Image Classification

- Why classify?
- Make sense of a landscape
  - Place landscape into categories (classes)
    - Forest, Agriculture, Water, etc
- Classification scheme = structure of classes
  - Depends on needs of users

Example Uses

- Provide context
  - Landscape planning or assessment
  - Research projects
- Drive models
  - Global carbon budgets
  - Meteorology
  - Biodiversity

Example: Near Mary's Peak

- Derived from a 1988 Landsat TM image
- Distinguish types of forest

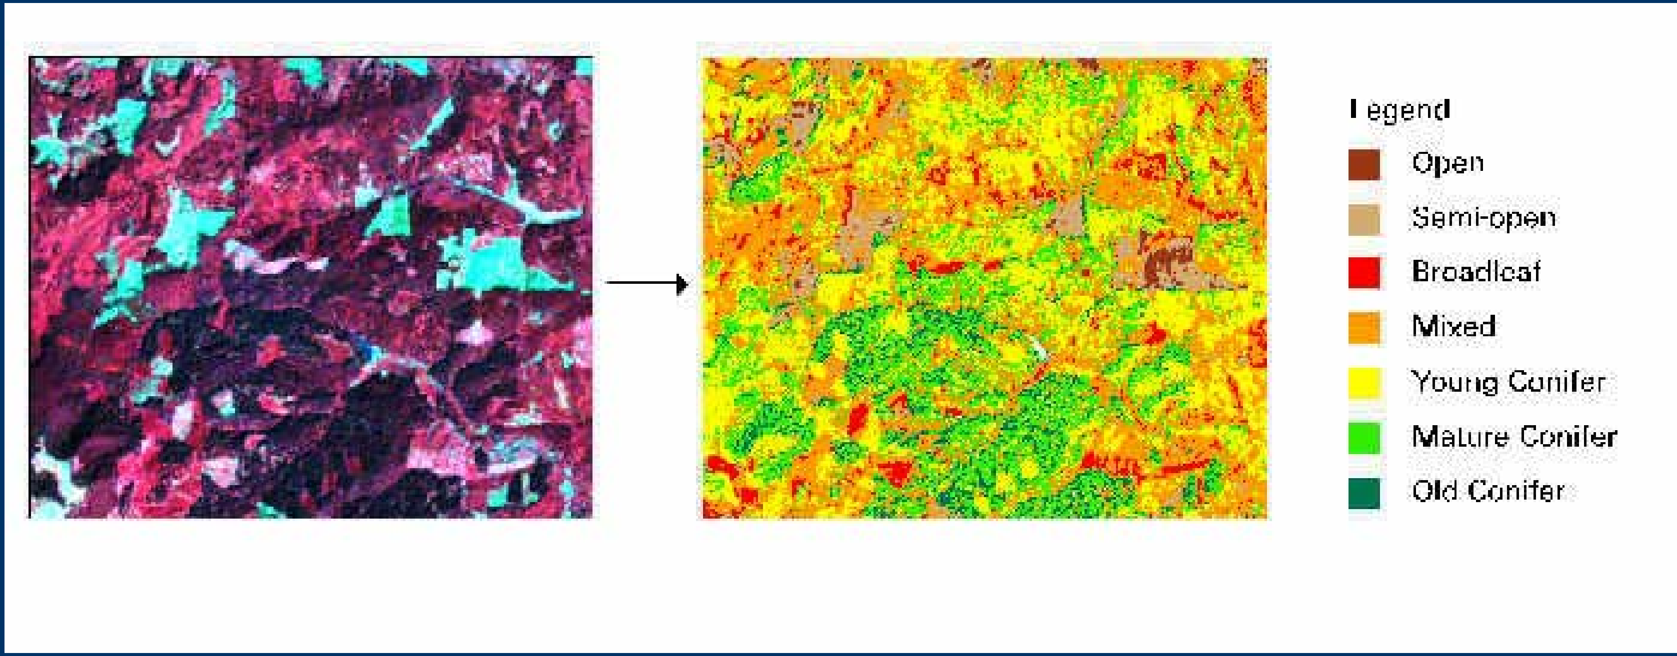

Classification: Critical Point

- LAND COVER not necessarily equivalent to LAND USE
  – We focus on what’s there: LAND COVER
  – Many users are interested in how what’s there is being used: LAND USE
- Example
  – Grass is land cover; pasture and recreational parks are *land uses* of grass

Classification

TODAY’S PLAN

- Basic strategy for classifying remotely-sensed images using spectral information
- Supervised Classification
- Unsupervised Classification
- Lab 4

Next class: Important considerations when classifying; improving classifications; assessing accuracy of classified maps

Basic Strategy: How do you do it?

- Use radiometric properties of remote sensor
- Different objects have different spectral signatures

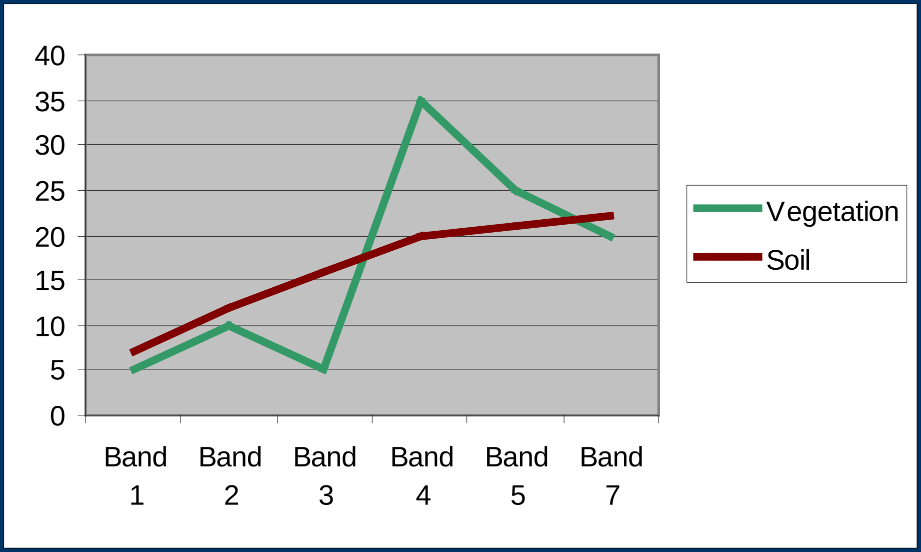

Basic Strategy: How do you do it?

- In an easy world, all “Vegetation” pixels would have exactly the same spectral signature
- Then we could just say that any pixel in an image with that signature was vegetation
- We’d do the same for soil, etc. and end up with a map of classes

Basic Strategy: How do you do it?

But in reality, that isn't the case. Looking at several pixels with vegetation, you'd see variety in spectral signatures.

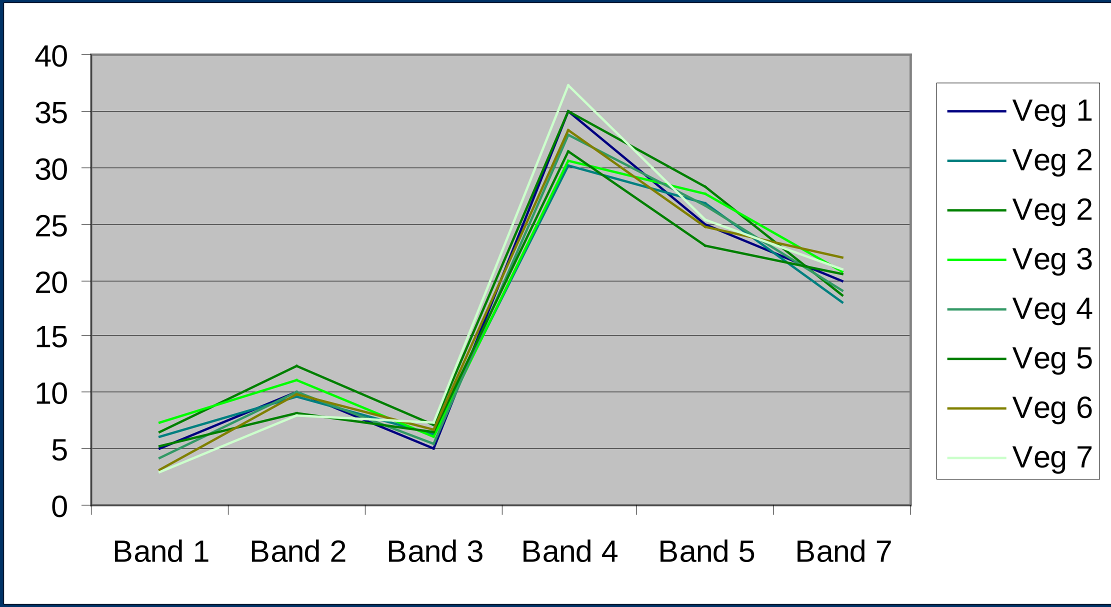

The same would happen for other types of pixels, as well.

The Classification Trick: Deal with variability

- Different ways of dealing with the variability lead to different ways of classifying images
- To talk about this, we need to look at spectral signatures a little differently

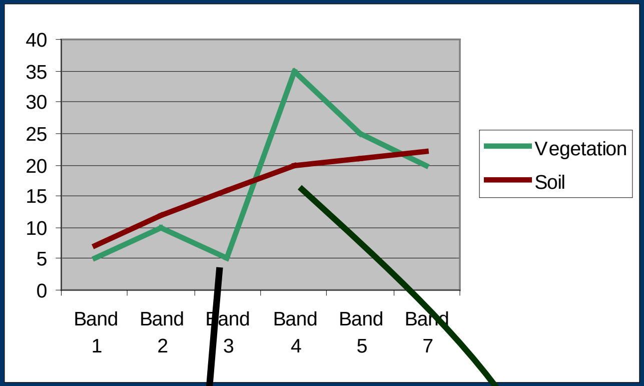

Think of a pixel’s reflectance in 2-dimensional space. The pixel occupies a point in that space.

The vegetation pixel and the soil pixels occupy different points in 2-d space

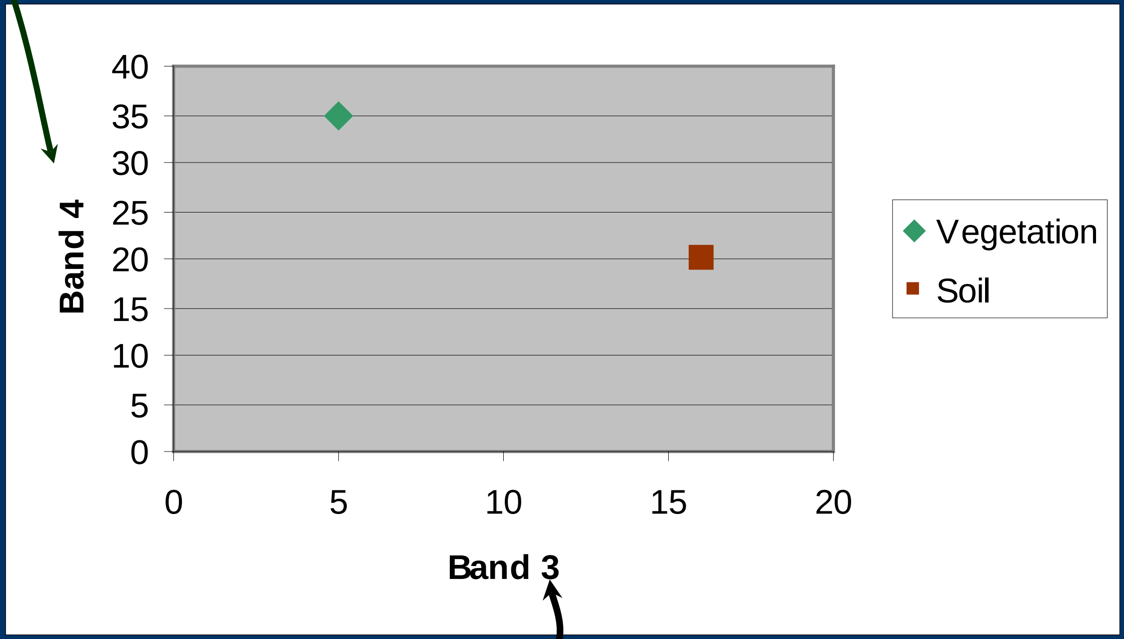

- In a Landsat scene, instead of two dimensions, we have six spectral dimensions
- Each pixel represents a point in 6-dimensional space
- To be generic to any sensor, we say **“n-dimensional”** space
- For examples that follow, we use 2-d space to illustrate, but principles apply to any n-dimensional space

Feature space image

- A graphical representation of the pixels by plotting 2 bands vs. each other
- For a 6-band Landsat image, there are 15 feature space images

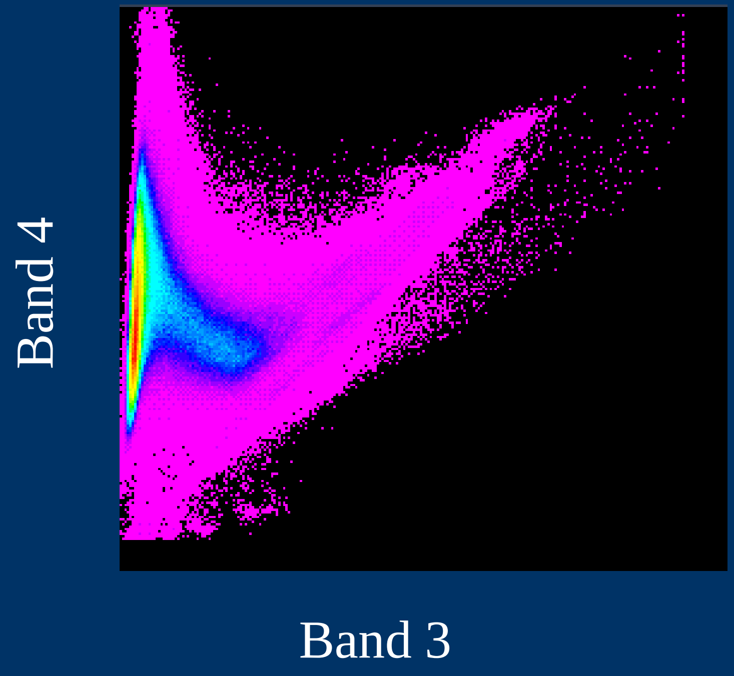

Basic Strategy: Dealing with variability

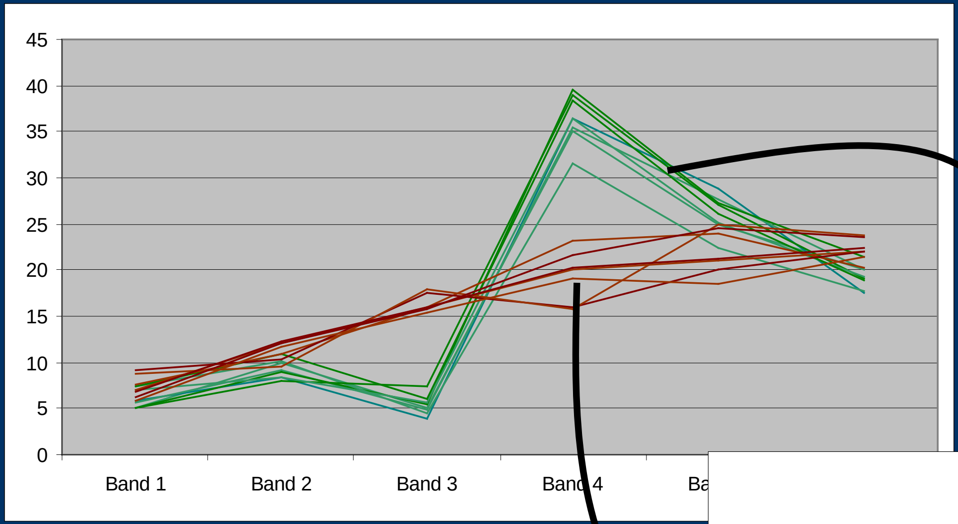

With variability, the vegetation pixels now occupy a region, not a point, of n-dimensional space

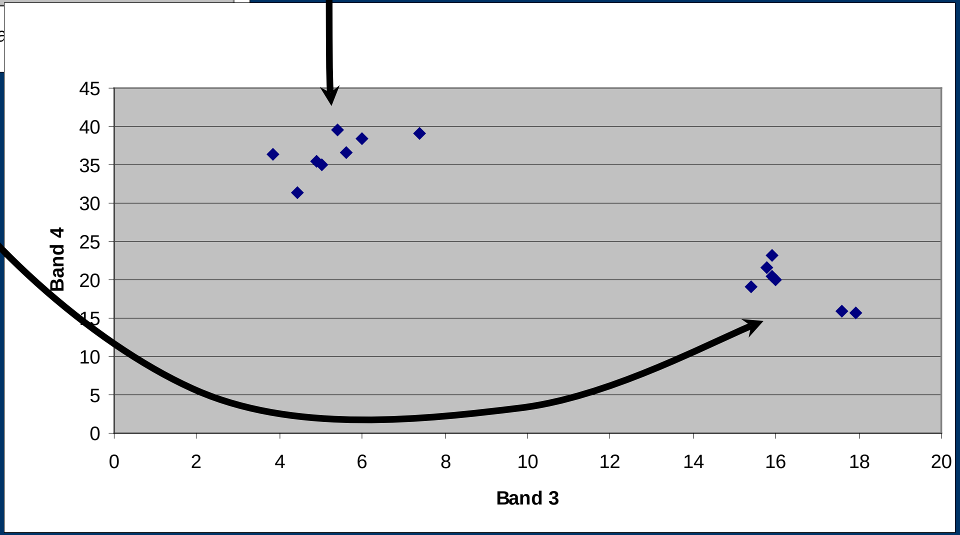

Soil pixels occupy a different region of n-dimensional space

Basic strategy: Dealing with variability

- Classification:
  - Delineate boundaries of classes in n-dimensional space
  - Assign class names to pixels using those boundaries

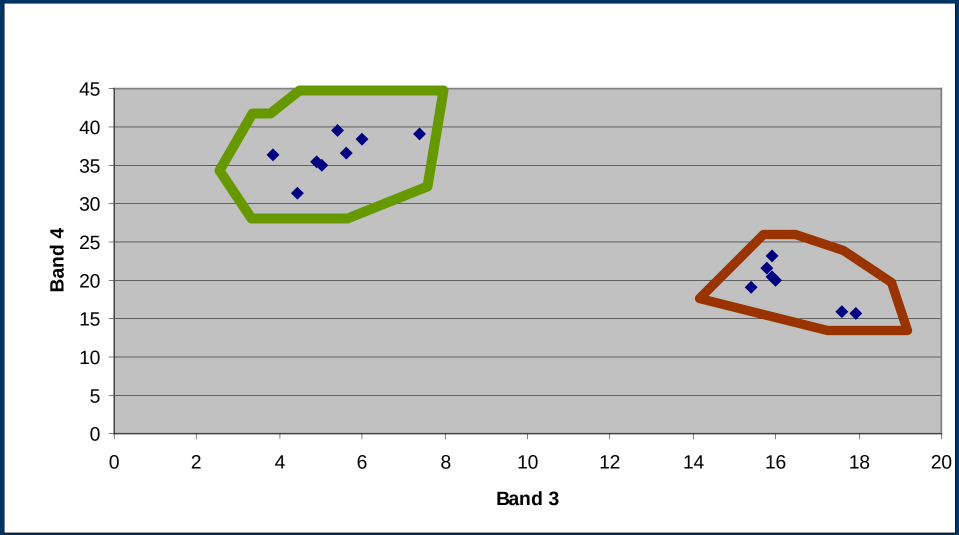

Classification Strategies

- Two basic strategies
  - Supervised classification
    - We impose our perceptions on the spectral data
  - Unsupervised classification
    - Spectral data imposes constraints on our interpretation

Supervised Classification

Supervised classification requires the analyst to select training areas where he/she knows what is on the ground and then digitize a polygon within that area…

The computer then creates…

Mean Spectral Signatures

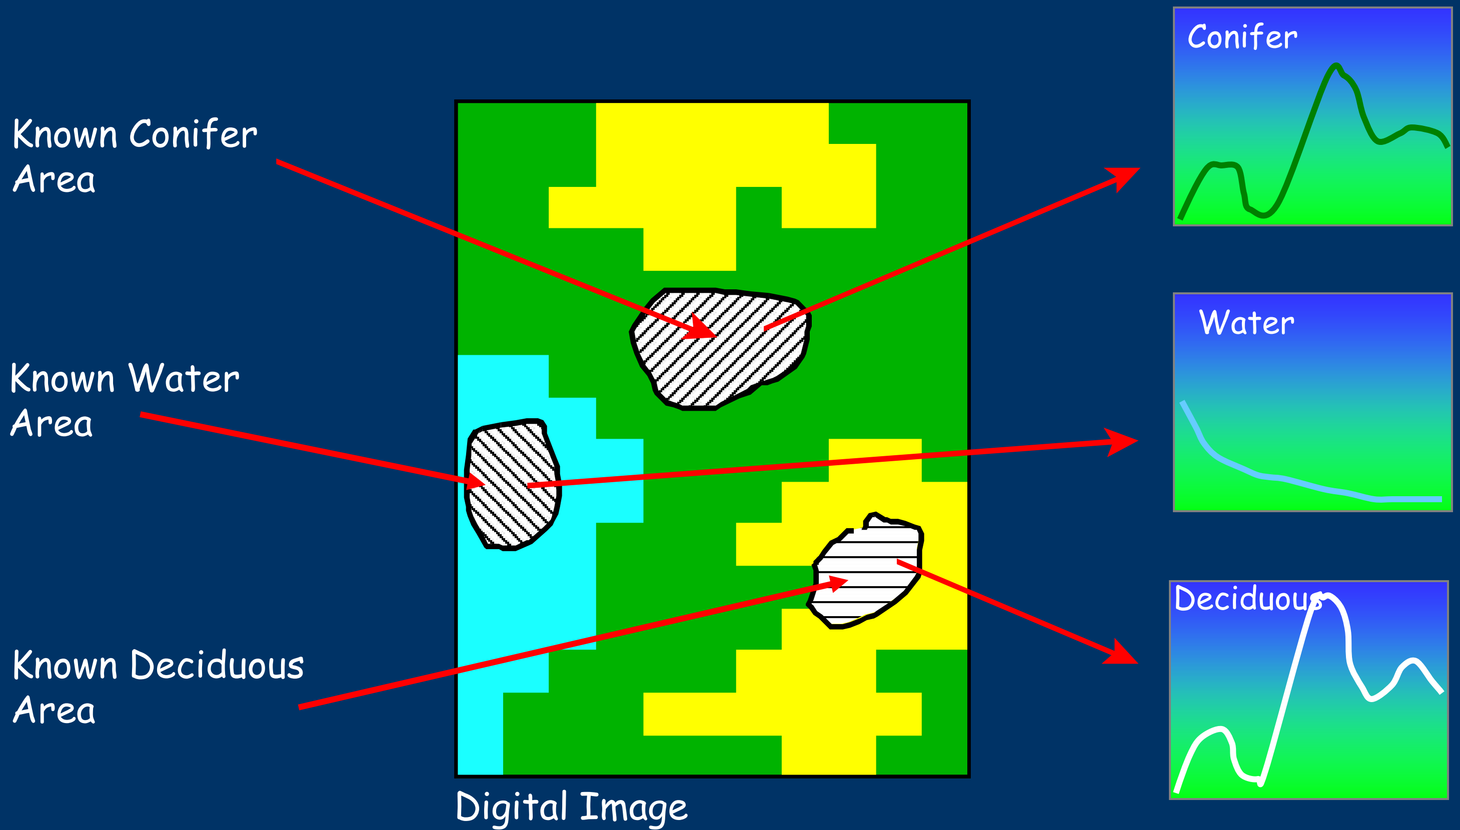

Supervised Classification

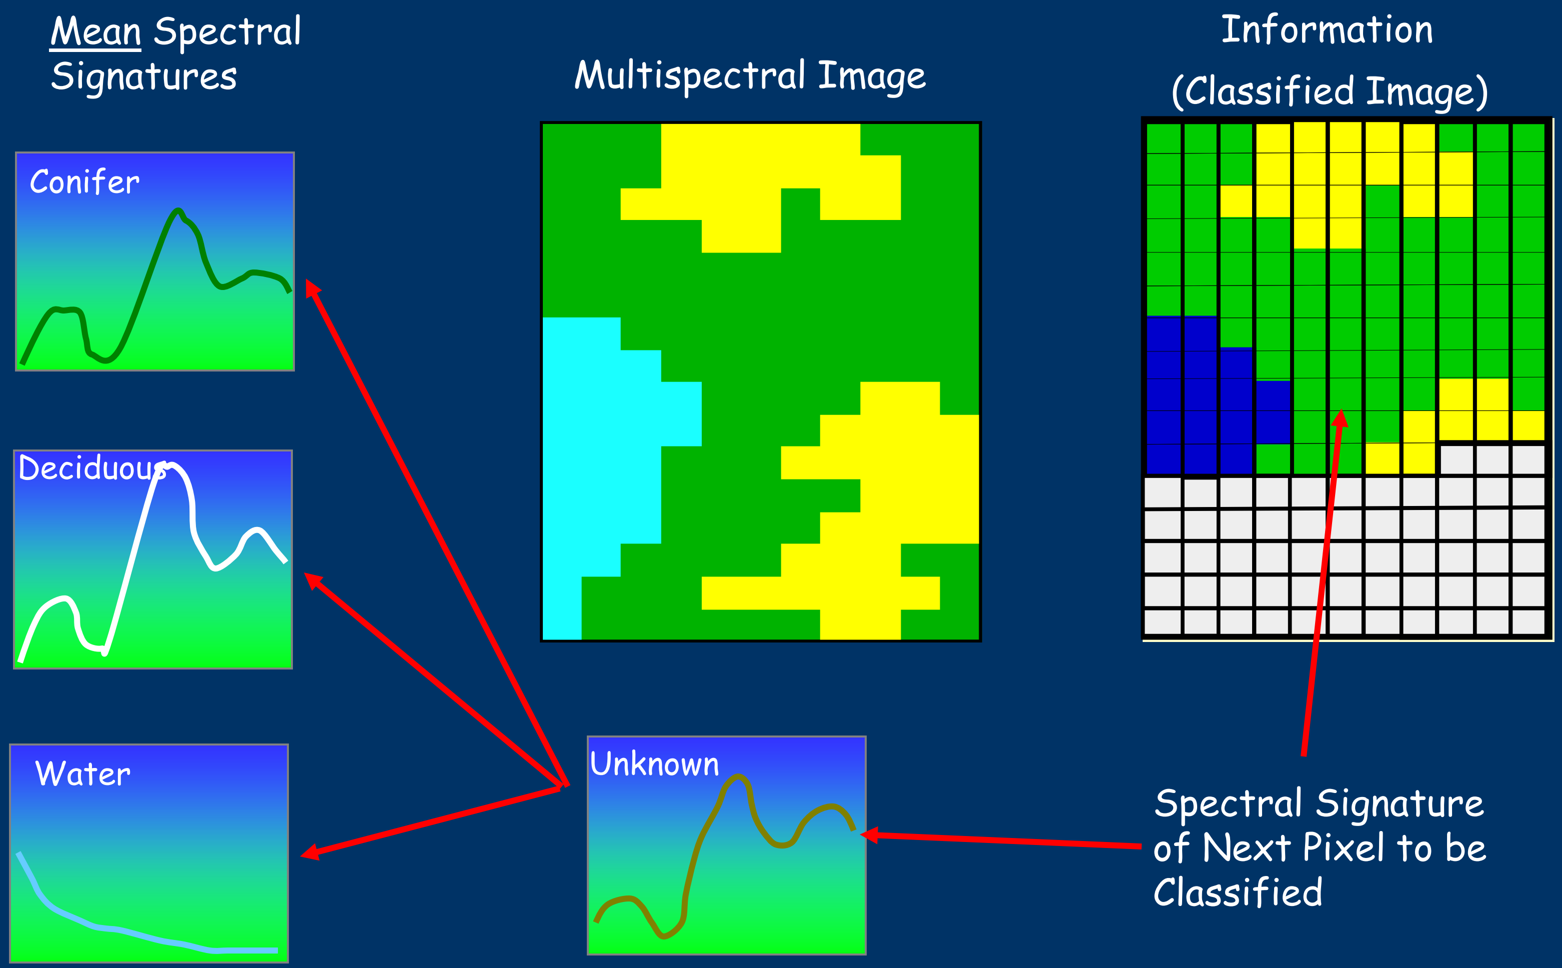

The Result is Information--in this case a Land Cover map...

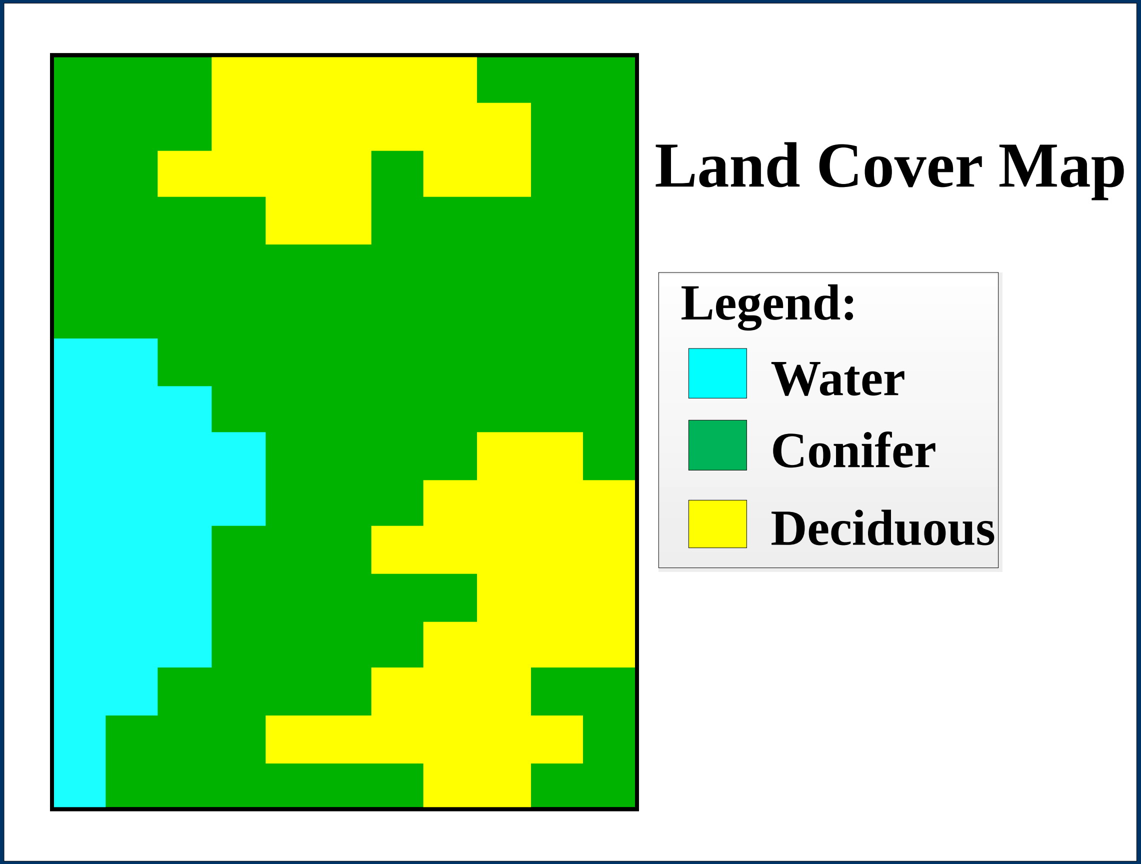

Supervised Classification

- Common Classifiers:
  - Parallelpiped
  - Minimum distance to mean
  - Maximum likelihood

Supervised Classification

- Parallelepiped Approach
- Pros:
  - Simple
  - Makes few assumptions about character of the classes

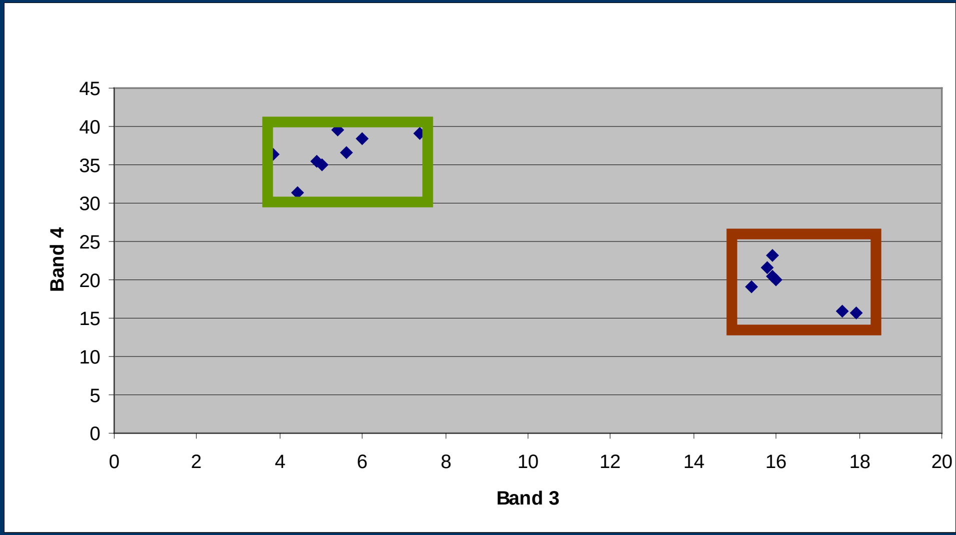

Supervised Classification

Cons: When we look at all the pixels in image, we find that they cover a continuous region in n-dimensional space: the parallelepiped approach may not be able to classify those regions

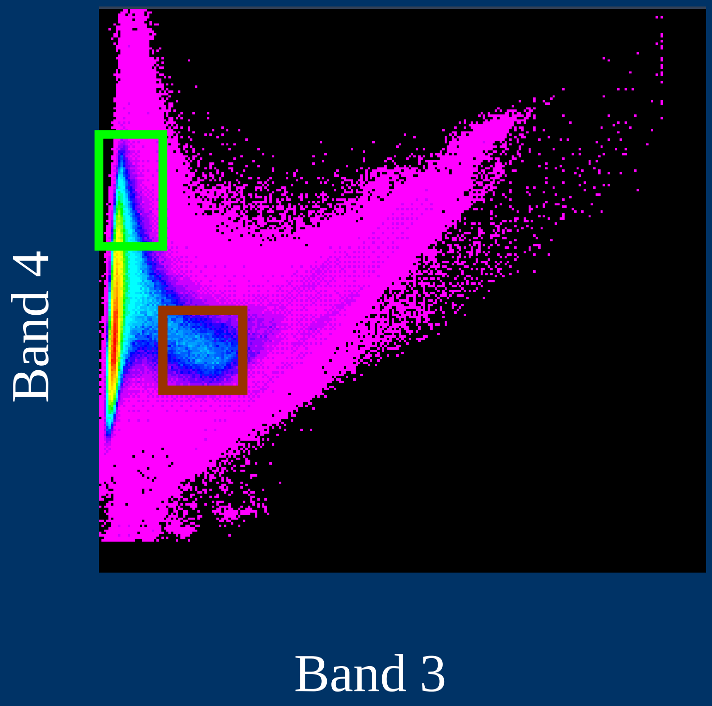

Supervised Classification

- Cons: Parallelepipeds are rectangular, but spectral space is “diagonal,” so classes may overlap

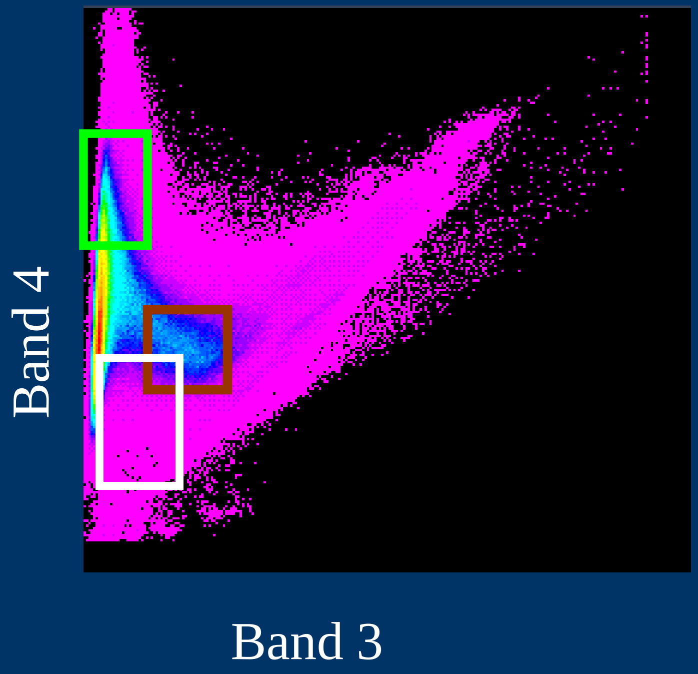

Supervised Classification: Statistical Approaches

- Minimum distance to mean
  - Find mean value of pixels of training sets in n-dimensional space
  - All pixels in image classified according to the class mean to which they are closest

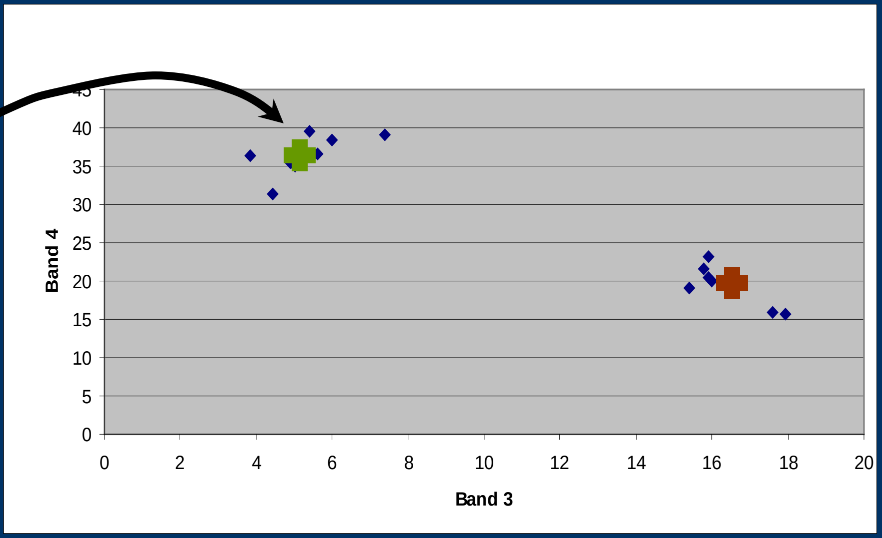

Supervised Classification: Minimum Distance

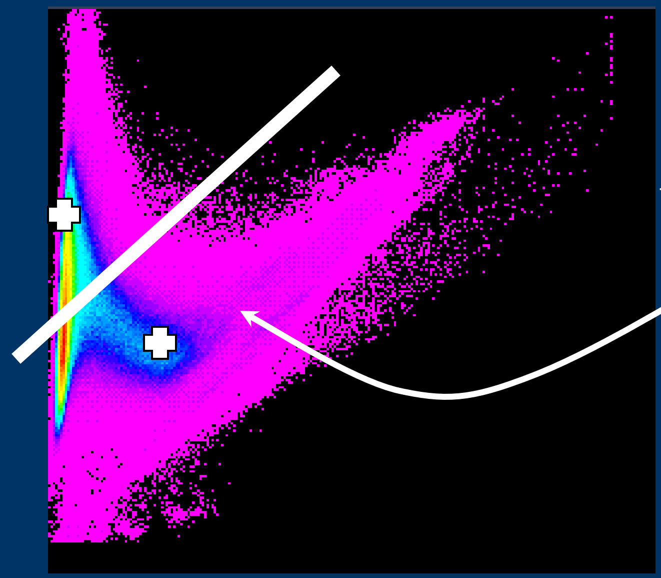

Supervised Classification: Minimum Distance

- Minimum distance
  - Pros:
    - All regions of n-dimensional space are classified
    - Allows for diagonal boundaries (and hence no overlap of classes)

Supervised Classification

- Minimum distance
  - Con:
    - Assumes that spectral variability is same in all directions, which is not the case

Band 4

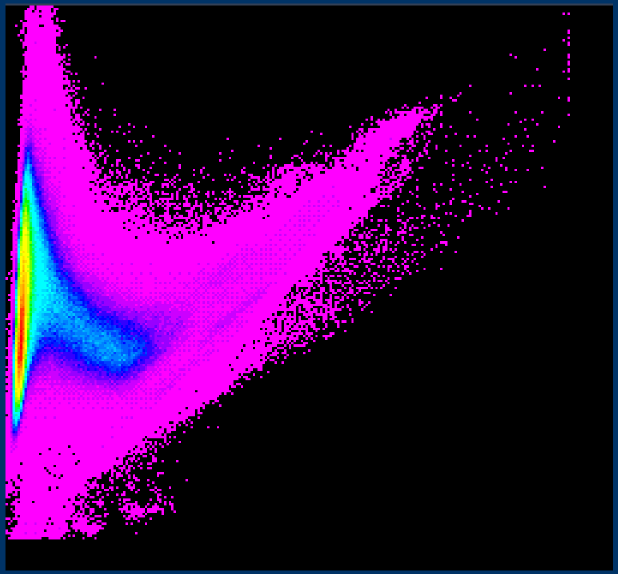

Band 3

For most pixels, Band 4 is much more variable than Band 3

Supervised Classification: Maximum Likelihood

- Maximum likelihood classification: another statistical approach
- Assume multivariate normal distributions of pixels within classes
- For each class, build a discriminant function
  - For each pixel in the image, this function calculates the probability that the pixel is a member of that class
  - Takes into account mean *and* covariance of training set
- Each pixel is assigned to the class for which it has the highest probability of membership

Maximum Likelihood Classifier

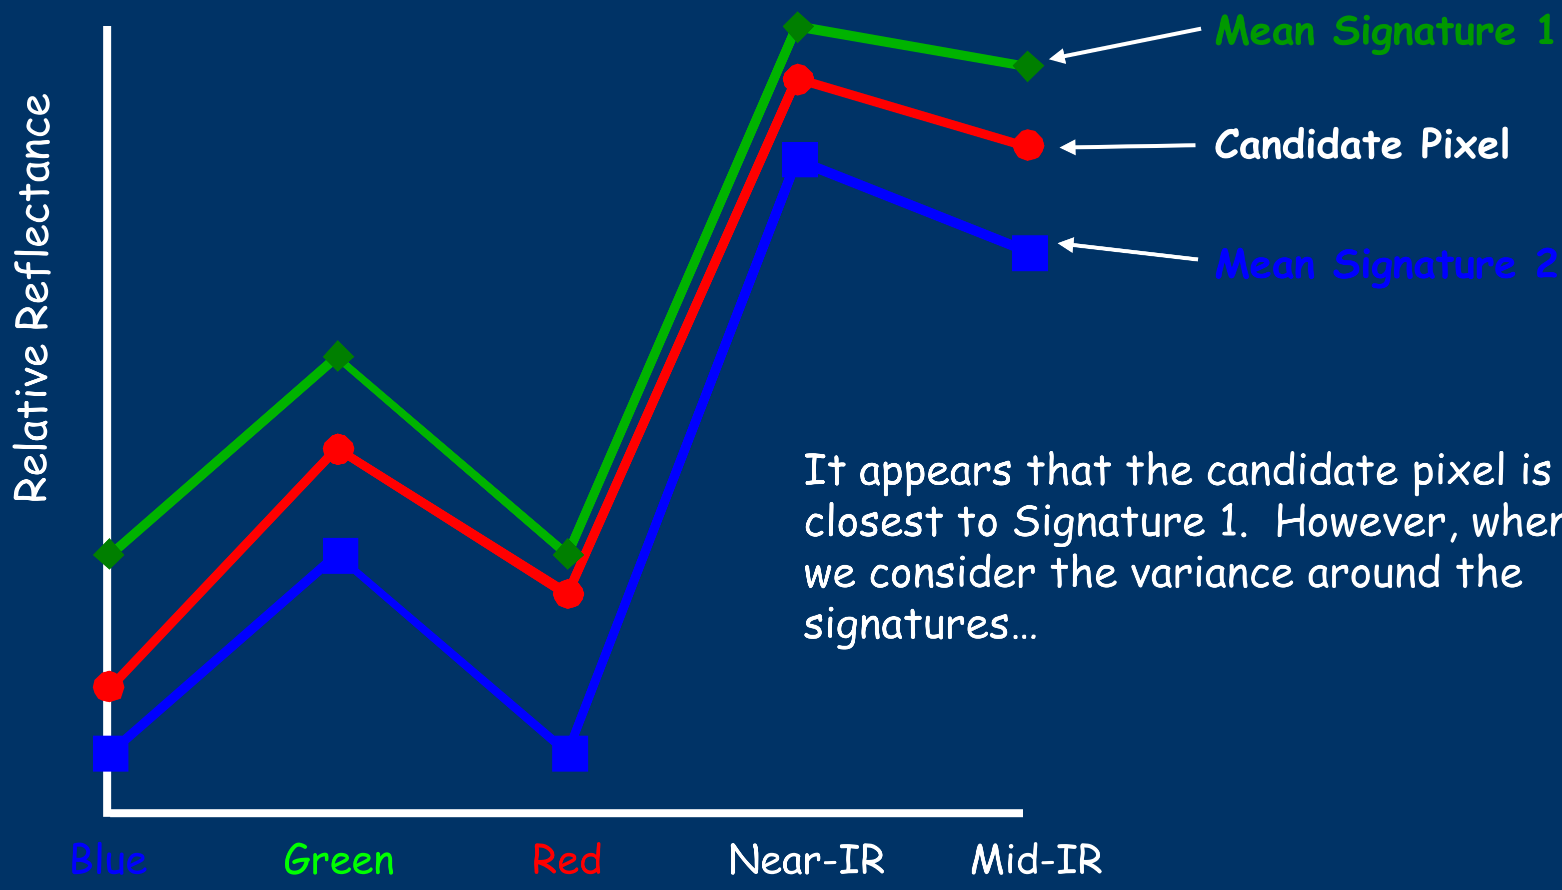

It appears that the candidate pixel is closest to Signature 1. However, when we consider the variance around the signatures…

Maximum Likelihood Classifier

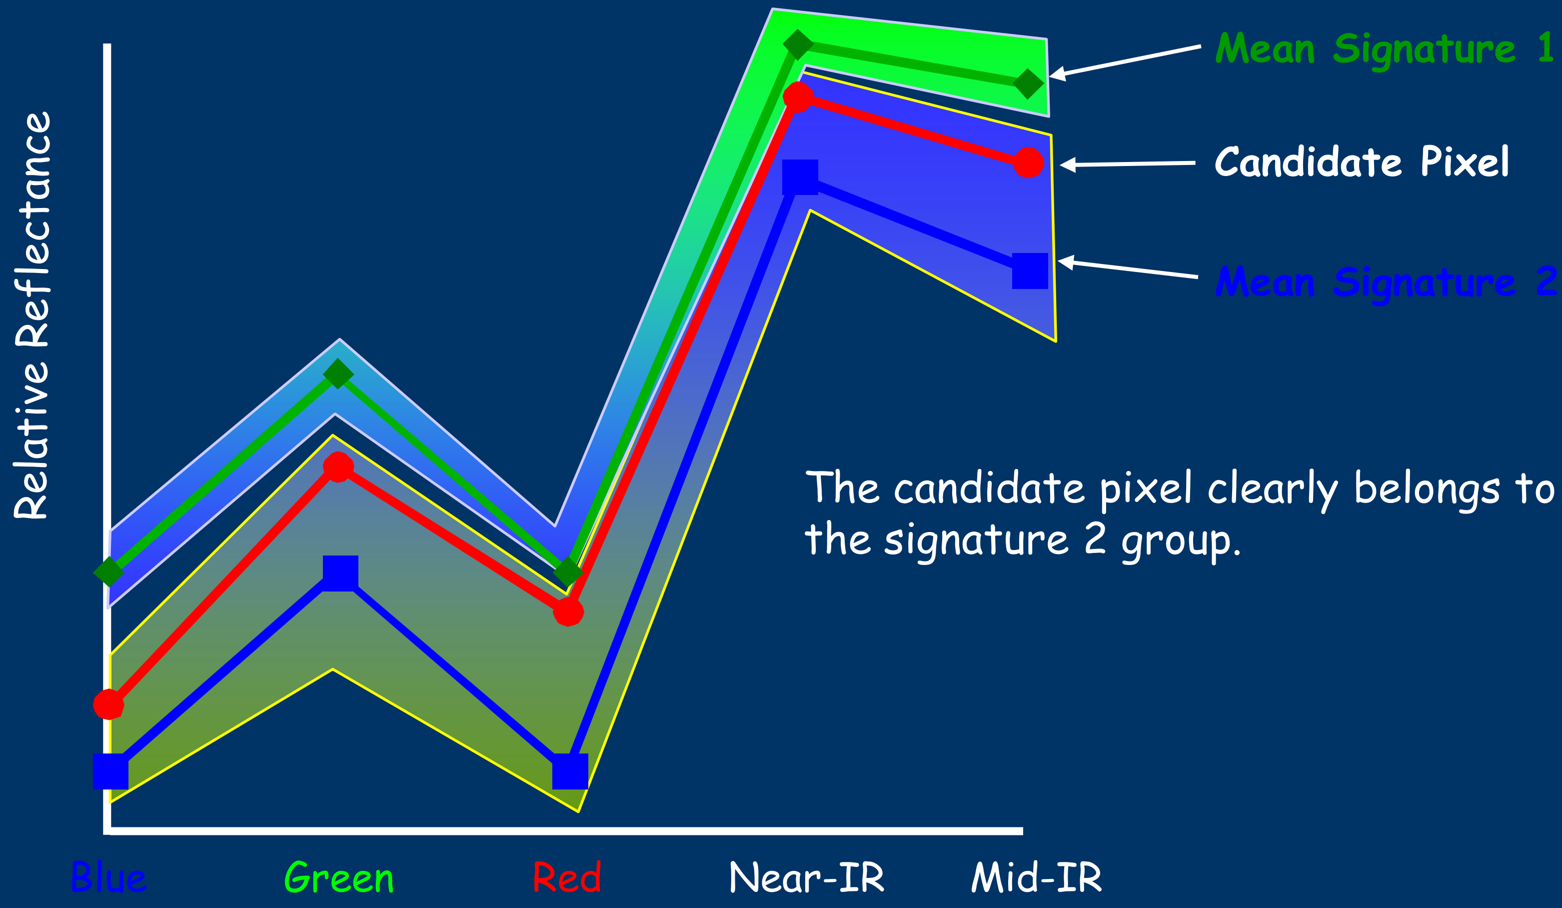

The candidate pixel clearly belongs to the signature 2 group.

Supervised Classification

- Maximum likelihood
  - Pro:
    - Most sophisticated; achieves good separation of classes
  - Con:
    - Requires strong training set to accurately describe mean and covariance structure of classes

Supervised Classification

- In addition to classified image, you can construct a “distance” image
  - For each pixel, calculate the distance between its position in n-dimensional space and the center of class in which it is placed
  - Regions poorly represented in the training dataset will likely be relatively far from class center points
    - May give an indication of how well your training set samples the landscape

Supervised Classification

- Some advanced techniques
  - Neural networks
    - Use flexible, not-necessarily-linear functions to partition spectral space
  - Contextual classifiers
    - Incorporate spatial or temporal conditions
  - Linear regression
    - Instead of discrete classes, apply proportional values of classes to each pixel; ie. 30% forest + 70% grass

Unsupervised Classification

- Recall: In unsupervised classification, the spectral data imposes constraints on our interpretation
- How? Rather than defining training sets and carving out pieces of n-dimensional space, we define *no* classes beforehand and instead use statistical approaches to divide the n-dimensional space into clusters with the *best separation*
- After the fact, we assign class names to those clusters

Unsupervised Classification

The analyst requests the computer to examine the image and extract a number of spectrally distinct clusters…

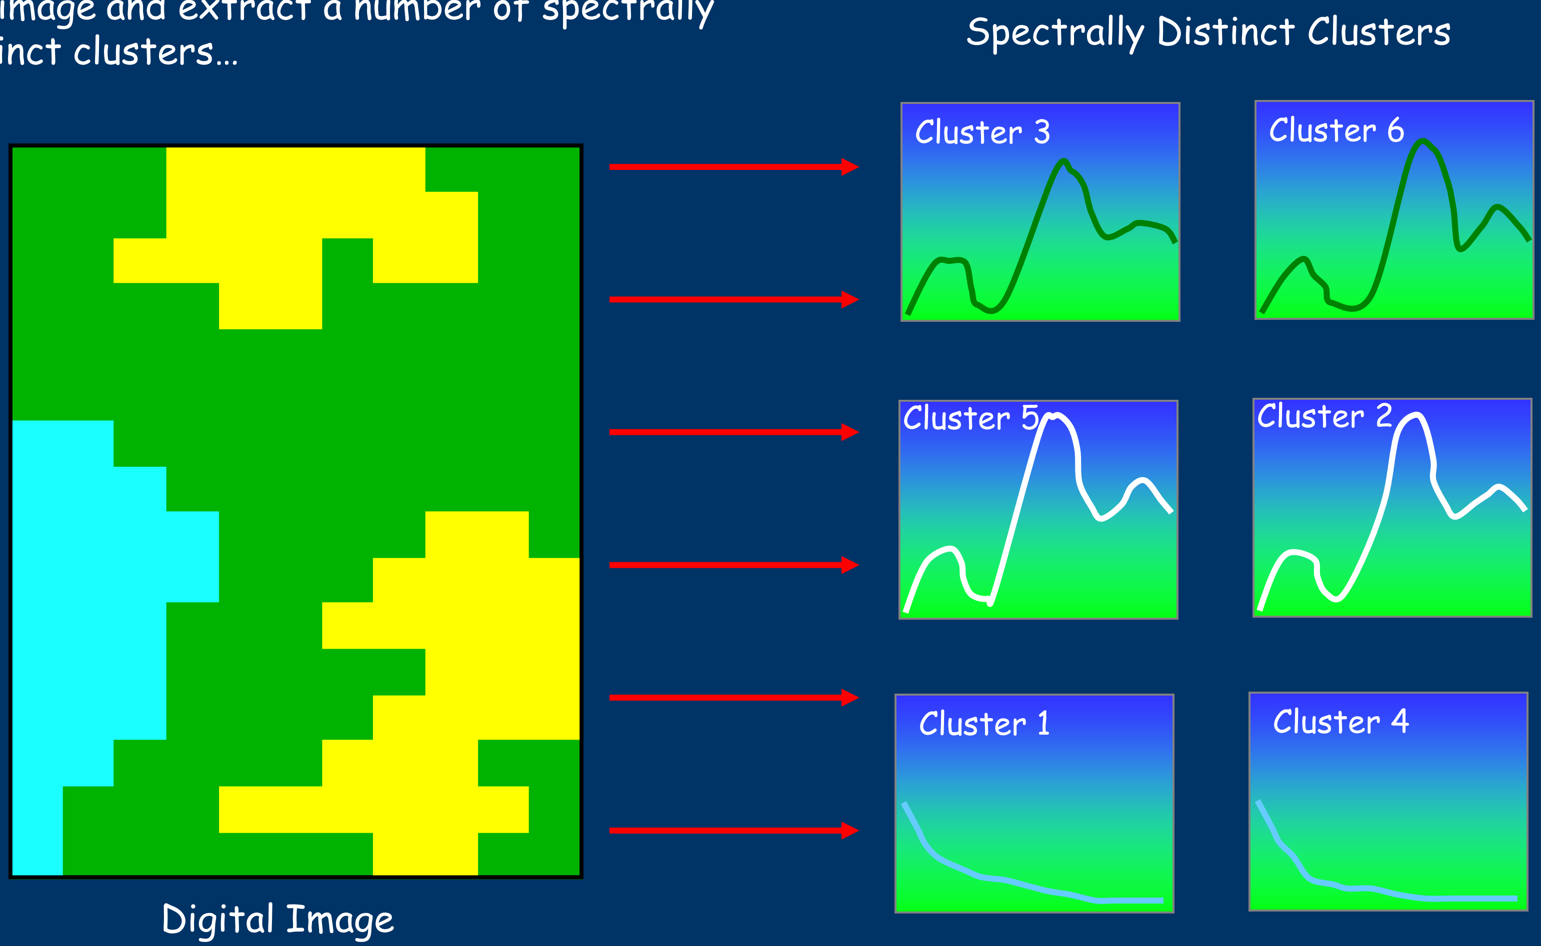

Unsupervised Classification

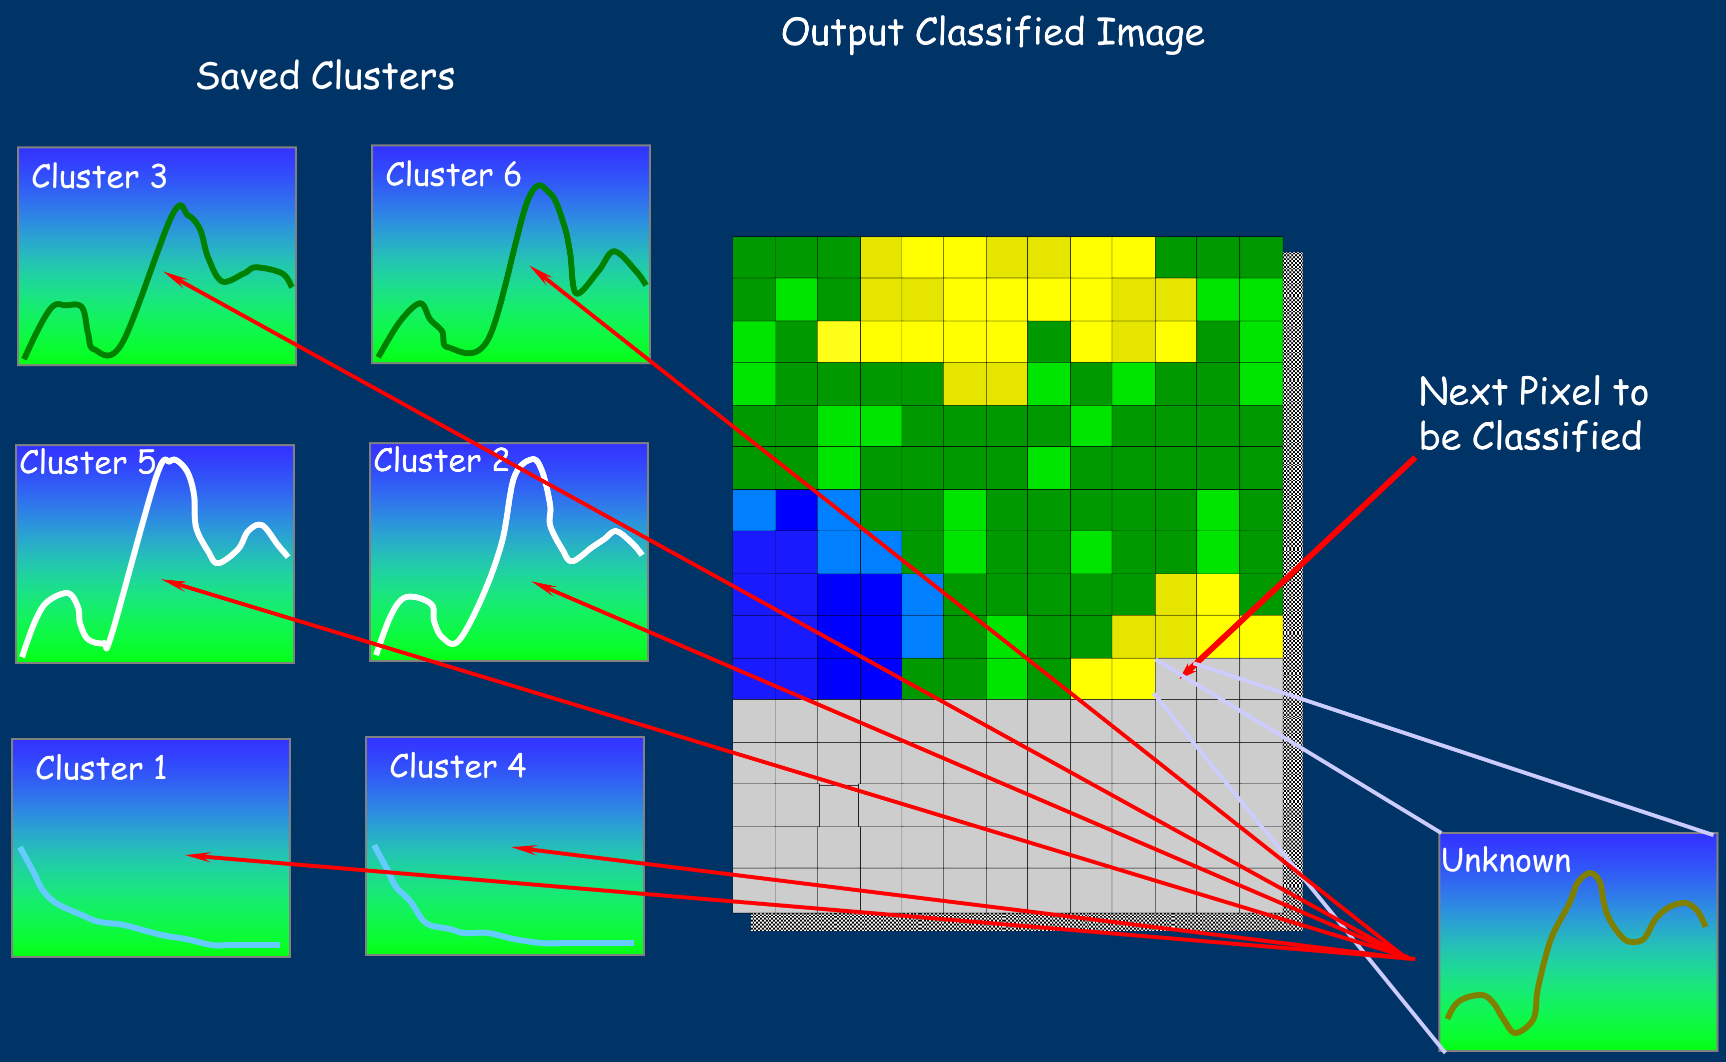

Unsupervised Classification

The result of the unsupervised classification is not yet information until…

The analyst determines the ground cover for each of the clusters…

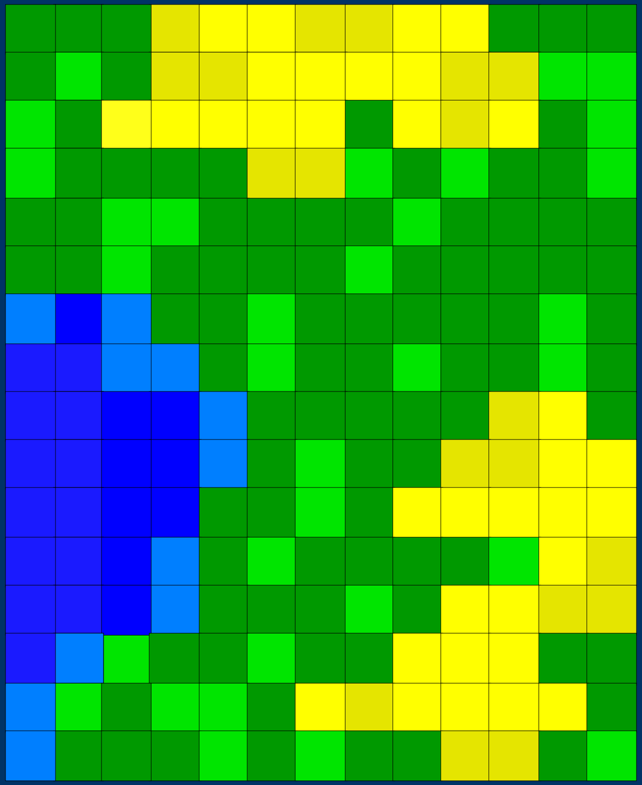

??? → Water

??? → Water

??? → Conifer

??? → Conifer

??? → Hardwood

??? → Hardwood

Unsupervised Classification

It is a simple process to regroup (recode) the clusters into meaningful information classes (the legend).

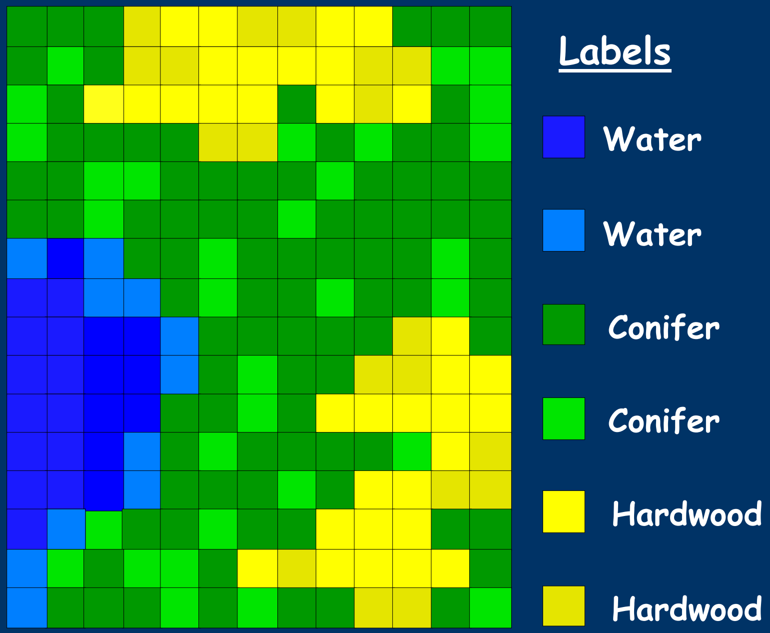

The result is essentially the same as that of the supervised classification:

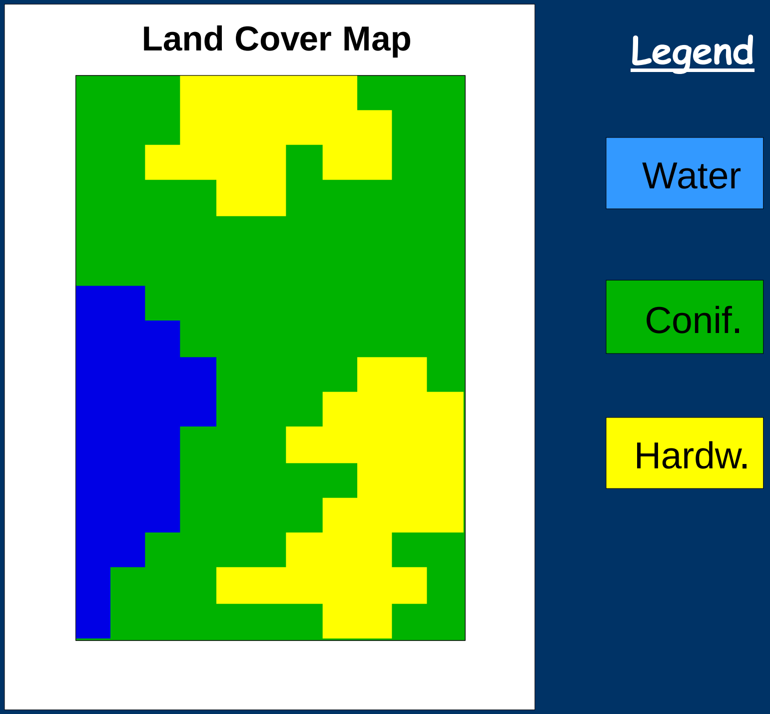

Unsupervised Classification

- Pros
  - Takes maximum advantage of spectral variability in an image
- Cons
  - The maximally-separable clusters in spectral space may not match our perception of the important classes on the landscape

ISODATA -- A Special Case of Minimum Distance Clustering

- “Iterative Self-Organizing Data Analysis Technique”
- Parameters you must enter include:
  - N - the maximum number of clusters that you want
  - T - a convergence threshold and
  - M - the maximum number of iterations to be performed.

ISODATA Procedure

- N arbitrary cluster means are established,
- The image is classified using a minimum distance classifier
- A new mean for each cluster is calculated
- The image is classified again using the new cluster means
- Another new mean for each cluster is calculated
- The image is classified again...

ISODATA Procedure

- After each iteration, the algorithm calculates the percentage of pixels that remained in the same cluster between iterations
- When this percentage exceeds T (convergence threshold), the program stops or…
- If the convergence threshold is never met, the program will continue for M iterations and then stop.

ISODATA Pros and Cons

- Not biased to the top pixels in the image (as sequential clustering can be)
- Non-parametric--data does not need to be normally distributed
- Very successful at finding the “true” clusters within the data if enough iterations are allowed
- Cluster signatures saved from ISODATA are easily incorporated and manipulated along with (supervised) spectral signatures
- Slowest (by far) of the clustering procedures.

Unsupervised Classification

- Critical issue: where to place initial *k* cluster centers

Along diagonal axis

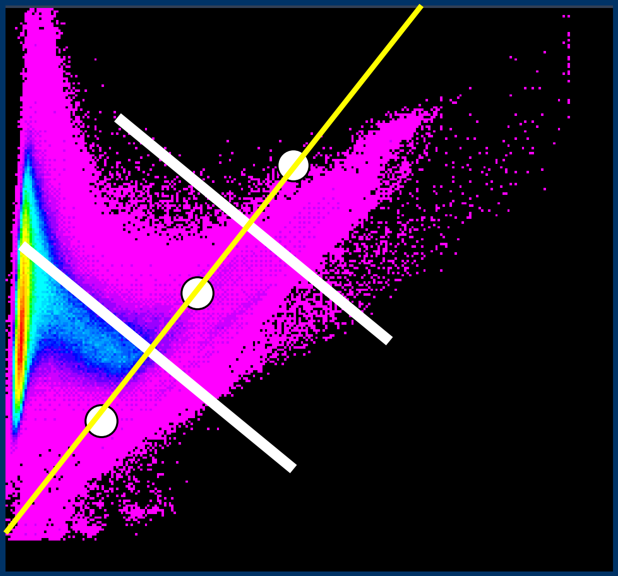

Along principal axis

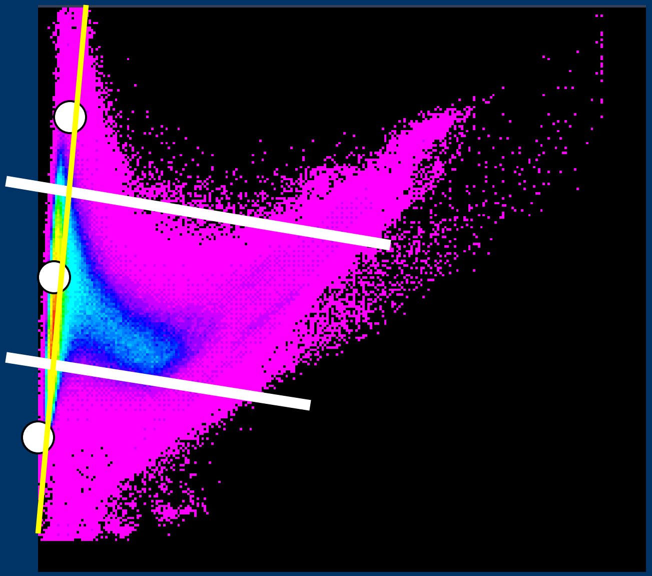

Unsupervised Classification

- Important issue: How to distribute cluster centers along axis

Distribute normally

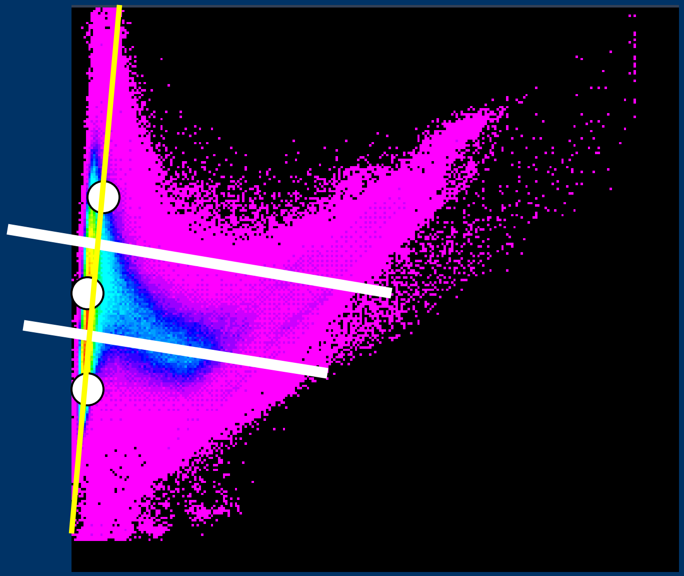

Distribute at tails of distribution

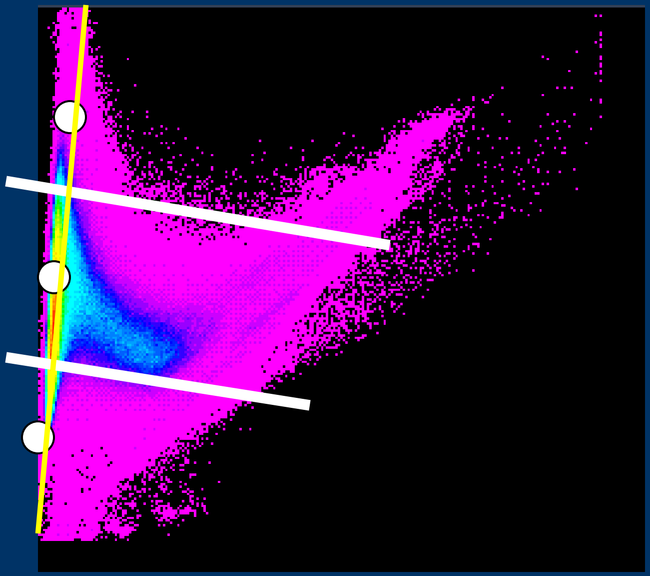

Unsupervised Classification

- After iterations finish, you're left with a map of distributions of pixels in the clusters
- How do you assign class names to clusters?
  - Requires some knowledge of the landscape
  - Ancillary data useful, if not critical (aerial photos, personal knowledge, etc.)
  - Covered in more depth in the Lab 4

Unsupervised Classification

- Alternatives to ISODATA approach
  - K-means algorithm
    - assumes that the number of clusters is known a priori, while ISODATA allows for different number of clusters
  - Non-iterative
    - Identify areas with “smooth” texture
    - Define cluster centers according to first occurrence in image of smooth areas
  - Agglomerative hierarchical
    - Group two pixels closest together in spectral space
    - Recalculate position as mean of those two; group
    - Group next two closest pixels/groups
    - Repeat until each pixel grouped

Classification: Summary

- Use spectral (radiometric) differences to distinguish objects
- Land cover not necessarily equivalent to land use
- Supervised classification
  - Training areas characterize spectral properties of classes
  - Assign other pixels to classes by matching with spectral properties of training sets
- Unsupervised classification
  - Maximize separability of clusters
  - Assign class names to clusters after classification

Spectral Clusters and Spectral Signatures

- Recall that clusters are spectrally distinct and signatures are informationally distinct
- When using the supervised procedure, the analyst must ensure that the informationally distinct signatures are spectrally distinct
- When using the unsupervised procedure, the analyst must supply the spectrally distinct clusters with information (label the clusters).

Spectrally Distinct Signatures

- Most image processing software have a set of programs which allow you to:
  - Graphically view the spectral signatures
  - Compute a distance matrix (measuring the spectral distance between all pairs of signature means)
  - Analyze statistics and histograms etc...
- After you analyze the signatures, the software should allow you to:
  - Modify merge or delete any signatures
  - Remember--they must be spectrally distinct!
- Finally, you can then classify the imagery (using a maximum likelihood classifier).

Evaluating Signatures--Signature Plots

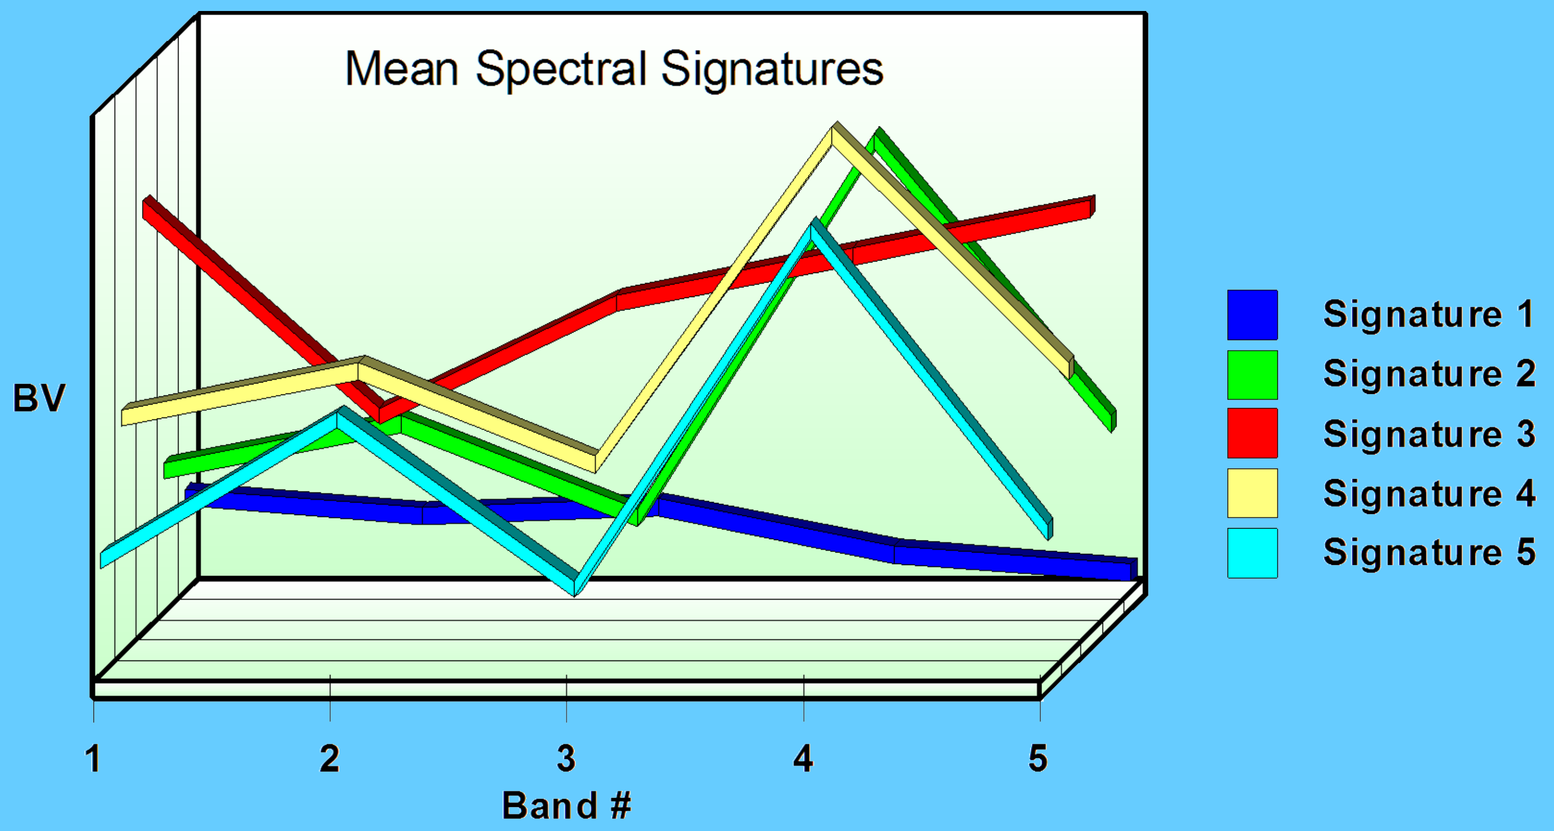

Ideally, each spectral signature must be separate from in at least one band from all other signatures. Note: signature variance is not shown here, but is a vital part of the signature...

Evaluating Signatures--Signature Ellipses

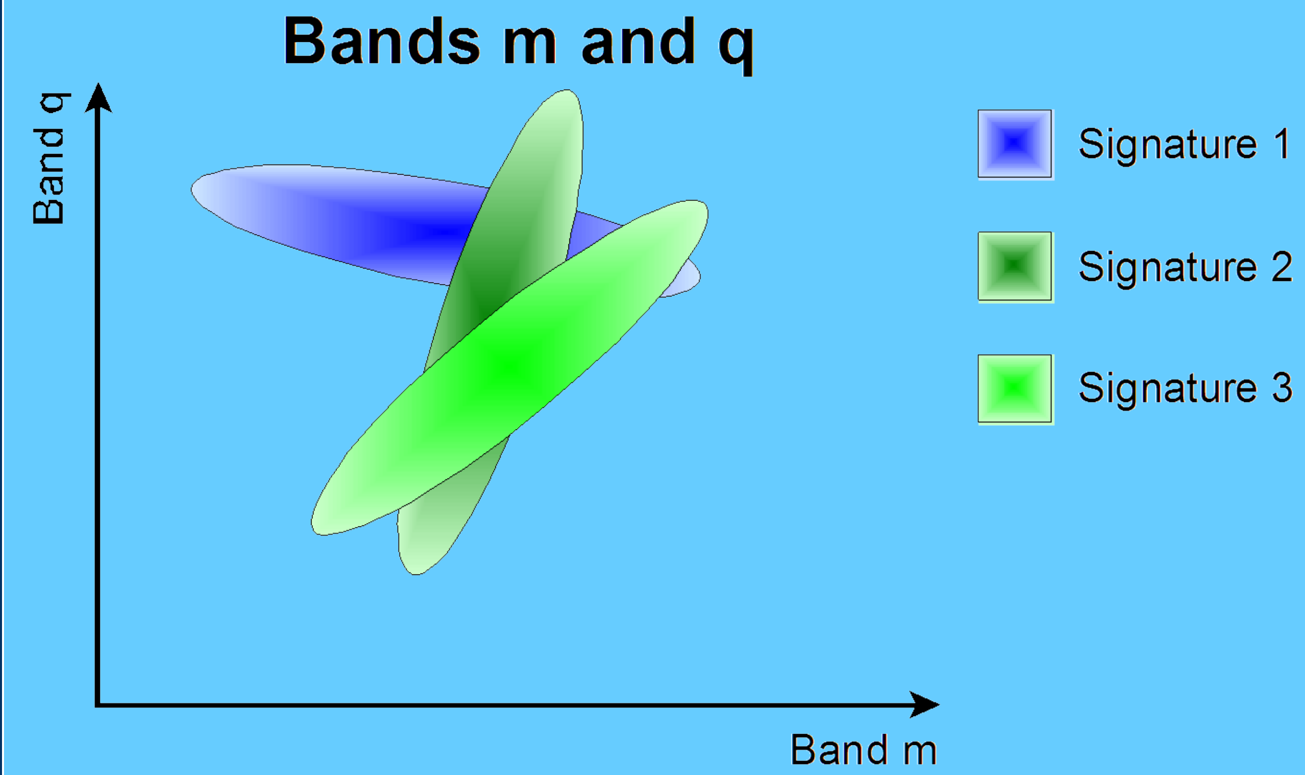

These three signature data ellipses overlap in bands m and q; however,...

Evaluating Signatures--Signature Ellipses

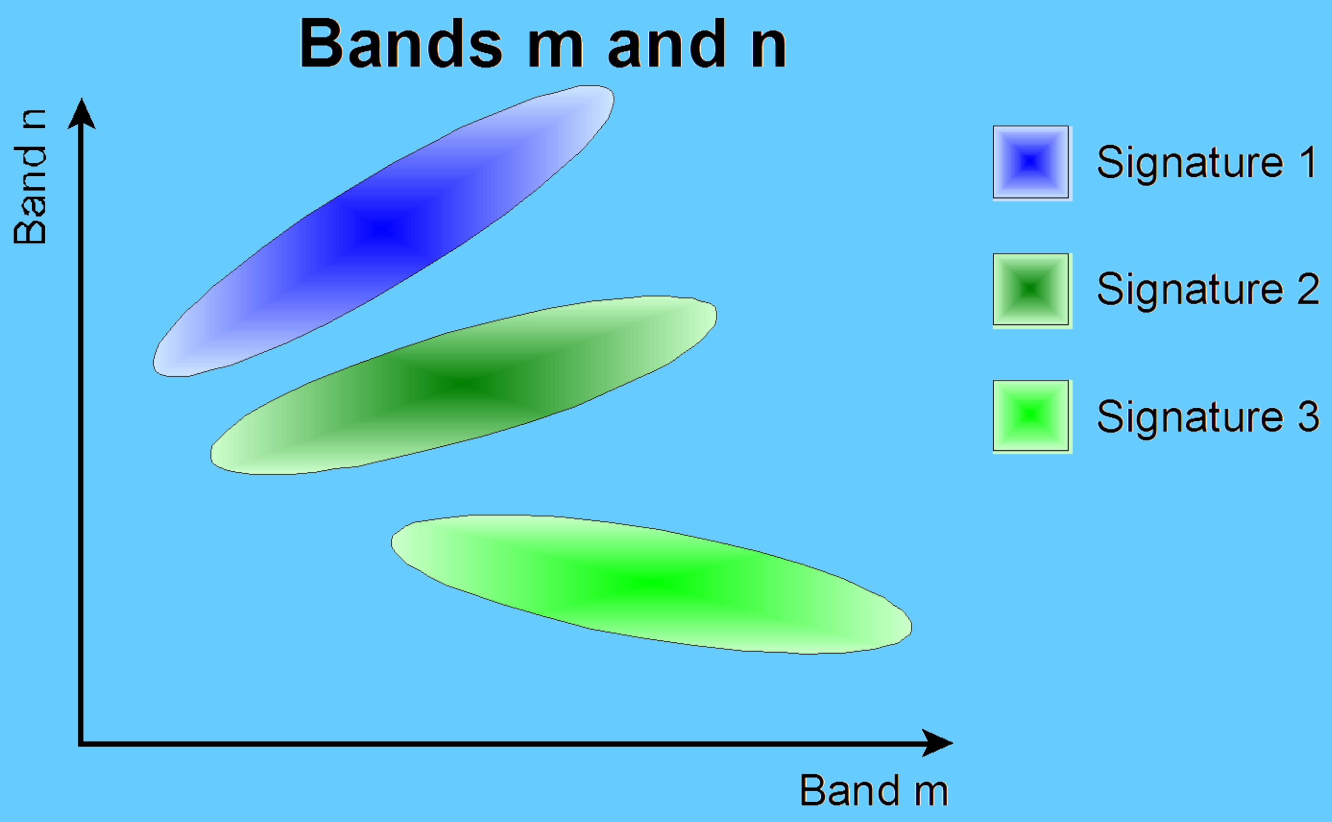

They don’t overlap in bands m and n.