

## Lesson 10: Dimensionality Reduction

Dimensionality reduction is a type of unsupervised learning technique that is used to reduce the number of features in a dataset while preserving as much information as possible. This can be particularly useful when working with high-dimensional data, as it can help to reduce noise and improve the efficiency of algorithms that work on the data.

There are two main types of dimensionality reduction techniques: feature selection and feature extraction. Feature selection involves selecting a subset of the original features that are most relevant to the task at hand, while feature extraction involves creating new features that are a combination of the original features.

One of the most commonly used dimensionality reduction techniques is principal component analysis (PCA), which involves projecting the data onto a lower-dimensional space while preserving as much of the original variation as possible. PCA works by identifying the principal components of the data, which are linear combinations of the original features that capture the most variation in the data.

Another popular dimensionality reduction technique is t-distributed stochastic neighbor embedding (t-SNE), which is particularly useful for visualizing high-dimensional data in two or three dimensions. t-SNE works by first computing pairwise similarities between the data points, and then optimizing a cost function that minimizes the difference between the pairwise similarities in the high-dimensional space and the pairwise similarities in the low-dimensional space.

While dimensionality reduction can be a powerful tool for improving the efficiency and accuracy of machine learning algorithms, it is important to be aware of its limitations. In particular, dimensionality reduction can lead to the loss of important information in the data, and it can be difficult to interpret the meaning of the new features that are created.

Examples of real-world applications of dimensionality reduction include image and video processing, text mining, and bioinformatics. In image and video processing, dimensionality reduction can be used to extract important features from the data, such as color, texture, and shape. In text mining, dimensionality reduction can be used to extract the most important words or topics from a corpus of text, while in bioinformatics, dimensionality reduction can be used to identify patterns in large datasets of genetic data.

## Principal Component Analysis

Principal Component Analysis (PCA) is a widely used dimensionality reduction technique that aims to reduce the number of dimensions in a dataset while preserving as much of the original information as possible. The main idea behind PCA is to find a new set of variables, called principal components, that are linear combinations of the original variables and capture the most variation in the data.

PCA works by first centering the data around its mean and then computing the covariance matrix of the data. The covariance matrix contains information about the relationships between the variables in the data, and it can be calculated using the following formula:

$$\begin{bmatrix} \text{cov}(x_1, x_1) & \text{cov}(x_1, x_n) & \text{cov}(x_2, x_1) & \text{cov}(x_2, x_n) & \vdots & \vdots & \text{cov}(x_n, x_1) & \text{cov}(x_n, x_n) \end{bmatrix}$$

PCA then finds the eigenvectors and eigenvalues of the covariance matrix. The eigenvectors represent the principal components, and the corresponding eigenvalues represent the amount of variation explained by each principal component. The principal components can be calculated using the following formula:

$$\begin{bmatrix} \text{PC}_1 & \text{PC}_2 & \vdots & \text{PC}_n \end{bmatrix} = \begin{bmatrix} e_{11} & e_{12} & \cdots & e_{1n} & e_{21} & e_{22} & \cdots & e_{2n} & \vdots & \vdots & \ddots & \vdots & e_{n1} & e_{n2} & \cdots & e_{nn} \end{bmatrix} \begin{bmatrix} x_1 - \mu_1 & x_2 - \mu_2 & \vdots & x_n - \mu_n \end{bmatrix}$$

PCA can be used for various purposes, such as data compression, visualization, and noise reduction. In data compression, PCA can be used to reduce the dimensionality of the data while retaining most of the information. In visualization, PCA can be used to project high-dimensional data onto a lower-dimensional space for visualization purposes. In noise reduction, PCA can be used to remove noise from the data by filtering out the components with low eigenvalues.

One limitation of PCA is that it is a linear technique and may not be able to capture nonlinear relationships in the data. In such cases, nonlinear dimensionality reduction techniques, such as t-SNE, may be more appropriate. Nonetheless, PCA is a powerful tool that can be used to extract meaningful insights from high-dimensional datasets.


## t-SNE

t-SNE (t-Distributed Stochastic Neighbor Embedding) is a nonlinear dimensionality reduction technique that is often used for visualizing high-dimensional data in a low-dimensional space. It was first introduced by Laurens van der Maaten and Geoffrey Hinton in 2008.

t-SNE works by minimizing the divergence between the distribution of pairwise similarities in the high-dimensional space and the distribution of pairwise similarities in the low-dimensional space. It does this by modeling the high-dimensional data points as a probability distribution over the pairwise similarities, and then modeling the low-dimensional data points as another probability distribution over the pairwise similarities. The two probability distributions are then compared using the Kullback-Leibler divergence, and the low-dimensional representation of the data is iteratively adjusted until the divergence is minimized.

One of the advantages of t-SNE over other dimensionality reduction techniques is its ability to preserve the local structure of the data. This makes it particularly useful for visualizing clusters or groups of similar data points. However, t-SNE can be computationally expensive and may require careful tuning of hyperparameters.

t-SNE has been used in various applications, such as visualizing gene expression data, analyzing text data, and exploring high-dimensional images in computer vision. It has also been used in anomaly detection and clustering.



---

## EXAMPLE CODE

In this example, we load the digits dataset and perform PCA and t-SNE on it. We first perform PCA with two components and visualize the result using a scatter plot. Then, we perform t-SNE with two components and visualize the result using another scatter plot. The colors of the points correspond to the digit labels in the dataset.

```
from sklearn.datasets import load_digits
from sklearn.decomposition import PCA
from sklearn.manifold import TSNE
import matplotlib.pyplot as plt

# Load the dataset
digits = load_digits()

# Perform PCA with two components
pca = PCA(n_components=2)
pca_result = pca.fit_transform(digits.data)

# Visualize the PCA result
plt.scatter(pca_result[:, 0], pca_result[:, 1], c=digits.target)
plt.title('PCA Result')
plt.show()

# Perform t-SNE with two components
tsne = TSNE(n_components=2)
tsne_result = tsne.fit_transform(digits.data)

# Visualize the t-SNE result
plt.scatter(tsne_result[:, 0], tsne_result[:, 1], c=digits.target)
plt.title('t-SNE Result')
plt.show()
```