Multi target tracking via PHD filter

Equipe ALEA
http://alea.bordeaux.inria.fr

F. Caron, P. Del Moral, P. Legrand, M. Pace

{Francois.Caron,Pierre.Del-Moral,Michele.Pace}@inria.fr,legrand@sm.u-bordeaux2.fr

Institut de Mathématiques de Bordeaux

Contribution

Filtering is the problem of estimating the state of a system as a set of observations becomes available in time, and it has gained a great importance in many fields of science, engineering and finance. To solve it, one begins by modeling the evolution of the system and the uncertainty in the measurements. The resulting models typically exhibit complex non-linearities and non-Gaussian distributions, thus precluding analytical solution. Sequential Monte Carlo methods, also known as particle filters can be used to estimate the a posteriori distribution of the system state. They are mainly composed by two steps: the evolution of initial samples using some dynamic equation, and the resampling of particles according to their likelihood (in a Bayesian sense), with respect to the current observation. Modern lines of research in different fields like target tracking, computer vision, financial mathematics, and biology make wide use of these techniques; nevertheless they develop and adopt each concept rather autonomously, with very few connections with each other and with the mathematical foundation of the field. We try to illustrate one of the possible links among these disciplines, more precisely how Bayesian filtering applied to multi-target tracking can be completely and fruitfully cast into the framework of genetic algorithms.

Bayesian filtering

The target motion is modeled by a transition probability density:

$$x_k \sim f_{k|k-1}(.|x_{k-1}) \tag{1}$$

The target state x_k generates an observation y_k according to the observation model:

$$y_k \sim g_k(.|x_k) \tag{2}$$

The aim is to construct the posterior probability density function of the target state given the observations. The Bayes filter is a recursion based on two steps: **prediction** (3) and **update** (4):

$$p_{k|k-1}(x_k|y_{1:k-1}) = \int f_{k|k-1}(x_k|x)p_{k-1}(x|y_{1:k-1})dx \tag{3}$$

$$p_k(x_k|y_{1:k}) = \frac{g_k(y_k|x_k)p_{k|k-1}(x_k|y_{1:k-1})}{\int g_k(y_k|x)p_{k|k-1}(x|y_{1:k-1})dx} \tag{4}$$

The analytical solution to the Bayes filter is generally intractable, with the exception of the linear Gaussian case whose exact solution is given by Kalman filter. Extended Kalman filter (EKF) and Unscented Kalman filter (UKF) have been proposed to address non linear models. A more general solution strategy is to use sequential Monte Carlo methods, also known as particle filters.

Multi-target tracking and PHD filter

Multi-target tracking involves jointly estimating the number and the states of a finite, time varying number of targets given a set of finite and time varying measurements of uncertain origins.

The PHD Filter:

- is a suboptimal but tractable alternative to the multi-target Bayes filter.
- It propagates the first order moments of the targets random finite sets (RFS) known as intensity (or PHD) function.
- Given a RFS X on $\mathcal{X} \subseteq \mathbb{R}^{n_x}$ with probability $\mathcal{P}$ the intensity function is a non-negative function v on $\mathcal{X}$ such that, for any closed subset $S \subseteq \mathcal{X}$:

  $$\int_S v(x)dx = \int |X \cap S| \mathcal{P}(dX)$$

  providing an estimate about the number of elements of a RFS in the region considered.
- In the SMC-PHD filter the PHD function is approximated by the particle representation $\hat{v}_k(x) = \sum_{i=1}^{N_k} w_k^{(i)} \delta_{x_k^{(i)}}(x)$.
- Importance sampling and resampling are applied to propagate the particle-approximated intensity function through the recursion

  $$v_{k|k} = (\psi_k \circ \Phi_{k|k-1})v_{k-1|k-1}$$

  as specified in equations 8 and 9.

References

[1] Del Moral P., Kallel L. and Rowe J. Modeling genetic algorithms with interacting particle systems In *Revista de Matematica, Teoria y aplicaciones,Vol.8,No. 2, July (2001).*

[2] Del Moral P. Feynman-Kac Formulae Genealogical and Interacting Particle Systems with Applications. *Springer New York.*

[3] Ba-Ngu Vo, Sumeetpal Singh, Arnaud Doucet Sequential Monte Carlo methods for Multi-target Filtering with Random Finite Sets In *IEEE Transaction on Aerospace and Electronic Systems, Vol. X, No. XXX, June 2005*

Particle Filtering and Genetic Algorithms

Genetic Algorithms

In mathematical terms a genetic algorithm is a Markov chain $X_n = (X_n^i)_{1\le i\le N}$ defined on the product space E^N. The evolution of the chain is composed by a phase of selection, during which individuals are selected according to certain criteria (fitness) and by a phase of state exploration.

$$X_n = (X_n^i)_{1\le i\le N} \xrightarrow{\text{selection}} \widehat{X}_n = (\widehat{X}_n^i)_{1\le i\le N} \xrightarrow{\text{mutation}} X_{n+1} = (X_{n+1}^i)_{1\le i\le N}$$

- For the selection phase a positive function (fitness) is given and each new individual is selected proportionally to his fitness.

  $$\psi(\hat{x}_n^{(i)}) = \sum_{i=1}^N \frac{G(X_n^i)}{\sum_{j=1}^N G(X_n^j)} \delta_{X_n^i}$$

- During the mutation phase the state space is explored according to the Markov transition $M(x_{n+1}^{(i)}, \hat{x}_n^{(i)})$

  $$x_{n+1}^{(i)} \sim M(x_{n+1}^{(i)}, \hat{x}_n^{(i)})$$

Filtering

In filtering problems, the correction/prediction phases of the optimal filter:

$$\text{Law}(X_n|Y_0\ldots Y_{n-1}) = \eta_n \xrightarrow{\text{correction}} \text{Law}(X_n|Y_0\ldots Y_n) = \Psi_{G_n}(\eta_n) = \widehat{\eta}_n \xrightarrow{\text{prediction}} \eta_{n+1} = \widehat{\eta}_n M_{n+1}$$

are estimated using a Markov process defined on the product space E^N:

$$\xi_n^{(N)} = \left(\xi_n^{(N,i)}\right)_{1\le i\le N} \xrightarrow{\text{selection}} \widehat{\xi}_n^{(N)} = \left(\widehat{\xi}_n^{(N,i)}\right)_{1\le i\le N} \xrightarrow{\text{mutation}} \xi_{n+1}^{(N)}$$

This genetic algorithm is a particle filter as the occupation measures of the **predicted** and **updated** populations (or **selected** and **mutated** in terms of genetic algorithms) converge almost surely to the optimal predictor and filter.

$$\eta_n^N := \frac{1}{N}\sum_{i=1}^N \delta_{\xi_n^{(N,i)}} \quad , \quad \widehat{\eta}_n^N := \frac{1}{N}\sum_{i=1}^N \delta_{\widehat{\xi}_n^{(N,i)}} \quad \text{et} \quad \lim_{N\to\infty} \eta_n^N = \eta_n \quad , \quad \lim_{N\to\infty} \widehat{\eta}_n^N = \widehat{\eta}_n$$

Feynman Kac Measures

Let's consider the probability distribution flow $(\eta_n)_{n\ge 0}$ defined by equations:

$$\eta_n = \Phi_n(\eta_{n-1}) := \Psi_{G_{n-1}}(\eta_{n-1})M_n \tag{5}$$

The measures η_n and the corresponding updated measures $\widehat{\eta}_n := \Psi_{G_n}(\eta_n)$ can be expressed as weighted path integrals known as **Feynman-Kac formulae**. These functional representations are defined on test functions f_n by the following equations:

$$\eta_n(f_n) = \gamma_n(f_n)/\gamma_n(1) \quad \text{et} \quad \widehat{\eta}_n(f_n) = \widehat{\gamma}_n(f_n)/\widehat{\gamma}_n(1) \tag{6}$$

With γ_n and $\widehat{\gamma}_n$ being unnormalized measures defined by:

$$\gamma_n(f_n) = \mathbb{E}[f_n(X_n)\textstyle\prod_{0\le k<n} G_k(X_k)] \quad \text{et} \quad \widehat{\gamma}_n(f_n) = \gamma_n(f_n G_n) \tag{7}$$

PHD filter

The PHD prediction and update operators $\Phi_{k|k-1}$ and Ψ_k used to approximate the first moment of the posterior distribution are respectively defined as

$$(\Phi_{k|k-1}\alpha)(x) = \int [P_{S,k|k-1}(x)f_{k|k-1}(x|\xi) + b_{k|k-1}(x|\xi)]\alpha(\xi)\lambda(d\xi) + \gamma_k(x) \tag{8}$$

$$(\Psi_k\alpha)(x) = \left[(1-P_D(x)) + \sum_{z\in Z_k} \frac{P_D(x)g_k(z|x)}{\kappa_k(z) + \langle P_D(x)g_k(z|x), \alpha\rangle}\right]\alpha(x) \tag{9}$$

Simulations

Simulation 1: Particle filter approximating the posterior distribution $Law(X_n|Y_0\ldots Y_n)$ of a target moving on a square. Three sensors report only angular measurements, the observation noise is Gaussian, centered with variance $\sigma^2 = 0.1rad$. A random walk is used as a priori model for target evolution. 1000 particles are used.

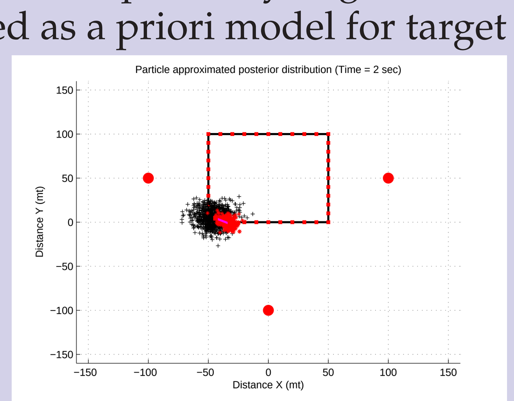

Fig.1: Time = 2 sec.

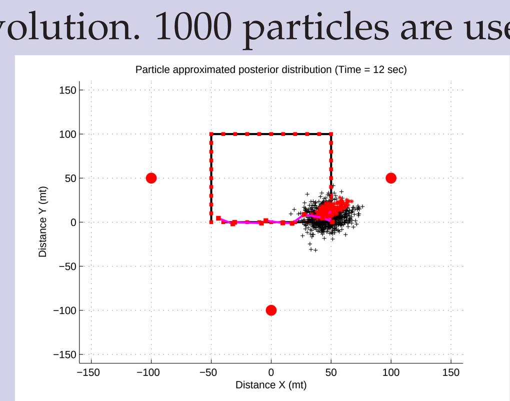

Fig.2: Time = 12 sec.

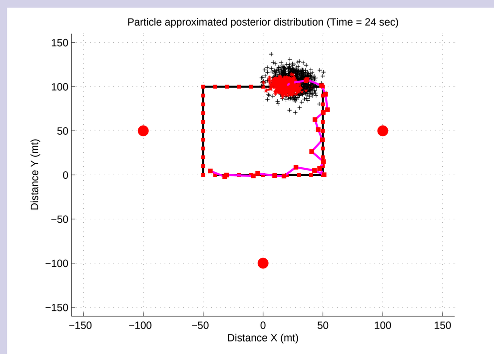

Fig.3: Time = 24 sec.

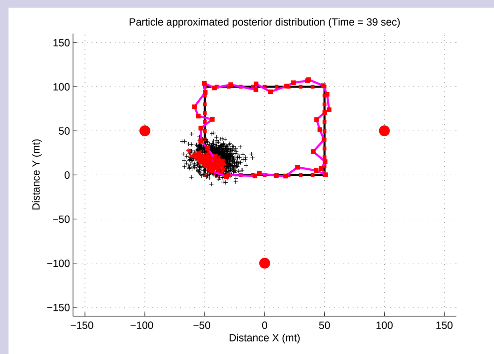

Fig.4: Time = 39 sec.

Simulation 2: PHD Filter estimating the posterior intensity function of multiple targets moving over a 1D surveillance zone with false alarms and detection uncertainty.

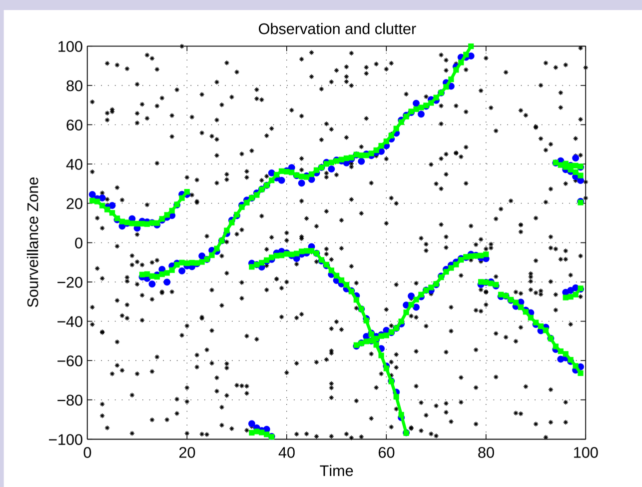

Fig.5: Targets (green), measurements (blue) and clutter (black)

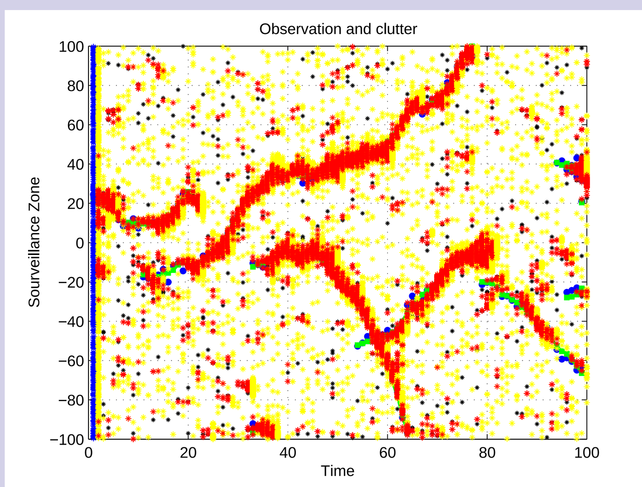

Fig.6: SMC-PHD filter output.

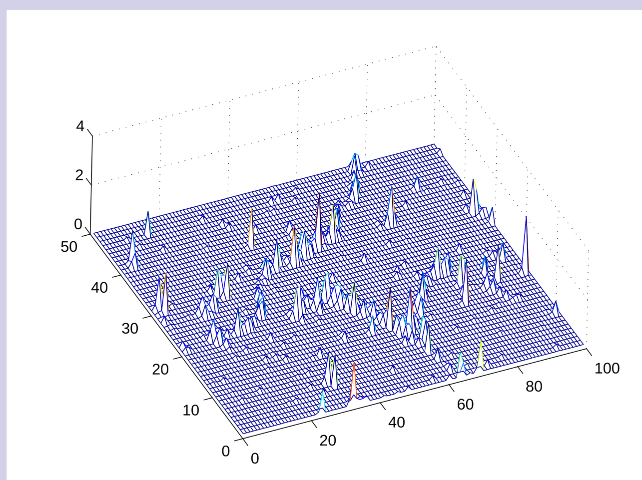

Fig.7: 3D plot of the posterior intensity.

Simulation 3: PHD Filter applied to a realistic 3D scenario, radar-like measurements containing false alarms are used to estimate the position of multiple targets (airplanes and boats) maneuvering in the surveillance zone.

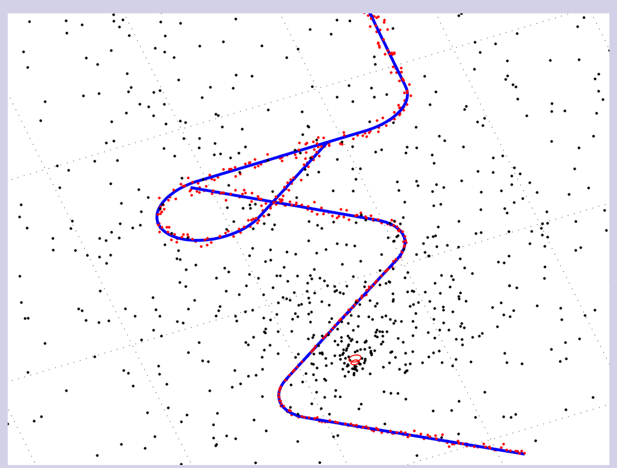

Fig.8: Maneuvering targets (blue), measurements (red) and clutter (black).

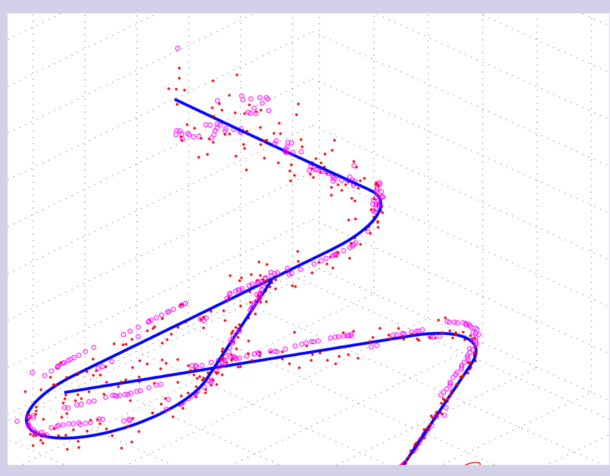

Fig.9: SMC-PHD Filter estimated target positions.

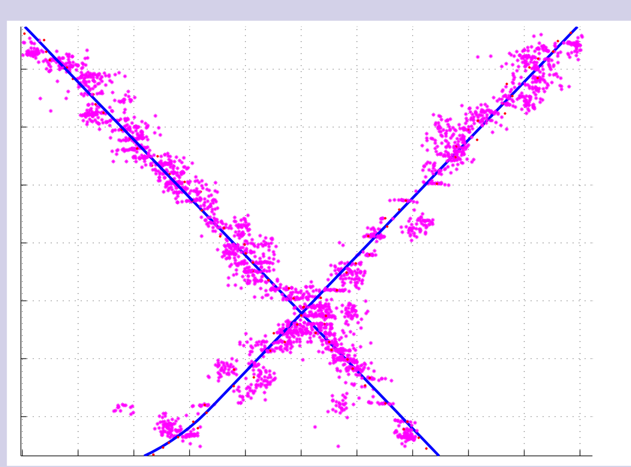

Fig.10: PHD particles surrounding true target trajectories during filtering.

Equipe ALEA: Advanced Learning Evolutionary Algorithms

Joint team with the IMB (CNRS and university Bordeaux 1 et 2). INRIA Bordeaux - Sud Ouest. http://alea.bordeaux.inria.fr/