



# **CS 540 Introduction to Artificial Intelligence**

## **Course Overview**

University of Wisconsin-Madison

Spring 2023



**Hello, welcome. I'm your 540 professor Jerry Zhu.**

**I cannot stay after class today, so if you want to chat, come to the podium now :-)**





# Today's outline

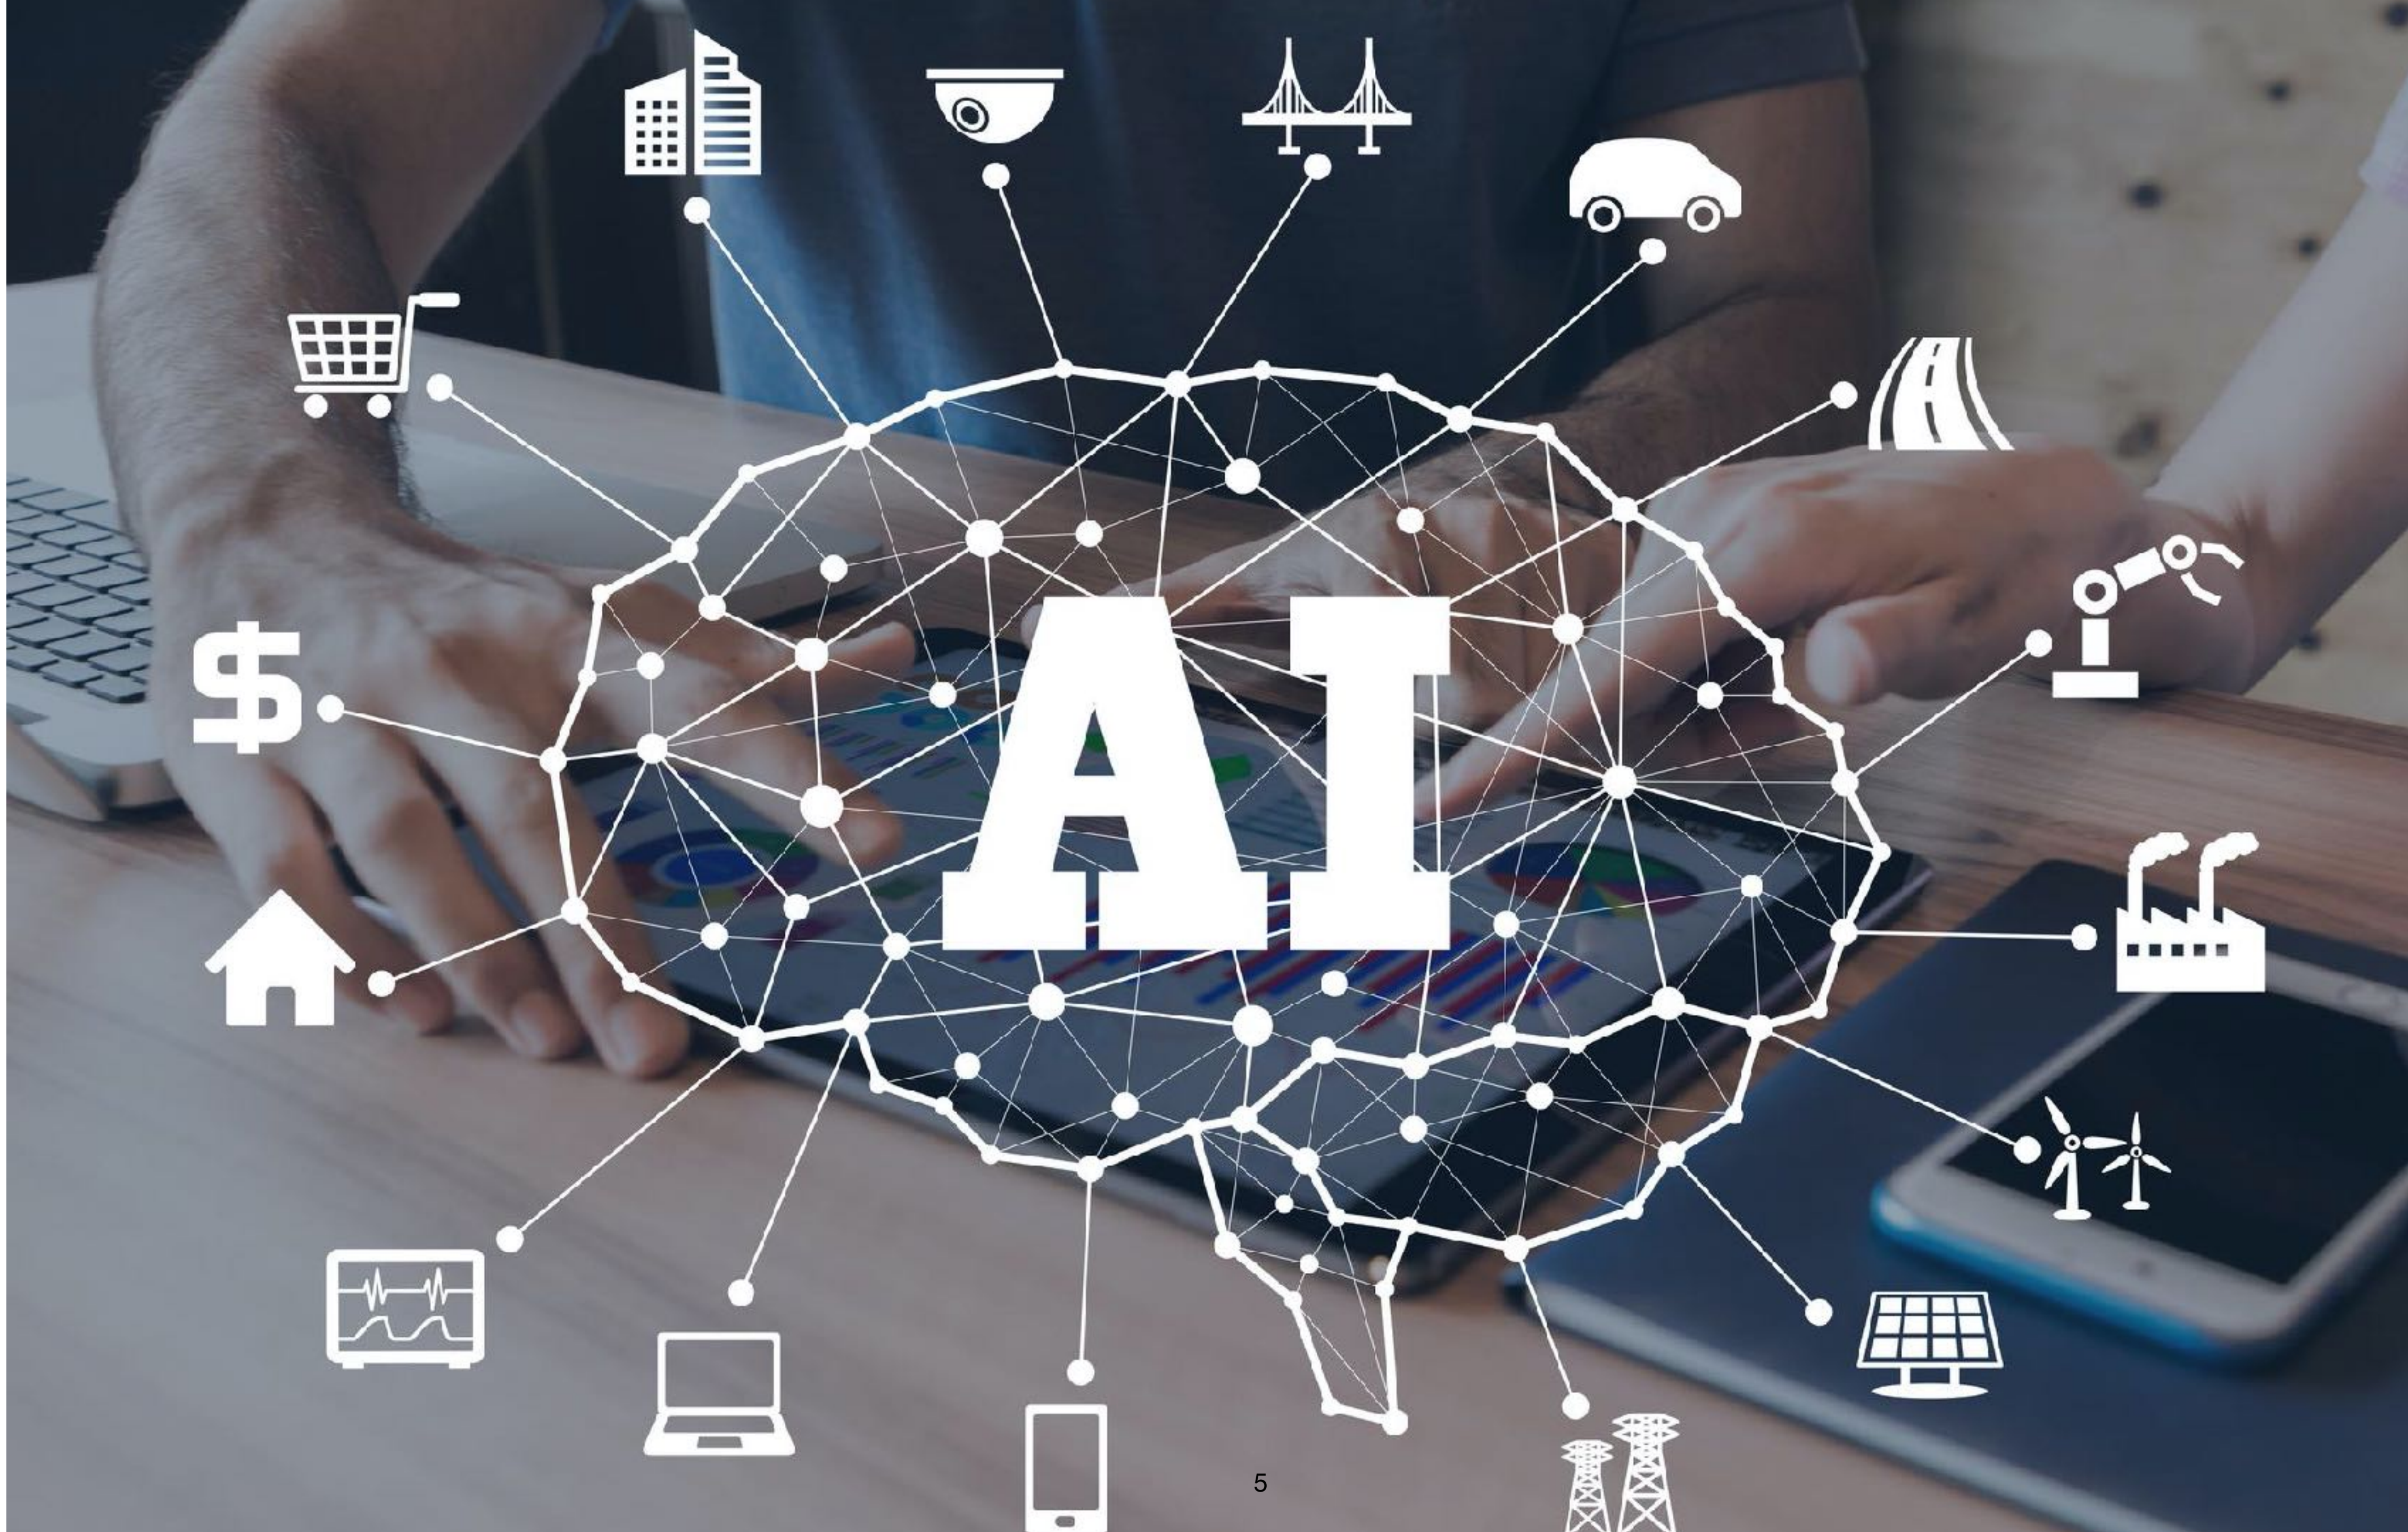
- What's in CS540
- Course logistics
- Software





# Part I: Course overview







# Classify Images

<http://www.image-net.org/>

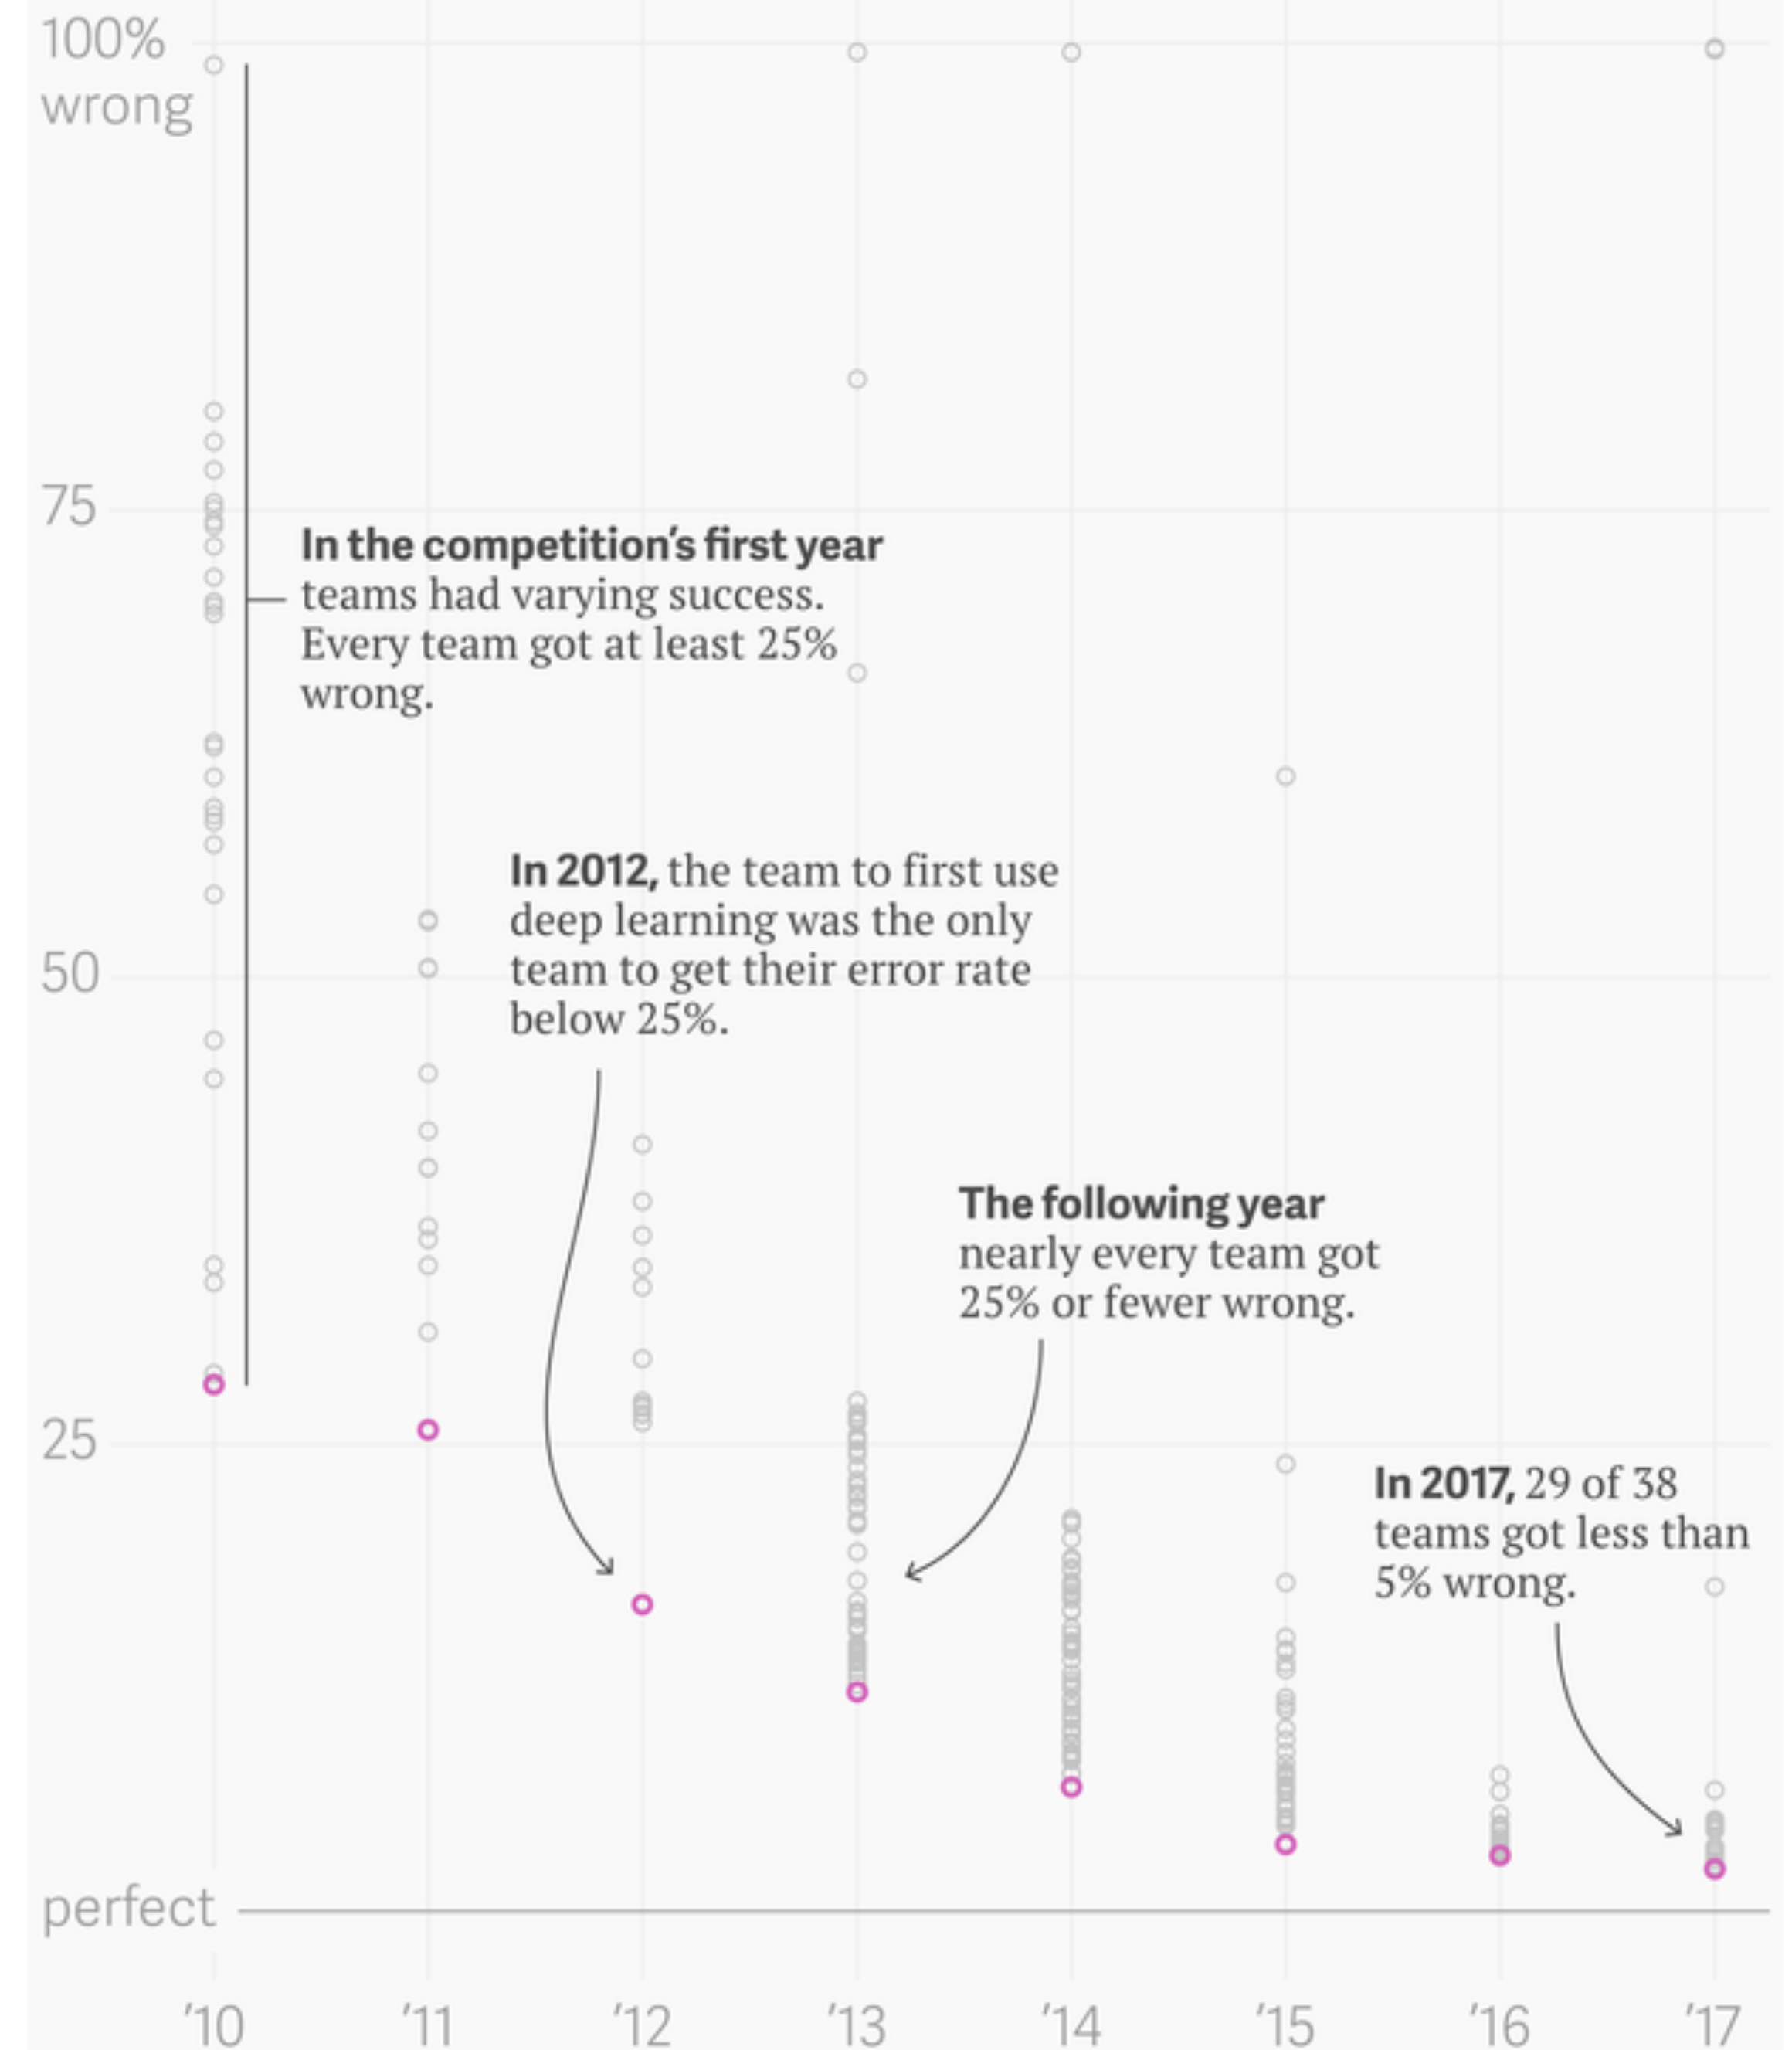




# Classify Images



ImageNet Large Scale Visual Recognition Challenge results



<https://qz.com/1034972/the-data-that-changed-the-direction-of-ai-research-and-possibly-the-world/>



# Detect and Segment Objects

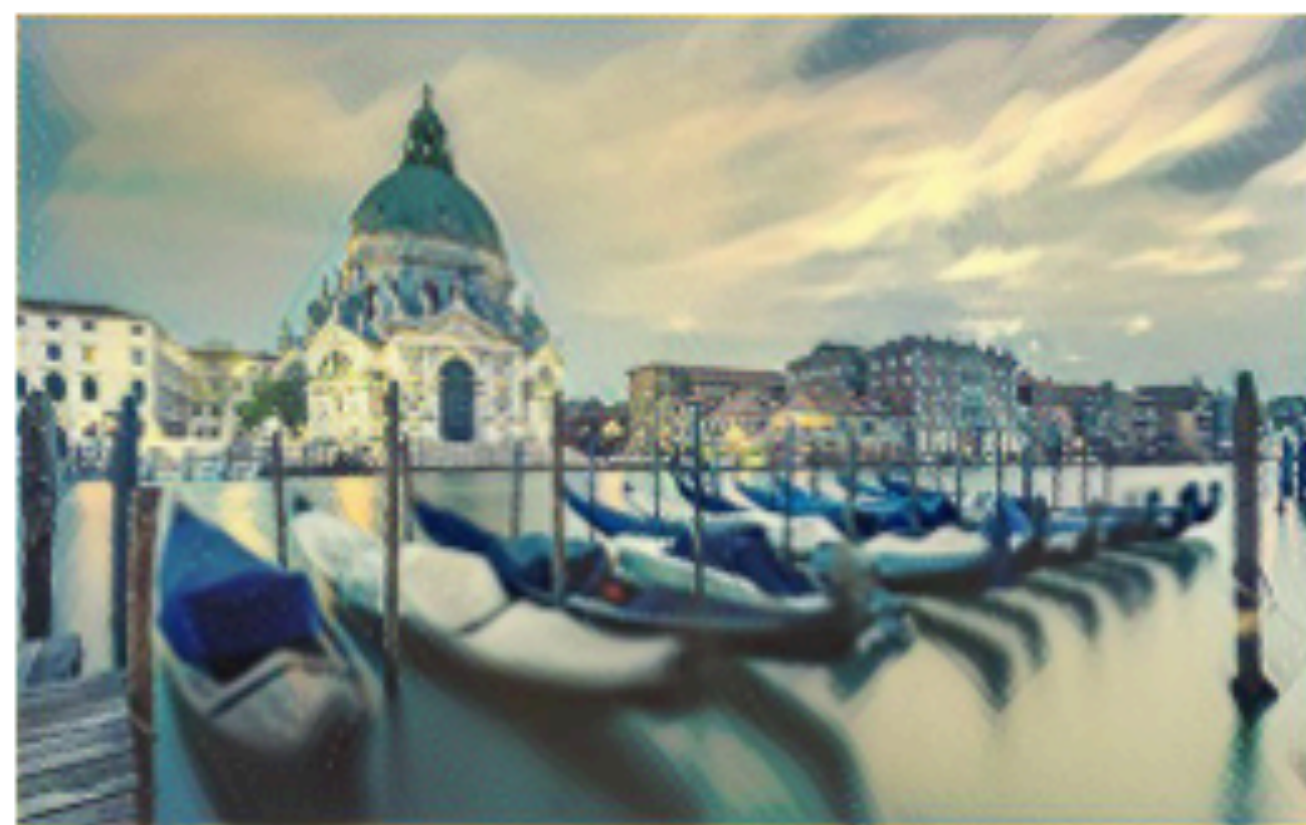


[https://github.com/matterport/Mask\\_RCNN](https://github.com/matterport/Mask_RCNN)



# Style Transfer

<https://github.com/StacyYang/MXNet-Gluon-Style-Transfer>





# Generative Modeling

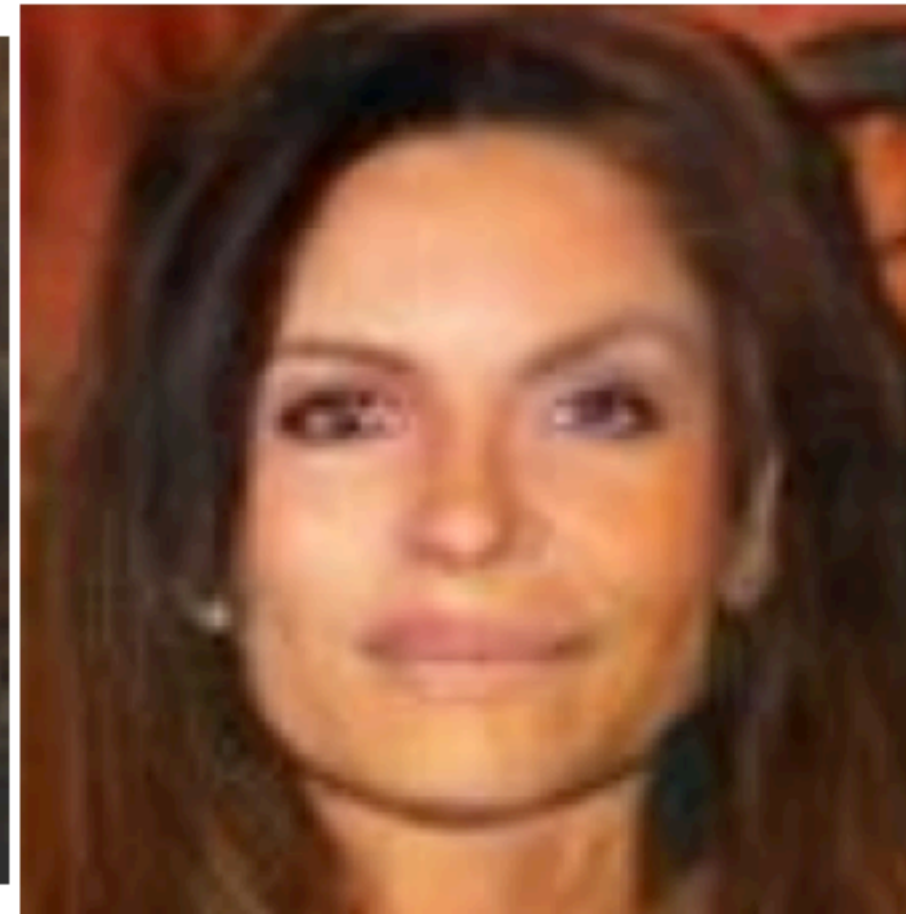
4.5 years of face generation



2014



2015



2016



2017



2018



# Artificial Image Synthesis

Synthesize the images





# Machine Translation

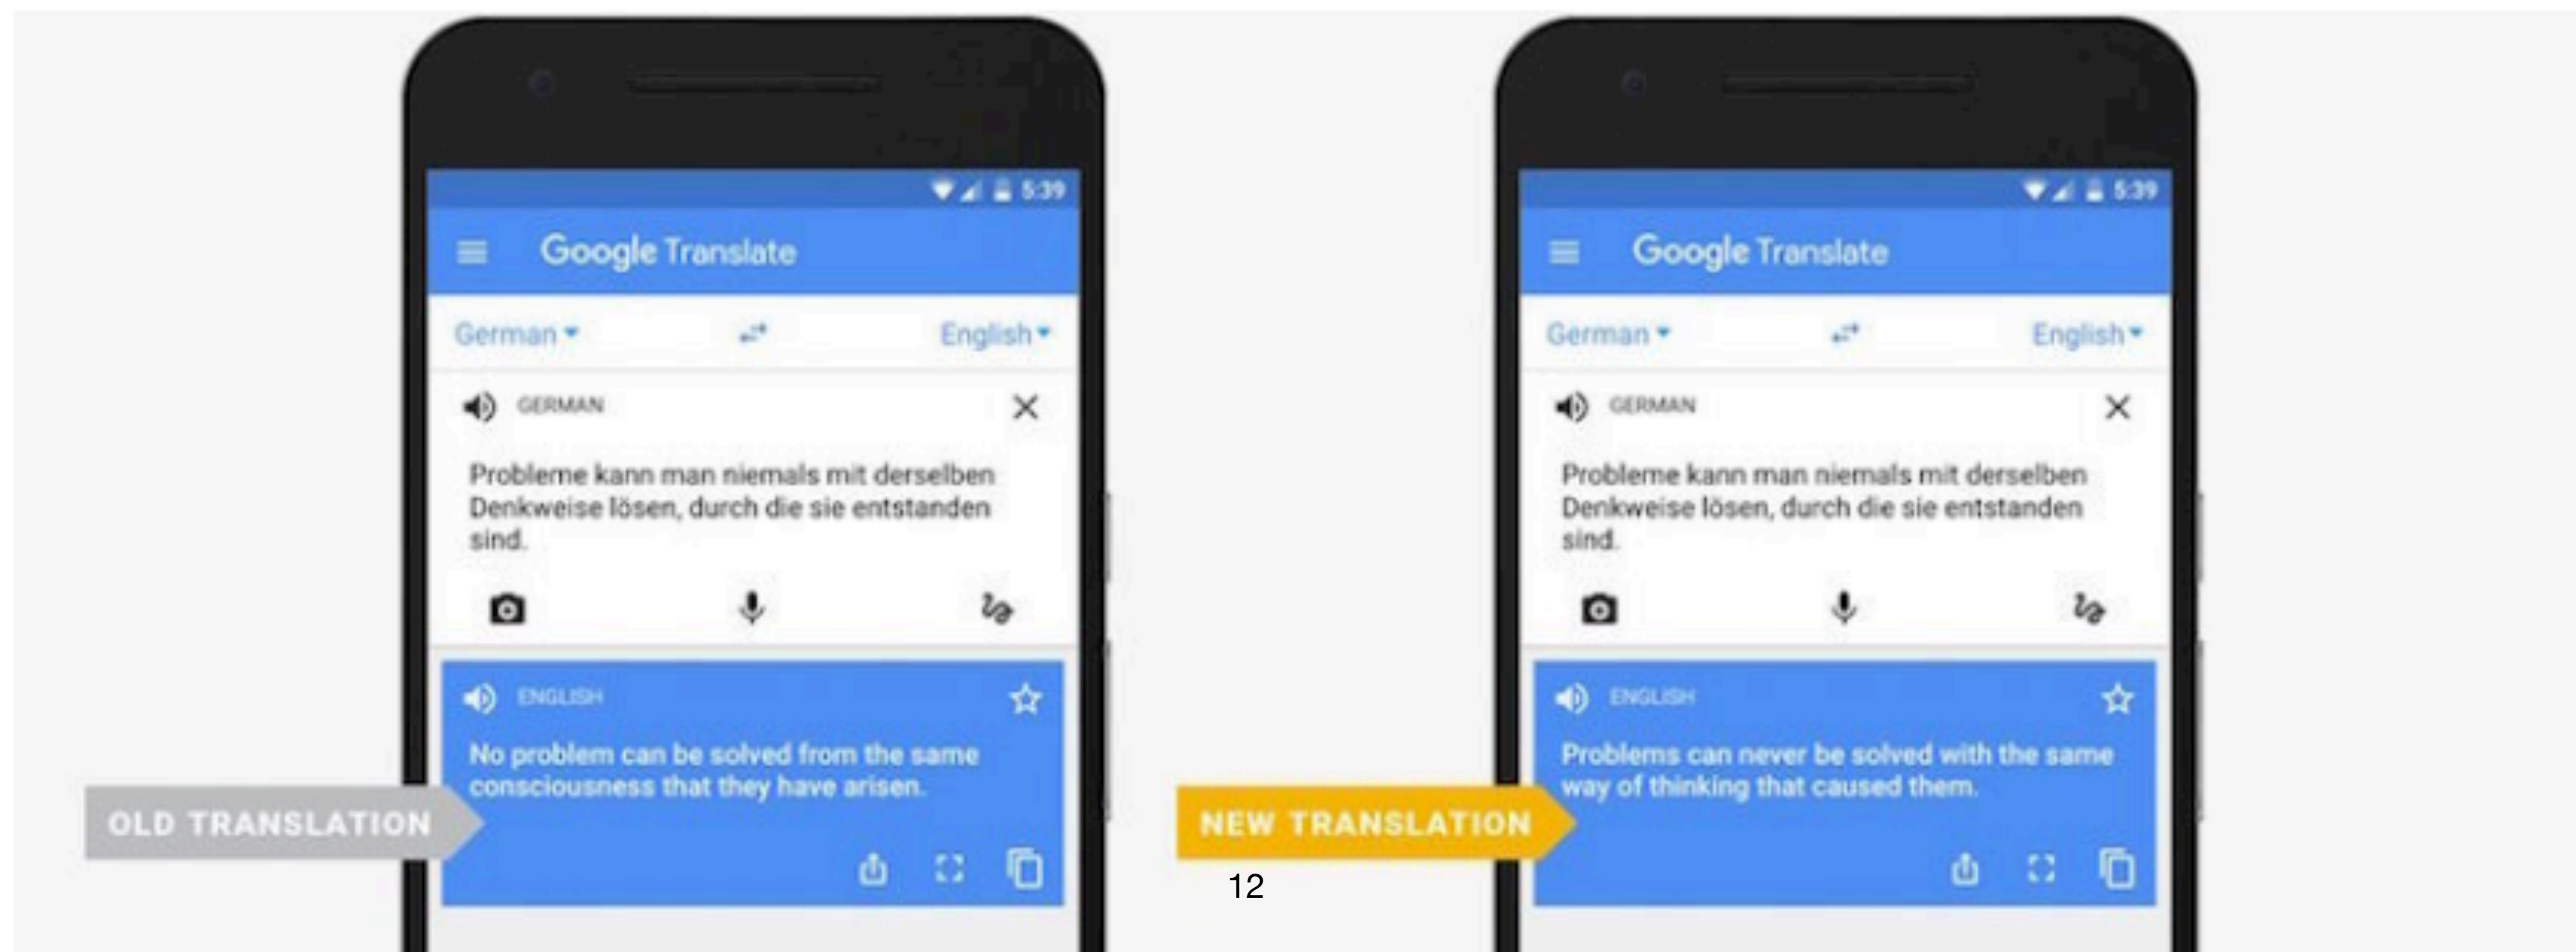
<https://www.pcmag.com/news/google-expands-neural-networks-for-language-translation>

## *Google Expands Neural Networks for Language Translation*

The new system can translate whole sentences at a time, rather than just phrases.



By Tom Brant November 15, 2016





# Text Synthesis

Content: Two dogs play by a tree.

Style: happily, love



RNN

Two dogs **in love** play **happily** by a tree.

Li et al, NACCL, 2018

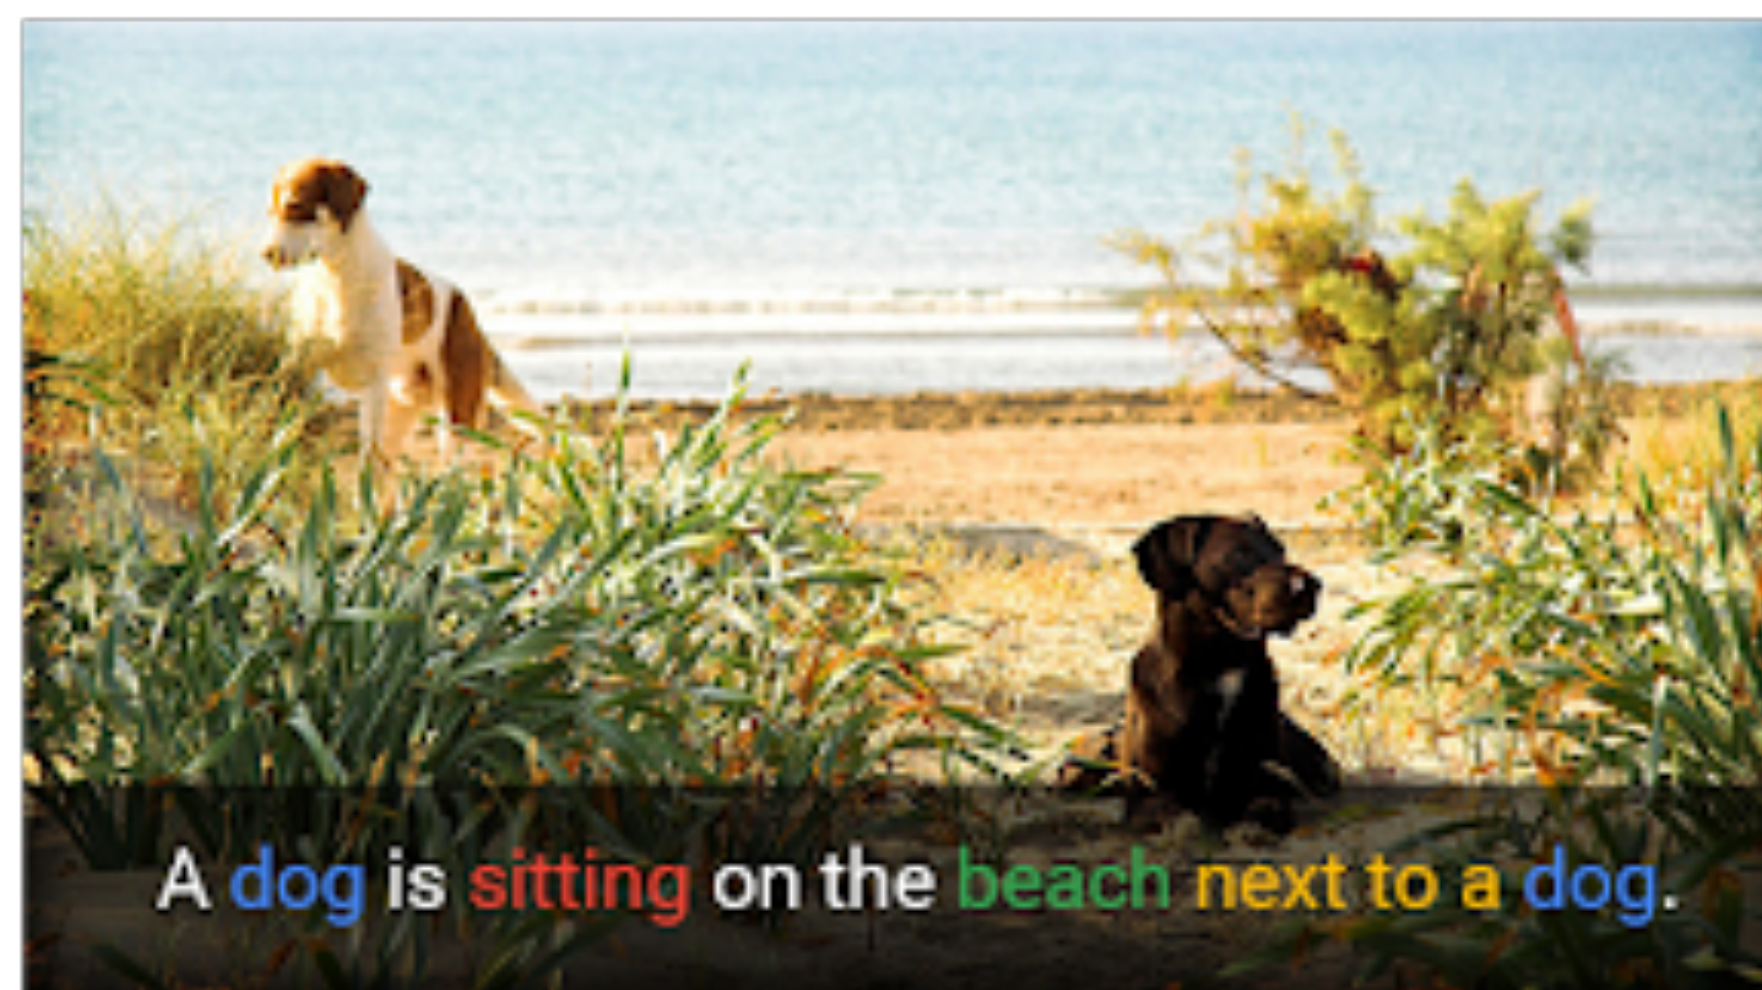


# Image Captioning

Human captions from the training set



Automatically captioned



<https://ai.googleblog.com/2016/09/show-and-tell-image-captioning-open.html>



## OpenAI ChatGPT:

me: why is carbon dioxide highly sought after?

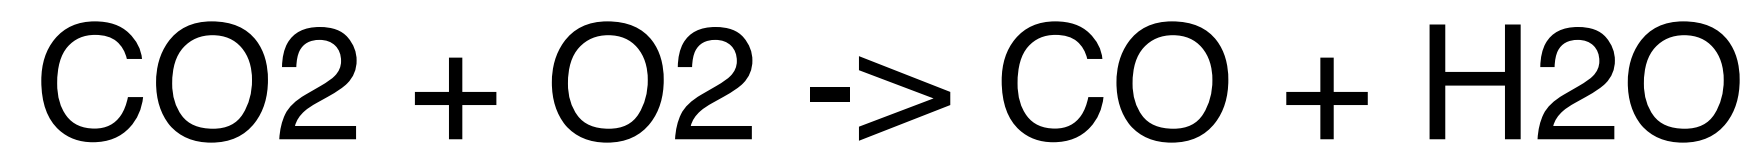
ChatGPT: In recent years, CO<sub>2</sub> has also become highly sought after as a potential source of renewable energy. When burned, CO<sub>2</sub> releases a large amount of heat, which can be harnessed to generate electricity. It can also be used to power fuel cells, which are a type of clean energy technology that converts chemical energy into electricity.



## OpenAI ChatGPT:

me: which molecules are produce when burning co2?

ChatGPT: When carbon dioxide (CO<sub>2</sub>) is burned, it reacts with oxygen in the air to produce carbon monoxide (CO) and water (H<sub>2</sub>O). The chemical reaction for this process can be written as:



The exact products of the combustion of CO<sub>2</sub> depend on the conditions under which it is burned, such as the temperature, oxygen concentration, and presence of other chemicals.

However, in general, CO and H<sub>2</sub>O are the primary products of the burning of CO<sub>2</sub>.



# OpenAI ChatGPT:

me: why is salt highly sought after in the desalination process?

ChatGPT: Salt, or sodium chloride, is highly sought after in the desalination process because it is a readily available and relatively inexpensive compound that can be used to remove dissolved salts from seawater. Desalination is a process that is used to produce fresh water from seawater, and it is often necessary in areas where there is a shortage of fresh water.

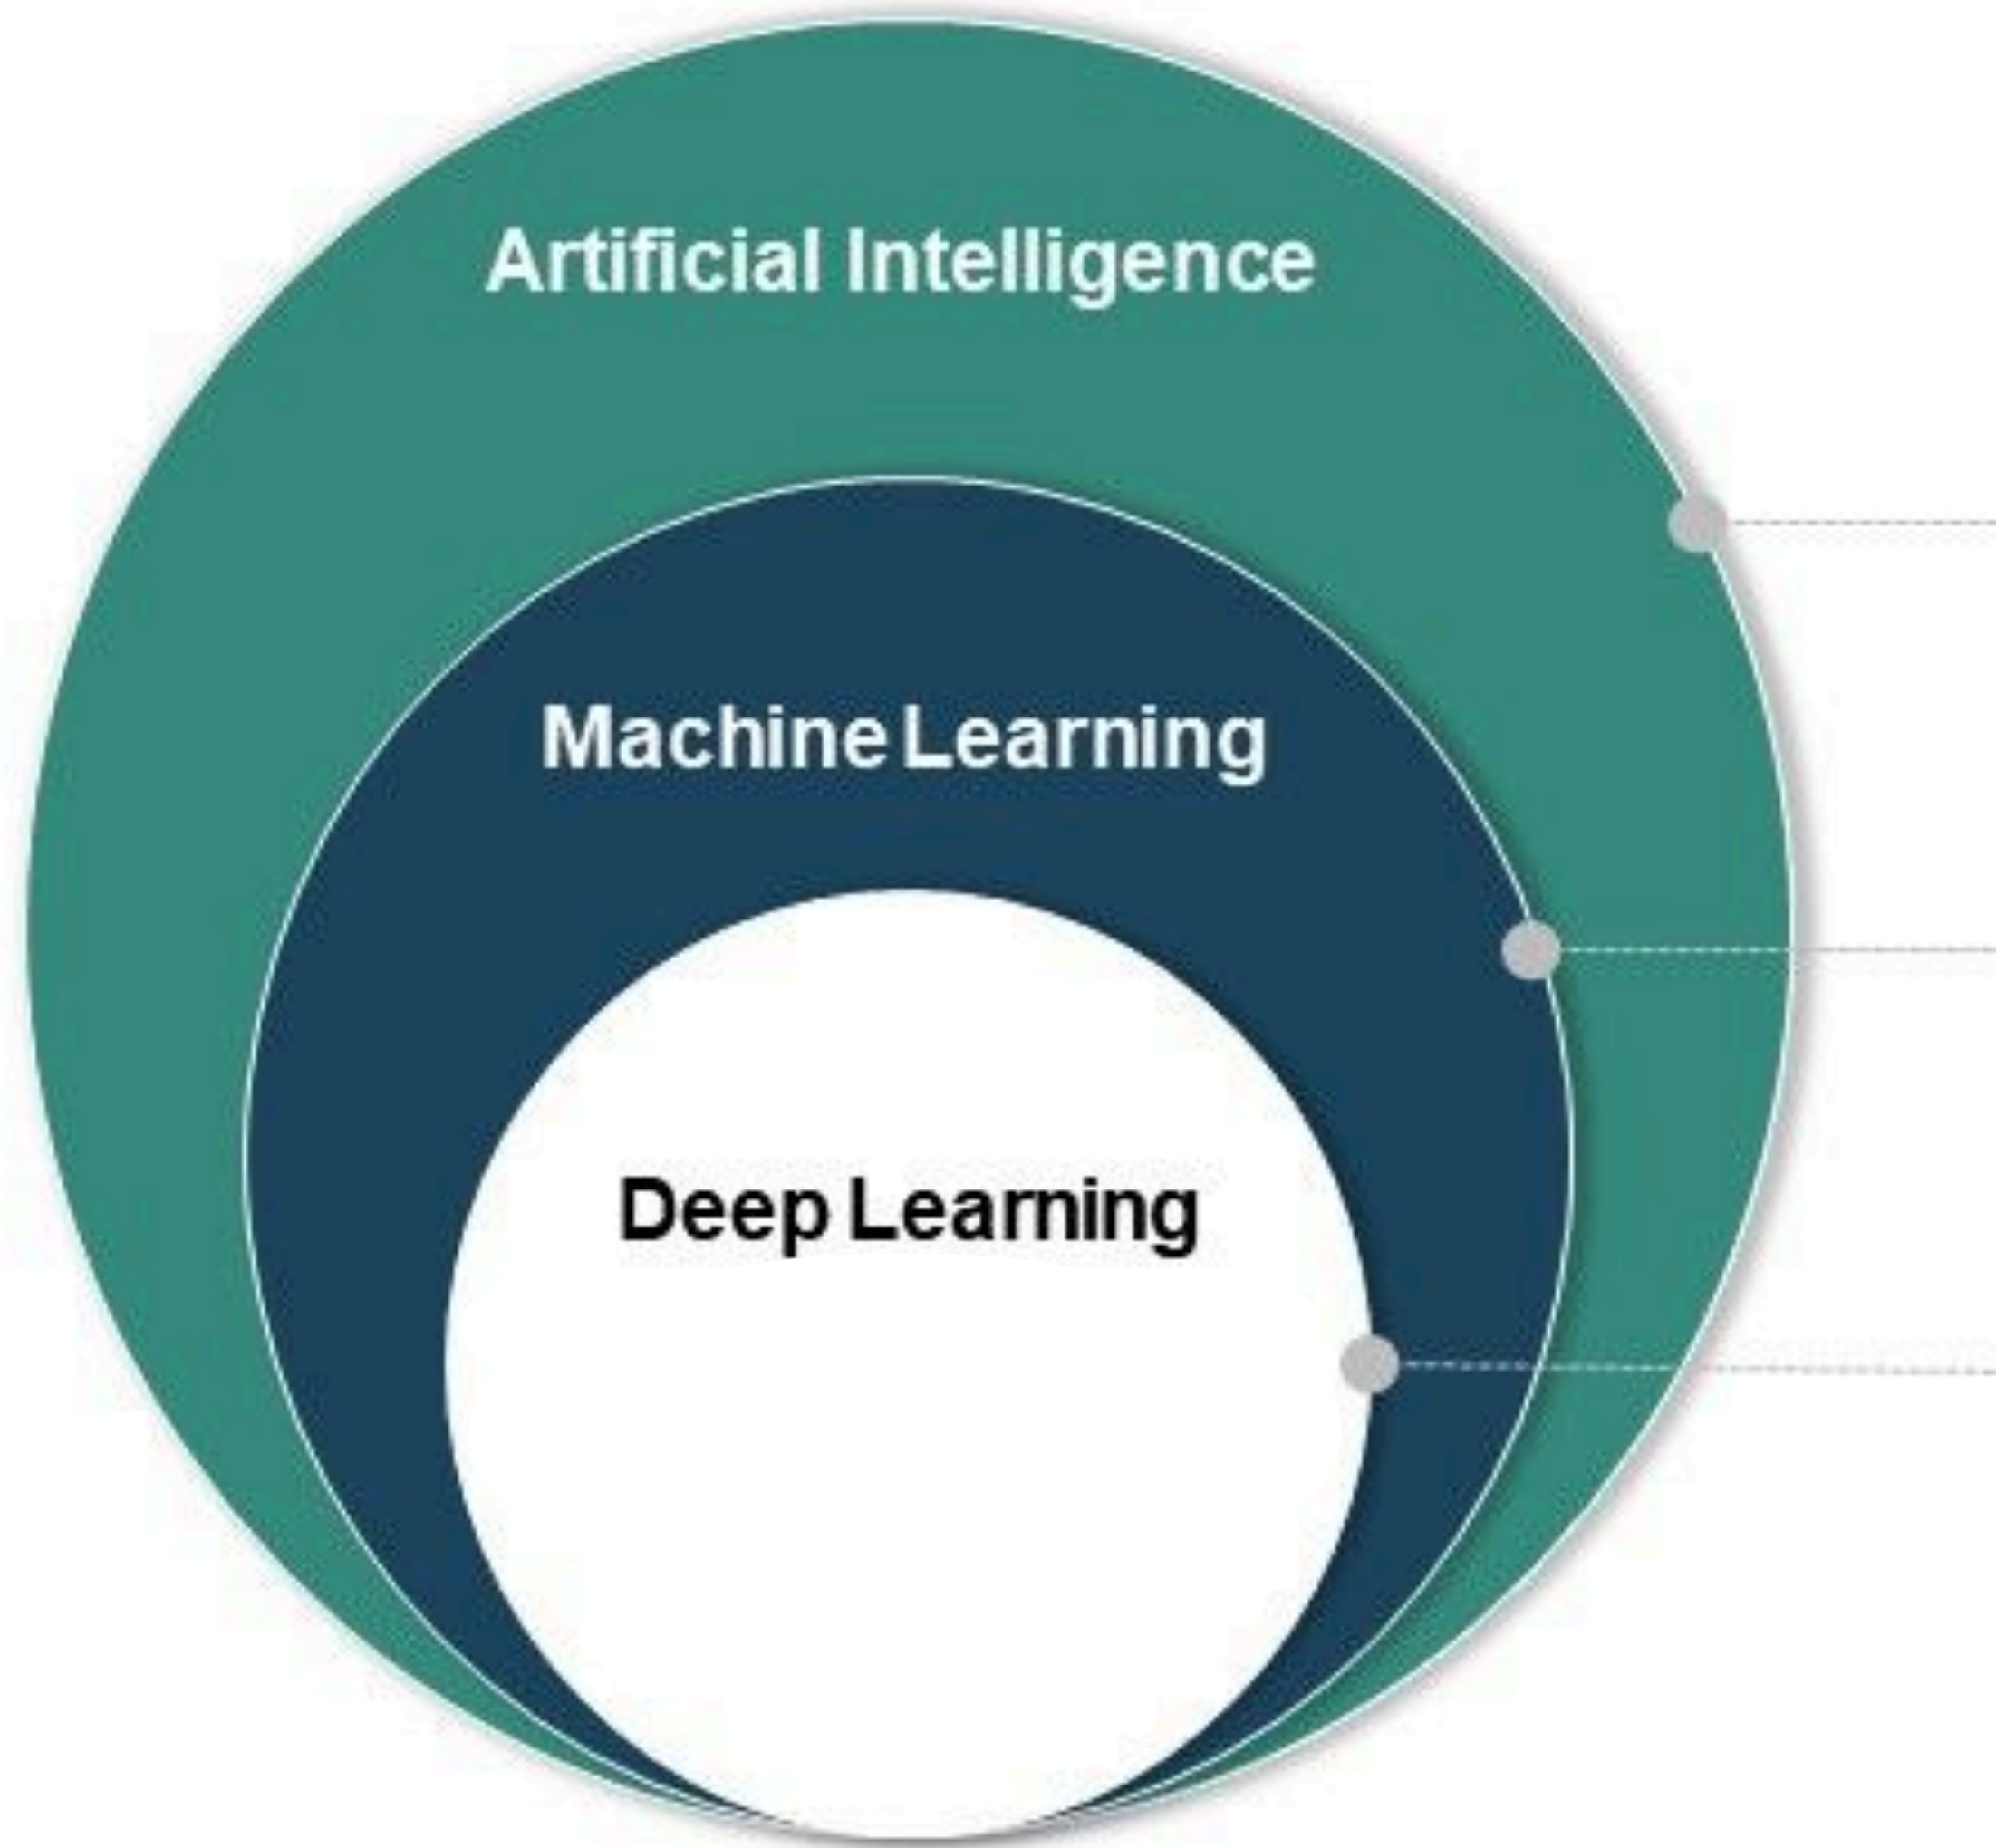
Salt is used in the desalination process because it has a unique chemical property that allows it to readily dissolve in water. When seawater is mixed with salt, the salt ions will dissolve and become evenly distributed throughout the water. This creates a solution that has a higher concentration of salt ions than the seawater.

The higher concentration of salt ions in the solution makes it possible to use a process called reverse osmosis to remove the dissolved salts from the seawater. In reverse osmosis, the salty solution is forced through a membrane that only allows water molecules to pass through, effectively separating the salt from the water. The resulting water is fresh and drinkable, and the salt can be recovered and used again in the desalination process.

Overall, salt is highly sought after in the desalination process because it is a key component in the reverse osmosis process, which is the most commonly used method for removing dissolved salts from seawater.



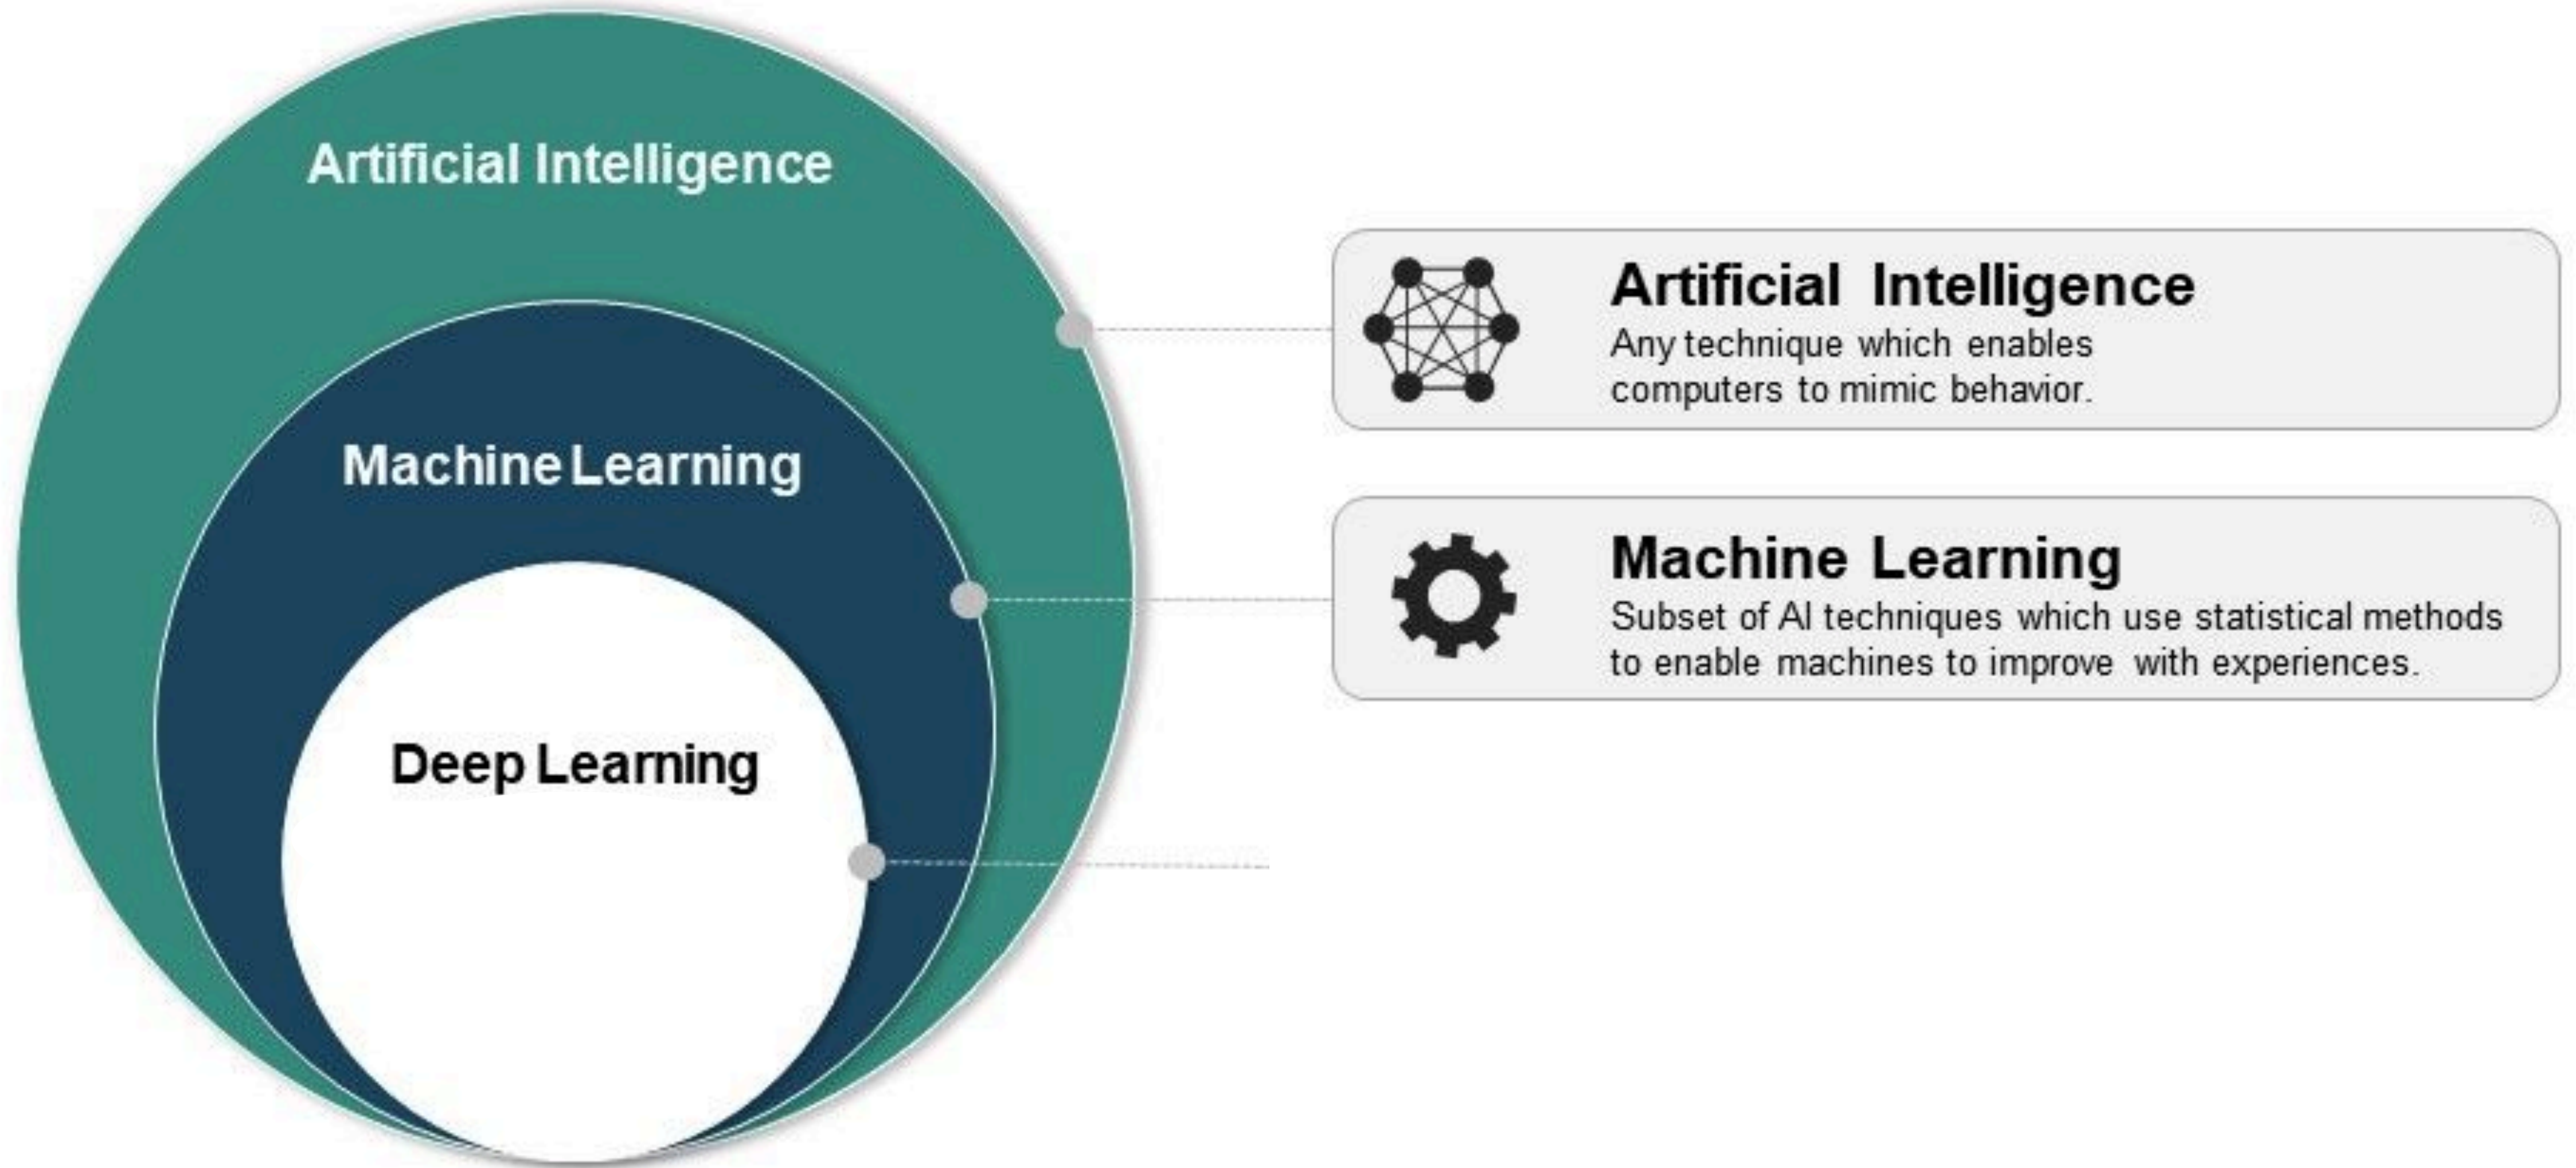
# What is Artificial Intelligence (AI)?



**Artificial Intelligence**  
Any technique which enables computers to mimic behavior.

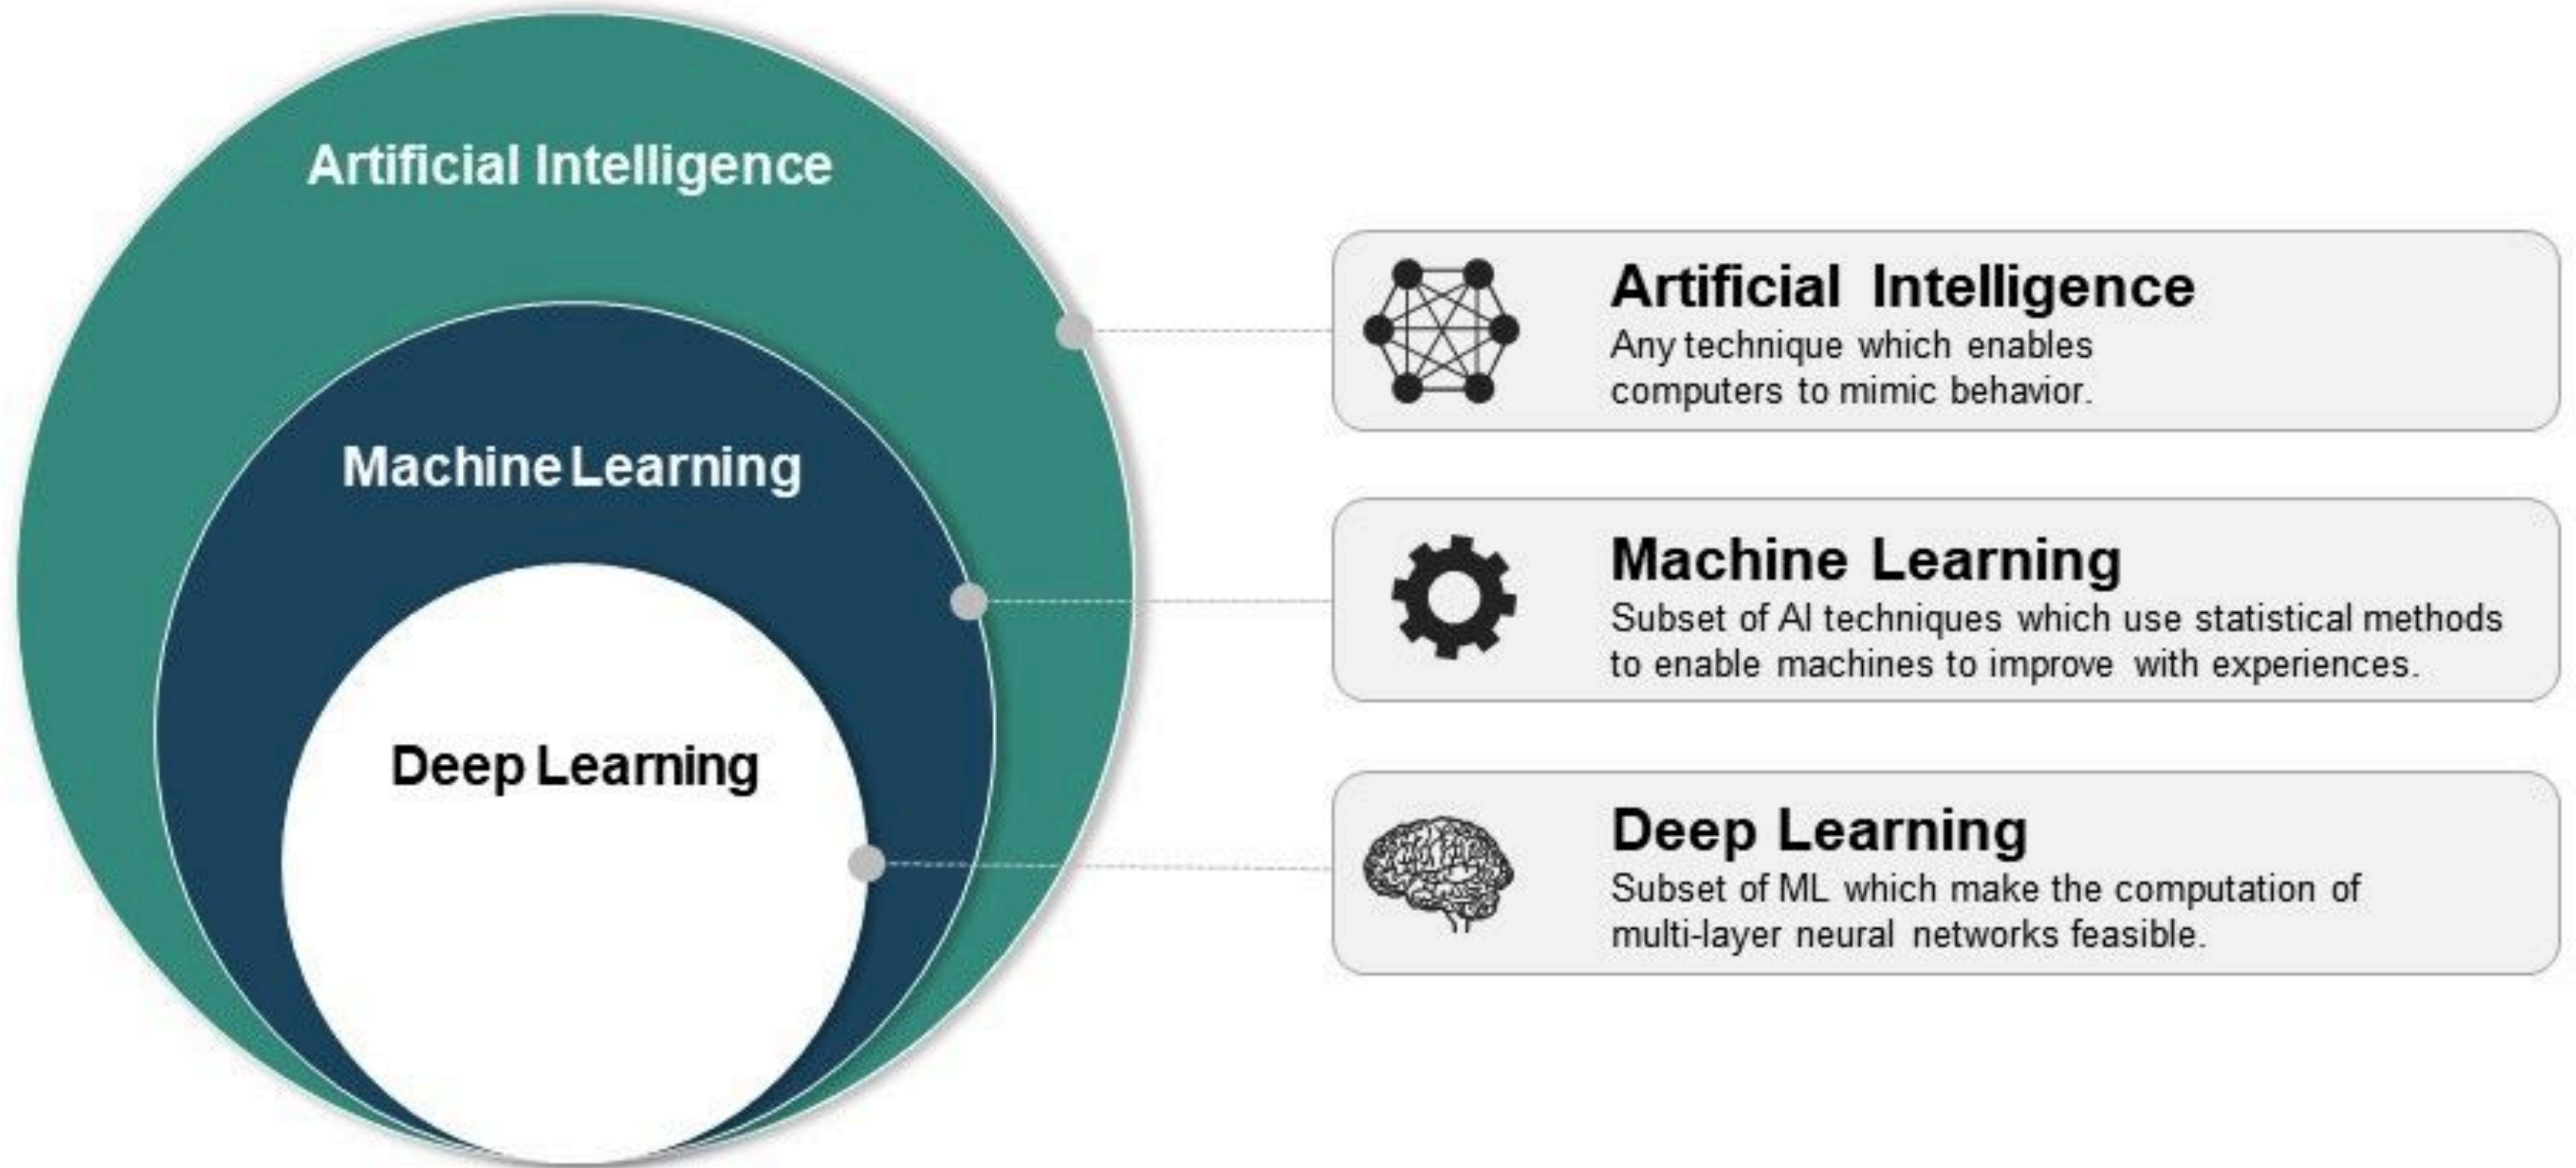


# What is Artificial Intelligence (AI)?





# What is Artificial Intelligence (AI)?





# Artificial Intelligence is not Magic

They rely on **fundamental** techniques in:

- Algorithms
- Mathematics
- Logic
- Probability and Statistics
- Optimization



# What you can learn from CS540?

- Foundational tools in **Machine Learning** and **Artificial Intelligence**: Linear algebra, Probability, Logic, and elements of Statistics.
- Core techniques in **Natural Language Processing (NLP)**, including bag-of-words, tf-idf, n-Gram Models, and Smoothing.
- Basics of **Machine Learning**. supervised learning vs. unsupervised learning
- **Neural Networks and Deep Learning**: Network Architecture, Training, Backpropagation, Stochastic Gradient Descent.
- Fundamentals of **Game Theory**.
- **Search and Reinforcement Learning**
- **Artificial Intelligence** and **Machine Learning** in Real-World settings and the Ethics of Artificial Intelligence.



# What you can learn from CS540?

- Foundational tools in **Machine Learning** and **Artificial Intelligence**: Linear algebra, Probability, Logic, and elements of Statistics.
- Core techniques in **Natural Language Processing (NLP)**, including bag-of-words, tf-idf, n-Gram Models, and Smoothing.

**TL;DR Lots of useful stuff, theory and practice in AI**

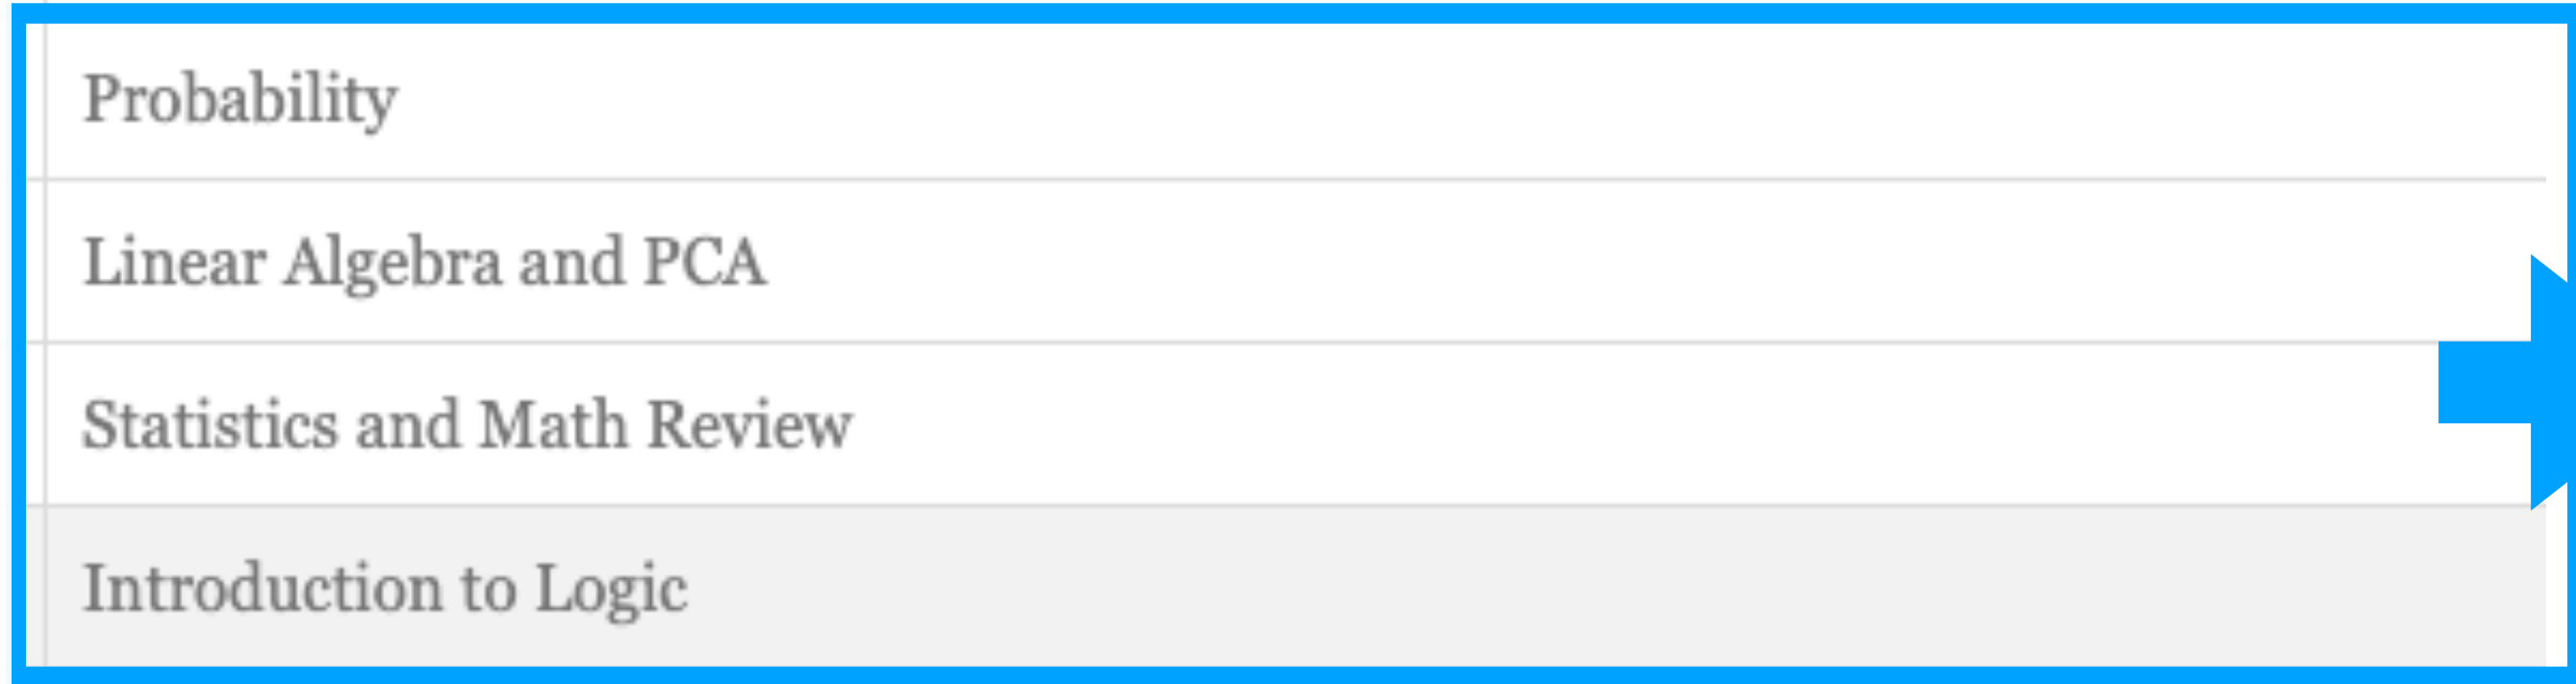
- **Basics of Machine Learning**: supervised learning vs. unsupervised learning
- **Neural Networks and Deep Learning**: Network Architecture, Training, Backpropagation, Stochastic Gradient Descent.
- Fundamentals of **Game Theory**.
- **Search and Reinforcement Learning**
- **Artificial Intelligence and Machine Learning** in Real-World settings and the Ethics of Artificial Intelligence.



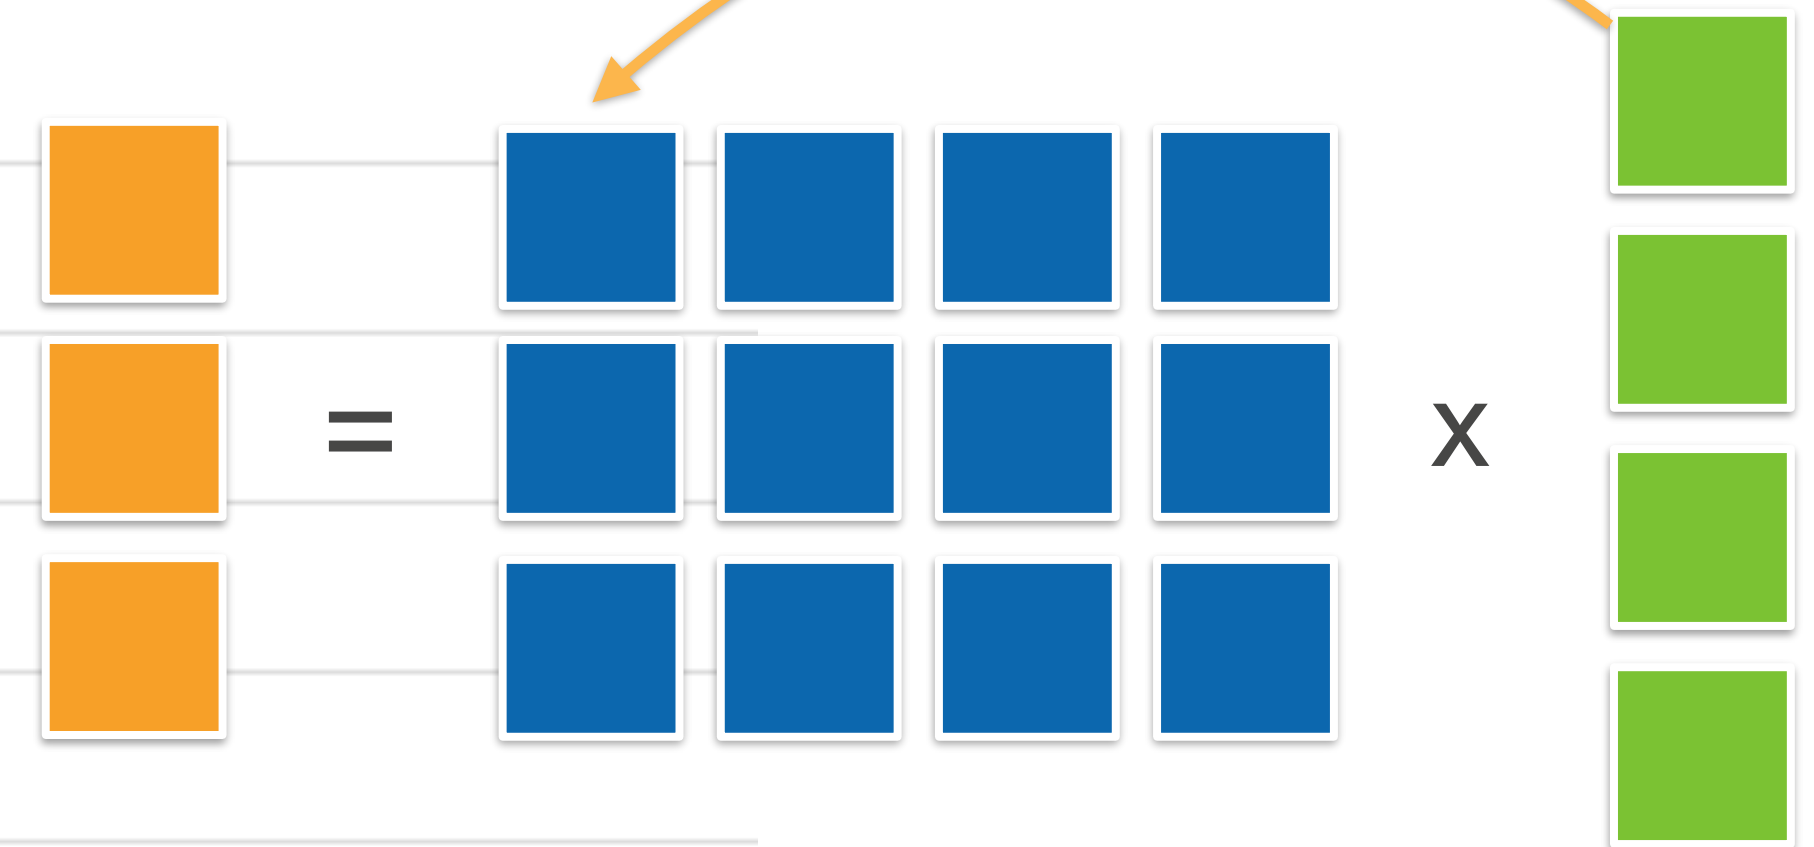
# What you can learn from CS540?

[https://pages.cs.wisc.edu/~sharonli/courses/cs540\\_spring2022/index.html](https://pages.cs.wisc.edu/~sharonli/courses/cs540_spring2022/index.html)

Date	Topic
Tuesday, Jan 26	Welcome and Introduction to Python
Thursday, Jan 28	Probability
Tuesday, Feb 2	Linear Algebra and PCA
Thursday, Feb 4	Statistics and Math Review
Tuesday, Feb 9	Introduction to Logic
Thursday, Feb 11	Natural Language Processing
Tuesday, Feb 16	Machine Learning: Introduction
Thursday, Feb 18	Machine Learning: Unsupervised Learning I
Tuesday, Feb 23	Machine Learning: Unsupervised Learning II
Thursday, Feb 25	Machine Learning: Linear regression
Tuesday, March 2	Machine Learning: K - Nearest Neighbors
Thursday, March 4	Machine Learning: Perceptron
Tuesday, March 9	Machine Learning: Neural Network I



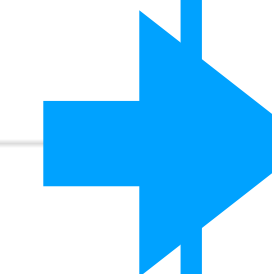
**Foundations**





# What you can learn from CS540?

Date	Topic
Tuesday, Jan 26	Welcome and Introduction to Python
Thursday, Jan 28	Probability
Tuesday, Feb 2	Linear Algebra and PCA
Thursday, Feb 4	Statistics and Math Review
Tuesday, Feb 9	Introduction to Logic
Thursday, Feb 11	Natural Language Processing
Tuesday, Feb 16	Machine Learning: Introduction
Thursday, Feb 18	Machine Learning: Unsupervised Learning I
Tuesday, Feb 23	Machine Learning: Unsupervised Learning II
Thursday, Feb 25	Machine Learning: Linear regression
Tuesday, March 2	Machine Learning: K - Nearest Neighbors
Thursday, March 4	Machine Learning: Perceptron
Tuesday, March 9	Machine Learning: Neural Network I



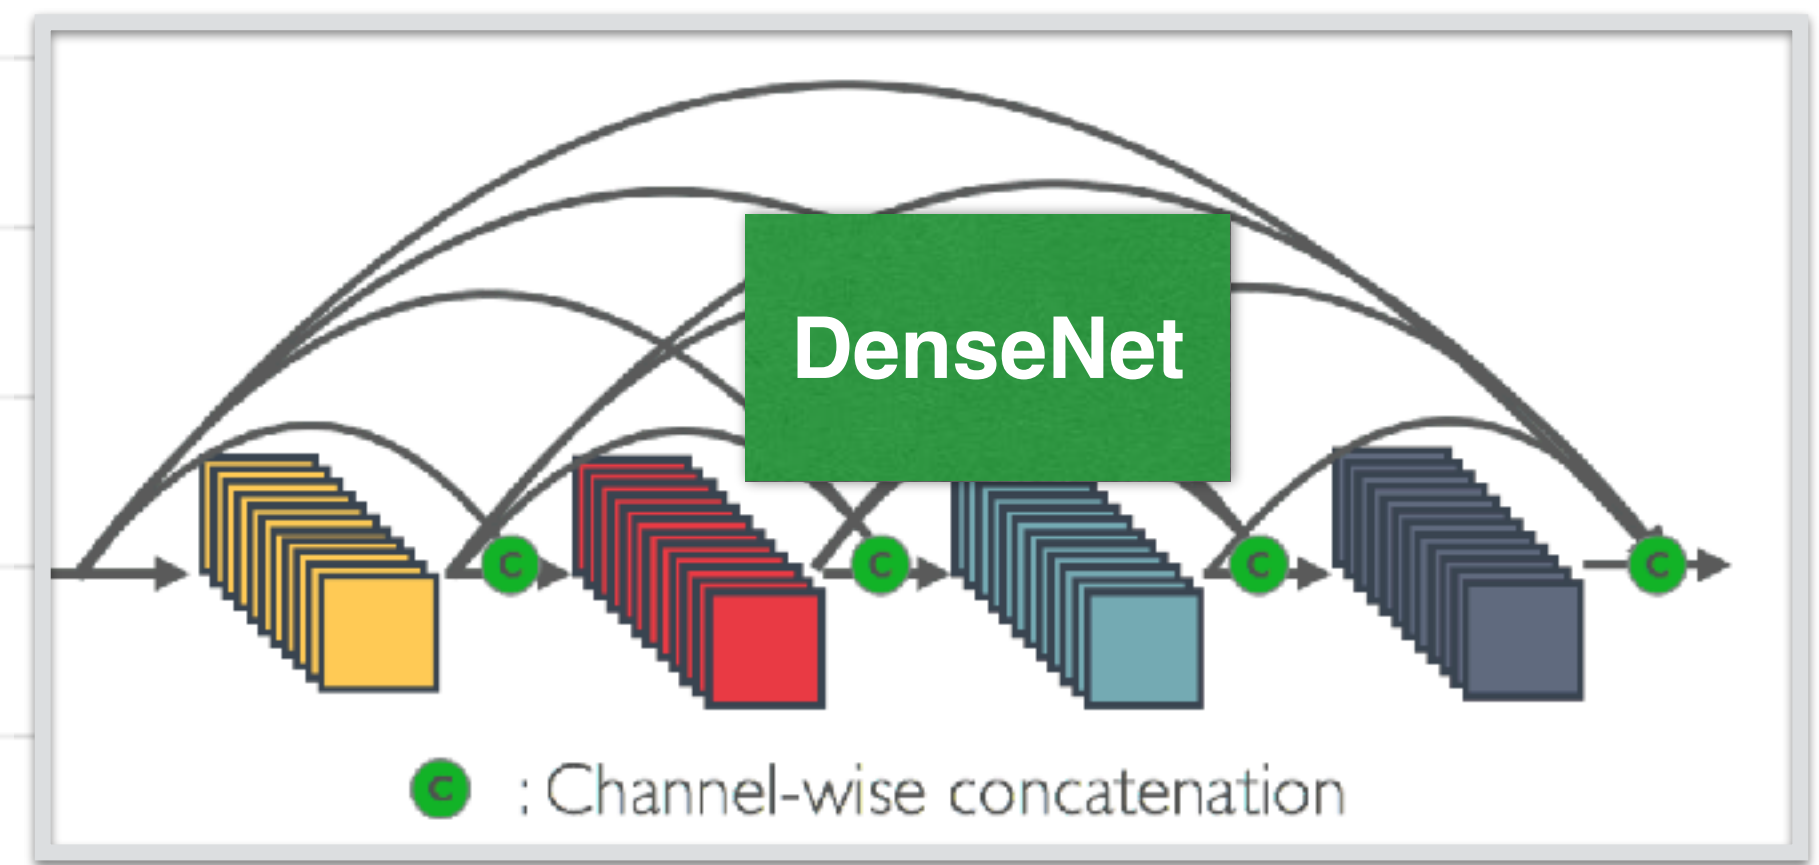
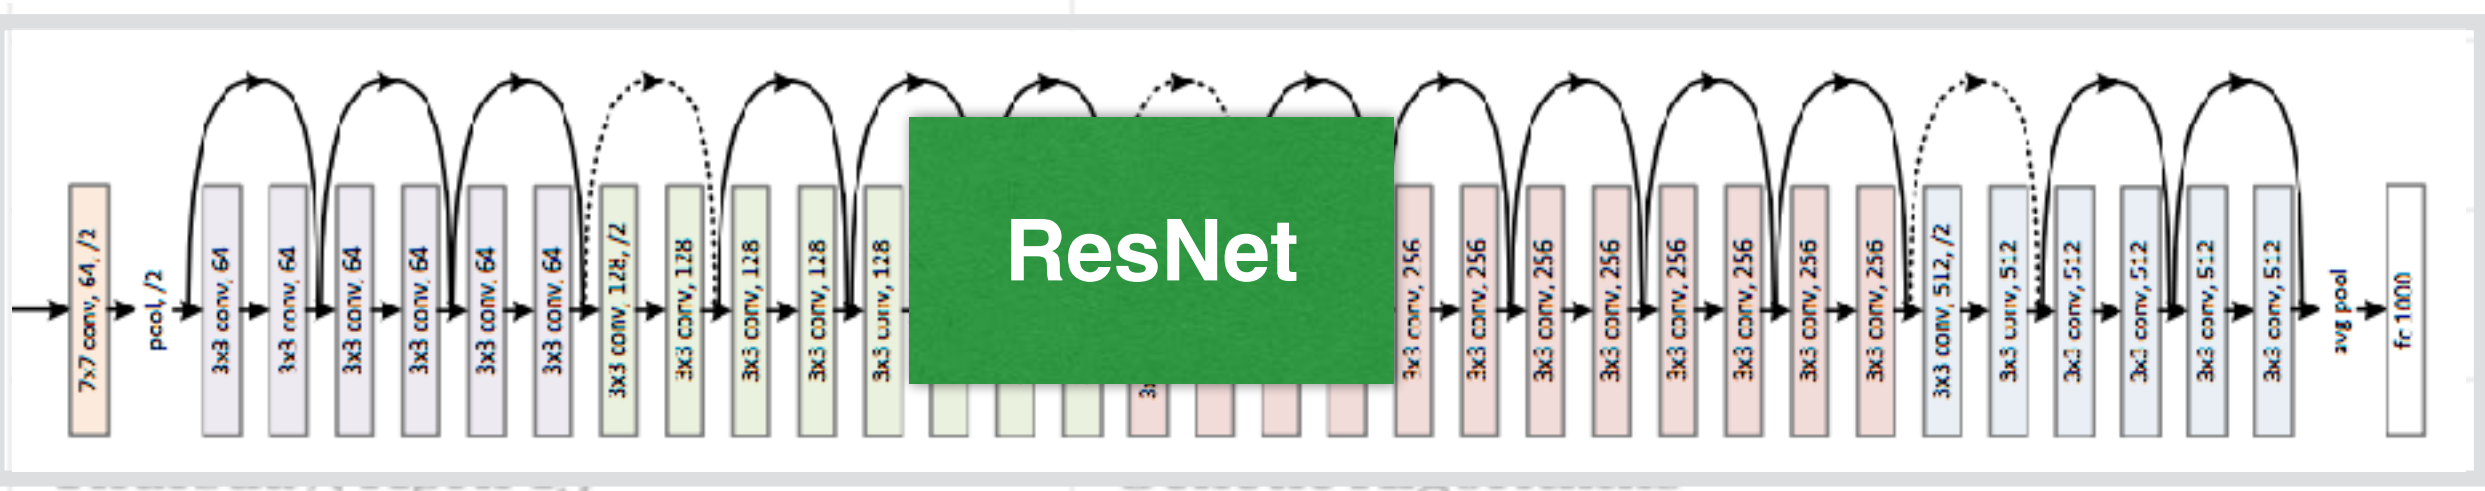
**Machine learning**



# What you can learn from CS540?

Tuesday, March 9	Machine Learning: Neural Network I
Thursday, March 11	Machine Learning: Neural Network II
Tuesday, March 16	Machine Learning: Neural Network III
Thursday, March 18	Machine Learning: Deep Learning I
Tuesday, March 23	Machine Learning: Deep Learning II
Thursday, March 25	Machine Learning: Deep Learning III
Tuesday, March 30	Machine Learning: Deep Learning and Neural Network's Summary

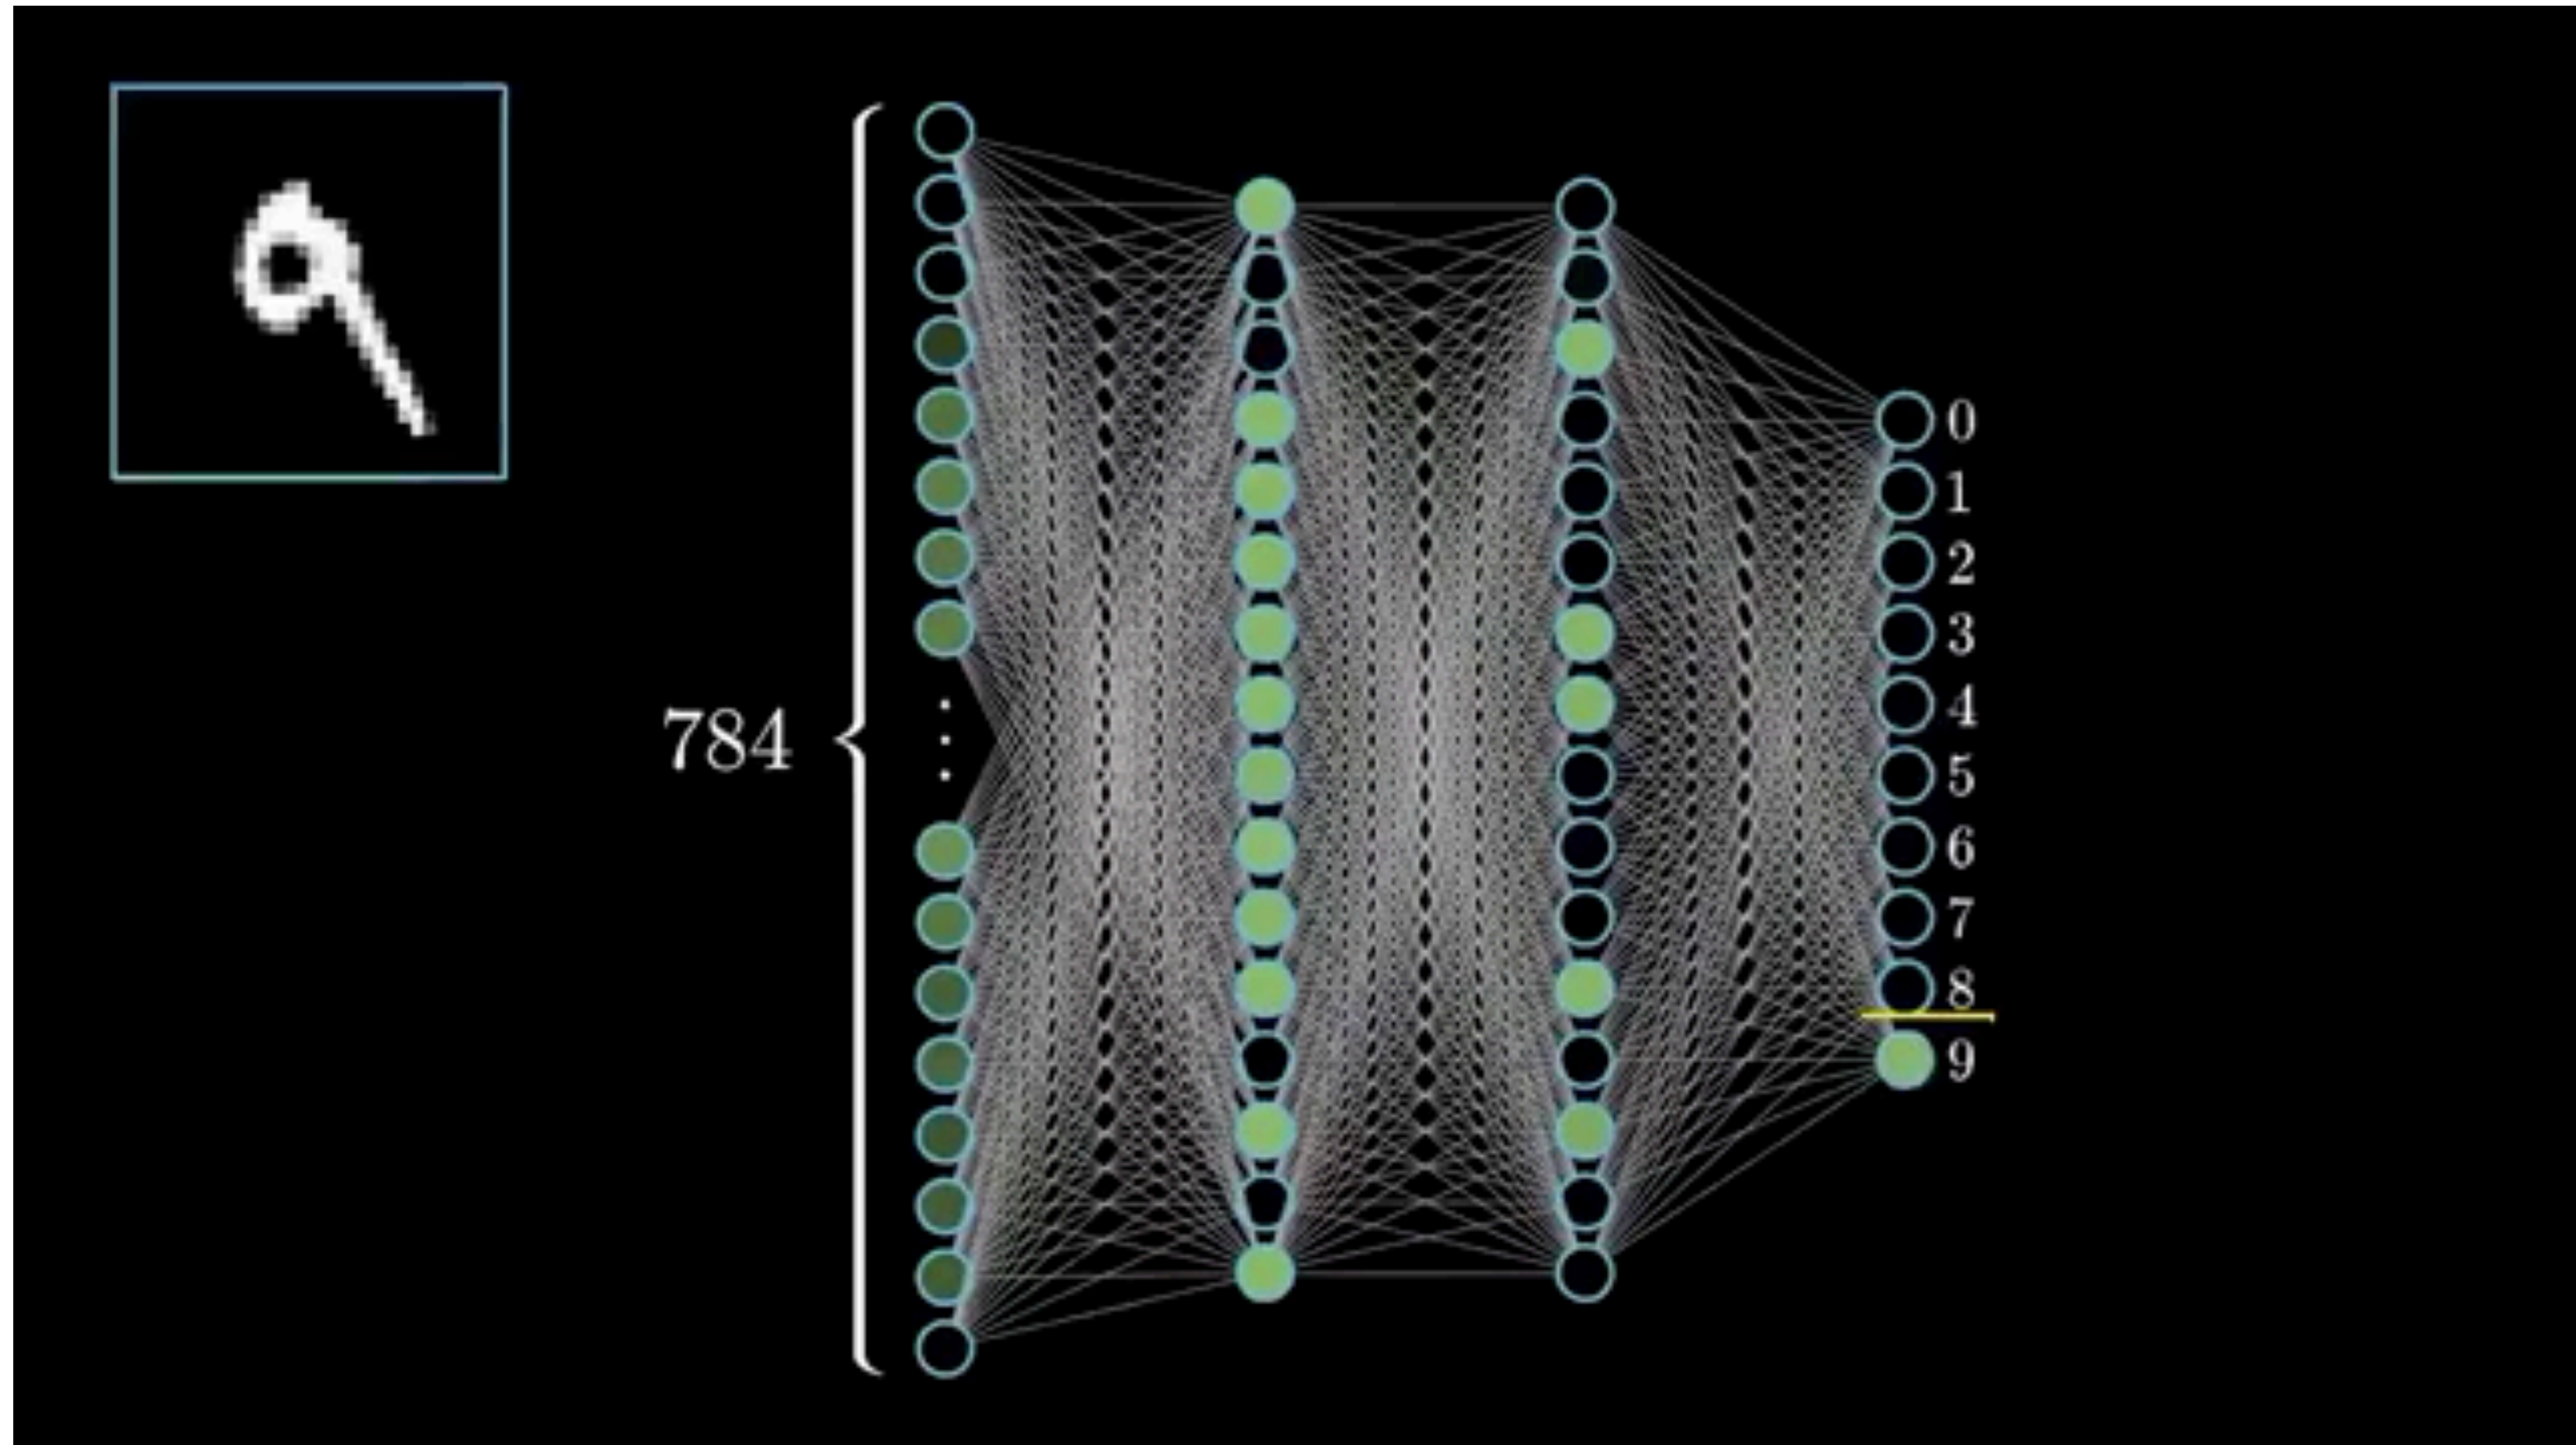
**Deep Learning**



Tuesday, April 20	Introduction to Reinforcement Learning
-------------------	--



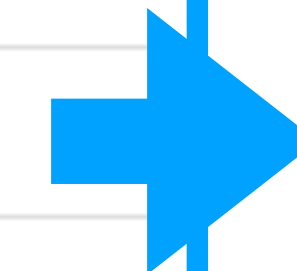
# What you can learn from CS540?





# What you can learn from CS540?

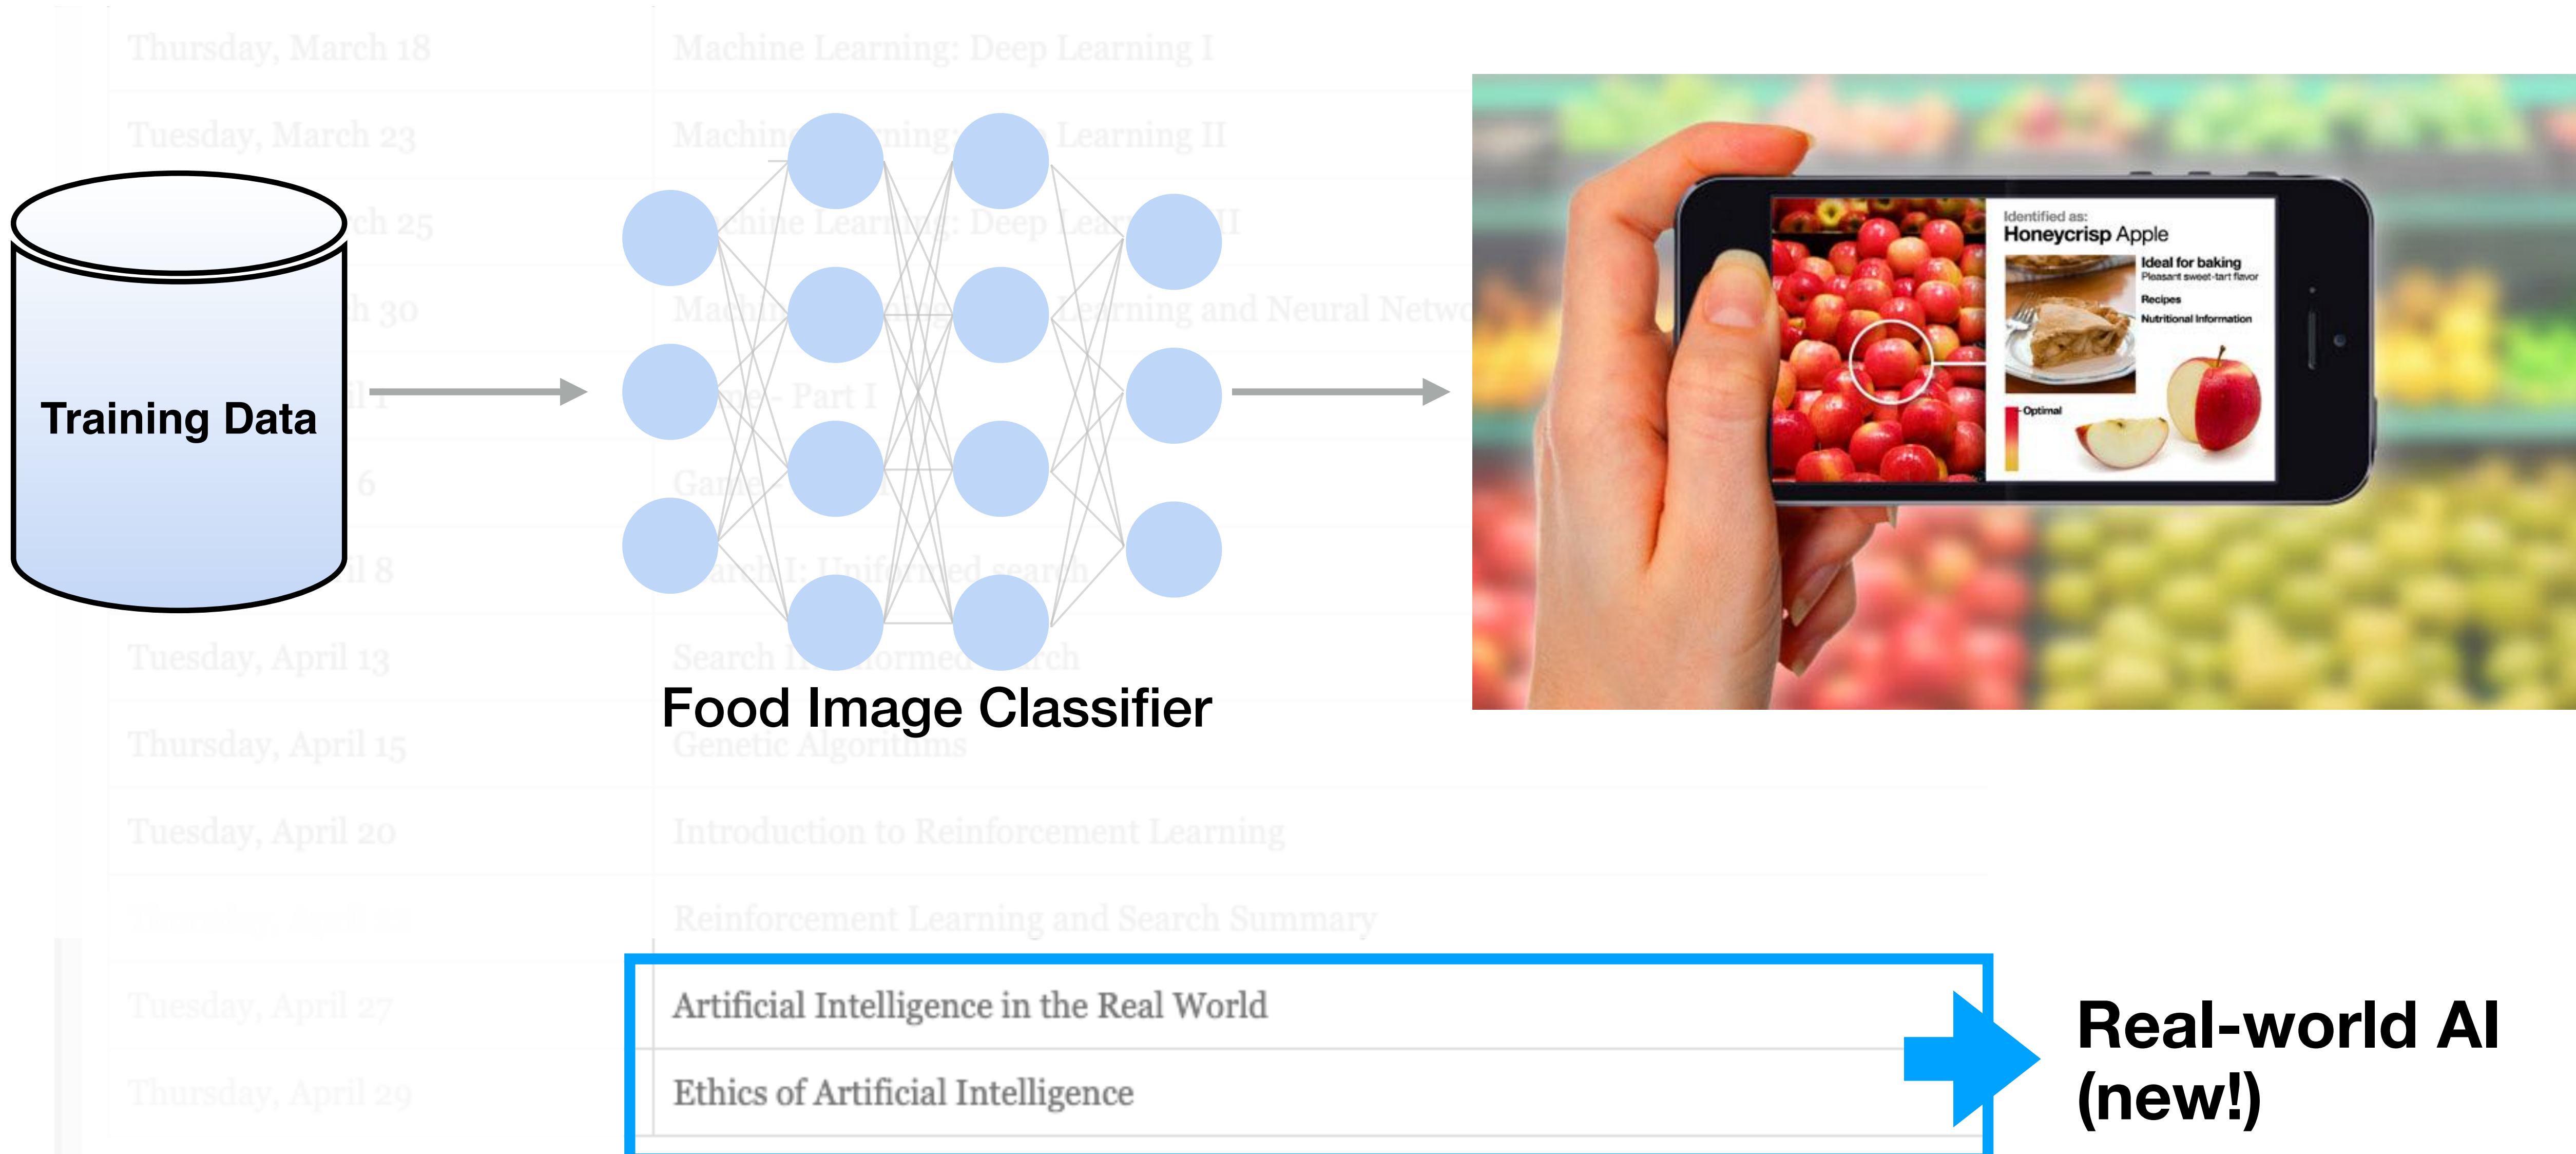
Thursday, March 18	Machine Learning: Deep Learning I
Tuesday, March 23	Machine Learning: Deep Learning II
Thursday, March 25	Machine Learning: Deep Learning III
Tuesday, March 30	Machine Learning: Deep Learning and Neural Network's Summary
Thursday, April 1	Game - Part I
Tuesday, April 6	Game - Part II
Thursday, April 8	Search I: Uniformed search
Tuesday, April 13	Search II: Informed search
Thursday, April 15	Genetic Algorithms
Tuesday, April 20	Introduction to Reinforcement Learning
Thursday, April 22	Reinforcement Learning and Search Summary
Tuesday, April 27	Artificial Intelligence in the Real World
Thursday, April 29	Ethics of Artificial Intelligence



**Game, search  
and  
Reinforcement  
Learning**



# What you can learn from CS540?







# Part II: Course Logistics



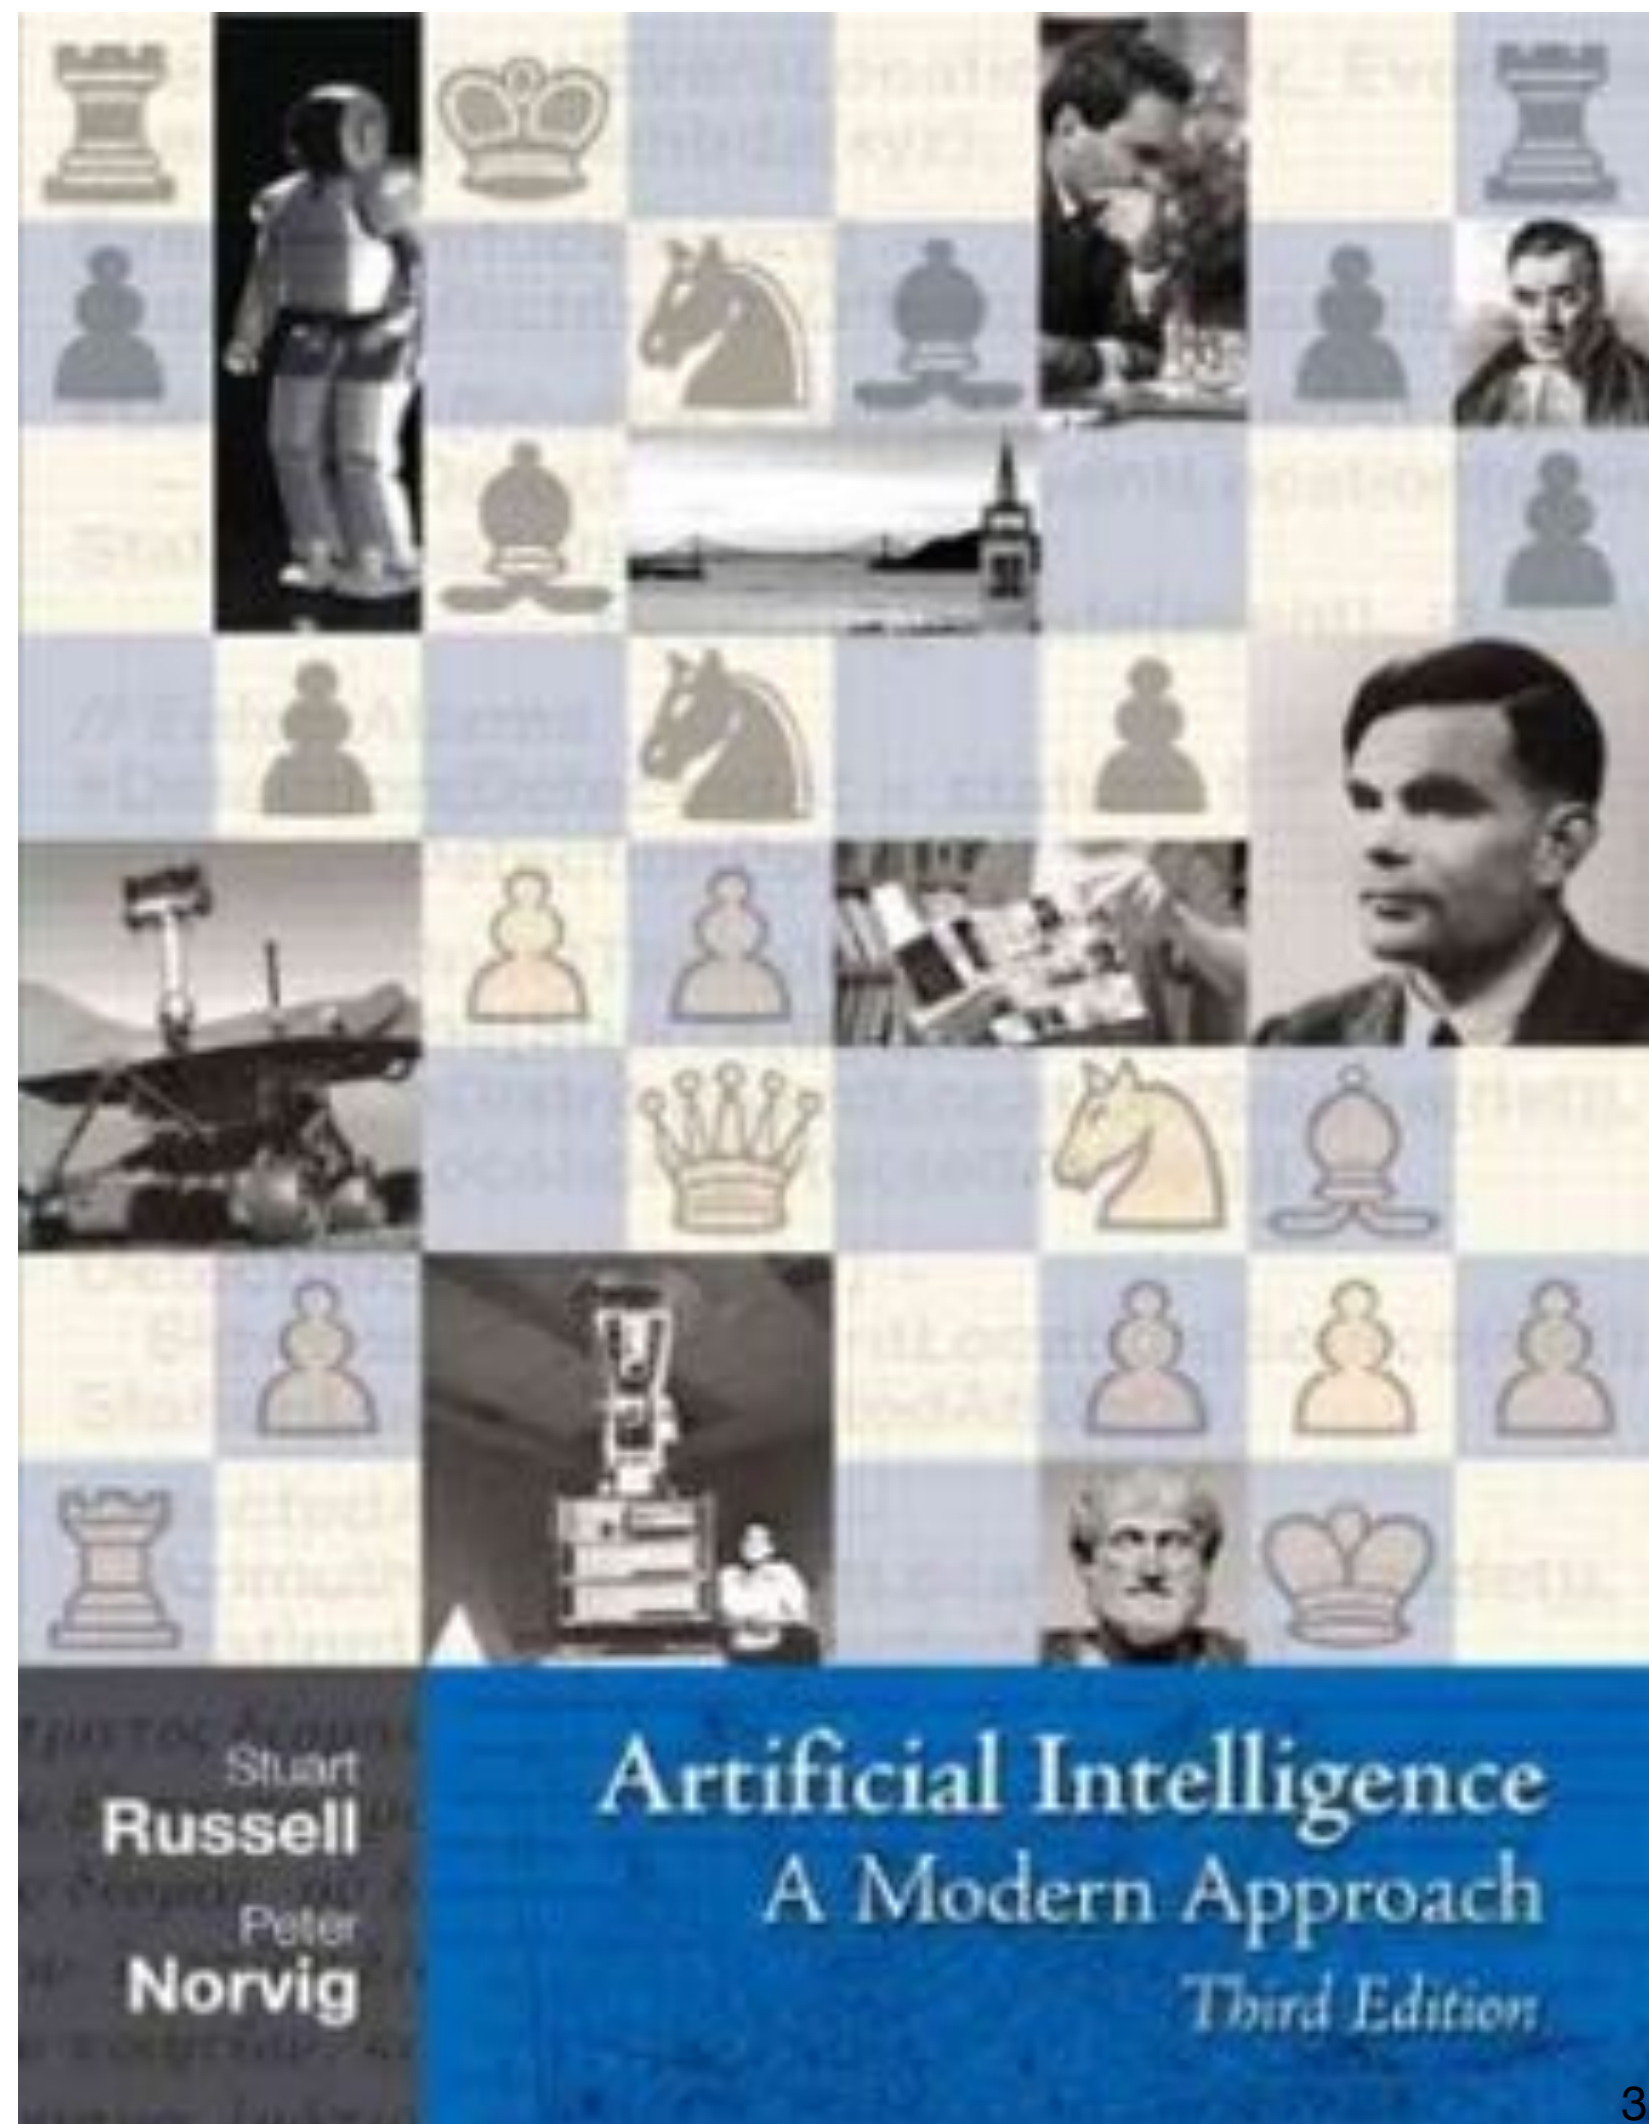
# Where to find everything?

- **Canvas** - [Pointer to everything](#)
  - Assignments, submissions, grades (private materials that should not be shared)
- **Course website** - public materials
  - [https://pages.cs.wisc.edu/~jphanna/teaching/2023spring\\_cs540/index.html](https://pages.cs.wisc.edu/~jphanna/teaching/2023spring_cs540/index.html)
  - Slides, schedule, policies
- **Piazza** <https://piazza.com/wisc/spring2023/cs540>
  - Discussion, questions, announcements



# Textbook

Artificial Intelligence: A Modern Approach (4th edition). Stuart Russell and Peter Norvig. Pearson, 2020. ISBN 978-0134610993. **(textbook is optional, but may be useful)**





# Instruction Team

(See course webpage)

Merged across sections:

- Teaching Assistants (TAs): hold office hours, grade your homework
- Peer Mentors: hold office hours
- Graders: grade your homework



# Office Hours

- Available on the course website
- All office hours are **merged** across sections, you can go to anyone
- Use TA and Peer Mentor hours for detailed-level questions (e.g. coding related), and use professor office hours for conceptual level questions



# Grading scheme

- **Midterm Exam:** 15% (around March 10, evening)
- **Final Exam:** 15% (TBA)
- **Homework Assignments:** 70% (10 HWs)

**TWO** lowest homework scores are dropped from the final homework average calculation.  
This is for emergency, sickness, etc.

Homework is always due 9am on the specified date (mostly Tuesday).  
(Late submissions will not be accepted.)

Homework will be posted and submitted via Canvas.



# Regrade Request

Use Google Form (will be announced) for regrade request

Raised with the TAs within 72 hours after homework / exam is returned.



# Integrity

Just don't cheat at all. You'll be caught. It's not worth it.

You are encouraged to discuss with your peers, the TA or the instructors ideas, approaches and techniques broadly. However, all examinations, programming assignments, and written homeworks must be written up individually. For example, code for programming assignments must not be developed in groups, nor should code be shared. Make sure you work through all problems yourself, and that your final write-up is your own. If you feel your peer discussions are too deep for comfort, declare it in the homework solution: "I discussed with X,Y,Z the following specific ideas: A, B, C; therefore our solutions may have similarities on D, E, F...".

You may use books or legit online resources to help solve homework problems, but you must always credit all such sources in your writeup and you must never copy material verbatim.

We are aware that certain websites host previous years' CS540 homework assignments and solutions against the wish of instructors. Do not be tempted to use them: the solutions may contain "poisonous berries" previous instructors planted intentionally to catch cheating. If we catch you copy such solutions, you automatically fail.

Do not bother to obfuscate plagiarism (e.g. change variable names, code style, etc.) One application of AI is to develop sophisticated plagiarism detection techniques!

Cheating and plagiarism will be dealt with in accordance with University procedures (see the [UW-Madison Academic Misconduct Rules and Procedures](#))



# Quiz

1. Where can I find all the 540 stuff, if I didn't write down the URL?
2. I feel sick, should I still show up to class?
3. I can't finish my homework because I was traveling, I was sick, my dog ate it, etc. Can I ask for an extension?
4. Can I do homework with a group?



# Answers

1. Where can I find all the 540 stuff, if I didn't write down the URL? **Your Canvas has the main link.**
2. I feel sick, should I still show up to class? **NO. Study materials online.**
3. I can't finish my homework because I was traveling, I was sick, my dog ate it, etc. Can I ask for an extension? **No. But we discard 2 lowest hw scores.**
4. Can I do homework with a group? **Yes (and encouraged) for high level discussions. No for exact solutions.**





# Part III: Software



# Tools



- **Python**

- Everyone is using it in machine learning & data science
- Conda package manager (for simplicity)

- **Jupyter**

- So much easier to keep track of your experiments
- Obviously you should put longer code into modules



# Python for Java Pros (slides available on website)

## A Crash Course in Python

---

1. Why are we doing this in Python?
2. Where do I write Python code? How do I run it?
  - a. Online
  - b. Offline
3. What are the big differences between Java and Python



# Colab

- Go to [colab.research.google.com](https://colab.research.google.com)
- Activate the GPU supported runtime (this is a K80 GPU)



An aerial photograph of a server rack, showing a dense grid of server units with gold-colored components. The text "ML loves matrices and vectors" is overlaid in white, bold font across the center of the image.

**ML loves matrices and vectors**



# Access Elements

An element: [1, 2]

	0	1	2	3
0	1	2	3	4
1	5	6	7	8
2	9	10	11	12
3	13	14	15	16

A row: [1, :]

	0	1	2	3
0	1	2	3	4
1	5	6	7	8
2	9	10	11	12
3	13	14	15	16

A column:[:, 2]

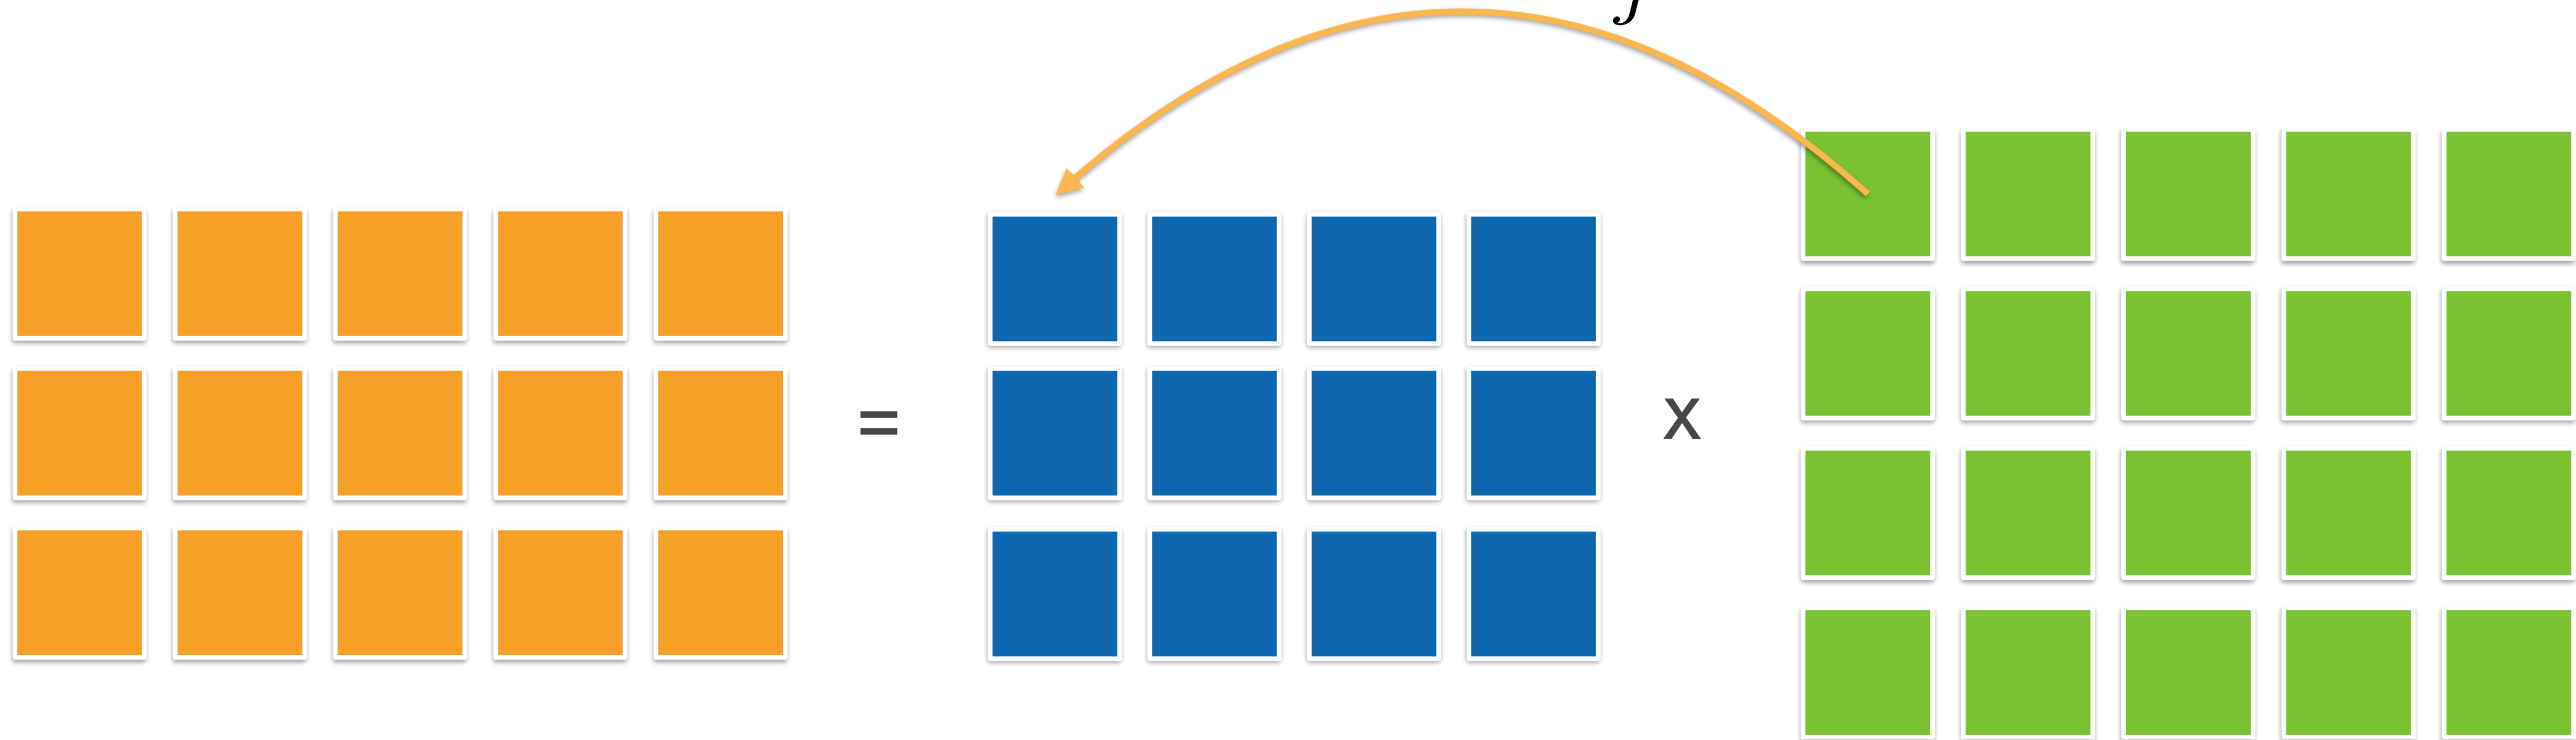
	0	1	2	3
0	1	2	3	4
1	5	6	7	8
2	9	10	11	12
3	13	14	15	16

	0	1	2	3
0	1	2	3	4
1	5	6	7	8
2	9	10	11	12
3	13	14	15	16



# Coming up: Probability and Linear Algebra Review

$$C = AB \text{ where } C_{ik} = \sum_j A_{ij} B_{jk}$$







# Recap

- What's in CS540
- Course logistics
- Software





**Thanks!**