Control Meets AI: Automation, Autonomy, and Intelligent Machines

Control 2026
UKACC
Newcastle, UK
June 23, 2026

Pramod P. Khargonekar
EECS Department
University of California, Irvine

UCI University of California, Irvine

UCI Samueli School of Engineering
University of California, Irvine

Driving Question:

How can the controls field contribute to a positive future in the AI era and thereby thrive?

Outline

- Intelligent Machines
  - Historical Origins
  - Cognitive Cyber-Physical Systems
  - Physical AI
  - AI and Control

- Human-Machine Collaboration
  - Automation-Autonomy Framework
  - Insights from Cognitive Science

- Physical AI Safety

- Concluding Thoughts

Wiener and Macy Conferences: Cybernetics and Intelligent Machines

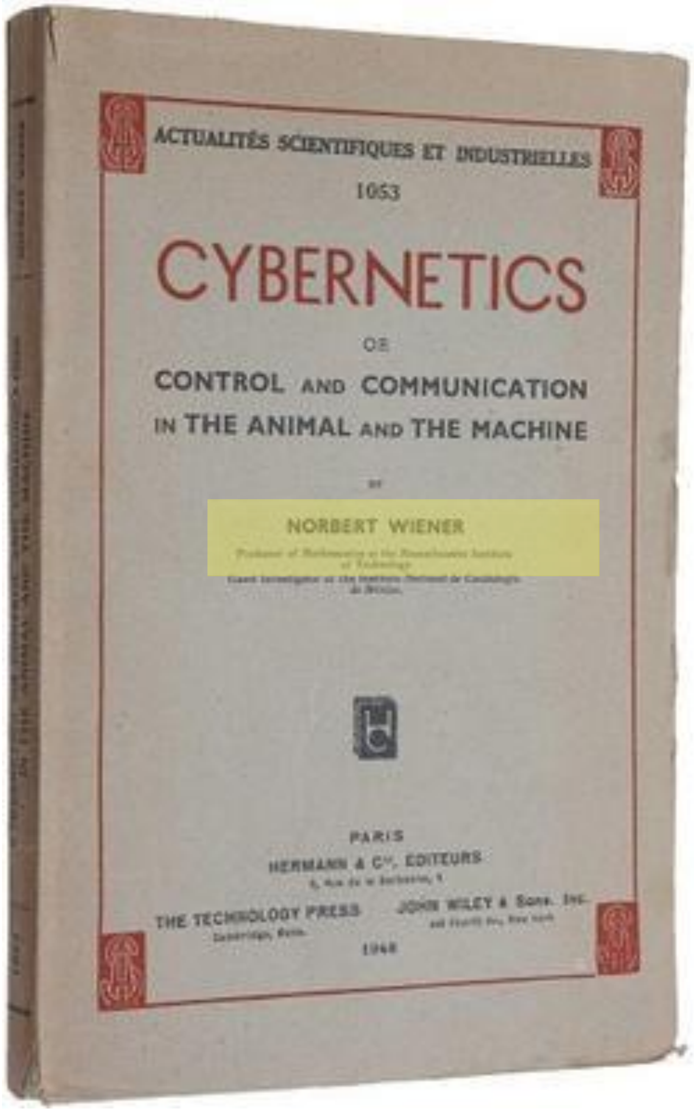

ACTUALITÉS SCIENTIFIQUES ET INDUSTRIELLES
1053

CYBERNETICS
OR
CONTROL AND COMMUNICATION
IN THE ANIMAL AND THE MACHINE

BY

NORBERT WIENER

PARIS
HERMANN & Cie, EDITEURS

THE TECHNOLOGY PRESS
JOHN WILEY & Sons, Inc.

1948

R. Kline, [Cybernetics, Automata Studies, and the Dartmouth Conference on Artificial Intelligence](), IEEE Annals of the History of Computing, 2011.

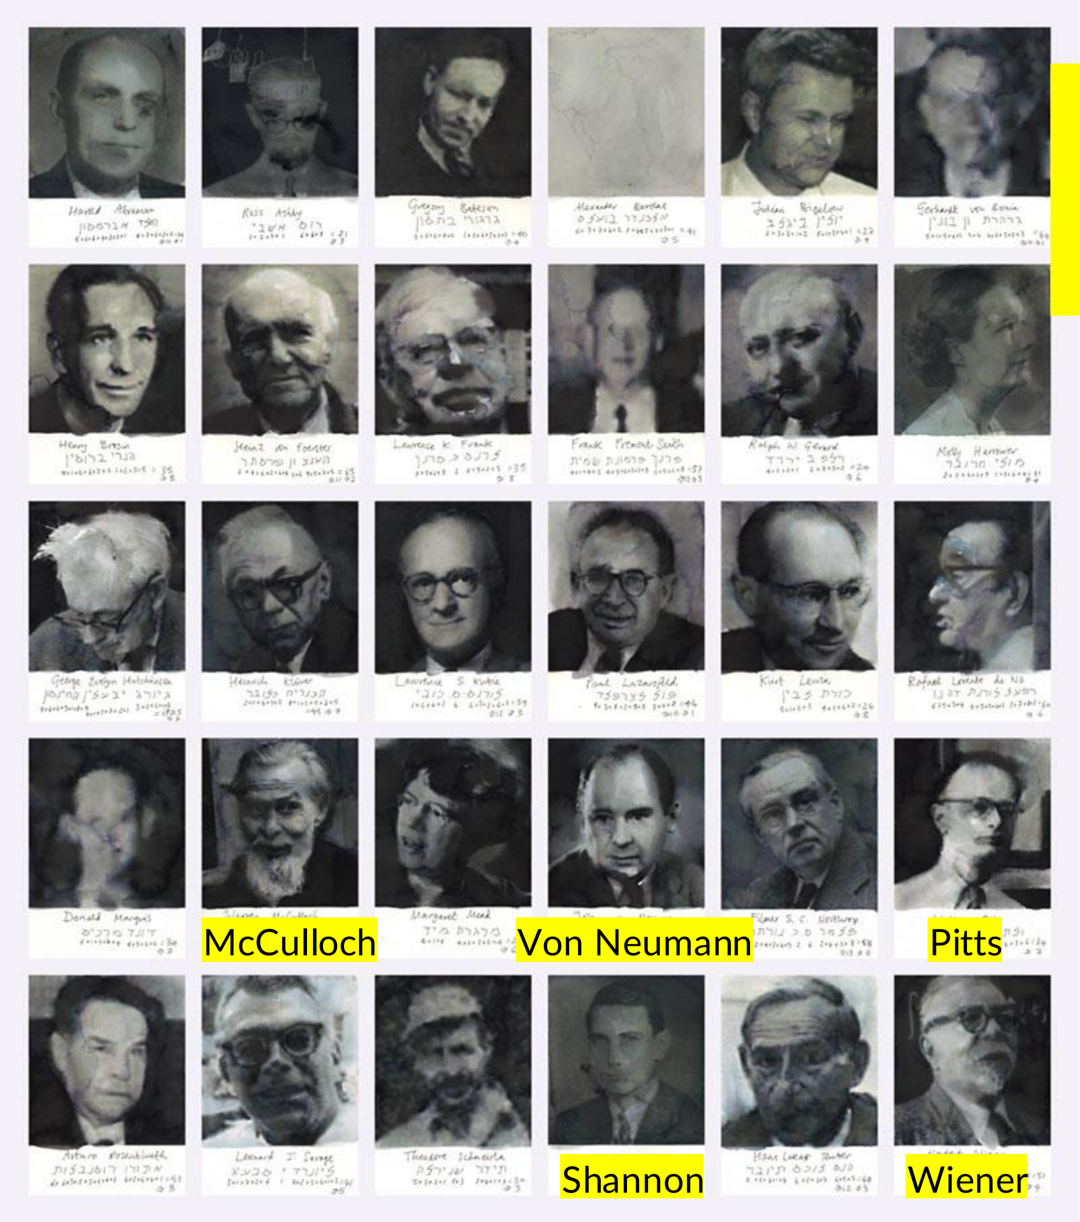

Macy Conferences 1943-53

Wiener
McCulloch
Pitts
von Neumann
Shannon

Wiener’s Foresight: The Age of Intelligent Machines

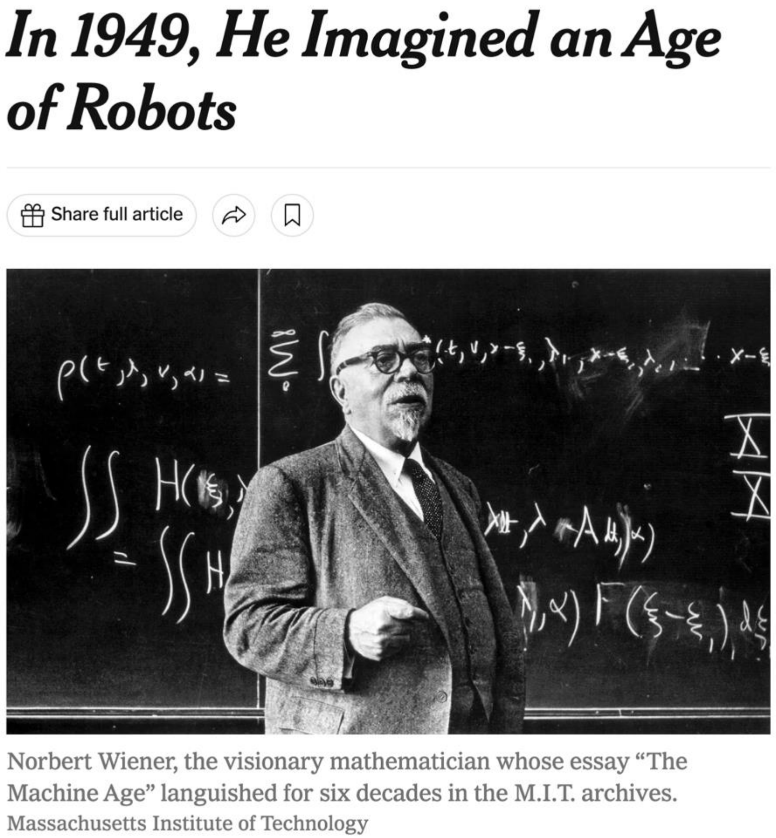

In 1949, He Imagined an Age of Robots

Share full article

Norbert Wiener, the visionary mathematician whose essay “The Machine Age” languished for six decades in the M.I.T. archives.
Massachusetts Institute of Technology

Source: J. Markoff, NY Times, May 20, 2013

“We have so far spoken of the computing machine as an analogue to the human nervous system rather than to the whole of the human organism. **Machines much more closely analogous to the human organism** are well understood, and are now on the verge of being built.”

“Finally the machines will do what we ask them to do and not **what we ought to ask them to do**.”

“Moreover, if we move in the direction of making machines which learn and whose **behavior is modified by experience**, we must face the fact that every degree of independence we give the machine is a degree of **possible defiance of our wishes**.”

Wiener, 1949

Intelligent Machines: AI + Engineered Cyber-Physical Systems

Definitions of AI

- “The science and engineering of making **intelligent machines**, especially intelligent computer programs.” — John McCarthy
- “A machine-based system that, **for explicit or implicit objectives, infers, from the input it receives**, how to generate outputs such as **predictions, content, recommendations, or decisions that can influence physical or virtual environments**. Different AI systems vary in their **levels of autonomy and adaptiveness** after deployment.” — Organization for Economic Cooperation and Development
- **Levels of AGI: Operationalizing Progress on the Path to AGI**, Morris et al, Google DeepMind, 2024

Jensen Huang @NVIDIA

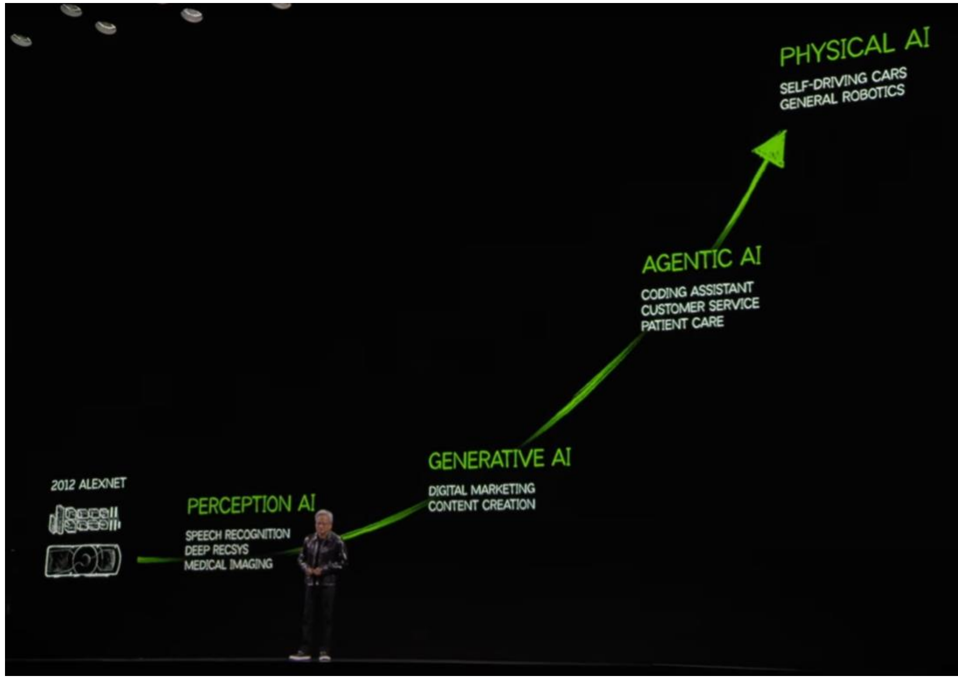

CES, 2025

Contextualizing Control: Cyber-Physical Systems

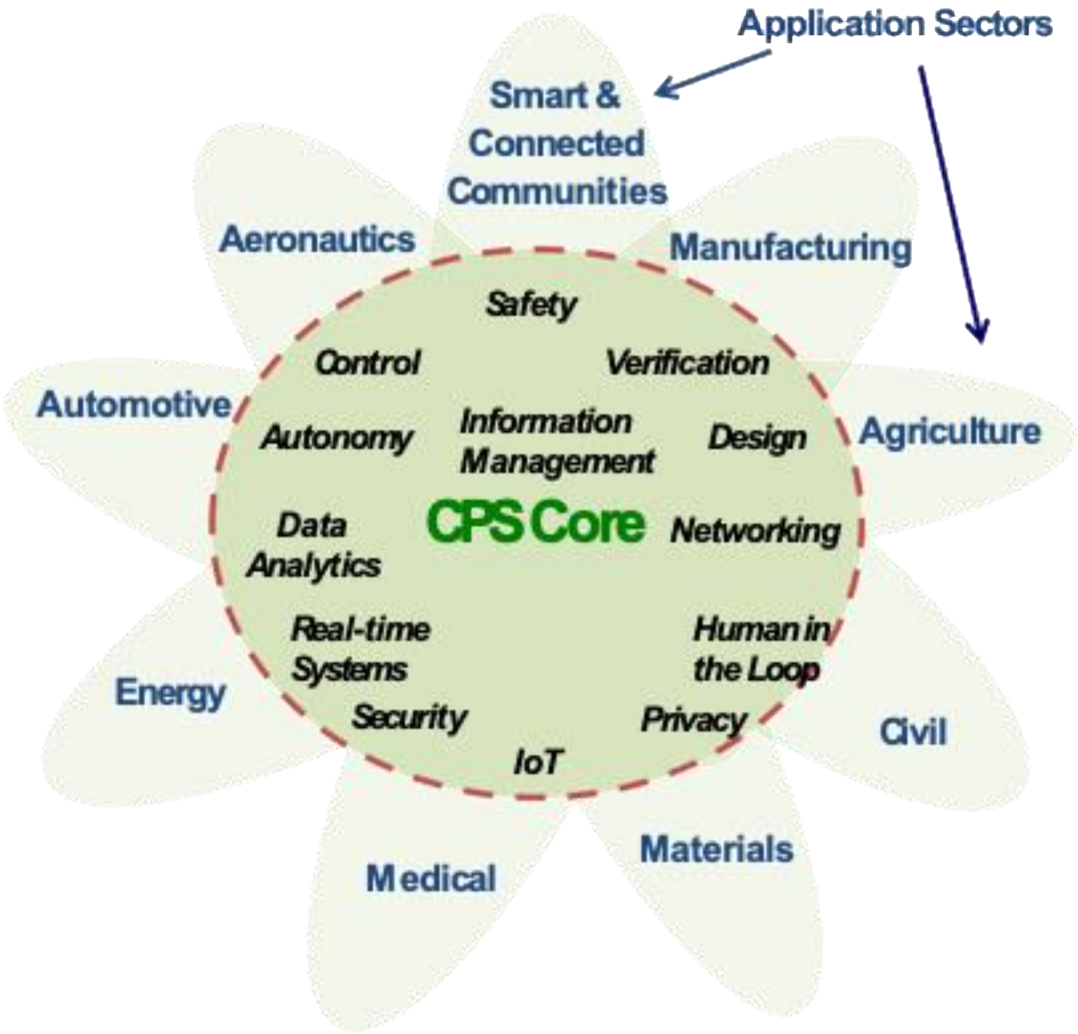

Application Sectors

Smart & Connected Communities
Aeronautics
Manufacturing
Automotive
Agriculture
Energy
Civil
Medical
Materials

CPS Core
Safety
Control
Verification
Autonomy
Information Management
Design
Data Analytics
Networking
Real-time Systems
Human in the Loop
Security
Privacy
IoT

Application Domains

Transportation

- Faster and safer vehicles (airplanes, cars, etc)
- Improved use of airspace and roadways
- Energy efficiency
- Manned and un-manned

Energy and Industrial Automation

- Homes and offices that are more energy efficient and cheaper to operate
- Distributed micro-generation for the grid

Healthcare and Biomedical

- Increased use of effective in-home care
- More capable devices for diagnosis
- New internal and external prosthetics

Critical Infrastructure

- More reliable power grid
- Highways that allow denser traffic with increased safety

Source: NSF

Emerging Visions: Smart, Intelligent, Autonomous X

- Smart Manufacturing (Industry 5.0)
- Smart Electric Grids
- Smart and Connected Health
- Intelligent Transportation
- Smart (Precision) Agriculture

- Self-driving Cars
- Unmanned Air Vehicles
- Unmanned Underwater Vehicles
- Robots

Cyber-Physical Systems are Systems of Systems

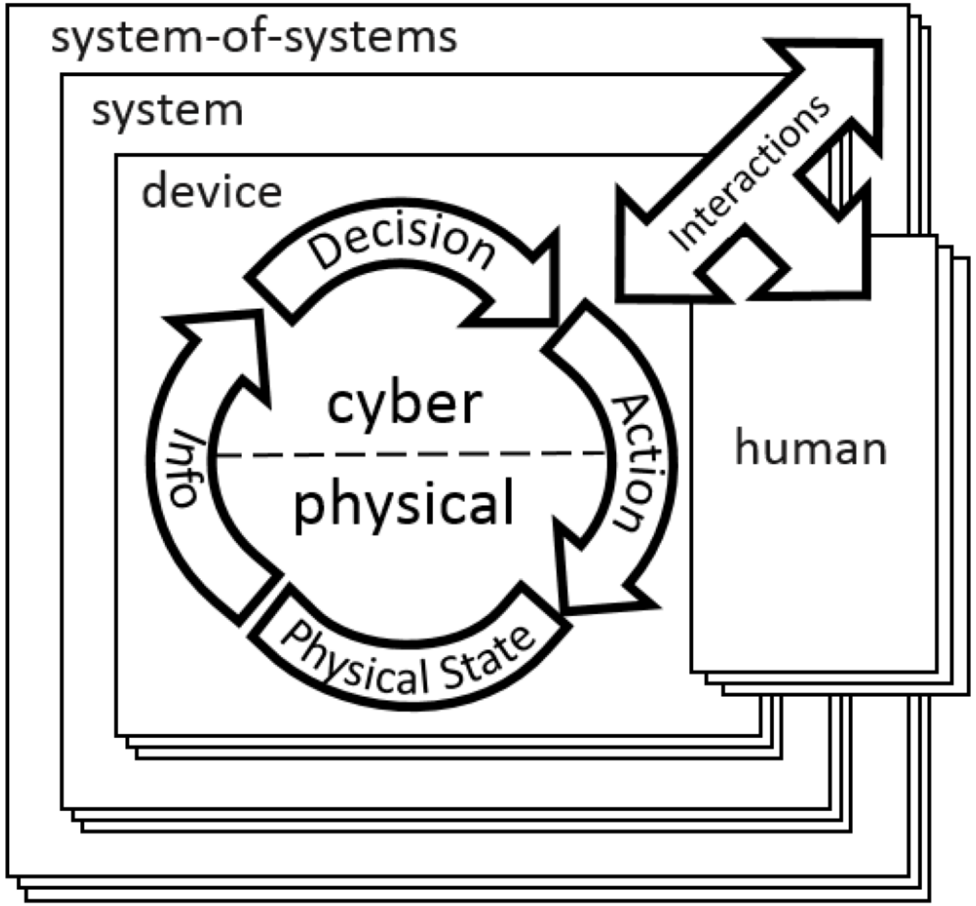

Figure 1: CPS Conceptual Model

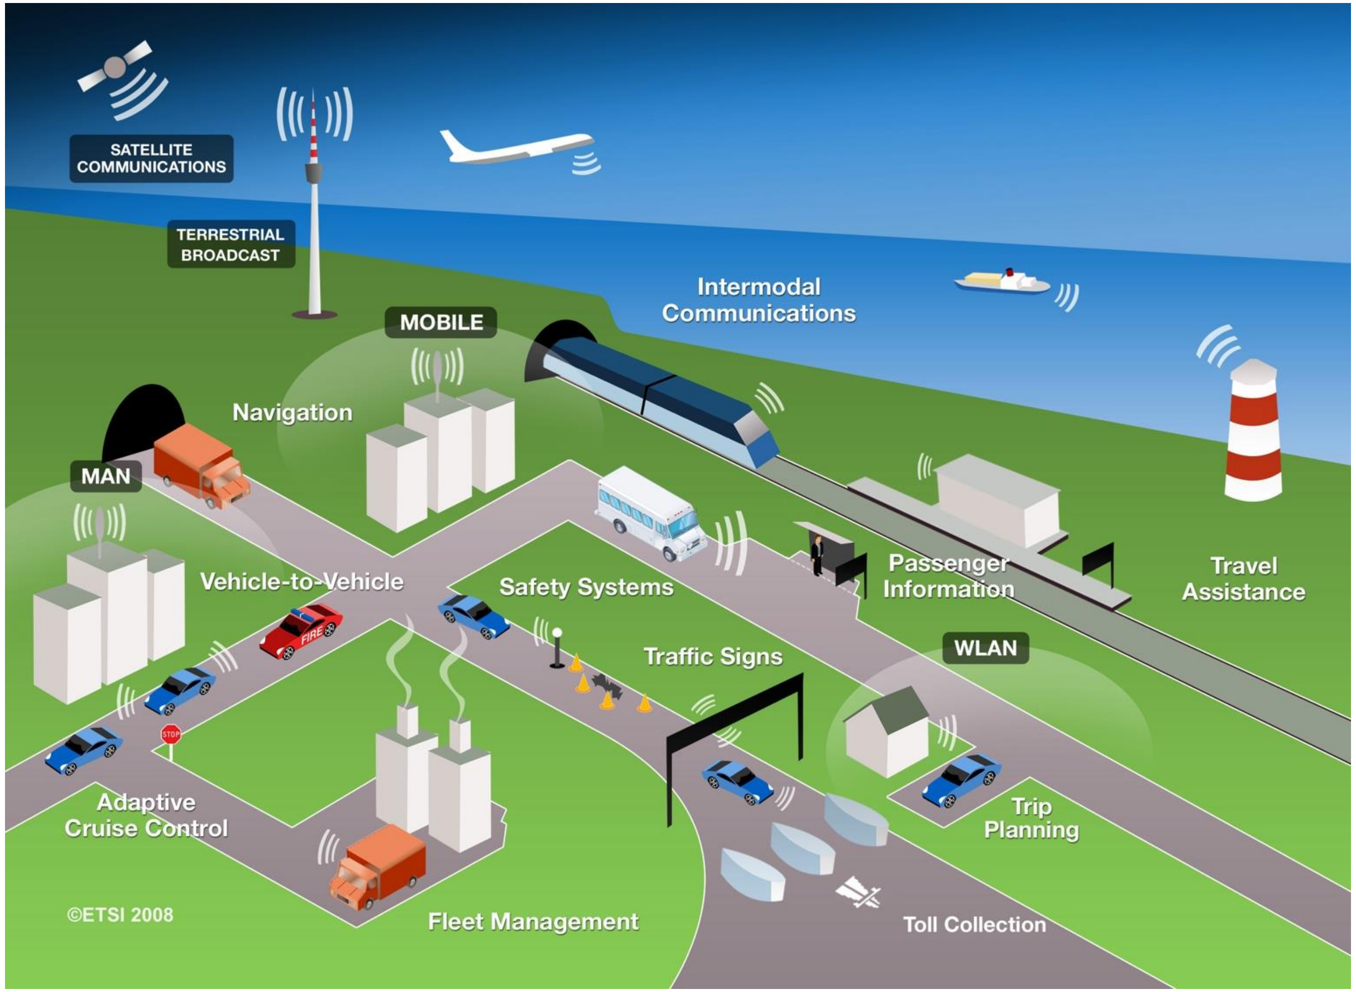

Source: NIST, 2016

Source: ETSI, 2008

Cyber-Physical Systems Meet Humans and AI

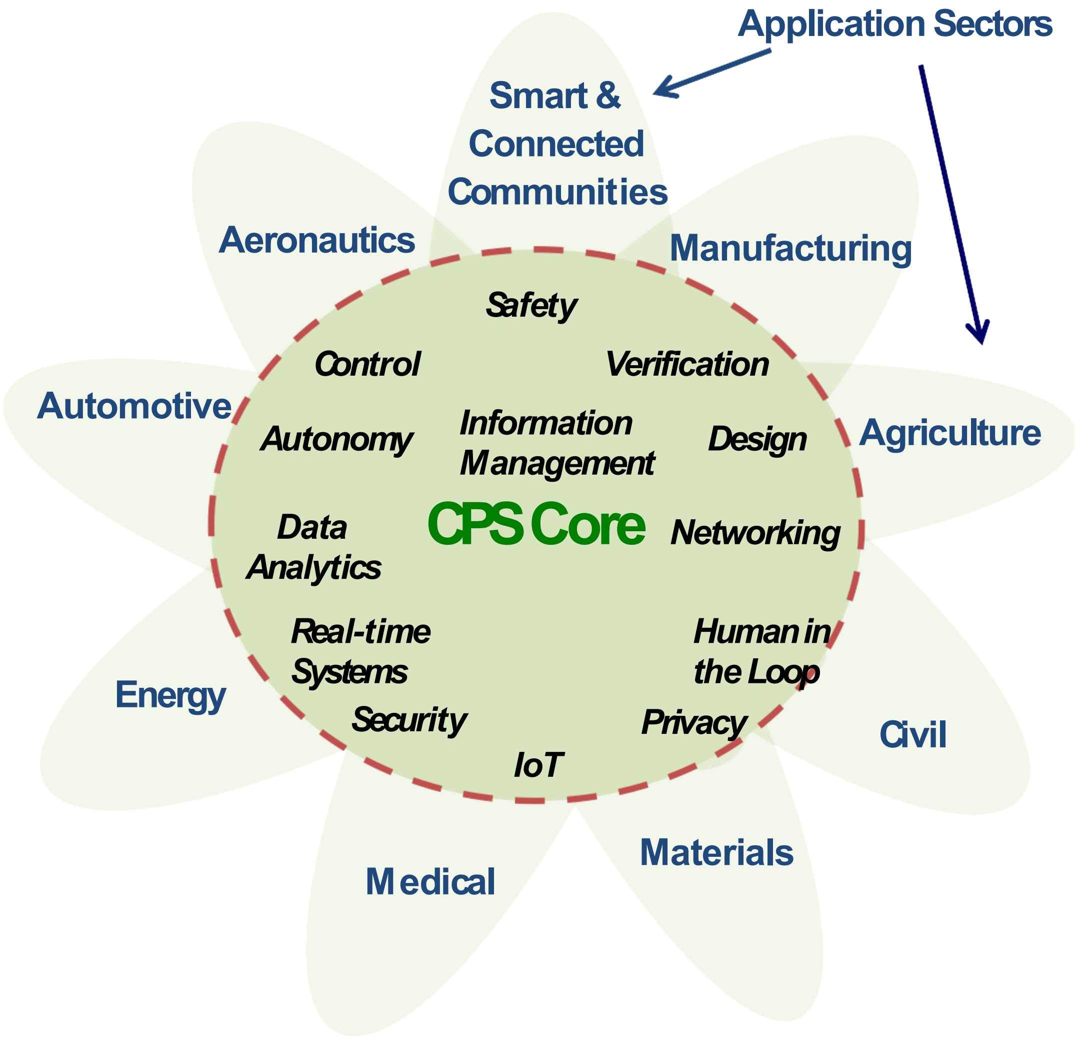

Application Sectors

Smart & Connected Communities

Aeronautics

Manufacturing

Automotive

Agriculture

Energy

Civil

Medical

Materials

Safety

Control

Verification

Autonomy

Information Management

Design

Data Analytics

CPS Core

Networking

Real-time Systems

Human in the Loop

Security

Privacy

IoT

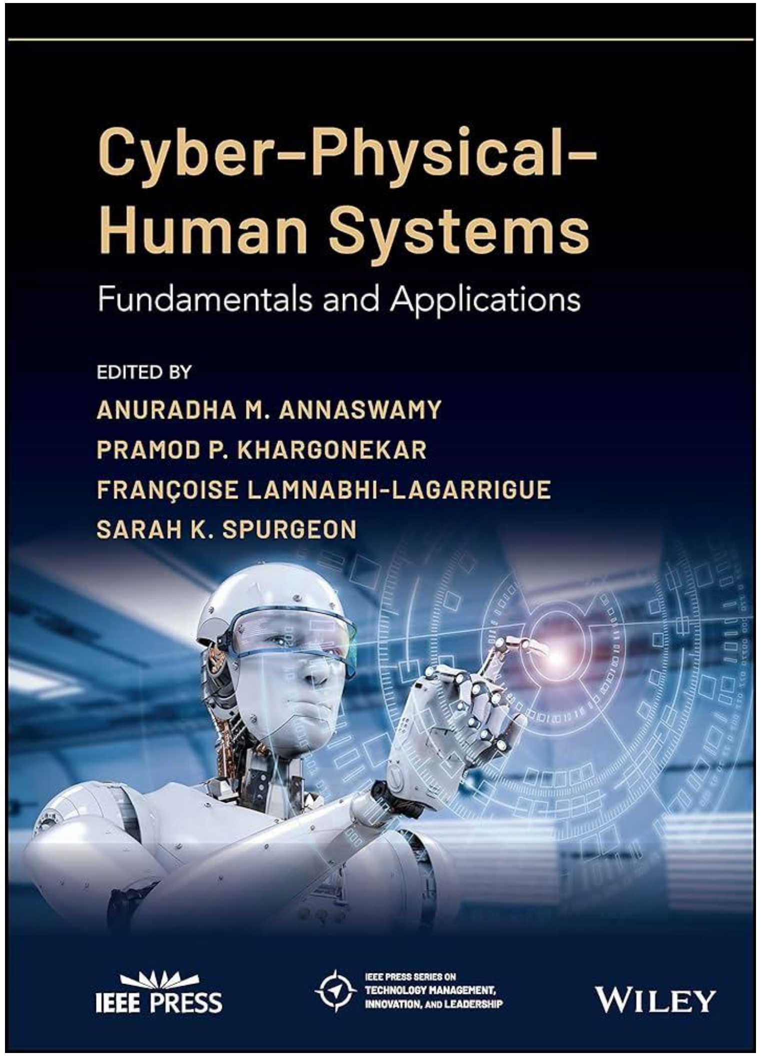

Cyber–Physical–Human Systems

Fundamentals and Applications

EDITED BY

ANURADHA M. ANNASWAMY

PRAMOD P. KHARGONEKAR

FRANÇOISE LAMNABHI-LAGARRIGUE

SARAH K. SPURGEON

IEEE PRESS

IEEE PRESS SERIES ON TECHNOLOGY MANAGEMENT, INNOVATION, AND LEADERSHIP

WILEY

\+ AI

Source: NSF

AI Engineering Vision

AI Engineering: Convergence of AI and Engineering as a bi-directional flow leading to a virtuous cycle:

- **Engineering for AI:** Engineering disciplines enabling safer, more reliable, and trustworthy AI-integrated cyber-physical systems

- **AI for Engineering:** AI technology and advances impacting engineering across the board: in research, design, manufacturing, and operations

erVa

NSF Engineering Research Visioning Alliance

AI Engineering

A Strategic Research Framework to Benefit Society

Visioning Event Report

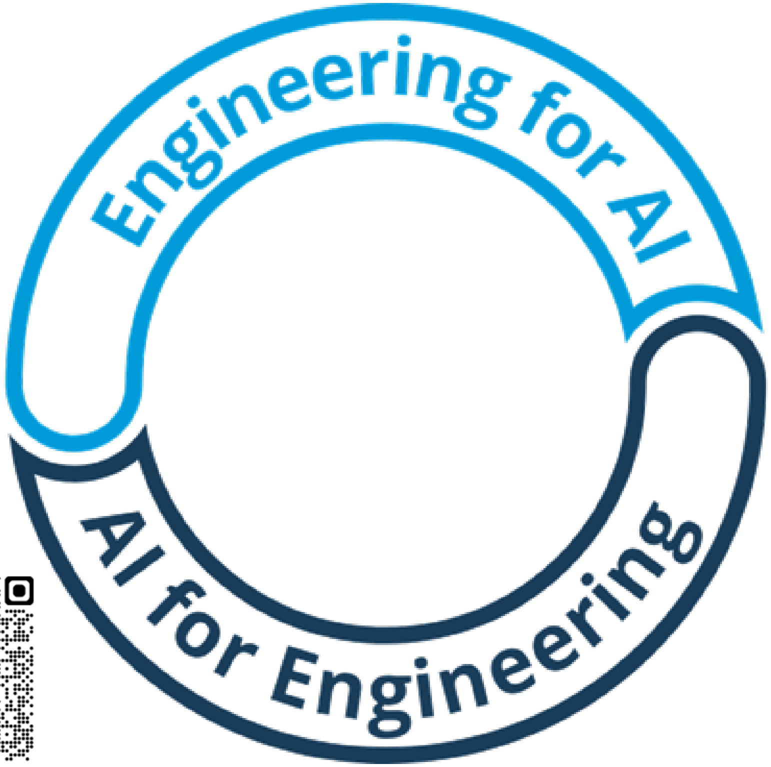

https://www.ervacommunity.org/visioning-report/report-ai-engineering-a-strategic-research-framework-to-benefit-society/

Cognitive CPS: A Concept for Intelligent Machines

Working Definition: Cognitive CPS =
Cyber-physical systems with cognitive properties or capabilities

ELSEVIER

Annual Reviews in Control
Volume 45, 2018, Pages 1-4

Vision article

Advancing systems and control research in the era of ML and AI

Pramod P. Khargonekar, Munther A. Dahleh

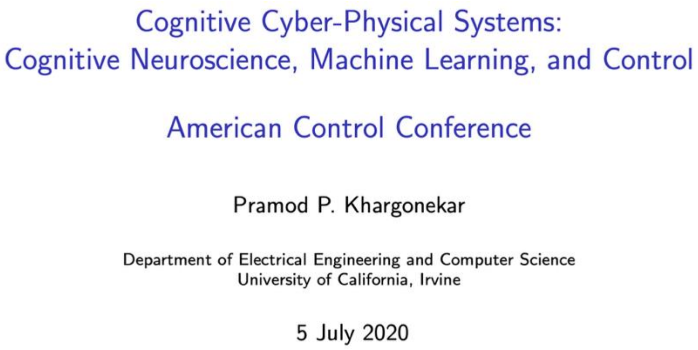

Cognitive Cyber-Physical Systems:
Cognitive Neuroscience, Machine Learning, and Control

American Control Conference

Pramod P. Khargonekar

Department of Electrical Engineering and Computer Science
University of California, Irvine

5 July 2020

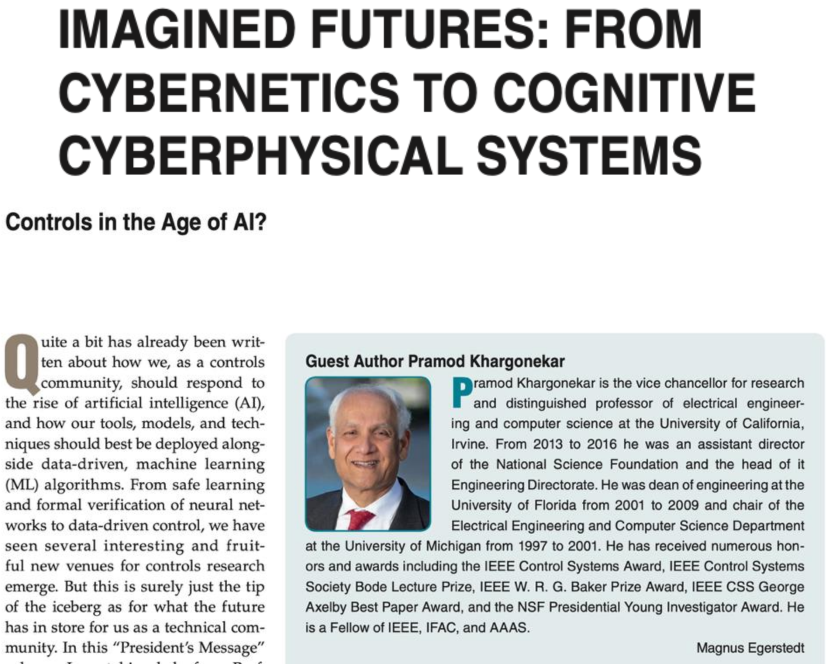

IMAGINED FUTURES: FROM CYBERNETICS TO COGNITIVE CYBERPHYSICAL SYSTEMS

Controls in the Age of AI?

Quite a bit has already been written about how we, as a controls community, should respond to the rise of artificial intelligence (AI), and how our tools, models, and techniques should best be deployed alongside data-driven, machine learning (ML) algorithms. From safe learning and formal verification of neural networks to data-driven control, we have seen several interesting and fruitful new venues for controls research emerge. But this is surely just the tip of the iceberg as for what the future has in store for us as a technical community. In this "President's Message"

Guest Author Pramod Khargonekar

Pramod Khargonekar is the vice chancellor for research and distinguished professor of electrical engineering and computer science at the University of California, Irvine. From 2013 to 2016 he was an assistant director of the National Science Foundation and the head of it Engineering Directorate. He was dean of engineering at the University of Florida from 2001 to 2009 and chair of the Electrical Engineering and Computer Science Department at the University of Michigan from 1997 to 2001. He has received numerous honors and awards including the IEEE Control Systems Award, IEEE Control Systems Society Bode Lecture Prize, IEEE W. R. G. Baker Prize Award, IEEE CSS George Axelby Best Paper Award, and the NSF Presidential Young Investigator Award. He is a Fellow of IEEE, IFAC, and AAAS.

Magnus Egerstedt

IEEE CSM, 2023

https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10186028

What is Cognition?

- “all processes by which the sensory input is transformed, reduced, elaborated, stored, recovered, and used.” – Neisser, Cognitive Psychology, 1967
- “study of higher mental processes such as **attention, language use, memory, perception, problem solving, and thinking**.” – American Psychological Association
- Cognition is a **dynamic process**
- Connections to and inspiration from neuroscience

How Might We Create Cognitive CPS?

- Design cognitive capabilities into CPS
- CPS “learns” over time to possess cognitive capabilities
- Combination of the above

Integration of AI into CPS to build cognitive CPS may take many different trajectories!

Perception

- CPS with suite of sensors
- Learning + state estimation + sensor fusion
- Major area of research in ML/AI and autonomous systems

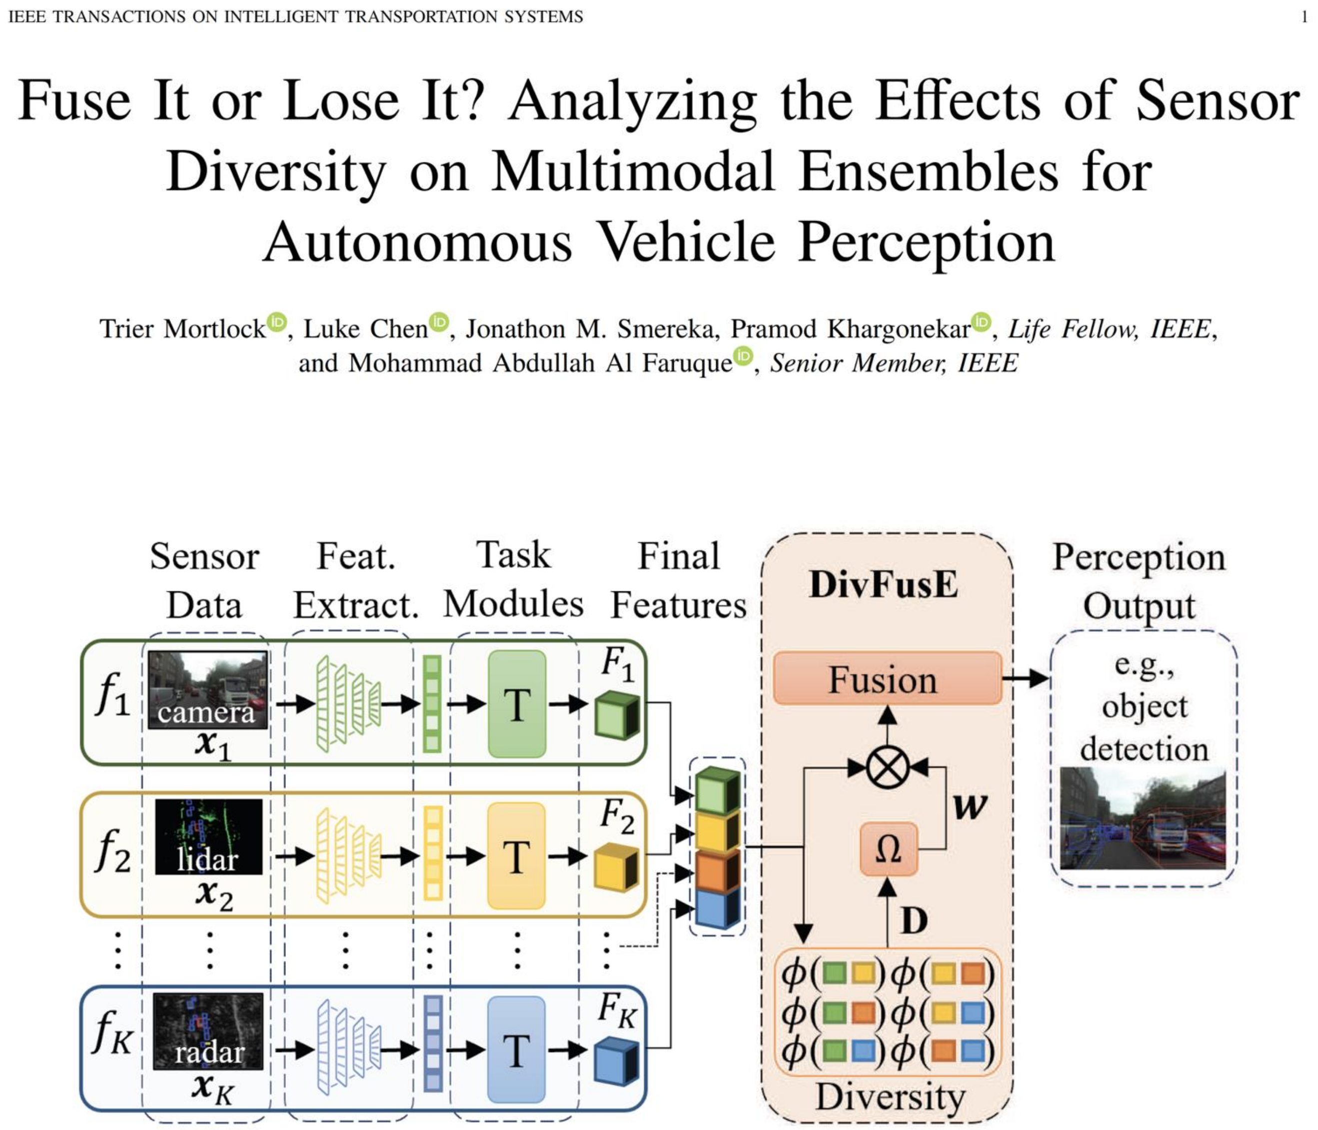

IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS 1

Fuse It or Lose It? Analyzing the Effects of Sensor Diversity on Multimodal Ensembles for Autonomous Vehicle Perception

Trier Mortlock, Luke Chen, Jonathon M. Smereka, Pramod Khargonekar, *Life Fellow, IEEE*,
and Mohammad Abdullah Al Faruque, *Senior Member, IEEE*

Memory

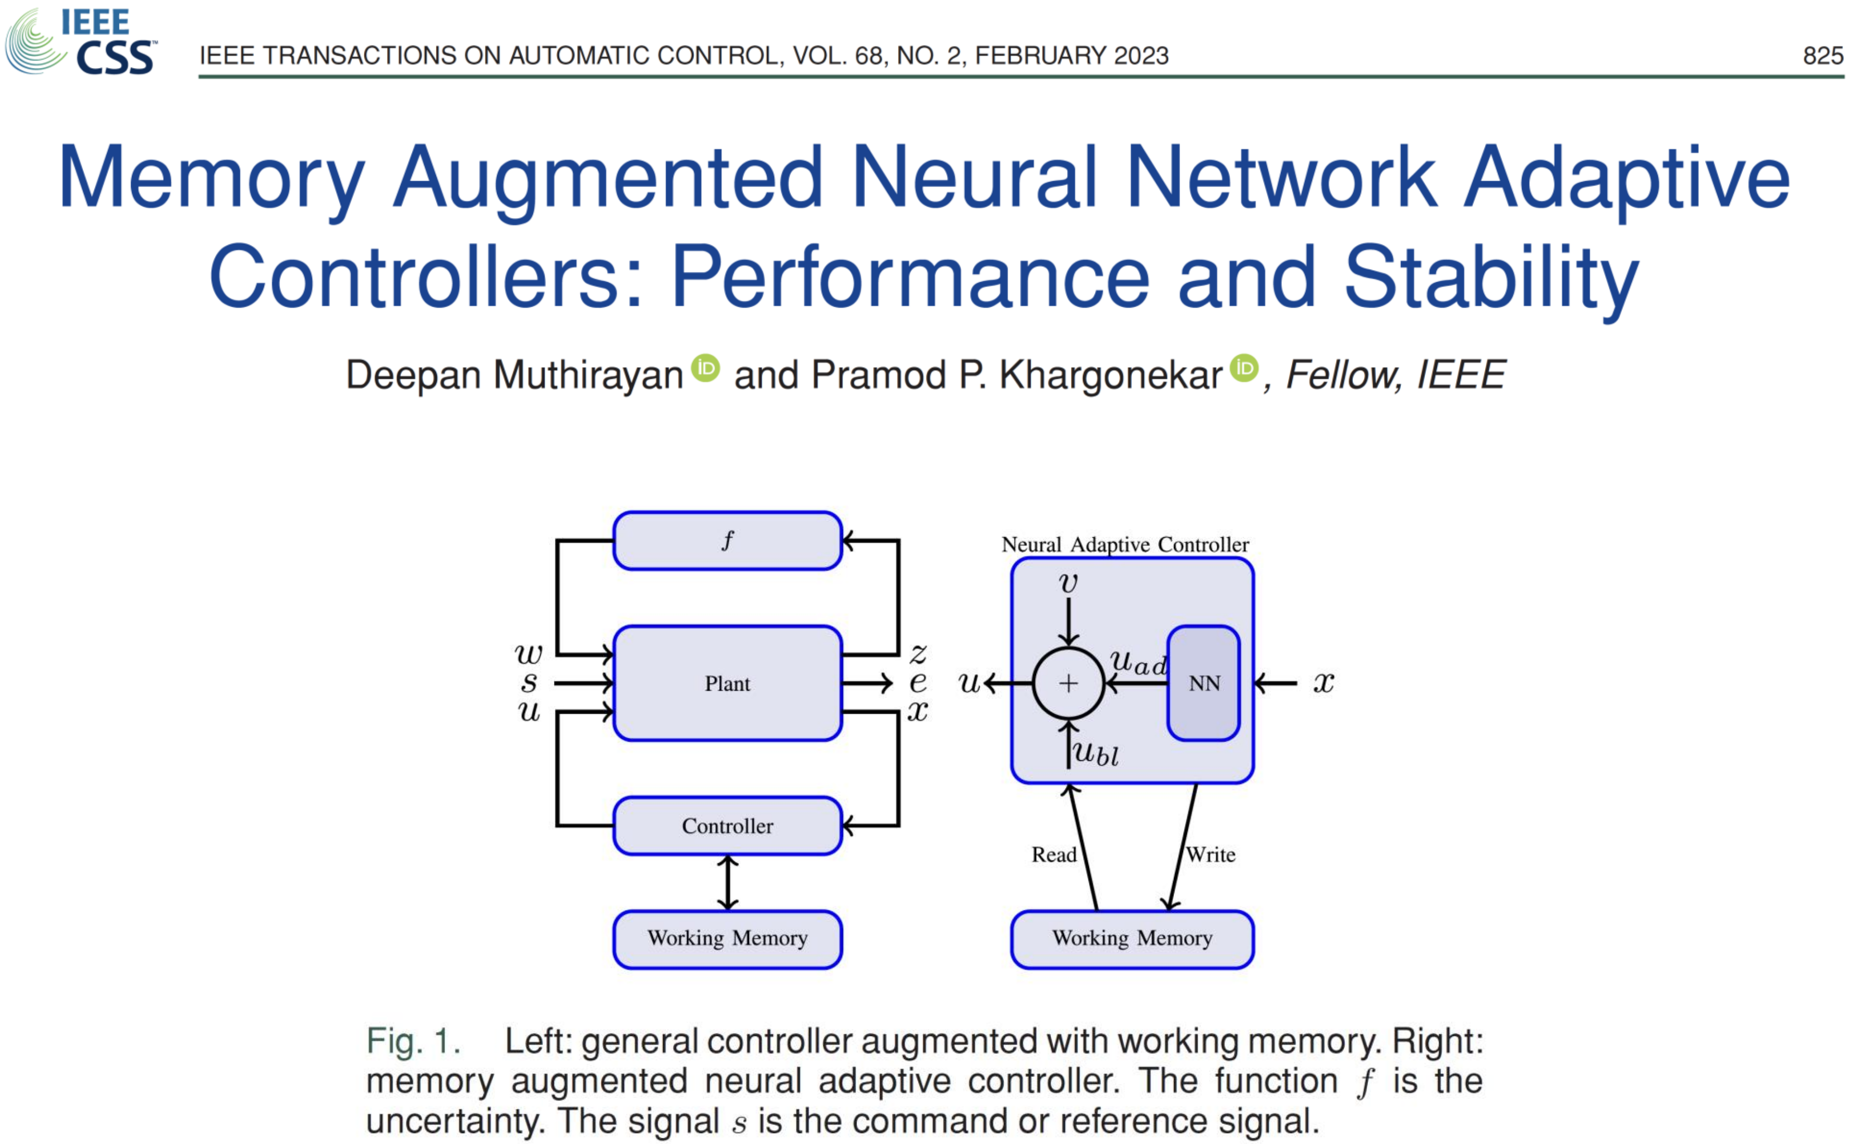

IEEE CSS IEEE TRANSACTIONS ON AUTOMATIC CONTROL, VOL. 68, NO. 2, FEBRUARY 2023 825

Memory Augmented Neural Network Adaptive Controllers: Performance and Stability

Deepan Muthirayan and Pramod P. Khargonekar, *Fellow, IEEE*

Fig. 1. Left: general controller augmented with working memory. Right: memory augmented neural adaptive controller. The function f is the uncertainty. The signal s is the command or reference signal.

Intelligent Machines: Physical AI

Jensen Huang @NVIDIA

In collaboration with
Boston Consulting Group

WORLD
ECONOMIC
FORUM

Physical AI: Powering the New Age of Industrial Operations

WHITE PAPER

SEPTEMBER 2025

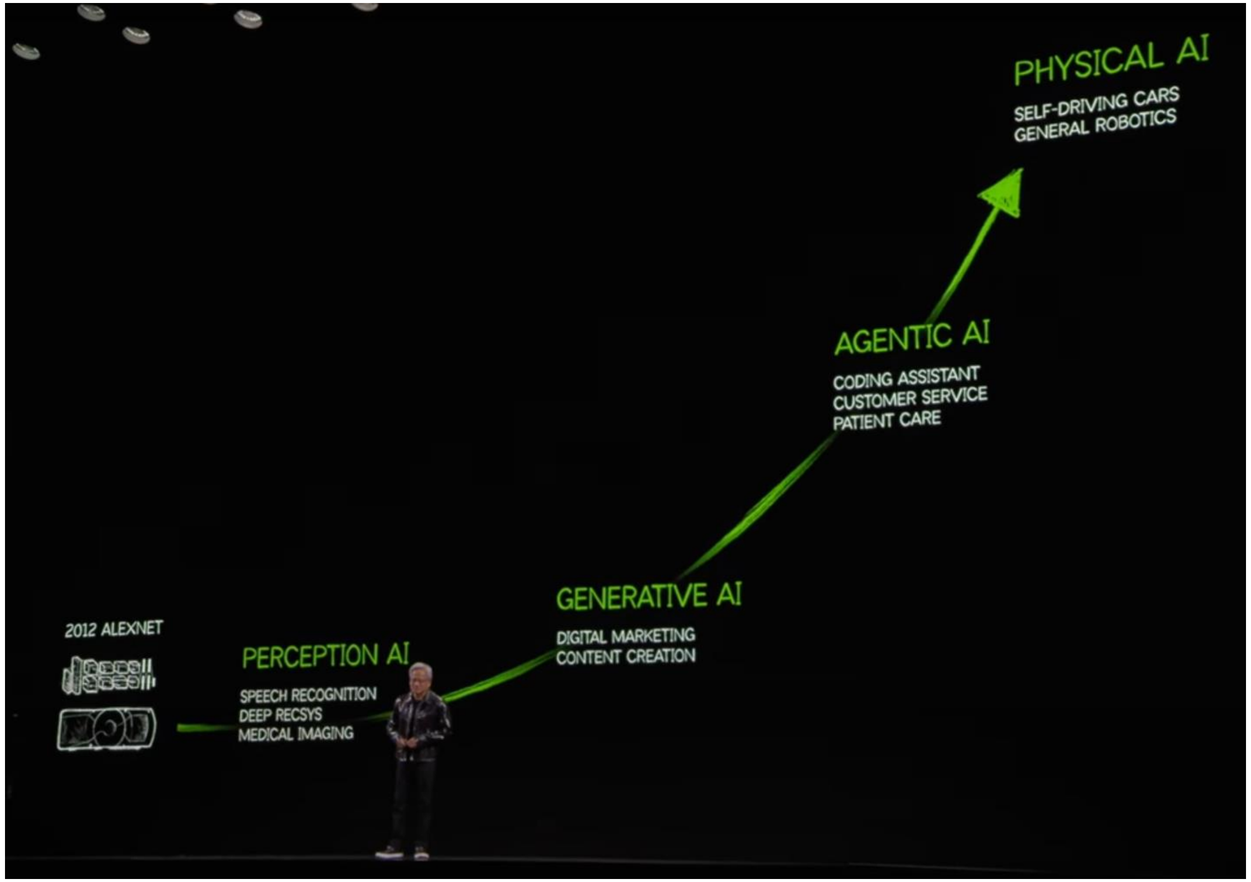

Physical AI: Manufacturing

“Using an “AI Brain,” — powered by software-defined automation and industrial operations software, combined with NVIDIA Omniverse™ libraries and NVIDIA AI infrastructure, **factories can continuously analyze their digital twins, test improvements virtually and turn validated insights into operational changes on the shopfloor.**”

Siemens and NVIDIA Expand Partnership to Build the Industrial AI Operating System, January 6, 2026

Physical AI: Robotics

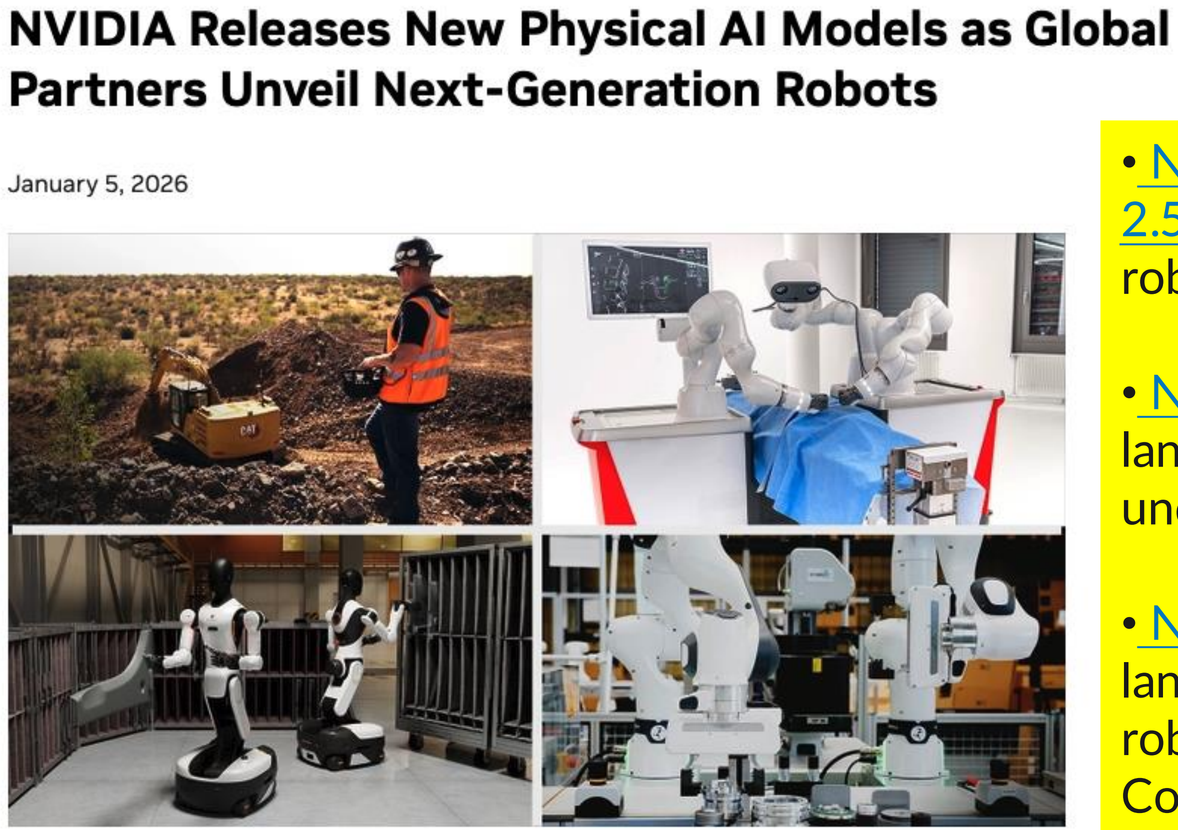

NVIDIA Releases New Physical AI Models as Global Partners Unveil Next-Generation Robots

January 5, 2026

- NVIDIA Cosmos™ Transfer 2.5 and NVIDIA Cosmos Predict 2.5 — enable physically based synthetic data generation and robot policy evaluation in simulation for physical AI.

- NVIDIA Cosmos Reason 2, an open reasoning vision language model (VLM) that enables intelligent machines to see, understand and act in the physical world like humans.

- NVIDIA Isaac™ GR00T N1.6, an open reasoning vision language action (VLA) model, purpose-built for humanoid robots, that unlocks full body control and uses NVIDIA Cosmos Reason for better reasoning and contextual understanding.

COSMOS World Foundation Model for Physical AI

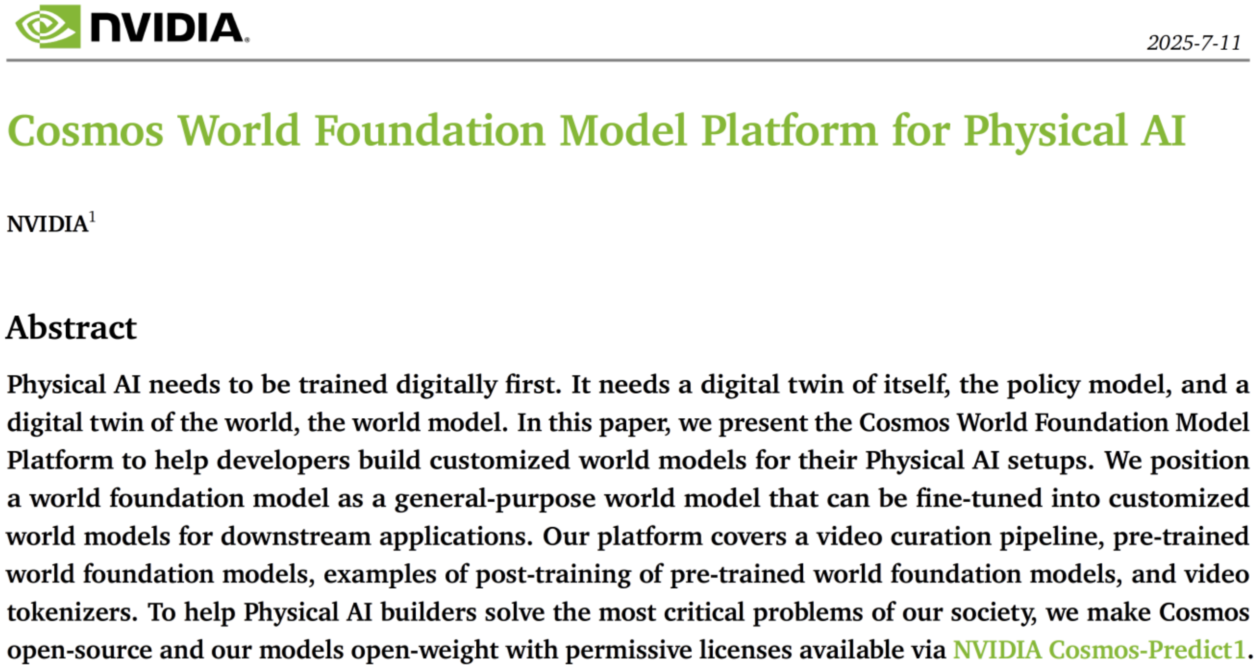

2025-7-11

Cosmos World Foundation Model Platform for Physical AI

NVIDIA[1]

Abstract

Physical AI needs to be trained digitally first. It needs a digital twin of itself, the policy model, and a digital twin of the world, the world model. In this paper, we present the Cosmos World Foundation Model Platform to help developers build customized world models for their Physical AI setups. We position a world foundation model as a general-purpose world model that can be fine-tuned into customized world models for downstream applications. Our platform covers a video curation pipeline, pre-trained world foundation models, examples of post-training of pre-trained world foundation models, and video tokenizers. To help Physical AI builders solve the most critical problems of our society, we make Cosmos open-source and our models open-weight with permissive licenses available via NVIDIA Cosmos-Predict1.

Figure 3: A world foundation model (WFM) $\mathcal{W}$ is a model that generates the future state of the world x_{t+1} based on the past observations $x_{0:t}$ and current perturbation c_t.

Source: Agarwal, et al, 2025

Robotics: Classical Approach

- Robotics is an old field going back to the 60’s

- **Traditional approach to robotics control:** first principles modeling, forward kinematics, inverse kinematics, trajectory planning and generation, nonlinear robust control for trajectory following, higher level logic and discrete event control

- Integration into larger systems to perform valuable tasks in applications such as manufacturing, biomedical, construction, hazardous environments, search and rescue, …

Robot Control: Changing Nature of Robotics

- The classical model-based control approach has been successful but has important limitations

- **Systems composition and integration challenges:** engineering overhead, compounding of errors, in **adapting or modifying an existing set of “building blocks” to new tasks or robotic platforms**

- **Planning and Tasks Related Challenges:** limited scalability to **multimodal (RGB, depth, tactile, audio) data**; adoption of per-task planners, goal parameterizations,
- and safety constraints lead to a **large increase in design and validation options**

- **Modeling Related Challenges:** contact, friction, and compliance, deformability, …

- **Limited Ways for Leveraging Data: w**hile available data may become increasingly available, there are limits to which they can be leveraged

Robot Learning: A Tutorial by Capuano et al

“This shift from classical, model-based methods to data-driven, learning-based paradigms is unlocking unprecedented capabilities in autonomous systems”

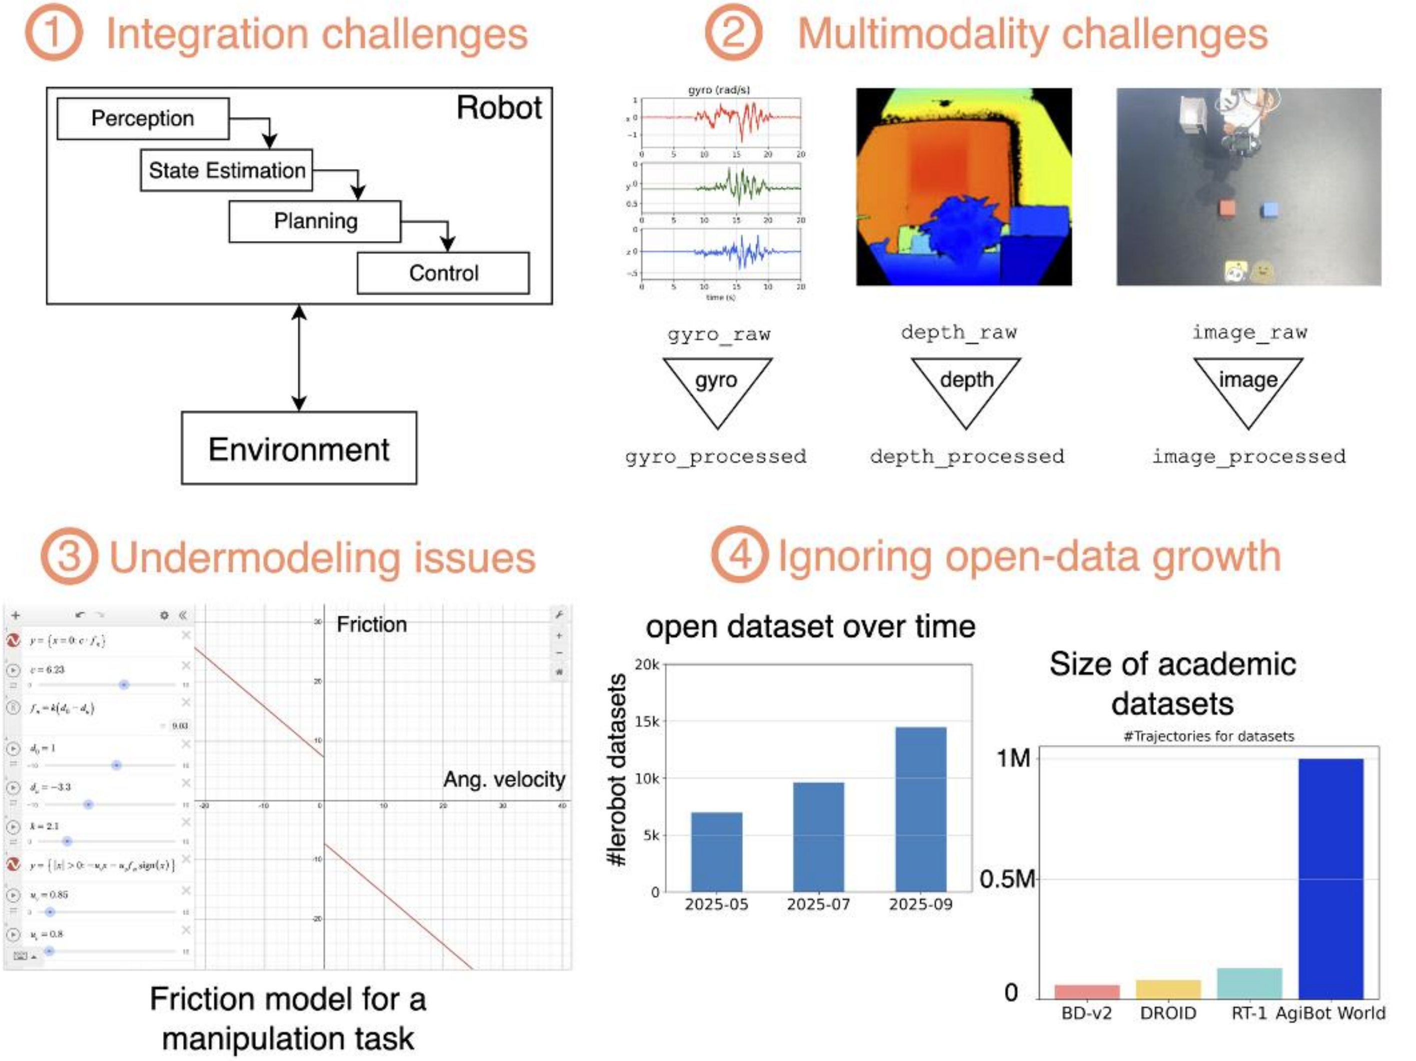

Figure 8 | Dynamics-based approaches to robotics suffer from several limitations: (1) orchestrating multiple components poses integration challenges; (2) the need to develop custom processing pipelines for the sensing modalities and tasks considered hinders scalability; (3) simplified analytical models of physical phenomena (here friction at the gripper; credits to Antonova et al. (2017)) limit real-world performance. Lastly, (4) dynamics-based methods overlook trends in the availability and growth of robotics data.

Capuano et al, Robot Learnin, arXiv:2510.12403v1 [cs.RO] 14 Oct 2025

Recent Trends in Robot Learning

- Sim-to-Real Transfer: Domain Randomization
- Sample Efficient Reinforcement Learning
- Behavior Cloning and Imitation learning
- Diffusion based Methods for Robot Learning
- Foundation Models for Robotics: Gemini Robotics, NVIDIA platform, Embedded Multimodal Large Models (EMLMs), …

“I'm very confident that our knowledge of dynamics and control will help us penetrate the new vistas enabled by these high-capacity models.” R. Tedrake, Imitation Learning, Chapter 21, Underactuated Robotics, 2025.

How will Humans Collaborate with “Intelligent Machines”?

Human-Machine Relationship Across Technological Eras

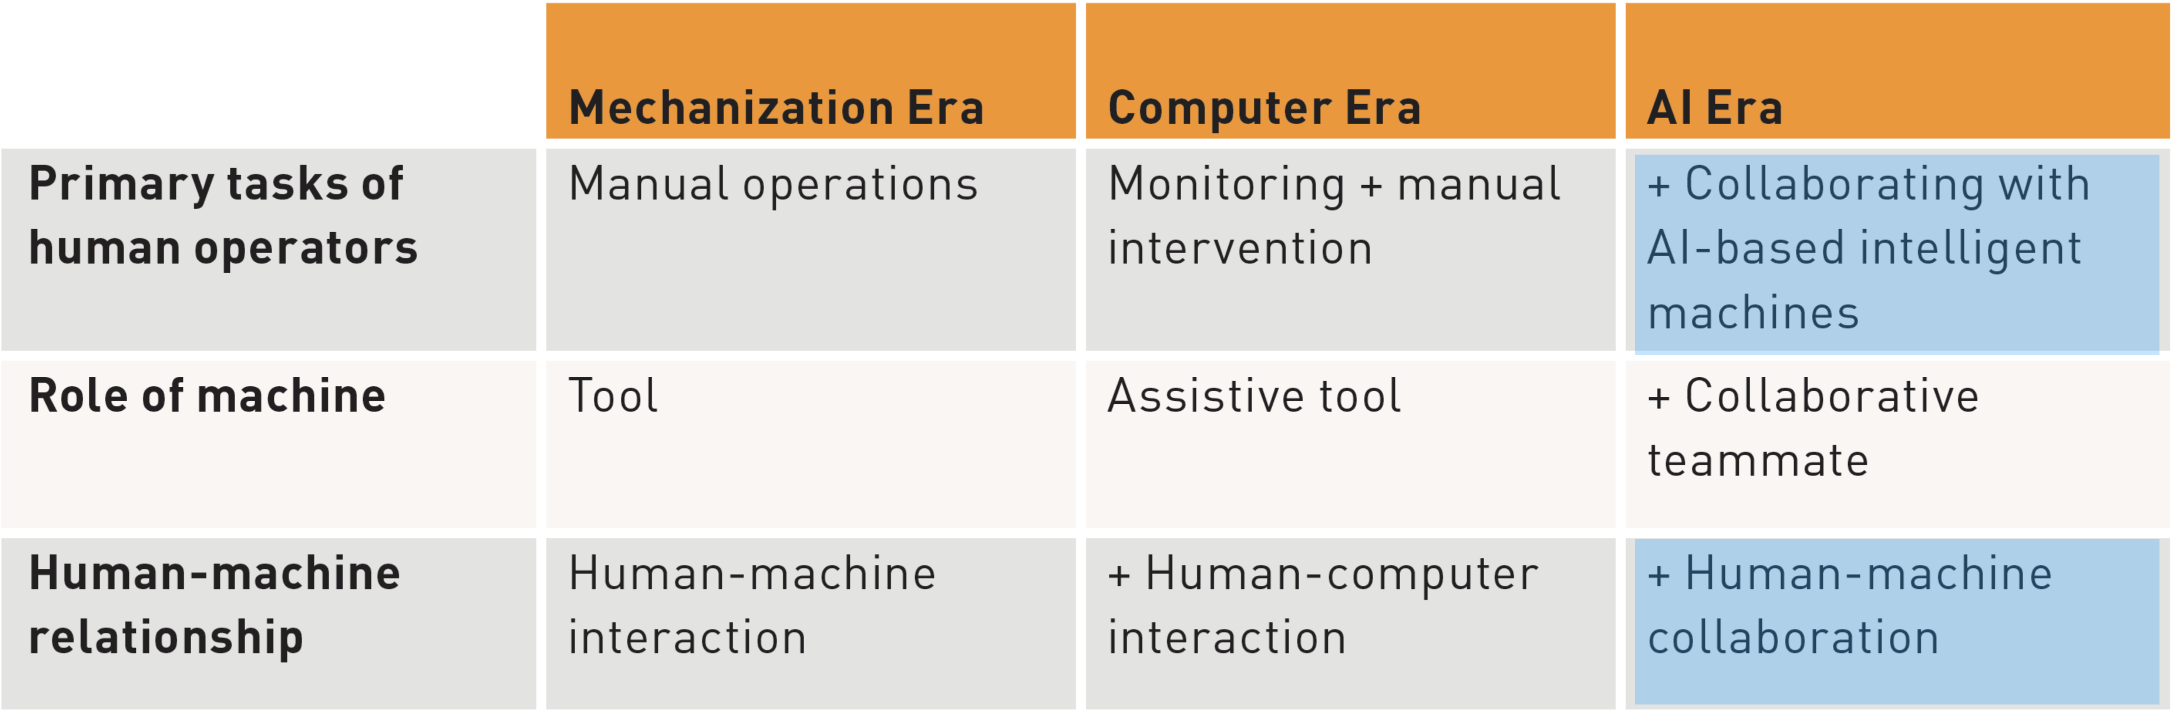

	Mechanization Era	Computer Era	AI Era
Primary tasks of human operators	Manual operations	Monitoring + manual intervention	+ Collaborating with AI-based intelligent machines
Role of machine	Tool	Assistive tool	+ Collaborative teammate
Human-machine relationship	Human-machine interaction	+ Human-computer interaction	+ Human-machine collaboration

Table 2. Human-machine relationships across technological eras.

Source: [Wei Xu, ACM Interactions, 2021]

Central Issue for Human-Intelligent Machine Collaboration: Placement of Intelligent Machines on the Autonomy-Automation Axes

Automation

	Transparent Deterministic	Non-Transparent Deterministic	Non-Transparent Non-Deterministic Closed	Adaptable Closed	Adaptable Open
Fully Automated					
Informs Human					
Human Veto					
Human Approval					
Human Decision					

Technical Autonomy

Simmler and Frischknecht, A taxonomy of human–machine collaboration: capturing automation and technical autonomy, AI and Society, 2021

Third Axis: The CPS Level

- Automation-Autonomy axes need to be complemented with a third dimension
- CPS Level: Device, System, and Systems of Systems
- Example: Smart cities as systems of systems

"Arguably for the first time in history, artificially intelligent urban technologies are taking the management of urban services out of the hands of humans, operating the city in an autonomous way."
Cugurullo (2020)

"It is a complex system designed by humans to manage and sustain human lives, and it is therefore imperative that human stakeholders do not cede their decision-making capabilities to nonhuman intelligences" – Cugurullo and Xu (2025)

F. Cugurullo, Urban Artificial Intelligence, 2020

Cugurullo and Xu, When AIs become oracles, 2025

How might insights from cognitive science help?

Humans and Machines as Thought Partners

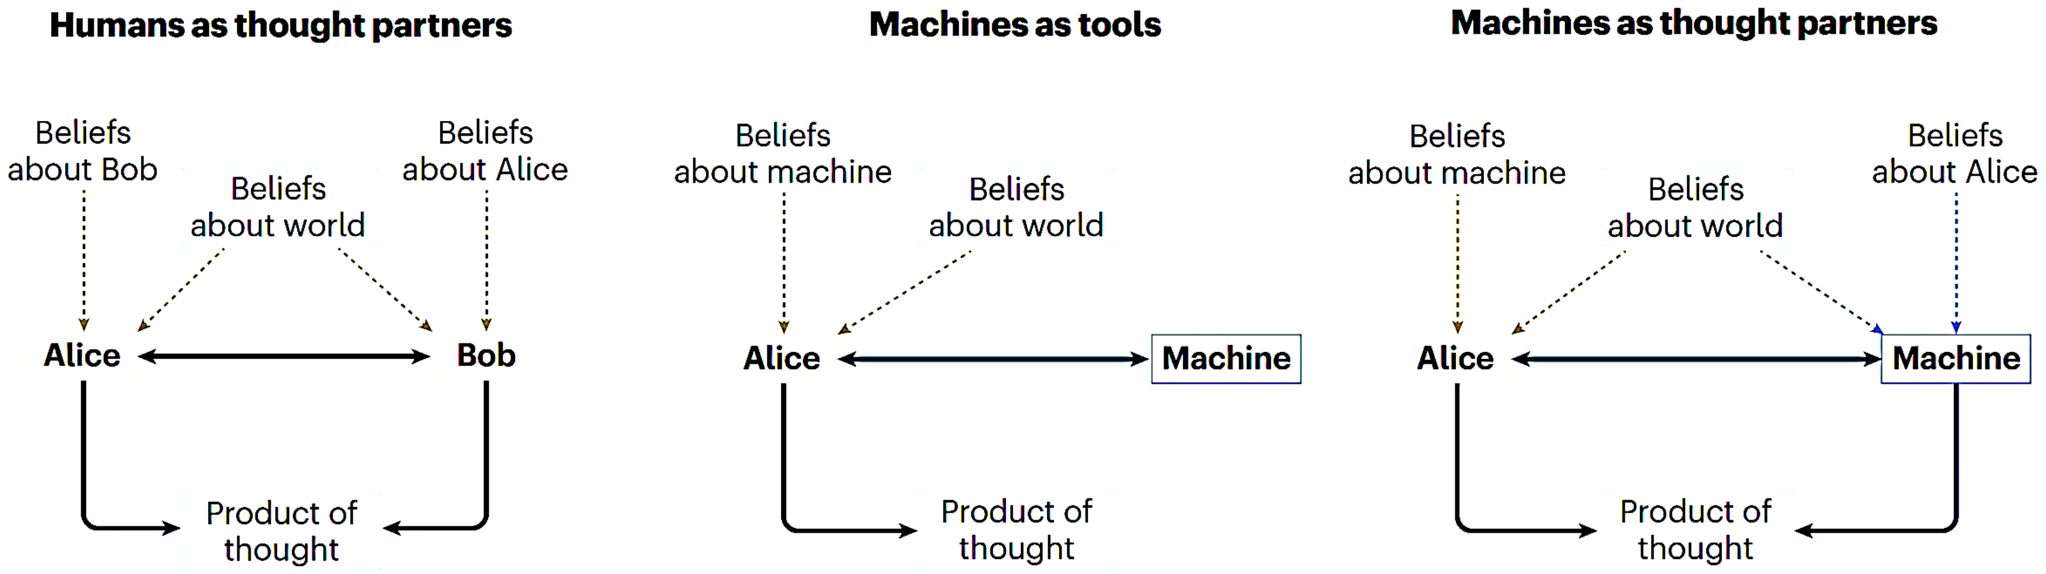

Fig. 1 | Examples of ecosystems for thinking. Humans have long thought together. Machines have expanded the efficiency of human thinking. Now, machines—powered by AI—open up new realms of computational thought partnership with humans.

Source: [Building machines that learn and think with people, Collins et al, Nature Human Behavior](), 2024

Mutual Modeling for Collaboration

- Human should have a mental model of the intelligent machine
- Intelligent machine should have a computational model of the human
- These models should help predict actions of the partner

Source: Building machines that learn and think with people, Collins et al, Nature Human Behavior, 2024

Exploring Modes of Collaborative Thinking

- **Collaborative Planning & Strategy:** Value & Intent Alignment, Scalable Multi-Agent Planning
- **Collaborative Problem Solving & Learning:** Robust Problem-Solving Abilities, Knowledge Gaps
- **Collaborative Deliberation & Reasoning:** Verifiable Reasoning, Opinion Diversity
- **Collaborative Sense-Making, Creation and Ideation:** Fidelity to the Real World, Generation Diversity, Style Consistency

Source: Building machines that learn and think with people, Collins et al, Nature Human Behavior, 2024

Tools for Building Models: Computational Cognitive Science Motifs

- **Probabilistic models of cognition:** Bayesian inference through structured probabilistic generative world models
- **Theory of mind and communication:** Rational speech act framework
- **Resource-rationality and tractable theory-building:** Approximate inference or bounded planning
- **Probabilistic programming:** mathematics for hybrids of optimization, dynamic programming, and Monte Carlo inference

Source: Building machines that learn and think with people, Collins et al, Nature Human Behavior, 2024

Control to Enable Physical AI: Major Challenges

- Lots of progress in multiple directions of ongoing work, e.g., **data-driven control, learning for decision and control,** etc.

- Frameworks and techniques for combining first principles, model-based control and AL/ML data-driven approaches are inadequate

- Assessing and ensuring safety and reliability of AI/ML integrated CPS remains a major unsolved problem

- Architectures (hierarchical?) and associated algorithms for integrating AI/ML into CPS are not well developed.

- Fundamental, principles-based understanding of AI/ML so that they can be reliably and efficiently engineered into current and future CPS needs much advancement

Physical AI Safety

- Safety is a critical requirement in all intelligent machines
- Safety is conceptually and technically connected to fundamentals ideas in control such as stability, constraints, robustness, and performance
- Safety in self driving cars is the most developed
- Safety in other physical AI domains will also be important: manufacturing, energy, transportation, biomedical, agriculture, …
- Human-Intelligent Machine collaboration will be founded on **safety as an underlying bedrock**

Safety in Self Driving Cars: Absence of Unreasonable Risk

Three Layers

- **Architectural:**
  - Example: undesired presence of blind spots, due to sensors' typology and placement;
- **Behavioral:**
  - Example: undesired degree of proximity to surrounding road users;
- **In-service operational:**
  - Example: improper securing of cargo or undesired access to the vehicle from a malicious actor.

Safety as Dynamic Process

- **Emergent property:**
  - safety-by-design practices
- **Acceptable prediction/observation:**
  - Actual and simulated performance
- **Continuous confidence growth:**
  - Statistical confidence in the AUR determination

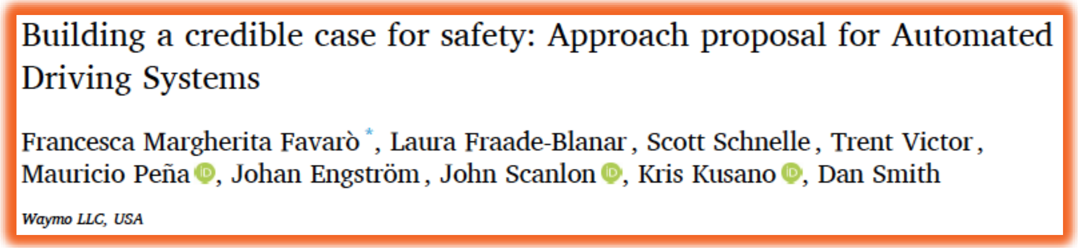

Building a credible case for safety: Approach proposal for Automated Driving Systems

Francesca Margherita Favarò *, Laura Fraade-Blanar, Scott Schnelle, Trent Victor, Mauricio Peña, Johan Engström, John Scanlon, Kris Kusano, Dan Smith

Waymo LLC, USA

Source: Favaro et al, [Journal of Safety Research 96 (2026) 181–196]

Waymo Safety Record

Compared to an average human driver over the same distance in our operating cities, the Waymo Driver had

Overall crash reduction

92%
Fewer serious injury or worse crashes
35 FEWER

83%
Fewer airbag deployment in any vehicle crashes
230 FEWER

82%
Fewer injury-causing crashes
544 FEWER

Crash reductions involving injuries to Vulnerable Road Users

92%
Fewer pedestrian crashes with injuries
62 FEWER

85%
Fewer cyclist crashes with Injuries
39 FEWER

81%
Fewer motorcycle crashes with injuries
25 FEWER

https://waymo.com/safety/impact/

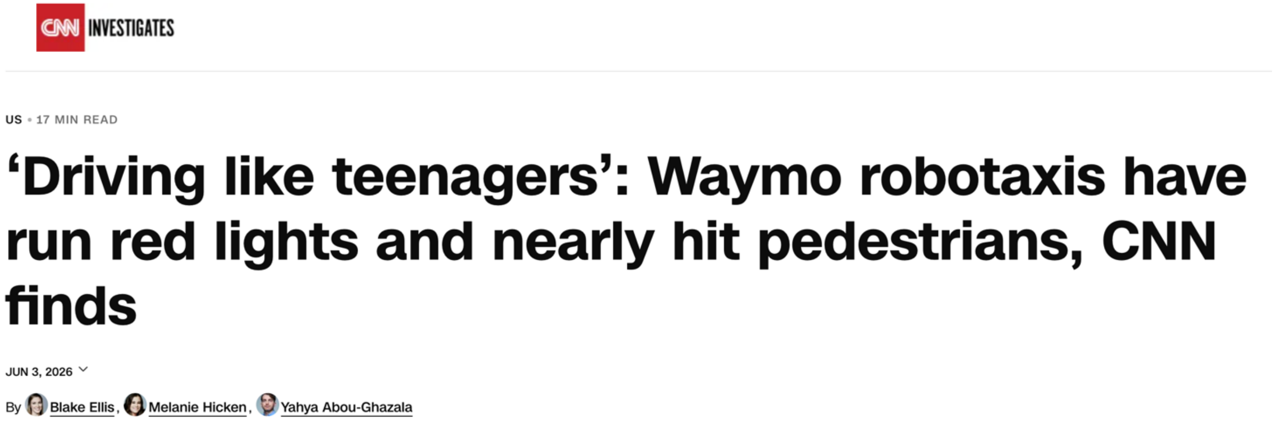

US • 17 MIN READ

‘Driving like teenagers’: Waymo robotaxis have run red lights and nearly hit pedestrians, CNN finds

JUN 3, 2026

By Blake Ellis, Melanie Hicken, Yahya Abou-Ghazala

Waymo pauses robotaxis in five US cities after cars drive into flooded roads

Physical AI Safety: Salient Challenges for a New Field

- **Uncertainty Quantification:** ability for foundation (or related) models to reliably quantify uncertainty
- **Scenario-Based Development:** Redesigning learning and control algorithms and components to maintain verified levels of accuracy throughout their entire operational domain
- **Formal Verification:** Improving the scaling of methods to verify the desired properties of large neural networks or learning control algorithms
- **Safe Failing**: For various realizations of intelligent machines, developing protocols to "fail safely" to avoid harming humans or damaging the system itself.
- **Continual Learning:** Implementing "outer loops" where novelties are detected, data is collected, and intelligent machines are retrained to handle new environmental factors.
- **System-of-Systems and Safety:** connecting safety at the system level to the safety at the system-of-systems level

NASEM, Machine Learning for Safety-Critical Applications: Opportunities, Challenges, and a Research Agenda, 2025

Three Paths to Building Future Intelligent Machines

- Scaling and expanding the currently dominant approach: foundation models leveraging multi-modal datasets
- Learning and intelligence through interactions with physical and social reality - biomimetic
- Learning with humans in the loop

“A compelling hypothesis is that if **embodied AI systems are trained—at least in part—in social settings through ongoing interaction with people**, they may develop aligned representations from the outset: shared, negotiated and embedded in human norms.”

Source: Pezzulo et al, Shared Worlds, Shared Minds, 2025

Issues and Questions

- How will safety, security, reliability, and other critical requirements continue to be met as intelligent machines learn, adapt, and evolve?
- Will intelligence in “intelligent machines” be centralized or distributed/decentralized or a mix?
- How will “intelligence” at the three CPS levels: devices, systems, and systems of systems be coordinated?
- How will recent AI developments affect human-AI relationships?
- Which of the three paths to building future AI systems will be taken? Is there path dependence in this evolution?

Parting Thought

“The question is no longer whose models hit technical benchmarks, but who can build and sustain an ecosystem that embeds AI into everyday products and services,” writes Angela Huyue Zhang, law professor at the University of Southern California, in a recent article for Project Syndicate. In other words, **the race is as much about embodying AI into physical environments through sensing, control and decision-making** as generating text and images. This includes intelligent manufacturing, humanoid robots and applications in other devices, such as cars, phones and wearables.”

China will Clinch the AI Race, Tej Parikh,
Financial Times, January 18, 2026

Concluding Comments

- AI is being integrated into CPHS at all levels: devices, systems, and systems-of-systems resulting in a wide variety of “intelligent machines”
- Design and evolution of these “intelligent machines” will shape and be shaped by human-machine collaboration
- Tools, techniques, and insights from cognitive science can help in creating a more effective and fruitful human-machine collaboration
- How will we rise to meet Wiener’s challenge: will we ask machines to do what they ought to do?
- **Fertile ground for the controls field to contribute to a positive future of intelligent machines serving human society**

Comments

Ideas

Questions?

pramod.Khargonekar@uci.edu
http://faculty.sites.uci.edu/khargonekar/