Introduction to Mobile Robotics

Information Gain-Based Exploration

Wolfram Burgard, Cyrill Stachniss,

Maren Bennewitz, Kai Arras

Tasks of Mobile Robots

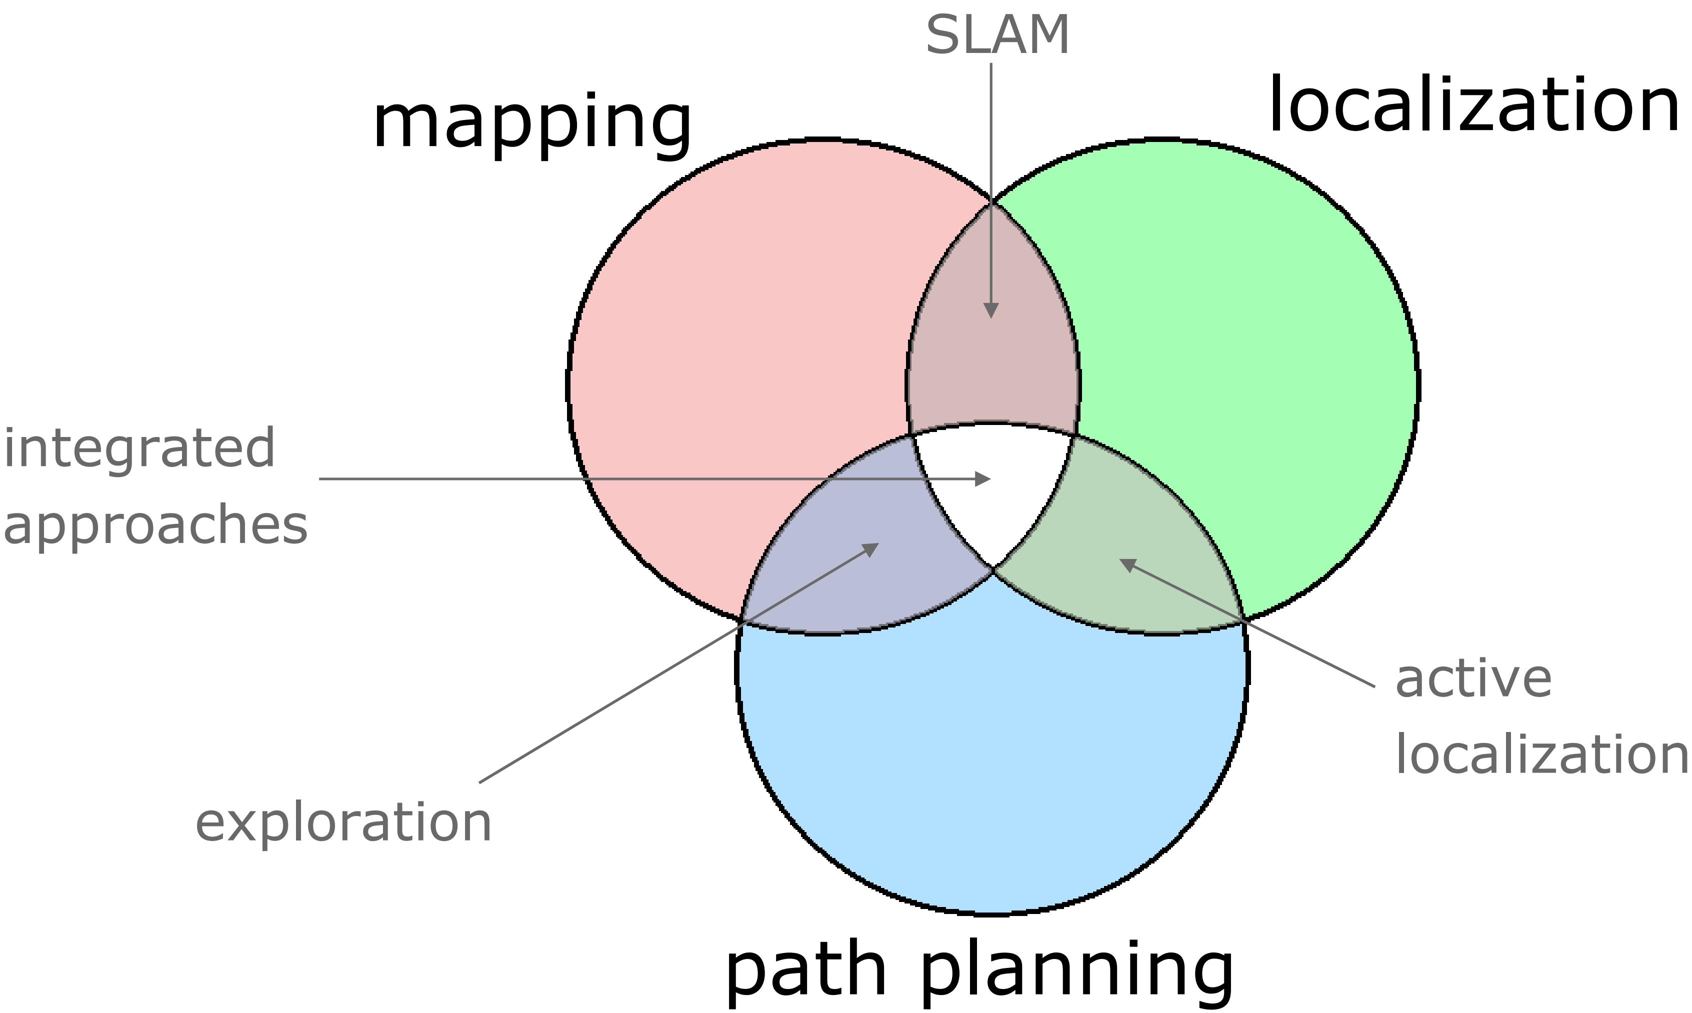

Exploration and SLAM

- 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

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

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

Mapping with known poses

Particle filter representing trajectory hypotheses

Example: Particle Filter for Mapping

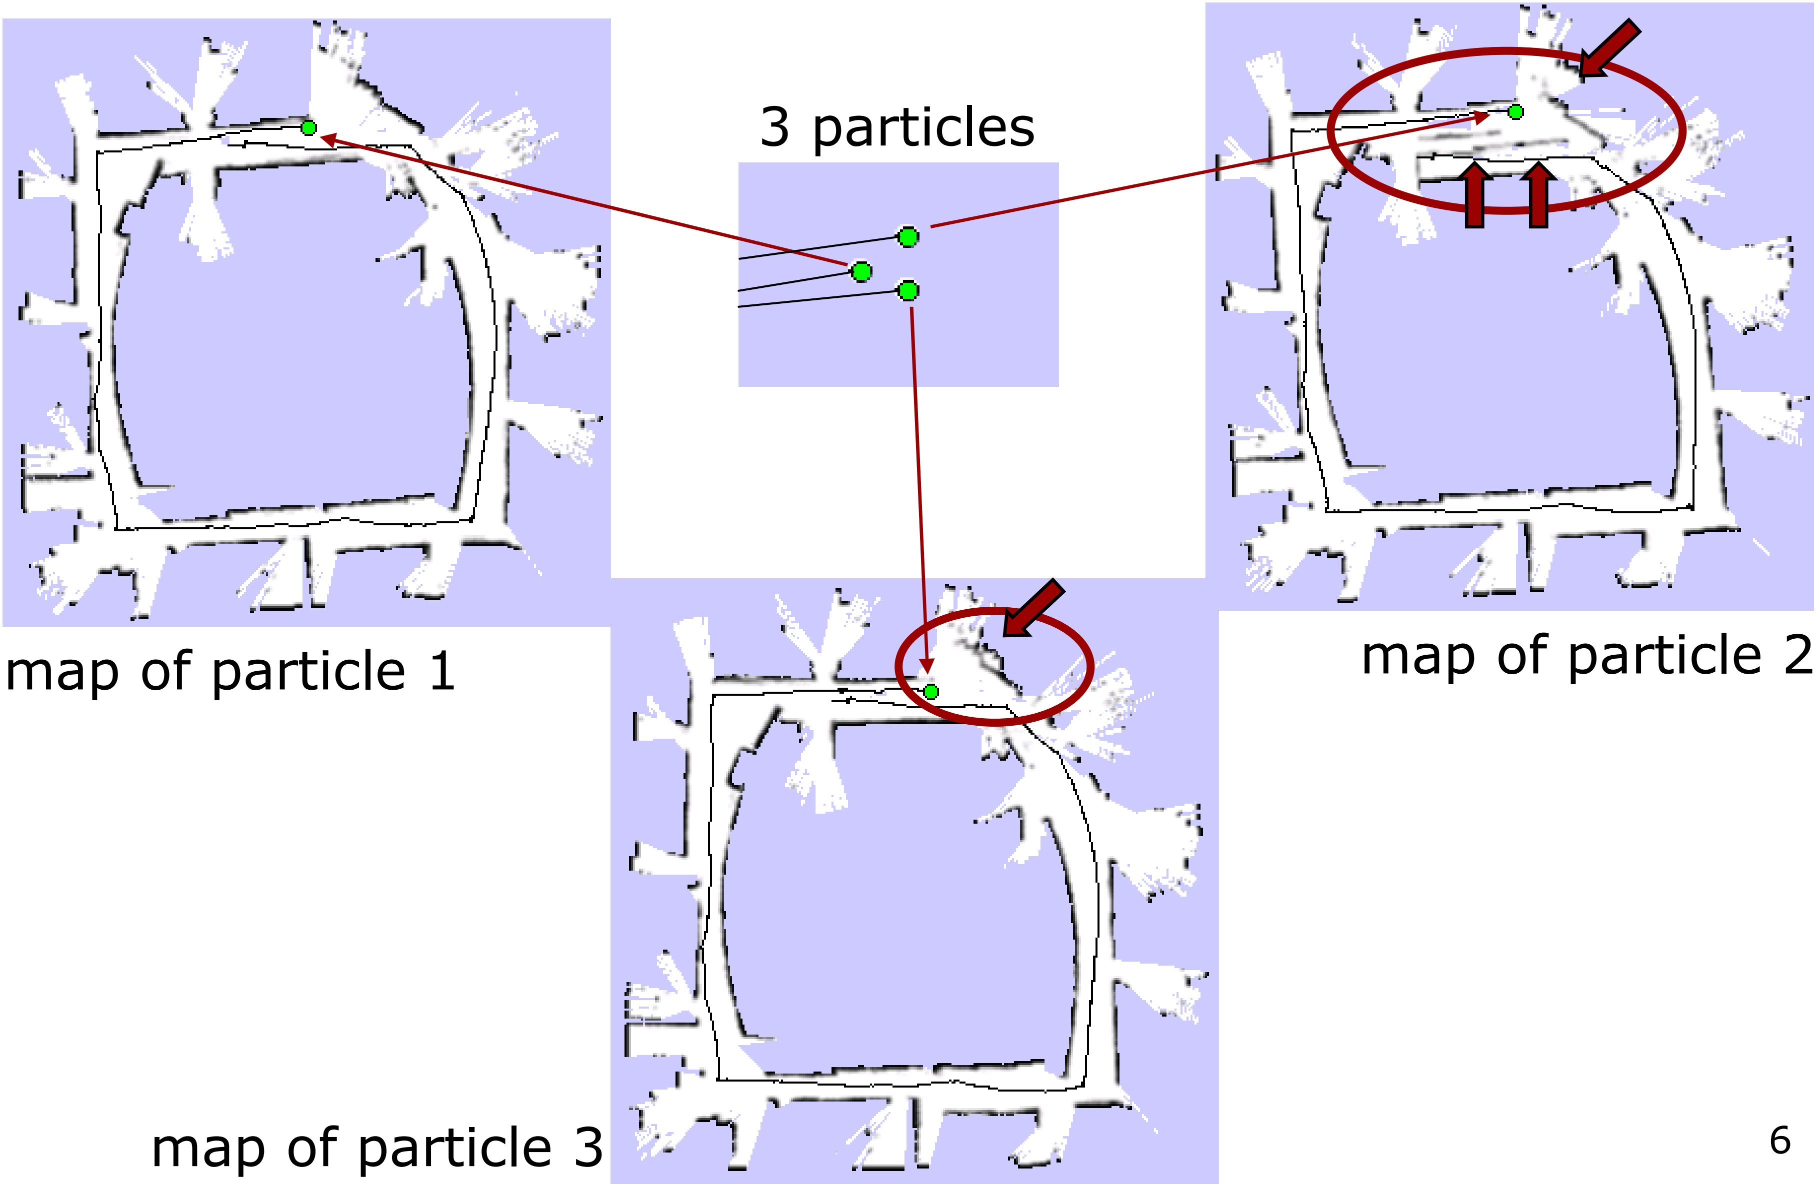

Outdoor Campus Map

- **30 particles**
- 250x250m^2
- 1.75 km (odometry)
- 20cm resolution during scan matching
- 30cm resolution in final map

Combining Exploration and SLAM

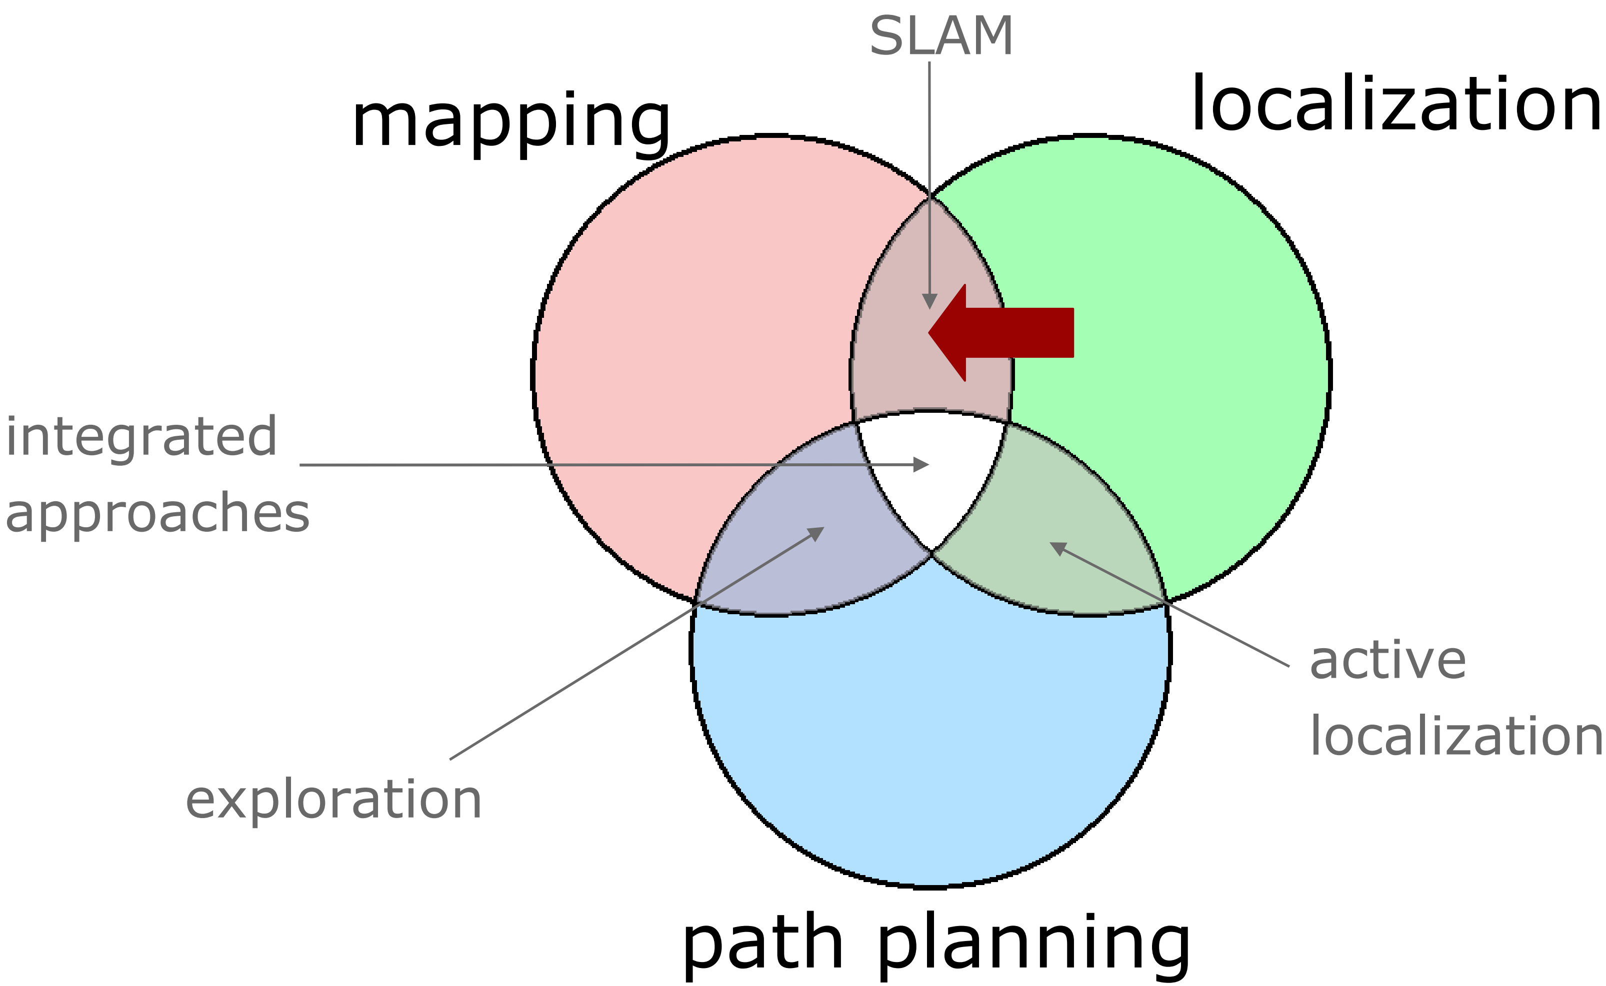

Exploration

- The 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?

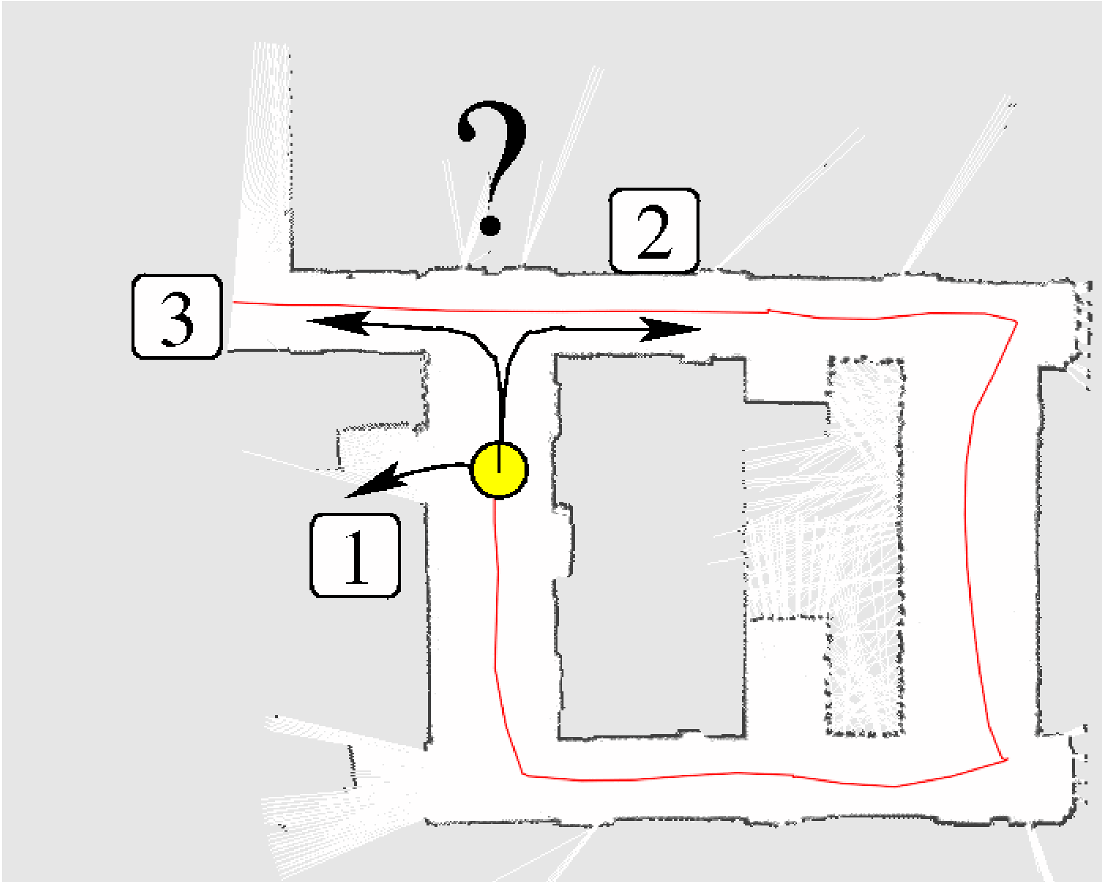

Decision-Theoretic Approach

- 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

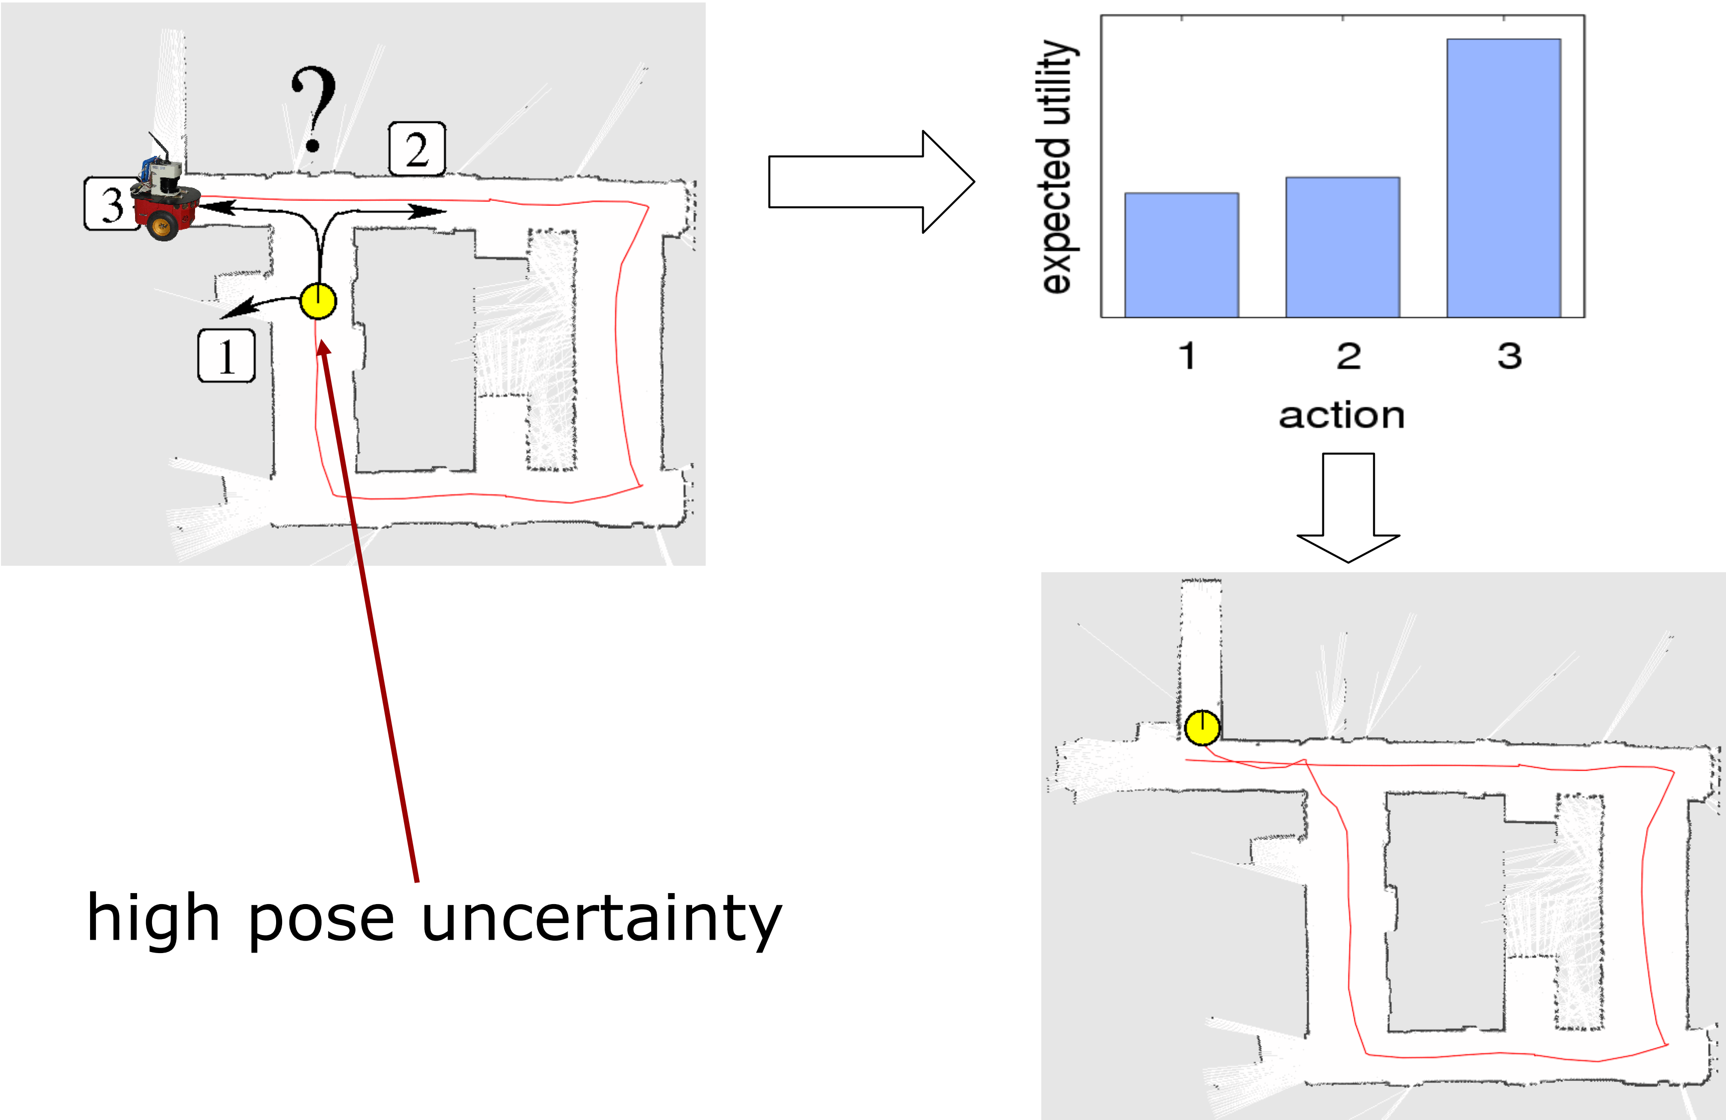

high pose uncertainty

The Uncertainty of a Posterior

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

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

- Information Gain = Uncertainty Reduction

$$I(t+1 \mid t) \;=\; H(p(x_t)) - H(p(x_{t+1}))$$

Entropy Computation

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

Computing the Map and Pose Uncertainty

data (laser and odometry)

$$H(p(x, m \mid d))$$

$$= \quad H(p(x \mid d)) + \int_x p(x \mid d) H(p(m \mid x, d))\, dx$$

$$\approx \quad H(p(x \mid d)) + \sum_{i=1}^{\#particles} \omega^{[i]} H(p(m^{[i]} \mid x^{[i]}, d))$$

trajectory uncertainty

particle weight

map uncertainty

Computing the Entropy of the Map Posterior

Occupancy Grid map *m*:

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

map uncertainty

grid cells

probability that the cell is occupied

Map Entropy

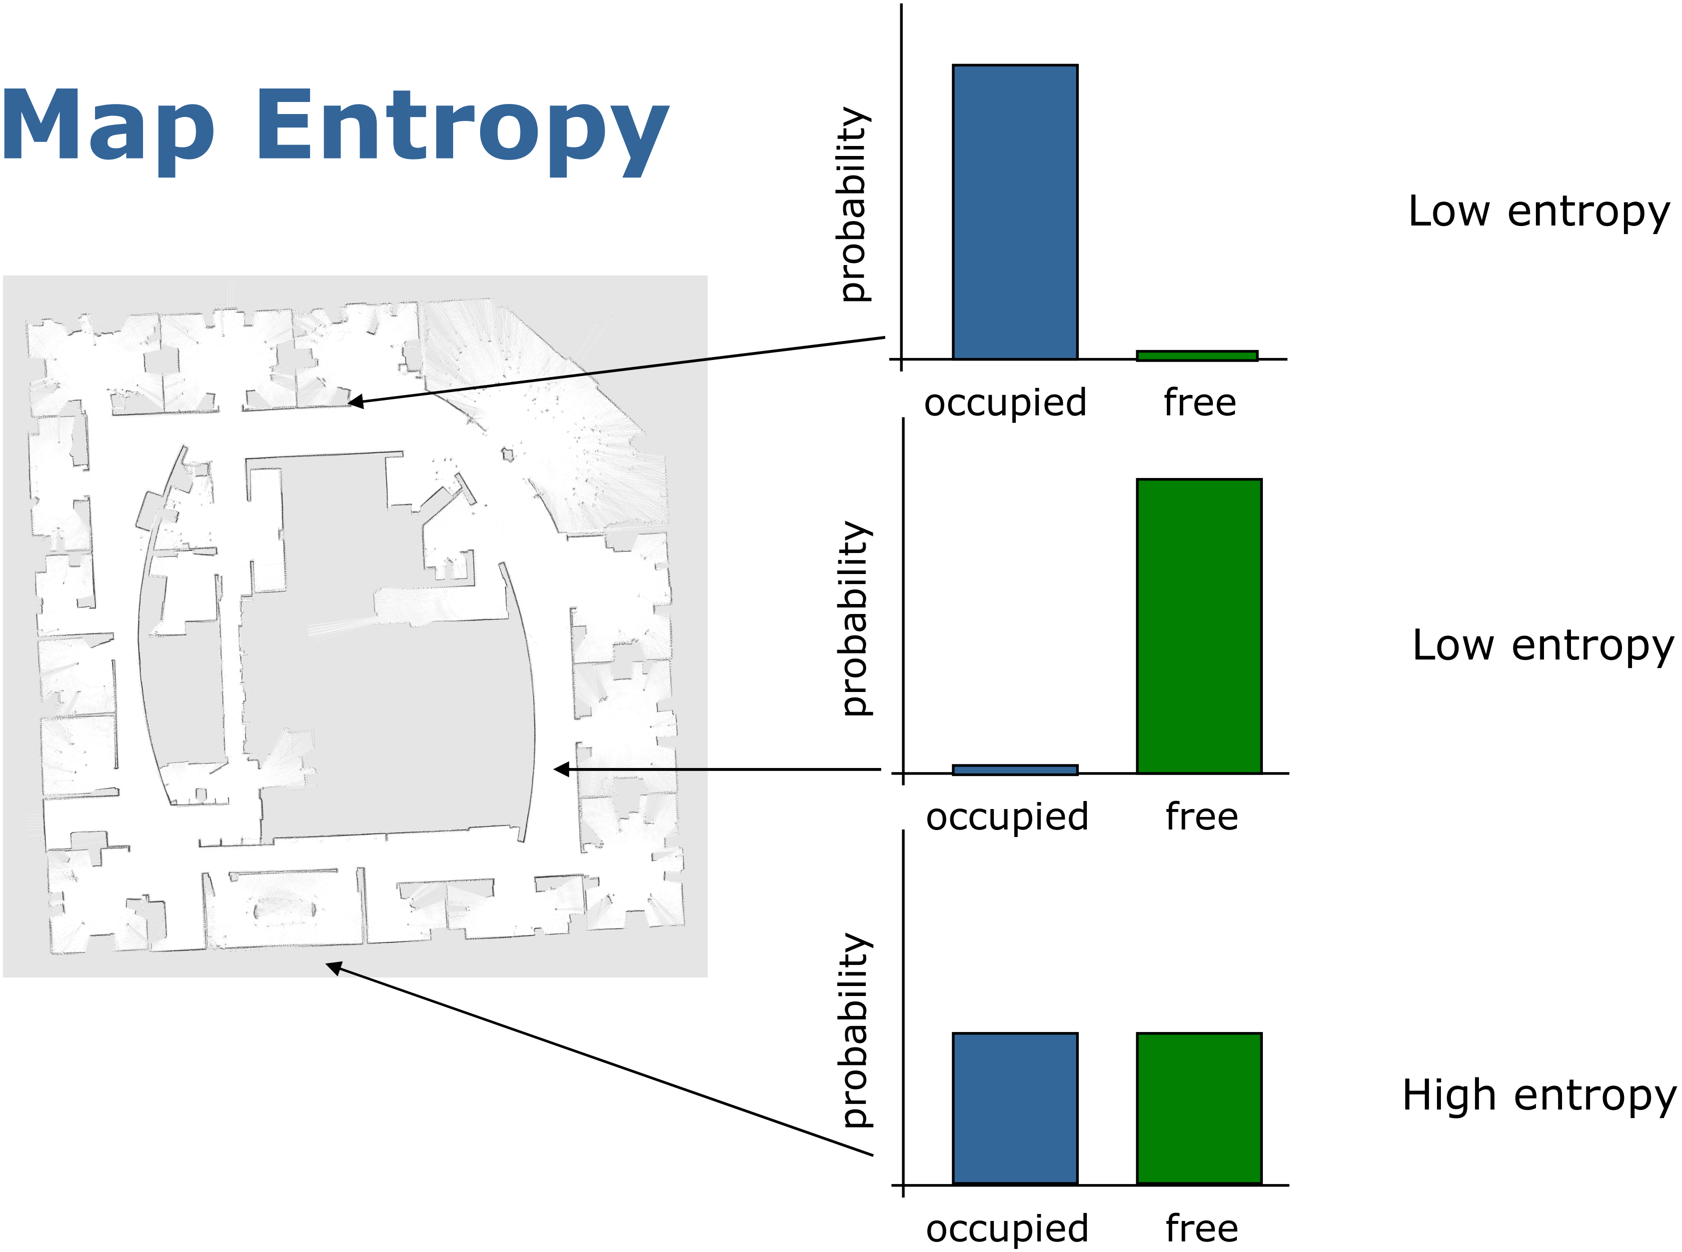

The overall entropy is the sum of the individual entropy values

Computing the Entropy of the Trajectory Posterior

1. High-dimensional Gaussian

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

reduced rank for sparse particle sets

2. Grid-based approximation

$$H(p(x \mid d)) \;\rightsquigarrow\; const.$$

for sparse particle clouds

Approximation of the Trajectory Posterior Entropy

Average pose entropy over time:

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

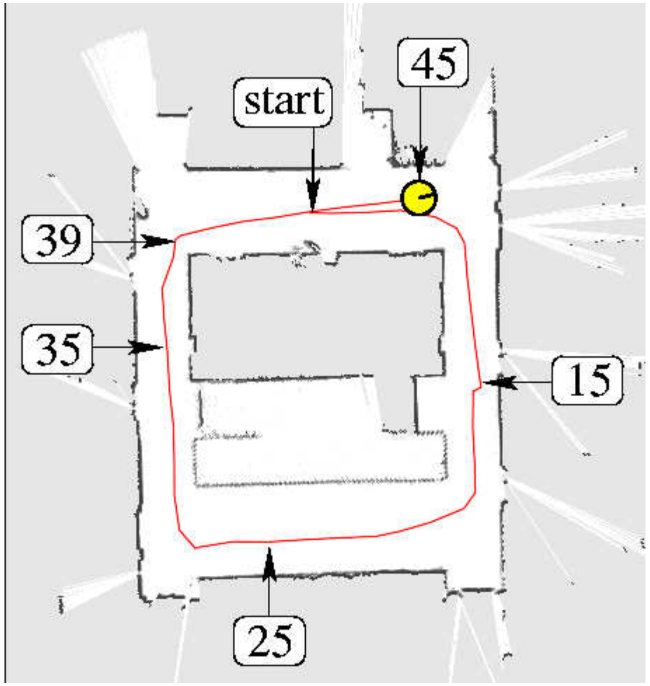

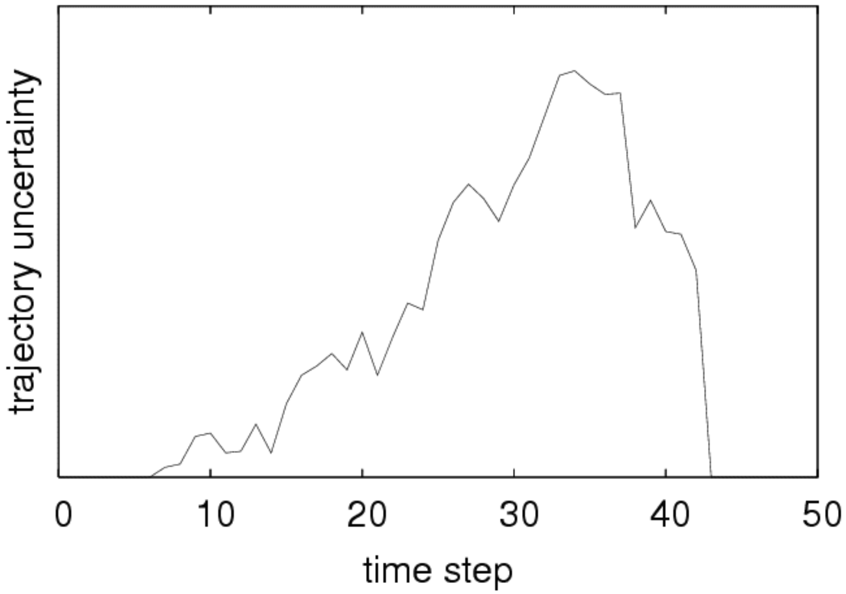

Information Gain

- The reduction of entropy in the model

observations
to be obtained

action

H before action
is carried out

$$I(\widehat{z}, a) = \\ H(p(m, x \mid d)) - \\ H(p(m, x, \widehat{x} \mid d, a, \widehat{z}))$$

new poses introduced
by action

H after action is
carried out

Computing the Expected Information Gain

- To compute the information gain one needs to know the observations obtained when carrying out an action
- This quantity is not known! Reason about potential measurements

$$E[I(a)] \;=\; \int_{\widehat{z}} p(\widehat{z} \mid a, d) \cdot I(\widehat{z}, a)\, d\widehat{z}$$

Reasoning about Measurements

- The filter represents a posterior about possible maps
- Use these maps to reason about possible observation
- Simulate laser measurements in the maps of the particles

$$E[I(a)] = \int_{\hat{z}} p(\hat{z} \mid a, d) \cdot I(\hat{z}, a) \, d\hat{z}$$

measurement sequences simulated in the maps

likelihood (particle weight)

Reasoning about Measurements

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

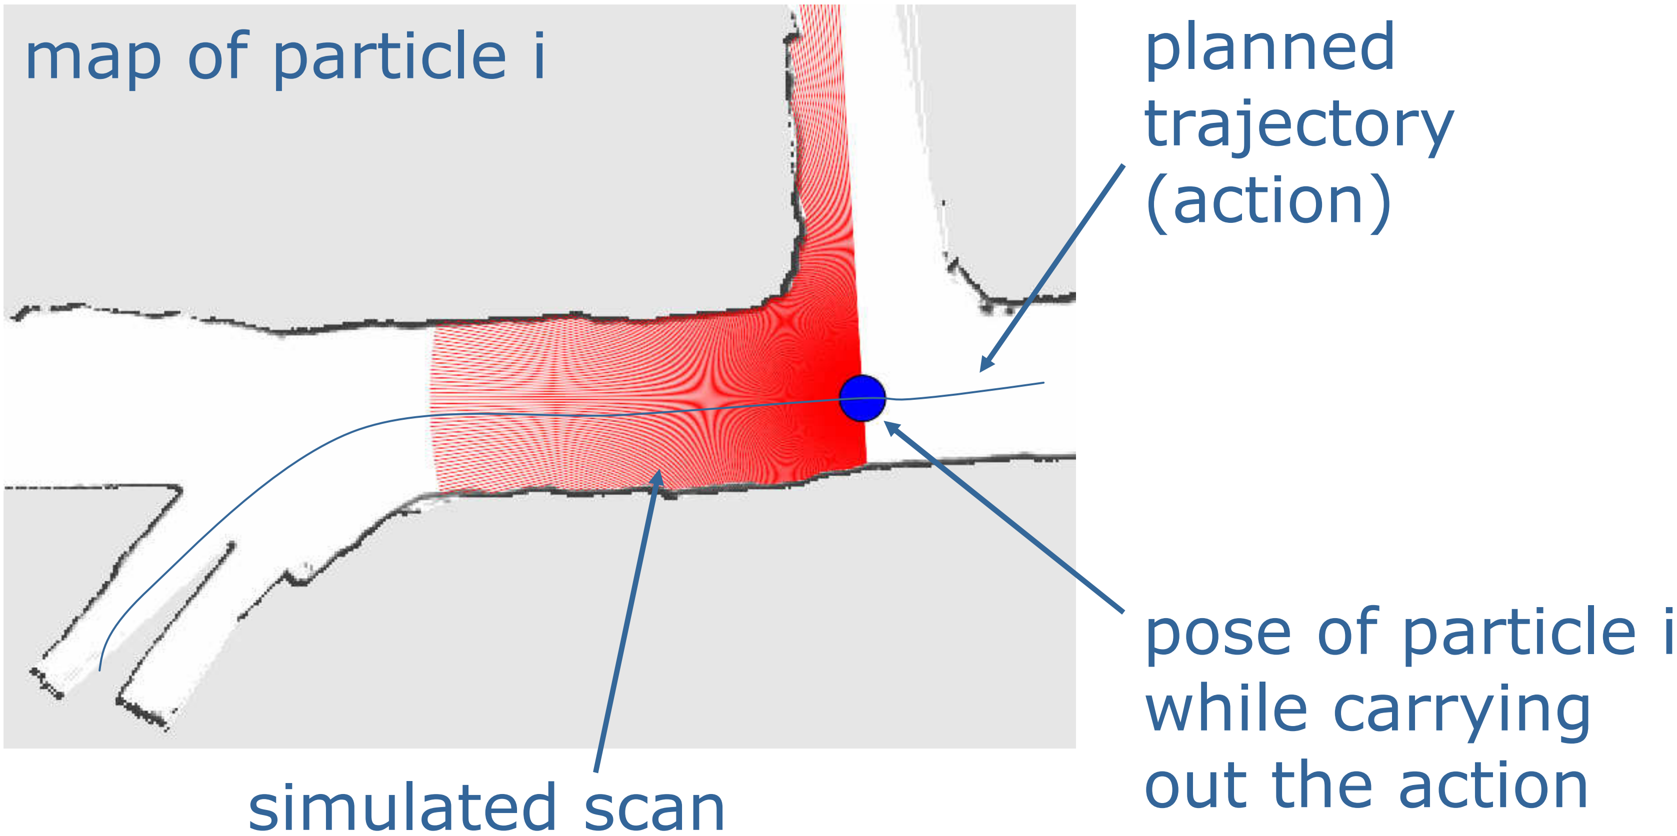

The Utility

- To take into account the cost of an action, we compute a utility

$$E[U(a)] \;=\; I(a) - \alpha \cdot cost(a)$$

- Select the action with the highest expected utility

$$a^* \;=\; \operatorname*{argmax}_a \{E[U(a)]\}$$

Focusing on Specific Actions

To efficiently sample actions we consider

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

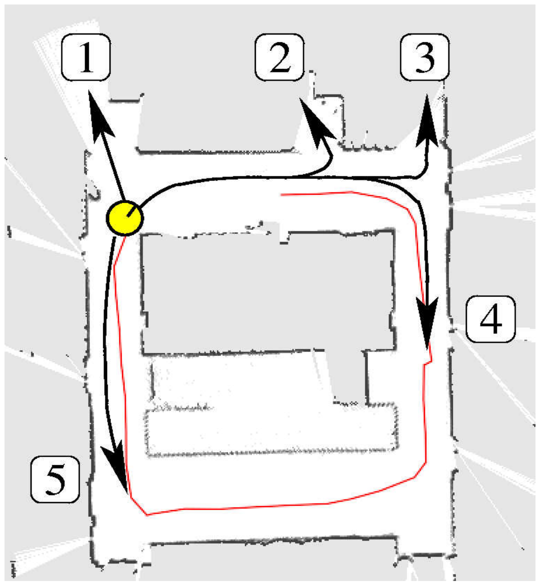

Dual Representation for Loop Detection

- **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

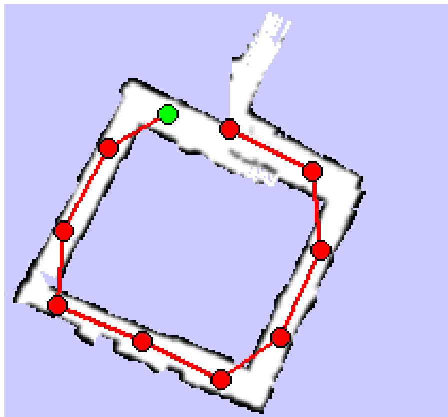

Example: Trajectory Graph

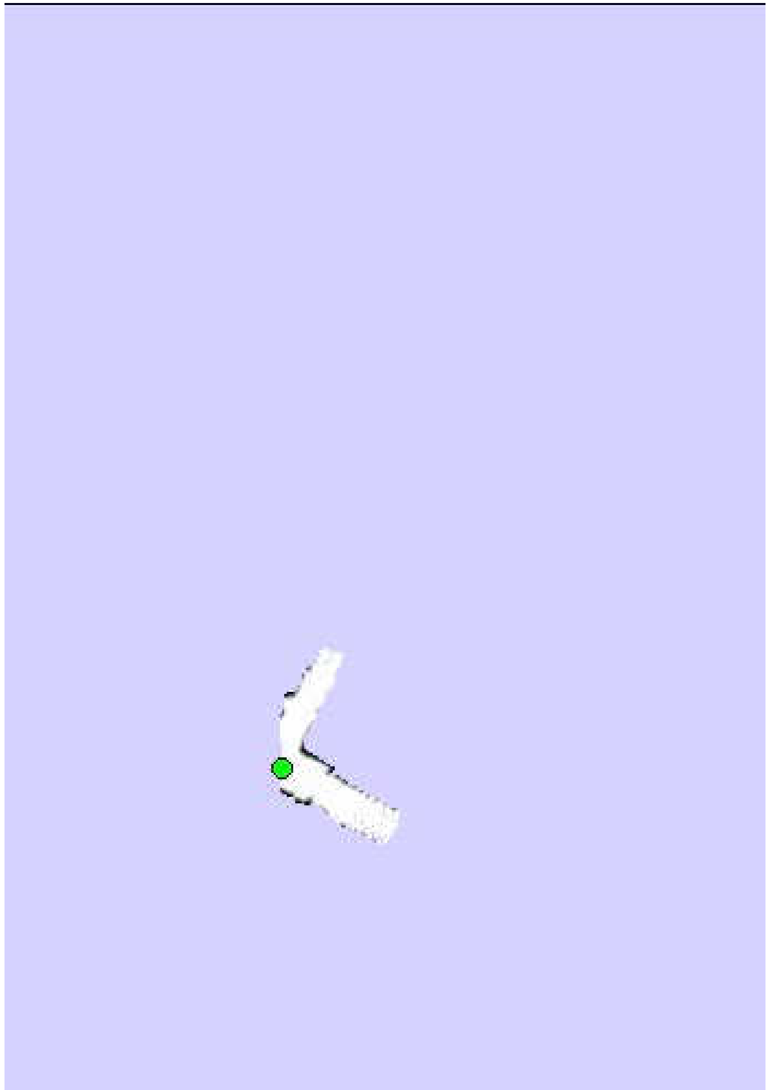

Application Example

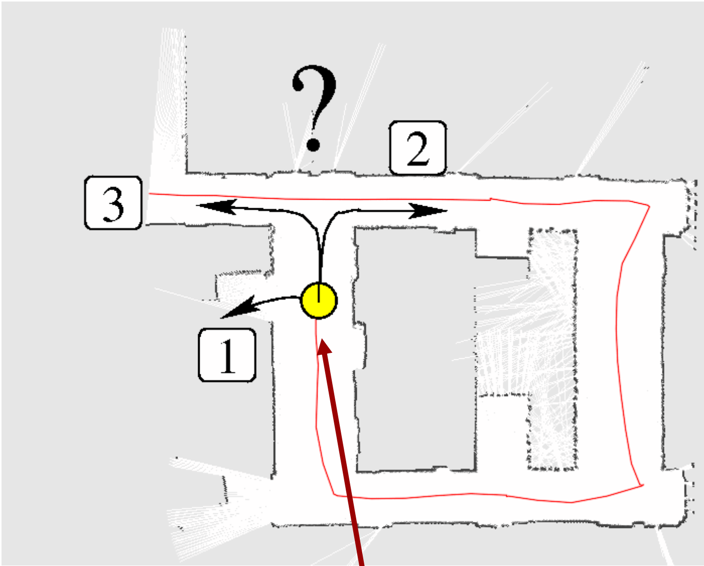

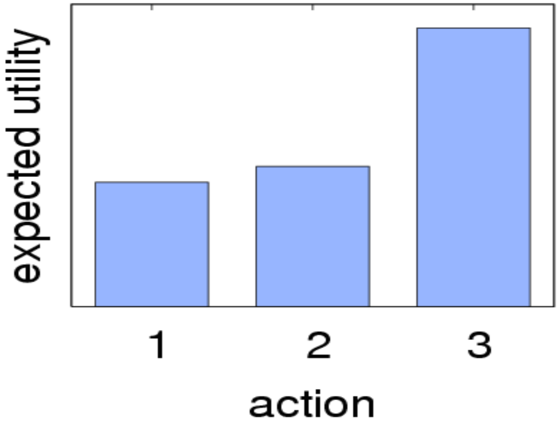

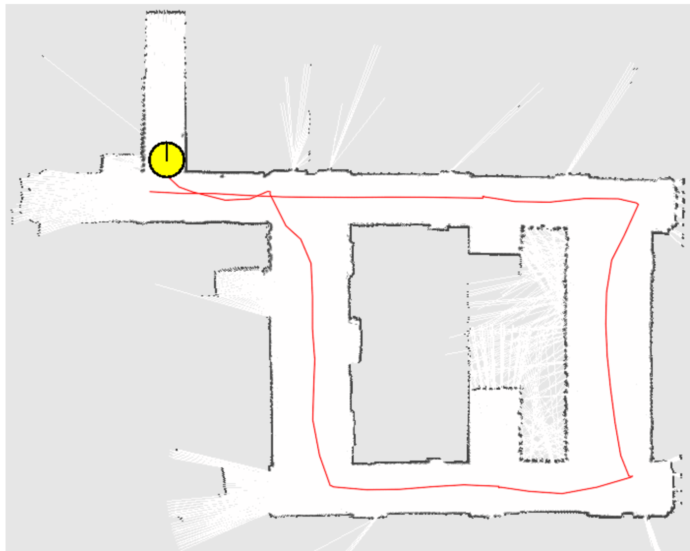

high pose uncertainty

Example: Possible Targets

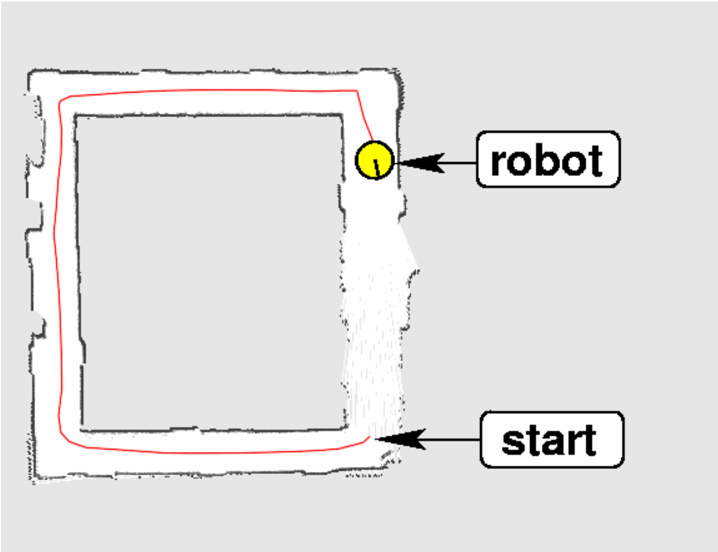

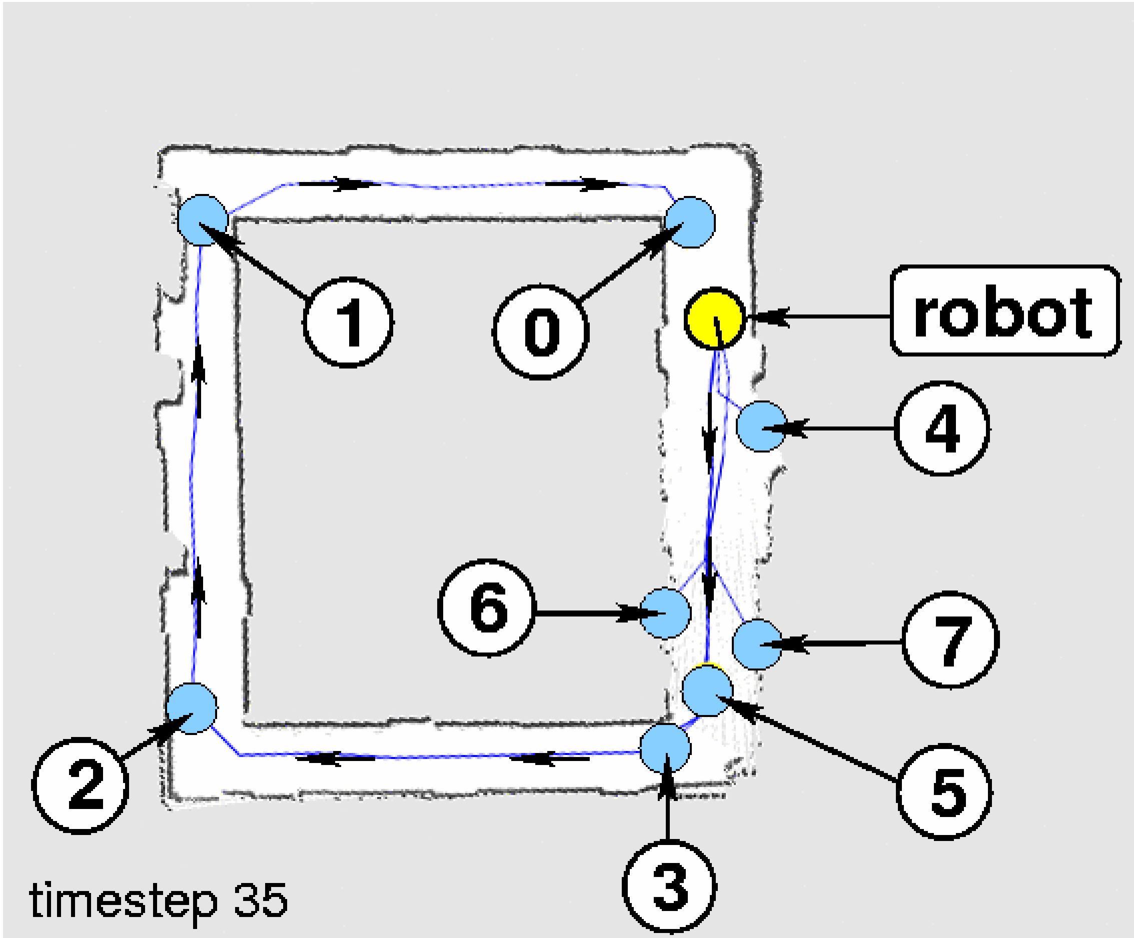

Example: Evaluate Targets

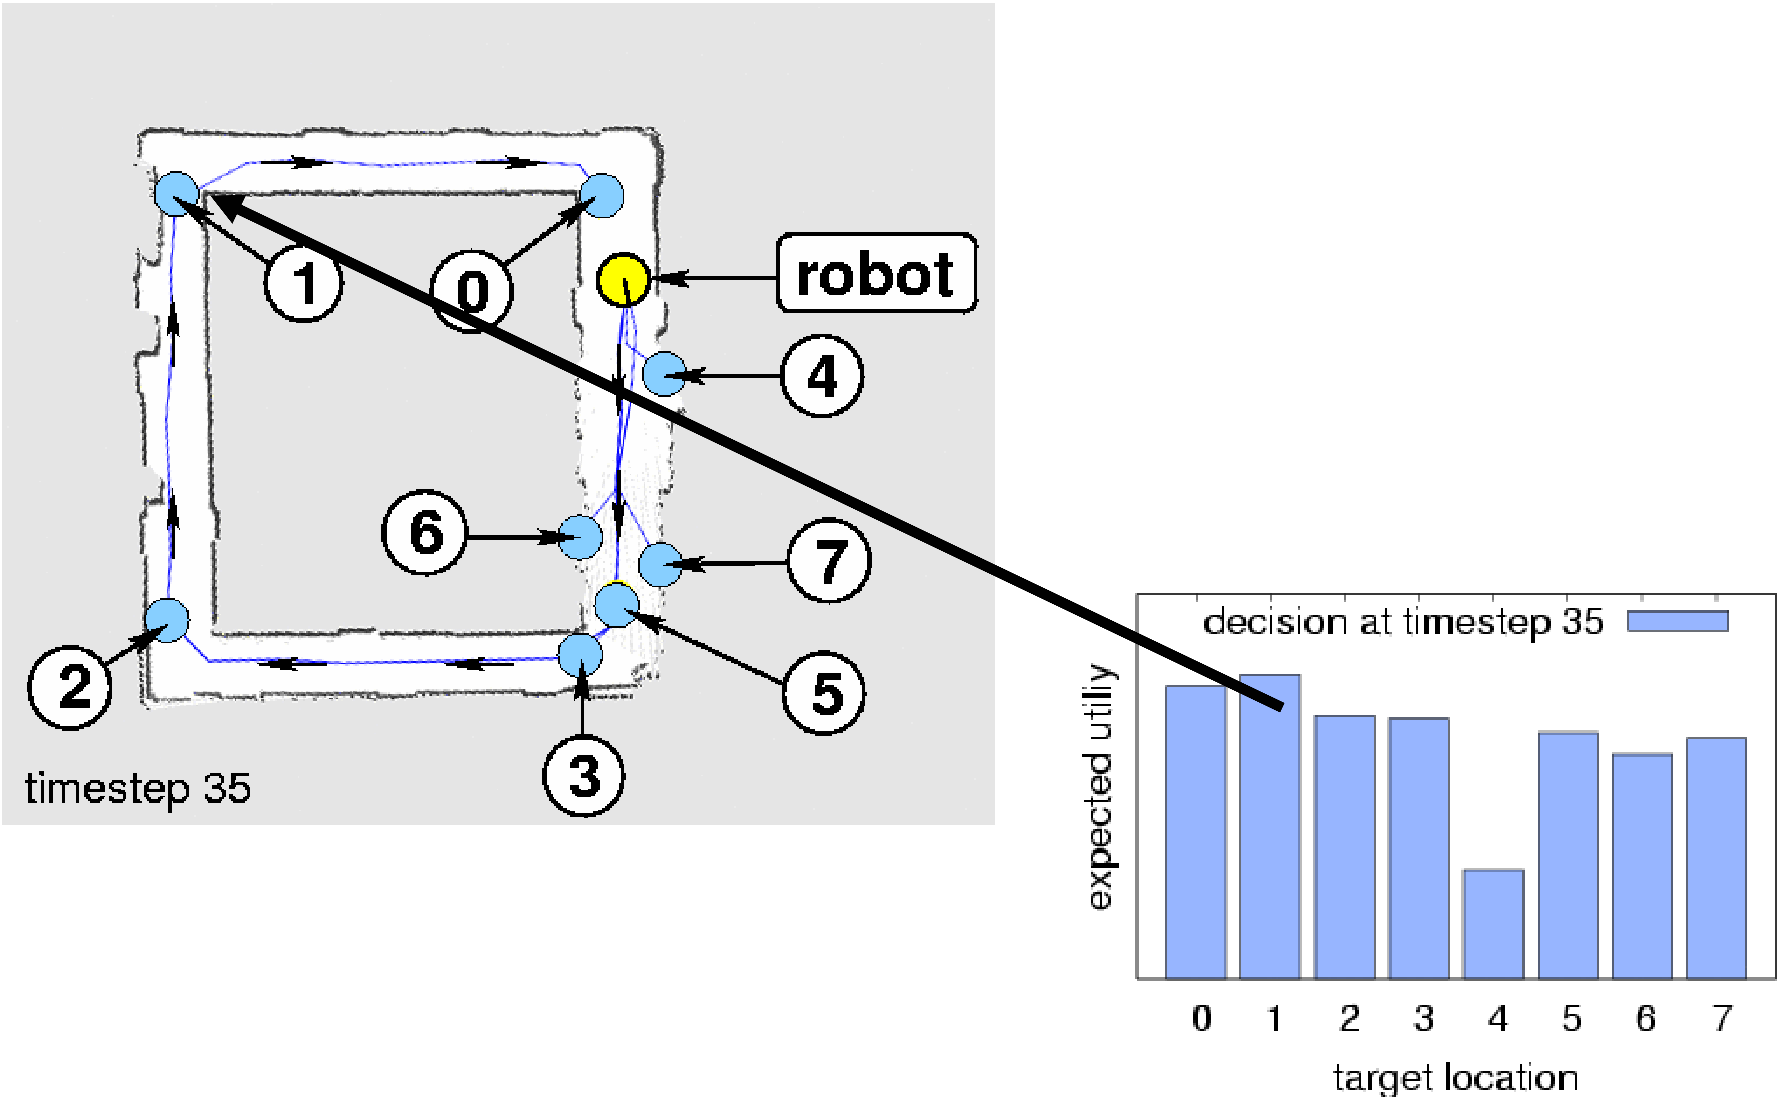

Example: Move Robot to Target

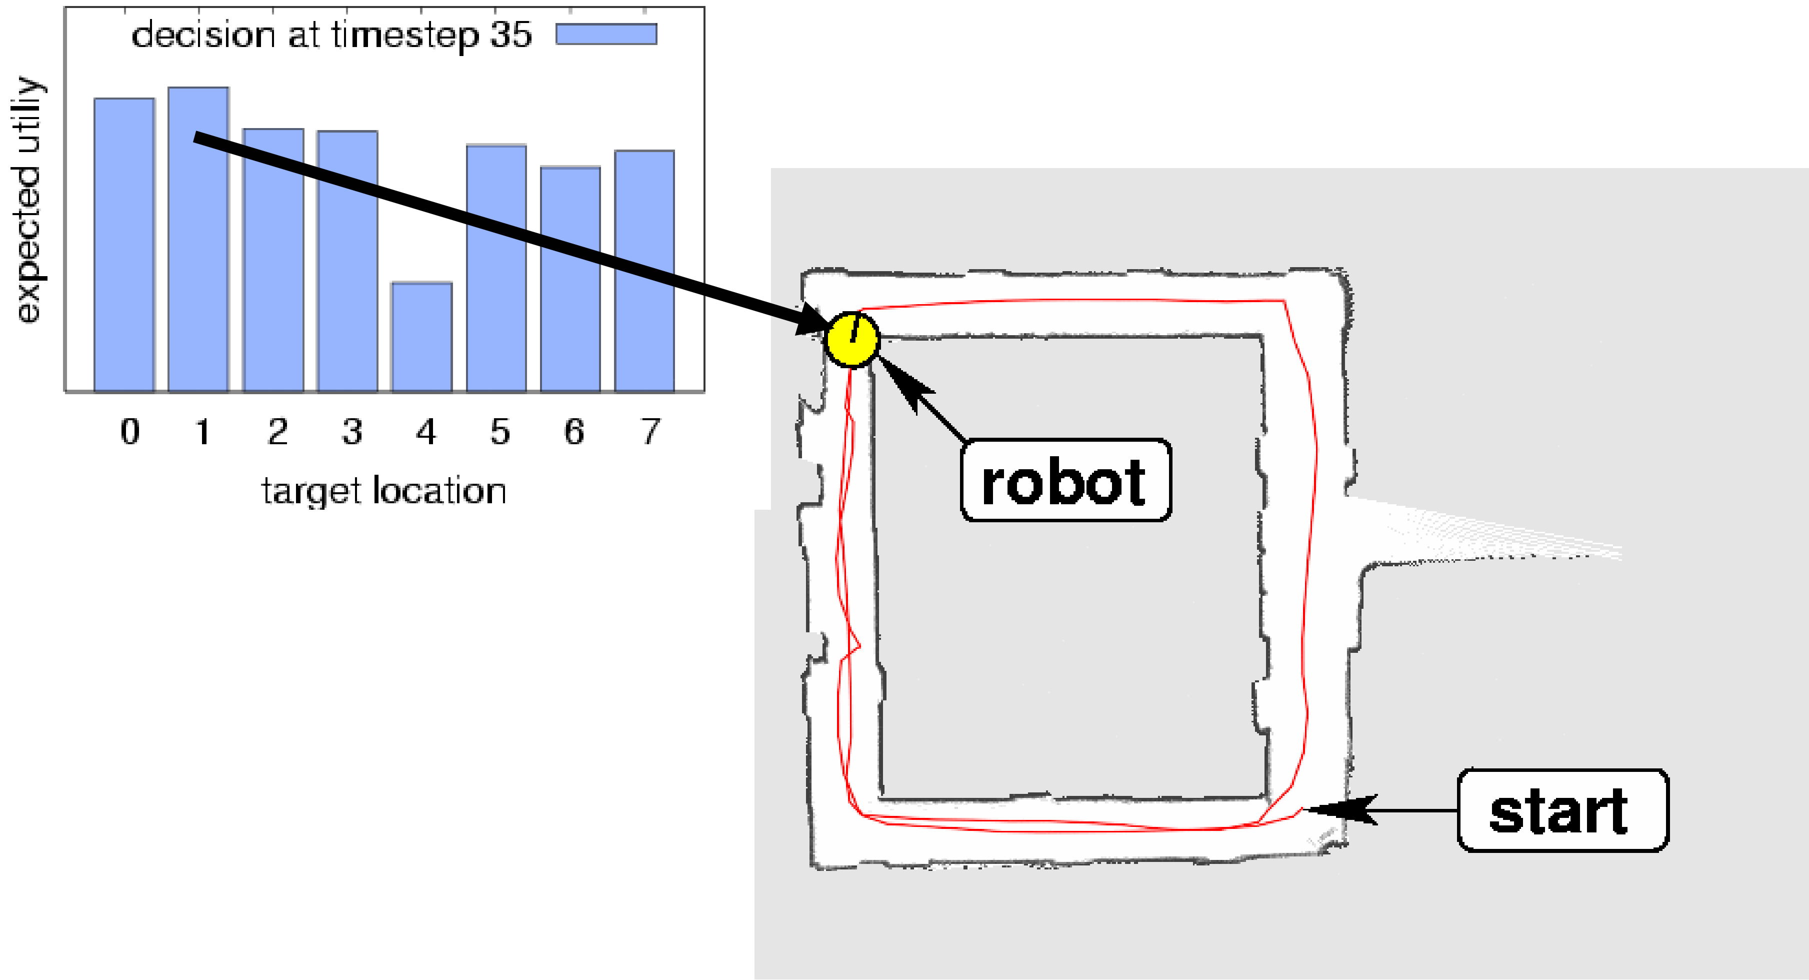

Example: Evaluate Targets

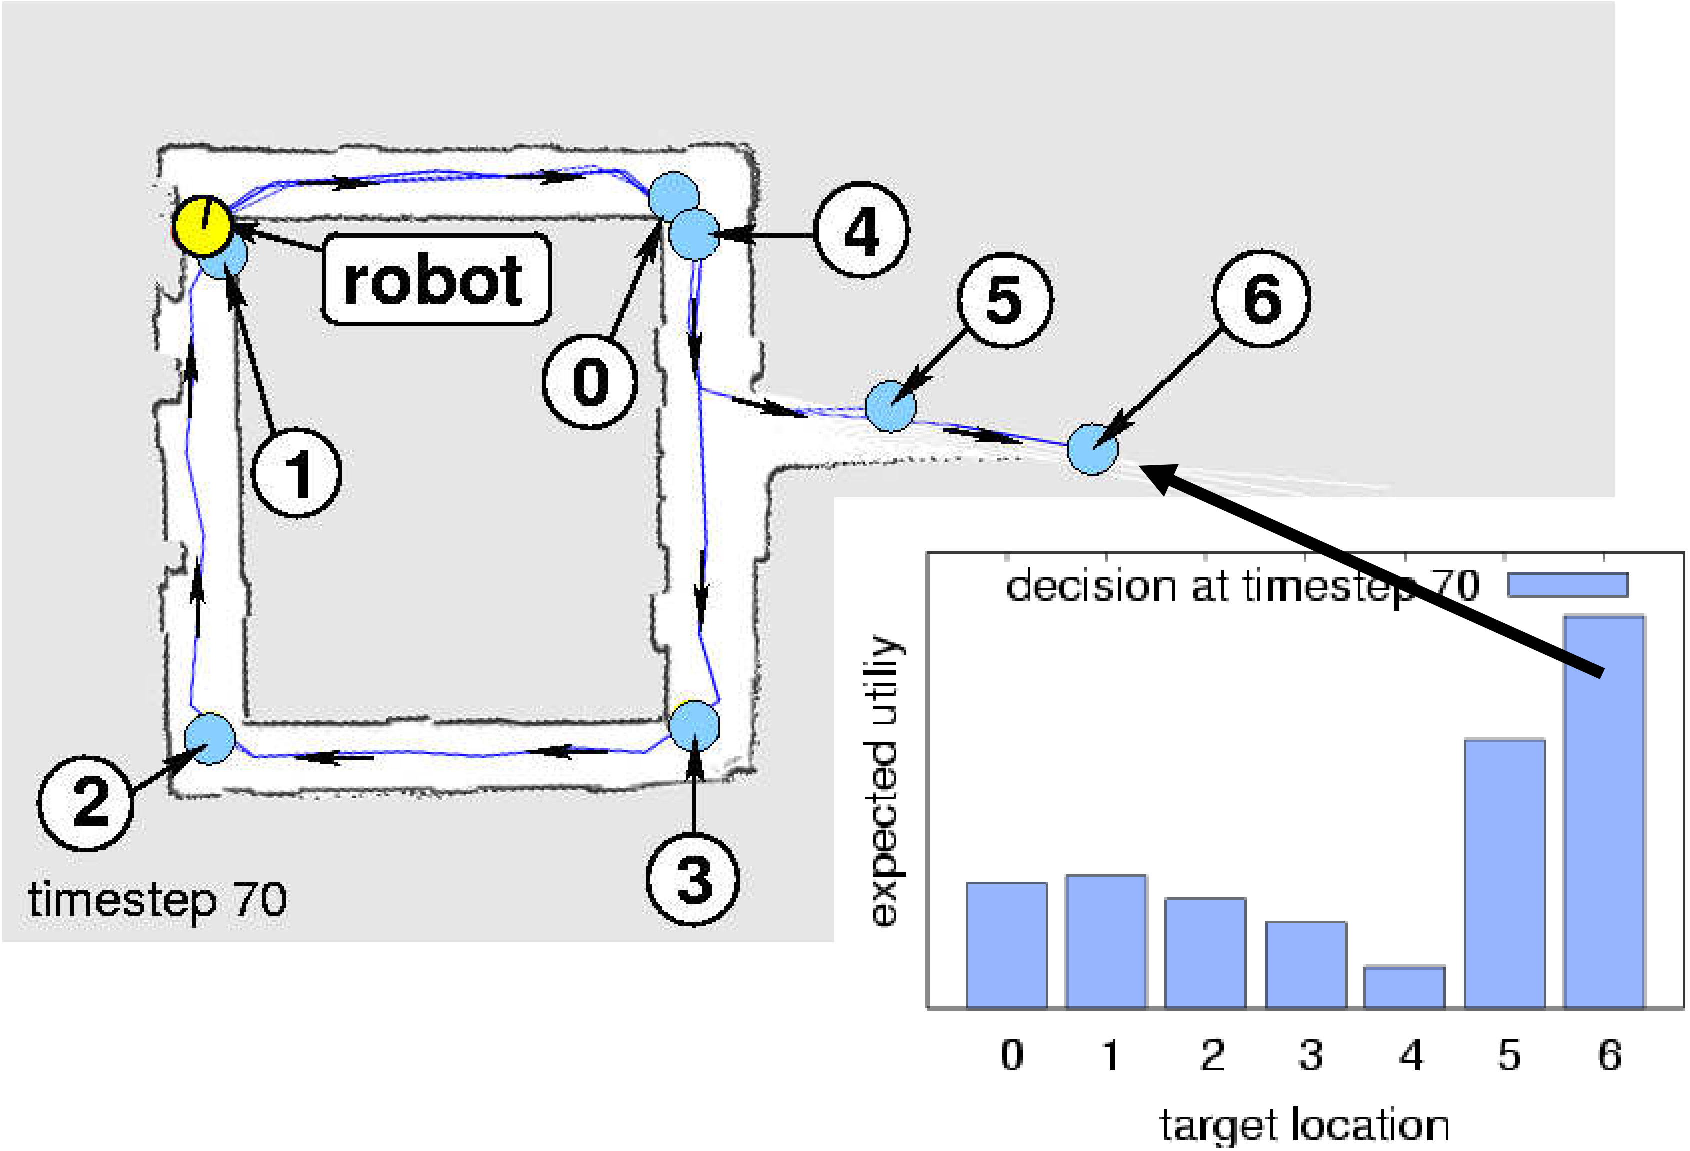

Example: Move Robot

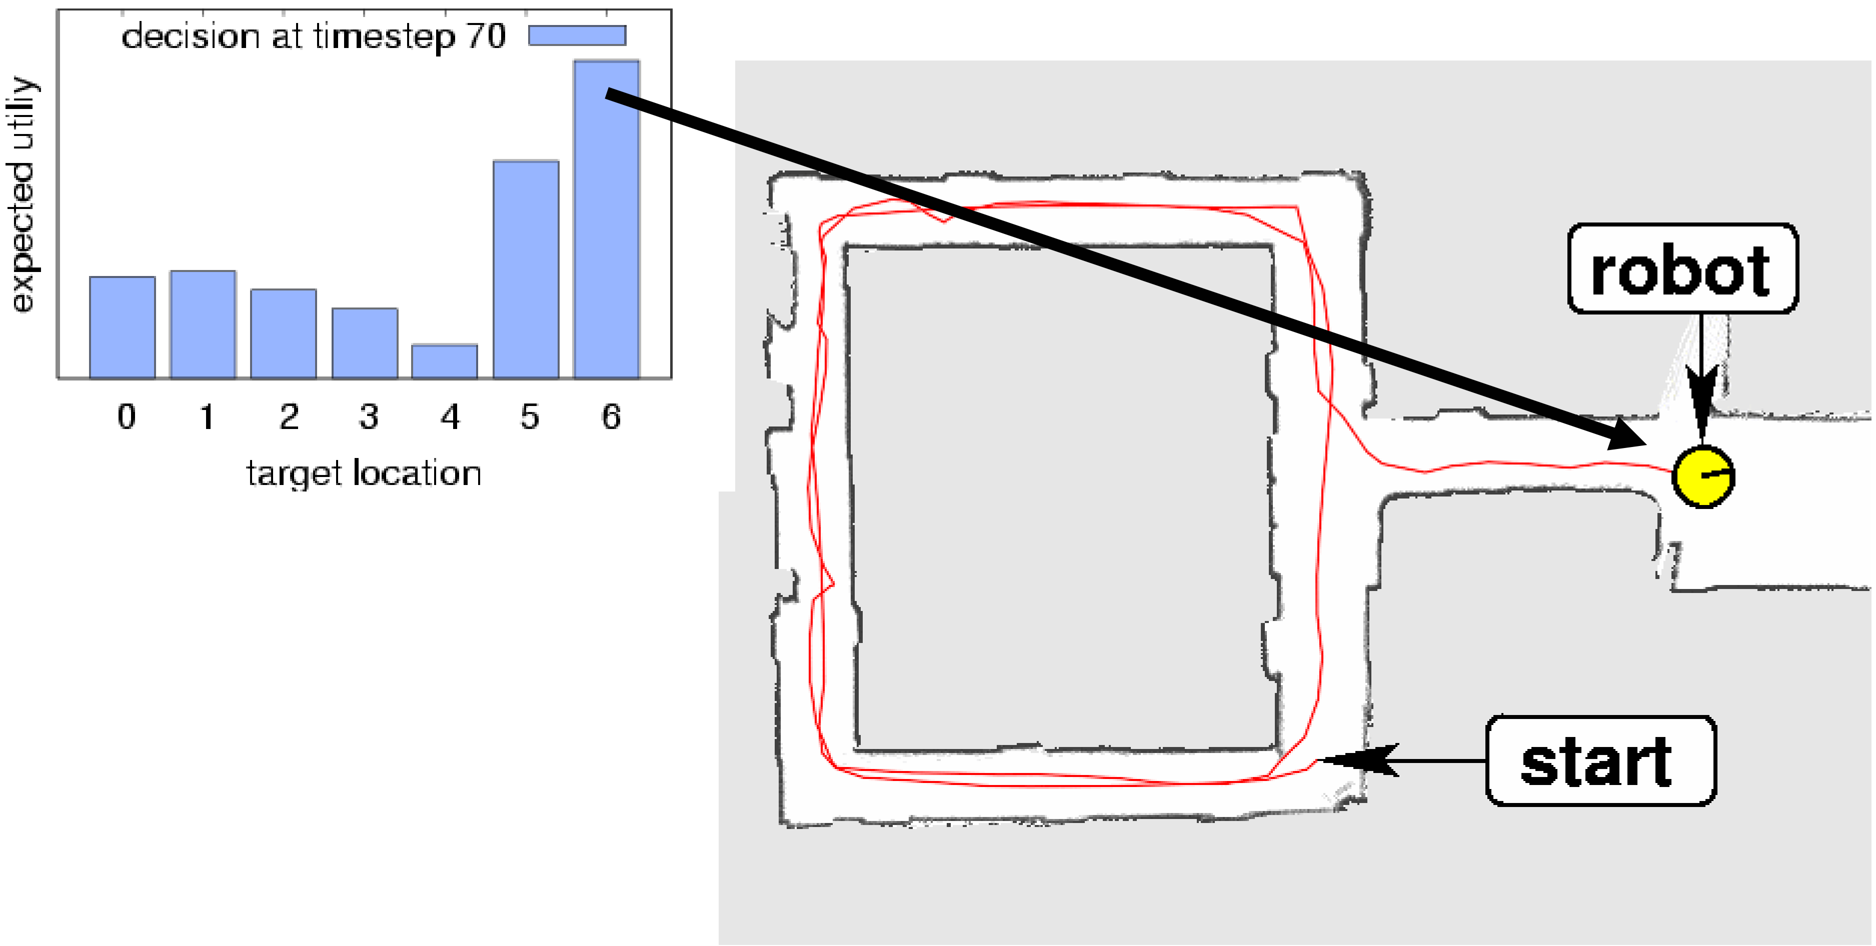

... continue ...

Example: Entropy Evolution

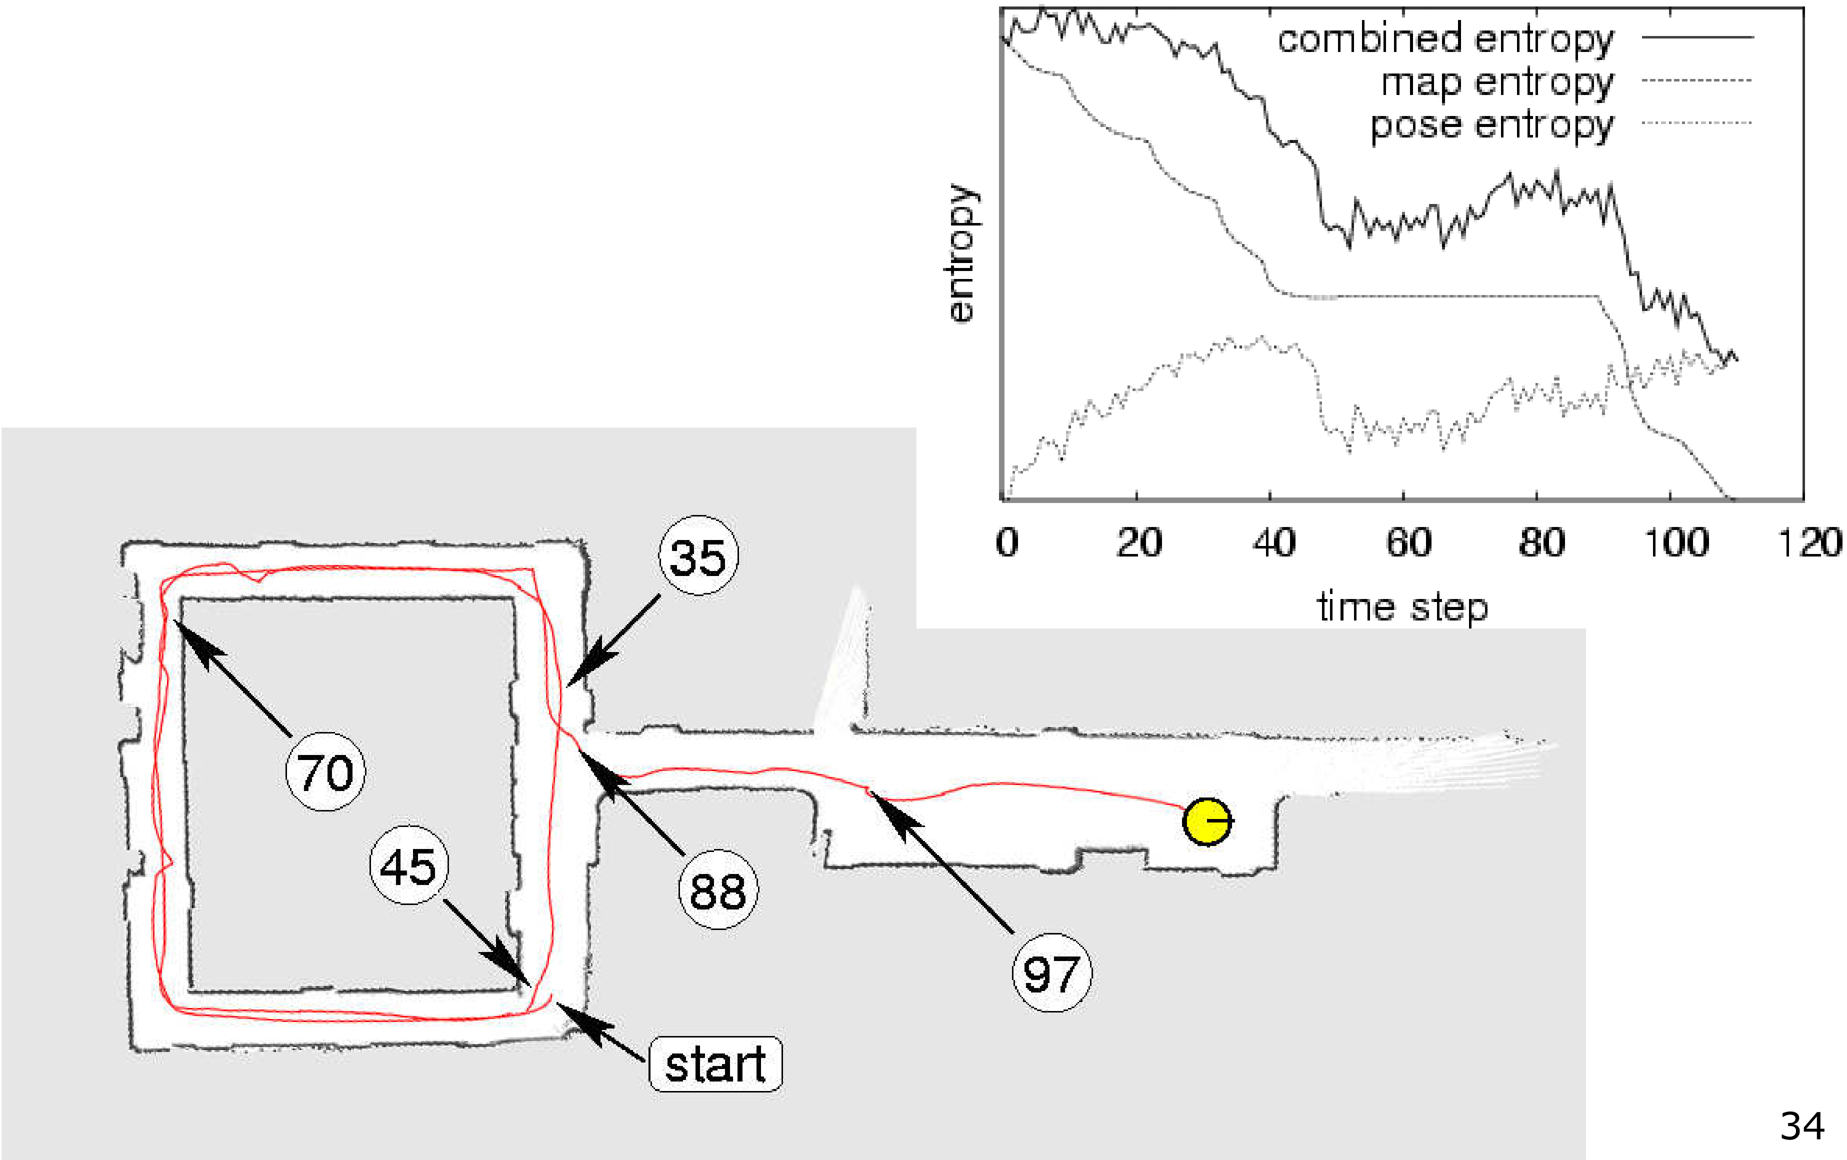

Comparison

Map uncertainty only:

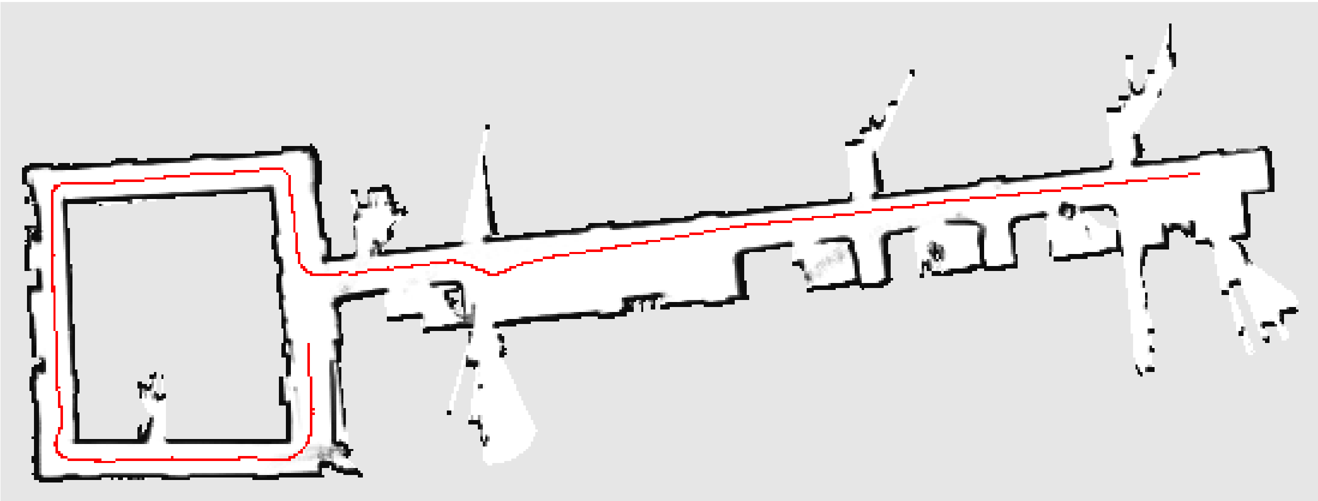

After loop closing action:

Real Exploration Example

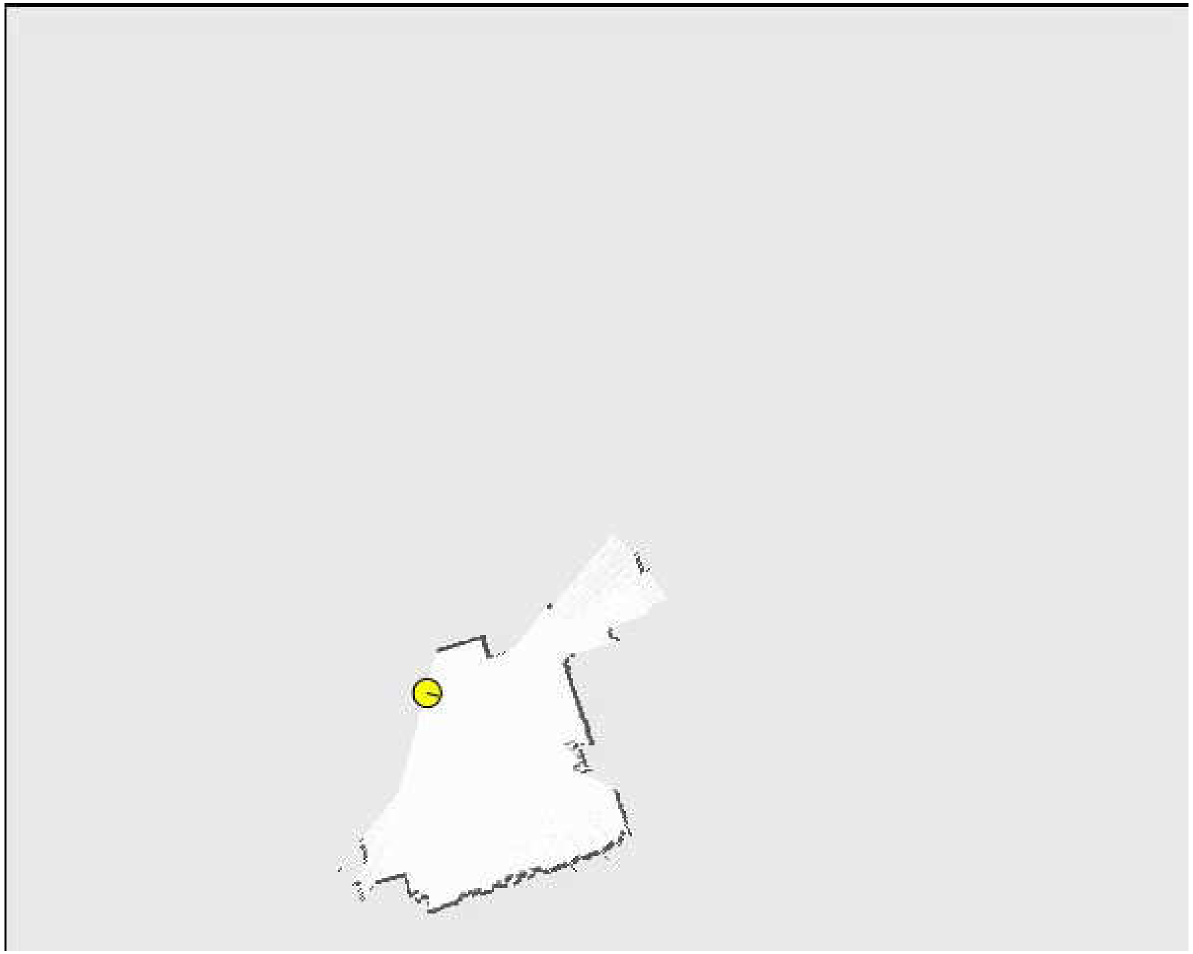

Selected target location 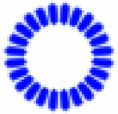

Corridor Exploration

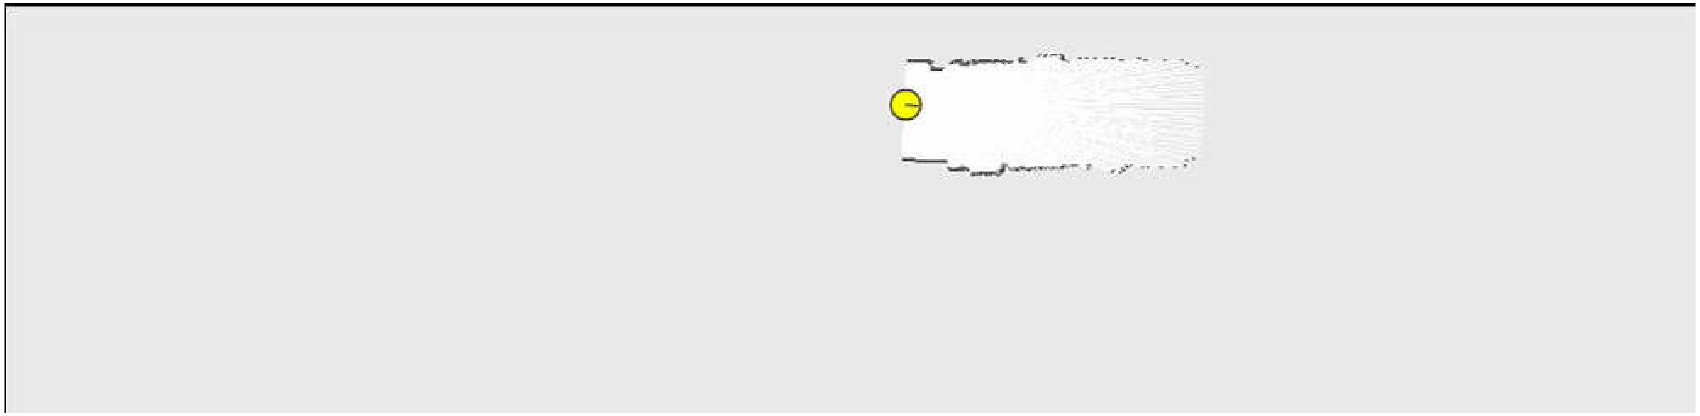

- The decision-theoretic approach leads to **intuitive behaviors:** “re-localize before getting lost”
- Some animals show a similar behavior (dogs marooned in the tundra of north Russia)

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
- Experimental results demonstrate the usefulness of the overall approach