



CS 540 Introduction to Artificial Intelligence

Course Overview

Josiah Hanna
University of Wisconsin-Madison

09/09/2021

Slides created by Sharon Li [modified by Josiah Hanna]



Today's outline

- Introduction to “Introduction to Artificial Intelligence”
- Course logistics
- Python overview

Instructor

- Prof. Josiah Hanna

Email: jphanna@cs.wisc.edu

Office: 5391 Computer Sciences

Virtual office hours: Tuesday after class or
by appointment

Use Piazza for questions:

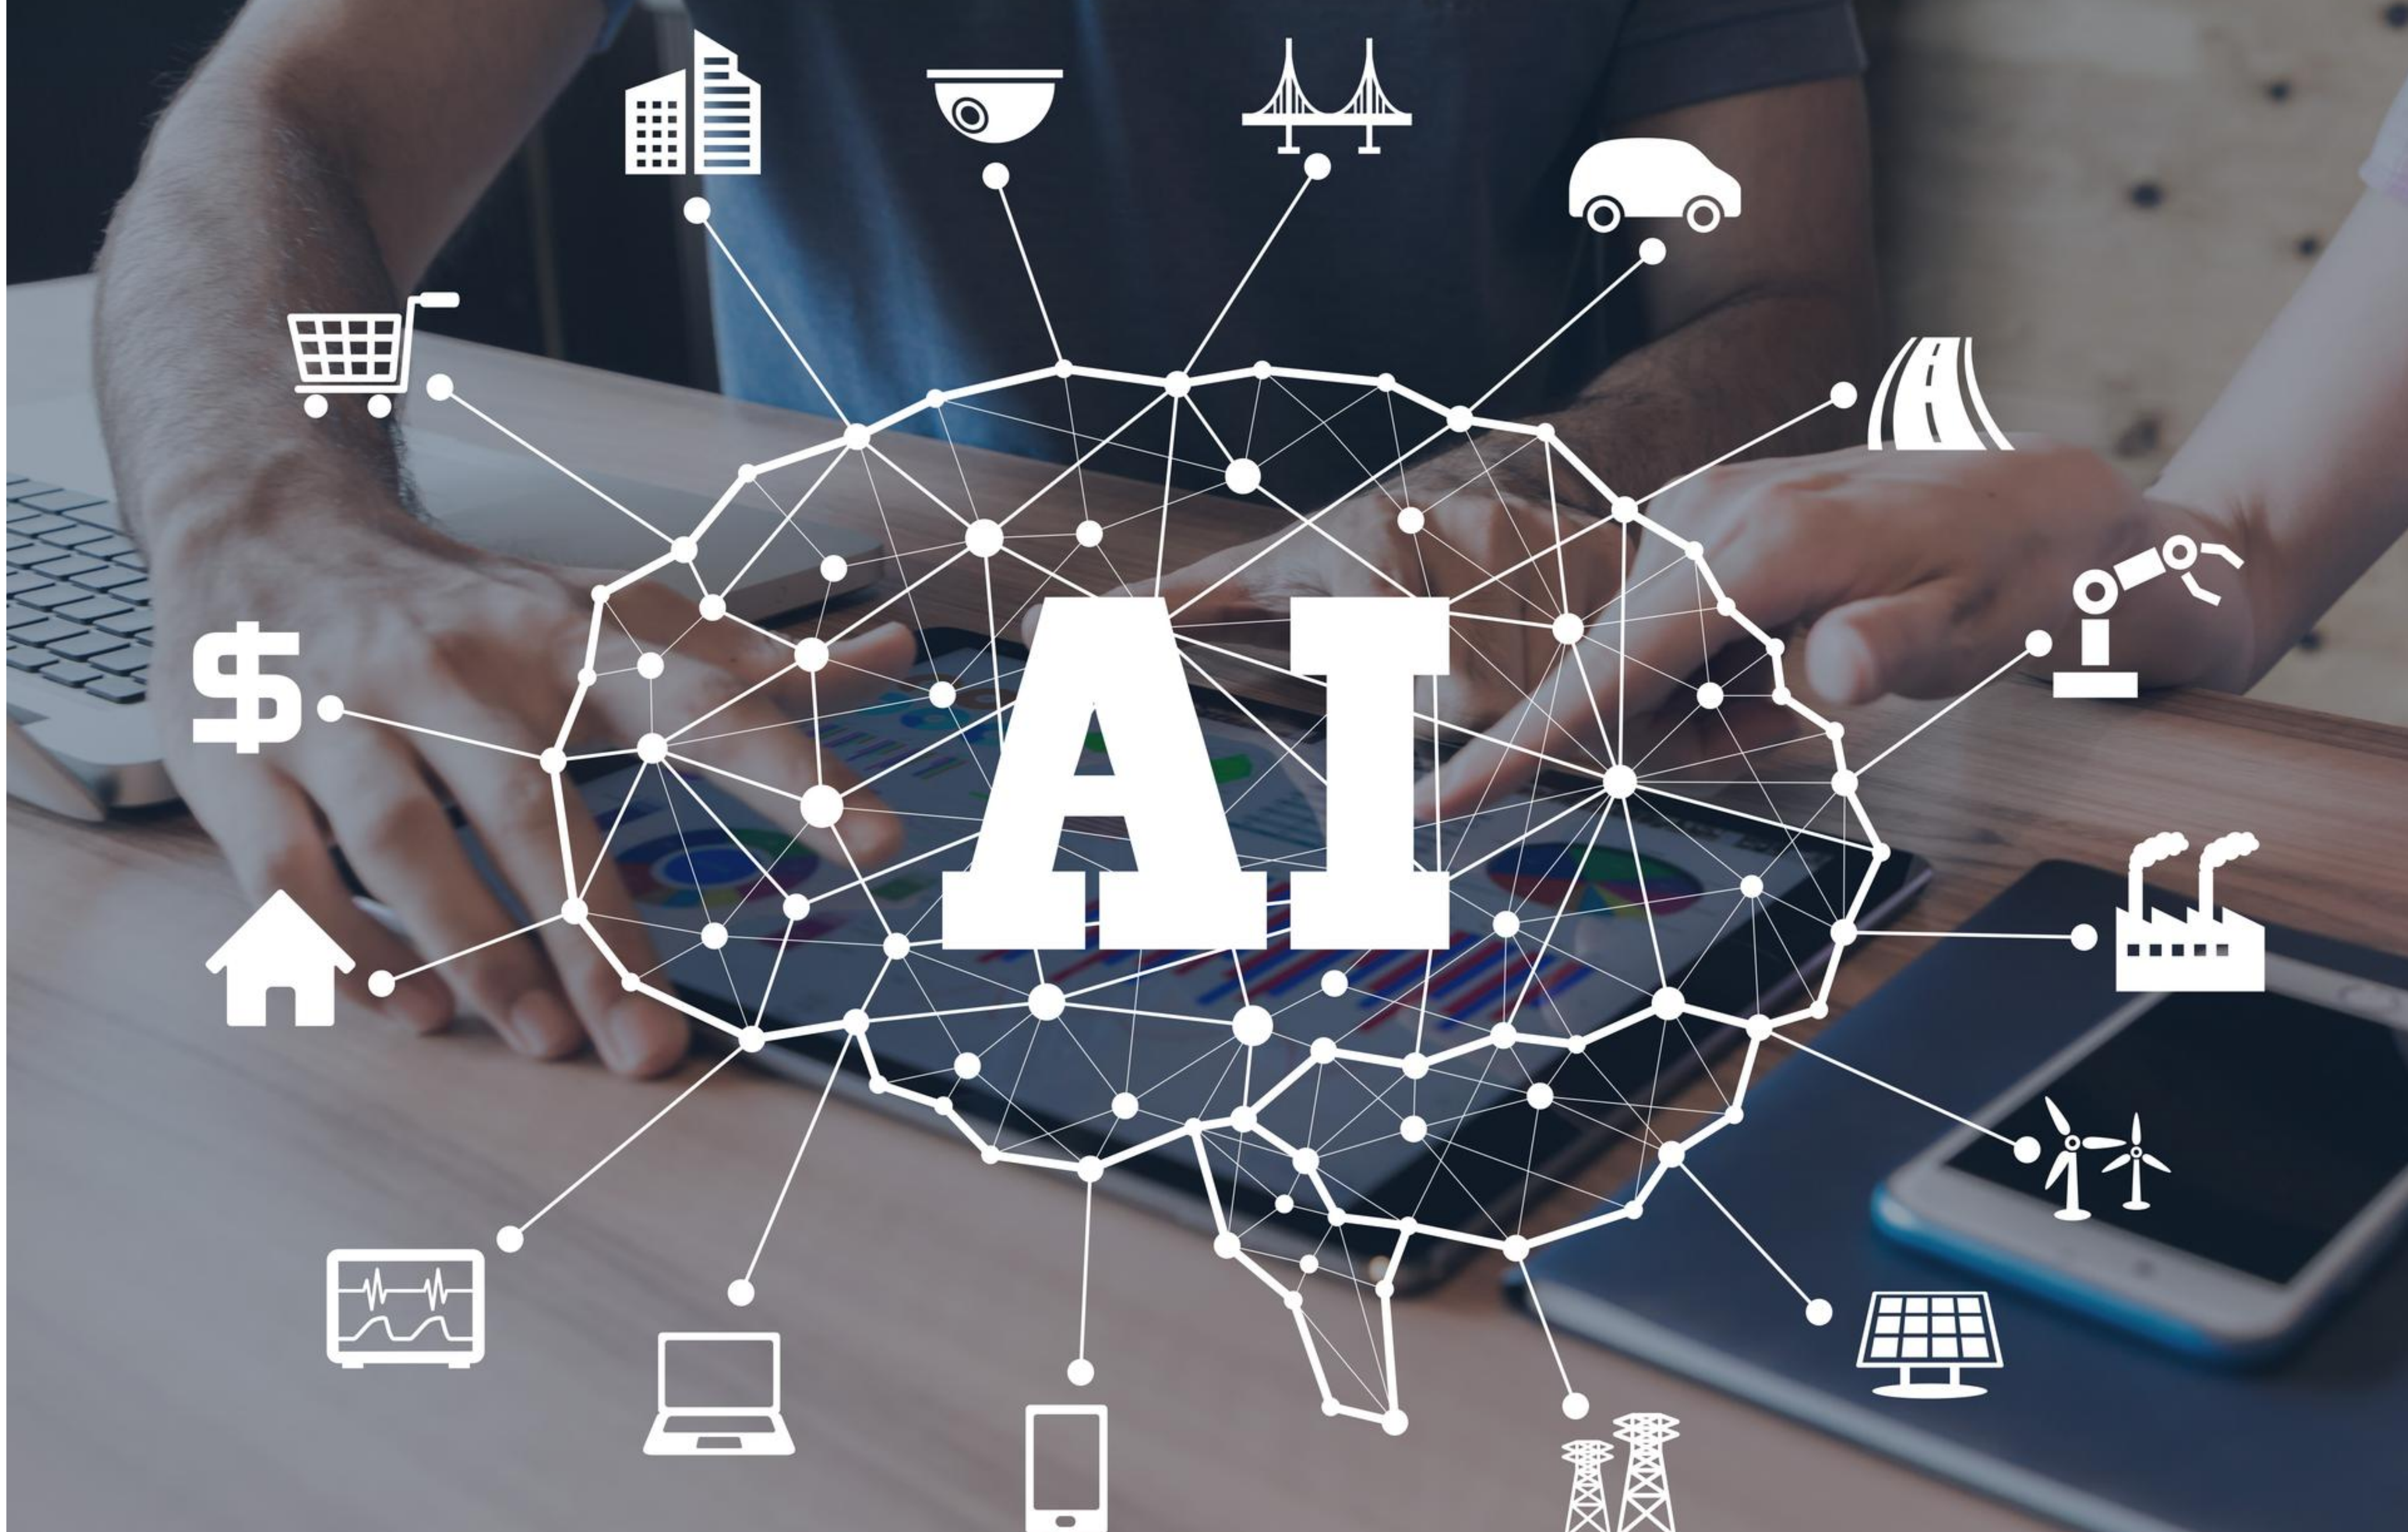
piazza.com/wisc/fall2021/cs5402/home

For emails, please include [**CS540**] in the subject title :)





Part I: Course overview



“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.”

-Proposal for the Dartmouth Workshop 1956

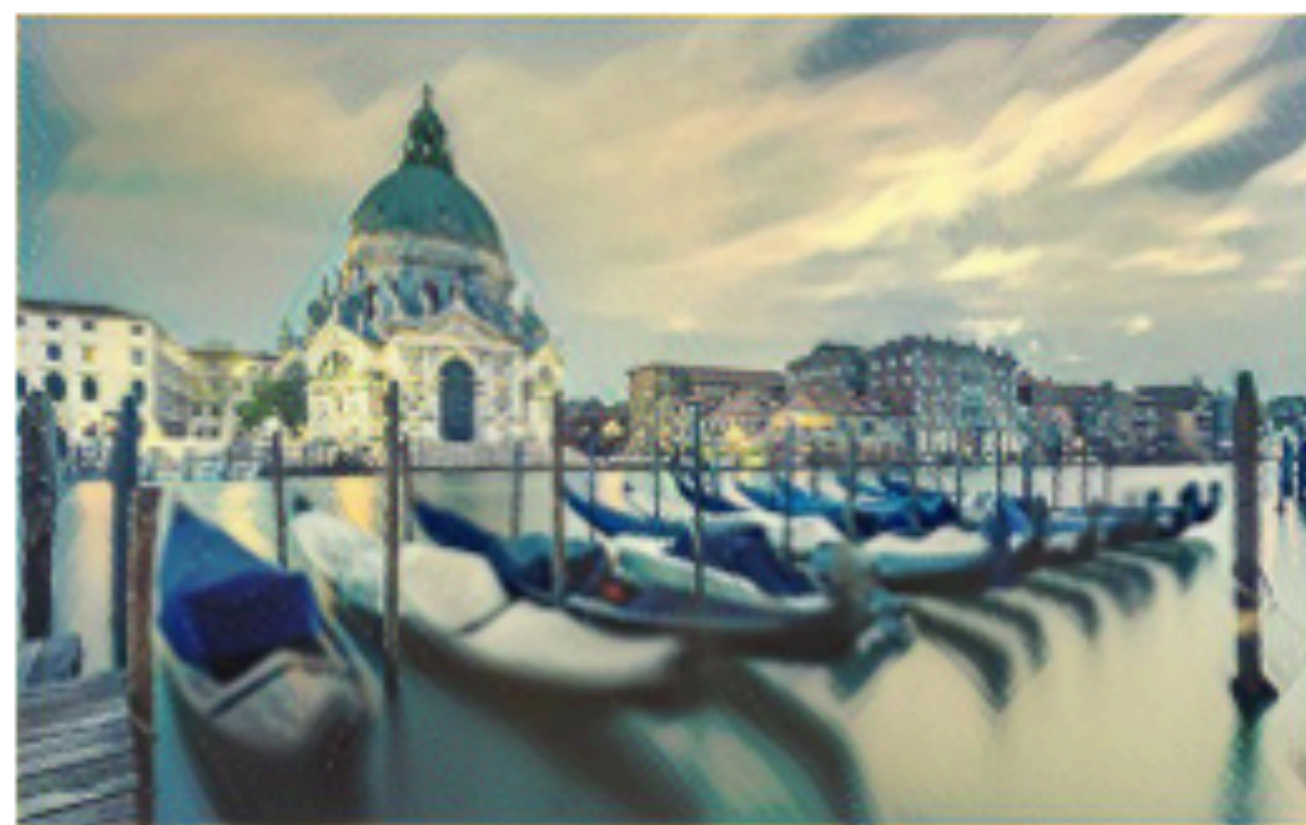
Classify Images

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



Style Transfer

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



Generative Modeling

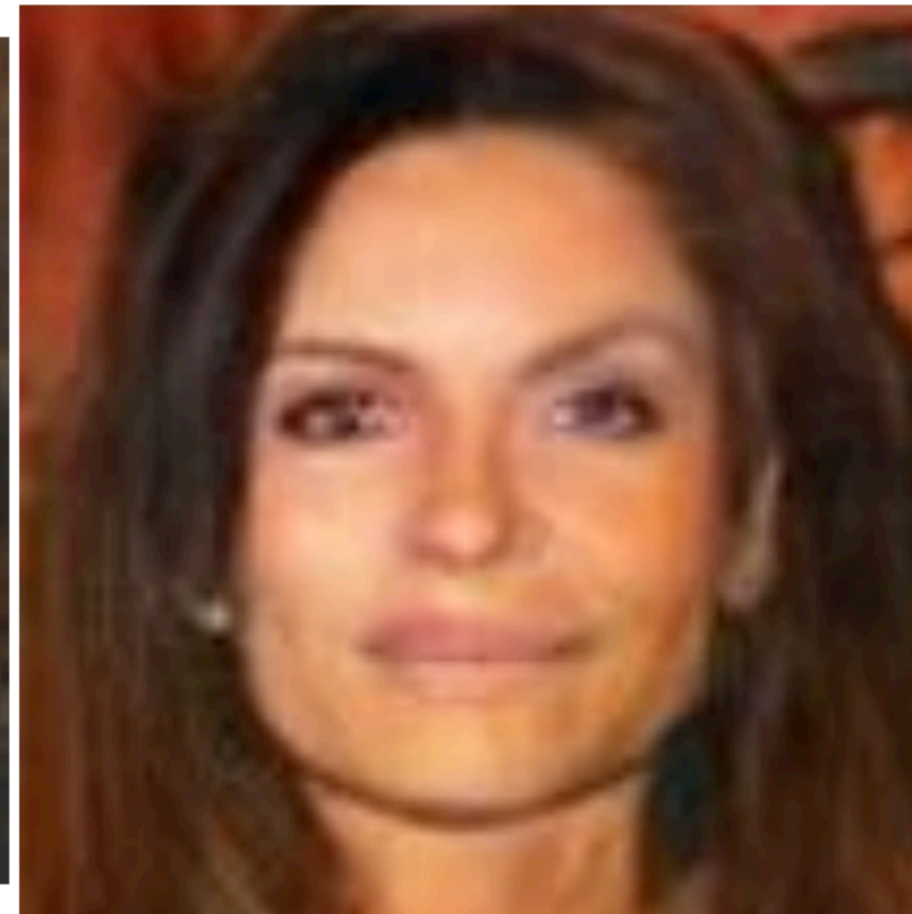
4.5 years of face generation



2014



2015



2016



2017



2018

Artificial Image Synthesis

Synthesize images



Machine Translation

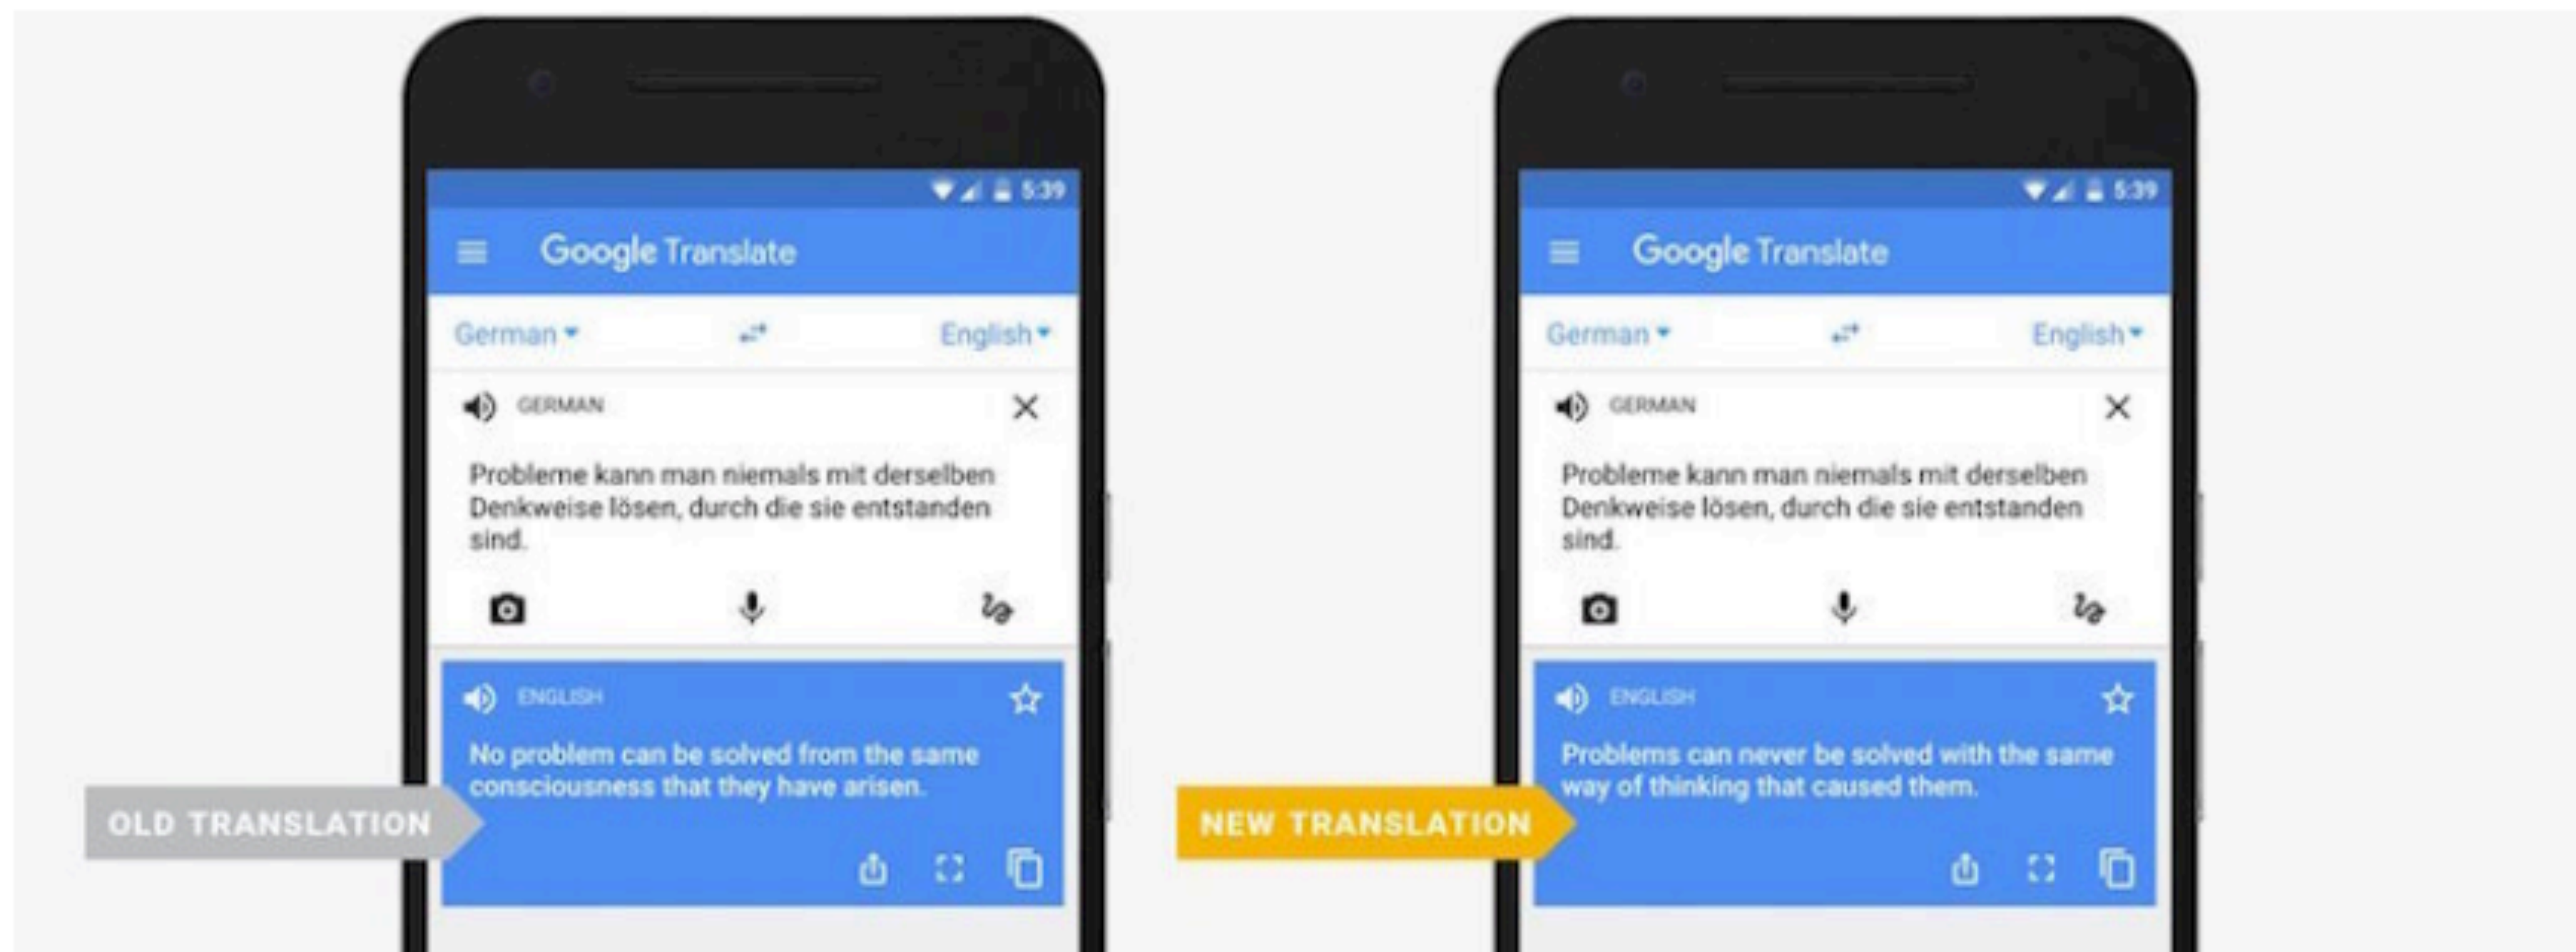
<https://www.pcmag.com/news/google-expands-neural-networks-for-language-translation>

Google Expands Neural Networks for Language Translation

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



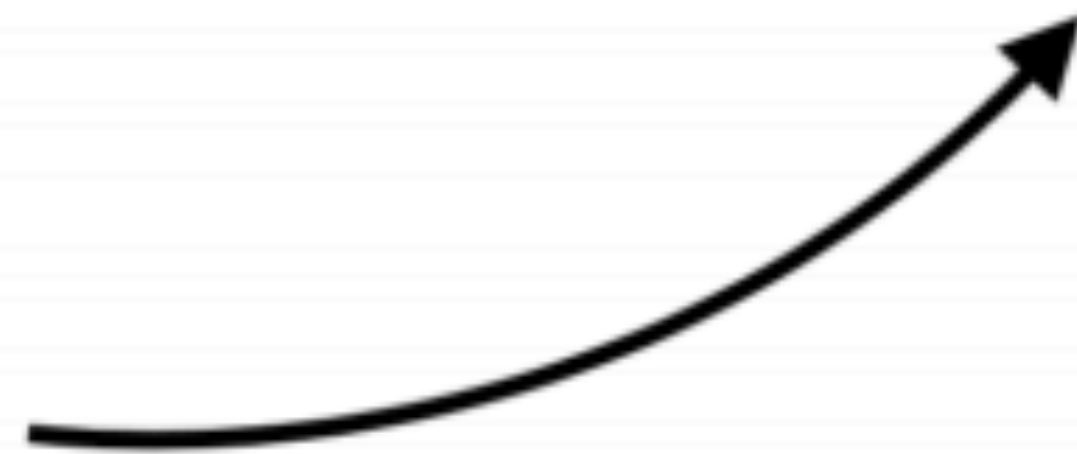
By Tom Brant November 15, 2016



Text Synthesis

Content: Two dogs play by a tree.

Style: happily, love



RNN

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

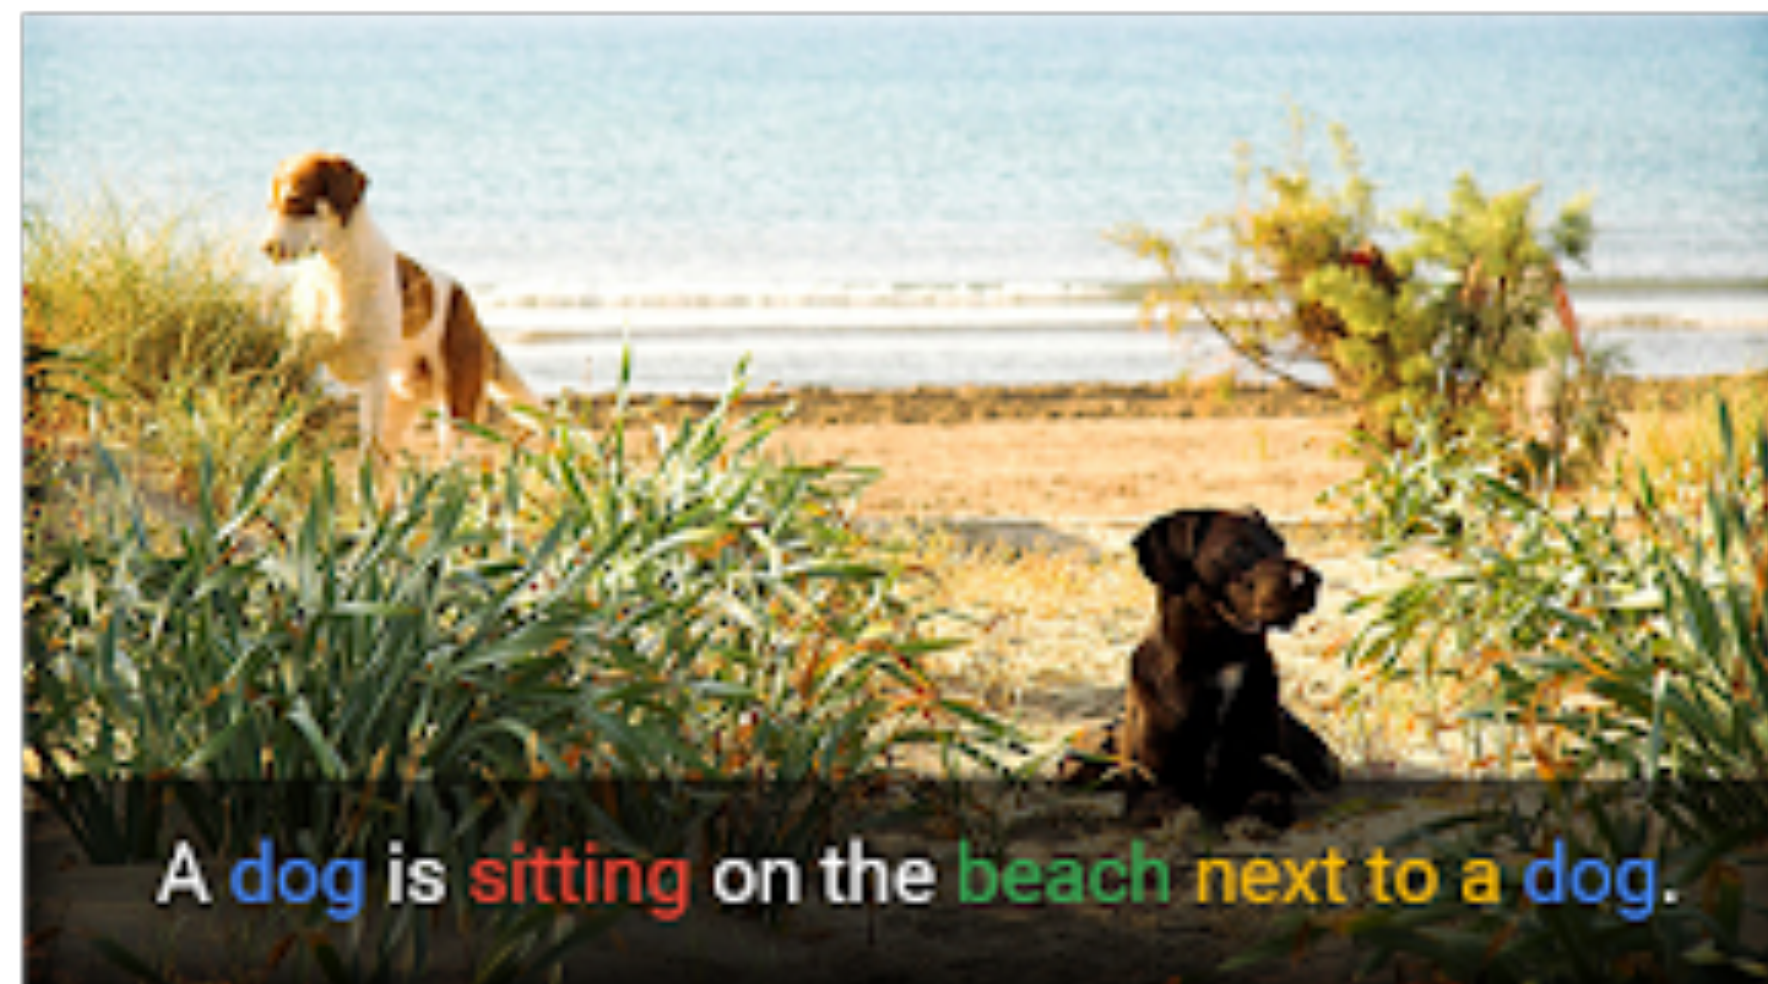
Li et al, NACCL, 2018

Image Captioning

Human captions from the training set



Automatically captioned



Defeat Human Players at Board Games



<https://qz.com/639952/googles-ai-won-the-game-go-by-defying-millennia-of-basic-human-instinct/>

Complex Robot Behavior



Boston Dynamics — This is likely more control theory than machine learning.

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.

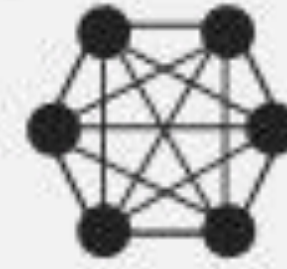
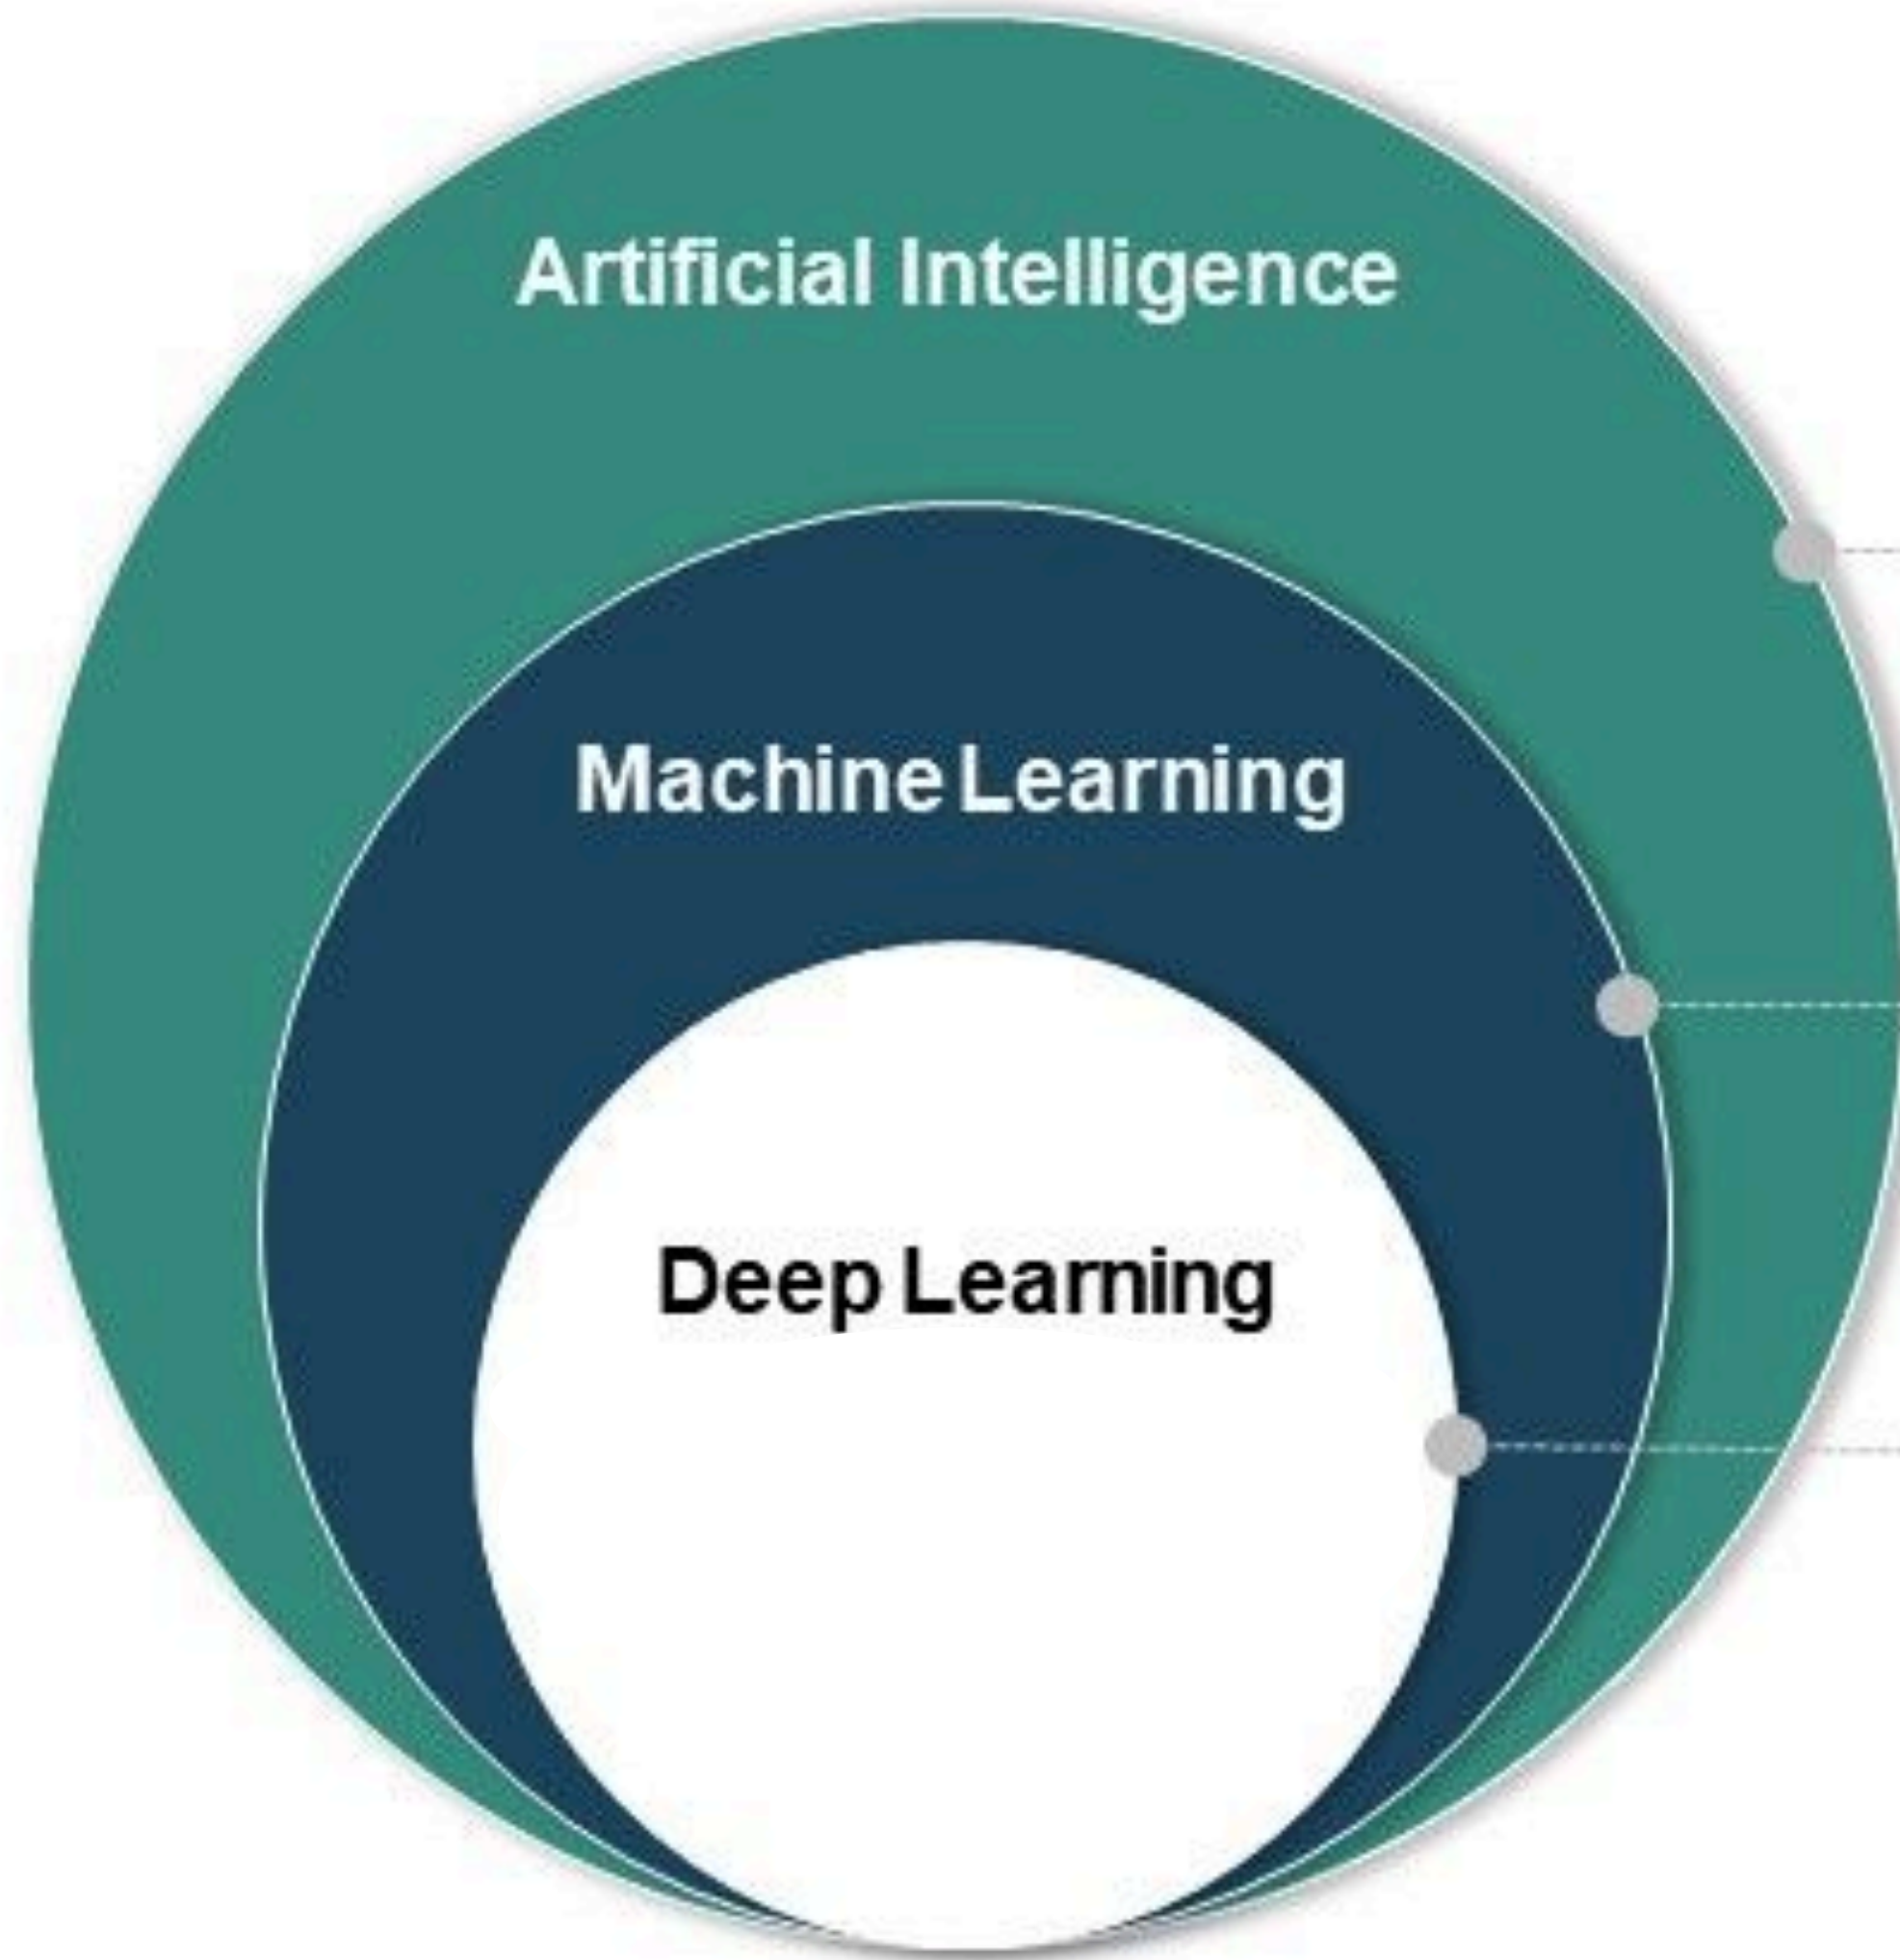
Q: Are there any animals with three legs?

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

It relies 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

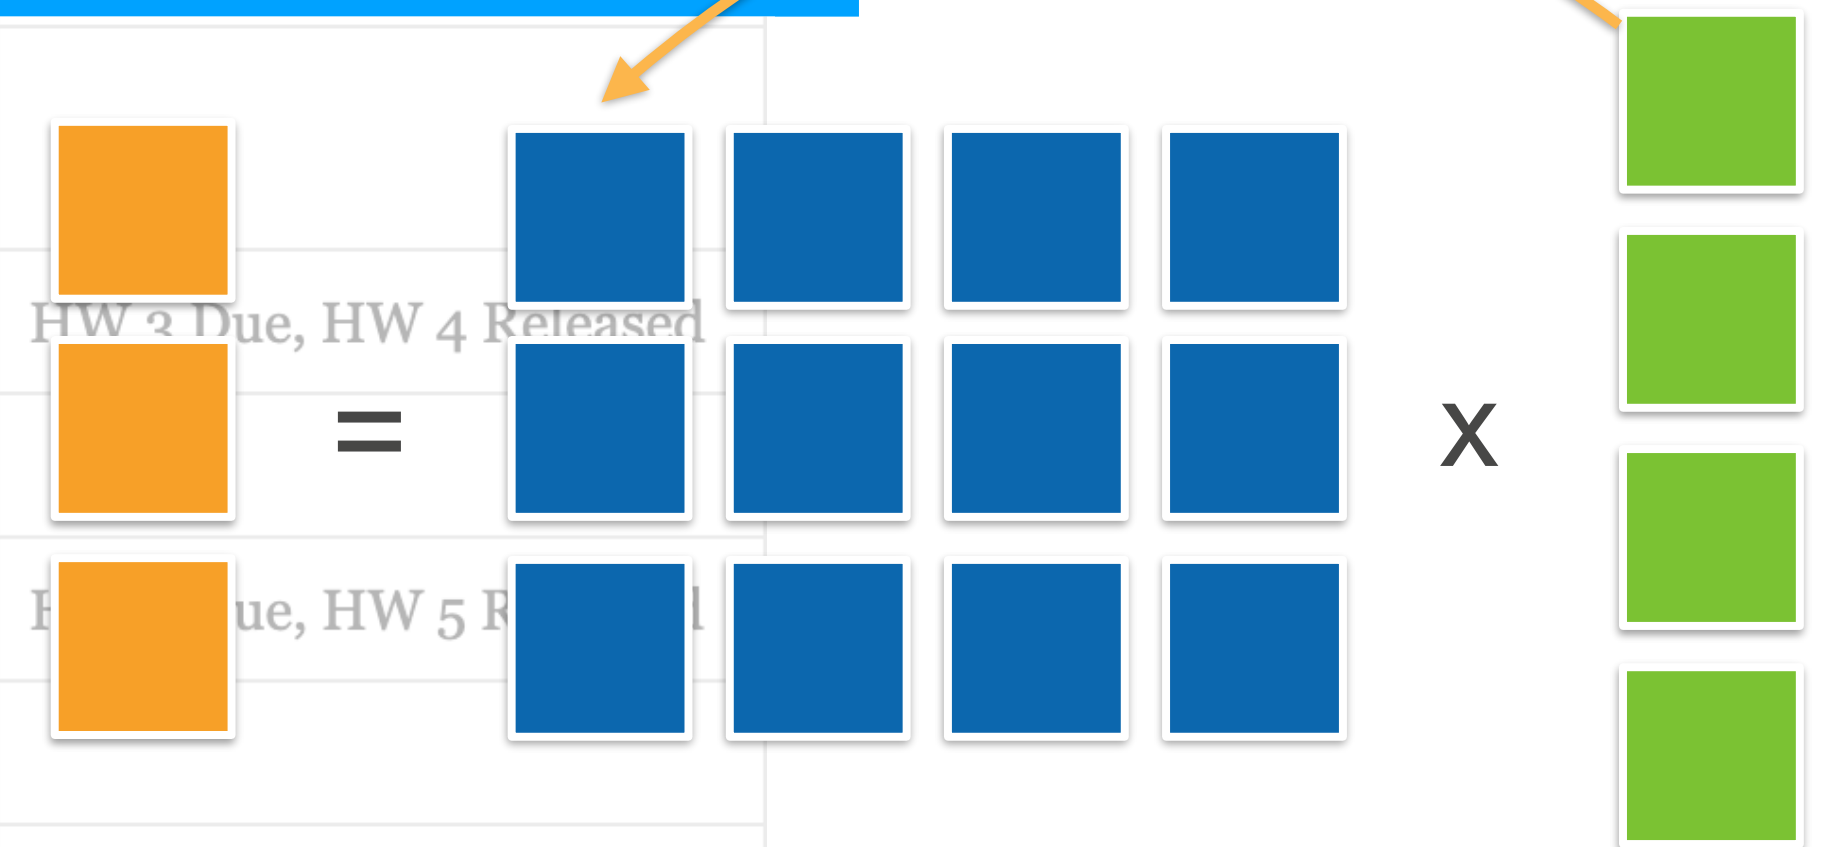
- **Basics of Machine Learning**: supervised learning vs. unsupervised learning
- **Neural Networks and Deep Learning**: Network Architecture, Training, Backpropagation, Stochastic Gradient Descent.
- Fundamentals of **Game Theory**.
- **Search and Reinforcement Learning**
- **Artificial Intelligence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

What you can learn from CS540?

https://pages.cs.wisc.edu/~jphanna/teaching/2021fall_cs540/schedule.html

Date	Topic	Reading materials	Assignments
Thursday, Sept 9	Welcome and Course Overview		
Everything below here is tentative and subject to change.			
Tuesday, Sept 14	Probability		HW 1 Released
Thursday, Sept 16	Linear Algebra and PCA		
Tuesday, Sept 21	Statistics and Math Review		HW 1 Due, HW 2 Released
Thursday, Sept 23	Introduction to Logic		
Tuesday, Sept 28	Natural Language Processing		HW 2 Due, HW 3 Released
Thursday, Sept 30	Machine Learning: Introduction		
Tuesday, Oct 5	Machine Learning: Unsupervised Learning I		HW 3 Due, HW 4 Released
Thursday, Oct 7	Machine Learning: Unsupervised Learning II		
Tuesday, Oct 12	Machine Learning: Linear Regression		HW 4 Due, HW 5 Released
Thursday, Oct 14	Machine Learning: K-Nearest Neighbors & Naive Bayes		
Tuesday, Oct 19	Machine Learning: Neural Network I (Perceptron)		HW 5 Due, HW 6 Released
Thursday, Oct 21	Machine Learning: Neural Network II		

Foundations



What you can learn from CS540?

https://pages.cs.wisc.edu/~jphanna/teaching/2021fall_cs540/schedule.html

Date	Topic	Reading materials	Assignments
Thursday, Sept 9	Welcome and Course Overview		
Everything below here is tentative and subject to change.			
Tuesday, Sept 14	Probability		HW 1 Released
Thursday, Sept 16	Linear Algebra and PCA		
Tuesday, Sept 21	Statistics and Math Review		HW 1 Due, HW 2 Released
Thursday, Sept 23	Introduction to Logic		
Tuesday, Sept 28	Natural Language Processing		HW 2 Due, HW 3 Released
Thursday, Sept 30	Machine Learning: Introduction		
Tuesday, Oct 5	Machine Learning: Unsupervised Learning I		HW 3 Due, HW 4 Released
Thursday, Oct 7	Machine Learning: Unsupervised Learning II		
Tuesday, Oct 12	Machine Learning: Linear Regression		HW 4 Due, HW 5 Released
Thursday, Oct 14	Machine Learning: K-Nearest Neighbors & Naive Bayes		
Tuesday, Oct 19	Machine Learning: Neural Network I (Perceptron)		HW 5 Due, HW 6 Released
Thursday, Oct 21	Machine Learning: Neural Network II		

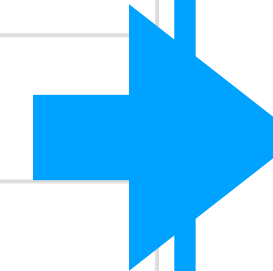


Machine learning

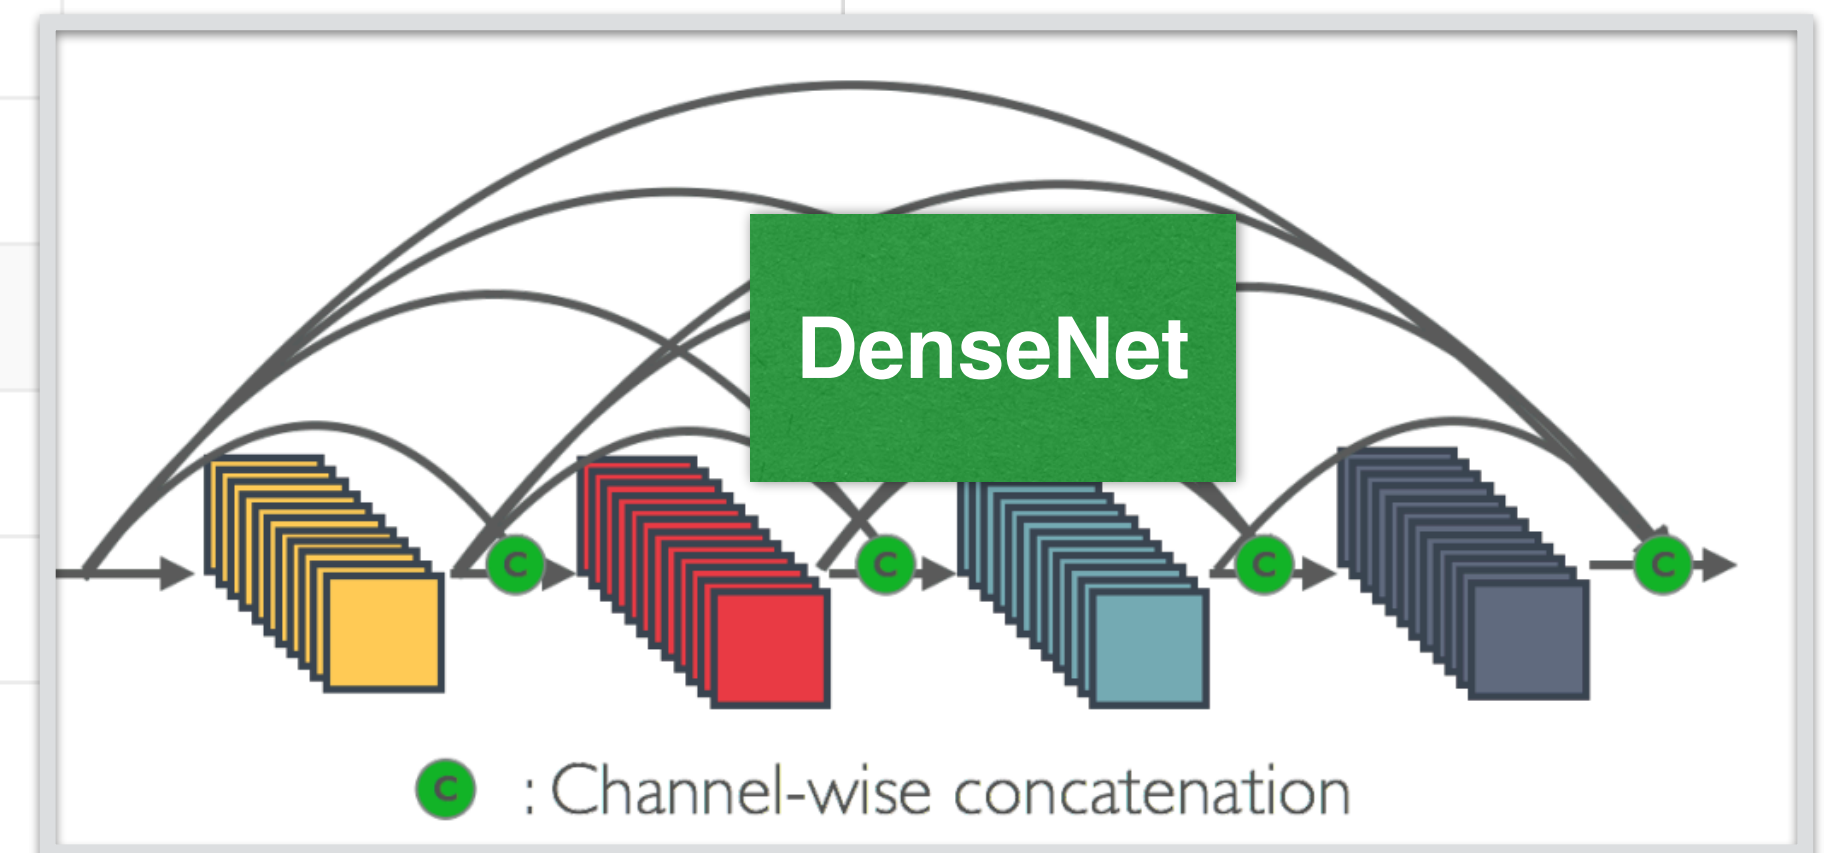
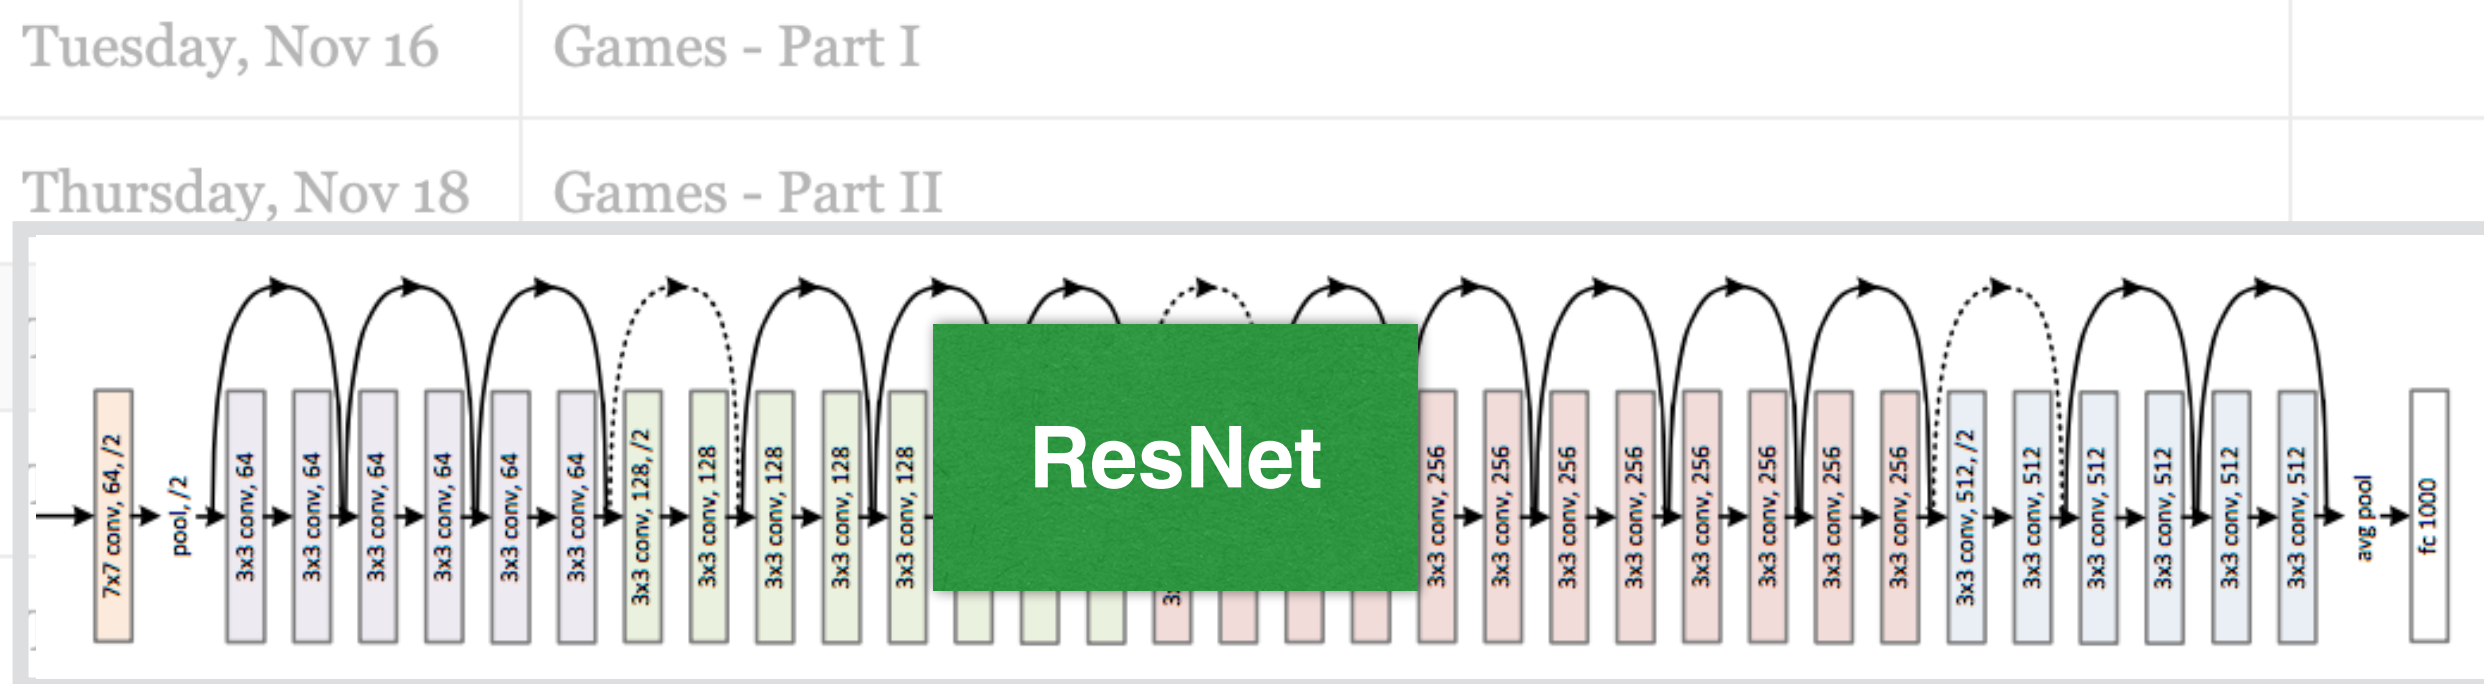
What you can learn from CS540?

https://pages.cs.wisc.edu/~jphanna/teaching/2021fall_cs540/schedule.html

Tuesday, Oct 19	Machine Learning: Neural Network I (Perceptron)	HW 5 Due, HW 6 Released
Thursday, Oct 21	Machine Learning: Neural Network II	
Tuesday, Oct 26	Machine Learning: Neural Network III	
MIDTERM EXAM October 28		
Tuesday, Nov 2	Machine Learning: Deep Learning I	
Thursday, Nov 4	Machine Learning: Deep Learning II	HW 6 Due, HW 7 Released
Tuesday, Nov 9	Machine Learning: Deep Learning III	
Thursday, Nov 11	Machine Learning: Deep Learning and Neural Network's Summary	HW 7 Due



Deep Learning



Thursday, Dec 2 Reinforcement Learning I

What you can learn from CS540?

https://pages.cs.wisc.edu/~jphanna/teaching/2021fall_cs540/schedule.html

Tuesday, Nov 2	Machine Learning: Deep Learning I		
Thursday, Nov 4	Machine Learning: Deep Learning II		HW 6 Due, HW 7 Released
Tuesday, Nov 9	Machine Learning: Deep Learning III		
Thursday, Nov 11	Machine Learning: Deep Learning and Neural Network's Summary		HW 7 Due
Tuesday, Nov 16	Games - Part I		
Thursday, Nov 18	Games - Part II		HW 8 Released
Tuesday, Nov 23	Search I: Un-Informed search		
Thursday, Nov 25	Search II: Informed search		HW 8 Due, HW 9 Released
Tuesday, Nov 30	Advanced Search		
Thursday, Dec 2	Reinforcement Learning I		HW 9 Due, HW 10 Released
Tuesday, Dec 7	Reinforcement Learning II		
Thursday, Dec 9	Reinforcement Learning and Search Summary		
Tuesday, Dec 14	Ethics and Trust in AI		HW10 Due



**Game, search
and RL**

What you can learn from CS540?

https://pages.cs.wisc.edu/~jphanna/teaching/2021fall_cs540/schedule.html

Tuesday, Nov 2	Machine Learning: Deep Learning I		
Thursday, Nov 4	Machine Learning: Deep Learning II		HW 6 Due, HW 7 Released
Tuesday, Nov 9	Machine Learning: Deep Learning III		
Thursday, Nov 11	Machine Learning: Deep Learning and Neural Network's Summary		HW 7 Due
Tuesday, Nov 16	Games - Part I		
Thursday, Nov 18	Games - Part II		HW 8 Released
Tuesday, Nov 23	Search I: Un-Informed search		
Thursday, Nov 25	Search II: Informed search		HW 8 Due, HW 9 Released
Tuesday, Nov 30	Advanced Search		
Thursday, Dec 2	Reinforcement Learning I		HW 9 Due, HW 10 Released
Tuesday, Dec 7	Reinforcement Learning II		
Thursday, Dec 9	Reinforcement Learning and Search Summary		
Tuesday, Dec 14	Ethics and Trust in AI		HW10 Due



**Ethics and
Real World AI!**



Part II: Course Logistics

Course Schedule

- Time: **Tuesday** and **Thursday** 11:00am-12:15pm CT
- Location: **SMI 331**
- Schedule is available on the course website:

https://pages.cs.wisc.edu/~jphanna/teaching/2021fall_cs540/schedule.html

- Slides online on website

Lecture

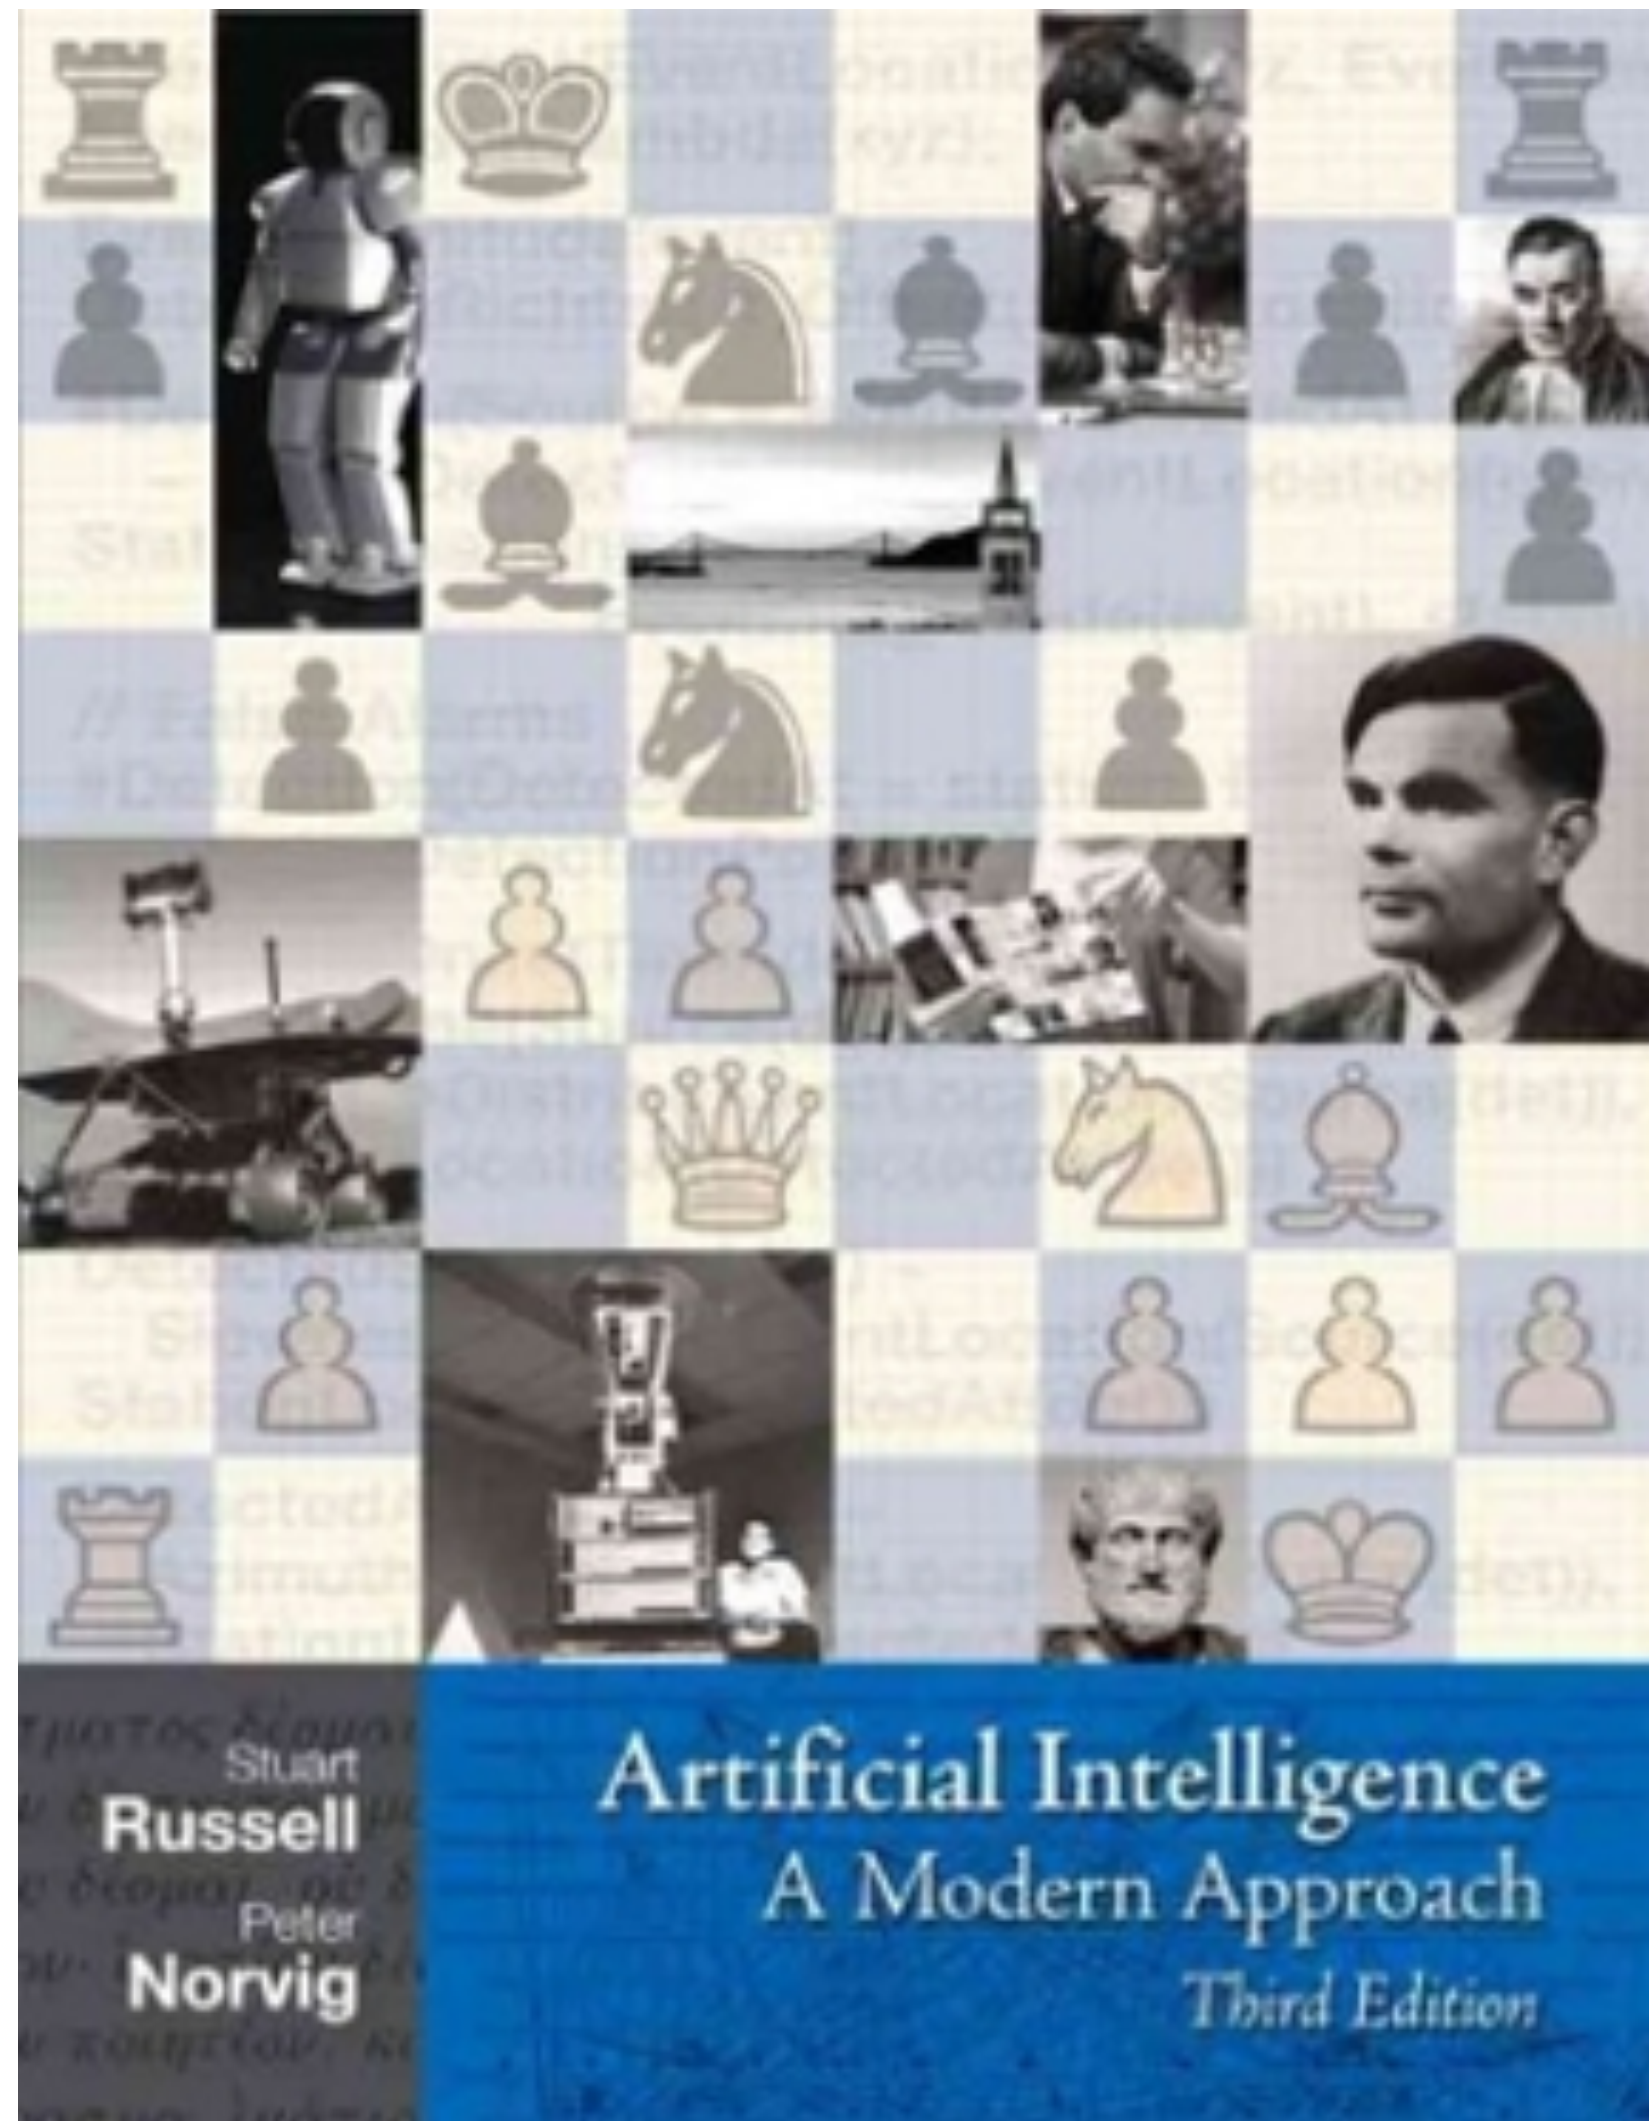
- Each lecture will be a series of **3** short mini-lectures.
- In each block, I will:
 - Lecture using slides or a whiteboard.
 - Pause for interactive Q&A.
 - Deliver short quiz questions to clear up any confusion before proceeding to the next block.
- **If you're experiencing any symptoms of Covid or have been exposed to someone who tested positive for Covid, please do not attend lecture.**
- No penalty for not attending lecture.

Where to find content?

- **Piazza** piazza.com/wisc/fall2021/cs5402/home
 - Discussion, questions
 - Announcements
- **Canvas** - private materials *that should not be shared*
 - Videos
 - Assignments
 - Grades
- **Course website** - public materials
https://pages.cs.wisc.edu/~jphanna/teaching/2021fall_cs540/schedule.html
 - Slides
 - Schedule
 - Policies

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)**



Teaching Assistants

Shuyao Li

- shuyao.li@wisc.edu
- Office hours: TBD

Hugh Liu

- ylin785@wisc.edu
- Office hours: TBD

Jeremy McMahan

- jmcmahan@wisc.edu
- Office hours: TBD

Peer Mentors

- See course website
- Shared across course sections

Office Hours

- Available on the course website: https://pages.cs.wisc.edu/~jphanna/teaching/2021fall_cs540/schedule.html
- Instructor office hours: default by zoom; In-person office hours available upon request.
- TA / Peer Mentor office hours: up to individual.
- Use Peer Mentor hours for detailed-level questions (e.g. coding related), and use TA office hours for conceptual level questions.

Grading scheme

- **Midterm Exam:** 15% (October 28)
- **Final Exam:** 15% (December 23)
- **Homework Assignments:** 70% (10 HWs)

TWO lowest homework scores are dropped from the final homework average calculation.

Homework is always due the minute before class starts on the due date.

(Late submissions will not be accepted.)

Homework will be posted and submitted via Canvas.

Regrade Request

Use Google Form for regrade request

Raised with the TAs within 72 hours after it is returned.

Entire assignment will be regraded.

Integrity

https://pages.cs.wisc.edu/~jphanna/teaching/2021fall_cs540/about.html

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](#))

Diversity

https://pages.cs.wisc.edu/~jphanna/teaching/2021fall_cs540/about.html

Respect for Diversity: It is my intent that students from all diverse backgrounds and perspectives be well served by this course, that students' learning needs be addressed both in and out of class, and that the diversity that students bring to this class be viewed as a resource, strength and benefit. It is my intent to present materials and activities that are respectful of diversity: gender, sexuality, disability, age, socioeconomic status, ethnicity, race, and culture. Your suggestions are encouraged and appreciated. Please let me know ways to improve the effectiveness of the course for you personally or for other students or student groups. In addition, if any of our class meetings conflict with your religious events, please let me know so that we can make arrangements for you.

Please, commit to helping create a climate where we treat everyone with dignity and respect. Listening to different viewpoints and approaches enriches our experience, and it is up to us to be sure others feel safe to contribute. Creating an environment where we are all comfortable learning is everyone's job: offer support and seek help from others if you need it, not only in class but also outside class while working with classmates.



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 view of a server rack, showing a grid of server units with gold-colored connectors. 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: [1, :]

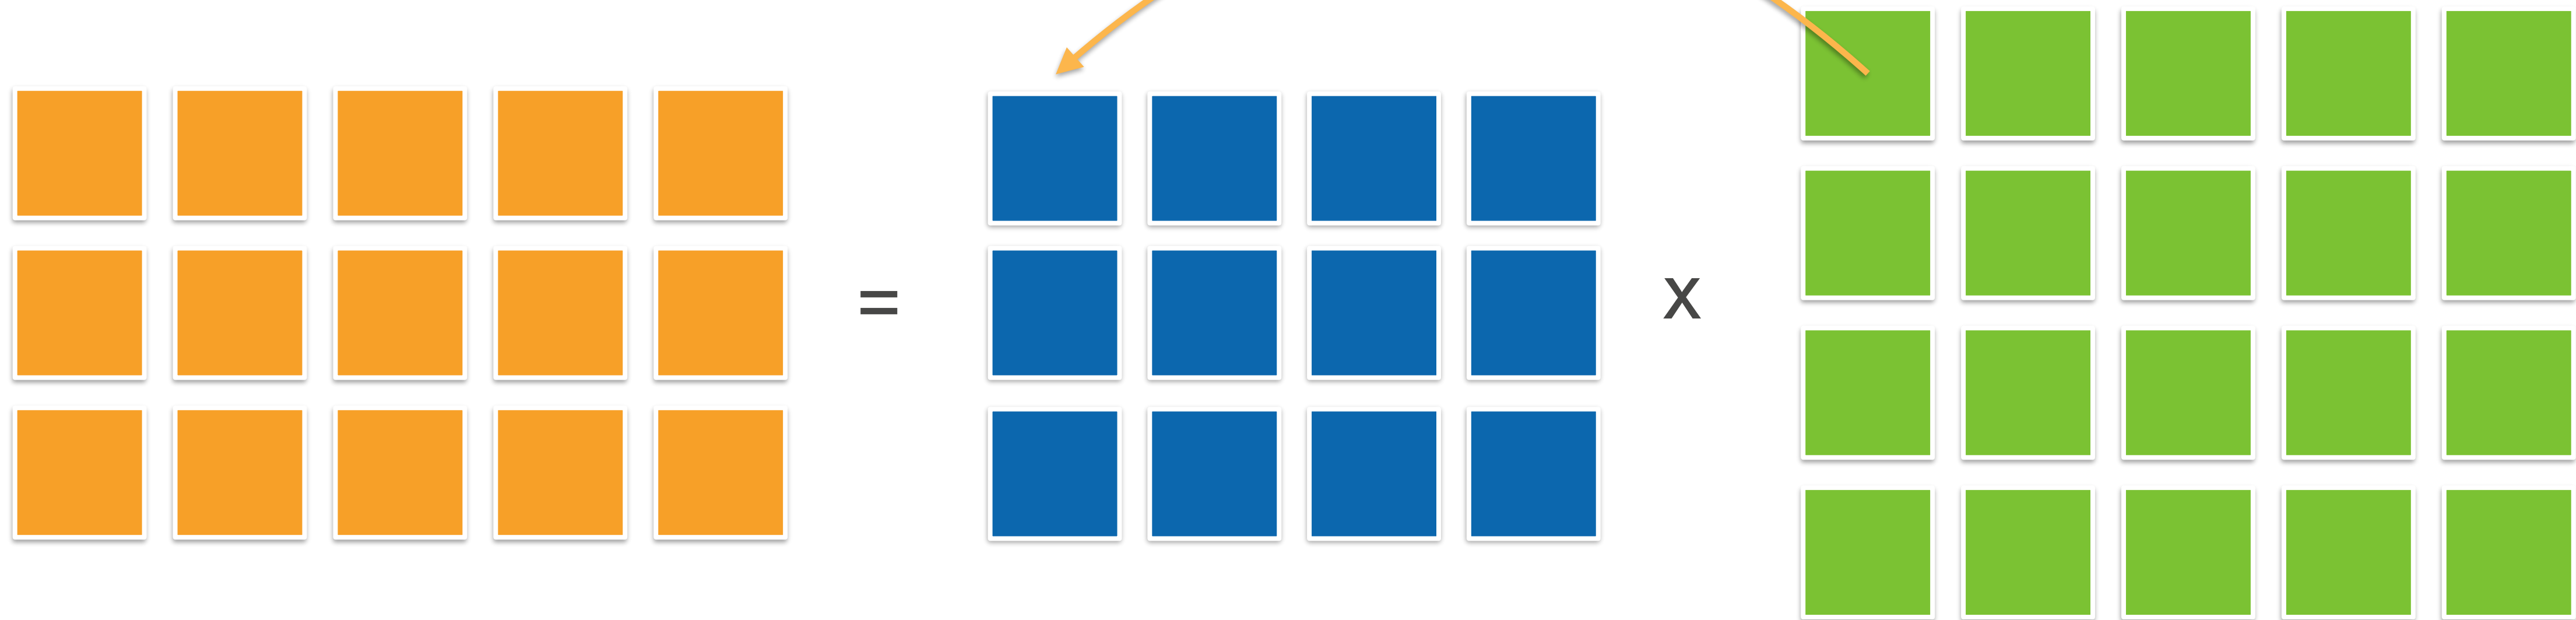
	0	1	2	3
0	1	2	3	4
1	5	6	7	8
2	9	10	11	12
3	13	14	15	16

A slice: [1:3, 1:3]

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

Next week: Linear Algebra and PCA

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





Recap

- Artificial Intelligence
 - Overview and agenda
- Course logistics
 - Homework, exam, OH etc
- Software



Thanks!