Efficient Belief Space Planning using Sparse Approximations

Khen Elimelech and Vadim Indelman
Autonomous Navigation and Perception Lab (ANPL), Technion - Israel Institute of Technology
{khen,vadim.indelman}@technion.ac.il

I. Motivation

In this era, intelligent autonomous agents and robots can be found all around us. These agents share the same fundamental goal – to autonomously plan and execute their actions. These agents are required to account for real-world uncertainty when planning their actions, in order to achieve reliable and robust performance. There are multiple possible sources for such uncertainty, including dynamic environments, in which unpredictable events might occur; noisy or limited measurements, such as an imprecise GPS signal; and inaccurate delivery of actions. Also, problems such as long-term autonomous navigation and sensor placement over large areas, often involve optimization of numerous variables. These settings are translated to high-dimensional probabilistic states, known as "beliefs". Appropriately, the corresponding problem is known as Belief Space Planning (BSP). Relevant instantiations include active Simultaneous Localization and Mapping (SLAM), sensor placement and active sensing, robotic arm manipulation, and recently more profound problems, such as dialogue management. The objective in these planning problems is to select "safe" actions, which reduce the uncertainty in the belief. When planning under uncertainty, one should evaluate the propagation of the belief's uncertainty, considering *multiple* courses of action; yet, proper uncertainty measures, such as differential entropy, are expensive to calculate. Overall, the computational complexity of the problem can turn exceptionally high, thus making it challenging for online systems, or when having a limited processing power.

II. Problem Definition

For conciseness, let provide here a simplified definition to the BSP problem. No limiting assumptions on our approach should be deduced from this definition. Consider a sequential BSP scenario, at time-step k, the agent transitions from pose x_{k-1} to pose x_k, using a control u_{k-1}. It then receives an observation of the world z_k, based on this updated state. The transition and observation are both probabilistic operations, with some Gaussian noise. At each time-step, the agent maintains the posterior distribution over its current state vector X_k, given the controls and observations until that time; this distribution is also known as its *belief*:

$$b_k \doteq \mathbb{P}(X_k \mid u_{0:k-1}, z_{0:k}) = \mathcal{N}(\hat{X}_k, \Lambda_k^{-1}), \tag{1}$$

where $u_{0:k-1} \doteq \{u_0, \ldots, u_{k-1}\}$ and $z_{0:k} \doteq \{z_0, \ldots, z_k\}$. The agent's state vector consists of the series of poses, and may also include external variables, which are introduced by the observations (such as landmarks in a full-SLAM scenario).

To describe belief b_k, we can use the information matrix Λ_k, the inverse of the covariance matrix Σ_k. We can now reason about an updated belief at time $k+1$, after performing a control u_k and taking an observation z_{k+1}. The information matrix can be updated according to the following rule:

$$\Lambda_{k+1} = \Lambda_k + U_k^T U_k \tag{2}$$

where the *collective Jacobian* U_k encapsulates the new information regarding the control and the succeeding observation (for the full derivation see, for example, [10]). Each control can be described using a collective Jacobian of this form. Thanks to the additivity of the information, we can easily examine the information matrix after applying a sequence of T controls $u \doteq u_{k:k+T-1}$; the respective collective Jacobians of each control can simply be stacked to yield the collective Jacobian U of the entire sequence.

Now, considering an initial belief b, and a set $\mathcal{U}$ of such candidate control sequences, we wish to select the one which minimizes the expected uncertainty in the future belief, which is measured as (differential) entropy. Thus, we can define the following objective function:

$$J(b,u) = \frac{1}{2} \cdot \left(\ln\left|\Lambda + U^T U\right| - N \cdot \ln(2\pi e) \right), \tag{3}$$

where Λ is the information matrix of the prior belief b, and U is the collective Jacobian of u. An alternative way to represent the belief is using the upper triangular square root R of the information matrix Λ, given (e.g.) by the Cholesky decomposition: $\Lambda = R^T R$. This representation is used in prominent state-of-the-art SLAM algorithms (such as in[11]). To solve this decision problem, we should select u^*, s.t.

$$u^* = \underset{u \in \mathcal{U}}{\operatorname{argmax}}\, J(b,a). \tag{4}$$

Our goal is to allow an efficient solution to this problem.

III. Our Approach

A traditional solution requires calculation of the objective function for each candidate control action. We point out a key observation: to solve the decision problem, we should only *sort* the candidate actions in terms of their objective function value; when two problems maintain the same order of actions, their solution is equivalent. In this case, we can simply say that the two problems are *action consistent*. We thus suggest to identify and solve a simplified yet analogous decision problem, which results in the same (or similar) action selection, but for which the solution is more computationally efficient. In this work, we focus on simplifying the initial belief b.

Generally, solving a simplified problem may lead to loss in the quality of solution, when the selected action is not the real optimal action. Most often it is indeed possible to settle for a sub-optimal action selection in order to reduce the complexity of the problem; yet, to provide valuable results, it is important to set bounds over this loss. This can sometimes be done using the solution of the simplified problem.

As we saw, in BSP, calculation of the objective function involves calculation of the determinant of the posterior information matrix (Eq. 3). The cost of this calculation depends directly on the number of non-zero elements in the matrix, and is significantly lower for sparse matrices. Thus, sparsifying the posterior information matrices shall reduce the cost of solving the problem, as desired. Thanks to the additivity of the information, sparsifying the prior information matrix Λ essentially leads to sparser posterior information matrix for every candidate action. Notably such sparsification of the prior is only calculated once, for any number of actions. We also note that in many problems, especially in navigation problems, the collective Jacobians are also sparse; hence, even after adding the new information, the posterior information matrix shall remain sparse. An appropriate belief sparsification algorithm is given in [7]. The algorithm depends on a pre-selection S of state variables, and wisely removes elements which correspond to these variables from the belief's information matrix (or its triangular root). Approximations of different degrees can be generated using different variable selections.

Considering a specific action, a state variable is *involved* if applying the action directly impacts the variable; i.e., if the relevant transition and observation models are a function of this variable. Practically, in the collective Jacobian of an action, each of the columns represents a variable of the state vector; a variable is involved if at least one of the entries in its matching column is non-zero; uninvolved variables correspond to columns of zeros. In a navigation scenario, for example, the observed landmarks, and variables of the current pose are *involved*; variables referring to landmarks from the past, which are not observed anymore, are *uninvolved*. In this scenario, after a while, most landmarks are expected to be out of the visible range, and only a small portion of nearby landmarks shall remain relevant.

We claim that sparsification of uninvolved variables does not affect the posterior information determinant. Hence, when sparsifying variables which are uninvolved for all the actions, the objective function values are unaffected, and the resulting approximation is *action consistent* (and therefore, inducing no loss in the quality of action selection). It is possible to sparsify also involved variables, but then action consistency is not guaranteed. In that case a bound over the loss can be derived, as done in [8]. In conducted experiments, we demonstrated that even when sparsifying all the variables, the quality of solution is still well preserved, while achieving a significant improvement in decision making time (as shown in Fig. 1).

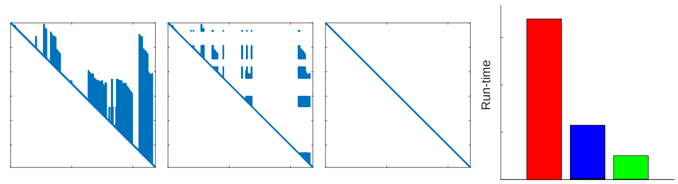

(a) **(b)**

Figure 1: Data is taken from a realistic active-SLAM scenario. (a) Comparison of an original information root matrix and its sparse approximations, for different variable selections S. On the left – the original matrix; in the center – the matrix after partial sparsification of about half of the variables; on the right – the matrix after full sparsification. (b) Decision making time, using the corresponding versions of the matrix. The action selection (i.e. quality of solution) was unaffected by these approximations.

We can now summarize the presented paradigm for the efficient solution of BSP problems: start with an initial belief b; (optionally) identify uninvolved variables and select a set S of variables to be sparsified; find a sparse approximation b_s using the suggested algorithm; easily calculate the objective values of all candidate actions using b_s, and select the optimal candidate action; (optionally) bound the loss to guarantee the quality of solution; finally, apply the selected action on the *original* belief b. Most importantly, we notice that this simplification method is completely separated from the state inference problem, and therefore does not compromise its accuracy.

IV. Comparison to Existing Methods

Several works [e.g. 23, 9, 4, 3, 15, 17] consider sparsification for the state inference problem, in order to limit the state size and allow long-term operation. However, these methods do not examine sparsification in the context of planning problems (influence over action selection, computational benefits, etc.). The novelty in our approach is the exploitation of sparse approximations exclusively for efficient decision making – a concept which we were the first to introduce.

Furthermore, the main advantage of our approach is that, in fact, the simplified problem can be solved in any desired manner, making our approach complementary to other methods. The research community extensively investigated BSP solution methods to provide better scalability in real-world problems. These include, for example, point-based value iteration [e.g. 19, 18], sampling-based motion planning [e.g. 20, 12, 13], and direct trajectory optimization approaches [e.g. 10, 24]. Approaches that focus on autnomous navigation problems, such as active SLAM, have been also widely examined [e.g. 22, 2, 6, 14, 5].

More closely related to our approach, a few other works [e.g. 16, 1, 21] examine approximation of the state or the objective function in order to reduce the planning complexity. These approaches, however, suffer some significant drawbacks. They either consider limiting assumptions on the planning scenario, are unclear regarding the quality of action selection, or do not demonstrate improvement in efficiency.

References

[1] Shaunak D Bopardikar, Brendan Englot, Alberto Speranzon, and Jur van den Berg. Robust belief space planning under intermittent sensing via a maximum eigenvalue-based bound. *IJRR*, 35(13):1609–1626, 2016. doi: 10.1177/0278364916653816.

[2] M. Bryson and S. Sukkarieh. Observability analysis and active control for airborne SLAM. *IEEE Trans. Aerosp. Electron. Syst.*, 44:261–280, 2008.

[3] N. Carlevaris-Bianco and R. M. Eustice. Conservative edge sparsification for graph slam node removal. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 854–860, 2014.

[4] Nicholas Carlevaris-Bianco, Michael Kaess, and Ryan M Eustice. Generic node removal for factor-graph SLAM. *IEEE Trans. Robotics*, 30(6):1371–1385, 2014.

[5] Stephen M Chaves and Ryan M Eustice. Efficient planning with the Bayes tree for active SLAM. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 4664–4671. IEEE, 2016.

[6] J. Du, L. Carlone, M. Kaouk Ng, B. Bona, and M. Indri. A comparative study on active SLAM and autonomous exploration with particle filters. In *Proc. of the IEEE/ASME Int. Conf. on Advanced Intelligent Mechatronics*, pages 916–923, 2011.

[7] K. Elimelech and V. Indelman. Scalable sparsification for efficient decision making under uncertainty in high dimensional state spaces. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, September 2017.

[8] K. Elimelech and V. Indelman. Efficient decision making and belief space planning using sparse approximations. *Intl. J. of Robotics Research*, 2018. Submitted, available at http://www.khen.io.

[9] G. Huang, M. Kaess, and J.J. Leonard. Consistent sparsification for graph optimization. In *Proc. of the European Conference on Mobile Robots (ECMR)*, pages 150 – 157, 2012.

[10] V. Indelman, L. Carlone, and F. Dellaert. Planning in the continuous domain: a generalized belief space approach for autonomous navigation in unknown environments. *Intl. J. of Robotics Research*, 34(7):849–882, 2015.

[11] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert. iSAM2: Incremental smoothing and mapping using the Bayes tree. *Intl. J. of Robotics Research*, 31:217–236, Feb 2012.

[12] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *Intl. J. of Robotics Research*, 30(7):846–894, 2011.

[13] L.E. Kavraki, P. Svestka, J.-C. Latombe, and M.H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot. Automat.*, 12(4):566–580, 1996.

[14] D. Kopitkov and V. Indelman. No belief propagation required: Belief space planning in high-dimensional state spaces via factor graphs, matrix determinant lemma and re-use of calculation. *Intl. J. of Robotics Research*, 36 (10):1088–1130, August 2017.

[15] H. Kretzschmar and C. Stachniss. Information-theoretic compression of pose graphs for laser-based SLAM. *Intl. J. of Robotics Research*, 31(11):1219–1230, 2012.

[16] David A McAllester and Satinder Singh. Approximate planning for factored pomdps using belief state simplification. In *UAI*, pages 409–416. Morgan Kaufmann Publishers Inc., 1999.

[17] Beipeng Mu, Liam Paull, Ali-Akbar Agha-Mohammadi, John J Leonard, and Jonathan P How. Two-stage focused inference for resource-constrained minimal collision navigation. *IEEE Trans. Robotics*, 33(1):124–140, 2017.

[18] J. Pineau, G. J. Gordon, and S. Thrun. Anytime point-based approximations for large POMDPs. *J. of Artificial Intelligence Research*, 27:335–380, 2006.

[19] J. M. Porta, N. Vlassis, M. T. Spaan, and P. Poupart. Point-based value iteration for continuous pomdps. *J. of Machine Learning Research*, 7:2329–2367, 2006.

[20] S. Prentice and N. Roy. The belief roadmap: Efficient planning in belief space by factoring the covariance. *Intl. J. of Robotics Research*, 28(11-12):1448–1465, 2009.

[21] Nicholas Roy, Geoffrey J Gordon, and Sebastian Thrun. Finding approximate pomdp solutions through belief compression. *J. Artif. Intell. Res.(JAIR)*, 23:1–40, 2005.

[22] C. Stachniss, D. Haehnel, and W. Burgard. Exploration with active loop-closing for FastSLAM. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2004.

[23] S. Thrun, Y. Liu, D. Koller, A.Y. Ng, Z. Ghahramani, and H. Durrant-Whyte. Simultaneous localization and mapping with sparse extended information filters. *Intl. J. of Robotics Research*, 23(7-8):693–716, 2004.

[24] J. Van Den Berg, S. Patil, and R. Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *Intl. J. of Robotics Research*, 31 (11):1263–1278, 2012.