Deep Learning (1470)

Randall Balestriero

Jan 20th, 2026

Who Am I?

Who Am I?

Who Am I?

Who Am I?

- 2016 -> 2021 

- 2021 -> 2023 facebook AI Research

- 2024 -> 

- 2025 -> 

Who Am I?

Who Am I?

- 2016 -> 2021 

- 2021 -> 2023 facebook AI Research
- 2024 -> 

- 2025 -> 

- Theory everything!
- Foundation Models!
- World Models!
- Self Supervised Learning!
- Signal Processing!

Who Am I?

Who Am I?

- 2016 -> 2021 

- 2021 -> 2023 facebook AI Research
- 2024 -> 

- 2025 -> 

- Theory everything!
- Foundation Models!
- World Models!
- Self Supervised Learning!
- Signal Processing!
- Office Hours @ CIT453:
  - Monday: 11am to noon
  - Wednesday: 11am to noon

Why Deep Learning?

Why Deep Learning?

Why Deep Learning?

Why Deep Learning?

- I love(d) programming (C++/Theano)

Why Deep Learning?

Why Deep Learning?

- I love(d) programming (C++/Theano)
- I am lazy

Why Deep Learning?

Why Deep Learning?

- I love(d) programming (C++/Theano)
- I am lazy
- -> let’s code a machine to do my work

Why Deep Learning?

Why Deep Learning?

- I love(d) programming (C++/Theano)
- I am lazy
- -> let’s code a machine to do my work
- -> edge cases everywhere, ugly code…

Why Deep Learning?

Why Deep Learning?

- I love(d) programming (C++/Theano)
- I am lazy
- -> let’s code a machine to do my work
- -> edge cases everywhere, ugly code…
- -> we need “intelligent programs”

Why Deep Learning?

Why Deep Learning?

- I love(d) programming (C++/Theano)
- I am lazy
- -> let’s code a machine to do my work
- -> edge cases everywhere, ugly code…
- -> we need “intelligent programs”

- What about you?
  - No wrong answers!

Why Deep Learning?

Why Deep Learning?

- I love(d) programming (C++/Theano)
- I am lazy
- -> let’s code a machine to do my work
- -> edge cases everywhere, ugly code…
- -> we need “intelligent programs”

- What about you?
  - No wrong answers!

Think Big!


ALPHAGO
00:20:49
Google DeepMind
Challenge Match
LEE SEDOL
00:36:22


nature
DeepMind


10:24
How do you say I do not want sour cream in my burrito in Spanish
Tap to Edit
osxdaily.com
TRANSLATION
English
I do not want sour cream in my burrito
Spanish
No quiero crema agria en mi burrito

Think Big!

Think Big!

Think Big!

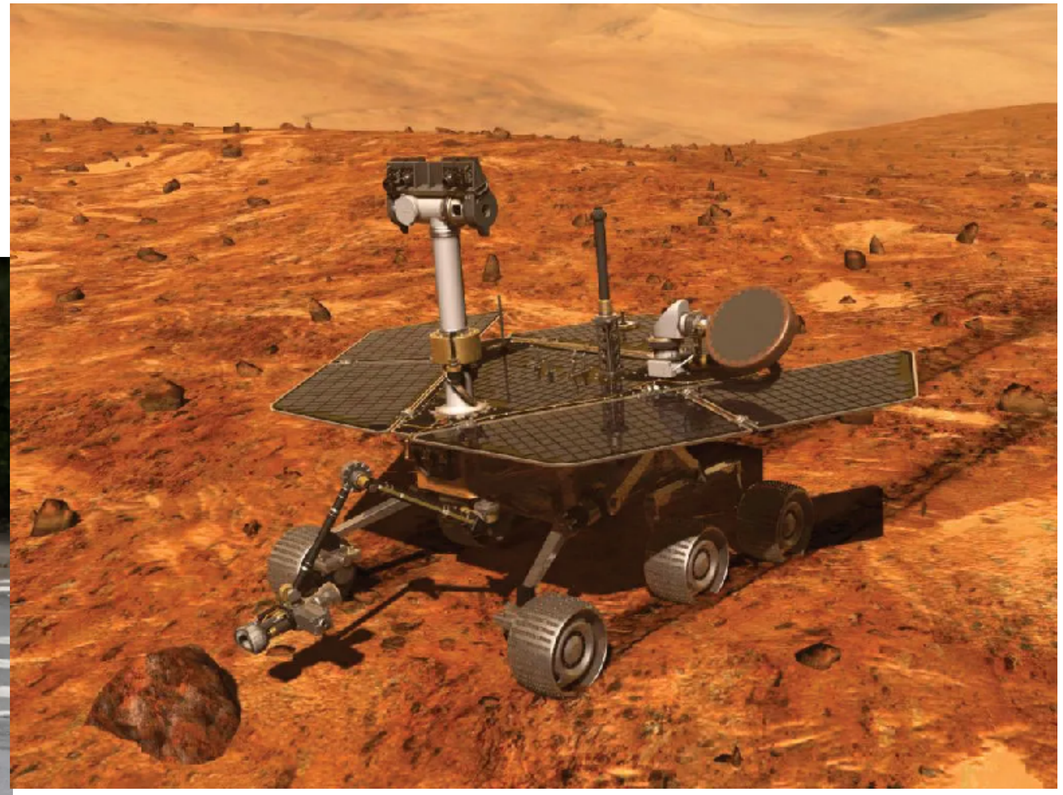

Our goal is to show you the way!

- What is Deep Learning:
  - The models
  - The use-cases
  - The code
- How to iterate/debug
- How to approach a problem

Our goal is to show you the way!

- What is Deep Learning:
  - The models
  - The use-cases
  - The code
- How to iterate/debug
- How to approach a problem

How?

Amazing HTAs

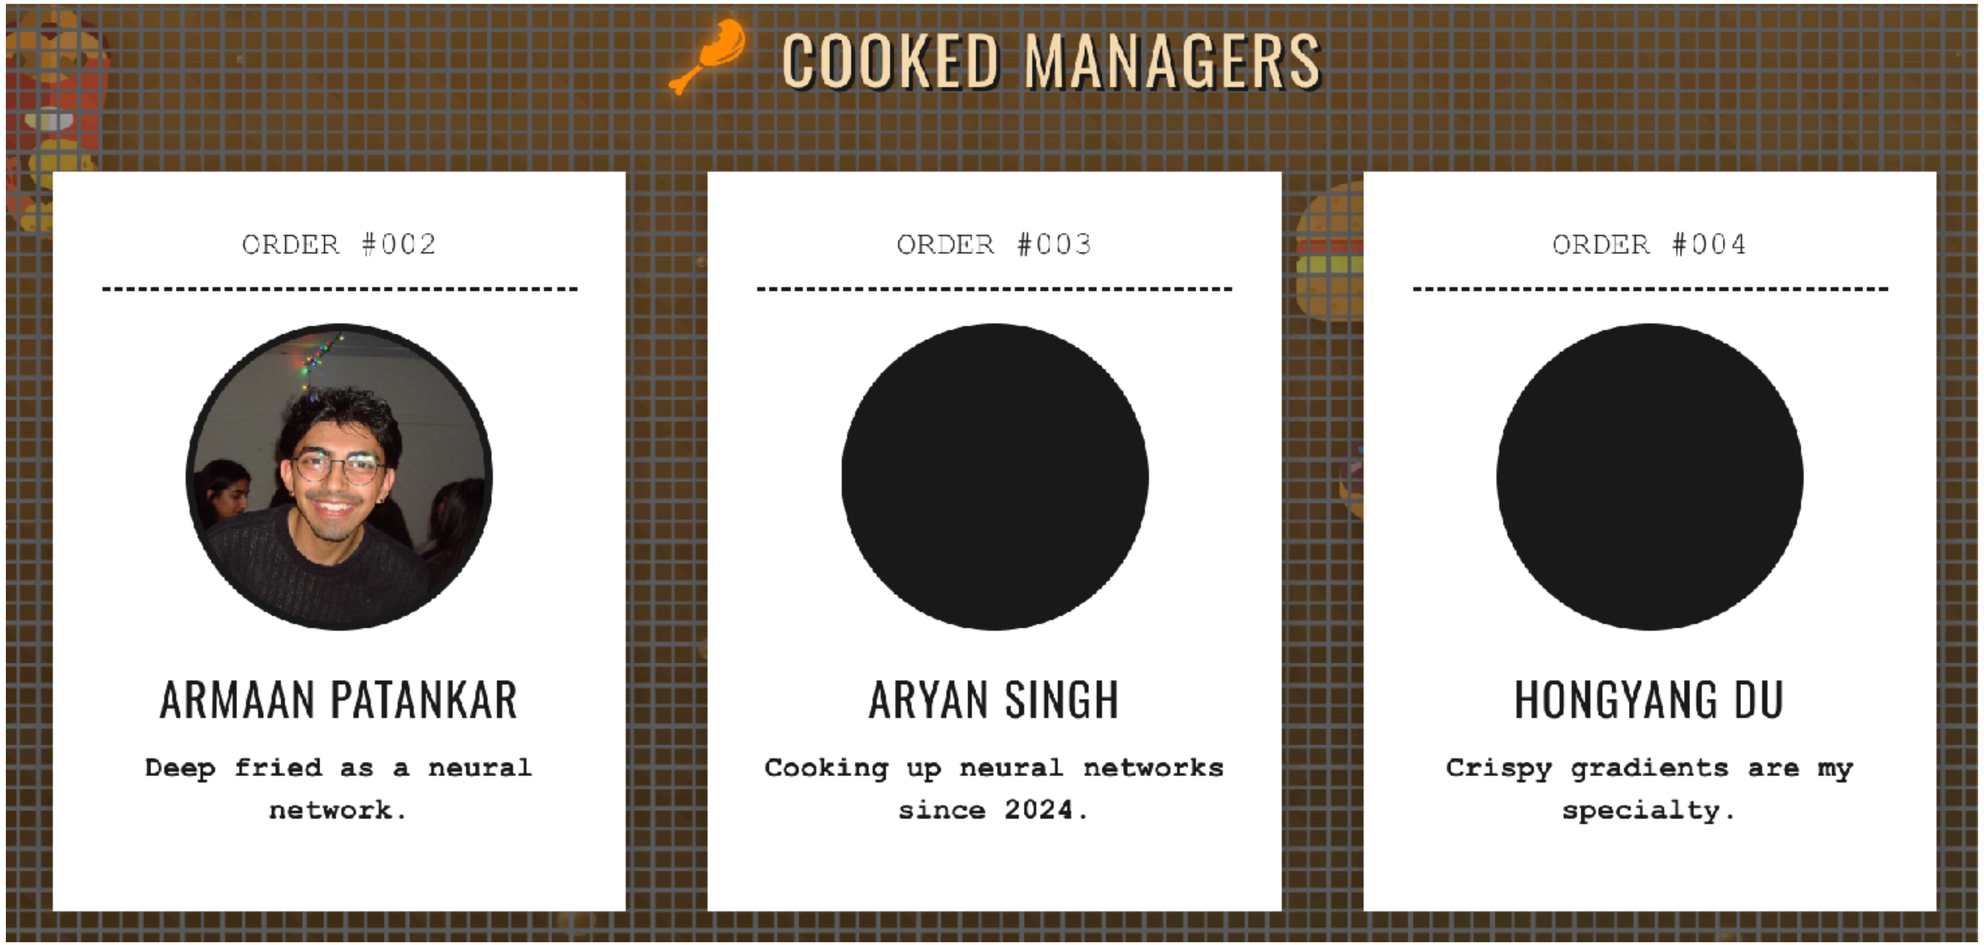

How?

And amazing TAs

ORDER #100

JOHNNY ELIAS

Crispy gradients are my specialty.

ORDER #101

DURU HUSYENI

Master of the deep fried descent.

ORDER #102

JASON LIN

Serving hot batches of tensors.

ORDER #103

DANIEL ZHU

Seasoning models with extra bias.

ORDER #104

BUMJIN JOO

Frying up features layer by layer.

ORDER #105

PETER POPESCU

Extra crispy loss functions only.

ORDER #106

DEVESH KUMAR

Grilling weights until they converge.

How?

And amazing materials

Acknowledgements

Prof. Eric Ewing
(last semesters)

Ritambhara Singh
(taught in Spring 2024)

Professor Chen Sun
(taught 2470 in Fall 2024)

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

How?

And making you “do” stuff!

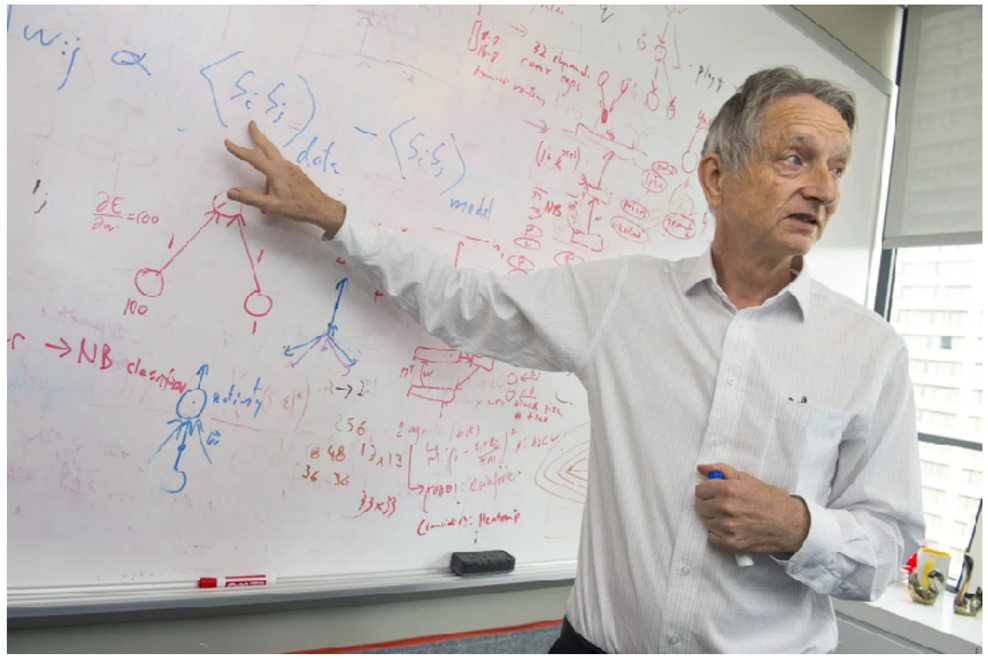

Today’s (light) Goals

- What is Machine Learning?
- What about Deep Learning?
- What is NOT Deep Learning?

What is Machine Learning

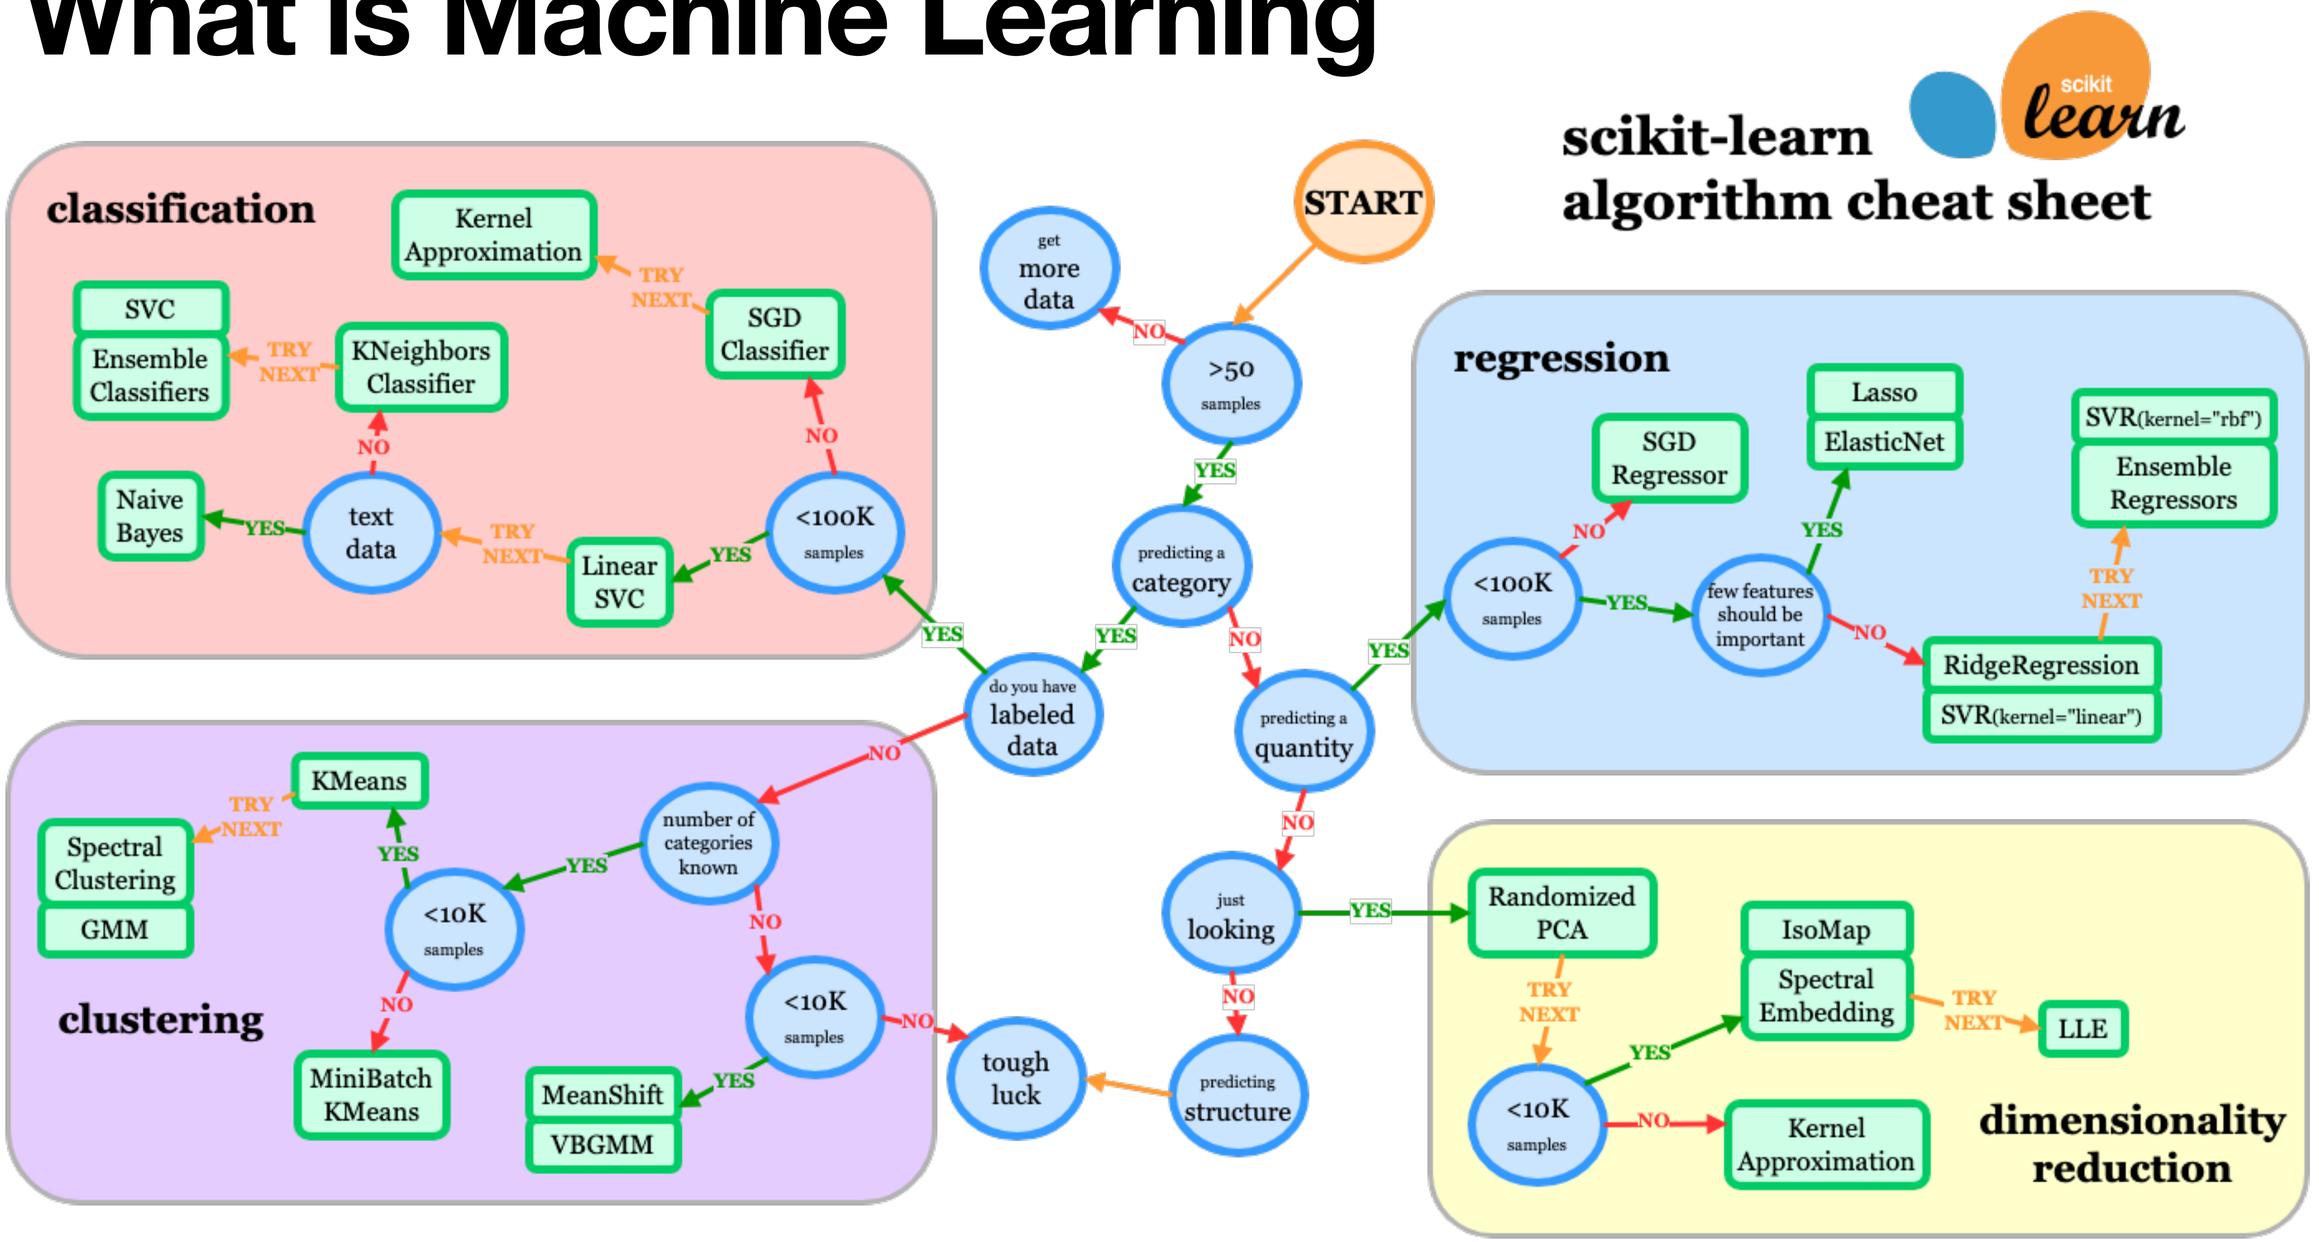

What is Machine Learning

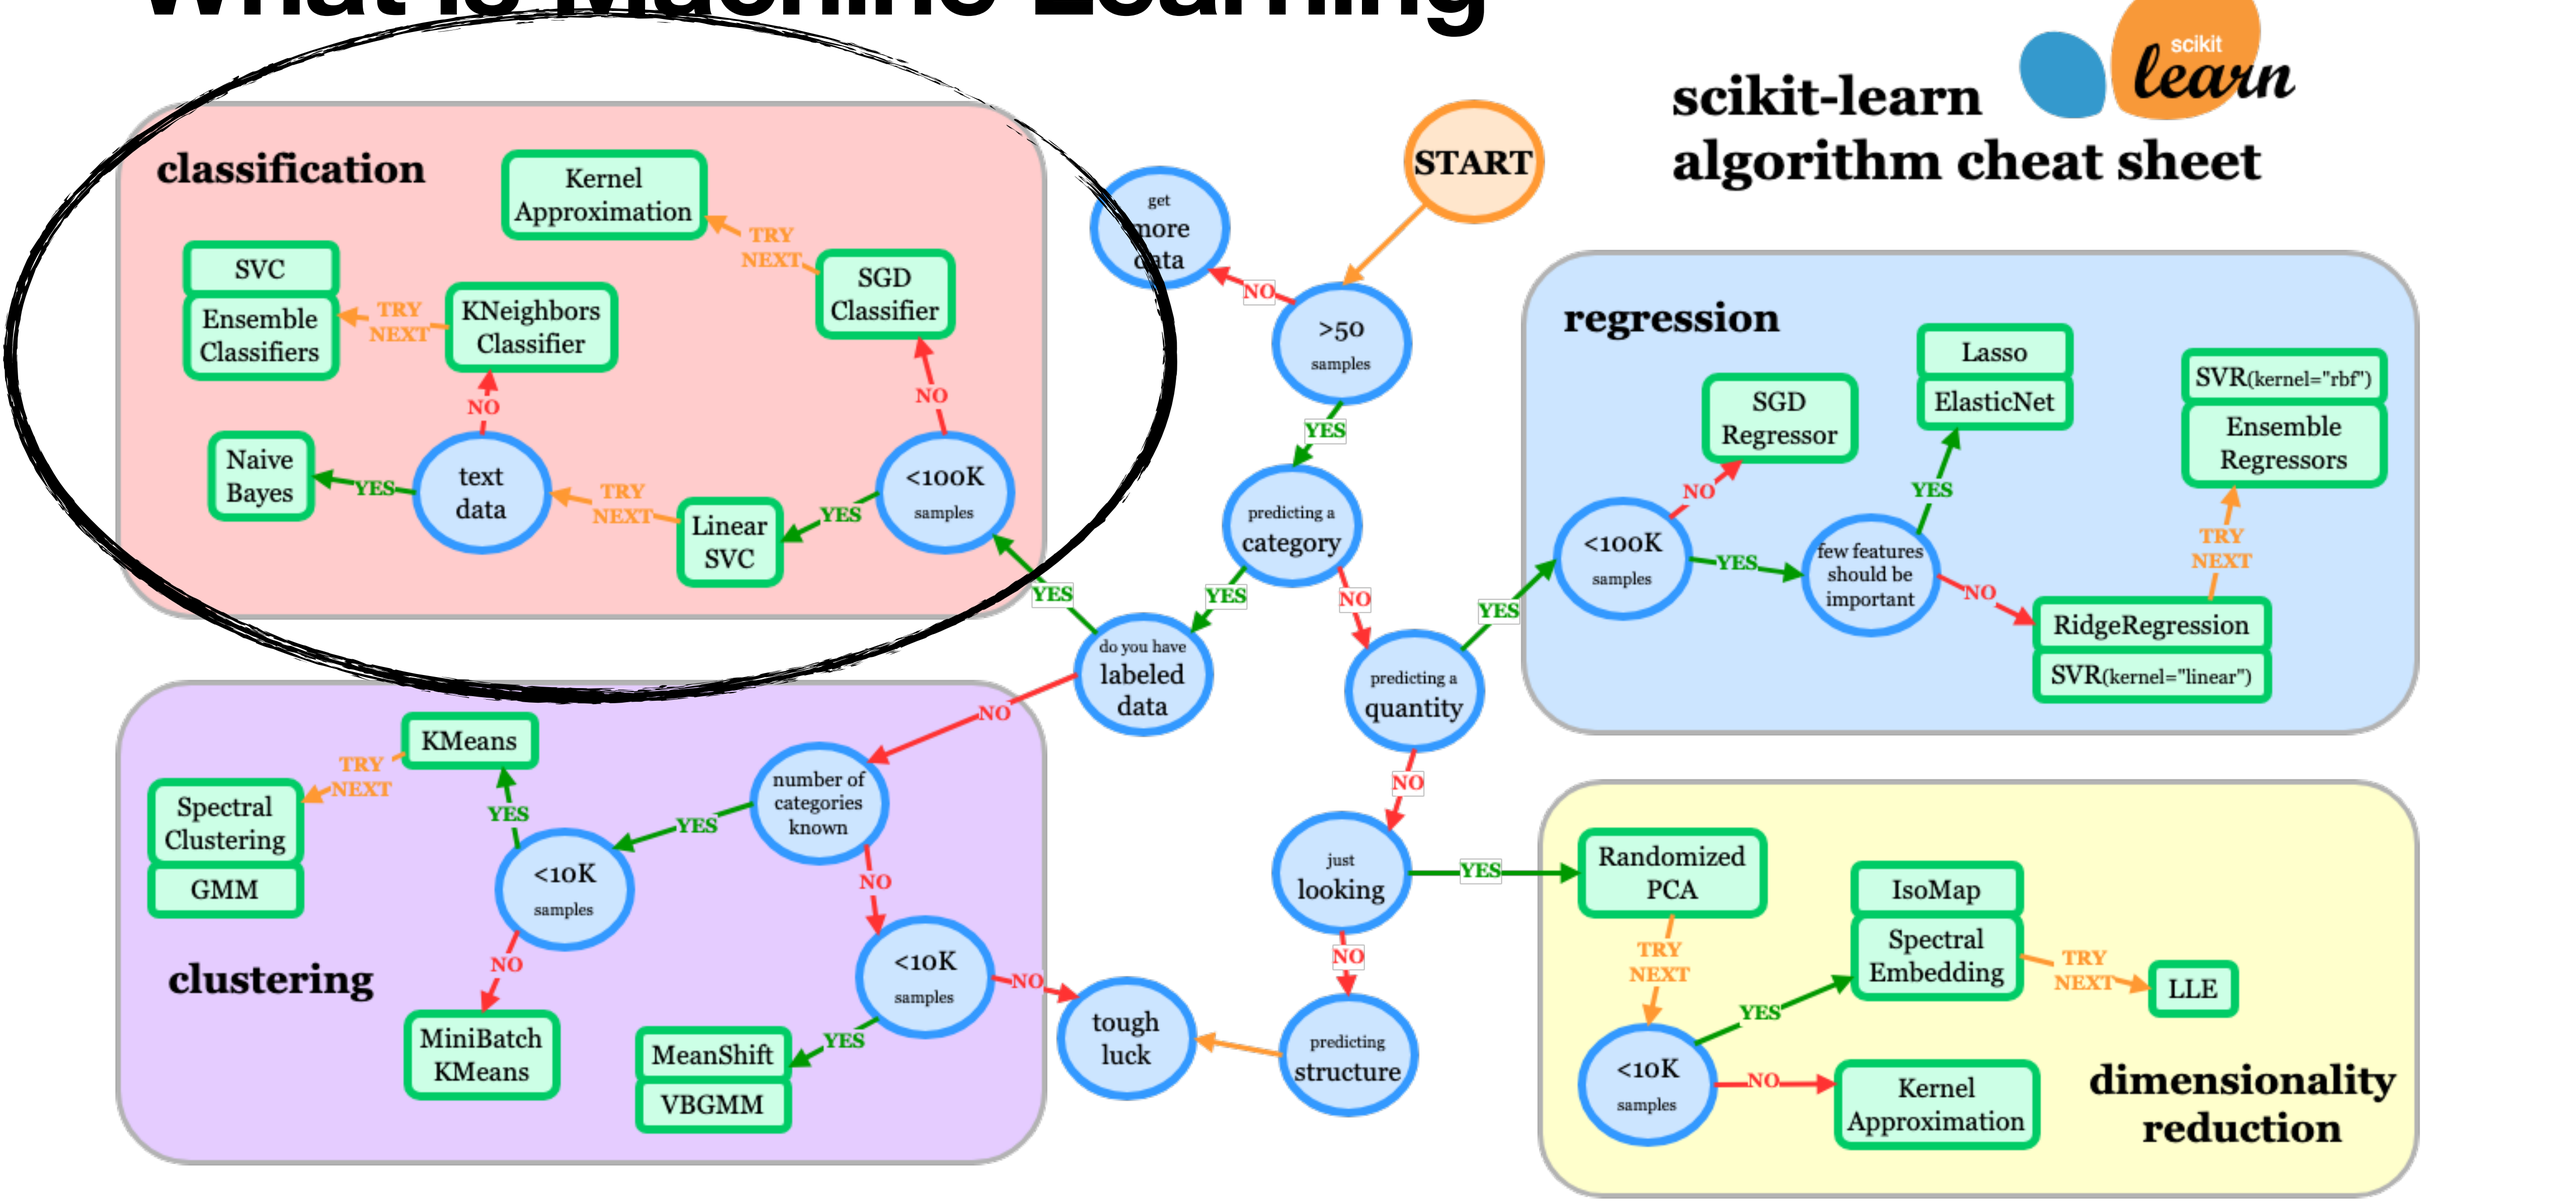

What is Machine Learning/Supervised Learning

- You have a goal (task) in mind

What is Machine Learning/Supervised Learning

- You have a goal (task) in mind

Input: X

Output: Y

”Cooking?”

What is Machine Learning/Supervised Learning

- You have a goal (task) in mind

Input: X

Output: Y

"Cooking?"

Binary classification task

What is Supervised Learning

- You have a goal (task) in mind

Input: X

Output: Y

”Cooking?”

Input: X

f

What is Supervised Learning

- The function/algorithm f is learned from many paired samples

What is Supervised Learning

- The function/algorithm f is learned from many paired samples
- The same principle works for any supervised task

Input: X

I do not want sour cream in my burrito

Learned function: f

Output: Y

No quiero crema agrea en mi burrito

What is Deep Learning?

What is Deep Learning?

- It is a particular choice of f

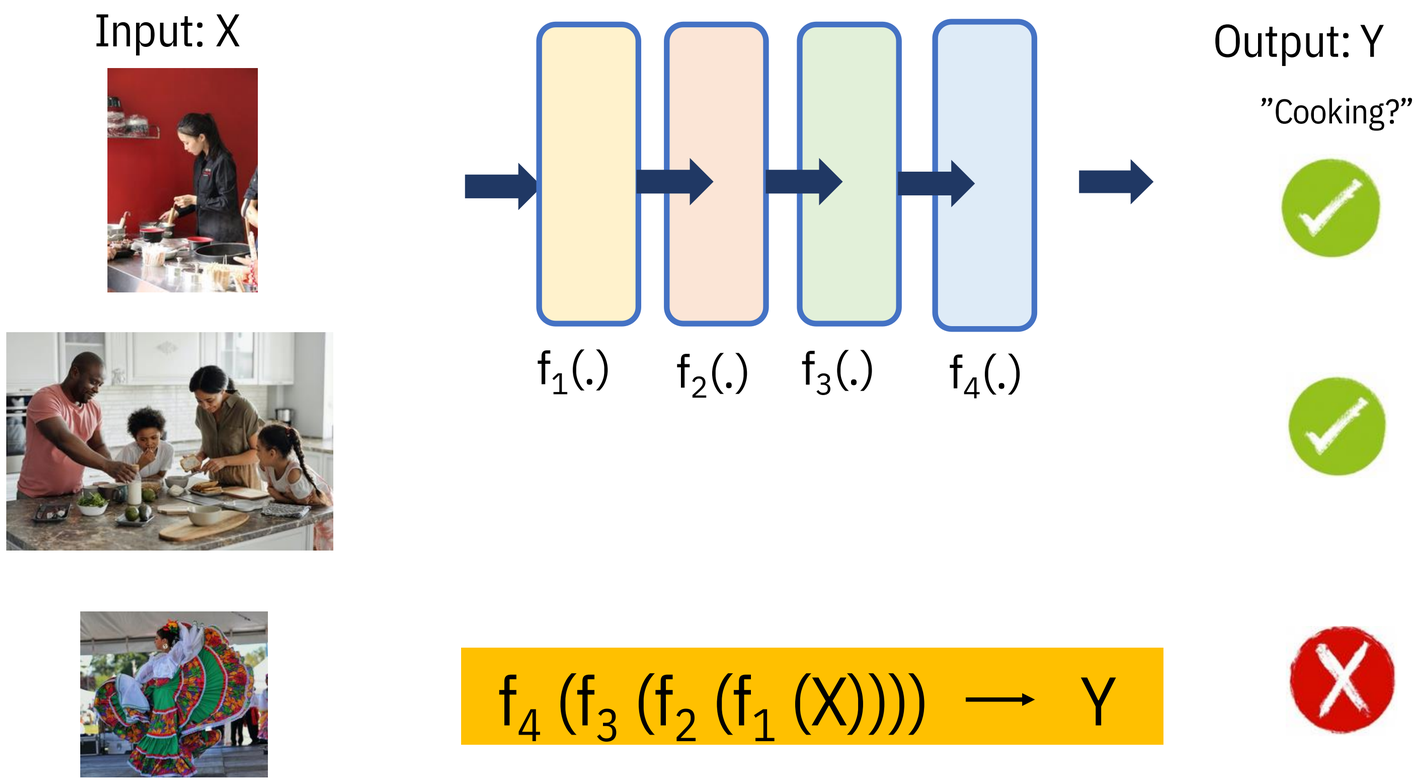

What is NOT Deep Learning?

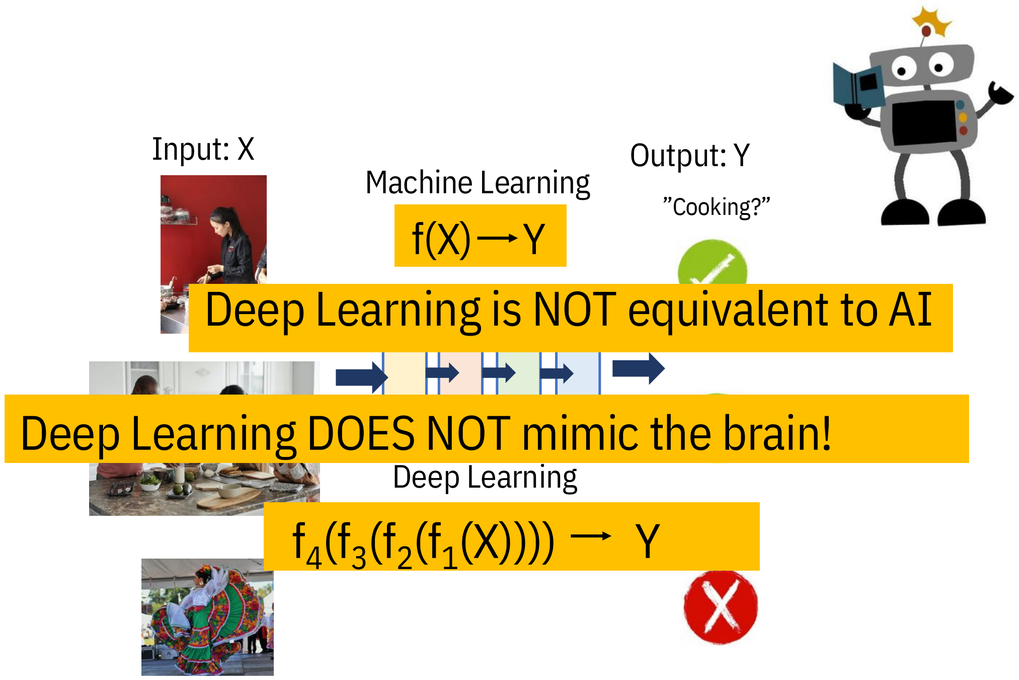

What is NOT Deep Learning?

Input: X

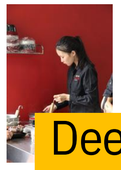

Machine Learning

$f(X) \rightarrow Y$

Output: Y

"Cooking?"

Deep Learning is NOT equivalent to AI

Deep Learning DOES NOT mimic the brain!

Deep Learning

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

What is NOT Deep Learning?

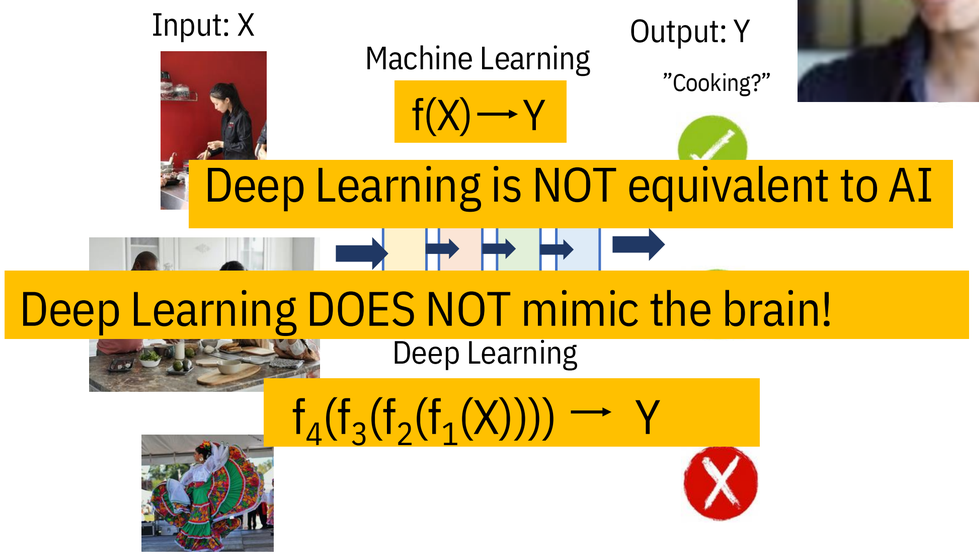

No AI without Deep Learning

Strong alignment between Deep Learning models and Human brain activation patterns

Many “old-school” algorithm can be written as Deep Learning models

Questions?

Turing Award Winner?

Turing Award Winner?

Yes

Yes

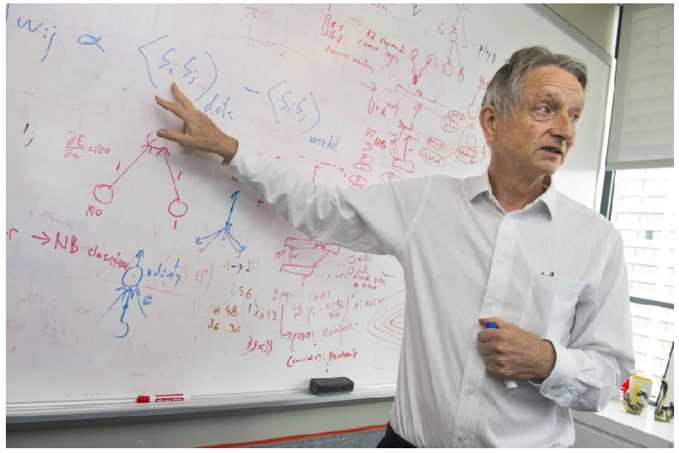

Yes

Turing Award Winner?

Yes

Yes

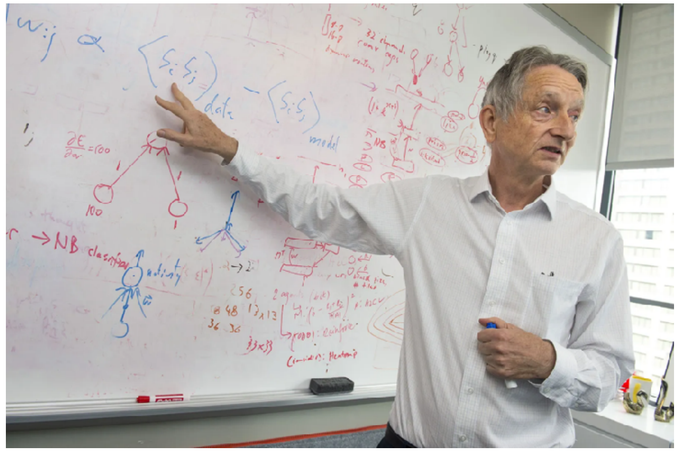

Yes

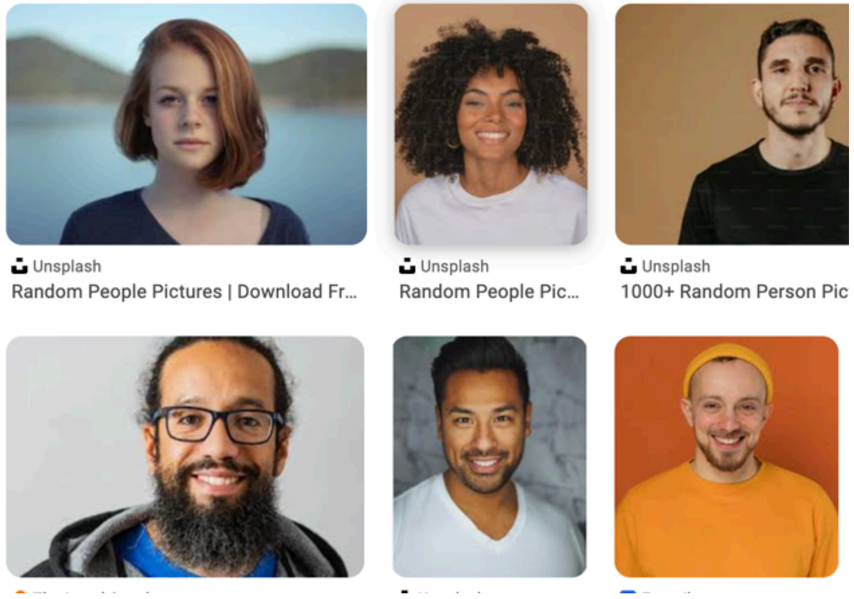

No

Turing Award Winner?

Yes

Yes

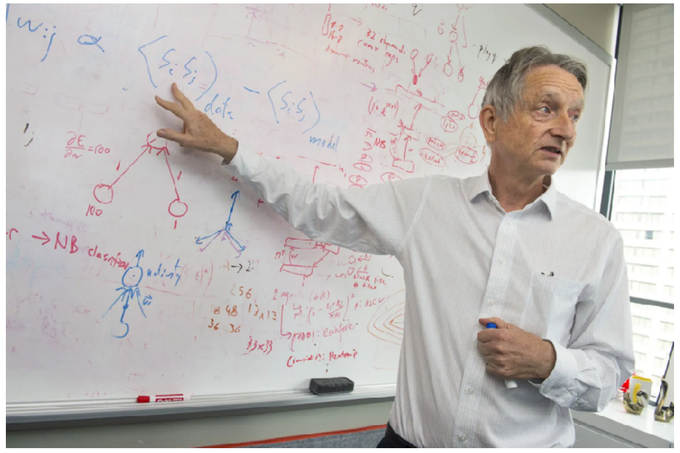

Yes

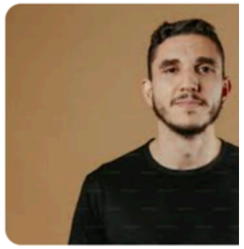

No

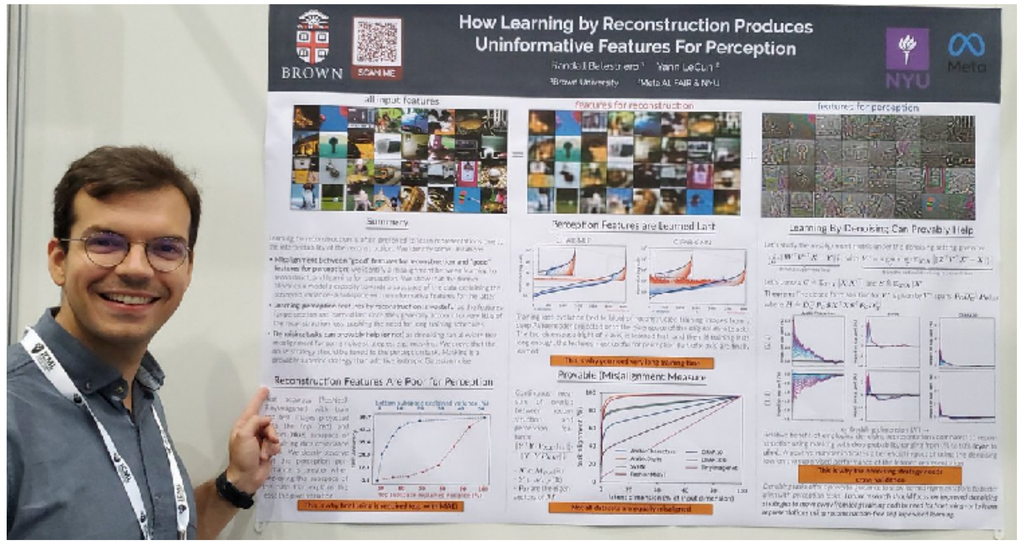

Turing Award Winner?

Yes

Yes

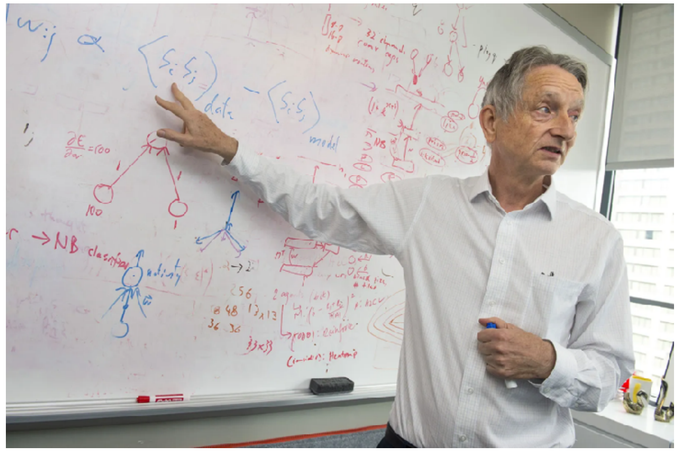

Yes

Unsplash
Random People Pictures | Download Fr...

Unsplash
Random People Pic...

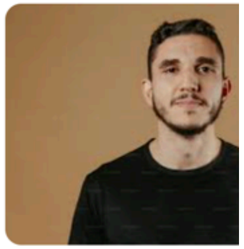

Unsplash
1000+ Random Person Pic

No

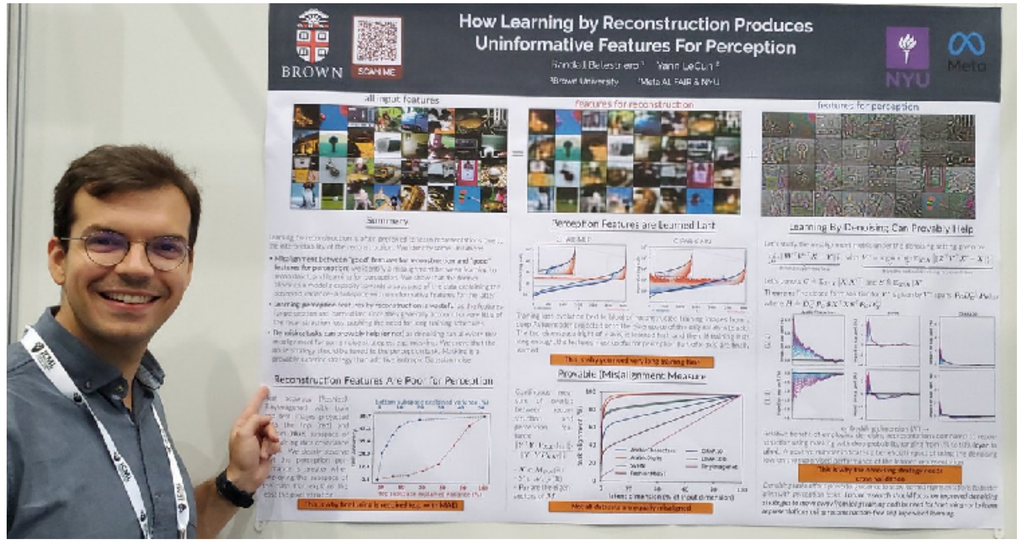

Yes!

Turing Award Winner?

Yes

Yes

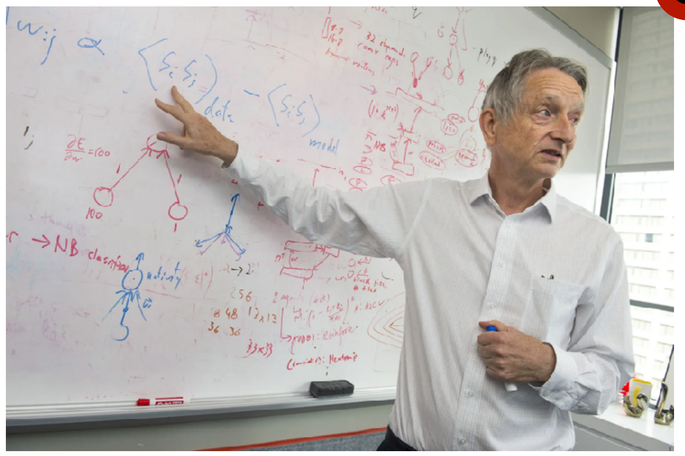

Yes

Unsplash
Random People Pictures | Download Fr...

Unsplash
Random People Pic...

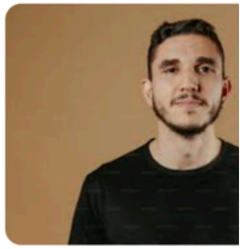

Unsplash
1000+ Random Person Pic

No

Spurious Correlations!

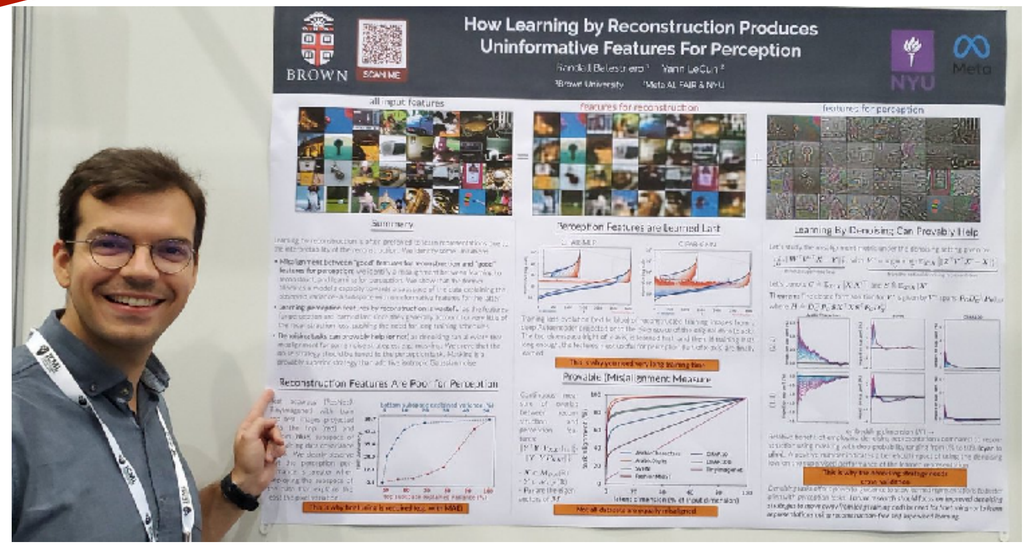

Yes!

Logistic Details: Canvas

Your access to:

- Ed Discussion
- GradeScope
- Weekly quizzes

Logistic Details: Website

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

Logistic Details: Classes

- 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!

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	

Logistic Details: HW + Grading

- Homework 0 released today
  - Review of relevant math and probability concepts
  - Setting up programming environment
  - Graded for completion only

- Assignments
  - Get stencils via Github Classroom
  - Submission via Gradescope
  - Due Wednesdays 10pm

Logistic Details: HW + Grading

- SRC Discussion sections focus on societal impact of deep learning and AI
- Workshop focus on skills/applications of deep learning that are useful for your final project (and your future you)
  - How to read/reproduce academic research paper
  - Write your own visualization tooling
  - Applications (finance, biology, medical)
- Each is an hour long with many time slots offered
- Requires to attend 2 of each

Brown Deep Learning Day

- Course final project
- In-person mini conference!
- Poster sessions and presentations
  - Grouped by theme: e.g. vision, language, robotics, …

Details forthcoming!

Deep Learning Day (Spring 2022)

Questions?

Brainstorm with the person seating next to you: what do you want to solve?