Research statement and proposal

Tin Lai

School of Computer Science,
University of Sydney, NSW 2006, Australia
https://cs.tinyiu.com
tin.lai@sydney.edu.au

November 29, 2021

My research interests span the areas of machine learning, robotic systems, motion planning and probabilistic inference. A common thread in my research is in understanding the theory and design of scalable architectures and accelerated approaches in a learning-based robotic system. In particular, machine learning allows robots to take a data-driven approach to leverage prior experience. Robotics is an interdisciplinary branch of computer science and engineering that historically involve physical hardware and optimal control. As a learning-based approach, I have resorted to mathematical methods of proof borrowed from the areas of Machine Learning, Statistics, Set Theory, Combinatorics *&* Probability. Broadly speaking, my research intersects *Machine learning*—which is a branch of Artificial Intelligence that focuses on leverging data to imitate the way that humans learn, and *Robotic System*—which enables computational machines to interact with the real-world.

1 Background and Current Work

My Ph.D. dissertation focuses on modelling the planning problem as learning-based multi-objective optimisation problems that arise in most realistic robotic scenarios. Planning enables computational machines to reason about the possible approaches for the agent to interact with the world [LM20]. Typically, more controllable actuators imply more freedom for a robot to execute its actions. However, the infamous *curse of dimensionality* typically arise due to the exponential growth of search space [ES14]. Sampling-based motion planners (SBPs) are a class of robust methods that takes a probabilistic approach to tackle this issue [KKL96]. SBPs uses a graph structure [KF11] such as roadmap [KKL96] or tree [LaV98] to explore the search space. They are *probabilistic complete* and *almost-surely asymptotic optimal* [GS21]—that is, converge to the optimal solution with probability one. However, traditional approaches suffers the narrow passages issue [YJSL05, WAS99, HJRS03, SHJK05] and often recreate motion plan from scratch for every new tasks, which throws away useful information.

On the other hand, utilising space information have been proven to be helpful. For example, with an expansive tree [HLM97] or exploting local information [RBK08]. Multiple search trees [KOH+15] or local trees [Str04] can often speed up the exploration process; bridge test [QMI+13] can also helps to identify narrow passages [WZX18]. However, these are algorithmic approaches that do not utilise prior experience in a data-driven approach [LS02]. The benefit of a learning-based approach lies in the utilisation of information observed from previous examples [ABSC12, Che15] to improve future planning.

2 My Research Focus

My research aims to bridge the gap between the learning-based and algorithmic-based approaches. Traditional approaches, such as optimal control, often provide theoretical guarantees in bounds by rigorous theorems that ensure safety when robots interact with the world. In contrast, learning-based approaches often fail to provide strict guarantees due to their data-driven nature. However, they offer an excellent prior belief to robots on how the world works.

Combining Machine Learning and Algorithmic Approach: An innovative contribution of my work is establishing a new perspective to incorporating learning into the algorithmic planning framework while maining theoretical guarantees. My research on robot planning algorithms is two-fold. Firstly, I address the issue of the narrow passage by decoupling the planning with a disjointed-tree approach to simultaneously explore multiple regions [LRF19]. This approach mimics an MCMC random walk, borrowing from statistics literature. I then propose to use a Bayesian approach in the sampling distribution to update our prior belief [LMRF20] on the distribution as we observe more data from the environment [LR21c]. A Bayesian view on planning problems allows us to exploit prior domain knowledge on the problem setup while updating our posterior as we observe more outcomes. This is especially essential in robotic applications as they are often interdisciplinary with various existing domain-specific knowledge and theorems. Rather than reinventing the wheels with a purely learning-based approach, we should instead better utilise existing domain knowledge to combine the strength from both fields.

Reasoning on Past Experience: Replanning is often necessary when the robot's internal world representation differs from its surrounding environment changes. This is especially true in many real-world applications as the world is inherently dynamic with spatiotemporal uncertainty. Therefore, constructing a robust agent calls for the need of leveraging experience to reduce recomputation in subsequent episodes [LR21b]. Moreover, we can often exploit structure and patterns within the problem setup to learn to plan more optimally in subsequent episodes [LR20]. I have proposed several algorithms that leveraged learning-based approaches in robot planning. Importantly, we can still maintain theoretical results that support the guarantees of completeness and solution optimality for the planning horizon.

Formualise Sampling Distribution for Optimal Control and Planning: Complementary to previous works, I also focus on improving sampling distribution with diffeomorphic bijective mappings which provides unique mathematical properties. One of the main classes of methods for planning and control is the sampling-based approach. Rather than explicitly representing the entire search space (which is often exponential to the dimensionality of the space), sampling-based approaches can be thought of as an approximation method. Planing and control algorithms perform reasoning on the search space through sampling, which can often provide a rough estimation of the actual problem. The performance of many existing algorithms relies on the quality of the samples, which I had attempted to address with normalising flow [LR21a] for learning a near-optimal distribution. Another follow-up work constructs a diffeomorphic distribution to draw samples that are near obstacles-free [LZHR21]. These distributions provides a grounded informative prior based on the robot's past experience under similar environments [LR20].

Probabilistic Planning and Reasoning: Robots should also hold a probabilistic view of the world as realistic environments constantly evolve as time goes by. For example, we can build a predictive framework [LZR19] for predicting human trajectories [ZLOR21] while avoiding collisions with pedestrians [ZLO+21]. This perspective creates a new class of planners that utilises prior experience to improve future decision making. These methods highlight the gap within literature and call for the need for a unified framework to integrate a learning-based approach in planning. I help to facilitate the research discussion among this line of work by open-sourcing a unified testing environment [Lai21] that allows easy integration with existing machine learning techniques for motion planning in a popular language.

3 Motivation for future research

Research on the border of computer science and artificial intelligence attracts many scholars from various communities. Traditional planning theories provide robust algorithmic procedures to tackle the complexity within robot control. However, the growing number of machine learning and data science literature indicates that we can extract much more from the readily available data for robots to perform inferences—akin to allow robots learn through experiences.

I believe the current interests in this field are the results of both intellectual and practical motives. The lowering cost of hardware and readily available sensors enable us to access a sufficient amount of data to construct models that are learnable from consumer hardware. This trend motivates researchers and industry partners to better utilise their domain-specific data for their application use-cases. I conclude this statement with a number of issues which, in my opinion, will motivate future research in this area.

Unified theories for learning-based complexity: Complexity issues are inherent to algorithmic research and are currently absent from these disciplines. Empirically, we can evaluate the performance gain by applying certain machine learning techniques. However, there is often a lack of theoretical guarantees on how effective are these approaches. Unifying these very different fields is a serious and fascinating challenge and will source many intriguing problems.

Uncertainty quantification on learned policy: Data-driven approaches tend to provide a point estimate of the objective function that we are modelling. However, since the prediction comes from some learned policy, it is often challenging to provide theoretical bounds on subsequent actions that depend on such a prediction. A possible approach is to compose some learning-based policies that provide upper and lower bounds on such estimation through convex optimisation. The benefit of such a policy is that the robot can then create robust planning and reasoning that accounts for the uncertainty due to the lack of data for rare events.

Probabilistic planning and reasoning: In addition to uncertainty that arises from learning-based policy, the world is also inherently stochastic. Modelling the stochastic nature of and accounting for the uncertainty should be an integrated part of robot planning. However, approaches commonly take a static snapshot of the world and plans without considering the stochasticity. They rely heavily on the controller to address the environmental noises. This becomes problematic when the internal belief about the world is also stochastic and imprecise. The planner will create plans that are infeasible for the underlying controller to follow. We should instead take a probabilistic view during planning to better account for the stochastic nature and minimise regrets during execution.

References

[ABSC12] I. Al-Bluwi, T. Siméon, and J. Cortés, *Motion planning algorithms for molecular simulations: A survey*, Computer Science Review **6** (2012), no. 4, 125–143. 00057.

[Che15] W. Chen, *Motion Planning with Monte Carlo Random Walks*, PhD Thesis, 2015. 00000.

[ES14] M. Elbanhawi and M. Simic, *Sampling-based robot motion planning: A review*, IEEE Access **2** (2014), 56–77. 00129.

[GS21] J. D. Gammell and M. P. Strub, *Asymptotically Optimal Sampling-Based Motion Planning Methods*, Annual Review of Control, Robotics, and Autonomous Systems **4** (May 2021), no. 1, 295–318 (en).

[HJRS03] D. Hsu, T. Jiang, J. Reif, and Z. Sun, *The bridge test for sampling narrow passages with probabilistic roadmap planners*, Proceedings of IEEE International Conference on Robotics and Automation, 2003, pp. 4420–4426. 00332.

[HLM97] D. Hsu, J.-C. Latombe, and R. Motwani, *Path planning in expansive configuration spaces*, Proceedings of IEEE International Conference on Robotics and Automation, 1997, pp. 2719–2726. 00725.

[KF11] S. Karaman and E. Frazzoli, *Sampling-based algorithms for optimal motion planning*, The International Journal of Robotics Research **30** (2011), no. 7, 846–894. 01458.

[KKL96] L. E. Kavraki, M. N. Kolountzakis, and J.-C. Latombe, *Analysis of probabilistic roadmaps for path planning*, Proceedings of IEEE International Conference on Robotics and Automation, 1996, pp. 3020–3025. 00000.

[KOH+15] S. Klemm, J. Oberländer, A. Hermann, A. Roennau, T. Schamm, J. M. Zollner, and R. Dillmann, *RRT*-Connect: Faster, asymptotically optimal motion planning*, 2015 IEEE International Conference on Robotics and Biomimetics (ROBIO), 2015, pp. 1670–1677.

[Lai21] T. Lai, *`sbp-env`: A python package for sampling-based motion planner and samplers*, Journal of Open Source Software **6** (2021), no. 66, 3782.

[LaV98] S. M. LaValle, *Rapidly-exploring random trees: A new tool for path planning*, TR 98-11, Computer Science Dept., Iowa State University (1998). 01934.

[LM20] T. Lai and P. Morere, *Robust Hierarchical Planning with Policy Delegation*, arXiv:2010.13033 [cs] (October 2020). arXiv: 2010.13033.

[LMRF20] T. Lai, P. Morere, F. Ramos, and G. Francis, *Bayesian Local Sampling-Based Planning*, IEEE Robotics and Automation Letters **5** (2020), no. 2, 1954–1961. 00000 Publisher: IEEE.

[LR20] T. Lai and F. Ramos, *Learning to Plan Optimally with Flow-based Motion Planner*, arXiv:2010.11323 [cs.RO] (October 2020). arXiv: 2010.11323.

[LR21a] ——, *PlannerFlows: Learning Motion Samplers with Normalising Flows*, IEEE/RSJ Proceedings of The International Conference on Intelligent Robots and Systems (IROS), 2021.

[LR21b] ——, *Rapid Replanning in Consecutive Pick-and-Place Tasks with Lazy Experience Graph*, arXiv preprint arXiv:2109.10209 (2021).

[LR21c] ——, *Rapidly-exploring Random Forest: Adaptively Exploits Local Structure with Generalised Multi-Trees Motion Planning*, IEEE Robotics and Automation Letters (RA-L) (2021).

[LRF19] T. Lai, F. Ramos, and G. Francis, *Balancing Global Exploration and Local-connectivity Exploitation with Rapidly-exploring Random disjointed-Trees*, Proceedings of The International Conference on Robotics and Automation, 2019. 00000.

[LS02] T.-Y. Li and Y.-C. Shie, *An incremental learning approach to motion planning with roadmap management*, Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292), May 2002, pp. 3411–3416 vol.4.

[LZHR21] T. Lai, W. Zhi, T. Hermans, and F. Ramos, *Parallelised Diffeomorphic Sampling-based Motion Planning*, arXiv preprint arXiv:2108.11775 (2021).

[LZR19] T. Lai, W. Zhi, and F. Ramos, *Occ-traj120: Occupancy maps with associated trajectories*, Computing Research Repository (CoRR) (2019).

[QMI+13] A. H. Qureshi, S. Mumtaz, K. F. Iqbal, B. Ali, Y. Ayaz, F. Ahmed, M. S. Muhammad, O. Hasan, W. Y. Kim, and M. Ra, *Adaptive potential guided directional-RRT*, Robotics and Biomimetics (ROBIO), 2013 IEEE International Conference on, 2013, pp. 1887–1892. 00009.

[RBK08] M. Rickert, O. Brock, and A. Knoll, *Balancing exploration and exploitation in motion planning*, Proceedings of IEEE International Conference on Robotics and Automation, 2008, pp. 2812–2817. 00000.

[SHJK05] Z. Sun, D. Hsu, T. Jiang, and H. Kurniawati, *Narrow passage sampling for probabilistic roadmap planning*, IEEE Transactions on Robotics **21** (December 2005), no. 6, 1105–1115. 00119.

[Str04] M. Strandberg, *Augmenting RRT-planners with local trees*, Proceedings of IEEE International Conference on Robotics and Automation, 2004, pp. 3258–3262 Vol.4 (en). 00092.

[WAS99] S. A. Wilmarth, N. M. Amato, and P. F. Stiller, *MAPRM: A probabilistic roadmap planner with sampling on the medial axis of the free space*, Proceedings of IEEE International Conference on Robotics and Automation, 1999, pp. 1024–1031. 00000.

[WZX18] W. Wang, L. Zuo, and X. Xu, *A Learning-based Multi-RRT Approach for Robot Path Planning in Narrow Passages*, Journal of Intelligent & Robotic Systems **90** (2018), no. 1-2, 81–100. 00000.

[YJSL05] A. Yershova, L. Jaillet, T. Siméon, and S. M. LaValle, *Dynamic-domain RRTs: Efficient exploration by controlling the sampling domain*, Proceedings of IEEE International Conference on Robotics and Automation, 2005, pp. 3856–3861. 00000.

[ZLOR21] W. Zhi, T. Lai, L. Ott, and F. Ramos, *Anticipatory Navigation in Crowds by Probabilistic Prediction of Pedestrian Future Movements*, Proceedings of the international conference on robotics and automation (icra), 2021, pp. 8459–8464.

[ZLO+21] W. Zhi, T. Lai, L. Ott, G. Francis, and F. Ramos, *Trajectory generation in new environments from past experiences*, Ieee/rsj proceedings of the international conference on intelligent robots and systems (iros), 2021.