Continually Learning Robots that Reason about Uncertainty and Consequences

Cherie Ho, Carnegie Mellon University

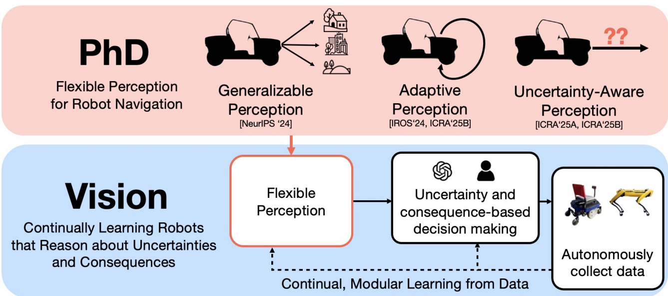

Fig. 1. I propose to develop "Continually Learning Robots" that can autonomously explore, safely collect high-quality data, and adapt to new scenarios. Building on my PhD research in generalizable, adaptive, and uncertainty-aware perception, I aim to extend these capabilities across the entire robotic stack, enabling robots that can operate at-scale in new situations.

I. Introduction

In my PhD, I have worked on three robotics systems for wide-ranging real-world applications: outdoor drone cinematography systems [3, 1, 2], high-speed off-road driving [6, 9, 5], and autonomous wheelchairs [9]. Although these systems cater to different needs and pose unique technical challenges, they have one common thread: each required years of research and continual development by humans, and required significant additional engineering even with small changes in domain and problem definition. Through my research, I aim to develop robots that can *continuously learn* in new scenarios (Fig. 1) to break the need for immense human engineering when encountering new situations and environments. This will pave the way for robots capable of operating effectively and at scale in diverse scenarios. However, robot deployments are not only challenging and costly but also high-stakes, where minor failures can have serious consequences. My vision is to create robots capable of autonomously curating high-quality data safely and at scale while learning efficiently from this data. To achieve this, I propose that ***robots need to reason about uncertainty and consequences within a modular architecture****, allowing them to learn in a risk-adjusted manner when facing new situations*. It will continuously learn by identifying knowledge gaps and actively gathering information.

II. Prior Work: Flexible Perception for Robot Navigation

Toward this goal, my PhD research has laid a foundation by developing perception modules designed for generalizability, adaptability, and enabling proactive exploration in areas of uncertainty. My first research focus is the development of generalizable perception modules that exhibit effective zero-shot performance across diverse settings [7]. While more generalizable perception provides a solid foundation, there is a strong need for real-time adaptability. In this direction, my second line of work explores how a perception module can learn and improve using newly collected data, enhancing its performance over time [6, 9]. My third area of research examines how a continuously learning agent can estimate uncertainty to identify where it should gather information [8, 4].

A. Data Engine for Generalizable Robot Perception

Robot perception modules require achieving scene understanding across varied scenarios, however, the datasets for training these modules remain scarce. An example data-scarce application is Bird's Eye View (BEV) map prediction, an important perception task for ground robots [21, 15]. To address the lack of data, I introduced `MIA`, a scalable data engine [7] where we challenge a conventional assumption - there's not enough data and we need to wait for autonomous driving companies to release more. Inspired by a related map task [16], we posit the data is already out there, hidden in disparate crowd-sourced mapping platforms that each contain worldwide-scale data. `MIA` automates the curation of large-scale BEV data from such crowd-sourced platforms.

Implications: More "anywhere" mapping in unseen locations. Using our data engine, we easily curated a diverse 1.2 million-pair dataset that includes varied geographies, conditions and capture setups. The resulting model achieved a significant improvement in zero-shot performance over baselines trained with conventional datasets. Additionally, the scalable data engine enabled the creation of a benchmark that is significantly more challenging than previous ones, supporting research aimed at the universal deployment of BEV mapping.

Implications: Empowering new perception tasks. The engine's scalable approach unlocks previously infeasible map-related tasks by addressing the scarcity of large, real-world datasets. I am collaborating with others to enable long-range mapping (over 200 meters) and large-scale cross-modal representation learning by extending to satellite imagery and embodied robot datasets. These works can pave the way for new opportunities in strategic decision-making.

B. Adaptive Perception that Improve by Itself

Robots frequently face unfamiliar scenarios not covered in training data. To ensure safety and performance, robots must adapt in real-time. However, standard online fine-tuning may fall short of meeting the demands of high-performance autonomy. For instance, it may not be robust to sensor noise, and errors can build up and cause the system to fail, and it may not be able to remember scenarios where it visited a long time ago, namely catastrophic forgetting [11]. My research has advanced adaptive methods that are robust in the face of sensor failures and recall previous long-term experiences. I developed `ALTER`, an online learning method for high-speed offroad perception that rapidly adapts visual perception models to new environments by itself [6]. `ALTER` uses self-supervised learning in a model selection framework to adapt long-range visual models from near-range LiDAR feedback.

Implications: Long-range, robust perception for offroad driving. This framework significantly improved long-range traversability prediction compared to LiDAR-only approaches and state-of-the-art visual semantic methods [13]. Additionally, I demonstrated the algorithm improved robustness to sensor failures by leveraging alternative sensors.

Implications: Safe high-speed autonomy with minimal human input. Recently, my collaborators and I developed `SALON`, a framework that quickly adapts to new environments and leverages anomaly detection to avoid out-of-distribution terrains [9]. Within seconds of collected experience, we demonstrated comparable navigation performance over kilometer-scale offroad courses as methods trained on 100-1000x more data.

C. Collecting Information on the Right Things: Exploring by Predicting What We Don't Know

For robots capable of learning continuously, it is crucial to prioritize gathering high-value information. However, this is challenging because predicting which area will yield the most valuable information is difficult. My key idea is to predict a state's information gain by reasoning on the potential sensor coverage and uncertainties from multiple explicit environment hypotheses. Specifically, I focus on autonomous exploration in structured indoor environments that are often predictable. I introduced a new exploration framework that uses multiple predicted maps to form a probabilistic sensor model for information gain estimation to guide exploration [8].

Implications: More informative data gathering. Traditional exploration algorithms often do not explicitly leverage the predictability of the environment [20]. While recent works leverage map prediction, they tend to be sensitive to the predicted map quality and often overlook sensor coverage [12, 17]. Our framework can better reason about the total uncertainty that can be resolved when visiting a given state. This resulted in significant improvement compared to representative map-prediction based methods.

Implications: Longer-horizon, multi-agent exploration. My work on exploration with predicted maps has extended in multiple directions. My collaborators and I incorporated `MapEx` in a human-inspired reinforcement learning (RL) framework to learn longer-horizon decisions[4] and to compute longer-horizon information gain [10].

III. Vision: Continually Learning Robots that Reason about Uncertainty and Consequences

To bridge the gap between current robot systems that work in specialized settings to ones that can operate ubiquitously, I envision building robots that actively reason about *what they don't know* (uncertainties) and consequences of *what they do know*. This reasoning capability will allow them to safely and efficiently learn to operate in new situations. I will build upon my previous research to develop such robots following three design principles for quick adaptation and safety: modular, uncertainty-aware and learns from humans.

A. Uncertainty and consequence-aware decision making

Common-sense reasoning on action consequences. *All uncertainties are not made equal.* For example, in my research, our anomaly detector flagged a water bottle and a human as having similar uncertainty levels, yet the consequences of traversing them are drastically different [9]. LLMs have emerged as a promising method for bringing common sense reasoning into agents [19]. Our preliminary results show that LLM can be prompted to generate reasonable traversability scores over various terrains.

Intentional data gathering under full-stack uncertainties. Robots must act *intentionally* in the face of uncertainty. In my previous work, we treated uncertainty as a black-box measure, for example pixels identified as snow on a wheelchair robot are marked as simply high-uncertainty, we generally choose to avoid it. But uncertainty needs *reasoning*, I plan to build algorithms that can, for example, reason why snow is uncertain (i.e., the depth of snow) and act intentionally to resolve the uncertainty, for example by slowly moving forward to examine the thickness of snow, or looking at previous weather.

B. Modular full-stack learning from humans and experience

Learning from humans in long-tail cases. While a robot can continuously operate by reasoning on uncertainties, there are many long-tail cases where it is best to ask for human help and learn from it [14]. I aim to build on my uncertainty-aware perception representation [8, 9] to determine when to defer to humans, present them with options to choose from, and query them for their reasoning. Developing such a human deferral mechanism can enable robots to be deployed at scale with minimal supervision. In addition, these robots act as data flywheels, collecting data for aligning LLM reasoning engines with human decision-making processes.

Modular learning with rich signals. My proposed robot system will be able to collect a wealth of embodied data and human preference data from multiple robot platforms. How should we best learn from this large cross-embodied dataset? I intend to develop modular, multi-task algorithms [18] that can improve individual modules quickly. I will borrow insights from my previous online adaptation work [6, 9] to research best meta-representations for fast adaptation.

Summary: Overall, I envision this work will further enable robots that operate across diverse scenarios without extensive human engineering.

References

[1] Towards a robust aerial cinematography platform: Localizing and tracking moving targets in unstructured environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

[2] Autonomous aerial cinematography among unstructured environments with online artistic decision-making. *Journal of Field Robotics*, January 2020.

[3] 3d human reconstruction in the wild with collaborative aerial cameras. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, September 2021.

[4] Mapexrl: Human-inspired indoor exploration with predicted environment context and reinforcement learning., 2024.

[5] Deep bayesian future fusion for self-supervised, high-resolution, off-road mapping., 2024.

[6] Learning-on-the-drive: Self-supervised adaptation of visual offroad traversability models. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.

[7] Map it anywhere: Empowering bev map prediction using large-scale public datasets. In *Advances in Neural Information Processing Systems*, 2024.

[8] Mapex: Indoor structure exploration with probabilistic information gain from global map predictions. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.

[9] Salon: Self-supervised adaptive learning for off-road navigation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.

[10] Seungjae Baek, Brady Moon, Seungchan Kim, Muqing Cao, Cherie Ho, Sebastian Scherer, and Jeong hwan Jeon. Pipe planner: Pathwise information gain with map predictions for indoor robot exploration, 2025. URL https://arxiv.org/abs/2503.07504.

[11] Robert M. French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135, 1999. ISSN 1364-6613.

[12] Georgios Georgakis, Bernadette Bucher, Anton Arapin, Karl Schmeckpeper, Nikolai Matni, and Kostas Daniilidis. Uncertainty-driven planner for exploration and navigation. In *IEEE Conference on Robotics and Automation (ICRA)*, 2022.

[13] Tianrui Guan, Divya Kothandaraman, Rohan Chandra, Adarsh Jagan Sathyamoorthy, Kasun Weerakoon, and Dinesh Manocha. Ga-nav: Efficient terrain segmentation for robot navigation in unstructured outdoor environments. *IEEE Robotics and Automation Letters*, 7(3):8138–8145, 2022. doi: 10.1109/LRA.2022.3187278.

[14] Allen Z. Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu, Leila Takayama, Fei Xia, Jake Varley, Zhenjia Xu, Dorsa Sadigh, Andy Zeng, and Anirudha Majumdar. Robots that ask for help: Uncertainty alignment for large language model planners. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2023.

[15] Avishkar Saha, Oscar Mendez, Chris Russell, and Richard Bowden. Translating images into maps. In *2022 International conference on robotics and automation (ICRA)*, pages 9200–9206. IEEE, 2022.

[16] Paul-Edouard Sarlin, Daniel DeTone, Tsun-Yi Yang, Armen Avetisyan, Julian Straub, Tomasz Malisiewicz, Samuel Rota Bulo, Richard Newcombe, Peter Kontschieder, and Vasileios Balntas. OrienterNet: Visual Localization in 2D Public Maps with Neural Matching. In *CVPR*, 2023.

[17] Rakesh Shrestha, Fei-Peng Tian, Wei Feng, Ping Tan, and Richard Vaughan. Learned map prediction for enhanced mobile robot exploration. In *Proc. of International Conference on Robotics and Automation (ICRA)*, pages 1197–1204, 2019.

[18] Samuel Triest, Mateo Guaman Castro, Parv Maheshwari, Matthew Sivaprakasam, Wenshan Wang, and Sebastian Scherer. Learning risk-aware costmaps via inverse reinforcement learning for off-road navigation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 924–930, 2023. doi: 10.1109/ICRA48891.2023.10161268.

[19] Sai Vemprala, Rogerio Bonatti, Arthur Bucker, and Ashish Kapoor. Chatgpt for robotics: Design principles and model abilities, 2023. URL https://arxiv.org/abs/2306.17582.

[20] Brian Yamauchi. Frontier-based exploration using multiple robots. In *Proceedings of the second international conference on Autonomous agents*, pages 47–53, 1998.

[21] Xiyue Zhu, Vlas Zyrianov, Zhijian Liu, and Shenlong Wang. Mapprior: Bird's-eye view map layout estimation with generative models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8228–8239, 2023.