

The blue and green colors are actually the same



<http://blogs.discovermagazine.com/badastronomy/2009/06/24/the-blue-and-the-green/>

Machine Learning: Overview



Computer Vision
James Hays, Brown

Slides: Isabelle Guyon,
Erik Sudderth,
Mark Johnson,
Derek Hoiem

Machine learning: Overview

- Core of ML: Making predictions or decisions from Data.
- This overview will not go in to depth about the statistical underpinnings of learning methods. We're looking at ML as a tool. Take *CSCI 1950-F: Introduction to Machine Learning* to learn more about ML.

Impact of Machine Learning

- Machine Learning is arguably the greatest export from computing to other scientific fields.

Machine Learning Applications

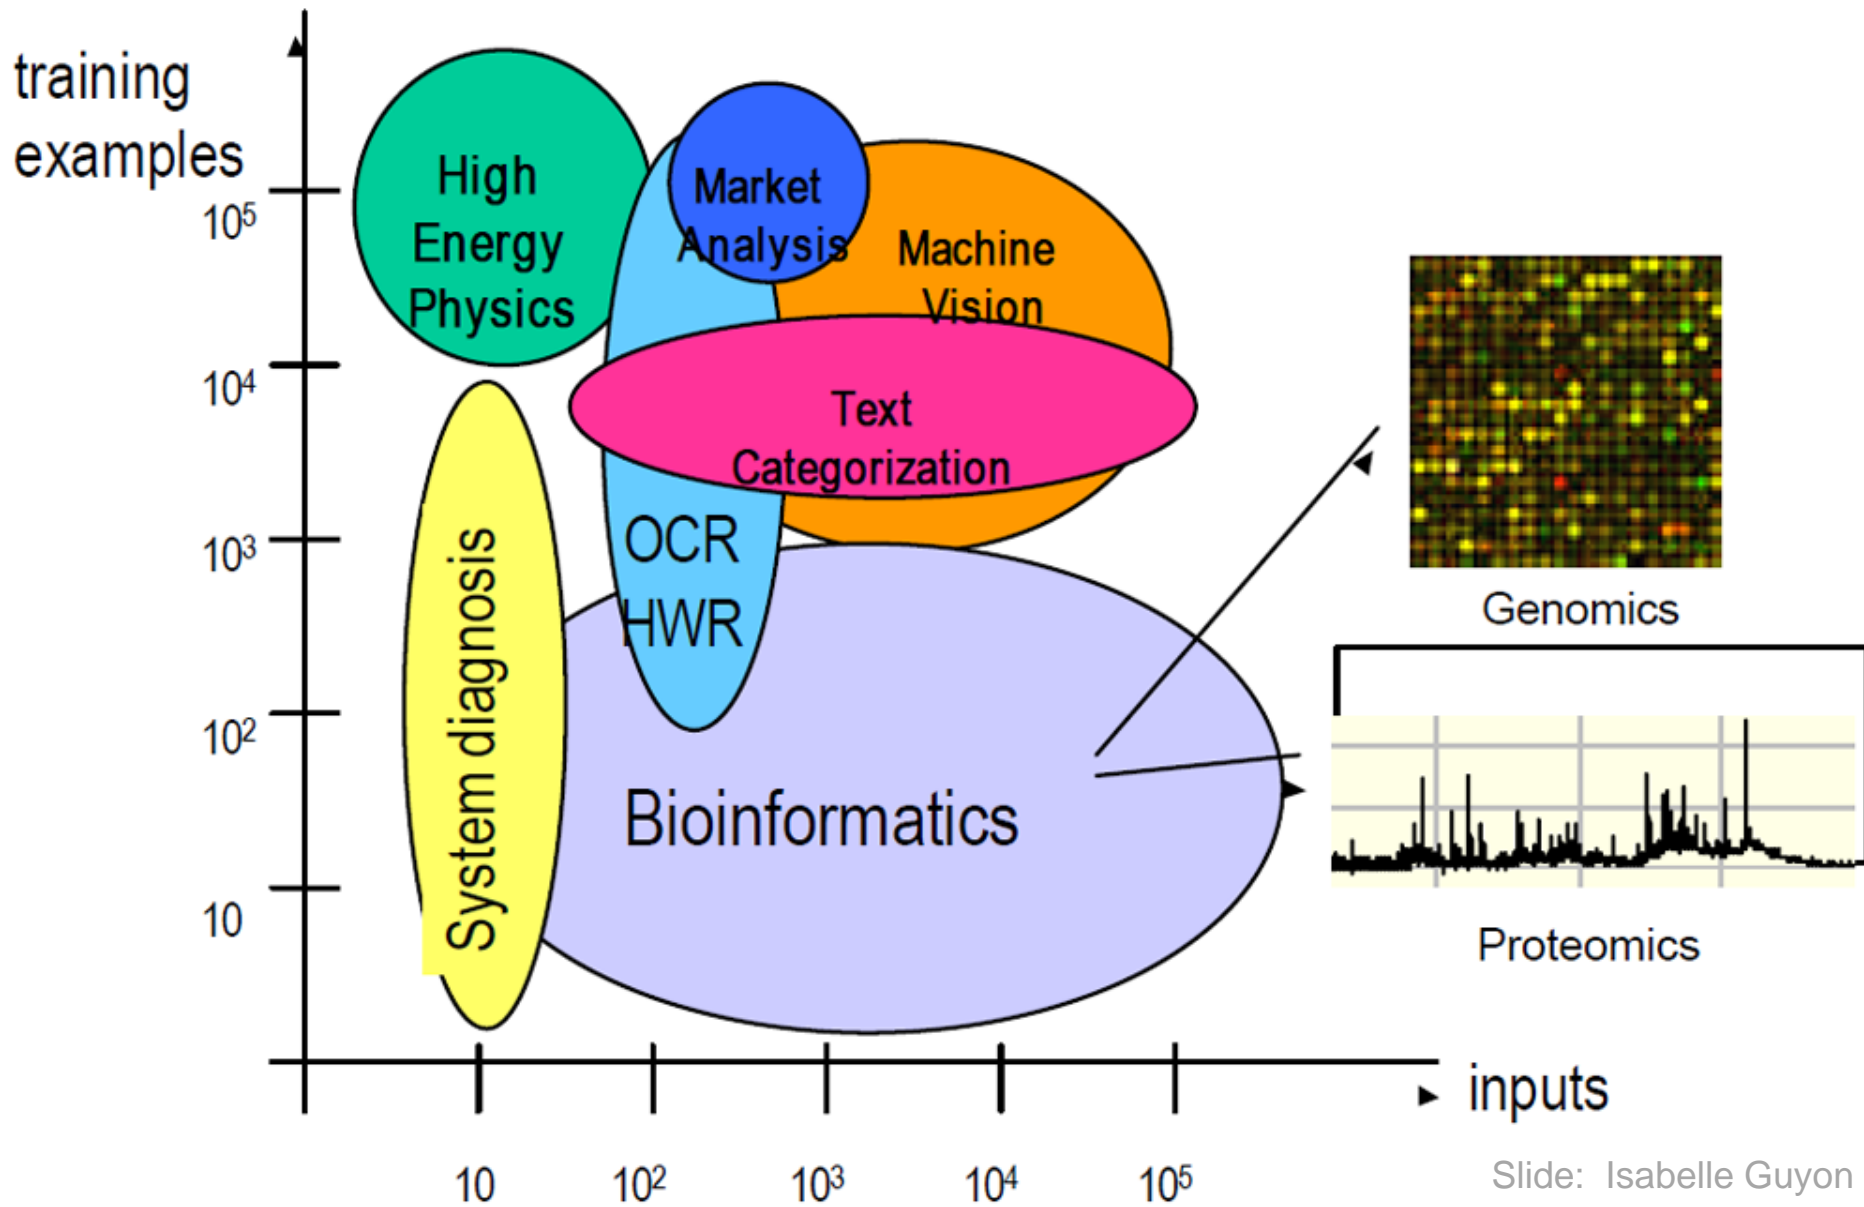


Image Categorization

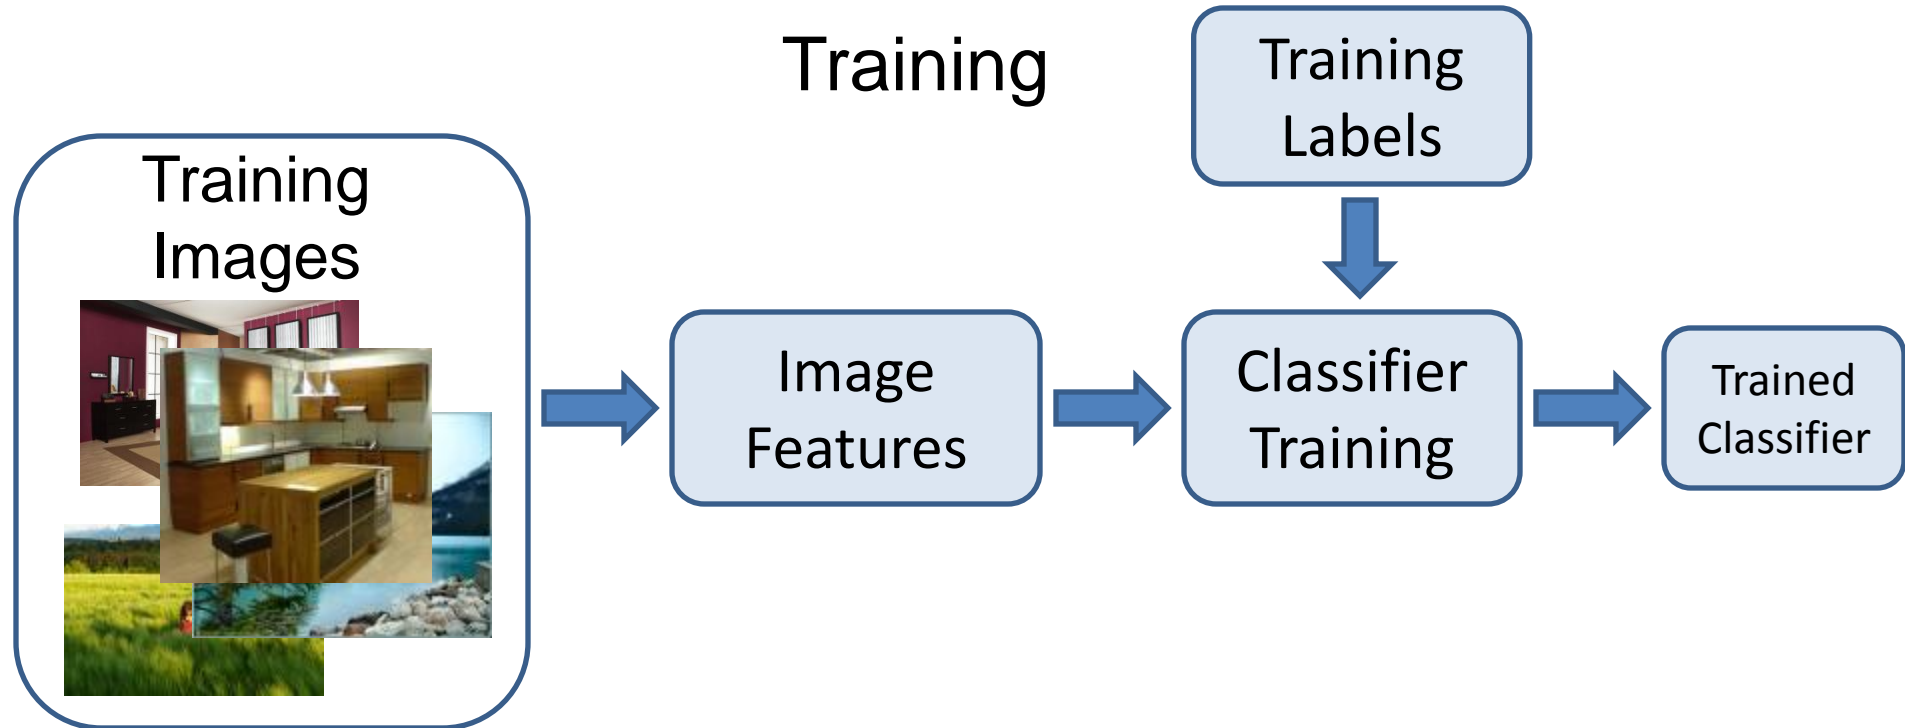
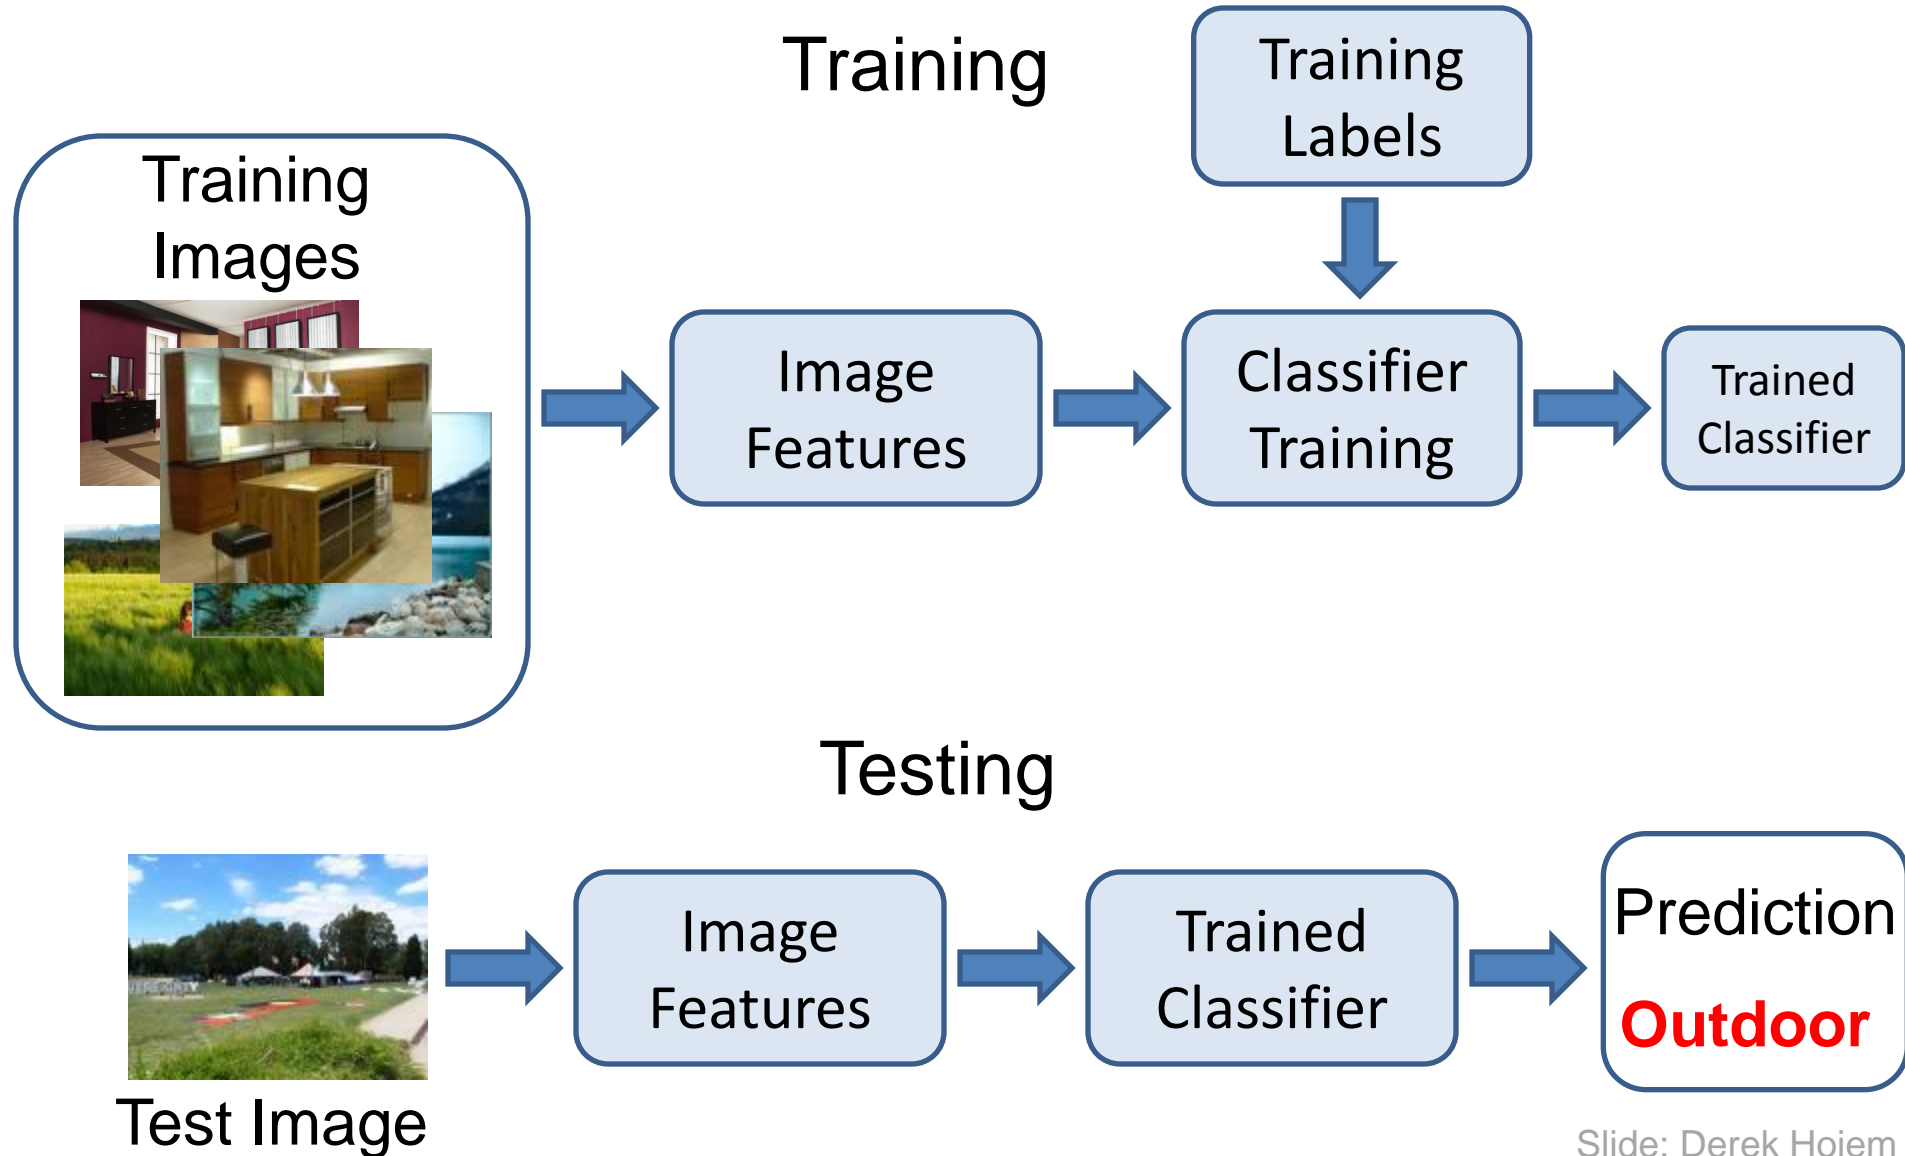


Image Categorization



Claim:

The decision to *use* machine learning is more important than the choice of a *particular* learning method.

If you hear somebody talking of a specific learning mechanism, be wary (e.g. YouTube comment "Oooh, we could plug this in to a Neural network and blah blah blah")

Example: Boundary Detection

- Is this a boundary?


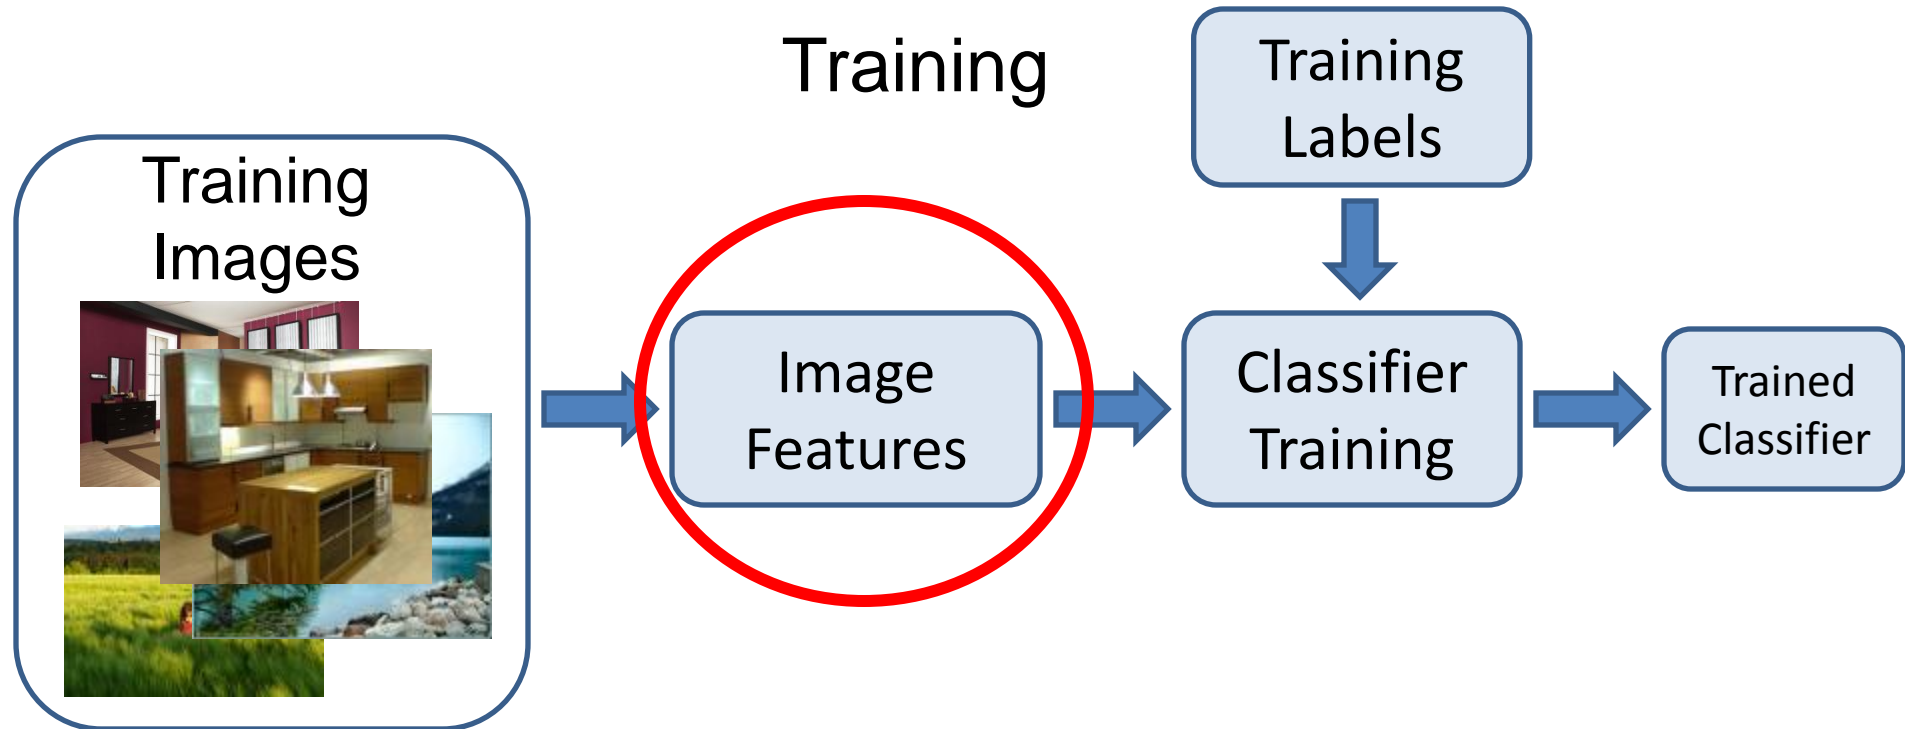


Image features



General Principles of Representation

- Coverage
 - Ensure that all relevant info is captured
- Concision
 - Minimize number of features without sacrificing coverage
- Directness
 - Ideal features are independently useful for prediction





Image Intensity

Image representations

- Templates

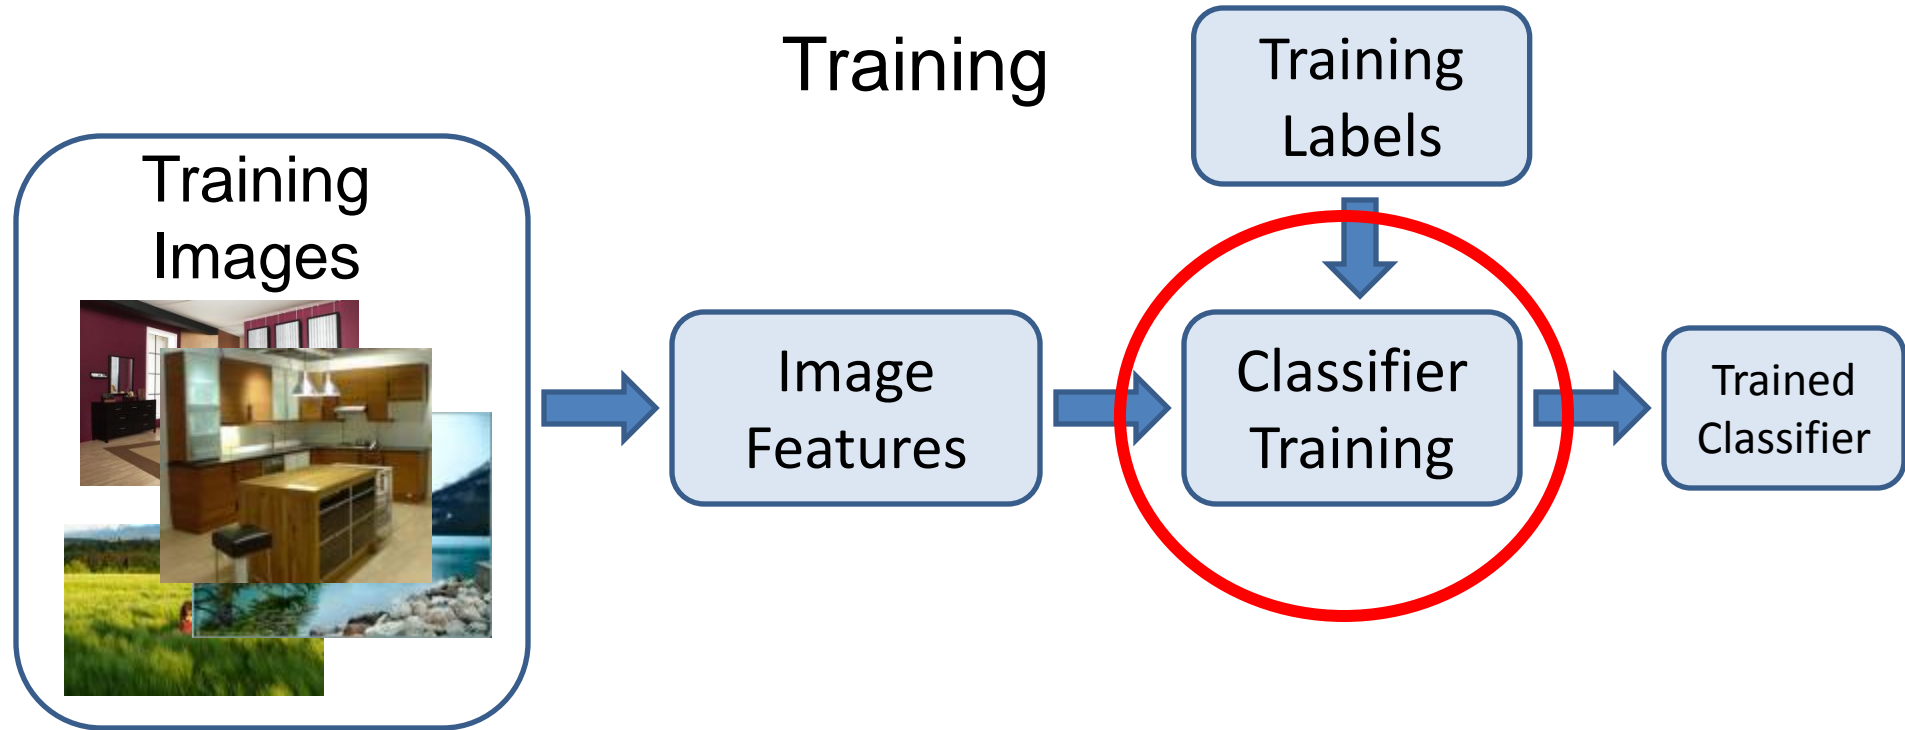
- Intensity, gradients, etc.



- Histograms


- Color, texture, SIFT descriptors, etc.

Classifiers



Learning a classifier

Given some set of features with corresponding labels, learn a function to predict the labels from the features



One way to think about it...

- Training labels dictate that two examples are the same or different, in some sense
- Features and distance measures define visual similarity
- Classifiers try to learn weights or parameters for features and distance measures so that visual similarity predicts label similarity

Machine Learning Problems

Supervised Learning

Unsupervised Learning

Discrete
Continuous

classification or
categorization


clustering

regression

dimensionality
reduction

Dimensionality Reduction

- **PCA, ICA, LLE, Isomap**
- PCA is the most important technique to know. It takes advantage of correlations in data dimensions to produce the best possible lower dimensional representation, according to reconstruction error.
- PCA should be used for dimensionality reduction, not for discovering patterns or making predictions. Don't try to assign semantic meaning to the bases.



Many classifiers to choose from

- SVM
- Neural networks
- Naïve Bayes
- Bayesian network
- Logistic regression
- Randomized Forests
- Boosted Decision Trees
- K-nearest neighbor
- RBMs
- Etc.

Which is the best one?

Next Two Lectures:

- Friday we'll talk about clustering methods (k-means, mean shift) and their common usage in computer vision -- building "bag of words" representations inspired by the NLP community. We'll be using these models for projects 2 and 3.
- Monday we'll focus specifically on classification methods, e.g. nearest neighbor, naïve-Bayes, decision trees, linear SVM, Kernel methods. We'll be using these for projects 3 and 4 (and optionally 2).