Planning Under Uncertainty in the Continuous Domain: a Generalized Belief Space Approach

Vadim Indelman, Luca Carlone and Frank Dellaert

Institute of Robotics and Intelligent Machines (IRIM)
Georgia Institute of Technology

Georgia Institute of Technology

Robotics@GT & Intelligent Machines

International Conference on Robotics and Automation (ICRA), June 2014

Introduction

- Key components for autonomous operation include
  - **Perception**: Where am I? What is the surrounding environment?
  - **Planning**: What to do next?

Localization and mapping given robot motion

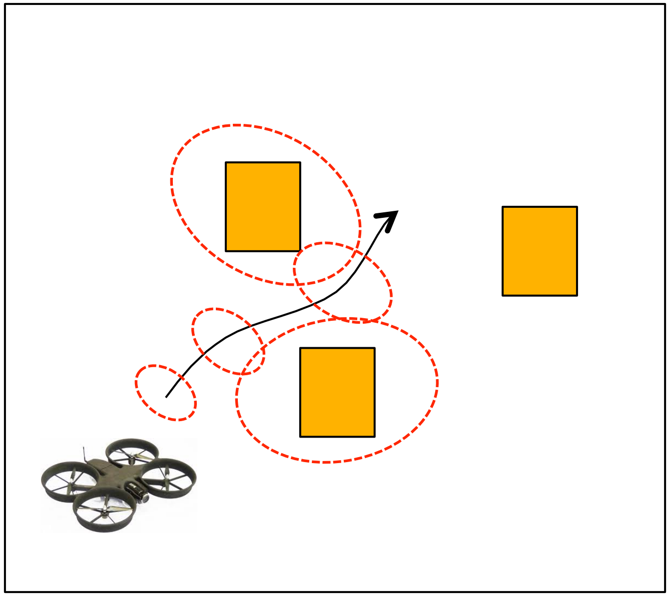

Coupled problems

Planning (e.g. reach a goal)

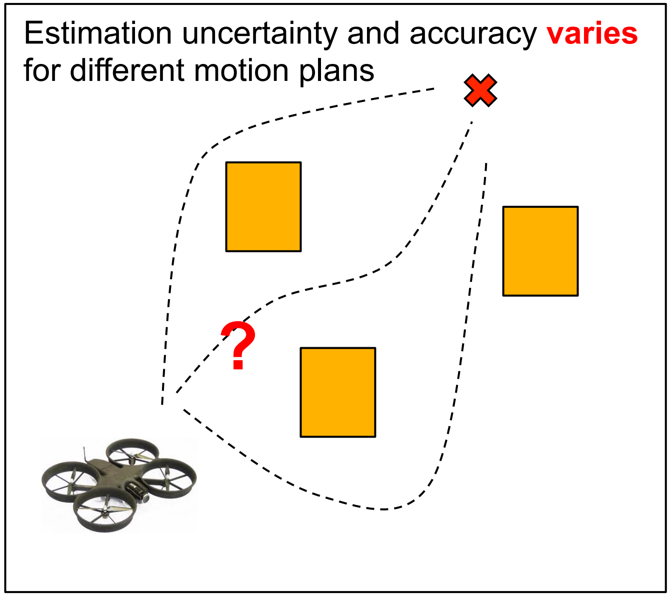

Introduction – Motivating Example

- **Autonomous navigation** to different goals in unknown environment

Not accounting for uncertainty in planning

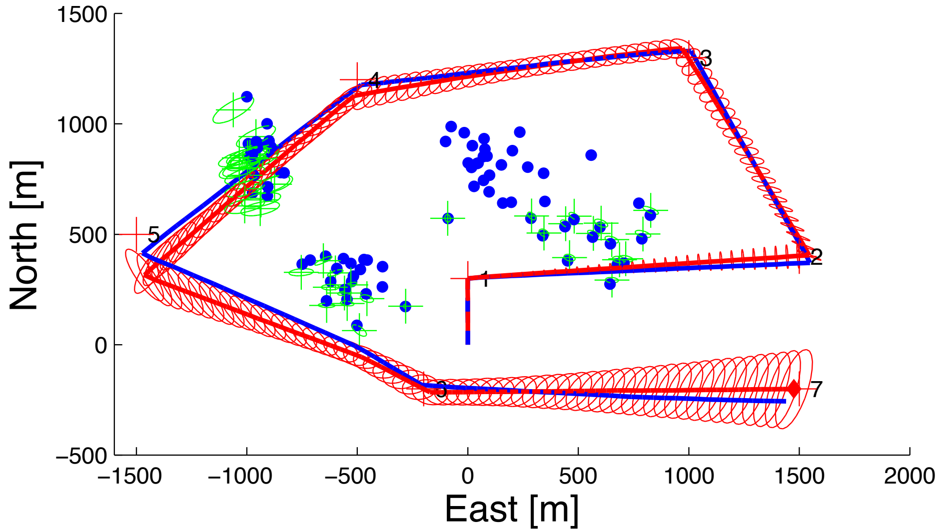

Accounting for uncertainty in planning

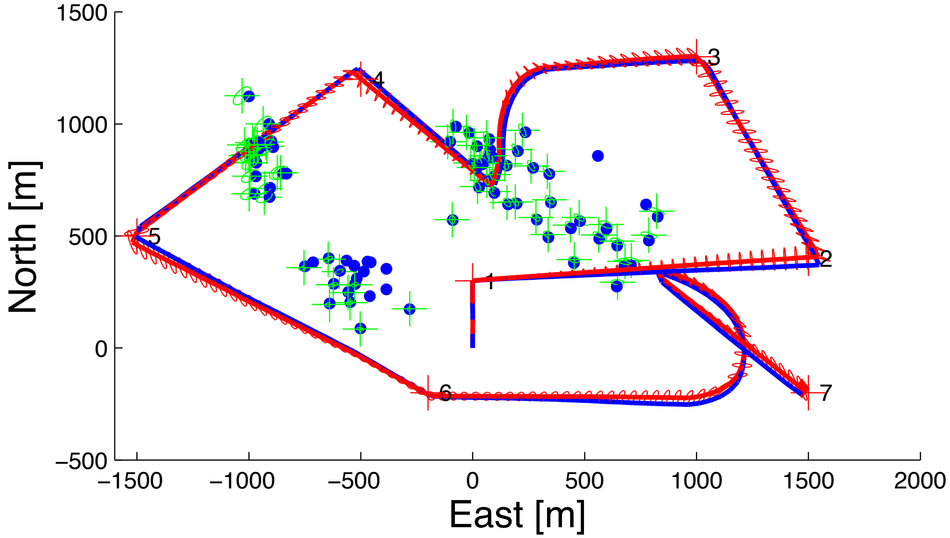

Georgia Institute of Technology

Robotics@GT & Intelligent Machines

Related Work

- Existing approaches often
  - Assume known environment (e.g. map)
    [Prentice and Roy 2009], [Van den Berg et al. 2012], [Hollinger et al. 2013]
  - Discretize state and control space - performance depends on grid resolution
    [Stachniss et al. 2004], [Bryson and Sukkarieh 2008], [Valencia et al. 2012], [Kim and Eustice 2013]
  - Assume maximum likelihood observations
    [Miller et al. 2009], [Platt et al. 2010], [Patil et al. 2014]

- Planning in the continuous domain - Generalized Belief Space (GBS)
  - Probabilistic description of the **robot** and the **environment** states
  - Direct trajectory optimization approach (provides locally-optimal trajectories)

- Environment is unknown/uncertain
- Maximum likelihood observations assumption is avoided
- Model the probability of acquiring a future observation (extends [Indelman et al. 2013])

Notations and Probabilistic Formulation

- **Joint** state vector $X_k \doteq \{\underbrace{x_0, \ldots, x_k}_{\text{Past \& current robot states}}, \underbrace{L_k}_{\text{Mapped environment}}\}$

- **Joint** probability distribution function $p\left(X_k|\mathcal{Z}_k, \mathcal{U}_{k-1}\right)$

$$p\left(X_k|\mathcal{Z}_k, \mathcal{U}_{k-1}\right) = priors \cdot \prod_{i=1}^{k} p\left(x_i|x_{i-1}, u_{i-1}\right) \underbrace{p\left(z_i|X_i^o\right)}_{\text{General observation model } X_i^o \subseteq X_i}$$

- Computationally-efficient maximum a posteriori inference e.g. [Kaess et al. 2012]

$$p\left(X_k|\mathcal{Z}_k, \mathcal{U}_{k-1}\right) \sim N\left(X_k^*, \Sigma_k\right)$$

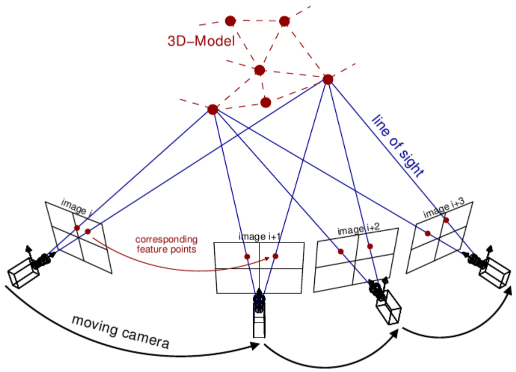

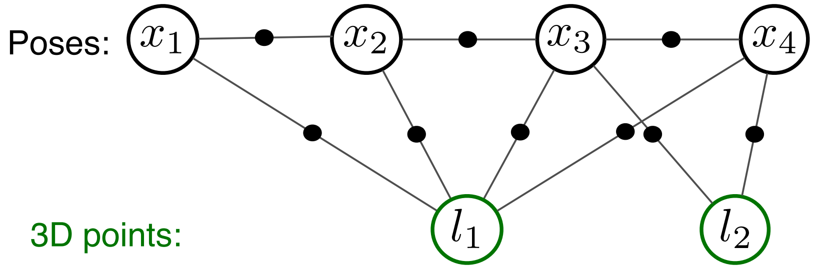

Planning in the Generalized Belief Space

- Plan (locally) optimal control sequence over L look-ahead steps: $u^*_{k:k+L-1}$
  - By minimizing an objective function
  - Operating over the **generalized belief**
  - Model predictive control framework

X_k : Joint state at time t_k
$\mathcal{U}_{k-1}$: All past controls

- **What is the generalized belief?**
  - Probabilistic description of the **robot** <u>and</u> the **environment** states
  - Generalized belief at planning time t_k : $gb(X_k) \doteq p\left(X_k | \mathcal{Z}_k, \mathcal{U}_{k-1}\right) \sim N\left(X^*_k, \Sigma_k\right)$

Known (from perception)

Generalized belief at planning time = joint pdf

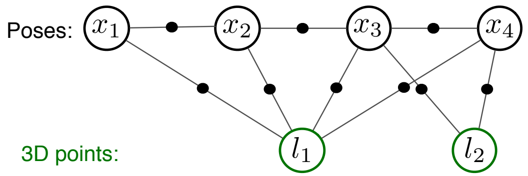

Planning in the Generalized Belief Space

- Generalized belief at the l-th look-ahead step
  - Describes the joint pdf (robot and environment states) at that time

$$gb\left(X_{k+l}\right) \doteq p\left(X_{k+l} | \mathcal{Z}_k, \mathcal{U}_{k-1}, Z_{k+1:k+l}, u_{k:k+l-1}\right)$$

Joint state at the *l*-th look ahead step

Past controls & measurements

Controls & measurements at the first *l* look-ahead steps

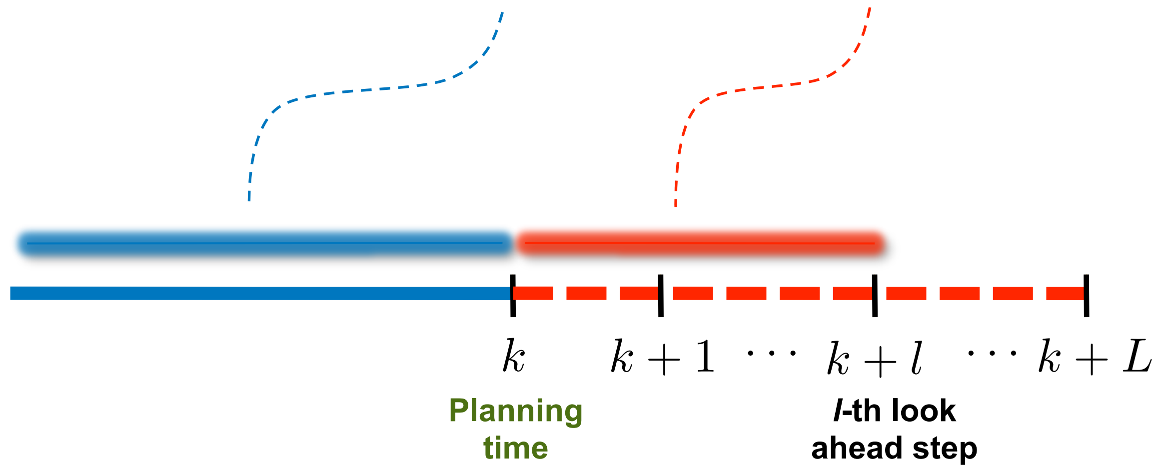

Planning in the Generalized Belief Space

- Generalized belief at the l-th look-ahead step
  - Describes the joint pdf (robot and environment states) at that time

$$gb\left(X_{k+l}\right) \doteq p\left(\underbrace{X_{k+l}}_{\text{Joint state at the } l\text{-th look ahead step}} \middle| \underbrace{\mathcal{Z}_k, \mathcal{U}_{k-1}}_{\text{Past controls \& measurements}}, \underbrace{Z_{k+1:k+l}, u_{k:k+l-1}}_{\text{Controls \& measurements at the first } l \text{ look-ahead steps}}\right)$$

- Objective function can now involve **uncertainty** (e.g. covariance) in robot and environment states

$$J_k\left(u_{k:k+L-1}\right) \doteq \underset{Z_{k+1:k+L}}{\mathbb{E}} \left\{ \sum_{l=0}^{L-1} \underline{c_l} \left(gb\left(X_{k+l}\right), u_{k+l}\right) + \underline{c_L}\left(gb\left(X_{k+L}\right)\right) \right\}$$

- For example, plan motion to minimize uncertainty in robot state

Generalized Belief Space (Cont.)

$$gb\left(X_{k+l}\right) \doteq p\left(X_{k+l} | \underbrace{\mathcal{Z}_k, \mathcal{U}_{k-1}}_{\text{Past}}, \underbrace{Z_{k+1:k+l}, u_{k:k+l-1}}_{\text{Future}}\right)$$

- Modeling **future** observations $Z_{k+1:k+l}$:
  - Treated as random variables [Van den Berg et al. 2012]
  - Will a future observation be actually acquired?
    - Depends on the *true* state at *future* time
    - e.g. is a 3D point within sensing range
  - Model probabilistically acquisition of $Z_{k+1:k+l}$ by random binary variables $\Gamma_{k+1:k+l}$
- Marginalize out $\Gamma_{k+1:k+l}$ to get $gb\left(X_{k+l}\right)$

$$gb\left(X_{k+l}\right) = \sum_{\Gamma_{k+1:k+l}} p\left(X_{k+l}, \Gamma_{k+1:k+l} | \mathcal{Z}_k, \mathcal{U}_{k-1}, Z_{k+1:k+l}, u_{k:k+l-1}\right)$$

- Expensive! Instead – Expectation Maximization (details soon)

Planning in the Generalized Belief Space

$$J_k\left(u_{k:k+L-1}\right) \doteq \underset{Z_{k+1:k+L}}{\mathbb{E}} \left\{ \sum_{l=0}^{L-1} c_l\left(gb\left(X_{k+l}\right), u_{k+l}\right) + c_L\left(gb\left(X_{k+L}\right)\right) \right\}$$

- How to calculate locally-optimal control?

$$u^{\star}_{k:k+L-1} = \underset{u_{k:k+L-1}}{\arg\min}\, J_k\left(u_{k:k+L-1}\right)$$

Dual-layer iterative optimization

Outer layer - Inference over $u_{k:k+L-1}$

Starting from initial guess $u^{(0)}_{k:k+L-1}$

At **each** iteration:

- Compute $\Delta u_{k:k+L-1}$
- Update control

$$u^{(i+1)}_{k:k+L-1} = u^{(i)}_{k:k+L-1} + \Delta u_{k:k+L-1}$$

Inner layer - Inference over the **belief**

For **each** look-ahead step, for a **given** control

$$gb(X_{k+l}) \sim N\left(X^*_{k+l}, \Sigma_{k+l}\right)$$

- As a function of random variables $Z_{k+1:k+l}$
- EM formulation to avoid marginalizing over the latent variables $\Gamma_{k+1:k+l}$

Outer Layer: Inference over the Control

Iterative optimization over the nonlinear objective function $J_k\left(u_{k:k+L-1}\right)$

$$J_k\left(u_{k:k+L-1}\right) \doteq \underset{Z_{k+1:k+L}}{\mathbb{E}}\left\{\sum_{l=0}^{L-1} c_l\left(gb\left(X_{k+l}\right), u_{k+l}\right) + c_L\left(gb\left(X_{k+L}\right)\right)\right\}$$

- Involves:
  - Calculating gradient ∇J_k
  - Evaluating objective function J_k for different control values

Outer layer - Inference over $u_{k:k+L-1}$

Starting from initial guess $u_{k:k+L-1}^{(0)}$

At **each** iteration:

- Compute $\Delta u_{k:k+L-1}$
- Update control

$$u_{k:k+L-1}^{(i+1)} = u_{k:k+L-1}^{(i)} + \Delta u_{k:k+L-1}$$

Inner layer - Inference over the **belief**

For **each** look-ahead step, for a **given** control

$$gb(X_{k+l}) \sim N\left(X_{k+l}^{*}, \Sigma_{k+l}\right)$$

- As a function of random variables $Z_{k+1:k+l}$
- EM formulation to avoid marginalizing over the latent variables $\Gamma_{k+1:k+l}$

Inner Layer: Inference Over the Belief

Given **current** controls $u_{k:k+L-1}$, for each look ahead step l :

- Compute the Gaussian approximation X^*_{k+l}, Σ_{k+l} such that

$$gb(X_{k+l}) \sim N\left(X^*_{k+l}, \Sigma_{k+l}\right)$$

- EM formulation:

$$X^\star_{k+l} = \underset{X_{k+l}}{\arg\min} \underset{\Gamma_{k+1:k+l}|\bar{X}_{k+l}}{\mathbb{E}} \left[-\log p\left(X_{k+l}, \Gamma_{k+1:k+l} | \mathcal{Z}_k, \mathcal{U}_{k-1}, Z_{k+1:k+l}, u_{k:k+l-1}\right)\right]$$

- Gauss Newton method

Outer layer - Inference over $u_{k:k+L-1}$

Starting from initial guess $u^{(0)}_{k:k+L-1}$

At **each** iteration:

- Compute $\Delta u_{k:k+L-1}$
- Update control

$$u^{(i+1)}_{k:k+L-1} = u^{(i)}_{k:k+L-1} + \Delta u_{k:k+L-1}$$

Inner layer - Inference over the **belief**

For **each** look-ahead step, for a **given** control

$$gb(X_{k+l}) \sim N\left(X^*_{k+l}, \Sigma_{k+l}\right)$$

- As a function of random variables $Z_{k+1:k+l}$
- EM formulation to avoid marginalizing over the latent variables $\Gamma_{k+1:k+l}$

Inner Layer: Inference Over the Belief

Given **current** controls $u_{k:k+L-1}$, for each look ahead step l :

- Compute the Gaussian approximation X^*_{k+l}, Σ_{k+l} such that

$$gb(X_{k+l}) \sim N\left(X^*_{k+l}, \Sigma_{k+l}\right)$$

- EM formulation:

$$X^\star_{k+l} = \underset{X_{k+l}}{\arg\min} \underset{\Gamma_{k+1:k+l}|\bar{X}_{k+l}}{\mathbb{E}} \left[-\log p\left(X_{k+l}, \Gamma_{k+1:k+l} | \mathcal{Z}_k, \mathcal{U}_{k-1}, Z_{k+1:k+l}, u_{k:k+l-1}\right)\right]$$

- Gauss Newton method

- Next - we show the above formulation:
  - Guides the robot towards **informative** distant 3D points (**outside** sensing range)
    - Loop closures to reduce uncertainty
  - Alternative formulation using signed distance function [Patil et al. 2014]

Limited Sensing Range

- Illustrative (toy) example
- Without modeling probability of acquiring future observations:
  - 3D points outside sensing range contribute **zero** gradient to ∇J_k
  - Robot will not be guided to re-observe these points (i.e. no loop closures)

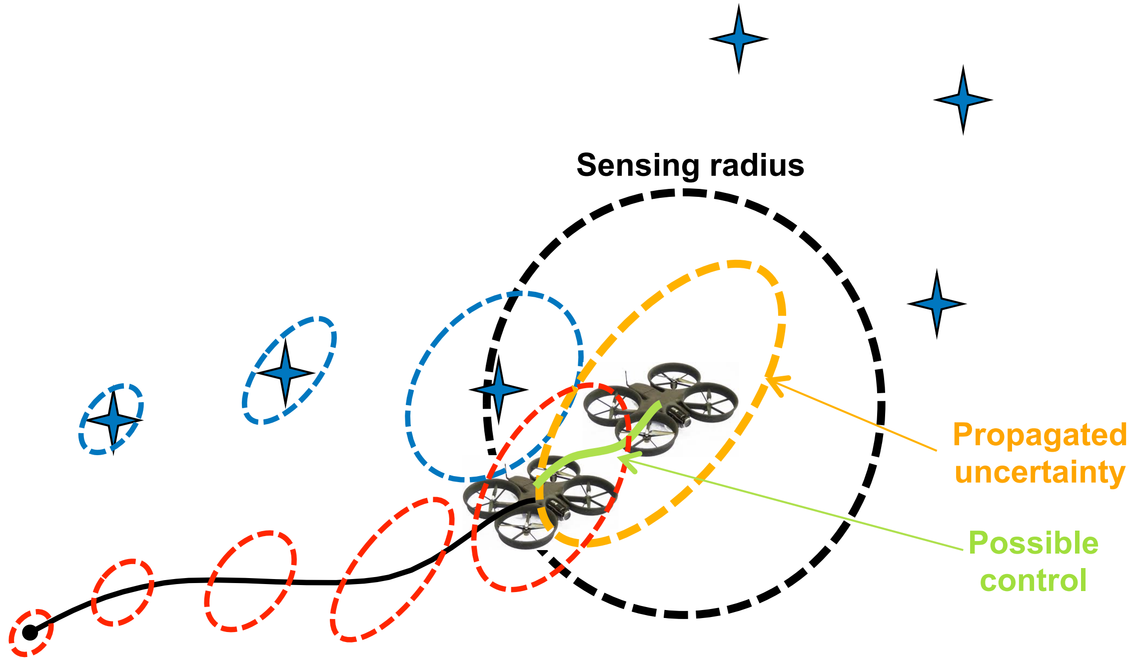

… Back to Inner Layer

$$X^{\star}_{k+l} = \underset{X_{k+l}}{\arg\min} \underset{\Gamma_{k+1:k+l}|\bar{X}_{k+l}}{\mathbb{E}} \left[-\log p\left(X_{k+l}, \Gamma_{k+1:k+l} | \mathcal{Z}_k, \mathcal{U}_{k-1}, Z_{k+1:k+l}, u_{k:k+l-1}\right)\right]$$

- Joint pdf:

$$p(X_k|\mathcal{Z}_k, \mathcal{U}_{k-1}) \prod_{i=1}^{l} p\left(x_{k+i}|x_{k+i-1}, u_{k+i-1}\right) p\left(Z_{k+i}, \Gamma_{k+i}|X^o_{k+i}\right)$$

- Inference over the belief:

$$X^{\star}_{k+l} = \underset{X_{k+l}}{\arg\min} \left\|X_k - X^{\star}_k\right\|^2_{I_k} + \sum_{i=1}^{l} \left\|x_{k+i} - f\left(x_{k+i-1}, u_{k+i-1}\right)\right\|^2_{\Omega_w}$$

$$+ \sum_{i=1}^{l} \sum_{j=1}^{n_i} \underline{p\left(\gamma_{k+i,j} = 1|\bar{X}_{k+l}\right)} \left\|z_{k+i,j} - h\left(X^o_{k+i,j}\right)\right\|^2_{\Omega_v^{ij}}$$

Probability of observing the j-th 3D point

… Back to Inner Layer

$$X^{\star}_{k+l} = \underset{X_{k+l}}{\arg\min} \underset{\Gamma_{k+1:k+l}|\bar{X}_{k+l}}{\mathbb{E}} \left[-\log p\left(X_{k+l}, \Gamma_{k+1:k+l} | \mathcal{Z}_k, \mathcal{U}_{k-1}, Z_{k+1:k+l}, u_{k:k+l-1}\right) \right]$$

- Joint pdf:

$$p(X_k|\mathcal{Z}_k, \mathcal{U}_{k-1}) \prod_{i=1}^{l} p\left(x_{k+i}|x_{k+i-1}, u_{k+i-1}\right) p\left(Z_{k+i}, \Gamma_{k+i}|X^o_{k+i}\right)$$

- Inference over the belief:

$$X^{\star}_{k+l} = \underset{X_{k+l}}{\arg\min} \left\| X_k - X^{\star}_k \right\|^2_{I_k} + \sum_{i=1}^{l} \left\| x_{k+i} - f\left(x_{k+i-1}, u_{k+i-1}\right) \right\|^2_{\Omega_w}$$

$$+ \sum_{i=1}^{l} \sum_{j=1}^{n_i} p\left(\gamma_{k+i,j} = 1 | \bar{X}_{k+l}\right) \left\| z_{k+i,j} - h\left(X^o_{k+i,j}\right) \right\|^2_{\Omega^{ij}_v}$$

- Equivalent to <u>weighting</u> the measurement covariance matrix

$$\bar{\Omega}^{ij}_v = p\left(\gamma_{k+i,j} = 1 | \bar{X}_{k+l}\right) \Omega^{ij}_v \qquad \left\| z_{k+i,j} - h\left(X^o_{k+i,j}\right) \right\|^2_{\bar{\Omega}^{ij}_v}$$

  – E.g., probability to observe 3D points decreases with distance
  – Contributes to ∇J_k, becomes dominant if information gain is substantial

Results

- **Autonomous navigation** to different goals in an unknown environment
  - **Objective function**: penalize **control usage, uncertainty and distance to goal**
  - No absolute information
  - Onboard sensors: camera and range sensor
  - Control: heading angle

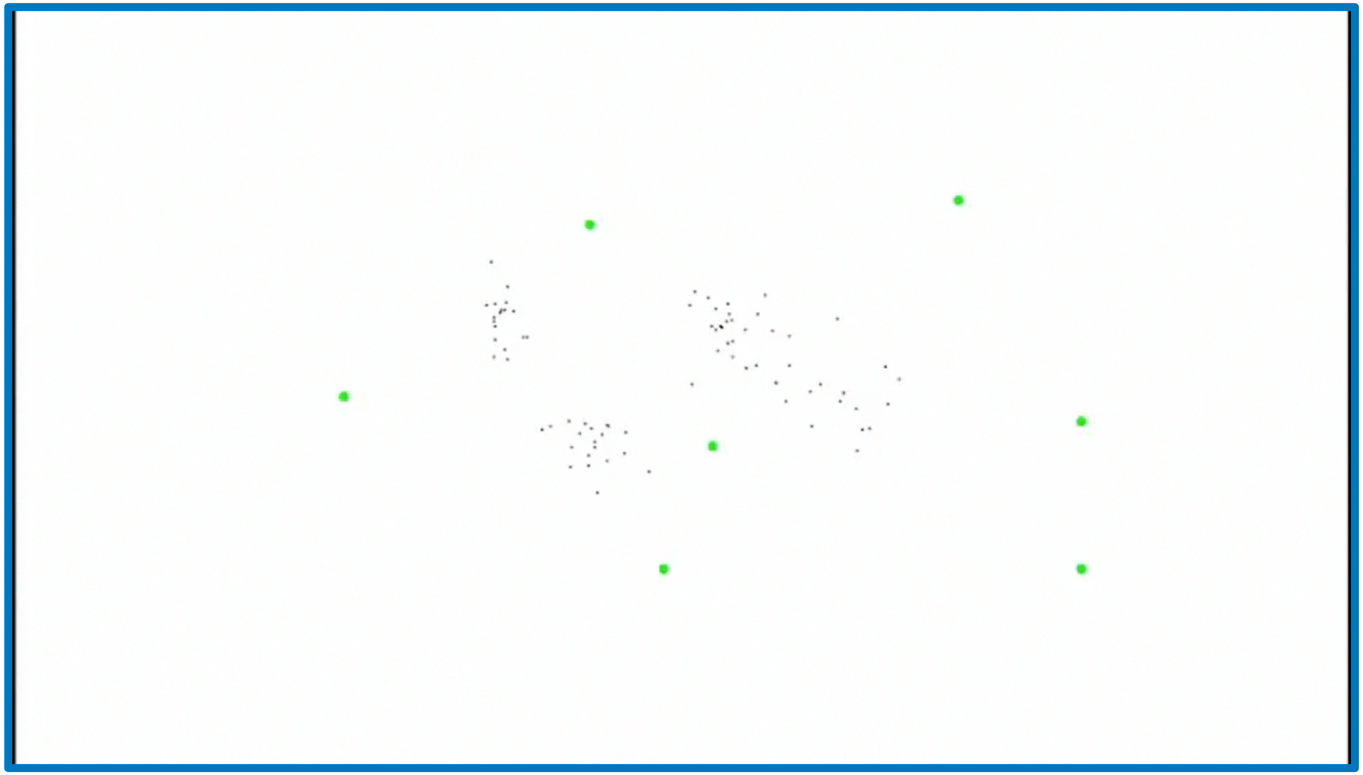

Results

- **Autonomous navigation** to different goals in an unknown environment
  - **Objective function**: penalize **control usage, uncertainty and distance to goal**
  - Compared methods:
    - Planning in GBS
    - Planning in GBS, no uncertainty
    - Discrete planning - A*, adaptation of [Kim and Eustice 2013]

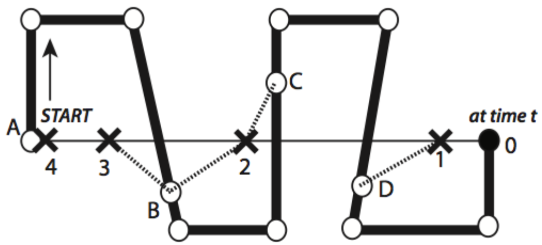

[Kim and Eustice 2013]

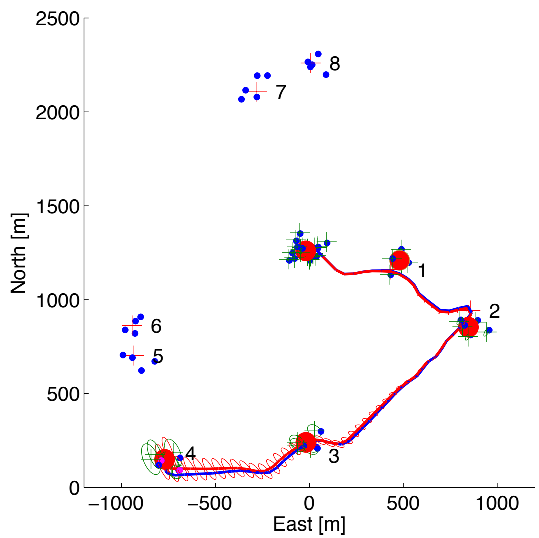

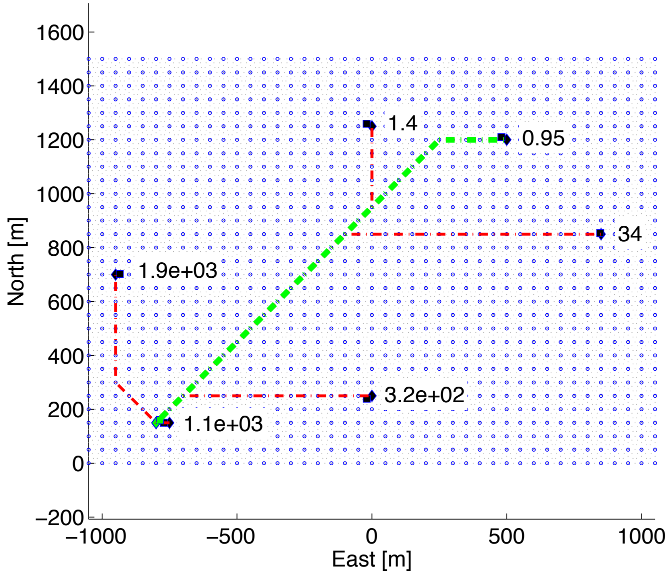

Discrete Planning (adaptation of [Kim and Eustice 2013])

Results

- **Autonomous navigation** to different goals in an unknown environment
  - **Objective function**: penalize **control usage, uncertainty and distance to goal**
  - Compared methods:
    - Planning in GBS
    - Planning in GBS, no uncertainty
    - Discrete planning - A*, adaptation of [Kim and Eustice 2013]

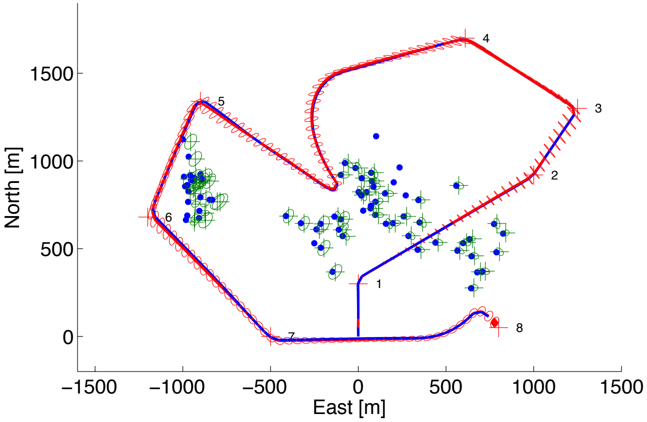

Planning in GBS

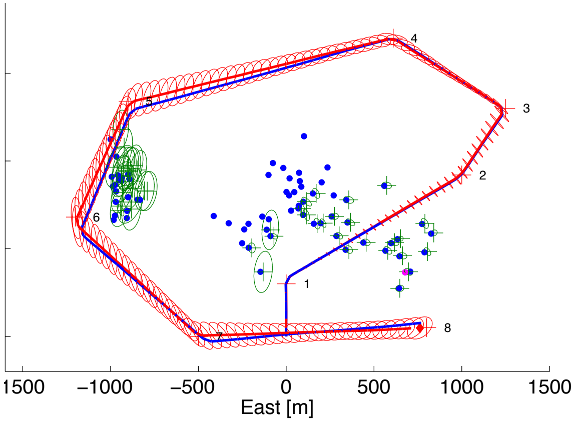

Planning in GBS, no uncertainty

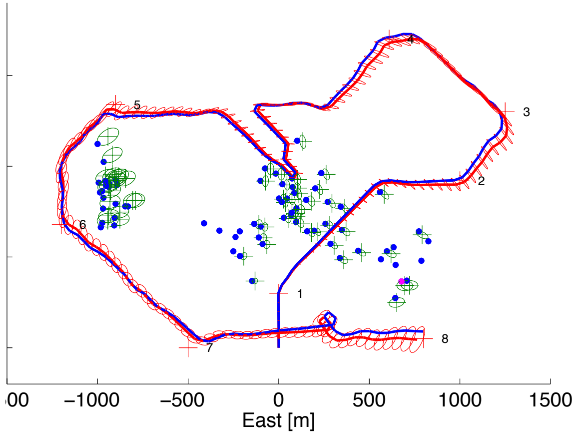

Discrete planning

Results

- **Autonomous navigation** to different goals in an unknown environment
  - **Objective function**: penalize **control usage, uncertainty and distance to goal**
  - Compared methods:
    - Planning in GBS
    - Planning in GBS, no uncertainty
    - Discrete planning - A*, adaptation of [Kim and Eustice 2013]

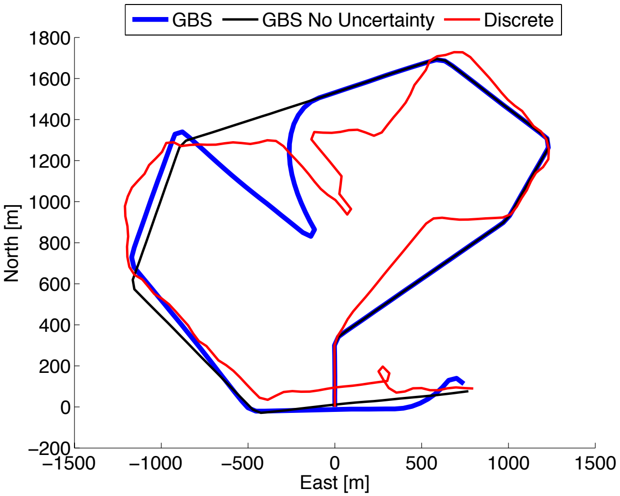

Trajectories

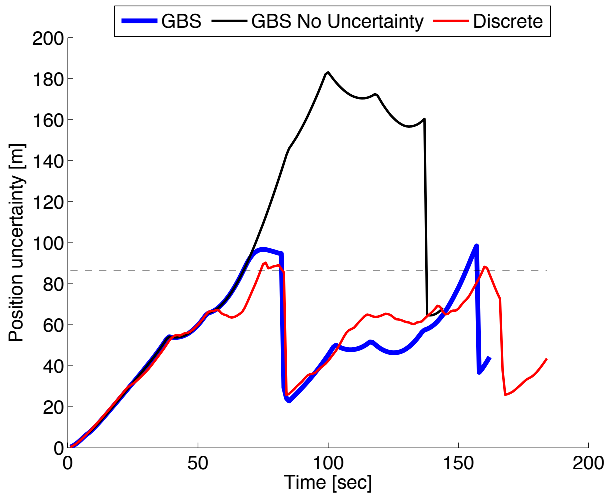

Uncertainty

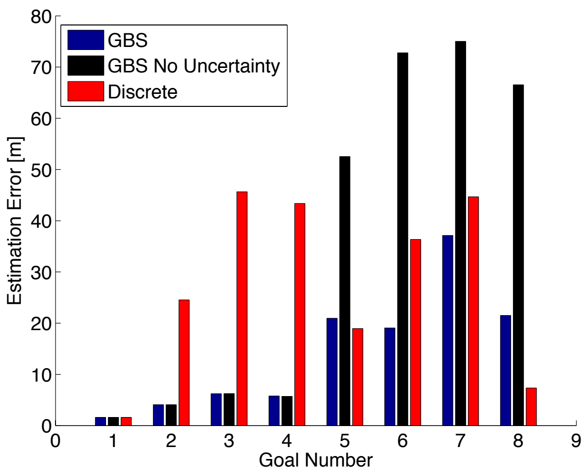

Est. errors @ goals

Conclusions

- **Planning in the continuous domain - Generalized Belief Space**
  - General framework for planning under uncertainty (incl. uncertainty in environment)
  - Addresses 3 limitations of state of the art:

Discretization of state or control space	✓
Maximum likelihood observations assumption	✓
Exact prior knowledge regarding environment	✓

  - Limited sensing range
    - Latent variables to model acquisition probability of future observations
    - Allow planning loop closures outside sensing range
  - Produces smooth trajectories with reduced control effort

Extras

Discrete Planning

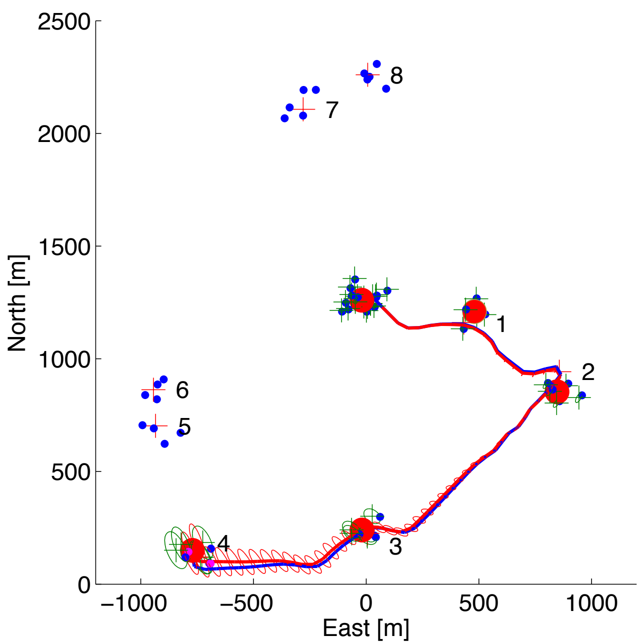

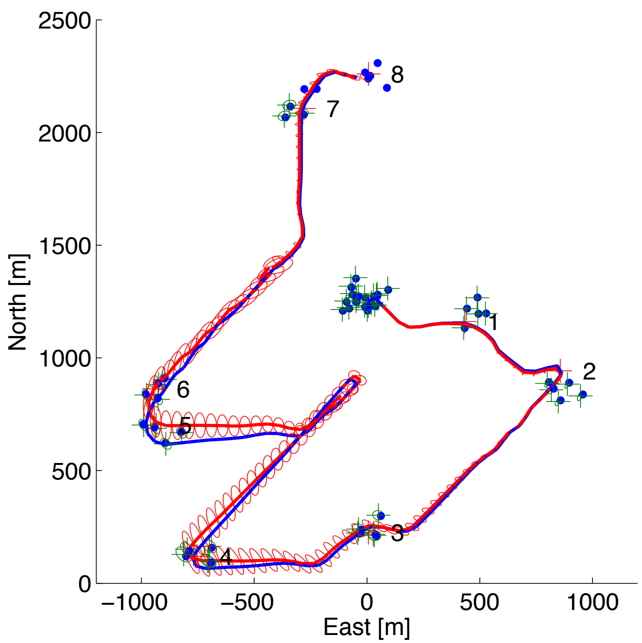

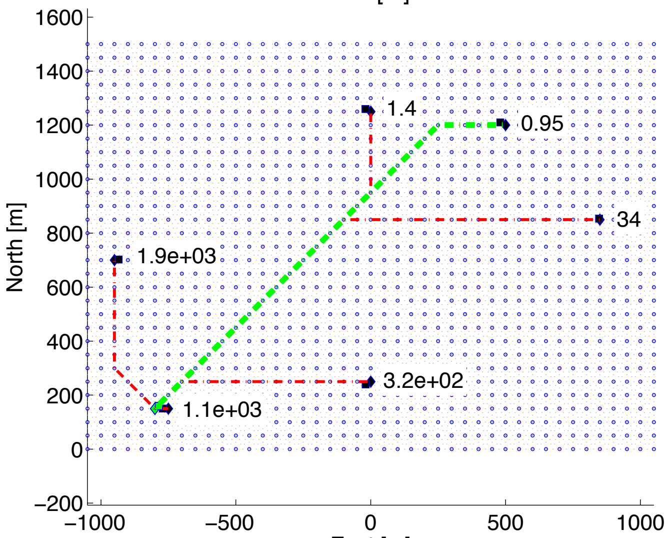

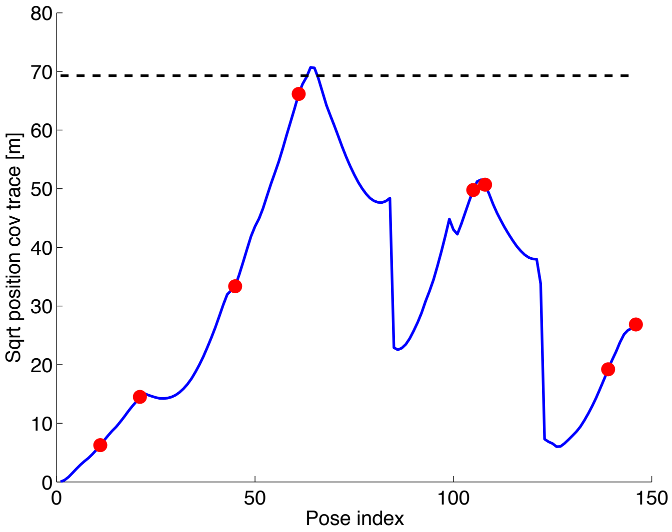

Comparison Between Different Methods