

CMSC 475/675 Neural Networks

Lecture 0: `__init__()`



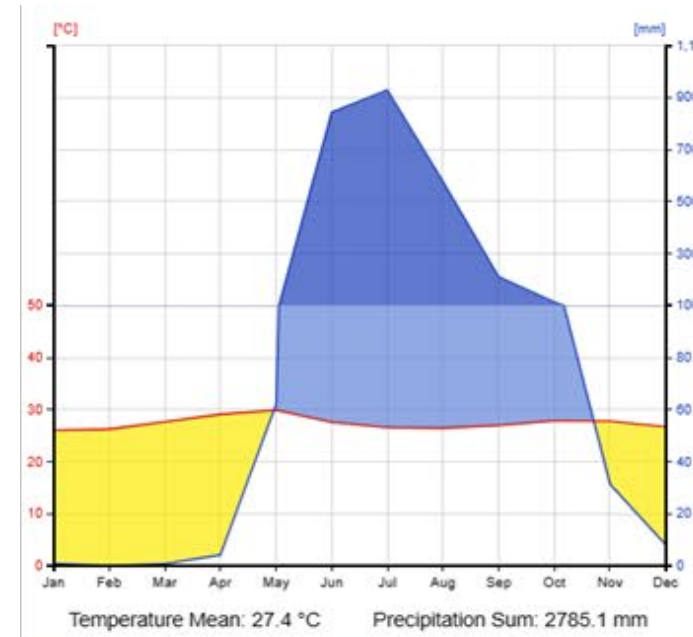
Tejas Gokhale ~

Tejas:
Gokhale:

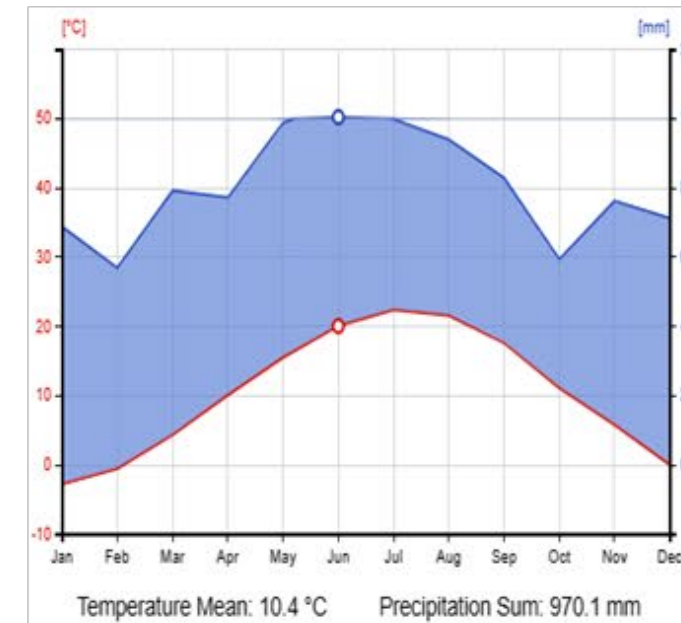
Tay + Juss
Go + Clay



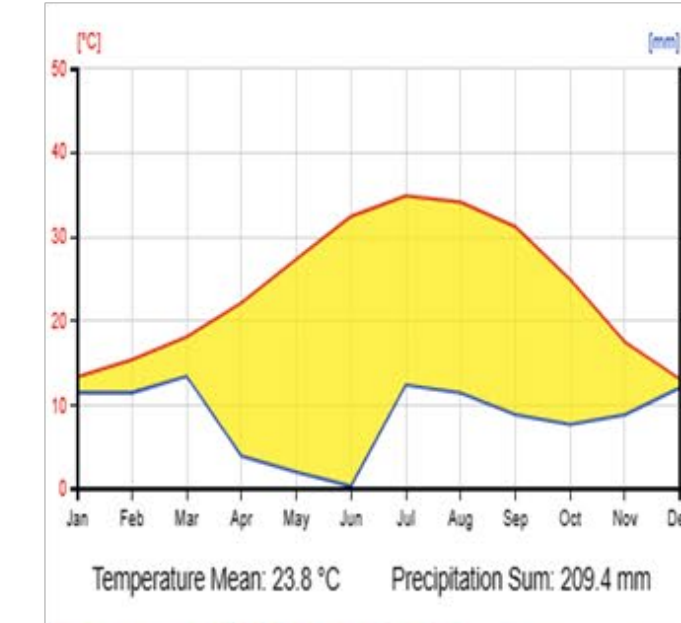
Assistant Professor
Computer Science
University of Maryland, Baltimore County



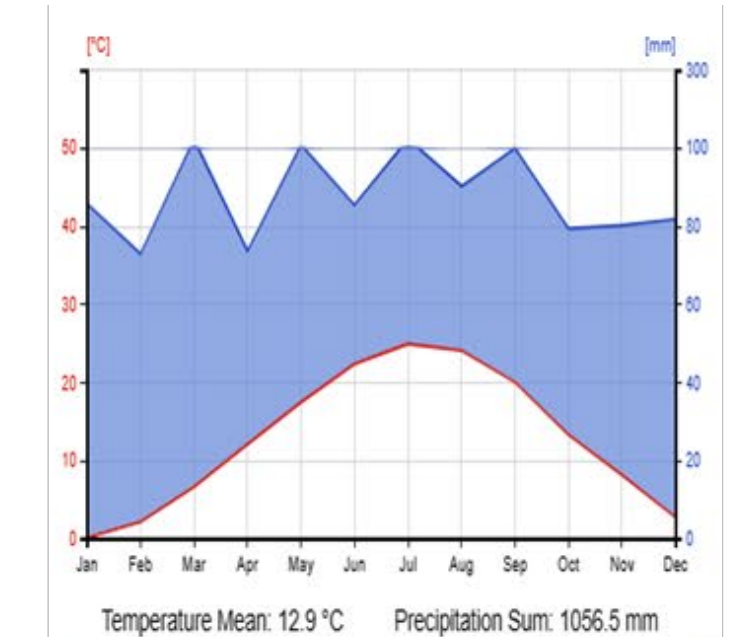
2011—2015
B.E. (Honours)
BITS Pilani
(Goa)



2016-2018:
M.S.
Carnegie Mellon
University



2018—2023:
Ph.D.
Arizona State
University



2023—present
Assistant Professor
University of Maryland
Baltimore County

<https://www.tejasgokhale.com/>

Course Staff

Instructor: Tejas Gokhale
Assistant Professor, CSEE



Wednesday 2:30 – 3:30 PM

ITE 214

gokhale@umbc.edu

TA: Ziwei Zhang
Ph.D. Researcher, CSEE



Monday 1:00 – 2:00 PM

Tuesday 2:30 – 3:30 PM

ITE 334

ziweiz1@umbc.edu

Office Hours

Class Website

- This is your primary source for class material (slides, notes, readings, references, etc.)
- Assignments & Grades will be released on Blackboard

<https://courses.cs.umbc.edu/graduate/675/>

Quick Round of Introductions

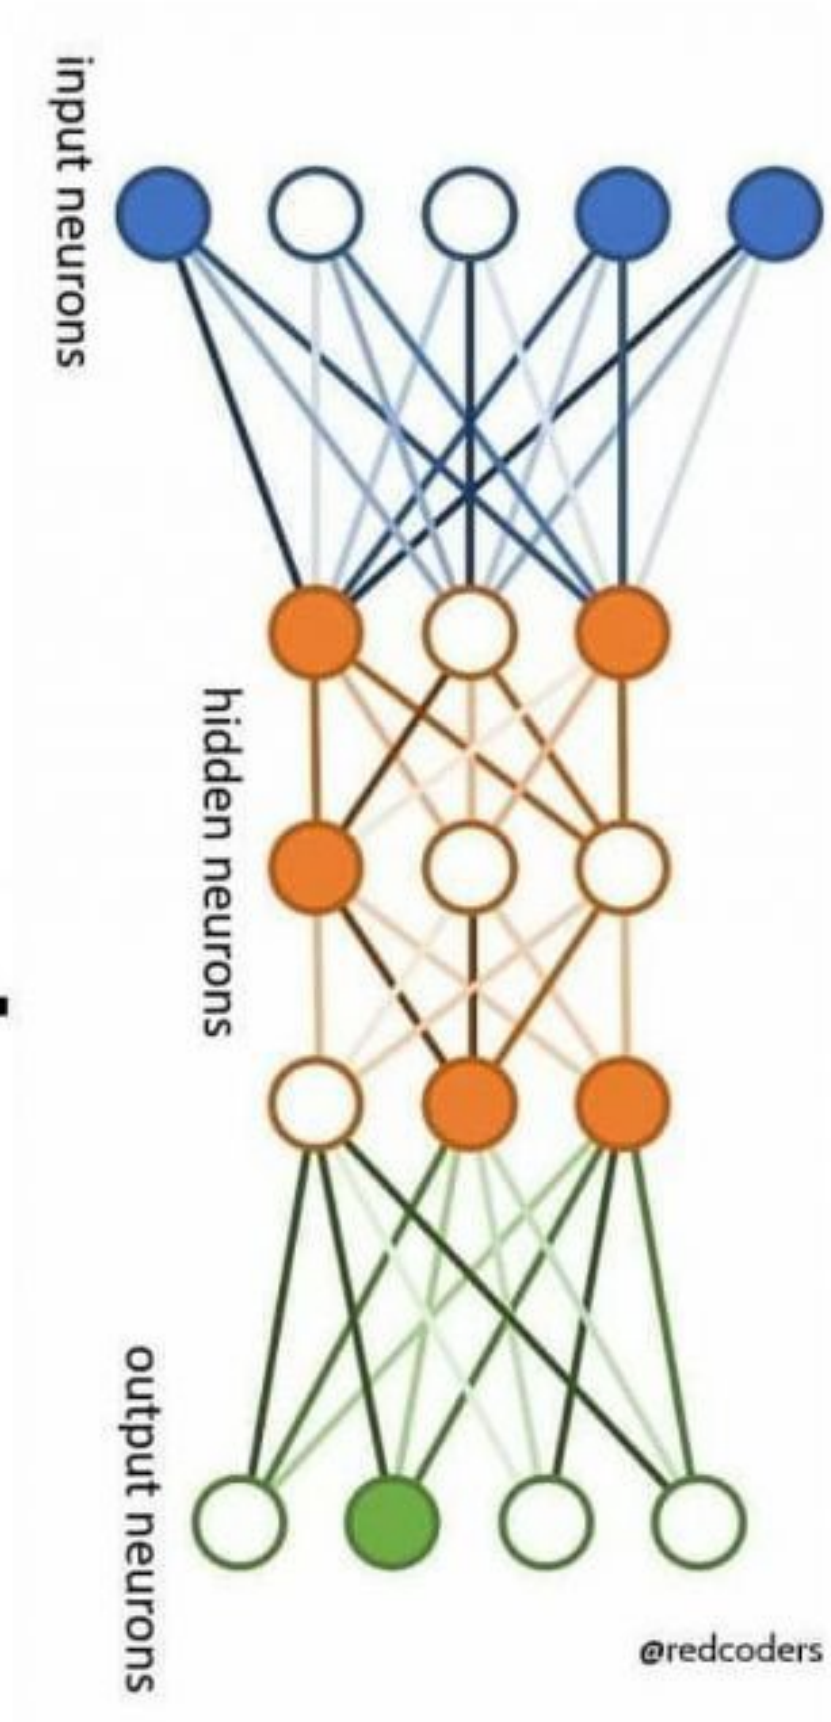
1. Name
2. Major (e.g. CS, CE, EE, etc.)
3. Level (e.g. BS, MS, Ph.D. etc.)
4. Why are you taking this class?



THIS IS A NEURAL NETWORK.

**IT MAKES MISTAKES.
IT LEARNS FROM THEM.**

BE LIKE A NEURAL NETWORK.



**How Neural Network works?
Neurons:**



What this course **is NOT** about ...

- ❌ A “review” of the state-of-the-art:
 - we will talk about some recent work, but the main goal is to learn fundamentals

- ❌ Applications-focused or domain-focused:
 - CMSC 472/672: Computer Vision (next semester?)
 - CMSC 473/673: Natural Language Processing

- ❌ Broad introduction to “AI”: **CMSC 471/671: Artificial Intelligence**

- ❌ “Hands-on” coding
 - We expect you to be proficient in Python programming
 - The TA will organize a PyTorch tutorial (during regular class hours)
 - Homework and Projects: implementation. Lectures: mathematical concepts

What this course **is about** ...

- Mathematical Models of N.N. Architectures and N.N. Optimization Routines
- Supervised, Unsupervised, Self-Supervised Paradigms of Learning with N.N.
- Relationships between Deep Learning and “classical” models
- Generalization, Robustness, and Reliability issues
- ...

In short, this course expects:

- mathematical rigor
- interest in fundamentals
- programming proficiency



Machine Learning: The Success Story



IS "DEEP LEARNING" A REVOLUTION IN ARTIFICIAL INTELLIGENCE?



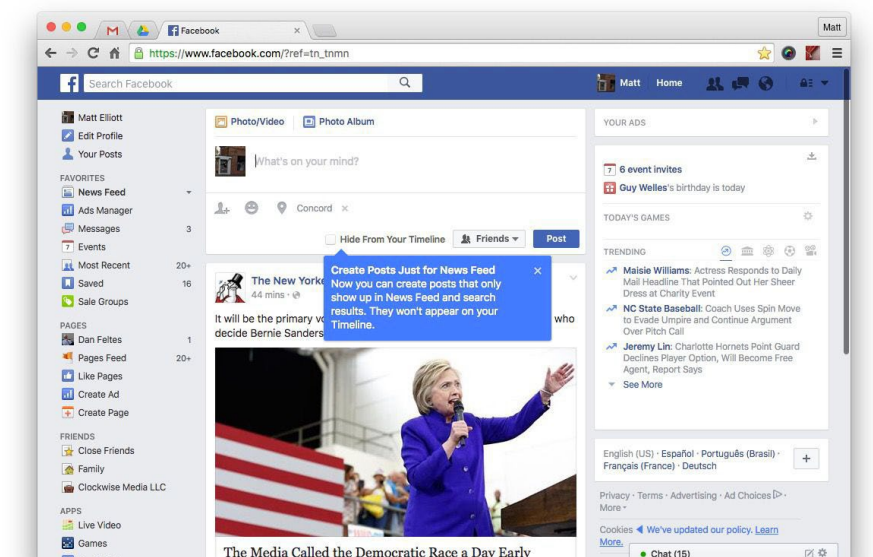
Andrew Ng
@AndrewYNg

Follow

"AI is the new electricity!" Electricity transformed countless industries; AI will now do the same.

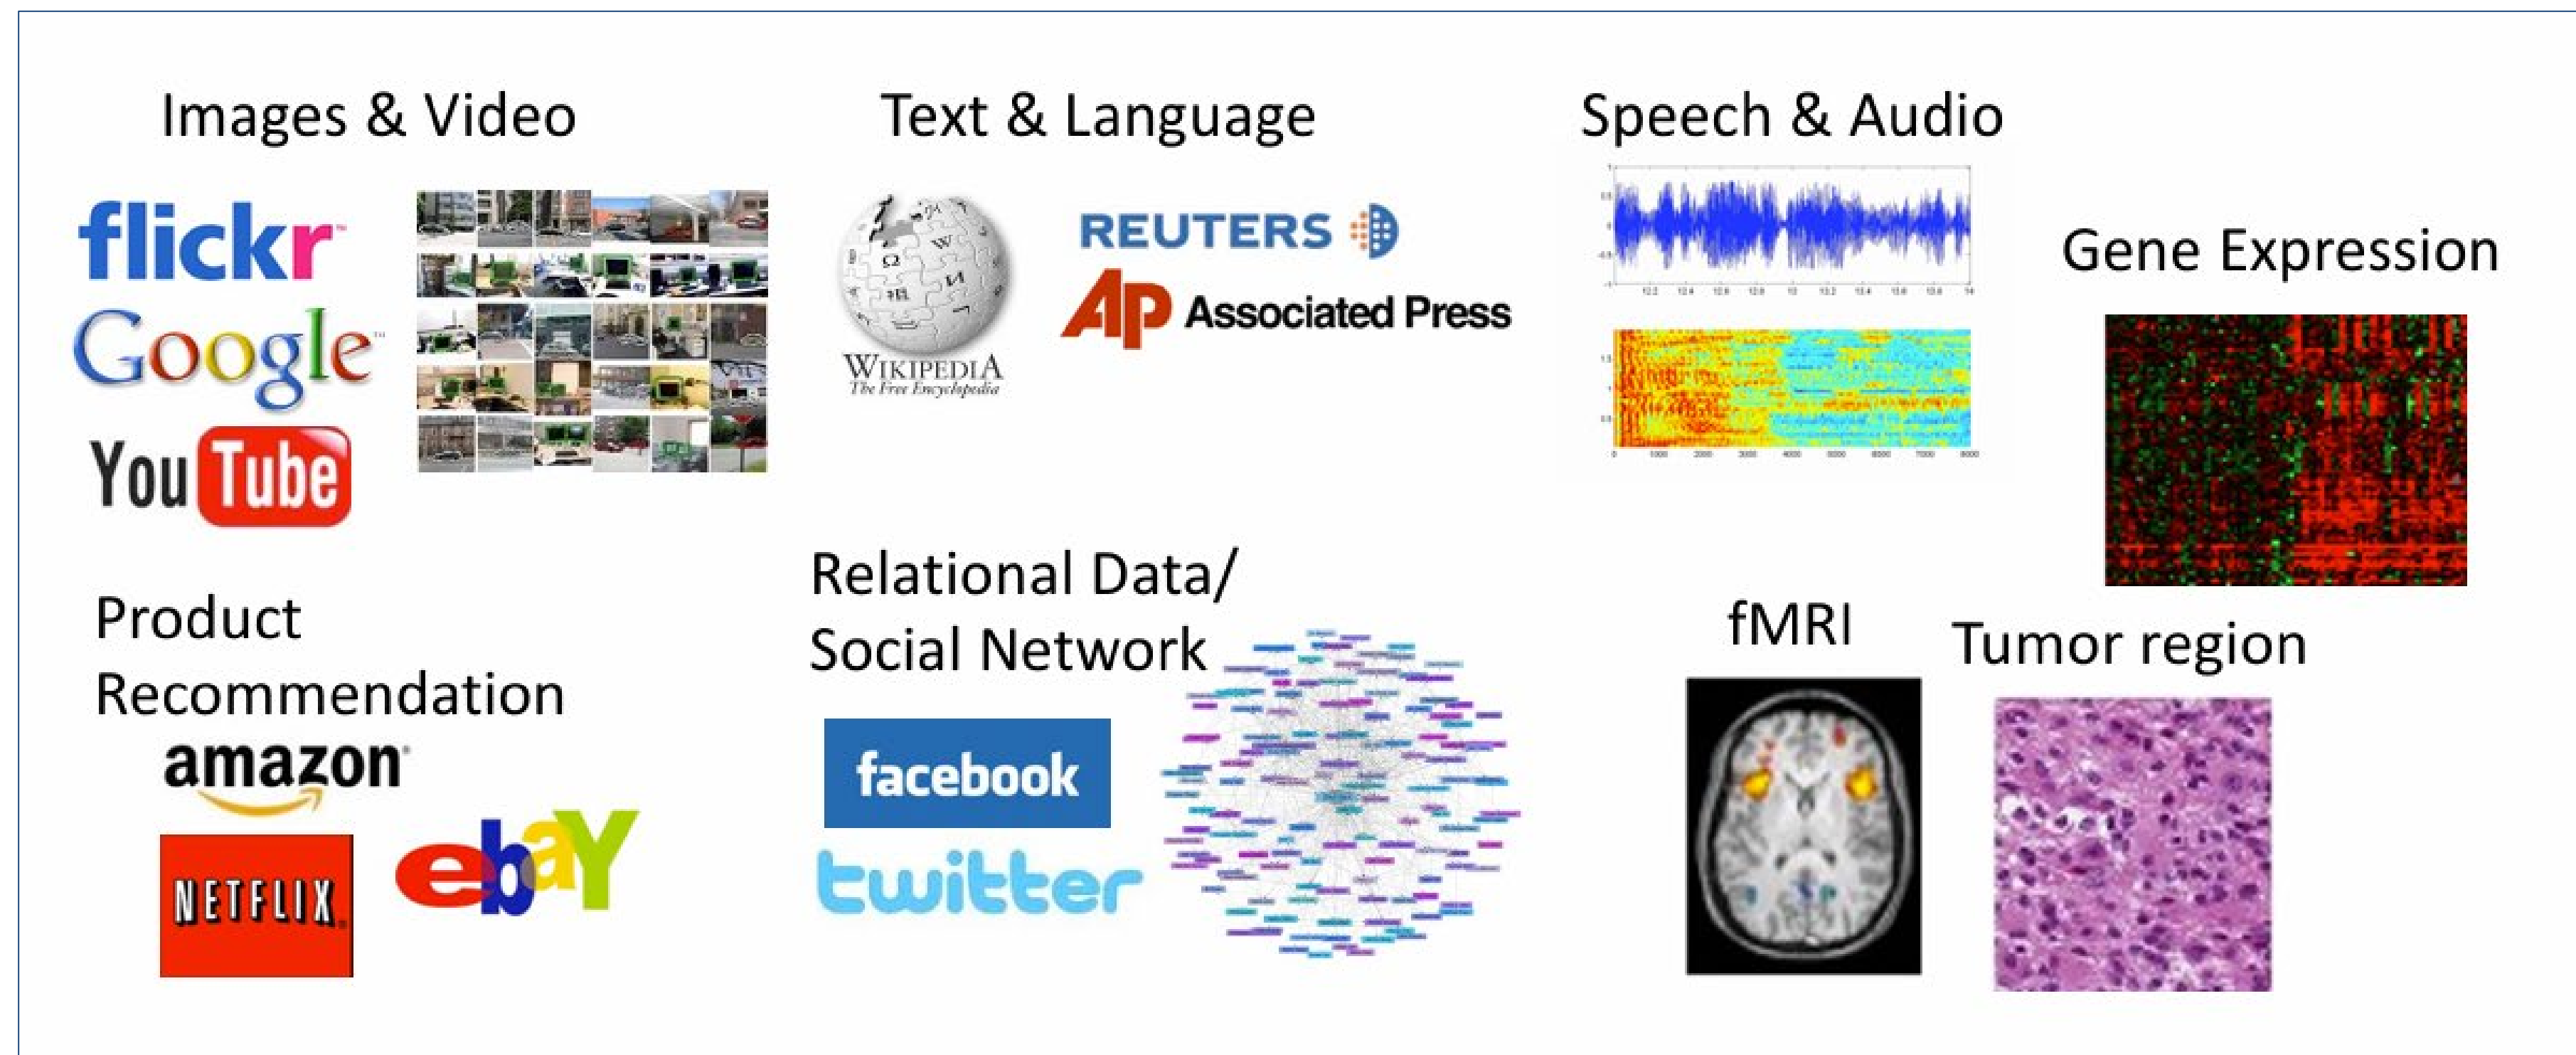
2016: The Year That Deep Learning Took Over

WHY DEEP LEARNING IS SUDDENLY CHANGING YOUR LIFE



Structure in the data ...

- 21st century: massive increase in amount of digital data available from the internet, cameras, microphones, scientific measurements



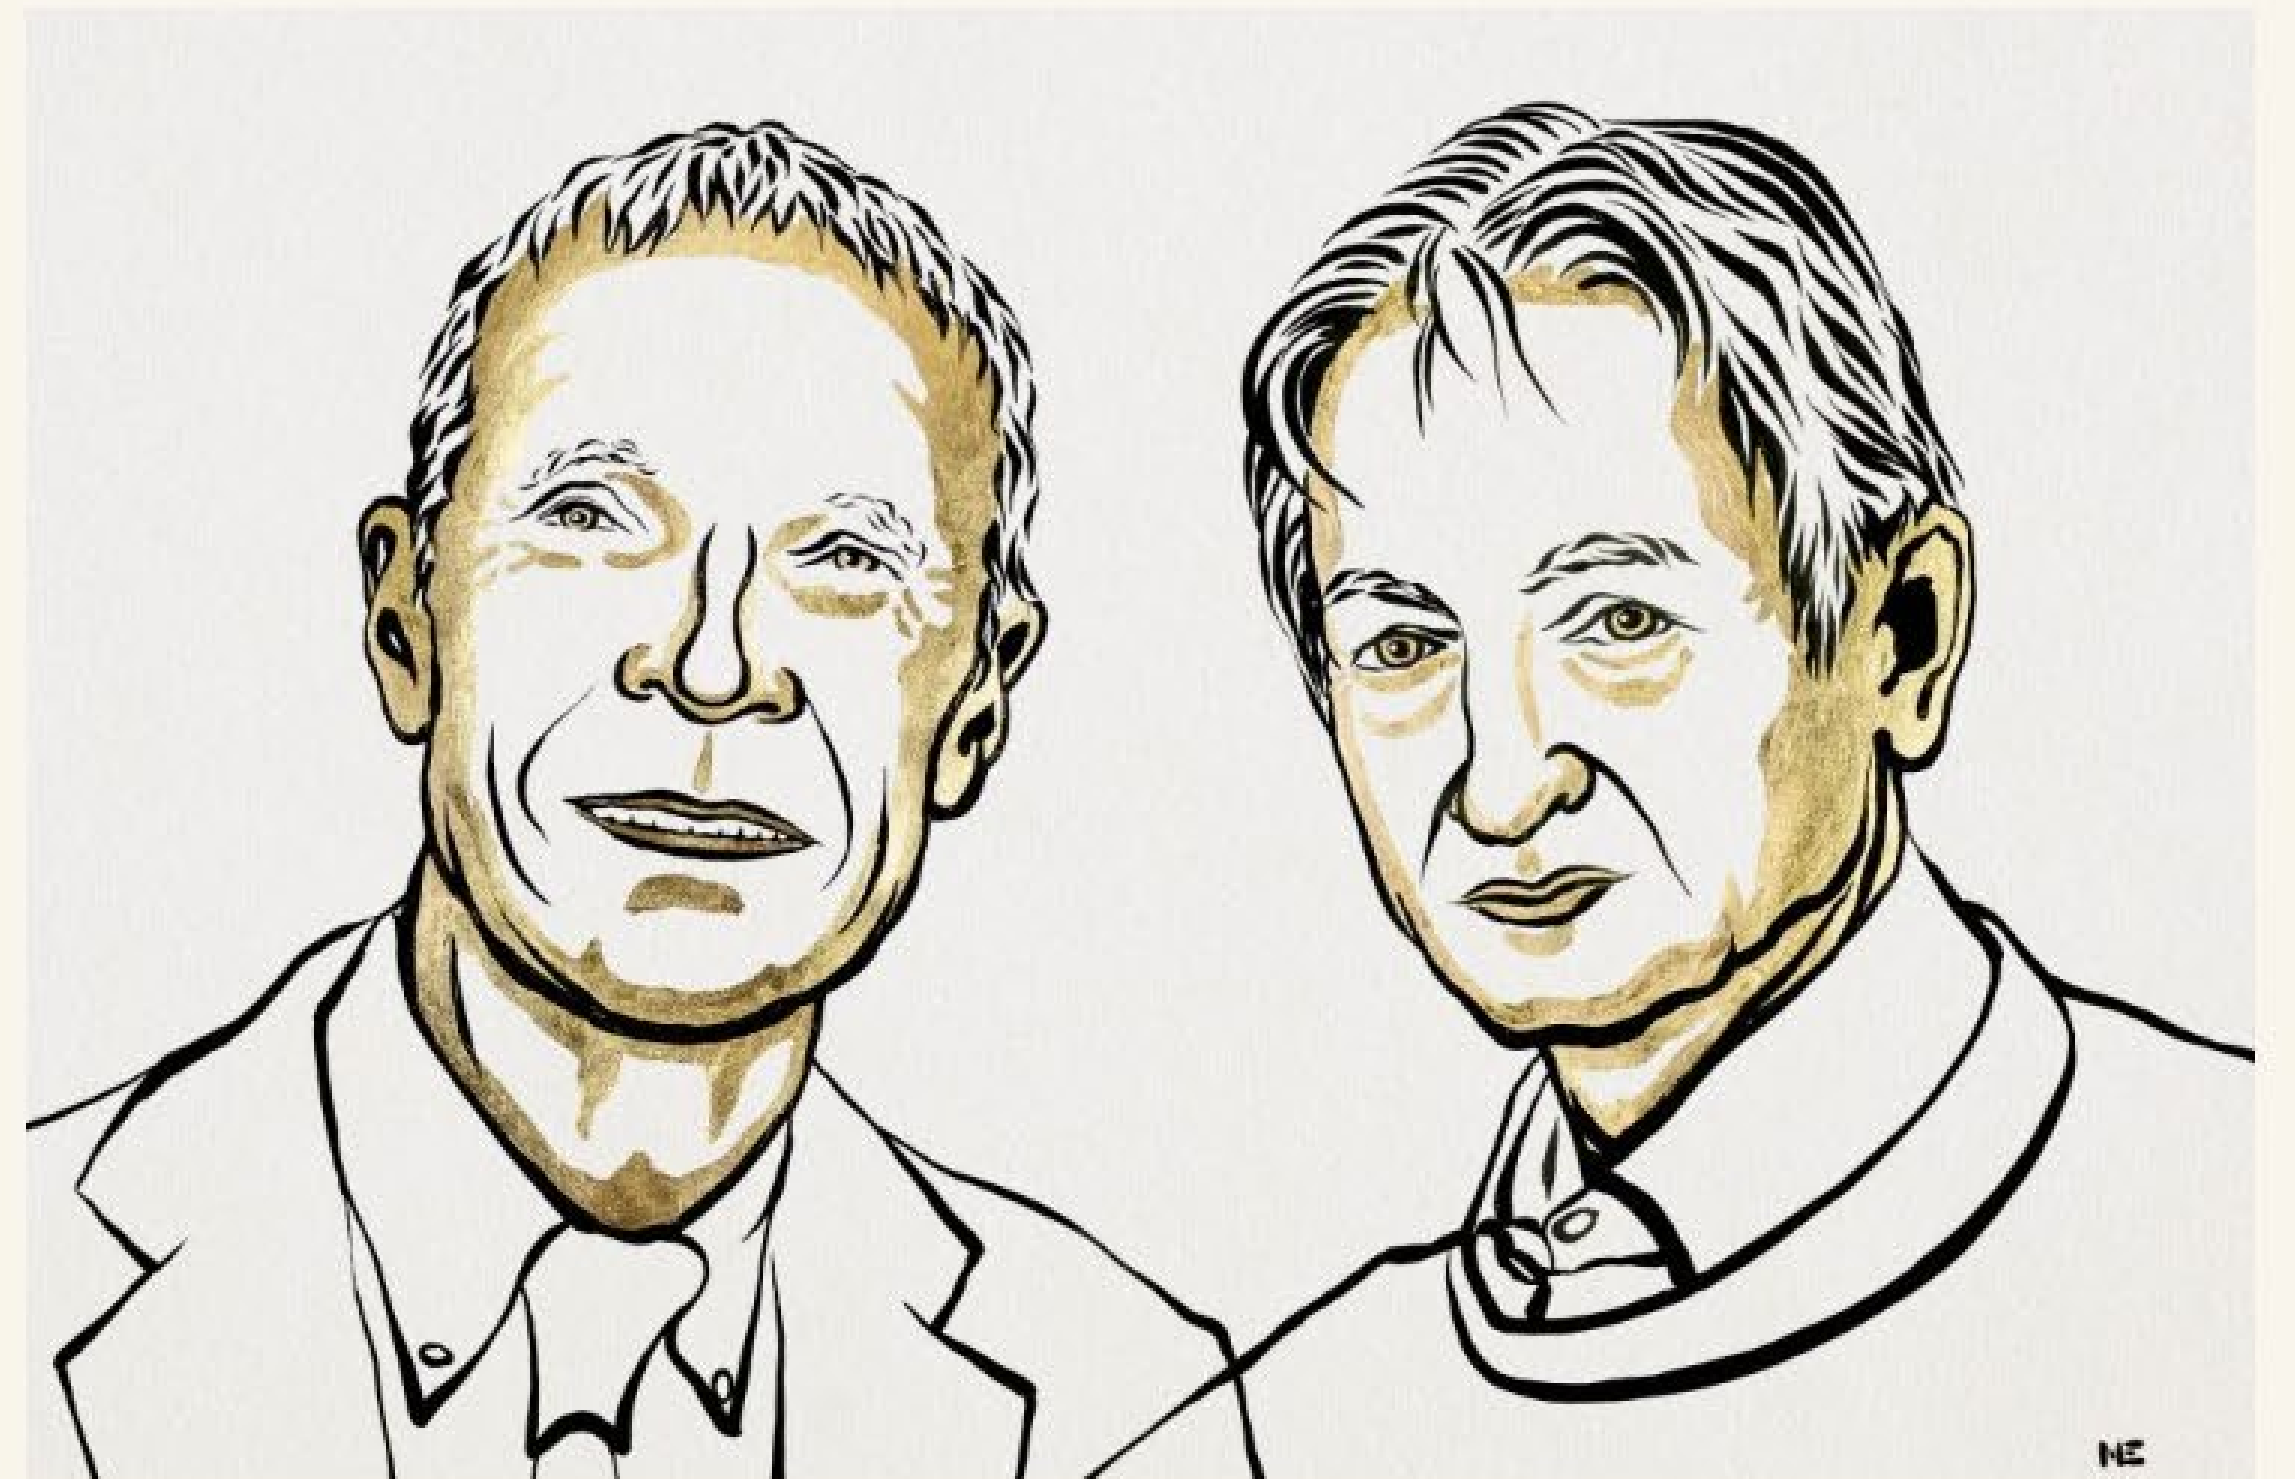
- **NN/ Deep Learning** has been revolutionary for finding structure in data

How “Revolutionary” ?

The 2024 physics laureates

The Nobel Prize in Physics 2024 was awarded to [John J. Hopfield](#) and [Geoffrey Hinton](#) “for foundational discoveries and inventions that enable machine learning with artificial neural networks.”

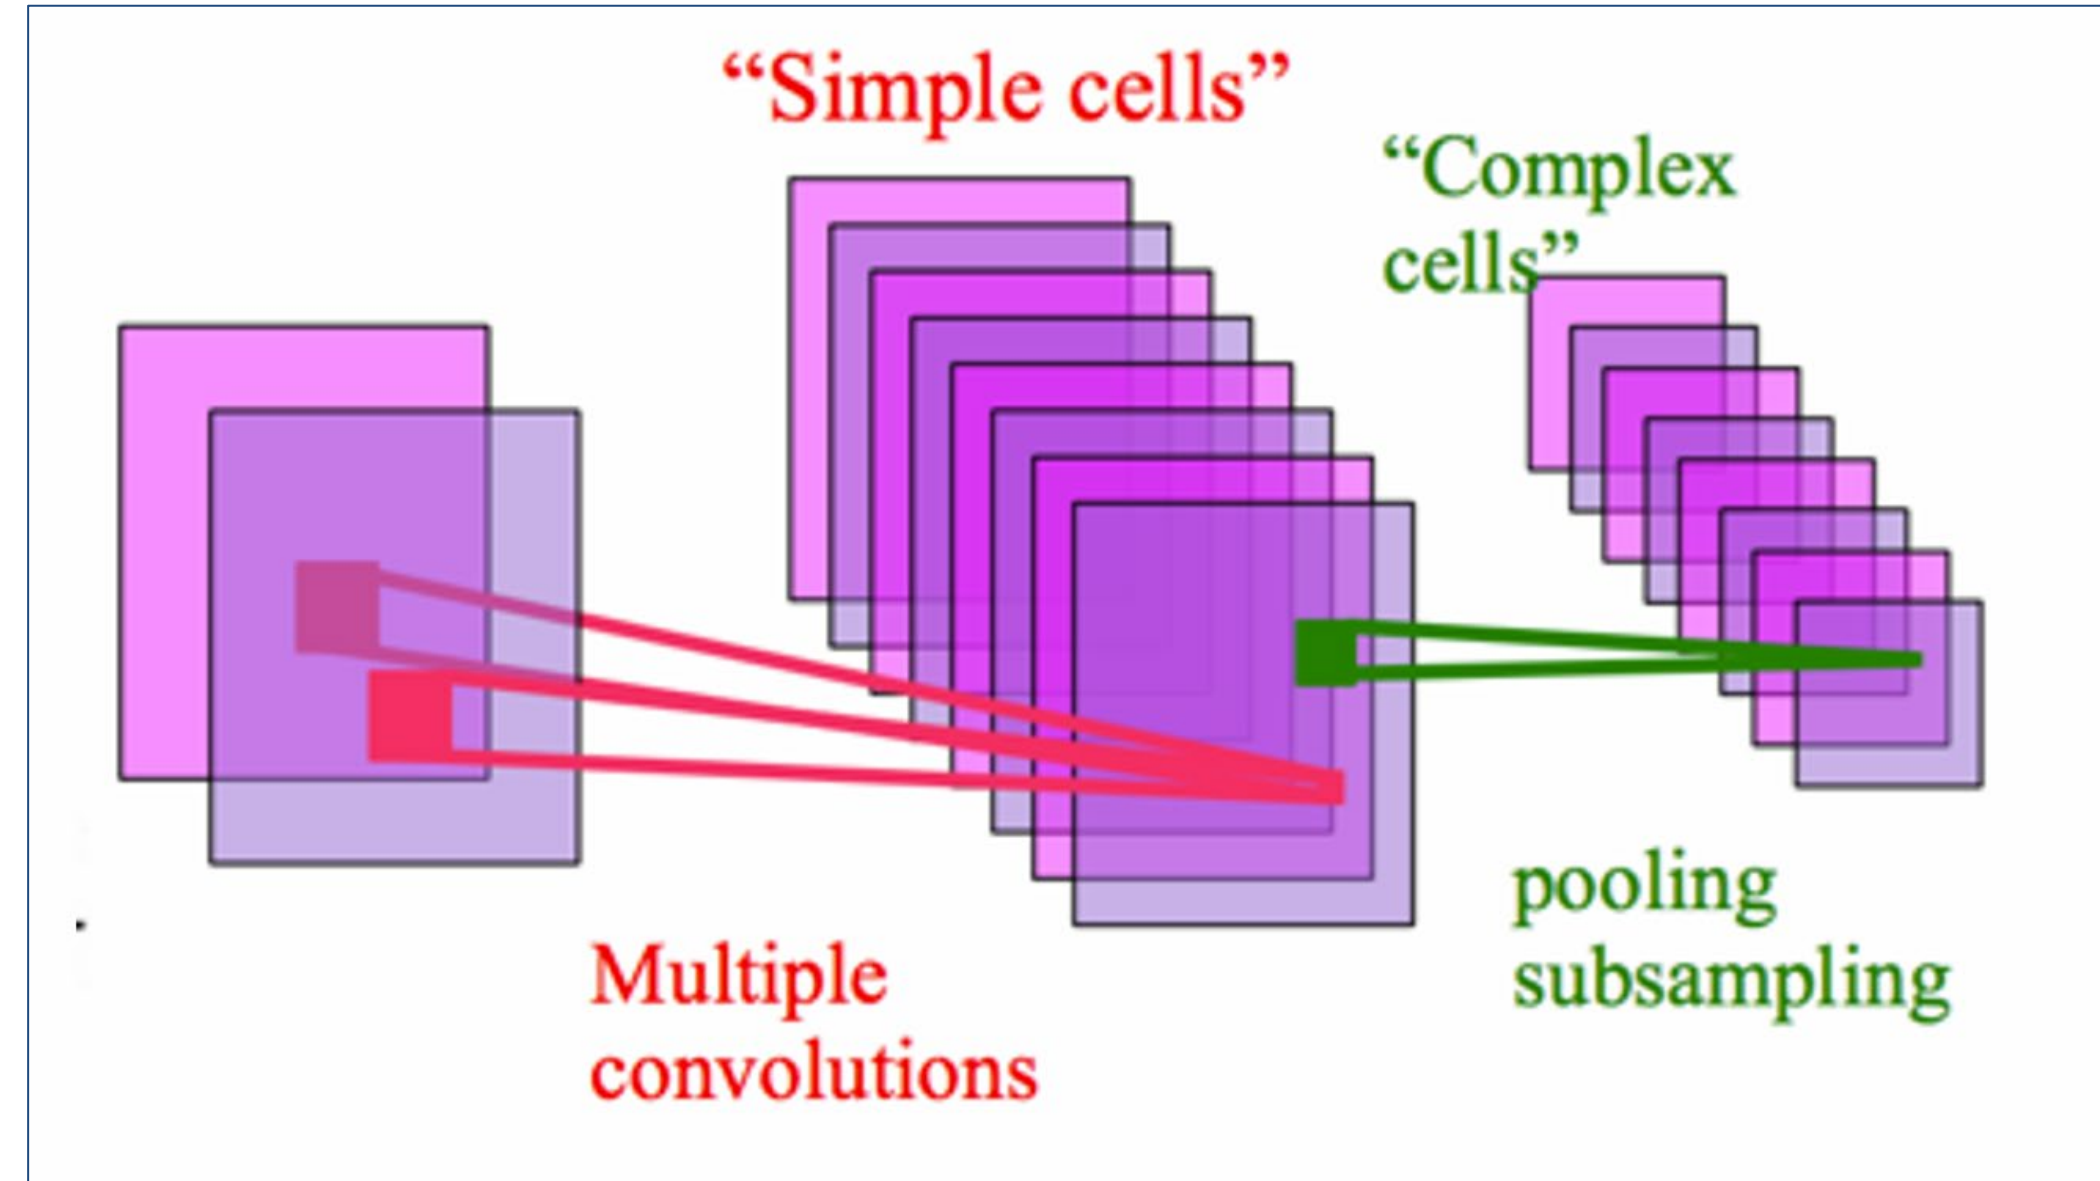
John Hopfield created an associative memory that can store and reconstruct images and other types of patterns in data. Geoffrey Hinton invented a method that can autonomously find properties in data, and so perform tasks such as identifying specific elements in pictures.



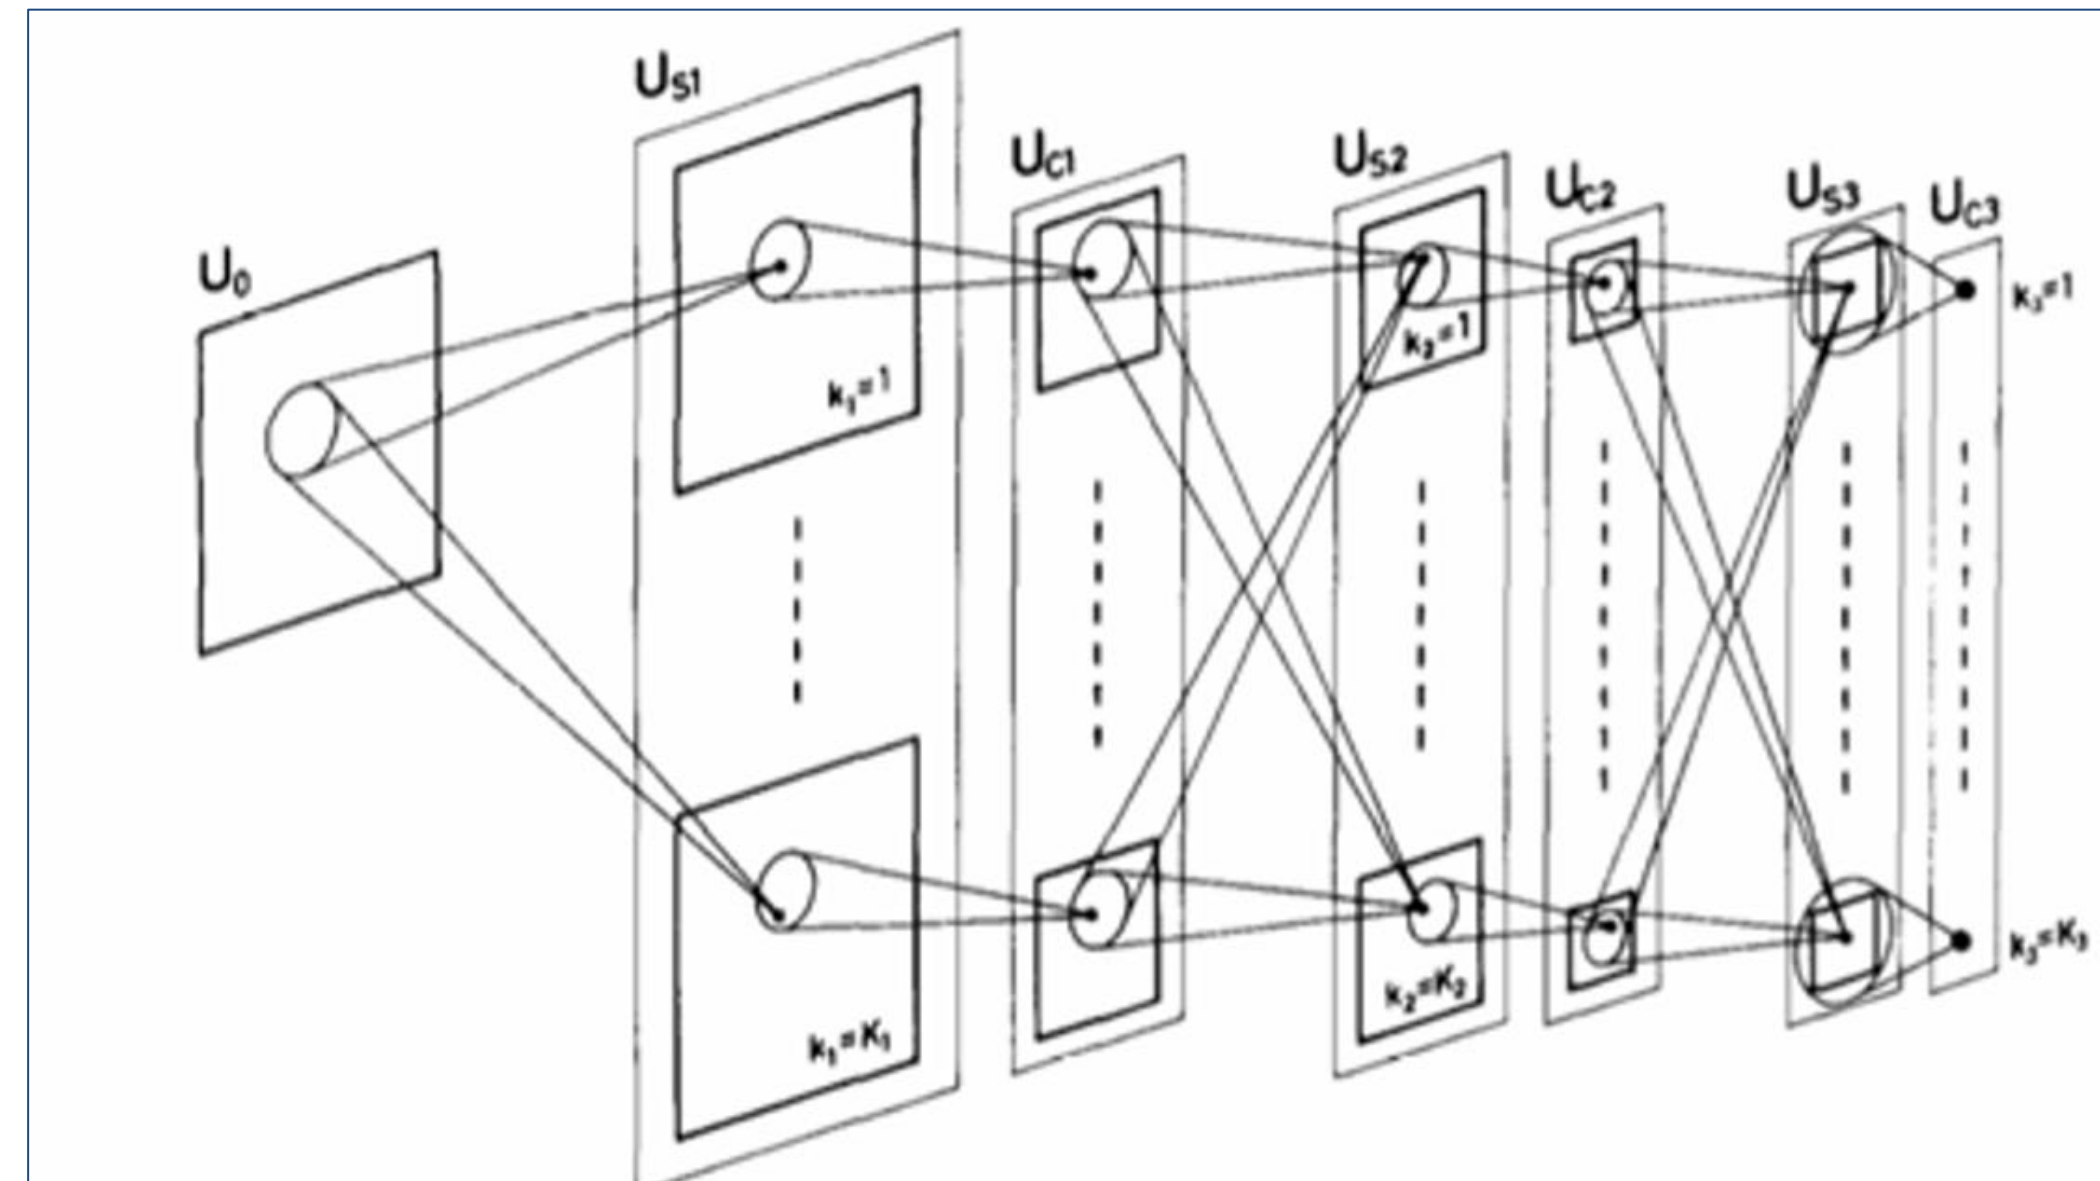
John J. Hopfield and Geoffrey Hinton. Ill. Niklas Elmehed © Nobel Prize Outreach

Early efforts

- Hubel & Wiesel [60s]
 - Simple and Complex cells architecture



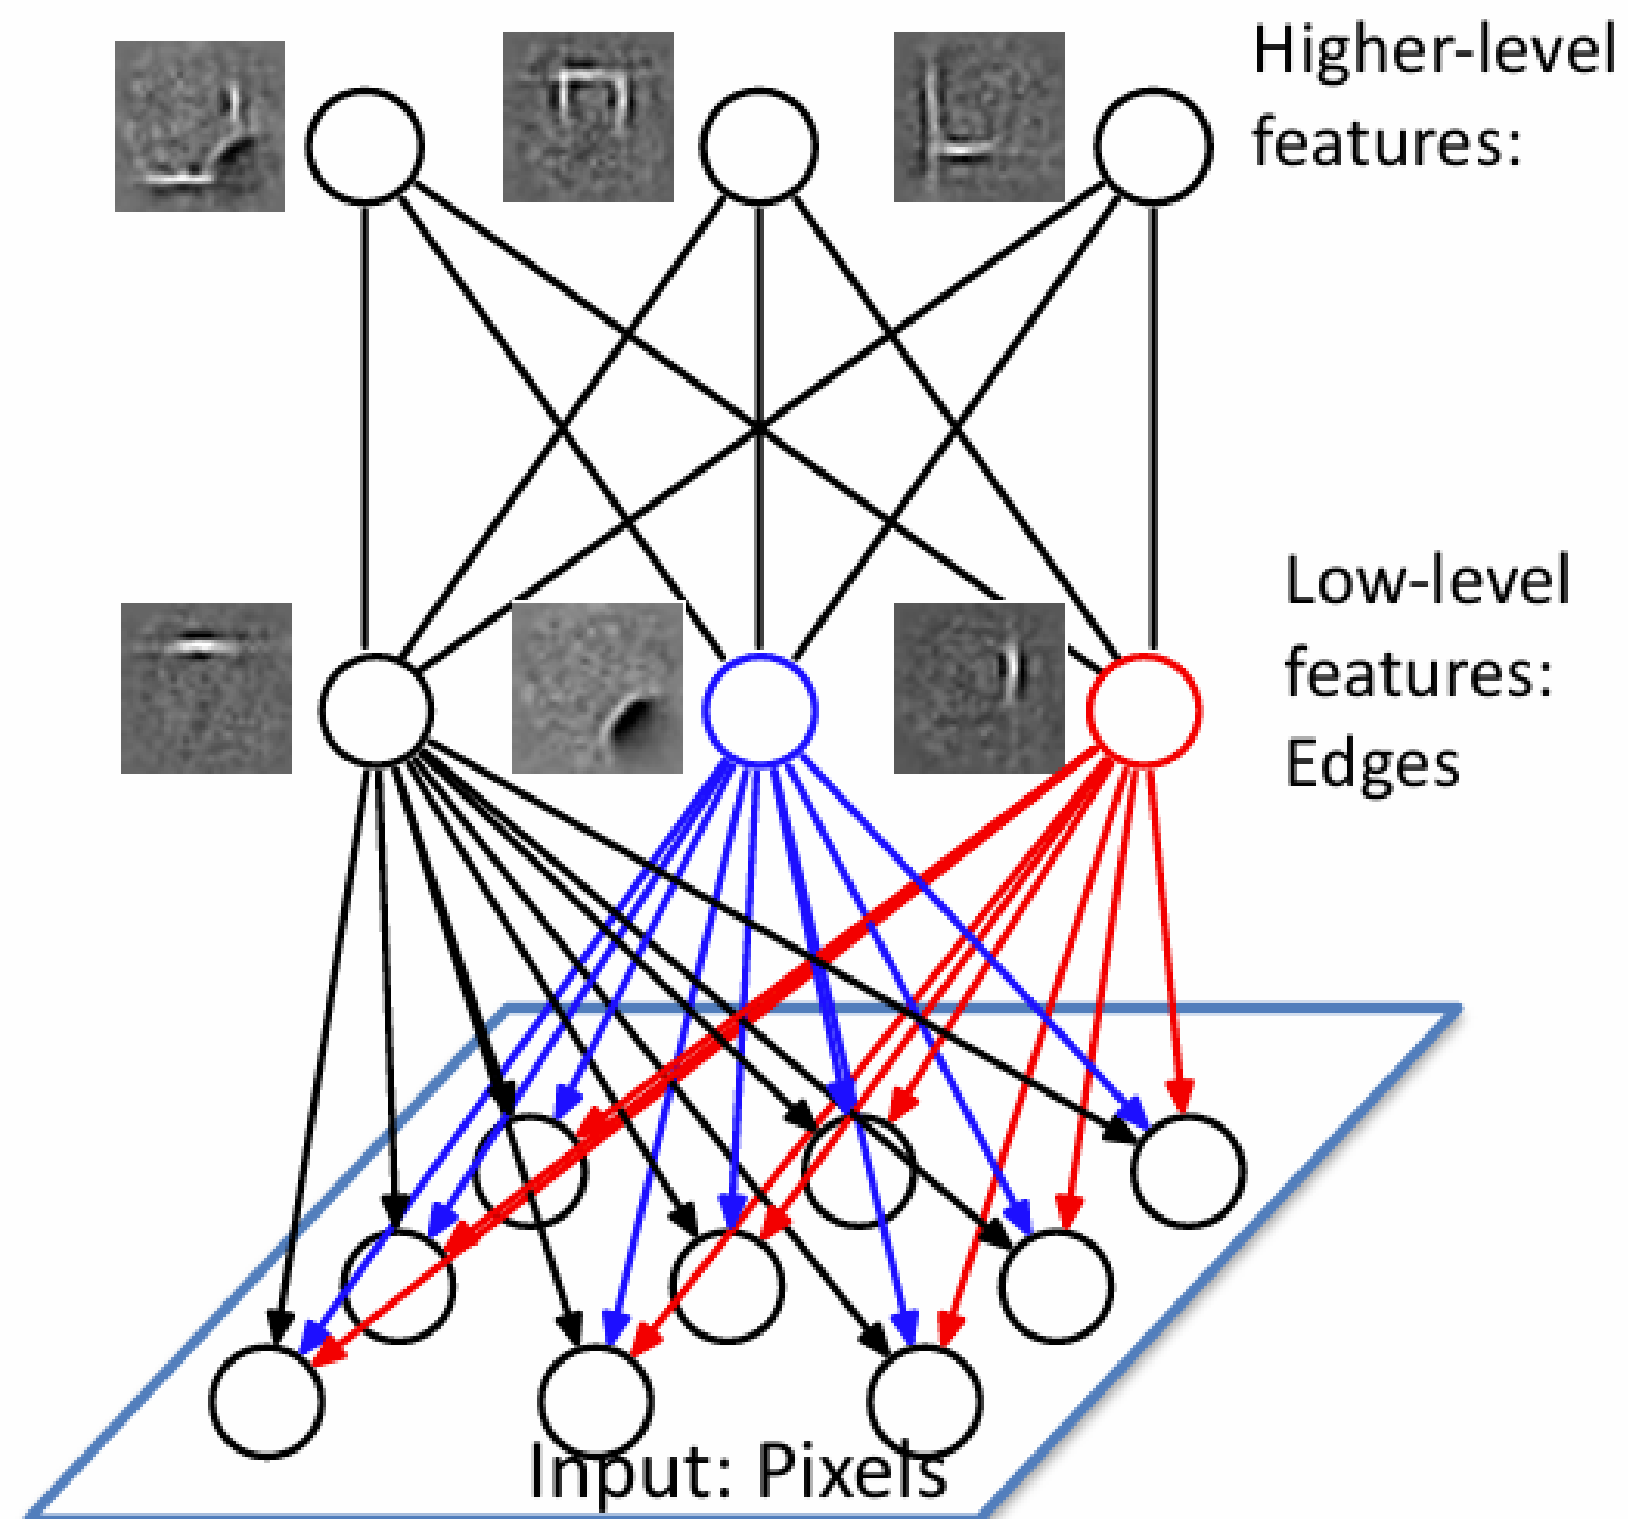
- Fukushima [70s] Neocognitron



Early Breakthroughs

- **Deep Belief Networks, 2006 (Unsupervised)**

Hinton, G. E., Osindero, S. and Teh, Y., A fast learning algorithm for deep belief nets, Neural Computation, 2006.



Theoretical Breakthrough:

- Adding additional layers improves variational lower-bound.

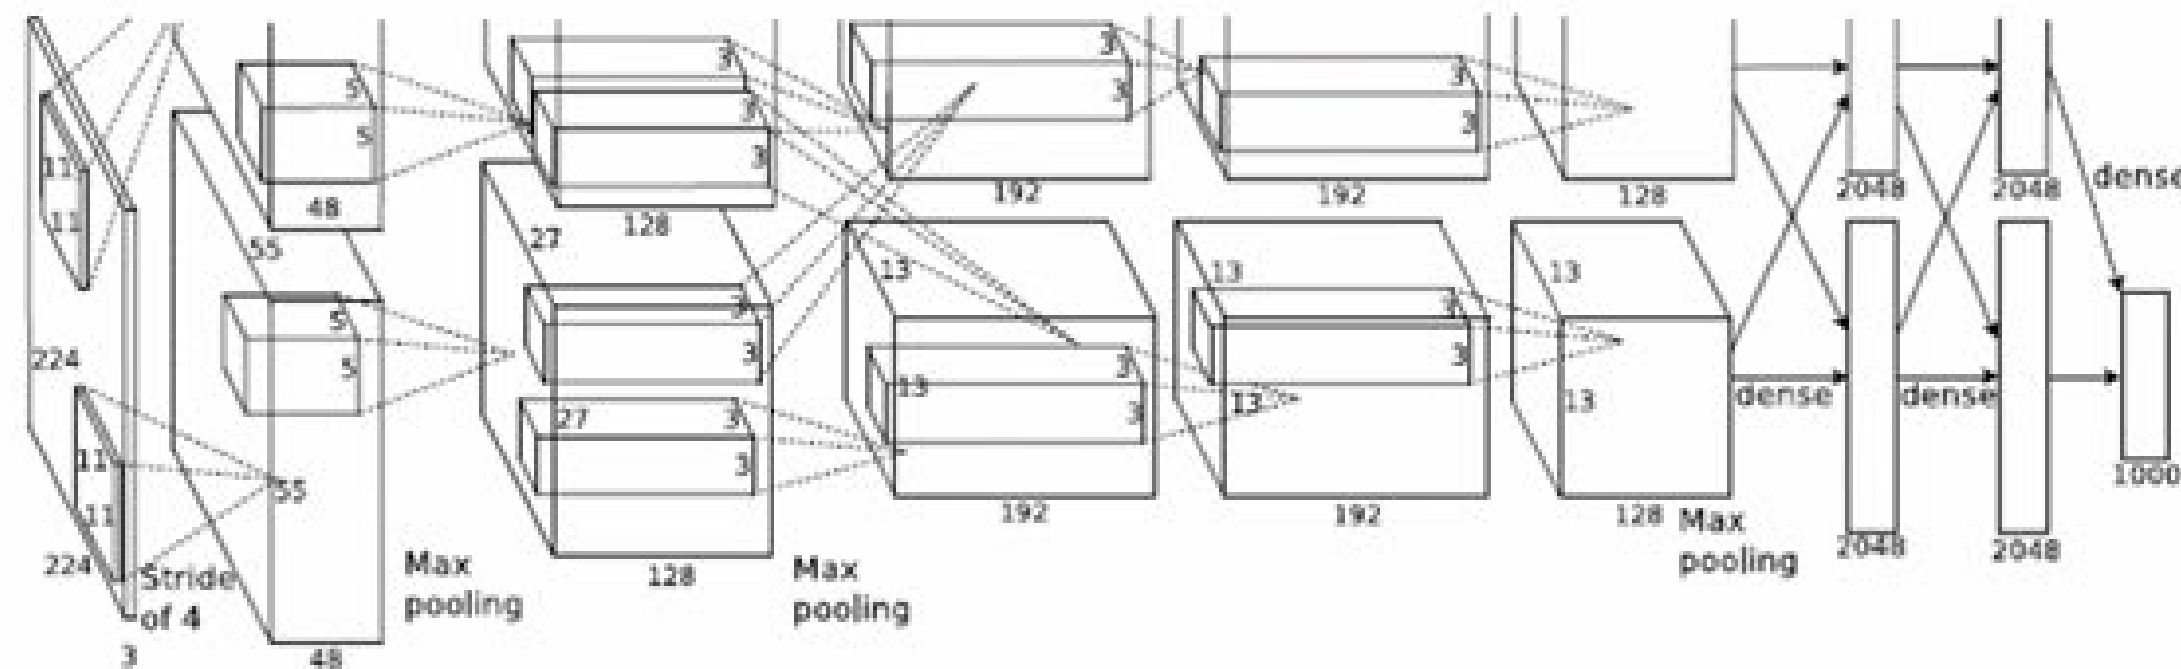
Efficient Learning and Inference with multiple layers:

- Efficient greedy layer-by-layer learning algorithm.
- Inferring the states of the hidden variables in the top most layer is easy.

Breakthroughs: Convolutional Neural Networks

- Convolutional Networks for Vision

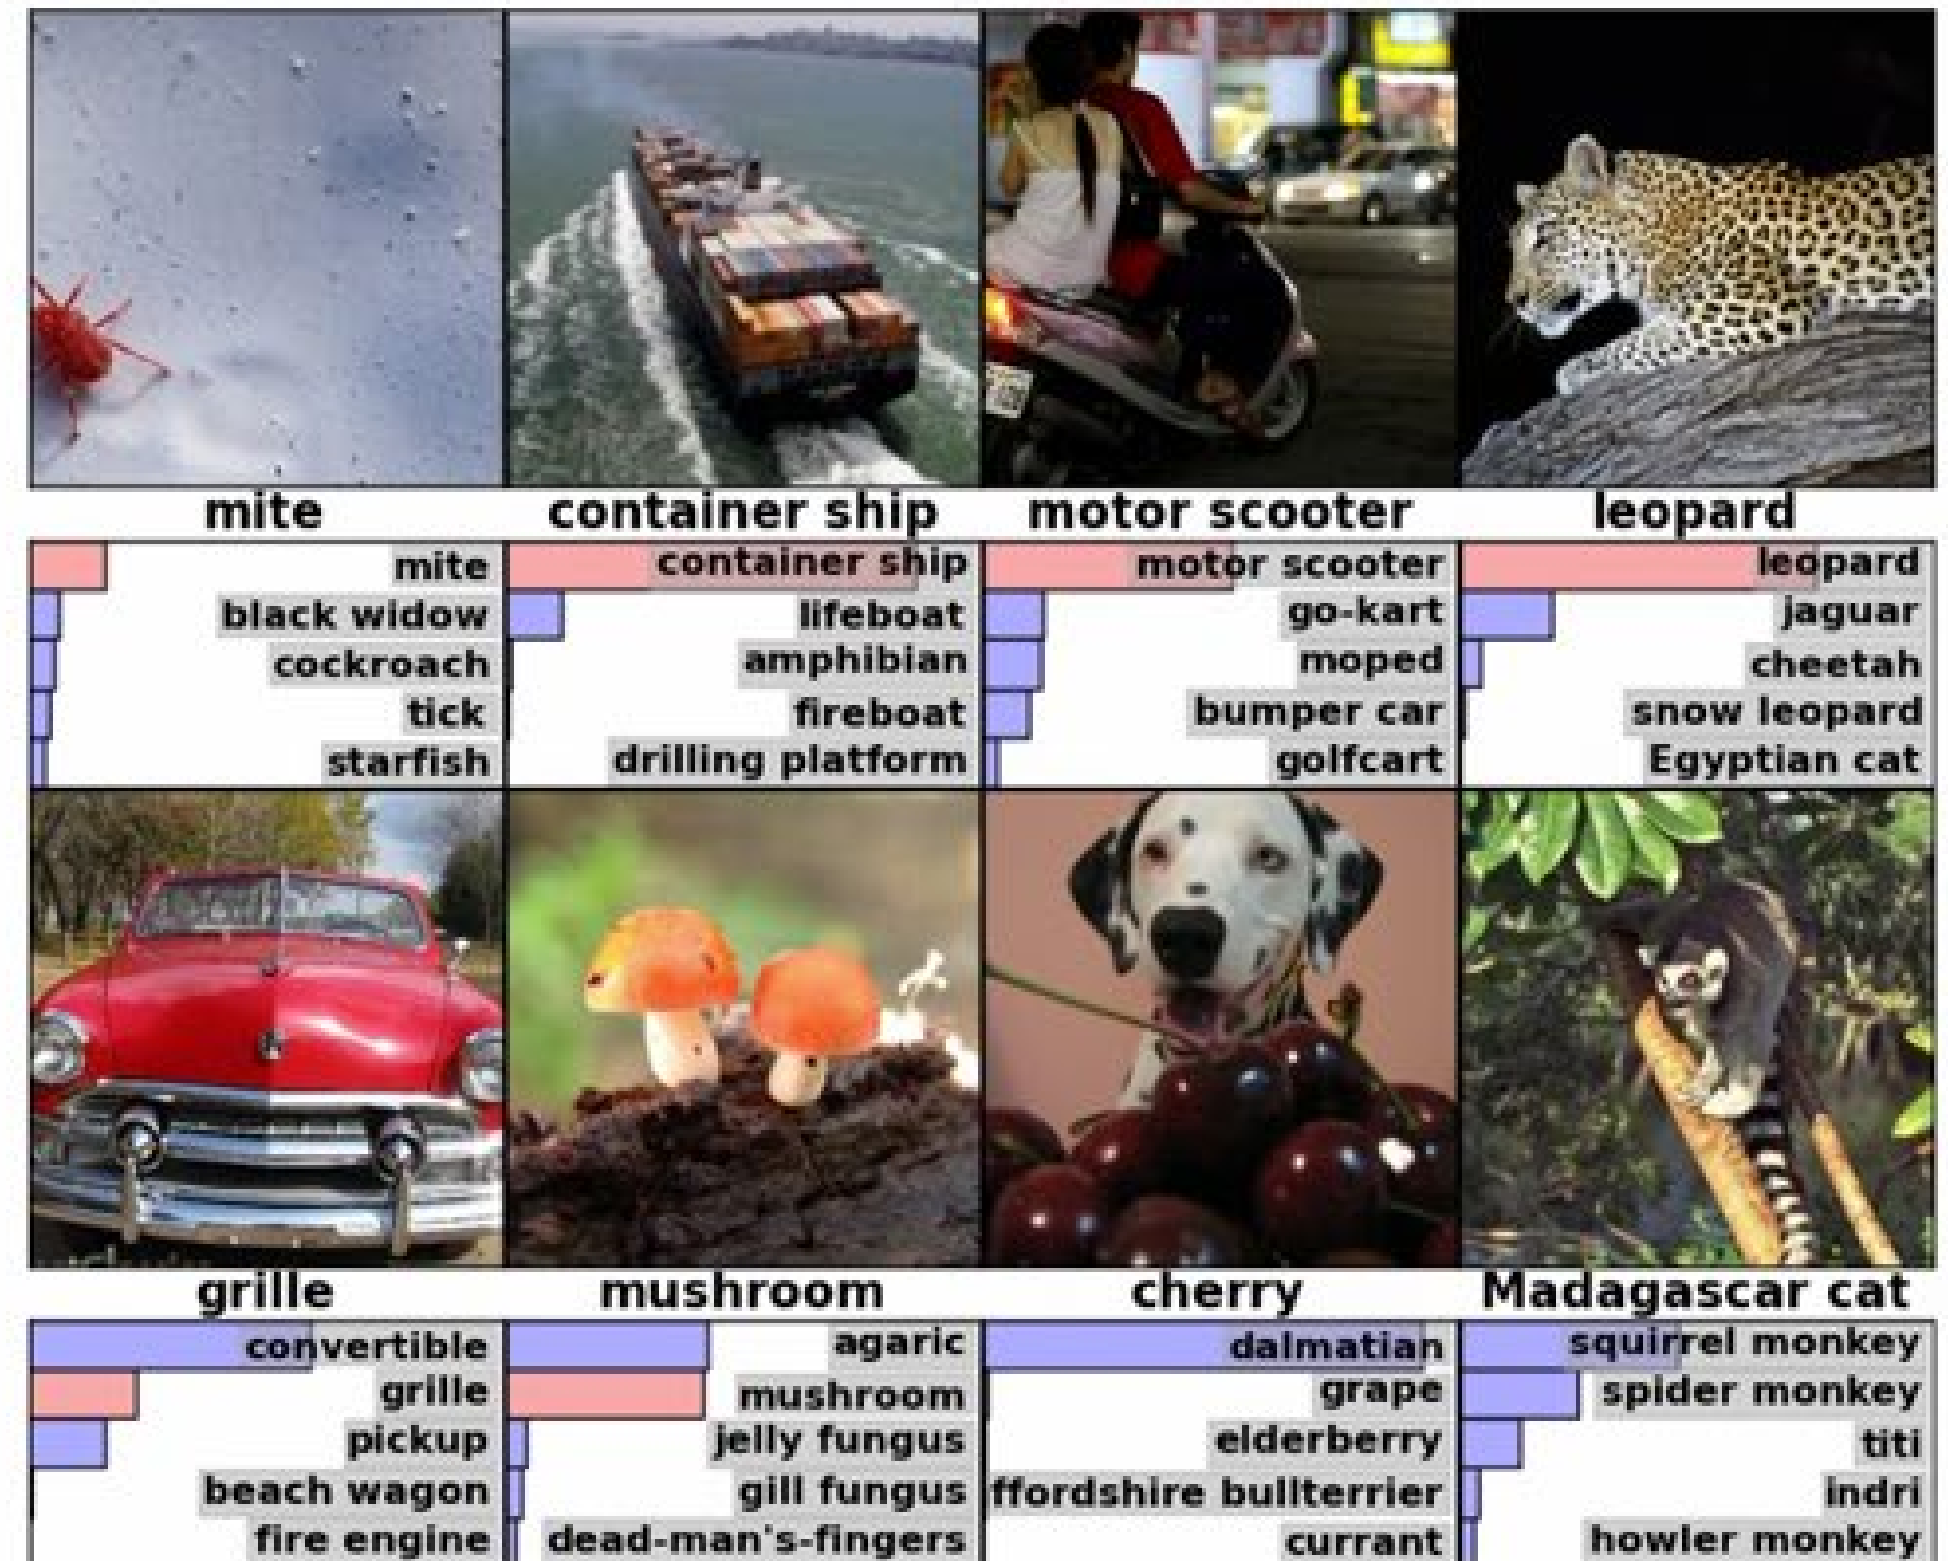
- Krizhevsky, A., Sutskever, I. and Hinton, G. E., ImageNet Classification with Deep Convolutional Neural Networks, NIPS, 2012.



IMAGENET

1.2 million training images

1000 classes

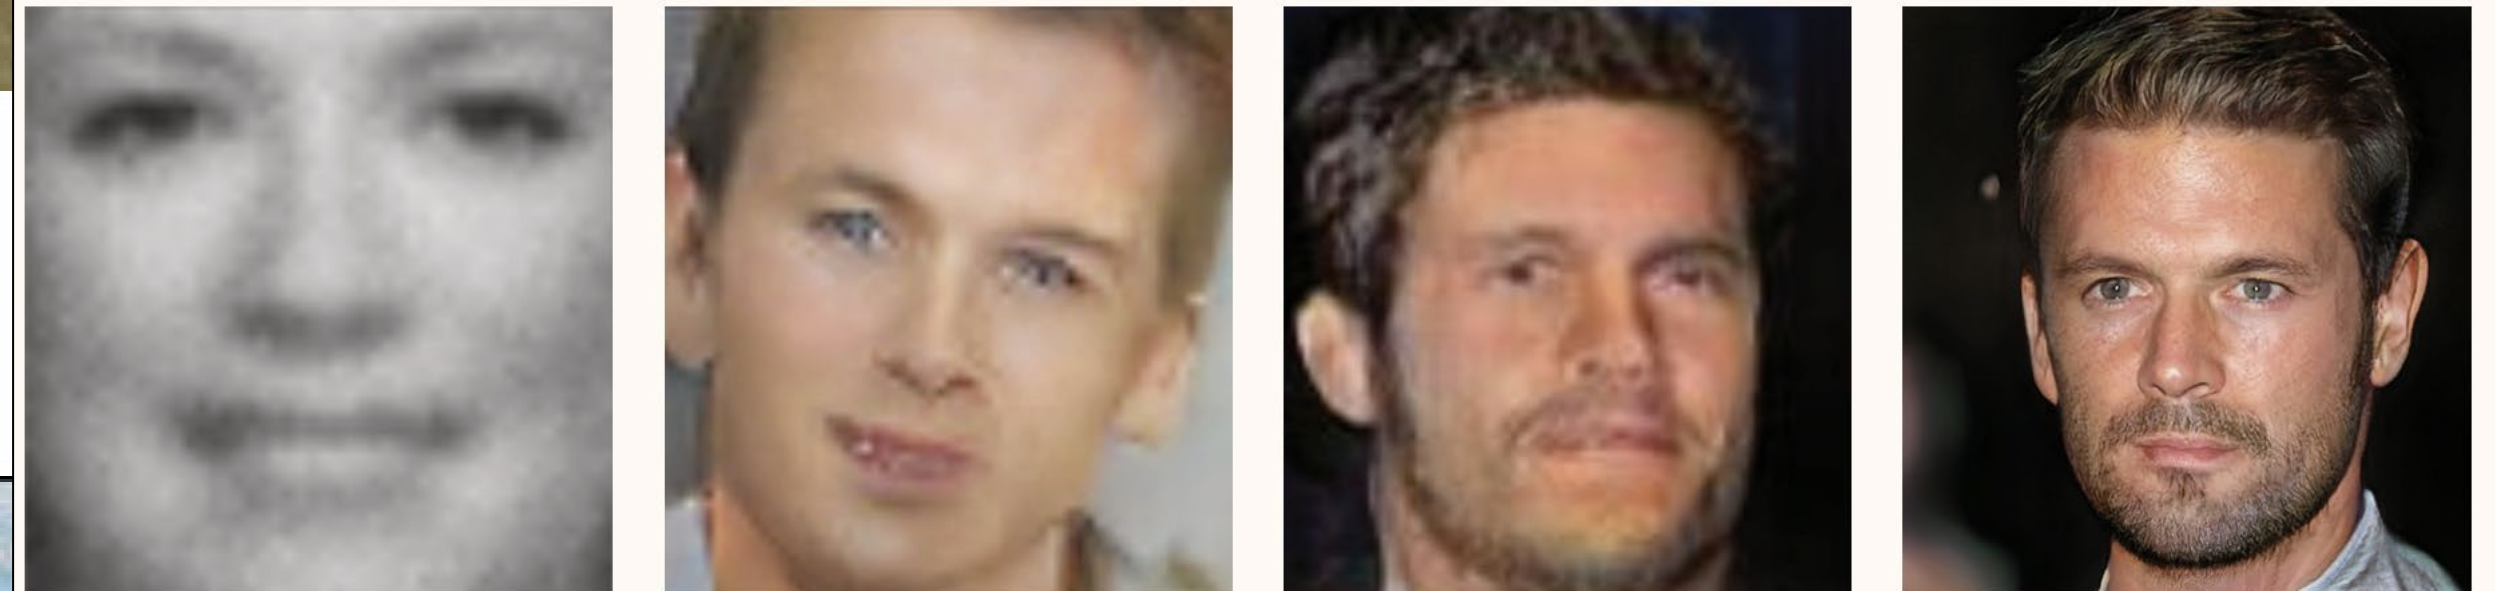


Breakthroughs: Generative Modeling

- Conditional generative model $P(\text{zebra images} | \text{horse images})$



4 years of progression on Faces



2014

2015

2016

2017

► Style Transfer



Input Image

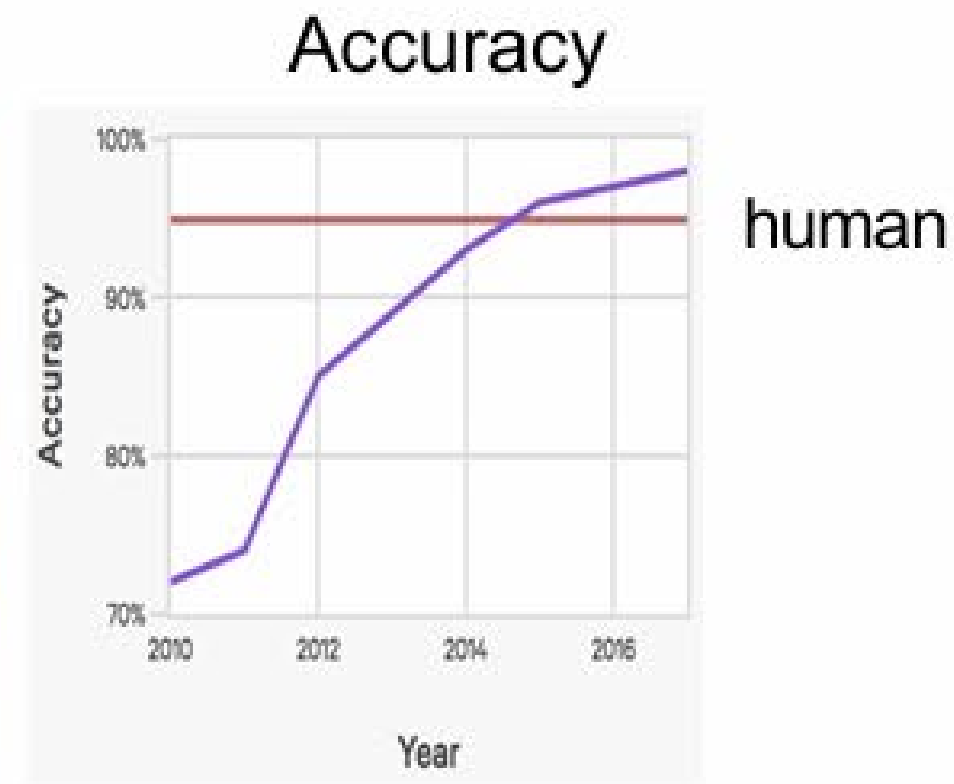
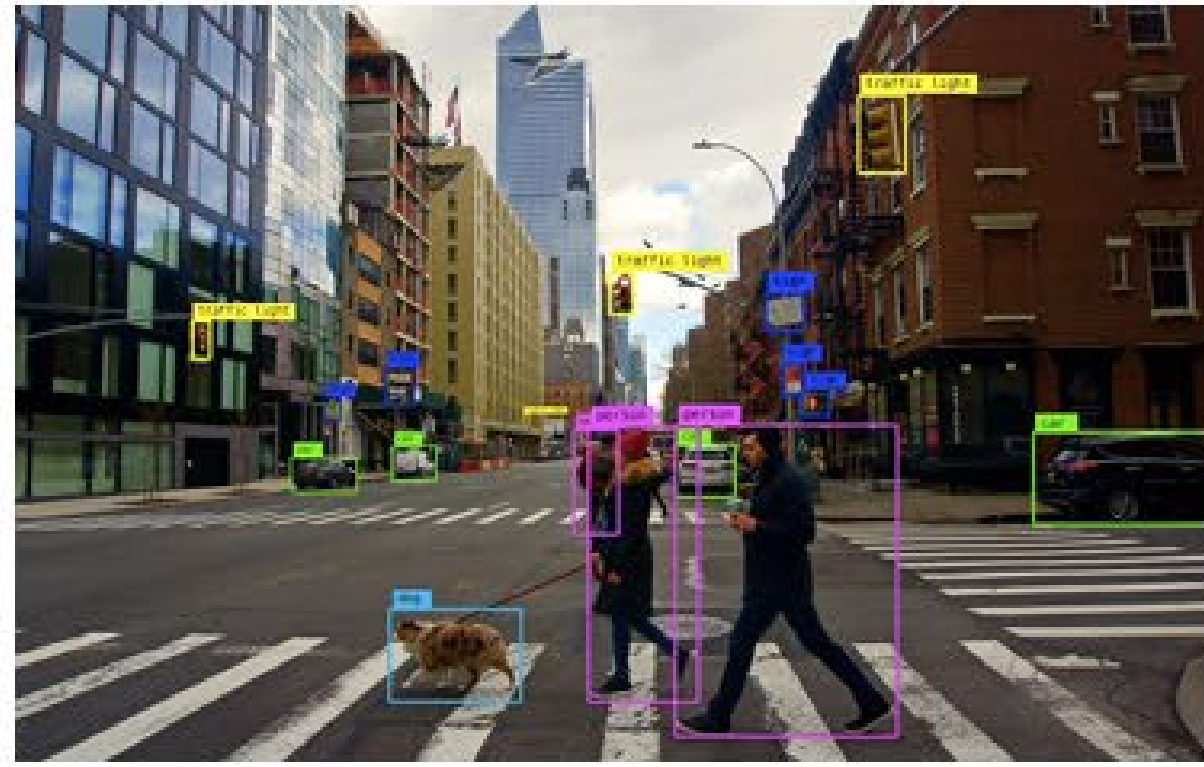


Monet

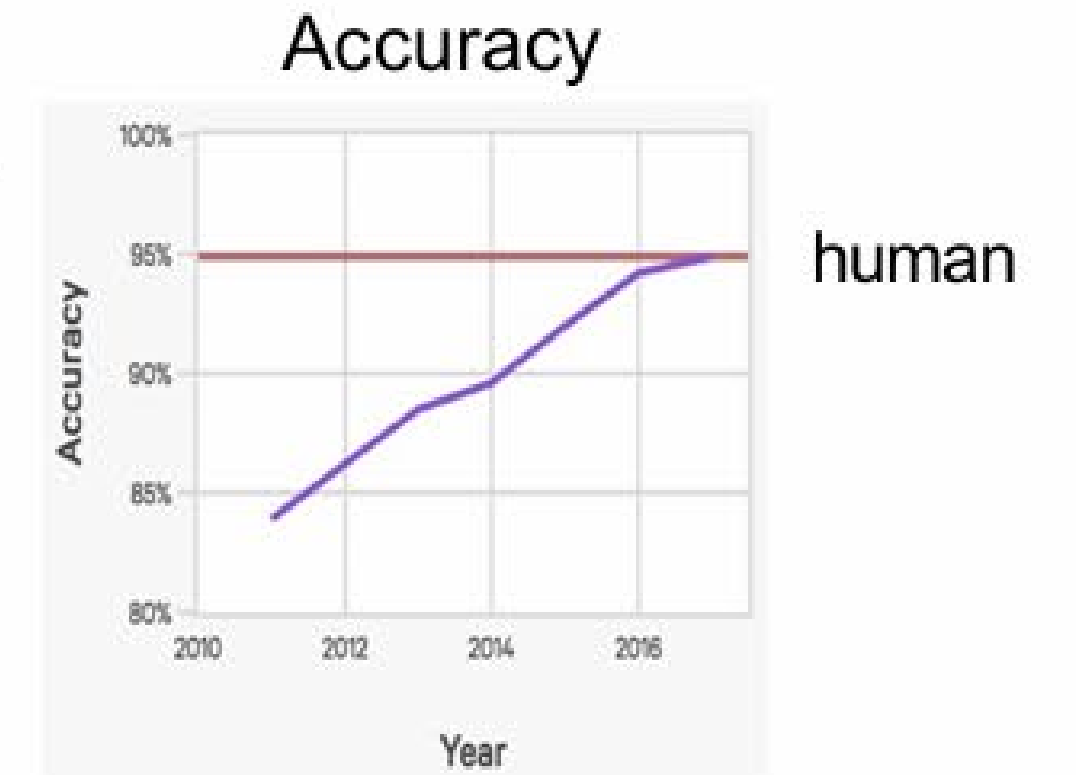
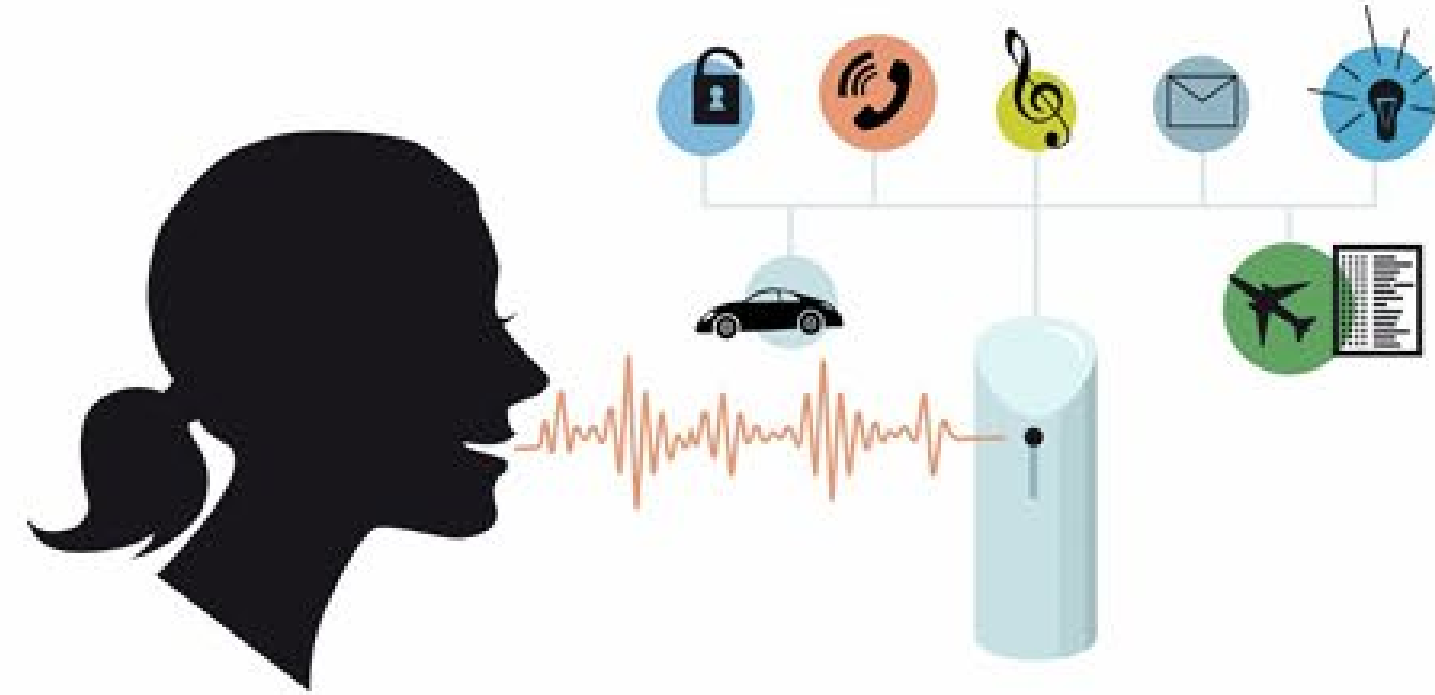


Van Gogh

Computer Vision



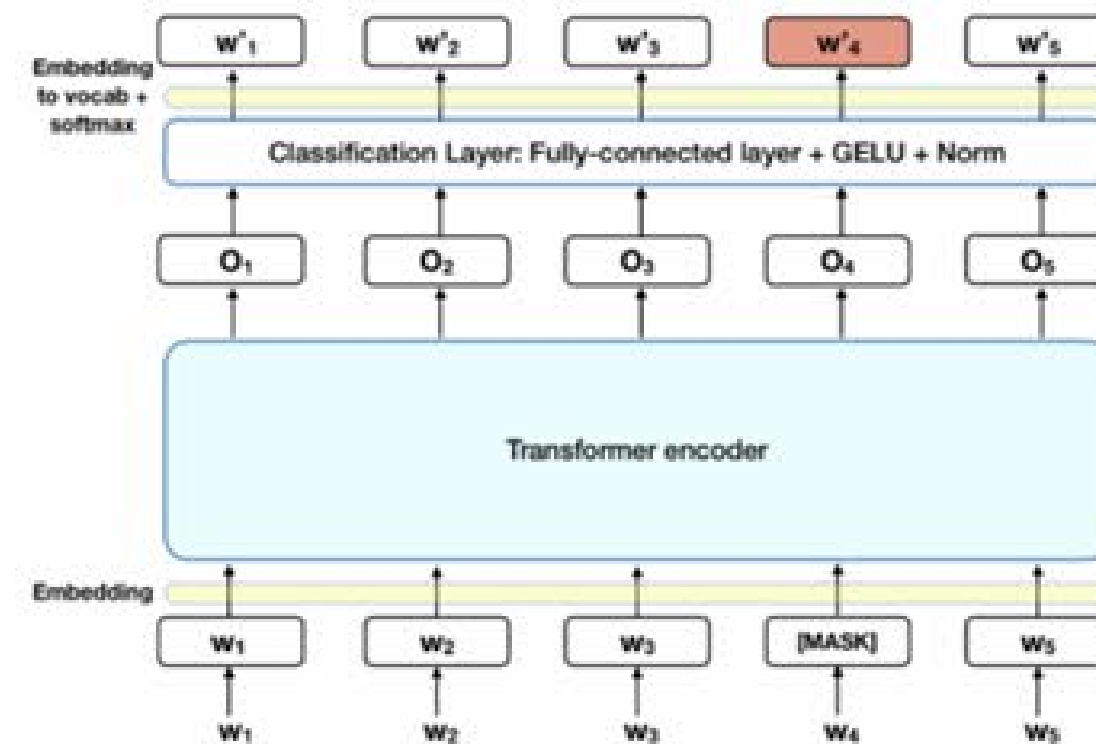
Speech Recognition



Reasoning & Planning



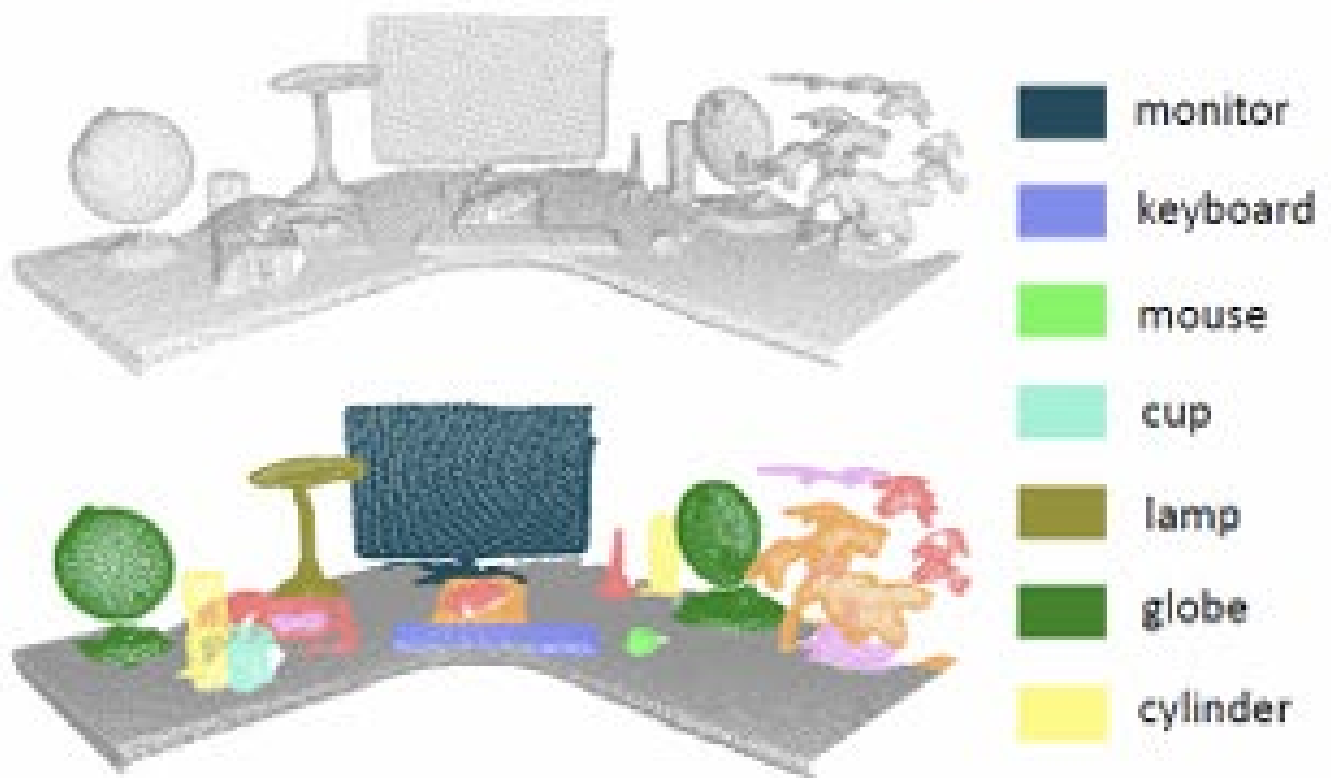
Language Understanding



Robotics



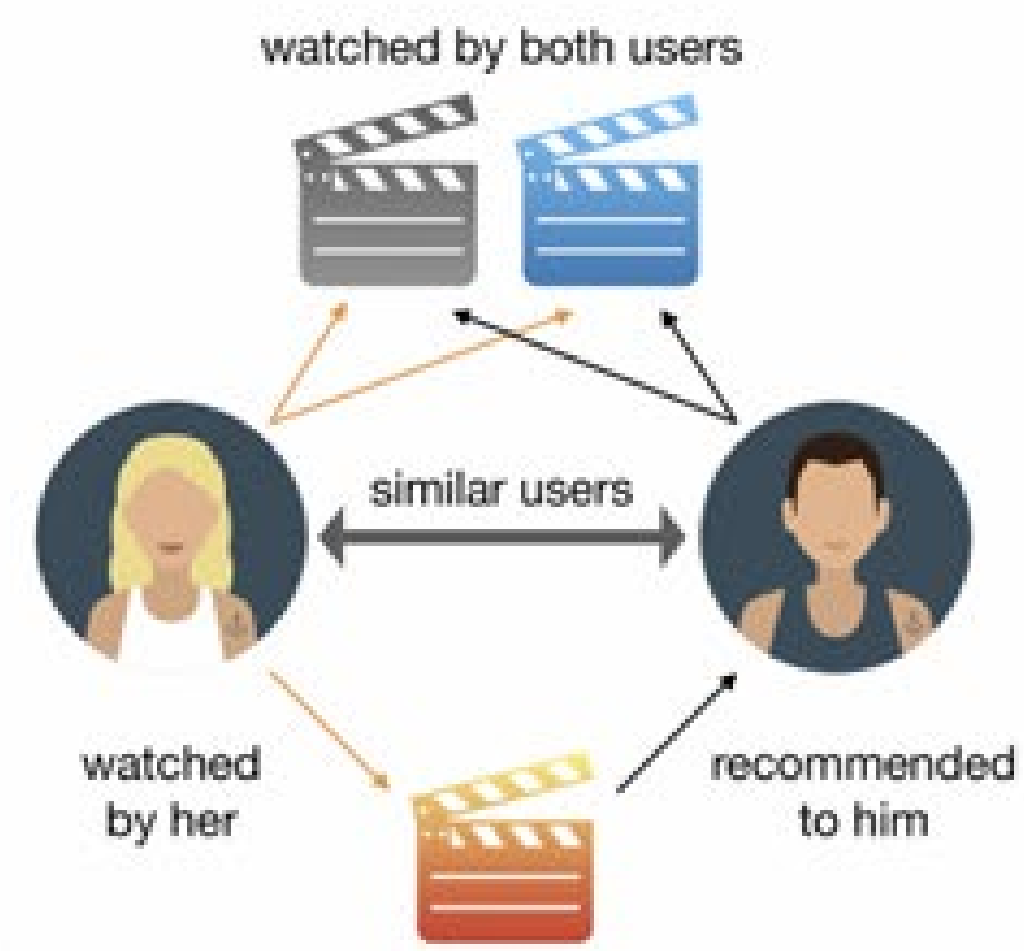
Perception and 3D Scene Understanding



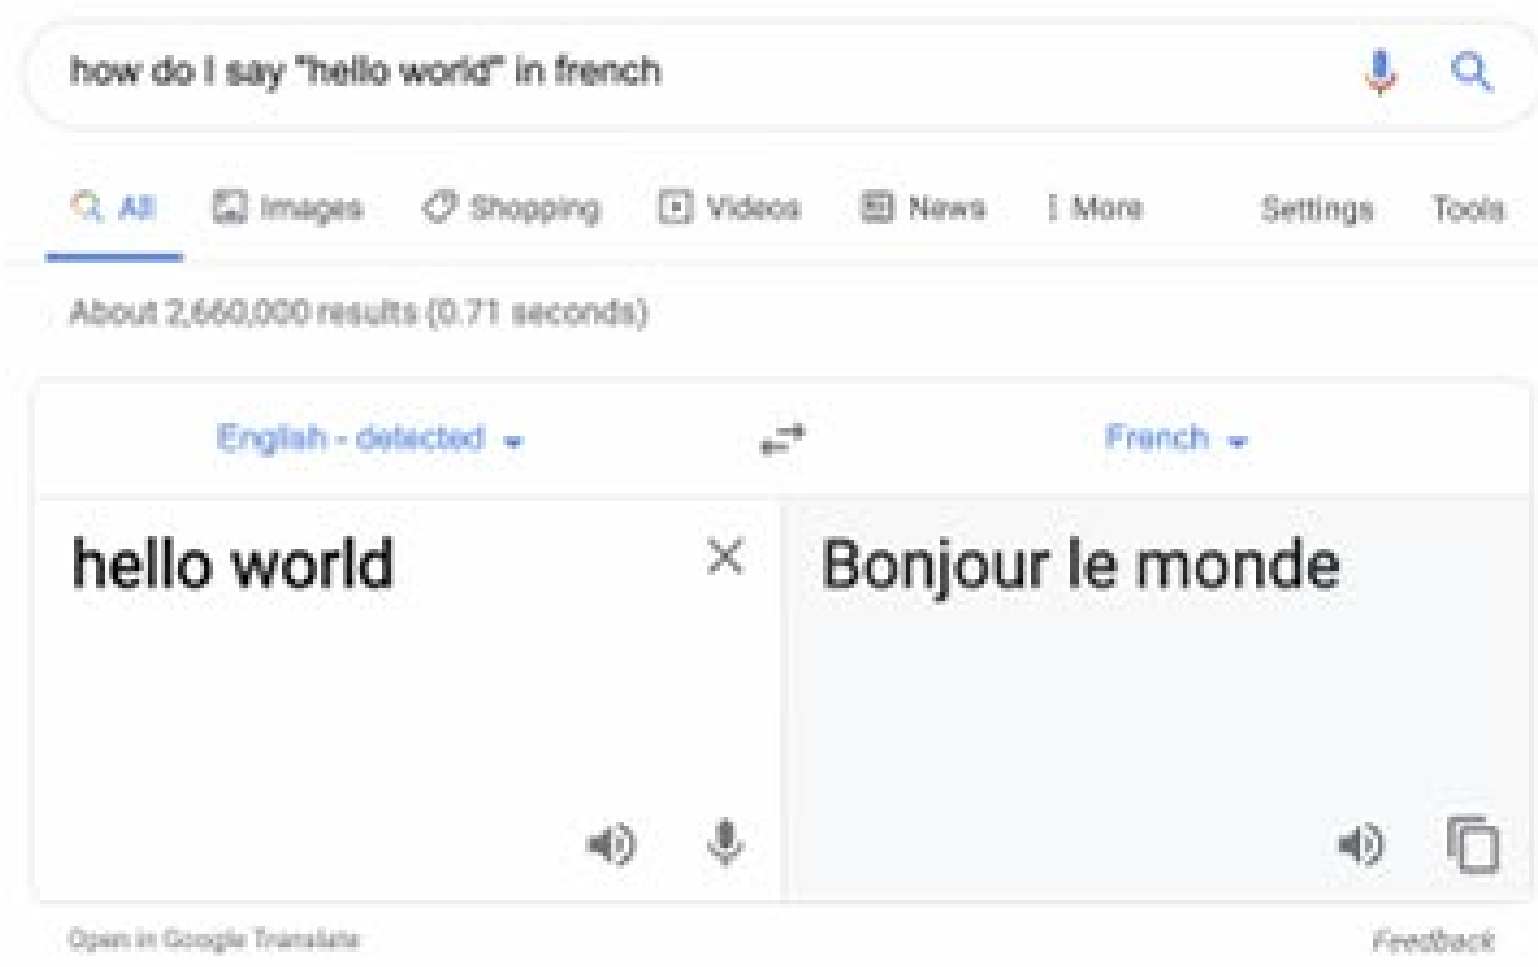
Conversational Agents



Recommendation Systems



Machine Translation



Medical Image Analysis

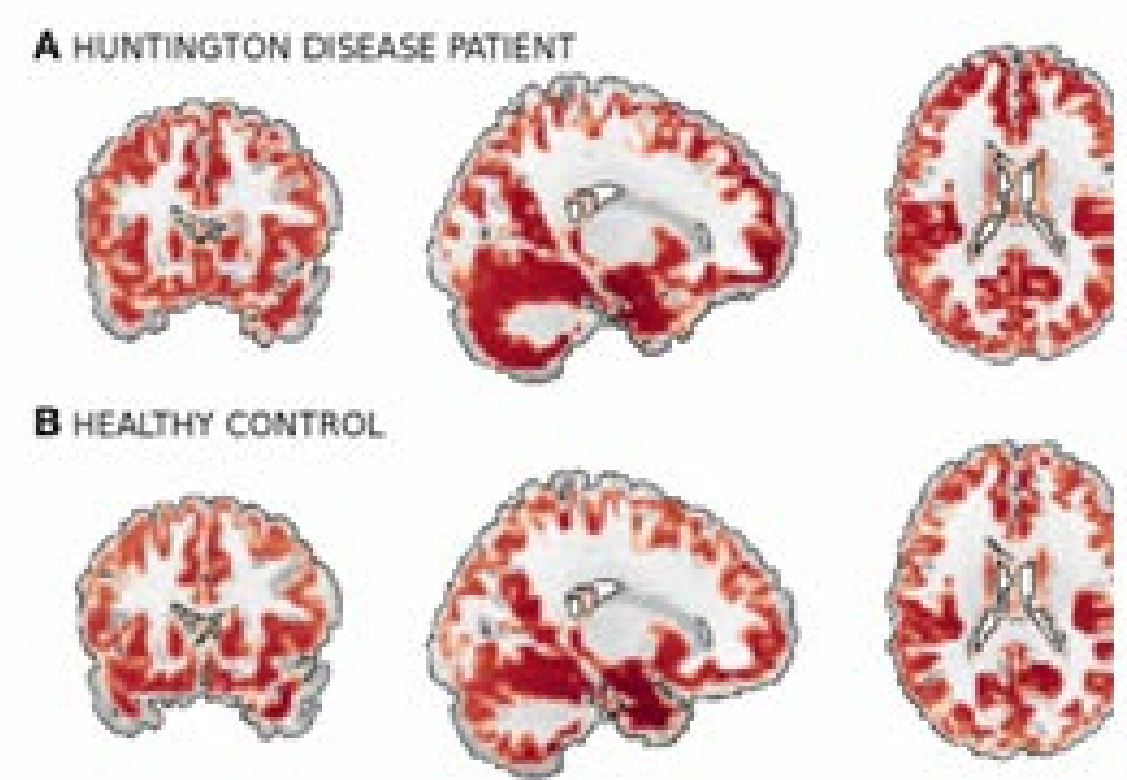


FIGURE 7 | A gray matter of MRI scans of an HD patient (A) and a healthy control (B).

Self Driving Cars



Language Models

Human: Hey, how are you doing?

AI: I'm good! What would you like to chat about?

Human: Have you done any coding with Python?

AI: Yep! I've done a lot of Python coding as it is my favorite language.

Human: Nice, it's my favorite too! What have you done with it?

AI: Oh, a lot. I have done everything from creating Python modules to creating micro-websites with it.

Human: What web framework do you use?

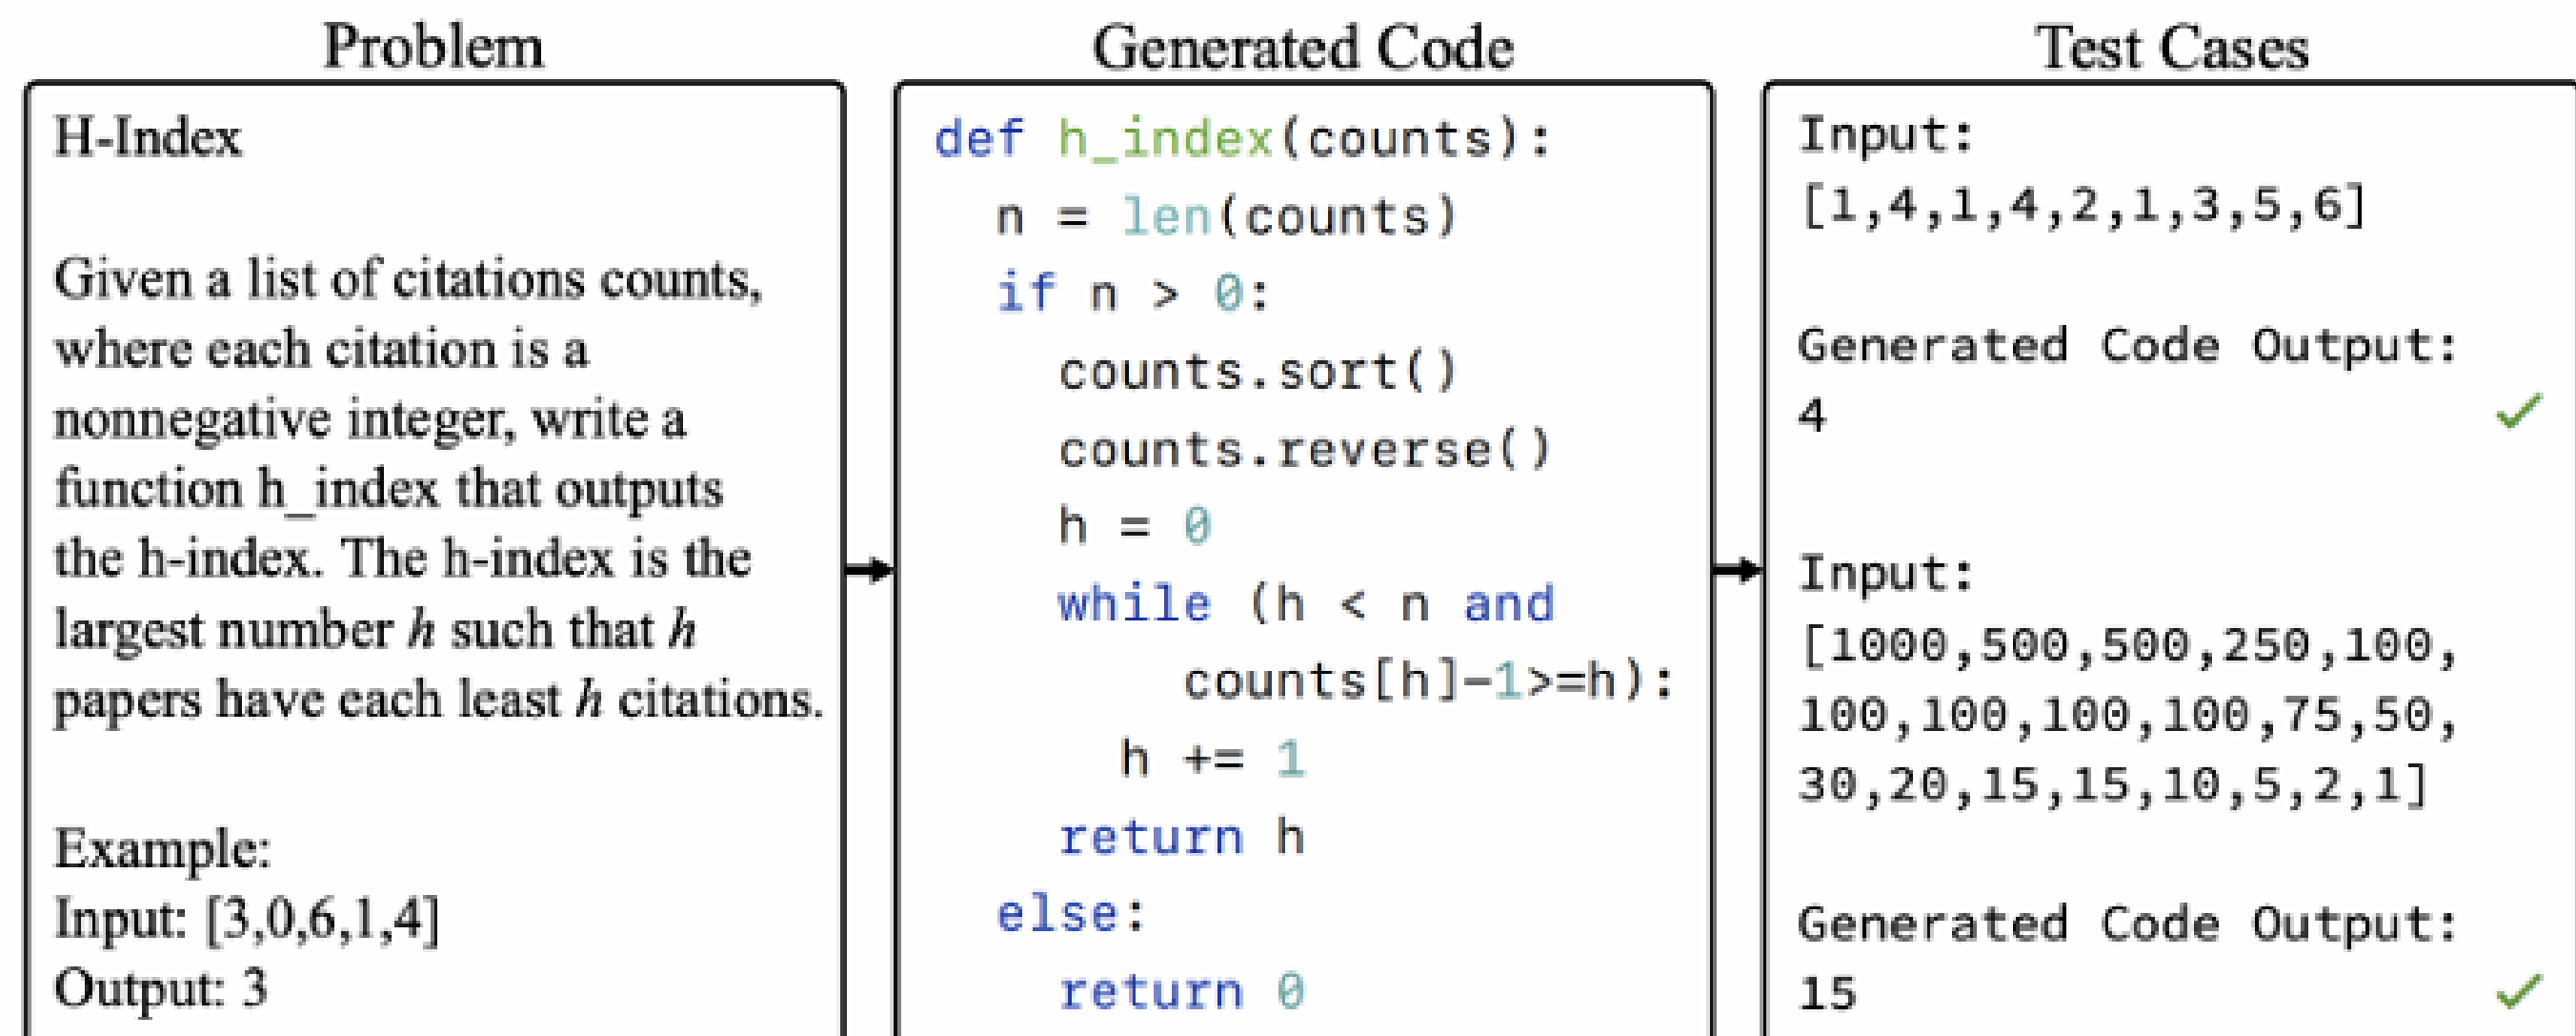
AI: I used to use Django, but it was too difficult for me. I have found Flask is a better option as it's lighter and simpler.

Human: Wow, I agree 100%. Have you used WebSocket at all?

AI: Yes, I actually have written a chatbot using it!

Human:

- Chatbots
- Question-Answering Systems
- Information Retrieval
- Code Generation



Diffusion Models

Text-to-Image Synthesis on LAION. 1.4B Model.

'A zombie in the style of Picasso'

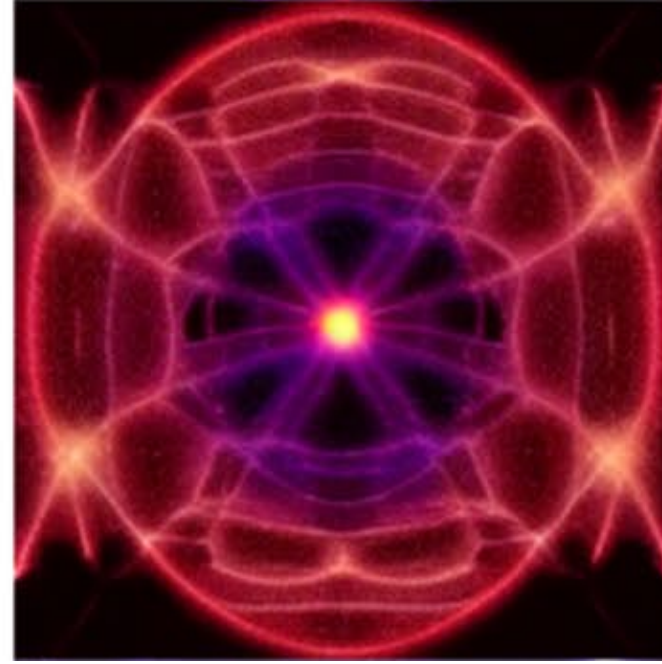
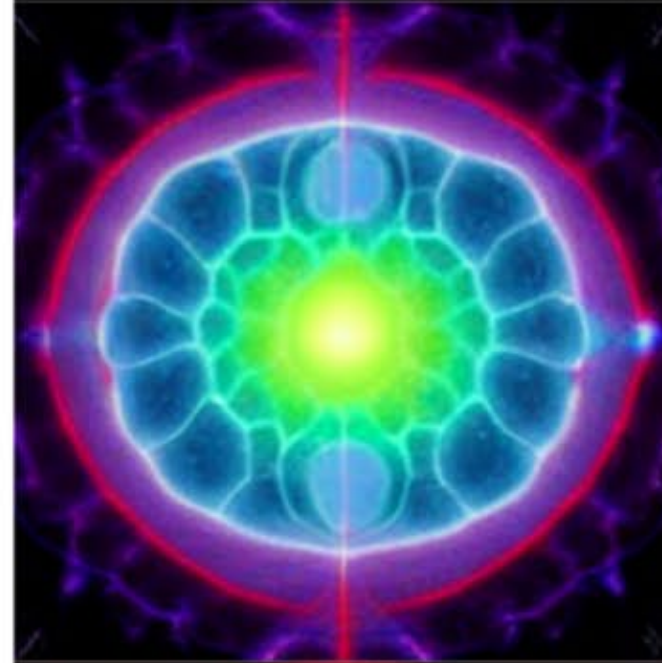
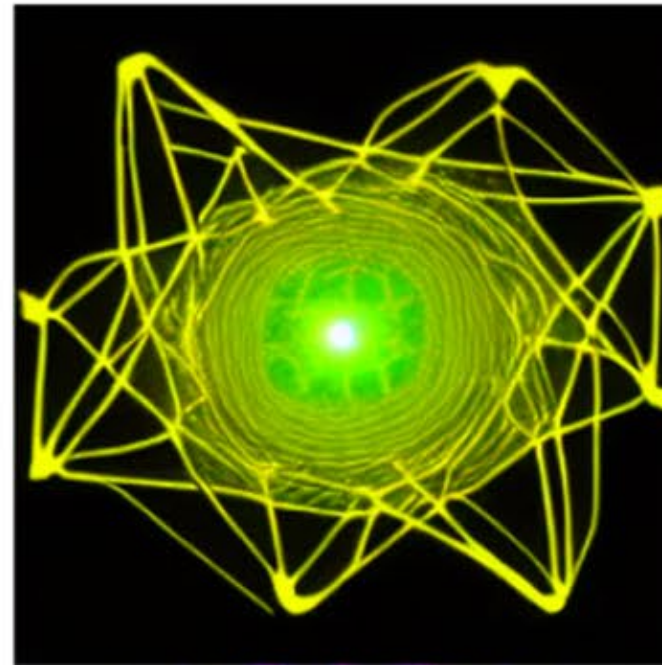
'An image of an animal half mouse half octopus'

'An illustration of a slightly conscious neural network.'

'A painting of a squirrel eating a burger.'

'A watercolor painting of a chair that looks like an octopus.'

'A shirt with the inscription: "I love generative models!"'



Course Organization

Class Website

- This is your primary source for class material (slides, notes, readings, references, etc.)
- Assignments & Grades will be released on Blackboard

<https://courses.cs.umbc.edu/graduate/675/>

~~ List of Topics

Phase 1: INTRODUCTION

- Prerequisites Review: Linear algebra, probability distributions, some basics of optimization
- Machine Learning Review (Regression, Classification, etc.)
- Feedforward Neural Networks, Backpropagation Algorithm
- Optimization and Regularization for Neural Networks

Phase 2: ARCHITECTURES

- Convolutional Neural Networks
- Autoencoders
- Self-Supervised Learning
- Contrastive Learning
- Recurrent Neural Networks
- Transformers
- Multimodal Learning
- Generative Adversarial Networks
- Diffusion Models
- Implicit Neural Representations

Phase 3: Challenges

- Evaluation
- Robustness
- Guarantees
- Tradeoffs

Grading



Grading

There is no Final Exam

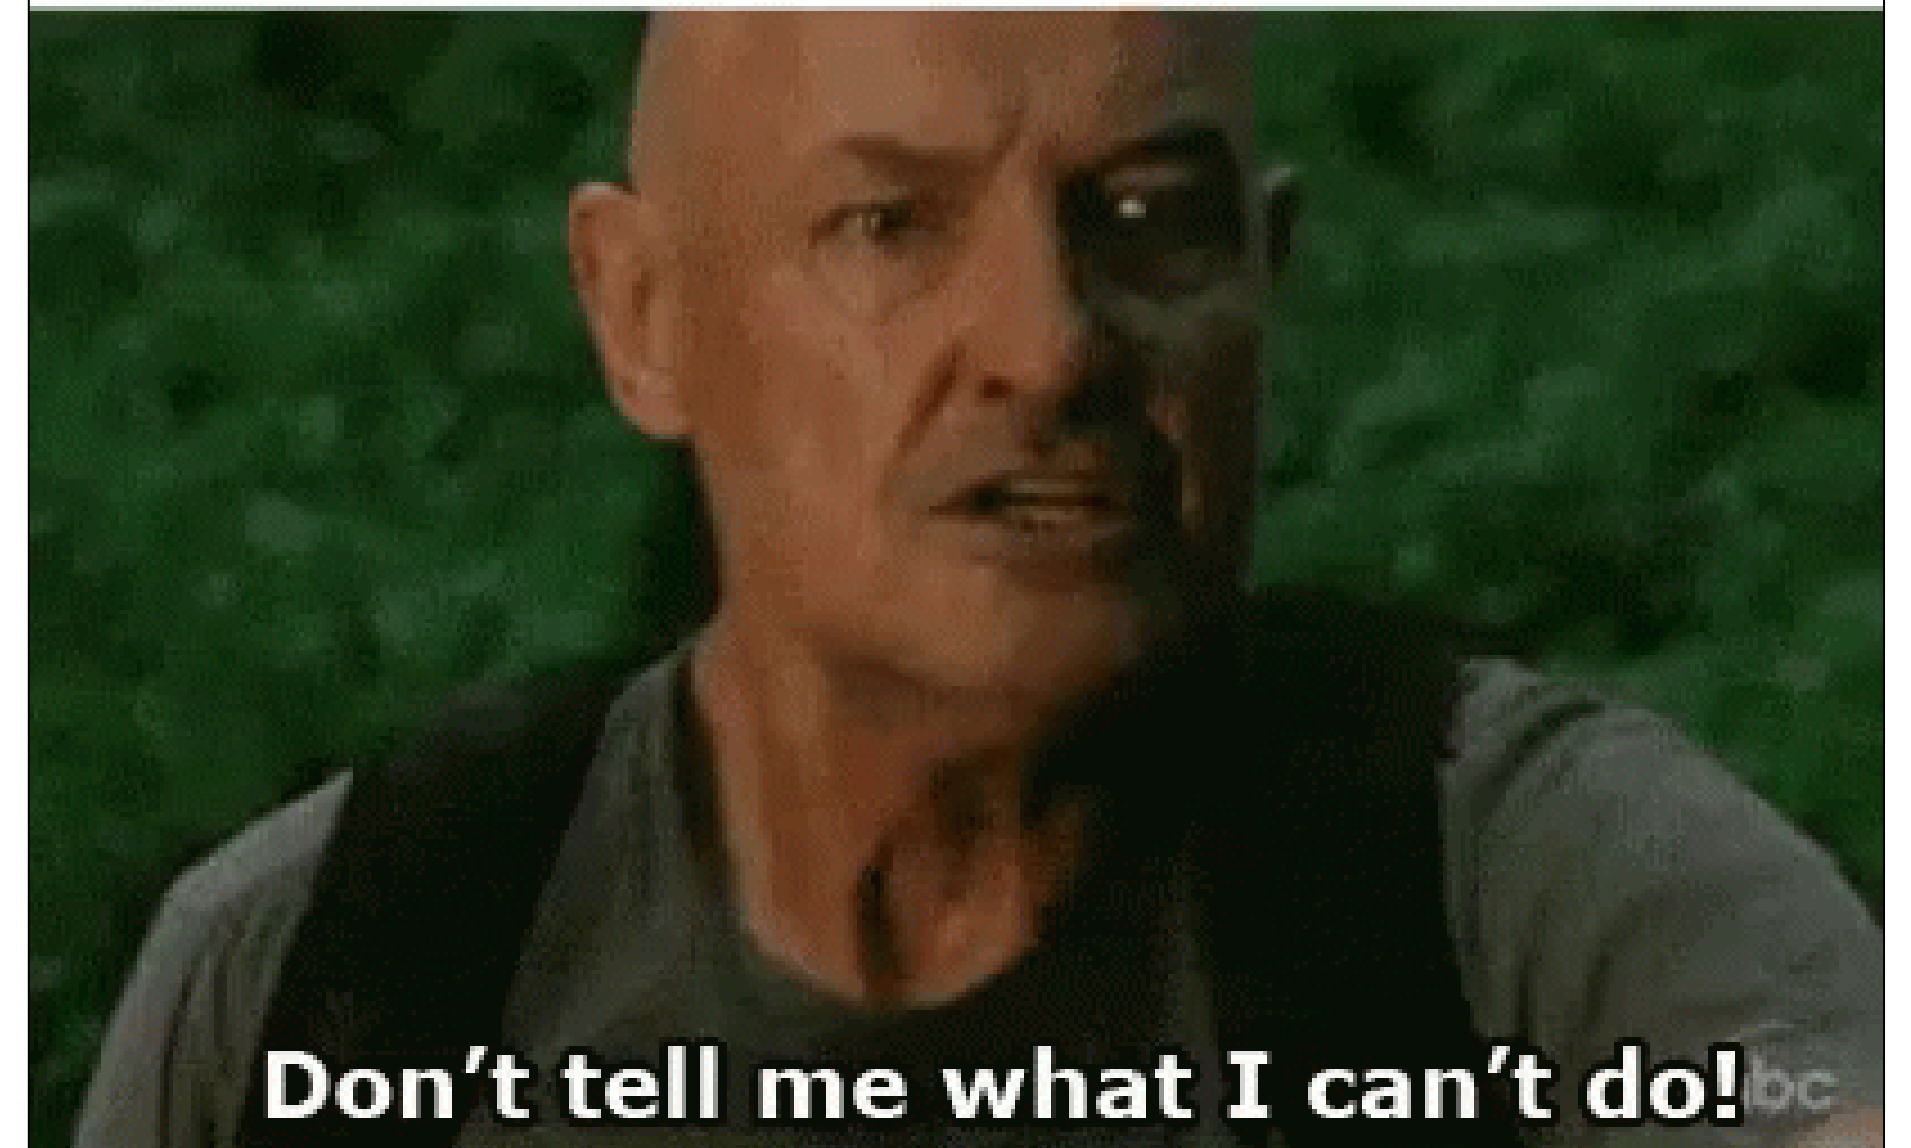
Homework	3-5 assignments including conceptual questions, exercises, and Python implementation and system design.	30%
Project	Course project in groups of 3-4 (smaller or individual groups only for PhD students with the professor's consent). <ul style="list-style-type: none">• 5% for project proposal• 5% for midterm video• 10% for final presentation• 5% for final report• 5% for summarizing other projects	30%
Quizzes	short quizzes during class	15%
Scribing	Take notes and typeset them for one lecture in the semester	5%
Midterm Exam	March 31, 2025 during class (4PM – 5PM)	20%

There will be opportunities for extra credit (max 10% of total grade)
Open-ended questions, writing assignments, creative assignments etc.

Homework

- Two major types of tasks:
 - Conceptual Questions (pen & paper)
 - Programming mini-projects in Python/Pytorch
- Conceptual:
 - Will be proofs / derivations
 - Points for steps (your thought process)
 - Show all of your work when answering
- Programming:
 - Most will be open-ended (exploratory). No right or wrong answers! Do your best!
 - Some tasks will result in bad results (by design). Evaluation will be based on your analysis and conclusions drawn from successful as well as failure cases.

**When the teacher says
you won't be able to do all
the homework in one
night**



Quizzes (in class)

- Multiple-choice / True-False / “draw a diagram” / short answer questions about previous lecture(s)
- Take notes during class (in your favorite format)
 - The very act of writing stuff down while you hear something new *really really* helps with information retention
- We will consider your “top k ” grades out of “ q ” quizzes, with $q - k \geq 2$
 - E.g. $q = 12$; $k = 10$
 - There will be no make-up quizzes if you miss a class

Scribing



- Scribing == Taking high-quality detailed notes during a lecture and typesetting them using Overleaf/LaTeX
(CVPR template is available on the Class Website).
- All students are required to scribe **at least once** during the semester.
- You can sign-up for your preferred lectures *(signup sheet: QR code)*
- ***Due Dates:***
 - *Notes for Monday lectures are due before class next Monday*
 - *Notes for Wednesday lectures are due before class next Wednesday*
- ***Submit notes by emailing both me and the TA with the title:***
[CMSC NN] Scribing <DATE>

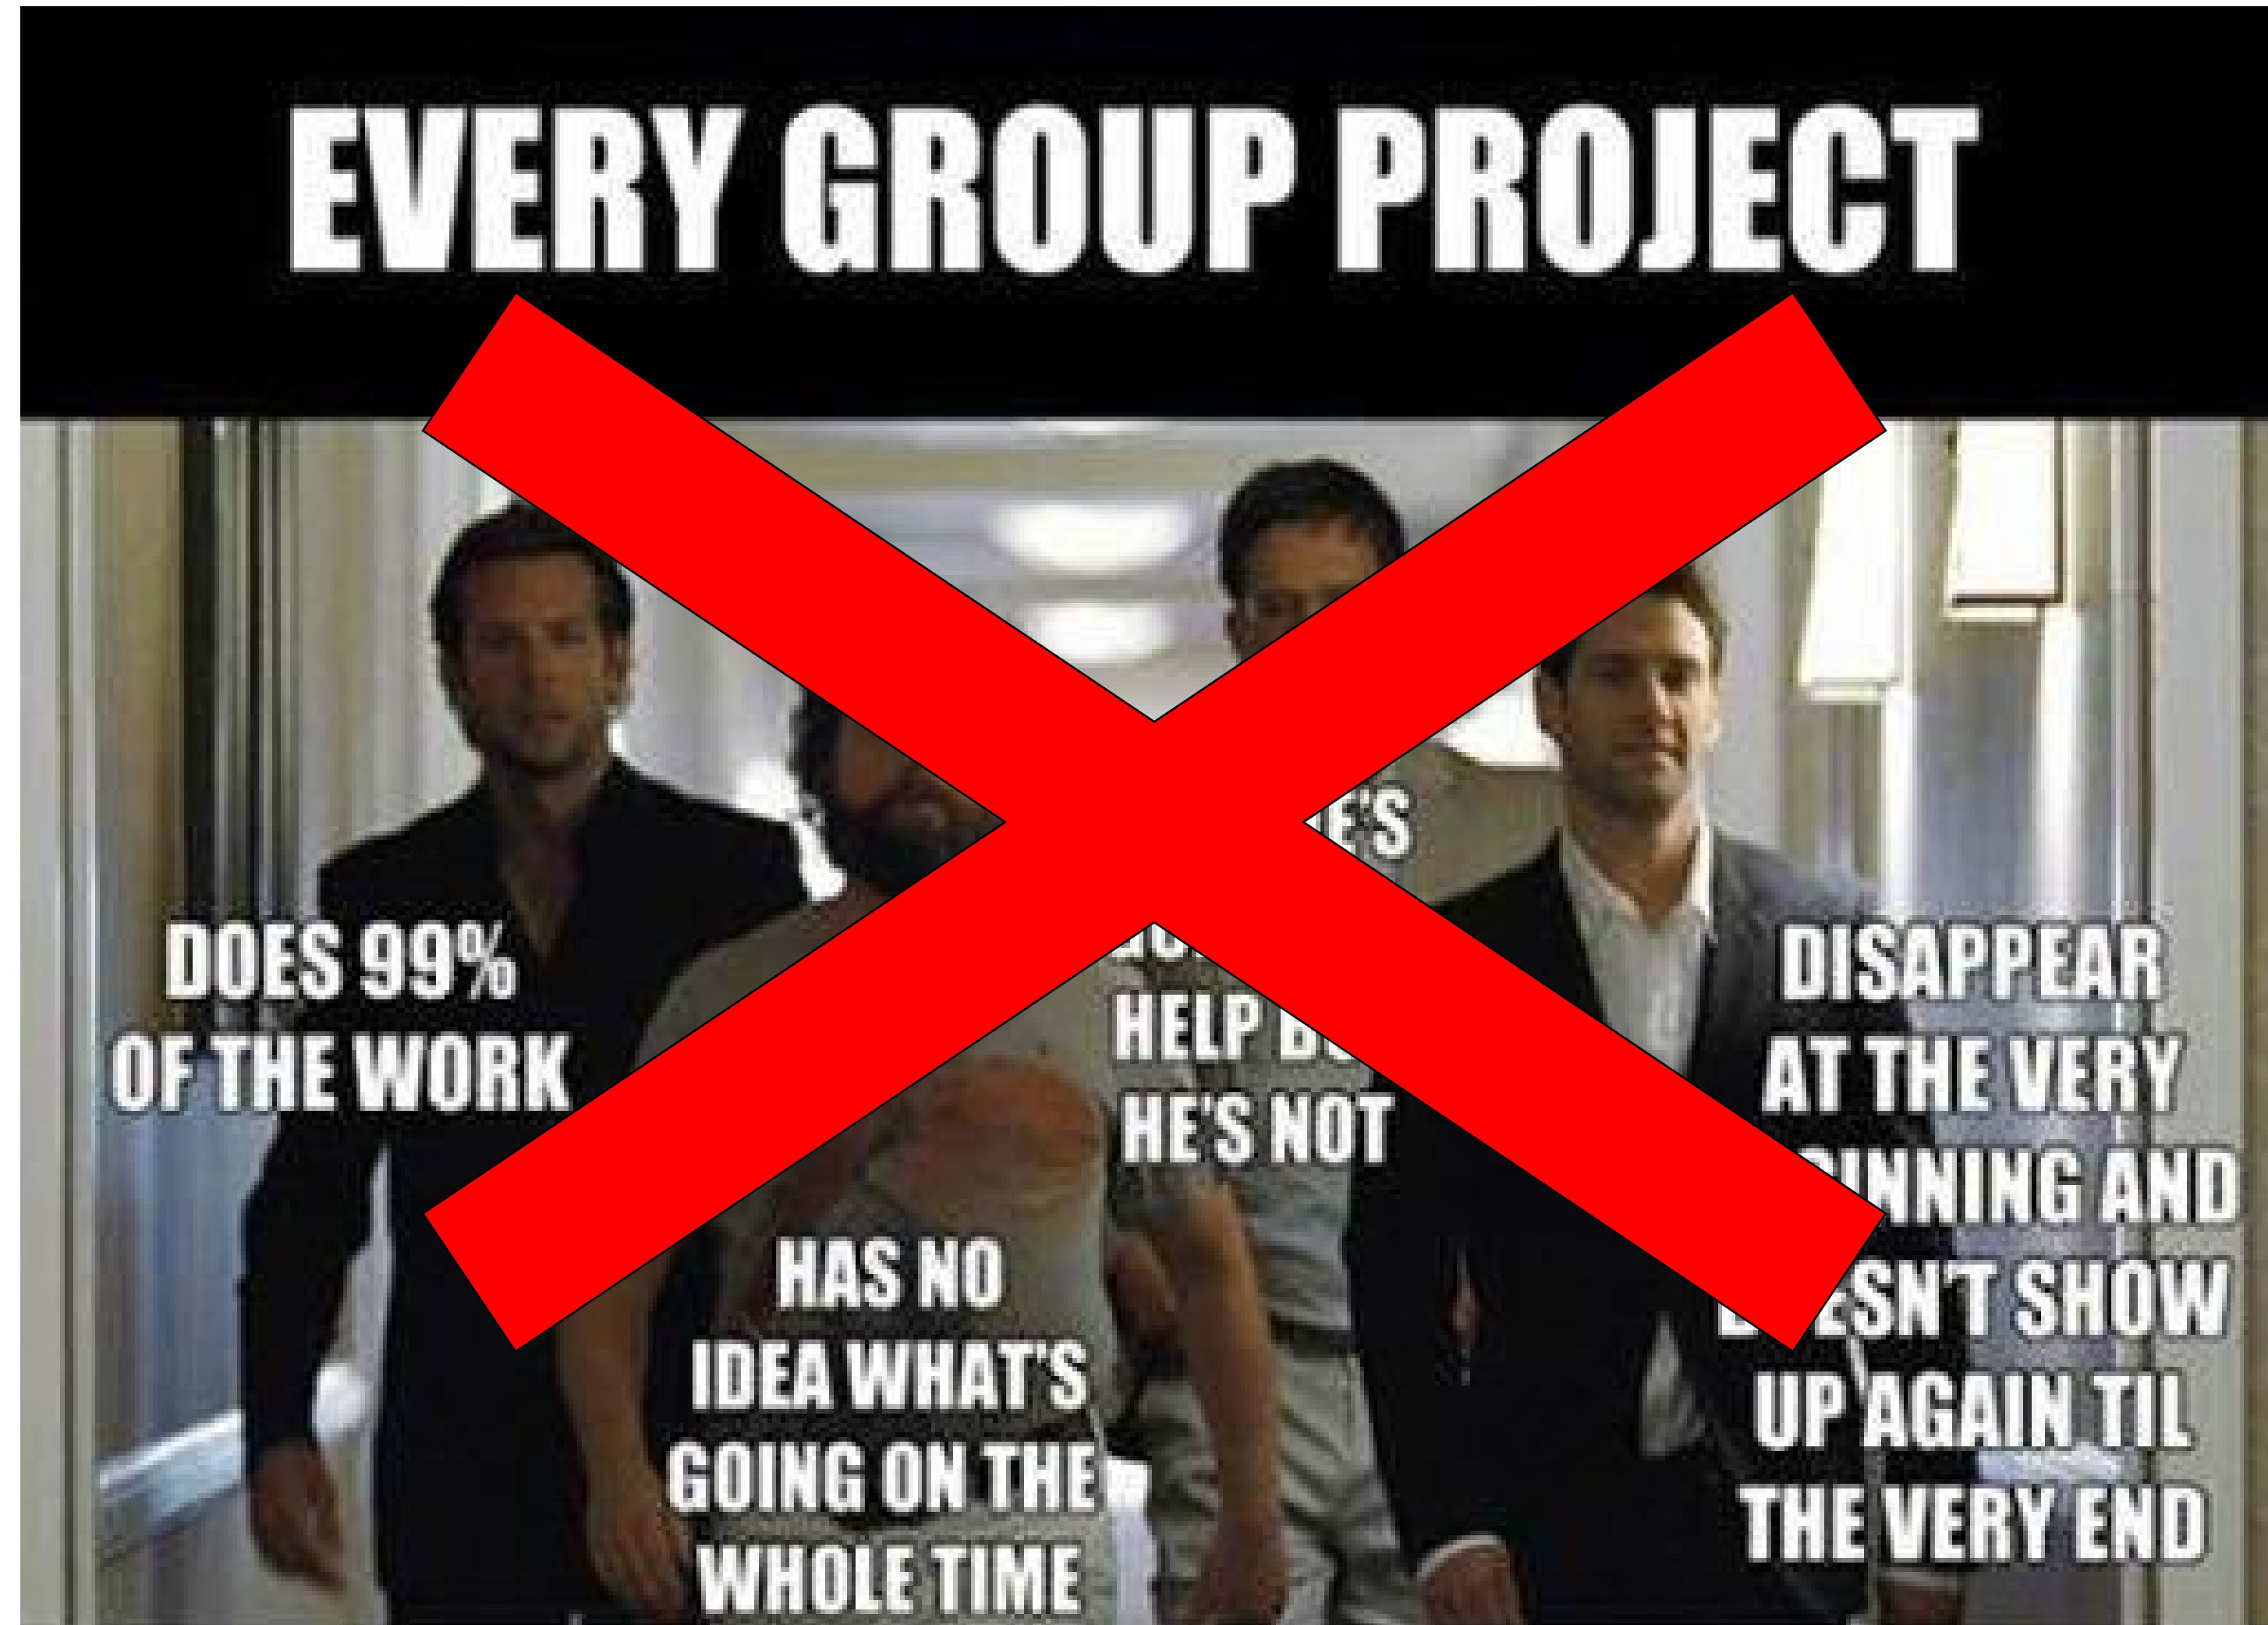
Extra Credit

- HW will have some optional problems (but for extra credit!)
 - Open ended questions or tasks where your creativity is required
 - We may showcase best outcomes of these in class
- Extra credit is capped at 10%
 - 5 extra points and 91 in the rest of the class → your final grade is 96
 - 18 extra points and 91 in the rest of the class → your final grade is 101

Project

- **Intention:** To give you an opportunity to do original research in deep learning and coherently writing up your result.
- **Expectation:** A simple original idea that you can *describe clearly*, relate the idea to existing work, *implement and test* the idea on some real-world problem(s)
- How?
 - Write code, run it on some data, make some figures/tables to present results
 - read relevant background papers, collect some references
 - write a report describing your model, algorithm, and results.

Project



Project

Each student will be graded separately

Project Topic

- Pick from our list
- Choose your own *

* *needs approval from Tejas*

Group Size: 3 to 4 students

- *Declare group by March 1*
- *PhD students: can work alone
(approval granted on a case-by-case basis after consultation with Tejas)*

Deliverables

(1) Group Formation	02/11	(2) Project Proposal	02/28
(3) Midterm Update (Video)	04/04	(4) Final Presentation (in class)	Last week
(5) Submit Slides (PDF)	05/16	(6) Submit Report (8 pg CVPR format)	05/16

Deadlines & Late Days

- Late Days: each student will get 10 late days (for projects and HW only)
 - Each late day extends the deadline by 24 hours
 - Using a late day does not influence the grade.
 - Late submissions turned in *after all 10 late days have been exhausted* **will receive 0 points**
- Projects: *if a group uses 1 late day → 1 late day deducted from all group members!*
- Quizzes: *we will select your best **10** scores out of **10 + d**.*
No make up quizzes.
- Late days are provided to help you deal with illness or injury, emergencies, paper deadlines, conference travel, interviews, computer problems, or other personal reasons.

Do not use the late days as an excuse for procrastination 😊

Attending Classes



- Attendance is mandatory
 - Exceptions: health reasons and personal emergencies.
 - Impossible to do a good job at homework, quizzes, and midterm (70% of your grade) without your attendance and attention in class.
- Perks:
 - Classes have (relatively easy) quizzes.
 - Meet your future project team – new friends – I'm still in touch with my NN project teammates from 2017!

-Do you really want to miss these memes?

Academic Integrity

I take academic integrity very seriously. You should too.

If you're unsure about something, ask us.

You are at a top-tier (R1) research university.

Use this privilege to learn. A good grade will follow.

Don't throw this opportunity away by cheating.

- Presentations, Survey Papers, Quizzes must be done independently.
- Do not plagiarize. Consequences will not be pleasant.
- Do not use “AI” assistants for any part of any assignment. Consequences will not be pleasant.
- Familiarize yourself with UMBC policy on plagiarism and other forms of cheating:
<https://academicconduct.umbc.edu/resources-for-students/>
- Read the syllabus for consequences of academic integrity violations.



Academic Integrity

I take academic integrity very seriously. You should too.

If you're unsure about something, ask us.

You are at a top-tier (R1) research university.

Use this privilege to learn. A good grade will follow.

Don't throw this opportunity away by cheating.

3.4 Use of AI Assistants

This is a class that teaches you how to design AI systems. Using AI systems to do class work is therefore an obvious form of cheating. The use of AI systems or AI assistants (including but not limited to language models such as ChatGPT) for completing *any* part of assignments in this class is considered cheating. There are no exceptions to this rule.

Academic Integrity

I take academic integrity very seriously. You should too.

If you're unsure about something, ask us.

You are at a top-tier (R1) research university.

Use this privilege to learn. A good grade will follow.

Don't throw this opportunity away by cheating.

3.6 Good Practices

If the integrity of your work in this course is challenged, you are responsible for demonstrating proof that the work submitted is your own. A good starting point is to enable versioning/tracking in Google Docs, Word, Pages, or other software so that your writing activities/progress during the semester can be logged if necessary. Keeping copies of research notes, scribbles, and related material may be helpful, too.

3.7 Viva or Oral Defense of Flagged Submissions

To ensure academic and professional integrity, I reserve the right to hold a one-on-one oral viva (defense) of submissions deemed questionable, to determine your knowledge and mastery of the topic/resources versus the material submitted. Failing that viva will result in an 'F' on the assignment and an Academic Integrity violation report filed with the Graduate School.

Academic Integrity

I take academic integrity very seriously. You should too.

If you're unsure about something, ask us.

You are at a top-tier (R1) research university.

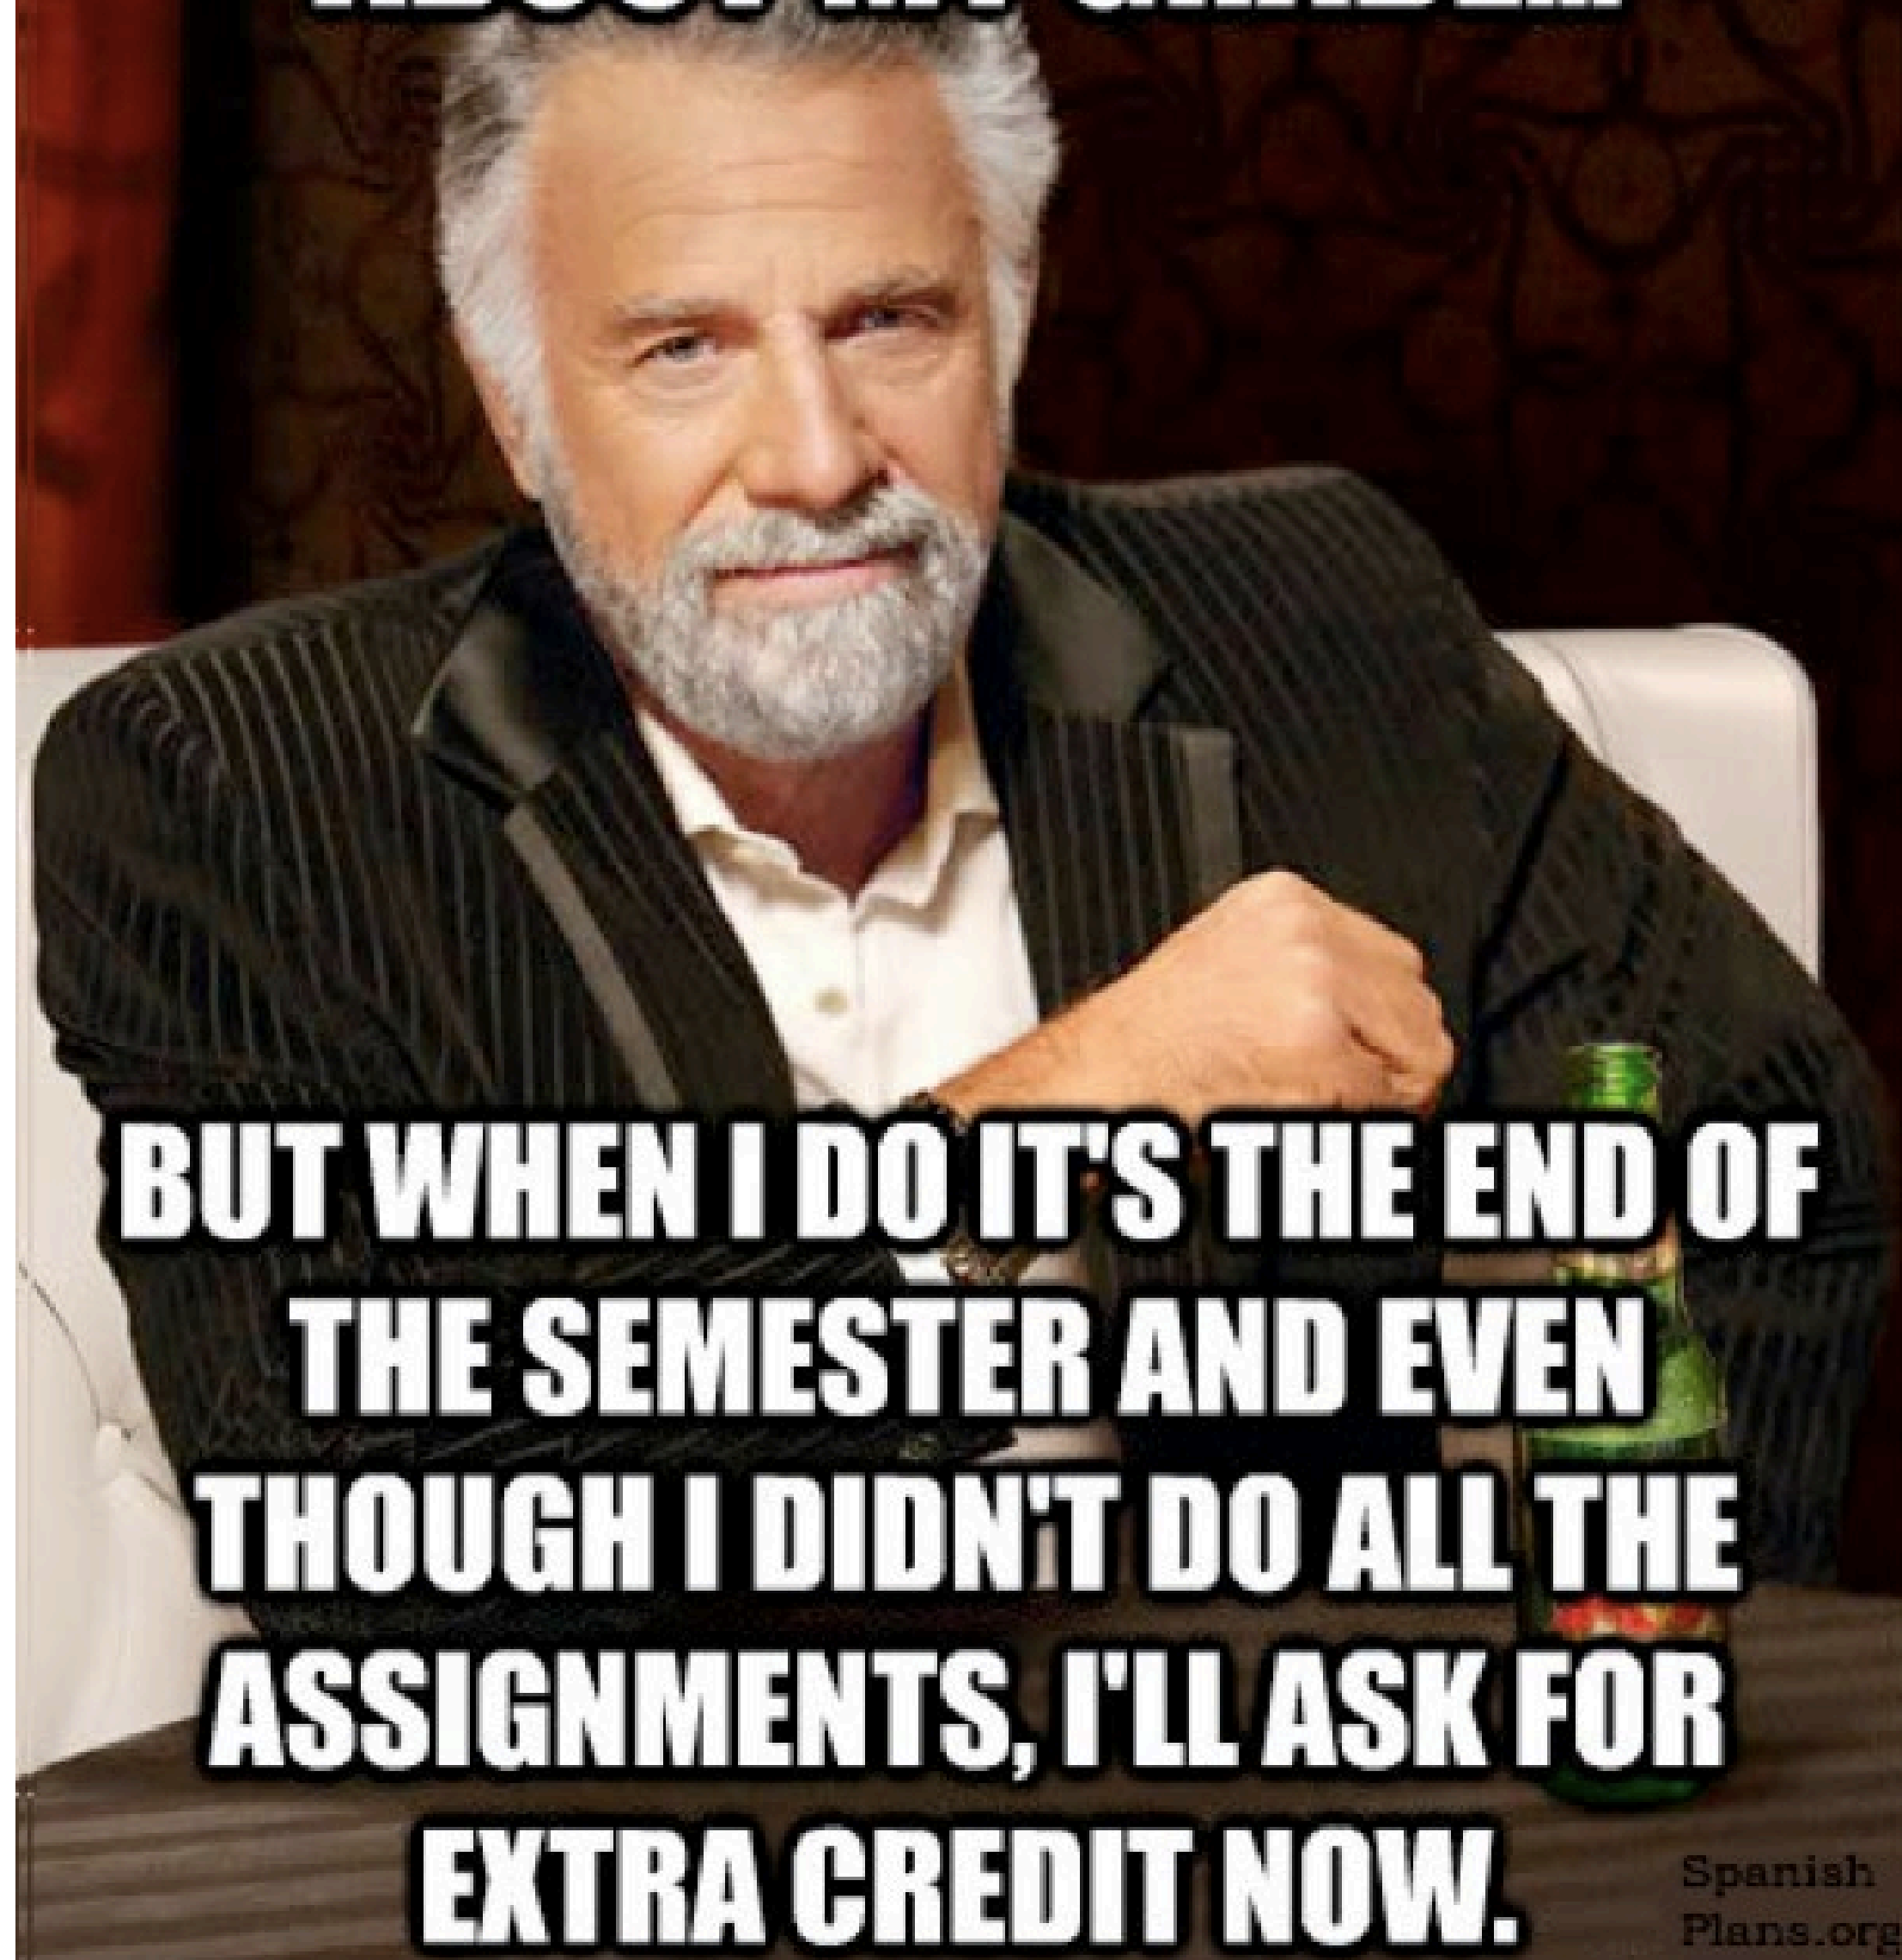
Use this privilege to learn. A good grade will follow.

Don't throw this opportunity away by cheating.

3.8 Penalties

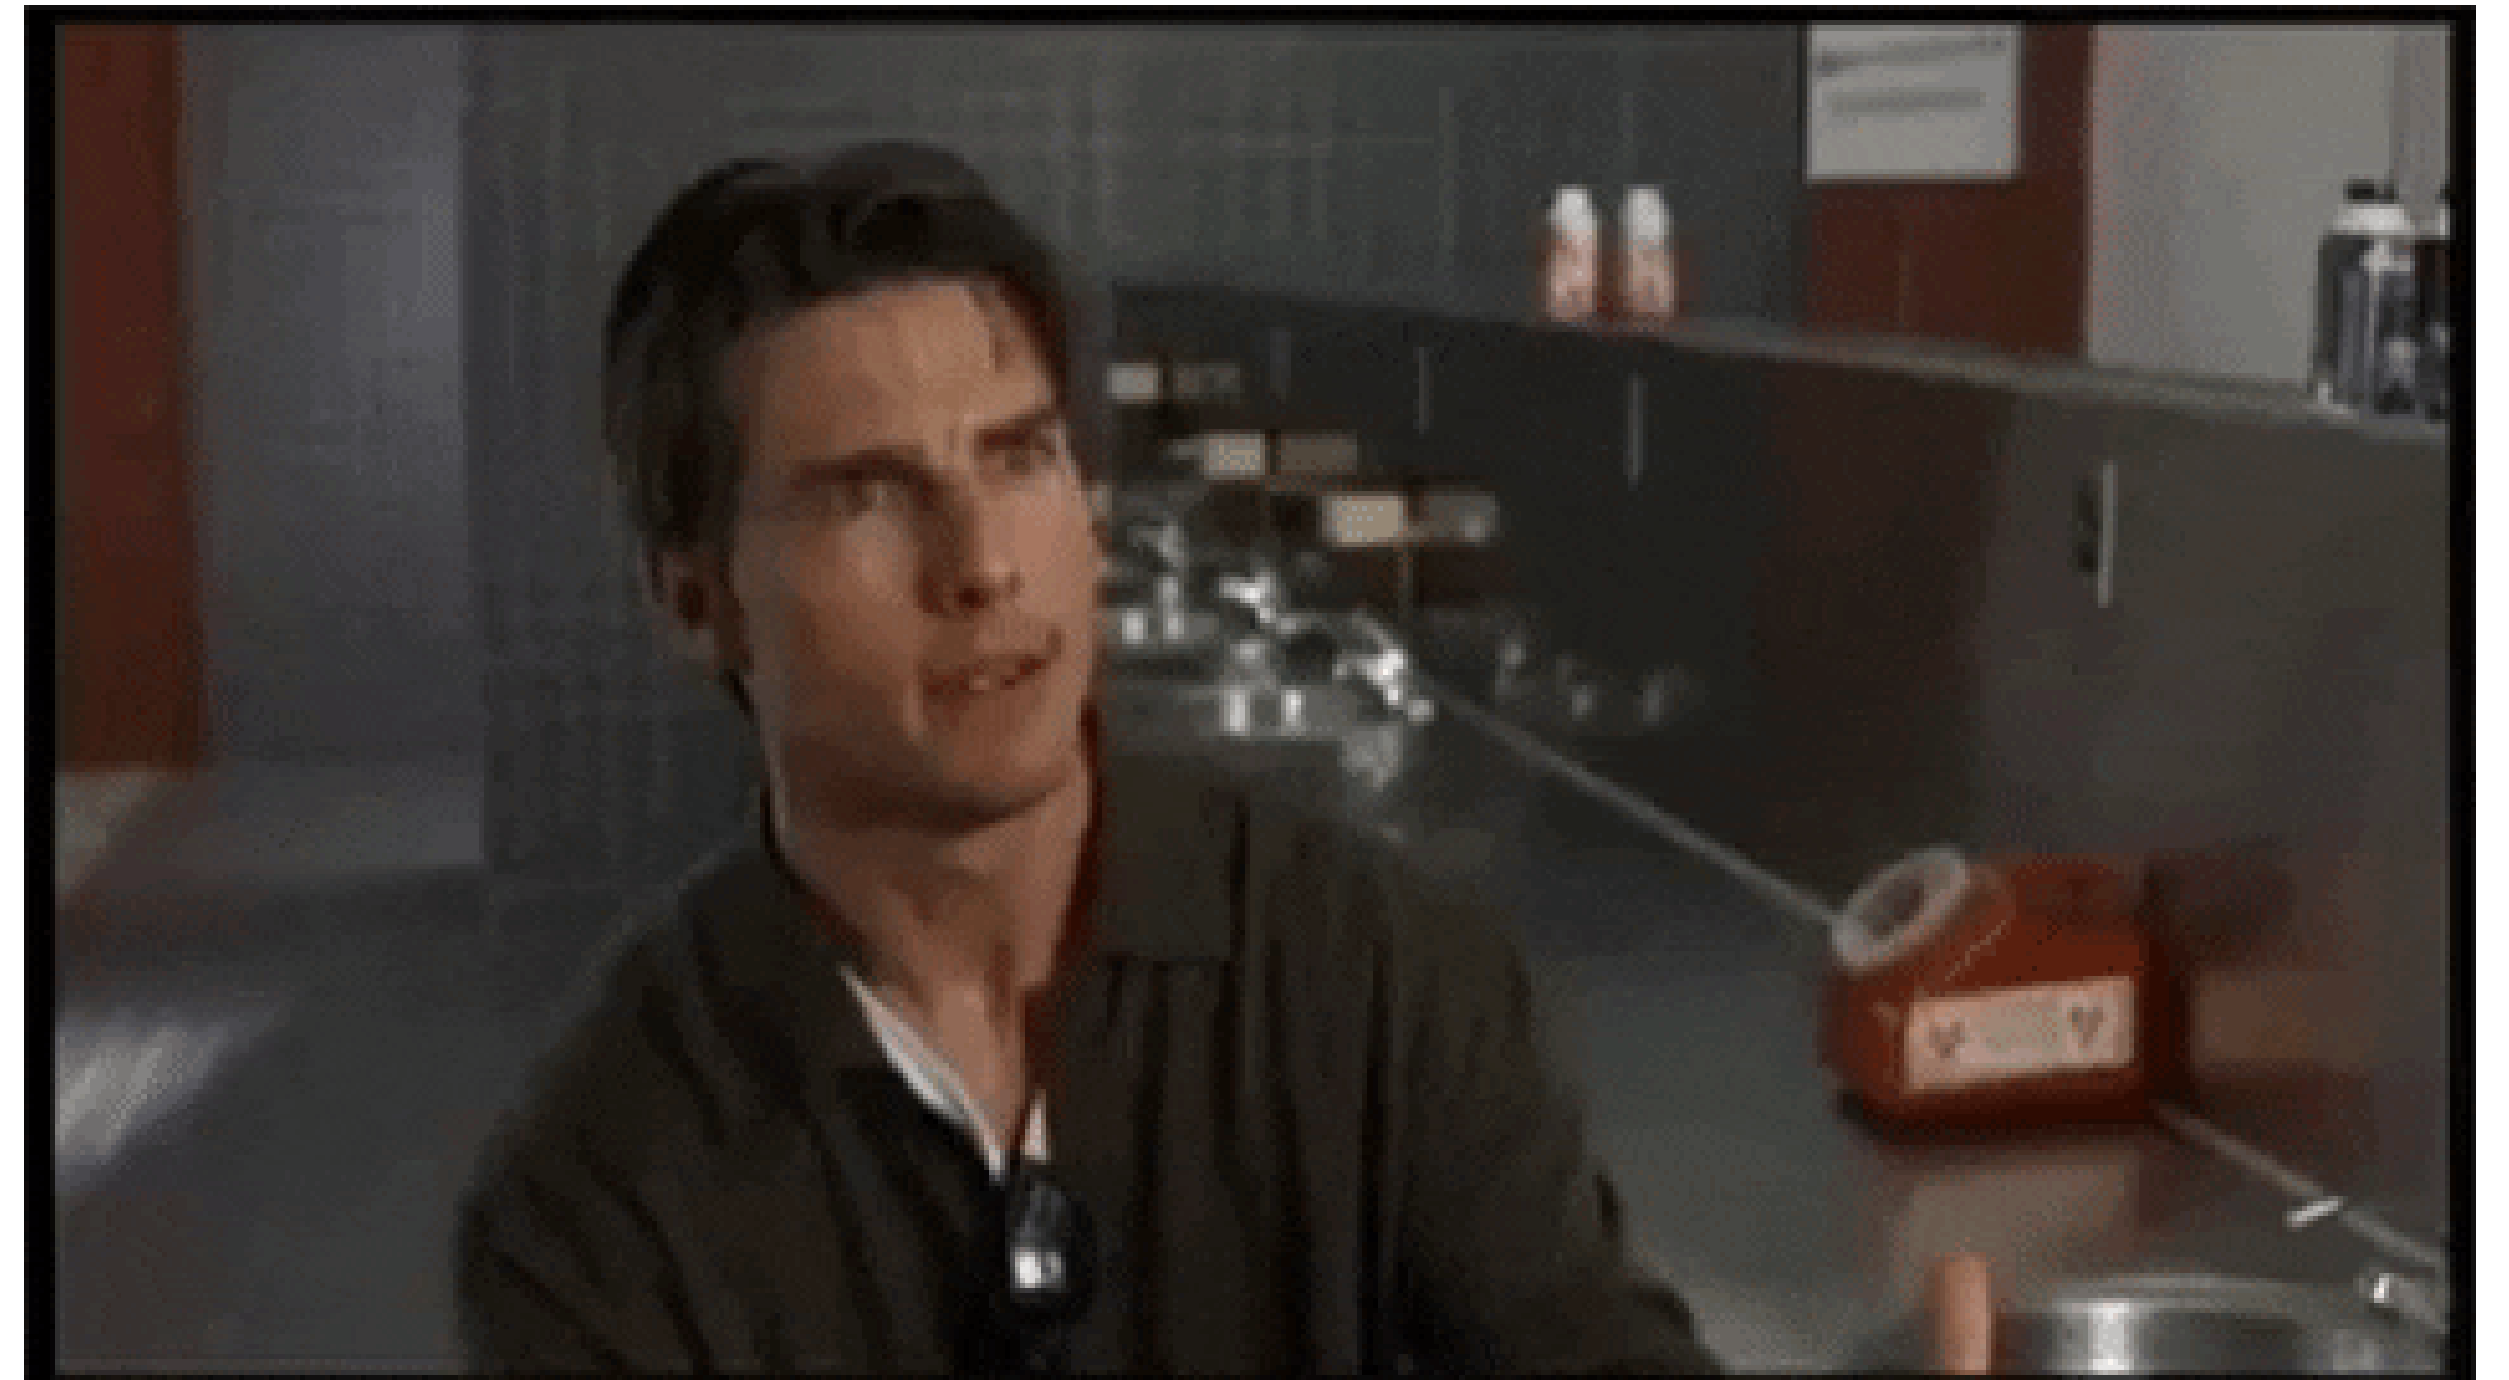
Academic misconduct could result in disciplinary action that may include, but is not limited to, suspension or dismissal. The **absolute minimum penalty** for a first offense of academic dishonesty in this course is a grade of zero on the assignment and a one-letter-grade reduction in the final class grade. However, depending on the nature of the offense, the penalty may be more severe, including but not limited to an F for the course, suspension, or expulsion. The minimum penalty for a second offense of academic dishonesty is an F for the course without possibility of dropping, but may be more severe.

**I DON'T ALWAYS CARE
ABOUT MY GRADE...**



Spanish
Plans.org

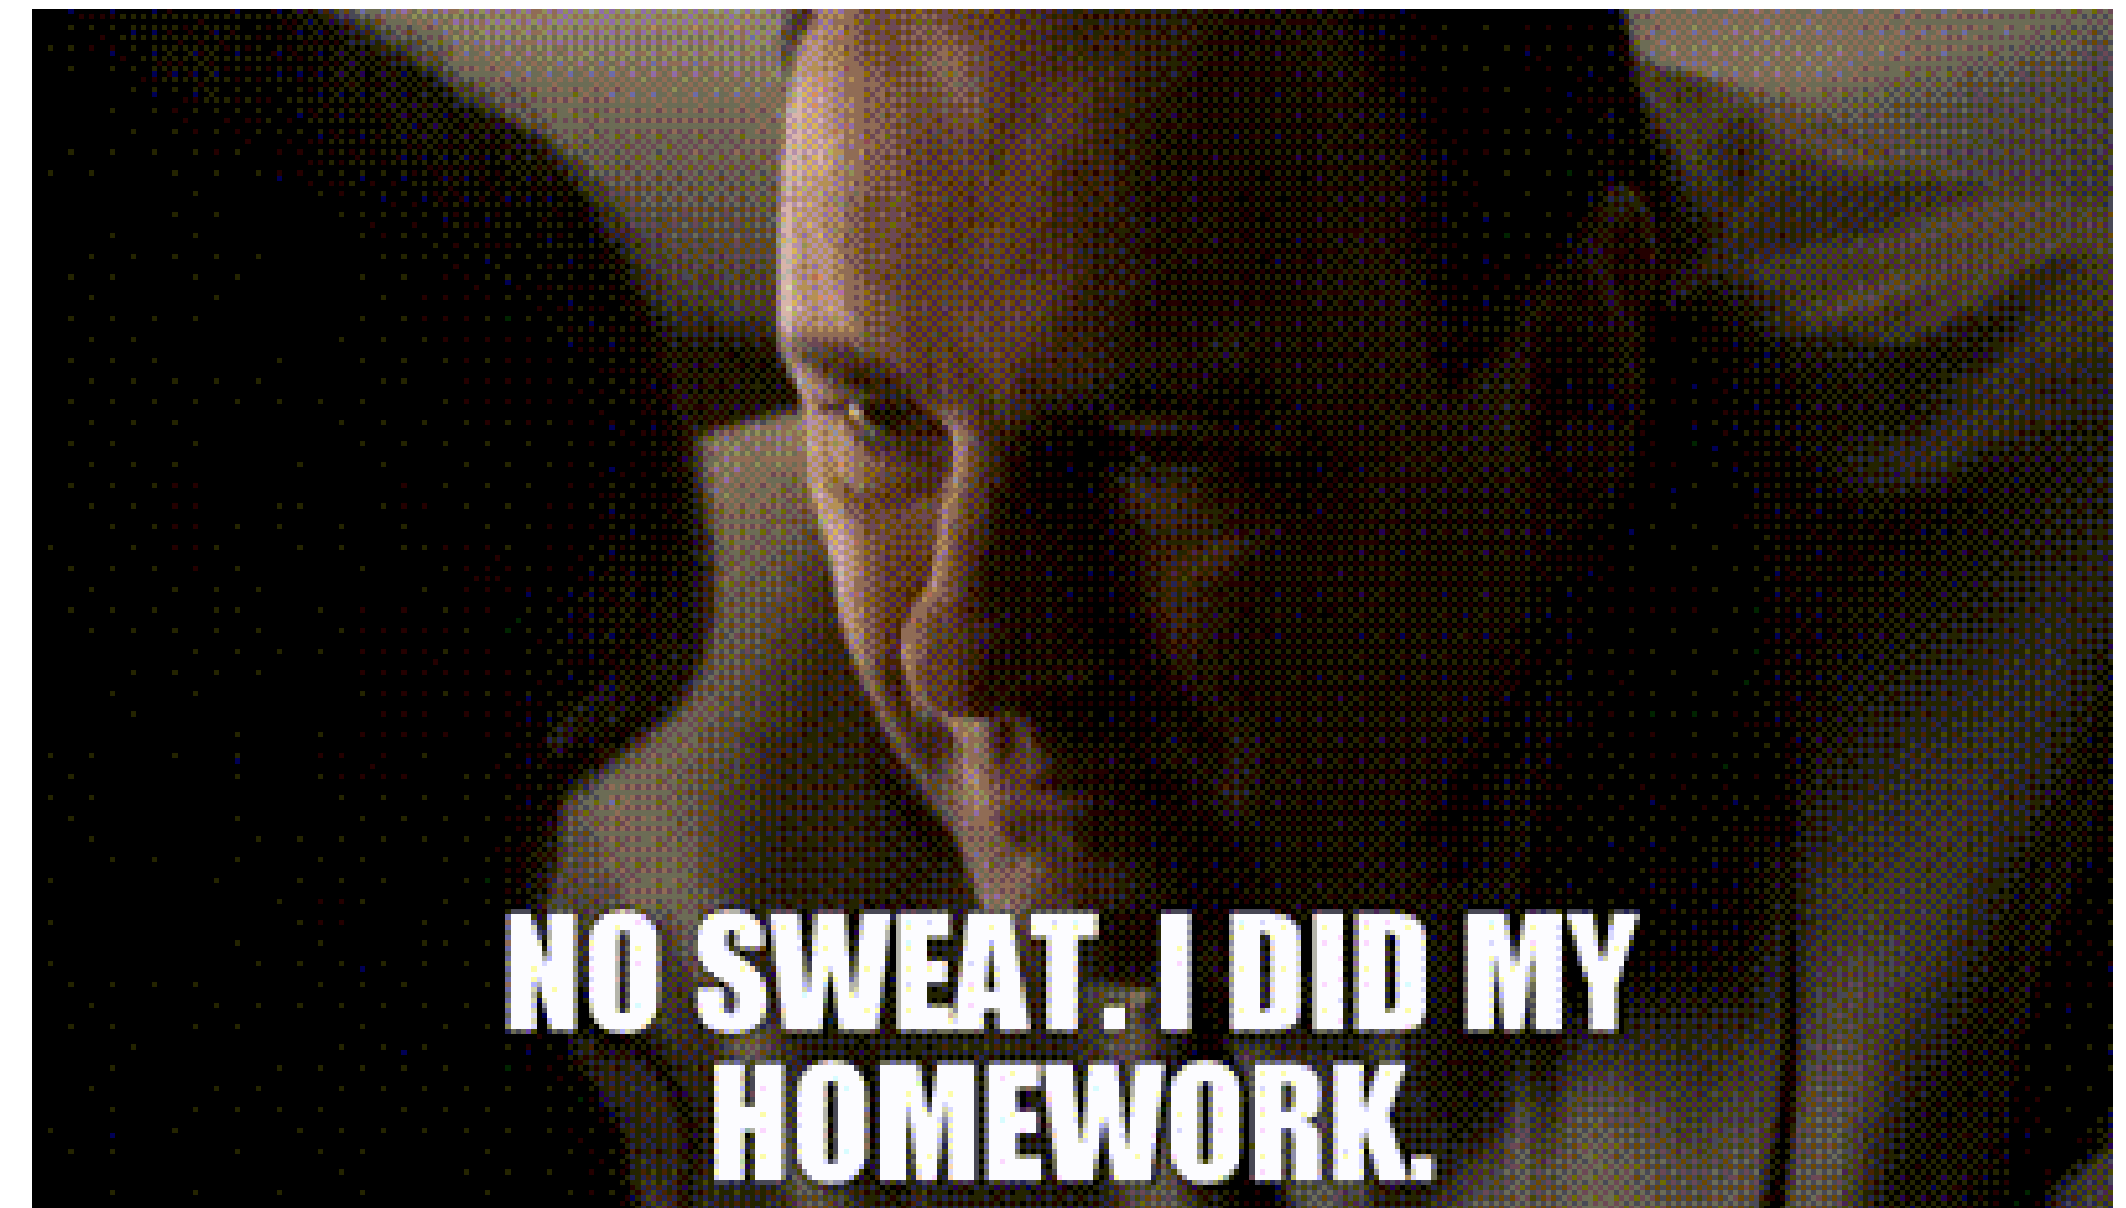
Seek Help Early!



Help us help you.

Homework 0 (Optional)

- This is optional. Due next Wednesday (03/05)
- An opportunity for you to self-evaluate whether you have the pre-requisites for this class (linear algebra, probability, stats, etc.)
- Grading:



SCORE > 80 ?

NO

You should either:

- (1) Brush up on pre-reqs
(resources on next slide / website)
- (2) Consider dropping the class

YES

SCORE =
100% ?

NO

+1 Extra Credit

YES

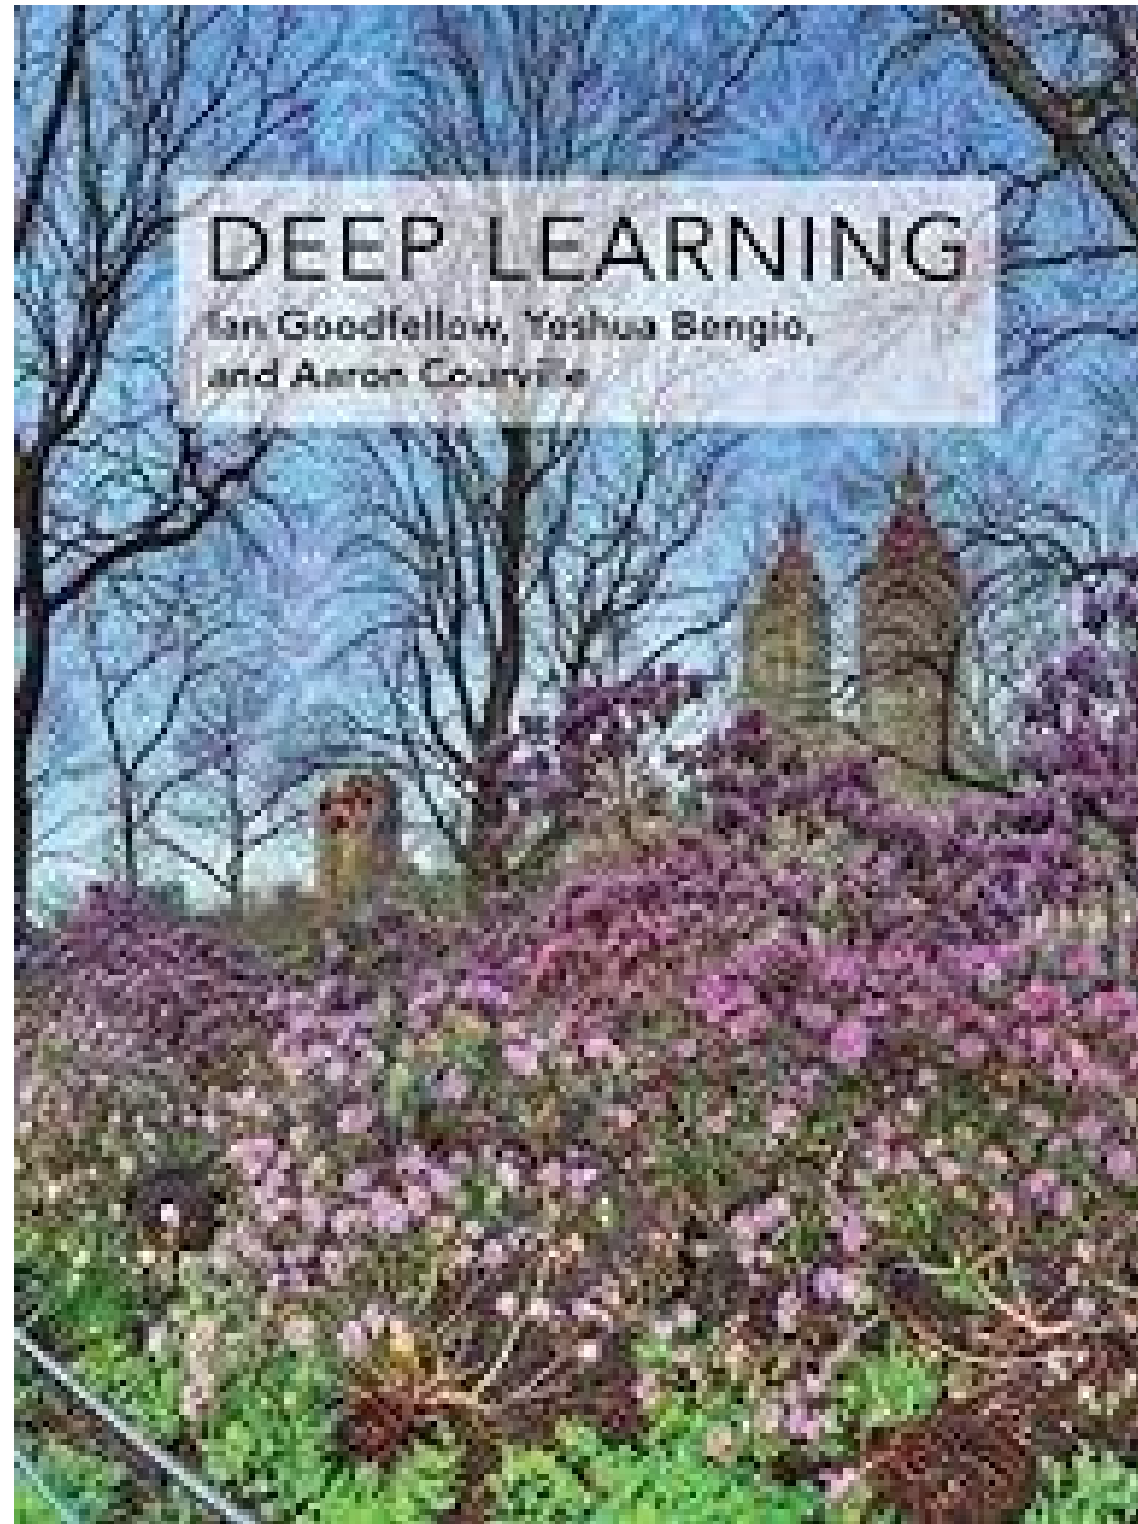
+2 Extra Credit

Recommended Background

- Linear algebra + calculus + geometry + prob/stats (required)
 - 475/675 should not be your first introduction to these topics
 - Without these tools, you are likely to struggle with the course.
- Python programming with numerical libraries like `numpy`
 - TA will give a tutorial on computer vision with PyTorch
- Useful resources to brush up on these topics
 - [deeplearningbook.org/contents/linear_algebra.html](https://www.deeplearningbook.org/contents/linear_algebra.html)
 - https://www.deeplearningbook.org/slides/02_linear_algebra.pdf
 - [deeplearningbook.org/contents/prob.html](https://www.deeplearningbook.org/contents/prob.html)
 - https://www.deeplearningbook.org/slides/03_prob.pdf

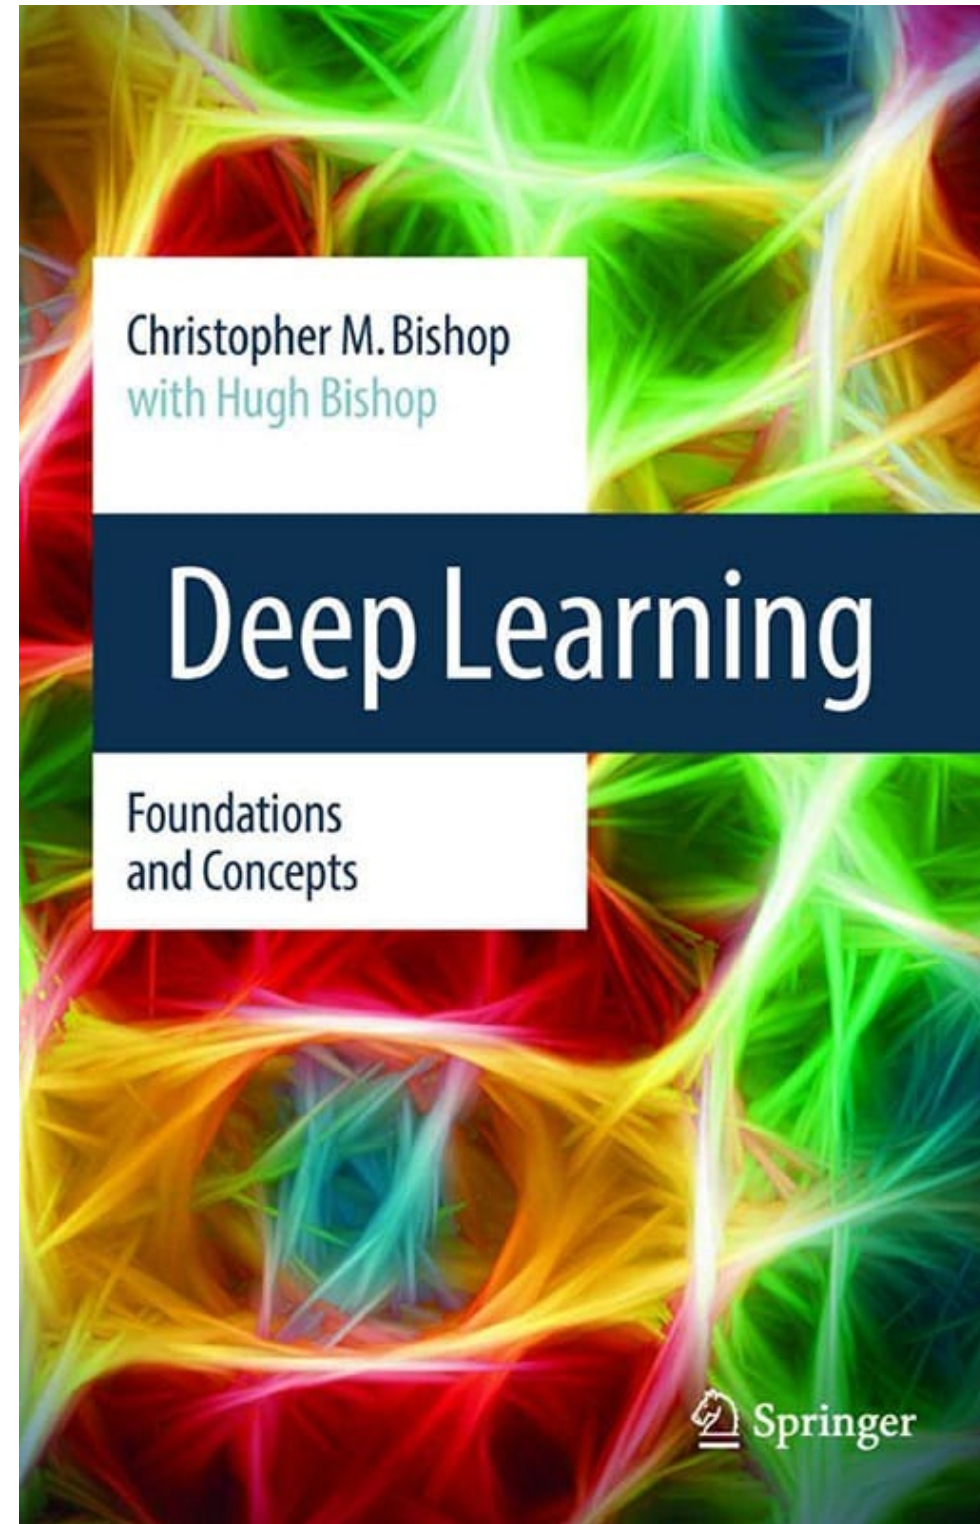
Readings

Topic-specific reading will be provided (pdf)
Other useful resources:



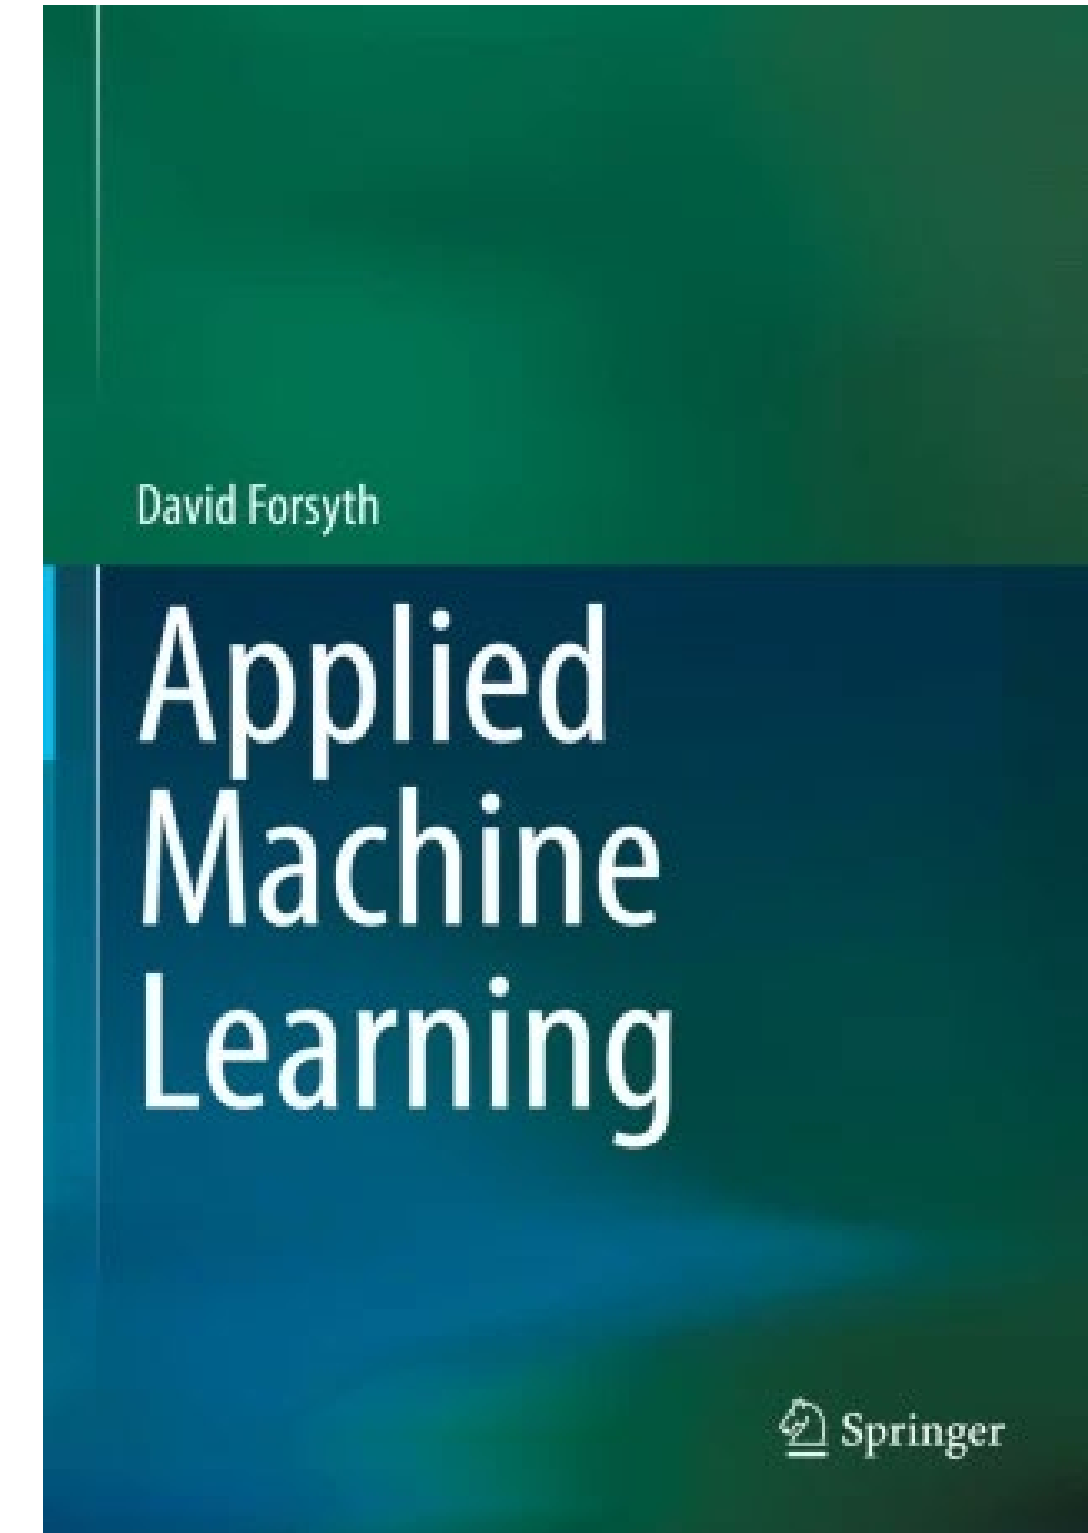
www.deeplearningbook.org/

Free download



<https://www.bishopbook.com/>

Free download



<https://link.springer.com/book/10.1007/978-3-030-18114-7>

Free download via UMBC

475 (undergrad) vs 675 (grad)

- We are confident that 475 students are at least as capable than their 675 classmates
- ***No difference in class materials, exams, quizzes, majority of homework***
- Homework: additional parts for 675
- Grad projects will be evaluated at a higher standard

Grad student and undergrad sitting in the same class



475 (undergrad) vs 675 (grad)

Main difference: projects (scope and novelty)

- Projects will be graded in terms of “relative growth”
 - *You may have previous research experience*
 - *You may be taking this class to get research experience*
- Grad projects:
 - Original & unique research hypothesis with a potential for publication
- Undergrads projects can be:
 - Original & unique research hypothesis with a potential for publication
 - Working on an idea that we provide (i.e. you get to skip “ideation”)
 - Innovative applications or combination of existing work

Grad student and undergrad sitting in the same class



FAQ: I'm waitlisted – what can I do?

1. I'm sorry this is happening to you. Thanks for your patience I (like you) am also bound by UMBC's waitlist policy.

01/31/2025	Waitlist Deadline	Last day to add yourself to a waitlist for a course
02/03/2025	Waitlists Deactivated	Waitlists are deactivated and purged
02/07/2025	Schedule Adjustment Deadline	Last day to make changes to your schedule
		Last day to change grade method
		Last day to drop a course without a W grade

2. If you're still interested AND are on the waitlist, please see me AFTER class on 01/29. **No promises, but I'll try**
3. Consider CMSC 678 (Machine Learning) as an alternative.

FAQ: Access to Practice Quizzes and Midterm?

- I am revamping UMBC's neural networks course
 - The course hasn't been offered for more than a decade.
- Therefore, no practice / past midterms or quizzes are available!
- This is a graduate-level class – we will test your concepts (and not your ability to do fast arithmetic / memorization)
 - Make sure you understand the concepts
 - Read to the reference books / materials

FAQ: Can I join your research lab?

- The Cognitive Vision Group (CVG) broadly works on “conceptual characterization of visual scenes” with some goals including:
 - interpretation of visual data in presence of incomplete information,
 - recognizing and adapting to novelty and variations,
 - leveraging external knowledge and reasoning modules to generalize to new contexts, domains, environments, and tasks,
 - acquiring visual knowledge and communicating it to other machines and humans.
- Joining (See FAQ/note on my website)
 - Take this class and talk to me during office hours about your interests
 - https://www.tejasgokhale.com/docs/cvg_starter_pack.pdf



CVG

Other Questions?



Scribing /
Project Team
Signup Sheet



Class Website

