THE NEW ERA: AI

Dr. Şeyda Ertekin

Associate Professor

ODTÜ Department of Computer Engineering

Director

ODTÜ CENG Smart Systems Lab.

ODTÜ Digital Innovation Center – AI and Big Data Analytics Lab.

Email: sertekin@metu.edu.tr

@theaiqueen

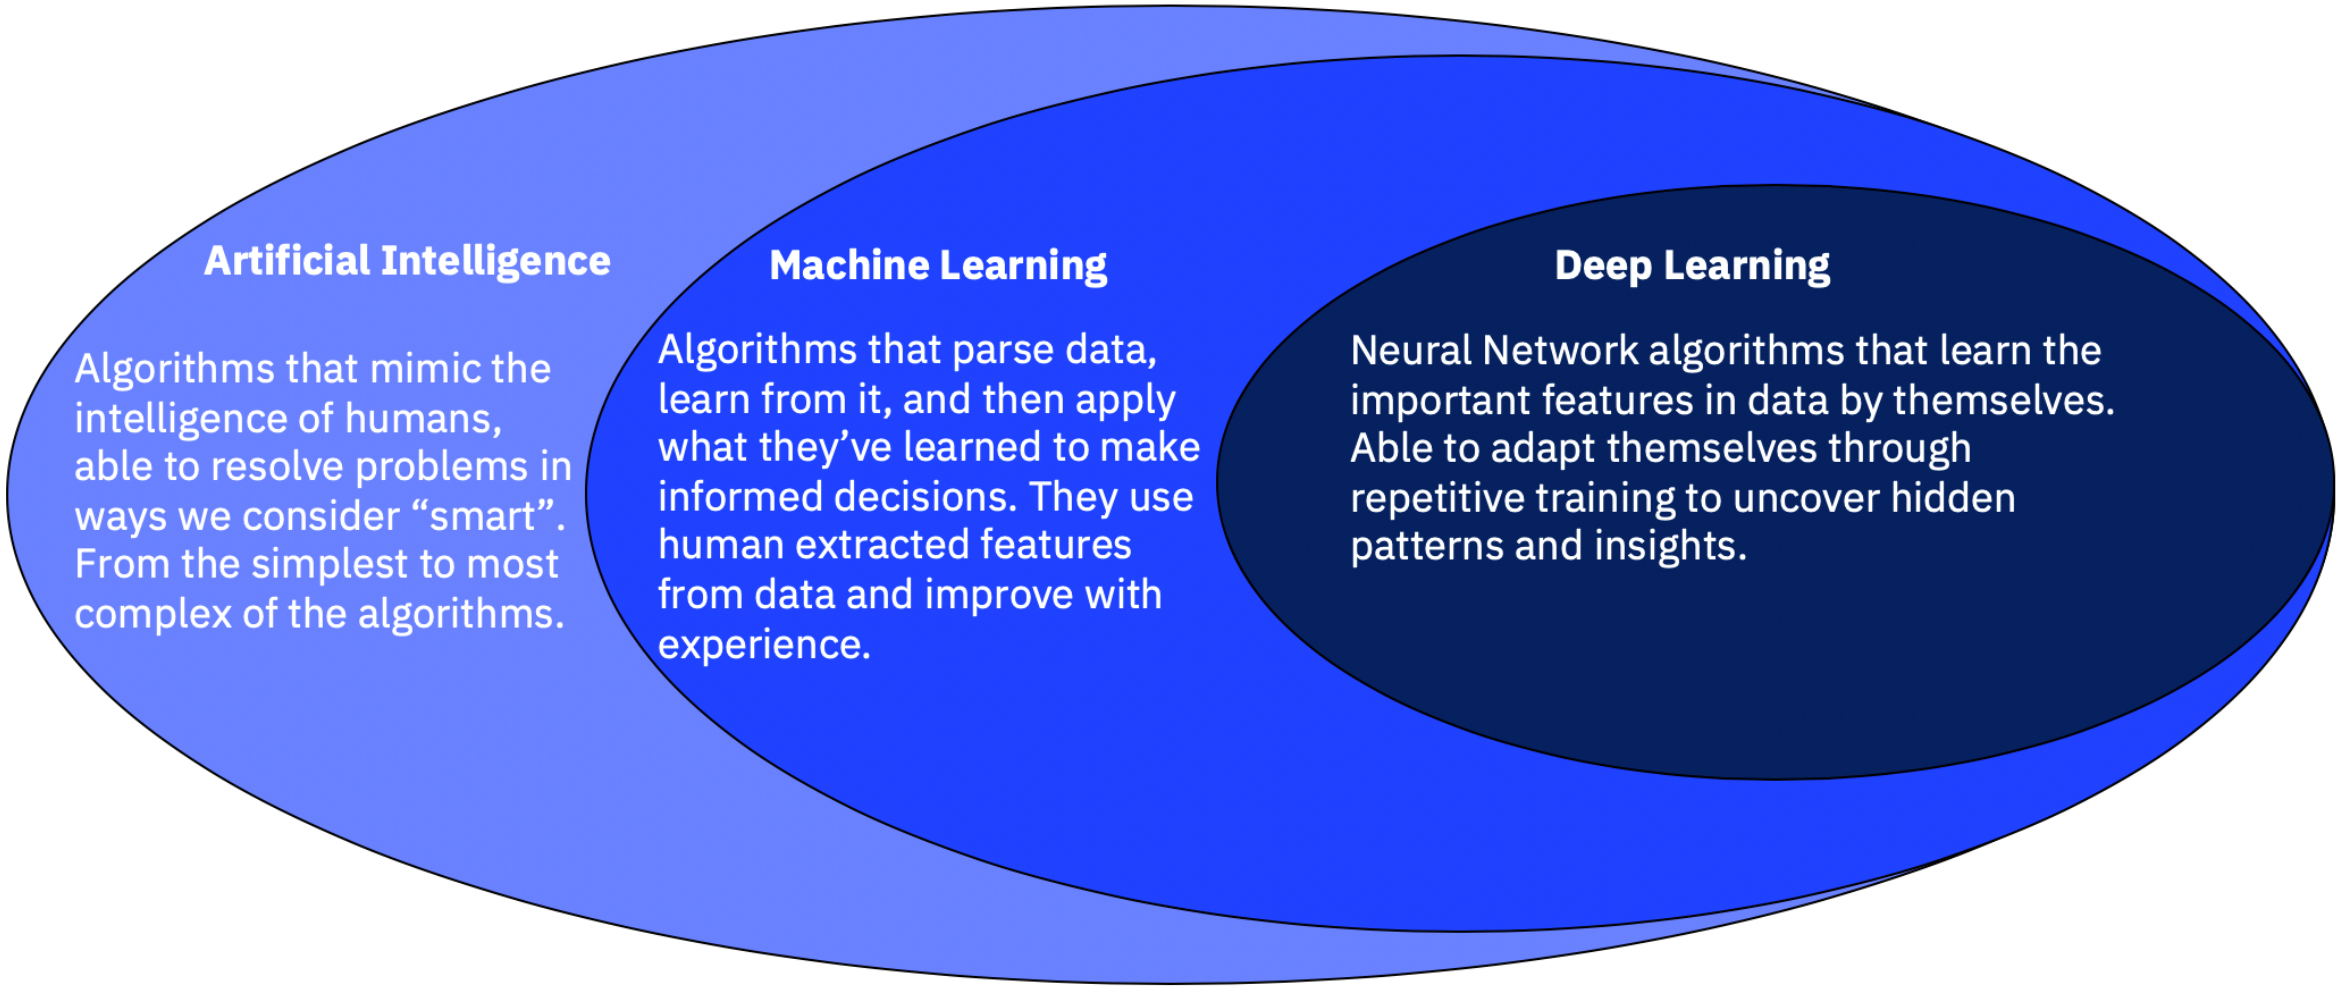
Artificial Intelligence
Algorithms that mimic the intelligence of humans, able to resolve problems in ways we consider “smart”. From the simplest to most complex of the algorithms.
Machine Learning
Algorithms that parse data, learn from it, and then apply what they’ve learned to make informed decisions. They use human extracted features from data and improve with experience.
Deep Learning
Neural Network algorithms that learn the important features in data by themselves. Able to adapt themselves through repetitive training to uncover hidden patterns and insights.

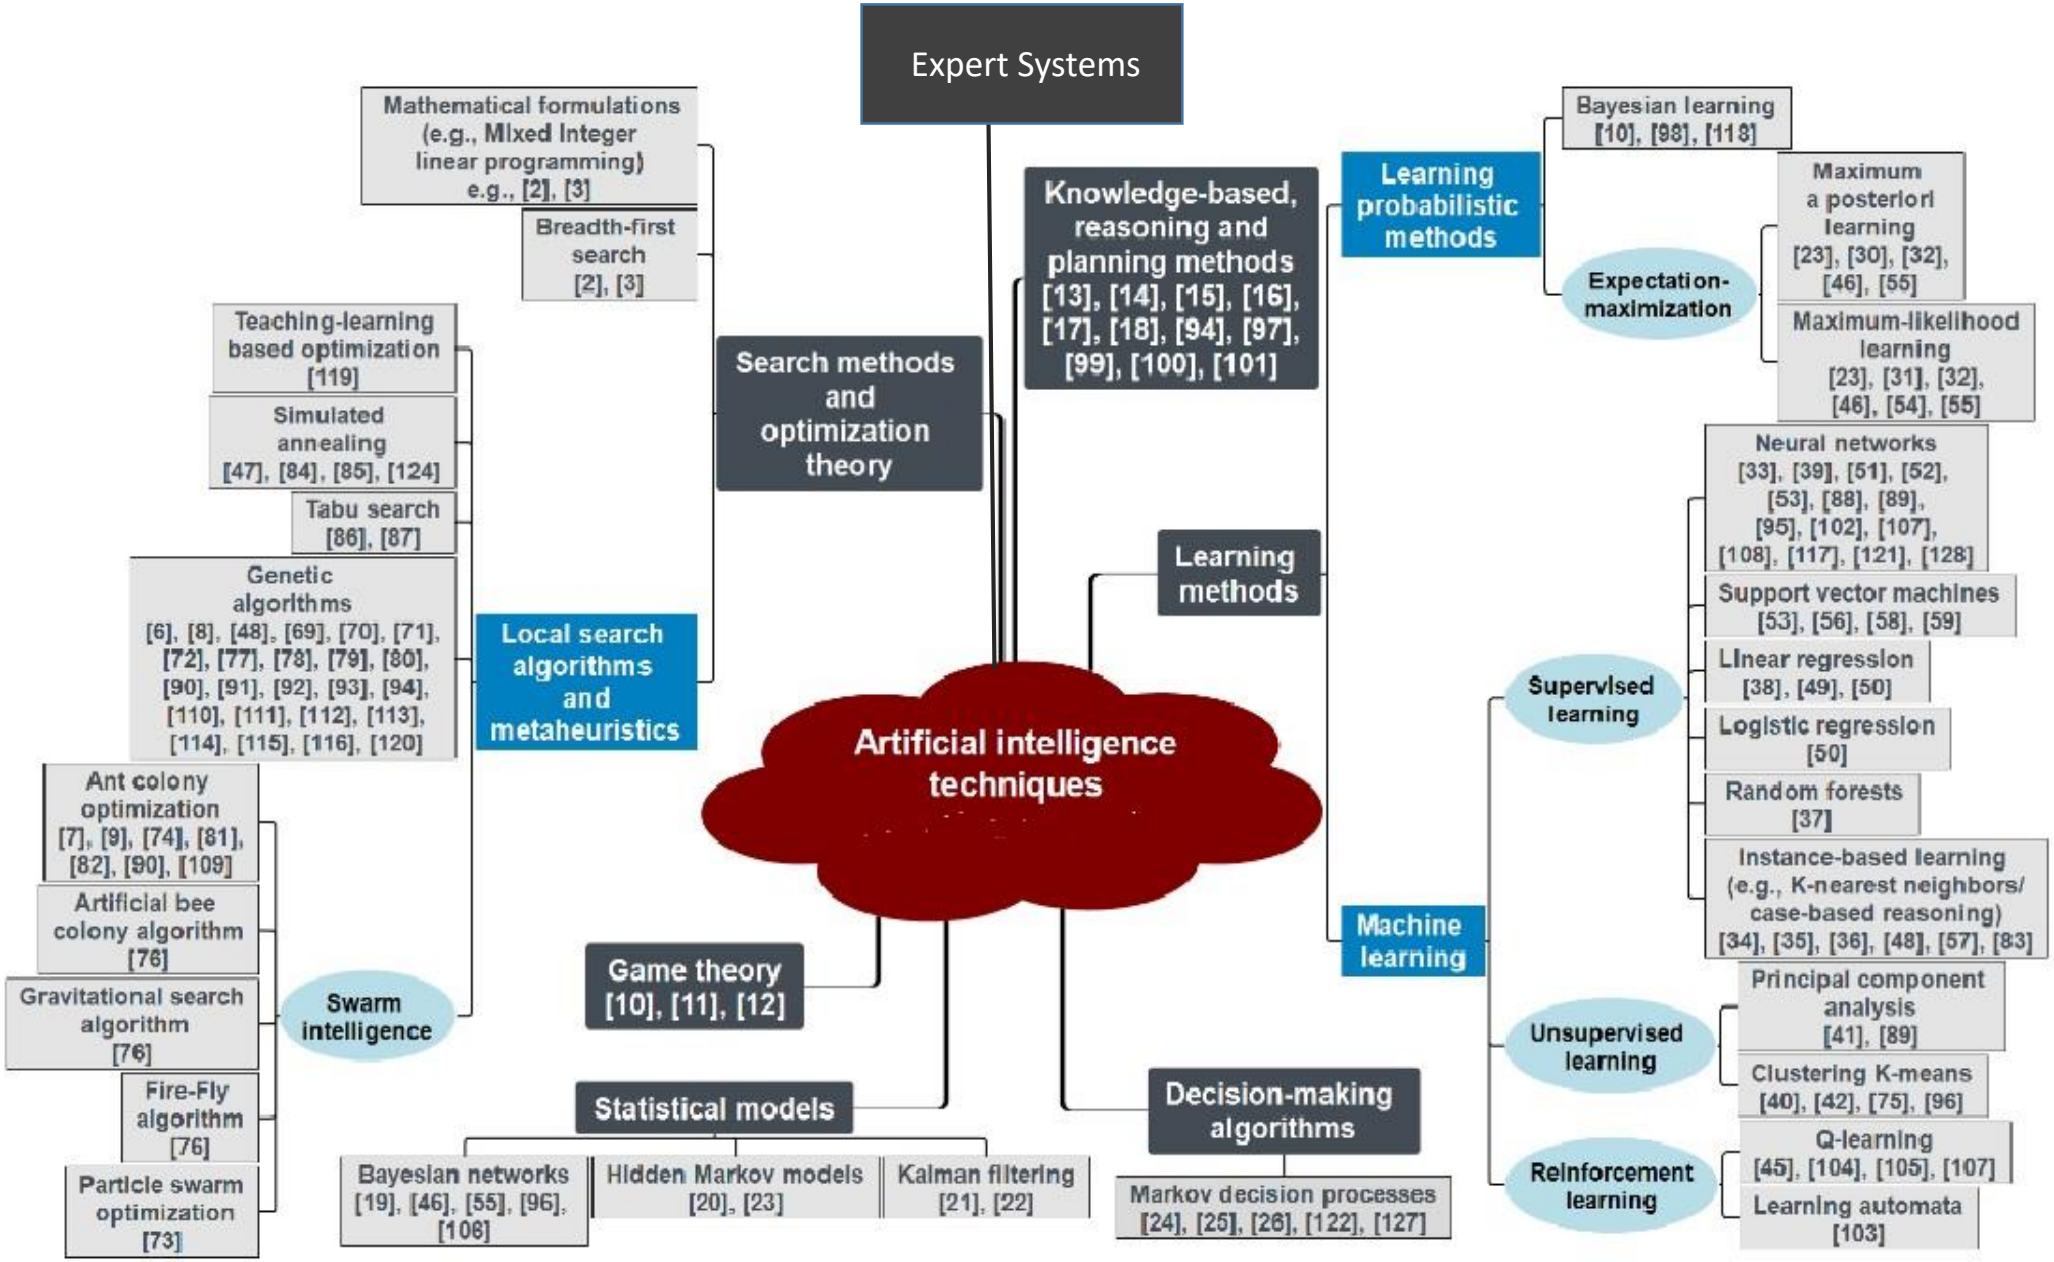
Expert Systems
Artificial intelligence techniques
Search methods and optimization theory
Mathematical formulations (e.g., Mixed Integer linear programming) e.g., [2], [3]
Breadth-first search [2], [3]
Local search algorithms and metaheuristics
Teaching-learning based optimization [119]
Simulated annealing [47], [84], [85], [124]
Tabu search [86], [87]
Genetic algorithms [6], [8], [48], [69], [70], [71], [72], [77], [78], [79], [80], [90], [91], [92], [93], [94], [110], [111], [112], [113], [114], [115], [116], [120]
Swarm intelligence
Ant colony optimization [7], [9], [74], [81], [82], [90], [109]
Artificial bee colony algorithm [76]
Gravitational search algorithm [76]
Fire-Fly algorithm [76]
Particle swarm optimization [73]
Knowledge-based, reasoning and planning methods [13], [14], [15], [16], [17], [18], [94], [97], [99], [100], [101]
Learning methods
Machine learning
Learning probabilistic methods
Bayesian learning [10], [98], [118]
Expectation-maximization
Maximum a posteriori learning [23], [30], [32], [46], [55]
Maximum-likelihood learning [23], [31], [32], [46], [54], [55]
Supervised learning
Neural networks [33], [39], [51], [52], [53], [88], [89], [95], [102], [107], [108], [117], [121], [128]
Support vector machines [53], [56], [58], [59]
Linear regression [38], [49], [50]
Logistic regression [50]
Random forests [37]
Instance-based learning (e.g., K-nearest neighbors/ case-based reasoning) [34], [35], [36], [48], [57], [83]
Unsupervised learning
Principal component analysis [41], [89]
Clustering K-means [40], [42], [75], [96]
Reinforcement learning
Q-learning [45], [104], [105], [107]
Learning automata [103]
Game theory [10], [11], [12]
Statistical models
Bayesian networks [19], [46], [55], [96], [106]
Hidden Markov models [20], [23]
Kalman filtering [21], [22]
Decision-making algorithms
Markov decision processes [24], [25], [26], [122], [127]

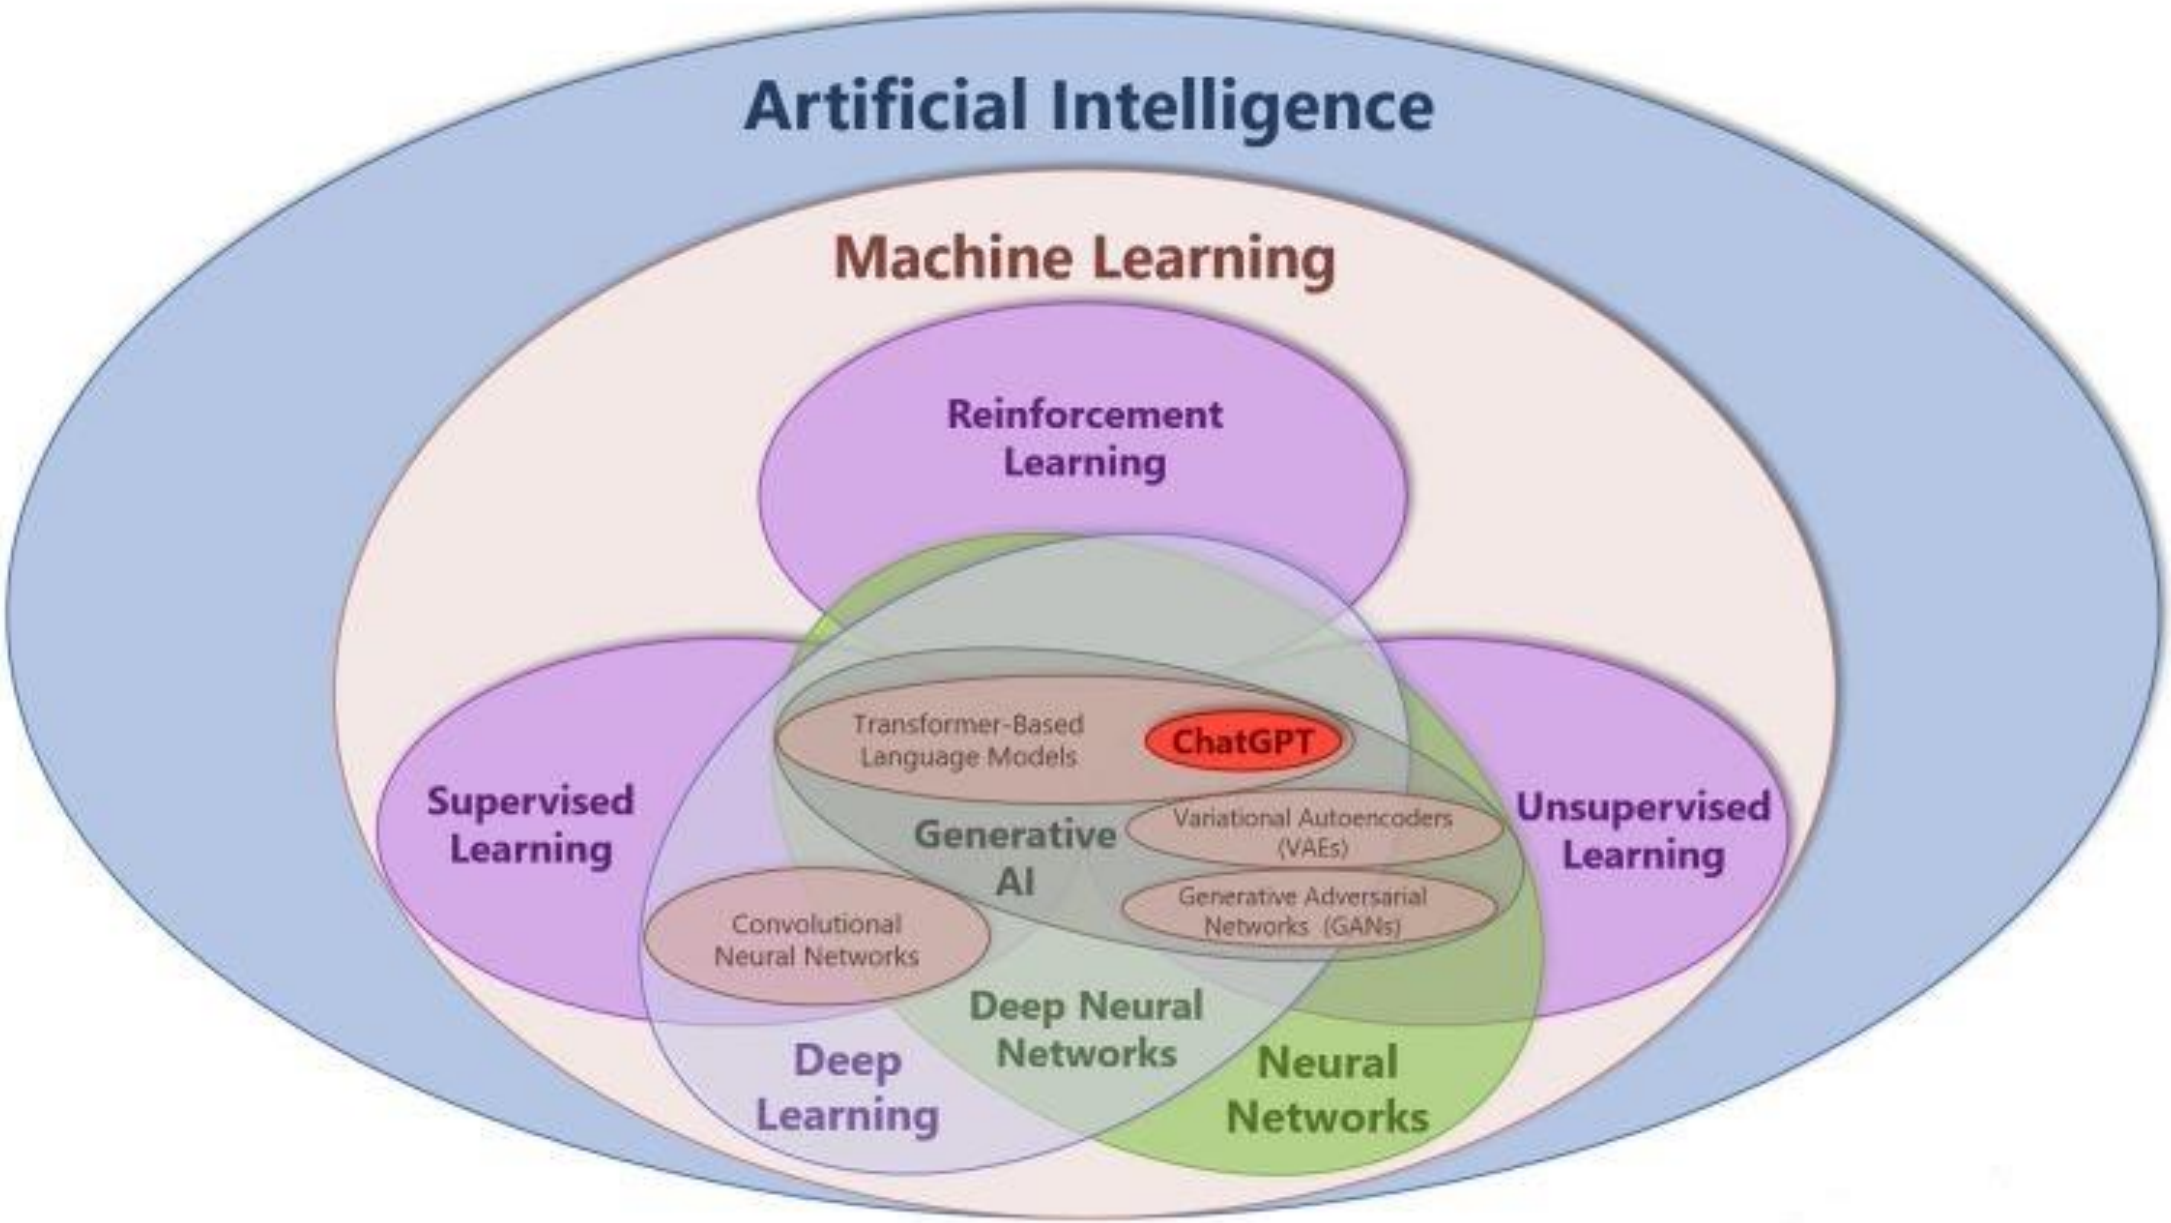
Artificial Intelligence
Machine Learning
Reinforcement Learning
Supervised Learning
Unsupervised Learning
Transformer-Based Language Models
ChatGPT
Generative AI
Variational Autoencoders (VAEs)
Generative Adversarial Networks (GANs)
Convolutional Neural Networks
Deep Neural Networks
Deep Learning
Neural Networks

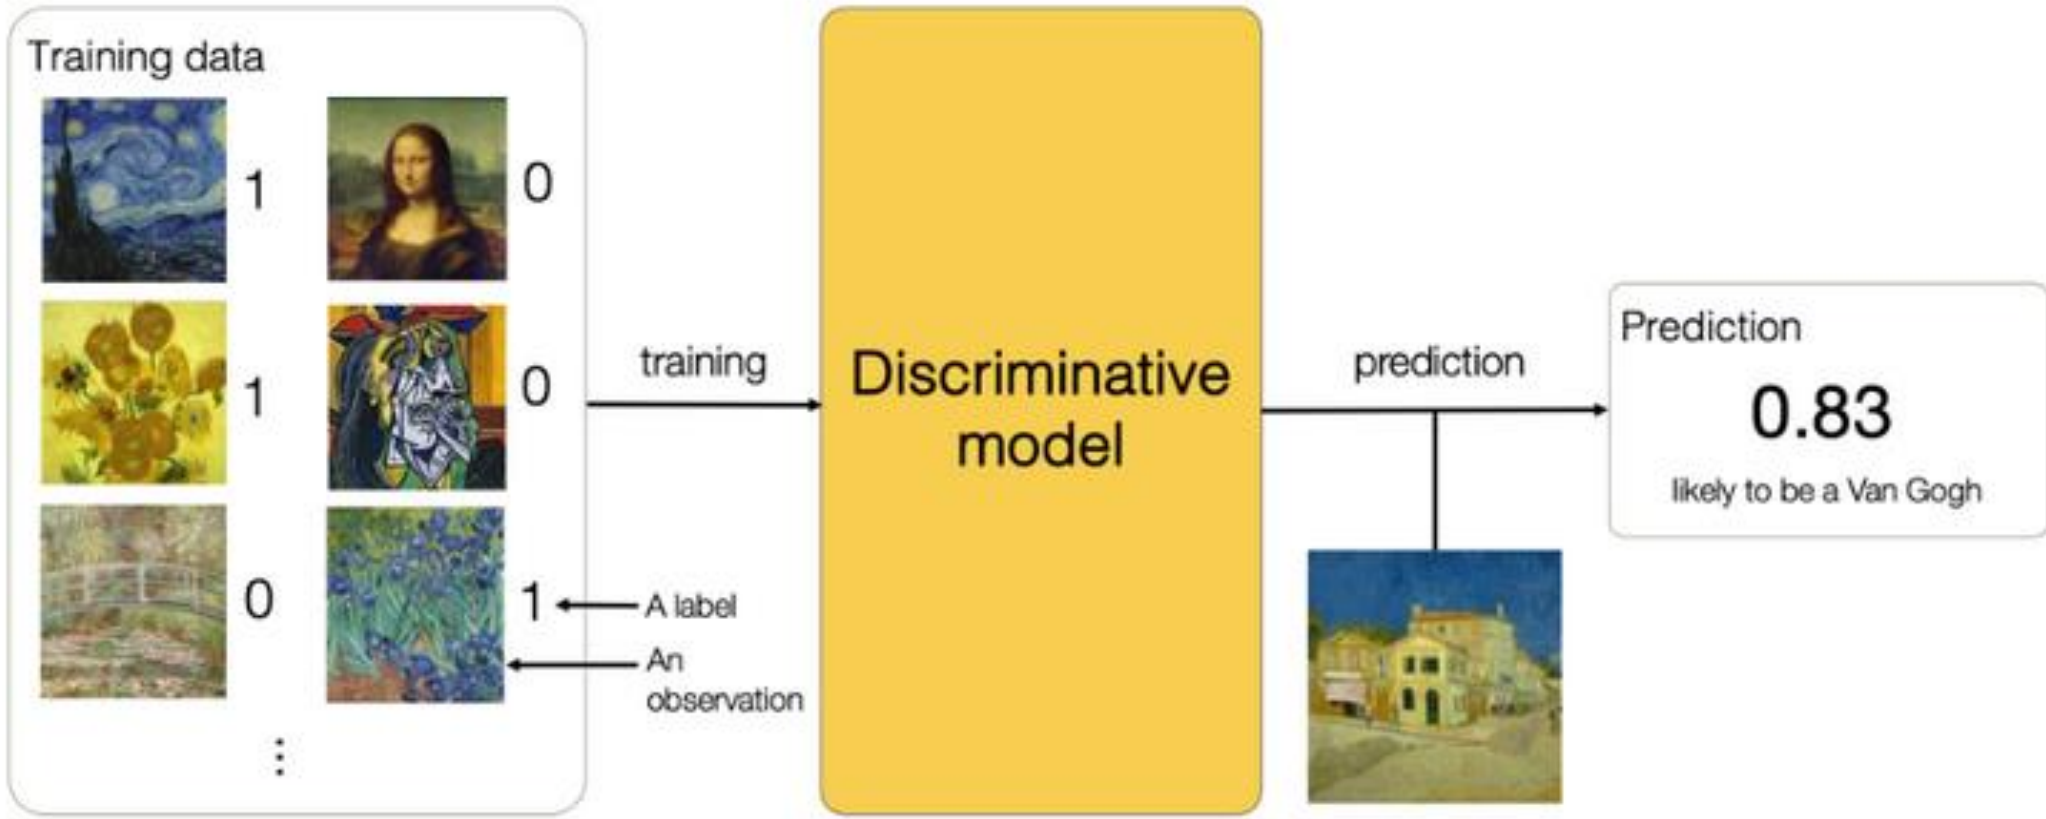
Training data
1
0
1
0
0
1
A label
An observation
training
Discriminative model
prediction
Prediction
0.83
likely to be a Van Gogh

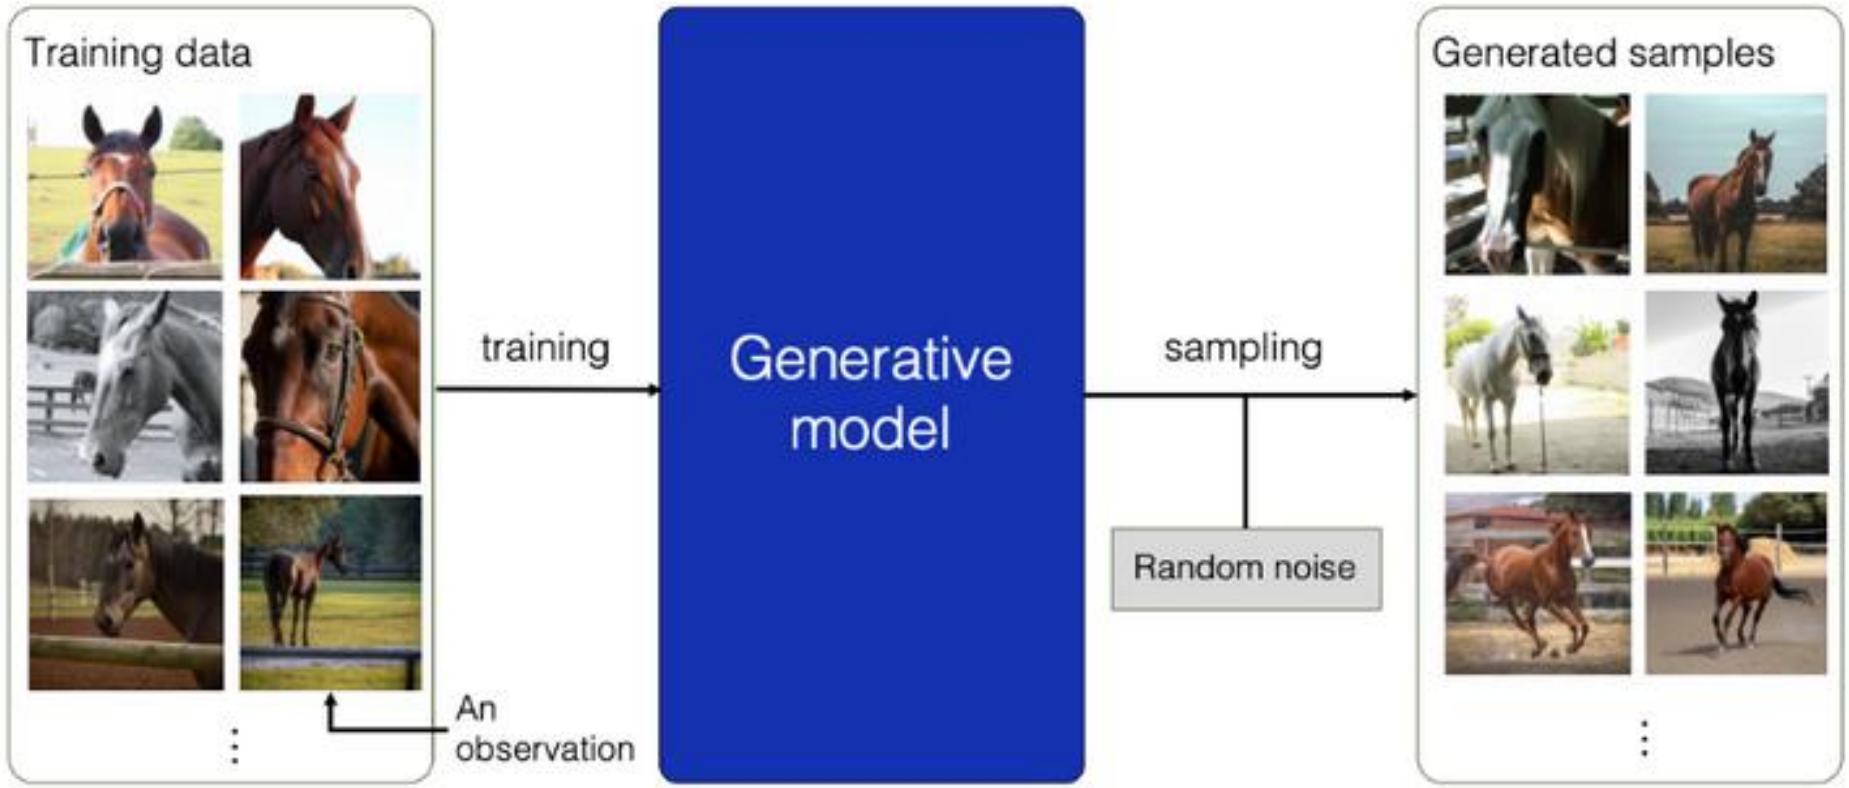
Training data
training
Generative model
sampling
Random noise
Generated samples
An observation

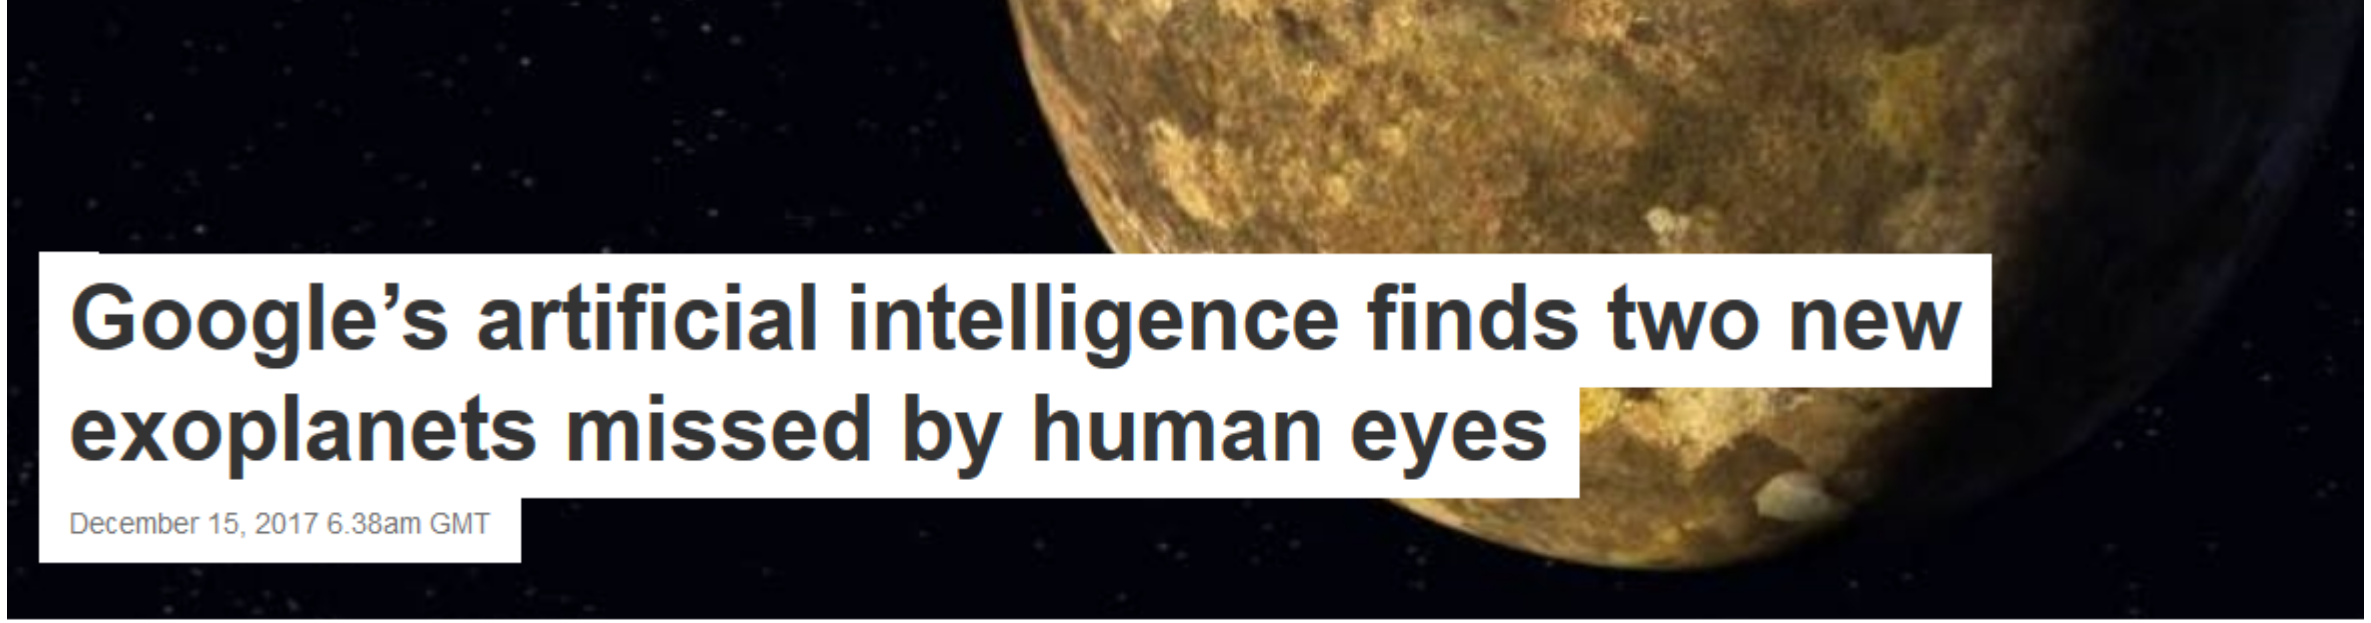

Google's artificial intelligence finds two new exoplanets missed by human eyes

December 15, 2017 6.38am GMT

Artist impression of Kepler-90i, the eighth planet discovered orbiting around Kepler-90. NASA

Two new exoplanets have been discovered thanks to NASA's collaboration with Google's artificial intelligence (AI). One of those in today's announcement is an eighth planet – Kepler-90i – found orbiting the Sun-like star Kepler-90. This makes it the first system discovered with an equal number of planets to our own Solar system.

A mere road trip away, at 2,545 light-years from Earth, Kepler-90i orbits its host star every 14.4 Earth days, with a sizzling surface temperature similar to Venus of 426°C.

The new exoplanets are added to the growing list of known worlds found orbiting other stars.

Health & Medicine

AI Can Diagnose Heart Disease and Lung Cancer More Accurately Than Doctors

IN BRIEF

A pair of recently developed AI systems can diagnose lung cancer and heart disease more accurately than human doctors. These AIs have the potential to save billions of dollars and countless lives if widely adopted.

SHARE

WRITTEN BY

Kyree Leary


← Pay at —
Muni-Meter

Examples of images generated by training a generative adversarial network (GAN) with portraits from the last 500 years of Western art. The distorted faces are the algorithm's attempts to imitate those inputs. Images generated at Art & Artificial Intelligence Laboratory, Rutgers.

thispersondoesnotexist.com

Generative AI Applications

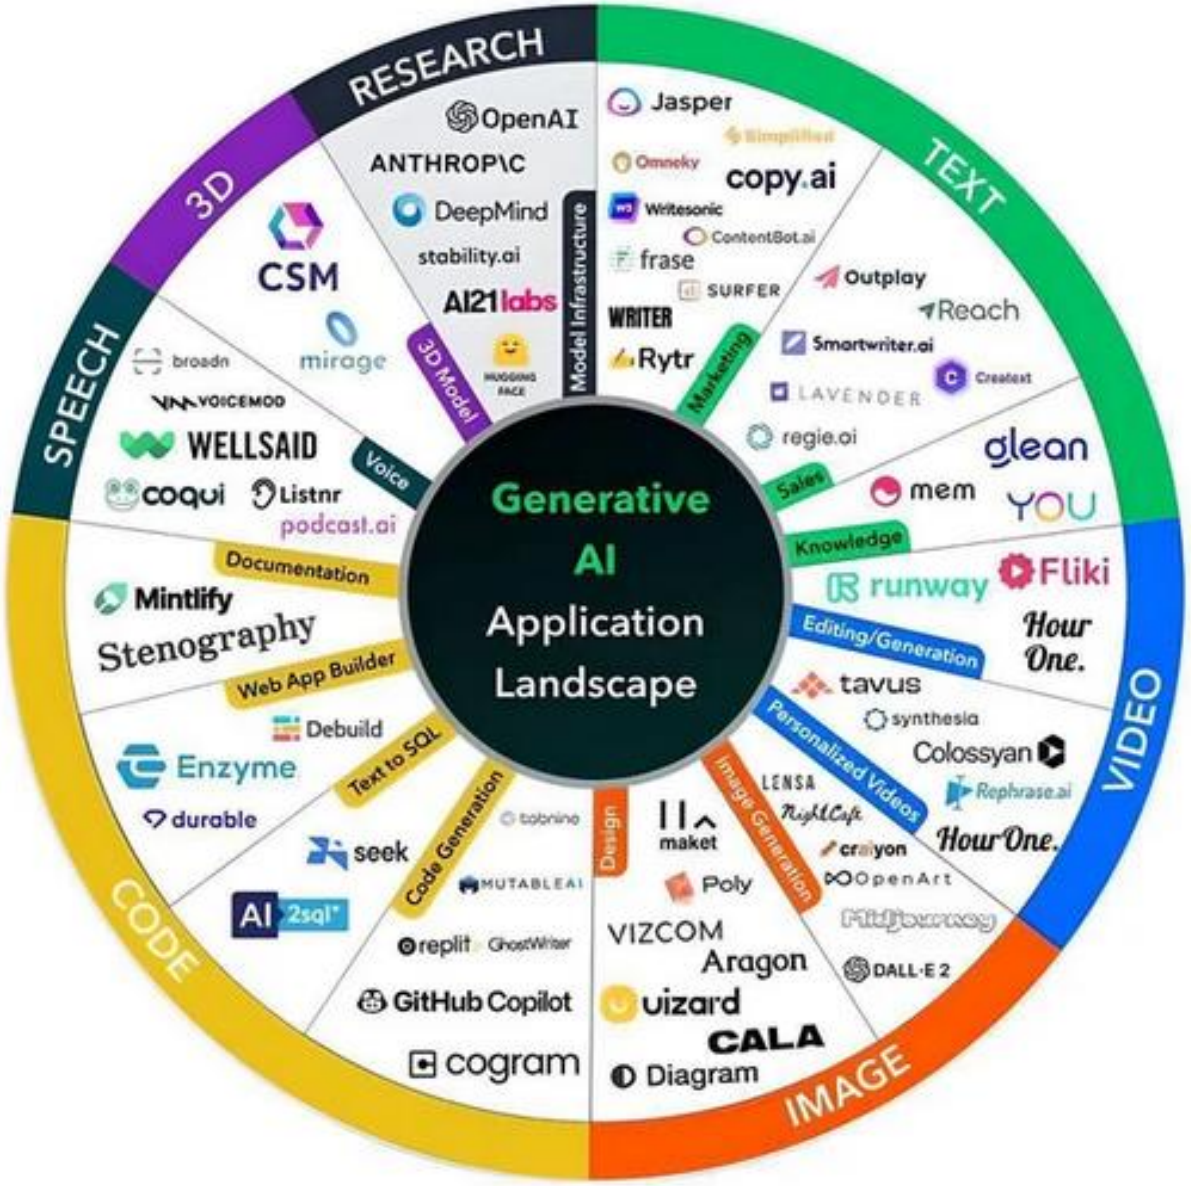

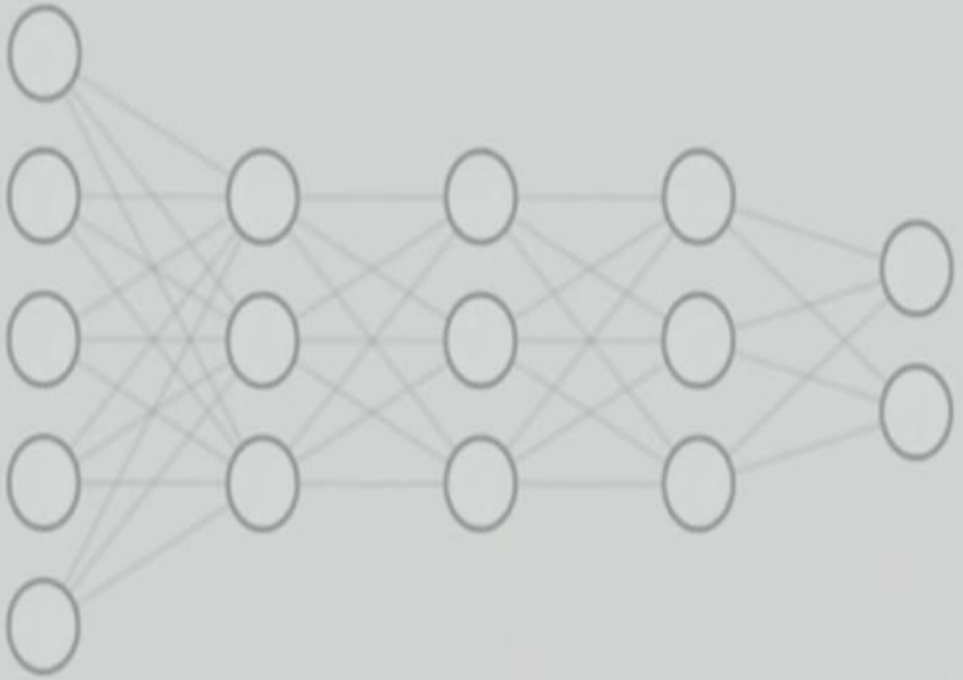

BIRD

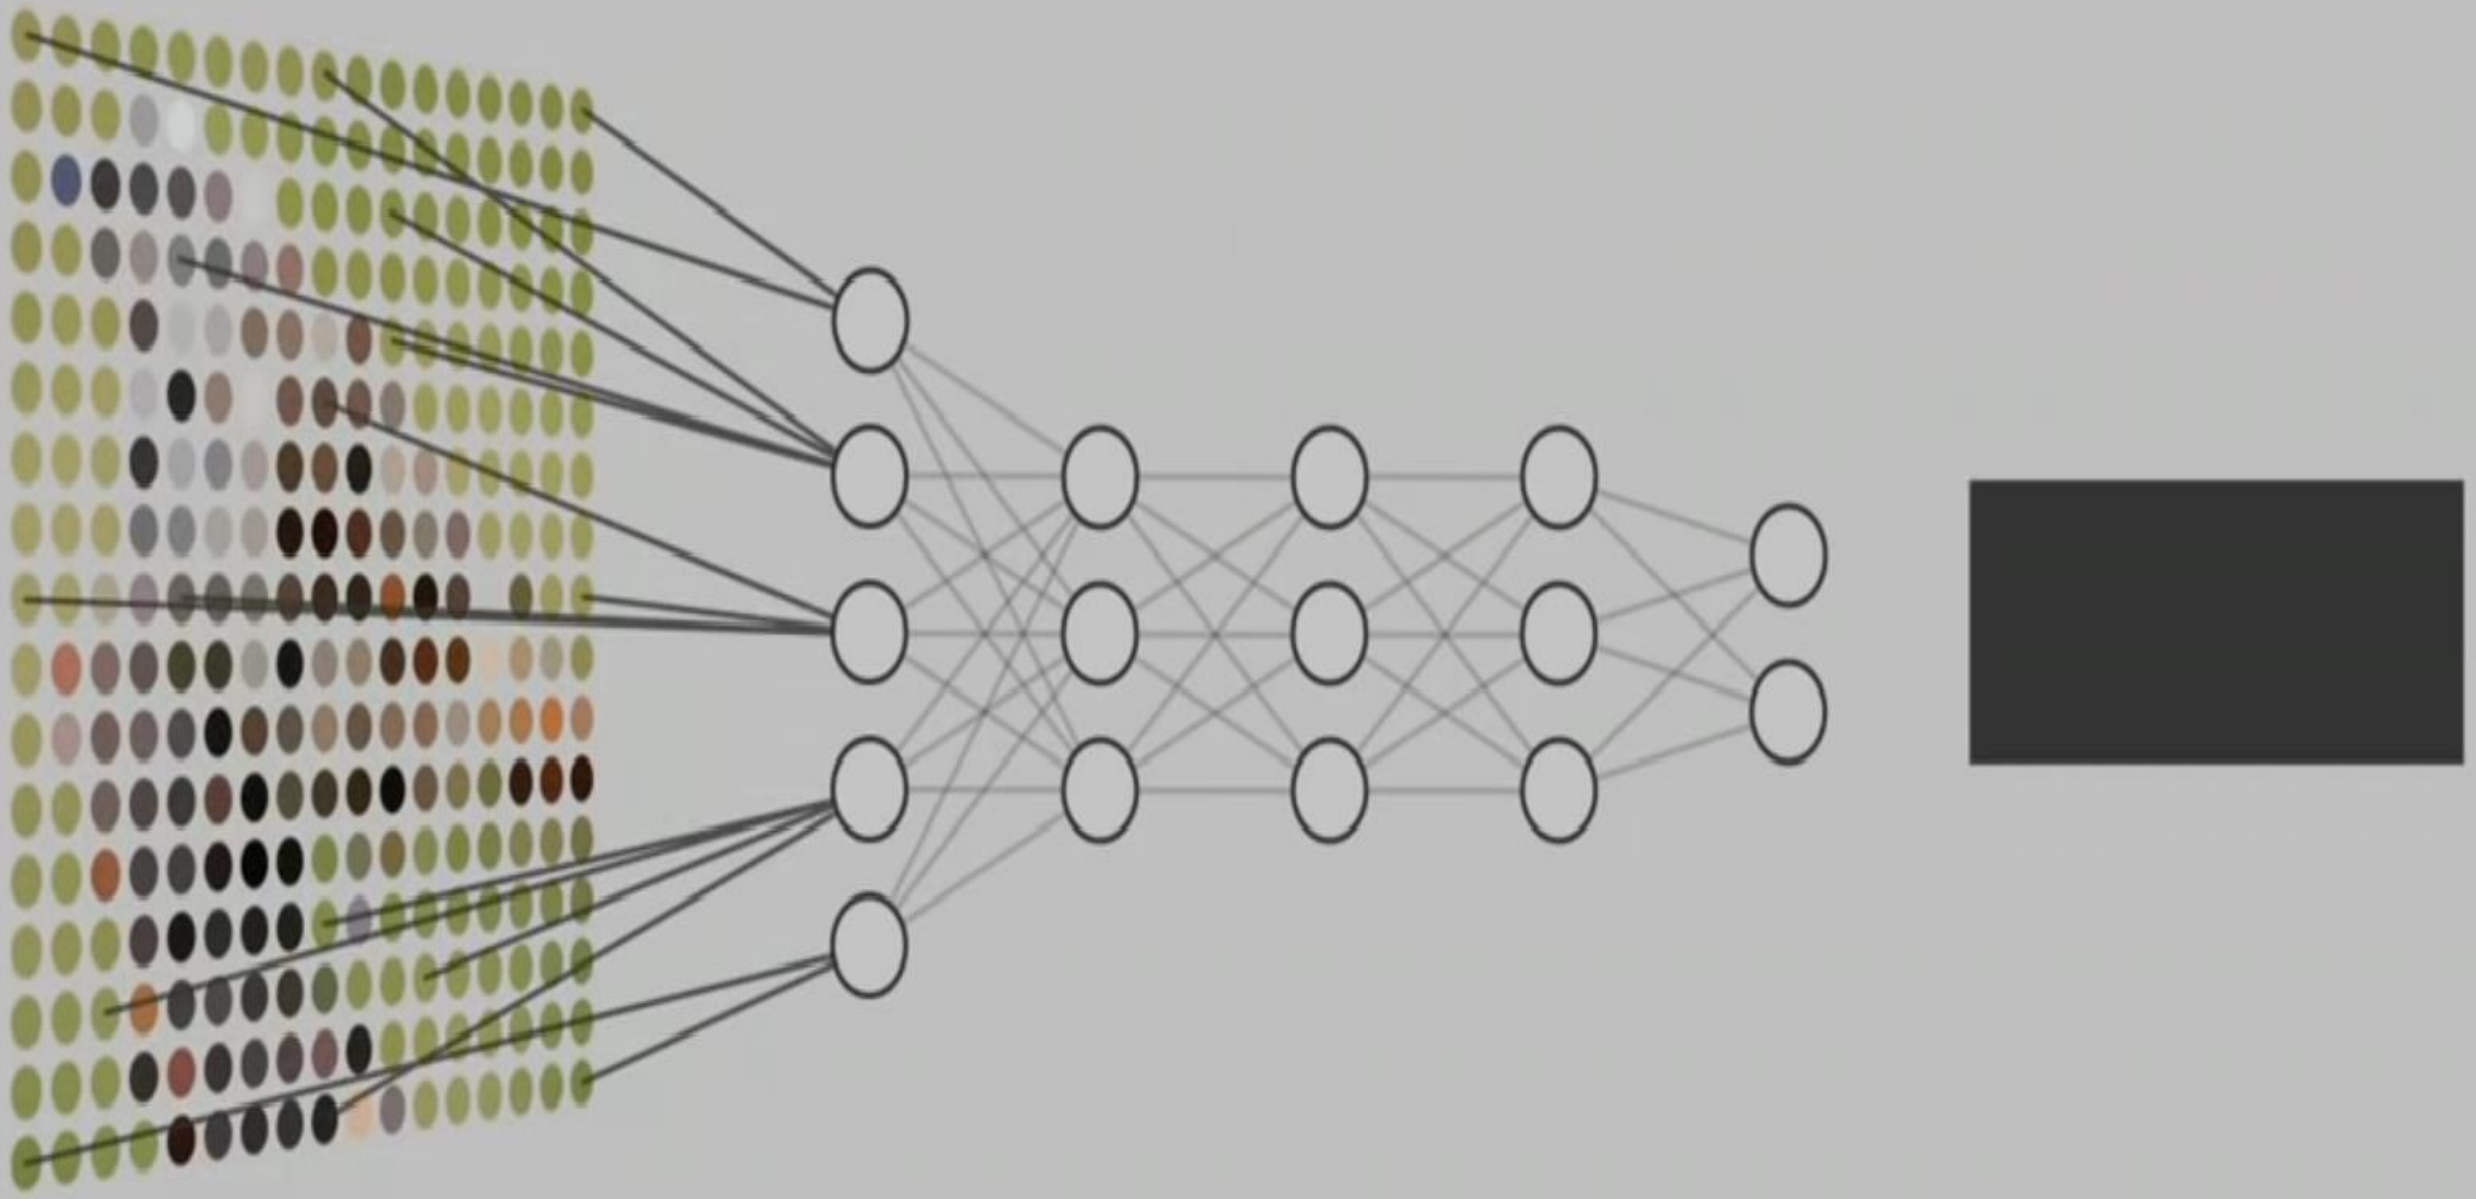

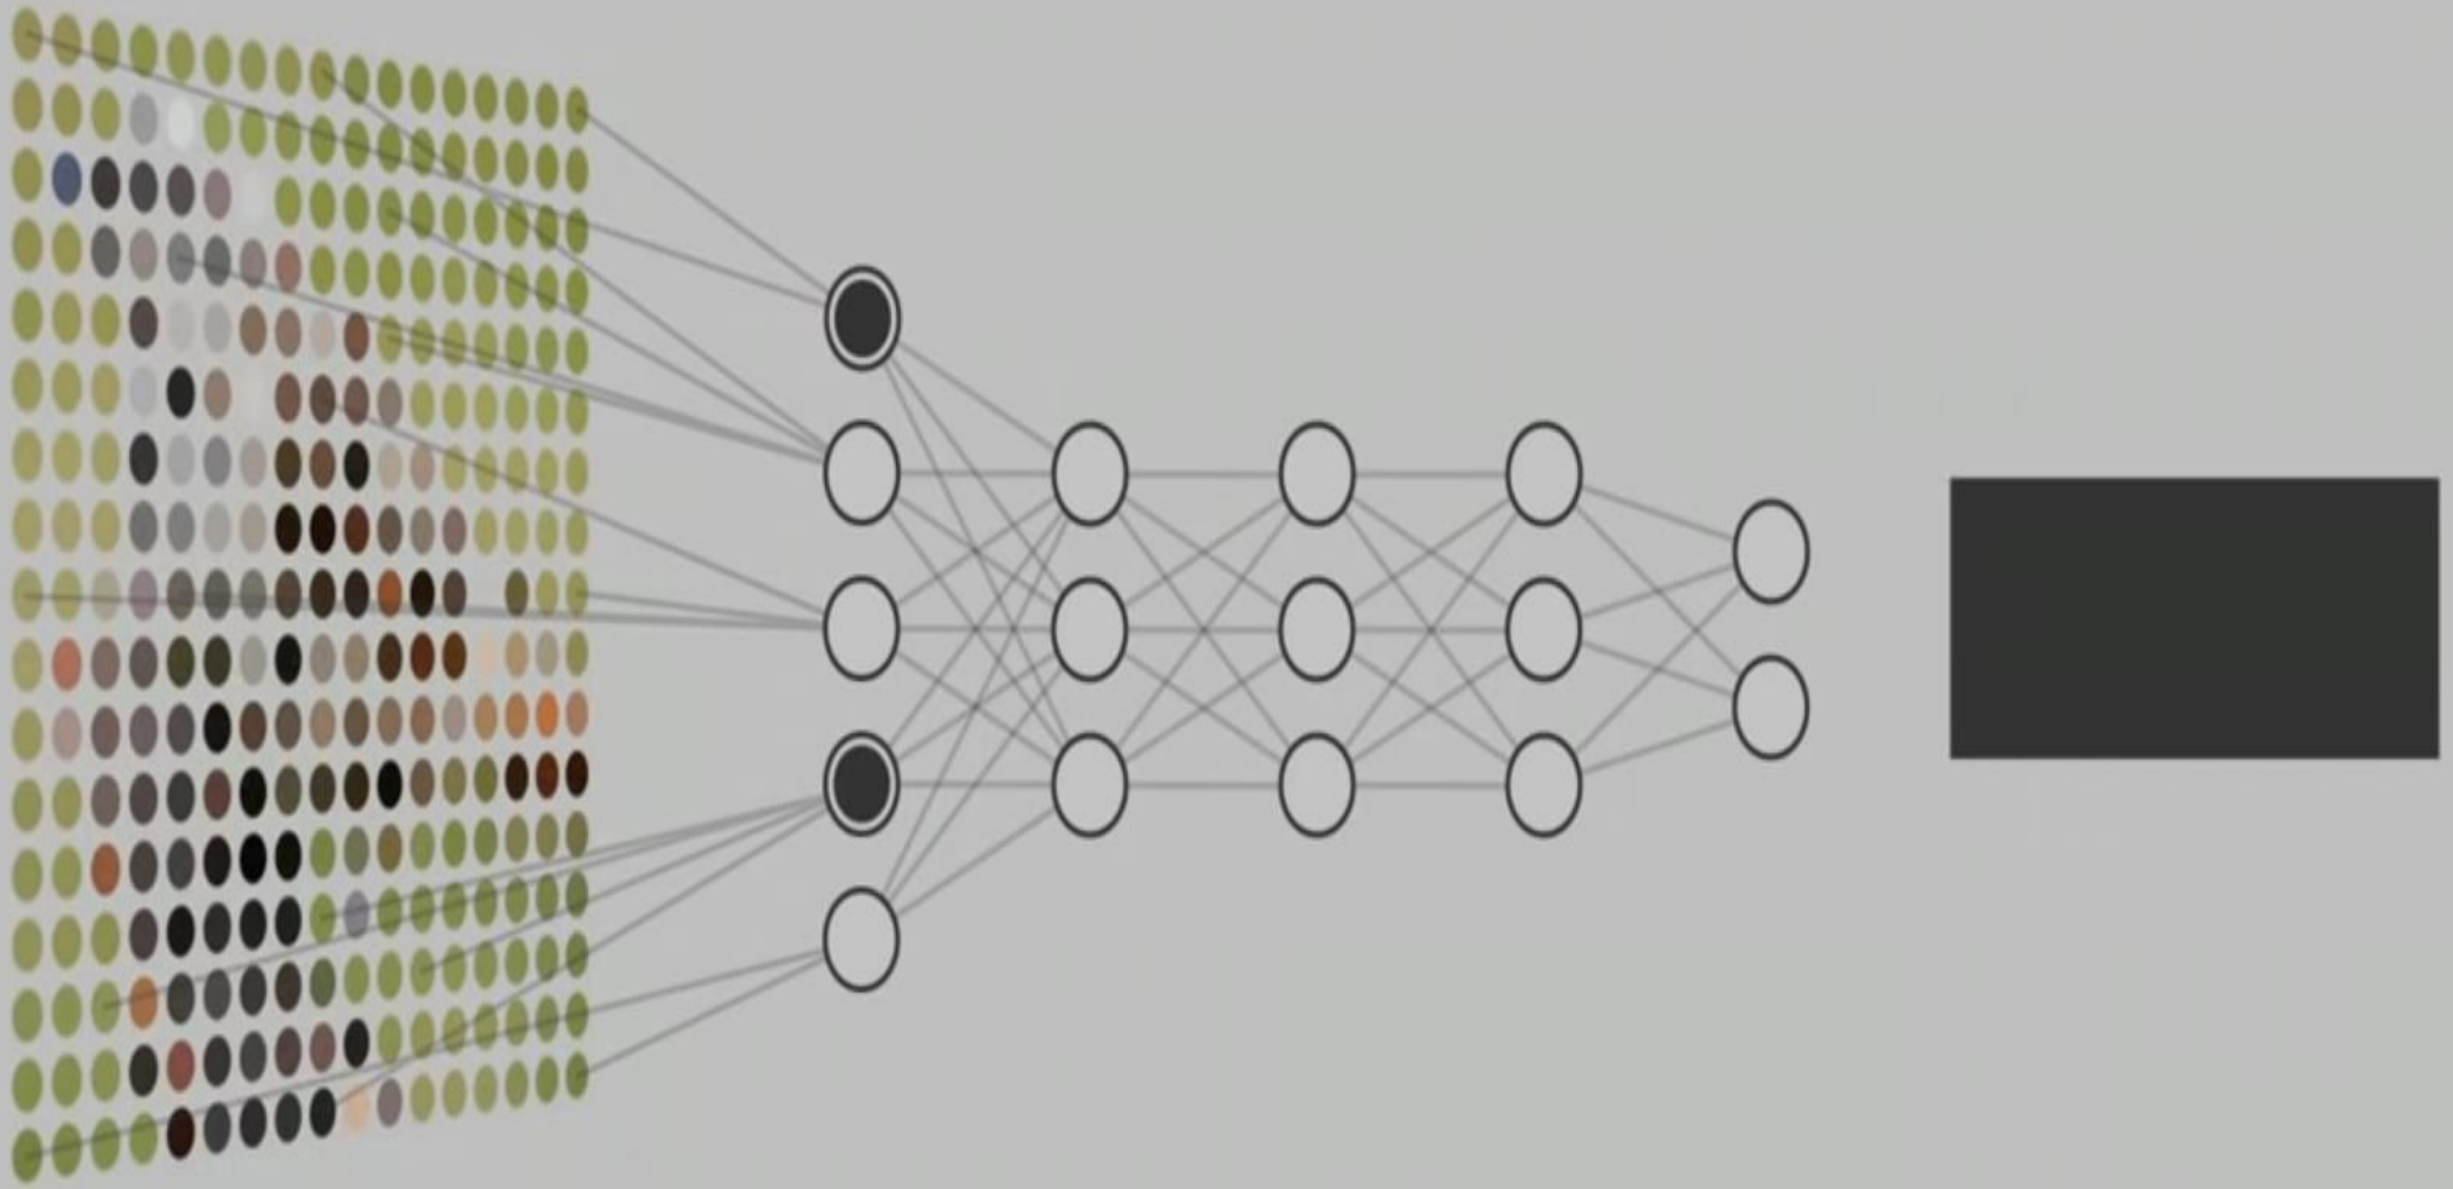

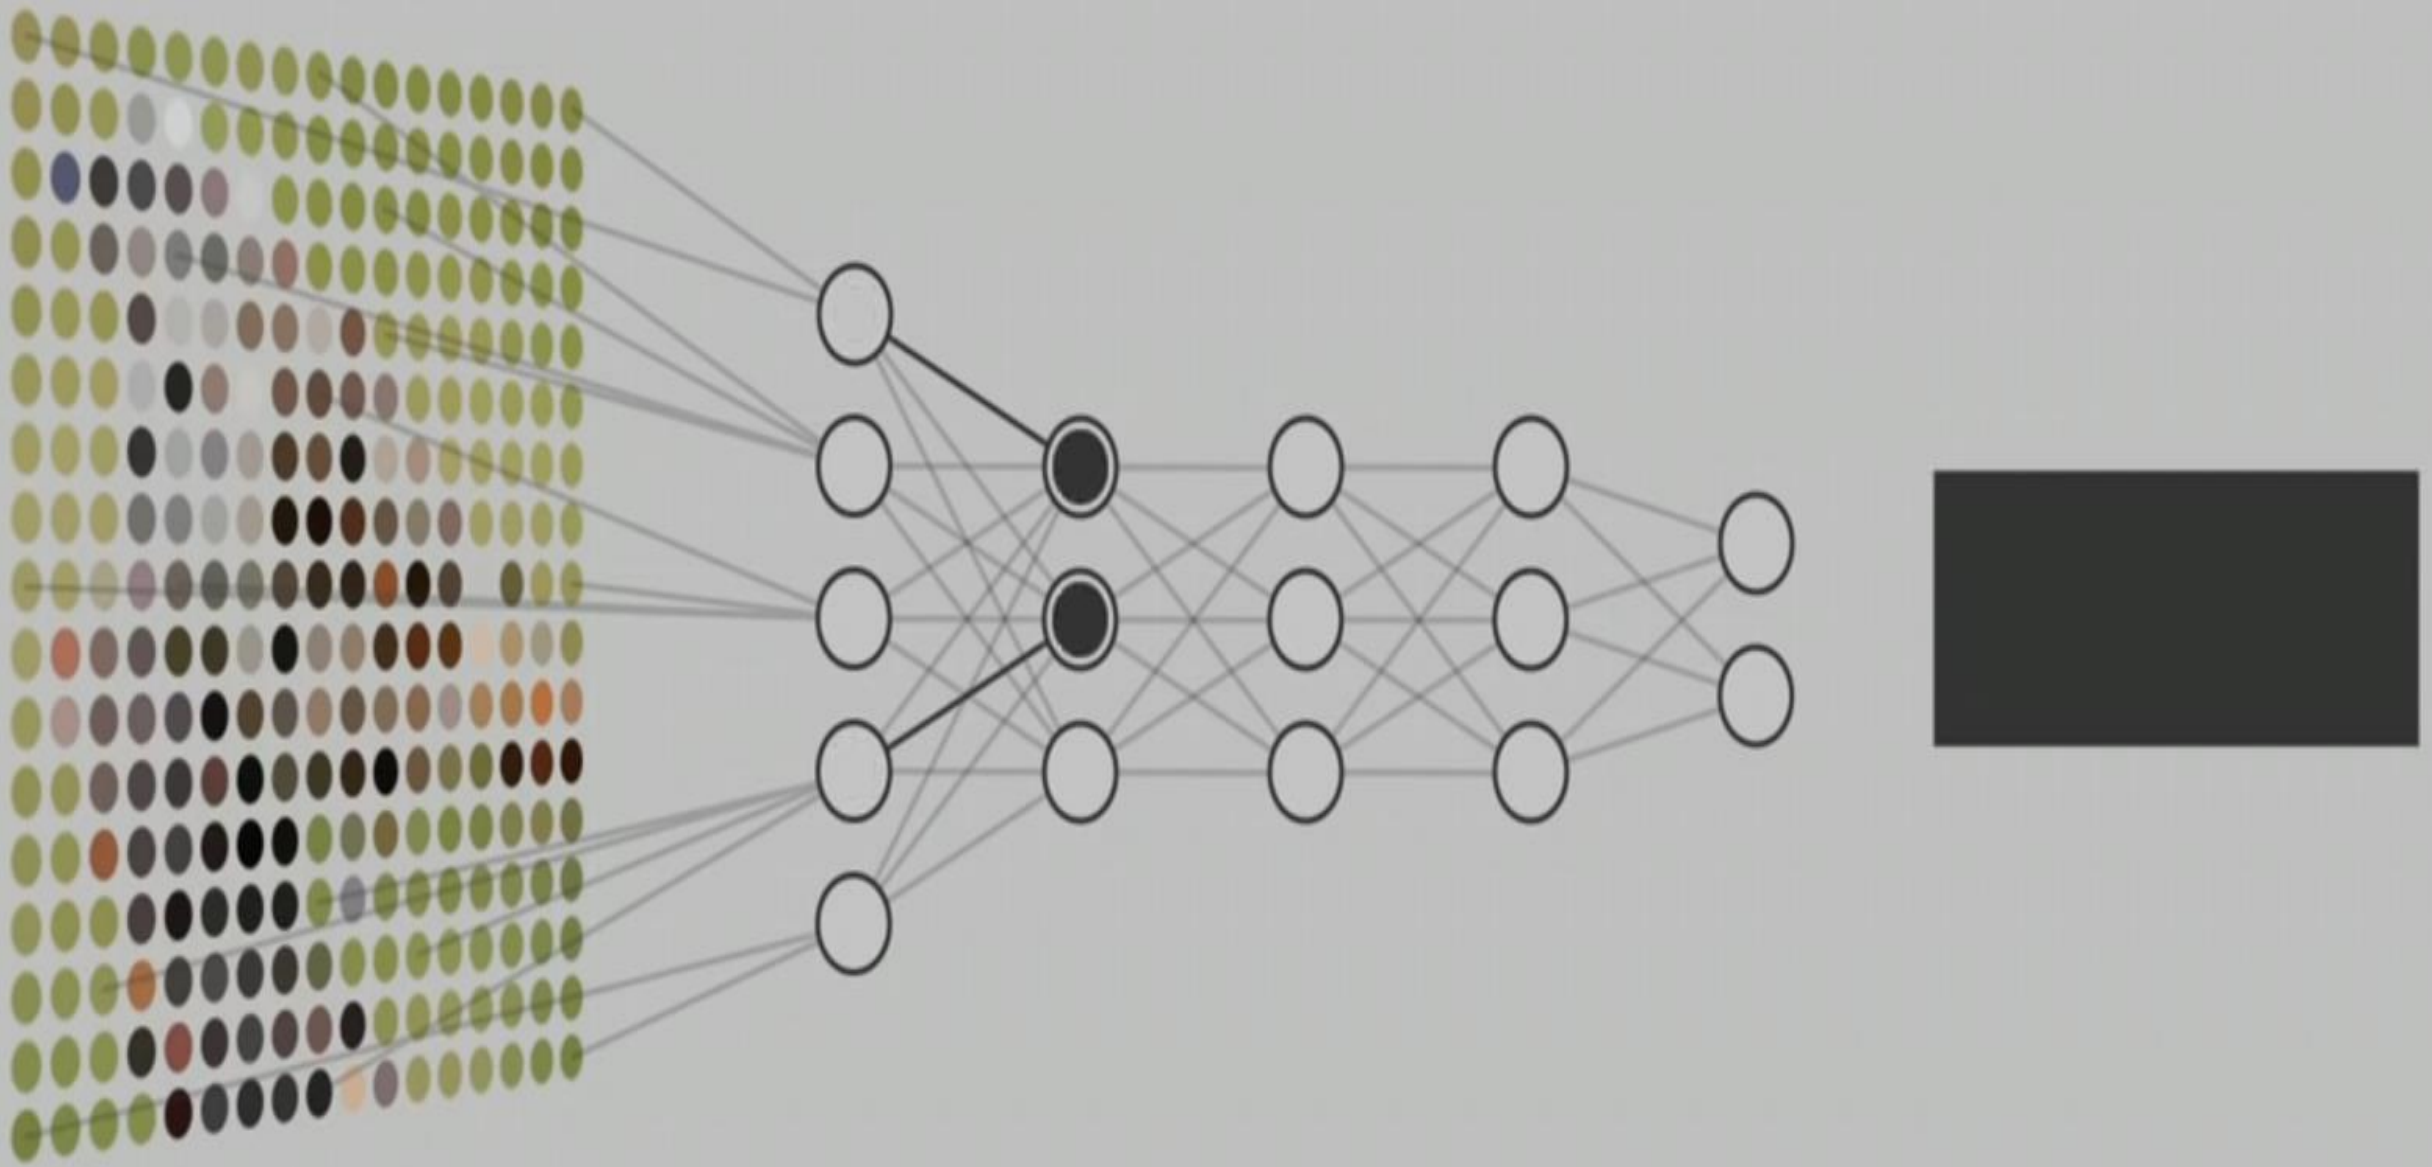

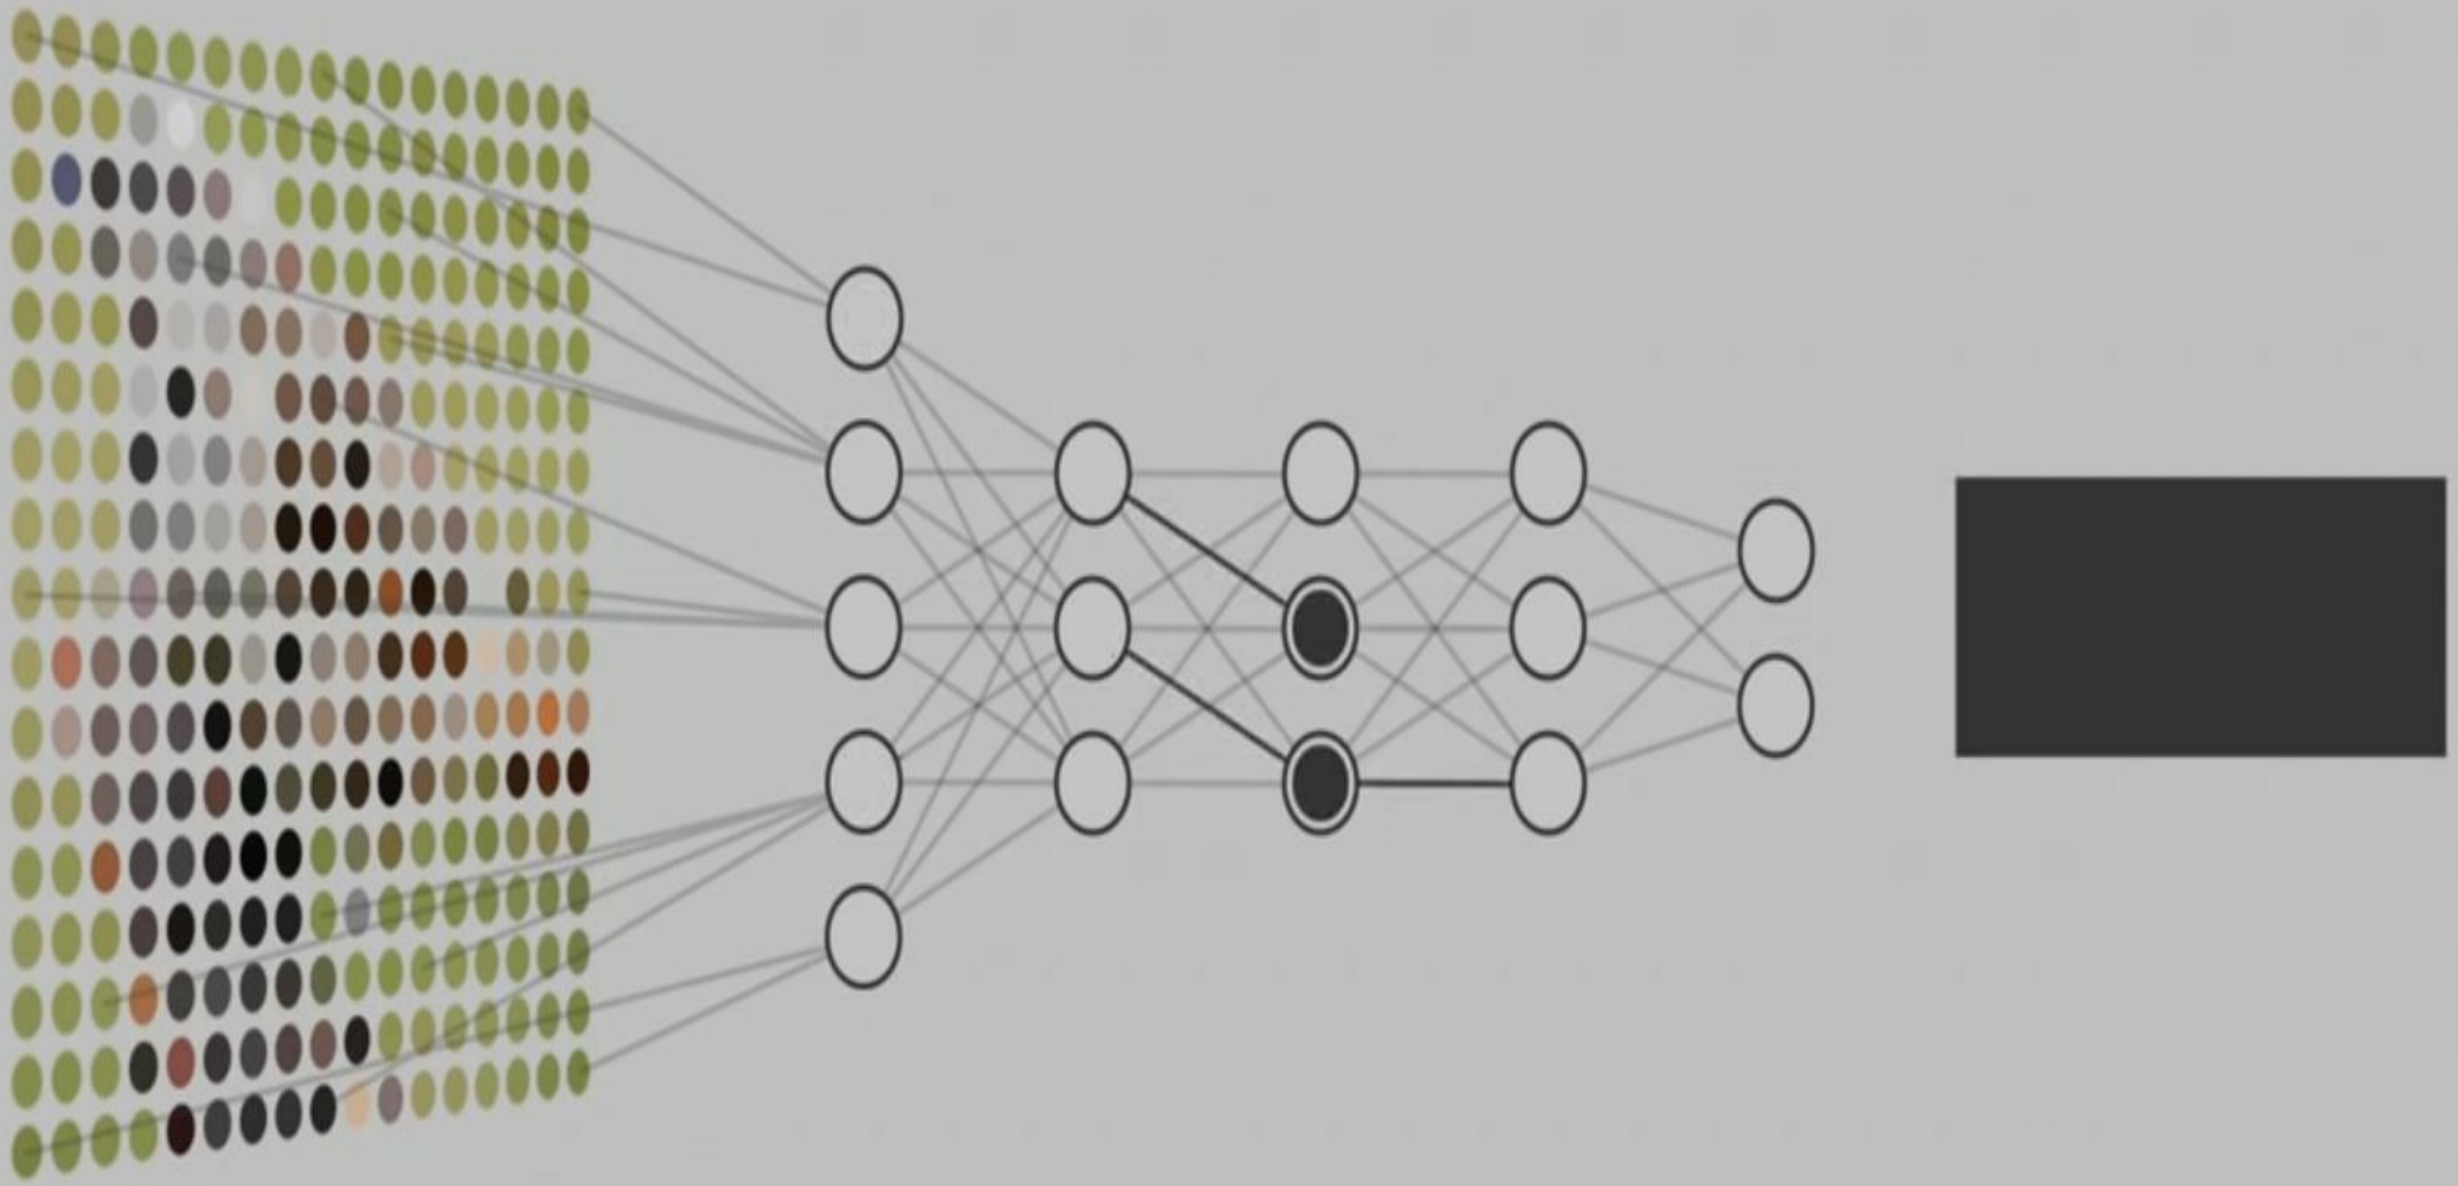

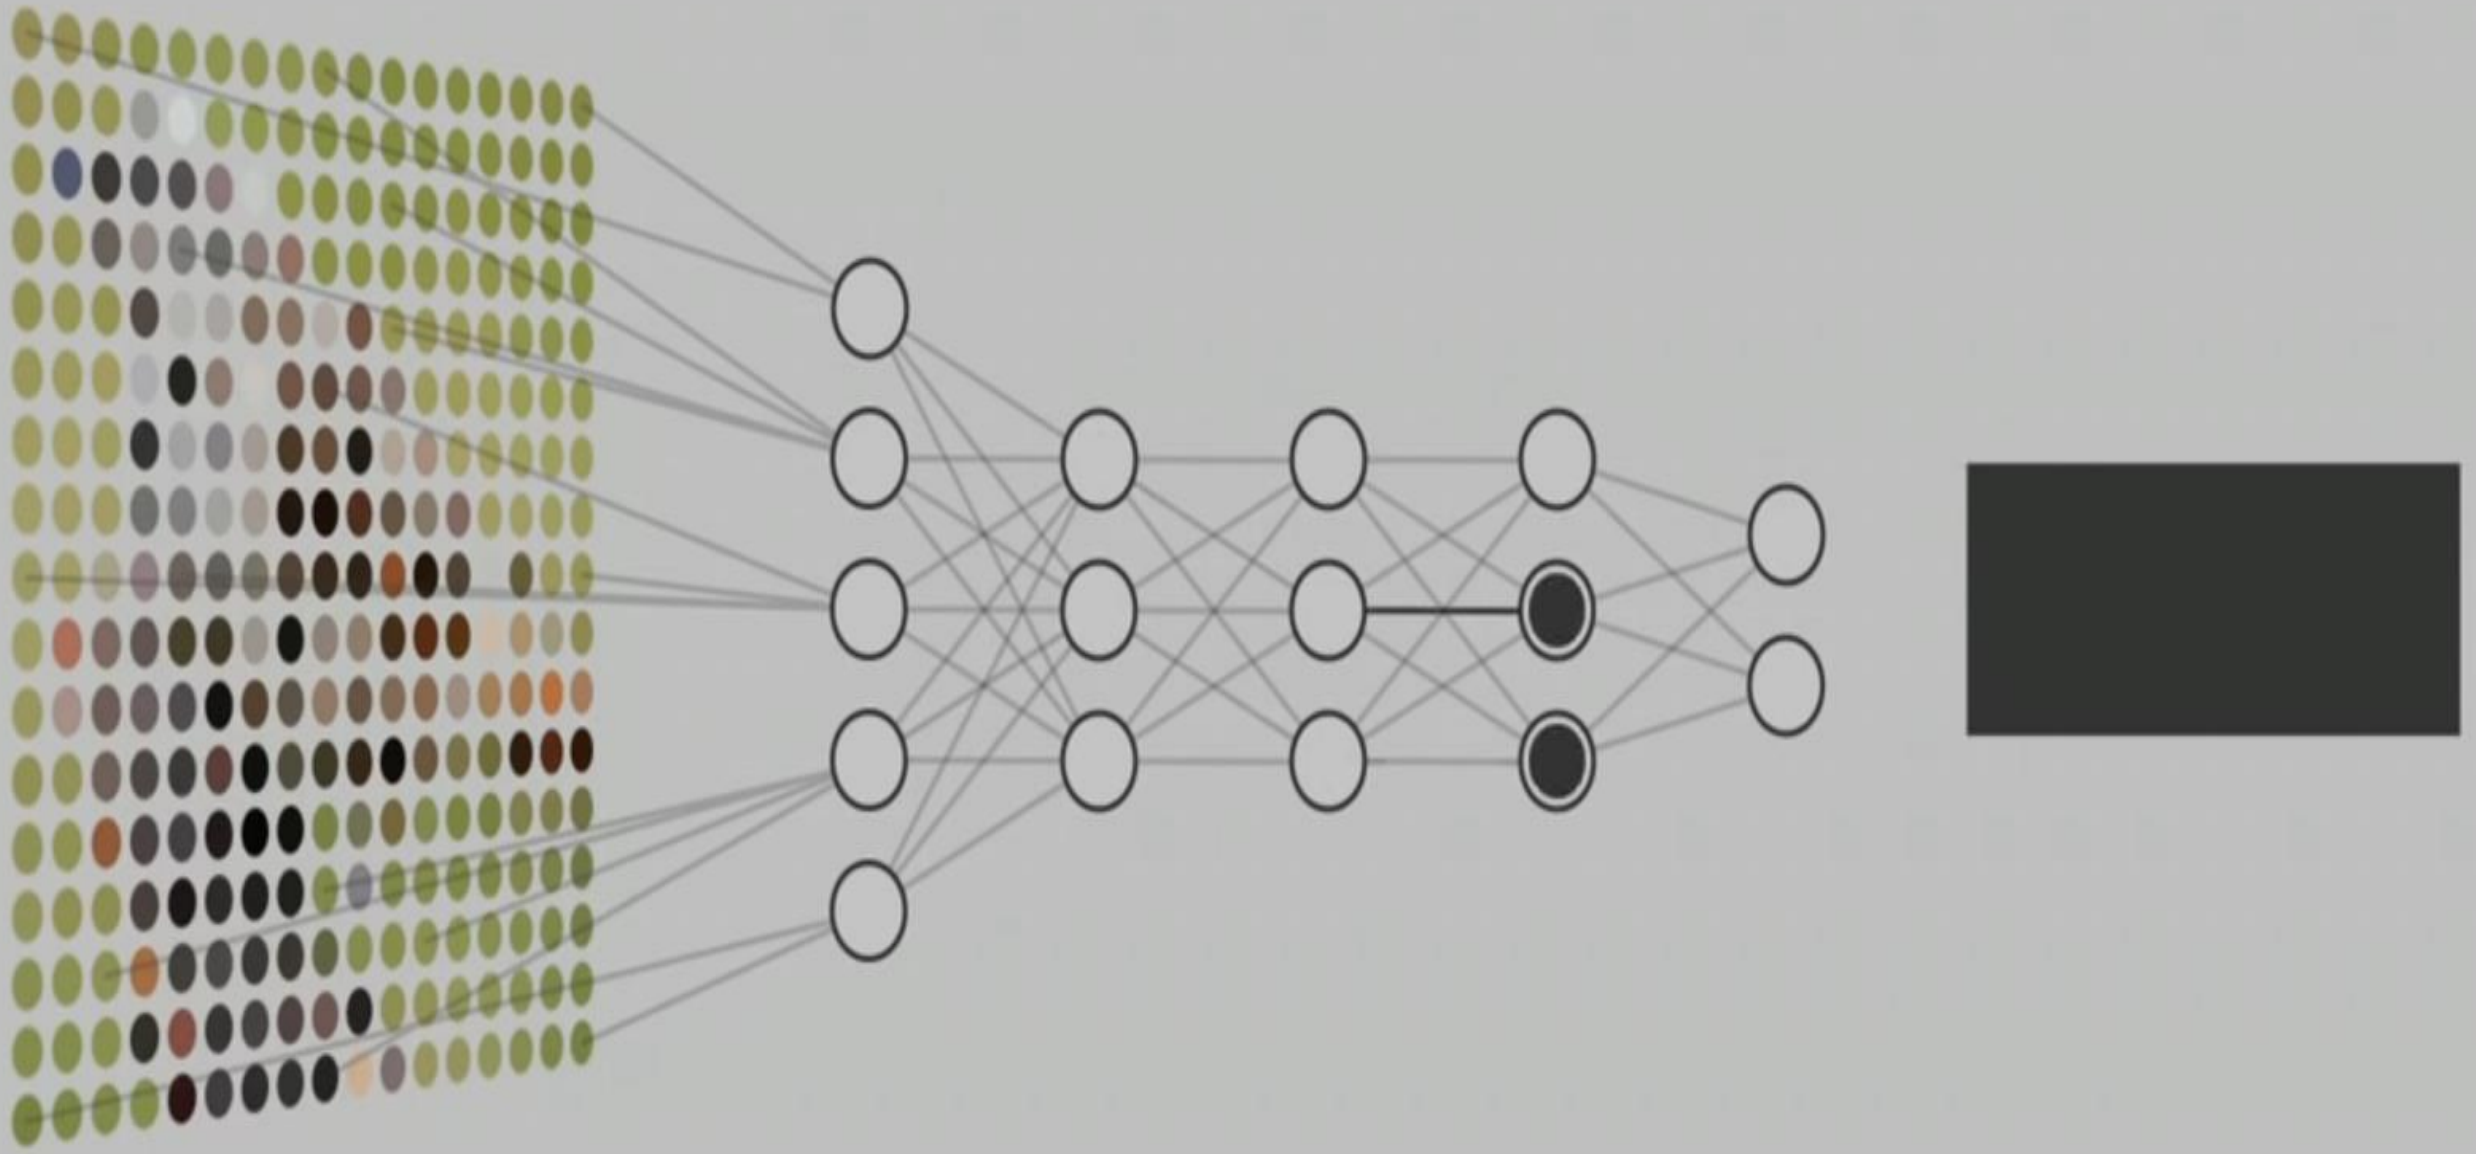

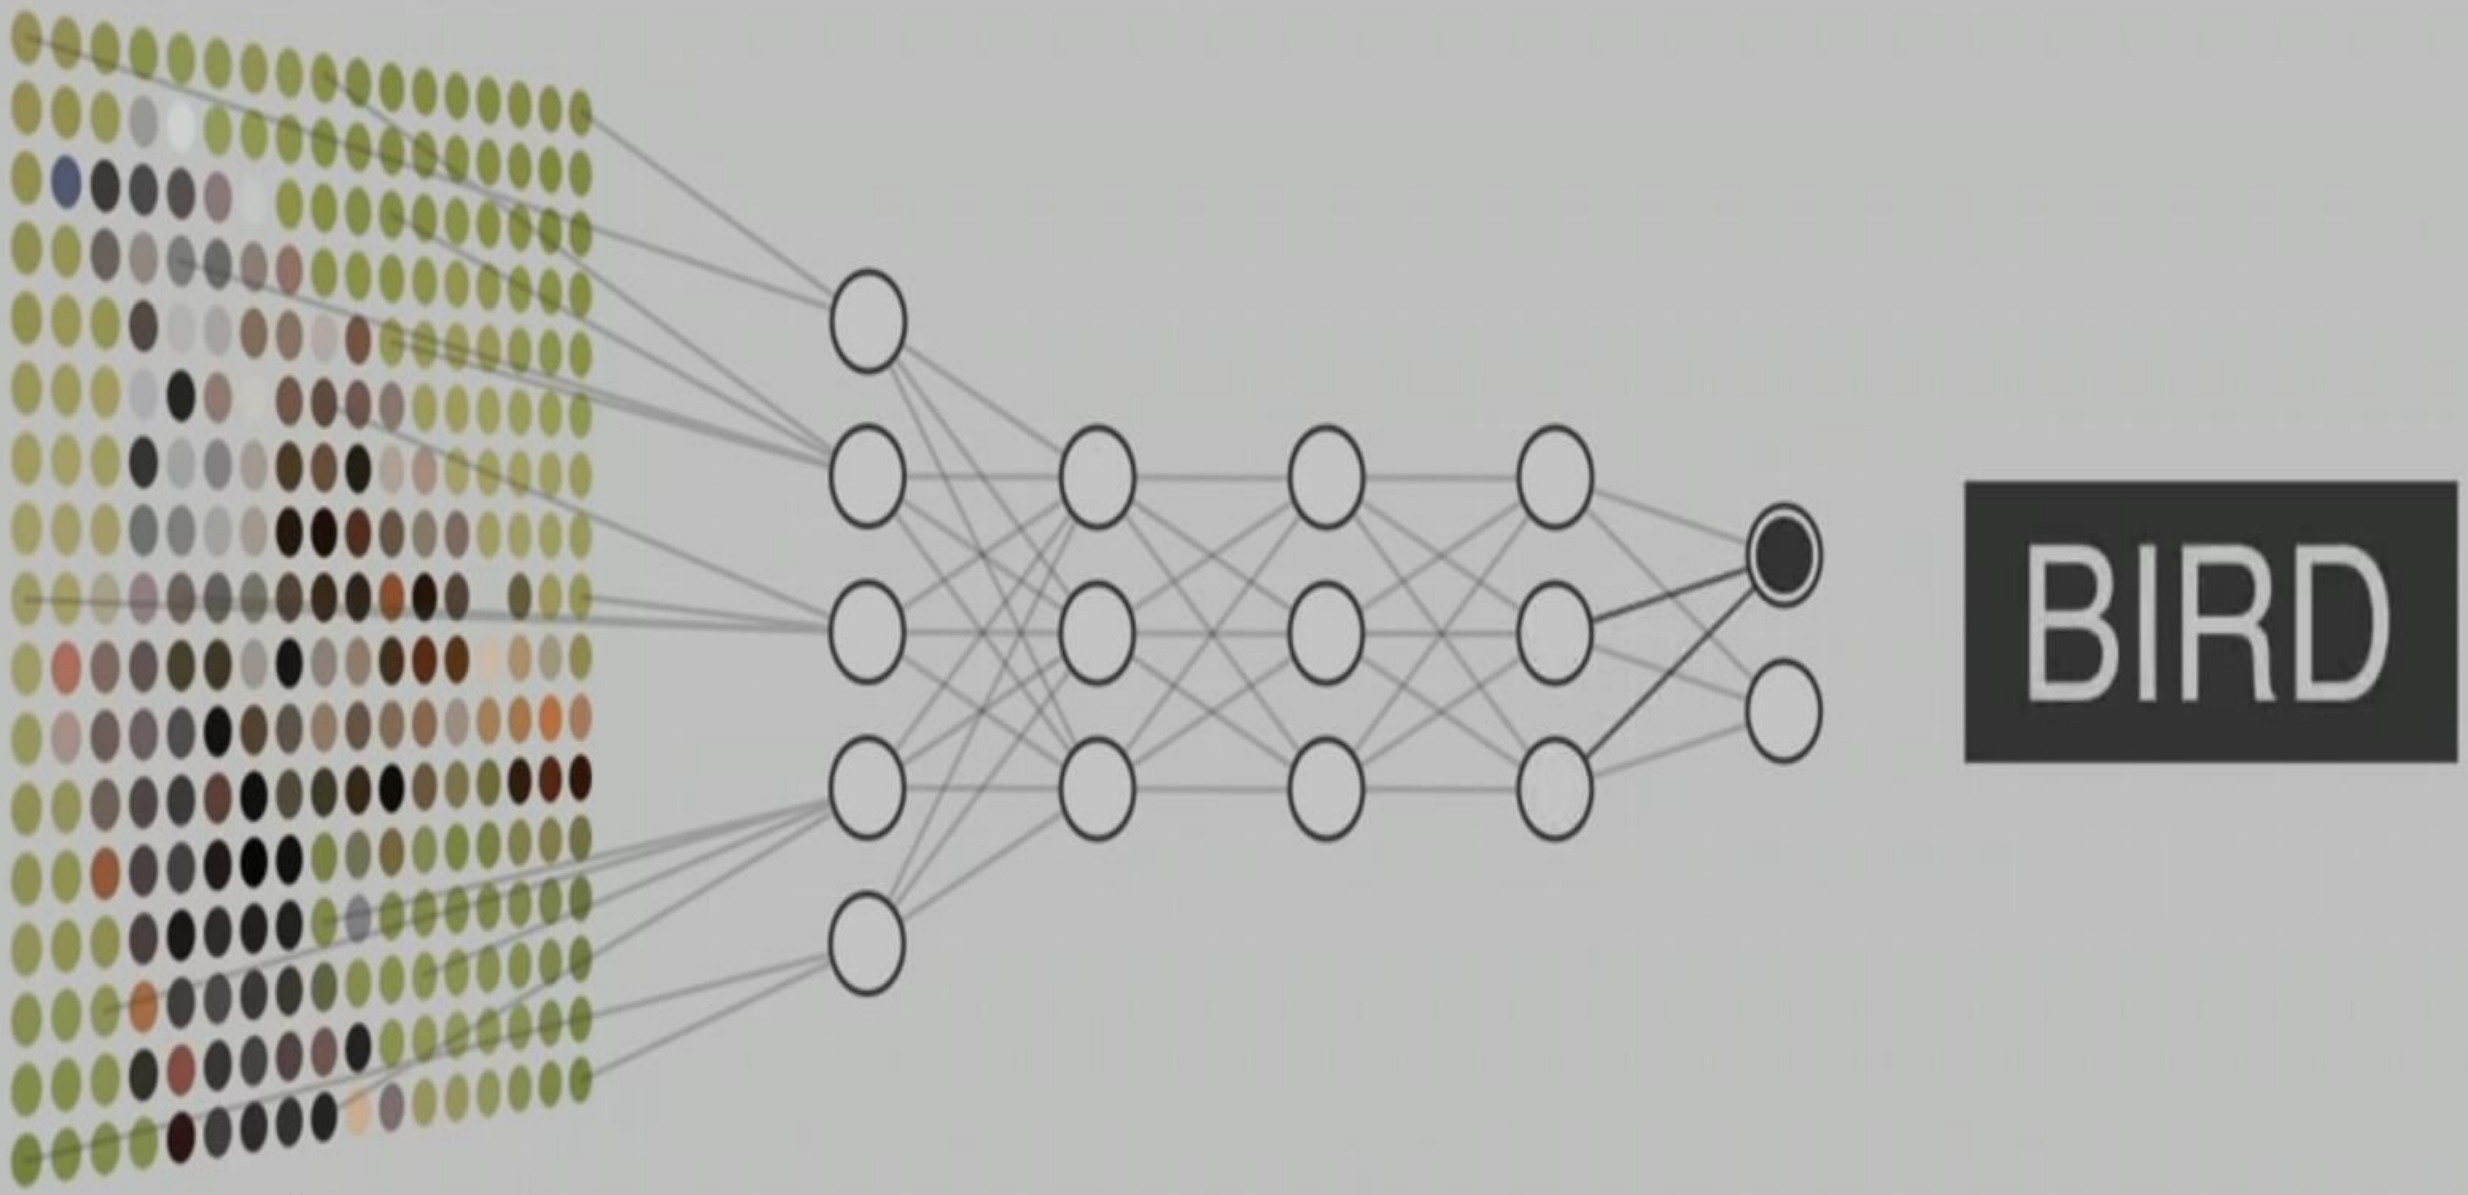
BIRD

x_1
x_2
x_3
x_4
x_5
x_6
x_7
x_8
x_9
x_{10}
x_{11}
x_{12}
x_{13}
$\vdots$

w_1
w_2
w_3
w_4
w_5
w_6
w_7
w_8
w_9
w_{10}
w_{11}
w_{12}
w_{13}
$\vdots$

y_1
y_2
y_3
y_4


x “×” W = y

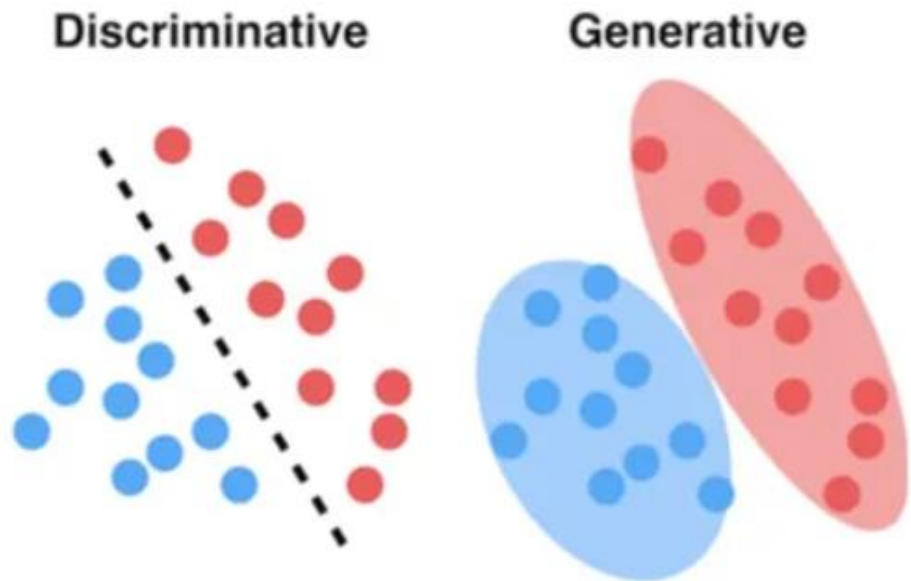

- Discriminative Model

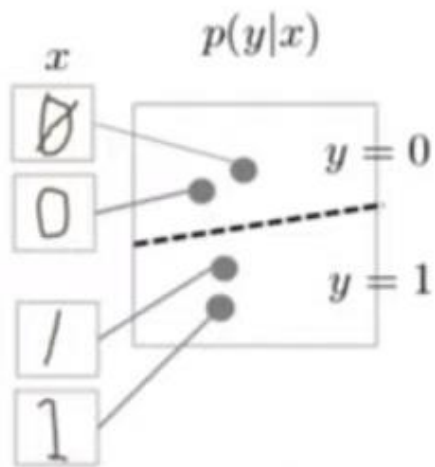

- Generative Model

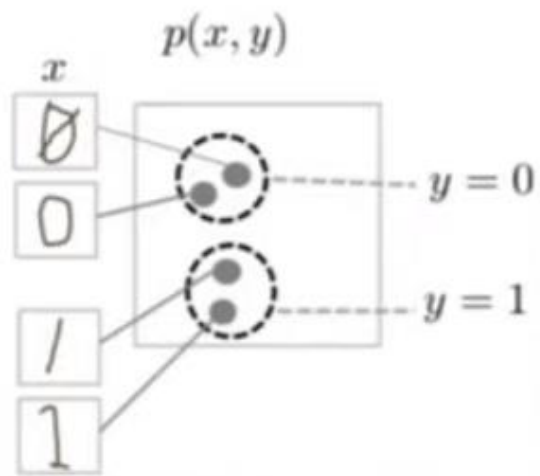

Veri kümesi = { (1,0), (1,0), (2,0), (2,1) }		
Ayırt edici model p(y|x)'i öğrenir.		
x verildiğinde y'nin olasılığı (probability of y given x)		
	y=0	y=1
x=1	1	0
x=2	½	½

Ayırt edici modeller şartlı olasılığı hesaplar

Veri kümesi = { (1,0), (1,0), (2,0), (2,1) }		
Üretici model p(x,y)'i öğrenir.		
x ve y'nin birlikte olasılığı (probability of x and y)		
	y=0	y=1
x=1	½	0
x=2	¼	¼

Üretici modeller birlikte gelme olasılıklarını hesaplar

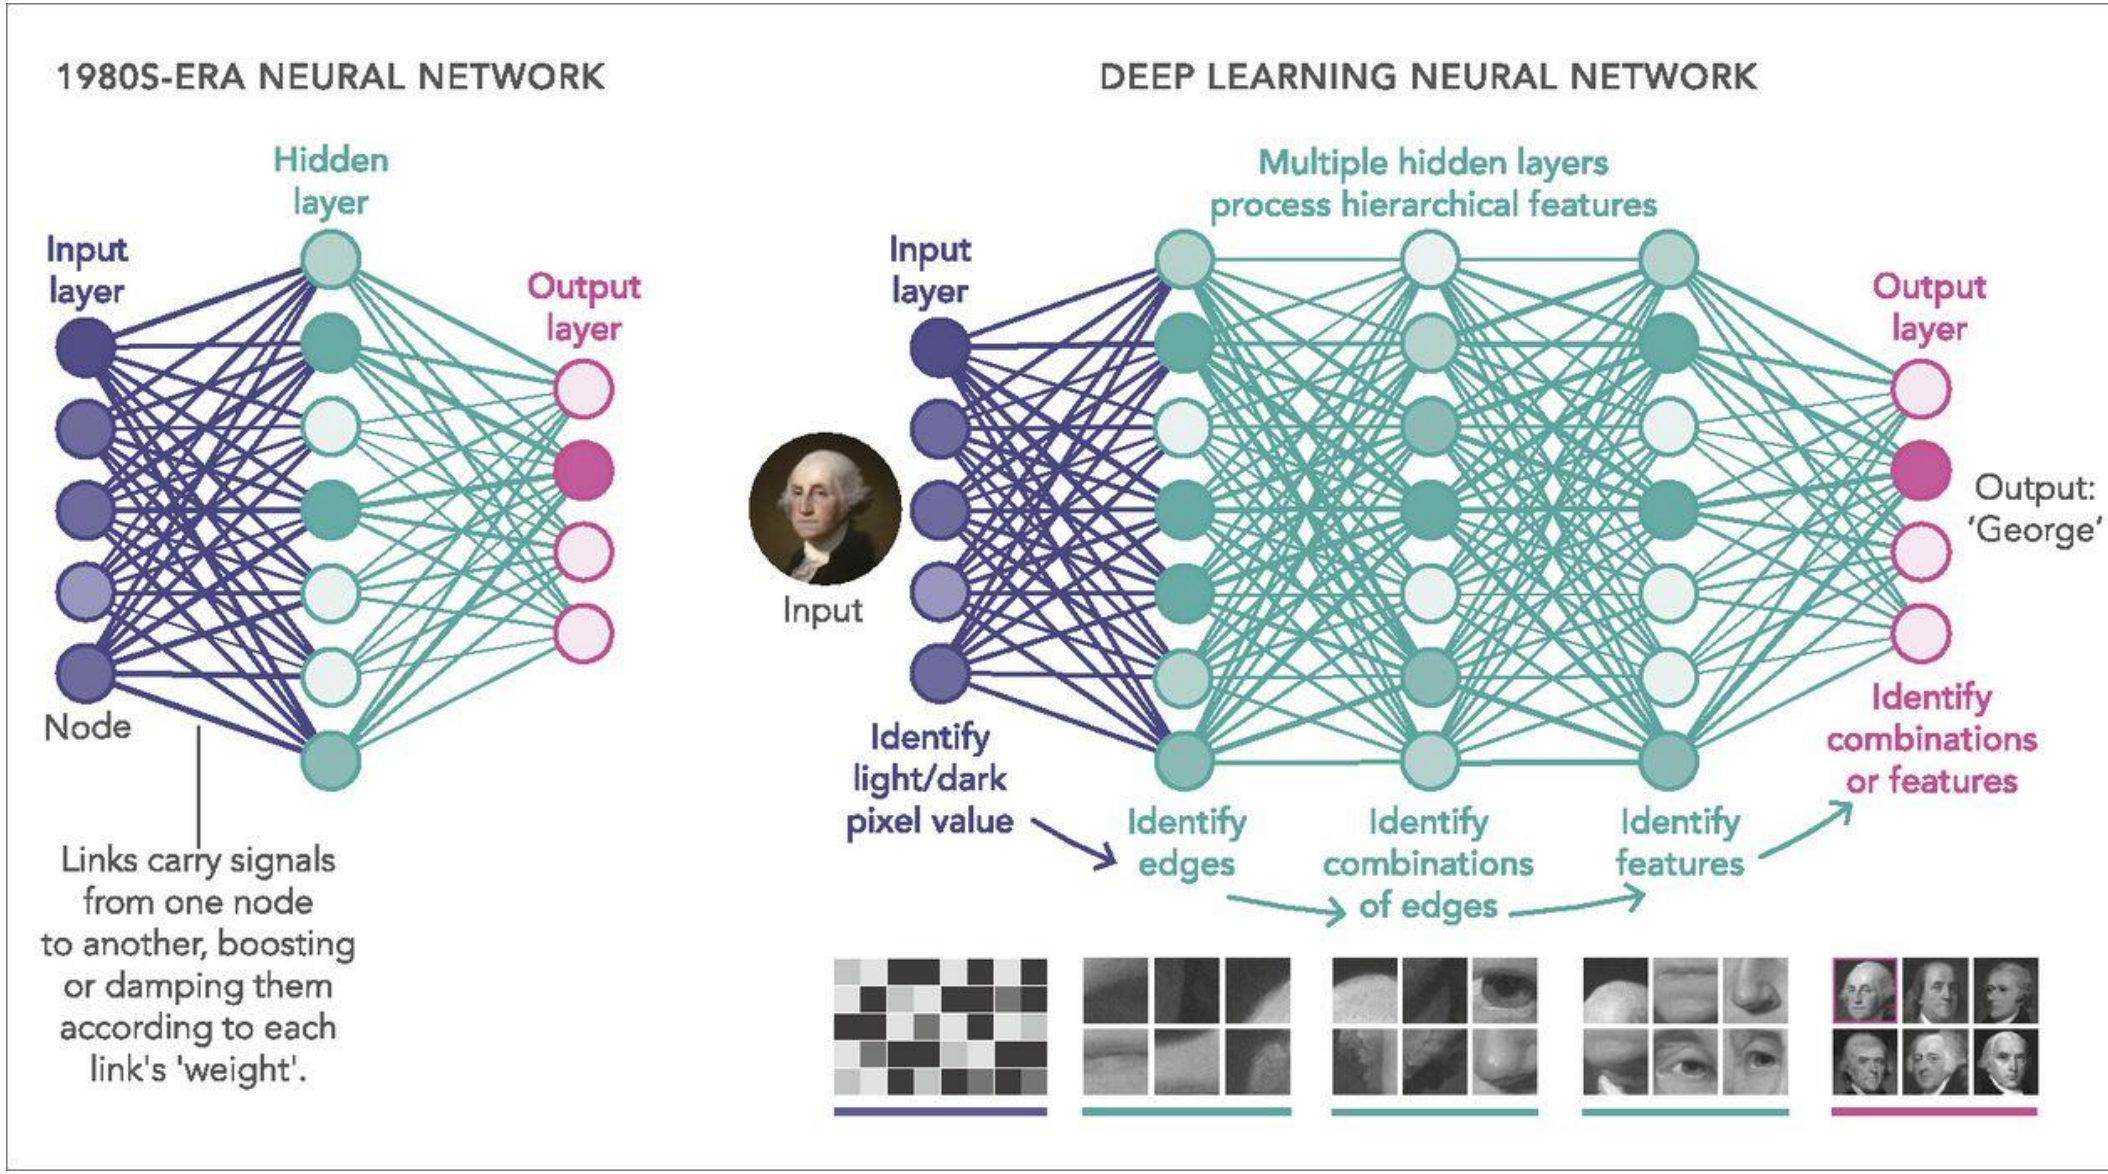
1980S-ERA NEURAL NETWORK
Hidden layer
Input layer
Output layer
Node
Links carry signals from one node to another, boosting or damping them according to each link's 'weight'.
DEEP LEARNING NEURAL NETWORK
Multiple hidden layers process hierarchical features
Input layer
Output layer
Input
Output: 'George'
Identify light/dark pixel value
Identify edges
Identify combinations of edges
Identify features
Identify combinations or features

GENERATIVE MODELS

- Autoencoders
- Adversarial Networks
- Sequence Models

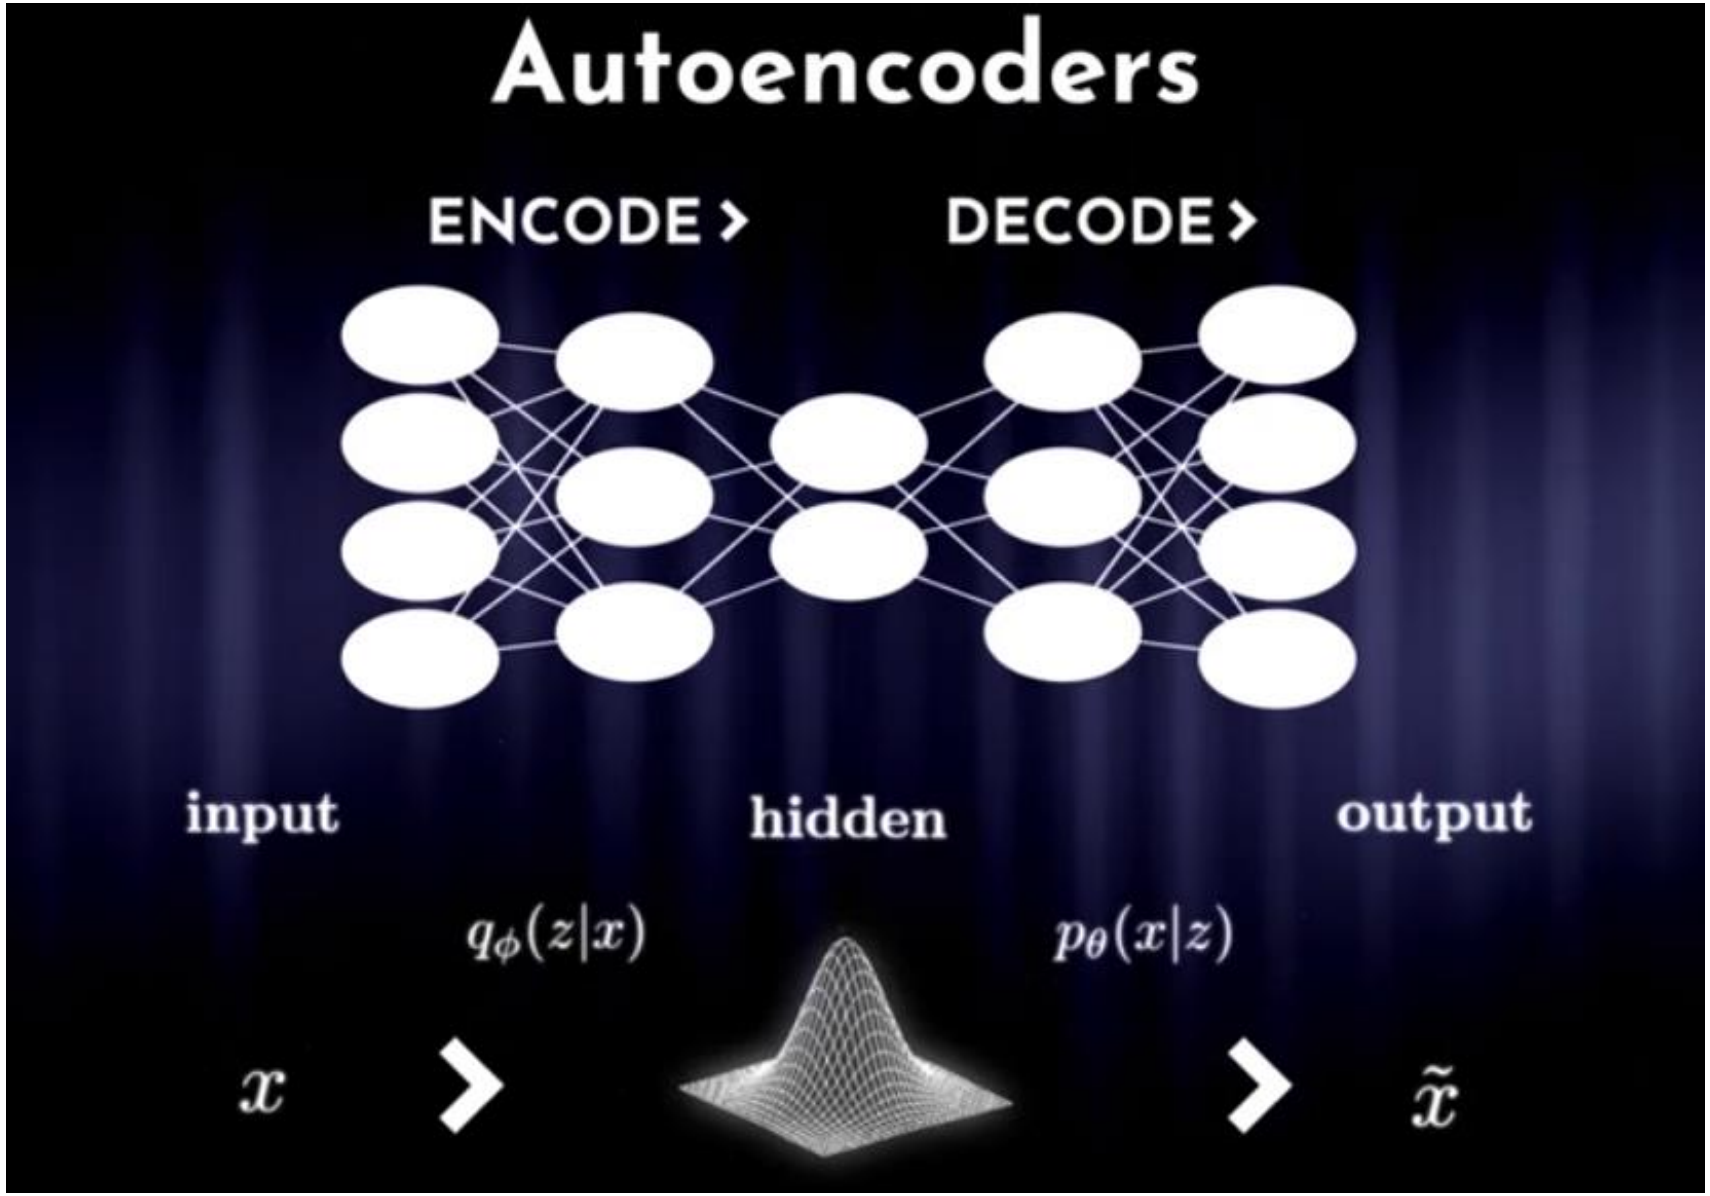
Autoencoders
ENCODE >
DECODE >
input
hidden
output
$q_\phi(z|x)$
$p_\theta(x|z)$
x
$\tilde{x}$

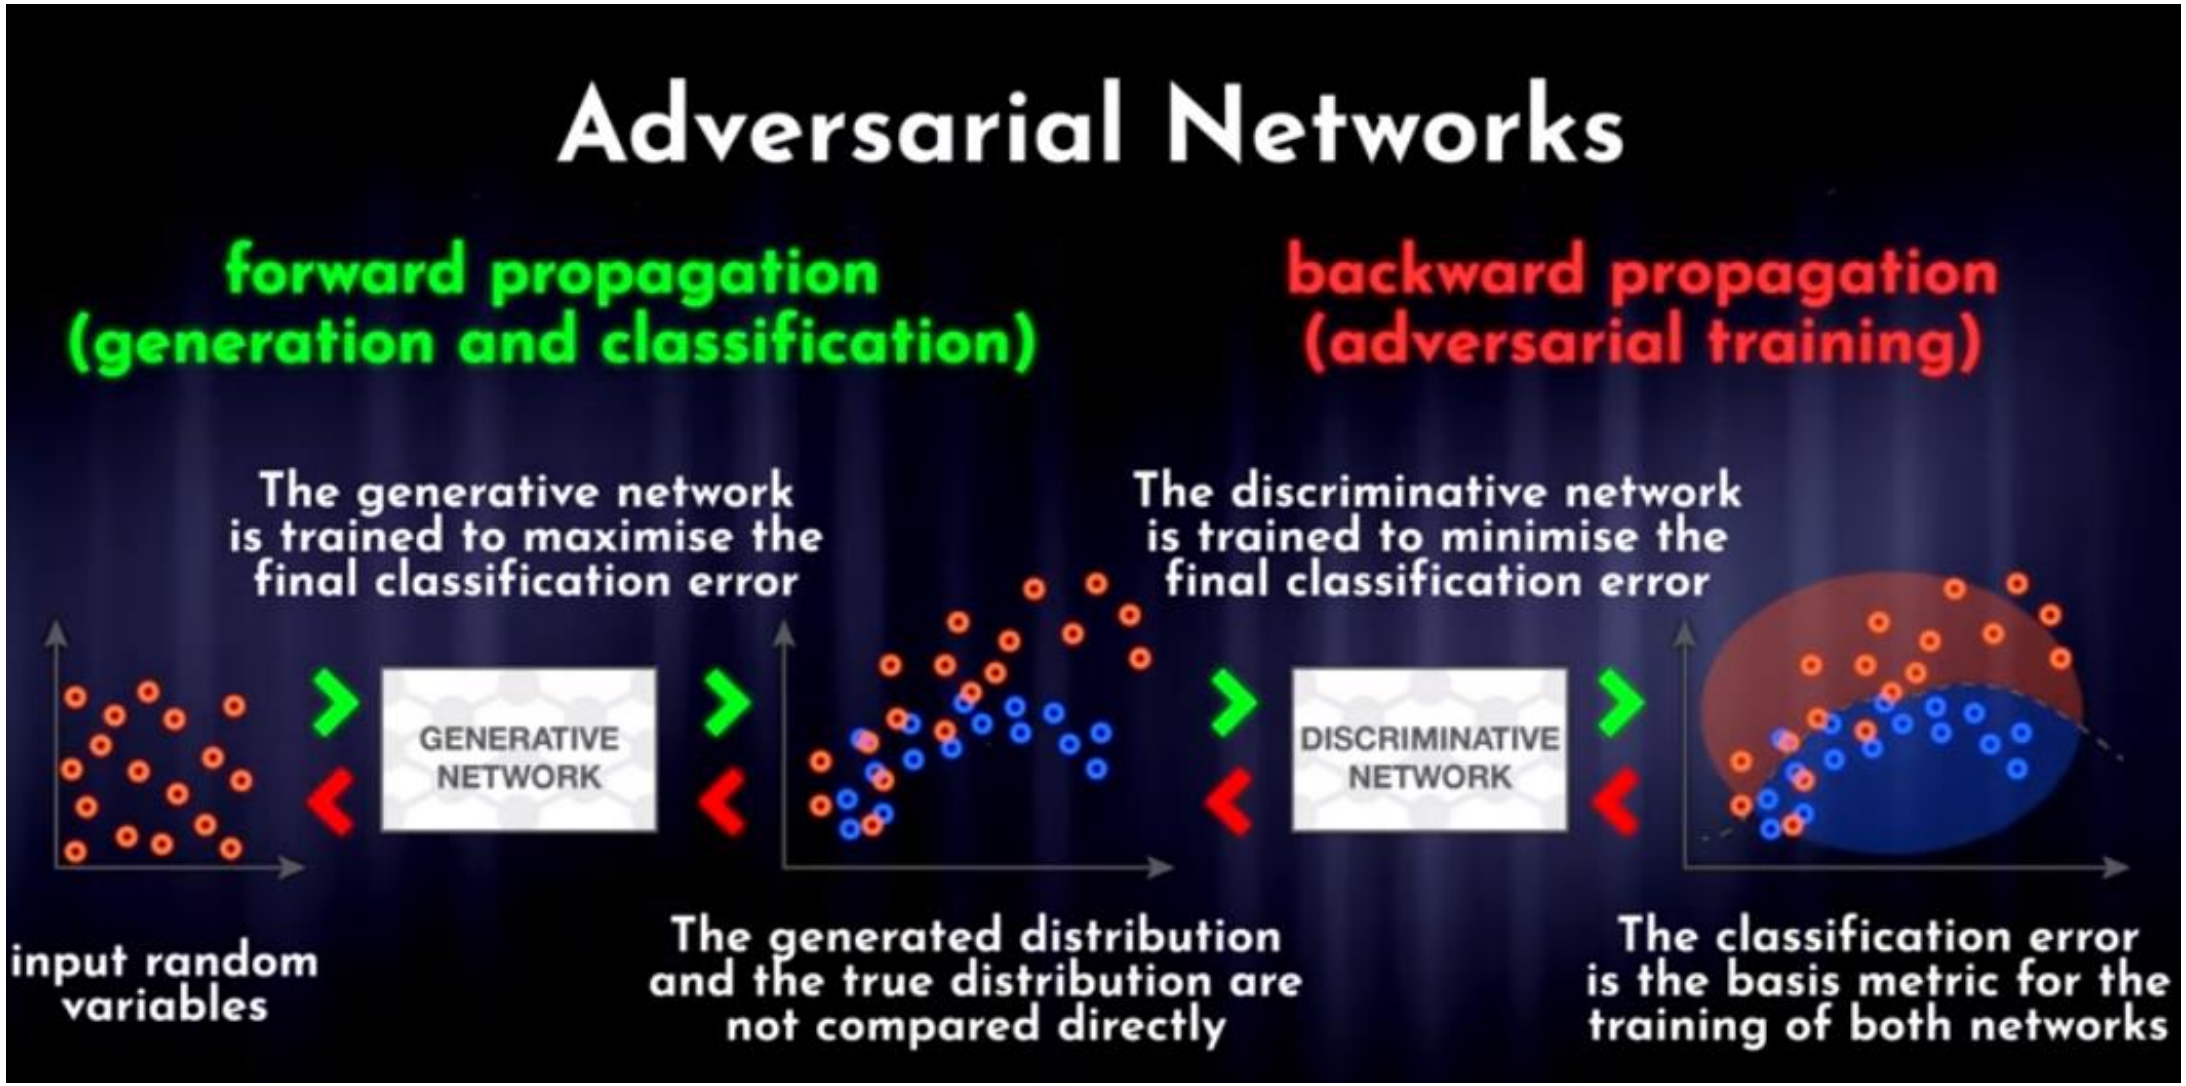
Adversarial Networks
forward propagation
(generation and classification)
backward propagation
(adversarial training)
The generative network
is trained to maximise the
final classification error
The discriminative network
is trained to minimise the
final classification error
GENERATIVE
NETWORK
DISCRIMINATIVE
NETWORK
input random
variables
The generated distribution
and the true distribution are
not compared directly
The classification error
is the basis metric for the
training of both networks

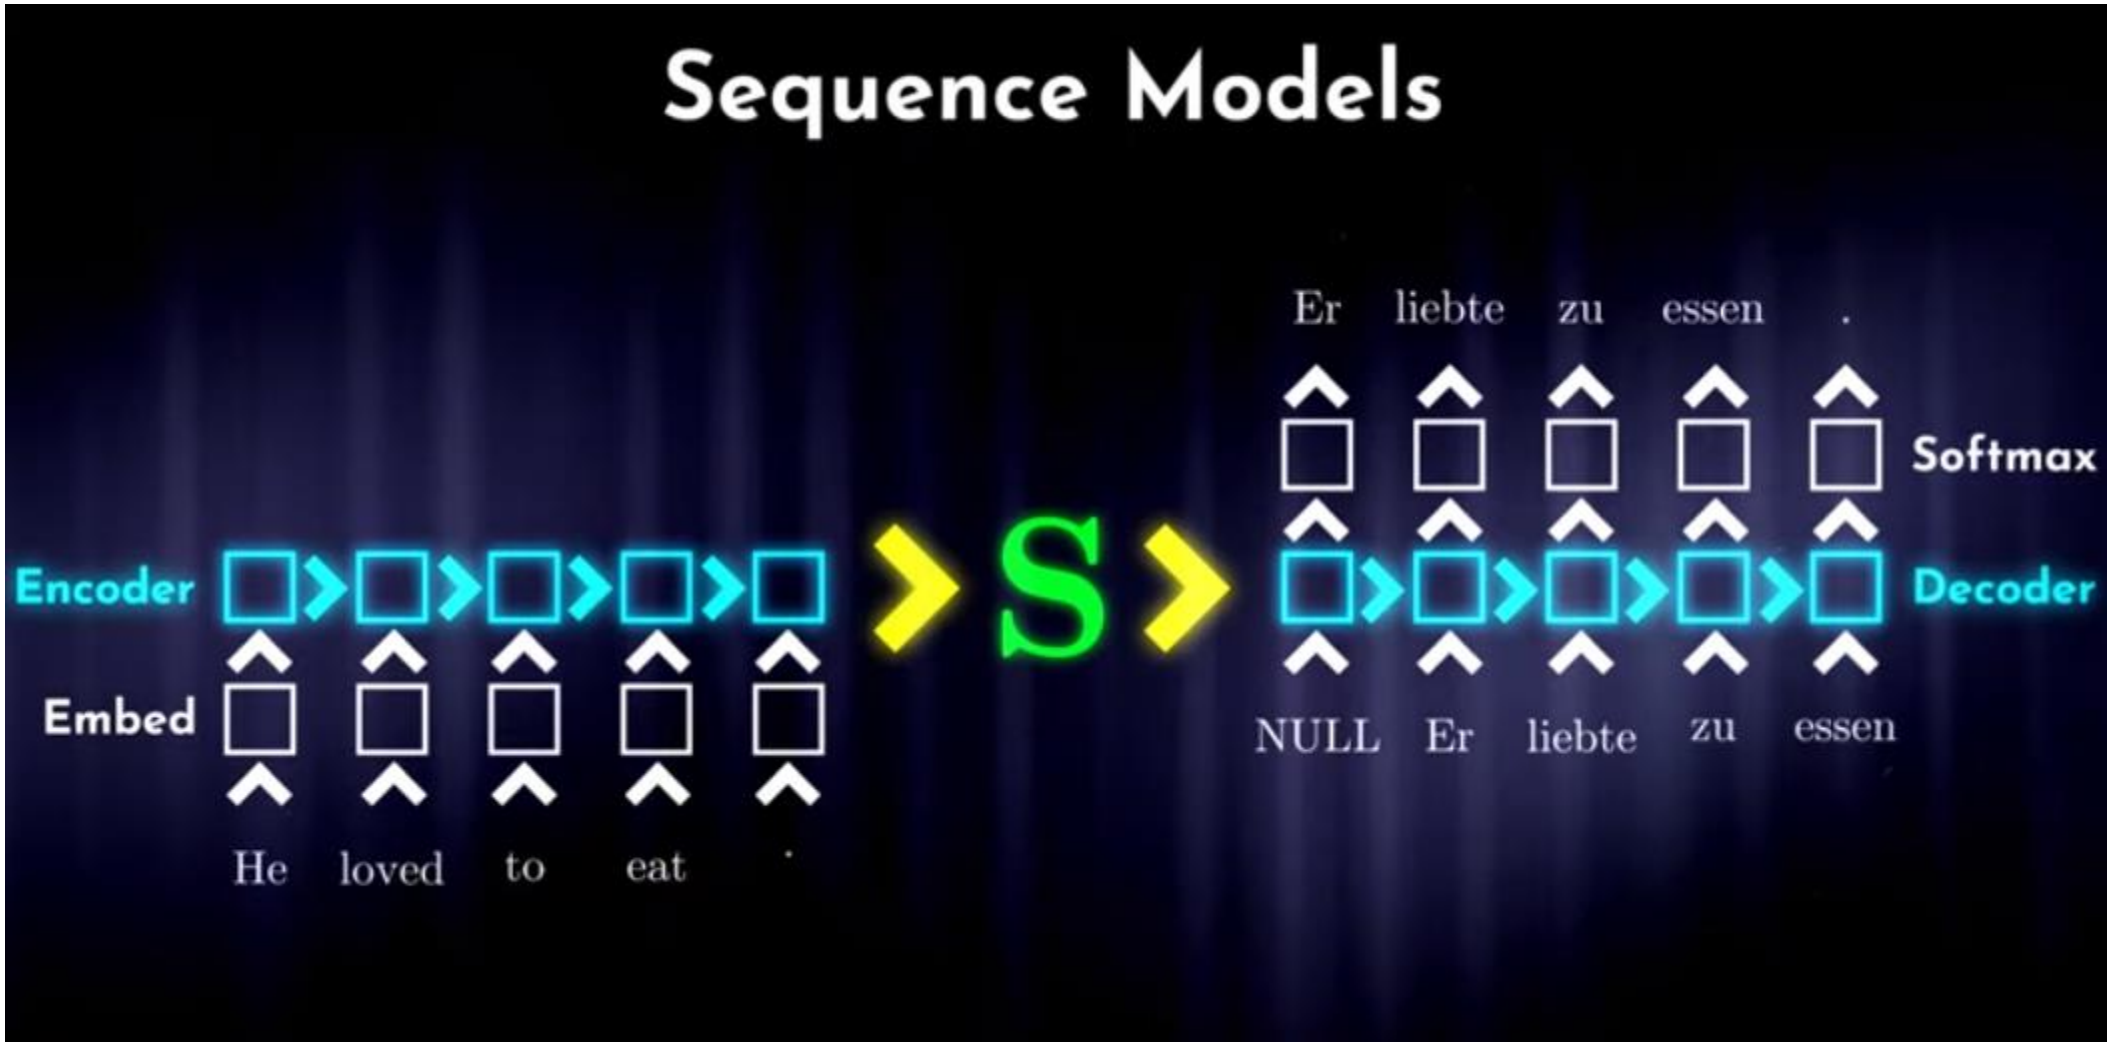
Sequence Models
Er liebte zu essen .
Softmax
Encoder
S
Decoder
Embed
He loved to eat .
NULL Er liebte zu essen

ChatGPT

1) Generative pretraining

2) Supervised fine-tuning

3) Reinforcement learning from human feedback

“I am learning about ChatGPT right now.”


Segmentation
Tokenization
Stop Words
Stemming
Named Entity Recognition
Speech Tagging
Lemmatization
0
1
2
3
4
i
be
learn
chatgpt
now
pronoun
auxillary verb
verb
noun
adverb
Software Entity

```
[
    [1, 0, 0, 0],     # pronoun
    [0, 1, 0, 0],     # auxiliary verb
    [0, 0, 1, 0],     # verb
    [0, 0, 0, 1],     # proper noun
    [0, 0, 0, 1],     # proper noun
    [0, 0, 0, 0],     # stop word
    [0, 0, 1, 0],     # verb
    [0, 0, 0, 0],     # stop word
    [0, 0, 0, 1]      # noun
]
```

GPT
GENERATIVE PRETRAINED TRANSFORMER

I am learning about GPTChat right now

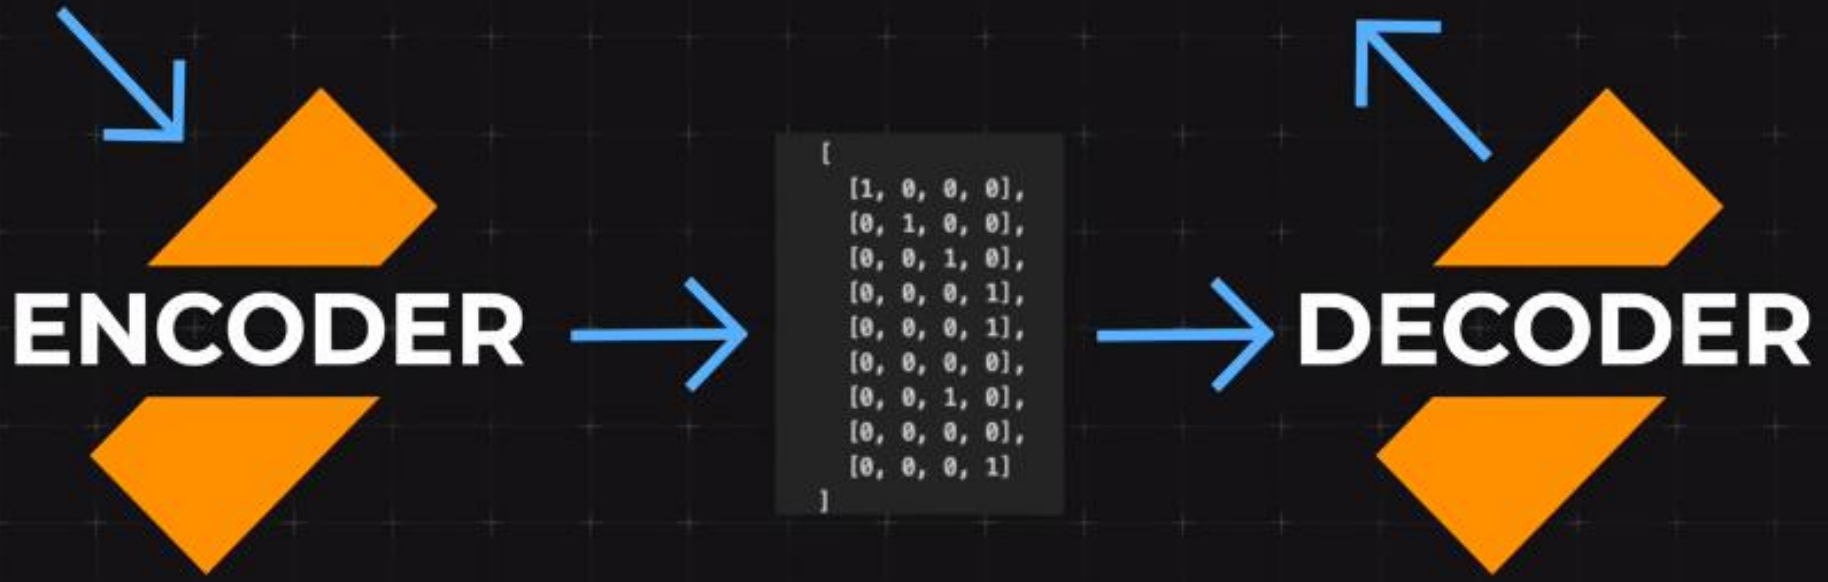

As a result, I now have a better understanding of how ChatGPT works and what it can do.

Sequence to Sequence Transformation

Attention Is All You Need

Ashish Vaswani*
Google Brain
avaswani@google.com

Noam Shazeer*
Google Brain
noam@google.com

Niki Parmar*
Google Research
nikip@google.com

Jakob Uszkoreit*
Google Research
usz@google.com

Llion Jones*
Google Research
llion@google.com

Aidan N. Gomez* †
University of Toronto
aidan@cs.toronto.edu

Łukasz Kaiser*
Google Brain
lukaszkaiser@google.com

Illia Polosukhin* ‡
illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while

INITIAL MODEL TRAINED WITH SUPERVISED FINE-TUNING

HUMAN WRITES PROMPT

e.g. What is a car?

HUMAN ANSWERS PROMPT

e.g. A vehicle moving on wheels

REWARD MODEL

Human writes prompt.
e.g. ***What is a car?***

A	B	C	D
a vehicle moving on wheels	a Disney movie	something similar to a bycicle	a tool typically found in the bathroom
0.98	0.43	0.56	0.13

A > C > B > D

OPTIMIZING A POLICY AGAINST THE REWARD MODEL

PPO: Proximal Policy Optimizer

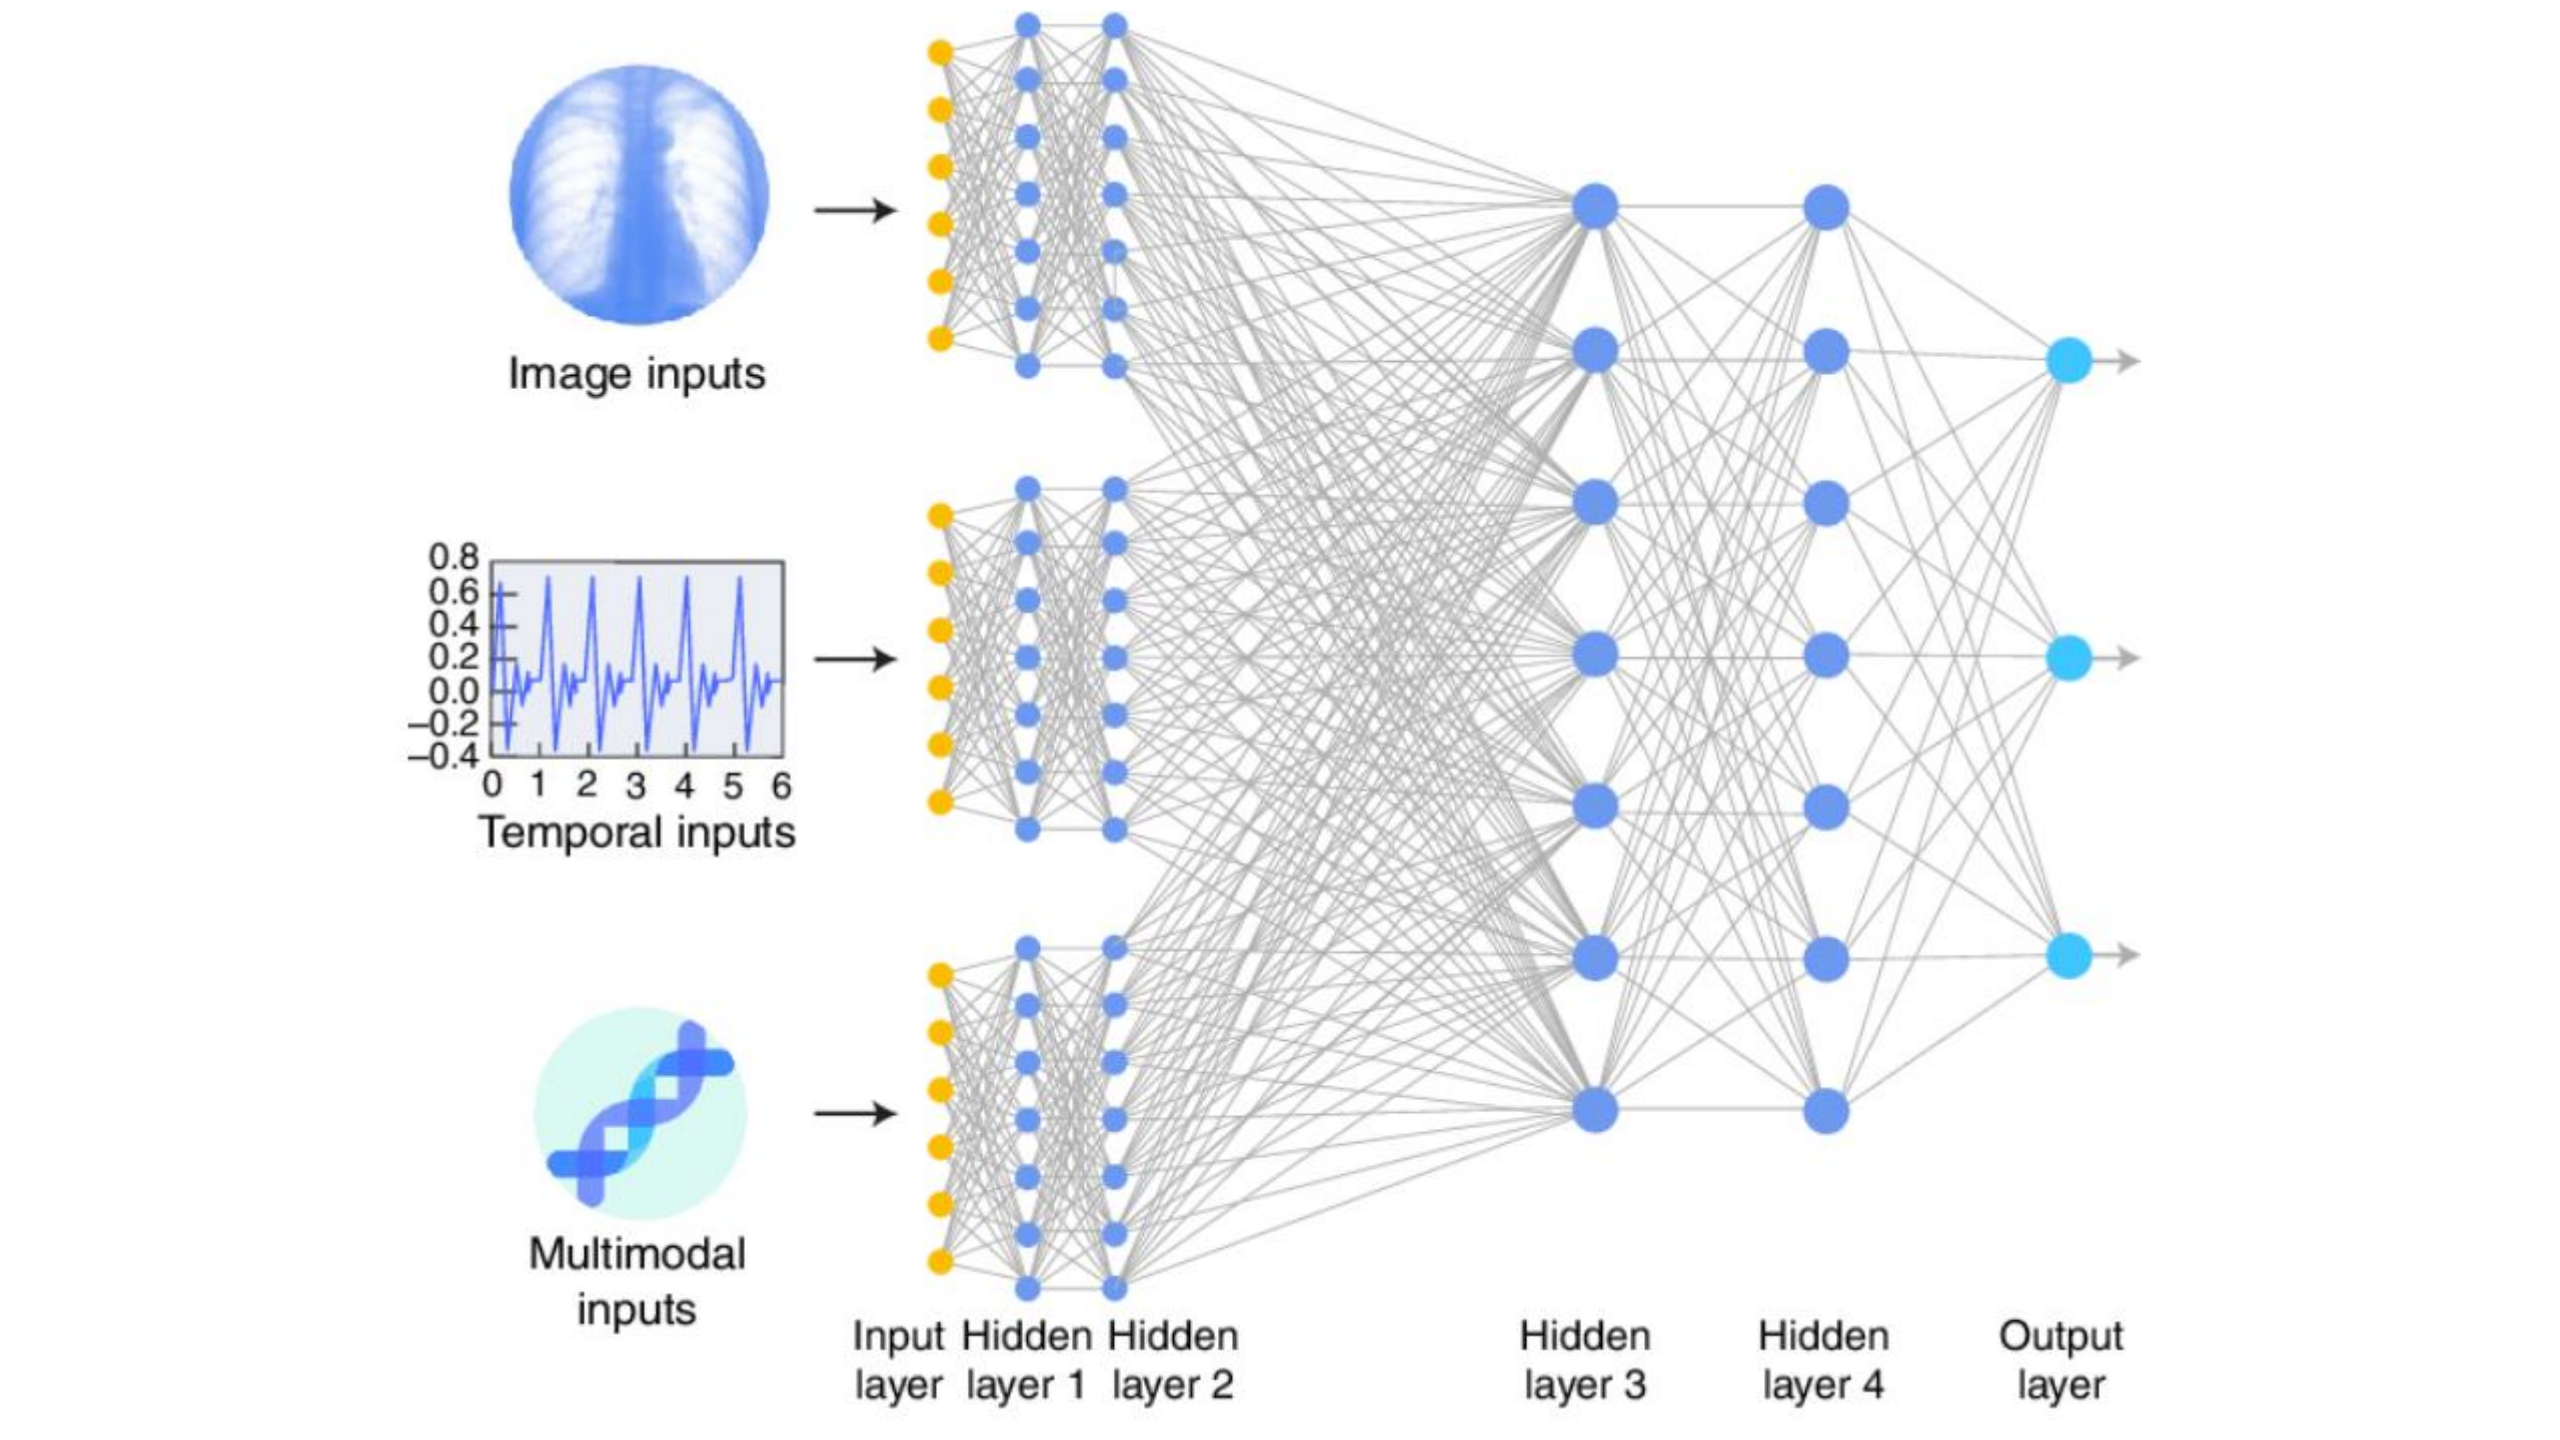
Image inputs
0.8
0.6
0.4
0.2
0.0
−0.2
−0.4
0 1 2 3 4 5 6
Temporal inputs
Multimodal inputs
Input layer
Hidden layer 1
Hidden layer 2
Hidden layer 3
Hidden layer 4
Output layer

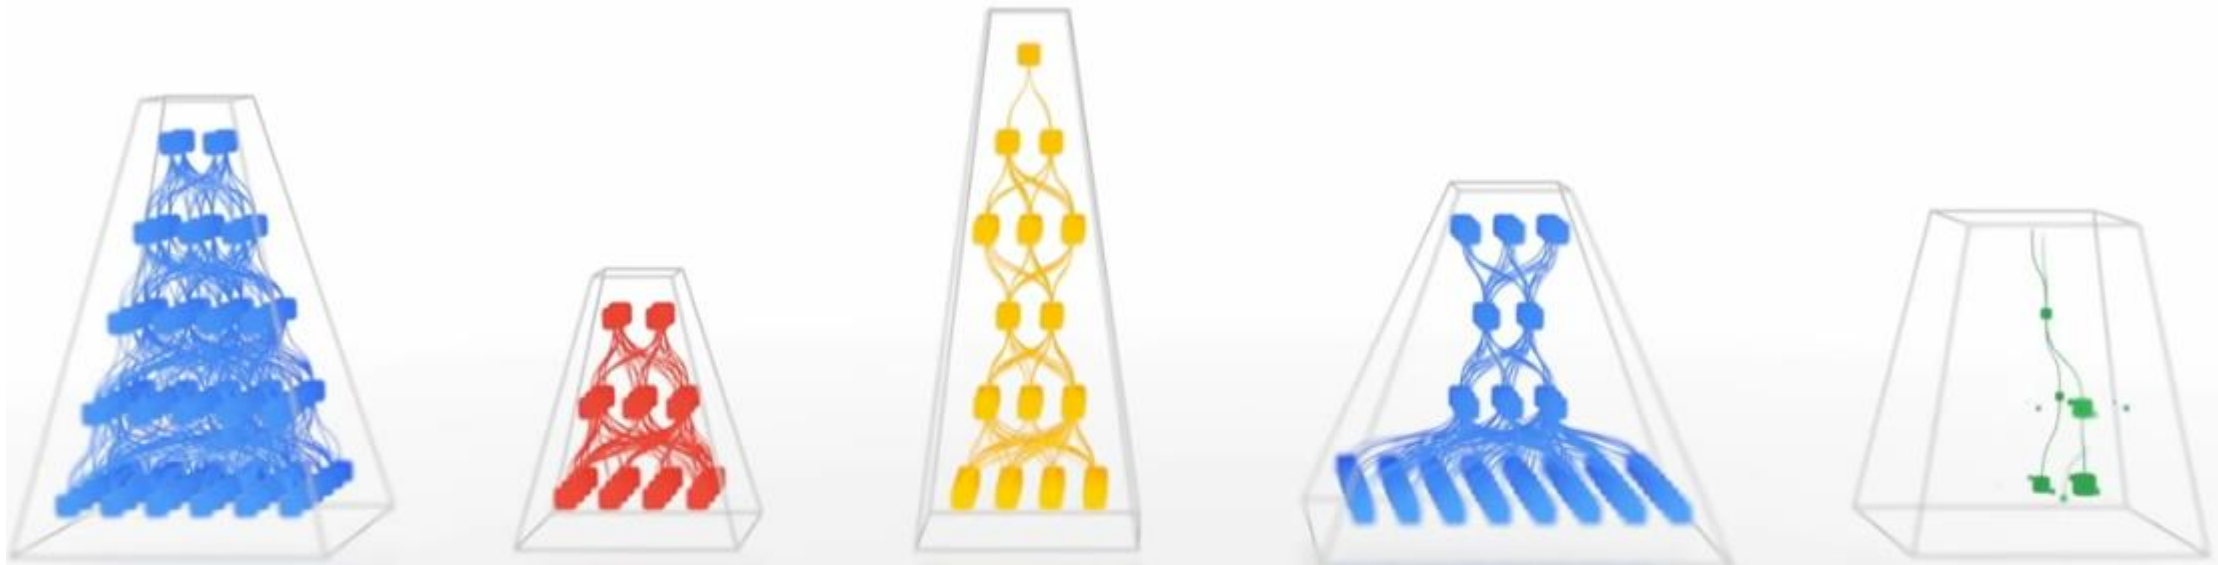

Separate models for thousands of tasks

How Cheap Labor Drives China's A.I. Ambitions

Workers at the headquarters of Ruijin Technology Company in Jiaxian, in central China's Henan Province. They identify objects in images to help artificial intelligence make sense of the world. Yan Cong for The New York Times

Optical Adversarial Attack Can Change the Meaning of Road Signs

Published 1 year ago on August 16, 2021
By **Martin Anderson**

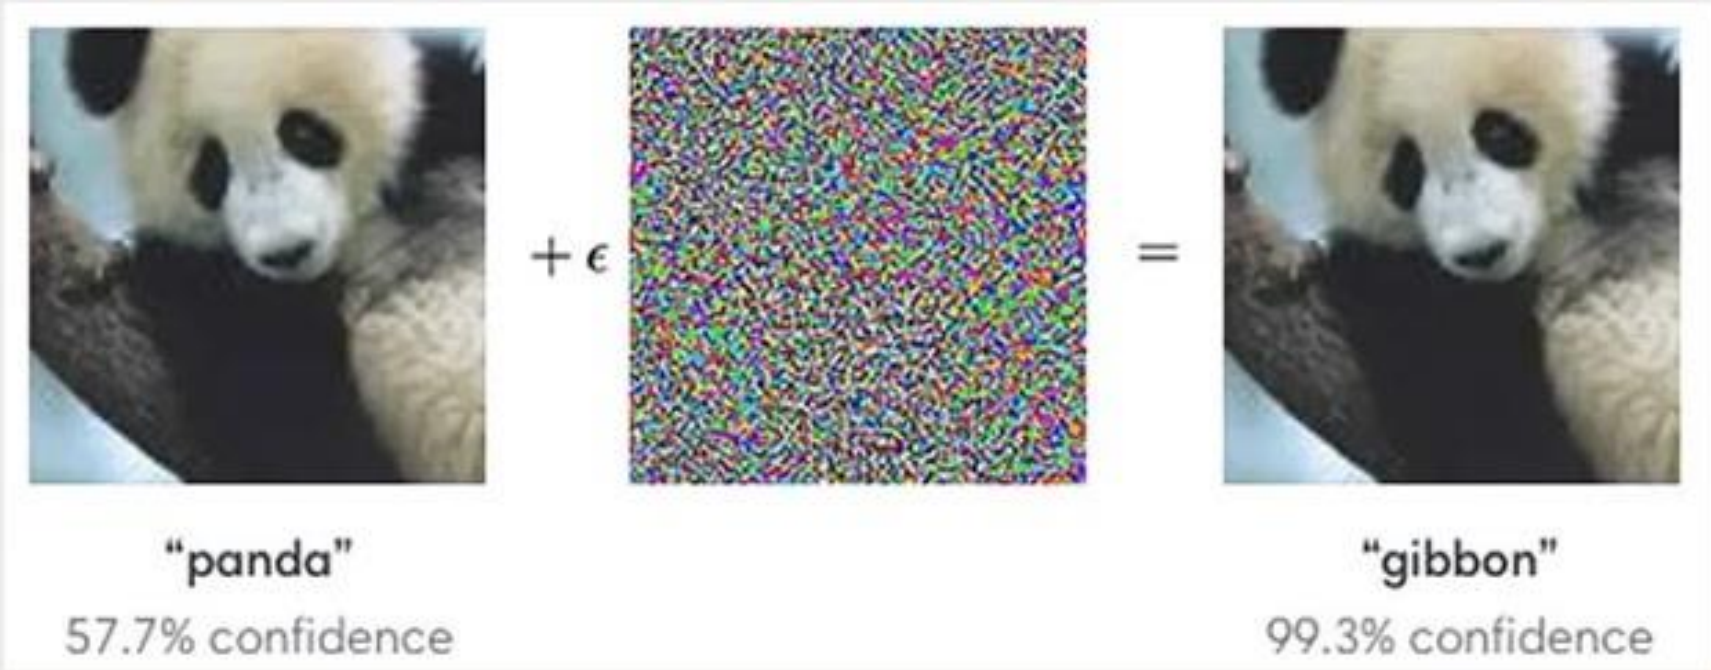

An image of a panda, when combined with an adversarial input, can convince a classifier that it’s looking at a gibbon. IMAGE: OPENAI

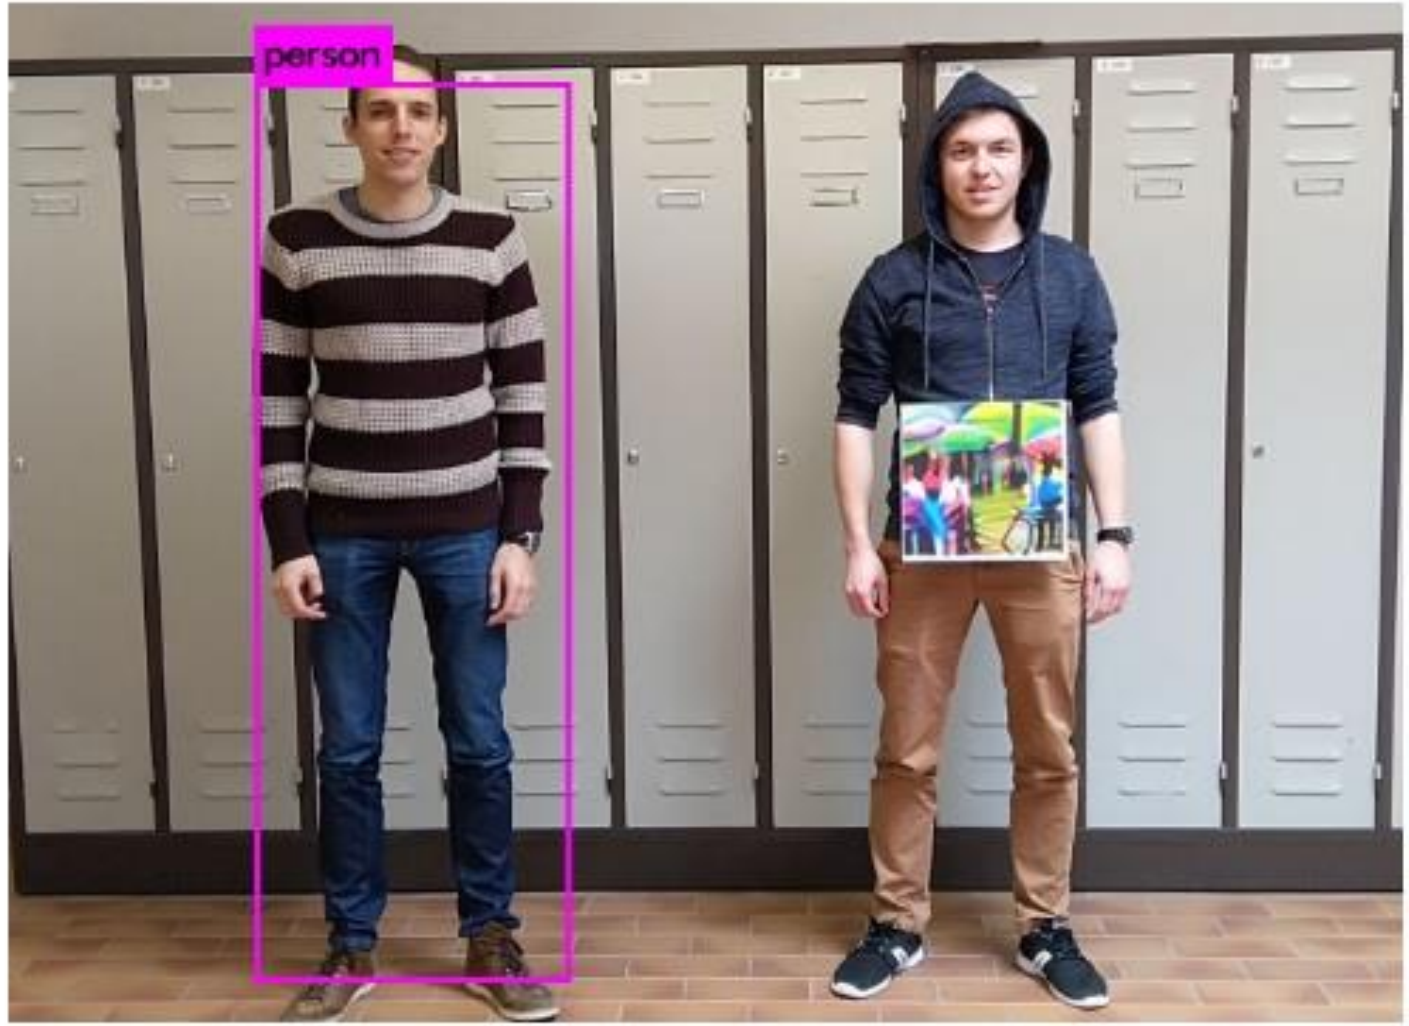

Figure 1: We create an adversarial patch that is successfully able to hide persons from a person detector. Left: The person without a patch is successfully detected. Right: The person holding the patch is ignored.