# CSC580: Probabilistic Graphical Models 

Final Exam Review

Jason Pacheco

## Final Exam

- Similar format to Midterm but take-home
- Some students need to take it early so I will release in the next day or two (Due: 12/13)
-6+1 Questions
- 1 of these is only for CSC580 students
- No coding (might use minimal code for one problem)


## Random Variables

(Informally) A random variable is an unknown quantity that maps events to numeric values.

Example X is the sum of two dice with values,

$$
X \in\{2,3,4, \ldots, 12\}
$$

Example Flip a coin and let random variable Y represent the outcome,

$$
Y \in\{\text { Heads, Tails }\}
$$

## Random Variables and Probability

Capitol letters represent random variables

Lowercase letters are realized values
$X=x$ is the event that X takes the value x

Example Let X be the random variable ( RV ) representing the sum of two dice with values,

$$
X \in\{2,3,4, \ldots, 12\}
$$

$X=5$ is the event that the dice sum to 5 .

## Probability Mass Function

A function $p(X)$ is a probability mass function (PMF) of a discrete random variable if the following conditions hold:
(a) It is nonnegative for all values in the support,

$$
p(X=x) \geq 0
$$

(b) The sum over all values in the support is 1,

$$
\sum_{x} p(X=x)=1
$$

Intuition Probability mass is conserved, just as in physical mass. Reducing probability mass of one event must increase probability mass of other events so that the definition holds...

## Joint Probability

Definition Two (discrete) RVs X and Y have a joint PMF denoted by $p(X, Y)$ and the probability of the event $\mathrm{X}=\mathrm{x}$ and $\mathrm{Y}=\mathrm{y}$ denoted by $p(X=x, Y=y)$ where,
(a) It is nonnegative for all values in the support,

$$
p(X=x, Y=y) \geq 0
$$

(b) The sum over all values in the support is 1,

$$
\sum_{x} \sum_{y} p(X=x, Y=y)=1
$$

## Joint Probability

Let $X$ and $Y$ be binary $R V$ s. We can represent the joint PMF $p(X, Y)$ as a $2 \times 2$ array (table):


All values are nonnegative

## Joint Probability

Let $X$ and $Y$ be binary $R V$ s. We can represent the joint PMF $p(X, Y)$ as a $2 \times 2$ array (table):


The sum over all values is 1 :
$0.04+0.36+0.30+0.30=1$

## Joint Probability

Let $X$ and $Y$ be binary $R V$ s. We can represent the joint PMF $p(X, Y)$ as a $2 \times 2$ array (table):


$$
P(X=1, Y=0)=0.30
$$

## Tabular Method

Let $X, Y$ be binary $R V$ s with the joint probability table
$P\left(y_{1}\right)=P\left(x_{1}, y_{1}\right)+P\left(x_{2}, y_{1}\right)$
$P\left(y_{2}\right)=P\left(x_{1}, y_{2}\right)+P\left(x_{2}, y_{2}\right)$
[i.e., sum down columns]

For Binomial use K-by-K probability table.

Y

## Tabular Method



## Tabular Method



## Graphical Models

A variety of graphical models can represent the same probability distribution


Bayes Network


Directed Models


Factor Graph


Markov Random Field

Undirected Models

## Graphical Models

A variety of graphical models can represent the same probability distribution


Factor Graph
Markov Random Field

Directed Models

## From Probabilities to Pictures

A probabilistic graphical model allows us to pictorially represent a probability distribution

Graphical Model:

## Probability Model:

$p\left(x_{1}, x_{2}, x_{3}\right)=$

$$
p\left(x_{1}\right) p\left(x_{2}\right) p\left(x_{3} \mid x_{1}, x_{2}\right)
$$



Conditional distribution on each RV is dependent on its parent nodes in the graph

## Discriminative vs Generative modeling

## Discriminative model:

- Only models $P(y \mid x, \theta)$-- i.e. doesn't model data $x$
- Recall linear regression: $y \mid x ; \theta \sim N\left(x^{\top} \theta, \sigma^{2}\right)$
- Logistic regression: $y \mid x ; \theta \sim \operatorname{Bernoulli}\left(\sigma\left(x^{\top} \theta\right)\right)$


Observations


Generative model:

- Models everything including data: $P(k, y)=P(k) P(y \mid k, \theta)$
- e.g., Gaussian mixture model (GMM)
- $\theta=\left(\pi_{k}, \mu_{k}, \Sigma_{k}\right)_{k=1}^{K}$
- $k \sim \operatorname{Categorical}(\pi)$ (hidden), i.e. $P(k=l)=\pi_{l}$
- $y \mid k \sim N\left(\mu_{k}, \Sigma_{k}\right)$



## Barbershop Example

Suppose you go to a barbershop at every last Friday of the month. You want to be able to predict the waiting time. You have collected 12 data points (i.e., how long it took to be served) from the last year: $S=\left\{x_{1}, \ldots, x_{12}\right\}$

- 1. Modeling assumption: $x_{i} \sim$ Gaussian distribution $N(\mu, 1)$
- $p(x ; \mu)=\frac{1}{\sqrt{2 \pi}} \exp \left(-\frac{(x-\mu)^{2}}{2}\right)$
- Observation: this distribution has mean $\mu$
-2. Find the MLE $\hat{\mu}$ from data $S$
- (2.1) write down the neg. log likelihood of the sample

$$
L_{n}(\mu)=-\ln P\left(x_{1}, \ldots, x_{n} ; \mu\right)=12 \ln \sqrt{2 \pi}+\frac{1}{2} \sum_{i=1}^{12}\left(x_{i}-\mu\right)^{2}
$$



Is this a generative or discriminative model?

## Generative model: basic example I (cont'd)

2. Find the MLE $\hat{\mu}$ from data S

- (2.2) compute the first derivative, set it to 0 , solve for $\lambda$ (be sure to check convexity)

$$
L_{n}^{\prime}(\mu)=\sum_{i=1}^{12}\left(x_{i}-\mu\right)=0 \Rightarrow \mu=\frac{x_{1}+\cdots x_{12}}{12} \text { Sample Mean }
$$

3. The learned model $N(\hat{\mu}, 1)$ is yours!


- Simple prediction: e.g., predict the next wait time by $\mathbb{E}_{X \sim N(\hat{\mu}, 1)}[X]$
- which is $\hat{\mu}=\frac{x_{1}+\cdots x_{12}}{12}$

4. (Optional: Model Checking) Generate some data... Does it look realistic?

## Conditional Independence

Recall two RVs $X$ and $Y$ are conditionally independent given $Z$ (or $X \perp Y \mid Z$ ) iff:

$$
p(X \mid Y, Z)=p(X \mid Z)
$$

Idea Apply chain rule with ordering that exploits conditional independencies to simplify the terms

Ex. Suppose $x_{4} \perp x_{1} \mid x_{3}$ and $x_{2} \perp x_{4} \mid x_{1}$ then:

$$
\begin{aligned}
p(x) & =p\left(x_{3}\right) p\left(x_{1} \mid x_{3}\right) p\left(x_{4} \mid x_{1}, x_{3}\right) p\left(x_{2} \mid x_{1}, x_{3}, x_{4}\right) \\
& =p\left(x_{3}\right) p\left(x_{1} \mid x_{3}\right) p\left(x_{4} \mid x_{3}\right) p\left(x_{2} \mid x_{1}, x_{3}\right)
\end{aligned}
$$

Can visualize conditional dependencies using directed acyclic graph (DAG)

What is the joint factorization?


## $\mathbf{p}(\mathbf{a}, \mathbf{b}, \mathbf{c})=\mathbf{p}(\mathbf{a}) \mathbf{p}(\mathbf{b}) \mathbf{p}(\mathbf{c})$



$\circlearrowleft_{c}$

## Are $a$ and $b$ independent $(a \perp b)$ ?




$$
\mathbf{p}(\mathbf{a}, \mathbf{b}, \mathbf{c})=\mathbf{p}(\mathbf{a}) \mathbf{p}(\mathbf{b}) \mathbf{p}(\mathbf{c})
$$

$$
p(a, b, c)=p(a) p(b \mid a) p(c \mid a, b)
$$



Note there are no conditional independencies

## Case one where c is observed



Is $\mathrm{a} \perp \mathrm{b} \mid \mathrm{c} \quad$ ?

## Case one where c is observed



$$
\mathrm{a} \perp \mathrm{~b} \mid \mathrm{c}
$$

$$
\begin{array}{ll}
p(a, b, c)=p(c) p(a \mid c) p(b \mid c) & (\text { what the graph represents in general) } \\
p(a, b \mid c)=p(a \mid c) p(b \mid c) & (\text { with } c \text { observed) }
\end{array}
$$

This is the definition of $a \perp b \mid c$

## Shading \& Plate Notation

Convention: Shaded nodes are observed, open nodes are latent/hidden/unobserved


Plates denote replication of random variables

## Naïve Bayes for supervised learning

- Motivation: supervised learning for classification
- high-dimensional $x=(x(1), \ldots, x(F))$, modeling $P(x \mid y)$ can be tricky
- In general, $P(x \mid y)=P(x(1) \mid y) \cdot P(x(2) \mid x(1), y) \cdot \ldots \cdot P(x(F) \mid x(1), \ldots, x(F-1), y)$
- A modeling assumption: $x(1), \ldots, x(F)$ are conditionally independent given $y$
i.e. for all $i$

$$
x(i) \Perp(x(1), \ldots, x(i-1), x(i+1), \ldots, x(F)) \mid y
$$

(Conditional independence notation: $A \Perp B \mid C$ )

- Equivalently $P(x \mid y)=P(x(1) \mid y) \cdot \ldots P(x(F) \mid y)$



## Recall : Class Preference Prediction

## Define the labeled training dataset $S=\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{m}$

To make this a binary classification we set "Liked" $=\{+2,+1,0\}$ "Nah" $=\{-1,-2\}$


## Naïve Bayes: binary-valued features

Training Data $S=\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{n}$,

$$
x_{i} \in\{0,1\}^{F}
$$



## Generative Story

$y \sim \operatorname{Bernoulli}(\pi)$; for all $j \in[F], x(j) \mid y=c \sim \operatorname{Bernoulli}\left(\theta_{c, j}\right)$ \#parameters $=1+2 F$
Training (denote by $\theta=\left\{\theta_{c, j}\right\}$ )


$$
\begin{gathered}
\max _{\pi, \theta} \sum_{i=1}^{n} \ln P\left(x_{i}, y_{i} ; \pi, \theta\right)=\sum_{i=1}^{n} \ln P\left(y_{i} ; \pi\right)+\sum_{i=1}^{n} \ln P\left(x_{i} \mid y_{i} ; \theta\right) \\
=\max _{\pi} \sum_{i=1}^{n} \ln P\left(y_{i} ; \pi\right)+\max _{\left\{\theta_{0, j}\right\}} \sum_{i: y_{i}=0} \ln P\left(x_{i} \mid y_{i} ; \theta\right)+\max _{\left\{\theta_{1, j}\right\}} \sum_{i: y_{i}=1} \ln P\left(x_{i} \mid y_{i} ; \theta\right)
\end{gathered}
$$

Key observation: optimal $\pi$, optimal $\left\{\theta_{0, j}\right\}$, optimal $\left\{\theta_{1, j}\right\}$ can be found separately
Optimal $\pi: \max _{\pi} \sum_{i=1}^{n} \ln P\left(y_{i} ; \pi\right)=\max _{\pi} n_{0} \ln (1-\pi)+n_{1} \ln (\pi)=>\hat{\pi}=\frac{n_{1}}{n}$

## Learning / Training



Model random data with hyperparameters $\theta$ :

$$
y \sim p(y \mid \theta)
$$

Sometimes we use:

$$
p(y ; \theta)
$$

Given training data:

$$
\left\{y_{i}\right\}_{i=1}^{n} \stackrel{\text { i.i.d. }}{\sim} p(y \mid \theta)
$$

Learn parameters, e.g. via maximum likelihood estimation:

$$
\hat{\theta}^{\mathrm{MLE}}=\arg \max _{\theta} \log p\left(y_{1}, \ldots, y_{n} \mid \theta\right)
$$

Other estimators are possible:

```
We will talk more about MLE in coming weeks
```

- Maximum a posteriori (MAP)
- Minimum mean squared error (MMSE)
- Etc.


## Likelihood (Intuitively)

Suppose we observe $N$ data points from a Gaussian model and wish to estimate model parameters...


Likelihood Principle Given a statistical model, the likelihood function describes all evidence of a parameter that is contained in the data.

## Likelihood Function

Suppose $x_{i} \sim p(x ; \theta)$, then what is the joint probability over N independent identically distributed (iid) observations $x_{1}, \ldots, x_{N}$ ?

$$
p\left(x_{1}, \ldots, x_{N} ; \theta\right)=\prod_{i=1}^{N} p\left(x_{i} ; \theta\right)
$$

- We call this the likelihood function, often denoted $\mathcal{L}_{N}(\theta)$
- It is a function of the parameter $\theta$, the data are fixed
- Measures how well parameter $\theta$ describes data (goodness of fit)

How could we use this to estimate a parameter $\theta$ ?

## Maximum Likelihood

Maximum Likelihood Estimator (MLE) as the name suggests, maximizes the likelihood function.

$$
\hat{\theta}^{\mathrm{MLE}}=\arg \max _{\theta} \mathcal{L}_{N}(\theta)=\prod_{i=1}^{N} p\left(x_{i} ; \theta\right)
$$

Question How do we find the MLE?
Answer Remember calculus...



## Maximum Likelihood

Maximizing log-likelihood makes the math easier (as we will see) and doesn't change the answer (logarithm is an increasing function)

$$
\hat{\theta}^{\mathrm{MLE}}=\arg \max _{\theta} \log \mathcal{L}_{N}(\theta)=\sum_{i=1}^{N} \log p\left(x_{i} ; \theta\right)
$$

Derivative is a linear operator so,

$$
\frac{d}{d \theta} \log \mathcal{L}_{N}(\theta)=\underbrace{\sum_{i=1}^{N} \frac{d}{d \theta} \log p\left(x_{i} ; \theta\right)}_{\begin{array}{c}
\text { One term per data point } \\
\text { Can be computed in parallel } \\
\text { (big data) }
\end{array}}
$$



## Maximum Likelihood

## Example Suppose we have N coin

 tosses with $X_{1}, \ldots, X_{n} \sim \operatorname{Bernoulli}(p)$ but we don't know the coin bias $p$. The likelihood function is,$$
\mathcal{L}_{n}(p)=\prod_{i=1}^{n} p^{x_{i}}(1-p)^{1-x_{i}}=p^{S}(1-p)^{n-S}
$$

where $S=\sum_{i} x_{i}$. The log-likelihood is,


Likelihood function for Bernoulli with $n=20$ and $\sum_{i} x_{i}=12$ heads

$$
\log \mathcal{L}_{n}(p)=S \log p+(n-S) \log (1-p)
$$

Set the derivative of $\log \mathcal{L}_{n}(p)$ to zero and solve,

$$
\hat{p}^{\mathrm{MLE}}=S / n=\frac{1}{n} \sum_{i=1}^{n} x_{i}
$$

Maximum likelihood is equivalent to sample mean in Bernoulli

## Marginal Likelihood

More often, we have a joint distribution with observations $y$, unknown variables $z$, and parameters $\theta$

Marginal likelihood is normalizer of posterior:
$p(z \mid y)=\frac{p(z) p(y \mid z)}{p(y)}$
Bayes' Rule

$$
p(z, y \mid \theta)=p(z \mid \theta) p(y \mid z, \theta)
$$



Prior Likelihood

Need to marginalize out unknown variables, hence the name marginal likelihood:

$$
p(y \mid \theta)=\int p(z \mid \theta) p(y \mid z, \theta) d z=\mathcal{L}(\theta)
$$

Typically, this integral lacks a closed-form solution...so we need to compute approximate MLE solutions

## Example: Gaussian Mixture Model

Model is often specified in terms of unknown parameters


GMM
Low Likelihood


How likely are parameters for observed data?
$\theta=\left\{\pi, \mu_{1}, \sigma_{1}, \ldots, \mu_{K}, \sigma_{K}\right\}$

$$
\mathcal{Y}=\left\{y_{1}, \ldots, y_{N}\right\}
$$

Marginal Likelihood (likelihood function):

$$
p(\mathcal{Y} \mid \theta)=\underbrace{\sum_{z_{1}} \ldots \sum_{z_{N}}} p\left(z_{1}, \ldots, z_{N}, \mathcal{Y} \mid \theta\right)
$$



## Sum over all possible $\mathrm{K}^{\mathrm{N}}$ assignments,

 which we cannot computeIntuition Learn / estimate parameters that assign highest probability (under the model) to data we've observed.

## Lower Bounding Marginal Likelihood

Conditionally-independent model with partial observations...

## Unknowns



$$
\log p(\mathcal{Y} \mid \theta)=\log \sum_{z_{1}} \ldots \sum_{z_{N}} p\left(z_{1}, \ldots, z_{N}, \mathcal{Y} \mid \theta\right)
$$

$$
(\text { Multiply by } \mathbf{q}(\mathbf{z}) / \mathbf{q}(\mathbf{z})=1) \quad=\log \sum_{z} p(z, \mathcal{Y} \mid \theta)\left(\frac{q(z)}{q(z)}\right)
$$

$$
\begin{aligned}
\text { (Definition of Expected Value ) } & =\log \mathbf{E}_{q}\left[\frac{p(z, \mathcal{Y} \mid \theta)}{q(z)}\right] \\
\text { (Jensen's Inequality ) } & \geq \mathbf{E}_{q}\left[\log \frac{p(z, \mathcal{Y} \mid \theta)}{q(z)}\right]
\end{aligned}
$$

$$
1
$$

## Jensen’s Inequality

$$
f(\mathbb{E}[X]) \leq \mathbb{E}[f(X)]
$$

Valid for both discrete (expectations are sums) and continuous (expectations are integrals) random variables, for any convex function $f$.

$\ln (\mathbb{E}[X]) \geq \mathbb{E}[\ln (X)]$


## Expectation Maximization

Find tightest lower bound of marginal likelihood,

$$
\max _{\theta} \log p(\mathcal{Y} \mid \theta) \geq \max _{q, \theta} \mathbf{E}_{q}\left[\log \frac{p(z, \mathcal{Y} \mid \theta)}{q(z)}\right] \equiv \mathcal{L}(q, \theta)
$$

Solve by coordinate ascent...

Initialize Parameters: $\theta^{(0)}$
At iteration $t$ do:
E-Step: $\quad q^{(t)}=\arg \max _{q} \mathcal{L}\left(q, \theta^{(t-1)}\right)$
M-Step: $\quad \theta^{(t)}=\arg \max _{\theta} \mathcal{L}\left(q^{(t)}, \theta\right)$
Until convergence

## E-Step

$$
q^{(t)}(z)=\arg \max _{q} \mathcal{L}\left(q, \theta^{(t-1)}\right) \equiv \mathbf{E}_{q}\left[\log \frac{p\left(z, y \mid \theta^{(t-1)}\right)}{q(z)}\right]
$$

Concave (in $q(z)$ ) and optimum occurs at,

$$
q^{(t)}(z)=p\left(z \mid y, \theta^{(t-1)}\right) \quad \begin{gathered}
\text { Set } q(z) \text { to posterior with } \\
\text { current parameters }
\end{gathered}
$$

Initialize Parameters: $\theta^{(0)}$
At iteration t do:
E-Step:

$$
q^{(t)}(z)=p\left(z \mid y, \theta^{(t-1)}\right)
$$

M-Step: $\quad \theta^{(t)}=\arg \max _{\theta} \mathcal{L}\left(q^{(t)}, \theta\right)$
Until convergence






## Example: Gaussian Mixture Model

$$
\frac{\left.\log p(\mathcal{Y} \mid \pi, \mu, \Sigma) \geq \sum_{n=1}^{N} \sum_{k=1}^{K} q\left(z_{n}=k\right) \log \left\{\pi_{k} \mathcal{N}\left(y_{n} \mid \mu_{k}, \Sigma_{k}\right)\right\}=\mathcal{L}(q, \theta),{ }_{\sim}^{n}\right)}{}
$$

E-Step: $q^{\text {new }}=\arg \max \mathcal{L}\left(q, \theta^{\text {old }}\right)$

$$
q^{\text {new }}\left(z_{n}=k\right)=p\left(z_{n}=k \mid \mathcal{Y}, \mu^{\text {old }}, \Sigma^{\text {old }}, \pi^{\text {old }}\right)
$$

$$
\begin{aligned}
& =\frac{p\left(z_{n}=k, \mathcal{Y} \mid \mu^{\text {old }}, \Sigma^{\text {old }}, \pi^{\text {old }}\right)}{\sum_{j=1}^{K} p\left(z_{n}=j, \mathcal{Y} \mid \mu^{\text {old }}, \Sigma^{\text {old }}, \pi^{\text {old }}\right)} \\
& =\frac{\pi_{k}^{\text {old }} \mathcal{N}\left(y_{n} \mid \mu_{k}^{\text {old }}, \Sigma_{k}^{\text {old }}\right)}{\sum_{j=1}^{K} \pi_{j} \mathcal{N}\left(y_{n} \mid \mu_{j}^{\text {old }}, \Sigma_{j}^{\text {old }}\right)}
\end{aligned}
$$

Commonly refer to $q\left(Z_{n}\right)$ as responsibility

## Example: Gaussian Mixture Model



M-Step: $\quad \theta^{\text {new }}=\arg \max _{\theta} \mathcal{L}\left(q^{\text {new }}, \theta\right)$
Start with mean parameter $\mu_{k}$,

$$
0=\nabla_{\mu_{k}} \mathcal{L}\left(q^{\text {new }}, \theta\right)
$$



$$
\begin{gathered}
0=\sum_{n=1}^{N} \nabla_{\mu_{k}} \mathbf{E}_{z_{n} \sim q^{\text {new }}}\left[\log \mathcal{N}\left(y_{n} \mid \mu_{z_{n}}, \Sigma_{z_{n}}\right)\right] \\
0=-\sum_{n=1}^{N} q^{\text {new }}\left(z_{n}=k\right) \Sigma_{k}\left(y_{n}-\mu_{k}\right) \\
\mu_{k}^{\text {new }}=\frac{1}{N_{k}} \sum_{n=1}^{N} q^{\text {new }}\left(z_{n}=k\right) y_{n} \text { where } N_{k}=\sum_{n=1}^{N} q\left(z_{n}=k\right)
\end{gathered}
$$

## Example: Gaussian Mixture Model

$$
\overbrace{n=1}^{\pi} \log p(\mathcal{Y} \mid \pi, \mu, \Sigma) \geq \sum_{n=1}^{N} \sum_{k=1}^{K} q\left(z_{n}=k\right) \log \left\{\pi_{k} \mathcal{N}\left(y_{n} \mid \mu_{k}, \Sigma_{k}\right)\right\}=\mathcal{L}(q, \theta)
$$

M-Step: $\quad \theta^{\text {new }}=\arg \max _{\theta} \mathcal{L}\left(q^{\text {new }}, \theta\right)$
Repeat for remaining parameters,

$$
\begin{gathered}
\Sigma_{k}^{\text {new }}=\frac{1}{N_{k}} \sum_{n=1}^{N} q\left(z_{n}=k\right)\left(y_{n}-\mu_{k}^{\text {new }}\right)\left(y_{n}-\mu_{k}^{\text {new }}\right)^{T} \\
\pi_{k}^{\text {new }}=\frac{N_{k}}{N}
\end{gathered}
$$

- Solving for mixture weights requires a bit more work
- Need constraint $\sum_{k} \pi_{k}=1$
- Use Lagrange multiplier approach


## Example: Gaussian Mixture Model

M-Step: $\quad \theta^{\text {new }}=\arg \max _{\theta} \mathcal{L}\left(q^{\text {new }}, \theta\right)$
Repeat for remaining parameters,

$$
\begin{gathered}
\Sigma_{k}^{\text {new }}=\frac{1}{N_{k}} \sum_{n=1}^{N} q\left(z_{n}=k\right)\left(y_{n}-\mu_{k}^{\text {new }}\right)\left(y_{n}-\mu_{k}^{\text {new }}\right)^{T} \\
\pi_{k}^{\text {new }}=\frac{N_{k}}{N}
\end{gathered}
$$

- Solving for mixture weights requires a bit more work
- Need constraint $\sum_{k} \pi_{k}=1$
- Use Lagrange multiplier approach


## K-Means Clustering



## Clustering

- Input: $k$ : the number of clusters (hyperparameter)

$$
S=\left\{x_{1}, \ldots, x_{n}\right\}
$$

- Output
- partition $\left\{G_{i}\right\}_{i=1}^{k}$ s.t. $S=U_{i} G_{i}$ (disjoint union).
- often, we also obtain 'centroids'
- Q: what would be a reasonable definition of centroids?



## $k$-means clustering

- Idea: to partition the data, it would be great if someone gives us $k$ reasonable centroids $c_{1}, \ldots, c_{k}$, since then we can partition the data with them.

$$
A(x)=\arg \min _{j \in[k]}\left\|x-c_{j}\right\|_{2}
$$

- But we don't have those centroids => Let's find them with an optimization formulation.

$$
\underset{c_{1}, \ldots, c_{k}}{\operatorname{minimize}} f\left(c_{1}, \ldots, c_{k}\right) \text {, where } f\left(c_{1}, \ldots, c_{k}\right)=\sum_{i=1}^{n} \min _{j \in[k]}\left\|x-c_{j}\right\|_{2}^{2}
$$



## Special case: $k=1$

- $\min _{c_{1}, \ldots, c_{k}} \sum_{i=1}^{n} \min _{j \in[k]}\left\|x_{i}-c_{j}\right\|_{2}^{2}=>\min _{c} \sum_{i=1}^{n}\left\|x_{i}-c\right\|_{2}^{2}$
- Let $F(c)=\sum_{i=1}^{n}\left\|x_{i}-c\right\|_{2}^{2}$ convex; minimizer $c^{*}$ satisfies that $\nabla F\left(c^{*}\right)=0$
$\Rightarrow \sum_{i=1}^{n}\left(x_{i}-c^{*}\right)=0 \Rightarrow c^{*}=\frac{1}{n} \sum_{i=1}^{n} x_{i}$


## For $k \geq 2$

- $\operatorname{minimize}_{c_{1}, \ldots, c_{k}} f\left(c_{1}, \ldots, c_{k}\right)$, where $f\left(c_{1}, \ldots, c_{k}\right)=\sum_{i=1}^{n} \min _{j \in[k]}\left\|x-c_{j}\right\|_{2}^{2} \Rightarrow$ NP-hard even when $d=2$
- K-means algorithm: solve it approximately (heuristic)
- Observation: The chicken-and-egg problem.
- Cluster center location depends on the cluster assignment
- Cluster assignment depends on cluster location
- Very common heuristic (that may or may not be the best thing to do)



## Initialization



Arbitrary/random initialization of $c_{1}$ and $c_{2}$

## Iteration 1


(A) update the cluster assignments.

(B) Update the centroids $\left\{c_{j}\right\}$

## Iteration 2


(A) update the cluster assignments.

(B) Update the centroids $\left\{c_{j}\right\}$

## Iteration 3


(A) update the cluster assignments.

(B) Update the centroids $\left\{c_{j}\right\}$

## Iteration 4


(A) update the cluster assignments.

(B) Update the centroids $\left\{c_{j}\right\}$

## K-means clustering

Input: $k$ : num. of clusters, $S=\left\{x_{1}, \ldots, x_{n}\right\}$
[Initialize] Pick $c_{1}, \ldots, c_{k}$ as randomly selected points from $S$ (see next slides for alternatives) For $\mathrm{t}=1,2, \ldots$, max_iter

- [Assignments] $\forall x \in S, \quad a_{t}(x)=\arg \min _{j \in[k]}\left\|x-c_{j}\right\|_{2}^{2}$
- If $\mathrm{t} \neq 1$ AND $a_{t}(x)=a_{t-1}(x), \forall x \in S$
- break

- [Centroids] $\quad \forall j \in[k], \quad c_{j} \leftarrow \operatorname{average}\left(\left\{x \in S: a_{t}(x)=j\right\}\right)$

Output: $c_{1}, \ldots, c_{k}$ and $\left\{a_{t}\left(x_{i}\right)\right\}_{i \in[n]}$


# Dimensionality Reduction and Principal Component Analysis (PCA) 

## Motivation

## Data often have a lot of redundant information...



Example A dataset consisting of a hand-drawn 3 at random locations and rotations in a $100 \times 100$ pixel image.

Data Dimension $100 \times 100=10,000$ Intrinsic Dimension 3 (X-position, Y-position, Rotation)

## Example : Iris Dataset

Recall that the Iris dataset has 4 features:
sepal length / width, petal length / width...







## Example : Iris Dataset



Data still cluster in a two-dimensional subspace

We can fit model in 2D to reduce complexity, visualize results, etc.

## Linear Dimensionality Reduction



Project data onto a line or plane...

...one of the simplest dimensionality reduction approaches

First, let's review some linear algebra...

## Linear Dimensionality Reduction



Projecting data onto a vector is a simple inner product,

$$
\tilde{x}_{n}=u^{T} x_{n}
$$

We call $u$ the linear subspace

## Linear Dimensionality Reduction

Which choice of subspace is best? And why?


Idea Choose the subspace that captures the most variation in the original data

## Principal Component Analysis (PCA)

Identify directions of maximum variation as subspaces...


...we call each direction a principal component
[ Source: Bishop, C. ]

## Principal Component Analysis (PCA)

First, center the data by subtracting the sample mean,

$$
\bar{x}=\frac{1}{N} \sum_{n=1}^{N} x_{n}
$$

Variance of projected subspace,

$$
\frac{1}{N} \sum_{n=1}^{N}\left(u^{T} x_{n}-u^{T} \bar{x}\right)^{2}
$$



## Maximum Variance Formulation

A little algebra...

$$
\begin{aligned}
\frac{1}{N} \sum_{n=1}^{N}\left(u^{T} x_{n}-u^{T} \bar{x}\right)^{2} & =\frac{1}{N} \sum_{n=1}^{N}\left\{u^{T}\left(x_{n}-\bar{x}\right)\right\}^{2} \quad \text { Pull out } u \\
\text { Quadratic form } & =\frac{1}{N} \sum_{n=1}^{N} u^{T}\left(x_{n}-\bar{x}\right)\left(x_{n}-\bar{x}\right)^{T} u
\end{aligned}
$$

Define: $\quad S=\frac{1}{N} \sum_{n=1}^{N}\left(x_{n}-\bar{x}\right)\left(x_{n}-\bar{x}\right)^{T}$
Then: $\quad \frac{1}{N} \sum_{n=1}^{N}\left(u^{T} x_{n}-u^{T} \bar{x}\right)^{2}=u^{T} S u$

## Maximum Variance Formulation

Find u so that projected variance is maximal...

$$
\max _{u} u^{T} S u
$$

Don't want to cheat with large magnitude $u$, so we add penalty,

$$
\max _{u} u^{T} S u-\lambda u^{T} u
$$

Set the derivative (gradient) to zero and solve...

$$
\begin{array}{ll}
S u-\lambda u=0 & \text { What equation is this? } \\
S u=\lambda u & \mathrm{u} \text { is an eigenvector with } \\
\text { eigenvalue } \lambda
\end{array}
$$

## Recap of Concepts

- Learning a reduced intrinsic dimension is useful for a bunch of reasons
- The easiest approach is to find a linear subspace
- PCA defines the linear subspace as that which maximizes variance of the projected data
- The set of subspaces are defined by the eigenvectors,

$$
\max _{u} u^{T} S u-\lambda u^{T} u
$$

$$
S u=\lambda u
$$

## Eigenstuff

Eigenvectors / values of a matrix solve the equation

$$
S u=\lambda u
$$

- Matrix S may have multiple eigenvectors / values that solve the above equation
- For D-dimensional u can find all vectors in $O\left(D^{3}\right)$ time
- PCA finds M<D vectors with largest eigenvalues
- Can find $\mathrm{M}<\mathrm{D}$ sorted eigenvectors in $\mathrm{O}\left(\mathrm{MD}^{2}\right)$ time
- Note that D can be large!


## Principal Component Analysis (PCA)

How much variance is captured by just the first principal component (i.e. eigenvector with largest eigenvalue)?


Let $u_{1}$ be the first principal component, then variance of first PC is,

$$
\frac{1}{N} \sum_{n}\left\{u_{1}^{T}\left(x_{n}-\bar{x}\right)\right\}^{2}
$$

How much in the second PC?

$$
\frac{1}{N} \sum_{n}\left\{u_{2}^{T}\left(x_{n}-\bar{x}\right)\right\}^{2}
$$

## Explained Variance

How much variance is captured in M < D principal components?


$$
\frac{1}{N} \sum_{m=1}^{M} \sum_{n}\left\{u_{m}^{T}\left(x_{n}-\bar{x}\right)\right\}^{2}
$$

We call this the explained variance of the first M principal components

Divide by total variance to find percentage of the total variance explained by the subspace

## Fully Connected Neural Networks

## Multilayer Perceptron



This is the quintessential Neural Network...
...also called Feed Forward Neural Net or Artificial Neural Net
[ Source: http://neuralnetworksanddeeplearning.com ]

## Handwritten Digit Classification

Classifying handwritten digits is the "Hello World" of NNs


Modified National Institute of Standards and Technology (MNIST) database contains 60k training and 10k test images

Each character is centered in a $28 \times 28=784$ pixel grayscale image



Each image pixel is a number in [0,1] indicated by highlighted color


## Feedforward Procedure



Each node computes a weighted combination of nodes at the previous layer...

$$
w_{1} x_{1}+w_{2} x_{2}+\ldots+w_{n} x_{n}
$$

Then applies a nonlinear function to the result

$$
\sigma\left(w_{1} x_{1}+w_{2} x_{2}+\ldots+w_{n} x_{n}+b\right)
$$

## Multilayer Perceptron



Final layer is typically a linear model...for classification this is a Logistic Regression

$$
\sigma\left(w^{T} x+b\right)=\frac{1}{1+e^{-\left(w^{T} x+b\right)}}
$$

Vector of activations from previous layer

Recall that for multiclass logistic regression with K classes,

$$
p(\mathrm{Class}=k \mid x) \propto \sigma\left(w_{k}^{T} x+b_{k}\right)
$$



$$
\begin{gathered}
784 \times 16+16 \times 16+16 \times 10 \\
\text { weights }
\end{gathered}
$$

$$
\begin{gathered}
16+16+10 \\
\text { biases }
\end{gathered}
$$

## 13,002

Each parameter has some impact on the output...need to tweak (learn) all parameters simultaneously to improve prediction accuracy

## Convolutional Neural Networks

## Convolutional neural networks (CNN)

- A.K.A. ConvNet architecture
- A set of neural network architecture that consists of
- convolutional layers
- pooling layers
- fully-connected (FC) layers



## Convolution for single-channel images

Consider one filter with weights $\left\{w_{i, j}\right\}$ with size $\mathrm{F} \times \mathrm{F}$

- For every F x F region of the image, perform inner product (= element wise product, then sum them all)
- Q: given a w x h image, after convolution with a F x F filter, what is the size of the resulting image?
- Terminologies: filter size, receptive field size, kernel.




## Convolution: Some Intuition

Define the convolution of filter $f$ on image I as:

$$
(I * f)(x)=\sum_{m} \sum_{n} f(x-m, y-n) I(m, n)
$$

Many ML libraries actually implement cross-correlation:

$$
(f * I)(x)=\sum_{m} \sum_{n} f(x, y) I(x+m, y+n)
$$

Learning finds good values for the convolution filter...

## Convolutional layer for multi-channel images

Input: w (width) x h (height) x c (\#channels)

- E.g. $32 \times 32 \times 3$
- 3 channels: $R, G$, and $B$

A convolutional filter on such image is of shape $F$ xFxc

- Only spatial structure in the first two dimensions

- Denoted by $\left\{w_{i, j, k}\right\}$


## Convolutional layer: visual explanation

- Consider one filter with weights $\left\{w_{i, j, k}\right\}$ with $5 \times 5 \times 3$
- Imagine a sliding 3D window.
- Convolution: For every $5 \times 5$ region of the image, perform inner product (= element wise product, then sum them all)
- Then apply the activation function (e.g., ReLU)
- Results in $28 \times 28 \times 1$ - called activation map.
- Now, we can do $K$ of these filters but with different weights $\left\{w_{i, j, k}^{(\ell)}\right\}$ for $\ell \in[K]$ => output is $28 \times 28 \times K$




## Convolutional Layer: More Details

Stride length S

- Skip input regions; Move the sliding window of a filter not by 1 but by S .
- E.g., S=2 means skipping every other 5 by 5 region.

Zero-padding P : add P number of artificial pixels with value 0 around the input image on both sides

- To ensure the spatial dimension is maintained (otherwise, patterns at the corners are not detected well)
- If we use $P=1$, then the activation map will be $30 \times 30$, not $28 \times 28$ in our example!



## Example

Filter

## Padding = Same




Output

| 0.5 | 0 | 0.25 | 0.25 |
| :---: | :---: | :---: | :---: |
| 0 | 1.25 | 0.5 | 0.5 |
| 0 | 0.5 | 0.75 | 1.5 |
| 0.5 | 0.25 | 1.25 | 1 |

outDim $=($ inpDim $) /$ strideDim

## Conyolytional Layer: More Details Stridelenyths Stride length S

- Skip input regions; Move the sliding window of a filter not by 1 but by S .
- E.g., S=2 means skipping every other 5 by 5 region.


## Zero-padding P: add P number of artificial pixels with value 0 arour

 image.

- To ensure the spatial dimension is maintained (otherwise, patterns at the corners are not detected well)
- If we use $P=2$, then the activation map will be 32 by 32 not 28 by 28 in our example!


## Rules (same goes for height)

- W: input volume width, F: filter width (usually, the filter has the same width and height)
- The output width $K=$ floor $((W-F+2 P) / S)+1$
- E.g., $W=32, F=5, P=0, S=1 \quad K=28$
- E.g., $W=32, F=5, P=2, S=1 \quad \Rightarrow \quad K=32$


## Strides and padding: animations

Strides only

Padding only

Strides + Padding


## Supervised learning setup: putting it together



## $k$-nearest neighbors ( $k$-NN): main concept

Training set: $S=\left\{\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right)\right\}$

Inductive bias: given test example $x$, its label should resemble the labels of nearby points

## Function

- input: $x$
- find the $k$ nearest points to $x$ from $S$; call their indices $N(x)$
- output: the majority vote of $\left\{y_{i}: i \in N(x)\right\}$

- For regression, the average.


## k-NN classification example



## $k$-NN classification: pseudocode

- Training is trivial: store the training set
- Test:

Algorithm 3 KNN-Predict(D, $K, \hat{x}$ )
list $\longrightarrow$ : $S \leftarrow$ []
z: for $n=1$ to $N$ do
append to list $\longrightarrow \quad S: S \oplus\left\langle\mathrm{~d}\left(x_{n}, \hat{x}\right), n\right\rangle \quad$ // store distance to training example $n$ ${ }_{4}$ end for
sort in first coordinate ${ }_{5: S} S \leftarrow \operatorname{sorT}(S) \quad$ // put lowest-distance objects first
6: $\hat{y} \leftarrow 0$
7: for $k=1$ to $K$ do
8: $\quad\langle$ dist,$n\rangle \leftarrow S_{k} \quad / / n$ this is the $k$ th closest data point
9: $\quad \hat{y} \leftarrow \hat{y}+y_{n} \quad / /$ vote according to the label for the $n$th training point
${ }_{\text {10: }}$ end for
Majority vote of $\left\{y_{i}: i \in N(x)\right\} \longrightarrow$ return $\operatorname{sIGN}(\hat{y}) \quad / /$ return +1 if $\hat{y}>0$ and -1 if $\hat{y}<0$

- Time complexity (assuming distance calculation takes $O(d)$ time)
- $O(m d+m \log m+k)=O(m(d+\log m))$
- Faster nearest neighbor search: k-d trees, locality sensitive hashing


## Variations

- Classification
- Recall the majority vote rule: $\hat{y}=\arg \max _{y \in\{1, \ldots, C\}} \sum_{i \in N(x)} 1\left\{y_{i}=y\right\}$
- Soft weighting nearest neighbors: $\hat{y}=\arg \max _{y \in\{1, \ldots, C\}} \sum_{i=1}^{m} w_{i} 1\left\{y_{i}=y\right\}$, where $w_{i} \propto \exp \left(-\beta d\left(x, x_{i}\right)\right)$, or $\propto \frac{1}{1+d\left(x, x_{i}\right)^{\beta}}$
- Class probability estimates
- $\hat{P}(Y=y \mid x)=\frac{1}{k} \sum_{i \in N(x)} 1\left\{y_{i}=y\right\}$
- Useful for "classification with rejection"



## Inductive Bias



How would you label the test examples?

## Overfitting vs Underfitting



High training error High test error


Low training error
Low test error

Overfit
(high variance)


Low training error
High test error

## Bayes optimal classifier

Theorem $f_{B O}$ achieves the smallest 0-1 error among all classifiers.

$$
f_{B O}(x)=\arg \max _{y \in \mathcal{Y}} P_{D}(X=x, Y=y)=\arg \max _{y \in \mathcal{Y}} P_{D}(Y=y \mid X=x), \forall x \in \mathcal{X}
$$

Example Iris dataset classification:


Iris Setosa


Iris Versicolor


Iris Virginica


## Bayes error rate: alternative form

$$
\begin{aligned}
L_{D}\left(f_{B O}\right) & =P_{D}\left(Y \neq f_{B O}(X)\right) \\
& =\sum_{x} P_{D}\left(Y \neq f_{B O}(x) \mid X=x\right) P_{D}(X=x) \\
& =\sum_{x}\left(1-P_{D}\left(Y=f_{B O}(x) \mid X=x\right)\right) P_{D}(X=x) \\
& =\sum_{x}\left(1-\max _{y} P_{D}(Y=y \mid X=x)\right) P_{D}(X=x) \\
& =\mathrm{E}\left[1-\max _{y} P_{D}(Y=y \mid X)\right]
\end{aligned}
$$

- Special case: binary classification

$$
\text { - } \begin{aligned}
L_{D}\left(f_{B O}\right)= & \sum_{x} P_{D}\left(Y \neq f_{B O}(x), X=x\right) \\
& =\sum_{x} \min \left(P_{D}(Y=+1, X=x), P_{D}(Y=-1, X=x)\right)
\end{aligned}
$$




## When is the Bayes error rate nonzero?

$$
L_{D}\left(f_{B O}\right)=\sum_{x} \min \left(P_{D}(Y=+1, X=x), P_{D}(Y=-1, X=x)\right)
$$

- Limited feature representation
- Noise in the training data
- Feature noise
- Label noise
- Sensor failure

- Typo in reviews for sentiment classification
- May not be a single "correct" answer

- Inductive bias of the model / learning algorithm


## New measures of classification performance

- True positive rate (TPR)
$=\frac{\mathrm{TP}}{\mathrm{P}}=\frac{P(\hat{y}=+1, y=+1)}{P(y=+1)}$
(aka recall, sensitivity)
- True negative rate $(T N R)=\frac{\mathrm{TN}}{\mathrm{N}}$ (specificity)
- False positive rate (FPR) $=\frac{\mathrm{FP}}{\mathrm{N}}$
- False negative rate $(F N R)=\frac{\mathrm{FN}}{\mathrm{P}}$


$$
P=T P+F N \quad N=F P+T N
$$

- Precision $=\frac{\mathrm{TP}}{\mathrm{P}-\text { called }}=\frac{P(\hat{y}=+1, y=+1)}{P(\hat{y}=+1)}, \mathrm{P}-$ called $=\mathrm{TP}+\mathrm{FP}$


## Linear Regression



Regression Learn a function that predicts outputs from inputs,

$$
y=f(x)
$$

Outputs y are real-valued

Linear Regression As the name suggests, uses a linear function:

$$
y=w^{T} x+b
$$

We will add noise later...

## Linear Regression

Input-output mapping is not exact, so we will add zero-mean Gaussian noise,

$$
y=w^{T} x+\epsilon \quad \text { where } \quad \epsilon \sim \mathcal{N}\left(0, \sigma^{2}\right)
$$

This is equivalent to the likelihood function,

$$
p(y \mid w, x)=\mathcal{N}\left(y \mid w^{T} x, \sigma^{2}\right)
$$



INPUT: X

Because Adding a constant to a Normal RV is still a Normal RV,

$$
z \sim \mathcal{N}(m, P) \quad z+c \sim \mathcal{N}(m+c, P)
$$

In the case of linear regression $z \rightarrow \epsilon$ and $c \rightarrow w^{T} x$

## Great, we're done right?

> We need to fit it to data by learning the regression weights

How to do this?
What makes good weights?
Data - We have this


Don't know these; need to learn them

## Learning Linear Regression Models

There are several ways to think about fitting regression:

- Intuitive Find a plane/line that is close to data
- Functional Find a line that minimizes the least squares loss
- Estimation Find maximum likelihood estimate of parameters

They are all the same thing...

## Learning Linear Regression Models

There are several ways to think about fitting regression:

- Intuitive Find a plane/line that is close to data
- Functional Find a line that minimizes the least squares loss
- Estimation Find maximum likelihood estimate of parameters

They are all the same thing...

## MLE for Linear Regression

Given training data $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{N}$ likelihood function is given by,

$$
\log \prod_{i=1}^{N} p\left(y_{i} \mid x_{i}, w\right)=\sum_{i=1}^{N} \log p\left(y_{i} \mid x_{i}, w\right)
$$

Recall that the likelihood is Gaussian:

$$
p(y \mid w, x)=\mathcal{N}\left(y \mid w^{T} x, \sigma^{2}\right)
$$



So MLE maximizes the log-likelihood over the whole data as,

$$
w^{\mathrm{MLE}}=\arg \max _{w} \sum_{i=1}^{N} \log \mathcal{N}\left(y_{i} \mid w^{T} x_{i}, \sigma^{2}\right)
$$

## MLE of Gaussian Mean

Assume data are i.i.d. univariate Gaussian,

$$
p(\mathcal{Y} \mid \mu)=\prod_{i=1}^{N} \mathcal{N}\left(y_{i} \mid \mu, \sigma^{2}\right)
$$

Log-likelihood function:

$$
\mathcal{L}(\mu)=\sum_{i=1}^{N} \log \left(\frac{1}{\sqrt{2 \pi \sigma^{2}}} \exp \left(-\frac{1}{2}\left(y_{i}-\mu\right)^{2} \sigma^{-2}\right)\right)
$$

$\begin{aligned} & \text { Constant doesn't } \\ & \text { depend on mean }\end{aligned} \longleftarrow$ const. $-\frac{1}{2} \sum_{i=1}^{N}\left(\left(y_{i}-\mu\right)^{2} \sigma^{-2}\right), ~(1) ~$
MLE doesn't change when we:

1) Drop constant terms (in $\mu$ )
2) Minimize negative log-likelihood

MLE estimate is least squares estimator:

$$
\mu^{\mathrm{MLE}}=-\frac{1}{2 \sigma^{2}} \arg \max _{\mu} \sum_{i=1}^{N}\left(y_{i}-\mu\right)^{2}=\arg \min _{\mu} \sum_{i=1}^{N}\left(y_{i}-\mu\right)^{2}
$$

## MLE of Linear Regression



## Substitute linear regression prediction into MLE solution and we have,

$$
\min _{w} \sum_{i=1}^{N}\left(y_{i}-w x_{i}\right)^{2}
$$

So for Linear Regression,
MLE = Least Squares Estimation

## MLE of Linear Regression

Using previous results, MLE is equivalent to minimizing squared residuals,

$$
\min _{w} \sum_{i=1}^{N}\left(y_{i}-w^{T} x_{i}\right)^{2}=\left\|\mathbf{y}-w^{T} \mathbf{X}\right\|^{2}
$$

Some slightly more advanced linear algebra gives us a solution,
[ Image: Murphy, K. (2012) ]


$$
w=\left(\mathbf{X}^{T} \mathbf{X}\right)^{-1} \mathbf{X}^{T} \mathbf{y}
$$

Derivation a bit involved for lecture but...

- We know it has a closed-form and why
- We can evaluate it

Ordinary Least Squares (OLS) solution

## Basis Functions

- A basis function can be any function of the input features X
- Define a set of $m$ basis functions $\phi_{1}(x), \ldots, \phi_{m}(x)$
- Fit a linear regression model in terms of basis functions,

$$
y=\sum_{i=1}^{m} w_{i} \phi_{i}(x)=w^{T} \phi(x)
$$

- Regression model is linear in the basis transformations
- Model is nonlinear in the data $X$


## Kernel Functions

A kernel function is an inner-product of some basis function computed on two inputs

$$
k\left(x, x^{\prime}\right)=\boldsymbol{\phi}(x)^{\mathrm{T}} \boldsymbol{\phi}\left(x^{\prime}\right)=\sum_{i=1}^{M} \phi_{i}(x) \phi_{i}\left(x^{\prime}\right)
$$

A consequence is that kernel functions are non-negative realvalued functions over a pair of inputs,

$$
\kappa\left(x, x^{\prime}\right) \in \mathbb{R} \quad \kappa\left(x, x^{\prime}\right) \geq 0
$$

Kernel functions can be interpreted as a measure of distance between two inputs

## Kernel Functions

Example The linear basis $\phi(x)=x$ produces the kernel,

$$
\kappa\left(x, x^{\prime}\right)=\phi(x)^{T} \phi\left(x^{\prime}\right)=x^{T} x^{\prime}
$$

It is often easier to directly specify the kernel rather than the basis function...

Example Gaussian kernel models similarity according to an unnormalized Gaussian distribution,

$$
\kappa\left(x, x^{\prime}\right)=\exp \left(-\frac{1}{2 \sigma^{2}}\left(x-x^{\prime}\right)^{2}\right)
$$

Note Despite the name, this is not a Gaussian probability density.

Also called a radial basis function (RBF)

## Kernel Functions

Given any set of data $\left\{x_{i}\right\}_{i=1}^{n}$ a necessary and sufficient condition of a valid kernel function is that the nxn gram matrix,

$$
\mathbf{K}=\left(\begin{array}{cccc}
\kappa\left(x_{1}, x_{1}\right) & \kappa\left(x_{1}, x_{2}\right) & \ldots & \kappa\left(x_{1}, x_{n}\right) \\
\kappa\left(x_{2}, x_{1}\right) & \kappa\left(x_{2}, x_{2}\right) & \ldots & \kappa\left(x_{2}, x_{n}\right) \\
\vdots & \vdots & \vdots & \vdots \\
\kappa\left(x_{n}, x_{1}\right) & \kappa\left(x_{n}, x_{2}\right) & \ldots & \kappa\left(x_{n}, x_{n}\right)
\end{array}\right)
$$

Is a symmetric positive semidefinite matrix.

## Kernel Ridge Regression

## Kernel representation requires inversion of NxN matrix

## Primal

$\boldsymbol{\Phi}=\left(\begin{array}{cccc}1 & \phi_{1}\left(x_{1}\right) & \ldots & \phi_{M}\left(x_{1}\right) \\ 1 & \phi_{1}\left(x_{2}\right) & \ldots & \phi_{M}\left(x_{2}\right) \\ \vdots & \vdots & \vdots & \vdots \\ 1 & \phi_{1}\left(x_{N}\right) & \ldots & \phi_{M}\left(x_{N}\right)\end{array}\right)$
$w=\left(\boldsymbol{\Phi}^{T} \boldsymbol{\Phi}+\lambda I\right)^{-1} \boldsymbol{\Phi}^{T} \mathbf{y}$

MxM Matrix Inversion $\mathrm{O}\left(\mathrm{M}^{3}\right)$

## Dual

$$
\mathbf{K}=\left(\begin{array}{cccc}
\kappa\left(x_{1}, x_{1}\right) & \kappa\left(x_{1}, x_{2}\right) & \ldots & \kappa\left(x_{1}, x_{n}\right) \\
\kappa\left(x_{2}, x_{1}\right) & \kappa\left(x_{2}, x_{2}\right) & \ldots & \kappa\left(x_{2}, x_{n}\right) \\
\vdots & \vdots & \vdots & \vdots \\
\kappa\left(x_{n}, x_{1}\right) & \kappa\left(x_{n}, x_{2}\right) & \ldots & \kappa\left(x_{n}, x_{n}\right)
\end{array}\right)
$$

$$
y(x)=\mathbf{k}(\mathbf{x})^{T} \underbrace{(\mathbf{K}+\lambda I)^{-1}} \mathbf{y}
$$

NxN Matrix Inversion $\mathrm{O}\left(\mathrm{N}^{3}\right)$

Number of training data $N$ greater than basis functions $M$

