# 9.913 Pattern Recognition for Vision 

Class VI - Density Estimation<br>Yuri Ivanov

## Road Map



## Generative vs. Discriminative

There are two schools of thought in Machine Learning:

1. Generative:

- Estimate class models from data
- Compute the discriminant function
- Plug in your data - get the answer

2. Discriminative:

- Estimate the discriminant from data $\succ$ Last class
- Plug in your data - get the answer


## Density Estimation

Density Estimation is at the core of generative Pattern Recognition

$$
P(a<x<b)=\int_{a}^{b} p(x) d x
$$

mean : $E[x]=\int x p(x) d x$
covariance : $E\left[(x-E[x])(x-E[x])^{T}\right]$

$$
=\int\left[(x-E[x])(x-E[x])^{T}\right] p(x) d x
$$

function mean : $E[f(x)]=\int f(x) p(x) d x$
conditional mean $: E[y \mid x]=\int y p(y \mid x) d x$

## Refresher

Minimum expected risk:

$$
R^{*}=\int \min _{\omega}[R(\alpha \mid x)] p(x) d x
$$

... is based on conditional risk:

$$
\omega_{i}=\underset{\omega}{\arg \min } R(\alpha \mid x)
$$

... which is computed from the posterior:

$$
R(\alpha \mid x)=L(\alpha \mid \omega) P(\omega \mid x)
$$

... which depends on the likelihood:

$$
P(\omega \mid x)=\frac{p(x \mid \omega) P(\omega)}{p(x)}
$$

Data:

$$
D=\left\{D_{i}\right\}_{i=1}^{C}
$$

Assume that $D_{j}$ contains noinformation about $\omega_{i}, \forall i \neq j$
NOTATIONALLY - we abandon the class label:

$$
p\left(x \mid \phi x_{j}\right) \longleftrightarrow p(x)
$$

Keep in mind:

$$
p\left(x \mid \omega_{i}\right) \neq p(x)
$$

Goal:
model the probability density function $p(x)$, given a finite number of data points, $x_{1}, x_{2}, \ldots, x_{N}$, drawn from it.

## Three Methods

1. Parametric

- Good: small number of parameters
- Bad: choice of the parametric form

2. Non-parametric

- Good: data "dictates" the approximator
- Bad: large number of parameters

3. Semi-parametric

- Good: combine the best of both worlds
- Bad: harder to design
- Good again: design can be subject to optimization


## Parametric Density Estimation

Estimate the density from a given functional family

Given:

$$
p(x \mid \theta)=f(x, \theta)
$$

Find:
$\theta$

Two methods of parameter estimation:

1. Maximum Likelihood method

- Parameters are viewed as unknown but fixed values

2. Bayesian method

- Parameters are random variables that have their distributions


## Normal (Gaussian) Density Function

A common assumption - Gaussian $\quad \theta=(\mu, \Sigma)$

$$
p(x \mid \theta)=\frac{1}{(2 \pi)^{d / 2}|\Sigma|^{1 / 2}} \exp (-\frac{1}{2} \underbrace{\left((x-\mu)^{T} \Sigma^{-1}(x-\mu)\right.}))
$$

Number of "Volume" dimensions of the covariance

## Squared

Mahalanobis
distance
$\mu=E[x] \quad-d$ parameters
$\Sigma=E\left[(x-\mu)^{T}(x-\mu)\right]$
$-d(d+1) / 2$ parameters

Normal Density

$$
\mu=\left[\begin{array}{l}
0 \\
0 \\
0
\end{array}\right] \quad \Sigma=\left[\begin{array}{ccc}
1 & 0 & .5 \\
0 & 1 & .3 \\
.5 & .3 & 1
\end{array}\right]
$$

Constant density, $(x-\mu)^{T} \Sigma^{-1}(x-\mu)=C \quad$ - quadratic surface


## Whitening Transform

Define:

$$
\begin{array}{ll}
\Lambda=\operatorname{diag}(\operatorname{eigval}(\Sigma)) & - \text { Scaling matrix } \\
\Phi=\operatorname{eigvec}(\Sigma) & - \text { Rotation matrix }
\end{array}
$$

Then:

$$
W=\Lambda^{-1 / 2} \Phi^{T}
$$

-"Unscales" and "unrotates" the data
For all:

$$
x \sim N(\mathbf{0}, \Sigma)
$$

$$
W x \sim N(\mathbf{0}, I)
$$

## You Are Here



## Maximum Likelihood

Parameters are fixed but unknown.

$$
D \equiv\left\{x^{1}, x^{2}, \ldots, x^{N}\right\} \quad \text { a data set, drawn from } p(x)
$$

Notationally, we make density explicitly dependent on parameters:

$$
p(x) \Longleftrightarrow p(x \mid \theta)
$$

Assuming that the data is drawn independently (i.i.d.):

$$
L(\theta) \equiv p(D \mid \theta)=\prod_{n=1}^{N} p\left(x^{n} \mid \theta\right) \quad \text { - a likelihood function }
$$

To find $\theta$ Maximize $L(\theta)$ w.r.t. parameters.

Maximizing $L(\theta)$ is equivalent to maximizing log-likelihood function:

$$
l(\theta) \equiv \log L(\theta)=\log \prod_{n=1}^{N} p\left(x^{n} \mid \theta\right)=\sum_{n=1}^{N} \log p\left(x^{n} \mid \theta\right)
$$

To find $\theta$ set the derivative to 0 :

$$
\nabla_{\theta} l(\theta)=\sum_{n=1}^{N} \nabla_{\theta} \log p\left(x^{n} \mid \theta\right)=0
$$

And solve for $\theta$

## Quick Summary - ML Parameter Estimation



Solving a Maximum Likelihood Problem
Fixed covariance:




## Maximum Likelihood Example

In d-dimensions:
$\nabla_{\theta} l(\theta)=\sum_{n} \nabla_{\theta}\left\{-\frac{d}{2} \log [2 \pi]-\frac{1}{2} \log [|\Sigma|]-\frac{1}{2}\left(x^{n}-\mu\right)^{T} \Sigma^{-1}\left(x^{n}-\mu\right)\right\}$

Solving for the mean:

$$
\begin{aligned}
\nabla_{\mu} l(\theta) & =-\frac{1}{2} \sum_{n} \Sigma^{-1}\left(x^{n}-\hat{\mu}\right)=0 \Rightarrow \\
\hat{\mu} & =\frac{1}{N} \sum_{n=1}^{N} x^{n} \quad \text { - arithmetic average of samples }
\end{aligned}
$$

$$
\nabla_{\theta} l(\theta)=\sum_{n} \nabla_{\theta}\left\{-\frac{d}{2} \log [2 \pi]-\frac{1}{2} \log [|\Sigma|]-\frac{1}{2}\left(x^{n}-\mu\right)^{T} \Sigma^{-1}\left(x^{n}-\mu\right)\right\}
$$

Solving for the covariance:
For symmetric $M: \quad \frac{d|M|}{d M}=|M| M^{-1} \quad$ and $\quad \frac{d\left(a^{T} M^{-1} b\right)}{d M}=M^{-1} a b^{T} M^{-1} \quad \Longrightarrow$

$$
\nabla_{\Sigma} l(\theta)=-\frac{1}{2} \sum_{n}\left\{\hat{\Sigma}^{-1}-\hat{\Sigma}^{-1}\left(x^{n}-\hat{\mu}\right)\left(x^{n}-\hat{\mu}\right)^{T} \hat{\Sigma}^{-1}\right\}=0 \Rightarrow
$$

$$
\text { biased } \quad \hat{\Sigma}=\frac{1}{N} \sum_{n=1}^{N}\left(x^{n}-\hat{\mu}\right)\left(x^{n}-\hat{\mu}\right)^{T} \quad \begin{gathered}
\text {-arithmetic average of } \\
\text { indiv. covariances }
\end{gathered}
$$

## Recursive ML

## What if data comes one sample at a time?

$$
\begin{aligned}
\hat{\mu}_{N}= & \frac{1}{N} \sum_{n=1}^{N} x^{n}=\frac{1}{N}\left[x^{N}+\sum_{n=1}^{N-1} x^{n}\right] \\
& =\frac{1}{N}\left[x^{N}+(N-1) \hat{\mu}_{N-1}\right]=\hat{\mu}_{N-1}+\frac{1}{N}\left[x^{N}-\hat{\mu}_{N-1}\right]
\end{aligned}
$$

This estimate "stiffens" with more data (as it should).

One idea - fix the fraction. Then the estimate can track a non-stationary process

## Recursive ML

Fix the update rate and retrace the steps:

$$
v_{N}=v_{N-1}+\gamma\left[x^{N}-v_{N-1}\right]=(1-\gamma) v_{N-1}+\gamma x^{N}
$$

$$
=(1-\gamma)^{2} v_{N-2}+(1-\gamma) \gamma x^{N-1}+\gamma x^{N}
$$



$$
=(1-\gamma)^{M} v_{N-M}+\sum_{k=1}^{M}(1-\gamma)^{M-k} \gamma x^{k}
$$



## Simple Example

Several images from a static camera:


How much noise is in it?

$$
\begin{aligned}
& x=\operatorname{vec}\left(I_{t}-I_{t-1}\right) \\
& \mu=0 \\
& \sigma=1.2
\end{aligned}
$$

Now we can set a threshold that will statistically distinguish pixel noise from an object


## Problems with ML

## We are given two estimates:



Which one do we believe?
ML gives a single solution, regardless of uncertainty.

## You Are Here



In classification our goal so far has been to estimate $P(\omega \mid x)$

Let's make the dependency on the data explicit:

$$
P\left(\omega_{i} \mid x, D\right)=\frac{p\left(x \mid \omega_{i}, D\right) P\left(\omega_{i} \mid D\right)}{p(x \mid D)}
$$

$P\left(\omega_{i} \mid D\right) \quad$ - this is easy to compute
$P(x \mid D) \quad$ - this is easy to compute by marginalization

What about $p\left(x \mid \omega_{i}, D\right) ?$

## Bayesian Parameter Estimation

$$
P\left(\omega_{i} \mid x, D\right)=\frac{p\left(x \mid \omega_{i}, D\right) P\left(\omega_{i} \mid D\right)}{p(x \mid D)}
$$

This is a supervised problem so far:

$$
\begin{gathered}
D=\left\{D_{1}, D_{2}, \ldots, D_{N}\right\} \\
\begin{aligned}
p\left(x \mid \omega_{i}, D\right) & =p\left(x \mid \omega_{i},\left\{D_{j}\right\}_{j=1 \ldots N}\right) \\
& =p\left(x \mid \omega_{i}, D_{i},\{D\}_{j \neq i}\right)=p\left(x \mid \omega_{i}, D_{i}\right) \\
P\left(\omega_{i} \mid x, D\right) & =\frac{p\left(x \mid \omega_{i}, D_{i}\right) P\left(\omega_{i} \mid D\right)}{p(x \mid D)}
\end{aligned}
\end{gathered}
$$

We will assume that we can obtain "labeled" data, so again:

Notationally:

$$
p\left(x \mid{ }_{p}, D_{i}\right) \Longleftrightarrow p(x \mid D)
$$

Now our problem is to compute density for $x$ given the data $D$.

We assume the form of $p(x)$ - the source density for $D$ :

$$
p(x) \Longleftrightarrow p(x \mid \theta)
$$

$\ldots$ and treat $\theta$ as a random variable

## Bayesian Parameter Estimation

Instead of choosing a value for a parameter, we use them all:

$$
p(x \mid D)=\int p(x, \theta \mid D) d \theta=\int p(x \mid \theta, \nsim p(\theta \mid D) d \theta
$$

Data predicts the new sample $x$ is independent of $D$ given $\theta$

$$
\begin{gathered}
\quad=\int \underbrace{p(x \mid \theta)}_{\neq} \underbrace{p(\theta \mid D)}_{\begin{array}{c}
\searrow \\
\text { We chose the form } \\
\text { of this }
\end{array}} d \theta \\
\text { What is this? }
\end{gathered}
$$

Average densities $p(x \mid \theta)$ for ALL possible values of $\theta$ weighted by its posterior probability

## Bayesian Parameter Estimation

Computing the posterior probability for $\theta$ :

## Quick Summary - Bayesian Parameter Estimation



Bayesian Parameter Estimation
For $\theta=\mu$ :
Parameter prior



## Bayesian Parameter Estimation - Example

First let's deal with the parameter:

Likelihood:

$$
\begin{aligned}
& p(x \mid \mu)=\mathbf{N}\left(\mu, \sigma^{2}\right) \quad \text { fixed } \\
& p(\mu)=\mathbf{N}\left(\mu_{0}, \widehat{\sigma_{0}^{2}}\right)
\end{aligned}
$$

Need to find: $\quad p(\mu \mid D)$

Bayes rule again:

$$
p_{N}(\mu \mid D)=\frac{p(D \mid \mu) p(\mu)}{p(D)}=\alpha \underbrace{\left[\prod_{n=1}^{N} p\left(x^{n} \mid \mu\right)\right] \mathbf{N}\left(\mu_{0}, \sigma_{0}^{2}\right)}_{\text {This is a Gaussian }}=\mathbf{N}\left(\mu_{N}, \sigma_{N}\right)
$$

## Bayesian Parameter Estimation - Example

So, the posterior is a Gaussian

$$
p_{N}(\mu \mid D)=\mathbf{N}\left(\mu_{N}, \sigma_{N}\right)
$$

After some algebra and identifying the terms:

$$
\frac{1}{\sigma_{N}^{2}}=\frac{1}{\sigma^{2}} N+\frac{1}{\sigma_{0}^{2}} \quad-\text { when Gaussians multiply - precisions add }
$$

$\ldots$ and

$$
\mu_{N}=\frac{N \sigma_{0}^{2}}{N \sigma_{0}^{2}+\sigma^{2}} \bar{x}+\frac{\sigma^{2}}{N \sigma_{0}^{2}+\sigma^{2}} \mu_{0}
$$

With increasing $N$ covariance of the posterior decreases and the prior becomes unimportant.

## Bayesian Parameter Estimation - Example

Now the integral:

$$
\begin{aligned}
& p(x \mid D)=\int p(x \mid \theta) p(\theta \mid D) d \theta \\
& =\int \mathbf{N}\left(\mu, \sigma^{2}\right) \mathbf{N}\left(\mu_{N}, \sigma_{N}^{2}\right) d \mu=\mathbf{N}\left(\mu_{N}, \sigma^{2}+\sigma_{N}^{2}\right) \\
& \text { You can show that it } \\
& \text { is also a Gaussian }
\end{aligned}
$$

Any guesses about why Gaussian is such a common assumption?

## Recursive Bayes

For $N$-point likelihood:

$$
\begin{aligned}
& p\left(D^{N} \mid \theta\right)=\prod_{n=1}^{N} p\left(x^{n} \mid \theta\right) \\
& \quad=p\left(x^{N} \mid \theta\right) \prod_{n=1}^{N-1} p\left(x^{n} \mid \theta\right)=p\left(x^{N} \mid \theta\right) p\left(D^{N-1} \mid \theta\right)
\end{aligned}
$$

From this the recursive relation for the posterior:

$$
\begin{aligned}
p\left(\theta \mid D^{N}\right)= & \frac{p\left(x^{N} \mid \theta\right) p\left(D^{N-1} \mid \theta\right) p(\theta)}{p\left(D^{N}\right)} \\
& =\frac{p\left(x^{N} \mid \theta\right) p\left(\theta \mid D^{N-1}\right)}{\int p\left(x^{N} \mid \theta\right) p\left(\theta \mid D^{N-1}\right) d \theta}
\end{aligned}
$$

Again:

$$
p\left(\theta \mid D^{N}\right)=\frac{p\left(x^{N} \mid \theta\right) p\left(\theta \mid D^{N-1}\right)}{\int p\left(x^{N} \mid \theta\right) p\left(\theta \mid D^{N-1}\right) d \theta}
$$

- 1-point update.

Setting $N=1$ :

$$
\begin{aligned}
& \frac{1}{\sigma_{n}^{2}}=\frac{1}{\sigma^{2}}+\frac{1}{\sigma_{n-1}^{2}} \\
& \mu_{n}=\frac{\sigma_{n-1}^{2}}{\sigma_{n-1}^{2}+\sigma^{2}} x+\frac{\sigma^{2}}{\sigma_{n-1}^{2}+\sigma^{2}} \mu_{n-1}
\end{aligned}
$$

Recursive Bayes (cont.)


## Problems with Bayesian Method

1. Integration is difficult
2. Analytic solutions are only available for restricted class of densities
3. Technicality: If the true $p(x \mid \theta)$ is NOT what we assume it is, the prior probability of any parameter setting is 0 !
4. Integration is difficult
5. Did I mention that the integration is hard?

Relation between Bayesian and ML Inference

$$
\begin{aligned}
p(\theta \mid D) \propto p & D \mid \theta) p(\theta) \\
& =\left[\prod_{n} p\left(x^{n} \mid \theta\right)\right] p(\theta)=L(\theta) p(\theta)
\end{aligned}
$$

If the peak is sharp and $p(\theta)$ is flat, then:

$$
\begin{aligned}
& p(x \mid D)=\int p(x \mid \theta) p(\theta \mid D) d \theta \\
& \simeq \int p(x \mid \hat{\theta}) p(\theta \mid D) d \theta=p(x \mid \hat{\theta}) \int p(\theta \mid D) d \theta=p(x \mid \hat{\theta}) \\
& \quad \operatorname{As} N \rightarrow \infty, p(x \mid D) \leftrightarrow p(x \mid \hat{\theta})
\end{aligned}
$$

Non-Parametric Methods for Density Estimation

Non-parametric methods do not assume any particular form for $p(x)$

1. Histograms
2. Kernel Methods
3. K-NN method

## You Are Here



## $\hat{P}(x)$ is a discrete approximation of $p(x)$

- Count a number of times that $x$ lands in the $i$-th bin

$$
H(i)=\sum_{j=1}^{N} I\left(x \in R_{i}\right), \quad \forall i=1,2, \ldots, M
$$

- Normalize

$$
\hat{P}(i)=\frac{H(i)}{\sum_{j=1}^{M} H(j)}
$$

## Histograms

How many bins?


## Histograms

Good:

- Once it is constructed, the data can be discarded
- Quick and intuitive


## Bad:

- Very sensitive to number of bins, $M$
- Estimated density is not smooth
- Poor generalization in higher dimensions


## Aside: Curse of dimensionality (Bellman, ‘61):

- Imagine we build a histogram of a 1-d feature (say, Hue)
- 10 bins
- 1 bin $=10 \%$ of the input space
- need at least 10 points to populate every bin
- We add another feature (say, Saturation)
- 10 bins again
- 1 bin $=1 \%$ of the input space
- we need at least 100 points to populate every bin
-We add another feature (say, Value)
- 10 bins again
- 1 bin $=0.1 \%$ of the input space
- we need at least 1000 points to populate every bin

$$
N=b^{d} \quad \text { - number of points grows exponentially }
$$

## Aside: Curse Continues

Volume of a cube in $R^{d}$ with side $l$ :

$$
V_{l}=l^{d}
$$

Volume of a cube with side $l-\varepsilon$ :

$$
V_{\varepsilon}=(l-\varepsilon)^{d}
$$

Volume of the $\varepsilon$-shell:


$$
\Delta=V_{l}-V_{\varepsilon}=l^{d}-(l-\varepsilon)^{d}
$$

Ratio of the volume of the $\varepsilon$-shell to the volume of the cube:

$$
\frac{\Delta}{V_{l}}=\frac{l^{d}-(l-\varepsilon)^{d}}{l^{d}}=1-\left(1-\frac{\varepsilon}{l}\right)^{d} \rightarrow 1 \text { as } d \rightarrow \infty!!!!!!
$$

## Aside: Lessons of the curse

In generative models:

- Use as much data as you can get your hands on
- Reduce dimensionality as much as you can get away with
<End of Digression>


## General Reasoning

## By definintion:

$$
P(x \in R)=P=\int_{R} p\left(x^{\prime}\right) d x^{\prime}
$$

If we have $N$ i.i.d. points drawn form $p(x)$ :
$B(N, P)$ is a binomial distribution of $k$

## General Reasoning (cont.)

Mean and variance of $B(N, P)$ :
Mean:

$$
\mu=E[k]=N P \quad \Rightarrow P=E[k / N]
$$

Variance: $\quad \sigma^{2}=E\left[(k-\mu)^{2}\right]=N P(1-P)$

$$
\Rightarrow E\left[(k / N-P)^{2}\right]=\left(\frac{\sigma}{N}\right)^{2}=P(1-P) / N
$$

That is:

- $\quad E[k / N]$ is a good estimate of $P$
- $\quad P$ is distributed around this estimate with vanishing variance

So:

$$
P \simeq k / N
$$

## General Reasoning (cont.)

So:

$$
P \simeq k / N
$$

On the other hand, under mild assumptions:

$$
P=\int_{R} p\left(x^{\prime}\right) d x^{\prime} \simeq p(x) V
$$


... which leads to:

$$
p(x) \simeq \frac{k}{N V}
$$

## General Reasoning (cont.)

Now, given $N$ data points - how do we really estimate $p(x)$ ?

$$
p(x) \simeq \frac{k}{N V}
$$

Fix $k$ and vary $V$ until it encloses $k$ points

K-Nearest Neighbors (KNN)

Fix $V$ and count how many points (k) it encloses

Kernel methods

## You Are Here



## Kernel Methods of Density Estimation

We choose $V$ by specifying a hypercube with a side $h$ :

$$
V=h^{d}
$$

Mathematically:

$$
H(\mathbf{y})= \begin{cases}1 & \left|y_{j}\right|<1 / 2 \quad j=1, \ldots, d \\ 0 & \text { otherwise }\end{cases}
$$


kernel function:

$$
H(\mathbf{y}) \geq 0, \quad \forall \mathbf{y} \quad \text { and } \quad \int H(\mathbf{y}) d \mathbf{y}=1
$$

## Parzen Windows

## Then

$$
H\left(\left(\mathbf{x}-\mathbf{x}^{n}\right) / h\right) \quad \text { - a hypercube with side } h \text { centered at } \boldsymbol{x}^{n}
$$

$H$ can help count the points in a volume $V$ around any $x$ :

$$
k(x)=\sum_{n=1}^{N} H\left(\frac{x-x^{n}}{h}\right) \xrightarrow{\begin{array}{l}
\text { No contribution } \\
\text { to the count at } x
\end{array}}
$$

## Rectangular Kernel

So the number of points in h-neighborhood of $x$ :

$$
k(x)=\sum_{n=1}^{N} H\left(\frac{x-x^{n}}{h}\right)
$$

... is easily converted to the density estimate:

$$
\tilde{p}(x)=\frac{k(x)}{N V}=\frac{1}{N} \sum_{n=1}^{N} \frac{1}{h^{d}} H\left(\frac{x-x^{n}}{h}\right) \quad K\left(x, x^{n}\right)
$$

Subtle point:

$$
\begin{gathered}
\int\left[\frac{1}{N} \sum_{n=1}^{N} K\left(x, x^{n}\right)\right] d x=\frac{1}{N} \sum_{n=1}^{N}\left[\int K\left(x, x^{n}\right) d x\right]=1 \\
\Rightarrow \int \tilde{p}(x) d x=1
\end{gathered}
$$

## Example



## Smoothed Window Functions

The problem is as in histograms - it is discontinuous
We can choose a smoother function, s.t.:

$$
\tilde{p}(x) \geq 0, \quad \forall x \quad \text { and } \quad \int \tilde{p}(x) d x=1
$$

Ensured by kernel conditions
Eg: <loud cheer> a (spherical) Gaussian:

$$
\begin{aligned}
& K\left(x, x^{n}\right)=\frac{1}{(\sqrt{2 \pi} h)^{d}} \exp \left(-\frac{\left\|x-x^{n}\right\|}{2 h^{2}}\right) \\
& \tilde{p}(x)=\frac{1}{N} \sum_{n=1}^{N} \frac{1}{(\sqrt{2 \pi} h)^{d}} \exp \left(-\frac{\left\|x-x^{n}\right\|}{2 h^{2}}\right)
\end{aligned}
$$

... so:

## Example






## Some Insight

Interesting to look at expectation of the estimate with respect to all possible datasets:

$$
\begin{aligned}
& E[\tilde{p}(x)]=E\left[\frac{1}{N} \sum_{n=1}^{N} K\left(x, x^{n}\right)\right]=E\left[K\left(x, x^{\prime}\right)\right] \\
& \quad=\int K\left(x-x^{\prime}\right) p\left(x^{\prime}\right) d x^{\prime}-\text { convolution with true density }
\end{aligned}
$$

That is:

$$
\tilde{p}(x)=p(x) \text { if } K\left(x, x^{\prime}\right)=\delta\left(x, x^{\prime}\right)
$$

But not for the finite data set!

## Conditions for Convergence

How small can we make $h$ for a given $N$ ?

$$
\lim _{N \rightarrow \infty} h_{N}^{d}=0
$$

$\lim _{N \rightarrow \infty} N h_{N}^{d}=\infty$

- But slower than $1 / N$

Based on the similar analysis of variance of estimates

Eg:

$$
\begin{aligned}
& h_{N}^{d}=h_{1}^{d} / \sqrt{N} \\
& h_{N}^{d}=h_{1}^{d} / \log (N)
\end{aligned}
$$

Note that the choice of $h_{1}^{d}$ is still up to us.

## Problems With Kernel Estimation

- Need to choose the width parameter, $h$
- Can be chosen empirically
- Can be adaptive, eg. $h_{j}=h d_{j k}$ - where $d_{j k}$ the distance from $x_{j}$ to $k$-th nearest neighbor
- Need to store all data to represent the density
- Leads to Mixture Density Estimation


## You Are Here



## K-Nearest Neighbors

## Recall that:

$$
\tilde{p}(x)=\frac{k}{N V}
$$

Now we fix $k$ (typically $k=\sqrt{N}$ ) and expand $V$ to contain $k$ points

This is not a true density!
Eg.: choose $N=1, k=1$. Then:

$$
\tilde{p}(x)=\frac{1}{1 \cdot\left\|x-x_{1}\right\|} \quad \text { Oops! }
$$

BUT it is useful for a number of theoretical and practical reasons.

## K-NN Classification Rule

Let's try classification with K-NN density estimate

Data: $\quad N \quad$ - total points
$N_{j} \quad$ - points in class $\omega_{j}$

Need to find the class label for a query, $x$

Expand a sphere from $x$ to include $K$ points

$$
\begin{array}{ll}
K & \text { - number of neighbors of } x \\
K_{j} & \text { - points of class } \omega_{j} \text { among } K
\end{array}
$$



## KNN Classification

Then class priors are given by: $\quad p\left(\omega_{j}\right)=\frac{N_{j}}{N}$
We can estimate conditional and marginal densities around any $x$ :

$$
p\left(x \mid \omega_{j}\right)=\frac{K_{j}}{N_{j} V} \quad p(x)=\frac{K}{N V}
$$

By Bayes rule: $\quad p\left(\omega_{j} \mid x\right)=\frac{K_{j}}{N_{j} V} \frac{N_{j}}{N} \frac{N V}{K}=\frac{K_{j}}{K}$
Then for minimum error rate classification:

$$
C=\underset{j}{\arg \max } K_{j}
$$

## KNN Classification

## Important theoretical result:

In the extreme case, $K=1$, it can be shown that:

$$
\begin{aligned}
\text { for } \quad P=\lim _{N \rightarrow \infty} P_{N}^{N} \text {-sample error rate } \\
P^{*} \leq P \leq P^{*}\left(2-\frac{c}{c-1} P^{*}\right)
\end{aligned}
$$

That is, using just a single neighbor rule, the error rate is at most twice the Bayes error!!!

## Problems with Non-parametric Methods

- Memory: need to store all data points
- Computation: need to compute distances to all data points every time
- Parameter choice: need to choose the smoothing parameter


## You Are Here



## Mixture Density Model

Mixture model - a linear combination of parametric densities

$$
\begin{aligned}
& \text { Number of components } \\
& p(x)=\sum_{j=1}^{M} p(x \mid j) P(j) \\
& P(j) \geq 0, \quad \forall j \quad \text { and } \quad \sum_{j=1}^{\text {Component }} \quad P(j)=1
\end{aligned}
$$

Uses MUCH less "kernels" than kernel methods
Kernels are parametric densities, subject to estimation

## Example



Using ML principle, the objective function is the log-likelihood:

$$
l(\theta) \equiv \log \prod_{n=1}^{N} p\left(x^{n}\right)=\sum_{n=1}^{N} \log \left\{\sum_{j=1}^{M} p\left(x^{n} \mid j\right) P(j)\right\}
$$

Differentiate w.r.t. parameters:

$$
\begin{aligned}
& \nabla_{\theta_{j}} l(\theta)=\sum_{n=1}^{N} \frac{\partial}{\partial \theta_{j}} \log \left\{\sum_{k=1}^{M} p\left(x^{n} \mid k\right) P(k)\right\} \\
&=\sum_{n=1}^{N} \frac{1}{\sum_{k=1}^{M} p\left(x^{n} \mid k\right) P(k)} \frac{\partial}{\partial \theta_{j}} p\left(x^{n} \mid j\right) P(j)
\end{aligned}
$$

Again let's assume that $p(x \mid \omega)$ is a Gaussian
We need to estimate $M$ priors, and $M$ sets of means and covariances

$$
\frac{\partial l(\theta)}{\partial \mu_{j}}=\sum_{n=1}^{N} P\left(j \mid x^{n}\right)\left[\Sigma_{j}^{-1}\left(x^{n}-\hat{\mu}_{j}\right)\right]
$$

Setting it to 0 and solving for $\mu_{j}$ :

$$
\hat{\mu}_{j}=\frac{\sum_{n=1}^{N} P\left(j \mid x^{n}\right) x^{n}}{\sum_{n=1}^{N} P\left(j \mid x^{n}\right)} \quad \text { - convex sum of all data }
$$

## Mixture Density

Similarly for the covariances:

$$
\frac{\partial l(\theta)}{\partial \sigma_{j}^{2}}=\sum_{n=1}^{N} P\left(j \mid x^{n}\right)\left[\hat{\mathbf{S}}_{j}^{-1}-\hat{\mathbf{S}}_{j}^{-1}\left(x^{n}-\hat{\mu}_{j}\right)\left(x^{n}-\hat{\mu}_{j}\right)^{T} \hat{\mathbf{S}}_{j}^{-1}\right]
$$

Setting it to 0 and solving for $\Sigma_{i}$ :

$$
\hat{\mathbf{S}}_{j}=\frac{\sum_{n=1}^{N} P\left(j \mid x^{n}\right)\left(x^{n}-\hat{\mu}_{j}\right)\left(x^{n}-\hat{\mu}_{j}\right)^{T}}{\sum_{n=1}^{N} P\left(j \mid x^{n}\right)}
$$

## Mixture Density

A little harder for $P(j)$ - optimization is subject to constraints:

$$
\sum_{j=1}^{M} P(j)=1 \quad \text { and } \quad P(j) \geq 0, \forall j
$$

Here is a trick to enforce the constraints:

$$
\begin{gathered}
P(j)=\frac{\exp \left(\gamma_{j}\right)}{\sum_{k=1}^{M} \exp \left(\gamma_{k}\right)} \\
\frac{\partial P(i)}{\partial \gamma_{j}}=\delta(i-j) P(j)-P(i) P(j)
\end{gathered}
$$

Using the chain rule:

$$
\begin{aligned}
& \nabla_{\gamma_{j}} l(\theta)=\sum_{k=1}^{M} \frac{\partial l(\theta)}{\partial P(k)} \frac{\partial P(k)}{\partial \gamma_{j}}=\sum_{k=1}^{M} \sum_{n=1}^{N} \frac{p\left(x^{n} \mid k\right)}{P(x)}\left(\delta_{j k} P(j)-P(j) P(k)\right) \\
& \quad=\sum_{n=1}^{N}\left\{\frac{p\left(x^{n} \mid j\right)}{P(x)} P(j)-\sum_{k=1}^{M} \frac{p\left(x^{n} \mid k\right)}{P(x)} P(j) P(k)\right\} \\
& =\sum_{n=1}^{N}\left\{P\left(j \mid x^{n}\right)-P(j) \sum_{k=1}^{M} p\left(k \mid x^{n}\right)\right\}=\sum_{n=1}^{N}\left\{P\left(j \mid x^{n}\right)-P(j)\right\}=0
\end{aligned}
$$

The last expression gives the value at the extremum:

$$
P(j)=\frac{1}{N} \sum_{n=1}^{N} P\left(j \mid x^{n}\right)
$$

## Mixture Density

What's the problem?

$$
\sum^{N} P\left(j \mid x^{n}\right) x^{n}
$$

$$
\begin{aligned}
& P(j)=\frac{1}{N} \sum_{n=1}^{N} P\left(j \mid x^{n}\right) \quad \hat{\mu}_{j}=\frac{\sum_{n=1}^{N}}{\sum_{n=1}^{N} P\left(j \mid x^{n}\right)} \\
& \hat{\mathbf{S}}_{j}=\frac{\sum_{n=1}^{N} P\left(j \mid x^{n}\right)\left(x^{n}-\hat{\mu}_{j}\right)\left(x^{n}-\hat{\mu}_{j}\right)^{T}}{\sum_{n=1}^{N} P\left(j \mid x^{n}\right)}
\end{aligned}
$$

We can't compute these directly!

Solution - EM algorithm. We will study it in Clustering.

## You Are Here



