High-Assurance Design of Learning-Enabled Cyber-Physical Systems with Deep Contracts

Pierluigi Nuzzo, Department of Electrical and Computer Engineering, University of Southern California

nuzzo@usc.edu

https://descyphy.usc.edu/research/cyber-physical-system-design/

Award ID#: 1846524

Learning-Enabled Cyber-Physical Systems

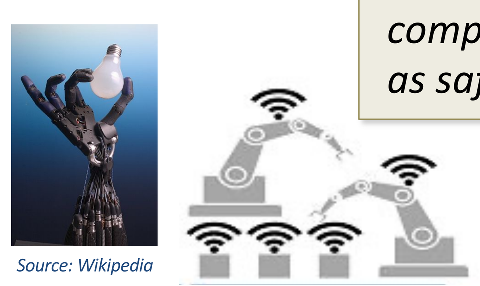

Source: Wikipedia

*- Modern AI techniques enable **adaptiveness** and **resilience** of cyber-physical systems, but also bring more **complexity**, **heterogeneity**, **approximations** and **uncertainty** in the design*
*- **Requirements** are **not rigidly defined:** How to relate component-level robustness to system-level objectives, such as safety, reliability, performance, cost?*

Goal: *A **holistic framework** including **modeling techniques, specification formalisms, and scalable algorithms** for the design and analysis of **intelligent, autonomous, cyber-physical systems including AI-enabled components** with **high guarantees of correctness** in a **modular** way*

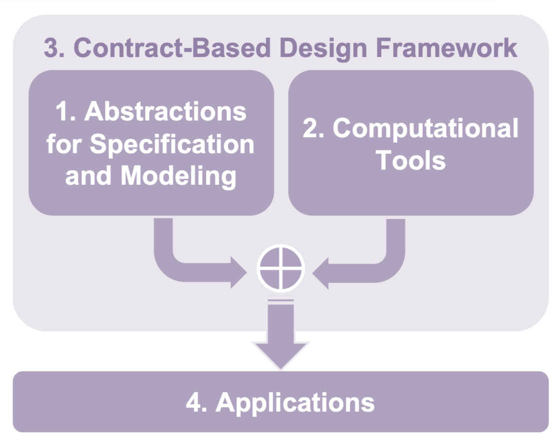

Research Organization

Reasoning with Deep Contracts

Contract $C=(V,A,G)$:
Set $V = I \cup O$ of **variables**
Set A of **assumptions**
Set G of **guarantees**

A, G: **behaviors** over V

An **implementation** M satisfies a contract if $M \cap A \subseteq G$
An **environment** E satisfies a contract if $E \subseteq A$

(A, G) is **compatible** iff $A \neq \emptyset$

(A, G) is **consistent** iff $G \neq \emptyset$

Composition $C_1 \otimes C_2$
$A = (A_1 \cap A_2) \cup \neg G_1 \cup \neg G_2$
$G = G_1 \cap G_2$

Conjunction $C_1 \wedge C_2$
$A = A_1 \cup A_2$
$G = G_1 \cap G_2$

Refinement $C_1 \preccurlyeq C_2$
$A_1 \supseteq A_2$
$G_1 \subseteq G_2$

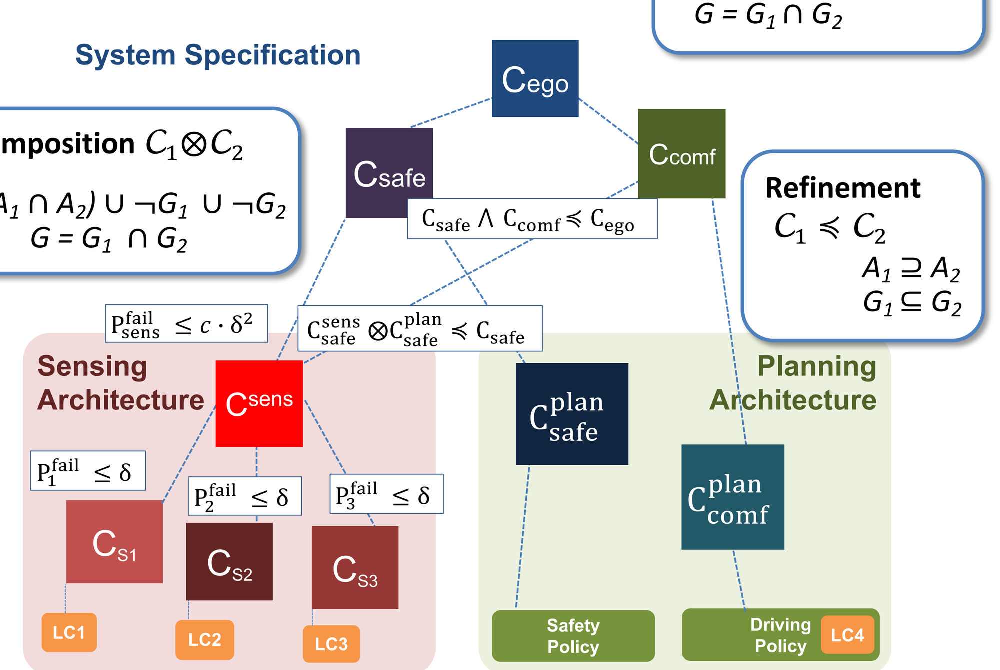

Existing contract frameworks (e.g., [Benveniste et al. ‘12, Nuzzo et al. ‘15, ’18, ‘19]) enable **modular verification**, **hierarchical refinement**, and **design reuse** based on a rigorous calculus, but fall short of *effectively capturing uncertainty*, often leading to *pessimistic solutions (over-design)* or *intractable representations*

Deep Contracts *for **compositional reasoning** about **probabilistic system behaviors**:*
- ***Context-aware:*** *describe components conditioned to their environment and overall system goals*
- ***Stochastic:*** *express and propagate uncertainty at different abstraction layers*
- ***Vertically-integrated:*** *bridge heterogeneous models and architectures across the design hierarchy*
- ***Pervasive:*** *offers mechanisms to monitor requirements for continual assurance*

Stochastic Contracts for Requirement Analysis and Computer-Aided Construction of Assurance Cases

Leverage *Stochastic Signal Temporal Logic (StSTL)* to express assumptions and guarantees on real-time, real-valued, stochastic signals and formulate *quantitative* verification and synthesis problems as *robust* StSTL satisfiability problems

P. Nuzzo et al., “Stochastic Assume-Guarantee Contracts for Cyber-Physical System Design,” *TECS’19*
C. Oh et al., “Optimizing Assume-Guarantee Contracts for Cyber-Physical System Design,” *DATE’19*
C. Oh et al., “Quantitative Verification and Design Space Exploration Under Uncertainty with Parametric Stochastic Contracts,” *ICCAD’22*

Synthesis and **validation** of **assurance cases** as networks of stochastic contracts: Contracts offer the **semantic foundation** to capture claims, premises, and confidence

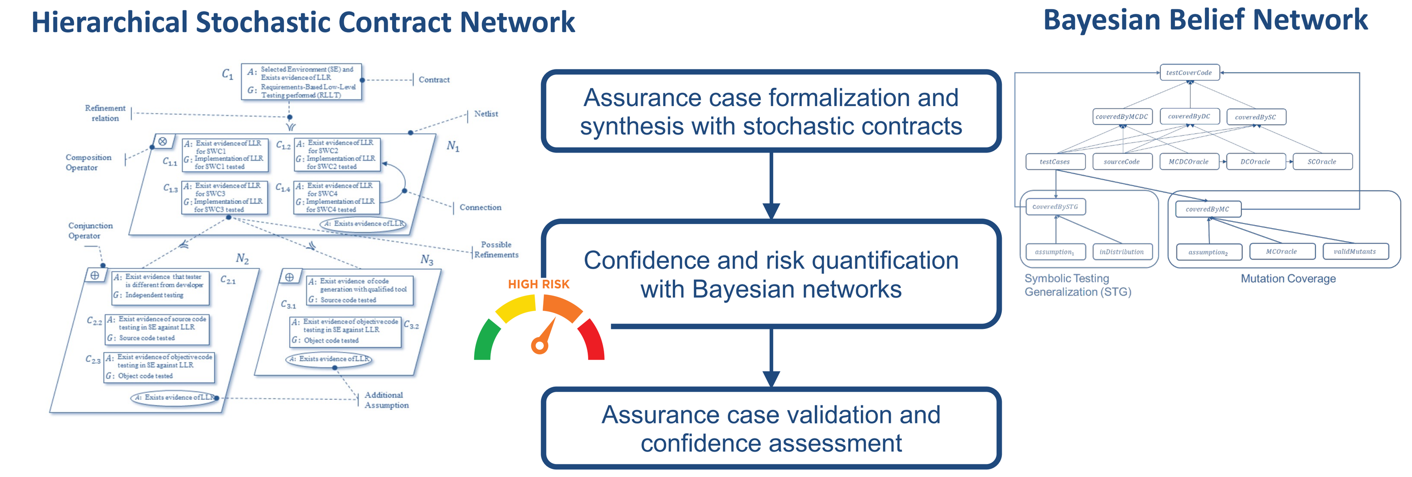

Z. Daw et al., “Computer-Aided Evaluation for Argument-Based Certification,” *DASC’23*, Best-in-Session Paper Award
Z. Daw et al., “AACE: Automated Assurance Case Environment for Aerospace Certification,” *DASC’23*, Best-in-Session Paper Award
C. Oh et al. “ARACHNE: Automated Validation of Assurance Cases with Stochastic Contract Networks,” *SAFECOMP’22*
T. Wang et al., “Hierarchical Contract-Based Synthesis for Assurance Cases,” *NFM’22*
T. Wang et al., “Computer-Aided Generation of Assurance Cases,” *SAFECOMP’23*

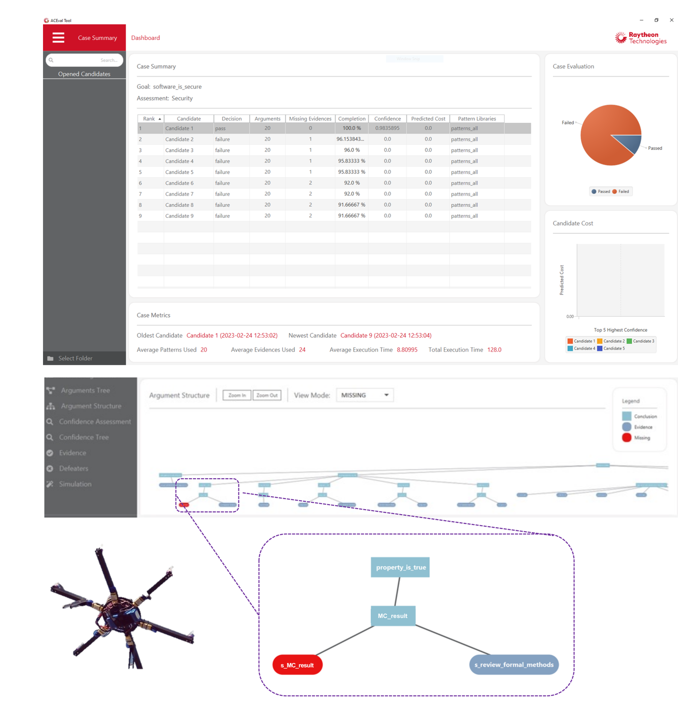

Execution of AACE (Automated Assurance Case Environment) on ArduCopter Rotorcraft

Generated over 10^5 arguments in <100 min for an industrial case study

Computer-aided, compositional *construction of* ***assurance cases*** *helps transition from process-based to* ***property-based*** *and* ***continuous certification***

Verification of AI-Enabled Systems

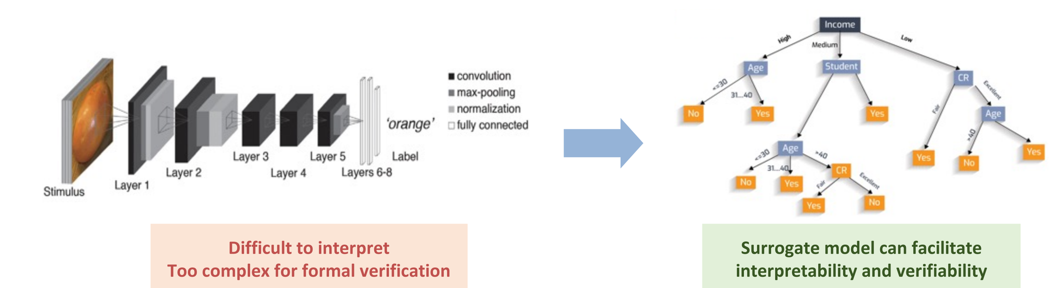

Exact and cost-effective transformation to SDT:
Up to 20x improvement in verification time (MountainCar)

Theorem 1 (Equivalence): Any discrete-output neural network (NN) has an equivalent soft decision tree (SDT) in terms of input-output behavior

Theorem 2 (Complexity): SDT size scales exponentially with the number of inputs but polynomially with the width of the maximum hidden layer

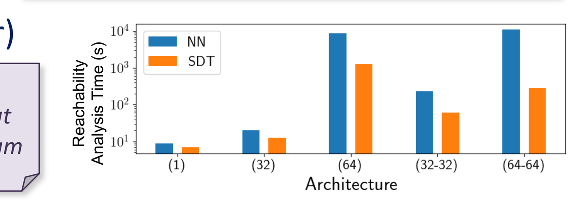

K. Chang et al., “Exact and Cost-Effective Automated Transformation of Neural Network Controllers to Decision Tree Controllers,” *CDC’23*

Robustness Contracts: Compositional verification of closed-loop systems with deep reinforcement learning controllers against perception errors

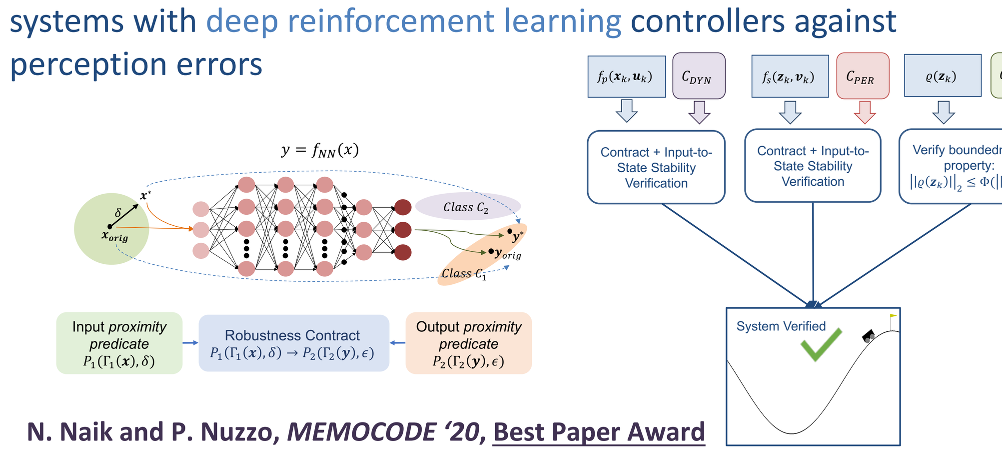

N. Naik and P. Nuzzo, *MEMOCODE ‘20*, Best Paper Award

Co-Design of Topology, Scheduling, and Planning in Automated Warehouses

- Given a warehouse layout, a list of products and a time limit, find a motion plan for a team of robots which brings every product to a packaging station within a given timeframe
- Contract-based approach outperforms alternative methods and scales to real-world problems involving teams of hundreds of robots transporting a million products.

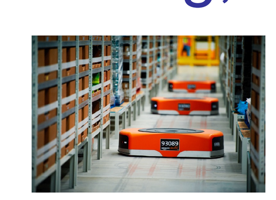

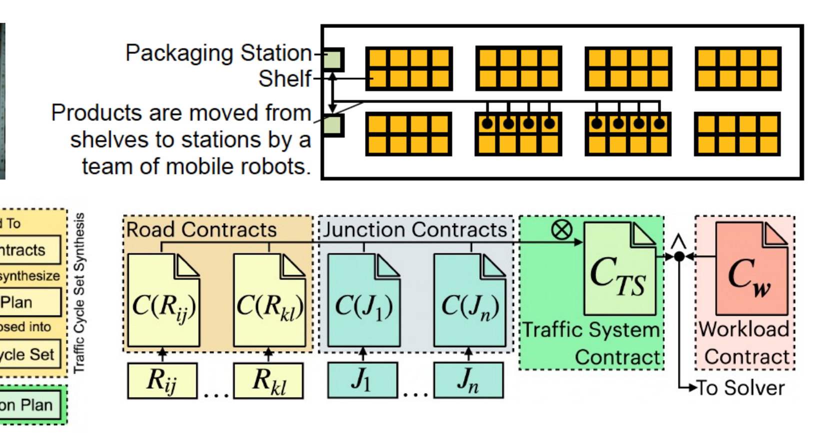

C. Leet et al., “Task Assignment, Scheduling and Motion Planning for Automated Warehouses for Million Product Workloads,” IROS 2023
C. Leet et al., “Co-Design of Topology, Scheduling, and Path Planning in Automated Warehouses,” DATE 2023

Logically Constrained Decision Making Under Uncertainty

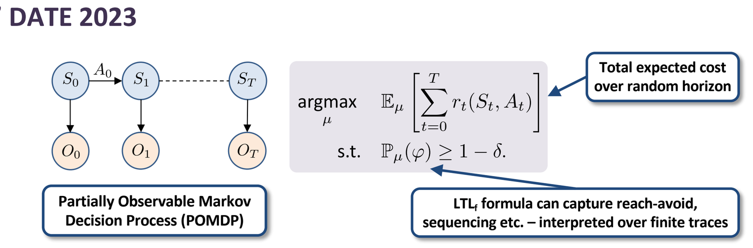

- Optimal planning via reinforcement learning for Markov decision processes (MDPs) under signal temporal logic (STL) specifications
- Optimal planning for partially observable MDPs (POMDPs) under finite linear temporal logic (LTL_f) specifications

K. Kalagarla et al., “Optimal Control of Partially Observable Markov Decision Processes with Finite Linear Temporal Logic Constraints, ” UAI 2022
K. Kalagarla et al., “Model-Free Reinforcement Learning for Optimal Control of Markov Decision Processes Under Signal Temporal Logic Specifications,” *CDC’21*
K. Kalagarla et al., “A Sample-Efficient Algorithm for Episodic Finite-Horizon MDP with Constraints”, *AAAI’21*
K. Kalagarla et al., “Optimal Control of Discounted-Reward Markov Decision Processes Under Linear Temporal Logic Specifications, ” *ACC’21*

Impact on Society and Education

- Provide the foundations for **rapid, compositional, certified design and operation of adaptive and resilient learning-enabled cyber-physical systems** for a broad range of applications: autonomous vehicles, robotics, industrial automation, medical devices, …
- Research outcomes are part of an **educational program** focusing on systems engineering concepts and multidisciplinary methods to realize safe and cost-effective intelligent systems interacting with people
  - **Pre-college:** via the USC Viterbi SHINE Program
  - **Undergraduate and graduate:** via new labs and collateral initiatives such as the **USC AutoDRIVE Lab**, the **USC Autonomous Vehicles Club**, and the USC autonomous driving **RaceOn!** competition

SHINE Summer High School Intensive in Next-Generation Engineering

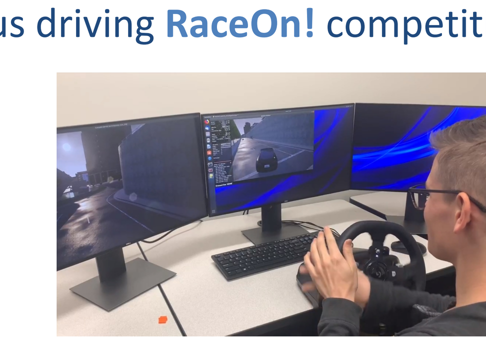

AutoDRIVE LAB

USC Viterbi School of Engineering | RTRC RTX Technology Research Center | Collins Aerospace