



CS 540 Introduction to Artificial Intelligence

Course Overview

University of Wisconsin-Madison

Fall 2022



Today's outline

- What's in CS540
- Course logistics
- Software



Part I: Course overview

\$.



AI



Classify Images

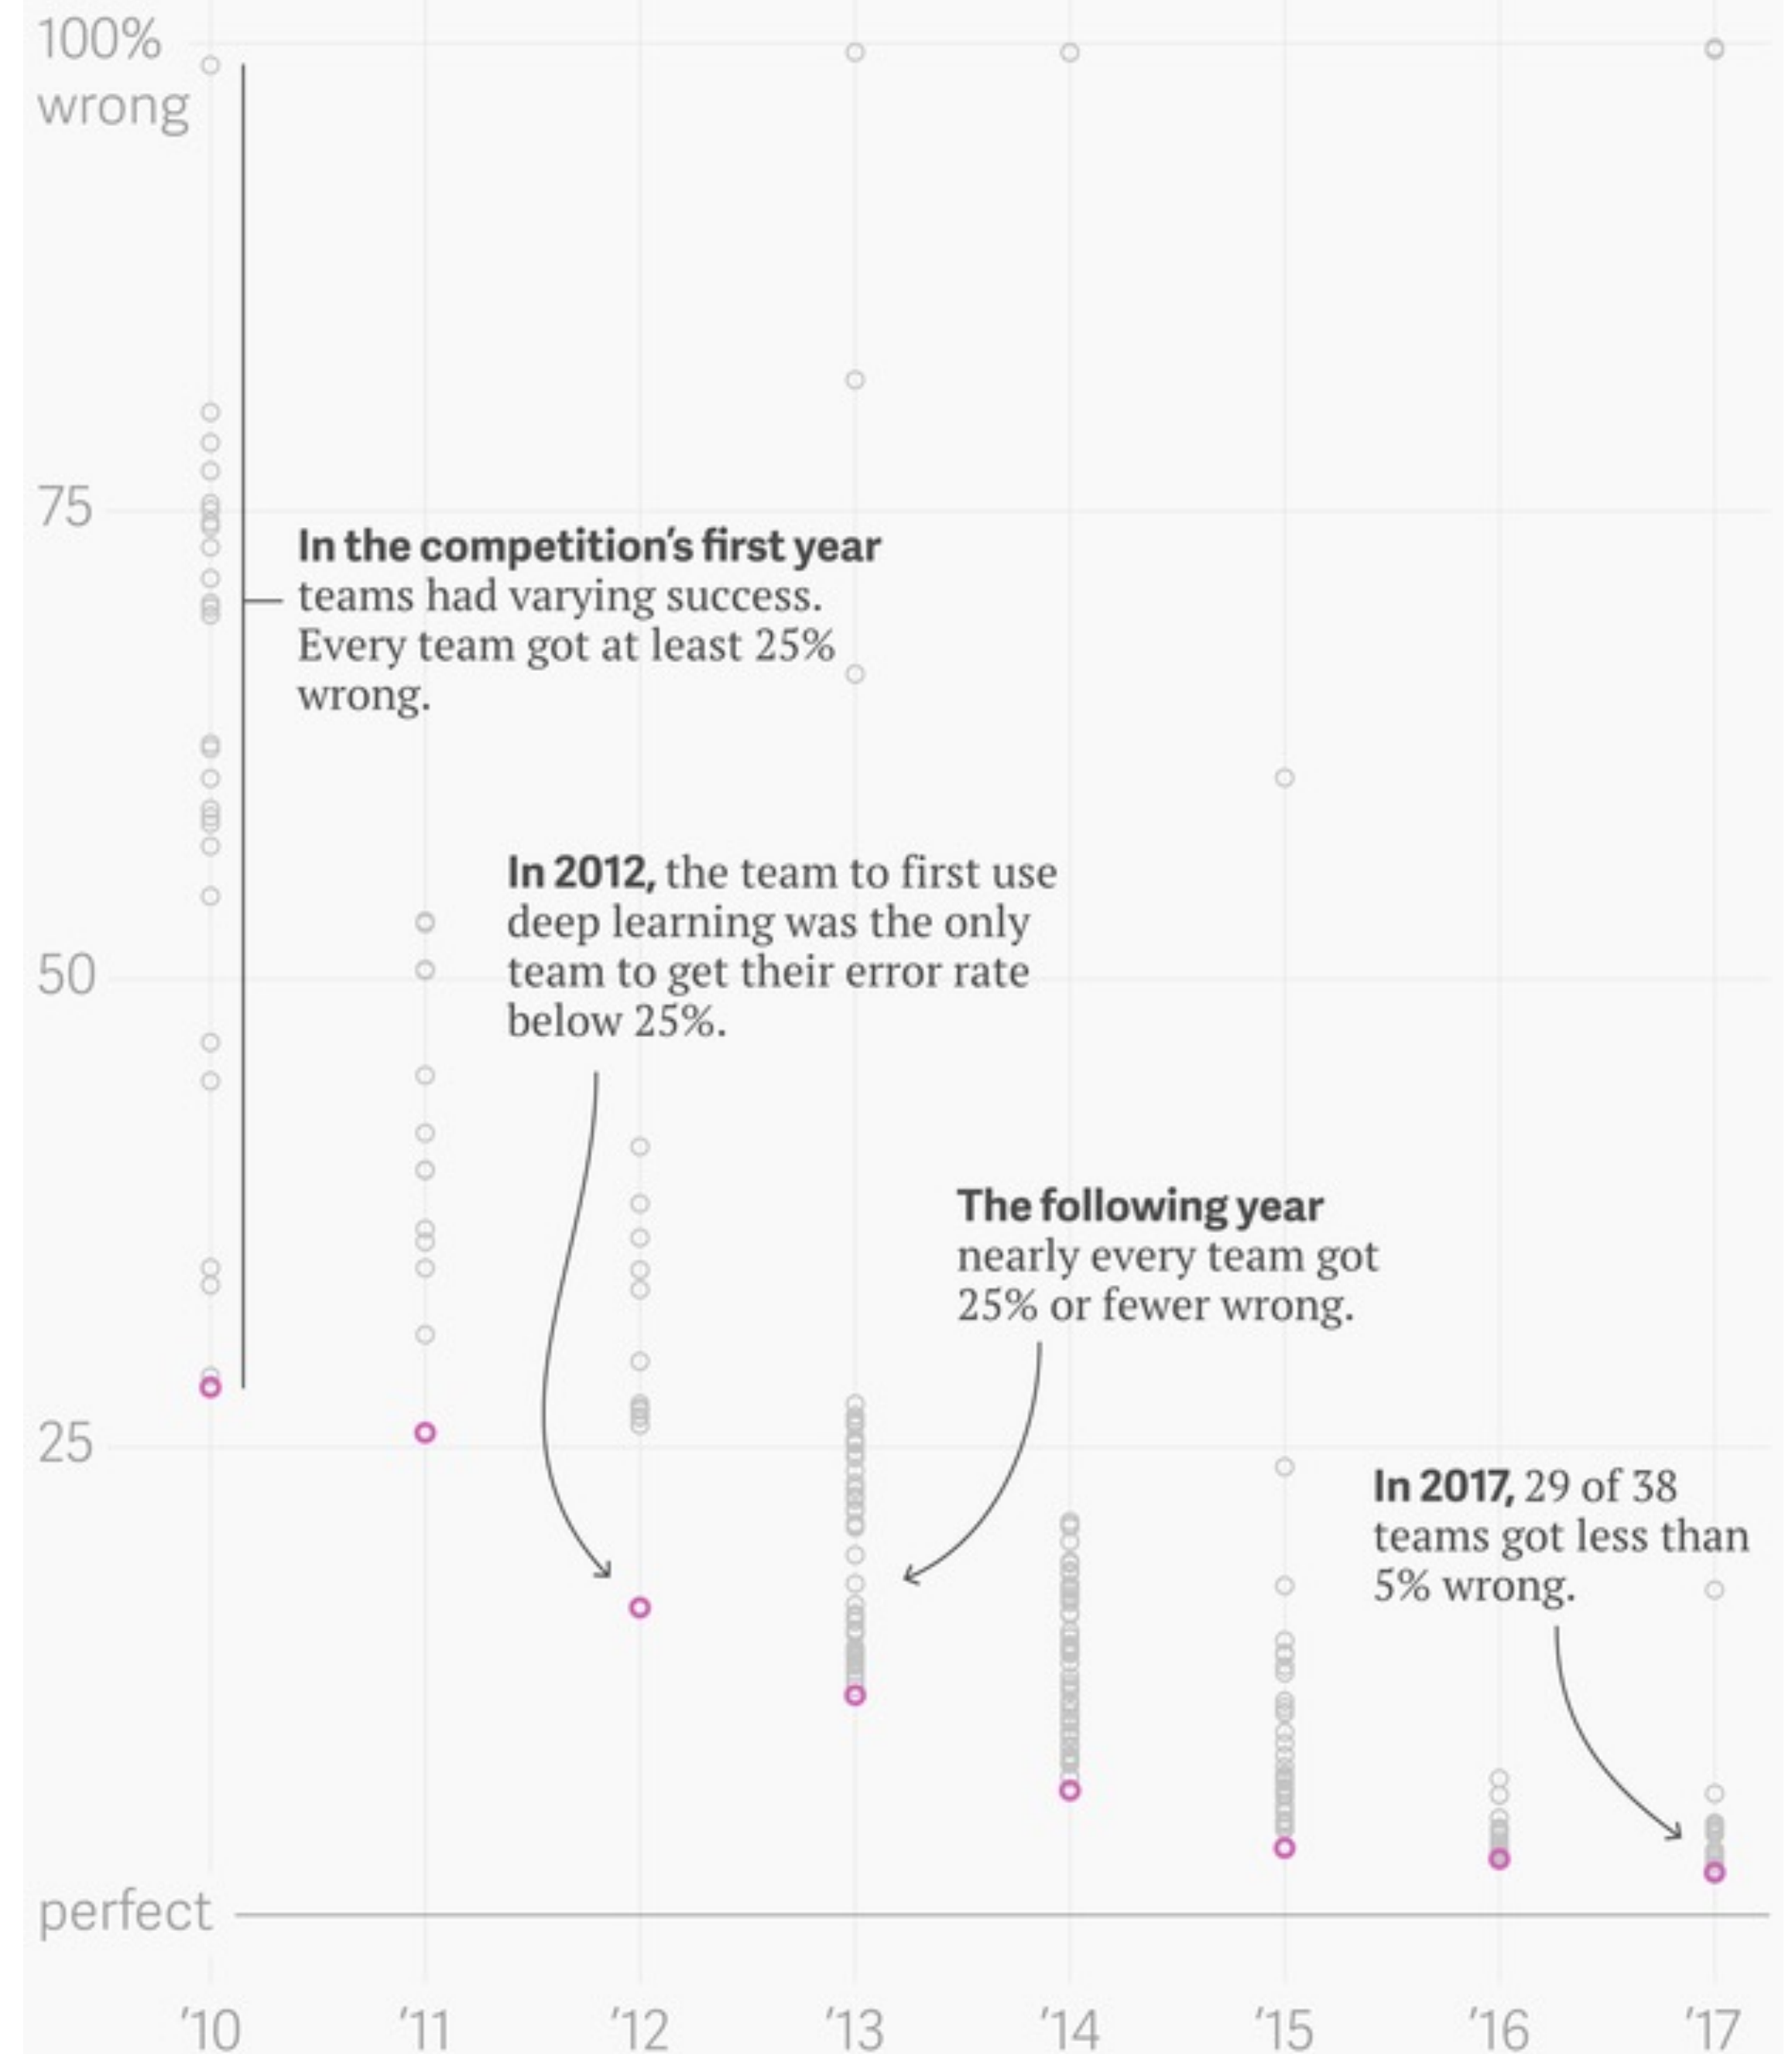
<http://www.image-net.org/>



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

Style Transfer

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



Generative Modeling

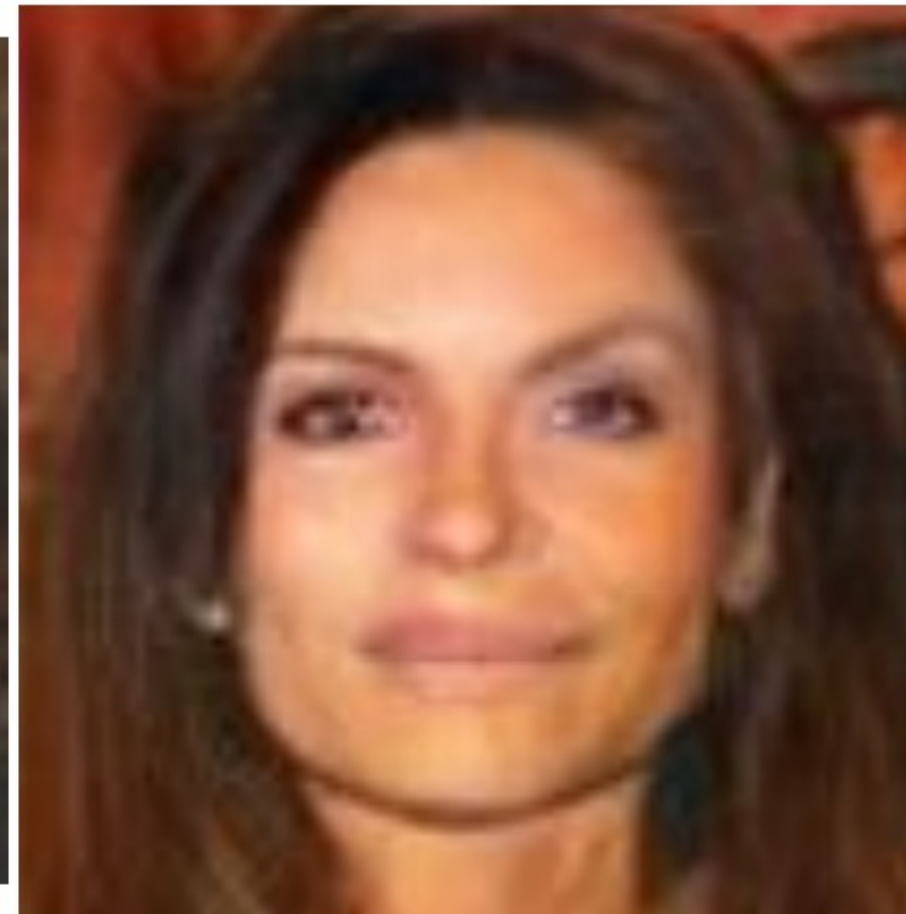
4.5 years of face generation



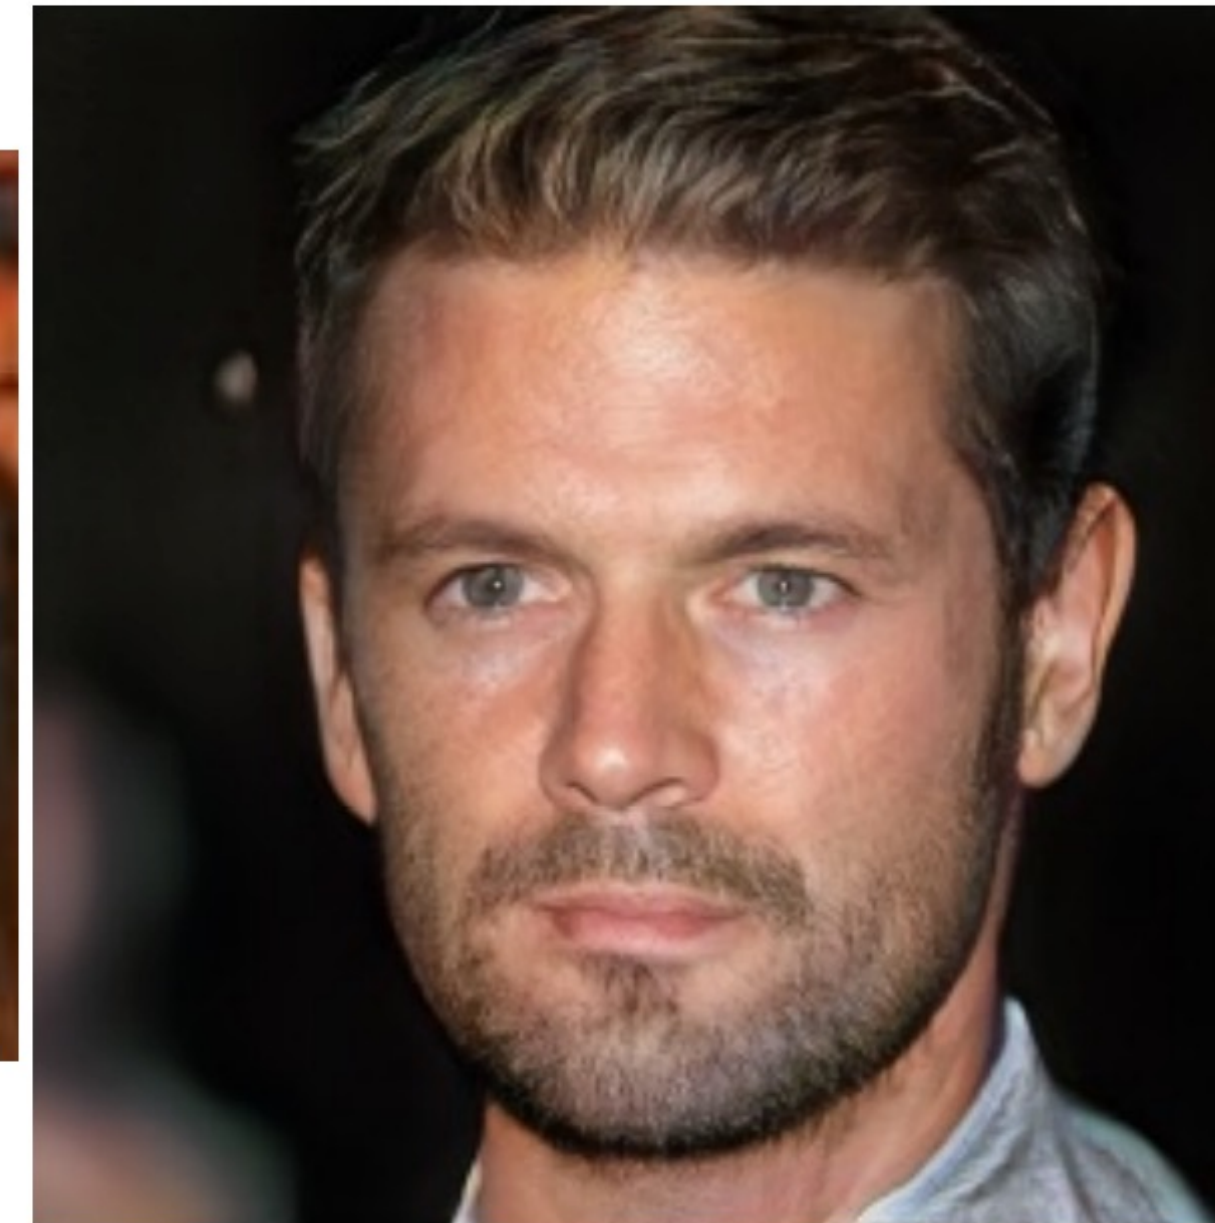
2014



2015



2016



2017



2018

Artificial Image Synthesis

Synthesize the images



Text to Image Generative Models

Stable Diffusion, etc.



Machine Translation

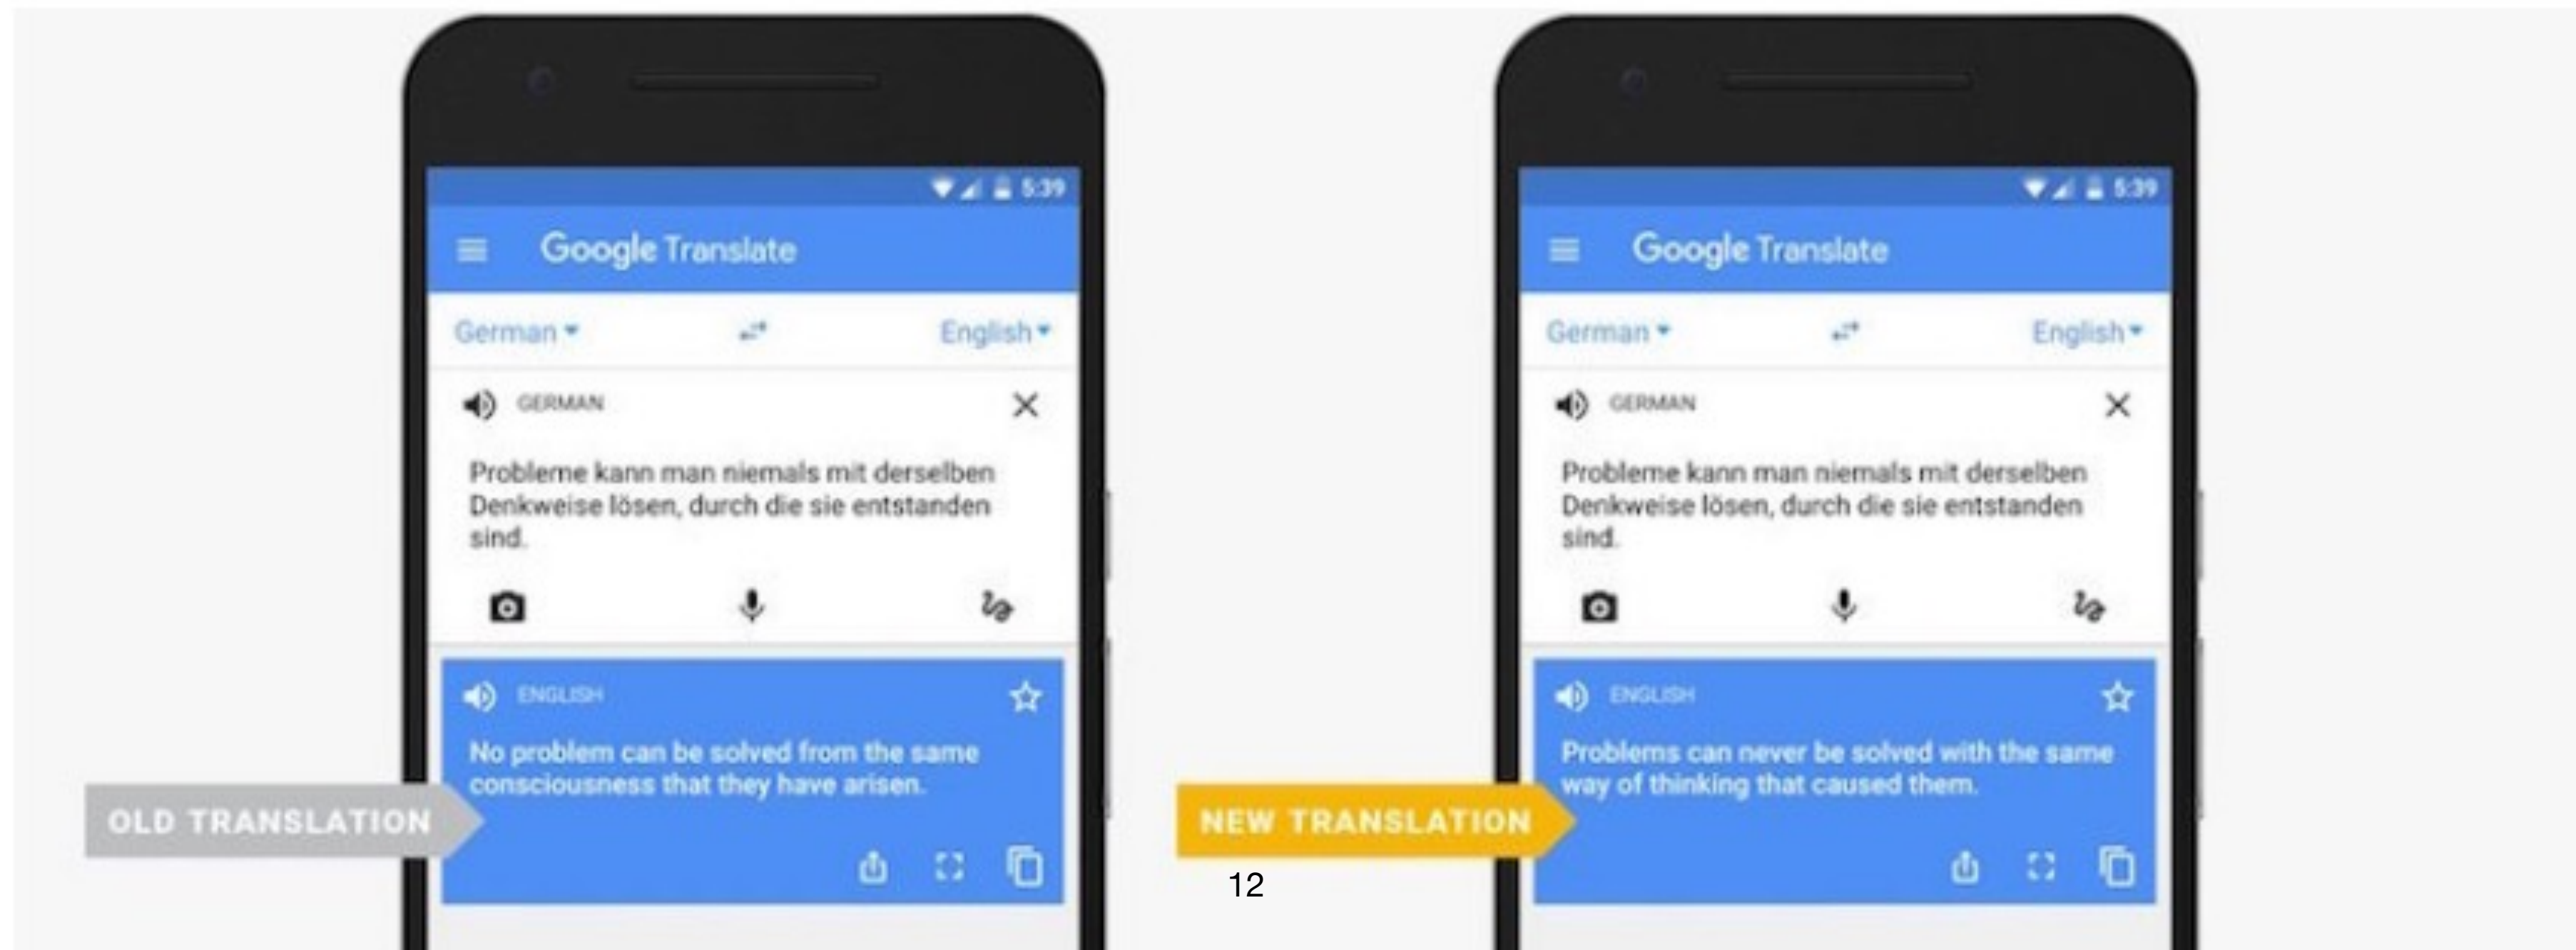
<https://www.pcmag.com/news/google-expands-neural-networks-for-language-translation>

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>

Open AI GPT-3: [Giving GPT-3 a Turing Test](#)

Q: What is your favorite animal?

A: My favorite animal is a dog.

Q: Why?

A: Because dogs are loyal and friendly.

Open AI GPT-3: [Giving GPT-3 a Turing Test](#)

Q: What is your favorite animal?

A: My favorite animal is a dog.

Q: Why?

A: Because dogs are loyal and friendly.

Q: What are two reasons that a dog might be in a bad mood?

A: Two reasons that a dog might be in a bad mood are if it is hungry or if it is hot.

Q: How many eyes does a giraffe have?

A: A giraffe has two eyes.

Q: How many legs does a frog have?

A: A frog has four legs.

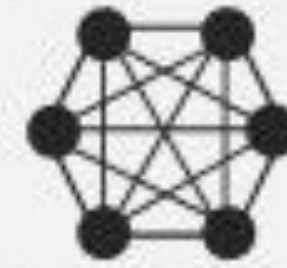
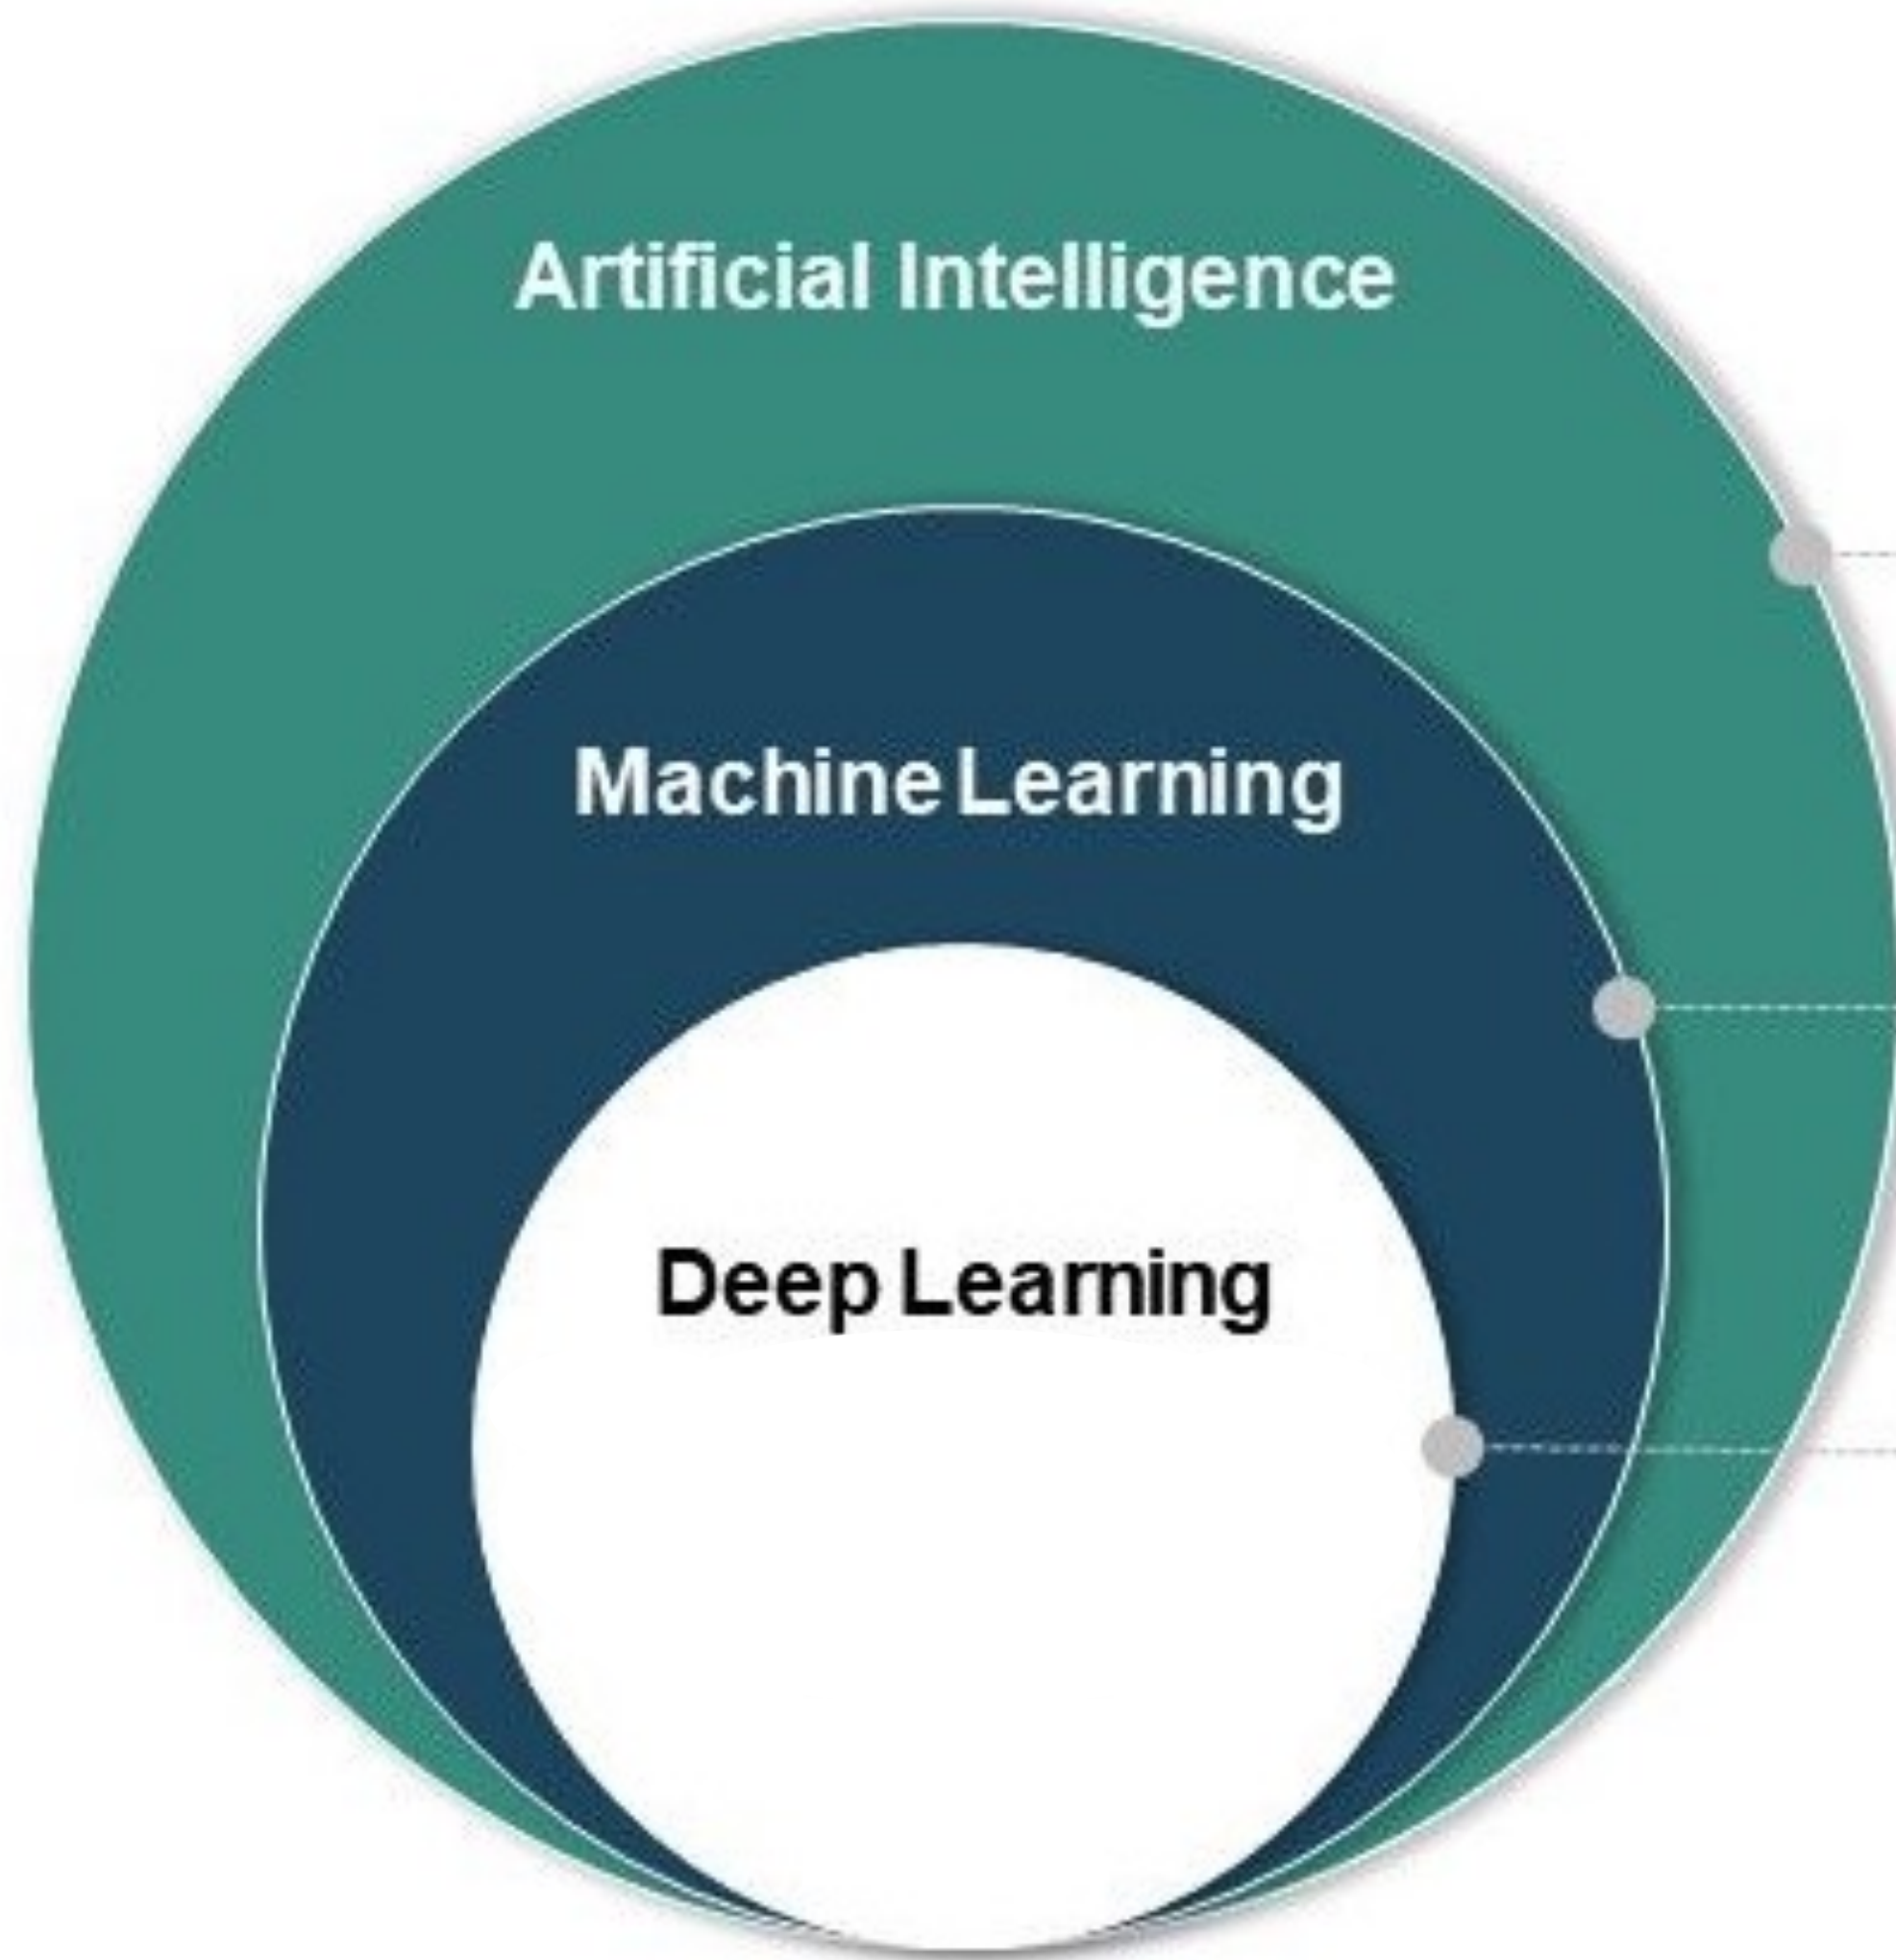
Q: Are there any animals with three legs?

A: No, there are no animals with three legs.

Q: Why don't animals have three legs?

A: Animals don't have three legs because they would fall over.

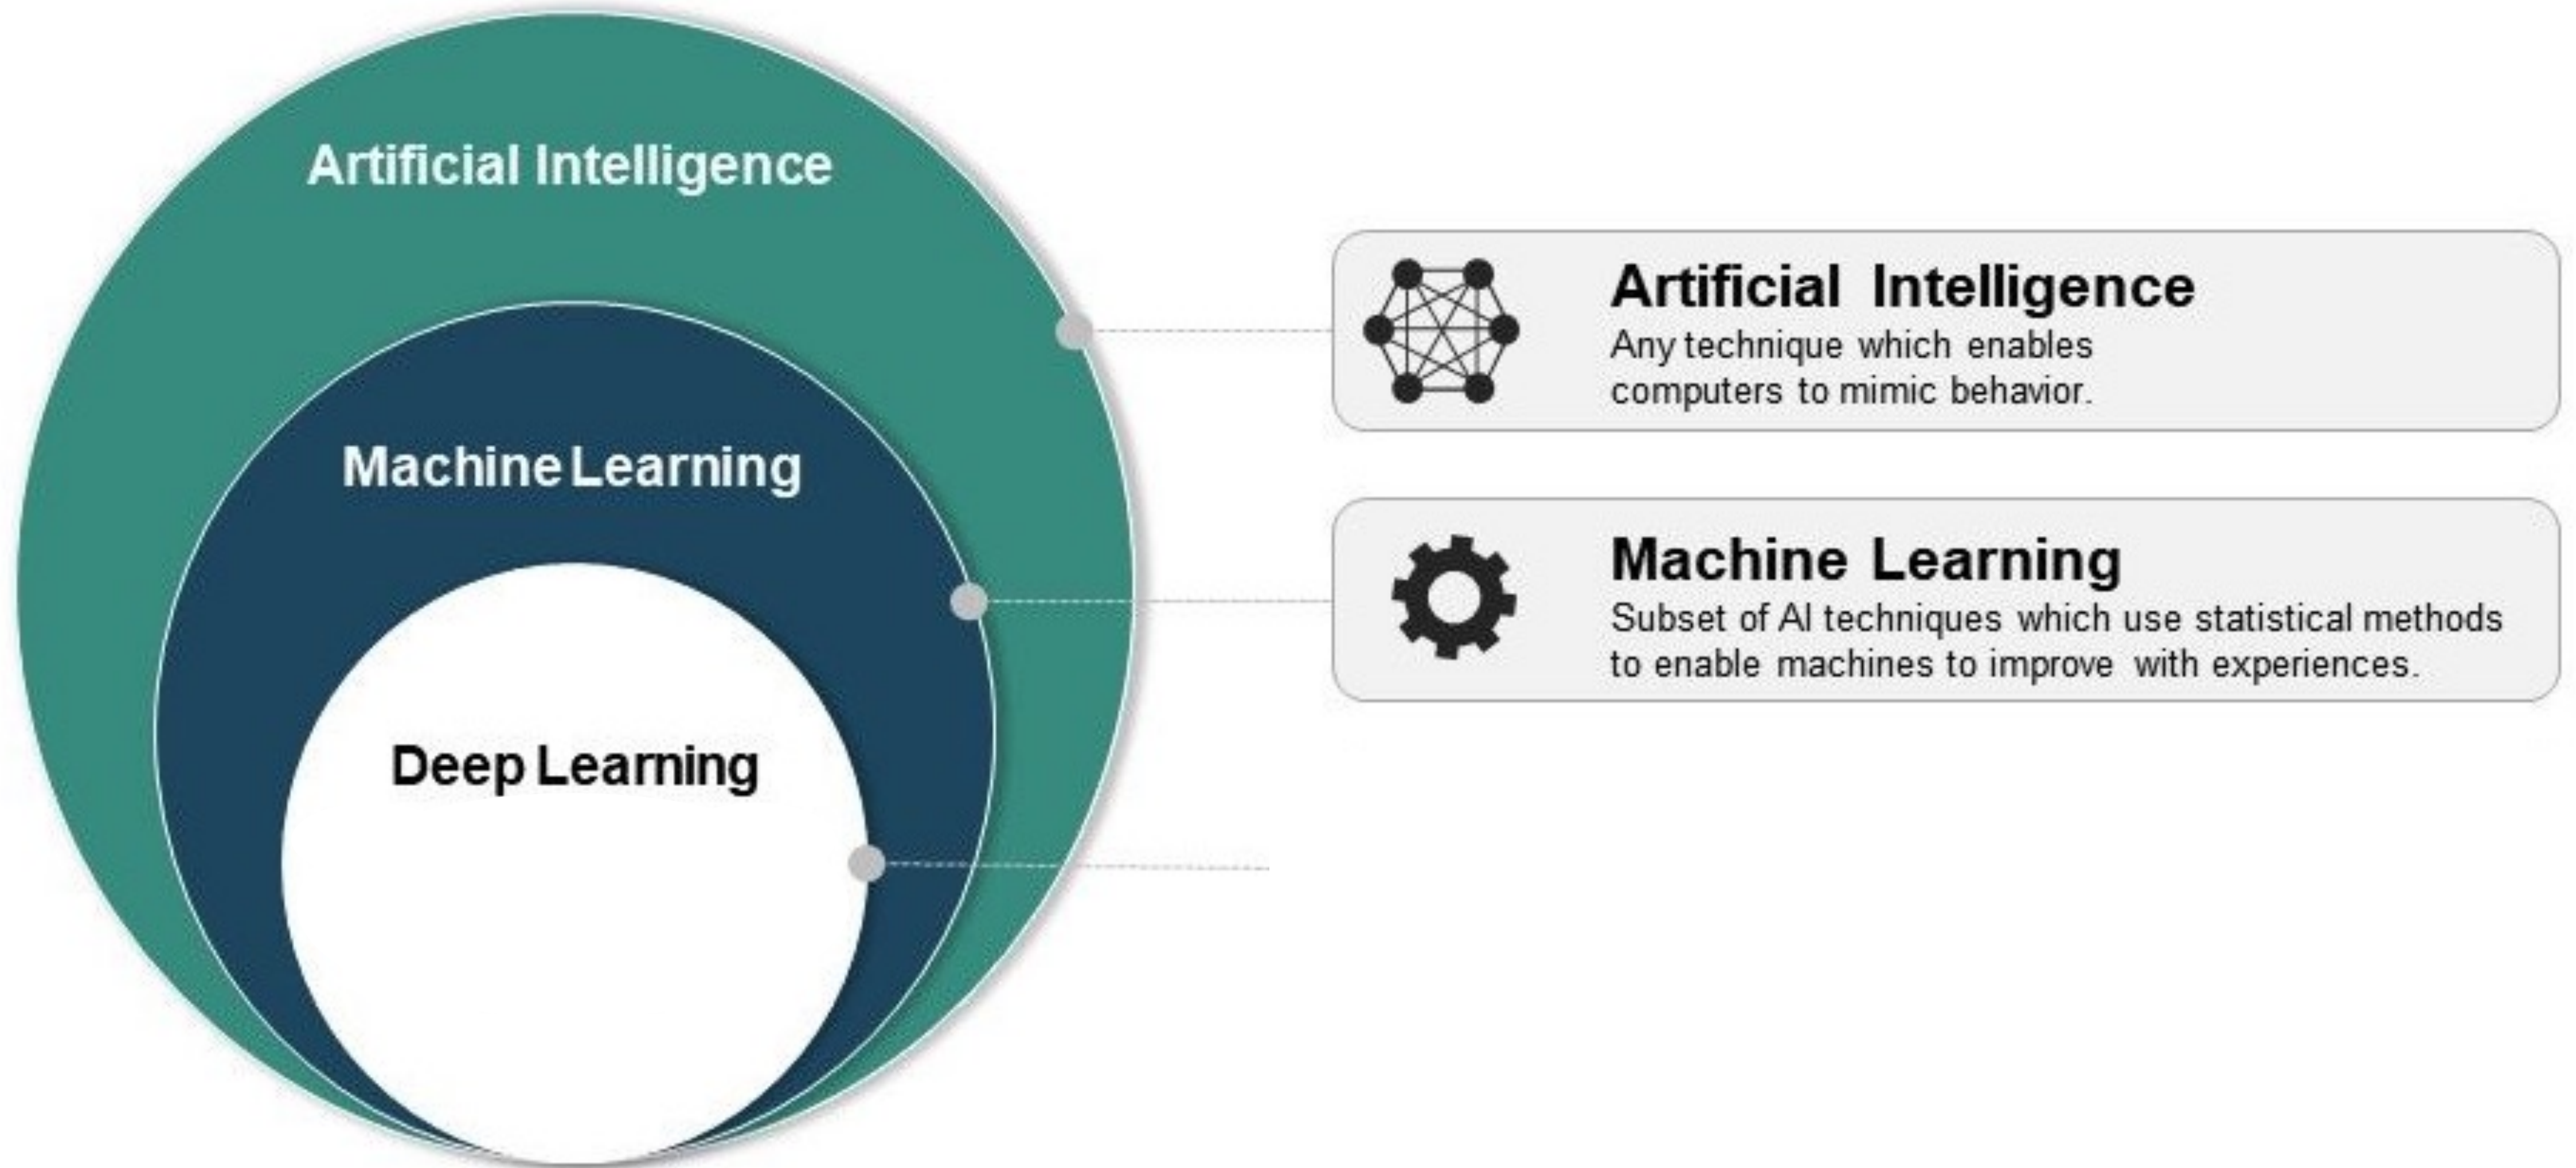
What is Artificial Intelligence (AI)?



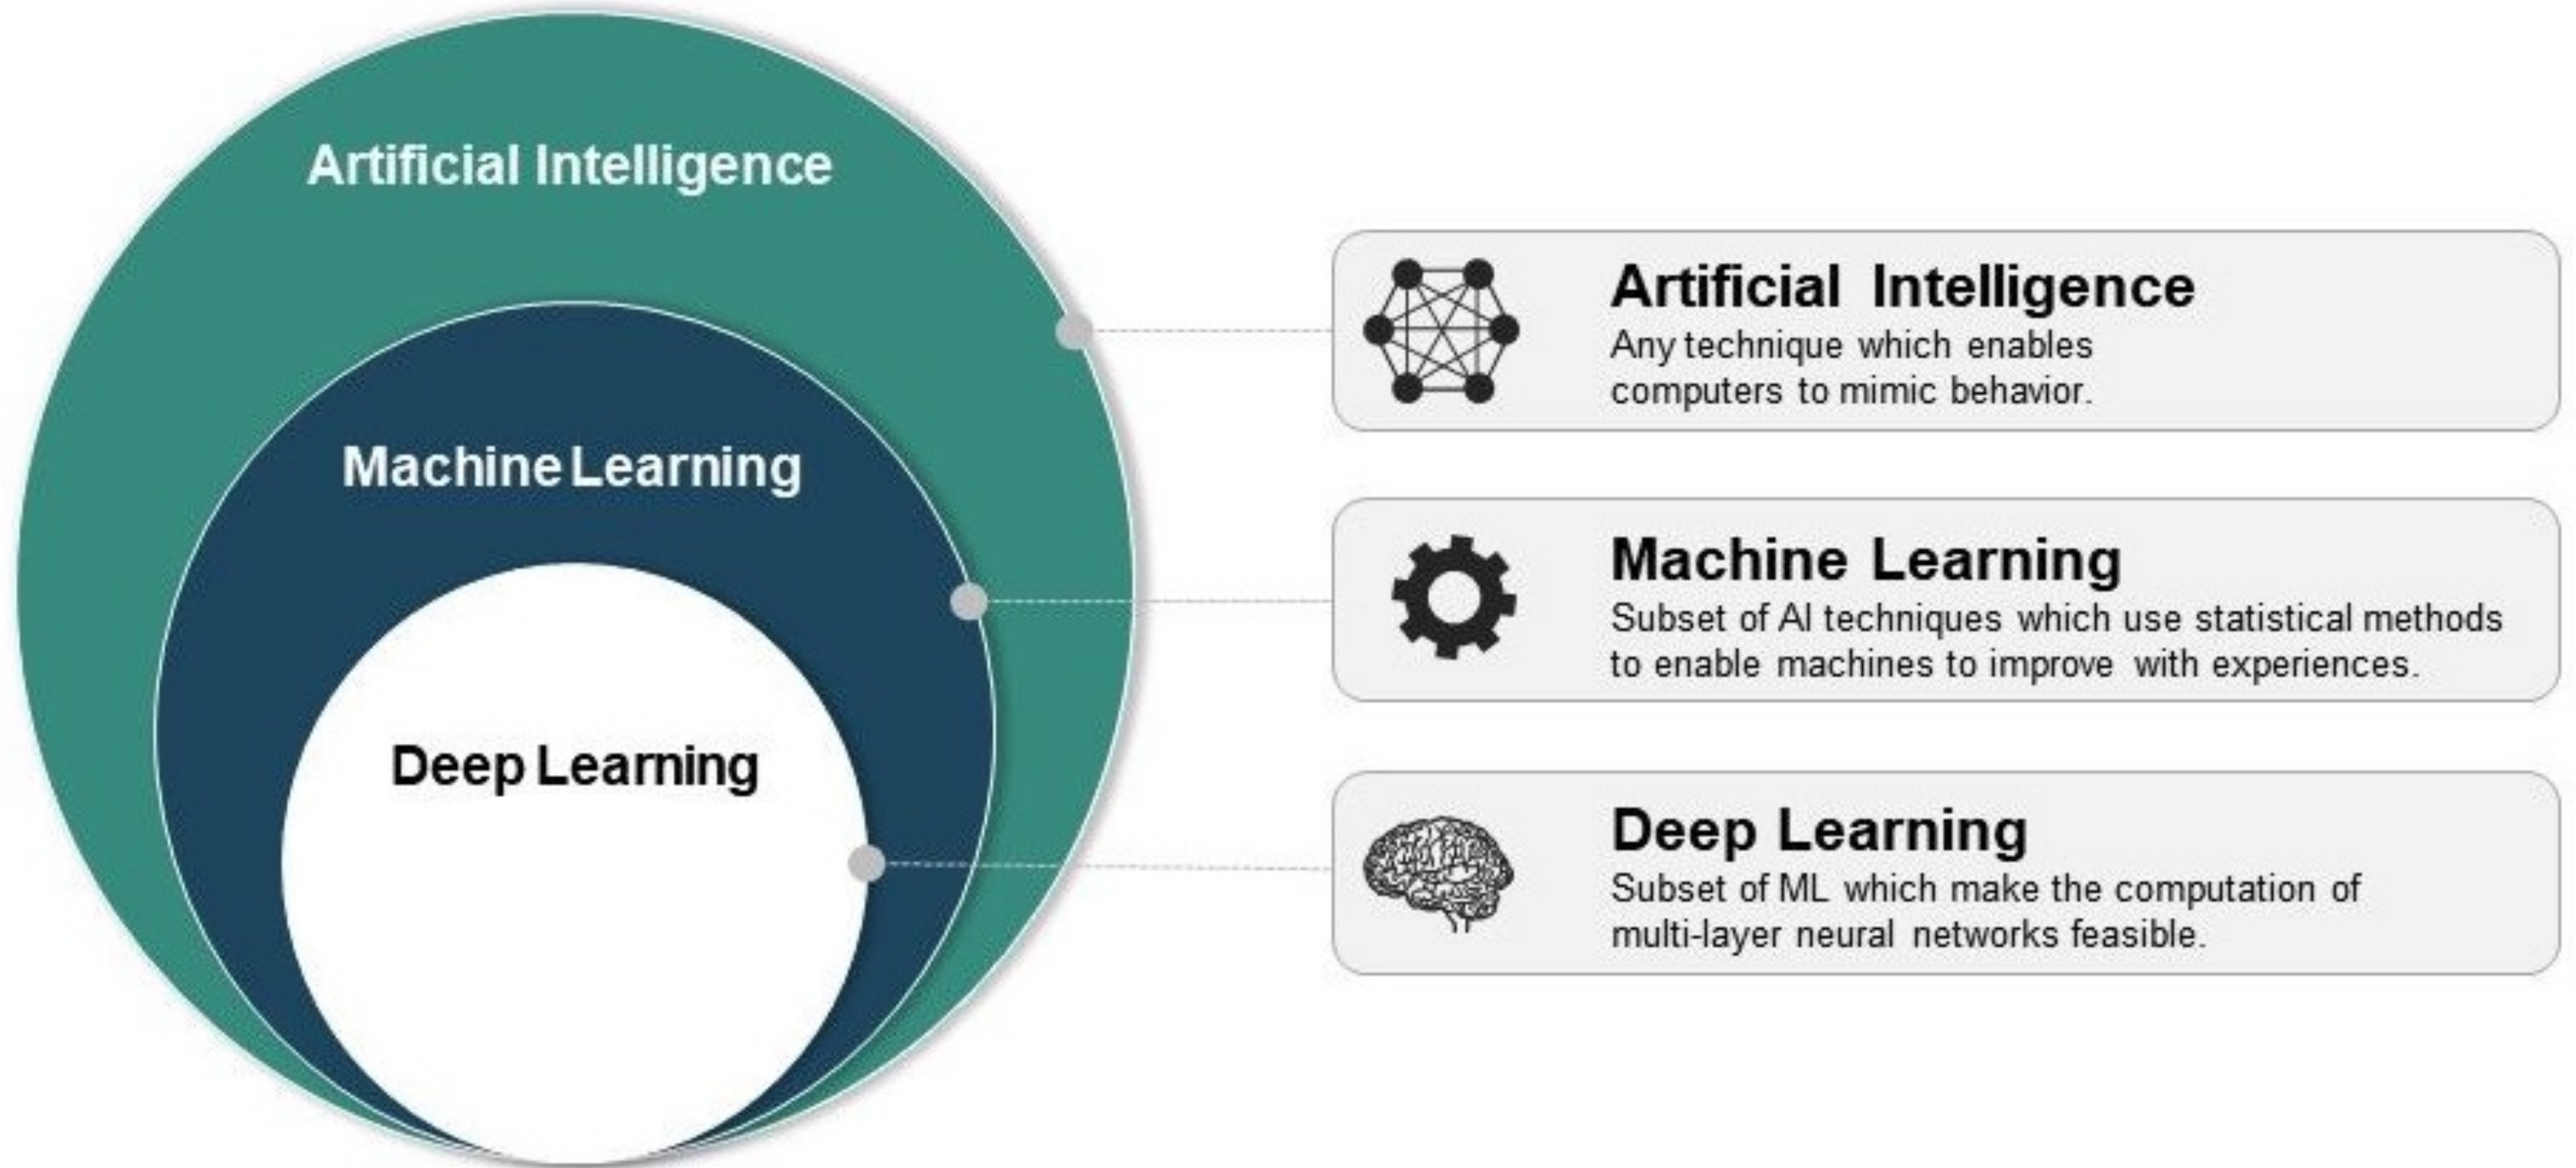
Artificial Intelligence

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/~kandasamy/courses/22fall-cs540/schedule.html>

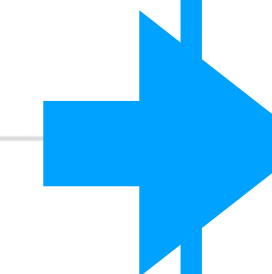
Date	Topic
Thursday, Sep 8	Welcome and Course Overview
Tuesday, Sep 13	Probability
Thursday, Sep 15	Linear Algebra and PCA
Tuesday, Sep 20	Statistics and Math Review
Thursday, Sep 22	Introduction to Logic
Tuesday, Sep 27	Natural Language Processing
Thursday, Sep 29	Machine Learning: Introduction
Tuesday, Oct 4	Machine Learning: Unsupervised Learning I
Thursday, Oct 6	Machine Learning: Unsupervised Learning II
Tuesday, Oct 11	Machine Learning: Linear regression
Thursday, Oct 13	Machine Learning: K - Nearest Neighbors & Naive Bayes
Tuesday, Oct 18	Machine Learning: Neural Network I (Perceptron)

Foundations

The diagram illustrates a matrix multiplication operation. On the left is a single orange square. This is followed by an equals sign, then a 3x4 grid of blue squares. To the right of this grid is an 'X' symbol, followed by a vertical column of four green squares. An orange arrow points from the word 'Foundations' to the top-right corner of the blue grid.

What you can learn from CS540?

Date	Topic
Thursday, Sep 8	Welcome and Course Overview
Tuesday, Sep 13	Probability
Thursday, Sep 15	Linear Algebra and PCA
Tuesday, Sep 20	Statistics and Math Review
Thursday, Sep 22	Introduction to Logic
Tuesday, Sep 27	Natural Language Processing
Thursday, Sep 29	Machine Learning: Introduction
Tuesday, Oct 4	Machine Learning: Unsupervised Learning I
Thursday, Oct 6	Machine Learning: Unsupervised Learning II
Tuesday, Oct 11	Machine Learning: Linear regression
Thursday, Oct 13	Machine Learning: K - Nearest Neighbors & Naive Bayes
Tuesday, Oct 18	Machine Learning: Neural Network I (Perceptron)

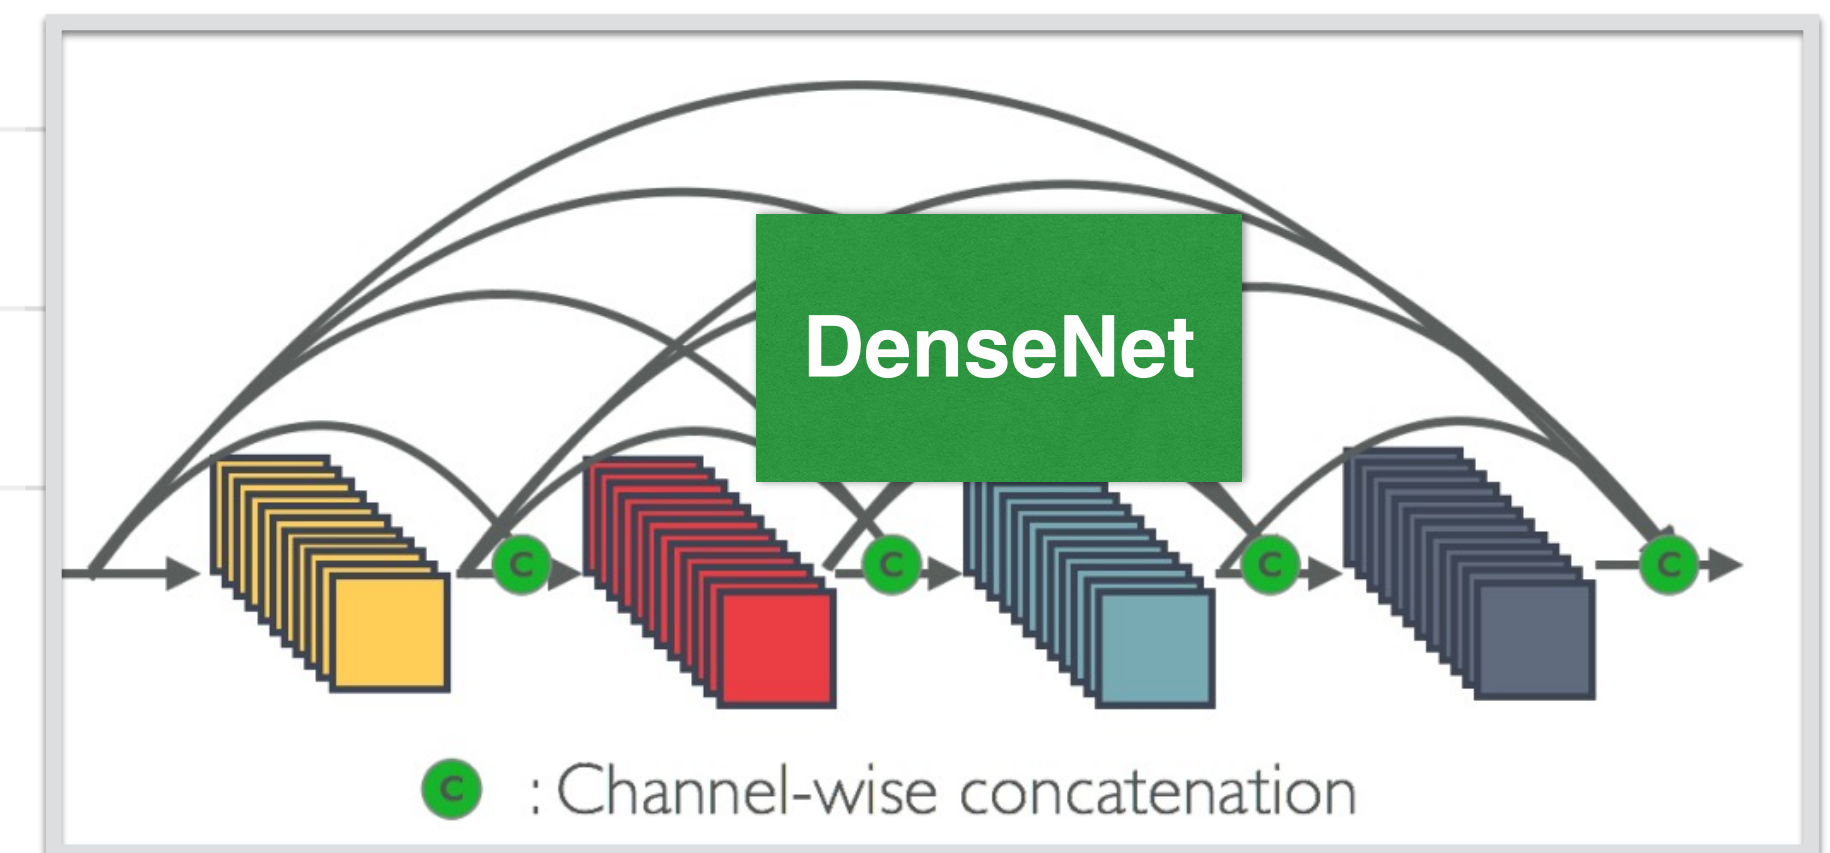
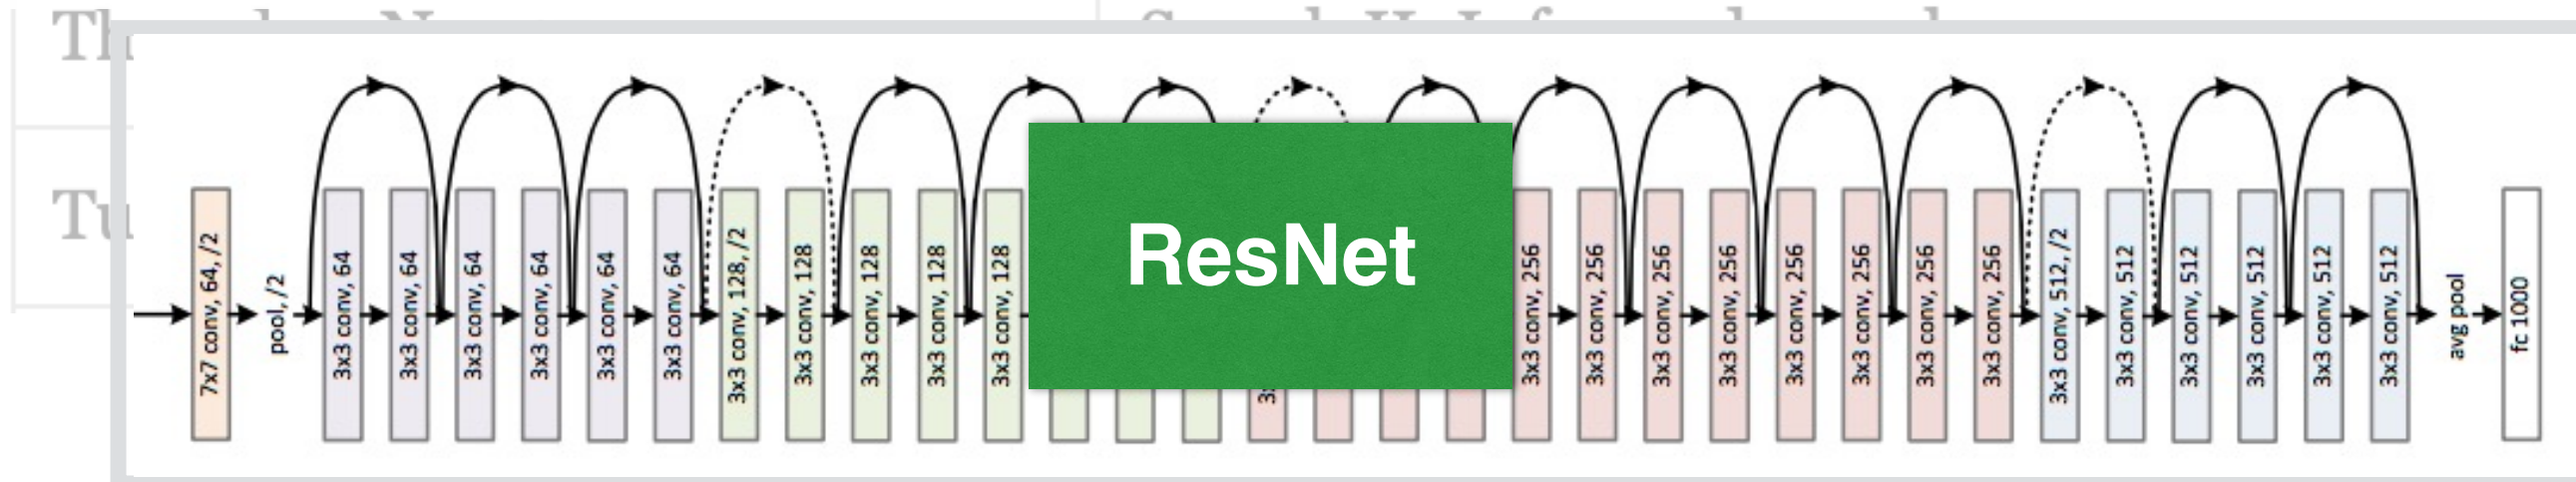


Machine learning

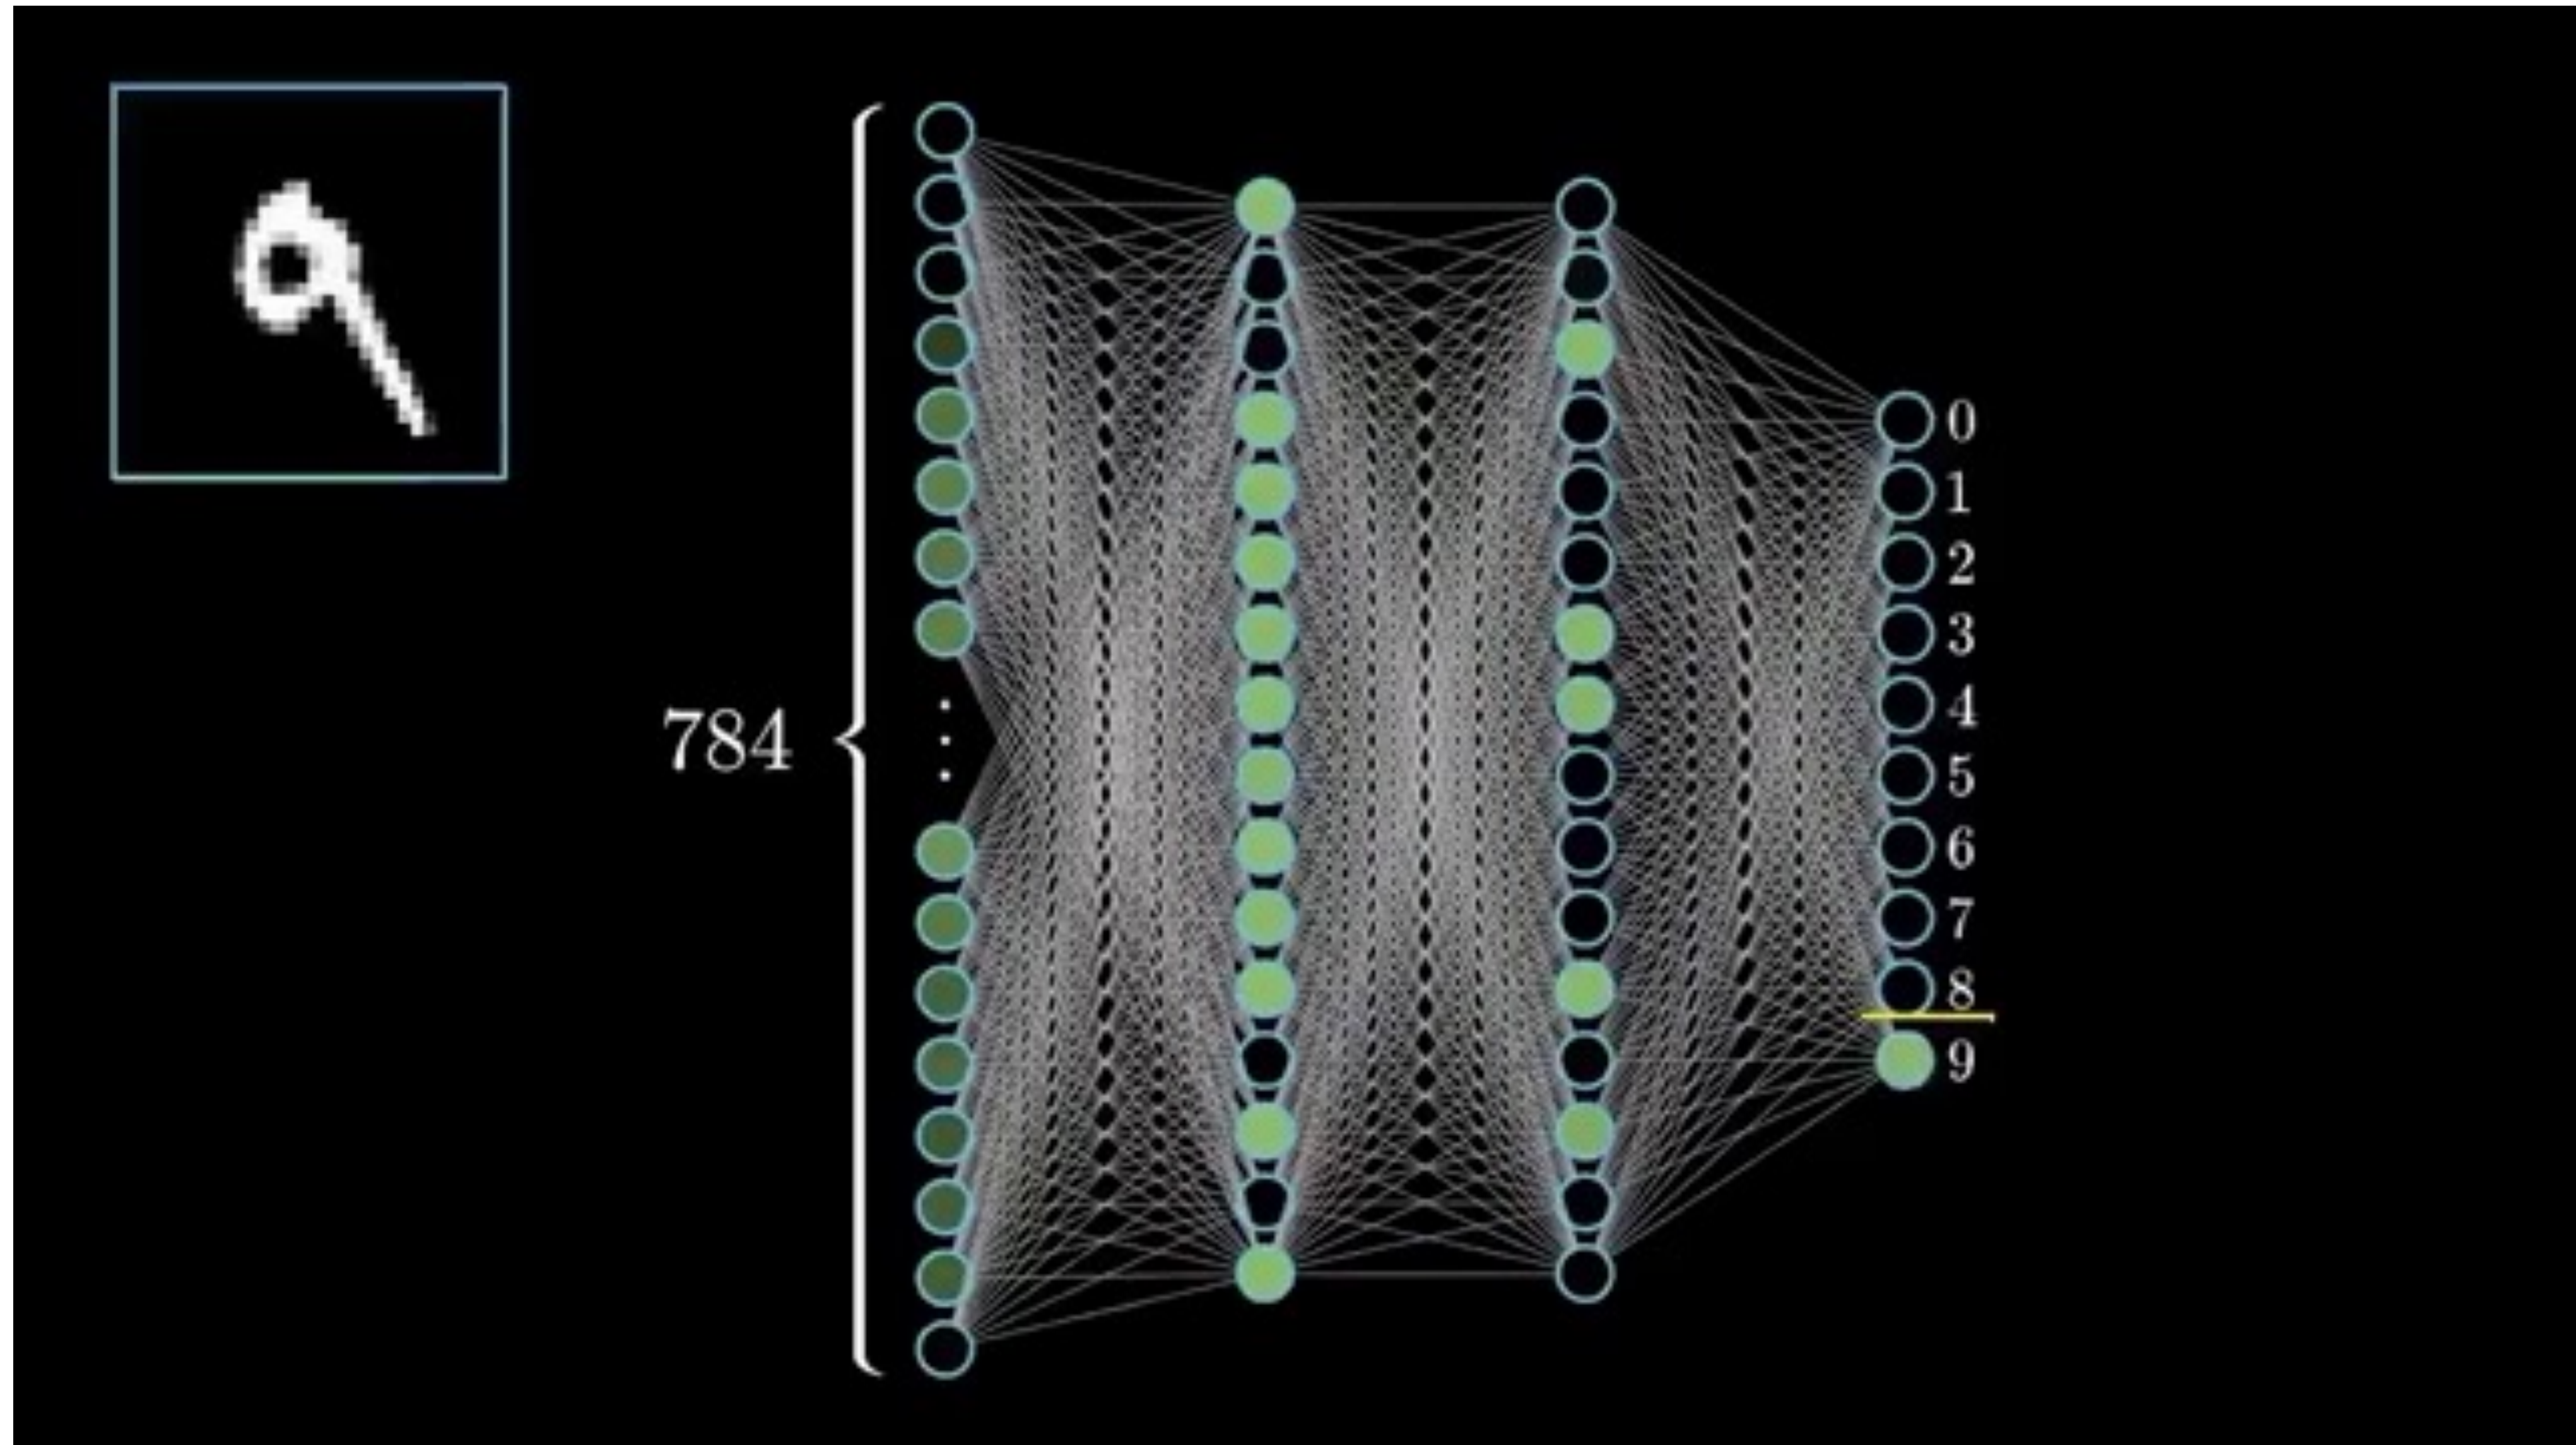
What you can learn from CS540?

Tuesday, Oct 18	Machine Learning: Neural Network I (Perceptron)
Thursday, Oct 20	Machine Learning: Neural Network II
Tuesday, Oct 25	Machine Learning: Neural Network III
Tuesday, Nov 1	Machine Learning: Deep Learning I
Thursday, Nov 3	Machine Learning: Deep Learning II
Tuesday, Nov 8	Machine Learning: Deep Learning III
Thursday, Nov 10	Machine Learning: Deep Learning and Neural Network's Summary
Tuesday, Nov 15	Search I: Un-Informed search

Deep Learning

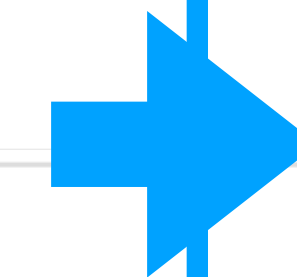


What you can learn from CS540?



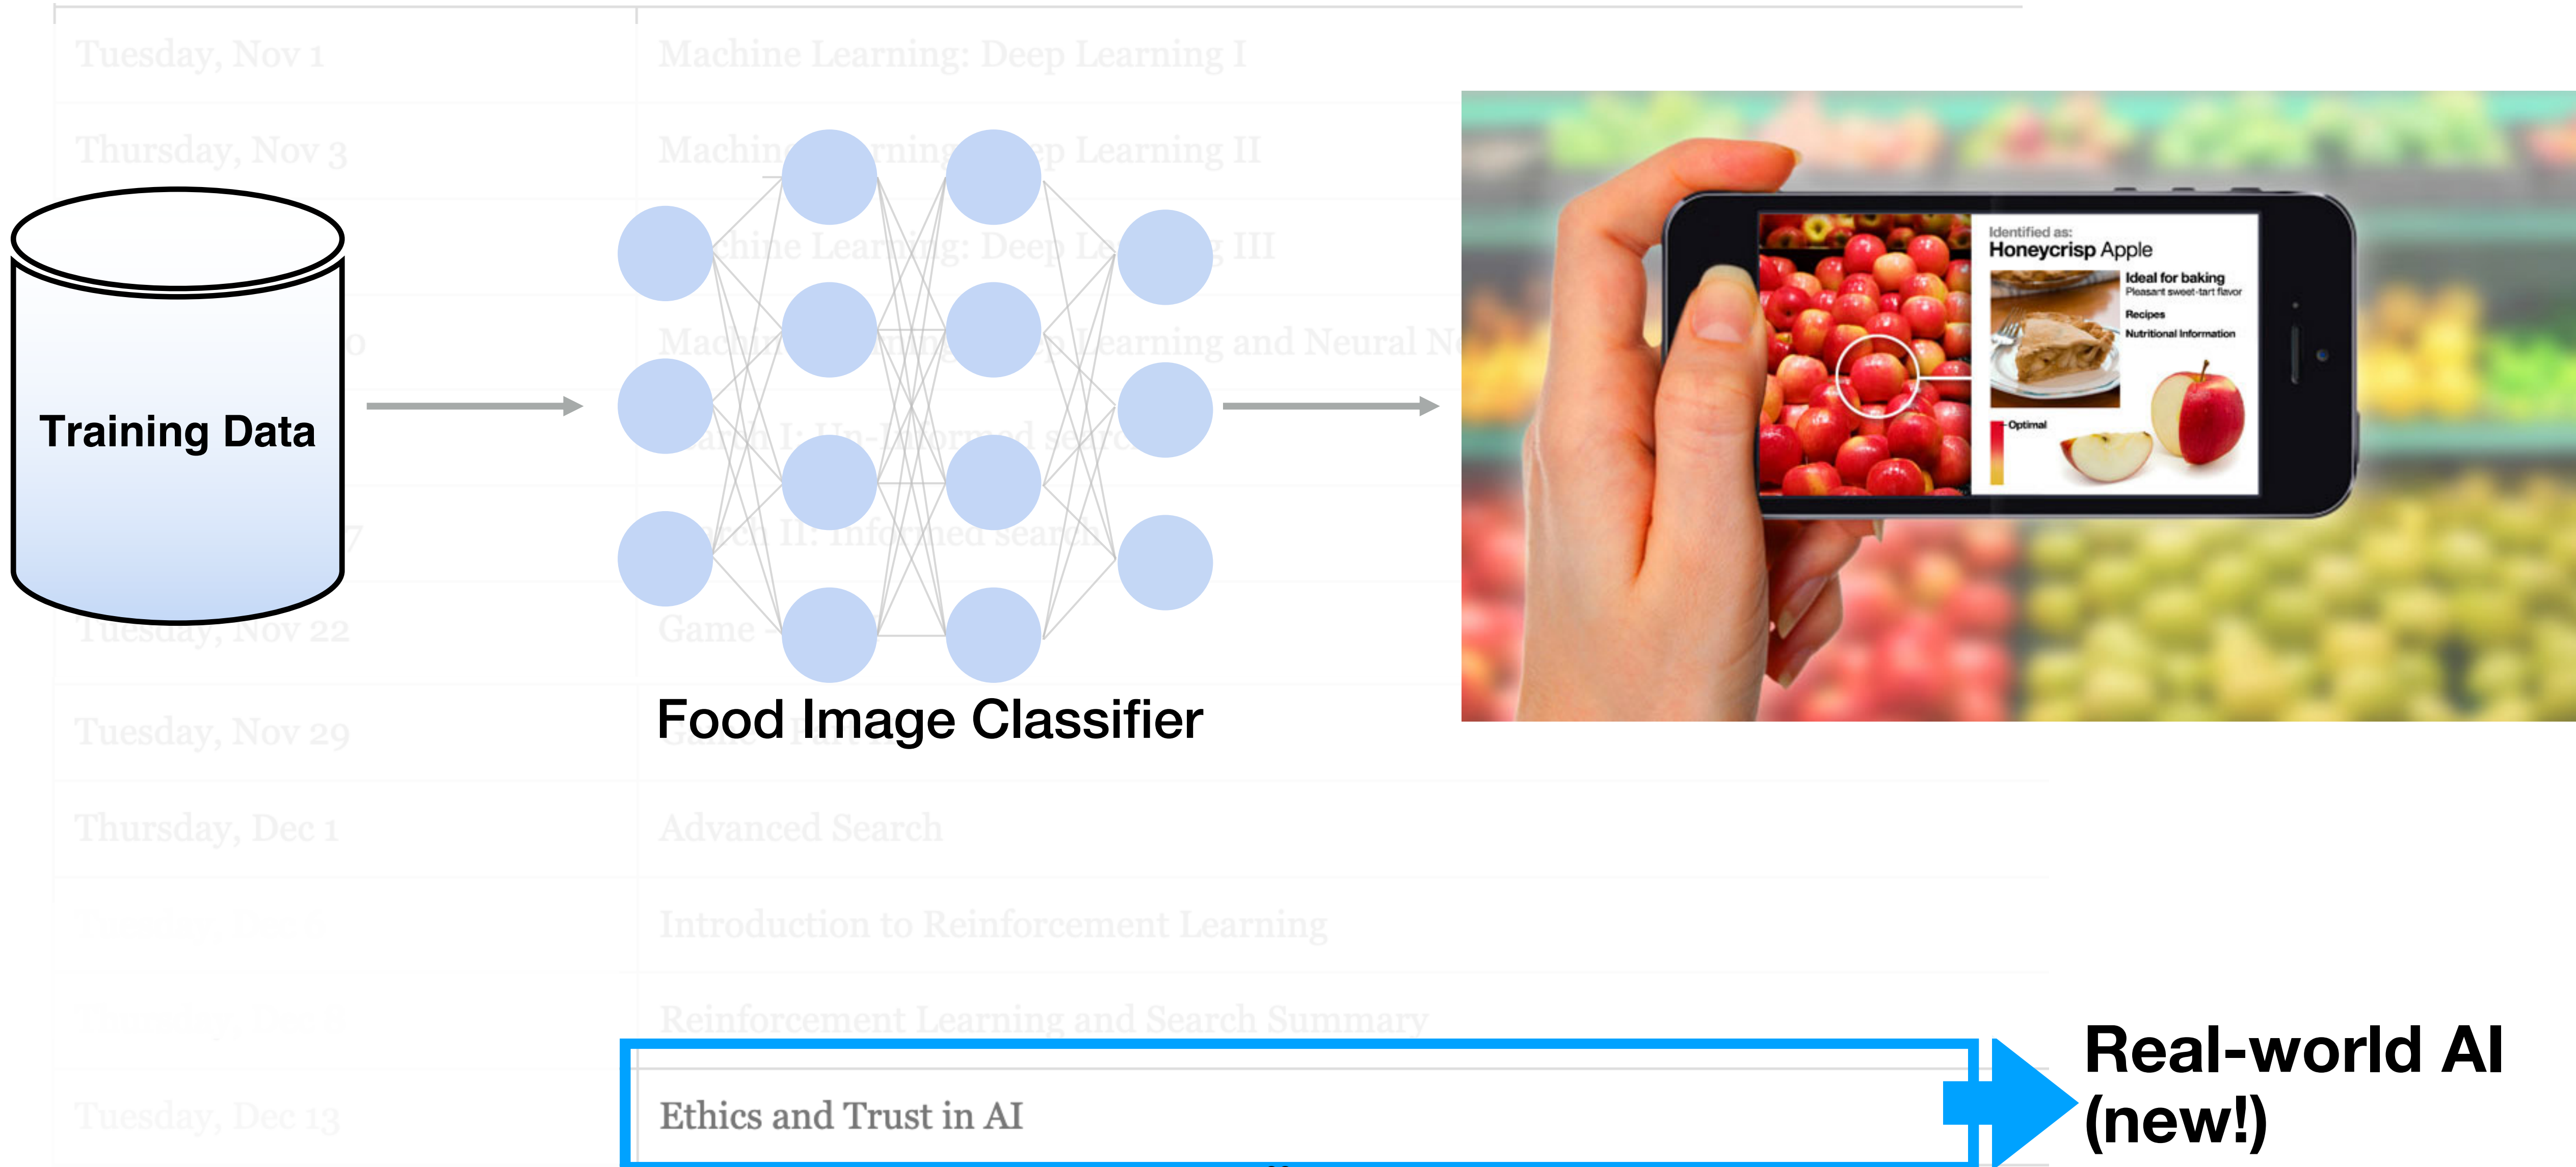
What you can learn from CS540?

Tuesday, Nov 1	Machine Learning: Deep Learning I
Thursday, Nov 3	Machine Learning: Deep Learning II
Tuesday, Nov 8	Machine Learning: Deep Learning III
Thursday, Nov 10	Machine Learning: Deep Learning and Neural Network's Summary
Tuesday, Nov 15	Search I: Un-Informed search
Thursday, Nov 17	Search II: Informed search
Tuesday, Nov 22	Game - Part I
Tuesday, Nov 29	Game - Part II
Thursday, Dec 1	Advanced Search
Tuesday, Dec 6	Introduction to Reinforcement Learning
Thursday, Dec 8	Reinforcement Learning and Search Summary
Tuesday, Dec 13	Ethics and Trust in AI



**Game, search
and
Reinforcement
Learning**

What you can learn from CS540?





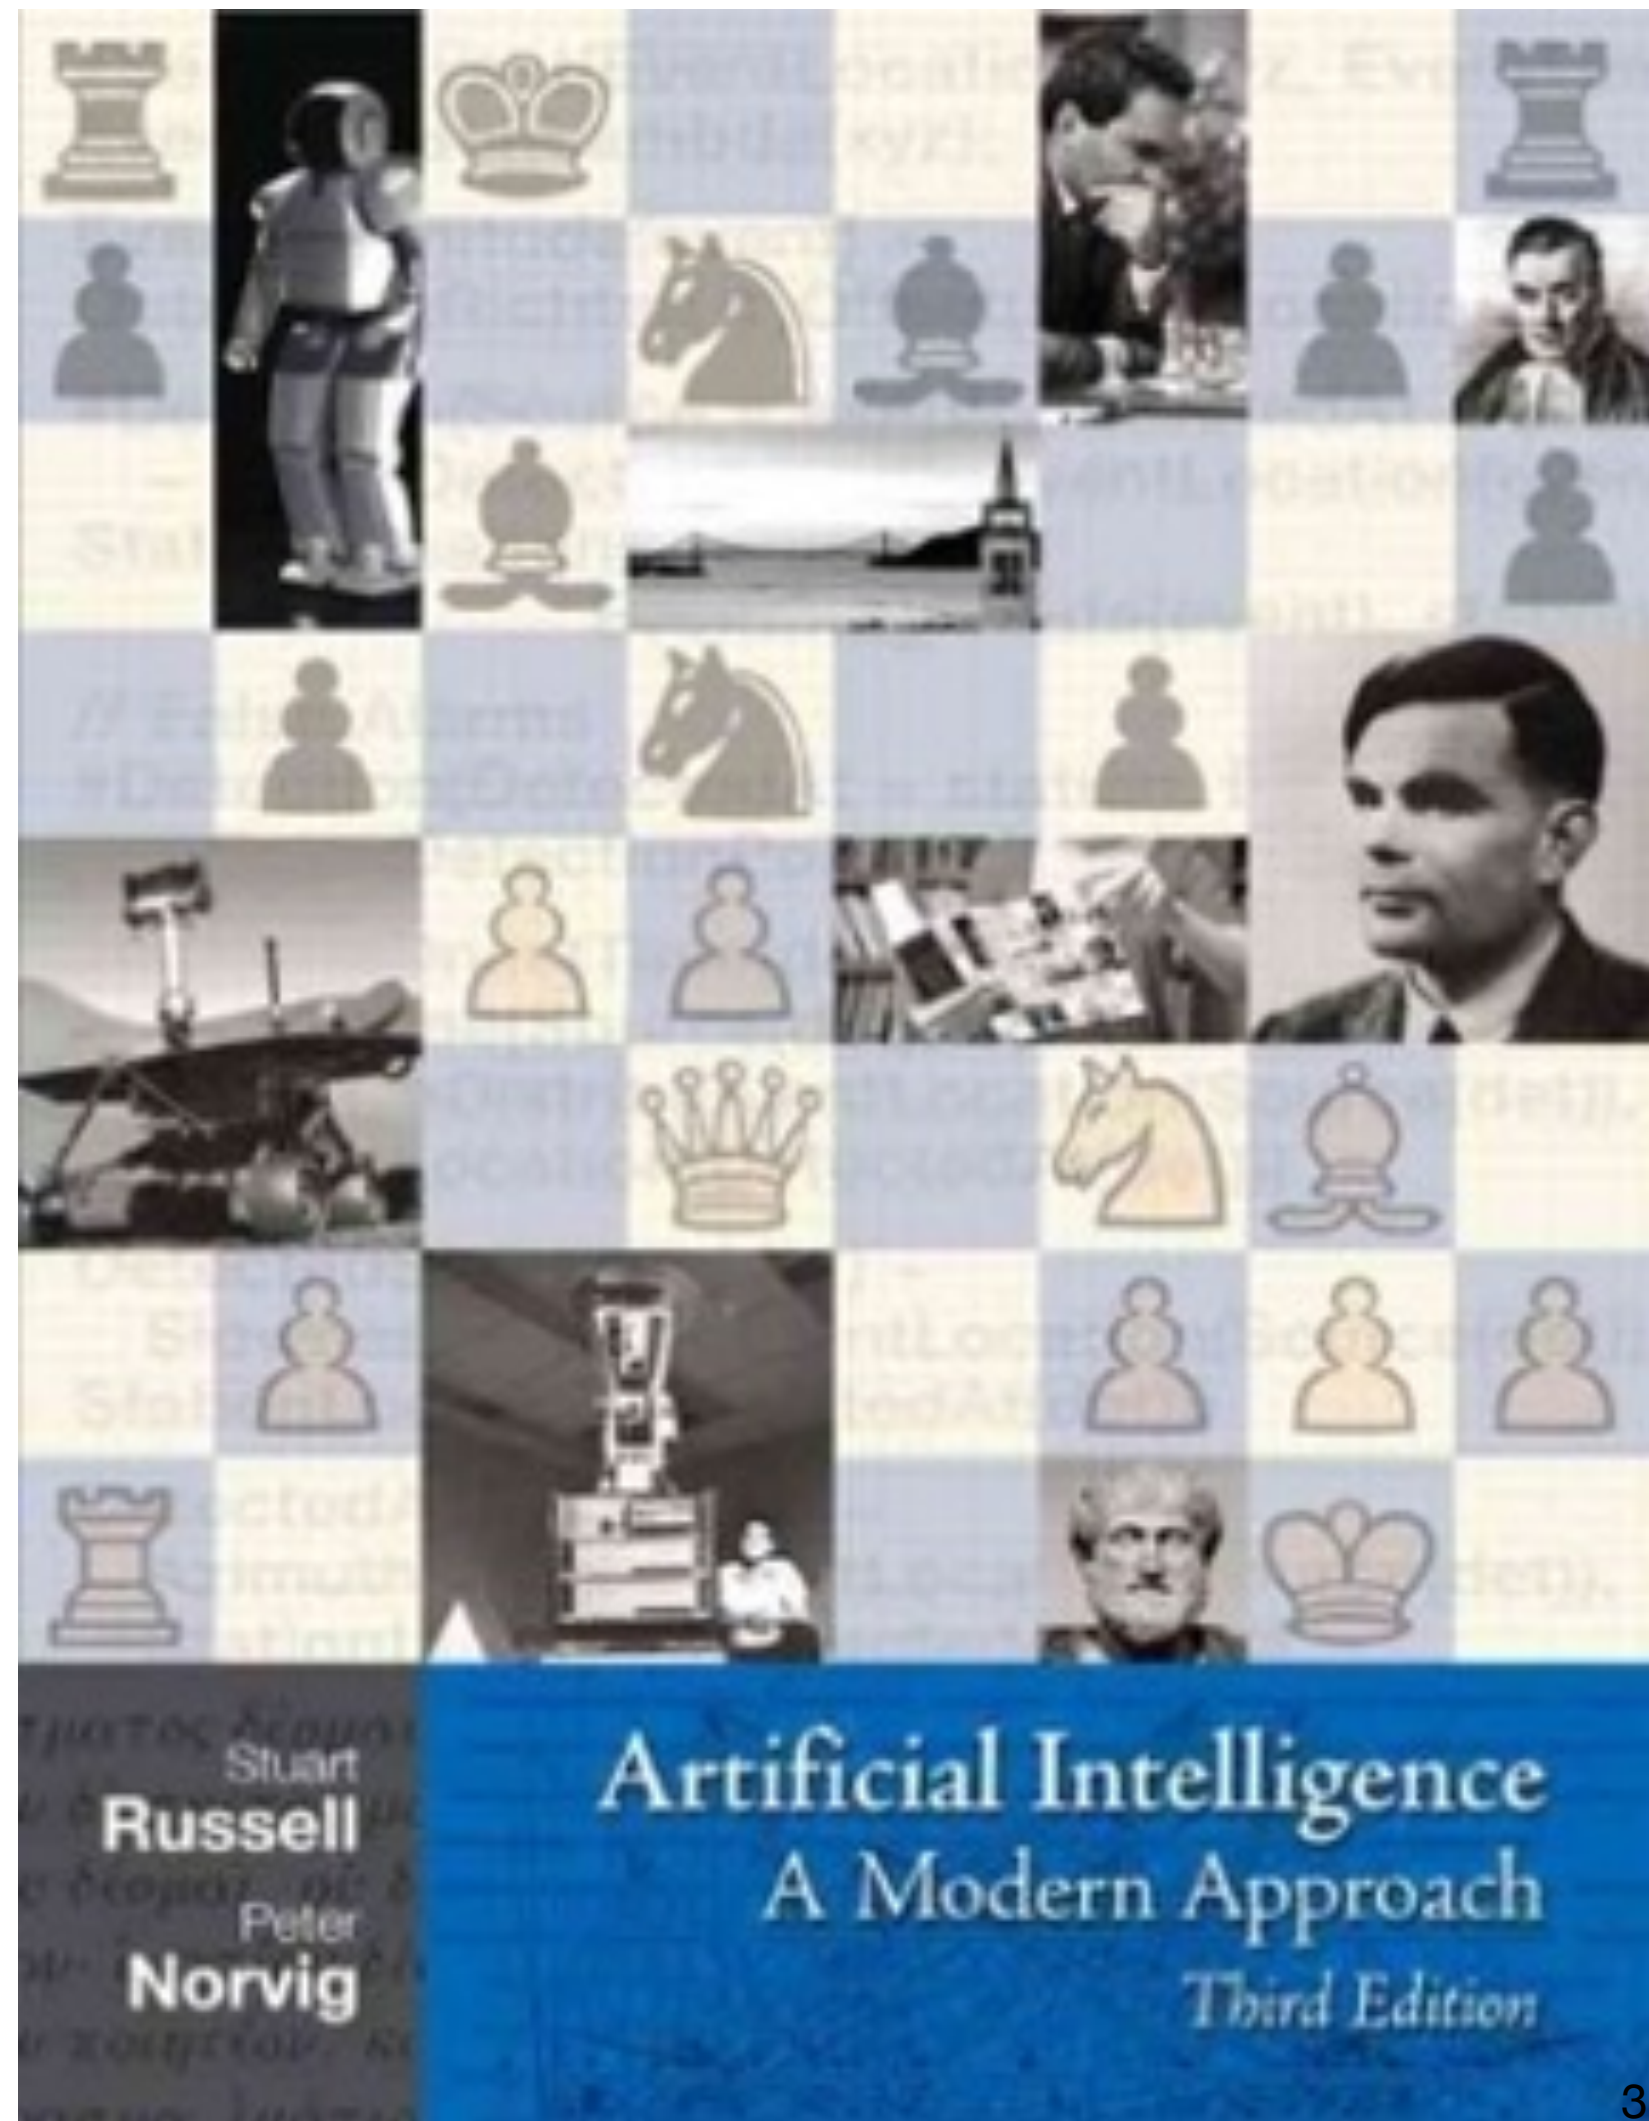
Part II: Course Logistics

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/~kandasamy/courses/22fall-cs540/index.html>
 - Slides, schedule, policies
- **Piazza** <https://piazza.com/wisc/fall2022/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% (timing TBA in near future)
- **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
- Activate the GPU supported runtime (this is a K80 GPU)

An aerial photograph of a server farm, showing a dense grid of server racks. The racks are arranged in long, straight rows, and the perspective is from directly above, looking down on the facility. The lighting is somewhat dim, highlighting the metallic surfaces and the repetitive patterns of the server units.

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

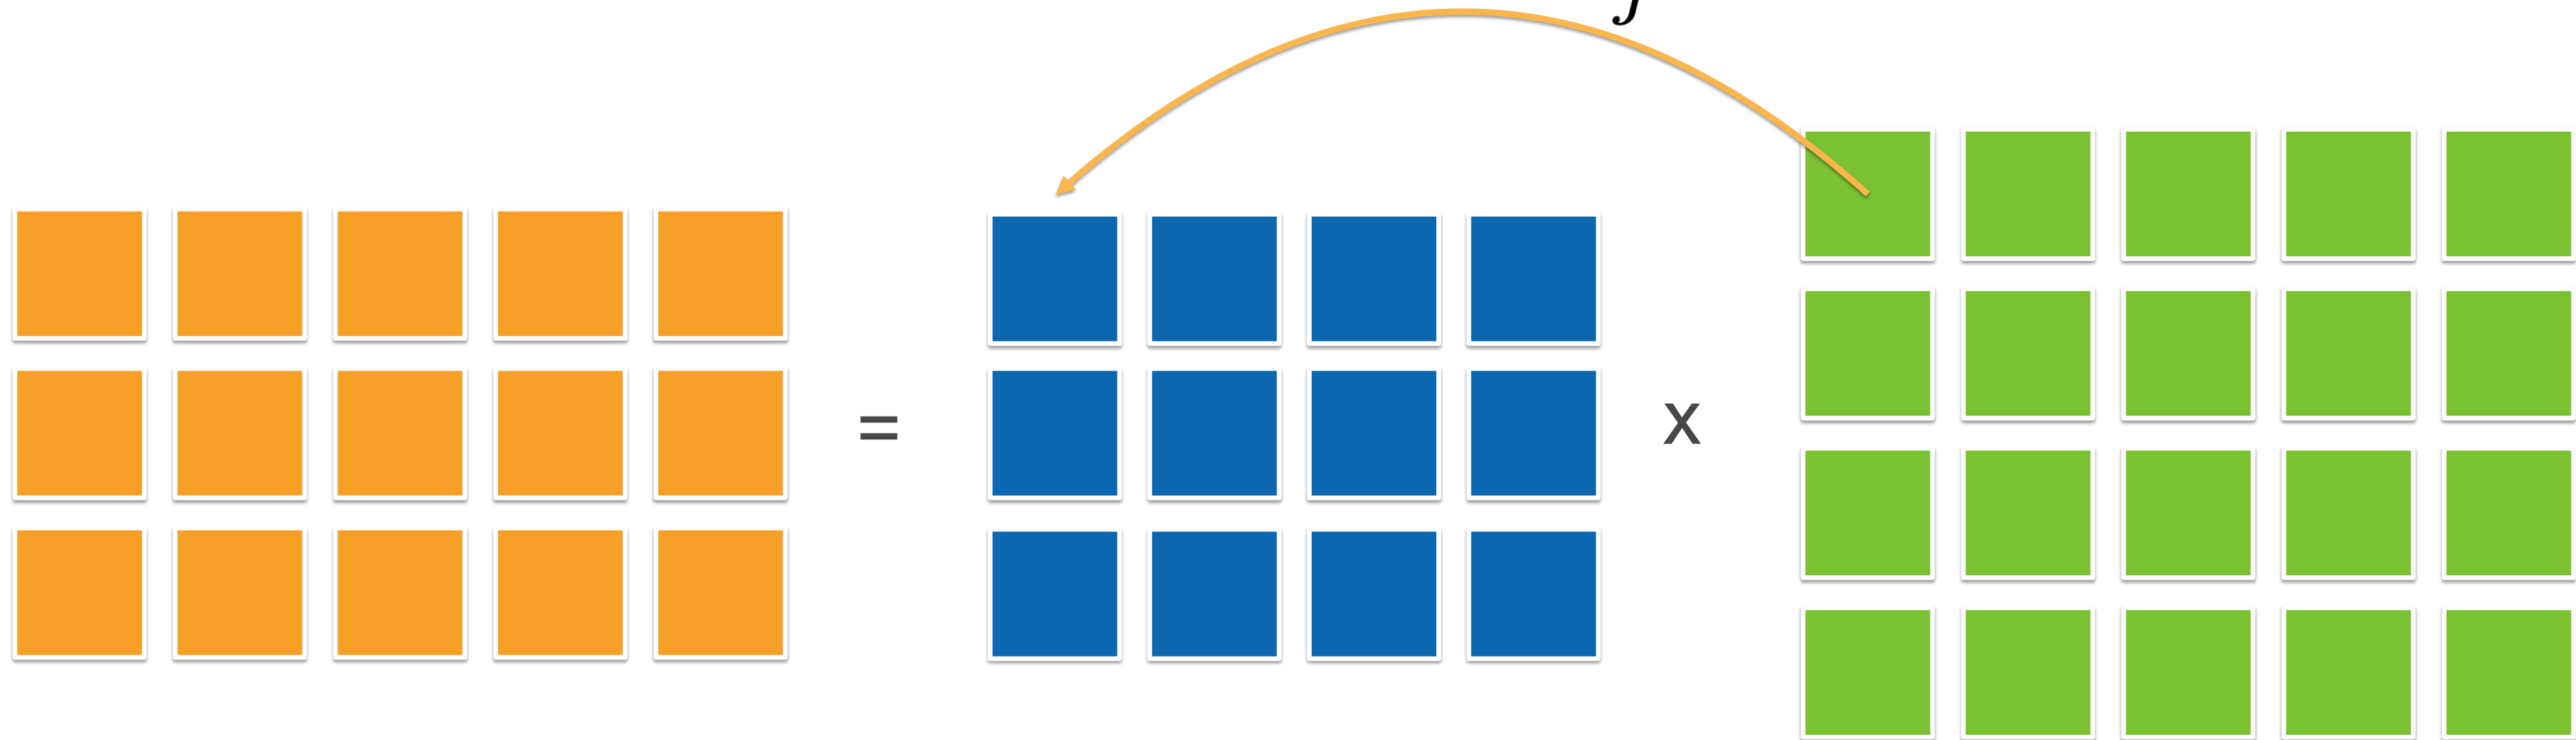
A column:[:, 2]

	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!