Xuesu Xiao

Research Statement

Mobile robots have the unrealized potential to assist or substitute for human rescuers after disasters during initial response, restoration, reconstruction, and betterment. My research is motivated by the goal of enabling robots to dramatically improve our ability to mount such after-disaster missions quickly and safely, so as to maximize our ability to save victims, restore basic facilities, reconstruct infrastructures, and improve preparedness for future disasters, while minimizing the risk to rescuers. To accomplish this objective, future mobile robots need to be (1) highly capable of reliably moving through those challenging and most likely adversarial environments, and (2) highly intelligent so that involvement of human rescuers, both physically and intellectually, can be effectively minimized. I envision future after-disaster missions to be efficiently conducted by fully autonomous robots, which are aware of the risk and constraints from the field, collaborative with other robotic (or human) teammates, adaptive and robust when facing new scenarios, and eventually achieve high performance in real-world autonomous deployment.

However, current disaster robots still lack such capabilities. For example, the ongoing decommissioning in response to the Fukushima Daiichi nuclear disaster still completely relies on teleoperation: Multiple human rescuers have to slowly and cautiously drive a robot due to mobility and manipulation challenges. To make matters worse, this practice even requires a second teleoperated visual assistant robot to give the operators a better external viewpoint, causing problems such as difficulty in coordination between teleoperators and manually-chosen suboptimal viewpoints. These current practices are inefficient and require extensive human involvement.

Disaster is among many of the potential applications of mobile robots, which also include automated agriculture, infrastructure inspection, and scientific exploration. Motivated by all these applications and the current status of how robots are being used, my research goal is to **develop highly capable and intelligent mobile robots that are robustly deployable in the real world with minimal human supervision**. As a roboticist with unique expertise evenly grounded in motion planning and machine learning, and vast experience working on real-world problems in the field with disaster responders, I build advanced robot platforms, develop complex sensing and actuation systems, design sophisticated motion planning algorithms, and set up standardized testbeds and metrics in order to create highly capable and intelligent robots to locomote on land, in air, and at sea.

Risk-Aware, Constraint-Oriented, Collaborative Robotics Deployed in Humanitarian Crises (Pre-Doctoral)

Achieving my research goal first requires creating physical robots with advanced intelligence and locomotion capabilities which can reason about real-world challenges and adversaries, address constraints imposed by real-world applications, and cooperate with robot teammates with different motion and sensing modalities. These research topics have been addressed by my Ph.D. thesis and other pre-doctoral research.

Aiming at automating and robustifying the aforementioned teleoperated visual assistance in the Fukushima Daiichi nuclear decommissioning, my Ph.D. thesis was to develop **a robot motion risk reasoning framework for unstructured or confined spaces, a risk-aware motion planner, and an entire motion suite for a tethered aerial visual assistant** under challenging motion constraints (Fig. 1). At least two challenges prevent an autonomous aerial visual assistant from being applied in Fukushima: (1) flying in those unstructured or confined spaces entails extensive motion risk, which most existing robots do not have the ability to reason about, and (2) energy considerations and safety precautions for such mission-critical tasks require tethered flight (tethered to the primary robot), as a real-world constraint imposed on existing free-flying robots. To address (1), I used propositional logic and probability theory to derive a reasoning framework for risk-awareness and discovered that the risk a moving robot faces is not simply a function of where the robot is, but also depends on its entire motion history [9]. This discovery contradicts most existing simplified motion risk/cost assumptions and makes the risk-aware planning problem PSPACE-complete. I also developed a risk-aware motion planning paradigm that can effectively trade off risk history and computation. Not only suitable for the visual assistance problem in Fukushima, the risk-awareness framework and risk-aware planner is also general to most mobile robots working in unstructured or confined spaces in the real world. To address (2), I developed a full motion suite for tethered aerial robots flying in cluttered spaces [6], including tether-based localization [13], motion primitives [7], tether contact planning [15], and

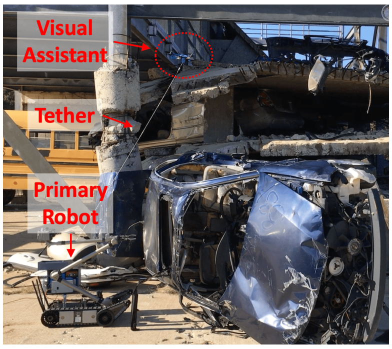

Fig. 1: Tethered Risk-Aware Visual Assistance in Disaster Environment

visual servoing [8]. This tethered motion suite opens up a new regime for resilient indoor aerial locomotion under energy and safety constraints stemming from real-world applications.

During my time as a Ph.D. student at Texas A&M University, I also developed other robotic technologies that were directly deployed in real-world search and rescue missions, including Hurricanes Harvey and Irma, the Greece refugee crisis, and other emergency response exercises world-wide. At the Center for Robot-Assisted Search and Rescue, I developed **cooperative motion planning techniques for a heterogeneous Unmanned Aerial/Surface Vehicle team (UAV/USV)**, in which a USV can fully autonomously navigate to drowning victims [17] with the overheard visual guidance from a UAV using motion-based viewpoint stabilization [1] (Fig. 2). The UAV/USV team has been deployed for marine mass casualty incident response in search and rescue exercises conducted by the United States Coast Guard and Galveston Fire Department during Summer Institute 2016 in Galveston, TX; Italian Coast Guard during 2016 exercise in Genoa, Italy; Brazos County Fire Department and Grimes County Emergency Management during Brazos Valley Search and Rescue Exercise 2017 in Gibbons Creek, TX; Los Angeles County Fire Department Lifeguards during 2017 exercise in Los Angeles, CA; and Department of Homeland Security during 2017 CAUSE V exercise in Bellingham, WA.

Fig. 2: UAV/USV Team in Hurricane Harvey Deployment

In the early stages of my research career at Carnegie Mellon University, I investigated **ground vehicle energetic models for long-range missions in remote spaces under real-world energy constraints** [3, 4]. I also researched **locomotive reduction techniques for hyper redundant locomotors** so that a snake robot with 16 Degrees of Freedom (DoFs) can be effectively controlled as a 2-DoF differential drive car to reach constrained spaces inaccessible to humans and conventional wheeled or treaded robot platforms [14]. The mechatronics background acquired through my undergraduate studies, where I built several robotics systems from scratch, has also given me the mindset of a problem-solver that can apply scientific research to address real-world mechatronics or robotics problems.

Machine Learning for Adaptive and Robust Autonomy with Minimal Human Involvement (Post-Doctoral)

After developing several highly capable mobile robots to reduce *physical* involvement of human rescuers in challenging or adversarial tasks, my real-world deployment experiences suggest that most mobile robots still lack sufficient intelligence to minimize human's *intellectual* involvement during deployment. To adapt to various deployment scenarios and to achieve robust performance, robots still require extensive expert knowledge in the form of manual teleoperation, parameter tuning, or human supervision as a "safety officer". I see potential in the state-of-the-art machine learning techniques, which can be used on top of the well-engineered systems so that modern robots can actively *learn*, instead of being passively *engineered*, to be robustly deployable in the real world. My post-doctoral research at the University of Texas at Austin therefore focuses on **using machine learning techniques in conjunction with classical motion planners** to improve autonomous robot locomotion through the combination of the best of both worlds, rather than creating pure engineering or end-to-end learning solutions. Classical methods enhanced with learning enjoy crucial benefits such as safety and explainability, both of which are important properties for mobile robots interacting with the real world.

Adaptive Planner Parameter Learning (APPL) is one contribution of my post-doctoral research, in which I used machine learning to create highly intelligent robots and therefore to reduce human involvement during deployment. Most classical motion planners are capable of moving robots from one point to another in a safe and explainable manner. However, when facing new environments, the planner parameters (e.g. maximum speed, sampling rate, inflation radius) need to be properly fine-tuned by a robotics expert. Therefore human expert knowledge is required onsite during deployment. In this line of work, I devised **learning components, which interact with an underlying classical motion planner through automatic parameter tuning**. These learning components are able to learn how to dynamically adjust planner parameters based on the current situation through a non-expert human demonstration (APPLD, Fig. 3) [12] or even a small amount of interventions (APPLI) [5] through Imitation Learning. Furthermore, I proposed the concept of a *parameter policy*, which works in a meta-environment composed of both the physical world and the underlying classical planner under a Markov Decision Process framework. I used Reinforcement Learning to train the parameter policy so it can select appropriate planner parameters at each time step and reason about the future consequences of using these parameters (APPLR) [18]. These variants of the APPL paradigm

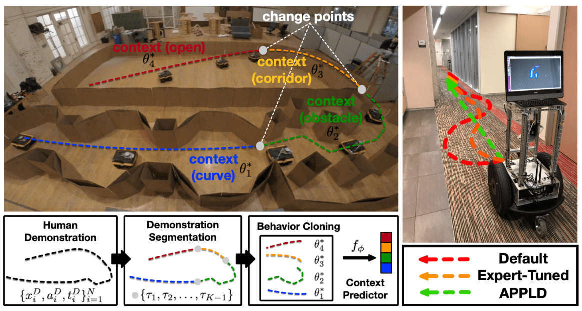

Fig. 3: APPLD on Two Robots

have been implemented on different classical navigation planners on different ground mobile robots and have achieved better navigation performance than pure classical and end-to-end learning approaches. From simply requiring a non-expert demonstration (APPLD), to needing only four human interventions (APPLI), to finally not relying on any human interaction at all (APPLR), APPL is a general paradigm to combine the benefits of emerging learning and classical planning approaches to gradually minimize human involvement.

Another research thrust in minimizing human involvement by combining classical planning with learning is to use learning to enable performance improvement from past (suboptimal) experiences achieved by classical motion planners. Improvement from experience is not possible for most classical approaches, especially without expert knowledge, so they tend to repeat the same mistake, regardless of how many times they have executed the same suboptimal motion in the same situation. I developed **a Lifelong Learning for Navigation (LLfN) framework [2] that allows mobile robots to achieve self-supervised in-environment improvement (with their own unsupervised experience), and cross-environment adaptation (without the notorious catastrophic forgetting for many learning systems)**. LLfN is able to self-identify suboptimal motion plans and gradually eliminate those plans through learning from similar self-supervised experience. In contrast to learning an end-to-end motion planner from scratch with hours of training time and millions of training data/steps, LLfN leverages classical motion planners and is therefore extremely efficient. It is the first learning for navigation approach that is completely implemented onboard a mobile robot with very limited onboard memory and computational resources during deployment. By allowing robots to improve from their own mistakes and experiences in and across different environments, human involvement is no longer necessary when robots encounter problems during deployment.

The aforementioned postdoctoral research work exposed the limitations of many classical motion planners. To outperform what classical motion planners can achieve, I formulated **a novel "dual" problem of motion planning called *hallucination***: instead of finding the optimal motion plan for a given obstacle configuration, the robot can easily *hallucinate* obstacle configurations, where a certain motion plan is optimal. Solving this relatively easier "dual" problem allows us to generate a lot of training data for learning algorithms and creates a new Learning from Hallucination (LfH) paradigm [16, 10] to learn high-performance motion planners. One conundrum of learning safe motion planners is that in order to produce safe motions in obstacle-occupied spaces, a robot needs to first learn in those dangerous spaces without the ability of planning safe motions. Therefore, it either requires a good demonstration (e.g. from a classical planner or a human), or exploration based on trial-and-error, both of which become costly in highly-constrained and therefore dangerous spaces. LfH addressed this problem by allowing the robot to safely explore in a completely open environment without any obstacles, and to *hallucinate* the obstacle configurations to the robot perception, where the motions executed in the open training environment are optimal. In this way, a lot of training data can be generated from which a motion planner can be effectively learned. The learned motion planner is combined with classical global planning and model predictive collision checking to assure safety during deployment. It can produce agile maneuvers in highly-constrained spaces without slowing down or requiring extra computation, as most classical motion planners do. In addition to minimizing human involvement during *deployment*, LfH further reduces engineering effort during *development* by allowing the robot to *learn* a motion planner in a safe manner all by itself.

Despite roboticists' devotion to developing fully autonomous robots, most state-of-the-art robots still require human supervision, i.e. a human "safety officer" during real-world deployment. From my point of view, the best remedy for this lack of confidence in deploying fully autonomous robots in the real world is through extensive testing. Therefore, I am an advocate for **creating standardized testing methods and metrics to objectively benchmark robot performance and research progress**. Although the ultimate touchstone for good robotics research is the challenging and adversarial real world without any human supervision, e.g. remote areas, inaccessible environments, and search and rescue scenarios, an intermediate testbed for the real world is a general, comprehensive, and systematic benchmark with a set of unbiased metrics, motivated by the real-world use case. With my experience on snake robots, I conducted a survey on current (ad hoc) testbeds for snake robot locomotion and provided recommendations for building a general-purpose testbed [11]. I also built a suite of 300 (simulated) navigational environments to benchmark ground navigation capabilities, which is publicly available for use by the whole research community. Since these benchmarks and metrics serve as an accurate proxy for the real world and are independent of specific robot platforms and motion planners, robots performing well on these tests come closer to adaptive and robust autonomy during deployment.

Reliable, Resilient, Task-Efficient Robots towards Full Autonomy (Future Research Agenda)

Building upon my past research experience, my future work will be continually driven by my research goal to develop highly capable and intelligent mobile robots that are robustly deployable in the real world with minimal human supervision. Although my previous research on risk-aware, constraint-oriented, collaborative robots and machine learning for adaptive robust autonomy has led to many robots that have already been deployed in disaster applications, most robots are still at the "proof-of-concept" stage and are therefore far from being a mature solution to improving human capabilities in real-world missions. I posit three major reasons that currently preclude wide adoption of autonomous robots in dirty, dull, and dangerous environments: unreliability, inefficiency, and human-dependency. These three problems are the driving forces for my future research.

First, most autonomous robots are still not reliable in the field, in contrast to the controlled lab environments where these robots are developed. This unreliability is not only due to hardware or engineering challenges, but also due to a lack of consideration and quantification of real-world challenges and adversaries, and more importantly, a lack of resilience to problems or failures when they occur. My previous work on risk reasoning only scratched the surface of how to consider and quantify those challenges and adversaries. Furthermore, only reasoning about risk is not enough. One must also consider how to minimize risk (partially addressed by my previous risk-aware planning work) and what to do *after* a risky event actually happens. While my constraint-oriented tethered motion suite represented a groundbreaking advance with respect to resilient indoor flight, it only aimed to achieve the most rudimentary level of resilience: to manually retrieve the UAV using the teleoperated ground robot after a crash. A truly resilient tethered UAV should be able to bootstrap the tether to recover from a collision or even a crash. My future research will be geared towards **developing hardened, reliable, and resilient robots**, both in terms of hardware and algorithms, that can reason about real-world challenges and adversaries, make intelligent decisions to reliably conduct their missions, and withstand and then recover quickly from difficult conditions. I will continue to adopt a similar field methodology to that of my pre-doctoral research to identify problems or failures robots encounter in the field, create resilient mechatronics solutions including novel perception, actuation, control, and mechanisms, and develop corresponding intelligence adaptive to the underlying robotic systems and real-world challenges and adversaries. Developing such reliability will be the first step towards creating mature and trustworthy robotic solutions for real-world missions.

Second, most autonomous robots are still inefficient at the task-level, in contrast to the mobility-level. My previous work has enabled a fleet of versatile mobile robots, which can achieve "faster, higher, and stronger" locomotion performance compared to human rescuers, or even reach places inaccessible to humans. However, creating high-performance locomotors is not the final goal of developing intelligent robots: they need to utilize their superior mobility to eventually accomplish different tasks. Currently, most autonomous task execution is still less efficient than manual teleoperation. Building upon the high mobility developed in my previous work, my future research will focus on **how to create highly capable and intelligent robots that are efficient at the task-level**. I will investigate task-level efficiency metrics beyond mobility, the relationship or trade off between task-execution and locomotion, as well as planning and learning techniques to achieve overall mission success. In addition to the traditional motion planning used in my previous work, I plan to utilize multi-model motion planning and task and motion planning to improve task-level efficiency. Notably, many of these tasks vary in different deployments, and my research will also create robots that are adaptive to different tasks by exploiting machine learning.

Third, most autonomous robots are still heavily dependent on human. For example, most robotic deployment still requires extensive setup, adjustment, tuning, or even teleoperation from humans onsite. I envision that in future robot deployments, a human will only need to "press a button and walk away", while the robot can conduct its mission "out-of-the-box" in a reliable and efficient manner. To approach this vision, my previous APPL line of work investigated how human-dependency can be minimized through machine learning to achieve high mobility performance during deployment. In my future work, I plan to investigate **methods that most efficiently utilize different modalities of human interactions under a Cycle-of-Learning scheme to gradually achieve full autonomy**. Teleoperation, expert tuning, human demonstration, corrective interventions, and evaluative feedback should not target separated deployment scenarios, while reinforcement learning through trial-and-error should not be isolated from different modalities of human interactions either. An intelligent robot that aims at eventually achieving full autonomy without any human-dependency needs to effectively utilize every piece of precious human interaction along the way towards full autonomy and learn from it to benefit all future deployments. The learned knowledge can take the form of an expanding library of parameters, an increasing set of motion primitives, an improving motion or parameter policy, or a task-level knowledge graph. I will investigate a systematic methodology to leverage every single human interaction and a comprehensive knowledge representation to improve capability and intelligence towards achieving full autonomy with techniques such as continual/lifelong learning, transfer learning, curriculum learning, and a combination of imitation with reinforcement learning.

In summary, I see myself as a robotics researcher who uses scientific approaches to solve real-world robotics problems, in order to push the boundary of how robots can be used to work in the real world on behalf of human. Building on the history of my past achievements, I am looking forward to pursuing the numerous goals I laid out for future research in collaboration with other roboticists from Mechanical, Electrical, Aerospace Engineering, and Computer Science.

References

[1] Jan Dufek, Xuesu Xiao, and Robin Murphy. "Visual pose stabilization of tethered small unmanned aerial system to assist drowning victim recovery". In: *2017 IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*. IEEE. 2017, pp. 116–122.

[2] Bo Liu, Xuesu Xiao, and Peter Stone. "Lifelong Navigation". In: *arXiv preprint arXiv:2007.14486* (2020).

[3] Kshitij Tiwari, Xuesu Xiao, and Nak Young Chong. "Estimating achievable range of ground robots operating on single battery discharge for operational efficacy amelioration". In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2018, pp. 3991–3998.

[4] Kshitij Tiwari et al. "A unified framework for operational range estimation of mobile robots operating on a single discharge to avoid complete immobilization". In: *Mechatronics* 57 (2019), pp. 173–187.

[5] Zizhao Wang et al. "APPLI: Adaptive Planner Parameter Learning From Interventions". In: *arXiv preprint arXiv:2011.00400* (2020).

[6] Xuesu Xiao, Jan Dufek, and Robin Murphy. "Autonomous visual assistance for robot operations using a tethered uav". In: *12th Conference on Field and Service Robotics (FSR)*. 2019.

[7] Xuesu Xiao, Jan Dufek, and Robin Murphy. "Benchmarking tether-based uav motion primitives". In: *2019 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2019, pp. 51–55.

[8] Xuesu Xiao, Jan Dufek, and Robin Murphy. "Visual servoing for teleoperation using a tethered uav". In: *2017 IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*. IEEE. 2017, pp. 147–152.

[9] Xuesu Xiao, Jan Dufek, and Robin R Murphy. "Robot risk-awareness by formal risk reasoning and planning". In: *IEEE Robotics and Automation Letters* 5.2 (2020), pp. 2856–2863.

[10] Xuesu Xiao, Bo Liu, and Peter Stone. "Agile Robot Navigation through Hallucinated Learning and Sober Deployment". In: *arXiv preprint arXiv:2010.08098* (2020).

[11] Xuesu Xiao and Robin Murphy. "A review on snake robot testbeds in granular and restricted maneuverability spaces". In: *Robotics and Autonomous Systems* 110 (2018), pp. 160–172.

[12] Xuesu Xiao et al. "APPLD: Adaptive Planner Parameter Learning From Demonstration". In: *IEEE Robotics and Automation Letters* 5.3 (2020), pp. 4541–4547.

[13] Xuesu Xiao et al. "Indoor uav localization using a tether". In: *2018 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2018, pp. 1–6.

[14] Xuesu Xiao et al. "Locomotive reduction for snake robots". In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2015, pp. 3735–3740.

[15] Xuesu Xiao et al. "Motion planning for a uav with a straight or kinked tether". In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2018, pp. 8486–8492.

[16] Xuesu Xiao et al. "Toward Agile Maneuvers in Highly Constrained Spaces: Learning from Hallucination". In: *arXiv preprint arXiv:2007.14479* (2020).

[17] Xuesu Xiao et al. "UAV assisted USV visual navigation for marine mass casualty incident response". In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2017, pp. 6105–6110.

[18] Zifan Xu et al. "APPLR: Adaptive Planner Parameter Learning from Reinforcement". In: *arXiv preprint arXiv:2011.00397* (2020).