

Deep Generative Models (AI314)

Master of Science in AI



Dr. Nibras Abdullah

*Assistant professor
Faculty of Informatics
Midocean University*



Lecture 1

What Is Generative Modeling?

Generative Versus Discriminative Modeling

The Rise of Generative Modeling

Types of Generative Models

Generative modeling can be broadly defined as follows:

Generative modeling is a branch of machine learning that involves training a model to produce new data that is similar to a given dataset.

What Is Generative Modeling?





Figure 1-1. A generative model trained to generate realistic photos of horses


What Is Generative Modeling?



What are Generative Models?

Generating New Data


Generative models are a class of machine learning algorithms that can learn to generate new data samples that are similar to the training data, such as images, text, or audio. They aim to capture the underlying patterns and distributions in the data.



What are Generative Models?

Learning Latent Representations


These models learn a latent, or hidden, representation of the input data, which can then be used to generate new samples. This allows them to capture the essential features and structure of the data in a compact form.



What are Generative Models?

Adversarial Training

Generative models like Generative Adversarial Networks (GANs) employ an adversarial training process, where a generator model competes with a discriminator model to produce increasingly realistic samples.



What are Generative Models?

A generative model must also be probabilistic rather than deterministic, because we want to be able to sample many different variations of the output, rather than get the same output every time.

If our model is merely a fixed calculation, such as taking the average value of each pixel in the training dataset, it is not generative.

Generative Versus Discriminative Modeling

In order to truly understand what generative modeling aims to achieve and why this is important, it is useful to compare it to its counterpart, discriminative modeling.

Generative Versus Discriminative Modeling




Figure 1-2. A discriminative model trained to predict if a given image is painted by Van Gogh

Generative Versus Discriminative Modeling




Figure 1-2. A discriminative model trained to predict if a given image is painted by Van Gogh

Generative modeling doesn't require the dataset to be labeled because it concerns itself with generating entirely new images, rather than trying to predict a label of a given image.

Generative Versus Discriminative Modeling

Let's define these types of modeling formally, using mathematical notation:

Discriminative modeling estimates $p(y | \mathbf{x})$.

That is, discriminative modeling aims to model the probability of a label y given some observation $\mathbf{x}$.

Generative Versus Discriminative Modeling

Let's define these types of modeling formally, using mathematical notation:

Generative modeling estimates $p(\mathbf{x})$.

That is, generative modeling aims to model the probability of observing an observation $\mathbf{x}$. Sampling from this distribution allows us to generate new observations.

Generative Versus Discriminative Modeling

Let's define these types of modeling formally, using mathematical notation:

Conditional Generative Models

Note that we can also build a generative model to model the conditional probability $p(\mathbf{x} | y)$ —the probability of seeing an observation $\mathbf{x}$ with a specific label y .

The Rise of Generative Modeling



The Rise of Generative Modeling





Figure 1-3. Face generation using generative modeling has improved significantly over the last decade (adapted from *Brundage et al., 2018*)₁

Generative Artificial Intelligence



Generative Modeling: The The Key to Unlocking Advanced AI

Generative modeling represents a significant advancement in the field of artificial intelligence, offering the potential to drive progress in various domains. This powerful technique goes beyond the limitations of traditional discriminative modeling, providing a more complete understanding of data distributions and paving the way for more sophisticated AI systems.

Theoretical Foundations of Generative Modeling

1 Capturing the Full Picture

Generative modeling is not simply about categorizing data; it aims to understand the underlying distribution of the data, going beyond the constraints of specific labels or classes.

f Probability Dist eristics, Examples,

PROBABILITY DISTRIBUTION




Discrete distributions have finite number of different possible outcomes.




Continuous distributions have infinite many consecutive possible values.

Expected values might not be achievable
 $P(Y \leq y) = P(Y < y + 1)$



$P(Y = y) = 0$ for any distinct value y .
 $P(Y < y) = P(Y \leq y)$



DatabaseTown.com

Theoretical Foundations of Generative Modeling

2

Comprehensive Understanding

By training models to capture the complete data distribution, we can gain a more thorough understanding of the data, leading to more robust and adaptable AI systems.

f Probability Dist eristics, Examples,

PROBABILITY DISTRIBUTION




Discrete distributions have finite number of different possible outcomes.




Continuous distributions have infinite many consecutive possible values.

Expected values might not be achievable
 $P(Y \leq y) = P(Y < y + 1)$



$P(Y = y) = 0$ for any distinct value y .
 $P(Y < y) = P(Y \leq y)$



DatabaseTown.com

Theoretical Foundations of Generative Modeling

3 Unlocking New Possibilities

This comprehensive understanding of the data distribution opens the door to a wide range of applications and advancements in the field of artificial intelligence.

f Probability Dist eristics, Examples,

PROBABILITY DISTRIBUTION




Discrete distributions have finite number of different possible outcomes.




Continuous distributions have infinite many consecutive possible values.

Expected values might not be achievable
 $P(Y \leq y) = P(Y < y + 1)$



$P(Y = y) = 0$ for any distinct value y .
 $P(Y < y) = P(Y \leq y)$



DatabaseTown.com

Advancing AI with Generative Modeling

Reinforcement

Generative modeling is now being used to drive progress in other areas of AI, such as reinforcement learning. By capturing the underlying distribution of data, generative models can be used to simulate and train agents in complex environments, improving their ability to optimize for a given goal.

Advancing AI with Generative Modeling

Robotic Control

For example, in the case of training a robot to walk across a given terrain, generative modeling can be used to create realistic simulations of the environment, allowing the robot to test and refine its strategies before deployment in the real world.

Advancing AI with Generative Modeling

Expanding AI

By combining generative modeling with other AI techniques, such as reinforcement learning, we can develop more powerful and adaptable systems that can tackle a wide range of problems, pushing the boundaries of what is possible in the field of artificial intelligence.

Human Brain Vs Artificial Intelligence



Generative Modeling and Human-level Intelligence

Capturing Complexity

To truly achieve a level of artificial intelligence comparable to a human's, we must move beyond the limitations of discriminative modeling and embrace the power of generative modeling. This approach allows us to capture the inherent complexity and nuances of the world around us.




Human Brain Vs Artificial Intelligence



Generative Modeling and Human-level Intelligence

Emulating Human Cognition

By training models to understand the underlying structure and distribution of data, we can begin to emulate the way humans perceive and make sense of the world, paving the way for AI systems that can truly think, reason, and problem-solve in a more human-like manner.



Human Brain Vs Artificial Intelligence



Generative Modeling and Human-level Intelligence

Towards Artificial General Intelligence

The integration of generative modeling into the development of artificial intelligence is a crucial step towards achieving a level of artificial general intelligence (AGI) that can rival and potentially surpass human capabilities in a wide range of tasks and domains.

1

2

3

Hello World!

Hello World!




Figure 1-4. A set of points in two dimensions, generated by an unknown rule p_{data}

Types of Generative Models

Variational Autoencoders (VAEs)

VAEs are a type of deep generative model that learn a low-dimensional latent representation of input data. They can generate new samples by sampling from the learned latent space.

Generative Adversarial Networks (GANs)
(GANs)
GANs pit two neural networks - a generator and a discriminator - against each other in a game-like framework, leading to the generation of highly realistic synthetic data.

Autoregressive Models

Autoregressive models, like Transformers and Recurrent Neural Networks, generate new data one element at a time, using the previously generated elements to predict the next.

Diffusion Models

Diffusion models learn to reverse a noising process, starting from pure noise and gradually removing it to produce new, realistic samples. They have shown impressive results in image and audio generation.

Variational Autoencoders (VAEs)


Variational Autoencoders (VAEs) are a type of generative model that use neural networks to learn a compressed, probabilistic representation of the input data. VAEs can generate new, realistic-looking samples by sampling from this learned latent space. VAEs work by training an encoder network to map the input data to a compact latent representation, and a decoder network to reconstruct the original input from the latent code. The latent space is structured to follow a standard normal distribution, allowing for smooth interpolation and generation of new samples.



Generative Adversarial Networks (GANs)

GANs are a powerful class of generative models that pit two neural networks against each other in a competitive game. One network, the generator, learns to create realistic samples, while the other network, the discriminator, learns to distinguish real samples from the generator's output.


The interplay between the generator and discriminator networks allows GANs to generate highly realistic and diverse samples, making them particularly useful for tasks such as image and text generation.



Autoregressive Models

Autoregressive models are a type of generative model that generate new data by predicting the next element in a sequence, based on the previous elements. These models learn the underlying statistical patterns in the training data and use them to generate plausible new samples.

Autoregressive models, such as Recurrent Neural Networks (RNNs) and Transformers, have been widely used for tasks like language modeling, speech recognition, and time series forecasting.



Applications of Generative Models



Healthcare

Generative models have enabled the synthesis of realistic medical images, assisting in diagnosis, treatment planning, and clinical decision-making.

Applications of Generative Models



Art & Design

Generative models can be used to create unique and visually stunning artwork, fashion designs, and other creative outputs.

Applications of Generative Models



Natural Language

Generative models excel at tasks such as text generation, language translation, and dialogue systems, revolutionizing natural language processing.

Challenges and Limitations

Data Availability

Generative models often require large, high-quality datasets for effective training. Limited data can hinder their performance and versatility.

Challenges and Limitations

Data Availability

Generative models often require large, high-quality datasets for effective training. Limited data can hinder their performance and versatility.

Model Complexity

The intricate architectures of some generative models can make them computationally intensive and challenging to train and deploy efficiently.

Challenges and Limitations

Data Availability

Generative models often require large, high-quality datasets for effective training. Limited data can hinder their performance and versatility.

Model Complexity

The intricate architectures of some generative models can make them computationally intensive and challenging to train and deploy efficiently.

Mode Collapse

GANs in particular can suffer from mode collapse, where the generator learns to produce limited, repetitive outputs instead of diverse samples.

Challenges and Limitations

Data Availability

Generative models often require large, high-quality datasets for effective training. Limited data can hinder their performance and versatility.

Model Complexity

The intricate architectures of some generative models can make them computationally intensive and challenging to train and deploy efficiently.

Mode Collapse

GANs in particular can suffer from mode collapse, where the generator learns to produce limited, repetitive outputs instead of diverse samples.

Lack of Control

Generating high-fidelity, controllable outputs remains an active area of research, as users often desire more fine-grained control over the generative process.

Thank You