



SNS COLLEGE OF TECHNOLOGY

(Autonomous)



DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

AUTOENCODERS

Autoencoders have emerged as one of the technologies and techniques that enable computer systems to solve data compression problems more efficiently.


They became a popular solution for reducing noisy data.

Simple autoencoders provide outputs that are the same or similar to the input data—only compressed. In the case of variational autoencoders, often discussed in the context of large language models, the output is newly generated content.

are a kind of artificial neural network that is utilized for unsupervised learning in order to create effective data representations. The objective of an autoencoder is to acquire an encoding for a dataset, often used for reducing dimensionality, cleaning data, or learning features.

This is how an autoencoder operates:

- **Encoding:** The input information undergoes encoding by a neural network, which condenses the information into a latent-space representation that is typically of a smaller dimension compared to the initial input.
- **Latent Representation:** The most crucial characteristics of the input data are captured by the latent representation. It is a compact, concentrated depiction that ideally retains the most important information.
- **Decoding:** The hidden representation is later fed through a decoder neural network that tries to recreate the initial input data from the hidden representation.



Basic architecture of autoencoders

Types of Autoencoders

The types of autoencoders are as follows:

Vanilla Autoencoder

Vanilla Autoencoder is a simple yet powerful framework for unsupervised learning tasks. It comprises of two primary components: an encoder and a decoder. Together, these components function together compress input data into a lower-dimensional representation and then reconstruct it.

- Training a Vanilla Autoencoder involves minimizing the reconstruction error, which quantifies the disparity between the input data and the reconstructed output.

- This is typically achieved through a process known as backpropagation, where gradients of the reconstruction error with respect to the model parameters are computed and used to update the weights of the neural network layers. The objective is to optimize the parameters such that the reconstructed output closely matches the original input.

Applications of Vanilla Autoencoders:

Vanilla Autoencoders find applications across a wide spectrum of domains, owing to their versatility and simplicity. Some notable applications include:

- **Data Compression:** By learning a compressed representation of the input data, Vanilla Autoencoders excel in data compression tasks, facilitating efficient storage and transmission of information.
- **Feature Learning:** Vanilla Autoencoders are adept at capturing salient features of the input data, making them valuable for feature learning tasks in domains such as computer vision, natural language processing, and sensor data analysis.
- **Anomaly Detection:** The ability of Vanilla Autoencoders to reconstruct input data accurately makes them well-suited for anomaly detection tasks. Deviations between the original input and the reconstructed output can signal anomalies or outliers in the data.

Sparse Autoencoder

In neural networks, Sparse Autoencoders emerge as a remarkable variant, distinguished by their emphasis on sparsity within the latent representation. While traditional autoencoders aim to faithfully reconstruct input data in a lower-dimensional space, Sparse Autoencoders introduce constraints that encourage only a small subset of neurons to activate, leading to a sparse and efficient representation.

- **Regularization:** Introducing penalties or constraints on the activation of neurons in the latent space encourages sparsity. Common regularization techniques include L1 regularization, which penalizes the absolute values of the weights, and dropout, which randomly deactivates neurons during training.
- **Kullback-Leibler (KL) Divergence:** KL divergence is a measure of the difference between two probability distributions. In the context of Sparse Autoencoders, it can be used as an additional loss term to encourage the latent distribution to match a predefined sparse target distribution.

Denoising Autoencoder

Denoising Autoencoders emerge as a formidable solution for handling noisy input data. Unlike Denoising Autoencoders offer a powerful solution for handling noisy input data, enabling robust feature learning and data reconstruction in the presence of noise. By intentionally corrupting input data with noise and training the autoencoder to recover the clean underlying representation, Denoising Autoencoders effectively filter out noise and enhance the quality of reconstructed data. Their applications span diverse domains, from image and signal processing to data preprocessing and beyond, making them invaluable tools in the arsenal of machine learning practitioners striving for robust and reliable solutions in the face of noisy data.

- Training Denoising Autoencoders involves optimizing the model parameters to minimize the reconstruction error between the clean input data and the output reconstructed by the decoder.
- However, since the input data is intentionally corrupted during training, the autoencoder learns to filter out the noise and recover the underlying clean representation. This process encourages the autoencoder to focus on capturing meaningful features while disregarding the noise present in the input data.

Contractive Autoencoder

Contractive Autoencoders are a specialized type of autoencoder designed to learn robust and stable representations of input data. Unlike traditional autoencoders that focus solely on reconstructing input data, contractive autoencoders incorporate an additional regularization term in the training objective. This term penalizes the model for producing representations that are sensitive to small variations in the input data, encouraging the learning of more stable and invariant features.

- Training contractive autoencoders involves optimizing the model parameters to minimize both the reconstruction error and the additional regularization term.
- The regularization term penalizes the model for producing representations that are sensitive to small changes in the input data, encouraging the learning of stable and invariant features.
- This is typically achieved through backpropagation and gradient descent, where the weights of the neural network layers are adjusted iteratively to minimize the combined loss function.

Applications of Contractive Autoencoders:

Contractive autoencoders find applications in domains where stability and robustness in feature representation are critical. Some notable applications include:

- **Representation Learning:** Contractive autoencoders excel in learning stable and invariant representations of input data, making them valuable for tasks such as classification and clustering.
- **Transfer Learning:** The stable representations learned by contractive autoencoders can be transferred to downstream tasks, where they serve as effective feature vectors for tasks with limited labeled data.
- **Data Augmentation:** Contractive autoencoders can be used to generate augmented data by perturbing the input data and reconstructing it with stable representations, thereby increasing the diversity of the training dataset.

Applications of Denoising Autoencoders:

Denoising Autoencoders find applications across a wide range of domains where input data is prone to noise or corruption. Some notable applications include:

- **Image Denoising:** In computer vision tasks, Denoising Autoencoders are used to remove noise from images, enhancing image quality and improving the performance of subsequent image processing algorithms.
- **Signal Denoising:** In signal processing applications such as audio processing and sensor data analysis, Denoising Autoencoders can effectively filter out noise from signals, improving the accuracy of signal detection and analysis.
- **Data Preprocessing:** Denoising Autoencoders can be employed as a preprocessing step in machine learning pipelines to clean and denoise input data before feeding it into downstream models. This helps improve the robustness and generalization performance of the overall system.