ECE209AS (Winter 2021)

Lecture 1: Course Introduction

Mani Srivastava
mbs@ucla.edu
Networked & Embedded Systems Lab
ECE & CS Departments
UCLA

About Me

- Name: Mani Srivastava
  - Professor in CS and ECE Departments
  - Office: 6730-E Boelter Hall / Lab: 1762 Boelter Hall
  - Office Hours: Tu 9-10PM, Th 12-1PM LA Time (tentative)
    - https://ucla.zoom.us/j/92303243429 (password as announced)
  - Email: mbs@ucla.edu
    - please use private note on Piazza instead of email
    - if you do email please put “ECE 209AS” as subject
  - More information:
    - https://www.ee.ucla.edu/mani-srivastava/ [Me]
    - https://scholar.google.com/citations?user=X2Qs7XYAAAAJ [My research]
    - http://www.nesl.ucla.edu [My lab]

About the course

- New special topics course
- On AI/ML methods for data-driven CPS and IoT systems
- Zoom for lectures and office hours
  - https://ucla.zoom.us/j/92303243429 (password announced in email to class)
- Piazza for course material and discussion
  - https://piazza.com/ucla/winter2021/ece209as (access code announced in email to class)
- Assignment submission via Gradescope
  - Please log in to Gradescope via the institutional login and you will see ECE209AS-1
- CCLE not used for anything

CPS and IoT Systems Fueled by Data + AI/ML

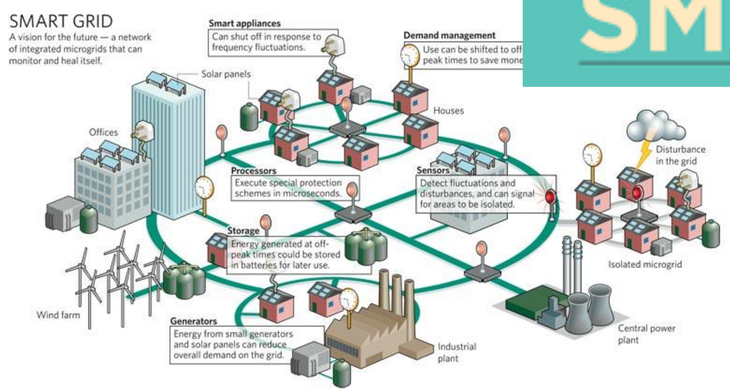

Guide decisions, Control things, Nudge behaviors….

CPS and IoT Systems Fueled by Data + AI/ML

mHealth: New Biomarkers and Just-in-Time Adaptive Interventions

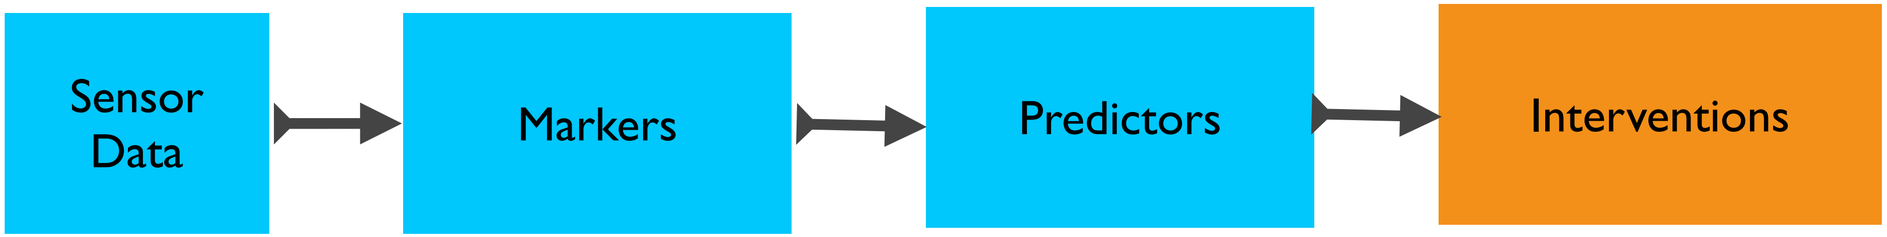

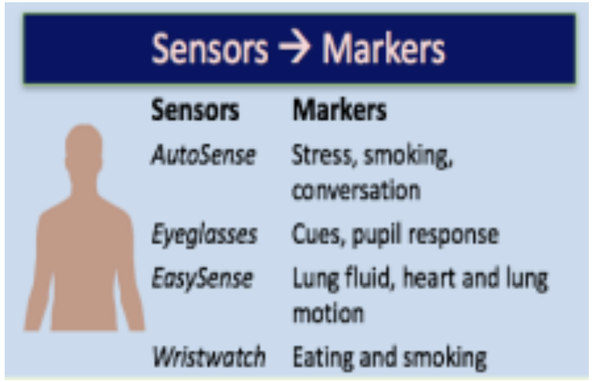

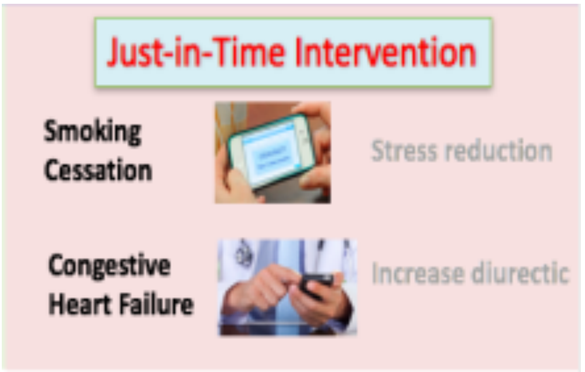

CPS and IoT Systems Fueled by Data + AI/ML

Example: BCG for Monitoring of Heart Failure Patients in Daily Life
(Continual Assessment of Blood Volume, i.e. Hemodynamic Congestion)

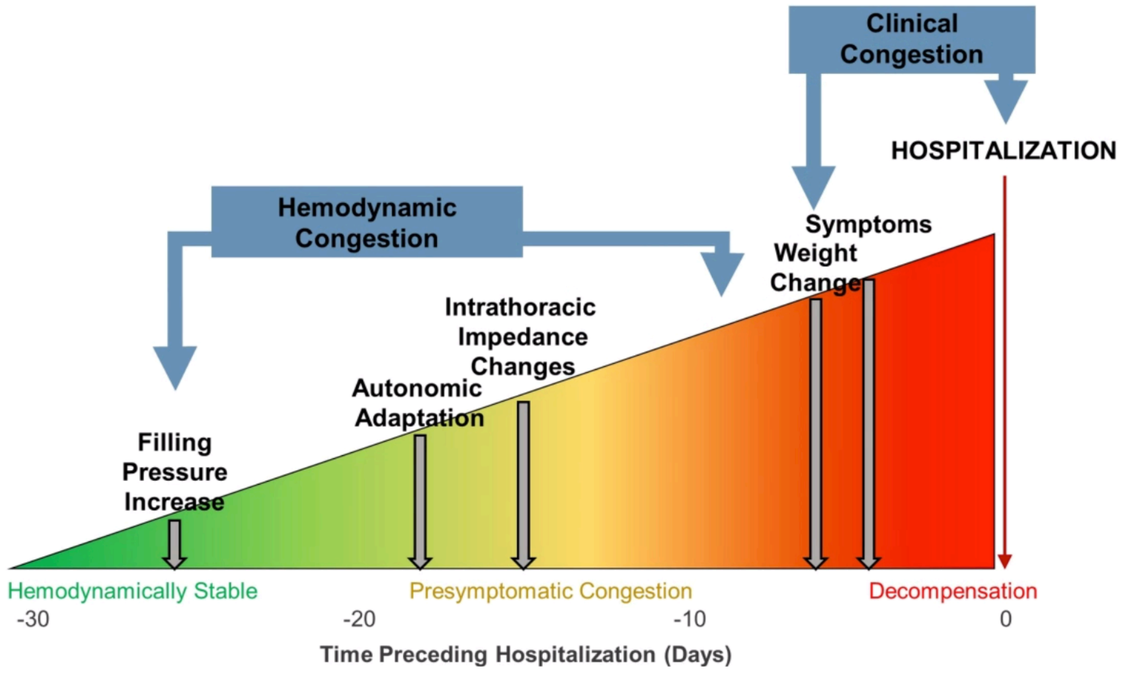

Progression from Chronic Compensated to Acute Decompensated Heart Failure

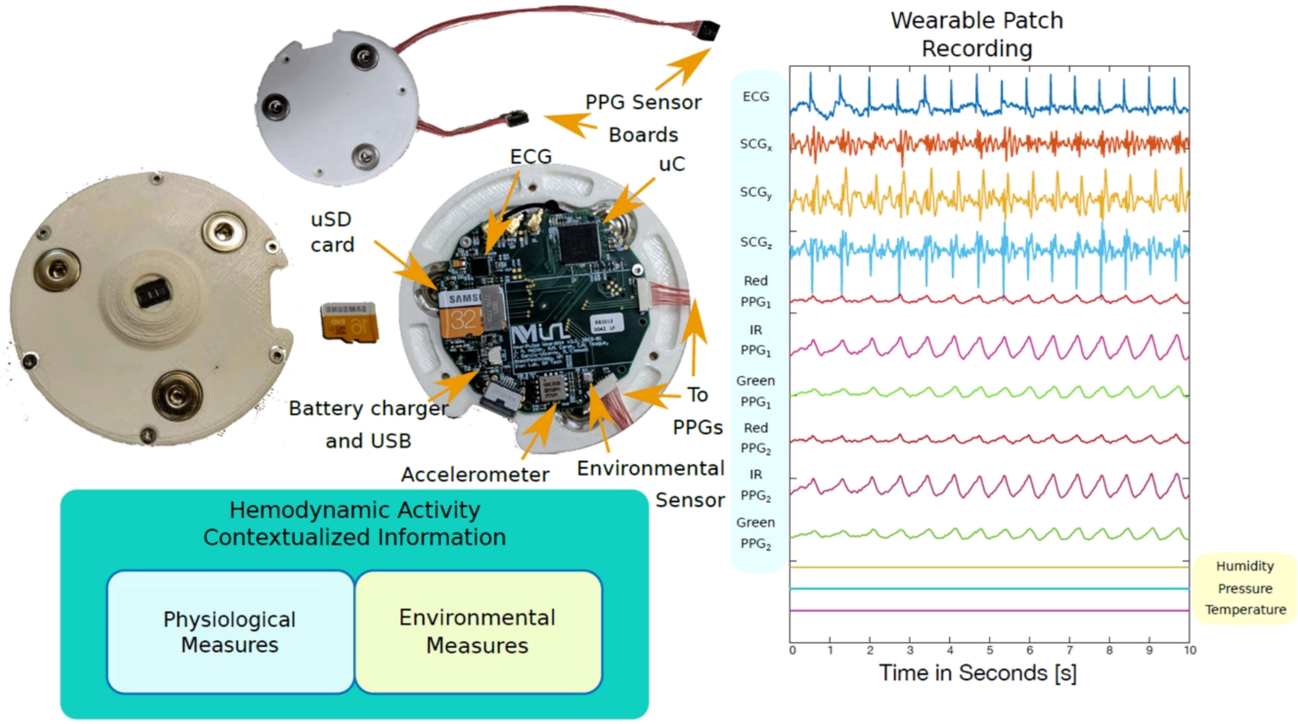

Wearable Patch for Multimodal Cardiovascular Sensing

Reference: Research at Prof. Inan's Laboratory @ GaTech

CPS and IoT Systems Fueled by Data + AI/ML

Example: BCG for Monitoring of Heart Failure Patients in Daily Life
(Continual Assessment of Blood Volume, i.e. Hemodynamic Congestion)

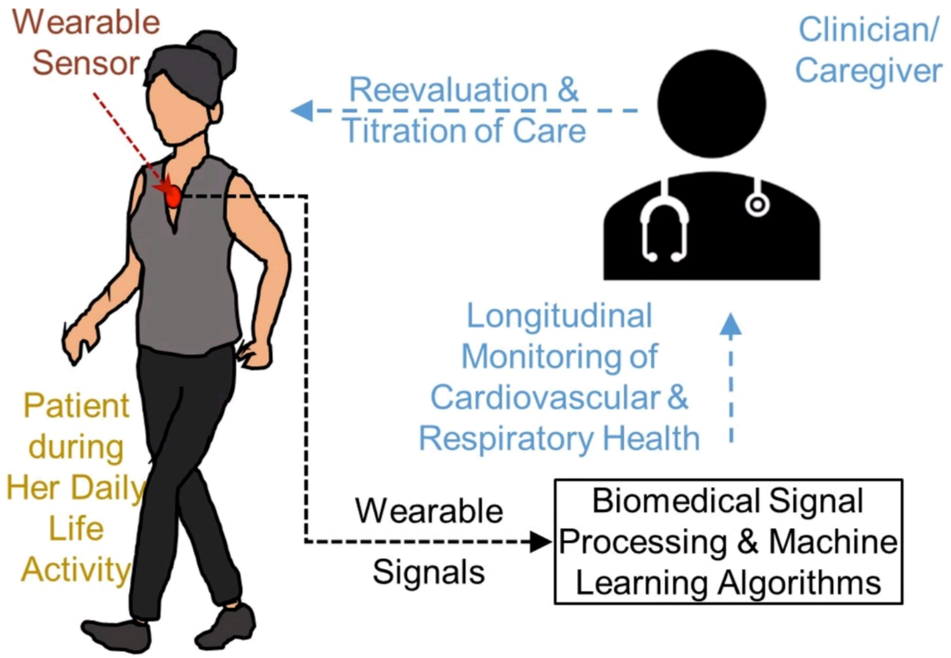

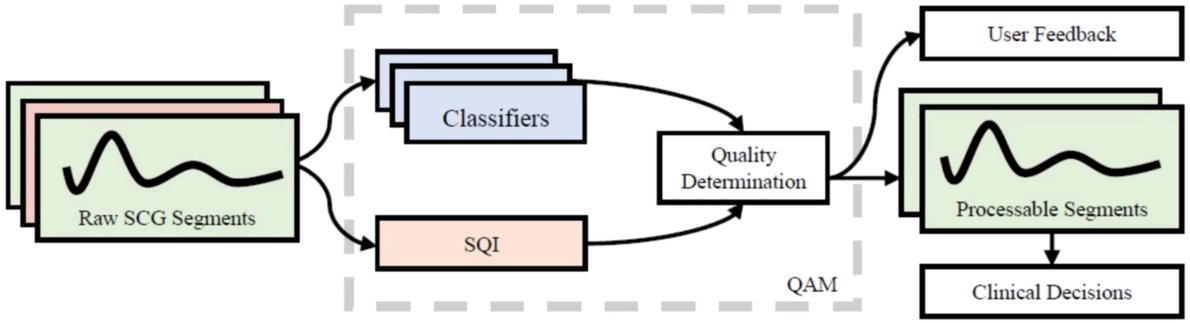

Overall Framework

Reference: Research at Prof. Inan's Laboratory @ GaTech

CPS and IoT Systems Fueled by Data + AI/ML

Smart Built Environments: Sustainability, Occupant Comfort & Safety

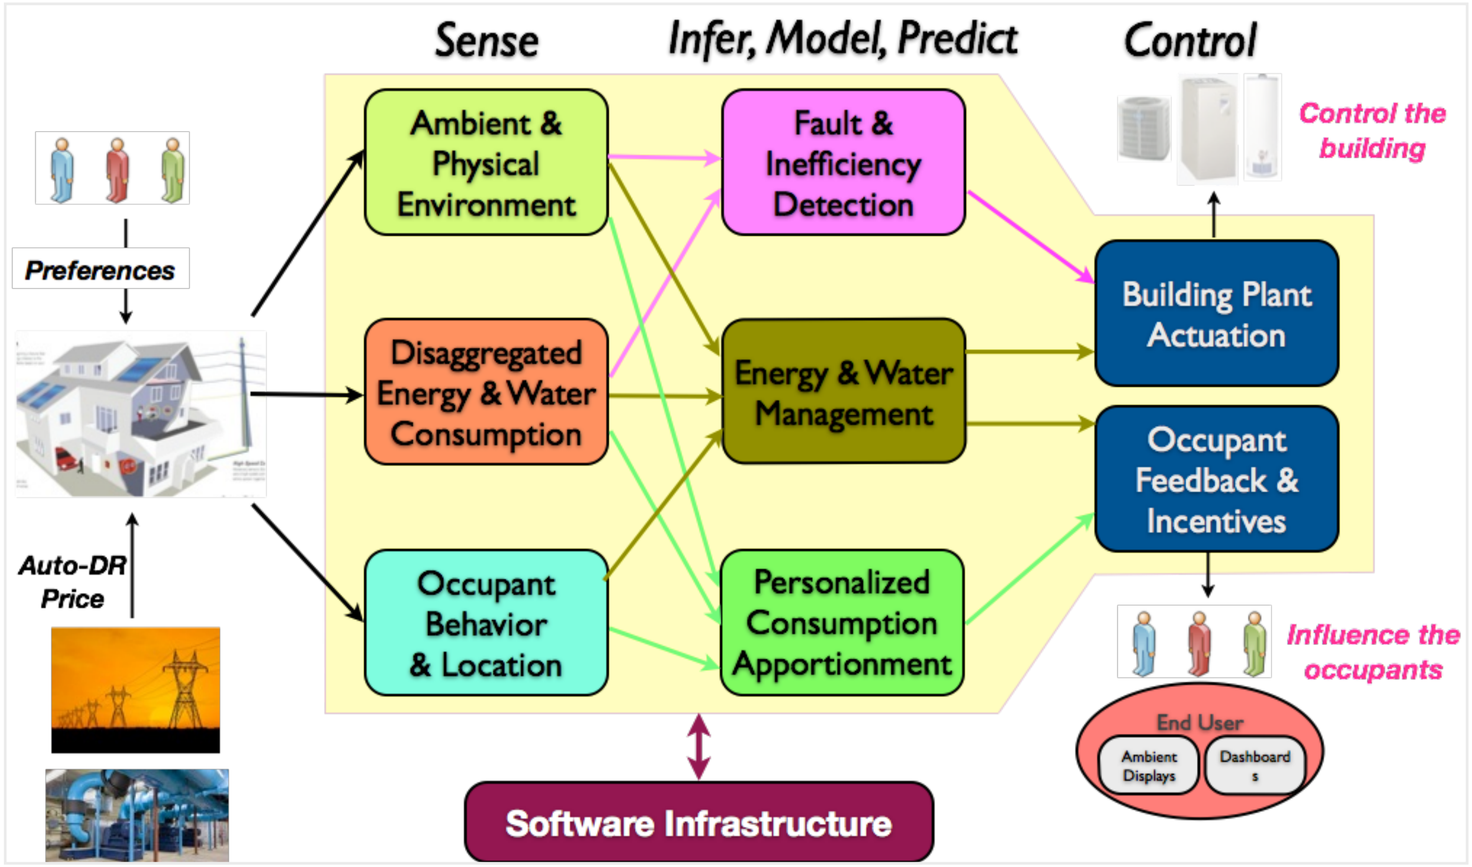

CPS and IoT Systems Fueled by Data + AI/ML

Data-driven & Learning-enabled City Services and Planning

Sensor Data
Social Media

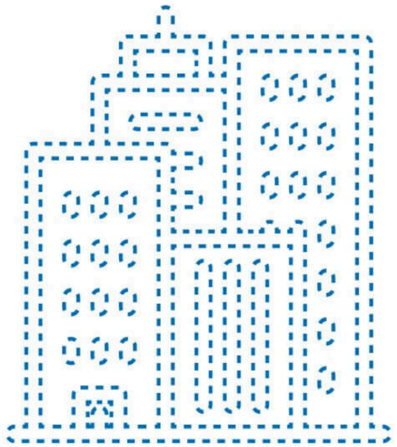

Digital Twin

Cross-reality Model
Predictive Analytics
Decision Making

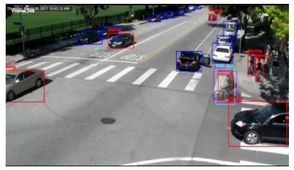

CPS and IoT Systems Fueled by Data + AI/ML

Urban Robotics

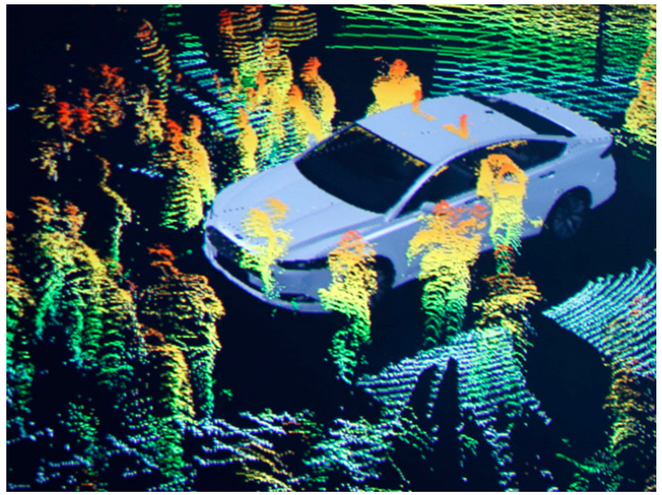

GNSS/GPS

Ultrasound

Camera

mmWave Radar

LIDAR

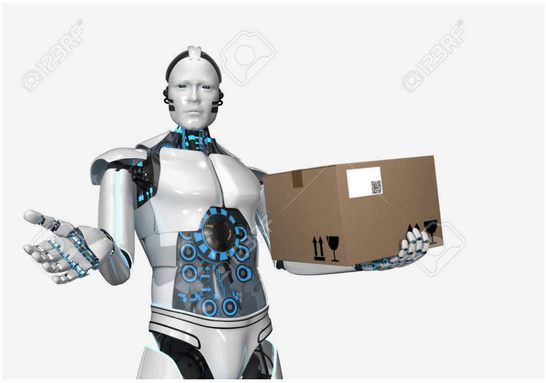

CPS and IoT Systems Fueled by Data + AI/ML

Augmented Vehicular Reality for Enhanced Situational Awareness

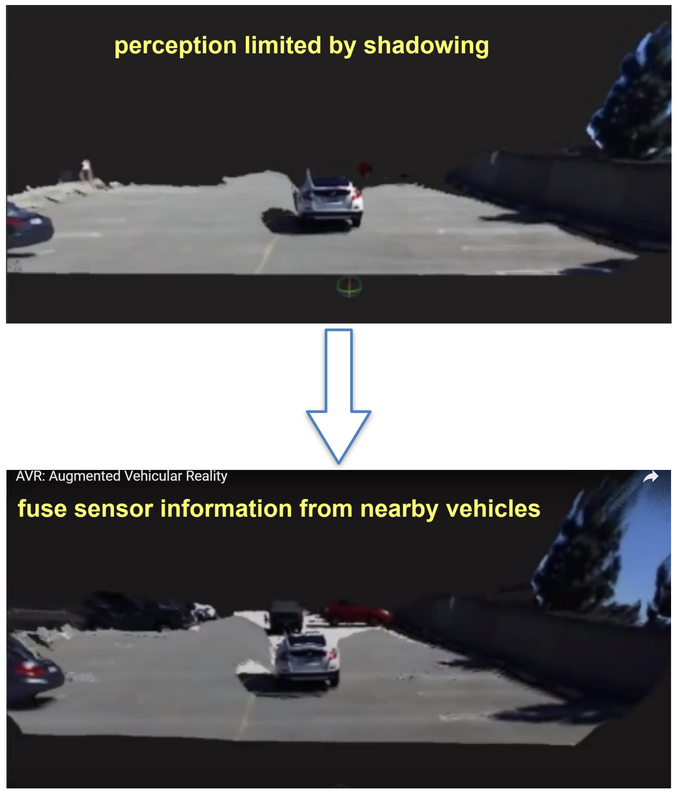

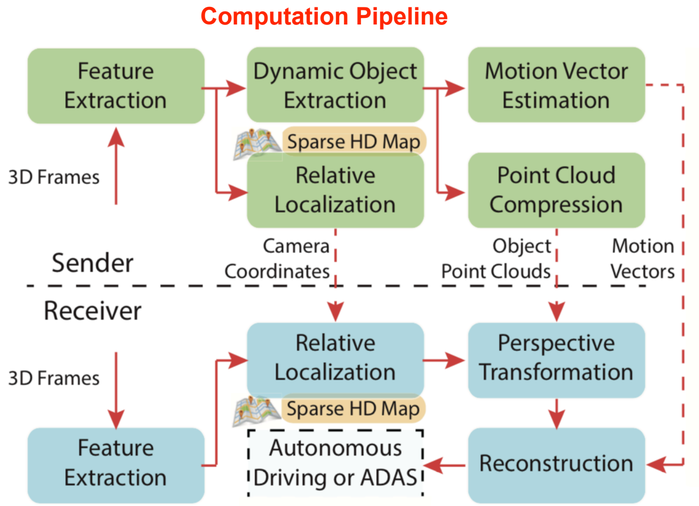

CPS and IoT Systems Fueled by Data + AI/ML

National Defense and Security

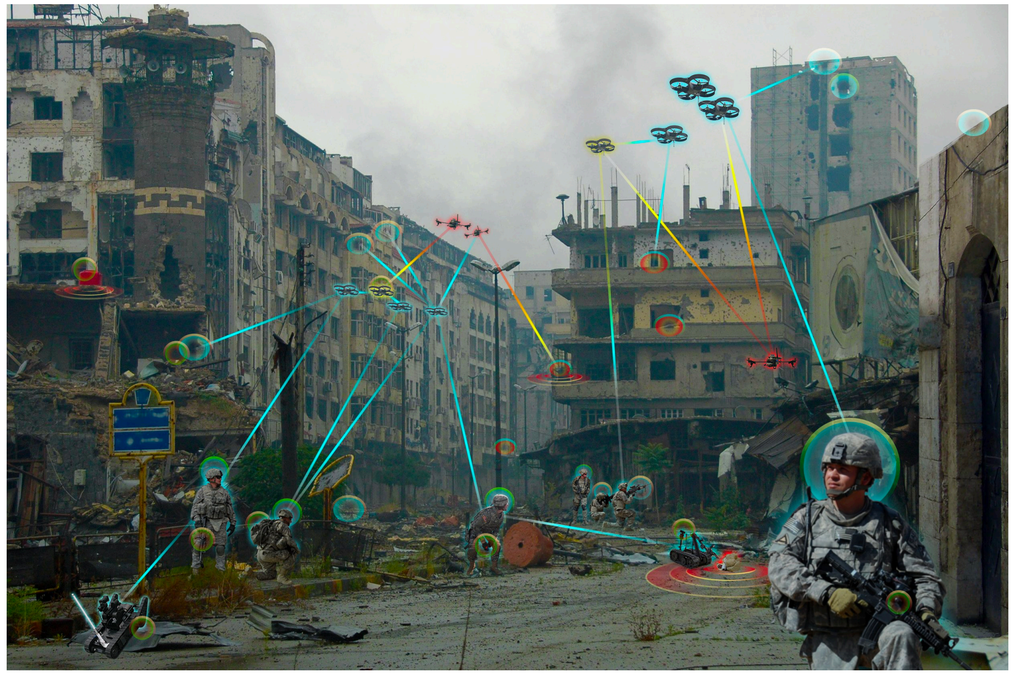

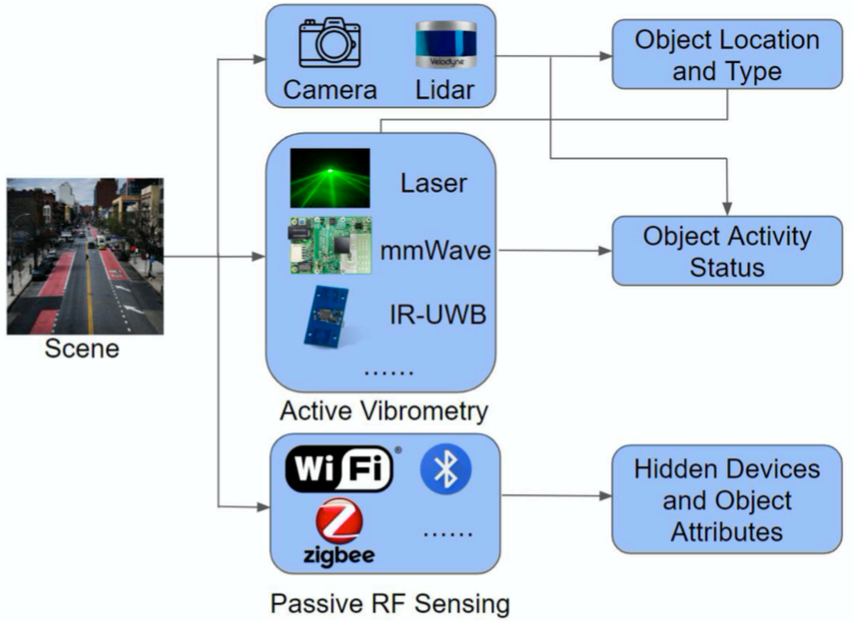

Intelligent Agent Interacting with Dynamic Environment

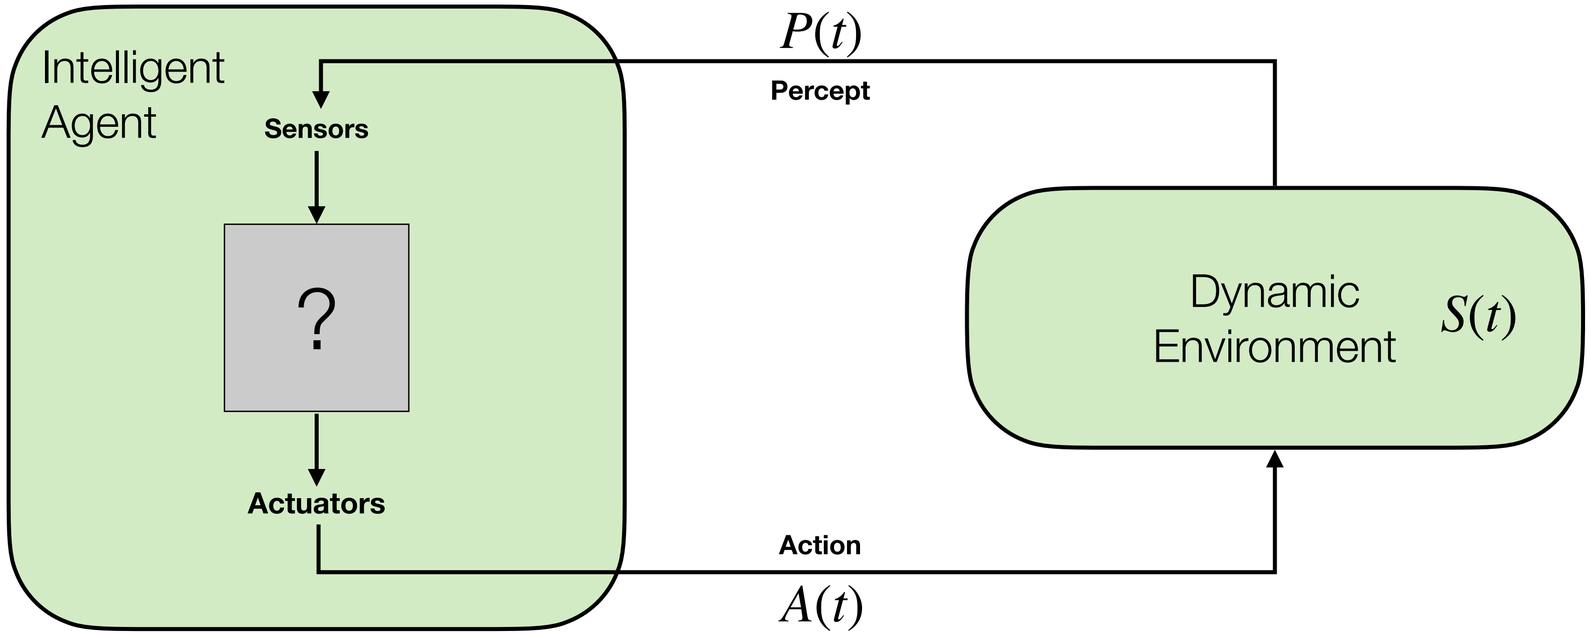

- Agents include humans, IoT devices, robots, cyber-physical systems, …
- The agent has a function $f : P^*(t) \rightarrow A(t)$ that maps from percept histories to actions to ensure that the trajectory of environment state $S(t)$ meet some desired properties or specifications
  - e.g. keep a human in a healthy state, a building not consume excessive energy, etc.

Russell, Stuart J., and Peter Norvig. "Artificial Intelligence: A Modern Approach"

Perception-Cognition-Action Loop

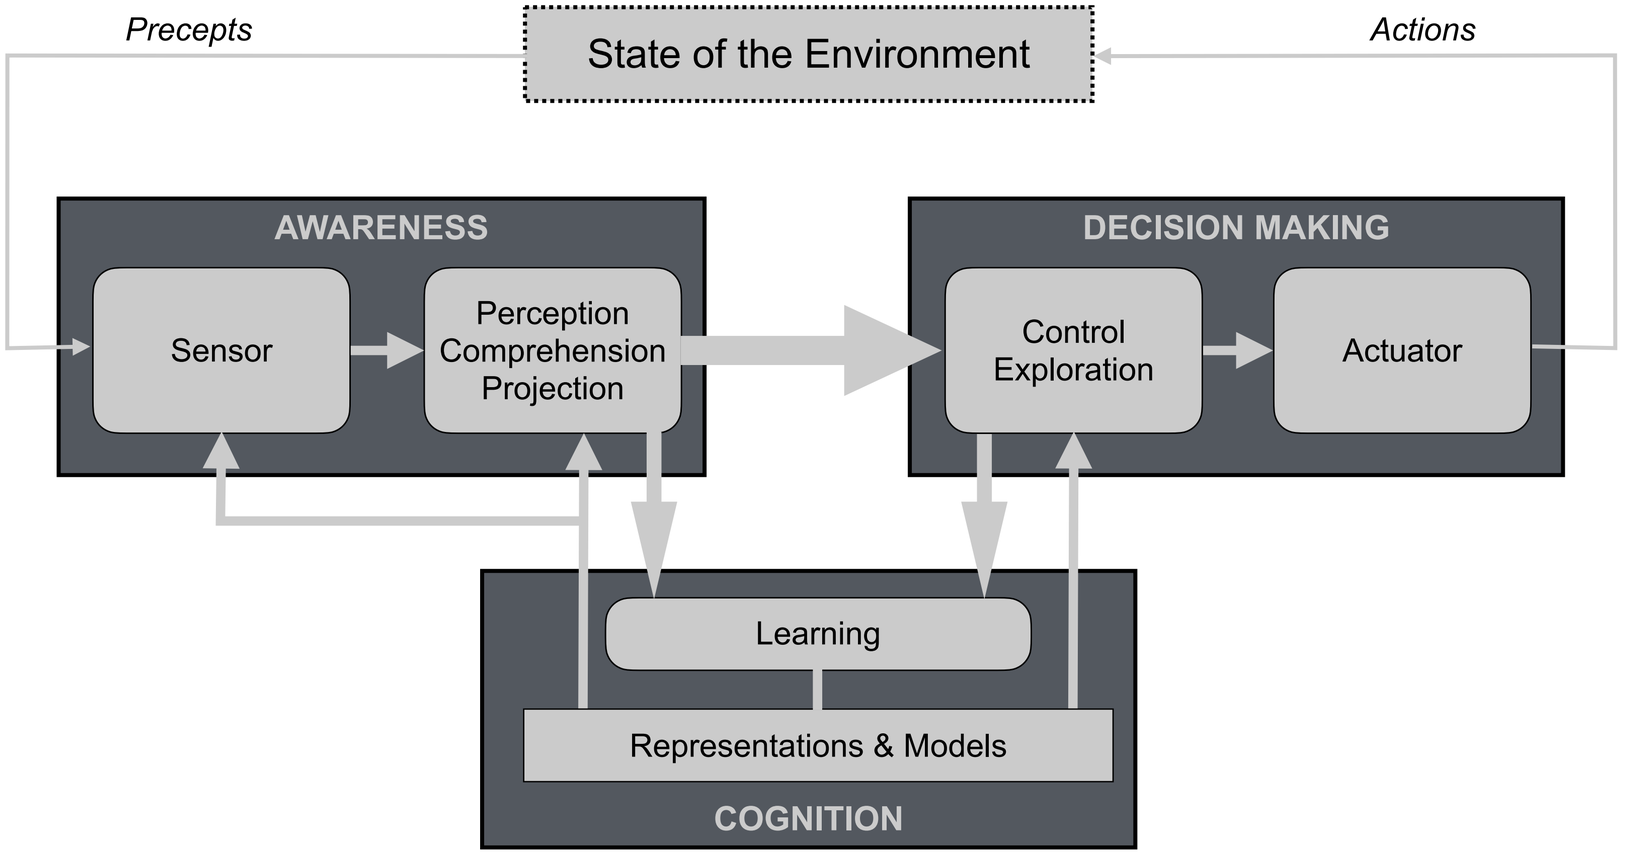

Perception-Cognition-Action Loop

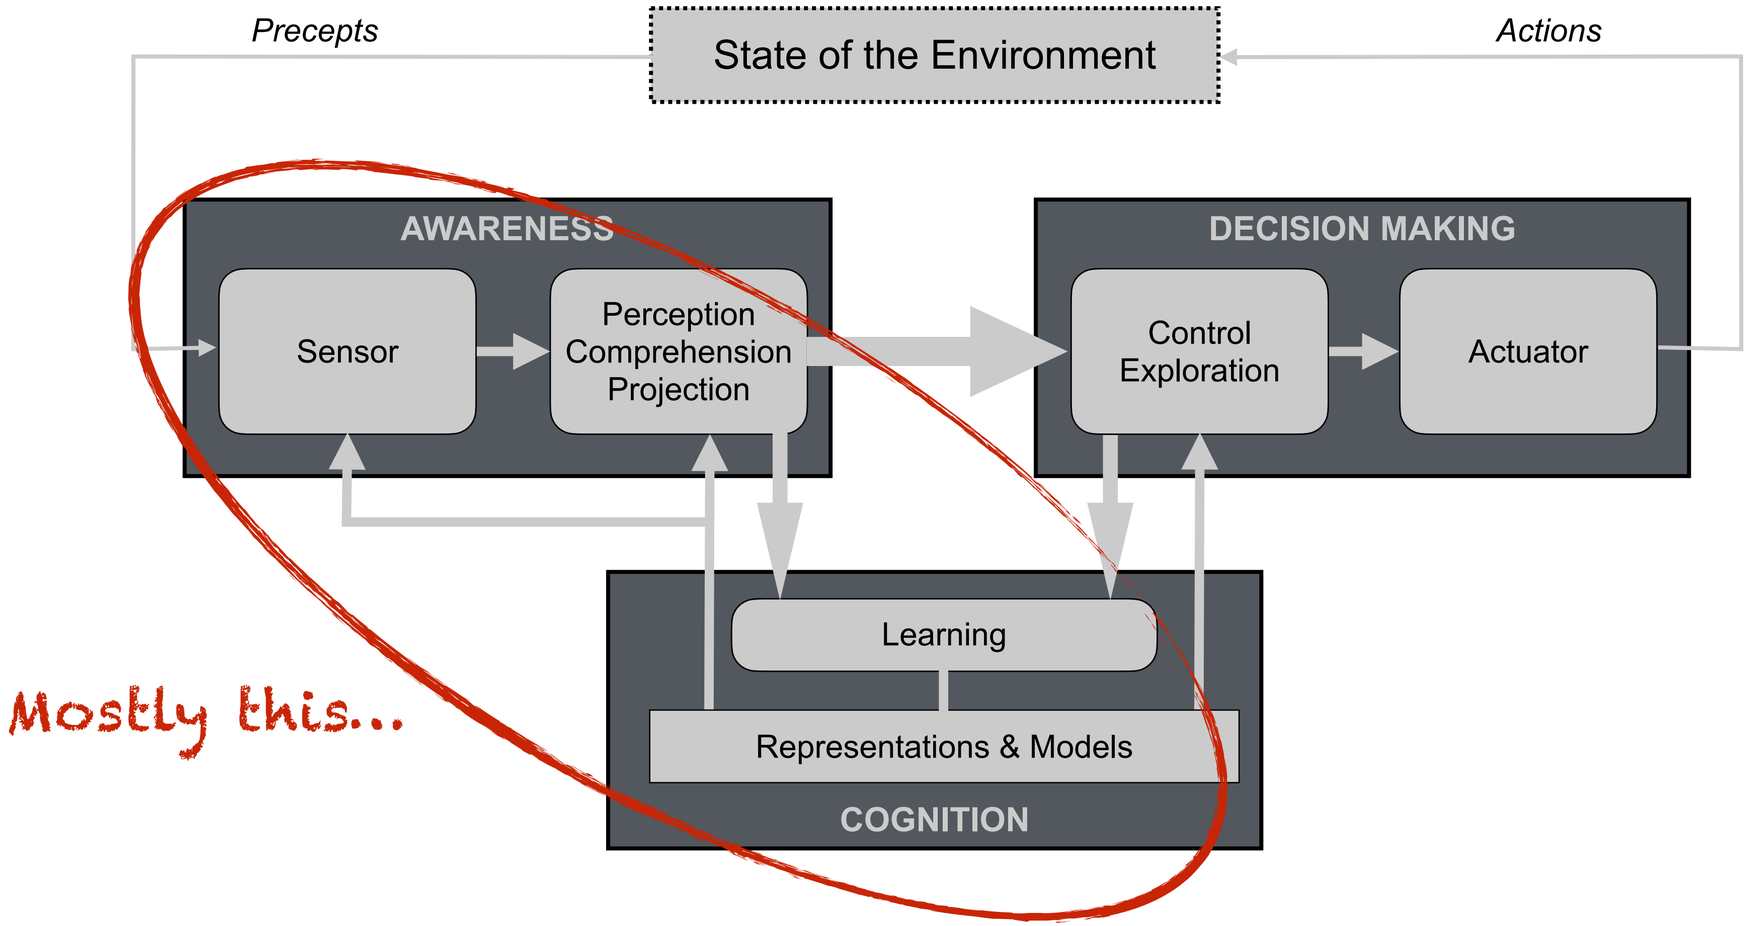

Perception-Cognition-Action Loop

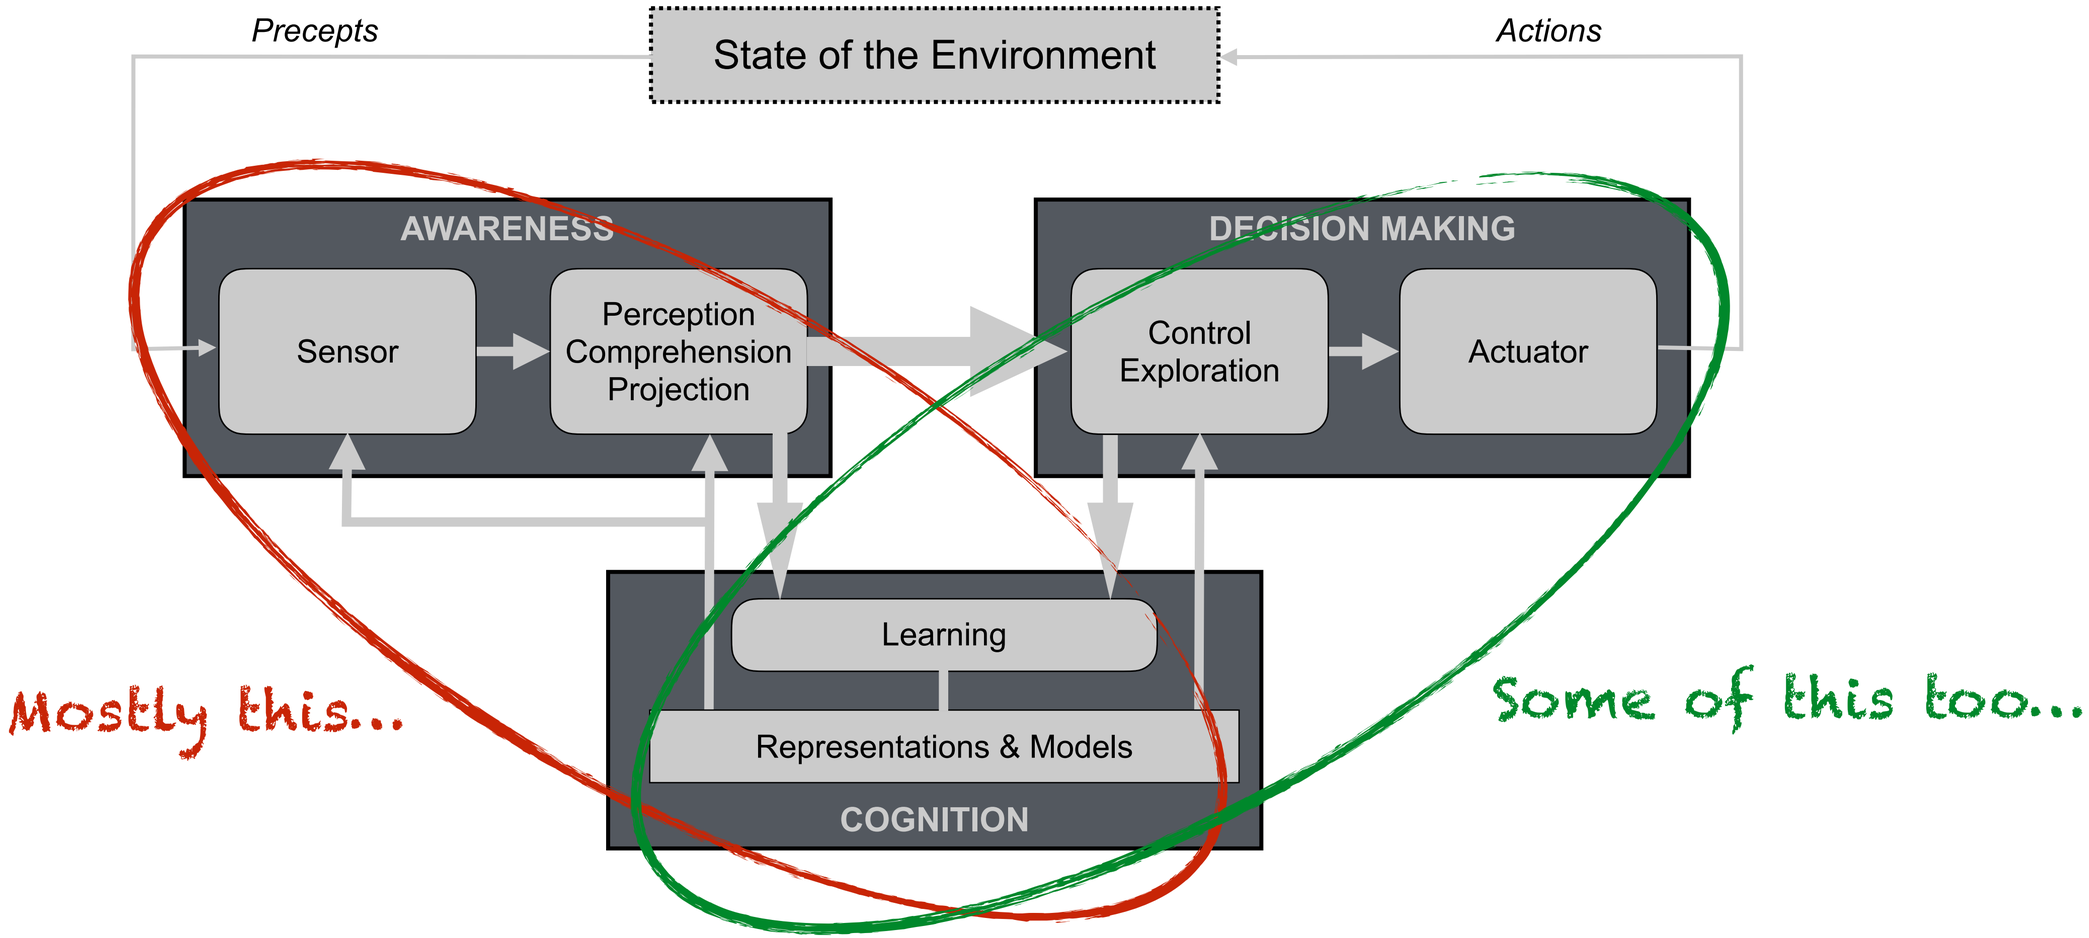

Driven by Technology Trends

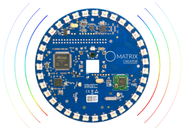

Acoustic Array

Camera

mmWave Radar

LIDAR

UAVs

Rich Sensing & Actuation

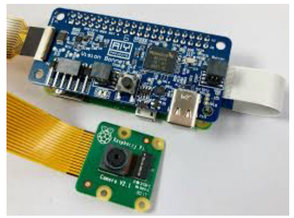

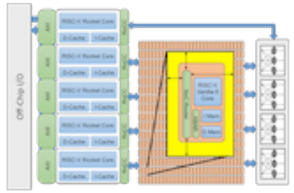

Compute Accelerators

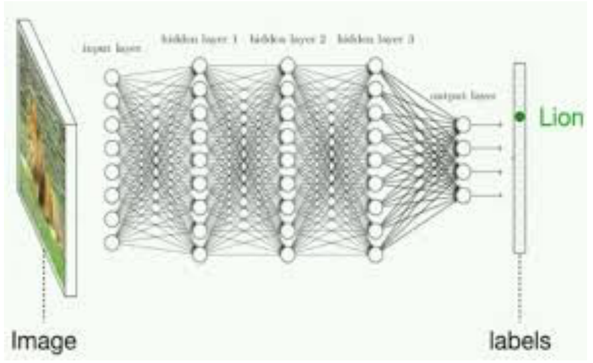

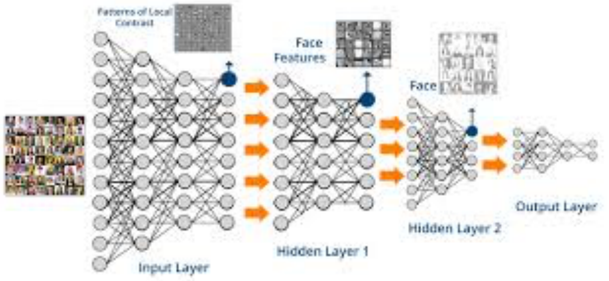

Deep Neural Networks

Topics of Interest in this Course

- **Architectures for perception and control in IoT/CPS**
  - Spatiotemporal, multimodal/crossmodal
  - Classical / probabilistic, deep learning, neurosymbolic, data+physics approaches, deep RL
  - Various inferences tasks and queries
    - detection, classification, localization of events/activities
    - forecast, superresolution, anomaly detection, …
    - control
- **Real world issues and challenges**
  - Missing, noisy and incomplete data; noisy and imprecise labels
  - Irregular sampling, timestamp misalignment
  - Domain shift, sensor adaptation, Sim2Real
- **System-level issues and challenges**
  - Resource availability, dynamics and resource constrains
- **Human/societal issues and challenges**
  - Explainability and interpretability
  - Safety, privacy, FATE

- We will explore these via a combination of lectures, paper presentation, assignments, and project

Topics Explored

- Inference / prediction with temporal data
- Advanced issues
  - Multimodal / cross-modal
  - Irregular sampling
  - Point cloud
  - Spatially distributed sensors
- Data quality
  - Missing data values, time alignment
- Data vs Physics
- Domain shift
- Distributing across edge-cloud
- Control with Deep RL and Sim2Real

Course Components

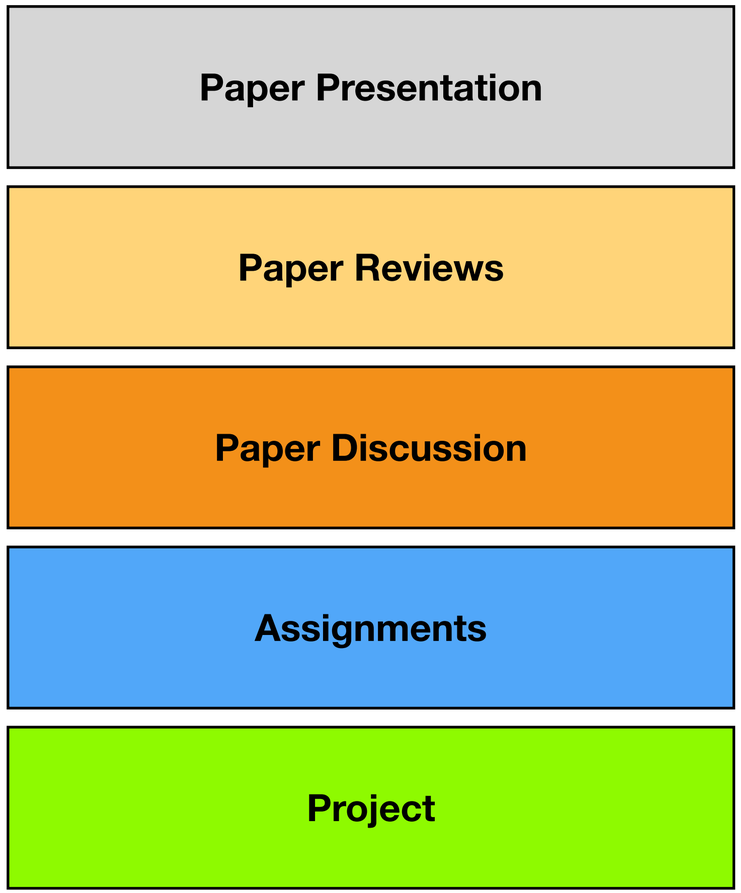

Individual
Present assigned paper (20 min)

15%

Individual
Written reviews of assigned papers (~10)

15%

Individual
Participate in discussion on papers (in class or on Piazza)

10%

Individual or Team as specified
Concepts and analysis

15%

Individual or Team (max 2)
(Website + Video Demo/Talk)

45%

Grading

- Some components of the course will be graded on a numerical scale while others directly on a letter grade scale (A+=4.3, A=4, A-=3.7, … D-=0.7, F=0)
- Expectation for letter grade:
  - A+/A/A- = correct / high-quality; B+/B/B- = some-flaws / acceptable; C+/C/C- = significant-flaws / below-average-quality; D+/D/C- = mostly-incorrect / unacceptable-quality; F = did not make a serious attempt
- Numerical scores will be mapped to letter grade as follows
  - Scale to 100
  - If median<80, normalize so that highest=100 and median=80
  - Mapping: scores ≥95 = A, ≥90=A-; ≥85=B+, ≥80=B, and so on.
- Course grade
  - Compute weighted course score on scale of 4
  - Convert to letter grade using mid-points between adjacent letter grades as thresholds
    - A+=(4,4.3], A=[3.85,4], A-=[3.5,3.85), etc.
    - thresholds may be selective nudged lower at instructor's discretion to account for unusual clustering
  - S grade means letter grade of B or better [UCLA rules
- Historical: Median course grade A-; Mean between B+ and A-

Research Paper Presentations

- List of research papers from literature provided by me
- Each student has to present once during the course
  - Most papers will be for individual presentation
  - But some long ones may be designated for two person presentation
- Presentation slot:
  - Main presentation: 20 min talk per paper + 10-min Q&A
- Slide preparation
  - Expected to be original overall, but may incorporate elements from slides obtained from authors
  - Have slides at the end giving your own perspectives on strengths, weaknesses, and additional directions
- Participation by everyone
  - Reviews of ~10 (maximum of 12) papers assigned to each student by the instructor
    - 4-slide file on summary, strength, weakness, and future work
  - Participation in Q&A (during lecture or on Piazza)
- Starting Week 3

Course Project Expectations

- What should be your goal?
  - Dig deep into a focus area on your own as lectures would provide a “broad” coverage
  - Something useful, cool, or risky ideas leading to negative results are fine
  - The key is to keep the project simple, and focused
  - Aim for something concrete and tangible, even if minor, and high quality
- What type of projects are not acceptable?
  - Application, i.e. collect data (or use available data), train model(s), and report performance, without any associated methodological novelty
  - Literature survey
- The project must have a clear hypothesis and associated questions that it will explore
- While teams of 2 are allowed, it is with the understanding that if your partner drops the course then you would need to finish the project by yourself, and other than giving you extra time the instructor will not be able to provide any other consideration

Thinking about Project Ideas: Heilmeier Catechism

- George H. Heilmeier, a former DARPA director (1975-1977), crafted a set of questions known as the "Heilmeier Catechism" that anyone proposing a research project or product development effort should be able to answer

  - What are you trying to do? Articulate your objectives using absolutely no jargon.
  - How is it done today, and what are the limits of current practice?
  - What is new in your approach and why do you think it will be successful?
  - Who cares? If you are successful, what difference will it make?
  - What are the risks?
  - How much will it cost?
  - How long will it take?
  - What are the mid-term and final “exams” to check for success?

- Of course not all are relevant to a course project setting…

Course Project Timeline and Grading

- Week 2-3: Mandatory initial discussion with instructor about project ideas [2%]
  - Document and summarize meeting via a direct non-anonymous message on Piazza to instructor
- Week 3-4: Mandatory second discussion with instructor about project ideas [2%]
  - Document and summarize meeting via a direct non-anonymous message on Piazza to instructor
- Week 5: Initial project repo and website ready on GitHub [6%]
  - Submit project title, team information, GitHub repo URL, and GitHub website URL via web form
  - Website must have initial sections for goals, specific aims,
- Week 7: Midterm checkpoint [20%]
  - YouTube video of 8 min max of presentation plus any initial demo [80%]
  - Website updated and also add initial text for data, tools, methods, and results [20%]
- Finals Week: Submission [70%]
  - YouTube video of 15 min max of presentation plus any demo [80%]
  - Website finalized [20%]
- Grading will consider
  - problem and approach: difficulty (relative to team size), novelty
  - execution: architecture, analysis, evaluation, effort, elegance
  - presentation: quality, organization, discussion

No Textbook

But, Some Useful Books…

- Cook, Diane J., and Narayanan C. Krishnan. Activity learning: discovering, recognizing, and predicting human behavior from sensor data. John Wiley & Sons, 2015.
  - Available for free via UCLA network
    - https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781119010258

- Brunton, Steven L., and J. Nathan Kutz. Data-driven science and engineering: Machine learning, dynamical systems, and control. Cambridge University Press, 2019.
  - Available for free via UCLA network
    - https://www.cambridge.org/core/books/datadriven-science-and-engineering/77D52B171B60A496EAFE4DB662ADC36E
  - Great videos by Brunton on YouTube

- Russell, Stuart J., and Peter Norvig. "Artificial Intelligence: A Modern Approach" (4th edition, 2020)
  - Book not available but you can find slides on various chapters at http://aima.cs.berkeley.edu/index.html

- Warden, Pete. TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low. 2020.
  - https://tinymlbook.com
  - Free from UCLA library

Desired Student Background

- Exposure to ML methods, both classical (linear regression, decision tree, SVM, clustering etc.) and deep learning based (MLP, CNN, RNN/LSTM etc.)
- Comfortable with linear algebra and probability as most papers use them heavily
- Experience with Python including scikit-learn and Tensorflow/Keras or Pytorch
  - equivalent experience with MATLAB also acceptable, but you will be on your own as I don’t have expertise in MATLAB
- For projects you may (depending on project you select) benefit from software expertise in mobile, wearable, or embedded platforms
- It is fine if you don’t have some of the above-mentioned expertise
  - there are plenty of web resources for self-learning, on Wikipedia, YouTube etc.

Plagiarism, Cheating, and other Disallowed Behaviors

- Discussing solution approaches to problems with classmates is okay but showing or sharing your writeup, code etc. is never okay
- Using material from external sources (even if provided by me as reference) is acceptable only if explicitly allowed but must be acknowledged by a clear citation
- Generally:
  - For paper presentations: you need to make your own but can use graphics from the paper or from slides obtained from authors, with explicit acknowledgement
  - For assignment: the work needs to be your own, though you can draw inspiration from approach from provided references and external sources (if allowed) and must cite them.
  - For project: your are free to use any legally available code, data, and hardware provided you cite with explicit acknowledgement
  - If in doubt, discuss with me

Post any questions on Piazza!!!