CSSE463: Image Recognition Day 31

- Today: Bayesian classifiers
- Tomorrow: project meetings.

- Questions?

Bayesian classifiers

- Use training data
  - Assume that you know probabilities of each feature.
- If 2 classes:
  - Classes ω_1 and ω_2
  - Say, circles vs. non-circles
  - A single feature, x
  - Both classes equally likely
  - Both types of errors equally bad
- Where should we set the threshold between classes? Here?
- Where in graph are 2 types of errors?

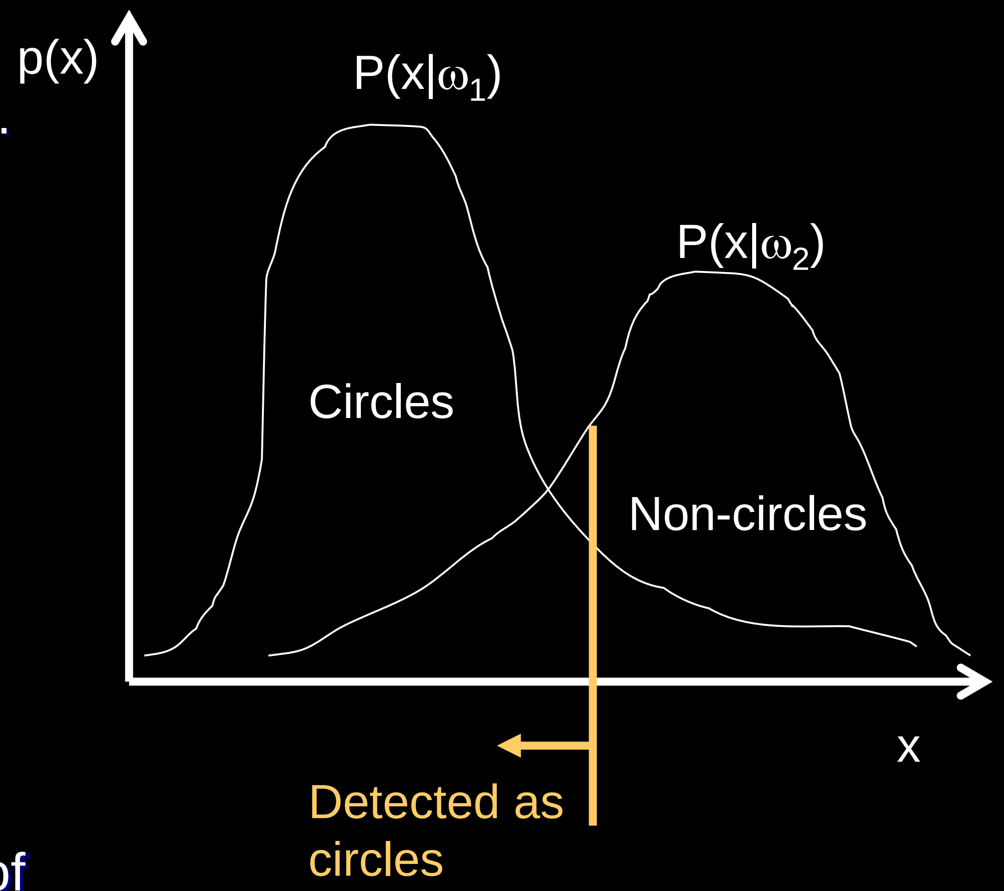

What if we have prior information?

- Bayesian probabilities say that if we only expect 10% of the objects to be circles, that should affect our classification

Bayesian classifier in general

- Bayes rule:
  - Verify with example
- For classifiers:
  - x = feature(s)
  - ω_i = class
  - $P(\omega|x)$ = posterior probability
  - $P(\omega)$ = prior
  - $P(x)$ = unconditional probability
  - Find best class by *maximum a posteriori (MAP)* priniciple. *Find class i that maximizes $P(\omega_i|x)$.*
    - Denominator doesn't affect calculations
  - Example:
    - indoor/outdoor classification

$$p(a \mid b) = \frac{p(b \mid a)\, p(a)}{p(b)}$$

$$p(\omega_i \mid x) = \frac{p(x \mid \omega_i)\, p(\omega_i)}{p(x)}$$

Fixed

Learned from examples (histogram)

Learned from training set (or leave out if unknown)

Bayes rule is used in prediction of disease

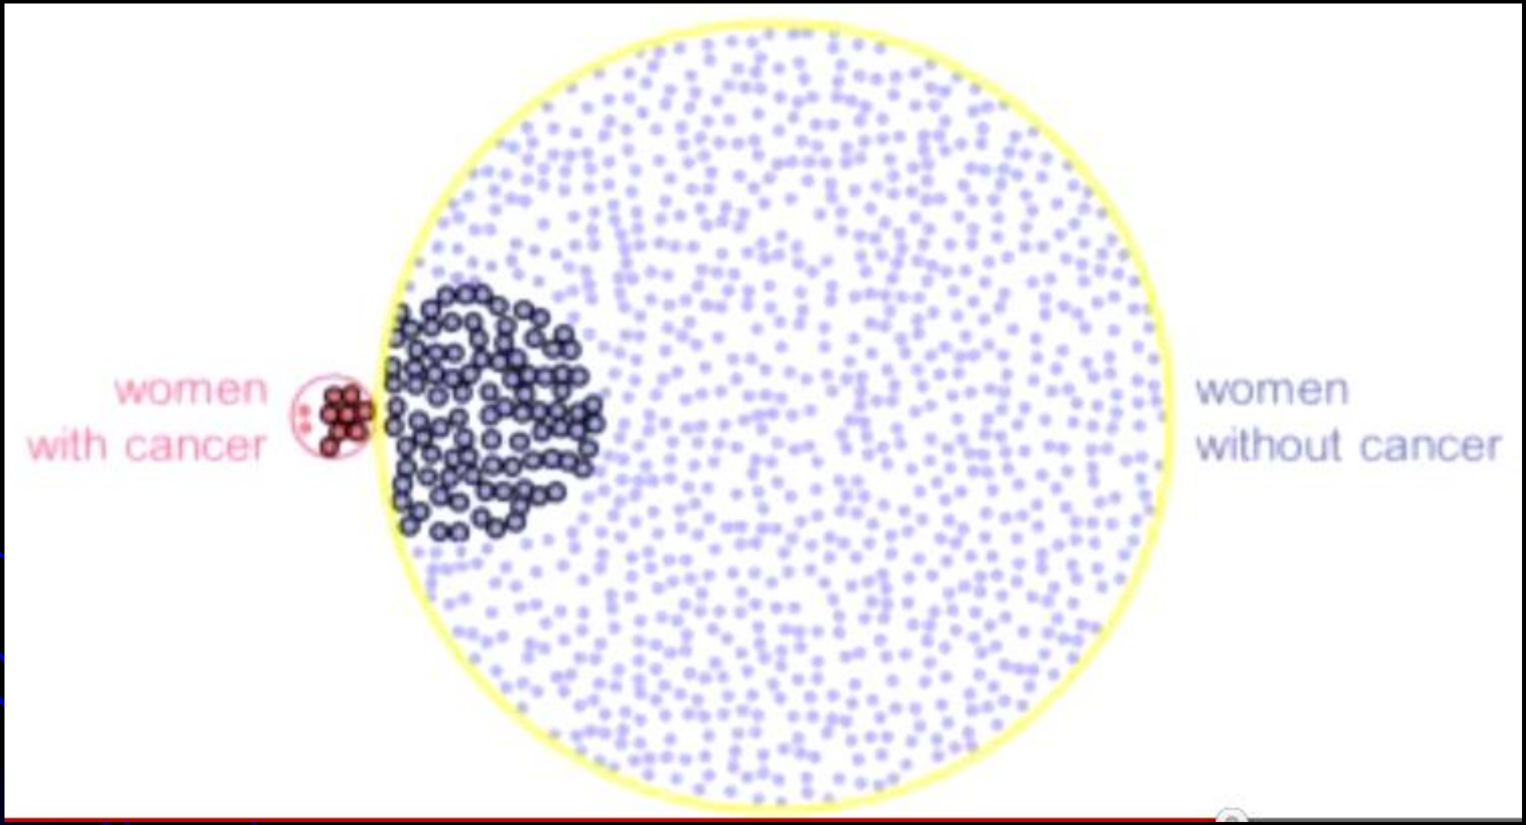

http://www.youtube.com/watch?v=D8VZqxcu0I0

Can you verify the approximation they found?

Indoor vs. outdoor classification

- I can use low-level image info (color, texture, etc)
- But there’s another source of really helpful info!

Camera Metadata Distributions

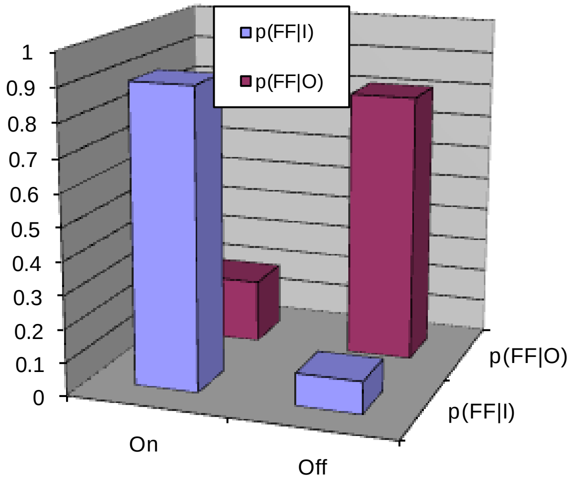

Flash

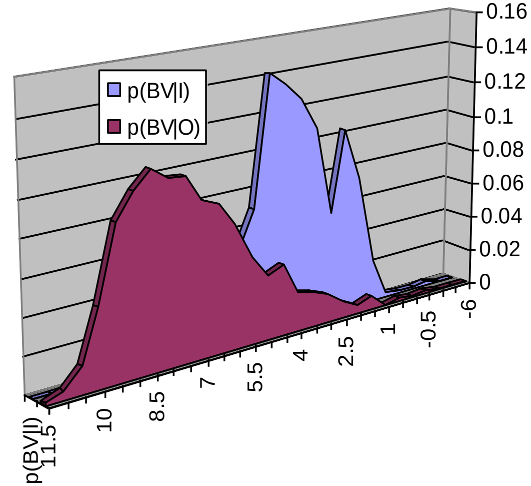

Scene Brightness

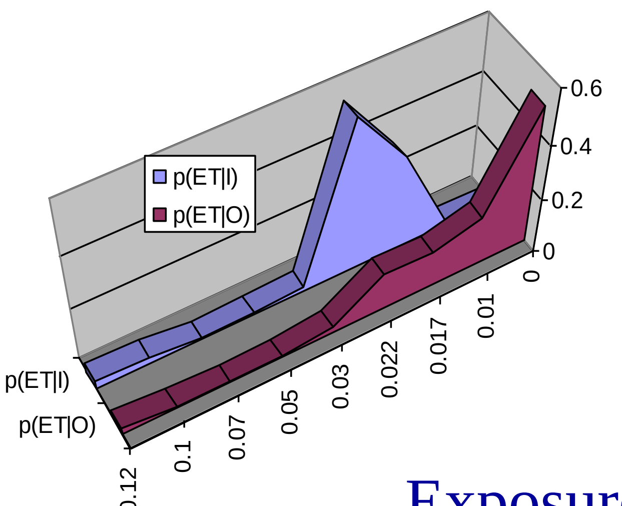

Exposure Time

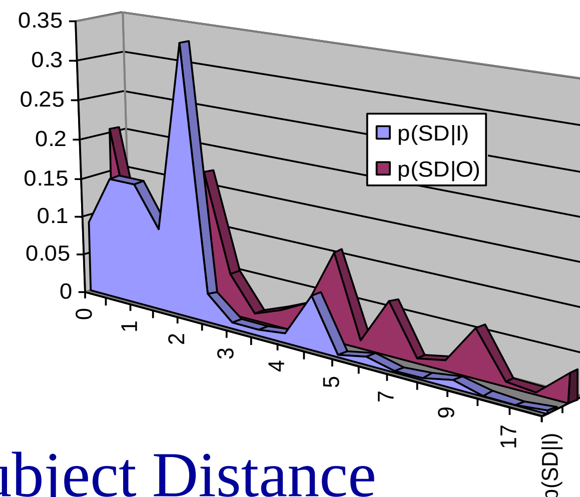

Subject Distance

Why we need Bayes Rule

Problem:

We know conditional probabilities like P(*flash was on* | indoor)

We want to find conditional probabilities like

P(indoor | flash was on, exp time = 0.017, sd=8 ft, SVM output)

Let ω = class of image, and x = all the evidence.

More generally, we know P(x | ω) from the training set (why?)

But we want P(ω | x)

$$p(\omega_i \mid x) = \frac{p(x \mid \omega_i)\, p(\omega_i)}{p(x)}$$

Using Bayes Rule

$P(\omega|x) = P(x|\omega)P(\omega)/P(x)$

The denominator is constant for an image, so

$P(\omega|x) = \alpha P(x|\omega)P(\omega)$

Using Bayes Rule

$P(\omega|x) = P(x|\omega)P(\omega)/P(x)$

The denominator is constant for an image, so

$P(\omega|x) = \alpha P(x|\omega)P(\omega)$

We have two types of features, from image metadata (M) and from low-level features, like color (L)

Conditional independence means $P(x|\omega) = P(M|\omega)P(L|\omega)$

$$P(\omega|X) = \alpha P(M|\omega)\ P(L|\omega)\ P(\omega)$$

From histograms **From SVM** **Priors (initial bias)**

Bayesian network

- Efficient way to encode conditional probability distributions and calculate marginals
- Use for classification by having the classification node at the root
  - Examples
    - Indoor-outdoor classification
    - Automatic image orientation detection

Indoor vs. outdoor classification

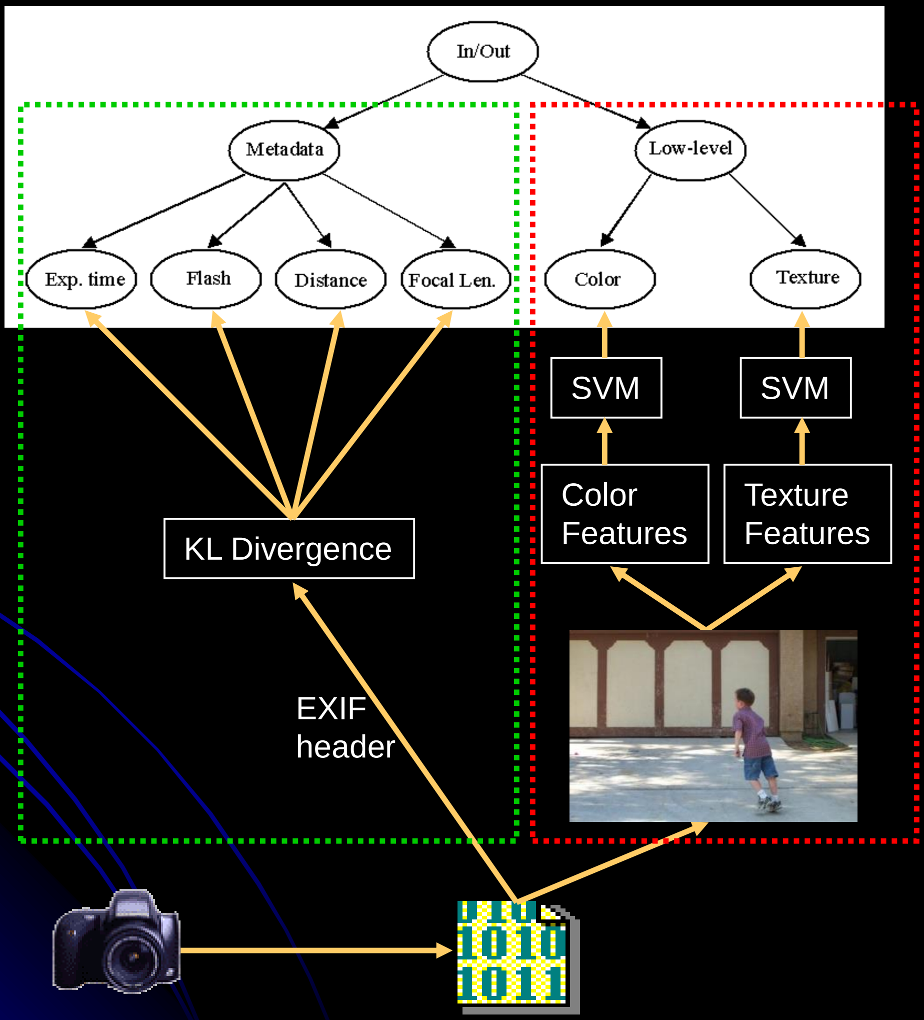

Each edge in the graph has an associated matrix of conditional probabilities

Effects of Image Capture Context

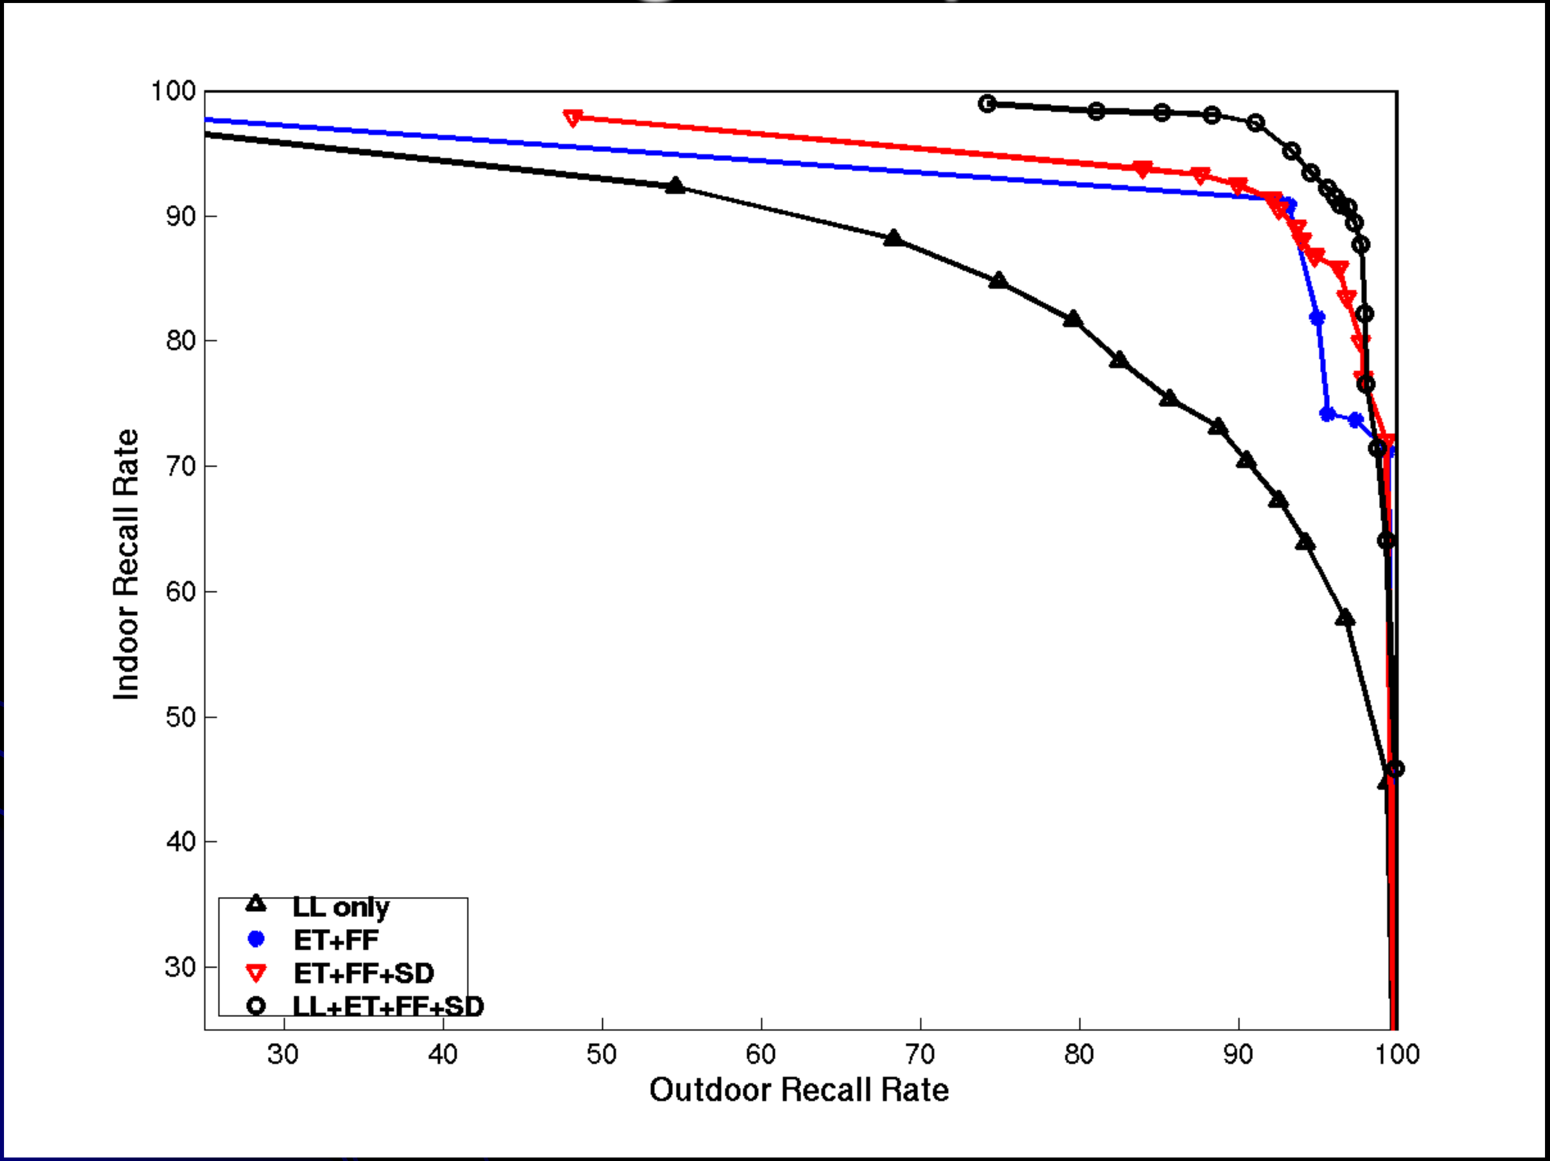

Recall for a class C is fraction of C classified correctly

Orientation detection

- See IEEE TPAMI paper
  - Hardcopy or posted
- Also uses single-feature Bayesian classifier (answer to #1-4)
- Keys:
  - 4-class problem (North, South, East, West)
  - Priors **really** helped here!

- You should be able to understand the two papers (both posted)