

## GENERATIVE ADVERSARIAL NETWORKS

Vincent Barra

LIMOS, UMR 6158 CNRS, Université Clermont Auvergne



# LATENT MODELS

A latent variable model relates a set of observable variables  $\mathbf{x} \in X$  to a set of latent variables  $\mathbf{h} \in H$

$$\mathbf{p}(\mathbf{x}, \mathbf{h}) = \mathbf{p}(\mathbf{x}|\mathbf{h})\mathbf{p}(\mathbf{h})$$

if  $\mathbf{h}$  are causal factors for  $\mathbf{x} \Rightarrow$  sampling from  $\mathbf{p}(\mathbf{x}|\mathbf{h}) =$  generative process from  $H$  to  $X$ .

## Inference

Inference: given  $\mathbf{p}(\mathbf{x}, \mathbf{h})$ , compute

$$\mathbf{p}(\mathbf{h}|\mathbf{x}) = \frac{\mathbf{p}(\mathbf{x}|\mathbf{h})\mathbf{p}(\mathbf{h})}{\mathbf{p}(\mathbf{x})}$$

Intractable

# GAN

## GAN

- ▶ Generative Adversarial Network
- ▶ Two-player game between a discriminator  $D$  and a Generator  $G$
- ▶  $G$  and  $D$ : neural networks

# GAN

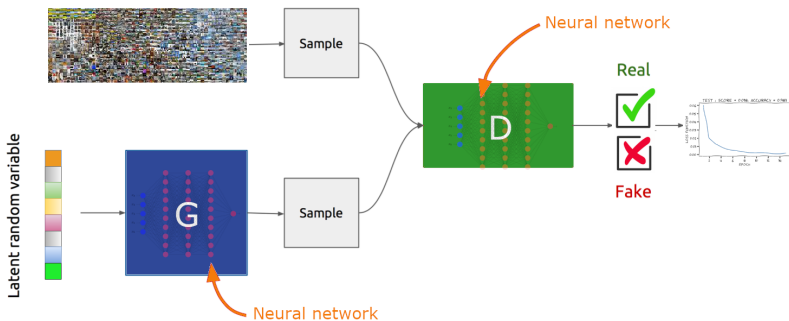
## GAN

- ▶ Generative Adversarial Network
- ▶ Two-player game between a discriminator  $D$  and a Generator  $G$
- ▶  $G$  and  $D$ : neural networks

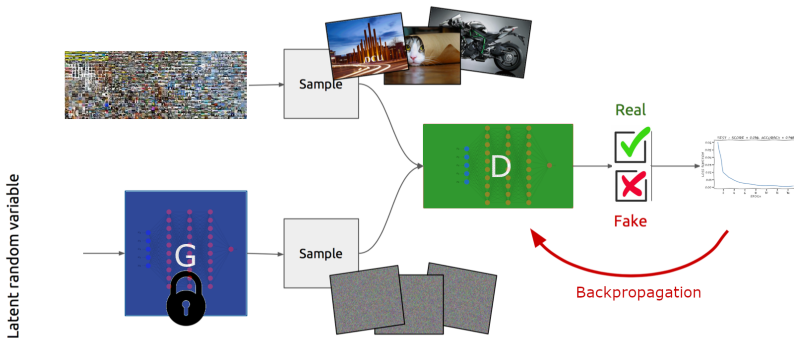
## Game

- ▶  $G$  tries to generate synthetic data close to real ones, and aims at fooling  $D$
- ▶  $D$  tries to discriminate between real and fake images.
- ▶ Adversarial:  $G$  and  $D$  have antagonistic objectives.

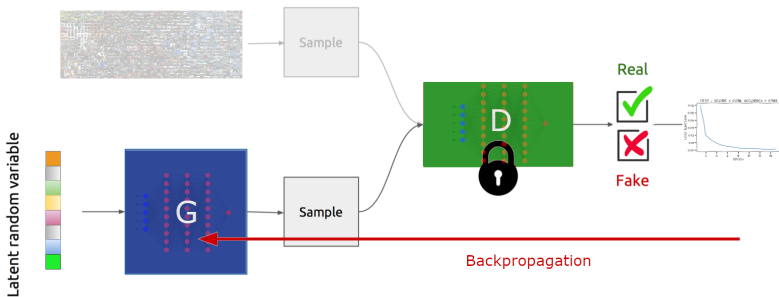
# OVERVIEW



## OVERVIEW



## OVERVIEW



# GENERATOR AND DISCRIMINATOR

$$G : \mathbb{R}^d \rightarrow \mathbb{R}^n$$

- ▶ Function from the latent space to the data space
- ▶ MLP, CNN, RNN...
- ▶ trained so that for  $\mathbf{h} \in \mathbb{R}^d$ ,  $G(\mathbf{h}) \sim \mathbf{p}_{data}$



# GENERATOR AND DISCRIMINATOR

$$G : \mathbb{R}^d \rightarrow \mathbb{R}^n$$

- ▶ Function from the latent space to the data space
- ▶ MLP, CNN, RNN...
- ▶ trained so that for  $\mathbf{h} \in \mathbb{R}^d$ ,  $G(\mathbf{h}) \sim \mathbf{p}_{data}$

$$D : \mathbb{R}^n \rightarrow [0, 1]$$

- ▶ Function from the data space producing a probability
- ▶ MLP, CNN, RNN...
- ▶ trained so that for  $\mathbf{x} \in \mathbb{R}^n$ , predicts if  $\mathbf{x} = G(\mathbf{h})$  or real data

# TRAINING

If  $G$  is fixed, training  $D$  is easy

- 1 pick up  $p$  real data  $\mathbf{x}_1 \cdots \mathbf{x}_p \in \mathbb{R}^n$
- 2 generate  $p$  fakes  $G(\mathbf{h}_i), i \in \llbracket 1, p \rrbracket, \mathbf{h}_i \sim \mathbf{p}_G$
- 3 build a training set =  $Z = \{(G(\mathbf{h}_i), 0), (\mathbf{x}_i, 1), i \in \llbracket 1, p \rrbracket\}$
- 4 train  $D$  by minimizing the binary cross-entropy

$$\begin{aligned} \mathcal{L}(Z) &= -\frac{1}{2p} \left( \sum_{i=1}^p [\log D(\mathbf{x}_i) + \log(1 - D(G(\mathbf{h}_i)))] \right) \\ &= -\frac{1}{2} (\mathbb{E}_{X \sim \mathbf{p}_{data}} \log(D(X)) + \mathbb{E}_{X \sim \mathbf{p}_G} \log(1 - D(X))) \end{aligned}$$

# TRAINING

If  $G$  is fixed, training  $D$  is easy

- 1 pick up  $p$  real data  $\mathbf{x}_1 \cdots \mathbf{x}_p \in \mathbb{R}^n$
- 2 generate  $p$  fakes  $G(\mathbf{h}_i), i \in \llbracket 1, p \rrbracket, \mathbf{h}_i \sim \mathbf{p}_G$
- 3 build a training set =  $Z = \{(G(\mathbf{h}_i), 0), (\mathbf{x}_i, 1), i \in \llbracket 1, p \rrbracket\}$
- 4 train  $D$  by minimizing the binary cross-entropy

$$\begin{aligned} \mathcal{L}(Z) &= -\frac{1}{2p} \left( \sum_{i=1}^p [\log D(\mathbf{x}_i) + \log(1 - D(G(\mathbf{h}_i)))] \right) \\ &= -\frac{1}{2} (\mathbb{E}_{X \sim \mathbf{p}_{data}} \log(D(X)) + \mathbb{E}_{X \sim \mathbf{p}_G} \log(1 - D(X))) \end{aligned}$$

- ▶  $\text{Min} - \log D(\mathbf{x}_i)$ : maximizing the recognition of true data
- ▶  $\text{Min} - \log(1 - D(G(\mathbf{h}_i)))$ : maximizing the recognition of fake data

# TRAINING

But...  $G$  wants to fool  $D$  and has to be optimized to maximize  $D$ 's loss.

## Loss of $G$

$$\mathcal{L}_G(D, G) = \mathbb{E}_{X \sim \mathbf{p}_{data}} \log(D(X)) + \mathbb{E}_{X \sim \mathbf{p}_G} \log(1 - D(X))$$

- ▶ high if  $D$  is right
- ▶ low if  $G$  fools  $D$  often

# TRAINING

But...  $G$  wants to fool  $D$  and has to be optimized to maximize  $D$ 's loss.

## Loss of $G$

$$\mathcal{L}_G(D, G) = \mathbb{E}_{X \sim \mathbf{p}_{data}} \log(D(X)) + \mathbb{E}_{X \sim \mathbf{p}_G} \log(1 - D(X))$$

- ▶ high if  $D$  is right
- ▶ low if  $G$  fools  $D$  often

## Find an optimal generator $G^*$ fooling any $D$

$$G^* = \arg \min_G \max_D \mathcal{L}_G(D, G) = \arg \min_G \mathcal{L}_G(D_G^*, G)$$

where  $D_G^* = \arg \max_D \mathcal{L}_G(D, G)$

⇒ Find a generator  $G$  whose loss against the best discriminator is low.

# TRAINING

## Nash equilibrium

$$(\forall \mathbf{x}) D_G^*(\mathbf{x}) = \frac{\mathbf{p}_{data}(\mathbf{x})}{\mathbf{p}_{data}(\mathbf{x}) + \mathbf{p}_G(\mathbf{x})}$$

and thus

$$\begin{aligned} \mathcal{L}_G(D_G^*, G) &= \mathbb{E}_{X \sim \mathbf{p}_{data}} \log(D_G^*(X)) + \mathbb{E}_{X \sim \mathbf{p}_G} \log(1 - D_G^*(X)) \\ &= \mathbb{E}_{X \sim \mathbf{p}_{data}} \log\left(\frac{\mathbf{p}_{data}(X)}{\mathbf{p}_{data}(X) + \mathbf{p}_G(X)}\right) + \mathbb{E}_{X \sim \mathbf{p}_G} \log\left(\frac{\mathbf{p}_{latent}(X)}{\mathbf{p}_{data}(X) + \mathbf{p}_G(X)}\right) \\ &= KL\left(\mathbf{p}_{data} \parallel \frac{\mathbf{p}_{data} + \mathbf{p}_G}{2}\right) + KL\left(\mathbf{p}_G \parallel \frac{\mathbf{p}_{data} + \mathbf{p}_G}{2}\right) - \log(4) \\ &= 2JS(\mathbf{p}_{data}, \mathbf{p}_G) - \log(4) \quad (\text{JS: Jensen-Shannon divergence}) \end{aligned}$$

# TRAINING

In practice  $D$  is not fully optimized when optimizing  $G$   
 Alternating gradient step for  $G$  and  $D$

---

**Algorithm 1** Minibatch stochastic gradient descent training of generative adversarial nets. The number of steps to apply to the discriminator,  $k$ , is a hyperparameter. We used  $k = 1$ , the least expensive option, in our experiments.

---

**for** number of training iterations **do**

**for**  $k$  steps **do**

- Sample minibatch of  $m$  noise samples  $\{\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(m)}\}$  from noise prior  $p_g(\mathbf{z})$ .
- Sample minibatch of  $m$  examples  $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}\}$  from data generating distribution  $p_{\text{data}}(\mathbf{x})$ .
- Update the discriminator by ascending its stochastic gradient:

$$\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^m \left[ \log D(\mathbf{x}^{(i)}) + \log \left( 1 - D(G(\mathbf{z}^{(i)})) \right) \right].$$

**end for**

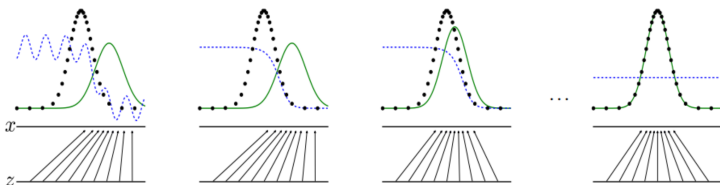
- Sample minibatch of  $m$  noise samples  $\{\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(m)}\}$  from noise prior  $p_g(\mathbf{z})$ .
- Update the generator by descending its stochastic gradient:

$$\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^m \log \left( 1 - D(G(\mathbf{z}^{(i)})) \right).$$

**end for**

Source: Goodfellow et al., 2014

## ILLUSTRATION



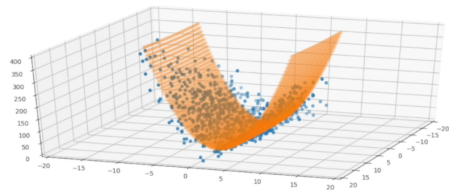
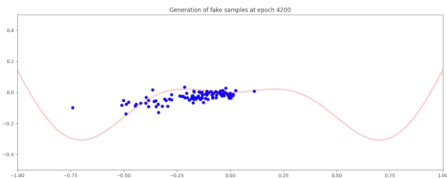
Source: Goodfellow et al., 2014

- ▶ black, dotted line : the data generating distribution
- ▶ green solid line : the generator distribution
- ▶ blue, dashed line : the discriminator
- ▶  $z$  is sampled uniformly over the domain described by the lower lines

- 1 initial values of the data,  $G$  and  $D$  distributions.
- 2 Convergence  $D \rightarrow D^*$
- 3  $G$  updating: gradient of  $D$  guides  $G(h)$  to high probability regions for the original data
- 4 After several epochs



# EXAMPLES



## EXAMPLES



## EXAMPLES



# TRAINING PROBLEMS

## 1 Oscillation between generator and discriminator loss

- ▶ competitive loss between  $G$  and  $D$
- ▶ no guarantee that the loss will decrease

# TRAINING PROBLEMS

- 1 Oscillation between generator and discriminator loss
- 2 Mode collapse

- ▶  $G$  tries to fool  $D$ .
  - ▶ When  $G$  is trained without updating  $D$ ,  $G$  produces mode  $x^*$  that fools  $D$  the most
  - ▶ When  $D$  is training, the most effective way to detect generated images is to detect this single mode
- ⇒ generator produces examples of a particular kind only
- ▶ hard problem to solve
  - ▶ one possible solution: Wasserstein loss (see next)

# TRAINING PROBLEMS

- 1 Oscillation between generator and discriminator loss
- 2 Mode collapse
- 3 Discriminator is too strong, such that the gradient for the generator vanishes and the generator can't keep up

- ▶ fake samples can be initially so "fake" that the response of  $D$  saturates
- ▶  $\log(1 - D(X))$  far in the exponential tail of the sigmoid of  $D \Rightarrow$  null gradient
- ▶  $\mathbb{E}_{X \sim p_G} \log(1 - D(X)) \longrightarrow -\mathbb{E}_{X \sim p_G} \log(D(X))$

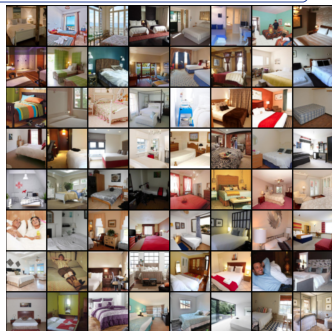
# A COOKING RECIPE FOR IMAGE GENERATION

“After extensive model exploration we identified a family of architectures that resulted in stable training across a range of datasets and allowed for training higher resolution and deeper generative models”

Radford et al., 2015

## DC-GAN

- ▶ pooling layers in  $D$  → strided convolutions
- ▶ pooling layers in  $G$  → strided transposed convolutions in  $G$
- ▶ use batchnorm in both  $D$  and  $G$
- ▶ remove fully connected hidden layers
- ▶ use ReLU in  $G$  except for the output ( $\tanh$ )
- ▶ use LeakyReLU activation in  $D$  for all layers.



Real bedrooms

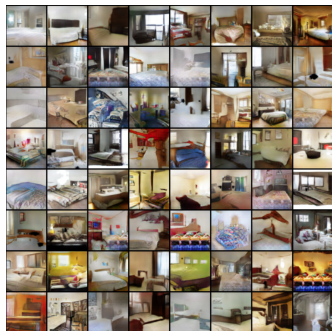
# A COOKING RECIPE FOR IMAGE GENERATION

“After extensive model exploration we identified a family of architectures that resulted in stable training across a range of datasets and allowed for training higher resolution and deeper generative models”

Radford et al., 2015

## DC-GAN

- ▶ pooling layers in  $D$  → strided convolutions
- ▶ pooling layers in  $G$  → strided transposed convolutions in  $G$
- ▶ use batchnorm in both  $D$  and  $G$
- ▶ remove fully connected hidden layers
- ▶ use ReLU in  $G$  except for the output ( $\tanh$ )
- ▶ use LeakyReLU activation in  $D$  for all layers.



epoch 1



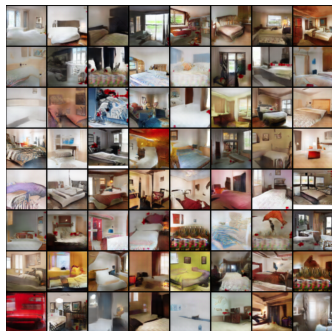
## A COOKING RECIPE FOR IMAGE GENERATION

“After extensive model exploration we identified a family of architectures that resulted in stable training across a range of datasets and allowed for training higher resolution and deeper generative models”

Radford et al., 2015

### DC-GAN

- ▶ pooling layers in  $D$  → strided convolutions
- ▶ pooling layers in  $G$  → strided transposed convolutions in  $G$
- ▶ use batchnorm in both  $D$  and  $G$
- ▶ remove fully connected hidden layers
- ▶ use ReLU in  $G$  except for the output ( $\tanh$ )
- ▶ use LeakyReLU activation in  $D$  for all layers.



epoch 5

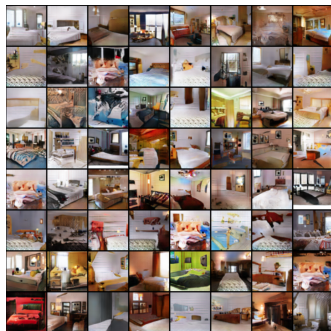
# A COOKING RECIPE FOR IMAGE GENERATION

“After extensive model exploration we identified a family of architectures that resulted in stable training across a range of datasets and allowed for training higher resolution and deeper generative models”

Radford et al., 2015

## DC-GAN

- ▶ pooling layers in  $D$  → strided convolutions
- ▶ pooling layers in  $G$  → strided transposed convolutions in  $G$
- ▶ use batchnorm in both  $D$  and  $G$
- ▶ remove fully connected hidden layers
- ▶ use ReLU in  $G$  except for the output ( $\tanh$ )
- ▶ use LeakyReLU activation in  $D$  for all layers.



epoch 10

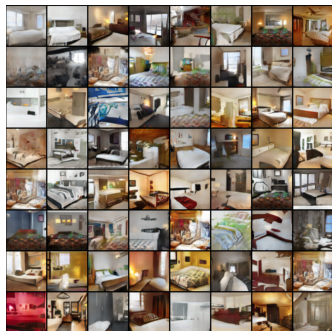
## A COOKING RECIPE FOR IMAGE GENERATION

“After extensive model exploration we identified a family of architectures that resulted in stable training across a range of datasets and allowed for training higher resolution and deeper generative models”

Radford et al., 2015

### DC-GAN


- ▶ pooling layers in  $D$  → strided convolutions
- ▶ pooling layers in  $G$  → strided transposed convolutions in  $G$
- ▶ use batchnorm in both  $D$  and  $G$
- ▶ remove fully connected hidden layers
- ▶ use ReLU in  $G$  except for the output ( $\tanh$ )
- ▶ use LeakyReLU activation in  $D$  for all layers.



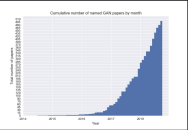
epoch 20

Since 2014, A LOT of models have been developed.

### The GAN Zoo



Every week, new GAN papers are coming out and it's hard to keep track of them all, not to mention the incredibly creative ways in which researchers are naming these GANs! So, here's a list of what started as a fun activity compiling all named GANs!



You can also check out the same data in a tabular format with functionality to filter by year or do a quick search by title [here](#).

Contributions are welcome. Add links through pull requests in `gans.tsv` file in the same format or create an issue to let me know something I missed or to start a discussion.

Check out [Deep Hunt](#) - my weekly AI newsletter for this repo as blogpost and follow me on [Twitter](#).

- [3D-ED-GAN](#) - Shape inpainting using 3D Generative Adversarial Network and Recurrent Convolutional Networks
- [3D-GAN](#) - Learning a Probabilistic Latent Space of Object Shapes via 3D Generative-Adversarial Modeling ([github](#))
- [3D-IWGAN](#) - Improved Adversarial Systems for 3D Object Generation and Reconstruction ([github](#))
- [3D-PhysNet](#) - 3D-PhysNet: Learning the Intuitive Physics of Non-Rigid Object Deformations

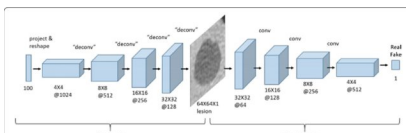
<https://github.com/hindupuravinash/the-gan-zoo>

# SOME MODELS

## Some models

A (non exhaustive) selection:

- ▶ DCGAN: from MLP to CNN (Radford).



# SOME MODELS

## Some models

A (non exhaustive) selection:

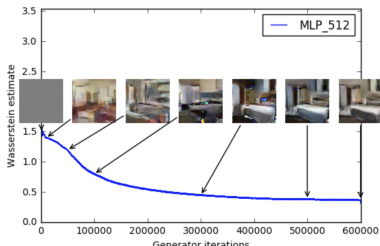
- ▶ WGAN: Wasserstein GAN .
  - ▶ JS divergence → Earth-Mover's distance:

$$W(\mathbf{p}_{data}, \mathbf{p}_G) = \inf_{\gamma \in \Pi(\mathbf{p}_{data}, \mathbf{p}_G)} \mathbb{E}_{x, y \sim \gamma} [\|x - y\|]$$

$\Pi(\mathbf{p}_{data}, \mathbf{p}_G)$ : set of all joint distributions  $\gamma(x, y)$  whose marginals are  $\mathbf{p}_{data}$  and  $\mathbf{p}_G$

⇒  $\gamma(x, y)$ : how much "mass" must be transported from  $x$  to  $y$  in order to transform  $\mathbf{p}_{data}$  to  $\mathbf{p}_G$ .

- ▶ More stable training
- ▶ less mode collapsing



## SOME MODELS

## Some models

A (non exhaustive) selection:

- ▶ CGAN: Conditional GAN

$$\mathcal{L}_G(D, G) = \mathbb{E}_{x \sim \mathbf{p}_{data}} \log(D(x|y)) + \mathbb{E}_{h \sim \mathbf{p}_G} \log(1 - D(G(h|y)), y)$$



## SOME MODELS

## Some models

A (non exhaustive) selection:

- ▶ InfoGAN: CGAN trained in an unsupervised way = GAN + maximization of the mutual information between a small subset of the latent variables and the observation

$$\min_g \max_D \mathcal{L}_G(D, G) - \lambda I(y|G(h, y)), \quad I(X, Y) = H(X) - H(X|Y)$$

(a) Varying  $c_1$  on InfoGAN (Digit type)(b) Varying  $c_1$  on regular GAN (No clear meaning)(c) Varying  $c_2$  from -2 to 2 on InfoGAN (Rotation)(d) Varying  $c_2$  from -2 to 2 on InfoGAN (Width)



## SOME MODELS

### Some models

A (non exhaustive) selection:

- ▶ BiGAN:  $G$  maps from  $H$  to  $X$  and from  $X$  to  $H \Rightarrow$  Adversarial Feature Learning

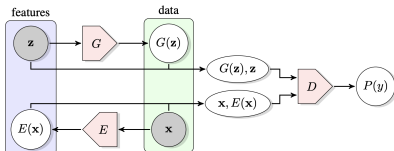


Figure 1: The structure of Bidirectional Generative Adversarial Networks (BiGAN).

Source: Donahue et al., 2017

# SOME MODELS

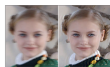
## Some models

A (non exhaustive) selection:

- ▶ Paired or unpaired image to image translation:  $h$  is replaced by an image (SRGAN, Pix2Pix, SimGAN, CycleGAN, DiscoGAN, CoVAE-GAN...)



Low-res to high-res



Blurry to sharp



Thermal to color



Synthetic to real



LDR to HDR



Noisy to clean



Image to painting



Day to night



Summer to winter

- Bad weather to good weather
- Greyscale to color
- ...

Source: CVPR 2017 Tutorial

```

#Define generator
def G():
    g = Sequential()
    ...

    return g

#Define discriminator
def D():
    d = Sequential()
    ...

    return d

#Define loss function

g.compile(loss='binary_crossentropy', optimizer=adam, metrics=['accuracy'])
d.compile(loss='binary_crossentropy', optimizer=adam, metrics=['accuracy'])

#Define GAN
inputs = Input(shape=(z_dim, ))
h = G(inputs)
output = D(h)
gan = Model(inputs, output)
gan.compile(loss='binary_crossentropy', optimizer=adam, metrics=['accuracy'])

```

```

for e in range(1, epochs+1):
    for _ in range(batchCount):
        # batch creation
        image_batch = x_train[np.random.randint(0, x_train.shape[0], size=BATCH_SIZE)]
        # Latent variables
        noise = np.random.normal(0, 1, size=(BATCH_SIZE, z_dim))

        # Fake images
        generated_images = g.predict(noise)
        X = np.concatenate((image_batch, generated_images))
        # Labels
        y = np.zeros(2*BATCH_SIZE)
        y[:BATCH_SIZE] = 1

        # Train discriminator
        d.trainable = True
        d_loss = d.train_on_batch(X, y)

        # Train generator
        noise = np.random.normal(0, 1, size=(BATCH_SIZE, z_dim))
        y2 = np.ones(BATCH_SIZE)
        d.trainable = False
        g_loss = gan.train_on_batch(noise, y2)

```



```

# training
train_dataset = tf.data.Dataset.from_tensor_slices(x_train).shuffle(BUFFER_SIZE).batch(BATCH_SIZE)
for epoch in range(epochs):

    # for batches
    for images in train_dataset:
        # latent variables
        noise = tf.random.normal([BATCH_SIZE, z_dim])

        # Gradient recording and updating
        with tf.GradientTape() as G_tape, tf.GradientTape() as D_tape:
            # Fake
            g_images = g(noise, training=True)

            # Output on real images
            real_output = d(images, training=True)

            # Output on fake images
            fake_output = d(g_images, training=True)

            # losses
            gen_loss = G_loss(fake_output)
            disc_loss = D_loss(real_output, fake_output)

        # Gradient computation
        G_gradients = G_tape.gradient(G_loss, g.trainable_variables)
        D_gradients = D_tape.gradient(D_loss, d.trainable_variables)

        # Updating
        G_optimizer.apply_gradients(zip(G_gradients, g.trainable_variables))
        D_optimizer.apply_gradients(zip(D_gradients, d.trainable_variables))

```