Toward Structured and Adaptive Robot Physical Reasoning

Yixuan Huang

1 Introduction

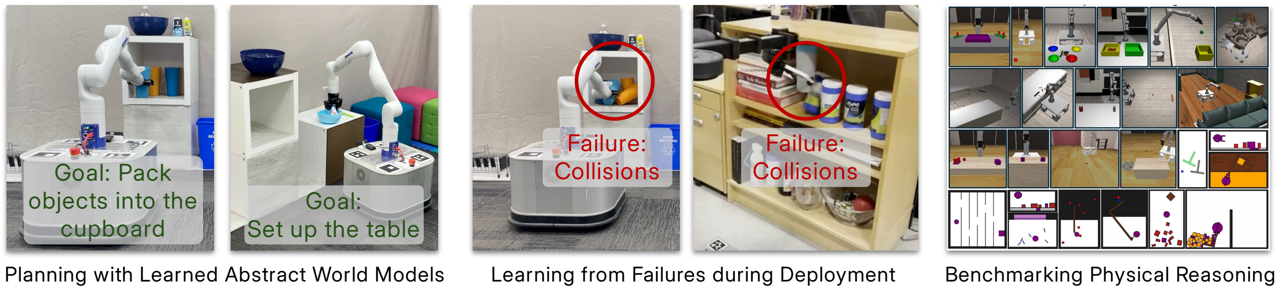

Figure 1: Toward Structured and Adaptive Robot Physical Reasoning. I propose planning with learned abstract world models, enabling structured reasoning, enabling robots to improve after deployment through adaptive learning, and developing benchmarks that reveal the remaining challenges in robot physical reasoning.

My research aims to develop general-purpose robotic systems capable of structured and adaptive physical reasoning. Robots that interact with the physical world must reason about the kinematic and dynamic constraints imposed by their embodiment, their environment, and the task at hand. These often-entangled constraints can turn semantically simple tasks into challenging puzzles. For example, to place a book on a shelf, a robot must determine whether there is a clear path to grasp the book, whether the book needs to be set down and re-grasped before placement is feasible, and whether it should gently push other books to create space using its arm. While recent advances in large-scale policy learning and foundation models have significantly expanded robotic capabilities, today's systems remain far from general-purpose robots. First, collecting large-scale interaction data in the physical world is extremely expensive and time-consuming. Second, robots must continuously adapt to out-of-distribution environments with novel physical constraints while generalizing to new task goals, states, and embodiments, capabilities that remain challenging for current end-to-end policy learning approaches. As robots transition from curated laboratory settings to unstructured everyday environments, the ability to explicitly reason about novel physical constraints and adapt to unseen environments will become central to scalable robot intelligence.

A key challenge is that robots must reason over long horizons while operating under partial observability. Long-horizon reasoning often requires hierarchical planning, where high-level decisions guide sequences of low-level robot actions. At the same time, robots perceive the world through incomplete and noisy sensory observations, such as partial-view point clouds rather than fully known object states. My research develops algorithmic frameworks that integrate perception, reasoning, learning, and planning within structured representations that address these challenges. Concretely, I developed **the first methods for learning models that enable hierarchical planning directly from partial-view point clouds** [ICRA23, T-RO24, ICRA25]. To enable continual improvement after real-world deployment, I designed systems that detect failures, recover from them, and leverage them as signals for adaptation, allowing robots to reduce future errors over time [CoRL25]. To systematically study the capabilities and limitations of existing state-of-the-art approaches, including imitation learning (IL), task and motion planning (TAMP), reinforcement learning (RL), and foundation models (FM), I developed a benchmark that systematically studies robot physical reasoning across **25 environments** [InSub26a], highlighting key challenges in reasoning about kinematic and dynamic constraints.

Looking forward, I aim to scale structured and adaptive robot physical reasoning throughout a robot's lifetime. My future research will focus on three directions: (1) enabling robots to build structured memory and reason about uncertainty (**Day 1**), (2) improving robot hierarchical models through structured failure understanding and adaptive execution (**Month 1**), and (3) scaling robot physical reasoning that generalizes across diverse robotic platforms for everyday tasks (**Year 1 and beyond**).

2 Past Work

2.1 Toward Structured Robot Physical Reasoning

A central step toward structured physical reasoning is the development of abstract, object-centric world models, referred to as *relational dynamics models* in my prior work, that capture dynamic interactions between robots and their environments. To this end, I developed **the first methods for learning such models to enable hierarchical planning from partial-view point clouds** [ICRA23, T-RO24, ICRA25].

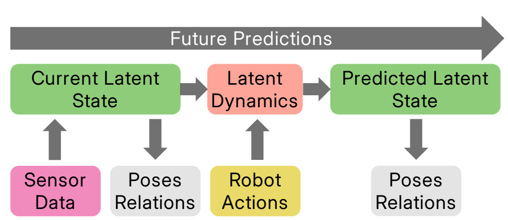

Figure 2: Structured physical reasoning with abstract, structured world models.

A key question I addressed is how to learn a representation directly from partial-view point clouds that can capture both spatial relations (e.g., the book is in the drawer) and temporal effects of robot actions (e.g., which relations would change if the robot sweeps the trash on the table). My research aims to answer this question by integrating spatial-temporal reasoning over a learned latent representation from sensory data shown in Fig. 2. This representation captures both geometric and symbolic effects of actions within a shared latent space, and the learned latent dynamics model can predict future latent states to enable planning. I deployed the learned world models on real robotic platforms, leveraging sampling-based methods to achieve goal-directed behavior specified either logically or through language. My approach achieves **a 38% improvement** in success rate over the strongest prior baselines across different numbers of objects.

2.2 Toward Adaptive Robot Physical Reasoning

An important advantage of structured world models is that they enable principled adaptation during real-world deployment. Because these models predict the effects of actions, discrepancies between predicted and observed outcomes can be detected and interpreted as signals of model mismatch. For example, a robot may predict that pushing an object will bring it into contact with a bowl, but instead observes that the object becomes blocked by another bowl. Such mismatches between predicted and observed states provide informative signals to adapt the underlying dynamics model to improve future predictions.

Building on this insight, I developed a failure-driven adaptation framework [CoRL25] that treats execution failures as informative signals for improving the learned model. The system detects inconsistencies between predicted and observed states, triggers recovery through replanning, and performs targeted real-to-sim data generation when model inaccuracies are identified. The framework then finetunes the world model using selected data to maximize information gain, reducing future prediction errors. This iterative loop of failure detection, recovery, and targeted model updates improves performance by **over 60%** in out-of-distribution environments, demonstrating the promise of refining structured robot models during deployment.

2.3 Benchmarking Robot Physical Reasoning

My work on structured world models and adaptive robot reasoning revealed a broader challenge: we currently lack systematic ways to evaluate how well robots reason about physical constraints. While my previous systems demonstrated improved planning and adaptation, existing evaluation protocols make it difficult to compare different approaches or understand where each method fails. Existing work on this problem spans multiple paradigms, including IL, TAMP, RL, and FM, each using different abstractions and evaluation protocols, making systematic comparison difficult.

To address this gap, I introduced KinDER [InSub26a], a benchmark for **Kinematic and Dynamic Embodied Reasoning**. KinDER provides 25 procedurally generated environments, a Gymnasium-compatible Python library with parameterized skills and demonstrations, and a standardized evaluation suite with eight baselines spanning TAMP, RL, IL, and FM. Empirical results show that current approaches struggle across many environments, revealing a substantial gap between existing methods and structured physical reasoning. These findings motivate my future research on enabling robots to learn during their deployment lifetime.

3 Future Directions and Long-term Vision

Looking forward, I aim to develop robots that continuously learn to reason novel constraints and adapt to out-of-distribution environments during their lifetime, shown in Fig. 3. Rather than viewing training as a one-time pre-deployment phase, I envision a deployment-centric learning paradigm in which structured reasoning provides a strong initialization, and real-world experience drives continual adaptation. Robots will begin deployment with a hierarchical model including high-level reasoning of goals, abstract world models, predictive models of physical transitions, and parameterized skills.

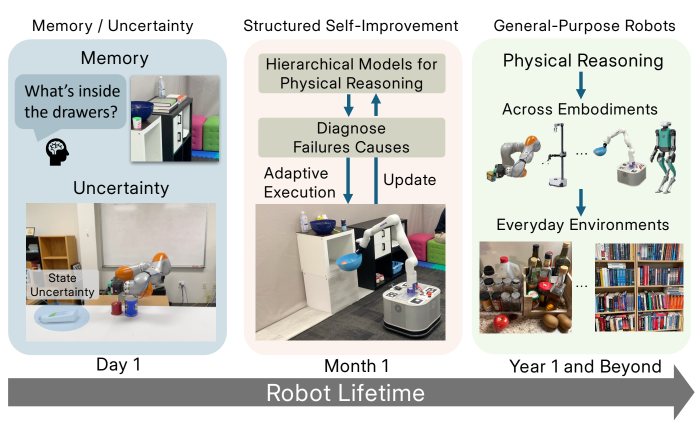

Figure 3: Looking forward, I envision robots continuously improving their models during their deployment lifetime.

3.1 Deployment as the Core Learning Phase

Day 1: Memory and Uncertainty. At the beginning of deployment, a robot must develop calibrated awareness of what it does not know. Two foundational capabilities are therefore required: structured memory and uncertainty estimation. Memory allows the robot to accumulate structured knowledge about its environment, such as previously observed object locations and relational structures, while uncertainty estimation determines the reliability of this knowledge and the robot's predictive models, guiding when the robot should explore, collect additional data, or refine its world model. Existing robotic systems typically rely on static world models learned in simulation that lack persistent structure. As a result, they struggle to distinguish between consistent structure in the environment and uncertain or poorly understood dynamics. Building on my prior work on structured robot memory [ICRA24] and uncertainty-aware world model learning [CoRL25w], I aim to develop models that jointly maintain structured memories while quantifying predictive uncertainty. This will enable the robot to reason about what it knows and what it does not: identifying stable aspects of the environment, detecting unreliable predictions, and adapting its behavior accordingly. Rather than simply executing pre-trained policies, the robot can explore when necessary, act cautiously when appropriate, and update its models in a principled manner during deployment.

Month 1: Structured Failure Understanding and Adaptive Execution. Over time, the robot will encounter failures across multiple layers, including perception errors, skill execution mismatches, world model inaccuracies, and high-level reasoning breakdowns. Leveraging its hierarchical structure, the robot will reason about the source of failure and update the appropriate module. While my prior work focused on detecting failures and recovering through replanning and model updates [CoRL25], the next step is to develop structured mechanisms for diagnosing why failures occur. Rather than treating failures as undifferentiated signals, the robot will infer whether each failure arises from perception, skill execution, world model inaccuracies, or high-level reasoning errors. This enables adaptive policy execution: reasoning more deeply when uncertainty is high, but acting efficiently when confidence is sufficient. The robot thus becomes not merely reactive, but diagnostically aware and self-improving during the deployment lifetime.

Year 1 and Beyond: Cross-Embodiment and Environment Generalization. With accumulated experience, the robot should begin to disentangle what is universal from what is embodiment- or environment-specific. High-level reasoning can transfer across robotic platforms, while skill parameterizations and failure-aware execution strategies adapt locally. This opens the door to cross-embodiment transfer and scalable adaptation across homes, warehouses, laboratories, and other real-world settings, including agriculture, earth, and ocean environments. My prior work demonstrates the feasibility of this direction. I have validated structured physical reasoning across multiple embodiments, including two 7-DOF manipulators (Kuka iiwa [ICRA23, T-RO24] and Franka [ICRA26]), three mobile manipulators (TidyBot [ICRA25], TidyBot++ [InSub26a], and Stretch RE2 [CoRL25]), and a surgical robotic platform [ISMR21]. In parallel, my work on vision-language models [InSub26b] and vision-language-action models [InSub26c] demonstrates the potential of cross-embodiment transfer. These experiences position me to systematically study how structured physical reasoning can serve as the invariant layer across embodiments, enabling robots to adapt efficiently while preserving transferable abstractions. A key principle underlying this vision is that data scaling should be structured rather than brute-force. Instead of collecting indiscriminate experience, the robot should gather data driven by uncertainty, failure signals, and information gain. Structured reasoning guides where new data is most valuable, enabling efficient continual learning while mitigating catastrophic forgetting.

3.2 Long-Term Vision

Ultimately, I envision robots that are not static models deployed into the world, but evolving physical reasoners. Structured physical reasoning, through hierarchical model learning, structured abstractions, and compositional skills, provides a strong foundation for scalable robot intelligence. Deployment then becomes the primary driver of improvement, where uncertainty estimation, structured memory, and failure-aware adaptation shape continual learning. Such robots will progressively improve over days, months, and years and become more robust, more transferable across embodiments, and more capable in diverse environments. Lifelong structured physical reasoning is therefore not merely about accumulating more data over time, but about learning in a principled and adaptive manner throughout deployment.

4 Funding and Collaboration

During my Ph.D., I had the opportunity to engage with DARPA-funded research activities and present my work directly to the DARPA TIAMAT program manager, which provided valuable exposure to federally funded robotics research. During my postdoc, I am currently involved in the preparation of an NSF proposal, gaining experience in collaborative proposal writing. As a faculty member, I plan to seek research funding from federal agencies such as NSF, DARPA, and ONR. My research on structured and adaptive robot physical reasoning aligns closely with programs such as the NSF Foundational Research in Robotics program, the DARPA TIAMAT program, and the ONR Young Investigator Program. I will also actively seek industry funding from NVIDIA, Google, Amazon Robotics, and Toyota Research Institute.

References

[ISMR21] **Yixuan Huang**, Michael Bentley, Tucker Hermans, Alan Kuntz. *Toward Learning Context-Dependent Tasks from Demonstration for Tendon-Driven Surgical Robots*. International Symposium on Medical Robotics (ISMR), 2021. **Best Paper Award–Finalist**

[ICRA23] **Yixuan Huang**, Adam Conkey, Tucker Hermans. *Planning for Multi-Object Manipulation with Graph Neural Network Relational Classifiers*. IEEE International Conference on Robotics and Automation (ICRA), 2023.

[T-RO24] **Yixuan Huang**, Nichols Crawford Taylor, Adam Conkey, Weiyu Liu, and Tucker Hermans. *Latent Space Planning for Multi-Object Manipulation with Environment-Aware Relational Classifiers*. IEEE Transactions on Robotics (T-RO), 2024.

[ICRA24] **Yixuan Huang**, Jialin Yuan, Chanho Kim, Pupul Pradhan, Bryan Chen, Li Fuxin, Tucker Hermans. *Out of Sight, Still in Mind: Reasoning and Planning about Unobserved Objects with Video Tracking Enabled Memory Models*. IEEE International Conference on Robotics and Automation (ICRA), 2024.

[ICRA25] **Yixuan Huang**, Christopher Agia, Jimmy Wu, Tucker Hermans, Jeannette Bohg. *Points2Plans: From Point Clouds to Long-Horizon Plans with Composable Relational Dynamics*. IEEE International Conference on Robotics and Automation (ICRA), 2025.

[CoRL25] **Yixuan Huang**, Novella Alvina, Mohanraj Devendran Shanthi, Tucker Hermans. *Fail2Progress: Learning from Real-World Robot Failures with Stein Variational Inference*. Conference on Robot Learning (CoRL), 2025.

[CoRL25w] Siyeon Kim, Mohanraj Devendran Shanthi, **Yixuan Huang**, Tucker Hermans. *Learning Multimodal Probabilistic Models of Manipulation Skill Effects*. Workshop on Learning to Simulate Robot Worlds, Conference on Robot Learning (CoRL), 2025.

[ICRA26] Mingen Li, Houjian Yu, **Yixuan Huang**, Youngjin Hong, Changhyun Choi. *Hierarchical DLO Routing with Reinforcement Learning and In-Context Vision-language Models*. IEEE International Conference on Robotics and Automation (ICRA), 2026.

[InSub26a] **Yixuan Huang***, Bowen Li*, Vaibhav Saxena*, Yichao Liang, Utkarsh Aashu Mishra, Liang Ji, Lihan Zha, Jimmy Wu, Nishanth Kumar, Sebastian Schere, Danfei Xu, Tom Silver. *KinDER: A Physical Reasoning Benchmark for Robot Learning and Planning*. In submission to Robotics: Science and Systems (RSS), 2026.

[InSub26b] Shuangyu Xie, Kaiyuan Chen, Ziyang Chen, Simeon Adebola, **Yixuan Huang**, Zehan Ma, Tianshuang Qiu, Wentao Yuan, Dhruv Shah, Pannag R Sanketi, Ken Goldberg. *RoboVista: Evaluating Vision Language Models for Diverse Robot Applications*. In submission to Robotics: Science and Systems (RSS), 2026.

[InSub26c] Lihan Zha, Asher James Hancock, Mingtong Zhang, Tenny Yin, **Yixuan Huang**, Dhruv Shah, Allen Z. Ren, Anirudha Majumdar. *LAP: Language-Action Pre-training Enables Zero-Shot Cross-Embodiment Transfer*. In submission to Robotics: Science and Systems (RSS), 2026.