IS396 ***(Selected topics in Software Engineering "Data Science")***

… & …

CS361 ***(Artificial Intelligence)***

Lecture 1
Introduction Artificial Intelligence & Learning

Amr S. Ghoneim
(Assistant Professor, Computer Science Dept.)

Helwan University
Fall 2017

Lecture is based on its counterparts in the following courses:

- *Intelligent Systems*, **University of British Columbia** (Dept. of Computer Science)
- *Introduction to Artificial Intelligence*, **University of Wisconsin-Madison**
- *Artificial Intelligence*, **University of Illinois at Urbana-Champaign**
- *Artificial Intelligence*, **University of California, Berkeley**

Resources for this lecture

- **This lecture covers the following chapters:**
  - **Chapter 1 (Introduction)** *from* **Stuart J. Russell and Peter Norvig, "Artificial Intelligence: A Modern Approach," Third Edition (2010), by Pearson Education Inc.**

  … AND …

  - **Chapter 1 (AI: History and Applications)** *from* **George F. Luger, "Artificial Intelligence: Structures and strategies for complex problem solving, " Fifth edition (2005), Pearson Education Limited.**

Outline

- What is Intelligence?
- What is Artificial Intelligence?
  - Systems that act like humans
    - Turing Test; the Imitation Game …
    - The Chinese room argument
    - Strong vs. Weak AI Hypotheses
  - Systems That Think Like Humans
  - Systems that Think Rationally
  - Systems that Act Rationally
    - AI as Study and Design of Intelligent Agents
    - Intelligent Agents
    - Pac-Man as an Intelligent Agent
    - Intelligent Agents in the World
- Learning
  - Learning from Labelled Examples
  - Supervised Machine Learning: Task Overview
    - Standard Approach for Constructing a Machine Learning Dataset for a Task
    - Concept Learning
  - Generic Machine Learning Models

Some Foundations of Artificial Intelligence

Philosophy
- Can formal rules be used to draw valid conclusions?
- How does the mind arise from a physical brain?
- Where does knowledge come from?
- How does knowledge lead to action?

Mathematics
- What are the formal rules to draw valid conclusions?
- What can be computed?
- How do we reason with uncertain information?

Neuroscience
- How do brains process information?

Psychology
- How do humans and animals think and act?

Economics
- How should we make decisions so as to maximize payoff?
- How should we do this when others may not go along?
- How should we do this when the payoff may be far in the future?

Computer engineering
- How can we build an efficient computer?

Control theory and cybernetics
- How can artifacts operate under their own control?

Linguistics
- How does language relate to thought?

What is Intelligence?

Intelligence:

- **Judgment, otherwise called "good sense," "practical sense," "initiative," the faculty of adapting one's self to circumstances ... auto-critique** ~ *Alfred Binet (July 8, 1857 – October 18, 1911) was a French psychologist who invented the first practical intelligence test (An intelligence quotient (IQ); a total score derived from one of several standardized tests designed to assess human intelligence)*
- **"...the resultant of the process of acquiring, storing in memory, retrieving, combining, comparing, and using in new contexts information and conceptual skills."** *~Lloyd G. Humphreys (December 12, 1913 – September 7, 2003) was an American psychologist*
- **"the capacity to learn and solve problems"** (Websters dictionary)
  - in particular,
    - *the ability to solve novel problems*
    - *the ability to act rationally*
    - *the ability to act like humans*

What is Artificial Intelligence?

Thinking humanly

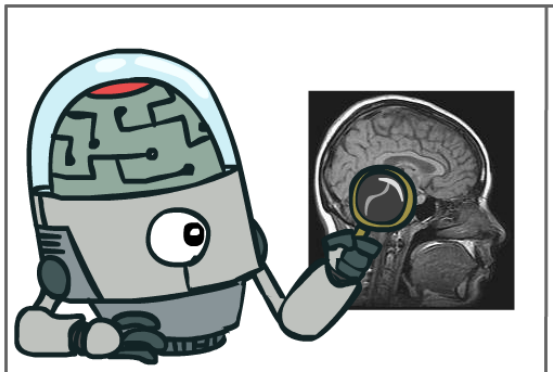

Acting humanly

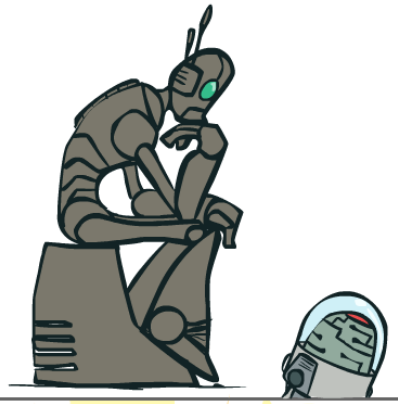

Thinking rationally

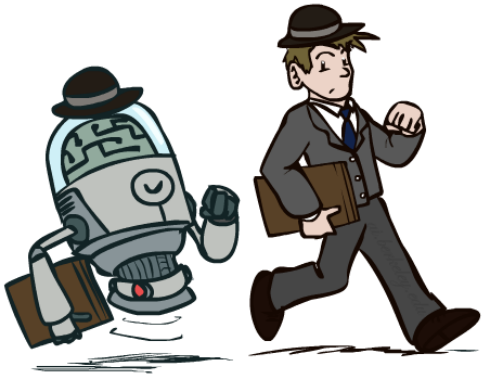

Acting rationally

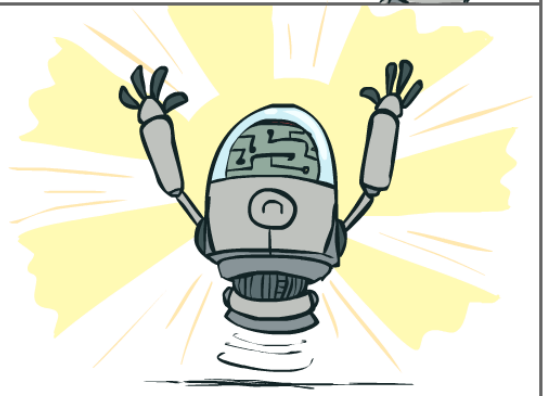

What is Artificial Intelligence?

Systems that act like humans	Systems that think rationally
"The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991) "The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)	"The study of mental faculties through the use of computational models" (Charniack and McDermott, 1985). "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Systems that think like humans	**Systems that act rationally**
"The automation of activities that we associate with human thinking, such as decision making, problem solving, learning" (Bellman, 1978) "The exciting new effort to make computers think … machines with minds, in the full and literal sense." (Haugeland, 1985)	"AI .. is concerned with intelligent behavior in artifacts (Nilsson, 1998) "Computational Intelligence is the study of the design of intelligent agents." (Poole et al., 1998)

What is Artificial Intelligence?

Systems that act like humans	Systems that think rationally
"The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991) "The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)	"The study of mental faculties through the use of computational models" (Charniack and McDermott, 1985). "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Systems that think like humans	**Systems that act rationally**
"The automation of activities that we associate with human thinking, such as decision making, problem solving, learning" (Bellman, 1978) "The exciting new effort to make computers think … machines with minds, in the full and literal sense." (Haugeland, 1985)	"AI .. is concerned with intelligent behavior in artifacts (Nilsson, 1998) "Computational Intelligence is the study of the design of intelligent agents." (Poole et al., 1998)

Systems that act like humans … ?!

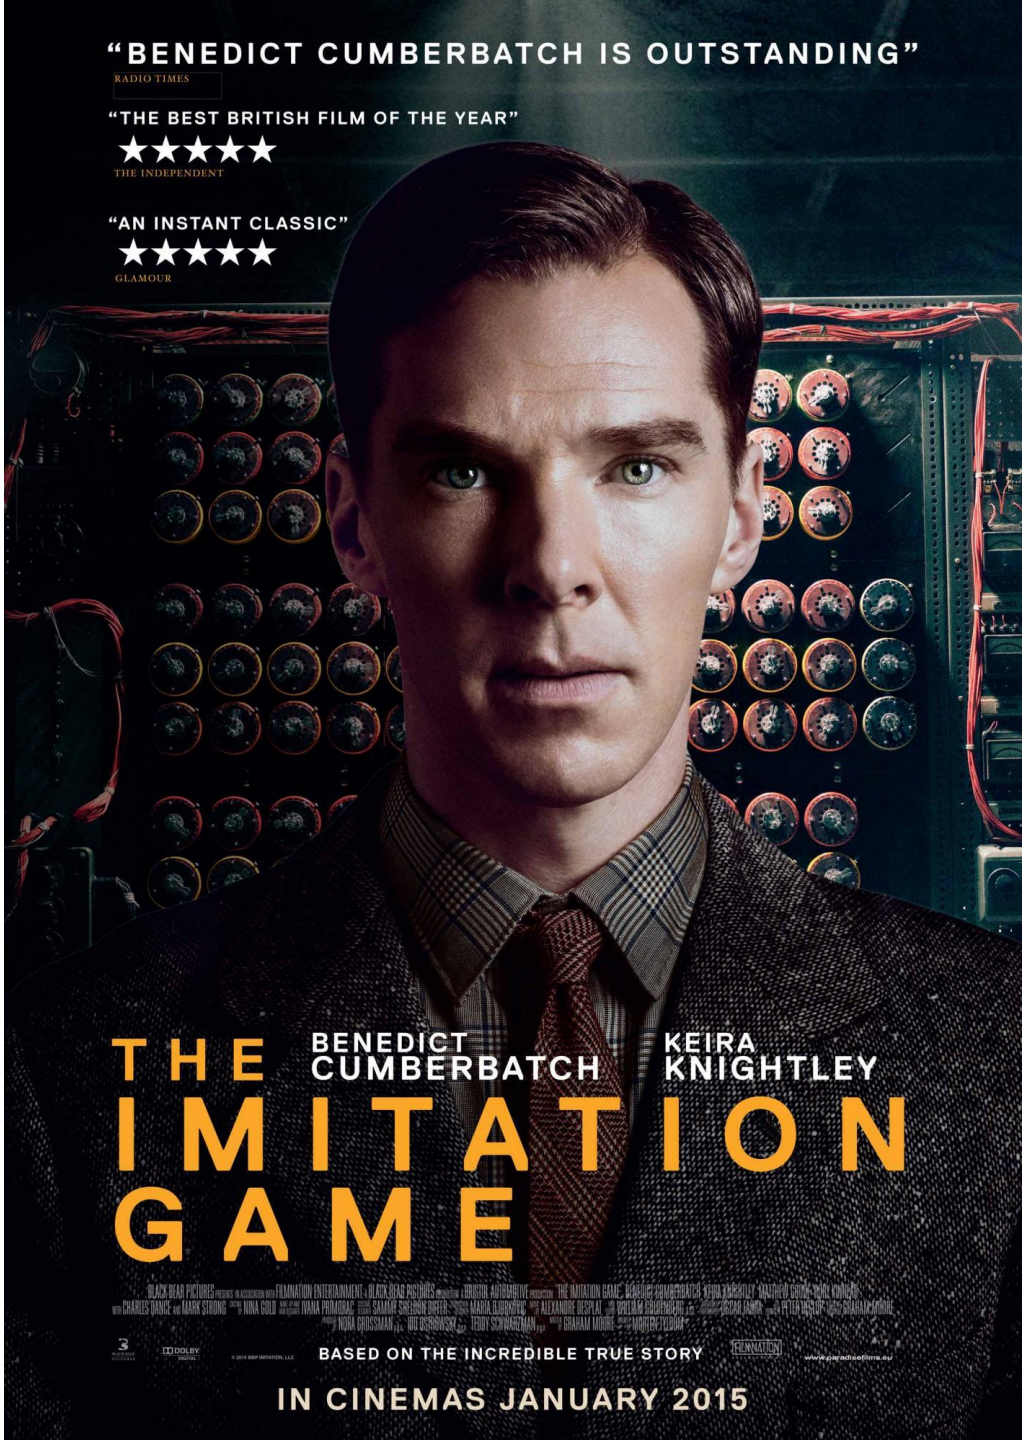

Systems that act like humans

Turing Test; *the Imitation Game …*

Turing (1950) "Computing machinery and intelligence":

- "Can machines think ?" −→ "Can machines behave intelligently ?"
- Operational test for intelligent behavior: the Imitation Game

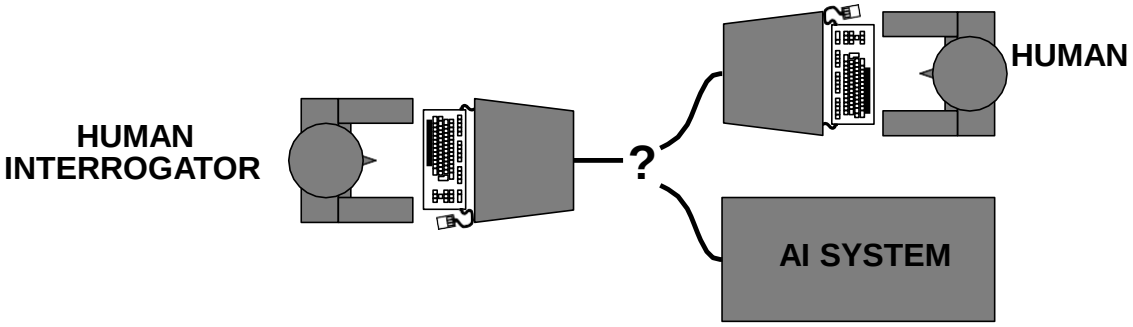

Systems that act like humans

Turing Test; *the Imitation Game …*

- **Turing test (1950)**: Can a human interrogator tell whether (written) responses to her (written) questions come from a human or a machine?
  - Natural Language Processing
  - Knowledge Representation
  - Automated Reasoning
  - Machine Learning
- **Total Turing Test** *(extended to include physical aspects of human behavior):*
  - Computer Vision
  - Robotic

But why do we want an intelligent system to act like a human?

- Because for many tasks, humans are still the Gold Standard.

Systems that act like humans

The Chinese room argument

John Rogers Searle (born July 31, 1932) is an American philosopher

“Searle's thought experiment begins with this hypothetical premise: suppose that artificial intelligence research has succeeded in constructing a computer that behaves as if it understands Chinese. It takes Chinese characters as input and, by following the instructions of a computer program, produces other Chinese characters, which it presents as output. Suppose, says Searle, that this computer performs its task so convincingly that it comfortably passes the Turing test: it convinces a human Chinese speaker that the program is itself a live Chinese speaker. To all of the questions that the person asks, it makes appropriate responses, such that any Chinese speaker would be convinced that they are talking to another Chinese-speaking human being.”

The question Searle wants to answer is this: does the machine *literally* "understand" Chinese? Or is it merely *simulating* the ability to understand Chinese? Searle calls the first position "**strong AI**" and the latter "**weak AI**".

Systems that act like humans

The Chinese room argument

(Continued)

Searle then supposes that he is in a closed room and has a book with an English version of the computer program, along with sufficient paper, pencils, erasers, and filing cabinets. Searle could receive Chinese characters through a slot in the door, process them according to the program's instructions, and produce Chinese characters as output. If the computer had passed the Turing test this way, it follows, says Searle, that he would do so as well, simply by running the program manually.
Searle asserts that there is no essential difference between the roles of the computer and himself in the experiment. Each simply follows a program, step-by-step, producing a behaviour which is then interpreted as demonstrating intelligent conversation. However, Searle would not be able to understand the conversation.

Searle argues that without "understanding" (or "intentionality"), we cannot describe what the machine is doing as "thinking" and since it does not think, it does not have a "mind" in anything like the normal sense of the word. Therefore, he concludes that "strong AI" is false.

Systems that act like humans
The Chinese room argument
(Continued)

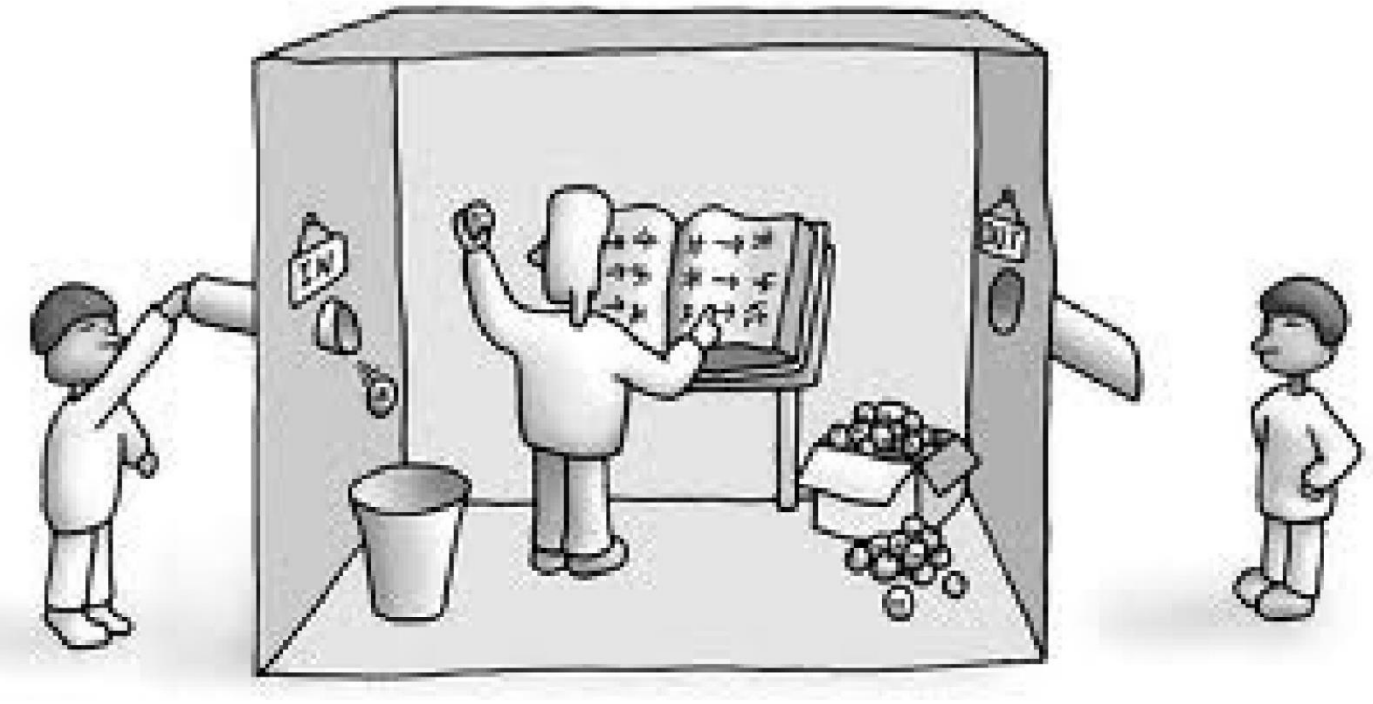

If person inside does a great job of answering questions, can we say he or she *understands*?

Even if she or he is only blindly following rules?

(Of course the 'person inside' is acting like an AI program)

Strong vs. Weak AI Hypotheses

WEAK AI Hypothesis
we can accurately *simulate* animal/human intelligence in a computer

STRONG AI Hypothesis
we can create algorithms that *are* intelligent (conscious?)

What is Artificial Intelligence?

Systems that act like humans	Systems that think rationally
"The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991) "The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)	"The study of mental faculties through the use of computational models" (Charniack and McDermott, 1985). "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Systems that think like humans	**Systems that act rationally**
"The automation of activities that we associate with human thinking, such as decision making, problem solving, learning" (Bellman, 1978) "The exciting new effort to make computers think … machines with minds, in the full and literal sense." (Haugeland, 1985)	"AI .. is concerned with intelligent behavior in artifacts (Nilsson, 1998) "Computational Intelligence is the study of the design of intelligent agents." (Poole et al., 1998)

Systems That Think Like Humans

- Need to study the brain as an information processing machine, … in other words …
- Use Computational Models to Understand the Actual Workings of Human Mind
  - Devise/Choose a sufficiently precise theory of the mind.
  - Express it as a computer program.
  - Check match between program and human behavior (actions and timing) on similar tasks.
- Tight connections with *Cognitive Science and Neuroscience.*
- Also known as *descriptive approaches* to AI.

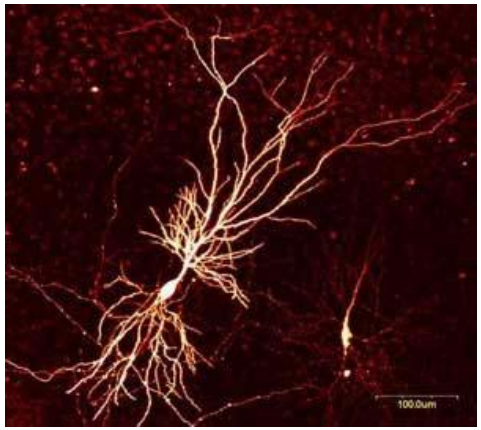

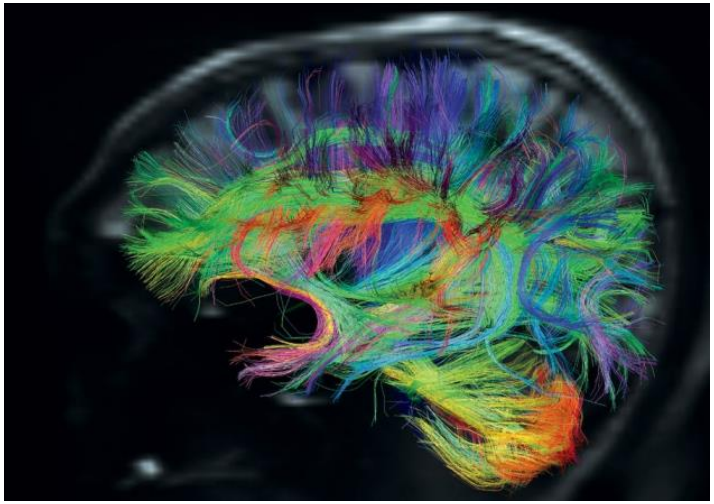

What is Artificial Intelligence?

Systems that act like humans	Systems that think rationally
"The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991) "The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)	"The study of mental faculties through the use of computational models" (Charniack and McDermott, 1985). "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Systems that think like humans	**Systems that act rationally**
"The automation of activities that we associate with human thinking, such as decision making, problem solving, learning" (Bellman, 1978) "The exciting new effort to make computers think … machines with minds, in the full and literal sense." (Haugeland, 1985)	"AI .. is concerned with intelligent behavior in artifacts (Nilsson, 1998) "Computational Intelligence is the study of the design of intelligent agents." (Poole et al., 1998)

Systems that Think Rationally

•***Logic*:** formalize *idealized* or *right* thinking, i.e. irrefutable reasoning processes. That is; patterns of argument that always yield correct conclusions when supplied with correct premises:

"Socrates is a man; all men are mortal; therefore Socrates is mortal."

- Logistic tradition in AI aims to build computational frameworks based on logic, that is, describe problem in formal logical notation and apply general deduction procedures to solve it.
- Then use these frameworks to build intelligent systems.

- Some examples are (Propositional Logic) and (Logic Programming).
- More advanced logic-based representations:
  - Semantic Networks

Systems that Think Rationally

- Main Research Problems/Challenges
  - Describing real-world problems and knowledge in logical notation.
  - Proving *Soundness* and *Completeness* of various formalisms.
  - How to represent often *informal* and *uncertain* domain knowledge and formalize it in logic notation *(i.e., dealing with Uncertainty)*.
  - Computational Complexity of finding a solution.
  - A lot of “rational” behavior has nothing to do with logic.

What is Artificial Intelligence?

Systems that act like humans	Systems that think rationally
"The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991) "The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)	"The study of mental faculties through the use of computational models" (Charniack and McDermott, 1985). "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Systems that think like humans	**Systems that act rationally**
"The automation of activities that we associate with human thinking, such as decision making, problem solving, learning" (Bellman, 1978) "The exciting new effort to make computers think … machines with minds, in the full and literal sense." (Haugeland, 1985)	"AI .. is concerned with intelligent behavior in artifacts (Nilsson, 1998) "Computational Intelligence is the study of the design of intelligent agents." (Poole et al., 1998)

Systems that Act Rationally

- The "*think rationally*" approach focuses on correct inference.
- But more is needed for *rational behavior,* e.g.
  - How to behave when there is no provably correct thing to do (i.e. reasoning under uncertainty).
  - Fully reactive behavior (instinct vs. reason).

AI as the Study & Design of Intelligent Agents
(Poole and Mackworth, 1999)

- An ***intelligent agent*** is such that:
  - Its ***actions*** are ***appropriate*** for its goals and circumstances.
  - It is ***flexible*** to changing environments and goals.
  - It ***learns*** from experience.
  - It makes ***appropriate choices*** given ***perceptual limitations*** and ***limited resources*** (**bounded rationality** or **bounded optimality**).
- This definition drops the constraint of *cognitive plausibility;*
  - Same as building flying machines by understanding general principles of flying (aerodynamic) vs. by reproducing how birds fly.

Thus, a **rational agent** acts to optimally achieve its goals (does the right thing). The right thing: that which is expected to maximize goal achievement, given the available information

Intelligent Agents

- In AI, artificial agents that have a physical presence in the world are usually known as *Robots*.
- Robotics is the field primarily concerned with the implementation of the physical aspects of a robot (i.e. perception of the physical environment, actions on the environment).
- Another class of artificial agents include *interface agents,* for either stand alone or Web-based applications (e.g. intelligent desktop assistants, recommender systems, intelligent tutoring systems).
- Interface agents don't have to worry about interaction with the physical environment, but share all other fundamental components of intelligent behavior with robots.
- We will focus on these agents in this course.

Pac-Man as an Intelligent Agent

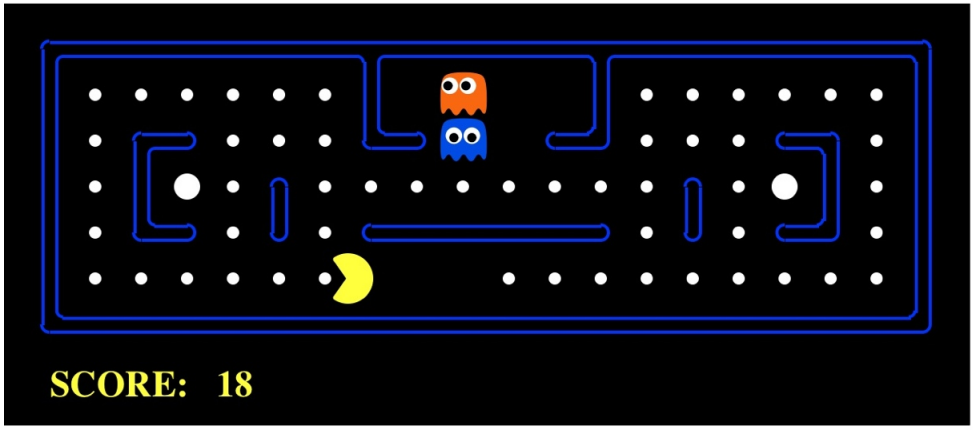

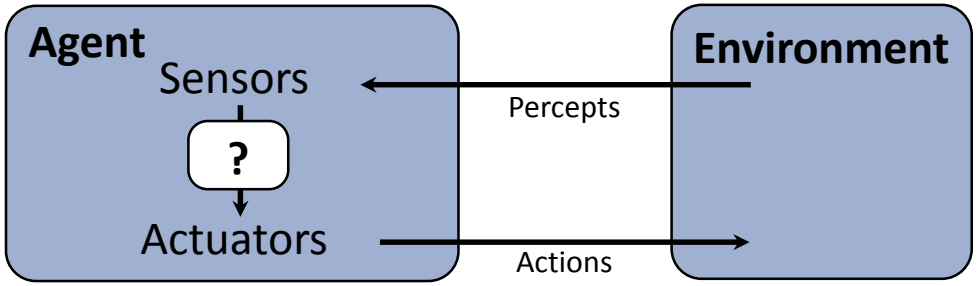

Intelligent Agents in the World

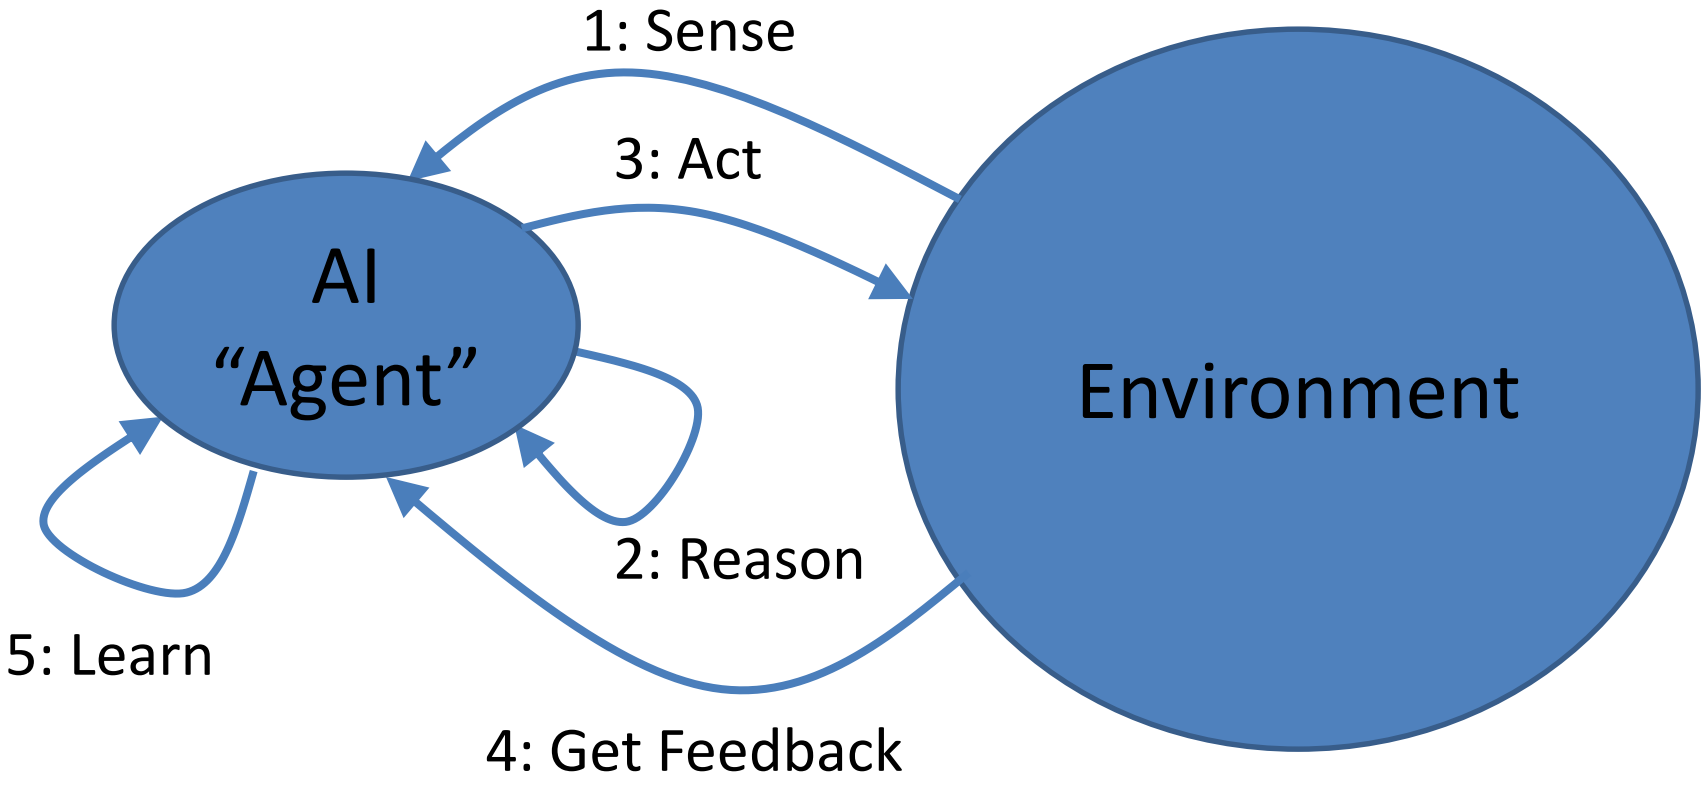

Intelligent Agents in the World

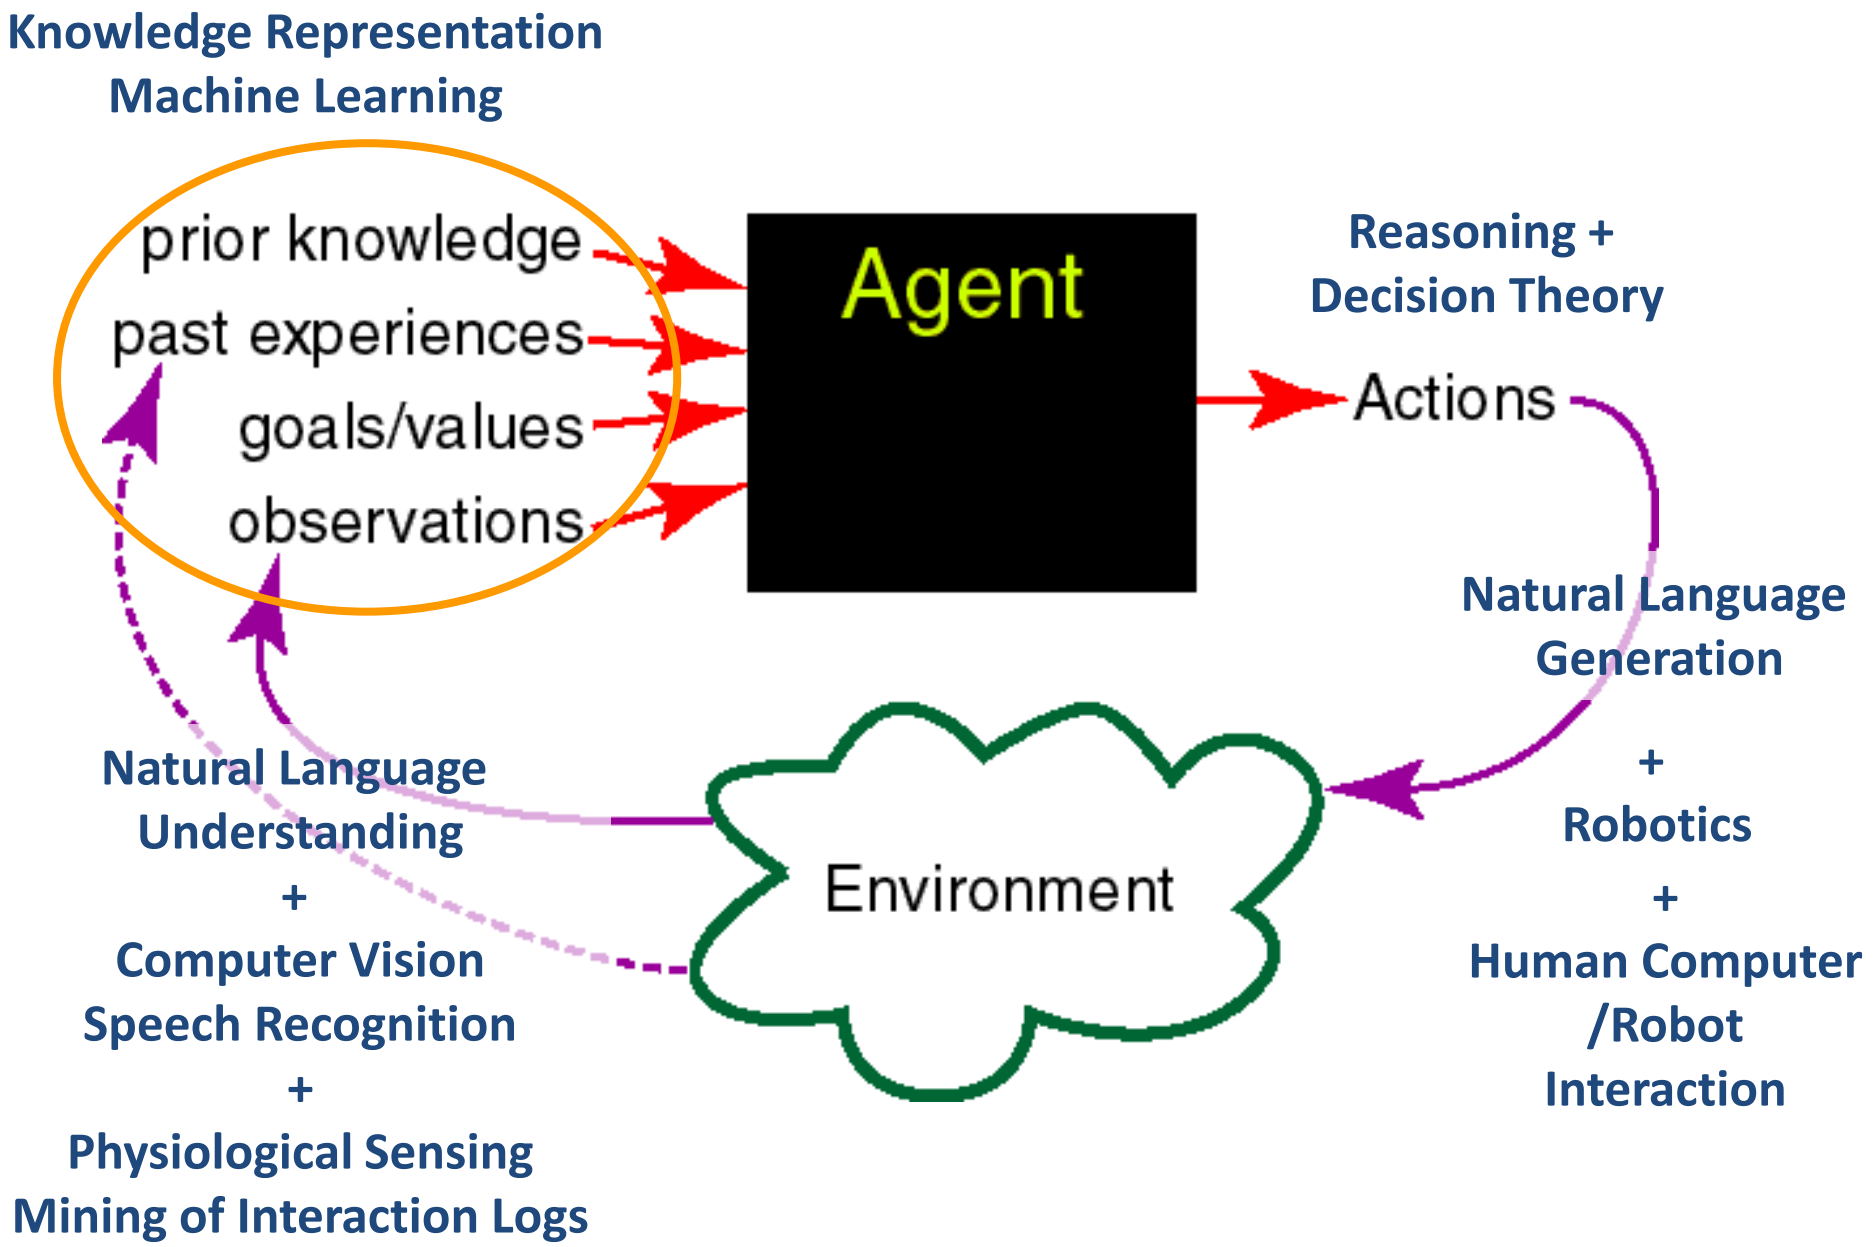

Outline

- What is Intelligence?
- What is Artificial Intelligence?
  - Systems that act like humans
    - Turing Test; the Imitation Game …
    - The Chinese room argument
    - Strong vs. Weak AI Hypotheses
  - Systems That Think Like Humans
  - Systems that Think Rationally
  - Systems that Act Rationally
    - AI as Study and Design of Intelligent Agents
    - Intelligent Agents
    - Pac-Man as an Intelligent Agent
    - Intelligent Agents in the World
- **Learning**
  - **Learning from Labelled Examples**
  - **Supervised Machine Learning: Task Overview**
    - **Standard Approach for Constructing a Machine Learning Dataset for a Task**
    - **Concept Learning**
  - **Generic Machine Learning Models**

Learning

"Learning denotes changes in the system that
… enable the system to do the same task …
more effectively the next time."

- **Herbert Simon**

"Learning is making useful changes in our minds."

- **Marvin Minsky**

Learning from Labelled Examples

Positive Examples

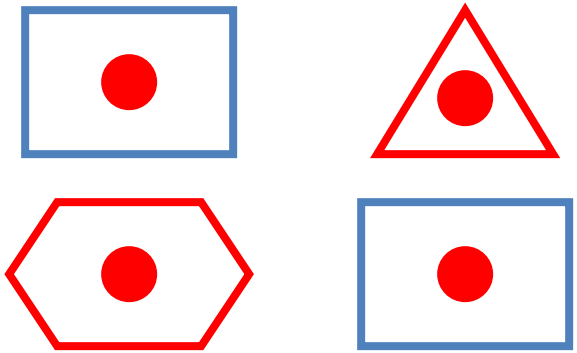

Negative Examples

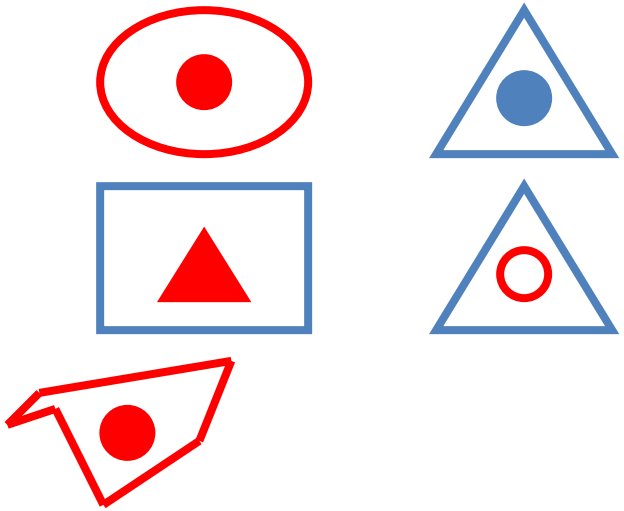

Category of this example?

Learning from Labelled Examples

Positive Examples

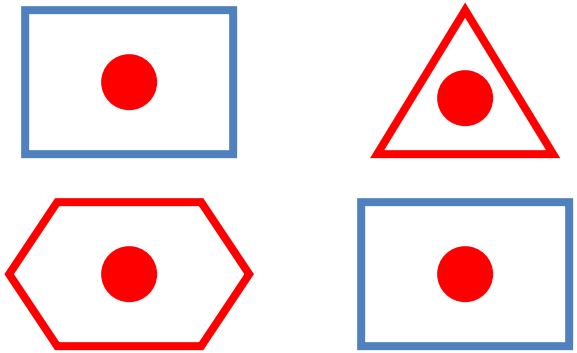

Negative Examples

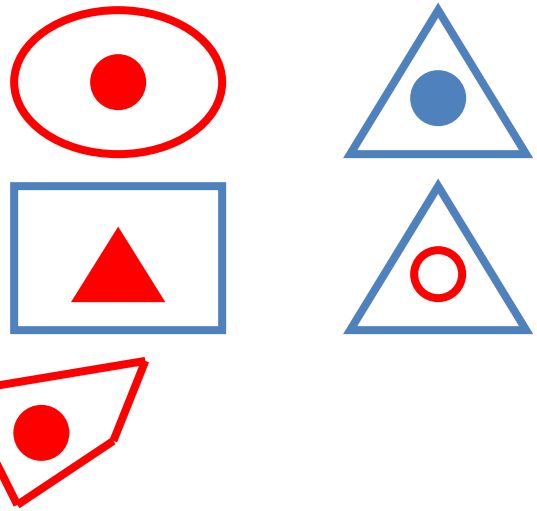

Category of this example?

Concept

Solid Red Circle in a (Regular?) Polygon

Learning from Labelled Examples

Positive Examples

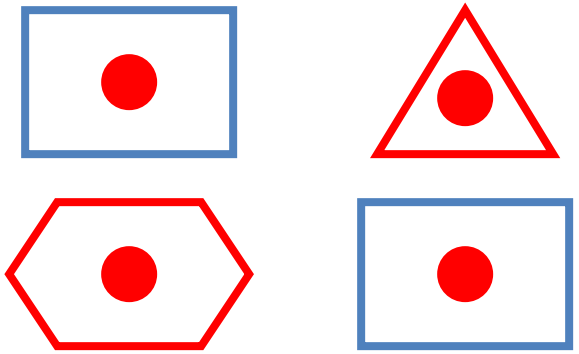

Negative Examples

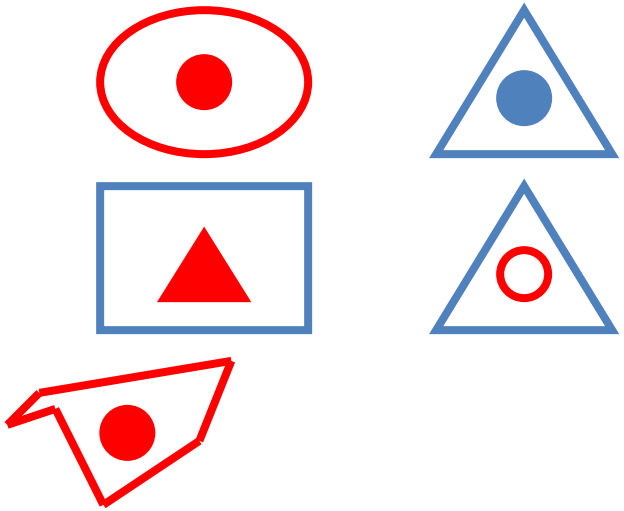

Category of this example?

Concept

Solid Red Circle in a (Regular?) Polygon

What about?

Figures on left side of page

Supervised Machine Learning: Task Overview

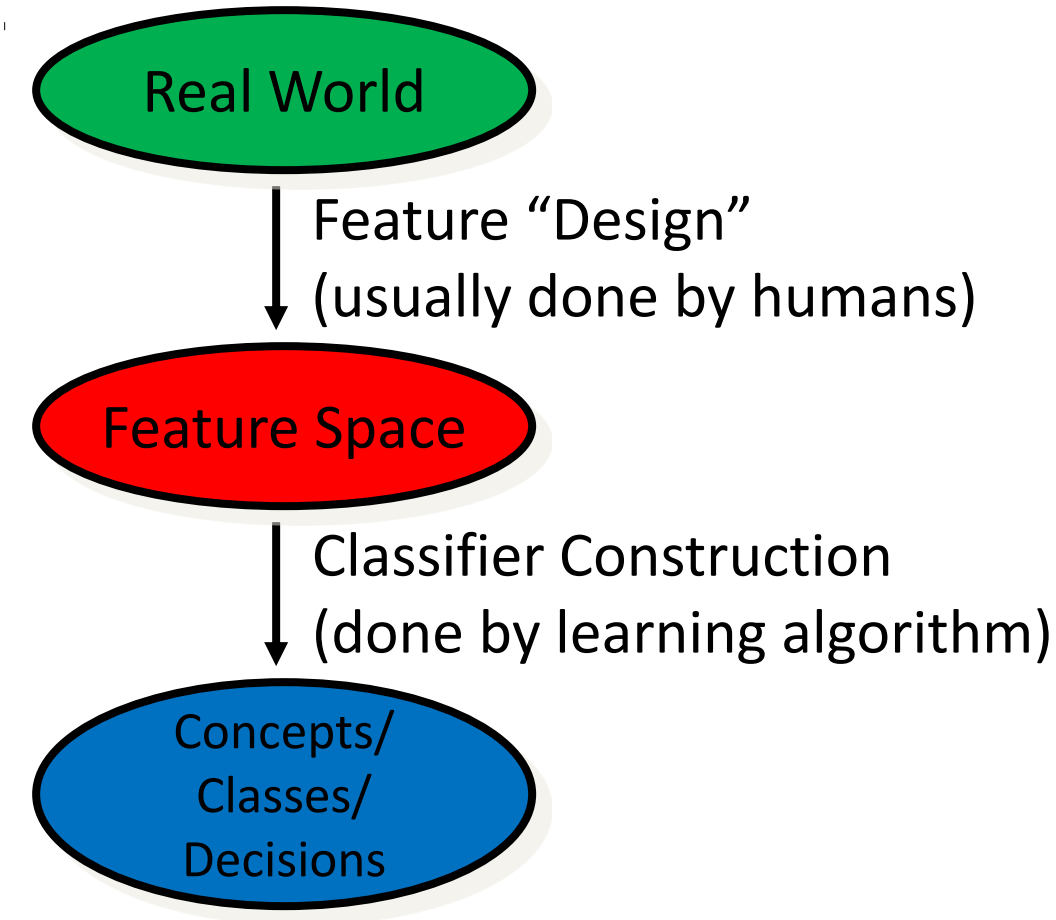

Standard Approach for Constructing a Machine Learning Dataset for a Task

Step 1: Choose a *feature space*

We will use fixed-length feature vectors

Choose N features

Each feature has V_i possible values

} **Defines a space**

Each example is represented by a vector of N feature values (is a point in the **feature space**)

eg <red, 50, round>

color weight shape

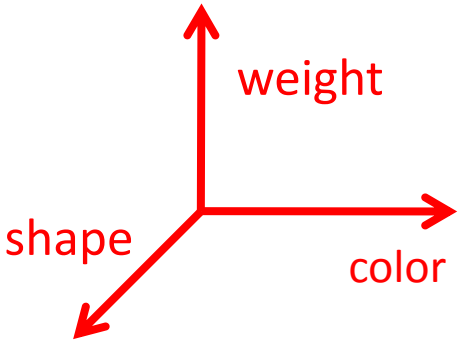

Standard Approach for Constructing a Machine Learning Dataset for a Task

Step 2: Collect examples ("I/O" pairs)

	Feature 1	Feature 2	. . .	Feature *N*	Output Category
Example 1	0.0	small		red	**true**
Example 2	9.3	medium		red	**false**
Example 3	8.2	small		blue	**false**
. . .					
Example *M*	5.7	medium		green	**true**

Concept Learning

Have selected ‘concept’ to learn ✓

Have chosen features to rep examples ✓

Have created at least *(for instance)* 100 labeled examples ✓

Next: learn a ‘model’ that can predict output for NEW examples

Concept Learning

Learning systems differ in how they represent concepts:

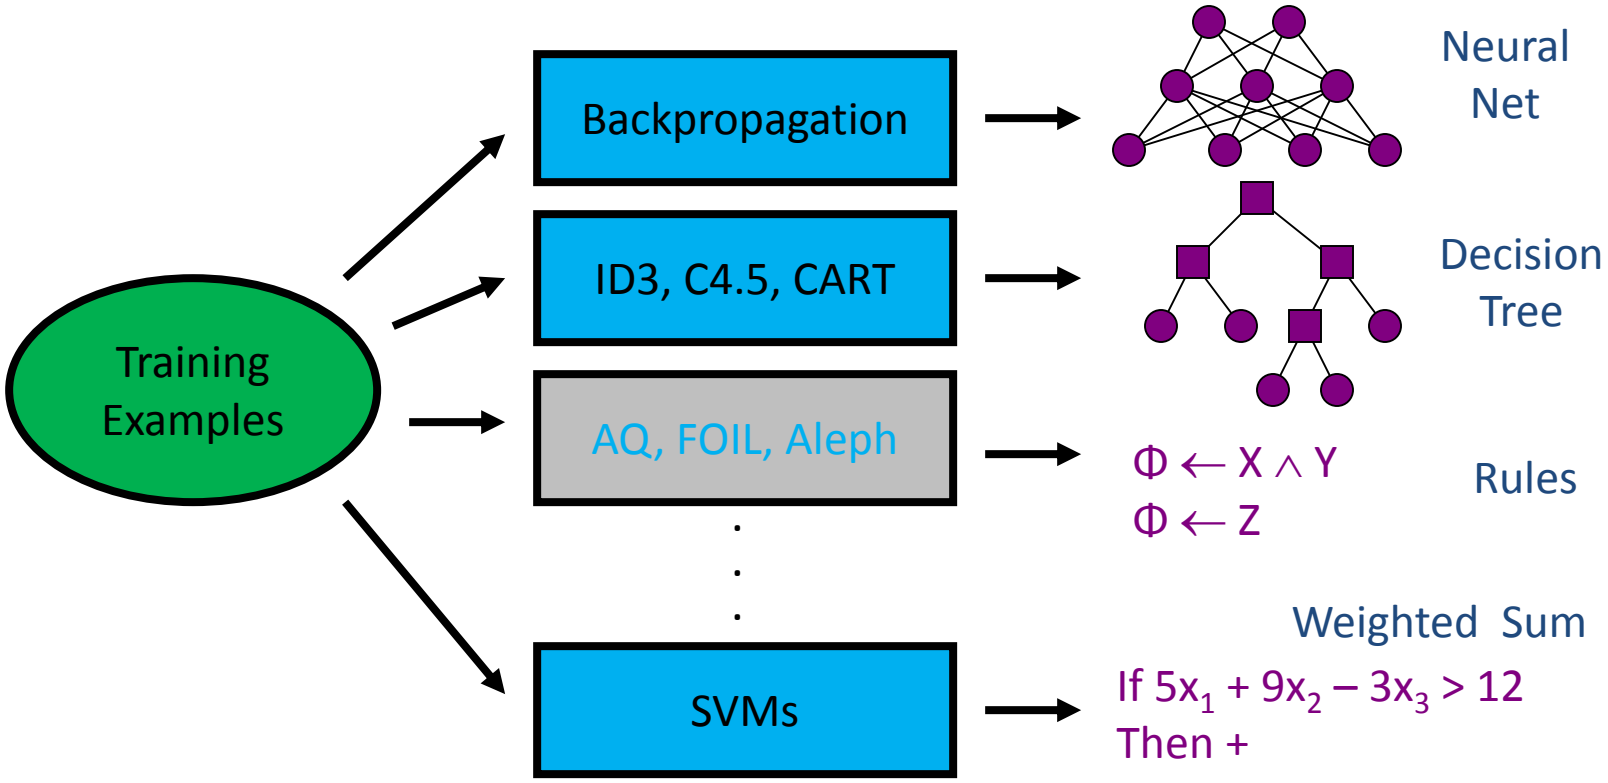

Generic Machine Learning Models

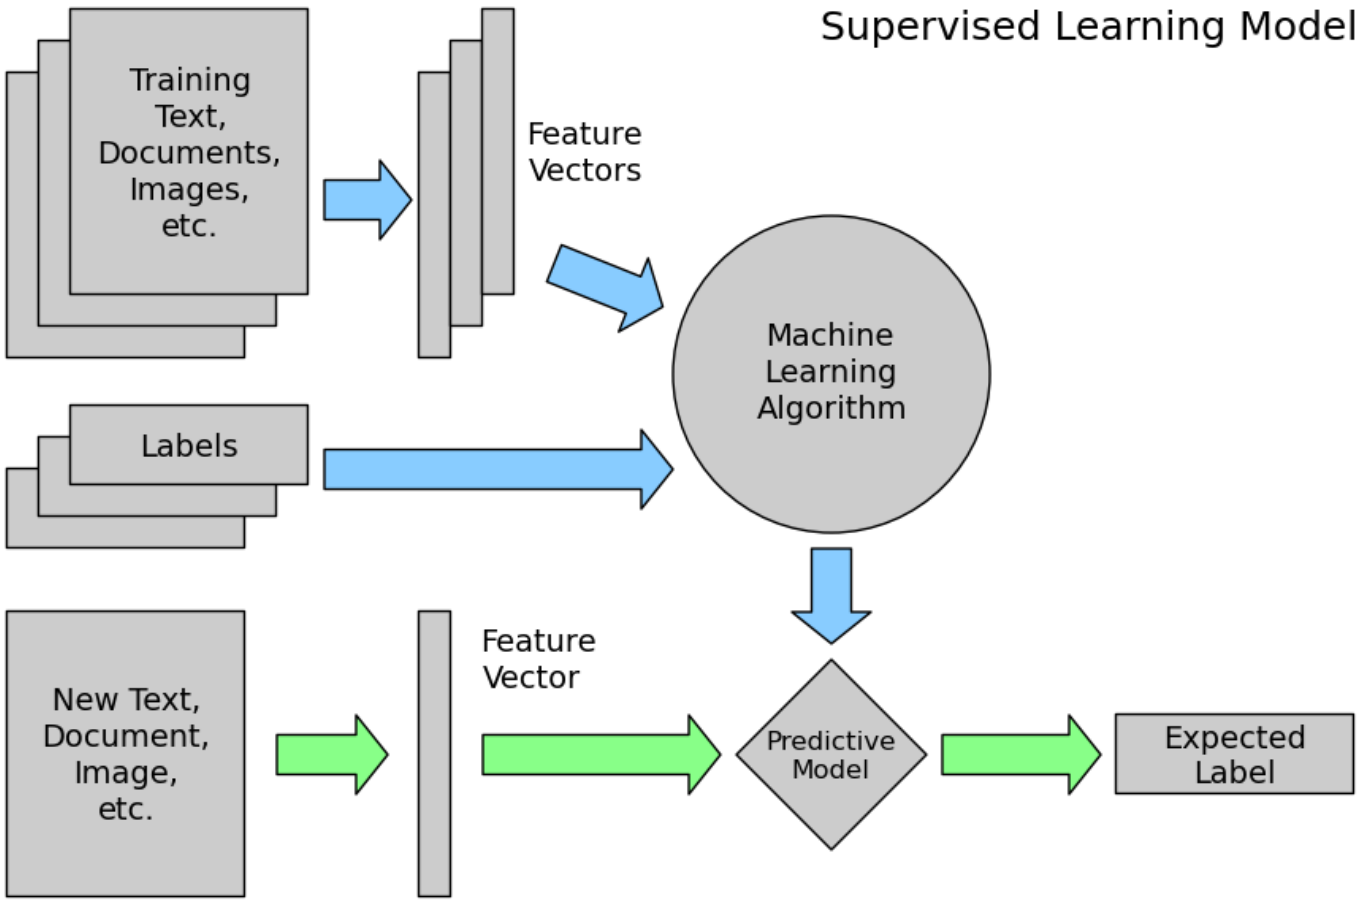

Generic Machine Learning Models

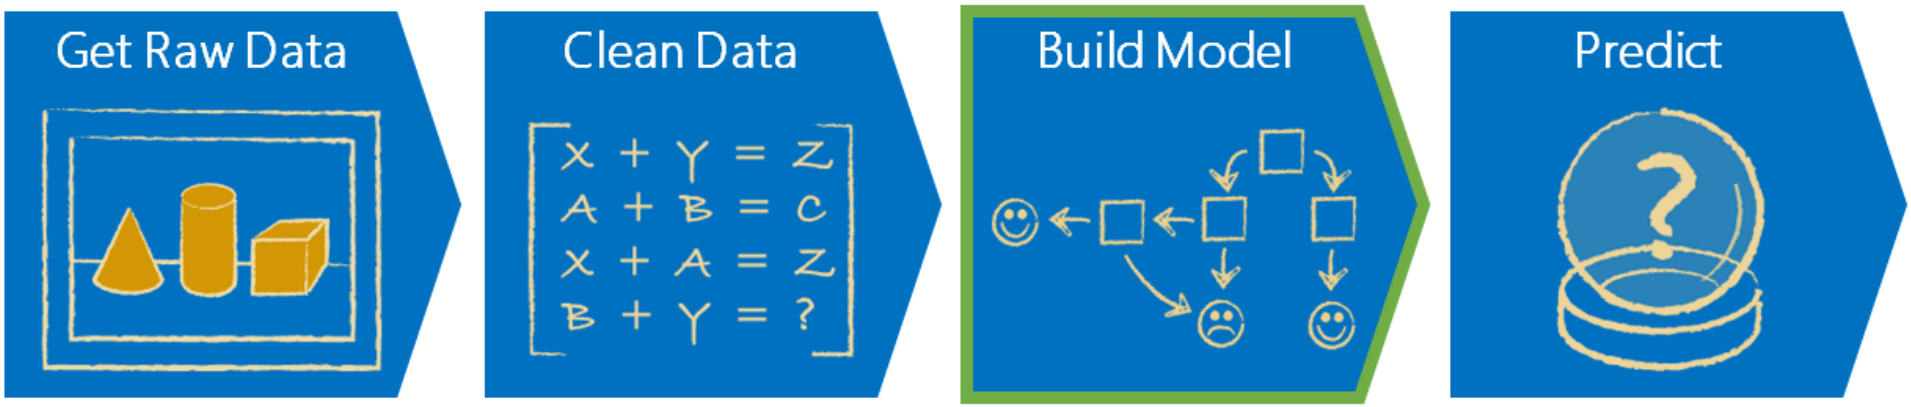

Thanks! … *Questions?*