What actually *is* Artificial Intelligence?

And how does it relate to astronomy?

https://i.imgur.com/1AraF4i.jpg

Nick Konz, ERIRA 2022

About me

(Here) circa ERIRA 2019

UNC Chapel Hill
Double major in Astrophysics & Math, class of 2020

Research **with Dan**
in computational statistical techniques for astronomy
started ERIRA in 2018 as participant

Duke University (now)
Ph.D. student in Machine Learning (going into 4th semester)

Research in medical image analysis with deep learning

Duke | DEPARTMENT OF Electrical & Computer Engineering

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

AI in popular fiction

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

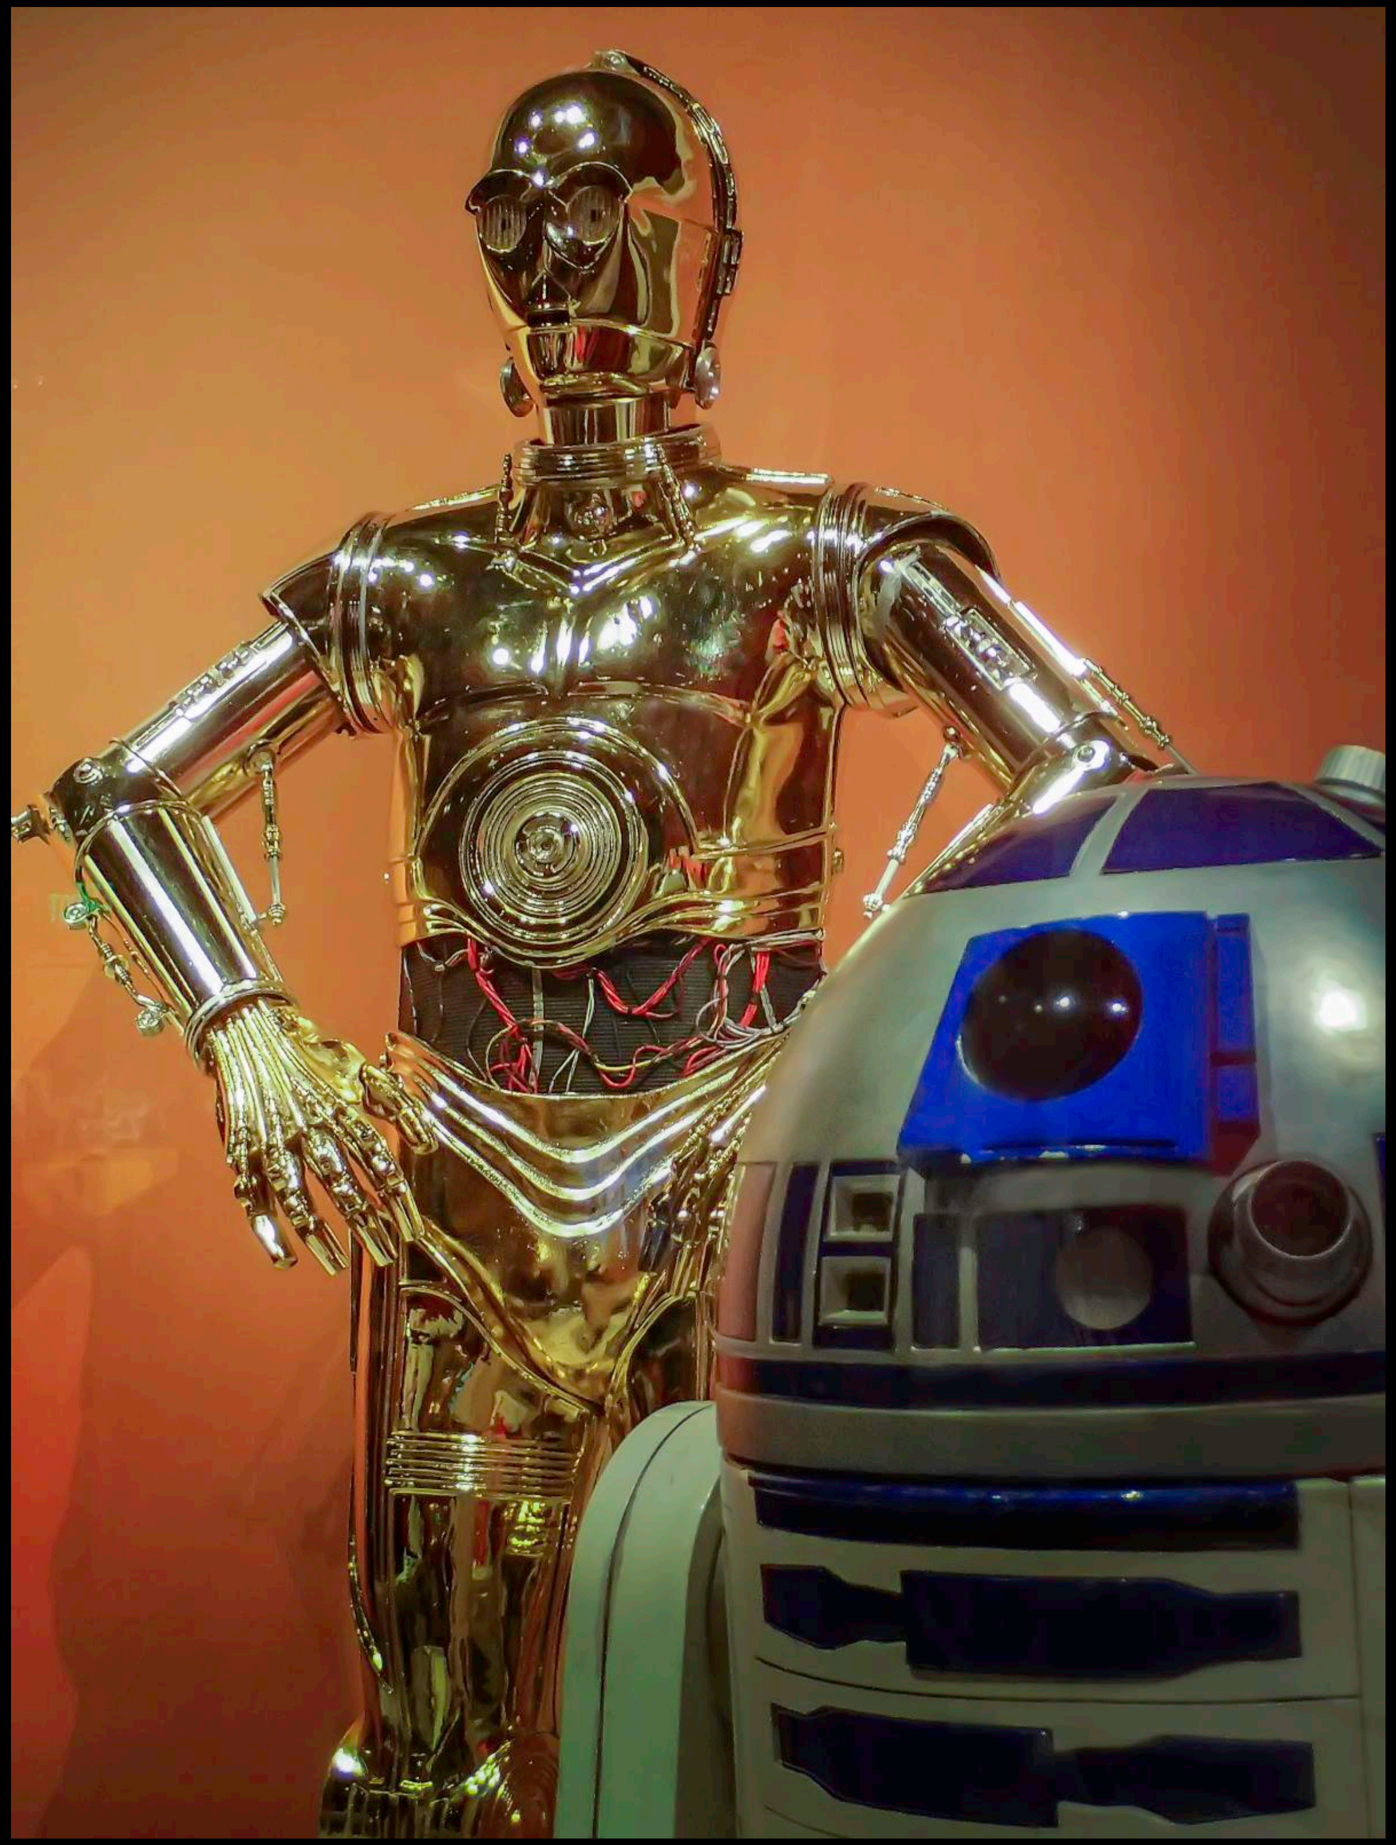

https://www.flickr.com/photos/mharrsch/16446792154

- Usually ***artificial general intelligence****: understanding or learning any intellectual task that a human being can.*
- Way beyond the realm of current tech. Unanswered questions:
    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?
- The goalposts keep shifting (for decades); only possible in fiction? Or maybe our definition of “ultimate” intelligence is too human-centric?
- Current AI advancements are within single-task systems, mainly:
    1. visual understanding
    2. language processing

So what can modern AI *actually* do?

Classical AI vs. Modern AI

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

1. For a long time, AI systems were completely (or almost completely) human-designed:
    1. “Expert”-designed decision rules/algorithms
    2. “Expert”-designed knowledge representations
2. However, these were very brittle and only worked in “toy” scenarios:
    1. The issue: how long would it take to manually model for all possible phenomena/**features** seen in the real world? This could not ***generalize*** well to new situations

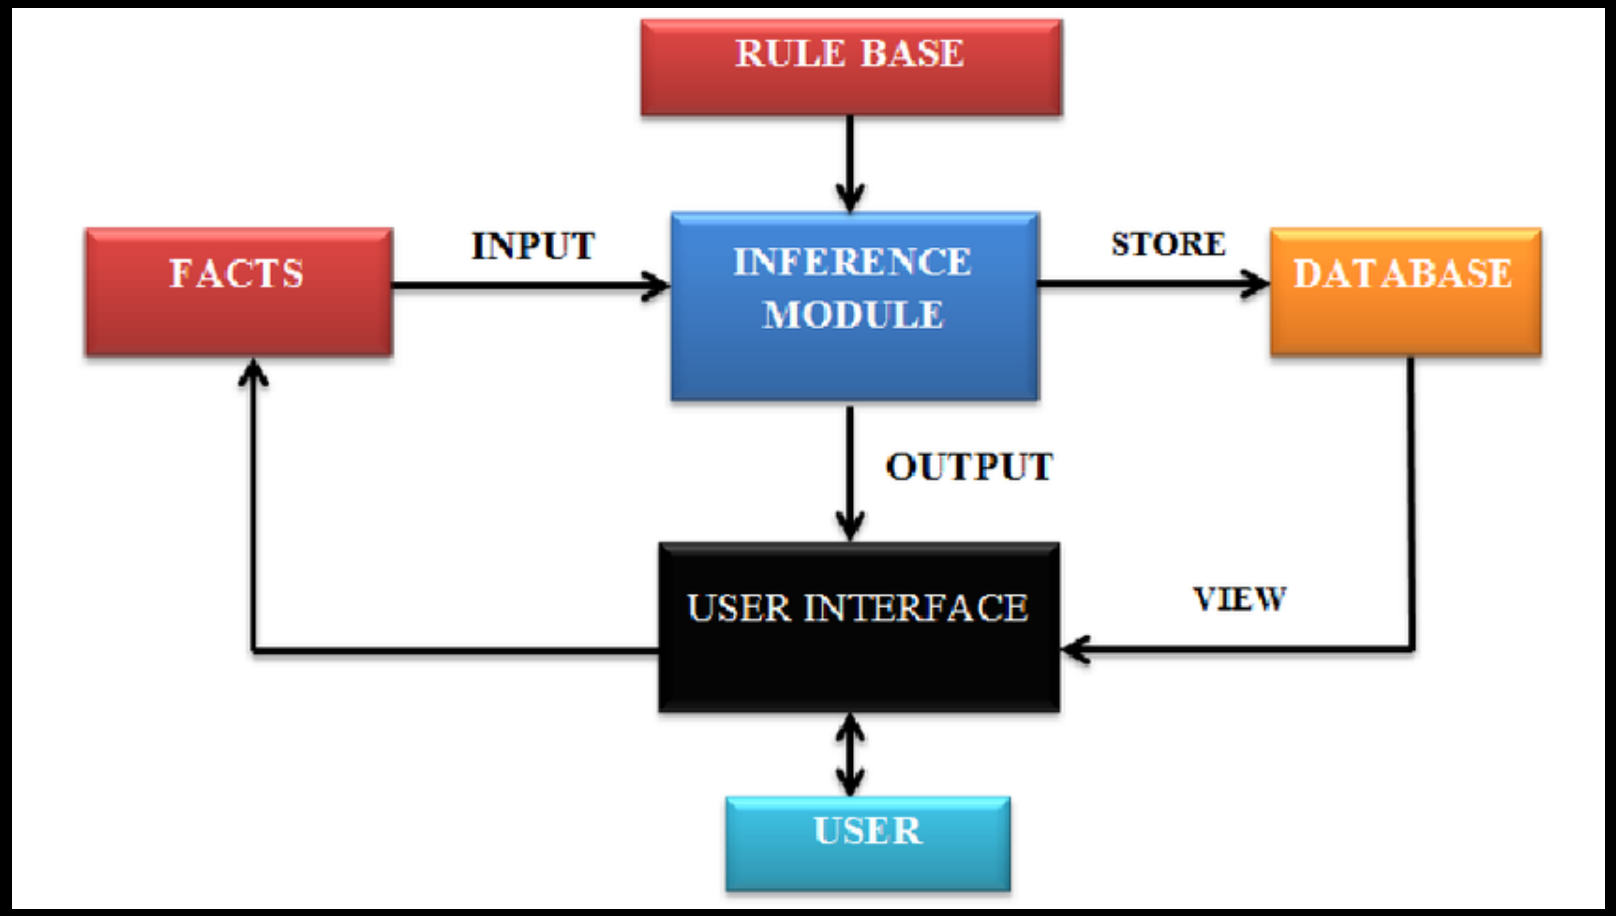

https://studiousguy.com/wp-content/uploads/2021/07/Rule-based-Production-Systems.jpeg

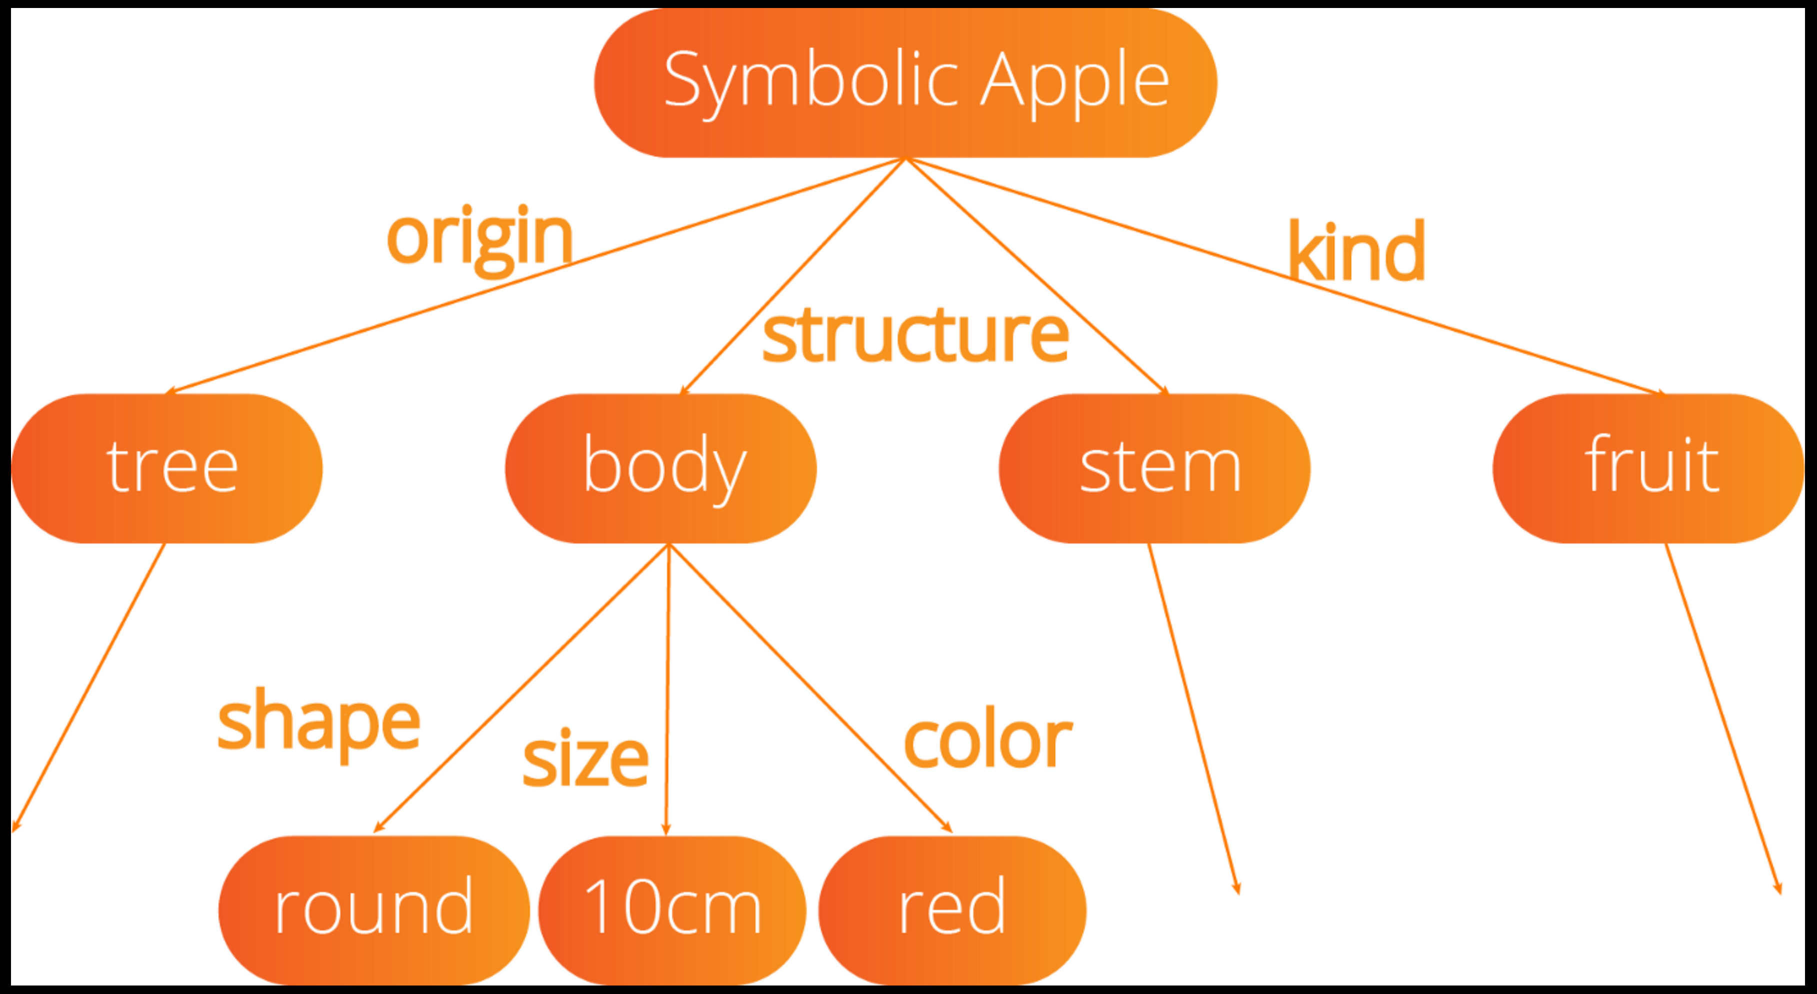

https://miro.medium.com/max/1838/1*xYZrAXZi6lq3c1Z_CxW6ZA.png

Modern AI: automatic learning from data

1. **Deep learning** revolutionized AI, by making knowledge representations and decision algorithms **learned from data**, rather than **hand-designed**
2. AI models that didn't require human-designed decision rules and knowledge representations could easily be **scaled up** to a complexity far greater than what we could design: thousands, millions or **billions** of parameters that are **learned automatically from data**
3. Deep learning models are called **artificial neural networks**
4. All that we need is enough data for the AI to **learn from.**
5. Two main applications of deep learning: **Computer Vision** and **Natural Language Processing**. For this presentation I'll focus on computer vision, which typically use **convolutional neural networks**

Classical AI vs. Machine Learning vs. Deep Learning

- Key differences by color:
    1. **Manually-designed**
    2. **Automatically learned from data**
- Thus, deep learning allowed scaling to far more complex models
- **features** (for images) are shapes, colors, objects, etc…

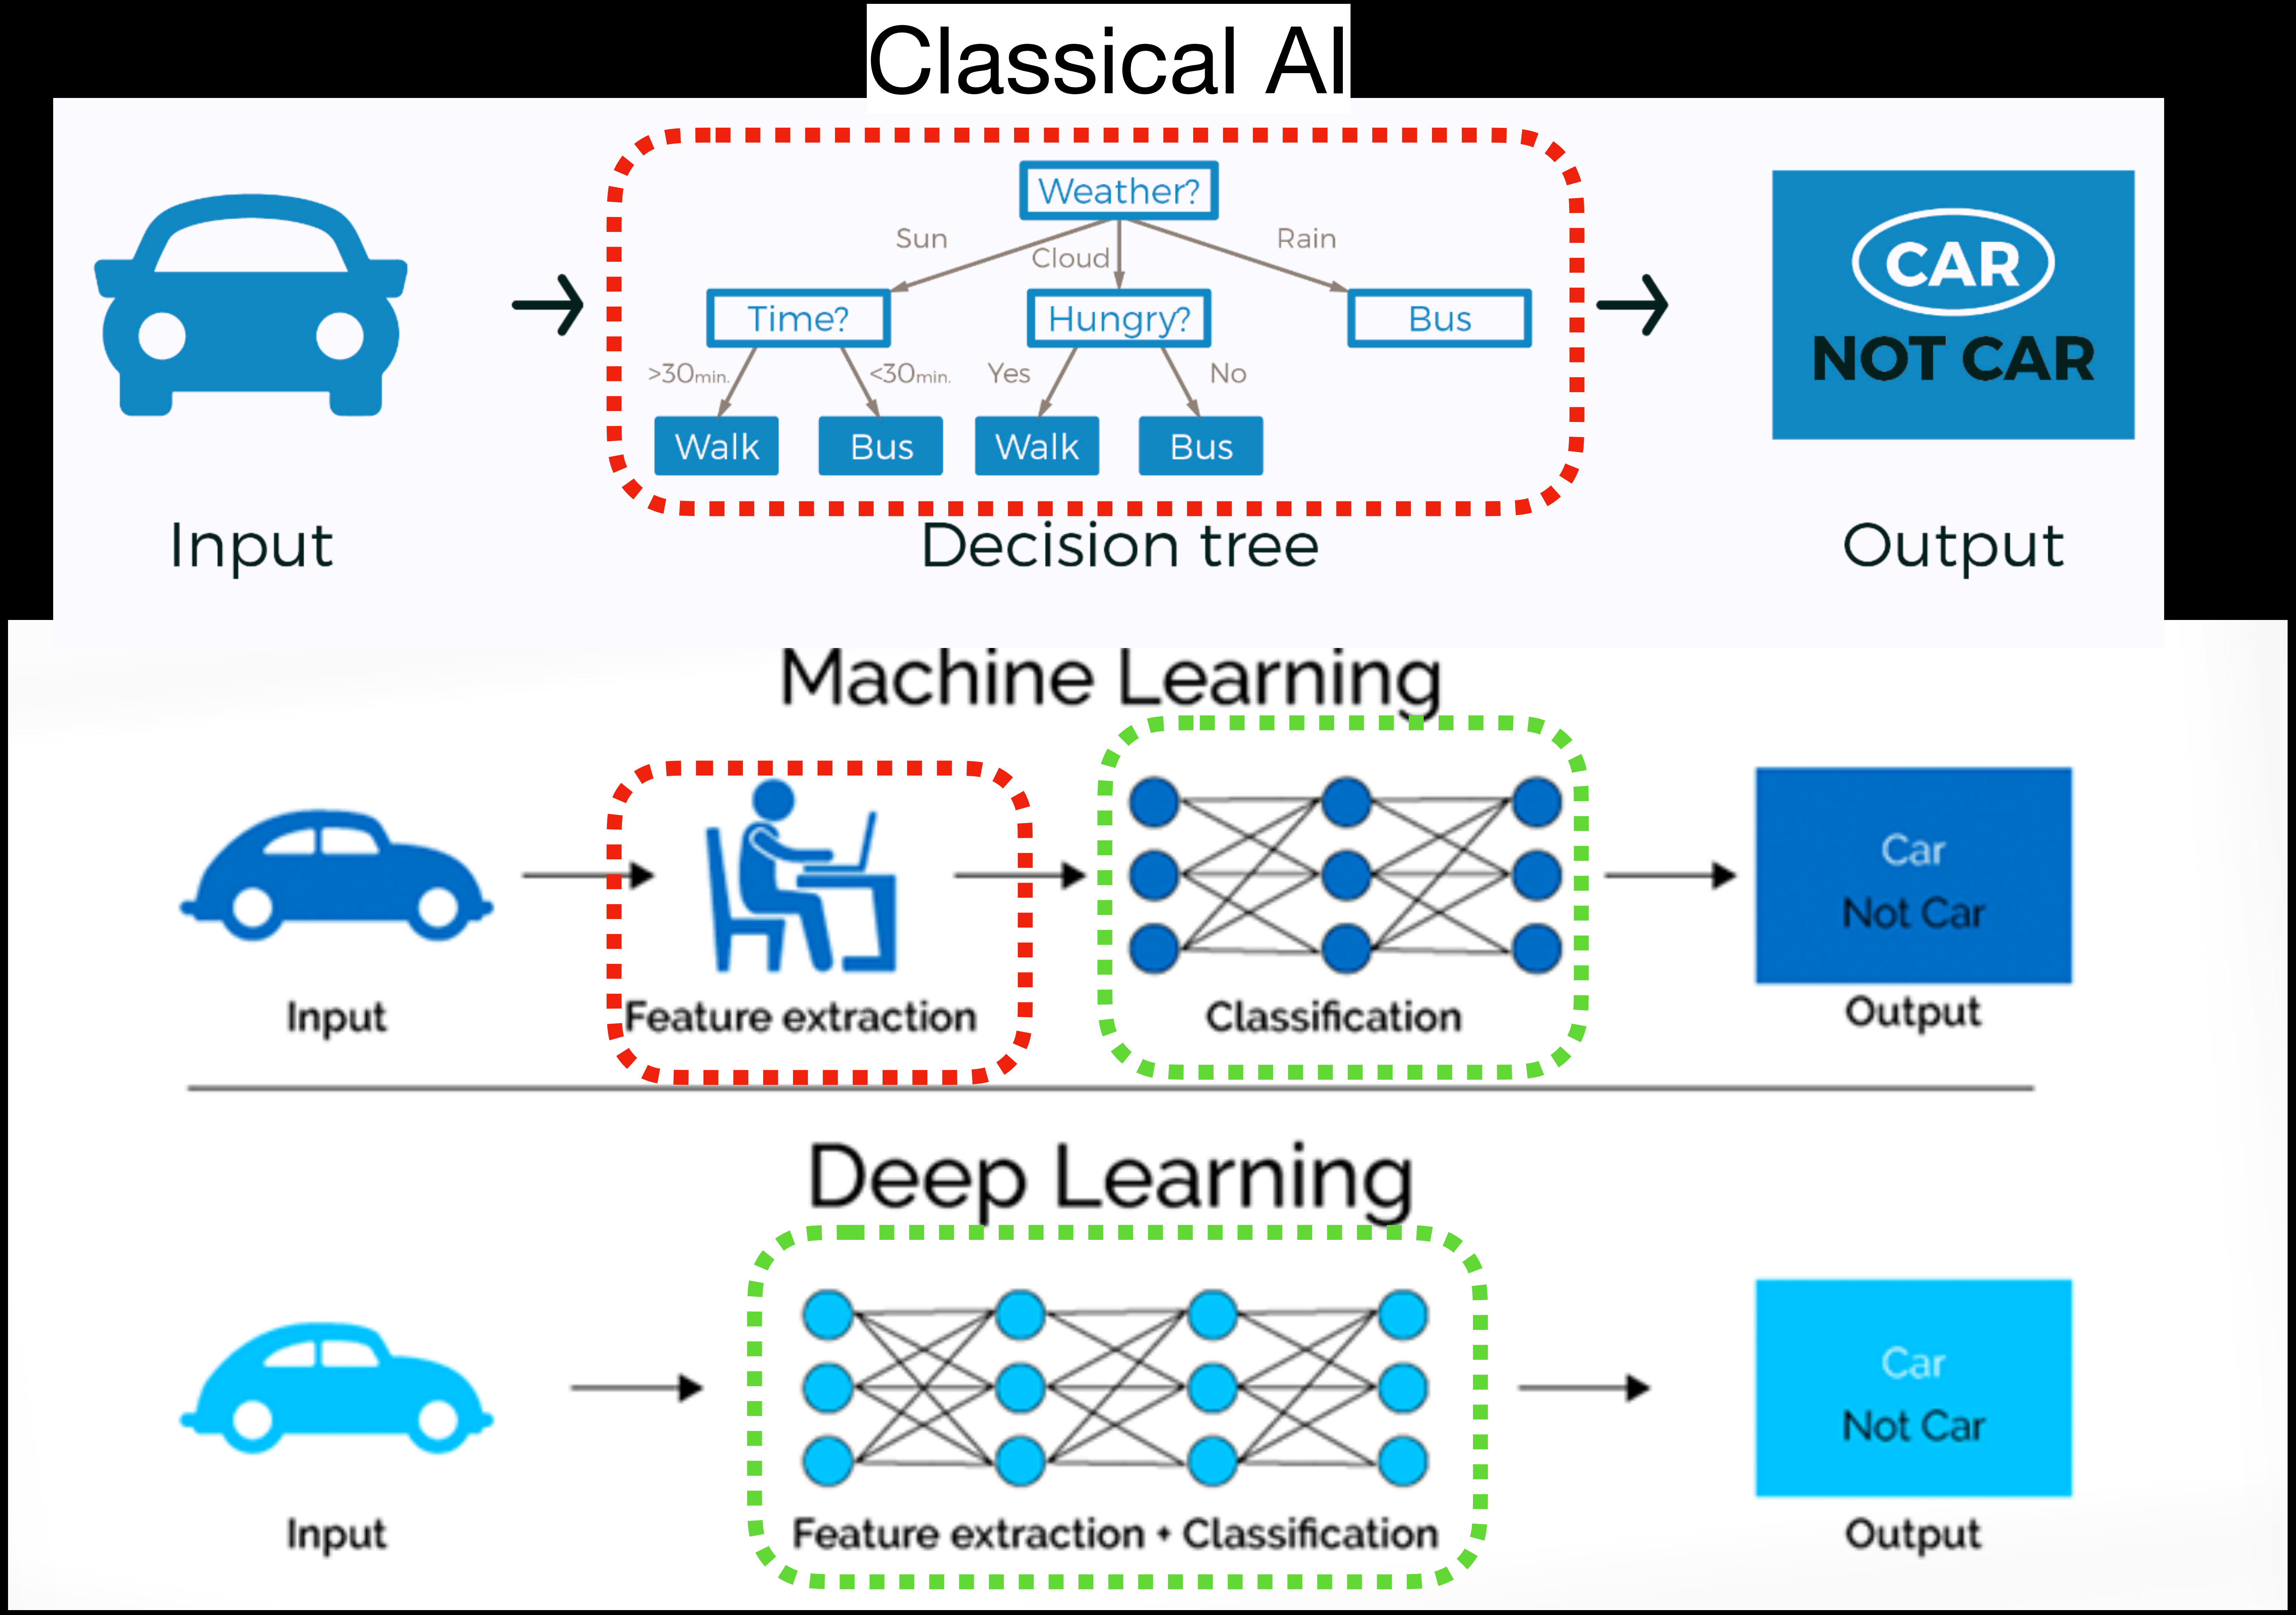

https://miro.medium.com/max/1838/0*KJOhHjSJXo0RMFuC.png

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

Why “*deep*” learning?

- Shallow: extract hand-designed features from data and use these to generate the output prediction

- Deep: many sequential layers that allow for more and more abstract visual features to be modeled and learned from the data directly:

  - neural networks take raw images as input and “figure out” which features to extract

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

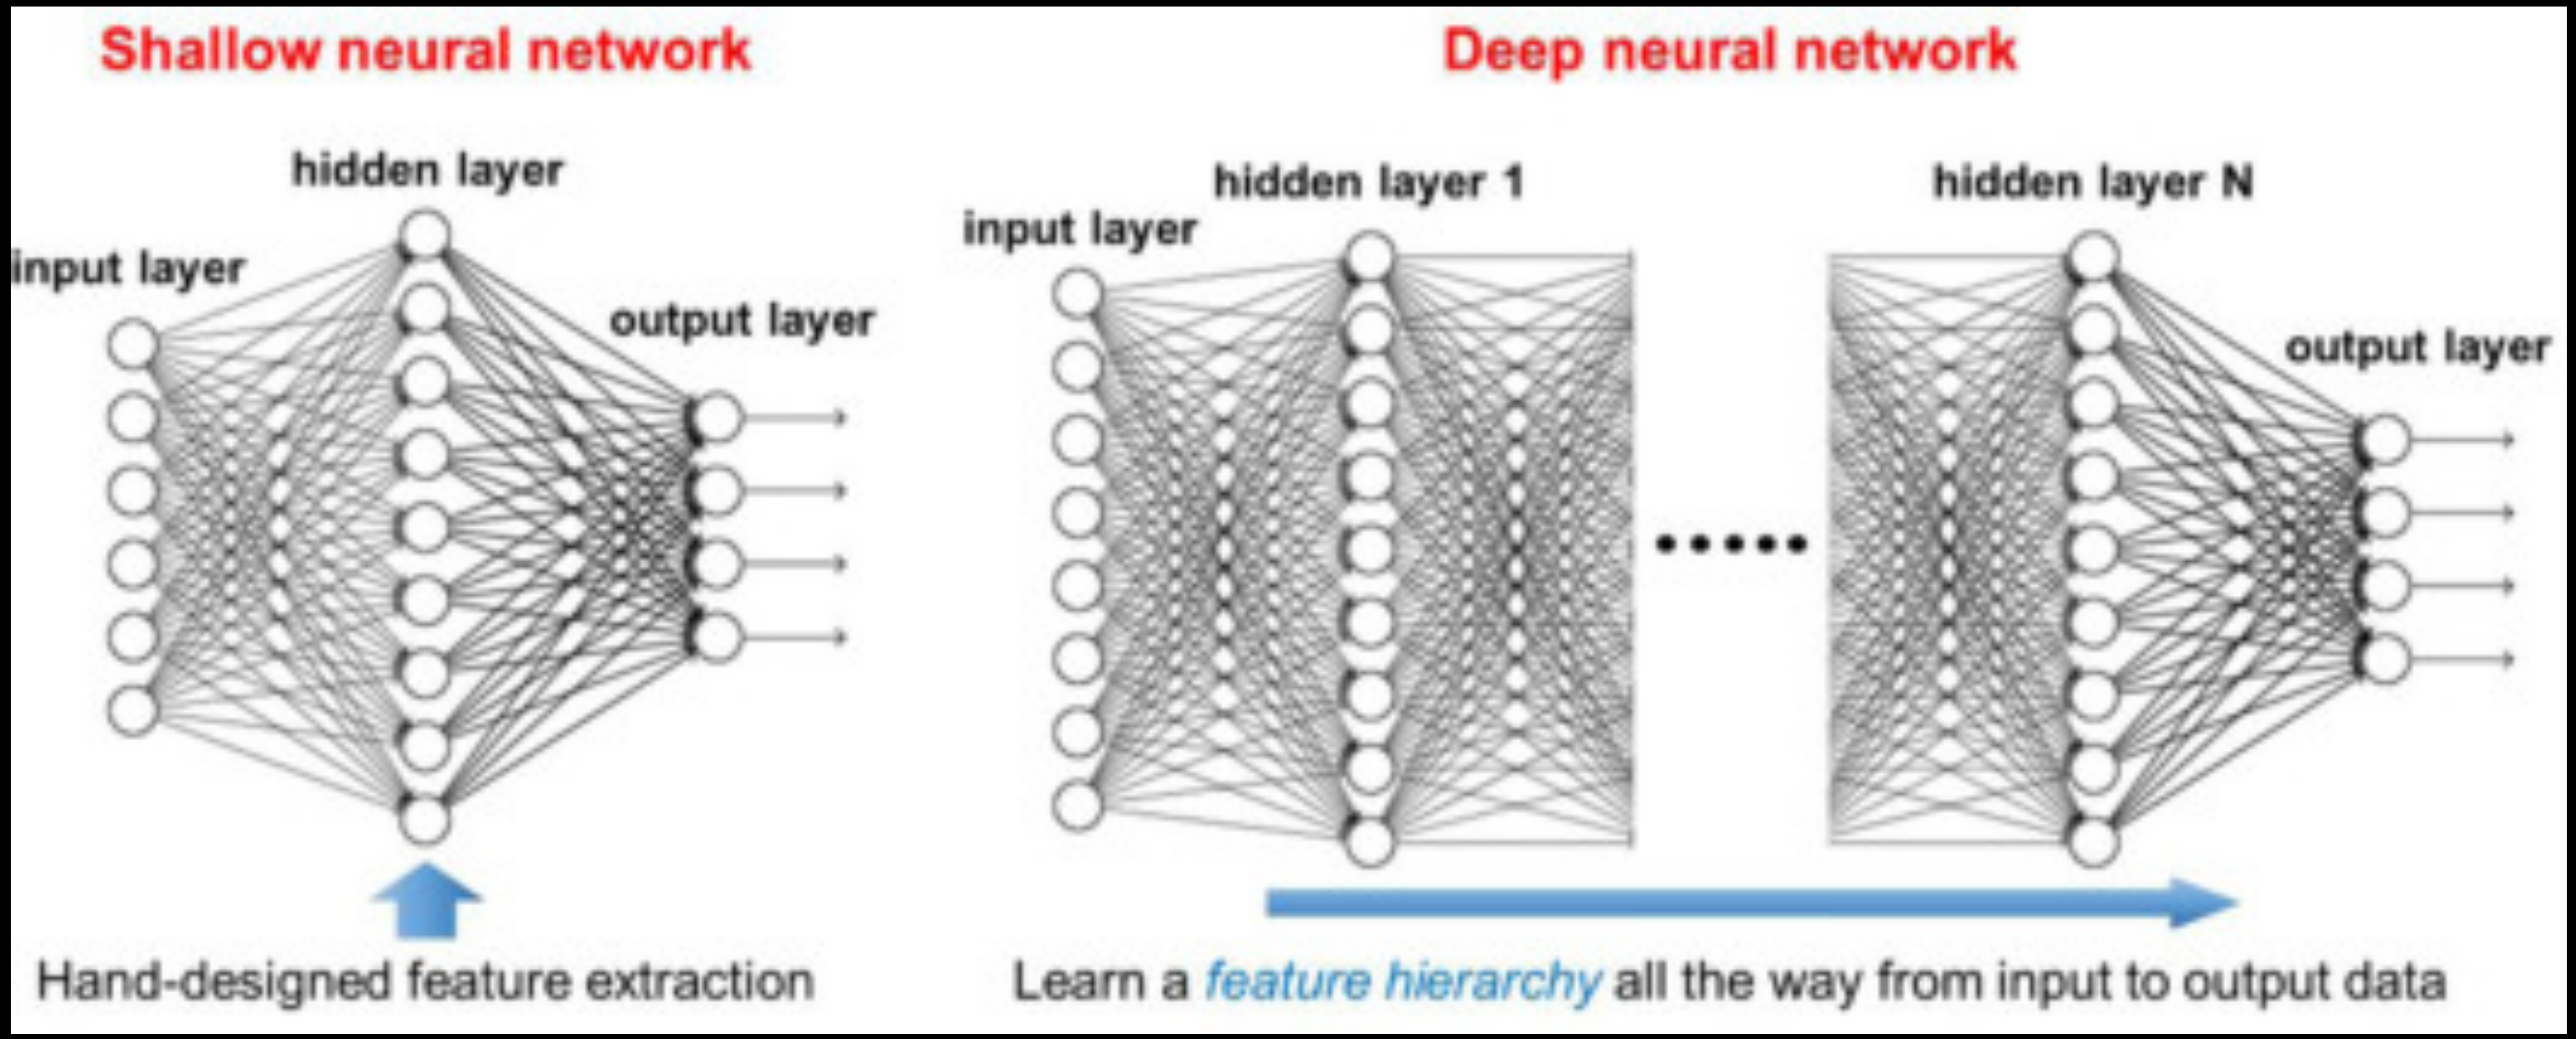

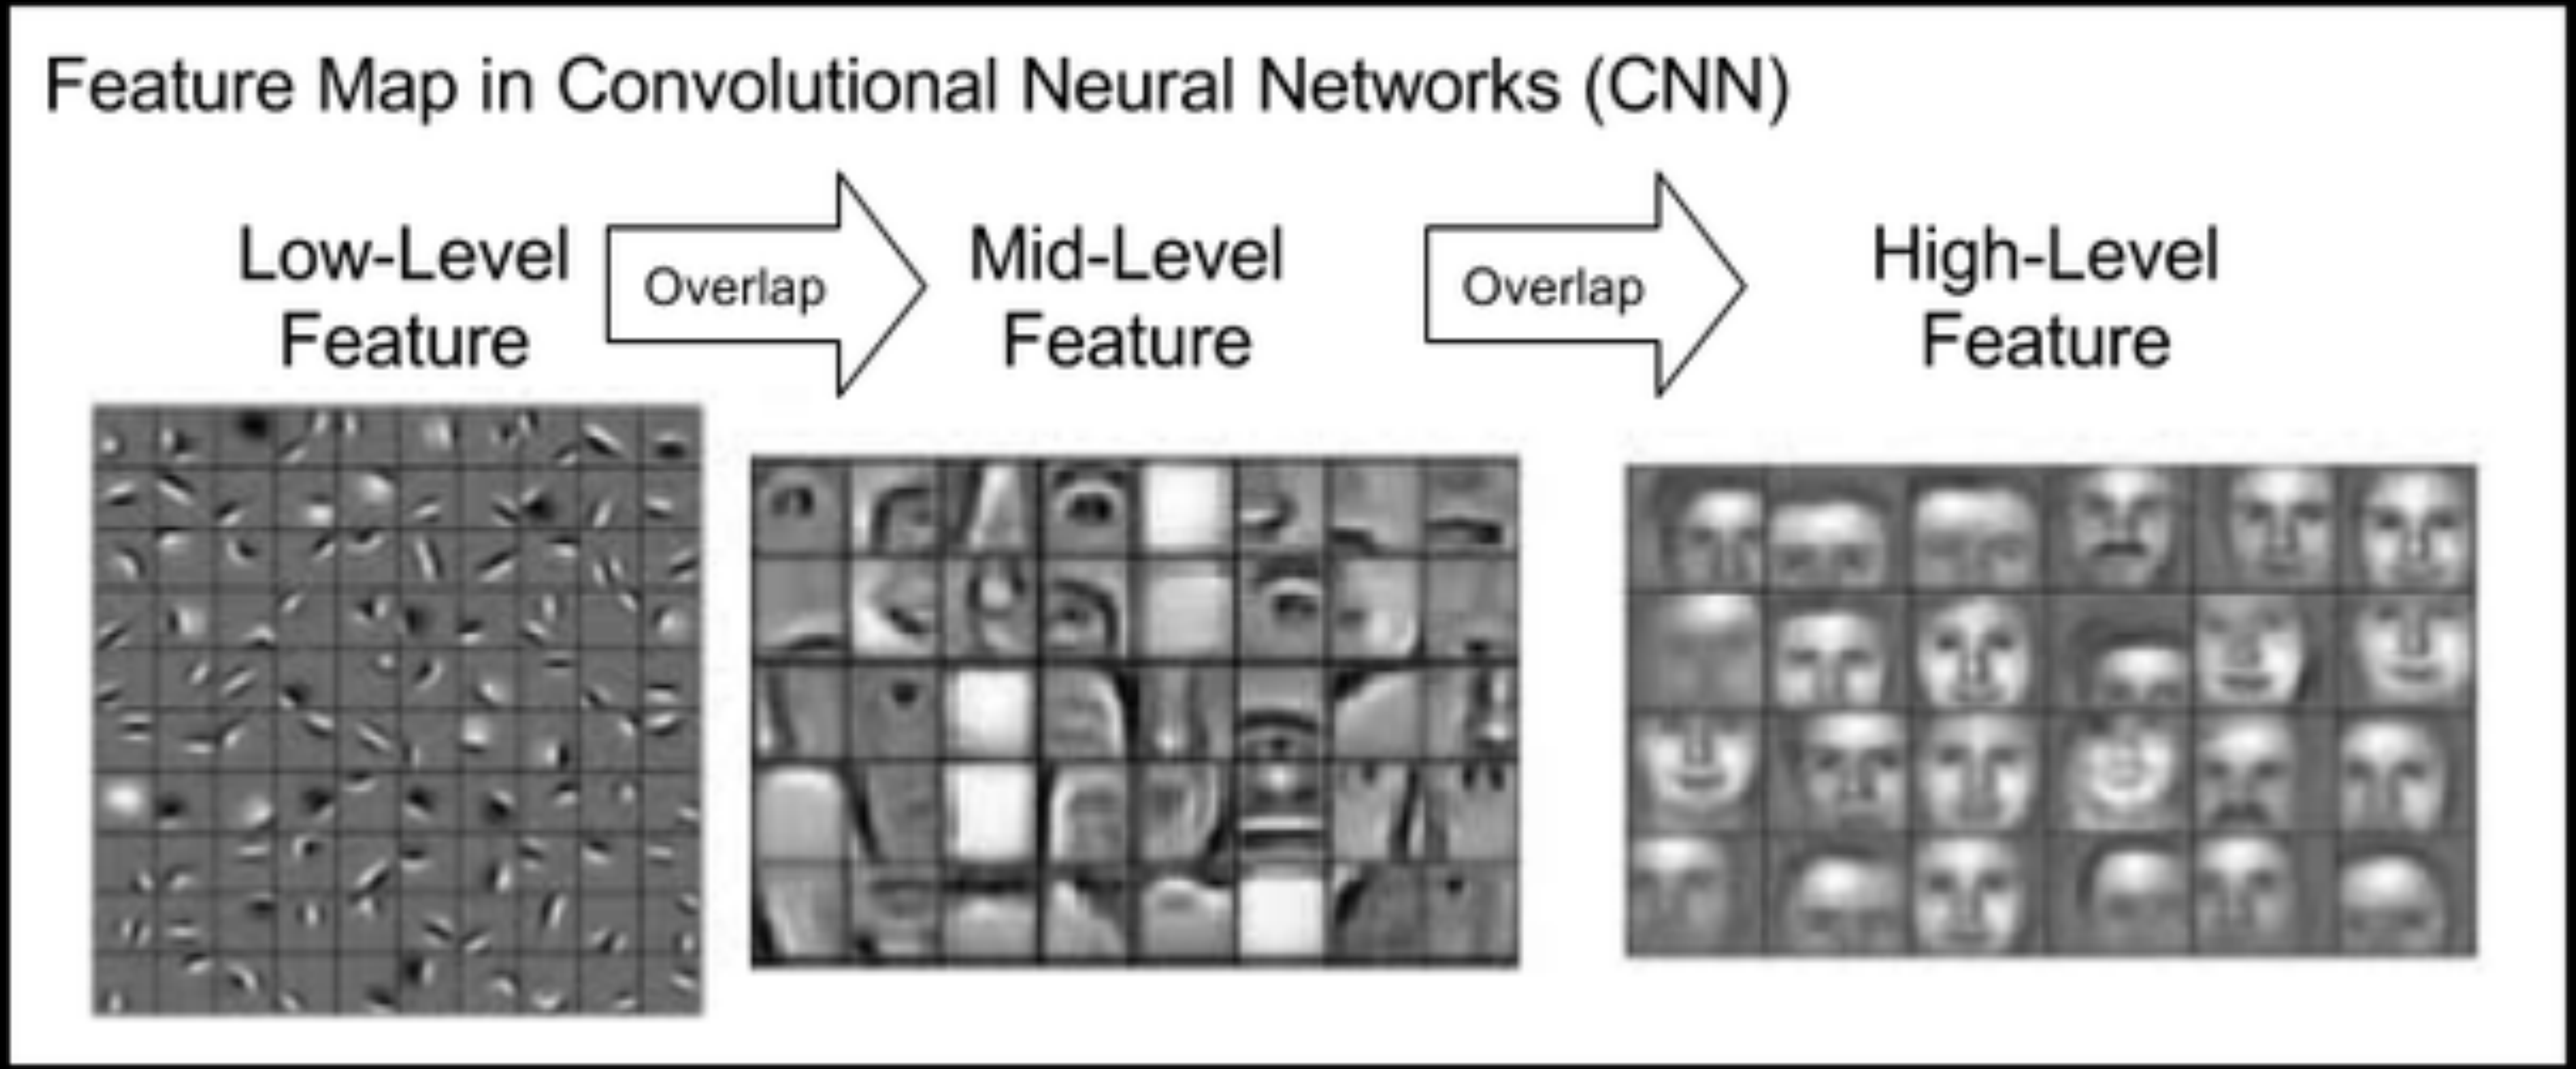

https://miro.medium.com/max/1838/0*KJOhHjSJXo0RMFuC.png

How does visual feature recognition actually work?

Visual Recognition with Classical AI

- **Classical AI:** Manually-designed filter ran over image to *detect edges:*

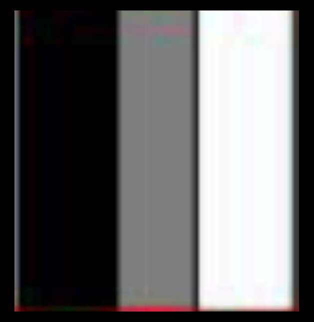

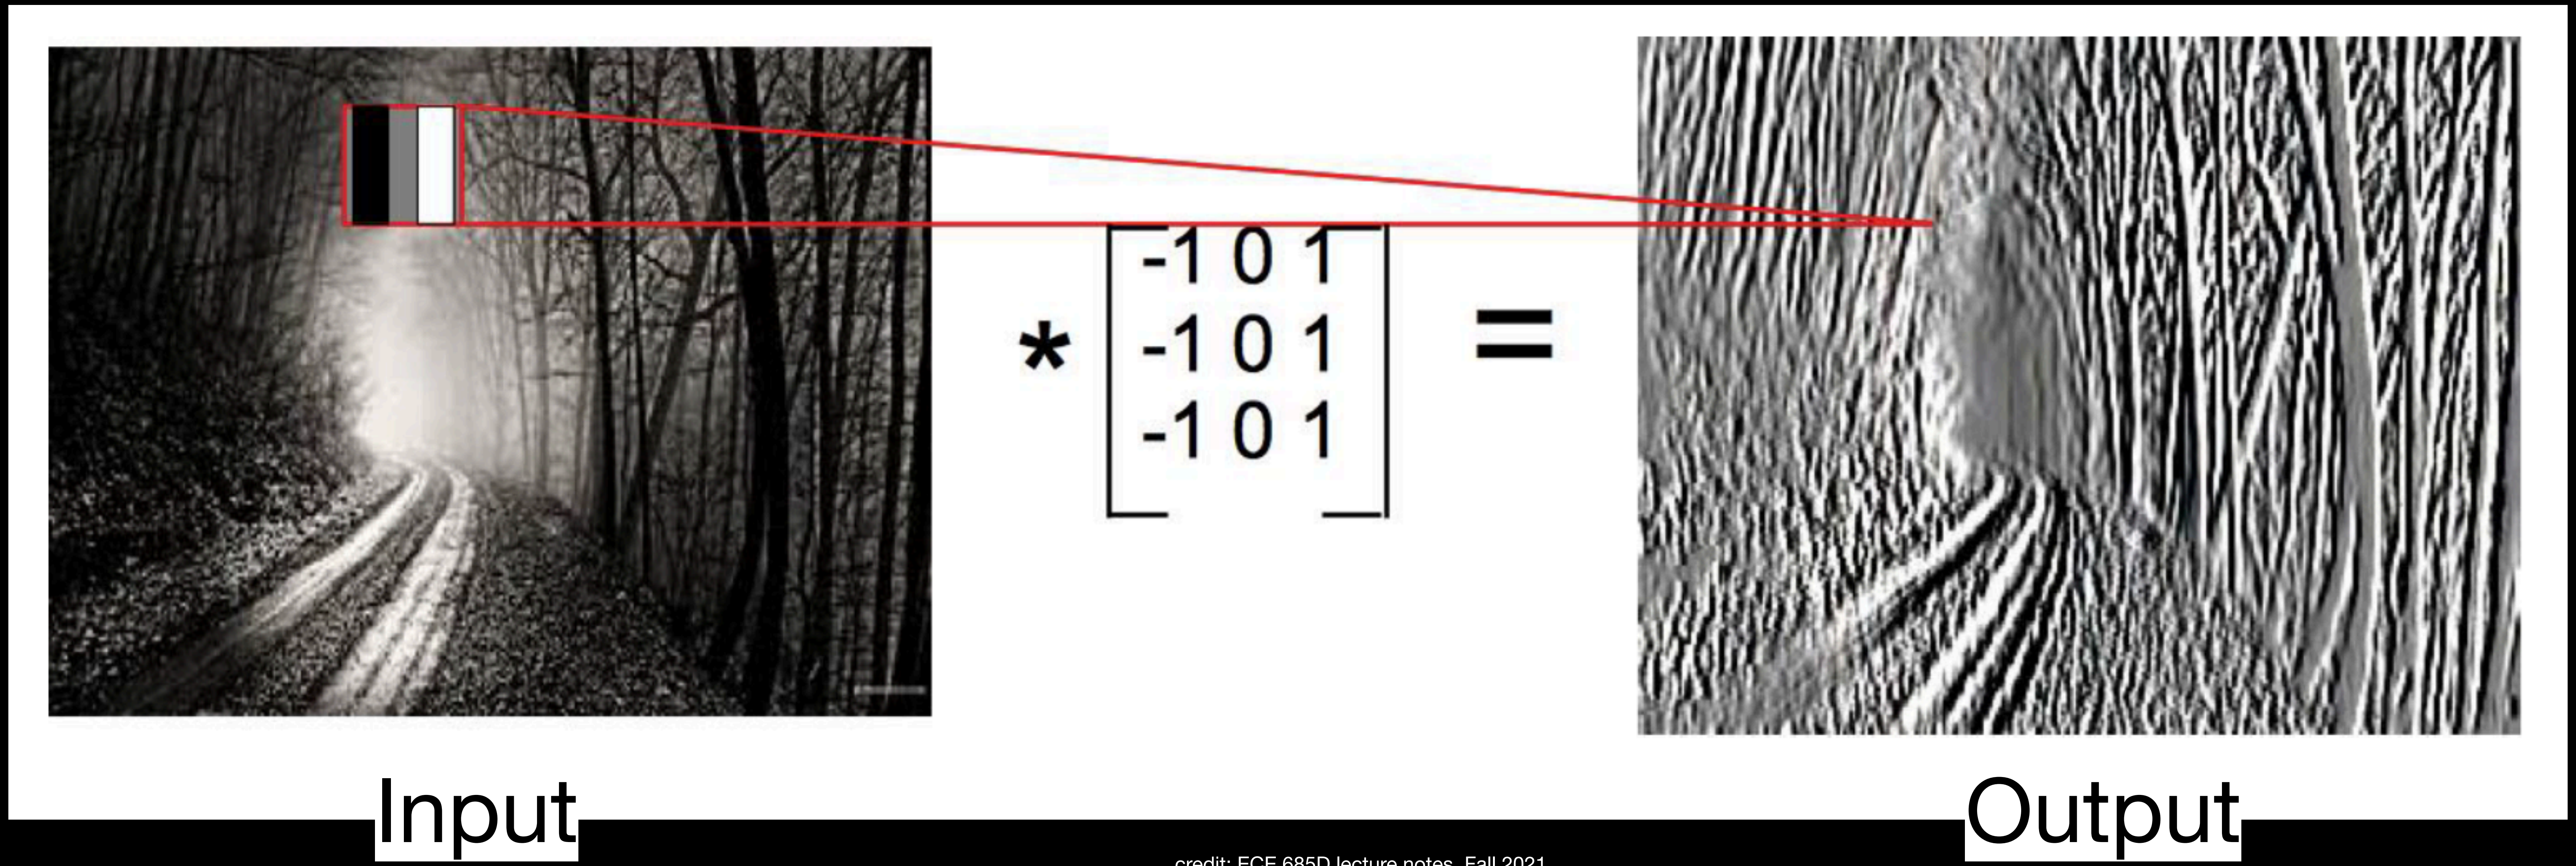

credit: ECE 685D lecture notes, Fall 2021

- These edge **features** could then be used for further hand-designed predictions

Visual Recognition with Deep Learning

- **Deep Learning:** many hierarchical filters *learned from data*, to process data into increasingly abstract features that are useful for making predictions

Input (image of my cat)

https://miro.medium.com/max/1026/1*Ji5QhY9QXBIpNNLH4qAcNA.png

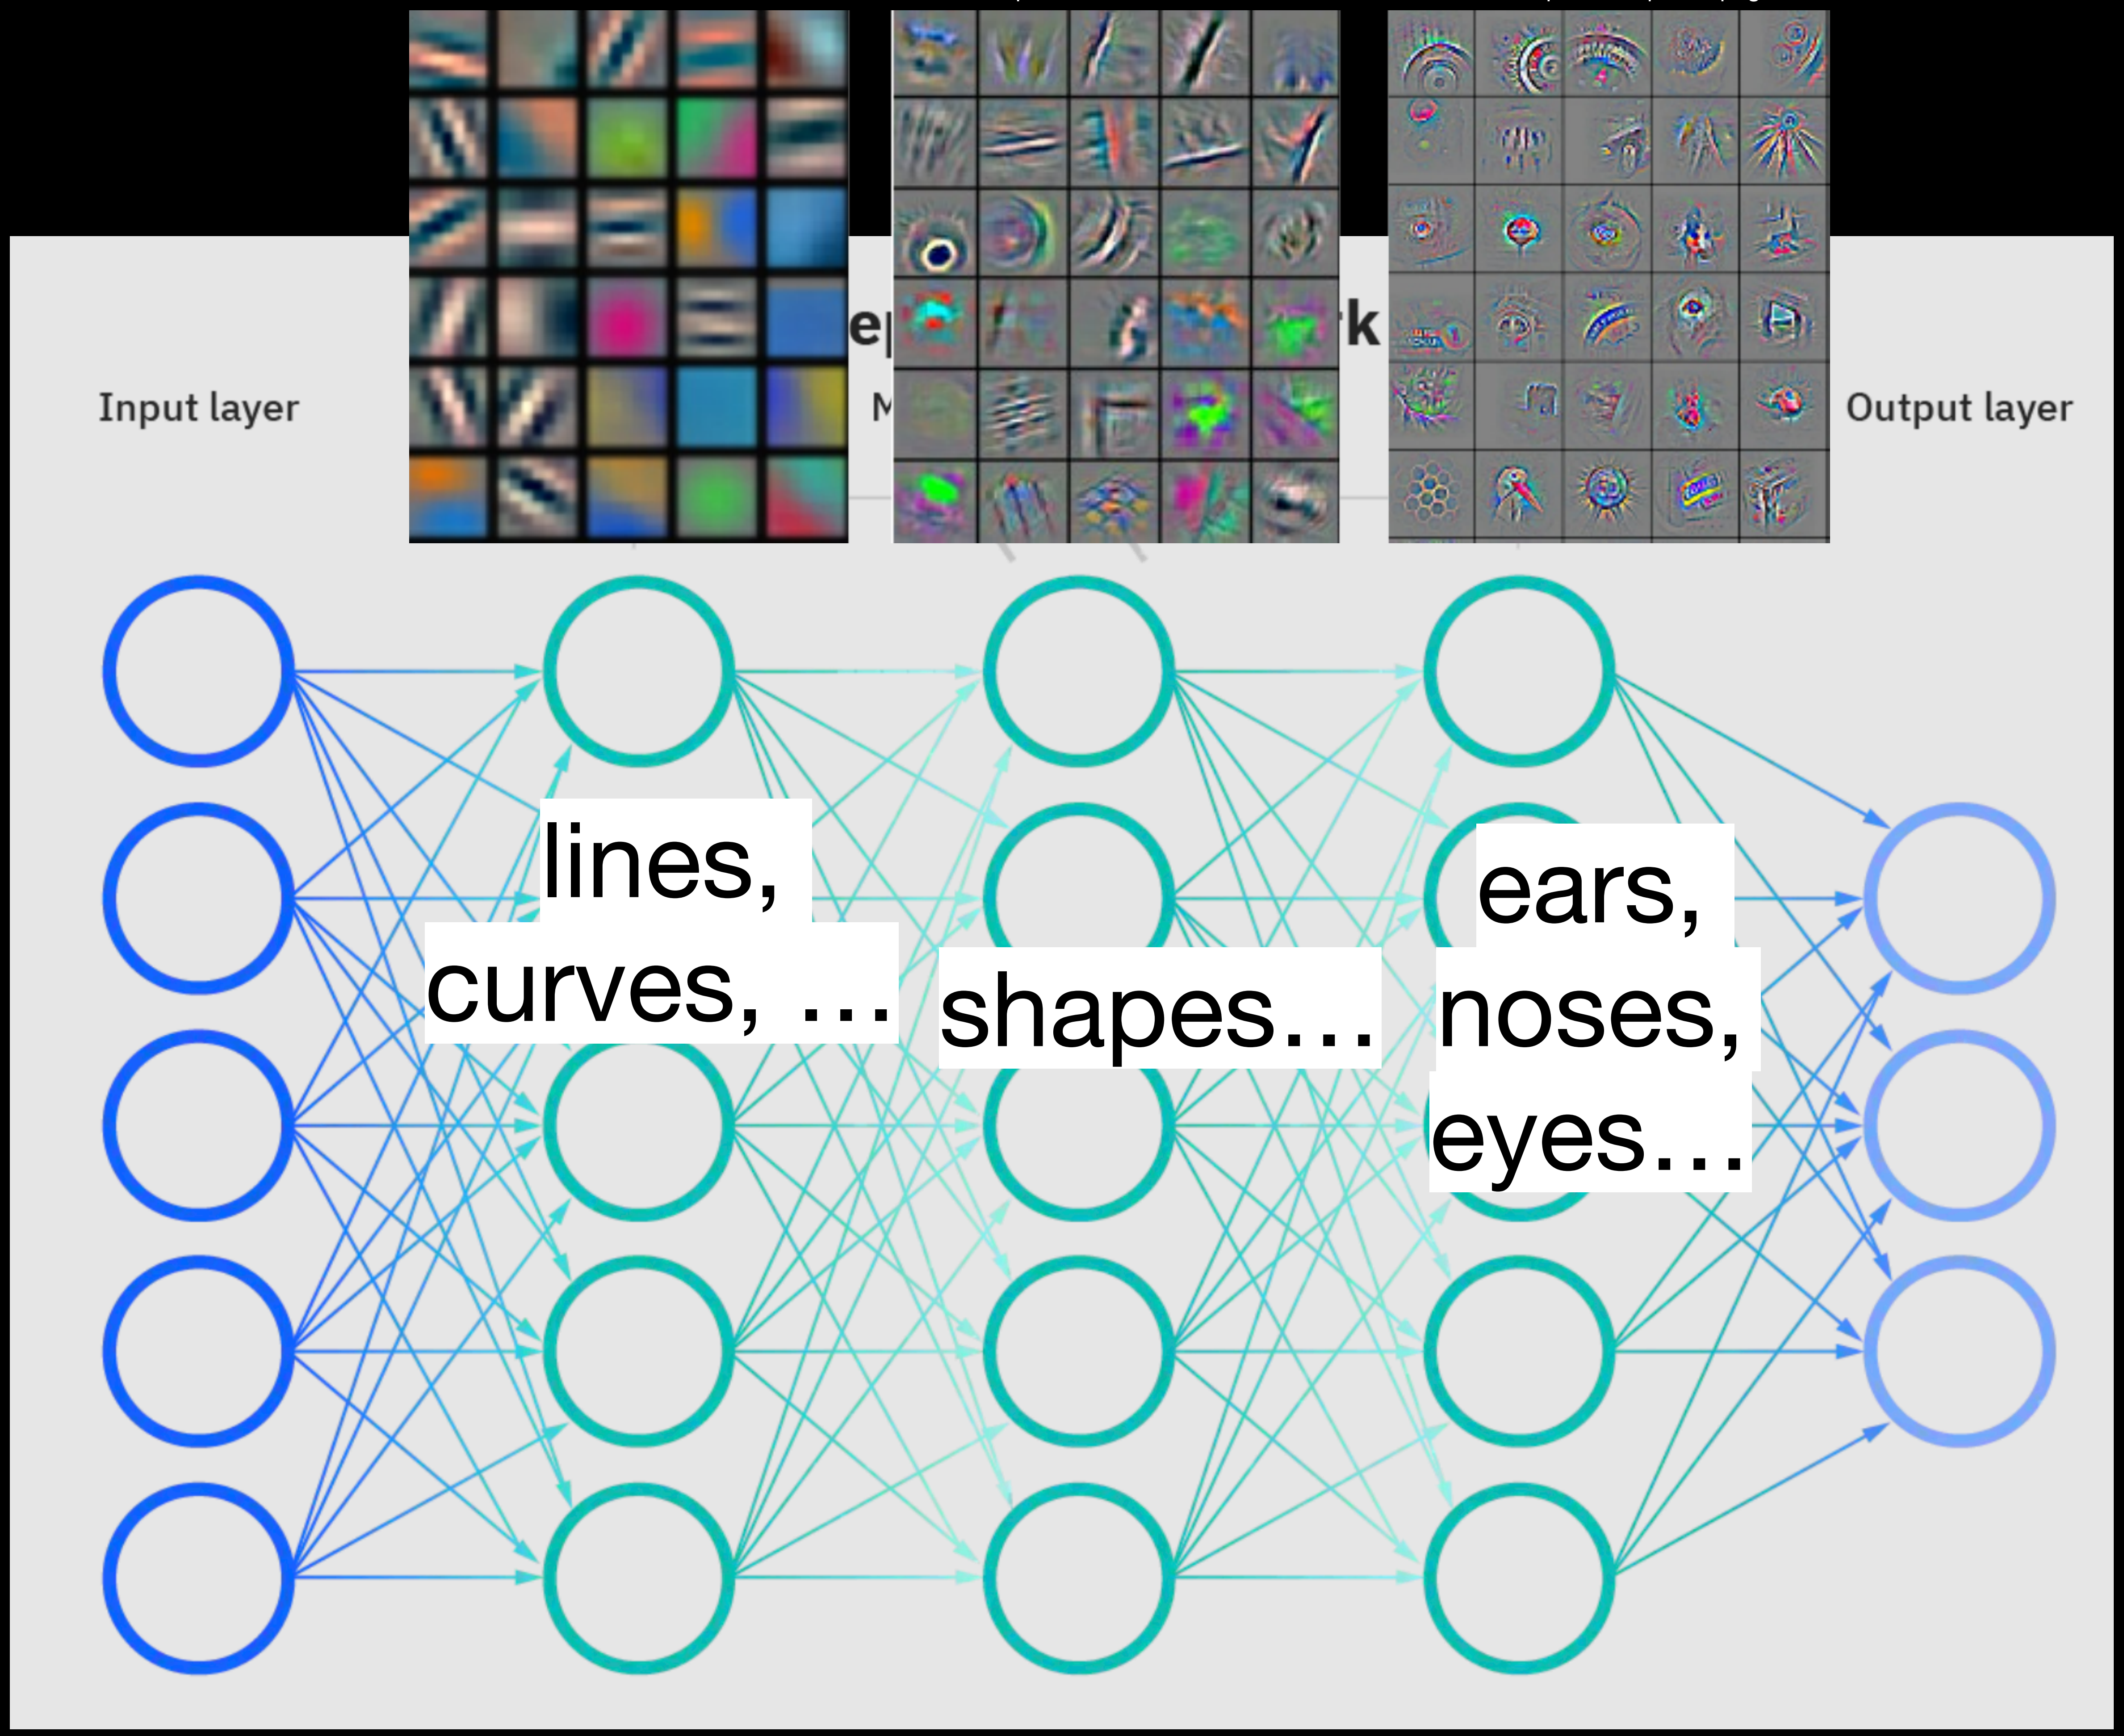

https://velog.velcdn.com/images/tjswodud/post/a5149a8a-18e4-4fd2-9c39-e7a14a9a4a57/image.png

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?

- These extremely complex *neural networks* would be useless if we couldn't *train them*
- Learning is just the adjusting of (thousands, millions or billions) of model parameters to minimize the error of the model's predictions on data
    - This error is minimized using a **gradient descent algorithm**, like finding the lowest point of the "hill" of the error with respect to the parameters
- This is made computationally feasible by various technical advancements such as parallel computing, automatic differentiation, etc.

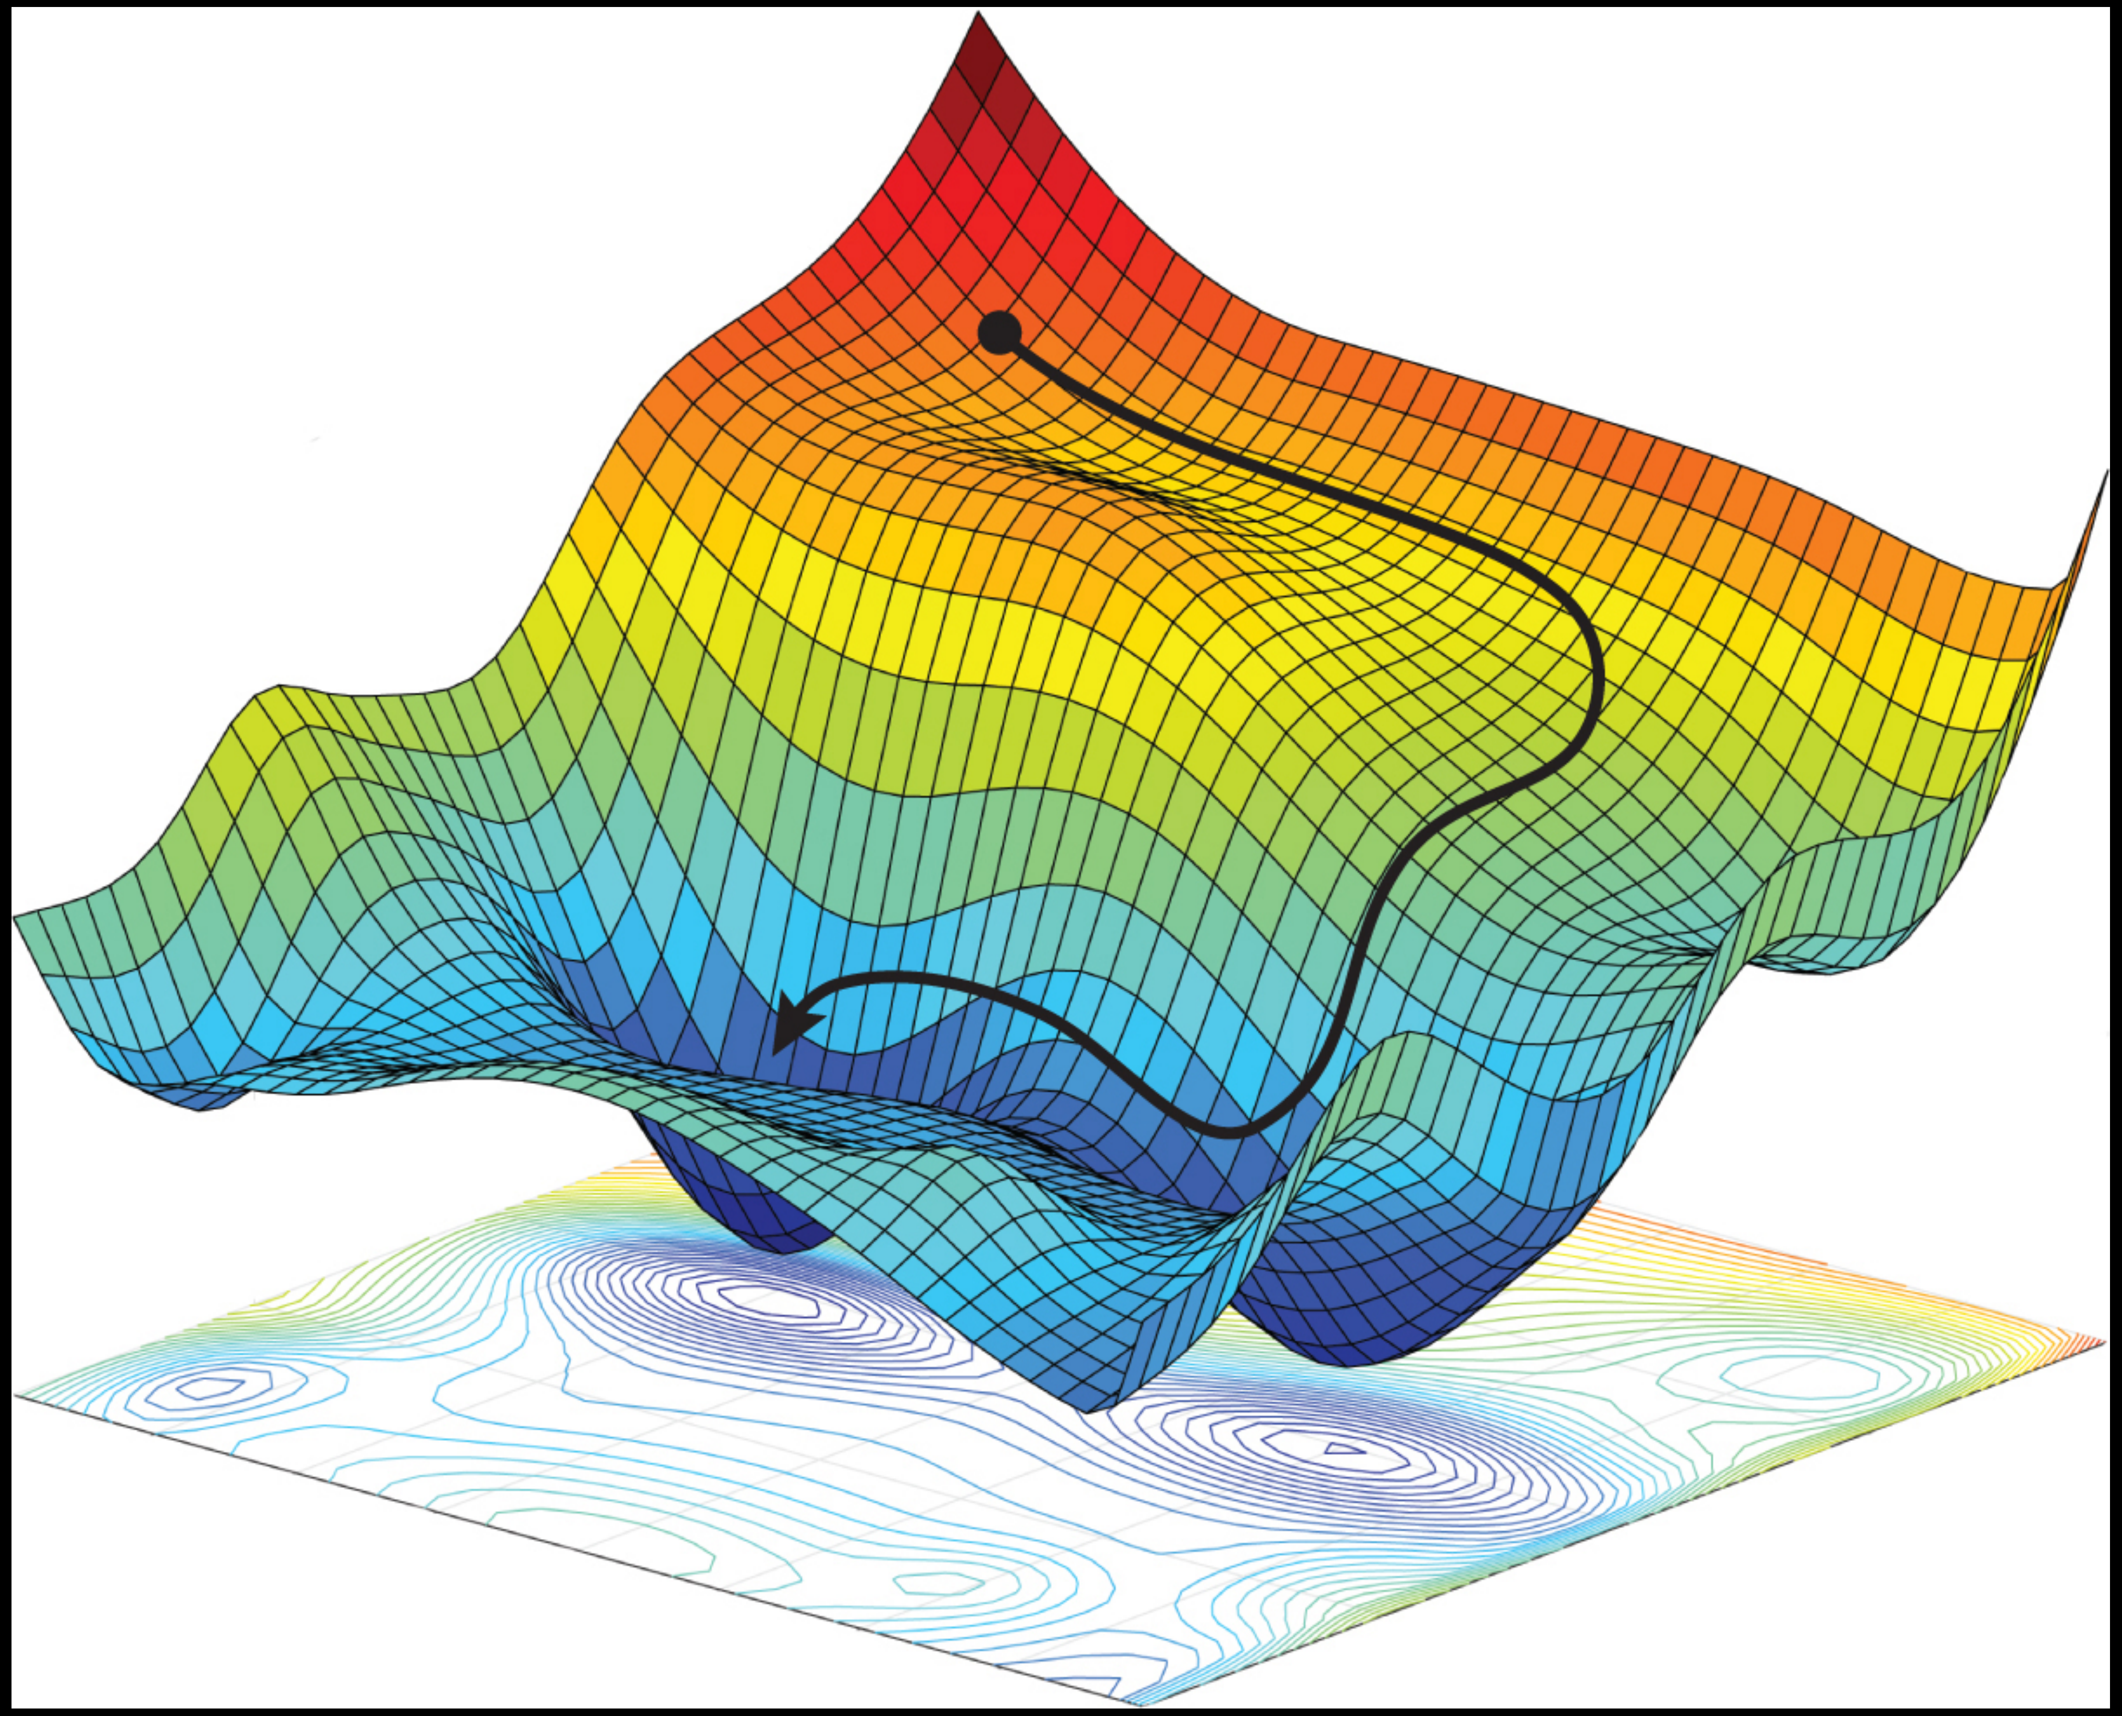

https://www.researchgate.net/profile/Alexander-Amini/publication/325142728/figure/fig1/AS:766109435326465@1559666131320/Non-convex-optimization-We-utilize-stochastic-gradient-descent-to-find-a-local-optimum.jpg

From Linear Regression to Neural Networks

- Neural networks are just sequential high dimensional nonlinear functions!
- In brief: start with your regular old linear regression model $\overrightarrow{y} = M\overrightarrow{x} + \overrightarrow{b}$
- Output $\overrightarrow{y}$ into a nonlinear function like $\tanh$, send that through another linear function, and repeat…
- Enough layers of this (modern neural networks have 10s or 100s), and neural networks can *approximate any function* (Universal Approximation Theorem)
- For linear regression, you fit the model (slope and intercept) to the data; same for all of the weights/connections of a neural network! **This is the "learning"**

https://swayanshu.medium.com/correlation-of-linear-regression-with-neural-networks-595eae15a322

From Linear Regression to Neural Networks

A single-layer (shallow) neural network

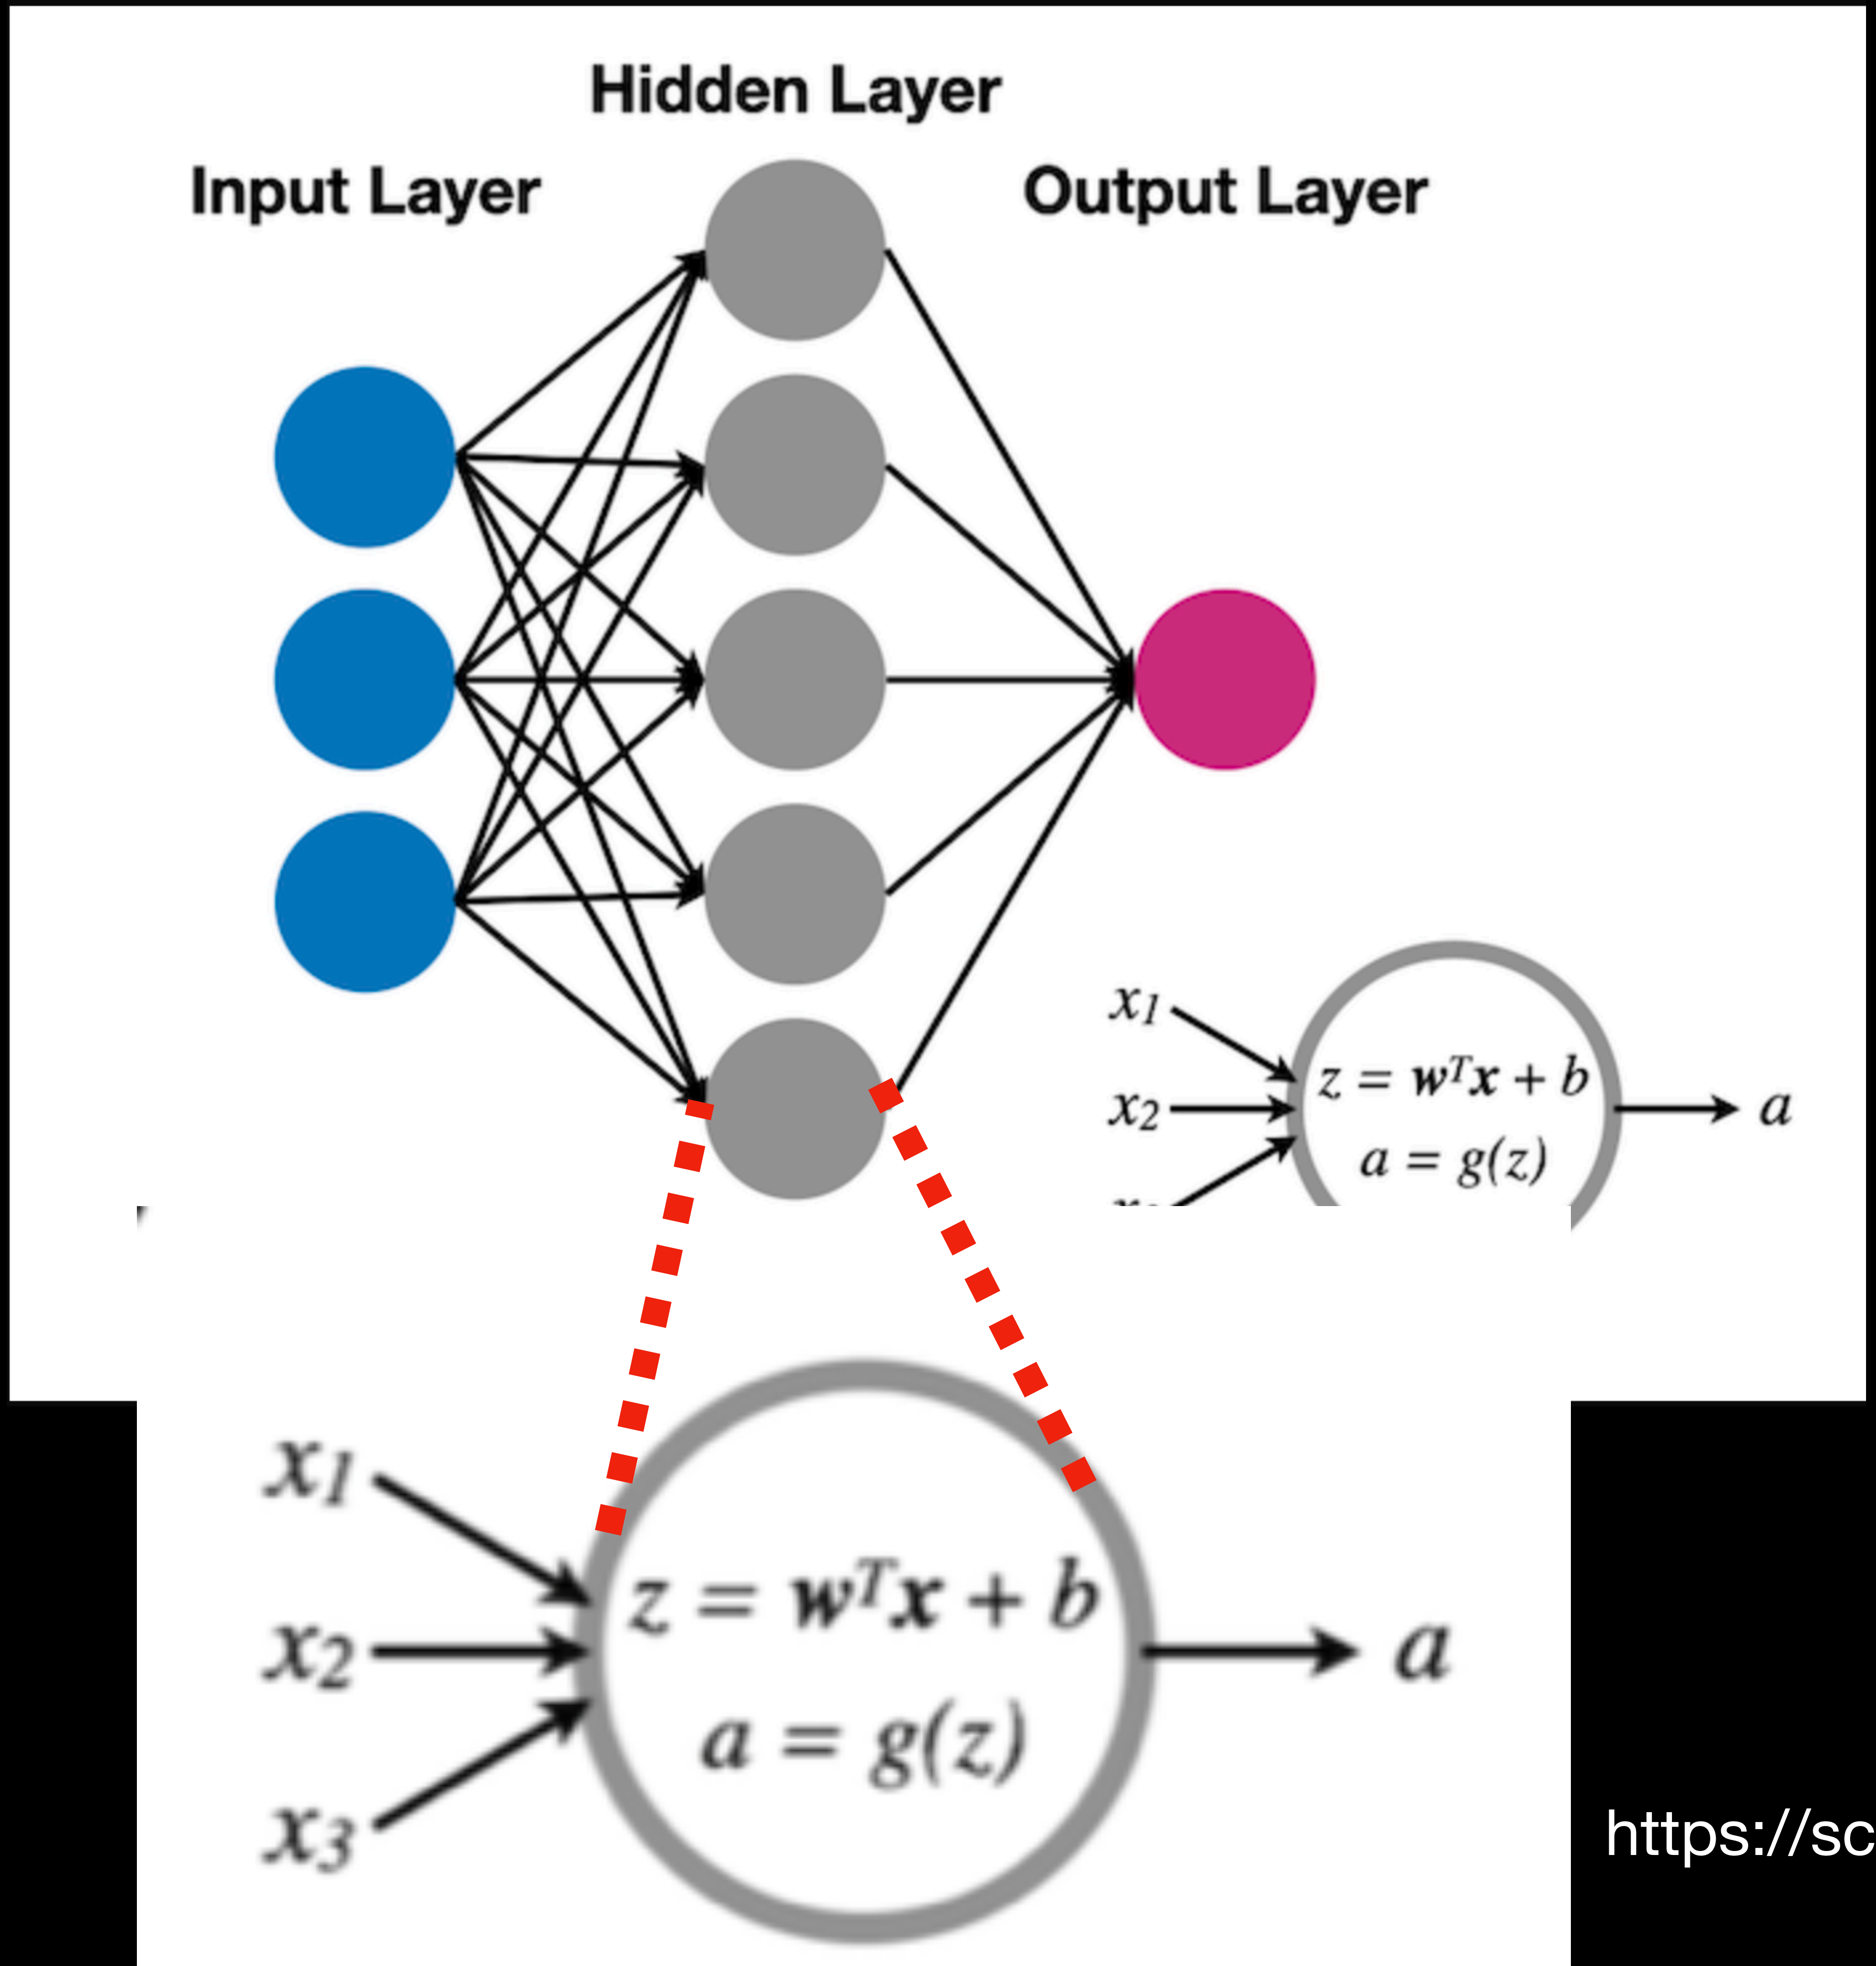

One *neuron*

A multi-layer (deep) neural network

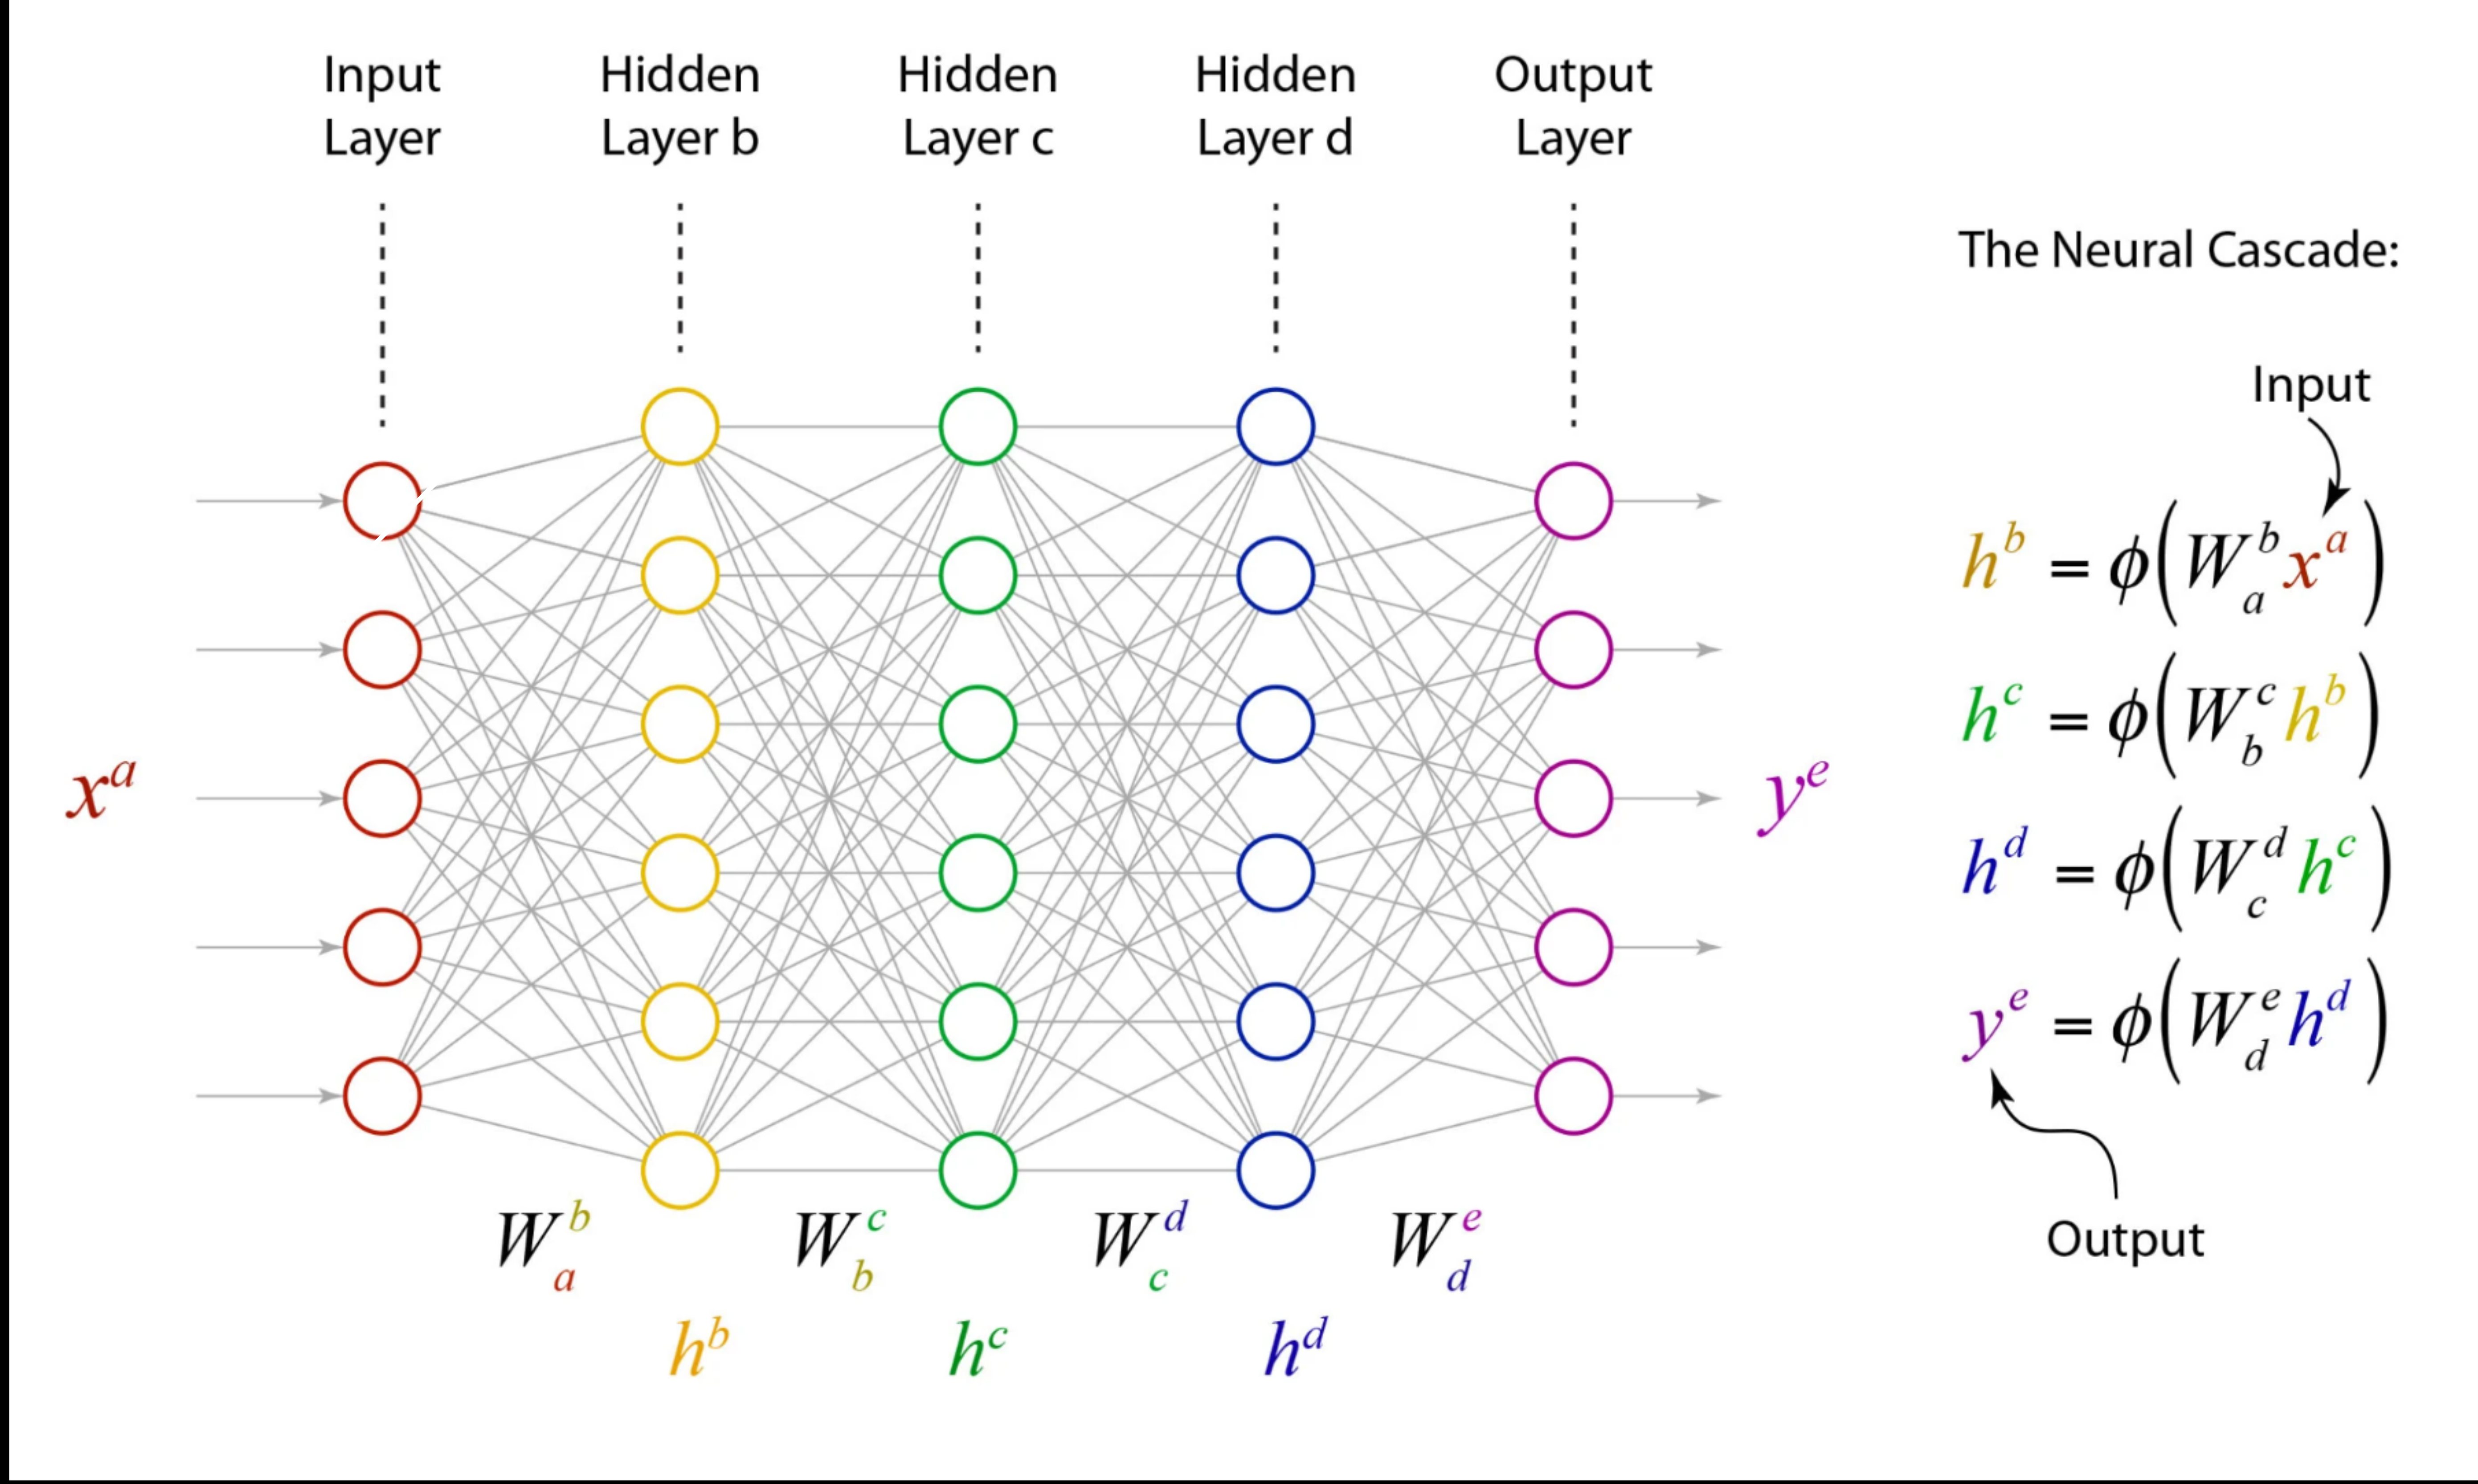

https://scipython.com/static/media/uploads/blog/shallow-neural-net/snn.png

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

Relating back to astronomy

Deep learning can be applied to many fields; here are some applications in astronomy

Big Data enables Big Models

- As mentioned before, deep models need lots of data to learn from due to them having so many parameters. The more complex the model, the more data needed.
- One application is light curve/intensity-over-time analysis, for:
    - exoplanet detection
    - supernova classification
    - etc…

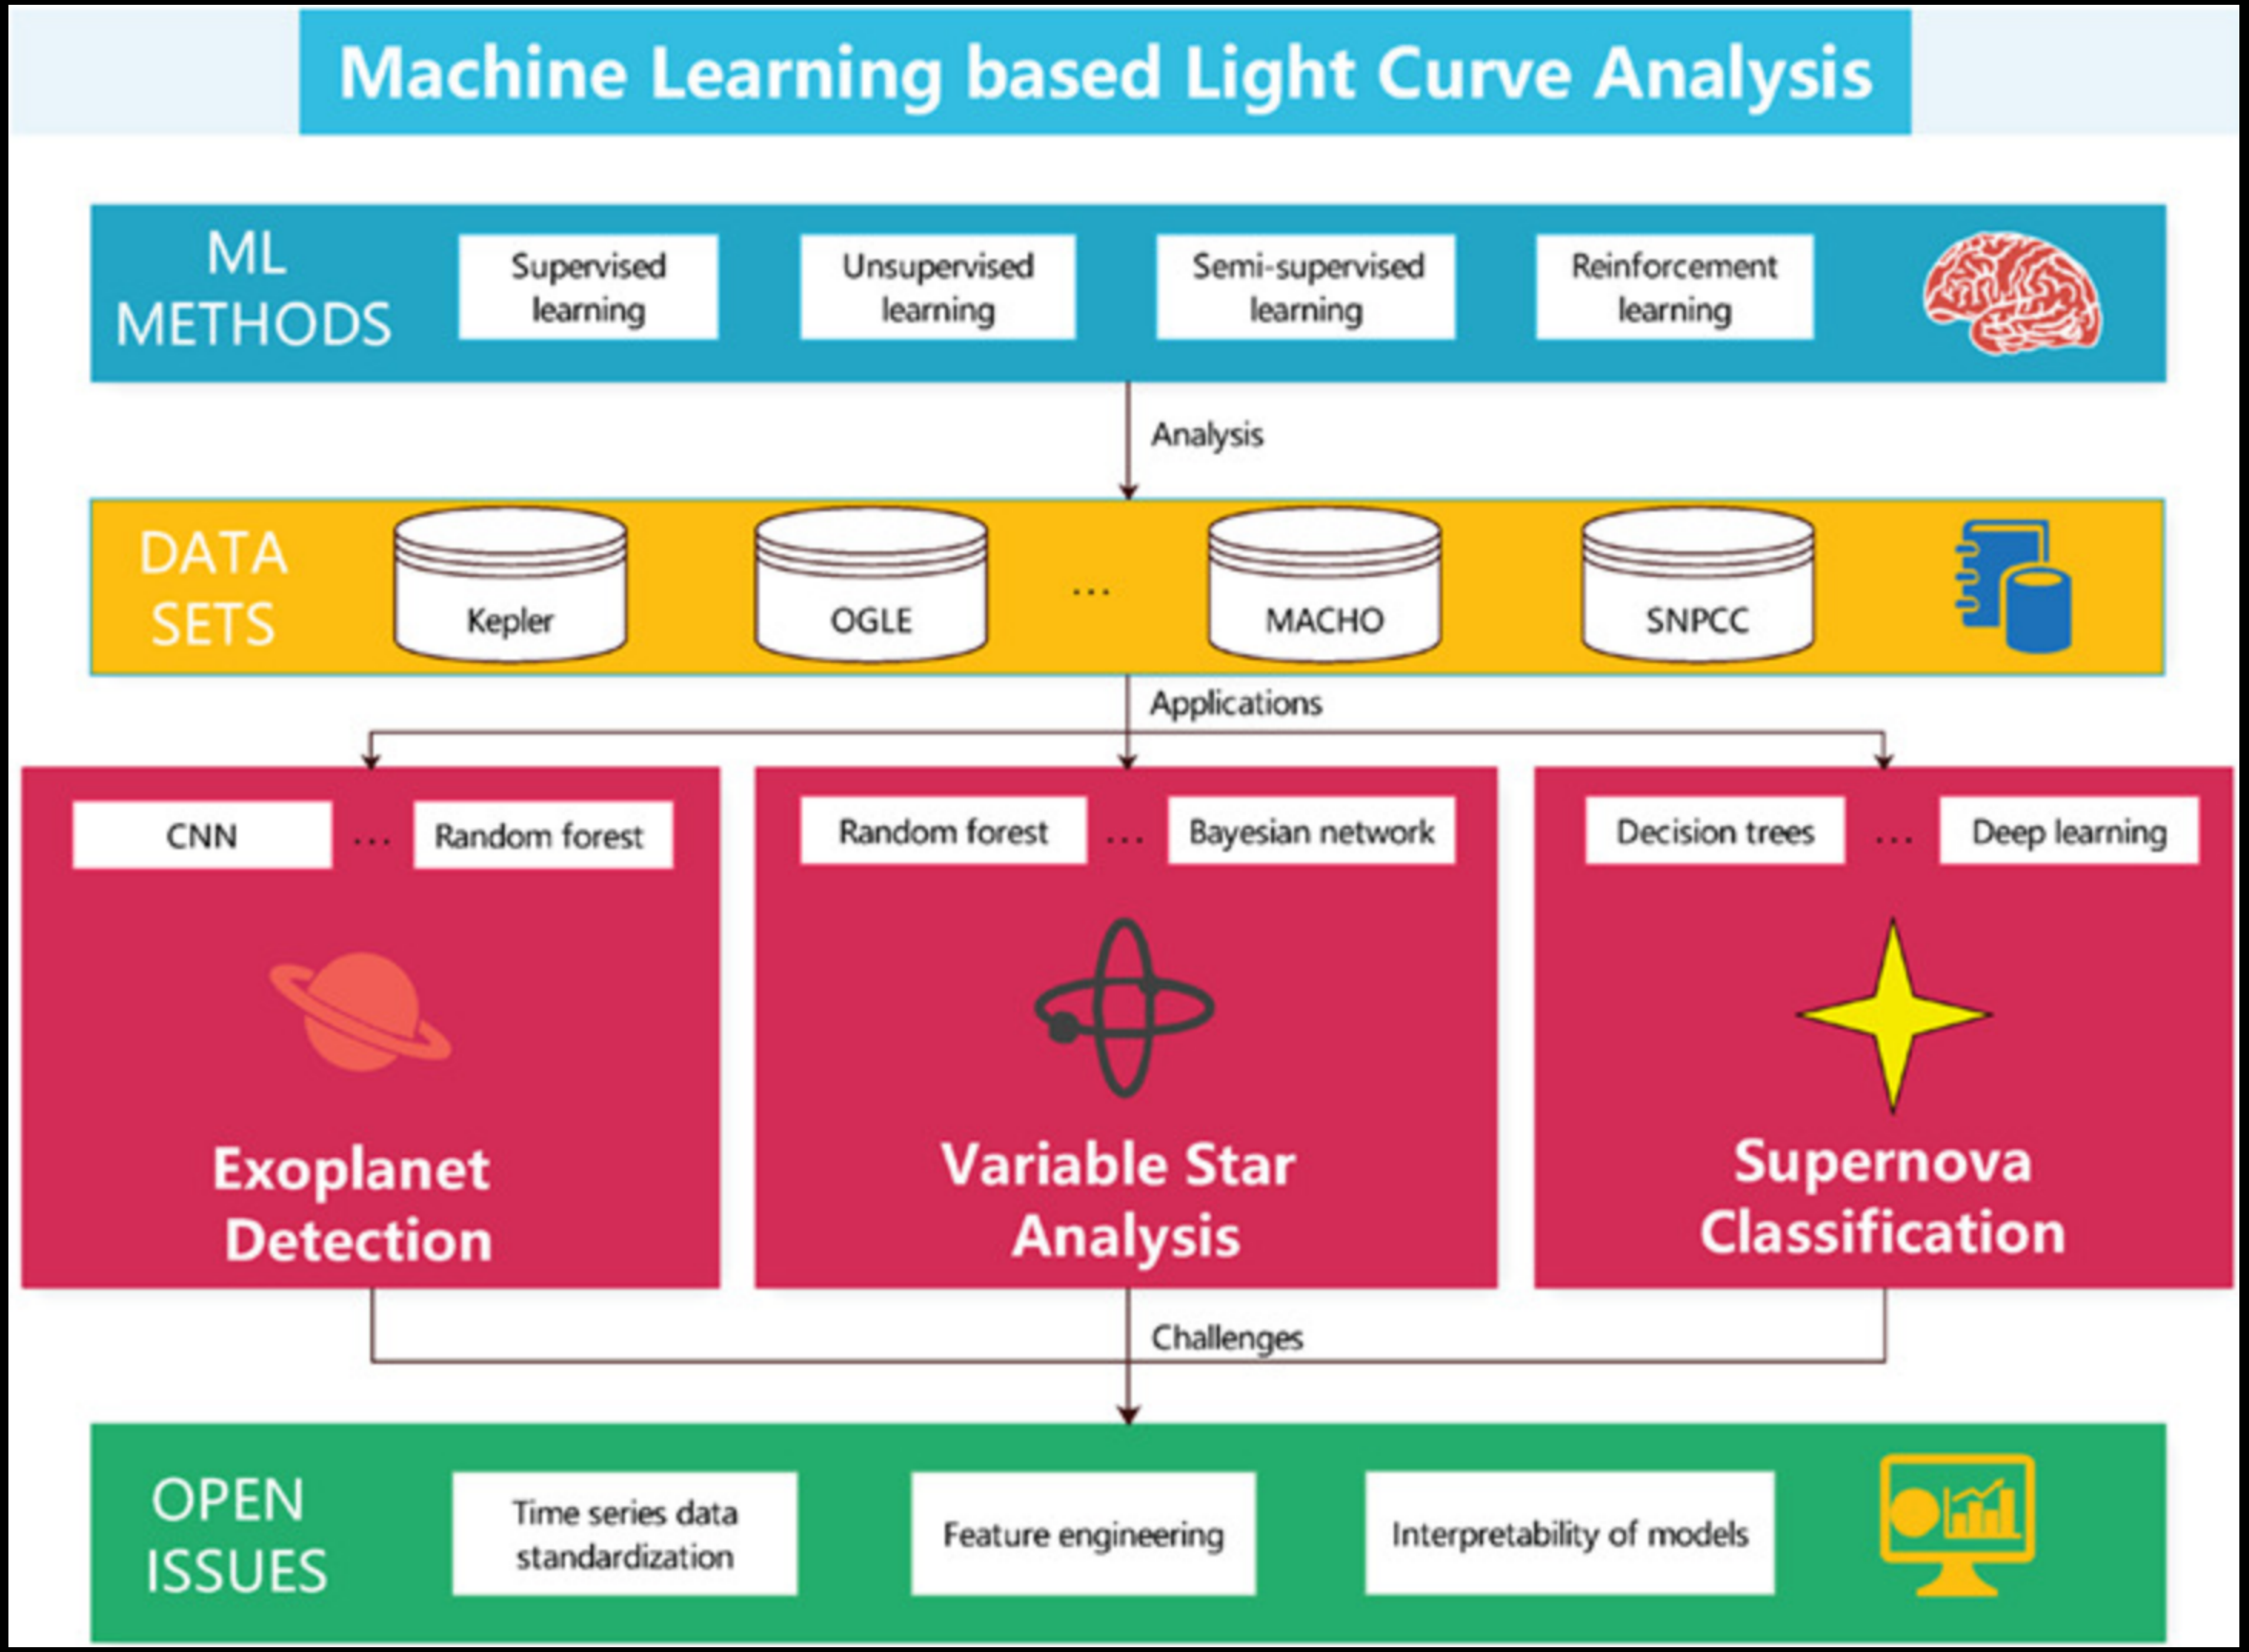

Yu, Ce, et al. "A survey on machine learning based light curve analysis for variable astronomical sources." *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 11.5 (2021): e1425.

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

- A neural network was trained on almost 40,000 images of galaxies from the Galaxy Zoo dataset, and tested on a another ~10,000 from the SDSS and DES survey datasets, for the task of classifying galaxy images as spiral or elliptical

- Achieved an accuracy on the test datasets of $\geq 99.6\,\%$

- My team project is doing something similar!

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

- Train a neural network to remove noise from images!

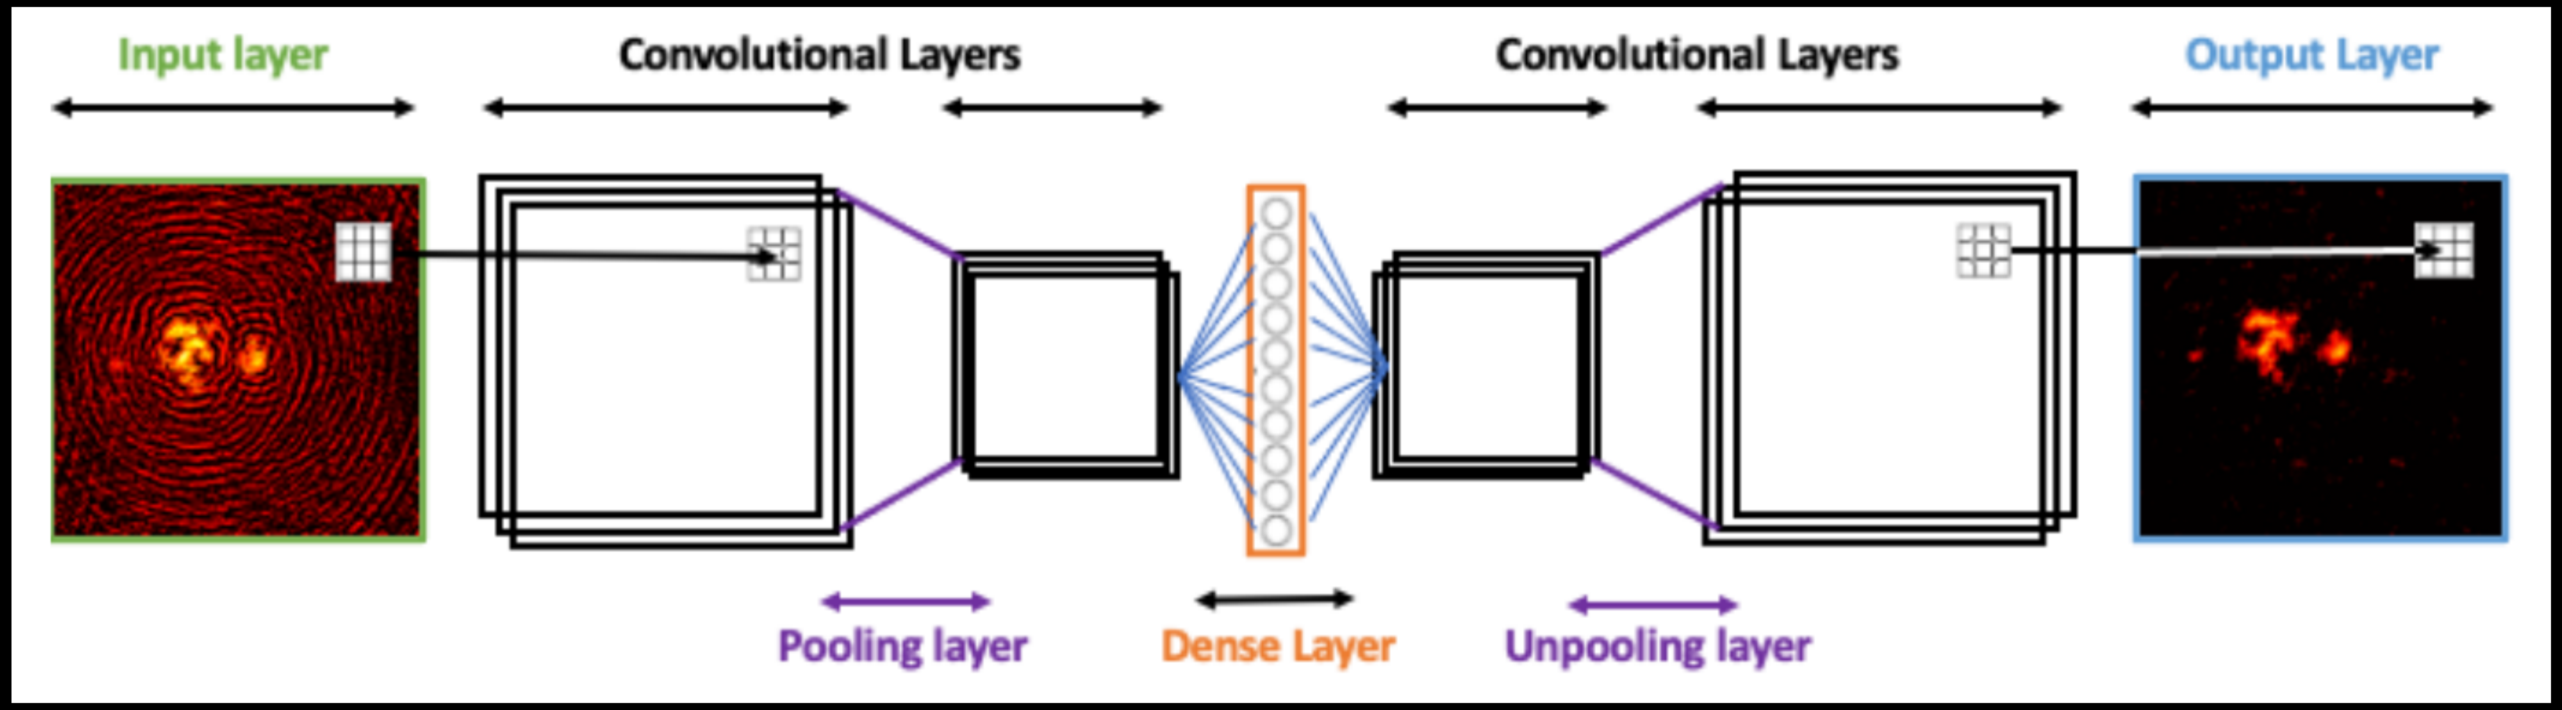

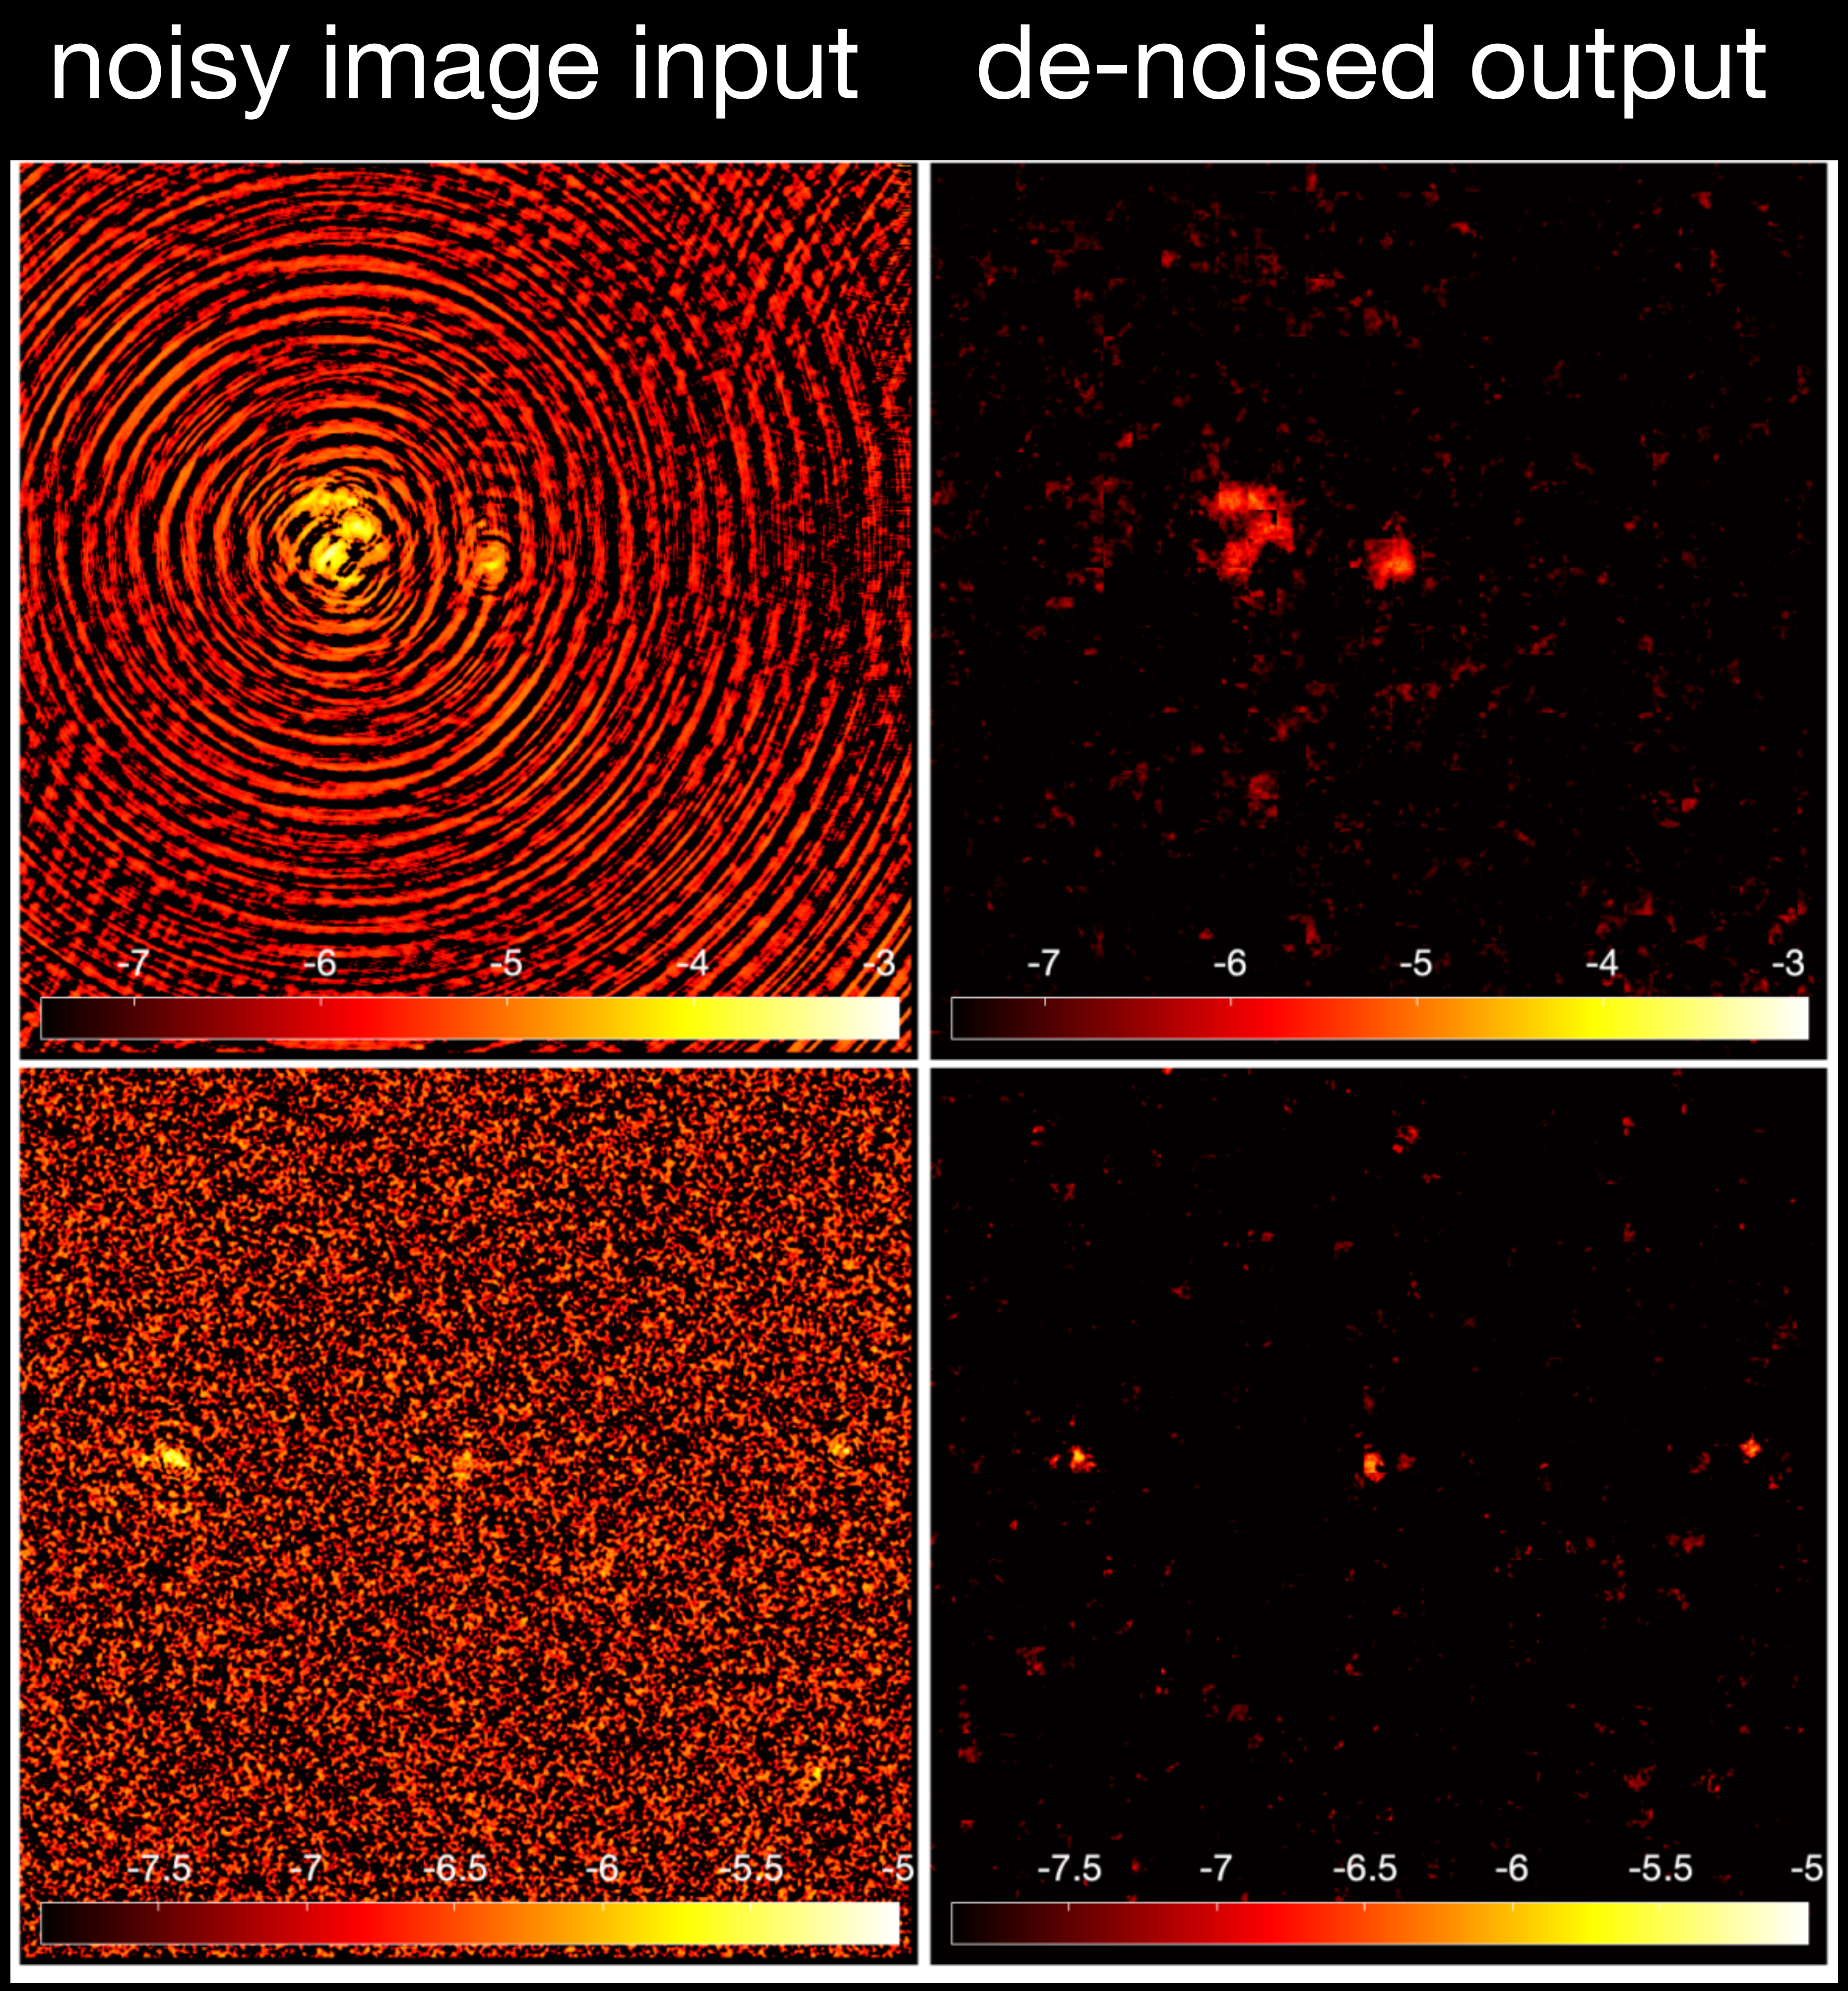

Gheller, Claudio, and Franco Vazza. "Convolutional deep denoising autoencoders for radio astronomical images." Monthly Notices of the Royal Astronomical Society 509.1 (2022): 990-1009.

And many others…

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

Recognize this person? One of these photos isn’t real

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

1. Computer Vision

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

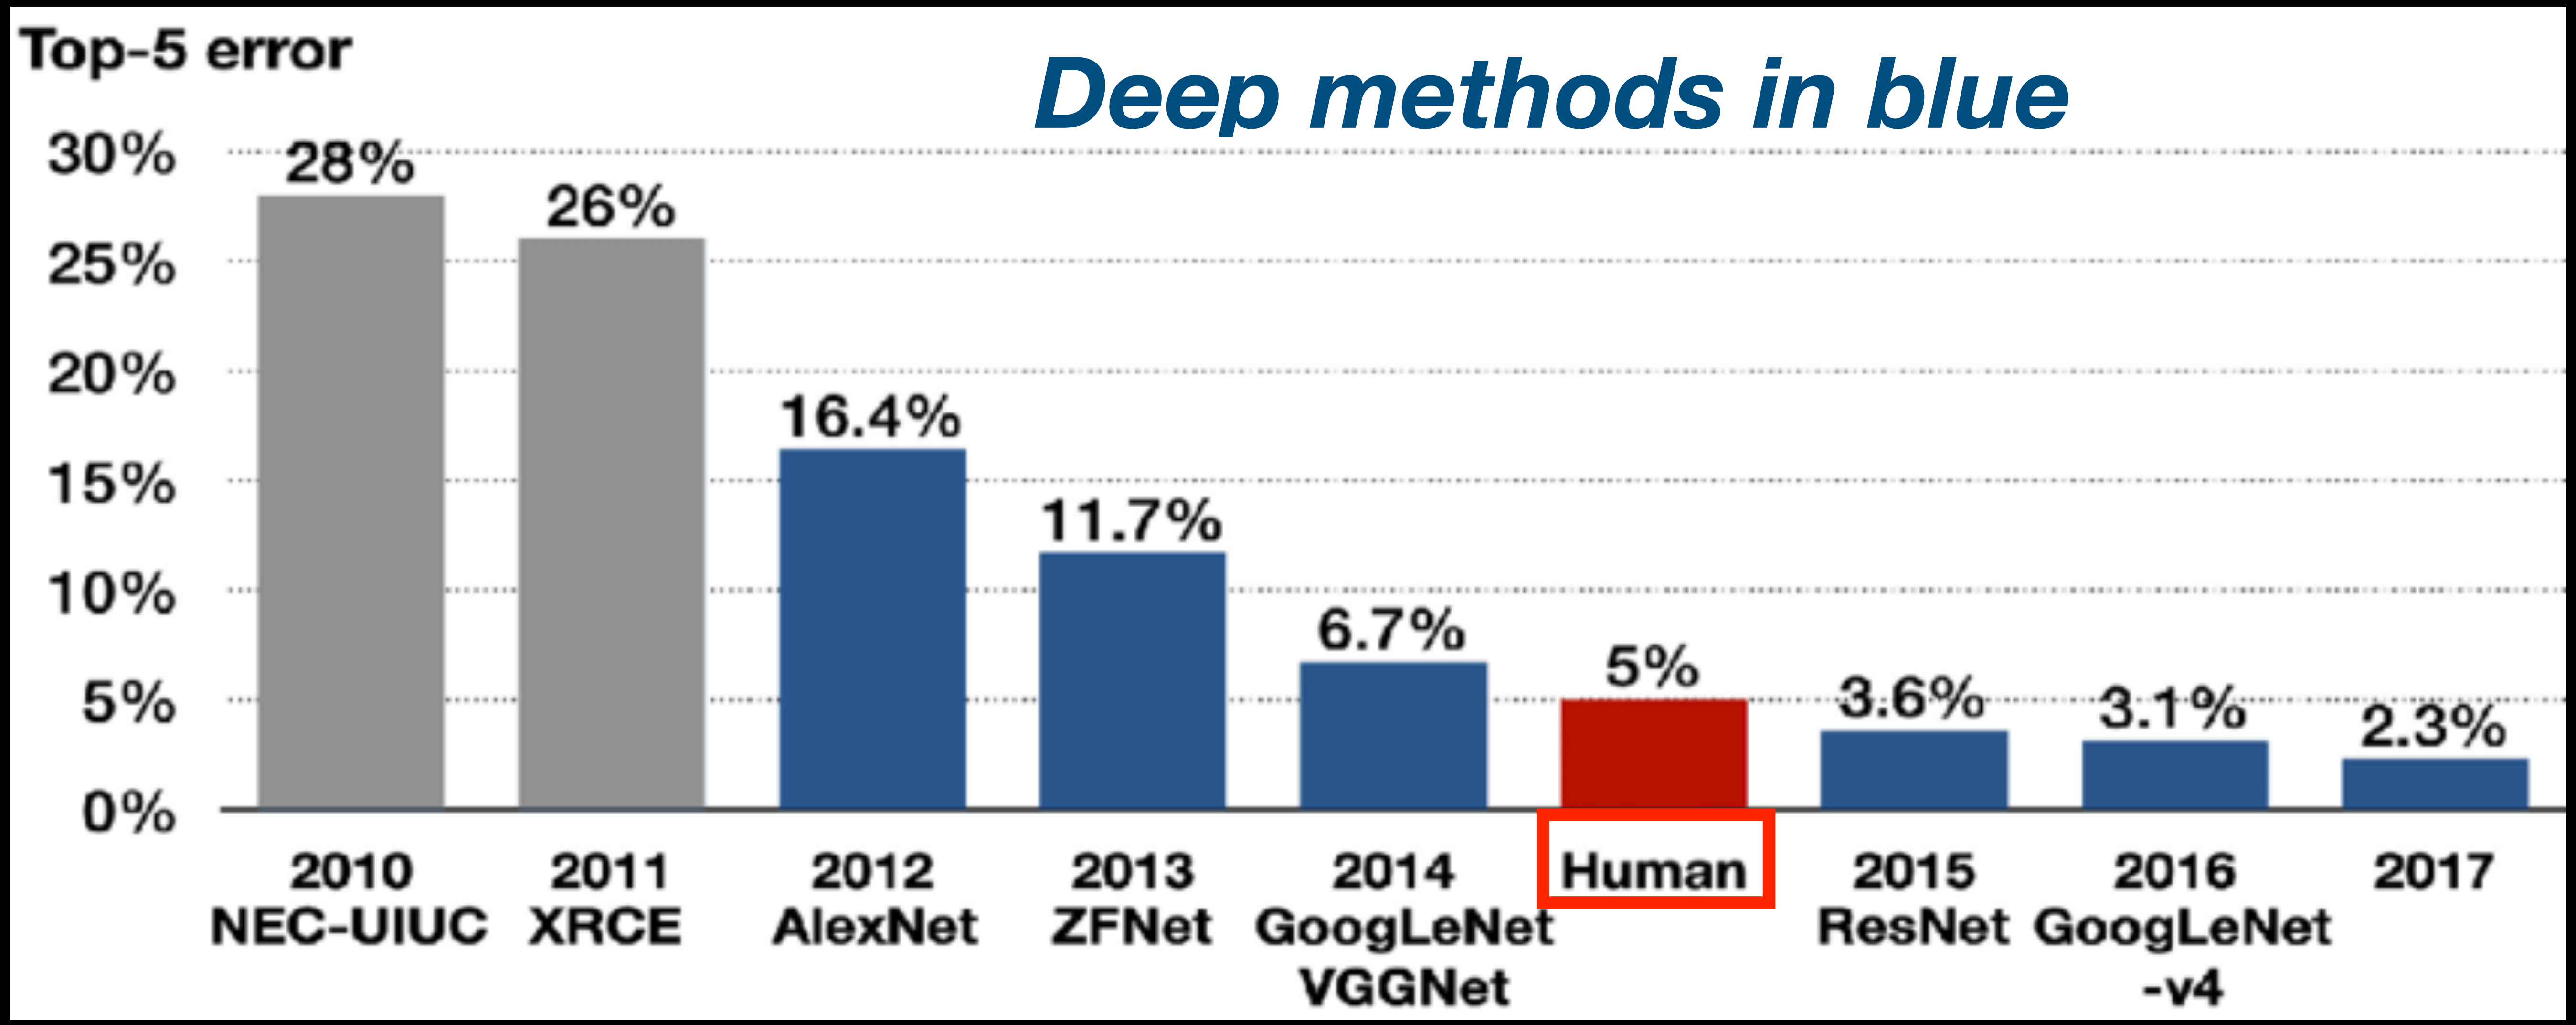

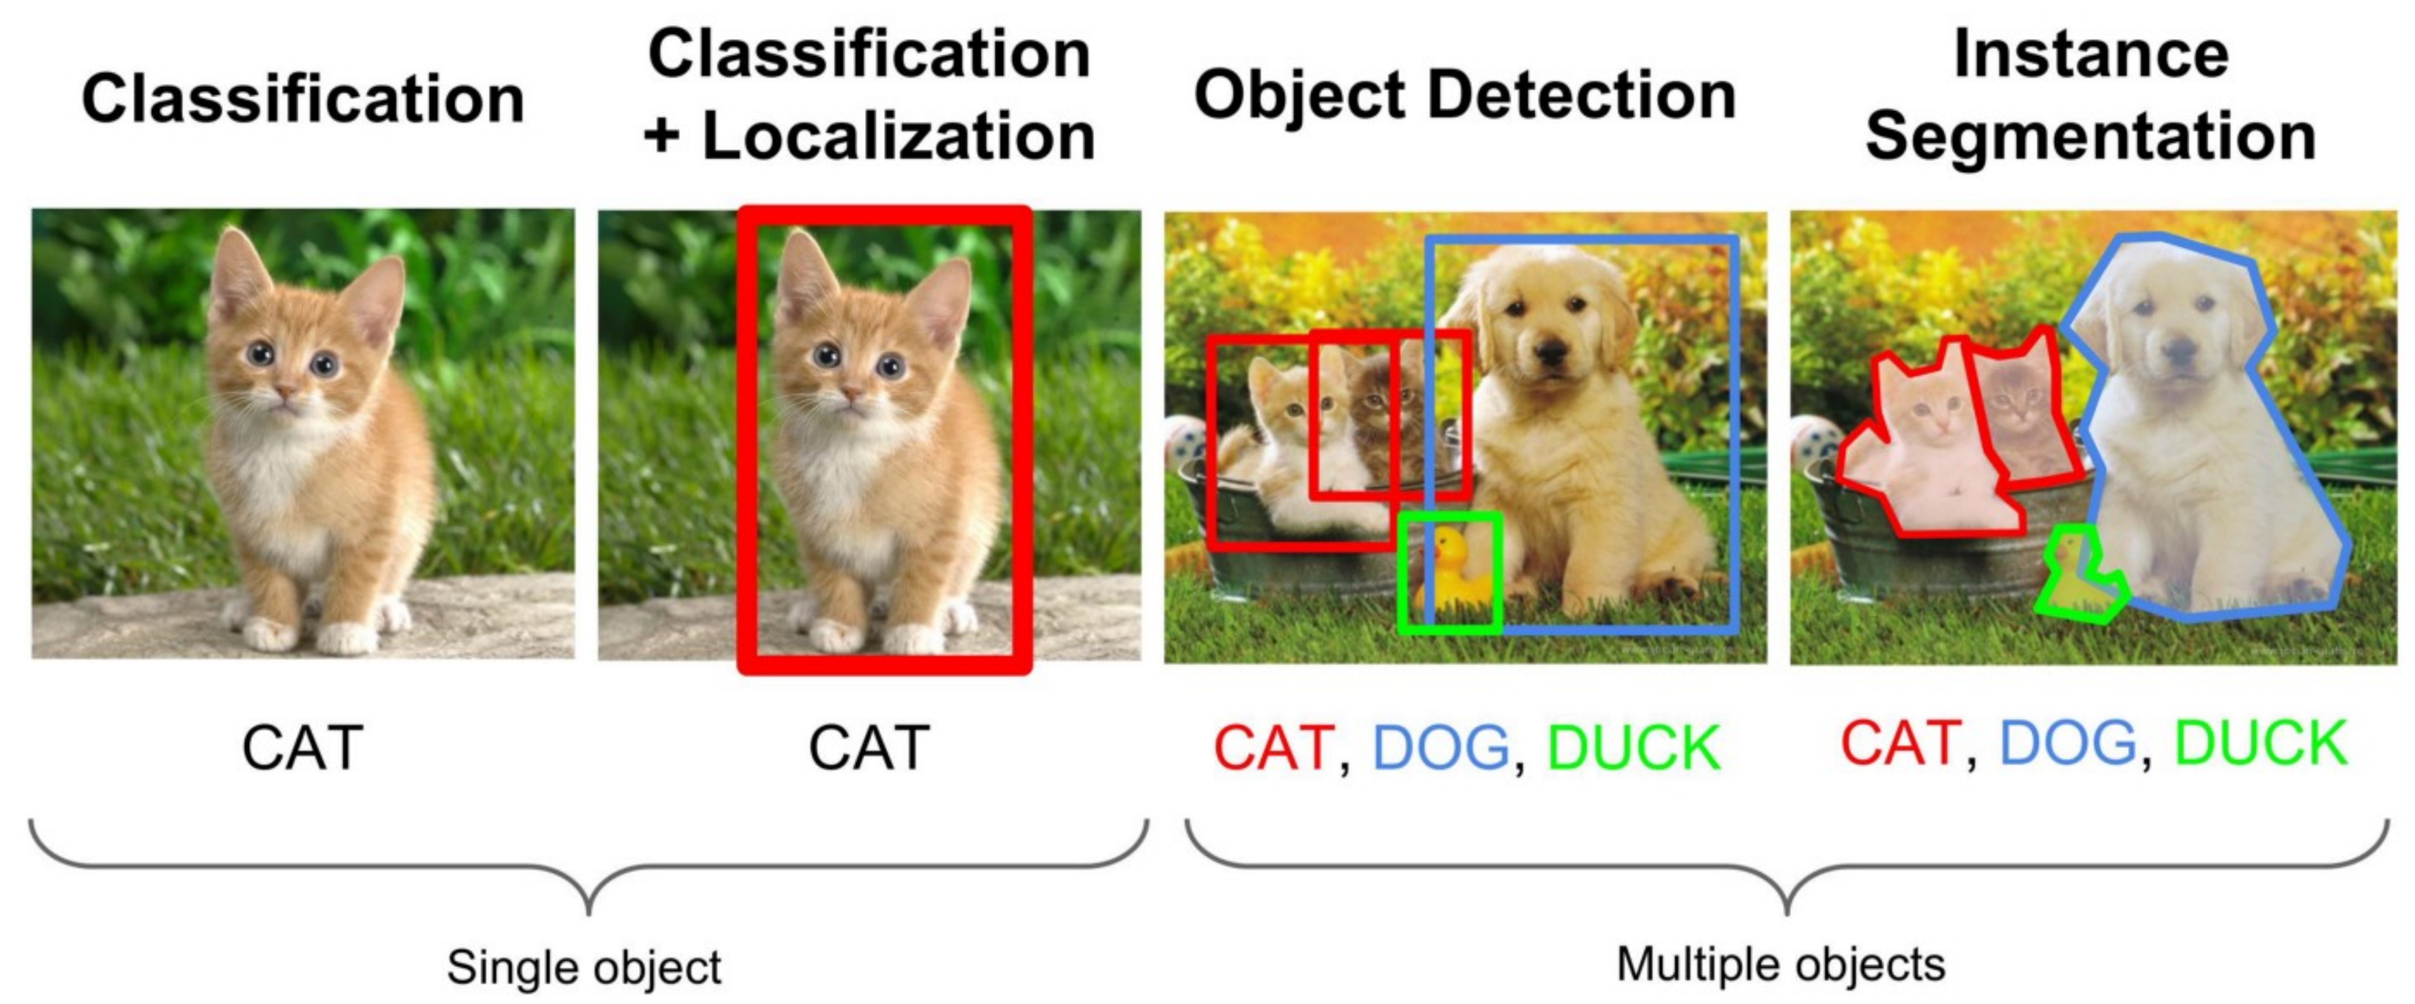

2. image/art generation

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

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

2. image/art generation (I created these with DALL-E @ https://labs.openai.com/)

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

Prompt: “A photo of a bear doing jiu jitsu with a monkey in a dojo”

Prompt for DALL-E:

"'Dan Reichart devouring his pizza', painting by Francisco Goya"

The Good

Deep learning has created *huge* advancements in (continued):

3. natural language processing (speech recognition, text generation, etc
4. Drug discovery and toxicology
5. Recommendation systems
6. Bioinformatics and medical image analysis (my field)
7. Fraud detection
8. Solving partial differential equations in physics

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

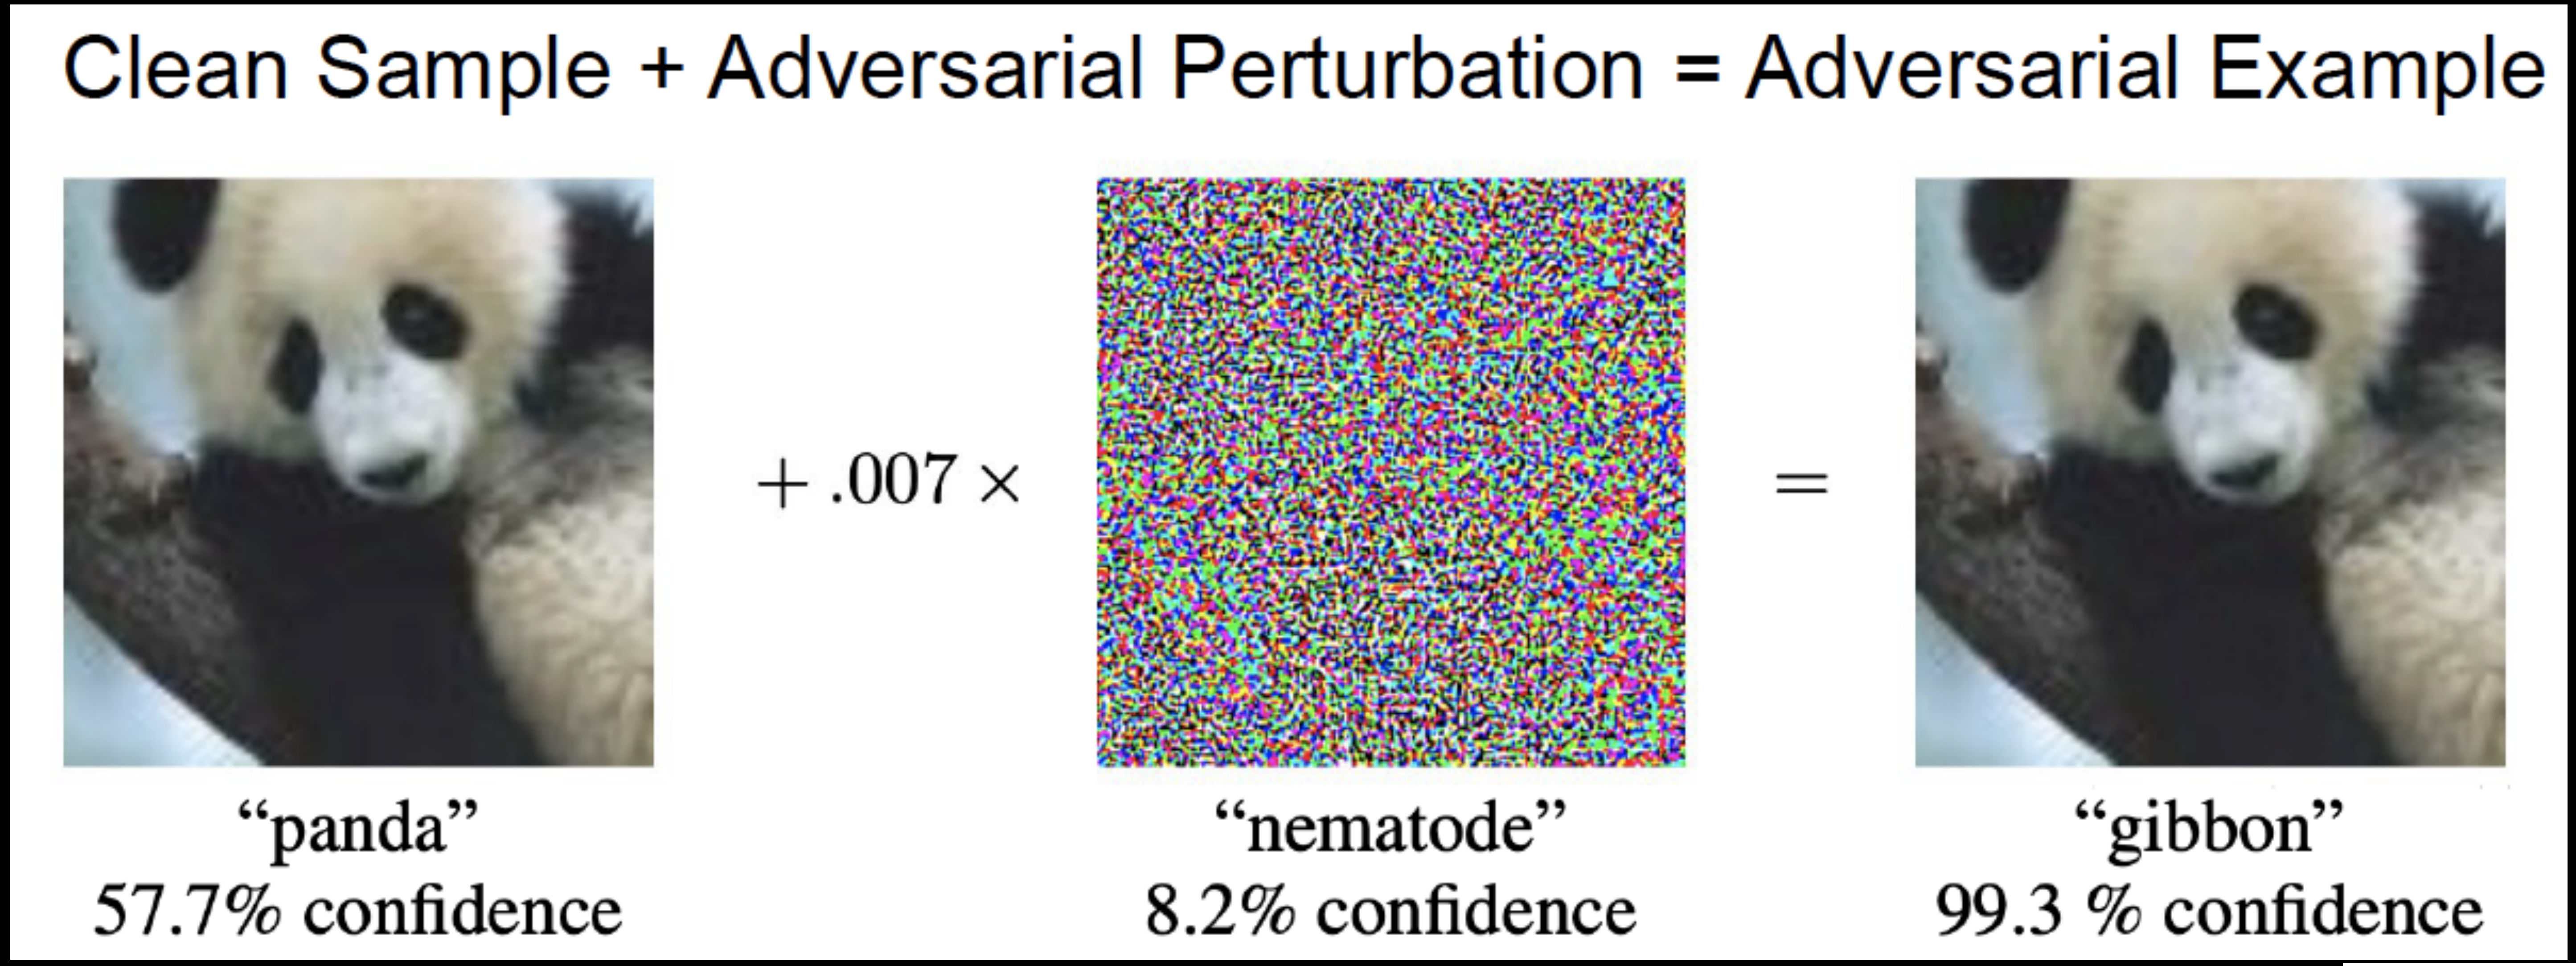

Goodfellow, Ian J., Jonathon Shlens, and Christian Szegedy. "Explaining and harnessing adversarial examples." arXiv preprint arXiv:1412.6572 (2014).

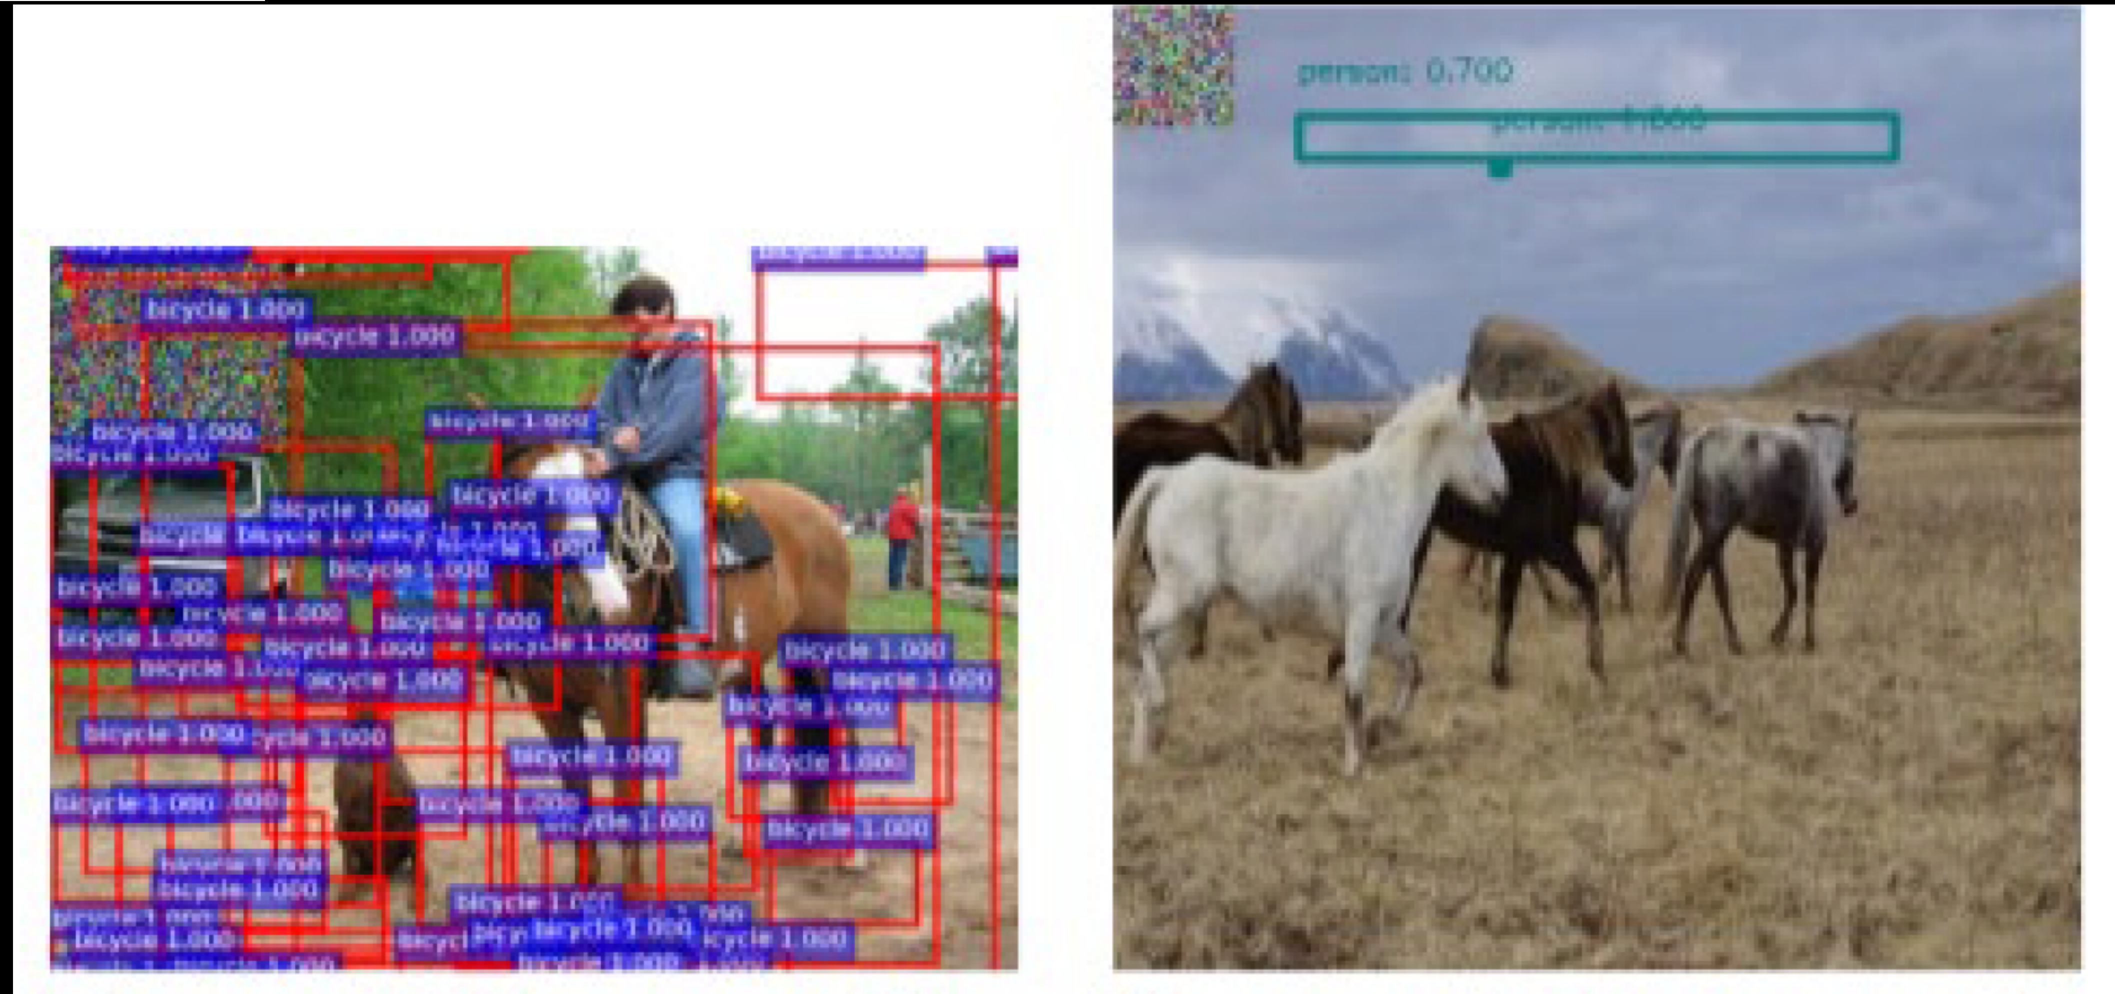

(a) targeted DPATCH attacking Faster R-CNN (b) targeted DPATCH attacking YOLO

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

1. Deepfakes
2. Mass surveillance/facial recognition of “undesirable” populations
3. Human-mistakable natural language and image synthesis

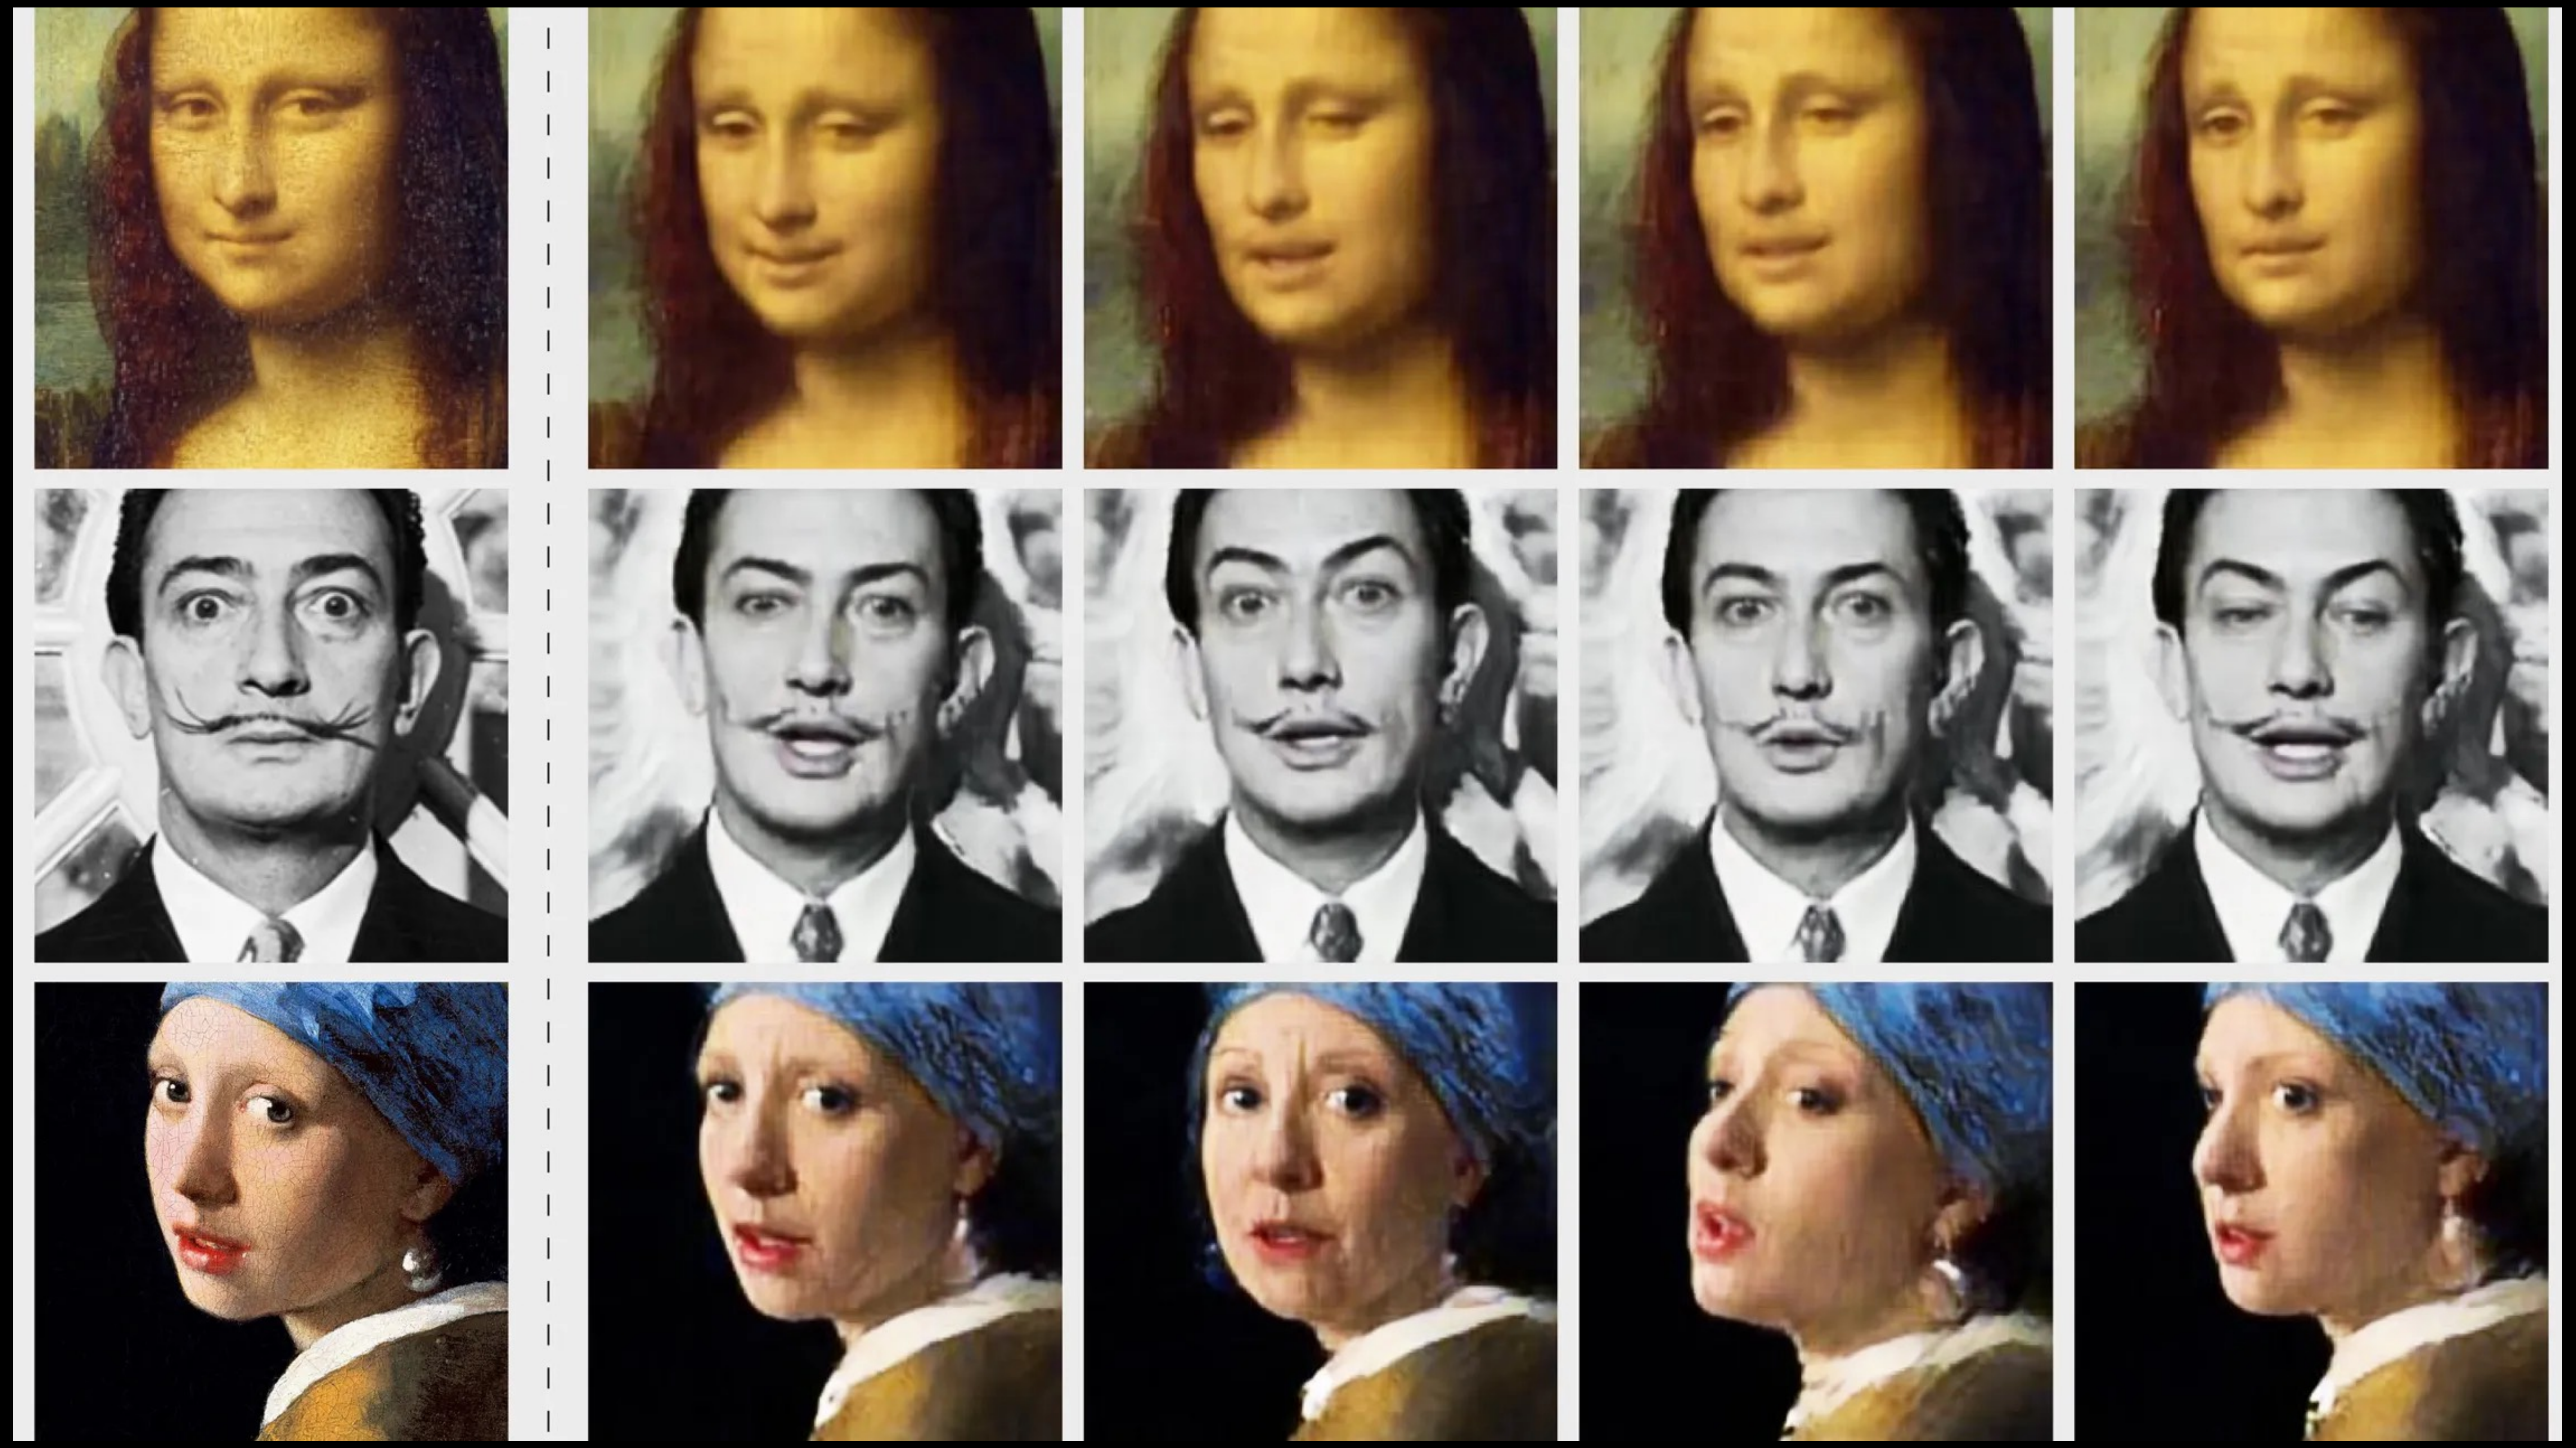

https://cdn.thegeekherald.com/wp-content/uploads/2019/06/New-AI-deepfake-app-creates-nude-images-of-women-in-seconds-1-e1561666459843.jpg

Thanks for Listening! Questions?