



UNIVERSITY OF TARTU



Generative Adversarial Networks

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 857627 (CIPHR)



J. Amudhavel

SENIOR RESEARCHER

Monday, February 17, 2025



This Project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 857627 (CIPHR)



J. Amudhavel

SENIOR RESEARCHER

Introductory Goals

1. Understand the different types of machine learning.
2. Understand the key concepts of supervised machine learning.
3. Learn how solving problems with ML is different from traditional approaches.

Types of ML Systems

1. ML systems fall into three distinct categories based on how they learn to make predictions:

- ❖ Supervised learning
- ❖ Unsupervised learning
- ❖ Reinforcement learning

Supervised Learning

Supervised learning models can make predictions after seeing lots of data.

Ex: Student Learning.

Supervised Learning

Regression

A regression model predicts a numeric value. For example, a weather model that predicts the amount of rain, in inches or millimeters, is a regression model.

Classification

Classification models predict the likelihood that something belongs to a category.

Supervised Learning



Regression



What will be the temperature tomorrow?

84°



Fahrenheit

Classification



Will it be hot or cold tomorrow?

COLD

HOT



Fahrenheit

Supervised Learning

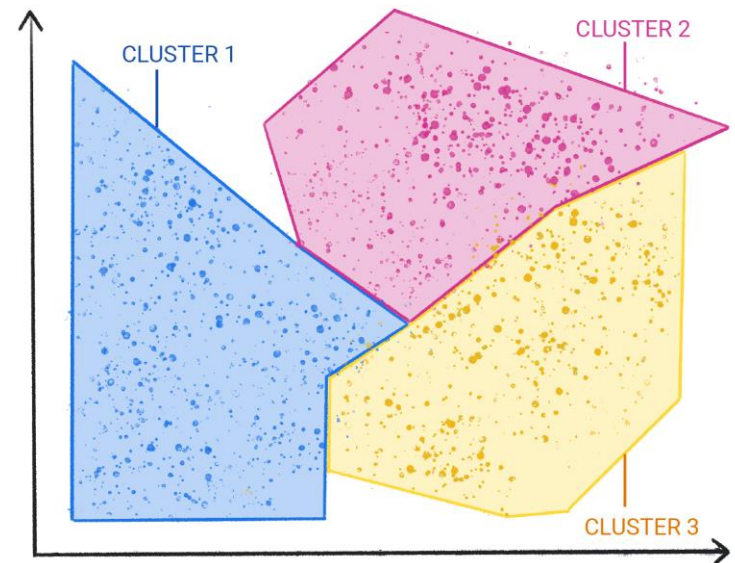
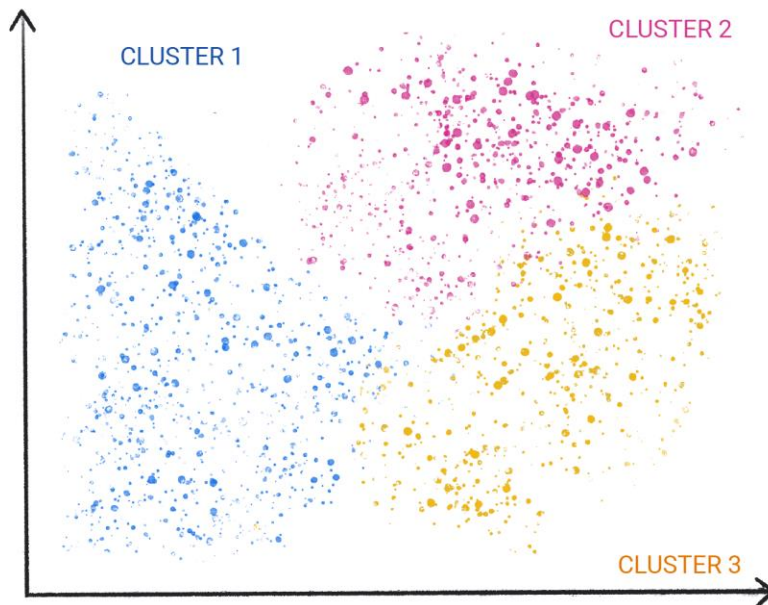
Classification models are divided into two groups:

Binary classification and

Multiclass classification.

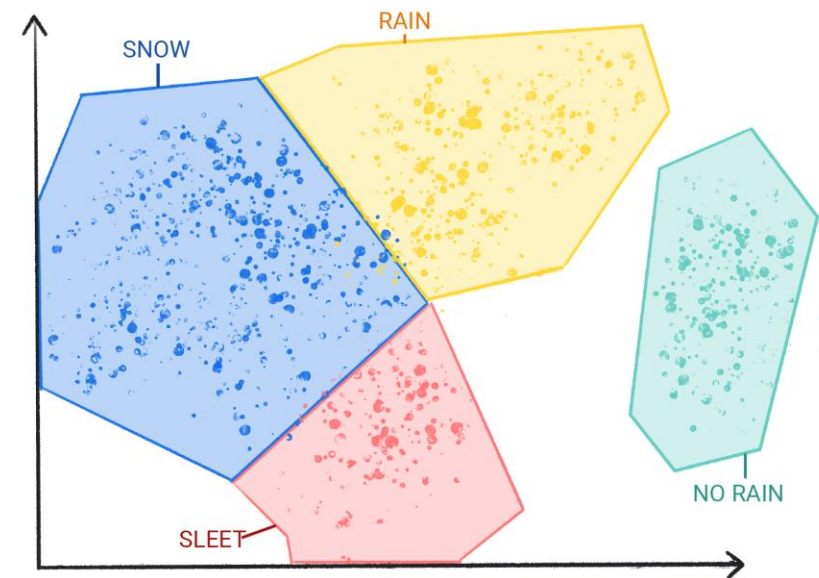
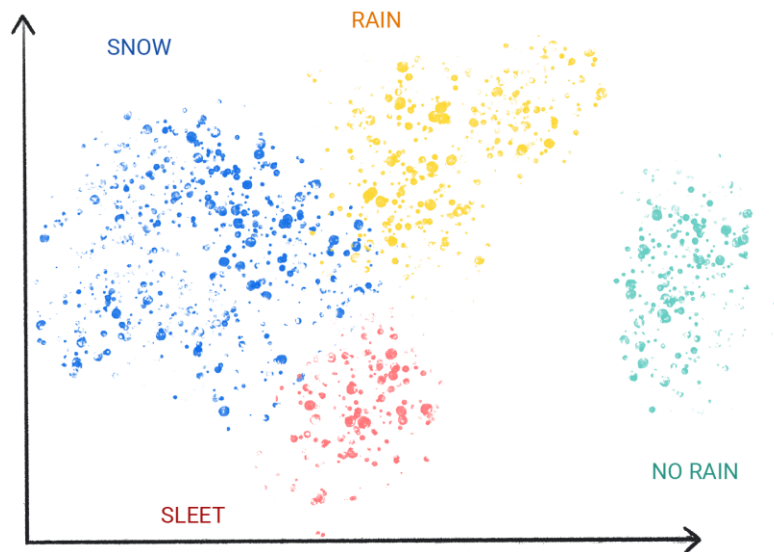
Unsupervised Learning

Unsupervised learning model's goal is to identify meaningful patterns among the data.



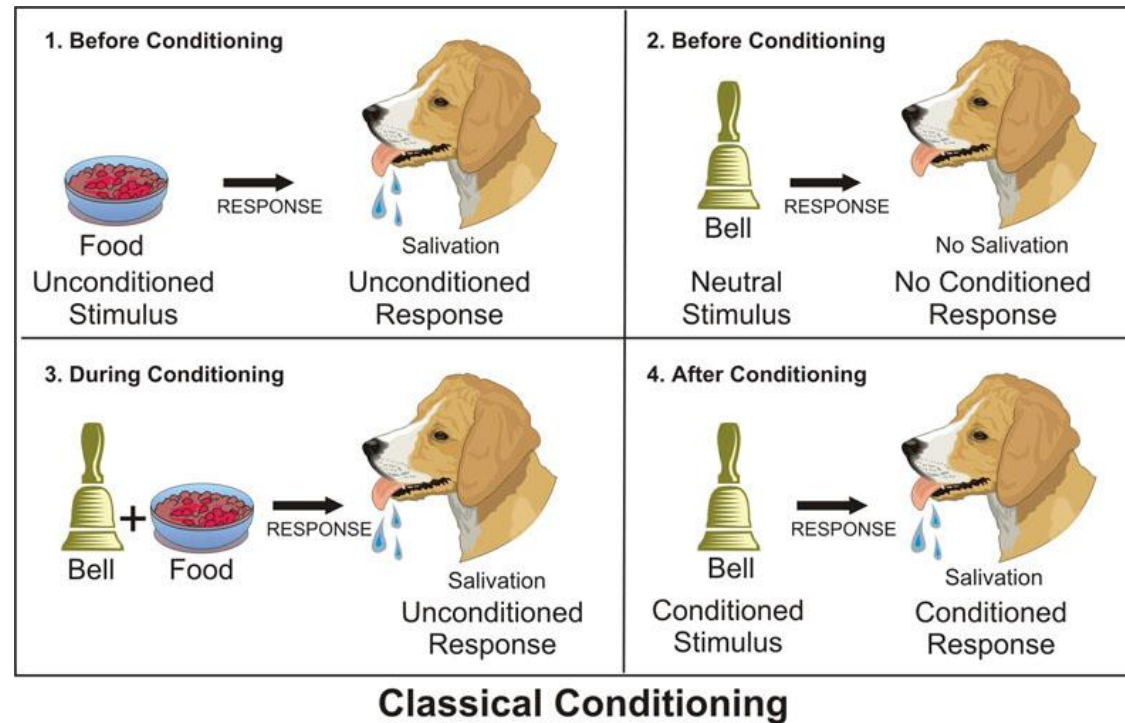
Unsupervised Learning

Unsupervised learning model's goal is to identify meaningful patterns among the data.



Reinforcement Learning

Reinforcement learning models make predictions by getting rewards or penalties based on actions performed within an environment.



Workaround

https://colab.research.google.com/github/google/eng-edu/blob/main/ml/cc/exercises/linear_regression_with_synthetic_data.ipynb?utm_source=mlcc&utm_campaign=colab-external&utm_medium=referral&utm_content=linear_regression_synthetic_tf2-colab&hl=en#scrollTo=_GMGgR6O54IN

https://colab.research.google.com/github/google/eng-edu/blob/main/ml/cc/exercises/linear_regression_with_a_real_dataset.ipynb?utm_source=mlcc&utm_campaign=colab-external&utm_medium=referral&utm_content=linear_regression_real_tf2-colab&hl=en#scrollTo=xRfxp_3yofe3


Intermediate Goals

1. To Understand the difference between generative and discriminative models.
2. Understand the roles of the generator and discriminator in a GAN system.
3. To understand loss functions in GAN Training.
4. Use the TF GAN library to make a GAN. – Hands on


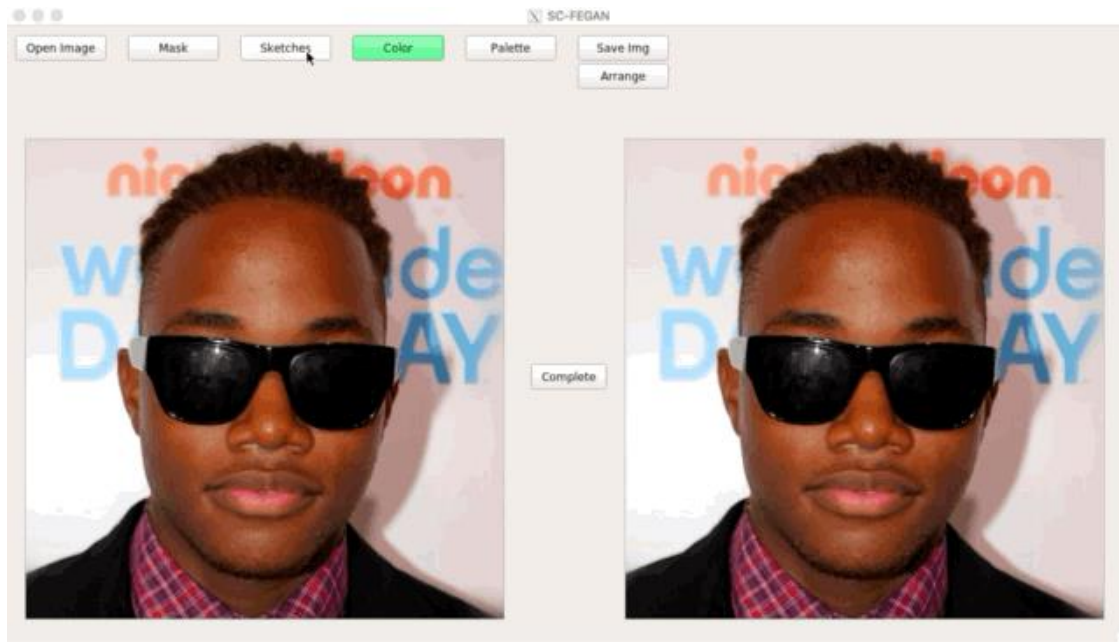
Identify the Celebrities



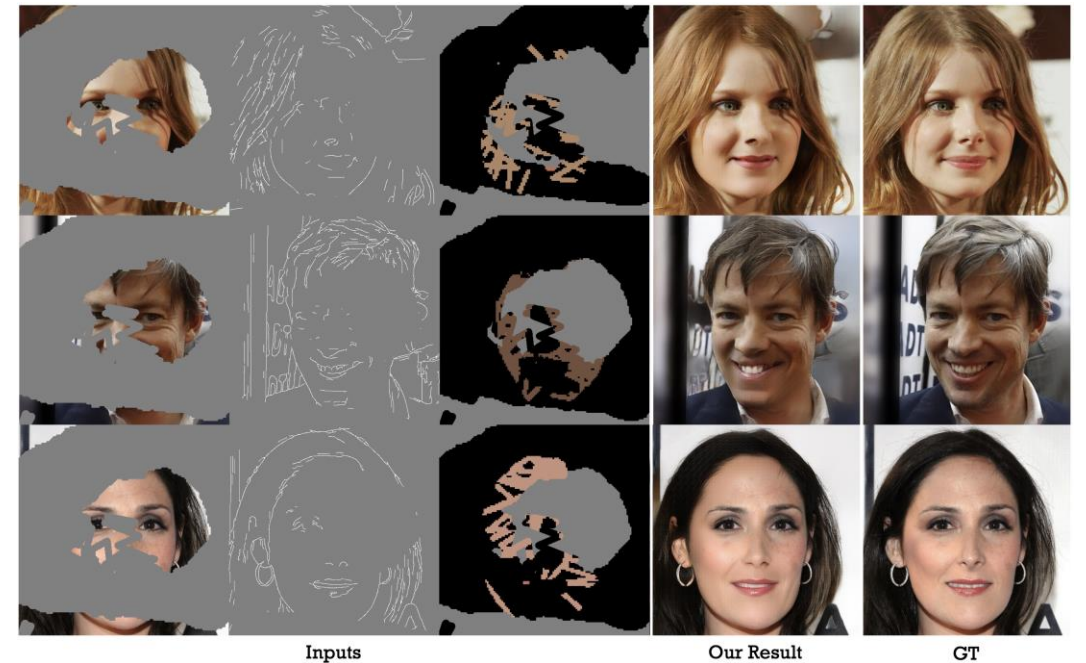

GAN - Progressive




GAN - Progressive



GAN - Progressive



GAN - Progressive



Free-form input

Our result

Target GT

GAN - Progressive

<https://t.ly/nDZLf>

<https://thispersondoesnotexist.com/>

<https://thisartworkdoesnotexist.com/>

<https://thiscatdoesnotexist.com/>

<https://thischemicaldoesnotexist.com/>

GAN - ?

Generative adversarial networks (GANs) are *generative* models that create new data instances that resemble your training data.

GANs can create images that look like photographs of human faces, even though the faces don't belong to any real person.

Generative Model

1. Generative models can generate new data instances.
2. Discriminative models discriminate between different kinds of data instances.
 - A generative model could generate new photos of animals that look like real animals, while a discriminative model could tell a dog from a cat.

GAN - Generative Model?

- GANs are just one kind of generative model.
- More formally, given a set of data instances X and a set of labels Y :
- **Generative** models capture the joint probability $p(X, Y)$, or just $p(X)$ if there are no labels.
- **Discriminative** models capture the conditional probability $p(Y | X)$.
- A generative model includes the distribution of the data itself and tells you how likely a given example is.
- A discriminative model ignores the question of whether a given instance is likely, and just tells you how likely a label is to apply to the instance.

Generative Models – Hard?

Generative models tackle a more difficult task than analogous discriminative models. Generative models have to model *more*.

Why it is Hard?

A generative model for images might capture correlations like "things that look like boats are probably going to appear near things that look like water" and "eyes are unlikely to appear on foreheads." These are very complicated distributions.

In contrast, a discriminative model might learn the difference between "sailboat" or "not sailboat" by just looking for a few tell-tale patterns. It could ignore many of the correlations that the generative model must get right.

Generative Models –Brainstorming

1. A model returns a probability when you give it a data instance. Is this model a generative model or a discriminative model? - **5 Minutes**

GAN Structure

- A generative adversarial network (GAN) has two parts:
- The **generator** learns to generate plausible data. The generated instances become negative training examples for the discriminator.
- The **discriminator** learns to distinguish the generator's fake data from real data. The discriminator penalizes the generator for producing implausible results.
- When training begins, the generator produces obviously fake data, and the discriminator quickly learns to tell that it's fake

Generated Data



Discriminator


FAKE

REAL

Real Data



GAN Training – Process



FAKE

REAL



GAN Training – Process



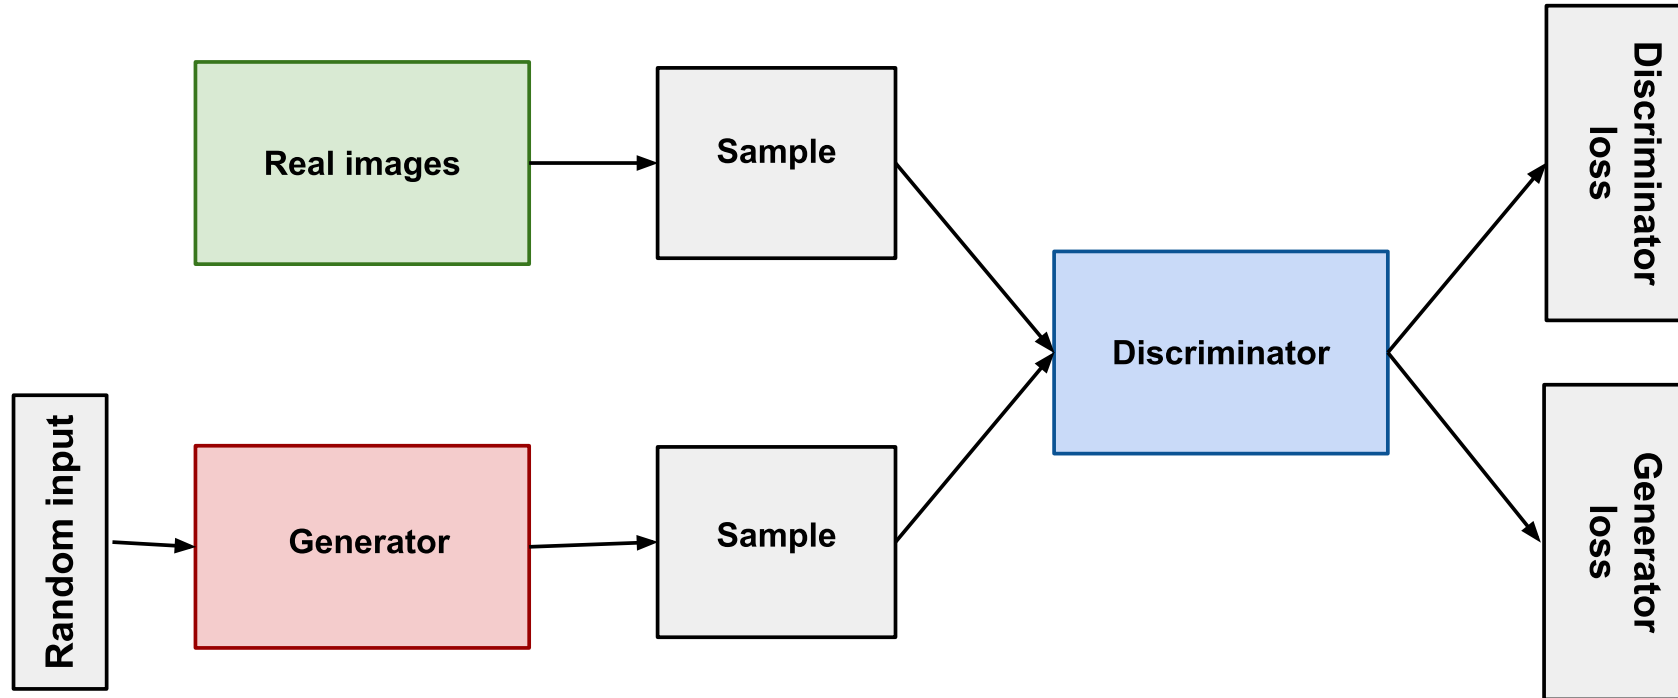
REAL

REAL

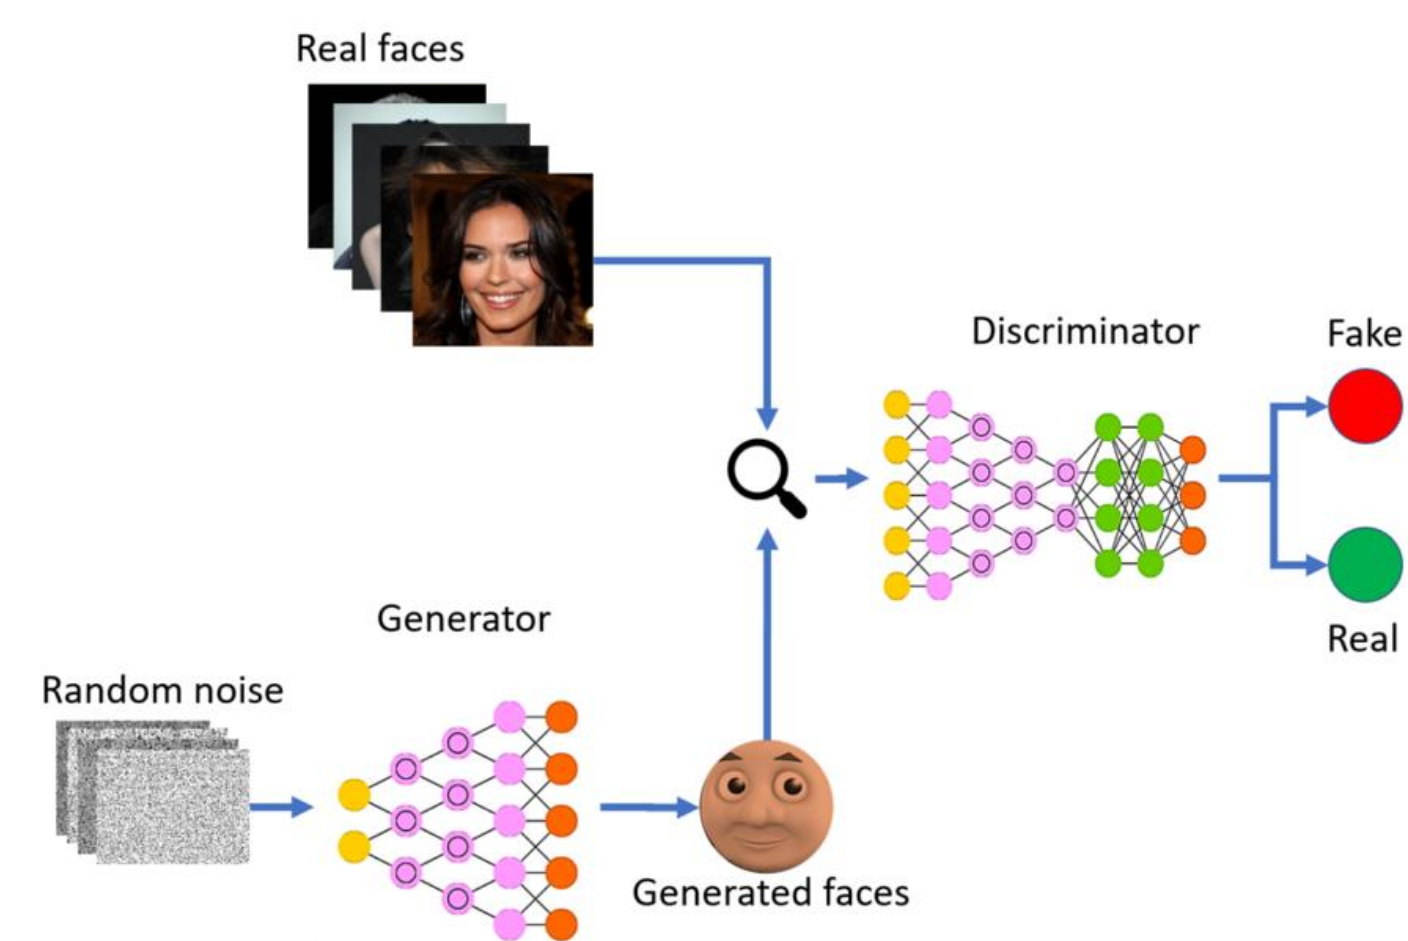


GAN Training – Process

Structure of GAN



Structure of GAN



Generator vs Discriminator

- A generative adversarial network (GAN) has two parts:
- The **generator** learns to generate plausible data. The generated instances become negative training examples for the discriminator.
- The **discriminator** learns to distinguish the generator's fake data from real data. The discriminator penalizes the generator for producing implausible results.
- When training begins, the generator produces obviously fake data, and the discriminator quickly learns to tell that it's fake

The Discriminator

The discriminator in a GAN is simply a classifier. It tries to distinguish real data from the data created by the generator. It could use any network architecture appropriate to the type of data it's classifying.

- 1) Random input
- 2) Generator network, which transforms the random input into a data instance
- 3) Discriminator network, which classifies the generated data
- 4) Discriminator output
- 5) Generator loss, which penalizes the generator for failing to fool the discriminator

Generator

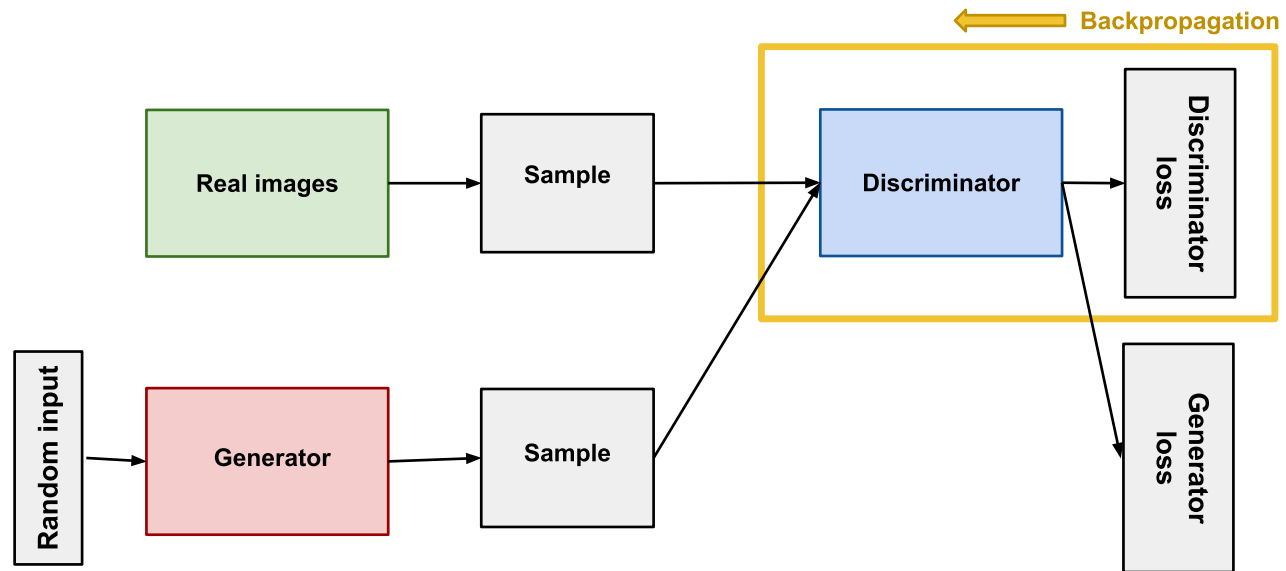


Figure 1: Backpropagation in discriminator training.

Discriminator Training Data

- ✓ The discriminator's training data comes from two sources:
- ✓ Real data instances, such as real pictures of people. The discriminator uses these instances as positive examples during training.
- ✓ Fake data instances created by the generator. The discriminator uses these instances as negative examples during training.

Training the Discriminator

1. The discriminator classifies both real data and fake data from the generator.
2. The discriminator loss penalizes the discriminator for misclassifying a real instance as fake or a fake instance as real.
3. The discriminator updates its weights through backpropagation from the discriminator loss through the discriminator network.

Generator

The generator part of a GAN learns to create fake data by incorporating feedback from the discriminator. It learns to make the discriminator classify its output as real.

Generator training requires tighter integration between the generator and the discriminator than discriminator training requires. The portion of the GAN that trains the generator includes:

- 1) Random input
- 2) Generator network, which transforms the random input into a data instance
- 3) Discriminator network, which classifies the generated data
- 4) Discriminator output
- 5) Generator loss, which penalizes the generator for failing to fool the discriminator

Generator

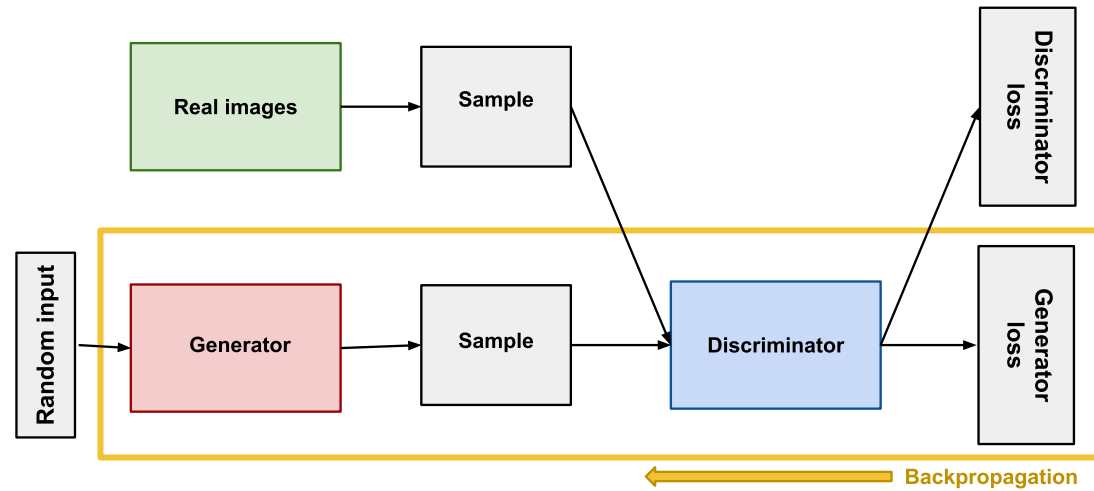


Figure 1: Backpropagation in generator training.

Using the Discriminator to Train the Generator

1. Sample random noise.
2. Produce generator output from sampled random noise.
3. Get discriminator "Real" or "Fake" classification for generator output.
4. Calculate loss from discriminator classification.
5. Backpropagate through both the discriminator and generator to obtain gradients.
6. Use gradients to change only the generator weights.

Workaround

<https://www.youtube.com/watch?v=vScjWYSFtmw>

<https://semiconductor.withgoogle.com/> - Active Participant

<https://artsandculture.google.com/experiment/viola-the-bird/nAEJVwNkp-FnrQ>

<https://artsandculture.google.com/experiment/paint-with-music/YAGuJyDB-XbbWg?cp=eyJhdG1vc3E>

<https://artsandculture.google.com/experiment/3d-pottery/nwHg1D0riJ1ItA>

<https://artsandculture.google.com/experiment/assisted-melody/cAGcgh18Zi7DqQ?cp=eyJ1cmwiOiI2M>

<https://artsexperiments.withgoogle.com/living-archive/map?token=1696070289&sequence=124525-1>

<https://creatability.withgoogle.com/clarion-lite/>

<https://creatability.withgoogle.com/body-synth/>

<https://quickdraw.withgoogle.com/>

<https://www.autodraw.com/>

<https://experiments.withgoogle.com/ai/ai-duet/view/>

<https://experiments.withgoogle.com/ai/beat-blender/view/>

<https://experiments.withgoogle.com/interplay-mode/view/>

<https://magic-sketchpad.glitch.me/>

<https://huggingface.co/spaces/AlphaDragon/Voice-Clone>

https://huggingface.co/spaces/mithril-security/blind_chat

<https://huggingface.co/spaces/facebook/MusicGen>

Workaround

<https://teachablemachine.withgoogle.com/train/image>

https://colab.research.google.com/drive/19haFgvhNwQQYGdqvoV9X_WLZjmlgM4Oi#scrollTo=83-azWpoYsDg

https://colab.research.google.com/github/google/eng-edu/blob/main/ml/cc/exercises/linear_regression_with_synthetic_data.ipynb?utm_source=mlcc&utm_campaign=colab-external&utm_medium=referral&utm_content=linear_regression_synthetic_tf2-colab&hl=en#scrollTo=_GMGgR6054IN

https://colab.research.google.com/github/google/eng-edu/blob/main/ml/cc/exercises/linear_regression_with_a_real_dataset.ipynb?utm_source=mlcc&utm_campaign=colab-external&utm_medium=referral&utm_content=linear_regression_real_tf2-

References

Theoretical:

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

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

<https://ai.google/>



UNIVERSITY OF TARTU



unitartu



tartuuniversity

