

# Generative Models

Your Name \*

Sihan Wei

What is the main difference between an **Autoencoder (AE)** and a **Variational Autoencoder (VAE)**?

1 point

- A. AE is generative, VAE is not
- B. AE learns deterministic mappings, VAE learns a probabilistic latent space
- C. VAE does not use a decoder
- D. AE requires labeled data

Which of the following are **reasons for using denoising autoencoders (DAEs)**?

1 point

- A. To improve robustness against noisy inputs
- B. To force the model to learn better feature representations
- C. To remove dropout from the network
- D. To reconstruct inputs from corrupted versions

In a **GAN**, the **discriminator** is trained to:

1 point

- A. Generate realistic samples
- B. Classify inputs as real or fake
- C. Minimize reconstruction loss
- D. Produce a latent code

In a **denoising autoencoder**, the corruption step is applied:

1 point

- A. During both training and testing
- B. Only during training
- C. Only during testing
- D. Never – DAE uses clean data only

Which of the following statements about **generative capabilities** is correct?

1 point

- A. A standard autoencoder is a generative model
- B. A VAE is a generative model
- C. A GAN is a generative model
- D. A DAE cannot generate samples without additional assumptions

In a **GAN**, what is the primary goal of the generator?

1 point

- A. Maximize classification accuracy
- B. Produce samples indistinguishable from real data to the discriminator
- C. Minimize KL divergence between real and fake data
- D. Encode data into a latent representation

Which of the following are common uses of autoencoders?

1 point

- A. Dimensionality reduction
- B. Image denoising
- C. Data compression
- D. Supervised classification

In a GAN's **minimax objective**, the discriminator tries to:

1 point

- A. Minimize  $-\log D(x)$  for real samples and  $-\log(1 - D(G(z)))$  for fake samples
- B. Maximize  $\log D(G(z))$  for fake samples
- C. Predict the latent code for each sample
- D. Maximize the generator's loss

In a standard autoencoder, what happens if the bottleneck dimension is too large?

1 point

- A. The model may overfit and learn the identity function
- B. The model will always underfit
- C. The KL divergence will explode
- D. The generator will fail to converge

Which of the following best describes the **latent variable**  $z$  in a VAE?

1 point

- A. It is a deterministic encoding of the input
- B. It is a random variable sampled from a learned distribution
- C. It is the output of the decoder
- D. It is the label assigned to the input

This content is neither created nor endorsed by Google.

Google Forms