



Monitoring Pistachio Orchards

ADVANCED REMOTE SENSING AND EARTH OBSERVATION COURSE

Project name:

VET Demonstration Partnership for Pistachio Farm in Armenia:
Geospatial Approach

Project acronym: VETfarm

Project number: 101128732

Erasmus+ Capacity Building in the field of Vocational Education and Training

ERASMUS-EDU-2023-CB-VET

ERASMUS+ VET PROGRAMME

June 2025

Course Description

Advanced Remote Sensing and Earth Observation Course

VETfarm Project

Table of Contents

Contents

1. Introduction	3
2. Image classification principles	3
3. Remote Sensing Image Classification.....	5
3.1 Image space	6
3.2 Feature space	7
3.3 Spectral features	9
3.3.1 Cube Spaces	9
4. Key Concepts.....	11
5. The Classification Process	16
6. Image Classification: Practical.....	25
7. Q&A.....	29
8. Essay Format Questions.....	32
9. Glossary of Key Terms.....	33
10. Conclusion: Classification methodology	36
11. Basic literature	36

1. Introduction

The course is designed for vocational training in information extraction from satellite data, with a focus on image classification and segmentation. It covers the fundamental concepts in remote sensing image classification, including the concepts of image and feature space, supervised and unsupervised approaches, various classification algorithms like maximum likelihood, and the crucial step of validating results using error or confusion matrices.

Image classification is utilised in various regional projects. The European Commission requires national governments to verify farmers' claims regarding crop subsidies. To fulfil these requirements, national governments hire companies to conduct an initial assessment using image classification techniques, which is subsequently followed by field checks.

This course also covers key concepts, methods, and challenges in remote sensing image classification.

Special attention is also given to free, open-source systems for viewing and interpreting satellite imagery. It enables participants to efficiently collect up-to-date information about their projects and monitor changes over time.

The main objectives are:

- To provide the knowledge, skills, and understanding necessary for information extraction from satellite data.
- To develop an understanding of the fundamentals of Remote Sensing image classification.

2. Image classification principles

Goal: classify **scene objects** or estimate object model parameters.

Scene: It encompasses the physical surface under study. In the context of remote sensing, these are the actual physical entities or objects, along with their surroundings, that the sensor observes. Object surfaces are sampled through the IFOV of sensor arrays. For example, in an agricultural context, the object surface could contain, e.g. pistachio trees, canals, roads, grasses, and backgrounds such as soils and rocks.

Scene objects: Scene objects are physical objects and their surroundings that exist in the real world (3D) and are the subject of study. The surfaces of these objects are sampled through the IFOV of sensor arrays. The intersection of the "solid" IFOV cone with an object surface is projected onto a sensor element,

and the photon counts from this intersection are represented as image elements or "pixels".

Figure 1 illustrates how an imaging system captures photons from a specific area of an object. The "IFOV angle" (Instantaneous Field of View) represents the narrow cone of vision that a single "pixel" or "image element" on a sensor array sees. When this cone intersects with a physical "object surface," the photon from that specific point on the object (referred to as "Ground Resolution Cell - GRC") is directed to and registered by a single sensor element.




Figure 1. This diagram illustrates the intersection of the "solid" IFOV cone with an object surface, which is projected onto a sensor element of an array of sensors ($AB \rightarrow b,a$). The area AB , which represents the sampled area of the object surface, is also referred to as the Ground Resolution Cell (GRC). Photon counts on ab are represented as image elements or "pixels" in the plane a',b' . The IFOV is not a constant, but rather changes based on the sensor's position relative to the focal point, indicating a dynamic relationship between the optical setup and the captured image data.

Image elements (Pixels): Arrays of sensor data can be represented as arrays of pixels (image elements).

Spectral features: Sensor arrays count photons (also referred to as 'Digital Numbers - DN's') in distinct spectral bands, which can be converted into spectral features.

Spatial features: Connected (adjacent) image samples form regions with spatial features (properties).

Object classes: Object classes or parameters are estimated from prior knowledge (GIS) or previous results, as well as spectral and spatial features.

Estimation is based on causal and/or statistical models requiring supervised sampling.

Training samples: Training samples are collected using training areas, also known as Regions of Interest (ROIs). ROI are selected on homogeneous areas, and can be labelled with class labels or a cluster index.

Clusters: Clusters are defined as adjacent points in a feature space (Spectral features).

Parametric or nonparametric: The density of points in feature space is modelled by statistical distributions (parametric or nonparametric).

Image regions: Image regions have spatial properties (features) that can be used for selection.

Training sessions: In training sessions, there is interaction between spectral and spatial features.

3. Remote Sensing Image Classification

Remote sensing involves capturing data about the object surface using sensors, which is then processed to extract meaningful information, often in the form of maps or thematic classifications. Image classification is a core technique in this process, transforming raw image data into object classes or thematic maps, such as land cover maps. Classified images serve as input for Geographic Information Systems (GIS), supporting various analyses and decision-making processes. For example, the European Commission mandates that national governments verify crop subsidy claims, which involves initial inventories using image classification, followed by field checks.

The application problem in this course is to identify scene objects, such as pistachio crop fields and roads, classify them, and estimate their parameters.

Remote sensing image classification is a process used to extract information from remotely sensed images by assigning pixels to specific object class labels or thematic classes based on their characteristics. It is one of the techniques within the domain of machine learning that allows a computer to perform an interpretation according to defined conditions. This process transforms remote sensing images into land cover types (classes), such as soil type or crop disease maps, including those related to pistachio trees.

The basic assumption for image classification is that a specific part of the feature space corresponds to a specific class, meaning spectrally similar pixels group together to form compact clusters.

3.1 Image space

In remote sensing, **radiometric measurements** are primarily the outputs of a sensor. For this purpose, the sensor may be an RGB camera or a satellite-borne multi-spectral scanner, such as Sentinel-2.

A digital image is a 2-dimensional matrix in which a typical element is referred to by (x_i, y_j) where $i=1, 2 \dots, r$ and $j=1, 2 \dots, c$ (with r : number of rows, c : number of columns). At each element in the **image**, a set of photon counts is available in several spectral bands (referred to as 'Image Space' – Figure 2).

Let $\mathbf{d}(x_i, y_j)$ denote samples corresponding to the N -dimensional measurement space \mathbf{D} at image sample position (x_i, y_j) , where N is the number of spectral bands. The sensor's output is a set of N measurements, each corresponding to a single channel of the scanner. The set of all measurements referring to the same sample is combined in the measurement vector

$$\mathbf{d} = [d_1 \ d_2 \ \dots \ d_N]^T$$

where T is the matrix transpose operator. Hereafter, the measurement vector $\mathbf{d}(x_i, y_j)$ is abbreviated to \mathbf{d} . The vectors are denoted in boldface type (for instance, \mathbf{d} or \mathbf{p}).

The dimension N of \mathbf{d} determines the number of bands that are used for classification.

$$N = \dim(\mathbf{d}) = \text{number of measurements (or bands)}$$

The measurement vector \mathbf{d} points to a single point in the N -dimensional measurement space \mathbf{D} . Measurements are represented as numerical data with a physical dimension of photon counts.

In conclusion, \mathbf{d} is the collection of remotely sensed multispectral data at a fixed time.

Plotting all the feature vectors of an image pair yields a 2D scatter plot of many scatter plot points (Figure 4). A 2D scatterplot provides information about pixel value pairs that occur within a two-band image. Note that some combinations will occur more frequently.




Figure 3. Plotting of the pixel values of a GRC in the feature space for a two-band and three-band image.

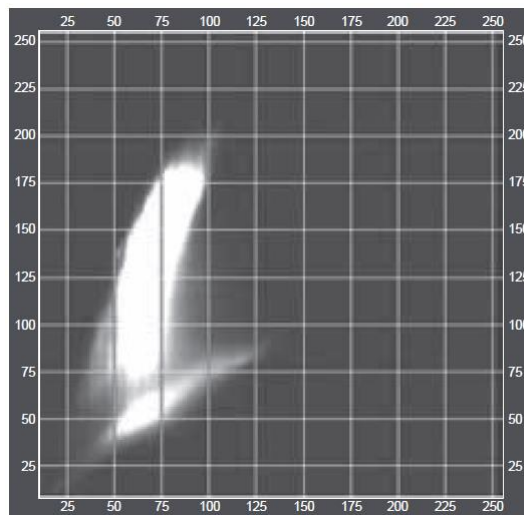


Figure 4. Scatterplot of two bands of an RS image. Note the units along the x- and y-axes. The intensity at a point in the feature space is related to the number of cells at that point.

3.3 Spectral features

Image classification aims to assign pixels to specific object classes based on their spectral features (i.e., their feature vectors) by comparing them to predefined clusters in feature space.

Mapping to RGB Channels for Display: Computer image displays typically convert image data into a continuous, analog image using red, green, and blue (RGB) primary screen colours. To achieve this, three bands of a multispectral image are processed by hardware Look-Up Tables (LUTs) that convert photon counts (their integer DNs) into the Red, Green, and Blue display channels. With a 24-bit display system, each band is assigned to one of three 8-bit integers corresponding to the display colours (R, G, B). However, RGB is **not perceptually uniform** — similar numerical differences don't necessarily mean similar visual differences.

When considering an image with three RGB spectral bands, these three bands form a three-dimensional feature vector. This vector can be visualised and plotted in a three-dimensional graph, which is referred to as a feature space or a scatter plot. This concept is directly analogous to an "RGB cube space."

3.3.1 Cube Spaces

Images are often represented in the **RGB** colour space (Red, Green, Blue). Each pixel has three values.

The RGB colour cube is a conceptual representation of the RGB (Red, Green, Blue) data. It functions as a dataspace for images that are captured by 3-band sensors, such as RGB cameras. Here's a breakdown of what the colour cube entails:

The RGB cube space is a three-dimensional feature space specifically designed for visualising and processing three-band remote sensing data, where each axis represents the scaled photon counts of the Red, Green, or Blue bands.

Dimensions and Data Cells When RGB data is scaled to a byte range (0 to 255), the colour cube can be thought of as containing 256^3 , or 2^{24} , data cells or 3-dimensional bins.

This "colour cube" visually represents all possible combinations of these Red, Green, and Blue photon count values. The axes of this 3D space correspond to the intensity of the red, green, and blue components. For an 8-bit per channel system, this cube ranges from (Red=0, Green=0, Blue=0) (black) to (Red=255, Green=255, Blue=255). Black (RGB = 0) represents a common point at the origin of this cube, and the maximum values (e.g., 255 for each channel) represent saturation levels determined by scaling the photon

counts. This "colour cube" is a specific type of feature space for three-band data, allowing the visualisation of spectral information in three dimensions, similar to how a two-band image can be plotted in a 2D scatter plot.

The fundamental principle underlying the use of such a space for image classification is that different materials (or land cover types) exhibit distinct spectral characteristics. Consequently, spectrally similar pixels tend to group, forming compact clusters within this multi-dimensional feature space, including the RGB cube space.

Mapping to Other Colour Spaces:

To improve segmentation and clustering performance, images are often converted from RGB to alternative colour spaces (Table 1).

The content and surface of the RGB colour cube can be mapped into other coordinate systems. For instance, it can be separated into:

Intensity (I): This is a multiplicative factor, typically the sum of the Red, Green, and Blue photon counts (scaled to a numerical range), such as $(\text{Red} + \text{Green} + \text{Blue})/3$. Illumination conditions, such as shadows and shading, can influence this "intensity" factor.

Orthogonal Colour Features (m1, m2): These are two orthogonal axes that define a colour triangle within the sum-normalised colour space. This m1m2I transform is preferred over systems like HSV (Hue, Saturation, Value) for quantitative analysis because m1 and m2 are mathematically independent and do not suffer from noise amplification at low saturation, unlike hue. The concept of using orthogonal axes is generalizable to any number of spectral bands.

Purpose in Analysis: The transformation of raw RGB data (which can be correlated and redundant) from the colour cube into features like m1, m2, and I is a form of feature extraction. The goal is to separate possible information from redundant data and provide a more stable and robust representation of an object's relative reflectance characteristics, which is crucial for tasks like object classification.

Table 1. To improve segmentation and clustering performance, images are often converted from RGB to alternative colour spaces:

Colour Space	Components	Why Use It in Segmentation
I	Intensity (I)	This is a multiplicative factor, typically the sum of the Red, Green, and Blue photon counts (scaled to a numerical range), such as (Red+Green+Blue)/3.
HSV	Hue, Saturation, Value	Separates colour from intensity; Hue is great for identifying colours regardless of lighting.
m1m2I	Orthogonal Colour Features (m1, m2), and Intensity (I)	The m1m2 features are a pair of orthogonal colour features that, together with Intensity (I), form an alternative coordinate system for representing RGB (Red, Green, Blue) data.
Normalized RGB	$r = R/(R+G+B)$, etc.	A small constant (e.g., 0.001) is often added to the denominator to prevent division by zero.

4. Key Concepts

Radiometric classes: Image pixel classification deals with samples (or pixels) to be classified. It is assumed that there exists a finite set of possible radiometric classes (class labels)

$$\Omega = \{\omega_1, \omega_2, \dots, \omega_k, \dots, \omega_K\}$$

The elements of ω_k of Ω are called radiometric classes, and Ω is the set of radiometric classes. The number of classes to be distinguished is

$$K = \text{card}(\Omega) = \text{number of classes}$$

The K classes ω_k are represented by a distinct name or class label, defined by the community of GIS users.

The example of agricultural crop classification in a satellite image of a farming area in the Flevopolder, the Netherlands, illustrates the idea of a set of classes. The image is recorded by the Thematic Mapper sensor of the Landsat satellite, which measures reflected sunlight in six spectral bands (visible and infrared) with a spatial resolution of 30 m (Figure 5). Let Ω be the seven dominant crops defining the classes of the classification, i.e.

$$\Omega = \{\text{Grass, Wheat, Potatoes, Sugar-beets, Beans, Peas, and Onions}\}$$

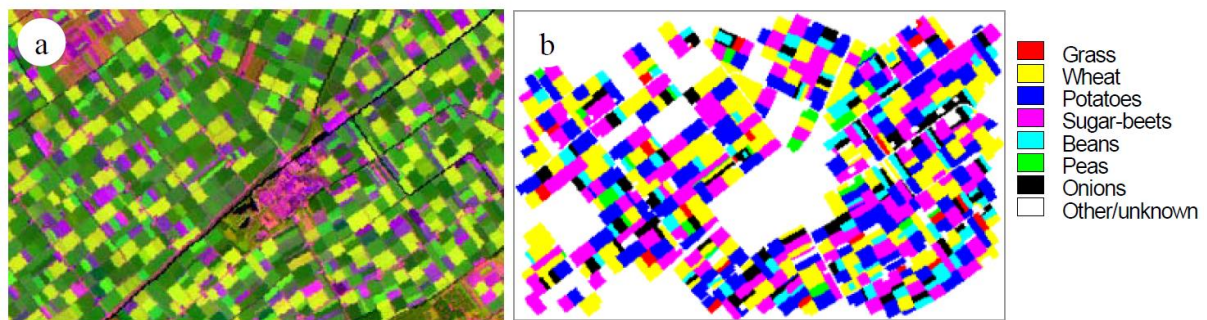


Figure 5. Detection and classification of agricultural fields in RS images. (a) The Flevopolder Landsat-TM images, obtained on July 7, 1987, of the Biddinghuizen area, in Bands 5, 4, and 3 (RGB). (b) Land cover map of 1987, showing seven dominant crops.

The image classification task is then the task of classifying and recognising crops in the area. This specific definition of Ω excludes every other class; a village, a few canals, some forested areas, roads, and farmhouses, although present in the area, are not included in the classification results. Together, these constitute the unknown class ω_0 .

The set Ω of classes may be defined quite differently depending on the application. The definition of Ω is predominantly a question of application.

Thematic maps: Image classification is the process used to produce 'thematic maps' from imagery. The themes can range, for example, from categories or object classes such as soil, vegetation, and surface water (also referred to as the **Radiometric class** above). Therefore, to produce a thematic map, the class 'Labels' at each pixel replaces photon counts.

Similarity between pixels, or groups of pixels: In classification, for example, we want to label areas on the ground that have similar physical characteristics. We do this by grouping data with similar characteristics.

Parametric or Nonparametric Classification: Classification algorithms can be categorised into two types: parametric and nonparametric. Parametric algorithms assume a particular statistical distribution of class, commonly the normal distribution, and require estimates of the distribution parameters from the training sample, such as the mean vector and covariance matrix, for classification. Nonparametric algorithms make no assumptions about the probability distribution and are often considered robust because they may work well for a wide variety of class distributions; however, they require a considerable number of training samples.

Formation of Clusters: The core idea is that pixels originating from spectrally similar areas or land cover types (classes) form compact clusters or groups. These groupings are known as "clusters" and are defined as adjacent points in

the feature space. For example, feature vectors of water areas typically form a compact cluster, as do those for grass or trees.

Training for defining these distinct clusters: Typically, distinct clusters that correspond to different radiometric classes are identified in the feature space through a "training" process. Each feature vector from a multi-band image can be plotted within this space. This principle enables a comparison of each pixel against these predefined clusters, allowing for assignment to the class that fits best.

The definition of the clusters is an interactive process and is carried out during the training process. Comparison of individual pixels with clusters is performed using classifier algorithms.

Training involves selecting pixels to train the classifier to recognise the desired classes, and classification involves determining decision boundaries that partition the feature space according to the properties of the training pixels. This step is either supervised by the analyst or unsupervised with the aid of a computer algorithm.

In principle, the training sample is a collection of $[\mathbf{d}, \omega]$ representative of the application, i.e. measurement vectors of known identity, which are assumed to be representative of the classes of interest. **Adequate training** samples must be available to estimate the parameters of the probability distribution.

Training samples are collected using regions of interest (ROI). ROI are selected on:

- Define the number of classes.
- Class-homogeneous sites, without mixtures among classes.
- Representation of full within-class variability.
- Often, more than one site or area is needed per class.
- Some classes may have a small spatial extent, e.g., an 'asphalt' road.
- Analyse training samples before classification.

A multi-dimensional graph where measurement vectors \mathbf{d} are plotted. Each pixel corresponds to one point in feature space. Clusters of points in feature space correspond to distinct object classes.

Figure 6 illustrates a 2D feature space where clusters of five specific land cover types (classes), such as Water and Trees, are plotted.

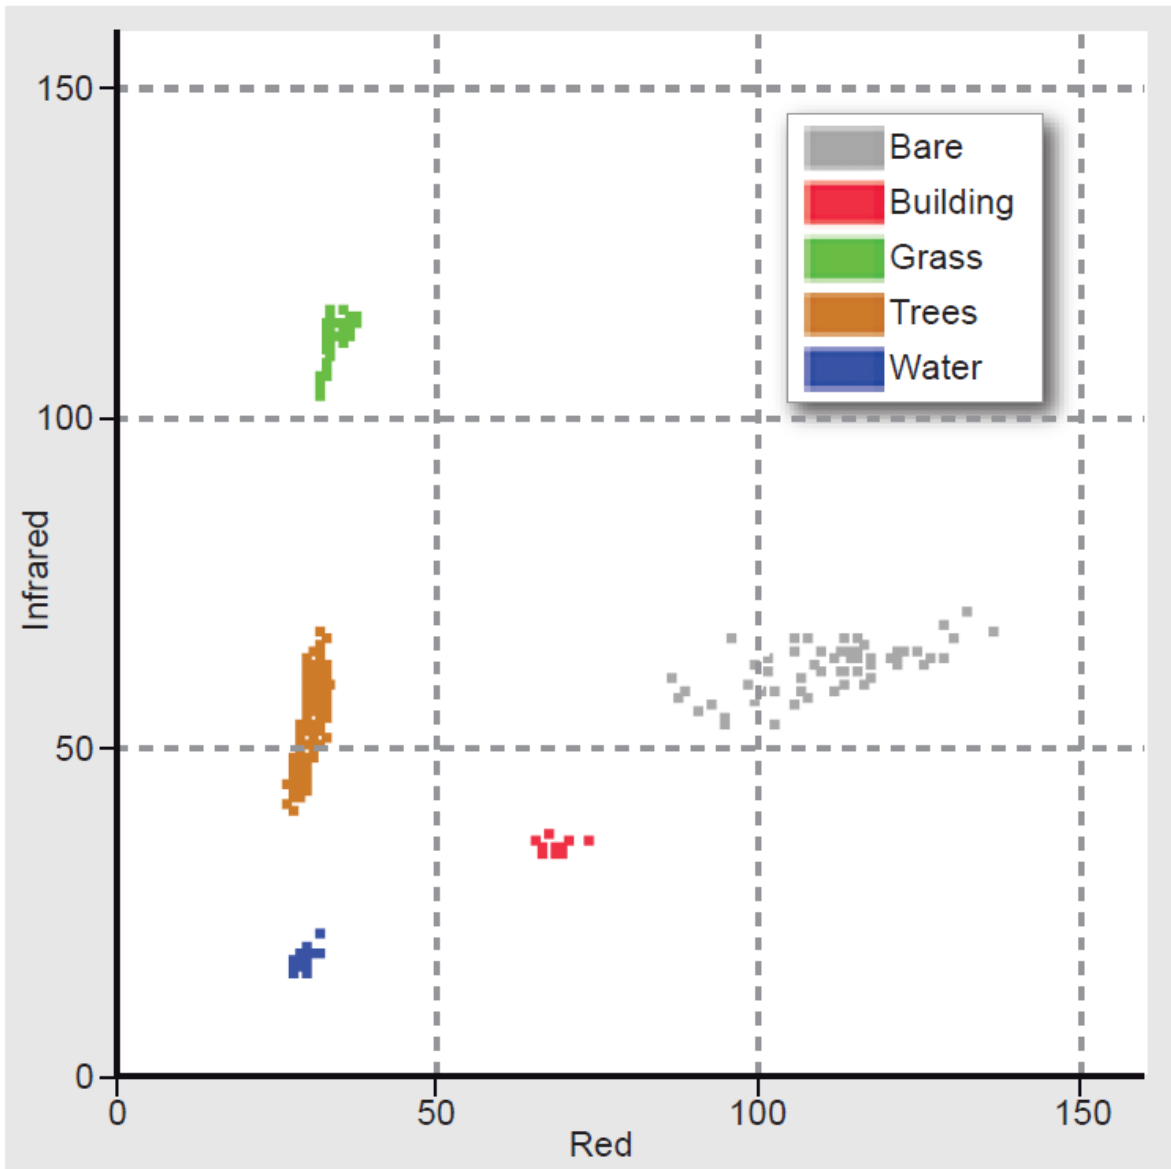


Figure 6. A 2D feature space showing the respective clusters of five classes; note that each class occupies a limited area in the feature space.

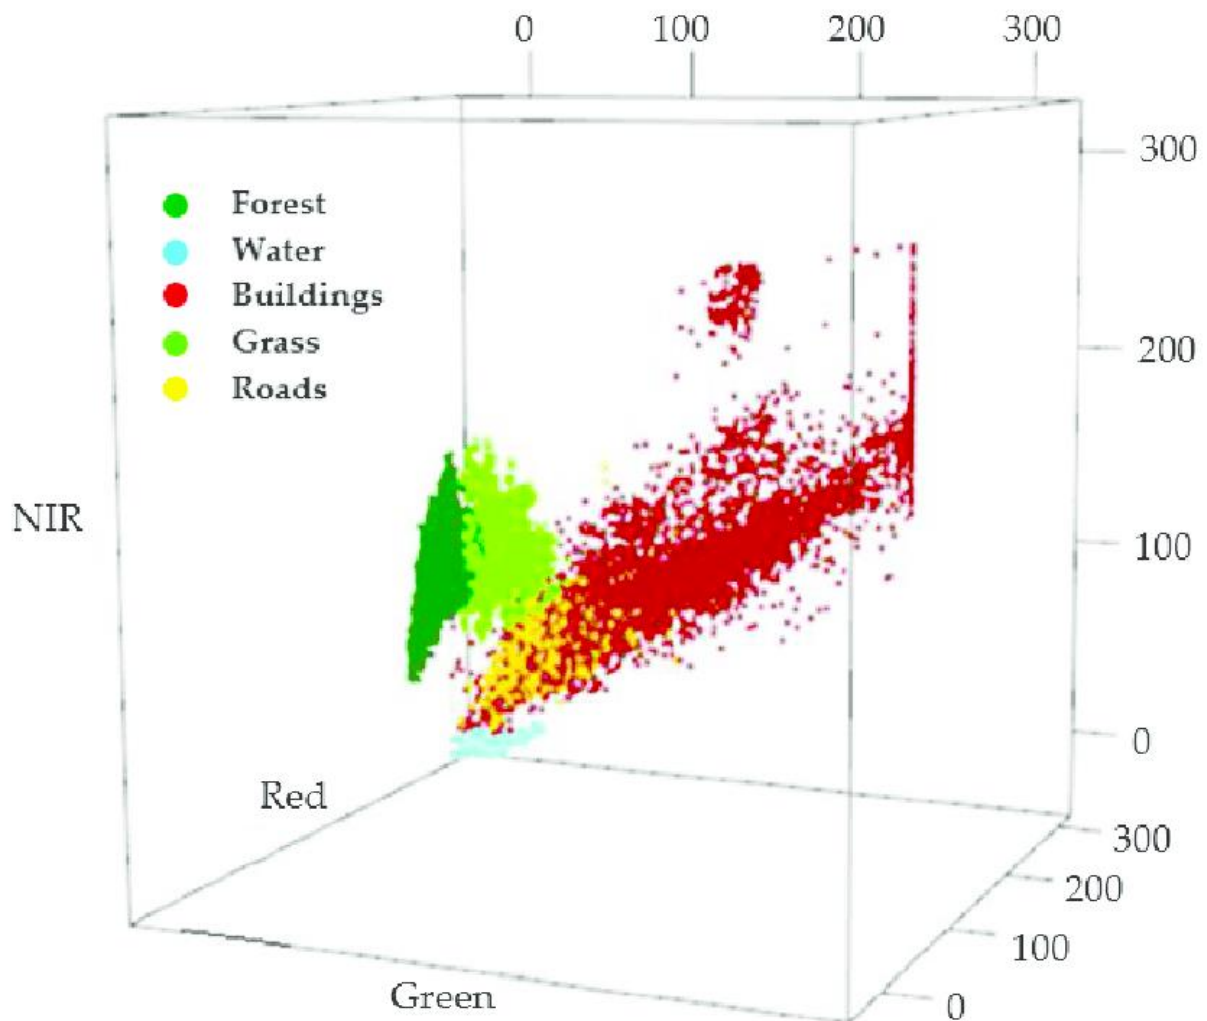


Figure 7. illustrates a 3D feature space where the training pixels for five land cover classes in the Green, Red, and NIR are plotted.

Level-Slice Classifier: This classifier, also known as a **box or parallelepiped classifier**, is the simplest classification method. A set of K-dimensional boxes, centred at the estimated class mean vectors, are placed in N-dimensional feature space (N is the number of spectral bands). If an unlabeled pixel vector lies within one of the boxes, it is assigned that class label. Specification of the box limits is typically in terms of the data extent in each dimension, for example, ± 1 standard deviation about the mean in each band. The delineation of the boxes can also be done directly by the analyst in feature space in an interactive manner. Since the boxes are aligned with the data axes, classification labelling of the whole image can be achieved quickly with hardware or software look-up tables (LUTs), and the resulting map can be viewed simultaneously while manipulating the feature space boxes. A complication occurs if a pixel vector falls within two or more boxes (the boxes can overlap unless that is explicitly prohibited). A decision on a pixel's label must then be made with another algorithm, such as the nearest mean.

By its nature, the level-slice algorithm also creates an “unlabeled” class referred to as the Unclassified pixels, consisting of all pixel vectors that do not fall within any of the designated boxes.

Hard versus Soft Classification: Labelling of pixels accomplished by partitioning the feature space.

- **Hard classification** results in one class per pixel.
- **Soft classification** results in multiple classes per pixel, each with an associated likelihood.
- Soft classification is more accurate and descriptive of reality, accommodating within- and between-class variation, as well as class mixing.

Density slicing: classification using a single band. In theory, it is possible to base a classification on a single spectral band, using single-band classification. DS is a technique whereby the photon counts distributed along the horizontal axis of an image histogram are divided into a series of user-specified intervals or slices. The number of slices and the boundaries between the slices depend on the different object classes in the study area.

Multi-spectral classification: classification using many bands.

5. The Classification Process

The process of image classification typically involves five steps (Figure 8):

- a) Selection and Preparation of RS Images: Choosing appropriate sensor, acquisition date(s), and wavelength bands based on the classification goal. This includes consideration of band correlation, where similar spectral reflection in two bands can lead to redundant information and disturb the classification.
- b) Definition of Clusters in Feature Space: This can be done via supervised or unsupervised methods. In supervised classification, the operator defines the clusters during the training process. In contrast, unsupervised classification involves a clustering algorithm that automatically identifies and determines the number of clusters in the feature space.
- c) Selection of the Classification Algorithm: Once the spectral classes have been defined in the feature space, the operator needs to decide on how the pixels (based on their feature vectors) are to be assigned to the classes. The assignment can be based on different criteria.
- d) Running the Actual Classification: Once the training data have been established and the classifier algorithm selected, the actual

classification can be carried out. This means that, based on its DN's, each "multi-band pixel" in the image is assigned to one of the predefined classes (Figure 9).

- e) Validation of the Result: Once the classified image has been produced, its quality is assessed by comparing it to reference data (ground truth). This requires the selection of a sampling technique, the generation of an error matrix, and the calculation of error parameters

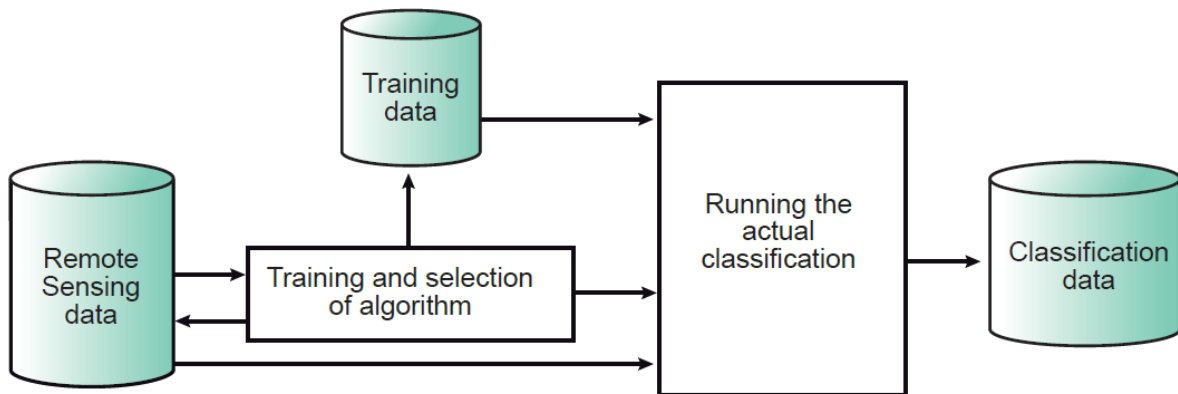


Figure 8. The classification process; the most important component is the training, in combination with the selection of the algorithm.




Figure 9. The result of the classification of a multi-band image (a) is an image in which each pixel is assigned to some thematic class (b).

Supervised Classification: An operator defines the spectral characteristics of classes by identifying "sample areas (training areas)". This requires the

operator to be "familiar with the area of interest" and know where to find the classes, often through field observations. The principle is that a pixel is assigned to a class by comparing its feature vector to these predefined clusters in feature space. Training samples are collected using "regions of interest (ROI)" which are selected in homogeneous areas.

In supervised classification, the statistical data of classes play a pivotal role in characterising clusters within the feature space, which is the foundation upon which image classification operates. The statistical data (such as means and covariance matrices) are estimated from training samples. These training samples should be representative of the class and cover its within-class variability.

The clusters should not, or only partially, overlap with one another; otherwise, a reliable separation is not possible. For a specific data set, some classes may have significant spectral overlap, which, in principle, means that these classes cannot be discriminated by image classification. Solutions include adding other spectral bands and/or images acquired at different times (Figure 10).





Figure 10. Visualising Clusters and Overlap. Shows the spectral overlaps of the clusters in the feature space for supervised classification. Spectral overlap occurs when these clusters, representing different classes, are not distinctly separated in the feature space. This means that some pixel value

combinations for one class might fall into the region predominantly occupied by another class, making their differentiation challenging.


In the supervised classification, the derived cluster statistics are then used to classify the complete image using a selected classification algorithm.

Unsupervised Classification is used when knowledge of the area is insufficient or classes are undefined. Clustering algorithms automatically "partition the feature space into several clusters" based on spectral similarities. A common approach involves the user defining the maximum number of clusters; the computer then locates arbitrary mean vectors as cluster centres, assigns pixels to the nearest centre, and recalculates centres iteratively until they stabilise. Clusters with too few points can be eliminated or merged.


Figure 11 illustrates the results of a clustering algorithm applied to a dataset. Note that the cluster centres coincide with the high-density areas in the feature space.



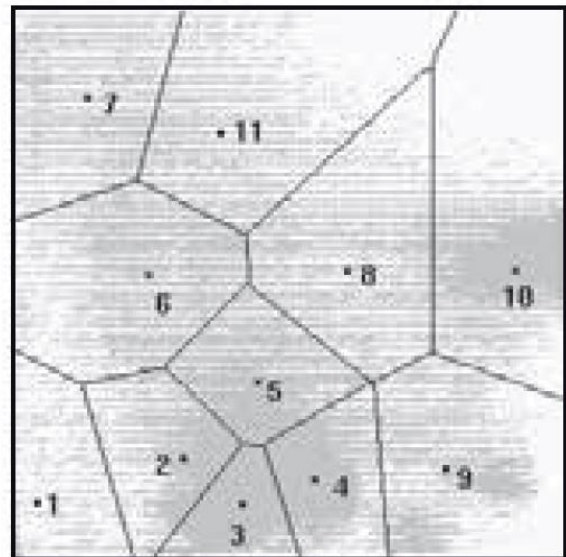
Iteration 0



Iteration 1



Iteration 3



Iteration 10

Figure 11. The subsequent results of an iterative clustering algorithm on a sample dataset.

Classification Algorithms:

After the training sample sets have been defined, classification of the image can be carried out by applying a classification algorithm. Several classification algorithms exist. The choice of the algorithm depends on the purpose of the classification and the characteristics of the image and training data. The operator needs to decide if a reject or unknown class is allowed. In the following, three classifier algorithms are described.

First the box classifier is explained—its simplicity helps in understanding the principle. In practice, the box classifier is hardly ever used, however; minimum

distance to mean and the maximum likelihood classifiers are most frequently used.

Box Classifier (Level-Slice/Parallelepiped): Simplest, defines upper and lower limits for each band and class, creating box-like areas in feature space. Pixels falling within a box are assigned its label. Can result in "unknown" or "reject" classes, and overlap can cause arbitrary assignment. The number of boxes depends on the number of classes. During classification, every feature vector of an input (two-band) image will be checked to see if it falls in any of the boxes. If so, the cell will get the class label of the box it belongs to. Cells that do not fall inside any of the boxes will be assigned the "unknown class", sometimes also referred to as the "reject class" (Figure 12).

The disadvantage of the box classifier is the overlap between classes. In such a case, a cell is arbitrarily assigned the label of the first box it encounters.

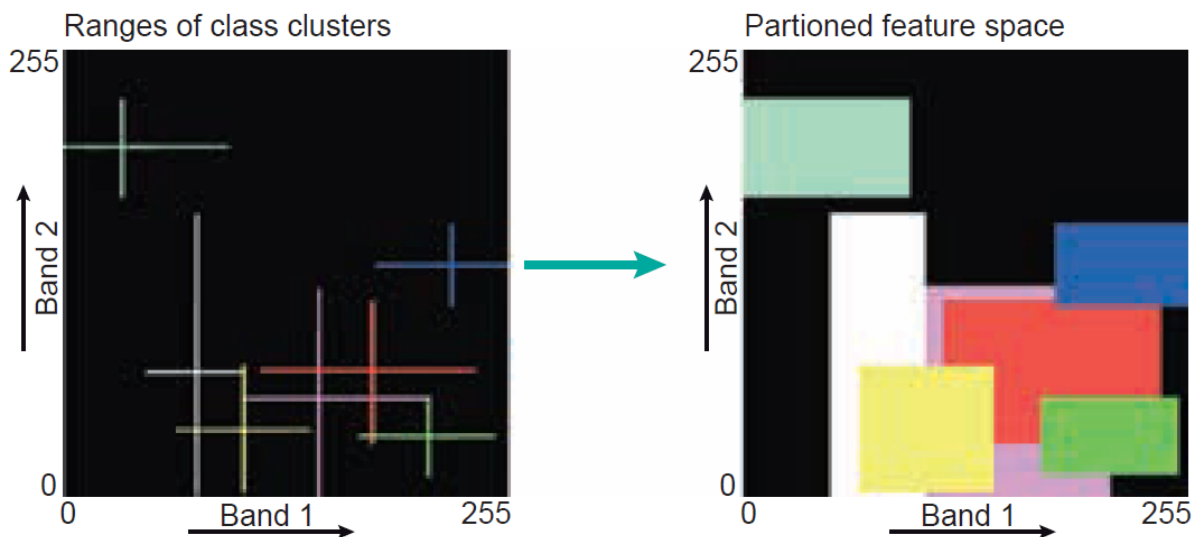


Figure 12. Principle of the box classification in a case of two-dimensional feature space partitioning.

Minimum Distance to Mean (MDM): Assigns a pixel to the class whose cluster centre (mean) is closest. Does not account for class variability. A threshold distance can be set to prevent the assignment of distant points.

Figure 6.14 illustrates how a feature space is partitioned based on the cluster centres. One of the disadvantages of the MDM classifier is that points that are at a large distance from a cluster centre may still be assigned to this centre.

This problem can be overcome by defining a threshold value that limits the search distance. Figure 13 illustrates the effect; the threshold distance to the centre is shown as a circle.

A further disadvantage of the MDM classifier is that it does not account for class variability: some clusters are small and dense, while others are large and dispersed. Maximum likelihood classification, however, does take class variability into account.

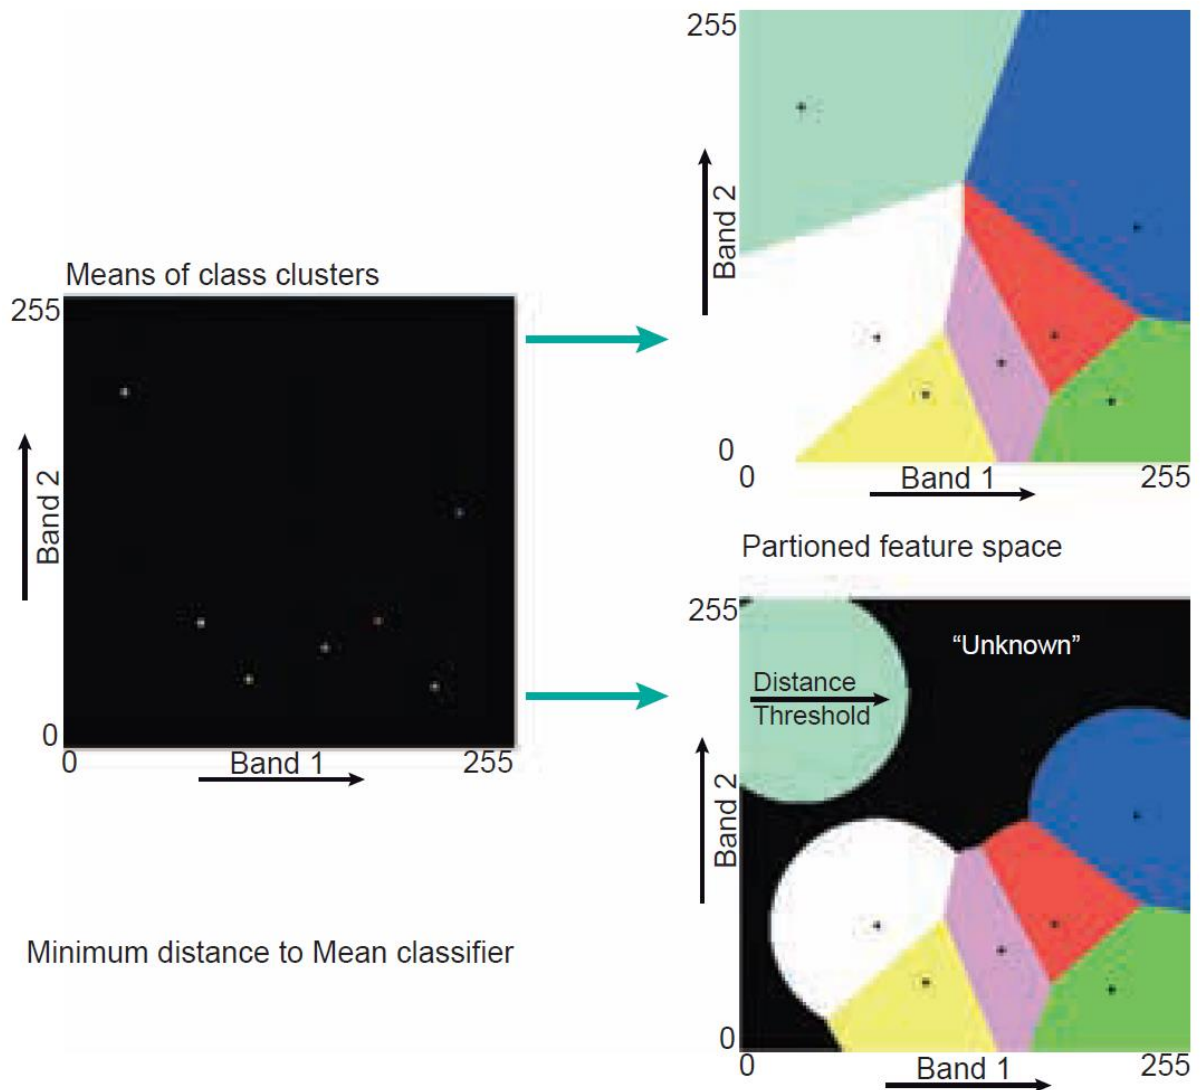


Figure 13. Principle of the minimum distance to mean classification in a two-dimensional situation. The decision boundaries are shown for a situation without threshold distance (upper right) and one with threshold distance (lower right).

Maximum Likelihood (ML): Most frequently used, considers not only cluster centres but also shape, size, and orientation of clusters using probability. Assigns pixel to class with the highest probability. Most ML classifiers assume that the statistics of the clusters follow a normal (Gaussian) distribution.

Maximum likelihood also allows the operator to define a threshold distance by specifying a maximum probability value. A small ellipse centred on the

mean describes the values with the highest probability of membership of a class. Progressively larger ellipses surrounding the centre represent contours of probability of membership to a class, with the likelihood decreasing the further away from the centre. Figure 14 illustrates the decision boundaries for a scenario with and without a threshold distance.

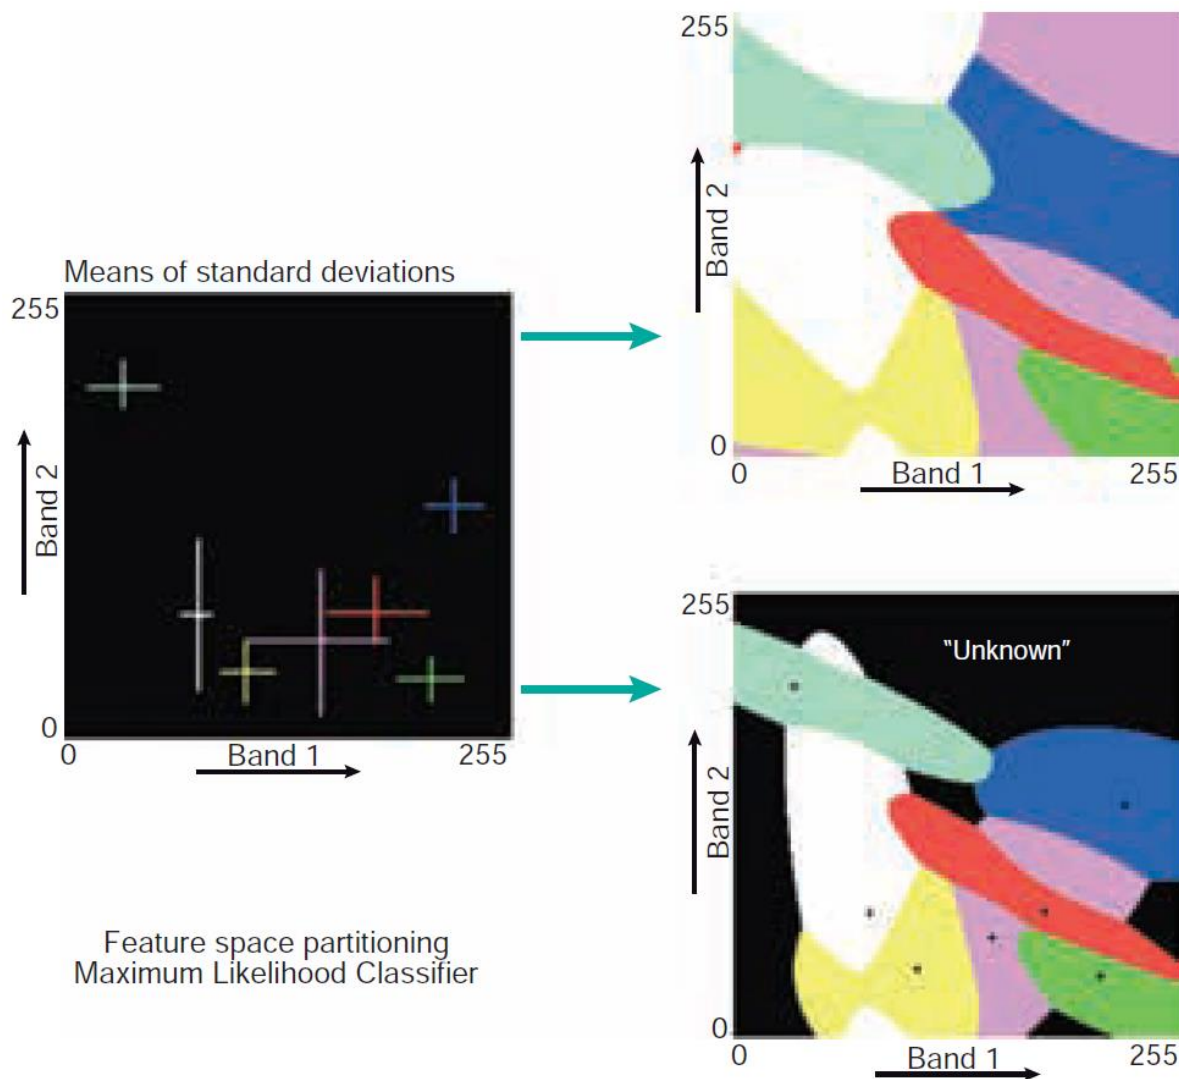


Figure 14. Principle of maximum likelihood classification. The decision boundaries are shown for a situation without threshold distance (upper right) and one with threshold distance (lower right).

Iterative Optimisation (Migrating Means/ISODATA): Estimates initial cluster centres, assigns pixels to nearest centre, recalculates means, and repeats until centres stabilise. The user specifies the number of clusters. Allows for merging and splitting of clusters based on criteria like minimum points or elongated shapes.

Validation of the classification results

Purpose: To check the actual quality and reliability of the classification result.

Methodology: Comparing classified output to "true class" data (ground truth or higher accuracy sources like aerial photos).

Sampling Schemes: Strategies for selecting pixels to test, including simple random sampling and stratified random sampling. Choices involve design, number of samples (determined by sampling theory), and sample unit area.

Error Matrix (Confusion/Contingency Matrix): A comparison is made by creating an error matrix from which various accuracy measures can be calculated. The actual class is preferably derived from field observations. Sometimes, sources with higher accuracy, such as drone images, are used as a reference.

Error matrix is a table comparing classified classes to reference classes.

Once sampling for validation has been carried out and the data collected, an error matrix, also sometimes called a confusion matrix or a contingency matrix, can be established (Table 6.2). In the table, four classes (A, B, C, D) are listed. A total of 163 samples were collected. The table shows that, for example, 53 cases of A were found in the real world ('reference'), while the classification result yielded 61 cases of a; in 35 cases, they agree.

The first and most commonly cited measure of mapping accuracy is the overall accuracy, or proportion correctly classified (PCC). Overall accuracy is the number of correctly classified pixels (i.e. the sum of the diagonal cells in the error matrix) divided by the total number of pixels checked. In Table 2, the overall accuracy is $(35 + 11 + 38 + 2) / 163 = 53\%$. The overall accuracy yields one value for the result as a whole.

- Overall Accuracy: Total correctly classified samples divided by total samples.
- Error of Omission (Type II Error): Samples from a reference class that were incorrectly omitted from that class in the classification (misclassified as another class). Calculated per column. Corresponds to Producer Accuracy.
- Error of Commission (Type I Error): Samples classified into a particular class that belong to a different reference class. Calculated per row. Corresponds to User Accuracy.
- User Accuracy: Probability that a pixel labelled as a particular class belongs to that class.
- Producer Accuracy: Probability that a reference class is correctly mapped as that class.

Table 2. The error matrix with derived errors and accuracy expressed as percentages. A, B, C and D refer to the reference classes; a, b, c and d refer to the classes in the classification result. Overall accuracy is 53%.

	A	B	C	D	Total	Error of Commission (%)	User Accuracy (%)
a	35	14	11	1	61	43	57
b	4	11	3	0	18	39	61
c	12	9	38	4	63	40	60
d	2	5	12	2	21	90	10
Total	53	39	64	7	163		
Error of Omission	34	72	41	71			
Producer Accuracy	66	28	59	29			

6. Image Classification: Practical

3.1 A brief introduction to the problem and the task addressed

The study area for pistachio image classification is located at the UNDER SUN Company farm in Armenia. This section introduces the study area, along with the available information and datasets.

The Armenian plateau is an important centre of cultivated plant origin, with a comparatively large number of crop wild relatives, numerous varieties of different species occurring in the country, as well as endemic species. Among the diverse plants of economic importance, nut crops hold a special role due to their unique and interspecific variability in their wild forms. The only wild species of Pistachio reported in the literature and grown in Armenia is *Pistacia mutica* Fisch. & Mey (*Pistacia atlantica* Desf. subsp. *mutica* (Fisch. et C.A.Mey.) Rech.f. (Anacardiaceae)). The scientific explorations attest to the cultivation of pistachios in Armenia's territory dating back to ancient times. Results of geo-mineralogical research of cuts of the boreholes put in Masriksky plain, and natural exposures of Sevan intermountain hollow have revealed prints of leaves and remains of pistachio plants (along with remains and prints of such plants as liana, juniper, oak, beech, pine) that allow to date them as Upper Sarmatian.

In various parts of Armenia, pistachio groves were established in the 1970s, which had significant soil protection benefits. The established areas are still preserved, despite their neglected condition. The "Under Sun" company restores a part of these areas, creating nurseries from various pistachio varieties.

In recent years, industrial pistachio orchards have been established in Armenia in the following regions: the Syunik region (10 ha), the Ararat region

(20 ha), and the Armavir region (approximately 1000 ha). As well as small gardens in the Tavush and Kotayk regions.

"UNDER SUN" company started the establishment of a pistachio orchard in 2021 on an area of 150 hectares. Currently, the number of planted trees exceeds 60,000. "Pistachio natural" or "Vera" was chosen as the rootstock, which is maximally adapted to the local climatic conditions. This is evidenced by the previously planted almost 50-year-old pistachio soil protection groves. Currently, the trees are in the phase of continuous grafting. The orchard is based on a 3.5 x 6 m scheme with a 10% pollinator distribution. The first harvest is expected in 2025-2026.

According to soil types, the region includes sandy, sandy-clayey, clayey-sandy and clayey soils. Previously, the lands were not cultivated, so deep reclamation works were carried out.

According to the chemical composition of the soil, they are characterised by:

- With a very low content of organic matter, humus, not exceeding 0.8%.
- With a low content of nutrients, except for K, which in some cases reaches the minimum required level.
- With a high pH reaching up to 9.
- With a high content of dissolved salts (EC).

Pistachio orchard is prone to some diseases, which are:

- Septoria - expressed by drying of leaves and early leaf fall.
- Gray spotting - leaf spotting occurs.
- Curling - causes yellow-brown spots on the leaves.
- Powdery mildew - produces a white cotton-like coating.
- Root rot - manifested by a white or grey covering of the roots, especially in the root nodule, and dead cells.
- Verticillium Wilt is expressed by partial or complete drying of the plant. In the cross-section of the dried branches, there are strongly pronounced black marks.
- Blossom blight - expressed by the drying of young branches and leaves.
- Bacterial disease expressed by the drying of young branches, leaves and other parts.

Considering the mentioned diseases, preventive measures are planned during the entire vegetation period, which includes root and ecto-root treatments. Increasing the natural immune system of plants with proper nutrition and fertilisation is also a preventive measure.

For this purpose, considering the environment in which the plant grows and develops, fertilisation and nutrition processes are carried out using organic

substances, such as humic acids, fulvic acids, and amino acids, as well as mineral fertilisers and microelements in the form of chelated compounds.

As a result of proper feeding and fertilisation, the plants grow intensively, allowing the garden to be harvested 1-2 years earlier.

Monitoring is also carried out in the garden, by which the following care activities are planned and implemented:

- Irrigation - according to the phenological stage of the tree, soil type, and climatic conditions.
- Weed control is carried out mechanically and chemically using herbicides.
- Mechanical works of the land - loosening, levelling.
- Trimming and shaping - as needed.

- **With this background information, a key deliverable of this project is to prepare a map depicting the pistachio crop fields, which includes identification, classification, and area measurement using recent drone and satellite images for further analysis.**

This case study aims to create an accurate crop map by utilising image classification techniques, drawing on data from Earth observation satellites, UAVs, and ground sensors, including images captured by cell phones.

Additionally, various data from the test area were stored in a GIS (including Google Earth). For example, information about soil, pistachio crop types or varieties. We also have polygon boundaries for each crop field and its corresponding crop type. In this case study, the available information may be used for:

- Training field selection for the classification of the satellite image
- Extraction of the prior knowledge
- Validation of the final results

Note: Any existing errors in the data and/or information used may impact the final results.

3.2 Conduct practical exercises: Selection of images, bands and preparation of data:

- Define the information requirements.

- Mission planning: discuss mission planning to ensure quality data acquisition, processing and analysis. Address the limitations and complexities involved in operations by technicians.
- Access to a growing archive of satellite imagery and platforms, particularly Sentinel-2 data.
- Critical datasets selection for experimentation; the availability of new tools to facilitate image processing and interpretation, including improved EO data selection, visualisation, review, and analysis functions.
- Processing EO Satellite imagery, band selection, cloud masking, and image enhancement.
- QGIS exercises and Downloading Sentinel-2 satellite Imagery. The Copernicus Data Space Ecosystem Browser serves as a central hub for accessing, exploring and utilising the wealth of Earth observation and environmental data provided by the Copernicus Sentinel constellations. To download Sentinel-2 multispectral imagery, see the following tutorial from the Copernicus dataspace: <https://documentation.dataspace.copernicus.eu/Applications/Browser.html>

3.3 Conduct practical exercises in image classification using selected data

- Feature selection
- Remote sensing indices.
- Supervised image classification (Object classes and Training set selection, Maximum likelihood)
- Clustering and unsupervised image classification.
- QGIS exercises and case studies: Learn to create training samples, implement classification, assess classification accuracy, and apply these skills to practical case studies.

3.4 Generate independent, robust, and consistent reporting needs for image classification and estimate pistachio orchards/tree areas and their changes over time at regional, farm, and plot levels.

- Accuracy assessment
- Generate error maps and evaluate differences
- Comparison of supervised and unsupervised results.

3.5 Collaborate with Pistachio growers to do knowledge-based classification.

Measurements of the pistachio canopy's height, length, and width: Data collection should preferably be minimal and hypothesis-based. The hypotheses are supported by a method for collecting and processing the data, as well as feature extraction. The RGB, Depth cameras, or Drone would reduce the amount of work and improve the reliability of the required data.

7. Q&A

Describe the role of "photon counts" in remote sensing data acquisition. How are these physical measurements eventually represented as "Digital Numbers - DN" in an image? Photon counts represent the number of photons captured by a sensor element during an exposure time, indicating the amount of electromagnetic radiation reflected or emitted from a ground resolution cell. These photon counts generate an accumulated charge in the sensor element, which is converted into an analog voltage. This analog voltage is then transformed into a discrete "Digital Number" (DN) by an Analog-to-Digital Converter (ADC) for digital image representation.

Define "image space" and "feature space" in the context of image classification. How do these concepts relate to the photon counts or Digital Numbers (DNs) recorded by an RGB or multispectral sensor? Image space refers to the 2D array of pixels where the spatial distribution of DNs defines the image. Feature space is a multi-dimensional graph where the DNs from different spectral bands for a single Ground Resolution Cell (GRC) are plotted as a "feature vector." Essentially, image space shows the spatial arrangement of DNs, while feature space shows the spectral relationships among them.

Given the challenge of visualising multi-dimensional feature vectors in remote sensing when dealing with a high number of spectral bands, what is a practical solution for displaying the distribution of pixel values? The 2D scatterplots are often used as a practical solution for visualising data when dealing with four or more bands, by plotting all possible combinations of two bands separately.

What is digital image classification in remote sensing, and what are its primary applications?

Digital image classification is a process of assigning pixels in an image to specific classes based on their spectral characteristics. It essentially creates a map where each pixel represents a particular land cover type (class). This is done by analysing the photon counts or DNs of pixels, which represent the

electromagnetic radiation reflected or emitted from object surfaces or ground resolution cells (GRCs) across different spectral bands.

Its primary applications are diverse. For example, it's used in projects to verify farmers' claims for crop subsidies, often followed by field checks to ensure accuracy. The goal is either to classify scene objects (e.g., crops, water bodies, urban areas) or to estimate parameters related to these objects.

Briefly describe the core principle behind the Minimum Distance to Mean (MDM) classification algorithm. What is one notable disadvantage of this method, and how can it be mitigated? The MDM classifier assigns a pixel to the class whose cluster centre (mean feature vector) is spectrally closest to the pixel's feature vector. A disadvantage is that it does not account for the variability or shape of the clusters, potentially assigning points far from a cluster's core to that class. This can be mitigated by defining a threshold distance, beyond which a pixel will not be assigned to any class.

What are the differences between supervised and unsupervised image classification methods?

Supervised Image Classification:

Requires prior knowledge: The operator must be familiar with the area of interest and know where to find the specific classes (e.g., "water," "forest," "urban") within the scene. This knowledge often comes from field observations or existing maps.

Training areas: The operator manually selects "training areas" (Regions of Interest or ROIs) on the image that are representative of each desired class. These ROIs define the spectral characteristics and boundaries of the clusters in the feature space.

Operator-driven: The quality of the classification heavily depends on the operator's skill in selecting accurate and representative training samples.

Direct class mapping: The output directly maps to predefined, user-named classes.

Unsupervised Image Classification:

Does not require prior knowledge (initially): Useful when detailed knowledge of the area or specific classes is insufficient.

Automated clustering: Clustering algorithms automatically identify "natural" groupings or clusters of spectrally similar pixels in the feature space. The user typically specifies the maximum number of clusters.

Iterative process: Algorithms like the iterative optimisation (migrating means) technique (also known as ISODATA) start with arbitrary cluster centres and iteratively refine them by assigning pixels to the nearest cluster and recalculating cluster means until stability is achieved.

Post-classification labelling: The resulting clusters are initially statistical groupings. The user must then interpret and label these clusters (e.g., "Cluster 1 is water," "Cluster 2 is vegetation") based on external information or visual inspection.

Spectral groupings: Primarily produces spectral classes, which may then need to be aggregated into meaningful land cover types (classes).

In essence, supervised classification is "training by example," while unsupervised classification is "clustering by similarity."

Explain the concept of "feature vector" and "clusters" in the context of image classification.

A **feature vector** is a multi-dimensional representation of a single pixel's spectral characteristics. For a multispectral image with 'N' spectral bands, each pixel has 'N' Digital Numbers (DNs), one for each band. This set of 'N' DN's forms a feature vector (e.g., [DN_band1, DN_band2, ..., DN_bandN]). This vector can be plotted as a single point in an N-dimensional "feature space."

Clusters are groupings of these feature vectors (points) in the feature space. The fundamental assumption in image classification is that pixels belonging to the same class (e.g., water, forest, bare soil) will have similar spectral properties. When plotted in feature space, these spectrally similar pixels will naturally group together, forming distinct "clusters." For example, all pixels representing water might form a compact cluster in a specific region of the feature space, while pixels representing bare soil form another, often separate, cluster. The goal of classification algorithms is to identify these clusters and then assign every pixel in the image to the class associated with the cluster its feature vector falls into or is closest to.

How is the quality of image classification results validated and measured?

The quality of image classification results is crucial and is typically assessed through a rigorous validation process, often referred to as accuracy assessment. This involves:

Reference Data (Ground Truth): Comparing the classified image with "ground truth" data. This reference data is preferably derived from direct field observations or from higher-accuracy sources like drone images.

Sampling Scheme: Selecting a statistically representative sample of pixels from the classified image for comparison. Common sampling strategies

include simple random sampling or stratified random sampling, considering factors like the number of samples and the size of the sample unit (point or area).

Error Matrix (Confusion Matrix): Once samples are collected and compared, an error matrix is constructed. This table cross-tabulates the classified classes against the reference (true) classes. For example, if there are classes A, B, C, D, the matrix shows how many pixels truly belonging to A were classified as A, B, C, or D, and so on.

Accuracy Measures: Various accuracy measures are derived from the error matrix:

Overall Accuracy: The total percentage of correctly classified pixels.

These measures provide quantitative insights into the reliability and quality of the classification output.

8. Essay Format Questions

How are image space and feature space fundamental to image classification?

What is the application of 2D scatterplots?

Explain the five main steps involved in the image classification process. For each step, describe its purpose and critical considerations or decisions that an operator must make.

Discuss the different types of classification algorithms (Box Classifier, Minimum Distance to Mean, Maximum Likelihood) used in image classification. Compare their underlying principles, advantages, and disadvantages, and explain in which scenarios each might be preferred.

Elaborate on the importance of validation in remote sensing image classification. Describe the components of an error matrix and explain how the overall accuracy is measured and interpreted.

Analyse the progression from "Measurement" to "Information" to "Knowledge" in remote sensing. Provide specific examples from the provided texts to illustrate each stage and discuss how the reliability of derived "knowledge" (e.g., a classification result) is evaluated.

Explain the term "Overall Accuracy" in the context of an error matrix.

What is an "unknown class" in image classification? Give an example. How is it handled in Classification Algorithms?

9. Glossary of Key Terms

Sensor and Data Concepts

- **Image Elements/Pixels:** Digital representation of sampled object surfaces through the Instantaneous Field of View (IFOV) of sensor arrays.
- **Digital Numbers (DNs):** Photon counts recorded by sensor arrays in distinct spectral bands, which can be converted into spectral features. Range from 0 to 255 for 8-bit data.
- **Ground Resolution Cell (GRC):** The area on the ground represented by a single pixel.
- **IFOV:** The angle of the solid cone from the sensor that intersects an object surface, projecting onto a sensor element. Not constant; it depends on Y (the height of the sensor element) and F (the focal length).

Photon Capture and Digital Numbers

- **Photon Stream/Flux:** The number of photons falling onto or into a sensor element during an exposure time.
- **Electro-Optical Devices:** Imaging sensors are "arrays of cells that convert photons into electronic signals".
- **Digital Numbers (DNs):** The sensor counts photons, and these counts are "Digital Numbers" in distinct spectral bands. This photon count is converted to an analog voltage (from accumulated charge in a capacitor), which is then converted to digital data via an Analog-to-Digital Converter (ADC). For an 8-bit record, DN's range from 0 to 255.
- **Measurement Vector:** For a given image sample position, the set of N measurements from different spectral bands forms an N -dimensional "measurement vector" $\mathbf{d} = [d_1 \ d_2 \ \dots \ d_N]^T$.

Radiometric Properties and Classes

- **Radiometric Measurements:** The primary outputs of a sensor, such as an RGB camera or a multispectral scanner like Sentinel-2.
- **Radiometric Classes:** A finite set of possible class labels ($\Omega = \{\omega_1, \omega_2, \dots, \omega_K\}$) that pixels are assigned to. These are defined by the GIS user community (e.g., {Grass, Wheat, Potatoes, Sugar-beets, Beans, Peas, Onions} for agricultural classification). An "unknown class" (ω_0) encompasses features not included in the defined classes.
- **Box Classifier (Level-Slice Classifier/Parallelepiped Classifier):** The simplest classification method that defines upper and lower limits for each band and class, creating box-like areas in feature space. Pixels falling within a box are assigned its class.

- **Cluster:** A group of adjacent points in feature space that are spectrally similar, representing a distinct radiometric or land cover class.
- **Error Matrix (Confusion Matrix/Contingency Matrix):** A table that compares the results of an image classification to reference (ground truth) data, characterised classification accuracy.
- **Feature Space:** A multi-dimensional graph or scatter plot where the Digital Numbers (DNs) from multiple spectral bands for each Ground Resolution Cell (GRC) are plotted as a single point (feature vector).
- **Feature Vector:** A vector composed of the Digital Numbers (DNs) from different spectral bands for a single Ground Resolution Cell (GRC).
- **Ground Resolution Cell (GRC):** The area on the ground represented by a single pixel in a remote sensing image.
- **Hard Classification:** An image classification method that assigns each pixel exclusively to one single class.
- **Image Space:** The 2D array of pixels that constitutes a digital image, where the spatial distribution of Digital Numbers (DNs) defines the image.
- **Instantaneous Field of View (IFOV):** The angle or solid cone from a remote sensing sensor that projects onto a single sensor element, determining the area on the ground represented by a pixel.
- **Iterative Optimisation (Migrating Means/ISODATA):** An unsupervised clustering algorithm that iteratively adjusts cluster centres (means) and reassigns pixels until cluster centres stabilise.
- **Kappa (κ) Coefficient:** A statistical measure of classification accuracy that takes into account agreement occurring by chance, providing a more robust assessment than simple overall accuracy.
- **Knowledge:** In remote sensing, based on hypotheses and measured evidence, often expressed as the probability of a hypothesis given the evidence.
- **Maximum Likelihood (ML) Classifier:** A commonly used classification algorithm that assigns a pixel to the class for which it has the highest probability, considering not only cluster means but also the shape, size, and orientation of clusters based on their statistical distributions.
- **Measurement:** The interaction between two objects in space-time that produces data about that interaction, such as a photon sensor measuring photon counts.
- **Multispectral Data:** Remote sensing data collected across multiple distinct spectral bands of the electromagnetic spectrum.
- **Overall Accuracy:** The total number of correctly classified pixels divided by the total number of pixels in the validation sample, typically derived from an error matrix.
- **Pattern:** The spatial arrangement of objects and the characteristic repetition of certain forms or relationships in an image.
- **Photon Flux:** The number of photons falling on or into a sensor element per unit time.
- **Photons:** Fundamental particles of light, carrying electromagnetic energy.

- **Pixel (Image Element):** The smallest unit of a digital image, representing a Ground Resolution Cell (GRC) and containing Digital Number (DN) values for each spectral band.
- **Pixel-based Classification:** An image classification method that processes and assigns each individual pixel to a class based primarily on its spectral information.
- **Radiometric Class:** A distinct category or label assigned to pixels during image classification, based on their spectral characteristics.
- **Region of Interest (ROI):** Sample areas selected on homogeneous regions of an image during supervised classification, used to collect training samples for defining class spectral characteristics.
- **Remote Sensing (RS) Image:** An image acquired by sensors on platforms (e.g., satellites, aircraft) that record electromagnetic radiation reflected or emitted from the Earth's surface.
- **Soft Classification:** An image classification method that assigns multiple classes to each pixel, each with an associated likelihood or probability, to better account for within- and between-class variation and mixed pixels.
- **Spectral Characteristics:** The unique ways in which different materials interact with electromagnetic radiation across various wavelengths (e.g., reflectance, absorption).
- **Spectral Class:** A class defined purely by its spectral characteristics as identified in the feature space, which may or may not directly correspond to a desired land cover type.
- **Spectroscopic Data:** Data that measures the intensity of electromagnetic radiation at different wavelengths, providing detailed spectral signatures of materials.
- **Supervised Classification:** An image classification approach where the operator defines the spectral characteristics of classes by identifying sample areas (training areas) in the image.
- **Training Areas:** Sample areas selected by an operator in supervised classification, representing known land cover types, used to train the classification algorithm.
- **True Class:** The actual class or category of a feature on the ground, used as a reference for validating image classification results.
- **Unsupervised Classification:** An image classification approach where clustering algorithms automatically find spectral groupings (clusters) in the data without prior knowledge or training samples from the operator.
- **Validation:** The process of assessing the quality and reliability of image classification results by comparing them against independent reference data.

Definition of clusters in feature space: Achieved through supervised or unsupervised classification.

Selection of the classification algorithm: Deciding how pixels are assigned to classes.

Running the actual classification: Assigning each multi-band pixel to a predefined class.

Validation of the result: Assessing classification quality against reference data.

Band Selection: Avoiding highly correlated bands (redundant information) and considering hardware/software limitations.

10. Conclusion: Classification methodology

- 1) Information requirements
 - a. Study Area & Data and/or Knowledge: different GIS and Google Earth information and maps
 - b. Application problem and required object classes (understanding objects in the study area, size/area distribution of different classes, spectral overlap, etc.)
- 2) Data acquisition for mapping land cover
 - a. New vs. Existing RS data!!
- 3) Field work
- 4) Analysis & Basic Assumptions
 - a. Selection of image classification method
 - b. Available prior knowledge
- 5) Digital Image Classification
 - a. Training stage
 - b. Classification
 - c. Classification Enhancement (post classification, re-classification, Filters)
- 6) Classification accuracy assessment
 - a. Comparing Classification results and 'Ground Truth'
 - b. calculating various measures of error (e.g., Confusion Matrix)
 - c. Challenges in image classification (e.g., mixed pixels).

Ultimately, the material prepares students to understand, process, and interpret remotely sensed data for various applications, such as precision agriculture and environmental monitoring.

11. Basic literature

Bakker, W. H., Gorte, B. G. H., Horn, J. A., Hussin, Y. A., Huurneman, G. C., Janssen, L. L. F., Reeves, C. V., Shrestha, D. P., Skidmore, A. K., Soeters, R., Weir, M. J. C., & Wodai, T. (1999). Principles of remote sensing: ITC core module:

Version 1: syllabus. International Institute for Geo-Information Science and Earth Observation.

Tolpekin, V., & Stein, A. (Eds.) (2012). The core of GIScience: a systems-based approach. (ITC Educational Textbook Series). University of Twente, Faculty of Geo-Information Science and Earth Observation (ITC).

Abkar, A.A. (1994), "Knowledge-Based Classification Method for Crop Inventory Using High Resolution Satellite Data", MSc. Thesis, ITC, Enschede, The Netherlands.

Tutorial from the Copernicus dataspace:

- <https://documentation.dataspace.copernicus.eu/Applications/Browsing.html>

NASA ARSET: Introduction to Land Cover Classification and QGIS, Conducting Unsupervised and Supervised Land Cover Classifications, Part 1 and 2:

- <https://www.youtube.com/watch?v=ZtSlipKIH8E>
- <https://www.youtube.com/watch?v=LtM4X5J4Oul>