PURDUE UNIVERSITY

# INTRO TO ARTIFICIAL INTELLIGENCE

**ECE 47300, Spring 2024**

LILY G126

**Instructor: David I. Inouye**

**Website:** https://www.davidinouye.com/course/ece47300-spring-2024/

# PAIRS INTERACTION

- Find a nearby buddy (ideally pairs but okay with group of 3)
- Format of short interaction
  - 1st person answers question
  - “Switch”
  - 2nd person answers question
  - I will double clap to end
  - When I ask for answers, you can only answer with what the ***other*** person said
- First activity: Introduce yourself and say why you chose 47300 this semester

PURDUE UNIVERSITY

# WHAT IS ARTIFICIAL INTELLIGENCE?

- Discuss in pairs

PURDUE UNIVERSITY

# WHAT IS ARTIFICIAL INTELLIGENCE?

- Merriam-Webster Dictionary
  - “a branch of computer science dealing with the **simulation of intelligent behavior** in computers”
- Oxford Dictionary
  - “the theory and development of computer systems able to perform tasks that **normally require human intelligence**, such as visual perception, speech recognition, decision-making, and translation between languages.”

PURDUE UNIVERSITY

# POSSIBLY A.I. IS A MOVING TARGET

- Are these A.I.?
  - Chess solvers
  - TurboTax
  - Chatbot
- What about these?
  - Speech recognition
  - Face recognition
  - Personalized recommendations
  - Self-driving cars
- “AI is anything that humans can do that computers cannot **yet** do.”

PURDUE UNIVERSITY

# A MORE PRAGMATIC DEFINITION

"AI is that which appears in academic conferences on AI" *

* From slides by Prof. Zico Kolter at CMU: http://www.cs.cmu.edu/~./15780/slides/intro.pdf)

# A.I. IS A VERY BROAD FIELD

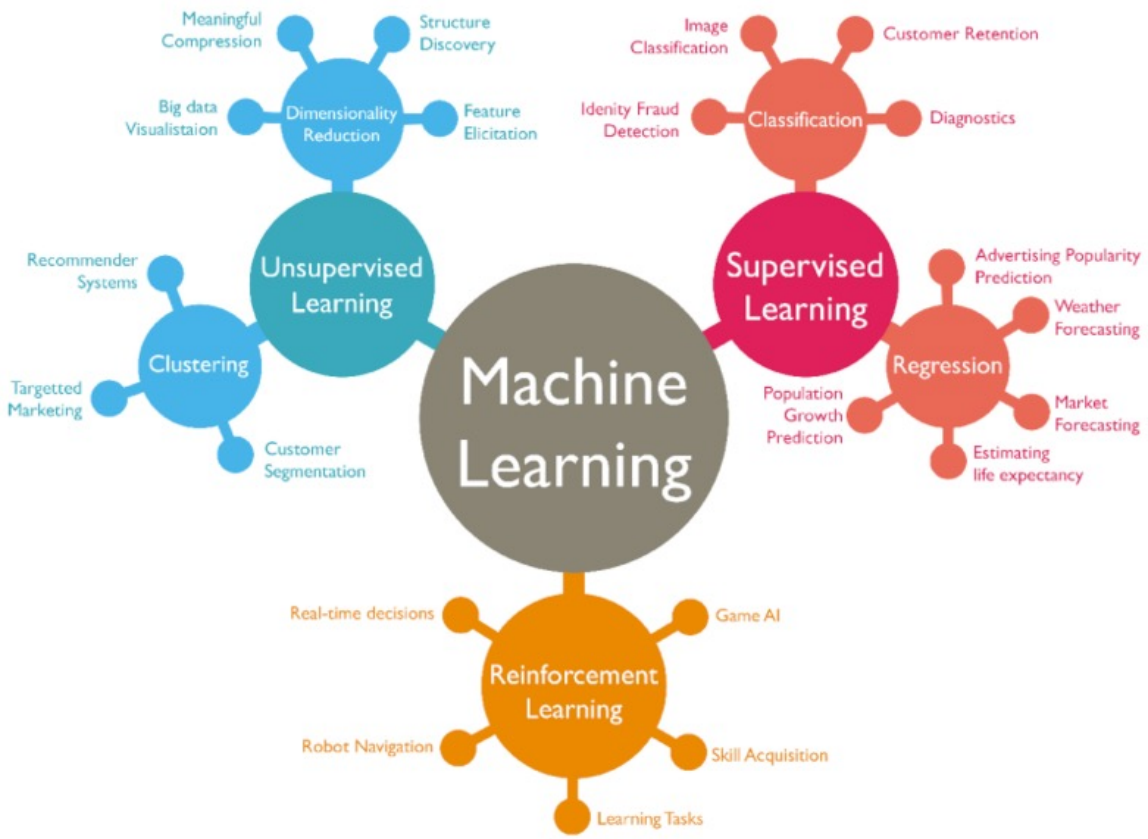


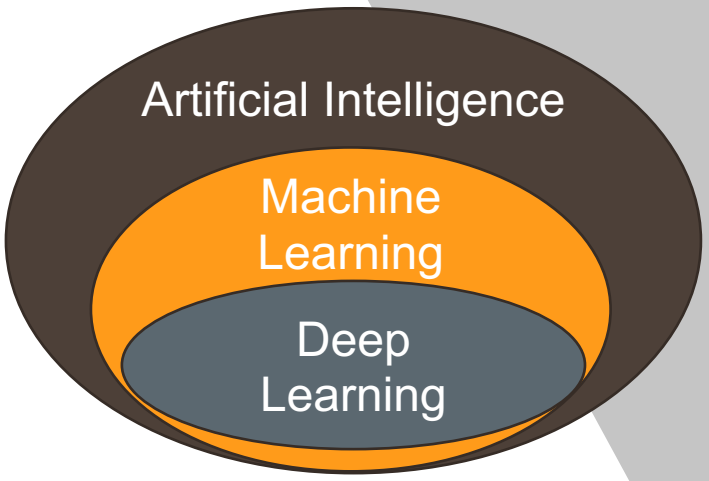


Major applications

- Computer vision
- Natural language processing
- Robotics

Original credit unknown, retrieved from https://www.sharper.ai/taxonomy-ai/

PURDUE UNIVERSITY

# THIS COURSE WILL ONLY COVER A SMALL SET OF TOPICS

1. Introduction to artificial intelligence
2. Machine learning basics
3. Deep learning basics
4. Natural language processing basics
5. Dimensionality reduction
6. Generative models
7. Markov decision processes
8. Special topics

PURDUE UNIVERSITY

# OTHER RELATED CLASSES

- ECE 57000: Artificial Intelligence (usually fall, project-based)
- ECE 50024: Machine Learning I (fall – online, spring – in-person) by Prof. Stanley Chan and Prof. Qi Quo
- ECE 59500: Intro. to Data Mining by Prof. Jing Gao
- ECE 59500 : Reinforcement Learning by Prof. Mahsa Ghasemi
- ECE 60146: Deep Learning (spring semester) by Prof. Kak and Prof. Bouman
- ECE 69500: Machine Learning in Bioinformatics and Healthcare by Prof. Joy Wang
- ECE 69500: Probabilistic Causal Inference by Prof. Murat Kocaoglu
- ECE 69500: Optimization in deep learning by Prof. Abolfazl Hashemi

PURDUE UNIVERSITY

# AGENDA

1. Announcements
   - Emergency preparedness
2. Syllabus
3. Logistics
4. (time-permitting) Defining AI

# EMERGENCY PREPAREDNESS

As we begin this semester, I want to take a few minutes and discuss emergency preparedness. Purdue University is a very safe campus and there is a low probability that a serious incident will occur here at Purdue. However, just as we receive a “safety briefing” each time we get on an aircraft, we want to emphasize our emergency procedures for evacuation and shelter-in-place incidents. Our preparedness will be critical IF an unexpected event occurs!

# EMERGENCY PREPAREDNESS

Emergency preparedness is your personal responsibility. Purdue University is actively preparing for natural disasters or human-caused incidents with the ultimate goal of maintaining a safe and secure campus. Let's review the following procedure

- For any emergency text or call 911.
- There are more than 300 Emergency Telephones (aka blue lights) throughout campus that connect directly to the Purdue Police Department (PUPD). If you feel threatened or need help, push the button and you will be connected right away.
- If we hear a fire alarm, we will immediately evacuate the building and proceed to the grassy area **north of Smith Hall**. Do not use the elevator.

# EMERGENCY PREPAREDNESS

- If we are notified of a Shelter in Place requirement for a **tornado warning, active shooter, or hazardous waste,** we will stop classroom or research activities and shelter in **this classroom**.

PURDUE UNIVERSITY

# SYLLABUS!

All significant updates to the syllabus or course schedule will be posted on **Piazza**.

See **course website** for syllabus and schedule.

# NOW BACK TO AI...

PURDUE UNIVERSITY

# A MORE PRAGMATIC DEFINITION

"AI is that which appears in academic conferences on AI" *

* From slides by Prof. Zico Kolter at CMU: http://www.cs.cmu.edu/~./15780/slides/intro.pdf)

PURDUE UNIVERSITY

# NOT "MOVIE" A.I.

*C-3PO and R2D2 in Star Wars.*


*The rogue A.I. HAL9000 from the movie 2001: A Space Odyssey.*


*Skynet from Terminator.*

PURDUE UNIVERSITY

# COMPUTERS DON'T "THINK" LIKE WE DO

```
[[5.49e-04 7.15e-04 5.60e-03 1.35e-02 9.42e-03 1.65e-03 4.38e-04 8.92e-04]
[9.64e-04 3.83e-04 1.38e-02 1.55e-02 1.06e-02 1.59e-02 5.07e-03 8.71e-05]
[2.02e-05 3.83e-03 1.58e-02 2.87e-03 9.79e-04 1.18e-02 8.46e-03 7.81e-04]
[1.18e-04 4.64e-03 1.21e-02 9.45e-04 5.22e-04 8.41e-03 8.26e-03 7.74e-04]
[4.56e-04 5.57e-03 8.02e-03 6.18e-04 6.12e-04 9.62e-03 8.94e-03 6.82e-04]
[3.60e-04 4.44e-03 1.17e-02 6.02e-05 1.67e-03 1.27e-02 7.21e-03 1.29e-04]
[3.15e-04 2.36e-03 1.46e-02 5.44e-03 1.10e-02 1.21e-02 2.09e-04 1.61e-04]
[6.53e-04 2.53e-04 6.47e-03 1.32e-02 1.02e-02 1.10e-04 6.56e-04 1.38e-04]]
```

A matrix of numbers as the computer "sees".
Do you know what this matrix represents?

PURDUE UNIVERSITY

# COMPUTERS DON'T "THINK" LIKE WE DO

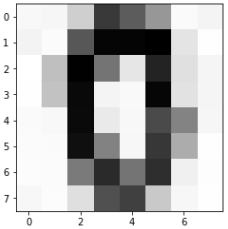

The same matrix of numbers displayed as an image.

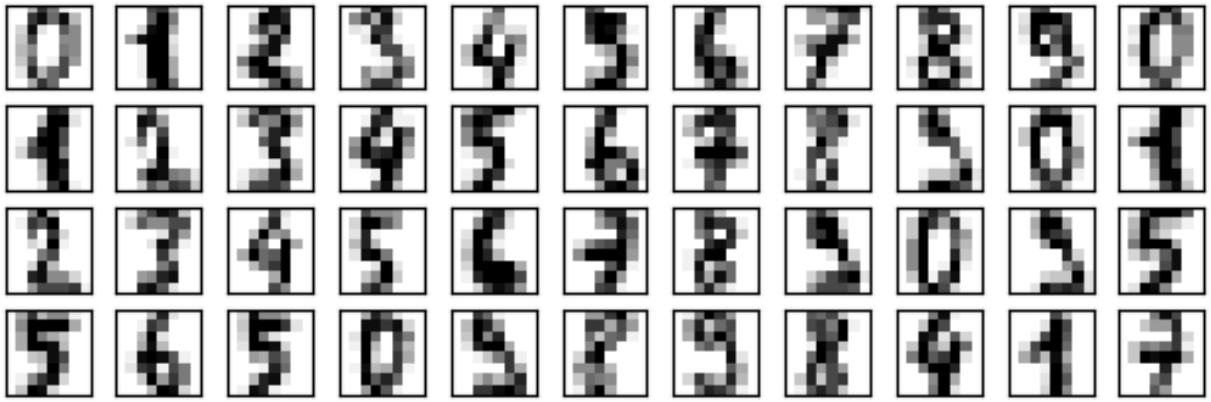

Other examples from this dataset.
Do you know what the numbers represent now?

# COMPUTERS DON'T "THINK" LIKE WE DO

$x$

"panda"
57.7% confidence

$+ .007 \times$

$\text{sign}(\nabla_x J(\boldsymbol{\theta}, \boldsymbol{x}, y))$

"nematode"
8.2% confidence

$=$

$\boldsymbol{x} + \epsilon\text{sign}(\nabla_x J(\boldsymbol{\theta}, \boldsymbol{x}, y))$

"gibbon"
99.3 % confidence

Real gibbon…

* From Explaining and Harnessing Adversarial Examples by Goodfellow et al.

# BUT DOESN'T CHATGPT SOLVE THIS?

## Still debated in the community!

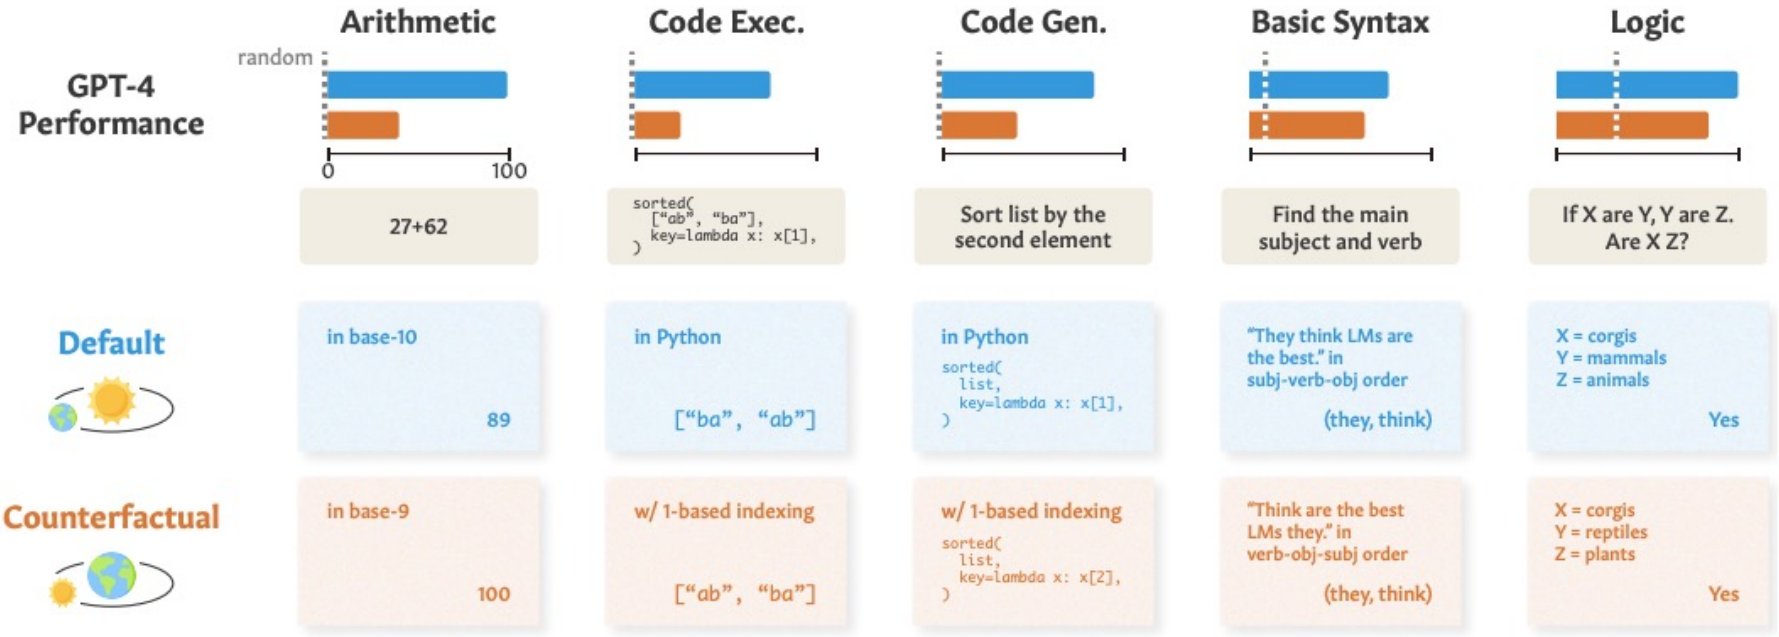


Figure from: Wu, Z., Qiu, L., Ross, A., Akyürek, E., Chen, B., Wang, B., ... & Kim, Y. (2023). Reasoning or reciting? exploring the capabilities and limitations of language models through counterfactual tasks. *arXiv preprint arXiv:2307.02477*.

# NOR IS HUMAN IMITATION NECESSARILY THE GOAL

- Consider flight

Imitation

Understanding + Engineering (Underlying principles)

PURDUE UNIVERSITY

# DARPA'S PERSPECTIVE ON AI

- First wave – Handcrafted Knowledge
- Second wave – Statistical Learning
- Third wave (future?) – Contextual Adaptation

Excellent DARPA video (16 min) on AI (content above based on this):
https://www.youtube.com/watch?v=-O01G3tSYpU

# 1ST WAVE: HANDCRAFTED KNOWLEDGE

- Examples
  - Auto-scheduling systems
  - Chess
  - TurboTax
  - Simple medical diagnosis
- Good at
  - Logical reasoning from rules
- Bad at
  - Perceiving
  - Learning

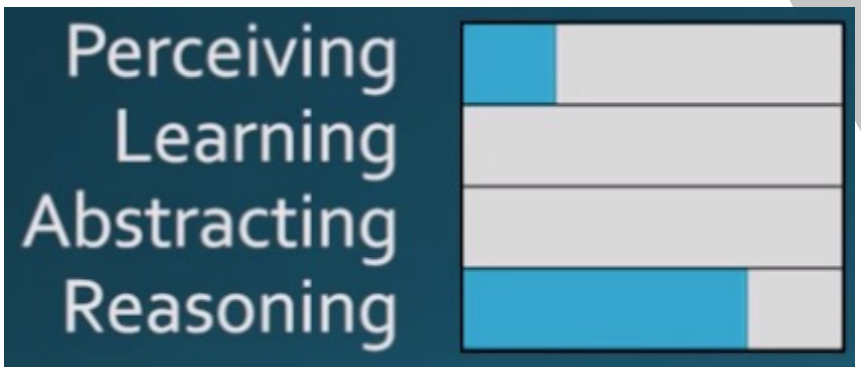


Excellent DARPA video (16 min) on AI (content above based on this):
https://www.youtube.com/watch?v=-O01G3tSYpU

PURDUE UNIVERSITY

# ONE KEY PROBLEM:
# (FORMALLY) STATING WHY IS HARD
# BUT LABELING IS EASY

- What is this a photo of?

https://unsplash.com/search/photos/cute-cat

- **Why** is this an image of a cat?
- Can you state a rule for all cats?

# 2ND WAVE: STATISTICAL LEARNING

- Examples
  - Voice/face recognition
  - Personalized recommendations
- Good at
  - Perceiving
  - Learning
- Bad at
  - Logical reasoning
  - Robustness


Excellent DARPA video (16 min) on AI (content above based on this):
https://www.youtube.com/watch?v=-O01G3tSYpU

# KEY PROBLEMS WITH 2ND WAVE

- Lack of robustness / fragile systems

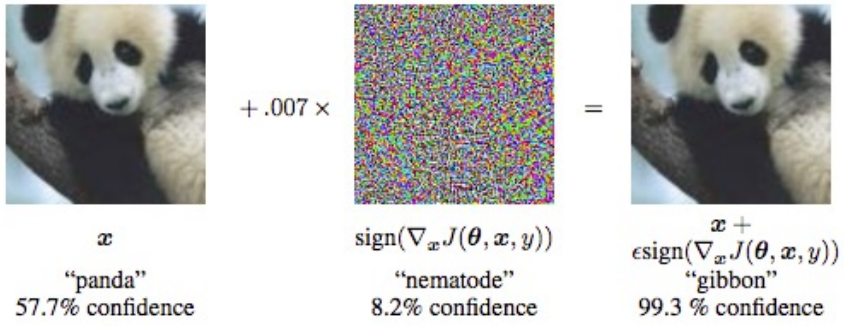



- Lack of explanations

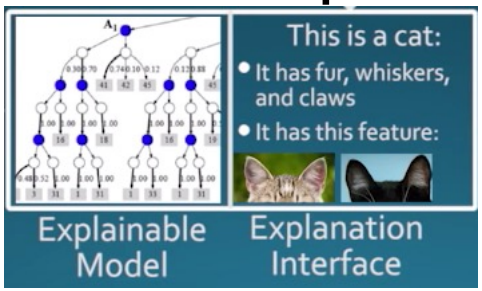


I understand why
I understand why not
I know when you'll succeed
I know when you'll fail
I know when to trust you

Excellent DARPA video (16 min) on AI (content above from): https://www.youtube.com/watch?v=-O01G3tSYpU

# 3RD WAVE (FUTURE?): CONTEXTUAL ADAPTATION

- Combination of previous two waves
- System will construct **explanatory models**
  - Causation
  - Some abstraction
  - Explainable
- Examples
  - Incorporate handwriting knowledge for recognizing new characters
  - ? - (Maybe you can work on this)
- Good at
  - Perceiving
  - Learning
  - Reasoning
- Slightly better at
  - Abstracting

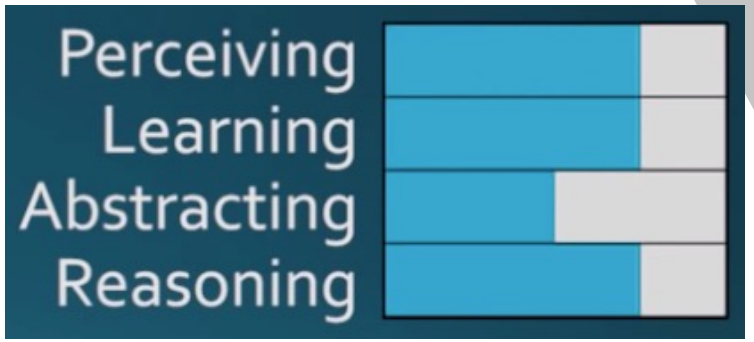


Excellent DARPA video (16 min) on AI (content above based on this):
https://www.youtube.com/watch?v=-O01G3tSYpU

# MODELING, INFERENCE, AND LEARNING PARADIGM OF AI

1. **Modeling** – How do you model the real-world problem?
   - How should you represent the **input to the model**? (e.g., string, vectors)
   - How to represent the **model functions**? (e.g., explicit, implicit, deep)
   - How do you represent the **world**? (e.g., simulation, probabilistic model)

   Critical for robotics and reinforcement learning

2. **Inference** – Given a model, how do you perform various tasks or answer questions based on the model?
   - How do you categorize photos?
   - How do you generate novel fake images?
   - How do you determine the next action to take to reach a goal?

   Increasingly important for AI using foundation models (e.g. ChatGPT, Dall-E)

3. **Learning** – How do you train/adapt the models given experience or data?
   - Which objective function are you optimizing?
   - What are the learning constraints?
   - Which algorithm do you use to find the model parameters?

   Critical component for modern advances and focus of this class

PURDUE
UNIVERSITY

# QUESTIONS?