



Introduction to Artificial Intelligence

Lecture 1 · Course 16:198:440 · Fall 2026

Kowndinya Boyalakuntla · Rutgers University

Teaching team



Kowndinya Boyalakuntla

Instructor

kowndinya.boyalakuntla@rutgers.edu

Office hours: Tuesday 3:00–4:00 PM
One Spring Street, 3rd floor



Doga Diren

Teaching Assistant

doga.diren@rutgers.edu

Office hours: TBD



Guoning Zhang

Teaching Assistant

guoning.zhang@rutgers.edu

Office hours: TBD

About me

- 3rd Year Ph.D. Student, Computer Science, Rutgers University advised by Prof. Abdeslam Boularias
- M.S. in Computer Science, Rutgers University
- B.Tech. in Computer Science & Engineering, IIT Tirupati advised by Prof. Sridhar Chimalakonda and Prof. Kalidas Yeturu
- Research Interests: Humanoid Robot Learning, Reinforcement Learning

Course logistics

Lectures & recitations

Lectures: Tuesday and Friday, 12:10–1:30 PM, SEC 118, Busch Campus.

Recitations (Fridays, Busch):

Sec 03 — 5:55–6:50 PM — SEC 205

Sec 04 — 2:15–3:10 PM — SEC 202

Sec 08 — 2:15–3:10 PM — SEC 208

Grading

Homework 50%

Midterm 25%

Final exam 25%

Both exams are closed-book. Some homework involves programming; Python is recommended.

Questions

Lecture, homework and exam questions go on Piazza, not e-mail:
piazza.com/rutgers/fall2026/cs440

Enrolment

No more Special Permission Numbers (SPNs).

Textbook: Russell & Norvig, Artificial Intelligence: A Modern Approach, 4th ed. (2021). Suggested: Burkov, The Hundred-Page Machine Learning Book; Bishop, Pattern Recognition and Machine Learning.

AI. AI. AI. AI. AI.



AI. AI. AI. AI. AI.

Artificial intelligence is everywhere.

But what exactly is it?

When companies say "AI," what are they actually talking about?



AI. AI. AI. AI. AI.



These are intelligent for completely different reasons.

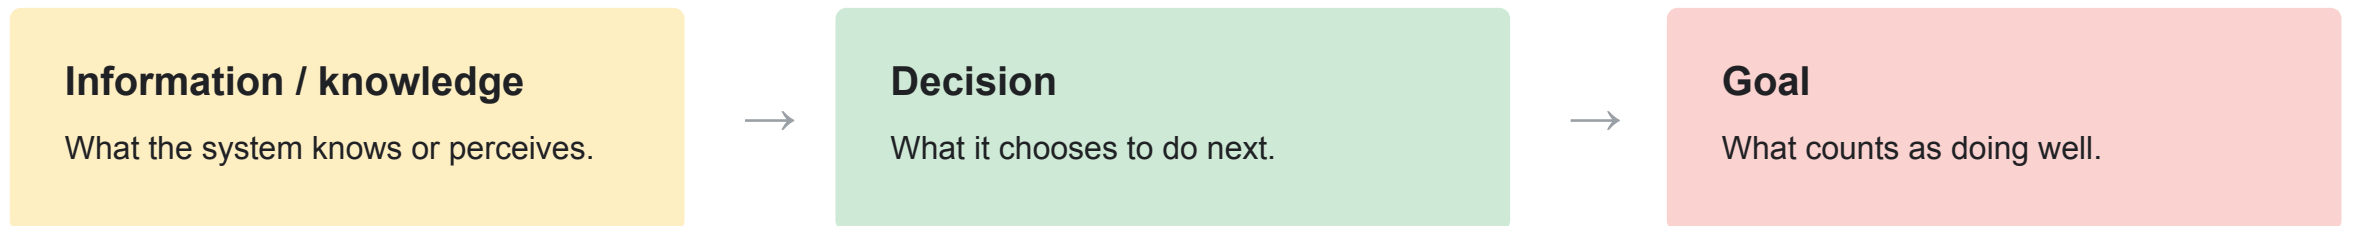
What is intelligence?

What is intelligence?

Intelligence is the ability to acquire and use information to make decisions that accomplish a goal.


What is intelligence?

Intelligence is the ability to acquire and use information to make decisions that accomplish a goal.




Intelligence is not simply having knowledge. The interesting part is using what you know to decide what to do.

Is a calculator intelligent?




Is a calculator intelligent?



- It receives information.
- It performs computation / reasoning.
- It produces the correct answer.
- So why should we not call it intelligent?

Is a calculator intelligent?



- It receives information.
- It performs computation / reasoning.
- It produces the correct answer.
- So why should we not call it intelligent?

Does it understand mathematics? Does it understand matter?

Is intelligence binary or a collection of capabilities?



Competence

How well can a system perform one particular task?
A calculator has extraordinary competence in a tiny domain.

Generality / adaptability

How broadly does that competence transfer to tasks it was never built for?

Let's talk about Chess



Chess is deterministic and solvable!! But, ..



White moves	1	20
White + Black	2	400
3 half-moves	3	8,902
2 full moves	4	197,281
5 half-moves	5	4,865,609
3 full moves	6	119,060,324
7 half-moves	7	3,195,901,860
4 full moves	8	84,998,978,956
9 half-moves	9	2,439,530,234,167

Legal board positions reachable from initial setup is between 10^{46} to 10^{47}

What about a chess engine?

Is Stockfish (strongest chess computer in the world) intelligent?

- It plays chess better than essentially every human being.
- Does it know that it is playing chess?
- Does it know what a king is?
- Does that matter, if its behaviour is extraordinarily competent?



What about a chess engine?

Is Stockfish (strongest chess computer in the world) intelligent?

- It plays chess better than essentially every human being.
- Does it know that it is playing chess?
- Does it know what a king is?
- Does that matter, if its behaviour is extraordinarily competent?



Intelligent behaviour $\neq$ necessarily human-like understanding

What about a chess engine?

Is Stockfish (strongest chess computer in the world) intelligent?

- It plays chess better than essentially every human being.
- Does it know that it is playing chess?
- Does it know what a king is?
- Does that matter, if its behaviour is extraordinarily competent?



Intelligent behaviour $\neq$ necessarily human-like understanding

Computer science gives us something much easier to measure than "true understanding": behaviour and performance. That is the move modern AI makes - fix a performance measure, then compare algorithms against it.

Does an AI need to understand what it is doing?

View 1 — Behavioural

If a machine consistently exhibits intelligent behaviour - solves problems, adapts appropriately, generalises within its domain - then calling it intelligent is useful.

Does an AI need to understand what it is doing?

View 1 - Behavioural

If a machine consistently exhibits intelligent behaviour - solves problems, adapts appropriately, generalises within its domain - then calling it intelligent is useful.

View 2 - Internal understanding

Perhaps genuine intelligence requires internal representations, grounding, semantic understanding, consciousness - something beyond observable behaviour.

Does an AI need to understand what it is doing?

View 1 — Behavioural

If a machine consistently exhibits intelligent behaviour — solves problems, adapts appropriately, generalises within its domain — then calling it intelligent is useful.

View 2 — Internal understanding

Perhaps genuine intelligence requires internal representations, grounding, semantic understanding, consciousness — something beyond observable behaviour.

There is no universally accepted answer.

Does an AI need to understand what it is doing?

View 1 — Behavioural

If a machine consistently exhibits intelligent behaviour — solves problems, adapts appropriately, generalises within its domain — then calling it intelligent is useful.

View 2 — Internal understanding

Perhaps genuine intelligence requires internal representations, grounding, semantic understanding, consciousness — something beyond observable behaviour.

There is no universally accepted answer.

A modern language model can solve many tasks, explain concepts, write programs, summarise papers and plan actions.

Does successful behaviour imply understanding?



Artificial Intelligence

Artificial Intelligence is the study and engineering of computational systems that perceive information, reason or learn from it, and choose actions to accomplish goals.

PERCEIVE



THINK / LEARN



ACT

Artificial

Implemented by a machine.

Intelligence

Using information to make effective decisions.

This definition is deliberately broad. By the end of the semester, I want you to be able to criticise it.

Where is the intelligence in a self-driving car?



Where is the intelligence in a self-driving car?



Camera?

Object detector?

Trajectory planner?

Neural network?

Controller?

The whole system?

There may be no single magical "AI algorithm".

Sensors

Perception

World model


Prediction

Planning

Control

Intelligence often emerges from a system of components, each solving a different problem.

An Intelligent Agent



An agent chooses its actions based only on its history of observations and actions — mathematically, a function from that history to an action.

What can be an agent?

Robot

Human


Chess engine

Recommender

Self-driving car

LLM with tools

An Intelligent Agent



An agent chooses its actions based only on its history of observations and actions - mathematically, a function from that history to an action.

What can be an agent?

Robot

Human

Chess engine

Recommender

Self-driving car

LLM with tools

Much of AI asks one question: given what I know right now, what should I do next?

Rational behaviour

Doing the "right" thing requires defining what "right" means.

Rational behaviour

Doing the "right" thing requires defining what "right" means.

Rational agent. An agent that chooses actions expected to produce the best outcome according to some performance measure.

Chess	Win probability
Navigation	Path length + safety + time
Recommendation	Engagement? Satisfaction? Purchases?
Robot	Task success – energy – collisions – time
LLM	Helpfulness? Correctness? Instruction following? Safety?

One idea connects almost all of AI

We define an objective and try to make good decisions.

Search

Which actions reach the goal?

Game playing

Which action is best against an opponent?

Probabilistic reasoning

What should I believe given partial evidence?

Planning / MDPs

Which actions maximise long-term utility?

Machine learning

Which function best explains the data?

Reinforcement learning


Which policy maximises long-term reward?

Generative AI

Which output fits the context and instructions?

The algorithms change. The underlying questions recur.

Why can't AlphaGo drive a car?



**AlphaGo defeated world champions like Lee Sedol.
Why can the same AI not simply drive a car?**

Go

Precisely defined world:

- State space
- Action space
- Fixed rules
- One clear objective

Driving

Completely different:

- Observations
- Actions
- Dynamics
- Uncertainty
- Competing objectives

Intelligence in one domain does not automatically transfer to another.

Narrow AI vs. General Intelligence

Narrow / specialised

Designed or trained for one particular family of tasks.

Chess engine · Protein-structure prediction · Object detector · Autonomous-driving stack

General

Acquires and transfers capabilities across substantially different tasks without being rebuilt for each one.

What would it take for a machine to learn a completely new task?

Memory

Learning

Reasoning

Abstraction

Transfer

Planning

Language

World models

Exploration

What is AGI?

There is no universally agreed technical definition of AGI.

One useful intuition: a generally intelligent system should be able to learn, reason, adapt and transfer knowledge across many different tasks and environments.

Driving system → chess?

Chess system → cooking?

Language model → robot manipulation?

Humans transfer knowledge constantly. What representations make transfer possible?

A genuinely open research question - not a settled definition to memorise.

Humans: one system, many capabilities



Perception

What is happening?



Reasoning / memory / learning

What do I know? What might happen? What should I do?



Action / control

Execute the decision.

Humans integrate perception, reasoning, learning and action remarkably well. That integration is what general machine intelligence is chasing.

Modern robotics



Vision

What am I looking at?

Language

What is the task?

World knowledge

How do objects
behave?

Planning

What sequence of
actions?

Control

How do I execute it?

Robotics exposes a central challenge of AI: intelligence has to survive contact with the real world.

Why AI requires different algorithms

If the world is

Known + deterministic

→ Search / planning may be enough

If the world is

Adversarial

→ Game-tree search

If the world is

Uncertain

→ Probability

If the world is

Sequential + stochastic

→ Markov decision processes

If the world is

Unknown

→ Learning

If the world is

Too complex to model explicitly

→ Machine learning, RL, approximate planning

Different AI algorithms are answers to different assumptions about the world.

The course in three questions

I.

What should I do when I understand the world?

SEARCH & PLANNING

BFS · DFS · A*

Heuristics

Adversarial search

Constraint satisfaction

II.

What should I believe and do when the world is uncertain?

PROBABILISTIC AI

Bayes

Bayesian networks

Temporal models · HMMs

Kalman filters · MDPs

III.

What if I don't know the solution and must learn it?

MACHINE LEARNING

Regression · Classification

Neural networks

Representation learning

RL · Perception · Foundation models

Search is still AI

If modern AI is so powerful, why are we learning search?



Because intelligence still requires choosing among possible futures.

Navigation

Robot planning

Game playing

Scheduling

Program synthesis

Theorem proving

Tool-use planning

Inference in modern AI

Instead of knowing the answer, systematically explore possible answers.

How did we get here?

