ELEG 5633 Detection and Estimation
Introduction

Jingxian Wu

Department of Electrical Engineering
University of Arkansas

January 14, 2019

Goal

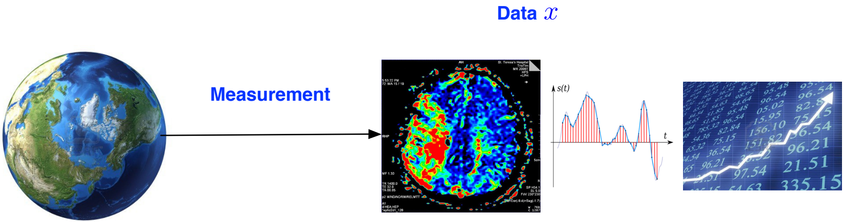

Statistical signal processing involves three processes:

- measurement: $\{x\}$
- modeling: $\{p(x|\theta)\}_{\theta\in\Theta}$
- inference: which value of θ fits the data best?

Goal: Infer value of unknown state of nature based on noisy measurements mathematically, optimally.

Why Probabilistic Models?

- Nonperfect observations: The observations or measurements are often impure and contaminated by effects unknown to us.
  We call those effects noise.
- Nonperfect model: Even the best choice of θ may not perfectly predict new observations.
  We call those effects bias.

How do we model uncertain errors including noise and bias?
$\Rightarrow$ We need a calculus for uncertainty!

The probabilistic framework appears to be the most successful, and in many situations it is physically plausible as well.

Elements of Statistical Signal Processing

Four fundamental inference problems in statistical signal processing:

- Detection: θ may be one of a finite number of values $\{\theta_1, \ldots, \theta_M\}$, we must decide among the M models.
- Parameter Estimation: θ belongs to an infinite set. An extension of detection to infinite model classes.
- Signal Estimation/Prediction: predict the value of a signal y given an observation of another related signal x. The relationship between x and y is characterized by a joint probability distribution, $p(x, y)$.
- Learning: Sometimes we dont know a good model for the relationship between x and y, but we do have a number of "training examples", say $\{(x_i, y_i)\}_{i=1}^n$, that give us some indication of the relationship. The goal of learning is to design a good prediction rule for y given x using these examples, instead of $p(y|x)$.

A Detection Example

010001000100

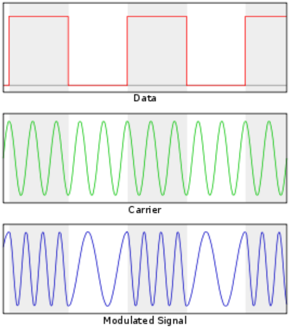

$$0 : s_0(t) = \sin(\omega_0 t)$$

$$1 : s_1(t) = \sin(\omega_1 t)$$

$$x(t) = \begin{cases} s_0(t) + n(t) & \text{if '0' sent} \\ s_1(t) + n(t) & \text{if '1' sent} \end{cases}$$

$$H_0 : \text{'0' is sent}$$

$$H_1 : \text{'1' is sent}$$

$$\hat{H}(x) = H_0? \quad H_1?$$

Another Detection Example

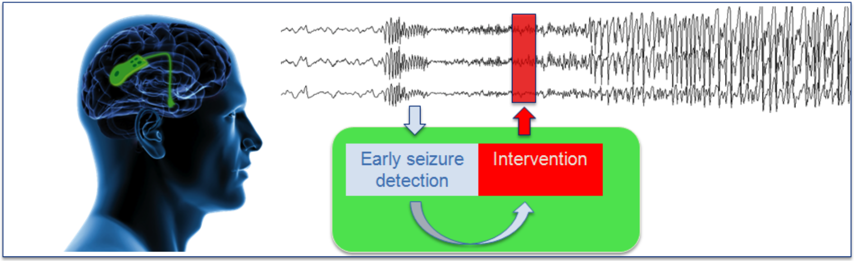

Early seizure detection triggers an intervention.
N different channels: $x_i(t), i = 1, 2, \ldots, N$.
θ has two possible values, corresponding to 'seisure' or 'not seisure'.

A Parameter Estimation Example

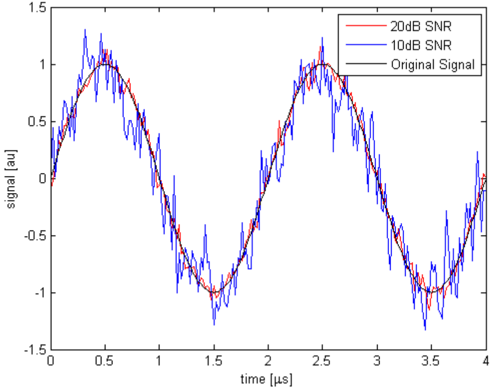

Sampling a noisy sinusoidal signal, we get

$$x_k = A \sin(\omega k + \phi) + w_k,$$

for $k = 1, \ldots, n$.

$\theta = [A, \omega, \phi]$ is unknown.

Specifying a probability distribution for the noises (say Gaussian), would yield a probability distribution for our data $p(x|\theta)$. Given the data $x = [x_1, x_2, ..., x_n]$, how would you estimate θ?

A Signal Estimation Example

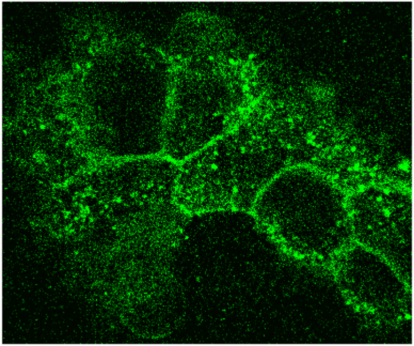

(A) Raw fibroblast cells and microbeads

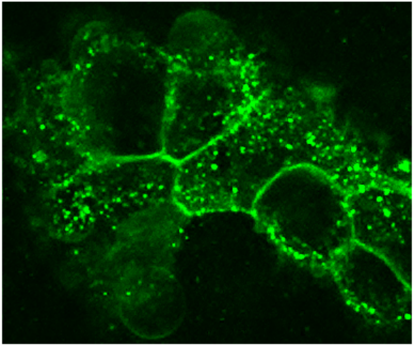

(B) Denoised fibroblast cells and microbeads

We can model the collected data by $y = Hx + w$,

- x: a vector representing the ideal image we wish to recover
- H: a known model of the distortion (represented as a matrix)
- w: a vector of noise

How to "restore the image, i.e., to recover x?

A learning example

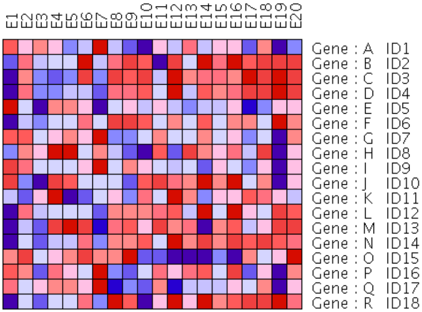

$x \in \mathbb{R}^m$: gene expression levels
$y \in \{+1, -1\}$: indicates whether or not the patient has the disease.
$p(x, y)$ unknown!

Given $\{(x_i, y_i)\}_{i=1}^n$, and a new test vector x, how to predict the value of y?

Course Outline

1. Introduction
   - Review of Probability and Statistics
   - Review of Linear Algebra
2. Detection and Classification
   - Binary Bayesian Hypothesis Test
   - Neyman-Pearson test
   - M-ary Hypothesis Testing
   - Composite Hypothesis Testing
   - Deterministic and Random Signal Detection
3. Estimation Theory
   - Bayesian Estimation
   - Classical Estimation
4. Advanced Topics in Signal Processing and Inference
   - Wiener and Kalman Filtering