

Leveraging Machine Learning for Advancing Smart-X Systems and Control

GE Edge and Controls Symposium

Pramod P. Khargonekar

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

10 September 2019

Why Machine Learning, Artificial Intelligence and Data Science in
Control?

Control Systems: Diverse Application Domains

- ▶ Aerospace
- ▶ Energy and power
- ▶ Manufacturing
- ▶ Chemical processes
- ▶ Automotive
- ▶ Transportation
- ▶ Water, food, and agriculture
- ▶ Biomedical

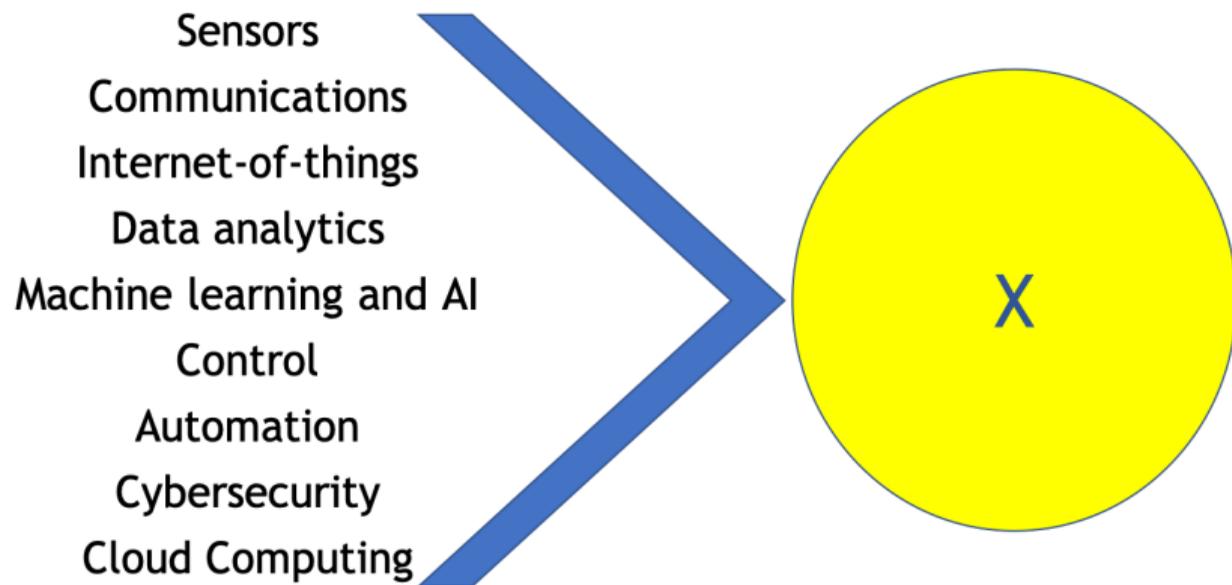
Control Systems : Strong Theoretical Foundations

- ▶ Stability theory
- ▶ Optimal control
- ▶ Linear multivariable control
- ▶ Robust control
- ▶ Nonlinear control
- ▶ Adaptive Control
- ▶ Stochastic control
- ▶ Distributed control

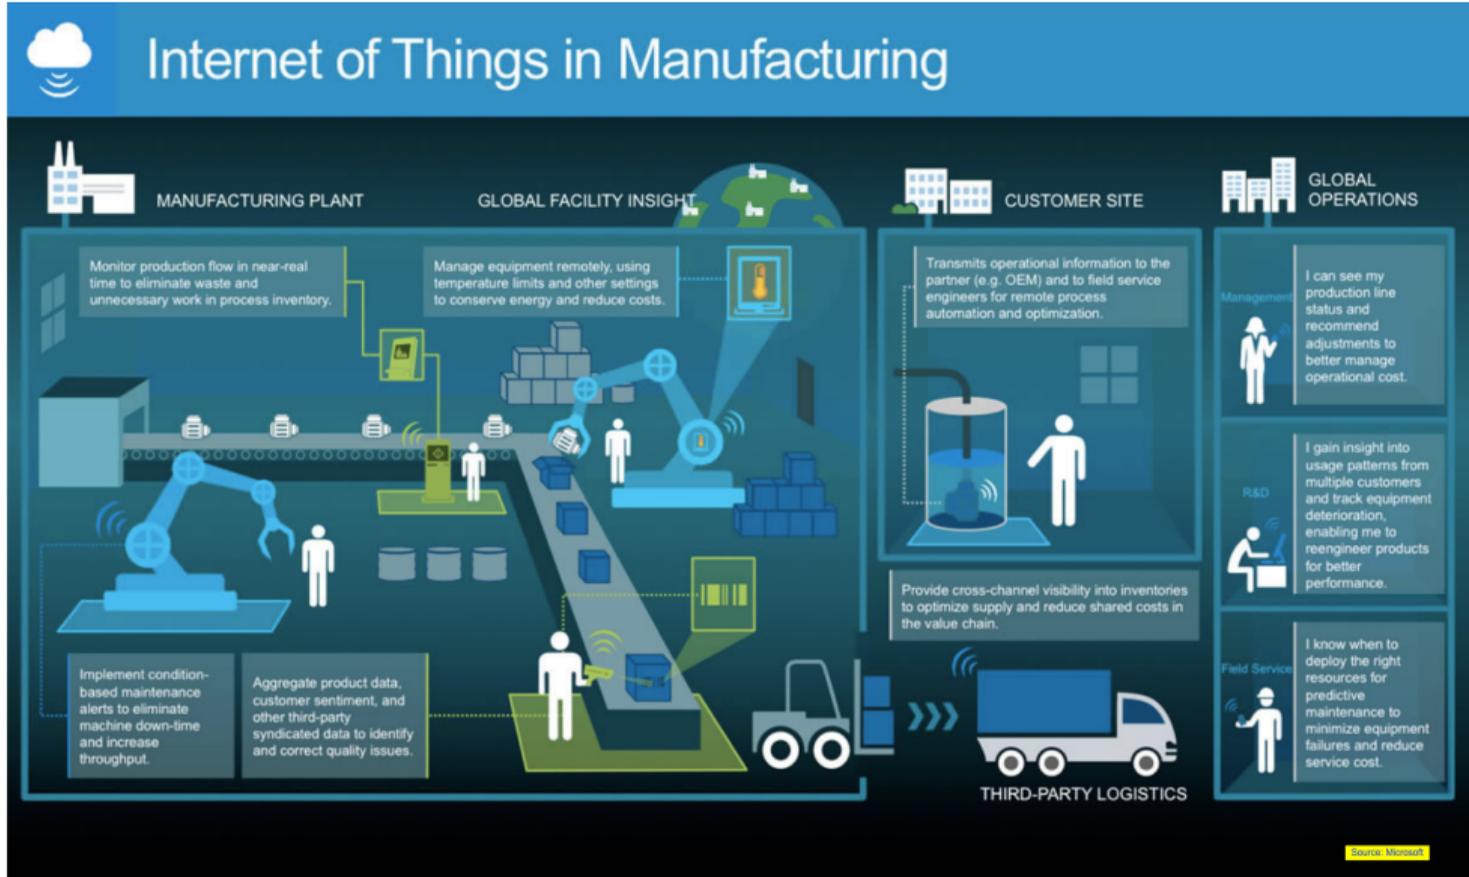
Aspirational and Emerging Applications: Examples

- ▶ Smart-X
 1. Smart manufacturing
 2. Smart electric grid
 3. Smart homes
 4. Smart cities
 5. Smart transportation
 6. Smart agriculture
 7. Smart health
- ▶ Autonomous systems
 1. Unmanned air vehicles
 2. Self-driving cars
 3. Autonomous robots

Smart-X: Conceptual View



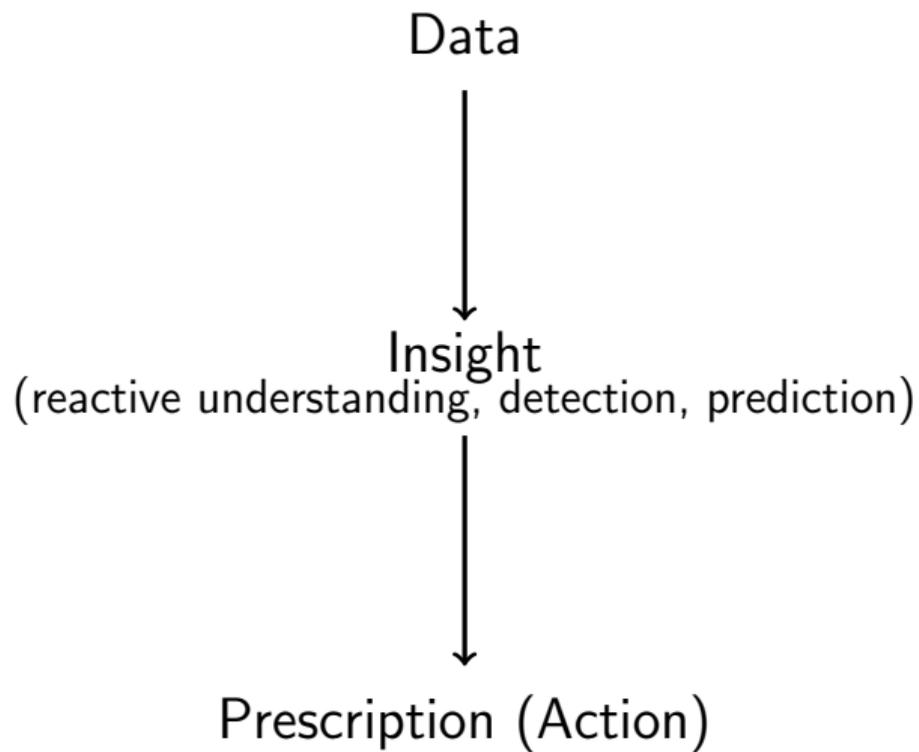
Example: Smart Manufacturing



Data: Distributed, Big and Streaming

- ▶ **Internet-of-Things (IoT)** into industrial and commercial settings: manufacturing, aerospace, chemical, electric grid, transportation, . . .
- ▶ **Cheap and ubiquitous sensors:** cameras, microphones, GPS, touch, health and fitness, . . .
- ▶ **User generated data:** social media, citizen science, . . .
- ▶ **Enterprise data:** manufacturing, healthcare, pharmaceutical, transportation, retail, energy and power, . . .
- ▶ **Scientific data:** genomics, proteomics, brain imaging, telescopes, weather, satellites, . . .
- ▶ **Government data**

Data to Action is a Form of Control



Challenge: Design, operation, management and control of large, distributed, heterogeneous, complicated, interconnected techno-socio-economic systems.

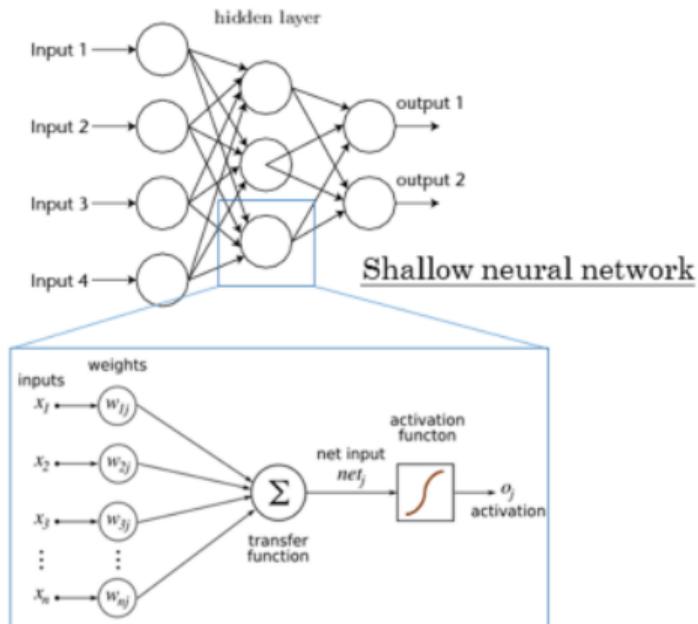
Vision: Control systems will play an important role but will need to integrate with cyber-physical-human systems, data science, machine learning, and artificial intelligence.

Recap of Recent Machine Learning Breakthroughs

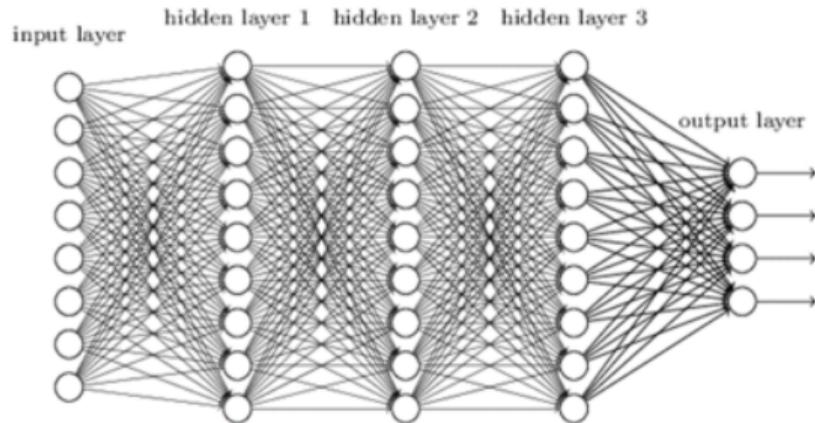
Computational Intelligence: Pattern Recognition or Model Building

- ▶ Two fundamentally different perspectives on learning from data:
 1. Statistical pattern recognition from data for prediction and control.
 2. Using data to build causal models to understand, predict and control.
- ▶ Possible to combine these two approaches.
- ▶ Causality a critical issue.

Deep vs Shallow Neural Networks



Deep neural network



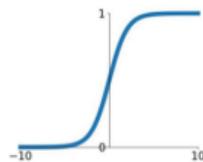
Building Block: A Single Artificial Neuron Unit

- ▶ Inputs: x_1, x_2, \dots, x_n
- ▶ Weights: w_1, w_2, \dots, w_n
- ▶ An activation function σ
- ▶ Examples of activation functions:
- ▶ Output given by

$$a = \sum_{j=1}^n w_j x_j$$
$$y = \sigma(a)$$

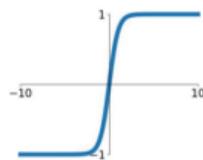
Sigmoid

$$\sigma(x) = \frac{1}{1+e^{-x}}$$



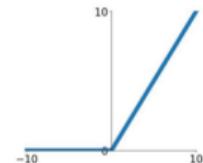
tanh

$$\tanh(x)$$



ReLU

$$\max(0, x)$$



Key Advantage of Deep Networks

*“ ... shallow classifiers require a good feature extractor ... one that produces representations that are selective to the aspects of the image that are important for discrimination ... The conventional option is to hand design good feature extractors, which requires a considerable amount of engineering skill and domain expertise. But this can all be avoided if **good features can be learned automatically** ... **This is the key advantage of deep learning.**”*

Deep Learning, LeCun, Bengio, and Hinton, Nature, 2015.

Major DL Innovations

- ▶ Convolutional neural networks
- ▶ Training and optimization of extremely large networks
- ▶ Long Short Term Memory (LSTM) for sequential data
- ▶ Use of graphics processors for computation

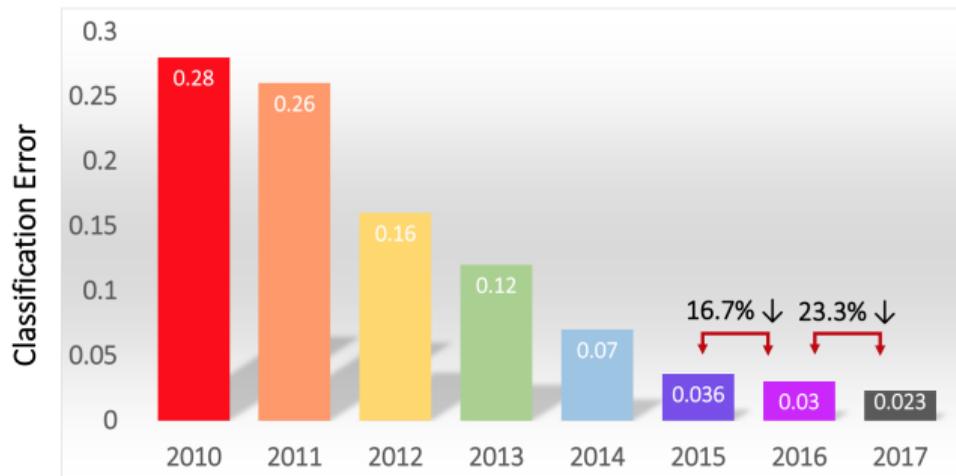
Major Applications

- ▶ Image recognition
- ▶ Object detection
- ▶ Segmentation
- ▶ Speech Processing
- ▶ Machine Translation

Breakthrough in Vision: ImageNet Competition

ImageNet Classification with Deep Convolutional Neural Networks, Krizhevsky, Sutskever, and Hinton, 2012

Classification Results (CLS)



Recurrent Neural Networks

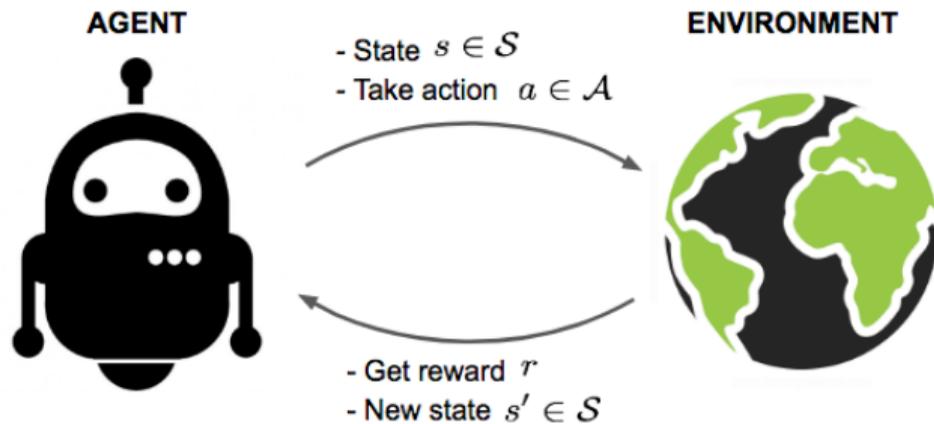
- ▶ Recurrent neural networks (RNNs): neural network models with the ability to pass information across time steps
- ▶ Suitable for modeling data that are
 - ▶ Sequential and dependent.
 - ▶ Of varying input lengths.
- ▶ RNNs: natural choice for time series and other sequential applications.
- ▶ Long Short Term Memory (LSTM) Networks: the state-of-the-art RNNs.

Google Neural Machine Translator: Results on Production Data

Table 10: Mean of side-by-side scores on production data

	PBMT	GNMT	Human	Relative Improvement
English → Spanish	4.885	5.428	5.504	87%
English → French	4.932	5.295	5.496	64%
English → Chinese	4.035	4.594	4.987	58%
Spanish → English	4.872	5.187	5.372	63%
French → English	5.046	5.343	5.404	83%
Chinese → English	3.694	4.263	4.636	60%

RL Framework



The “agent” is the controller and the “environment” includes the plant, uncertainty, disturbances, noise, etc.

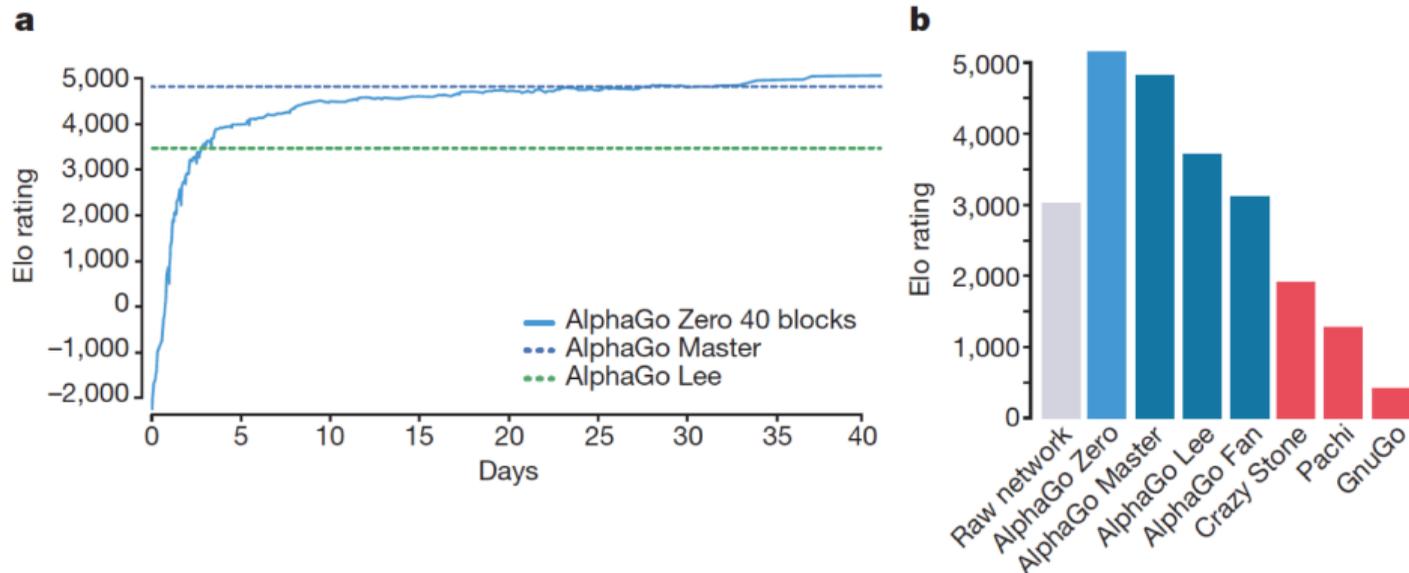
Reinforcement Learning: General Setup

- ▶ At each time step, agent observes the state, takes action, and receives a reward.
- ▶ Goal for the agent: choose actions to maximize total discounted reward.
- ▶ Optimal action policy is a form of control law.
- ▶ Can the agent learn the optimal policy by suitable use of state and reward data?
- ▶ RL: A general machine learning paradigm to solve problems and attain goals.

Key Ideas and Building Blocks

- ▶ Bellman's optimality principle: *Tail of an optimal policy must be optimal.*
- ▶ Function $Q(x, a)$: optimal policy given by maximizing with respect to a .
- ▶ One approach: Learn the Q -function.
- ▶ Recent innovations in modern RL
 1. Deep Reinforcement Learning: Use deep neural networks to approximate Q (DQN)
 2. Experience replay to reuse past data
 3. Asynchronous and parallel RL
 4. Rollouts based planning for RL
 5. Self-play for faster learning
 6. Techniques for data efficiency
 7. Techniques for continuous action spaces

AlphaGo Zero achieves State-of-the-Art Performance



Despite learning by itself from zero prior knowledge, it learns and outperforms all other algorithms.

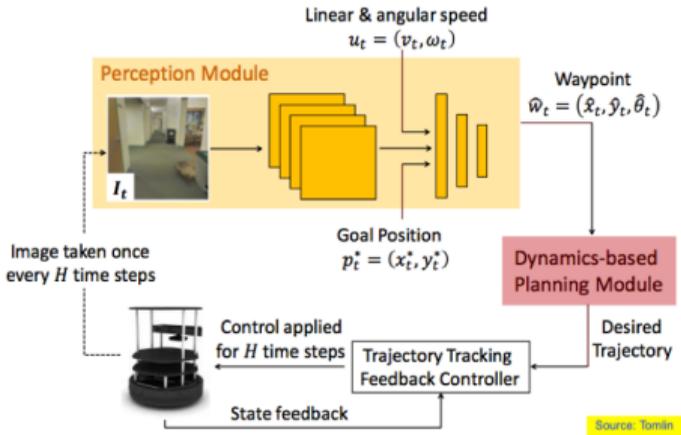
Critical Recap of ML Breakthroughs

- ▶ DL establishing itself as a major new technology.
- ▶ Insufficient theory of DL but progress on both approximation and generalization.
- ▶ Major investments in DL hardware that will make it cheaper to implement.
- ▶ Deep reinforcement learning — breakthrough performance in board games.
- ▶ Applications of DRL to physical systems at very early stages.
- ▶ DL and RL depend on large amounts of data.
- ▶ DL and (much of) RL are model-free.
- ▶ Numerous novel and promising research directions in DL, RL, and ML.
- ▶ Enormous global interest in private, academic, government sectors.

Leveraging ML Advances - Perception

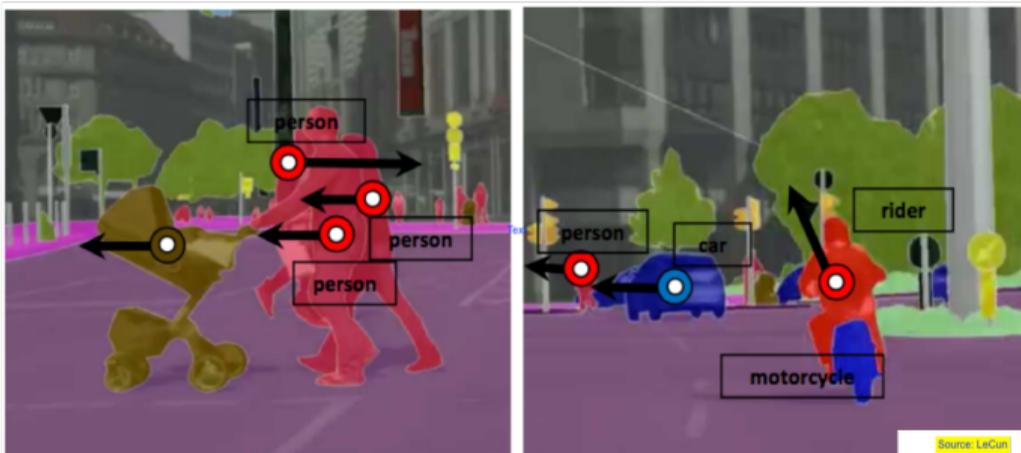
- ▶ New dimensionality reduction techniques to deal with high-dimensional data.
- ▶ DL to deal with image/video data for integration into Smart-X systems.
- ▶ New paths to integrate vision sensors in control systems.
 - ▶ Example: vision sensors in manufacturing, transportation, . . .
 - ▶ Example: image analysis for situational change and awareness
- ▶ DL to integrate multiple sensor modalities for failure detection, predictive analytics, recommendations, control and decision making.

Examples



Perception module generates waypoint

Planning module generates trajectory



ML algorithm predicts the motion of people and objects for situational awareness

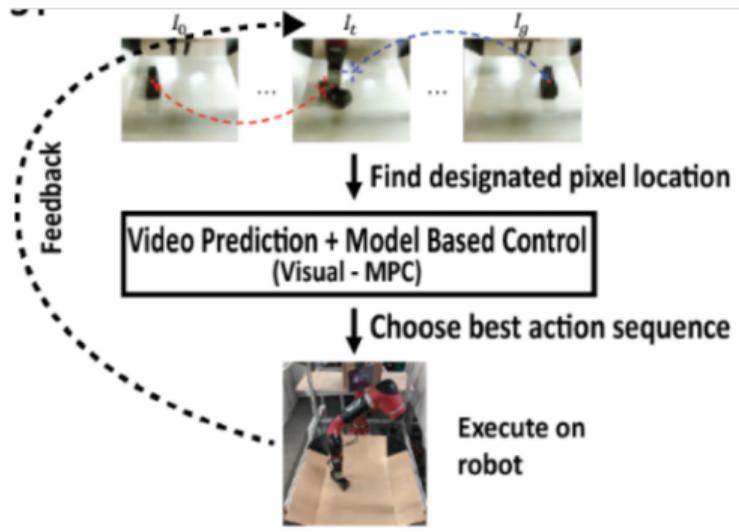
Leveraging ML Advances - Streaming Data

- ▶ Recurrent neural networks for drawing inference from streaming data.
- ▶ Exploit state-of-the-art LSTM structures and algorithms for extracting useful information from streaming data.
 - ▶ Example: LSTM designed for analysis of IoT streaming data for equipment state.
 - ▶ Example: LSTMs as soft-sensors for process control applications
- ▶ New paths to integrate audio sensors using LSTM advances in speech processing.

Leveraging ML Advances - Complementing Model Based Approaches

- ▶ Traditional engineered systems and control rely on model-based approaches.
- ▶ Leverage DL (and RL) based technologies for the “hard to model” parts of the system.
- ▶ Develop new techniques for integrating model-based and model-free technologies.
 - ▶ Example: Combine model-predictive control with RL based approaches.
 - ▶ Example: ML based planning with model-based feedback control.

Example: ML Based Planning + MPC



Video prediction model generates multiple planning trajectories

Model predictive control using the plan which is then repeated.

Leveraging ML Advances - Closing the Loop

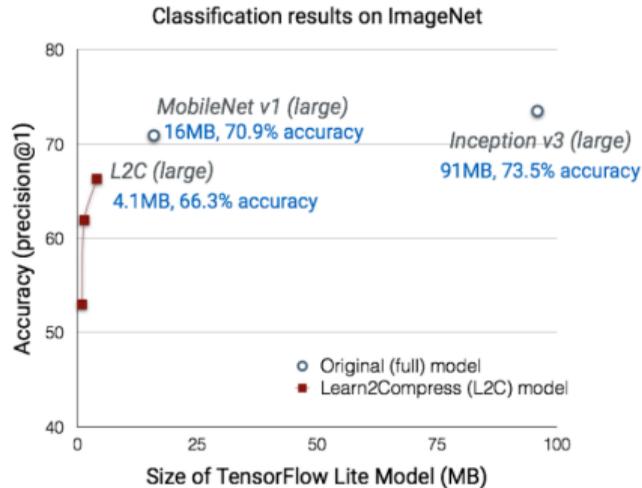
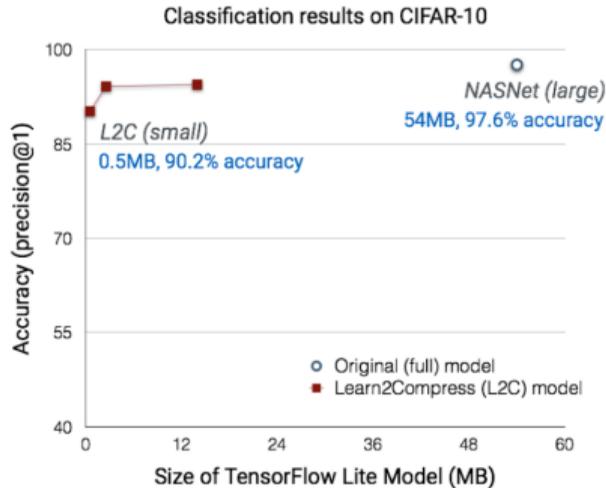
- ▶ Stability, safety, physical constraints as major issues when closing the loop.
- ▶ New techniques such as Safe-RL hold much promise for control.
- ▶ Traditional model-based control for lower-level, fast loops and RL types of techniques for higher-level, slower control decisions.

Leveraging ML Advances - Other Ideas

- ▶ Leverage innovations in training, optimization, data reuse, etc. from ML into systems and control.
- ▶ Active sensing using concepts of attention and perceptual loop.
- ▶ Exploit inherent parallelism for rapid spread of learning and adaptation.
 - ▶ Example: learning and adaptation in multiple copies of the same system in manufacturing.
- ▶ “Learning to learn” as a paradigm for controlling “new systems” from the control of known systems.
- ▶ Speculative: “Cognitive CPS” for improved CPS-Human collaboration.

Edge Intelligence: ML meets Edge - Current Paradigm

- ▶ Current paradigm: DL training in the data center. DL inference at the edge.
- ▶ Technical tools: distillation, compression, transfer learning
- ▶ Specialized hardware for energy and computational efficiency for edge



Edge Intelligence: ML meets Edge - Future

- ▶ Future: Device-edge-cloud hybrid and coordinated approaches
- ▶ Key challenge: trade-offs in energy, bandwidth, latency, privacy, and optimization
- ▶ Emerging approaches:
 1. Federated learning: local models at the edge, data center aggregates local models
 2. Gradient compression to reduce communication overhead
 3. DNN splitting for preserving user privacy

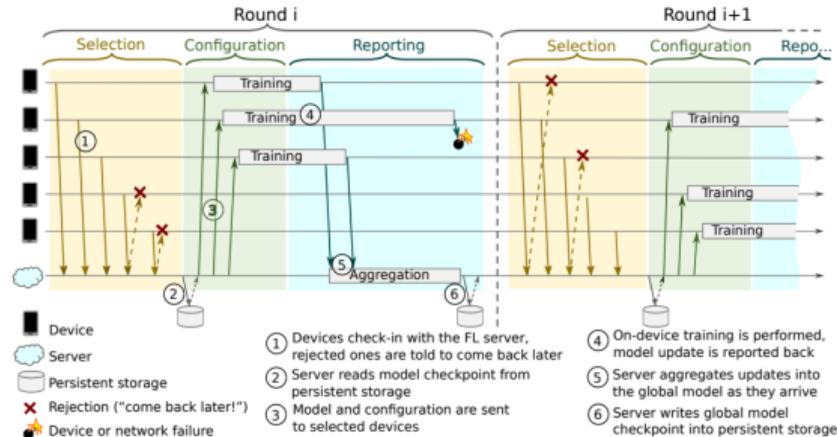


Figure 1: Federated Learning Protocol

Source: Google

Machine Learning and Control: Our Recent Work

- ▶ Long-term learning of adaptive controllers using sparse neural networks.
- ▶ Robust deep recurrent neural network controllers.
- ▶ Sparse, recurrent neural network adaptive controllers.
- ▶ External working memory to enhance neural adaptive controllers.
- ▶ Publications and presentations available on my website.

Concluding Remarks: ML Advancing Smart-X and Controls

- ▶ Machine learning expected to be the next major general purpose technology.
- ▶ Algorithmic, architectural, and hardware advances from ML into Smart-X and control systems.
- ▶ Necessary to deal with safety and other physical constraints.
- ▶ Smart-X systems to enhance human flourishing.

We are in the early stages of this exciting journey.

Thanks to D. Muthirayan for his help in preparation of this presentation.

Thank you!

Email: pramod.khargonekar@uci.edu

Website: <https://faculty.sites.uci.edu/khargonekar/>