

## Lesson 9: Unsupervised Learning

Unsupervised learning is a type of machine learning that involves discovering patterns or relationships in data without the use of explicit labels or supervision. Unlike supervised learning, where the algorithm is provided with labeled examples to learn from, unsupervised learning involves working with unstructured or unlabeled data and finding hidden structures or relationships within it.

In unsupervised learning, the goal is to uncover underlying patterns or structures in the data, such as clusters, outliers, or latent factors, that can be used for further analysis or decision making. Unsupervised learning algorithms include clustering, anomaly detection, dimensionality reduction, and association rule mining, among others.

Unsupervised learning is useful in a wide range of applications, including customer segmentation, anomaly detection, image and speech recognition, and recommendation systems. It is a powerful tool for exploring and understanding large, complex datasets, and can often reveal insights and patterns that may not be immediately apparent through other methods.

### Clustering Basics

Clustering is a fundamental technique in unsupervised learning that involves grouping similar data points together into clusters. Unlike supervised learning, where the data is labeled, unsupervised learning involves working with unlabeled data, and clustering is one of the most common approaches for analyzing such data.

The goal of clustering is to find patterns and structure in the data, by grouping similar data points together based on some similarity metric. The similarity metric can be based on various measures, such as distance, similarity, or density, depending on the specific clustering algorithm.

Clustering algorithms can be broadly divided into three categories: partitioning, hierarchical, and density-based. Partitioning algorithms, such as k-means, partition the data into a fixed number of clusters, whereas hierarchical algorithms, such as agglomerative clustering, create a hierarchy of nested clusters. Density-based algorithms, such as DBSCAN, group together data points that are within a certain density threshold.

Clustering has many applications in various fields, such as image processing, natural language processing, and customer segmentation. By identifying groups of similar data points, clustering can help to reveal underlying patterns and relationships in the data, which can be useful for various tasks, such as anomaly detection, recommendation systems, and data visualization.

## k-Means Clustering

k-Means Clustering is a popular unsupervised learning technique used to cluster data into a pre-specified number of groups or clusters. The algorithm is based on the concept of minimizing the sum of squared distances between the data points and their assigned cluster centers.

The k-means algorithm works by first selecting  $k$  initial cluster centers, either randomly or through some other initialization method. Each data point is then assigned to the closest cluster center, based on a distance metric such as Euclidean distance. The mean of each cluster is then computed and used as the new cluster center. This process is repeated iteratively, with data points being reassigned to new clusters and cluster centers being updated, until convergence or until some stopping criterion is met.

k-means clustering has several advantages, including its simplicity and ease of implementation. It is also computationally efficient and can be used with large datasets. However, k-means can be sensitive to the initial cluster center selection and can converge to suboptimal solutions.

There are several extensions and variations of k-means, including fuzzy k-means and hierarchical k-means. These variations offer increased flexibility and can often produce better clustering results, but can be more computationally expensive.

k-Means clustering is one of the most commonly used unsupervised learning algorithms for clustering tasks. It is a simple but powerful algorithm that can partition a dataset into  $k$  clusters based on the similarity between data points. The algorithm works by iteratively optimizing a clustering objective function, which measures the similarity between data points and their assigned cluster centers.

Here is a brief overview of the k-means algorithm:

1. Choose the number of clusters  $k$ .
2. Initialize  $k$  cluster centers randomly.
3. Assign each data point to the nearest cluster center.
4. Recalculate the cluster centers as the mean of the data points assigned to each cluster.
5. Repeat steps 3-4 until convergence (i.e., until the cluster centers no longer change).

To better understand how the k-means algorithm works, let's take a look at an example implementation on a dataset. In this example, we will use the scikit-learn library to perform k-means clustering on the Iris dataset.

```
# import libraries
from sklearn.datasets import load_iris
from sklearn.cluster import KMeans

# load dataset
iris = load_iris()
X = iris.data

# initialize k-means object with 3 clusters
kmeans = KMeans(n_clusters=3)

# fit the k-means object to the data
kmeans.fit(X)

# retrieve the labels assigned to each data point
labels = kmeans.labels_

# retrieve the coordinates of the cluster centers
centers = kmeans.cluster_centers_
```

In this code example, we first load the Iris dataset and extract the input data. Then, we initialize a k-means object with 3 clusters and fit it to the data. Finally, we retrieve the labels assigned to each data point and the coordinates of the cluster centers.

We can visualize the results of the clustering using a scatter plot, where each data point is colored according to its assigned cluster.

```
import matplotlib.pyplot as plt

# plot the results
plt.scatter(X[:, 0], X[:, 1], c=labels)
plt.scatter(centers[:, 0], centers[:, 1], marker='*', s=200,
c='#050505')
plt.show()
```

In this visualization, we can see that the k-means algorithm has successfully clustered the Iris dataset into 3 distinct clusters based on the similarity between the data points. The cluster centers are marked with a star symbol.

Overall, the k-means algorithm is a powerful tool for clustering tasks and can be easily implemented in Python using libraries like scikit-learn.

## Hierarchical Clustering

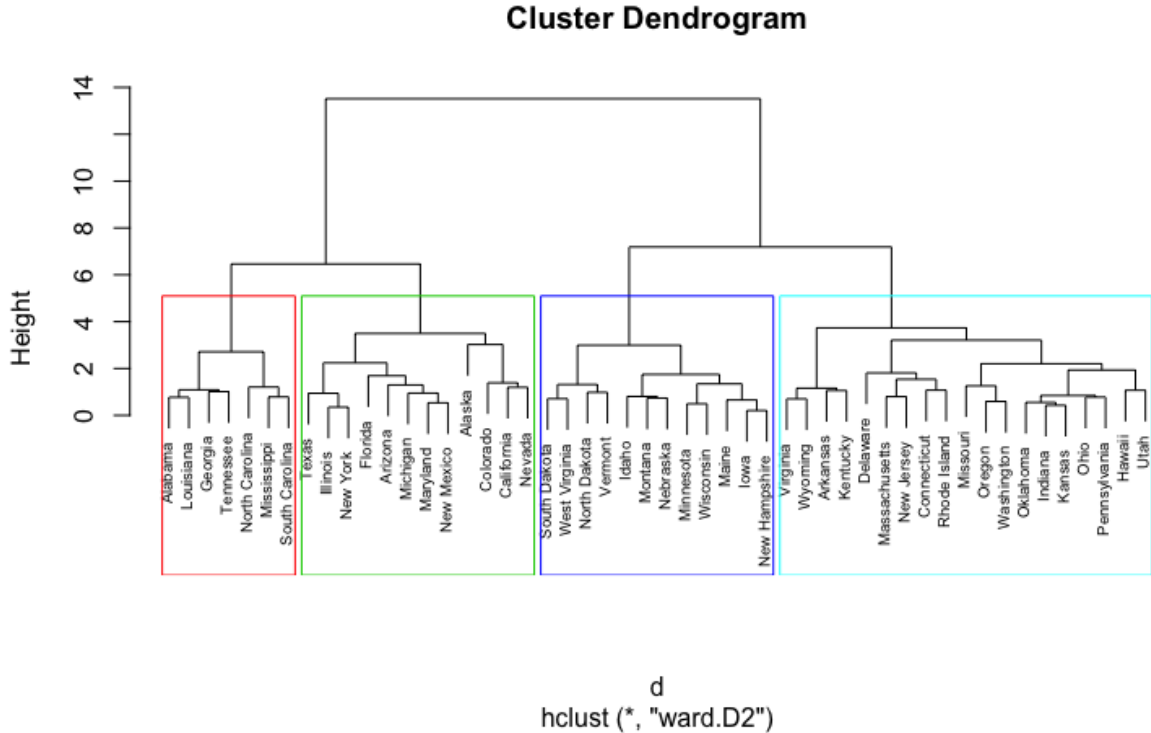
Hierarchical clustering is a type of clustering algorithm that creates a hierarchy of clusters by iteratively merging or splitting existing clusters. There are two main types of hierarchical clustering: agglomerative and divisive.

In agglomerative hierarchical clustering, each data point initially forms its own cluster, and the algorithm proceeds by iteratively merging the two closest clusters. The distance between two clusters is typically defined as the distance between the closest pair of data points from each cluster. This process continues until all data points belong to a single cluster.

In divisive hierarchical clustering, the process is the opposite of agglomerative clustering: all data points initially belong to a single cluster, and the algorithm proceeds by iteratively splitting the cluster into smaller subclusters. The distance between two subclusters is typically defined as the distance between the farthest pair of data points from each subcluster. This process continues until each subcluster consists of a single data point.

One advantage of hierarchical clustering is that it can produce a dendrogram, which is a tree-like diagram that shows the hierarchical relationships between clusters. This can be useful for visualizing and interpreting the results of the clustering.

Hierarchical clustering can be used for a wide range of applications, including image segmentation, gene expression analysis, and customer segmentation in marketing. However, it can be computationally expensive and may not be suitable for large datasets or datasets with high dimensionality.



## Density-Based Clustering


Density-based clustering is a type of unsupervised learning algorithm that groups together data points based on their density. It works by identifying regions of high density in the data space, and then assigning each point to a cluster based on its proximity to these high-density regions. This approach is particularly useful for data that has irregular shapes and varying densities.

One of the most popular density-based clustering algorithms is DBSCAN (Density-Based Spatial Clustering of Applications with Noise). DBSCAN starts with a

random point in the dataset and then finds all nearby points within a certain distance, called the epsilon neighborhood. If the number of points in the neighborhood is above a certain threshold, the point is considered to be in a dense region, and all neighboring points are added to the same cluster. The algorithm then moves to the next unvisited point and repeats the process until all points have been visited.

Another density-based clustering algorithm is OPTICS (Ordering Points To Identify the Clustering Structure), which is an extension of DBSCAN that allows for more flexible clustering. OPTICS computes a reachability distance for each point, which measures the distance between the point and its neighbors in the dataset. Points with low reachability distances are considered to be in dense regions, while points with high reachability distances are considered to be in sparse regions.

Density-based clustering algorithms are useful in a variety of applications, such as image segmentation, anomaly detection, and identifying groups of similar objects in large datasets. However, they can be sensitive to the choice of hyperparameters, such as the epsilon neighborhood size, and may not perform well on datasets with varying densities or complex structures.



## Anomaly Detection

Anomaly detection is a technique in unsupervised learning that involves identifying unusual or rare data points that deviate significantly from the norm. Anomaly detection is used in a variety of applications, including fraud detection, network intrusion detection, and equipment failure prediction.

The basic idea behind anomaly detection is to define a notion of what is normal or expected in the data and then identify any data points that do not conform to that notion. There are several approaches to anomaly detection, including statistical methods, machine learning techniques, and rule-based systems.

Statistical methods for anomaly detection involve modeling the distribution of the data and identifying data points that are unlikely to occur under that distribution. Machine learning techniques for anomaly detection involve training a model on a dataset of normal data points and then using the model to identify any data points that deviate significantly from the norm. Rule-based systems for anomaly detection involve defining a set of rules or thresholds for what is considered normal behavior and flagging any data points that violate those rules.

Anomaly detection can be a challenging problem, as anomalous data points may be rare or difficult to define. Furthermore, the cost of false positives and false negatives can vary greatly depending on the application, and it is important to balance the two appropriately.

Here is an example code for anomaly detection using the Isolation Forest algorithm in Python:

```
import pandas as pd
from sklearn.ensemble import IsolationForest

# Load data
data = pd.read_csv('data.csv')

# Define the outlier fraction
outlier_frac = 0.05
```

```
# Select features for analysis
features = ['feature_1', 'feature_2', 'feature_3']

# Subset the data to the selected features
X = data[features]

# Initialize the IsolationForest algorithm
isoforest = IsolationForest(contamination=outlier_frac)

# Fit the model to the data
isoforest.fit(X)

# Predict the outlier scores for the data points
outlier_scores = isoforest.decision_function(X)

# Add the outlier scores to the original data frame
data['outlier_score'] = outlier_scores

# Identify the anomalies using a threshold on the outlier score
anomalies = data[data['outlier_score'] < -0.5]
```

In this example, we load the data from a CSV file and select a subset of features for analysis. We then initialize an instance of the IsolationForest algorithm with a specified outlier fraction, fit the model to the selected features, and predict the outlier scores for the data points. Finally, we add the outlier scores to the original data frame and identify the anomalies using a threshold on the outlier score.

Note that the threshold value for identifying anomalies may need to be adjusted based on the specific data set and problem domain.



## Evaluating Clustering Performance

Evaluating clustering performance is essential to determine the quality and effectiveness of the clustering algorithm applied. There are various metrics and techniques available for evaluating clustering performance, depending on the type of data, the algorithm used, and the objective of the analysis.

One of the most commonly used metrics for evaluating clustering performance is the Silhouette score, which measures the similarity of an object to its own cluster compared to other clusters. The Silhouette score ranges from -1 to 1, with higher scores indicating better clustering results.

Another popular metric for evaluating clustering performance is the Calinski-Harabasz index, which measures the ratio of between-cluster variance to within-cluster variance. Higher Calinski-Harabasz scores indicate better clustering performance.

In addition to these metrics, visual inspection of the clustering results can also provide valuable insights into the effectiveness of the algorithm. This can be done by creating scatter plots or heatmaps to visualize the clusters and the distribution of the data points within each cluster.

It is important to note that there is no one-size-fits-all approach to evaluating clustering performance, and the appropriate metrics and techniques will depend on the specific problem and data being analyzed. Therefore, it is important to consider multiple evaluation metrics and techniques and to choose the ones that are most appropriate for the specific analysis.