CS 860 Spectral graph theory. Spring 2019, Waterloo.

Lecture 9: High dimensional expanders

We study a definition of high dimensional expanders through the notions of simplicial complexes and links.

We see the theorem of Oppenheim relating the expansion of top links to the expansion of bottom links, and use it to show that the natural simplicial complexes from matroids are high dimensional expanders.

High dimensional expanders

As we have seen, expander graphs have nice combinatorial, probabilistic and algebraic properties, and this is why they have a rich theory and found connections and applications in diverse areas.

It is easy to generalize the definition of expander graphs to higher dimensions for some properties (e.g. the combinatorial expansion in hypergraphs), but it is not easy to find a definition in higher dimension that generalizes all nice properties of expander graphs.

There are various definitions of high dimensional expanders using concepts from algebraic topology; see [Lubotzky] for a survey with motivations and applications.

We will study a more recent and elementary definition proposed by Kaufman and Mass [KM], with motivations from random walks.

Simplicial complexes

A simplicial complex is a high dimension generalization of a graph.

Definition (simplicial complex) A set system is a pair $X=(U,\mathcal{F})$ with U as the ground set and $\mathcal{F}$ is a set of subsets of U. A simplicial complex is a set system that is downward closed, i.e. if $\sigma \in \mathcal{F}$ and $\sigma' \subset \sigma$, then $\sigma' \in \mathcal{F}$. (We use Greek letters σ, τ, η etc for subsets in $\mathcal{F}$.)

Definition (face, dimension, pure simplicial complex) Any subset $\sigma \in \mathcal{F}$ is called a face of the complex. A face σ is of dimension i if $|\sigma| = i+1$, e.g. a 0-dimensional face is a singleton (a vertex), a 1-dimensional face is a pair (an edge), a 2-dimensional face is a triple, etc.

Given a simplicial complex $X=(U,\mathcal{F})$, we use $X(i)$ to denote the set of faces of dimension i.

A simplicial complex is d-dimensional if the maximum face size is $d+1$.

A d-dimensional simplicial complex is pure if every maximal face is of size $d+1$.

Definition (Balanced weights) We can assign a positive weight $w(\sigma)$ to each face $\sigma \in \mathcal{F}$.

The weight function $w: X \to \mathbb{R}_{>0}$ is balanced if the following holds for every face $\tau \in X(k)$ for every $0 \le k \le d-1$:

$$w(\tau) = \sum_{\sigma \in X(k+1) : \sigma \supset \tau} w(\sigma).$$

For a pure simplicial complex, we can define a balance weight function by assigning arbitrary positive weights to maximal faces and use the above equation to define the weights of lower dimensional faces recursively.

Then, for any $\tau \in X(k)$, we have $w(\tau) = (d-k)! \sum_{\sigma \in X(d) : \sigma \subset \tau} w(\sigma)$.

A natural choice is to assign weight one to each face of dimension d. In this case, for $\tau \in X(k)$, $w(\tau)$ is simply $(d-k)!$ times the number of maximal faces containing τ. This is the weight function that we are going to use next time.

It is a very general definition. We can associate a simplicial complex to many combinatorial objects.

Example (simplicial complex from spanning trees) Given an graph $G = (V, E)$, we can define a simplicial complex $X = (E, \mathcal{F})$ where the ground set in X is the edge set in G.

A subset of edges $E' \subseteq E$ is in $\mathcal{F}$ if and only if E' forms an acyclic subgraph in G.

It should be clear that X is a pure simplicial complex.

If G is connected, then the maximal faces are of size $|V|-1$ and so X is of dimension $|V|-2$.

Example (simplicial complex from matroids). A matroid naturally corresponds to a simplicial complex.

A matroid $M = (U, \mathcal{I})$ is a set system where U is the ground set and $\mathcal{I}$ is a set of subsets of U which satisfies the following two properties:

① $\mathcal{I}$ is downward closed, i.e. $S \in \mathcal{I}$ and $T \subseteq S$ implies $T \in \mathcal{I}$.

② if $S, T \in \mathcal{I}$ and $|T| > |S|$, then there exists $i \in T \setminus S$ such that $S \cup \{i\} \in \mathcal{I}$.

So, by ①, $M = (U, \mathcal{I})$ is a simplicial complex. By ②, $M = (U, \mathcal{I})$ is a pure simplicial complex.

The sets in $\mathcal{I}$ are usually called the independent sets, and the maximal sets are called bases.

It is not difficult to check that the simplicial complex from spanning tree is a matroid.

A more general example is linear matroids of a matrix. Given a matrix $A \in \mathbb{F}^{m \times n}$, the linear matroid is defined as $M = ([n], \mathcal{I})$ where the ground set $[n]$ is the set of columns in A and a subset S of columns is independent if and only if they are linearly independent.

It is not difficult to check that it is a matroid and includes the matroid from spanning trees as a special case.

Links and skeletons

The following are some important concepts from simplicial complexes that are crucial in solving problems. These concepts have natural meanings in combinatorial complexes, but we may not look at these concepts if we don't look at the combinatorial objects from the perspective of simplicial complexes.

Definition (Links) Given a simplicial complex $X=(U,\mathcal{F})$ and a face $\tau \in \mathcal{F}$, the link of τ is denoted by X_τ, which is defined by the faces in $\mathcal{F}_\tau = \{\sigma \setminus \tau \mid \sigma \in \mathcal{F} \text{ and } \sigma \supset \tau\}$.

In words, X_τ is defined by the faces η that can be used to extend τ, i.e. $\eta \cup \tau \in \mathcal{F}$.

If X is a pure d-dimensional simplicial complex and $\tau \in X(k)$, then X_τ is a pure $(d-k-1)$-dimensional simplicial complex (the empty set is a face of dimension -1).

In the spanning tree complex $X=(E,\mathcal{I})$, given a subset of (acyclic) edges $F \in \mathcal{I}$, the link X_F is defined such that a subset of edges F' is a face in X_τ iff $F \cup F' \in \mathcal{I}$, i.e. the faces in X_F are the set of subsets of edges that can extend F to form an acyclic subgraph.

In matroid terminology, the link X_F is obtained from X by "contracting" the elements in F.

Definition (k-skeletons) Given $X=(U,\mathcal{F})$, the k-skeleton of X is the simplicial complex $X_k = (U, \mathcal{F}_k)$ where $\mathcal{F}_k$ is the set of faces of $\mathcal{F}$ with dimension at most k.

The special case of 1-skeleton will be of particular interest, as it could be thought of as the underlying graph of the simplicial complex.

When there are weights, we use the same weight for every face in $\mathcal{F}_k$.

Definition (Adjacency / Correlation matrix) Given $X=(U,\mathcal{F})$ with a balanced weight function, the adjacency or the correlation matrix $A \in \mathbb{R}^{|U|\times|U|}$ is defined as $A(i,j) = \begin{cases} w(ij) & \text{if } ij \in \mathcal{F} \\ 0 & \text{otherwise} \end{cases}$.

In words, it is the adjacency matrix of the 1-skeleton of the simplicial complex.

When w is a balanced weight function of a d-dimensional simplicial complex when $w(\sigma)=1$ for every maximal face (or basis, in matroid terminology) σ, then $w(i,j)$ is $(d-1)!$ times the number of maximal faces containing both i and j (e.g. number of spanning trees containing two edges).

In this setting, we can think of the adjacency matrix as a correlation matrix.

Definition (Normalized adjacency / correlation matrix) Let $X=(U,\mathcal{F})$ with a balanced weight function $w:\mathcal{F}\to\mathbb{R}_{>0}$.

The i-th row sum of the adjacency / correlation matrix A is $\sum_{j=1}^{|U|} w(ij) = w(i)$, as w is balanced.

Let $W \in \mathbb{R}^{|U|\times|U|}$ be the diagonal matrix with $W(i,i) = w(i)$ for $1 \le i \le |U|$.

The normalized adjacency / correlation matrix is defined as $\mathcal{A} = W^{-\frac{1}{2}} A W^{-\frac{1}{2}}$.

As seen in L03, the eigenvalues of $\mathcal{A}$ satisfy $1=\alpha_1 \ge \alpha_2 \ge \dots \ge \alpha_n \ge -1$, with $v_1 = W^{\frac{1}{2}}\vec{1} / \|W^{\frac{1}{2}}\vec{1}\|$.

A closely related matrix is the random walk matrix P defined as $P = W^{-1}A$.

As seen in L07, P is similar to $\mathcal{A}$ and it follows that they have the same spectrum.

The random walk matrix will play an important role in our study.

Link expanders

Finally, we can state the definition of high dimensional expander proposed by Kaufman and Mass.

Definition (Link expanders) Let $X=(U,\mathcal{F})$ be a pure d-dimensional simplicial complex with a balanced weight function $w:\mathcal{F}\to\mathbb{R}_{>0}$.

We say that X is an α-link expander if for any face $\sigma \in \mathcal{F}$ of dimension at most $d-2$, the second eigenvalue of the normalized adjacency / correlation matrix of X_σ is at most α.

So, if we view the 1-skeleton of a link as a weighted graph, the definition requires that it is an expander graph, with large conductance by Cheeger's inequality.

Or, the random walk matrix P_σ for any link X_σ up to dimension $d-2$ is fast mixing.

Oppenheim's theorem

To show that a simplicial complex is an α-link expander, we need to bound the second largest eigenvalue of the normalized adjacency / correlation matrix for every link up to dimension $d-2$.

It is usually easier to work with the correlation matrix of the "top" links (with dimension $d-2$), because they are unweighted if the weight of every maximal face is the same (e.g. $w(\sigma)=1\ \forall\sigma$).

In a way, the top links seem to be the most important, as they are the finest structures of the simplicial complex. So, if they are all good expanders, then it should be the case for all links as well.

Oppenheim's theorem provides a way to bound the second largest eigenvalue of the "bottom" links from the second largest eigenvalue of the "top" links.

Theorem Let $X=(U,\mathcal{F})$ be a d-dimensional simplicial complex with a balanced weight function $w: \mathcal{F} \to \mathbb{R}_{\geq 0}$.
Suppose the graph $G=(X(0),X(1))$ is connected and for all $v \in X(0)$, the second largest eigenvalue of the normalized adjacency/correlation matrix A_v of X_v is at most α.
Then the second largest eigenvalue of the normalized adjacency/correlation matrix A of X is at most $\frac{\alpha}{1-\alpha}$.

Suppose we showed that the second largest eigenvalue of every link of dimension $d-2$ is at most α, and furthermore that every link is connected (i.e. the second largest eigenvalue of every link is strictly smaller than one).
Then we can apply the above theorem recursively (to links of dimension $d-3, d-4, \dots, -1$) to get the following.

Theorem Let $X=(U,\mathcal{F})$ be a d-dimensional simplicial complex with a balanced weight function $w: \mathcal{F} \to \mathbb{R}_{>0}$.
Suppose that the 1-skeleton of every link is connected and the second largest eigenvalue of the normalized adjacency/correlation matrix of every link of dimension $d-2$ is at most α.
Then X is a $\frac{\alpha}{1-(d-1)\alpha}$-link expander.

Before we see a proof of Oppenheim's theorem, we first apply it to matroid complexes and show that they are 0-link expanders (i.e. second largest eigenvalue is non-positive).

Theorem Let $M=(U,\mathcal{I})$ be a simplicial complex satisfying the matroid axioms, with a balanced weight function where all maximal faces are of weight one.
Then M is a 0-link expander.

Proof To apply Oppenheim's theorem, we only need to prove that 1-skeleton of every link is connected, and that the second largest eigenvalue of the normalized adjacency matrix of the top links is at most 0.
The first claim follows from the second axiom of matroids, and is left as a simple exercise.
For the second claim, we first consider the (un-normalized) adjacency matrix A_σ of a face σ of dimension $d-2$.
Since the maximal faces are of weight one, $A_\sigma(i,j) = \begin{cases} w(\sigma \cup i \cup j) = 1 & \text{if } \sigma \cup i \cup j \text{ is a maximal face} \\ 0 & \text{otherwise.} \end{cases}$

We would like to argue that A has at most one positive eigenvalue, and this would imply that the normalized adjacency matrix $\mathcal{A}$ has at most one positive eigenvalue by Courant-Fischer (exercise).

To argue that A has at most one positive eigenvalue, let us start with the spanning tree complex.

In a spanning tree complex, the maximal faces are of size $n-1$, where n is the number of vertices.

Given a face F of size $n-3$ (of dimension $d-2$), the subgraph formed by the edges in F has exactly three components left.

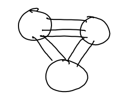

The edges remained in the link X_F are the edges with endpoints in different components.

Two edges e,f in X_F form a face of size 2 iff $F \cup e \cup f$ form a spanning tree iff e and f are not parallel edges if we contract the three components into single vertices.

In other words, the edges can be partitioned three equivalent classes E_1, E_2, E_3 such that two edges e,f form a face of size 2 in X_F iff they do not belong to the same class.

So, the adjacency matrix A can be written as $J - \chi_{E_1}\chi_{E_1}^T - \chi_{E_2}\chi_{E_2}^T - \chi_{E_3}\chi_{E_3}^T$, which is a rank one matrix minus three positive semidefinite matrices.

	E_1	E_2	E_3
E_1	0	1	1
E_2	1	0	1
E_3	1	1	0

It follows from Courant-Fischer that A has at most one positive eigenvalue (exercise), and this concludes the spanning tree case.

The same proof works for linear matroids, where two columns i,j do not form a face of size 2 iff they are parallel (in the linear algebraic sense), and so again the columns can be partitioned into equivalence classes $E_1, E_2, \ldots, E_k$ (k not necessarily equal to 3) so that $A = J - \sum_{i=1}^{k} \chi_{E_i}\chi_{E_i}^T$.

In general, this holds for arbitrary matroids and is known as the matroid partition property and the same proof works (details omitted). $\square$

Proof of Oppenheim's theorem

Inner product and Rayleigh quotient

For the proof, it will be more convenient to work with the random walk matrix P and a different inner product, so let us set up these notations first.

Since $\mathcal{A} = W^{-\frac{1}{2}} A W^{-\frac{1}{2}} \in \mathbb{R}^{n \times n}$ is real symmetric, there is an orthonormal set of eigenvectors $v_1, v_2, \ldots, v_n$, with eigenvalues $1 = \alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n \geq -1$.

As seen in L07, the random walk matrix $P = W^{-1}A$ is similar to $\mathcal{A}$ and has the same spectrum $1 \geq \alpha_2 \geq \ldots \geq \alpha_n \geq -1$.

with eigenvectors $u_1 = W^{-\frac{1}{2}} v_1,\ u_2 = W^{-\frac{1}{2}} v_2, \ldots, u_n = W^{-\frac{1}{2}} v_n$.

The vectors $u_1, \ldots, u_n$ are not orthonormal using the inner product, but if we define a new inner product $\langle u_i, u_j \rangle_W \overset{\text{def}}{=} \sum_{k=1}^{n} W(k) \cdot u_i(k) \cdot u_j(k) = \langle u_i, W u_j \rangle$, then $\langle u_i, u_j \rangle_W = 0$ for $i \neq j$ and $\langle u_i, u_i \rangle_W = 1$,

and so the vectors $u_1, u_2, \ldots, u_n$ are W-orthonormal with respect to this inner product.

Recall from L02 that the eigenvalues $\alpha_1, \ldots, \alpha_n$ of $\mathcal{A}$ are characterized by the Rayleigh quotient

$$\frac{x^T \mathcal{A} x}{x^T x} = \frac{x^T W^{-\frac{1}{2}} A W^{-\frac{1}{2}} x}{x^T x} = \frac{y^T A y}{y^T W y} \quad \text{by letting } y = W^{-\frac{1}{2}} x \text{ and so } x = W^{\frac{1}{2}} y$$

$$= \frac{y^T W W^{-1} A y}{y^T W y} = \frac{y^T W P y}{y^T W y} = \frac{\langle y, Py \rangle_W}{\langle y, y \rangle_W}$$

By doing this change of basis $y = W^{-\frac{1}{2}} x$ and $u_i = W^{-\frac{1}{2}} v_i$ and using the inner product $\langle \cdot, \cdot \rangle_W$, the Courant-Fischer theorem can be reformulated as $\alpha_k = \max_{S:\dim(S)=k} \min_{y \in S} \frac{\langle y, Py \rangle_W}{\langle y, y \rangle_W}$.

The advantage of working with P is that we know $u_1 = \vec{1} / \|\vec{1}\|_W$ as $P\vec{1} = \vec{1}$.

We can write any $y \in \mathbb{R}^n$ as $y = c_1 u_1 + \ldots + c_n u_n$ with $c_i = \langle y, u_i \rangle_W$ as $u_1, \ldots, u_n$ are W-orthonormal,

and in particular $c_1 = \langle y, u_1 \rangle_W = \langle y, \vec{1} \rangle_W / \|\vec{1}\|_W$

<u>The proof</u>

Let $A \in \mathbb{R}^{n \times n}$ be the adjacency matrix of the empty link, i.e. $A(i,j) = \begin{cases} w(i \cup j) & \text{if } i \cup j \in X(1) \\ 0 & \text{otherwise.} \end{cases}$

Let $A_k \in \mathbb{R}^{n \times n}$ be the (extended) adjacency matrix of the link $\{k\}$, i.e. $A_k(i,j) = \begin{cases} w(i \cup j \cup k) & \text{if } i \cup j \cup k \in X(2) \\ 0 & \text{otherwise.} \end{cases}$

Note that A_k is essentially the same matrix as we defined earlier, but extend it back to include the zero rows and zero columns so that it is of the same dimension as A.

Since the weight function $w: \mathcal{F} \to \mathbb{R}_{>0}$ is balanced, we have $A = \sum_{k=1}^{n} A_k$, as $w(i \cup j) = \sum_{\substack{k=1 \\ i \cup j \cup k \in X(2)}}^{n} w(i \cup j \cup k)$.

For each k, $W_k \in \mathbb{R}^{n \times n}$ is the diagonal matrix with $W_k(i,i) = w(k \cup i)$ for $k \cup i \in X(1)$ and zero otherwise,

$w_k \in \mathbb{R}^n$ is a vector with $w_k(i) = w(k \cup i)$ for $k \cup i \in X(1)$ and zero otherwise, and $P_k = W_k^{-1} A_k$ is the random walk matrix for the link $\{k\}$.

Let y be an eigenvector of P with eigenvalue λ.

As discussed above, $\lambda \|y\|_W^2 = \langle y, Py \rangle_W = \langle y, Ay \rangle = \sum_{k=1}^{n} \langle y, A_k y \rangle = \sum_{k=1}^{n} \langle y, P_k y \rangle_{W_k}$.

We would like to bound λ by α, the second largest eigenvalue of P_k.

Note that the first eigenvector of P_K is still $\vec{1}/\|1\|_{W_K}$.

For each term $\langle y, P_K y\rangle_{W_K}$, we write $y = c_1 \frac{\vec{1}}{\|\vec{1}\|_{W_K}} + y_K^{\perp}$, where $c_1 = \langle y, \frac{\vec{1}}{\|\vec{1}\|_{W_K}}\rangle_{W_K}$ and $\langle \frac{\vec{1}}{\|\vec{1}\|_{W_K}}, y_K^{\perp}\rangle_{W_K} = 0$.

So, $y = \frac{\langle y, \vec{1}\rangle_{W_K}}{\langle \vec{1}, \vec{1}\rangle_{W_K}} \vec{1} + y_K^{\perp}$ where $y_K^{\perp}$ is W_K-orthonormal to the first eigenvector of P_K, and has

$\langle y_K^{\perp}, P_K y_K^{\perp}\rangle \leq \alpha \|y_K^{\perp}\|_{W_K}^2$ by the assumption of the second largest eigenvalue of P_K.

Now, continue with $\lambda \|y\|_W^2 = \sum_{K=1}^{n} \langle y, P_K y\rangle_{W_K} = \sum_{K=1}^{n} \langle y_K^1 + y_K^{\perp}, P_K(y_K^1 + y_K^{\perp})\rangle_{W_K}$ where $y_K^1 = \frac{\langle y, \vec{1}\rangle_{W_K}}{\langle \vec{1}, \vec{1}\rangle_{W_K}} \vec{1}$

$$= \sum_{K=1}^{n} \left(\langle y_K^1, P_K y_K^1\rangle_{W_K} + \langle y_K^{\perp}, P_K y_K^{\perp}\rangle_{W_K} \right),$$

where the last equality holds because $\langle y_K^1, P_K y_K^{\perp}\rangle_{W_K} = \langle y_K^{\perp}, P_K y_K^1\rangle_{W_K} = 0$ as $P_K \vec{1} = \vec{1}$ and $\langle y_K^{\perp}, 1\rangle_{W_K} = 0$.

As argued before, $\langle y_K^{\perp} P_K y_K^{\perp}\rangle_{W_K} \leq \alpha \|y_K^{\perp}\|_{W_K}^2$, and so

$$\sum_{K=1}^{n} \langle y_K^{\perp}, P_K y_K^{\perp}\rangle_{W_K} \leq \alpha \sum_{K=1}^{n} \|y_K^{\perp}\|_{W_K}^2 = \alpha \sum_{K=1}^{n} \left(\|y\|_{W_K}^2 - \|y_K^1\|_{W_K}^2\right) = \alpha \|y\|_W^2 - \alpha \sum_{K=1}^{n} \|y_K^1\|_{W_K}^2,$$

where the last equality holds as $\sum_{K=1}^{n} \|y\|_{W_K}^2 = \sum_{K=1}^{n} \sum_{i=1}^{n} W_K(i)^2 y(i)^2 = \sum_{i=1}^{n} W(i)^2 y(i)^2 = \|y\|_W^2$,

where the second last equality uses that W is a balanced weight function.

Now, notice that $\sum_{K=1}^{n} \langle y_K^1, P_K y_K^1\rangle_{W_K} = \sum_{K=1}^{n} \|y_K^1\|_{W_K}^2$ as $\vec{1}$ is an eigenvector of P_K with eigenvalue 1,

and so $\lambda \|y\|_W^2 \leq \alpha \|y\|_W^2 + (1-\alpha) \sum_{K=1}^{n} \|y_K^1\|_{W_K}^2$.

It remains to compute $\sum_{K=1}^{n} \|y_K^1\|_{W_K}^2 = \sum_{K=1}^{n} \frac{\langle y, \vec{1}\rangle_{W_K}^2}{\|\vec{1}\|_{W_K}^2}$ since $y_K^1 = \frac{\langle y, \vec{1}\rangle_{W_K}}{\langle \vec{1}, \vec{1}\rangle_{W_K}} \vec{1}$.

Note that $\|\vec{1}\|_{W_K}^2 = \sum_{i=1}^{n} W_K(i) = \sum_{i=1}^{n} W(K \cup i) = W(K)$ as W is a balanced weight function,

and $\langle y, \vec{1}\rangle_{W_K} = \sum_{i=1}^{n} y(i) W_K(i) = \sum_{i=1}^{n} y(i) W(K \cup i) = (Ay)_K$ where A is the adjacency matrix.

So, $\sum_{K=1}^{n} \|y_K^1\|_{W_K}^2 = \sum_{K=1}^{n} \frac{\langle y, \vec{1}\rangle_{W_K}^2}{\|\vec{1}\|_{W_K}^2} = \sum_{K=1}^{n} \frac{(Ay)_K^2}{W(K)} = \langle Ay, W^{-1}Ay\rangle = \langle W^{-1}Ay, W^{-1}Ay\rangle_W = \|Py\|_W^2 = \lambda^2 \|y\|_W^2$.

Therefore, we have $\lambda \|y\|_W^2 \leq \alpha \|y\|_W^2 + (1-\alpha)\lambda^2 \|y\|_W^2$, and thus $\lambda - \lambda^2 \leq \alpha(1-\lambda^2)$.

Finally, we use the assumption that the empty link is connected to obtain that $\lambda < 1$, and so we can divide both sides by $1-\lambda$ to conclude that $\lambda \leq \alpha(1+\lambda)$ and thus $\lambda \leq \frac{\alpha}{1-\alpha}$. □

References

- High dimensional expanders, by Lubotzky.
- High dimensional combinatorial random walks and colorful expansion, by Kaufman and Mass.
- Local spectral expansion approach to high dimensional expanders part I: Descent of spectral gap, by Oppenheim.

- High dimensional combinatorial random walks and colorful expansion. by Kaufman and Mass.
- Local spectral expansion approach to high dimensional expanders part I: Descent of spectral gap. by Oppenheim.
- Log-concave polynomial II: High-dimensional walks and an FPRAS for counting bases of a matroid,

  by Anari, Liu, Oveis Gharan, Vinzart.