## 10-701

## Machine Learning

http://www.cs.cmu.edu/~epxing/Class/10701-15F/

## Organizational info

- All up-to-date info is on the course web page (follow links from my page).
- Instructors
- Eric Xing
- Ziv Bar-Joseph
- TAs: See info on website for recitations, office hours etc.
- See web page for contact info, office hours, etc.
- Piazza would be used for questions / comments. Make sure you are subscribed.



## Zhiting Hu

- Research: large scale machine learning and their applications in NLP/CV.
- Homepage: http://www.cs.cmu.edu/~zhitingh/
- Contact: zhitinghu@gmail.com


# Mrinmaya Sachan (mrinmays@cs.cmu.edu) 



## GHC 8013

Office Hours:
Thu 11AM-12Noon

I am interested in:

- Structured Prediction
- NLP



## Yuntian Deng

- Research: large scale machine learning.
- Contact:
yuntiand@cs.cmu.edu


Xun Zheng (xunzheng@cs.cmu.edu)
I work on...

- MCMC
- Distributed machine learning


## Hao Zhang

## (hao@cs.cmu.edu)



Find me: GHC 8116
Office Hours:
Friday 3.30 pm - 4.30 pm
Interest:
Distributed Machine
Learning
Deep Learning
Applications in computer vision

## Yan Xia

## 2nd year ML Masters student



Research Interests:

- Machine learning applications in drug discovery and development
- Identifying and modeling biological interactions
- 9/3 Intro to probability, MLE
- 9/8 No class
- 9/10 Classification, KNN
- 9/15 No class, Jewish new year
- 9/17 Decision trees - PS1 out
- 9/22 Naïve Bayes
- 9/24 Linear ranrocsion
- 9/26 Logis $11 / 17$ (Monday): Midterm
- 10/1 Perceptron, Neural networks - PS1 due / PS2 out
- 10/6 Deep learning, SVM1
- 10/10 SVM 2
- 10/13 Evaluating classifiers , Bias - Variance decomposition
- 10/15 Ensemble learning - Boosting, RF PS2 due / PS3 out
- 10/20 Unsupervised learning - clustering
- 10/22 Unsupervised learning - clustering / project proposal due
- 10/27 Semi-supervised learning
- 10/29 Learning theory 1 - PS3 due / PS4 out
- 11/3 PAC learning
- 11/5 Graphical models, BN
- 11/10 - BN
- 11/12 - Undirected graphical models / PS4 due
- 11/17 - Midterm
- 11/19 - HMM - PS5 out
- 11/24 - HMM inference
- 12/1 - MDPs / Reinforcement learning / ps5 due
- $12 / 3$ - Topic models-
- 12/4 - Project poster session
- 12/8 -Computational Biology
- $12 / 10$ - no class

Intro and classification (A.K.A. 'supervised learning')

## Probabilistic representation and modeling ('reasoning under uncertainty')

Applications of ML

## Grading

- 5 Problem sets -40\%
- Project - 35\%
- Midterm - 25\%


## Class assignments

- 5 Problem sets
- Most containing both theoretical and programming assignments
- Projects
- Groups of 1-2
- Open ended. Would have to submit a proposal based on your interest. We will also provide suggestions on the website.

Recitations

- Twice a week (same content in both)
- Expand on material learned in class, go over problems from previous classes etc.


## What is Machine Learning?

Easy part: Machine
Hard part: Learning

- Short answer: Methods that can help generalize information from the observed data so that it can be used to make better decisions in the future


## What is Machine Learning?

Longer answer: The term Machine Learning is used to characterize a number of different approaches for generalizing from observed data:

- Supervised learning
- Given a set of features and labels learn a model that will predict a label to a new feature set
- Unsupervised learning
- Discover patterns in data
- Reasoning under uncertainty
- Determine a model of the world either from samples or as you go along
- Active learning
- Select not only model but also which examples to use


## Paradigms of ML

- Supervised learning
- Given $D=\left\{X_{i}, Y_{i}\right\}$ learn a model (or function) $F: X_{k} \rightarrow>Y_{k}$
- Unsupervised learning

Given $D=\left\{X_{i}\right\}$ group the data into Y classes using a model (or function) $F: X_{i} \rightarrow Y_{j}$

- Reinforcement learning (reasoning under uncertainty)

Given $\mathrm{D}=\{$ environment, actions, rewards $\}$ learn a policy and utility functions:
policy: $F 1:\{e, r\}->a$
utility: $F 2$ : $\{a, e\}->R$

- Active learning
- Given $D=\left\{X_{i}, Y_{i}\right\},\left\{X_{j}\right\}$ learn a function $F 1:\left\{X_{j}\right\}$-> $x_{k}$ to maximize the success of the supervised learning function $F 2$ : $\left\{X_{i}, x_{k}\right\}$-> $Y$


## Recommender systems



## NELL：Never－Ending Language Learning

Can computers learn to read？We think so．＂Read the Web＂is a research project that attempts to create a computer system that learns over time to read the web． Since January 2010，our computer system called NELL（Never－Ending Language Learner）has been running continuously，attempting to perform two tasks each day：
－First，it attempts to＂read，＂or extract facts from text found in hundreds of millions of web pages（e．g．，playsInstrument（George＿Harrison， guitar））．
－Second，it attempts to improve its reading competence，so that tomorrow it can extract more facts from the web，more accurately．

## Semi supervised learning

At present，NELL has accumulated a knowledge base of 967,123 beliefs that it has read from various web pages．It is not perfect， but NELL is learning．You can track NELL＇s progress below or＠cmunell on Twitter，browse and download its knowledge base，read more about our technical approach，or join the discussion group．

Recently－Learned Facts twitter

| instance | iteration | date learned | confidence |
| :---: | :---: | :---: | :---: |
| robert＿trent＿jones＿sr is an Australian person | 473 | 27－dec－2011 |  |
| quality gift is a character trait | 475 | 29－dec－2011 | 99.5 倉 䇡 |
| confectioners＿sugar is a food | 473 | 27－dec－2011 |  |
| st＿petersburg＿times is a newspaper | 472 | 26－dec－2011 | 100.0 各 筞 |
| scott＿olynek is a Canadian person | 473 | 27－dec－2011 | 94.1 尔 |
| perth is a city that lies on the river Swan＿river | 472 | 26－dec－2011 | 99.2 为 |
| florida＿state＿university is a sports team also known as state＿university | 472 | 26－dec－2011 | 100.0 多 客 |
| press＿enterprise is a newspaper in the city riverside | 472 | 26－dec－2011 | 98.4 令 客 |

# Driveless cars 

Supervised and<br>reinforcement learning

## Helicopter control

Reinforcement learning

## Biology

ACGCTGAGCAATTCGATAGCAATTCGATAACGCTGAGCAACGCTGAGCAATTCGATAGCAATTC GATAACGCTGAGCAATCGGATAACGCTGAGCAATTCGATAGCAATTCGATAACGCTGAGCAACG CTGAGCAATTCGATAGCAATTCGATAACGCTGAGCAATCGGATATCGATAGCAATTCGATAAATC GGATAACGCTGAGCAATTCGATAGCAATTCGATAACGCTGAGCAACGCTGAGCAATTCGATAGC AATTCGATAACGCTGAGCAATCGGATATCGATAGCAATTCGATAACGCTGAGCAACGCTGAGCA ATTCGATAGCAATTCGATAACGCTGAGCAATCGGATAACGCTGAGCAATTCGATAGCATTCGAT AACGCTGAGCAA QUCTGAGCAATTCGATAGCAATTCGATAACGCTGAGCAATCGGATAACGCTG AGCAATTCGATAGCAATTCGATAACGCTGAGCTGAGCAATTCGATAGCAA GCAATCGGATAACGCTGAGC GATAGCAATTCGATAGCAA GATAGCAATTCGATAACGCTGAGCAACGCTGAGCAATTCGAT AGCAATTCGATAACGCTGA GAGCAACGCTGAGCAATTC -ATC GGATAACGCTGAGCAATTCGATAGCAATTCGATAACGCT CGATAACGCTGAGCAACGTGAGCAATTCGATAGCAATTCGATAACGCTGAGCAATCGGATAAC G( こGCTGAGCTGAGCAATTCGATAGCAATTCGATAACG ст Which partis the gene? zGATAGcaATTCGATAACGCTGAGCAACGCTGAGCA A 1 AATCGGATATCGATAGCAATTCGATAACGCTGAGCA ACGCTGAGCAATTCGATAGCAATTCGATAACGCTGAGCAATCGGATAACGCTGAGCAATTCGAT AGCATTCGATAACGCTGAGCAACGCTGAGCAATTCGATAGCAATTCGATCGGATAACGCTGAGC AATTCGATAGCAATTCGATAACGCTGAGCTGAGCAATTCGATAGCAATTCGATAACGCTGAGCA ATCGGATAACGCTGAGCAATTCGATAGCA AGCAATTCGATAACGCTGAGCAATCGGAT GCAATTCGATAGCAATTCGATAACGCTGA GATAACGCTGAGCAACGCTGAGCAATTCG CTGAGCAATTCGATAGCAATTCGATAACGI

## Supervised and

 unsupervised learning (can also use active learning) GAGCAATTCGAT GAGCAACGCTGA TCGATAGCATTC ンAATCGGATAACG -ATTCGATAACGC TGAGCAATCGGATAACGCTGAGCAATTCGATAGCAATTCGATAACGCTGAGCAACGCTGAGCAA TTCGATAGCAATTCGATAGCAATTCGATAGCAATTCGATAACGCTGAGCAACGCTGAGCAATTC GATAGCAATTCGATAACGCTGAGCAATCGGATAACGCTGAGCAATTCGATAGCAATTCGATAAC GCTGAGCAACGCTGAGCAATTCGATAGCAATTCGATAACGCTGAGCAATCGGATATCGATAGCA ATTCGATAACGCTGAGCAACGCTGAGCAATTCGATAGCAATTCGATAACGCTGAGCAATCGGAT AACGCTGAGCAATTCGATAGCAATTCGATAACGCTGAGCTGAGCAATTCGATAGCAATTCGATA ACGCTGAGCAATCGGA
## Common Themes

- Mathematical framework
- Well defined concepts based on explicit assumptions
- Representation
- How do we encode text? Images?
- Model selection
- Which model should we use? How complex should it be?
- Use of prior knowledge
- How do we encode our beliefs? How much can we assume?


## (brief) intro to probability

## Basic notations

- Random variable
- referring to an element / event whose status is unknown:

A = "it will rain tomorrow"

- Domain (usually denoted by $\Omega$ )
- The set of values a random variable can take:
- "A = The stock market will go up this year": Binary
- "A = Number of Steelers wins in 2012": Discrete
- "A = \% change in Google stock in 2012": Continuous


## Axioms of probability (Kolmogorov's axioms)

A variety of useful facts can be derived from just three axioms:

1. $0 \leq P(A) \leq 1$
2. $P($ true $)=1, P($ false $)=0$
3. $P(A \cup B)=P(A)+P(B)-P(A \cap B)$

There have been several other attempts to provide a foundation for probability theory. Kolmogorov's axioms are the most widely used.

## Priors

Degree of belief in an event in the absence of any other information

## No rain



P (rain tomorrow) $=0.2$
P (no rain tomorrow $)=0.8$

## Conditional probability

- $P(A=1 \mid B=1)$ : The fraction of cases where $A$ is true if $B$ is true

$$
\mathrm{P}(\mathrm{~A}=0.2)
$$

$$
P(A \mid B=0.5)
$$



## Conditional probability

- In some cases, given knowledge of one or more random variables we can improve upon our prior belief of another random variable
- For example:
$p($ slept in movie $)=0.5$
$p($ slept in movie | liked movie) $=1 / 4$
$p($ didn't sleep in movie | liked movie $)=3 / 4$

| Slept | Liked |
| :--- | :--- |
| 1 | 0 |
| 0 | 1 |
| 1 | 1 |
| 1 | 0 |
| 0 | 0 |
| 1 | 0 |
| 0 | 1 |
| 0 | 1 |

## Joint distributions

- The probability that a set of random variables will take a specific value is their joint distribution.
- Notation: $P(A \wedge B)$ or $P(A, B)$
- Example: P(liked movie, slept)

If we assume independence then

$$
P(A, B)=P(A) P(B)
$$

However, in many cases such an assumption maybe too strong (more later in the class)

## Joint distribution (cont)

Evaluation of classes

```
\(\mathrm{P}(\) class size \(>20)=0.6\)
\(\mathrm{P}(\) summer \()=0.4\)
\(\mathrm{P}(\) class size \(>20\), summer \()=\) ?
```

| Size | Time | Eval |
| :--- | :--- | :--- |
| 30 | R | 2 |
| 70 | R | 1 |
| 12 | S | 2 |
| 8 | S | 3 |
| 56 | R | 1 |
| 24 | S | 2 |
| 10 | R | 3 |
| 23 | R | 3 |
| 9 | R | 2 |
| 45 |  | 1 |

## Joint distribution (cont)

Evaluation of classes

```
\(\mathrm{P}(\) class size \(>20)=0.6\)
\(\mathrm{P}(\) summer \()=0.4\)
\(\mathrm{P}(\) class size \(>20\), summer \()=0.1\)
```

| Size | Time | Eval |
| :--- | :--- | :--- |
| 30 | R | 2 |
| 70 | R | 1 |
| 12 | S | 2 |
| 8 | S | 3 |
| 56 | R | 1 |
| 24 | S | 2 |
| 10 | R | 3 |
| 23 | R | 3 |
| 9 | R | 2 |
| 45 |  | 1 |

## Joint distribution (cont)

$\mathrm{P}($ class size $>20)=0.6$
$\mathrm{P}($ eval $=1)=0.3$
$P($ class size $>20$, eval $=1)=0.3$

| Size | Time | Eval |
| :--- | :--- | :--- |
| 30 | R | 2 |
| 70 | R | 1 |
| 12 | S | 2 |
| 8 | S | 3 |
| 56 | R | 1 |
| 24 | S | 2 |
| 10 | R | 3 |
| 23 | R | 3 |
| 9 | R | 2 |
| 45 |  | 1 |

## Joint distribution (cont)

Evaluation of classes
$\mathrm{P}($ class size $>20)=0.6$
$\mathrm{P}($ eval $=1)=0.3$
$P($ class size $>20$, eval $=1)=0.3$

| Size | Time | Eval |
| :--- | :--- | :--- |
| 30 | R | 2 |
| 70 | R | 1 |
| 12 | S | 2 |
| 8 | S | 3 |
| 56 | R | 1 |
| 24 | S | 2 |
| 10 | R | 3 |
| 23 | R | 3 |
| 9 | R | 2 |
| 45 |  | 1 |

## Chain rule

- The joint distribution can be specified in terms of conditional probability:

$$
P(A, B)=P(A \mid B)^{*} P(B)
$$

- Together with Bayes rule (which is actually derived from it) this is one of the most powerful rules in probabilistic reasening



## Bayes rule

- One of the most important rules for this class.
- Derived from the chain rule:

$$
P(A, B)=P(A \mid B) P(B)=P(B \mid A) P(A)
$$

- Thus,

$$
P(A \mid B)=\frac{P(B \mid A) P(A)}{P(B)}
$$



Thomas Bayes was an English clergyman who set out his theory of probability in 1764.

## Bayes rule (cont)

Often it would be useful to derive the rule a bit further:

$$
P(A \mid B)=\frac{P(B \mid A) P(A)}{P(B)}=\frac{P(B \mid A) P(A)}{\sum_{A} P(B \mid A) P(A)}
$$



$$
P(B, A=1)
$$

$P(B, A=0)$
This results from: $P(B)=\sum_{A} P(B, A)$


## Density estimation

## Density Estimation

- A Density Estimator learns a mapping from a set of attributes to a Probability



## Density estimation

- Estimate the distribution (or conditional distribution) of a random variable
- Types of variables:
- Binary
coin flip, alarm
- Discrete
dice, car model year
- Continuous
height, weight, temp.,


## When do we need to estimate densities?

- Density estimators can do many good things...
- Can sort the records by probability, and thus spot weird records (anomaly detection)
- Can do inference: P(E1|E2)

Medical diagnosis / Robot sensors

- Ingredient for Bayes networks and other types of ML methods


## Density estimation

- Binary and discrete variables:


## Easy: Just count!

- Continuous variables:


## Harder (but just a bit): Fit a model

# Learning a density estimator for discrete variables 

$$
\hat{P}\left(x_{i}=u\right)=\frac{\# \text { records in which } x_{i}=u}{\text { total number of records }}
$$

A trivial learning algorithm!

But why is this true?

## Maximum Likelihood Principle

We can define the likelihood of the data given the model as follows:
$\hat{P}($ dataset $\mid M)=\hat{P}\left(x_{1} \wedge x_{2} \ldots \wedge x_{n} \mid M\right)=\prod_{k=1}^{n} \hat{P}\left(x_{k} \mid M\right)$
For example M is

- The probability of 'head' for a coin flip
- The probabilities of observing 1,2,3,4 and 5 for a dice
- etc.


## Maximum Likelihood Principle

$$
\hat{P}(\text { dataset } \mid M)=\hat{P}\left(x_{1} \wedge x_{2} \cdots \wedge x_{n} \mid M\right)=\prod^{n} \hat{P}\left(x_{k} \mid M\right)
$$

- Our goal is to determine the values for the parameters in $M$
- We can do this by maximizing the probability of generating the observed samples
- For example, let $\Theta$ be the probabilities for a coin flip
- Then

$$
L\left(x_{1}, \ldots, x_{n} / \Theta\right)=p\left(x_{1} \mid \Theta\right) \ldots p\left(x_{n} \mid \Theta\right)
$$

- The observations (different flips) are assumed to be independent
- For such a coin flip with $P(H)=q$ the best assignment for $\Theta_{h}$ is

$$
\operatorname{argmax}_{q}=\# H / \# \text { samples }
$$

-Why?

## Maximum Likelihood Principle: Binary variables

- For a binary random variable $A$ with $P(A=1)=q$ $\operatorname{argmax}_{\mathrm{q}}=$ \#1/\#samples
- Why?

Data likelihood: $\quad P(D \mid M)=q^{n}(1-q)^{n_{2}}$
We would like to find: $\quad \arg _{\max }^{q} q^{n_{1}}(1-q)^{n_{2}}$

Omitting terms that
 do not depend on $q$

## Maximum Likelihood Principle

Data likelihood: $\quad P(D \mid M)=q^{m_{1}}(1-q)^{n_{2}}$
We would like to find: $\quad \arg \max _{q} q^{n}(1-q)^{n_{2}}$

$$
\begin{aligned}
& \frac{\partial}{\partial q} q^{n_{1}}(1-q)^{n_{2}}=n_{1} q^{n_{1}-1}(1-q)^{n_{2}}-q^{n_{1}} n_{2}(1-q)^{n_{2}-1} \\
& \frac{\partial}{\partial q}=0 \Rightarrow \\
& n_{1} q^{n_{1}-1}(1-q)^{n_{2}}-q^{n_{1}} n_{2}(1-q)^{n_{2}-1}=0 \Rightarrow \\
& q^{n_{1}-1}(1-q)^{n_{2}-1}\left(n_{1}(1-q)-q n_{2}\right)=0 \Rightarrow \\
& n_{1}(1-q)-q n_{2}=0 \Rightarrow \\
& n_{1}=n_{1} q+n_{2} q \Rightarrow \\
& q=\frac{n_{1}}{n_{1}+n_{2}}
\end{aligned}
$$

## Log Probabilities

When working with products, probabilities of entire datasets often get too small. A possible solution is to use the log of probabilities, often termed 'log likelihood'

$$
\log \hat{P}(\text { dataset } \mid M)=\log \prod_{k=1}^{n} \hat{P}\left(x_{k} \mid M\right)=\sum_{k=1}^{n} \log \hat{P}\left(x_{k} \mid M\right)
$$

Maximizing this likelihood function is the
same as maximizing P (dataset | M )


## Density estimation

- Binary and discrete variables:
- Continuolis variahles.



## How much do grad students sleep?

- Lets try to estimate the distribution of the time students spend sleeping (outside class).


## Possible statistics

- X

Sleep time
-Mean of X :

$$
\begin{aligned}
& E\{X\} \\
& 7.03
\end{aligned}
$$

- Variance of X :
$\operatorname{Var}\{X\}=E\left\{(X-E\{X\})^{\wedge} 2\right\}$ 3.05

Sleep


## Covariance: Sleep vs. GPA

-Co-Variance of X1,

## X2:

Covariance $\{X 1, X 2\}=$ $E\{(X 1-E\{X 1\})(X 2-E\{X 2\})\}$ $=0.88$


## Statistical Models

- Statistical models attempt to characterize properties of the population of interest
- For example, we might believe that repeated measurements follow a normal (Gaussian) distribution with some mean $\mu$ and

$$
\text { variance } \sigma^{2}, x \sim \mathrm{~N}\left(\mu, \sigma^{2}\right)
$$

$$
p(x \mid \Theta)=\frac{1}{\sqrt{2 \pi \sigma^{2}}} e^{\frac{-(x-\mu)^{2}}{2 \sigma^{2}}}
$$

and $\Theta=\left(\mu, \sigma^{2}\right)$ defines the parameters (mean and variance) of the model.

## The Parameters of Our Model

- A statistical model is a collection of distributions; the parameters specify individual



## The Parameters of Our Model

- A statistical model is a collection of distributions; the parameters specify individual distributions $\mathrm{x} \sim \mathrm{N}\left(\mu, \sigma^{2}\right)$
- We need to adjust the
parameters so that the resulting distribution fits the data well



## Computing the parameters of our model

- Lets assume a Guassian distribution for our sleep data
- How do we compute the parameters of the model?



## Maximum Likelihood Principle

- We can fit statistical models by maximizing the probability of generating the observed samples:
$L\left(x_{1}, \ldots, x_{n} \mid \Theta\right)=p\left(x_{1} \mid \Theta\right) \ldots p\left(x_{n} \mid \Theta\right)$
(the samples are assumed to be independent)
- In the Gaussian case we simply set the mean and the variance to the sample mean and the sample variance:

$$
\bar{\mu}=\frac{1}{n} \sum_{i=1}^{n} x_{i} \quad \overline{\sigma^{2}}=\frac{1}{n} \sum_{i=1}^{n}\left(x_{i}-\bar{\mu}^{2}\right.
$$

Why?

## Important points

- Random variables
- Chain rule
- Bayes rule
- Joint distribution, independence, conditional independence
- MLE


## Probability Density Function

- Discrete distributions
$\square \square \square \square \square \sum_{i} P\left(X=x_{i}\right)=1$
$\begin{array}{llllll}1 & 2 & 3 & 4 & 5 & 6\end{array}$
- Continuous: Cumulative Density Function (CDF): F(a)



## Cumulative Density Functions

- Total probability $\quad P(\Omega)=\int_{-\infty}^{\infty} f(x) d x=1$
- Probability Density Function (PDF)

$$
\frac{d}{d x} F(x)=f(x)
$$

- Properties:

$$
P(a \leq x \leq b)=\int_{b}^{a} f(x) d x=F(b)-F(a)
$$

$\lim _{x \rightarrow-\infty} F(x)=0$

$$
\lim _{x \rightarrow \infty} F(x)=1
$$

$F(a) \geq F(b) \forall a \geq b$

## Expectations

- Mean/Expected Value:

$$
E[x]=\bar{x}=\int x f(x) d x
$$

- Variance:

$$
\operatorname{Var}(x)=E\left[(x-\bar{x})^{2}\right]=E\left[x^{2}\right]-(\bar{x})^{2}
$$

- In general:

$$
\begin{aligned}
E\left[x^{2}\right] & =\int x^{2} f(x) d x \\
E[g(x)] & =\int g(x) f(x) d x
\end{aligned}
$$

## Multivariate

- Joint for ( $\mathrm{x}, \mathrm{y}$ )

$$
P((x, y) \in A)=\iint_{A} f(x, y) d x d y
$$

- Marginal:

$$
f(x)=\int f(x, y) d y
$$

- Conditionals:

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

- Chain rule:

$$
f(x, y)=f(x \mid y) f(y)=f(y \mid x) f(x)
$$

## Bayes Rule

- Standard form:

$$
f(x \mid y)=\frac{f(y \mid x) f(x)}{f(y)}
$$

- Replacing the bottom:

$$
f(x \mid y)=\frac{f(y \mid x) f(x)}{\int f(y \mid x) f(x) d x}
$$

## Binomial

- Distribution:

$$
x \sim \operatorname{Binomial}(p, n)
$$

$$
P(x=k)=\binom{n}{k} p^{k}(1-p)^{n-k}
$$

- Mean/Var:

$$
\begin{gathered}
E[x]=n p \\
\operatorname{Var}(x)=n p(1-p)
\end{gathered}
$$

## Uniform

- Anything is equally likely in the region $[a, b]$
- Distribution:

$$
x \sim U(a, b)
$$

- Mean/Var

$$
\begin{gathered}
f(x)= \begin{cases}\frac{1}{b-a} & a \leq x \leq b \\
0 & \text { otherwise }\end{cases} \\
E[x]=\frac{a+b}{2} \\
\operatorname{Var}(x)=\frac{a^{2}+a b+b^{2}}{3} \\
\frac{\mathrm{a}}{\mathrm{a}}
\end{gathered}
$$

## Gaussian (Normal)

- If I look at the height of women in country xx, it will look approximately Gaussian
- Small random noise errors, look Gaussian/Normal
- Distribution:

$$
x \sim N\left(\mu, \sigma^{2}\right) \quad f(x)=\frac{1}{\sqrt{2 \pi} \sigma} e^{\frac{-(x-\mu)^{2}}{2 \sigma^{2}}}
$$

- Mean/var

$$
\begin{gathered}
E[x]=\mu \\
\operatorname{Var}(x)=\sigma^{2}
\end{gathered}
$$



## Why Do People Use Gaussians

- Central Limit Theorem: (loosely)
- Sum of a large number of IID random variables is approximately Gaussian


## Multivariate Gaussians

- Distribution for vector x

$$
x=\left(x_{1}, \ldots, x_{N}\right)^{T}, \quad x \sim N(\mu, \Sigma)
$$

- PDF:

$$
f(x)=\frac{1}{(2 \pi)^{\frac{N}{2}}|\Sigma|^{\frac{1}{2}}} e^{-\frac{1}{2}(x-\mu)^{T} \Sigma^{-1}(x-\mu)}
$$

$$
E[x]=\mu=\left(E\left[x_{1}\right], \ldots, E\left[x_{N}\right]\right)^{T}
$$

$$
\operatorname{Var}(x) \rightarrow \Sigma=\left(\begin{array}{cccc}
\operatorname{Var}\left(x_{1}\right) & \operatorname{Cov}\left(x_{1}, x_{2}\right) & \ldots & \operatorname{Cov}\left(x_{1}, x_{N}\right) \\
\operatorname{Cov}\left(x_{2}, x_{1}\right) & \operatorname{Var}\left(x_{2}\right) & \ldots & \operatorname{Cov}\left(x_{2}, x_{N}\right) \\
\vdots & & \ddots & \vdots \\
\operatorname{Cov}\left(x_{N}, x_{1}\right) & \operatorname{Cov}\left(x_{N}, x_{2}\right) & \ldots & \operatorname{Var}\left(x_{N}\right)
\end{array}\right)
$$

## Multivariate Gaussians

$$
\begin{gathered}
f(x)=\frac{1}{(2 \pi)^{\frac{N}{2}}|\Sigma|^{\frac{1}{2}}} e^{-\frac{1}{2}(x-\mu)^{T} \Sigma^{-1}(x-\mu)} \\
E[x]=\mu=\left(E\left[x_{1}\right], \ldots, E\left[x_{N}\right]\right)^{T} \\
\operatorname{Var}(x) \rightarrow \Sigma=\left(\begin{array}{cccc}
\operatorname{Var}\left(x_{1}\right) & \operatorname{Cov}\left(x_{1}, x_{2}\right) & \ldots & \operatorname{Cov}\left(x_{1}, x_{N}\right) \\
\operatorname{Cov}\left(x_{2}, x_{1}\right) & \operatorname{Var}\left(x_{2}\right) & \ldots & \operatorname{Cov}\left(x_{2}, x_{N}\right) \\
\vdots & & \ddots & \vdots \\
\operatorname{Cov}\left(x_{N}, x_{1}\right) & \operatorname{Cov}\left(x_{N}, x_{2}\right) & \ldots & \operatorname{Var}\left(x_{N}\right)
\end{array}\right) \\
\operatorname{cov}\left(X_{1}, \boldsymbol{X}_{2}\right)=\frac{1}{n} \sum_{i=1}^{n}\left(x_{1, i}-\mu_{1}\right)\left(x_{2, i}-\mu_{2}\right)
\end{gathered}
$$

## Covariance examples



Covariance: -9.2

Correlated


Covariance: 18.33

Covariance: 0.6

## Sum of Gaussians

- The sum of two Gaussians is a Gaussian:

$$
\begin{aligned}
& x \sim N\left(\mu, \sigma^{2}\right) y \sim N\left(\mu_{y}, \sigma_{y}^{2}\right) \\
& a x+b \sim N\left(a \mu+b,(a \sigma)^{2}\right) \\
& x+y \sim N\left(\mu+\mu_{y}, \sigma^{2}+\sigma_{y}^{2}\right)
\end{aligned}
$$

