## Lectures 1\&2: Introduction and Course Overview

Prof. Krishna R. Pattipati BoT Distinguished Professor \& UTC Chair Professor in Systems Engineering Dept. of Electrical and Computer Engineering University of Connecticut

Contact: krishna.pattipati@uconn.edu; (860) 486-2890

> Fall 2018
> August 27 \& September 10, 2018

## Reading List

- Chapters $1 \& 2$ of Bishop
- Chapter 1 of Model-based Machine Learning by Winn and Bishop (Online)
- Chapters $1 \& 2$ of Murphy
- Chapter 2 of Theodoridis
- Class Notes
- Science:
> Special Issue on AI, July 17, 2015, Pages 248-278
> AI Transforms Science, July 7, 2017, Pages 16-30
- Web
$>$ https://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html
$>$ https://waitbutwhy.com/2015/01/artificial-intelligence-revolution-2.html
- YouTube (Bishop's talk on AI: The History and Future)
$>$ https://www.youtube.com/watch?v=8FHBh OmdsM


## Why Do Machines Need to Learn?

- Deluge of data (trillions of web pages, 10 years of video content in a single day on YouTube, retail sales, Facebook, twitter, blogs,...)
- Don't know Concise I/O or cause-effect relationships
- Are there hidden relationships and correlations ("Patterns") in large data?
- Need to adapt to an unknown environment ("Uncertainty")
- New knowledge is constantly being rediscovered ("Dynamic")
$\square$ Entering knowledge by humans is tedious



## 

## Machine Learning Applications

- Computational Biology
- Classifying protein sequences, identifying protein-coding genes, genetic networks from gene expression data, ...
- Medical Imaging
- Computer-aided diagnosis, image-guided therapy
- EEG classification, multi-model image fusion
- Information Retrieval/Natural Language Processing
- Text/audio/image retrieval
- Parsing/translation/text analysis/document classification/spam filters
- Robotics
- Autonomous driving, planning, control
- Speech processing
- Voice identification, speech recognition, speech $\rightarrow$ text $\rightarrow$ speech (in another language)
- Recommender Systems
- Collaborative filtering, market basket analysis,....
- Scientific Data Analysis
- E.g., NASA missions, Hadron Collider,...
- Financial Prediction and Automated Trading
- Computer Games, Gambling,....



## AI versus ML versus Deep Learning

Classic AI is based on deductive logic (Gen 1)

- Rules are based on human ingenuity

AI: The study of the Computations that make it possible to Perceive, Reason and Act
Machine Learning is based on learning "models" from data (Gen 2)

- Rules/Models derived from data are used for predictions and decisions

Deep learning seeks to "mimic" the biological brain (Gen 3)

- Learn and recognize patterns of input data using multiple layers of abstraction
- Applications:
- Image classification: Facial recognition
- Speech recognition in phones: Siri, Google assistant
- Game playing: Chess, Atari, Alpha GO
- Disease prediction
- Self-driving cars




## What Made Deep Learning Feasible?

Deluge of data and ability to crowd source the labeling process for supervisory learning ...Availability of large scale (and often pre-trained) labeled images (e.g., Alexnet, ImageNet, VGG16) and advances in cross-modal (image, text/data, video and audio) representations.... "Big Data"

- Availability of Graphics Processing Units (GPUs) that enable training of very large image/text/speech datasets from several weeks to a few days
- New nonlinearities (e.g., rectified linear units) and signal normalization that avoid numerical problems associated with gradient computations
- Convolution and max-pooling operations that exploit local connectivity to reduce the dimension of the weight space, control overfitting and make the network robust
- Concept of dropout to realize an exponentially large ensemble of networks from a single network
$\square$ Advances in stochastic optimization, adversarial training and regularization for robust network training
- More layers capture more invariances (Information-theoretic insights)
- Much better software tools (Keras, TensorFlow, Theano, Torch,...)
$\square$ Result:Features learned rather than hand-crafted, better accuracy (but easily fooled!)


## Types of Learning

$\square$ Learning from Data: Model Parameters, $W$ (weights) $W=f(D) \quad D \rightarrow$ data used to learn (training data) There are at least four forms of learning (others: active, transfer)
a) Supervised Learning (or learning with a teacher)
b) Unsupervised Learning
c) Semi-supervised Learning
d) Reinforcement Learning (like learning with a critic)
a) Supervised Learning:

$$
w=f\left[\left\{\underline{x}_{n}, \underline{z}_{n}\right\}_{n=1}^{N}\right] \text { (know inputs and correct outputs) }
$$

- Classification
- Regression
- Time series prediction
b) Unsupervised Learning:

$$
w=f\left[\left\{\underline{x}_{n}\right\}_{n=1}^{N}\right] \quad \text { (has input data only) }
$$

- Clustering; Density Estimation
- Outlier Detection
- Compression
c) Semi-supervised Learning: Practical setting; Active learning $w=\left\{f\left[\left\{\underline{x}_{n}, \underline{z}_{n}\right\}_{n=1}^{N_{1}}\right] ;\left\{f\left(\underline{x}_{n}\right)\right\}_{n=N_{1}+1}^{N}\right\}_{\left(N_{l} \ll N \text { : Few outputs are labeled, but most are not) }\right.}$
d) Reinforcement Learning: Useful in adaptive control and decision making $w=f\left[\left\{\underline{x}_{n}\right\}_{n=1}^{N}\right.$; correctness/effectiveness of outputs ] (learning with a critic)


## Things to Remember in ML Model Building

- Make Sure Training Data is Adequate ("Data Matters More than Algorithms")
- Make Sure Data is Representative ("No sampling Bias")
- Clean up the Data ("Detect and Fix/Remove Outliers", "Ignore/Estimate Missing Data", "Minimize Noise Effects in the Data")
- Feature Engineering ("Feature Selection", Feature Extraction", "Seek new Information")
- Avoid Overfitting and Underfitting ("Kiss" Principle, "Occam's Razor", "Theory of Parsimony", "Bias-Variance Tradeoff")
$\square$ Understand the difference between Training Error and Generalization (Out-of-sample) Error ("Split Data into Training, Validation and Test Data Sets", "No data snooping", "Test and Validate the Model", "Cross Validation", "Bootstrap", "Model Selection", "Bayesian Model Averaging")
- Exploit Domain Knowledge ("Relationships and Constraints", "No Free Lunch Theorem")


## Some Interesting Datasets

| Dataset | Accessible Location |
| :---: | :---: |
| IRIS | http://www.statlab.uni-heidelberg.de/data/iris/ |
| MNIST | http://yann.lecun.com/exdb/mnist/ |
| 20Newsgroups | http://cs.nyu.edu/~roweis/data.html |
| Olivetti Faces | http://cs.nyu.edu/~roweis/data.html |
| Alarm Network | http://www.bnlearn.com/bnrepository/ |
| UCI Machine Learning Repository: $\mathrm{http} / / /$ archive.ics.uci.edu/ml/index.php |  |
| Kaggle Data Sets: https://www.Kaggle.com/datasets |  |
| Amazon AWS: http://aws.amazon.com/fr/datasets |  |
| Open Source Data Portals: http://dataportals.org ; http://opendatamonitor.eu ; http://quandl.com; |  |
| Wiki: https://en.wikipedia.org/wiki/List of datasets for_machine_learning_research |  |
| Deep Learning Datasets: http://deeplearning.net/datasets/ |  |
| Blog: https://blog.algorithmia.com/machine-learning-datasets-for-data-scientists/ |  |



## Sample Data of Iris Dataset

| Sepal.Length | Sepal.Width | Petal.Length | Petal.Width | Species |
| ---: | ---: | ---: | ---: | :--- | :--- |
| 5.1 | 3.5 | 1.4 | 0.2 | setosa |
| 4.9 | 3 | 1.4 | 0.2 | setosa |
| 4.7 | 3.2 | 1.3 | 0.2 | setosa |
| 4.6 | 3.1 | 1.5 | 0.2 | setosa |
| 5 | 3.6 | 1.4 | 0.2 | setosa |
| 7 | 3.2 | 4.7 | 1.4 | versicolor |
| 6.4 | 3.2 | 4.5 | 1.5 | versicolor |
| 6.9 | 3.1 | 4.9 | 1.5 | versicolor |
| 5.5 | 2.3 | 4 | 1.3 | versicolor |
| 6.5 | 2.8 | 4.6 | 1.5 | versicolor |
| 6.3 | 3.3 | 6 | 2.5 | virginica |
| 5.8 | 2.7 | 5.1 | 1.9 | virginica |
| 7.1 | 3 | 5.9 | 2.1 | virginica |
| 6.3 | 2.9 | 5.6 | 1.8 | virginica |
| 6.5 | 3 | 5.8 | 2.2 | virginica |

Both petal length and width clearly distinguish setosa; but the other two are harder to distinguish

1. R. A. Fisher (1936). "The use of multiple measurements in taxonomic problems." Annals of Eugenics. 7 (2): 179-188.
2. Edgar Anderson (1936). "The species problem in Iris." Annals of the Missouri Botanical Garden. 23 (3): 457-509.

## Get to Know the Data

Continuous Features: For each feature, tabulate
> Count
> \% Missing
> Cardinality
> Minimum and Maximum
> $1^{\text {st }}$ and $3^{\text {rd }}$ Quartiles
> Mean and Median
> Variance and Standard Deviation
> Skewness
> Kurtosis

- Discrete (Categorical) Features
> Count
$>$ \% Missing
> Cardinality
> First and Second Mode, Mode Frequency and \% Mode Frequency


## Data Statistics: Mean

Data Matrix $X$ : an $N$ by $p$ matrix
> $N=$ Number of data points (data samples)
> $p=$ Number of features (e.g., attributes, measurements, variables)

- Central Tendencies
> Mean of each column feature (MATLAB: mean (X) )

$$
\underline{\hat{\mu}}=\frac{X^{T} \underline{e}}{N} ; \quad \underline{e} \sim N \text { column vector of } 1^{s}
$$

$>$ Iris Dataset:
Mean Removed Data Matrix
$X_{1}=X-\underline{e} \underline{\hat{\mu}}^{T}=\underbrace{\left(I-\frac{e e^{T}}{N}\right)}_{\text {Projection Matrix, } P} X$
Note $: P^{n}=P ; n=2,3, .$.

- Overall mean: $\mu^{T}=\left[\begin{array}{llll}5.8433 & 3.0573 & 3.7580 & 1.1993\end{array}\right]$
- Setosa: $\mu_{l}{ }^{T}=\left[\begin{array}{lllll}5.0060 & 3.4280 & 1.4620 & 0.246\end{array}\right]$
- Versicolor: $\mu_{2}{ }^{T}=\left[\begin{array}{lllll}5.9360 & 2.7700 & 4.2600 & 1.3260\end{array}\right]$
- Virginica: $\mu_{3}{ }^{T}=\left[\begin{array}{lllll}6.5880 & 2.9740 & 5.5520 & 2.0260\end{array}\right]$
> Setosa has smaller petal length and petal width. Easy to classify
> Versicolor has smaller petal width than Virginica; somewhat harder to distinguish between Versicolor and Virginica


## Data Statistics: Median

## Central Tendencies

$>$ Median (median (X))

- Sort the values in each column $j$ in increasing order: $x_{[1], j} \leq x_{[2], j} \leq \ldots \leq x_{[N], j}$
- If $N$ is odd

$$
M d_{j}=x_{\left[\frac{N+1}{2}\right], j}
$$

- If $N$ is even

$$
M d_{j}=\frac{1}{2}\left[x_{\left[\frac{N}{2}\right], j}+x_{\left[\frac{N}{2}+1\right], j}\right]
$$

> Iris Dataset

- Overall Median: $\quad\left[\begin{array}{llll}5.8000 & 3.0000 & 4.3500 & 1.3000\end{array}\right]$
- Setosa: $\quad\left[\begin{array}{lllll}5.0000 & 3.4000 & 1.5000 & 0.2000\end{array}\right]$
- Versicolor: $\quad\left[\begin{array}{lllll}5.9000 & 2.8000 & 4.3500 & 1.3000\end{array}\right]$
- Virginica: $\quad\left[\begin{array}{lllll}6.5000 & 3.0000 & 5.5500 & 2.0000]\end{array}\right.$
$>$ Similar conclusions as with mean


## Data Statistics: Mode

## Central Tendencies

$>$ Mode of each feature $\equiv$ most frequent value in each column

- Good for discrete features: (mode (X))
- For continuous features, do a histogram or density estimate first and calculate the peak of the estimate (histogram (X $(:, j))$ or ksdensity $(X(:, j)) ; j=1,2, . ., p)$

Histogram


Sepal Length of Setosa

Density


Sepal Length of Setosa

## Data Statistics: Variations -1

## Variations

Minimum: (min (X) )

$$
x_{j, \min }=\min _{1 \leq i \leq N}\left(x_{i j}\right) ; j=1,2, \ldots, p
$$

$>$ Maximum: $(\max (\mathrm{X}))$

$$
x_{j, \max }=\max _{1 \leq i \leq N}\left(x_{i j}\right) ; \quad j=1,2, . ., p
$$

$>$ Range: (range (X))

$$
r_{j}=x_{j, \max }-x_{j, \min } ; j=1,2, . ., p
$$

Sample Variance for iid samples: ( $\operatorname{Var}(\mathrm{X})$ )

$$
i i d=\text { independent and }
$$ identically distributed

Bernoulli: $\sigma_{j}^{2}=p(1-p)$

$$
P(x=1)=p
$$

Show: $E\left[\hat{\sigma}_{j}^{2}\right]=\sigma_{j}^{2}$
$\sigma_{j}^{2}=\operatorname{var}\left(x_{i j}\right)$

Why? You can also compute class-conditional variance

$$
\xrightarrow[\mathbf{o}]{\hat{\sigma}_{j}^{2}}=\frac{1}{N-1} \sum_{i=1}^{N}\left(x_{i j}-\hat{\mu}_{j}\right)^{2} ; \hat{\mu}_{j}=\frac{1}{N} \sum_{i=1}^{N} x_{i j} ; j=1,2, . ., p
$$

$>$ An unbiased estimator of the covariance matrix (class-conditioned)

Computed from data of each class

$$
\hat{\Sigma}=\frac{X_{1 c}^{T} X_{1 c}}{n_{c}-1}=\frac{X_{c}^{T}\left(I-\frac{e-e^{T}}{n_{c}}\right) X_{c}}{n_{c}-1}
$$

For common (class-independent) covariance matrix across $C$ classes, replace ( $N-1$ ) by ( $N-$ $C)$ for unbiasedness! See lectures 4-6.

## Data Statistics: Variations - 2

$>$ Correlation Coefficient Matrix $(\operatorname{corrcoef}(X))$

$$
C_{r}=[\operatorname{Diag}(\hat{\Sigma})]^{-1 / 2} \hat{\Sigma}[\operatorname{Diag}(\hat{\Sigma})]^{-1 / 2}
$$


http://www.statlab.uni-heidelberg.de/data/iris/

## $>$ Iris Dataset

$$
\hat{\Sigma}_{\text {versicolor }}=\left[\begin{array}{cccc}
0.2664 & 0.0852 & 0.1829 & 0.0558 \\
0.0852 & 0.0985 & 0.0827 & 0.0412 \\
0.1829 & 0.0827 & 0.2208 & 0.0731 \\
0.0558 & 0.0412 & 0.0731 & 0.0391
\end{array}\right] \quad C_{r, \text { versicolor }}=\left[\begin{array}{cccc}
1.0000 & 0.5259 & 0.7540 & 0.5465 \\
0.5259 & 1.0000 & 0.5605 & 0.6640 \\
0.7540 & 0.5605 & 1.0000 & 0.7867 \\
0.5465 & 0.6640 & 0.7867 & 1.0000
\end{array}\right]
$$

$>1^{\text {st }}$ and $3^{\text {rd }}$ Quartiles of feature $j$

- Sort data in increasing order: $x_{[1], j} \leq x_{[2], j} \leq \ldots \leq x_{[N], j}$
- The lower half of data: $x_{[1], j} \leq x_{[2], j} \leq \ldots \leq x_{\left[\frac{N-1}{2}\right], j}$ for $N$ odd; $x_{[1], j} \leq x_{[2]], j} \leq \ldots \leq x_{\left[\frac{N}{2}\right], j}$ for $N$ even
- The upper half of data: $x_{\left[\frac{N+1}{2}+1\right], j} \leq x_{\left[\frac{N+1}{2}+2\right], j} \leq \ldots . \leq x_{[N], j}$ for $N$ odd; $x_{\left[\frac{N}{2}+1\right], j} \leq x_{\left.\left[\frac{N}{2}+2\right]\right], j} \leq \ldots . \leq x_{[N], j}$ for $N$ even
- The first quartile $Q_{1, j}$ is the median of the lower half of the data. This means that about $25 \%$ of the numbers in the data set lie below $Q_{1, j}$ and about $75 \%$ lie above $Q_{1, j}$
- The third quartile, denoted $Q_{3, j}$, is the median of the upper half of the data set. This means that about $75 \%$ of the numbers in the data set lie below $Q_{3, j}$ and about $25 \%$ lie above $Q_{3, j}$.

$$
\begin{aligned}
& \text { For } N=9: Q_{1, j}=\frac{x_{[2], j}+x_{[3], j}}{2} ; Q_{3, j}=\frac{x_{[7], j}+x_{[8], j}}{2} \\
& \text { For } N=10: Q_{1, j}=x_{[3], j} ; Q_{3, j}=x_{[8], j}
\end{aligned}
$$

## Data Statistics: Skewness and Kurtosis

Bias Corrected Skewness (skewness $(X, 0)$ )

$$
s_{j}=\left[\frac{\sqrt{N(N-1)}}{(N-2)}\right]\left[\frac{\frac{1}{N} \sum_{i=1}^{N}\left(x_{i j}-\hat{\mu}_{j}\right)^{3}}{\left(\frac{1}{N} \sum_{i=1}^{N}\left(x_{i j}-\hat{\mu}_{j}\right)^{2}\right)^{3 / 2}}\right] ; \hat{\mu}_{j}=\frac{1}{N} \sum_{i=1}^{N} x_{i j} ; j=1,2, . ., p ; N \geq 3
$$

$>s_{j}<0 \Rightarrow$ left tail is longer; $s_{j}>0 \Rightarrow$ right tail is longer; $s_{j}=0 \Rightarrow$ symmetric
$>$ Skewness for versicolor: [[0.1022 -0.3519 -0.5882 -0.0302$]$
$>$ Returns from financial stocks have larger-than-Normal skew and kurtosis

- Bias Corrected Kurtosis (kurtosis(X,0)) ... useful for detecting outliers

$$
\kappa_{j}=\left[\frac{(N-1)}{(N-2)(N-3)}\right]\left[(N+1)\left(\frac{\frac{1}{N} \sum_{i=1}^{N}\left(x_{i j}-\hat{\mu}_{j}\right)^{4}}{\left(\frac{1}{N} \sum_{i=1}^{N}\left(x_{i j}-\hat{\mu}_{j}\right)^{2}\right)^{2}}\right)-3(N-1)\right]+3 ; \hat{\mu}_{j}=\frac{1}{N} \sum_{i=1}^{N} x_{i j} ; j=1,2, . ., p ; N \geq 4
$$

$>\kappa_{j}=3 \Rightarrow$ Gaussian and Mesokurtic distributions; Excess Kurtosis $=\kappa_{j}-3$
$>\kappa_{j}>3 \Rightarrow$ Leptokurtic or super-Gaussian distributions (e.g., Logistic, Laplace, Student-t,...)
$>\kappa_{j}<3 \Rightarrow$ Platykurtic or sub-Gaussian distributions (e.g., Bernoulli with $p=1 / 2$, uniform)



## Data Scaling Methods - 1

10 For each Feature

1. Zero mean and unit variance

$$
x_{i j} \rightarrow s_{i j}=s_{i j}\left(x_{i j}\right)=\frac{\left(x_{i j}-\hat{\mu}_{j}\right)}{\hat{\sigma}_{j}}
$$

In Matrix form : $\left(X-\underline{e} \underline{\mu}^{T}\right) \operatorname{Diag}\left(1 / \hat{\sigma}_{j}\right)$
2. Scale it to $[0,1]$ or $[-1,1]$
(Often times, $[\varepsilon, 1-\varepsilon]$ or $[-1+\varepsilon, 1-\varepsilon]) \varepsilon \approx 0.15 \sim 0.20$ to avoid the saturation of nonlinearities used in neural networks (e.g., sigmoid and tanh))

$$
x_{i j} \rightarrow s_{i j}=s_{i j}\left(x_{i j}\right)=(1-2 \varepsilon) \frac{\left(x_{i j}-x_{\min , j}\right)}{\left(x_{\max , j}-x_{\min , j}\right)}+\varepsilon \Rightarrow[\varepsilon, 1-\varepsilon]
$$

$$
x_{i j} \rightarrow s_{i j}=s_{i j}\left(x_{i j}\right)=2(1-\varepsilon) \frac{\left(x_{i j}-x_{\min , j}\right)}{\left(x_{\max , j}-x_{\min , j}\right)}+(\varepsilon-1) \quad \Rightarrow[\varepsilon-1,1-\varepsilon]
$$

## Data Scaling Methods - 2

3. If small values are disproportionately close together and large values are spread out, use log transformations to obtain uniformized $x^{s}$

$$
x_{i j} \rightarrow s_{i j}=s_{i j}\left(x_{i j}\right)=\ln \left[\frac{\left(x_{i j}-x_{\min , j}\right)}{\left(x_{\max , j}-x_{\min , j}\right)}+1\right] / \ln 2
$$

4. If large values are close together and small values are further apart, use exponential transformations (spreads out large values and pushes small values together)

$$
x_{i j} \rightarrow s_{i j}=s_{i j}\left(x_{i j}\right)=\frac{1-e^{-\left[\frac{\left(x_{i j}-x_{\min , j}\right)}{\left(x_{\text {max }, j}-x_{\min , j}\right)}\right]}}{1-e^{-1}}
$$

- Coding Classes (Targets): Coding [1,2,..,C] is okay for some (e.g., decision trees)

$$
\underline{z}_{k}=\underline{e}_{k}=\left[\begin{array}{c}
0 \\
\vdots \\
1 \\
\vdots \\
0
\end{array}\right] \Rightarrow z=k \quad \text { or } \quad \underline{z}_{k}=\left[\begin{array}{c}
-1 /(C-1) \\
\vdots \\
1 \\
\vdots \\
-1 /(C-1)
\end{array}\right] \quad \text { or } \quad \underline{z}_{k}=\left[\begin{array}{c}
-1 \\
\vdots \\
1 \\
\vdots \\
-1
\end{array}\right] \text { Good for }
$$

## Data Transformation/ Reduction/Cleaning

- Principal Component Analysis (PCA) is the best known and oldest technique for data reduction (Pearson, 1901)
- Singular Value Decomposition (SVD) if done on data directly
- Data compression and feature (indicator) selection
- Data Matrix $X$ : an $N^{\prime}$ by $p$ matrix
> $N^{\prime}=$ Number of data points (data samples)
$>p=$ Number of features (e.g., attributes, measurements)
- Remove $N a N$ (Not a Number) elements from the matrix $X$
> Use MATLAB function: $\mathrm{X}(\operatorname{any}(\mathrm{isnan}(\mathrm{X}), 2),:$ ) = [];
> $N=$ Number of data points after removing $N a N$ elements
> New $X$ matrix: an $N$ by $p$ matrix
- Compute the mean of each column feature

$$
\underline{\hat{\mu}}=\frac{X^{T} \underline{e}}{N} ; \quad \underline{e} \sim N \text { column vector of } 1^{s}
$$

## Principal Component Analysis Explained

- Remove mean from each column and form a new data matrix $X_{1}$

$$
X_{1}=X-\underline{e} \underline{\hat{\mu}}^{T}=\left(I-\frac{\underline{e} \underline{e}^{T}}{N}\right) X
$$

- Compute the Covariance Matrix

$$
P=\frac{X_{1}^{T} X_{1}}{N-1}=\frac{1}{N-1}\left[X^{T} X-N \underline{\hat{\mu}} \underline{\hat{\mu}}^{T}\right]=\frac{1}{N-1}\left[\sum_{n=1}^{N}\left(\underline{x}_{n} \underline{x}_{n}^{T}-\underline{\hat{\mu}}_{\hat{\mu}} \underline{\hat{\mu}}^{T}\right)\right] ; \text { a } p \text { by } p \text { matrix }
$$

- Perform Eigen decomposition of $P$ and select the Eigen vectors corresponding to the top $k$ eigen values $\ni$

$$
\begin{equation*}
\frac{\sum_{i=1}^{k} \lambda_{i}}{\operatorname{trace}(P)}>T h \tag{0.95}
\end{equation*}
$$

$$
\begin{aligned}
& P=V \Lambda V^{T} \\
& \lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{k} \geq \lambda_{k+1} \geq \cdots \geq \lambda_{p}
\end{aligned}
$$

- Reduced Data Matrix
$X_{2}=X_{1} V_{k} \quad$ an $N \times k$ matrix
$V_{k}=$ first $k$ columns of $V$; a $p \times k$ matrix
- Residual Matrix and Covariance

$$
R=X_{1}\left[I_{p}-V_{k} V_{k}^{T}\right] \quad \text { an } N \times p \text { matrix ; } \Sigma_{r}=\left[I_{p}-V_{k} V_{k}^{T}\right] P\left[I_{p}-V_{k} V_{k}^{T}\right] \text { a } p \times p \text { matrix }
$$

## Iris Data Visualization using PCA

First principal component explains $92.46 \%$ of variability in the data $\square$ First 2 principal components explain $97.8 \%$ of variability in the data


- Setosa is easy to distinguish; some ambiguity between Versicolor \& Virginica


## Machine Learning Tasks

Machine Learning seeks to solve many generic tasks of interest.

1. Classification
2. Regression / Interpolation
3. Density Estimation
4. Forecasting / Prediction ... basically regression
5. Optimization under uncertainty
6. Adaptive Control
7. Adaptive Communications

## Problem 1: Pattern Classification

## - Pattern Classification

* Pattern: Pattern means something exhibiting certain regularities, something able to serve as a model, something representing the concept of what was observed. A pattern is a collection of observations connected in time or space or both.
* Classification: Assign an input pattern (e.g., speech waveform, handwritten character, etc.) represented by a feature vector $\underline{x}=\left[\begin{array}{llll}x_{1} & x_{2} & \ldots & x_{p}\end{array}\right]^{T}$ to one or more specified classes or concepts


Pattern classification is the process of inferring $z$ from $\underline{x}: \underline{x} \rightarrow z$
$\underline{x}$ can be documents, images, DNA sequences, or graphs,....
$z$ can be classes, ranks, real values, part of speech tagging, graph,....

## Pattern Classification Process

- Typical Process of Pattern Classification



FIGURE 1.1. The objects to be classified are first sensed by a transducer (camera), whose signals are preprocessed. Next the features are extracted and finally the classification is emitted, here either "salmon" or "sea bass." Although the information flow is often chosen to be from the source to the classifier, some systems employ information flow in which earlier levels of processing can be altered based on the tentative or preliminary response in later levels (blue arrows). Yet others combine two or more stages into a unified step, such as simultaneous segmentation and feature extraction.

## Classification using Width and Lightness




FIGURE 1.4. The two features of lightness and width for sea bass and salmon. The dark line could serve as a decision boundary of our classifier. Overall classification error on the data shown is lower than if we use only one feature as in Fig. 1.3, but there will still be some errors. From: Richard O. Duda, Peter E. Hart, and David G. Stork, Pattern Classification. Copyright ©c) 2001 by John Wiley \& Sons, Inc.

- Linear (or) Hyperplane Classifier

From Duda, Hart and Stork

- Simple, but not necessarily the best for this problem


## How to Select the Right Decision Boundary?



FIGURE 1.5. Overly complex models for the fish will lead to decision boundaries that are complicated. While such a decision may lead to perfect classification of our training samples, it would lead to poor performance on future patterns. The novel test point marked ? is evidently most likely a salmon, whereas the complex decision boundary shown leads it to be classified as a sea bass. From: Richard O. Duda, Peter E. Hart, and David G. Stork, Pattern Classification. Copyright © 2001 by John Wiley \& Sons, Inc.

- Tradeoff between overtraining and generalizability (accuracy in use)
- Complex classifiers are not necessarily the best ("Occam's Razor")


## How to Select the Right Decision Boundary?




## Model Selection Problem

- Akaike Information Criterion (AIC)
- Bayesian Information Criterion (BIC)
- Minimum Description Length (MDL)
- Cross validation
- Regularization

FIGURE 1.6. The decision boundary shown might represent the optimal tradeoff between performance on the training set and simplicity of classifier, thereby giving the highest accuracy on new patterns. From: Richard O. Duda, Peter E. Hart, and David G. Stork, Pattern Classification. Copyright © 2001 by John Wiley \& Sons, Inc.

- Tradeoff between overtraining and generalizability
- How to find the "right" tradeoff?


## IRIS Classification Results using MATLAB APP

| Classifier | \% Accuracy <br> with 4 <br> Features | \% Accuracy <br> with 2 PCA- <br> Reduced <br> Features | \% Accuracy <br> with 8 <br> Features* | \%Accuracy <br> with 3 PCA- <br> reduced <br> Features* |
| :--- | :---: | :---: | :---: | :---: |
| Decision Tree | $94.7 \%$ | $93.3 \%$ | $96 \%$ | $92.7 \%$ |
| Weighted <br> kNN | $96.0 \%$ | $96.7 \%$ | $95.3 \%$ | $97.3 \%$ |
| Linear SVM | $96.0 \%$ | $96 \%$ | $94.7 \%$ | $98 \%$ |
| Medium <br> Gaussian <br> SVM | $96.7 \%$ | $97.3 \%$ | $95.3 \%$ | $96.7 \%$ |
| Bagged Trees | $95.3 \%$ | $92 \%$ | $96.7 \%$ | $97.3 \%$ |

- Add sepal and petal areas and ratios as 4 additional features


## Problem 2: Clustering/Categorization - 1

- Clustering / Categorization -- also known as unsupervised classification
- Explore similarity between the patterns and place similar things in a group

- Derive a statistical model of the process

- Density estimation



Is there a structure in the data?

## Problem 2: Clustering/Categorization - 2

a K-means clustering of height and weight data


Run kmeansHeightWeight in pmtk3-master\demos of Murphy to get these



- How to select K? Model selection problem again!



## Problem 3: Function Approximation - 1

- Function Approximation (Regression)

Given $\left(x_{1}, z_{1}\right),\left(x_{2}, z_{2}\right) \ldots,\left(x_{N}, z_{N}\right), N$ data points
What is $z \approx f(x)$ ?

- Very low order models are not good
$\Rightarrow$ too much bias
- Very high order models are not good either $\Rightarrow$ too much variability
"BIAS-VARIANCE DILEMMA"
- How do we select the "right" model?? AIC, BIC, MDL, CV, ....
- Applications of function approximation
- Regression
- System Identification/Parameter Estimation
- Neuro-dynamic Programming (NDP)/Approximate DP (ADP)
- What does $\approx$ mean? $\Rightarrow$ Criterion for optimization of weights

MSE: $\frac{1}{N} \sum_{i=1}^{N}\left[z_{i}-f\left(x_{i}\right)\right]^{2}$

This alone is not adequate! Why?

## Problem 3: Function Approximation - 2

$\square$ What is the best polynomial for a 21 point noisy data?

linear

quadratic

degree 10

degree 20

Run linregPolyVsDegree in pmtk3-masterldemos of Murphy to get these



Model selection is a difficult problem!

## J $\vdots$ $\vdots$ $\vdots$ $\vdots$ $\vdots$

## Problem 4: Estimation \& Prediction - 1

$\square$ Prediction/Forecasting (Regression)
Given $\left\{z\left(t_{1}\right), z\left(t_{2}\right), \ldots \ldots \ldots \ldots \ldots z\left(t_{N}\right)\right\}$, predict $\hat{z}\left(t_{N+m}\right)$

- Areas: Estimation theory, Adaptive Filtering
- Typical Applications
- Stock market prediction
- Weather forecasting
- Prognostics for system health management $\begin{array}{llllll}t_{1} & t_{2} & & t_{N} & t_{N+m} & t\end{array}$
- Ad selection systems
- Medicine
- Fraud detection
- Power Demand Prediction


## Problem 4: Estimation \& Prediction - 2

- How to detect an occasionally dishonest Casino? "Classification"


Run casinoDemo in pmtk3-masterldemos of Murphy to get these


- Learn the model based on data \& use it to forecast future behavior of Casino
- Optimization
$" \min f(\underline{x})$ s.t. $\underline{x} \in \Omega "$

$$
\Omega\left\{\begin{array}{cl}
h(\underline{x})=\underline{0} & \\
g(\underline{x}) \leq \underline{0} & \text { Non-Convex } \\
\underline{x}_{L B} \leq \underline{x} \leq \underline{x}_{U B} & \text { Non-Differentiable } \\
x_{i} \text { integer } & \\
\hline
\end{array}\right.
$$

Run surfaceFitDemo in pmtk3-masterldemos of Murphy to get these



Temperature as a function of $(x, y)$ location in a room

## Problem 6: Adaptive Control

- Adaptive (Neuro) Control (Regression, Classification and Optimization)


Design of a controller when have uncertain info. about the system


Lots of tuning required to make PID work

## Problem 7: Adaptive Communications

- Deep Neural Networks for Physical Layer Design*

$$
s \bullet \text { Transmitter } \xrightarrow{\mathbf{X}} \text { Channel } \xrightarrow{\mathbf{y}} \text { Receiver } \xrightarrow{\longrightarrow} \hat{s}
$$


*T. O'Shea and J. Hoydis, "An Introduction to Deep Learning for the Physical Layer," IEEE Trans. on Cognitive Communications and Networking, Vol. 3, No. 4, pp. 563-575, Dec 2017.

## Model Selection Problem is Ubiquitous

- How many clusters in the Data?
$\square$ What is the order of the system?
What is the intrinsic dimension of the system?
- Is this feature relevant?
- How many states in a Hidden Markov Model (HMM)?
- What is the structure of a graphical model?
- How many independent sources in the input?

Bayesian Methods provide a Principled Approach to Model Selection

## Bayes' Theorem and Inference

$\square$ Basic Axioms of Probability

- Probability of event $A, P(A) \in[0,1]$
- $\quad P(A)=1 \Leftrightarrow \mathrm{~A}$ is certain
- Sum Rule (Total Probability Theorem) $P(A)=\sum_{B} P(A, B)$
- Product Rule
$P(A, B)=P(B \mid A) P(A)=P(A \mid B) P(B)$
$-\quad P(A \cup B)=P(A$ or $B)=P(A)+P(B)-P(A B)$
- Bayes' theorem
- $P(A B)=P(A \mid B) P(B)=P(B \mid A) P(A)$
- $\quad P(A \mid B)=\frac{P(B \mid A) P(A)}{P(B)}$

We are inferring A given data B

- Begin with a priori probability $P(A)$ for our belief about $A$
- Observe B
- Bayes' theorem provides the revised belief about $A$, that is, the posterior probability $P(A \mid B)$
- Likelihood of $A$ : The quantity $P(B \mid A)$, as a function of varying $A$ for a fixed $B$
- posterior $\propto$ prior $\times$ likelihood $\Rightarrow P(A \mid B) \propto P(A) . P(B \mid A)$

| $\bullet P \sim$ Probability |
| :--- |
| $\bullet$ |
|  |

## Bayes' Theorem and Learning

$\square$ Data $D$; Model Parameters, $\underline{w}$; model set $M \in\{1,2, . ., m, \ldots, \mathrm{M}\}$

$$
p(\underline{w} \mid D, m)=\frac{p(D \mid \underline{w}, m) p(\underline{w} \mid m)}{p(D \mid m)}\left[\begin{array}{l}
p(D \mid \underline{w}, m)=\text { Likelihood of } \underline{w} \text { under model } m \\
p(\underline{w} \mid m)=\text { prior of } \underline{w} \text { under model } m \\
p(\underline{w} \mid D, m)=\text { posterior of } \underline{w} \text { given data } D \text { under model } m
\end{array}\right.
$$

- Prediction

$$
p(\underline{x} \mid D, m)=\int_{\underline{w}} p(\underline{x}, \underline{w} \mid D, m) d \underline{w}=\int_{\underline{w}} p(\underline{x} \mid \underline{w}, m) p(\underline{w} \mid D, m) d \underline{w}
$$

$\square$ Model Comparison

$$
P(m \mid D)=\frac{p(D \mid m) P(m)}{\sum_{k=1}^{M} p(D \mid k) P(k)} ; p(D \mid m)=\int_{\underline{w}} p(D \mid \underline{w}, m) p(\underline{w} \mid m) d \underline{w} \sim \text { Model Evidence }
$$

Interpretations of Model Evidence:

- Pdf that randomly selected model parameters from the prior generate $D$
- $-\log _{2} p(D \mid m)$ is the number of bits of surprise at observing $D$ under model $m$
- Bayesian model averaging .... Computationally intractable in general


## Bayesian Recursive Learning and Inference

Learning and Inference are not different in a Bayesian approach

- Recursion

$$
p(\underline{w} \mid D, \underline{x}, m)=\frac{p(\underline{w}, D, \underline{x} \mid m)}{p(D, \underline{x} \mid m)}=\frac{p(\underline{x} \mid \underline{w}, m) p(\underline{w} \mid D, m)}{\int p(\underline{x} \mid \underline{w}, m) p(\underline{w} \mid D, m) d \underline{w}}
$$

- Provides a principled method for Model Selection
- Very Simple model structures are unlikely to generate observed data ("large bias")
- Very Complex model structures are unlikely to generate the particular data set observed because they can generate many possible data sets ("large variance")
- The Right model is one with large model evidence for the observed data set D.... Large $p(D \mid m)$



## Graphical Modeling of Statistical Dependencies - 1

$\square$ Example 1 (Common Cause, Fork)


When $A$ is not observed, $B \& C$ are dependent.
When A is observed, B \&C are conditionally independent!

- Naïve Bayes
- iid observations
- prediction (sufficient statistics)
"tail-to-tail" dependency

$$
\begin{array}{|l}
B \perp C \mid A \\
" B \text { is independent of } C \text { given } A " \\
\text { Why? } \\
P(B, C)=\sum_{A} P(B \mid A) P(C \mid A) P(A) \neq P(B) P(C) \\
P(B, C \mid A)=\frac{P(A, B, C)}{P(A)}=P(B \mid A) P(C \mid A)
\end{array}
$$

Examples: 1. Smoking (A) causes Bronchitis (B) and Cancer (C).
2. Learning the probability of heads $P(A=1)$, given the outcomes of a number of coin tosses (in this case, two ( $\mathrm{B} \& \mathrm{C}$ ))

## Graphical Modeling of Statistical Dependencies - 2

- Example 2 (Indirect Cause, Chain)

"head-to-tail" dependency

$$
\begin{array}{ll}
P(A B C)=P(B) P(A \mid B) P(C \mid A) & \begin{array}{l}
\{A B\},\{C A\} \text { "cliques"" } \\
A \text { "separator" } \\
P(A B), P(C A) \text { are "clique potentials" } \\
P(A) \text { "separator potential"" }
\end{array} \\
\text { Also, } P(A B C)=\frac{P(A B) P(C A)}{P(A)} & \begin{array}{l}
\text { "ser }
\end{array}
\end{array}
$$

When $A$ is not observed, $B \& C$ are dependent.
When $A$ is observed, $B \& C$ are conditionally independent!
$B \perp C \mid A$
" $B$ is independent of $C$ given $A$ "
Why?
$P(B, C)=P(B) \sum_{A} P(A \mid B) P(C \mid A) \neq P(B) P(C)$
$P(B, C \mid A)=\frac{P(A, B, C)}{P(A)}=\frac{P(A \mid B) P(B) P(C \mid A)}{P(A)}=P(B \mid A) P(C \mid A)$

Example: B: Clouds; A: Rain; C: Grass is wet

## Graphical Modeling of Statistical Dependencies - 3

Example 3 (Common Effect, Collider)


> | $B \perp C \mid \varnothing$ |
| :--- |
| $" B$ is independent of $C$ given no evidence" |
| Why? |
| $P(B, C)=\sum_{A} P(B) P(C) P(A \mid B, C)=P(B) P(C)$ |

When $A$ is not observed, $A$ blocks the path from $B$ to $C$. However, when $A$ is observed, it unblocks the path from $B$ to $C \Rightarrow$ they become dependent

They are not independent given $A$

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

Example: $A$ : Engine Does Not Start; B: Battery; $C$ : Fuel

## Noisy ("Soft") XOR

| $\boldsymbol{B}$ | $\boldsymbol{C}$ | $P(A=1 \mid \dot{B}, C)$ | $P\left(A=0 \mid \boldsymbol{B}_{2} C\right)$ |
| :---: | :---: | :---: | :---: |
| 0 | 0 | 0.10 | 0.90 |
| 0 | 1 | 0.99 | 0.01 |
| 1 | 0 | 0.80 | 0.20 |
| 1 | 1 | 0.25 | 0.75 |

$$
\begin{aligned}
& \mathrm{P}(\mathrm{~B}=1)=0.65 \\
& \mathrm{P}(\mathrm{C}=1)=0.77
\end{aligned}
$$

Suppose we observed $A=0$. What is $P(B=1 \mid A=0)$ ?

$$
\begin{aligned}
P(B=1 \mid A=0) & =\frac{P(A=0 \mid B=1) P(B=1)}{P(A=0)}=\frac{P(A=0 \mid B=1) P(B=1)}{P(A=0 \mid B=1) P(B=1)+P(A=0 \mid B=0) P(B=0)} \\
P(A=0 \mid B=1) & =\sum_{C=0}^{1} P(A=0, C \mid B=1)=\sum_{C=0}^{1} P(A=0 \mid B=1, C) P(C \mid B=1) \\
& =\sum_{C=0}^{1} P(A=0 \mid B=1, C) P(C)=0.20 * 0.23+0.75 * 0.77=0.6235
\end{aligned}
$$

Similarly, $P(A=0 \mid B=0)=\sum_{C=0}^{1} P(A=0 \mid B=0, C) P(C)=0.90 * 0.23+0.01 * 0.77=0.2147$
$P(B=1 \mid A=0)=\frac{0.6235 * 0.65}{0.6235 * 0.65+0.2147 * 0.35}=0.8436$
$P(B=0 \mid A=0)=1-P(B=1 \mid A=0)=0.1564$

## Exploiting Dependency Structure: Graphical Models- 1

$\square$ Chain rule (general dependency)
$p\left(x_{1}, x_{2}, x_{3}, \ldots, x_{p-1}, x_{p}\right)=p\left(x_{1}\right) p\left(x_{2} \mid x_{1}\right) p\left(x_{3} \mid x_{1}, x_{2}\right) \ldots p\left(x_{p} \mid x_{1}, x_{2}, \ldots, x_{p-1}\right)$

- Storage complexity $O\left(K^{p}\right)$, assuming each variable has $K$ states
$\square$ Conditional Independence
- Naïve Bayes

$$
p(\underline{x}, z)=p(z)\left(\prod_{i=1}^{p} p\left(x_{i} \mid z\right)\right)
$$

- Storage complexity $[(K-1)(p C+1)]$

$$
\begin{aligned}
& x_{i} \in\{1,2, \ldots, K\} \\
& z \in\{1,2, \ldots, C\}
\end{aligned}
$$



- Classifiers based on Naïve Bayes can work quite well


## Exploiting Dependency Structure：Graphical Models－ 2

$\square$ Conditional Independence（continued）
－Markov chain（used to model sequential first order dependencies）


$$
p\left(x_{1: T}\right)=p\left(x_{1}\right) \prod_{t=2}^{T} p\left(x_{t} \mid x_{t-1}\right)
$$

－Homogeneous：$K^{2}$
－Non－homogeneous：$T K^{2}$

## Applications：

－Sentence completion
－Data compression
－Text classification
－Automatic essay writing
－Google＇s page rank
$-n^{\text {th }}$ order Markov chain（（ $\left.\mathbf{n}+\mathbf{1}\right)$－gram models）

$$
p\left(x_{1: T}\right)=p\left(x_{1}, x_{2}, . ., x_{n}\right) \prod_{t=n+1}^{T} p\left(x_{t} \mid x_{t-1}, x_{t-2}, \ldots, x_{t-n}\right) ; T>n
$$

 trigram

## Exploiting Dependency Structure: Graphical Models - 3

$\square$ Conditional Independence (continued)

- Hidden Markov model (HMM)

$p\left(z_{1: T}, x_{1: T}\right)=p\left(z_{1: T}\right) p\left(x_{1: T} \mid z_{i: T}\right)=p\left(z_{1}\right) \cdot\left(\prod_{t=2}^{T} p\left(z_{t} \mid z_{t-1}\right)\right) \cdot\left(\prod_{t=1}^{T} p\left(x_{t} \mid z_{t}\right)\right)$
- Applications: speech recognition, activity recognition, gene finding, protein sequence alignment, diagnosis,...
- Algorithms: Forwards-backwards algorithms, Viterbi decoding, Baum-Welch algorithm for learning HMM parameters from data
- Generalizations: Higher order, Factorial, Coupled, semi-Markov, Hierarchical, I/O, Auto-regressive, Buried, state space, Dynamic Bayesian networks (DBNs),...


## Three Broad Classes of Graphical Models - 1

$\square$ Bayesian Networks

- Represented in terms of directed acyclic graphs (DAGs)

Source code $\xrightarrow{\underline{u}}$ States of encoder

$$
p(\underline{u}, s, \underline{v}, \underline{y})=p(\underline{u}) p(s \mid \underline{u}) p(\underline{v} \mid s, \underline{u}) p(\underline{y} \mid \underline{v})
$$

## Encoded signal $\underset{\sim}{\boldsymbol{y}} \longrightarrow$ (V) Received signal

- In general,

$$
\begin{aligned}
& z=\left[z_{1} z_{2} \cdots \cdots z_{N}\right] \\
& P\left(z_{k} \mid a_{k}\right) \quad a_{k}=\text { parents of } z_{k}=p a\left(z_{k}\right) \\
& P(z)=\prod_{k=1}^{N} P\left(z_{k} \mid p a\left(z_{k}\right)\right)=\prod_{k=1}^{N} P\left(z_{k} \mid a_{k}\right)
\end{aligned}
$$

- Exploit graph properties
- d-separation, moralization, cliques,...
- Algorithms
- Variable elimination
- Junction tree
- Variational inference
- Belief propagation (BP)
- Expectation propagation (EP)
- Markov Chain Monte Carlo


## Three Broad Classes of Graphical Models - 2

$\square$ Markov random fields (MRF) and conditional MRFs...
Undirected graphical model (UGM) or Markov network

$\square$ Factor Graphs

- Useful when dealing with large graphs
- Both BNs and UGMs can be converted into factor graphs


$$
p(\underline{u}, s, \underline{v}, \underline{y})=\underbrace{p(\underline{u})}_{\text {factor }} \underbrace{p(s \mid \underline{u})}_{\text {factor }} \underbrace{p(\underline{v} \mid s, \underline{u})}_{\text {factor }} \underbrace{p(\underline{y} \mid \underline{v})}_{\text {factor }}
$$

Propagate beliefs between variables $\rightarrow$ factors \& factors $\rightarrow$ variables

## Efficient Inference via Message Passing - 1

Assume Discrete Variables, Each taking $K$ values


Suppose want marginal distribution $P\left(x_{2}\right)$. By sum formula or Total Probability Theorem $P\left(x_{2}\right)=\sum_{x_{1}} \sum_{x_{3}} \sum_{x_{4}} f_{a}\left(x_{1}, x_{2}\right) f_{b}\left(x_{2}, x_{3}\right) f_{c}\left(x_{2}, x_{4}\right)$

$$
=\underbrace{\left(\sum_{x_{1}} f_{a}\left(x_{1}, x_{2}\right)\right.}_{\mu_{f_{a} \rightarrow x_{2}}\left(x_{2}\right)} \underbrace{\left(\sum_{x_{3}} f_{b}\left(x_{2}, x_{3}\right)\right.}_{\mu_{f_{b} \rightarrow x_{2}}\left(x_{2}\right)} \underbrace{\left(\sum_{x_{4}} f_{c}\left(x_{2}, x_{4}\right)\right)}_{\mu_{f_{c} \rightarrow x_{2}}\left(x_{2}\right)}=\mu_{f_{a} \rightarrow x_{2}}\left(x_{2}\right) \mu_{f_{b} \rightarrow x_{2}}\left(x_{2}\right) \mu_{f_{c} \rightarrow x_{2}}\left(x_{2}\right)
$$

Note: $2 K^{3}$ multiplications and $K^{3}$ additions versus 2 multiplications and $3 K$ additions for each $x_{2}$ ! In fact, can compute any marginal distribution by exchanging messages from variables to factors and factors to variables.

Variables


Corresponding Bipartite Graph

$$
P\left(x_{1}, x_{2}, x_{3}, x_{4}\right)=f_{a}\left(x_{1}, x_{2}\right) f_{b}\left(x_{2}, x_{3}\right) f_{c}\left(x_{2}, x_{4}\right)
$$

## Efficient Inference via Message Passing - 2

| Leaves $\rightarrow$ Root |
| :--- | :--- |
| $\mu_{x_{3} \rightarrow f_{b}}\left(x_{3}\right)=1 \forall x_{3}$ |
| $\mu_{f_{b} \rightarrow x_{2}}\left(x_{2}\right)=\sum_{x_{3}} f_{b}\left(x_{2}, x_{3}\right) \mu_{x_{3} \rightarrow f_{b}}\left(x_{3}\right)$ |
| $\mu_{x_{4} \rightarrow f_{c}}\left(x_{4}\right)=1 \forall x_{4}$ |
| $\mu_{f_{c} \rightarrow x_{2}}\left(x_{2}\right)=\sum_{x_{4}} f_{c}\left(x_{2}, x_{4}\right) \mu_{x_{4} \rightarrow f_{c}}\left(x_{4}\right)$ |
| $\mu_{x_{2} \rightarrow f_{a}}\left(x_{2}\right)=\mu_{f_{b} \rightarrow x_{2}}\left(x_{2}\right) \mu_{f_{c} \rightarrow x_{2}}\left(x_{2}\right)$ |
| $\mu_{f_{a} \rightarrow x_{1}}\left(x_{1}\right)=\sum_{x_{2}} f_{a}\left(x_{1}, x_{2}\right) \mu_{x_{2} \rightarrow f_{a}}\left(x_{2}\right)$ |$\quad$| Root $\rightarrow$ Leaves |
| :--- |
| $\mu_{x_{1} \rightarrow f_{a}}\left(x_{1}\right)=1 \forall x_{1}$ |
| $\mu_{f_{a} \rightarrow x_{2}}\left(x_{2}\right)=\sum_{x_{1}} f_{a}\left(x_{1}, x_{2}\right) \mu_{x_{1} \rightarrow f_{a}}\left(x_{1}\right)$ |
| $\mu_{x_{2} \rightarrow f_{b}}\left(x_{2}\right)=\mu_{f_{a} \rightarrow x_{2}}\left(x_{2}\right) \mu_{f_{c} \rightarrow x_{2}}\left(x_{2}\right)$ |
| $\mu_{x_{2} \rightarrow f_{c}}\left(x_{2}\right)=\mu_{f_{a} \rightarrow x_{2}}\left(x_{2}\right) \mu_{f_{b} \rightarrow x_{2}}\left(x_{2}\right)$ |
| $\mu_{f_{b} \rightarrow x_{3}}\left(x_{3}\right)=\sum_{x_{2}} f_{b}\left(x_{2}, x_{3}\right) \mu_{x_{2} \rightarrow f_{b}}\left(x_{2}\right)$ |
| $\mu_{f_{c} \rightarrow x_{4}}\left(x_{4}\right)=\sum_{x_{2}} f_{c}\left(x_{2}, x_{4}\right) \mu_{x_{2} \rightarrow f_{c}}\left(x_{2}\right)$ |

Assume Discrete Variables, Each taking $K$ values

$P\left(x_{1}\right)=\mu_{f_{a} \rightarrow x_{1}}\left(x_{1}\right) ; P\left(x_{3}\right)=\mu_{f_{b} \rightarrow x_{3}}\left(x_{3}\right) ; P\left(x_{4}\right)=\mu_{f_{c} \rightarrow x_{4}}\left(x_{4}\right)$ $P\left(x_{2}\right)=\mu_{f_{a} \rightarrow x_{2}}\left(x_{2}\right) \mu_{f_{b} \rightarrow x_{2}}\left(x_{2}\right) \mu_{f_{c} \rightarrow x_{2}}\left(x_{2}\right)$

Copyright K.R. Pattipati, 2001-2018

To compute all marginals, form an arbitrary root node (say $x_{1}$ ), propagate and store messages from the leaves to the root, and then propagate and store messages from the root to the leaves. Marginals are products of the factor messages received at each variable nodes. Normalize marginals Exact for trees. For arbitrary graphs, iterate (approximate belief propagation (ABF)). In fact, you can pass messages asynchronously from nodes $\rightarrow$ factors and factors $\rightarrow$ nodes until convergence.

Evidence is easy to include

路

## Bayes Rule as Message Passing

$$
f_{a}=P(x) \quad x \quad f_{b}=P(y \mid x)^{y}
$$

$P(x)=\left[\begin{array}{l}0.8 \\ 0.2\end{array}\right] ; x \in\{1,2\}$
$P(y \mid x)=\left[\begin{array}{ccc}0.05 & 0.65 & 0.30 \\ 0.8 & 0.10 & 0.10\end{array}\right] ; y \in\{1,2,3\}$
Want $P(y)$ and $P(x \mid y=1)$

Marginal $P(y)$
$\mu_{f_{a} \rightarrow x}(x)=P(x) \Rightarrow \mu_{f_{a} \rightarrow x}(x=1)=0.8 ; \mu_{f_{a} \rightarrow x}(x=2)=0.2 \quad$ Bishop's "Who done it" Example
$\mu_{x \rightarrow f_{b}}(x)=\mu_{f_{a} \rightarrow x}(x)$
$\mu_{f_{b} \rightarrow y}(y)=\sum_{x} f_{b}(x, y) \mu_{f_{a} \rightarrow x}(x) \Rightarrow \mu_{f_{b} \rightarrow y}(y=1)=0.2 ; \mu_{f_{b} \rightarrow y}(y=2)=0.54 ; \mu_{f_{b} \rightarrow y}(y=3)=0.26$
$P(y)=\mu_{f_{b} \rightarrow y}(y)$
Evidence, $y=1$. What is $P(x \mid y=1)$ ?
Treat $x$ as the root.

$$
\begin{aligned}
& \mu_{y \rightarrow f_{b}}(y=1)=1 ; \mu_{y \rightarrow f_{b}}(y=2)=0 ; \mu_{y \rightarrow f_{b}}(y=3)=0 \\
& \mu_{f_{b} \rightarrow x}(x)=\sum_{y=1}^{3} f_{b}(x, y) \mu_{y_{b} \rightarrow f_{b}}(y) \Rightarrow \mu_{f_{b} \rightarrow x}(x=1)=0.05 ; \mu_{f_{b} \rightarrow x}(x=2)=0.8 \\
& \mu_{f_{a} \rightarrow x}(x)=P(x) \Rightarrow \mu_{f_{a} \rightarrow x}(x=1)=0.8 ; \mu_{f_{a} \rightarrow x}(x=2)=0.2 \\
& P(x \mid y=1)=\mu_{f_{a} \rightarrow x}(x) \mu_{f_{b} \rightarrow x}(x) \Rightarrow\left[\begin{array}{c}
\mu_{f_{a} \rightarrow x}(1) \mu_{f_{b} \rightarrow x}(1) \\
\mu_{f_{a} \rightarrow x}(2) \mu_{f_{b} \rightarrow x}(2)
\end{array}\right]=\left[\begin{array}{l}
0.04 \\
0.16
\end{array}\right] \Rightarrow \text { Normalize }:\left[\begin{array}{l}
0.20 \\
0.80
\end{array}\right]
\end{aligned}
$$

## Sum-Product Belief Propagation Algorithm

- Factors $\rightarrow$ Variables Messages (SUM)
messages sent by a factor node to a variable node involves multiplying all the incoming messages (except variable node $x$ ) with the factor and summing over all the variables except $x$

$$
\begin{aligned}
& \mu_{f_{a} \rightarrow x_{1}}\left(x_{1}\right)=\sum_{x_{2}} f_{a}\left(x_{1}, x_{2}\right) \mu_{x_{2} \rightarrow f_{a}}\left(x_{2}\right) ; \mu_{f_{a} \rightarrow x_{2}}\left(x_{2}\right)=\sum_{x_{1}} f_{a}\left(x_{1}, x_{2}\right) \mu_{x_{1} \rightarrow f_{a}}\left(x_{1}\right) \\
& \mu_{f_{b} \rightarrow x_{2}}\left(x_{2}\right)=\sum_{x_{3}} f_{b}\left(x_{2}, x_{3}\right) \mu_{x_{3} \rightarrow f_{b}}\left(x_{3}\right) ; \mu_{f_{b} \rightarrow x_{3}}\left(x_{3}\right)=\sum_{x_{2}} f_{b}\left(x_{2}, x_{3}\right) \mu_{x_{2} \rightarrow f_{b}}\left(x_{2}\right) \\
& \mu_{f_{c} \rightarrow x_{2}}\left(x_{2}\right)=\sum_{x_{4}} f_{c}\left(x_{2}, x_{4}\right) \mu_{x_{4} \rightarrow f_{b}}\left(x_{4}\right) ; \mu_{f_{c} \rightarrow x_{4}}\left(x_{4}\right)=\sum_{x_{2}} f_{c}\left(x_{2}, x_{4}\right) \mu_{x_{2} \rightarrow f_{c}}\left(x_{2}\right)
\end{aligned}
$$

- Variables $\rightarrow$ Factors Messages (PRODUCT)

Message sent by a variable node to a factor node is the product of all the incoming messages along all of the other links (factors)

$$
\begin{aligned}
& \mu_{x_{1} \rightarrow f_{a}}\left(x_{1}\right)=1 ; \mu_{x_{3} \rightarrow f_{b}}\left(x_{3}\right)=1 ; \mu_{x_{4} \rightarrow f_{c}}\left(x_{3}\right)=1 \\
& \mu_{x_{2} \rightarrow f_{a}}\left(x_{2}\right)=\mu_{f_{b} \rightarrow x_{2}}\left(x_{2}\right) \mu_{f_{c} \rightarrow x_{2}}\left(x_{2}\right) ; \\
& \mu_{x_{2} \rightarrow f_{b}}\left(x_{2}\right)=\mu_{f_{a} \rightarrow x_{2}}\left(x_{2}\right) \mu_{f_{c} \rightarrow x_{2}}\left(x_{2}\right) ; \\
& \mu_{x_{2} \rightarrow f_{c}}\left(x_{2}\right)=\mu_{f_{a} \rightarrow x_{2}}\left(x_{2}\right) \mu_{f_{b} \rightarrow x_{2}}\left(x_{2}\right)
\end{aligned}
$$



## Exponential Family

## $\square$ Exponential family:

Gaussian, Poisson, Gamma, Binomial, Exponential, Beta, Weibull,...
$p(\underline{x} \mid \underline{w})^{\dagger}=\frac{1}{Z(\underline{w})} h(\underline{x}) \exp \left\{\underline{w}^{T} \underline{T}(\underline{x})\right\}=h(\underline{x}) \exp \left\{\underline{w}^{T} \underline{T}(\underline{x})-A(\underline{w})\right\}$
$Z(\underline{w})=\int_{\underline{x}} h(\underline{x}) \exp \left\{\underline{w}^{T} \underline{T}(\underline{x})\right\} d \underline{x} \quad \ldots .$. called partition function
$A(\underline{w})=\ln Z(\underline{w}) \quad . . . \log$ partition function or cumulant function
Gradients of $A(\underline{w})$ give moments of density
$\underline{w} \sim$ canonical parameters
$\underline{T}(\underline{x})=$ "sufficient statistics"
$h(\underline{x})=$ scaling constant
$\dagger$ most often $\underline{\theta}$ is used to denote model parameters. Since Neural Networks use $\underline{w}$ for weights, we use $\underline{w}$ here.

## Exponential Family and Conjugate Priors

## $\square$ Conjugate Prior

$$
p(\underline{w})=g(\eta, \underline{v}) \exp \left\{\underline{w}^{T} \underline{v}-\eta A(\underline{w})\right\} ; \eta, \underline{v} \text { are hyper parameters }
$$

Data $D=\left\{\underline{x}_{1}, \underline{x}_{2}, \ldots, \underline{x}_{N}\right\}$
$p\left(\underline{x}_{n} \mid \underline{w}\right)=h\left(\underline{x}_{n}\right) \exp \left\{\underline{w}^{T} \underline{T}\left(\underline{x}_{n}\right)-A(\underline{w})\right\}$
$p(\underline{w} \mid D)=g\left(\eta+N, \underline{v}+\sum_{n=1}^{N} \underline{T}\left(\underline{x}_{n}\right)\right) \exp \left\{\underline{w}^{T}\left[\underline{v}+\sum_{n=1}^{N} \underline{T}\left(\underline{x}_{n}\right)\right]-(\eta+N) A(\underline{w})\right\}$
Posterior belongs to the exponential family
Examples of Likelihoods and Conjugate Priors

- (Bernoulli or Binomial or Geometric) -Beta
- Scalar Normal-(Normal, Inverse-Gamma)
- Multivariate Normal-(Normal, Inverse Wishart)
- (Poisson, Pareto, Exponential)-Gamma


## Bernoulli Distribution

$\square$ Bernoulli

$$
p(x)=p^{x}(1-p)^{1-x}=\exp [x \ln p+(1-x) \ln (1-p)]=\exp \left[x \ln \frac{p}{1-p}+\ln (1-p)\right]
$$

$$
\text { Let } \quad T(x)=x ; w=\ln \frac{p}{1-p} \Rightarrow p=\frac{1}{1+e^{-w}}=\sigma(w)
$$

$$
\Rightarrow 1-p=\frac{1}{1+e^{w}}=\sigma(-w)=1-\sigma(w)
$$

$$
\text { so, } p(x)=\exp \left[x \ln \frac{p}{1-p}-\ln \left(1+e^{w}\right)\right]
$$

$$
A(w)=\ln \left(1+e^{w}\right)
$$

$$
\frac{d A(w)}{d w}=\frac{e^{w}}{1+e^{w}}=\frac{1}{1+e^{-w}}=p=E(x)=E[T(x)]
$$



$$
\frac{d^{2} A(w)}{d w^{2}}=p(1-p)=\operatorname{Var}(x)=\operatorname{Var}[T(x)]
$$

## Binomial Likelihood -Beta Prior for Learning $p$

$\square$ Multiple trials, each with $x_{i} \in\{0,1\}$
$s_{n}=\sum_{i=1}^{n} x_{i} ; s_{n+1}=s_{n}+x_{n+1} ; s_{0}=0$
Beta Function: $B(\alpha, \beta)=\int_{0}^{1} q^{\alpha-1}(1-q)^{\beta-1} d q=\frac{\Gamma(\alpha) \Gamma(\beta)}{\Gamma(\alpha+\beta)}=\frac{(\alpha-1)!(\beta-1)!}{(\alpha+\beta-1)!}$
$p\left(s_{n} \mid p\right)=\binom{n}{s} p^{s_{n}}(1-p)^{n-s_{n}}=\frac{p^{s_{n}}(1-p)^{n-s_{n}}}{(n+1) B\left(s_{n}+1, n+1-s_{n}\right)}$
$p(p)=\frac{p^{\alpha-1}(1-p)^{\beta-1}}{B(\alpha, \beta)}=\operatorname{Beta}(\alpha, \beta)$
$p\left(p \mid s_{n}\right)=\frac{p^{s_{n}+\alpha-1}(1-p)^{n-s_{n}+\beta-1}}{\int_{0}^{1} q^{s_{n}+\alpha-1}(1-q)^{n-s_{n}+\beta-1} d q}=\frac{p^{s_{n}+\alpha-1}(1-p)^{n-s_{n}+\beta-1}}{B\left(s_{n}+\alpha, n-s_{n}+\beta\right)}=\operatorname{Beta}\left(s_{n}+\alpha, n-s_{n}+\beta\right)$
Estimate: $\hat{p}_{n}=\frac{s_{n}+\alpha}{n+\alpha+\beta}$
Recursion: $\hat{p}_{n+1}=\frac{s_{n+1}+\alpha}{n+\alpha+\beta+1}=\frac{s_{n}+x_{n+1}+\alpha}{n+\alpha+\beta+1}=\frac{n+\alpha+\beta}{n+\alpha+\beta+1} \hat{p}_{n}+\frac{1}{n+\alpha+\beta+1} x_{n+1}$
So, $\hat{p}_{n+1}=\hat{p}_{n}+\frac{1}{n+\alpha+\beta+1}\left(x_{n+1}-\hat{p}_{n}\right)$ Similar to Least Squares and Kalman Filter

## Univariate Gaussian (Normal) Distribution

## $\square$ Univariate Gaussian Distribution

$$
p(\underline{x} \mid \underline{w})=\frac{1}{Z(\underline{w})} h(\underline{x}) \exp \left\{\underline{w}^{T} \underline{T}(\underline{x})\right\}=h(\underline{x}) \exp \left\{\underline{w}^{T} \underline{T}(\underline{x})-A(\underline{w})\right\}
$$

$$
p\left(x \mid \mu, \sigma^{2}\right)=\frac{1}{\sqrt{2 \pi} \sigma} \exp \left\{-\frac{1}{2 \sigma^{2}}(x-\mu)^{2}\right\}
$$

$$
=\exp \left\{-\frac{1}{2 \sigma^{2}}(x-\mu)^{2}-\ln \sigma-\frac{1}{2} \ln (2 \pi)\right\}
$$

$$
=\exp \left\{\frac{\mu}{\sigma^{2}} x-\frac{1}{2 \sigma^{2}} x^{2}-\frac{\mu^{2}}{2 \sigma^{2}}-\ln \sigma-\frac{1}{2} \ln (2 \pi)\right\}
$$

$$
\sigma^{2}: \text { variance }
$$

$$
={\underset{n}{h(x)}}_{1} \exp \{\underbrace{\binom{\frac{\mu}{\sigma^{2}}}{-\frac{1}{2 \sigma^{2}}}^{T}}_{\underline{\underline{w}}}\binom{x}{x^{2}}-\underbrace{\left(\frac{\mu^{2}}{2 \sigma^{2}}+\ln \sigma+\frac{1}{2} \ln (2 \pi)\right)}_{\underline{\underline{T}(x)}}\}
$$

$$
1 / \sigma^{2}: \text { "information" }
$$

## Log Partition Function: UVN

$$
\begin{aligned}
& \underline{T}(x)=\left[\begin{array}{c}
x \\
x^{2}
\end{array}\right] ; \underline{w}=\left[\begin{array}{c}
\frac{\mu}{\sigma^{2}} \\
-\frac{1}{2 \sigma^{2}}
\end{array}\right]=\left[\begin{array}{l}
w_{1} \\
w_{2}
\end{array}\right] \\
& A(\underline{w})=\frac{\mu^{2}}{2 \sigma^{2}}+\ln \sigma+\frac{1}{2} \ln (2 \pi) \\
& =-\frac{w_{1}^{2}}{4 w_{2}}-\frac{1}{2} \ln \left(-2 w_{2}\right)+\frac{1}{2} \ln (2 \pi) \\
& \begin{aligned}
\frac{\partial A(\underline{w})}{\partial w_{1}} & =-\frac{w_{1}}{2 w_{2}}=-\frac{\mu}{\sigma^{2}} \cdot\left(-\sigma^{2}\right)=\mu=E(x) \\
\frac{\partial A(\underline{w})}{\partial w_{2}} & =\frac{w_{1}^{2}}{4 w_{2}^{2}}-\frac{1}{2} \cdot \frac{1}{\left(-2 w_{2}\right)} \cdot(-2) \\
& =\frac{w_{1}^{2}}{4 w_{2}^{2}}-\frac{1}{2 w_{2}} \\
& =\mu^{2}+\sigma^{2}=E\left(x^{2}\right)
\end{aligned} \\
& \nabla_{\underline{w}} A(\underline{w})=E[\underline{T}(\underline{x})] \\
& \frac{\partial^{2} A(\underline{w})}{\partial w_{1}^{2}}=-\frac{1}{2 w_{2}}=\sigma^{2}=\operatorname{var}(x) \\
& \frac{\partial^{2} A(\underline{w})}{\partial w_{1} \partial w_{2}}=\frac{\partial^{2} A(\underline{w})}{\partial w_{2} \partial w_{1}}=\frac{w_{1}}{2 w_{2}^{2}}=\frac{1}{2} \cdot \frac{\mu}{\sigma^{2}} \cdot 4 \sigma^{4}=2 \mu \sigma^{2}=E\left\{(x-\mu)\left(x^{2}-\mu^{2}-\sigma^{2}\right)\right\} \\
& \begin{array}{l}
E\left(x^{3}\right)=\mu^{3}+3 \mu \sigma^{2} \\
E\left(x^{4}\right)=\mu^{4}+6 \mu^{2} \sigma^{2}+3 \sigma^{4}
\end{array} \\
& \frac{\partial^{2} A(\underline{w})}{\partial w_{2}^{2}}=-\frac{w_{1}^{2}}{2 w_{2}^{3}}+\frac{1}{2 w_{2}^{2}}=\frac{1}{2} \cdot \frac{\mu^{2}}{\sigma^{4}} \cdot 8 \sigma^{6}+2 \sigma^{4}=4 \mu^{2} \sigma^{2}+2 \sigma^{4}=\operatorname{var}\left(x^{2}\right) \\
& \nabla_{\underline{w}}^{2} A(\underline{w})=\left[\begin{array}{cc}
\sigma^{2} & 2 \mu \sigma^{2} \\
2 \mu \sigma^{2} & 4 \mu^{2} \sigma^{2}+2 \sigma^{4}
\end{array}\right]=\operatorname{Cov}(T(\underline{x}))
\end{aligned}
$$

## Multivariate Gaussian (Normal) Distribution

## $\square$ Multivariate Gaussian Distribution

$$
\begin{aligned}
& p(\underline{x} \mid \underline{w})=\frac{1}{Z(\underline{w})} h(\underline{x}) \exp \left\{\underline{w}^{T} \underline{T}(\underline{x})\right\}=h(\underline{x}) \exp \left\{\underline{w}^{T} \underline{T}(\underline{x})-A(\underline{w})\right\}_{\text {temase }} \\
& p(\underline{x} \mid \underline{\mu}, \Sigma)=\frac{1}{(2 \pi)^{p / 2}|\Sigma|^{1 / 2}} \exp \left\{-\frac{1}{2}(\underline{x}-\underline{\mu})^{T} \Sigma^{-1}(\underline{x}-\underline{\mu})\right\} \\
& =\exp \left\{\operatorname{tr}\left[-\frac{1}{2} \Sigma^{-1}(\underline{x}-\underline{\mu})(\underline{x}-\underline{\mu})^{T}-\frac{1}{2} \ln |\Sigma|-\frac{p}{2} \ln (2 \pi)\right\}\right. \\
& =\exp \left\{\operatorname{tr}\left[\underline{\mu}^{T} \Sigma^{-1} \underline{x}-\frac{1}{2} \Sigma^{-1} \underline{x}^{T} \underline{x}^{T}\right]-\frac{1}{2} \Sigma^{-1} \underline{\mu}_{\underline{\mu}} \underline{T}^{T}-\frac{1}{2} \ln |\Sigma|-\frac{p}{2} \ln (2 \pi)\right\} \\
& ={\underset{h(\underline{x})}{1} \exp \left\{\operatorname{tr}\left[\begin{array}{c}
\left.\Sigma^{-1} \underline{\mu}{ }^{( }\right)^{T} \\
-\frac{1}{2} \Sigma^{-1}
\end{array}\right)^{\left(\begin{array}{c}
\underline{x} \\
\underline{x}^{T} \\
\underline{x}^{T}
\end{array}\right)}\right]-\underbrace{\left(\frac{1}{2} \underline{\mu}^{T} \Sigma^{-1} \underline{\mu}+\frac{1}{2} \ln |\Sigma|+\frac{p}{2} \ln (2 \pi)\right)}_{\underline{T}(\underline{x})}\}}_{A(\underline{w})}^{(\underbrace{2})} \\
& \square \text { Information form } \\
& p\left(\underline{x} ; \underline{\mu}, J^{-1}\right)=\frac{|J|^{1 / 2}}{(2 \pi)^{p / 2}} \exp \left\{-\frac{1}{2}(\underline{x}-\underline{\mu})^{T} J(\underline{x}-\underline{\mu})\right\}
\end{aligned}
$$

## Log Partition Function：MVN

He $T(\underline{x})=\left[\begin{array}{c}\underline{x} \\ \underline{x}^{T}\end{array}\right] ; \underline{w}=\left[\begin{array}{c}\Sigma^{-1} \underline{\mu} \\ -\frac{1}{2} \Sigma^{-1}\end{array}\right]=\left[\begin{array}{l}\underline{w}_{1} \\ W_{2}\end{array}\right]$

$$
\begin{aligned}
& A(\underline{w})=\frac{1}{2} \underline{\mu}^{T} \Sigma^{-1} \underline{\mu}+\frac{1}{2} \ln |\Sigma|+\frac{p}{2} \ln (2 \pi)=-\frac{1}{4}\left(\underline{\mu}^{T} \Sigma^{-1}\right)\left(-\frac{1}{2} \Sigma^{-1}\right)^{-1}\left(\Sigma^{-1} \underline{\mu}\right)-\frac{1}{2} \ln \left|\Sigma^{-1}\right|+\frac{p}{2} \ln (2 \pi) \\
& =-\frac{1}{4} \underline{w}_{1}^{T} W_{2}^{-1} \underline{w}_{1}-\frac{1}{2} \ln \left|-2 W_{2}\right|+\frac{p}{2} \ln (2 \pi) \\
& \nabla_{w_{1}} A(\underline{w})=-\frac{1}{2} W_{2}^{-1} \underline{w}_{1}=-\frac{1}{2}(-2 \Sigma) \Sigma^{-1} \underline{\mu}=\underline{\mu}=E(\underline{x}) \\
& \nabla_{W_{2}} A(\underline{w})=\nabla_{W_{2}}\left(-\frac{1}{4} \operatorname{tr}\left[W_{2}^{-1} \underline{w}_{1} \underline{w}_{1}^{T}\right]-\frac{1}{2} \ln \left|-2 W_{2}\right|\right) \\
& \nabla_{\underline{w}}\left(\underline{w}^{T} A \underline{w}\right)=\nabla_{\underline{w}}\left(\operatorname{tr}\left[A \underline{w} \underline{w}^{T}\right]\right)=2 A \underline{w} \\
& \nabla_{A}[\ln |A|]=A^{-1} \\
& =\nabla_{W_{2}}\left(-\frac{1}{4} \operatorname{tr}\left[W_{2}^{-1} \underline{w}_{1} \underline{w}_{1}^{T}\right]-\frac{1}{2} \ln \left[(-2)^{p}\right]-\frac{1}{2} \ln \left|W_{2}\right|\right. \\
& =\frac{1}{4}\left(W_{2}^{-1} \underline{w}_{1} \underline{w}_{1}^{T} W_{2}^{-1}\right)-\frac{1}{2}\left(W_{2}\right)^{-1} \\
& \left.=\frac{1}{4}\left((-2 \Sigma) \Sigma^{-1} \underline{\mu}_{\underline{\mu}} \underline{\Sigma}^{-1}(-2 \Sigma)\right)-\frac{1}{2}\left(W_{2}\right)^{-1}=\underline{\mu}_{\underline{\mu}} \underline{T}^{T}+\Sigma=E \underline{x} \underline{x}^{T}\right\}
\end{aligned}
$$

## Examples of Sufficient Statistics

Normal ：$\sum_{i=1}^{N} x_{i} ; \sum_{i=1}^{N} x_{i}^{2}$ for scalars；$\sum_{i=1}^{N} x_{i} ; \sum_{i=1}^{N} x_{i} \underline{X}_{i}^{T}$ for vectors
（or）sample mean and sample covariance
Bernoulli：$\sum_{i=1}^{N} x_{i}$
Poisson：$\sum_{i=1}^{N} x_{i}$
Sufficient statistics help in experimenting with data transformations

Exponential：$\sum_{i=1}^{N} x_{i}$
Uniform： $\max \left(x_{i}\right)$
Gamma ：$\sum_{i=1}^{N} x_{i} ; \sum_{i=1}^{N} \ln x_{i}$（or）$\sum_{i=1}^{N} x_{i} ; \prod_{i=1}^{N} x_{i}$
Beta $: \sum_{i=1}^{N} \ln x_{i} ; \sum_{i=1}^{N} \ln \left(1-x_{i}\right)($ or $) \prod_{i=1}^{N} x_{i} ; \prod_{i=1}^{N}\left(1-x_{i}\right)$

- Want to learn the mean $\underline{\mu}$ (covariance $\Sigma$ known)

$$
p\left(\underline{x}^{n} \mid \underline{\mu}\right)=N\left(\underline{\mu}, \Sigma_{v}\right) \text { and } p(\underline{\mu})=N\left(\underline{\hat{\mu}}^{0}, \hat{\Sigma}^{0}\right)
$$

$\square$ Data: Let $D^{n}=\left\{\underline{x}^{1}, \underline{x}^{2}, \ldots \ldots, \underline{x}^{n-1}, \underline{x}^{n}\right\}=\left\{D^{n-1}, \underline{x}^{n}\right\}$

$$
p\left(D^{n} \mid \underline{\theta}\right)=p\left(\underline{x}^{n} \mid \underline{\theta}\right) p\left(D^{n-1} \mid \underline{\theta}\right) \text {...conditionally i.i.d. }
$$

$\square$ Gaussian density reproduces itself $\Rightarrow$ Posterior density is also Gaussian

$$
p\left(\underline{x}^{n}\right)=N\left(\underline{\hat{\mu}}^{n-1}, \hat{\Sigma}^{n-1}+\Sigma_{v}\right)
$$

$$
p\left(\underline{\mu} \mid D^{n-1}\right)=N\left(\underline{\hat{\mu}}^{n-1}, \hat{\Sigma}^{n-1}\right) \text { and } p\left(\underline{x}^{n} \mid \underline{\mu}\right)=N\left(\underline{\mu}, \Sigma_{v}\right)
$$

$$
\begin{aligned}
& \text { Model: } \\
& \underline{x}^{n}=\underline{\mu}+\underline{v}^{n} ; \\
& \underline{v}^{n}=N\left(\underline{0}, \Sigma_{v}\right)
\end{aligned}
$$

$$
\underline{\hat{\mu}}^{n}=\underline{\hat{\mu}}^{n-1}+\hat{\Sigma}^{n-1}\left[\hat{\Sigma}^{n-1}+\Sigma_{v}\right]^{-1}\left(\underline{x}^{n}-\underline{\hat{\mu}}^{n-1}\right) \Rightarrow \text { RLS or a simple Kalman filter }
$$

$$
\hat{\Sigma}^{n}=\left[\left(\hat{\Sigma}^{n-1}\right)^{-1}+\Sigma_{v}^{-1}\right]^{-1}=\hat{\Sigma}^{n-1}-\hat{\Sigma}^{n-1}\left[\hat{\Sigma}^{n-1}+\Sigma_{v}\right]^{-1} \hat{\Sigma}^{n-1}
$$

$$
=\hat{\Sigma}^{n-1}\left[\hat{\Sigma}^{n-1}+\Sigma_{v}\right]^{-1} \Sigma_{v} \quad \text { Add \& subtract } \Sigma_{v}
$$

$$
\begin{aligned}
& =\Sigma_{v}\left[\hat{\Sigma}^{n-1}+\Sigma_{v}\right]^{-1} \hat{\Sigma}^{n-1} ; \hat{\Sigma}^{0} \text { initial value } \\
& \text { t if mean drifts? KF,MM, IMM, Information filter }
\end{aligned}
$$

## Recursive Estimation of Constant Parameters



Measurement can be a linear combination of parameters corrupted by noise

## Recursive Estimation of Drifting Parameters

$$
\begin{array}{ll}
\text { Model: } \\
\underline{\mu}^{n+1}=\underline{\mu}^{n}+\underline{\delta}^{n} \Delta+\underline{w}_{1}^{n} ; \Delta=\text { sampling interval } ; \underline{w}_{1}^{n} \sim N\left(\underline{0}, Q_{1}\right) ; \underline{\delta}^{n}=\text { drift rate at sample } n \\
\underline{\delta}^{n+1}=\underline{\delta}^{n}+\underline{w}_{2}^{n} ; \underline{w}_{2}^{n} \sim N\left(\underline{0}, Q_{2}\right) & \text { 2-state nearly-constant velocity mode } \\
\underline{x}^{n}=\underline{\mu}^{n}+\underline{v}^{n} ; \underline{v}^{n}=N\left(\underline{0}, \Sigma_{v}\right) &
\end{array}
$$

kalmanTrackingDemo from Murphy, Page 632


## Entropy and Conditional Entropy

■ Suppose $X \in\left\{x_{1}, x_{2}, \ldots, x_{m}\right\}$ and $Y \in\left\{y_{1}, y_{2}, \ldots, y_{n}\right\}$ are discrete random variables with the joint probability mass function $p(x, y)$

- Example: In the text categorization problem, $X$ is a random variable over the set of terms, and $Y$ is a random variable over the set of documents. Evidently, $p(x, y)$ is an $m$ by $n$ matrix, termed a two-dimensional contingency table or a two-way frequency table; this is often estimated from the co-occurrence data.
- Entropy of $\boldsymbol{X}$

Gaussian: $H(X)=\log (\sigma \sqrt{2 \pi e})$

$$
H(X)=E\left[-\log _{2} p(X)\right]=-\sum_{x \in\left\{x_{1}, x_{2}, \ldots, x_{m}\right\}} p(x) \log _{2} p(x)
$$

- A measure of the average amount of information (in bits) required to describe the random variable.
- Conditional Entropy $\boldsymbol{H}(\boldsymbol{X} \mid \boldsymbol{Y})$
- Entropy of random variable $X$ given another random variable, $Y$
- Entropy of a pair of random variables $(X, Y)$ is the entropy of one plus the conditional entropy of the other, i.e.,

$$
H(X, Y)=H(X)+H(Y \mid X)=H(Y)+H(X \mid Y)
$$

We find it convenient to use $\ln$ instead of $\log _{2}$

## Mutual Information (MI)

- Mutual Information (Information Gain, MI)
- Reduction in uncertainty due to knowledge of another random variable

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

- $X$ and $Y$ are independent, i.e., $p(x, y)=p(x) p(y) \Rightarrow$ mutual information is zero
- When there is a deterministic one-to-one relationship between $X$ and $Y$, mutual information is the largest (i.e., equal to the entropy of the random variable)
- In general, the mutual information is bounded by

$$
0 \leq I(X ; Y) \leq \min (H(X), H(Y))
$$

$$
\begin{aligned}
& p(\underline{x}, \underline{y})=N\left(\left[\begin{array}{l}
\underline{\mu}_{x} \\
\underline{\mu}_{y}
\end{array}\right] ;\left[\begin{array}{ll}
\Sigma_{x x} & \Sigma_{x y} \\
\Sigma_{x y}^{T} & \Sigma_{y y}
\end{array}\right]\right) \\
& \Rightarrow p(\underline{x})=N\left(\underline{\mu}_{x}, \Sigma_{x x}\right) ; p(\underline{y})=N\left(\underline{\mu}_{y}, \Sigma_{y y}\right) \\
& p(\underline{x} \mid \underline{y})=N\left(\underline{\mu}_{x}+\Sigma_{x y} \Sigma_{y y}^{-1}\left(\underline{y}-\underline{\mu}_{y}\right), \Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right) \\
& I(X ; Y)=H(X)+H(Y)-H(X, Y)
\end{aligned}
$$

A greedy forward/ backward feature selection algorithm based on (1) and (2) for regression.

1. Remove redundant features (large $\rho$ ).
2. Pick features $\{x\}$ with large correlation with respect to dependent variable (output) $y$.
$H(X)=E_{p(\underline{x})}[-\ln p(\underline{x})]=\frac{1}{2}\left[n_{x} \ln (2 \pi e)+\ln \left|\Sigma_{x x}\right|\right] ; H(Y)=\frac{1}{2}\left[n_{y} \ln (2 \pi e)+\ln \left|\Sigma_{y y}\right|\right]$
$H(X, Y)=\frac{1}{2}\left[\left(n_{x}+n_{y}\right) \ln (2 \pi e)+\ln \left|\Sigma_{x x}\right|+\ln \left|\Sigma_{y y}-\Sigma_{x y}^{T} \Sigma_{x x}^{-1} \Sigma_{x y}\right|\right]$
$=\frac{1}{2}\left[\left(n_{x}+n_{y}\right) \ln (2 \pi e)+\ln \left|\Sigma_{y y}\right|+\ln \left|\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right|\right]$
$I(X ; Y)=\frac{1}{2}\left[\ln \left|\Sigma_{x x}\right|-\ln \left|\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right|\right]=\frac{1}{2}\left[\ln \left|\Sigma_{y y}\right|-\ln \left|\Sigma_{y y}-\Sigma_{x y}^{T} \Sigma_{x x}^{-1} \Sigma_{x y}\right|\right]$
$=-\frac{1}{2} \ln \left|I_{\mathrm{dim}(x)}-\Sigma_{x x}^{-1} \Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right|=-\frac{1}{2} \ln \left|I_{\mathrm{dim}(y)}-\Sigma_{y y}^{-1} \Sigma_{x y}^{T} \Sigma_{x x}^{-1} \Sigma_{x y}\right|$
$x, y$ scalars $: \Sigma_{x x}=\sigma_{x}^{2} ; \Sigma_{x y}=\rho_{x y} \sigma_{x} \sigma_{y} ; \Sigma_{y y}=\sigma_{y}^{2}$
$I(X ; Y)=\frac{1}{2}\left[\ln \sigma_{x}^{2}-\ln \left(\sigma_{x}^{2}\left(1-\rho_{x y}^{2}\right)\right)\right]=-\frac{1}{2} \ln \left(1-\rho_{x y}^{2}\right)=\ln \frac{1}{\sqrt{1-\rho_{x y}^{2}}}$


## Mutual Information for Classification

$z \in\{1,2, . ., C\} ; \underline{x} \in R^{p} ; p(\underline{x})=\sum_{i=1}^{C} P[\underline{x}, z=i]=\sum_{i=1}^{C} P(z=i) p(\underline{x} \mid z=i)=\sum_{i=1}^{c} \pi_{i} N\left(\underline{x}^{\prime}, \underline{\mu}_{i},,_{i}\right)$
$I(\underline{x}, z)=H(z)-H(z \mid \underline{x})=H(\underline{x})-H(\underline{x} \mid z)$
It is not possible to compute $H(\underline{x})$ or $H(z \mid \underline{x})$ in closed-form. $H(z)$ and $H(\underline{x} \mid z)$ are easy to compute.
$H(\underline{x} \mid z)=\frac{1}{2} \sum_{i=1}^{C} \pi_{i}\left(p \ln (2 \pi e)+\ln \left|\Sigma_{i}\right|\right)$
Upper Bound on $H(\underline{x}): H(\underline{x})=E_{p(\underline{x})}\left\{-\ln \left(\sum_{i=1}^{C} \pi_{i} N\left(\underline{x} ; \underline{\mu}_{i}, \Sigma_{i}\right)\right)\right\} \leq \sum_{i=1}^{C} \pi_{i} E_{p(\underline{x})}\left\{-\ln \left(N\left(\underline{x} ; \underline{\mu}_{i}, \nu_{i}\right)\right)\right\}$
$=\frac{1}{2} \sum_{i=1}^{C} \sum_{j=1}^{C} \pi_{i} \pi_{j}\left[p \ln (2 \pi)+\ln \left|\Sigma_{i}\right|+\left(\underline{\mu}_{j}-\underline{\mu}_{i}\right)^{T} \Sigma_{i}^{-1}\left(\underline{\mu}_{j}-\underline{\mu}_{i}\right)+\operatorname{tr}\left(\Sigma_{i}^{-1} \Sigma_{j}\right)\right]$
So, $I(\underline{x}, z) \leq \frac{1}{2} \sum_{i=1}^{C} \sum_{j=1}^{C} \pi_{i} \pi_{j}\left[\left(\underline{\mu}_{j}-\underline{\mu}_{i}\right)^{T} \Sigma_{i}^{-1}\left(\underline{\mu}_{j}-\underline{\mu}_{i}\right)+\operatorname{tr}\left(\Sigma_{i}^{-1} \Sigma_{j}\right)-p\right]$
When $\Sigma_{i}=\Sigma$ for all $i, I(\underline{x}, z) \leq \frac{1}{2} \sum_{i=1}^{C} \sum_{j=1}^{C} \pi_{i} \pi_{j}\left[\left(\underline{\mu}_{j}-\underline{\mu}_{i}\right)^{T} \Sigma^{-1}\left(\underline{\mu}_{j}-\underline{\mu}_{i}\right)\right]$
Best single feature: $k=\arg \max _{l \in\{1,2, \ldots, p\}} \frac{1}{2}\left(\sum_{i=1}^{c} \sum_{j=1}^{c} \pi_{i} \pi_{j}\left[\frac{\left(\mu_{j l}-\mu_{i l}\right)^{2}}{\sigma_{l}^{2}}\right]\right)$
For binary classes: $k=\arg \max _{l \in\{1,2, ., p\}}\left(\frac{\left(\mu_{1 l}-\mu_{2 l}\right)^{2}}{\sigma_{l}^{2}}\right) \sim d_{k}^{\prime 2}$

| A greedy |
| :--- |
| forward/backward |
| feature selection |
| algorithm based on |
| (1) and (2) for |
| classification |

Discriminability measure

A greedy forward/backward feature selection algorithm based on
(1) and (2) for classification

## Measures of Information: Venn Diagrams

- Information (Venn) Diagram illustrates relationships among measures of information: Entropy, Joint Entropy, Conditional Entropy and MI

- $H(Y \mid X)$ is the novel information in $Y$
- $I(X ; Y)$ is the Redundant information between $X$ and $Y$


## Joint Mutual Information Between Gaussian Variables - 1

Consider three Gaussian random variables $x, y, z$
$p(x, y, z)=N\left(\left[\begin{array}{l}\mu_{x} \\ \mu_{y} \\ \mu_{z}\end{array}\right] ;\left[\begin{array}{lll}\sigma_{x}^{2} & \sigma_{x y} & \sigma_{x z} \\ \sigma_{x y} & \sigma_{y}^{2} & \sigma_{y z} \\ \sigma_{x z} & \sigma_{y z} & \sigma_{z}^{2}\end{array}\right]\right)=N\left(\left[\begin{array}{c}\mu_{x} \\ \mu_{y} \\ \mu_{z}\end{array}\right] ;\left[\begin{array}{ccc}\sigma_{x}^{2} & \rho_{x y} \sigma_{x} \sigma_{y} & \rho_{x z} \sigma_{x} \sigma_{z} \\ \rho_{x y} \sigma_{x} \sigma_{y} & \sigma_{y}^{2} & \rho_{y z} \sigma_{y} \sigma_{z} \\ \rho_{x z} \sigma_{x} \sigma_{z} & \rho_{y z} \sigma_{y} \sigma_{z} & \sigma_{z}^{2}\end{array}\right]\right)$
Recall $I(X ; Y)=\frac{1}{2}\left[\ln \left|\Sigma_{x x}\right|-\ln \left|\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right|\right]=\frac{1}{2}\left[\ln \left|\Sigma_{y y}\right|-\ln \left|\Sigma_{y y}-\Sigma_{x y}^{T} \Sigma_{x x}^{-1} \Sigma_{x y}\right|\right]$
$x, y$ scalars $: \Sigma_{x x}=\sigma_{x}^{2} ; \Sigma_{x y}=\rho_{x y} \sigma_{x} \sigma_{y} ; \Sigma_{y y}=\sigma_{y}^{2}$
$\Rightarrow I(X ; Y)=\frac{1}{2}\left[\ln \sigma_{x}^{2}-\ln \left(\sigma_{x}^{2}\left(1-\rho_{x y}^{2}\right)\right)\right]=-\frac{1}{2} \ln \left(1-\rho_{x y}^{2}\right)$
Recall $I(X ; Y ; Z)=I(X ; Y)-I(X ; Y \mid Z)=-\frac{1}{2} \ln \left(1-\rho_{x y}^{2}\right)-I(X ; Y \mid Z)$
Now, $I(X ; Y \mid Z)=H(X \mid Z)+H(Y \mid Z)-H(X, Y \mid Z) \left\lvert\, \begin{aligned} & \text { Simpler one: } \\ & I(X ; Y \mid Z)=H(X \mid Z)-H(X \mid Y, Z) \ldots \text { Try it! }\end{aligned}\right.$
$p(x \mid z)=N\left(\mu_{x}+\frac{\sigma_{x z}}{\sigma_{z}^{2}}\left(z-\mu_{z}\right) ; \sigma_{x}^{2}-\frac{\sigma_{x z}^{2}}{\sigma_{z}^{2}}\right) ; p(y \mid z)=N\left(\mu_{y}+\frac{\sigma_{y z}}{\sigma_{z}^{2}}\left(z-\mu_{z}\right) ; \sigma_{y}^{2}-\frac{\sigma_{y z}^{2}}{\sigma_{z}^{2}}\right)$
$p(x, y \mid z)=N\left(\left[\begin{array}{c}\mu_{x}+\frac{\sigma_{x z}}{\sigma_{z}^{2}}\left(z-\mu_{z}\right) \\ \mu_{y}+\frac{\sigma_{y z}}{\sigma_{z}^{2}}\left(z-\mu_{z}\right)\end{array}\right] ;\left[\begin{array}{cc}\sigma_{x}^{2} & \sigma_{x y} \\ \sigma_{x y} & \sigma_{y}^{2}\end{array}\right]-\frac{1}{\sigma_{z}^{2}}\left[\begin{array}{c}\sigma_{x z} \\ \sigma_{y z}\end{array}\right]\left[\begin{array}{ll}\sigma_{x z} & \sigma_{y z}\end{array}\right]\right)$

## Joint Mutual Information Between Gaussian Variables - 2

$$
\begin{aligned}
I(X ; Y \mid Z) & =\frac{1}{2}\left[\ln (2 \pi e)+\ln \left(\sigma_{x}^{2}-\frac{\sigma_{x z}^{2}}{\sigma_{z}^{2}}\right)+\ln (2 \pi e)+\ln \left(\sigma_{y}^{2}-\frac{\sigma_{y z}^{2}}{\sigma_{z}^{2}}\right)\right. \\
-2 \ln (2 \pi e) & -\ln \left\lvert\,\left[\begin{array}{cc}
\sigma_{x}^{2} & \sigma_{x y} \\
\sigma_{x y} & \sigma_{y}^{2}
\end{array}\right]-\frac{1}{\sigma_{z}^{2}}\left[\begin{array}{c}
\sigma_{x z} \\
\sigma_{y z}
\end{array}\right]\left[\begin{array}{ll}
\sigma_{x z} & \left.\left.\sigma_{y z}\right] \mid\right]
\end{array}\right.\right. \\
& =\frac{1}{2} \ln \frac{\left(\sigma_{x}^{2}-\frac{\sigma_{x z}^{2}}{\sigma_{z}^{2}}\right)\left(\sigma_{y}^{2}-\frac{\sigma_{y z}^{2}}{\sigma_{z}^{2}}\right)}{\sigma_{x}^{2} \sigma_{y}^{2}-\left(\frac{\sigma_{x z}^{2} \sigma_{y}^{2}+\sigma_{y z}^{2} \sigma_{x}^{2}}{\sigma_{z}^{2}}\right)+\frac{\sigma_{x z}^{2} \sigma_{y z}^{2}}{\sigma_{z}^{4}}-\left(\sigma_{x y}-\frac{\sigma_{x z} \sigma_{y z}}{\sigma_{z}^{2}}\right)^{2}} \\
= & \frac{1}{2} \ln \underbrace{}_{\text {determin ant of } 3 x 3 \text { correlation coeff.matrix } \geq 0} \frac{\left(1-\rho_{x z}^{2}\right)\left(1-\rho_{y z}^{2}\right)}{1-\rho_{x y}^{2}-\rho_{y z}^{2}-\rho_{x z}^{2}+2 \rho_{x y} \rho_{x z} \rho_{y z}} \quad \begin{array}{l}
I(X ; Y \mid Z)=0 \text { if } \rho_{x z}=1 \\
\text { and } I(X ; Y ; Z)=I(X ; Y)
\end{array}
\end{aligned}
$$

So, $I(X ; Y ; Z)=I(X ; Y)-I(X ; Y \mid Z)=-\frac{1}{2} \ln \frac{\left(1-\rho_{x y}^{2}\right)\left(1-\rho_{x z}^{2}\right)\left(1-\rho_{y z}^{2}\right)}{1-\rho_{x y}^{2}-\rho_{y z}^{2}-\rho_{x z}^{2}+2 \rho_{x y} \rho_{x z} \rho_{y z}}$
By symmetry, $I(X ; Z \mid Y)=\frac{1}{2} \ln \frac{\left(1-\rho_{x y}^{2}\right)\left(1-\rho_{y z}^{2}\right)}{1-\rho_{x y}^{2}-\rho_{y z}^{2}-\rho_{x z}^{2}+2 \rho_{x y} \rho_{x z} \rho_{y z}}$

$$
I(Y ; Z \mid X)=\frac{1}{2} \ln \frac{\left(1-\rho_{x y}^{2}\right)\left(1-\rho_{x z}^{2}\right)}{1-\rho_{x y}^{2}-\rho_{y z}^{2}-\rho_{x z}^{2}+2 \rho_{x y} \rho_{x z} \rho_{y z}}
$$

## Kullback-Leibler (KL) Divergence

■ Kullback-Leibler (K-L) Divergence ("Relative Entropy")

- A measure of the distance between two pmfs, $p(x)$ and $q(x)$

$$
K L(p(x) \| q(x))=D(p(x) \| q(x))=E_{p(x)}\left[\log _{2} \frac{p(X)}{q(X)}\right]=\sum_{x \in\left\{x, x_{2}, \ldots x_{n 1}\right\}} p(x) \log _{2} \frac{p(x)}{q(x)}
$$

- A measure of inefficiency of assuming that the pmf is $q(x)$ when the true pmf is $p(x)$
- K-L divergence is non-negative, but not necessarily symmetric
- Relation to Mutual Information

$$
\begin{aligned}
I(X, Y) & =E_{p(x, y)}\left[\log _{2} \frac{p(X, Y)}{p(X) p(Y)}\right]=D(p(x, y) \| p(x) p(y))=\sum_{x, y} p(x, y) \log _{2} \frac{p(x, y)}{p(x) p(y)} \\
& =\sum_{x} p(x) \sum_{y} p(y \mid x) \log _{2} \frac{p(y \mid x)}{p(y)}=E_{p(x)}\left[D(p(y \mid x) \| p(y)]=E_{p(y)}[D(p(x \mid y) \| p(x)]\right.
\end{aligned}
$$

## K-L Divergence Between Gaussian PDFs

$K L(p \| q)=-\int p(x) \ln \left\{\frac{q(x)}{p(x)}\right\} d x=-E_{p}\{\ln q(x)\}-H_{p}(x)$
$p(x)=N\left(\mu, \sigma^{2}\right) ; q(x)=N\left(m, s^{2}\right)$
$\ln p(x)=-\frac{1}{2} \ln (2 \pi)-\ln \sigma-\frac{(x-\mu)^{2}}{2 \sigma^{2}}$
$\ln q(x)=-\frac{1}{2} \ln (2 \pi)-\ln s-\frac{(x-m)^{2}}{2 s^{2}}$
$\Rightarrow K L(p \| q)=-E_{p}\left\{\ln \frac{\sigma}{s}+\frac{(x-\mu)^{2}}{2 \sigma^{2}}-\frac{(x-m)^{2}}{2 s^{2}}\right\}=\ln \frac{s}{\sigma}-\frac{1}{2}+\frac{\mu^{2}+\sigma^{2}-2 \mu m+m^{2}}{2 s^{2}}$
$=\ln \frac{s}{\sigma}-\frac{1}{2}+\frac{(\mu-m)^{2}+\sigma^{2}}{2 s^{2}}=\frac{1}{2}\left[\ln \frac{s^{2}}{\sigma^{2}}+\frac{(\mu-m)^{2}}{s^{2}}+\frac{\sigma^{2}}{s^{2}}-1\right]$
Multivariate case ( $n$ vectors): $\frac{1}{2}\left[\ln \frac{\operatorname{det}(S)}{\operatorname{det}(\Sigma)}+(\underline{\mu}-\underline{m})^{T} S^{-1}(\underline{\mu}-\underline{m})+\operatorname{tr}\left(S^{-1} \Sigma\right)-n\right]$
Minimum when $\underline{\mu}=\underline{m}$ and $S=\Sigma \Rightarrow$ moment matching; $p(x)=q(x)$

## K-L Divergence and MI for MVN Variables

$$
\begin{aligned}
& p(\underline{x}, \underline{y})=N\left(\left[\begin{array}{l}
\underline{\mu}_{x} \\
\underline{\mu}_{y}
\end{array}\right] ;\left[\begin{array}{ll}
\Sigma_{x x} & \Sigma_{x y} \\
\Sigma_{x y}^{T} & \Sigma_{y y}
\end{array}\right]\right) \\
& \Rightarrow p(\underline{x})=N\left(\underline{\mu}_{x}, \Sigma_{x x}\right) ; p(\underline{y})=N\left(\underline{\mu}, \Sigma_{y y}\right) \\
& p(\underline{x} \mid \underline{y})=N\left(\underline{\mu}_{x}+\Sigma_{x y} \Sigma_{y y}^{-1}\left(\underline{y}-\underline{\mu}_{y}\right), \Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right) \\
& I(X ; Y)=E_{p(\underline{y})}[K L(p(\underline{x} \mid \underline{y}) \| p(\underline{x}))] \\
& =\frac{1}{2} E_{p(\underline{y})}\left[\ln \frac{\left|\Sigma_{x x}\right|}{\left|\left(\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \sum_{x y}^{T}\right)\right|}+\left\|\Sigma_{x y} \Sigma_{y y}^{-1}\left(\underline{y}-\underline{\mu}_{y}\right)\right\|_{\Sigma_{x x}-1}+\operatorname{tr}\left(\Sigma_{x x}^{-1}\left(\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right)\right)-\operatorname{dim}(x)\right] \\
& =\frac{1}{2}\left[\ln \frac{\left|\Sigma_{x x}\right|}{\left|\left(\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right)\right|}+\operatorname{tr}\left(\Sigma_{y y}^{-1} \Sigma_{x y}^{T} \Sigma_{x x}^{-1} \Sigma_{x y}\right)+\operatorname{tr}\left(I_{\mathrm{dim}(x)}-\Sigma_{x x}^{-1} \Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right)-\operatorname{dim}(x)\right] \\
& =\frac{1}{2} \ln \frac{\left|\Sigma_{x x}\right|}{\left|\left(\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right)\right|}=-\frac{1}{2} \ln \left|\left(I_{\operatorname{dim}(x)}-\Sigma_{x x}^{-1} \Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{x y}^{T}\right)\right| \because \operatorname{tr}(A B)=\operatorname{tr}(B A) \& \operatorname{tr}\left(I_{\operatorname{dim}(x)}\right)=\operatorname{dim}(x) \\
& x, y \operatorname{scalars}: \Sigma_{x x}=\sigma_{x}^{2} ; \Sigma_{x y}=\rho_{x y} \sigma_{x} \sigma_{y} ; \Sigma_{y y}=\sigma_{y}^{2} \\
& \left.I(X ; Y)=-\frac{1}{2}\left[\ln \left(1-\rho_{x y}^{2}\right)\right)\right]=\ln \frac{1}{\sqrt{1-\rho_{x y}^{2}}} \\
& \begin{array}{l}
I(X ; Y \mid Z)=E_{p(z)}\{K L(p(x, y \mid z) \| p(x \mid z) p(y \mid z))\} \\
\operatorname{Try} \text { it for Gaussian! }
\end{array}
\end{aligned}
$$

## Classification vs. Regression

- Classification vs. Regression
- Set of input variables $\underline{x}$....... observed datalfeatures
- Target variables $\underline{z}$.... What you are interested in inferring

$\underline{x}$ can be discrete or continuous features or data
The goal is to evaluate:
$\bullet$ In classification $P(z=k \mid \underline{x}) \Rightarrow$ It is a probability (posterior probability)
$\bullet$ In regression $p(\underline{z} \mid \underline{x}) \quad \Rightarrow \mathrm{It}$ is a density (posterior density)


## Binary Classification

$\square$ Let us consider a binary classification case first:

$$
z \in(0,1) \quad \text { Likelihood Function }
$$

$$
P(z=1 \mid \underline{x})=\frac{p(\underline{x} \mid z=1) P(z=1)}{p(\underline{x})} \quad \text { (Here, } \underline{x} \text { is continuous.) }
$$

$$
=\frac{p(\underline{x} \mid z=1) P(z=1)}{p(\underline{x} \mid z=1) P(z=1)+p(\underline{x} \mid z=0) P(z=0)}
$$

\[

\]

## Sigmoid Function

$$
p(z=1 \mid \underline{x})=\frac{1}{1+e^{-y}}=g(y) \quad \text { where } \quad y=\ln \frac{p(\underline{x} \mid z=1)}{p(\underline{x} \mid z=0)}+\ln \frac{P(z=1)}{P(z=0)}
$$

If $y=\underline{w}^{T} \underline{x}$, then this corresponds to a simple neuron with a logistic sigmoid nonlinearity. More generally,


$$
g(y, \beta)=\frac{1}{1+e^{-\beta y}}
$$

- What happens as $\beta \rightarrow \infty$ or $\beta \rightarrow 0$ ?
- Is $g(y, \beta)$ convex? NO. Why?

We will show that for the so-called exponential family of densities, this form for $y=\underline{w}^{T} \underline{x}$ is valid.

## First and Second Derivative of Sigmoid



$$
\begin{aligned}
& g(y, \beta)=\frac{1}{1+e^{-\beta y}} \\
& \frac{d g(y, \beta)}{d y}=\beta g(1-g) \\
& \frac{d^{2} g(y, \beta)}{d y^{2}}=\beta^{2} g(1-g)(1-2 g)
\end{aligned}
$$

$$
\begin{aligned}
& \frac{d g(y, \beta)}{d y} \text { nonnegative } \Rightarrow \text { monotonic in } y \\
& \frac{d^{2} g(y, \beta)}{d y^{2}} \text { can be positive or negative } \Rightarrow \text { not convex }
\end{aligned}
$$

## Discriminant Function

- $y$ is also called the discriminant function.

$$
p(z=1 \mid \underline{x})=\frac{1}{1+e^{-y}}=g(y)
$$

Any function that can be used to decide a class membership

$$
y=0 \Rightarrow P(z=1 \mid \underline{x})=P(z=0 \mid \underline{x}) \Rightarrow \text { decision boundary }
$$



Even for a non-exponential family, it is OK to use a logistic function, but they have a nonlinear form of $y$

Multi - layer Perceptrons (MLP)
$y=g(\underline{x}, \underline{w}) \Rightarrow$
Radial Basis Functions (RBF)

## Regression

$\square$ Regression
$p(\underline{z} \mid \underline{x}) \Rightarrow$ MMSE estimate is the conditional mean.

$$
E(\underline{z} \mid \underline{x})=\int_{\underline{z}} \underline{z} p(\underline{z} \mid \underline{x}) d \underline{z}
$$

- For Linear Regression: $E(\underline{z} \mid \underline{x})=W \underline{x}=\underline{y}=\underline{\hat{z}}$

$$
\underline{\underline{\hat{z}}}=\underline{y}=\sum_{j=0}^{p} \underline{w}_{j} x_{j} \Rightarrow \hat{z}_{i}=y_{i}=\sum_{j=0}^{p} w_{i j} x_{j} ; i=1,2, . ., m
$$

- ADALINE: Adaptive Linear Element
- In general, we use nonlinear regression functions

$$
\underline{y}=\sum_{j=0}^{p} \underline{w}_{j} \phi_{j}(\underline{x}) \Rightarrow y_{i}=\sum_{j=0}^{p} w_{i j} \phi_{j}(\underline{x}) ; i=1,2, \ldots, m
$$

* How to select basis functions or kernels, $\phi_{j}(\underline{x})$ ?


## RBF Networks \& MLP

$\square$ RBF Networks
Radial basis function networks using Gaussian mixtures

$$
\phi_{j}(\underline{x})=\exp \left\{-\frac{1}{2}(\underline{x}-\underline{\mu})^{T} \Sigma^{-1}(\underline{x}-\underline{\mu})\right\}
$$

- Select $\mu_{j}, \Sigma_{j}$
- Optimize over $\left\{w_{i j}\right\}$
$\Rightarrow$ Local Approximation
$\square$ MLP: Generalize single layer perceptron

$$
\begin{array}{rlr}
y_{i}(\underline{x}) & =g\left\{\sum_{j=1}^{M} w_{i j} g\left(\underline{\tilde{w}}_{j}^{T} \underline{x}\right)\right\} \quad \underline{x}=\left[-1, x_{1}, x_{2}, \ldots x_{p}\right]^{T} \\
& =g\left\{w_{i}^{\mathrm{T}} \underline{g}(\underline{x}, \underline{\tilde{w}})\right\} &
\end{array}
$$

Model of this form can approximate continuous functions to arbitrary accuracy -- universal function approximators.

## Example Nonlinearities - 1

- Rectifier or rectified linear unit (ReLU). Popular in deep learning

$$
g(y)=\max (0, y)=(y)^{+}
$$

- Hyperbolic tangent

$$
g(y)=\tanh (y)=\left(e^{y}-e^{-y}\right) /\left(e^{y}+e^{-y}\right)=\left(e^{2 y}-1\right) /\left(e^{2 y}+1\right)=\sigma(2 y)-\sigma(-2 y)
$$

- Sigmoid

$$
g(y)=1 /\left(1+e^{-y}\right)=\sigma(y)
$$

- Softmax (for multiple classes as an output unit)

$$
g\left(y_{i}\right)=e^{v_{i}} / \sum_{k=1}^{c} e^{y_{k}} ; \sum_{i=1}^{c} g\left(y_{i}\right)=1 ; g\left(y_{i}\right)>0
$$

- RBF unit: used with features in SVM, RBF,...

$$
\phi_{j}(\underline{x})=\exp \left\{-\frac{1}{2}(\underline{x}-\underline{\mu})^{T} \Sigma^{-1}(\underline{x}-\underline{\mu})\right\}
$$

## Example Nonlinearities - 2

- Softplus (function approximation and Restricted Boltzmann Machine (RBM))

$$
\begin{aligned}
& g(y)=\ln \left(1+e^{y}\right)=-\ln \sigma(-y)=-\ln (1-\sigma(y)) \\
& \text { soft approximation to } y^{+}=\max (0, y)
\end{aligned}
$$

- Hard tanh or unit saturation block

$$
g(y)=\max (-1, \min (1, y))
$$

- Absolute value

$$
\mathrm{g}(y)=|y|
$$

- Maxout (used with features)

$$
g(\underline{x})=\max \left(\underline{w}_{i}^{T} \underline{x}+w_{o i}\right)
$$

## RBF Networks \& MLP

These can be extended to multiple layers
RBF: Mixture of exponentials

## MLP

RBF and MLP belong to the category of Feedforward Neural Networks $\Rightarrow$ Acyclic Graphs
$\otimes$ Can we get insight into what is going on within MLP and RBF?
$\otimes$ Can we extract rules from MLP and RBF?
$\otimes$ Alternately, can we represent $y_{\mathrm{i}}$ as a function of non-overlapping basis functions

## 1

- Decision tree classifiers
- Instance-based classifiers (case-based AI reasoning)


## Recurrent NNs \& Density Estimation

$\square$ Recurrent Neural Networks:
What if NNs have cycles? ( $\Rightarrow$ recurrent Neural Networks)

- Hopfield NNs
- Boltzmann and Helmholtz Machines
- Dynamic Neural Networks
$\square$ Density Estimation:

$$
\begin{aligned}
& p(\underline{z} \mid \underline{x})=\frac{p(\underline{x}, \underline{z})}{p(\underline{x})}=\frac{p(\underline{x} \mid \underline{z}) p(\underline{z})}{p(\underline{x})} \\
& p(\underline{x})=\int_{\underline{z}} p(\underline{x}, \underline{z}) d \underline{z} \text { for continuous } \underline{z} ;\left(\text { or) } \sum_{\underline{z}} p(\underline{x}, \underline{z}) \text { for discrete } \underline{z}\right.
\end{aligned}
$$

Two key questions:

- Should we model $p(\underline{x}, z)$ or $p(z \mid \underline{x})$ directly?
- What are the weights? How do you learn them?


## Generative versus Discriminative

## Generative

- Models $p(\underline{x}, \underline{z})$ from which $p(\underline{x}), p(\underline{z})$ and $p(\underline{x} \mid \underline{z})$ follow
- Easy to fit
- Easy to add new classes
- Handles missing data easily
- Handles unlabeled training data
- Can infer $p(\underline{x} \mid \underline{z}) \ldots$ inverse problem
- Hard to preprocess features
- Can give extreme predictions
- Less accurate in practice


## Discriminative

- Models $p(\underline{z} \mid \underline{x})$ directly ... flexibility
- Needs optimization
- Retrain to add new classes
- Not easy to handle missing data
- Harder to handle unlabeled training data
- Not possible to infer $p(\underline{x} \mid \underline{z})$
- Can preprocess features (e.g., kernels)
- Provides calibrated predictions
- More accurate in practice


## Generative \& Discriminative Methods

## Generative

- Classifiers
- Discriminant Analysis (DA)
- Naïve Bayes (NB)
- Tree-augmented NB
- K-nearest neighbor
- Infinite Mixture DA


## Discriminative

- Classifiers (C) \& Regression (R)
- Logistic Regression (LR).. C only
- Sparse Linear/Logistic Regression
- Mixture of experts
- MLP/RBF
- Conditional Random Fields (CRF) ..C
- Decision \& Regression Trees
- Boosting
- Sparse kernelized linear/logistic
- Relevance vector machine (RVM)
- Support vector machine (SVM)
- Gaussian processes (GP)
- Smoothing splines ... R only


## Popular Models of $p(\underline{x} \mid \underline{z})$ and $p(\underline{z})$

$\square p(\underline{x})$ is a Gaussian mixture

$$
p(\underline{x} \mid \underline{z}) \text { is multivariate Gaussian and } p(\underline{z}) \text { is discrete }
$$

$\square p(\underline{x})$ is a mixture of multinomials
$p(\underline{x} \mid \underline{z})$ is a product of discrete distributions and $p(\underline{z})$ is discrete
$\square$ Factor analysis
$p(\underline{x} \mid \underline{z})$ is a product of Gaussians and $p(\underline{z})$ is a product of Gaussians
$\square$ Probabilistic independent component analysis (ICA) $p(\underline{x} \mid \underline{z})$ is a product of Gaussians and $p(\underline{z})$ is a product of Laplace
$\square$ Multinomial PCA
$p(\underline{x} \mid \underline{z})$ is a product of discrete and $p(\underline{z})$ is a product of Gaussian
$\square$ Sigmoid belief net
$p(\underline{x} \mid \underline{z})$ is a product of Bernoulli and $p(\underline{z})$ is a product of Bernoulli

## MAP \& ML Learning

So, among the learning processes, the difference is only in the data used for learning.
$D=$ data
MAP: $\quad w=\arg \max p(w \mid D)$
ML: $\quad w=\arg \max _{w} p(D \mid w)$
$p(w \mid D)=\frac{p(D \mid w) p(w)}{p(D)}$
$J(w)=-\ln p(w \mid D)$
$=-\ln p(D \mid w)-\ln p(w)+\ln p(D)$
Independent of $w$
$J(w)=-\ln p(D \mid w)-\ln p(w)$
$p(w)$ uniform $\Rightarrow M A P=M L$

## Supervised Learning

For Supervised Learning (assuming independent data sets)

$$
\begin{aligned}
& p(D \mid w)=\prod_{n=1}^{\mathrm{N}} p\left(z_{n}, \underline{x}_{n} \mid w\right) \\
&=\prod_{\mathrm{n}=1}^{\mathrm{N}} p\left(z_{n} \mid \underline{x}_{n}, w\right) p\left(\underline{x}_{n} \mid w\right) \\
&=\prod_{\mathrm{n}=1}^{\mathrm{N}} p\left(z_{n} \mid \underline{x}_{n}, w\right) p\left(\underline{x}_{n}\right) \\
&- \ln [p(D \mid w)]=-\sum_{\mathrm{n}=1}^{\mathrm{N}}\left\{\ln p\left(z_{n} \mid \underline{x}_{n}, w\right)+\ln p\left(\underline{x}_{n}\right)\right\} \\
& J(w)=-\ln \mathrm{p}(\mathrm{D} \mid \mathrm{w})=\frac{1}{2} \sum_{n=1}^{N}\left\|z_{n}-g\left(\underline{x}_{n}, w\right)\right\|^{2} \Leftarrow \text { Gaussian Unity Covariance }
\end{aligned}
$$

Thus, determining weights in supervised learning is a function minimization problem!

## Classification Problem

For classification problems:

$$
\begin{aligned}
\ln P\left(z_{n} \mid \underline{x}_{n}, \underline{w}\right) & =z_{n} \ln \underbrace{P\left(z_{n}=1 \mid \underline{x}_{n}, \underline{w}\right)}_{g\left(y_{n}\right)}+\left(1-z_{n}\right) \ln \underbrace{P\left(z_{n}=0 \mid \underline{x}_{n}, \underline{w}\right)}_{1-g\left(y_{n}\right)} \\
& =z_{n} \ln g\left(y_{n}\right)+\left(1-z_{n}\right) \ln \left(1-g\left(y_{n}\right)\right) \\
g\left(y_{n}\right)=P\left(z_{n}\right. & \left.=1 \mid \underline{x}_{n}, \underline{w}\right)=\frac{1}{1+e^{-\beta y_{n}}} ; y_{n}=\underline{w}^{T} \underline{x}_{n} \quad \begin{array}{l}
\text { Cross-entropy } \\
\text { error function }
\end{array}
\end{aligned}
$$

$$
\frac{\partial g\left(y_{n}\right)}{\partial y_{n}}=\frac{\partial}{\partial y_{n}}\left(\frac{1}{1+e^{-\beta y_{n}}}\right)=\frac{\beta e^{-\beta y_{n}}}{\left(1+e^{-\beta y_{n}}\right)^{2}}=\beta g\left(y_{n}\right)\left(1-g\left(y_{n}\right)\right)=\beta g_{n}\left(1-g_{n}\right)
$$

Gradient has a nice (and unified) form: for regression

$$
\frac{\partial J}{\partial w}=\nabla_{w} J=-\sum_{n=1}^{N}\left\{z_{n}-g\left(\underline{x}_{n}, w\right)\right\} \nabla_{w} g=-\beta \sum_{n=1}^{N}\left(z_{n}-g_{n}\right) g_{n}\left(1-g_{n}\right) \underline{x}_{n}
$$

## Gradient for Classification

## $\square$ Classification

$$
\begin{aligned}
& \frac{\partial J}{\partial w}=-\sum_{n=1}^{N}\left[\frac{z_{n}}{g_{n}} \frac{\partial g_{n}}{\partial y_{n}} \frac{\partial y_{n}}{\partial w}-\frac{\left(1-z_{n}\right)}{\left(1-g_{n}\right)} \frac{\partial g_{n}}{\partial y_{n}} \frac{\partial y_{n}}{\partial w}\right] \\
&=-\sum_{n=1}^{N} \frac{z_{n}-g_{n}}{g_{n}\left(1-g_{n}\right)} \frac{\partial g_{n}}{\partial y_{n}} \frac{\partial y_{n}}{\partial w} \\
& \frac{\partial g_{n}}{\partial y_{n}}=\frac{\partial}{\partial y_{n}}\left(\frac{1}{1+e^{-\beta y_{n}}}\right)=\frac{\beta e^{-\beta y_{n}}}{\left(1+e^{-\beta y_{n}}\right)^{2}}=\beta g_{n}\left(1-g_{n}\right) \\
& \text { So, } \frac{\partial J}{\partial w}=-\beta \sum_{n=1}^{N}\left\{z_{n}-g_{n}\right) \frac{\partial y_{n}}{\partial w} \\
& \text { If } \quad y_{n}=\underline{w}^{T} \underline{x}_{n} \quad \frac{\partial J}{\partial w}=-\beta \sum_{n=1}^{N}\left(z_{n}-g_{n}\right) \underline{x}_{n} \quad \begin{array}{c}
\text { Minor difference with } \\
\text { regression case }
\end{array} \\
& \Rightarrow \text { Calculation of gradient is simple } \\
& \Rightarrow \text { Note the summation } \\
& \Rightarrow \text { With minor variations, leads to back propagation algorithm }
\end{aligned}
$$

## Hessian

－Classification

$$
\frac{\partial J}{\partial w}=\nabla_{w} J=-\beta \sum_{n=1}^{N}\left\{z_{n}-g_{n}\right) \underline{x}_{n} \Rightarrow \frac{\partial^{2} J}{\partial w^{2}}=\nabla_{w}^{2} J=\beta^{2} \sum_{n=1}^{N} g_{n}\left(1-g_{n}\right) \underline{x}_{n} \underline{x}_{n}^{T}
$$

$\square$ Regression

## Cross－entropy error

 function is convex$$
\begin{aligned}
& \frac{\partial J}{\partial w}=\nabla_{w} J=-\beta \sum_{n=1}^{N}\left(z_{n}-g_{n}\right) g_{n}\left(1-g_{n}\right) \underline{x}_{n} \\
& \frac{\partial^{2} J}{\partial w^{2}}=\nabla_{w}^{2} J=\beta^{2} \sum_{n=1}^{N}\left[\left(z_{n}-g_{n}\right) g_{n}\left(1-g_{n}\right)\left(2 g_{n}-1\right)+\left(g_{n}\left(1-g_{n}\right)\right)^{2}\right] \underline{x}_{n} \underline{x}_{n}^{T}
\end{aligned}
$$

## - Optimization Algorithms - 1

$\square$ Optimization Algorithms

$$
\begin{aligned}
& \underline{w}_{0} \rightarrow \underline{w}_{1} \rightarrow \ldots \ldots \ldots \rightarrow \underline{w}^{*} \\
& \underline{w}_{k+1}=\underline{w}_{k}-\eta H \nabla_{w} J
\end{aligned}
$$

$H=\left\{\begin{array}{l}I \Rightarrow \text { SD or Gradient method } \\ {\left[\nabla_{w}^{2} J\right]^{-1}}\end{array} \quad \Rightarrow\right.$ Newton


$$
\approx\left(\nabla_{w}^{2} J+v I\right)^{-1}
$$

## Optimization Algorithms - 2

- Batch versus incremental algorithms
- Would like to estimate $\operatorname{Var}\left(y_{i}\right)$
- Prune connections to simplify the architecture
$\otimes$ How long to converge, how many training patterns,...
- Optimization for ill-conditioned problems (ridge regression) Use regularization

$$
\begin{aligned}
& \tilde{J}=J+\frac{1}{2} v \underline{w}^{T} \underline{w} \quad \text { Regularization combats overfitting } \\
\Rightarrow & \nabla^{2} \tilde{J}=\nabla_{w}^{2} J+v I \ldots \text { (Levenberg-Marquardt) }
\end{aligned}
$$

- $L_{l}$-regularization (good for feature selection)

$$
\tilde{J}=J+v\|\underline{w}\|_{1} \bullet \quad \text { Basis pursuit denoising (BPDN) }
$$

- Lasso (Least Absolute Shrinkage \& Selection Operator)


## Optimization Algorithms - 3

- Combined $L_{l^{-}}$and $L_{2}$-regularization

$$
\tilde{J}=J+v_{1}\|\underline{w}\|_{1}+\frac{1}{2} v_{2} \underline{w}^{T} \underline{w} \quad \text { Elastic Net }
$$

- Bridge regression

$$
\tilde{J}=J+v_{1} \sum_{i}\left|w_{i}\right|^{b} ; b>0
$$

- Coordinate descent or probabilistic relaxation methods (EM, variational Bayes, belief propagation, empirical Bayes (Type II ML,...))
- Numerous local minima $\Rightarrow$ use global search techniques

Global Search Techniques


| Unsupervised Learning | Reinforcement Learning |
| :--- | :---: |
| Density Estimation |  |
| Non-parametric techniques | Application of Approximate |
| Kohonen SOFM, LVQ, ART | Dynamic Programming |
| Clustering (K-means, GMM, |  |
| Hierarchical,...) |  |

## Summary

- Provide a systematic account of the major topics in ML based on fundamental mathematical principles
- Focus on pattern classification and optimization applications
$\square$ Treat classification, regression, density estimation problems in a unified way
- Generative \& Discriminative classification \& regression
- Learning from data: Supervised, unsupervised, semisupervised, reinforcement, active
- Use of a variety of optimization techniques for learning weights

