



HAL
open science

Sojourn time estimation in partially observed piecewise deterministic Markov processes - application to myeloma modeling

Alice Cleynen, Benoîte de Saporta, Amélie Vernay

► To cite this version:

Alice Cleynen, Benoîte de Saporta, Amélie Vernay. Sojourn time estimation in partially observed piecewise deterministic Markov processes - application to myeloma modeling. 55^e journées de statistique de la SFdS, Société Française de Statistique, May 2024, Bordeaux, France. hal-04702656

HAL Id: hal-04702656

<https://hal.umontpellier.fr/hal-04702656v1>

Submitted on 19 Sep 2024

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

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

SOJOURN TIME ESTIMATION IN PARTIALLY OBSERVED PIECEWISE DETERMINISTIC MARKOV PROCESSES — APPLICATION TO MYELOMA MODELING

Alice Cleynen^{1,2} & Benoîte de Saporta¹ & Amélie Vernay¹

¹*Univ Montpellier, CNRS, Montpellier, France, alice.cleynen@umontpellier.fr, benoite.de-saporta@umontpellier.fr, amelie.vernay@umontpellier.fr*

²*John Curtin School of Medical Research, The Australian National University, Canberra, ACT, Australia*

Résumé. Nous considérons un problème d'estimation du temps de rechute dans des Processus de Markov Déterministes par Morceaux (PDMP) dont la composante euclidienne est un biomarqueur de substitution pour l'état de patients atteints de myélome. L'une des principales difficultés du problème réside dans le fait que notre processus n'est que partiellement observé, ce que peu de travaux ont pris en compte jusqu'à présent. Nous proposons une méthode basée sur de la régression itérative pour estimer les paramètres d'un PDMP observé en temps discret et avec du bruit. Nous évaluons les performances de notre procédure à travers une étude de simulation et discutons des limites de notre approche.

Mots-clés. Processus de Markov Déterministes par Morceaux, Temps de rechute, Observations partielles, Survie

Abstract. We consider a problem of relapse time estimation in Piecewise Deterministic Markov Processes (PDMPs) whose Euclidean component is a proxy biomarker for the status of myeloma patients. One of the main difficulties of the problem lies in the fact that our process is only partially observed, which few works have considered until now. We provide a method based on iterative regression for estimating the parameters of a PDMP observed in discrete time and through noise. We assess the performances of our procedure through a simulation study and discuss the limitations of our approach.

Keywords. Piecewise Deterministic Markov Processes, Relapse time, Partial observations, Survival

1 Introduction

When monitoring patients suffering from a disease, it is common to measure the level of a specific biomarker as an indicator of the pathological process or the action of a treatment. The biomarker may evolve in a deterministic manner during different phases of the disease and be disrupted by random jumps that indicate a change in the patient's condition. In this work, we propose to model this behaviour through Piecewise Deterministic Markov Processes (PDMPs) (Davis, 1984) in the case of patients followed after developing myeloma. Our aim is to estimate patients' relapse time, or survival time, based on their biomarker levels measured during follow-up visits. One of the difficulties of the problem lies in the fact that our observations are noisy and only partially observed: we have a measurement of a biomarker at discrete visit dates, but its true value at and between dates is unknown. Such frameworks are considered e.g. in (Cleynen and de Saporta, 2018) where the authors propose a numerical scheme to approximate the value function of a

change-point detection problem for a partially observed PDMP in discrete time and through noise. In this paper, we focus on relapse time estimation. We develop an estimation procedure to recover the parameters of our model and evaluate its performance through a simulation study. We discuss the impact of parameter estimation errors on the overall relapse time estimation.

2 The model

2.1 PDMPs

Piecewise Deterministic Markov Processes (PDMPs) introduced by Davis in the 80's are a general class of stochastic processes, including almost all non-diffusion models found in applied probability (Davis, 1984). These continuous-time processes are used to describe deterministic motions punctuated by random jumps.

Let $(X_t)_{t \geq 0}$ be a PDMP defined on a state space $E \subset \mathbb{R}^d$. The trajectories of $(X_t)_{t \geq 0}$ are determined by the behavior of the process between jumps, as well as when and where the jumps occur. These aspects are described by a flow Φ , a jump intensity λ and a Markov kernel Q , respectively. The flow $\Phi: E \times \mathbb{R}_+ \rightarrow E$ is a continuous function satisfying the semi-group property: $\forall x \in E, \forall t, s \in \mathbb{R}_+, \Phi(x, t+s) = \Phi(\Phi(x, t), s)$. Starting from $x \in E$, $\Phi(x, t)$ gives the position of the process after some time t if no jump has occurred (see Figure 1).




Figure 1 – Starting from x at time 0, the process follows its flow up to time t and ends up at $\Phi(x, t)$, assuming no jump has occurred. Going on up to time $t + s$ is the same as starting from $\Phi(x, t)$ and following the flow for a time s .

The process can jump deterministically or randomly. Deterministic jumps occur when the flow reaches the boundary ∂E of E . Given a starting point $x \in E$, this happens after a time $t^*(x) = \inf \{t > 0: \Phi(x, t) \in \partial E\}$. Random jumps are governed by the jump intensity $\lambda: E \rightarrow \mathbb{R}_+$ — also known as the hazard rate — which is a measurable function such that $\forall x \in E, \exists \varepsilon > 0: \int_0^\varepsilon \lambda(\Phi(x, s)) ds < \infty$. That is, jumps cannot occur instantaneously (and therefore there cannot be several jumps at the same time). The jump times of a PDMP are obtained by taking the minimum between deterministic jumps and stochastic ones. Given a starting point $x_0 \in E$, for all $t \in \mathbb{R}_+$, the first jump time T_1 satisfies

$$\mathbb{P}_{X_0=x_0}(T_1 > t) = \begin{cases} \mathbb{P}_{X_0=x_0}(T_1 > t) = e^{-\int_0^t \lambda(\Phi(x_0, s)) ds} & \text{if } t < t^*(x_0) \\ 0 & \text{if } t \geq t^*(x_0) \end{cases}. \quad (1)$$

For both deterministic and random jumps, the new location of the PDMP is drawn from the Markov kernel $Q: \bar{E} \times \mathcal{B}(\bar{E}) \rightarrow E$, where $\mathcal{B}(\bar{E})$ is the set of Borels of \bar{E} . When the process starts from $x \in \bar{E}$, we have that $\forall A \in \mathcal{B}(\bar{E}), Q(x, A) = \mathbb{P}(X_{T_1} \in A \mid X_{T_1^-} = x)$, where T_1^- denotes the time just before the first jump. For Q to be a Markov kernel, we need $x \mapsto Q(x, A)$ to be measurable $\forall A \in \mathcal{B}(\bar{E})$ and $A \mapsto Q(x, A)$ to be a probability density function $\forall x \in \bar{E}$. The Markov kernel Q satisfies $\mathbb{P}(X_t = x \mid X_{t^-} = x) = 0, \forall t \in \mathbb{R}_+$. In other words, each jump must involve a real change of location.

It is common practice to separate the state space E into a hybrid one made up of a discrete component and a continuous one, such that $X_t = (m_t, \zeta_t) \in E \subset M \times \mathbb{R}^d$, where m_t corresponds to a discrete mode and ζ_t to a continuous variable. Furthermore, the state space can be specific to each mode: $\forall m \in M, E_m \subset \mathbb{R}^{d_m}$. The mode-specific flow Φ_m is such that $\forall m \in M, \Phi_m: E_m \times \mathbb{R}_+ \rightarrow E_m$ and $\Phi((m, \zeta), t) = (m, \Phi_m(\zeta, t))$.

2.2 Our model

We consider subjects undergoing medical follow-up after developing myeloma. During medical visits, their serum-M protein levels, a proxy for the progression of their disease, are measured. At the start of monitoring they are administered a treatment, the effect of which is to reduce the serum-M protein level exponentially. If the level falls below a certain fixed threshold ζ_0 , the patient is considered to be in remission. In the event of a relapse, the serum level rises again exponentially. The horizon H of follow-up is different for each patient. We now link the notations introduced in 2.1 with our specific model. There are three possible modes for the subjects in the study. They can either be sick under treatment ($m = -1$), sick without treatment ($m = 1$) or in remission ($m = 0$) — for simplicity, we assume that the process characteristics in remission mode are the same with or without treatment and do not differentiate the two cases. We thus have $M = \{-1, 0, 1\}$ and $E_{-1} =]\zeta_0, +\infty[$, $E_0 = \{\zeta_0\} \times \mathbb{R}_+$ and $E_1 = [\zeta_0, +\infty[$. In mode $m = 0$, a time variable $u \in \mathbb{R}_+$ is added to the state space to allow more flexibility in the jump intensity while ensuring the Markov property holds. Under treatment, the serum spike decreases exponentially with a slope $v_{-1} < 0$. During relapse, it increases exponentially with a slope $v_1 > 0$. For all $\zeta \in \mathbb{R}$ and for all $u, t \in \mathbb{R}_+$ we have

$$\begin{cases} \Phi_{-1}(\zeta, t) = \zeta e^{v_{-1}t}, \\ \Phi_0(\zeta, t, u) = \zeta = \zeta_0, \\ \Phi_1(\zeta, t) = \zeta e^{v_1t}. \end{cases} \quad (2)$$

In mode $m = -1$, the jump occurs when the subject reaches the fixed remission threshold ζ_0 . This is therefore a deterministic jump at the boundary and $\lambda_{-1}(\zeta) = 0$. The jump time $t_{-1}^*(\zeta)$ is the solution of $\Phi_{-1}(\zeta, t) = \zeta_0$. That is, $t_{-1}^*(\zeta) = \frac{1}{v_{-1}} \log(\frac{\zeta_0}{\zeta})$. With the additional time variable in mode $m = 0$, we have $\Phi_0((\zeta_0, u, t), t) = (\zeta_0, u + t)$ and $t_0^*(\zeta, u) = +\infty$. The jump intensity $\lambda_0 > 0$ is unknown. For practical reasons explained in Section 4, we decide to approximate it with a Weibull distribution. In our model, we consider that once the process reached mode $m = 1$, no more jump can occur.

For the mode-specific Markov kernels, we have

$$\begin{cases} Q_{-1}(m', \zeta', u' | \zeta) = \mathbb{1}_{\zeta=\zeta_0} \times \mathbb{1}_{\zeta'=\zeta_0} \times \mathbb{1}_{m'=0} \times \mathbb{1}_{u'=0} \\ Q_0(m', \zeta' | \zeta, u) = \mathbb{1}_{\zeta=\zeta_0=\zeta'} \times \mathbb{1}_{m'=1} \end{cases} \quad (3)$$

3 Real data

Our data comes from a study carried out by the Inter-Groupe Francophone du Myélome (IFM) in 2009. The author consider the effect of lenalidomide, bortezomib and dexamethasone (RVD) therapy alone versus RVD therapy plus autologous stem cell transplantation on disease progression (Attal et al., 2017). About 700 patients with newly diagnosed myeloma were randomly divided into two groups and followed up after receiving their respective therapy. Their serum M-protein levels were measured at different frequencies depending on the phase of the trial. Patients may remain in remission, suffer a relapse or leave the study for various reasons. The length of follow-up therefore varies from one individual to another. In this work, we are interested in estimating the relapse time of patients. For the time being, for the sake of simplicity we do not take into account the difference between patient groups, nor any other covariate, and we consider relapse times in a general way. The raw data had been preprocessed to remove observations unsuitable for model fitting. The detailed preprocessing procedure is given in Appendix A. The post-processed dataset consists of 479 serum M protein peak trajectories over time.

4 Estimating model parameters

In this section, we explain the estimation procedure for the parameters of our model based on the data. This involves first estimating the process jump times, and then finding the parameters of the relapse time distribution. Note that our observations here are not in continuous time, so the jump times are hidden (see Figure 2). In what follows, we let $t_0 := T_1$ the time of the first jump, from mode $m = -1$ to mode $m = 0$ and $T_0 := T_2$ the second jump, if any, from mode $m = 0$ to mode $m = 1$.

4.1 Jump time estimation

Our estimation method for t_0 and T_0 is an iterative optimisation process based on regression. It is illustrated in Figure 2. Let $(X_t)_{t \geq 0}$ be a PDMP as defined in 2.2 and let $(X_k)_{k \in \mathbb{N}} = (X_{d_k})_{k \in \mathbb{N}}$ be the process at the observation dates $(d_k)_{k \in \mathbb{N}}$ that generates our trajectories. The observations are defined as $Y_k = F(X_k)e^{\varepsilon_k}$, where $\varepsilon_k \sim \mathcal{N}(0, \sigma^2)$ is a Gaussian noise and where $F: E \rightarrow \mathbb{R}_+$ is a function that returns the second component of the PDMP. Hence, $Y_k = \zeta_k e^{\varepsilon_k}$. We use a multiplicative noise both to match the exponential growth and decay of biomarker level and to simplify the estimation procedure described hereafter. We start with the estimation of t_0 . Let y_0, y_1, \dots, y_j be the first j values of an N -length trajectory of M-protein levels recorded at dates d_0, d_1, \dots, d_j , respectively. The biomarker level has an exponential form, so we use least squares to fit a linear function $f: x \mapsto ax + b$ to the logarithm of our data, where $a^* = \widehat{v_{-1}}$ and $e^{b^*} = \widehat{x_0}$ are the optimal solutions of the problem. We use a logarithmic transformation to prevent errors at the beginning of the trajectory from having too much weight on the overall error.

We then estimate t_0 as the solution for t of $\widehat{x_0} e^{\widehat{v_{-1}} t} = \zeta_0$ (see Figure 2). This gives us an approximation of the jump time from $m = -1$ to $m = 0$ and we compute a general regression error as the sum of two errors: one between the points falling before $\widehat{t_0}$ and the fitted curve, and the other on the remaining part of the trajectory. Note that all the above estimates depend on j , which is omitted for clarity. This process is repeated for $j \in \{3, \dots, N\}$ until a stopping criterion is met. This results in optimal estimates for the entry time into remission t_0 and the slope v_{-1} in mode $m = -1$. Note that $\widehat{v_{-1}}$ is only used to estimate t_0 and will not be used to estimate the relapse time afterwards. Details of the estimation procedure can be found in Algorithm 1.

We can then use the same process again on the remaining part of the trajectory¹ — that is, on $(y_k)_{k=\widehat{t_0}, \dots, N}$ — to obtain $\widehat{T_0}$ and $\widehat{v_1}$. It is important to note that if the trajectory is very flat, it is not relevant to look for a jump time, plus the algorithmic minimization could fail due to numerical instability. In such cases, we assume that no change in mode occurred. This happens mainly when subjects do not relapse within their follow up time. They are considered "censored subjects" and are discussed hereafter.

¹ From a computational point of view, if $k_{\widehat{t_0}}$ is the index of time $\widehat{t_0}$ in the vector $d \in \mathbb{R}^N$ of visit dates associated with the trajectory, we apply Algorithm 1 with $d_{k_{\widehat{t_0}}:N}$ in reverse order of its elements as vector of visit dates. Same for the corresponding vector of spikes.

Algorithm 1: Jump time estimation

input: $d \in \mathbb{R}^N$ vector of visit dates, $y \in \mathbb{R}^N$ vector of spikes at visit dates, $\zeta_0 \in \mathbb{R}$ theoretical threshold for remission mode
init: $\Delta_{\text{tmp}} = \Delta_{\text{min}} = \infty$
for $j = 3, \dots, N$ **do**
 $d_{\text{tmp}} = d_{0:j}$ //slicing of the first j coordinates of d
 $y_{\text{tmp}} = y_{0:j}$
 find \hat{a} and \hat{b} optimal solutions when fitting $f : x \mapsto ax + b$ to $\log(y_{\text{tmp}})$ using least squares
 $t_{0\text{tmp}} = (\log(\hat{\zeta}_0)) - \log(\hat{b})/\hat{a}$ //solve $\hat{b}e^{t\hat{a}} = \zeta_0$ for t
 $k = |i : d_i \leq t_{0\text{tmp}}|$
 $n_1 = \|y_{0:[k]} - \hat{b}e^{-d_{0:[k]} \times \hat{a}}\|_2^2$ //error between the first $[k]$ points and the fitted curve
 $n_2 = \|y_{[j]+1:N} - \zeta_0\|_2^2$ //error on the remaining part of the trajectory
 $\Delta = \Delta_{\text{tmp}}$
 $\Delta_{\text{tmp}} = n_1 + n_2$
 if $\Delta_{\text{tmp}} \leq \Delta_{\text{min}}$ **AND** $t_{0\text{tmp}} > 0$ **then**
 $t_0 = t_{0\text{tmp}}$
 $\Delta_{\text{min}} = \Delta_{\text{tmp}}$
 $a^* = \hat{a}$
 $b^* = \hat{b}$
 if $\Delta_{\text{tmp}} > \Delta$ **AND** $t_0 > 0$ **AND** $j > 15$ **then break**
return t_0, a^*, b^*




Figure 2 – Illustration of the estimation process. The true PDMP starts in mode $m = -1$ from an initial point x_0 and follows a deterministic trajectory along its flow unit the first jump occurs at time t_0 . The mode changes to become $m = 0$ and the flow equals ζ_0 until a new jump occurs at time T_0 . The mode switches to $m = 1$ and the trajectory rises exponentially along $\Phi_1(\zeta, t)$. The black crosses represents the observations and we seek to recover the model parameters.

4.2 Survival time before relapse

Having calculated \widehat{t}_0 and \widehat{T}_0 for every trajectory, we now have access to the survival times of subjects. That is, the number of days between the remission start and the beginning of the relapse, if any, or the follow up time otherwise. Following the terms of survival analysis, an *event* is defined as the occurrence of a relapse. A patient is considered *censored* if no event has occurred until the end of its follow up. The survival function $S(t) = \mathbb{P}(T_0 - t_0 > t)$ gives us the probability that a patient remain in remission beyond a time t after remission entry, and we seek to recover its parameters. We assume that the hazard rate — the event rate at time t conditional on survival up to time t or later — increases with time. This leads us to choose a Weibull distribution to model our relapse time, as is conventionally done in such cases in survival analysis. We will see in Section 5 that this assumption is quite reasonable. The probability density function of the Weibull distribution is given by

$$f(x) = \left(\frac{\alpha}{\beta}\right) \left(\frac{x}{\beta}\right)^{\alpha-1} e^{-\left(\frac{x}{\beta}\right)^\alpha}, \quad (4)$$

for $x > 0$ and where α is a shape parameter and β is a scale parameter. Its hazard function is

$$h(x) = \left(\frac{\alpha}{\beta}\right) \left(\frac{x}{\beta}\right)^{\alpha-1}. \quad (5)$$

Note that a shape parameter $\alpha < 1$ (resp. $\alpha > 1$) means that the failure rate decreases (resp. increases) over time. If $\alpha = 1$, this rate is constant. To estimate the parameters of the Weibull distribution, we fit a parametric survival regression model with the survival time as a response variable together with an censoring indicator. This gives us the estimations $\widehat{\alpha}$ and $\widehat{\beta}$ of the shape and scale parameters, respectively.

5 Simulation study

In this section, we assess the performances of the estimation procedure described in Section 4 on simulated data. The trajectories are generated from the model presented in 2.2. We compare the estimates with the ground-truth and discuss the impact of jump times errors on relapse time estimation. The parametric survival regressions are performed with the R `survival` package (Therneau, 2023).

5.1 Simulation process

We use Algorithm 2 to simulate patient trajectories. The input parameters are chosen as follows. When patients enter the clinical trial, they can have a serum M-protein level anywhere between a remission threshold and some upper threshold with equiprobability. This leads us to opt for a uniform distribution to generate the first point $x_0 = \zeta_{t=0}$ of a trajectory. Similarly, patients follow-up may end at any time within a certain date range since they did not enter the study at the same time. We thus choose H uniformly distributed as well. Based on real data analysis, we pick $l_0 = 15$, $u_0 = 55$ and $l_H = 900$, $u_H = 1900$ as lower and upper bounds for $\zeta_{t=0}$ and H , respectively. To select the model parameters α , β , v_{-1} and v_1 for the simulation, we apply the estimation procedure on our real data. We take v_{-1} to be the average of all the estimated slopes in mode $m = -1$. The same is done for v_1 but only considering trajectories for which a relapse is predicted. This gives us $v_{-1} = -0.046$ and $v_1 = 0.012$. The parametric survival regression gives $\widehat{\alpha} = 4.69$ and $\widehat{\beta} = 1650$ which we choose as our α and β inputs. Note that $\widehat{\alpha} > 1$, which is consistent with the fact that the risk of a relapse increases over time. Figure 3 shows the fitted survival curve and the overlapping Weibull survival function plotted with $\widehat{\alpha}$ and $\widehat{\beta}$. The Weibull distribution seems to fit the data fairly well, although its survival curve is slightly shifted compared to the survival curve adjusted on the data. Both the shape and the scale estimates are obtained with significant p-values (2×10^{-16}). On clinical trials, visit frequency often varies during follow-up depending on the stage of the study and the patient status. For the sake of

simplicity, we only consider fixed time intervals $\delta \in \mathbb{N}$ between visits. However, we study several different values for δ to assess the impact of visit frequency on estimation errors. Finding an optimal frequency is an important point: visits must be frequent enough to detect a relapse as early as possible, while avoiding a burdensome and restrictive monitoring for patients. We consider values of δ in $\{10, 20, 30, 40, 50, 60\}$. The remission threshold ζ_0 is set to 1 according to medical criteria. For practical reasons, we use additive noise to simulate our data, although we have assumed in 4.1 that the noise is multiplicative. This provides trajectories that more closely resemble those of the real data. We set the standard deviation σ to 1.

Algorithm 2: Simulation of one trajectory from a PDMP

```

input:  $l_0, u_0 \in \mathbb{R}, l_H, u_H \in \mathbb{R}$ , lower and upper bounds for starting point and follow up time
          distribution,  $\alpha, \beta \in \mathbb{R}$  shape and scale parameters for the Weibull distribution,  $v_{-1}, v_1 \in \mathbb{R}$ 
          slopes for mode  $m = -1$  and  $m = 1$ ,  $\delta \in \mathbb{N}$  number of days between two visit dates,  $\zeta_0 \in \mathbb{R}$ 
          theoretical threshold for remission mode

init:  $\zeta_{t=0} \sim \mathcal{U}_{[l_0, u_0]}, H \sim \mathcal{U}_{[l_H, u_H]}$ 
 $t_0 \leftarrow (\log(\zeta_0) - \log(\zeta_{t=0}))/v_{-1}$ ;  $w \sim \mathcal{W}(\alpha, \beta)$ 
 $T_0 \leftarrow w + t_0$ ;  $c \leftarrow \mathbb{1}_{\{H \leq T_0\}}$  //c censoring indicator
 $\delta_{\text{end}} \leftarrow \lfloor H/\delta \rfloor$  //last visit date
for  $k = 0, \dots, \delta_{\text{end}}$  do //visit dates at regular time intervals until H
   $d_k \leftarrow k\delta$ 
   $m_k \leftarrow -\mathbb{1}_{\{d_k < t_0\}} + \mathbb{1}_{\{d_k > T_0\}} \mathbb{1}_{\{c\}}$ 
  if  $m_k = -1$  then
     $\zeta_k \leftarrow \Phi_{-1}(\zeta_0, d_k)$ 
  else if  $m_k = 0$  then
     $\zeta_k \leftarrow \zeta_0$ 
  else
     $\zeta_k \leftarrow \Phi_1(\zeta_0, d_k - T_0)$ 
   $y_k = \zeta_k + \mathcal{N}(0, \sigma^2)$ 
return  $y = (y_k)_k$ 

```




Figure 3 – Survival curve fitted on real data and 95% confidence interval. Crosses signify censored events. The black dashed line shows the Weibull survival function with parameters estimated from survival regression.

5.2 Results

We evaluate our method through 100 repetitions of 500 PDMP trajectories.

Figure 4 presents the distributions of absolute errors on \hat{t}_0 and \hat{T}_0 depending on visit frequency. Unsurprisingly, with longer time intervals between visits, the mean error and the variability increase for both jump times. For $\delta \leq 30$, the mean error on \hat{t}_0 is less than the one on \hat{T}_0 , whereas the opposite occurs for larger values of δ . Note that for $\delta = 60$, the average number of days of error is approximately equal to the time interval itself: the estimate \hat{T}_0 is one visit apart from the actual jump time T_0 . As a complement to Figure 4, Table 1 presents the mean number days of error on relapse time estimation $|(T_0 - t_0) - (\hat{T}_0 - \hat{t}_0)|$ depending on visit frequency. One can see that the errors on \hat{t}_0 and \hat{T}_0 do not compensate: roughly speaking, the mean error on relapse time estimation is the sum of the two, especially for larger values of δ . Figure 5 shows the distributions of relative errors on Weibull shape and scale parameter estimates. The average error on the scale parameter β increases with the time between visits. This is fairly consistent with the results in Table 1, since the increase in δ degrades the estimate of relapse time and spreads out its distribution, leading to a less accurate estimation. Conversely, the average relative error on $\hat{\alpha}$ decreases with δ , but this behaviour is less clear to us to interpret. However, it can be said that a compromise needs to be found on the frequency of visits for the estimation of the distribution parameters. The curves of the probability density functions of the Weibull distributions obtained with the mean estimates are shown in Figure 6, together with the ground-truth density used to simulate the trajectories. The modes of the distributions are closer to the true mode for smaller values of δ . However, the estimated distributions are less spread out than the true one, whatever the visit frequency. On the whole, we tend to underestimate the relapse time. Note however that of all the trajectories simulated, about 66% were censored, which adds a significant difficulty to the problem. Table 2 provides the average proportions of false censoring and false relapse predictions. A false censoring occurs when the T_0 estimates falls after the time horizon H . Such errors tend to appear more often as the visit frequency decreases: the longer we wait before checking a patient again, the more likely we are to miss a relapse. Our estimation procedure never predicts a relapse when there is none. This is probably due to the low noise level we have chosen for the simulations, and it would be interesting to consider noisier trajectories to see if this behaviour is maintained.




Figure 4 – Boxplots of absolute number of days of error on estimated jump times \hat{t}_0 and \hat{T}_0 over 100 repetitions with 500 trajectory samples, depending on visit frequency (in days). The distributions are only calculated on trajectories for which relapse occurred and is correctly predicted.

Visit frequency	$ (T_0 - t_0) - (\widehat{T}_0 - \widehat{t}_0) $
10	10.745 ± 1.156
20	16.494 ± 1.738
30	24.429 ± 2.931
40	38.650 ± 4.552
50	62.341 ± 6.471
60	86.816 ± 5.793

Table 1 – Average absolute number of days error and standard deviation on estimated survival times over 100 repetitions with 500 trajectory samples, depending on visit frequency. The mean is only calculated on trajectories for which relapse occurred and is correctly predicted. Values are rounded to the nearest 10^{-3} .

Visit frequency	False censoring	False relapse
10	0.092	0.0
20	0.096	0.0
30	0.101	0.0
40	0.109	0.0
50	0.124	0.0
60	0.144	0.0

Table 2 – Average proportion of false censoring and false relapse prediction over 100 repetitions with 500 trajectory samples, depending on visit frequency. Values are rounded to the nearest 10^{-3} .

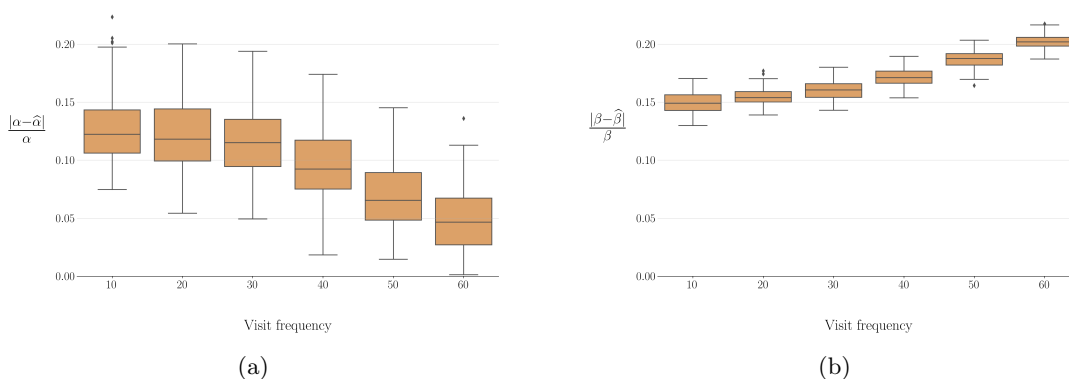


Figure 5 – Boxplots of relative error on Weibull (a) shape and (b) scale parameter estimates over 100 repetitions with 500 trajectory samples, depending on visit frequency (in days).

6 Conclusion and further work

We have presented a proof of concept for the estimation of relapse time in PDMPs with noisy and partially observed trajectories. The results are encouraging, but there is still considerable room for improvement. For the moment, we have only considered a simplistic simulation framework and it would seem appropriate to choose the initial parameters differently and to control the censoring rate. We have chosen to fully exploit the assumptions of our model. It would be interesting to see what happens if one or more of them are no longer verified. What would happen if our noise were additive rather than multiplicative? Or if the growth and decay of the process were not exponential? We could also consider the case where the slopes are no longer fixed but chosen at random. Finally, it would be interesting to compare our estimation procedure with methods for which no or almost no assumptions about the model are required, such as moving average methods or hidden Markov models. This comparison is planned as future work.

7 Acknowledgement

We acknowledge the support of European Union’s Horizon 2020 research and innovation program (<https://marie-sklodowska-curie-actions.ec.europa.eu>, Marie Skłodowska-Curie grant agreement No 890462, to Alice Cleyen), and of the French National Research Agency (ANR), under grant ANR-20-CHIA-0001 (<https://anr.fr/>, project CAMELOT, to Amélie Vernay).




Figure 6 – Curves of the Weibull probability density functions with average parameters estimates over 100 repetitions with 500 trajectory samples, depending on visit frequency (in days). The dashed curve shows the ground-truth density.

Bibliography

Michel Attal, Valerie Lauwers-Cances, Cyrille Hulin, Xavier Leleu, Denis Caillot, Martine Escoffre, Bertrand Arnulf, Margaret Macro, Karim Belhadj, Laurent Garderet, et al. Lenalidomide, bortezomib, and dexamethasone with transplantation for myeloma. *New England Journal of Medicine*, 376(14):1311–1320, 2017.

Alice Cleynen and Benoîte de Saporta. Change-point detection for piecewise deterministic markov processes. *Automatica*, 97:234–247, 2018. ISSN 0005-1098. doi: <https://doi.org/10.1016/j.automatica.2018.08.011>. URL <https://www.sciencedirect.com/science/article/pii/S0005109818304011>.

MHA. Davis. Piecewise-deterministic Markov processes: a general class of nondiffusion stochastic models. *J. Roy. Statist. Soc. Ser. B*, 46(3):353–388, 1984. With discussion.

Terry M Therneau. *A Package for Survival Analysis in R*, 2023. URL <https://CRAN.R-project.org/package=survival>. R package version 3.5-7.

A Preprocessing real data

The raw data had been preprocessed to remove observations unsuitable for model fitting. Apart from observations for which the quantity of interest is missing or incomplete, the following data have been deleted, in this order:

1. first observations whose spike is lower than that of the second; the process is repeated iteratively until the first spike is higher than the next one
2. trajectories whose first spike is lower than a threshold of $5\mu\text{g/L}$ (value at which the level is considered negligible)
3. trajectories with less than two spikes below $5\mu\text{g/L}$: we assume that the associated subjects never reached remission
4. observations with spike equalling $0\mu\text{g/L}$ surrounded by spikes above $5\mu\text{g/L}$: we assume these correspond to insignificant isolated zeros
5. trajectories with less than 10 visits.