CS 343
Overview and Conclusion

Profs. Peter Stone and Yuke Zhu

The University of Texas at Austin

TEXAS
The University of Texas at Austin

Announcements

- Please fill out the course survey
  - Feedback to both instructors and TAs
  - Positive and negative points are useful
  - Post on Piazza/Email your completion screenshot teaching staff as a form of participation
- Capture the Flag contest results!

CTF Contest

6 teams participated, 4 qualified

Rankings

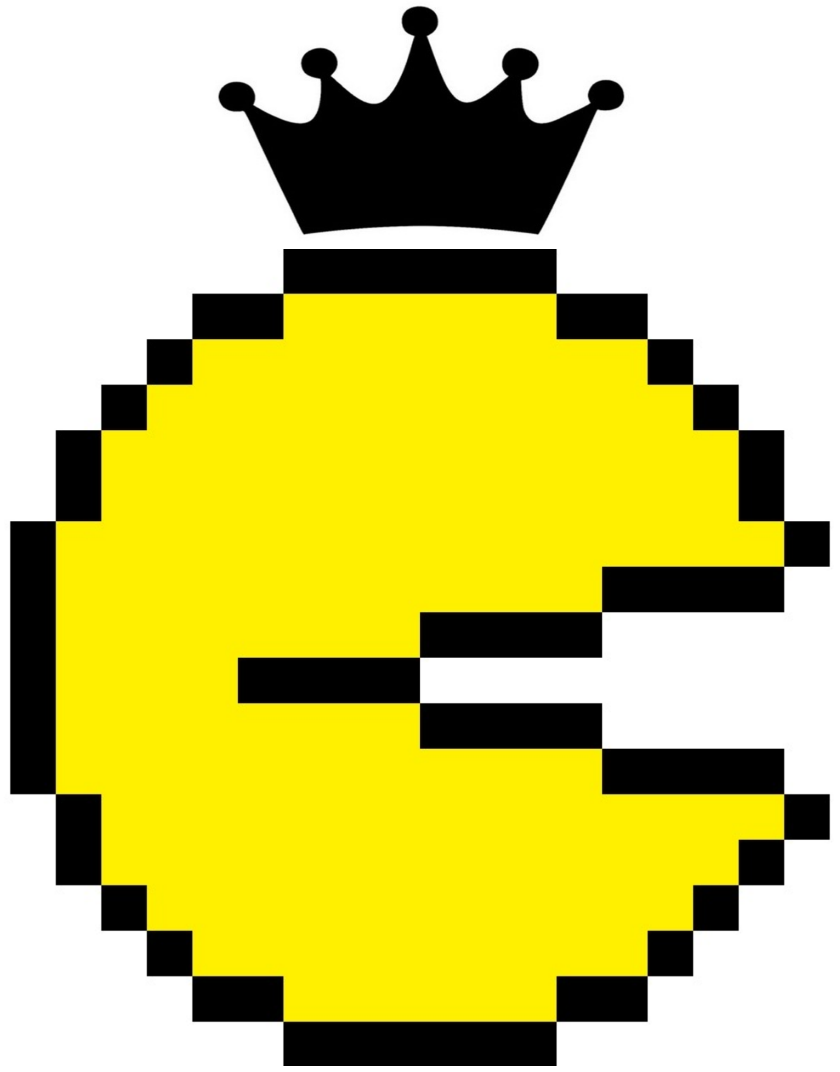

CTF Contest

6 teams participated, 4 qualified

Rankings

4th place: ReinforcedChaos (Nalin Mahajan, Vineeth Bandi)

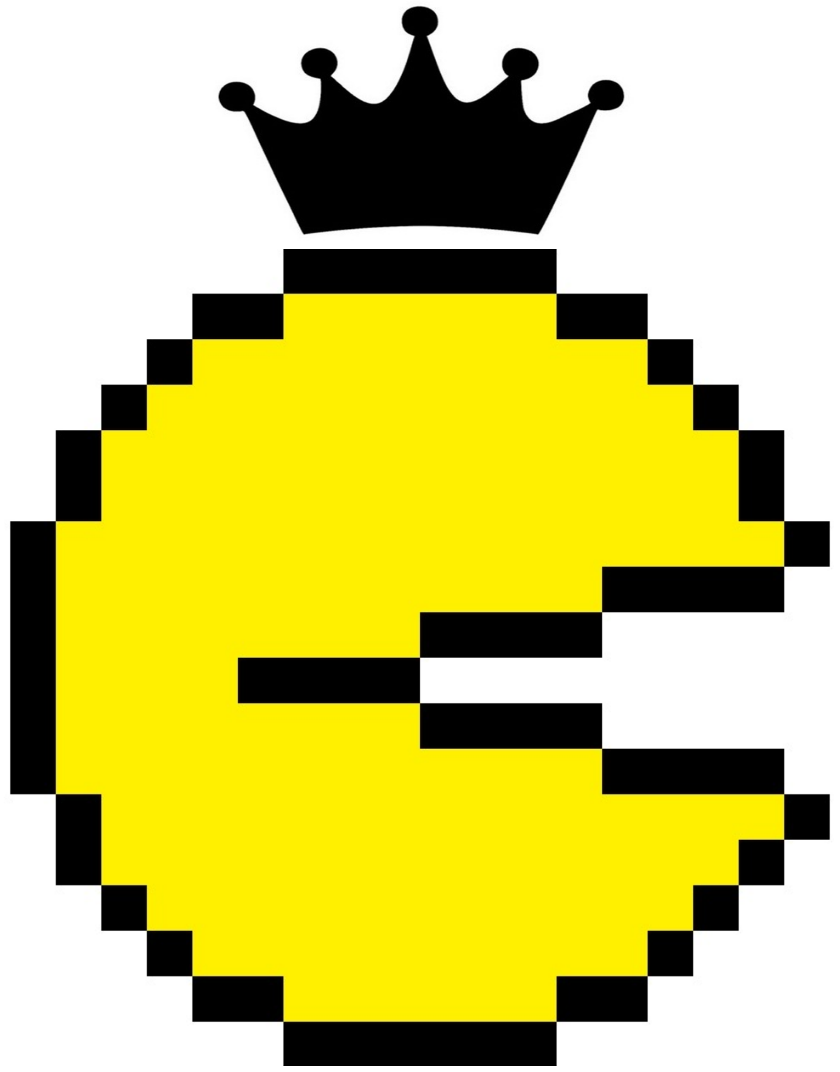

CTF Contest

6 teams participated, 4 qualified

Rankings

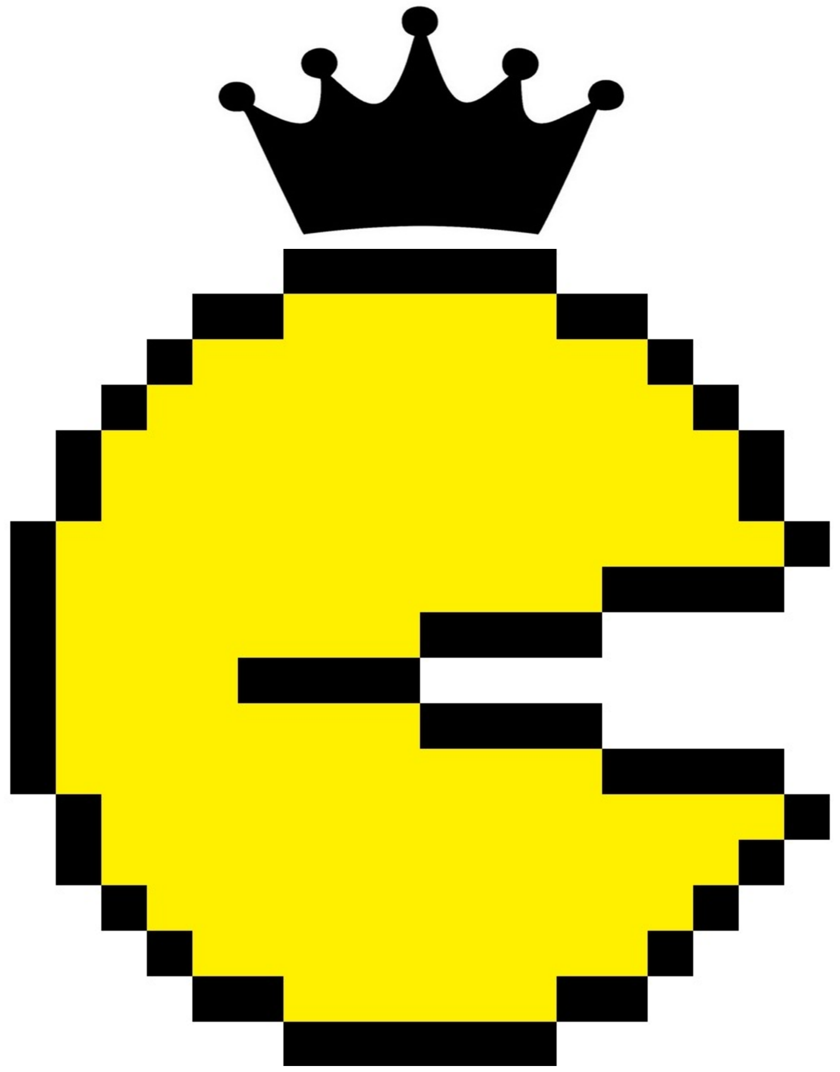

3rd place: KentFoxes (Xuefei Zhao, Yuhan Zheng)

4th place: ReinforcedChaos (Nalin Mahajan, Vineeth Bandi)

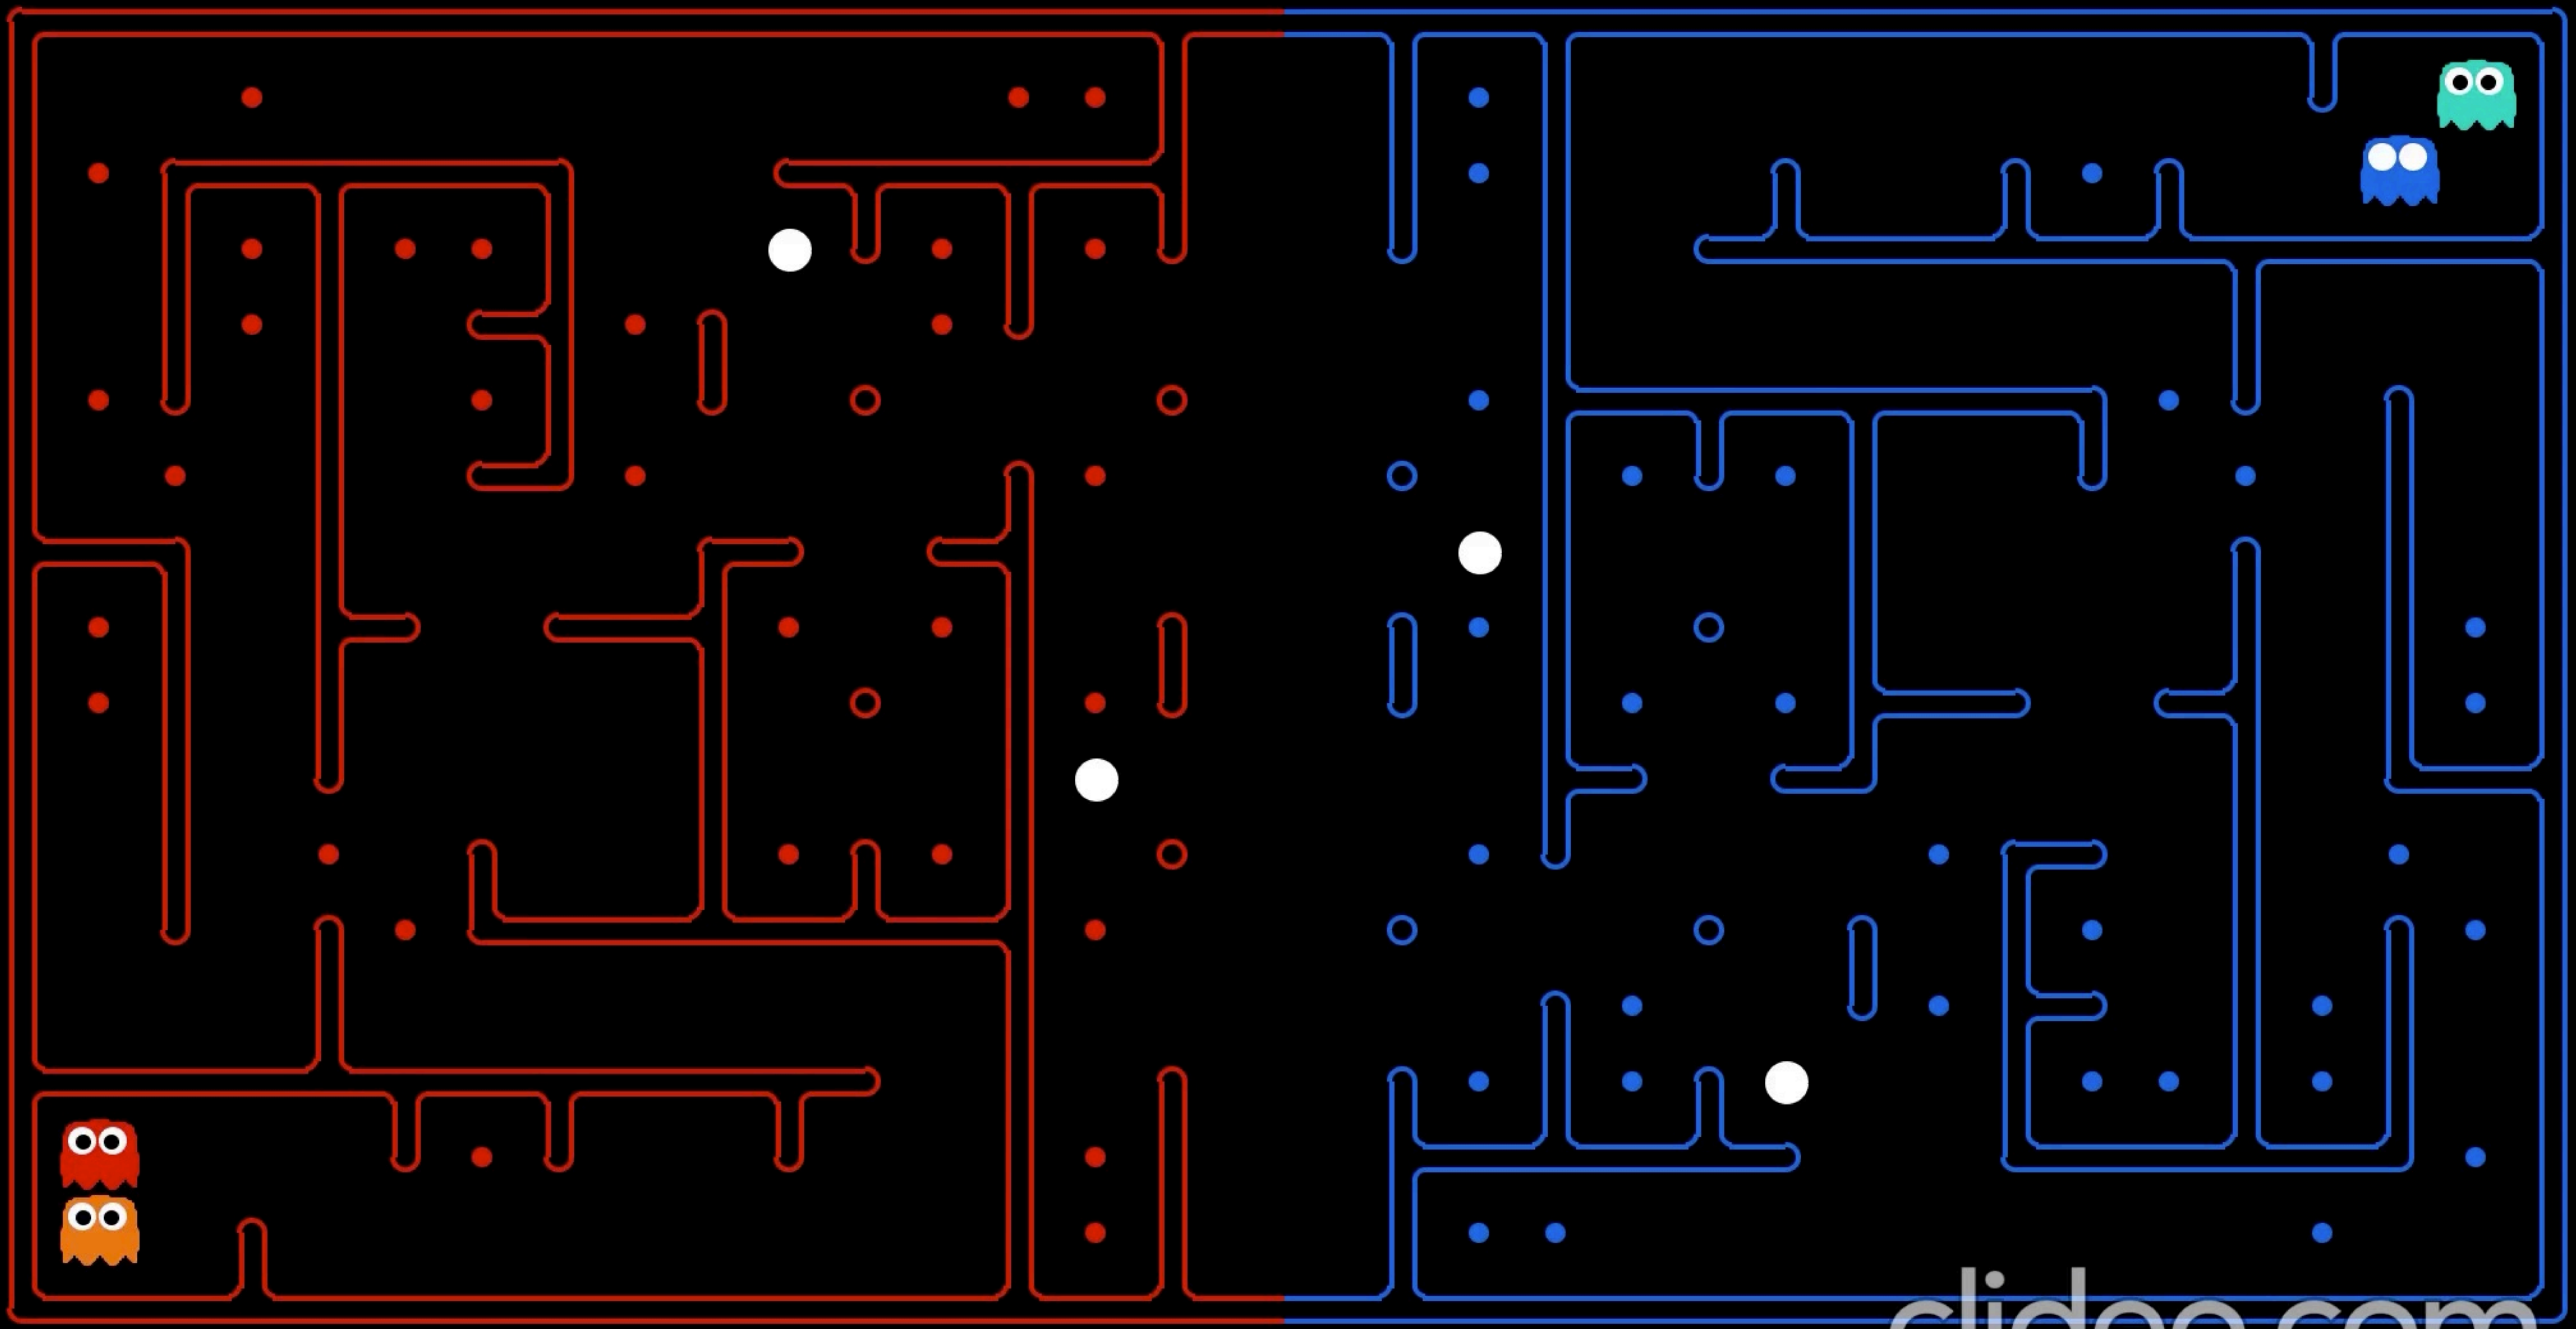

SCORE: 0

CTF Contest

6 teams participated, 4 qualified

Rankings

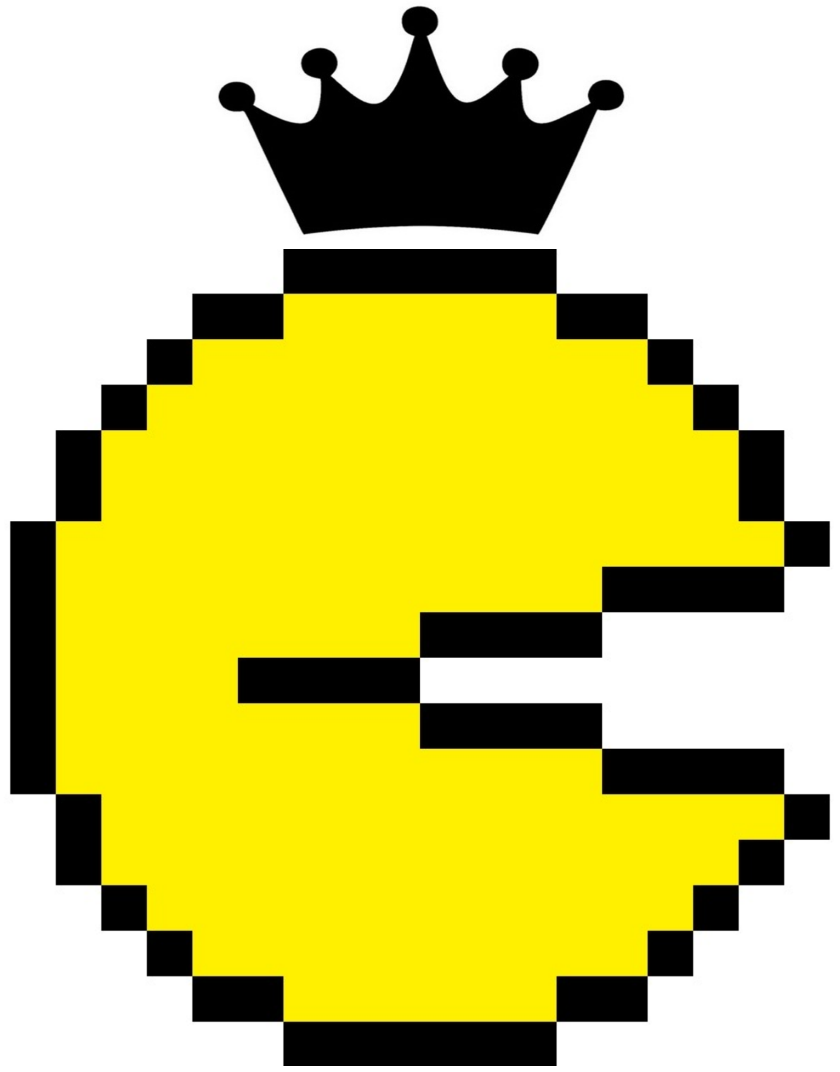

2^{nd} place: EasiestDifficulty (Joseph Muffoletto)

3^{rd} place: KentFoxes (Xuefei Zhao, Yuhan Zheng)

4^{th} place: ReinforcedChaos (Nalin Mahajan, Vineeth Bandi)

CTF Contest

6 teams participated, 4 qualified

Rankings

- 1st place: oranges (Ramya Prasad, Jennifer Suriadinata)
- 2nd place: EasiestDifficulty (Joseph Muffoletto)
- 3rd place: KentFoxes (Xuefei Zhao, Yuhan Zheng)
- 4th place: ReinforcedChaos (Nalin Mahajan, Vineeth Bandi)

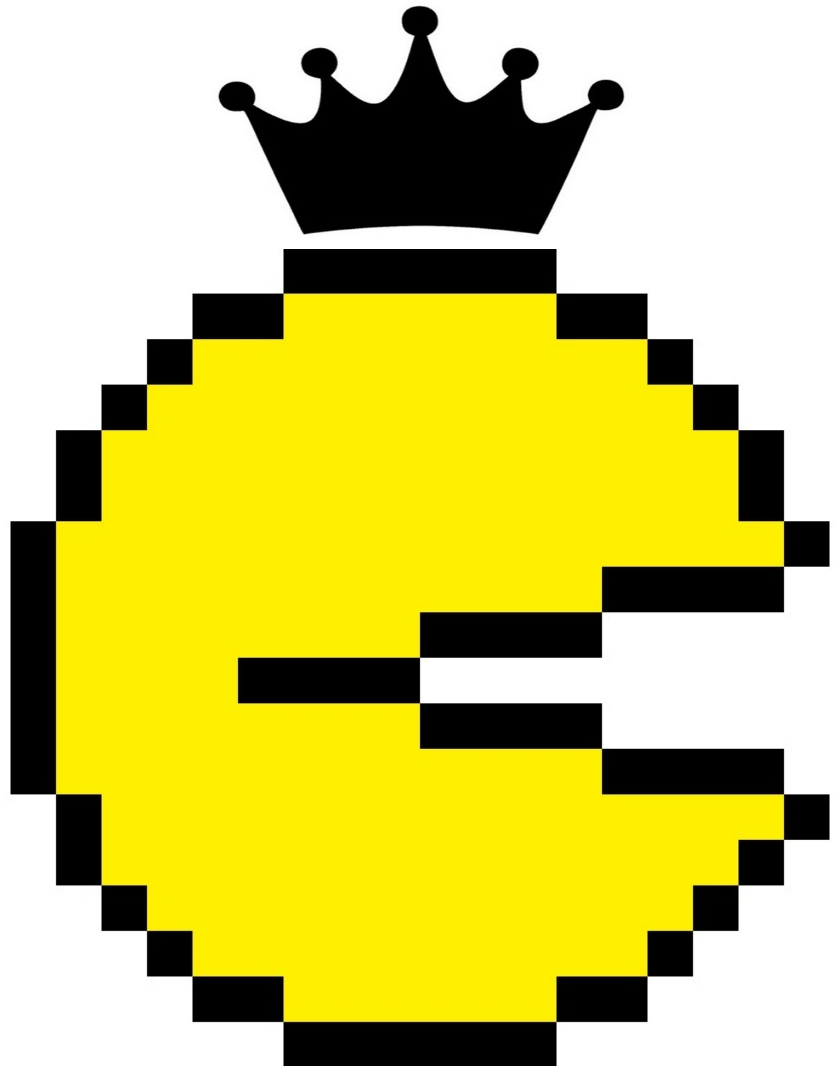

Congratulations to All!

Overview of AI Topics

Search / Planning

- Uninformed Search
- A* Search
- CSPs
- Local Search
- Minimax
- Expectimax
- MDPs
- Classic Planning

Machine Learning

- Reinforcement Learning
- Probability Theory
- Bayes Nets
- HMMs
- Particle Filters
- Decision Diagrams
- Naive Bayes
- Perceptrons
- Neural Networks
- Kernels
- Clustering
- VPI

Overview of Machine Learning

Supervised Learning

Discriminative Models

Perceptrons

Neural Networks

Generative Models

Bayes Nets

Naive Bayes

HMMs

Reinforcement Learning

MDPs

Value Iteration

Policy Iteration

Q Learning

Unsupervised Learning

K-Means Clustering

Hierarchical Agglomeration

Maximize Your Expected Utility

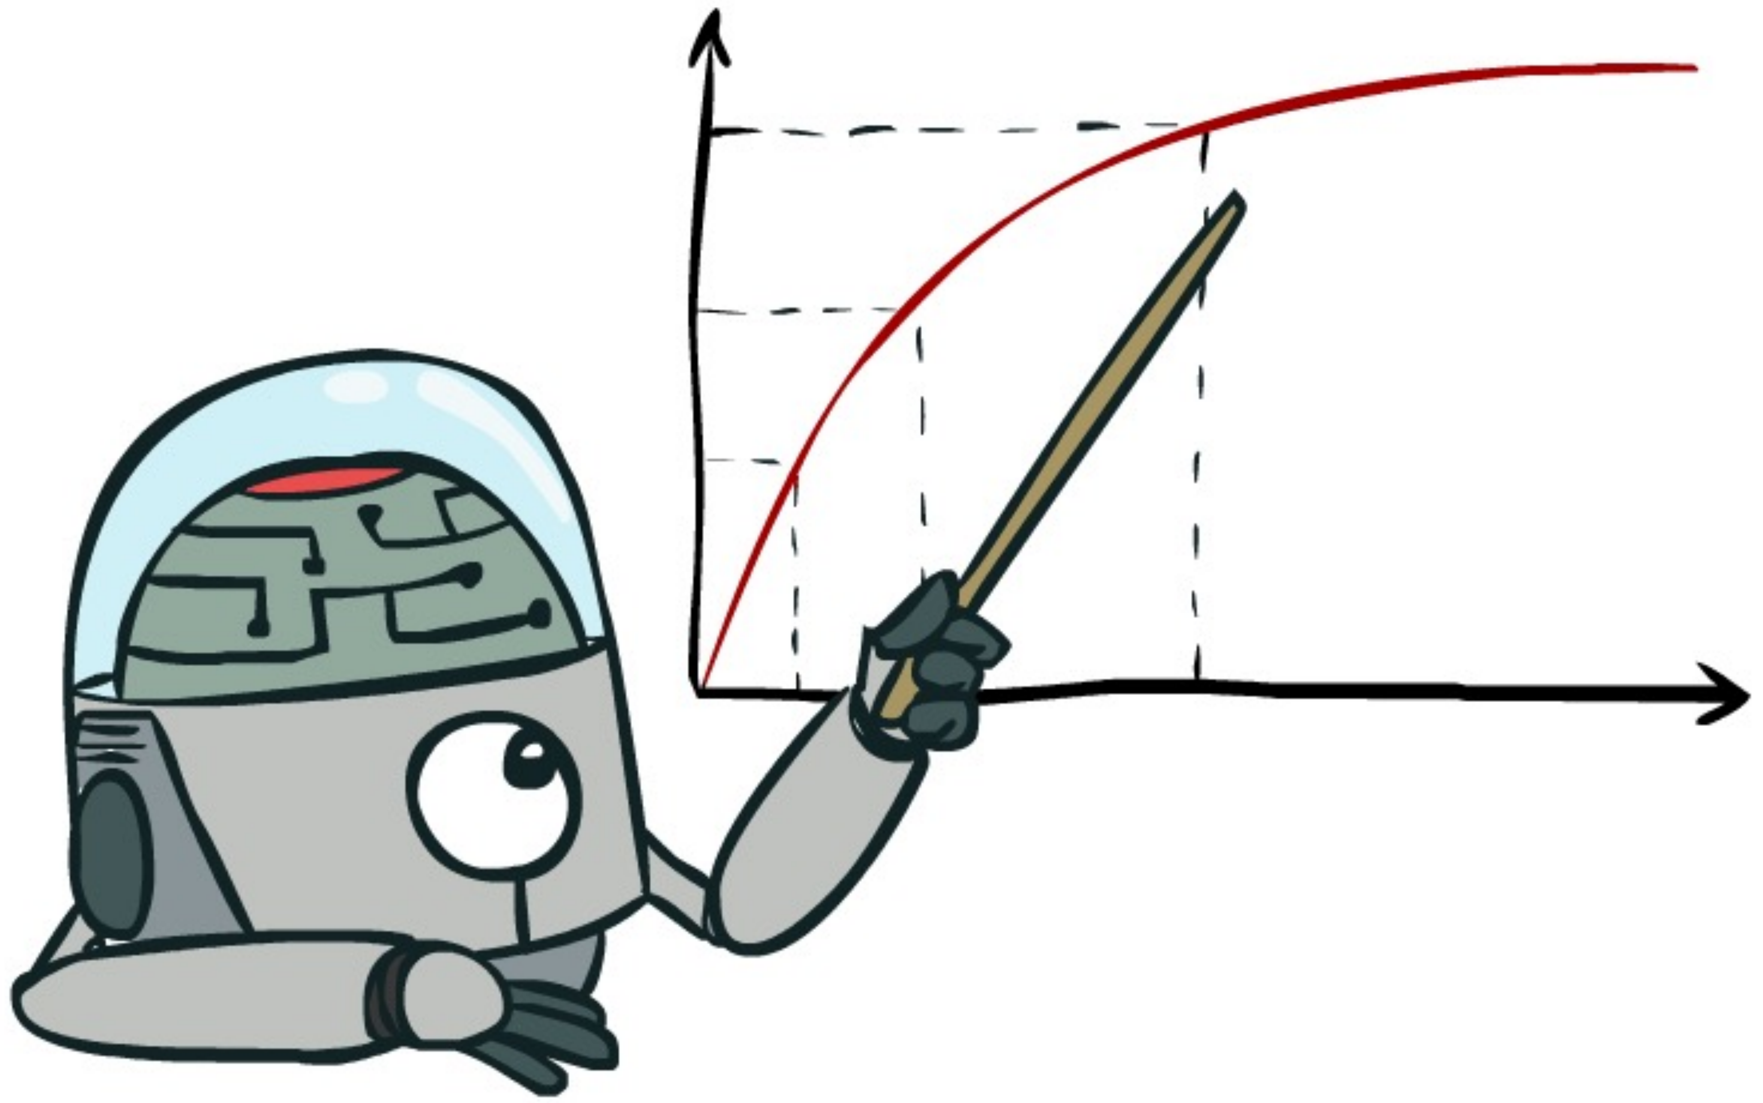

How Do AI Systems Maximize Utility?

Constraint satisfaction: searching intelligently for legal solutions

8			4		6			7
						4		
	1					6	5	
5		9		3		7	8	
				7				
	4	8		2		1		3
	5	2					9	
		1						
3			9		2			5

Example: Sudoku

Utility: Does the solution satisfy the rules / constraints?

Assumptions: We can write down the rules / constraints of the problem

How Do AI Systems Maximize Utility?

Planning: reasoning with models

Example: Robot navigation

Utility: Path length, collisions, surfaces, energy, social factors

Assumptions: We have a model of the world and the effects of the agent's actions

How Do AI Systems Maximize Utility?

Supervised Learning: learning from labeled examples

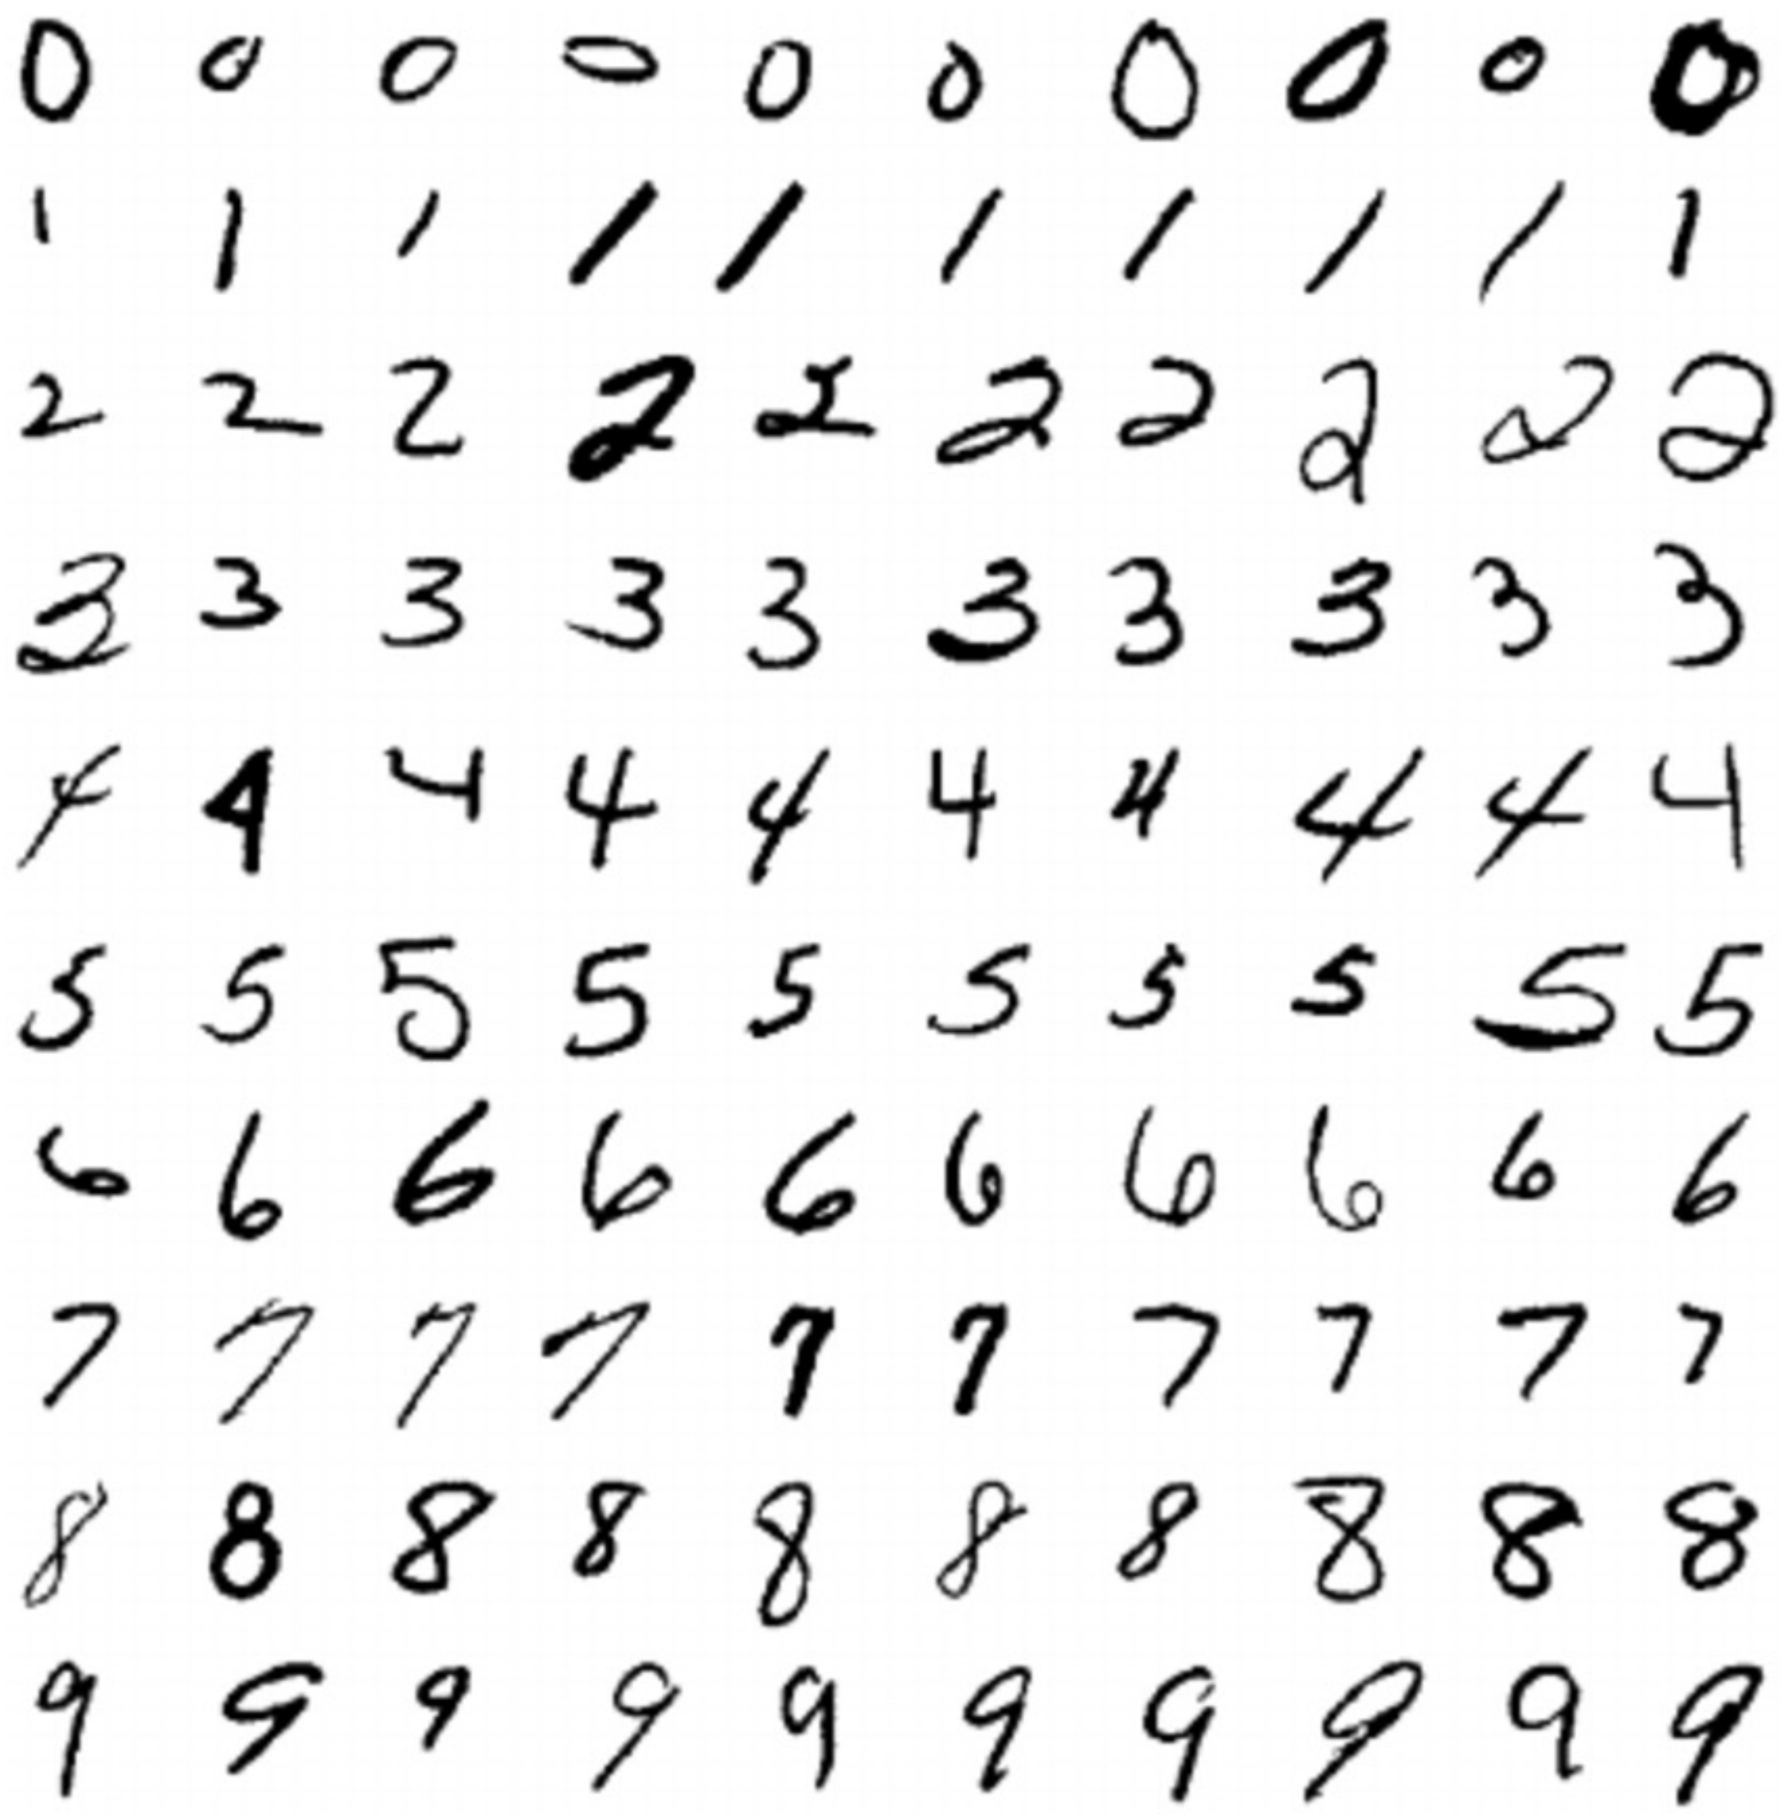

Example: Image classification

Utility: Classification accuracy on images not seen during training

Assumptions: We have access to a (usually large) labeled data set

How Do AI Systems Maximize Utility?

Unsupervised Learning: discovering patterns in unlabeled data

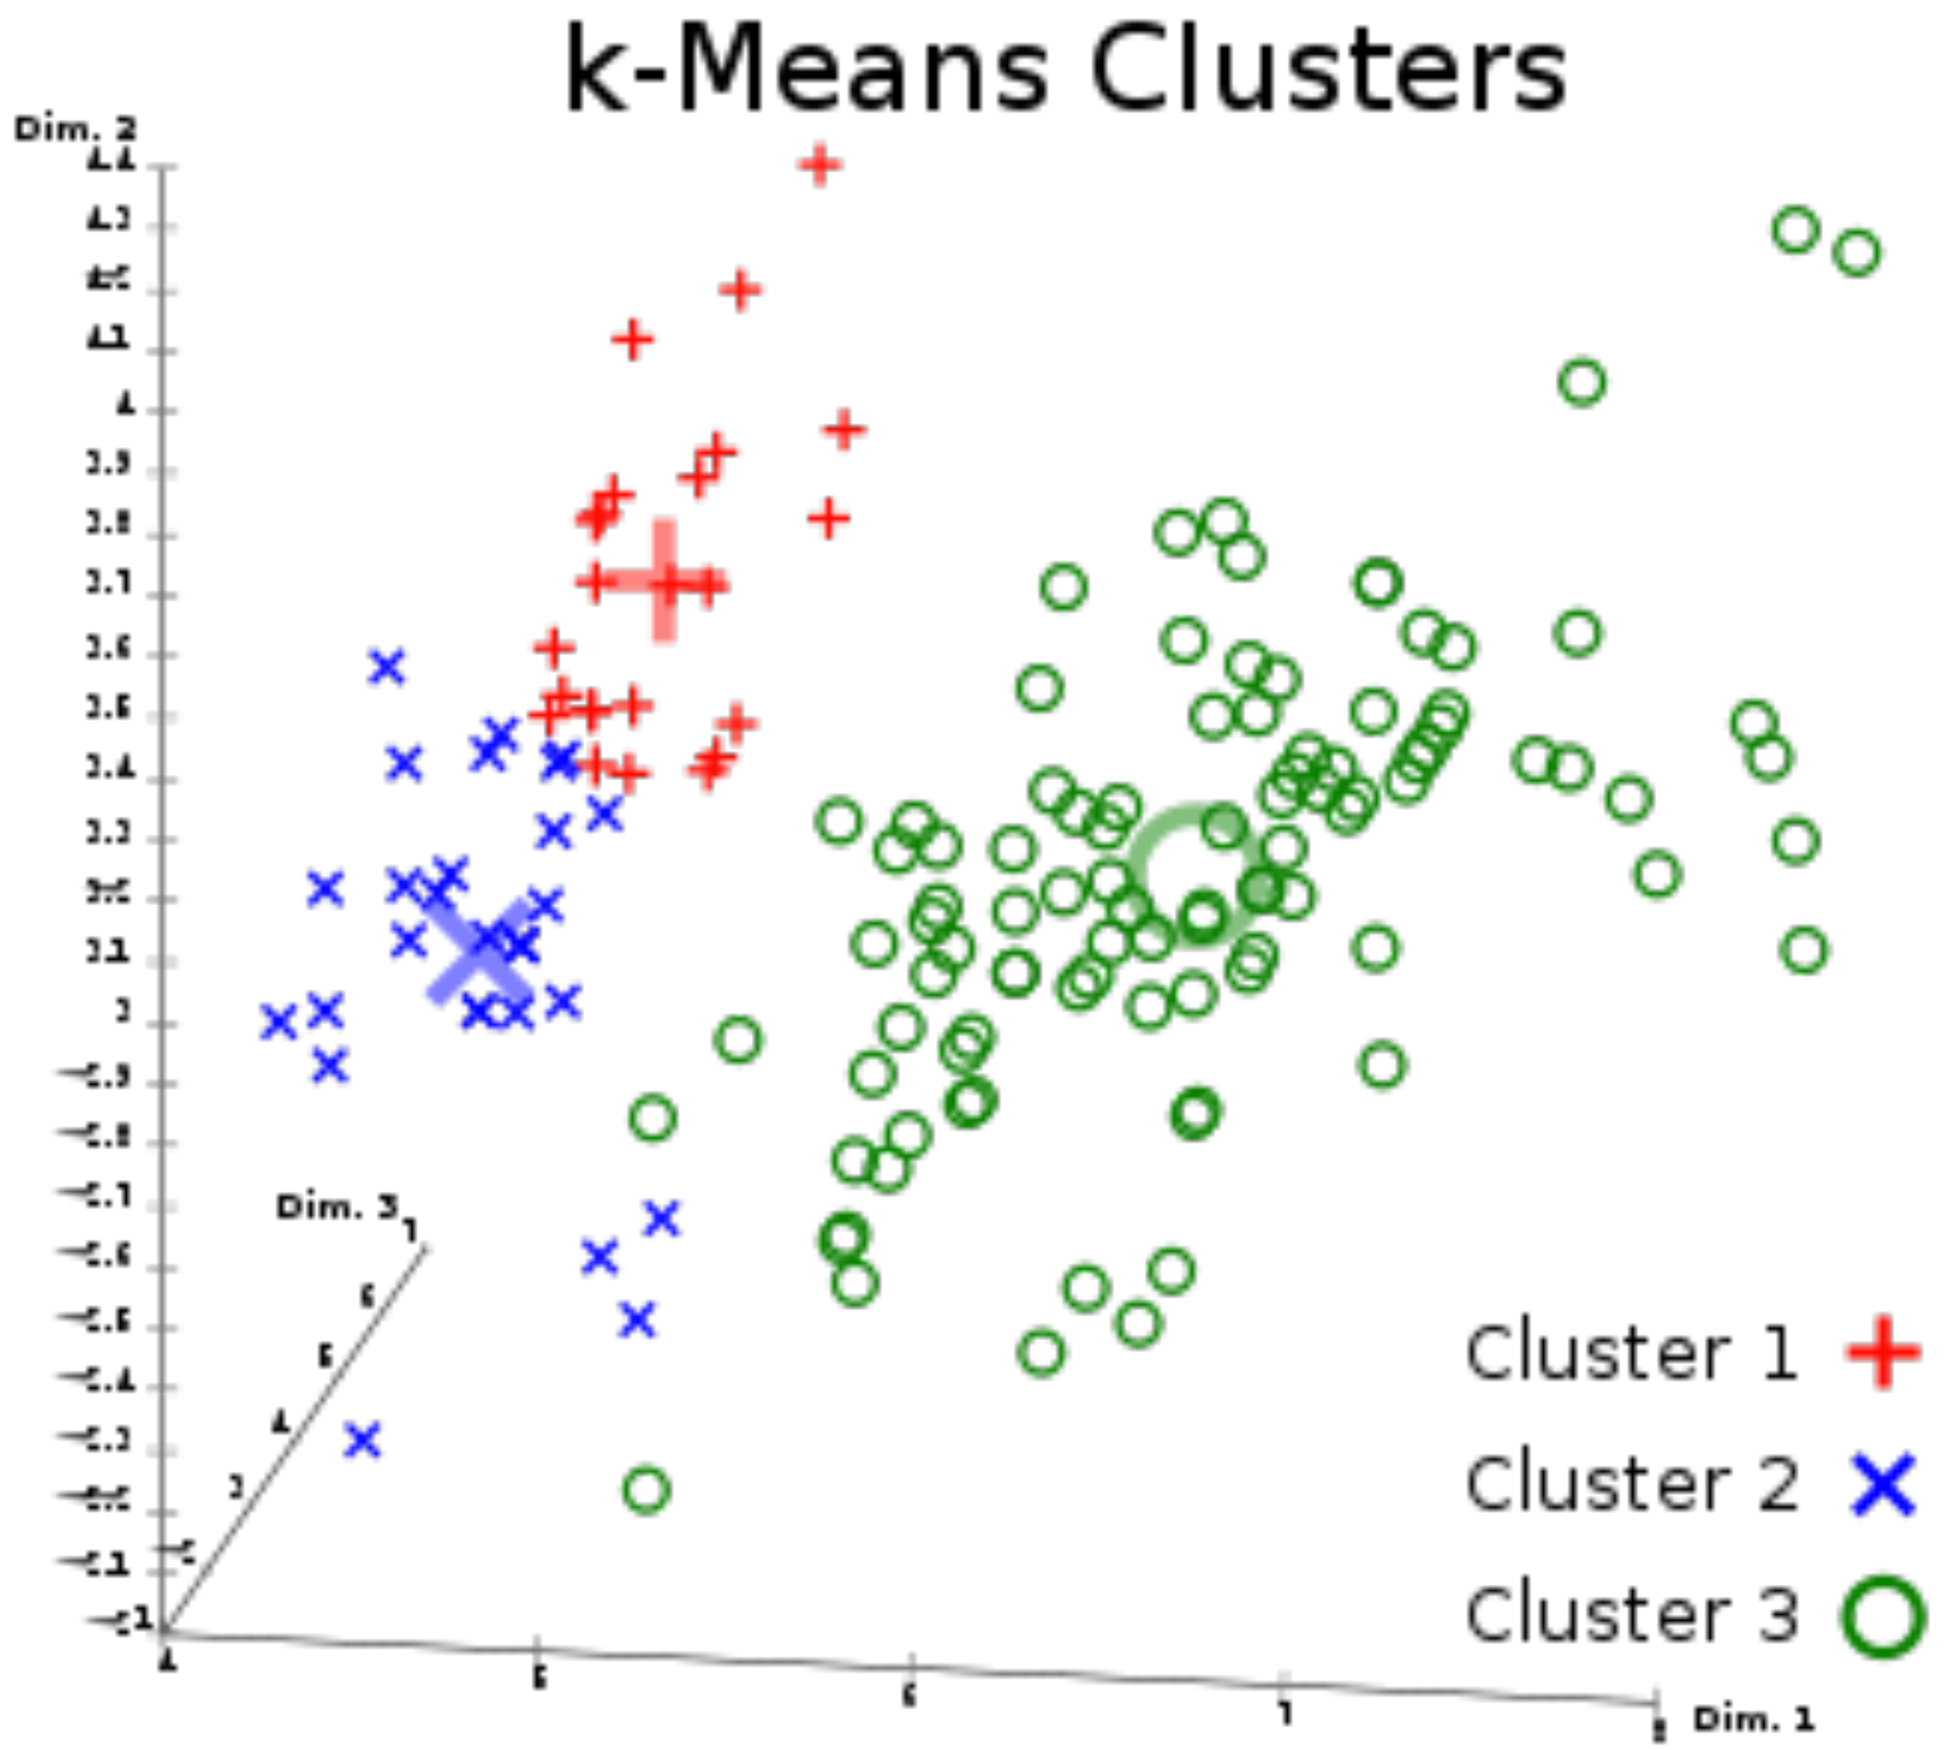

Example: Clustering species

Utility: "Best" explanation of data

Assumptions: Data points that should be clustered together are "close" together

How Do AI Systems Maximize Utility?

Reinforcement Learning: learning from experience

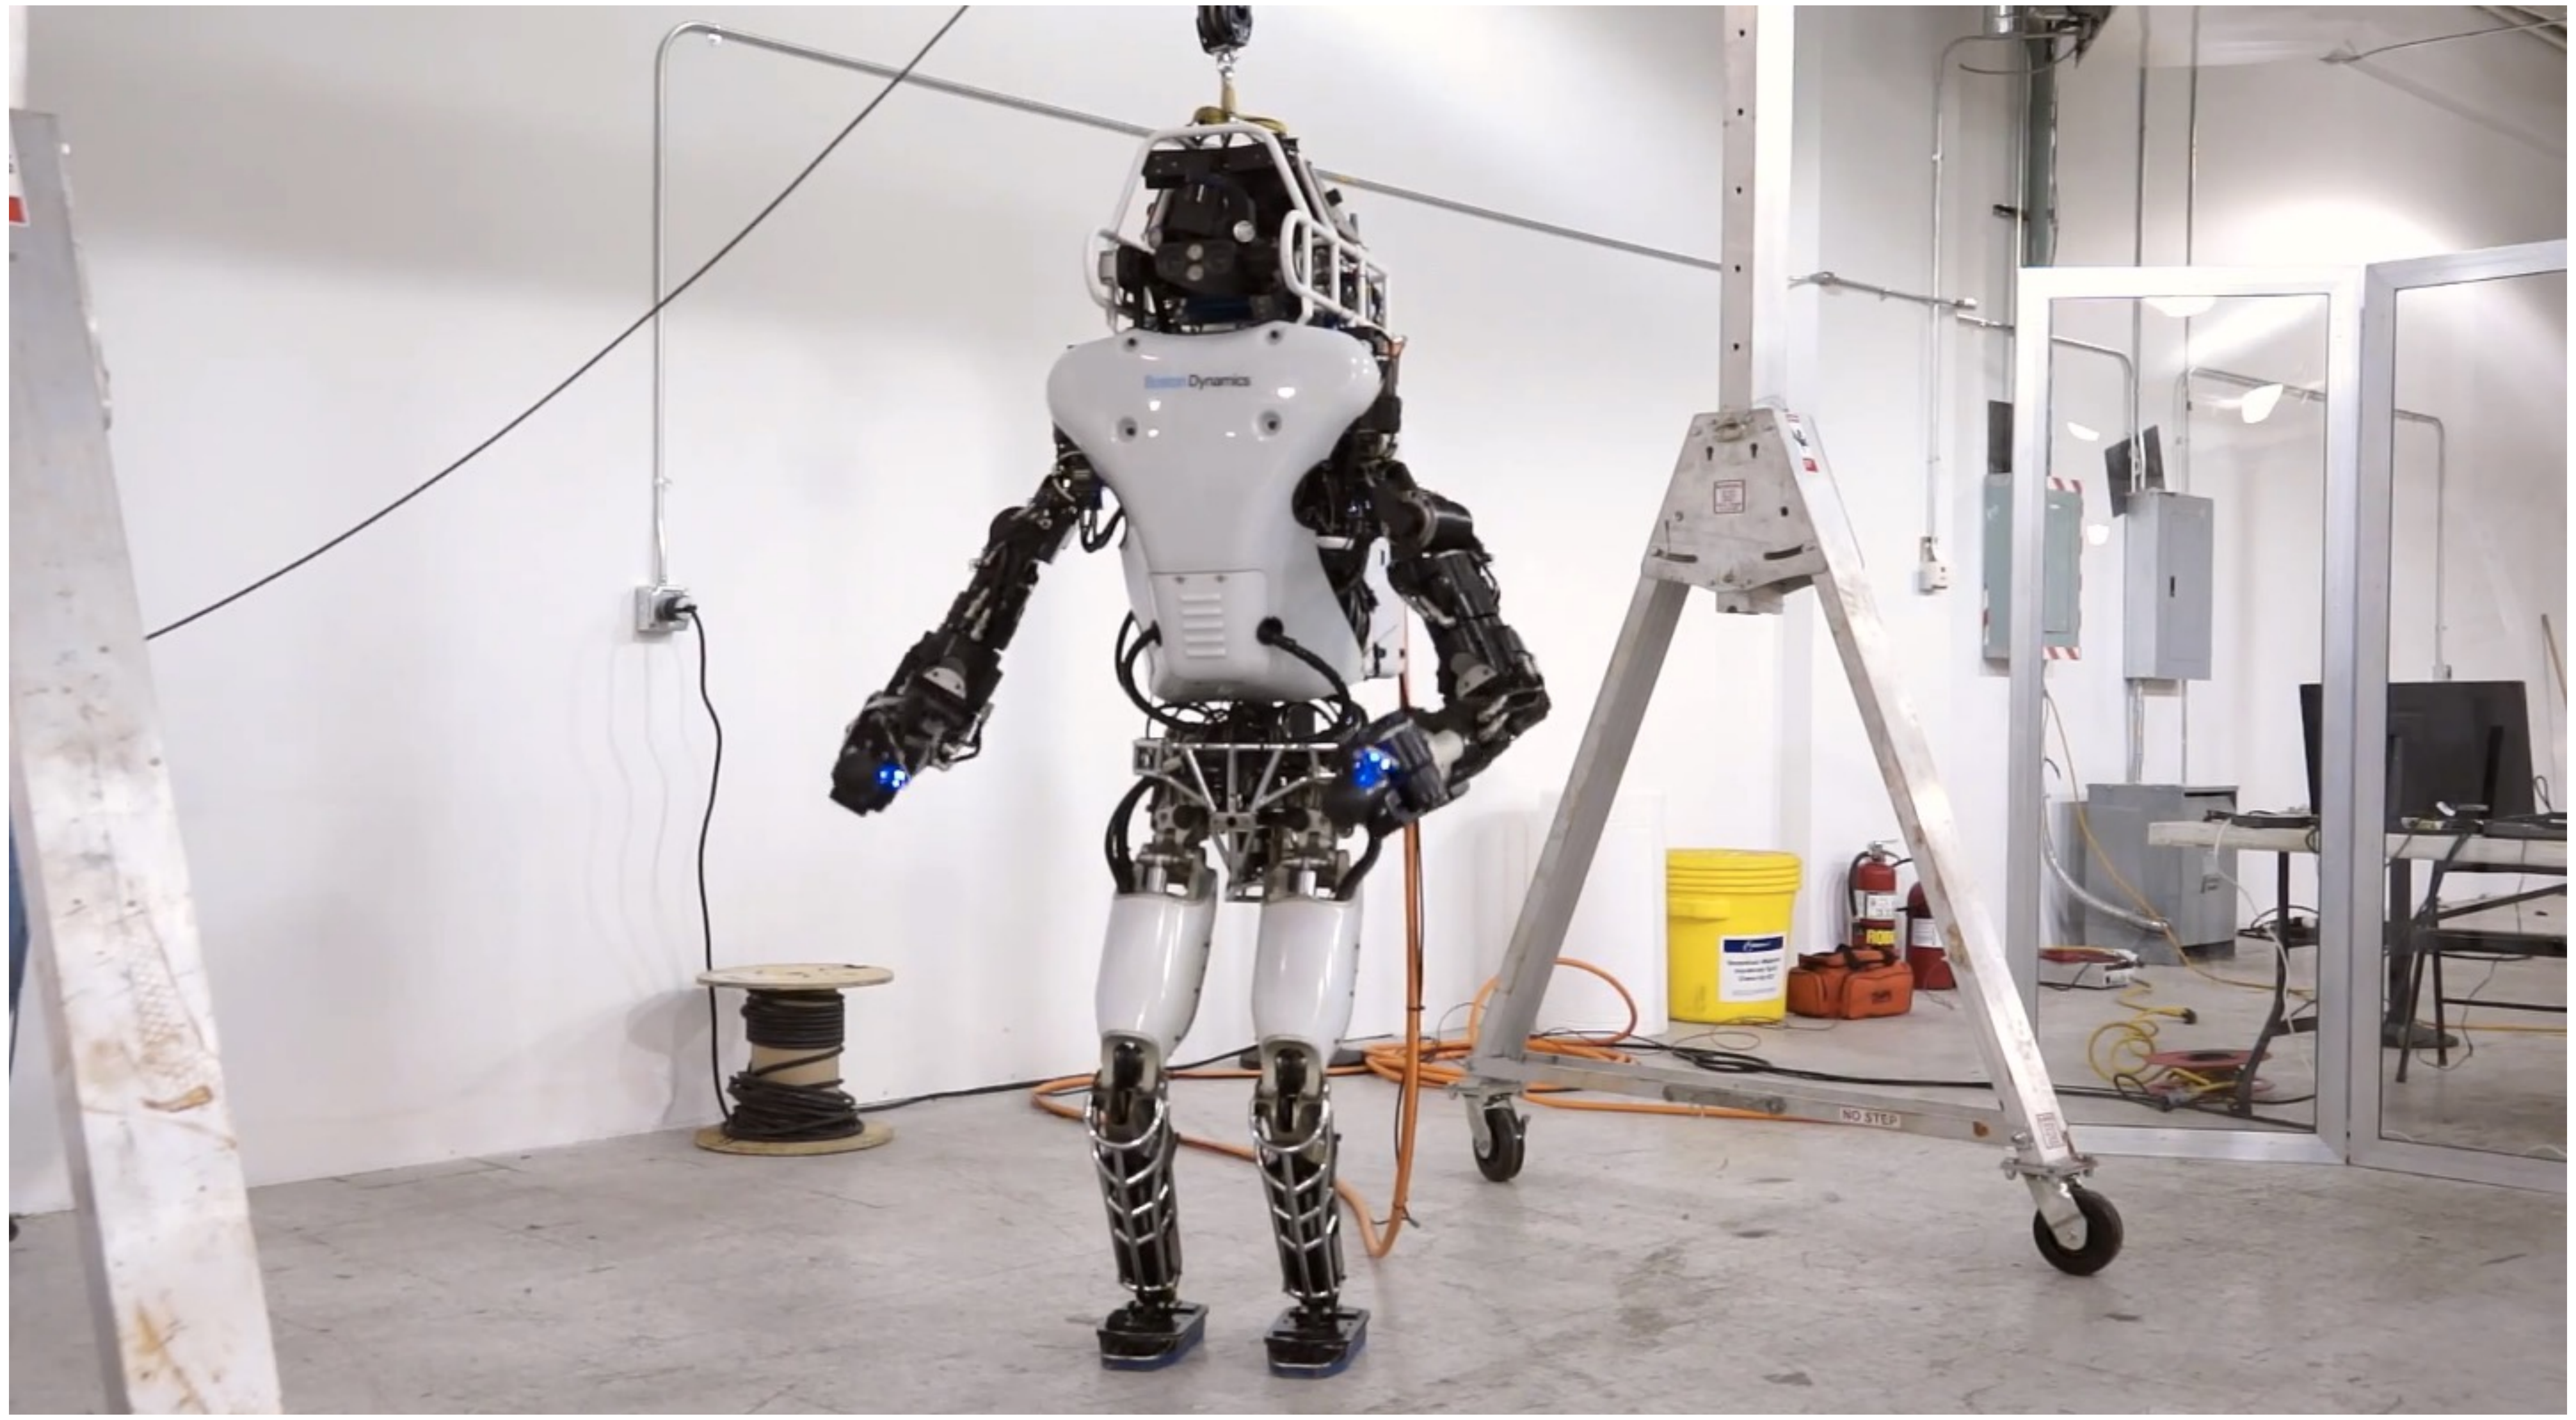

Example: Robot walking

Utility: Time until fall, speed, energy efficiency

Assumptions: We can "reward" and "punish" good and bad performance, but don't know what the correct action at each step should be

Properties of task environment

- Fully observable vs. partially observable
- Single-agent vs. multi-agent
- Deterministic vs. stochastic
- Episodic vs. sequential
- Static vs. dynamic
- Discrete vs. continuous
- Known vs. unknown

Single agent vs. multi-agent

- Not multi-agent if other agents can be considered part of the environment
- Only considered to be multi-agent if the agents are maximizing a performance metric that depends on other agents' behavior
- Single agent example: Pacman with randomly moving ghosts
- Multi-agent example: Pacman with ghosts that use a planner to follow him

Single / Multi Agent

Single	Multi
Uninformed Search	Minimax
A* Search	Expectimax
Local Search	MDPs
CSPs	Reinforcement Learning

Deterministic vs. stochastic

- Deterministic: next state of environment is completely determined by the current state and the action executed by the agent
- Stochastic: actions have probabilistic outcomes
- Strongly related to partial observability — most apparent stochasticity results from partial observation of a deterministic system
- Example: Coin flip

Determinism

Deterministic	Stochastic
Uninformed Search	Expectimax
A* Search	MDPs
Local Search	Reinforcement Learning
CSPs	Decision Diagrams
Minimax	

Fully observable vs. partially observable

- Fully observable: agent's sensors give it access to complete state of the environment at all times
- Can be partially observable due to noisy and inaccurate sensors, or because parts of the state are simply missing from the sensor data
- Example: Perfect GPS vs noisy pose estimation
- Example: IKEA assembly while blindfolded

Almost everything in the real world is partially observable

Observability

Fully Observable

Uninformed Search
A* Search
Local Search
CSPs

Minimax
Expectimax
MDPs
Reinforcement Learning

Partially Observable

POMDPs
Bayes Nets
HMMs
Decision Diagrams

Known vs. unknown

- Agent’s state of knowledge about the “rules of the game” / “laws of physics”
- Known environment: the outcomes for all actions are given
- Unknown: agent has to learn how it works to make good decisions
- Possible to be partially observable but known (solitaire)
- Possible to be fully observable but unknown (video game)

Model of the World

Known

Uninformed Search
A* Search
Local Search
CSPs
Classic Planning

Minimax
Expectimax
MDPs
Value Iteration
Decision Diagrams

Unknown

Q Learning
Learning parameters of Bayes Net

LARG's Research Highlights

Discussion Question

When will AI reach human-level intelligence?

- In the next 10 years
- In the next 50 years
- In the next 100 years
- Later or Never

Discussion Question

When will AI reach human-level intelligence?

- In the next 10 years
- In the next 50 years
- In the next 100 years
- Later or Never

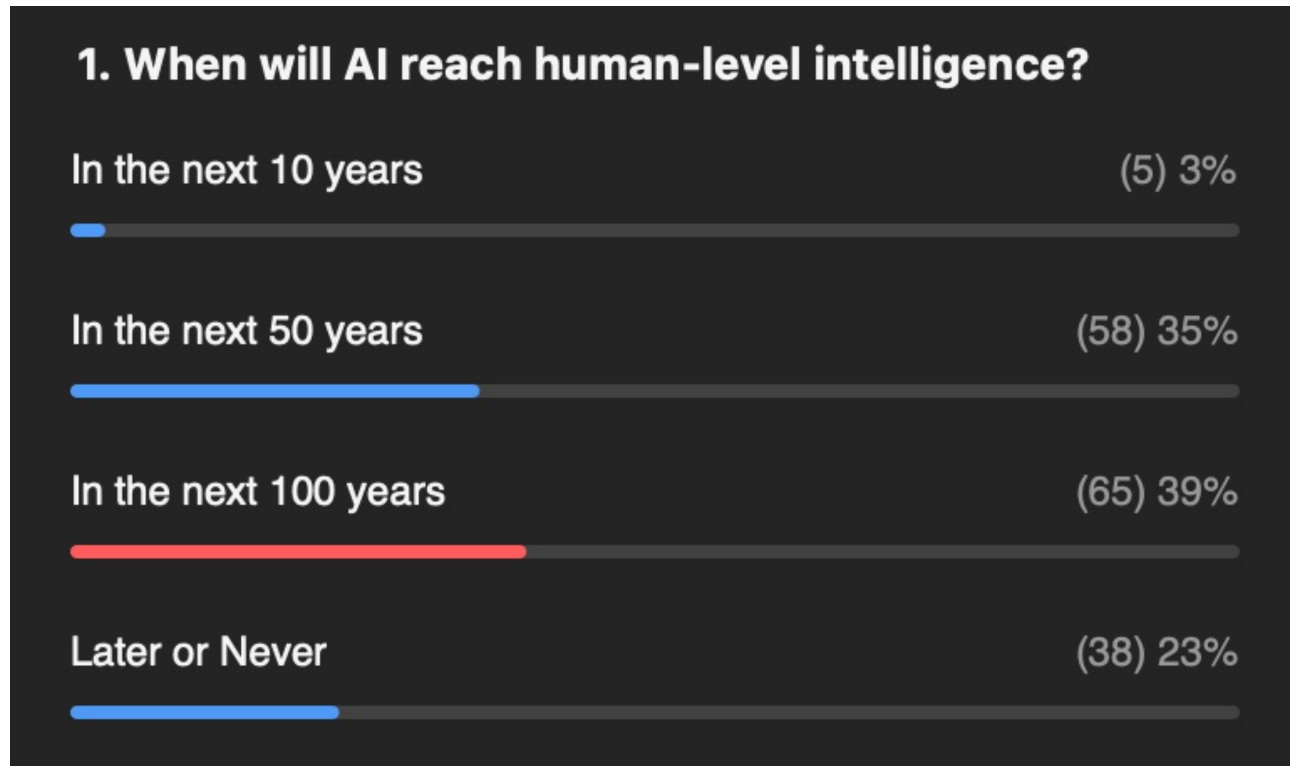

First class before discussion

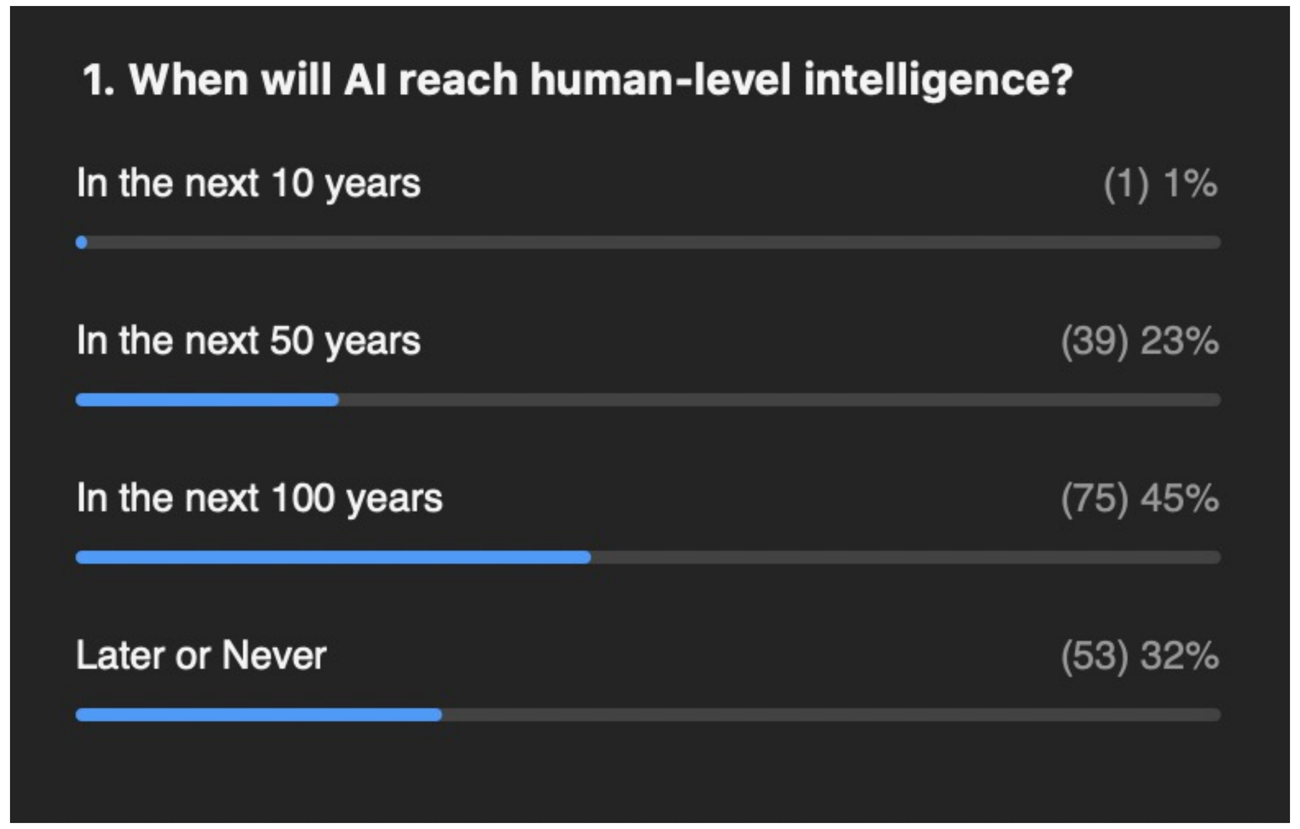

First class after discussion

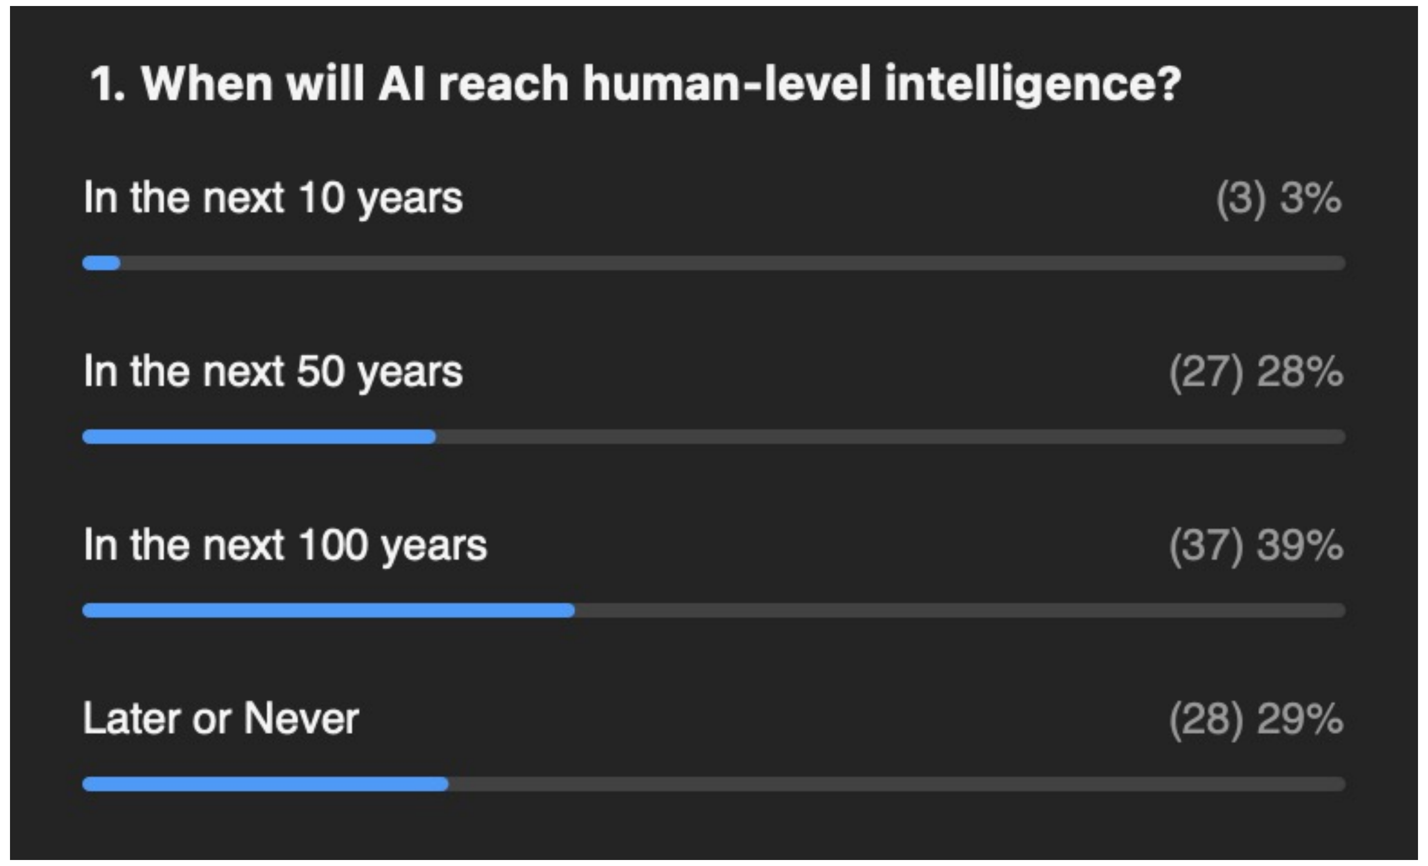

Discussion Question

Do you care if our "descendants" are human?

How do we develop AI to its fullest while keeping the human spirit unique to Homo Sapiens, or do we even care about this distinction? (Cole Kauppinen)

Final Remarks

- We have come a long way! Thank you!
- We are very proud that you have made it to the end of this demanding course, especially during this special time!
- We are impressed by your ingenuity and critical thinking in the reading responses, in-class discussions, and assignments!
- Thanks to Yifeng, Soroush, Zizhao, and Xiang for handling the course logistics.
- If this course helps you kickstart your future endeavors in AI, please email us and let us know!