What actually *is* Artificial Intelligence?

And how does it relate to astronomy?

Nick Konz, ERIRA 2024

https://imgs.xkcd.com/comics/machine_learning_2x.png

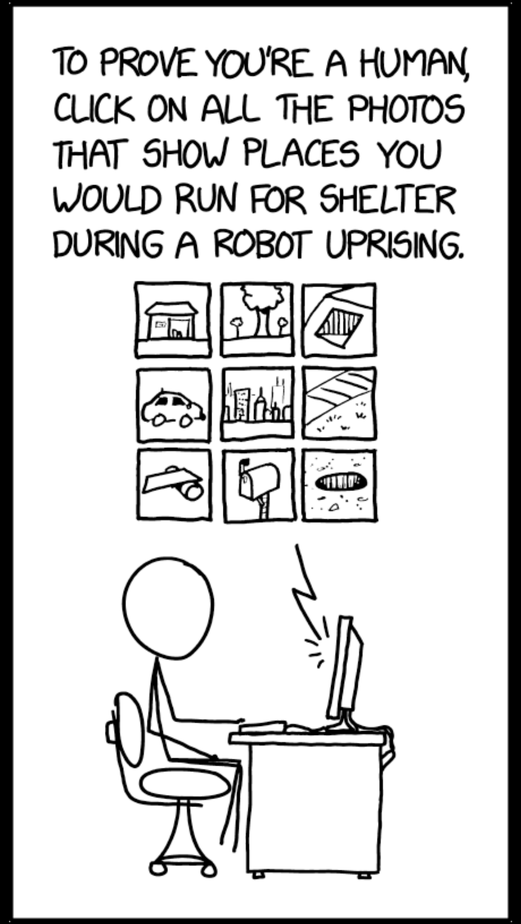

https://imgs.xkcd.com/comics/machine_learning_captcha_2x.png

About me

Duke | DEPARTMENT OF Electrical & Computer Engineering

http://people.ee.duke.edu/~mbrooke/ECE-HorizontalLogo-Print-Blue.jpg

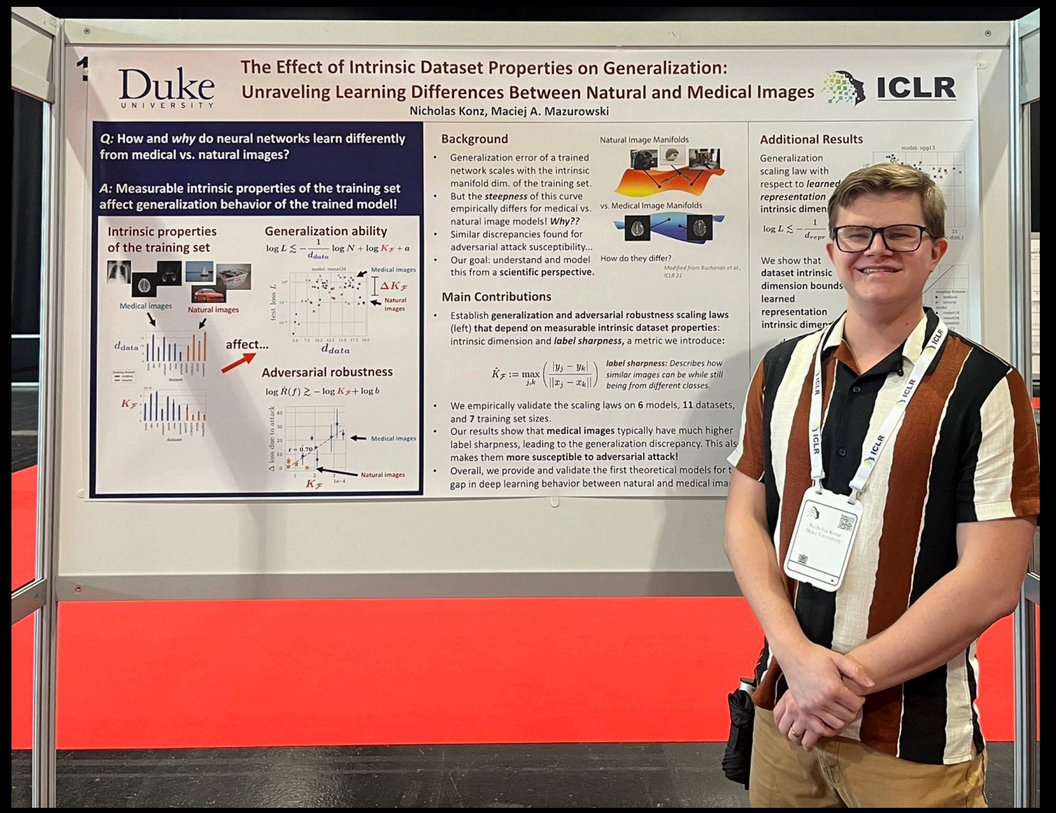

AI in popular fiction

MOVIEWALLPAPERS.NET

https://live.staticflickr.com/5052/5392319221_b622a82d0a_b.jpg

https://static.wikia.nocookie.net/ironman/images/d/d9/Ultron_EW_Poster.png/revision/latest?cb=20191203212946

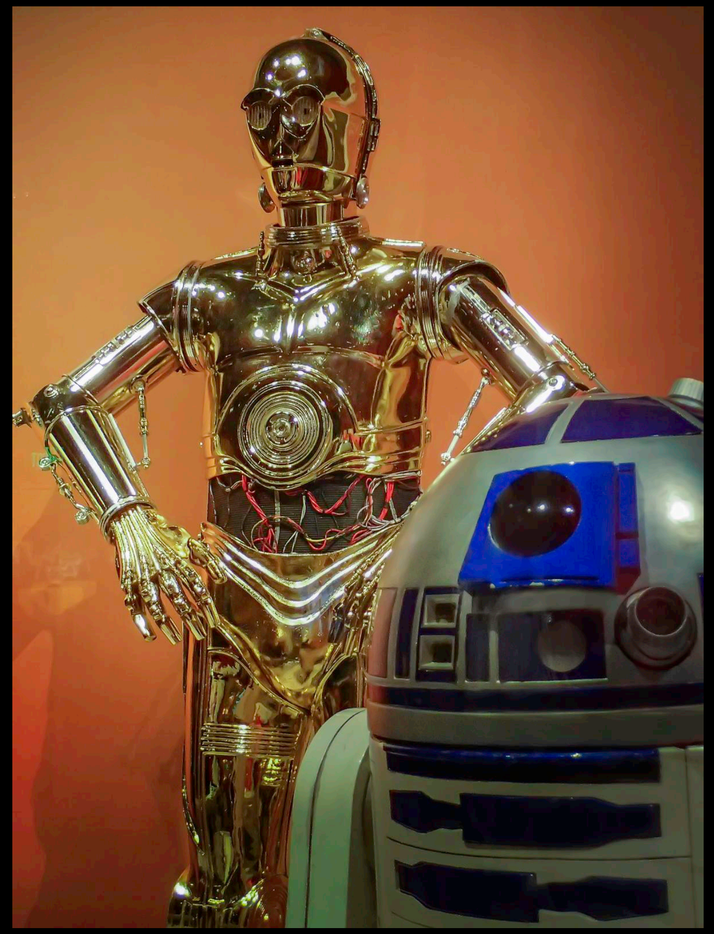

https://www.flickr.com/photos/mharrsch/16446792154

1. how can we encode common sense?
2. how can humans learn from so few examples in totally new contexts?
3. how can knowledge be represented best and distributed between many different systems?

Rapid Advancement, and the Bitter Lesson

http://www.incompleteideas.net/IncIdeas/BitterLesson.html

The Bitter Lesson

Rich Sutton

March 13, 2019

The biggest lesson that can be read from 70 years of AI research is that general methods that leverage computation are ultimately the most effective, and by a large margin. The ultimate reason for this is Moore's law, or rather its generalization of continued exponentially falling cost per unit of computation. Most AI research has been conducted as if the computation available to the agent were constant (in which case leveraging human knowledge would be one of the only ways to improve performance) but, over a slightly longer time than a typical research project, massively more computation inevitably becomes available. Seeking an improvement

1.76 trillion

GPT-4 is the latest language model developed by OpenAI, and its parameter count is a staggering **1.76 trillion**. By comparison, GPT-3, the current largest language model, has 175 billion parameters, while GPT-2 has 1.5 billion parameters. Mar 19, 2023

The Cost of Training GPT-4

OpenAI has revealed that it cost them $100 million and took 100 days, utilizing 25,000 **NVIDIA A100** GPUs. Oct 8, 2023

NVIDIA Tesla A100 Ampere 40 GB Graphics Processor Accelerator - PCIe 4.0 x16 - Dual Slot

Visit the Dell Store

4.4 ★★★★½ 3 ratings | Search this page

$8,099.99

Or **$320.57**/mo (48 mo). Select from 1 plan

Eligible for Return, Refund or Replacement within 30 days of receipt

Rapid Advancement, and the Bitter Lesson

how can we encode common sense?

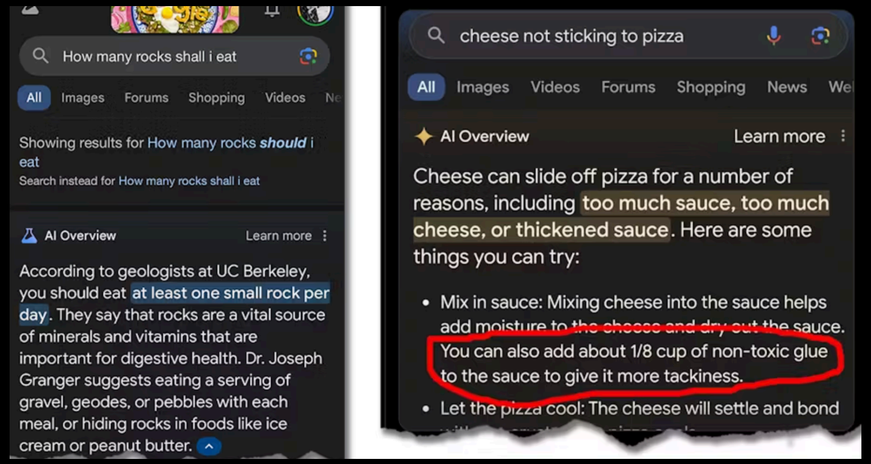

how can humans learn from so few examples in totally new contexts?

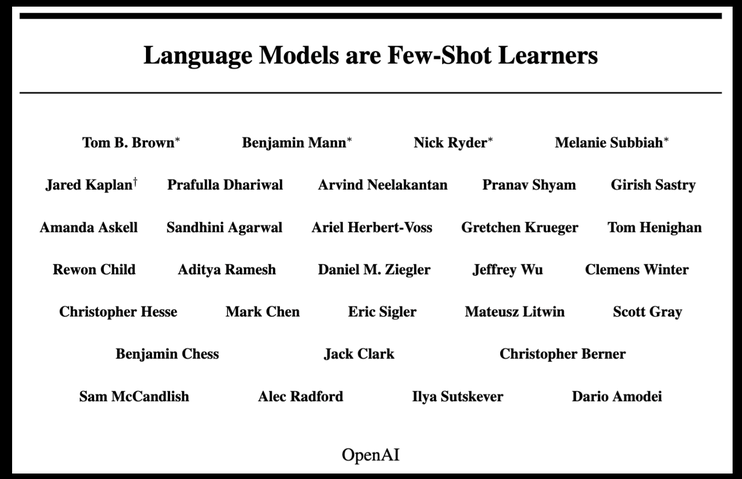

Language Models are Few-Shot Learners

Tom B. Brown* Benjamin Mann* Nick Ryder* Melanie Subbiah*

Jared Kaplan† Prafulla Dhariwal Arvind Neelakantan Pranav Shyam Girish Sastry

Amanda Askell Sandhini Agarwal Ariel Herbert-Voss Gretchen Krueger Tom Henighan

Rewon Child Aditya Ramesh Daniel M. Ziegler Jeffrey Wu Clemens Winter

Christopher Hesse Mark Chen Eric Sigler Mateusz Litwin Scott Gray

Benjamin Chess Jack Clark Christopher Berner

Sam McCandlish Alec Radford Ilya Sutskever Dario Amodei

OpenAI

how can knowledge be represented best and distributed between many different systems?

So how does AI actually work?

Classical AI vs. Modern AI

Good old fashioned AI: rule-based/“expert” systems

“Expert-designed” decision rules/ algorithms

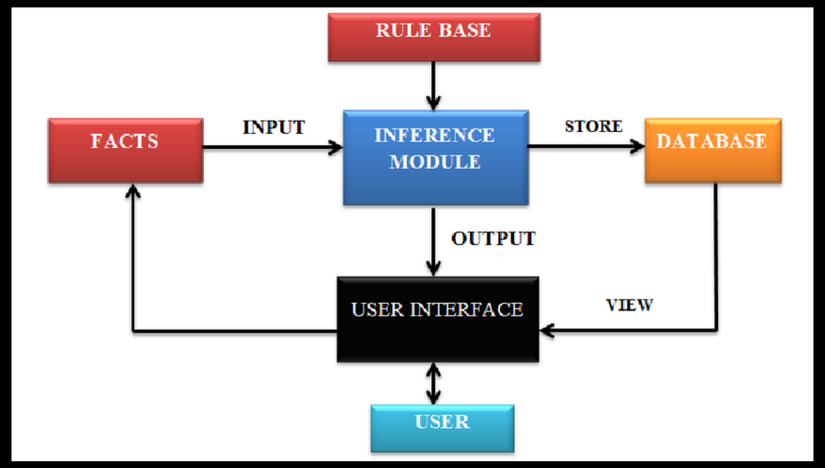

https://studiousguy.com/wp-content/uploads/2021/07/Rule-based-Production-Systems.jpeg

“Expert-designed” knowledge representations

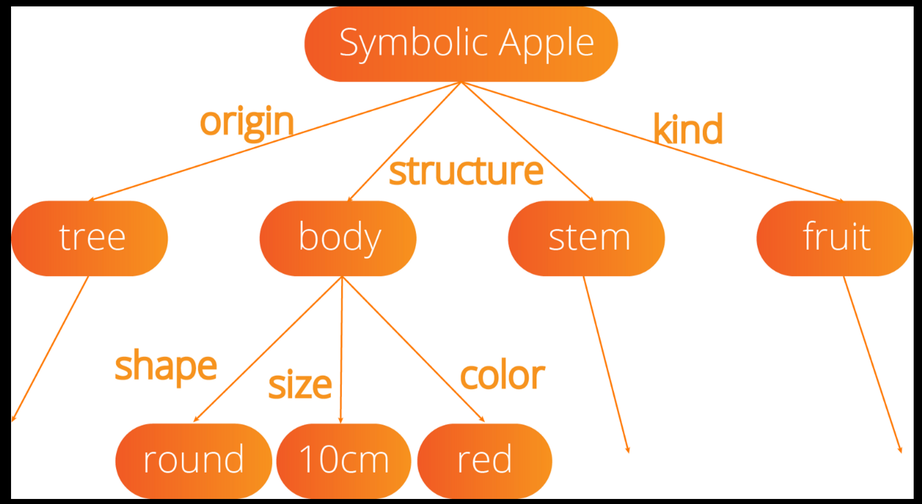

https://miro.medium.com/max/1838/1*xYZrAXZi6Iq3c1Z_CxW6ZA.png

Classical AI vs. Machine Learning vs. Deep Learning

Manually-designed

vs.

Automatically learned from data (Modern AI)

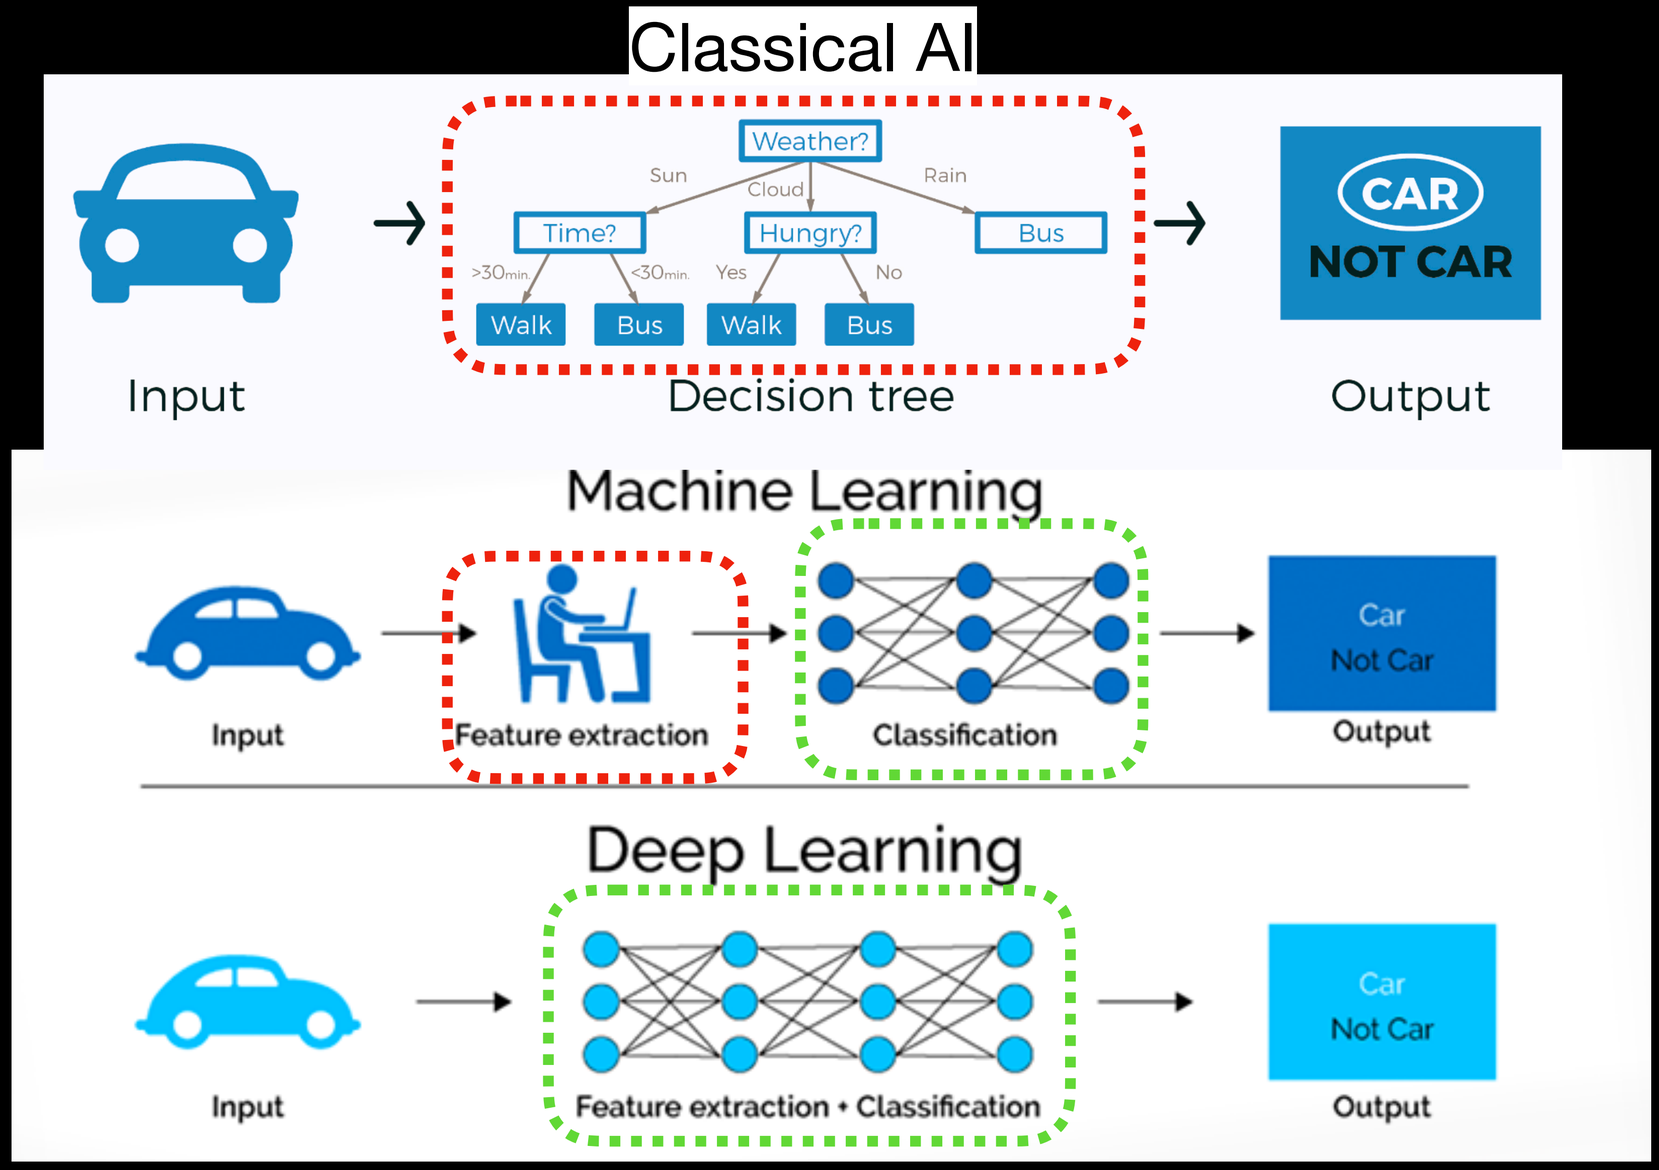

https://miro.medium.com/max/1838/0*KJOhHjSJXo0RMFuC.png

https://smartpeme.depo.gal/documents/10180/0/DL2+noticia.jpg/48fcaf55-dacb-4ff4-98d1-825dffdeca6f?t=1581678754738

Why “*deep*” learning?

https://miro.medium.com/max/1838/0*KJOhHjSJXo0RMFuC.png

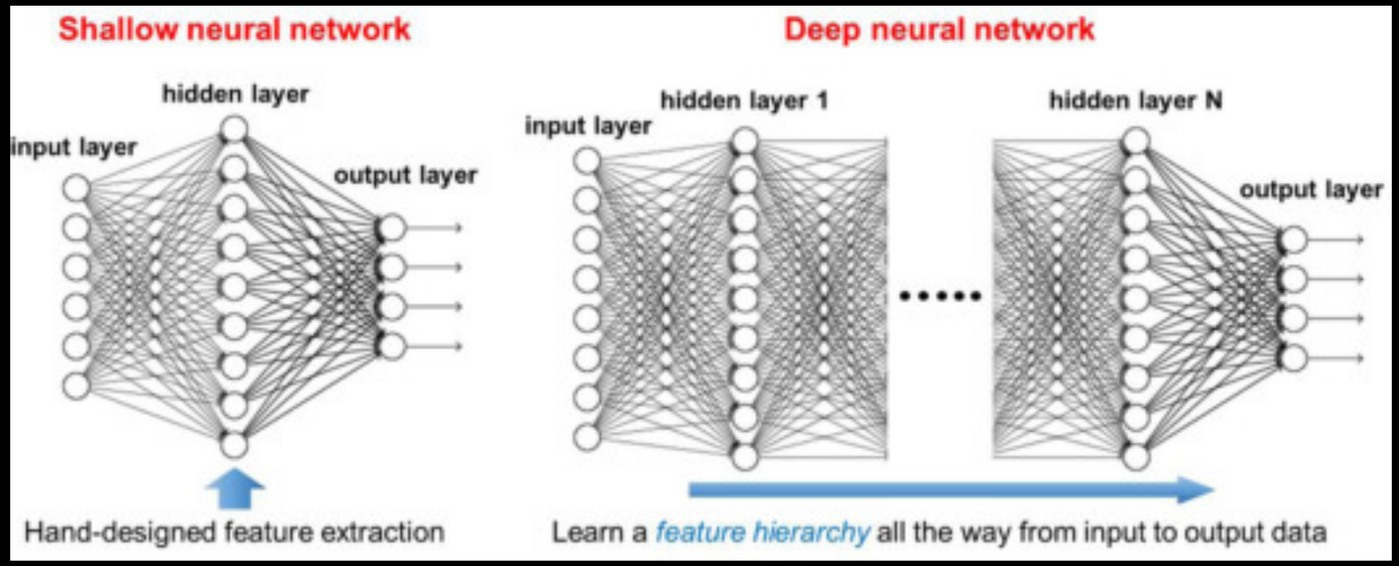

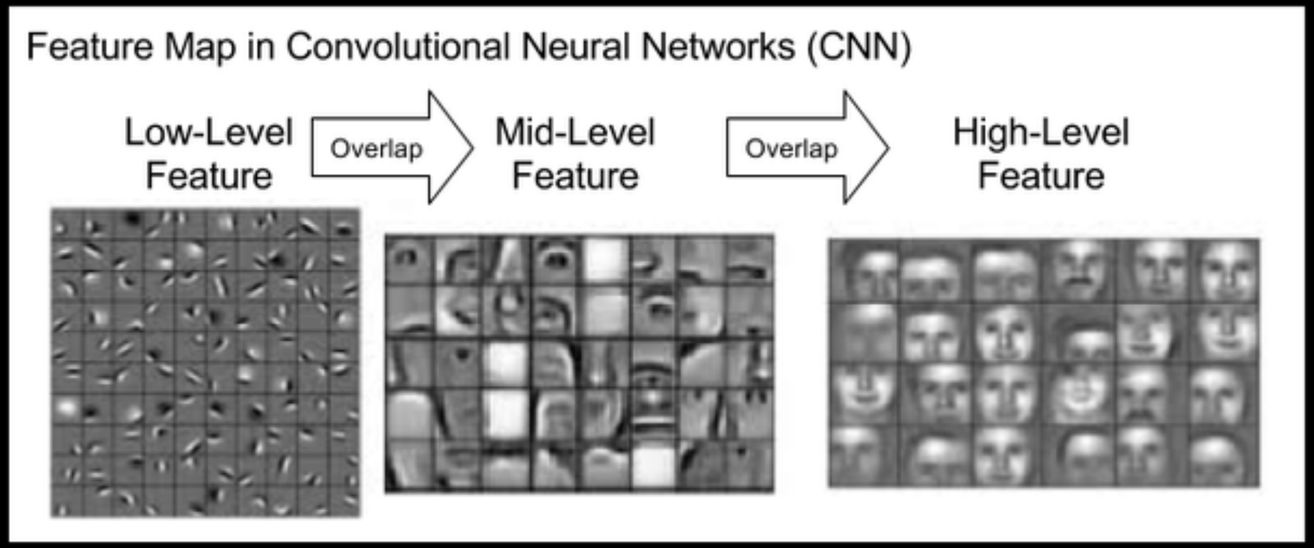

https://miro.medium.com/max/1838/0*KJOhHjSJXo0RMFuC.png

How does visual feature recognition actually work?

Visual Recognition with Classical AI

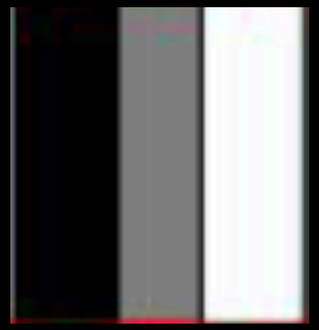

$$= \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$$

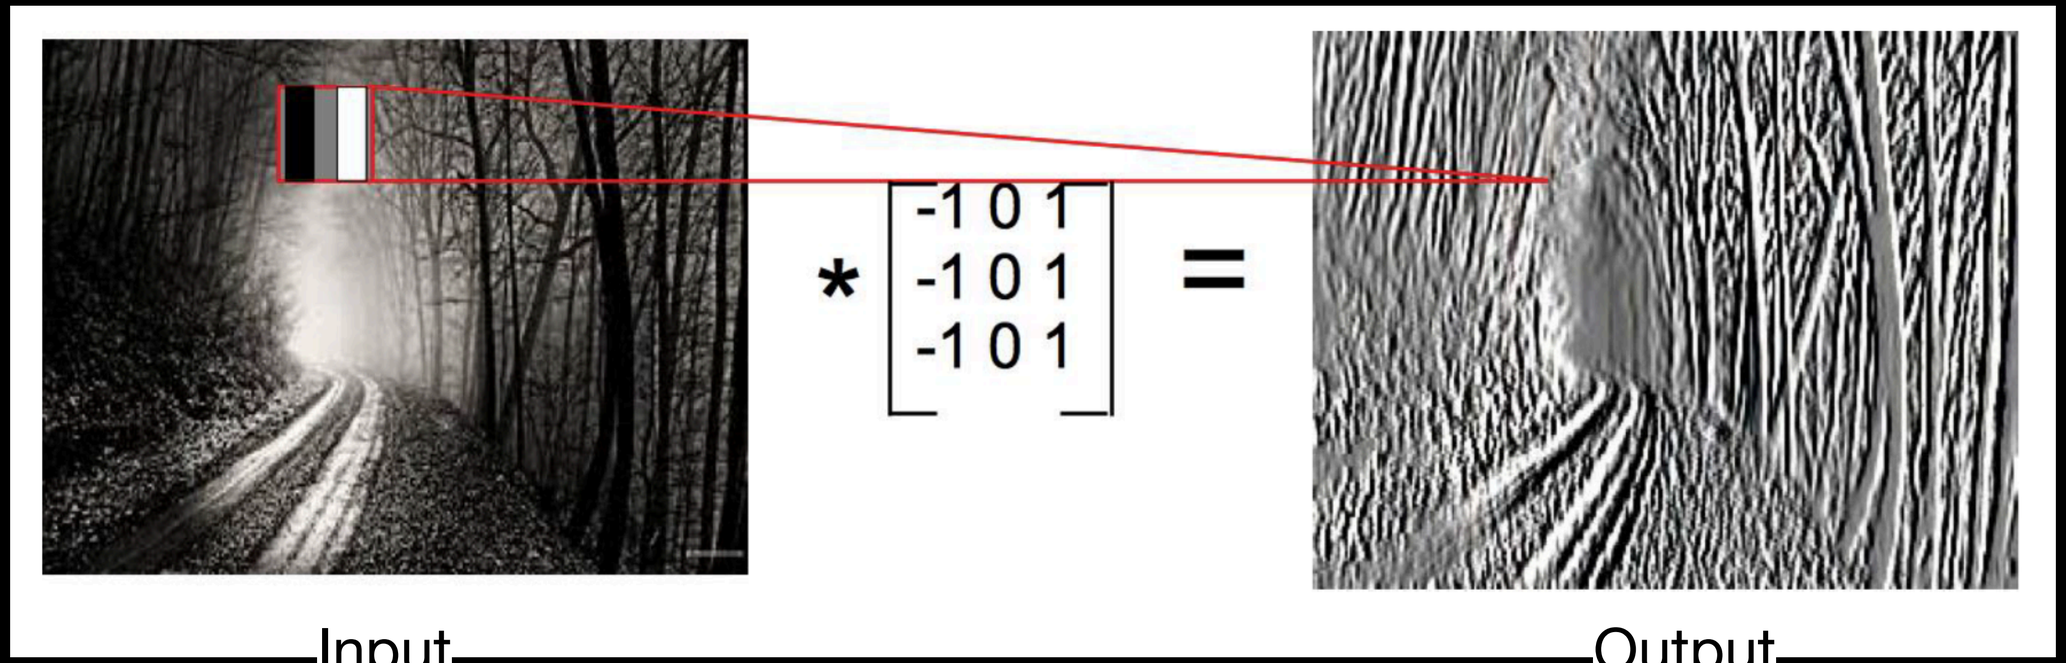

credit: ECE 685D lecture notes, Fall 2021

Visual Recognition with Deep Learning

Input

https://miro.medium.com/max/1026/1*Ji5QhY9QXBlpNNLH4qAcNA.png

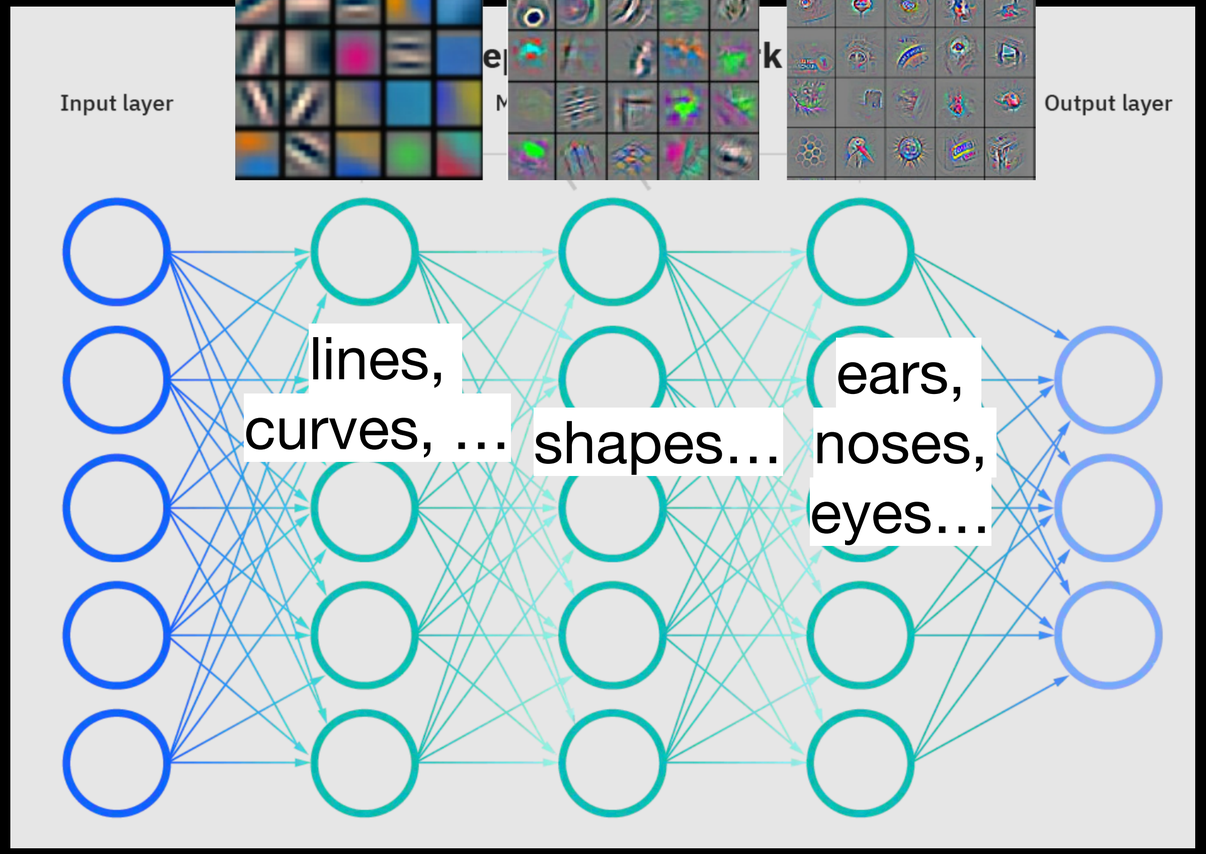

https://velog.velcdn.com/images/tjswodud/post/a5149a8a-18e4-4fd2-9c39-e7a14a9a4a57/image.png

Prediction:
95% probability of cat

4% probability of squirrel

1% probability of dog

How do neural networks actually “learn”?

How does the *learning* actually work in deep learning?

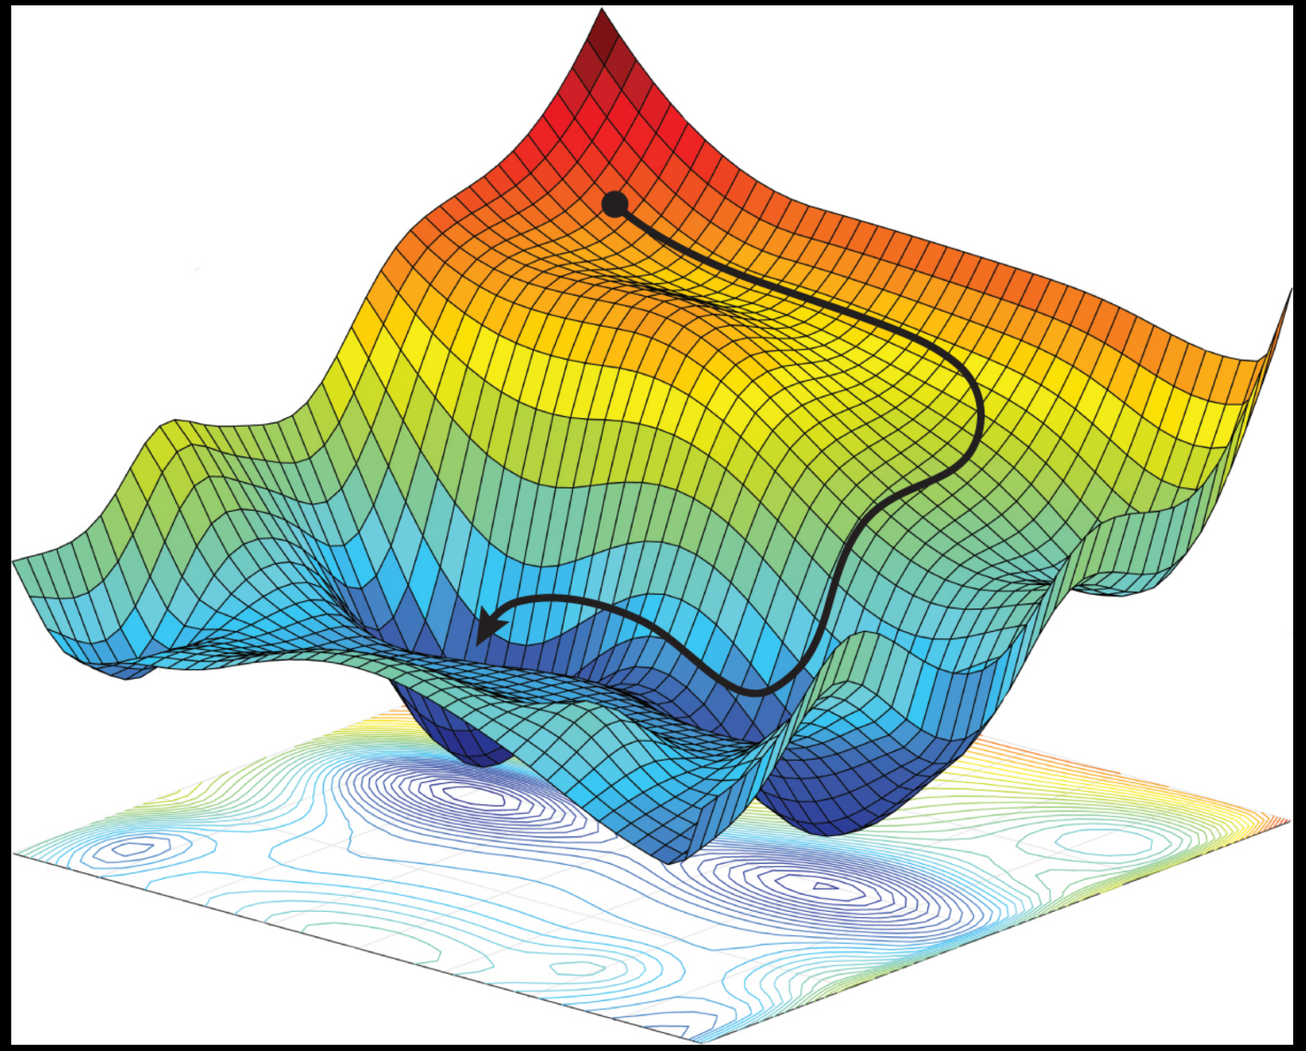

https://www.researchgate.net/profile/Alexander-Amini/publication/325142728/figure/fig1/AS:766109435326465@1559666131320/Non-convex-optimization-We-utilize-stochastic-gradient-descent-to-find-a-local-optimum.jpg

From Linear Regression to Neural Networks

$$y = W\vec{x} + \vec{b}$$

single linear “layer”

$$y = \sigma(W\vec{x} + \vec{b})$$

add a nonlinearity

$$y = \sigma(W_2\sigma(W_1\vec{x} + \vec{b}_1) + \vec{b}_2)$$

add more linear layers
+ nonlinearities…

From Linear Regression to Neural Networks

A single-layer (shallow) neural network

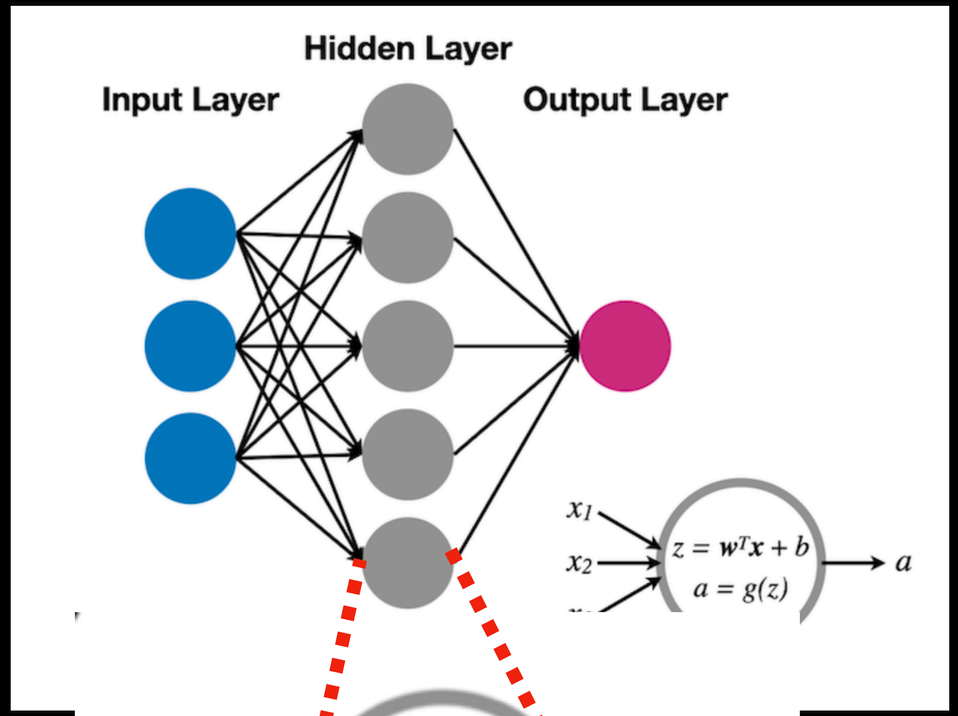

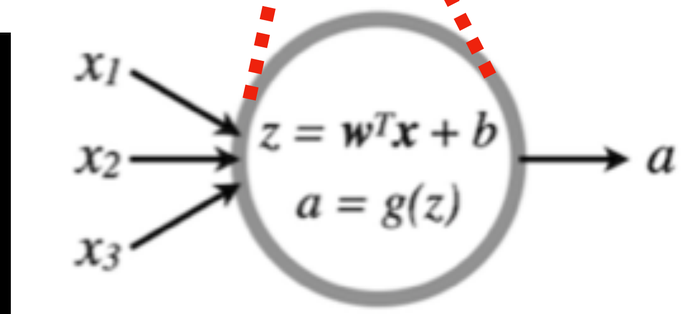

One *neuron*

A multi-layer (deep) neural network

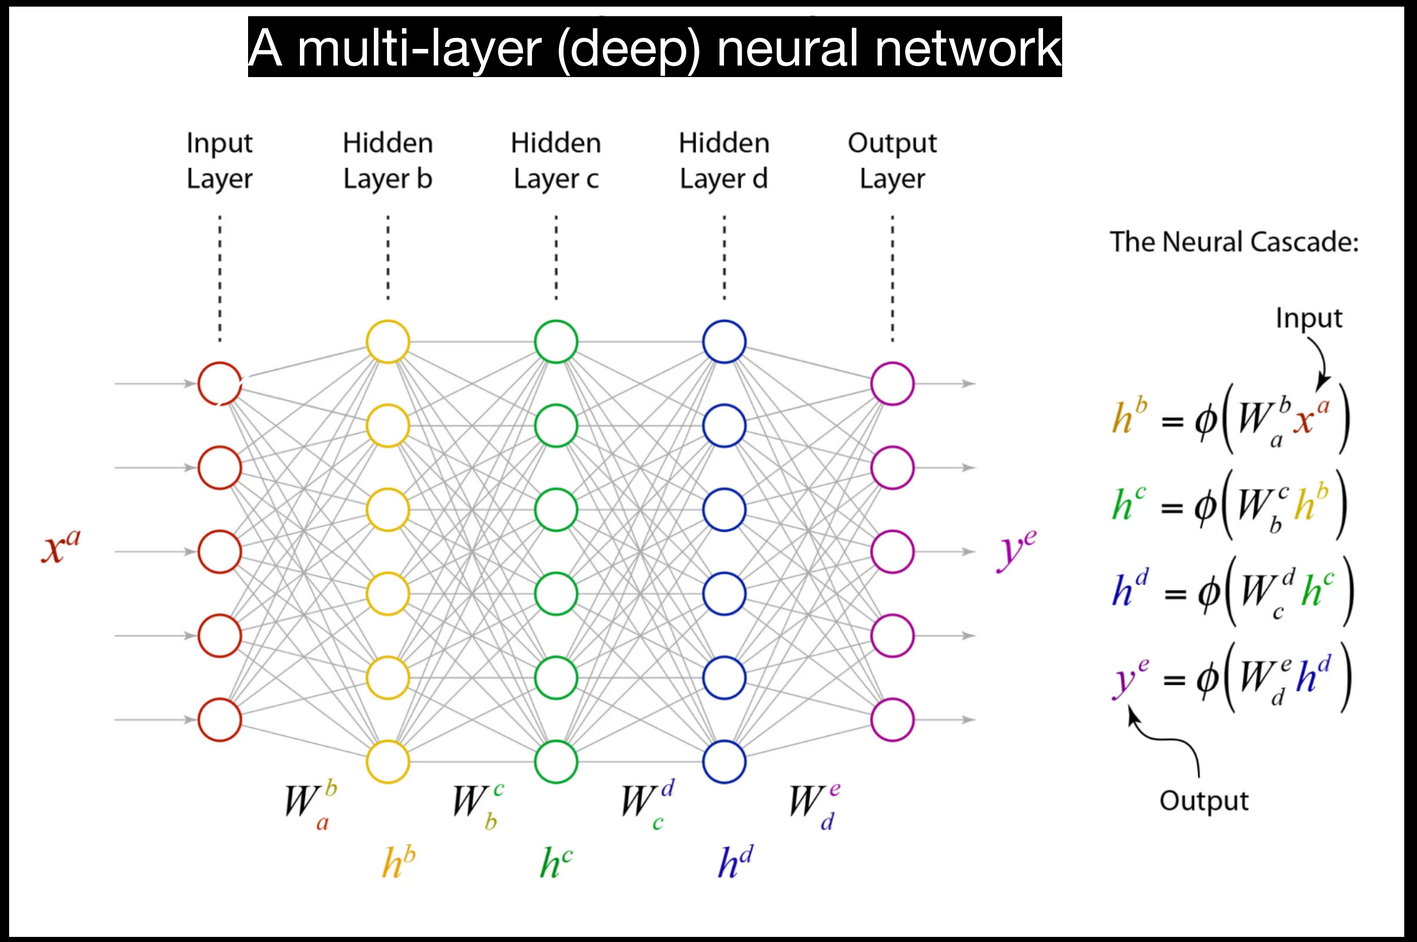

https://scipython.com/static/media/uploads/blog/shallow-neural-net/snn.png

https://galileo-unbound.blog/2022/04/18/post-modern-machine-learning-the-deep-revolution/

Relating back to astronomy

Big Data enables Big Models

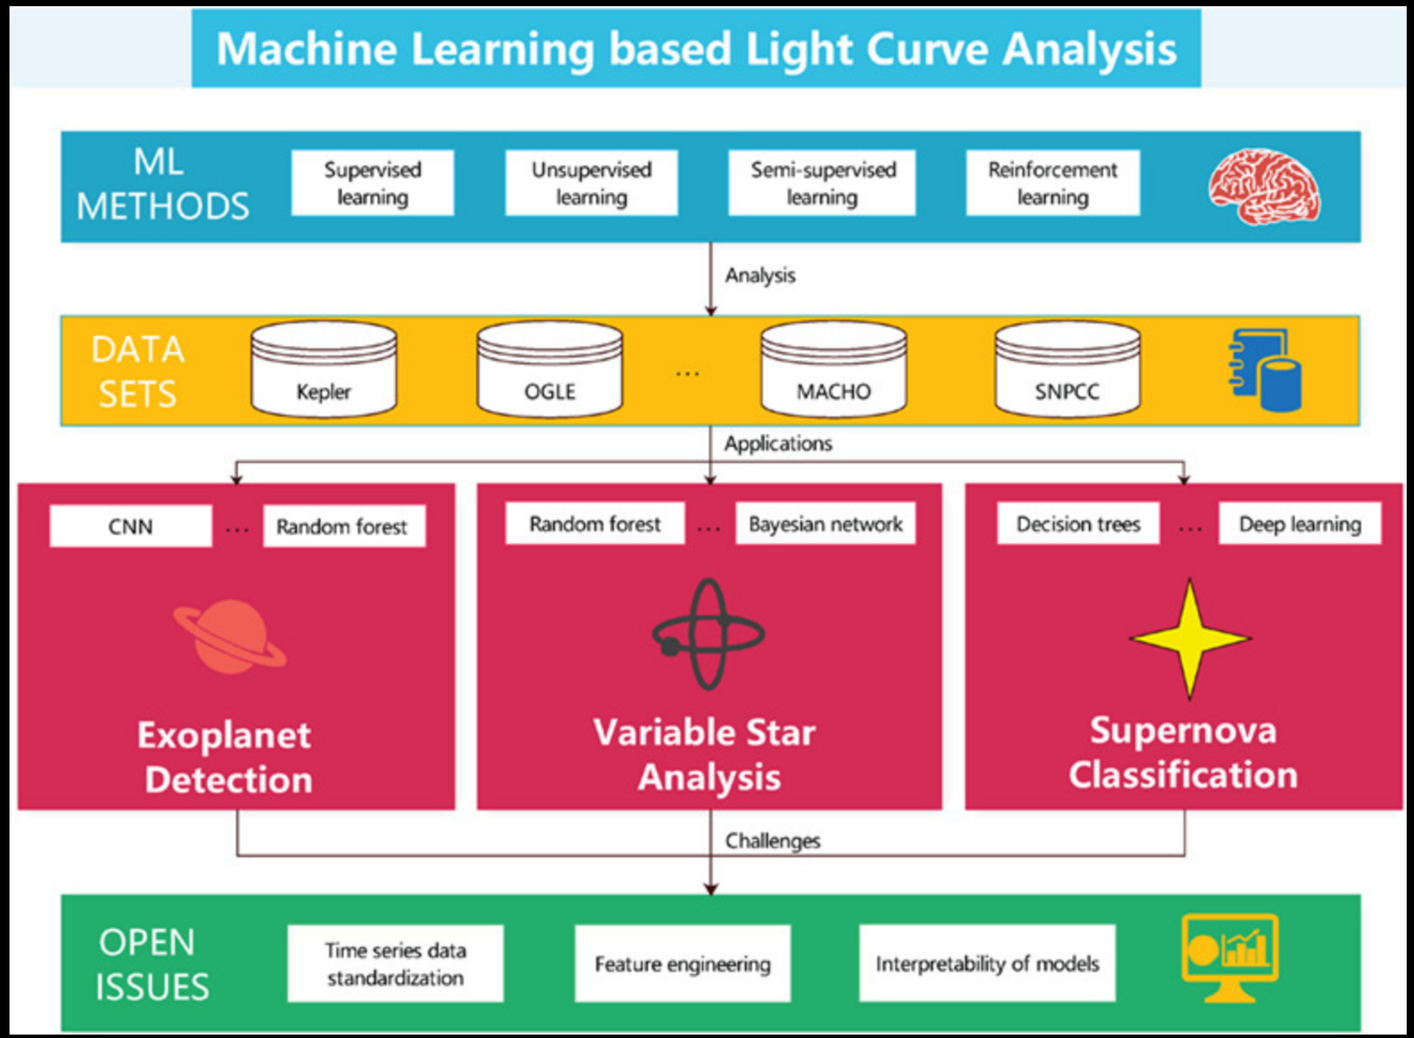

Yu, Ce, et al. "A survey on machine learning based light curve analysis for variable astronomical sources." *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 11.5 (2021): e1425.

Automated galaxy cataloging from surveys

Physics Letters B

Deep learning at scale for the construction of galaxy catalogs in the Dark Energy Survey

Asad Khan [a,b,*], E.A. Huerta [a,c], Sibo Wang [a], Robert Gruendl [a,c], Elise Jennings [d], Huihuo Zheng [d]

[a] *National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA*
[b] *Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA*
[c] *Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA*
[d] *Argonne National Laboratory, Leadership Computing Facility, Lemont, IL 60439, USA*

Khan, Asad, et al. "Deep learning at scale for the construction of galaxy catalogs in the Dark Energy Survey." *Physics Letters B* 795 (2019): 248-258.

Radio astronomy image de-noising

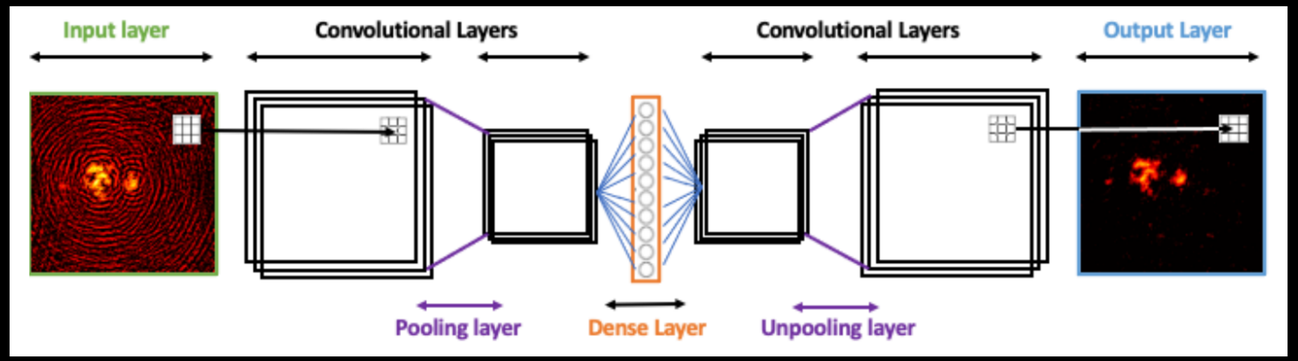

noisy image input de-noised output

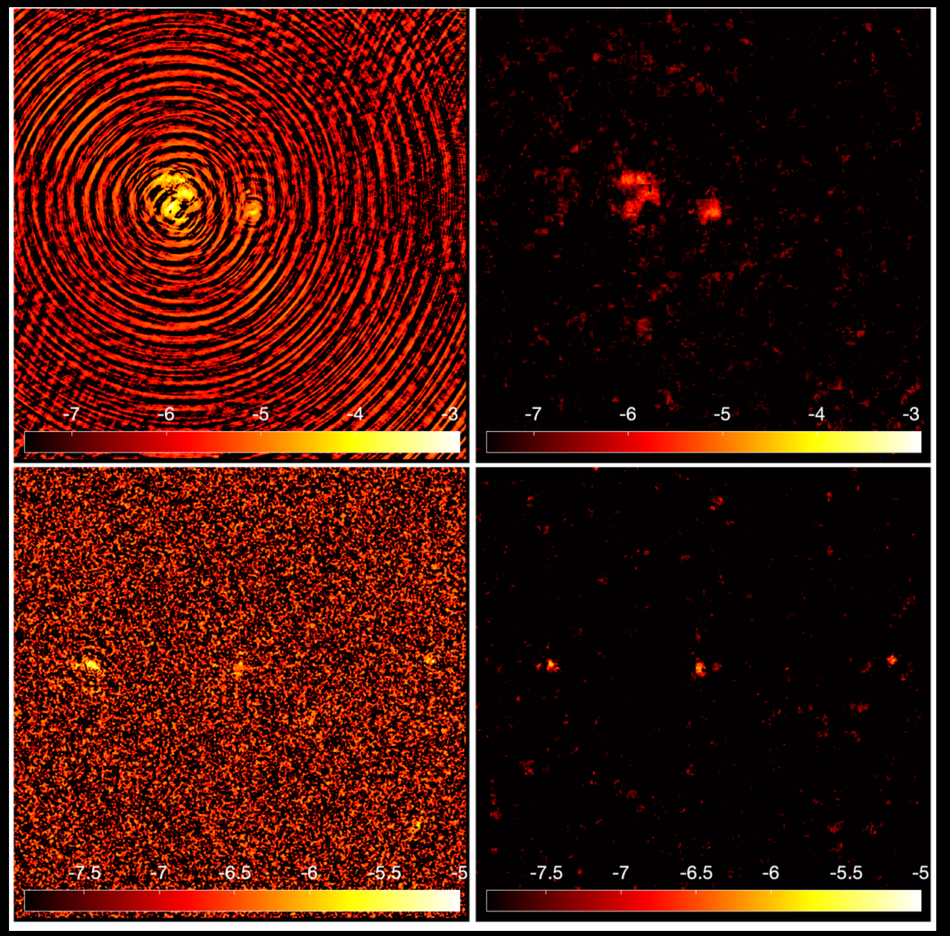

Gheller, Claudio, and Franco Vazza. "Convolutional deep denoising autoencoders for radio astronomical images." Monthly Notices of the Royal Astronomical Society 509.1 (2022): 990-1009.

And many others…

Including my team project for this year 😄

The Good, Bad and Spooky Capabilities of Modern AI/Deep Learning

The Good: Deep learning has created *huge* advancements in:

1. Computer Vision

2. image/art/text/etc. generation

https://www.researchgate.net/profile/Dae-Young-Kang/publication/346091812/figure/fig2/AS:979480482938881@1610537753860/Algorithms-that-won-the-ImageNet-Large-Scale-Visual-Recognition-Challenge-ILSVRC-in.png

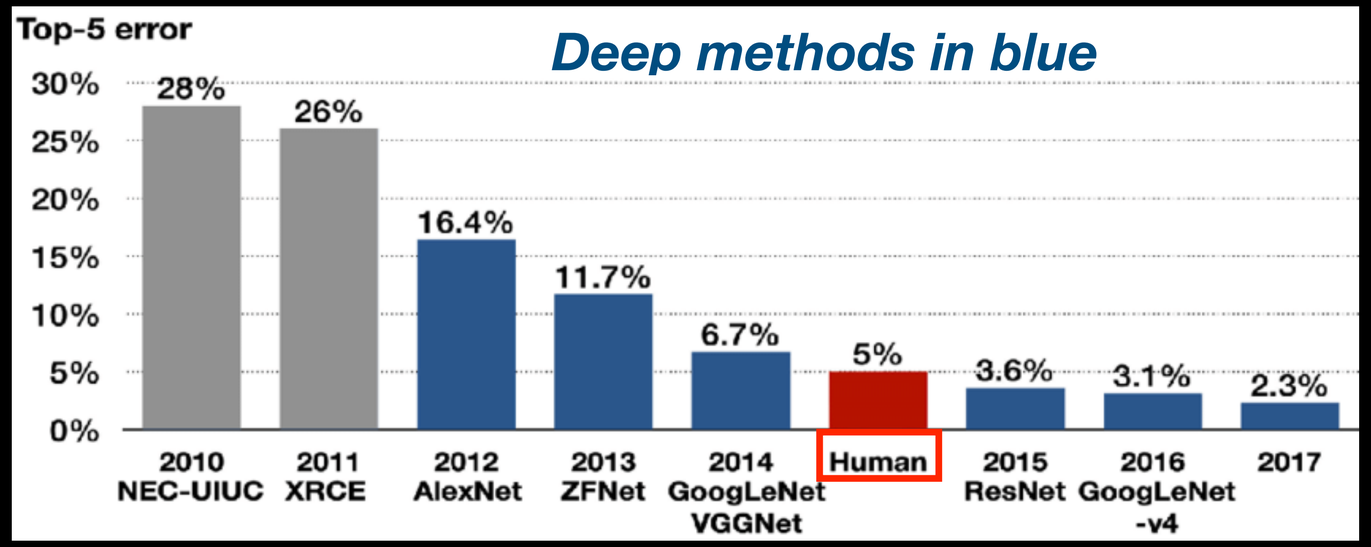

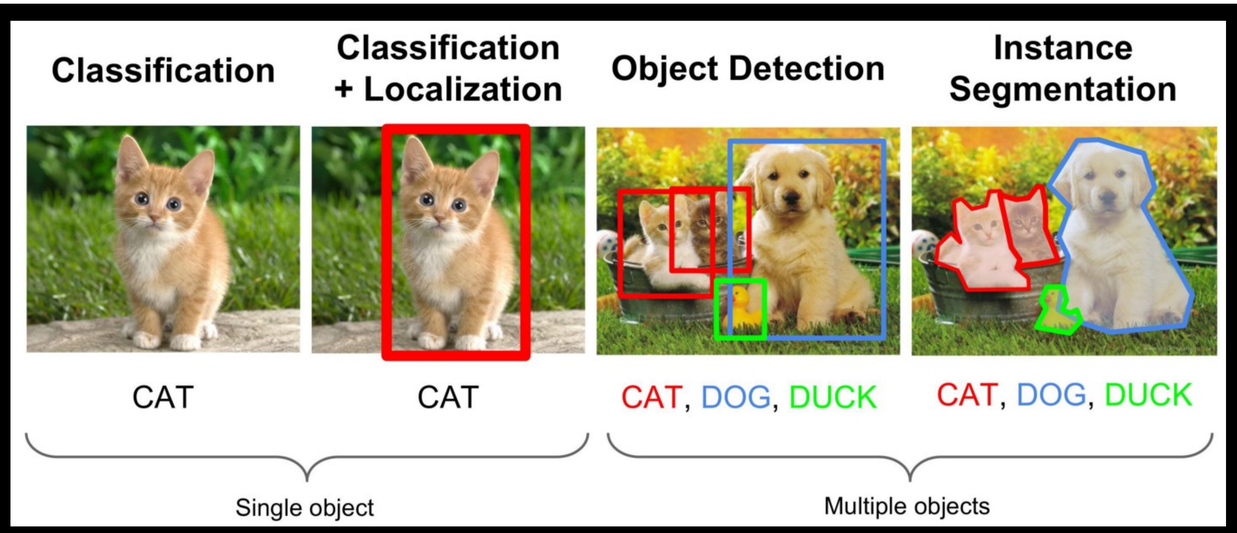

https://this-person-does-not-exist.com/en

https://miro.medium.com/max/2000/1*TwcMmXXuumsDRvgaY2OCQA.png

2. image/art generation

Prompt provided to DALL-E (2022): “An expressive oil painting of a cat as a fisherman”

2. image/art generation

Same prompt, for ChatGPT-4o (2024)

2. image/art generation

An image of fourteen kittens in a boat
playing and looking at fish

2. image/art generation

Depict this person in the National Radio Astronomy Observatory in Green Bank, West Virginia, very strong and holding up a massive radio telescope with his bare hands.

2. text generation

Write an epic, Shakespearean poem about a professor named Dan Reichart conducting a week-long program for education research in radio astronomy, at Green Bank Radio Observatory, WV

The Starry Quest of Dan Reichart: An Epic Shakespearean Poem

Act I: The Summons

Hark! Attend, ye scholars, to this noble tale,
Of Dan Reichart, whose wisdom doth prevail.
In fields where radio waves do softly sing,
He seeks the knowledge that the heavens bring.

Green Bank, that noble observatory grand,
Doth lie in West Virginia's verdant land.
There, Dan, a professor wise and true,
Assembles minds, their thirst for truth to woo.

Act II: The Ascent

From distant lands the eager students came,
Drawn by the beacon of Dan Reichart's name.
A week-long venture, filled with learning's li
To delve in depths of radio's cosmic sight

The Bad: Neural Networks can also be easy to fool…

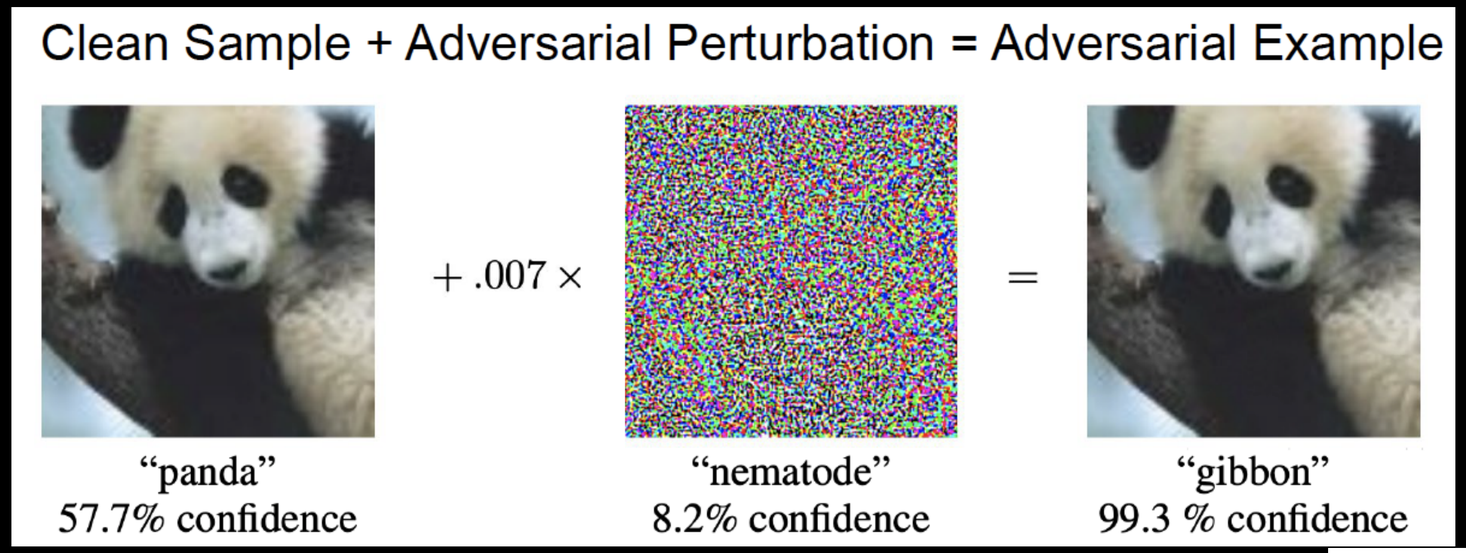

Goodfellow, Ian J., Jonathon Shlens, and Christian Szegedy. "Explaining and harnessing adversarial examples." arXiv preprint arXiv:1412.6572 (2014).

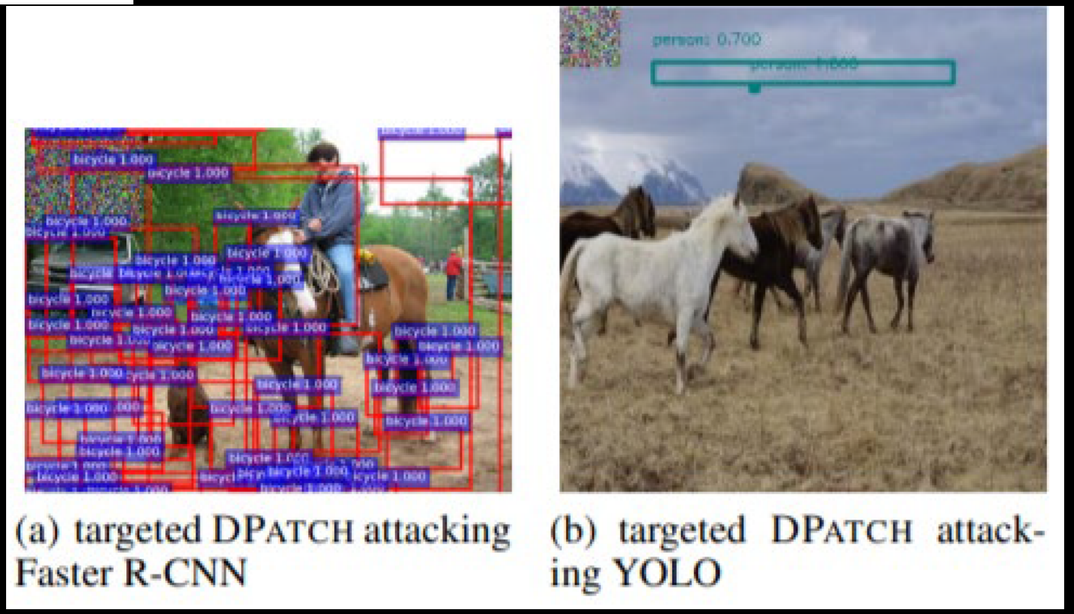

Liu, Xin, et al. "Dpatch: An adversarial patch attack on object detectors." arXiv preprint arXiv:1806.02299 (2018).

The Spooky

AI can be *(and has already been)* used for unethical applications

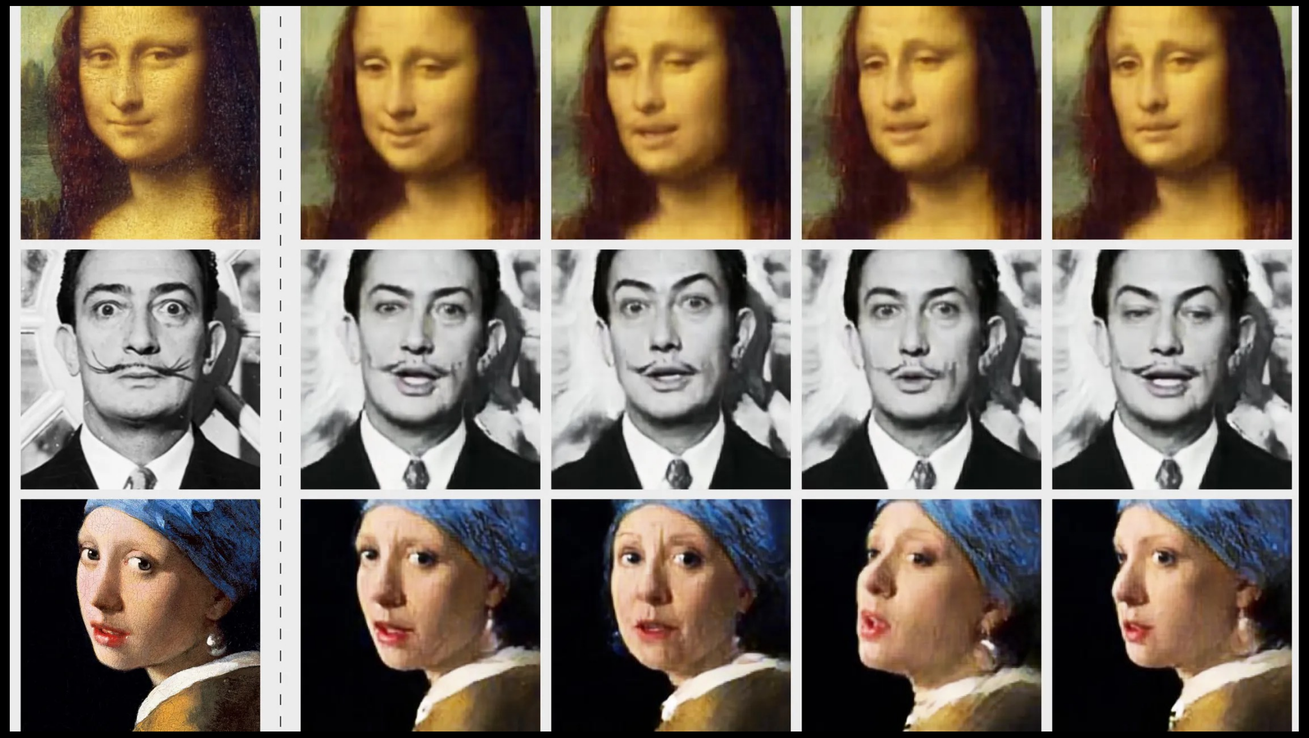

https://cdn.thegeekherald.com/wp-content/uploads/2019/06/New-AI-deepfake-app-creates-nude-images-of-women-in-seconds-1-e1561666459843.jpg

Thanks for Listening! Questions?
Live demo if we have time!