

Robot Mapping


Short Introduction to Particle Filters and Monte Carlo Localization

**Gian Diego Tipaldi, Luciano Spinello,
Wolfram Burgard**

Gaussian Filters


- The Kalman filter and its variants can only model **Gaussian distributions**

$$p(x) = \det(2\pi\Sigma)^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu)\right)$$




Motivation

- Goal: approach for dealing with **arbitrary distributions**



Key Idea: Samples

- Use **multiple samples** to represent arbitrary distributions



Particle Set

- Set of weighted samples

$$\mathcal{X} = \left\{ \left\langle x^{[j]}, w^{[j]} \right\rangle \right\}_{j=1, \dots, J}$$

**state
hypothesis**



**importance
weight**

- The samples represent the posterior

$$p(x) = \sum_{j=1}^J w^{[j]} \delta_{x^{[j]}}(x)$$

Particles for Approximation

- Particles for function approximation




- The more particles fall into a region, the higher the probability of the region

How to obtain such samples?

Closed Form Sampling is Only Possible for a Few Distributions

- Example: Gaussian




$$x \leftarrow \frac{1}{2} \sum_{i=1}^{12} \text{rand}(-\sigma, \sigma)$$


How to sample from **other** distributions?

Importance Sampling Principle

- We can use a different distribution g to generate samples from f
- Account for the “differences between g and f ” using a weight $w = f / g$
- target f
- proposal g
- Pre-condition:
 $f(x) > 0 \rightarrow g(x) > 0$



Importance Sampling Principle



Particle Filter

- Recursive Bayes filter
- Non-parametric approach
- Models the distribution by samples
- Prediction: draw from the proposal
- Correction: weighting by the ratio of target and proposal

**The more samples we use,
the better is the estimate!**

Particle Filter Algorithm

1. Sample the particles using the proposal distribution

$$x_t^{[j]} \sim \pi(x_t \mid \dots)$$

2. Compute the importance weights

$$w_t^{[j]} = \frac{\text{target}(x_t^{[j]})}{\text{proposal}(x_t^{[j]})}$$

- Resampling: Draw sample i with probability $w_t^{[i]}$ and repeat J times

Particle Filter Algorithm

Particle_filter($\mathcal{X}_{t-1}, u_t, z_t$):

1: $\bar{\mathcal{X}}_t = \mathcal{X}_t = \emptyset$

2: *for* $j = 1$ *to* J *do*

3: *sample* $x_t^{[j]} \sim \pi(x_t)$

4: $w_t^{[j]} = \frac{p(x_t^{[j]})}{\pi(x_t^{[j]})}$

5: $\bar{\mathcal{X}}_t = \bar{\mathcal{X}}_t + \langle x_t^{[j]}, w_t^{[j]} \rangle$

6: *endfor*

7: *for* $j = 1$ *to* J *do*

8: *draw* $i \in 1, \dots, J$ *with probability* $\propto w_t^{[i]}$

9: *add* $x_t^{[i]}$ *to* \mathcal{X}_t

10: *endfor*

11: *return* \mathcal{X}_t

Monte Carlo Localization

- Each particle is a pose hypothesis
- Proposal is the motion model

$$x_t^{[j]} \sim p(x_t \mid x_{t-1}, u_t)$$

- Correction via the observation model


$$w_t^{[j]} = \frac{\text{target}}{\text{proposal}} \propto p(z_t \mid x_t, m)$$

Particle Filter for Localization

Particle_filter($\mathcal{X}_{t-1}, u_t, z_t$):

- 1: $\bar{\mathcal{X}}_t = \mathcal{X}_t = \emptyset$
- 2: *for* $j = 1$ *to* J *do*
- 3: *sample* $x_t^{[j]} \sim p(x_t \mid u_t, x_{t-1}^{[j]})$
- 4: $w_t^{[j]} = p(z_t \mid x_t^{[j]})$
- 5: $\bar{\mathcal{X}}_t = \bar{\mathcal{X}}_t + \langle x_t^{[j]}, w_t^{[j]} \rangle$
- 6: *endfor*
- 7: *for* $j = 1$ *to* J *do*
- 8: *draw* $i \in 1, \dots, J$ *with probability* $\propto w_t^{[i]}$
- 9: *add* $x_t^{[i]}$ *to* \mathcal{X}_t
- 10: *endfor*
- 11: *return* \mathcal{X}_t


Application: Particle Filter for Localization (Known Map)




Resampling

- Draw sample i with probability $w_t^{[i]}$. Repeat J times.
- Informally: “Replace unlikely samples by more likely ones”
- Survival of the fittest
- “Trick” to avoid that many samples cover unlikely states
- Needed as we have a limited number of samples

Resampling



- Roulette wheel
- Binary search
- $O(J \log J)$




- Stochastic universal sampling
- Low variance
- $O(J)$

Low Variance Resampling


Low_variance_resampling($\mathcal{X}_t, \mathcal{W}_t$):

```
1:    $\bar{\mathcal{X}}_t = \emptyset$ 
2:    $r = \text{rand}(0; J^{-1})$ 
3:    $c = w_t^{[1]}$ 
4:    $i = 1$ 
5:   for  $j = 1$  to  $J$  do
6:      $U = r + (j - 1)J^{-1}$ 
7:     while  $U > c$ 
8:        $i = i + 1$ 
9:        $c = c + w_t^{[i]}$ 
10:    endwhile
11:    add  $x_t^{[i]}$  to  $\bar{\mathcal{X}}_t$ 
12:  endfor
13:  return  $\bar{\mathcal{X}}_t$ 
```




initialization

Courtesy: Thrun, Burgard, Fox 19




observation

Courtesy: Thrun, Burgard, Fox 20




resampling

Courtesy: Thrun, Burgard, Fox 21



motion update

Courtesy: Thrun, Burgard, Fox 22




measurement

Courtesy: Thrun, Burgard, Fox 23




weight update

Courtesy: Thrun, Burgard, Fox 24



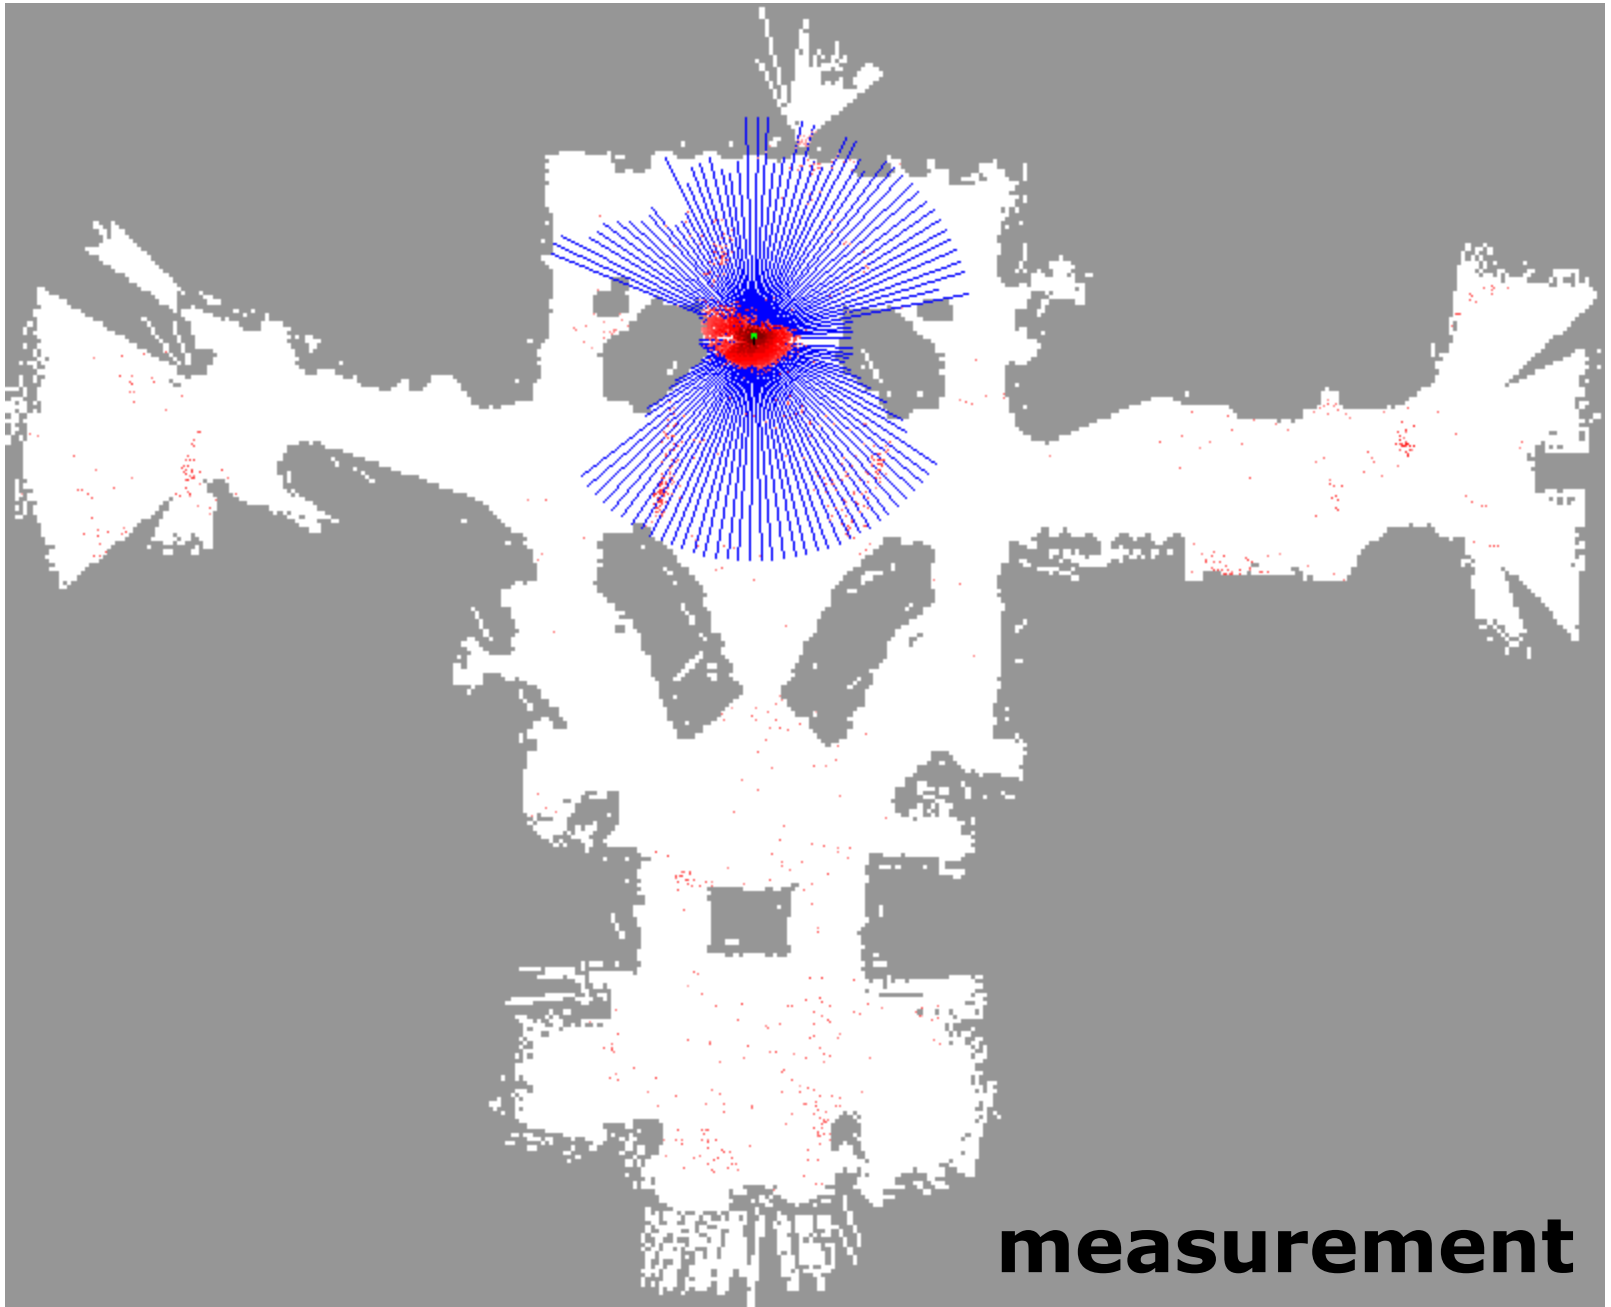
resampling

Courtesy: Thrun, Burgard, Fox 25



motion update

Courtesy: Thrun, Burgard, Fox 26




measurement

Courtesy: Thrun, Burgard, Fox 27




weight update

Courtesy: Thrun, Burgard, Fox 28




Courtesy: Thrun, Burgard, Fox 29




motion update

Courtesy: Thrun, Burgard, Fox 30




measurement

Courtesy: Thrun, Burgard, Fox 31




weight update

Courtesy: Thrun, Burgard, Fox 32




Courtesy: Thrun, Burgard, Fox 33



motion update

Courtesy: Thrun, Burgard, Fox 34



measurement

Courtesy: Thrun, Burgard, Fox 35

Summary – Particle Filters

- Particle filters are non-parametric, recursive Bayes filters
- Posterior is represented by a set of weighted samples
- Proposal to draw the samples for $t+1$
- Weight to account for the differences between the proposal and the target
- Work well in low-dimensional spaces

Summary – PF Localization

- Particles are propagated according to the motion model
- They are weighted according to the likelihood of the observation
- Called: Monte-Carlo localization (MCL)
- MCL is the gold standard for mobile robot localization today

Literature

On Monte Carlo Localization

- Thrun et al. "Probabilistic Robotics", Chapter 8.3

On the particle filter

- Thrun et al. "Probabilistic Robotics", Chapter 3

On motion and observation models

- Thrun et al. "Probabilistic Robotics", Chapters 5 & 6

Slide Information

- These slides have been created by Cyrill Stachniss as part of the robot mapping course taught in 2012/13 and 2013/14. I created this set of slides partially extending existing material of Edwin Olson, Pratik Agarwal, and myself.
- I tried to acknowledge all people that contributed image or video material. In case I missed something, please let me know. If you adapt this course material, please make sure you keep the acknowledgements.
- Feel free to use and change the slides. If you use them, I would appreciate an acknowledgement as well. To satisfy my own curiosity, I appreciate a short email notice in case you use the material in your course.
- My video recordings are available through YouTube:
http://www.youtube.com/playlist?list=PLgnQpQtFTOGQrZ4O5QzbIHgl3b1JHimN_&feature=g-list