AI & Robotics: Research
Exercise 7

Marc Toussaint
Learning & Intelligent Systems Lab, TU Berlin
Marchstr. 23, 10587 Berlin, Germany

Summer 2020

Please prepare written solutions to the following exercises that you can share by screen in our session. The notes can be brief, but esp. equations and derivations should be written precisely. Try to use LaTeX.

1 Filtering a 2D agent

Consider an agent with 2D state space $(x, v) \in \mathbb{R}^2$, and you try to track the state of the agent. (A bit like tracking Mister X in 'Scotland Yard', but with different motion and observation models.) The agent's motion is quite erratic in v-direction, but smooth in x-direction:

$$v' = v + \beta \ , \quad \beta \sim \mathcal{U}(\{-\gamma, \gamma\}) \ , \quad \gamma \sim \mathcal{U}(0.5, 1.) \tag{1}$$

$$x' = x + v + \epsilon \ , \quad \epsilon \sim \mathcal{N}(0, \sigma^2) \ . \tag{2}$$

(Note, this notation is equivalent to writing $P(x'|x, v) = \mathcal{N}(x + v, \sigma^2)$. Some intuition behind these equations: If v is interpreted as a velocity, the β adds erratic accelerations of $+\gamma$ or $-\gamma$, with γ uniform in $[0.5, 1]$ (a bit like noisy bang-bang) in each time step, while the position x changes with the velocity and small Gaussian noise with sdv σ. This is also called double integrator.)

Your observations of the agent are rather limited: You can observe the position only, with Gaussian observation noise

$$y = x + \xi \ , \quad \xi \sim \mathcal{N}(0, \widehat{\sigma}^2) \ . \tag{3}$$

Assume your initial belief is Gaussian, i.e., $p_0(x, v) = \mathcal{N}(x|0, 1)\ \mathcal{N}(v|0, 1)$.

a) Derive the Rao-Blackwellized particle filter to track the state of the agent, where your belief is factored $p(x, v) = p(v)p(x|v)$, and $p(v)$ is approximated as a particle distribution, while $p(x|v)$ is a Gaussian (conditional to v).

   Describe the RBPF using pseudo code as explicitly as possible; detailing step-by-step what has to be computed and making equations explicit.

   You may assume that methods for multiplying Gaussians and propagating Gaussians are given:
   $(L, c, C) \leftarrow \texttt{PROD}(a, A, b, B)$ computes $L\ \mathcal{N}(x \mid c, C) = \mathcal{N}(x \mid a, A)\ \mathcal{N}(x \mid b, B)$
   $(b, B) \leftarrow \texttt{PROP}(a, A, v, \sigma)$ computes $\mathcal{N}(x' \mid b, B) = \int_x \mathcal{N}(x' \mid x + v, \sigma^2)\ \mathcal{N}(x \mid a, A)\ dx$

b) Using a particle distribution to represent $p(v)$ is of course an approximation. But, *conditional to a particle*, is the Gaussian filtering of $p_t(x|v)$ also an approximation?

   In fact, what exactly does the Gaussian distribution that is attached to each particle represent? Is it $p_t(x|v_t^i)$, or $p_t(x|v_{0:t}^i)$? And what exactly is the difference?

c) Above, we used particles to represent $p(v)$, and Gaussians to represent $p(x|v)$. Could we swap this? What issues would arise if we swap this?

d) If we bin the v-values (discretize velocities into discrete bins), then we could use a histogram (i.e., multinomial) to represent $p(v)$. How could you compute an approximate predictive distribution $\widehat{p}_t(v)$ from an old belief $p_{t\text{-}1}(v)$, when $p_{t\text{-}1}(v)$ is a histogram over bins?