

School of Computer Science
McGill University

PARTICLE FILTER TUTORIAL

**Computer Robot Vision Conference
2010**

Fundamental Problems In Robotics

- How to Go From A to B ? (**Path Planning**)
- What does the world looks like? (**mapping**)
 - sense from various positions
 - integrate measurements to produce map
 - assumes perfect knowledge of position
- Where am I in the world? (**localization**)
 - Sense
 - relate sensor readings to a world model
 - compute location relative to model
 - assumes a perfect world model
- Together, the above two are called **SLAM**
(Simultaneous Localization and Mapping)



Localization

- Tracking: Known initial position
- Global Localization: Unknown initial position
- Re-Localization: Incorrect known position
 - (kidnapped robot problem)



Sensors

- **Proprioceptive Sensors**

(monitor state of vehicle-propagate)


- IMU (accels & gyros)
- Wheel encoders
- Doppler radar ...
 - **Noise**



- **Exteroceptive Sensors**

(monitor environment-update)

- Cameras (single, stereo, omni, FLIR ...)
- Laser scanner
- MW radar
- Sonar
- Tactile...
 - **Uncertainty**



Bayesian Filter

- Estimate state x from data Z
 - *What is the probability of the robot being at x ?*
- x could be robot location, map information, locations of targets, etc...
- Z could be sensor readings such as range, actions, odometry from encoders, etc...)
- This is a general formalism that does not depend on the particular probability representation
- Bayes filter **recursively** computes the posterior distribution:

$$Bel(x_T) = P(x_T | Z_T)$$



Iterating the Bayesian Filter

- Propagate the motion model:

$$Bel_{-}(x_t) = \int P(x_t | a_{t-1}, x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

Compute the current state estimate before taking a sensor reading by integrating over all possible previous state estimates and applying the motion model

- Update the sensor model:

$$Bel(x_t) = \eta P(o_t | x_t) Bel_{-}(x_t)$$

Compute the current state estimate by taking a sensor reading and multiplying by the current estimate based on the most recent motion history

Mobile Robot Localization

(Where Am I?)

- A mobile robot moves while collecting sensor measurements from the environment.
- Two steps, action and sensing:
 - **Prediction/Propagation**: what is the robots pose \mathbf{x} after action \mathbf{A} ? (X,Y, θ)
 - **Update**: Given measurement \mathbf{z} , correct the pose \mathbf{x}'
- What is the probability density function (*pdf*) that describes the uncertainty \mathbf{P} of the poses \mathbf{x} and \mathbf{x}' ?



State Estimation

- Propagation

$$P(x_{t+1}^- | x_t, \alpha)$$

- Update

$$P(x_{t+1}^+ | x_{t+1}^-, z_{t+1})$$

Traditional Approach Kalman Filter

- Optimal for linear systems with Gaussian noise
- Extended Kalman filter:
 - Linearization
 - Gaussian noise models
- Fast!



Monte-Carlo State Estimation

(Particle Filtering)

- Employing a Bayesian Monte-Carlo simulation technique for pose estimation.
- A particle filter uses N samples as a discrete representation of the probability distribution function (*pdf*) of the variable of interest:

$$S = [\vec{\mathbf{x}}_i, w_i : i = 1 \cdots N]$$

where \mathbf{x}_i is a copy of the variable of interest and w_i is a weight signifying the quality of that sample.

In our case, each particle can be regarded as an alternative hypothesis for the robot pose.



Particle Filter (cont.)

The particle filter operates in two stages:

- **Prediction:** After a motion (α) the set of particles S is modified according to the action model

$$S' = f(S, \alpha, \nu)$$

where (ν) is the added noise.

The resulting *pdf* is the prior estimate before collecting any additional sensory information.



Particle Filter (cont.)

- **Update:** When a sensor measurement (z) becomes available, the weights of the particles are updated based on the likelihood of (z) given the particle x_i

$$w'_i = P(z | \vec{x}_i) w_i$$

The updated particles represent the posterior distribution of the moving robot.

Remarks:

- **In theory**, for an infinite number of particles, this method models the true *pdf*.
- **In practice**, there are always a finite number of particles.

Resampling

For finite particle populations, we must focus population mass where the *PDF* is substantive.

- Failure to do this correctly can lead to divergence.
- Resampling needlessly, also has disadvantages.

One way is to estimate the need for resampling based on the variance of the particle weight distribution, in particular the coefficient of variance:

$$cv_t^2 = \frac{\text{var}(w_t(i))}{E^2(w_t(i))} = \frac{1}{M} \sum_{i=1}^M (Mw_t(i) - 1)^2$$

$$ESS_t = \frac{M}{1 + cv_t^2}$$




Prediction: Odometry Error Modeling

- **Piecewise linear motion**: a simple example.
- **Rotation**: Corrupted by Gaussian Noise.
- **Translation**: Simulated by multiple steps. Each step models translational and rotational error.

Single step:

Small *rotational* error (drift) before and after the translation.

Translational error proportional to the distance traveled.



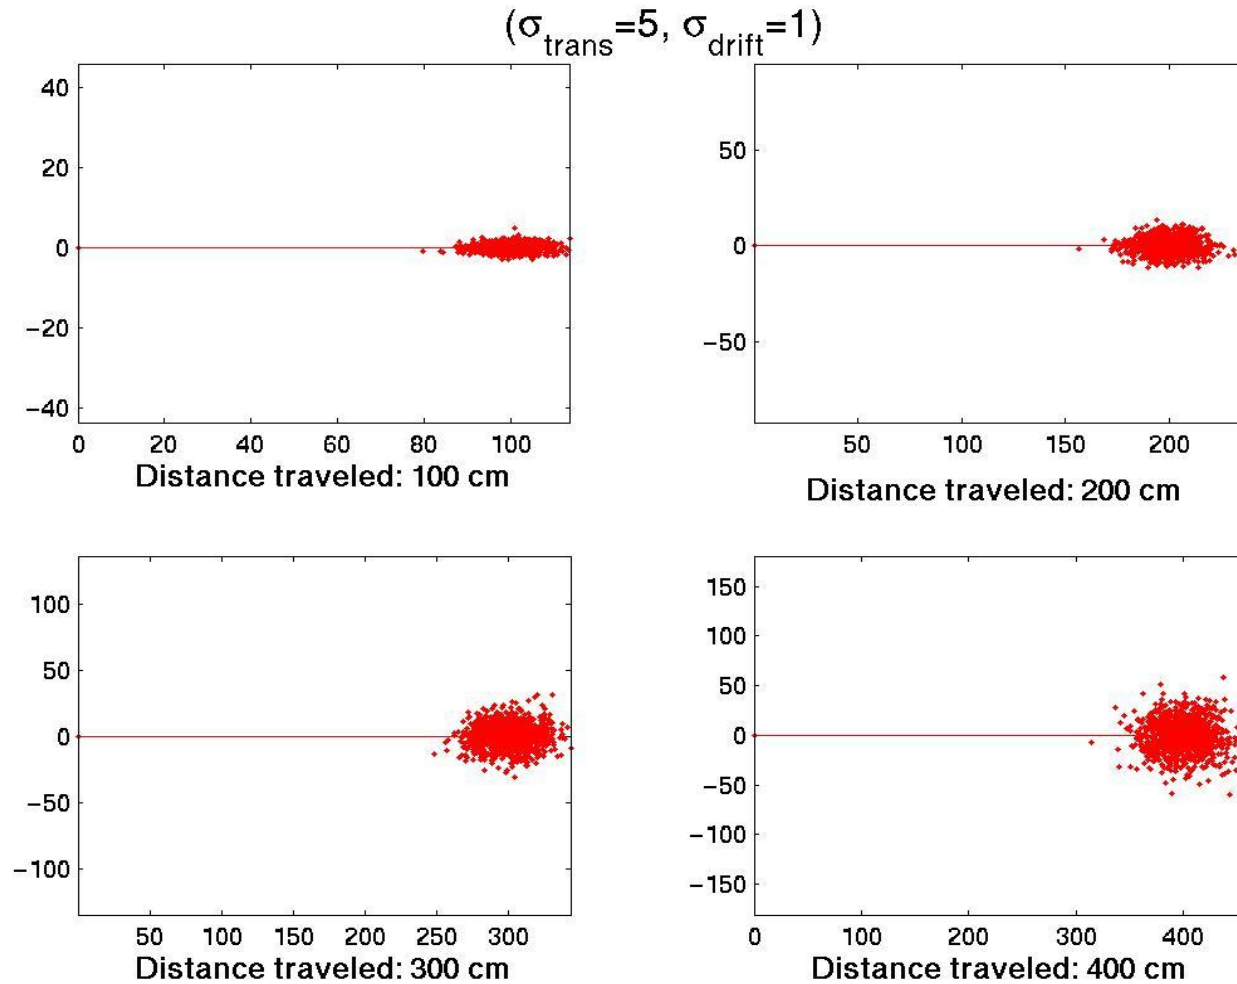
All errors drawn from a Normal Distribution.



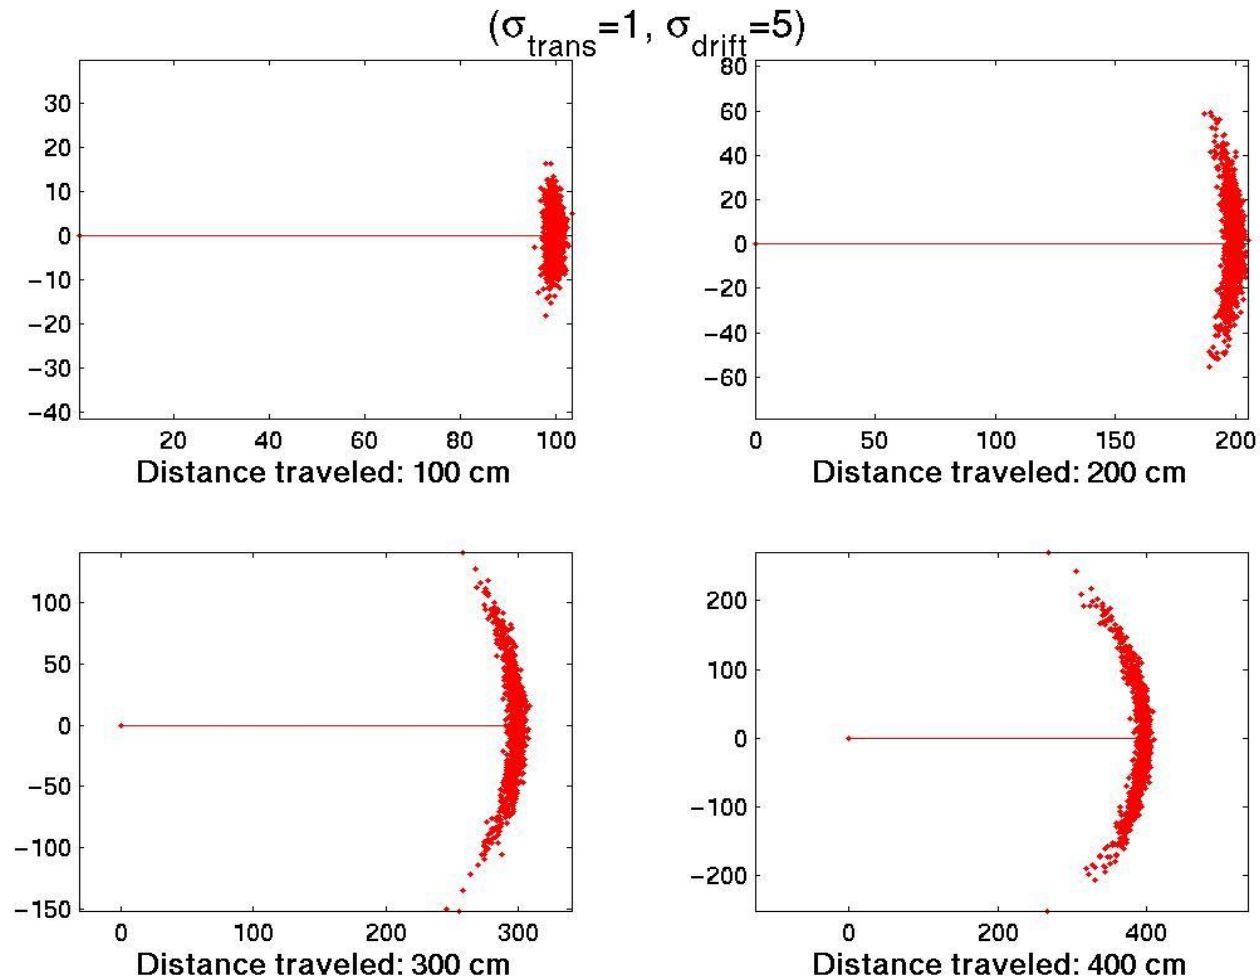
Odometry Error Modeling




Odometry Error Modeling




Odometry Error Modeling




Odometry Error Modeling



Odometry Error Modeling



Prediction-Only Particle Distribution



Propagation of a discrete time system

($\delta t=1$ sec)

$$x_i^{t+1} = x_i^t + (v_t + w_{v_t}) \delta t \cos \phi_i^t$$

$$y_i^{t+1} = y_i^t + (v_t + w_{v_t}) \delta t \sin \phi_i^t$$


$$\phi_i^{t+1} = \phi_i^t + (\omega_t + w_{\omega_t}) \delta t$$

Where w_{v_t} is the additive noise for the linear velocity, and


w_{ω_t} is the additive noise for the angular velocity

Continuous motion example

- $\Delta t = 1 \text{ sec}$
- Plotting 1 sample/sec all the particles every 5 sec
- Constant linear velocity
- Angular velocity changes randomly every 10 sec



Continuous motion example



Prediction Examples Using a PF


Piecewise linear motion

(Translation and Rotation)


- Command success 70%
- Start at $[-8,0,0]$
- Translate by 4m
- Rotate by 30°
- Translate by 6m




Start [-8,0,0°]




Translate by 4m



Rotate by 30°



Translate by 6m




Update Examples Using a PF




Environment with two red doors

(uniform distribution)




Environment with two red doors


(Sensing the red door)



Sensing four walls




Four possible areas



Cooperative Localization

- Pose of the moving robot is estimated relative to the pose of the stationary robot.
Stationary Robot observes the **Moving Robot**.



Robot Tracker Returns:

$$\langle \rho, \theta, \phi \rangle$$


$$\mathbf{x}_{m_{est}}(k+1) = \begin{pmatrix} x_{m_{est}} \\ y_{m_{est}} \\ \theta_{m_{est}} \end{pmatrix} = \begin{pmatrix} x_s + \rho \cos(\theta + \theta_s) \\ y_s + \rho \sin(\theta + \theta_s) \\ \pi - (\phi - (\theta + \theta_s)) \end{pmatrix}$$

Laser-Based Robot Tracker



Robot Tracker Returns:


$$\langle \rho, \theta, \phi \rangle$$




Tracker Weighting Function

Update


The pdf of the M-Robot using ρ




The pdf of the M-Robot using θ



The pdf of the M-Robot using ϕ



The pdf of the M-Robot using T




$$(\sigma_\rho = 3, \sigma_\theta = 3, \sigma_\phi = 2)$$


$$W = \frac{1}{\sqrt{2\pi\sigma_\rho^2}} e^{-\frac{(\rho - \rho_i)^2}{\sigma_\rho^2}} \frac{1}{\sqrt{2\pi\sigma_\theta^2}} e^{-\frac{(\theta - \theta_i)^2}{\sigma_\theta^2}} \frac{1}{\sqrt{2\pi\sigma_\phi^2}} e^{-\frac{(\phi - \phi_i)^2}{\sigma_\phi^2}}$$




Example: Prediction




Example: Update



Example: Prediction



Example: Update



Variations on PF

- Add some particles uniformly
- Add some particles where the sensor indicates
- Add some jitter to the particles after propagation
- Combine EKFs to track landmarks



Keep in Mind:

- The number of particles increases with the dimension of the state space



Complexity results for SLAM

- n =number of map features
- Problem: naïve methods have high complexity
 - EKF models $O(n^2)$ covariance matrix
 - PF requires prohibitively many particles to characterize complex, interdependent distribution
- Solution: exploit conditional independencies
 - Feature estimates are independent given robot's path



References

- **Ioannis Rekleitis**. [A Particle Filter Tutorial for Mobile Robot Localization](#). Technical Report TR-CIM-04-02, Centre for Intelligent Machines, McGill University, Montreal, Québec, Canada, 2004.
- **Ioannis M. Rekleitis**, Gregory Dudek and Evangelos Miliotis. Multi-robot Cooperative Localization: A study of Trade-offs Between Efficiency and Accuracy. In *Proc. of Int. Conf. on Intelligent Robots and Systems*, pp. 2690-2695, Lausanne, Switzerland, Oct. 2002.
- [Sequential Monte Carlo Methods in Practice](#). Arnaud Doucet - Nando de Freitas - Neil Gordon (eds). Springer-Verlag, 2001, ISBN 0-387-95146-6.
- Isard M. and Blake A. [CONDENSATION - conditional density propagation for visual tracking](#). *Int. J. Computer Vision*, 29, 1, 5-28, 1998.
- F. Dellaert, W. Burgard, D. Fox, and S. Thrun. Using the condensation algorithm for robust, vision-based mobile robot localization. In *Conf. on Computer Vision & Pattern Recognition*, 1999.
- M. Montemerlo and S. Thrun. Fastslam 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges. In *SODA '01: Proc. of the 12th annual ACM-SIAM symposium on Discrete algorithms*, pages 735–744, 2001.
- Doucet, A., de Freitas, N., Murphy, K., and Russell, S. 2000. Rao-Blackwellised particle filtering for dynamic Bayesian networks. In *Uncertainty in Artificial Intelligence*, pp. 176–183.
- [Sim, R.\[Robert\], Elinas, P.\[Pantelis\], Little, J.J.\[James J.\]](#), **A Study of the Rao-Blackwellised Particle Filter for Efficient and Accurate Vision-Based SLAM**, *IJCV(74)*, No. 3, September 2007, pp. 303-318.
- Doucet, A.; Johansen, A.M.; "[A tutorial on particle filtering and smoothing: fifteen years later](#)". *Technical report, Department of Statistics, University of British Columbia*. December 2008.
- Arulampalam, M.S., Maskell, S., Gordon, N. and Clapp, T. [A Tutorial on Particle Filters for nonlinear/non-Gaussian Bayesian Tracking](#). *IEEE Trans. Signal Processing*, Vol. 50, No. 2, 2002. p.174-188.
- [Sequential Monte Carlo Methods Homepage](#)
- [Monte-Carlo Localization-in-action page](#)



Questions

- For more information on PF:

<http://www.cim.mcgill.ca/~yiannis/ParticleTutorial.html>

- For offline questions:

yiannis@cim.mcgill.ca

