Research Statement | Zachary Kingston

ROBONAUT 2 (or R2, Fig. 1) is an endless well of inspiration: a complex and dexterous humanoid robot built by NASA to demonstrate the potential of robotic helpers in the dangers of space—this domain truly lets robots shine; outer space is filled with endless tedium and not fit for living beings. The problems I discovered working with R2 influenced how I envision the future of robotics and what limits us from a future with collaborative, fluent, and flexible agents that could be thrown into far-flung reaches of space or at problems closer to home.

In my research, I have worked on three critical limitations centered around the essential problem of *motion planning*: how to compute feasible motion for a robot to follow from a high-level specification. For R2, advances were necessary in the capability of planning algorithms to (1) handle the complexities of interacting with the environment, (2) to be sequenced together into long-horizon plans with multiple steps, and (3) to be applied to robots like R2—this system was too complicated, too flexible, and faced with problems more difficult than those previously considered. To this end, I have focused on developing unifying theories to tackle *previously intractable* problems efficiently through novel algorithmic frameworks, theoretical advances, and efficient and extensible implementations, turning the intractable tractable.

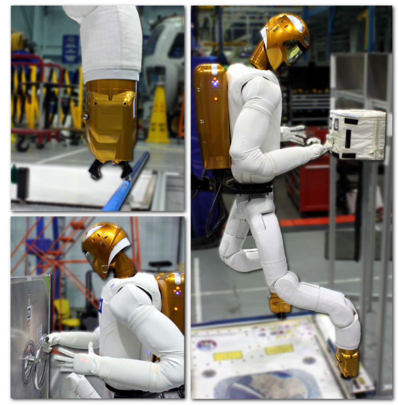

Fig. 1. R2 and a variety of tasks at NASA JSC [1].

§ Current and Ongoing Work

Planning to Satisfy Task Constraints Consider what R2 has to accomplish in Fig. 1: grasping a handrail requires keeping the leg facing the rail, the valve must rotate about its axis, and removing the cargo bag requires pulling straight outwards while keeping both feet planted. All of these require motion planning that respects the *constraints* of the task to find a feasible plan—by default, *sampling-based planning techniques* can be quick to find motions for high-dimensional (*i.e.*, ≥7 DoF) systems but do not consider task constraints. Adhering to these constraints is non-trivial: while the constraint may be expressed naturally in the Euclidean space around us, how this translates onto a robot's joints can be complicated and non-linear. To address these problems for high-dimensional systems such as R2, I developed a framework that unified the prior theory of *manifold-constrained sampling-based motion planning* into a core representation: an implicitly-defined space that facilitates any sampling-based method to plan with manifold constraints [2, 3]. For example, this representation permits specialized planners designed for high-dimensional problems, shown on the 168-degree-of-freedom system in Fig. 2, and is still used to move R2. While developing this framework, I also published a review article that has become a central reference in task-constrained planning [4] and contributed an invited chapter to the Encyclopedia of Robotics on *"Planning Under Manifold Constraints"* [5]. Furthering this work, I am investigating planning under other implicit constraints, *e.g.*, learned neural representations.

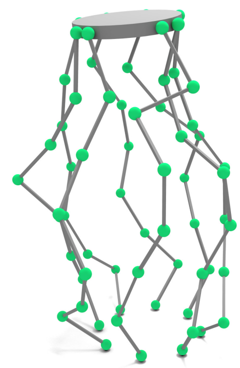

Fig. 2. Planning for a 168-DoF parallel robot.

Long-Horizon Locomotion and Manipulation Beyond singular constraints is the question of long-horizon reasoning—how does R2 know to walk across the handrails of the ISS, turn the valve, open the hatch, and finally pull the cargo bag from the rack? Sequencing and searching through the multitude of possible constraints results in the far more complex problem of *task and motion planning* (TAMP), which interleaves discrete decision-making and continuous motion planning—I collaborated on a key feedback-based approach to this problem [6, 7]. This feedback between the discrete and continuous is requisite to effective TAMP; *e.g.*, a bag cannot be reached if it is too far away, or a step cannot be traversed if an obstacle is in the way. However, using feedback in problems with continuous action choices is difficult: for example, where along the handrail should R2 grasp if it wants to continue, and how does this inform future choices? I made contributions to *multi-modal motion planning*, a class of methods for TAMP, that enabled long-horizon reasoning for systems with heavy geometric constraints by leveraging my manifold-constrained planning framework. The insights of this work showed how these methods could learn by reusing information from "similar" problems to accelerate search [8] and by informing what actions and motions should be attempted next [9]. These led to a planner capable of solving terribly hard problems: efficient autonomous manipulation (Fig. 3) and guiding R2 to walk across handrails (Fig. 4) [10].

I have extended my research in TAMP by mentoring younger students, collaborating on extensions to my prior work: how to improve the capabilities of the algorithm to solve problems that were too difficult to specify or solve before, and to efficiently solve these problems. Together, I have investigated using simulation to infer what actions are physically possible [11], how to use optimization to extract constraints from infeasible subsystems for highly-cluttered tasks [12], and finally, in collaboration advising a student at TU Berlin, how to find better solutions using novel space representations [13]. I plan to continue expanding the complexity and scope of problems these methods can solve, *e.g.*, multiple robots, execution robustness, and more.

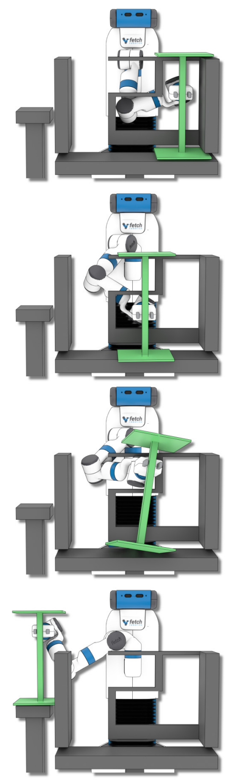

Fig. 3. Removing an object from a puzzle box; regrasping is required.

Simplifying, Democratizing, and Improving Robot Software I have a deep and vested interest in developing effective and efficient software for robotics research. I am the maintainer of the popular open-source Open Motion Planning Library (OMPL)[1], to which I contributed my manifold-constrained planning framework; this framework has been broadly adopted, *e.g.*, integrated into *MoveIt*, the *de facto* motion planner for ROS. A tool I developed to more fluently use *MoveIt*, *Robowflex* [14], was nominated for best paper in industrial robotics research for practicality at the IEEE/RSJ *International Conference on Intelligent Robots and Systems* (IROS) in 2022, one of the largest conferences in robotics. Recently, I collaborated on a hardware-accelerated planner [15] that outperforms state-of-the-art methods by several orders of magnitude, highlighting how integrating theory and implementation can bring about powerful advances. I also was invited to the NSF-funded Software Engineering for Robotics Workshop[2], which brought together experts to further the fundamental infrastructure of robotics.

Beyond implementation, proper evaluation requires good benchmarks and datasets. I worked with several universities to develop a widely-adopted dataset generator [16] and collaborated on a method to tune motion planning algorithms so non-experts can choose good configurations for evaluation [17]. Inspired by the above, I co-organized a successful workshop[3] at IROS 2022 to proselytize and discuss the evaluation of motion planning and give a venue to share benchmarks and tools that are not typically published.

§ Future Directions

I am confident in my continued ability to contribute to the state-of-the-art in robotics. R2 was the first step in removing the limitations that prevent robots from being effective, useful, and welcome. I have begun work in concrete directions that have revealed limitations arising from the computational constraints of the available hardware, the ability of robots to be life-long learners and adapt to novel environments, and the limitations of our sociotechnical perspective on the role robots have in our lives.

Rugged, Robust, and Real-time It is an unfortunate trend that the power and computational requirements of algorithms and learned models continue to increase exponentially as industry players push the envelope to deliver impressive results that research labs and individuals struggle to reproduce. In contrast, I want to push requirements lower toward rugged systems with limited computation, power, and time before an answer is necessary to keep the system safe. For example, robots in space will not have access to large clusters or the cloud—certifying algorithms to run on NASA's high-performance spaceflight computer[4] will require innovation

[1] https://ompl.kavrakilab.org/ [2] https://se4robotics.github.io/
[3] https://motion-planning-workshop.kavrakilab.org/ [4] https://www.nasa.gov/hpsc/

in both algorithms and software, using machine-sympathetic design to leverage the system's capabilities. Ongoing work of mine uses novel memory layouts and SIMD-accelerated subroutines for sampling-based planning, delivering performance improvements several orders of magnitude over baselines even on system-on-chip computers like a Raspberry Pi [15]. I plan to expand this work in several directions: for kinodynamic systems, uncertainty-aware planning, dynamic and changing environments, and for integrating learned heuristics to deliver the reliable planning necessary for resource-constrained systems. I am confident this work will enable my dream of deploying low-power systems in the wild (or even space), demonstrating advanced manipulation on platforms that were previously impossibly limited in capability.

Learned Heuristics and Implicit Representations My work has shown me the key to improving autonomy is not simply to improve the planner's capabilities; the planner must also learn from experience—a robot should be able to be a lifelong student. Two areas of interest to me are improving planning heuristics based on prior experience and learning abstractions or specifications to handle the inherent uncertainty of modeling the world. Learning heuristics and samplers for motion planners is a burgeoning area in which I have collaborated [18] (nominated best paper in cognitive robotics at the IEEE *International Conference of Robotics and Automation*, one of the premier robotics conferences). There has also been an explosion in learning implicit models of the world, given the advent of, *e.g.*, NERFs. A recent collaboration [19] has investigated using these models for quantifying the risk of collision with the environment. Models such as these could be extended to many other cases, for example, quantifying privacy, the danger of traversing rugged terrain, or modeling task constraints. There are endless opportunities for integrating these models in neurosymbolic frameworks to generalize planning capabilities to problems previously impossible, and there are interesting algorithmic and implementation questions on how to leverage these learned models best, understand their limitations, and leverage hardware for efficient implementations.

Fig. 4. R2 autonomously climbing handrails.

Privacy, Sensing, and Human Interaction Consider the outcome of improving robot autonomy; robotics is rapidly maturing, and systems are entering the home, workplace, and more, with ethical considerations trailing behind algorithmic advances. It is crucial to think of how the algorithms that imbue systems with intelligence are used, as it is not enough to be agnostic to human concerns; in this case, they will (un)intentionally be exploited by malicious, unethical, or ignorant actors, *e.g.*, by violating privacy, as investigated by my work [20]. This work showed that, with minimal changes, existing software could be used to covertly observe people while still achieving task objectives. I am pursuing the foundations laid by this work to perform an interdisciplinary investigation into the broader sociotechnical questions of robotics and privacy. Moreover, this work highlighted *planning while (not) observing*, which I wish to investigate further, *e.g.*, by extending my work on long-horizon planning to consider (un)observation. There are also exciting directions in considering human behavior (*e.g.*, [21])—how should robot team members act around their human partners?

§ Collaboration Opportunities and Funding

From my fruitful collaborations between members of my lab and other institutions (*e.g.*, NASA JSC, TU Berlin, ANU, Houston Methodist) I am positive I can contribute my expertise and connect my research to other faculty members, further improving the state-of-the-art and bringing planning over new horizons. For funding, I was awarded the NSF Graduate Student Research Fellowship, the NASA Space Technology Research Fellowship, and have co-authored successful NSF grants based on my research, *totaling $500,000 in funding*. I plan to apply for the NSF Career Award and to submit proposals to the NSF, NASA (including the NASA Early Career Faculty award), and other government sources. I will also strengthen my ties to cutting-edge robotics companies for industrial collaborations. Altogether, I will continue pushing the boundaries of robotic planning beyond current limits through advances in theory, algorithms, and implementation to bring deliberative, intelligent agents into reality.

§ References

1. W. Baker, **Zachary Kingston**, M. Moll, J. Badger, and L. E. Kavraki. "Robonaut 2 and You: Specifying and Executing Complex Operations". In: *IEEE Workshop on Advanced Robotics and its Social Impacts*. Austin, TX, Mar. 2017, pp. 1–8. DOI: 10.1109/ARSO.2017.8025204.
2. **Zachary Kingston**, M. Moll, and L. E. Kavraki. "Decoupling Constraints from Sampling-Based Planners". In: *Robotics Research*. Ed. by N. M. Amato, G. Hager, S. Thomas, and M. Torres-Torriti. Cham: Springer International Publishing, 2020, pp. 913–928. DOI: 10.1007/978-3-030-28619-4_62.
3. **Zachary Kingston**, M. Moll, and L. E. Kavraki. "Exploring Implicit Spaces for Constrained Sampling-Based Planning". In: *The International Journal of Robotics Research* 38.10–11 (2019), pp. 1151–1178. DOI: 10.1177/0278364919868530.
4. **Zachary Kingston**, M. Moll, and L. E. Kavraki. "Sampling-Based Methods for Motion Planning with Constraints". In: *Annual Review of Control, Robotics, and Autonomous Systems* 1.1 (2018), pp. 159–185. DOI: 10.1146/annurev-control-060117-105226.
5. **Zachary Kingston**. "Encyclopedia of Robotics". In: ed. by M. H. Ang, O. Khatib, and B. Siciliano. Springer Berlin Heidelberg, 2020. Chap. Planning Under Manifold Constraints, pp. 1–9. DOI: 10.1007/978-3-642-41610-1_174-1.
6. N. T. Dantam, **Zachary Kingston**, S. Chaudhuri, and L. E. Kavraki. "Incremental Task and Motion Planning: A Constraint-Based Approach". In: *Robotics: Science and Systems*. Ann Arbor, MI, June 2016. DOI: 10.15607/RSS.2016.XII.002.
7. N. T. Dantam, **Zachary Kingston**, S. Chaudhuri, and L. E. Kavraki. "An Incremental Constraint-Based Framework for Task and Motion Planning". In: *The International Journal of Robotics Research* 37.10 (2018), pp. 1134–1151. DOI: 10.1177/0278364918761570.
8. **Zachary Kingston**, C. Chamzas, and L. E. Kavraki. "Using Experience to Improve Constrained Planning on Foliations for Multi-Modal Problems". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2021, pp. 6922–6927. DOI: 10.1109/IROS51168.2021.9636236.
9. **Zachary Kingston**, A. M. Wells, M. Moll, and L. E. Kavraki. "Informing Multi-Modal Planning with Synergistic Discrete Leads". In: *IEEE International Conference on Robotics and Automation*. 2020, pp. 3199–3205. DOI: 10.1109/ICRA40945.2020.9197545.
10. **Zachary Kingston** and L. E. Kavraki. "Scaling Multi-Modal Planning: Using Experience and Informing Discrete Search". In: *IEEE Transactions on Robotics* 39.1 (2023), pp. 128–146. DOI: 10.1109/TRO.2022.3197080.
11. Y. Lee, W. Thomason, **Zachary Kingston**, and L. E. Kavraki. "Object Reconfiguration with Simulation-Derived Feasible Actions". In: *IEEE International Conference on Robotics and Automation*. 2023, pp. 8104–8111. DOI: 10.1109/ICRA48891.2023.10160377.
12. C. Quintero-Peña, **Zachary Kingston**, T. Pan, R. Shome, A. Kyrillidis, and L. E. Kavraki. "Optimal Grasps and Placements for Task and Motion Planning in Clutter". In: *IEEE International Conference on Robotics and Automation*. 2023, pp. 3707–3713. DOI: 10.1109/ICRA48891.2023.10161455.
13. S. B. Bayraktar, A. Orthey, **Zachary Kingston**, M. Toussaint, and L. E. Kavraki. "Solving Rearrangement Puzzles using Path Defragmentation in Factored State Spaces". In: 8.8 (2023), pp. 4529–4536. DOI: 10.1109/LRA.2023.3282788.
14. **Zachary Kingston** and L. E. Kavraki. "Robowflex: Robot Motion Planning with *MoveIt* Made Easy". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2022, pp. 3108–3114. DOI: 10.1109/IROS47612.2022.9981698.
15. W. Thomason†, **Zachary Kingston**†, and L. E. Kavraki. *Motions in Microseconds via Vectorized Sampling-Based Planning*. † Equal Contribution. Under Review. 2024.
16. C. Chamzas, C. Quintero-Peña, **Zachary Kingston**, A. Orthey, D. Rakita, M. Gleicher, M. Toussaint, and L. E. Kavraki. "MotionBenchMaker: A Tool to Generate and Benchmark Motion Planning Datasets". In: *IEEE Robotics and Automation Letters* 7.2 (2021), pp. 882–889. DOI: 10.1109/LRA.2021.3133603.
17. M. Moll, C. Chamzas, **Zachary Kingston**, and L. E. Kavraki. "HyperPlan: A Framework for Motion Planning Algorithm Selection and Parameter Optimization". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2021, pp. 2511–2518. DOI: 10.1109/IROS51168.2021.9636651.
18. C. Chamzas, **Zachary Kingston**, C. Quintero-Peña, A. Shrivastava, and L. E. Kavraki. "Learning Sampling Distributions Using Local 3D Workspace Decompositions for Motion Planning in High Dimensions". In: *IEEE International Conference on Robotics and Automation*. 2021, pp. 1283–1289. DOI: 10.1109/ICRA48506.2021.9561104.
19. C. Quintero-Peña, W. Thomason, **Zachary Kingston**, A. Kyrillidis, and L. E. Kavraki. *Stochastic Implicit Neural Signed Distance Functions for Safe Motion Planning under Sensing Uncertainty*. Under Review. 2024.
20. R. Shome, **Zachary Kingston**, and L. E. Kavraki. "Robots as AI Double Agents: Privacy in Motion Planning". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. To Appear. 2023.
21. A. M. Wells, **Zachary Kingston**, M. Lahijanian, L. E. Kavraki, and M. Y. Vardi. "Finite Horizon Synthesis for Probabilistic Manipulation Domains". In: *IEEE International Conference on Robotics and Automation* (2021), pp. 6336–6342. DOI: 10.1109/ICRA48506.2021.9561297.