Prospection for Robot Autonomy under Uncertainty

Ki Myung Brian Lee
Robotics Institute, University of Technology Sydney
Email: kmbrian.lee@uts.edu.au

I. Introduction

Greater autonomy is being asked of robots than ever before. Modern applications demand robots to 'just work' without human intervention in increasingly dynamic and uncertain environments. I am interested in achieving such a higher level of autonomy for robots under uncertainty.

Significant progress has been made toward robot autonomy, such that many commercial systems demonstrate reasonable *adaptation* to uncertain environments. Such adaptation is predominantly driven by the perception-action cycle, or the *separation principle* model, illustrated in Fig. 1a. In this model, perception and decision-making are separate modules that treat each other as independent black boxes. Perception processes sensor data to produce an environmental representation, which is subsequently consumed by decision-making to produce an action plan. Enhancing autonomy is achieved by improving these two modules in terms of robustness to uncertainty or adding capabilities such as higher-level user inputs. Can we, then, eventually create a truly autonomous robot by allowing for all practical ranges of uncertainty and user inputs?

I believe this is not the case, because an autonomous system based on the separation principle is inherently adaptive and hence reactive. Because decision-making treats perception as a black box, even non-myopic planners are rendered myopic in the sense that it does not account for how different actions lead to resolution of uncertainty at varying degrees. Instead, uncertainty is simply avoided through risk aversion (e.g., [10]), which is further exacerbated by perception modules generating conservative predictions of the environment in unseen areas that are of interest to the planners (e.g., [21, 29, 1]). I believe that the next generation of robots should embrace and explore environmental uncertainties, rather than simply avoiding it.

My research proposes to overcome such limitations of adaptive autonomy through *robotic prospection* [14]. The concept of prospection originates from cognitive psychology, where it refers to the generation and evaluation of possible future scenarios and outcomes during decision-making [7, 24]. That is, prospection equips an agent with foresight on possible future environmental states and outcomes.

To realise prospection in robot autonomy, I view the two main elements to be: 1) prospective planning and 2) prospective perception, illustrated in Fig. 1b. The *prospective perception* problem asks to generate useful predictions about the unseen environment given limited onboard sensor data, informed by prior data or domain knowledge. *Prospective planning* asks to devise a strategy for choosing control actions

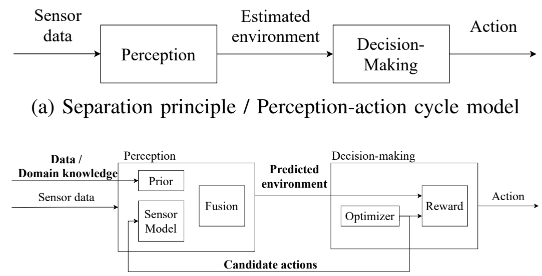

(a) Separation principle / Perception-action cycle model

(b) Prospection

Fig. 1. Comparison of (a) the standard and (b) the proposed autonomy architectures. Whereas the standard architecture passively processes sensory data into environmental estimate and subsequently into action, the planning module in the proposed architecture takes perception into account, and predicts unseen areas with domain knowledge or data.

that achieve a given task while gathering relevant information in doing so, through the balance of information gathering (exploration) and task completion (exploitation). In other words, prospective planning must consider what task-relevant information can be gathered from different actions, and how to balance this with task completion.

This perspective unifies several existing and emerging ideas. Prospective planning considers the effect of measurements as do partially observable Markov decision processes (POMDPs) [27, 13], but aims to side-step the need for exhaustive sensor simulation in POMDPs through the use of acquisition functions for exploration-exploitation trade-off, similar to Bayesian optimisation (BO) [25, 28] and mutli-armed bandits (MABs) [2, 6]. Meanwhile, much more general problem settings than BO and MABs are enabled through the use of information-theoretic quantities as used in active perception [12, 8, 4]. Some recent work already implements prospective perception in applications such as indoor modelling [22, 26]. Combination of these frameworks with prospective planning will be beneficial in terms of not only accurate feasibility or cost computation, but also better evaluation of the quality of measurements.

I anticipate that robotic prospection will form the basis of the next generation of robot autonomy frameworks, because proactivity will be necessary for upcoming application of robotics. In what follows, I detail the progress and results thus far, which support the claimed benefits of robotic prospection.

II. Progress to Date

Throughout my PhD, I developed a number of tools for prospective perception and planning in a general Bayesian setting. Perception is modelled as a Bayesian inference of some environmental parameters of arbitrary class (e.g. discrete or continuous), and the robot's task is modelled by a reward function dependent on these environmental parameters. Then, the prospective perception problem can be formalised as that of developing useful *predictive priors* on these environmental parameters, and the prospective planning problem as devising a strategy that outperforms the conventional expected reward maximisation one (i.e. expectimax).

A. Prospective Planning

For prospective planning in the most general Bayesian setting, I developed the mutual information upper confidence bound (MI-UCB) strategy in [17]. MI-UCB is an exploration-exploitation trade-off strategy that approximately maximises the *posterior expected reward* over a horizon given any possible future measurements, by instead maximising the weighted sum of Shannon information gain and prior expected reward before measurements. I have shown that the weighted sum of information gain and prior expected reward forms a valid probabilistic UCB on the posterior expected reward of a trajectory for any future measurements and any class of environmental parameters, whether discrete and continuous. This allows no-regret maximisation of the future posterior expected reward over a horizon without knowing the values of measurements ahead.

The weighted sum structure of the MI-UCB allowed scaling prospective planning to non-myopic coordination of a heterogeneous multi-robot system comprising scout- and task-robots, each equipped with sensing (i.e. information gathering) and/or task (i.e. reward-seeking) payloads. I demonstrated MI-UCB in a high-fidelity simulation of a heterogeneous multi-drone team performing a search-and-capture task [18], and showed that it outperforms the standard approach of maximising expected reward under current environmental belief by up to 134% in terms of targets captured.

I also studied special cases of prospective planning with a greater problem structure. Most recently, I considered complex tasks specified in temporal logic, which can specify tasks such as "visit target A and target B in any order" or "avoid target A until visiting target B". I proposed random signal temporal logic (RSTL) [19] to model such tasks with environmental uncertainty, and developed specialised acquisition functions for task-relevant information that outperform the general MI-UCB with up to two-fold reduction in task failure [9].

B. Prospective Perception

To develop *predictive priors* for prospective perception, my work thus far focused on incorporating physical laws of nature into Bayesian methods. Specifically, I devised several specialized Gaussian process (GP) regression schemes that satisfy governing partial differential equations (PDEs) by construction. This includes GPs satisfying 1) incompressibility constraint for ocean current estimation [16, 31], 2) advection-diffusion PDEs for chemical plumes in oceanic flows [15], and 3) Eikonal PDE for signed distance fields in obstacle mapping [34]. All of these perception algorithms exhibit much better prediction of unseen areas given limited measurements compared to conventional priors owing to the domain knowledge, which is useful for classical planning in many practical applications such as indoor and oceanic navigation [5, 35, 30]. More interestingly, the combination of prospective perception in GPs with prospective planning via BO [28] exhibited intelligent exploration-exploitation in plume source localisation [15].

III. Ongoing and Future Work

A. Prospection as a Framework for Autonomy

I aim to establish robotic prospection as a go-to framework for robotic systems. This will be an important step in robotics research toward proactive autonomy in uncertain environments. MI-UCB [17] already provides a good starting point, as it only requires the computation of Shannon information gain and prior expected reward regardless of the type or representation of the environmental uncertainty. I aim to apply MI-UCB in a greater variety of applications with environmental uncertainty, and welcome suggestions from collaborators.

Alternative formulations may be beneficial in doing so. MI-UCB yields probabilistic guarantees owing to its Bayesian foundation; however, deterministic error bounds and guarantees are often more useful than probabilistic ones in planning. To this end, I collaborated on deriving and working with deterministic worst-case error bounds in important perception problems such as SLAM [3] and spatial field reconstruction [32]. A part of my research will examine translating these bounds to acquisition functions for prospective planning.

B. Inductive Bias

My research on prospective perception has focused on imposing laws of nature that are *always* true; however, also useful would be to incorporate knowledge that is *usually* true, as long as it is evidenced by data. In other words, *inductive bias* could be as useful as laws of nature. This was explored in [33], where we propose a learnt model of topological information to accelerate trajectory prediction of crowds.

Deep generative models have already shown excellent applicability for modelling inductive bias in applications such as indoor structures [36, 22, 26]. The missing piece for prospection is the evaluation of uncertainty, quality of measurements and relevance to task at hand. In the simplest case, tractable computation of Shannon information in these models will allow using MI-UCB for prospective planning. I welcome collaborators in addressing this gap.

C. Multi-Robot Systems

Multi-robot systems be increasingly important in large-scale applications. Robotic prospection offers exciting insights for multi-robot coordination by viewing the actions of other teammates as a parameter to estimate, as is done in the control-as-inference literature [11, 20, 23]. In the context of multi-robot systems, this perspective will allow coordination with sparse or planned communication from which the teammates' actions are inferred. I am currently working on these ideas for physical human-robot collaboration in defence applications.

REFERENCES

[1] A. A. Agha-mohammadi, E. Heiden, K. Hausman, and G. Sukhatme. Confidence-rich grid mapping. *The Int. J. of Rob. Res.*, 38(12-13):1352–1374, 2019.

[2] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Mach. learn.*, 47(2-3):235–256, 2002.

[3] Y. Chen, L. Zhao, K. M. B. Lee, C. Yoo, S. Huang, and R. Fitch. Broadcast your weaknesses: cooperative active pose-graph slam for multiple robots. In *Proc. of IEEE ICRA*, 2020.

[4] J. Delmerico, S. Isler, R. Sabzevari, and D. Scaramuzza. A comparison of volumetric information gain metrics for active 3D object reconstruction. *Auton. Rob.*, 42: 197–208, 2018.

[5] G. D'urso, J. J. H. Lee, K. M. B. Lee, J. Shields, B. Leighton, O. Pizarro, C. Yoo, and R. Fitch. Field trial on ocean estimation for multi-vessel multi-float-based active perception. In *ICRA2021 1st Advanced Marine Robotics TC Workshop: Active Perception*, 2021.

[6] A. Garivier and O. Cappé. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proc. of Conf. Comput. Learn. Theory*, pages 359–376, 2011.

[7] D. Gilbert and T. Wilson. Prospection: Experiencing the future. *Science*, 351:1351–1354, 2007.

[8] G. A. Hollinger and G. S. Sukhatme. Sampling-based robotic information gathering algorithms. *Int. J. of Rob. Res.*, 33(9):1271–1287, 2014.

[9] R. Hull, K. M. B. Lee, J. Wakulicz, C. Yoo, J. McMahon, B. Clarke, S. Anstee, J. Kim, and R. Fitch. Decentralised active perception in continuous action spaces for the coordinated escort problem. In *Proc. of IEEE ICRA*, page To appear, 2023.

[10] M. Jun and R. D'Andrea. Path planning for unmanned aerial vehicles in uncertain and adversarial environments. In *Cooperative Control: Models, Applications and Algorithms*, page Ch. 6. Springer US, 2003.

[11] H. J. Kappen, V. Gómez, and M. Opper. Optimal control as a graphical model inference problem. *Mach. learn.*, 87(2):159–182, 2012.

[12] A. Krause and C. Guestrin. Near-optimal nonmyopic value of information in graphical models. In *Proc. of Uncertainty in Artificial Intelligence (UAI)*, pages 324–331, 2005.

[13] H. Kurniawati. Partially observable Markov decision processes and robotics. *Annual Review of Control, Rob. and Auton. Systems*, 5:253–277, 2022.

[14] K. M. B. Lee. *Prospection for Mobile Robots in Unknown Environments*. PhD thesis, University of Technology Sydney, 2023.

[15] K. M. B. Lee, J. J. H. Lee, C. Yoo, B. Hollings, and R. Fitch. Active perception for plume source localisation with underwater gliders. In *Proc. of Australasian Conference on Robotics and Automation (ACRA)*, page Best Student Paper Award, 2018.

[16] K. M. B. Lee, C. Yoo, B. Hollings, S. Anstee, S. Huang, and R. Fitch. Online estimation of ocean current from sparse GPS data for underwater vehicles. In *Proc. of IEEE ICRA*, pages 3443–3449, 2019.

[17] K. M. B. Lee, F. H. Kong, R. Cannizzaro, J. L. Palmer, D. Johnson, C. Yoo, and R. Fitch. An upper confidence bound for simultaneous exploration and exploitation in heterogeneous multi-robot systems. In *Proc. of IEEE ICRA*, 2021.

[18] K. M. B. Lee, J. J. H. Lee, C. Yoo, B. Hollings, and R. Fitch. Decentralised intelligence, surveillance, and reconnaissance in unknown environments with heterogeneous multi-robot systems. In *ICRA2021 Robot Swarms in the Real World: From Design to Deployment*, page Best Poster Award, 2021.

[19] K. M. B. Lee, C. Yoo, and R. Fitch. Signal temporal logic synthesis as probabilistic inference. In *Proc. of IEEE ICRA*, pages 5483–5489, 2021.

[20] S. Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.

[21] H. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proc. of IEEE ICRA*, pages 116–121, 1985.

[22] A. Pronobis and R. P. N. Rao. Learning deep generative spatial models for mobile robots. In *Proc. of IROS*, 2017.

[23] N. Rhinehart, R. McAllister, and S. Levine. Deep imitative models for flexible inference, planning, and control. In *Proc. of ICLR*, April 2020.

[24] M. E. Seligman, P. Railton, R. F. Baumeister, and C. Sripada. Navigating into the future or driven by the past. *Perspectives on psychological science*, 8(2):119–141, 2013.

[25] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proc. of the IEEE*, 104 (1):148–175, 2015.

[26] R. Shrestha, F.-P. Tian, W. Feng, P. Tan, and R. Vaughan. Learned map prediction for enhanced mobile robot exploration. In *Proc. of IEEE ICRA*, pages 1197–1204, 2019.

[27] D. Silver and J. Vennes. Monte Carlo planning in large POMDPs. In *Proc. of NIPS*, 2012.

[28] N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proc. of ICML*, ICML'10, pages 1015–1022, USA, 2010. Omnipress. ISBN 978-1-60558-907-7.

[29] S. Thrun. Learning occupancy grids with forward models. In *Proc. of IROS*, volume 3, pages 1676–1681, 2001.

[30] K. Y. C. To, K. M. B. Lee, C. Yoo, S. Anstee, and R. Fitch. Streamlines for motion planning in underwater currents. In *Proc. of IEEE ICRA*, pages 4619–4625, 2019.

[31] K. Y. C. To, F. H. Kong, K. M. B. Lee, C. Yoo, S. Anstee, and R. Fitch. Estimation of spatially-correlated ocean currents from ensemble forecasts and online measure-

ments. In *Proc. of IEEE ICRA*, pages 2301–2307, 2021.

[32] J. Wakulicz, K. M. B. Lee, C. Yoo, T. Vidal-Calleja, and R. Fitch. Informative planning for worst-case error minimisation in sparse Gaussian process regression. In *Proc. of IEEE ICRA*, pages 11066–11072, 2022.

[33] J. Wakulicz, K. M. B. Lee, C. Yoo, T. Vidal-Calleja, and R. Fitch. Topological trajectory prediction with homotopy classes. In *Proc. of IEEE ICRA*, page To appear, 2023.

[34] L. Wu, K. M. B. Lee, L. Liu, and T. Vidal-Calleja. Faithful Euclidean distance field from log-Gaussian process implicit surfaces. *IEEE Robot. and Automat. Lett.*, 6(2): 2461–2468, 2021.

[35] L. Wu, K. M. B. Lee, L. Liu, and T. Vidal-Calleja. Log-gpis-mop: A unified representation for mapping, odometry and planning. *To appear in IEEE Transactions on Robotics*, 2023.

[36] K. Zheng, A. Pronobis, and R. P. N. Rao. Learning graph-structured sum-product networks for probabilistic semantic maps. In *Proc. of Conference on Artificial Intelligence (AAAI)*, 2018.