

# What is "Artificial Intelligence"?

AI is anything that computers can't do yet

couldn't do 30 years ago

- \* AI is anything computers can't do yet.
  - \* Actually: anything they couldn't do 30 years ago.
- \* Survey; covers lots of stuff.
- \* This is a practical course: algorithms for many loosely-related tasks considered "AI".
- \* The course is not:
  - \* about ethics/philosophy
  - \* just about ML

# What is "Artificial Intelligence"?

**1) Systems that think like humans**

**2) Systems that think rationally**

**3) Systems that act like humans**

**4) Systems that act rationally**

1) Think like humans

\* 60s: collab between CS, psychology and neuroscience.

\* Problem: Our understanding is far from understanding or replicating the brain.

2) Think rationally

\* Logic

\* Problem: Humans don't always follow logical deliberation. Most practical problems are NP-hard.

3) Act like humans

\* Turing test (slide)

\* Non-problems:

\* Not useful: Ask the system to do something useful that an employee would be able to do.

\* "Surface-level imitation": Interrogate longer, ask a complex question.

\* Problem: Hard to do; not necessary for useful AI.

4) Act rationally / correctly


\* Doesn't necessarily need to involve thinking. Ex: Blinking.

\* We focus on this definition.

Strong AI: A system that can do everything a human can do

Weak AI: A system that can do \*something\* better or as well as a human can.

# Turing test



- Interrogator  $\leftrightarrow$  human; computer
- Question: Strong or weak? Which of the four? Strong. Type 3 = act like humans.

## History

- 60s-70s: Hype and promise of AI.
- 80s-90s: "AI Winter".
- 2000-today: Second wave of AI.

### 60s-70s:

- Focused on type-1 (think like humans); strong AI.
- Human-like neural networks.
- Expert systems: type-2 (think rationally)

### 2000-today:

- Focusing on type-4 (act rationally); weak AI.
- Probabilistic models, machine learning, neural networks (again)

## Course syllabus

- Search
  - Depth-first, breadth-first, informed
  - Games and adversarial search
  - Constraint satisfaction problems
  - Propositional logic
- Probability
  - Bayesian networks
  - Temporal models
- Machine learning
  - Decision trees
  - Neural networks
  - Training and evaluation

\* Survey; covers lots of stuff.

\* This is a practical course: algorithms for many loosely-related tasks considered "AI".

\* The course is not:

\* about ethics/philosophy

\* just about ML

## State of the art

- Which of the following can be done now?

**Yes** • Play a decent game of table tennis.

**Yes** • Drive safely along a curving mountain road.

**Maybe** • Drive safely along Hastings Street.

**Yes** • Buy a week's worth of groceries online.

**No** • Buy a week's worth of groceries at Nester's.

**Yes** • Play a decent game of bridge.

**Maybe** • Discover and prove a new mathematical theorem.

**Maybe** • Design and execute a research program in molecular biology.

**Yes** • Give competent legal advice in a specialized area of law.

**Yes** • Translate spoken English to spoken Swedish in real time.

**No** • Converse successfully with another person for an hour.

**Maybe** • Perform a complex surgical operation.

**No** • Unload a dishwasher and put everything away.