Research Statement

Cai Panpan `comp.nus.edu.sg/~caipp`
Research Fellow dcscaip@nus.edu.sg
National University of Singapore +65 98257466

Decision making or planning optimize the robots' behaviors to transit the real-world towards desired states. Planning for real-world tasks is extremely challenging. The world can be highly complex and dynamic, while robots only have limited sensing capabilities. Planning algorithms thus need to handle a plethora of uncertainties: imperfect robot control, noisy sensors, and fast-changing environments.

Planning becomes particularly challenging when involving many human participants interacting intensively with each other and the robot. A representative example is driving among a crowded traffic, where a robot vehicle interacts with many partially observable traffic participants. A large crowd induces a high-dimensional state space, highly-complex dynamics and uncertain human behaviours, which raises enormous difficulties in perception, prediction, and planning. To act safely and efficiently in such environments, sophisticated long-term planning is required, which additionally requires accurate models and efficient solvers.

I dedicate my research to large-scale decision making in complex and highly-interactive environments, especially those involving uncertainties and long-term planning. I aim to tackle crowded, chaotic environments and enable robots to accomplish complex tasks safely and efficiently. I have developed models and algorithms been validated on a variety of tasks: classical planning benchmarks, real-world industrial environments, and autonomous driving in crowded traffic. My work brings practical solutions to large-scale, long-term planning by focusing on three aspects: human behaviour modelling, real-time planning, and integration with learning. Specifically:

1. traffic agent motion models for accurate long-term predictions and realistic simulations;
2. massively-parallel planners for real-time planning in large-scale environments;
3. integration of planning and learning to to solve complex long-term planning tasks.

The following sections present the three aspects in detail.

Traffic motion models and crowd simulation

A "good" model not only needs to accurately model the complexity of real-world dynamics and human behaviours, but also needs to capture the intrinsic uncertainties in principled ways. We have formalize the interaction among human traffic participants as constrained optimization in the velocity space. Constraints of the problem encodes kinematic and collision avoidance constraints of traffic participants; Objective of the problem is to navigate efficiently towards the intended goals. Using this formulation, we proposed two traffic motion models: PORCA [9] for pedestrians and GAMMA [8] for mixed traffic. Both models can be solved using quadratic programming in linear time. The efficiency of these models enables integration to real-time crowd-driving algorithms. The power of using these models in planning has been demonstrated in both driving among pedestrians and urban driving [10].

Importantly, our motion models are conditioned on human factors such as intention, attention, willing to take responsibility. These factors are not observable, but can be effectively inferred using Bayesian filtering from the interaction history of traffic participants. This filtering process helps to produce accurate and highly-variable motion predictions.

Building upon these motion models, we have developed an driving simulator, SUMMIT [5], for simulating unregulated urban crowds (Fig. 1). The simulator parses online, real-world maps to construct realistic driving scenes and executes the motion models to simulate highly-interactive heterogeneous traffic. The purpose of the simulator is to provide unlimited amount of highly-interactive driving scenes to enable developing, training, and evaluating driving algorithms that tackle important problems such as perception, motion prediction, control, decision making, and end-to-end learning.

The model models has been published in RAL[9] and available on arXiv [8]. A collaborative work [10] presented in IROS 2019 demonstrates planning using the motion models in an real-world urban environment. The SUMMIT simulator has been presented in ICRA 2020 [5] with open-source code released on GitHub.

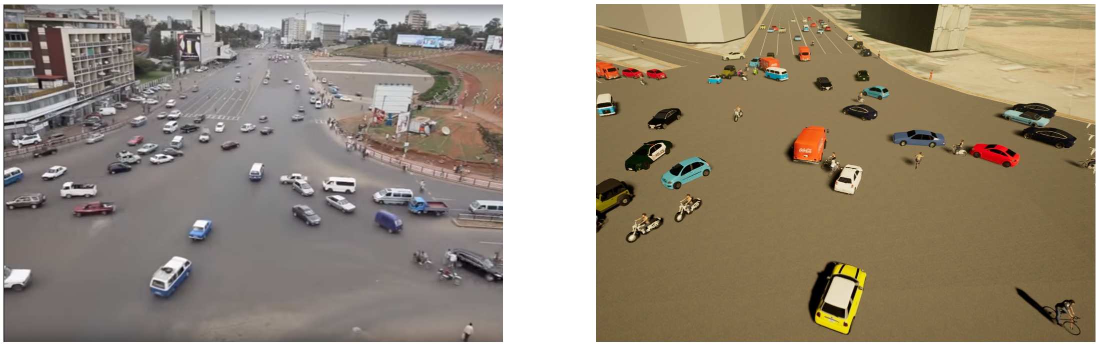

Figure 1: Driving among an unregulated traffic crowd. Left: a real-world intersection in Ethiopia, Africa; Right: the simulated scene in the SUMMIT simulator.

Large-scale planning under uncertainty

Real-world robots have to contend with a complex and uncertain environment. It often requires sophisticated long-term planning to achieve human-level performance. Unfortunately, long-term planning brings combinatorial complexities, also known as the "curse of dimensionality" and the "the curse of history": the complexity of optimal planning grows exponentially with the problem scale and the planning horizon. Practical planning algorithms are in urgent demand.

My idea to scale-up planning is massive parallelization. My early attempts focused on parallelizing motion planning in industrial environments. I developed real-time motion planners for crane-lifting in highly-complex industrial sites [2, 3]. The algorithms uses parallel Genetic Algorithms to plan globally-optimal paths and uses parallel pixel-space checking to detect collisions. Both the planner and the collision checker are integrated in a single hierarchical GPU parallelization scheme to achieve real-time performance.

This line of work has been published in IEEE Transactions on Industrial Informatics (TII) [3] and Automation in Construction (AIC) [2].

A greater challenge in robotics is to handle stochastic environments and partial observability in large-scale problems. Such problems are often solved in the belief space: the space of probability distributions over possible system states. State-of-the-art algorithms perform online belief tree search: at each time step, look-ahead from the current belief to search an optimal action. The system then execute the action, receive observations from sensors, update the belief, and enter the next planning cycle. Belief tree search offers a principled way to perform online planning under uncertainty. However, it still suffers from the combinatorial complexity and requires additional techniques to scale-up.

I have developed a massively-parallelized belief tree search algorithm, HyP-DESPOT [4], to scale up to large-scale problems. The core idea is to integrate CPU and GPU parallelization (Fig. 2): use CPU cores to parallelize irregular tasks, i.e., the tree search, and use GPU cores to parallelize regular tasks, i.e. roll-outs at leaf nodes. By doing so, HyP-DESPOT achieves hundreds of times of speed-up in various large-scale planning benchmarks, and enables a robot vehicle to drive among crowds of pedestrians safely and efficiently.

This line of research has been published in Robotics: Science Systems (RSS 2018) [4] and the International Journal of Robotics Research (IJRR). I have open-sourced the parallel planner with a general API for users to easily plugin their problem models and boost real-time planning for their own tasks.

Integrating planning and learning

Long-term planning has several inherent problems: accumulative model errors, combinatorial complexity w.r.t. the planning horizon, and exponentially decreasing coverage of Monte Carlo simulations. I seek to integrate planning with learning to address the above problems by exploiting both the robustness of explicit reasoning and the capability to learn from data.

I first proposed a strategy for integrating planning and learning: "think locally and learn globally". Specifically, it means to constrain search to short-term futures, and use learning to account for long-term

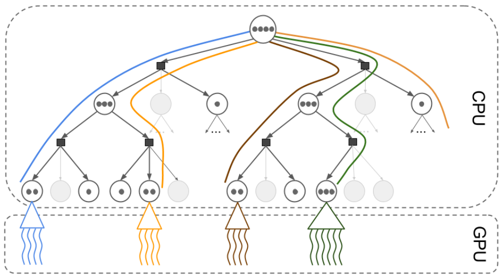

Figure 2: An illustration of the HyP-DESPOT algorithm, a massively parallelized belief tree search algorithm for large-scale planning under uncertainty.

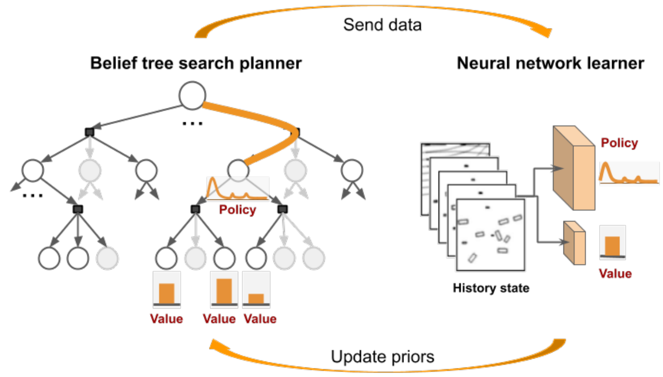

Figure 3: An illustration of LeTS-Drive: constrain search to short-term futures, and use learning to account for long-term futures.

futures. The concrete implementation is a crowd-driving algorithm, LeTS-Drive [6], that integrates offline imitation learning with online belief tree search (Fig. 3). LeTS-Drive learns two global priors from expert driving data in the offline stage: a policy function and a value function, both represented as neural networks. During online planning, LeTS-Drive uses the policy network to guide explorations within the belief tree, and applies the value network to initialize value estimations at leaf nodes. In effect, the learned priors encode an initial global policy and LeTS-Drive exploits online planning to improve this policy for particular problem instances. By integrating planning and learning, LeTS-Drive achieved superior driving performance among crowds of pedestrians, and outperforms either planning or learning alone.

The next step is to bring in feedback from the environment and enable the driving system to continuously learn from data. This is achieved by integrating planning and learning in a *closed-loop*: while the planner provides experiences to train the prior policy, the learned priors are also constantly fed back to the planner, thus closing the planning-learning loop. The new algorithm is flexible and takes advantage of both self-supervised and reinforcement learning. It quickly learns sophisticated driving skills among dense urban crowds and outperforms the previous two-stage integration scheme by a large margin.

In collaboration with my student, I have explored another possibility for integrating planning and learning: learning macro-actions for long-horizon planning. The idea is to learn situation-aware macro-actions and use the learned macro-actions to perform efficient, long-horizon planning. The core challenge is to learn a macro-action generator directly optimized for the down-stream planning. To achieve this, we additionally learn an auxiliary critic function as a differentiable approximation of the planner's value estimations. The critic informs the generator how good a macro-action set is for planning at the given belief, and serves as a surrogate objective to enable end-to-end training for the generator. The resulting algorithm, MAGIC [7], brings significant performance gains over planning on various long-horizon planning tasks.

This line of research has been published in RSS 2019 [6], RSS 2021 [7], and submitted to T-RO [1].

Summary

I aim to seek principled solutions to attack complex real-world problems. I work hard to propose mathematically-sound formulations of real-world problems and develop principled algorithms to solve them efficiently. I believe that explicit, sophisticated reasoning and its combination with learning are the key towards super-human intelligence. I have grounded this belief and dedicated my efforts to traffic motion modeling, real-time planning under uncertainty, and the integration of planning and learning.

References

[1] Panpan Cai and David Hsu. Closing the planning-learning loop with application to autonomous driving in a crowd, 2021.

[2] Panpan Cai, Yiyu Cai, Indhumathi Chandrasekaran, and Jianmin Zheng. Parallel genetic algorithm based automatic path planning for crane lifting in complex environments. *Automation in Construction*, 62:133–147, 2016.

[3] Panpan Cai, Indhumathi Chandrasekaran, Jianmin Zheng, and Yiyu Cai. Automatic path planning for dual-crane lifting in complex environments using a prioritized multiobjective pga. *IEEE Transactions on Industrial Informatics*, 14(3):829–845, 2017.

[4] Panpan Cai, Yuanfu Luo, David Hsu, and Wee Sun Lee. HyP-DESPOT: A hybrid parallel algorithm for online planning under uncertainty. In *Proc. Robotics: Science & Systems*, 2018.

[5] Panpan Cai, Yiyuan Lee, Yuanfu Luo, and David Hsu. Summit: A simulator for urban driving in massive mixed traffic. *To be presented at ICRA 2020*, 2019.

[6] Panpan Cai, Yuanfu Luo, Aseem Saxena, David Hsu, and Wee Sun Lee. Lets-drive: Driving in a crowd by learning from tree search. In *Proc. Robotics: Science & Systems*, 2019.

[7] Yiyuan Lee, Panpan Cai, and David Hsu. Magic: Learning macro-actions for online pomdp planning using generator-critic. In *Proc. Robotics: Science & Systems*, 2021.

[8] Yuanfu Luo and Panpan Cai. Gamma: A general agent motion prediction model for autonomous driving. *arXiv preprint arXiv:1906.01566*, 2019.

[9] Yuanfu Luo, Panpan Cai, Aniket Bera, David Hsu, Wee Sun Lee, and Dinesh Manocha. Porca: Modeling and planning for autonomous driving among many pedestrians. *IEEE Robotics and Automation Letters*, 3(4):3418–3425, 2018.

[10] M. Meghjani, Y. Luo, Q. H. Ho, P. Cai, S. Verma, D. Rus, and D. Hsu. Context and intention aware planning for urban driving. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.