An introduction to Generative Adversarial Networks and their uses

Stephen McGough
Newcastle University / Fellow at the Alan Turing Institute

Petroleum Exploration Society of Great Britain / Royal Statistical Society
Machine Learning SIG meeting
25th March 2021, Online

Overview

- Deep Learning Primer
- Generative Adversarial Network (GAN)
  - Generating data from random noise
  - E.g., generating traffic data
- Conditional GAN – Pix2Pix
  - Converting one dataset into another
  - E.g., Segmentation
- Style transfer – CycleGan
  - Transferring without matched data
  - E.g., Tracking Larvae

Deep Learning Primer

What is AI, ML and DL?

- Artificial Intelligence (AI)
  - A system which exhibits characteristics which could be seen as intelligent
- Machine Learning (ML)
  - A system which is able to learn and improve its ability
- Deep Learning (DL)
  - A system which uses (Deep) Neural Networks to exhibit ML

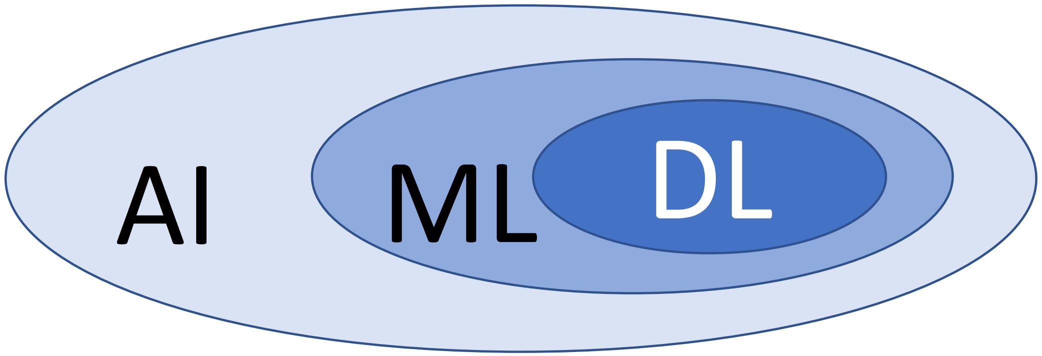

Machine Learning vs Deep Learning

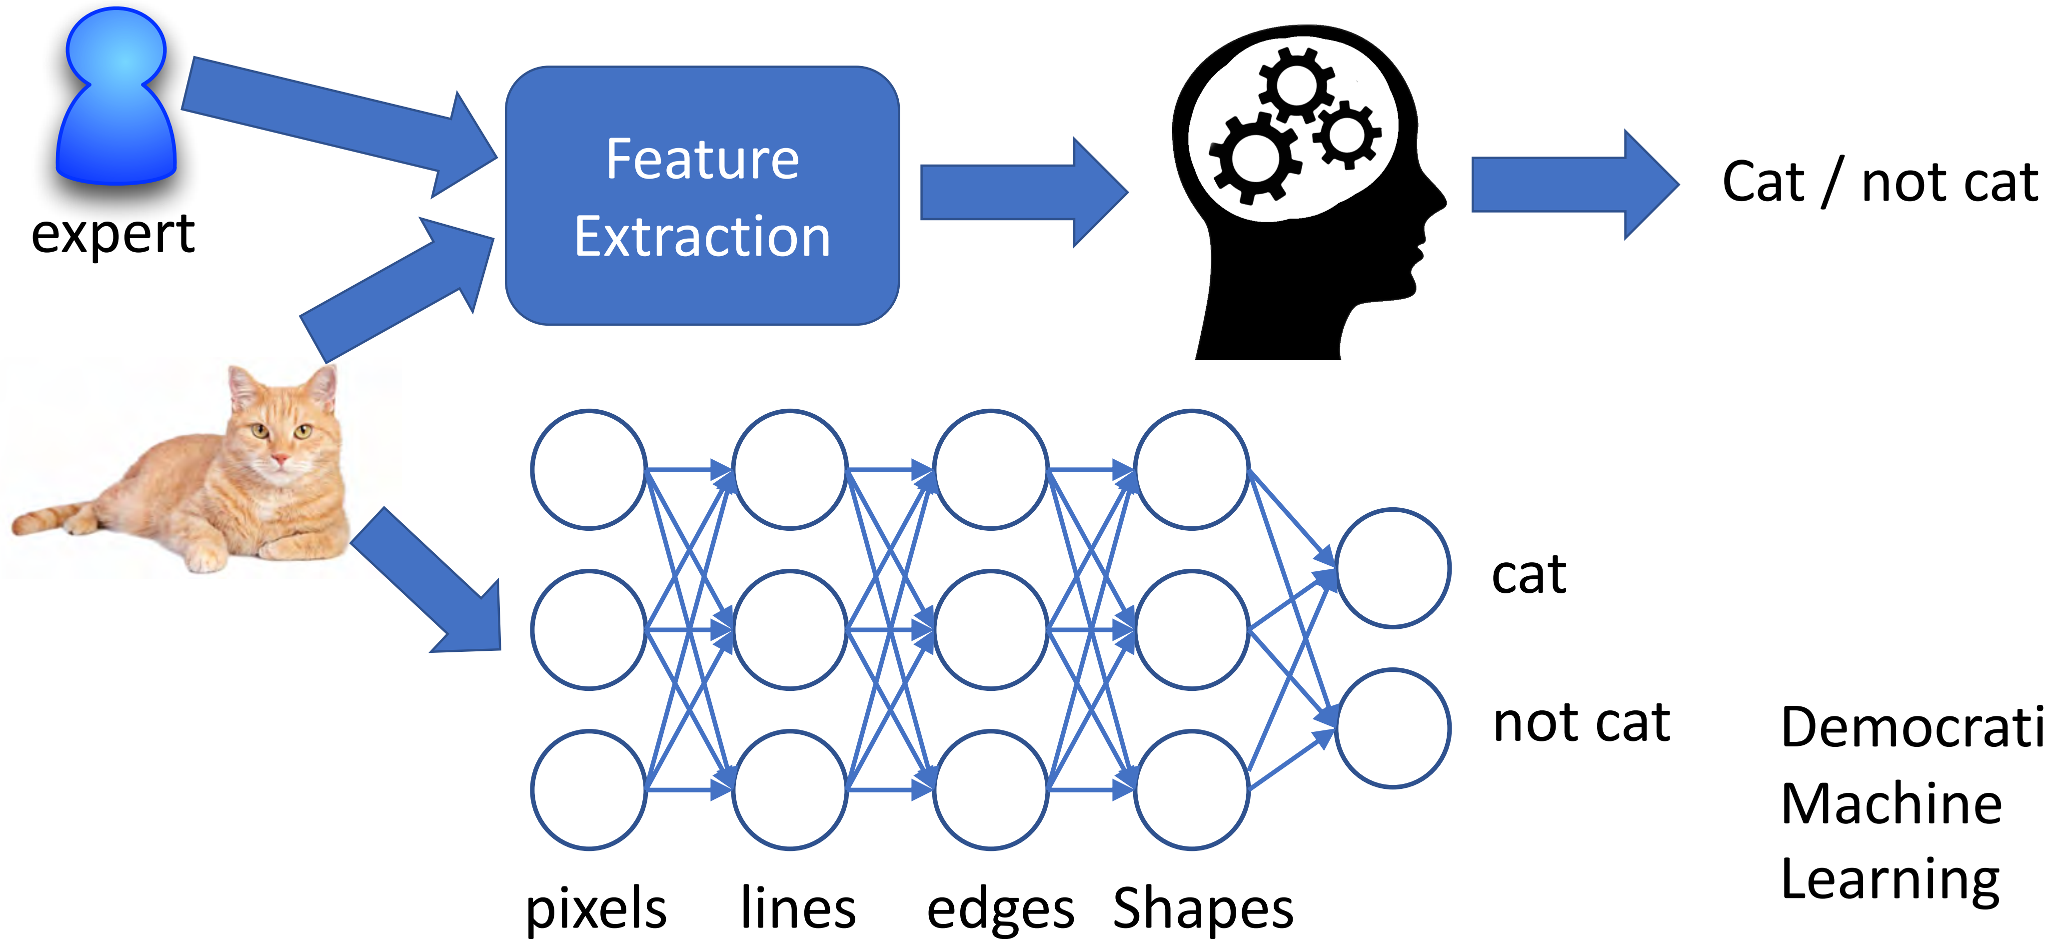

Democratizing Machine Learning

Basic building blocks:
The data

- Data is key here: Sample as 1D array of values

X[1]

X[2]

X[3]

X[4]

X[5]

X[6]

Basic building blocks: The neuron

- Sums up all of the input values
  - But that's not very clever – what if x[1] is more important than x[4]?

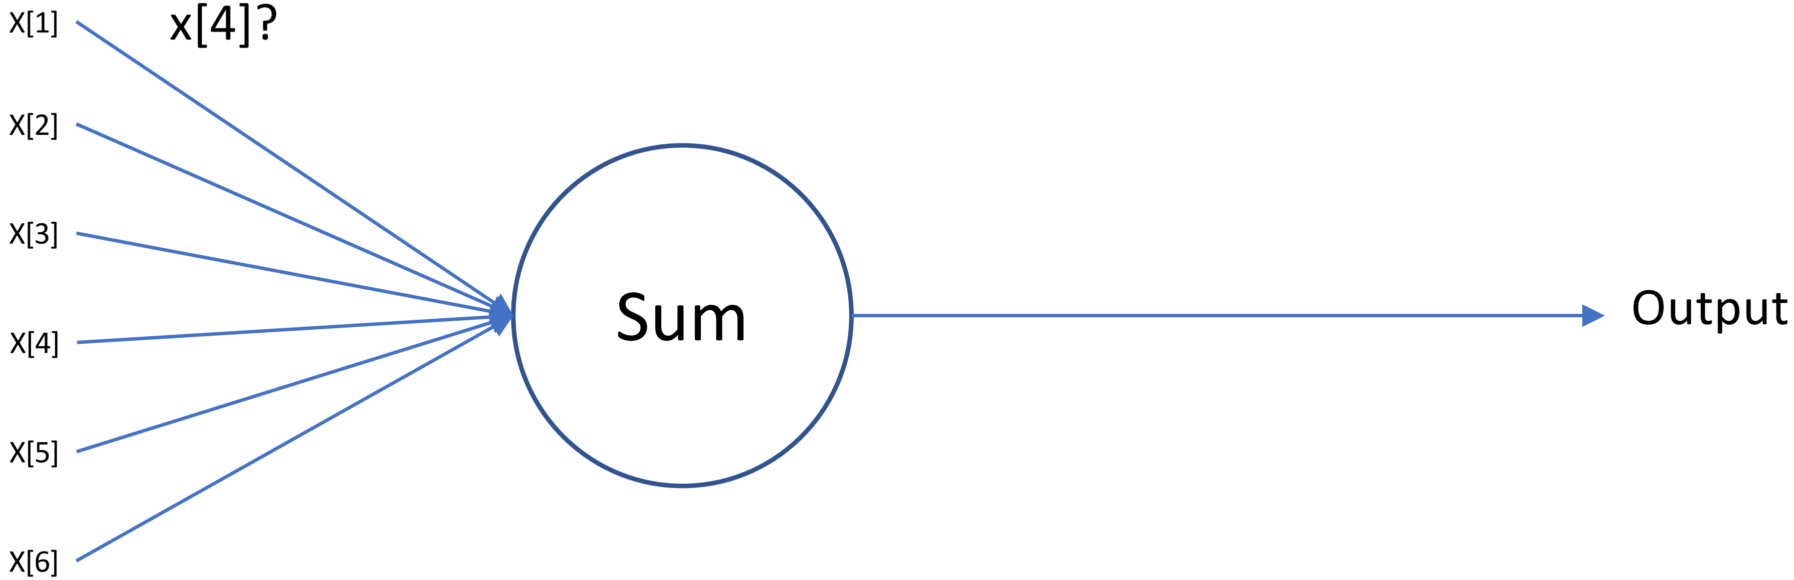

Basic building blocks: Weights

- The weights – so we can attribute importance to each x[]
  - Multiply x[] by w[] before adding them all together - still not that clever

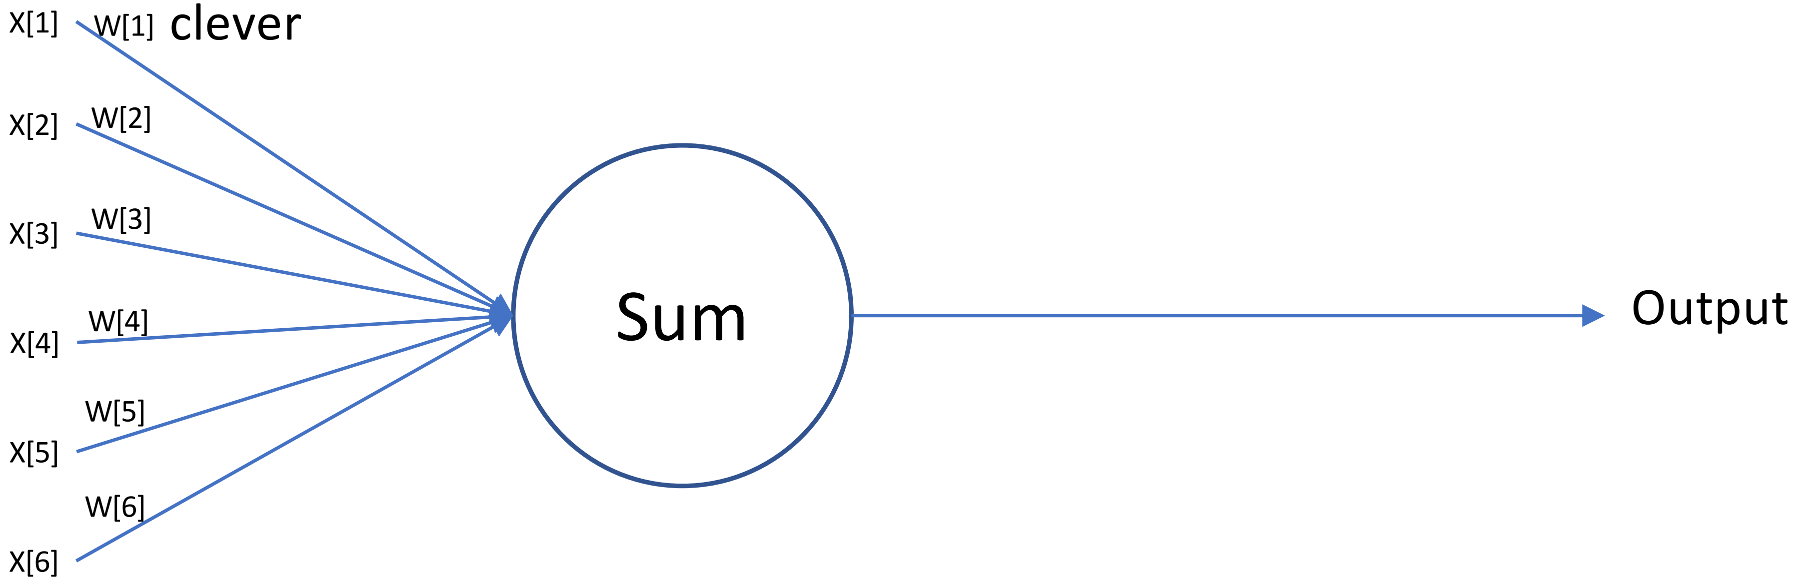

Basic building blocks:
Bias

- The bias is a value we add to the output
  - A constant 'fix' – cleverer – but still not good enough – everything is liner

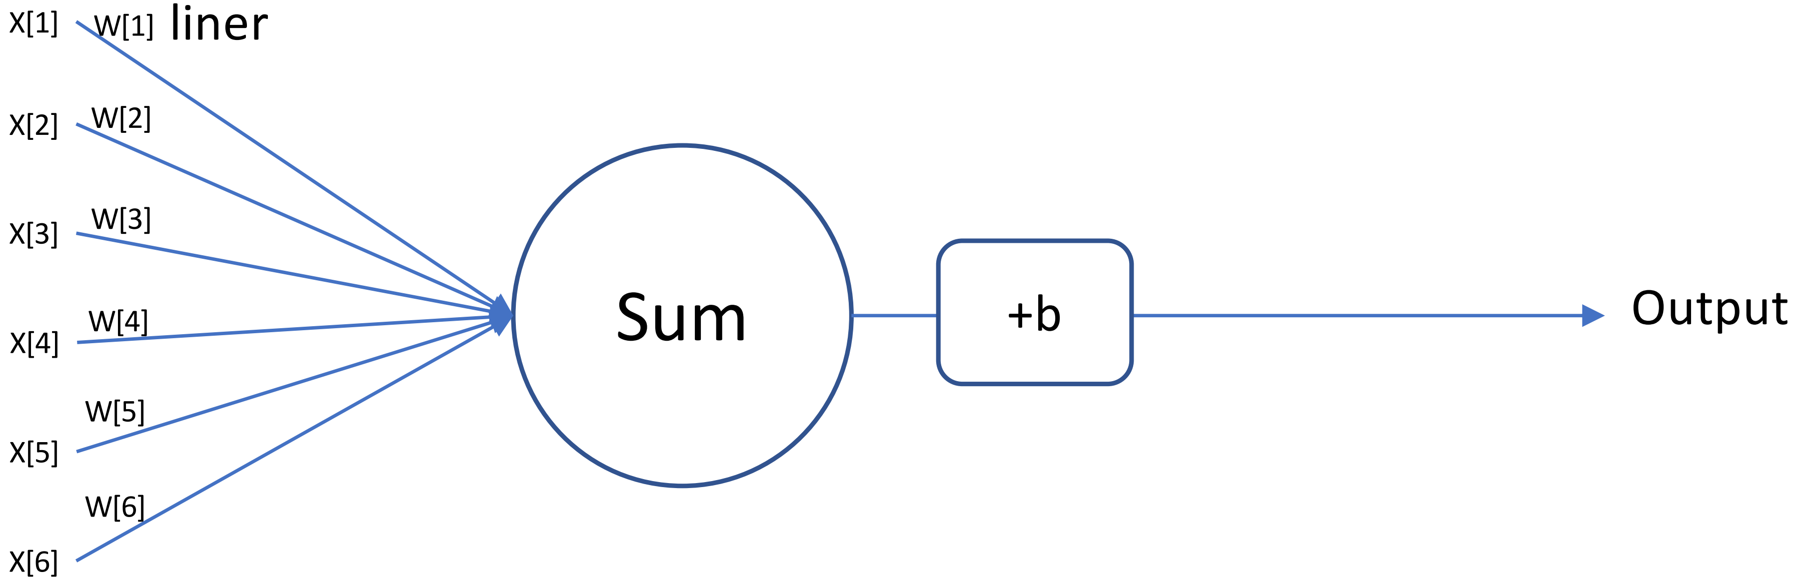

Basic building blocks: Activation function

- The activation function is a non-linear operation applied to the output
  - Allows much more complex things to be learnt

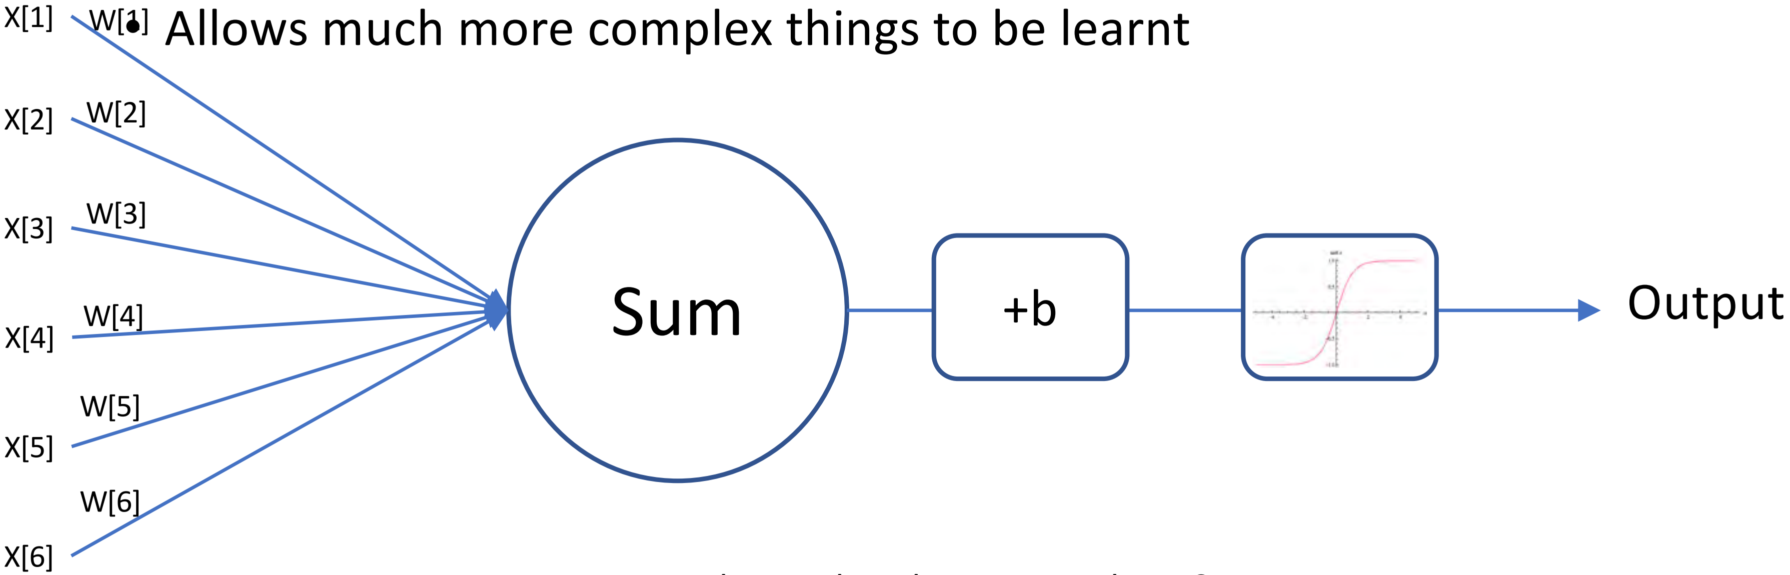

What makes the neuron clever?

- The right values of w[], b
- Trained by passing lots of examples through and modifying these values

Building a full Neural Network

- A single block on its own can't do much
  - So, we use a whole set of them to make a neural network

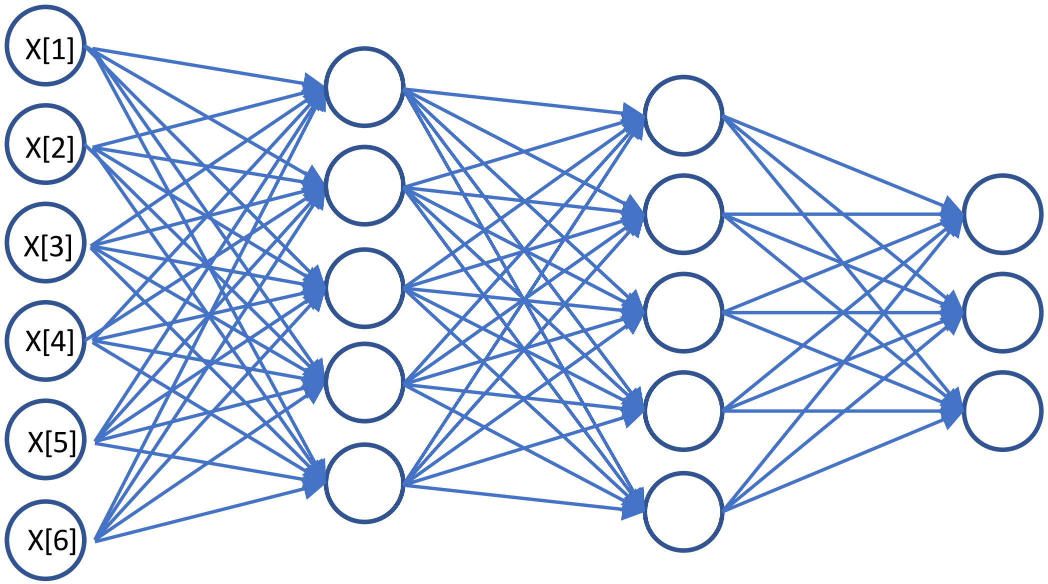

Building a full Neural Network

- A single block on its own can't do much
  - So, we use a whole set of them to make a neural network

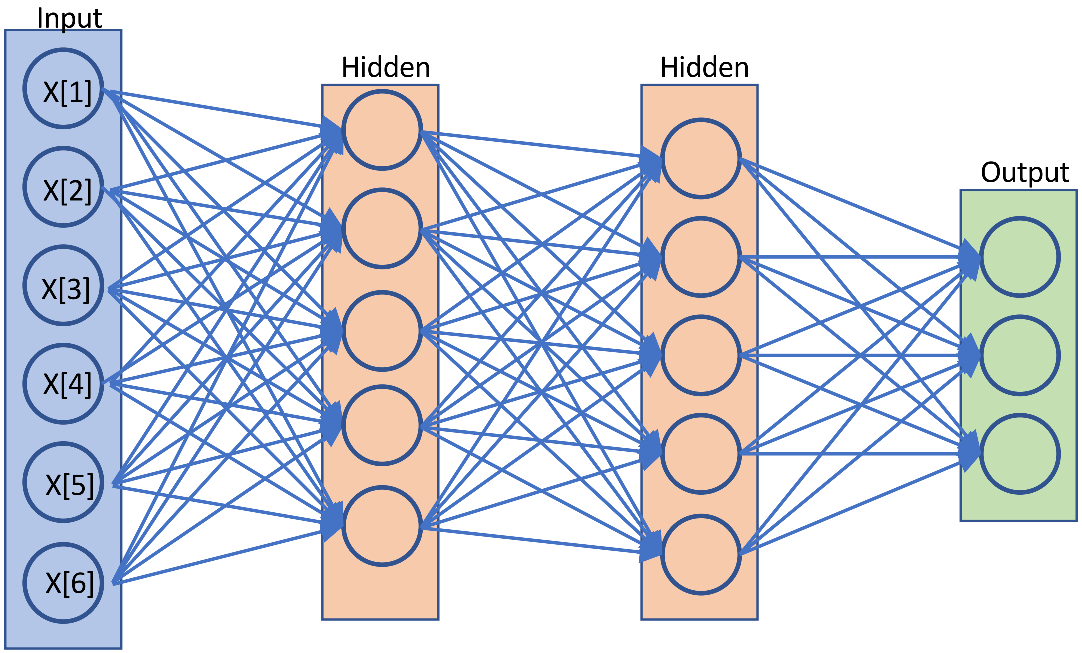

Building a full Neural Network

- A single block on its own can't do much
  - So, we use a whole set of them to make a neural network

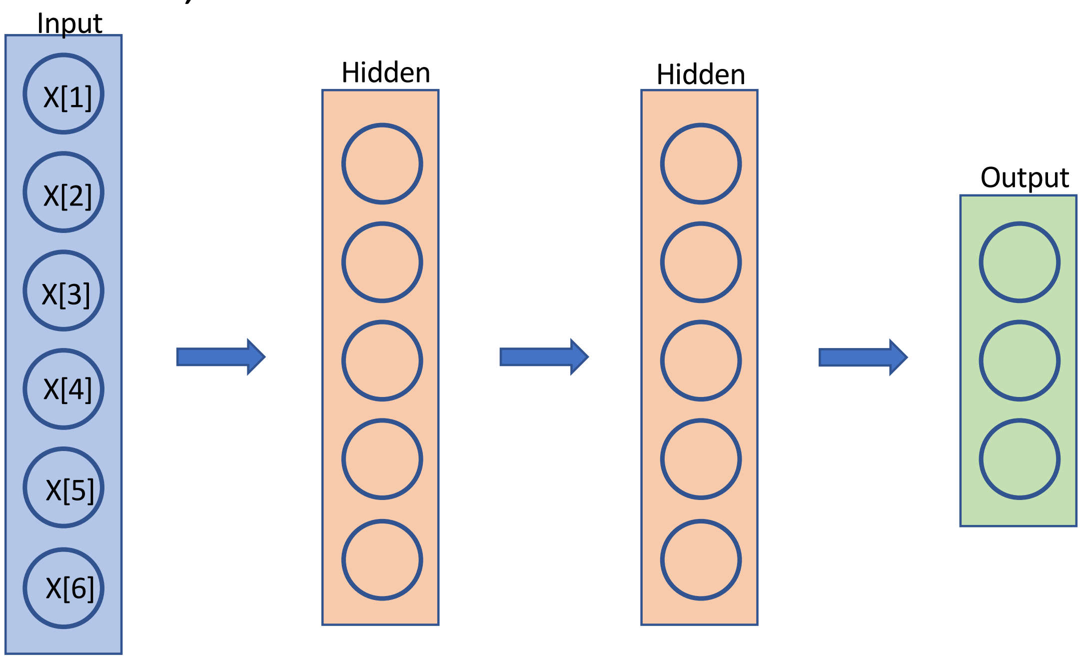

Building a full Neural Network

- A single block on its own can't do much
  - So, we use a whole set of them to make a neural network

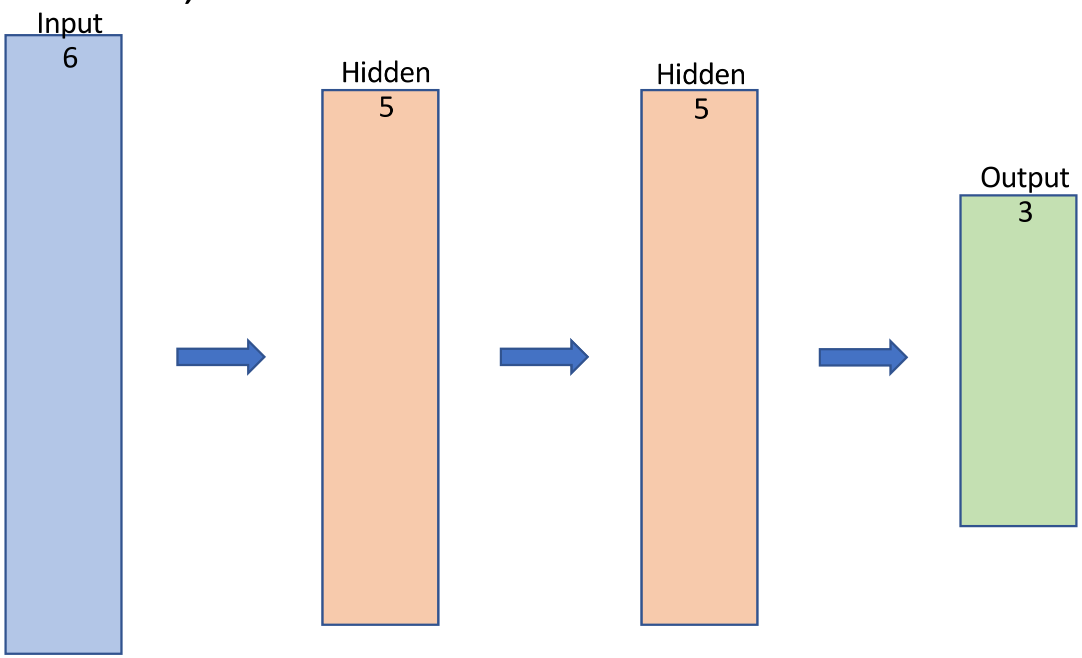

Called:

- Multi-Layer Perceptron (MLP)
- Fully Connected Layers
- Dense layers

So Why Deep Learning?

- Named due to the depth of the network
- Number of layers in the network
- ‘Real’ networks have 10’s, 100’s or 1000’s of layers
- These networks are often referred to as models
- Why now?
    - Data is available
    - Powerful (GPUs) to train them

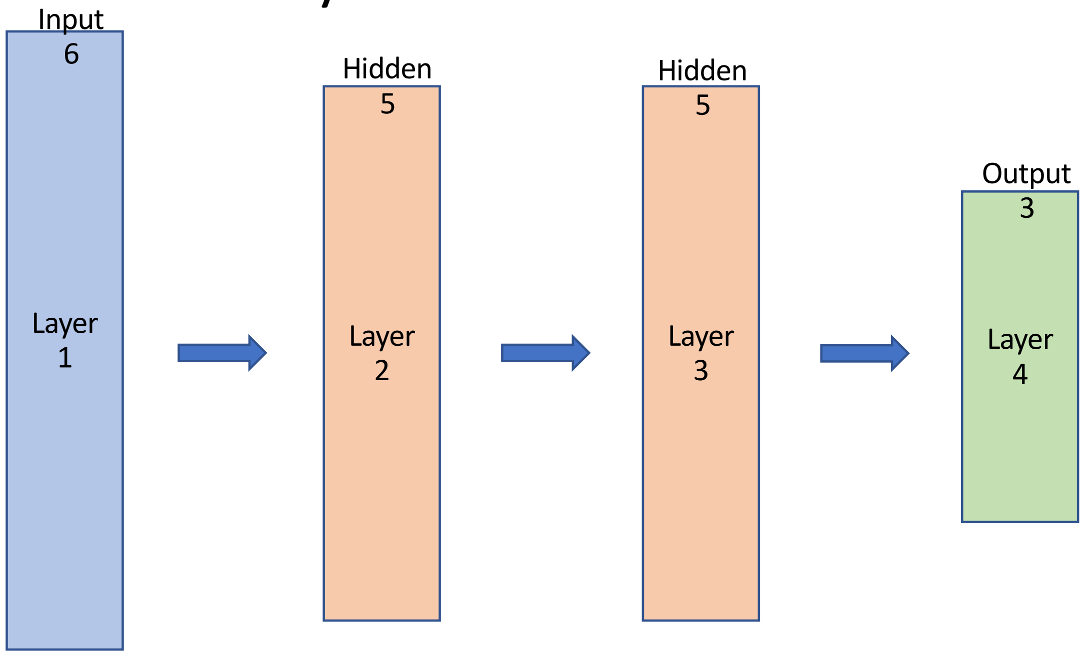

Further Reading

- Language: Python
- For Keras:
  - Deep Learning with Python, François Chollet
  - Read online at: https://www.manning.com/books/deep-learning-with-python
- For PyTorch:
  - Programming PyTorch for Deep Learning, Ian Pointer
  - Read online at: https://www.oreilly.com/library/view/programming-pytorch-for/9781492045342/
- If you want all the Deep Learning theory:
  - Deep Learning, Ian Goodfellow, Yoshua Bengio, Aron Courville
  - Read online at: https://www.deeplearningbook.org
- Platform
  - https://colab.research.google.com

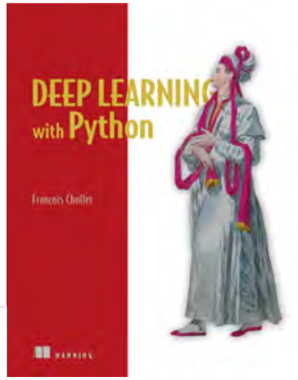

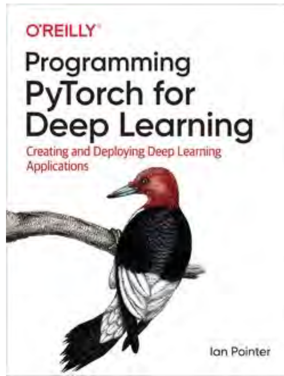

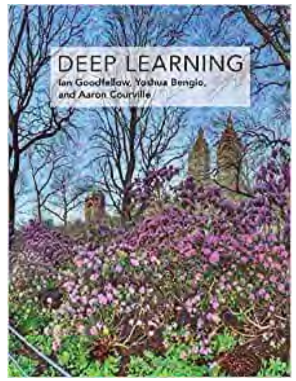

Generative Adversarial Network

Generative Adversarial Network (GAN)

- Main aim: generate fake samples from some input domain that are as close to the real data as possible. E.g., random input -> Italian Renaissance portraits

- Needs two components:

Generator

- Generates fake samples
- Tries to make the samples as 'real' as possible to fool the discriminator

Discriminator

- Identifies if a sample is fake
- Tries to identify if a sample is from the real set or a fake from the generator

Domain Adaptation

- What is the modelled distribution?

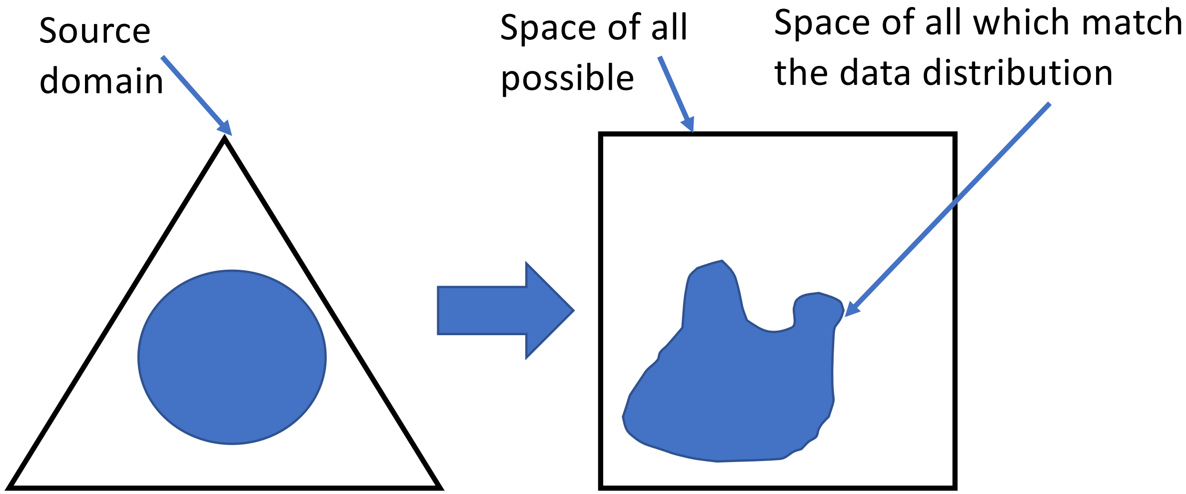

Real experiment output
Experiment Artifact
Input to GAN
Output from GAN

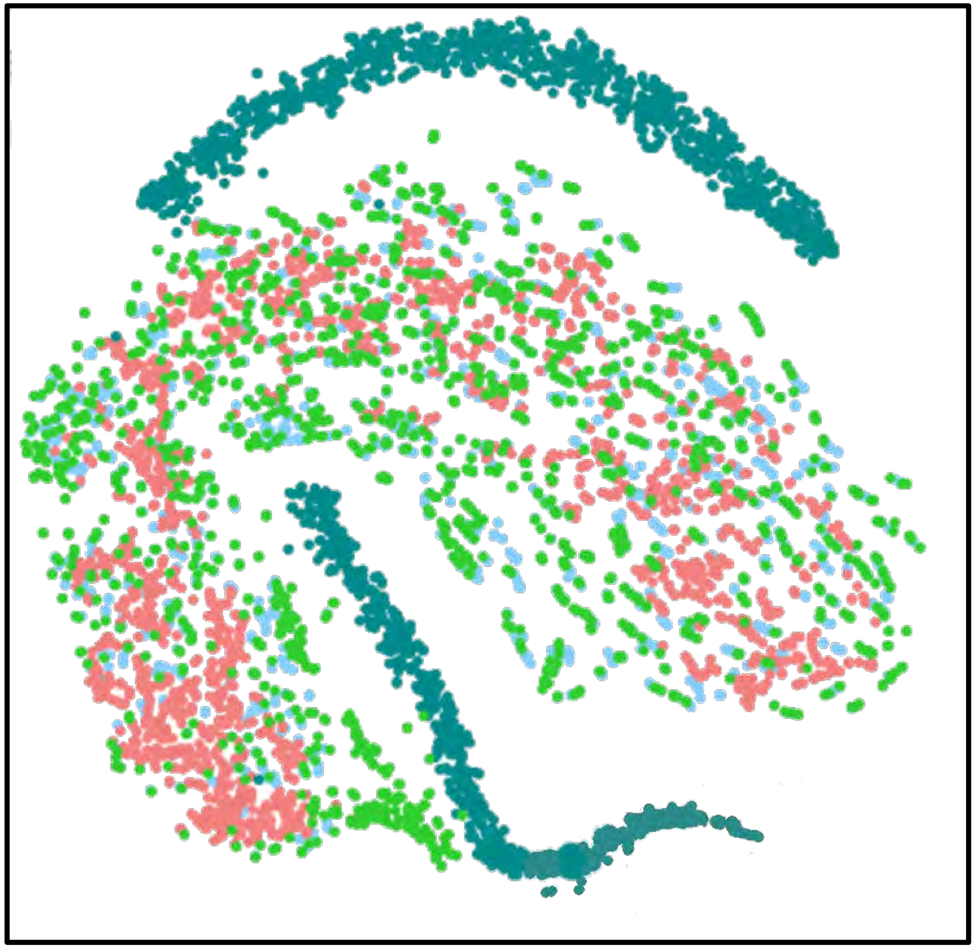

Generative Adversarial Network (GAN)

Generator

- Generates fake samples
- Forger (e.g. of art)

Discriminator

- Identifies if a sample is fake
- E.g., art critic

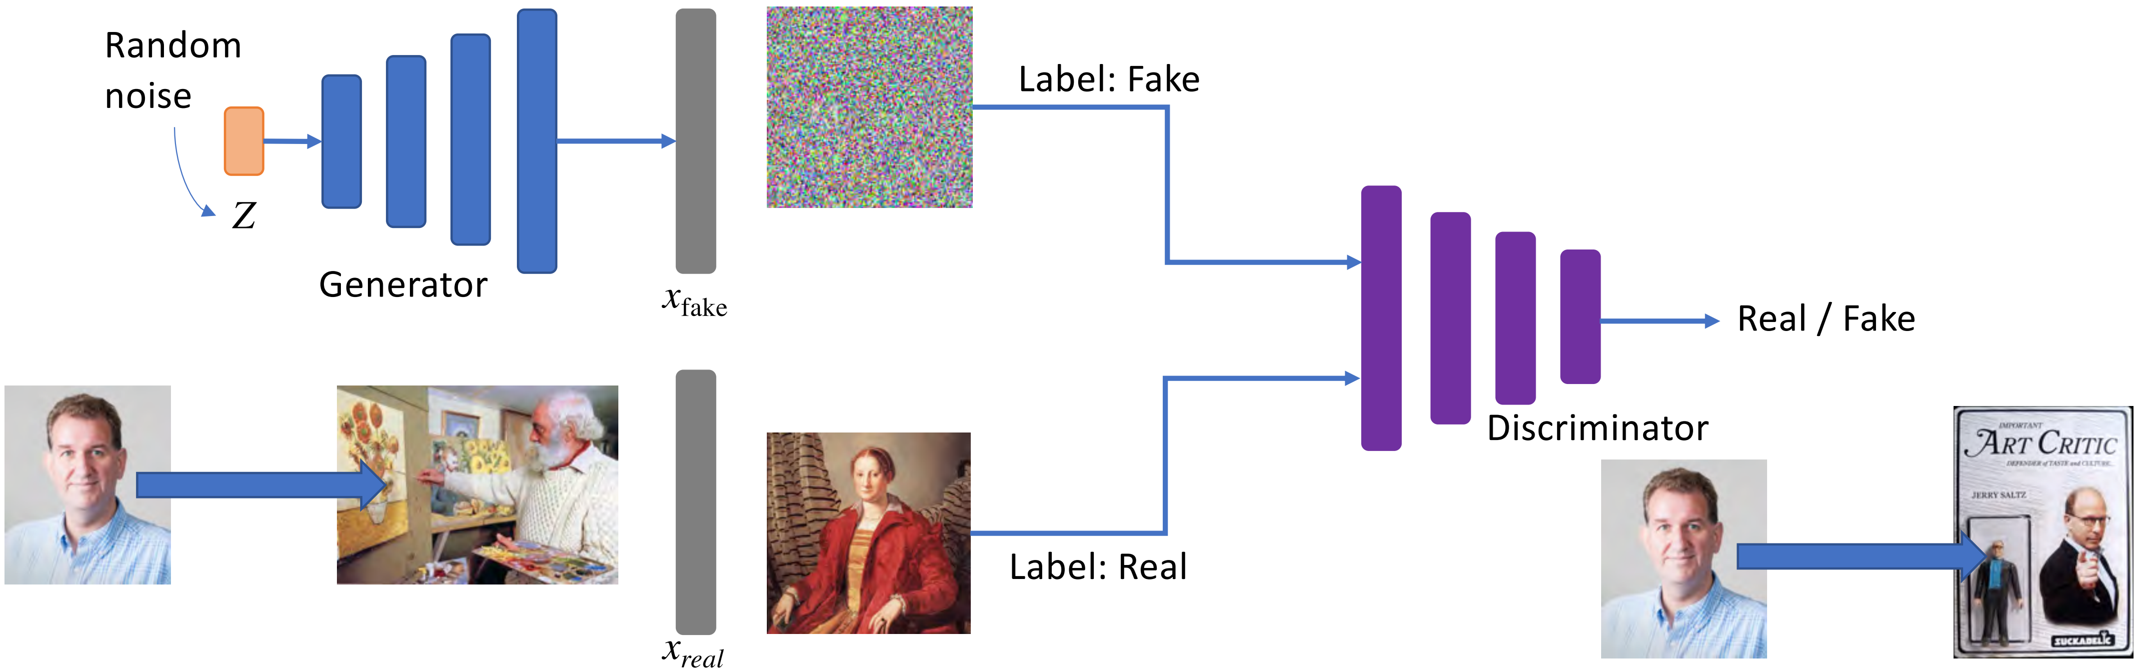

Generative Adversarial Network (GAN)

- Iteratively train **discriminator** and then generator

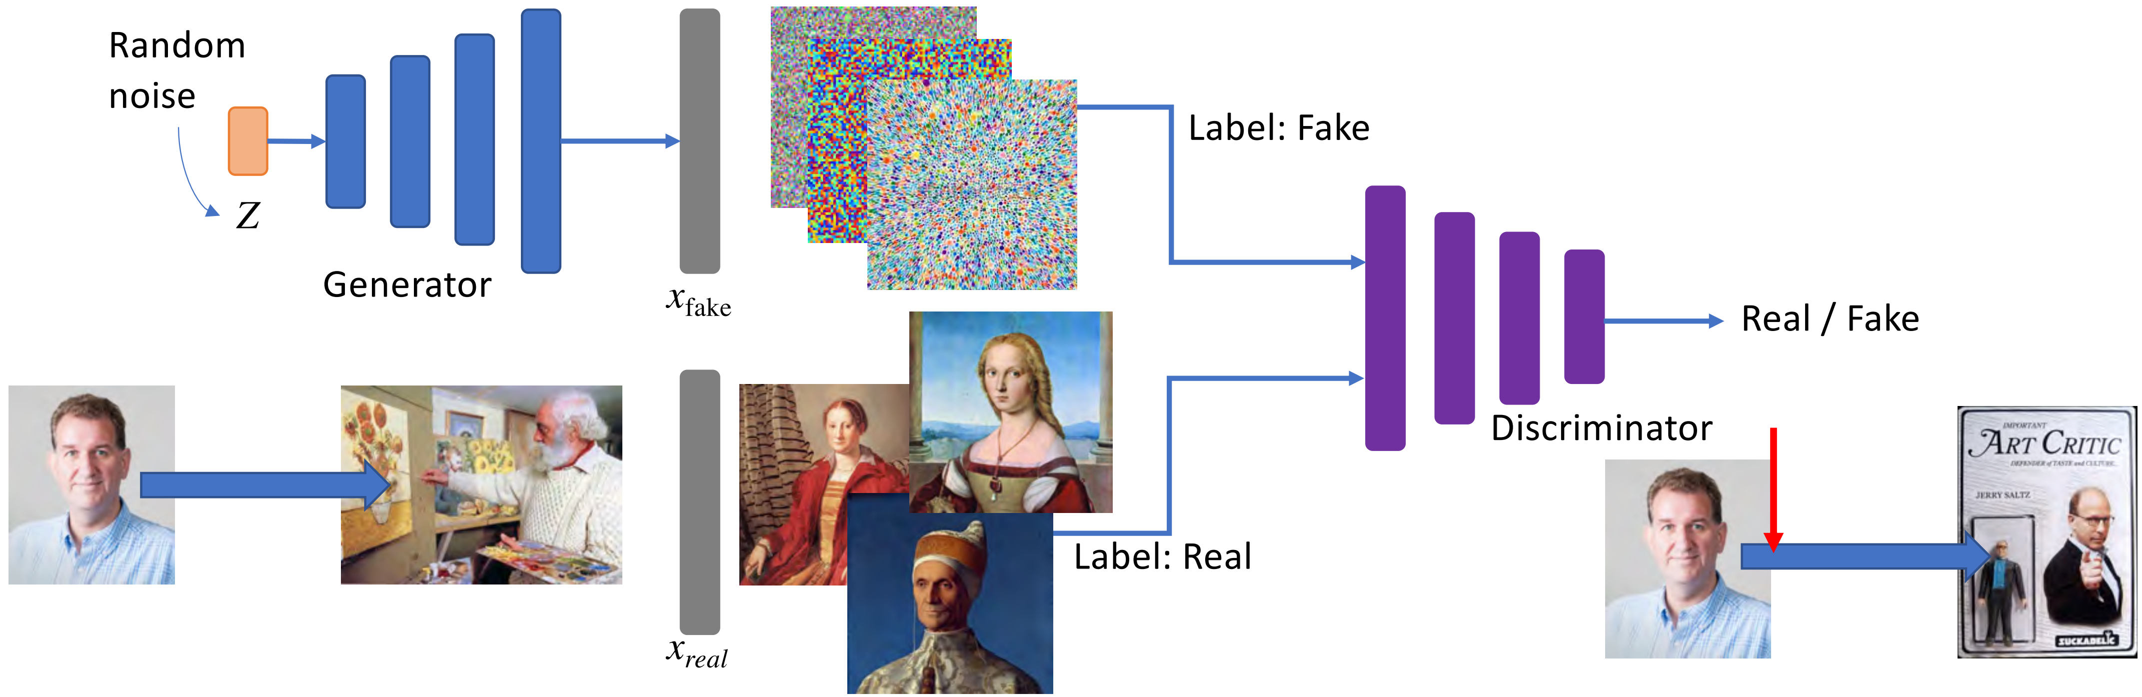

Generative Adversarial Network (GAN)

- Iteratively train discriminator and then **generator**

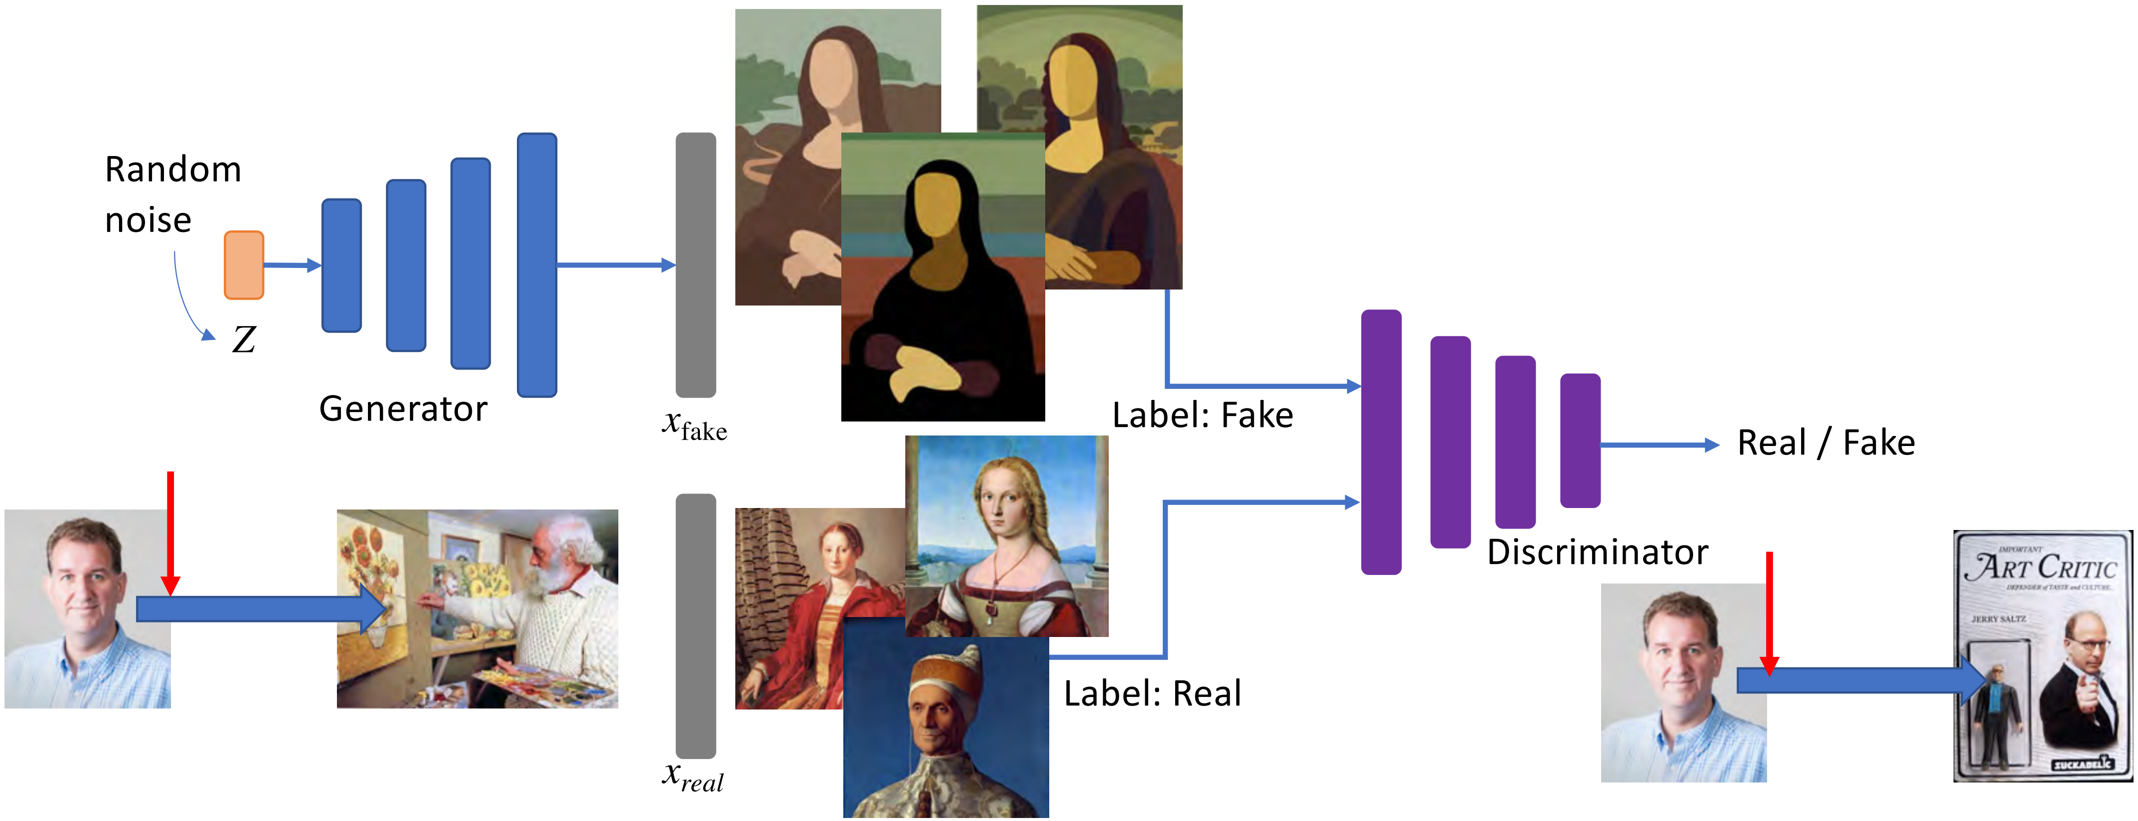

Generative Adversarial Network (GAN)

- Iteratively train **discriminator** and then generator

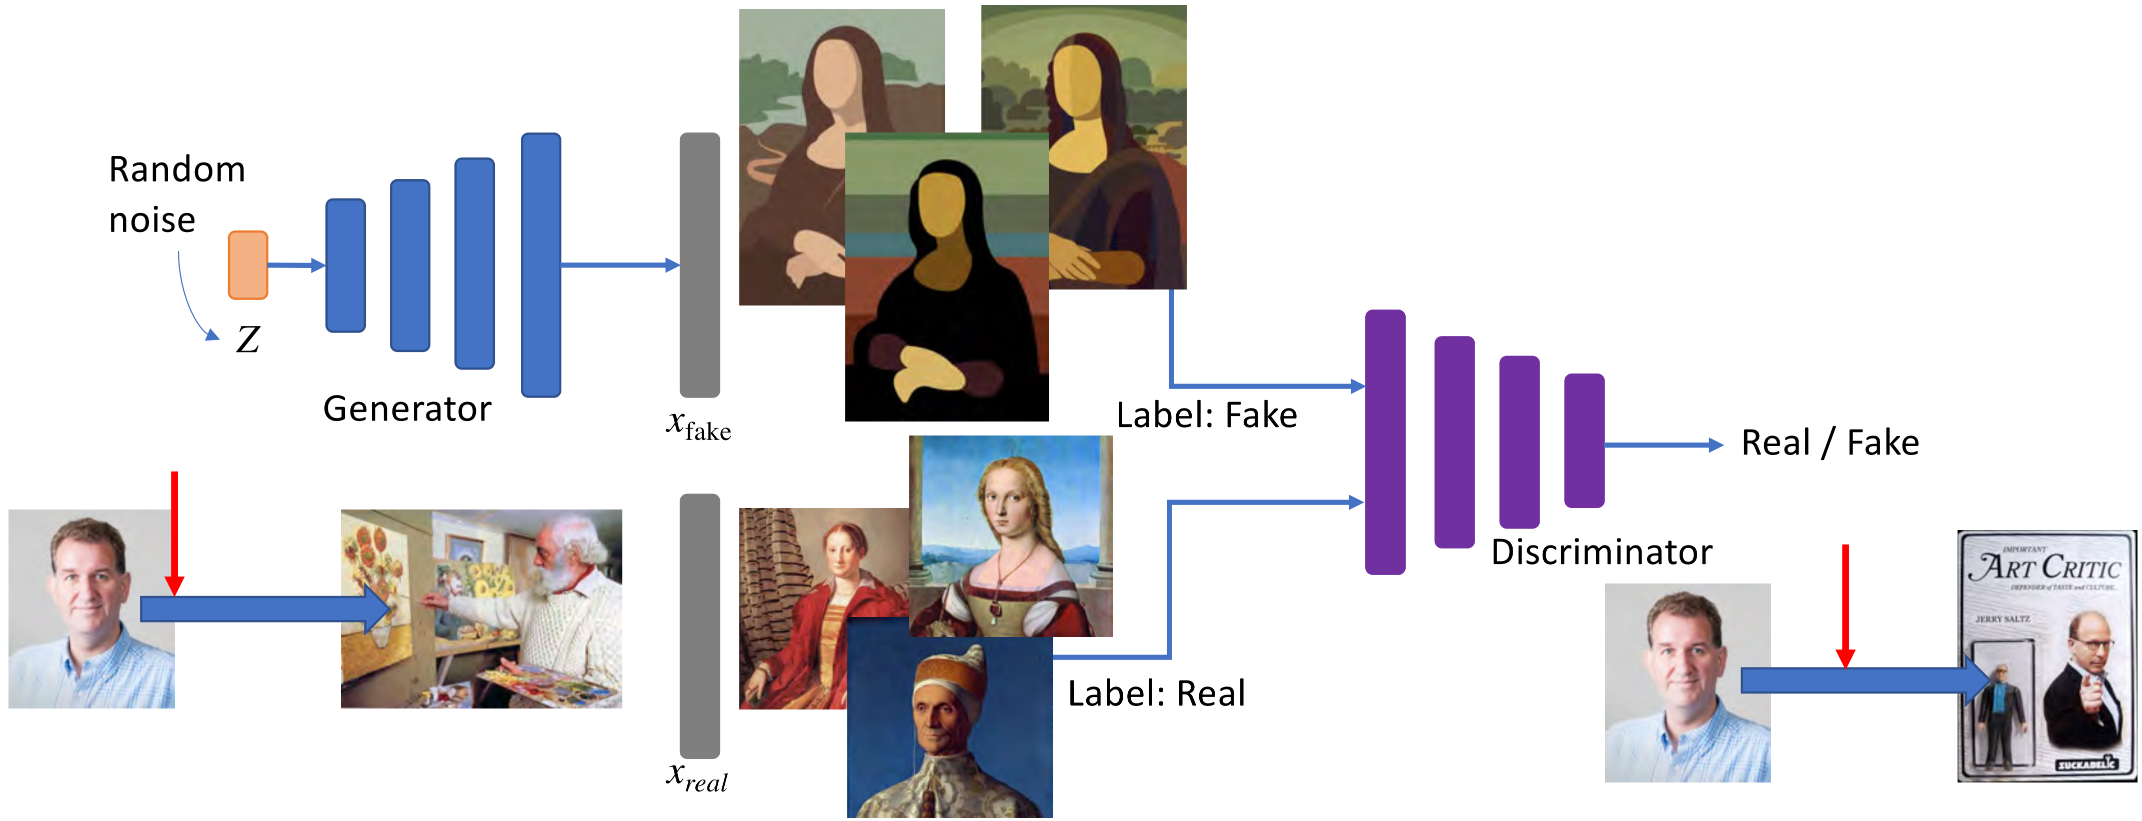

Generative Adversarial Network (GAN)

- Iteratively train discriminator and then **generator**

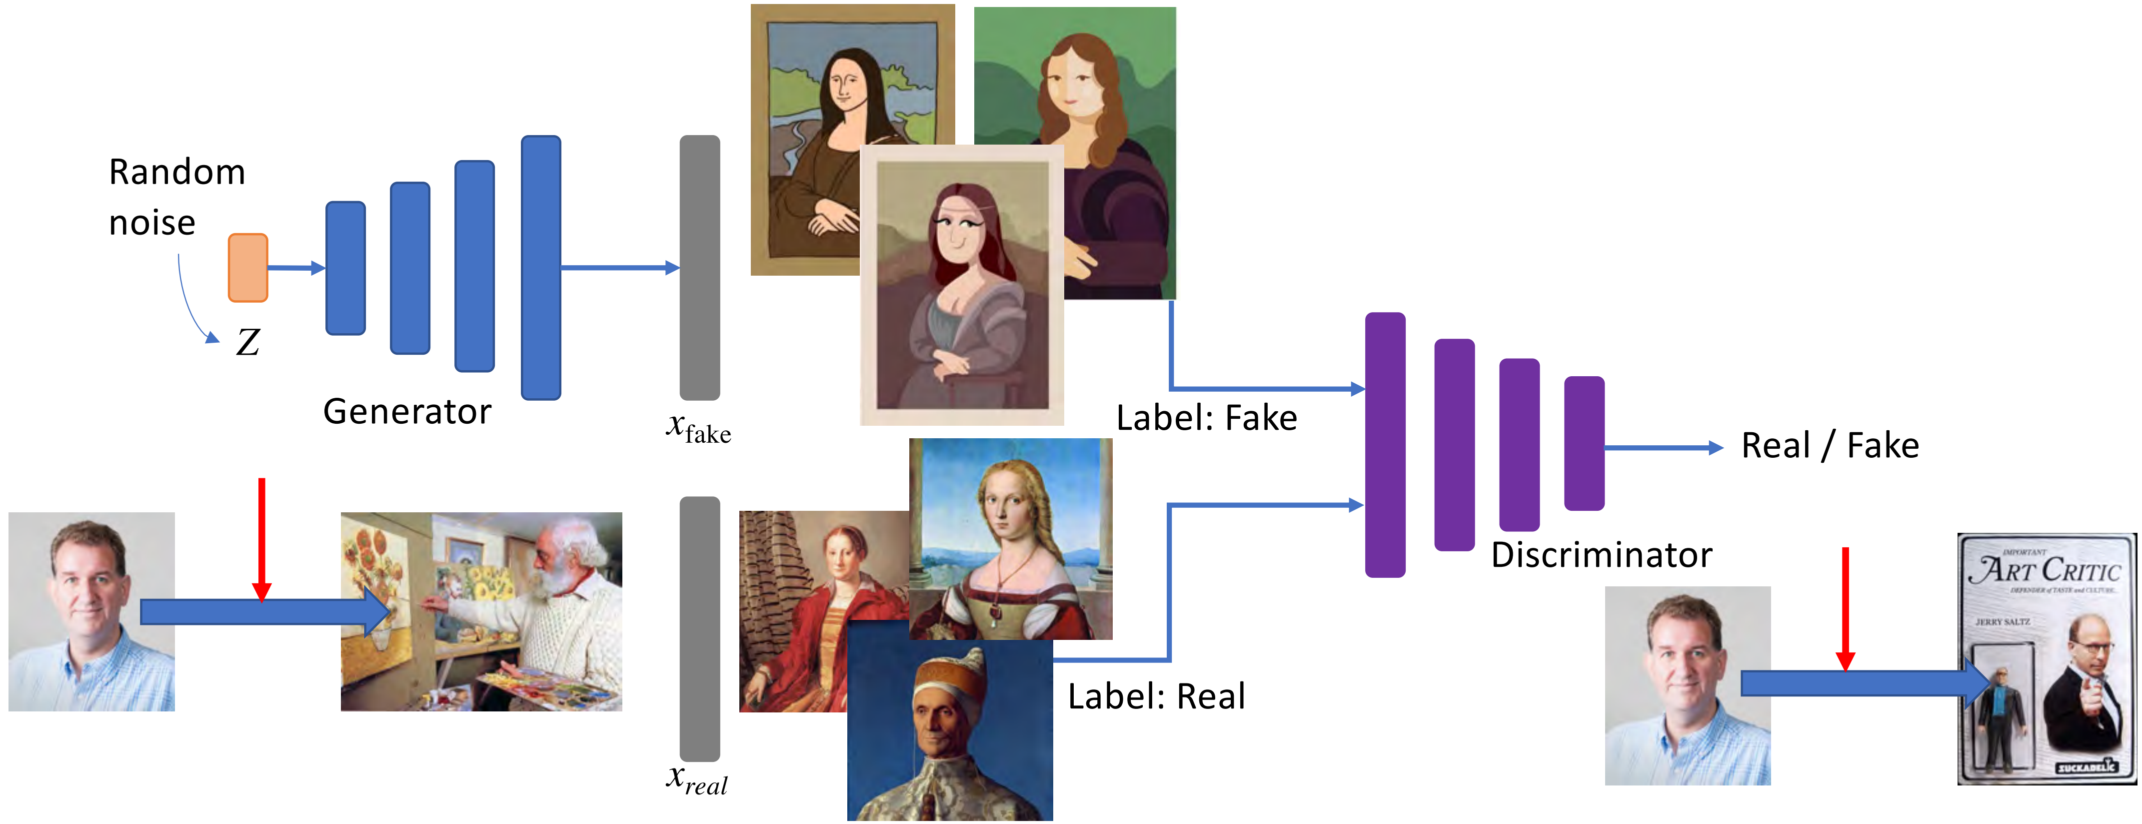

Generative Adversarial Network (GAN)

- Iteratively train **discriminator** and then generator

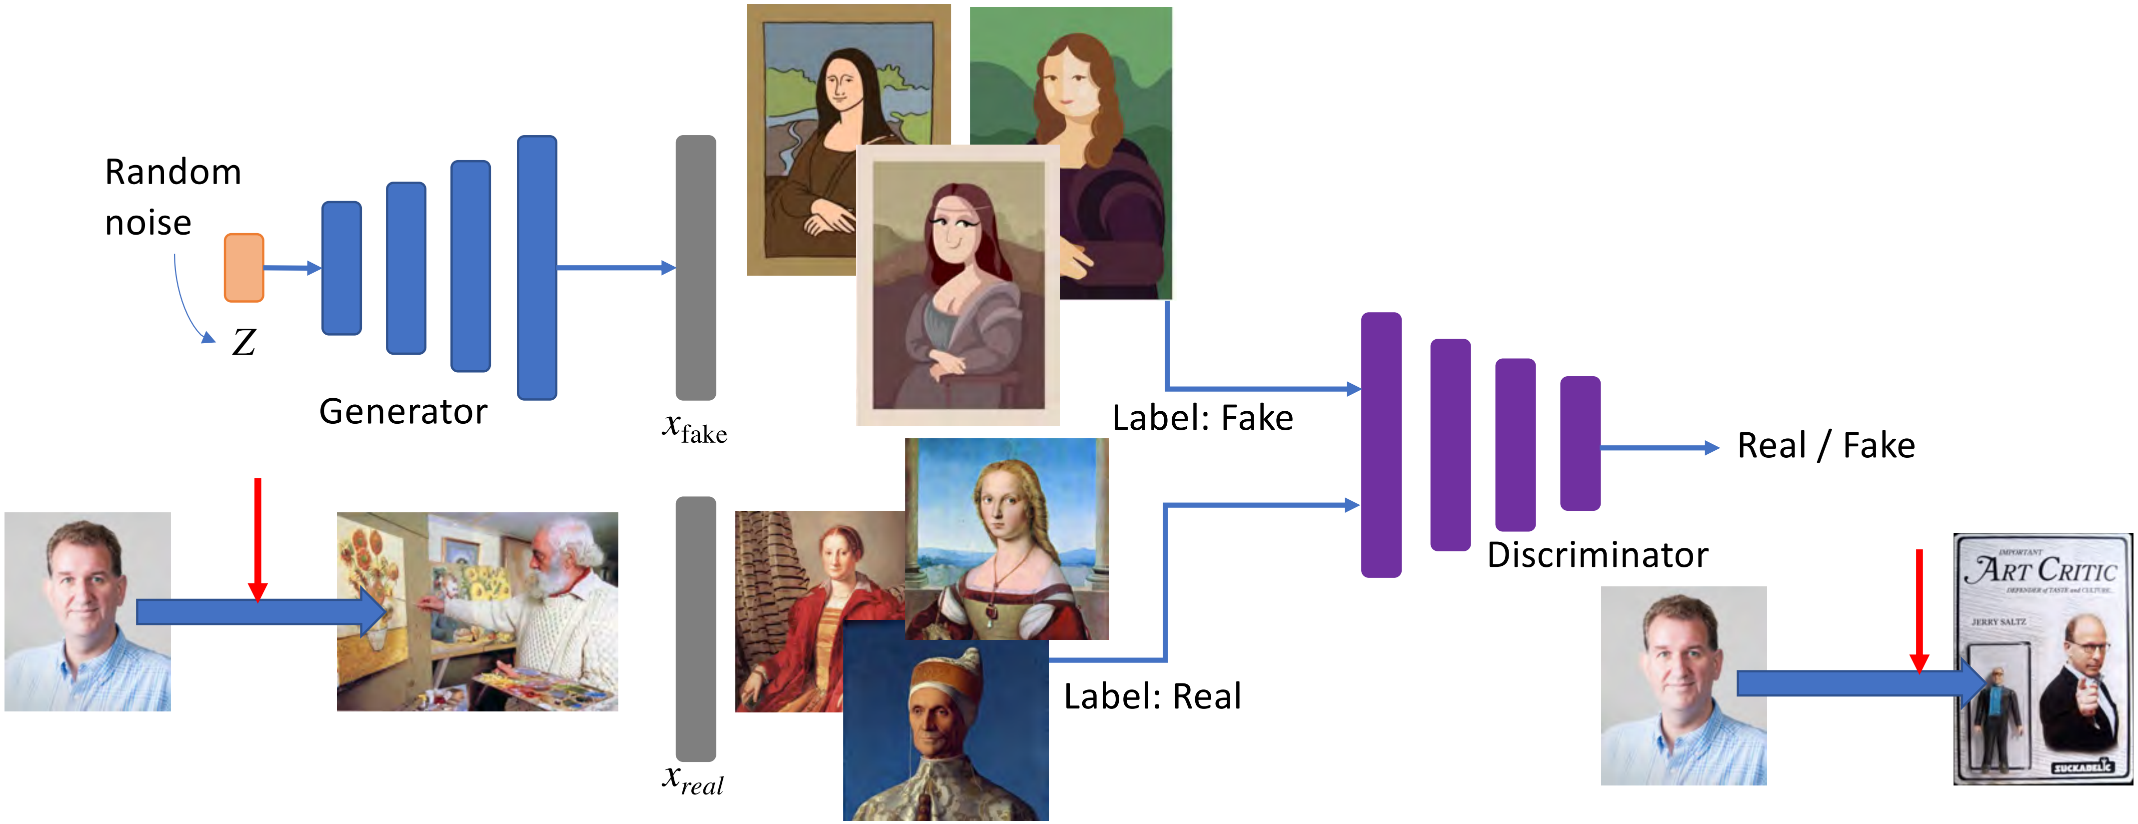

Generative Adversarial Network (GAN)

- Iteratively train discriminator and then **generator**

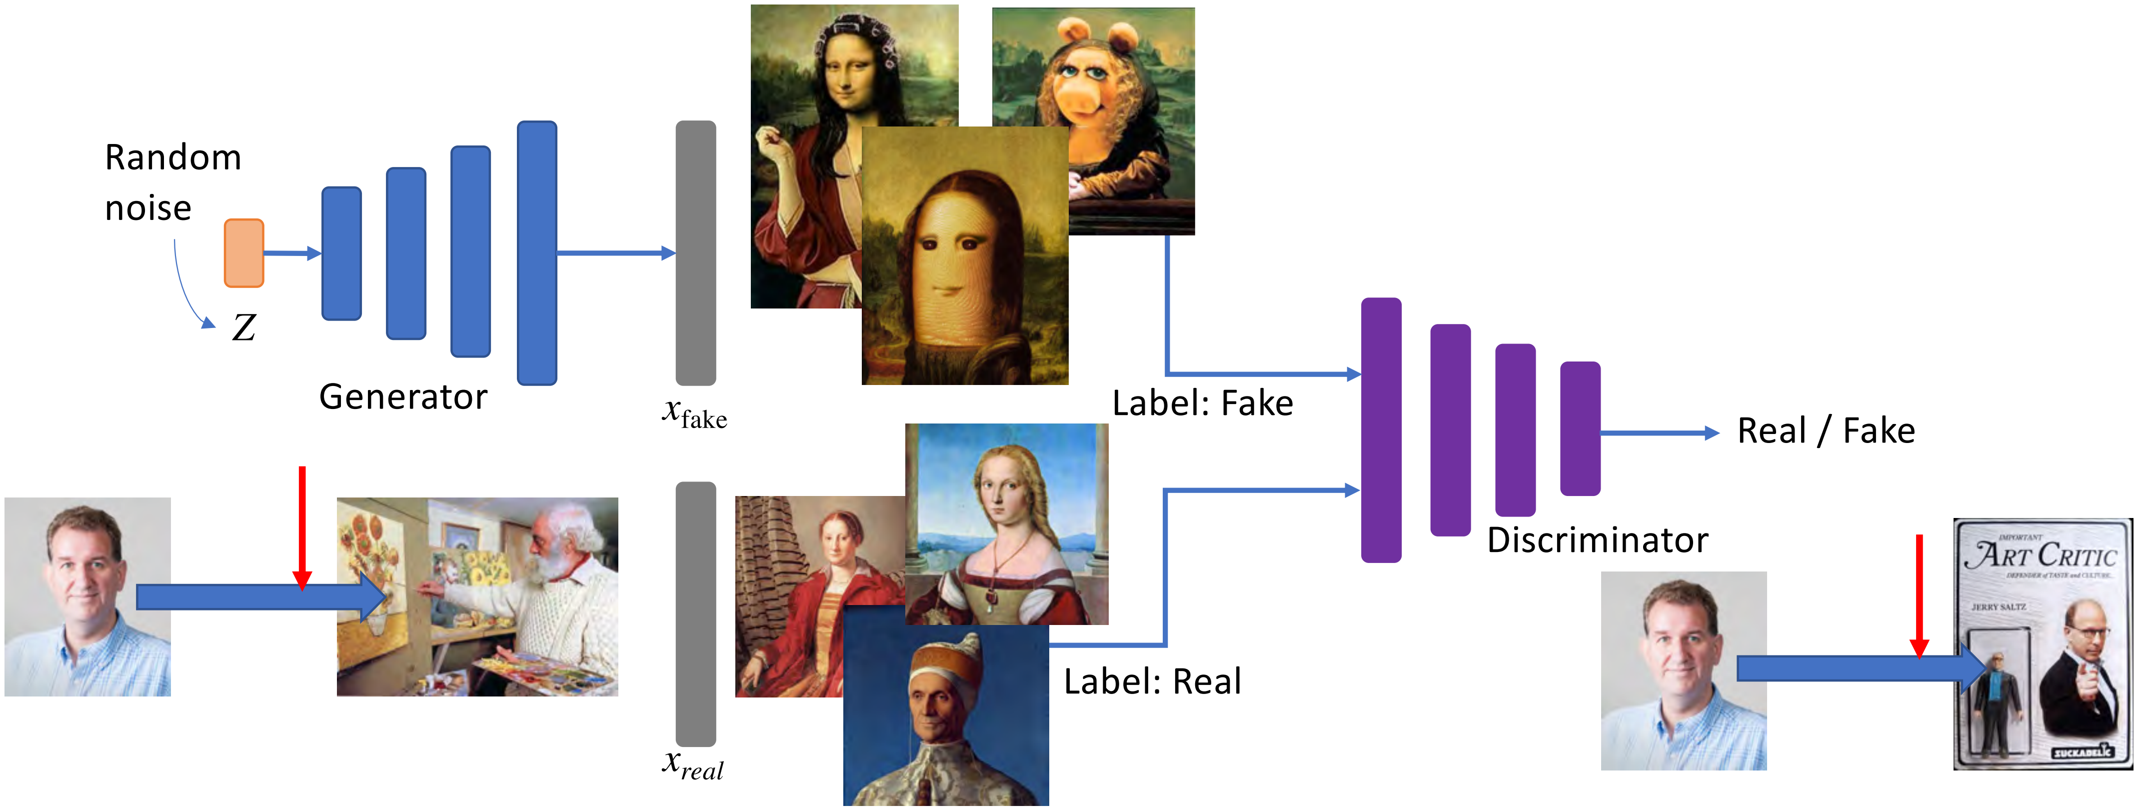

Generative Adversarial Network (GAN)

- How far can this go?

https://thispersondoesnotexist.com

Further Reading

- Original GAN paper
  - https://arxiv.org/abs/1406.2661
- Papers with Code
  - https://paperswithcode.com/task/image-generation
- A Gentle Introduction to Generative Adversarial Networks (GANs)
  - https://machinelearningmastery.com/what-are-generative-adversarial-networks-gans/

Generating synthetic data

Traffic data

Automatic Number Plate Recognition (ANPR)

- Used by most cities for traffic management
- Huge volumes of data
- License plate, camera, time
- Data could be used for more
- But can't release data!

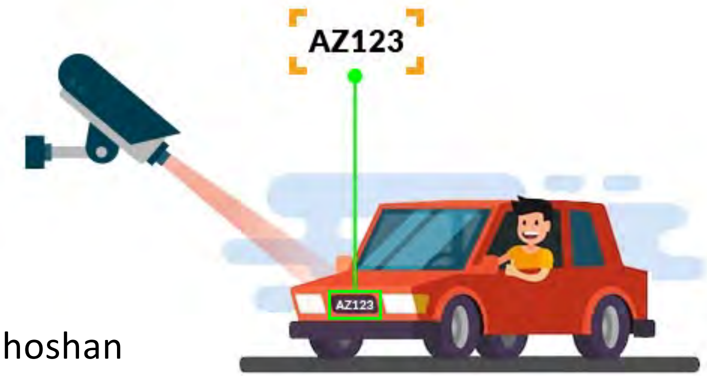

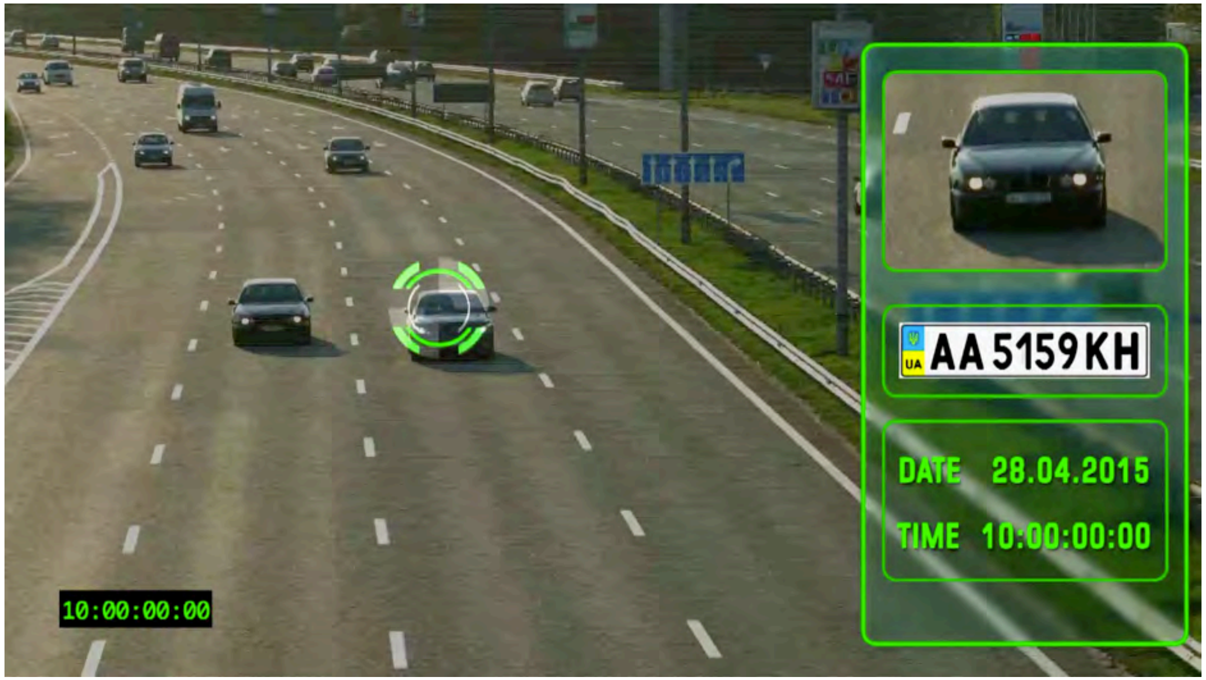

Areeb Alshoshan

https://www.theproche.com/2020/08/19/what-is-anpr/

https://www.youtube.com/watch?v=Et4x8bdpSqc

GANs

For generating license plate, camera and time-stamp

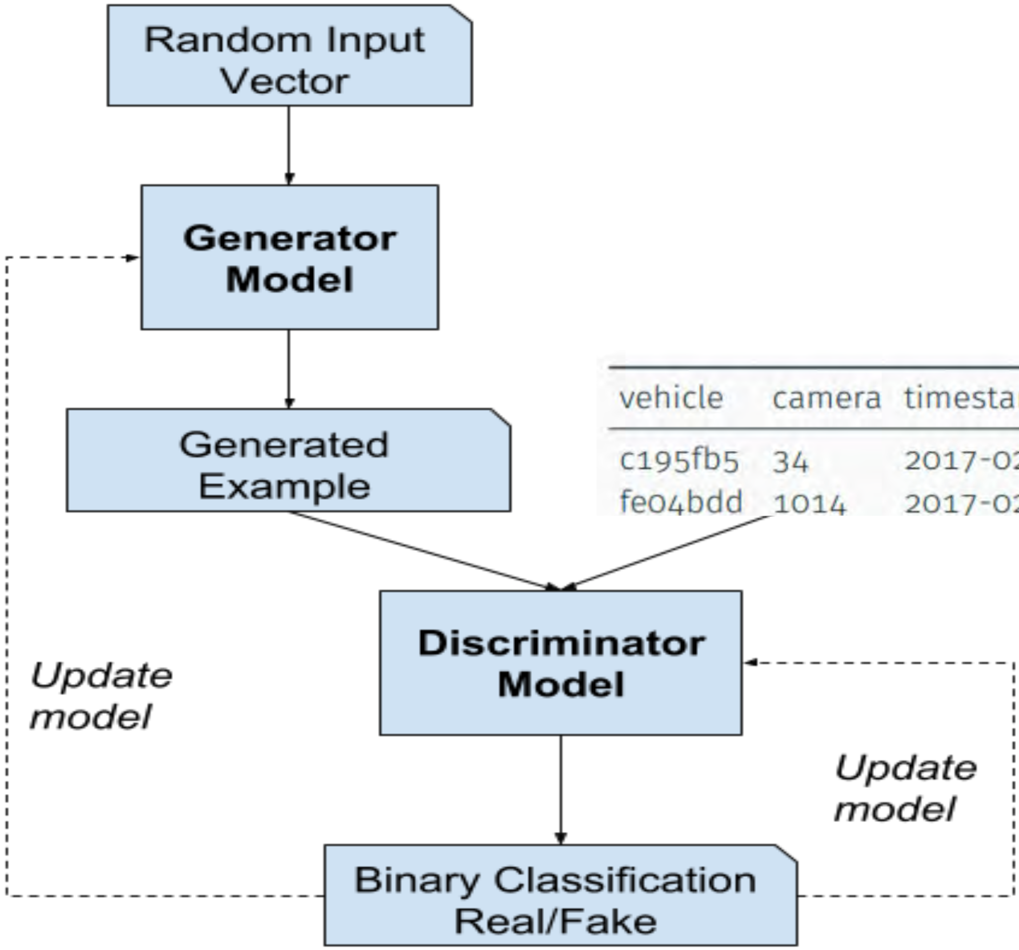

Figure from : http://machinelearningmastery.com/

Digital twin and GANs model

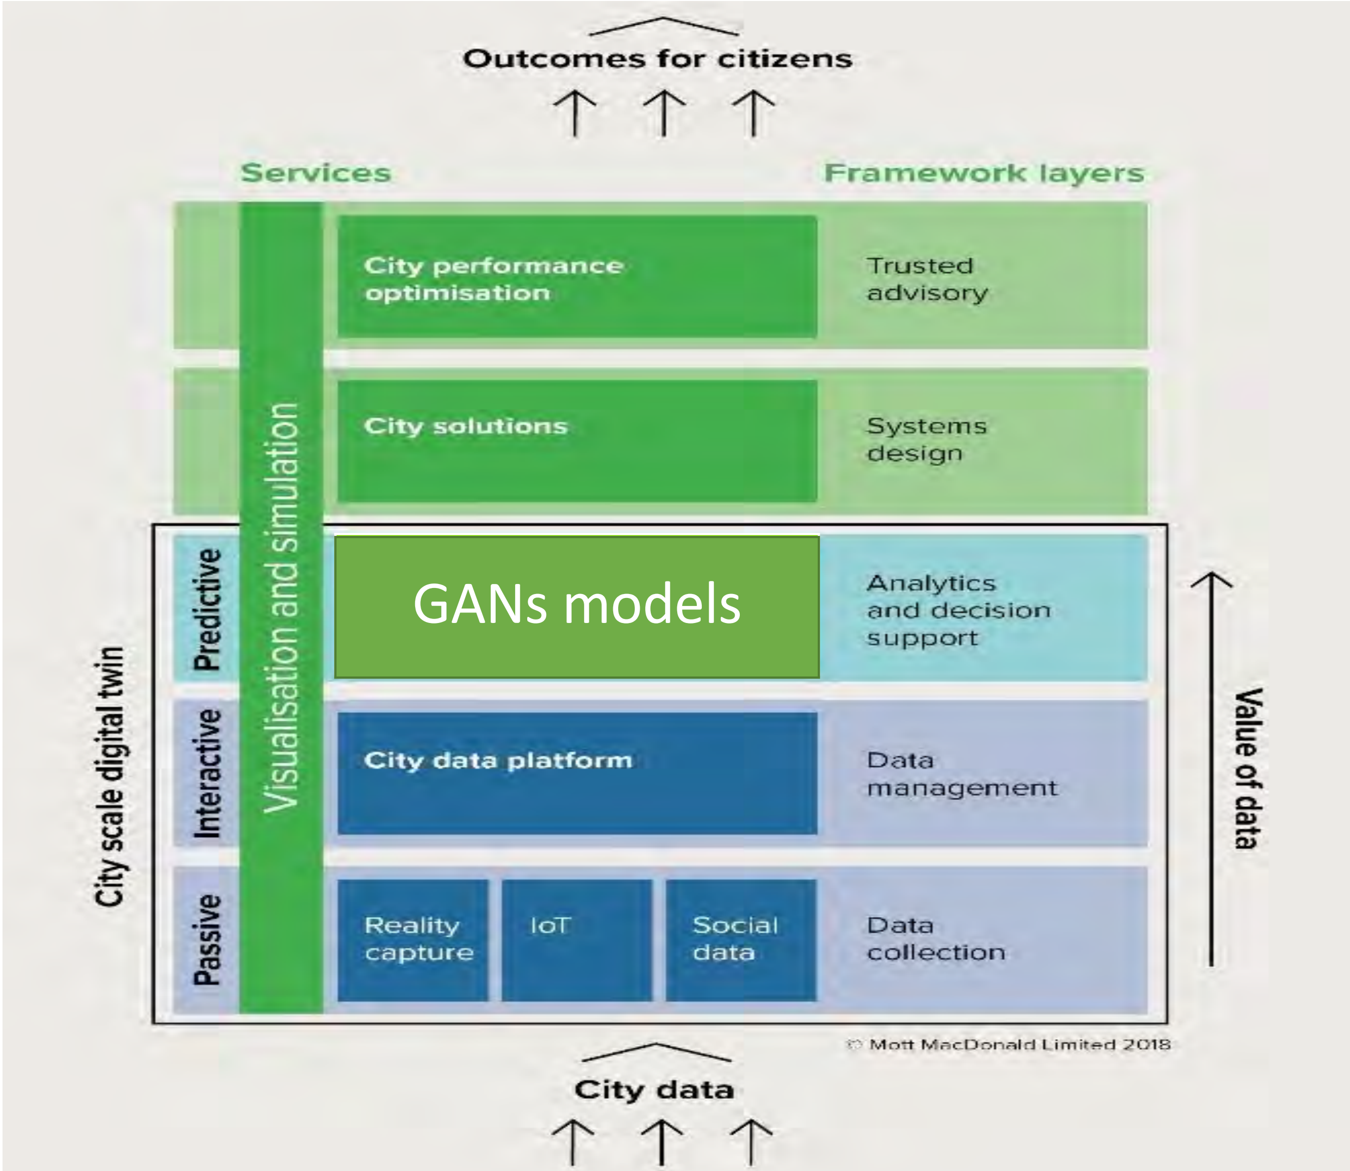

Areeb Alshoshan

Transforming data through GANs

Conditional GANs

- Takes data from one domain and maps it to data in a different domain
- E.g., satellite → map

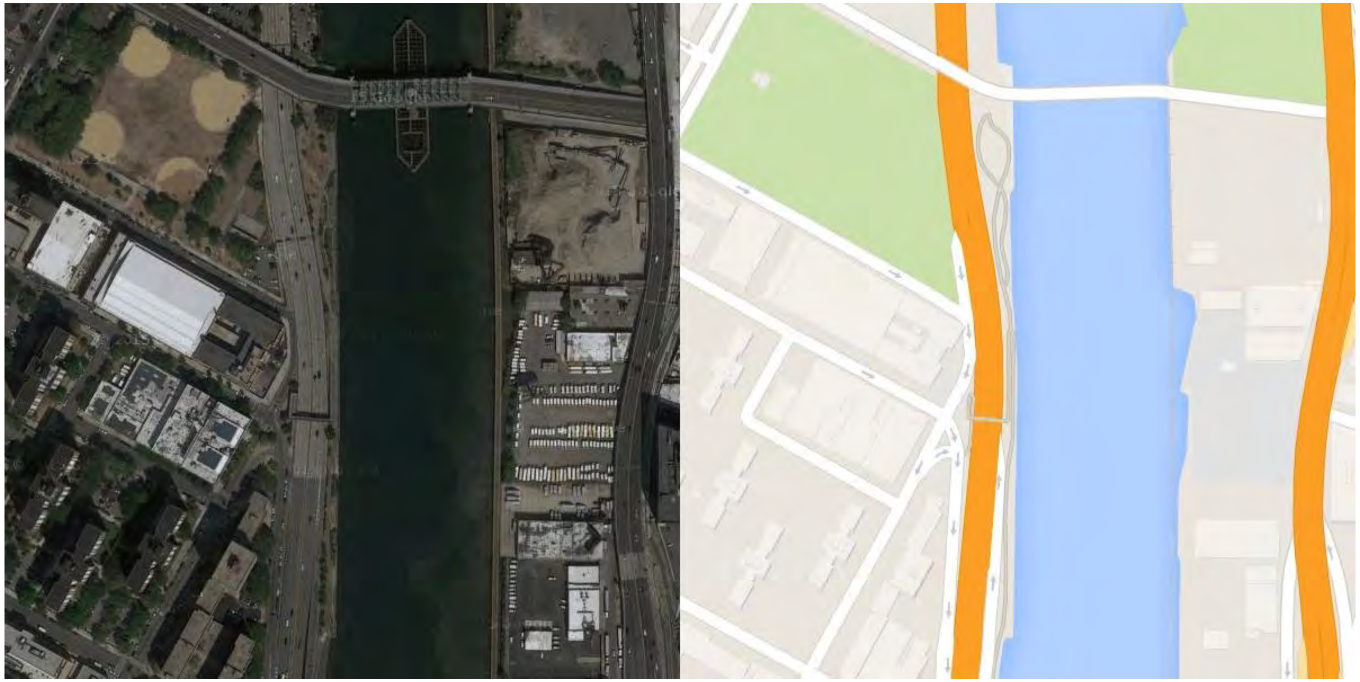

Conditional GAN

- Pix2Pix

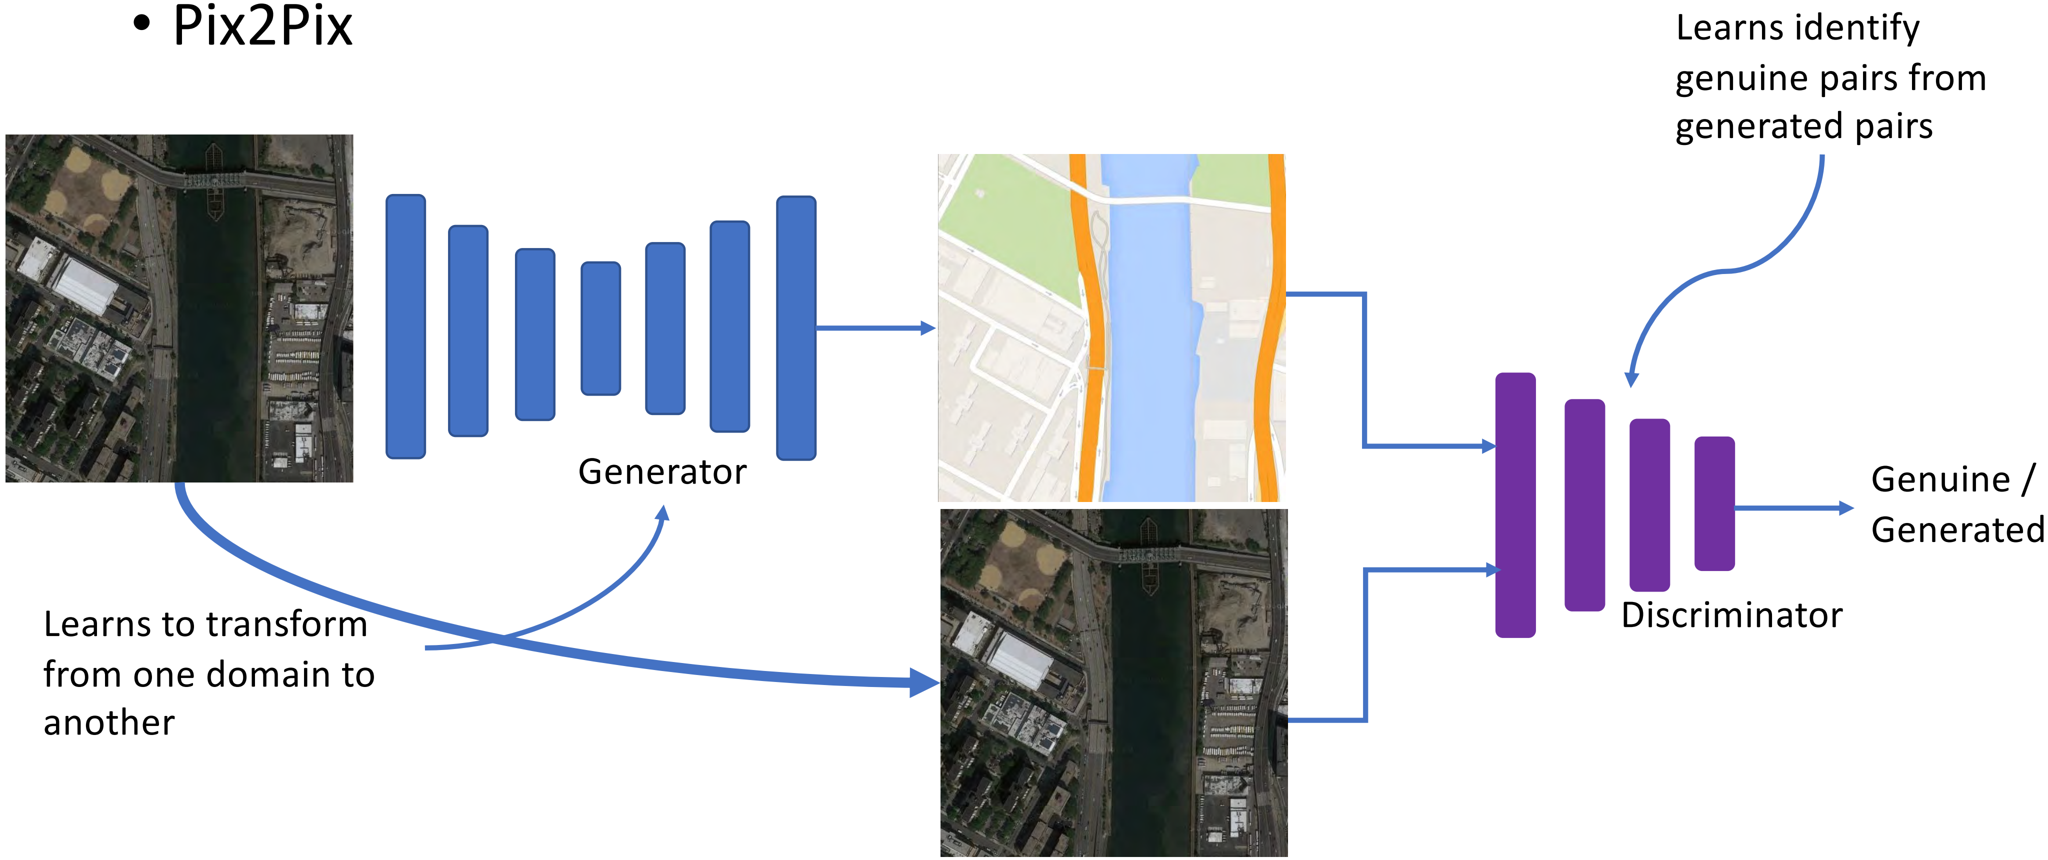

Further Reading

- Pix2Pix developer's page
  - https://phillipi.github.io/pix2pix/

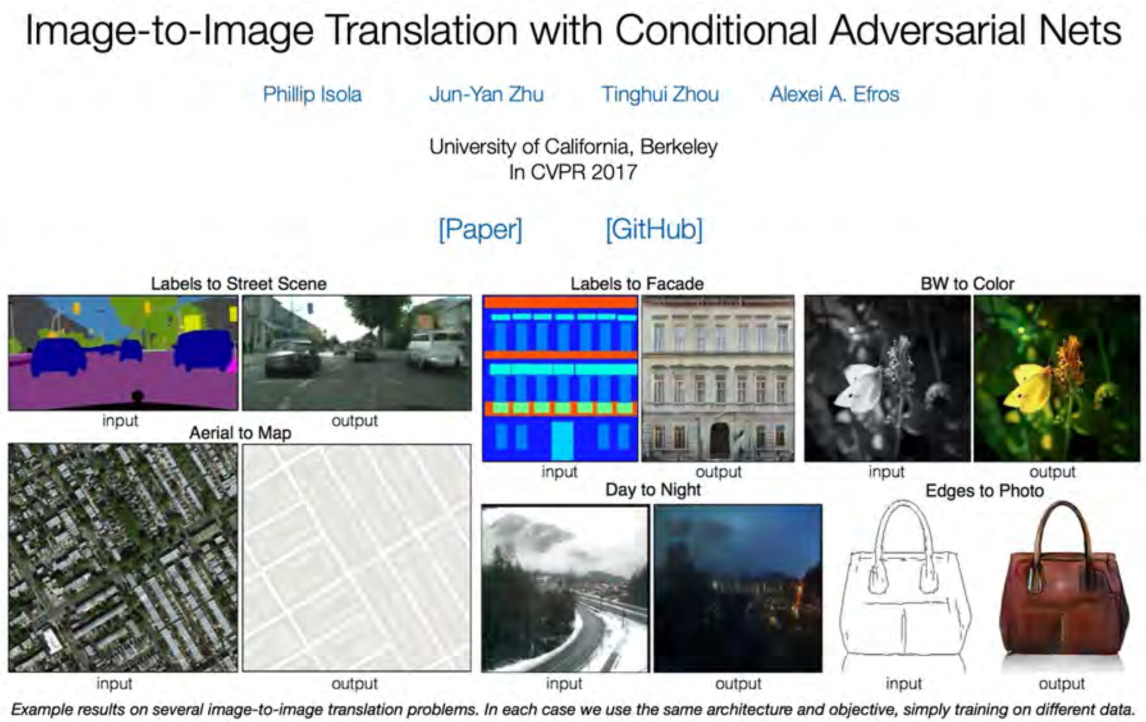

Image Segmentation through GANs

Identifying cells in images

- Given an image of a set of cells identify each cell in the image

Cell image Ground truth labels Generated Segmentation

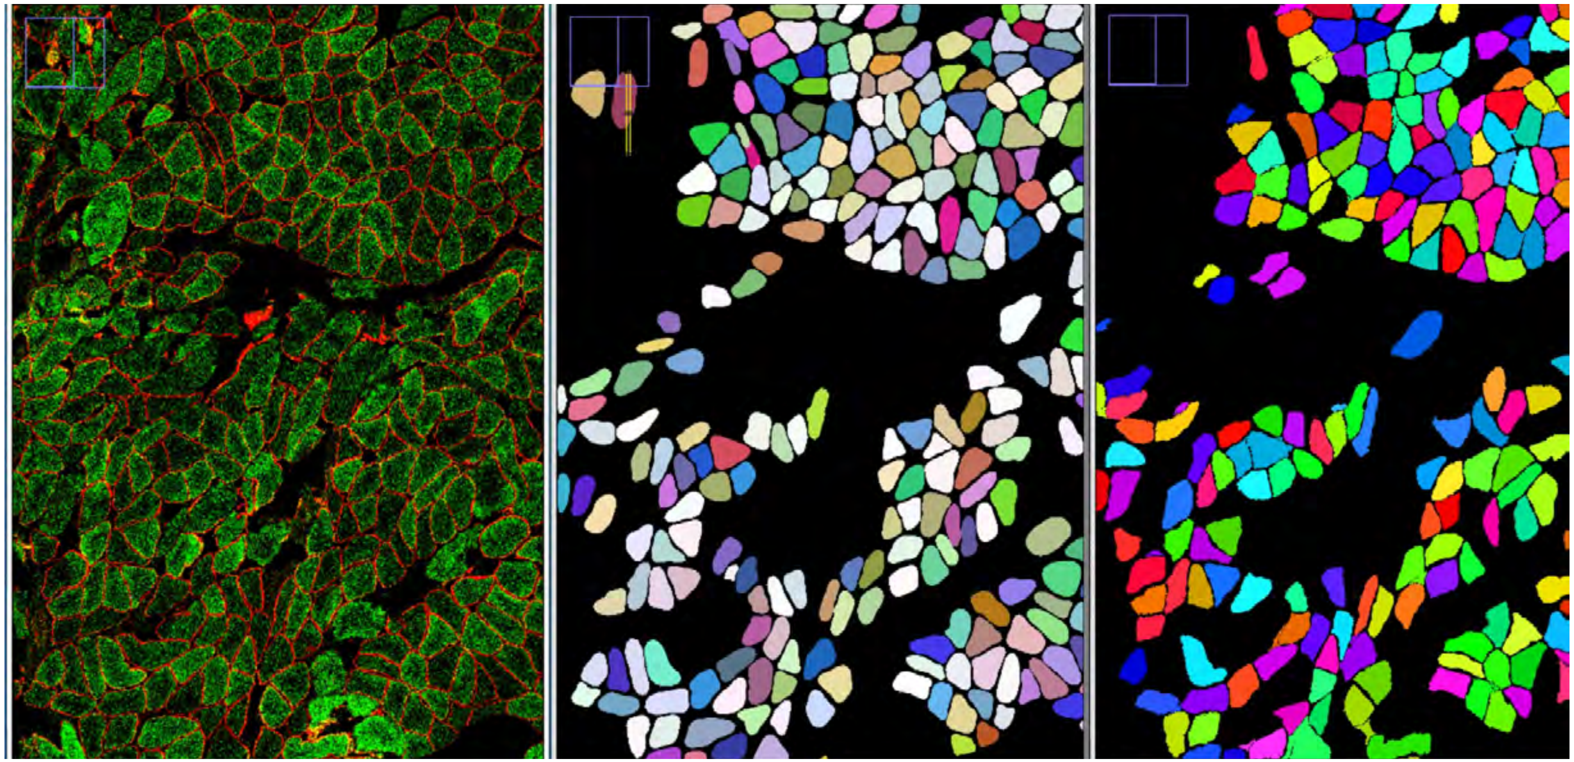

Atif Khan

Segmenting dolphins

- Identify individual dolphins from photographs
- Helps in determining population sizes

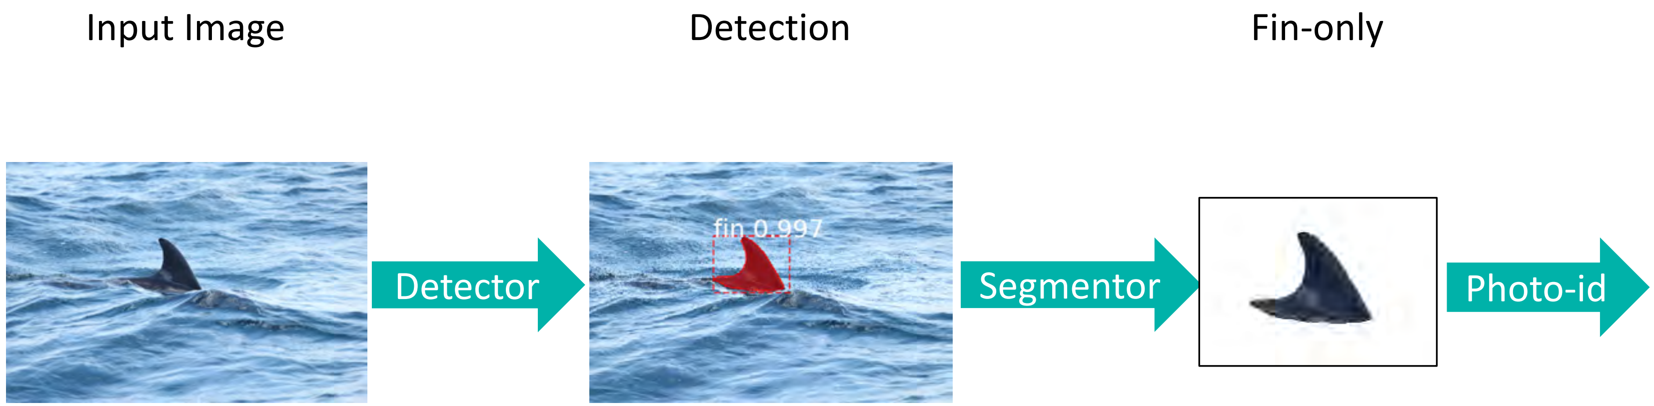

ID	Confidence
Per	90.2
Steve	4.7
Nick	3.1
Georgia	1.5
Matt	0.5

Cameron Trotter

Conditional GAN

Segmenting Dolphins

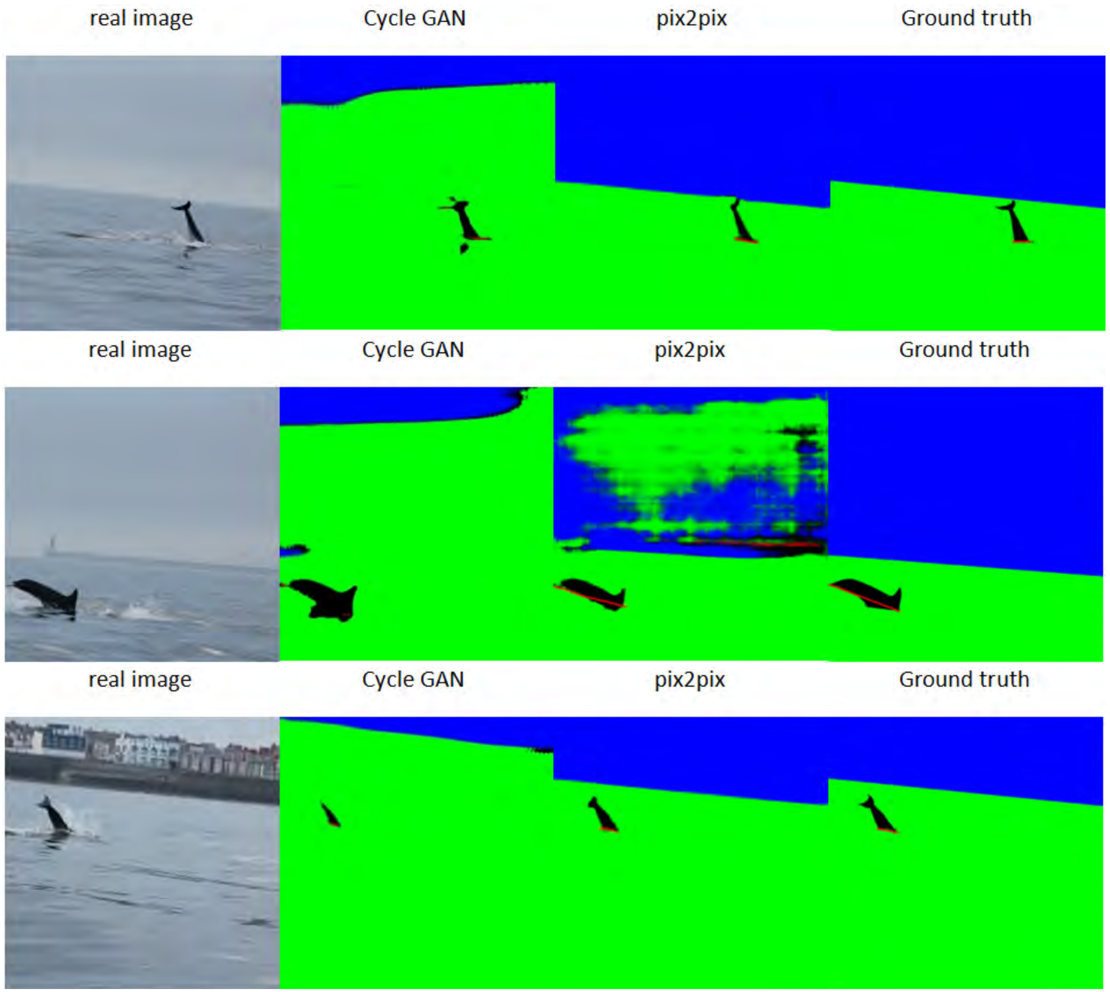

Zhenwen Luo

Style Transfer

Style Transfer Generative Adversarial Network CycleGAN

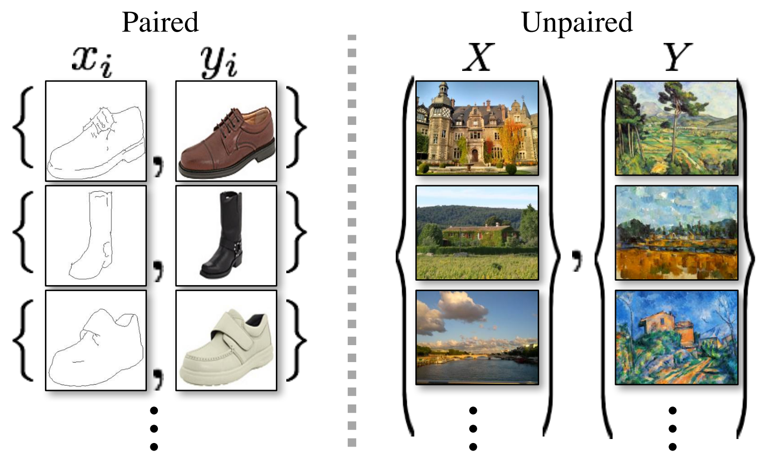

- Overcomes problem of needing paired data
- X → Y is a GAN, Y → X is a second GAN
- Map from domain X→ Y → X
- Look at how close points are in X

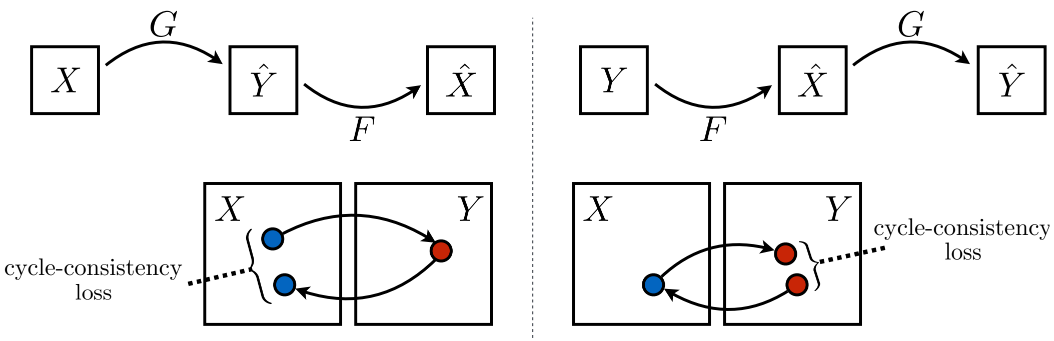

Further Reading

- CycleGAN developer's page
  - https://junyanz.github.io/CycleGAN/

Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks

Jun-Yan Zhu* Taesung Park* Phillip Isola Alexei A. Efros

UC Berkeley

In ICCV 2017

Paper | PyTorch code | Torch code

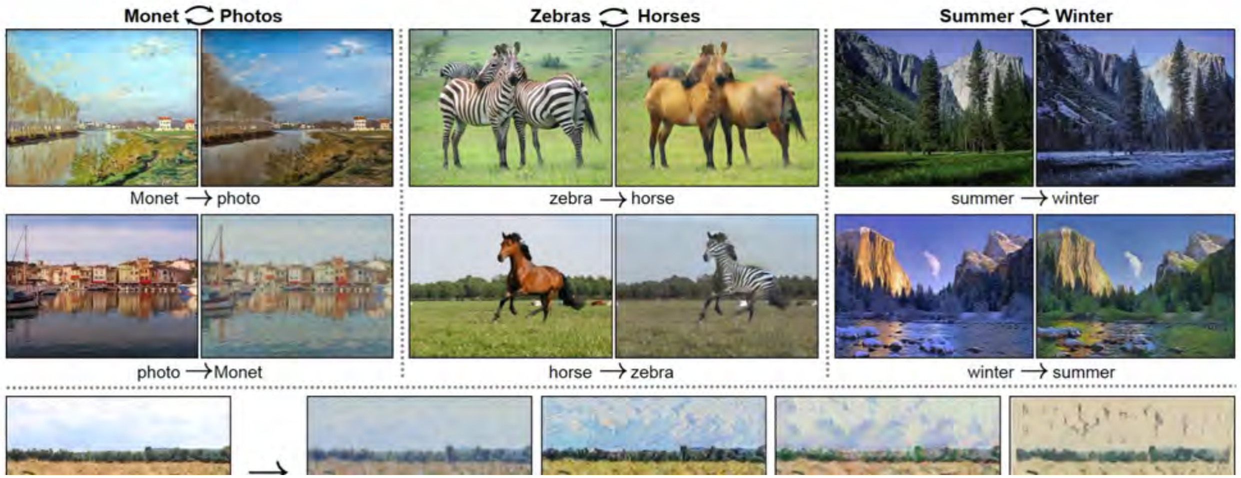

Tracking Larvae

Experiments and Simulations

Style Transfer Generative Adversarial Network
CycleGAN

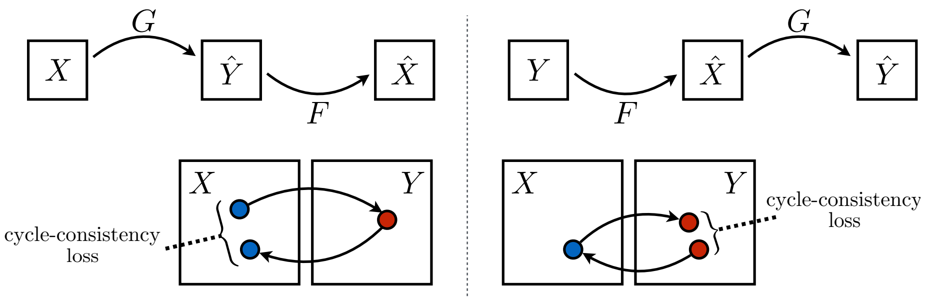

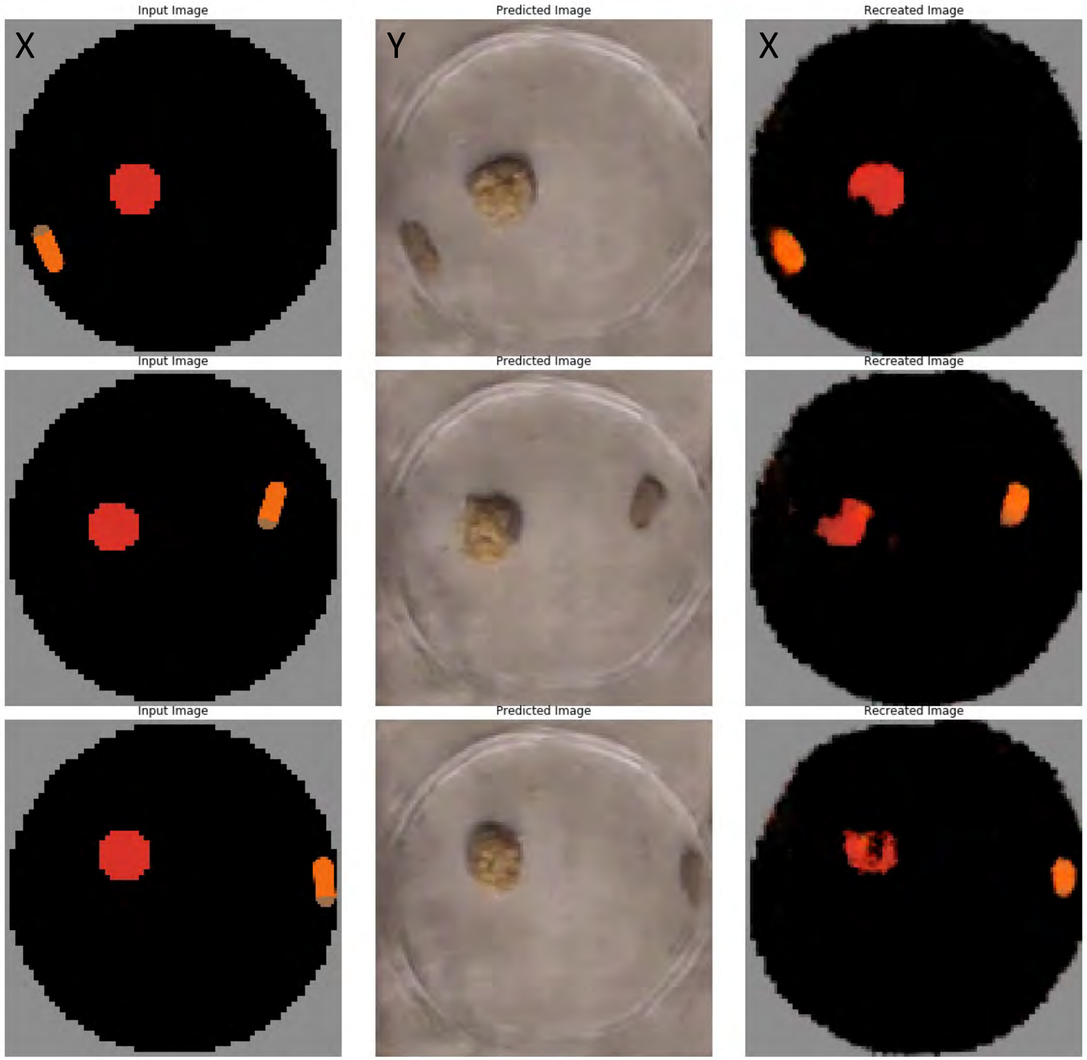

Style transfer for Larvae

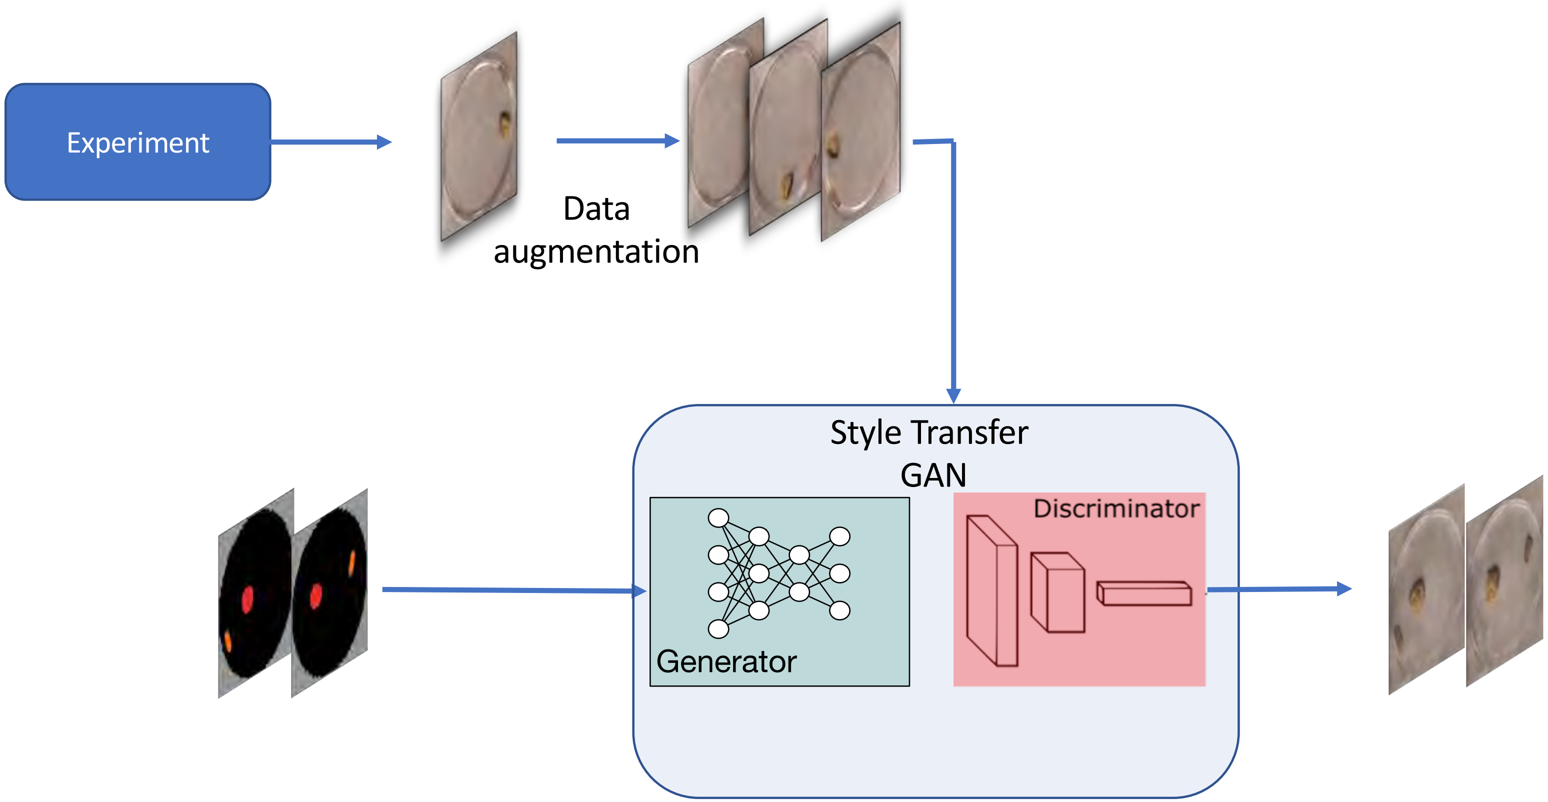

Data augmentation

- There’s not enough larva videos to train with – so create more…

Using CycleGANs : Tracking Larvae

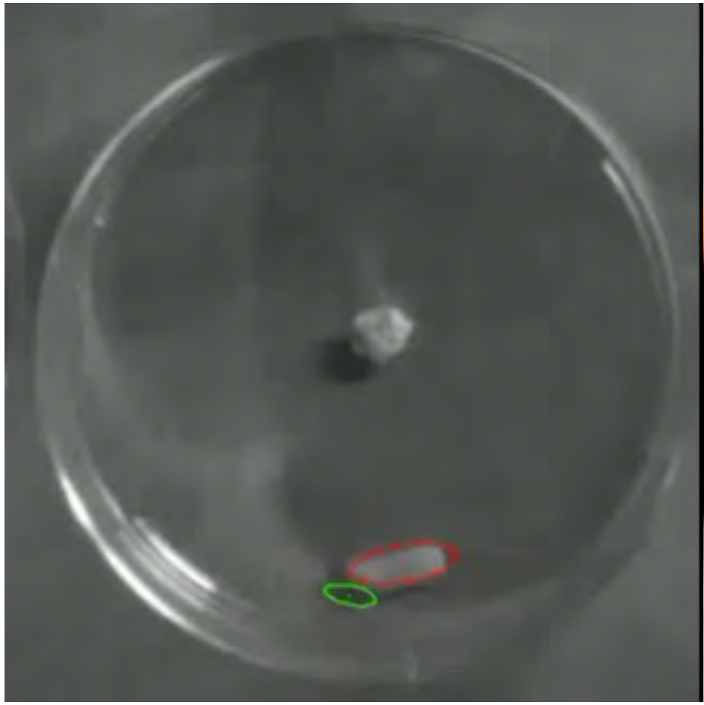

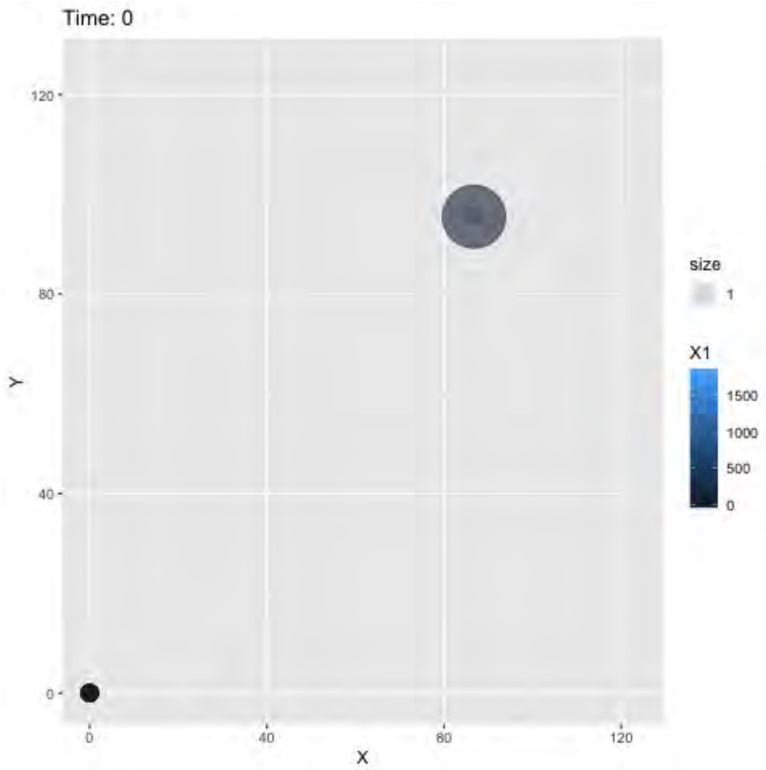

Other things you can do

- Arithmetic operations on data
  - https://arxiv.org/abs/1511.06434

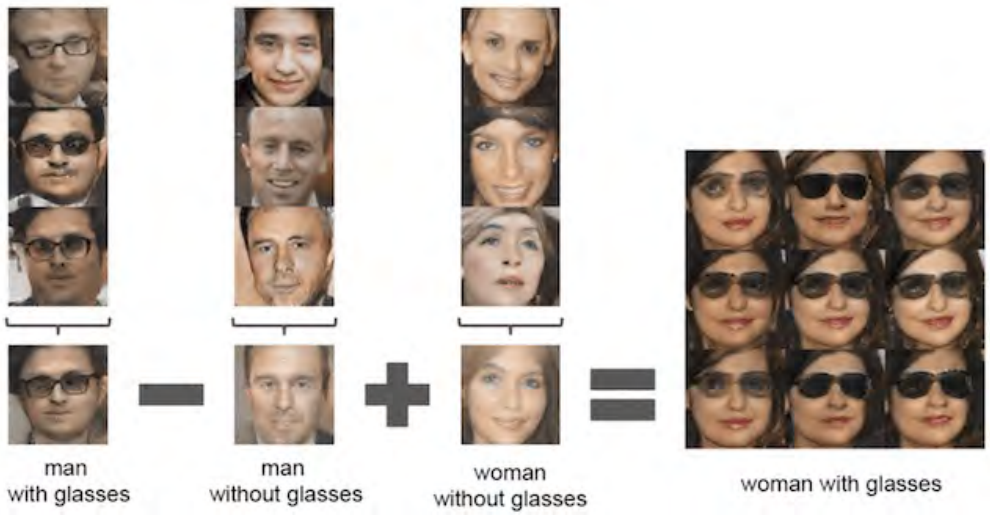

Other things you can do

- Arithmetic operations on data
  - https://arxiv.org/abs/1511.06434
- Manipulate images
  - https://arxiv.org/abs/1611.06355

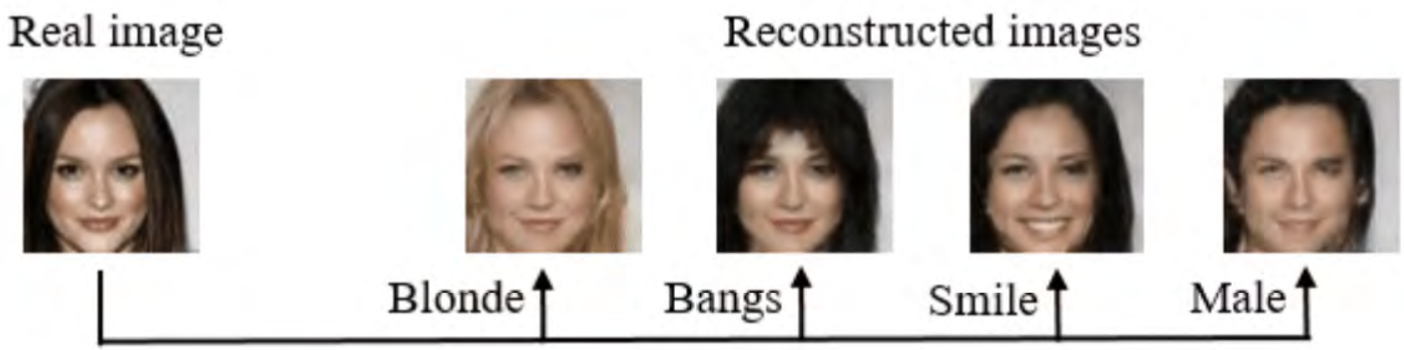

Other things you can do

- Arithmetic operations on data
  - https://arxiv.org/abs/1511.0643
- Manipulate images
  - https://arxiv.org/abs/1611.0635
- Image super-resolution
  - https://arxiv.org/abs/1609.04802

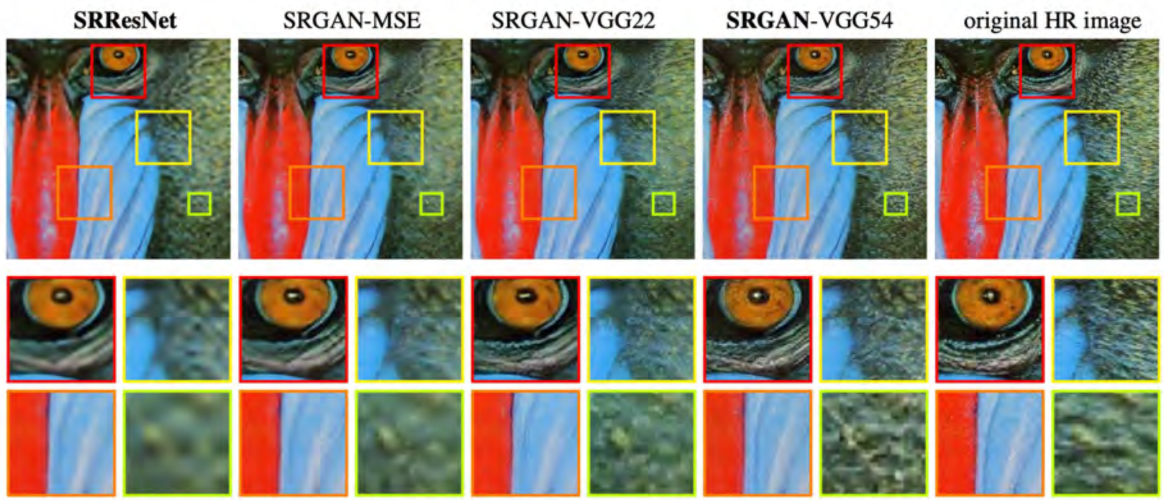

Other things you can

- Arithmetic operations on data
  - https://arxiv.org/abs/1511.0643
- Manipulate images
  - https://arxiv.org/abs/1611.0635
- Image super-resolution
  - https://arxiv.org/abs/1609.048
- Photo inpainting
  - https://arxiv.org/abs/1604.07379
- And more…

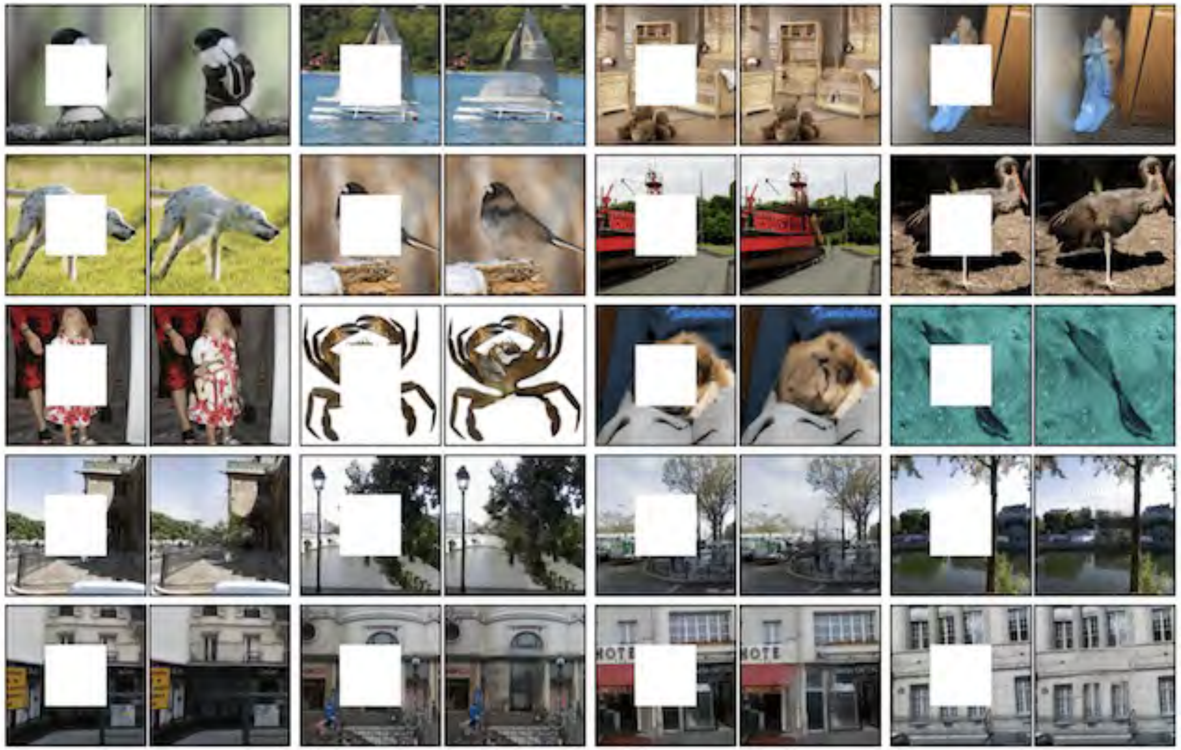

Conclusions

- Generative Adversarial Networks (GANs) map from one domain to another
- Use a Generator and a Discriminator
  - Work in tandem to train the GAN
- Simple GAN -> generates data from random values
- Conditional GAN -> takes some input and maps this to new domain
- Style Transfer GANS -> takes some input and maps, but doesn't need matching examples
- GANs are new – lot's of new examples all the time

stephen.mcgough@Newcastle.ac.uk