



A Bayesian framework for optimal motion planning with uncertainty

Andrea Censi, Daniele Calisi,
Giuseppe Oriolo, Alessandro De Luca



SAPIENZA
UNIVERSITÀ DI ROMA





Many formalizations, many approaches

- **Preimage-back chaining:** Lozano-Perez *et al.* (1984); Lazanas and Latombe (1992); Fraichard and Mermond (1998)
- **Sensor-based planning:** Bouilly *et al.* (1995); Khatib *et al.* (1997)
- **The Information Space approach:** Barraquand and Ferbach (1995); O’Kane and LaValle (2005); O’Kane (2006); O’Kane and LaValle (2006)
- **Sensor uncertainty fields (SUF):** Takeda and Latombe (1992); Takeda *et al.* (1994); Trahanias and Komninos (1996); Vlassis and Tsanakas (1998); Makarenko *et al.* (2002)
- **Set-membership approach:** Page and Sanderson (1995b,a)
- **Dynamic programming:** Blackmore *et al.* (2006); Blackmore (2006)
- **A^* , RRT:** Lambert and Fort-Piat (2000); Lambert and Gruyer (2003); Gonzalez and Stentz (2007)

Dimensions

- **How to represent uncertainty?**
 - Uncertainty is a bounded set.
 - Probabilistic ([isotropic] covariances, compressed information space, ...)
- **How does the uncertainty accumulate?** Bayesian, linearly with distance, ...
- **How does the uncertainty shrink?** Bayesian, “reset” to zero...
- **Which problem to solve?**
 - find a safe path, minimizing the execution time
 - find a safe path, minimizing the final covariance
 - maximize the collected information, with free final pose, ...
- **How to represent the plan/policy?**

Our approach – overview

- We work in the space poses \times covariances.
 - Already used in Lambert and Gruyer (2003)
 - We are more careful with assumptions.
 - We define transitions independently of localization algorithm.
- We consider two problems: minimize final time and final covariance.
- We develop two algorithms:
 - *forward*: A^* -like with propagation of states
 - *backward*: backprojection of constraints from target to goal
- Emphasis on **exploiting problem structure** with generic search framework based on **dominance relations**.

Motion planning **with uncertainty**

Find a continuous function $q^*(t)$ such that:

$$q^*(0) = q_{\text{start}} \qquad q(0) \sim p_0(q)$$

$$q^*(t) \in \mathcal{C}_{\text{free}} \qquad \mathbb{P}(q(t) \in \mathcal{C}_{\text{free}}) \geq 1 - \epsilon$$

$$q^*(t_f) \in \mathcal{C}_{\text{target}} \qquad \mathbb{P}(q(t_f) \in \mathcal{C}_{\text{target}}) \geq 1 - \epsilon$$

(+ kinematic/dynamic constraints)

(+ model for robot/sensors)

$$\min J(q^*, t_f)$$

$$\min E\{J(q^*, t_f)\}$$

- In general, the solution is a function from the space of probability distribution of the state to the space of actions.

Approach: PP with uncertainty \simeq PP in the pose \times covariance space

- We reduce the problem to deterministic planning in the space $\mathcal{S} = \text{pose} \times \text{covariance}$:

$$\mathbf{q}(0) \sim p_0(\mathbf{q}) \qquad \mathbf{s}_0 = \langle \mathbf{q}_0, \mathbf{\Sigma}_0 \rangle$$

$$\mathbb{P}(\mathbf{q}(t) \in \mathcal{C}_{\text{free}}) \geq 1 - \epsilon \qquad \mathbf{s}_t \in \mathcal{S}_{\text{free}}$$

$$\mathbb{P}(\mathbf{q}(t_f) \in \mathcal{C}_{\text{target}}) \geq 1 - \epsilon \qquad \mathbf{s}_{t_f} \in \mathcal{S}_{\text{target}}$$

- $\mathcal{S}_{\text{free}}$ e $\mathcal{S}_{\text{target}}$ are defined using bounds on the covariances:

$$\mathbf{s}_t \in \mathcal{S}_{\text{free}} \iff \mathbf{q}_t \in \mathcal{C}_{\text{free}} \quad \wedge \quad \mathbf{\Sigma}_t \leq \text{CONSTRAINTS}(\mathbf{q}_t)$$

$$\mathbf{s}_t \in \mathcal{S}_{\text{target}} \iff \mathbf{q}_t \in \mathcal{C}_{\text{target}} \quad \wedge \quad \mathbf{\Sigma}_t \leq \mathbf{M}$$

- The set $\text{CONSTRAINTS}(\mathbf{q}_t)$ depends on the geometry of the environment.

Evolution of uncertainty

- Let Σ_u be the odometry error, and $\mathcal{I}(q)$ the Fisher information matrix. Then for the covariance of the estimate:

$$\Sigma_k \star \left(\mathcal{I}(q_k) + (\Sigma_{k-1} + \Sigma_u)^{-1} \right)^{-1}$$

(note: simplified formula) where \star is:

- “=” in the linear case.
 - “ \geq ” is the Bayesian Cramér-Rao bound for unbiased estimators.
 - “ \simeq ” in practice, at least for range-finders (see ICRA’07 paper)
- Semi-formal assumptions:
 - The distribution is \simeq Gaussian during the optimal motion.
 - The localization algorithm is unbiased and \simeq efficient.
 - The uncertainty of the pose is small with respect to the complexity of the environment: $\mathcal{I}(q) \simeq \mathcal{I}(\hat{q})$.

Problems considered

We study two problems:

- Minimizing the final time.
- Minimizing the final covariance (with a bound on the time).


$$\min_{\leq} \Sigma(t_f) \quad \text{subject to } t_f \leq t_{\max}$$

Lots of differences with respect to standard motion planning:

- There is, in general, a continuity of solutions.
- Solutions are not reversible.
- Because sensors have specific frequency, time is important, not merely a parameterization.
- Much of the complexity comes from the fact that \leq is not a total order for covariances.

Planning by searching

- The generic search algorithm has two relations:
 - a partial order \succeq used for dominance (discarding nodes)
 - a total order \blacktriangleleft used for precedence (search direction)



- 1: Put n_0 in OPEN.
- 2: **while** OPEN is not empty **do**
- 3: Pop first (according to \blacktriangleleft) node n from OPEN.
- 4: **for all** s in $SUCCESSORS(n)$ **do**
- 5: Report success if $IS_GOAL(s)$.
- 6: Ignore s if it is \succeq -dominated in VISITED.
- 7: Discard nodes in VISITED \succeq -dominated by s .
- 8: Put s in VISITED.
- 9: Discard nodes in OPEN \succeq -dominated by s .
- 10: Put s in OPEN.
- 11: **end for**
- 12: **end while**
- 13: Report failure.

Forward approach


- Nodes are tuples $n = \langle q, \Sigma, t \rangle$:

“Can go from q_{start} to q in time t with final covariance Σ .”

- Search starts from initial pose: $n_0 = \langle q_{goal}, \Sigma_0, 0 \rangle$.
- Most of the work is the definition of dominance relations ($n_1 \supseteq n_2$) for discarding nodes. Basic example (there are more powerful ones):

$$(n_1 \supseteq n_2) \Leftrightarrow (q_1 = q_2) \wedge (t_1 \leq t_2) \wedge (\Sigma_1 \leq \Sigma_2)$$

- Example of two nodes that are not comparable:




Backward approach

- Nodes are tuples $n = \langle q_k, \{\mathbf{M}_i\}, tg \rangle$:


“If in q_k , and $\Sigma_k \leq \{\mathbf{M}_1, \mathbf{M}_2, \dots\}$, then I can arrive to q_{goal} in time tg .”

- Search starts from final pose: $n_0 = \langle q_{goal}, \text{CONSTRAINTS}(q_{goal}), 0 \rangle$.
- Constraints are back-propagated.




- Dominance relations are really ugly to show.

Example: Fisher information matrix

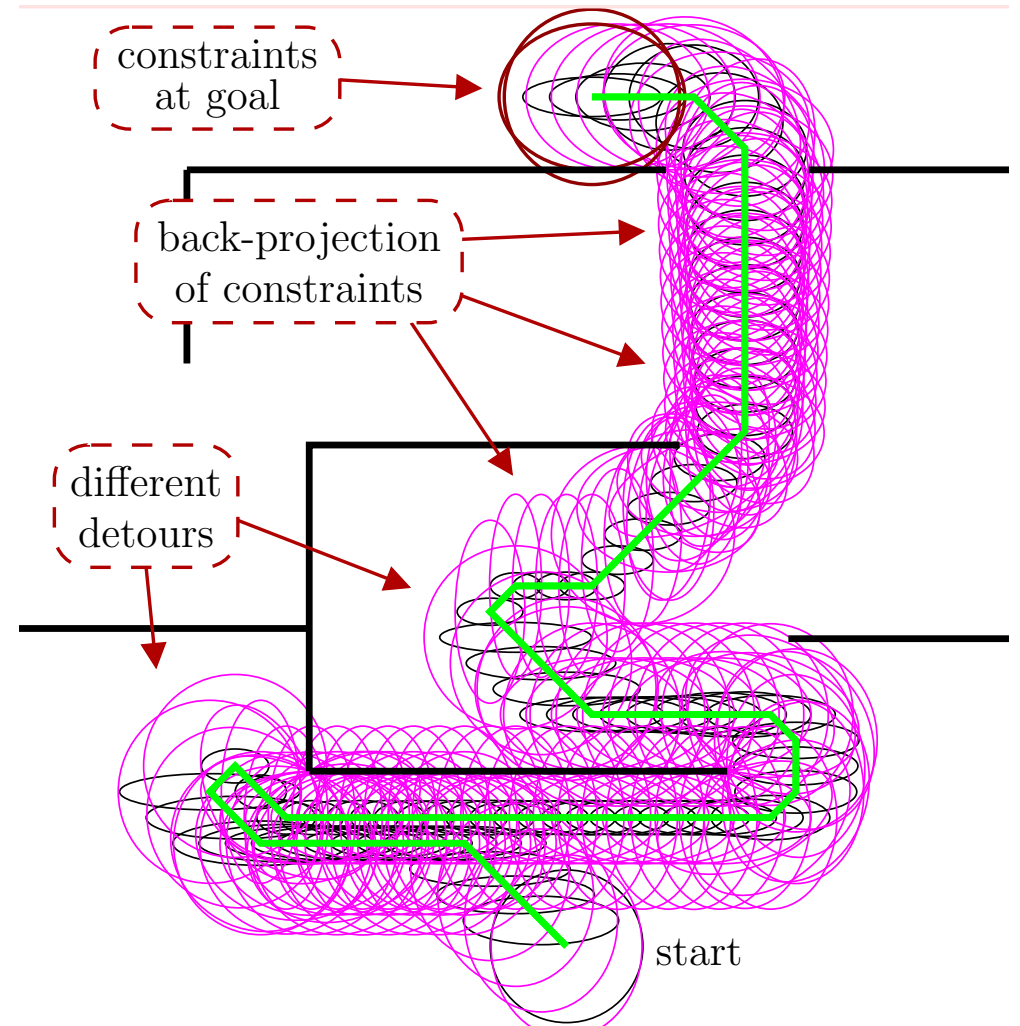


Example: minimum time path

forward algorithm




backward algorithm




Minimum time vs. minimum final covariance

minimum time



minimum final covariance



Conclusions (things we learned)

- On planning with uncertainty:
 - Nice formalization using Fisher information matrix.
 - Nice problem(s), with many peculiar properties.
 - Lots of structure that can be exploited.
- On solving through the two algorithms:
 - Dominance relations are great for analysis and implementation.
 - Two optimal algorithms giving different solutions.
- Bad ideas met along the way:
 - Do not discretize covariances! (lose correctness)
 - Representation matter (inverse of covariance might be better)
- Source code and cool animations available at my website.

Computational cost

A^* :	forward	backward	no uncertainty
created nodes	5'474	10'110	369
expanded nodes	4'211	7'321	229
nodes still active at the end	4'477	7'337	328
matrix comparisons	106'252	1'021'156	-
time - G4 1.5GHz	0.51	2.13	0.04
time - P4 2.8GHz	0.23	1.06	0.02

- Backward tree can be re-used if goal remains the same.

References

- Jérôme Barraquand and Pierre Ferbach. Motion planning with uncertainty: The Information Space approach. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, 1995.
- Lars Blackmore, Hui Li, and Brian Williams. A probabilistic approach to optimal robust path planning with obstacles. In *Proceedings of the AIAA Guidance, Navigation and Control Conference*, 2006.
- Lars Blackmore. A probabilistic particle control approach to optimal, robust predictive control. In *Proceedings of the AIAA Guidance, Navigation and Control Conference*, 2006.
- B. Bouilly, T. Simeon, R. Alami, and T. CNRS. A numerical technique for planning motion strategies of a mobile robot in presence of

uncertainty. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, volume 2, 1995.

T. Fraichard and R. Mermond. Path planning with uncertainty for car-like robots. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, pages 27–32, 1998.

Juan Pablo Gonzalez and Anthony Stentz. Planning with uncertainty in position using high-resolution maps. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, Rome, Italy, April 2007.

M. Khatib, B. Bouilly, T. Simeon, and R. Chatila. Indoor navigation with uncertainty using sensor-based motions. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, volume 4, pages 3379–3384, Albuquerque, NM, USA, April 1997.

Alain Lambert and Nadine Le Fort-Piat. Safe task planning integrating uncertainties and local maps federations. *International Journal of Robotics Research*, 19(6):597–611, Jun 2000.

A. Lambert and D. Gruyer. Safe path planning in an uncertain-configuration space. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, volume 3, pages 4185–4190, September 2003.

A. Lazanas and J.-C. Latombe. Landmark-based robot navigation. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*, pages 816–822, San Jose, California, 1992. AAAI Press.

Tomas Lozano-Perez, Matthew Mason, and Russell H. Taylor. Automatic synthesis of fine-motion strategies for robots. *International Journal of Robotics Research*, 3(1), 1984.

Alexei Makarenko, Stefan Williams, Frederic Bourgault, and Hugh Durrant-Whyte. An experiment in integrated exploration. In *Proceedings of the IEEE/RSJ International conference on Intelligent Robots and Systems (IROS)*, Lausanne, Switzerland, 2002.

Jason M. O’Kane and Steven M. LaValle. Almost-sensorless localization. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, 2005.

Jason M. O’Kane and Steven M. LaValle. Localization with limited sensing. Conditionally accepted to *IEEE Transactions on Robotics*, September 2006.

Jason M. O’Kane. Global localization using odometry. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, 2006.

L. A. Page and A. C. Sanderson. A path-space search algorithm for motion planning with uncertainties. In *Proceedings of the IEEE*

International Symposium on Assembly and Task Planning, pages 334–340, Pittsburgh, PA, USA, August 1995.

- L. A. Page and A. C. Sanderson. Robot motion planning for sensor-based control with uncertainties. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, volume 2, pages 1333–1340, Nagoya, Japan, May 1995.
- H. Takeda and J.-C. Latombe. Sensory Uncertainty Field for mobile robot navigation. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, pages 2465–2472, Nice, France, May 1992.
- Haruo Takeda, Claudio Facchinetti, and Jean-Claude Latombe. Planning the motions of a mobile robot in a Sensory Uncertainty Field. *IEEE Trans. Pattern Anal. Mach. Intell.*, 16(10):1002–1017, 1994.
- Panos E. Trahanias and Yiannis Komninos. Robot motion planning:

Multi-Sensory Uncertainty Fields enhanced with obstacle avoidance.
In *Proceedings of the IEEE/RSJ International conference on Intelligent Robots and Systems (IROS)*, 1996.

N. A. Vlassis and P. Tsanakas. A Sensory Uncertainty Field model for unknown and non-stationary mobile robot environments. In *Proceedings of the IEEE International conference on Robotics & Automation (ICRA)*, 1998.