Shared Autonomy for Unstructured Underwater Environments through Vision and Language

PIs: Matthew Walter (TTIC), Richard Camilli (WHOI), & Matthew Johnson-Roberson (UofM)
Students: Charles Schaff (TTIC) & Gideon Billings (UofM)
Research Fellow: Oscar Pizarro (ACFR/WHOI)

Motivation

Underwater science and exploration is **currently limited to non-contact surveys** with autonomous underwater vehicles **and to simple autonomous intervention in controlled laboratory settings**

Interaction with unstructured (underwater) environments **relies on humans to perform low-level teleoperation and to interpret sensor data** via manned or remotely operated vehicles

Teleoperation is a limiting factor in underwater science due to significant cognitive load placed on the operator, restriction to simple behaviors, lack of environment structure, limited field-of-view, 6-DOF motion, and communication constraints over latent, bandwidth-limited channels

Challenges of underwater environment **make marine science a surrogate for many domains** that would benefit from **shared autonomy that exploits the complementary nature of humans and robots**

Pilot performs teleoperated sampling of Kolumbo volcano with NUI HROV

Objectives

Learn representations for shared autonomy that **capture the richness of robot manipulation** in challenging unstructured underwater environments

Exploit multiple modalities including language, vision, proprioception, and manual control input **to explore the space between operator and robot**

Shared autonomy framework **learns an assistive policy that reasons over these multiview representations** to **improve the safety, efficiency, and effectiveness of human-collaborative manipulation** in complex environments

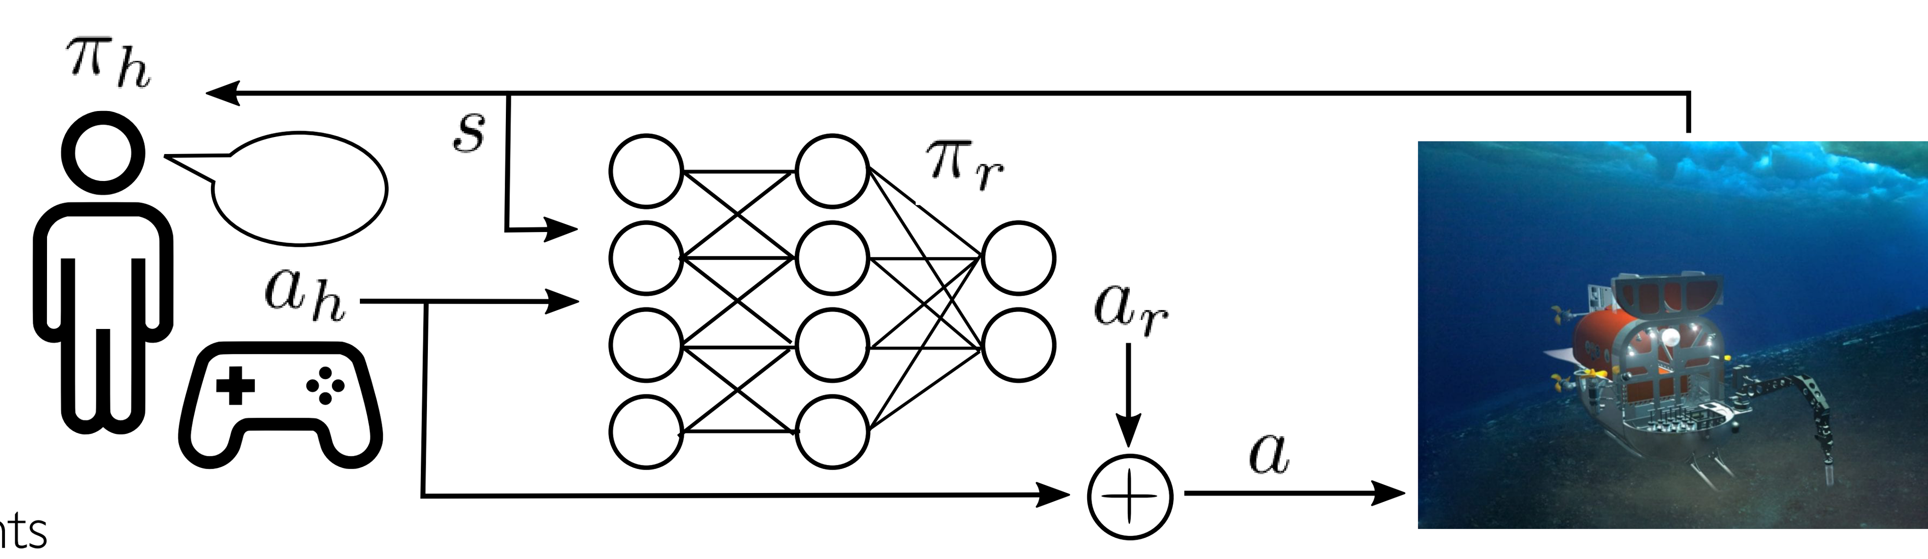

Shared autonomy via model-free multiview residual policy learning

Current Progress

Residual Policy Learning for Shared Autonomy

Developed **model-free** reinforcement learning algorithm for shared autonomy

Agent learns a residual policy that **minimally corrects human input** as necessary to ensure that a **task-agnostic reward** (safety) constraint is satisfied

Algorithm **requires no knowledge** of the environment model, the human's goal or the space of goals, or the task, and generalizes across multiple tasks

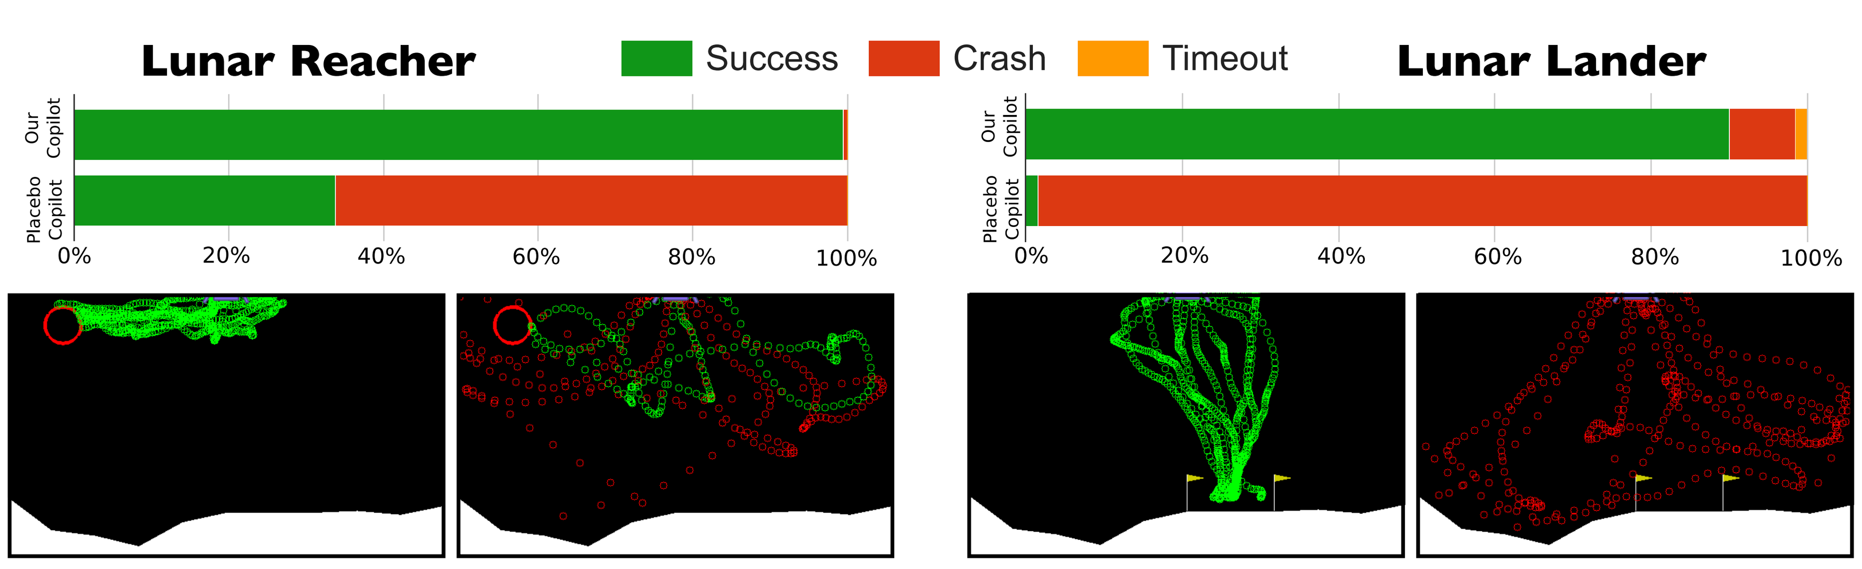

User study performance (top) & human pilot trajectories (bottom) with our assistant (left) & without (right)

We evaluate our method on three control tasks with complex dynamics

We train the assistant to satisfy a crash rate constraint, and include a residual penalty that discourages adjustments to the human-provided control actions

We train the assistant using human policies that were trained to imitate diverse demonstrations provided by multiple people via behavioral cloning

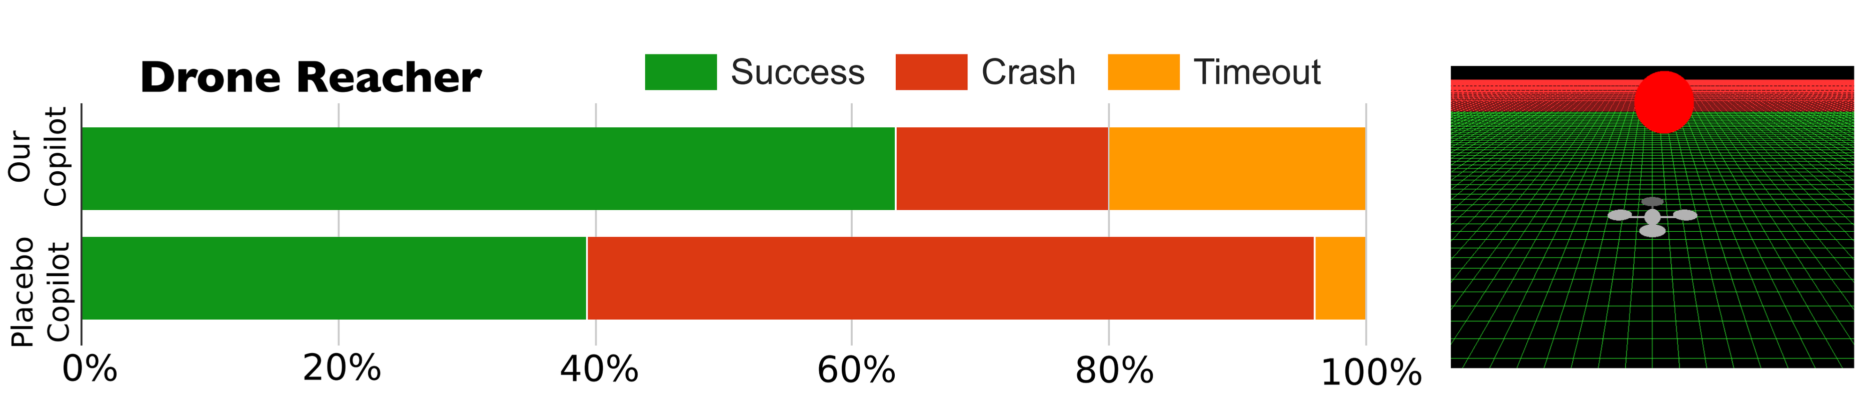

User study performance for the Drone Reacher experiments

We conducted user studies to quantitatively and qualitatively evaluate the effectiveness of our shared autonomy agent in the three control domains

Despite having no knowledge of the task or model, **our method significantly improved performance** and **received higher user ratings** over the baseline

Multigrid Neural Memory

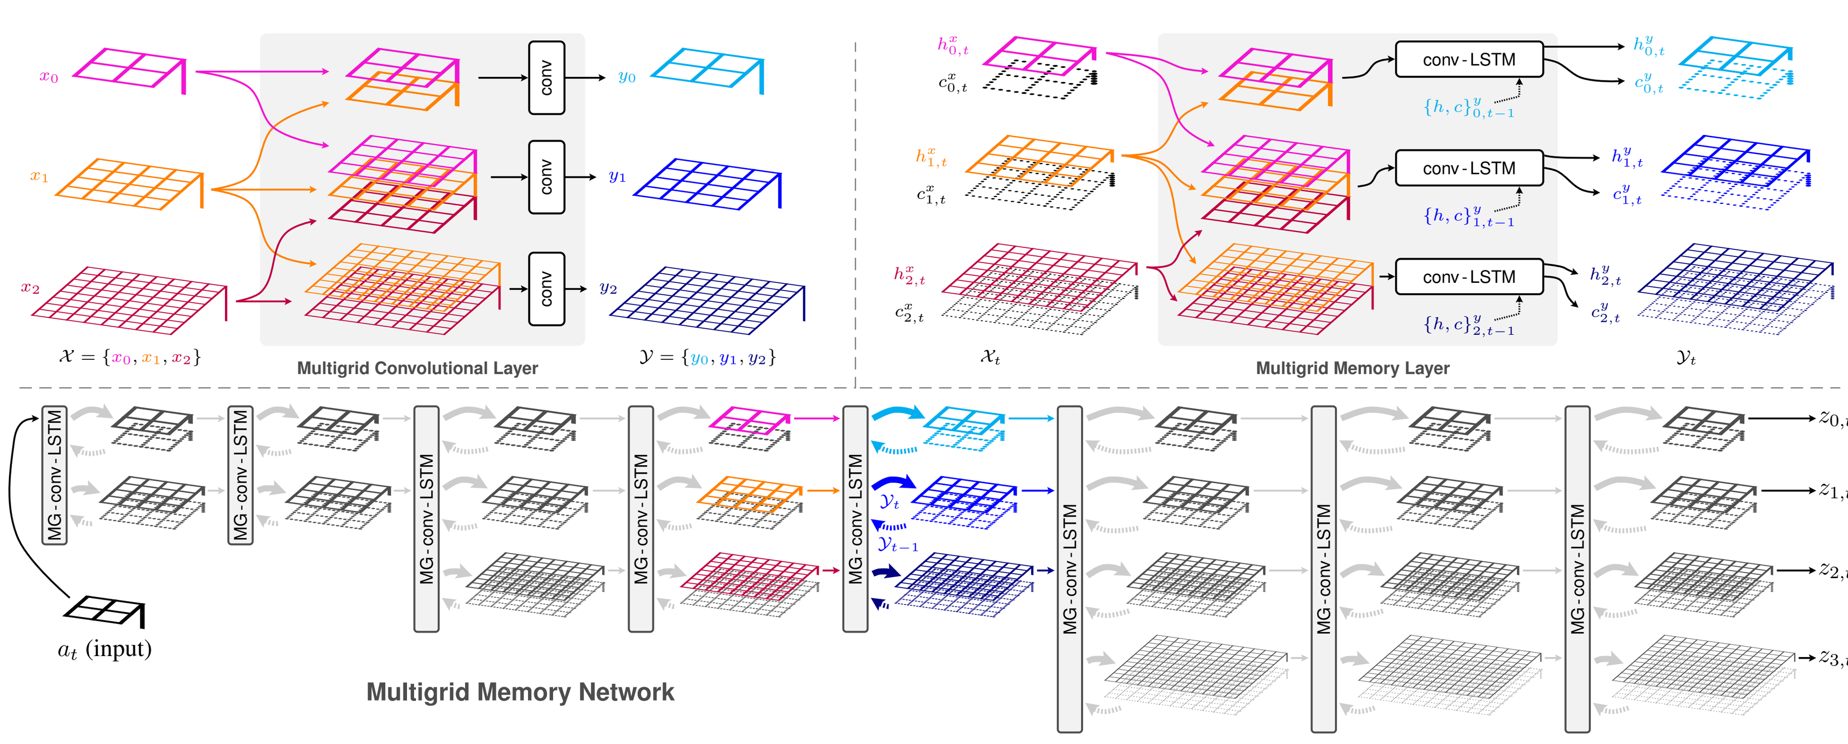

Our multigrid memory architecture incorporates multigrid convolutional layers (top-left) into a convolutional LSTM variant (top-right). Connecting such layers across time yields a radical new memory architecture (bottom)

Developed a **radical new approach** to endowing neural networks with coherent **long-term and large-scale memory** that **emerges** during training

Multigrid wiring provides **exponentially more efficient routing** that enables **implicit memory addressing** with the parameter-sharing efficiency of CNNs

Multigrid memory networks provide **distributed, co-located memory and compute**, and can be connected as freely and flexibly as CNNs and LSTMs

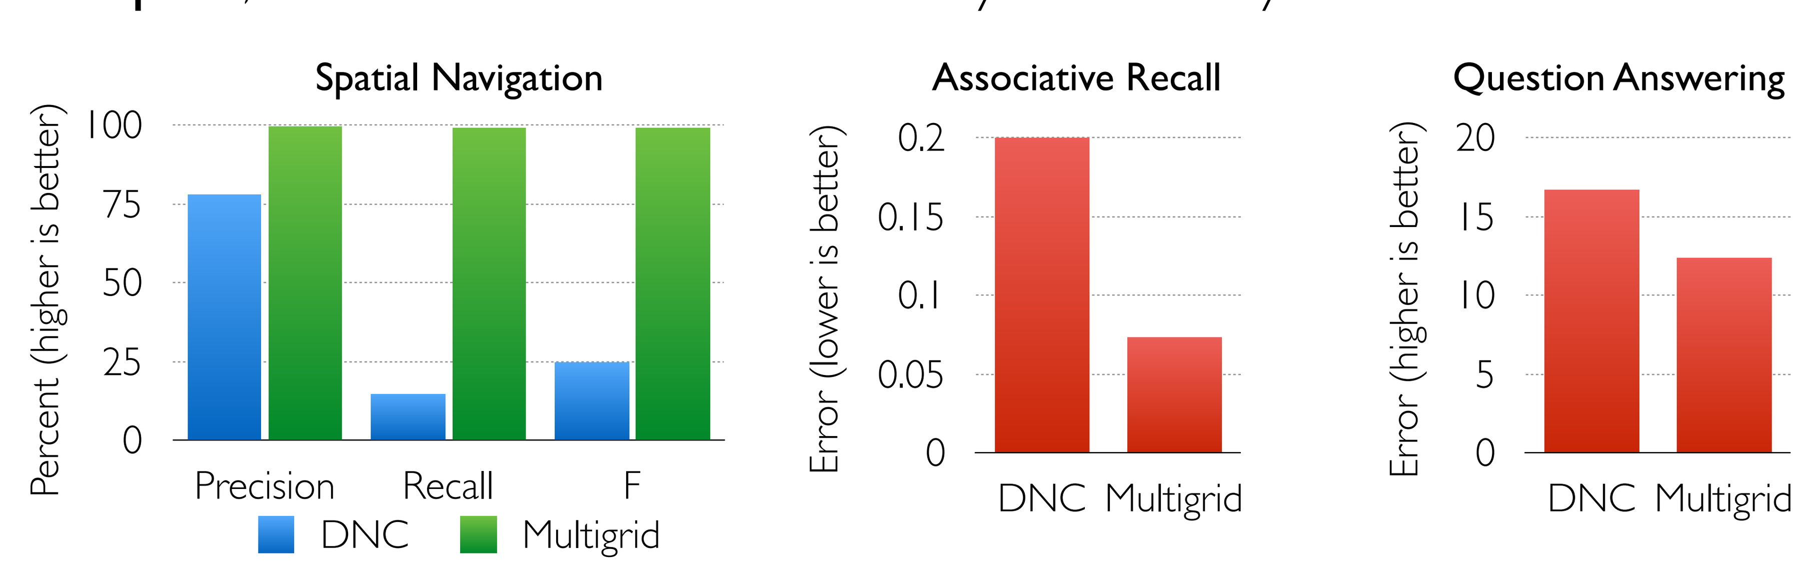

Multigrid memory networks are **general**, **outperforming state-of-the-art** DNC on navigation, algorithmic, and natural language processing tasks

Acknowledgements

This work was supported by the National Science Foundatiion under grants IIS-1830660, IIS-1830500, and IIS-1830345