

Deep Generative Models

Prepared by Dr. Sherif Mostafa

AI314

Master of Science in Artificial Intelligence

Lecture 2



Deep Generative Models

What is Generative Modeling?

Generative Versus Discriminative Modeling

Rise of Generative Modeling

Types of Generative Models

What is generative modeling?

Generative modeling can be broadly defined as:

Deep generative models are a class of machine learning models designed to generate new data points that resemble a given training dataset as images, text or audio. They learn the underlying distribution of the training data and can produce new samples that are similar in structure and characteristics to the input data.


What is generative modeling?

Deep generative models are multi-layer nonlinear neural networks that are used to simulate data dependency.

Generative models learn and capture the inner probabilistic distribution of training samples to produce new samples similar to it. In other word, they model simulated observations that are drawn from probability density functions.

Generative models aims to capture the underlying patterns and distributions in the data.

What is generative modeling?



A generative model trained to generate realistic photos of horses

What is generative modeling?



What is generative modeling?

Learn Latent Representations


Latent Representations are the abstract features or variables that summarize the essential information of the input data while discarding irrelevant details. They are not directly observable but are inferred through the model.

Generative 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 more precise and compact form.

What is generative modeling?

Adversarial training

Generative models like Generative Adversarial Networks (GANs) employ an adversarial training process, where a Generator model compete with a Discriminator to produce increasingly more realistic samples.



What is generative modeling?


Generative models 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 trial.

In other words, generative data is not just a fixed calculations to get the mean value of each pixel in a dataset.

What is generative modeling?

It is important to know the main difference between the ordinary discriminative models and the generative models. That generative models 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.

What is generative modeling?



A discriminative model trained to predict if a given image is painted by Van Gogh

What is generative modeling?



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

Generative Modeling: the key to unlock the 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 Foundation of Generative Modeling

Capturing the full picture

Generative modeling is not just about categorizing data; it also aiming to comprehend the underlying distribution of data exceeding the constraints of specific labels or classes

Theoretical Foundation of Generative Modeling

Comprehensive Understanding

We will get a more through understanding of the data through training models to get and capture the fully data distribution leading to more robust acceptable AI systems.

Theoretical Foundation of Generative Modeling

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.

Advancing AI with Generative Modeling

Reinforcement Learning

Generative modeling is now being used to drive progress in other areas of AI, such as reinforcement learning. Through 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 target.

Advancing AI with Generative Modeling

Robotic Control

In cases like 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 just before the deployment in the realistic world.

Advancing AI with Generative Modeling

Expanding AI capabilities:

Through combining generative modeling with other AI techniques, like 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

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.

Generative Modeling and Human-Level intelligence

Emulating Human Cognition

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

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 that can rival a potentially surpass human capabilities in a wide range of tasks and domains.

Types of Generative Models

Variational Autoencoder (VAEs)	Generative Adversarial Networks (GANs)	Autoregressive Models	Diffusion Models
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	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 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 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 specific type of deep learning model used for generating new data that use neural networks to learn a compressed, probabilistic representation of the input data.

VAEs belong to a class of models designed to generate new data that resembles a training dataset by learning how to represent the underlying structure of the data.

VAEs encode input data (like images or text) into a simpler, lower-dimensional form called a latent representation. This represents essential features of the data while ignoring noise and irrelevant details.

Variational Autoencoders (VAEs)



The basic scheme of a variational autoencoder. The model receives x as input. The encoder compresses it into the latent space. The decoder receives as input the information sampled from the latent space and produces $\hat{x}$ as similar as possible to x

Variational Autoencoders (VAEs)

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 to allow for smooth interpolation and generation of new samples.

Once the VAE has learned this latent space, it can create new data by sampling from this space. This means you can generate new instances that share characteristics with the original dataset.


Generative Adversarial Networks (GANs)

GANs are powerful generative models which place two neural networks against each other in a competitive game.

One of these network, the generator, learn to create realistic samples, while the other network, the discriminator, learns to distinguish real samples from 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.

Generative Adversarial Networks (GANs)



Autoregressive Models

Autoregressive models are a form of generative model that creates new data by anticipating the next element in a series using the preceding element. These models discover the underlying statistical patterns in training data and apply them to generate convincing fresh new samples.

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

Transformer-based Models

Transformer-based models, such as OpenAI's GPT (Generative Pre-trained Transformer) series, have revolutionized natural language processing. These models utilize attention mechanisms to capture long-range dependencies in text, enabling them to generate coherent and contextually appropriate language.

Transformer-based Models

The success of transformer-based models can be attributed to their ability to process input sequences in parallel, making them efficient and capable of handling large-scale text data. By pre-training on vast amounts of text data, these models acquire a strong understanding of language and context, which is then fine-tuned on specific downstream tasks.

Advancing AI with Generative Modeling

Models	Key Concept	Pros	Cons
Variational Autoencoders (VAEs)	An encoder-decoder framework via probabilistic graphical models, where the lower bound is maximized on the log-likelihood of the data	Simultaneously perform both generation and inference with latent variables	VAE-generated image samples tend to be slightly blurry. KL vanishing issues in sequence modeling
Generative Adversarial Networks (GANs)	A generator-discriminator framework via an adversarial training game, where samples of the data are directly generated	Generate the sharpest image samples	More difficult to optimize due to unstable training dynamic
Autoregressive models (e.g. Neural Language Model, or NLM)	Factorize the joint distribution of data into the conditional distributions, modeling every individual dimension given previous dimensions	Simple and stable training, yielding the best log-likelihood	NLMs are inefficient during sampling and don't easily provide low-dimensional features

Applications of generative models

Natural Language Processing

Natural language processing is being revolutionized by generative models, which are excellent at tasks like dialogue systems, language translation, and text production.

Healthcare

Generative models have allowed for the creation of realistic medical visualizations, which aid in diagnosis, treatment planning, and clinical decision making.

Applications of generative models

Art & Design

Generative models can be used to create distinctive and visually attractive artwork, fashion design, and other creative outputs that mimic and resemble human creativity.

Content Creation

They can produce new images, music, or texts that can be used for artistic purposes or content development across various fields.

Applications of generative models

Healthcare

Generative models have allowed for the creation of realistic medical visualizations, which aid in diagnosis, treatment planning, and clinical decision making.

Data Augmentation

They can add artificial data to training datasets to increase their size and enhance the performance of other machine learning models.

Applications of generative models

Simulation and Gaming

They can create believable characters, settings, or situations for video games or digital simulations

Scenario Generation

They can create fictitious scenarios or other potential outcomes to assist in decision-making or strategic planning.

Challenges and Limitations

Data Availability

Generative models typically require big, high-quality datasets for effective training. Limited data may negatively impact their performance and flexibility.

Model Complexity

Certain generative models can be computationally demanding and difficult to train and scale effectively because to their extensive structure.

Challenges and Limitations

Mode Collapse

GANs, in particular, are vulnerable to mode collapse, in which the generator learns to produce limited repeating output rather than diversified samples.

Lack of Control

Research on producing highly controllable outputs is still ongoing since customers frequently want more precise control over the generative process. There is still more work to be done in this area and development required to meet this requirement.

Conclusion

In conclusion, generative AI models represent a significant leap forward in our ability to harness artificial intelligence for creative endeavors. Whether generating realistic images, composing music, or crafting compelling stories, these models reshape industries and provide new avenues for human expression.

With continued research and responsible implementation, generative AI models hold immense potential to push the boundaries of human imagination and innovation.



Thank You