

CMSC 471

Intro to Artificial Intelligence

Fall 2023

KMA Solaiman

ksolaima@umbc.edu

Instructor: KMA Solaiman (Salvi)

ITE 201C/Remote

ksolaima@umbc.edu

Wednesday: 1-2 pm

Monday: 1-2pm (if needed)

by appointment

Multimodal Information
Retrieval

Vision & language processing

Learning with low-to-no
supervision

Novelties in Learning Models

TA

Prasanna Bollineni

[LinkedIn](#)

lakshmb1@umbc.edu

Office hours:

Mon 11:30-12:30pm,
and by appointment

Location: ITE 334

Shashank Sacheti

[LinkedIn](#)

ssachet1@umbc.edu

Office hours:

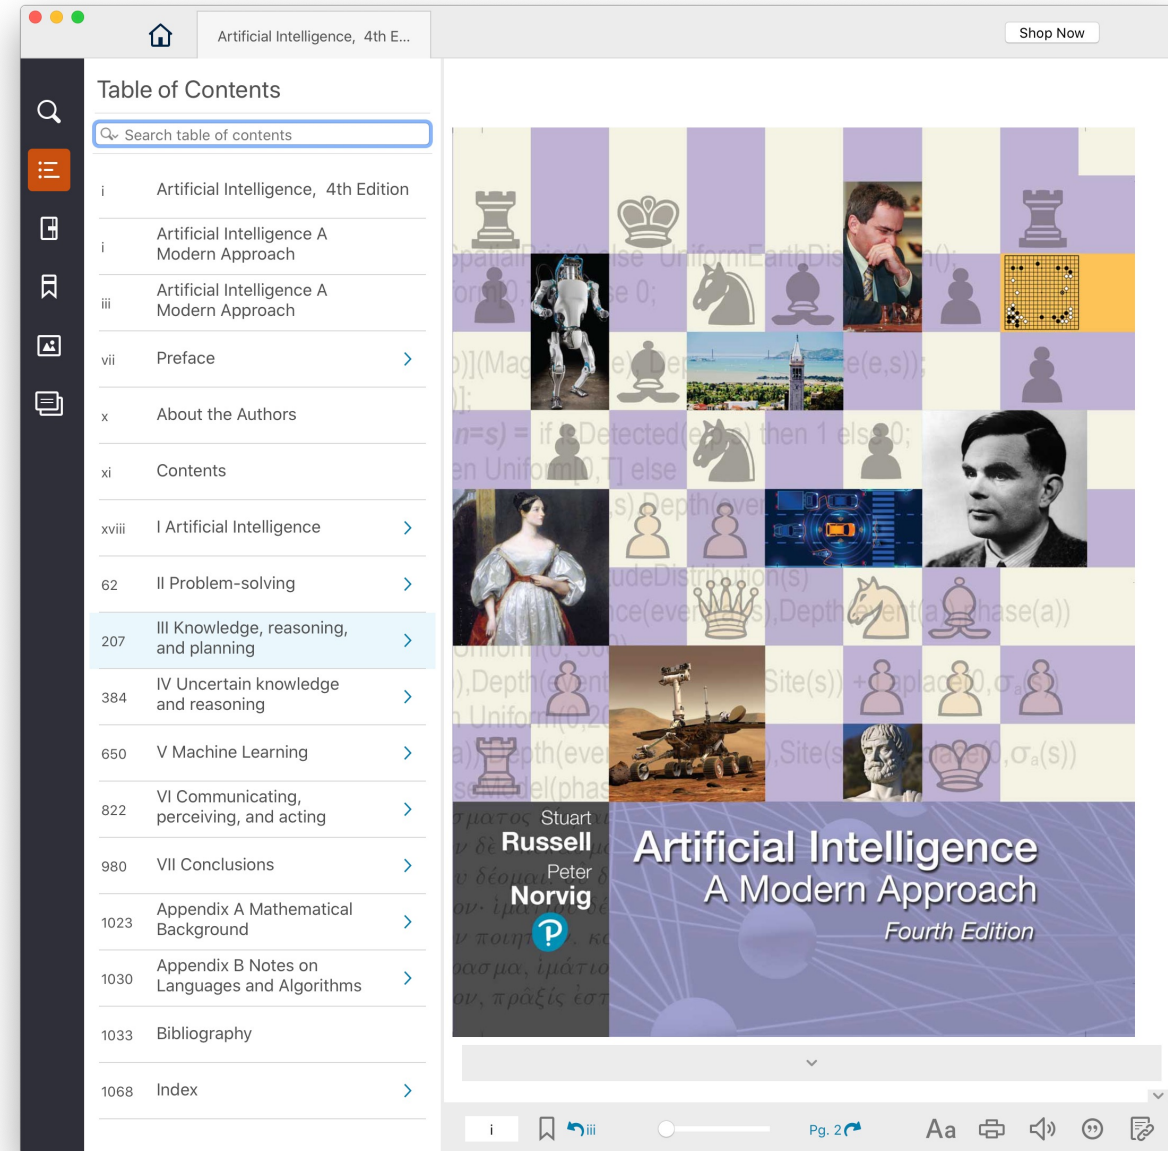
Wed 12-1pm,
and by appointment

Location: ITE 334

Administrivia

Text, CMI

- 4th edition of AIMA (2020) has lots of new material since the 2009 3rd edition
- UMBC CMI program charges
- \$?? for a digital version, \$160 on Amazon for hardcopy!
- Access on Blackboard and/or download to computer/phone
- May access copy on your system for many years
- Opt-out of CMI via Blackboard



Course Website

WWW

Schedule, slides,
assignments, readings,
materials, syllabus here

<https://umbc-cmsc471.github.io/fall2023-public/>

<https://campuswire.com/c/G38B2595B/feed>



- **Course announcements**, Q&A, discussion board here
- No public code, follow posted rules and etiquette

Course Website

WWW

Schedule, slides,
assignments, readings,
materials, syllabus here

<https://umbc-cmsc471.github.io/fall2023-public/>

<https://campuswire.com/c/G38B2595B/feed>

This Week



- **Course announcements**, Q&A, discussion board here
- No public code, follow posted rules and etiquette

Course Syllabus (On the Website)

CMSC 471 — Introduction to Artificial Intelligence

Fall 2023

- [CMSC 471 — Introduction to Artificial Intelligence](#)
 - [Logistics](#)
 - [Course Description](#)
 - [CMI Text Book](#)
 - [Prerequisites](#)
 - [Course Schedule](#)
 - [Assignments](#)
 - [Midterm and Final Exams](#)
 - [Course Evaluation](#)
 - [Policies](#)
 - [Due Dates](#)
 - [Extensions and Late Policy](#)
 - [Academic Honesty](#)
 - [Accommodations](#)
 - [Students with Accommodation Needs](#)
 - [Sexual Assault, Sexual Harassment, and Gender-based Violence and Discrimination](#)
 - [Hate, Bias, Discrimination, and Harassment](#)
 - [Acknowledgements](#)

Academic Integrity

- Super important: I take it **very** seriously
- **You** are responsible for your (& your group's) own work: if in doubt, ask!
- Penalties could include 0 on the assignment, course failure, suspension, or expulsion (not exhaustive)

Course Evaluation

Components	Percentage
Quiz (?)	10%
Assignments	45%
Exams (Midterm + Final)	40%
Course Engagement	5%

Course Evaluation

On Blackboard if we have them, mostly to motivate you to keep up with the reading and classes

Components	Percentage
Quiz (?)	10%
Assignments	45%
Exams (Midterm + Final)	40%
Course Engagement	5%

Course Evaluation

On Blackboard if we have them, mostly to motivate you to keep up with the reading and classes

Components	Percentage
Quiz (?)	10%
Assignments	45%
Exams (Midterm + Final)	40%
Course Engagement	5%

What is engagement?

Final Grades

\geq	Letter
90	A
80	B
70	C
60	D
0	F

Running the Assignments

A "standard" x86-64 Linux machine, like `g1` or the platform Google Colab `g1`

A passable amount of memory (2GB-4GB) [this is definitely the upper range]

Modern but not necessarily cutting edge software

If in doubt, ask first

Programming Languages for **Assignments**

Python, though individual assignments could vary

Remember: programming languages are *tools*. Don't get too caught up in not "knowing" a language. This course will not be grading software engineering prowess.

Libraries: Assignment dependent. Generally OK, as long as you don't use their implementation of what you need to implement

If in doubt, ask first

Late Policy

Everyone has a budget of 10 *late days*, maximum 3 per assignment

Late Policy

Everyone has a budget of 10 *late days*, maximum 3 per assignment

If you have them left: assignments turned in after the deadline will be graded and recorded, no questions asked

Late Policy

Everyone has a budget of 10 *late days*, maximum 3 per assignment

If you have them left: assignments turned in after the deadline will be graded and recorded, no questions asked

If you don't have any left: still turn assignments in. They could count in your favor in borderline cases

Late Policy

Everyone has a budget of 10 *late days*, maximum 3 per assignment

Use them as needed throughout the course

They're meant for personal reasons and **emergencies**

Do not procrastinate

Late Policy


Everyone has a budget of 10 *late days*, maximum 3 per assignment

Contact me privately if an extended absence will occur

You must know how
many you've used

Q: What is AI
(artificial intelligence)?

Sci-Fi AI?



Partially Adapted from: Pieter Abbeel

What is AI?

Q. What is artificial intelligence?

A. It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.

<http://www-formal.stanford.edu/jmc/whatisai/>

Courtesy: Tim Finin

Ok, so what is intelligence?

Q. Yes, but what is intelligence?

<http://www-formal.stanford.edu/jmc/whatisai/>

Courtesy: Tim Finin

Ok, so what is intelligence?

Q. Yes, but what is intelligence?


A. Intelligence is the computational part of the ability to **achieve goals** in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines

<http://www-formal.stanford.edu/jmc/whatisai/>

Courtesy: Tim Finin

Possible Approaches


The science of making machines that:



Possible Approaches

The science of making machines that:


Think like people




Possible Approaches

The science of making machines that:

Think like people




Act like people




Possible Approaches

The science of making machines that:


Think like people



Think rationally




Act like people




Possible Approaches

The science of making machines that:


Think like people




Think rationally



Act like people



Act rationally



Rational Decisions

We'll use the term **rational** in a very specific, technical way:

- Rational: maximally achieving pre-defined goals
- Rationality only concerns what decisions are made
(not the thought process behind them)
- Goals are expressed in terms of the **utility** of outcomes
- Being rational means **maximizing your expected utility**

Rational Decisions

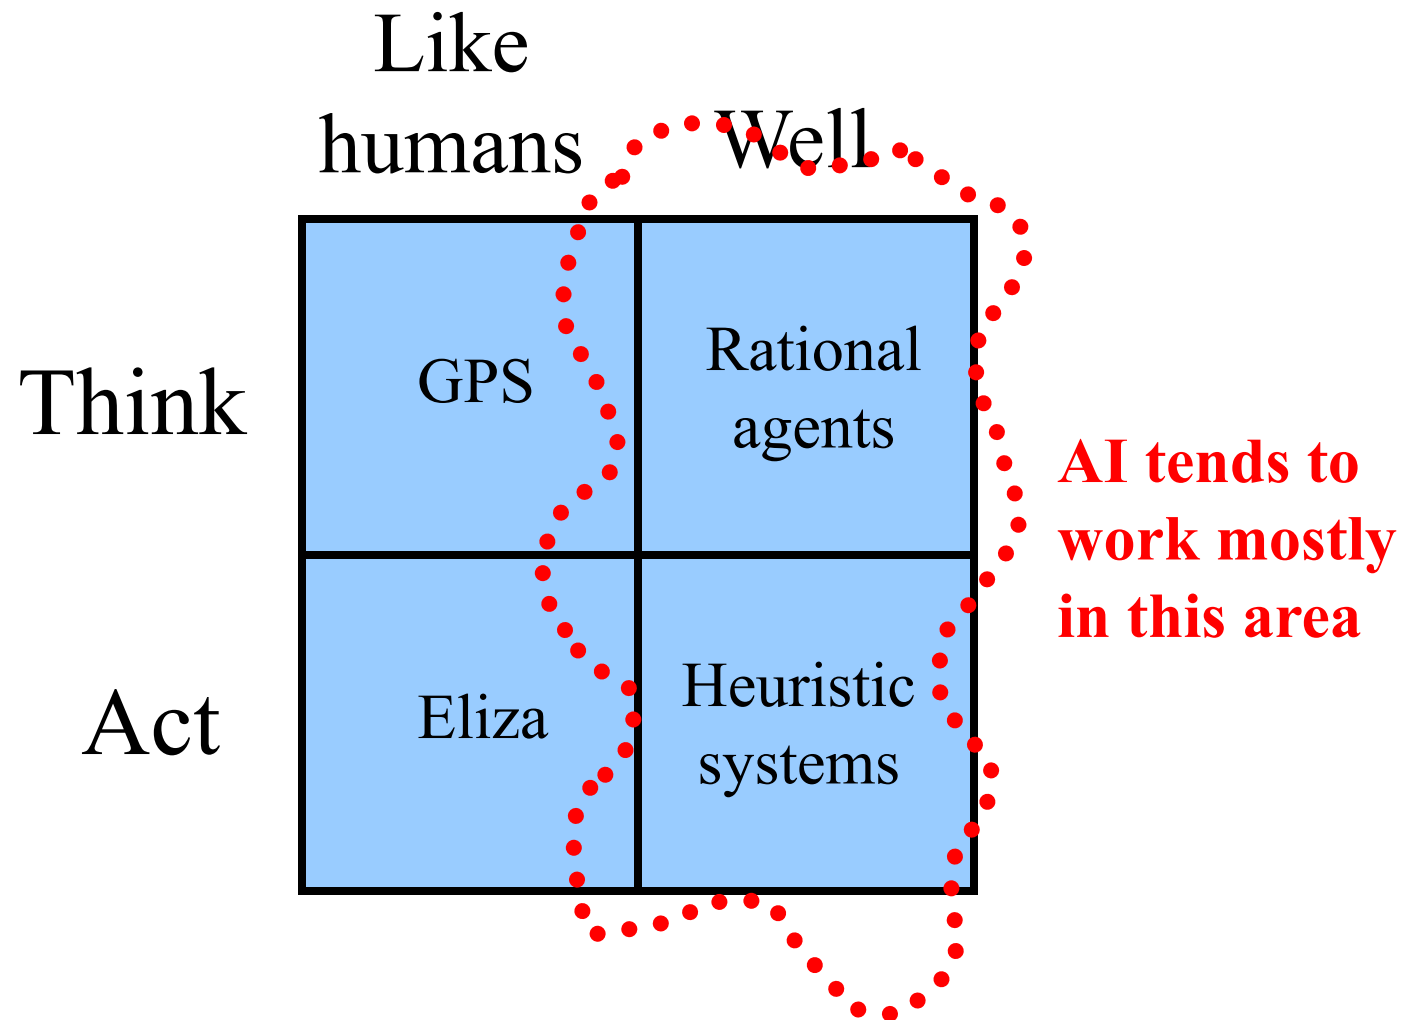
We'll use the term **rational** in a very specific, technical way:

- Rational: maximally achieving pre-defined goals
- Rationality only concerns what decisions are made
(not the thought process behind them)
- Goals are expressed in terms of the **utility** of outcomes
- Being rational means **maximizing your expected utility**

Possible Approaches

	Like humans	Well
Think	GPS	Rational agents
Act	Eliza	Heuristic systems

Possible Approaches



Courtesy: Tim Finin

Big questions

- Can machines think?
- Must/should they think like humans do?
- Can they learn from their experience?
- Can they have common sense?
- If so, how? If not, why not?
- What does this say about human beings?
- What does this say about the mind?



Why AI?

1 Engineering: get machines to do useful things

e.g., understand spoken natural language, recognize individual people in visual scenes, autonomous vehicles, better cybersecurity, ...


2 Cognitive Science: model and understand how natural minds and mental phenomena work

e.g., visual perception, memory, learning, language, decision making, ...

3 Philosophy: explore basic, interesting and important philosophical questions


e.g., mind-body problem, what's consciousness, free will, ...

A (Short) History of AI




A (Short) History of AI

- 1940-1950: Early days
 - 1943: McCulloch & Pitts: Boolean circuit model of brain
 - 1950: Turing's "Computing Machinery and Intelligence"




A (Short) History of AI

- 1940-1950: Early days
 - 1943: McCulloch & Pitts: Boolean circuit model of brain
 - 1950: Turing's "Computing Machinery and Intelligence"
- 1950—70: Excitement: Look, Ma, no hands!
 - 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
 - **1956: Dartmouth meeting: "Artificial Intelligence" adopted**
 - 1965: Robinson's complete algorithm for logical reasoning




A (Short) History of AI

- 1940-1950: Early days
 - 1943: McCulloch & Pitts: Boolean circuit model of brain
 - 1950: Turing's "Computing Machinery and Intelligence"
- 1950—70: Excitement: Look, Ma, no hands!
 - 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
 - **1956: Dartmouth meeting: "Artificial Intelligence" adopted**
 - 1965: Robinson's complete algorithm for logical reasoning
- 1970—90: Knowledge-based approaches
 - 1969—79: Early development of knowledge-based systems
 - 1980—88: Expert systems industry booms
 - 1988—93: Expert systems industry busts: "AI Winter"




A (Short) History of AI

- 1940-1950: Early days
 - 1943: McCulloch & Pitts: Boolean circuit model of brain
 - 1950: Turing's "Computing Machinery and Intelligence"
- 1950—70: Excitement: Look, Ma, no hands!
 - 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
 - **1956: Dartmouth meeting: "Artificial Intelligence" adopted**
 - 1965: Robinson's complete algorithm for logical reasoning
- 1970—90: Knowledge-based approaches
 - 1969—79: Early development of knowledge-based systems
 - 1980—88: Expert systems industry booms
 - 1988—93: Expert systems industry busts: "AI Winter"
- 1990—: Statistical approaches
 - Resurgence of probability, focus on uncertainty
 - General increase in technical depth
 - Agents and learning systems... "AI Spring"?



A (Short) History of AI

- 1940-1950: Early days
 - 1943: McCulloch & Pitts: Boolean circuit model of brain
 - 1950: Turing's "Computing Machinery and Intelligence"
- 1950—70: Excitement: Look, Ma, no hands!
 - 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
 - **1956: Dartmouth meeting: "Artificial Intelligence" adopted**
 - 1965: Robinson's complete algorithm for logical reasoning
- 1970—90: Knowledge-based approaches
 - 1969—79: Early development of knowledge-based systems
 - 1980—88: Expert systems industry booms
 - 1988—93: Expert systems industry busts: "AI Winter"
- 1990—: Statistical approaches
 - Resurgence of probability, focus on uncertainty
 - General increase in technical depth
 - Agents and learning systems... "AI Spring"?
- 2000—: Where are we now?




1956 Dartmouth Project

“We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.”

<http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>

Turing Test¹

- Three rooms: person, computer, interrogator
- Interrogator communicates with others via a ‘teleprinter’
- Interrogator tries to distinguish person and machine
- Machine tries to fool interrogator into believing it is the person
- Person acts naturally
- If machine succeeds, we conclude it can think



¹ [Turing, Alan](#) (October 1950), "Computing Machinery and Intelligence", *Mind*. (236): 433–460.


Hot Topics Today

- Neural networks again: Large Language Models, ChatGPT
 - Machine learning, data science
 - Exploiting big data
 - Autonomous vehicles, robotics
 - Text mining, spoken natural language understanding, Siri
 - Conversational AI, DuerOS
 - Computer vision, DALL-E
- **GPT-3 (OpenAI)**
 - **GPT-4 (OpenAI)**
 - **AlphaGo (Google DeepMind)**
 - **Watson (IBM)**
 - **Sophia (Hanson Robotics)**
 - **Tesla Autopilot (Tesla Inc)**

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Play a decent game of Jeopardy?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Walmart?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person **for an hour**?
- ✓ Perform a surgical operation?
- ✓ Put away the dishes and fold the laundry?
- ✗ Translate spoken Chinese into spoken English in real time?
 - Write an intentionally funny story?




Unintentionally Funny Stories

Unintentionally Funny Stories

- One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe walked to the oak tree. He ate the beehive. The End.


Unintentionally Funny Stories

- One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe walked to the oak tree. He ate the beehive. The End.




Unintentionally Funny Stories

- One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe walked to the oak tree. He ate the beehive. The End.
- Henry Squirrel was thirsty. He walked over to the river bank where his good friend Bill Bird was sitting. Henry slipped and fell in the river. Gravity drowned. The End.





Unintentionally Funny Stories

- One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe walked to the oak tree. He ate the beehive. The End.
- Henry Squirrel was thirsty. He walked over to the river bank where his good friend Bill Bird was sitting. Henry slipped and fell in the river. Gravity drowned. The End.
- Once upon a time there was a dishonest fox and a vain crow. One day the crow was sitting in his tree, holding a piece of cheese in his mouth. He noticed that he was holding the piece of cheese. He became hungry, and swallowed the cheese. The fox walked over to the crow. The End.



Unintentionally Funny Stories


- One day Joe Bear was hungry. He went to Irving Bird where some honey was. There was a beehive in the oak tree. He ate the bees.
- Henry Squirrel was thirsty. He went to the river bank where his good friend was. Henry slipped and fell in the water. The End.
- Once upon a time there was a fox. He was sitting in his tree, holding a piece of cheese. He noticed that he was holding the piece of cheese. The fox walked over.



vain crow. One day the crow was sitting in his tree, holding a piece of cheese. He noticed that he was holding the piece of cheese. The fox walked over.

Unintentionally Funny Stories


- One day Joe Bear was hungry. He went to Irving Bird where some honey was. There was a beehive in the tree. He ate the bees.
- Henry Squirrel was thirsty. He went to the river bank where his good friend was. Henry slipped and fell in the water. The End.
- Once upon a time there was a fox. He was sitting in his tree, holding a piece of cheese. He was holding the piece of cheese. The fox walked over.



vain crow. One day the crow was sitting in his tree, holding a piece of cheese. He noticed that he was holding the piece of cheese. The fox walked over.


Natural Language

- Speech technologies (e.g. Siri)
 - Automatic speech recognition (ASR)
 - Text-to-speech synthesis (TTS)
 - Dialog systems




Natural Language

- Speech technologies (e.g. Siri)
 - Automatic speech recognition (ASR)
 - Text-to-speech synthesis (TTS)
 - Dialog systems
- Language processing technologies




Natural Language

- Speech technologies (e.g. Siri)
 - Automatic speech recognition (ASR)
 - Text-to-speech synthesis (TTS)
 - Dialog systems
- Language processing technologies
 - Question answering



Natural Language

- Speech technologies (e.g. Siri)
 - Automatic speech recognition (ASR)
 - Text-to-speech synthesis (TTS)
 - Dialog systems
- Language processing technologies
 - Question answering
 - Machine translation



"Il est impossible aux journalistes de rentrer dans les régions tibétaines"

Bruno Philip, correspondant du "Monde" en Chine, estime que les journalistes de l'AFP qui ont été expulsés de la province tibétaine du Qinghai "n'étaient pas dans l'illégalité".

Les faits Le dalaï-lama dénonce l'"enfer" imposé au Tibet depuis sa fuite, en 1959

Vidéo Anniversaire de la rébellion tibétaine - la Chine sur ses gardes



"It is impossible for journalists to enter Tibetan areas"

Philip Bruno, correspondent for "World" in China, said that journalists of the AFP who have been deported from the Tibetan province of Qinghai "were not illegal."


Facts The Dalai Lama denounces the "hell" imposed since he fled Tibet in 1959

Video Anniversary of the Tibetan rebellion: China on guard



Natural Language

- Speech technologies (e.g. Siri)
 - Automatic speech recognition (ASR)
 - Text-to-speech synthesis (TTS)
 - Dialog systems



- Language processing technologies
 - Question answering
 - Machine translation

"Il est impossible aux journalistes de rentrer dans les régions tibétaines"

Bruno Philip, correspondant du "Monde" en Chine, estime que les journalistes de l'AFP qui ont été expulsés de la province tibétaine du Qinghai "n'étaient pas dans l'illégalité".

Les faits Le dalaï-lama dénonce l'"enfer" imposé au Tibet depuis sa fuite, en 1959

Vidéo Anniversaire de la rébellion tibétaine - La Chine sur ses gardes



"It is impossible for journalists to enter Tibetan areas"

Philip Bruno, correspondent for "World" in China, said that journalists of the AFP who have been deported from the Tibetan province of Qinghai "were not illegal."

Facts The Dalai Lama denounces the "hell" imposed since he fled Tibet in 1959



Video Anniversary of the Tibetan rebellion: China on guard



- Web search
- Text classification, spam filtering, etc...

Vision (Perception)

- Object and face recognition
- Scene segmentation
- Image classification



Demo: [YOLO obj-detect](#)

Images from Erik Sudderth (left), wikipedia (right)

Courtesy: Pieter Abbeel

Robotics



- Robotics
 - Part mech. eng.
 - Part AI
 - Reality much harder than simulations!
- Technologies
 - Vehicles
 - Rescue
 - Soccer!
 - Lots of automation...
- In this class:
 - We ignore mechanical aspects
 - Methods for planning
 - Methods for control

Demo 1: [ROBOTICS – soccer.avi](#)

Demo 4: [ROBOTICS – laundry.avi](#)

Demo 3: [ROBOTICS – gcar.avi](#)


Demo 5: [ROBOTICS – petman.avi](#)



Images from UC Berkeley, Boston Dynamics, RoboCup, Google


Logic

- Logical systems
 - Theorem provers
 - NASA fault diagnosis
 - Question answering
- Methods:
 - Deduction systems
 - Constraint satisfaction
 - Satisfiability solvers (huge advances!)




Game Playing

- Classic Moment: May, '97: Deep Blue vs. Kasparov
 - First match won against world champion
 - “Intelligent creative” play
 - 200 million board positions per second
 - Humans understood 99.9 of Deep Blue's moves
 - Can do about the same now with a PC cluster




Game Playing

- Classic Moment: May, '97: Deep Blue vs. Kasparov
 - First match won against world champion
 - “Intelligent creative” play
 - 200 million board positions per second
 - Humans understood 99.9 of Deep Blue's moves
 - Can do about the same now with a PC cluster
- Open question:
 - How does human cognition deal with the search space explosion of chess?
 - Or: how can humans compete with computers at all??



Game Playing

- Classic Moment: May, '97: Deep Blue vs. Kasparov
 - First match won against world champion
 - “Intelligent creative” play
 - 200 million board positions per second
 - Humans understood 99.9 of Deep Blue's moves
 - Can do about the same now with a PC cluster
- Open question:
 - How does human cognition deal with the search space explosion of chess?
 - Or: how can humans compete with computers at all??
- 1996: Kasparov Beats Deep Blue
“I could feel --- I could smell --- a new kind of intelligence across the table.”



Game Playing

- Classic Moment: May, '97: Deep Blue vs. Kasparov
 - First match won against world champion
 - “Intelligent creative” play
 - 200 million board positions per second
 - Humans understood 99.9 of Deep Blue's moves
 - Can do about the same now with a PC cluster
- Open question:
 - How does human cognition deal with the search space explosion of chess?
 - Or: how can humans compete with computers at all??
- 1996: Kasparov Beats Deep Blue
“I could feel --- I could smell --- a new kind of intelligence across the table.”



Game Playing

- Classic Moment: May, '97: Deep Blue vs. Kasparov
 - First match won against world champion
 - “Intelligent creative” play
 - 200 million board positions per second
 - Humans understood 99.9 of Deep Blue's moves
 - Can do about the same now with a PC cluster
- Open question:
 - How does human cognition deal with the search space explosion of chess?
 - Or: how can humans compete with computers at all??
- 1996: Kasparov Beats Deep Blue
“I could feel --- I could smell --- a new kind of intelligence across the table.”
- 1997: Deep Blue Beats Kasparov
“Deep Blue hasn't proven anything.”



Game Playing

- Classic Moment: May, '97: Deep Blue vs. Kasparov
 - First match won against world champion
 - “Intelligent creative” play
 - 200 million board positions per second
 - Humans understood 99.9 of Deep Blue's moves
 - Can do about the same now with a PC cluster
- Open question:
 - How does human cognition deal with the search space explosion of chess?
 - Or: how can humans compete with computers at all??
- 1996: Kasparov Beats Deep Blue
“I could feel --- I could smell --- a new kind of intelligence across the table.”
- 1997: Deep Blue Beats Kasparov
“Deep Blue hasn't proven anything.”
- Huge game-playing advances recently, e.g. in Go!



Game Playing

- Classic Moment: May, '97: Deep Blue vs. Kasparov
 - First match won against world champion
 - “Intelligent creative” play
 - 200 million board positions per second
 - Humans understood 99.9 of Deep Blue's moves
 - Can do about the same now with a PC cluster
- Open question:
 - How does human cognition deal with the search space explosion of chess?
 - Or: how can humans compete with computers at all??
- 1996: Kasparov Beats Deep Blue
“I could feel --- I could smell --- a new kind of intelligence across the table.”
- 1997: Deep Blue Beats Kasparov
“Deep Blue hasn't proven anything.”
- Huge game-playing advances recently, e.g. in Go!


In March 2016, it beat Lee Sedol in a five-game match, the first time a computer Go program has beaten a 9-dan professional without handicap



Decision Making

- Applied AI involves many kinds of automation

- Scheduling, e.g. airline routing, military
- Route planning, e.g. Google maps
- Medical diagnosis
- Web search engines
- Spam classifiers
- Automated help desks
- Fraud detection
- Product recommendations
- ... Lots more!



What can't AI systems do yet? (Fully? Completely?)

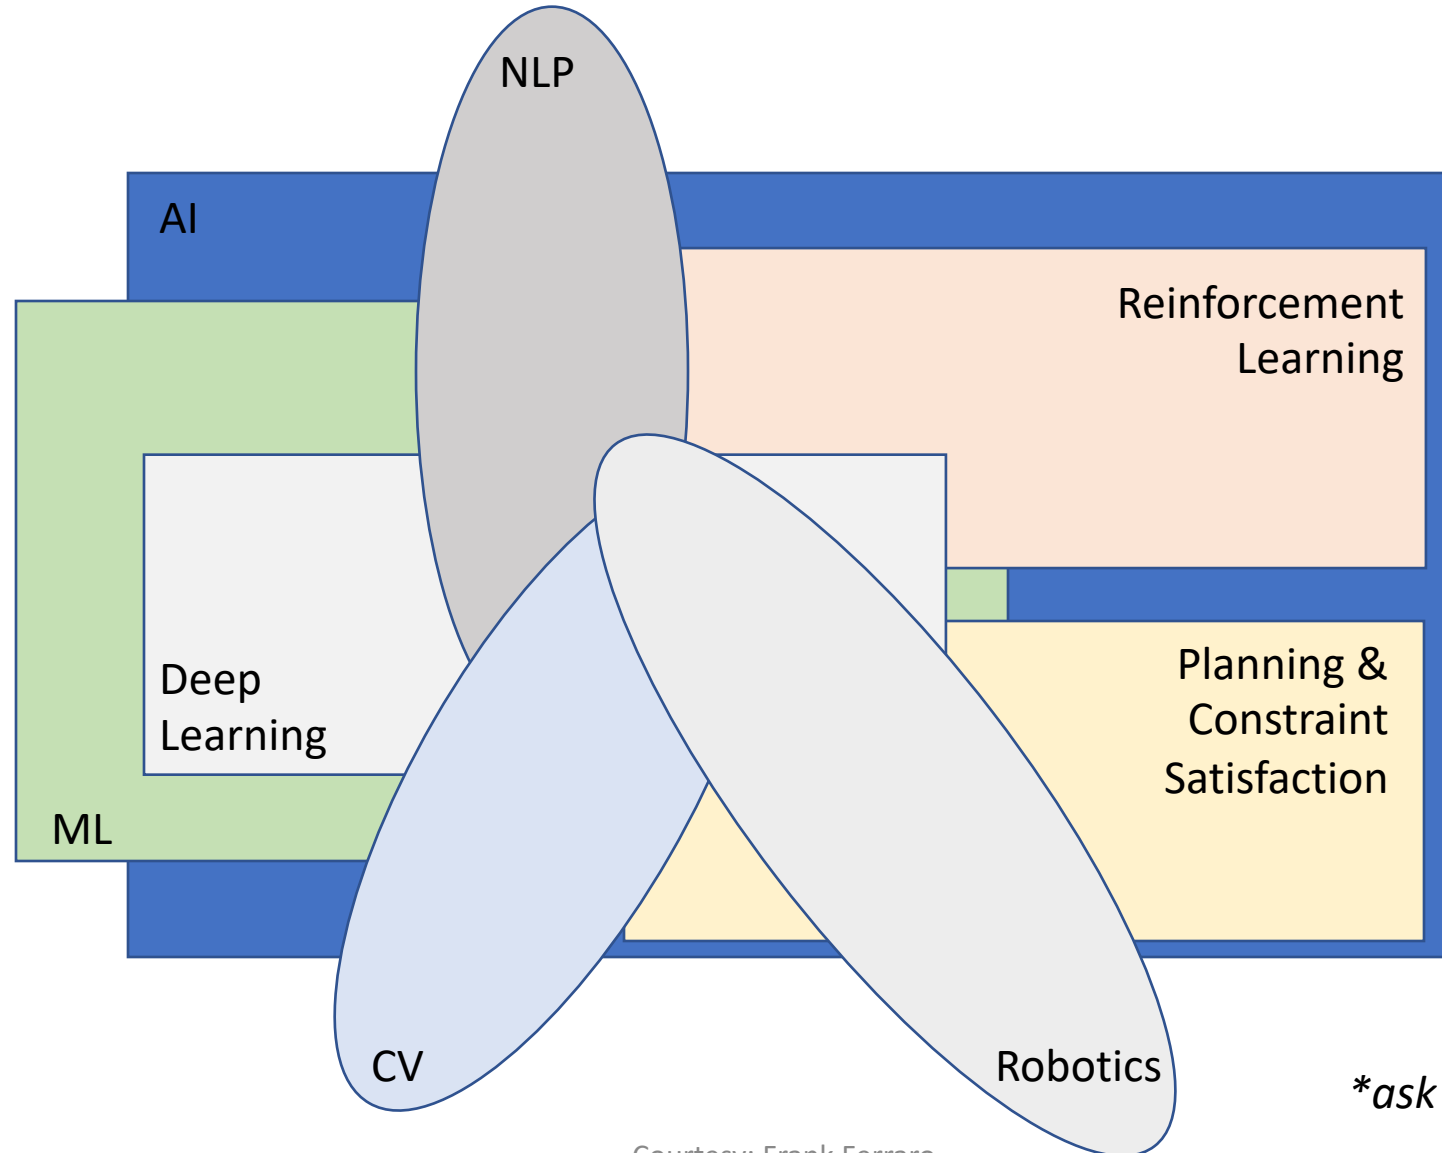
- Understand natural language robustly (e.g., read and understand articles in a newspaper)
- Surf the web and find interesting knowledge
- Interpret an arbitrary visual scene
- Learn a natural language
- ~~Play Go well~~
- Construct plans in dynamic real-time domains
- Refocus attention in complex environments
- Perform life-long learning

What can't AI systems do yet? (Fully? Completely?)

- Understand natural language robustly (e.g., read and understand articles in a newspaper)
- Surf the web and find interesting knowledge
- Interpret an arbitrary visual scene
- Learn a natural language
- ~~Play Go well~~
- Construct plans in dynamic real-time domains
- Refocus attention in complex environments
- Perform life-long l

**Exhibit true autonomy
and intelligence!**

A Partial, Incomplete View of AI, as a diagram of keywords*



Courtesy: Frank Ferraro

**ask 10 people, get
11+ opinions*

Course Goals

- Be introduced to some of the core problems and solutions of AI (big picture)
- Learn different ways that success and progress can be measured in AI

Course Goals

- Be introduced to some of the core problems and solutions of AI (big picture)
- Learn different ways that success and progress can be measured in AI
- **Relate to other CS fields/disciplines**

Course Goals

- Be introduced to some of the core problems and solutions of AI (big picture)
- Learn different ways that success and progress can be measured in AI
- Relate to other CS fields/disciplines
- **Implement AI programs**

Survey

- There will be a survey posted in Blackboard
- The purpose of this **anonymous survey** is to find what materials we should focus on and how would I design later part of the syllabus
- Sample questions:
 - Why are you taking this course?
 - Have you developed AI/ML projects before?
 - Which courses are you taking from AI and ML this semester?