Automatica 97 (2018) 234–247

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Change-point detection for piecewise deterministic Markov processes ☆

Alice Cleynen [a], Benoîte de Saporta [a,b,*]

[a] IMAG, Univ Montpellier, CNRS, Montpellier, France
[b] Inria Bordeaux Sud-Ouest, team CQFD, France

ARTICLE INFO

Article history:
Received 27 September 2017
Received in revised form 26 June 2018
Accepted 11 July 2018

Keywords:
Convergence of numerical methods
Decision making
Discretization
Dynamic programming
Filtering
Markov models
Markov decision processes
Numerical methods
Optimal control
Quantization

ABSTRACT

We consider a change-point detection problem for a simple class of Piecewise Deterministic Markov Processes (PDMPs). A continuous-time PDMP is observed in discrete time and through noise, and the aim is to propose a numerical method to accurately detect both the date of the change of dynamics and the new regime after the change. To do so, we state the problem as an optimal stopping problem for a partially observed discrete-time Markov decision process taking values in a continuous state space and provide a discretization of the state space based on quantization to approximate the value function and build a tractable stopping policy. We provide error bounds for the approximation of the value function and numerical simulations to assess the performance of our candidate policy.

1. Introduction

Piecewise Deterministic Markov processes (PDMPs) are a general class of non-diffusion processes introduced by M. Davis in the 80s (Davis, 1984) covering a wide range of applications from workshop optimization, queuing theory (Davis, 1993), internet networks (Bardet, Christen, Guillin, Malrieu, & Zitt, 2013), reliability (de Saporta, Dufour, & Zhang, 2016), insurance and finance (Bäuerle & Rieder, 2011) or biology (Doumic, Hoffmann, and Krell, & Robert, 2015; Riedler & Thieullen, 2015; Riedler, Thieullen, & Wainrib, 2012) for instance. PDMPs are continuous time hybrid processes with a discrete component called mode or regime and a Euclidean component. The process follows a deterministic trajectory punctuated by random jumps. In the special case where the Euclidean component is continuous, the jumps correspond to a change of regime. For many applications, the regime is not observed and the Euclidean variable is measured in discrete-time, through noise. It may be e.g. a degradation or failure of some component of a system, see Baysse, Bihannic, Gégout-Petit, Prenat, and Saracco (2014) where the Euclidean component is some cool-down time that increases with the degradation of the system, or the cancer cell load of remission patients monitored through proxy tumor markers at regular follow-up blood tests to detect relapse (Abbott & Michor, 2006). The aim of this paper is to propose a fully computable discretization of the value function of the optimal stopping problem corresponding to the change-point detection, and derive error bounds for this approximation. We also use the approximation to build a computable candidate strategy that should be close to optimality. We assess its performance on numerical examples.

The general problem of change-point detection can be seen as an impulse control problem if there are multiple changes in regime. This is a very difficult problem. Although the optimal control of PDMPs has attracted a lot of attention since the 80s, see e.g. Cohen, Madan, Siu, and Yang (2012), Costa and Dufour (2013), Davis (1993), de Saporta and Dufour (2012), de Saporta, Dufour, and Geeraert (2017), Dempster and Ye (1995), Gatarek (1992) and Lenhart (1989), very few works consider such models under partial observations. In Brandejsky, de Saporta, and Dufour (2013), the authors consider an optimal stopping problem for PDMPs

☆ The work was partially supported by Région Languedoc-Roussillon and FEDER under grant *Chercheur(se)s d'Avenir*, project PROMMECE. The material in this paper was partially presented at the SIAM Conference on Control and Its Applications, July 10–12, 2017, Pittsburg, Pennsylvania, USA. This paper was recommended for publication in revised form by Associate Editor Hyeong Soo Chang under the direction of Editor Ian R. Petersen.

* Corresponding author at: Institut Montpelliérain Alexander Grothendieck, CNRS, University Montpellier, France.

E-mail addresses: alice.cleynen@umontpellier.fr (A. Cleynen), benoite.de-saporta@umontpellier.fr (B. de Saporta).

https://doi.org/10.1016/j.automatica.2018.08.011
0005-1098/

where the jump times are perfectly observed and the post-jump locations are observed through noise. They derive the dynamic programming equations of the problem, as well as a numerical approximation of the value function and a computable ϵ-optimal stopping time. In Shi, Elliott, and Chen (2016), the authors consider the problem of state-estimation for a discrete-time and discrete-space process observed only if threshold conditions are reached. In Bian and Jiang (2016) the authors introduce an adaptive dynamic programming algorithm for fully observed linear continuous-time systems. In Bäuerle and Lange (2018), the authors consider a general continuous control problem where both the jump times and post-jump locations are observed through noise. They reduce the problem to a discrete-time Markov Decision Process (MDP) and prove the existence of optimal policies, but provide no numerical approximation of the value function or optimal strategies.

In the present paper, we make a first step towards solving the difficult problem of change-point detection of PDMPs when the jumps are not observed at all. We address the simple case where there is only one change of regime to detect. Unlike in Lin, Loxton, Teo, and Wu (2012) where the authors control the process to avoid it reaching a stopping boundary, the goal of this paper is to detect a jump without controlling the process. The problem can thus be formulated as an optimal stopping problem for PDMPs under partial observations. However, unlike Bäuerle and Lange (2018) and Brandejsky et al. (2013) we do not suppose that the observations are made at or around the jump times. Instead, we suppose that the observations times are deterministic and on a regular grid of step size δ. This enables us to formulate the problem as an optimal stopping problem for a discrete-time partially observed MDP. The equivalent fully-observed MDP for the filter process is still in discrete-time but on an infinite state space. We then propose a two-step discretization of this MDP, following an idea from de Saporta, Dufour, and Nivot (2016). The first step is a time-dependent discretization of the state space of the original PDMP. The second step is a joint discretization of the approximate filter thus obtained together with an approximation of the observation process. Note that unlike Bäuerle and Rieder (2011) or de Saporta, Dufour, Nivot, and (2016), we do not make the assumption that the MDP kernel has a density with respect to some fixed probability measure. We thus obtain a computable approximation of the value function and prove that it converges to the value function of the original change-point detection problem when the discretization parameters are suitably chosen. Based on this approximate value function, we also propose a computable stopping strategy. The optimality of this strategy remains an open problem out of the scope of the present paper, instead we evaluate its performance on various numerical examples.

The main discretization tool we use is optimal quantization. The quantization of a random variable X consists in finding a finite grid such that the projection $\widehat{X}$ of X on this grid minimizes some L_p norm of the difference $X-\widehat{X}$. There exists an extensive literature on quantization methods for random variables and processes. The interested reader may for instance consult Gray and Neuhoff (1998) and Pagès, Pham and Printems (2004b) and the references therein. Quantization methods have been developed recently in numerical probability or optimal stochastic control with applications in finance, see e.g. Bally and Pagès (2003), Bally, Pagès, and Printems (2005), Pagès (1998) and Pagès et al. (2004b).

The paper is organized as follows. In Section 2, we introduce our continuous-time PDMP model as well as the observation model. We define the change-point detection problem as an optimal stopping problem under partial observations and give the equivalent fully observed dynamic programming equations for the filter process. In Section 3, we propose a two-step discretization approach by quantization to numerically solve the optimization problem and build a tractable strategy. In Section 4, we fully detail the practical implementation of our procedure. In Section 5, we investigate the performance of our candidate strategy and compare our approach to moving average and Kalman filtering when possible. A conclusion is given in Section 6. Proofs of our main statements are postponed to Appendix

2. Model and problem setting

In this section, we present the special class of PDMPs we focus on, namely switching flows, define the observation process and state the change-point detection problem as an optimal stopping problem under partial observation. We then derive the filter recursive equation and state the equivalent fully observed optimal stopping problem as well as the corresponding dynamic programming equations.

2.1. Continuous-time PDMP model

We consider the problem of detecting a change-point in the dynamic of a special class of PDMPs which is observed with noise on discrete observation times. The process $\boldsymbol{X}_t = (m_t, x_t, u_t)$ is defined on a state space $\boldsymbol{E} = \mathcal{M} \times \mathbb{K} \times \mathbb{R}_+$, where $\mathcal{M} = \{0, 1, \dots, d\}$ is the finite set of modes, $\mathbb{K}$ is a compact subset of $\mathbb{R}$ representing the position of the process and the third coordinate is the running time since the last jump, needed to ensure the process is Markovian.

Starting from point $(0, x, 0)$ in $\boldsymbol{E}$, which means starting from mode 0 and position x at time 0, the first (and only) jump time T of the process has distribution

$$\mathbb{P}(T > t|\boldsymbol{X}_0 = (0, x, 0)) = e^{-\int_0^t \lambda(s)ds},$$

where λ is a measurable function from $\mathbb{R}_+$ onto $\mathbb{R}_+$ representing the jump intensity. For $t \in [0, T)$, $\boldsymbol{X}_t = (0, x_t, t)$ with $x_t = \Phi_0(x, t)$ for some flow $\Phi_0 : \mathbb{K} \times \mathbb{R}_+ \to \mathbb{K}$ being the solution of an ordinary differential equation.

At time T, the process selects a new mode $i \in \{1, \dots, d\}$ with positive probability π_i.

For $t \geq T$, $\boldsymbol{X}_t = (i, x_t, t - T)$, where $x_t = \Phi_i(x_T, t - T)$ for some flow $\Phi_i : \mathbb{K}\times\mathbb{R}_+ \to \mathbb{K}$ being the solution of an ordinary differential equation.

The assumption that the flows Φ_m do not depend on the running time is made only to keep notation simple and is not actually required, see Example 2. As we will see in the sequel, solving the change-point detection problem is not straightforward, even for such simple dynamics.

We suppose that the observation times $(t_n)_{n\in\mathbb{N}}$ are deterministic and on a regular grid of step size δ until a finite horizon $N\delta$, and that a noisy observation of x_{t_n} is available at each time t_n:

$$Y_n = F(\boldsymbol{X}_{t_n}) + \varepsilon_n = F(x_{t_n}) + \varepsilon_n, \tag{1}$$

where F is a deterministic function, (ε_n) are *i.i.d.* real-valued random variables with density f with respect to the Lebesgue measure on $\mathbb{R}$ and independent from the process $(\boldsymbol{X}_t)$. We assume that Y takes its values in $\mathbb{Y}$, subset of $\mathbb{R}$. We will further denote $\mathbb{X} = \mathcal{M} \times \mathbb{K}$.

2.2. Examples

The following toy examples will be extensively investigated numerically in Section 5. In all examples, the jump intensity is of the form $\lambda(t) = t$ so that the probability to jump gets higher as time goes by. The distribution $\pi = (\pi_1, \dots, \pi_d)$ is the uniform distribution. The distribution of the noise is a centered Gaussian with variance σ^2 truncated at $[-s, s]$ for some $s \in \mathbb{R}$. We investigate several forms for the flow. The function linking the process and the observations will be either $F(x) = x$ or $F(x) = x^{-1}$.