Introduction to Mobile Robotics

Information Driven Exploration

Wolfram Burgard, Michael Ruhnke,

Bastian Steder

Tasks of Mobile Robots

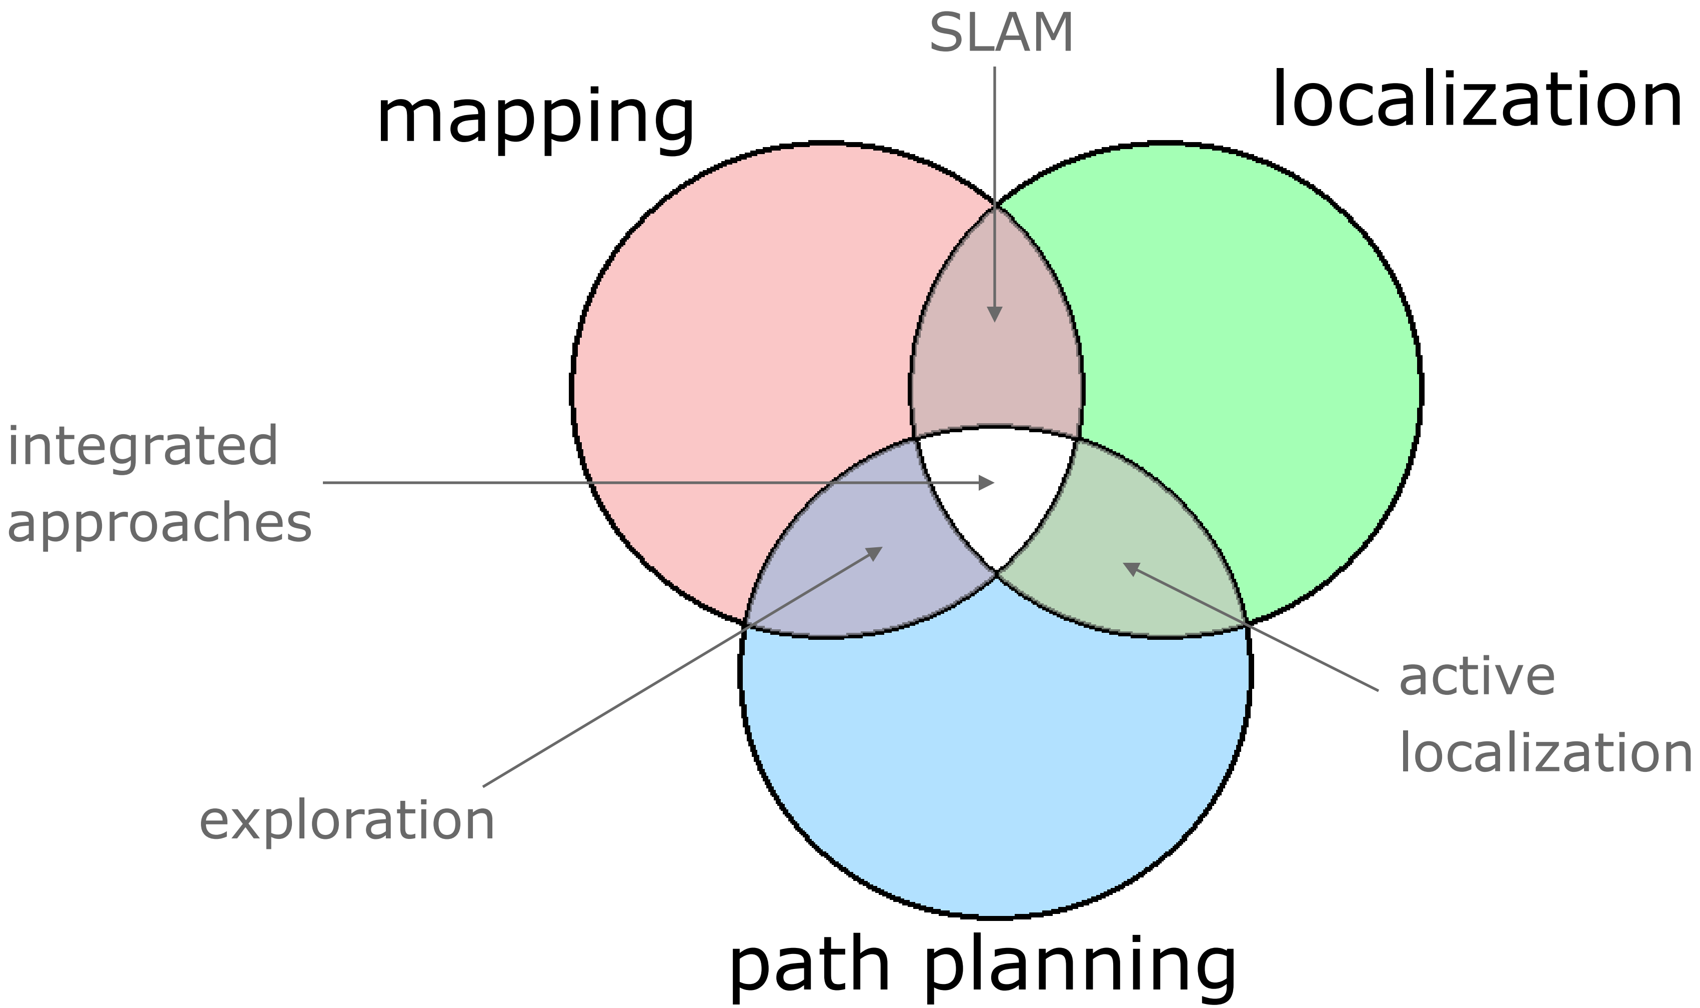

Exploration and SLAM

- SLAM is typically **passive**, because it consumes incoming sensor data
- Exploration **actively guides the robot** to cover the environment with its sensors
- Exploration in combination with SLAM: **Acting under pose and map uncertainty**
- Uncertainty should/needs to be taken into account when selecting an action

Mapping with Rao-Blackwellized Particle Filter (Brief Summary)

- Each particle represents a possible trajectory of the robot
- Each particle
  - maintains its own map and
  - updates it upon "mapping with known poses"
- Each particle survives with a probability proportional to the likelihood of the observations relative to its own map

Factorization Underlying Rao-Blackwellized Mapping

poses map observations & odometry

$$p(x, m \mid z, u)$$

$$= \quad p(m \mid x, z, \cancel{u})\, p(x \mid z, u)$$

Mapping with known poses

Particle filter representing trajectory hypotheses

Example: Particle Filter for Mapping

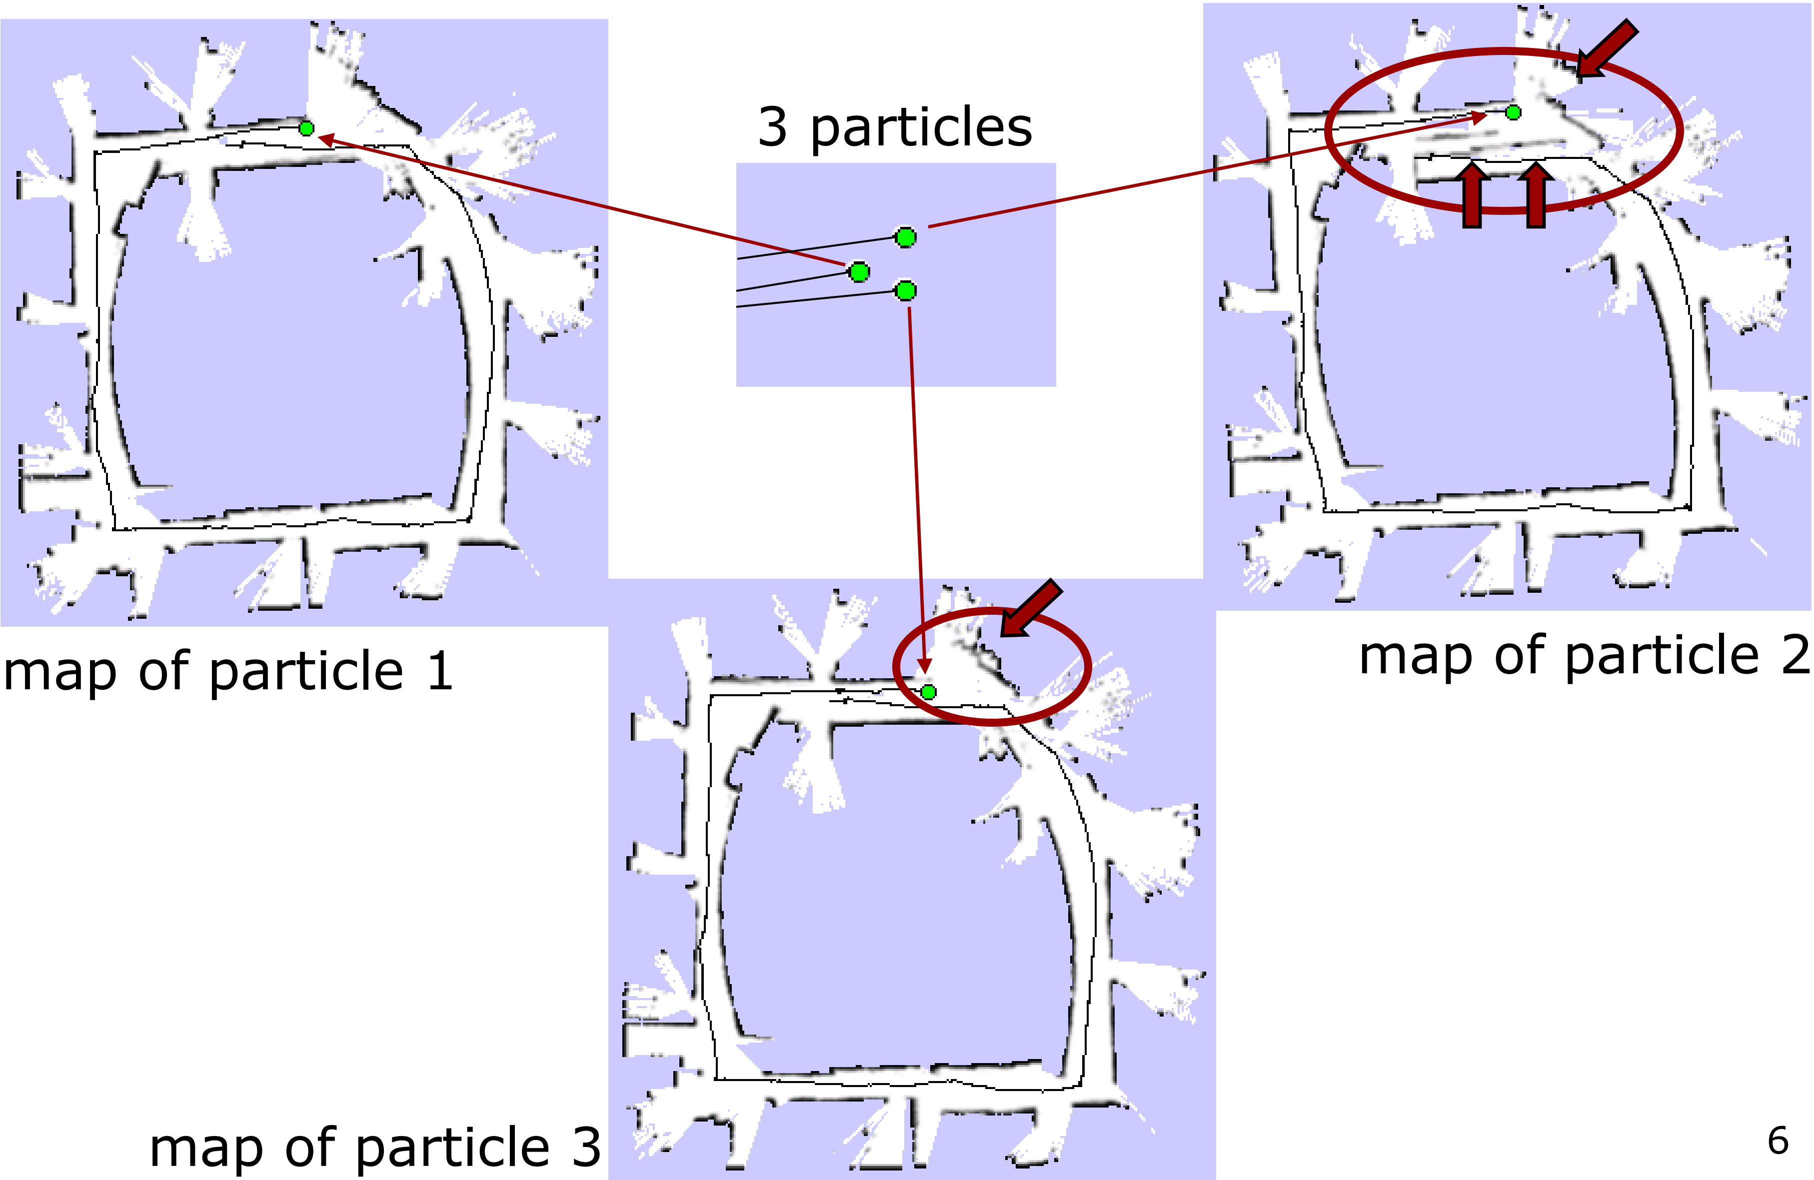

Combining Exploration and SLAM

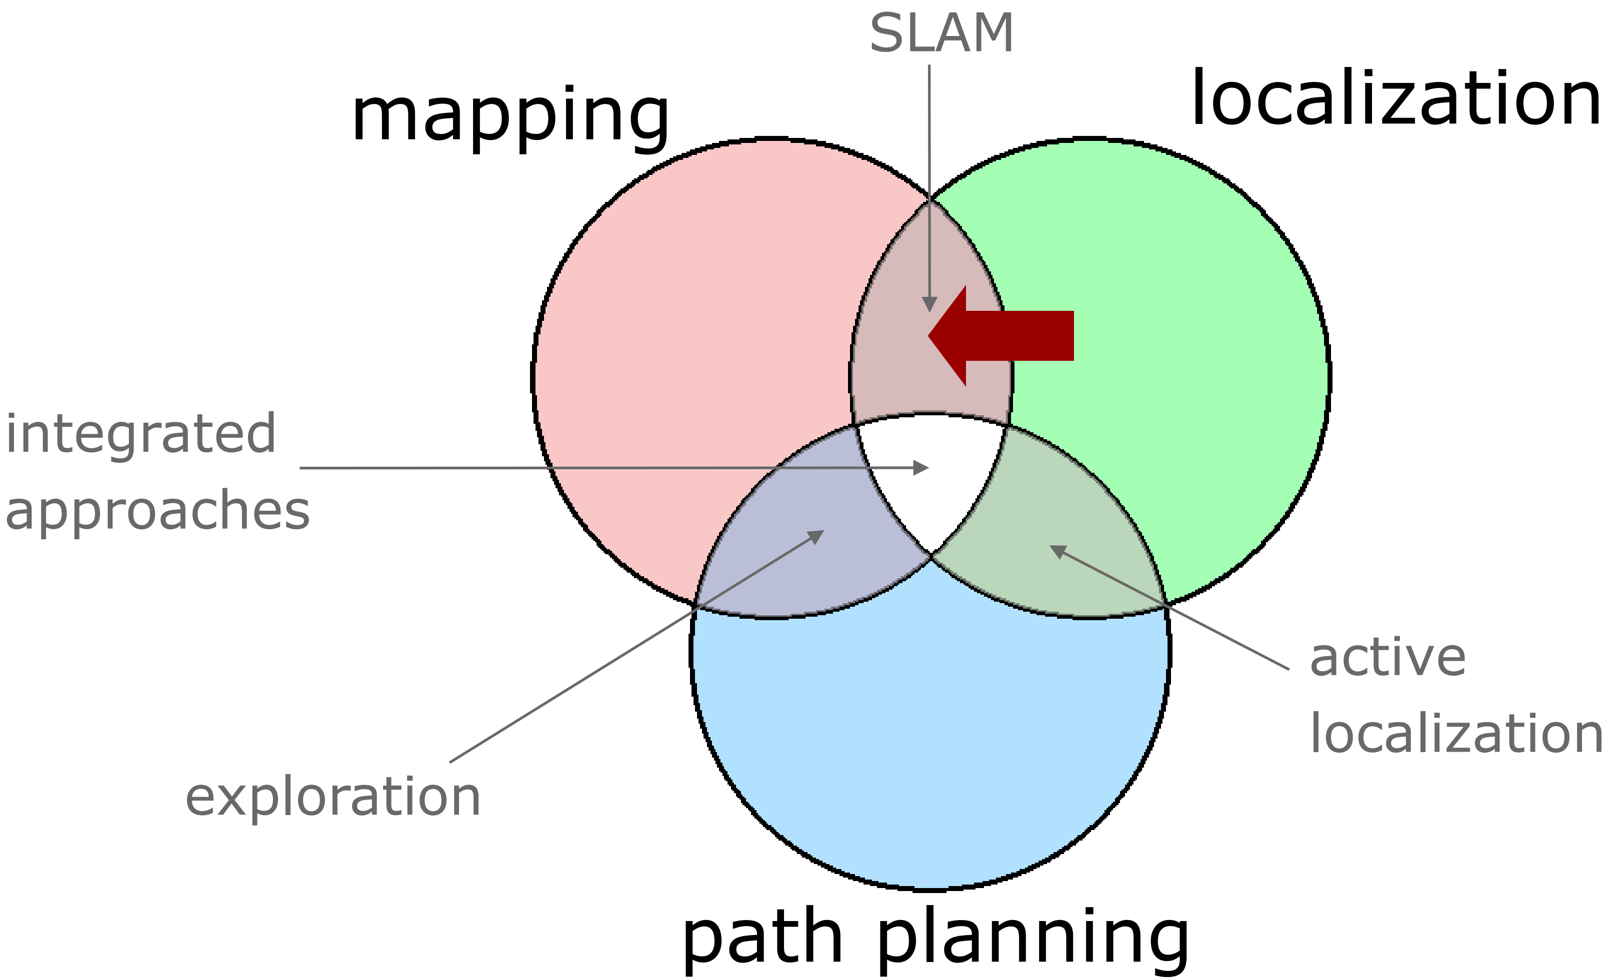

Exploration

- SLAM approaches seen so far are purely passive
- By reasoning about control, the mapping process can be made much more effective
- Question: **Where to move next?**

Where to Move Next?

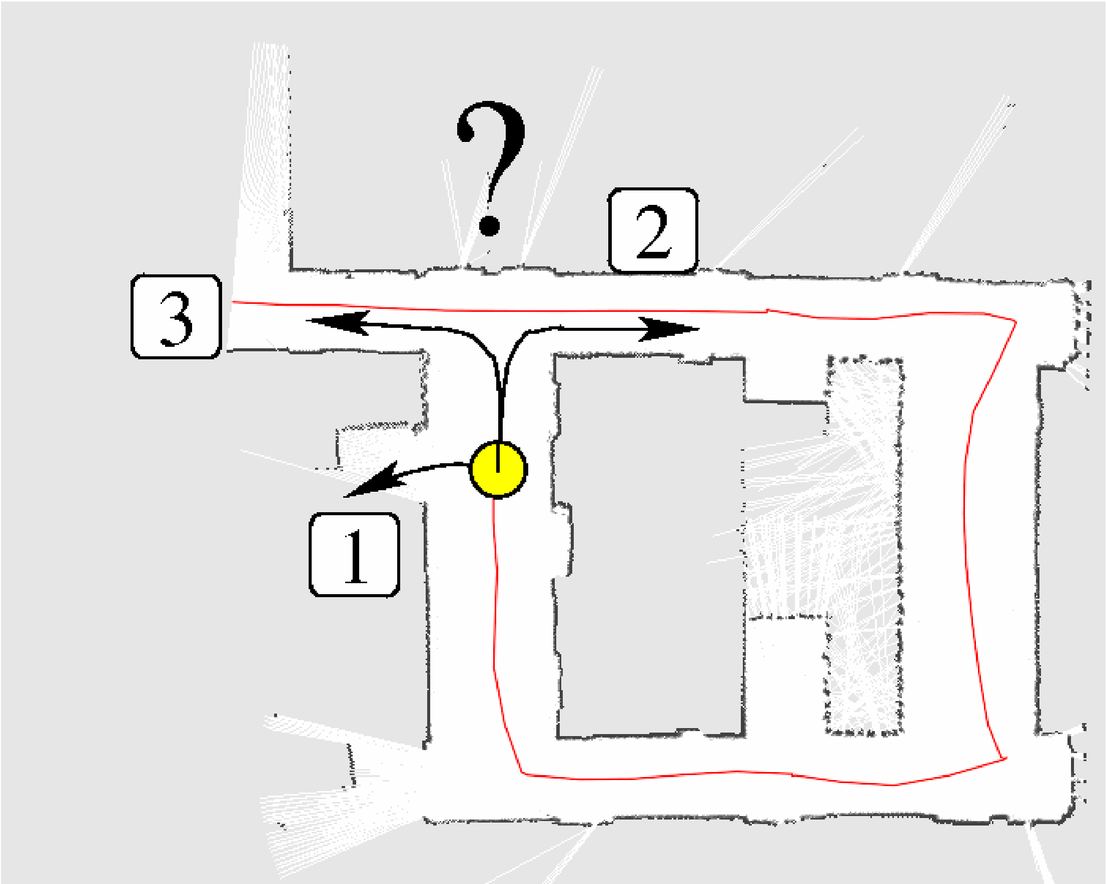

Decision-Theoretic Approach

- Learn the map using a Rao-Blackwellized particle filter
- Consider a set of potential actions
- Apply an exploration approach that minimizes the overall uncertainty

Utility = uncertainty reduction - cost

Example

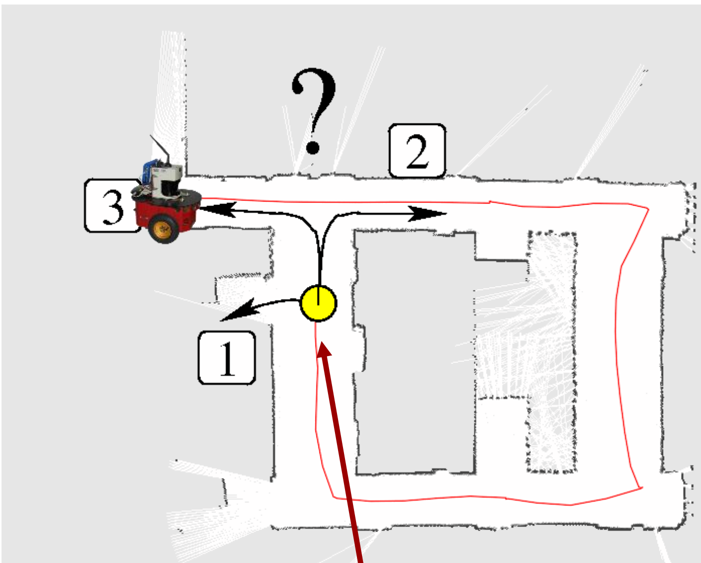

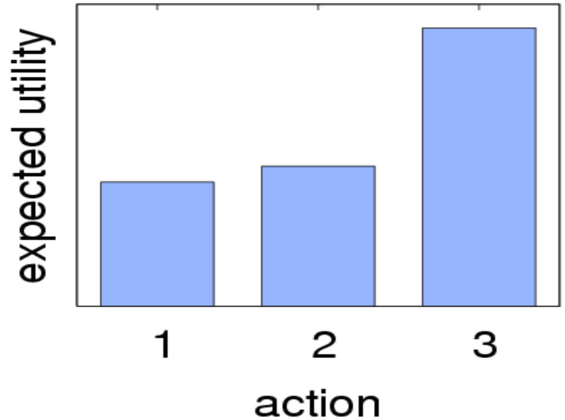

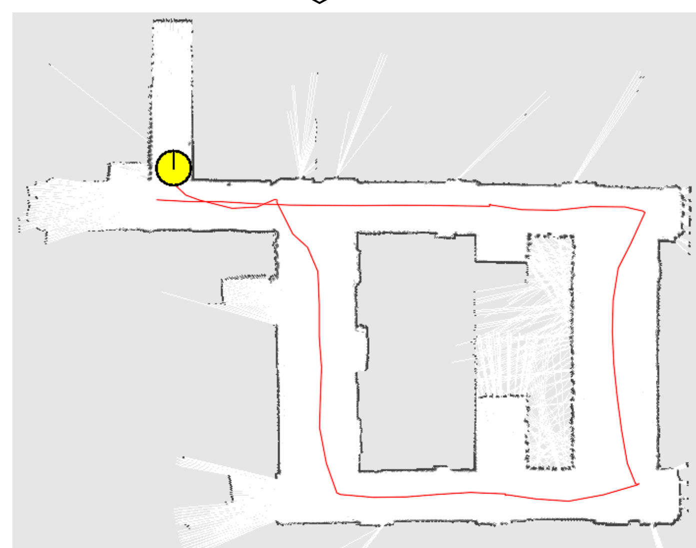

high pose uncertainty

The Uncertainty of a Posterior

- **Entropy** is a general measure for the uncertainty of a posterior

$$\begin{aligned} H(X) &= -\int_x p(X=x)\log p(X=x)\,dx \\ &= E_X[-\log(p(X))] \end{aligned}$$

- **Conditional Entropy**

$$H(X \mid Y) = \int_y p(Y=y)H(X \mid Y=y)\,dy$$

Mutual Information

- **Expected Information Gain** or **Mutual Information** = Expected Uncertainty Reduction

$$I(X;Y) = H(X) - H(X \mid Y)$$

$$I(X;Y) = H(Y) - H(Y \mid X)$$

$$I(X;Y \mid z = c_k) = H(X \mid z = c_k) - H(X \mid Y, z = c_k)$$

$$I(X;Y \mid Z) = H(X \mid Z) - H(X \mid Y, Z)$$

Entropy Computation

$$
\begin{aligned}
&H(X,Y)\\
&= E_{X,Y}[-\log p(X,Y)]\\
&= E_{X,Y}[-\log(p(X)\ p(Y \mid X))]\\
&= E_{X,Y}[-\log p(X)] + E_{X,Y}[-\log p(Y \mid X)]\\
&= H(X) + \int_{x,y} -p(x,y) \log p(y \mid x)\ dx\ dx\\
&= H(X) + \int_{x,y} -p(y \mid x) p(x) \log p(y \mid x)\ dx\ dy\\
&= H(X) + \int_{x} p(x) \int_{y} -p(y \mid x) \log p(y \mid y)\ dy\ dx\\
&= H(X) + \int_{x} p(x) H(Y \mid X = x)\ dx\\
&= H(X) + H(Y \mid X)
\end{aligned}
$$

The Uncertainty of the Robot

- The uncertainty of the RBPF:

$$H(X,M) \;=\; H(X) + \int_x p(x) H(M \mid X = x)\, dx$$

$$\Downarrow$$

$$H(X,M) \;=\; \underbrace{H(X)}_{\text{trajectory uncertainty}} + \sum_{i=1}^{\#particles} \underbrace{\omega^{[i]}}_{\text{particle weights}} \underbrace{H(M^{[i]}}_{\text{map uncertainty}} \mid X^{[i]} = x^{[i]})$$

Computing the Entropy of the Map Posterior

Occupancy Grid map *m*:

$$H(M) \;=\; -\sum_{c\in M} p(c)\log p(c) + (1-p(c))\log(1-p(c))$$

map uncertainty

grid cells

probability that the cell is occupied

Map Entropy

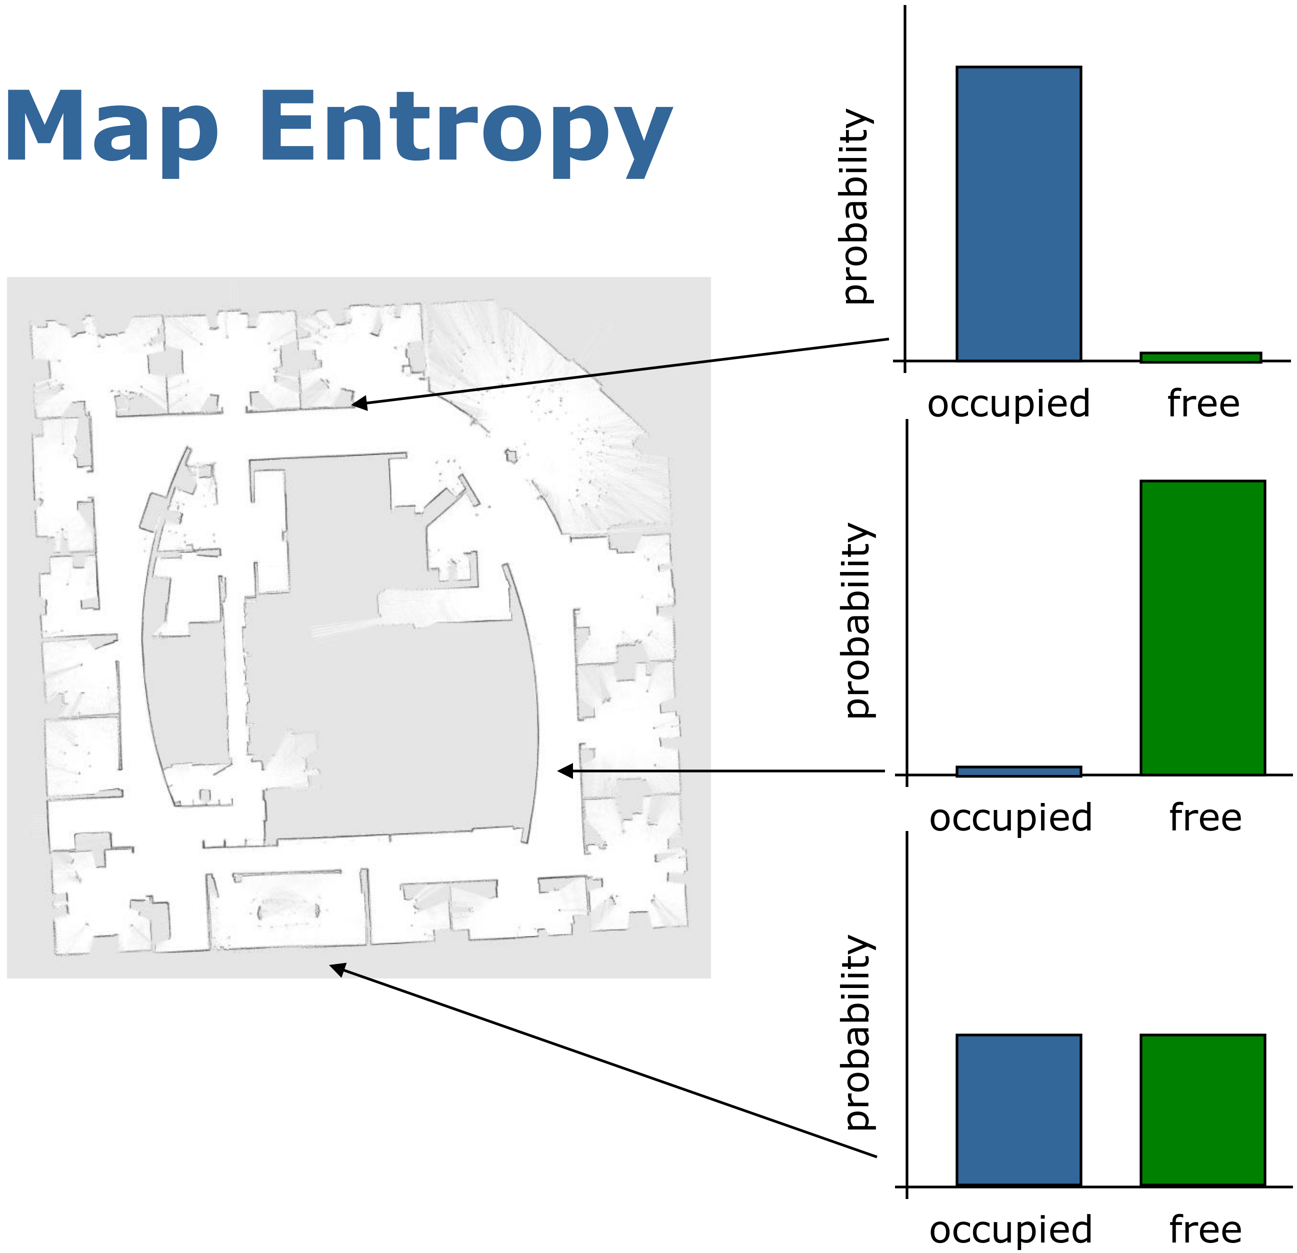

Low entropy

Low entropy

High entropy

The overall entropy is the sum of the individual entropy values

Computing the Entropy of the Trajectory Posterior

1. High-dimensional Gaussian

$$H(\mathscr{G}(\mu, \Sigma)) \quad = \quad \log((2\pi e)^{(n/2)} |\Sigma|)$$

   reduced rank for sparse particle sets

2. Grid-based approximation

$$H(X) \quad \leadsto \quad const.$$

   for sparse particle clouds

Approximation of the Trajectory Posterior Entropy

Average pose entropy over time:

$$H(X_{1:t} \mid d) \approx \frac{1}{t} \sum_{t'=1}^{t} H(X_{t'} \mid d)$$

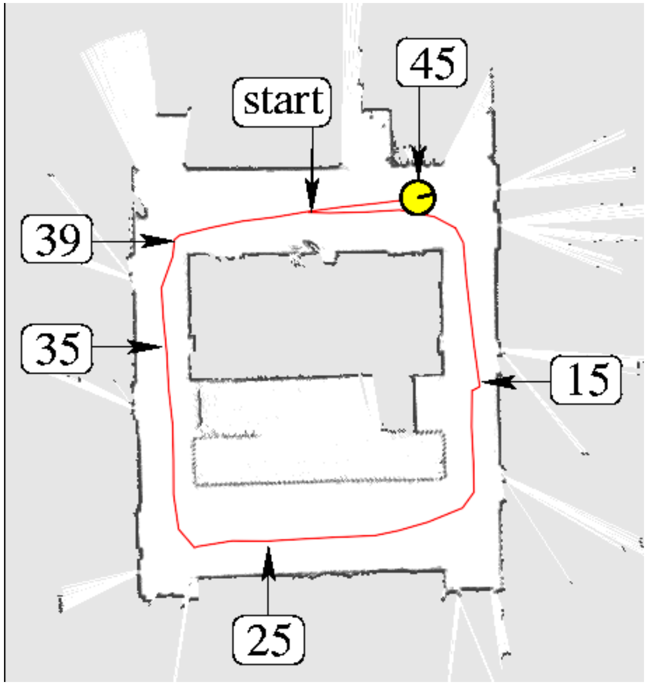

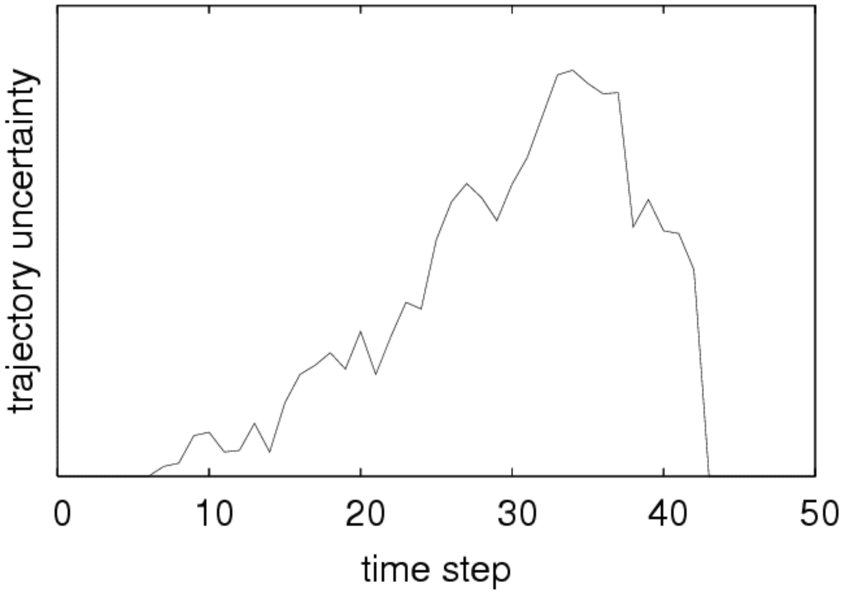

Mutual Information

- The mutual information I is given by the expected reduction of entropy in the belief

action to be carried out

$I(X,M;Z^a) =$ "uncertainty of the filter" –
"uncertainty of the filter after carrying out action *a*"

Integrating Over Observations

- Computing the mutual information requires to integrate over potential observations

$$I(X,M;Z^a) = H(X,M) - H(X,M \mid Z^a)$$

$$H(X,M \mid Z^a) = \int_z p(z \mid a) H(X,M \mid Z^a = z)\, dz$$

potential observation sequences

Approximating the Integral

- The particle filter represents a posterior about possible maps

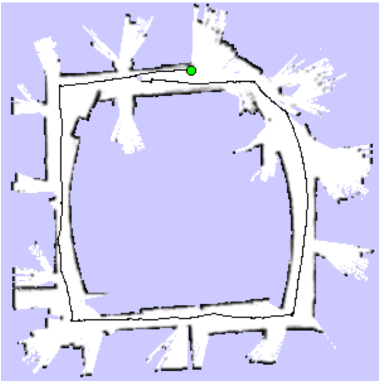

map of particle 1

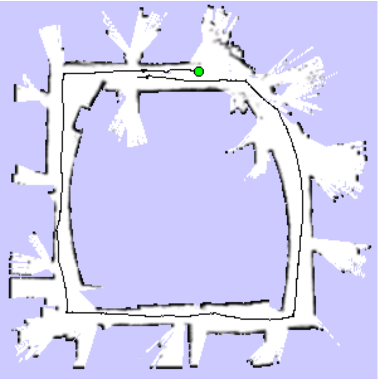

map of particle 2

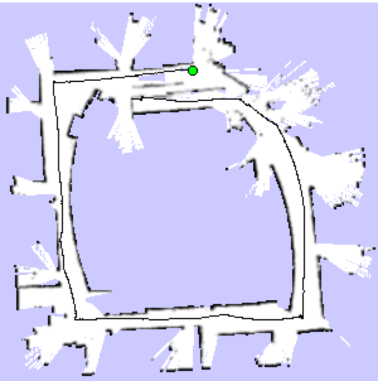

map of particle 3

...

Approximating the Integral

- The particle filter represents a posterior about possible maps
- Simulate laser measurements in the maps of the particles

$$H(X,M \mid Z^a) = \sum_z p(z \mid a) H(X,M \mid Z^a = z)$$

measurement sequences simulated in the maps

likelihood (particle weight)

$$= \sum_i \omega^{[i]} H(X,M \mid Z^a = z_{sim_a}^{[i]})$$

Simulating Observations

- Ray-casting in the map of each particle to generate observation sequences

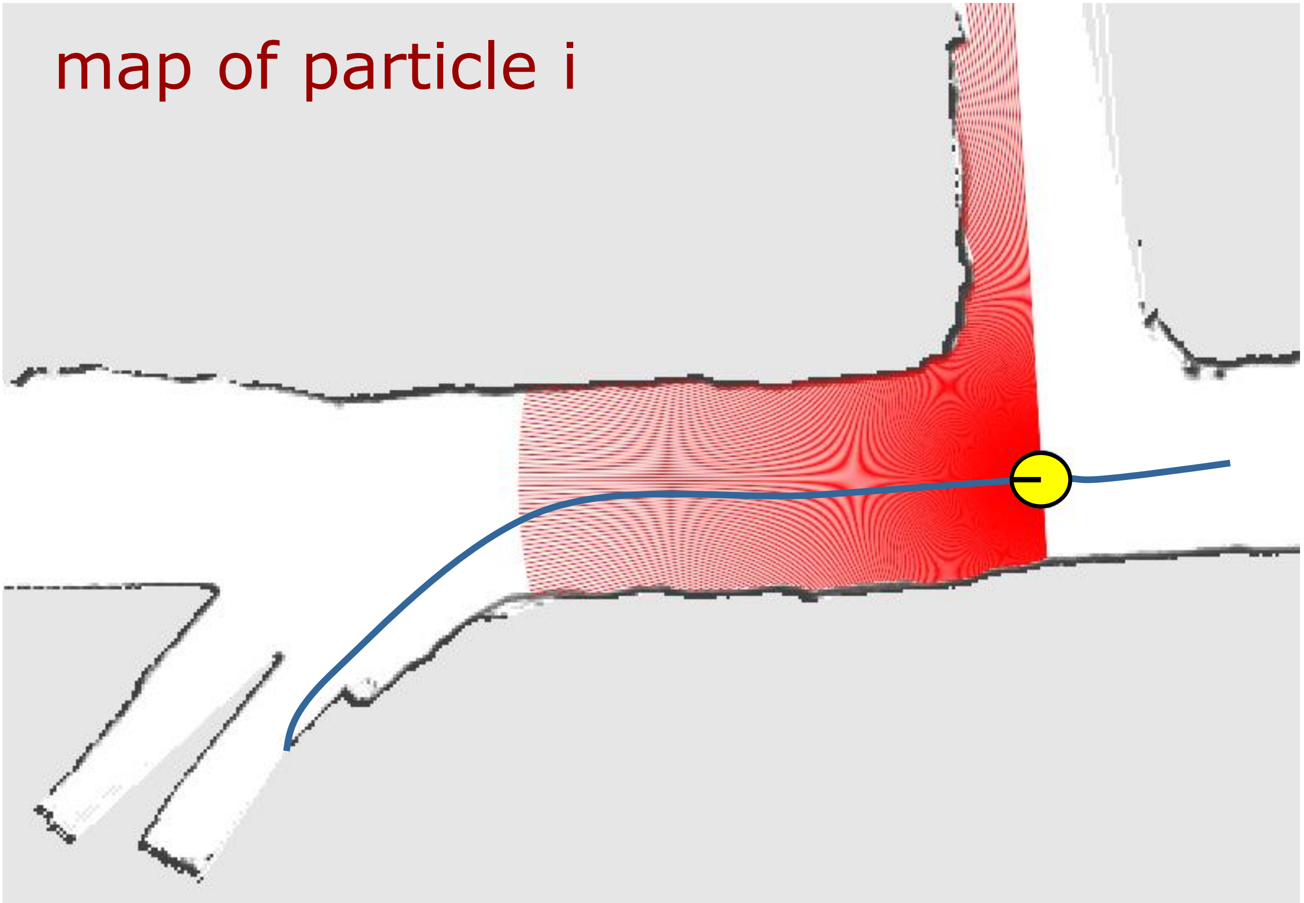

The Utility

- We take into account the cost of an action: mutual information ➡ utility U

- Select the action with the highest utility

$$a^* = \underset{a}{\operatorname{argmax}} I(X, M; Z^a) - cost(a)$$

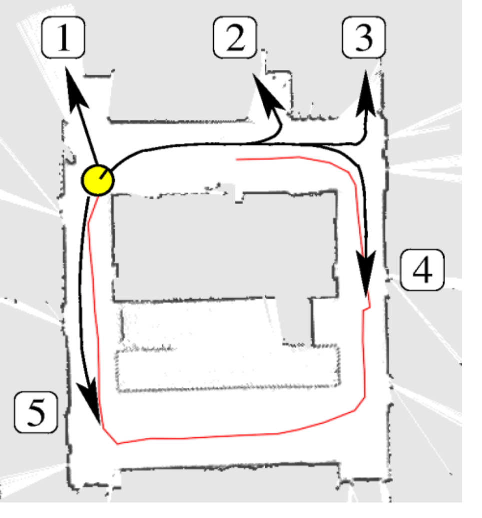

Focusing on Specific Actions

To efficiently sample actions we consider

- **exploratory actions (1-3)**
- **loop closing actions (4)** and
- **place revisiting actions (5)**

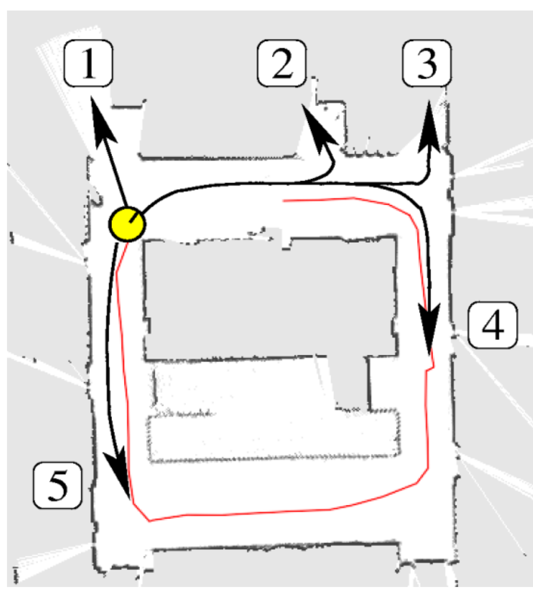

Dual Representation for Loop Detection

- **Trajectory graph** (“topological map”) stores the path traversed by the robot
- **Occupancy grid** map represents the space covered by the sensors
- **Loops** correspond to long paths in the trajectory graph and short paths in the grid map

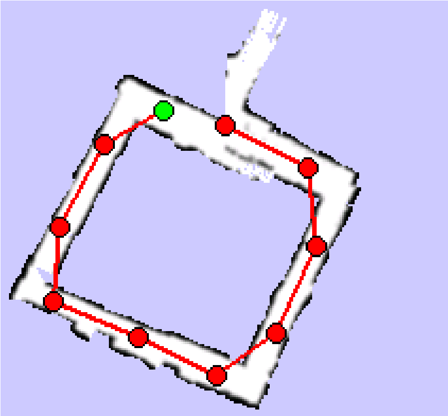

Example: Trajectory Graph

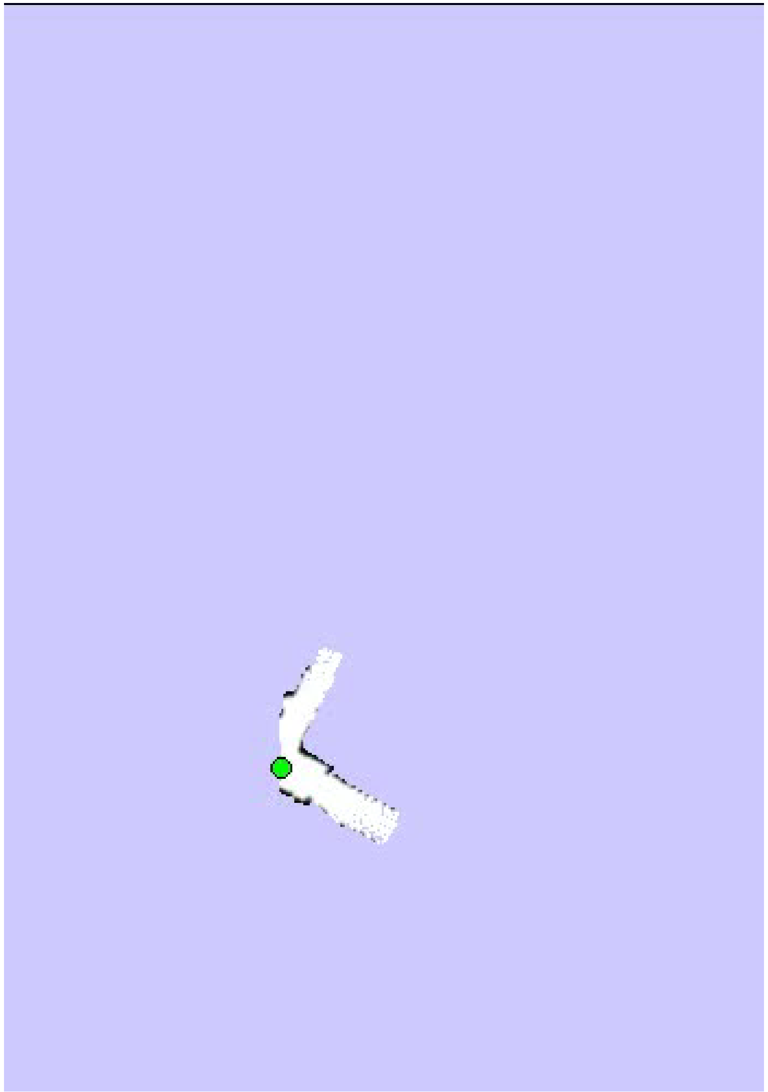

Application Example

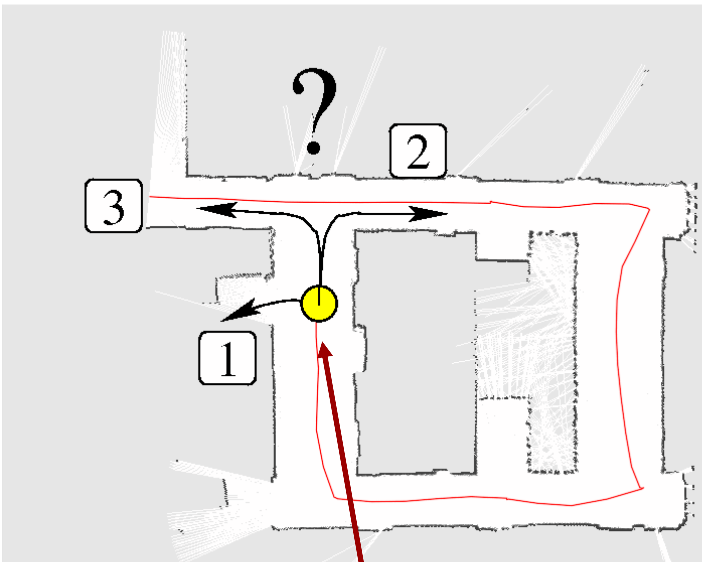

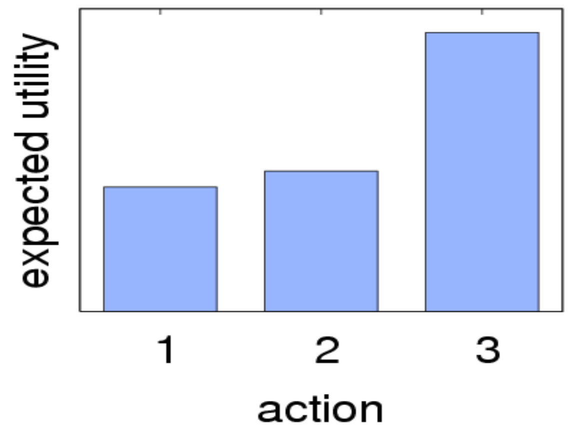

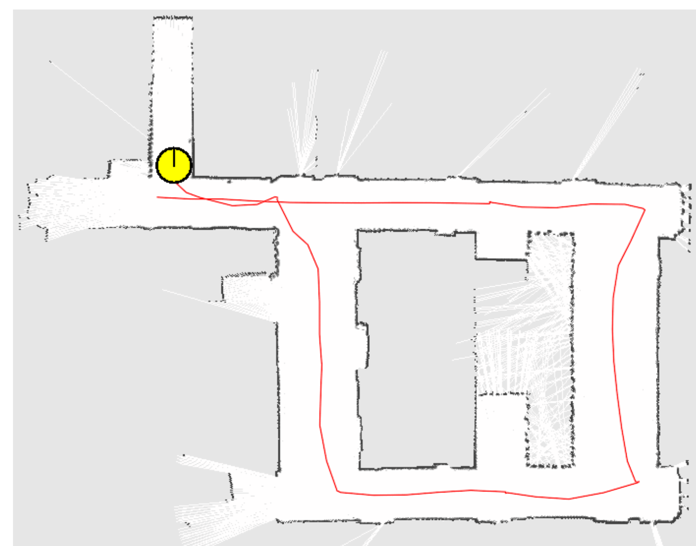

high pose uncertainty

Example: Possible Targets

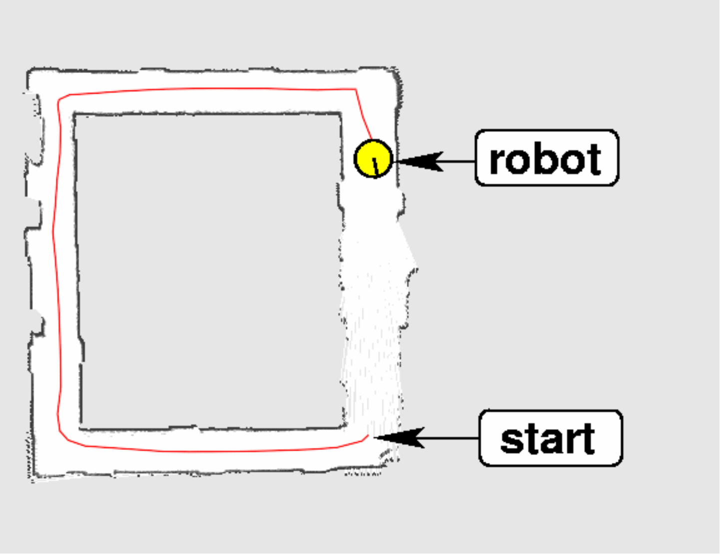

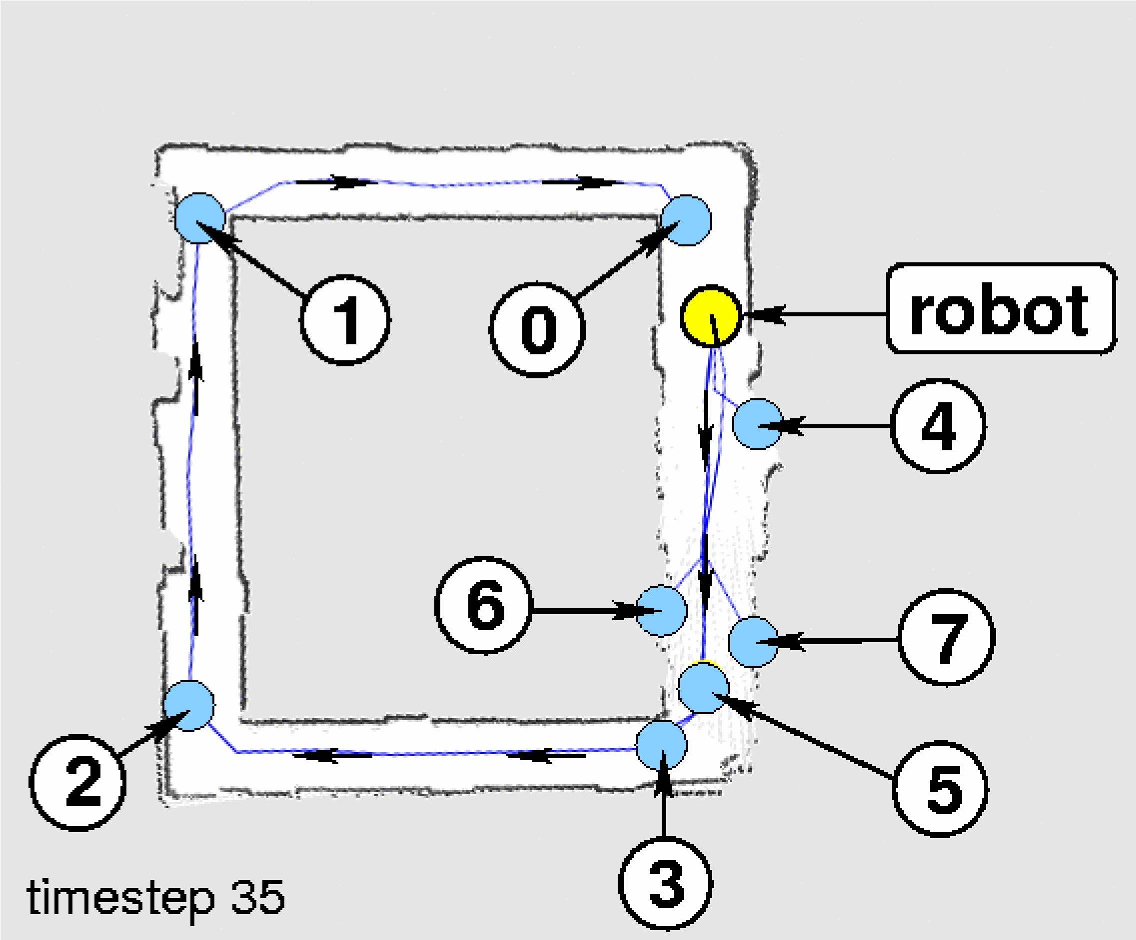

Example: Evaluate Targets

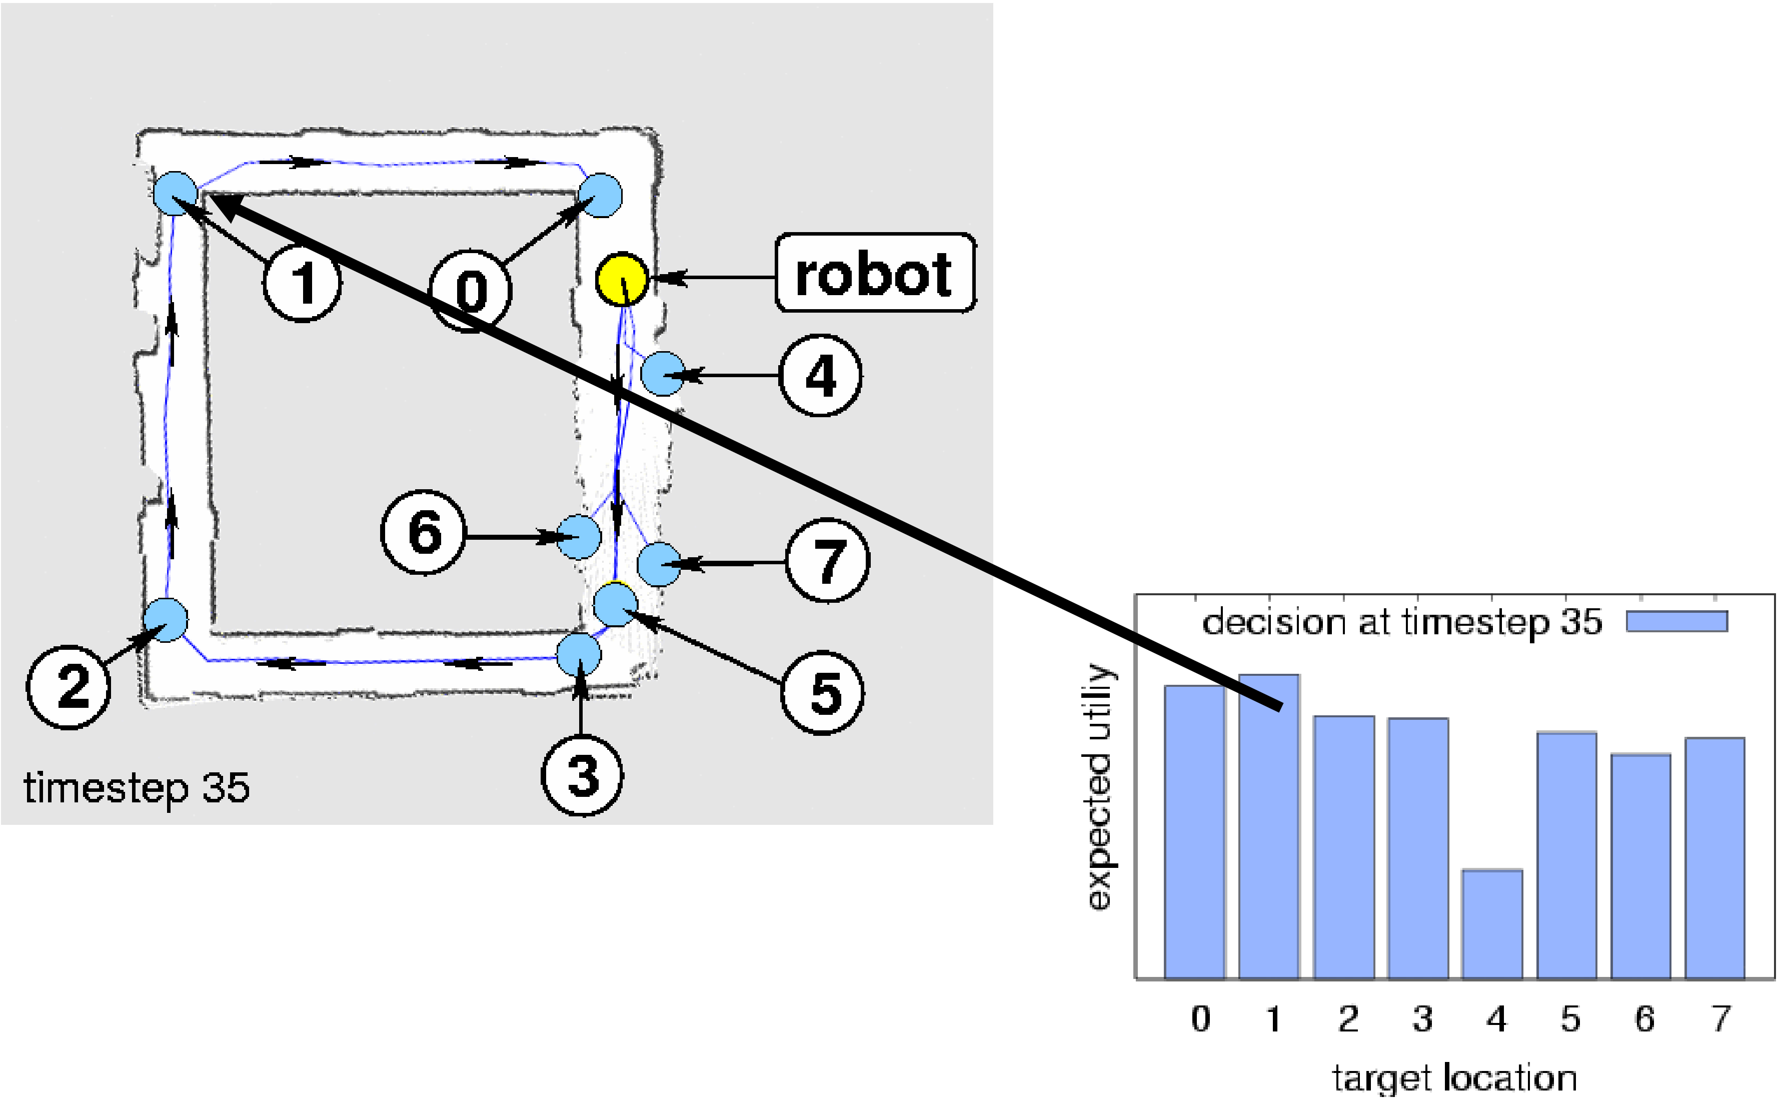

Example: Move Robot to Target

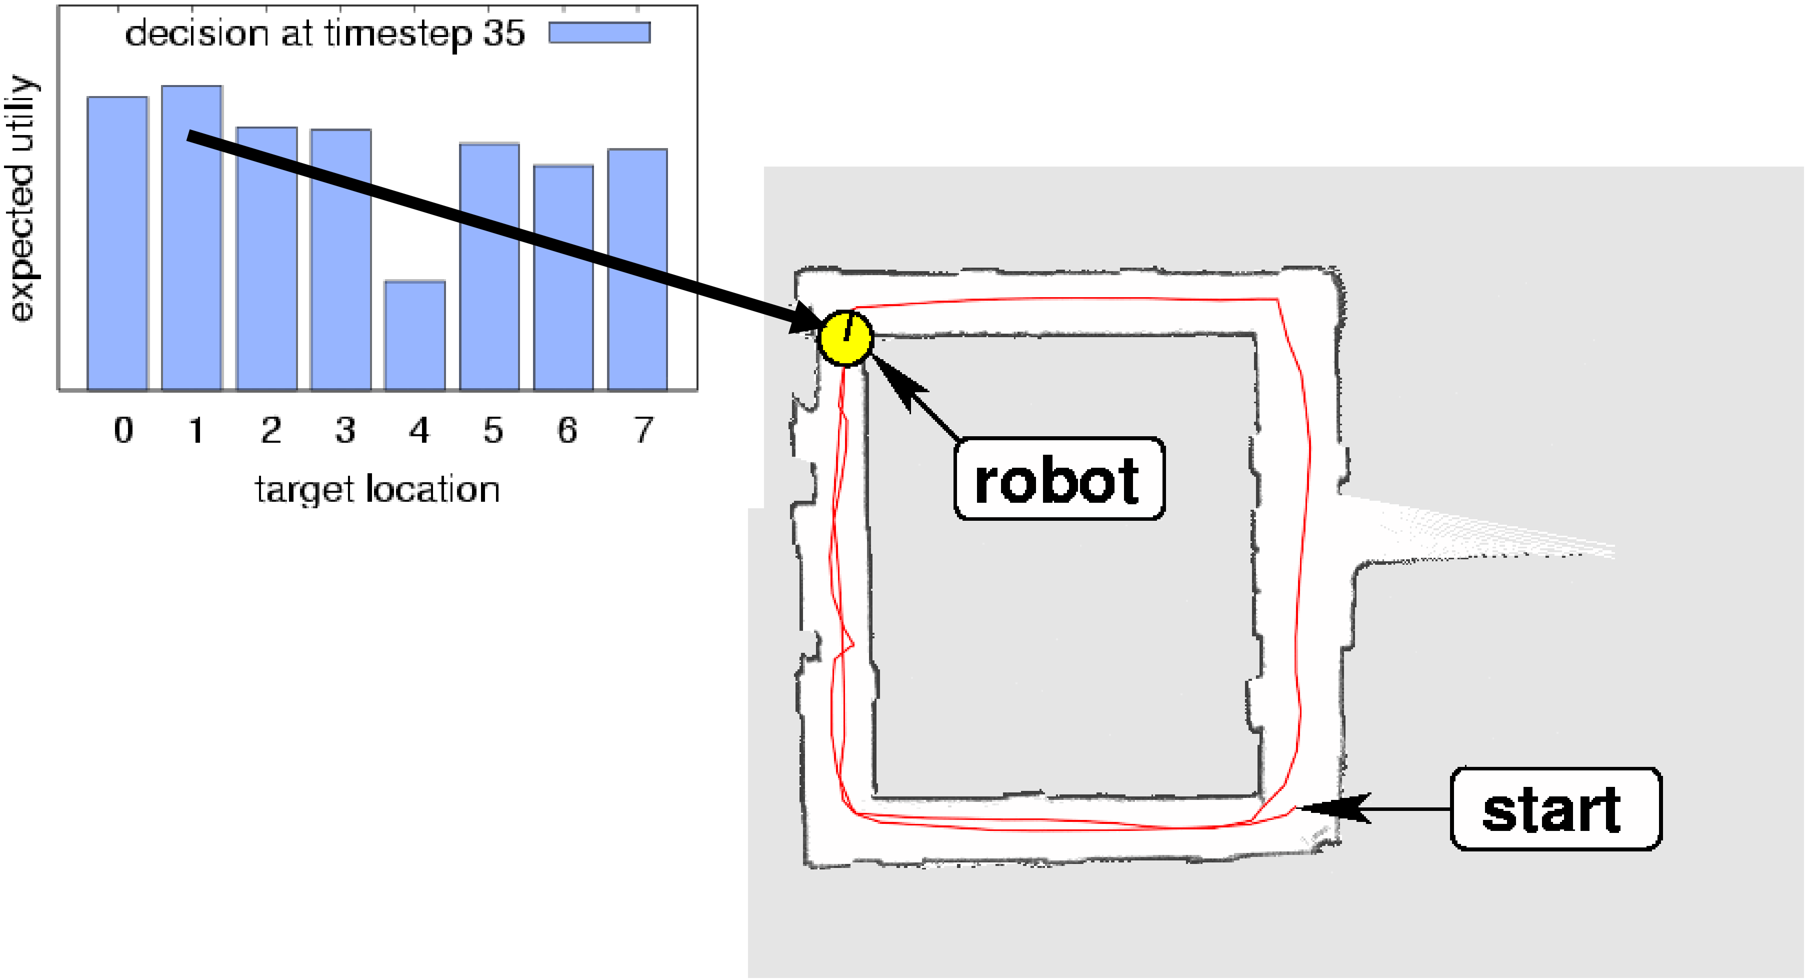

Example: Evaluate Targets

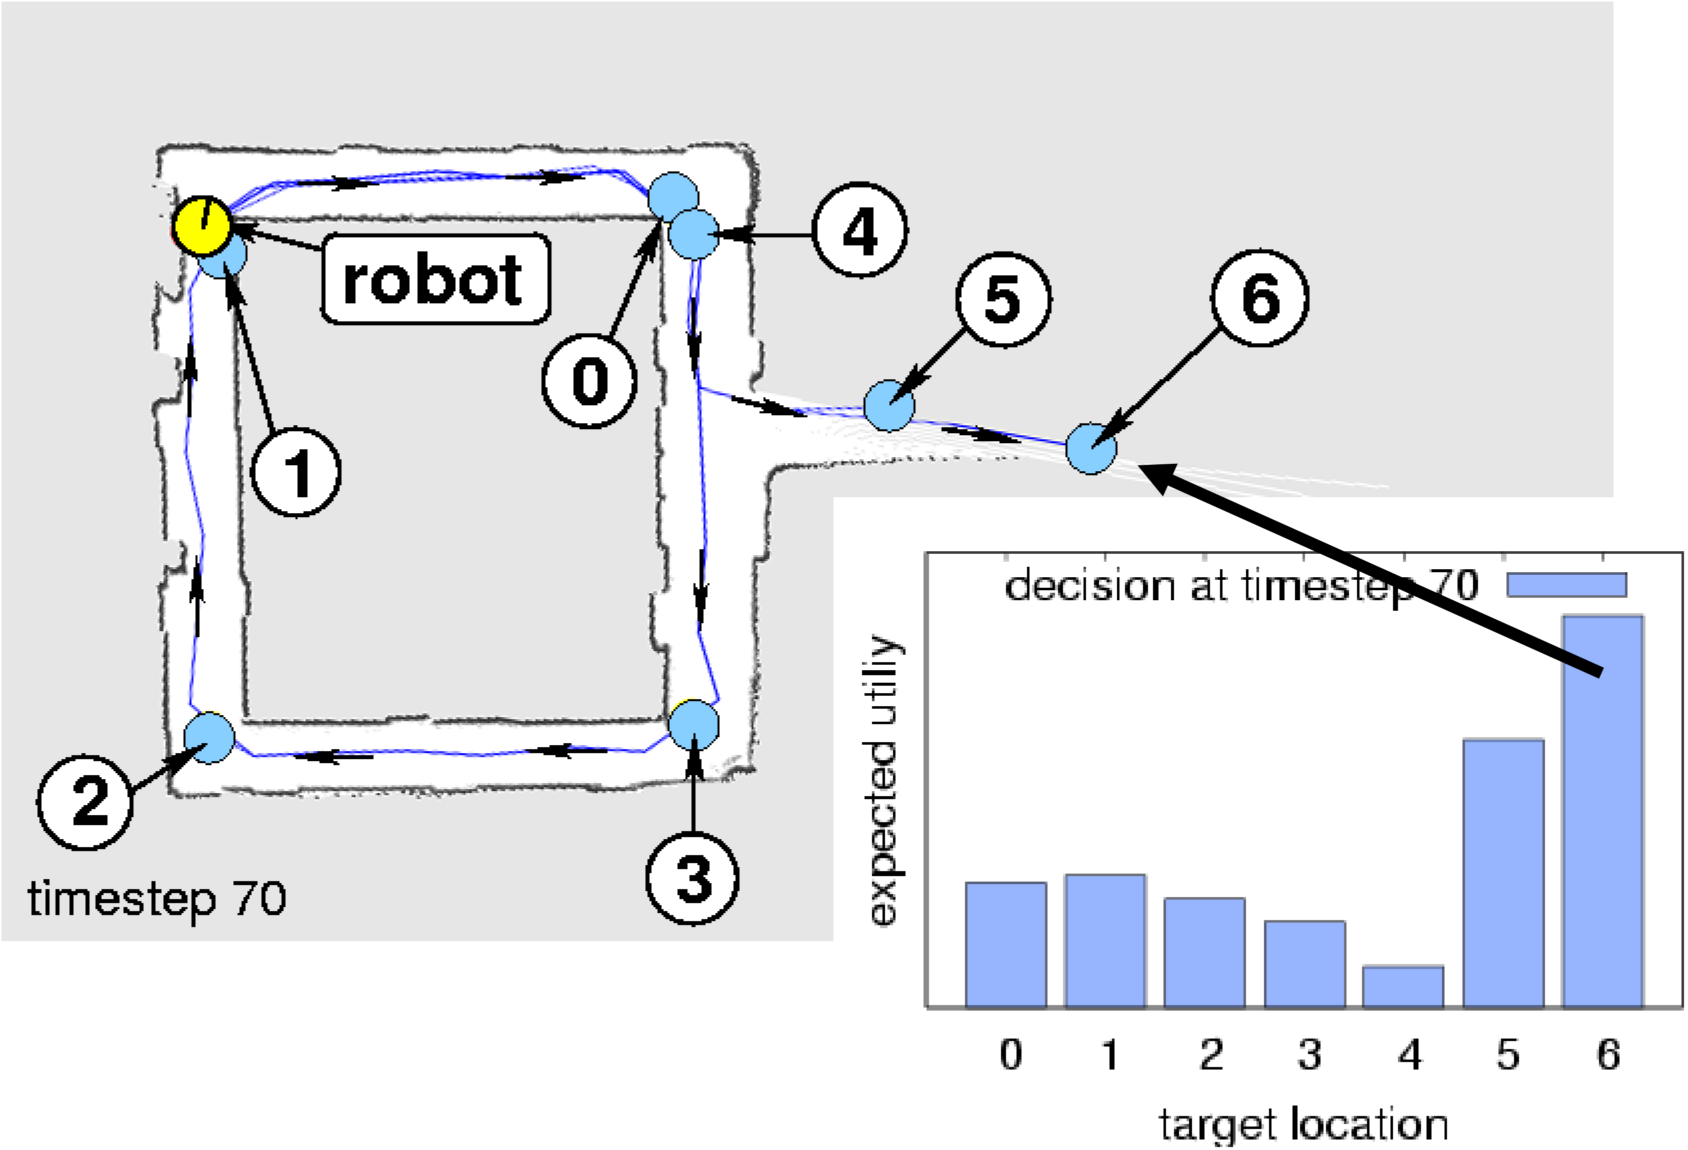

Example: Move Robot

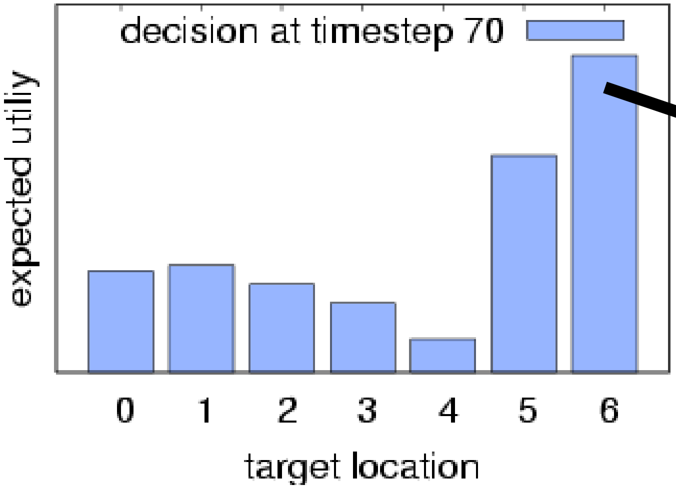

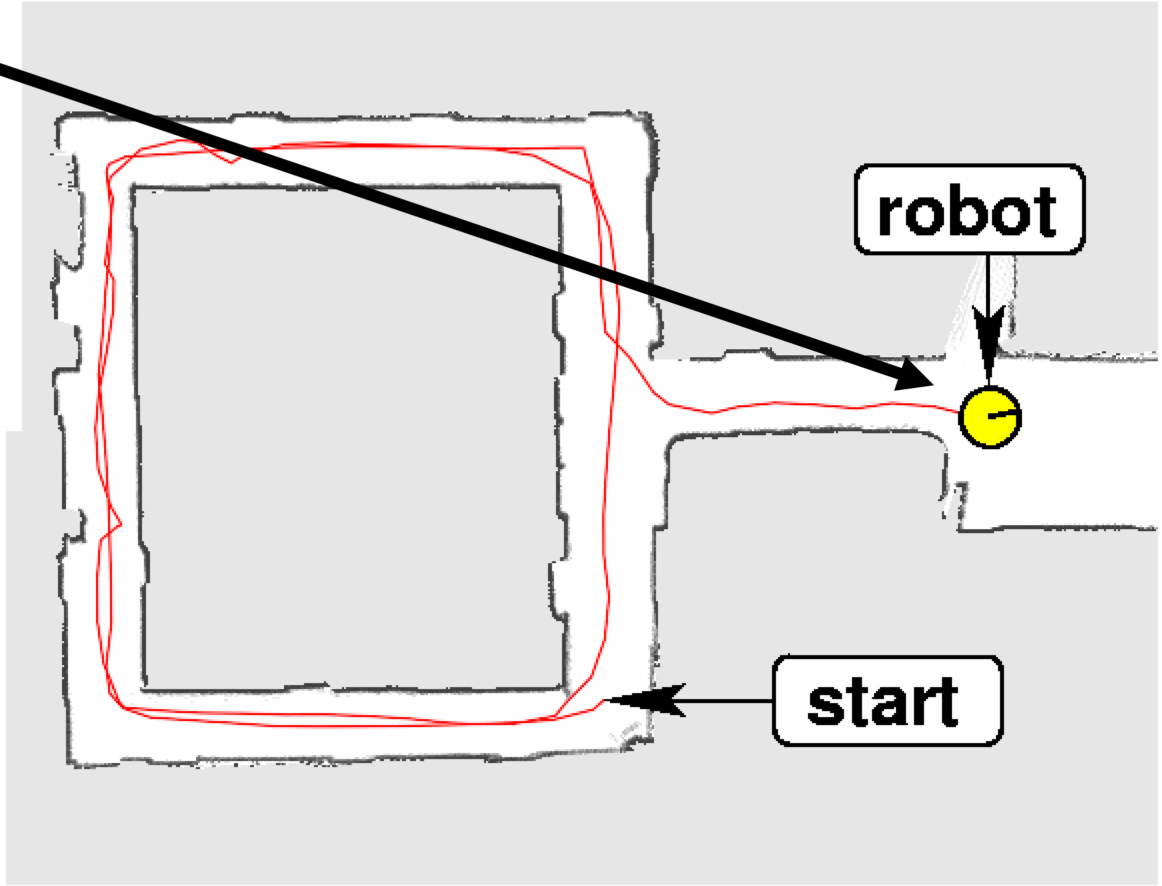

... continue ...

Example: Entropy Evolution

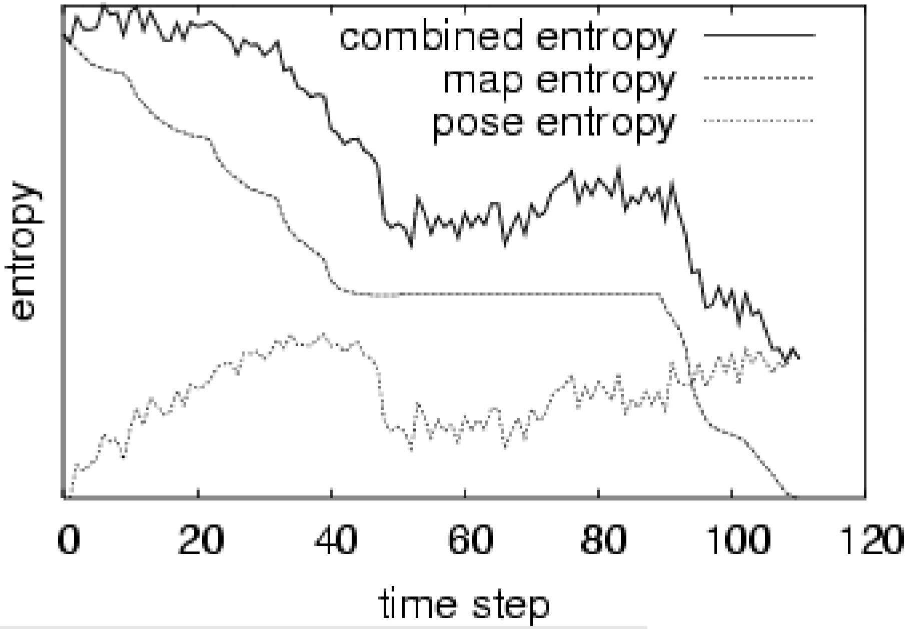

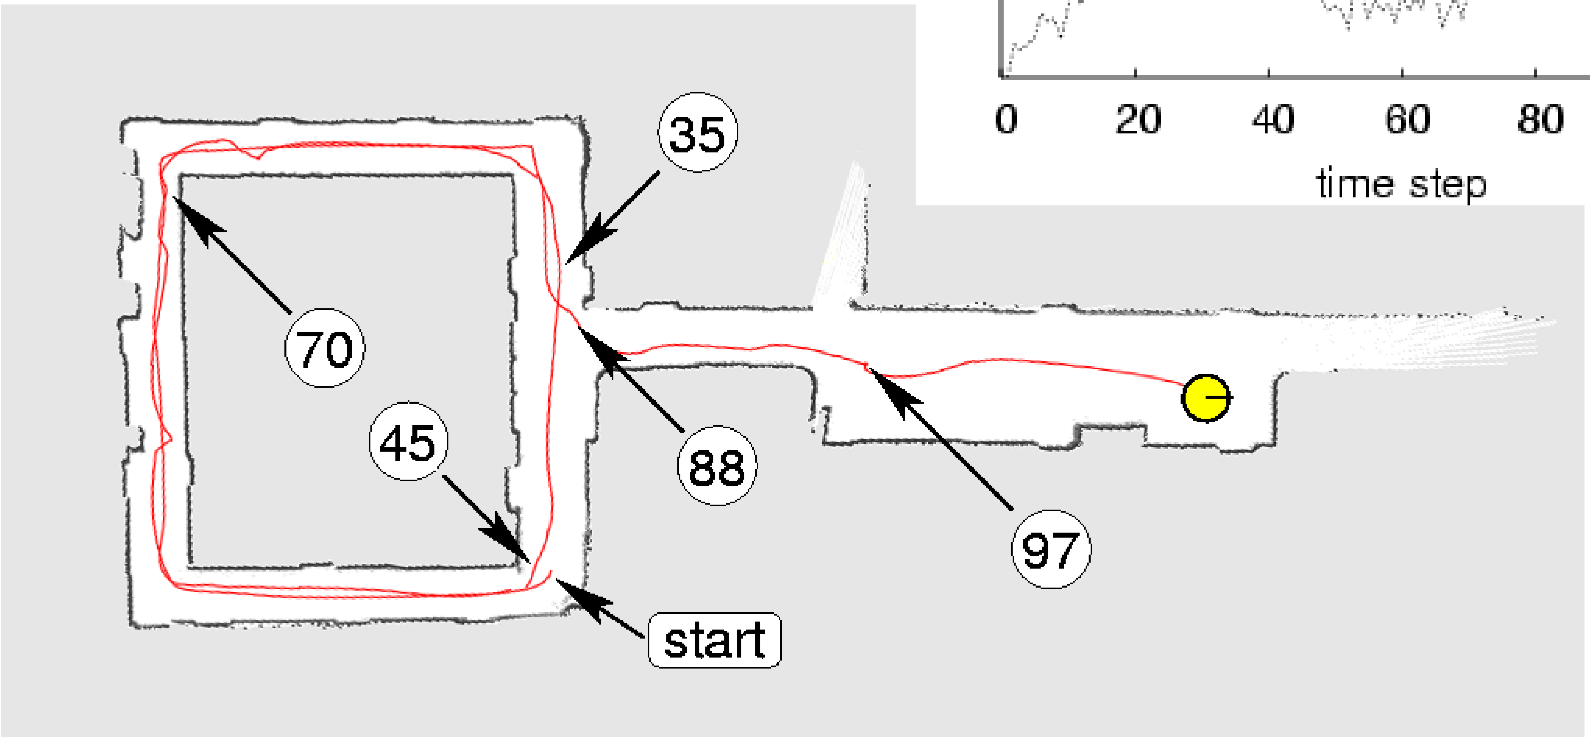

Comparison

Map uncertainty only:

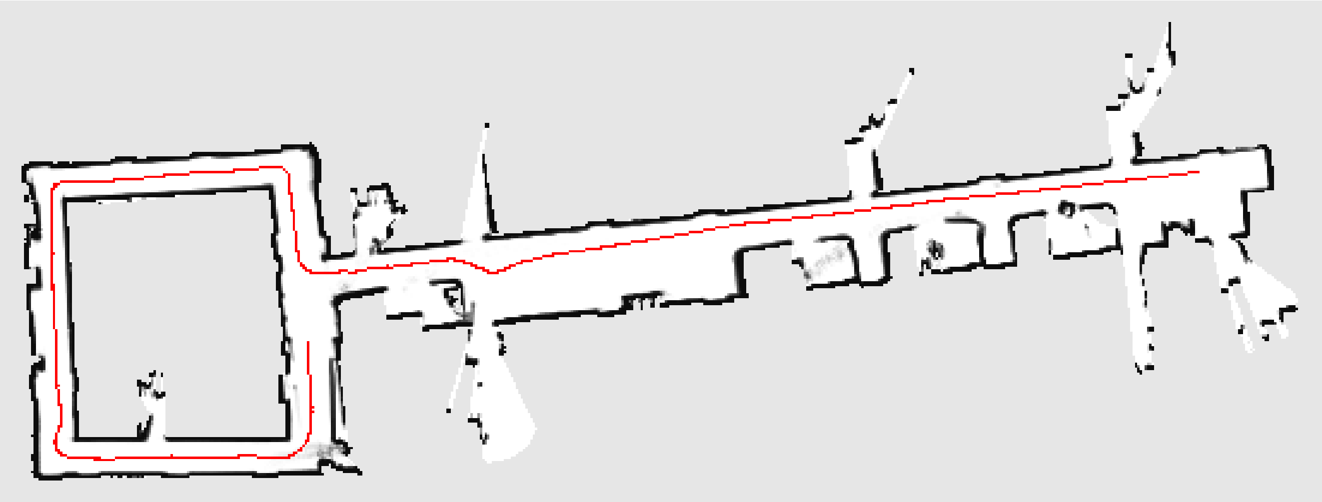

After loop closing action:

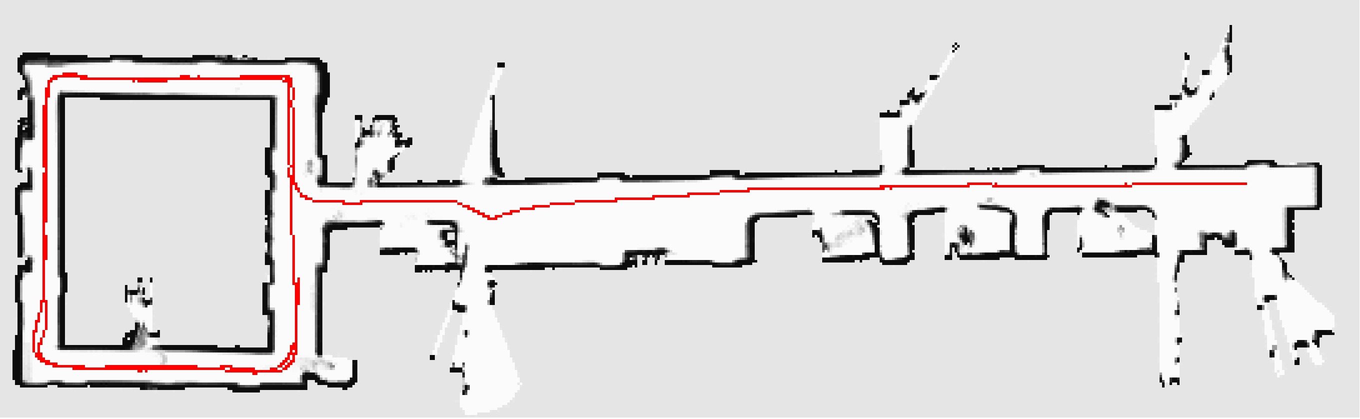

Real Exploration Example

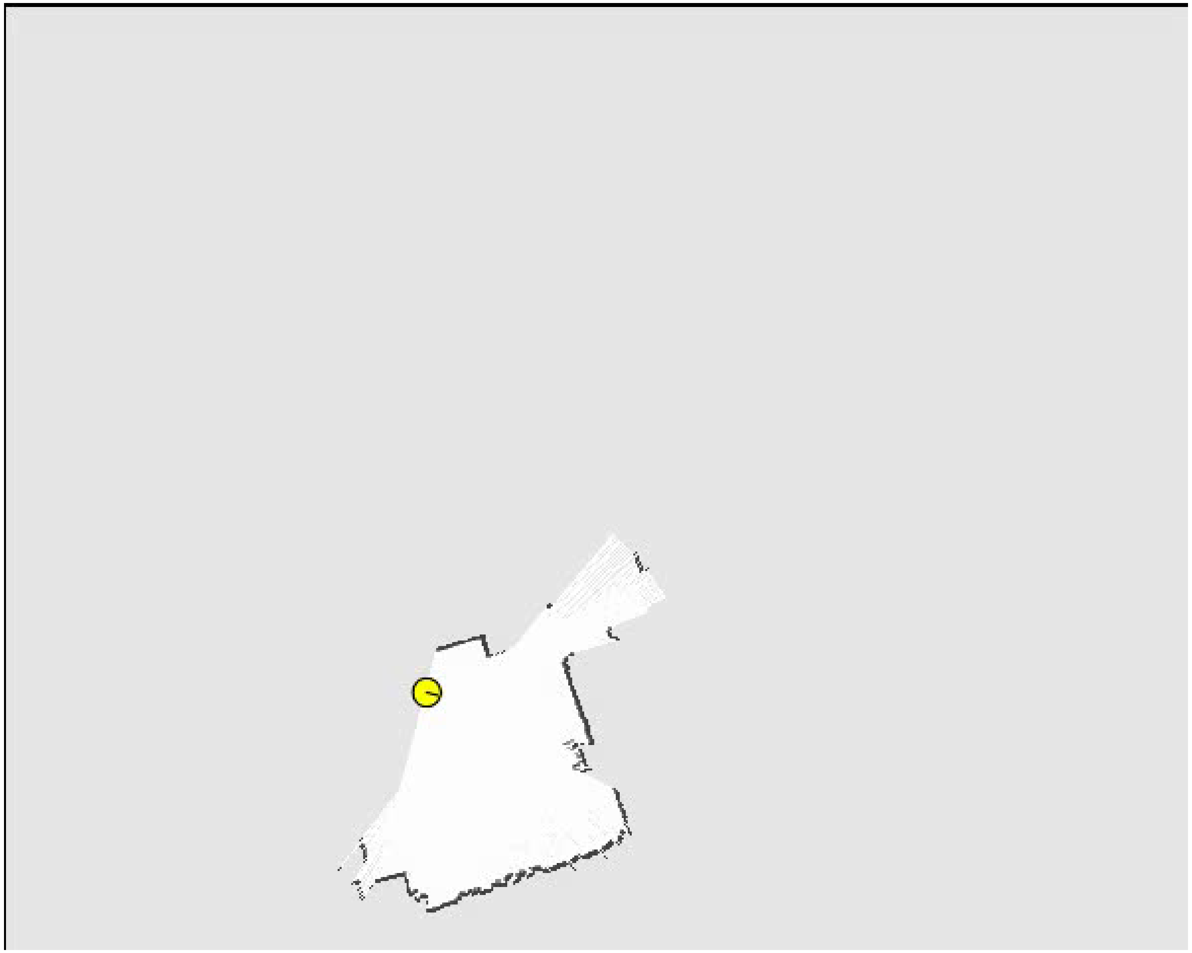

Selected target location

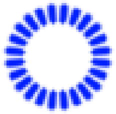

Corridor Exploration

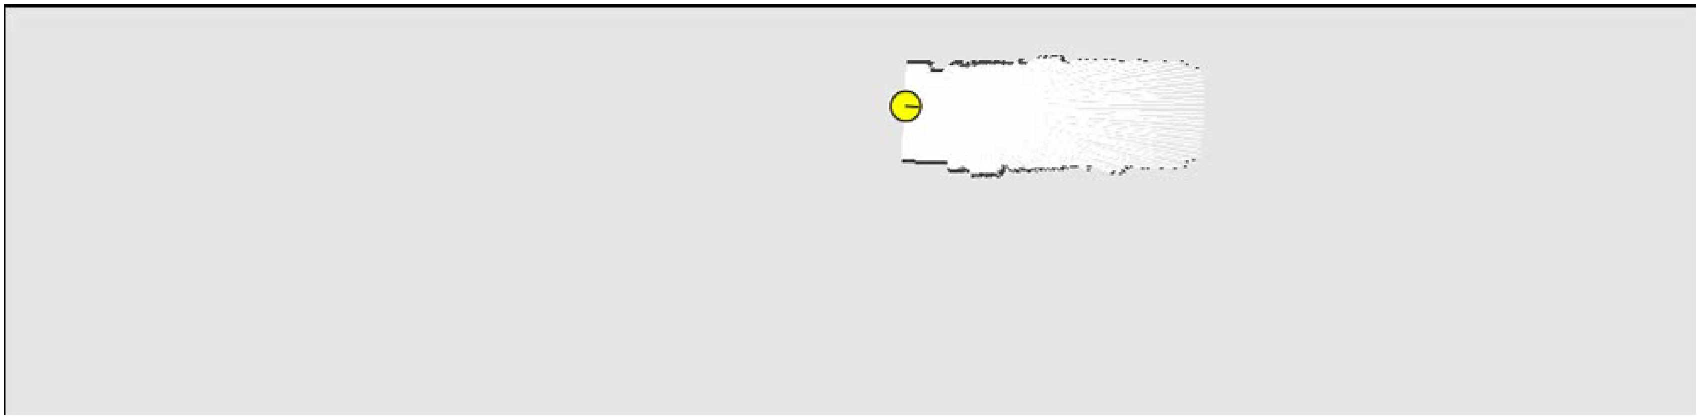

- The decision-theoretic approach leads to **intuitive behaviors:** “re-localize before getting lost”
- Some animals show a similar behavior (dogs marooned in the tundra of north Russia)

Summary

- A decision-theoretic approach to exploration in the context of RBPF-SLAM
- The approach utilizes the factorization of the Rao-Blackwellization to efficiently calculate the expected information gain
- Reasons about measurements obtained along the path of the robot
- Considers a reduced action set consisting of exploration, loop-closing, and place-revisiting actions