CSCI 1470/2470
Spring 2022

Welcome ☺

Deep Learning

Ritambhara Singh

January 24, 2024

Wednesday

ChatGPT prompt “minimalist landscape painting of a deep underwater scene with a blue tang fish in the bottom right corner”

About your instructor!

5th year CS and Data Science faculty and a Center for Computational Molecular Biology (CCMB) member!

Research Interests

Machine Learning, Data Mining, Computational Biology, Health Sciences

B.E
208-2012

Ph.D.
2012-2018

Postdoc
2018-2019

Office Location

Room 313, Data Science Institute (DSI)
3rd Floor, 164 Angell Street

Office Hours (Starting this week!)

Thursdays, 2:00-4:00 PM or by appointment

Where: Room 375, 3rd Floor, 164 Angell Street ← Not CIT!

Email: cs_deeplearning@brown.edu ← Please email here!

Website: www.ritambharasingh.com

I also teach:
CSCI 2952-G: Deep Learning in Genomics

Introducing...

Your Awesome Course Staff!

Your Graduate TA ☺ Your HTAs!

Michal Golovanevsky
she/her

Raymond Dai
he/him

Erica Song
she/her

Joe Dodson
he/him

Karan Kashyap
he/him

Pranav Mahableshwarkar
he/him

Earth Mokkamakkul
he/him

Your TAs!

Julian Dai
he/him

Calvin Eng
he/him

Taj Gillin
he/him

Spandan Goel
he/him

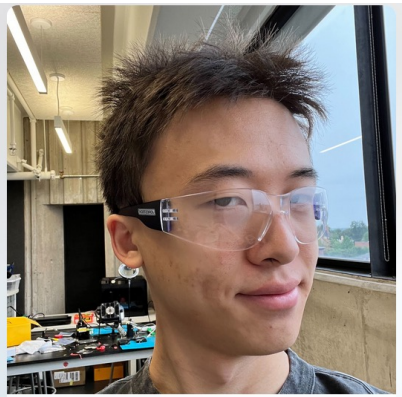

Naicheng (Arnie) He
he/him

Amanda Hernandez Sandate
she/her

Woody Hulse
he/him

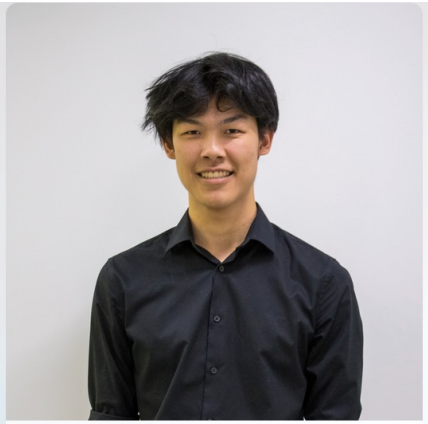

Kelvin Jiang
he/him

Bumjin Joo
he/him

Sophia Qiming
she/her

Aayush Setty
he/him

Jason Silva
he/him

Your TAs!

Kyle Lam
he/him

Jennifer Li
she/her

Alyssa Loo
she/her

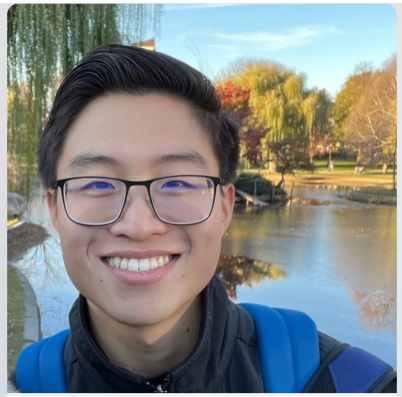

Michael Lu
he/him

Ben Maizes
he/him

Ken Ngamprasertsith
he/him

Preetish Juneja
he/him

Mohammed Khan
he/him

Philip LaDuca
he/him

Aryan Singh
he/him

Sameer Sinha
he/him

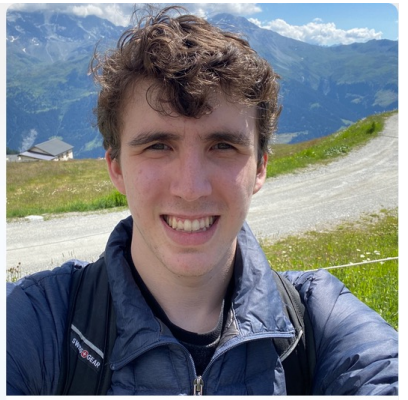

Quinn Straus
he/him

Your TAs!

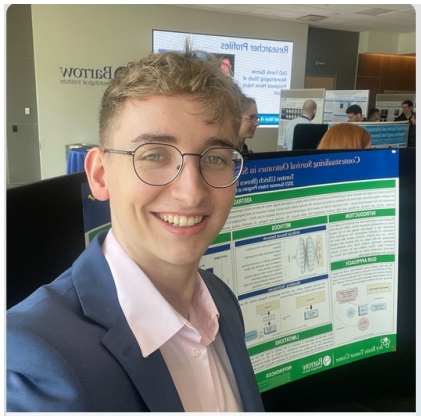

Torsten Ullrich
he/him

Mikayla Walsh
he/him

Emily Wang
she/her

Xilin (Rice) Wang
he/him

Ray Xu
he/him

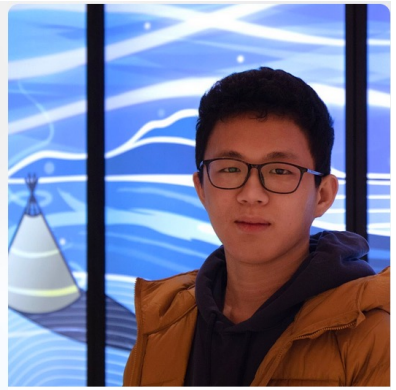

Enyan Zhang
he/him

Alex Zheng
he/him

Alex Zhou
he/him

Your STAs!

Naphat Permpredanun
he/him • STA

Kyle Yeh
he/him • UTA-STA

Lingze Zhang
he/him • UTA-STA

Why take this course?

You may have heard of “Deep Learning” or “Artificial Intelligence (AI)”

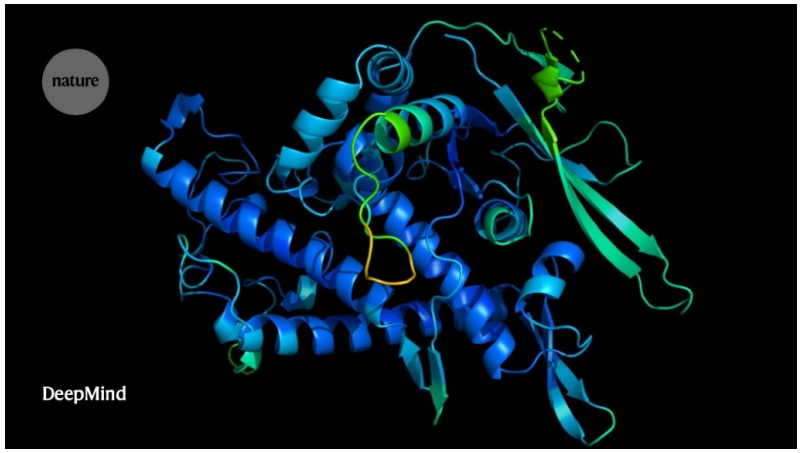

You may have heard of “Deep Learning” or “Artificial Intelligence (AI)”

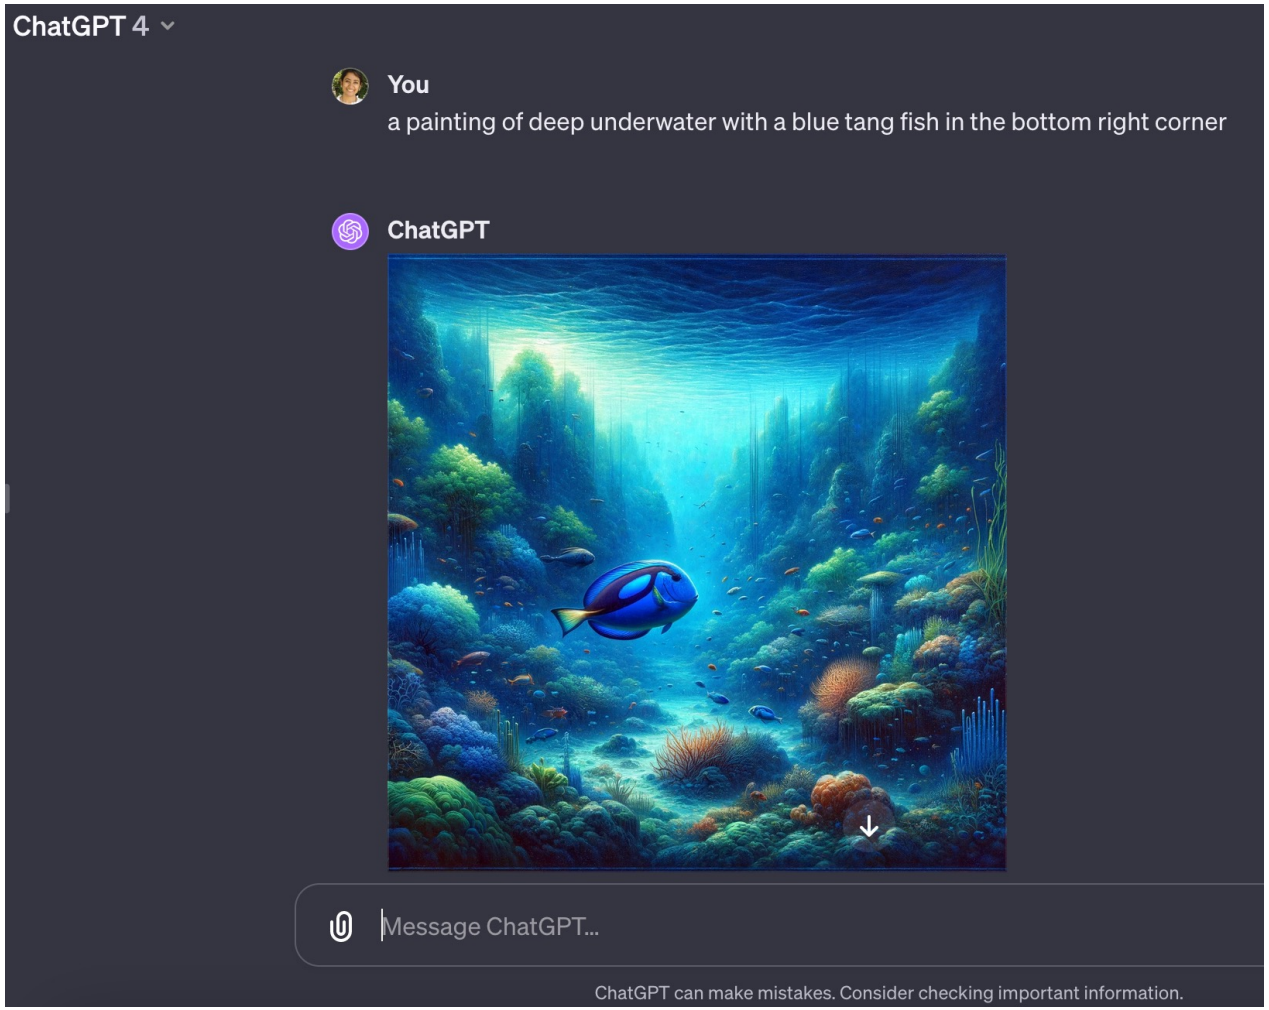

https://chat.openai.com/

You may have heard of “Deep Learning” or “Artificial Intelligence (AI)”

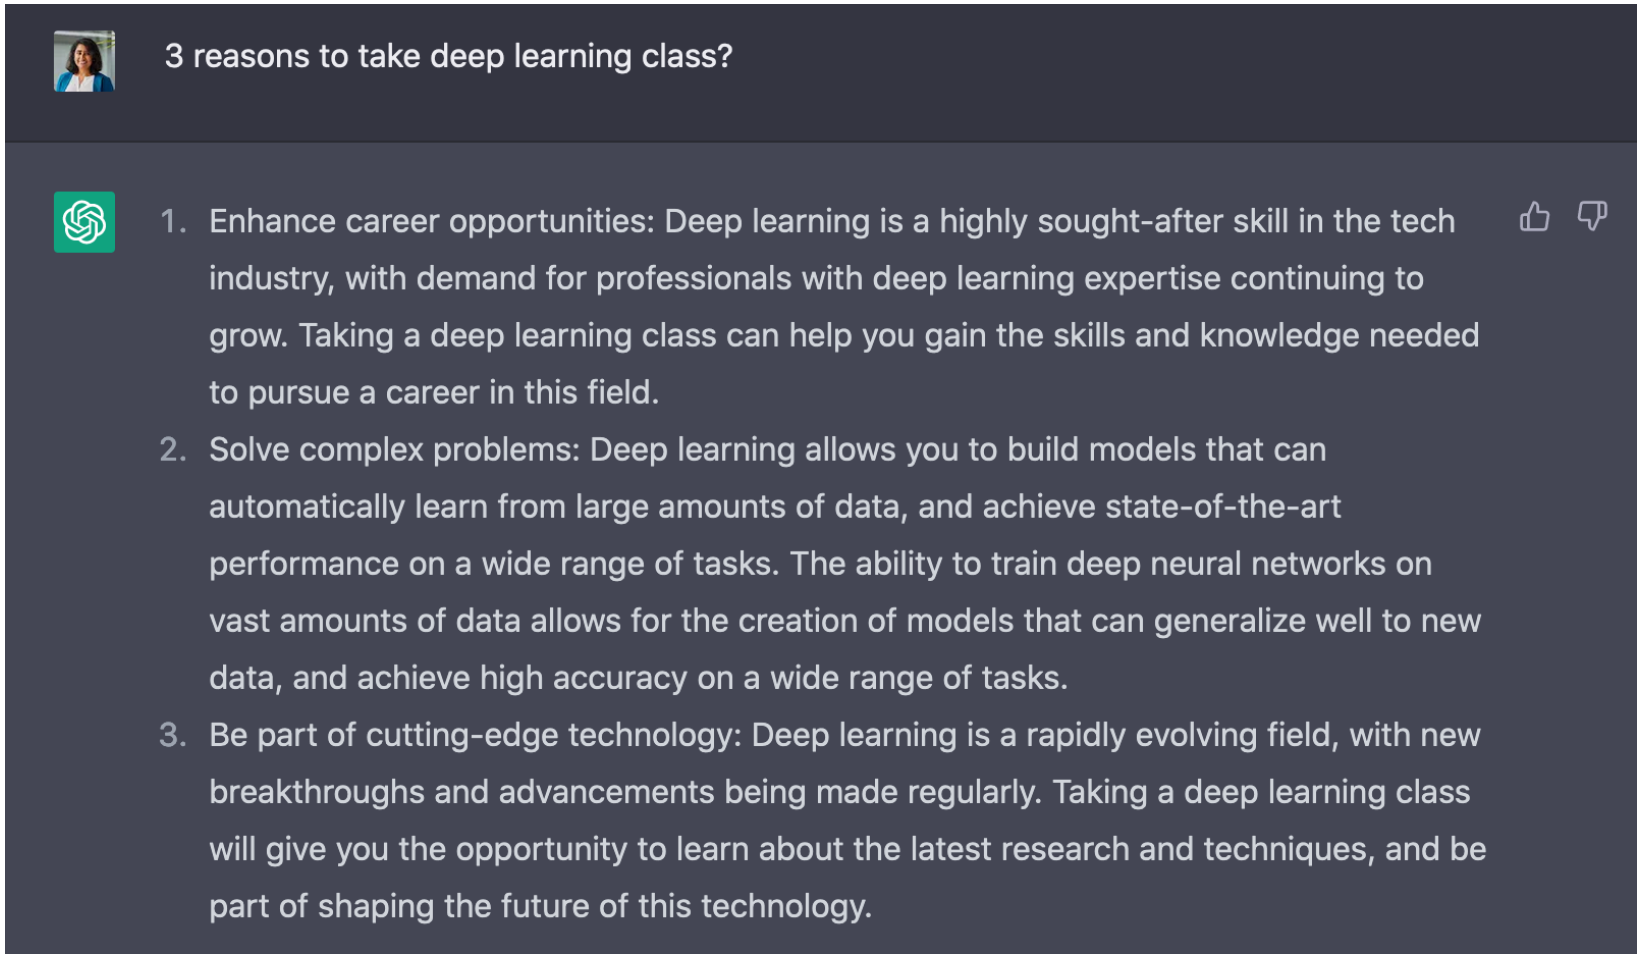

https://openai.com/blog/chatgpt/

You may have heard of “Deep Learning” or “Artificial Intelligence (AI)”

Gender Classifier	Darker Male	Darker Female	Lighter Male	Lighter Female	Largest Gap
Microsoft	94.0%	79.2%	100%	98.3%	20.8%
FACE++	99.3%	65.5%	99.2%	94.0%	33.8%
IBM	88.0%	65.3%	99.7%	92.9%	34.4%

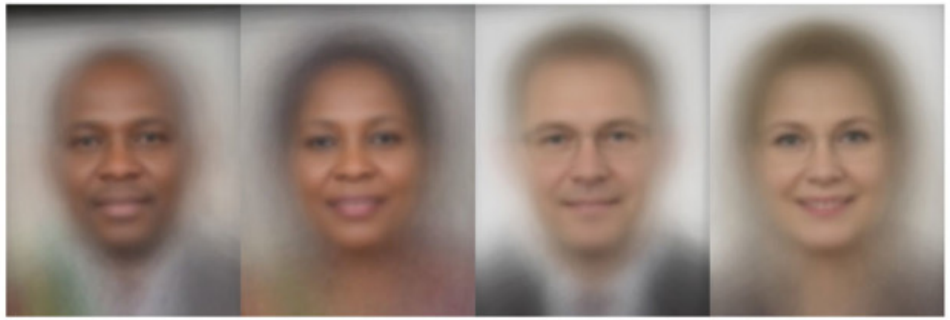

You may have heard of “Deep Learning” or “Artificial Intelligence (AI)”

Artificial intelligence / Machine learning

Training a single AI model can emit as much carbon as five cars in their lifetimes

Deep learning has a terrible carbon footprint.

by **Karen Hao** June 6, 2019

In review of fatal Arizona crash, U.S. agency says Uber software had flaws

By David Shepardson 4 MIN READ

WASHINGTON (Reuters) - An Uber self-driving test vehicle that struck and killed an Arizona woman in 2018 had software flaws, the National Transportation Safety Board said Tuesday as it disclosed the company’s autonomous test vehicles were involved in 37 crashes over the prior 18 months.

You may have heard of “Deep Learning” or “Artificial Intelligence (AI)”

Bloomberg US Edition

● Live Now Markets Economics Industries Tech AI Politics Wealth Pursuits Opinion Businessweek Equality Green

Technology
AI

Dall-E 3 Is So Good It’s Stoking an Artist Revolt Against AI Scraping

Artists are worried AI will take their jobs — so they’re getting creative.

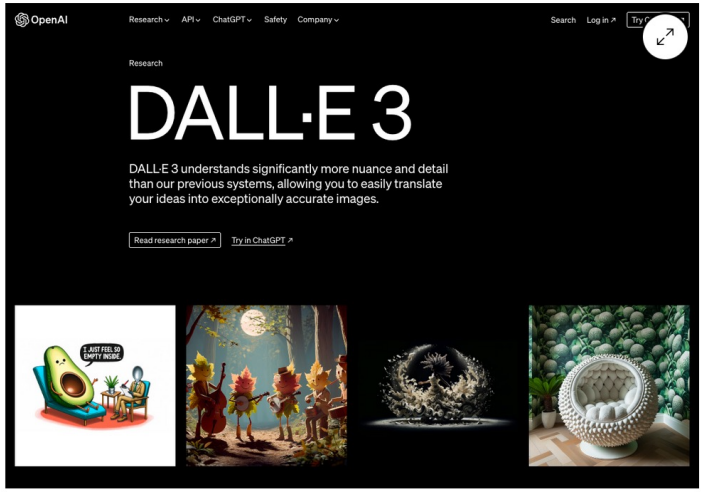

The Dall-E 3 website.

SANDER VAN DER LINDEN IDEAS JAN 22, 2024 7:00 AM

AI-Generated Fake News Is Coming to an Election Near You

Targeted, AI-generated political misinformation is already out there—and humans are falling for it.

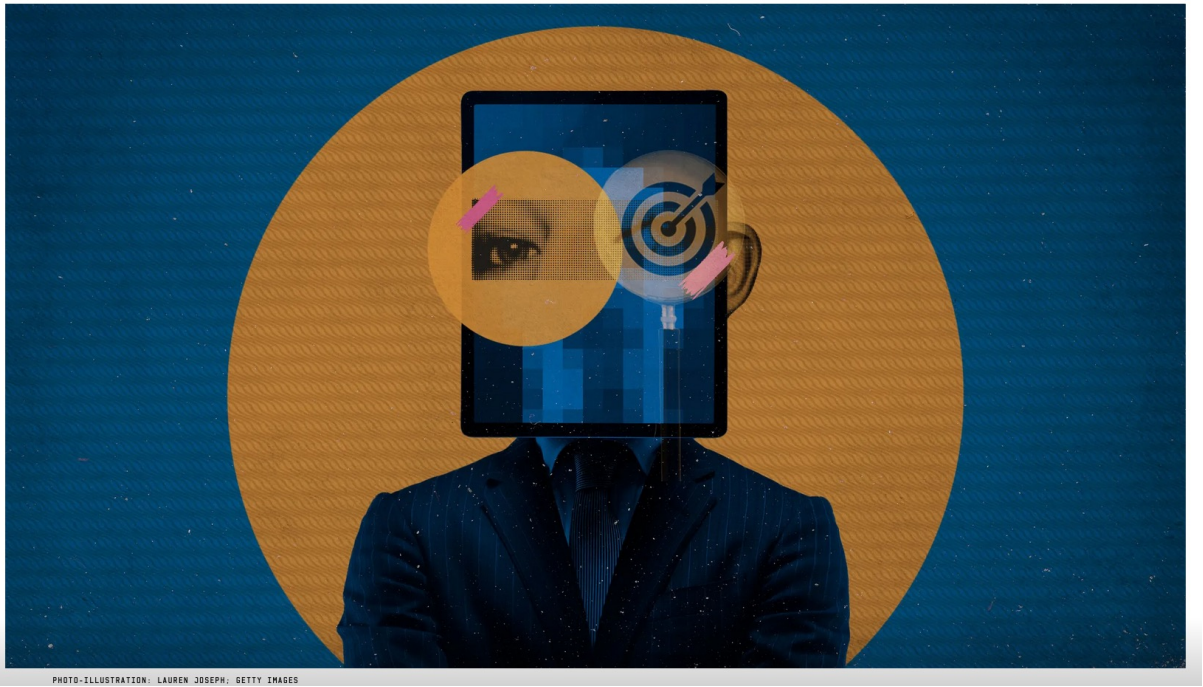

PHOTO-ILLUSTRATION: LAUREN JOSEPH; GETTY IMAGES

Our goal is to answers some important questions

- What is deep learning?
- What are the the different types of deep learning models?
- How to implement a deep learning models?
- What models are appropriate for different applications?
- Will our approach improve our understanding of the data or the problem?
- What are the ethical considerations when using deep learning models?

DALL-E 2 prompt “a cartoon of a penguin with a question mark”

Next time when you come across “Deep Learning” you will know:

What is Deep Learning?

(1) What is Machine Learning?

(2) How does it connect to Deep Learning?

(3) What is NOT Deep Learning?

What is Machine Learning?

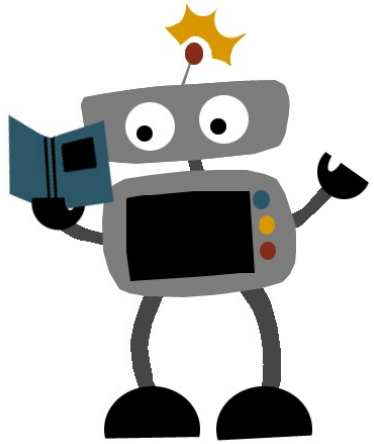

Input: X

Function: f

f(X)→Y

Output: Y

”Cooking?”

What is Machine Learning?

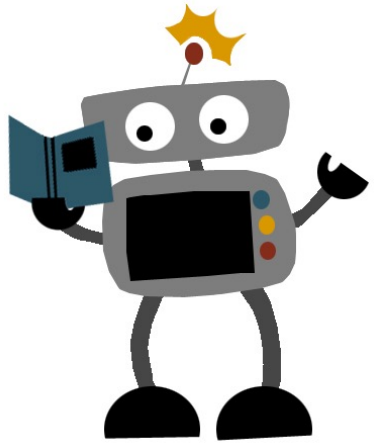

Supervised Learning

Input: X

Learned function: f

Output: Y

"Cooking?"

$f(X) \rightarrow Y$

What is Machine Learning?

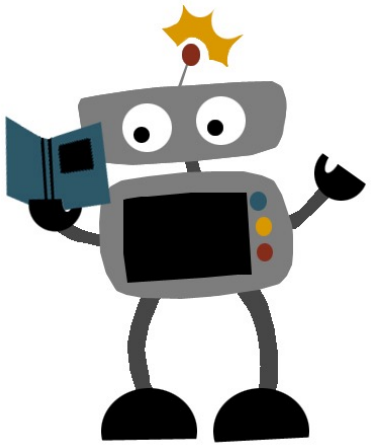

Input: X

I do not want sour cream in my burrito

Learned function: f

Output: Y

No quiero crema agrea en mi burrito

$f(X) \rightarrow Y$

What is Deep Learning?

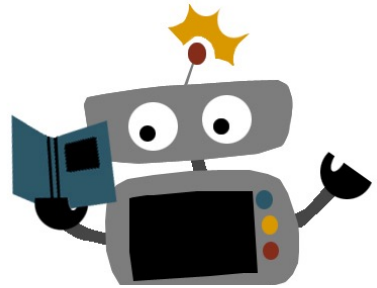

Input: X

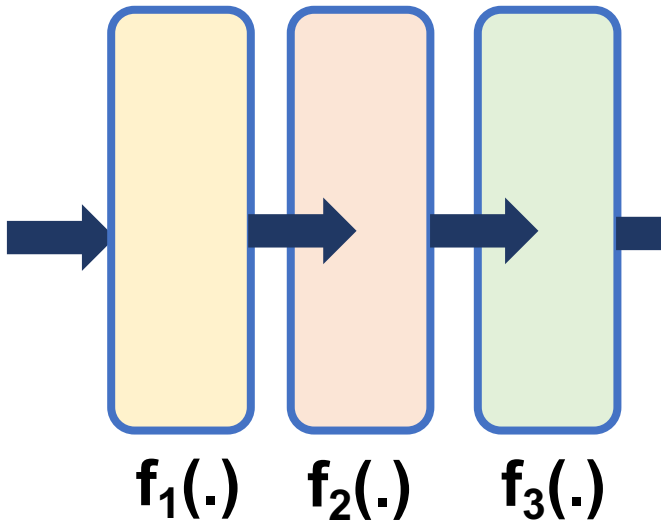

$$f_4\,(f_3\,(f_2\,(f_1\,(X))))$$

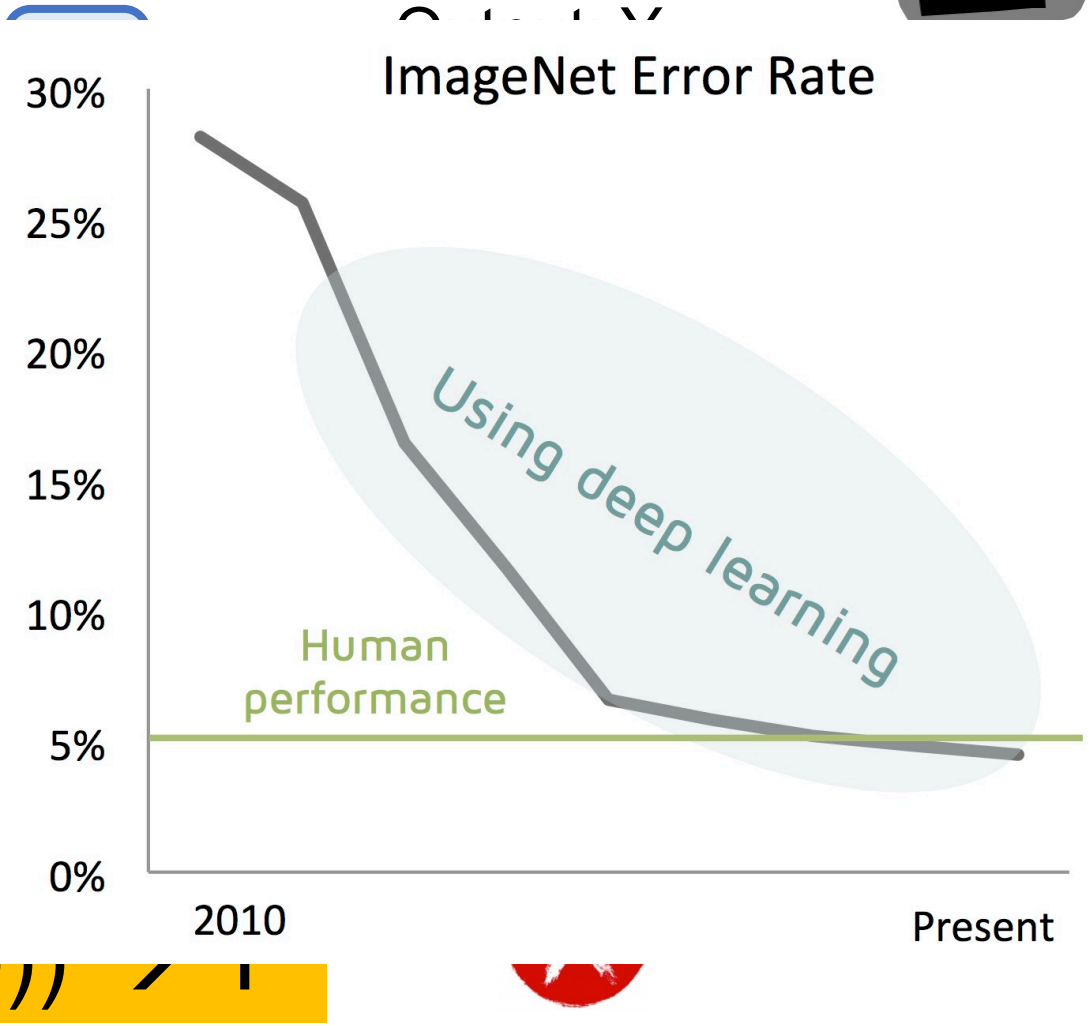

What is Deep Learning?

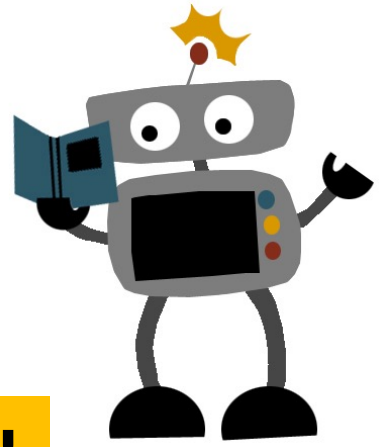

Input: X

Output: Y

”Cooking?”

Deep Learning DOES NOT mimic the brain!

TURN ANY PHOTO INTO AN ARTWORK – FOR FREE!

We use an algorithm inspired by the human brain. It uses the stylistic elements of one image to draw the content of another. Get your own artwork in just three steps.

[https://deepart.io]

WHAT IS DEEP LEARNING?

A newly re-invigorated form of machine learning, which is itself a subset of artificial intelligence, deep learning employs powerful computers, massive data sets, “supervised” (trained) neural networks and an algorithm called back-propagation (backprop for short) to recognize objects and translate speech in real time by mimicking the layers of neurons in a human brain’s neocortex.

[https://builtin.com/artificial-intelligence/deep-learning]

”Cooking?”

What is NOT Deep Learning?

Deep Learning is NOT AI

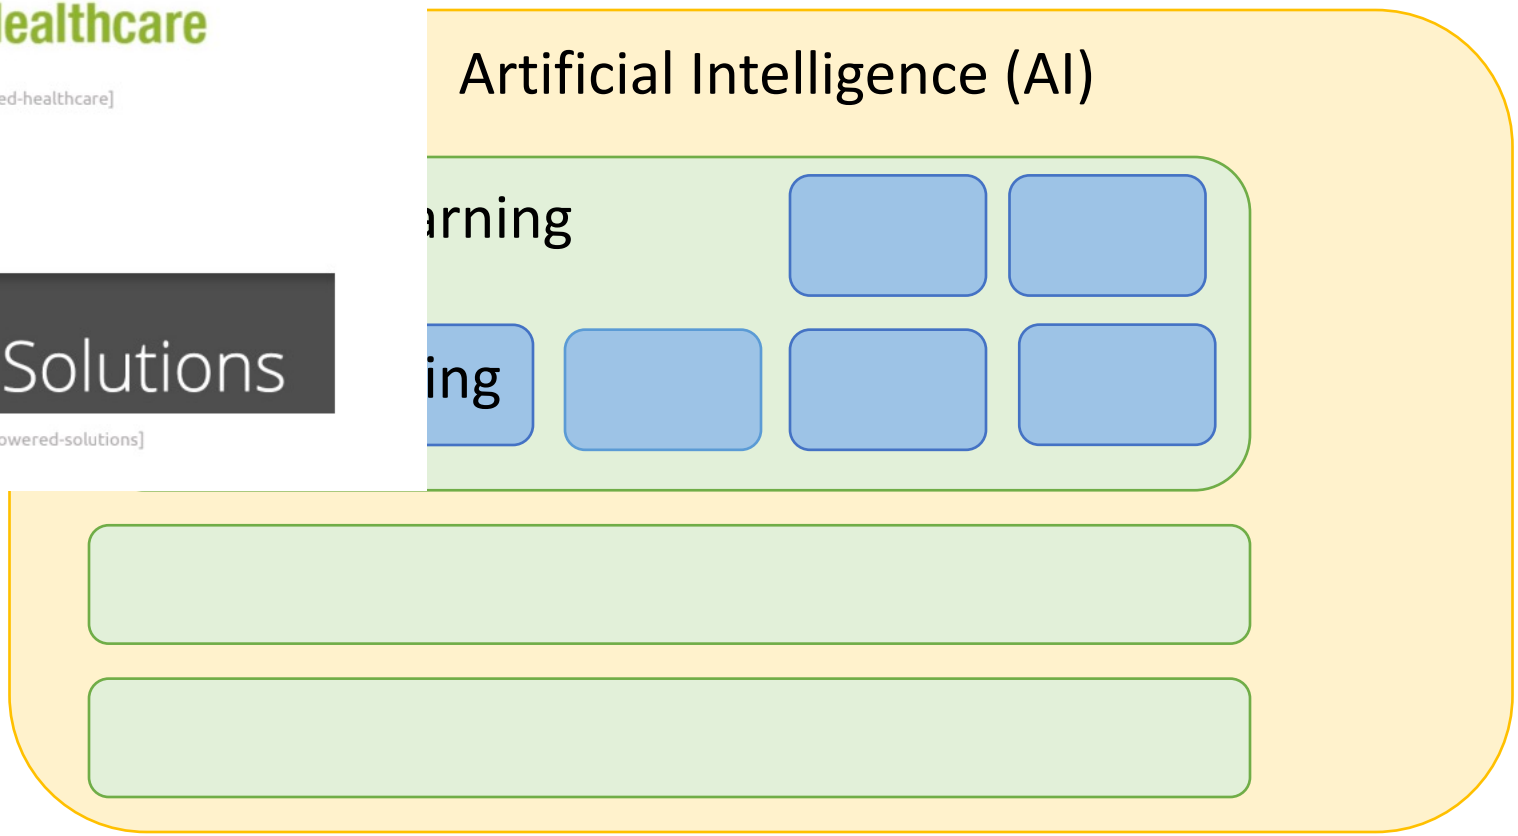

Recap

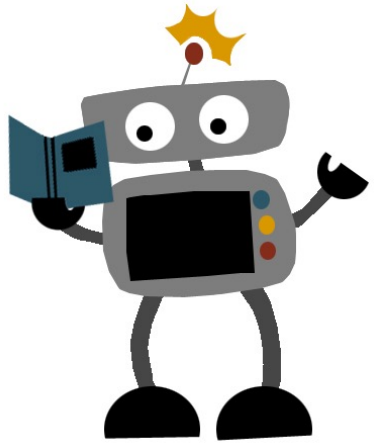

Input: X

Machine Learning

$f(X) \rightarrow Y$

Output: Y

"Cooking?"

Deep Learning is NOT AI

Deep Learning DOES NOT mimic the brain!

Deep Learning

$f_4 (f_3 (f_2 (f_1 (X)))) \rightarrow Y$

Questions?

Ice-breaker
(a.k.a “please-don’t-make-me-do-this” activity)

- Turn to the person sitting next to you and introduce yourself!

- **What do you hope to learn/be able to do by the end of this course?**

Join at menti.com | use code 35 33 01 2

ChatGPT prompt “minimalist landscape painting of a deep underwater scene with a blue tang fish in the bottom right corner”

Course Logistics

The Course Website

http://cs.brown.edu/courses/csci1470

- Your one-stop-shop for:
  - Syllabus
  - Lecture, lab, & assignment schedules
  - Links to important forms, etc.
  - …

The Canvas Website

https://canvas.brown.edu/courses/1094502

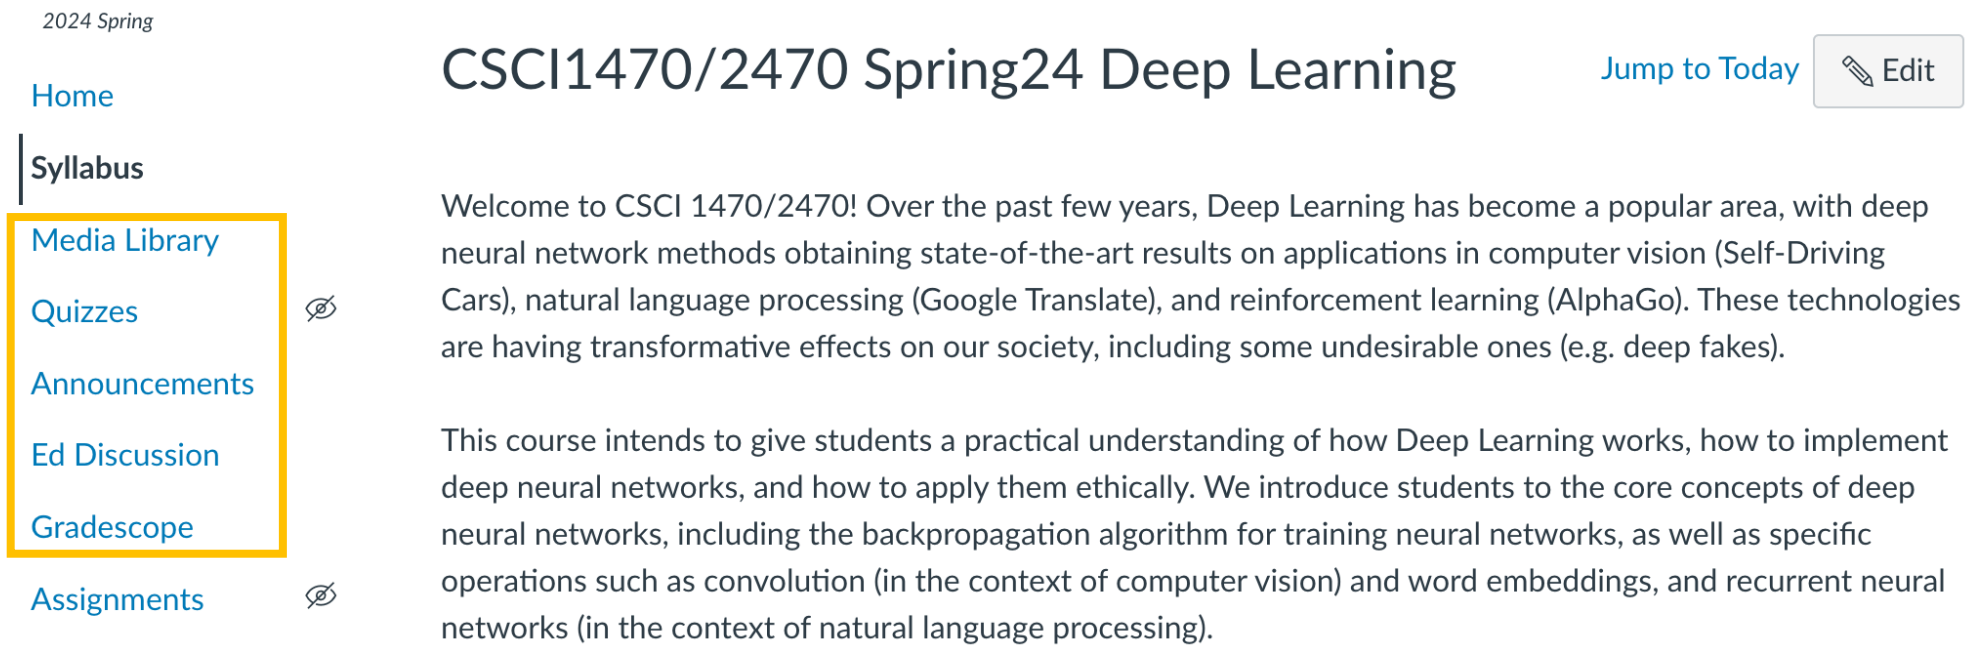

- Your access to:
  - Ed Discussion
  - GradeScope
  - Lectures
  - Weekly quizzes
  - Weekly course announcements from instructor

Which Version of the Course Should I Take?

CS 1470

- Undergrads + grads
- Lectures
- Labs
- Assignments (Code + Written)
- Group final project
  - Implement existing research paper
  - Poster presentation

CS 2470

- Grad students only
- Same Lectures
- Same Labs
- Same Assignments, plus:
  - Additional required features
  - Additional written questions
- Group final project
  - Try something new
  - Oral presentation

Which Version of the Course Should I Take?

CS 1470

- Undergrads + grads
- Lectures
- Labs
- Assignments (Code + Written)
- Group final project
  - Implement existing research paper
  - Poster presentation

CS 2470

- Grad students only
- Same Lectures
- Same Labs
- Same Assignments, plus:
  - Additional required features
  - Additional written questions
- Group final project
  - Try something new
  - Oral presentation

Six Awesome Assignments

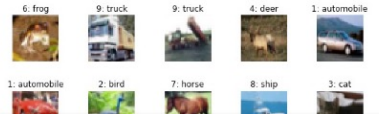

TensorFlow 2.0

Language modeling

Image Captioning

Variational Autoencoders

Brown Deep Learning Day!

- Course final project
- In-person mini conference!
- Poster sessions and presentations
  - Grouped by theme: e.g. vision, language, robotics, ...
- **Tentative Date: May 6-7, 2024**
- Details forthcoming!

Deep Learning Day (Spring 2022)

Lectures and class participation

- In-person Lectures
  - Lecture recordings available
  - Recordings posted to Canvas (Media Library)
- Weekly quiz on Canvas
  - Released on Wednesday (starts next week!)
  - Due on Thursday
  - Minimum time/effort if you attend class or watch lectures regularly
  - No deadline extensions!

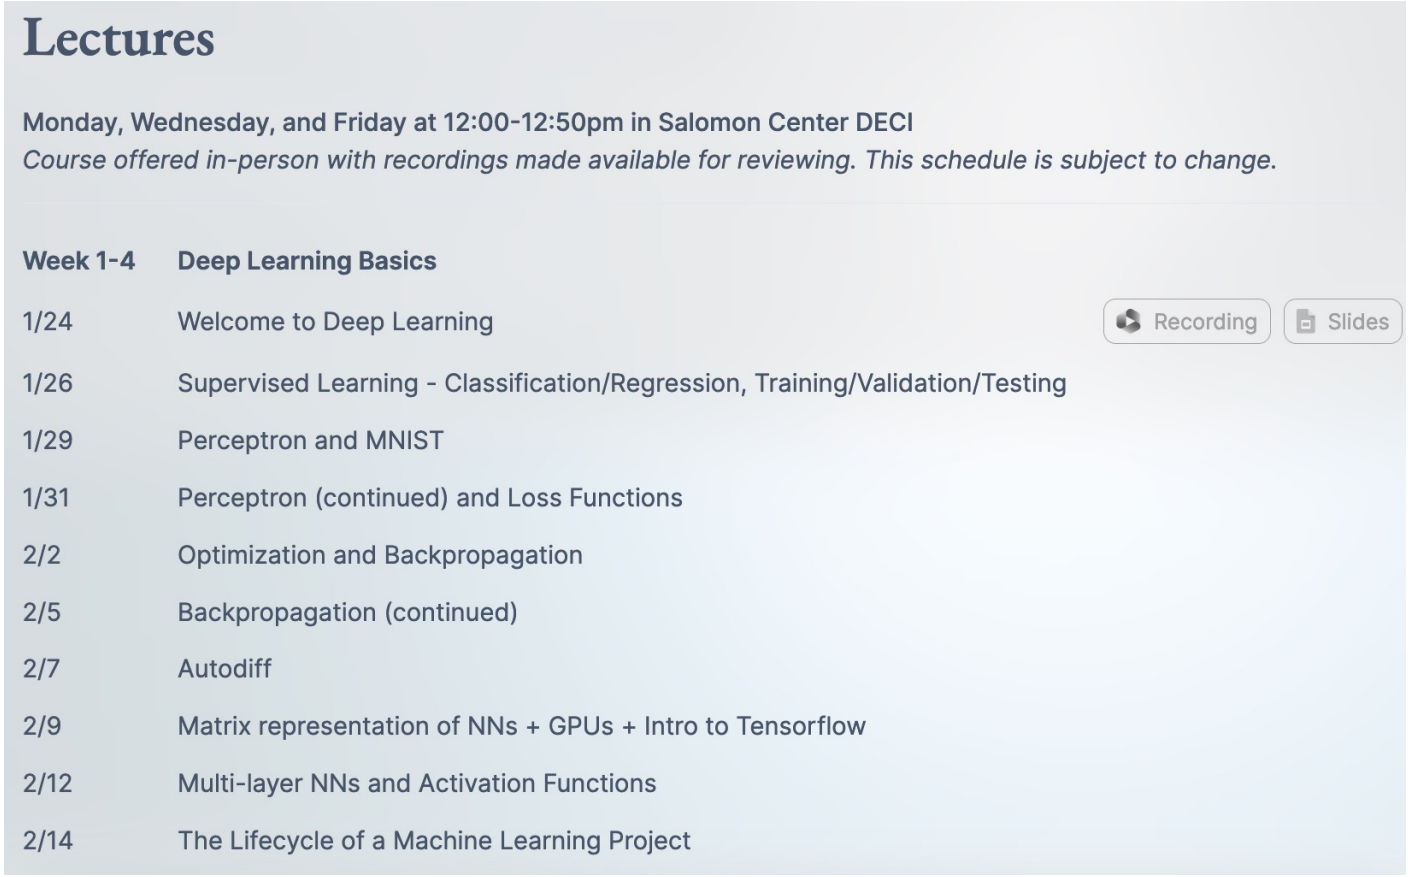

Lectures

Monday, Wednesday, and Friday at 12:00-12:50pm in Salomon Center DECI
Course offered in-person with recordings made available for reviewing. This schedule is subject to change.

Week 1-4	**Deep Learning Basics**	
1/24	Welcome to Deep Learning	Recording Slides
1/26	Supervised Learning - Classification/Regression, Training/Validation/Testing	
1/29	Perceptron and MNIST	
1/31	Perceptron (continued) and Loss Functions	
2/2	Optimization and Backpropagation	
2/5	Backpropagation (continued)	
2/7	Autodiff	
2/9	Matrix representation of NNs + GPUs + Intro to Tensorflow	
2/12	Multi-layer NNs and Activation Functions	
2/14	The Lifecycle of a Machine Learning Project	

Labs and office hours

- In-person lab and office hours
  - Will take a week or so to finalize
  - **Team work highly encouraged for lab hours!**
  - Might have remote options (in the works!)
- **NEW FORMAT** for office hours
  - Conceptual hours (group-based help/discussion on concepts)
  - Collab hours (group-based help/discussion on code)
  - **TAs will only look at your code in the first week of assignment release (including 1-1 debugging) during collab hours**

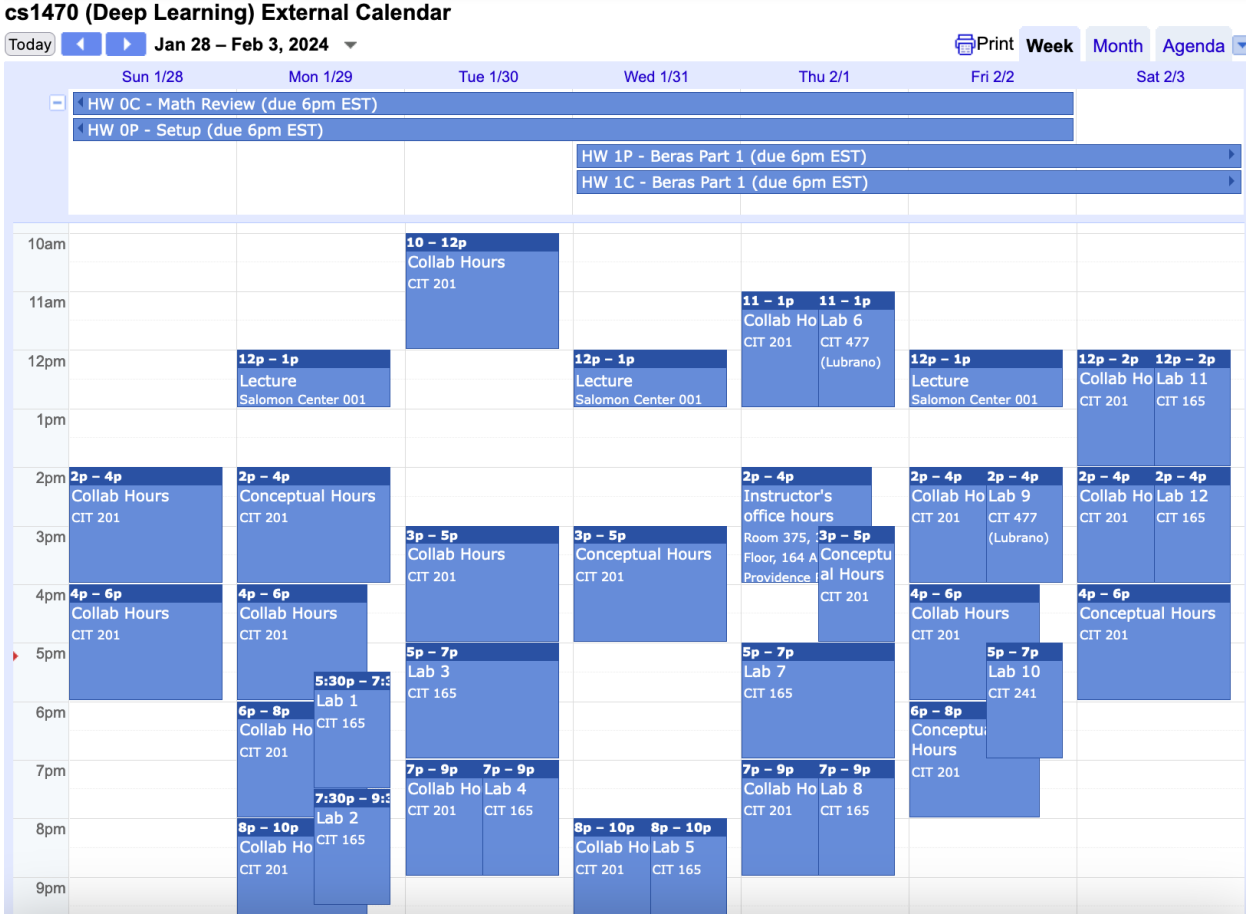

Assignment logistics

- Assignments
  - Get stencils via Github Classroom
  - Submission via Gradescope

Somewhat new!

Homework and Lab 0 + SRC discussions

- Homework 0 (will be released today!)
  - Review of relevant math and probability concepts
  - Setting up programming environment
  - Points for completion only (**deadline Feb 02**)
- Lab 0 (will be released today!)
  - Review of python and numpy
  - Complete on your own (**preferably by Jan 29**)
- SRC discussion sessions
  - **Sign up to attend 2 sessions for this semester**
  - STAs will provide prompts related to that week's homework

The only thing set in stone is our excitement to learn!

- ❑ Will try things for class engagement
- ❑ Due dates might move around
- ❑ We might have make-up classes/labs
- ❑ The schedule will remain flexible till the end
- ❑ Suggestions are welcome!

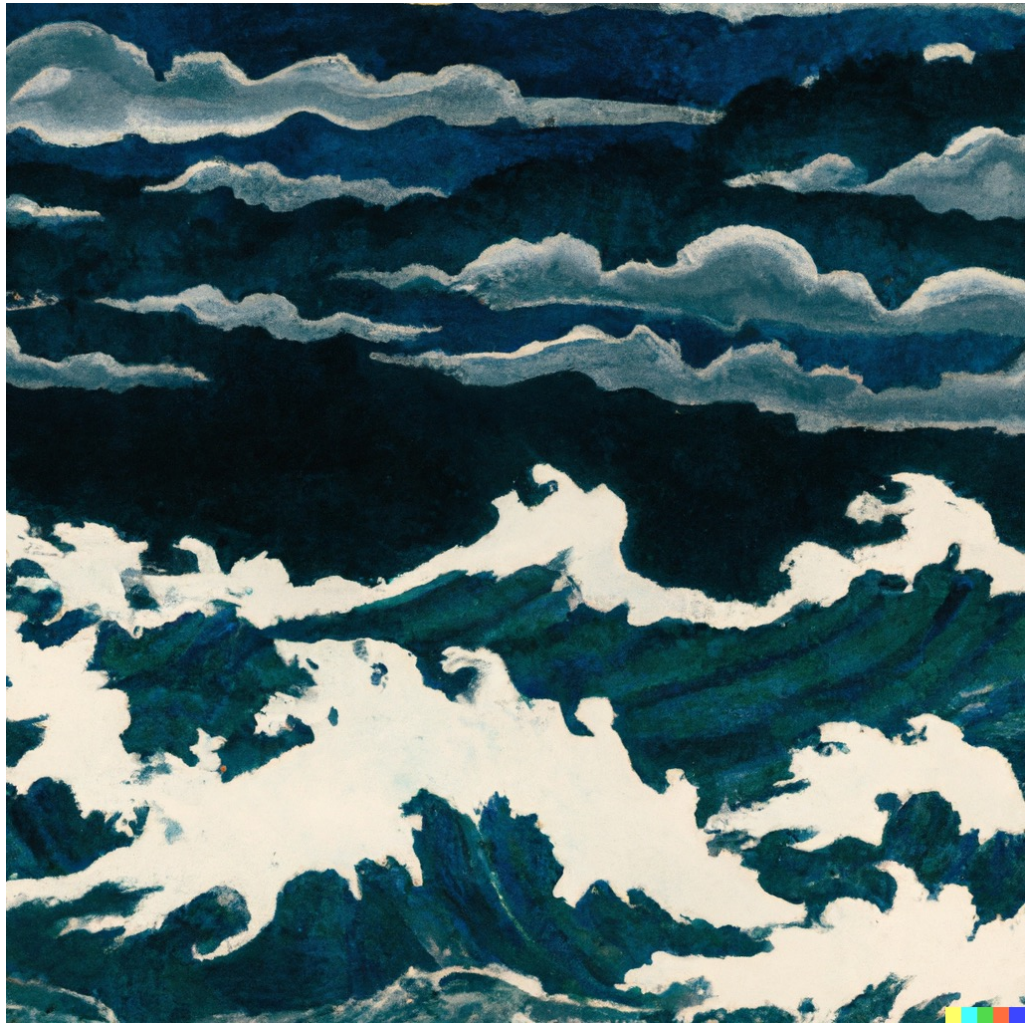

DALL-E 2 prompt “japanese painting of a stormy ocean”

This course is also offered in the Fall!

Professor Chen Sun
(taught in Fall 2022)

Acknowledgements

Original course material developed by
Professor Daniel Ritchie
and previous FABULOUS TA staff

Immediate Action Items

- **Read the course missive**; sign the collaboration policy form
- Make sure you can access the course Ed Discussion page via Canvas (all announcements will happen there)
- Create a GitHub account (if you don't have one already)
- Make sure you can access the course GradeScope via Canvas (all assignments will be submitted there)
- **Start working on HW0 – Submission deadline Feb 02 (Friday 6PM EST)**
- **Do Lab 0 (Especially if you've never used Python and/or numpy before!) – Finish before lab 1!**
- If you want to do work on CS department machines over SSH, you'll need a CS login
  - All enrolled students automatically have one (it's the same as your Brown ID)
  - If you are not yet enrolled, you'll need to email problem@cs.brown.edu
  - More information about CS accounts can be found here.
  - *Note that a CS login is **not required** for this course, as we are not using handin or any other department infrastructure to manage files or submissions*

Questions?