Certifiable Robot Autonomy Under Uncertainty

Kehan Long
University of California, San Diego k3long@ucsd.edu

I. Introduction and Motivation

Guaranteeing safe and efficient control for autonomous robotic systems is a long-standing challenge. Unlike humans, who rely on intuition and experience without explicit models, robots can leverage multi-modal sensor data, prior knowledge of their own dynamics, and computational reasoning to make informed decisions. This capability has the potential to not only enable reliable performance but also generate interpretable, robust, and theoretically grounded control policies. The overarching vision of my research is to design robots that match and ultimately surpass humans in terms of motion and manipulation capabilities in complex and dynamic environments, where real-time sensory data integration is essential, and various sources of uncertainty are inevitable.

Despite considerable progress in robotics and control theory, existing techniques for motion planning and control have inherent limitations. Classical model-based approaches (e.g., certificate-based and optimal control) offer strong theoretical guarantees but often struggle to scale to high-dimensional systems and complex environments. In contrast, data-driven methods (e.g., reinforcement learning (RL)) achieve impressive performance yet typically lack formal guarantees of safety, stability, and robustness, making their deployment in safety-critical applications challenging.

My research focuses on developing safe and robust control strategies that account for various sources of uncertainty, including model inaccuracies, sensor noise, state estimation errors, and uncertainties arising from neural representations, such as learned models of robot dynamics, environment, and robot geometry. I integrate control barrier functions (CBFs) and control Lyapunov functions (CLFs) with robust, probabilistic, and distributionally robust optimization to ensure real-time safety and stability under uncertainty. These methods have been validated on both ground mobile robots and 6-dimensional manipulators, demonstrating feasibility for real-world deployment (Fig. 1). At the intersection of learning and control, I develop learning-based control policies with formal stability guarantees. I utilize a generalized notion of Lyapunov function to certify modern RL policies by augmenting their value functions and enforcing multi-step decrease conditions

Ultimately, my research seeks to unify control-theoretic guarantees with learning-based adaptability, enabling high-performance autonomous robots in real-world environments while providing safety and stability guarantees.

II. Past and Current Research

My research merges control-theoretic rigor with learning-based representations to address the fundamental challenge of safe, efficient robotic **control under uncertainty**. In particular, I extend classical certificate functions (e.g., CBFs, CLFs) to robust, probabilistic, and distributionally robust formulations, and demonstrate their efficacy across various robotic systems in complex and dynamic environments.

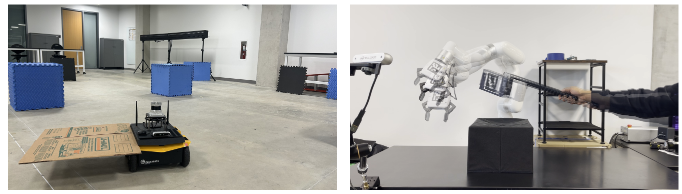

(a) Ground Robot with Complex Geometry (b) 6-DoF Manipulator Navigation

Fig. 1: Demonstrations of safe control under uncertainty for a ground mobile robot and a 6-DoF manipulator, both utilizing neural shape representations of their bodies.

A. Safe Control under Uncertainty

Ensuring safe control under uncertainty remains a core challenge for autonomous robot systems. Control barrier functions (CBFs) provide a principled, **real-time** method for enforcing safety in control-affine systems by solving a quadratic program that adjusts nominal inputs to maintain forward invariance of a safe set [3, 2]. Due to their computational efficiency, CBFs have been successfully deployed across robotic domains [48, 45, 39, 15]. However, most early work assumes that the CBF, system dynamics, and robot states are precisely known, which is often invalid in real-world deployments with noisy onboard sensors. My research addresses these limitations by systematically incorporating sensor noise, dynamic inaccuracies, and state estimation errors into the CBF framework.

A central challenge in safe control synthesis is obtaining accurate, efficient geometric representations of both the environment and the robot itself. In [28], I tackled this by using neural networks to learn a signed distance function (SDF) [41] of obstacles from onboard sensor data. Unlike conventional approaches that assume a priori known barriers, my method incrementally learns the SDF using range measurements. By explicitly taking the estimation error bounds of the learned SDF and its gradient into account, I formulated a robust control barrier constraint, which led to a novel **second-order cone programming** formulation for safe control synthesis. Building on this, I extended my research to consider uncertainty not only in environment perception but also in neural system dynamics approximations [29]. I developed both probabilistic and robust formulations of CBF constraints, allowing the control synthesis problem to handle uncertainty in both barrier function and system dynamics estimates.

Although recent work has considered uncertainty in CBF-based safety filters [6, 10, 46, 4, 9, 23, 1, 17], they often address a single source (e.g., model or barrier error) in isolation. In contrast, real-world robots face **compounded uncertainties** from various sources (e.g., localization, sensor, geometry estimations), which interact in nonlinear ways and are hard to model explicitly. This motivates my use of distributionally robust optimization (DRO), which handles multiple uncertainties without requiring precise bounds or distributions.

DRO [20, 12, 47] circumvents the need for explicit uncertainty models by operating directly on sampled data, such as LiDAR hits or states from standard estimators [21, 11]. By enforcing a chance constraint over a Wasserstein ambiguity set [19], it avoids restrictive distributional assumptions. At the same time, simplistic robot shapes (e.g., circles or spheres) often over-constrain feasible motions, particularly for manipulators. Recent work instead uses neural SDFs [24, 27, 33] and configuration-space distance functions [26] to capture complex geometries more accurately, but these **neural representations** introduce uncertainties that are difficult to quantify, reinforcing the need for a distributionally robust approach.

In my work [31, 36, 35], I propose a novel **distributionally robust control barrier constraint** that systematically accounts for uncertainties in state estimation, system dynamics, sensor measurements, and neural shape representations. Critically, the resulting safe control synthesis problem can be reformulated as a quadratic program. I validate this approach on both ground mobile robots and 6-dimensional manipulators (Fig. 1), demonstrating safe, efficient control under uncertainty in cluttered, dynamic environments.

B. Stability Certification for Neural Policies

Lyapunov functions (LFs) are a fundamental tool for stability analysis of nonlinear systems [44]. Sum-of-squares (SOS) programming enabled the systematic search for polynomial LFs [42, 40], but scales poorly in high dimensions and require polynomial dynamics. Recently, neural networks have been used to approximate LFs and stabilizing policies [5, 7, 14], significantly expanding the function space beyond SOS polynomials. However, most existing work assumes deterministic models, leaving a gap in analyzing Lyapunov stability and synthesizing stable controllers under model uncertainty.

My research addresses this gap by integrating Lyapunov-based stability principles with DRO, enabling the synthesis of neural controllers and certificates that remain valid under model uncertainty. In [30], I introduce the concept of a **distributionally robust Lyapunov function** (DR-LF) for closed-loop systems with parametric uncertainty. The DR-LF search is formulated as both an SOS programming and a neural network-based optimization problem, offering a scalable and efficient framework for certifying stability in probability for uncertain systems. Building on this, I extend the approach in [32] to jointly learn neural stabilizing controllers and Lyapunov certificates for uncertain nonlinear systems.

Most recently, I study the problem of certifying the stability of closed-loop systems under policies derived from optimal control and reinforcement learning (RL). We make two key observations: first, the value function of a given policy can be augmented with a **residual term** to construct a valid stability certificate; second, to certify stability, the classical step-wise Lyapunov decrease condition can be relaxed to a multi-step, weighted criterion. Building on these insights, my work [34] leverages the notion of a **generalized Lyapunov function** from [13] to certify the stability of neural policies. In particular, the value function of a given policy is augmented with a residual neural network, and the generalized Lyapunov decrease condition is enforced by jointly optimizing state-dependent weights. This relaxation enables certification for modern RL policies [16, 43, 18]. The same formulation also supports joint training of neural controllers and certificates using a multi-step Lyapunov loss, resulting in significantly larger certified regions of attraction.

III. Future Work

Optimality and Stability: While optimal control and RL optimize long-term performance, they typically lack stability guarantees, even within a region of interest. Lyapunov-based methods offer formal stability certification but are often hard to construct and scale. My work [34] bridges this gap by introducing a generalized Lyapunov framework that augments an RL policy's value function with a residual neural network and verifies stability via a relaxed, multi-step decrease condition. Building on this, my future work will use certificate structures to guide policy refinement and reward design, aiming to incorporate stability constraints into learning so that resulting policies are both high-performing and certifiably stable.

Safety in Contact-Rich Tasks: Conventional safety filters focus on collision avoidance but often overlook broader constraints such as force, pressure, and compliance requirements [38]. I have recently begun applying our uncertainty-aware CBF formulations to soft robots operating in anatomical environments, ensuring safe and reliable contact with vessels. Moving forward, I will extend these formulations to incorporate richer constraints, such as contact force limits, and apply them to contact-rich tasks and dexterous manipulation where physical interaction is critical.

Whole-Body Motion Planning: Traditional sampling-based methods [25, 22] often struggle in high-dimensional spaces due to the exponential growth of possible configurations. My recent work introduces configuration-space bubbles [35], where each graph node represents a set rather than a single configuration, enabling more efficient planning for manipulators in cluttered environments while relying on raw sensor data. Building on this, I aim to develop efficient and sensor-driven convex decomposition of the safe configuration space, enabling the integration of recent advances in optimization-based motion planning [37, 8]. This direction aims to create scalable, sensor-driven motion planning for high-dimensional robots in dynamic and cluttered environments.

References

[1] Hossein Abdi, Golnaz Raja, and Reza Ghabcheloo. Safe control using vision-based control barrier function (V-CBF). In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 782–788, 2023.

[2] Aaron Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *European Control Conference (ECC)*, pages 3420–3431, 2019.

[3] Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8):3861–3876, 2017.

[4] Joseph Breeden and Dimitra Panagou. Robust control barrier functions under high relative degree and input constraints for satellite trajectories. *Automatica*, 155: 111109, 2023.

[5] Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural Lyapunov control. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[6] Andrew Clark. Control barrier functions for complete and incomplete information stochastic systems. In *2019 American Control Conference (ACC)*, pages 2928–2935, 2019. doi: 10.23919/ACC.2019.8814901.

[7] Hongkai Dai, Benoit Landry, Lujie Yang, Marco Pavone, and Russ Tedrake. Lyapunov-stable neural-network control. In *Proceedings of Robotics: Science and Systems*, Virtual, July 2021.

[8] Hongkai Dai, Alexandre Amice, Peter Werner, Annan Zhang, and Russ Tedrake. Certified polyhedral decompositions of collision-free configuration space. *The International Journal of Robotics Research*, 43(9):1322–1341, 2024.

[9] Ersin Das and Joel W Burdick. Robust control barrier functions using uncertainty estimation with application to mobile robots. *arXiv preprint arXiv:2401.01881*, 2024.

[10] Vikas Dhiman, Mohammad Javad Khojasteh, Massimo Franceschetti, and Nikolay Atanasov. Control barriers in bayesian learning of system dynamics. *IEEE Transactions on Automatic Control*, 68(1):214–229, 2023.

[11] Petar M Djuric, Jayesh H Kotecha, Jianqui Zhang, Yufei Huang, Tadesse Ghirmai, Mónica F Bugallo, and Joaquin Miguez. Particle filtering. *IEEE signal processing magazine*, 20(5):19–38, 2003.

[12] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*, 171:115–166, 2018.

[13] Annika Fürnsinn, Christian Ebenbauer, and Bahman Gharesifard. Relaxed feasibility and stability criteria for flexible-step MPC. *IEEE Control Systems Letters*, 7: 2851–2856, 2023. doi: 10.1109/LCSYS.2023.3289150.

[14] Nathan Gaby, Fumin Zhang, and Xiaojing Ye. Lyapunov-net: A deep neural network architecture for Lyapunov function approximation. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 2091–2096, 2022. doi: 10.1109/CDC51059.2022.9993006.

[15] Ruben Grandia, Andrew J. Taylor, Aaron D. Ames, and Marco Hutter. Multi-layered safety for legged robots via control barrier functions and model predictive control. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 8352–8358, 2021.

[16] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.

[17] Vahid Hamdipoor, Nader Meskin, and Christos G Cassandras. Safe control synthesis using environmentally robust control barrier functions. *European Journal of Control*, 74:100840, 2023. 2023 European Control Conference Special Issue.

[18] Nicklas Hansen, Xiaolong Wang, and Hao Su. Temporal difference learning for model predictive control. In *Proceedings of the 39th International Conference on Machine Learning*, pages 8387–8406. PMLR, 2022.

[19] Ashish Ranjan Hota, Ashish Kumar Cherukuri, and John Lygeros. Data-driven chance constrained optimization under Wasserstein ambiguity sets. In *American Control Conference (ACC)*, pages 1501–1506, 2019.

[20] Ruiwei Jiang and Yongpei Guan. Data-driven chance constrained stochastic program. *Mathematical Programming*, 158:291–327, 2016.

[21] R.E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82 (1):35–45, 1960.

[22] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation (T-RO)*, 12(4):566–580, 1996.

[23] Shaghayegh Keyumarsi, Made Widhi Surya Atman, and Azwirman Gusrialdi. Lidar-based online control barrier function synthesis for safe navigation in unknown environments. *IEEE Robotics and Automation Letters*, 9(2): 1043–1050, 2024.

[24] Mikhail Koptev, Nadia Figueroa, and Aude Billard. Neural joint space implicit signed distance functions for reactive robot manipulator control. *IEEE Robotics and Automation Letters*, 8(2):480–487, 2023. doi: 10.1109/LRA.2022.3227860.

[25] J.J. Kuffner and S.M. LaValle. Rrt-connect: An efficient approach to single-query path planning. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, volume 2, pages 995–1001 vol.2, 2000.

[26] Yiming Li, Xuemin Chi, Amirreza Razmjoo, and Sylvain Calinon. Configuration space distance fields for manipu-

lation planning. In *Robotics: Science and Systems (RSS)*, 2024.

[27] Yiming Li, Yan Zhang, Amirreza Razmjoo, and Sylvain Calinon. Representing robot geometry as distance fields: Applications to whole-body manipulation. In *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, pages 15351–15357, 2024.

[28] Kehan Long, Cheng Qian, Jorge Cortés, and Nikolay Atanasov. Learning barrier functions with memory for robust safe navigation. *IEEE Robotics and Automation Letters*, 6(3):4931–4938, 2021.

[29] Kehan Long, Vikas Dhiman, Melvin Leok, Jorge Cortés, and Nikolay Atanasov. Safe control synthesis with uncertain dynamics and constraints. *IEEE Robotics and Automation Letters*, 7(3):7295–7302, 2022.

[30] Kehan Long, Yinzhuang Yi, Jorge Cortes, and Nikolay Atanasov. Distributionally robust Lyapunov function search under uncertainty. In *Learning for Dynamics and Control Conference*, pages 864–877. PMLR, 2023.

[31] Kehan Long, Yinzhuang Yi, Jorge Cortés, and Nikolay Atanasov. Safe and stable control synthesis for uncertain system models via distributionally robust optimization. In *American Control Conference (ACC)*, pages 4651–4658, 2023.

[32] Kehan Long, Jorge Cortés, and Nikolay Atanasov. Distributionally robust policy and lyapunov-certificate learning. *IEEE Open Journal of Control Systems*, 3:375–388, 2024.

[33] Kehan Long, Hardik Parwana, Georgios Fainekos, Bardh Hoxha, Hideki Okamoto, and Nikolay Atanasov. Neural configuration distance function for continuum robot control. *arXiv preprint arXiv:2409.13865*, 2024.

[34] Kehan Long, Jorge Cortés, and Nikolay Atanasov. Certifying stability of reinforcement learning policies using generalized lyapunov functions. *arXiv preprint arXiv:2505.10947*, 2025.

[35] Kehan Long, Ki Myung Brian Lee, Nikola Raicevic, Niyas Attasseri, Melvin Leok, and Nikolay Atanasov. Neural configuration-space barriers for manipulation planning and control. *arXiv preprint arXiv:2503.04929*, 2025.

[36] Kehan Long, Yinzhuang Yi, Zhirui Dai, Sylvia Herbert, Jorge Cortés, and Nikolay Atanasov. Sensor-based distributionally robust control for safe robot navigation in dynamic environments. *The International Journal of Robotics Research*, 2025. To appear.

[37] Tobia Marcucci, Mark Petersen, David von Wrangel, and Russ Tedrake. Motion planning around obstacles with convex optimization. *Science robotics*, 8(84):eadf7843, 2023.

[38] Kensuke Nakamura, Lasse Peters, and Andrea Bajcsy. Generalizing safety beyond collision-avoidance via latent-space reachability analysis. *arXiv preprint arXiv:2502.00935*, 2025.

[39] Quan Nguyen, Ayonga Hereid, Jessy W. Grizzle, Aaron D. Ames, and Koushil Sreenath. 3d dynamic walking on stepping stones with control barrier functions. In *IEEE Conference on Decision and Control*, pages 827–834, 2016.

[40] Antonis Papachristodoulou and Stephen Prajna. On the construction of Lyapunov functions using the sum of squares decomposition. In *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, volume 3, pages 3482–3487, 2002.

[41] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019.

[42] Pablo Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. California Institute of Technology, 2000.

[43] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. URL https://arxiv.org/abs/1707.06347.

[44] Eduardo Daniel Sontag. A ‘universal’ construction of Artstein’s theorem on nonlinear stabilization. *Systems & Control Letters*, 13(2):117–123, 1989. ISSN 0167-6911.

[45] Li Wang, Aaron D. Ames, and Magnus Egerstedt. Safe certificate-based maneuvers for teams of quadrotors using differential flatness. *IEEE International Conference on Robotics and Automation*, pages 3293–3298, 2017.

[46] Yujie Wang and Xiangru Xu. Disturbance observer-based robust control barrier functions. In *American Control Conference (ACC)*, pages 3681–3687, 2023.

[47] Weijun Xie. On distributionally robust chance constrained programs with Wasserstein distance. *Math. Program.*, 186:115–155, 2021.

[48] Xiangru Xu, Thomas Waters, Daniel Pickem, Paul Glotfelter, Magnus Egerstedt, Paulo Tabuada, Jessy W Grizzle, and Aaron D Ames. Realizing simultaneous lane keeping and adaptive speed regulation on accessible mobile robot testbeds. In *IEEE Conference on Control Technology and Applications*, pages 1769–1775, 2017.