# Clustering 10/26-702 Spring 2016 

## 1 The Clustering Problem

In a clustering problem we aim to find groups in the data. Unlike classification, the data are not labeled, and so clustering is considered an example of unsupervised learning. In many cases, clustering methods are presented as heuristic algorithms without a clearly stated goal. We will cast clustering problems as estimating some population quantities.

Example 1 Figures 1 and 2 show some synthetic examples where the clusters are meant to be intuitively clear. In Figure 1 there are two blob-like clusters. Identifying clusters like this is easy. Figure 2 shows four clusters: a blob, two rings and a half ring. Identifying clusters with unusual shapes like this is not quite as easy. In fact, finding clusters of this type requires nonparametric methods.

## 2 Mode-Based Clustering

Let $p$ be the density of $X \in \mathbb{R}^{d}$. Assume that $p$ has modes $m_{1}, \ldots, m_{k_{0}}$ and that $p$ is a Morse function, which means that the Hessian of $p$ at each stationary point is non-degenerate. We can use the modes to define clusters as follows.

Given any point $x \in \mathbb{R}^{d}$, there is a unique gradient ascent path, or integral curve, passing through $x$ that eventually leads to one of the modes. We define the clusters to be the "basins of attraction" of the modes, the equivalence classes of points whose ascent paths lead to the same mode. Formally, an integral curve through $x$ is a path $\pi_{x}: \mathbb{R} \rightarrow \mathbb{R}^{d}$ such that $\pi_{x}(0)=x$ and

$$
\begin{equation*}
\pi_{x}^{\prime}(t)=\nabla p\left(\pi_{x}(t)\right) \tag{1}
\end{equation*}
$$

Integral curves never intersect (except at stationary points) and they partition the space.
Equation (1) means that the path $\pi$ follows the direction of steepest ascent of $p$ through $x$. The destination of the integral curve $\pi$ through a (non-mode) point $x$ is defined by

$$
\begin{equation*}
\operatorname{dest}(x)=\lim _{t \rightarrow \infty} \pi_{x}(t) \tag{2}
\end{equation*}
$$

It can then be shown that for all $x, \operatorname{dest}(x)=m_{j}$ for some mode $m_{j}$. That is: all integral curves lead to modes. For each mode $m_{j}$, define the sets

$$
\begin{equation*}
\mathcal{A}_{j}=\left\{x: \operatorname{dest}(x)=m_{j}\right\} . \tag{3}
\end{equation*}
$$



Figure 1: A synthetic example with two "blob-like" clusters.


Figure 2: A synthetic example with four clusters with a variety of different shapes.

These sets are known as the ascending manifolds, and also known as the cluster associated with $m_{j}$, or the basin of attraction of $m_{j}$. The $\mathcal{A}_{j}$ 's partition the space. See Figure 3. The collection of ascending manifolds is called the Morse complex.

Given data $X_{1}, \ldots, X_{n}$ we construct an estimate $\widehat{p}$ of the density. Let $\widehat{m}_{1}, \ldots, \widehat{m}_{k}$ be the estimated modes and let $\widehat{\mathcal{A}}_{1}, \ldots, \widehat{\mathcal{A}}_{k}$ be the corresponding ascending manifolds derived from $\widehat{p}$. The sample clusters $C_{1}, \ldots, C_{k}$ are defined to be $C_{j}=\left\{X_{i}: X_{i} \in \widehat{\mathcal{A}}_{j}\right\}$.

Recall that the kernel density estimator is

$$
\begin{equation*}
\widehat{p}(x) \equiv \widehat{p}_{h}(x)=\frac{1}{n} \sum_{i=1}^{n} \frac{1}{h^{d}} K\left(\frac{\left\|x-X_{i}\right\|}{h}\right) \tag{4}
\end{equation*}
$$

where $K$ is a smooth, symmetric kernel and $h>0$ is the bandwidth. ${ }^{1}$ The mean of the estimator is

$$
\begin{equation*}
p_{h}(x)=\mathbb{E}\left[\widehat{p}_{h}(x)\right]=\int K(t) p(x+t h) d t . \tag{5}
\end{equation*}
$$

To locate the modes of $\widehat{p}_{h}$ we use the mean shift algorithm which finds modes by approximating the steepest ascent paths. The algorithm is given in Figure 4. The result of this process is the set of estimated modes $\widehat{\mathcal{M}}=\left\{\widehat{m}_{1}, \ldots, \widehat{m}_{k}\right\}$. We also get the clustering for free: the mean shift algorithm shows us what mode each point is attracted to. See Figure 5.

A modified version of the algorithm is the blurred mean-shift algorithm (Carreira-Perpinan, 2006). Here, we use the data as the mesh and we replace the data with the mean-shifted data at each step. This converges very quickly but must be stopped before everything converges to a single point; see Figures 6 and 7.

[^0]

Figure 3: The left plot shows a function with four modes. The right plot shows the ascending manifolds (basins of attraction) corresponding to the four modes.

What we are doing is tracing out the gradient flow. The flow lines lead to the modes and they define the clusters. In general, a flow is a map $\phi: \mathbb{R}^{d} \times \mathbb{R} \rightarrow \mathbb{R}^{d}$ such that $\phi(x, 0)=x$ and $\phi(\phi(x, t), s)=\phi(x, s+t)$. The latter is called the semi-group property.

### 2.1 Choosing the Bandwidth

As usual, choosing a good bandwidth is crucial. You might wonder if increasing the bandwidth, decreases the number of modes. Silverman (1981) showed that the answer is yes if you use a Normal kernel.

Theorem 2 (Silverman 1981) Let $\widehat{p}_{h}$ be a kernel density estimator using a Gaussian kernel in one dimension. Then the number of modes of $\widehat{p}_{h}$ is a non-increasing function of $h$. The Gaussian kernel is the unique kernel with this property.

We still need a way to pick $h$. We can use cross-validation as before. One could argue that we should choose $h$ so that we estimate the gradient $g(x)=\nabla p(x)$ well since the clustering is based on the gradient flow.

How can we estimate the loss of the gradient? Consider, first the scalar case. Note that

$$
\int\left(\widehat{p}^{\prime}-p^{\prime}\right)^{2}=\int\left(\widehat{p}^{\prime}\right)^{2}-2 \int \widehat{p}^{\prime} p^{\prime}+\int\left(p^{\prime}\right)^{2}
$$

## Mean Shift Algorithm

1. Input: $\widehat{p}(x)$ and a mesh of points $A=\left\{a_{1}, \ldots, a_{N}\right\}$ (often taken to be the data points).
2. For each mesh point $a_{j}$, set $a_{j}^{(0)}=a_{j}$ and iterate the following equation until convergence:

$$
a_{j}^{(s+1)} \longleftarrow \frac{\sum_{i=1}^{n} X_{i} K\left(\frac{\left\|a_{j}^{(s)}-X_{i}\right\|}{h}\right)}{\sum_{i=1}^{n} K\left(\frac{\left\|a_{j}^{(s)}-X_{i}\right\|}{h}\right)} .
$$

3. Let $\widehat{\mathcal{M}}$ be the unique values of the set $\left\{a_{1}^{(\infty)}, \ldots, a_{N}^{(\infty)}\right\}$.
4. Output: $\widehat{\mathcal{M}}$.

Figure 4: The Mean Shift Algorithm.


Figure 5: A simple example of the mean shift algorithm.


Figure 6: The crescent data example. Top left: data. Top right: a few steps of mean-shift. Bottom left: a few steps of blurred mean-shift.


Figure 7: The Broken Ring example. Top left: data. Top right: a few steps of mean-shift. Bottom left: a few steps of blurred mean-shift.

We can ignore the last term. The first term is known. To estimate the middle term, we use integration by parts to get

$$
\int \widehat{p}^{\prime} p^{\prime}=-\int p^{\prime \prime} p
$$

suggesting the cross-validation estimator

$$
\int\left(\widehat{p}^{\prime}(x)\right)^{2} d x+\frac{2}{n} \sum_{i} \widehat{p}_{i}^{\prime \prime}\left(X_{i}\right)
$$

where $\widehat{p}_{i}^{\prime \prime}$ is the leave-one-out second derivative. More generally, by repeated integration by parts, we can estimate the loss for the $r^{\text {th }}$ derivative by

$$
\mathrm{CV}_{r}(h)=\int\left(\widehat{p}^{(r)}(x)\right)^{2} d x-\frac{2}{n}(-1)^{r} \sum_{i} \widehat{p}_{i}^{(2 r)}\left(X_{i}\right)
$$

Let's now discuss estimating derivatives more generally following Chacon and Duong (2013). Let

$$
\widehat{p}_{H}(x)=\frac{1}{n} \sum_{i=1}^{n} K_{H}\left(x-X_{i}\right)
$$

where $K_{H}(x)=|H|^{-1 / 2} K\left(H^{-1 / 2} x\right)$. Let $D=\partial / \partial x=\left(\partial / \partial x_{1}, \ldots, \partial / \partial x_{d}\right)$ be the gradient operator. Let $H(x)$ be the Hessian of $p(x)$ whose entries are $\partial^{2} p /\left(\partial x_{j} \partial x_{k}\right)$. Let

$$
D^{\otimes r} p=(D p)^{\otimes r}=\partial^{r} p / \partial x^{\otimes r} \in \mathbb{R}^{d^{r}}
$$

denote the $r^{\text {th }}$ derivatives, organized into a vector. Thus

$$
D^{\otimes 0} p=p, \quad D^{\otimes 1} p=D p, \quad D^{\otimes 2} p=\operatorname{vec}(H)
$$

where vec takes a matrix and stacks the columns into a vector.
The estimate of $D^{\otimes r} p$ is
$\widehat{p}^{(r)}(x)=D^{\otimes r} \widehat{p}_{H}(x)=\frac{1}{n} \sum_{i=1}^{n} D^{\otimes r} K_{H}\left(x-X_{i}\right)=\frac{1}{n} \sum_{i=1}^{n}|H|^{-1 / 2}\left(H^{-1 / 2}\right)^{\otimes r} D^{\otimes r} K\left(H^{-1 / 2}\left(x-X_{i}\right)\right.$.
The integrated squared error is

$$
L=\int\left\|D^{\otimes r} \widehat{p}_{H}(x)-D^{\otimes r} p(x)\right\|^{2} d x
$$

Chacon, Duong and Wand shows that $\mathbb{E}[L]$ is minimized by choosing $H$ so that each entry has order $n^{-2 /(d+2 r+4)}$ leading to a risk of order $O\left(n^{-4 /(d+2 r+4)}\right)$. In fact, it may be shown that

$$
\begin{aligned}
\mathbb{E}[L] & =\frac{1}{n}|H|^{-1 / 2} \operatorname{tr}\left(\left(H^{-1}\right)^{\otimes r} R\left(D^{\otimes r} K\right)\right)-\frac{1}{n} \operatorname{tr} R^{*}\left(K_{H} \star K_{H}, D^{\otimes r} p\right) \\
& +\operatorname{tr} R^{*}\left(K_{H} \star K_{H}, D^{\otimes r} p\right)-2 \operatorname{tr} R^{*}\left(K_{H}, D^{\otimes r} p\right)+\operatorname{tr} R\left(D^{\otimes r} p\right)
\end{aligned}
$$

where

$$
\begin{aligned}
R(g) & =\int g(x) g^{T}(x) d x \\
R^{*}(a, g) & =\int(a \star g)(x) g^{T}(x) d x
\end{aligned}
$$

and $(a \star g)$ is componentwise convolution.
To estimate the loss, we expand $L$ as

$$
L=\int\left\|D^{\otimes r} \widehat{p}_{H}(x)\right\|^{2} d x-2 \int\left\langle D^{\otimes r} \widehat{p}_{H}(x), D^{\otimes r} p(x)\right\rangle d x+\text { constant } .
$$

Using some high-voltage calculations, Chacon and Duong (2013) derived the following leave-one-out approximation to the first two terms:

$$
\mathrm{CV}_{r}(H)=(-1)^{r}|H|^{-1 / 2}\left(\operatorname{vec}\left(H^{-1}\right)^{\otimes r}\right)^{T} B(H)
$$

where

$$
B(H)=\frac{1}{n^{2}} \sum_{i, j} D^{\otimes 2 r} \bar{K}\left(H^{-1 / 2}\left(X_{i}-X_{j}\right)\right)-\frac{2}{n(n-1)} \sum_{i \neq j} D^{\otimes 2 r} K\left(H^{-1 / 2}\left(X_{i}-X_{j}\right)\right)
$$

and $\bar{K}=K \star K$ In practice, the minimization is easy if we restrict to matrices of the form $H=h^{2} I$.

A better idea is to used fixed (non-decreasing $h$ ). We don't need $h$ to go to 0 to find the clusters. More on this when we discuss persistence.

### 2.2 Theoretical Analysis

How well can we estimate the modes?

Theorem 3 Assume that $p$ is Morse with finitely many modes $m_{1}, \ldots, m_{k}$. Then for $h>0$ and not too large, $p_{h}$ is Morse with modes $m_{h 1}, \ldots, m_{h k}$ and (possibly after relabelling),

$$
\max _{j}\left\|m_{j}-m_{j h}\right\|=O\left(h^{2}\right) .
$$

With probability tending to $1, \widehat{p}_{h}$ has the same number of modes which we denote by $\widehat{m}_{h 1}, \ldots, \widehat{m}_{h k}$. Furthermore,

$$
\max _{j}\left\|\widehat{m}_{j h}-m_{j h}\right\|=O_{P}\left(\sqrt{\frac{1}{n h^{d+2}}}\right)
$$

and

$$
\max _{j}\left\|\widehat{m}_{j h}-m_{j}\right\|=O\left(h^{2}\right)+O_{P}\left(\sqrt{\frac{1}{n h^{d+2}}}\right)
$$

Remark: Setting $h \asymp n^{-1 /(d+6)}$ gives the rate $n^{-2 /(d+6)}$ which is minimax (Tsyabkov 1990) under smoothness assumptions. See also Romano (1988). However, if we take the fixed $h$ point if view, then we have a $n^{-1 / 2}$ rate.

Proof Outline. But a small ball $B_{j}$ around each $m_{j h}$. We will skip the first step, which is to show that there is one (and only one) local mode in $B_{j}$. Let's focus on showing

$$
\max _{j}\left\|\widehat{m}_{j h}-m_{j h}\right\|=O_{P}\left(\sqrt{\frac{1}{n h^{d+2}}}\right) .
$$

For simplicity, write $m=m_{j h}$ and $x=\widehat{m}_{j h}$. Let $g(x)$ and $H(x)$ be the gradient and Hessian of $p_{h}(x)$ and let $\widehat{g}(x)$ and $\widehat{H}(x)$ be the gradient Hessian of $\widehat{p}_{h}(x)$. Then

$$
(0, \ldots, 0)^{T}=\widehat{g}(x)=\widehat{g}(m)+(x-m)^{T} \int_{0}^{1} \widehat{H}(m+u(x-m)) d u
$$

and so

$$
(x-m)^{T} \int_{0}^{1} \widehat{H}(m+u(x-m)) d u=(g(m)-\widehat{g}(m))
$$

where we used the fact that $\mathbf{0}=g(m)$. Multiplying on the right by $x-m$ we have

$$
(x-m)^{T} \int_{0}^{1} \widehat{H}(m+u(x-m))(x-m) d u=(\widehat{g}(m)-\widehat{g}(m))^{T}(x-m) .
$$

Let $\lambda=\inf _{0 \leq u \leq 1} \lambda_{\text {min }}(H(m+u(x-m)))$. Then $\lambda=\lambda_{\text {min }}(H(m))+o_{P}(1)$ and

$$
(x-m)^{T} \int_{0}^{1} \widehat{H}(x+u(m-x))(x-m) d u \geq \lambda\|x-m\|^{2} .
$$

Hence, using Cauchy-Schwartz,
$\lambda\|x-m\|^{2} \leq\|\widehat{g}(m)-g(m)\|\|x-m\| \leq\|x-m\| \sup _{y}\|\widehat{g}(y)-\widehat{g}(y)\| \leq\|x-m\| O_{P}\left(\sqrt{\frac{1}{n h^{d+2}}}\right)$
and so $\|x-m\|=O_{P}\left(\sqrt{\frac{1}{n h^{d+2}}}\right)$.
Remark: If we treat $h$ as fixed (not decreasing) then the rate is $O_{P}(\sqrt{1 / n})$ independent of dimension.

## 3 Level Set Clustering

An alternative to mode clustering is level set clustering. Let $L_{t}=\left\{x: p_{h}(x)>t\right\}$ denote an upper level set of $p$. Suppose that $L_{t}$ can be decomposed into finitely many disjoint sets: $L_{t}=C_{1} \bigcup \cdots \bigcup C_{k_{t}}$. We call $\mathcal{C}_{t}=\left\{C_{1}, \ldots, C_{k_{t}}\right\}$ the level set clusters at level $t$.

Let $\mathcal{C}=\bigcup_{t \geq 0} \mathcal{C}_{t}$. The clusters in $\mathcal{C}$ form a tree: if $A, B \in \mathcal{C}$, the either (i) $A \subset B$ or (ii) $B \subset A$ or (iii) $A \cap B=\emptyset$. We call $\mathcal{C}$ the level set cluster tree.

The level sets can be estimated in the obvious way: $\widehat{L}_{t}=\left\{x: \widehat{p}_{h}(x)>t\right\}$. How do we decompose $\widehat{L}_{t}$ into its connected components? This can be done as follows. For each $t$ let

$$
\mathcal{X}_{t}=\left\{X_{i}: \widehat{p}_{h}\left(X_{i}\right)>t\right\} .
$$

Now construct a graph $G_{t}$ where each $X_{i} \in \mathcal{X}_{t}$ is a vertex and there is an edge between $X_{i}$ and $X_{j}$ if and only if $\left\|X_{i}-X_{j}\right\| \leq \epsilon$ where $\epsilon>0$ is a tuning parameter. Bobrowski et al (2104) show that we can take $\epsilon=h . G_{t}$ is a called a Rips graphs. The clusters at level $t$ are estimated by taking the connected components of the graph $G_{t}$. In summary:

1. Compute $\widehat{p}_{h}$.
2. For each $t$, let $\mathcal{X}_{t}=\left\{X_{i}: \widehat{p}_{h}\left(X_{i}\right)>t\right\}$.
3. Form a graph $G_{t}$ for the points in $\mathcal{X}_{t}$ by connecting $X_{i}$ and $X_{j}$ if $\left\|X_{i}-X_{j}\right\| \leq h$.
4. The clusters at level $t$ are the connected components of $G_{t}$.

A Python package, called DeBaCl , written by Brian Kent, can be found at
http://www.brianpkent.com/projects.html.
Fabrizio Lecci has written an R implementation, include in his R package: TDA (topological data analysis). You can get it at:
http://cran.r-project.org/web/packages/TDA/index.html
Two examples are shown in Figures 8 and 9 .

### 3.1 Theory

How well does this work? Define the Hausdorff distance between two sets by

$$
H(U, V)=\inf \{\epsilon: U \subset V \oplus \epsilon \text { and } V \subset U \oplus \epsilon\}
$$

where

$$
V \oplus \epsilon=\bigcup_{x \in V} B(x, \epsilon)
$$

and $B(x, \epsilon)$ denotes a ball of radius $\epsilon$ centered at $x$. We would like to say that $L_{t}$ and $\widehat{L}_{t}$ are close. In general this is not true. Sometimes $L_{t}$ and $L_{t+\delta}$ are drastically different even for small $\delta$. (Think of the case where a mode has height $t$.) But we can estimate stable level sets. Let us say that $L_{t}$ is stable if there exists $a>0$ and $C>0$ such that, for all $\delta<a$,

$$
H\left(L_{t-\delta}, L_{t+\delta}\right) \leq C \delta
$$



Figure 8: DeBaClR in two dimensions.


Figure 9: DeBaClR in three dimensions.

Theorem 4 Suppose that $L_{t}$ is stable. Then $H\left(\widehat{L}_{t}, L_{t}\right)=O_{P}\left(\sqrt{\log n /\left(n h^{d}\right)}\right)$.

Proof. Let $\left.r_{n}=\sqrt{\log n /\left(n h^{d}\right)}\right)$. We need to show two things: (i) for every $x \in L_{t}$ there exists $y \in \widehat{L}_{t}$ such that $\|x-y\|=O_{P}\left(r_{n}\right)$ and (ii) for every $x \in \widehat{L}_{t}$ there exists $y \in L_{t}$ such that $\|x-y\|=O_{P}\left(r_{n}\right)$. First, we note that, by earlier results, $\left\|\widehat{p}_{h}-p_{h}\right\|_{\infty}=O_{P}\left(r_{n}\right)$. To show (i), suppose that $x \in L_{t}$. By the stability assumption, there exists $y \in L_{t+r_{n}}$ such that $\|x-y\| \leq C r_{n}$. Then $p_{h}(y)>t+r_{n}$ which implies that $\widehat{p}_{h}(y)>t$ and so $y \in \widehat{L}_{t}$. To show (ii), let $x \in \widehat{L}_{t}$ so that $\widehat{p}_{h}(x)>t$. Thus $p_{h}(x)>t-r_{n}$. By stability, there is a $y \in L_{t}$ such that $\|x-y\| \leq C r_{n}$.

### 3.2 Persistence

Consider a smooth density $p$ with $M=\sup _{x} p(x)<\infty$. The $t$-level set clusters are the connected components of the set $L_{t}=\{x: p(x) \geq t\}$. Suppose we find the upper level sets $L_{t}=\{x: p(x) \geq t\}$ as we vary $t$ from $M$ to 0 . Persistent homology measures how the topology of $L_{t}$ varies as we decrease $t$. In our case, we are only interested in the modes, which correspond to the zeroth order homology. (Higher order homology refers to holes, tunnels etc.) The idea of using persistence to study clustering was introduced by Chazal, Guibas, Oudot and Skraba (2013).

Imagine setting $t=M$ and then gradually decreasing $t$. Whenever we hit a mode, a new level set cluster is born. As we decrease $t$ further, some clusters may merge and we say that one of the clusters (the one born most recently) has died. See Figure 10.

In summary, each mode $m_{j}$ has a death time and a birth time denoted by $\left(d_{j}, b_{j}\right)$. (Note that the birth time is larger than the death time because we start at high density and move to lower density.) The modes can be summarized with a persistence diagram where we plot the points $\left(d_{1}, b_{1}\right), \ldots,\left(d_{k}, b_{k}\right)$ in the plane. See Figure 10. Points near the diagonal correspond to modes with short lifetimes. We might kill modes with lifetimes smaller than the bootstrap quantile $\epsilon_{\alpha}$ defined by

$$
\begin{equation*}
\epsilon_{\alpha}=\inf \left\{z: \frac{1}{B} \sum_{b=1}^{B} I\left(\left\|\widehat{p}_{h}^{* b}-\widehat{p}_{h}\right\|_{\infty}>z\right) \leq \alpha\right\} . \tag{6}
\end{equation*}
$$

Here, $\widehat{p}_{h}^{* b}$ is the density estimator based on the $b^{\text {th }}$ bootstrap sample. This corresponds to killing a mode if it is in a $2 \epsilon_{\alpha}$ band around the diagonal. See Fasy, Lecci, Rinaldo, Wasserman, Balakrishnan and Singh (2014). Note that the starting and ending points of the vertical bars on the level set tree are precisely the coordinates of the persistence diagram. (A more precise bootstrap approach was introduced in Chazal, Fasy, Lecci, Michel, Rinaldo and Wasserman (2104).)


Figure 10: Starting at the top of the density and moving down, each mode has a birth time $b$ and $a$ death time $d$. The persistence diagram (right) plots the points $\left(d_{1}, b_{1}\right), \ldots,\left(d_{4}, b_{4}\right)$. Modes with a long lifetime are far from the diagonal.

## 4 k-means (Vector Quantization)

One of the oldest approaches to clustering is to find $k$ representative points, called prototypes or cluster centers, and then divide the data into groups based on which prototype they are closest to. The set of prototypes is called a codebook. For now, we asume that $k$ is given. Later we discuss how to choose $k$.

Let $X_{1}, \ldots, X_{n} \sim P$ where $X_{i} \in \mathbb{R}^{d}$. Let $C=\left\{c_{1}, \ldots, c_{k}\right\}$ where each $c_{j} \in \mathbb{R}^{d}$. We call $C$ a codebook. Let $\Pi_{C}[X]$ be the projection of $X$ onto $C$ :

$$
\begin{equation*}
\Pi_{C}[X]=\operatorname{argmin}_{c \in C}\|c-X\|^{2} . \tag{7}
\end{equation*}
$$

Define the empirical clustering risk of a codebook $C$ by

$$
\begin{equation*}
R_{n}(C)=\frac{1}{n} \sum_{i=1}^{n}\left\|X_{i}-\Pi_{C}\left[X_{i}\right]\right\|^{2}=\frac{1}{n} \sum_{i=1}^{n} \min _{1 \leq j \leq k}\left\|X_{i}-c_{j}\right\|^{2} . \tag{8}
\end{equation*}
$$

Let $\mathcal{C}_{k}$ denote all codebooks of length $k$. The optimal codebook $\widehat{C}=\left\{\widehat{c}_{1}, \ldots, \widehat{c}_{k}\right\} \in \mathcal{C}_{k}$ minimizes $R_{n}(C)$ :

$$
\begin{equation*}
\widehat{C}=\operatorname{argmin}_{C \in \mathcal{C}_{k}} R_{n}(C) \tag{9}
\end{equation*}
$$

The empirical risk is an estimate of the population clustering risk defined by

$$
\begin{equation*}
R(C)=\mathbb{E}\left\|X-\Pi_{C}[X]\right\|^{2}=\mathbb{E} \min _{1 \leq j \leq k}\left\|X-c_{j}\right\|^{2} \tag{10}
\end{equation*}
$$



Figure 11: The Voronoi tesselation formed by 10 cluster centers $c_{1}, \ldots, c_{10}$. The cluster centers are indicated by dots. The corresponding Voronoi cells $T_{1}, \ldots, T_{10}$ are defined as follows: a point $x$ is in $T_{j}$ if $x$ is closer to $c_{j}$ than $c_{i}$ for $i \neq j$.
where $X \sim P$. The optimal population quantization $C^{*}=\left\{c_{1}^{*}, \ldots, c_{k}^{*}\right\} \in \mathcal{C}_{k}$ minimizes $R(C)$. We can think of $\widehat{C}$ as an estimate of $C^{*}$. This method is called $k$-means clustering or vector quantization.

A codebook $C=\left\{c_{1}, \ldots, c_{k}\right\}$ defines a set of cells known as a tesselation. Let

$$
\begin{equation*}
T_{j}=\left\{x:\left\|x-c_{j}\right\| \leq\left\|x-c_{s}\right\|, \text { for all } s \neq j\right\} . \tag{11}
\end{equation*}
$$

The set $T_{j}$ is known as a Voronoi cell and consists of all points closer to $c_{j}$ than any other point in the codebook. See Figure 11.
The usual algorithm to minimize $R_{n}(C)$ and find $\widehat{C}$ is the $k$-means clustering algorithmalso known as Lloyd's algorithm - see Figure 12. The risk $R_{n}(C)$ has multiple minima. The algorithm will only find a local minimum and the solution depends on the starting values. A common way to choose the starting values is to select $k$ data points at random. We will discuss better methods for choosing starting values in Section 4.1.

Example 5 Figure 13 shows synthetic data inspired by the Mickey Mouse example from

1. Choose $k$ centers $c_{1}, \ldots, c_{k}$ as starting values.
2. Form the clusters $C_{1}, \ldots, C_{k}$ as follows. Let $g=\left(g_{1}, \ldots, g_{n}\right)$ where $g_{i}=\operatorname{argmin}_{j}\left\|X_{i}-c_{j}\right\|$. Then $C_{j}=\left\{X_{i}: g_{i}=j\right\}$.
3. For $j=1, \ldots, k$, let $n_{j}$ denote the number of points in $C_{j}$ and set

$$
c_{j} \longleftarrow \frac{1}{n_{j}} \sum_{i: X_{i} \in C_{j}} X_{i} .
$$

4. Repeat steps 2 and 3 until convergence.
5. Output: centers $\widehat{C}=\left\{c_{1}, \ldots, c_{k}\right\}$ and clusters $C_{1}, \ldots, C_{k}$.

Figure 12: The $k$-means (Lloyd's) clustering algorithm.
http://en.wikipedia.org/wiki/K-means_clustering.

The data in the top left plot form three clearly defined clusters. $k$-means easily finds in the clusters (top right). The bottom shows the same example except that we now make the groups very unbalanced. The lack of balance causes $k$-means to produce a poor clustering.

Example 6 We applied $k$-means clustering to the Topex data with $k=9$. The data are discretized so we treated each curve as one vector of length 70. The resulting nine clusters are shown in Figure 14.

Example 7 (Supernova Clustering) Figure 15 shows supernova data where we apply $k$ means clustering with $k=4$. The type Ia supernovae get split into two groups although the groups are very similar. The other type also gets split into two groups which look qualitatively different.
$k$-means works best when clusters are at least roughly spherical. Otherwise, it can produce very non-intuitive clusters.

Example 8 The top left plot of Figure 16 shows a dataset with two ring-shaped clusters. The remaining plots show the clusters obtained using $k$-means clustering with $k=2,3,4$. Clearly, $k$-means does not capture the right structure in this case.


Figure 13: Synthetic data inspired by the "Mickey Mouse" example from wikipedia. Top left: three balanced clusters. Top right: result from running $k$ means with $k=3$. Bottom left: three unbalanced clusters. Bottom right: result from running $k$ means with $k=3$ on the unbalanced clusters. $k$-means does not work well here because the clusters are very unbalanced.

### 4.1 Starting Values for $k$-means

Since $\widehat{R}_{n}(C)$ has multiple minima, Lloyd's algorithm is not guaranteed to minimize $R_{n}(C)$. The clustering one obtains will depend on the starting values. The simplest way to choose starting values is to use $k$ randomly chosen points. But this often leads to poor clustering.

Example 9 Figure 17 shows data from a distribution with nine clusters. The raw data are in the top left plot. The top right plot shows the results of running the $k$-means algorithm with $k=9$ using random points as starting values. The clustering is quite poor. This is because we have not found the global minimum of the empirical risk function. The two bottom plots show better methods for selecting starting values that we will describe below.

Hierarchical Starting Values. Tseng and Wong (2005) suggest the following method for choosing staring values for $k$-means. Run single-linkage hierarchical clustering (which we describe in Section 6) to obtains $p \times k$ clusters. They suggest using $p=3$ as a default. Now take the centers of the $k$-largest of the $p \times k$ clusters and use these as starting values. See the bottom left plot in Figure 17.
$k$-means ${ }^{++}$. Arthur and Vassilvitskii (2007) invented an algorithm called k -means ${ }^{++}$to get


Figure 14: The nine clusters found in the Topex data using $k$-means clustering with $k=9$. Each plot show the curves in that cluster together with the mean of the curves in that cluster.


Figure 15: Clustering of the supernova light curves with $k=4$.


Figure 16: Top left: a dataset with two ring-shaped clusters. Top right: $k$-means with $k=2$. Bottom left: $k$-means with $k=3$. Bottom right: $k$-means with $k=4$.


Figure 17: An example with 9 clusters. Top left: data. Top right: $k$-means with random starting values. Bottom left: $k$-means using starting values from hierarchical clustering. Bottom right: the $k$-means ${ }^{++}$algorithm.
good starting values. They show that if the starting points are chosen in a certain way, then we can get close to the minimum with high probability. In fact the starting points themselves - which we call seed points - are already close to minimizing $R_{n}(C)$. The algorithm is described in Figure 18. See the bottom right plot in Figure 17 for an example.

Theorem 10 (Arthur and Vassilvitskii, 2007). Let $C=\left\{c_{1}, \ldots, c_{k}\right\}$ be the seed points from the $k$-means ${ }^{++}$algorithm. Then,

$$
\begin{equation*}
\mathbb{E}\left(R_{n}(C)\right) \leq 8(\log k+2)\left(\min _{C} R_{n}(C)\right) \tag{12}
\end{equation*}
$$

where the expectation is over the randomness of the algorithm.

See Arthur and Vassilvitskii (2007) for a proof. They also show that the Euclidean distance can be replaced with the $\ell_{p}$ norm in the algorithm. The result is the same except that the constant 8 gets replaced by $2^{p+2}$. It is possible to improve the $k$-means ${ }^{++}$algorithm. Ailon, Jaiswal and Monteleoni (2009) showed that, by choosing $3 \log k$ points instead of one point, at each step of the algorithm, the $\log k$ term in (12) can be replaced by a constant. They call the algorithm, k-means\#.

1. Input: Data $X=\left\{X_{1}, \ldots, X_{n}\right\}$ and an integer $k$.
2. Choose $c_{1}$ randomly from $X=\left\{X_{1}, \ldots, X_{n}\right\}$. Let $C=\left\{c_{1}\right\}$.
3. For $j=2, \ldots, k$ :
(a) Compute $D\left(X_{i}\right)=\min _{c \in C}\left\|X_{i}-c\right\|$ for each $X_{i}$.
(b) Choose a point $X_{i}$ from $X$ with probability

$$
p_{i}=\frac{D^{2}\left(X_{i}\right)}{\sum_{j=1}^{n} D^{2}\left(X_{j}\right)} .
$$

(c) Call this randomly chosen point $c_{j}$. Update $C \longleftarrow C \cup\left\{c_{j}\right\}$.
4. Run Lloyd's algorithm using the seed points $C=\left\{c_{1}, \ldots, c_{k}\right\}$ as starting points and output the result.

Figure 18: The $k$-means ${ }^{++}$algorithm.

### 4.2 Choosing $k$

In $k$-means clustering we must choose a value for $k$. This is still an active area of research and there are no definitive answers. The problem is much different than choosing a tuning parameter in regression or classification because there is no observable label to predict. Indeed, for $k$-means clustering, both the true risk $R$ and estimated risk $R_{n}$ decrease to 0 as $k$ increases. This is in contrast to classification where the true risk gets large for high complexity classifiers even though the empirical risk decreases. Hence, minimizing risk does not make sense. There are so many proposals for choosing tuning parameters in clustering that we cannot possibly consider all of them here. Instead, we highlight a few methods.

Elbow Methods. One approach is to look for sharp drops in estimated risk. Let $R_{k}$ denote the minimal risk among all possible clusterings and let $\widehat{R}_{k}$ be the empirical risk. It is easy to see that $R_{k}$ is a nonincreasing function of $k$ so minimizing $R_{k}$ does not make sense. Instead, we can look for the first $k$ such that the improvement $R_{k}-R_{k+1}$ is small, sometimes called an elbow. This can be done informally by looking at a plot of $\widehat{R}_{k}$. We can try to make this more formal by fixing a small number $\alpha>0$ and defining

$$
\begin{equation*}
k_{\alpha}=\min \left\{k: \frac{R_{k}-R_{k+1}}{\sigma^{2}} \leq \alpha\right\} \tag{13}
\end{equation*}
$$

where $\sigma^{2}=\mathbb{E}\left(\|X-\mu\|^{2}\right)$ and $\mu=\mathbb{E}(X)$. An estimate of $k_{\alpha}$ is

$$
\begin{equation*}
\widehat{k}_{\alpha}=\min \left\{k: \frac{\widehat{R}_{k}-\widehat{R}_{k+1}}{\widehat{\sigma}^{2}} \leq \alpha\right\} \tag{14}
\end{equation*}
$$

where $\widehat{\sigma}^{2}=n^{-1} \sum_{i=1}^{n}\left\|X_{i}-\bar{X}\right\|^{2}$.
Unfortunately, the elbow method often does not work well in practice because there may not be a well-defined elbow.

Hypothesis Testing. A more formal way to choose $k$ is by way of hypothesis testing. For each $k$ we test
$H_{k}$ : the number of clusters is $k$ versus $H_{k+1}:$ the number of clusters is $>k$.
We begin $k=1$. If the test rejects, then we repeat the test for $k=2$. We continue until the first $k$ that is not rejected. In summary, $\widehat{k}$ is the first $k$ for which $k$ is not rejected.

Currently, my favorite approach is the one in Liu, Hayes, Andrew Nobel and Marron (2012). (JASA, 2102, 1281-1293). They simply test if the data are multivariate Normal. If this rejects, they split into two clusters and repeat. The have an R package sigclust for this. A similar procedure, called PG means is described in Feng and Hammerly (2007).

Example 11 Figure 19 shows a two-dimensional example. The top left plot shows a single cluster. The $p$-values are shown as a function of $k$ in the top right plot. The first $k$ for which the $p$-value is larger than $\alpha=.05$ is $k=1$. The bottom left plot shows a dataset with three clusters. The $p$-values are shown as a function of $k$ in the bottom right plot. The first $k$ for which the p-value is larger than $\alpha=.05$ is $k=3$.

Stability. Another class of methods are based on the idea of stability. The idea is to find the largest number of clusters than can be estimated with low variability.

We start with a high level description of the idea and then we will discuss the details. Suppose that $Y=\left(Y_{1}, \ldots, Y_{n}\right)$ and $Z=\left(Z_{1}, \ldots, Z_{n}\right)$ are two independent samples from $P$. Let $A_{k}$ be any clustering algorithm that takes the data as input and outputs $k$ clusters. Define the stability

$$
\begin{equation*}
\Omega(k)=\mathbb{E}\left[s\left(A_{k}(Y), A_{k}(Z)\right)\right] \tag{15}
\end{equation*}
$$

where $s(\cdot, \cdot)$ is some measure of the similarity of two clusterings. To estimate $\Omega$ we use random subsampling. Suppose that the original data are $X=\left(X_{1}, \ldots, X_{2 n}\right)$. Randomly


Figure 19: Top left: a single cluster. Top right: p -values for various $k$. The first $k$ for which the p-value is larger than .05 is $k=1$. Bottom left: three clusters. Bottom right: p-values for various $k$. The first $k$ for which the p -value is larger than .05 is $k=3$.
split the data into two equal sets $Y$ and $Z$ of size $n$. This process if repeated $N$ times. Denote the random split obtained in the $j^{\text {th }}$ trial by $Y^{j}, Z^{j}$. Define

$$
\widehat{\Omega}(k)=\frac{1}{N} \sum_{j=1}^{N}\left[s\left(A_{k}\left(Y^{j}\right), A_{k}\left(Z^{j}\right)\right)\right] .
$$

For large $N, \widehat{\Omega}(k)$ will approximate $\Omega(k)$. There are two ways to choose $k$. We can choose a small $k$ with high stability. Alternatively, we can choose $k$ to maximize $\widehat{\Omega}(k)$ if we somehow standardize $\widehat{\Omega}(k)$.

Now we discuss the details. First, we need to define the similarity between two clusterings. We face two problems. The first is that the cluster labels are arbitrary: the clustering $(1,1,1,2,2,2)$ is the same as the clustering $(4,4,4,8,8,8)$. Second, the clusterings $A_{k}(Y)$ and $A_{k}(Z)$ refer to different data sets.

The first problem is easily solved. We can insist the labels take values in $\{1, \ldots, k\}$ and then we can maximize the similarity over all permutations of the labels. Another way to solve the problem is the following. Any clustering method can be regarded as a function $\psi$ that takes two points $x$ and $y$ and outputs a 0 or a 1 . The interpretation is that $\psi(x, y)=1$ if $x$ and $y$ are in the same cluster while $\psi(x, y)=0$ if $x$ and $y$ are in a different cluster. Using this representation of the clustering renders the particular choice of labels moot. This is the approach we will take.

Let $\psi_{Y}$ and $\psi_{Z}$ be clusterings derived from $Y$ and $Z$. Let us think of $Y$ as training data and $Z$ as test data. Now $\psi_{Y}$ returns a clustering for $Y$ and $\psi_{Z}$ returns a clustering for $Z$. We'd like to somehow apply $\psi_{Y}$ to $Z$. Then we would have two clusterings for $Z$ which we could then compare. There is no unique way to do this. A simple and fairly general approach is to define

$$
\begin{equation*}
\psi_{Y, Z}\left(Z_{j}, Z_{k}\right)=\psi_{Y}\left(Y_{j}^{\prime}, Y_{k}^{\prime}\right) \tag{16}
\end{equation*}
$$

where $Y_{j}^{\prime}$ is the closest point in $Y$ to $Z_{j}$ and $Y_{k}^{\prime}$ is the closest point in $Y$ to $Z_{k}$. (More generally, we can use $Y$ and the cluster assignment to $Y$ as input to a classifier; see Lange et al 2004). The notation $\psi_{Y, Z}$ indicates that $\psi$ is trained on $Y$ but returns a clustering for $Z$. Define

$$
s\left(\psi_{Y, Z}, \psi_{Z}\right)=\frac{1}{\binom{n}{2}} \sum_{s \neq t} I\left(\psi_{Y, Z}\left(Z_{s}, Z_{t}\right)=\psi_{Z}\left(Z_{s}, Z_{t}\right)\right)
$$

Thus $s$ is the fraction of pairs of points in $Z$ on which the two clusterings $\psi_{Y, Z}$ and $\psi_{Z}$ agree. Finally, we define

$$
\widehat{\Omega}(k)=\frac{1}{N} \sum_{j=1}^{N} s\left(\psi_{Y^{j}, Z^{j}}, \psi_{Z^{j}}\right)
$$

Now we need to decide how to use $\widehat{\Omega}(k)$ to choose $k$. The interpretation of $\widehat{\Omega}(k)$ requires some care. First, note that $0 \leq \widehat{\Omega}(k) \leq 1$ and $\widehat{\Omega}(1)=\widehat{\Omega}(n)=1$. So simply maximizing $\widehat{\Omega}(k)$
does not make sense. One possibility is to look for a small $k$ larger than $k>1$ with a high stability. Alternatively, we could try to normalize $\widehat{\Omega}(k)$. Lange et al (2004) suggest dividing by the value of $\widehat{\Omega}(k)$ obtained when cluster labels are assigned randomly. The theoretical justification for this choice is not clear. Tibshirani, Walther, Botstein and Brown (2001) suggest that we should compute the stability separately over each cluster and then take the minimum. However, this can sometimes lead to very low stability for all $k>1$.

Many authors have considered schemes of this form, including Breckenridge (1989), Lange, Roth, Braun and Buhmann (2004), Ben-Hur, Elisseeff and Guyron (2002), Dudoit and Fridlyand (2002), Levine and Domany (2001), Buhmann (2010), Tibshirani, Walther, Botstein and Brown (2001) and Rinaldo and Wasserman (2009).

It is important to interpret stability correctly. These methods choose the largest number of stable clusters. That does not mean they choose "the true $k$." Indeed, Ben-David, von Luxburg and Pál (2006), Ben-David and von Luxburg Tübingen (2008) and Rakhlin (2007) have shown that trying to use stability to choose "the true $k$ " - even if that is well-defined - will not work. To explain this point further, we consider some examples from Ben-David, von Luxburg and Pál (2006). Figure 20 shows the four examples. The first example (top left plot) shows a case where we fit $k=2$ clusters. Here, stability analysis will correctly show that $k$ is too small. The top right plot has $k=3$. Stability analysis will correctly show that $k$ is too large. The bottom two plots show potential failures of stability analysis. Both cases are stable but $k=2$ is too small in the bottom left plot and $k=3$ is too big in the bottom right plot. Stability is subtle. There is much potential for this approach but more work needs to be done.

### 4.3 Theoretical Properties

A theoretical property of the $k$-means method is given in the following result. Recall that $C^{*}=\left\{c_{1}^{*}, \ldots, c_{k}^{*}\right\}$ minimizes $R(C)=\mathbb{E}\left\|X-\Pi_{C}[X]\right\|^{2}$.

Theorem 12 Suppose that $\mathbb{P}\left(\left\|X_{i}\right\|^{2} \leq B\right)=1$ for some $B<\infty$. Then

$$
\begin{equation*}
\mathbb{E}(R(\widehat{C}))-R\left(C^{*}\right) \leq c \sqrt{\frac{k(d+1) \log n}{n}} \tag{17}
\end{equation*}
$$

for some $c>0$.

Warning! The fact that $R(\widehat{C})$ is close to $R\left(C_{*}\right)$ does not imply that $\widehat{C}$ is close to $C_{*}$.
This proof is due to Linder, Lugosi and Zeger (1994). The proof uses techniques from a later lecture on VC theory so you may want to return to the proof later.


Figure 20: Examples from Ben-David, von Luxburg and Pál (2006). The first example (top left plot) shows a case where we fit $k=2$ clusters. Stability analysis will correctly show that $k$ is too small. The top right plot has $k=3$. Stability analysis will correctly show that $k$ is too large. The bottom two plots show potential failures of stability analysis. Both cases are stable but $k=2$ is too small in the bottom left plot and $k=3$ is too big in the bottom right plot.

Proof. Note that $R(\widehat{C})-R\left(C^{*}\right)=R(\widehat{C})-R_{n}(\widehat{C})+R_{n}(\widehat{C})-R\left(C^{*}\right) \leq R(\widehat{C})-R_{n}(\widehat{C})+$ $R_{n}\left(C^{*}\right)-R\left(C^{*}\right) \leq 2 \sup _{C \in \mathcal{C}_{k}}\left|R(\widehat{C})-R_{n}(\widehat{C})\right|$. For each $C$ define a function $f_{C}$ by $f_{C}(x)=$ $\left\|x-\Pi_{C}[x]\right\|^{2}$. Note that $\sup _{x}\left|f_{C}(x)\right| \leq 4 B$ for all $C$. Now, using the fact that $\mathbb{E}(Y)=$ $\int_{0}^{\infty} \mathbb{P}(Y \geq t) d t$ whenever $Y \geq 0$, we have

$$
\begin{aligned}
2 \sup _{C \in \mathcal{C}_{k}}\left|R(\widehat{C})-R_{n}(\widehat{C})\right| & =2 \sup _{C}\left|\frac{1}{n} \sum_{i=1}^{n} f_{C}\left(X_{i}\right)-\mathbb{E}\left(f_{C}(X)\right)\right| \\
& =2 \sup _{C}\left|\int_{0}^{\infty}\left(\frac{1}{n} \sum_{i=1}^{n} I\left(f_{C}\left(X_{i}\right)>u\right)-\mathbb{P}\left(f_{C}(Z)>u\right)\right) d u\right| \\
& \leq 8 B \sup _{C, u}\left|\frac{1}{n} \sum_{i=1}^{n} I\left(f_{C}\left(X_{i}\right)>u\right)-\mathbb{P}\left(f_{C}(Z)>u\right)\right| \\
& =8 B \sup _{A}\left|\frac{1}{n} \sum_{i=1}^{n} I\left(X_{i} \in A\right)-\mathbb{P}(A)\right|
\end{aligned}
$$

where $A$ varies over all sets $\mathcal{A}$ of the form $\left\{f_{C}(x)>u\right\}$. The shattering number of $\mathcal{A}$ is $s(\mathcal{A}, n) \leq n^{k(d+1)}$. This follows since each set $\left\{f_{C}(x)>u\right\}$ is a union of the complements of
$k$ spheres. By the VC Theorem,

$$
\begin{aligned}
\mathbb{P}\left(R(\widehat{C})-R\left(C^{*}\right)>\epsilon\right) & \leq \mathbb{P}\left(8 B \sup _{A}\left|\frac{1}{n} \sum_{i=1}^{n} I\left(X_{i} \in A\right)-\mathbb{P}(A)\right|>\epsilon\right) \\
& =\mathbb{P}\left(\sup _{A}\left|\frac{1}{n} \sum_{i=1}^{n} I\left(X_{i} \in A\right)-\mathbb{P}(A)\right|>\frac{\epsilon}{8 B}\right) \\
& \leq 4(2 n)^{k(d+1)} e^{-n \epsilon^{2} /\left(512 B^{2}\right)} .
\end{aligned}
$$

Now conclude that $\mathbb{E}\left(R(\widehat{C})-R\left(C^{*}\right)\right) \leq C \sqrt{k(d+1)} \sqrt{\frac{\log n}{n}}$.
A sharper result, together with a lower bound is the following.

Theorem 13 (Bartlett, Linder and Lugosi 1997) Suppose that $\mathbb{P}\left(\|X\|^{2} \leq 1\right)=1$ and that $n \geq k^{4 / d}$, $\sqrt{d k^{1-2 / d} \log n} \geq 15, k d \geq 8, n \geq 8 d$ and $n / \log n \geq d k^{1+2 / d}$. Then,

$$
\mathbb{E}(R(\widehat{C}))-R\left(C^{*}\right) \leq 32 \sqrt{\frac{d k^{1-2 / d} \log n}{n}}=O\left(\sqrt{\frac{d k \log n}{n}}\right)
$$

Also, if $k \geq 3, n \geq 16 k /\left(2 \Phi^{2}(-2)\right)$ then, for any method $\widehat{C}$ that selects $k$ centers, there exists $P$ such that

$$
\mathbb{E}(R(\widehat{C}))-R\left(C^{*}\right) \geq c_{0} \sqrt{\frac{k^{1-4 / d}}{n}}
$$

where $c_{0}=\Phi^{4}(-2) 2^{-12} / \sqrt{6}$ and $\Phi$ is the standard Gaussian distribution function.

See Bartlett, Linder and Lugosi (1997) for a proof. It follows that $k$-means is risk consistent in the sense that $R(\widehat{C})-R\left(C^{*}\right) \xrightarrow{P} 0$, as long as $k=o\left(n /\left(d^{3} \log n\right)\right)$. Moreover, the lower bound implies that we cannot find any other method that improves much over the $k$-means approach, at least with respect to this loss function.

The previous results depend on the dimension $d$. It is possible to get a dimension-free result at the expense of replacing $\sqrt{k}$ with $k$. In fact, the following result even applies to functions instead of vectors. In that case, we interpret $\|\cdot\|$ to be the norm in a Hilbert space.

Theorem 14 (Biau, Devroye and Lugosi 2008). Suppose that $\mathbb{P}\left(\left\|X_{i}\right\| \leq B\right)=1$. Then

$$
\mathbb{E}(R(\widehat{C}))-R\left(C^{*}\right) \leq \frac{12 B^{2} k}{\sqrt{n}}
$$

Proof. Define $W(C, P)=\mathbb{E}_{P}\left(\min _{1 \leq j \leq k}\left[-2\left\langle X, c_{j}\right\rangle+\left\|c_{j}\right\|^{2}\right]\right)$. Minimizing $R(C)$ is equivalent to minimizing $W(C, P)$ and minimizing $R_{n}(C)$ is equivalent to minimizing $W\left(C, P_{n}\right)$
where $P_{n}$ is the empirical measure that puts mass $1 / n$ at each $X_{i}$. Arguing as in the proof of Theorem 12,

$$
\mathbb{E}(W(\widehat{C}, P))-W\left(C^{*}, P\right) \leq 2 \mathbb{E}\left(\sup _{C} W(C, P)-W\left(C, P_{n}\right)\right)
$$

Let $\sigma_{1}, \ldots, \sigma_{n}$ be Rademacher random variables. That is, $\sigma_{1}, \ldots, \sigma_{n}$ are iid and $\mathbb{P}\left(\sigma_{i}=\right.$ $+1)=\mathbb{P}\left(\sigma_{i}=-1\right)=1 / 2$. Let $X_{1}^{\prime}, \ldots, X_{n}^{\prime}$ be a second independent sample. Let $\ell_{c}(x)=$ $-2\langle x, c\rangle+\|c\|^{2}$. Then,

$$
\begin{aligned}
\mathbb{E}\left(\sup _{C} W(C, P)-W\left(C, P_{n}\right)\right) \leq & \mathbb{E}\left(\sup _{C} \frac{1}{n} \sum_{i=1}^{n} \sigma_{i}\left[\min _{1 \leq j \leq n} \ell_{c_{j}}\left(X_{i}\right)-\min _{1 \leq j \leq n} \ell_{c_{j}}\left(X_{i}^{\prime}\right)\right]\right) \\
\leq & \mathbb{E}\left(\sup _{C} \frac{1}{n} \sum_{i=1}^{n} \sigma_{i}\left[\min _{1 \leq j \leq n} \ell_{c_{j}}\left(X_{i}\right)\right]\right) \\
& +\mathbb{E}\left(\sup _{C} \frac{1}{n} \sum_{i=1}^{n}\left(-\sigma_{i}\right)\left[\min _{1 \leq j \leq n} \ell_{c_{j}}\left(X_{i}\right)\right]\right) \\
= & 2 \mathbb{E}\left(\sup _{C} \frac{1}{n} \sum_{i=1}^{n} \sigma_{i}\left[\min _{1 \leq j \leq n} \ell_{c_{j}}\left(X_{i}\right)\right]\right) .
\end{aligned}
$$

An inductive argument shows that

$$
2 \mathbb{E}\left(\sup _{C} \frac{1}{n} \sum_{i=1}^{n} \sigma_{i}\left[\min _{1 \leq j \leq n} \ell_{c_{j}}\left(X_{i}\right)\right]\right) \leq 4 k\left[\mathbb{E} \sup _{c \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \sigma_{i}\left\langle X_{i}, c\right\rangle+\frac{B^{2}}{2 \sqrt{n}}\right]
$$

Also,

$$
\begin{aligned}
\mathbb{E}\left(\sup _{c \in \mathbb{R}^{d}} \frac{1}{n} \sum_{i=1}^{n} \sigma_{i}\left\langle X_{i}, c\right\rangle\right) & =\mathbb{E}\left(\sup _{c \in \mathbb{R}^{d}} \frac{1}{n}\left\langle\sum_{i=1}^{n} \sigma_{i} X_{i}, c\right\rangle\right)=\frac{B}{n} \mathbb{E}\left\|\sum_{i=1}^{n} \sigma_{i} X_{i}\right\| \\
& \leq \frac{B}{n} \sqrt{\mathbb{E}\left\|\sum_{i=1}^{n} \sigma_{i} X_{i}\right\|^{2}}=B \sqrt{\frac{\mathbb{E} \| X| |^{2}}{n}} \leq \frac{B^{2}}{\sqrt{n}} .
\end{aligned}
$$

The k-means algorithm can be generalized in many ways. For example, if we replace the $L_{2}$ norm with the $L_{1}$ norm we get $k$-medians clustering. We will not discuss these extensions here.

## 5 Geometric Graph Clustering

In geometric graph clustering we form a graph $G_{\epsilon}$ with one node for each observation $X_{i}$. We put an edge between $X_{i}$ and $X_{j}$ if and only if $d\left(X_{i}, X_{j}\right) \leq \epsilon$ where $\epsilon \geq 0$ is a tuning
parameter. Here, $d$ is nay measure of similarity. We will use the Euclidean metric $d\left(X_{i}, X_{j}\right)=$ $\left\|X_{i}-X_{j}\right\|$. The vertices and edges define the graph $G_{\epsilon}$. Define the clusters to be the connected components of $G_{\epsilon}$. We shall see that this type of clustering is related to hierarchical clustering (Section 6) and spectral clustering (Section 7).

To find the connected components of the graph, we mention two algorithms: depth-first search and the spectral method. The latter is discussed in Section 7. The depth-first algorithm is as follows:

1. Pick a starting node $u$.
2. Run traverse to find the connected component containing $u$.
3. Remove this component and repeat (until there are no more nodes).
traverse: In this routine, nodes are marked white (unvisited), gray (in progress) and black (finished).
4. Mark all nodes to white.
5. Mark $u$ as gray.
6. If there is an edge $(u, v)$, where $v$ is white: go to $v$ and mark $v$ gray.
7. Continue until you reach a node $w$ such that there is no edge $(w, x)$ where $x$ is white. Mark $w$ as black and backtrack to parent.
8. Continue until all nodes are black.

Example 15 Figure 21 shows a simple example with two clusters. The top row corresponds to $\epsilon=1$ and the bottom row corresponds to $\epsilon=3$. The left column shows the graph. The right column shows a ball of radius $\epsilon / 2$ around each $X_{i}$. In this way, the connected components of the graph show up as connected sets. The top plot of Figure 22 shows the number of connected components versus $\epsilon$. There is a fairly large range of values of $\epsilon$ that yields two connected components. The bottom plot shows the values of $\epsilon$ at which $k(\epsilon)$ changes.

Example 16 The top left plot of Figure 23 shows data with two ring-shaped clusters. The remaining plots show geometric graph clustering with $\epsilon=.2,1$ and 6 . Clearly, the clustering is quite sensitive to the choice of $\epsilon$.

Example 17 Figure 24 shows data with two ring-shaped clusters. In this case, we have also added background noise. The noise consist of points drawn from a uniform distribution. Here we see that geometric graph clustering fails.

As the last example shows, background noise makes geometric graph clustering fail. It is tempting to fix this by trying to remove the background noise first. But this is essentially what density based clustering does; see Section ??.


Figure 21: Geometric graph clustering. Top row: $\epsilon=1$. Bottom row: $\epsilon=3$. Left column: the graphs. Right column: a ball of radius $\epsilon / 2$ around each $X_{i}$.



Figure 22: Geometric graph clustering. The top plot shows the number of connected components $k(\epsilon)$ versus $\epsilon$. There is a fairly large range of values of $\epsilon$ that yields two connected components. The bottom plot shows the values of $\epsilon$ at which $k(\epsilon)$ changes.


Figure 23: Geometric graph clustering applied to two ring-shaped clusters. Top left: data. Top right: $\epsilon=.2$. Bottom left: $\epsilon=1.0$. Bottom right: $\epsilon=6$.


Figure 24: Geometric graph clustering applied to two ring-shaped clusters with background noise. Top left: data. Top right: $\epsilon=0.40$. Bottom left: $\epsilon=0.54$. Bottom right: $\epsilon=1.46$.

## 6 Hierarchical Clustering

Hierarchical clustering methods build a set of nested clusters at different resolutions. The are two types of hierarchical clustering: agglomerative (bottom-up) and divisive (top-down). With agglomerative clustering we start with some distance or dissimilarity $d(x, y)$ between points. We then extend this distance so that we can compute the distance $d(A, B)$ between to sets of points $A$ and $B$.

The three most common ways of extending the distance are:

| Single Linkage | $d(A, B)=\min _{x \in A, y \in B} d(x, y)$ |
| :--- | :--- |
| Average Linkage | $d(A, B)=\frac{1}{N_{A} N_{B}} \sum_{x \in A, y \in B} d(x, y)$ |
| Complete Linkage | $d(A, B)=\max _{x \in A, y \in B} d(x, y)$ |

The algorithm is:

1. Input: data $X=\left\{X_{1}, \ldots, X_{n}\right\}$ and metric $d$ giving distance between clusters.
2. Let $T_{n}=\left\{C_{1}, C_{2}, \ldots, C_{n}\right\}$ where $C_{i}=\left\{X_{i}\right\}$.
3. For $j=n-1$ to 1 :
(a) Find $j, k$ to minimize $d\left(C_{j}, C_{k}\right)$ over all $C_{j}, C_{k} \in T_{j+1}$.
(b) Let $T_{j}$ be the same as $T_{j+1}$ except that $C_{j}$ and $C_{k}$ are replaced with $C_{j} \cup C_{k}$.
4. Return the sets of clusters $T_{1}, \ldots, T_{n}$.

The result can be represented as a tree, called a dendogram. We can then cut the tree at different places to yield any number of clusters ranging from 1 to $n$. Single linkage often produces thin clusters while complete linkage is better at rounder clusters. Average linkage is in between.

Example 18 Figure 25 shows agglomerative clustering applied to data generated from two rings plus noise. The noise is large enough so that the smaller ring looks like a blob. The data are show in the top left plot. The top right plot shows hierarchical clustering using single linkage. (The tree is cut to obtain two clusters.) The bottom left plot shows average linkage and the bottom right plot shows complete linkage. Single linkage works well while average and complete linkage do poorly.

Let us now mention some theoretical properties of hierarchical clustering. Suppose that $X_{1}, \ldots, X_{n}$ is a sample from a distribution $P$ on $\mathbb{R}^{d}$ with density $p$. A high density cluster is a maximal connected component of a set of the form $\{x: p(x) \geq \lambda\}$. One might expect that single linkage clusters would correspond to high density clusters. This turns out not quite to be the case. See Hartigan (1981) for details. DasGupta (2010) has a modified version of hierarchical clustering that attempts to fix this problem. His method is very similar to density clustering.

Single linkage hierarchical clustering is the same as geometric graph clustering. Let $G=$ $(V, E)$ be a graph where $V=\left\{X_{1}, \ldots, X_{n}\right\}$ and $E_{i j}=1$ if $\left\|X_{i}-X_{j}\right\| \leq \epsilon$ and $E_{i j}=0$ if $\left\|X_{i}-X_{j}\right\|>\epsilon$. Let $C_{1}, \ldots, C_{k}$ denote the connected components of the graph. As we vary $\epsilon$ we get exactly the hierarchical clustering tree.

Finally, we let us mention divisive clustering. This is a form of hierarchical clustering where we start with one large cluster and then break the cluster recursively into smaller and smaller pieces.


Figure 25: Hierarchical clustering applied to two noisy rings. Top left: the data. Top right: two clusters from hierarchical clustering using single linkage. Bottom left: average linkage. Bottom right: complete linkage.

## $7 \quad$ Spectral Clustering

Spectral clustering refers to a class of clustering methods that use ideas related to eigenvector. An excellent tutorial on spectral clustering is von Luxburg (2006) and some of this section relies heavily on that paper. More detail can be found in Chung (1997).

Let $G$ be an undirected graph with $n$ vertices. Typically these vertices correspond to observations $X_{1}, \ldots, X_{n}$. Let $W$ be an $n \times n$ symmetric weight matrix. Say that $X_{i}$ and $X_{j}$ are connected if $W_{i j}>0$. The simplest type of weight matrix has entries that are either 0 or 1 . For example, we could define

$$
W_{i j}=I\left(\left\|X_{i}-X_{j}\right\| \leq \epsilon\right)
$$

as we did in Section 5. An example of a more general weight matrix is $W_{i j}=e^{-\left\|X_{i}-X_{j}\right\|^{2} /\left(2 h^{2}\right)}$. The degree matrix $D$ is the $n \times n$ diagonal matrix with $D_{i i}=\sum_{j=1}^{n} W_{i j}$. The graph Laplacian is

$$
\begin{equation*}
L=D-W \tag{18}
\end{equation*}
$$

The graph Laplacian has many interesting properties which we list in the following result. Recall that a vector $v$ is an eigenvector of $L$ if there is a scalar $\lambda$ such that $L v=\lambda v$ in which case we say that $\lambda$ is the eigenvalue corresponding to $v$. Let $\mathcal{L}(v)=\{c v: c \in \mathbb{R}, c \neq 0\}$ be the linear space generated by $v$. If $v$ is an eigenvector with eigenvalue $\lambda$ and $c$ is any nonzero
constant, then $c v$ is an eigenvector with eigenvalue $c \lambda$. These eigenvectors are considered equivalent. In other words, $\mathcal{L}(v)$ is the set of vectors that are equivalent to $v$.

Theorem 19 The graph Laplacian L has the following properties:

1. For any vector $f=\left(f_{1}, \ldots, f_{n}\right)^{T}$,

$$
f^{T} L f=\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} W_{i j}\left(f_{i}-f_{j}\right)^{2}
$$

2. L is symmetric and positive semi-definite.
3. The smallest eigenvalue of $L$ is 0 . The corresponding eigenvector is $(1,1, \ldots, 1)^{T}$.
4. L has $n$ non-negative, real-valued eigenvalues $0=\lambda_{1} \leq \lambda_{2} \leq \cdots \leq \lambda_{k}$.
5. The number of eigenvalues that are equal to 0 is equal to the number of connected components of $G$. That is, $0=\lambda_{1}=\ldots=\lambda_{k}$ where $k$ is the number of connected components of $G$. The corresponding eigenvectors $v_{1}, \ldots, v_{k}$ are orthogonal and each is constant over one of the connected components of the graph.

Part 1 of the theorem says that $L$ is like a derivative operator. The last part shows that we can use the graph Laplacian to find the connected components of the graph.

## Proof.

(1) This follows from direct algebra.
(2) Since $W$ and $D$ are symmetric, it follow that $L$ is symmetric. The fact that $L$ is positive semi-definite folows from part (1).
(3) Let $v=(1, \ldots, 1)^{T}$. Then

$$
L v=D v-W v=\left(\begin{array}{c}
D_{11} \\
\vdots \\
D_{n n}
\end{array}\right)-\left(\begin{array}{c}
D_{11} \\
\vdots \\
D_{n n}
\end{array}\right)=\left(\begin{array}{c}
0 \\
\vdots \\
0
\end{array}\right)
$$

which equals $0 \times v$.
(4) This follows from parts (1)-(3).
(5) First suppose that $k=1$ and thus that the graph is fully connected. We already know that $\lambda_{1}=0$ and $v_{1}=(1, \ldots, 1)^{T}$. Suppose there were another eigenvector $v$ with eigenvalue 0 . Then

$$
0=v^{T} L v=\sum_{i=1}^{n} \sum_{j=1}^{n} W_{i j}(v(i)-v(j))^{2}
$$

It follows that $W_{i j}(v(i)-v(j))^{2}=0$ for all $i$ and $j$. Since $G$ is fully connected, all $W_{i j}>0$. Hence, $v(i)=v(j)$ for all $i, j$ and so $v$ is constant and thus $v \in \mathcal{L}\left(v_{1}\right)$.

Now suppose that $K$ has $k$ components. Let $n_{j}$ be the number of nodes in components $j$. We can reliable the vertices so that the first $n_{1}$ nodes correspond to the first connected component, the second $n_{2}$ nodes correspond to the second connected component and so on. Let $v_{1}=(1, \ldots, 1,0, \ldots, 0)$ where the 1 's correspond to the first component. Let Let $v_{2}=(0, \ldots, 0,1, \ldots, 1,0, \ldots, 0)$ where the 1 's correspond to the second component. Define $v_{3}, \ldots, v_{k}$ similarly. Due to the re-ordering of the vertices, $L$ has block diagonal form:

$$
L=\left(\begin{array}{cccc}
L_{1} & & & \\
& L_{2} & & \\
& & \ddots & \\
& & & L_{k}
\end{array}\right)
$$

Here, each $L_{i}$ corresponds to one of th connected components of the graph. It is easy to see that $L V-j=0$ for $j=1, \ldots, k$. Thus, each $v_{j}$, for $j=1, \ldots, k$ is an eigenvector with zero eigenvalue. Suppose that $v$ is any eigenvector with 0 eigenvalue. Arguing as before, $v$ must be constant over some component and 0 elsewhere. Hence, $v \in \mathcal{L}\left(v_{j}\right)$ for some $1 \leq j \leq k$.

Example 20 Consider the graph

$$
X_{1} \longrightarrow X_{2} \quad X_{3} \longrightarrow X_{4} \longrightarrow X_{5}
$$

and suppose that $W_{i j}=1$ if and only if there is an edge between $X_{i}$ and $X_{j}$. Then

$$
W=\left(\begin{array}{ccccc}
0 & 1 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 \\
0 & 0 & 1 & 0 & 1 \\
0 & 0 & 0 & 1 & 0
\end{array}\right) \quad D=\left(\begin{array}{ccccc}
1 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 2 & 0 \\
0 & 0 & 0 & 0 & 1
\end{array}\right)
$$

and the Laplacian is

$$
L=D-W=\left(\begin{array}{ccccc}
1 & -1 & 0 & 0 & 0 \\
-1 & 1 & 0 & 0 & 0 \\
0 & 0 & 1 & -1 & 0 \\
0 & 0 & -1 & 2 & -1 \\
0 & 0 & 0 & -1 & 0
\end{array}\right)
$$

$0-0-0-0-0 \quad 0-0-0-0-0-0-0-0-0-0-0-0-0-0-0$


Figure 26: The top shows a simple graph. The remaining plots are the eigenvectors of the graph Laplacian. Note that the first two eigenvectors correspond to the two connected components of the graph.

The eigenvalues of $W$, from smallest to largest are $0,0,1,2,3$. The eigenvectors are

$$
v_{1}=\left(\begin{array}{l}
1 \\
1 \\
0 \\
0 \\
0
\end{array}\right) \quad v_{2}=\left(\begin{array}{l}
0 \\
0 \\
1 \\
1 \\
1
\end{array}\right) \quad v_{3}=\left(\begin{array}{c}
0 \\
0 \\
-.71 \\
0 \\
.71
\end{array}\right) \quad v_{4}=\left(\begin{array}{c}
-.71 \\
.71 \\
0 \\
0 \\
0
\end{array}\right) \quad v_{5}=\left(\begin{array}{c}
0 \\
0 \\
-.41 \\
.82 \\
-.41
\end{array}\right)
$$

Note that the first two eigenvectors correspond to the connected components of the graph.

Note $f^{T} L f$ measures the smoothness of $f$ relative to the graph. This means that the higher order eigenvectors generate a basis where the first few basis elements are smooth (with respect to the graph) and the later basis elements become more wiggly.

Example 21 Figure 26 shows a graph and the corresponding eigenvectors. The two eigenvectors correspond two the connected components of the graph. The other eignvectors can be thought of as forming bases vectors within the connected components.

One approach to spectral clustering is to set

$$
W_{i j}=I\left(\left\|X_{i}-X_{j}\right\| \leq \epsilon\right)
$$

for some $\epsilon>0$ and then take the clusters to be the connected components of the graph which can be found by getting the eigenvectors of the Laplacian $L$. This is exactly equivalent to geometric graph clustering from Section 5. In this case we have gained nothing except that we have a new algorithm to find the connected components of the graph. However, there are other ways to use spectral methods for clustering as we now explain.

The idea underlying the other spectral methods is to use the Laplacian to transform the data into a new coordinate system in which clusters are easier to find. For this purpose, one typically uses a modified form of the graph Laplacian. The most commonly used weights for this purpose are

$$
W_{i j}=e^{-\left\|X_{i}-X_{j}\right\|^{2} /\left(2 h^{2}\right)}
$$

Other kernels $K_{h}\left(X_{i}, X_{j}\right)$ can be used as well. We define the symmetrized Laplacian $\mathcal{L}=$ $D^{-1 / 2} W D^{-1 / 2}$ and the random walk Laplacian $\mathcal{L}=D^{-1} W$. (We will explain the name shortly.) These are very similar and we will focus on the latter. Some authors define the random walk Laplacian to be $I-D^{-1} W$. We prefer to use the definition $\mathcal{L}=D^{-1} W$ because, as we shall see, it has a nice interpretation. The eigenvectors of $I-D^{-1} W$ and $D^{-1} W$ are the same so it makes little difference which definition is used. The main difference is that the connected components have eigenvalues 1 instead of 0 .

Lemma 22 Let $L$ be the graph Laplacian of a graph $G$ and let $\mathcal{L}$ be the random walk Laplacian.

1. $\lambda$ is an eigenvalue of $\mathcal{L}$ with eigenvector $v$ if and only if $L v=(1-\lambda) D v$.
2. 1 is an eigenvalue of $\mathcal{L}$ with eigenvector $(1, \ldots, 1)^{T}$.
3. $\mathcal{L}$ is positive semidefinite with $n$ non-negative real-valued eigenvalues.
4. The number of eigenvalues of $\mathcal{L}$ equal to 1 equals the number of connected components of $G$. Let $v_{1}, \ldots, v_{k}$ denote the eigenvectors with eigenvalues equal to 1 . The linear space spanned by $v_{1}, \ldots, v_{k}$ is spanned by the indicator functions of the connected components.

Proof. Homework.
Let $\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{n}$ be the eigenvalues of $\mathcal{L}$ with eigenvectors $v_{1}, \ldots, v_{n}$. Define

$$
Z_{i} \equiv T\left(X_{i}\right)=\sum_{j=1}^{r} \sqrt{\lambda_{j}} v_{j}(i)
$$

The mapping $T: X \rightarrow Z$ transforms the data into a new coordinate system. The numbers $h$ and $r$ are tuning parameters. The hope is that clusters are easier to find in the new parameterization.

To get some intuition for this, note that $\mathcal{L}$ has a nice probabilistic interpretation (Coifman, Lafon, Lee 2006). Consider a Markov chain on $X_{1}, \ldots, X_{n}$ where we jump from $X_{i}$ to $X_{j}$ with probability

$$
\mathbb{P}\left(X_{i} \longrightarrow X_{j}\right)=\mathcal{L}(i, j)=\frac{K_{h}\left(X_{i}, X_{j}\right)}{\sum_{s} K_{h}\left(X_{s}, X_{j}\right)}
$$

The Laplacian $\mathcal{L}(i, j)$ captures how easy it is to move from $X_{i}$ to $X_{j}$. If $Z_{i}$ and $Z_{j}$ are close in Euclidean distance, then they are connected by many high density paths through the data. This Markov chain is a discrete version of a continuous Markov chain with transition probability:

$$
P(x \rightarrow A)=\frac{\int_{A} K_{h}(x, y) d P(y)}{\int K_{h}(x, y) d P(y)}
$$

The corresponding averaging operator $\widehat{A}: f \rightarrow \tilde{f}$ is

$$
(\widehat{A} f)(i)=\frac{\sum_{j} f(j) K_{h}\left(X_{i}, X_{j}\right)}{\sum_{j} K_{h}\left(X_{i}, X_{j}\right)}
$$

which is an estimate of $A: f \rightarrow \widetilde{f}$ where

$$
A f=\frac{\int_{A} f(y) K_{h}(x, y) d P(y)}{\int K_{h}(x, y) d P(y)}
$$

The lower order eigenvectors of $\mathcal{L}$ are vectors that are smooth relative to $P$. Thus, projecting onto the first few eigenvectors parameterizes in terms of closeness with respect to the underlying density.

The steps are:
Input: $n \times n$ similarity matrix $W$.

1. Let $D$ be the $n \times n$ diagonal matrix with $D_{i i}=\sum_{j} W_{i j}$.
2. Compute the Laplacian $\mathcal{L}=D^{-1} W$.
3. Find first $k$ eigenvectors $v_{1}, \ldots, v_{k}$ of $\mathcal{L}$.
4. Project each $X_{i}$ onto the eigenvectors to get new points $\widehat{X}_{i}$.
5. Cluster the points $\widehat{X}_{1}, \ldots, \widehat{X}_{n}$ using any standard clustering algorithm.

There is another way to think about spectral clustering. Spectral methods are similar to multidimensional scaling. However, multidimensional scaling attempts to reduce dimension while preserving all pairwise distances. Spectral methods attempt instead to preserve local distances.

Example 23 Figure 27 shows a simple synthetic example. The top left plot shows the data. We apply spectral clustering with Gaussian weights and bandwidth $h=3$. The top middle


Figure 27: Top left: data. Top middle: eigenvalues. Top right: second versus third eigenvectors. Remaining plots: first six eigenvectors.
plot shows the first 20 eigenvalues. The top right plot shows the the first versus the second eigenvector. The two clusters are clearly separated. (Because the clusters are so separated, the graph is essentially disconnected and the first eigenvector is not constant. For large $h$, the graph becomes fully connected and $v_{1}$ is then constant.) The remaining six plots show the first six eigenvectors. We see that they form a Fourier-like basis within each cluster. Of course, single linkage clustering would work just as well with the original data as in the transformed data. The real advantage would come if the original data were high dimensional.

Example 24 Figure 28 shows a spectral analysis of some zipcode data. Each datapoint is a $16 \times 16$ image of a handwritten number. We restrict ourselves to the digits 1, 2 and 3. We use Gaussian weights and the top plots correspond to $h=6$ while the bottom plots correspond to $h=4$. The left plots show the first 20 eigenvalues. The right plots show a scatterplot of the second versus the third eigenvector. The three colors correspond to the three digits. We see that with a good choice of $h$, namely $h=6$, we can clearly see the digits in the plot. The original dimension of the problem is $16 \times 16=256$. That is, each image can be represented by a point in $\mathbb{R}^{256}$. However, the spectral method shows that most of the information is captured


Figure 28: Spectral analysis of some zipcode data. Top: $h=6$. Bottom: $h=4$. The plots on the right show the second versus third eigenvector. The three colors correspond to the three digits 1,2 and 3 .
by two eignvectors so the effective dimension is 2. This example also shows that the choice of $h$ is crucial.

Spectral methods are interesting. However, there are some open questions:

1. There are tuning parameters (such as $h$ ) and the results are sensitive to these parameters. How do we choose these tuning parameters?
2. Does spectral clustering perform better than density clustering?

## 8 Variable Selection

If $X \in \mathbb{R}^{d}$ is high dimensional, then it makes sense to do variable selection before clustering. There are a number of methods for doing this. But, frankly, none are very convincing. This is, in my opinion, an open problem. Here are a couple of possibilities.

### 8.1 Marginal Selection

In marginal selection, we look for variables that marginally look 'clustery." This idea was used in Chan and Hall (2010) and Wasserman, Azizyan and Singh (2014). We proceed as follows:

## Test For Multi-Modality

1. Fix $0<\alpha<1$. Let $\widetilde{\alpha}=\alpha /(n d)$.
2. For each $1 \leq j \leq d$, compute $T_{j}=\operatorname{Dip}\left(F_{n j}\right)$ where $F_{n j}$ is the empirical distribution function of the $j^{\text {th }}$ feature and $\operatorname{Dip}(F)$ is defined in (19).
3. Reject the null hypothesis that feature $j$ is not multimodal if $T_{j}>c_{n, \widetilde{\alpha}}$ where $c_{n, \widetilde{\alpha}}$ is the critical value for the dip test.

Any test of multimodality may be used. Here we describe the dip test (Hartigan and Hartigan, 1985). Let $Z_{1}, \ldots, Z_{n} \in[0,1]$ be a sample from a distribution $F$. We want to test " $H_{0}: F$ is unimodal" versus " $H_{1}: F$ is not unimodal." Let $\mathcal{U}$ be the set of unimodal distributions. Hartigan and Hartigan (1985) define

$$
\begin{equation*}
\operatorname{Dip}(F)=\inf _{G \in \mathcal{U}} \sup _{x}|F(x)-G(x)| \tag{19}
\end{equation*}
$$

If $F$ has a density $p$ we also write $\operatorname{Dip}(F)$ as $\operatorname{Dip}(p)$. Let $F_{n}$ be the empirical distribution function. The dip statistic is $T_{n}=\operatorname{Dip}\left(F_{n}\right)$. The dip test rejects $H_{0}$ if $T_{n}>c_{n, \alpha}$ where the critical value $c_{n, \alpha}$ is chosen so that, under $H_{0}, \mathbb{P}\left(T_{n}>c_{n, \alpha}\right) \leq \alpha .^{2}$

Since we are conducting multiple tests, we cannot test at a fixed error rate $\alpha$. Instead, we replace $\alpha$ with $\widetilde{\alpha}=\alpha /(n d)$. That is, we test each marginal and we reject $H_{0}$ if $T_{n}>c_{n, \widetilde{\alpha}}$. By the union bound, the chance of at least one false rejection of $H_{0}$ is at most $d \widetilde{\alpha}=\alpha / n$.

There are more refined tests such as the excess mass test given in Chan and Hall (2010), building on work by Muller and Sawitzki (1991). For simplicity, we use the dip test in this paper; a fast implementation of the test is available in $R$.

Marginal selection can obviously fail. See Figure 29 taken from Wasserman, Azizyan and Singh (2014).

[^1]

Figure 29: Three examples, each showing two clusters and two features $X(1)$ and $X(2)$. The top plots show the clusters. The bottom plots show the marginal density of $X(1)$. Left: The marginal fails to reveal any clustering structure. This example violates the marginal signature assumption. Middle: The marginal is multimodal and hence correctly identifies $X(1)$ as a relevant feature. This example satisfies the marginal signature assumption. Right: In this case, $X(1)$ is relevant but $X(2)$ is not. Despite the fact that the clusters are close together, the marginal is multimodal and hence correctly identifies $X(1)$ as a relevant feature. This example satisfies the marginal signature assumption. (Figure from Wasserman, Azizyan and Singh, 2014).

### 8.2 Sparse $k$-means

Here we discuss the approach in Witten and Tibshirani (2010). Recall that in $k$-means clustering we choose $C=\left\{c_{1}, \ldots, c_{k}\right\}$ to minimize

$$
\begin{equation*}
R_{n}(C)=\frac{1}{n} \sum_{i=1}^{n}\left\|X_{i}-\Pi_{C}\left[X_{i}\right]\right\|^{2}=\frac{1}{n} \sum_{i=1}^{n} \min _{1 \leq j \leq k}\left\|X_{i}-c_{j}\right\|^{2} \tag{20}
\end{equation*}
$$

This is equivalent to minimizing the within sums of squares

$$
\begin{equation*}
\sum_{j=1}^{k} \frac{1}{n_{j}} \sum_{s, t \in A_{j}} d^{2}\left(X_{s}, X_{t}\right) \tag{21}
\end{equation*}
$$

where $A_{j}$ is the $j^{\text {th }}$ cluster and $d^{2}(x, y)=\sum_{r=1}^{d}(x(r)-y(r))^{2}$ is squared Euclidean distance. Further, this is equivalent to maximizing the between sums of squares

$$
\begin{equation*}
B=\frac{1}{n} \sum_{s, t} d^{2}\left(X_{s}, X_{t}\right)-\sum_{j=1}^{k} \frac{1}{n_{j}} \sum_{s, t \in A_{j}} d^{2}\left(X_{s}, X_{t}\right) \tag{22}
\end{equation*}
$$

1. Input $X_{1}, \ldots, X_{n}$ and $k$.
2. Set $w=\left(w_{1}, \ldots, w_{d}\right)$ where $w_{1}=\ldots=w_{d}=1 / \sqrt{d}$.
3. Iterate until convergence:
(a) Optimize (22) over $C$ holding $w$ fixed. Find $c_{1}, \ldots, c_{k}$ from the $k$-means algorithm using distance $d_{w}\left(X_{i}, X_{j}\right)$. Let $A_{j}$ denote the $j^{\text {th }}$ cluster.
(b) Optimize (22) over $w$ holding $c_{1}, \ldots, c_{k}$ fixed. The solution is

$$
w_{r}=\frac{s_{r}}{\sqrt{\sum_{t=1}^{d} s_{t}^{2}}}
$$

where

$$
\begin{gathered}
s_{r}=\left(a_{r}-\Delta\right)_{+} \\
a_{r}=\left[\frac{1}{n} \sum_{s, t} w_{r}\left(X_{s}(r)-X_{t}(r)\right)^{2}-\sum_{j=1}^{k} \frac{1}{n_{j}} \sum_{s, t \in A_{j}} w_{r}\left(X_{s}(r)-X_{t}(r)\right)^{2}\right]_{+}
\end{gathered}
$$

and $\Delta=0$ if $\|w\|_{1}<s$ otherwise $\Delta>0$ is chosen to that $\|w\|_{1}=s$.

Figure 30: The Witten-Tibshirani Sparse $k$-means Method

Witten and Tibshirani propose replace the Euclidean norm with the weighted norm $d_{w}^{2}(x, y)=$ $\sum_{r=1}^{d} w_{r}(x(r)-y(r))^{2}$. Then they propose to maximize

$$
\begin{equation*}
B=\frac{1}{n} \sum_{s, t} d_{w}^{2}\left(X_{s}, X_{t}\right)-\sum_{j=1}^{k} \frac{1}{n_{j}} \sum_{s, t \in A_{j}} d_{w}^{2}\left(X_{s}, X_{t}\right) \tag{23}
\end{equation*}
$$

over $C$ and $w$ subject to the constraints

$$
\|w\|^{2} \leq 1, \quad\|w\|_{1} \leq s, \quad w_{j} \geq 0
$$

where $w=\left(w_{1}, \ldots, w_{d}\right)$. The optimization is done iteratively by optimizing over $C$, optimizing over $w$ and repeating. See Figure ??.

The $\ell_{1}$ norm on the weights causes some of the components of $w$ to be 0 which results in variable selection. There is no theory that shows that this method works.

## 9 Finite Mixture Models

Simple cluster structure can be discovered using mixture models. We start with a simple example. We flip a coin with success probability $\eta$. If heads, we draw $X$ from a density $p_{1}(x)$. If tails, we draw $X$ from a density $p_{0}(x)$. Then the density of $X$ is

$$
p(x)=\eta p_{1}(x)+(1-\eta) p_{0}(x)
$$

which is called a mixture of two densities $p_{1}$ and $p_{0}$. Figure 31 shows a mixture of two Gaussians distribution.

Let $Z \sim \operatorname{Bernoulli}(\eta)$ be the unobserved coin flip. Then we can also write $p(x)$ as

$$
\begin{equation*}
p(x)=\sum_{z=0,1} p(x, z)=\sum_{z=0,1} p(x \mid z) p(z) \tag{24}
\end{equation*}
$$

where $p(x \mid Z=0):=p_{0}(x), p(x \mid Z=1):=p_{1}(x)$ and $p(z)=\eta^{z}(1-\eta)^{1-z}$. Equation (24) is called the hidden variable representation. A more formal definition of finite mixture models is as follows.
[Finite Mixture Models] Let $\left\{p_{\theta}(x): \theta \in \Theta\right\}$ be a parametric class of densities. Define the mixture model

$$
p_{\psi}(x)=\sum_{j=0}^{K-1} \eta_{j} p_{\theta_{j}}(x),
$$

where the mixing coefficients $\eta_{j} \geq 0, \sum_{j=0}^{K-1} \eta_{j}=1$ and $\psi=\left(\eta_{0}, \ldots, \eta_{K-1}, \theta_{0}, \ldots, \theta_{K-1}\right)$ are the unknown parameters. We call $p_{\theta_{0}}, \ldots, p_{\theta_{K-1}}$ the component densities.

Generally, even if $\left\{p_{\theta}(x): \theta \in \Theta\right\}$ is an exponential family model, the mixture may no longer be an exponential family.

### 9.1 Mixture of Gaussians

Let $\phi\left(x ; \mu_{j}, \sigma_{j}^{2}\right)$ be the probability density function of a univariate Gaussian distribution with mean $\mu_{j}$ and variance $\sigma_{j}^{2}$. A typical finite mixture model is the mixture of Gaussians. In one dimension, we have

$$
p_{\psi}(x)=\sum_{j=0}^{K-1} \eta_{j} \phi\left(x ; \mu_{j}, \sigma_{j}^{2}\right),
$$

which has $3 K-1$ unknown parameters, due to the restriction $\sum_{j=0}^{K-1} \eta_{j}=1$.
A mixture of $d$-dimensional multivariate Gaussians is

$$
p(x)=\sum_{j=0}^{K-1} \frac{\eta_{j}}{(2 \pi)^{d / 2}\left|\Sigma_{j}\right|^{1 / 2}} \exp \left\{-\frac{1}{2}\left(x-u_{j}\right)^{T} \Sigma_{j}^{-1}\left(x-u_{j}\right)\right\} .
$$



Figure 31: A mixture of two Gaussians, $p(x)=\frac{2}{5} \phi(x ;-1.25,1)+\frac{3}{5} \phi(x ; 2.95,1)$.
There are in total

$$
K(\underbrace{\frac{d(d+1)}{2}}_{\text {\# of parameters in } \Sigma_{j}}+\underbrace{d}_{\text {\# of parameters in } u_{j}})+\underbrace{(K-1)}_{\text {\# of mixing coefficients }}=\frac{K d(d+3)}{2}+K-1
$$

parameters in the mixture of $K$ multivariate Gausssians.

### 9.2 Mixture of Multinomials

Another important mixture mode is a mixture of multinomials, which is useful for modeling count data. Let $X=\left(X_{1}, \ldots, X_{m}\right)^{T} \in \mathbb{Z}^{m}$ be the observed variable and $Z \in\{0, \ldots, K-1\}$ be the latent variable with $Z \sim \operatorname{Multinomial}\left(1, \eta_{0}, \ldots, \eta_{K-1}\right)$. We assume that

$$
X \mid Z=j \sim \operatorname{Multinomial}\left(n, \theta_{j 1}, \ldots, \theta_{j m}\right)
$$

Let $\theta_{j}=\left(\theta_{j 1}, \ldots, \theta_{j m}\right)^{T}$ and $x=\left(x_{1}, \ldots, x_{m}\right)$, we have

$$
\begin{aligned}
p(x) & =\sum_{j=0}^{K-1} \mathbb{P}(X=x \mid Z=j) \mathbb{P}(Z=j) \\
& =\frac{n!}{x_{1}!\cdots x_{m}!} \sum_{j=0}^{K-1} \eta_{j}\left(\prod_{\ell=1}^{m} \theta_{j \ell}^{x_{\ell}}\right) .
\end{aligned}
$$

There are in total $K(m-1)+(K-1)=K m-1$ parameters.

### 9.3 More General Mixture Models

We can write a finite mixture model more generally as

$$
p_{\psi}(x)=\int p_{\theta}(x) d Q(\theta)
$$

where $Q$ is the distribution that puts mass $\eta_{j}$ on $\theta_{j}$. But, more genrally, $Q$ can also be a continuous distribution.

Sometimes we can formally write a complicated distribution as a mixture of simpler distributions. For example, let $p(x)$ be the density function for a $t$ random variable with $\nu$ degrees of freedom,

$$
p(x)=\frac{\Gamma((\nu+1) / 2)}{\sqrt{\pi \nu} \Gamma(\nu / 2)\left(1+x^{2}\right)^{(\nu+1) / 2}}
$$

Some calculations show that

$$
p(x)=\int \phi(x ; 0,1 / \sqrt{u}) d Q(u)
$$

where $Q$ is the distribution for a Gamma $(\nu / 2, \nu / 2)$ random variable. Thus, a $t$ distribution is a continuous mixture of Gaussians. However, this is a mixture of scale parameters instead of location parameters as in the previous section. We focus on finite mixture models.

### 9.4 Convexity for Mixture Models

Certain forms of mixture models result in convex estimation problems. It can be useful to recognize and exploit this convexity. First, suppose that we fix the density parameters $\theta_{j}$ in the parametric mixture and denote the component densities to be $p_{0}, p_{1}, \ldots, p_{K-1}$, the only unknown parameters in the finite mixture model

$$
p_{\eta}(x)=\sum_{j=0}^{K-1} \eta_{j} p_{j}(x)
$$

are the mixing coefficients $\eta_{0}, \ldots, \eta_{K-1}$. Let $X_{1}, \ldots, X_{n}$ be $n$ random samples with $X_{i}=$ $\left(X_{i 1}, \ldots, X_{i m}\right)^{T}$, the log-likelihood function

$$
\ell(\eta)=\sum_{i=1}^{n} \log \left(\sum_{j=0}^{K-1} \eta_{j} p_{j}\left(X_{i}\right)\right)
$$

is a concave function of $\eta$. This can be seen since $\log \left(\sum_{j=0}^{K-1} \eta_{j} p_{j}\left(X_{i}\right)\right)$ is the logarithm of an affine function of $\eta$, and the logarithmic function is concave. Thus, optimizing the mixing coefficients, holding the component parameters fixed, is a convex optimization problem. For the same reason, if we hold the mixing parameters $\eta$ fixed, and the model $p_{\theta_{j}}(x)$ is a multinomial, then the log-likelihood of $\theta_{0}, \ldots, \theta_{K-1}$

$$
\ell\left(\theta_{0}, \ldots, \theta_{K-1}\right)=\sum_{i=1}^{n} \log \left(\sum_{j=0}^{K-1} \eta_{j} \prod_{\ell} \theta_{j \ell}^{x_{i} \ell}\right)
$$

is concave in $\theta$. However, it is not jointly concave in $\eta$ and $\theta$.

### 9.5 Maximum Likelihood Estimation

A finite mixture model $p_{\psi}(x)$ has parameters $\psi=\left(\eta_{0}, \ldots, \eta_{K-1}, \theta_{0}, \ldots, \theta_{K-1}\right)$. The likelihood of $\psi$ based on the observations $X_{1}, \ldots, X_{n}$ is

$$
\mathcal{L}(\psi)=\prod_{i=1}^{n} p_{\psi}\left(X_{i}\right)=\prod_{i=1}^{n}\left(\sum_{j=0}^{K-1} \eta_{j} p_{\theta_{j}}\left(X_{i}\right)\right)
$$

and, as usual, the maximum likelihood estimator is the value $\widehat{\psi}$ that maximizes $\mathcal{L}(\psi)$. Usually, the likelihood is multimodal and one seeks a local maximum instead if a global maximum.

For fixed $\theta_{0}, \ldots, \theta_{K-1}$, the log-likelihood is often a concave function of the mixing parameters $\eta_{j}$. However, for fixed $\eta_{0}, \ldots, \eta_{K-1}$, the log-likelihood is not generally concave with respect to $\theta_{0}, \ldots, \theta_{K-1}$.
One way to find $\widehat{\psi}$ is to apply your favorite optimizer directly to the log-likelihood.

$$
\ell(\psi)=\sum_{i=1}^{n} \log \left(\sum_{j=0}^{K-1} \eta_{j} p_{\theta_{j}}\left(X_{i}\right)\right)
$$

However, $\ell(\psi)$ is not jointly convex with respect to $\psi$. It is not clear which algorithm is the best to optimize such a nonconvex objective function.

A convenient and commonly used algorithm for finding the maximum likelihood estimates of a mixture model (or the more general latent variable models) is the expectation-maximization (EM) algorithm. The algorithm runs in an iterative fashion and alternates between the "E-step" which computes conditional expectations with respect to the current parameter estimate, and the "M-step" which adjusts the parameter to maximize a lower bound on the likelihood. While the algorithm can be slow to converge, its simplicity and the fact that it doesn't require a choice of step size make it a convenient choice for many estimation problems.

On the other hand, while simple and flexible, the EM algorithm is only one of many numerical procedures for obtaining a (local) maximum likelihood estimate of the latent variable models. In some cases procedures such as Newton's method or conjugate gradient may be more effective, and should be considered as alternatives to EM. In general the EM algorithm converges linearly, and may be extremely slow when the amount of missing information is large,

In principle, there are polynomial time algorithms for finding good estimates of $\psi$ based on spectral methods (refxxxx) and the method of moments (refxxxx). It appears that, at least so far, these methods are not yet practical enough to be used in routine data analysis.

### 9.6 Expectation-Maximization (EM) Algorithm

The expectation-maximization (EM) algorithm is widely used for obtaining maximum likelihood estimates for models with latent variables or missing data. The advantages of EM are that it is conceptually simple and exploits the structure of latent variable models. For finite mixture models, the EM algorithm often leads to an easy-to-implement procedure that is guaranteed to increase the likelihood monotonically, without the need for line search methods. The monotonicity property can be extremely helpful for debugging the implementation.

In this section, we first introduce a deterministic optimization strategy, block coordinate ascent, as a motivation of the methodological development of the EM algorithm. We then use the mixture of Gaussians model as a special case to provide an informal derivation of the EM algorithm. A formal derivation with geometric interoperation is then provided.

### 9.7 Main Idea

The main intuition of the expectation-maximization (EM) algorithm comes from block coordinate descent optimization, which considers a maximization problem with an objective function

$$
\begin{equation*}
\max _{x, z} f(\underbrace{x_{1}, \ldots, x_{m}}_{x}, \underbrace{z_{1}, \ldots, z_{n}}_{z}) . \tag{25}
\end{equation*}
$$

Sometimes, maximizing the objection function $f(\cdot)$ jointly with respect to $x$ and $z$ is not easy. However, if $z$ is fixed, optimizing $f$ with respect to $x$ is easy; and if $x$ is fixed, optimizing $f$ with respect to $z$ is easy. Block coordinate descent is an iterative algorithm which alternately conducts two steps: (1) Optimizing the objective function with respect to $x$ with $z$ fixed; (2) optimizing the objective function with respect to $z$ with $x$ fixed. The block coordinate descent algorithm can be described as follows:

## The Block Coordinate Descent Algorithm

Initialize $x^{(0)}$ and $z^{(0)}$.
For $t=1,2, \ldots\{$

- $x^{(t+1)} \leftarrow \operatorname{argmax}_{x} f\left(x, z^{(t)}\right)$
- $z^{(t+1)} \leftarrow \operatorname{argmax}_{z} f\left(x^{(t+1)}, z\right)$
$\}$ until convergence.

It is easy to see that the objective function values of the block coordinate decent algorithm will be nondecreasing.

The EM algorithm is similar to (but not quite the same as) applying the block coordinate descent strategy to optimize the likelihood function of latent variable models. To do this, we need to examine the likelihood function of latent variable models and figure out what are the corresponding blocks of variables $x$ and $z$ as in Equation (25). We illustrate this intuition by considering a simple example of mixture of two univariate Gaussians.

The statistical model for a mixture of two Gaussians can be written as

$$
\begin{array}{r}
Z \sim \operatorname{Bernouli}(\eta) \\
X \mid Z=j \sim N\left(\mu_{j}, 1\right) \text { for } j=0,1 \tag{27}
\end{array}
$$

Define

$$
p(x \mid Z=1):=p_{\mu_{1}}(x) \text { and } p(x \mid Z=0):=p_{\mu_{0}}(x)
$$

Let $X_{1}, \ldots, X_{n}$ be the observed data and let $Z_{1}, \ldots, Z_{n}$ be the missing data. For a clustering problem, there are two types of unknowns: (i) the parameter vector $\psi=\left(\eta, \mu_{0}, \mu_{1}\right)^{T}$ and the latent samples $Z_{1}, \ldots, Z_{n}$. The latent variables $Z_{1}, \ldots, Z_{n}$ can be used for clustering while $\psi$ can be used for evaluating the likelihood. The EM algorithm is similar to a block coordinate ascent procedure, which aims to maximize the log-likelihod function by alternatively inferring the information of $Z_{1}, \ldots, Z_{n}$ (Expectation-step) and estimating the parameter vector $\psi$ (Maximization-step).

### 9.8 An Informal Derivation of the EM Algorithm

Using the mixture of Gaussians as an example, we illustrate the informal derivation of the EM algorithm. We start with the following question:

Question 1: If the missing values $Z_{1}, \ldots, Z_{n}$ are given, how do we find the parameter value $\psi$ that maximizes the log-likelihood?

With $Z_{1}, \ldots, Z_{n}$ given, the complete log-likelihood of $\psi$ under the mixture of two Gaussians
model is

$$
\begin{align*}
\ell^{C}(\psi)= & \sum_{i=1}^{n} \log p_{\psi}\left(X_{i}, Z_{i}\right)  \tag{28}\\
= & \sum_{i=1}^{n} I\left(Z_{i}=1\right) \log \left[p_{\mu_{1}}\left(X_{i}\right) \cdot \eta\right]+\sum_{i=1}^{n} I\left(Z_{i}=0\right) \log \left[p_{\mu_{0}}\left(X_{i}\right) \cdot(1-\eta)\right]  \tag{29}\\
= & -\sum_{i=1}^{n} I\left(Z_{i}=1\right)\left[\frac{X_{i}-\mu_{1}}{2}\right]^{2}+\log (\eta) \cdot \sum_{i=1}^{n} I\left(Z_{i}=1\right) \\
& -\sum_{i=1}^{n} I\left(Z_{i}=0\right)\left[\frac{X_{i}-\mu_{0}}{2}\right]^{2}+\log (1-\eta) \cdot \sum_{i=1}^{n} I\left(Z_{i}=0\right)
\end{align*}
$$

We calculate the partial derivative of $\ell^{C}(\psi)$ with respect to $\eta$ and set it to zero:

$$
\frac{\partial \ell^{C}(\psi)}{\partial \eta}:=\frac{\sum_{i=1}^{n} I\left(Z_{i}=1\right)}{\eta}-\frac{\sum_{i=1}^{n} I\left(Z_{i}=0\right)}{1-\eta}=0
$$

which implies that

$$
\widehat{\eta}=\frac{\sum_{i=1}^{n} I\left(Z_{i}=1\right)}{n} .
$$

We then calculate the partial derivate with respect to $\mu_{1}$ and set it to be 0 ,

$$
\frac{\partial \ell^{C}(\psi)}{\partial \mu_{1}}:=\sum_{i=1}^{n} I\left(Z_{i}=1\right) X_{i}-\mu_{1} \sum_{i=1}^{n} I\left(Z_{i}=1\right)=0 .
$$

This implies that

$$
\widehat{\mu}_{1}=\frac{\sum_{i=1}^{n} I\left(Z_{i}=1\right) X_{i}}{\sum_{i=1}^{n} I\left(Z_{i}=1\right)} .
$$

Similarly, we can get the updating rule for $\mu_{0}$ as

$$
\widehat{\mu}_{0}=\frac{\sum_{i=1}^{n} I\left(Z_{i}=0\right) X_{i}}{\sum_{i=1}^{n} I\left(Z_{i}=0\right)} .
$$

From the above discussion, we see that once the missing data are given, it is easy to update the parameters. The next question is, given the current parameter configuration, how shall we infer the missing data $Z_{1}, \ldots, Z_{n}$ ?

Question 2: Given the current estimate of the parameter $\psi$, how can we infer the missing values $Z_{1}, \ldots, Z_{n}$ ?

Question 2 is more subtle since $Z_{1}, \ldots, Z_{n}$ are random quantities. Since we are treating $\psi$ as known in ths step, we can find the distribtion of $Z_{i}$ given the observed data and then we
estimate $Z_{i}$ with its mean with respect to its distribution. The distribution of $Z_{i}$ given the observed data is

$$
\mathbb{P}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right):=\mathbb{E}\left[I\left(Z_{i}=1\right) \mid X_{1}, \ldots, X_{n}\right] .
$$

For this, we use the Bayes formula (recalling that the parameter values are taken as known in this step)

$$
\begin{aligned}
\mathbb{P}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right) & =\mathbb{P}\left(Z_{i}=1 \mid X_{i}\right) \\
& =\frac{p\left(X_{i} \mid Z_{i}=1\right) \mathbb{P}\left(Z_{i}=1\right)}{p\left(X_{i}\right)} \\
& =\frac{p_{\mu_{1}}\left(X_{i}\right) \eta}{p_{\mu_{1}}\left(X_{i}\right) \eta+p_{\mu_{0}}\left(X_{i}\right)(1-\eta)} .
\end{aligned}
$$

Similarly, we know that

$$
\mathbb{P}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right)=1-\mathbb{P}\left(Z_{i}=0 \mid X_{1}, \ldots, X_{n}\right)=\frac{p_{\mu_{0}}\left(X_{i}\right)(1-\eta)}{p_{\mu_{1}}\left(X_{i}\right) \eta+p_{\mu_{0}}\left(X_{i}\right)(1-\eta)}
$$

For the mixture of two Gaussians in (26) and (27), we have

$$
\mathbb{P}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right)=\frac{\eta \exp \left[-\frac{\left(X_{i}-\mu_{1}\right)^{2}}{2}\right]}{\eta \exp \left[-\frac{\left(X_{i}-\mu_{1}\right)^{2}}{2}\right]+(1-\eta) \exp \left[-\frac{\left(X_{i}-\mu_{0}\right)^{2}}{2}\right]}
$$

To summarize: given the current parameter configuration $\psi$, we infer the conditional distribution $\mathbb{P}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right)$. But to answer Question 1, we need the exact realizations of $Z_{1}, \ldots, Z_{n}$ to update the parameter estimate; there is a gap here. To be able to alternate between the two steps of estimating the parameter $\psi$ and inferring the missing values $Z_{1}, \ldots, Z_{n}$, we need to answer the next question:

Question 3: Instead of the exact realizations $Z_{1}, \ldots, Z_{n}$, we are only given the conditional distributions $\mathbb{P}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right)$ for $i=1, \ldots, n$, how shall we find the parameter value $\psi$ that maximizes the log-likelihood?

Compared with the setup in Question 1, we see that we can no longer evaluate the complete log-likelihood $\ell^{C}(\psi)$ since we do not have the exact realizations of $Z_{1}, \ldots, Z_{n}$ (or more specifically, we are not able to evaluate the class assignment $I\left(Z_{i}=1\right)$ ). However, we have the conditional distribution $\mathbb{P}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right)$, which corresponds to the conditional expectation of $I\left(Z_{i}=1\right)$ with respect to the conditional distribution of $Z_{i}$ given $X_{1}, \ldots, X_{n}$, i.e.,

$$
\mathbb{P}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right)=\mathbb{E}\left[I\left(Z_{i}=1\right) \mid X_{1}, \ldots, X_{n}\right]
$$

One simple idea is to replace the indicator $I\left(Z_{i}=1\right)$ in (29) by the available conditional expectation $\mathbb{E}\left[I\left(Z_{i}=1\right) \mid X_{1}, \ldots, X_{n}\right]$. This is exactly what the EM algorithm does. Later,
we will explain that this corresponds to finding $\psi$ that maximizes the expected complete log-likelihood with respect to the conditional distribution of $\mathbb{P}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right)$.

More specifically, the expected complete log-likelihood of $\psi$ under the mixture of two Gaussians model is

$$
\begin{align*}
\ell^{E C}(\psi)= & \sum_{i=1}^{n} \mathbb{E}_{Z_{i} \mid X_{i}} \log p_{\psi}\left(X_{i}, Z_{i}\right)  \tag{30}\\
= & -\sum_{i=1}^{n} \mathbb{P}\left(Z_{i}=1 \mid X_{i}\right)\left[\frac{X_{i}-\mu_{1}}{2}\right]^{2}+\log (\eta) \cdot \sum_{i=1}^{n} I\left(Z_{i}=1\right) \\
& -\sum_{i=1}^{n} \mathbb{P}\left(Z_{i}=0 \mid X_{i}\right)\left[\frac{X_{i}-\mu_{0}}{2}\right]^{2}+\log (1-\eta) \cdot \sum_{i=1}^{n} I\left(Z_{i}=0\right)
\end{align*}
$$

We calculate the partial derivative of $\ell^{E C}(\psi)$ with respect to $\eta$ and set it to zero:

$$
\frac{\partial \ell^{E C}(\psi)}{\partial \eta}:=\frac{\sum_{i=1}^{n} \mathbb{P}\left(Z_{i}=1 \mid X_{i}\right)}{\eta}-\frac{\sum_{i=1}^{n} \mathbb{P}\left(Z_{i}=0 \mid X_{i}\right)}{1-\eta}=0
$$

which implies that

$$
\widehat{\eta}=\frac{\sum_{i=1}^{n} \mathbb{P}\left(Z_{i}=1 \mid X_{i}\right)}{n}
$$

We then calculate the partial derivate of $\ell^{E C}(\psi)$ with respect to $\mu_{1}$ and $\mu_{0}$, set them to be 0 , and get

$$
\widehat{\mu}_{1}=\frac{\sum_{i=1}^{n} \mathbb{P}\left(Z_{i}=1 \mid X_{i}\right) X_{i}}{\sum_{i=1}^{n} \mathbb{P}\left(Z_{i}=1 \mid X_{i}\right)} \text { and } \widehat{\mu}_{0}=\frac{\sum_{i=1}^{n} \mathbb{P}\left(Z_{i}=0 \mid X_{i}\right) X_{i}}{\sum_{i=1}^{n} \mathbb{P}\left(Z_{i}=0 \mid X_{i}\right)}
$$

Combining the results from Question 2 and Question 3, we get the following procedure for iteratively optimizing the log-likelihood of mixture of two Gaussians in (26) and (27).

The Expectation-Maximization Algorithm for the Mixture of Two Gaussians

Initialize $\psi^{(0)}:=\left(\eta^{(0)}, \mu_{1}^{(0)}, \mu_{0}^{(0)}\right)^{T}$.
For $t=1,2, \ldots\{$

- Expectation-Step (E-Step): for $i=1, \ldots, n$, calculate

$$
\begin{aligned}
\gamma_{i}^{(t+1)} & :=\mathbb{P}_{\psi^{(t)}}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right) \\
& =\frac{\eta^{(t)} \exp \left[-\frac{\left(X_{i}-\mu_{1}^{(t)}\right)^{2}}{2}\right]}{\eta^{(t)} \exp \left[-\frac{\left(X_{i}-\mu_{1}^{(t)}\right)^{2}}{2}\right]+\left(1-\eta^{(t)}\right) \exp \left[-\frac{\left(X_{i}-\mu_{0}^{(t)}\right)^{2}}{2}\right]}
\end{aligned}
$$

- Maximization-Step (M-Step): Given $\gamma_{i}^{(t+1)}$, we update the parameter $\psi$ by

$$
\begin{array}{r}
\eta^{(t+1)} \leftarrow \frac{1}{n} \sum_{i=1}^{n} \gamma_{i}^{(t+1)} ; \\
\mu_{1}^{(t+1)} \leftarrow \frac{\sum_{i=1}^{n} \gamma_{i}^{(t+1)} X_{i}}{\sum_{i=1}^{n} \gamma_{i}^{(t+1)}} \text { and } \mu_{0}^{(t+1)} \leftarrow \frac{\sum_{i=1}^{n}\left(1-\gamma_{i}^{(t+1)}\right) X_{i}}{\sum_{i=1}^{n}\left(1-\gamma_{i}^{(t+1)}\right)} .
\end{array}
$$

\} until convergence.

### 9.9 The EM Algorithm for Mixtures of Multivariate Gaussians

Consider the mixture of multivariate Gaussians model

$$
\begin{array}{r}
Z \sim \text { Multinomial }\left(1, \eta_{0}, \ldots, \eta_{K-1}\right), \\
X \mid Z=j \sim N\left(u_{j}, \Sigma_{j}\right) \text { for } j=0,1, \ldots, K-1 \tag{32}
\end{array}
$$

The EM steps becomes
The Expectation-Maximization Algorithm for the Mixture of Multivariate Gaussians
Initialize $\psi^{(0)}:=\left(\eta_{0}^{(0)}, \ldots, \eta_{K-1}^{(0)}, u_{0}^{(0)}, \ldots, u_{K-1}^{(0)}, \Sigma_{0}^{(0)}, \ldots, \Sigma_{K-1}^{(0)}\right)^{T}$.
For $t=1,2, \ldots\{$

- Expectation-Step (E-Step): for $i=1, \ldots, n$ and $j=0, \ldots, K-1$, calculate

$$
\gamma_{i j}^{(t+1)}:=\mathbb{P}_{\psi^{(t)}}\left(Z_{i}=j \mid X_{1}, \ldots, X_{n}\right)=\frac{\eta_{j}^{(t)} p_{u_{j}^{(t)}, \Sigma_{j}^{(t)}}\left(X_{i}\right)}{\sum_{\ell=0}^{K-1} \eta_{\ell}^{(t)} p_{u_{\ell}^{(t)}, \Sigma_{\ell}^{(t)}}\left(X_{i}\right)}
$$

Here $p_{u_{\ell}^{(t)}, \Sigma_{\ell}^{(t)}}\left(X_{i}\right)$ is a Gaussian density with mean $u_{j}^{(t)}$ and covariance $\Sigma_{j}^{(t)}$.

- Maximization-Step (M-Step): Given $\gamma_{i j}^{(t+1)}$, we update the parameter $\psi$ by

$$
\begin{array}{r}
\text { For } j=0, \ldots, K-1: \\
\eta_{j}^{(t+1)} \leftarrow \frac{1}{n} \sum_{i=1}^{n} \gamma_{i j}^{(t+1)} ; \\
u_{j}^{(t+1)} \leftarrow \frac{\sum_{i=1}^{n} \gamma_{i j}^{(t+1)} X_{i}}{\sum_{i=1}^{n} \gamma_{i j}^{(t+1)}} ; \\
\Sigma_{j}^{(t+1)} \leftarrow \frac{\sum_{i=1}^{n} \gamma_{i j}^{(t+1)}\left(X_{i}-u_{j}^{(t+1)}\right)\left(X_{i}-u_{j}^{(t+1)}\right)^{T}}{\sum_{i=1}^{n} \gamma_{i j}^{(t+1)}},
\end{array}
$$

\} until convergence.

From the above description we see that in the M-step of the EM algorithm we calculate the weighted sample mean and covariance, where the weight $\gamma_{i j}$ is the posterior probability that example $i$ was generated from mixture component $j$. These weights are sometimes called "responsibilities," suggesting that it is the probability that the latent variable $j$ was responsible for example $i$.

### 9.10 A Geometric Perspective on the EM Algorithm

The EM algorithm is an iterative procedure that has two main steps. As in the mixture of Gaussians example, in the E-step, it tries to "guess" the values of the missing data $Z_{1}, \ldots, Z_{n}$. In the M-step, it updates the parameters of our model based on our guesses.


Figure 32: A geometric illustration of the EM Algorithm. Our goal is to find a local maxima of the $\log$-likelihood function $\ell(\psi)$. In each iteration of the EM algorithm, we find a lower bound function and try to maximize the lower bound function.

Figure 32 illustrates the geometric intuition of the EM algorithm. Using the mixture of Gaussians as an example, we see that our goal is to find a local maxima of a nonconcave likelihood function

$$
\ell(\psi):=\sum_{i=1}^{n} \log p_{\psi}\left(X_{i}\right)=\sum_{i=1}^{n} \log \left(\sum_{j=0}^{K-1} \eta_{j} p_{u_{j}, \Sigma_{j}}\left(X_{i} \mid Z_{i}=j\right)\right) .
$$

We first find an initial value $\psi^{(0)}$. Geometrically, the EM algorithm tries to find a lower bound function $F_{\psi(0)}$, which satisfies

$$
F_{\psi^{(0)}}(\psi) \leq \ell(\psi) \text { for all } \psi,
$$

and

$$
F_{\psi^{(0)}}\left(\psi^{(0)}\right)=\ell\left(\psi^{(0)}\right) .
$$

The EM algorithm requires that the lower bound function $F_{\psi^{(0)}}(\psi)$ is easy to be optimized. To proceed, we maximize this lower bound function and get $\psi^{(1)}$ :

$$
\psi^{(1)}=\underset{\psi}{\operatorname{argmax}} F_{\psi^{(0)}}(\psi) .
$$

We then continue with the same strategy to find a new lower bound function $F_{\psi^{1}}(\psi)$ which satisfies $F_{\psi^{(1)}}(\psi) \leq \ell(\psi)$ for all $\psi$ and $F_{\psi^{(1)}}\left(\psi^{(1)}\right)=\ell\left(\psi^{(1)}\right)$. The algorithm proceeds by maximizing the new lower bound function to get $\psi^{(2)}$ :

$$
\psi^{(2)}=\underset{\psi}{\operatorname{argmax}} F_{\psi^{(1)}}(\psi) .
$$

We iterate the above process until the algorithm converges. In the next section, we will formally prove that such a procedure is guaranteed to converge to a local maxima of $\ell(\psi)$. By examining the above procedure, it is obvious that the initial value $\psi^{(0)}$ affects the final solution. In applications, a common practice is to run the EM algorithm for different initial values $\psi^{(0)}$ and use the one that gets the largest value of $\ell(\psi)$.

### 9.11 A Formal Derivation of the EM Algorithm

Let $X_{1}, \ldots, X_{n}$ be a random sample. We want to fit a mixture model with the density

$$
p_{\psi}(x)=\sum_{j=0}^{K-1} p_{\theta_{j}}(x) \eta_{j}
$$

where $\eta_{j}=\mathbb{P}(Z=j)$ satisfying $\sum_{j=0}^{K-1} \eta_{j}=1$, and $p_{\theta_{j}}(x)=p(x \mid Z=j)$. Our goal is to estimate the parameters $\psi=\left\{\theta_{0}, \eta_{0}, \ldots, \theta_{K-1}, \eta_{K-1}\right\}$ by maximum likelihood. We begin with some terminology:

- $X$ : observed variable,
- $Z$ : latent variable/unobserved varaible,
- $\ell^{\mathrm{M}}(\psi)=\log p_{\psi}\left(X_{1}, \ldots, X_{n}\right)$ : marginal likelihood,
- $\ell^{\mathrm{C}}(\psi)=\log p_{\psi}\left(X_{1}, \ldots, X_{n}, Z_{1}, \ldots, Z_{n}\right)$ : complete likelihood.

We also define

$$
\gamma_{i j}=\mathbb{P}_{\psi}\left(Z_{i}=j \mid X_{i}\right) .
$$

It is easy to see that $\sum_{j=0}^{K-1} \gamma_{i j}=1$. Therefore, $\gamma_{i 1}, \ldots, \gamma_{i(K-1)}$ specify a probability distribution which indicates the posterior class assignment of a point $X_{i}$ under the parameter configuration $\psi$.

The formal derivation of the EM algorithm hinges on the concavity of the logarithm, and unfolds by bounding the log-likelihood from below using Jensen's inequality. We derive the EM algorithm according to the following steps:

Task 1: Given the current parameter configuration $\psi^{(o l d)}$, we want to derive a lower bound function $F(\psi)$ of $\ell^{\mathrm{M}}(\psi)$ such that

$$
F(\psi) \leq \ell^{\mathrm{M}}(\psi) \text { for all } \psi,
$$

and $F\left(\psi^{(\text {old })}\right)=\ell^{\mathrm{M}}\left(\psi^{(\text {old })}\right)$. We define such a function to be $F_{\psi^{\text {(old })}}(\psi)$.
The key of this derivation is Jensen's inequality for the logarithmic function, which states that

$$
\log \mathbb{E} X \geq \mathbb{E} \log X
$$

Let $\gamma=\left\{\gamma_{i j}\right\}_{i=1, \ldots, n ; j=0, \ldots, K-1}$. We have

$$
\begin{aligned}
\ell^{\mathrm{M}}(\psi) & =\sum_{i=1}^{n} \log p_{\psi}\left(X_{i}\right) \\
& =\sum_{i=1}^{n} \log \left[\sum_{j=0}^{K-1} p_{\psi}\left(X_{i}, Z_{i}=j\right)\right] \\
& =\sum_{i=1}^{n} \log \left[\sum_{j=0}^{K-1} \gamma_{i j} \frac{p_{\psi}\left(X_{i}, Z_{i}=j\right)}{\gamma_{i j}}\right] \quad \text { (Mean field trick) } \\
& \geq \sum_{i=1}^{n} \sum_{j=0}^{K-1} \gamma_{i j} \log \left[\frac{p_{\psi}\left(X_{i}, Z_{i}=j\right)}{\gamma_{i j}}\right] \quad \text { (Jensen's inequality) } \\
& =\sum_{i=1}^{n} \sum_{j=0}^{K-1} \gamma_{i j} \log p_{\psi}\left(X_{i}, Z_{i}=j\right)-\sum_{i=1}^{n} \sum_{j=0}^{K-1} \gamma_{i j} \log \gamma_{i j} \\
& :=Q(\psi, \gamma)
\end{aligned}
$$

Therefore, $Q(\psi, \gamma)$ is a function that lower bounds $\ell^{\mathrm{M}}(\psi)$ for any $\psi$ and $\gamma$, where $\psi$ is the unknown parameter and $\gamma$ carries the information of the missing data. The EM algorithm applies the block coordinate ascent strategy to maximize $Q(\psi, \gamma)$. More specifically, the EM algorithm iterates the following two steps until converge:

- E-step: $\gamma^{(\text {new })} \leftarrow \operatorname{argmax}_{\gamma} Q\left(\psi^{(\text {old })}, \gamma\right)$,
- M-step: $\psi^{(\text {(old })} \leftarrow \operatorname{argmax}_{\psi} Q\left(\psi, \gamma^{(\text {new })}\right)$.

To map the above updating rule back to that in Figure 32, we define

$$
\begin{equation*}
F_{\psi^{\text {(old) })}}(\psi):=Q\left(\psi, \gamma^{(\text {new })}\right) . \tag{33}
\end{equation*}
$$

It is easy to see that $F_{\psi^{\text {(old })}}(\psi) \leq \ell^{\mathrm{M}}(\psi)$ for any $\psi$ and $F_{\psi^{\text {(old) }}( }\left(\psi^{(\text {old })}\right)=\ell^{\mathrm{M}}\left(\psi^{\text {(old })}\right)$. In the finite mixture model setting, the next theorem shows that the E-step update has a closed form solution.

Theorem 25 For finite mixture models, the E-step update has a closed form solution

$$
\gamma_{i j}^{(\text {new })}=\mathbb{P}_{\psi^{\text {(odd) }}}\left(Z_{i}=j \mid X_{i}\right),
$$

where $\psi^{(\mathrm{old})}$ is the current parameter configuration.

Proof. In the above derivation, we know that $Q(\psi, \gamma) \leq \ell^{\mathrm{M}}(\psi)$ for any $\psi$ and $\gamma$. To show that $\gamma_{i j}^{(\text {new })}=\mathbb{P}_{\psi^{\text {(old })}}\left(Z_{i}=j \mid X_{i}\right)$ is the maximizer of $Q\left(\psi^{(\text {(old })}, \gamma\right)$, it suffices to prove that

$$
Q\left(\psi^{(\mathrm{old})}, \gamma^{(\mathrm{new})}\right)=\ell^{\mathrm{M}}\left(\psi^{(\mathrm{old})}\right)
$$

To verify this, we have

$$
\begin{aligned}
Q\left(\psi^{(\text {old })}, \gamma^{(\text {new })}\right) & =\sum_{i=1}^{n} \sum_{j=0}^{K-1} \gamma_{i j}^{(\text {new })} \log \left[\frac{p_{\psi}^{\text {(old) }}\left(X_{i}, Z_{i}=j\right)}{\gamma_{i j}^{(\text {new })}}\right] \\
& =\sum_{i=1}^{n} \sum_{j=0}^{K-1} \mathbb{P}_{\psi^{\text {(old) }}}\left(Z_{i}=j \mid X_{i}\right) \cdot \log \left[\frac{\left.\mathbb{P}_{\psi^{\text {(old) }}}\left(Z_{i}=j \mid X_{i}\right) p_{\psi^{\text {(old) })}}\left(X_{i}\right)\right)}{\mathbb{P}_{\psi^{\text {(old) }}}\left(Z_{i}=j \mid X_{i}\right)}\right] \\
& =\sum_{i=1}^{n} \log p_{\psi^{\text {(old) })}}\left(X_{i}\right)[\underbrace{\sum_{j=0}^{K-1} \mathbb{P}_{\psi^{\text {(old })}}\left(Z_{i}=j \mid X_{i}\right)}_{=1}] \\
& =\sum_{i=1}^{n} \log p_{\psi^{\text {(old) }}}\left(X_{i}\right) \\
& =\ell^{\mathrm{M}}\left(\psi^{\text {(old })}\right) .
\end{aligned}
$$

This finishes the proof. $\square$ Once the lower bound function $F(\psi)$ is derived, our next task is to optimize it with the parameter $\psi$. For this, we need to handle the next task:

Task 2: Given the current lower bound function $F_{\psi(\text { old })}(\psi)$ of $\ell^{\mathrm{M}}(\psi)$, we update the parameter $\psi$ by solving

$$
\psi^{(\text {new })}=\underset{\psi}{\operatorname{argmax}} F_{\psi(\text { old })}(\psi) .
$$

The next theorem shows that $\psi^{(\text {new })}$ can be updated by maximizing the expected complete log-likelihood with respect to the conditional distribution of $\mathbb{P}_{\psi^{(\text {old })}}\left(Z_{i}=1 \mid X_{1}, \ldots, X_{n}\right)$.

Theorem 26 Under a finite mixture model $p_{\psi}(x)=\sum_{j=0}^{K-1} p_{\theta_{j}}(x) \eta_{j}$, we have the following updating rule in the $M$-step:

$$
\begin{array}{r}
\theta_{j}^{(\text {new })}=\underset{\psi_{j}}{\operatorname{argmax}} \sum_{i=1}^{n} \sum_{j=0}^{K-1} \gamma_{i j}^{(\text {old })} \log p_{\theta_{j}^{\text {(old) }}}\left(X_{i}, Z_{i}=j\right) \\
\eta_{j}^{(\text {new })}=\frac{1}{n} \sum_{i=1}^{n} \gamma_{i j}^{(\text {new })} \tag{35}
\end{array}
$$

where $\gamma_{i j}^{(\text {new })}=\mathbb{P}_{\psi^{(\text {(old })}}\left(Z_{i}=j \mid X_{i}\right)$.

Proof. From (33), we know that

$$
\begin{aligned}
F_{\psi^{(\text {(old })}}(\psi) & =Q\left(\psi, \gamma^{(\text {new })}\right) \\
& =\sum_{i=1}^{n} \sum_{j=0}^{K-1} \gamma_{i j}^{(\text {new })} \log \left[\frac{p_{\theta_{j}}\left(X_{i}, Z_{i}=j\right)}{\gamma_{i j}^{\text {(new })}}\right] \\
& =\sum_{i=1}^{n} \sum_{j=0}^{K-1} \gamma_{i j}^{(\text {new })} \log \left[\frac{p_{\theta_{j}}\left(X_{i} \mid Z_{i}=j\right) \eta_{j}}{\gamma_{i j}^{(\text {new })}}\right]
\end{aligned}
$$

Optimizing $F_{\psi(\text { old })}(\psi)$ with respect to $\theta_{j}$, we get (34). To optimize $F_{\psi(\text { old })}(\psi)$ with respect to $\eta_{j}$, we need to solve a constrained optimization

$$
\begin{array}{r}
\max _{\eta_{0}, \ldots, \eta_{K-1}} \sum_{i=1}^{n} \sum_{j=0}^{K-1} \gamma_{i j}^{(\text {new })} \log \eta_{j}, \\
\text { subject to } \eta_{0}, \ldots, \eta_{K-1} \geq 0 \text { and } \sum_{j=0}^{K-1} \eta_{j}=1 .
\end{array}
$$

Solving the above optimization, we get the desired result.

### 9.12 Convergence of the EM Algorithm

The next theorem shows that the convergence of the EM algorithm to a local maximizer is guaranteed.

Theorem 27 (Convergence of the EM Algorithm ) When the marginal likelihood $\ell^{\mathrm{M}}(\psi)$ is bounded from above, the EM algorithm is guaranteed to converge.

Proof. From the derivation in the last section, we see that the EM algorithm is a block coordinate descent algorithm for maximizing the auxiliary function $Q(\psi, \gamma)$. The convergence of the EM algorithm follows from the convergence of the block coordinate descent algorithm under the bounded likelihood condition.

Remark 28 Since the marginal log-likelihood functions of the finite mixture models are in general nononvex, the EM algorithm may suffer from local maxima. In applications, we generally try different random initializations and pick the one that achieves the largest likelihood function.

### 9.13 Old Faithful Geyser Data Example

We study a version of the eruptions data "Old Faithful" geyser in Yellowstone national park. The geyser was named by the Washburn expedition of 1870 . They were impressed by its size and frequency. It is not the biggest or most regular geyser in Yellowstone but it is the biggest regular geyser. Furthermore, it has been erupting in nearly the same fashion throughout the recorded history of Yellowstone. Through the years, it has become one of the most studied geysers in the park.

The data in this study are continuous measurements from August 1 to August 15, 1985. There are two variables with 299 observations. The first variable ,"Duration", represents the numeric eruption time in minutes. The second variable, "waiting", represents the waiting time to next eruption. This data is believed to have two modes. We fit a mixture of two Gaussians using EM algorithm. To illustrate the EM step, we purposely choose a bad starting point. The EM algorithm quickly converges in six steps. Figure 33 illustrates the fitted densities for all the six steps. We see that even though the starting density is unimodal, it quickly becomes bimodal.


Figure 33: Fitting a mixture of two Gaussians on the Old Faithful Geyser data. The initial values are $\eta_{0}=\eta_{1}=0.5 . \quad u_{0}=(4,70)^{T}, u_{1}=(3,60)^{T}, \Sigma_{1}=\Sigma_{2}=\left(\begin{array}{ll}0.8 & 7 \\ 7 & 70\end{array}\right)$. We see that even though the starting density is not bimodal, the EM algorithm converges quickly to a bimodal density.

## 10 Examples

Example 29 Figures 1 and 2 shows some synthetic examples where the clusters are meant to be intuitively clear. In Figure 1 there are two blob-like clusters. Identifying clusters like this is easy. Figure 2 shows four clusters: a blob, two rings and a half ring. Identifying clusters with unusual shapes like this is not quite as easy. To the human eye, these certainly look like clusters. But what makes them clusters?

Example 30 (Gene Clustering) In genomic studies, it is common to measure the expression levels of d genes on $n$ people using microarrays (or gene chips). The data (after much simplification) can be represented as an $n \times d$ matrix $X$ where $X_{i j}$ is the expression level of gene $j$ for subject $i$. Typically $d$ is much larger than $n$. For example, we might have


Figure 34: Some curves from a dataset of 472 curves. Each curve is a radar waveform from the Topex/Poseidon satellite.
$d \approx 5,000$ and $n \approx 50$. Clustering can be done on genes or subjects. To find groups of similar people, regard each row as a data vector so we have $n$ vectors $X_{1}, \ldots, X_{n}$ each of length $d$. Clustering can then be used to place the subjects into similar groups.

Example 31 (Curve Clustering) Sometimes the data consist of a set of curves $f_{1}, \ldots, f_{n}$ and the goal is to cluster similarly shaped clusters together. For example, Figure 34 shows a small sample of curves a from a dataset of 472 curves from Frappart (2003). Each curve is a radar waveform from the Topex/Poseidon satellite which used to map the surface topography of the oceans. ${ }^{3}$ One question is whether the 472 curves can be put into groups of similar shape.

Example 32 (Supernova Clustering) Figure 35 shows another example of curve clustering. Briefly, each data point is a light curve, essentially brightness versus time. The top two plots show the light curves for two types of supernovae called "Type Ia" and "other." The

[^2]

Figure 35: Light curves for supernovae. The top two plots show the light curves for two types of supernovae. The bottom two plots show the results of clustering the curves into two groups, without using knowledge of their labels.
bottom two plots show what happens if we throw away the labels ("Type Ia" and "other") and apply a clustering algorithm ( $k$-means clustering). We see that the clustering algorithm almost completely recovers the two types of supernovae.

## 11 Summary

The main clustering methods are:

1. density-based
2. $k$-means
3. hierarchical
4. spectral.

Each method has strengths and weaknesses. They all involve tuning parameters that must be chosen carefully.

## 12 Bibliographic Remarks

$k$-means clustering goes back to Stuart Lloyd who apparently came up with the algorithm in 1957 although he did not publish it until 1982. See [?]. Another key reference is [?]. Similar ideas appear in [?]. The related area of mixture models is discussed at length in McLachlan
and Basford (1988). $k$-means is actually related to principal components analysis; see Ding and He (2004) and Zha, He, Ding, Simon and Gu (2001). The probabilistic behavior of random geometric graphs is discussed in detail in [?].


[^0]:    ${ }^{1}$ In general, we can use a bandwidth matrix $H$ in the estimator, with $\widehat{p}(x) \equiv \widehat{p}_{H}(x)=\frac{1}{n} \sum_{i=1}^{n} K_{H}\left(x-X_{i}\right)$ where $K_{H}(x)=|H|^{-\frac{1}{2}} K\left(H^{-\frac{1}{2}} x\right)$.

[^1]:    ${ }^{2}$ Specifically, $c_{n, \alpha}$ can be defined by $\sup _{G \in \mathcal{U}} P_{G}\left(T_{n}>c_{n, \alpha}\right)=\alpha$. In practice, $c_{n, \alpha}$ can be defined by $P_{U}\left(T_{n}>c_{n, \alpha}\right)=\alpha$ where $U$ is Unif(0,1). Hartigan and Hartigan (1985) suggest that this suffices asymptotically.

[^2]:    ${ }^{3}$ See http://topex-www.jpl.nasa.gov/overview/overview.html. The data are available at "Working Group on Functional and Operator-based Statistics" a web site run by Frederic Ferrarty and Philippe Vieu. The address is http://www.math.univ-toulouse.fr/staph/npfda/. See also http://podaac.jpl.nasa.gov/DATA_CATALOG/topexPoseidoninfo.html.

