

CSE-490R

Robotics

SLAM + Fast-SLAM

(Partial slides borrowed from Dieter Fox, Wolfram Burgard & Cyril Stachniss)

What is SLAM?

- **Localization:** Estimate current pose given map, controls and observations

$$p(x_t \mid u_{1:t}, z_{1:t}, m)$$

- **Mapping:** Build map given poses and observations

$$p(m \mid x_{1:t}, z_{1:t})$$

- **Simultaneous Localization And Mapping (SLAM):** Find poses and map given controls and observations

$$p(x_{1:t}, m \mid u_{1:t}, z_{1:t})$$

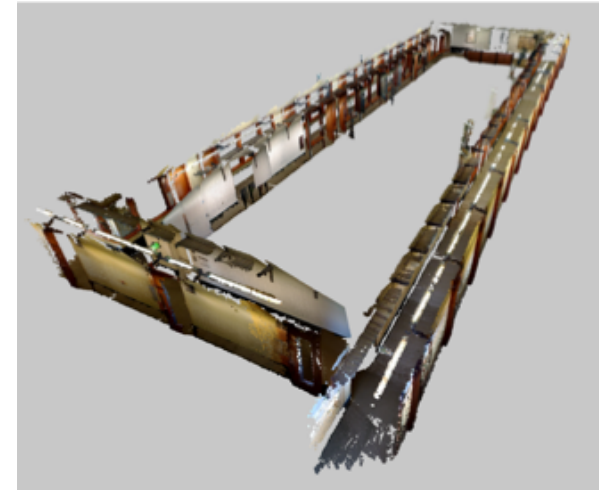
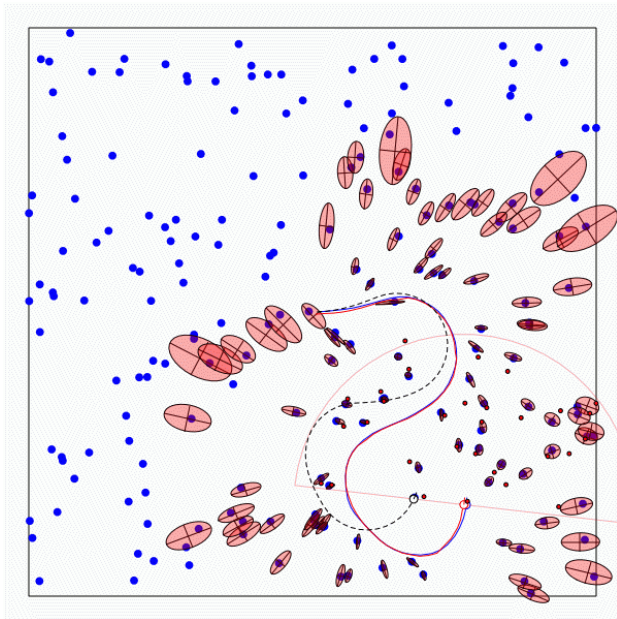
SLAM

- A robot moving through an unknown, static environment
- **Given:**
 - ▣ The robot's controls
 - ▣ Observations of nearby features
- **Estimate:**
 - ▣ Map of environment
 - ▣ Path of the robot



Map representations

- Typical representations are:
 - ▣ Feature-based
 - ▣ Grid maps (occupancy maps)
 - ▣ 3D representations (voxels, surfels, octrees etc)

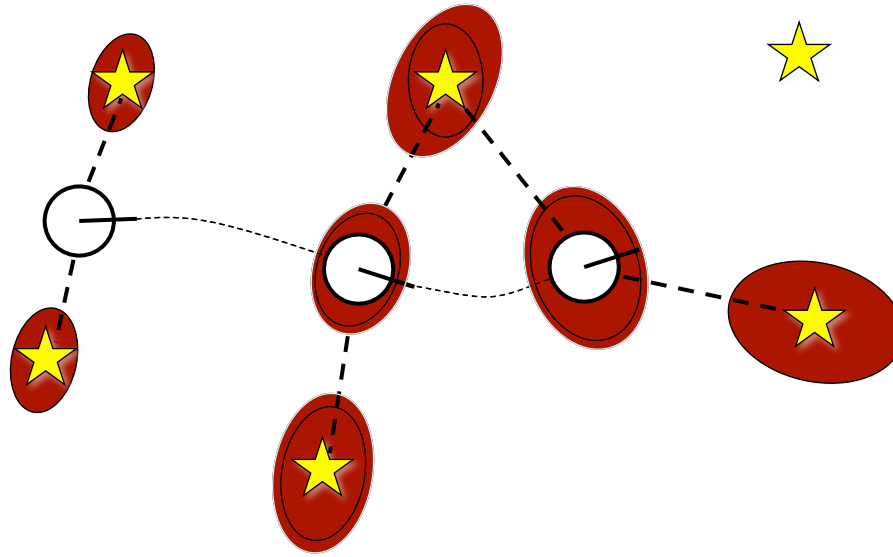


Why is SLAM hard?

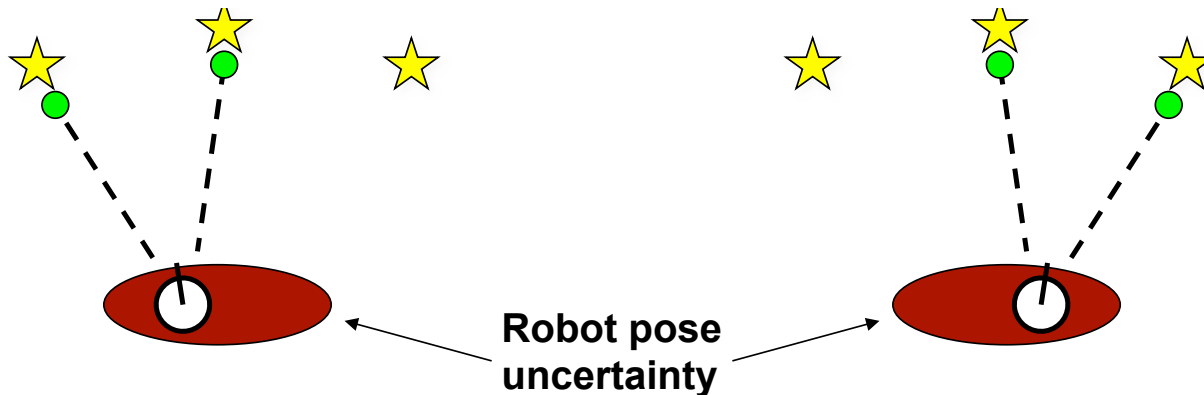
- SLAM is the task of building a map while estimating the pose of the robot relative to this map
- Chicken-or-egg problem:
 - ▣ A map is needed to localize the robot
 - ▣ A pose estimate is needed to build a map

Why is SLAM hard?

SLAM: robot path and map are both **unknown!**



Robot path error correlates errors in the map



Particle Filters

- Represent belief by random **samples**
- Sampling Importance Resampling (SIR) principle
 - ▣ Draw the new generation of particles
 - ▣ Assign an importance weight to each particle
 - ▣ Resampling
- Applications are localization, tracking, ...

Particle Filter Algorithm

1. Sample the particles from the proposal distribution

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

2. Compute the importance weights

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

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

Particle Filters for SLAM

- Localization: state space is $\langle x, y, \theta \rangle$
- SLAM: state space is $\langle x, y, \theta, map \rangle$
 - ▣ For grid maps: $\langle c_{11}, c_{12}, \dots, c_{1n}, c_{21}, \dots, c_{nm} \rangle$
 - ▣ For feature maps: $\langle l_1, l_2, \dots, l_m \rangle$
- **Problem:** The number of particles needed to represent a posterior grows exponentially with the dimension of the state space!

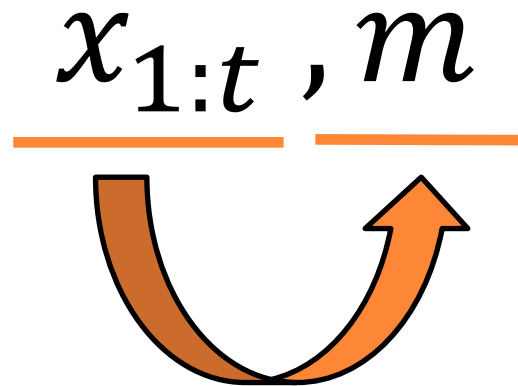
Dependencies

- Is there a dependency between the dimensions of the state space?
- If so, can we use the dependency to solve the problem more efficiently?
- In the SLAM context
 - ▣ The map depends on the **poses** of the robot.
 - ▣ We know how to build a map **given** the position of the sensor is **known**.

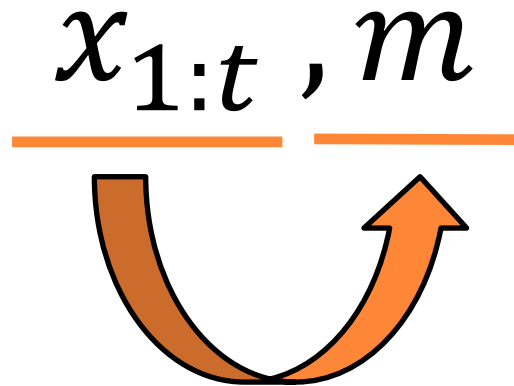
Can We Exploit Dependencies Between
the Different Dimensions of the State
Space?

$$x_{1:t}, m$$

If We Know the Poses of the Robot,
Mapping is Easy!



Key Idea



If we use the particle set only to model the robot's path, each sample is a path hypothesis. For each particle, we can compute an individual map using it's path.

Rao-Blackwellization

- Factorization to exploit dependencies between variables:


$$p(a, b) = p(b \mid a) p(a)$$

- If $p(b \mid a)$ can be computed efficiently, represent only $p(a)$ with samples and compute $p(b \mid a)$ for every sample

Rao-Blackwellization for SLAM

- Factorization of the SLAM posterior

poses map observations & controls


$$p(x_{1:t}, m \mid z_{1:t}, u_{1:t}) =$$

First introduced for SLAM by Murphy in 1999

Rao-Blackwellization for SLAM

- Factorization of the SLAM posterior

poses map observations & controls

$$p(x_{1:t}, m \mid z_{1:t}, u_{1:t}) = p(x_{1:t} \mid z_{1:t}, u_{1:t}) p(m \mid x_{1:t}, z_{1:t})$$

path posterior map posterior

First introduced for SLAM by Murphy in 1999

Rao-Blackwellization for SLAM

- Factorization of the SLAM posterior


$$p(x_{1:t}, m \mid z_{1:t}, u_{1:t}) = \underbrace{p(x_{1:t} \mid z_{1:t}, u_{1:t}) p(m \mid x_{1:t}, z_{1:t})}$$

Grid cells are conditionally independent given the poses

First exploited in FastSLAM by Montemerlo et al., 2002

Rao-Blackwellization for SLAM

- Factorization of the SLAM posterior

$$\begin{aligned} p(x_{1:t}, m \mid z_{1:t}, u_{1:t}) &= \\ p(x_{1:t} \mid z_{1:t}, u_{1:t}) & p(m \mid x_{1:t}, z_{1:t}) \\ p(x_{1:t} \mid z_{1:t}, u_{1:t}) & \prod_{i=1}^N p(m^i \mid x_{1:t}, z_{1:t}) \end{aligned}$$


**Grid cells are conditionally
independent given the poses**

First exploited in FastSLAM by Montemerlo et al., 2002

Rao-Blackwellization for SLAM

- Factorization of the SLAM posterior

$$p(x_{1:t}, m \mid z_{1:t}, u_{1:t}) =$$
$$\frac{p(x_{1:t} \mid z_{1:t}, u_{1:t})}{\text{particle filter for localization}} \frac{p(m \mid x_{1:t}, z_{1:t})}{\text{occupancy grid mapping}}$$
$$p(x_{1:t} \mid z_{1:t}, u_{1:t}) \prod_{i=1}^N p(m^i \mid x_{1:t}, z_{1:t})$$

First exploited in FastSLAM by Montemerlo et al., 2002

Modeling the Robot's Path

- Sample-based representation for

$$p(x_{1:t} | z_{1:t}, u_{1:t})$$

- Each sample is a path hypothesis

x_1
↑

starting location,
typically (0,0,0)

x_2
↑

pose hypothesis
at time $t=2$

x_3

...

- Past poses of a sample are not revised
- No need to maintain past poses in the sample set

FastSLAM

- Proposed by Montemerlo et al. in 2002 (for landmark based SLAM)
- Each particle has a pose and a map

Particle
1

x, y, θ

Occupancy grid map

Particle
2

x, y, θ

Occupancy grid map

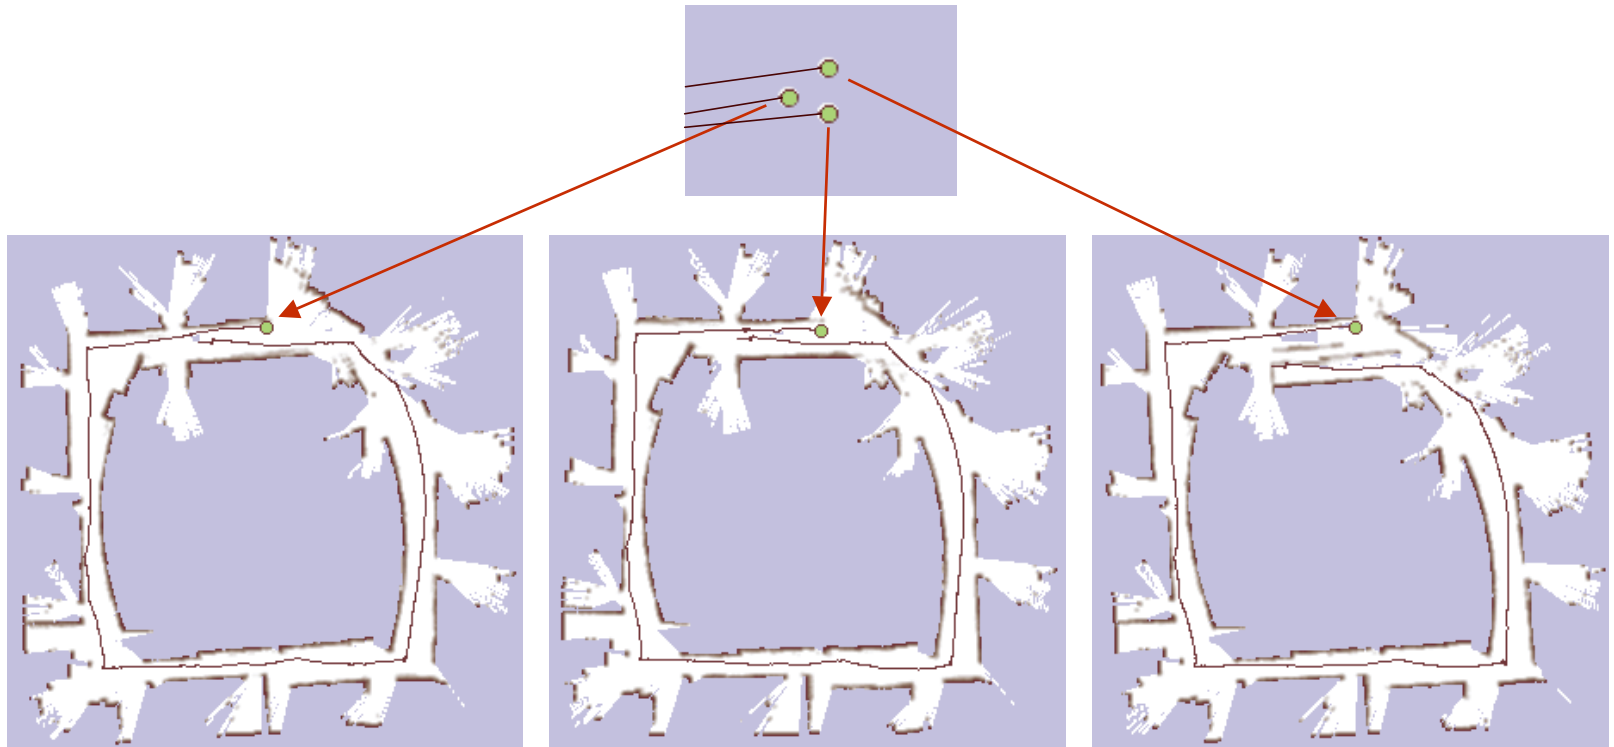
⋮

Particle
 N

x, y, θ

Occupancy grid map

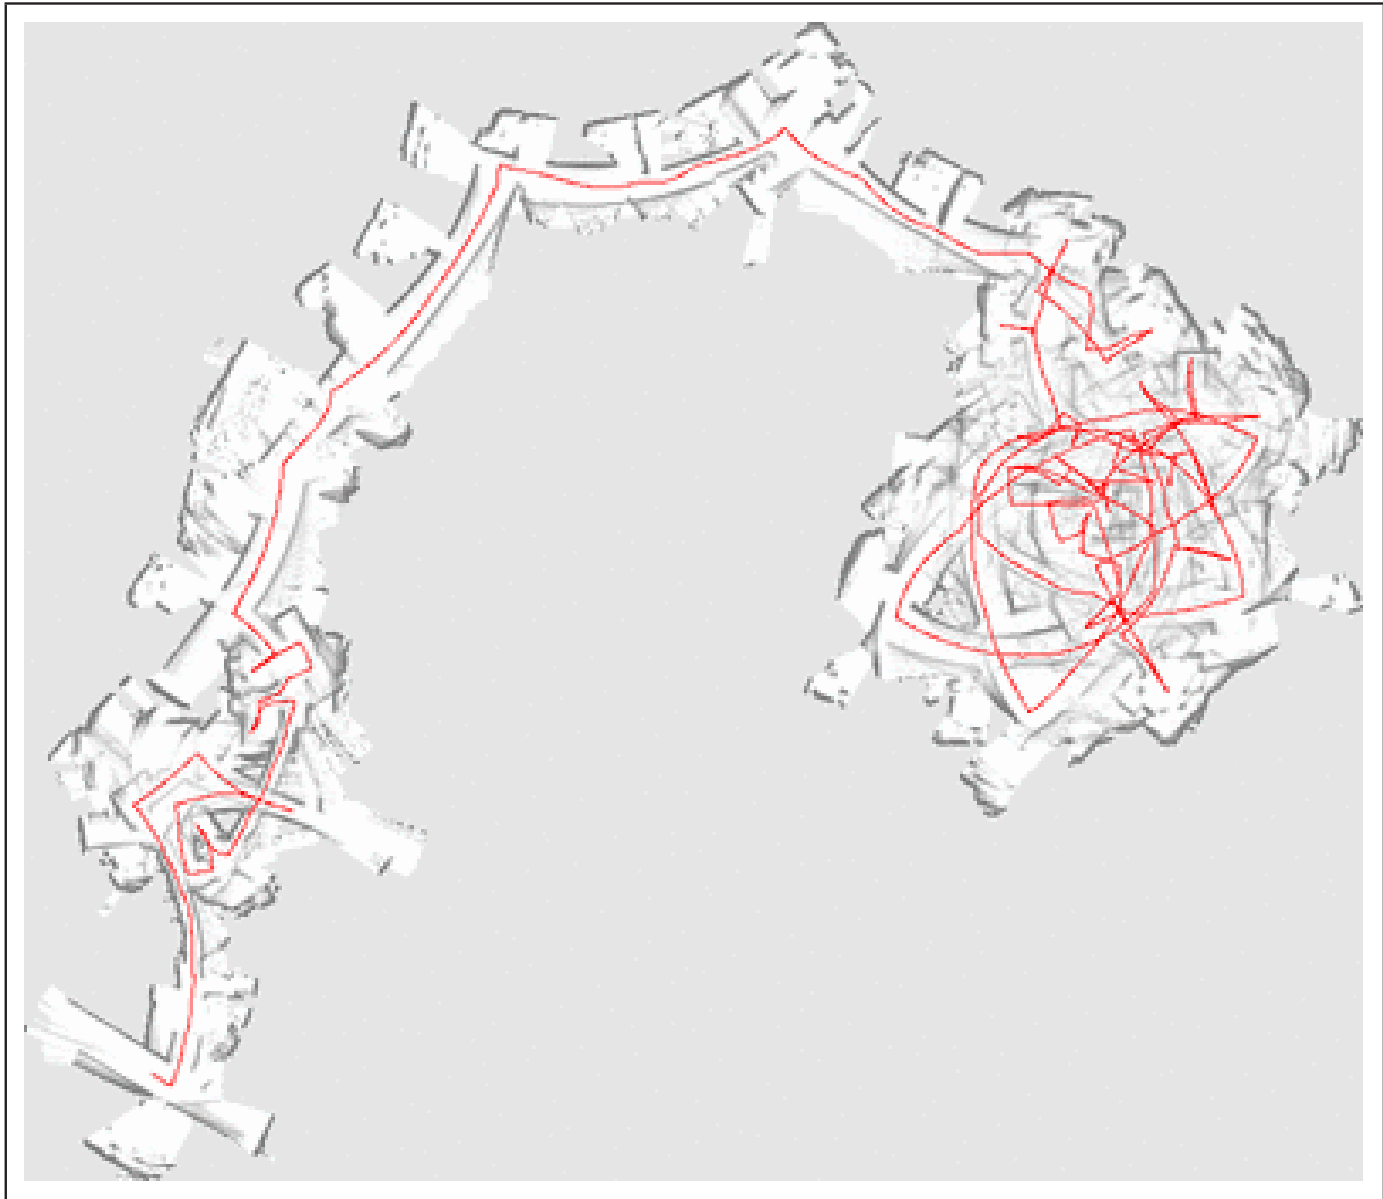
FastSLAM – Particle representation



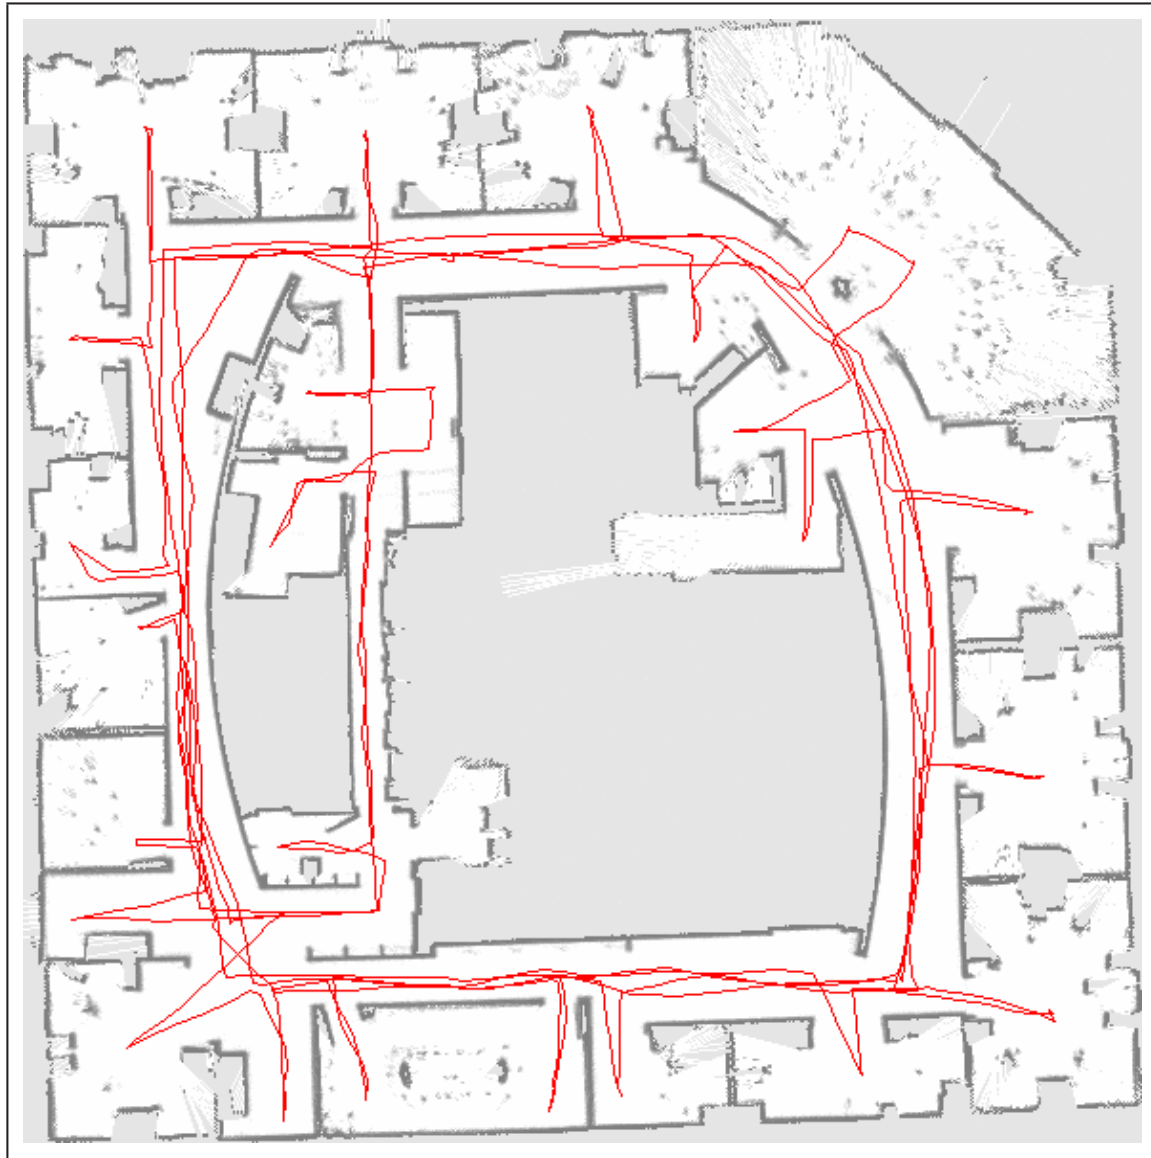
FastSLAM Algorithm

```
1:   Algorithm FastSLAM_occupancy_grids( $\mathcal{X}_{t-1}, u_t, z_t$ ):  
2:      $\bar{\mathcal{X}}_t = \mathcal{X}_t = \emptyset$   
3:     for  $k = 1$  to  $M$  do  
4:        $x_t^{[k]} = \text{sample\_motion\_model}(u_t, x_{t-1}^{[k]})$   
5:        $w_t^{[k]} = \text{measurement\_model\_map}(z_t, x_t^{[k]}, m_{t-1}^{[k]})$   
5:        $m_t^{[k]} = \text{updated\_occupancy\_grid}(z_t, x_t^{[k]}, m_{t-1}^{[k]})$   
6:        $\bar{\mathcal{X}}_t = \bar{\mathcal{X}}_t + \langle x_t^{[k]}, m_t^{[k]}, w_t^{[k]} \rangle$   
7:     endfor  
8:     for  $k = 1$  to  $M$  do  
9:       draw  $i$  with probability  $\propto w_t^{[i]}$   
10:      add  $\langle x_t^{[i]}, m_t^{[i]} \rangle$  to  $\mathcal{X}_t$   
11:    endfor  
12:    return  $\mathcal{X}_t$ 
```

Pure odometry



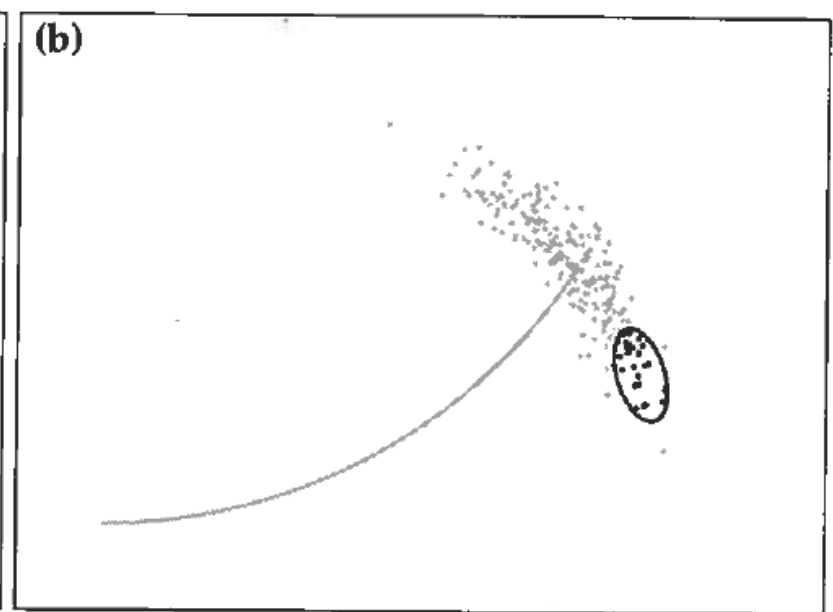
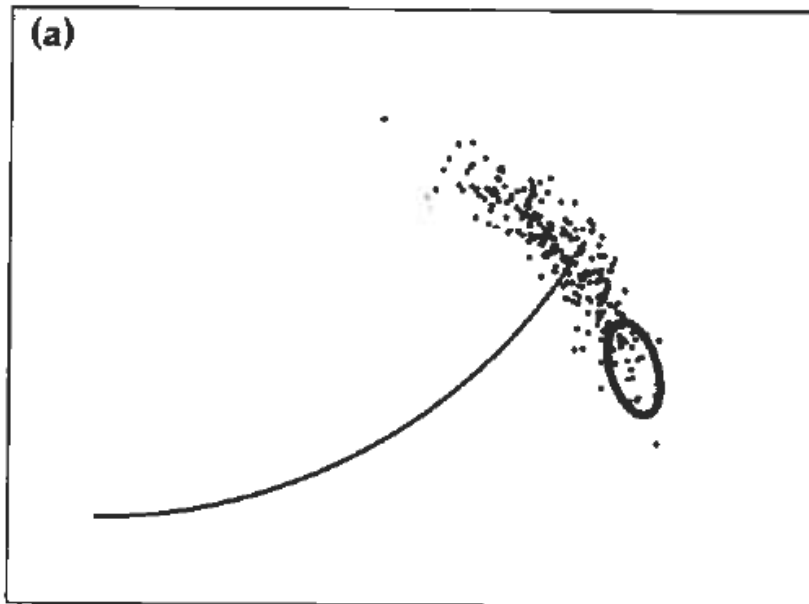
FastSLAM – Best particle



Weakness of FastSLAM 1.0

□ Proposal Distribution

□ Importance weighting



FastSLAM 1.0 to FastSLAM 2.0

- FastSLAM 1.0 uses the motion model as the proposal distribution

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

- FastSLAM 2.0 **considers also the measurements during sampling**
- Especially useful if an accurate sensor is used (compared to the motion noise)

FastSLAM 2.0 (Informally)

- FastSLAM 2.0 samples from

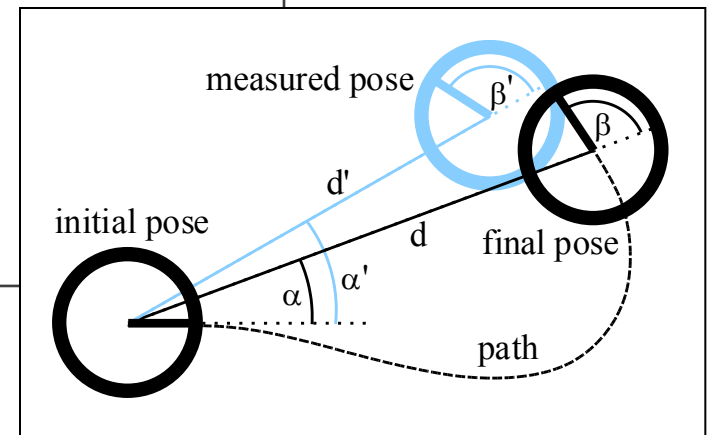
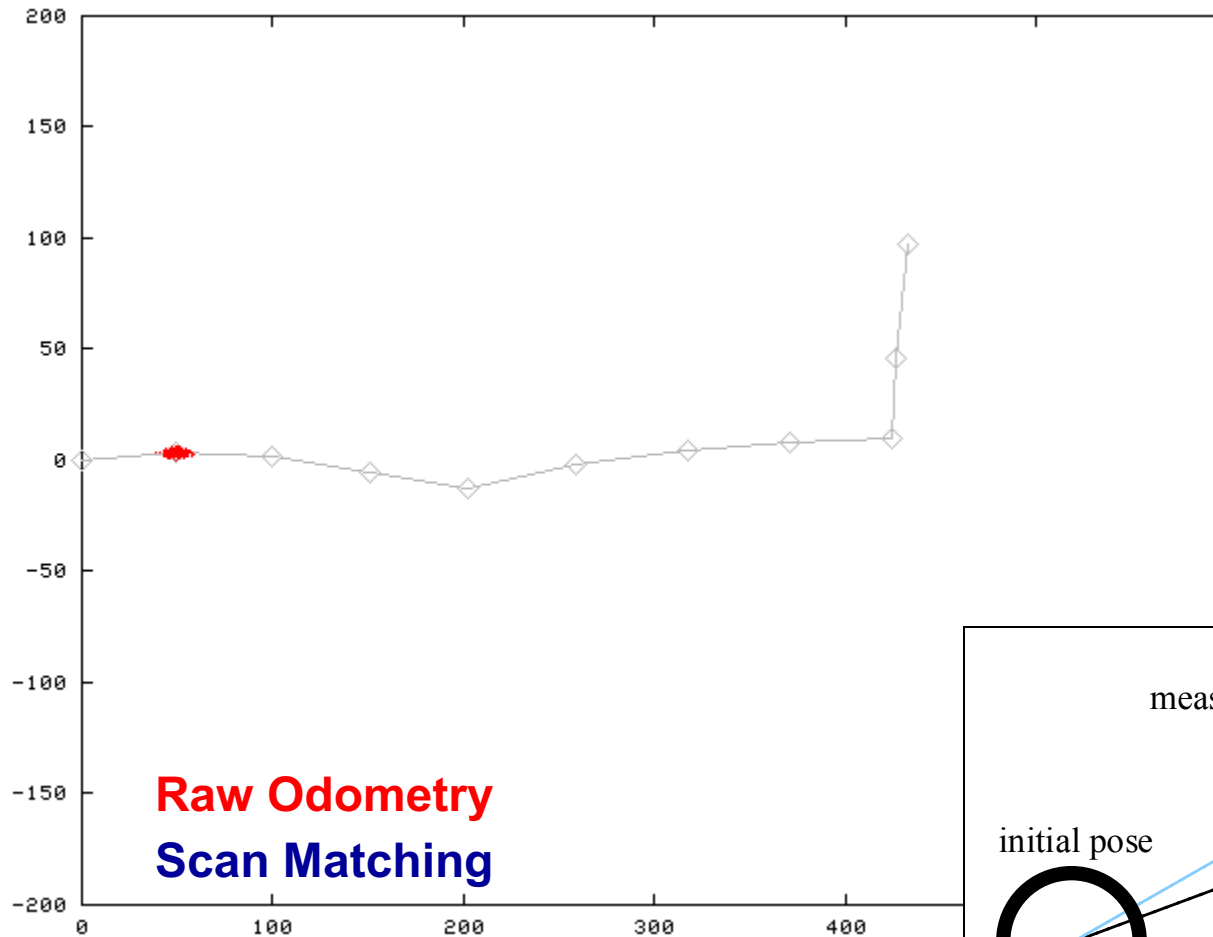
$$x_t^{[k]} \sim p(x_t \mid x_{1:t-1}^{[k]}, u_{1:t}, z_{1:t})$$

- Results in a more peaked proposal distribution
- Less particles are required
- More robust and accurate
- But more complex...

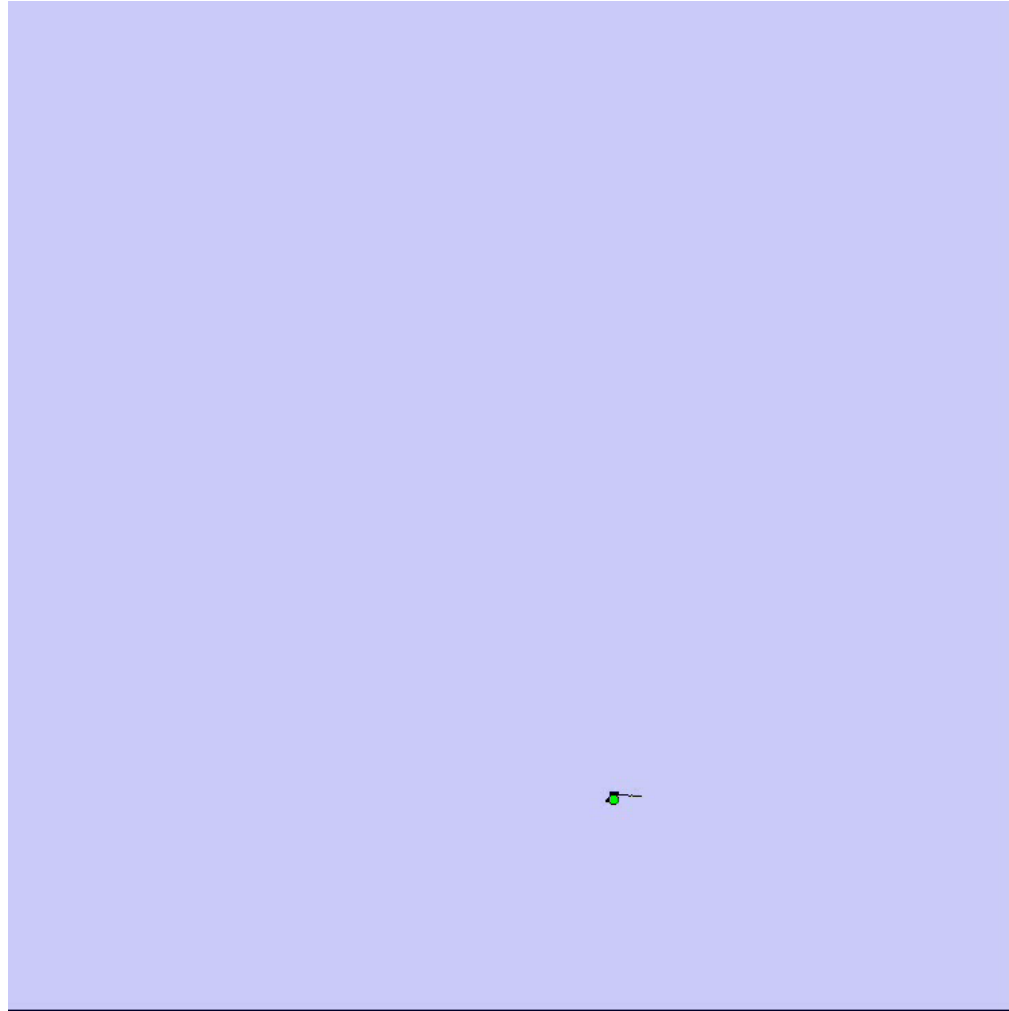
Generating better proposals

- Use scan-matching to compute highly accurate odometry measurements from consecutive range scans.
- Use the improved odometry in the prediction step to get highly accurate proposal distributions.

Motion Model for Scan Matching



Rao-Blackwellized Mapping with Scan-Matching



Map: Intel Research Lab Seattle

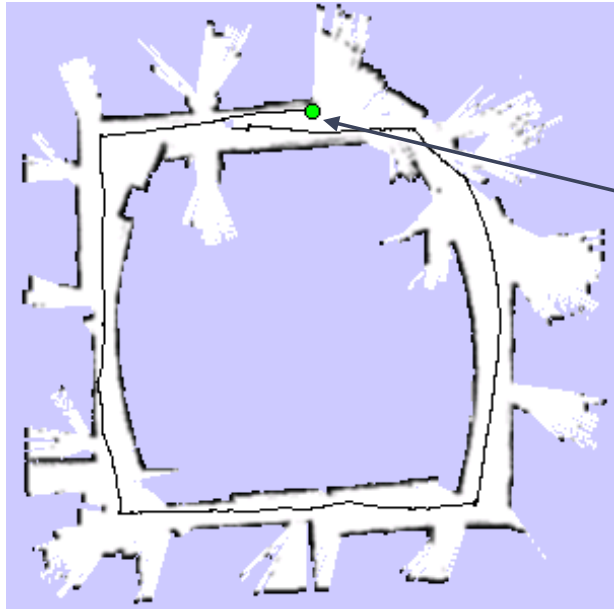
Loop Closure

- Loop closure involves
 - ▣ Recognizing when the robot has returned to a previously mapped region
 - ▣ Using this information to reduce the uncertainty in the map estimate
- Without loop closure, the uncertainty can grow without bounds

Loop Closure in FastSLAM

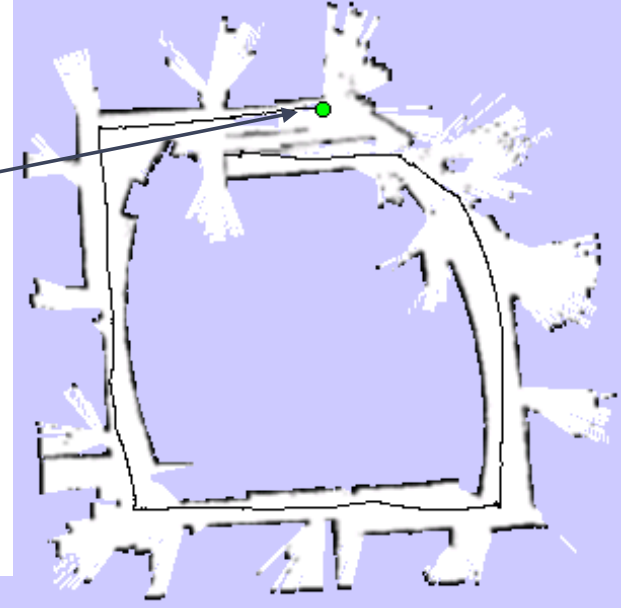
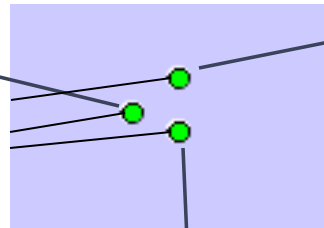
- Each particle has it's own map
- Maps which agree to closing the loop are weighed higher than others
- These maps are more likely to be resampled
- **Key:** Need diversity of paths/particles/maps

Loop Closure Example

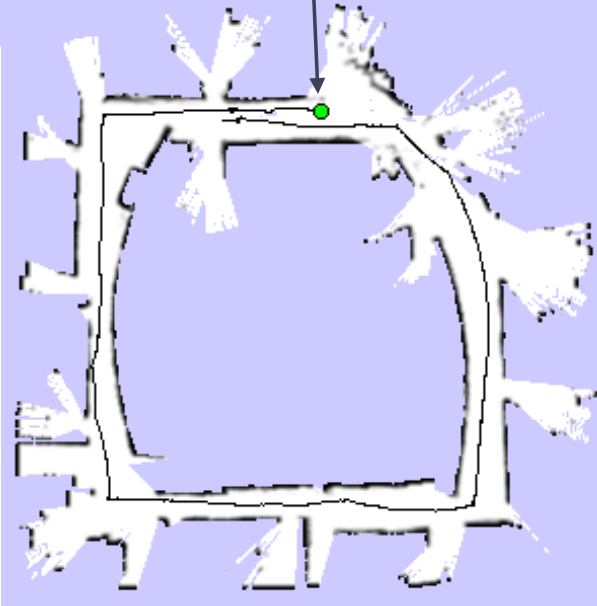


map of particle 1

3 particles

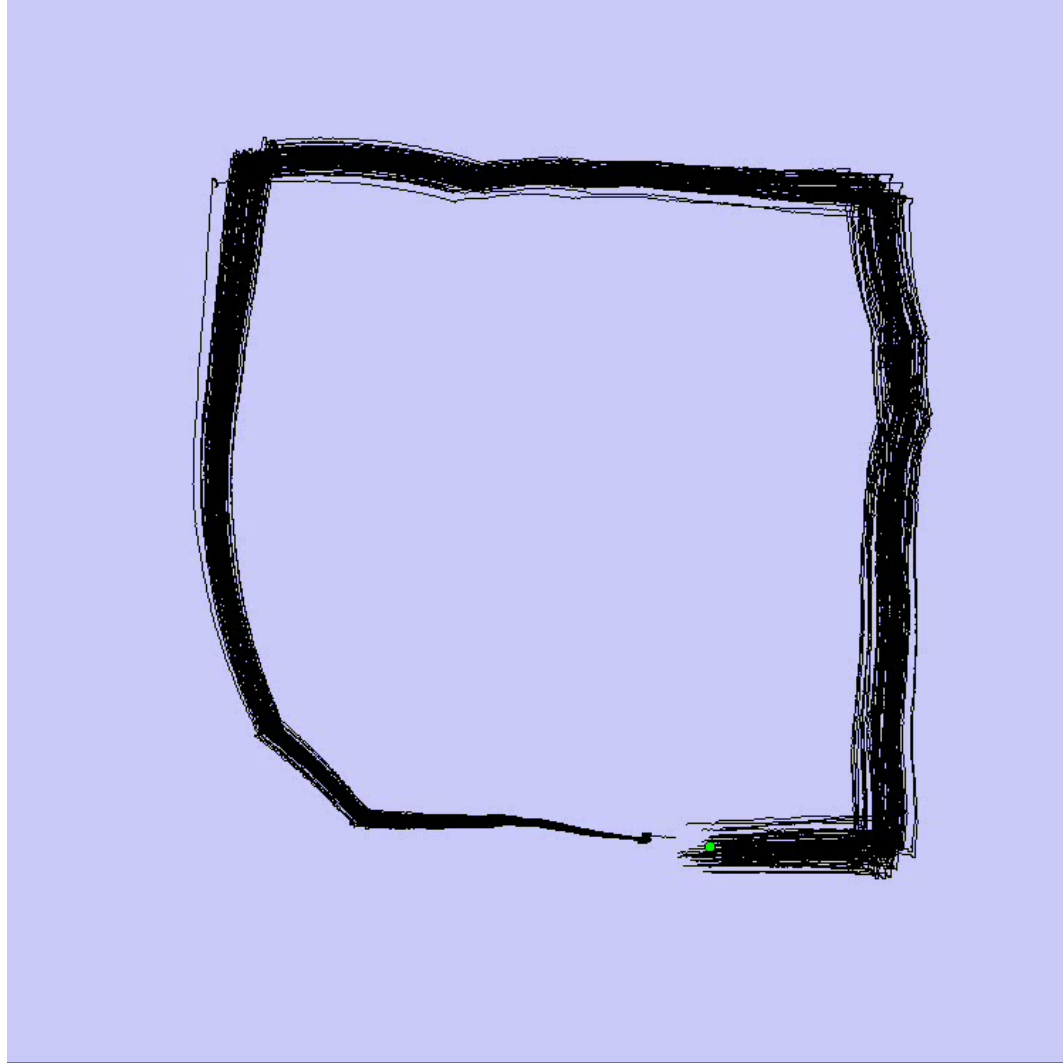


map of particle 3



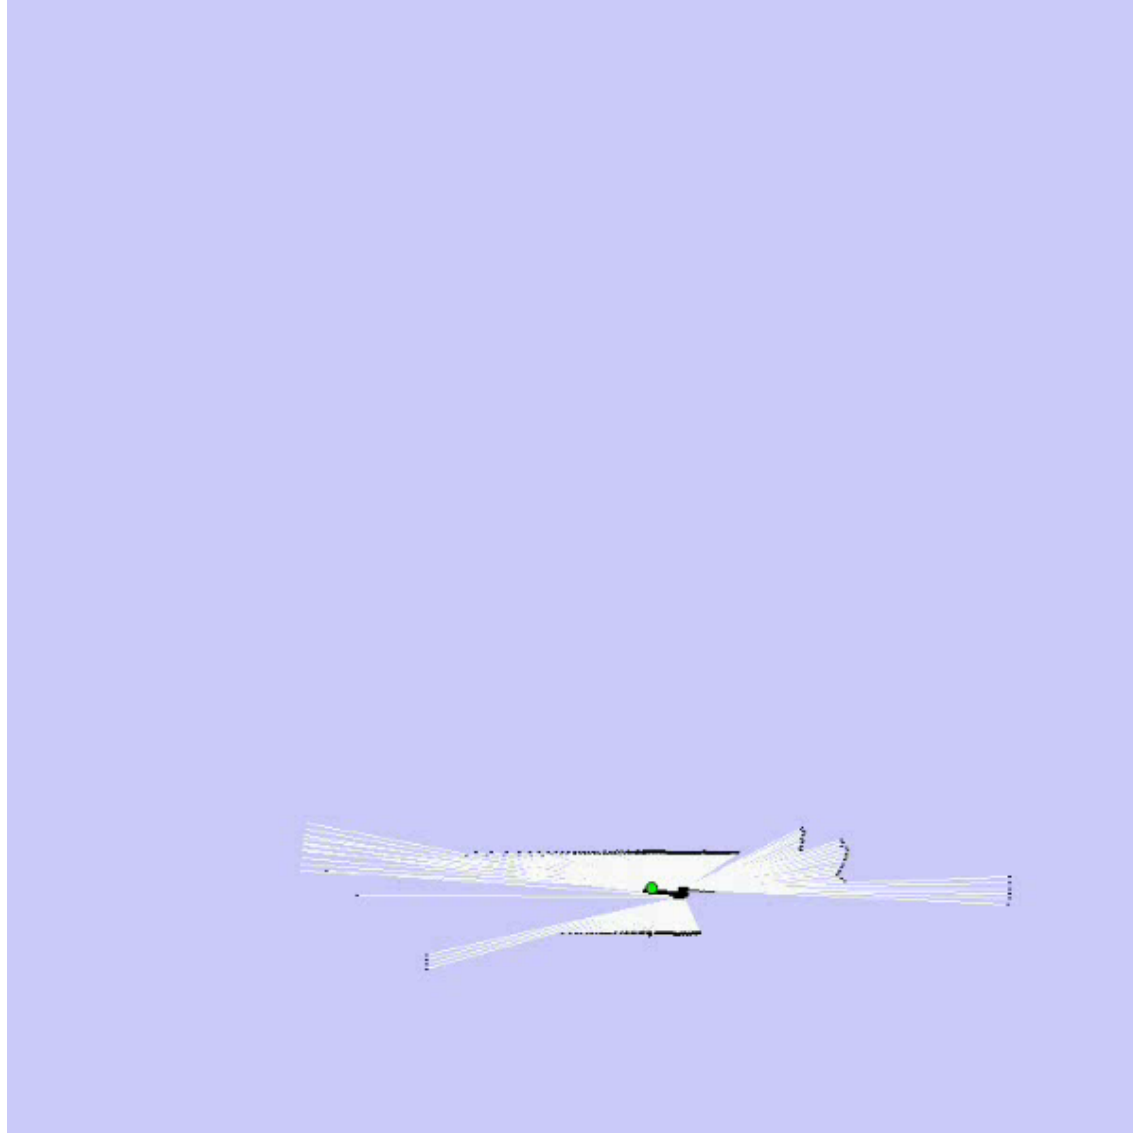
map of particle 2

Rao-Blackwellized Mapping with Scan-Matching



Map: Intel Research Lab Seattle

Rao-Blackwellized Mapping with Scan-Matching



Example (Intel Lab)



- **15 particles**
- four times faster than real-time P4, 2.8GHz
- 5cm resolution during scan matching
- 1cm resolution in final map

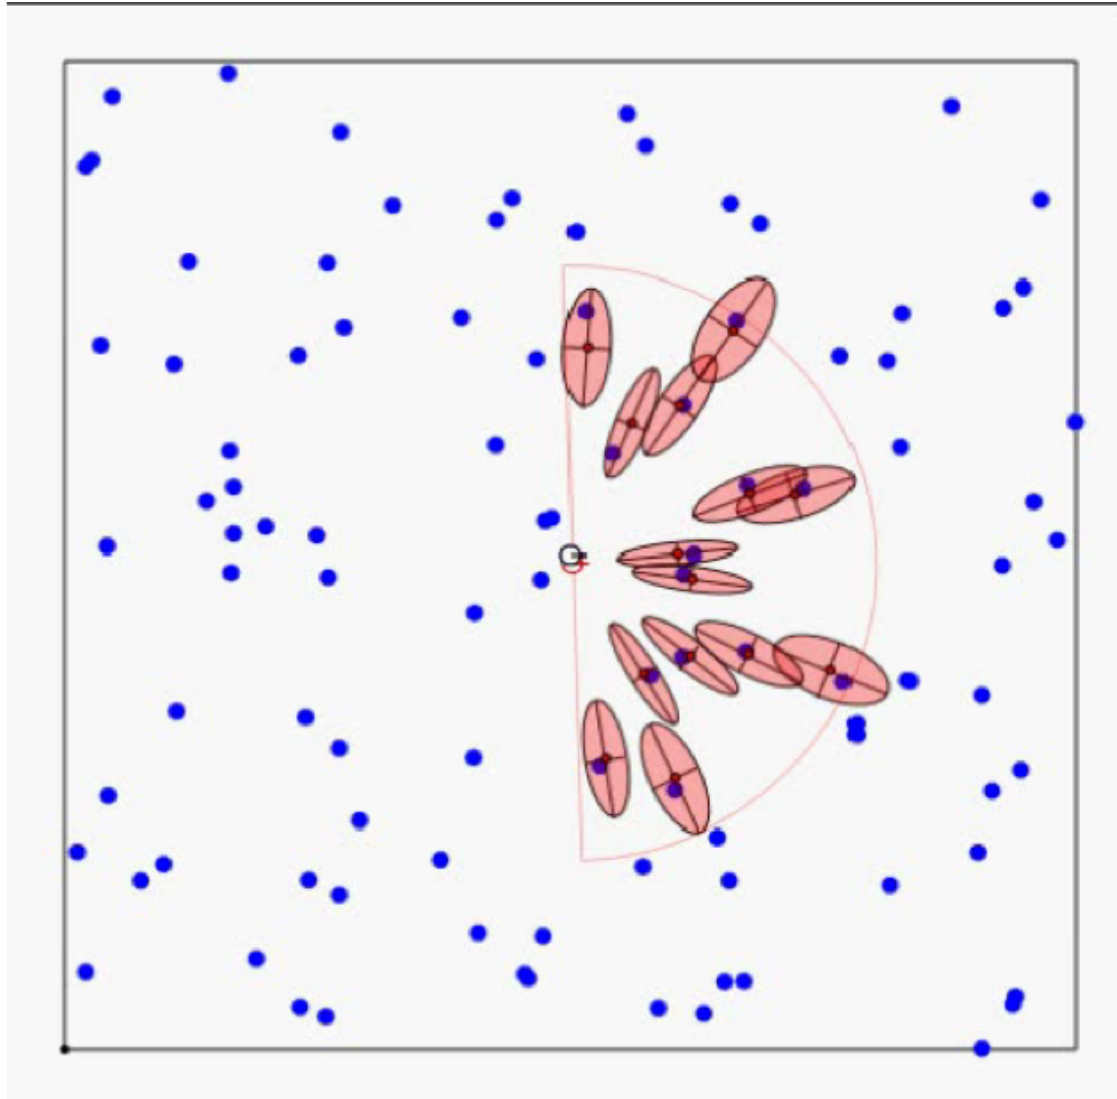
Outdoor Campus Map



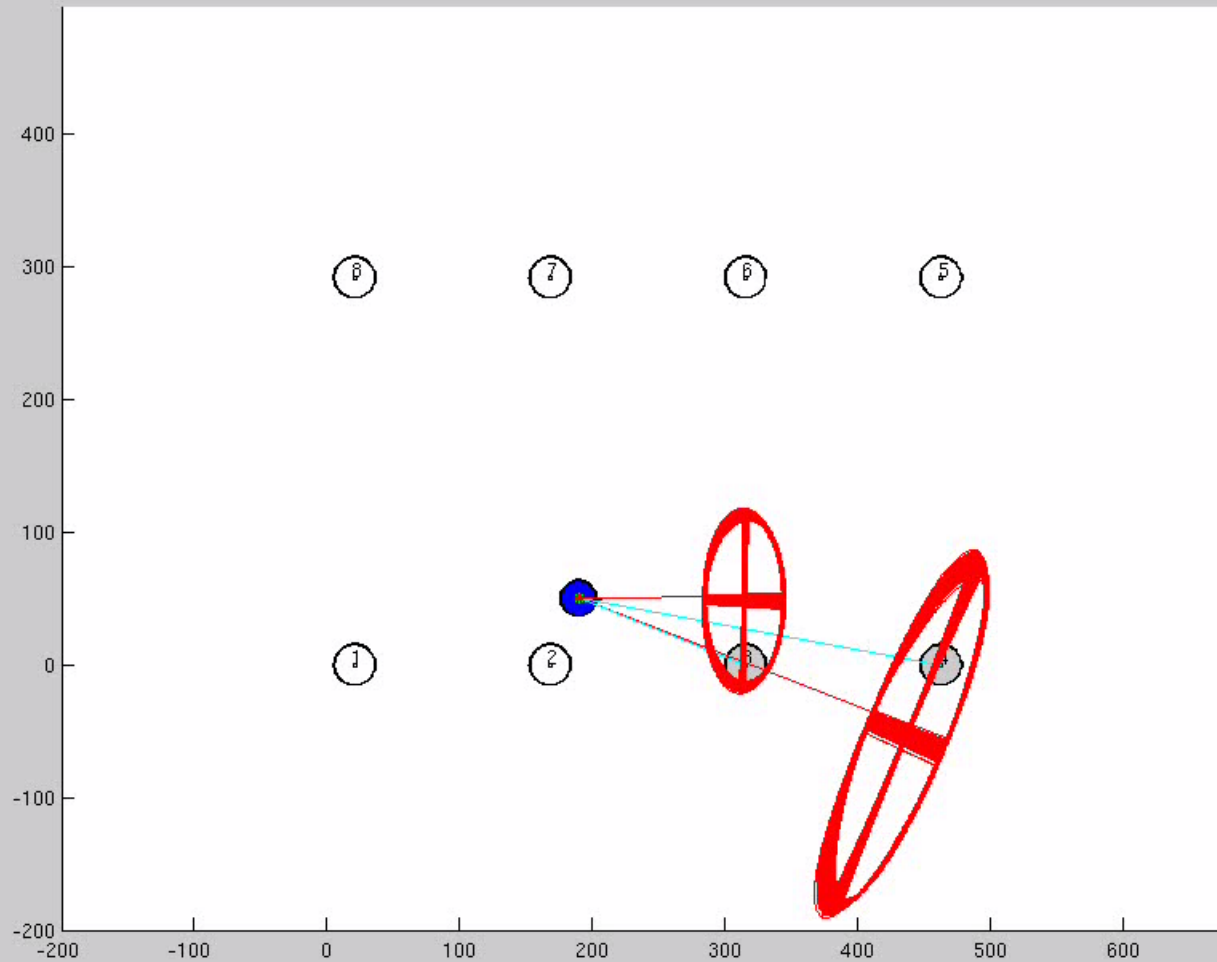
- **30 particles**
- 250x250m²
- 1.088 miles (odometry)
- 20cm resolution during scan matching
- 30cm resolution in final map

FastSLAM for feature-based maps

FastSLAM in Action



FastSLAM – Video – All Maps



Results – Victoria Park

- 4 km traverse
- < 2.5 m RMS position error
- 100 particles

Blue = GPS

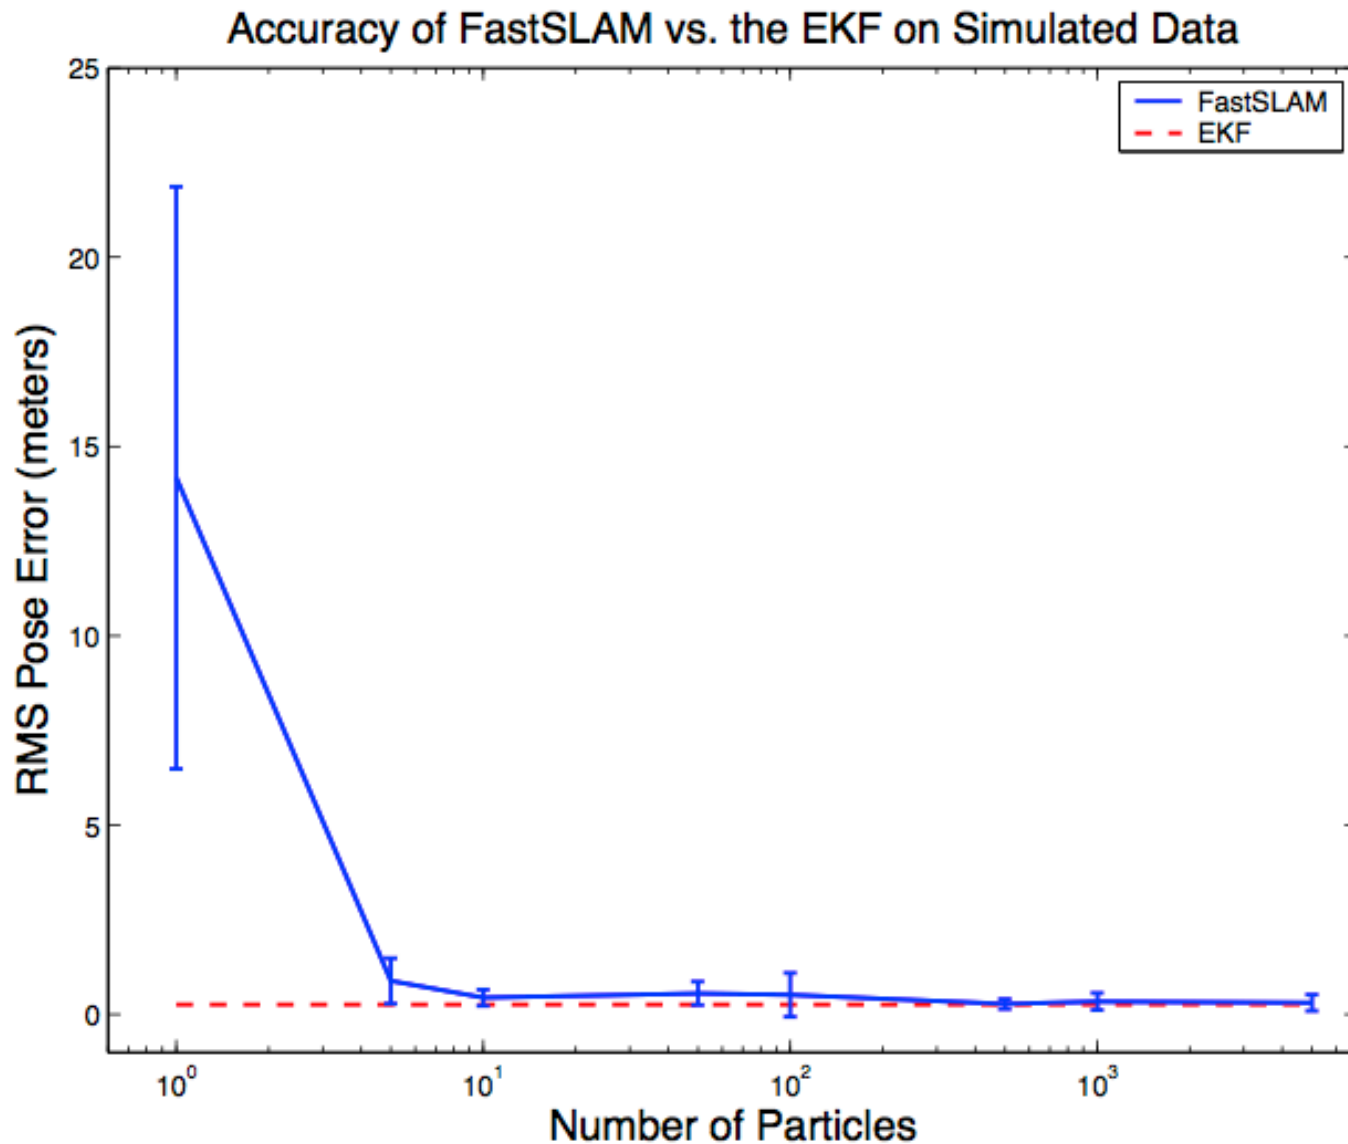
Yellow = FastSLAM



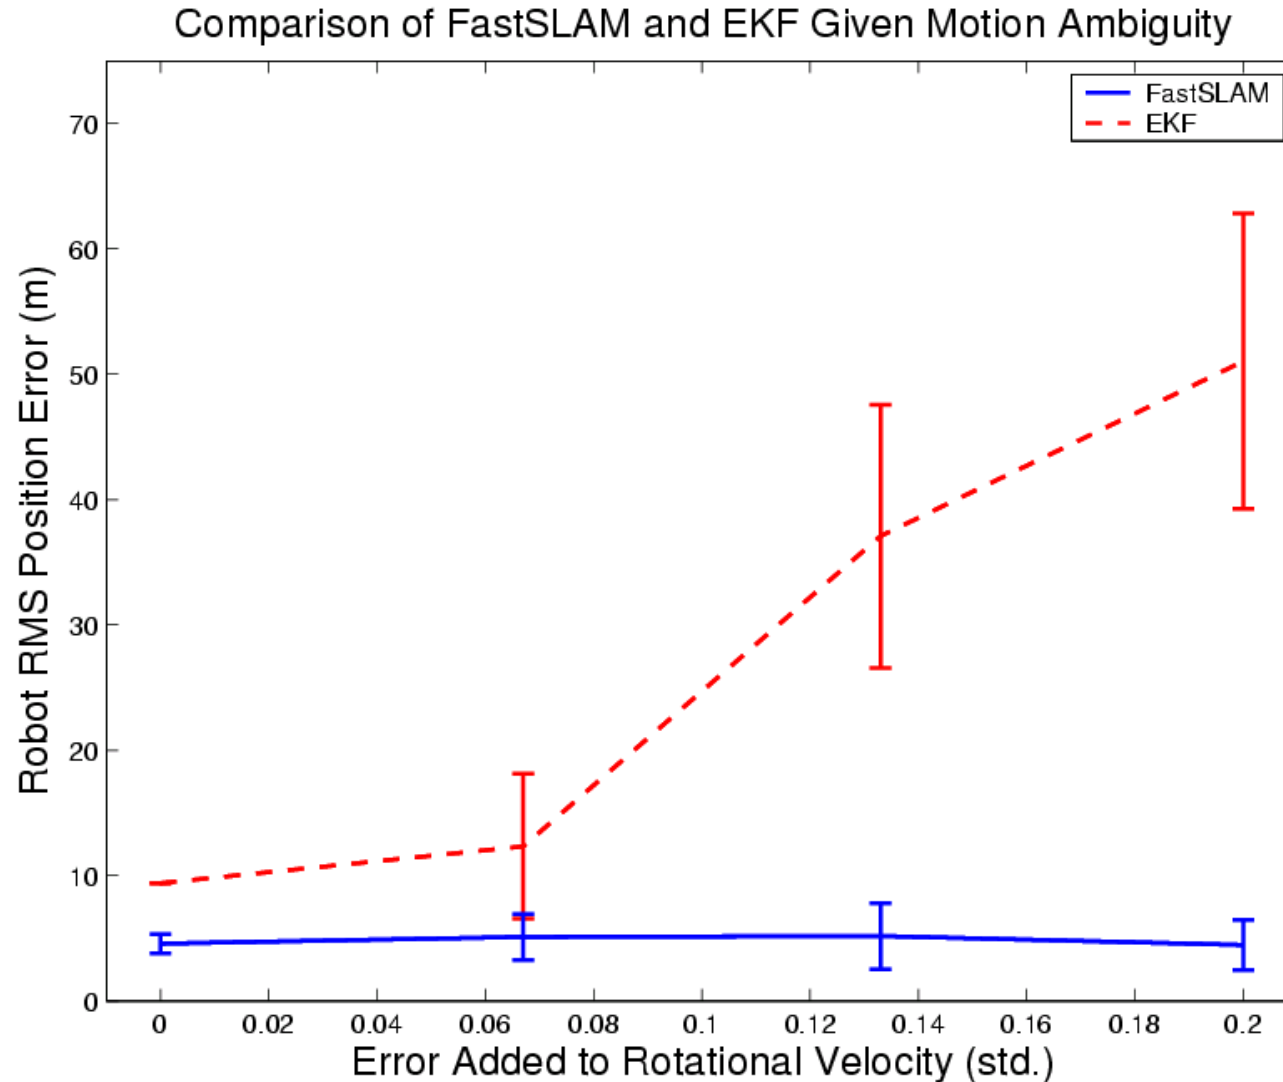
Results – Victoria Park (Video)



Results (Sample Size)

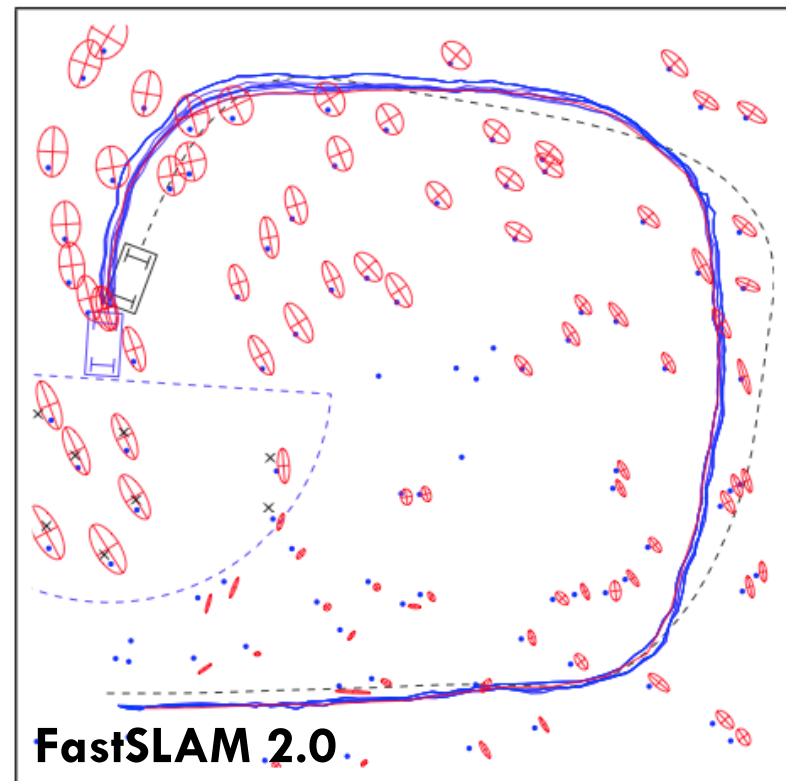
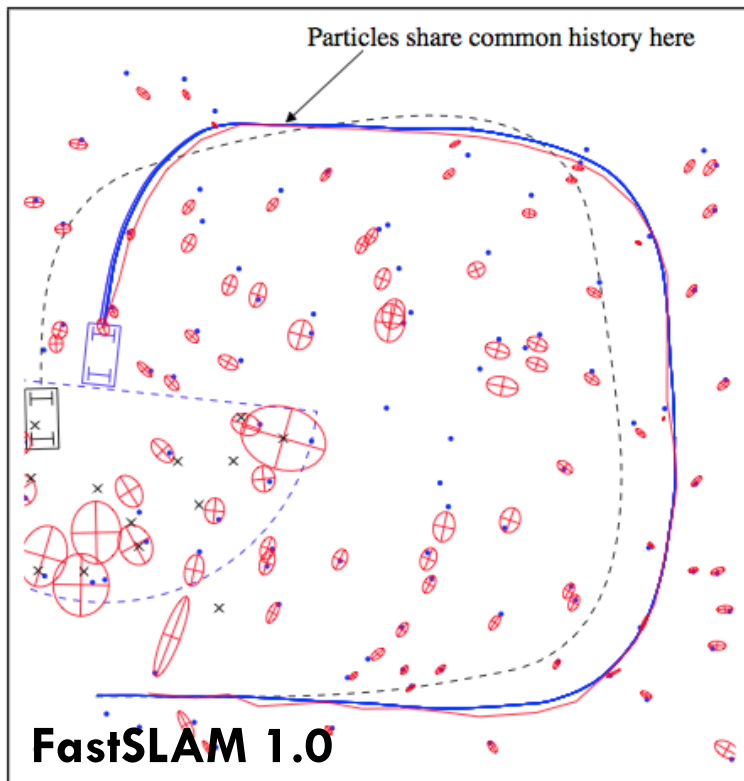


Results (Motion Uncertainty)

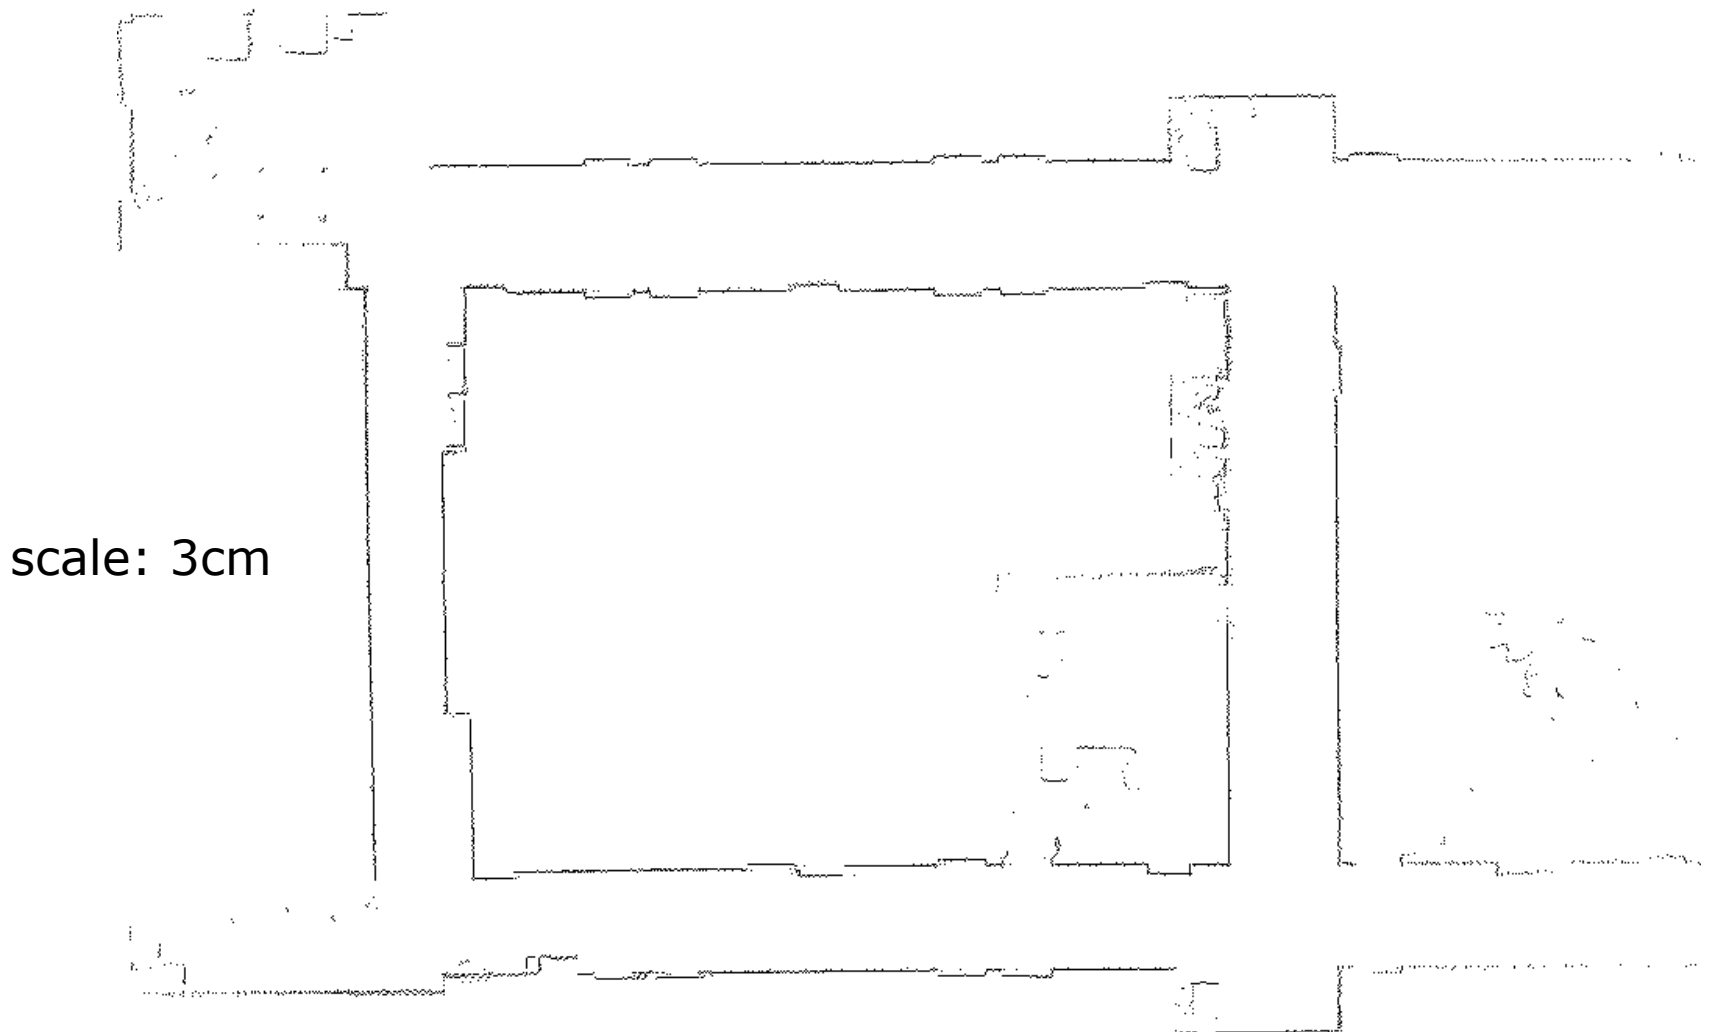


FastSLAM Problems

- How to determine the sample size?
- Particle deprivation, especially when closing (multiple) loops

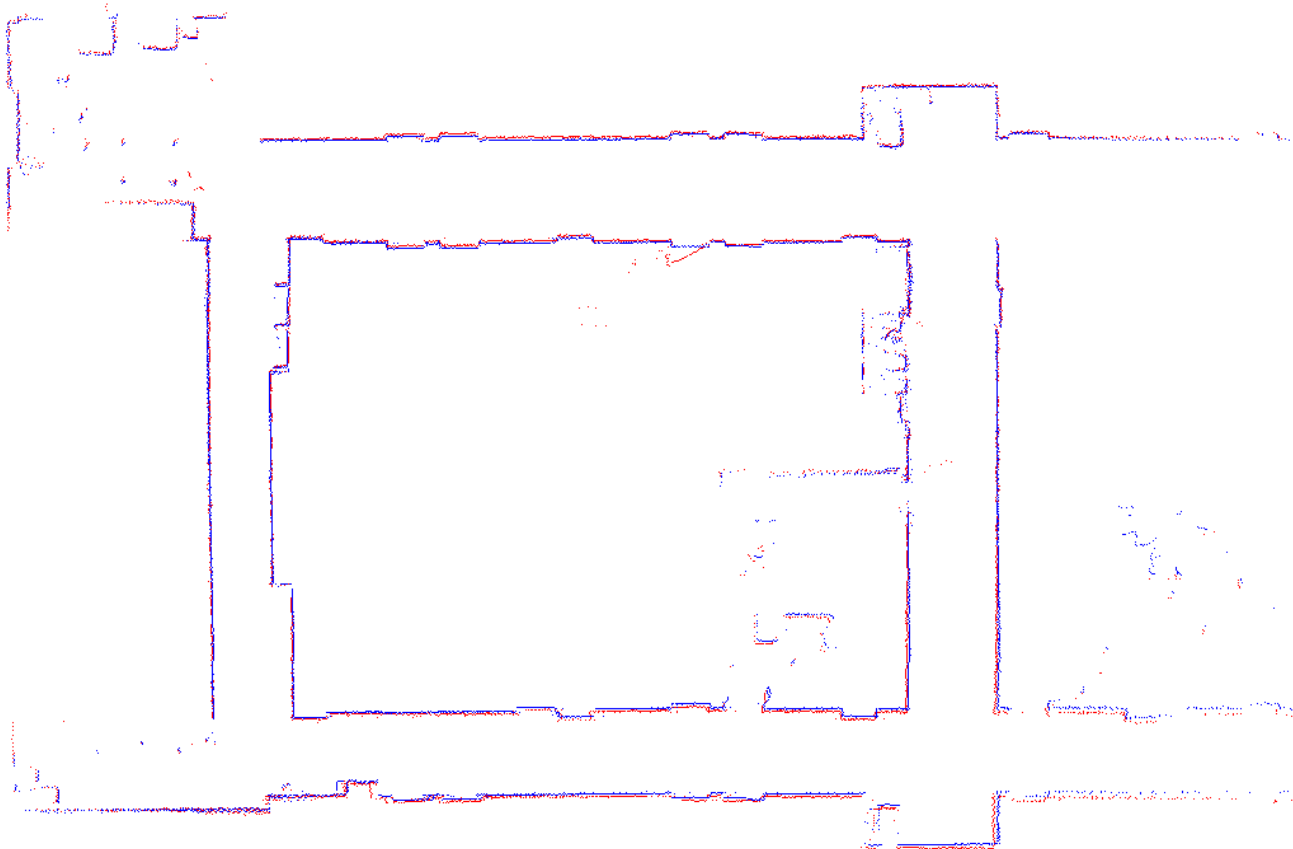


DP-SLAM: High-Res Fast-SLAM via History Sharing

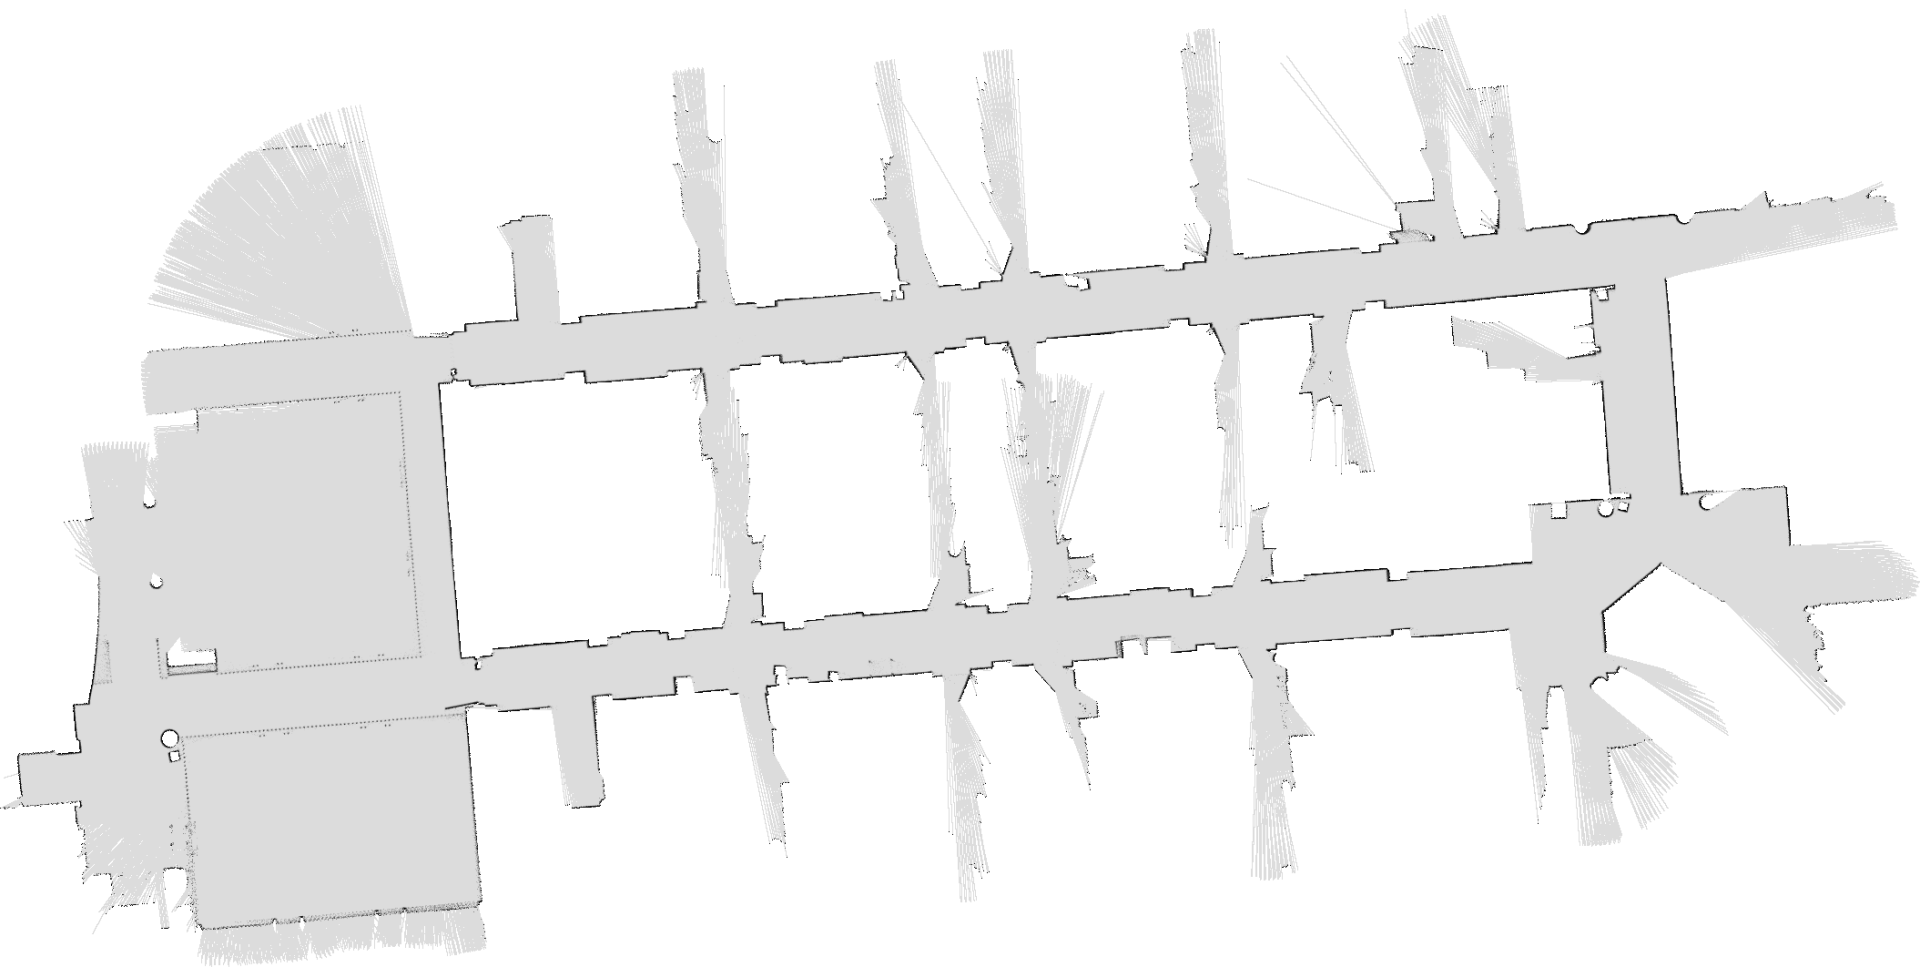


Run at real-time speed on 2.4GHz Pentium 4 at 10cm/s

Consistency



Results obtained with DP-SLAM 2.0 (offline)



Close up



End courtesy of Eliazar & Parr

FastSLAM Summary

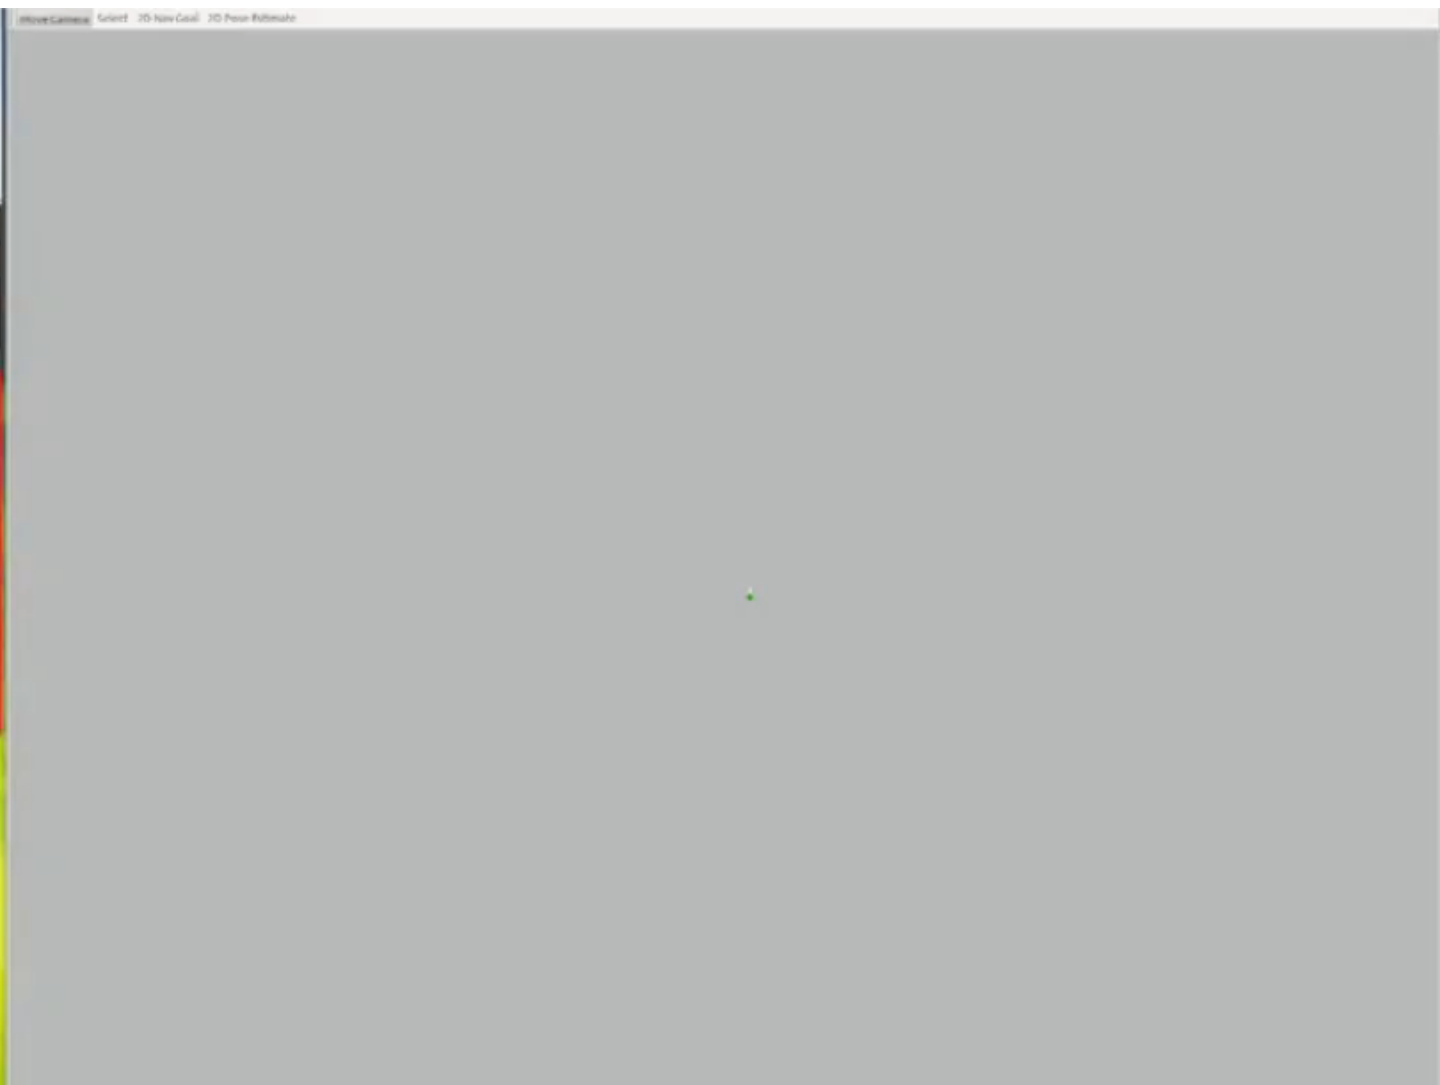
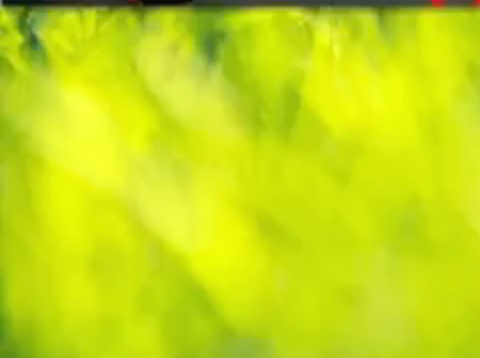
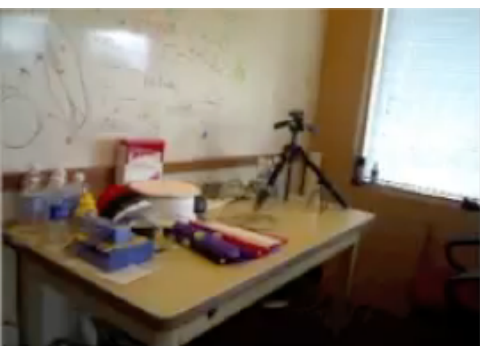
- Particle filter-based SLAM
- Rao-Blackwellization: model the robot's path by sampling and compute the landmarks given the poses
- Allow for per-particle data association
- FastSLAM 1.0 and 2.0 differ in the proposal distribution
- Complexity $\mathcal{O}(N \log M)$

Literature

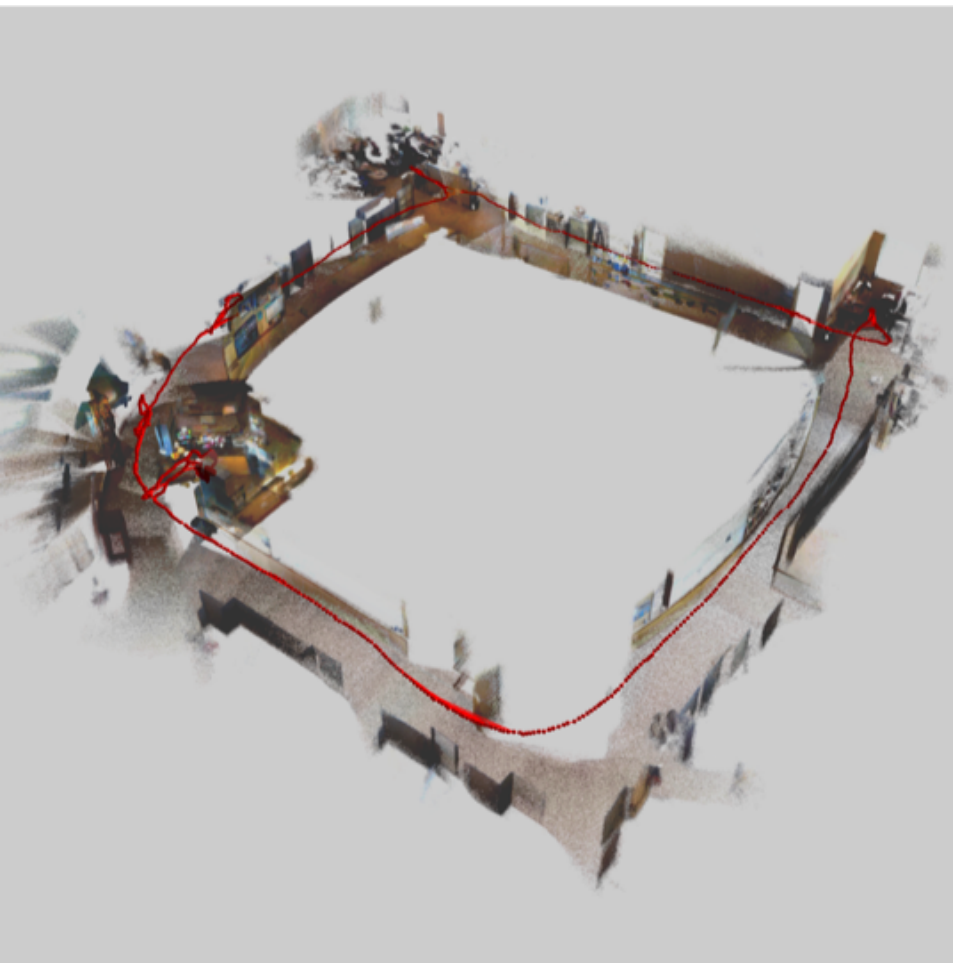
FastSLAM

- Thrun et al.: “Probabilistic Robotics”, Chapter 13.1 - 13.3 + 13.8 (see errata!)
- Montemerlo, Thrun, Kollar, Wegbreit: FastSLAM: A Factored Solution to the Simultaneous Localization and Mapping Problem, 2002
- Montemerlo and Thrun: Simultaneous Localization and Mapping with Unknown Data Association Using FastSLAM, 2003

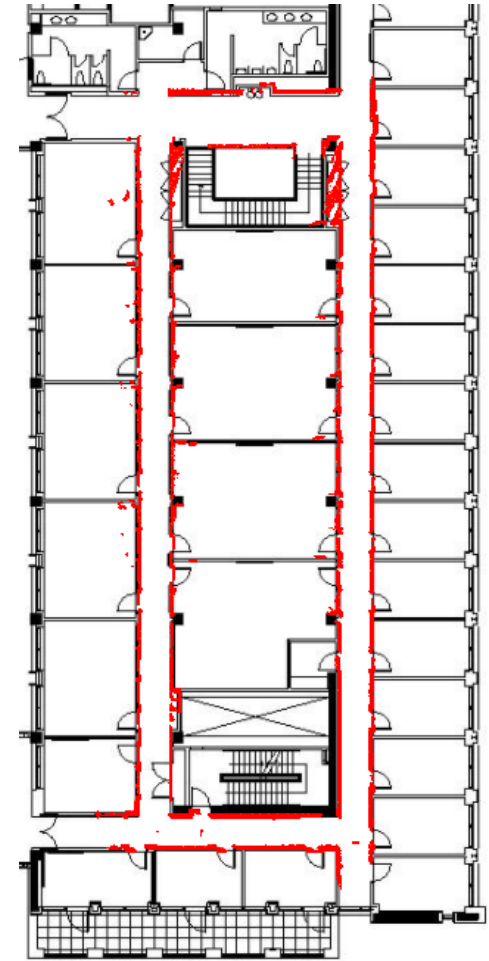
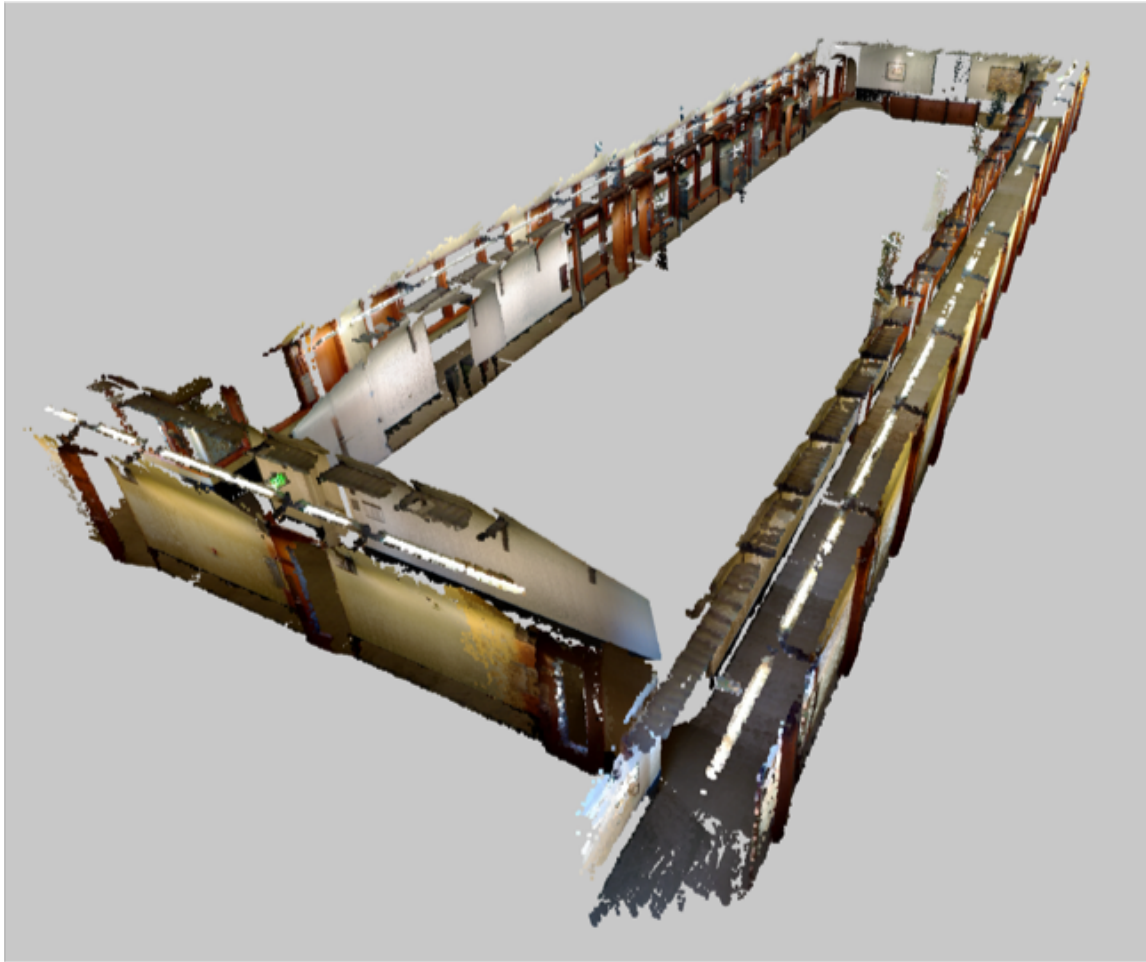
RGBD SLAM



Resulting Map



Experiments: Overlay 1



Experiments: Overlay 2

