

Generative Adversarial Networks

Talk & Live-Demo


26.11.2019 19:00

Linux User Group Frankfurt

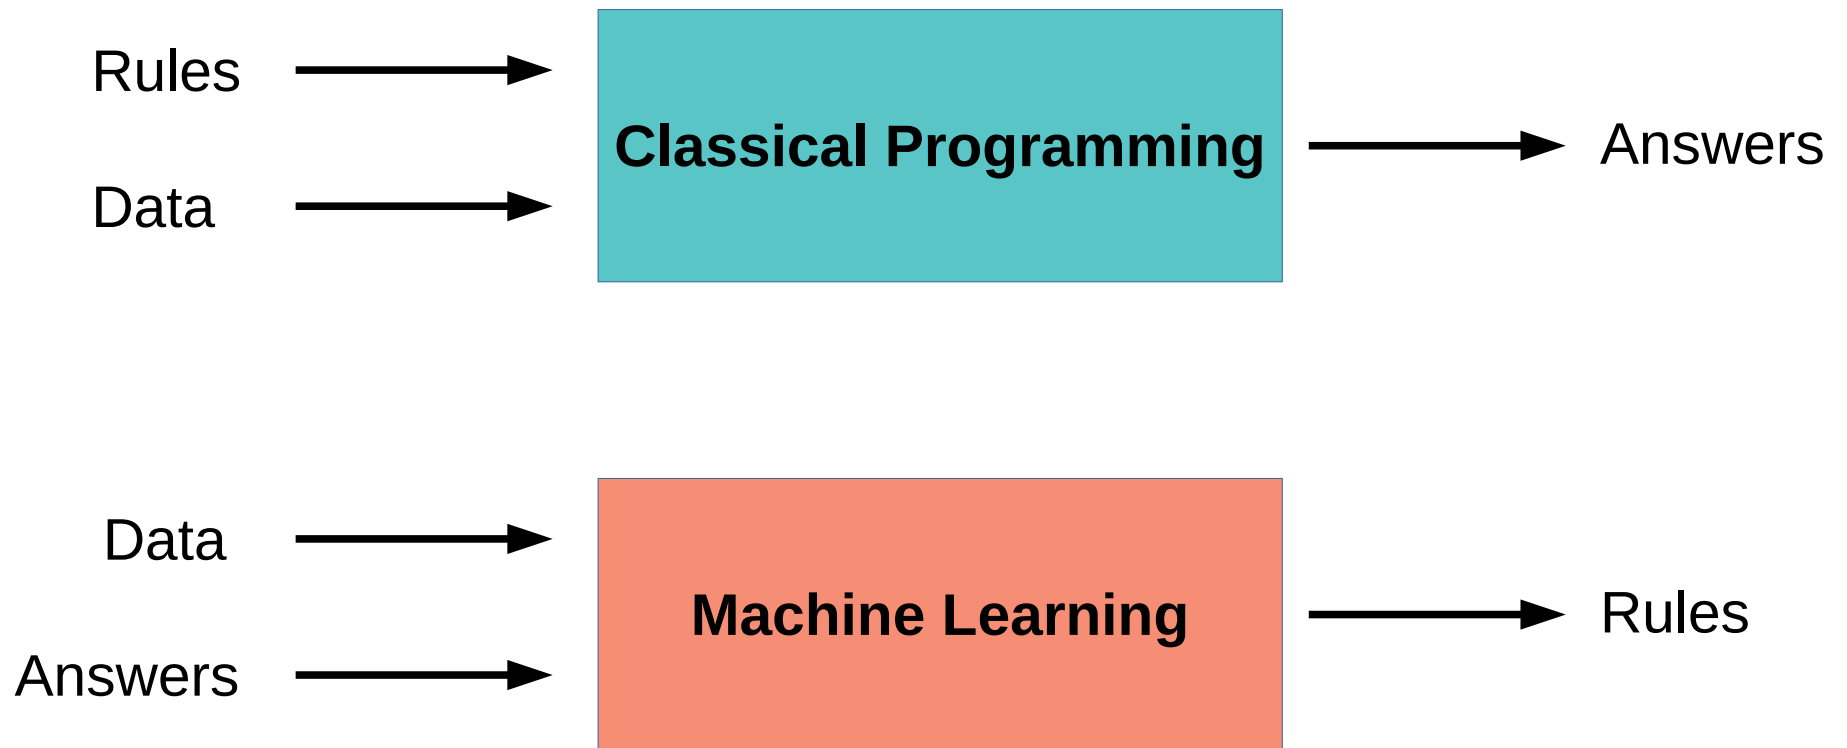
GAN application examples

- **“First AI-generated painting auctioned at Christie’s”**
<https://www.theverge.com/2018/10/23/18013190/ai-art-portrait-auction-christies-belamy-obvious-robbie-barrat-gans>
- **Automated coloring of black-white photographs**
<https://github.com/jantic/DeOldify>
- **<https://thispersondoesnotexist.com/>**
- **Cross-Domain transfer**
<https://hardikbansal.github.io/CycleGANBlog/>
- **Few-Shot transfer**
<https://medium.com/syncedreview/samsung-ai-makes-the-mona-lisa-speak-bea2b8362c38>
- **Deep Fakes**
<https://arxiv.org/abs/1909.11573>
- **‘Vishing’**
<https://thenextweb.com/security/2019/09/02/fraudsters-deepfake-ceos-voice-to-trick-manager-into-transferring-243000/>
- **Biosignal denoising**
<https://github.com/softserveinc-rnd/ecg-denoise>

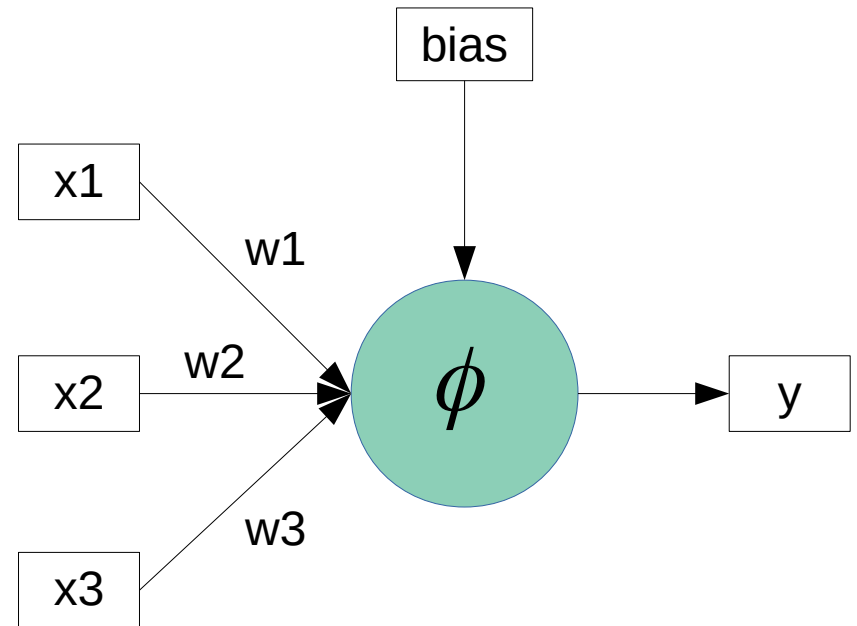

Artificial Intelligence, Machine Learning & Deep Learning



Classical Programming versus Machine Learning




Artificial neurons: Building blocks for Neural Networks




$$y = \phi\left(\left(\sum^n x_n \cdot w_n\right) + b\right)$$

Neural Networks: Common types


Dense Feedforward



Convolutional networks



Recurrent networks

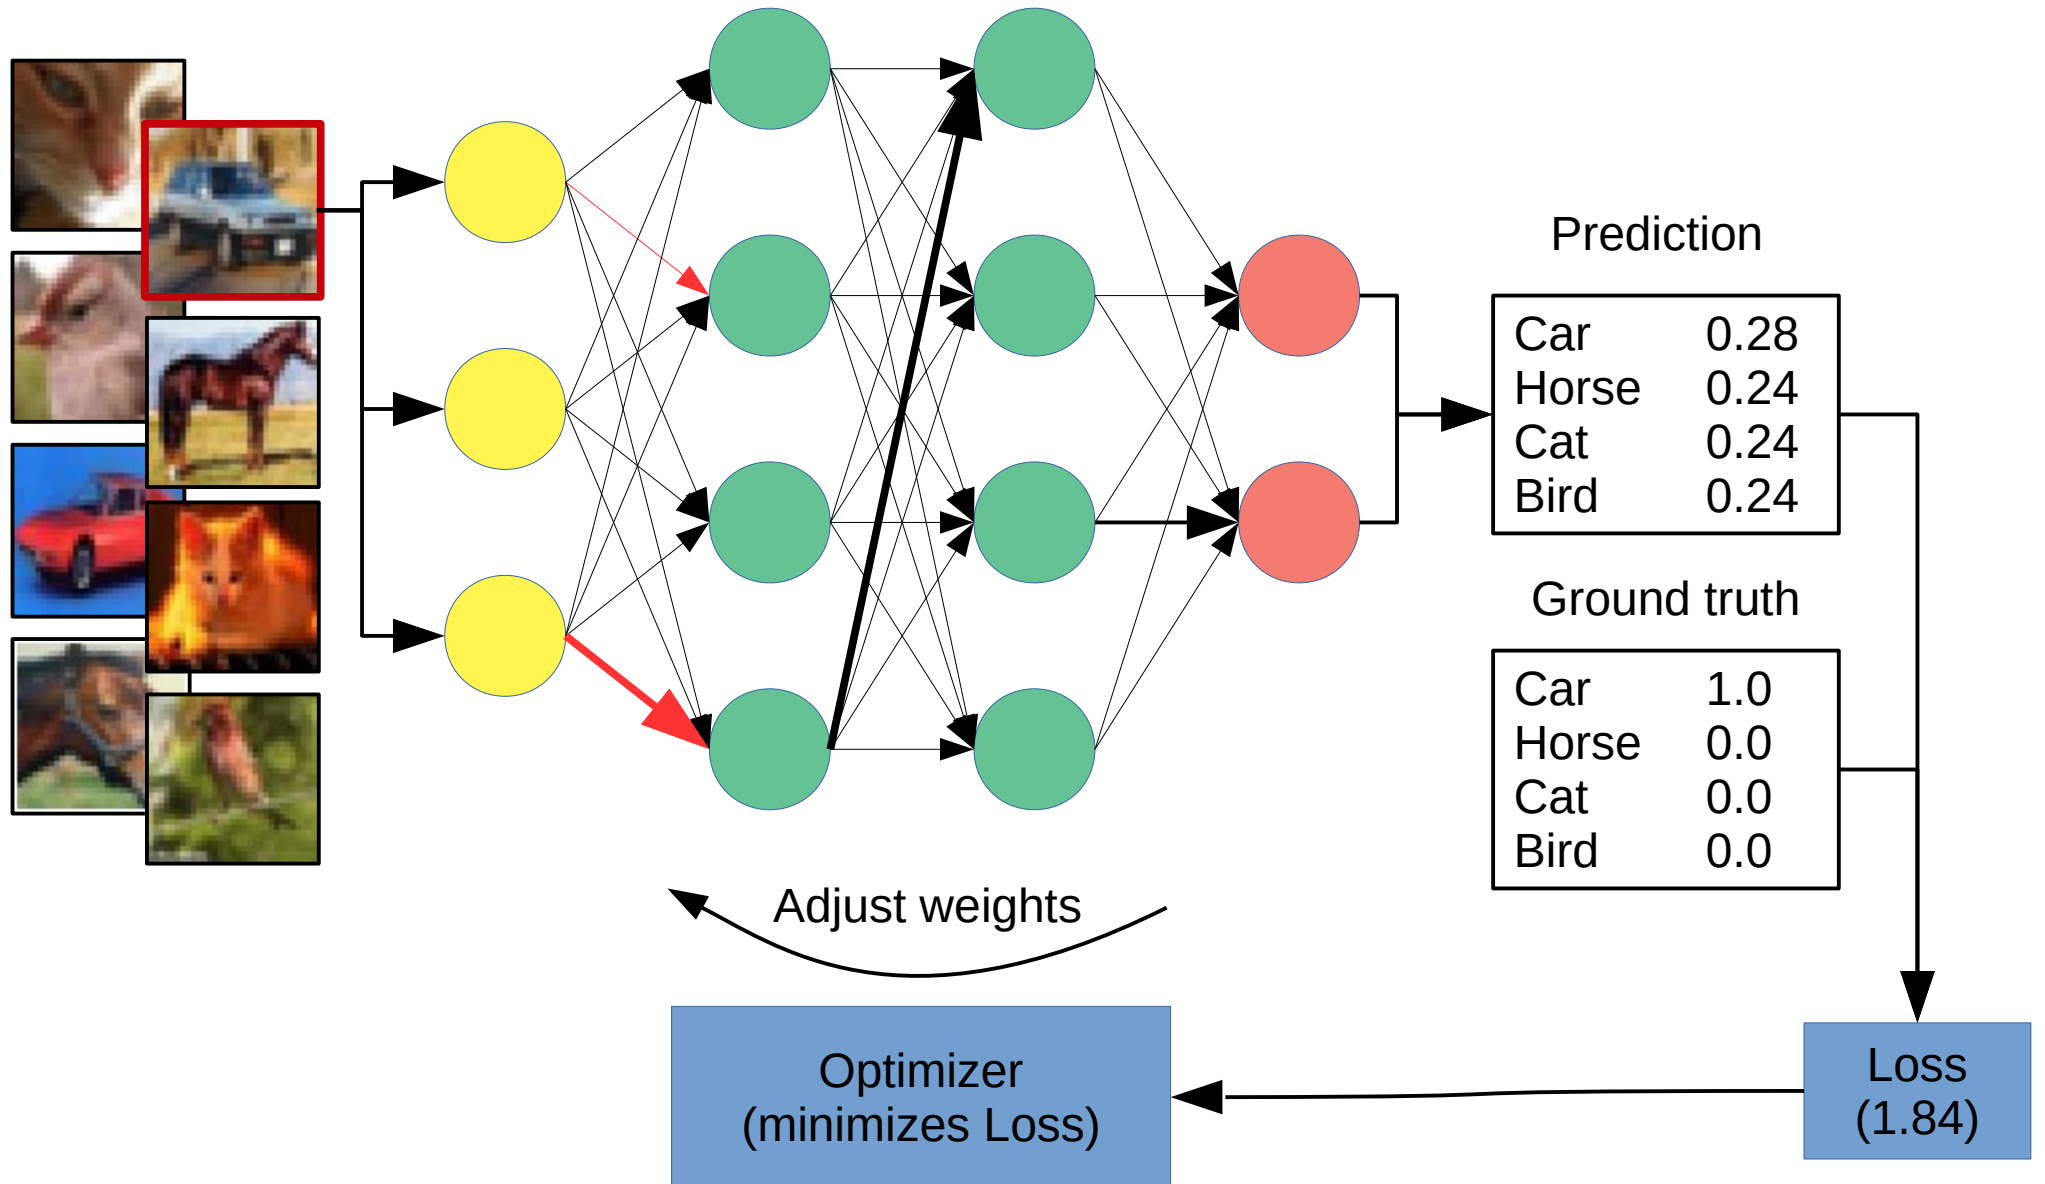


- All neurons connected between layers
- Data flow left to right
- Fixed-sized input
- General purpose for classification / regression problems


- Convolution operation on local features
- Weights are shared across space
- Can process input of any size
- Object recognition

- Feedback loops -> Remembers past data
- Weights are shared across time
- Can process input of any length
- Time series analysis; translation; speech recognition

Training a neural network - Supervised learning



Predicting from a trained neural network (,Inference‘)



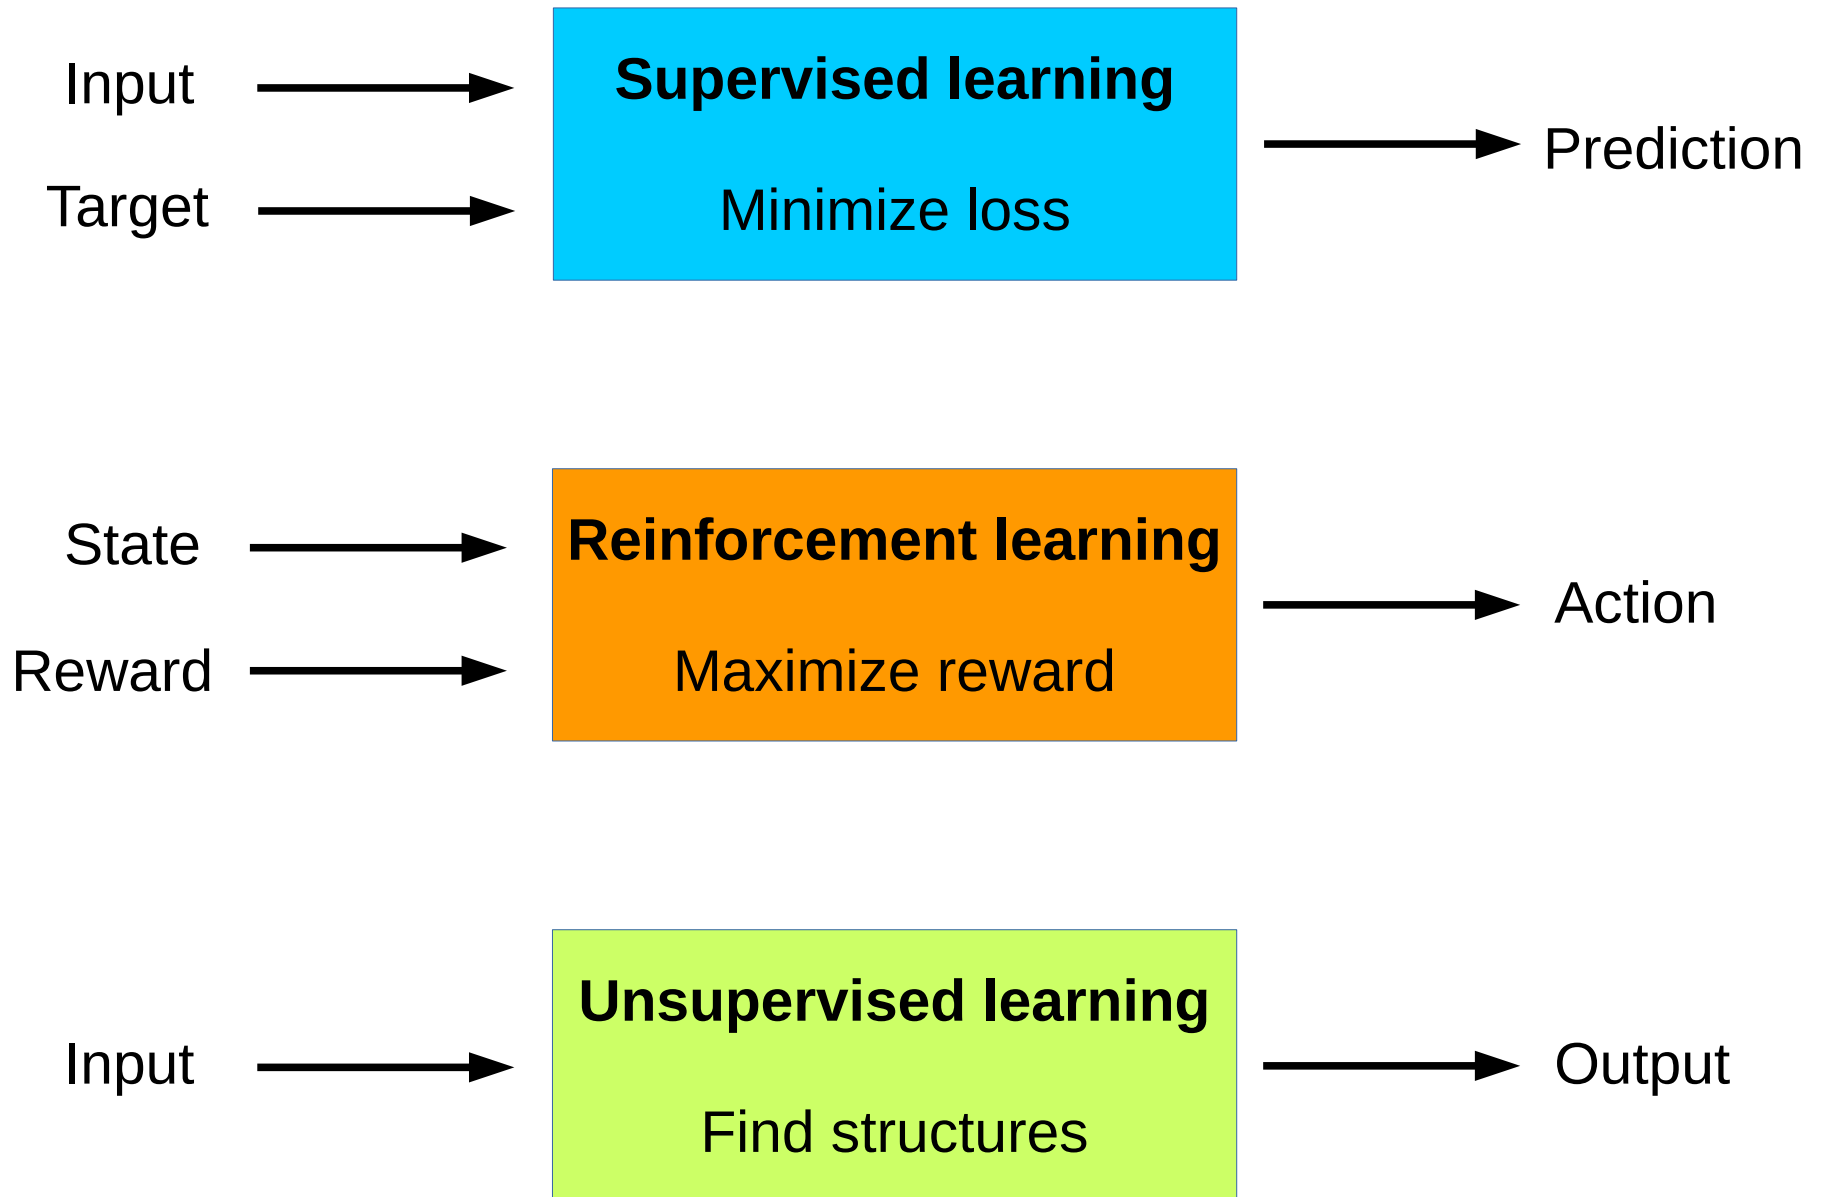
Production network with trained weights

A well-trained network can generalize new input

Learned knowledge is contained in the **weights**

Generative **A**dversarial **N**etworks

Learning methods in AI





A **GAN** learns to *mimic* a given *data distribution*
(-> unsupervised learning)

Source dataset samples



Generated samples



CIFAR10 dataset, class 7 'Horse'

DCGAN, after 80k epochs / 8h


1990 - 2013: Preliminary work

- ‘Artificial curiosity’, *Schmidhuber*, 1990
- Model vs. discriminator architecture, *Li / Gauci / Gross*, 2013

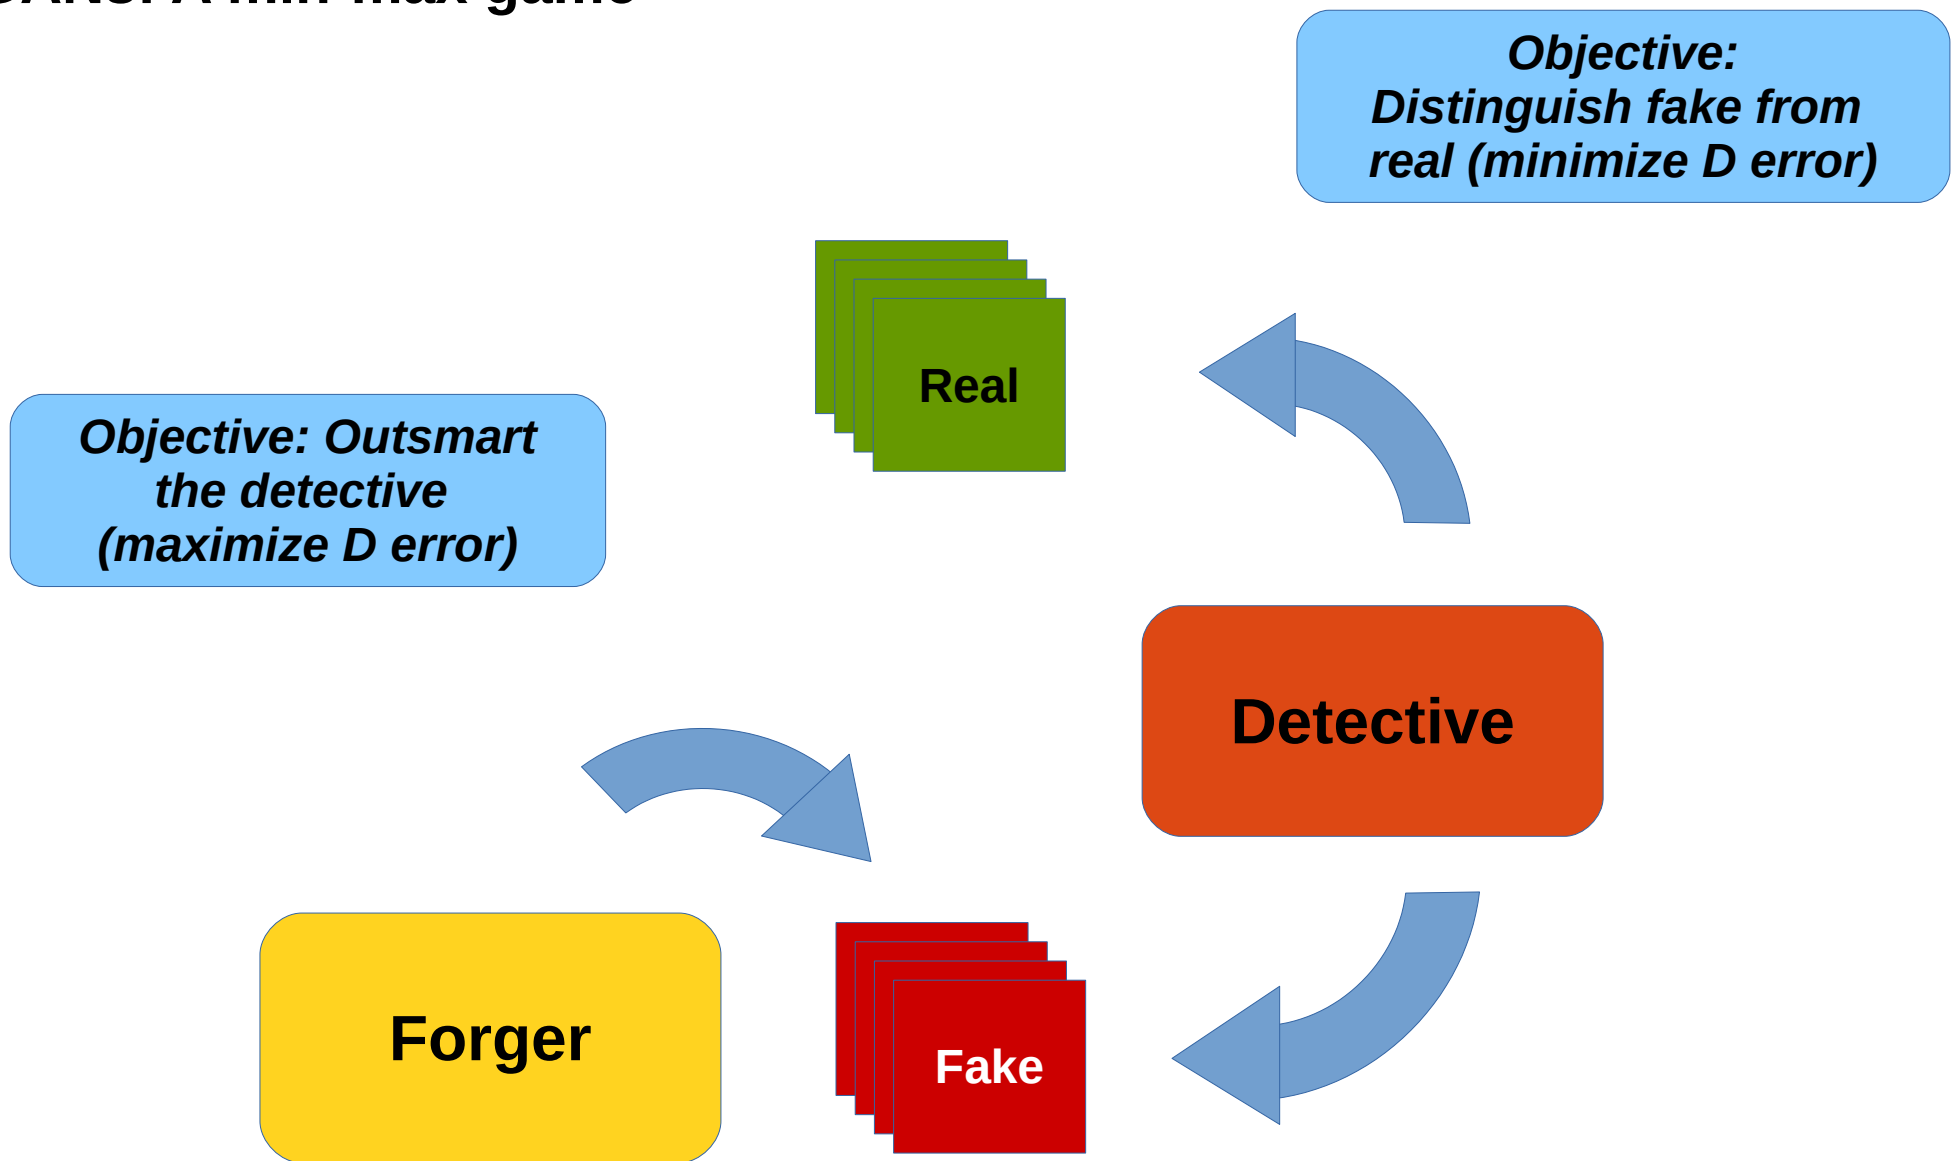
2014: ‘Generative Adversarial Nets’, Ian Goodfellow

- Breakthrough and name-defining paper

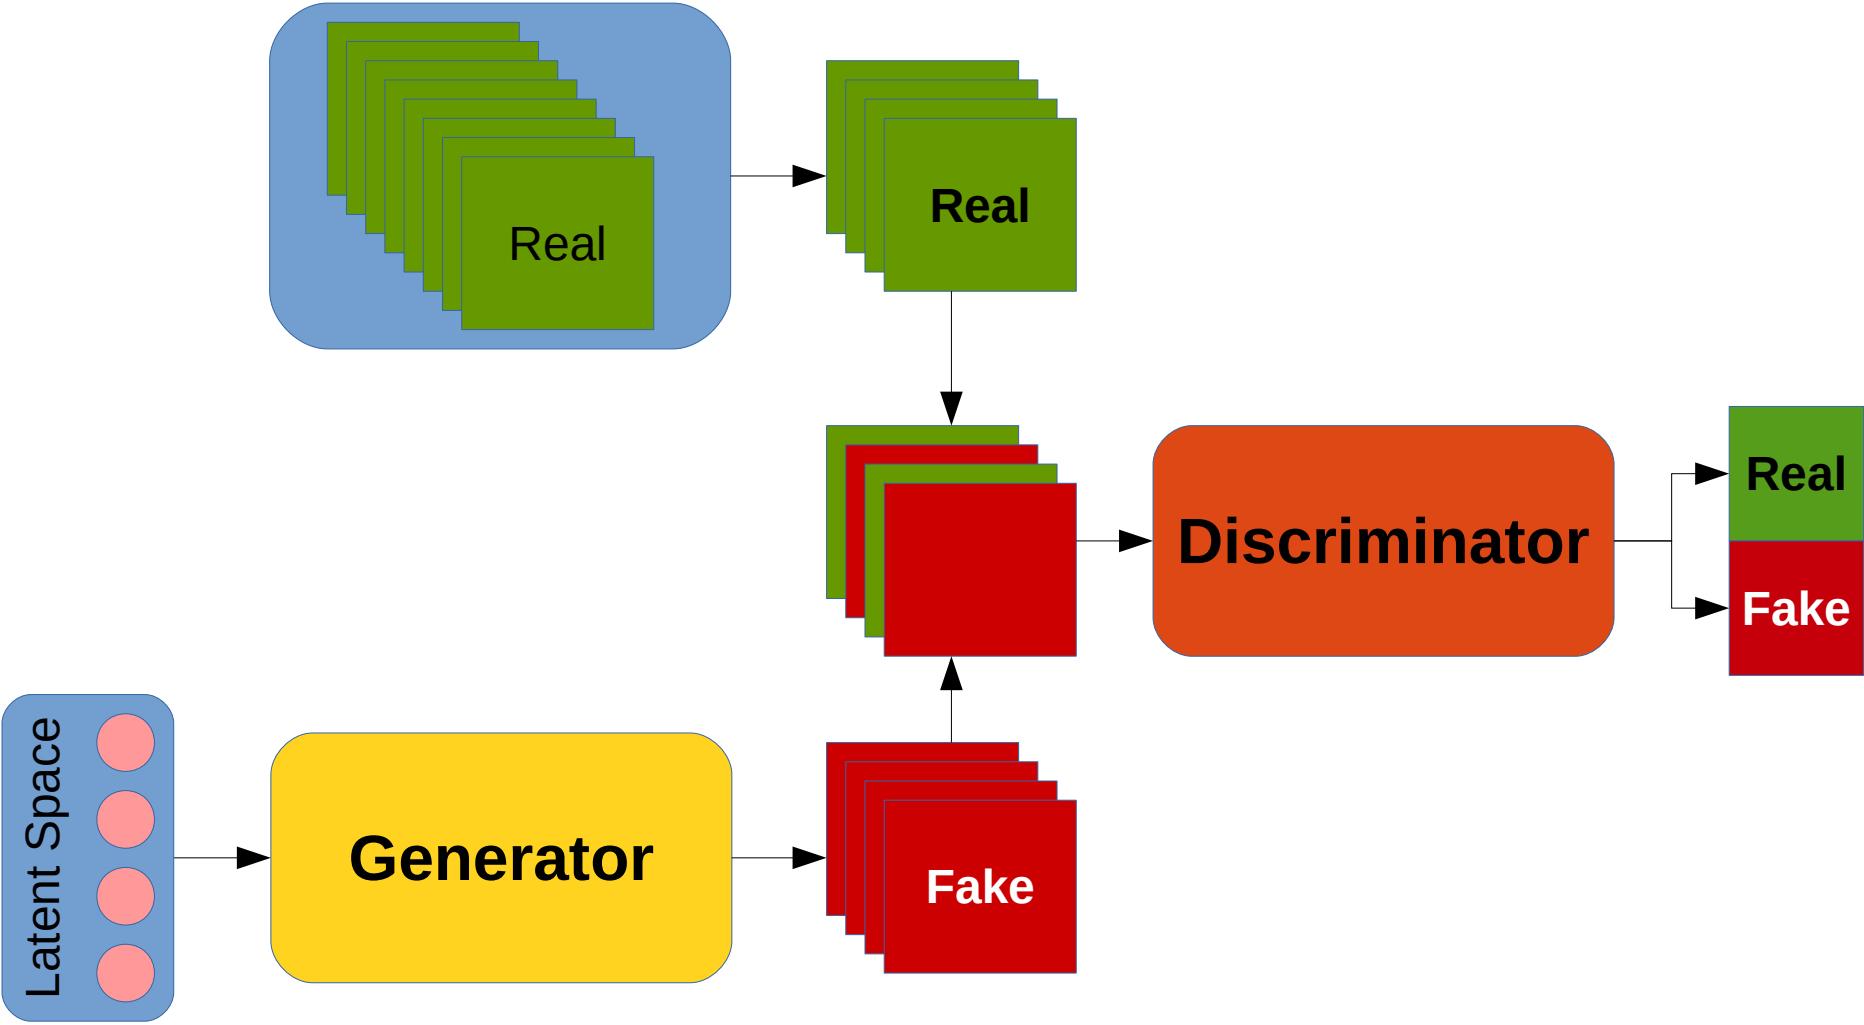
‘This, and the variations that are now being proposed is the most interesting idea in the last 10 years in ML, in my opinion.’ (Yann LeCun)



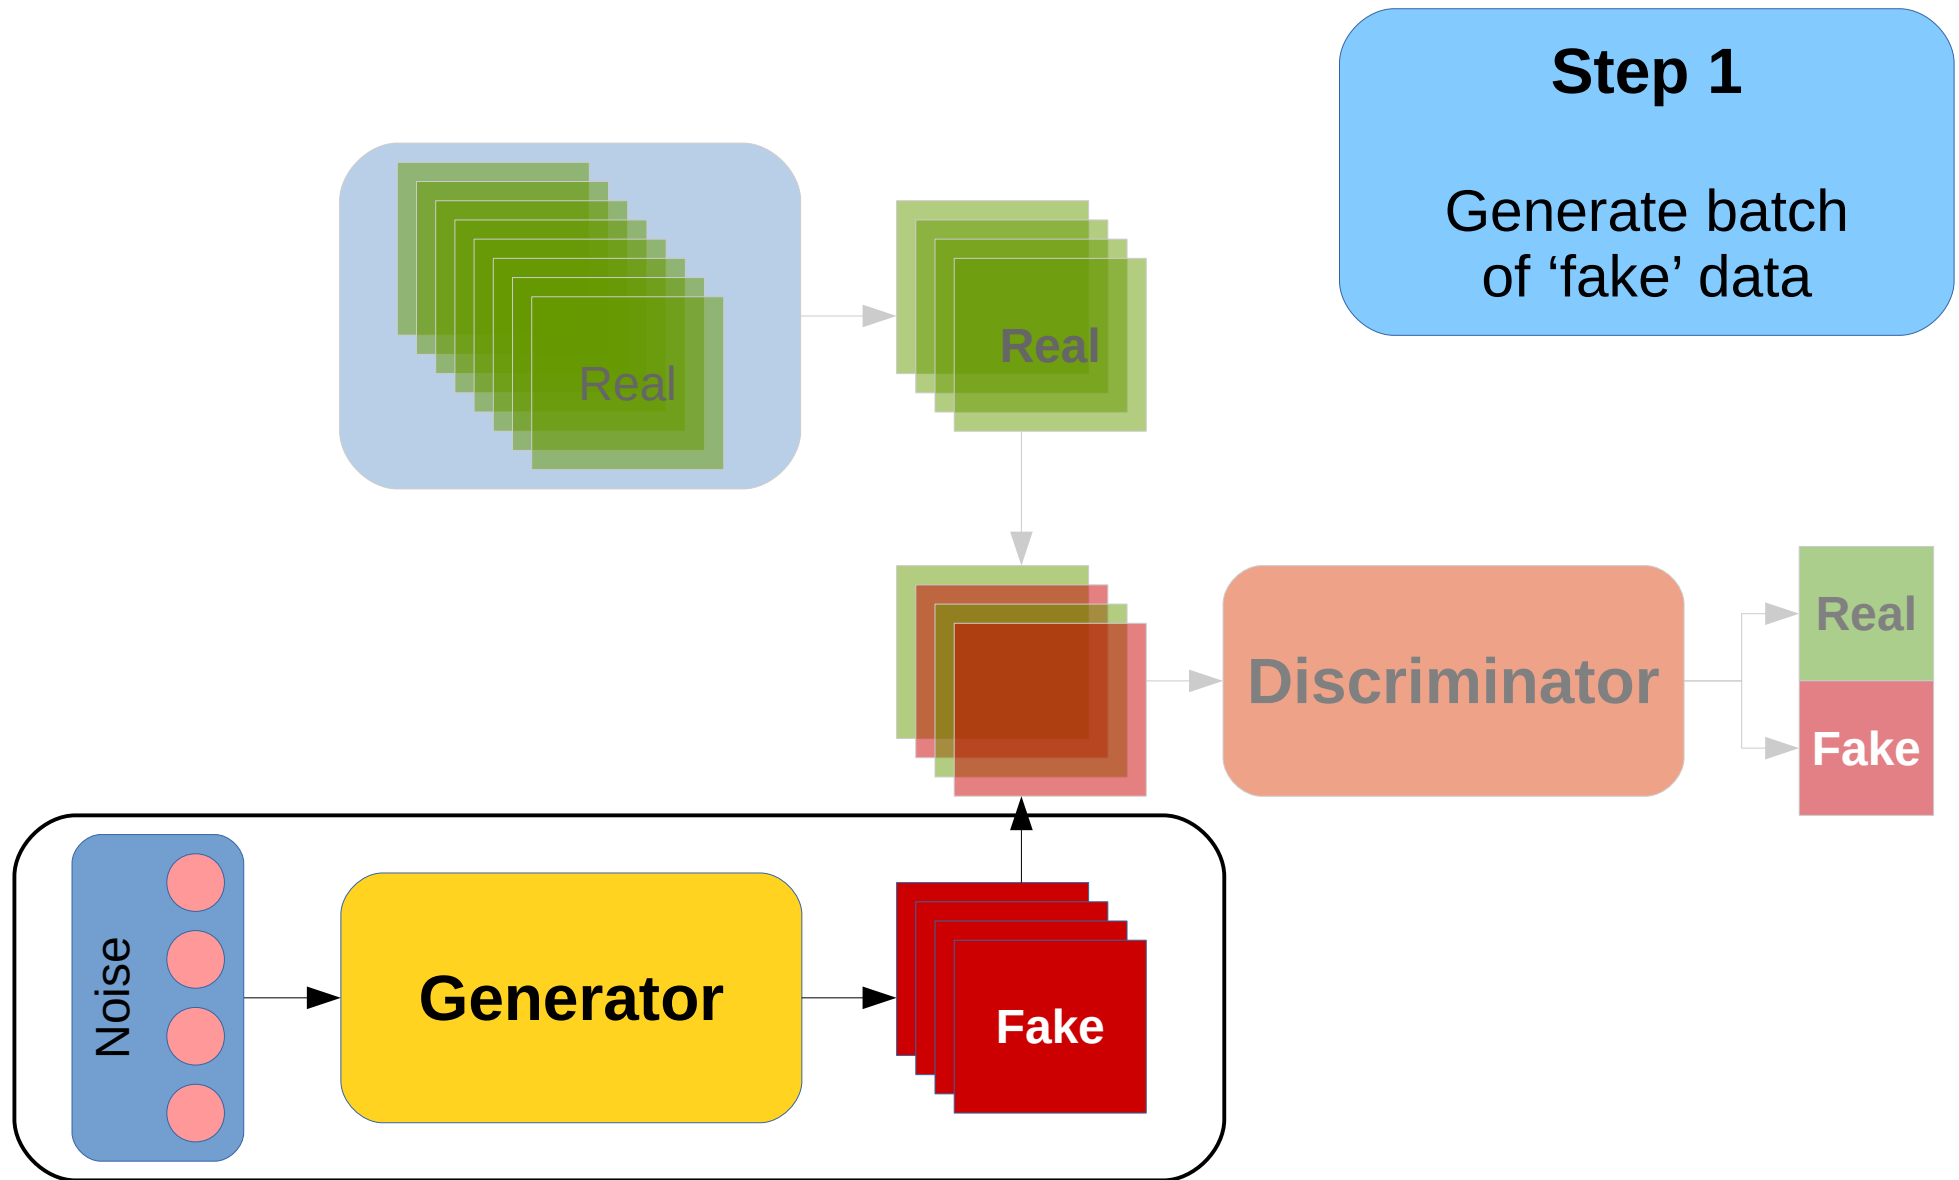
GANs: A min-max game



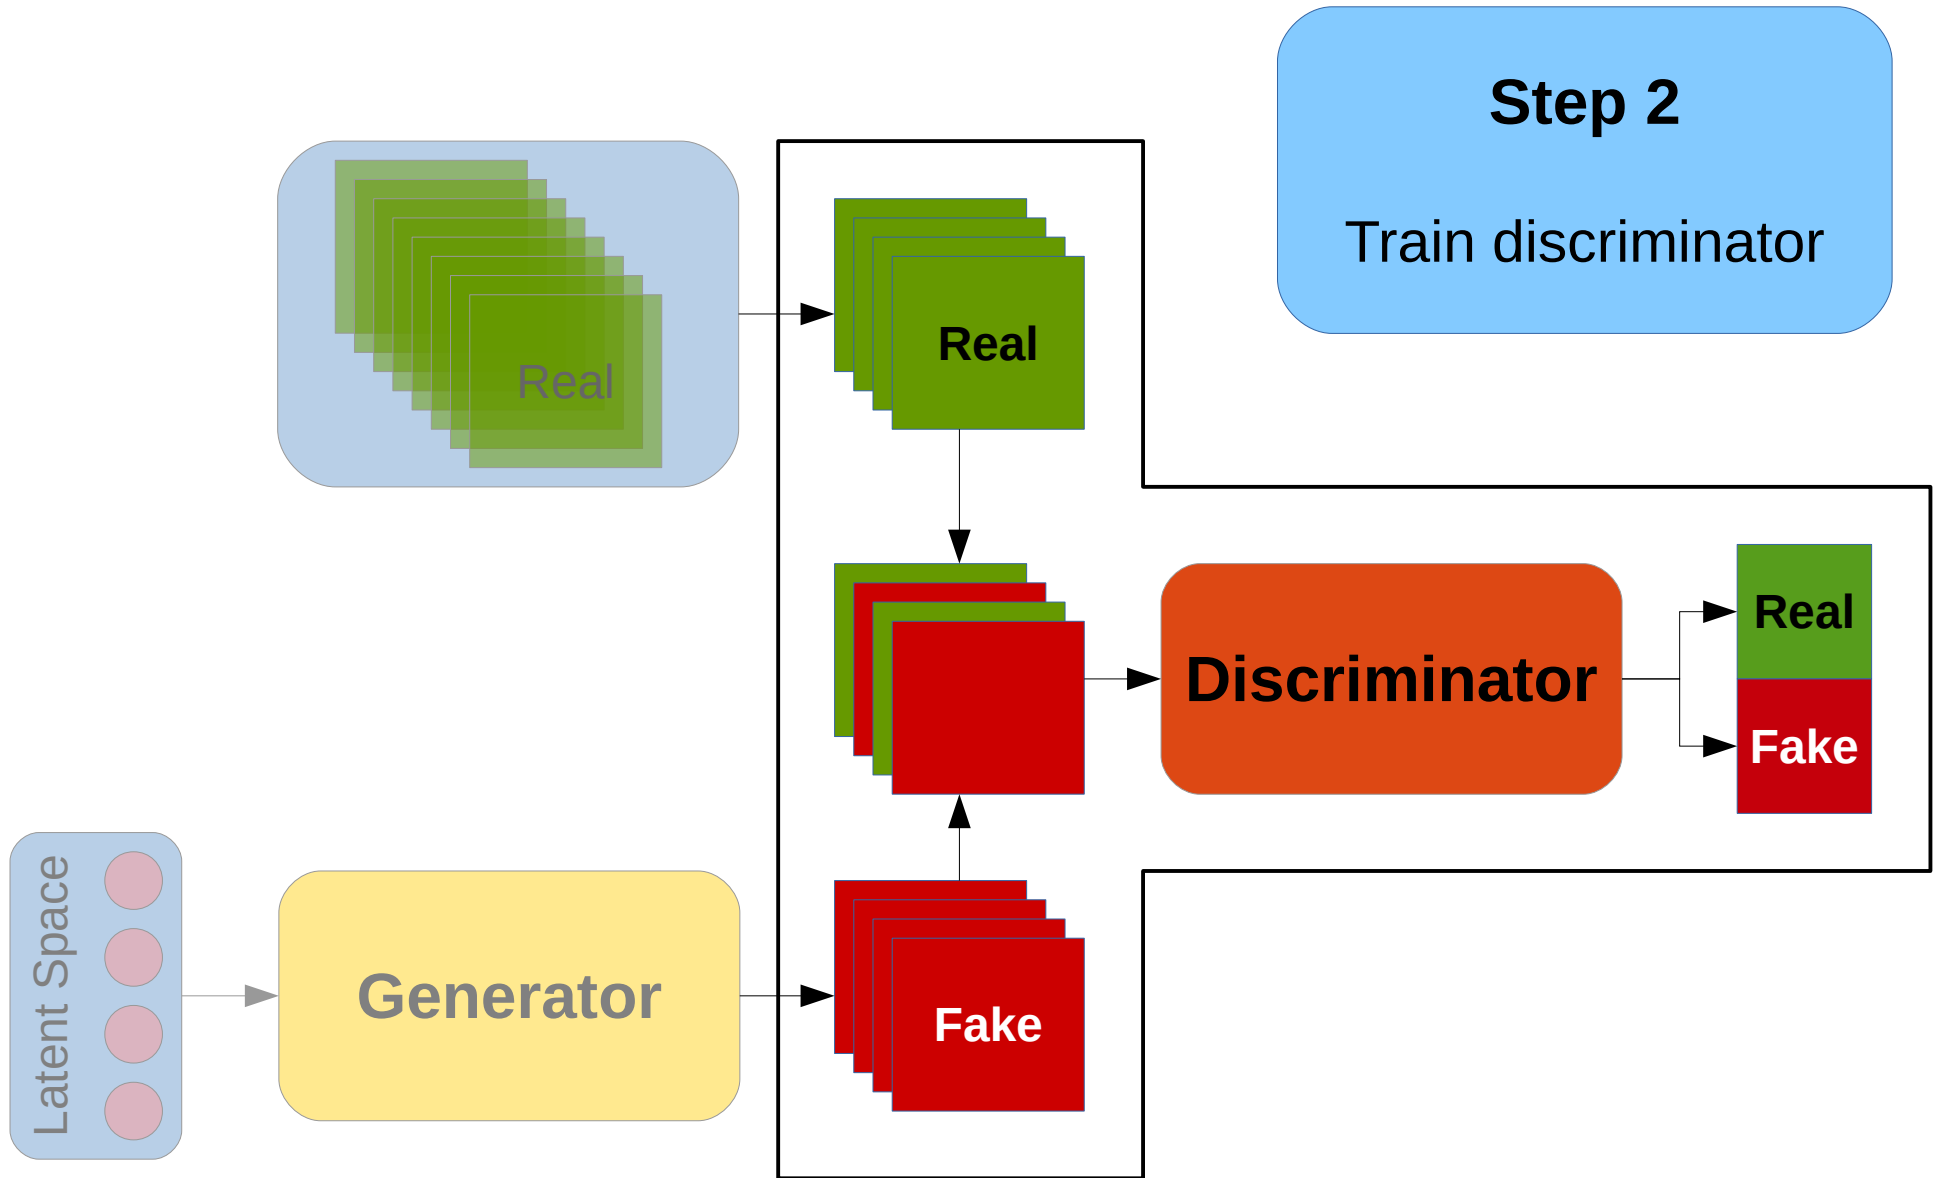
Basic GAN architecture



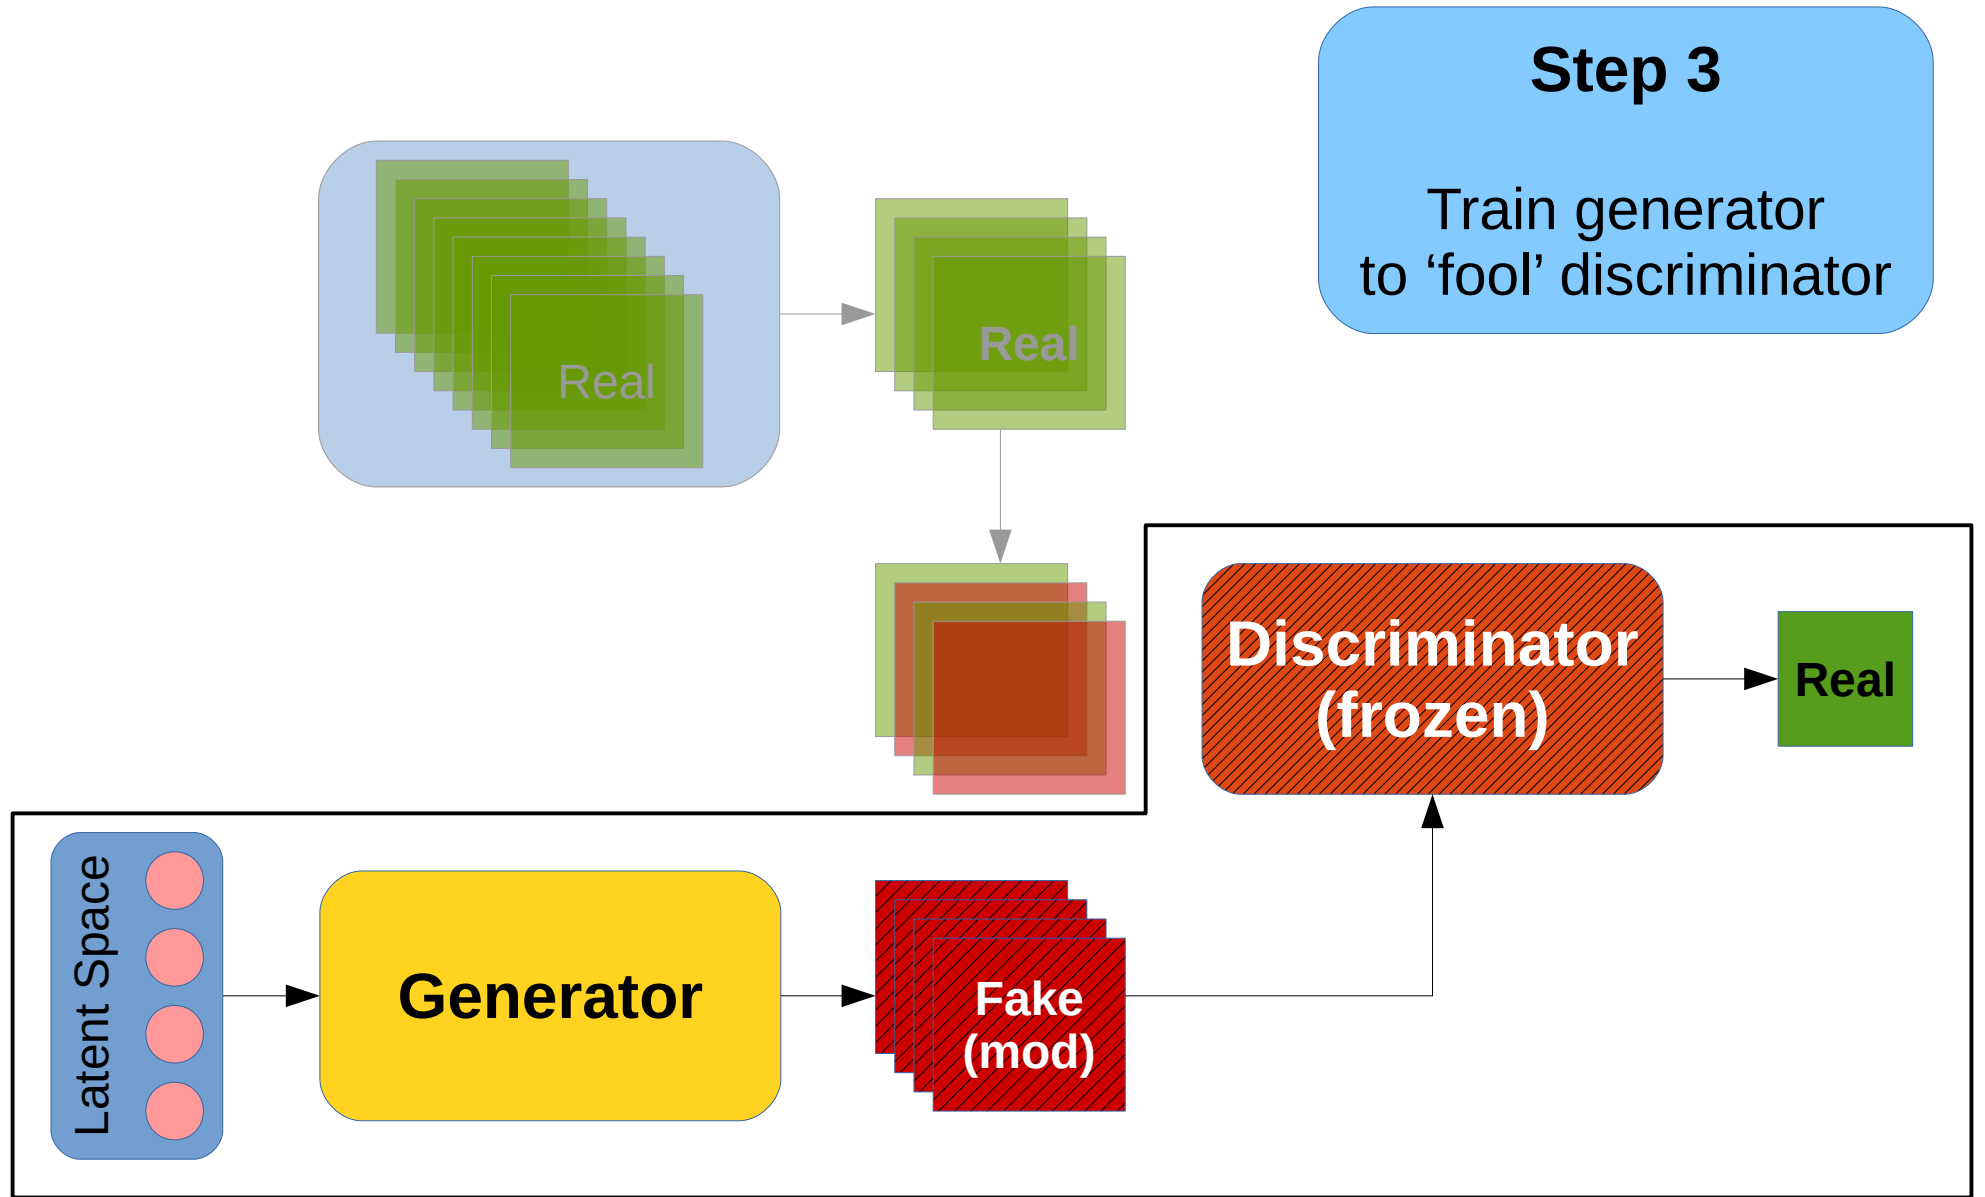
GANs: Training phase



GANs: Training phase



GANs: Training phase

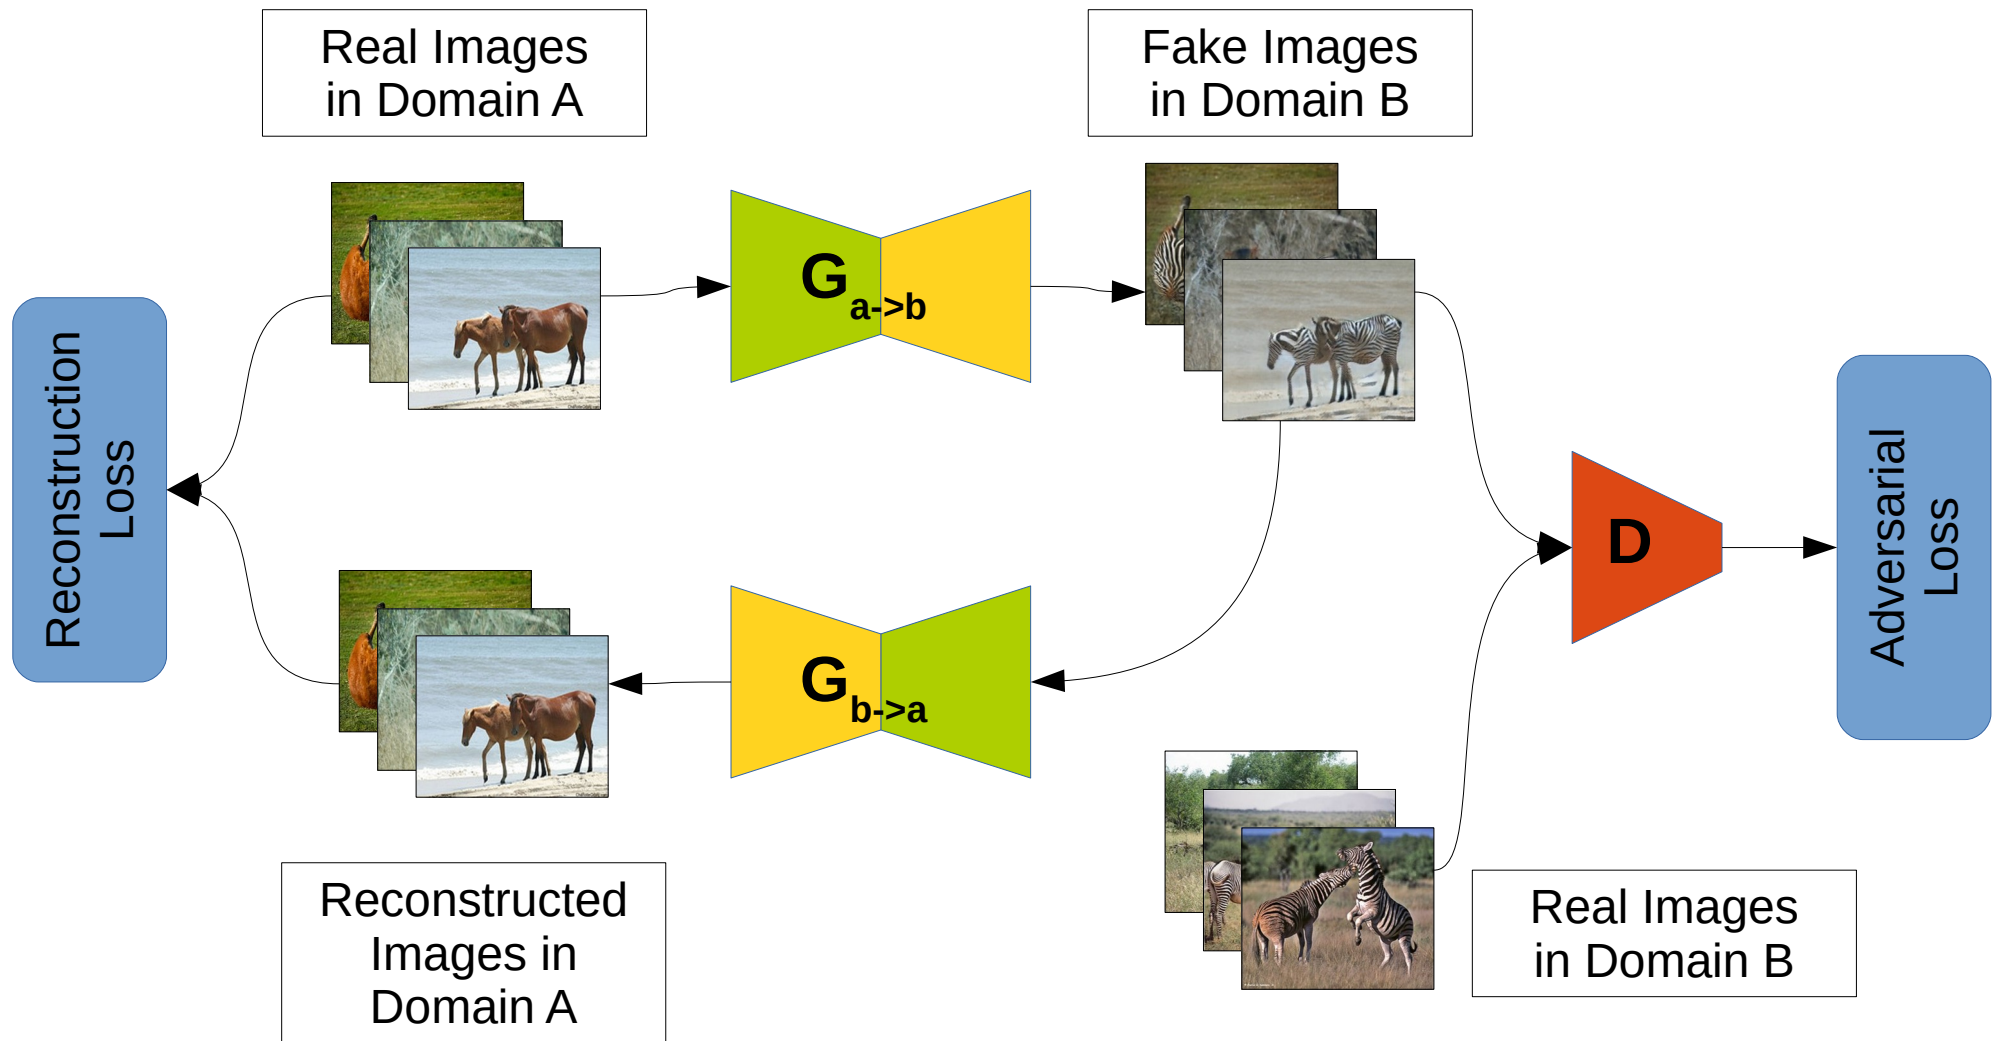


GANs: Using the generator

Create synthetic data with generator
Use latent space as parameter vector



Domain Transfer with a CycleGAN

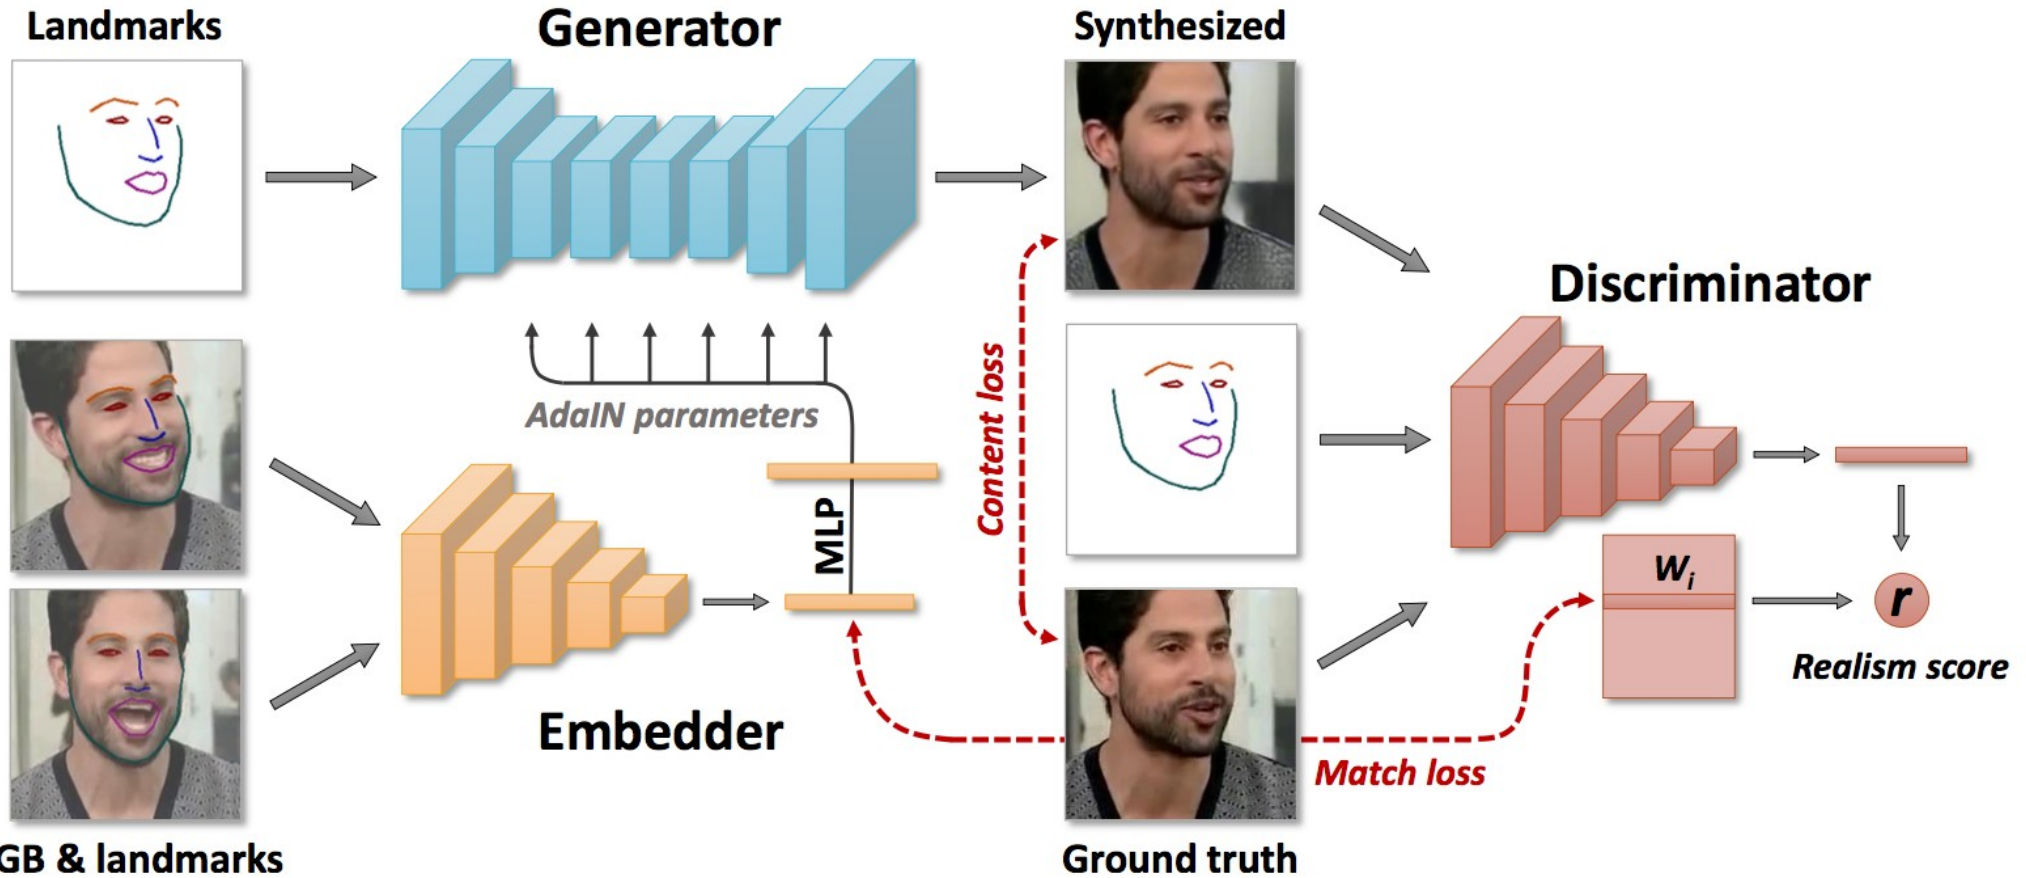


Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks
Jun-Yan Zhu, Taesung Park, Phillip Isola, Alexei A. Efros

<https://arxiv.org/abs/1703.10593>

<https://github.com/junyanz/CycleGAN>

Few-Shot GAN



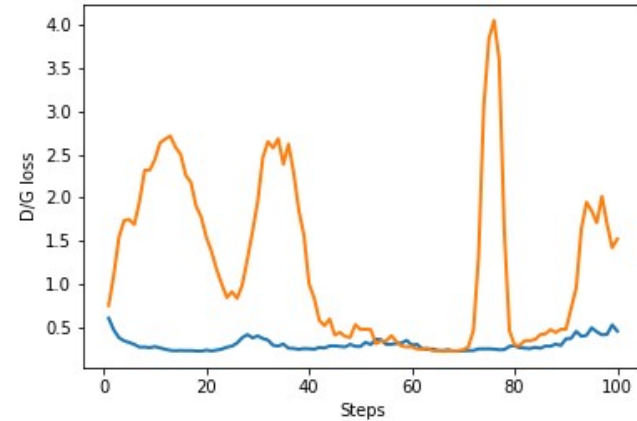
Few-Shot Adversarial Learning of Realistic Neural Talking Head Models
Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, Victor Lempitsky

<https://arxiv.org/abs/1905.08233>

GAN problems

Instability / Non-convergence


- Failure to approach *Nash equilibrium*
- *Vanishing gradient*: One network 'outsmarts' the other, no training progress



Mode collapse ('Helvetica scenario')

- Generator falls back to mode subset
- Discriminator rejects those modes
- Complete learning breakdown

Generated Distribution




Data Distribution

Global structure problems

('Cerberus effect')

- Caused by convnet spatial invariance
- Important to spot GAN-generated fakes
- New architectures try to solve this



GAN problems: Transfer fail



Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks
Jun-Yan Zhu, Taesung Park, Phillip Isola, Alexei A. Efros

<https://arxiv.org/abs/1703.10593>


<https://github.com/junyanz/CycleGAN>

Hot to spot GAN-created 'deep fakes'

General caution

Context? Political agenda? Rage-inducing?
Confirmed by other sources?

Checkerboard
& pixel
artifacts



Unnatural textures

Hair is difficult to reproduce

Global structure problems

'Cubist error'
Ear should not be visible,
Eyes misaligned

More information:

See also:

<https://blog.inten.to/welcome-to-the-simulation-dd0d8cb6534d>

<https://developers.google.com/machine-learning/gan>

Do-it-yourself in a webbrowser playground:

playground.tensorflow.org



poloclub.github.io/ganlab/

Jupyter notebooks used in the live demos:

<https://github.com/SteffenBauer/KerasTools/blob/master/Notebooks/04%20MNIST%20complete%20workflow.ipynb>

<https://github.com/SteffenBauer/KerasTools/blob/master/Notebooks/13%201%20MNIST%20DCGAN.ipynb>

Live Demo 1: Deep Learning



Live Demo 2: GAN

