MIT OpenCourseWare
http://ocw.mit.edu

HST.582J / 6.555J / 16.456J Biomedical Signal and Image Processing Spring 2007

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.

# Automated Decision Making Systems 

Probability, Classification, Model Estimation

## Information and Statistics

One the use of statistics:
"There are three kind of lies: lies, damned lies, and statistics"

- Benjamin Disraeli (popularized by Mark Twain)

On the value of information:
"And when we were finished renovating our house, we had only $\$ 24.00$ left in the bank only because the plumber didn't know about it."

- Mark Twain (from a speech paraphrasing one of his books)


## Elements of Decision Making Systems

1. Probability

- A quantitative way of modeling uncertainty.

2. Statistical Classification

- application of probability models to inference.
- incorporates a notion of optimality

3. Model Estimation

- we rarely (OK never) know the model beforehand.
- can we estimate the model from labeled observations.


Cite as: John Fisher. Course materials for HST.582J / 6.555J / 16.456J, Biomedical Signal and Image Processing, Spring 2007. MIT OpenCourseWare (http://ocw.mit.edu), Massachusetts Institute of Technology. Downloaded on [DD Month YYYY].

## Concepts

- In many experiments there is some element of randomness the we are unable to explain.
- Probability and statistics are mathematical tools for reasoning in the face of such uncertainty.
- They allow us to answer questions quantitatively such as
- Is the signal present or not?
- Binary : YES or NO
- How certain am I?
- Continuous : Degree of confidence
- We can design systems for which
- Single use performance has an element of uncertainty
- Average case performance is predictable


## Anomalous behavior (example)

- How do quantify our belief that these are anomalies?
- How might we detect them automatically?


April 07

## Detection of signals in noise

- In which of these plots is the signal present?
-Why are we more certain in some cases than others?



## Coin Flipping

- Fairly simple probability modeling problem
- Binary hypothesis testing
- Many decision systems come down to making a decision on the basis of a biased coin flip (or N -sided die)


## Bayes' Rule

- Bayes' rule plays an important role in classification, inference, and estimation.

$$
\begin{aligned}
P(A B) & =P(A \mid B) P(B) \\
& =P(B \mid A) P(A)
\end{aligned}
$$

$$
\begin{aligned}
P(A \mid B) & =\frac{P(B \mid A) P(A)}{P(B)} \\
& =\frac{P(B A)}{P(B)} \\
P(B \mid A) & =\frac{P(A \mid B) P(B)}{P(A)} \\
& =\frac{P(A B)}{P(A)}
\end{aligned}
$$

- A useful thing to remember is that conditional probability relationships can be derived from a Venn diagram. Bayes' rule then arises from straightforward algebraic manipulation.


## Heads/Tails Conditioning Example

- If I flip two coins and tell you at least one of them is "heads" what is the probability that at least one of them is "tails"?
- The events of interest are the set of outcomes where at least

|  |  | $2^{\text {nd }}$ |  |
| :---: | :---: | :---: | :---: |
|  | flip |  |  |
|  |  | $H$ | $T$ |
| $\stackrel{\rightharpoonup}{u}$ | $H$ | $H H$ | $H T$ |
| $\stackrel{\rightharpoonup}{\bar{\sigma}}$ | $T$ | TH | TT |
|  |  |  |  | one of the results is a head.

- The point of this example is two-fold
- Keep track of your sample space and events of interest.
- Bayes' rule tells how to incorporate information in order to adjust probability.


## Heads/Tails Conditioning Example

- The probability that at least one of the results is heads is $\frac{3}{4}$ by simple counting.
- The probability that both of the coins are heads is $\frac{1}{4}$

|  |  | $2^{\text {nd }}$ flip |  |
| :---: | :---: | :---: | :---: |
|  |  | H | T |
| $\stackrel{\rightharpoonup}{\square}$ | H | HH | HT |
| $\cdots$ | T | TH | TT |

$$
\begin{aligned}
A & =\text { the "other" coin is heads } \\
B & =\text { at least one of the coins is heads } \\
A B & =\text { both of the coins are heads } \\
P(A \mid B) & =\frac{P(B A)}{P(B)}
\end{aligned}
$$

- The chance of winning is 1 in 3
- Equivalently, the odds of
 winning are 1 to 2


## Defining Probability (Frequentist vs. Axiomatic)

The probability of an event is the number of times we expect a specific outcome relative to the number of times we conduct the experiment.

## Define:

- $N$ : the number of trials
- $N_{A}, N_{B}$ : the number of times events $A$ and $B$ are observed.
- Events $A$ and $B$ are mutually exclusive (i.e. observing one precludes observing the other).


## Empirical definition:

- Probability is defined as a limit over observations

| $P\{A\}$ | $=\lim _{N \rightarrow \infty}\left(\frac{N_{A}}{N}\right)$ |
| ---: | :--- |
| $P\{B\}$ | $=\lim _{N \rightarrow \infty}\left(\frac{N_{B}}{N}\right)$ |
| $P\{A+B\}$ | $=\lim _{N \rightarrow \infty}\left(\frac{N_{A}+N_{B}}{N}\right)$ |

Axiomatic definition:

- Probability is derived from its properties

$$
0 \leq P\{A\}, P\{B\} \leq 1
$$

$$
P\{\text { the certain event }\}=1
$$

$$
P\{A+B\}=P\{A\}+P\{B\}
$$



## 4 out of 5 Dentists...

-What does this statement mean?

- How can we attach meaning/significance to the claim?
- An example of a frequentist vs. Bayesian viewpoint
- The difference (in this case) lies in:
- The assumption regarding how the data is generated
- The way in which we can express certainty about our answer
- Asympotitically (as we get more observations) they both converge to the same answer (but at different rates).


## Sample without Replacement, Order Matters

Begin with N empty boxes

- each term represents the number of different choices we have at each stage

$$
N \times(N-1) \times(N-2) \times \cdots \times(N-k+1)
$$

- this can be re-written as

Choose the kth box from the $N-k+1$ remaining choices



- and then "simplified" to


At left: color indicates the order in which we filled the boxes. Any sample which fills the same boxes, but has a different color in any box (there will be at least 2) is considered a different sample.

## Sample without Replacement, Order doesn't Matter

- The sampling procedure is the same as the previous except that we don't keep track of the colors.
- The number of sample draws with the same filled boxes is equal to the number of ways we can re-order (permute) the colors.
- The result is to reduce the total number of draws by that factor.

$$
\frac{N!}{(N-k)!} \rightarrow \frac{N!}{(N-k)!k!}=\binom{N}{k}
$$



## Cumulative Distributions Functions (PDFs)

- cumulative distribution function (CDF) divides a continuous sample space into two events

$$
P_{X}(x)=\operatorname{Pr}\{X \leq x\} \quad 1-P_{X}(x)=\operatorname{Pr}\{X>x\}
$$

- It has the following properties

$$
\begin{aligned}
P_{X}(-\infty) & =0 \\
P_{X}(\infty) & =1 \\
0 & \leq P_{X}(x) \leq 1 \\
P_{X}(x+\Delta) & \geq P_{X}(x) \quad ; \quad \Delta \geq 0
\end{aligned}
$$

## Probability Density Functions (PDFs)

- probability density function (PDF) is defined in terms of the CDF

$$
\begin{aligned}
P_{X}(x) & =\int_{-\infty}^{x} p_{x}(u) d u \\
p_{X}(x) & =\frac{\partial}{\partial x} P_{X}(x)
\end{aligned}
$$

- Some properties which follow are:

$$
\begin{array}{r}
\int_{-\infty}^{\infty} p_{x}(u) d u=1 \\
p_{X}(x) \geq 0
\end{array}
$$

## Expectation

- Given a function of a random variable (i.e. $g(X)$ ) we define it's expected value as:

$$
\begin{aligned}
E\{g(X)\} & =\sum_{i=1}^{N} g\left(x_{i}\right) p_{x}\left(x_{i}\right) \\
& =\int_{\Omega_{X}} g(u) p_{x}(u) d u
\end{aligned}
$$

- For the mean, variance, and entropy (continous examples):

- Expectation is linear (see variance example once we've defined joint density function and statistical independence)

$$
\begin{aligned}
E\{\alpha f(x)+\beta g(x)\}= & \alpha E\{f(x)\} \\
& +\beta E\{g(x)\}
\end{aligned}
$$

- Expectation is with regard to ALL random variables within the arguments.
- This is important for multidimensional and joint random variables.


## Multiple Random Variables (Joint Densities)

We can define a density over multiple random variables in a similar fashion as we did for a single random variable.

1. We define the probability of the event $\{X \leq x$ AND $Y \leq y\}$ as a function of $x$ and $y$.
2. The density is the function we integrate to compute the probability.

$$
\begin{aligned}
P_{X Y}(x, y) & =\operatorname{Pr}\{X \leq x \text { AND } Y \leq y\} \\
& =\int_{-\infty}^{x} \int_{-\infty}^{y} p_{x y}(u, v) d u d v \\
p_{X Y}(x, y) & =\frac{\partial^{2}}{\partial x \partial y} P_{X Y}(x, y)
\end{aligned}
$$



## Conditional Density

Given a joint density or mass function over two random variables we can define the conditional density similar to conditional probability from Venn diagrams

$$
p_{x \mid y}(x \mid y)=\frac{p_{x y}(x, y)}{p_{y}(y)}
$$

This is, it is not of practical use unless we condition on $Y$ equal to a value versus letting it remain a variable (creating an actual density)

$$
p_{x \mid y}\left(x \mid y=y_{o}\right)=\frac{p_{x y}\left(x, y_{o}\right)}{p_{y}\left(y_{o}\right)}
$$

We also get the following relationship

$$
\begin{aligned}
p_{x \mid y}(x \mid y) p_{y}(y) & =p_{x y}(x, y) \\
& =p_{y \mid x}(y \mid x) p_{x}(x)
\end{aligned}
$$



$p_{x \mid y}\left(x \mid Y=y_{o}\right)$ is $p_{x y}\left(x, y_{o}\right)$ (a slice of the joint density, normalized by $p_{y}\left(y_{o}\right)$ so that it integrates to unity (i.e. it is a valid density)

## Bayes' Rule

- For continuous random variables, Bayes' rule is essentially the same (again just an algebraic manipulation of the definition of a conditional density).

$$
p_{X \mid Y}(x \mid y)=\frac{p_{Y \mid X}(y \mid x) p_{X}(x)}{p_{Y}(y)}
$$

- This relationship will be very useful when we start looking at classification and detection.


## Binary Hypothesis Testing (Neyman-Pearson) (and a "simplification" of the notation)

- 2-Class problems are equivalent to the binary hypothesis testing problem.

$$
\begin{aligned}
& \hline H_{1}: x \sim p_{X \mid H_{1}}\left(x \mid H_{1} \text { is true }\right) \\
& H_{0}: x \sim p_{X \mid H_{0}}\left(x \mid H_{0} \text { is true }\right) \\
& \hline
\end{aligned}
$$

The goal is estimate which Hypothesis is true (i.e. from which class our sample came from).

- A minor change in notation will make the following discussion a little simpler.

$$
\left.\begin{array}{l}
p_{1}(x)=p_{X \mid H_{1}}\left(x \mid H_{1} \text { is true }\right) \\
p_{0}(x)=p_{X \mid H_{0}}\left(x \mid H_{0} \text { is true }\right)
\end{array}\right\} \begin{aligned}
& \text { Probability density models for the } \\
& \text { measurement } x \text { depending on which } \\
& \text { hypothesis is in effect. }
\end{aligned}
$$



- Decision rules are functions which map measurements to choices.
- In the binary case we can write it as

$$
\phi(x)=\left\{\begin{array}{lll}
1 & ; & x \in R_{1} \\
0 & ; & x \in R_{0}
\end{array}\right.
$$



## Binary Hypothesis Testing (Bayesian)

- 2-Class problems are equivalent to the binary hypothesis testing problem.

$$
\begin{aligned}
H_{1} & : x \sim p_{X \mid H_{1}}\left(x \mid H_{1} \text { is true }\right) \\
H_{0} & : x \sim p_{X \mid H_{0}}\left(x \mid H_{0} \text { is true }\right)
\end{aligned}
$$

The goal is estimate which Hypothesis is true (i.e. from which class our sample came from).

- A minor change in notation will make

Marginal density of $X$

$$
p_{x}(x)=P_{1} p_{1}(x)+P_{0} p_{0}(x)
$$

Conditional probability of the hypothesis $H_{i}$ given $X$

$$
\begin{aligned}
P_{H_{i} \mid x}\left(H_{i} \mid x\right) & =\frac{P_{i} p_{i}(x)}{p_{x}(x)} \\
& =\frac{P_{i} p_{i}(x)}{P_{1} p_{1}(x)+P_{0} p_{0}(x)}
\end{aligned}
$$ the following discussion a little simpler.

$\left.\begin{array}{l}P_{1}=\operatorname{Pr}\left(H=H_{1}\right) \\ P_{0}=\operatorname{Pr}\left(H=H_{0}\right)\end{array}\right\}$ Prior probabilities of each class

$$
\left.\begin{array}{l}
p_{1}(x)=p_{X \mid H_{1}}\left(x \mid H_{1} \text { is true }\right) \\
p_{0}(x)=p_{X \mid H_{0}}\left(x \mid H_{0} \text { is true }\right)
\end{array}\right\} \begin{aligned}
& \text { Class-conditional probability density } \\
& \text { models for the measurement } \mathrm{x}
\end{aligned}
$$



- So given observations of $x$, how should select our best guess of $H_{i}$ ?
- Specifically, what is a good criterion for making that assignment?
- Which $H_{i}$ should we select before we observe $x$.


## Bayes Classifier



- A reasonable criterion for guessing values of $H$ given observations of $X$ is to minimize the probability of error.
- The classifier which achieves this minimization is the Bayes classifier.

- An error is comprised of two events


## $E_{1}: X$ falling in $R_{0}$ AND $H_{1}$ being correct

- These are mutually exclusive events so their joint probability is the sum of their individual probabilities

$$
\begin{aligned}
P_{E} & =\operatorname{Pr}\left\{E_{1}\right\}+\operatorname{Pr}\left\{E_{0}\right\} \\
& =P_{1} \operatorname{Pr}\left\{X \in R_{0} \mid H_{1}\right\}+P_{0} \operatorname{Pr}\left\{X \in R_{1} \mid H_{0}\right\} \\
& =P_{1} \int_{R_{0}} p_{1}(x) d x+P_{0} \int_{R_{1}} p_{0}(x) d x
\end{aligned}
$$

## Minimum Probability of Misclassification

- So now let's choose regions to minimize the probability of error.

$$
\begin{aligned}
P_{E} & =P_{1} \int_{R_{0}} p_{1}(x) d x+P_{0} \int_{R_{1}} p_{0}(x) d x \\
& =P_{1}\left(1-\int_{R_{1}} p_{1}(x) d x\right)+P_{0} \int_{R_{1}} p_{0}(x) d x \\
& =P_{1}+\int_{R_{1}}(\underbrace{P_{0} p_{0}(x)}_{\geq 0}-\underbrace{P_{1} p_{1}(x)}_{\geq 0}) d x
\end{aligned}
$$

- In the second step we just change the region over which integrate for one of the terms (these are complementary events).
- In the third step we collect terms and note that all underbraced terms in the integrand are non-negative.
- If we want to choose regions (remember choosing region 1 effectively chooses region 2) to minimize $P_{E}$ then we should se $\dagger$ region 1 to be such that the integrand is negative.


## Minimum Probability of Misclassification

- Consequently, for minimum probability of misclassification (which is the Bayes error), $R_{1}$ is defined as

$$
R_{1}=\left\{x: P_{1} p_{1}(x)>P_{2} p_{2}(x)\right\}
$$

- $R_{2}$ is the complement. The boundary is where we have equality.
- Equivalently we can write the condition as when the likelihood ratio for $\mathrm{H}_{1}$ vs $\mathrm{H}_{0}$ exceeds the PRIOR odds of $\mathrm{H}_{0}$ vs $\mathrm{H}_{1}$

$$
R_{1}=\left\{x: \frac{p_{1}(x)}{p_{0}(x)}>\frac{P_{0}}{P_{1}}\right\}
$$

## Risk Adjusted Classifiers

Suppose that making one type of error is more of a concern than making another. For example, it is worse to declare $H_{1}$ when $H_{2}$ is true then vice versa.

- This is captured by the notion of "cost".
$C_{i j}=$ cost of declaring $H_{i}$ when $H_{j}$ is correct
- In the binary case this leads to a cost matrix.

- The Risk Adjusted Classifier tries to minimize the expected "cost"


## Derivation

- We'll simplify by assuming that $C_{11}=C_{22}=0$ (there is zero cost to being correct) and that all other costs are positive.
- Think of cost as a piecewise constant function of $X$.
- If we divide $X$ into decision regions we can compute the expected cost as the cost of being wrong times the probability of a sample falling into that region.

$$
\begin{aligned}
E\{C(x, H)\} & =\int_{R_{0}} C_{01} P_{1} p_{1}(x) d x+\int_{R_{1}} C_{10} P_{0} p_{0}(x) d x \\
& =C_{01} P_{1}\left(1-\int_{R_{1}} p_{1}(x) d x\right)+C_{10} P_{0} \int_{R_{1}} p_{0}(x) d x \\
& =C_{01} P_{1}+\int_{R_{1}}(\underbrace{C_{10} P_{0} p_{0}(x)}_{\geq 0}-\underbrace{C_{01} P_{1} p_{1}(x)}_{\geq 0}) d x
\end{aligned}
$$

## Risk Adjusted Classifiers

Expected Cost is then

$$
E\{C(x, H)\}=C_{01} P_{1}+\int_{R_{1}}(\underbrace{C_{10} P_{0} P_{0}(x)}_{\geq 0}-\underbrace{C_{01} P_{1} p_{1}(x)}_{\geq 0}) d x
$$

- As in the minimum probability of error classifier, we note that all terms are positive in the integral, so to minimize expected "cost" choose $R_{1}$ to be:
$R_{1}=\left\{x: C_{01} P_{1} p_{1}(x)>C_{10} P_{0} p_{0}(x)\right\}$
- Alternatively

$$
R_{1}=\left\{x: \frac{p_{1}(x)}{p_{0}(x)}>\frac{C_{10} P_{0}}{C_{01} P_{1}}\right\}
$$

- If $C_{10}=C_{01}$ then the risk adjusted classifier is equivalent to the minimum probability of error classifier.
- Another interpretation of "costs" is an adjustment to the prior probabilities.

$$
\frac{P_{0}^{\mathrm{adj}}}{P_{1}^{\mathrm{adj}}}=\frac{C_{10} P_{0}}{C_{01} P_{1}}
$$

- Then the risk adjusted classifier is equivalent to the minimum probability of error classifier with prior probabilities equal to $P_{1}$ adj and $P_{0}{ }^{\text {adj, }}$, respectively.


Cite as: John Fisher. Course materials for HST.582J / 6.555J / 16.456J, Biomedical Signal and Image Processing, Spring 2007. MIT OpenCourseWare (http://ocw.mit.edu), Massachusetts Institute of Technology. Downloaded on [DD Month YYYY].

## Okay, so what.

All of this is great. We now know what to do in a few classic cases if some nice person hands us all of the probability models.

- In general we aren't given the models - What do we do?

Density estimation to the rescue.

- While we may not have the models, often we do have a collection of labeled measurements, that is a set of $\left\{x, H_{j}\right\}$.
- From these we can estimate the class-conditional densities. Important issues will be:
- How "close" will the estimate be to the true model.
- How does "closeness" impact on classification performance?
- What types of estimators are appropriate (parametric vs. nonparametric).
- Can we avoid density estimation and go straight to estimating the decision rule directly? (generative approaches versus discriminative approaches)

