Class 4: Unsupervised Learning: k-means
4.1 Unsupervised Learning

Course material
Unsupervised learning is a type of machine learning where the computer is given data that has **inputs but no correct output labels**. In supervised learning, each training example includes both the input and the correct answer. For example, an image may be labelled as "cat" or "dog", or an email may be labelled as "spam" or "not spam". In unsupervised learning, those labels are not provided. The machine must look for patterns in the data by itself.

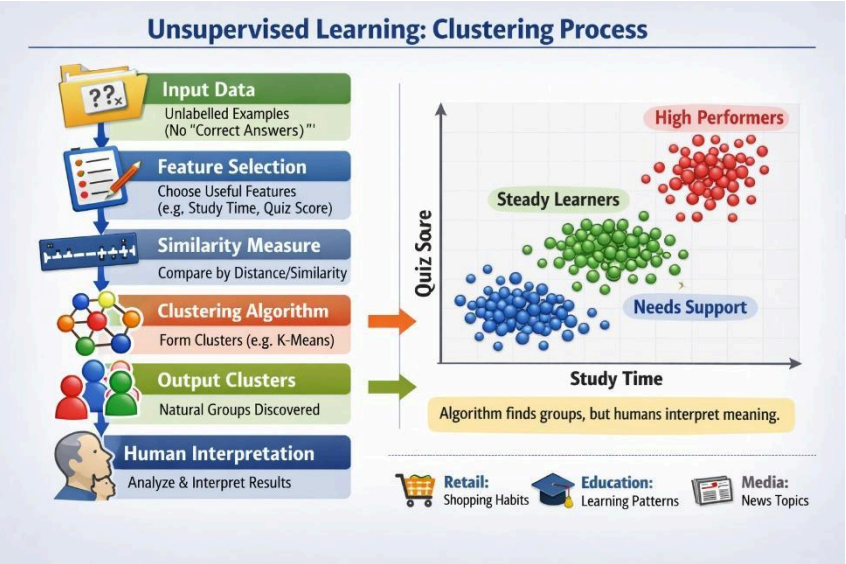

The goal of unsupervised learning is not to predict a known answer. Instead, the goal is to discover hidden structure in the data. This may involve finding natural groups, reducing the complexity of data, or discovering relationships between items. Because there are no correct labels, the algorithm does not know the "true meaning" of the patterns it finds. Humans must examine the results and decide whether they are useful.

A common unsupervised learning task is **clustering**. Clustering means grouping similar items together. Items inside the same cluster should be more similar to each other than to items in other clusters. For example, a shop may use clustering to group customers with similar buying behaviour. A school may use clustering to explore different learning patterns among students. A news website may group articles by topic, even if the topics were not labelled in advance.

Type of machine learning	Data used	Main goal	Example
Supervised learning	Labelled data	Predict a known output	Predict whether an email is spam
Unsupervised learning	Unlabelled data	Discover hidden patterns or groups	Group customers by shopping habits

In this class, the main focus is clustering. Clustering is useful when we want to explore data but do not already know the categories. For example, if a teacher has data about students' quiz scores, study time, and homework completion, clustering may show groups of students with similar learning patterns. However, the algorithm

does not automatically know what these groups mean. A human teacher must interpret the clusters carefully.

Clustering example	Possible data used	Possible cluster meaning
Customers	Items bought, amount spent, visit frequency	Budget shoppers, frequent buyers, luxury buyers
Students	Study time, quiz score, homework completion	Needs support, steady learners, high performers
Images	Colour values, shapes, visual features	Similar image types or styles
News articles	Words, topics, keywords	Politics, sports, technology, entertainment

It is important to understand that clustering does not begin with correct labels. The algorithm only looks at patterns in the data, usually by measuring similarity or distance. If the features are useful, the clusters may reveal meaningful groups. If the features are poorly chosen, the clusters may be confusing or misleading.

For example, clustering students using study time and quiz score may produce useful learning groups. However, clustering students using shoe size and student ID number would probably not produce meaningful academic groups. This shows that the quality of unsupervised learning depends strongly on the data features selected by humans.

Other uses of unsupervised learning

Clustering is one important type of unsupervised learning, but it is not the only one. Unsupervised learning can also be used for **dimensionality reduction**. Dimensionality reduction means reducing the number of features while trying to keep the most important information. This can make data easier to visualise, store, or process.

For example, a digital image may contain thousands or millions of pixel values. A dimensionality reduction method may try to represent the image using fewer important values. This can help computers process the image more efficiently or help humans visualise complex data.

Unsupervised learning can also be used for **representation learning**. Representation learning means learning useful internal features from data. For example, a system may learn patterns from many images, even if the images do not have human-written labels. These learned representations may later help with tasks such as classification, search, recommendation, or image comparison.

Unsupervised learning task	Meaning	Example
Clustering	Group similar items	Group customers by buying behaviour
Dimensionality reduction	Reduce the number of features while keeping important information	Represent a large image using fewer values
Association discovery	Find items or events that often occur together	Customers who buy bread may also buy butter
Representation learning	Learn useful internal features from data	Learn visual patterns from unlabelled images

For this class, students only need to understand that unsupervised learning is broader than clustering. However, k-means is a simple and useful starting point because it clearly shows how a machine can group data without being given correct labels.

--- End ---

