Summary of Research Report
Award Number: 80NSSC20K1720
Project Name: RASPBERRY SI: Resource Adaptive Software Purpose-Built for Extraordinary Robotic Research Yields — Science Instruments (USC)
PI: Pooyan Jamshidi (USC)
Program Managers: Erica Montbach (COLDTech) and Carolyn Mercer (ARROW)

Project Website: https://nasa-raspberry-si.github.io/raspberry-si/
Team Members: Pooyan Jamshidi (USC, PI), Bradley Schmerl (CMU, Co-I), David Garlan (CMU, Co-I), Javier Camara (York, Collaborator), Matt DeMinico (Aerospace Corporation, Co-I), Jianhai Su (USC, Graduate Student), Abir Hossen (USC, Graduate Student), Ellen Czaplinski (NASA, Science Consultant), Katherine Dzurilla (UArk, Science Consultant), Sonam Kharde (USC, Postdoc).

This collaborative project resulted in several **top-tier publications** and the development of general-purpose and open-source **tools** for enabling causal analysis in space explorations and robotics. This includes root cause identification, causal debugging, causal optimization, and causal-induced verification. Below we discussed each contribution as follows: **CaRE** (root cause identification), **CuRE** (causal debugging and optimization), **Causally Induced MDPs** (for scalable runtime verification), and **Causal-MAPE-K** (our end-to-end self-adaptation framework for fault-tolerant autonomy design for the ocean world lander system).

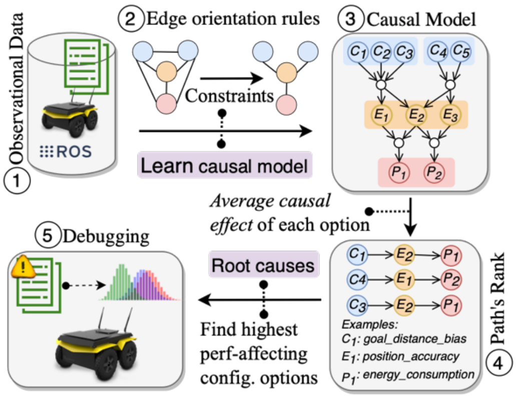

CaRE: Finding Root Causes of Configuration Issues in Highly Configurable Robots [published at IEEE/RSJ RA-L and iROS]**.** We have developed CaRE to find the Root Causes of Configuration Issues in highly configurable robots. Like other computer systems we have previously explored, robotic systems have several configurable subsystems and components, and thousands of possible software and hardware configuration options interacting non-trivially, such interactions depend on the environment the robots are operating, and such environments are typically dynamic and uncertain. The configurable parameters are set to target specific objectives, but they can cause functional faults when incorrectly configured. Finding the root cause of such faults is challenging due to the exponentially large configuration space and the dependencies between the robot's configuration settings and performance. CaRE is a method we have developed for diagnosing the root cause of functional faults through the lens of causality. Our approach abstracts the causal relationships between various configuration options and the robot's performance objectives by learning a causal structure and estimating the causal effects of options on robot performance indicators. We demonstrate CaRE's efficacy by finding the root cause of the observed functional faults and validating the diagnosed root cause by conducting experiments in physical robots (Husky and Turtlebot) and simulation~(Gazebo). Furthermore, we demonstrate that the causal models learned from robots in simulation (e.g., Husky in Gazebo) are transferable to physical robots across different platforms (e.g., Husky and Turtlebot).
Code: https://github.com/softsys4ai/care
Paper: https://ieeexplore.ieee.org/document/10137745/

CURE: Simulation-Augmented Auto-Tuning in Robotics [under review at IEEE Transactions on Robotics (T-RO)]**.** Robotic systems are typically composed of various subsystems, such as localization and navigation, each encompassing numerous configurable components (e.g., selecting different planning algorithms). Once an algorithm has been selected for a component, its associated configuration options must be set to the appropriate values. Configuration options across the system stack interact non-trivially. Finding optimal configurations for highly configurable robots to achieve desired performance poses a significant challenge due to the interactions between configuration options across software and hardware that result in an exponentially large and complex configuration space. These challenges are further compounded by the need for transferability between different environments and robotic platforms. Data efficient optimization algorithms (e.g., Bayesian optimization) have been increasingly employed to automate the tuning of configurable parameters in cyber-physical systems. However, such optimization algorithms converge at later stages, often after exhausting the allocated budget (e.g., optimization steps, allotted time) and lacking transferability. We developed CURE—a method that identifies causally relevant configuration options, enabling the optimization process to operate in a reduced search space, thereby enabling faster optimization of robot performance. CURE abstracts the causal relationships between various configuration options and robot performance objectives by learning a causal model in the source (a low-cost environment, such as the Gazebo simulator) and applying the learned knowledge to perform optimization in the target (e.g., Turtlebot 3 physical robot). We demonstrated the effectiveness and transferability of CURE by conducting experiments that involved varying degrees of deployment changes in both physical robots and simulation.

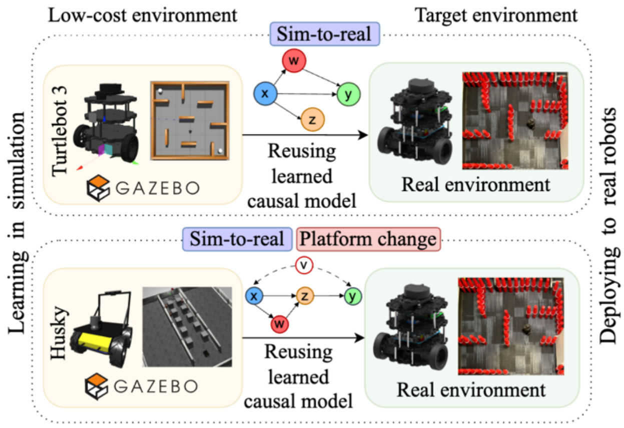

Code: https://github.com/softsys4ai/cure
Paper: https://arxiv.org/abs/2402.05399

From Roots to Probabilities: A Causal Exploration in the Probabilistic Model Checking. Probabilistic Model Checking (PMC) is a powerful technique for analyzing and verifying systems with probabilistic behavior. However, traditional PMC techniques with Markov Decision Processes (MDPs) models face several challenges, including state space explosion, complex reward structures, model uncertainty, computational complexity, and difficulties in synthesis. We address these challenges by developing the techniques that lie at the intersection of causality and PMC. At a high level, by incorporating causal models, we generate causally induced MDPs (CIMDPs) that capture the causal dependencies among (state, action, and reward) variables. While rewards and transition provide valuable information, understanding the underlying causal relationships in an MDP can offer deeper insights (e.g., identifying root

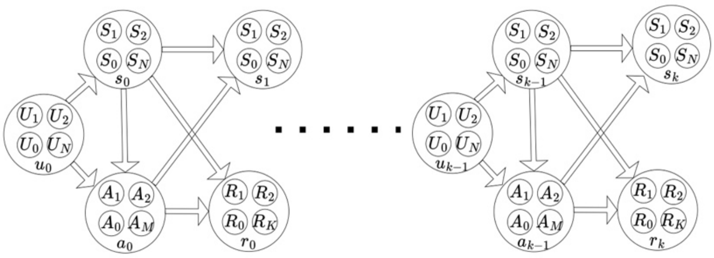

causes, finding optimal intervention variables, and predictions for unobserved behavior). We propose mathematical formulations and algorithms for incorporating causal information into the transition probabilities and rewards of the CIMDP and perform the verification over a new model. Our approach improves the verification processes and offers valuable insights such as enhanced accuracy, improved efficiency, better representation of system behavior, flexibility for model update, and improved decision-making in probabilistic systems. We demonstrated the effectiveness of our proposed approach via experimental evaluations and a case study of a self-driving car. The results highlight the benefits of incorporating causality in PMC with improved accuracy and scalability of runtime verification tasks, making it suitable for runtime planning in robotics.

Fault Tolerant Autonomy Design for Ocean World Lander System
We investigated the challenges of planetary exploration in space, specifically focusing on ocean worlds such as Europa and Enceladus, and highlighted the importance of autonomous robotics technology to overcome these challenges. Considering the limitations posed by unfamiliar environments, energy constraints, and communication delays, we developed a fault-tolerant autonomy framework to enable the applications of causal modeling and causal effect estimation within autonomous systems. Our autonomy framework, called Causal-MAPE-K, implements a cohesive and orchestrated set of components for enabling **Monitoring** (of the system states), **Analysis** (of structural and behavioral constraints to trigger re-planning), **Planning** (synthesizing a policy for generating actions), and **Execution** (that enacts the synthesized plan to the system under test). The proposed design utilizes the causally induced Markov Decision Process (MDP) to enable runtime analysis and decision-making. Causality plays a crucial role in identifying the root causes of faults and minimizing the search space in planner algorithms. Additionally, the efficacy of the proposed autonomy design is demonstrated on physical and virtual NASA testbeds.

Code: https://nasa-raspberry-si.github.io/raspberry-si

We presented the results at top-tier robotics conferences, like iROS, and published our research in top-tier robotics journals, such as IEEE Transactions on Robotics (T-RO) and IEEE Robotics and Automation Letters (RA-L). We have also presented our results in several NASA-related conferences, including the **Planetary Science Informatics and Data Analytics Conference (PSIDA 2022)** and the **53rd Lunar and Planetary Science Conference (LPSC 2022)**.

We have done several **media outreach**:
- **Breakthrough Star Award interview:**
https://sc.edu/uofsc/posts/2022/06/breakthrough_star_pooyan_jamshidi.php
- **University of South Carolina:**
https://www.sc.edu/study/colleges_schools/engineering_and_computing/news_events/news/2020/jamshidi_ai_space.php
- **The Post and Courier newsletter:**
https://www.postandcourier.com/columbia/news/usc-researcher-wants-to-train-robots-for-nasa-deep-space-missions/article_93d9bb3c-3afa-11eb-bd4c-7700ac496485.html

Finally, we have updated the project **website** with all presentations and research outcomes. The website serves as a medium for **disseminating** the results of our project: https://nasa-raspberry-si.github.io/raspberry-si/