

Course: Analytics, Machine
Learning, and the Digital Economy

Artificial Intelligence

Lecturer Radjabova Dilnora

Thinking Machines

- Are there tasks which cannot easily be automated? If so, what are the limitations?
- How do computers abilities compare to that of humans?

What is an AI?

Computers versus humans

- A computer can do some things better than a human can
 - Adding a thousand four-digit numbers
 - Drawing complex, 3D images
 - Store and retrieve massive amounts of data
- However, there are things humans can do much better.

Thinking Machines



A computer would have difficulty identifying the cat, or matching it to another picture of a cat.

Computer or human?

- Which of the following occupations could (or should) be performed by computers?
 - Postman
 - Bookstore clerk
 - Librarian
 - Doctor
 - Lawyer
 - Judge
 - Professor

Thinking Machines

Artificial intelligence (AI)

The study of computer systems that attempt to model and apply the intelligence of the human mind

For example, writing a program to pick out objects in a picture

First things first...

- Of course, first we have to understand why we use the term “intelligence” in regard to humans.
 - What defines “intelligence”?
 - Why is it that we assume humans are intelligent?
 - Are monkeys intelligent? Dogs? Ants? Pine trees?

Early History

- In 1950 English mathematician Alan Turing wrote a landmark paper titled “Computing Machinery and Intelligence” that asked the question: “Can machines think?”
- Further work came out of a 1956 workshop at Dartmouth sponsored by John McCarthy. In the proposal for that workshop, he coined the phrase a “study of artificial intelligence”

Can Machines Think?


- So Turing asked: “Can machines think?” He felt that such machines would eventually be constructed.
- But he also realized a bigger problem. How would we know if we’ve succeeded?

The Turing Test

Turing test

A test to empirically determine whether a computer has achieved intelligence

Figure 13.2
In a Turing test, the interrogator must determine which respondent is the computer and which is the human



The Turing Test

- Passing the Turing Test does not truly show that the machine was thinking. It simply shows that it generated behavior consistent with thinking.
- **weak equivalence:** the two systems (human and computer) are equivalent in **results** (output), but they do not necessarily arrive at those results in the same way
- **Strong equivalence:** the two systems use the same internal processes to produce results


The Turing Test

Loebner prize

The first formal instantiation of the Turing test, held annually

Chatbots

A program designed to carry on a conversation with



Knowledge Representation

- We want to compare the way that computers and humans work to see if we can better understand why each have their (computational) strengths.
 - Processing Models
 - Knowledge Representation
 - Reasoning

Semantic Networks

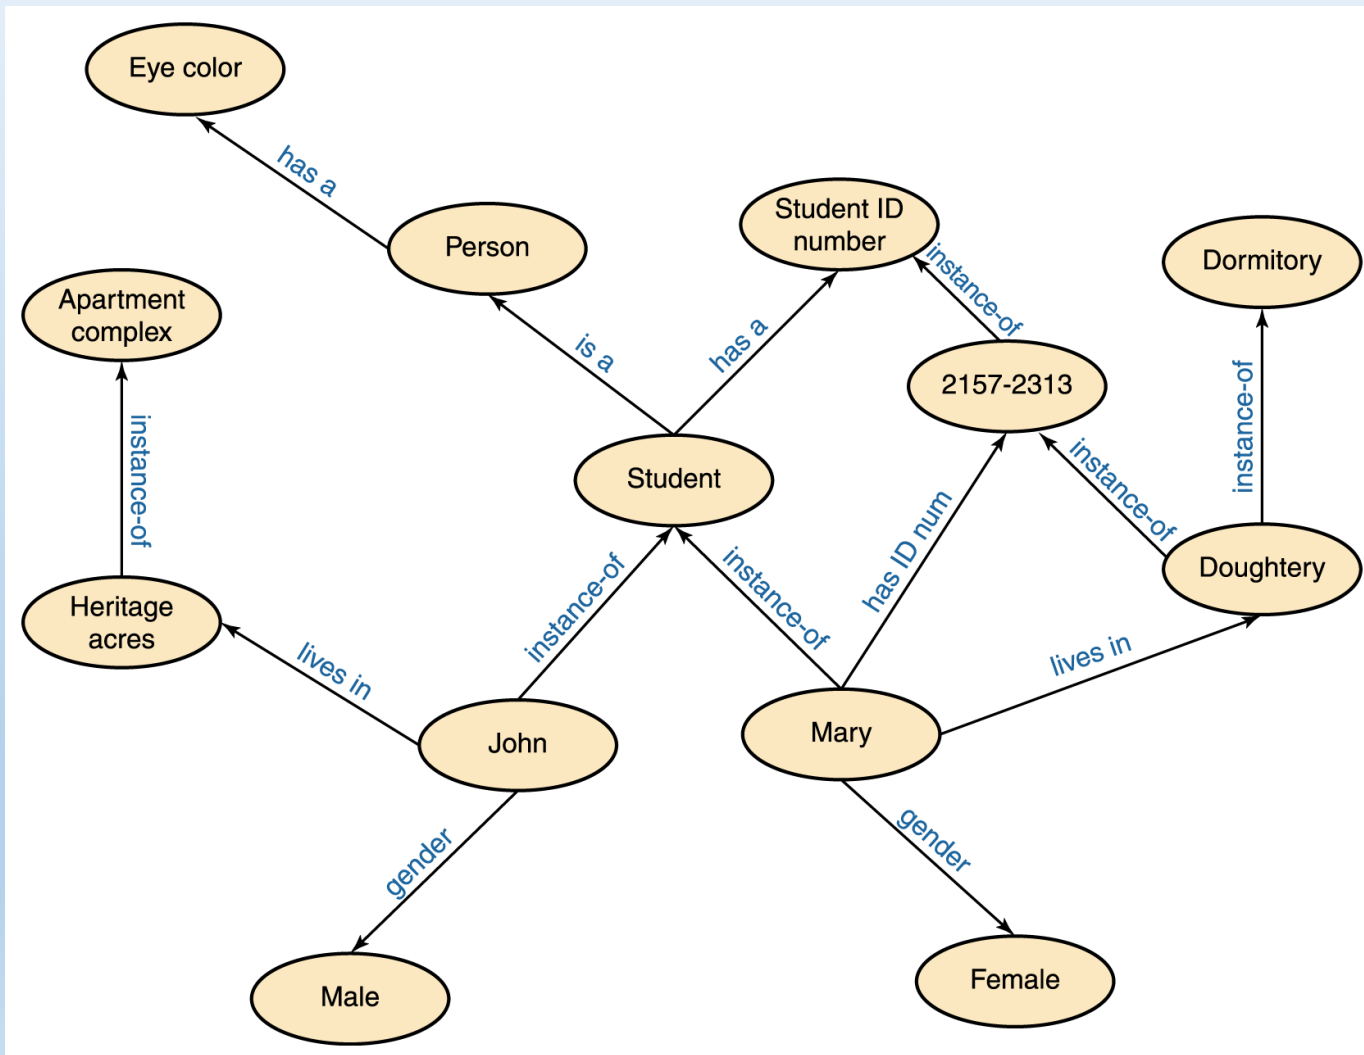
Semantic network

A knowledge representation technique that focuses on the relationships between objects

A directed graph is used to represent a semantic network or net

Remember directed graphs? (See Chapter 8.)

Semantic Networks



Semantic Networks

What questions can you ask about the data in Figure 13.3 (previous slide)?

What questions can you not ask?

Semantic Networks

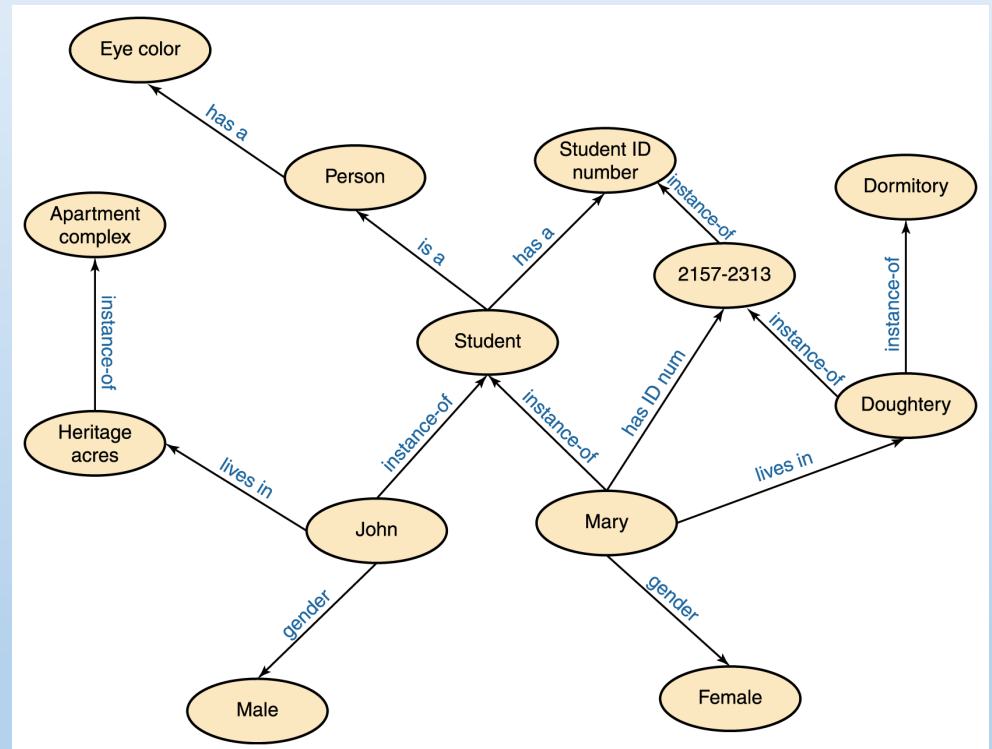
Network Design

- The objects in the network represent the objects in the real world that we are representing
- The relationships that we represent are based on the real world questions that we would like to ask
- That is, the types of relationships represented determine which questions are easily answered, which are more difficult to answer, and which cannot be answered

Difficult questions

For example, it would be difficult to ask “how many students are female?” or “who lives in Doughtery Hall?”

The answers are present - just difficult to find by searching here.



Search Trees

Search tree

A structure that represents alternatives in adversarial situations such as game playing

The paths down a search tree represent a series of decisions made by the players

*Remember trees?
(See Chapter 8.)*

Example: (simplified) Nim

- There are a row of spaces. The game alternates:
 - Player 1 may place 1, 2, or 3 x's in the leftmost spaces.
 - Player 2 may then place 1, 2, or 3 o's in the next spaces.
- Goal: place the last mark in the rightmost space.

Search Trees

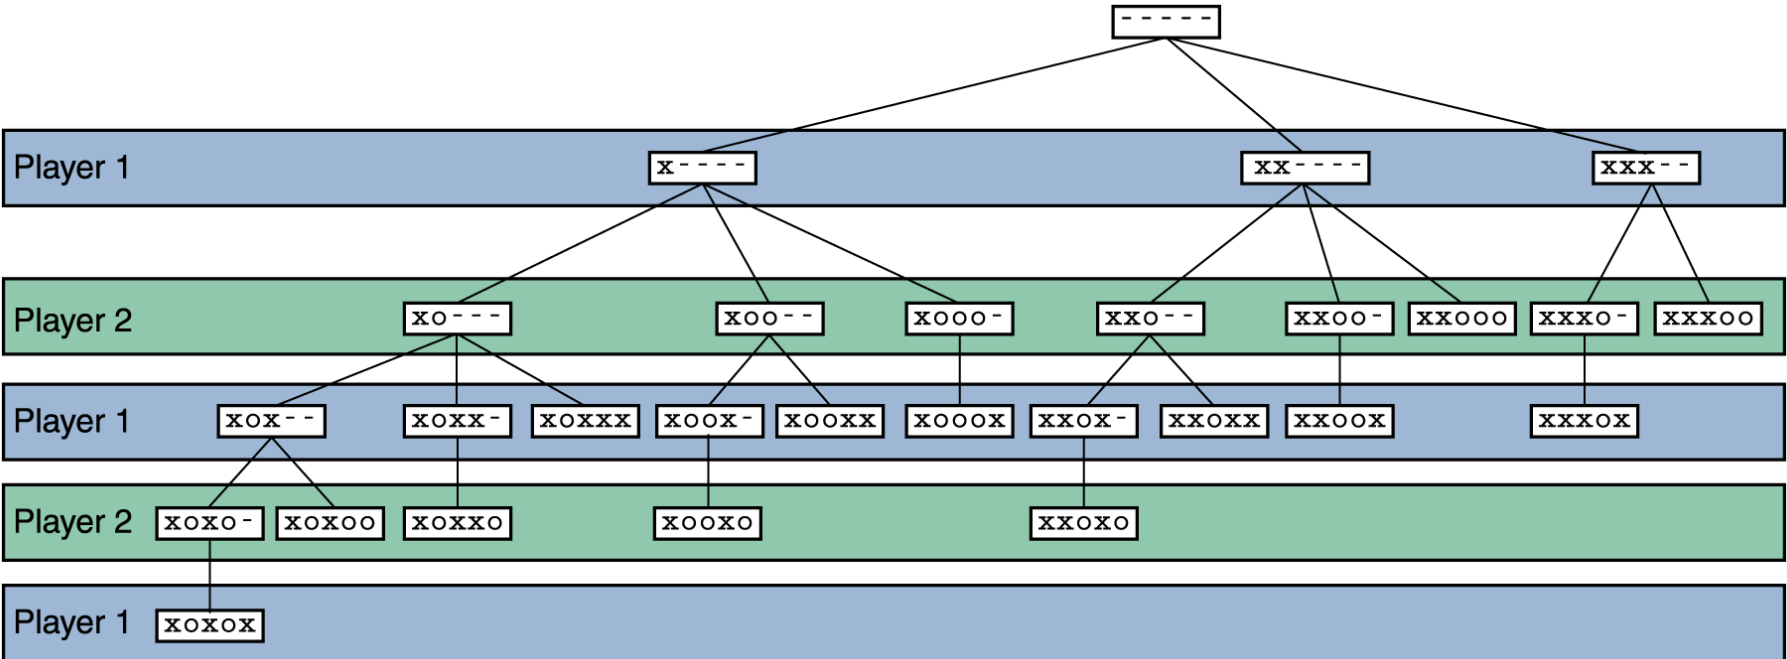


Figure 13.4 A search tree for a simplified version of Nim

Search Trees

Search tree analysis can be applied to other, more complicated games such as chess

However, full analysis of the chess search tree would take more than your lifetime to determine the first move

Because these trees are so large, only a fraction of the tree can be analyzed in a reasonable time limit, even with modern computing power

Therefore, we must find a way to prune the tree

Search Trees

Techniques for pruning search space

- **Depth-first** A technique that involves the analysis of selected paths all the way down the tree
- **Breadth-first** A technique that involves the analysis of all possible paths but only for a short distance down the tree

Search Trees

Depth-First Search

Breadth-first Search




Figure 13.5 Depth-first and breadth-first searches

Search Tree strategies

- Breadth-first strategies tend to be more successful (and computationally feasible).
- Search tree analysis can be applied nicely to other, more complicated games such as chess.
- In 1997, IBM's Deep Blue became the first program to defeat a master-level chess player.

Expert Systems

Knowledge-based system

Software that uses a specific set of information, from which it extracts and processes particular pieces

Expert system

A software system based the knowledge of human experts; it is a

- **Rule-based system**
- A software system based on a set of *if-then* rules
- **Inference engine**
- The software that processes rules to draw conclusions

Expert Systems



Gardner Expert System Example

Expert Systems

Named abbreviations that represent conclusions

- **NONE**—apply no treatment at this time
- **TURF**—apply a turf-building treatment
- **WEED**—apply a weed-killing treatment
- **BUG**—apply a bug-killing treatment
- **FEED**—apply a basic fertilizer treatment
- **WEEDFEED**—apply a weed-killing and fertilizer combination treatment

Expert Systems

Boolean variables needed to represent state of the lawn

- **BARE**—the lawn has large, bare areas
- **SPARSE**—the lawn is generally thin
- **WEEDS**—the lawn contains many weeds
- **BUGS**—the lawn shows evidence of bugs

Expert Systems

Data that is available

- **LAST**—the date of the last lawn treatment
- **CURRENT**—current date
- **SEASON**—the current season

Now we can formulate some rules for our gardening expert system Rules take the form of *if-then* statements

Expert Systems

Some rules

- if (**CURRENT** – **LAST** < 30) then **NONE**
- if (**SEASON** = winter) then not **BUGS**
- if (**BARE**) then **TURF**
- if (**SPARSE** and not **WEEDS**) then **FEED**
- if (**BUGS** and not **SPARSE**) then **BUG**
- if (**WEEDS** and not **SPARSE**) then **WEED**
- if (**WEEDS** and **SPARSE**) then **WEEDFEED**

Expert Systems

An execution of our inference engine

- **System:** Does the lawn have large, bare areas?
- **User:** No
- **System:** Does the lawn show evidence of bugs?
- **User:** No
- **System:** Is the lawn generally thin?
- **User:** Yes
- **System:** Does the lawn contain significant weeds?
- **User:** Yes
- **System:** You should apply a weed-killing and fertilizer combination treatment.

Artificial Neural Network

Artificial neural networks

A computer representation of knowledge that attempts to mimic the neural networks of the human body

Yes, but what is a human neural network?

Neural Network




Figure 13.6 A biological neuron

Neural Network

Neuron

A single cell that conducts a chemically-based electronic signal

At any point in time a neuron is in either an **excited** state or an **inhibited** state

Excited state

Neuron conducts a strong signal

Inhibited state

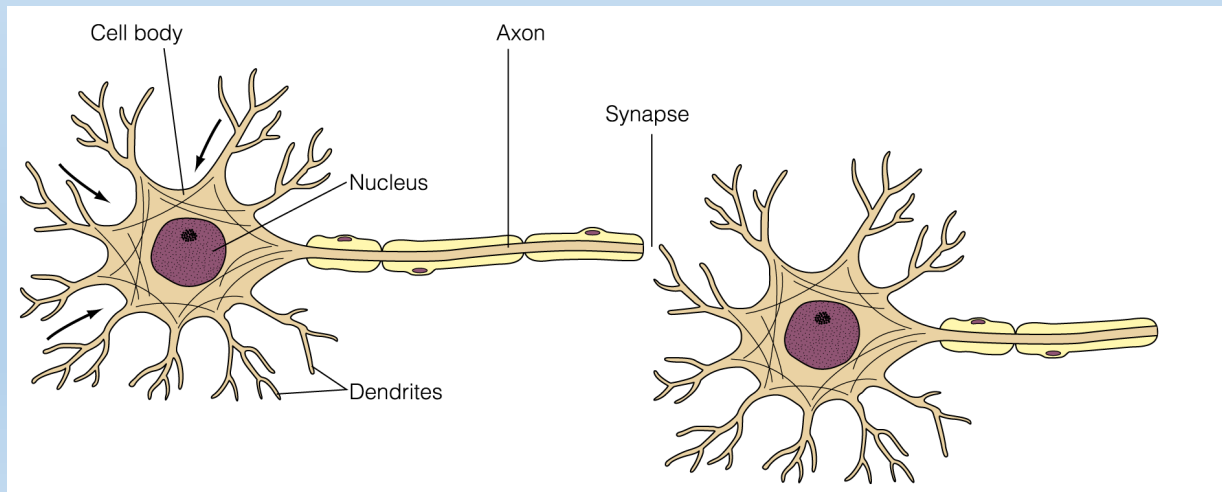
Neuron conducts a weak signal

Neural Network

- A series of connected neurons forms a pathway
- A series of excited neurons creates a strong pathway
- A biological neuron has multiple input tentacles called dendrites and one primary output tentacle called an axon
- The gap between an axon and a dendrite is called a synapse

Neural network


- A neuron accepts multiple input signals and then controls the contribution of each signal based on the “importance” the corresponding synapse gives to it



Neural Network

The pathways along the neural nets are in a constant state of flux

As we learn new things, new strong neural pathways in our brain are formed



Artificial Neural Networks

Each processing element in an **artificial neural net** is analogous to a biological neuron

- An element accepts a certain number of input values (dendrites) and produces a single output value (axon) of either 0 or 1
- Associated with each input value is a numeric weight (synapse)

Artificial Neural Networks

- The **effective weight** of the element is the sum of the weights multiplied by their respective input values

$$v1 * w1 + v2 * w2 + v3 * w3$$

- Each element has a numeric threshold value
- If the effective weight exceeds the threshold, the unit produces an output value of 1
- If it does not exceed the threshold, it produces an output value of 0

Artificial Neural Networks

- The process of adjusting the weights and threshold values in a neural net is called **training**
- A neural net can presumably be trained to produce whatever results are required
- But think about the complexity of this! Train a computer to recognize any cat in a picture, based on training run with several pictures.

Human vs. Computer

	Human Brain	Computer
speed	Neurotransmitters travel at about 1000 ft/second	Electrons at speed of light
memory	Roughly 100 billion neurons - about 50 trillion bits	Top super computers might approach this much memory
other	Each neuron connected to 1000 others (roughly)	Perhaps 100 parallel processors

Natural Language Processing

- There are three basic types of processing going on during human/computer voice interaction
 - Voice recognition—recognizing human words
 - Natural language comprehension—interpreting human communication
 - Voice synthesis—recreating human speech
- Common to all of these problems is the fact that we are using a natural language, which can be any language that humans use to communicate

Voice Synthesis

One Approach to Voice Synthesis

Dynamic voice generation

A computer examines the letters that make up a word and produces the sequence of sounds that correspond to those letters in an attempt to vocalize the word

Phonemes

The sound units into which human speech has been categorized

Voice Synthesis

Consonants			
Symbols	Examples	Symbols	Examples
p	pipe	k	kick, cat
b	babe	g	get
m	maim	ŋ	sing
f	fee, phone, rough	ʃ	shoe, ash, sugar
v	vie, love	ʒ	measure
θ	thin, bath	ç	chat, batch
ð	the, bathe	ʝ	jaw, judge, gin
t	tea, beat	d	day, bad
n	nine	ʔ	uh uh
l	law, ball	s	see, less, city
r	run, bar	z	zoo, booze

Semi Vowels	
w	we
h	he
j	you, beyond

Vowels	
Symbols	Examples
i	eel, sea, see
I	ill, bill
e	ale, aim, day
ɛ	elk, bet, bear
æ	at, mat
u	due, new, zoo
ʊ	book, sugar
o	own, no, know
ɔ	aw, crawl, law, dog
a	hot, bar, dart
ə	sir, nerd, bird
ʌ	cut, bun

Diphthongs	
aj	bite, fight
aw	out, cow
ɔj	boy, boil

Figure 13.7 Phonemes for American English

Voice Synthesis

Another Approach to Voice Synthesis

Recorded speech

A large collection of words is recorded digitally and individual words are selected to make up a message

Since words are pronounced differently in different contexts, some words may have to be recorded multiple times


- For example, a word at the end of a question rises in pitch compared to its use in the middle of a sentence

Voice Recognition

Problems with understanding speech

- Each person's sounds are unique
- Each person's shape of mouth, tongue, throat, and nasal cavities that affect the pitch and resonance of our spoken voice are unique
- Speech impediments, mumbling, volume, regional accents, and the health of the speaker are further complications

Voice Recognition



Voice Recognition

Other problems

- Humans speak in a *continuous, flowing* manner, stringing words together
- Sound-alike phrases like “ice cream” and “I scream”
- Homonyms such as “I” and “eye” or “see” and “sea”

Humans can often clarify these situations by the context of the sentence, but that processing requires another level of comprehension

Modern voice-recognition systems still do not do well with continuous, conversational speech

Voice Recognition

Voiceprint

The plot of frequency changes over time representing the sound of human speech

A human *trains* a voice-recognition system by speaking a word several times so the computer gets an average voiceprint for a word

Used to authenticate the declared sender of a voice message

Natural Language Comprehension

Natural language is ambiguous!

Lexical ambiguity

The ambiguity created when words have multiple meanings

Syntactic ambiguity

The ambiguity created when sentences can be constructed in various ways

Referential ambiguity

The ambiguity created when pronouns could be applied to multiple objects

Lexical Ambiguity

- A single word can have two meanings.
 - *“The bat slipped from his hand”*
 - *“Cinderella had a ball”*
 - *“Ron lies asleep in his bed”*
- Worse yet, those meanings may even constitute different parts of speech.
 - *“Time flies like an arrow”*
 - *“They are racing horses”*
 - *“Stampeding cattle can be dangerous”*

Syntactic Ambiguity

- Even if all words have a clear meaning, ambiguity may exist because the phrases can be combined in several ways when parsing.
 - *“I saw the Grand Canyon flying to New York”*
 - *“The clams are ready to eat”*
 - *“I saw the man in the park with the telescope”*

Referential Ambiguity

- Pronouns may cause ambiguity when it is not clear which noun is being referenced.
 - *“The brick fell on the computer but it is not broken”*
 - *“Jon met Bill before he went to the store”*

Rules of Conversation

- *“Do you know what time it is?”*
 - Presumably, a correct response is “Yes.”
 - Does this sentence’s meaning change if it is said by your boss when you walk into a meeting 30 minutes late?
- *“Do you know you have a flat tire?”*
- *“I’d like to ask everyone to raise their hand for two seconds.”*
 - How did you respond?

“Real world” knowledge

- “Norman Rockwell painted people.”
 - Did he do tatoos? Face painting?
 - Without background knowledge, difficult to parse this into real information.

Natural Language Comprehension

Lexical ambiguity

Stand up for your country

Take the street on the left

Syntactic ambiguity

I saw the bird watching from the corner

I ate the sandwich sitting on the table

Referential ambiguity

The bicycle hit the curb, but it was not damaged

John was mad at Bill, but he didn't care

Can you think
of
some others?

Robotics

Mobile robotics

The study of robots that move relative to their environment, while exhibiting a degree of autonomy

Sense-plan-act (SPA) paradigm

The world of the robot is represented in a complex semantic net in which the sensors on the robot are used to capture the data to build up the net




Figure 13.8 The sense-plan-act (SPA) paradigm

Subsumption Architecture

Rather than trying to model the entire world all the time, the robot is given a simple set of behaviors each associated with the part of the world necessary for that behavior

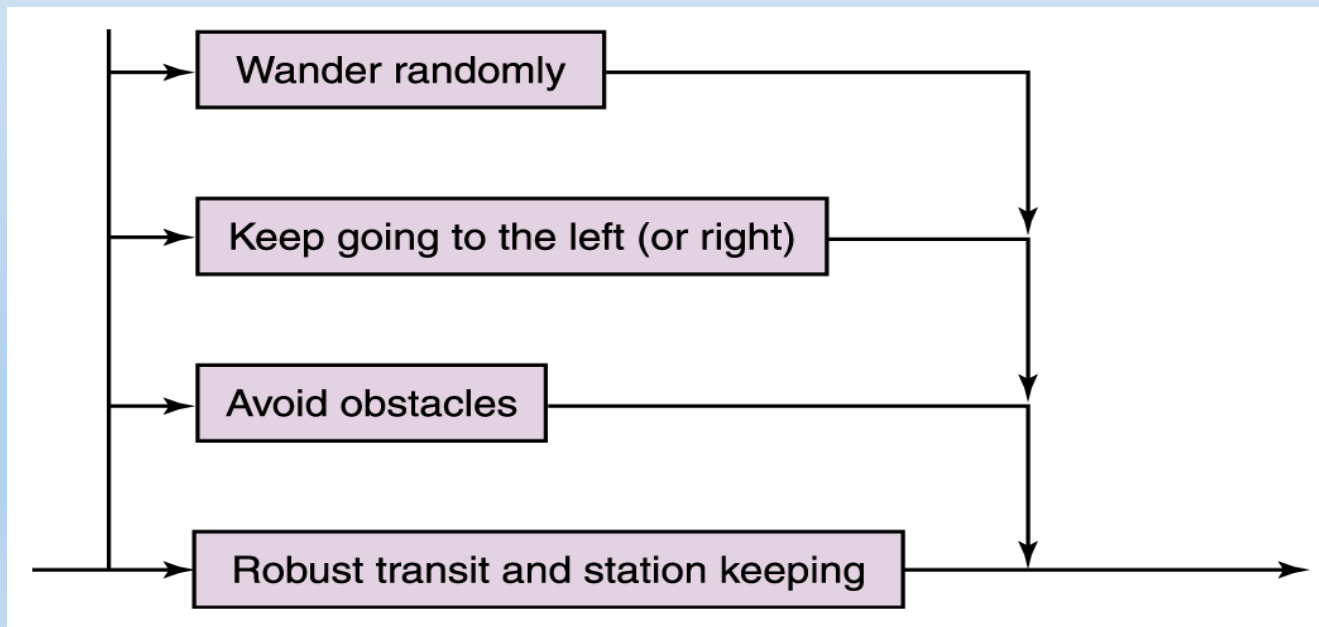


Figure 13.9
The new control paradigm

Subsumption Architecture

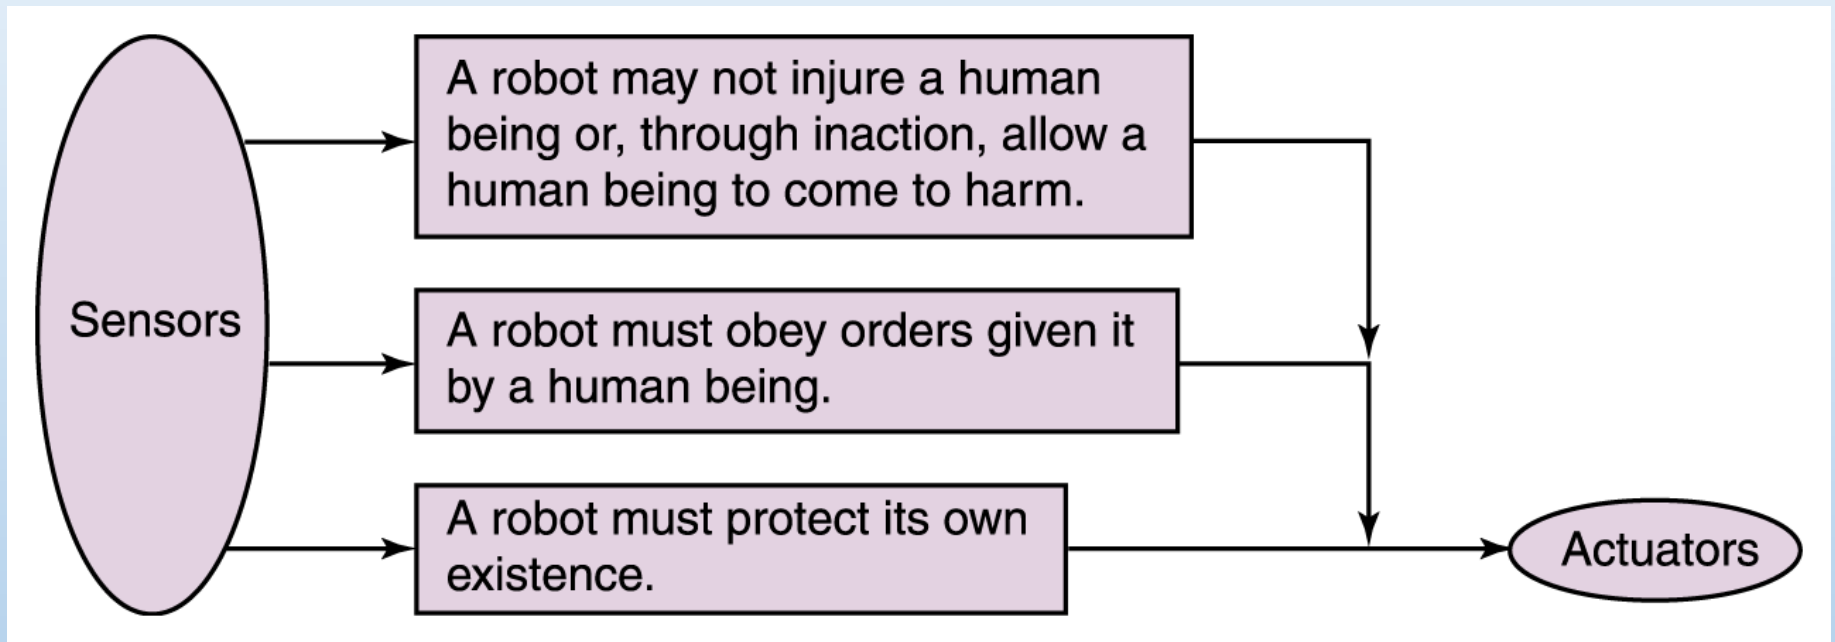



Figure 13.10 Asimov's laws of robotics are ordered.

Robots




Sony's Aibo

Robots



*Sojourner
Rover*

Robots



*Spirit or
Opportunity Rover*

Reference

- 1. Shalev-Shwartz and Ben-David. Understanding Machine Learning: From Theory to Algorithms (Cambridge University Press, 2014)
- 2. Daum´e. A Course in Machine Learning.
- 3. The Art of Statistics: How to Learn from Data by David Shpigelter
- 4. Learning From Data – January 1, 2012 by Yaser S. Abu-Mostafa (Author), Malik Magdon-Ismael (Author), Hsuan-Tien Lin (Author)
- 5. Statistics: The Art and Science of Learning from Data by Alan Agresti
- 6. Learning From Data: An Introduction To Statistical Reasoning by M.Glenber.
- 7. Statistics: Learning from Data (with JMP Printed Access Card) by Rocky Pek
- 8. The Elements of Statistical Learning by Gerim Garold
- 9. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems 2nd Edition
by Aurélien Géron (Author)