

Robot Learning

Introduction

Prof. Songhwai Oh

ECE, SNU

Intelligent Systems



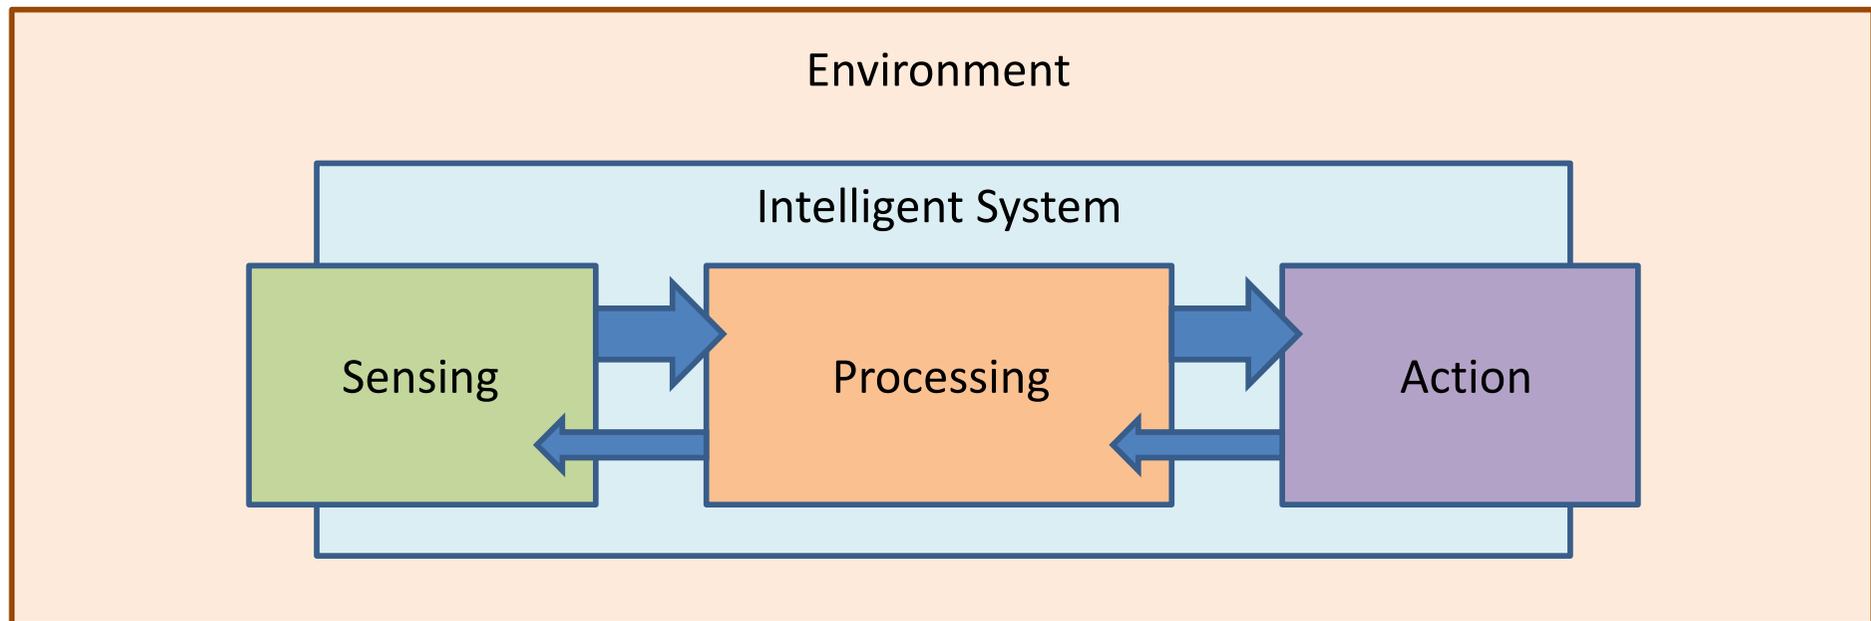
BigDog (2005)



PR2 (UC Berkeley, 2010)



Google Car (2014)



Learning

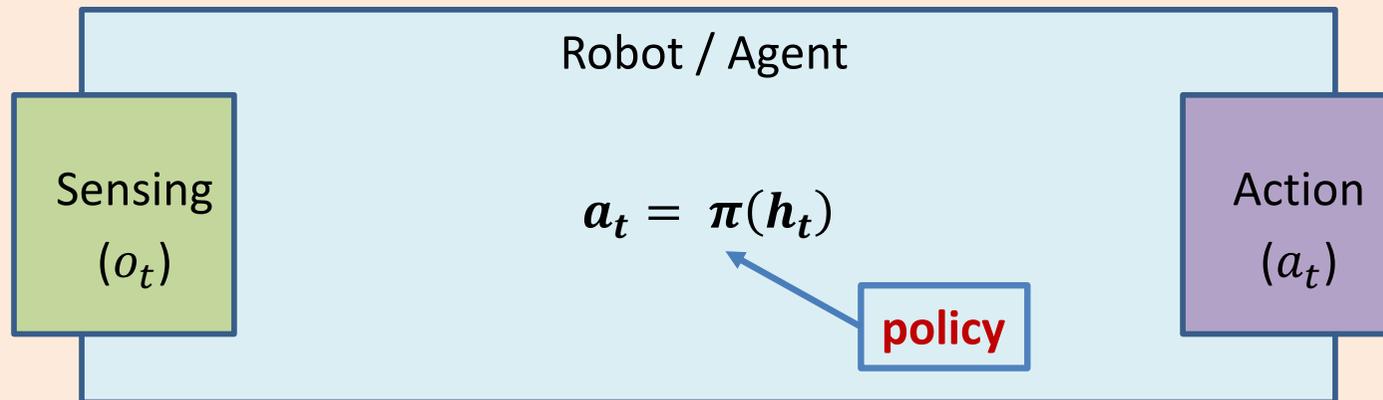
Google
(2016)



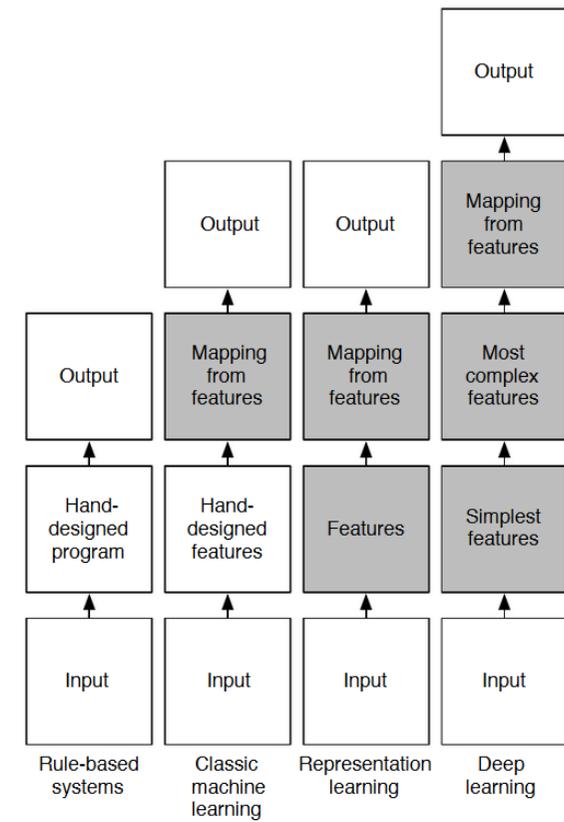
- 14 robots: 800,000 grasp attempts (3,000 robot hours)
- CNN grasp predictor
 - input: images, motor command
 - output: grasp success probability
- Cross entropy (important sampling) is used to find a good motor command

Robot Learning

Environment (Ω)



- Agent's action changes the environment
 - $o_t = \Omega(s_t, h_t, a_t)$ $s_t =$ environment's state
- Observations (o_t) are correlated over time
- History: $h_t = (o_1, a_1, o_2, a_2, \dots, o_{t-1}, a_{t-1}, o_t)$
- **Robot learning finds a policy from data (o_t, a_t)**



DEEP LEARNING FOR SENSING

ImageNet Large-Scale Visual Recognition Challenge, 2012

Tasks:

- Decide whether a given image contains a particular type of object or not. For example, a contestant might decide that there are cars in this image but no tigers.
- Find a particular object and draw a box around it. For example, a contestant might decide that there is a screwdriver at a certain position with a width of 50 pixels and a height of 30 pixels.

Image classification

Easiest classes

red fox (100) hen-of-the-woods (100) ibex (100) goldfinch (100) flat-coated retriever (100)



tiger (100)

hamster (100)

porcupine (100)

stingray (100)

Blenheim spaniel (100)



Hardest classes

muzzle (71) hatchet (68) water bottle (68) velvet (68) loupe (66)



hook (66)

spotlight (66)

ladle (65)

restaurant (64)

letter opener (59)



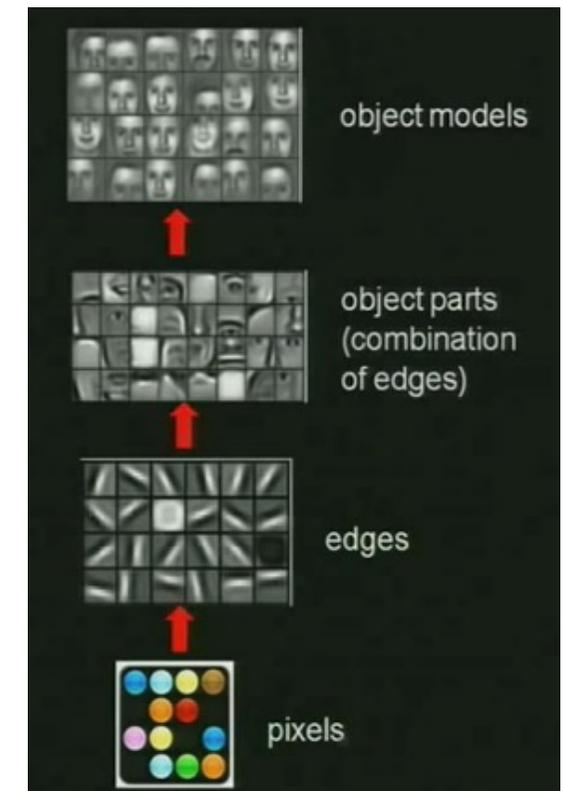
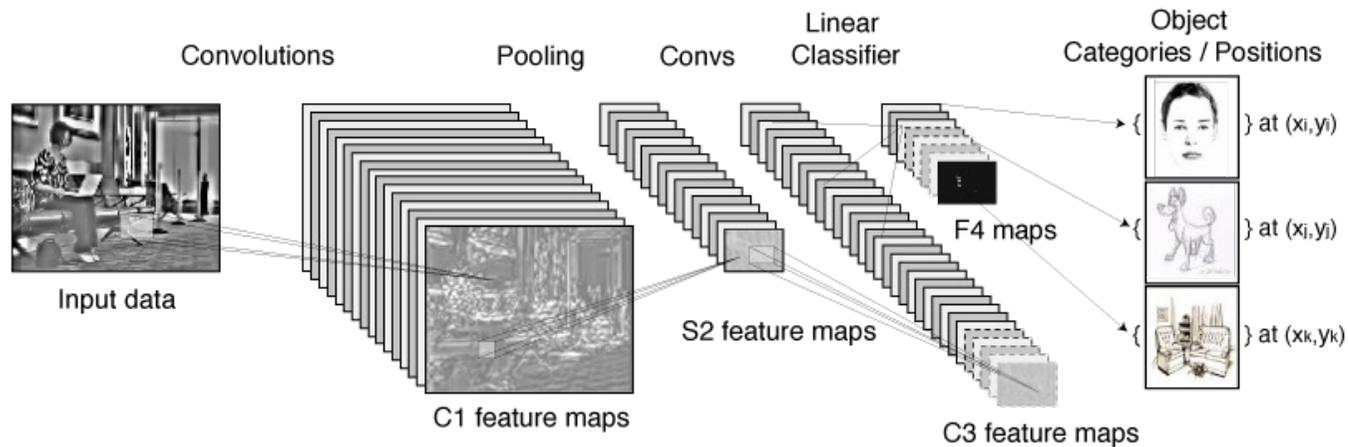
- 1000 different categories
- Over 1 million images
- Training set: 456,567 images

Year	Winning Error Rate
2010	28.2%
2011	25.8%
2012	16.4% (2 nd 25.2%)
2013	11.2%
2014	6.7%
2015	3.57%
Human	About 5.1%

ImageNet Large Scale Visual Recognition Challenge. Russakovsky et al. *arXiv preprint arXiv:1409.0575*. URL: <http://arxiv.org/abs/1409.0575v1>

Convolutional Neural Networks (CNNs)

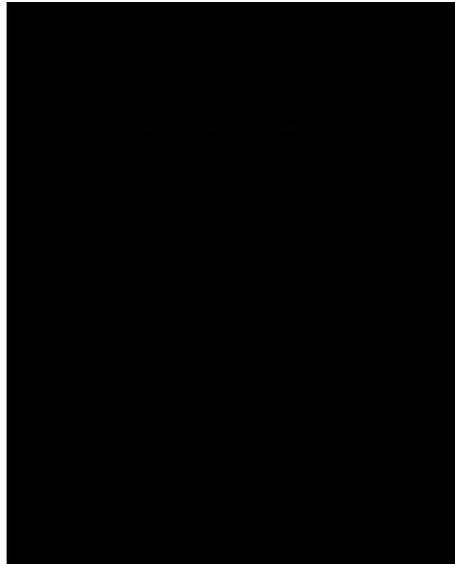
- SuperVision (2012)
 - Deep convolutional neural network
 - 650,000 neurons
 - 5 convolutional layers
 - Over 60 million parameters
- Clarifai (2013)
- GoogleLeNet (2014) – 22 layers
- ResNet (2015) – 152 layers



DEEP LEARNING FOR ACTION (AND SENSING): DEEP REINFORCEMENT LEARNING (RL)

Deep Q-Network (DQN), 2013

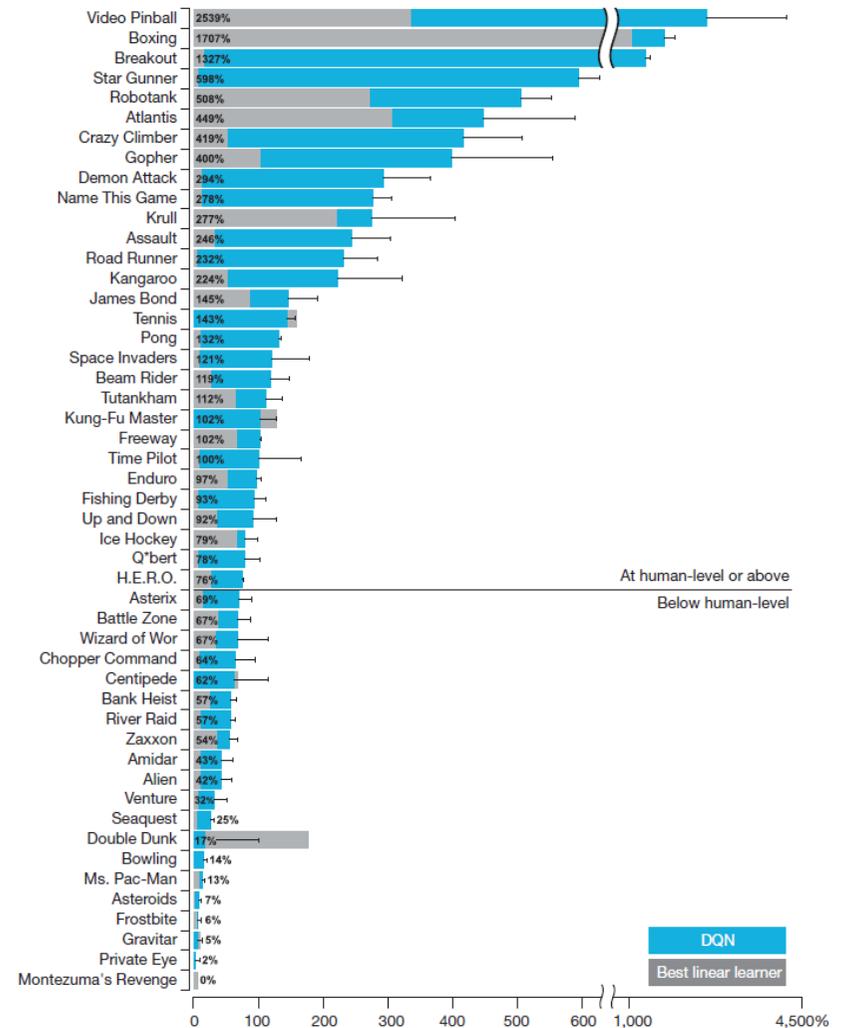
- Playing Atari games
- Input: Game screen shots
- Output: Control (left, right, shoot, ...)
- Convolutional neural networks (CNN)
- Reinforcement learning: Q-learning



Breakout



Space Invaders



AlphaGo, 2016

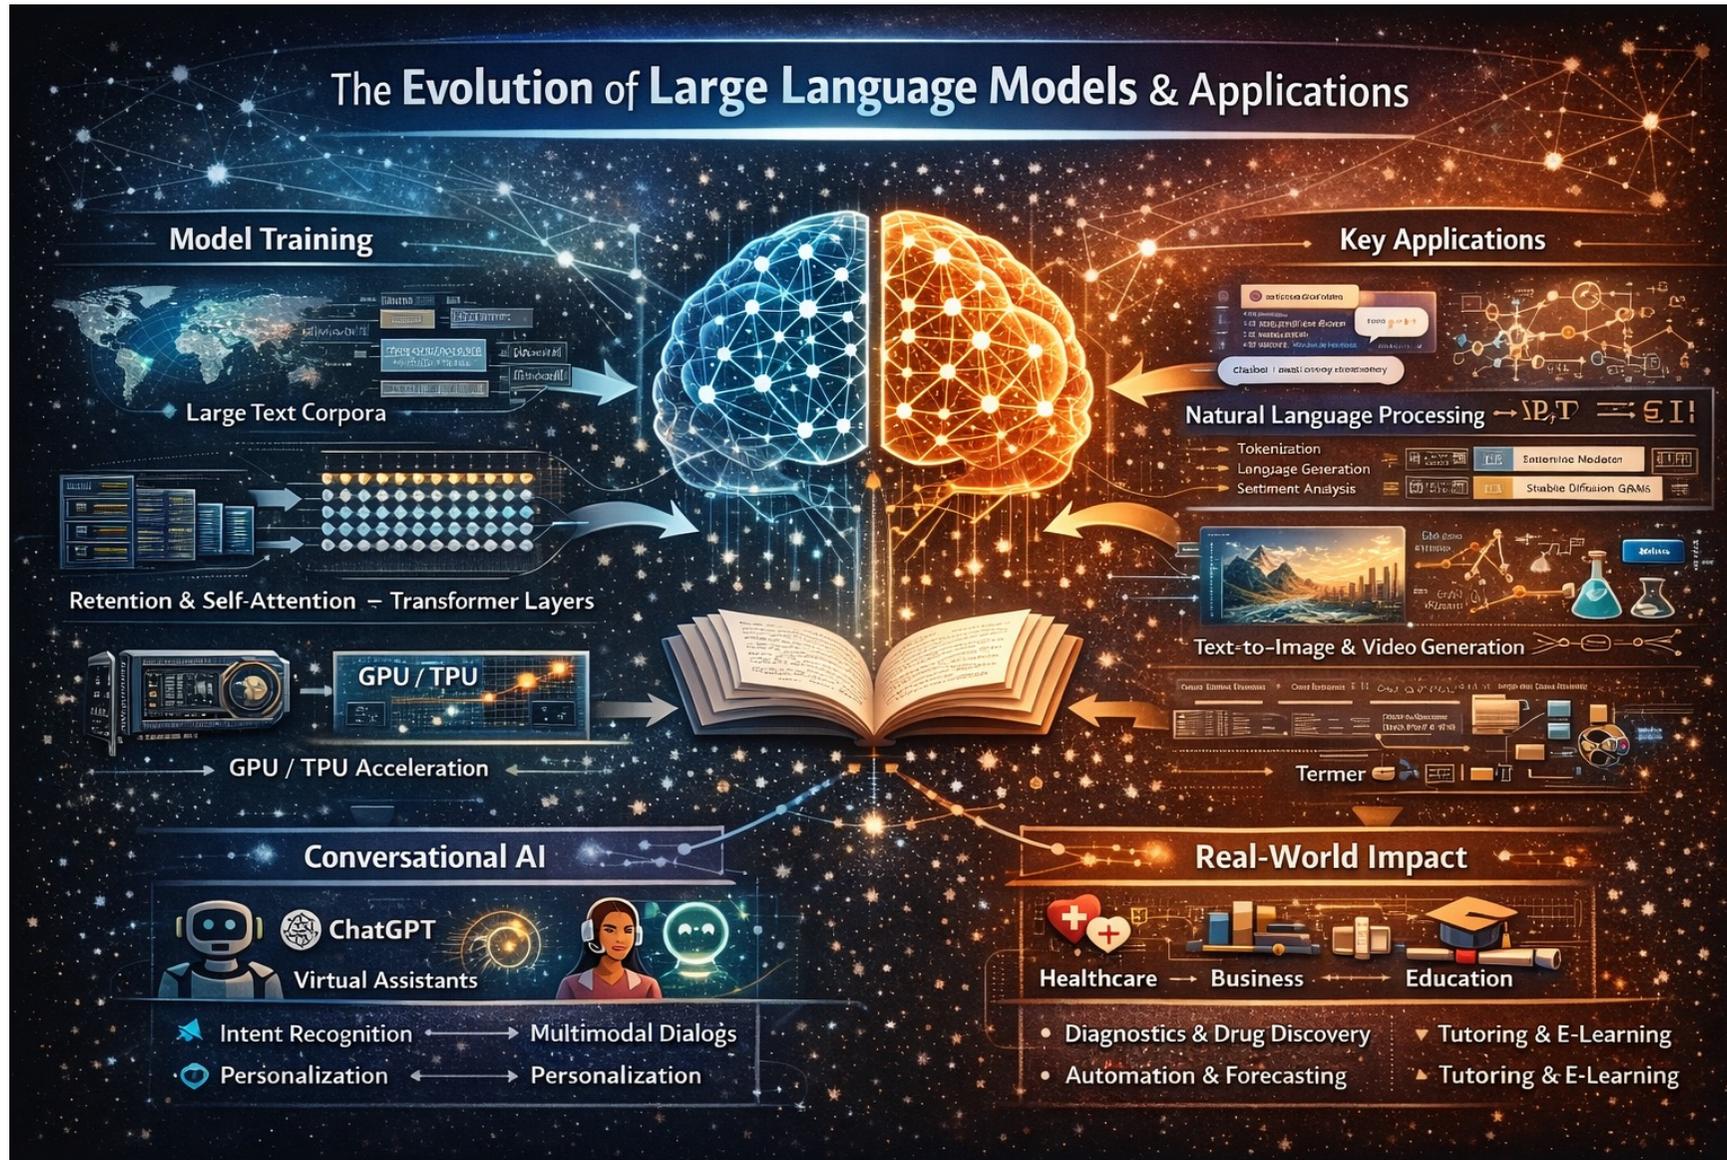
Google DeepMind's AlphaGo vs. Lee Sedol, March 2016

- Possible board positions of Go: 10^{170}
 - cf. Chess: 10^{47}
- Monte Carlo tree search
- Deep neural networks:
 - Value network
 - Policy network
- Reinforcement learning
- Trained from
 - 30 million human moves
 - Playing against itself
- 1,202 CPUs, 176 GPUs

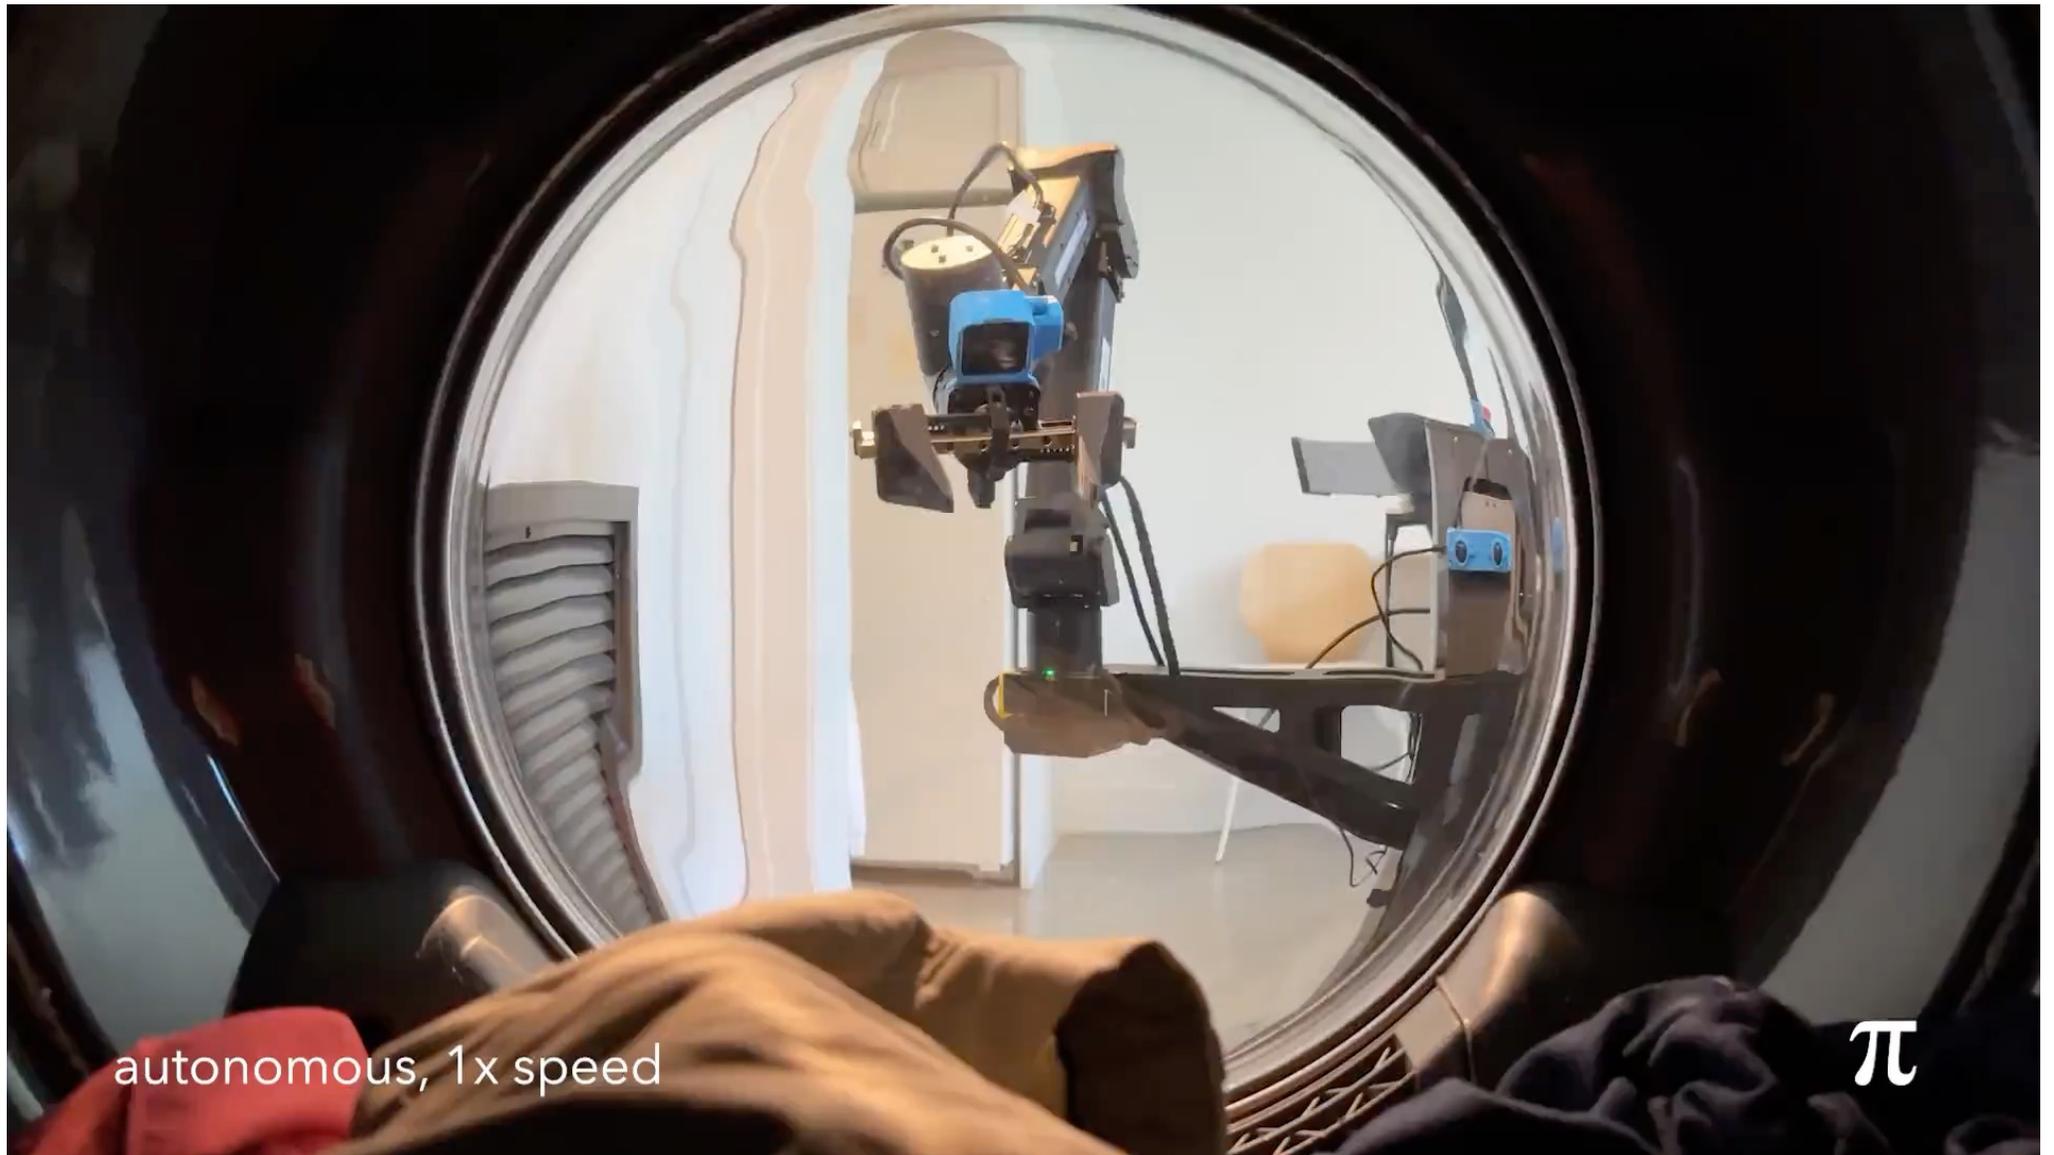


Game	Date	Black	White	Result	Moves
1	9 March 2016	Lee Sedol	AlphaGo	Lee Sedol resigned	186 Game 1
2	10 March 2016	AlphaGo	Lee Sedol	Lee Sedol resigned	211 Game 2
3	12 March 2016	Lee Sedol	AlphaGo	Lee Sedol resigned	176 Game 3
4	13 March 2016	AlphaGo	Lee Sedol	AlphaGo resigned	180 Game 4
5	15 March 2016	Lee Sedol ^[note 1]	AlphaGo	Lee Sedol resigned	280 Game 5
Result: AlphaGo 4 – 1 Lee Sedol					

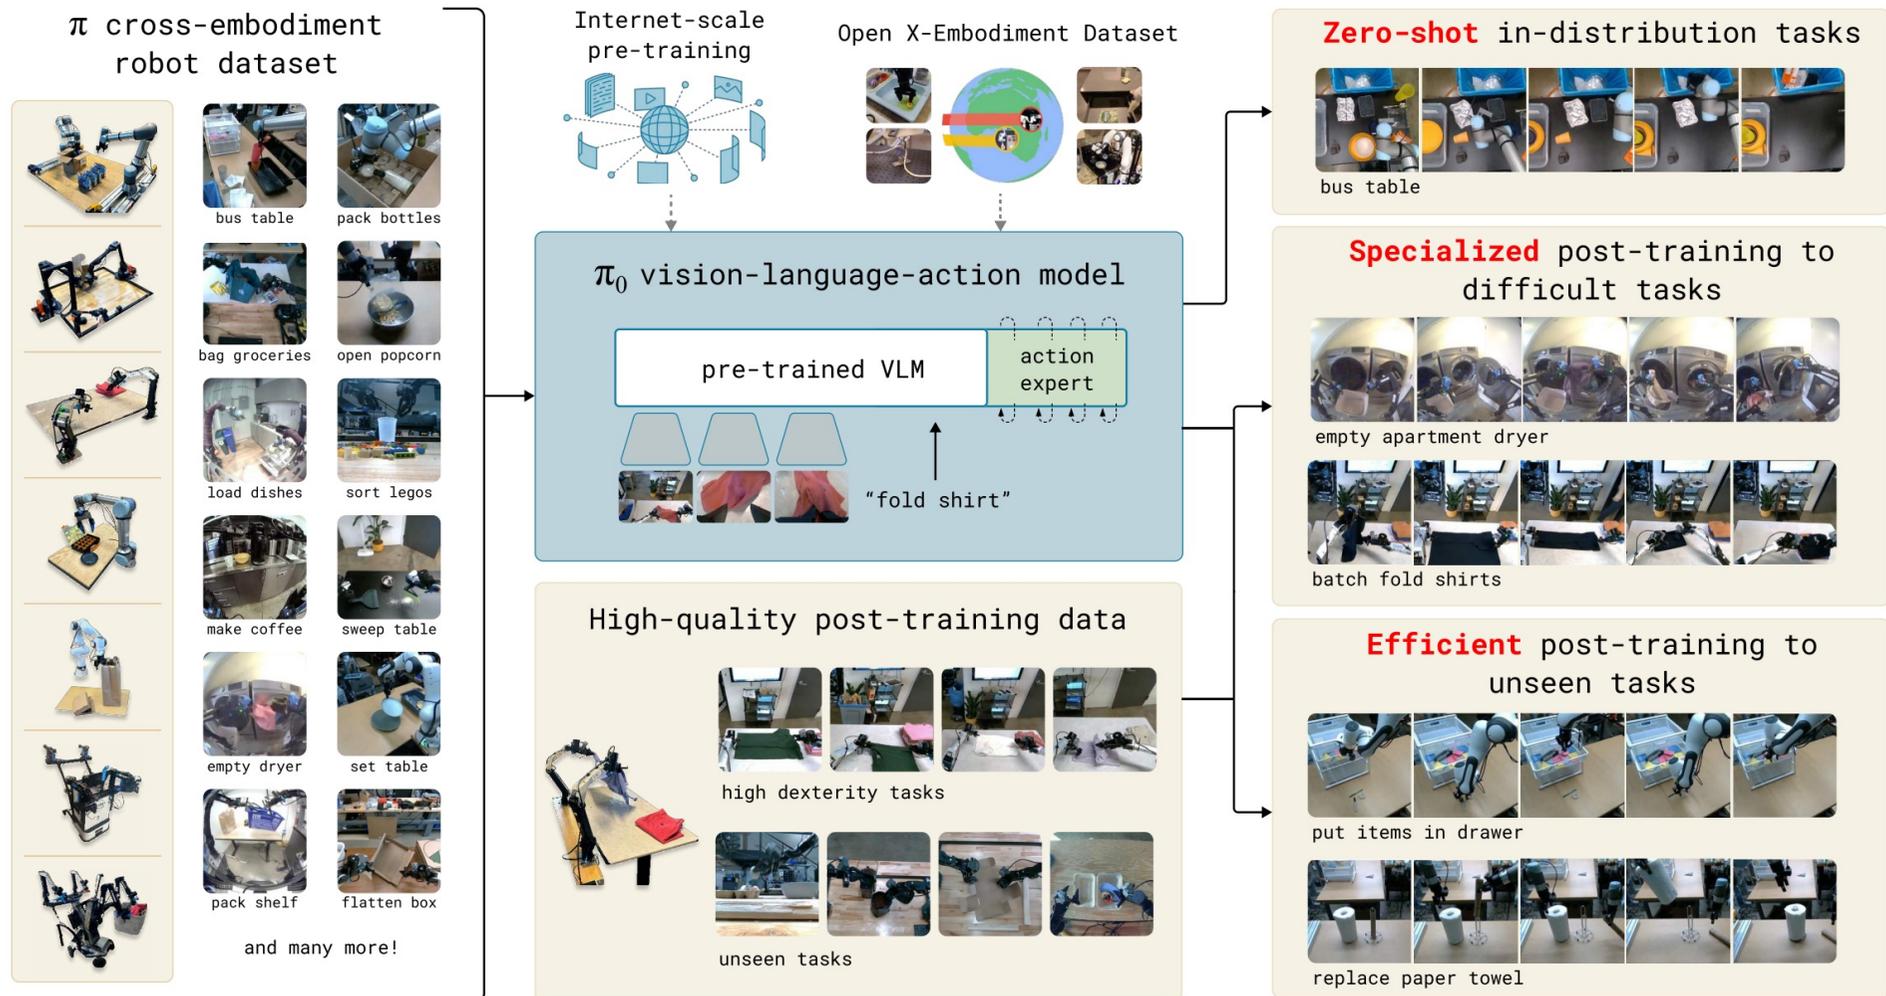
Large Language Models, 2022



Robot Foundation Model



π_0 : A Vision-Language-Action Flow Model for General Robot Control



source: <https://www.physicalintelligence.com/blog/pi0>

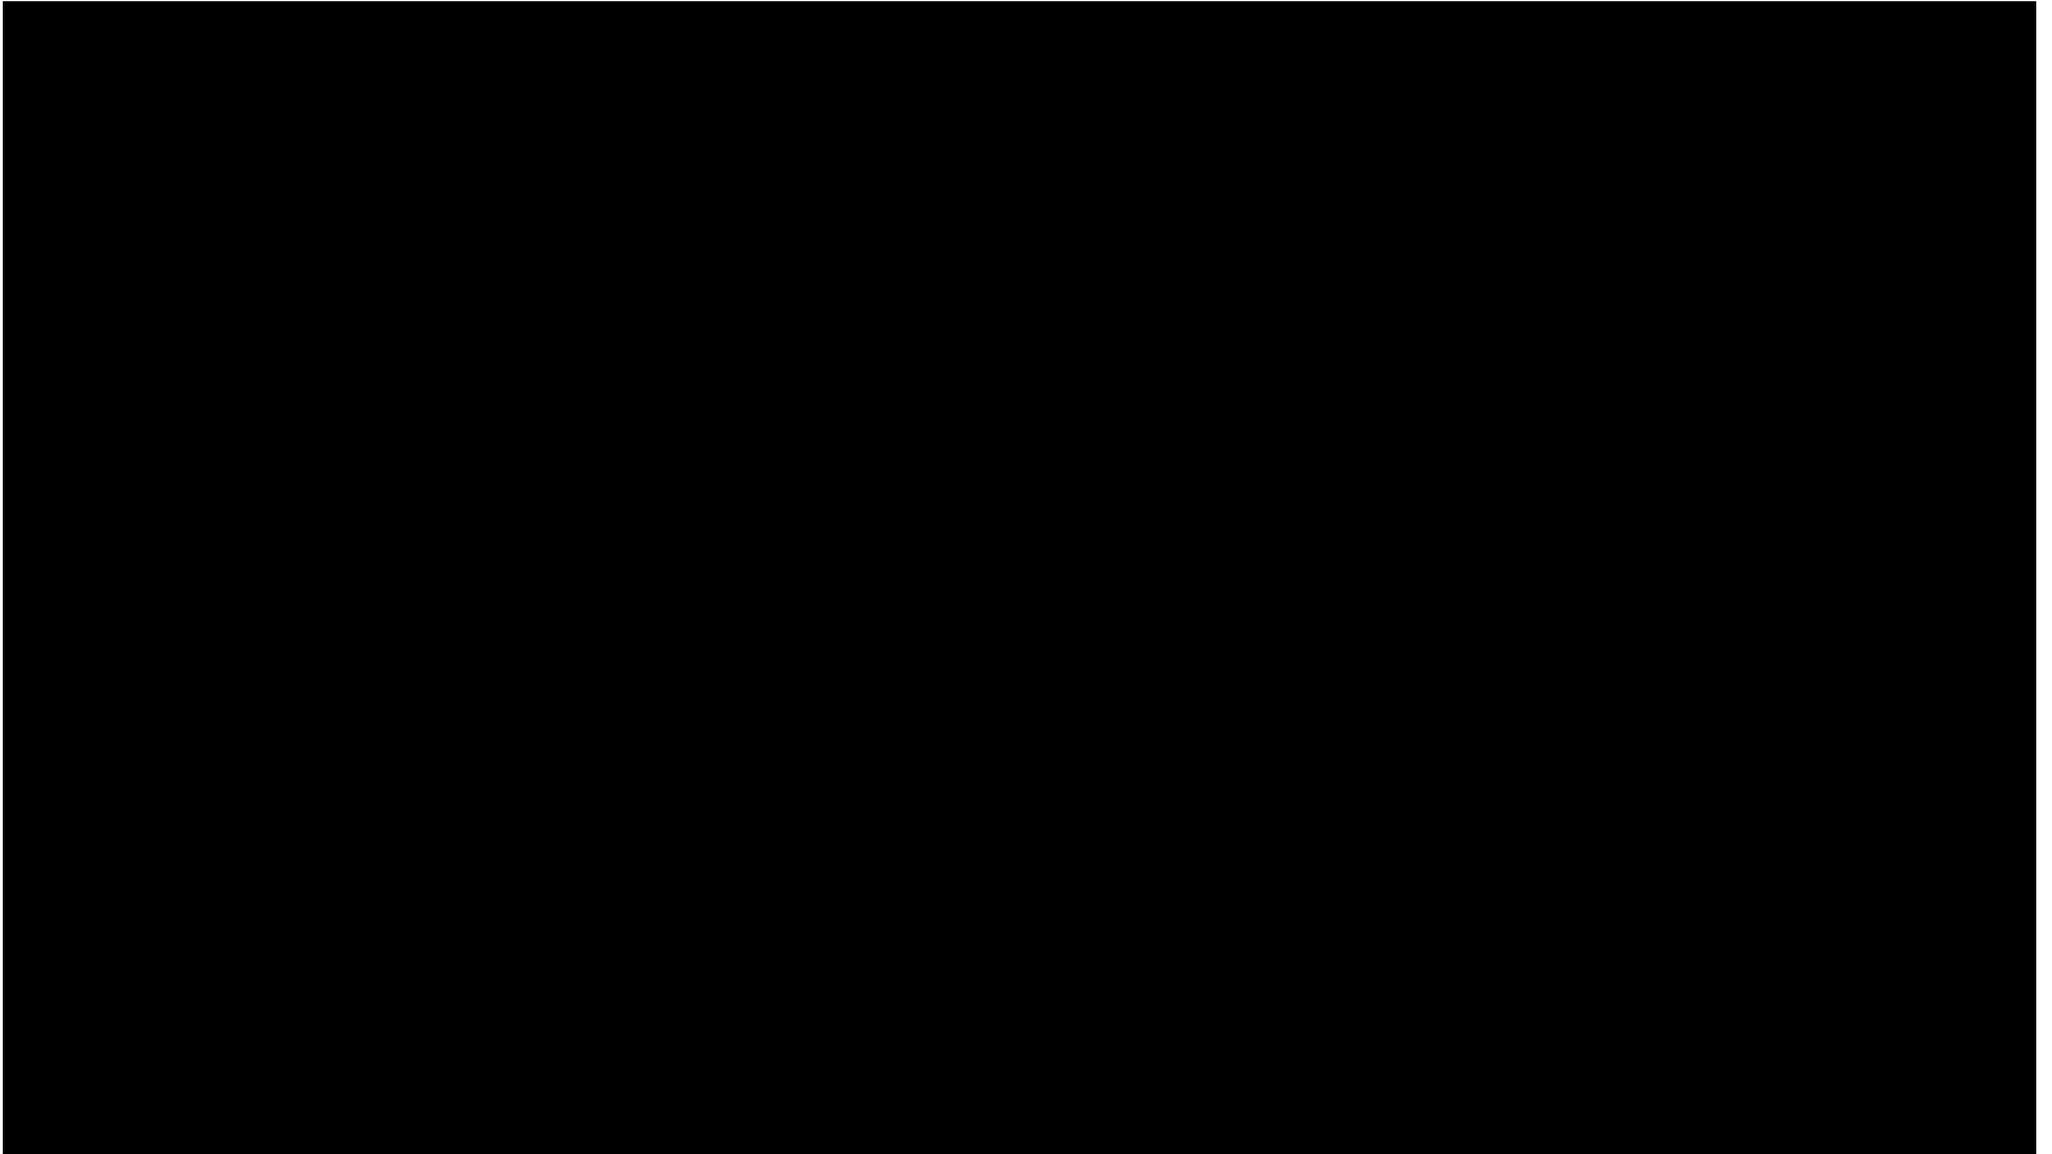
Humanoid: Figure AI + LLM



Humanoid: Tesla Optimus



Imitating Human Motion



source: <https://agile.human2humanoid.com/>

Why do animals have brains?

- Sea squirt



- Swims when young
- Implants on a rock and never leaves
- **Eats its own brain**

- Brain is for action

- To predict the consequences of actions
- The difference between the prediction and sensory input is due to external causes

- Action is one of the most challenging problems in AI

COURSE OUTLINE

Course Information

- Instructor: Songhwai Oh (songhwai@snu.ac.kr)
- TA: Geunje Cheon (geunje.cheon@rllab.snu.ac.kr)

- Homepage:
 - <https://rllab.snu.ac.kr/courses/robot-learning-2026>
- Class Board:
 - eTL
- Prerequisites
 - (430.457) Introduction to Intelligent Systems (지능시스템개론)
 - Requires strong background in algorithms, linear algebra, probability, and programming
- Topics
 - Review on MDPs, POMDPs, RL algorithms
 - Behavior cloning
 - Inverse reinforcement learning
 - Policy gradient
 - Deep Q-network (DQN)
 - Generative adversarial imitation learning (GAIL)
 - Maximum entropy reinforcement learning
 - Safe reinforcement learning
 - Offline reinforcement learning
 - Robot foundation models,

Grading

- Class participation
- 5-6 programming assignments
- Paper presentation (in English)
 - Will be assigned
 - Can be traded with classmates
- Reading recent research papers
 - About 2-3 papers per week (practice quick reading)
 - Submit 2 questions for each paper (2 days before)
- Term project [**approx. schedule**]
 - Project proposal (2 pages, 2 columns) [05/06]
 - Project presentation and poster session [06/10]
 - Project report (6 pages, 2 columns) [06/12]
 - Project topics from deep RL
 - comparing existing methods with **new** analyses
 - **new** algorithms
 - **new** theory
 - ... (**requires your new contributions**)