# Statistical Approaches to Learning and Discovery 

# Week 1: Some Basic Concepts from Statistics and Information Theory 

January 15, 2003

## Today's Agenda

- High-level view
- Sufficient statistics
- Data processing inequality (no free statistical lunch)
- Estimators: Bias, variance, Cramér-Rao
- Exponential families


## Supervised vs. Unsupervised Learning

Have a sequence (or set) of inputs $x_{1}, x_{2}, x_{3}, \ldots$, ,
"naturally" occuring, collected by hand, or generated by machine

Supervised Learning: Machine is given desired outputs $y_{1}, y_{2}, y_{3}, \ldots$, and goal is to learn to produce the correct output given a new input. This doesn't specify how "correct" should be assessed... Distinction between classification (discrete $y_{i}$ ) and regression (continuous $y_{i}$ ).

Unsupervised Learning: Goal is to build representations of $x$ that can be used for reasoning, decision making, predicting, communicating, etc. Task is often not well specified.

## Supervised vs. Unsupervised Learning (cont.).

Semi-Supervised: Same as supervised, but some of the values $y_{i}$ are missing in the training set, and the unlabeled $x_{i}^{\prime} s$ are incorporated.

## Inference vs. Learning

Estimation/Learning: Selecting parameters, a distribution over parameters, or a set of cdf's for a statistical problem based on data.

Inference: Making predictions, computing statistics, expectations, or marginal probabilities for a statistical model that has already been estimated/learned.

## Parameters

A statistical family with a finite collection of adjustable parameters is the starting point for a parametric estimation problem.

If there are an infinite number of adjustable parameterstypically entire functions or cdf's, then the problem (or approach) is said to be non-parametric.

## Parametric vs. Non-Parametric

This can be confusing, since often "non-parametric" problems seem to have many more "parameters" than a typical parametric problem.

Non-parametric approaches make fewer assumptions about the form or "shape" of the distribution being estimated.

However, the distinction is sometimes subtle (e.g., neural nets)

## A Simple Estimator

Suppose that $X_{1}, X_{2}, \ldots, X_{n} \sim \mathcal{N}(\theta, 1)$ (iid).
We want to determine $\theta$ from the sample. Two options:

1. $X_{1}$, since clearly $E\left(X_{1}\right)=\theta$
2. $\bar{X}_{n}=\frac{1}{n}\left(X_{1}+X_{2}+\cdots+X_{n}\right)$. Also mean $\theta$

Which is better? Well, depends what "good" means. In fact, $\bar{X}_{n}$ is the minimum mean squared error unbiased estimator.

Role of computation is not emphasized in classical statistics...

## Sufficiency

Suppose $X_{i} \sim f(\cdot \mid \theta)$, for $\theta \in \Theta \subset \mathbb{R}^{m}$.
A statistic is just a function of the sample: $T\left(X_{1}, \ldots, X_{n}\right)$. It's a random variable.

Supose there is a statistician and a computer scientist. The statistician has all of the data $X_{1}, \ldots, X_{n}$. The computer scientist only keeps a "hash" of the data $T\left(X_{1}, \ldots, X_{n}\right)$.

Who can make better estimates of $\theta$, or in general make better inferences?

## Sufficiency (cont.)

In general, the statistician can do better, but if $T$ is a sufficient statistic then the computer scientist will be able to do just as well.

In this case, intuitively, $T\left(X_{1}, \ldots, X_{n}\right)$ contains all of the "information" in the sample about $\theta$, and the individual values are irrelevant.
(We'll give a precise meaning to this later...)

## Example 1: Bernoulli

$X_{1}, X_{2}, \ldots, X_{n}$ are $n$ coin tosses. $X_{i} \sim$ Bernoulli $(\theta)$.
Given $n$, the number of "heads" is a sufficient statistic for $\theta$.

$$
\operatorname{Pr}\left(X_{i}=x_{i} \mid n, T(X)=k\right)= \begin{cases}\frac{1}{\binom{n}{k}} & \text { if } \sum_{i} x_{i}=k \\ 0 & \text { otherwise }\end{cases}
$$

More generally, for a multinomial $\theta=\left(p_{1}, p_{2}, \ldots, p_{t}\right)$, the vector of counts $\left(n_{1}, \ldots, n_{t}\right)$ is sufficient, where $n_{j}=$ $\sum_{i=1}^{n} \delta\left(x_{j}=i\right)$.

## Example 2: Gaussian

Take

$$
f_{\theta}(x)=\frac{1}{\sqrt{2 \pi}} e^{-\frac{1}{2}(x-\theta)^{2}}=\mathcal{N}(\theta, 1)
$$

A sufficient statistic is $\bar{X}_{n}=\frac{1}{n} \sum_{i} X_{i}$.
$\bar{X}_{n}$ and $\frac{1}{n} \sum_{i}\left(X_{i}-\bar{X}_{n}\right)^{2}$ are sufficient for $\mu$ and $\sigma^{2}$ if $\theta=$ $\left(\mu, \sigma^{2}\right)$.

## Example 2: Uniform

Take

$$
X_{i} \sim \operatorname{Uniform}(0, \theta)
$$

A sufficient statistic is $T\left(X_{1}, \ldots, X_{n}\right)=\max _{i} X_{i}$.

## Neyman Factorization Criterion

A statistic $T\left(X_{1}, \ldots, X_{n}\right)$ is sufficient for $\theta$ if and only if the joint pdf can be factored as

$$
f_{n}(\boldsymbol{x} \mid \theta)=u(\boldsymbol{x}) v(T((x), \theta))
$$

## Information

Now let's go back and give a precise meaning to "all of the relevant information about $\theta$ is in the sufficient statistic"

So far, we've only been thinking of $\mathcal{X}_{i}$ as random, not $\theta$. We'll now need to treat $\theta$ as a random variable.

## Data Processing Inequality

"No clever manipulation of the data can improve the inferences that can be made from the data."

Note: this is a statement about statistics, not computation

## Information Theory Concepts

For a discrete distribution $p_{1}, p_{2}, \ldots, p_{n}$, or random variable $X$ with $p\left(X=x_{i}\right)=p_{i}$, entropy

$$
H(p)=-\sum_{i=1}^{n} p_{i} \log _{2} p_{i}
$$

in bits of information.
Conditional entropy $H(X \mid Y)$ is

$$
\begin{aligned}
H(X \mid Y) & =\sum_{y} p(Y=y) H(X \mid Y=y) \\
& =-\sum_{y} p(y) \sum_{x} p(x \mid y) \log _{2} p(x \mid y)
\end{aligned}
$$

## Information Theory Concepts (cont.)

Mutual information $I(X ; Y)$

$$
\begin{aligned}
I(X ; Y) & =H(X)-H(X \mid Y) \\
& =H(Y)-H(Y \mid X) \\
& =\sum_{x, y} p(x, y) \log \frac{p(x, y)}{p(x) p(y)}
\end{aligned}
$$

Informally, "the average value of a hint." Amount by which knowing $X$ reduces the average code length needed to compress $Y$.

## Markov Chains

$X \longrightarrow Y \longrightarrow Z$ forms a Markov chain in case the conditional distribution of $Z$ is independent of $X$.

Equivalently, in case $X$ and $Z$ are conditionally independent given $Y$. Note: "time" symmetric
(Concept extends to spatial processes, or "random fields")

## Data Processing Inequality

If $X \longrightarrow Y \longrightarrow Z$ is a Markov chain, then

$$
I(X ; Y) \geq I(X ; Z)
$$

In particular, since $X \longrightarrow Y \longrightarrow g(Y)$,

$$
I(X ; Y) \geq I(X ; g(Y))
$$

## Sufficiency Revisited

Since $\Theta \longrightarrow X \longrightarrow T(X)$ is a Markov chain, we have that $I(\Theta ; X) \geq I(\Theta ; T(X))$.

However, if $\Theta \longrightarrow T(X) \longrightarrow X$ is a Markov chain also, i.e., $T(X)$ is sufficient, then we have equality:

$$
I(\Theta ; T(X))=I(\Theta ; X)
$$

(Historical note: Notion of sufficiency due to Fisher; Formulation in terms of mutual information due to Kullback.)

## Estimation: Basic Concepts

Point estimation: choose a single parameter $\hat{\theta}$ or cdf, or other prediction.
Note: $\hat{\theta}$ is a random variable, since it is a function of the data (which is random):

$$
\hat{\theta}_{n}=g\left(X_{1}, X_{2}, \ldots, X_{n}\right)
$$

where $g$ represents an algorithm for computing the point estimate.

## Bias

The bias of a point estimator is

$$
\operatorname{bias}\left(\hat{\theta}_{n}\right)=E_{F}\left[\hat{\theta}_{n}\right]-\theta
$$

An estimator is unbiased if

$$
E_{F}\left[\hat{\theta}_{n}\right]=\theta
$$

where $X_{1}, X_{2}, \ldots, X_{n}$ are iid $\sim F$.

## Consistency

A point estimate of a parameter $\theta$ is consistent if

$$
\hat{\theta}_{n} \longrightarrow \theta \quad \text { (in probability) }
$$

The standard error is the standard deviation of $\hat{\theta}_{n}$ :

$$
\operatorname{se}\left(\hat{\theta}_{n}\right)=\sqrt{E_{F}\left(\hat{\theta}_{n}-E_{F}\left(\hat{\theta}_{n}\right)\right)^{2}}
$$

For an unbiased estimator this is

$$
\operatorname{se}\left(\hat{\theta}_{n}\right)=\sqrt{E_{F}\left(\hat{\theta}_{n}-\theta\right)^{2}}
$$

Note that since the expectation is the "true" expectation over the data, this is in general impossible to compute.

## Example

Let $X_{i} \sim \operatorname{Bernoulli}(\theta)$.

$$
\hat{\theta}_{n}=\frac{1}{n} \sum_{i=1}^{n} X_{i}
$$

satisfies

$$
E\left[\hat{\theta}_{n}\right]=\frac{1}{n} \cdot n \theta=\theta
$$

so this is an unbiased estimate of $\theta$.

## Example (cont.)

The standard error is

$$
\begin{aligned}
\operatorname{se}\left(\hat{\theta}_{n}\right) & =\sqrt{E\left(\left(\frac{1}{n} \sum X_{i}\right)^{2}-\theta^{2}\right)} \\
& =\sqrt{\frac{\theta(1-\theta)}{n}}
\end{aligned}
$$

and so can't be computed. The estimated standard error is

$$
\hat{\mathrm{se}}=\sqrt{\frac{\hat{\theta}_{n}\left(1-\hat{\theta}_{n}\right)}{n}}
$$

## Mean Squared Error

The mean squared error (MSE) of an estimator is

$$
E\left[\left(\hat{\theta}_{n}-\theta\right)^{2}\right]
$$

Another way of looking at this is

$$
\begin{aligned}
M S E & =E\left[\left(\hat{\theta}_{n}-\theta\right)^{2}\right] \\
& =E\left[\left(\left(\hat{\theta}_{n}-E\left[\hat{\theta}_{n}\right]\right)^{2}+\left(E\left[\hat{\theta}_{n}\right]-\theta\right)\right)^{2}\right] \\
& =\operatorname{Var}\left(\hat{\theta}_{n}\right)+\operatorname{bias}^{2}\left(\hat{\theta}_{n}\right)
\end{aligned}
$$

Fundamental tradeoff.

## Asymptotically Normal

An estimator is asymptotically normal in case

$$
\frac{\hat{\theta}_{n}-\theta}{\operatorname{se}\left(\hat{\theta}_{n}\right)} \leadsto \mathcal{N}(0,1)
$$

## Point Estimation for Parametric Families

We have a family $\mathcal{F}=\left\{f_{\theta}(x), \theta \in \Theta\right\}$ and want to estimate certain parameters of interest.

## Maximum Likelihood

The most commonly used method for point estimation. Given a family $\mathcal{F}=\{f(x \mid \theta)\}$ and data $X_{1}, X_{2}, \ldots, X_{n}$, the likelihood function is defined as

$$
\mathcal{L}_{n}(\theta)=\prod_{i} f\left(X_{i} \mid \theta\right)
$$

and the log-likelihood function is given by

$$
\begin{aligned}
\ell_{n}(\theta) & =\log \mathcal{L}_{n}(\theta) \\
& =\sum_{i} \log f\left(X_{i} \mid \theta\right)
\end{aligned}
$$

## Maximum Likelihood

The maximum likelihood estimator is

$$
\hat{\theta}=\operatorname{argmax}_{\Theta} \ell_{n}(\theta)
$$

(whenever this exists)

## What is the Best Possible Estimator?

What is the minimum variance of an (unbiased) estimator of $\theta$ ?

Take $f(x \mid \theta)$ where $\theta \in \mathbb{R}$ (1-dimensional for simplicity).
Let's look at the change in log-likelihood as a function of $\theta$. The score $s(X, \theta)$ is defined as

$$
s(X, \theta)=\frac{\partial}{\partial \theta} \log f(X \mid \theta)
$$

This has mean zero (with respect to $f(\cdot \mid \theta)$ )

## Fisher Information and Cramér-Rao

Fisher information is the variance of the score:

$$
\begin{aligned}
J(\theta) & =E_{\theta}\left(s^{2}\right) \\
& =E_{\theta}\left(\frac{\partial}{\partial \theta} \log f(X \mid \theta)\right)^{2}
\end{aligned}
$$

Basic additivity property: The Fisher information of $n$ iid samples is $n J(\theta)$.

Cramér-Rao Inequality: The mean-squared error of any unbiased estimator $T(X)$ for $\theta$ satisfies

$$
E_{\theta}(T-\theta)^{2}=\operatorname{Var}(T) \geq \frac{1}{J(\theta)}
$$

## Example: Gaussian

Let $X_{1}, \ldots, X_{n} \sim \mathcal{N}\left(\theta, \sigma^{2}\right)$ where $\sigma$ is known.
It's easy to compute that $J(\theta)=\frac{1}{\sigma^{2}}$.
The sample mean meets the Cramér-Rao lower bound:

$$
E_{\theta}\left(\bar{X}_{n}-\theta\right)^{2}=\frac{\sigma^{2}}{n}=\frac{1}{J_{n}(\theta)}
$$

It is an efficient estimator

## Asymptotic Normality of the MLE

Under some regularity conditions, the MLE is asymptotically normal, with standard error given by the inverse Fisher information:

$$
\frac{(\widehat{\theta}-\theta)}{\sqrt{1 / n J(\theta)}} \leadsto \mathcal{N}(0,1)
$$

This enables us to compute asymptotic confidence intervals

## Different Emphasis for Estimation/Learning

| Traditional Statistics | Machine Learning |
| :---: | :---: |
| consistency | computational efficiency |
| bias | statistical efficiency |
| statistical efficiency | bias |
| computational efficiency | consistency |

