STATEMENT OF CAREER GOALS

My long-term research goal is to build algorithmic methodologies that enable robotic systems to think and make decisions quickly, reliably, and robustly despite the complexity of the world around them. I believe that the process of thinking relies strongly on the ability to search through the myriads of possible action sequences and their consequences on the world and the robot itself, taking advantage of the information that is known, being able to quickly react to new information, and learning offline and through experience how to improve the speed and reliability of the decision-making process.

In pursuit of such methodologies, my group, called Search-based Planning Lab, focuses on developing novel graph search-based planning approaches. Typically, graph search-based planning approaches operate by representing a given planning problem as a search through a graph for a least-cost path and then applying a graph search algorithm to find an optimal or close-to-optimal solution.

For the problems that robots face in the real world though, these graphs are prohibitively large and also exhibit uncertainty. Consequently, it was a common belief in the robotics community that graph search approaches could not provide real-time performance guarantees when applied to such problems. I believe my group's work has changed this thinking by developing graph search algorithms that *are* capable of solving challenging problems in robotics in real-time while maintaining all the positive properties of graph search algorithms such as generality, cost minimization and rigorous guarantees on completeness and quality of solutions. We then used these algorithms for real-time planning on physical robots performing such tasks as autonomous navigation, autonomous flight and landing, autonomous mobile manipulation, humanoid mobility, coordination of multi-robot systems, and many others.

Search-based planning algorithms that we have developed are now widely studied and used both in academia and industry. For example, on the academic side, algorithms such as D* Lite [16], ARA* [22] and Anytime D* [21] have received several awards including the Influential 10-year Paper Award at ICAPS'17, the top venue for the latest research on planning and are now taught in many classes on AI and Robotics across the world. Our much more recent work on Multi-Heuristic A* [4], the first heuristic search to handle multiple (possibly many) inadmissible heuristics without losing its theoretical guarantees, has been widely cited, extended by other groups, and even used for world-wide teaching of basic programming alongside Gaussian Elimination, Newton Method and other well-known simple but powerful techniques [1]. Furthermore, the algorithms we develop get utilized and investigated well-beyond robotics, for example in fields such as game development [19, 5], machine learning [24], and computer animation [33, 15]. At the same time, my group used our algorithms to

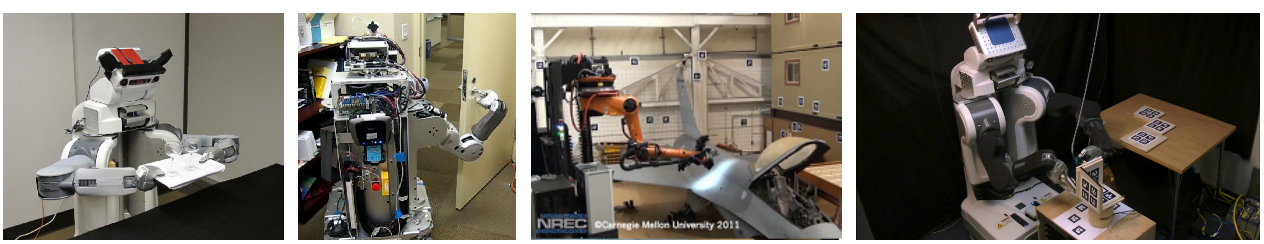

(a) dual-arm (b) full-body (c) industrial (d) assembly

Figure 1: Our work on developing novel graph search algorithms and compact graph representations and applying them to high-dimensional planning for a wide range of mobile manipulation tasks.

built planners for a highly diverse set of systems, for example, an industrial mobile manipulator used for paint stripping different aircrafts (Figure 1(c)) - the project that won the national Gold Edison award in 2013 - a full-size K-MAX helicopter (Figure 2(b)) used to fly autonomously at a speed of 100 knots, a full-size SUV (Figure 2(c)) that won the DARPA Urban Challenge race in 2007 [20], and a team of aerial vehicles performing border control, a system currently being commercialized by Mitsubishi Heavy Industries.

On the commercial side, the search-based planning algorithms my group developed are widely used in industry ranging from autonomous vehicles (e.g., by Nissan, Denso, Honda, etc.) to aerial vehicles (e.g., by Lockheed Martin, Mitsubishi, Dragonfly Pictures, etc.) to warehouse robotics (e.g., by Alert Innovation, Bossa Nova, etc.) to manipulators (e.g., by Mitsubishi, Honeywell, General Dynamics, etc.).

Furthermore, as an another example of the commercial success of our work on search-based planning, I have recently founded a company RobotWits devoted to developing search-based planning techniques for self-driving vehicles, grew the company to have a dozen of OEM and 1st tier-supplier customers, and then sold it in 2021 to Waymo, one of the leading self-driving companies in the world. With this acquisition, Waymo has now established additional offices in Pittsburgh and is actively growing its presence in the region, benefiting CMU and the city of Pittsburgh more generally.

Finally, it is a critical part of my goals to utilize the research outcomes we obtain in motivating the importance and study of mathematics and computer science among students that are in middle- and high-schools. Much of the research my lab does builds on basic concepts from mathematics and computer science and are accessible to anyone. I believe that we can and should use this to inspire the study of mathematics and computer science and to show the accessibility of pursuing STEM careers independently of the background. With this goal in my mind, in 2020 I co-produced a TV series (in conjunction with PBS TV network) called The Robot Doctor, that studies basic problems in robotics and shows how they can be translated into high-school level mathematics. The series has already been broadcasted and re-broadcasted across the state of Pennsylvania a number of times, has been nominated for a regional Grammy award, is now becoming available for nationwide broadcasting, and is being utilized by a number of educational institutions.

All in all, I am proud of the impact my group has made and continues to make on the field of robotics, AI and computer science more generally, and the broader impact we are making on the society we live in. In the following, I give few examples of our research. The first few examples describe several research themes that cut through most of the work that my group and I have done in the past, and the later examples describe some of the latest research directions that my group has been pursuing.

- **Graph search-based planning for solutions of bounded sub-optimality.** While finding a provably optimal path in a high-dimensional search-space is computationally intractable, for many planning problems in robotics, allowing even a small amount of sub-optimality in

(a) (b) (c) (d)

Figure 2: Our work on developing novel graph search algorithms and applying them to planning for autonomous flight and landing, autonomous navigation and control of small teams of robots.

the solution allows the search to quickly find high-quality solutions. We have exploited this observation to develop a number of graph search algorithms that allow the trade-off of solution quality for fast planning time including an anytime version of A*, ARA* [22], Anytime SIPP [28, 31] for planning in dynamic environments and Planning with Adaptive Dimensionality [8, 9]. Together with my students and colleagues, we have used these searches to build highly effective planners for high-dimensional robotic systems ranging from single-arm and dual-arm manipulation [7] to full-body manipulation on PR2 [6] (Figure 1(a,b)) and on a large mobile manipulator built to strip paint off airplanes autonomously (Figure 1(c)).

- **Incremental graph search algorithms.** Many problems in robotics require constant re-planning in response to the discovery of the environment, corrections in the localization of the robot, imperfect actuation and changes in the environment. Jointly with my students and collaborators, we have developed and continue to develop new incremental graph search algorithms that speed up repeated planning in such domains by re-using search efforts. Some of these algorithms include D* Lite [18, 16], Real-time Adaptive A* [17], Truncated Incremental Search [2], Anytime Tree-Restoring Weighted A* [10] and Anytime D* [21]. To the best of my knowledge, the algorithm Anytime D* we developed was the first heuristic search to be both anytime and incremental. We have used it to build motion planning for a variety of ground and aerial vehicles including a fully autonomous micro-aerial vehicle (Figure 2(a)) [23], a full-size K-MAX helicopter performing autonomous flight and landing (Figure 2(b)) and a full-size SUV (Figure 2(c)) that won the DARPA Urban Challenge race in 2007 [20].

- **Graph search algorithms capable of using multiple heuristics.** When planning for complex high-dimensional robotic systems, one can often derive multiple lower-dimensional planning problems that can be solved to provide guidance to the full-dimensional search in the form of heuristics (i.e., estimates on the cost-to-goal). Combining these multiple heuristic functions into a single heuristic function can often be highly ineffective. Furthermore, it is hard to ensure that all of these heuristic functions are admissible and consistent, the properties that are typically required to provide guarantees on completeness and solution quality. Motivated by these observations, we have developed a novel heuristic search, called Multi-Heuristic A* (MHA*) [4] and several extensions to it [32, 11]. MHA* takes in multiple, arbitrarily inadmissible heuristic functions in addition to a single consistent heuristic, and uses all of them simultaneously to search for a solution in a way that guarantees completeness and bounded sub-optimality. This methodology turned out to be highly effective for high-dimensional planning problems such as full-body mobile manipulation and humanoid mobility. The effectiveness of MHA* combined with its simplicitly and rigorous theoretical properties are typically what I strive to have the most in the algorithms my group and I develop.

- **Graph search algorithms that learn to improve their performance.** Robots are often used to perform similar tasks over and over again. It is therefore important for us to study how planning algorithms can improve their speed and robustness based on past planning experiences as well as demonstrations provided by humans and/or other robots. This is in contrast to incremental graph search algorithms that speed up re-planning within a *single* execution of a task but do not improve their performance over multiple executions of a task. To this end, my group has developed a new class of heuristic searches that *are* capable of improving their performance based on their previous experiences and demonstrated solutions [29, 30, 3]. In particular, we have developed a new approach to graph search-based planning that we call *Experience Graphs* [29]. Planning with Experience Graphs builds a faster-to-solve graph representation of the planning problem based on the solutions it has found previously or demonstrations provided by a person and utilizes this representation to focus the search for a solution in a way that preserves rigorous guarantees on completeness and bounded sub-optimality with respect to the original graph representation of the problem. Planning with Experience Graphs turned out to be highly beneficial in the variety of complex mobile manipulation tasks ranging from assembly (Figure 1(d)) to paint stripping (Figure 1(c)). To my knowledge, Experience Graphs is the first heuristic search method that "learns from its experience" a more compact graph representation that speeds up its future planning times and does it in a way that preserves rigorous guarantees on the solution quality. In fact, my student's thesis on Experience Graphs has been named a runner-up for the Best Thesis Award at the ICAPS'16 conference.

- **Deliberative Perception via Heuristic Search.** Another research direction that my group has been pursuing is the use of heuristic search algorithms to provide robust perception of a cluttered scene, in particular, identification and localization of objects whose 3D models are known. This task is common in manufacturing, warehouse and personal assistance domains, yet remains challenging when a scene is cluttered. Typical approaches for this task are based on scene-to-model feature matching or regression by statistical learners trained on a large database of annotated scenes. These approaches are fast but sensitive to occlusions, features, and/or training data. To overcome this brittleness, we developed a series of algorithms that seek to find the best explanation of the observed sensor data by hypothesizing possible scenes in a generative fashion [26, 25, 27]. In particular, these methods show how this problem can be transformed into a tree search problem and how this search can be made tractable for complex scenes by utilizing Multi-Heuristic A* search that we have previously developed. This line of research has been exciting for me not only because perception is a difficult problem but also because it illustrates a highly unorthodox and yet novel use of heuristic search algorithms.

- **Graph search-based provably Constant-time Motion Planning.** In manipulation applications, robots often have to do similar tasks over and over again, for example, picking up and putting down objects, opening and closing doors, cabinets and drawers, and moving chairs and other objects out of the way. Yet, typical planners solve these problems over and over again from scratch. While this may be less of an issue for tasks that are not time critical, many of the tasks in industry are time critical, for example, a robot working at a conveyor. For such tasks it becomes important to develop planning algorithms that can learn how to plan fast from offline and online experiences. To this end, my group has recently introduced the concept of provably Constant-Time Motion Planning (CTMP) [14]. Planning algorithms that fall within the class of CTMP algorithms provide provable guarantees on being able to generate the next action to execute for the robot within a (small) constant time, independently of the complexity and the size of the problem. Furthermore, the provided action is guaranteed to lie on a path that leads to the desired achievement of the task. To achieve these guarantees, CTMP algorithms are typically constructed by pre-processing the planning problem offline to construct a compact representation for constant-time lookup of a solution in query time (i.e., during online planning). Within this paradigm, my group has developed graph search-based CTMP algorithms for several domains including planning for pick and place tasks in static environments [14], planning manipulation tasks along a moving conveyor [13], and planning for pick and place in semi-static environments [12].

- **Graph search-based planning with inaccurate models.** One of the major causes of brittleness in fully autonomous robotic solutions is the inaccuracies in the a priori models used for planning. A model predicting how the world evolves as the robot interacts with it is rarely perfect. At the same time, relying on learning the true model during execution can often be prohibitively slow, and in many cases, just plain infeasible. Instead, it is important to develop planning algorithms that adapt their behavior to avoid the found discrepancies in the model and to exploit the portions of the models that are known to be correct. To this end, my group has recently developed CMAX algorithm [35], a graph search algorithm that adapts its behavior by comparing the outcomes of execution with the model given by the graph-based representation of the planning problem and biasing the search to avoid discovered discrepancies in the model. CMAX provide strong guarantees on the ability to achieve a task despite the model being inaccurate. CMAX has also been just extended in several ways including integration with model-free techniques such as Q-learning[36] to provide guarantees on convergence to provably optimal solutions in case of repeated executions of a task [34].

Forthcoming research. As some of my latest research outcomes show, my group is heavily leaning towards the development of planning methodologies that combine search-based planning with learning both offline and online. Some of the topics I plan to pursue further include exploring the question of interleaving learning the "right" graph representation with learning how to search the graph. I believe that finding the "right" representation for planning is typically as important as developing or using the "right" graph search algorithm itself. I am interested in developing algorithms that learn what planning dimensions are needed, how to construct appropriate graphs, and how to search them to provide strong guarantees on task achievement despite not knowing in advance a complete and accurate model. As part of this work, I am also very interested in understanding how to develop graph search algorithms that can exploit computationally very expensive but powerful models such as physics-based simulators and use these to compensate for inaccurate but computationally much cheaper models typically used for planning.

Students and collaborators. Much of the progress I made towards achieving my research goals is due to the hard work of my students and successful collaborations both internally and externally. I have been collaborating extensively with many external researchers including Vijay Kumar (UPenn), Sven Koenig (USC), Dieter Fox (UWash), Pere Ridao (University of Girona), Jianbo Shi (UPenn), Steve LaValle (UIUC), Gaurav Sukhatme (USC), Maren Bennewitz (University of Freiburg), Lydia Kavraki (Rice University), Isaac Kaminer (Naval Postgraduate School) and a number of other researchers across the world. With all of them I have had either joint proposals or joint publications. Internally at CMU, I have collaborated on joint projects and/or proposals with many faculty members on campus and at NREC including Manuela Veloso, Tony Stentz, Howie Choset, Katya Sycara, Phillip Gibson, Herman Herman, Al Kelly, Drew Bagnell, Dimi Apostolopoulos, Oliver Kroemer, and others. In terms of the students, I have been very fortunate to have a group of students and researchers that are very smart and highly motivated. The group currently includes eight PhD students, three Masters student, and several undergraduate students. I have graduated 12 PhD students - eight advised solely by me, one co-advised with Drew Bagnell, one co-advised with Manuela Veloso, and two co-advised with Vijay Kumar (while at UPenn) - and 22 Masters students.

My lab, the Search-based Planning Lab, is home to a dual-arm mobile manipulator PR2 (Figure 1(d)), acquired through an ARO DURIP grant I received in 2011, a mobile manipulation platform Roman built by JPL and General Dynamics, a UBTech dual-arm humanoid robot, and a number of fully autonomous robots my group built including a segbot (Figure 2(d)), two hexarotor micro-aerial vehicles (Figure 2(d)) and three quadrotor micro-aerial vehicles. In addition, I have just received another DURIP grant and am in the process of acquiring an industrial ABB arm, another mobile manipulation platform consisting of a UR10 arm and a mobile base, and a Husky outdoor ground robot. I believe that it is the hard work of my group combined with the infrastructure and funding I have acquired that enable me to pursue my research goals - developing methodologies for fast, reliable, and robust real-time decision-making in robots and using these to build planners that enable complex robotic systems to perform challenging tasks autonomously.

References

[1] *Algorithm Examples: the largest collection of algorithms across 13 programming languages.* https://python.algorithmexamples.com/web/graphs/multi_heuristic_astar.html.

[2] S. Aine and M. Likhachev. Truncated incremental search: Faster replanning by exploiting suboptimality. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI)*, 2013.

[3] S. Aine, C. Sharma, and M. Likhachev. Learning to search more efficiently from experience: A multi-heuristic approach. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*, 2015.

[4] S. Aine, S. Swaminathan, V. Narayanan, V. Hwang, and M. Likhachev. Multi-heuristic A*. *International Journal of Robotics Research (IJRR)*, 2015.

[5] Bobby Anguelov. Video game pathfinding and improvements to discrete search on grid-based maps. Master's thesis, University of Pretoria, 2011.

[6] S. Chitta, B. Cohen, and M. Likhachev. Planning for autonomous door opening with a mobile manipulator. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2010.

[7] B. Cohen, S. Chitta, and M. Likhachev. Heuristic search-based planning for manipulation. *International Journal of Robotics Research (IJRR)*, 2013.

[8] K. Gochev, B. Cohen, J. Butzke, A. Safonova, and M. Likhachev. Path planning with adaptive dimensionality. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*, 2011.

[9] K. Gochev, A. Safonova, and M. Likhachev. Incremental planning with adaptive dimensionality. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, 2013.

[10] K. Gochev, A. Safonova, and M. Likhachev. Anytime tree-restoring weighted A* graph search. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*, 2014.

[11] F. Islam, V. Narayanan, and M. Likhachev. Dynamic multi-heuristic A*. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015.

[12] F. Islam, C. Paxton, C. Eppner, B. Peele, M. Likhachev, and D. Fox. Alternative paths planner (app) for provably fixed-time manipulation planning in semi-structured environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

[13] F. Islam, O. Salzman, A. Agarwal, and M. Likhachev. Provably constant-time planning and replanning for real-time grasping objects off a conveyor belt. *International Journal of Robotics Research (IJRR)*, 2021.

[14] F. Islam, O. Salzman, and M. Likhachev. Provable indefinite-horizon real-time planning for repetitive tasks. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, 2019.

[15] Mubbasir Kapadia and Norman Badler. Navigation and steering for autonomous virtual humans. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4, 05 2013.

[16] S. Koenig and M. Likhachev. D* Lite. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI)*, 2002.

[17] S. Koenig and M. Likhachev. Real-time adaptive A*. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2006.

[18] S. Koenig, M. Likhachev, and D. Furcy. Lifelong planning A*. *Artificial Intelligence Journal (AIJ)*, 2004.

[19] Carina Krafft. *Implementation and comparison of pathfinding algorithms in a dynamic 3D space.* University of Applied Sciences Hamburg, Undergraduate thesis, 2019.

[20] M. Likhachev and D. Ferguson. Planning long dynamically-feasible maneuvers for autonomous vehicles. *International Journal of Robotics Research (IJRR)*, 2009.

[21] M. Likhachev, D. Ferguson, G. Gordon, A. Stentz, and S. Thrun. Anytime search in dynamic graphs. *Artificial Intelligence Journal (AIJ)*, 172(14), 2008.

[22] M. Likhachev, G. Gordon, and S. Thrun. ARA*: Anytime A* with provable bounds on sub-optimality. In *Advances in Neural Information Processing Systems (NIPS) 16*. Cambridge, MA: MIT Press, 2003.

[23] B. MacAllister, A. Kushleyev, J. Butzke, and M. Likhachev. Path planning for non-circular micro aerial vehicles in constrained environments. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2013.

[24] Brandon M. Malone. *Learning Optimal Bayesian Networks with Heuristic Search*. PhD thesis, Mississippi State University, 2012.

[25] V. Narayanan and Maxim Likhachev. Discriminatively-guided deliberative perception for pose estimation of multiple 3d object instances. In *Proceedings of the Robotics: Science and Systems Conference (RSS)*, 2016.

[26] V. Narayanan and Maxim Likhachev. Perch: Perception via search for multi-object recognition and localization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2016.

[27] V. Narayanan and Maxim Likhachev. Deliberative object pose estimation in clutter. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

[28] V. Narayanan, M. Phillips, and M. Likhachev. Anytime safe interval path planning for dynamic environments. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, 2012.

[29] M. Phillips, B. Cohen, S. Chitta, and M. Likhachev. E-graphs: Bootstrapping planning with experience graphs. In *Proceedings of the Robotics: Science and Systems Conference (RSS)*, 2012.

[30] M. Phillips, V. Hwang, S. Chitta, and M. Likhachev. Learning to plan for constrained manipulation from demonstrations. *Autonomous Robots (AURO)*, 2015.

[31] M. Phillips and M. Likhachev. Planning in domains with cost function dependent actions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2011.

[32] M. Phillips, V. Narayanan, S. Aine, and M. Likhachev. Efficient search with an ensemble of heuristics. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.

[33] A. Safonova and J. K. Hodgins. Construction and optimal search of interpolated motion graphs. In *ACM Transactions on Graphics Journal, SIGGRAPH Proceedings*, 2007.

[34] A. Vemula, J. A. Bagnell, and M. Likhachev. Cmax++: Leveraging experience in planning and execution using inaccurate models. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2021.

[35] A. Vemula, Y. Oza, J. A. Bagnell, and M. Likhachev. Planning and execution using inaccurate models with provable guarantees. In *Proceedings of the Robotics: Science and Systems Conference (RSS)*, 2020.

[36] C. Watkins and P Dayan. Q-learning. *Machine Learning*, 8(3–4):279–292, 1992.