Normalized Graph Cuts
Some Observations

Jean Gallier

CIS Department
University of Pennsylvania

jean@cis.upenn.edu

January 25, 2014

Figure 1: Dog Logic

1. Graph Clustering

Given a set of data, the goal of clustering is to partition the data into different groups according to their similarities.

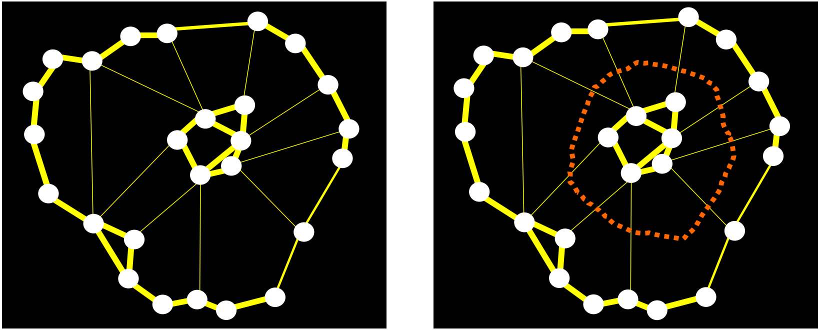

Figure 2: A weighted graph and its partition into two clusters.

When the data is given in terms of a similarity graph G, where the weight $w_{i\,j}$ between two nodes v_i and v_j is a measure of similarity of v_i and v_j, the problem can be stated as follows:

When the data is given in terms of a similarity graph G, where the weight $w_{i\,j}$ between two nodes v_i and v_j is a measure of similarity of v_i and v_j, the problem can be stated as follows:

Find a partition $(A_1, \ldots, A_K)$ of the set of nodes V into different groups such that the *edges between different groups have very low weight* (which indicates that the points in different clusters are dissimilar), and the *edges within a group have high weight* (which indicates that points within the same cluster are similar).

When the data is given in terms of a similarity graph G, where the weight $w_{i\,j}$ between two nodes v_i and v_j is a measure of similarity of v_i and v_j, the problem can be stated as follows:

Find a partition $(A_1, \ldots, A_K)$ of the set of nodes V into different groups such that the *edges between different groups have very low weight* (which indicates that the points in different clusters are dissimilar), and the *edges within a group have high weight* (which indicates that points within the same cluster are similar).

The above graph clustering problem can be formalized as an optimization problem, using the notion of *cut*.

2. Weigted Graphs, Cuts , Laplacians

Definition 1

A *weighted graph* is a pair $G = (V, W)$, where $V = \{v_1, \ldots, v_m\}$ is a set of *nodes* or *vertices*, and W is a symmetric matrix called the *weight matrix*, such that $w_{i\,j} \geq 0$ for all $i, j \in \{1, \ldots, m\}$, and $w_{i\,i} = 0$ for $i = 1, \ldots, m$. We say that a set $\{v_i, v_j\}$ is an edge iff $w_{i\,j} > 0$. The corresponding (undirected) graph (V, E) with $E = \{\{v_i, v_j\} \mid w_{i\,j} > 0\}$, is called the *underlying graph* of G.

2. Weigted Graphs, Cuts , Laplacians

Definition 1

A *weighted graph* is a pair $G = (V, W)$, where $V = \{v_1, \ldots, v_m\}$ is a set of *nodes* or *vertices*, and W is a symmetric matrix called the *weight matrix*, such that $w_{i\,j} \geq 0$ for all $i, j \in \{1, \ldots, m\}$, and $w_{i\,i} = 0$ for $i = 1, \ldots, m$. We say that a set $\{v_i, v_j\}$ is an edge iff $w_{i\,j} > 0$. The corresponding (undirected) graph (V, E) with $E = \{\{v_i, v_j\} \mid w_{i\,j} > 0\}$, is called the *underlying graph* of G.

We can think of the weight $w_{i\,j}$ of an edge $\{v_i, v_j\}$ as a degree of similarity (or affinity) in an image, or a cost in a network.

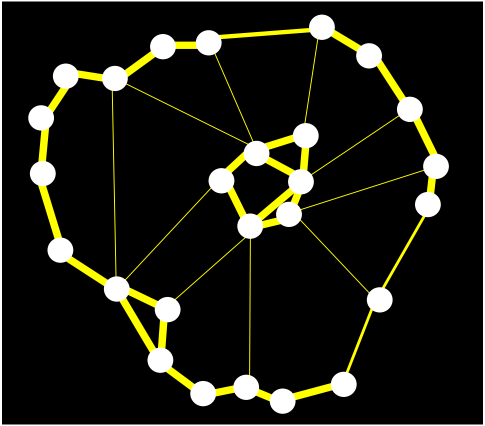

Figure 3: A weighted graph.

The thickness of an edge corresponds to the magnitude of its weight.

For every node $v_i \in V$, the *degree* $d(v_i)$ of v_i is the sum of the weights of the edges adjacent to v_i:

$$d(v_i) = \sum_{j=1}^{m} w_{i\,j}.$$

For every node $v_i \in V$, the *degree* $d(v_i)$ of v_i is the sum of the weights of the edges adjacent to v_i:

$$d(v_i) = \sum_{j=1}^{m} w_{i\,j}.$$

Note that in the above sum, only nodes v_j such that there is an edge $\{v_i, v_j\}$ have a nonzero contribution. Such nodes are said to be *adjacent* to v_i.

For every node $v_i \in V$, the *degree* $d(v_i)$ of v_i is the sum of the weights of the edges adjacent to v_i:

$$d(v_i) = \sum_{j=1}^{m} w_{i\,j}.$$

Note that in the above sum, only nodes v_j such that there is an edge $\{v_i, v_j\}$ have a nonzero contribution. Such nodes are said to be *adjacent* to v_i.

The *degree matrix* D is defined by $D = \mathrm{diag}(d(v_1), \ldots, d(v_m))$.

Given any subset of nodes $A \subseteq V$, we define the *volume* $\mathrm{vol}(A)$ of A as the sum of the weights of all edges adjacent to nodes in A:

$$\mathrm{vol}(A) = \sum_{v_i \in A} d(v_i) = \sum_{v_i \in A} \sum_{j=1}^{m} w_{i\,j}.$$

Given any subset of nodes $A \subseteq V$, we define the *volume* $\mathrm{vol}(A)$ of A as the sum of the weights of all edges adjacent to nodes in A:

$$\mathrm{vol}(A) = \sum_{v_i \in A} d(v_i) = \sum_{v_i \in A} \sum_{j=1}^{m} w_{i\,j}.$$

The notions of degree and volume are illustrated in Figure 4.

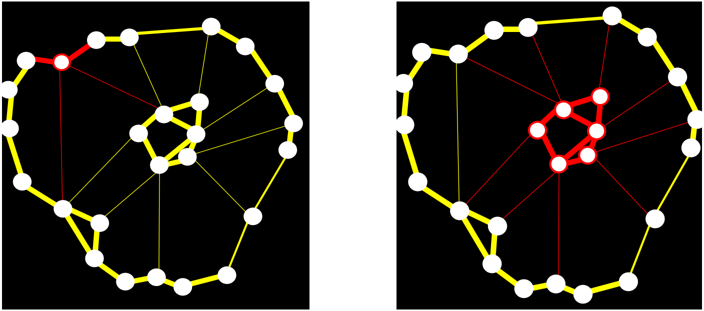

Figure 4: Degree and volume.

Observe that $\mathrm{vol}(A) = 0$ if A consists of isolated vertices, that is, if $w_{i\,j} = 0$ for all $v_i \in A$. Thus, it is best to assume that G does not have isolated vertices.

Observe that $\mathrm{vol}(A) = 0$ if A consists of isolated vertices, that is, if $w_{i\,j} = 0$ for all $v_i \in A$. Thus, it is best to assume that G does not have isolated vertices.

Given any two subset $A, B \subseteq V$ (not necessarily distinct), we define $\mathrm{links}(A, B)$ by

$$\mathrm{links}(A, B) = \sum_{v_i \in A, v_j \in B} w_{i\,j}.$$

Observe that $\mathrm{vol}(A) = 0$ if A consists of isolated vertices, that is, if $w_{i\,j} = 0$ for all $v_i \in A$. Thus, it is best to assume that G does not have isolated vertices.

Given any two subset $A, B \subseteq V$ (not necessarily distinct), we define $\mathrm{links}(A, B)$ by

$$\mathrm{links}(A, B) = \sum_{v_i \in A, v_j \in B} w_{i\,j}.$$

Since the matrix W is symmetric, we have

$$\mathrm{links}(A, B) = \mathrm{links}(B, A).$$

The quantity $\mathrm{links}(A, \overline{A}) = \mathrm{links}(\overline{A}, A)$, where $\overline{A} = V - A$ denotes the complement of A in V, *measures how many links escape from* A (and $\overline{A}$), and the quantity $\mathrm{links}(A, A)$ *measures how many links stay within* A *itself*.

The quantity $\mathrm{links}(A, \overline{A}) = \mathrm{links}(\overline{A}, A)$, where $\overline{A} = V - A$ denotes the complement of A in V, *measures how many links escape from A* (and $\overline{A}$), and the quantity $\mathrm{links}(A, A)$ *measures how many links stay within A itself*.

The quantity

$$\mathrm{cut}(A) = \mathrm{links}(A, \overline{A})$$

is often called the *cut* of A, and the quantity

$$\mathrm{assoc}(A) = \mathrm{links}(A, A)$$

is often called the *association* of A.

The quantity $\mathrm{links}(A, \overline{A}) = \mathrm{links}(\overline{A}, A)$, where $\overline{A} = V - A$ denotes the complement of A in V, *measures how many links escape from A* (and $\overline{A}$), and the quantity $\mathrm{links}(A, A)$ *measures how many links stay within A itself*.

The quantity

$$\mathrm{cut}(A) = \mathrm{links}(A, \overline{A})$$

is often called the *cut* of A, and the quantity

$$\mathrm{assoc}(A) = \mathrm{links}(A, A)$$

is often called the *association* of A.

Clearly,

$$\mathrm{cut}(A) + \mathrm{assoc}(A) = \mathrm{vol}(A).$$

The notions of cut is illustrated in Figure 5.

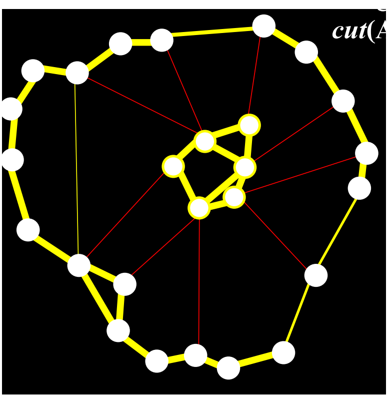

Figure 5: A Cut involving the set of nodes in the center and the nodes on the perimeter.

We now define the most important concept of this talk: The Laplacian matrix of a graph. Actually, as we will see, it comes in several flavors.

We now define the most important concept of this talk: The Laplacian matrix of a graph. Actually, as we will see, it comes in several flavors.

Definition 2

Given any weighted graph $G = (V, W)$ with $V = \{v_1, \ldots, v_m\}$, the *(unnormalized) graph Laplacian $L(G)$ of G* is defined by

$$L(G) = D(G) - W,$$

where $D(G) = \mathrm{diag}(d_1, \ldots, d_m)$ is the degree matrix of G (a diagonal matrix), with

$$d_i = \sum_{j=1}^{m} w_{i\,j}.$$

We now define the most important concept of this talk: The Laplacian matrix of a graph. Actually, as we will see, it comes in several flavors.

Definition 2

Given any weighted graph $G = (V, W)$ with $V = \{v_1, \ldots, v_m\}$, the *(unnormalized) graph Laplacian $L(G)$ of G* is defined by

$$L(G) = D(G) - W,$$

where $D(G) = \mathrm{diag}(d_1, \ldots, d_m)$ is the degree matrix of G (a diagonal matrix), with

$$d_i = \sum_{j=1}^{m} w_{i\,j}.$$

As usual, unless confusion arises, we write L instead of $L(G)$.

It is clear that each row of L sums to 0, so the vector $\mathbf{1}$ is the nullspace of L, but it is less obvious that L is positive semidefinite. An easy way to prove this is to evaluate the quadratic form $x^{\top} L x$.

It is clear that each row of L sums to 0, so the vector $\mathbf{1}$ is the nullspace of L, but it is less obvious that L is positive semidefinite. An easy way to prove this is to evaluate the quadratic form $x^\top L x$.

Proposition 1

For any $m \times m$ symmetric matrix W, if we let $L = D - W$ where D is the degree matrix of $W = (w_{ij})$, then we have

$$x^\top L x = \frac{1}{2} \sum_{i,j=1}^{m} w_{i\,j}(x_i - x_j)^2 \quad \text{for all } x \in \mathbb{R}^m.$$

Consequently, $x^\top L x$ does not depend on the diagonal entries in W, and if $w_{i\,j} \geq 0$ for all $i, j \in \{1, \ldots, m\}$, then L is positive semidefinite.

Proposition 1 immediately implies the following facts: For any weighted graph $G = (V, W)$,

Proposition 1 immediately implies the following facts: For any weighted graph $G = (V, W)$,

1. The eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_m$ of L are real and nonnegative, and there is an orthonormal basis of eigenvectors of L.
2. The smallest eigenvalue λ_1 of L is equal to 0, and $\mathbf{1}$ is a corresponding eigenvector.

It turns out that the dimension of the nullspace of L (the eigenspace of 0) is equal to the number of connected components of the underlying graph of G.

Proposition 1 immediately implies the following facts: For any weighted graph $G = (V, W)$,

1. The eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_m$ of L are real and nonnegative, and there is an orthonormal basis of eigenvectors of L.
2. The smallest eigenvalue λ_1 of L is equal to 0, and $\mathbf{1}$ is a corresponding eigenvector.

It turns out that the dimension of the nullspace of L (the eigenspace of 0) is equal to the number of connected components of the underlying graph of G.

Normalized variants of the graph Laplacian are needed, especially in applications to graph clustering.

These variants make sense only if G has no isolated vertices, which means that every row of W contains some strictly positive entry. In this case, the degree matrix D contains positive entries, so it is invertible and $D^{-1/2}$ makes sense; namely

$$D^{-1/2} = \mathrm{diag}(d_1^{-1/2}, \ldots, d_m^{-1/2}).$$

These variants make sense only if G has no isolated vertices, which means that every row of W contains some strictly positive entry. In this case, the degree matrix D contains positive entries, so it is invertible and $D^{-1/2}$ makes sense; namely

$$D^{-1/2} = \mathrm{diag}(d_1^{-1/2}, \ldots, d_m^{-1/2}).$$

Definition 3

Given any weighted directed graph $G = (V, W)$ with no isolated vertex and with $V = \{v_1, \ldots, v_m\}$, the *(normalized) graph Laplacians L_{sym} and L_{rw} of G* are defined by

$$\begin{aligned} L_{\mathrm{sym}} &= D^{-1/2} L D^{-1/2} = I - D^{-1/2} W D^{-1/2} \\ L_{\mathrm{rw}} &= D^{-1} L = I - D^{-1} W. \end{aligned}$$

Proposition 2

Let $G = (V, W)$ be a weighted graph without isolated vertices. The graph Laplacians, L, L_{sym}, and L_{rw} satisfy the following properties:

(1) *The matrix L_{sym} is symmetric, positive, semidefinite. In fact,*

$$x^\top L_{\mathrm{sym}} x = \frac{1}{2}\sum_{i,j=1}^{m} w_{i\,j}\left(\frac{x_i}{\sqrt{d_i}} - \frac{x_j}{\sqrt{d_j}}\right)^2 \quad \text{for all } x \in \mathbb{R}^m.$$

(2) *The normalized graph Laplacians L_{sym} and L_{rw} have the same spectrum $(0 = \nu_1 \leq \nu_2 \leq \ldots \leq \nu_m)$, and a vector $u \neq 0$ is an eigenvector of L_{rw} for λ iff $D^{1/2}u$ is an eigenvector of L_{sym} for λ.*

(3) *The graph Laplacians, L, L_{sym}, and L_{rw} are symmetric, positive, semidefinite.*

Proposition (continued)

(4) *A vector* $u \neq 0$ *is a solution of the generalized eigenvalue problem* $Lu = \lambda Du$ *iff* $D^{1/2}u$ *is an eigenvector of* L_{sym} *for the eigenvalue* λ *iff* u *is an eigenvector of* L_{rw} *for the eigenvalue* λ*.*

(5) *The graph Laplacians,* L *and* L_{rw} *have the same nullspace.*

(6) *The vector* $\mathbf{1}$ *is in the nullspace of* L_{rw}*, and* $D^{1/2}\mathbf{1}$ *is in the nullspace of* L_{sym}*.*

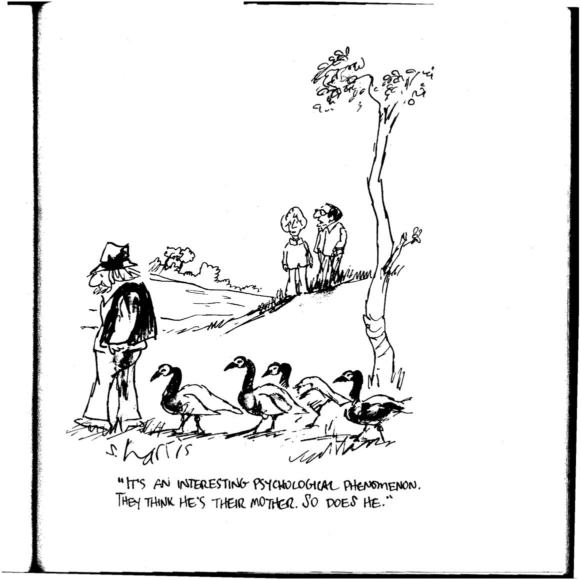

Figure 6: Are you my mother?

3. Back to Graph Clustering

If we want to partition V into K clusters, we can do so by finding a partition $(A_1, \ldots, A_K)$ that minimizes the quantity

$$\mathrm{cut}(A_1, \ldots, A_K) = \frac{1}{2} \sum_{1=1}^{K} \mathrm{cut}(A_i).$$

3. Back to Graph Clustering

If we want to partition V into K clusters, we can do so by finding a partition $(A_1, \ldots, A_K)$ that minimizes the quantity

$$\mathrm{cut}(A_1, \ldots, A_K) = \frac{1}{2} \sum_{1=1}^{K} \mathrm{cut}(A_i).$$

For $K = 2$, the mincut problem is a classical problem that can be solved efficiently, but in practice, it does not yield satisfactory partitions.

3. Back to Graph Clustering

If we want to partition V into K clusters, we can do so by finding a partition $(A_1, \ldots, A_K)$ that minimizes the quantity

$$\mathrm{cut}(A_1, \ldots, A_K) = \frac{1}{2} \sum_{1=1}^{K} \mathrm{cut}(A_i).$$

For $K = 2$, the mincut problem is a classical problem that can be solved efficiently, but in practice, it does not yield satisfactory partitions.

Indeed, in many cases, the mincut solution separates one vertex from the rest of the graph. What we need is to design our cost function in such a way that it keeps the subsets A_i "reasonably large" (reasonably balanced).

A way to get around this problem is to normalize the cuts by *dividing by some measure of each subset A_i*.

A way to get around this problem is to normalize the cuts by *dividing by some measure of each subset A_i*.

One possibility is to use the size (the number of elements) of A_i.

A way to get around this problem is to normalize the cuts by *dividing by some measure of each subset A_i*.

One possibility is to use the size (the number of elements) of A_i.

Another is to use the *volume* $\mathrm{vol}(A_i)$ of A_i. A solution using the second measure (the volume) (for $K = 2$) was proposed and investigated in a seminal paper of Shi and Malik.

A way to get around this problem is to normalize the cuts by *dividing by some measure of each subset A_i*.

One possibility is to use the size (the number of elements) of A_i.

Another is to use the *volume* $\mathrm{vol}(A_i)$ of A_i. A solution using the second measure (the volume) (for $K = 2$) was proposed and investigated in a seminal paper of Shi and Malik.

Subsequently, Stella Yu (in her dissertation) and Yu and Shi extended the method to $K > 2$ clusters.

The idea is to minimize the cost function

$$\mathrm{Ncut}(A_1, \ldots, A_K) = \sum_{i=1}^{K} \frac{\mathrm{links}(A_i, \overline{A_i})}{\mathrm{vol}(A_i)} = \sum_{i=1}^{K} \frac{\mathrm{cut}(A_i, \overline{A_i})}{\mathrm{vol}(A_i)}.$$

The idea is to minimize the cost function

$$\mathrm{Ncut}(A_1, \ldots, A_K) = \sum_{i=1}^{K} \frac{\mathrm{links}(A_i, \overline{A_i})}{\mathrm{vol}(A_i)} = \sum_{i=1}^{K} \frac{\mathrm{cut}(A_i, \overline{A_i})}{\mathrm{vol}(A_i)}.$$

We begin with the case $K = 2$, which is easier to handle.

4. 2-Way Clustering Using Normalized Cuts

Our goal is to express our optimization problem in matrix form.

4. 2-Way Clustering Using Normalized Cuts

Our goal is to express our optimization problem in matrix form.

In the case of two clusters, a single vector X can be used to describe the partition $(A_1, A_2) = (A, \overline{A})$.

4. 2-Way Clustering Using Normalized Cuts

Our goal is to express our optimization problem in matrix form.

In the case of two clusters, a single vector X can be used to describe the partition $(A_1, A_2) = (A, \overline{A})$.

It is desirable to choose the structure of this vector in such a way that

$$\mathrm{Ncut}(A, \overline{A}) = \frac{X^\top L X}{X^\top D X}$$

(the Rayleigh ratio).

4. 2-Way Clustering Using Normalized Cuts

Our goal is to express our optimization problem in matrix form.

In the case of two clusters, a single vector X can be used to describe the partition $(A_1, A_2) = (A, \overline{A})$.

It is desirable to choose the structure of this vector in such a way that

$$\mathrm{Ncut}(A, \overline{A}) = \frac{X^\top L X}{X^\top D X}$$

(the Rayleigh ratio).

It is also important to pick a vector representation which is invariant under multiplication by a nonzero scalar, because the Rayleigh ratio is *scale-invariant*, and it is crucial to take advantage of this fact to make the denominator go away.

Let $N = |V|$ be the number of nodes in the graph G. In view of the desire for a scale-invariant representation, it is natural to assume that the vector X is of the form

$$X = (x_1, \ldots, x_N),$$

where $x_i \in \{a, b\}$ for $i = 1, \ldots, N$, for any two distinct real numbers a, b.

Let $N = |V|$ be the number of nodes in the graph G. In view of the desire for a scale-invariant representation, it is natural to assume that the vector X is of the form

$$X = (x_1, \ldots, x_N),$$

where $x_i \in \{a, b\}$ for $i = 1, \ldots, N$, for any two distinct real numbers a, b.

This is an *indicator vector* in the sense that, for $i = 1, \ldots, N$,

$$x_i = \begin{cases} a & \text{if } v_i \in A \\ b & \text{if } v_i \notin A. \end{cases}$$

The correct interpretation is really to view X as a representative of a point in the *real projective space* $\mathbb{RP}^{N-1}$, namely the point $\mathbb{P}(X)$ of homogeneous coordinates $(x_1 : \cdots : x_N)$.

The correct interpretation is really to view X as a representative of a point in the *real projective space* $\mathbb{RP}^{N-1}$, namely the point $\mathbb{P}(X)$ of homogeneous coordinates $(x_1 : \cdots : x_N)$.

Therefore, from now on, we view X as a vector of homogeneous coordinates representing the point $\mathbb{P}(X) \in \mathbb{RP}^{N-1}$.

The correct interpretation is really to view X as a representative of a point in the *real projective space* $\mathbb{RP}^{N-1}$, namely the point $\mathbb{P}(X)$ of homogeneous coordinates $(x_1 : \cdots : x_N)$.

Therefore, from now on, we view X as a vector of homogeneous coordinates representing the point $\mathbb{P}(X) \in \mathbb{RP}^{N-1}$.

Let $d = \mathbf{1}^\top D\mathbf{1}$ and $\alpha = \mathrm{vol}(A)$. Then, $\mathrm{vol}(\overline{A}) = d - \alpha$, and using Proposition 1, we find that

The correct interpretation is really to view X as a representative of a point in the *real projective space* $\mathbb{RP}^{N-1}$, namely the point $\mathbb{P}(X)$ of homogeneous coordinates $(x_1 : \cdots : x_N)$.

Therefore, from now on, we view X as a vector of homogeneous coordinates representing the point $\mathbb{P}(X) \in \mathbb{RP}^{N-1}$.

Let $d = \mathbf{1}^\top D\mathbf{1}$ and $\alpha = \mathrm{vol}(A)$. Then, $\mathrm{vol}(\overline{A}) = d - \alpha$, and using Proposition 1, we find that

$$\begin{aligned} X^\top LX &= (a-b)^2 \,\mathrm{cut}(A, \overline{A}) \\ X^\top DX &= \alpha a^2 + (d-\alpha)b^2 \\ \mathrm{Ncut}(A, \overline{A}) &= \frac{d}{\alpha(d-\alpha)} \,\mathrm{cut}(A, \overline{A}). \end{aligned}$$

I proved that

$$\mathrm{Ncut}(A, \overline{A}) = \frac{X^{\top} L X}{X^{\top} D X}$$

iff we have the condition

$$a\alpha + b(d - \alpha) = 0. \qquad (\dagger)$$

I proved that

$$\mathrm{Ncut}(A, \overline{A}) = \frac{X^\top L X}{X^\top D X}$$

iff we have the condition

$$a\alpha + b(d - \alpha) = 0. \qquad (\dagger)$$

Note that condition $(\dagger)$ applied to a vector X whose components are a or b is equivalent to the fact that X is orthogonal to $D\mathbf{1}$.

Various choices for a choice of scale factor appear in the literature.

Various choices for a choice of scale factor appear in the literature.

von Luxburg picks

$$a = \sqrt{\frac{d - \alpha}{\alpha}}, \quad b = -\sqrt{\frac{\alpha}{d - \alpha}}.$$

Various choices for a choice of scale factor appear in the literature.

von Luxburg picks

$$a = \sqrt{\frac{d-\alpha}{\alpha}}, \quad b = -\sqrt{\frac{\alpha}{d-\alpha}}.$$

Shi and Malik use

$$a = 1, \quad b = -\frac{\alpha}{d-\alpha} = -\frac{k}{1-k},$$

with

$$k = \frac{\alpha}{d}.$$

Various choices for a choice of scale factor appear in the literature.

von Luxburg picks

$$a = \sqrt{\frac{d - \alpha}{\alpha}}, \quad b = -\sqrt{\frac{\alpha}{d - \alpha}}.$$

Shi and Malik use

$$a = 1, \quad b = -\frac{\alpha}{d - \alpha} = -\frac{k}{1 - k},$$

with

$$k = \frac{\alpha}{d}.$$

Belkin and Niyogi use

$$a = \frac{1}{\alpha}, \quad b = -\frac{1}{d - \alpha}.$$

However, there is no need to restrict solutions to be of either of these forms.

However, there is no need to restrict solutions to be of either of these forms.

So, let

$$\mathcal{X} = \big\{(x_1, \ldots, x_N) \mid x_i \in \{a, b\},\ a, b \in \mathbb{R} - \{0\},\ a \neq b\big\},$$

so that our solution set is

$$\mathcal{K} = \big\{X \in \mathcal{X} \mid X^\top D\mathbf{1} = 0\big\}.$$

However, there is no need to restrict solutions to be of either of these forms.

So, let

$$\mathcal{X} = \big\{(x_1, \ldots, x_N) \mid x_i \in \{a, b\},\ a, b \in \mathbb{R} - \{0\},\ a \neq b\big\},$$

so that our solution set is

$$\mathcal{K} = \big\{X \in \mathcal{X} \mid X^\top D\mathbf{1} = 0\big\}.$$

Actually, to be perfectly rigorous, we are looking for solutions in $\mathbb{RP}^{N-1}$, so our solution set is really

$$\mathbb{P}(\mathcal{K}) = \big\{(x_1 \colon \cdots \colon x_N) \in \mathbb{RP}^{N-1} \mid (x_1, \ldots, x_N) \in \mathcal{K}\big\}.$$

Consequently, our minimization problem can be stated as follows:

Consequently, our minimization problem can be stated as follows:

Problem PNC1

$$\begin{aligned} &\text{minimize} && \frac{X^\top L X}{X^\top D X} && \\ &\text{subject to} && X^\top D \mathbf{1} = 0, && X \in \mathcal{X}. \end{aligned}$$

Consequently, our minimization problem can be stated as follows:

Problem PNC1

$$\begin{aligned} &\text{minimize} && \frac{X^{\top} L X}{X^{\top} D X} && \\ &\text{subject to} && X^{\top} D \mathbf{1} = 0, && X \in \mathcal{X}. \end{aligned}$$

It is understood that the solutions are points $\mathbb{P}(X)$ in $\mathbb{RP}^{N-1}$.

Since the Rayleigh ratio and the constraints $X^\top D\mathbf{1} = 0$ and $X \in \mathcal{X}$ are scale-invariant, we are led to the following formulation of our problem:

Since the Rayleigh ratio and the constraints $X^\top D\mathbf{1} = 0$ and $X \in \mathcal{X}$ are scale-invariant, we are led to the following formulation of our problem:

Problem PNC2

$$\begin{aligned} &\text{minimize} && X^\top L X \\ &\text{subject to} && X^\top D X = 1, \qquad X^\top D\mathbf{1} = 0, \qquad X \in \mathcal{X}. \end{aligned}$$

Since the Rayleigh ratio and the constraints $X^\top D\mathbf{1} = 0$ and $X \in \mathcal{X}$ are scale-invariant, we are led to the following formulation of our problem:

Problem PNC2

$$
\begin{aligned}
&\text{minimize} && X^\top L X \\
&\text{subject to} && X^\top D X = 1, \qquad X^\top D\mathbf{1} = 0, \qquad X \in \mathcal{X}.
\end{aligned}
$$

Problem PNC2 is equivalent to problem PNC1 in the sense that if X is any minimal solution of PNC1, then $X/(X^\top D X)^{1/2}$ is a minimal solution of PNC2 (with the same minimal value for the objective functions), and if X is a minimal solution of PNC2, then λX is a minimal solution for PNC1 for all $\lambda \neq 0$ (with the same minimal value for the objective functions).

Since the Rayleigh ratio and the constraints $X^\top D\mathbf{1} = 0$ and $X \in \mathcal{X}$ are scale-invariant, we are led to the following formulation of our problem:

Problem PNC2

$$\begin{aligned} &\text{minimize} && X^\top L X \\ &\text{subject to} && X^\top D X = 1, \qquad X^\top D\mathbf{1} = 0, \qquad X \in \mathcal{X}. \end{aligned}$$

Problem PNC2 is equivalent to problem PNC1 in the sense that if X is any minimal solution of PNC1, then $X/(X^\top D X)^{1/2}$ is a minimal solution of PNC2 (with the same minimal value for the objective functions), and if X is a minimal solution of PNC2, then λX is a minimal solution for PNC1 for all $\lambda \neq 0$ (with the same minimal value for the objective functions).

Equivalently, problems PNC1 and PNC2 have the same set of minimal solutions as points $\mathbb{P}(X) \in \mathbb{RP}^{N-1}$ given by their homogeneous coordinates X.

Unfortunately, this is an NP-complete problem, as shown by Shi and Malik.

Unfortunately, this is an NP-complete problem, as shown by Shi and Malik.

As often with hard combinatorial problems, we can look for a *relaxation* of our problem, which means looking for an optimum in a larger continuous domain.

Unfortunately, this is an NP-complete problem, as shown by Shi and Malik.

As often with hard combinatorial problems, we can look for a *relaxation* of our problem, which means looking for an optimum in a larger continuous domain.

After doing this, the problem is to find a discrete solution which is close to a continuous optimum of the relaxed problem.

The natural relaxation of this problem is to allow X to be any nonzero vector in $\mathbb{R}^N$, and we get the problem:

$$\text{minimize} \quad X^\top L X \quad \text{subject to} \quad X^\top D X = 1, \quad X^\top D \mathbf{1} = 0.$$

As usual, let $Y = D^{1/2}X$, so that $X = D^{-1/2}Y$.

As usual, let $Y = D^{1/2}X$, so that $X = D^{-1/2}Y$.

We obtain the problem:

$$\text{minimize} \quad Y^\top D^{-1/2} L D^{-1/2} Y \quad \text{subject to} \quad Y^\top Y = 1, \quad Y^\top D^{1/2}\mathbf{1} = 0.$$

As usual, let $Y = D^{1/2}X$, so that $X = D^{-1/2}Y$.

We obtain the problem:

$$\text{minimize} \quad Y^\top D^{-1/2} L D^{-1/2} Y \quad \text{subject to} \quad Y^\top Y = 1, \quad Y^\top D^{1/2}\mathbf{1} = 0.$$

Because $L\mathbf{1} = 0$, the vector $D^{1/2}\mathbf{1}$ belongs to the nullspace of the symmetric Laplacian $L_{\mathrm{sym}} = D^{-1/2}LD^{-1/2}$.

As usual, let $Y = D^{1/2}X$, so that $X = D^{-1/2}Y$.

We obtain the problem:

$$\text{minimize} \quad Y^\top D^{-1/2} L D^{-1/2} Y \quad \text{subject to} \quad Y^\top Y = 1, \quad Y^\top D^{1/2}\mathbf{1} = 0.$$

Because $L\mathbf{1} = 0$, the vector $D^{1/2}\mathbf{1}$ belongs to the nullspace of the symmetric Laplacian $L_{\mathrm{sym}} = D^{-1/2} L D^{-1/2}$.

By the Rayleigh–Ritz theorem, minima are achieved by any unit eigenvector Y of the second eigenvalue ν_2 of L_{sym}.

Then, $Z = D^{-1/2}Y$ is a solution of our original relaxed problem.

The next question is to figure how close is Z to an exact solution in $\mathcal{X}$.

The next question is to figure how close is Z to an exact solution in $\mathcal{X}$.

Actually, because solutions are points in $\mathbb{RP}^{N-1}$, the correct statement of the question is:

The next question is to figure how close is Z to an exact solution in $\mathcal{X}$.

Actually, because solutions are points in $\mathbb{RP}^{N-1}$, the correct statement of the question is:

Find an exact solution $\mathbb{P}(X) \in \mathbb{P}(\mathcal{X})$ which is the closest (in a suitable sense) to the approximate solution $\mathbb{P}(Z) \in \mathbb{RP}^{N-1}$.

The next question is to figure how close is Z to an exact solution in $\mathcal{X}$.

Actually, because solutions are points in $\mathbb{RP}^{N-1}$, the correct statement of the question is:

Find an exact solution $\mathbb{P}(X) \in \mathbb{P}(\mathcal{X})$ which is the closest (in a suitable sense) to the approximate solution $\mathbb{P}(Z) \in \mathbb{RP}^{N-1}$.

However, because $\mathcal{X}$ is closed under the antipodal map, it can be shown that minimizing the distance $d(\mathbb{P}(X), \mathbb{P}(Z))$ on $\mathbb{RP}^{N-1}$ is equivalent to minimizing the Euclidean distance $\|X - Z\|_2$ (if we use the Riemannian metric on $\mathbb{RP}^{N-1}$ induced by the Euclidean metric on $\mathbb{R}^N$).

We may assume $b < 0$, in which case $a > 0$.

We may assume $b < 0$, in which case $a > 0$.

If all entries in Z are nonzero, due to the projective nature of the solution set, it seems reasonable to say that the partition of V is defined by the signs of the entries in Z.

We may assume $b < 0$, in which case $a > 0$.

If all entries in Z are nonzero, due to the projective nature of the solution set, it seems reasonable to say that the partition of V is defined by the signs of the entries in Z.

Thus, A will consist of nodes those v_i for which $x_i > 0$. Elements corresponding to zero entries can be assigned to either A or $\overline{A}$, unless additional information is available.

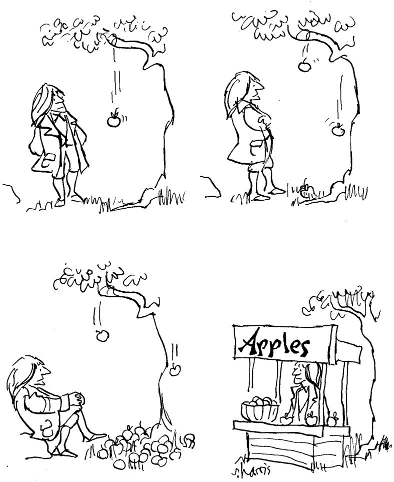

Figure 7: Newton goes to Wharton

5. K-Way Clustering Using Normalized Cuts

We describe a partition $(A_1, \ldots, A_K)$ of the set of nodes V by an $N \times K$ matrix $X = [X^1 \cdots X^K]$ whose columns $X^1, \ldots, X^K$ are indicator vectors of the partition $(A_1, \ldots, A_K)$.

5. K-Way Clustering Using Normalized Cuts

We describe a partition $(A_1, \ldots, A_K)$ of the set of nodes V by an $N \times K$ matrix $X = [X^1 \cdots X^K]$ whose columns $X^1, \ldots, X^K$ are indicator vectors of the partition $(A_1, \ldots, A_K)$.

Inspired by what we did in Section 4, we assume that the vector X^j is of the form

$$X^j = (x_1^j, \ldots, x_N^j),$$

where $x_i^j \in \{a_j, b_j\}$ for $j = 1, \ldots, K$ and $i = 1, \ldots, N$, and where a_j, b_j are any two distinct real numbers.

5. K-Way Clustering Using Normalized Cuts

We describe a partition $(A_1, \ldots, A_K)$ of the set of nodes V by an $N \times K$ matrix $X = [X^1 \cdots X^K]$ whose columns $X^1, \ldots, X^K$ are indicator vectors of the partition $(A_1, \ldots, A_K)$.

Inspired by what we did in Section 4, we assume that the vector X^j is of the form

$$X^j = (x_1^j, \ldots, x_N^j),$$

where $x_i^j \in \{a_j, b_j\}$ for $j = 1, \ldots, K$ and $i = 1, \ldots, N$, and where a_j, b_j are any two distinct real numbers.

The vector X^j is an indicator vector for A_j in the sense that, for $i = 1, \ldots, N$,

$$x_i^j = \begin{cases} a_j & \text{if } v_i \in A_j \\ b_j & \text{if } v_i \notin A_j. \end{cases}$$

Let $d = \mathbf{1}^\top D\mathbf{1}$ and $\alpha_j = \mathrm{vol}(A_j)$, so that $\alpha_1 + \cdots + \alpha_K = d$.

Let $d = \mathbf{1}^\top D\mathbf{1}$ and $\alpha_j = \mathrm{vol}(A_j)$, so that $\alpha_1 + \cdots + \alpha_K = d$.

Then, $\mathrm{vol}(\overline{A_j}) = d - \alpha_j$, and as in Section 4, we have

$$\begin{aligned}
(X^j)^\top L X^j &= (a_j - b_j)^2 \,\mathrm{cut}(A_j, \overline{A_j}), \\
(X^j)^\top D X^j &= \alpha_j a_j^2 + (d - \alpha_j) b_j^2.
\end{aligned}$$

Let $d = \mathbf{1}^\top D\mathbf{1}$ and $\alpha_j = \mathrm{vol}(A_j)$, so that $\alpha_1 + \cdots + \alpha_K = d$.

Then, $\mathrm{vol}(\overline{A_j}) = d - \alpha_j$, and as in Section 4, we have

$$\begin{aligned}(X^j)^\top LX^j &= (a_j - b_j)^2 \,\mathrm{cut}(A_j, \overline{A_j}),\\ (X^j)^\top DX^j &= \alpha_j a_j^2 + (d - \alpha_j) b_j^2.\end{aligned}$$

Since

$$\mathrm{Ncut}(A_1, \ldots, A_K) = \sum_{j=1}^{K} \frac{\mathrm{cut}(A_j, \overline{A_j})}{\mathrm{vol}(A_j)},$$

we would like to choose a_j, b_j so that

$$\frac{\mathrm{cut}(A_j, \overline{A_j})}{\mathrm{vol}(A_j)} = \frac{(X^j)^\top LX^j}{(X^j)^\top DX^j} \quad j = 1, \ldots, K.$$

We find that there are two possibilities:

1. $b_j = 0$.
2. $b_j \neq 0$, which yields
$$a_j = \frac{2\alpha_j - d}{2\alpha_j} b_j.$$

We find that there are two possibilities:

1. $b_j = 0$.
2. $b_j \neq 0$, which yields
$$a_j = \frac{2\alpha_j - d}{2\alpha_j} b_j.$$

The second choice is more complicated, so we opt for the choice $b_j = 0$, $j = 1, \ldots, K$.

We find that there are two possibilities:

1. $b_j = 0$.
2. $b_j \neq 0$, which yields
$$a_j = \frac{2\alpha_j - d}{2\alpha_j} b_j.$$

The second choice is more complicated, so we opt for the choice $b_j = 0$, $j = 1, \ldots, K$.

von Luxburg and Yu and Shi pick $b_j = 0$ and

$$a_j = \frac{1}{\sqrt{\alpha_j}} = \frac{1}{\sqrt{\mathrm{vol}(A_j)}}, \quad j = 1, \ldots, K.$$

When $N = 10$ and $K = 4$, an example of a matrix X representing the partition of $V = \{v_1, v_2, \ldots, v_{10}\}$ into the four blocks

$$\{A_1, A_2, A_3, A_4\} = \{\{v_2, v_4, v_6\}, \{v_1, v_5\}, \{v_3, v_8, v_{10}\}, \{v_7, v_9\}\},$$

is shown below:

When $N = 10$ and $K = 4$, an example of a matrix X representing the partition of $V = \{v_1, v_2, \ldots, v_{10}\}$ into the four blocks

$$\{A_1, A_2, A_3, A_4\} = \{\{v_2, v_4, v_6\}, \{v_1, v_5\}, \{v_3, v_8, v_{10}\}, \{v_7, v_9\}\},$$

is shown below:

$$X = \begin{pmatrix} 0 & a_2 & 0 & 0 \\ a_1 & 0 & 0 & 0 \\ 0 & 0 & a_3 & 0 \\ a_1 & 0 & 0 & 0 \\ 0 & a_2 & 0 & 0 \\ a_1 & 0 & 0 & 0 \\ 0 & 0 & 0 & a_4 \\ 0 & 0 & a_3 & 0 \\ 0 & 0 & 0 & a_4 \\ 0 & 0 & a_3 & 0 \end{pmatrix}.$$

We now consider the problem of finding necessary and sufficient conditions for a matrix X to represent a partition of V.

We now consider the problem of finding necessary and sufficient conditions for a matrix X to represent a partition of V.

When $b_j = 0$, the pairwise disjointness of the A_i is captured by the orthogonality of the X^i:

$$(X^i)^\top X^j = 0, \quad 1 \leq i, j \leq K, \; i \neq j. \tag{$*$}$$

We now consider the problem of finding necessary and sufficient conditions for a matrix X to represent a partition of V.

When $b_j = 0$, the pairwise disjointness of the A_i is captured by the orthogonality of the X^i:

$$(X^i)^\top X^j = 0, \quad 1 \leq i, j \leq K, \; i \neq j. \tag{$*$}$$

Now, because D is a diagonal matrix with positive entries and because the nonzero entries in each column of X have the same sign, for any $i \neq j$, the condition

$$(X^i)^\top X^j = 0$$

is equivalent to

$$(X^i)^\top D X^j = 0. \tag{$**$}$$

These conditions turn out to be more convenient.

Each A_j is nonempty iff $X^j \neq 0$, and the fact that the union of the A_j is V is captured by the fact that each row of X must have some nonzero entry (every vertex appears in some block).

Each A_j is nonempty iff $X^j \neq 0$, and the fact that the union of the A_j is V is captured by the fact that each row of X must have some nonzero entry (every vertex appears in some block).

It is not obvious how to state conveniently this condition in matrix form.

Each A_j is nonempty iff $X^j \neq 0$, and the fact that the union of the A_j is V is captured by the fact that each row of X must have some nonzero entry (every vertex appears in some block).

It is not obvious how to state conveniently this condition in matrix form.

Unltimately, we found that the following scale-invariant equation works:

$$X(X^\top X)^{-1} X^\top \mathbf{1} = \mathbf{1}. \qquad (\dagger)$$

Each A_j is nonempty iff $X^j \neq 0$, and the fact that the union of the A_j is V is captured by the fact that each row of X must have some nonzero entry (every vertex appears in some block).

It is not obvious how to state conveniently this condition in matrix form.

Unltimately, we found that the following scale-invariant equation works:

$$X(X^\top X)^{-1}X^\top \mathbf{1} = \mathbf{1}. \tag{$\dagger$}$$

Note that because the columns of X are linearly independent, $(X^\top X)^{-1}X^\top$ is the pseudo-inverse of X. Consequently, condition ($\dagger$), can also be written as

$$XX^+\mathbf{1} = \mathbf{1}.$$

If we let

$$\mathcal{X} = \left\{ [X^1 \ \ldots \ X^K] \mid X^j = a_j(x_1^j, \ldots, x_N^j), \ x_i^j \in \{1, 0\}, a_j \in \mathbb{R}, \ X^j \neq 0 \right\}$$

then the set of matrices representing partitions of V into K blocks is

If we let

$$\mathcal{X} = \Big\{ [X^1 \ \ldots \ X^K] \mid X^j = a_j(x_1^j, \ldots, x_N^j), \ x_i^j \in \{1, 0\}, a_j \in \mathbb{R}, \ X^j \neq 0 \Big\}$$

then the set of matrices representing partitions of V into K blocks is

$$\begin{aligned} \mathcal{K} = \Big\{ X = [X^1 \ \cdots \ X^K] \quad & \mid X \in \mathcal{X}, \\ & (X^i)^\top D X^j = 0, \quad 1 \leq i, j \leq K, \ i \neq j, \\ & X(X^\top X)^{-1} X^\top \mathbf{1} = \mathbf{1} \Big\}. \end{aligned}$$

If we let

$$\mathcal{X} = \left\{ [X^1 \ \ldots \ X^K] \mid X^j = a_j(x_1^j, \ldots, x_N^j),\ x_i^j \in \{1, 0\}, a_j \in \mathbb{R},\ X^j \neq 0 \right\}$$

then the set of matrices representing partitions of V into K blocks is

$$\begin{aligned} \mathcal{K} = \Big\{ X = [X^1 \ \cdots \ X^K] \quad & \mid X \in \mathcal{X}, \\ & (X^i)^\top D X^j = 0, \quad 1 \leq i, j \leq K,\ i \neq j, \\ & X(X^\top X)^{-1} X^\top \mathbf{1} = \mathbf{1} \Big\}. \end{aligned}$$

As in the case $K = 2$, to be rigorous, the solution are really K-tuples of points in $\mathbb{RP}^{N-1}$, so our solution set is really

$$\mathbb{P}(\mathcal{K}) = \left\{ (\mathbb{P}(X^1), \ldots, \mathbb{P}(X^K)) \mid [X^1 \ \cdots \ X^K] \in \mathcal{K} \right\}.$$

K-way Clustering of a graph using Normalized Cut, Version 1: Problem PNC1

$$
\begin{aligned}
&\text{minimize} && \sum_{j=1}^{K} \frac{(X^j)^\top L X^j}{(X^j)^\top D X^j} \\
&\text{subject to} && (X^i)^\top D X^j = 0, \quad 1 \leq i, j \leq K, \ i \neq j, \\
& && X(X^\top X)^{-1} X^\top \mathbf{1} = \mathbf{1}, \qquad X \in \mathcal{X}.
\end{aligned}
$$

K-way Clustering of a graph using Normalized Cut, Version 1: Problem PNC1

$$
\begin{aligned}
&\text{minimize} && \sum_{j=1}^{K} \frac{(X^j)^\top L X^j}{(X^j)^\top D X^j} \\
&\text{subject to} && (X^i)^\top D X^j = 0, \quad 1 \leq i, j \leq K,\ i \neq j, \\
& && X(X^\top X)^{-1} X^\top \mathbf{1} = \mathbf{1}, \qquad X \in \mathcal{X}.
\end{aligned}
$$

As in the case $K = 2$, the solutions that we are seeking are K-tuples $(\mathbb{P}(X^1), \ldots, \mathbb{P}(X^K))$ of points in $\mathbb{RP}^{N-1}$ determined by their homogeneous coordinates $X^1, \ldots, X^K$.

Our original formulation (PNC1) can be converted to a more convenient form, by chasing the denominators in the Rayleigh ratios, and by expressing the objective function in terms of the *trace* of a certain matrix.

Our original formulation (PNC1) can be converted to a more convenient form, by chasing the denominators in the Rayleigh ratios, and by expressing the objective function in terms of the *trace* of a certain matrix.

Let

$$\mu(X^1, \ldots, X^K) = \sum_{j=1}^{K} \frac{(X^j)^\top L X^j}{(X^j)^\top D X^j}.$$

Proposition 3

For any orthogonal $K \times K$ matrix R, any symmetric $N \times N$ matrix A, and any $N \times K$ matrix $X = [X^1 \ \cdots \ X^K]$, the following properties hold:

(1) $\mu(X) = \mathrm{tr}(\Lambda^{-1/2} X^\top L X \Lambda^{-1/2})$, *where*

$$\Lambda = \mathrm{diag}((X^1)^\top D X^1, \ldots, (X^K)^\top D X^K).$$

(2) *If* $(X^1)^\top D X^1 = \cdots = (X^K)^\top D X^K = \alpha^2$, *then* $\mu(X) = \mu(XR)$.

(3) *The condition* $X^\top A X = \alpha^2 I$ *is preserved if* X *is replaced by* XR.

(4) *The condition* $X(X^\top X)^{-1} X^\top \mathbf{1} = \mathbf{1}$ *is preserved if* X *is replaced by* XR.

K-way Clustering of a graph using Normalized Cut, Version 2: Problem PNC2

$$
\begin{aligned}
&\text{minimize} && \mathrm{tr}(X^\top L X) \\
&\text{subject to} && X^\top D X = I, \\
& && X(X^\top X)^{-1} X^\top \mathbf{1} = \mathbf{1}, \qquad X \in \mathcal{X}.
\end{aligned}
$$

K-way Clustering of a graph using Normalized Cut, Version 2: Problem PNC2

$$\begin{aligned} &\text{minimize} && \mathrm{tr}(X^\top L X) \\ &\text{subject to} && X^\top D X = I, \\ & && X(X^\top X)^{-1} X^\top \mathbf{1} = \mathbf{1}, \qquad X \in \mathcal{X}. \end{aligned}$$

Problem PNC2 is equivalent to problem PNC1 in the sense that for every minimal solution $(X^1, \ldots, X^K)$ of PNC1, $(((X^1)^\top D X^1)^{-1/2} X^1, \ldots, ((X^K)^\top D X^K)^{-1/2} X^K)$ is a minimal solution of PNC2 (with the same minimum for the objective functions), and that for every minimal solution $(Z^1, \ldots, Z^k)$ of PNC2, $(\lambda_1 Z^1, \ldots, \lambda_K Z^K)$ is a minimal solution of PNC1, for all $\lambda_i \neq 0$, $i = 1, \ldots, K$ (with the same minimum for the objective functions).

The main problem in finding a good relaxation of problem PNC2 is that it is very difficult to enforce the condition $X \in \mathcal{X}$.

The main problem in finding a good relaxation of problem PNC2 is that it is very difficult to enforce the condition $X \in \mathcal{X}$.

Also, the solutions X are not preserved under arbitrary rotations, but only by very special rotations which leave $\mathcal{X}$ invariant (they exchange the axes).

The main problem in finding a good relaxation of problem PNC2 is that it is very difficult to enforce the condition $X \in \mathcal{X}$.

Also, the solutions X are not preserved under arbitrary rotations, but only by very special rotations which leave $\mathcal{X}$ invariant (they exchange the axes).

The first natural relaxation of problem PNC2 is to drop the condition that $X \in \mathcal{X}$, and we obtain

Problem $(*_1)$

$$
\begin{aligned}
&\text{minimize} && \mathrm{tr}(X^\top L X)\\
&\text{subject to} && X^\top D X = I,\\
& && X(X^\top X)^{-1}X^\top \mathbf{1} = \mathbf{1}.
\end{aligned}
$$

Problem $(*_1)$

$$\begin{aligned} &\text{minimize} && \mathrm{tr}(X^\top L X) \\ &\text{subject to} && X^\top D X = I, \\ & && X(X^\top X)^{-1} X^\top \mathbf{1} = \mathbf{1}. \end{aligned}$$

By Proposition 3, for every orthogonal matrix $R \in \mathbf{O}(K)$ and for every X minimizing $(*_1)$, the matrix XR also minimizes $(*_1)$.

As a consequence, we can view the solutions of problem $(*_1)$ as elements of the *Grassmannian* $G(N, K)$, as explained next.

Recall that the *Stiefel manifold* $St(k, n)$ consists of the set of orthogonal k-frames in $\mathbb{R}^n$, that is, the k-tuples of orthonormal vectors $(u_1, \ldots, u_k)$ with $u_i \in \mathbb{R}^n$ ($St(n, n) = \mathbf{O}(n)$).

As a consequence, we can view the solutions of problem $(*_1)$ as elements of the *Grassmannian* $G(N, K)$, as explained next.

Recall that the *Stiefel manifold* $St(k, n)$ consists of the set of orthogonal k-frames in $\mathbb{R}^n$, that is, the k-tuples of orthonormal vectors $(u_1, \ldots, u_k)$ with $u_i \in \mathbb{R}^n$ ($St(n, n) = \mathbf{O}(n)$).

For $1 \leq n \leq n - 1$, the group $\mathbf{SO}(n)$ acts transitively on $St(k, n)$, and $St(k, n)$ is isomorphic to the coset manifold $\mathbf{SO}(n)/\mathbf{SO}(n - k)$.

As a consequence, we can view the solutions of problem $(*_1)$ as elements of the *Grassmannian* $G(N, K)$, as explained next.

Recall that the *Stiefel manifold* $St(k, n)$ consists of the set of orthogonal k-frames in $\mathbb{R}^n$, that is, the k-tuples of orthonormal vectors $(u_1, \ldots, u_k)$ with $u_i \in \mathbb{R}^n$ ($St(n, n) = \mathbf{O}(n)$).

For $1 \leq n \leq n - 1$, the group $\mathbf{SO}(n)$ acts transitively on $St(k, n)$, and $St(k, n)$ is isomorphic to the coset manifold $\mathbf{SO}(n)/\mathbf{SO}(n - k)$.

The *Grassmann manifold* $G(k, n)$ consists of all (linear) k-dimensional subspaces of $\mathbb{R}^n$.

As a consequence, we can view the solutions of problem $(*_1)$ as elements of the *Grassmannian* $G(N, K)$, as explained next.

Recall that the *Stiefel manifold* $St(k, n)$ consists of the set of orthogonal k-frames in $\mathbb{R}^n$, that is, the k-tuples of orthonormal vectors $(u_1, \ldots, u_k)$ with $u_i \in \mathbb{R}^n$ ($St(n, n) = \mathbf{O}(n)$).

For $1 \leq n \leq n - 1$, the group $\mathbf{SO}(n)$ acts transitively on $St(k, n)$, and $St(k, n)$ is isomorphic to the coset manifold $\mathbf{SO}(n)/\mathbf{SO}(n - k)$.

The *Grassmann manifold* $G(k, n)$ consists of all (linear) k-dimensional subspaces of $\mathbb{R}^n$.

Again, the group $\mathbf{SO}(n)$ acts transitively on $G(k, n)$, and $G(k, n)$ is isomorphic to the coset manifold $\mathbf{SO}(n)/S(\mathbf{SO}(k) \times \mathbf{SO}(n - k))$.

The group $\mathbf{O}(k)$ acts on the right on the Stiefel manifold $St(k, n)$ (by multiplication), and the orbit manifold $St(k, n)/\mathbf{O}(k)$ is isomorphic to the Grassmann manifold $G(k, n)$.

The group $\mathbf{O}(k)$ acts on the right on the Stiefel manifold $St(k,n)$ (by multiplication), and the orbit manifold $St(k,n)/\mathbf{O}(k)$ is isomorphic to the Grassmann manifold $G(k,n)$.

Furthermore, both $St(k,n)$ and $G(k,n)$ are *naturally reductive homogeneous manifolds* (for the Stiefel manifold, when $n \geq 3$), and $G(k,n)$ is even a *symmetric space*.

The group $\mathbf{O}(k)$ acts on the right on the Stiefel manifold $St(k, n)$ (by multiplication), and the orbit manifold $St(k, n)/\mathbf{O}(k)$ is isomorphic to the Grassmann manifold $G(k, n)$.

Furthermore, both $St(k, n)$ and $G(k, n)$ are *naturally reductive homogeneous manifolds* (for the Stiefel manifold, when $n \geq 3$), and $G(k, n)$ is even a *symmetric space*.

The upshot of all this is that to a large extent, *the differential geometry of these manifolds is completely determined by some subspace* $\mathfrak{m}$ *of the Lie algebra* $\mathfrak{so}(n)$, such that we have a direct sum

$$\mathfrak{so}(n) = \mathfrak{m} \oplus \mathfrak{h},$$

where $\mathfrak{h} = \mathfrak{so}(n - k)$ in the case of the Stiefel manifold, and $\mathfrak{h} = \mathfrak{so}(k) \times \mathfrak{so}(n - k)$ in the case of the Grassmannian manifold (some additional condition on $\mathfrak{m}$ is required).

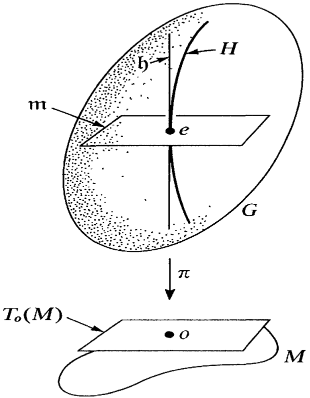

Figure 8: Reductive homogeneous space, from O'Neill

(In the above Figure, $G = \mathbf{SO}(n)$, $M = \mathbf{SO}(n)/H$).

In particular, the geodesics in both manifolds can be determined quite explicitly, and thus we obtain closed form formulae for distances, *etc*.

In particular, the geodesics in both manifolds can be determined quite explicitly, and thus we obtain closed form formulae for distances, *etc.*

The Stiefel manifold $St(k, n)$ can be viewed as the set of all $n \times k$ matrices X such that

$$X^\top X = I_k.$$

In our situation, we are considering $N \times K$ matrices X such that

$$X^\top D X = I.$$

In particular, the geodesics in both manifolds can be determined quite explicitly, and thus we obtain closed form formulae for distances, *etc.*

The Stiefel manifold $St(k, n)$ can be viewed as the set of all $n \times k$ matrices X such that

$$X^\top X = I_k.$$

In our situation, we are considering $N \times K$ matrices X such that

$$X^\top DX = I.$$

This is not quite the Stiefel manifold, but if we write $Y = D^{1/2}X$, then we have

$$Y^\top Y = I,$$

so the space of matrices X satisfying the condition $X^\top DX = I$ is the image $\mathcal{D}(St(K, N))$ of the Stiefel manifold $St(K, N)$ under the linear map $\mathcal{D}$ given by

$$\mathcal{D}(X) = D^{1/2}X.$$

Now, the right action of $\mathbf{O}(K)$ on $\mathcal{D}(St(K,N))$ yields a coset manifold $\mathcal{D}(St(K,N))/\mathbf{O}(K)$ which is obviously isomorphic to the Grassmann manidold $G(K,N)$.

Now, the right action of $\mathbf{O}(K)$ on $\mathcal{D}(St(K,N))$ yields a coset manifold $\mathcal{D}(St(K,N))/\mathbf{O}(K)$ which is obviously isomorphic to the Grassmann manidold $G(K,N)$.

Therefore, *the solutions of problem* $(*_1)$ *can be viewed as elements of the Grassmannian* $G(N,K)$.

Now, the right action of $\mathbf{O}(K)$ on $\mathcal{D}(St(K,N))$ yields a coset manifold $\mathcal{D}(St(K,N))/\mathbf{O}(K)$ which is obviously isomorphic to the Grassmann manidold $G(K,N)$.

Therefore, *the solutions of problem* $(*_1)$ *can be viewed as elements of the Grassmannian* $G(N,K)$.

We can take advantage of this fact to find a discrete solution of our original optimization problem PNC2 approximated by a continuous solution of $(*_1)$.

Using the linear map $\mathcal{D}$, we get the equivalent problem

Using the linear map $\mathcal{D}$, we get the equivalent problem

Problem $(**_1)$

$$
\begin{aligned}
&\text{minimize} && \mathrm{tr}(Y^\top D^{-1/2} L D^{-1/2} Y) \\
&\text{subject to} && Y^\top Y = I, \\
& && Y Y^\top D^{1/2}\mathbf{1} = D^{1/2}\mathbf{1}.
\end{aligned}
$$

Using the linear map $\mathcal{D}$, we get the equivalent problem

Problem $(**_1)$

$$
\begin{aligned}
&\text{minimize} && \mathrm{tr}(Y^\top D^{-1/2} L D^{-1/2} Y)\\
&\text{subject to} && Y^\top Y = I,\\
& && YY^\top D^{1/2}\mathbf{1} = D^{1/2}\mathbf{1}.
\end{aligned}
$$

This time, the matrices Y satisfying condition $Y^\top Y = I$ do belong to the Stiefel manifold $St(K, N)$, and again, *we view the solutions of problem* $(**_1)$ *as elements of the Grassmannian* $G(K, N)$.

We pass from a solution Y of problem $(**_1)$ in $G(K,N)$ to a solution Z of of problem $(*_1)$ in $G(K,N)$ by the linear map $\mathcal{D}^{-1}$; namely, $Z = \mathcal{D}(Y) = D^{-1/2}Y$.

We pass from a solution Y of problem $(**_1)$ in $G(K, N)$ to a solution Z of of problem $(*_1)$ in $G(K, N)$ by the linear map $\mathcal{D}^{-1}$; namely, $Z = \mathcal{D}(Y) = D^{-1/2}Y$.

The Rayleigh–Ritz Theorem tells us that if we temporarily ignore the second constraint, minima of problem $(**_1)$ are obtained by picking any K unit eigenvectors $(u_1, \ldots, u_k)$ associated with the smallest eigenvalues

$$0 = \nu_1 \leq \nu_2 \leq \ldots \leq \nu_K$$

of $L_{\mathrm{sym}} = D^{-1/2}LD^{-1/2}$.

We may assume that $\nu_2 > 0$, namely that the underlying graph is connected (otherwise, we work with each connected component), in which case $Y^1 = D^{1/2}\mathbf{1} / \left\| D^{1/2}\mathbf{1} \right\|_2$, because $\mathbf{1}$ is in the nullspace of L.

We may assume that $\nu_2 > 0$, namely that the underlying graph is connected (otherwise, we work with each connected component), in which case $Y^1 = D^{1/2}\mathbf{1}/\left\|D^{1/2}\mathbf{1}\right\|_2$, because $\mathbf{1}$ is in the nullspace of L.

Since $Y^1 = D^{1/2}\mathbf{1}/\left\|D^{1/2}\mathbf{1}\right\|_2$, the vector $D^{1/2}\mathbf{1}$ is in the range of Y, so the condition

$$YY^\top D^{1/2}\mathbf{1} = D^{1/2}\mathbf{1}$$

is also satisfied.

We may assume that $\nu_2 > 0$, namely that the underlying graph is connected (otherwise, we work with each connected component), in which case $Y^1 = D^{1/2}\mathbf{1}/\left\|D^{1/2}\mathbf{1}\right\|_2$, because $\mathbf{1}$ is in the nullspace of L.

Since $Y^1 = D^{1/2}\mathbf{1}/\left\|D^{1/2}\mathbf{1}\right\|_2$, the vector $D^{1/2}\mathbf{1}$ is in the range of Y, so the condition

$$YY^\top D^{1/2}\mathbf{1} = D^{1/2}\mathbf{1}$$

is also satisfied.

Then, $Z = D^{-1/2}Y$ with $Y = [u_1 \; \ldots \; u_K]$ yields a minimum of our relaxed problem $(*_1)$ (the second constraint is satisfied because $\mathbf{1}$ is in the range of Z).

Figure 9: Try and try again

The next step is to find an exact solution $(\mathbb{P}(X^1), \ldots, \mathbb{P}(X^K)) \in \mathbb{P}(\mathcal{K})$ which is the closest (in a suitable sense) to our approximate solution $(Z^1, \ldots, Z^K) \in G(K, N)$.

The next step is to find an exact solution $(\mathbb{P}(X^1), \ldots, \mathbb{P}(X^K)) \in \mathbb{P}(\mathcal{K})$ which is the closest (in a suitable sense) to our approximate solution $(Z^1, \ldots, Z^K) \in G(K, N)$.

The set $\mathcal{K}$ is not necessarily closed under all orthogonal transformations in $\mathbf{O}(K)$, so we can't view $\mathcal{K}$ as a subset of the Grassmannian $G(K, N)$.

The next step is to find an exact solution $(\mathbb{P}(X^1), \ldots, \mathbb{P}(X^K)) \in \mathbb{P}(\mathcal{K})$ which is the closest (in a suitable sense) to our approximate solution $(Z^1, \ldots, Z^K) \in G(K, N)$.

The set $\mathcal{K}$ is not necessarily closed under all orthogonal transformations in $\mathbf{O}(K)$, so we can't view $\mathcal{K}$ as a subset of the Grassmannian $G(K, N)$.

However, we can think of $\mathcal{K}$ as a subset of $G(K, N)$ by considering the subspace spanned by $(X^1, \ldots, X^K)$ for every $[X^1 \cdots X^K] \in \mathcal{K}$.

The next step is to find an exact solution $(\mathbb{P}(X^1), \ldots, \mathbb{P}(X^K)) \in \mathbb{P}(\mathcal{K})$ which is the closest (in a suitable sense) to our approximate solution $(Z^1, \ldots, Z^K) \in G(K, N)$.

The set $\mathcal{K}$ is not necessarily closed under all orthogonal transformations in $\mathbf{O}(K)$, so we can't view $\mathcal{K}$ as a subset of the Grassmannian $G(K, N)$.

However, we can think of $\mathcal{K}$ as a subset of $G(K, N)$ by considering the subspace spanned by $(X^1, \ldots, X^K)$ for every $[X^1 \cdots X^K] \in \mathcal{K}$.

Then, we have two choices of distances.

(1) We view $\mathcal{K}$ as a subset of $(\mathbb{RP}^{N-1})^K$. Because $\mathcal{K}$ is closed under the antipodal map, as minimizing the distance $d(\mathbb{P}(X^j), \mathbb{P}(Z^j))$ on $\mathbb{RP}^{N-1}$ is equivalent to minimizing the Euclidean distance $\left\|X^j - Z^j\right\|_2$, for $j = 1, \ldots, K$ (if we use the Riemannian metric on $\mathbb{RP}^{N-1}$ induced by the Euclidean metric on $\mathbb{R}^N$).

(1) We view $\mathcal{K}$ as a subset of $(\mathbb{RP}^{N-1})^K$. Because $\mathcal{K}$ is closed under the antipodal map, as minimizing the distance $d(\mathbb{P}(X^j), \mathbb{P}(Z^j))$ on $\mathbb{RP}^{N-1}$ is equivalent to minimizing the Euclidean distance $\left\|X^j - Z^j\right\|_2$, for $j = 1, \ldots, K$ (if we use the Riemannian metric on $\mathbb{RP}^{N-1}$ induced by the Euclidean metric on $\mathbb{R}^N$).

Then, minimizing the distance $d(X, Z)$ in $(\mathbb{RP}^{N-1})^K$ is equivalent to minimizing $\|X - Z\|_F$, where

$$\|X - Z\|_F^2 = \sum_{j=1}^{K} \left\|X^j - Z^j\right\|_2^2$$

is the Frobenius norm. This is implicitly the choice made by Yu.

(2) We view $\mathcal{K}$ as a subset of the Grassmannian $G(K, N)$. In this case, we need to pick a metric on the Grassmannian, and we minimize the corresponding Riemannian distance $d(X, Z)$. A natural choice is the metric on $\mathfrak{so}(n)$ given by

$$\langle X, Y \rangle = \mathrm{tr}(X^\top Y).$$

(2) We view $\mathcal{K}$ as a subset of the Grassmannian $G(K, N)$. In this case, we need to pick a metric on the Grassmannian, and we minimize the corresponding Riemannian distance $d(X, Z)$. A natural choice is the metric on $\mathfrak{so}(n)$ given by

$$\langle X, Y\rangle = \mathrm{tr}(X^\top Y).$$

This choice remains to be explored.

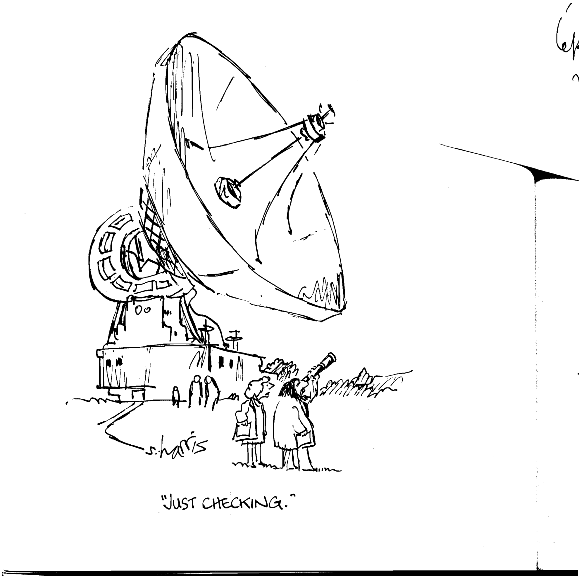

Figure 10: Just Checking!

6. Finding a Discrete Solution Close to a Continuous Approximation

Inspired by Yu and the previous section, given a solution Z_0 of problem $(*_1)$, we look for pairs $(X, R) \in \mathcal{K} \times \mathbf{O}(K)$ (where R is a $K \times K$ orthogonal matrix), with $\left\|X^j\right\| = \left\|Z_0^j\right\|$ for $j = 1, \ldots, K$, that minimize

$$\varphi(X, R) = \|X - Z_0 R\|_F .$$

6. Finding a Discrete Solution Close to a Continuous Approximation

Inspired by Yu and the previous section, given a solution Z_0 of problem $(*_1)$, we look for pairs $(X, R) \in \mathcal{K} \times \mathbf{O}(K)$ (where R is a $K \times K$ orthogonal matrix), with $\|X^j\| = \left\|Z_0^j\right\|$ for $j = 1, \ldots, K$, that minimize

$$\varphi(X, R) = \|X - Z_0 R\|_F .$$

The key to minimizing $\|X - ZR\|_F$ rests on the following equation:

$$\|X - ZR\|_F^2 = \mathrm{tr}(X^\top X) - 2\mathrm{tr}(R^\top Z^\top X) + \mathrm{tr}(Z^\top Z).$$

6. Finding a Discrete Solution Close to a Continuous Approximation

Inspired by Yu and the previous section, given a solution Z_0 of problem $(*_1)$, we look for pairs $(X, R) \in \mathcal{K} \times \mathbf{O}(K)$ (where R is a $K \times K$ orthogonal matrix), with $\left\|X^j\right\| = \left\|Z_0^j\right\|$ for $j = 1, \ldots, K$, that minimize

$$\varphi(X, R) = \|X - Z_0 R\|_F .$$

The key to minimizing $\|X - ZR\|_F$ rests on the following equation:

$$\|X - ZR\|_F^2 = \mathrm{tr}(X^\top X) - 2\mathrm{tr}(R^\top Z^\top X) + \mathrm{tr}(Z^\top Z).$$

Therefore, *minimizing* $\|X - ZR\|_F^2$ is equivalent to *maximizing* $\mathrm{tr}(R^\top Z^\top X)$.

This can be done by alternating steps during which we minimize $\varphi(X, R) = \|X - ZR\|_F$ with respect to X *holding R fixed*, and steps during which we minimize $\varphi(X, R) = \|X - ZR\|_F$ with respect to R *holding X fixed*.

This can be done by alternating steps during which we minimize $\varphi(X, R) = \|X - ZR\|_F$ with respect to X *holding R fixed*, and steps during which we minimize $\varphi(X, R) = \|X - ZR\|_F$ with respect to R *holding X fixed*.

For this second step, we use the following (known) proposition.

Proposition 4

For any $n \times n$ matrix A and any orthogonal matrix Q, we have

$$\max\{\mathrm{tr}(QA) \mid Q \in \mathbf{O}(n)\} = \sigma_1 + \cdots + \sigma_n,$$

where $\sigma_1 \geq \cdots \geq \sigma_n$ are the singular values of A. Furthermore, this maximum is achieved by $Q = VU^\top$, where $A = U\Sigma V^\top$ is any SVD for A.

As a corollary of Proposition 4 (with $A = Z^\top X$ and $Q = R^\top$), we get the following (known) result:

Proposition 5

For any two fixed $N \times K$ matrices X and Z, the minimum of the set

$$\{\|X - ZR\|_F \mid R \in \mathbf{O}(K)\}$$

is achieved by $R = UV^\top$, for any SVD decomposition $U\Sigma V^\top = Z^\top X$ of $Z^\top X$.

As a corollary of Proposition 4 (with $A = Z^\top X$ and $Q = R^\top$), we get the following (known) result:

Proposition 5

For any two fixed $N \times K$ matrices X and Z, the minimum of the set

$$\{\|X - ZR\|_F \mid R \in \mathbf{O}(K)\}$$

is achieved by $R = UV^\top$, for any SVD decomposition $U\Sigma V^\top = Z^\top X$ of $Z^\top X$.

Maximizing $\mathrm{tr}(R^\top Z^\top X) = \mathrm{tr}(X(ZR)^\top)$ holding R fixed can be done using a method of nonmaximal suppression, but some issues remain to be resolved.

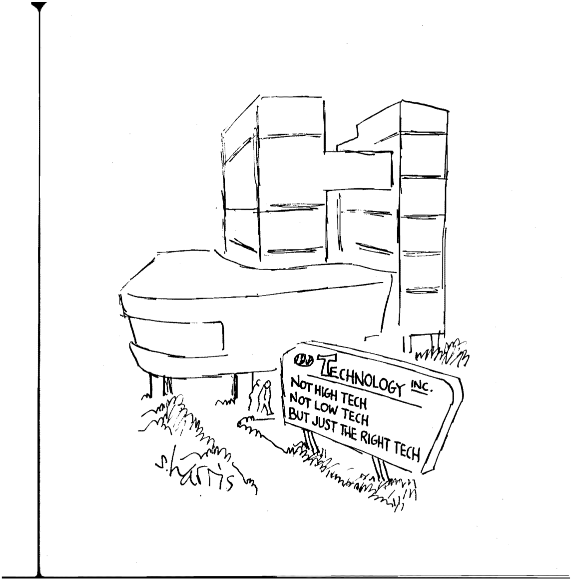

Figure 11: Just the right tech