pmatrix} 1 & 0 & 0 \end{pmatrix}$ 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 $\gamma^2_{0,\theta,\sigma}(t_i) = \triangle_i \sigma_1^2 > 0$.

Results are presented in Table 2. When $\sigma = (\sigma_1, \sigma_2)$ is known, the two estimators $\widehat{\theta}(\sigma)^{(1)}$ and $\widehat{\theta}(\sigma)^{(2)}$ are very close. This is not the case when σ is also estimated. This could be explained by the fact that the $G^{(1)}$ criteria uses the 2 dimensional control sequence in a direct way to construct an estimator of $(\sigma_1,\, \sigma_2)$. At the contrary the $G^{(2)}$ criteria only uses the observation to estimate the diffusion, the control sequence being used to estimate the parameters of the unobserved coordinate. As expected, the computational time is much larger when σ is unknown.

	γ	β	σ_1	σ_2	CPU
true value	1.5	0.8	0.1	0.3	time (s)
σ known, estimation procedure given in Definition 1					
$\widehat{\theta}^{(1)}$	1.52 (0.16)	0.82 (0.12)	–	–	84.3
$\widehat{\theta}^{(2)}$	1.53 (0.17)	0.82 (0.12)	–	–	86.3
σ unknown, estimation procedure given in Definition 2					
$\widehat{\theta}^{(1)}$ and $\hat{\sigma}^{(1)}$	1.51 (0.15)	0.81 (0.10)	0.06 (0.02)	0.18 (0.10)	12746
$\widehat{\theta}^{(2)}$ and $\hat{\sigma}^{(2)}$	1.57 (0.13)	0.85 (0.11)	0.04 (0.01)	0.12 (0.12)	2800

Table 2: Estimation results for the elliptic FitzHugh-Nagumo model. From a hundred simulated trajectories, mean and standard error (in brackets) of the different estimators of the estimators given in Definition 1 (σ known) and Definition 2 (σ unknown) obtained with the two functionals $G^{(1)}$ and $G^{(2)}$. Mean individual CPU times are given for each estimation procedure.

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. As explained before, the estimation procedures we propose are robust to the hypoellipticity and we thus apply the same four procedures.

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 the second equation [Pokern et al., 2009]. The model is defined as follows:

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

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

The trajectories are simulated with a Euler-Maruyama scheme with $\triangle = 0.02$. The parameters are set to $D = 4$, $\delta = 0.5$ and $\sigma = 0.5$. The objective is to estimate $\theta = (D, \delta, \sigma)$ from the discrete observations of the first coordinate $(Y_1, \ldots, Y_n)$.

The four estimation procedures are applied. 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 $\mathbf{A}_\theta(t_i) = \begin{pmatrix} 1 + \triangle_i & \triangle_i \\ -\triangle_i D & 1 - \triangle_i \delta \end{pmatrix}$. So $\Gamma_{\theta,\sigma}(0, t_i) = 0$, $\Gamma_{\theta,\sigma}(1, t_i) = \triangle_i^{\frac{3}{2}} \sigma$. We can deduce that $\gamma_{1,\theta}^2(t_i) = \triangle_i^3 \sigma^2 \neq 0$. 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 a invertible covariance matrix of the process. In that sense, it is close to what has recently been proposed by Ditlevsen and Samson [2017].

Results are presented in Table 3. When σ is known, the first approach gives better results than the second. Especially δ is estimated with bias by $\widehat{\theta}^{(2)}$ and not by $\widehat{\theta}^{(1)}$. This is the opposite when σ is also estimated. The second procedure performs better for δ and more convincingly for σ which is accurately estimated. As noted for the elliptic case, the computation time is larger when σ is estimated than when it is fixed at its true value.

	D	δ	σ	CPU
true value	4	0.5	0.5	time (s)
σ known, estimation procedure given in Definition 1				
$\widehat{\theta}(\sigma)^{(1)}$	4.12 (0.44)	0.52 (0.13)	–	119.3
$\widehat{\theta}(\sigma)^{(2)}$	4.06 (0.51)	0.16 (0.07)	–	117.5
σ unknown, estimation procedure given in Definition 2				
$\widehat{\theta}^{(1)}$ and $\hat{\sigma}^{(1)}$	4.11 (0.51)	0.10 (0.12)	0.15 (0.09)	2869
$\widehat{\theta}^{(2)}$ and $\hat{\sigma}^{(2)}$	4.03 (0.50)	0.19 (0.15)	0.50 (0.35)	3067

Table 3: Estimation results for the hypoelliptic Harmonic Oscillator model. From a hundred simulated trajectories, the mean and standard error (in brackets) of the estimators given in Definition 1 (σ known) and Definition 2 (σ unknown) obtained with the two functionals $G^{(1)}$ and $G^{(2)}$. Mean individual CPU times are given for each estimation procedure.

6.2.2 Hypoelliptic FitzHugh-Nagumo

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

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

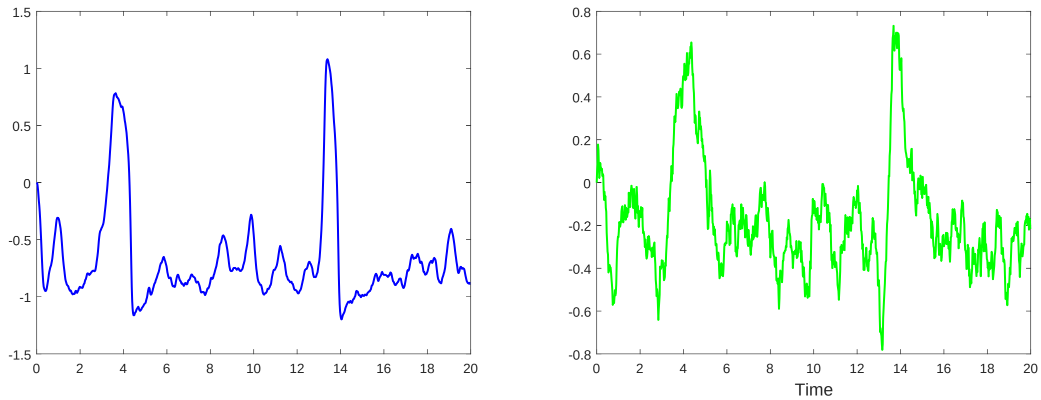

Figure 3: Hypoelliptic FitzHugh-Nagumo model: example of a simulation with the dynamics along time of the membrane potential (left) and the recovery variable (right)

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

The trajectories are simulated with a Euler-Maruyama scheme with $\triangle = 0.02$. The parameters are set to $\epsilon = 0.1$, $\gamma = 1.5$, $\beta = 0.8$, and $\sigma = 0.3$. An example of trajectory for the model (28) is presented in Figure 3. The time scale ϵ is difficult to estimate and as in the elliptic case, we consider it known (and set to 0.1) in the following. The objective is to estimate $\theta = (\gamma, \beta, \sigma)$ from the discrete observations of the first coordinate $(Y_1, \ldots, Y_n)$.

Model (28) is non-linear and time-inhomogenous. Thus, we have to consider a linear version, as done in the elliptic case. For the instability reason evoked in Section 6.1.2, we directly consider the linear discretized model:

$$\begin{array}{l} V_{t_{i+1}} = V_{t_i} + \triangle_i \frac{1}{\epsilon} \left(\left(1 - Y_i^2\right) V_{t_i} - U_{t_i} \right) \\ U_{t_{i+1}} = U_{t_i} + \triangle_i \left(\gamma V_{t_i} - U_{t_i} + \beta R_{t_i} \right) + \sigma (W_{t_{i+1}} - W_{t_i}) \\ R_{t_i+1} = 1 \\ (V_0, U_0, R_0) = (0, 0, 1) \end{array} \tag{29}$$

The four estimation procedures are applied on model (29). Let us give some details on the procedure $\hat{\sigma}^{(2)}$ defined by (19). We have $C = \begin{pmatrix} 1 & 0 & 0 \end{pmatrix}$, $B_\sigma = \begin{pmatrix} 0 \\ \sigma \\ 0 \end{pmatrix}$ and

$$\mathbf{A}_\theta(t_i) = \begin{pmatrix} 1 + \frac{\triangle_i}{\epsilon} \left(1 - Y(t)^2\right) & -\frac{\triangle_i}{\epsilon} & 0 \\ \triangle_i \gamma & 0 & \beta \triangle_i \\ 0 & 0 & 1 \end{pmatrix}.$$

So $\Gamma_{\theta,\sigma}(0, t_i) = 0$, $\Gamma_{\theta,\sigma}(1, t_i) = -\triangle_i^{\frac{3}{2}} \frac{\sigma}{\epsilon}$. We can deduce from that $\gamma_{1,\theta}^2(t_i) = \triangle_i^3 \frac{\sigma^2}{\varepsilon^2} \neq 0$. 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. The most striking and somewhat counter-intuitive result in the hypoelliptic case is the required computational time, in particular when σ needs to be estimated. This computational time is smaller than in 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 S is found. Since the dimension of σ is smaller in the hypoelliptic example than in the elliptic one, it explains the observed difference in the computational time.

true value	γ 1.5	β 0.8	σ 0.3	CPU time (s)
σ known, estimation procedure given in Definition 1				
$\widehat{\theta}(\sigma)^{(1)}$	1.49 (0.10)	0.81 (0.08)	–	147.9
$\widehat{\theta}(\sigma)^{(2)}$	1.49 (0.10)	0.81 (0.08)	–	152.2
σ unknown, estimation procedure given in Definition 2				
$\widehat{\theta}^{(1)}$ and $\hat{\sigma}^{(1)}$	1.46 (0.34)	0.81 (0.28)	0.38 (0.13)	3151
$\widehat{\theta}^{(2)}$ and $\hat{\sigma}^{(2)}$	1.48 (0.12)	0.79 (0.10)	0.29 (0.20)	3822

Table 4: Estimation results for the hypoelliptic FitzHugh-Nagumo model. From a hundred simulated trajectories, mean and standard error (in brackets) of the different estimators given in Definition 1 (σ known) and Definition 2 (σ unknown) obtained with the two functionals $G^{(1)}$ and $G^{(2)}$. Mean individual CPU times are given for each estimation procedure.

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 criteria focusing on estimating the Brownian motion realization given the observation rather than solely on the observation. By doing so, we manage to construct a criteria S_w well defined no matter the structure of B_σ. Another interest of the method is the computational time. The use of the discrete LQ theory allows to avoid the use of MCMC and/or stochastic approximation computationally costful. These two points enlighten the benefit of using control theory for dealing with statistical problems.

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, at the contrary 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 in its elliptic counterpart. A perspective could be to investigate a criteria which allows us to simultaneously estimate θ and σ instead of using our current nested procedure. However, so far we did not manage to find a criteria which fits in the framework of the discrete LQ theory. That 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). Our main challenge would be to extend our method to general nonlinear SDEs.

Acknowledgements

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

References

Y. Aït-Sahalia. Closed-form likelihood expansions for multivariate diffusions. *Ann. Statist.*, 36 (2):906–937, 2008.

V. Bally and D. Talay. The law of the euler scheme for stochastic differential equations i. convergence rate of the distribution function. *Monte Carlo Methods and Applications*, (2): 93–128, 1996.

D.P. Bertsekas. *Dynamic Programming and Optimal Control.* Athena Scientific, 2005.

A. Beskos, O. Papaspiliopoulos, G. O. Roberts, and P. Fearnhead. Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(3):333–382, 2006.

B.M. Bibby and M. Sorensen. Martingale estimating functions for discretely observed diffusion processes. *Bernoulli*, 1:17–39, 1995.

E. Bjørn. MCMC analysis of diffusion models with application to finance. *J. Bus. Econom. Statist.*, 19(2):177–191, 2001.

N. J-B. Brunel and Q. Clairon. A tracking approach to parameter estimation in linear ordinary differential equations. *Electronic Journal of Statistics*, 9:2903–2949, 2015.

P. Cattiaux, J.R. León, and C. Prieur. Estimation for stochastic damping hamiltonian systems under partial observation. ii. drift term. *ALEA*, 11:359–384, 2014a.

P. Cattiaux, J.R. León, and C. Prieur. Estimation for stochastic damping hamiltonian systems under partial observation. i. invariant density. *Stochastic Process. Appl.*, 124:1236–1260, 2014b.

Q. Clairon and N. Brunel. Parameter and state estimation of partially observed linear ordinary differential equations possibly misspecified. Technical report, 2016. submitted.

Q. Clairon and N. Brunel. Optimal control and additive perturbations help in estimating ill-posed and uncertain dynamical systems. 2017. JASA.

F. Comte, C. Prieur, and A. Samson. Adaptive estimation for stochastic damping hamiltonian systems under partial observation. *Stochastic processes and their applications*, 2017.

C.A. Cuenod, B. Favetto, V. Genon-Catalot, Y. Rozenholc, and A. Samson. Parameter estimation and change-point detection from dynamic contrast enhanced MRI data using stochastic differential equations. *Mathematical Biosciences*, 233:68–76, 2011.

S. Ditlevsen and P Greenwood. The morris-lecar neuron model embeds a leaky integrate-and-fire model. *Journal of Mathematical Biology*, 67:239–259, 2013.

S. Ditlevsen and A. Samson. Estimation in the partially observed stochastic morris-lecar neuronal model with particle filter and stochastic approximation methods. *Annals of Applied Statistics*, 2:674–702, 2014.

S. Ditlevsen and A. Samson. Hypoelliptic diffusions: discretization, filtering and inference from complete and partial information. *arXiv:1707.04235v1*, pages 1–33, 2017.

S. Ditlevsen, K.P. Yip, and N.H. Holstein-Rathlou. Parameter estimation in a stochastic model of the tubuloglomerular feedback mechanism in a rat nephron. *Mathematical Biosciences*, 194:49–69, 2005.

S. Donnet and A. Samson. A review on estimation of stochastic differential equations for pharmacokinetic - pharmacodynamic models. *Advanced Drug Delivery Reviews*, 65:929–939, 2013.

G. B. Durham and A. R. Gallant. Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *J. Bus. Econom. Statist.*, 20(3):297–338, 2002.

O. Elerian, S. Chib, and N. Shephard. Likelihood inference for discretely observed nonlinear diffusions. *Econometrica*, 69(4):959–993, 2001.

R. FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 6:445–466, 1961.

W. Gerstner and W. Kistler. *Spiking Neuron Models*. Cambridge University Press, 2002.

A. Gloter. Parameter estimation for a discretely observed integrated diffusion process. *Scand. J. Statist.*, 33(1):83–104, 2006.

A. Golightly and D. J. Wilkinson. Bayesian sequential inference for nonlinear multivariate diffusions. *Stat. Comput.*, 16(4):323–338, 2006.

A. Golightly and D. J. Wilkinson. Bayesian inference for nonlinear multivariate diffusion models observed with error. *Comput Stat & Data Analysis*, 52(3):1674–1693, JAN 1 2008. ISSN 0167-9473. doi: 10.1016/j.csda.2007.05.019.

M.S. Graham and A.J. Storkey. Asymptotically exact inference in differenciable generative models. *arXiv:1605.07826*, page 14, 2017.

A. Iolov, S. Ditlevsen, and A. Longtin. Optimal design for estimation in diffusion processes from first hitting times. *SIAM J. Uncertainty Quantification*, 5:88–110, 2017.

A. C. Jensen, S. Ditlevsen, M. Kessler, and O. Papaspiliopoulos. Markov chain Monte Carlo approach to parameter estimation in the FitzHugh-Nagumo model. *Physical Review E*, 86: 041114, 2012.

M. Kessler. Estimation of an ergodic diffusion from discrete observations. *Scand. J. Statist.*, 24(2):211–229, 1997.

R.S. Lipster and A.N. Shiryaev. *Statistics of random processes I : general theory.* Springer, 2001.

A.J. Lotka. *Elements of Mathematical Biology.* Dover, 1925.

X. Mao, G. Marion, and E. Renshaw. Environmental brownian noise suppresses explosions in populations dynamics. *Stochastic Processes and Their Applications*, 97:95–110, 2002.

X. Mao, S. Sabanis, and E. Renshaw. Asymptotic behaviour of the stochastic lotka-volterra model. *J. Math. Anal. Appl.*, 287:141–156, 2003.

C.F. Martin, S. Sun, and M. Egerstedt. Optimal control, statistics and path planning. *Mathematical and Computer Modelling*, 33:237–253, 2001.

J.C. Mattingly, A. M. Stuart, and D.J. Higham. Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stochastic Process. Appl.*, 101:185–232, 2002.

E. Meeds and M. Welling. Optimization monte carlo: Efficient and embarrassingly parallel likelihood-free inference. In *Advances in Neural Information Processing Systems*, 2015.

J. Nagumo, S. Animoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proc. Inst. Radio Eng.*, 50:2061–2070, 1962.

L. Paninski, Y. Ahmadian, D. G. Ferreira, S. Koyama, K. R. Rad, M. Vidne, J. Vogelstein, and W. Wu. A new look at state-space models for neural data. *Journal of Computational Neuroscience*, 29(1-2, Sp. Iss. SI):107–126, AUG-OCT 2010.

L. Paninski, M. Vidne, B. DePasquale, and D.G. Fereira. Inferring synaptic inputs given a noisy voltage trace via sequential Monte Carlo methods. *Journal of Computational Neuroscience*, 33(1):1–19, 2012.

A.R. Pedersen. A new approach to maximum likelihood estimation for stochastic differential equations based on discrete observations. *Scand. J. Statist.*, 22(1):55–71, 1995.

Y. Pokern, A.M. Stuart, and P. Wiberg. Parameter estimation for partially observed hypoelliptic diffusions. *J. Roy. Stat. Soc. B*, 71(1):49–73, 2009.

L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mischenko. *The Mathematical Theory of Optimal Processes.* Wiley-Interscience, 1962.

M. Pospischil, Z. Piwkowska, T. Bal, and A. Destexhe. Extracting synaptic conductances from single membrane potential traces. *Neurosci.*, 158(2):545–52, 2009.

A. Samson and M. Thieullen. A contrast estimator for completely or partially observed hypoelliptic diffusion. *Stochastic Processes and their Applications*, 122(7):2521–2552, 2012.

E. Sontag. *Mathematical Control Theory: Deterministic finite-dimensional systems.* Springer-Verlag (New-York), 1998.

H. Sørensen. Parametric inference for diffusion processes observed at discrete points in time: a survey. *Int Stat Rev*, 72(3):337–354, 2004.

E. Trelat. *Controle optimal : Theories et applications.* Vuibert, 2005.

F. van der Meulen and M. Schauer. Bayesian estimation of discretely observed multi-dimensional diffusion processes using guided proposals. *submitted*, 2016a.

F. van der Meulen and M. Schauer. Bayesian estimation of incompletely observed diffusions. *submitted*, 2016b.

A.W. van der Vaart. *Asymptotic Statistics.* Cambridge Series in Statistical and Probabilities Mathematics. Cambridge University Press, 1998.

H. Zhang, B. de Saporta, F. Dufour, D. Laneuville, and A. Negre. Quantization and stochastic ccontrol of trajectories of underwater vehicle in bearings-only tracking. *submitted*, 2017.

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:

$$J(u) = x_N^T S x_N + \sum_{i=0}^{N-1} x_i^T Q_i x_i + u_i^T R_i u_i \tag{30}$$

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:

$$x_{k+1} = A_k x_k + B_k u_k \ . \tag{31}$$

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 $i \in [\![0,\, N-1]\!]$, Q_i is positive semi–definite and R_i is positive definite. Then the cost (30) 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:

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