Optimal stopping for change-point detection of Piecewise Deterministic Markov Processes

Alice Cleynen, Benoîte de Saporta
Institut Montpelliérain Alexander Grothendieck
CNRS, Univ. Montpellier, France

Outline

Motivation

Change-point detection problem

Numerical approximation

Numerical results

Conclusion and perspectives

Piecewise deterministic Markov processes

Davis (80's)

General class of non-diffusion dynamic stochastic hybrid models: deterministic motion punctuated by random jumps.

Starting point

$$X_0 = (m, x)$$

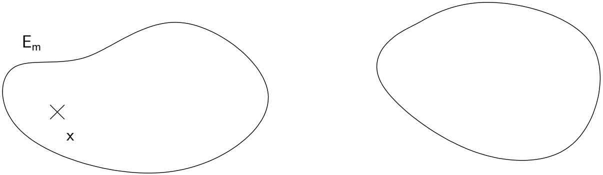

Piecewise deterministic Markov processes

Davis (80's)

General class of non-diffusion dynamic stochastic hybrid models: deterministic motion punctuated by random jumps.

X_t follows the deterministic flow until the first jump time $T_1 = S_1$

$$X_t = \big(m, \phi_m(x,t)\big), \quad \mathbb{P}_{(m,x)}(S_1 > t) = \mathrm{e}^{-\int_0^t \lambda_m\left(\phi_m(x,s)\right)ds}$$

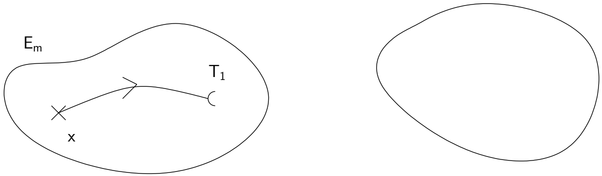

Piecewise deterministic Markov processes

Davis (80's)

General class of non-diffusion dynamic stochastic hybrid models: deterministic motion punctuated by random jumps.

Post-jump location (m_1, x_{T_1}) selected by the Markov kernel

$$Q_m(\phi_m(x, T_1), \cdot)$$

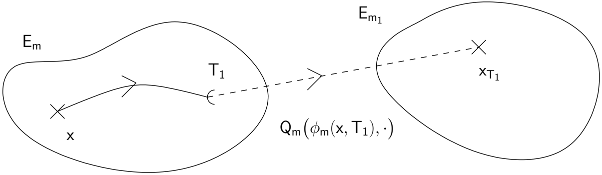

Piecewise deterministic Markov processes

Davis (80's)

General class of non-diffusion dynamic stochastic hybrid models: deterministic motion punctuated by random jumps.

X_t follows the flow until the next jump time $T_2 = T_1 + S_2$

$$X_{T_1+t} = \big(m_1, \phi_{m_1}(x_{T_1}, t)\big), \quad t < S_2$$

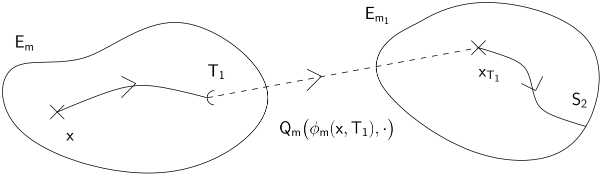

Piecewise deterministic Markov processes

Davis (80's)

General class of non-diffusion dynamic stochastic hybrid models: deterministic motion punctuated by random jumps.

Post-jump location (m_2, x_{T_2}) selected by Markov kernel

$$Q_{m_1}(\phi_{m_1}(x_{T_1}, S_2), \cdot) \ldots$$

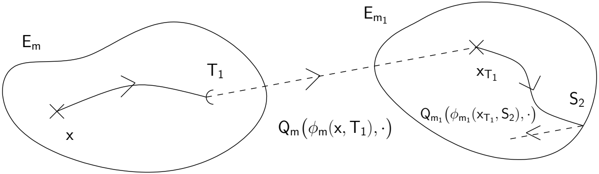

Applications

Applications of PDMPs

Engineering systems, operations research, management science, economics, internet traffic, dependability and safety, neurosciences, biology, ...

- mode: nominal, failures, breakdown, environment, number of individuals, response to a treatment, ...
- Euclidean variable: pressure, temperature, time, size, potential, protein level, ...

Impulse control problem

Impulse control

Select

- intervention dates
- new starting point for the process at interventions

to minimize a cost function

- repair a component before breakdown
- change treatment before relapse
- ...

[CD 89], [Davis 93], [dSDZ 14], ...

If the jump times are not observed?

- [BdSD 12] Optimal stopping
  - jump times observed
  - post-jump locations observed through noise

  Numerical approximation of the value function and ϵ-optimal stopping time

- [BL 17] Continuous control
  - jump times and post-jump locations observed through noise

  Optimality equation, existence of optimal policies

If the jump times are not observed?

- [BdSD 12] Optimal stopping
  - jump times observed
  - post-jump locations observed through noise

  Numerical approximation of the value function and ϵ-optimal stopping time

- [BL 17] Continuous control
  - jump times and post-jump locations observed through noise

  Optimality equation, existence of optimal policies

No information on the jump times $\Rightarrow$ very difficult problem

Change-point detection

Simplest special case

- only one jump of the mode variable
- discrete noisy observations of the continuous variable on a regular time grid

Optimal stopping = Change-point detection

Aim: numerical approximation to

- detect the change-point at best (not too early/late)
- estimate the new mode after the jump

Simple PDMP model

- State space $E \times \mathbb{R} = \{0, 1, \ldots, d\} \times \mathbb{R} \times \mathbb{R}$: mode, position, time
- Starting point $X_0 = (0, x, 0)$, flow Φ_0
- time-dependent Jump intensity $\lambda_0(x, u) = \lambda(u)$
- Jump kernel: position and time continuous, switch to mode i with probability p_i

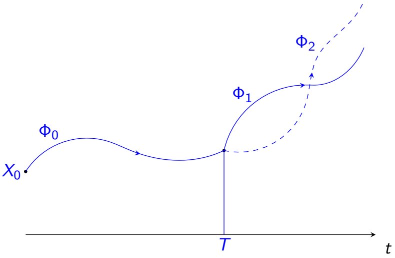

Observations

- Observation times $t_n = \delta n$
- Noisy observations of the positions $Y_n = F(x_{t_n}) + \epsilon_n$

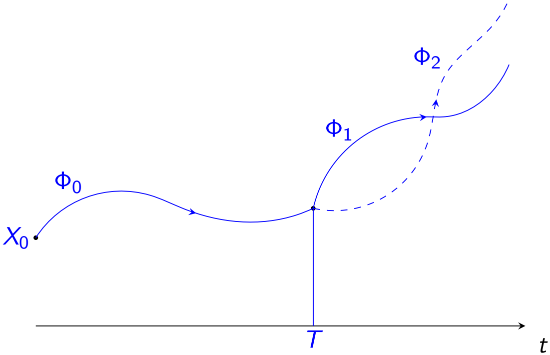

Observations

- Observation times $t_n = \delta n$
- Noisy observations of the positions $Y_n = F(x_{t_n}) + \epsilon_n$

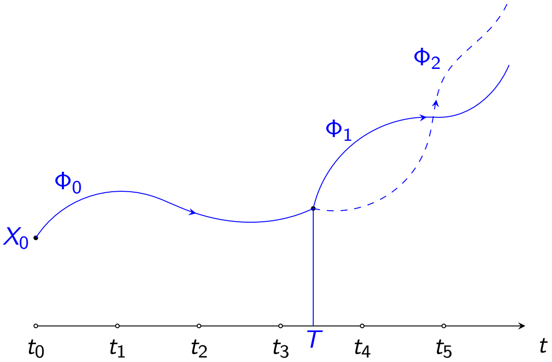

Observations

- Observation times $t_n = \delta n$
- Noisy observations of the positions $Y_n = F(x_{t_n}) + \epsilon_n$

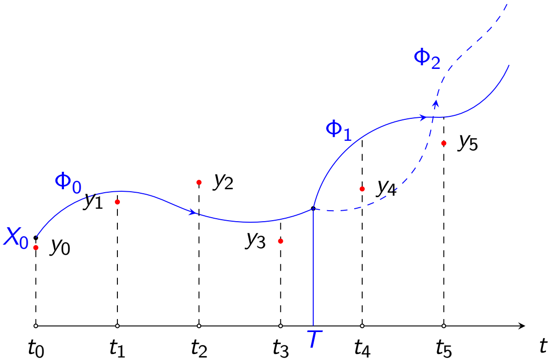

Partially observed optimal stopping problem

- Finite horizon δN

Partially observed optimal stopping problem

- Finite horizon δN
- Admissible stopping times τ: $\mathcal{F}^Y$-measurable
- Admissible decisions A: $\{0, 1, \ldots, d\}$ valued, $\mathcal{F}^Y_\tau$-measurable

Partially observed optimal stopping problem

- Finite horizon δN
- Admissible stopping times τ: $\mathcal{F}^Y$-measurable
- Admissible decisions A: $\{0, 1, \dots, d\}$ valued, $\mathcal{F}^Y_\tau$-measurable
- Cost per stage before stopping
  - $c(0, x, y) = 0$ rightfully not stopped
  - $c(m \neq 0, x, y) = \beta_i \delta$ lateness penalty
- Terminal cost at stopping
  - $C(m, x, y, 0) = c(m, x, y)$ no stopping before the horizon
  - $C(0, x, y, a \neq 0) = \alpha$ early stopping penalty
  - $C(m \neq 0, x, y, a = m) = 0$ good mode selection
  - $C(m \neq 0, x, y, a \neq 0, m) = \gamma$ wrong mode penalty

Partially observed optimal stopping problem

- Finite horizon δN
- Admissible stopping times τ: $\mathcal{F}^Y$-measurable
- Admissible decisions A: $\{0, 1, \dots, d\}$ valued, $\mathcal{F}^Y_\tau$-measurable
- Cost per stage before stopping
  - $c(0, x, y) = 0$ rightfully not stopped
  - $c(m \neq 0, x, y) = \beta_i \delta$ lateness penalty
- Terminal cost at stopping
  - $C(m, x, y, 0) = c(m, x, y)$ no stopping before the horizon
  - $C(0, x, y, a \neq 0) = \alpha$ early stopping penalty
  - $C(m \neq 0, x, y, a = m) = 0$ good mode selection
  - $C(m \neq 0, x, y, a \neq 0, m) = \gamma$ wrong mode penalty

Cost of admissible strategy (τ, A)

$$J(\tau, A, (m, x, y)) = \mathbb{E}_{(m,x,y)} \left[\sum_{n=0}^{(\tau-1)\wedge N} c(X_n, Y_n) + C(X_{\tau\wedge N}, Y_{\tau\wedge N}, A) \right]$$

Fully observed optimal stopping problem

- Filter process $\Theta_n(A \times B) = \mathbb{P}_{(0,x,y)}(X_{\delta n} \in A \times B | \mathcal{F}_n^Y)$
- (Θ_n, Y_n) time inhomogeneous Markov chain with explicit transition kernels R'_n on $\mathcal{P}(E) \times \mathbb{R}$
- cost functions $c'(\theta, y) = \int c(m, x, y) d\theta(m, x)$, $C'(\theta, y, a) = \int C(m, x, y, a) d\theta(m, x)$

Fully observed optimal stopping problem

Minimize over all admissible strategies (τ, a)

$$J'(\tau, A, (\theta, y)) = \mathbb{E}_{(\theta,y)}\left[\sum_{n=0}^{(\tau-1)\wedge N} c'(\Theta_n, Y_n) + C'(\Theta_{\tau\wedge N}, Y_{\tau\wedge N}, A)\right]$$

Aim of the talk

- numerical approximation of the value function
- computable strategy

Difficulties

- measure-valued filter process
- curse of dimensionality

Dynamic programming

Value function

$$V'(\theta, y) = \inf_{(\tau,A)} J'(\tau, A, (\theta, y))$$
$$= \inf_{(\tau,A)} \mathbb{E}_{(\theta,y)} \left[\sum_{n=0}^{(\tau-1)\wedge N} c'(\Theta_n, Y_n) + C'(\Theta_{\tau\wedge N}, Y_{\tau\wedge N}, A) \right]$$

Dynamic programming

$v'_N(\theta, y) = \min_{0\leq a\leq d} C'(\theta, y, a)$
$v'_k(\theta, y) = \min \left\{ \min_{1\leq a\leq d} C'(\theta, y, a); c'(\theta, y) + R'_k v'_{k+1}(\theta, y) \right\}$

$$v'_0 = V'$$

Discretization

$$X_t = (m_y, x_t, t)$$
$$E \times \mathbb{R},\ P$$

Discretization

$$X_t = (m_y, x_t, t)$$
$$E \times \mathbb{R},\ P$$
$$\downarrow$$
$$X_n = (m_{t_n}, x_{t_n})$$
$$E,\ P_n$$
observations $\downarrow$ $Y_n = F(X_n) + \varepsilon_n$
$$(X_n, Y_n)$$
$$E \times \mathbb{R},\ R_n$$

Discretization

$$X_t = (m_y, x_t, t)$$
$$E \times \mathbb{R},\ P$$

$$\downarrow$$

$$X_n = (m_{t_n}, x_{t_n})$$
$$E,\ P_n$$

observations $\downarrow$ $Y_n = F(X_n) + \varepsilon_n$

$$(X_n, Y_n)$$
$$E \times \mathbb{R},\ R_n$$

filtering $\downarrow$ Ψ

$$(\Theta_n, Y_n)$$
$$\mathcal{P}(E) \times \mathbb{R},\ R'_n$$

dynamic programming $\downarrow$

$$v'_n(\Theta_n, Y_n)$$

Discretization

$$X_t = (m_y, x_t, t)$$
$$E \times \mathbb{R},\ P$$

$\downarrow$

$$X_n = (m_{t_n}, x_{t_n}) \qquad (\bar{m}_{t_n}, \bar{x}_{t_n}) = \bar{X}_n$$
$$E,\ P_n \longrightarrow \Omega_n,\ \bar{P}_n$$

quantization

observations $\downarrow$ $Y_n = F(X_n) + \varepsilon_n$

$$(X_n, Y_n)$$
$$E \times \mathbb{R},\ R_n$$

filtering $\downarrow$ Ψ

$$(\Theta_n, Y_n)$$
$$\mathcal{P}(E) \times \mathbb{R},\ R'_n$$

dynamic programming $\downarrow$

$$v'_n(\Theta_n, Y_n)$$

Discretization

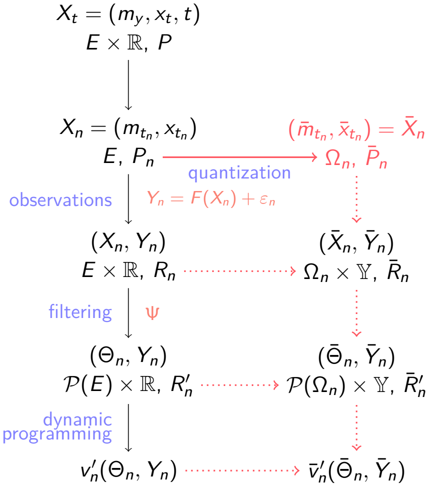

Discretization

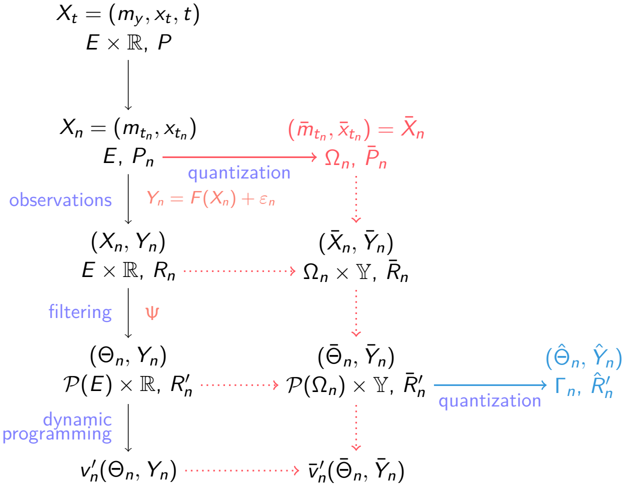

Discretization

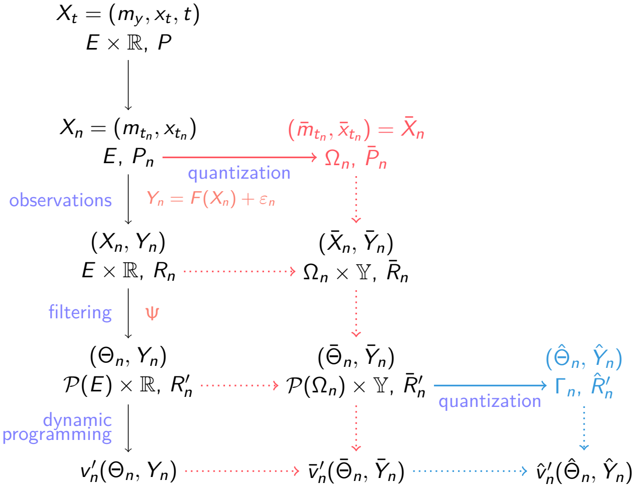

Quantization

[P 98], [PPP 04], [PRS05], ...

Quantization of a random variable $X \in L^2(\mathbb{R}^q)$

Approximate X by $\widehat{X}$ taking finitely many values such that $\|X - \widehat{X}\|_2$ is minimum

- Find a finite weighted grid Γ with $|\Gamma| = K$
- Set $\widehat{X} = p_\Gamma(X)$ closest neighbor projection

Asymptotic properties

If $E[|X|^{2+\eta}] < +\infty$ for some $\eta > 0$ then

$$\lim_{K\to\infty} K^{1/q} \min_{|\Gamma|\le K} \|X - \widehat{X}^\Gamma\|_2 = C$$

Algorithms

There exist algorithms providing

- Γ
- law of $\widehat{X}$
- transition probabilities for quantization of Markov chains

Example: $\mathcal{N}(0, I_2)$:

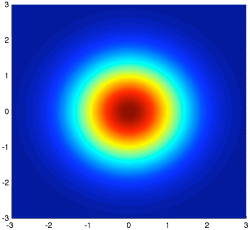

Algorithms

There exist algorithms providing

- Γ
- law of $\widehat{X}$
- transition probabilities for quantization of Markov chains

Example: $\mathcal{N}(0, I_2)$:

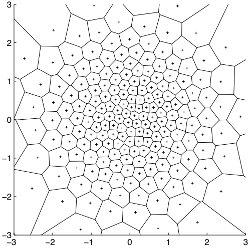

Grids construction

Model $\longrightarrow$ simulator of trajectories $\longrightarrow$ grids

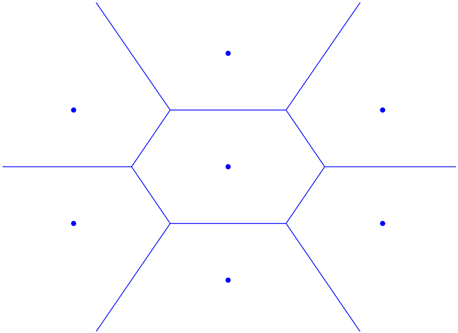

Grids construction

Model $\longrightarrow$ simulator of trajectories $\longrightarrow$ grids

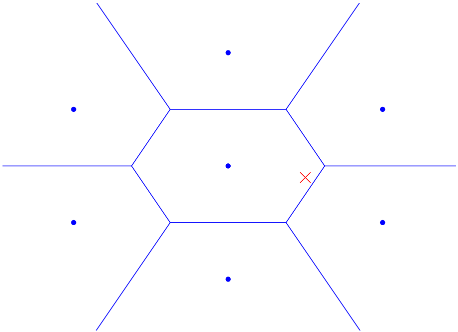

Grids construction

Model $\longrightarrow$ simulator of trajectories $\longrightarrow$ grids

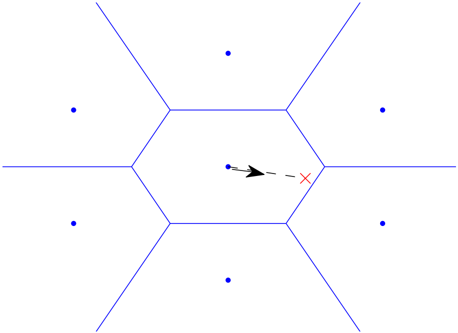

Grids construction

Model ⟶ simulator of trajectories ⟶ grids

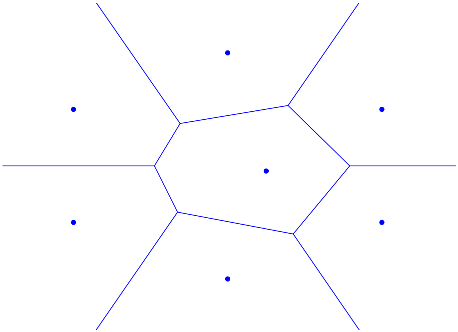

Assets and drawbacks of quantization

Assets

- a simulator of the target law is enough to build the grids
- automatic construction of grids
- convergence rate for $\mathbb{E}[|f(X) - f(\widehat{X})|]$ if f lipschitz
- empirical error measure by Monte Carlo

Drawbacks

- computation time
- curse of dimension
- open questions of convergence of the algorithms

Convergence

Technical assumptions

$$
\begin{aligned}
|v_0'(\delta_{(0,x_0)}, y_0) - \bar{v}_0'(\delta_{(0,x_0)}, y_0)| &\leq \sum_{n=0}^{N-1} a_n \mathbb{E}[|\bar{X}_n - X_n|] \\
&= O(N_\Omega^{-1})
\end{aligned}
$$

$$
\begin{aligned}
&|\hat{v}_0'(\delta_{(0,x_0)}, y_0) - \bar{v}_0'(\delta_{(0,x_0)}, y_0)| \\
&\qquad \leq \sum_{n=0}^{N} c_n \left(\mathbb{E}\left[\left| \hat{Y}_n - \bar{Y}_n \right| \right] + \mathbb{E}\left[\|\hat{\Theta}_n - \bar{\Theta}_n\|_{n,1} \right] \right) \\
&\qquad = O(N_\Gamma^{-1/N_\Omega})
\end{aligned}
$$

Candidate computable strategy

Dynamic programming

- $\hat{v}'_N(\hat{\theta}, \hat{y}) = \min_{0 \leq a \leq d} C'(\hat{\theta}, \hat{y}, a)$
- $\hat{v}'_k(\hat{\theta}, \hat{y}) = \min \left\{ \min_{1 \leq a \leq d} C'(\hat{\theta}, \hat{y}, a); c'(\hat{\theta}, \hat{y}) + \hat{R}'_k \hat{v}'_{k+1}(\hat{\theta}, \hat{y}) \right\}$

Set

- $r_N(\cdot) = 0$, $a_N(\cdot) = 0$ if $\hat{v}'_N(proj_{\Gamma_N}(\cdot)) = C'(proj_{\Gamma_N}(\cdot), 0)$
- $r_N(\cdot) = 1$, $a_N(\cdot) = i$ if $\hat{v}'_N(proj_{\Gamma_N}(\cdot)) = C'(proj_{\Gamma_N}(\cdot), i)$
- $r_n(\cdot) = 0$ if $\hat{v}'_n(proj_{\Gamma_n}(\cdot)) = \hat{R}'_n \hat{v}'_{n+1}(proj_{\Gamma_n}(\cdot))$
- $r_n(\cdot) = 1$, $a_n(\cdot) = i$ if $\hat{v}'_n(proj_{\Gamma_n}(\cdot)) = C'(proj_{\Gamma_n}(\cdot), i)$

Candidate computable strategy

$$
\begin{aligned}
n &\leftarrow 0 \\
y &\leftarrow y_0 \\
\bar{\theta} &\leftarrow \delta_{(0,x_0)} \\
r &\leftarrow r_0(\bar{\theta}, y)
\end{aligned}
$$

Observation y_0

Candidate computable strategy

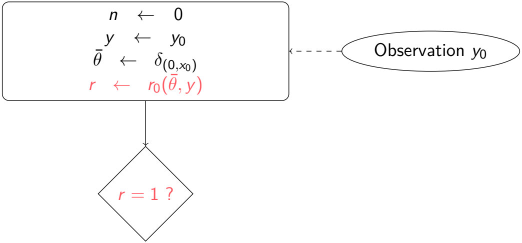

Candidate computable strategy

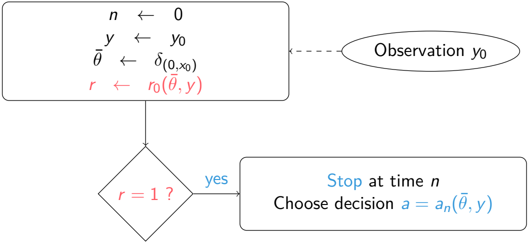

Candidate computable strategy

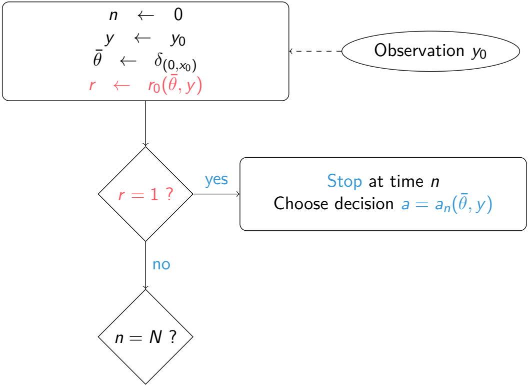

Candidate computable strategy

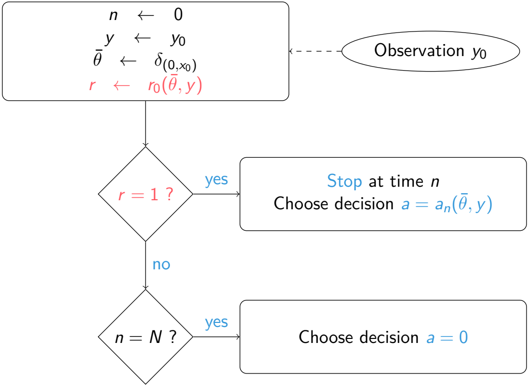

Candidate computable strategy

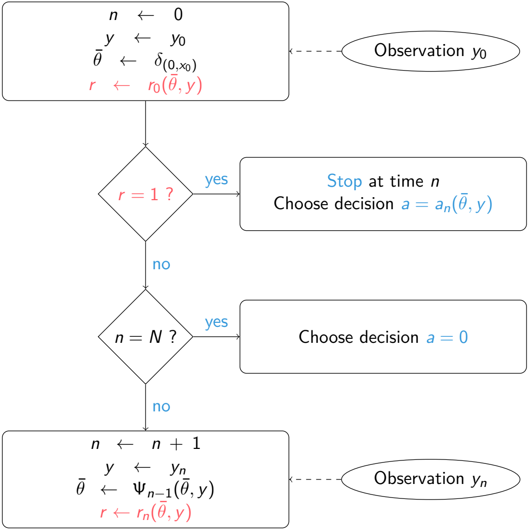

Candidate computable strategy

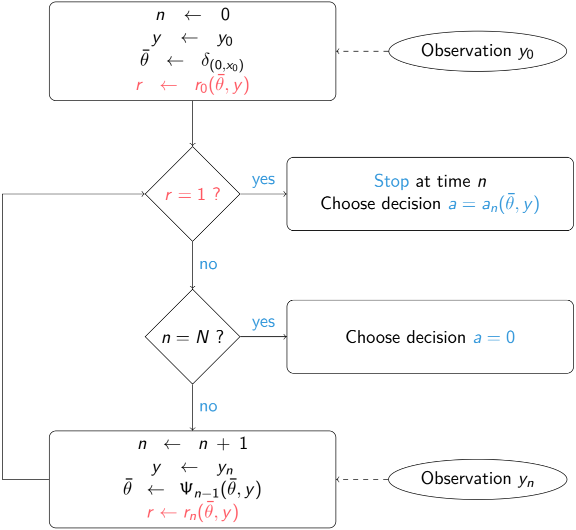

Example 1

- $d = 3$, $p_i = 1/3$, $x_0 = 1$
- $\Phi_0(x,t) = x$, $\Phi_1(x,t) = xe^{0.1t}$, $\Phi_2(x,t) = xe^{0.5t}$, $\Phi_3(x,t) = xe^{1t}$
- $\beta = 1$ (late detection), $\gamma = 1.5$ (wrong mode), $\delta = 1/6$

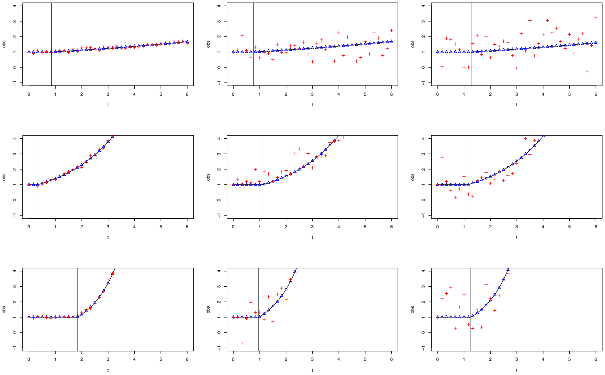

Example 1

- $d = 3$, $p_i = 1/3$, $x_0 = 1$
- $\Phi_0(x,t) = x$, $\Phi_1(x,t) = xe^{0.1t}$, $\Phi_2(x,t) = xe^{0.5t}$, $\Phi_3(x,t) = xe^{1t}$
- $\beta = 1$ (late detection), $\gamma = 1.5$ (wrong mode), $\delta = 1/6$

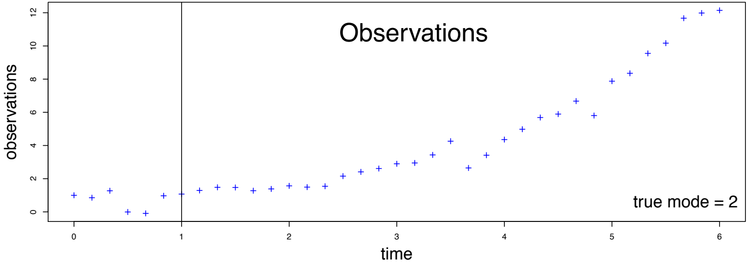

Example 1

- $d = 3$, $p_i = 1/3$, $x_0 = 1$
- $\Phi_0(x,t) = x$, $\Phi_1(x,t) = xe^{0.1t}$, $\Phi_2(x,t) = xe^{0.5t}$, $\Phi_3(x,t) = xe^{1t}$
- $\beta = 1$ (late detection), $\gamma = 1.5$ (wrong mode), $\delta = 1/6$

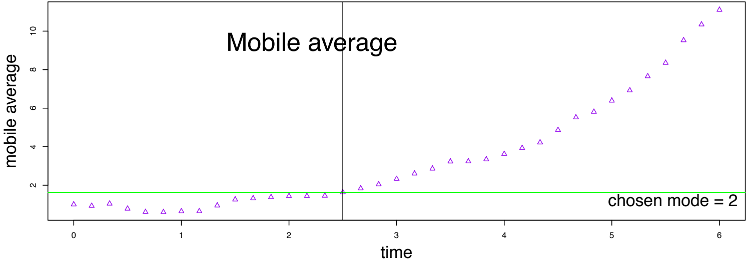

Example 1

- $d = 3,\ p_i = 1/3,\ x_0 = 1$
- $\Phi_0(x,t) = x,\ \Phi_1(x,t) = xe^{0.1t},\ \Phi_2(x,t) = xe^{0.5t},$ $\Phi_3(x,t) = xe^{1t}$
- $\beta = 1$ (late detection), $\gamma = 1.5$ (wrong mode), $\delta = 1/6$

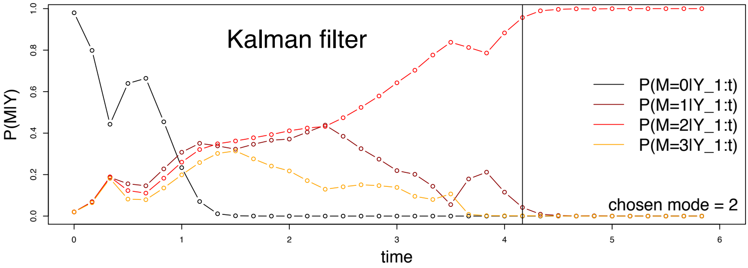

Example 1

- $d = 3$, $p_i = 1/3$, $x_0 = 1$
- $\Phi_0(x,t) = x$, $\Phi_1(x,t) = xe^{0.1t}$, $\Phi_2(x,t) = xe^{0.5t}$, $\Phi_3(x,t) = xe^{1t}$
- $\beta = 1$ (late detection), $\gamma = 1.5$ (wrong mode), $\delta = 1/6$

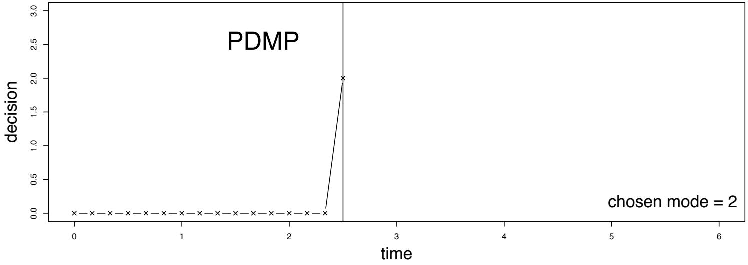

Example 1

- $d = 3$, $p_i = 1/3$, $x_0 = 1$
- $\Phi_0(x,t) = x$, $\Phi_1(x,t) = xe^{0.1t}$, $\Phi_2(x,t) = xe^{0.5t}$, $\Phi_3(x,t) = xe^{1t}$
- $\beta = 1$ (late detection), $\gamma = 1.5$ (wrong mode), $\delta = 1/6$

α	σ^2	Moving Average threshold=2				Kalman				PDMP			
		window				threshold				Nb grid points			
		2	3	4	5	0.5	0.75	0.9	cal	30	50	75	100
3	0.1	**0.40**	**0.40**	**0.40**	0.41	2.34	0.61	**0.42**	**0.42**	**0.70**	**0.70**	**0.70**	**0.70**
	0.5	0.93	0.81	0.76	**0.73**	1.44	0.54	0.51	**0.49**	0.78	0.79	0.77	**0.76**
	1	1.73	1.42	1.29	**1.16**	1.18	**0.58**	0.63	0.62	0.99	1.04	**0.98**	1.01
4	0.1	**0.40**	**0.40**	**0.40**	0.41	3.06	0.69	**0.42**	**0.42**	0.69	0.71	0.69	**0.68**
	0.5	0.95	0.81	0.76	**0.73**	1.76	0.56	0.51	**0.50**	0.73	**0.71**	0.72	0.72
	1	2.05	1.57	1.39	**1.22**	1.36	**0.60**	0.63	0.62	**0.92**	**0.92**	0.95	0.95
5	0.1	**0.40**	**0.40**	**0.40**	0.41	3.78	0.78	**0.42**	**0.42**	0.68	0.69	**0.67**	0.69
	0.5	0.97	0.81	0.76	**0.73**	2.08	0.59	0.51	**0.50**	0.72	**0.69**	0.72	0.72
	1	2.37	1.73	1.48	**1.28**	1.54	**0.61**	0.63	0.62	**0.92**	0.94	0.93	**0.92**
6	0.1	**0.40**	**0.40**	**0.40**	0.41	4.50	0.86	**0.42**	0.43	**0.68**	**0.68**	**0.68**	0.69
	0.5	0.98	0.82	0.76	**0.73**	2.40	0.62	0.51	**0.50**	0.70	0.70	0.70	**0.69**
	1	2.69	1.88	1.57	**1.35**	1.72	0.63	0.63	**0.62**	0.90	**0.89**	0.91	**0.89**

Example 2

- $d = 1$, $x_0 = (0, 0)$
- $\Phi_0((x, u), t) = (sin(3\pi(u + t)), u + t)$,
  $\Phi_1((x, u), t) = (sin(5\pi(u + t)), u + t)$
- $\delta = 1/6$, noise variance 1

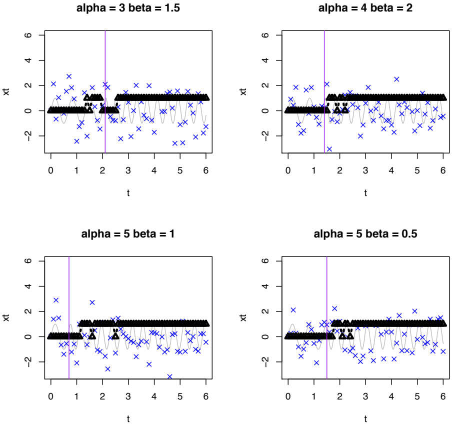

Example 3

- $d = 2$, $x_0 = (0,0)$
- $\Phi_0((x,u),t) = (sin(3\pi(u+t)), u+t)$,
  $\Phi_1((x,u),t) = (sin(3\pi(u+t)) + 0.5t, u+t)$,
  $\Phi_2((x,u),t) = (sin(3\pi(u+t)) + 1.5t, u+t)$
- $\delta = 1/6$, noise variance 1

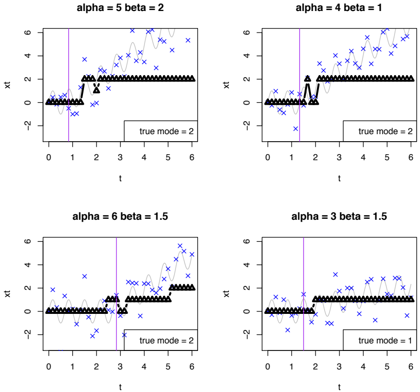

Conclusion and perspectives

- Change-point detection method for continuous-time jump dynamics, able to detect a jump and select the post-jump mode
- For general flows but dimension 1 (+ time)

Conclusion and perspectives

- Change-point detection method for continuous-time jump dynamics, able to detect a jump and select the post-jump mode
- For general flows but dimension 1 (+ time)

To be done

- Real data applications
- Theoretical validity of the stopping rule
- Allow to stop between observations
- Several jumps
- Stop and restart the process from a new point

Reference

[BL 17] N. Bäuerle, D. Lange *Optimal control of partially observed PDMPs*
[BdSD 12] A. Brandejsky, B. de Saporta, F. Dufour *Optimal stopping for partially observed PDMPs*
[CD 89] O. Costa, M. Davis *Impulse control of piecewise-deterministic processes*
[Davis 93] M. Davis, *Markov models and optimization*
[dSDZ 14] B. de Saporta, F. Dufour, H. Zhang *Numerical methods for simulation and optimization of PDMPs: application to reliability*
[P 98] G. Pagès *A space quantization method for numerical integration*
[PPP 04] G. Pagès, H. Pham, J. Printems *An optimal Markovian quantization algorithm for multi-dimensional stochastic control problems*
[PRS 05] H. Pham, W. Runggaldier, A.f Sellami *Approximation by quantization of the filter process and applications to optimal stopping problems under partial observation*