Head: Asst. Prof. Christos Mavridis
Cyber-Physical Systems Research Unit
Institute of Computer Engineering
Faculty of Informatics · TU Wien

Learning for Cyber-Physical Systems

Control · Optimization · Learning

LearnCPS develops theory and algorithms that enable **cyber-physical systems**, from robots and autonomous vehicles to infrastructure networks and medical devices, to **learn from data and interaction** while remaining **safe, robust, interpretable and efficient**.

Our research lies at the intersection of **control**, **optimization** and **machine learning**, unified by one premise: *recursive algorithms are dynamical systems*. Learning and optimization algorithms evolve in time and interact with the systems they act on; they can therefore be designed, analyzed and certified with the tools of systems and control theory. Conversely, learning gives controllers the ability to adapt, predict and generalize. We study this interplay in **cyber-physical-human systems**, where computation, physical dynamics and people form a single closed loop.

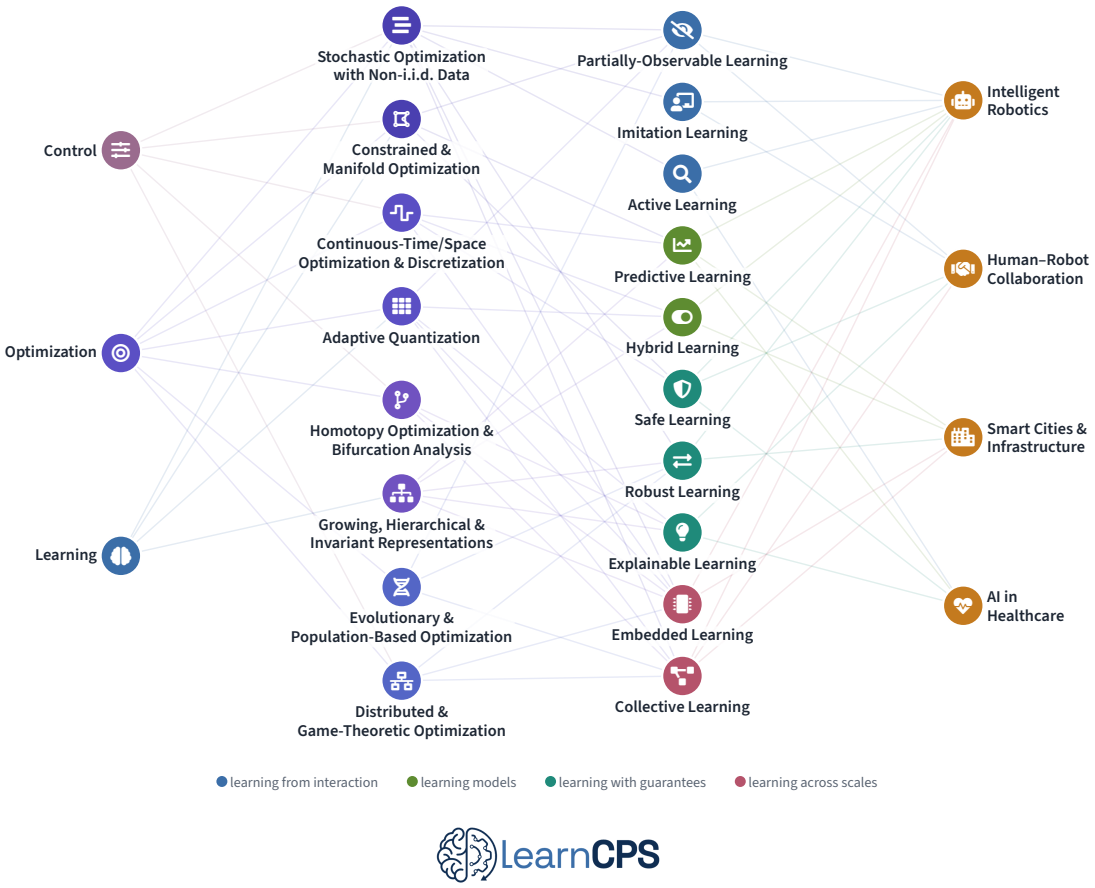


Figure 1. LearnCPS as a layered network. From left to right: disciplines, foundations, research directions and applications; each edge indicates that a node informs a node in the next layer.

Foundations

Our research directions rest on eight foundational problems in optimization and learning theory, grouped by the class of systems they address and each approached from the dynamical-systems viewpoint.

Constrained systems

Stochastic Optimization with Non-i.i.d. Data

Stochastic recursions driven by noisy, Markovian and non-stationary data, the regime of any learner in closed loop, analyzed through their limiting ODEs.

Constrained & Manifold Optimization

Optimization under constraints and on manifolds; learning the low-dimensional geometry of data and dynamics.

Hybrid systems

Continuous-Time/Space Optimization & Discretization

Algorithms derived as discretizations of continuous-time flows, with guarantees preserved under discretization.

Adaptive Quantization

Intelligent, task-aware quantization of data, states, actions, time, communication and models: adaptive codebooks that allocate resolution where it matters for the downstream decision, with rate–distortion guarantees.

Progressive systems

Homotopy Optimization & Bifurcation Analysis

Tracking solution paths as a homotopy (annealing) parameter varies; bifurcations indicate when and how model structure must change.

Growing, Hierarchical & Invariant Representations

Representations that grow in complexity only where the data demand it, organize into multi-resolution hierarchies, and respect system invariances.

Multi-agent systems

Evolutionary & Population-Based Optimization

Gradient-free, population-based global search for policy and controller design, combined with gradient-based refinement.

Distributed & Game-Theoretic Optimization

Optimization and learning over networks of agents with local information: consensus, distributed algorithms, and games among interacting decision makers.

Research Directions

We organize our research directions by the question they address: how systems *learn from interaction*, which *models* they learn, which *guarantees* learning must provide, and how learning operates *across scales*, from embedded devices to networks and swarms.

Learning from interaction

Partially-Observable Learning

Reinforcement learning under partial observability: belief representations and memory from observation streams, and **vision–language–action (VLA) models**.

POMDPs · RL · VLA

Imitation Learning

Learning policies from human or algorithmic **demonstrations**, capturing skills that are hard to encode as rewards.

Demonstrations · Human experts

Active Learning

Selecting **informative data**: information-seeking exploration and experiment design when interaction is costly.

Exploration · Experiment design

Learning models

Predictive Learning

Uncertainty-aware **world models** embedded in **learning-based model predictive control**, enabling planning over predicted futures and online adaptation.

World models · Learning-based MPC

Hybrid Learning

Learning **hybrid systems** that couple discrete decisions and switching dynamics with continuous control.

Hybrid systems · Switching dynamics

Learning with guarantees

Safe Learning

Learning-based control with **guarantees** of constraint satisfaction, invariance and stability, during learning and after deployment.

Constraints · Stability

Robust Learning

Performance under change: **transfer learning**, distribution shift, the **sim2real gap**, and **risk-sensitive** criteria.

Transfer · Sim2Real · Risk sensitivity

Explainable Learning

Models whose structure, such as prototypes and hierarchies, can be **inspected and verified** by humans.

Interpretability · Verification

Learning across scales

Embedded Learning

Lightweight, online and progressively growing models that run on **embedded hardware** under memory, energy and latency budgets.

Online learning · Real-time · Resource-aware

Collective Learning

Learning in networks and swarms: **inferring interaction laws** and the emergence of **leadership**, **trust**, **implicit communication** and **self-organization**.

Interaction laws · Leadership · Trust · Self-organization

Applications

Our methods target cyber-physical-human systems in four domains.

Intelligent Robotics

UAVs and swarms, autonomous cars, manipulators and industrial robots.

Human–Robot Collaboration

Learning from and adapting to human partners; brain–robot interfaces; safe physical interaction and shared autonomy.

Smart Cities & Infrastructure

Mobility, energy and infrastructure networks; large-scale multi-agent coordination.

AI in Healthcare

Data-efficient, interpretable decision support and assistive technology.

The Lab at a Glance

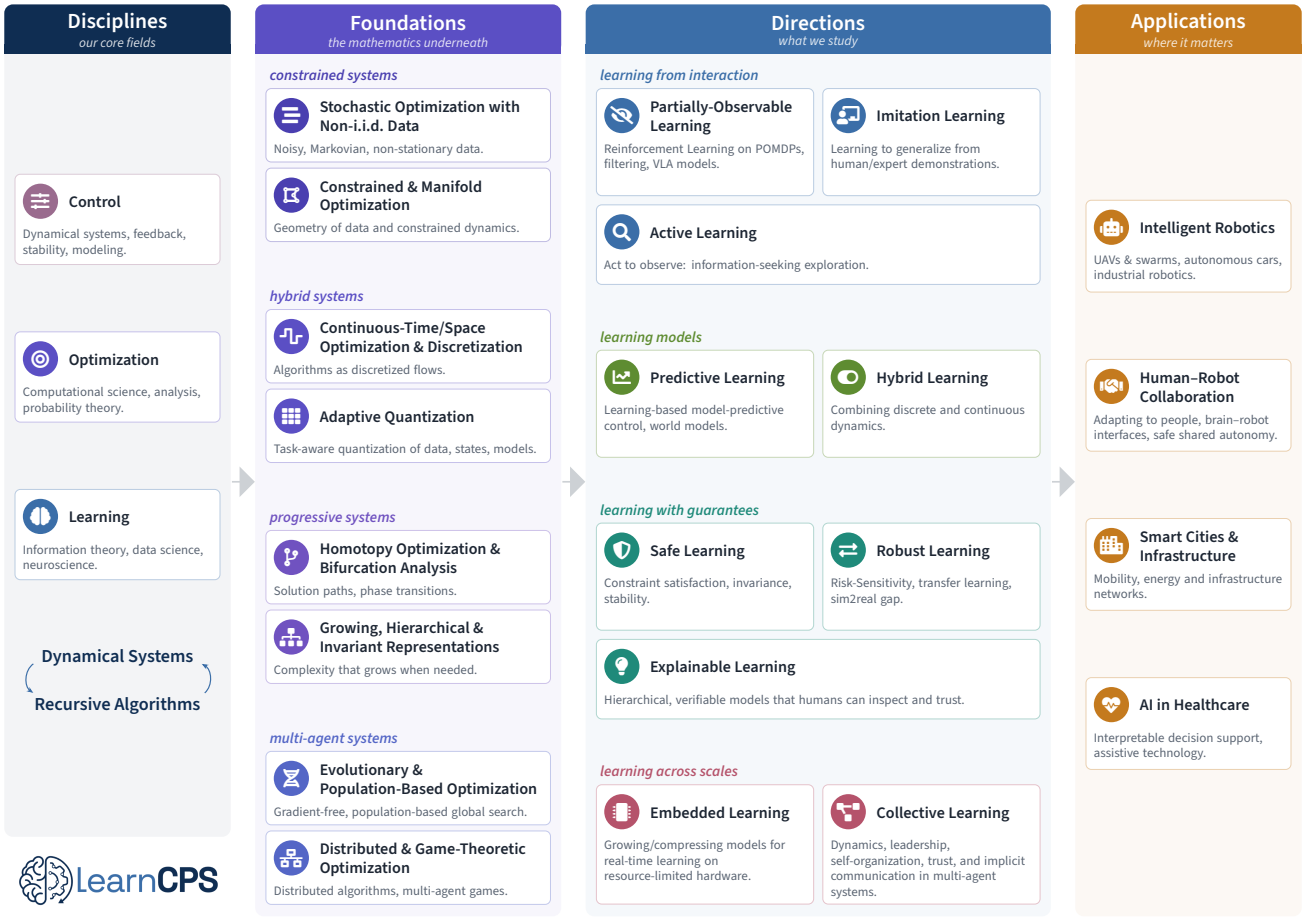


Figure 2. The LearnCPS research landscape: disciplines, foundations, research directions and applications.

Join us & collaborate learncps.eu