Generative Adversarial Networks

(GANs)

Prof. Kuan-Ting Lai 2022/6/6 CONTACTO TO LOS

VERIE DI STE II

MEDI SLAH CLARKE

Mare an everyment

Chief Inc. of Chief

Constrate and the second second

COMPACT OF ALL STREET, STREET, ST. L.

Star Gauge Children Star Star

DeepFake (Intro)



This Person does not Exist (thispersondoesnotexist.com)



Father of GAN: Ian Goodfellow



"The generative model can be thought of as analogous to a team of counterfeiters, trying to produce fake currency and use it without detection, while the discriminative model is analogous to the police, trying to detect the counterfeit currency.

Competition in this game drives both teams to improve their methods until the counterfeits are indistinguishable from the genuine articles." <u>Goodfellow et. al.</u>



Generative Adversarial Networks (GAN)

There are many interesting recent development in deep learning...The most important one, in my opinion, is adversarial training (also called GAN for Generative Adversarial Networks). This, and the variations that are now being proposed, is the most interesting idea in the last 10 years in ML. Vann LeCun

- <u>https://www.youtube.com/watch?v=9JpdAg6uMXs</u>
- <u>https://arxiv.org/abs/1701.00160</u>

What is a Generative Model?

- Informally:
 - Generative models can generate new data instances.
 - **Discriminative** models discriminate between different kinds of data instances.
- Formally
 - 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).

Training Generative Models is Hard!

- Discriminative Model
- Generative Model





https://developers.google.com/machine-learning/gan/generative

Generative Adversarial Networks (GAN)



GAN Structure



https://developers.google.com/machine-learning/gan/gan_structure

The Discriminator

- The discriminator loss penalizes the discriminator for misclassifying a real instance as fake or a fake instance as real.
- The discriminator updates its weights through backpropagation



The Generator



https://developers.google.com/machine-learning/gan/generator

The Generator Training Steps

- 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.

GAN Training

• Alternating Training

- Train the discriminator for one or more epochs.
- Train the generator for one or more epochs.
- Repeat steps 1 and 2 to continue to train the generator and discriminator networks.

• Convergence

- The discriminator performance gets worse
- If the generator succeeds perfectly, then the discriminator has a 50% accuracy.

GAN Loss Functions

• Minimax Loss

– The loss function used in the paper that introduced GANs.

- Wasserstein Loss
 - The default loss function for TF-GAN Estimators. First described in a <u>2017 paper</u>.

Minimax Loss

- D(x) is the discriminator's estimate of the probability that the real instance x is real.
- E_x is the expected value over all real data instances.
- G(z) is the generator's output when given noise z.
- D(G(z)) is the discriminator's estimate of the probability that a fake instance is real.
- E_z is the expected value over all generated fake instances G(z).
- The formula derives from the cross-entropy between the real and generated distributions.

$$\min_{\theta_g} \max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z))) \right]$$

https://developers.google.com/machine-learning/gan/loss

Wasserstein Loss

- In Wasserstein GAN, discriminator does not actually classify instances. The output does not have to be between 0 and 1
- Discriminator training just tries to make the output larger for real instances than for fake instances.
- Critic Loss: D(x) D(G(z))
 - The discriminator tries to maximize the difference between its output on real instances and fake instances.

• Generator Loss: D(G(z))

- The generator tries to maximize the discriminator's output for its fake instances.
- Use earth mover distance

Common Problems of Training GAN Networks

• Strong Discriminator

 if the discriminator is too good, then generator training can fail due to vanishing gradients

Mode Collapse

The generator may learn to produce only one output

• Failure to Converge

- Solution 1: adding noise to discriminator inputs
- Solution 2: penalizing discriminator weights

Bag of Tricks for Training GANs

- Use tanh as the last activation in the generator, instead of sigmoid
- Sample points from the latent space using a normal distribution
- Stochasticity is good to induce robustness. Introducing randomness during training helps prevent GAN to get stuck.
 - Use dropout in the discriminator
 - Add some random noise to the labels for the discriminator.
- Sparse gradients can hinder GAN training. There are two things that can induce gradient sparsity: 1) max pooling, 2) ReLU activations.
 - Use strided convolutions for downsampling
 - Use LeakyReLU, which allows small negative activation values.
- In generated images, it is common to see "checkerboard artifacts" caused by unequal coverage of the pixel space in the generator.
 - Use a kernel size that is divisible by the stride size

Checkerboard Artifacts

- Caused by unequal coverage of the pixel space in the generator
- Solution: Use a kernel size that is divisible by the stride size in *Conv2D* and *Conv2DTranspose*







Training GAN for Celebrity Faces

Code: <u>https://github.com/fchollet/deep-learning-with-python-notebooks/blob/master/chapter12_part05_gans.ipynb</u>
 Dataset: CelebA Dataset (cuhk.edu.hk)

Loading Images from Directory

• Unzip the dataset into "celeba_gan" folder

```
dataset = dataset.map(lambda x: x / 255.)
```

Generator

```
latent dim = 128
                                                 The latent space
                                                 will be made of 128-
           generator = keras.Sequential(
                                                 dimensional vectors.
   Revert the
                    keras.Input(shape=(latent dim,)),
Flatten layer of
                    layers.Dense(8 * 8 * 128),
  the encoder.
                    layers.Reshape((8, 8, 128)),
                    layers.Conv2DTranspose(128, kernel_size=4, strides=2, padding="same"),
                    layers.LeakyReLU(alpha=0.2),
   Revert the
Conv2D layers
                    layers.Conv2DTranspose(256, kernel size=4, strides=2, padding="same"),
of the encoder.
                    layers.LeakyReLU(alpha=0.2),
                    layers.Conv2DTranspose(512, kernel size=4, strides=2, padding="same"),
                    layers.LeakyReLU(alpha=0.2),
                    layers.Conv2D(3, kernel size=5, padding="same", activation="sigmoid"),
                ],
                                                                                  We use LeakyReLU
                name="generator",
                                                                                   as our activation.
```

Discriminator

```
from tensorflow.keras import layers
discriminator = keras.Sequential(
```

],

```
keras.Input(shape=(64, 64, 3)),
   layers.Conv2D(64, kernel_size=4, strides=2, padding="same"),
    layers.LeakyReLU(alpha=0.2),
    layers.Conv2D(128, kernel_size=4, strides=2, padding="same"),
    layers.LeakyReLU(alpha=0.2),
    layers.Conv2D(128, kernel_size=4, strides=2, padding="same"),
    layers.LeakyReLU(alpha=0.2),
    layers.Flatten(),
    layers.Dropout(0.2), # Important Trick!
    layers.Dense(1, activation="sigmoid"),
name="discriminator",
```

GAN Training Loop

- 1. Draw random points in the latent space (random noise).
- 2. Generate images with generator using this random noise.
- 3. Mix the generated images with real ones.
- 4. Train discriminator using these mixed images including real and fake (generated) images.
- 5. Fix the discriminator's weights
- 6. Draw new random points in the latent space.
- 7. Train the generator to fool the discriminator

The GAN Model

```
import tensorflow as tf
class GAN(keras.Model):
   def init (self, discriminator, generator, latent dim):
        super(). init ()
        self.discriminator = discriminator
        self.generator = generator
        self.latent dim = latent dim
        self.d loss metric = keras.metrics.Mean(name="d loss")
        self.g loss metric = keras.metrics.Mean(name="g loss")
    def compile(self, d_optimizer, g_optimizer, loss_fn):
        super(GAN, self).compile()
        self.d optimizer = d optimizer
        self.g optimizer = g optimizer
        self.loss_fn = loss_fn
   @property
   def metrics(self):
        return [self.d loss metric, self.g loss metric]
```

Sets up metrics to track the two losses over each training epoch

train_step() of the GAN Model (2-1)

```
def train step(self, real images):
                batch size = tf.shape(real images)[0]
                                                                 Samples random points
  Decodes
                random latent vectors = tf.random.normal(
                                                                 in the latent space
  them to
                     shape=(batch_size, self.latent dim))
fake images
                generated images = self.generator(random latent vectors)
                combined images = tf.concat([generated images, real images], axis=0)
 Combines
                labels = tf.concat(
                                                                                  Assembles labels,
 them with
                                                                                  discriminating real
                     [tf.ones((batch size, 1)), tf.zeros((batch size, 1))],
real images
                                                                                   from fake images
                     axis=0
                labels += 0.05 * tf.random.uniform(tf.shape(labels))
                                                                               Adds random
                                                                               noise to the
                with tf.GradientTape() as tape:
                                                                               labels-an
                    predictions = self.discriminator(combined images)
                                                                               important trick!
                    d loss = self.loss fn(labels, predictions)
     Trains the
                grads = tape.gradient(d loss, self.discriminator.trainable weights)
  discriminator
                self.d optimizer.apply gradients(
                     zip(grads, self.discriminator.trainable weights)
```

train_step() of the GAN Model (2-2)

Samples random points in the latent space

```
shape=(batch_size, self.latent_dim))
misleading labels = tf.zeros((batch size, 1))
```

random latent vectors = tf.random.normal(

Assembles labels that say "these are all real images" (it's a lie!)

```
with tf.GradientTape() as tape:
    predictions = self.discriminator(
        self.generator(random_latent_vectors))
    g_loss = self.loss_fn(misleading_labels, predictions)
    grads = tape.gradient(g_loss, self.generator.trainable_weights)
    self.g_optimizer.apply_gradients(
        zip(grads, self.generator.trainable_weights))
```

```
self.d_loss_metric.update_state(d_loss)
self.g_loss_metric.update_state(g_loss)
return {"d_loss": self.d_loss_metric.result(),
            "g_loss": self.g_loss_metric.result()}
```

Monitor the Training

```
class GANMonitor(keras.callbacks.Callback):
   def __init__(self, num_img=3, latent_dim=128):
        self.num_img = num_img
        self.latent_dim = latent_dim
   def on_epoch_end(self, epoch, logs=None):
        random_latent_vectors = tf.random.normal(shape=(self.num img,
                                                            self.latent dim))
        generated_images = self.model.generator(random_latent_vectors)
        generated_images *= 255
        generated images.numpy()
        for i in range(self.num img):
            img = keras.utils.array to img(generated images[i])
            img.save(f"generated img {epoch:03d} {i}.png")
```

Compiling and Training GAN

```
epochs = 100
```

```
gan.fit(
    dataset, epochs=epochs, callbacks=[GANMonitor(num_img=10,
        latent_dim=latent_dim)]
```

Training Results

Epoch 0



Epoch 21



Epoch 1



Epoch 26



Epoch 9



Epoch 27



Epoch 14



Epoch 28



Interesting GAN Applications

OUTPUT





pix2pix



Deepfake Round Table

https://www.creativeblog.com/features/deepfake-examples

This Person (and the Creepy Person) do not Exist!





thispersondoesnotexist.com

Evolution of GAN Face Generation



https://spectra.mathpix.com/article/2021.09.00009/gans

How to Detect Generated Faces?

- Detect fake faces by central positioning the eyes
- Stanford University researchers identify fake LinkedIn profiles using eye locations



https://www.capradio.org/news/npr/story?storyid=1088140809

Anime GAN (<u>https://make.girls.moe/</u>)





https://arxiv.org/abs/1708.05509

https://heartbeat.comet.ml/my-mangagan-building-my-first-generative-adversarial-network-2ec1920257e3

https://make.girls.moe/

MakeGirlsMoe Home History Tr	ransition Help v	E	nglish 🔻 Twitter Discord
	Options Defined Mode Model Camellia 256x256 Ver.171219 (9.9MB)		•
Generate	Hair Color Blonde	Hair Style	Eye Color Green
	Dark Skin Off Random On	Blush Off Random On	Smile Off Random On
	Open Mouth Off <mark>Random</mark> On	Hat Off Random On	Ribbon Off Random On
I(ウ+1 I(ウ-1	Glasses Off Random On	Style	
Share on Twitter	Noise Random Fixed	Current Noise	Noise Import/Export
	Operations Import Export Reset	WebGL Acceleration ③ Disabled Enabled	

Current Backend

WILCI

Super Resolution GAN

 Ledig et al., "Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network," 2017 (<u>https://arxiv.org/pdf/1609.04802.pdf</u>)



More GAN Models

- TensorFlow Tutorial / Generative
- 1. Pixel-2-Pixel
- 2. CycleGAN
- 3. Adversarial FGSM



• Phillip Isola et al., Image-to-Image Translation with Conditional Adversarial Networks, 2018



Training Conditional GAN

- Both the generator and discriminator observe the input edge map
- Use U-Net and PatchGAN discriminator



Image-to-Image Demo

<u>https://affinelayer.com/pixsrv/</u>



My Work (Cat Computer)



CycleGAN

- Zhu et al., <u>Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks</u>, 2018
- Learn to automatically "translate" an image from one into the other and vice versa



Photograph

Monet

Van Gogh

Cezanne

Ukiyo-e

Model of CycleGAN

- Two mapping functions $G : X \rightarrow Y$ and $F : Y \rightarrow X$
- Two cycle consistency losses:
 - Forward cycle-consistency loss: $x \rightarrow G(x) \rightarrow F(G(x)) \approx x$
 - Backward cycle-consistency loss: $y \rightarrow F(y) \rightarrow G(F(y)) \approx y$



TensorFlow CycleGAN Results

Input Image



Input Image



Predicted Image



Predicted Image



Input Image



Input Image



Predicted Image



Predicted Image





NVIDIA AI Playground



Al Demo

Paint Me a Picture: GauGAN2 Al Art Tool See Al Art in New Dimensions with Fresh Work from 4 Artists

Al in Art

Al Demo

Take Conversations Global with World-Class Speech Al

My Work by GauGAN 2



Adversarial Attack

- Goodfellow et al., Explaining and Harnessing Adversarial Examples, 2015
- Fast Gradient Signed Method (FGSM)



https://www.tensorflow.org/tutorials/generative/adversarial_fgsm

Fooling AI Surveillance Cameras



https://www.arxiv-vanity.com/papers/1904.08653/

Key Takeaways

- Generative Adversarial Networks has two sub-networks: Generator and Discriminator. They compete with each other.
- Training is done if the Discriminator fails to detect fake image. (ACC=0.5)
- GANs are hard to train. Introducing randomness during training is important.
- Max pooling and ReLu introduces sparsity and hinder GAN training
- Conditional GANs can accept conditional images as inputs than generate target outputs
- CycleGAN can translate image from one domain to the other and vice versa.

References

- <u>https://www.tensorflow.org/tutorials/generative/</u>
- <u>https://developers.google.com/machine-learning/gan</u>
- https://make.girls.moe/
- <u>https://www.creativebloq.com/features/deepfake-examples</u>
- <u>https://www.nvidia.com/en-us/research/ai-playground/</u>
- <u>https://www.projectpro.io/article/generative-adversarial-networks-gan-based-projects-to-work-on/530</u>
- <u>https://jonathan-hui.medium.com/gan-some-cool-applications-of-gans-</u> <u>4c9ecca35900</u>
- <u>https://affinelayer.com/pixsrv/</u>