

Global A-Optimal Robot Exploration in SLAM

Robert Sim and Nick Roy

Presented by Geoff Hollinger

CS599

Spring, 2011

SLAM

- Simultaneous localization and mapping (SLAM)
 - Build environmental models or maps from sensor data collected from a mobile robot
- Essentially a “chicken and egg” problem
 - Can’t map without localizing, can’t localize without map
 - Need to be estimated simultaneously
- Active SLAM
 - Ordering of measurements makes big difference

Problem formulation

- Quasi-holonomic robot in planar environment
- Translational and rotational displacement, observations

$$U_t = (\Delta d, \Delta \theta) \quad z_t = (r_1, b_1, r_2, b_2, \dots, r_n, b_n)$$

- Full state space

$$\xi = (x, y, \theta, x_1, y_1, x_2, y_2, \dots, x_n, y_n)$$

- Want to estimate

$$p(\xi | z_t, u_t, z_{t-1}, u_{t-1}, \dots, z_0, u_0)$$

- Gaussian assumption

$$p(\xi) = N(\mu, \Psi)$$

Entropy minimization

- Entropy of a distribution $H(p(\xi)) = \int_{\Xi} p(\xi) \log p(\xi)$

- Information gain of an action

$$\Delta I(a) = (H(p(\xi|a)) - H(p(\xi)))$$

- Expected information gain

$$p(\xi) = k \exp \left\{ -\frac{1}{2} (\Delta\xi)^T \Psi^{-1} (\Delta\xi) \right\} \quad (6)$$

- $\Rightarrow H(p(\xi)) = \frac{d}{2} (1 + \log 2\pi) + \frac{1}{2} \log \det(\Psi) \quad (7)$

$$\Rightarrow \Delta I_{t+1|t} \propto E_z [\log \det(\Psi_{t+1})] - \log \det(\Psi_t) \quad (8)$$

where $k = (2\pi)^{-d/2} \det(\Psi)^{-1/2}$ and $\Delta\xi = \xi - \mu$.

Entropy minimization

- For Gaussian:

$$p(\xi) = k \exp \left\{ -\frac{1}{2} (\Delta\xi)^T \Psi^{-1} (\Delta\xi) \right\} \quad (6)$$

$$\Rightarrow H(p(\xi)) = \frac{d}{2} (1 + \log 2\pi) + \frac{1}{2} \log \det(\Psi) \quad (7)$$

$$\Rightarrow \Delta I_{t+1|t} \propto E_z [\log \det(\Psi_{t+1})] - \log \det(\Psi_t) \quad (8)$$

where $k = (2\pi)^{-d/2} \det(\Psi)^{-1/2}$ and $\Delta\xi = \xi - \mu$.

- Maximizing information gain is equivalent to minimizing the determinant of the variance of the posterior!

A-Optimal vs D-Optimal

- D-Optimal: minimize the volume of the hyperellipsoid

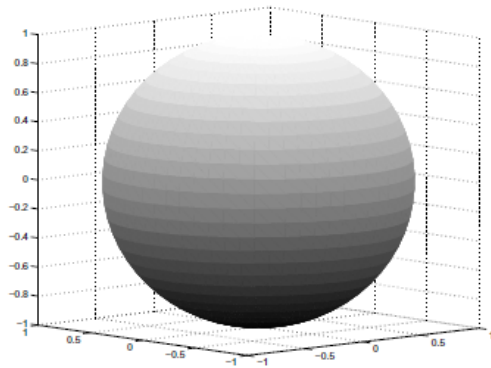
$$\det(\Psi_t) = \prod_{i=1}^n \lambda_i$$

- A-Optimal: maximize the likelihood of the data (mean-squared error)

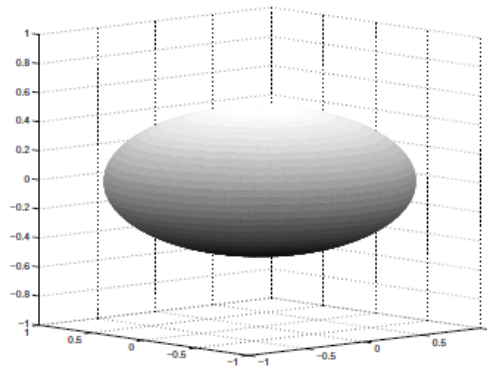
$$\begin{aligned} I(\Psi_t) &= \sum_{i=1}^n \lambda_i = \text{tr}(\Psi_t) \\ \Rightarrow \Delta I_{t+1|t} &= E_z [\text{tr}(\Psi_{t+1})] - \text{tr}(\Psi_t) \end{aligned}$$

Geometric interpretation

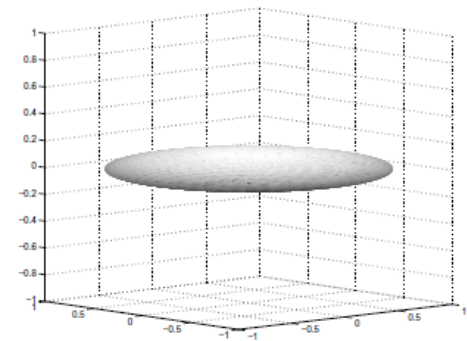
- D-optimal: (related to) volume of hyperellipsoid
- A-optimal: (related to) mean-squared error of the model
 - Possible issue: different scaling of parameters



(a) All eigenvalues at 1
 $\det \Psi = 1$
 $\text{tr}(\Psi) = 3$

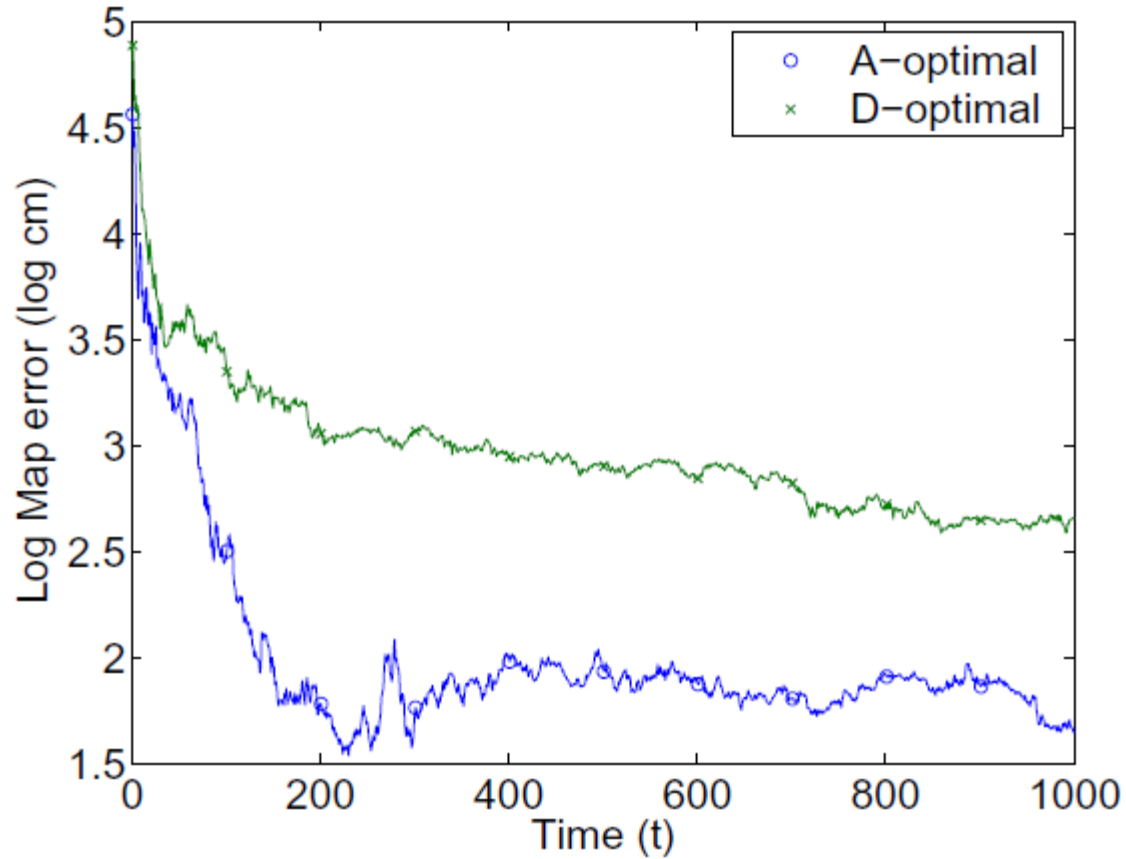


(b) One eigenvalue at 0.5
 $\det \Psi = 0.5$
 $\text{tr}(\Psi) = 2.5$



(c) One eigenvalue at 0.1
 $\det \Psi = 0.1$
 $\text{tr}(\Psi) = 2.1$

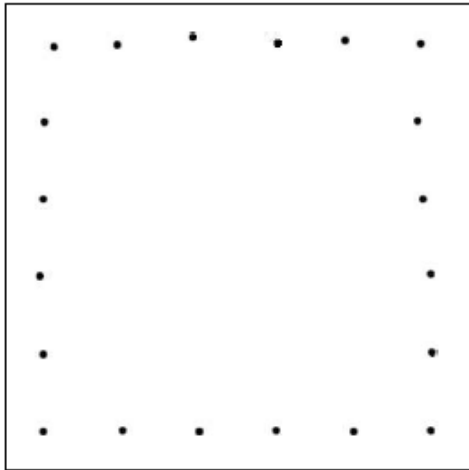
A-Optimal vs. D-Optimal



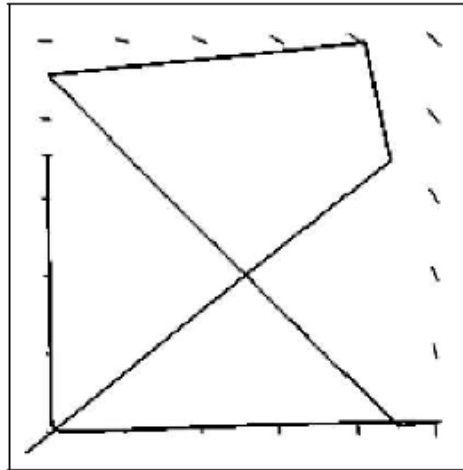
Planning approaches

- Greedy $\pi(\xi, \Psi) = \operatorname{argmin}_a E[\operatorname{tr}(\Psi)]$
- Global trajectory (POMDP, oh no!!)

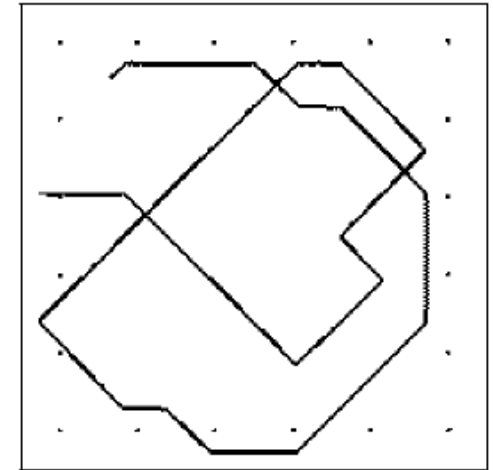
$$\pi(\xi, \Psi) = \operatorname{argmin}_{a_0, a_1, \dots, a_t} \operatorname{tr}(\Psi_t), \quad \pi(\xi, \Psi) = \operatorname{argmin}_{(\xi_0, \Psi_0), (\xi_1, \Psi_1), \dots, (\xi_t, \Psi_t)} \operatorname{tr}(\Psi_t).$$



(a) Initial Map



(b) Greedy Trajectory



(c) Global Trajectory

Pruned breadth first search

- Compact representation of state:
 - Discretize robot position into grid (similar to coastal nav)

$$(\xi, \Psi) \approx (x, y, \text{tr}(\Psi))$$

- Prune repeated states with worse information
 - If $x = x'$ and $y = y'$ and $E[\text{tr}(I)] < E[\text{tr}(I')]$, prune (x', y')
- Represents a 2D “information surface” over the x, y position of the robot
- Add a parent pointer to each state (x, y)
- Use priority queue to determine next state to update

Algorithm

- 1) Initialize all $I(x, y) = \infty$
- 2) Push current $\{x, y, \xi, \Psi\}$ onto Q , with priority $p = \text{tr}(\Psi)$
- 3) While Q not empty
 - a) Pop $\{x, y, \xi, \Psi\}$
 - b) For each neighbor (x', y') of (x, y) :
 - i) Compute ΔI from equation (11)
 - ii) If $I(x, y) + \Delta I < I(x', y')$ and $(x', y') \notin \{\phi(x, y), \phi(\phi(x, y)), \dots, (x_r, y_r)\}$ then
 - A) $I(x', y') = I(x, y) + \Delta I$
 - B) Push $\{x', y', \xi', \Psi'\}$ onto Q with priority $p = I(x', y')$
 - C) Set $\phi(x', y') = (x, y)$
- 4) $(x, y) = \text{argmin}_{(x, y)} I(x, y)$
- 5) while $(x, y) \neq (x_r, y_r)$
 - a) $(x', y') = (x, y)$
 - b) $(x, y) = \phi(x, y)$
- 6) Move to (x', y')

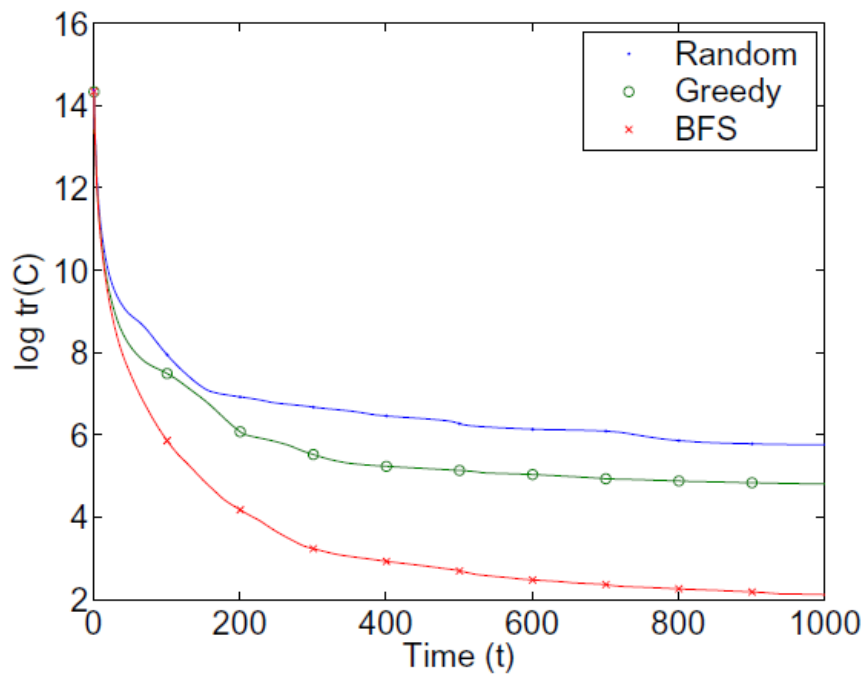
Caveats

- Requires:
 - Initial state estimate of features
 - Gaussian assumption to compute expected IG
- When to replan?
- Convergence guaranteed by monotonicity
- $O(s^2)$ state updates and $O(m)$ iterations
 - $O(s^3)$ if re-visiting is disallowed

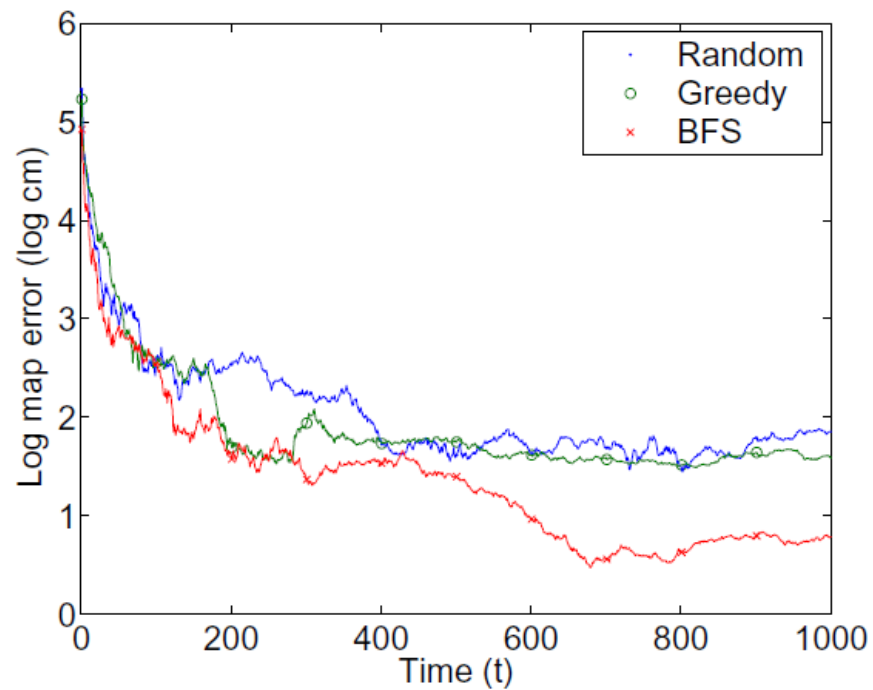
Experimental setup

- Random exploration
- Greedy exploration: robot drives to a single location that is maximally informative
- Breadth-first search: algorithm from previous slide
- 20 map features distributed randomly in 200m x 200m environment with robot that moves in 1m steps
- Infinite range sensor
 - Range variance proportional to distance

Simulations



(a) Trace of Posterior Covariance



(b) Map Accuracy

Extensions they mention

- When to re-plan?
- Using Sparse Information Filters for cases of viewing a subset of measurements
- Methods using scans of range data or images
- Continuous locations and non-looping trajectories

Additional related work

- ***J.P. Gonzalez, Planning with Uncertainty in Position Using High-Resolution Maps***, doctoral dissertation, tech. report CMU-RI-TR-08-02, Robotics Institute, Carnegie Mellon University, May, 2008
 - Slightly different problem: plan a path that best localizes a robot in a GPS-denied environment