Jonas Hübotter
based on Rasmus Kyng's lectures at ETH Zurich

Graph Algorithms and Optimization

Contents. Convex optimization and duality. Spectral Graph Theory. Combinatorial Graph Algorithms. Electrical Flows.

Contributions are welcome at `https://github.com/jonhue/graph-algorithms-and-optimization`.

ACKNOWLEDGEMENT

The contents of this summary are based on the lecture "Advanced Graph Algorithms and Optimization" given by Rasmus Kyng at ETH Zurich in the spring of 2022. Certain parts are similar/taken from the lecture notes.

Contents

PART I

Preliminaries

1

Electrical Flows

A classical graph problem is the flow of electrical currents through a network of resistors. Such a network $G = (V, E, \boldsymbol{r})$ can be described by a set of vertices V, set of wires (or edges) E, and resistances $\boldsymbol{r} \in \mathbb{R}^{|E|}_{>0}$ of wires. We are interested in finding the *electrical flow* $\tilde{\boldsymbol{f}} \in \mathbb{R}^{|E|}$ through the network, assigning to each wire the current that is transported per unit time. Alternatively, we can think of *electrical voltages* $\tilde{\boldsymbol{x}} \in \mathbb{R}^{|V|}$ at the vertices, which Ohm's law relates to the electrical flow.

By *Ohm's law*, we have that for any wire $e \in E$,

$$\tilde{\boldsymbol{f}}(e) = \frac{\tilde{\boldsymbol{x}}(e)}{\boldsymbol{r}(e)}, \quad \tilde{\boldsymbol{x}}(e) = \tilde{\boldsymbol{f}}(e) \cdot \boldsymbol{r}(e), \tag{1.1}$$

where $\tilde{\boldsymbol{x}}(\{u, v\}) \doteq \tilde{\boldsymbol{x}}(u) - \tilde{\boldsymbol{x}}(v)$ is the voltage difference of vertices u and v. For any *flow* $\boldsymbol{f} \in \mathbb{R}^{|E|}$, the *net flow* of current at a vertex $u \in V$ is given as,

$$\sum_{v \sim u} \boldsymbol{f}(v, u). \tag{1.2}$$

We use $v \sim u$ to denote all v that are adjacent to u.

We say that a flow routes *demand* $\boldsymbol{d} \in \mathbb{R}^{|V|}$ if the net flow at every vertex is $\boldsymbol{d}(v)$. The fact that at vertices with zero demand, the flow is conserved[1] is also known as *Kirchhoff's current law*.

[1] As much current is flowing into the vertex as is flowing out of it.

To keep track of the direction of flow on each edge, we assign an arbitrary direction to each edge (we "orient" G) and only consider non-negative flows, $\boldsymbol{f} \in \mathbb{R}^{|E|}_{\geq 0}$. Clearly, for any previously feasible flow, we can assign directions in such a way that the flow remains feasible.

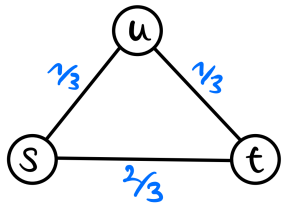

Figure 1.1: Example of an electrical flow (shown in blue) with voltages

$$\boldsymbol{x}(s) = 0, \quad \boldsymbol{x}(u) = 1, \quad \boldsymbol{x}(t) = 2$$

and unit resistances, routing demands

$$\boldsymbol{d}(s) = -1, \quad \boldsymbol{d}(u) = 0, \quad \boldsymbol{d}(t) = 1.$$

1.1 *The Laplacian Matrix*

Definition 1.1 (Adjacency matrix)**.** The *adjacency matrix* of a graph G, $\tilde{\boldsymbol{A}} \in \mathbb{R}^{|V| \times |V|}$, is defined as,[2]

[2] By $\boldsymbol{A}$ we will later denote the weighted adjacency matrix.

$$\tilde{\boldsymbol{A}}(u,v) \doteq \begin{cases} 1 & \text{if } u \sim v \\ 0 & \text{otherwise.} \end{cases} \tag{1.3}$$

Definition 1.2 (Incidence matrix)**.** The *incidence matrix* of an oriented graph G, $\boldsymbol{B} \in \mathbb{R}^{|V|\times|E|}$, is defined as,

$$\boldsymbol{B}(v,e) \doteq \begin{cases} 1 & \text{if } e = (u,v) \text{ for some } u \in V \\ -1 & \text{if } e = (v,u) \text{ for some } u \in V \\ 0 & \text{otherwise.} \end{cases} \tag{1.4}$$

Each column of $\boldsymbol{B}$ only has two non-zero entries and sums to one.

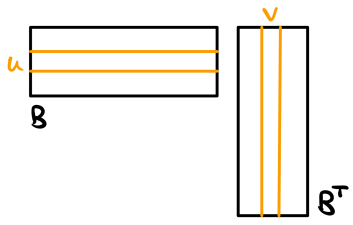

Figure 1.2: Illustration of the matrix product $\boldsymbol{B}\boldsymbol{B}^\top$.

Lemma 1.3. $\boldsymbol{B}\boldsymbol{B}^\top = \operatorname{diag}\{\deg(v)\}_{v\in V} - \tilde{\boldsymbol{A}}$.

Proof. The dot product of the rows, corresponding to the same vertex v, produces exactly $\deg(v)$. All other dot products between rows corresponding to vertices u and v are -1 iff $u \sim v$ and 0 otherwise. □

We can now also write the net flow constraint,

$$\boldsymbol{B}\boldsymbol{f} = \boldsymbol{d}. \tag{1.5}$$

We define $\boldsymbol{R} \doteq \operatorname{diag}\{\boldsymbol{r}(e)\}_{e\in E}$ and then have that Ohm's law can be expressed as,

$$\boldsymbol{B}^\top \tilde{\boldsymbol{x}} = \boldsymbol{R}\tilde{\boldsymbol{f}}, \quad \text{or equivalently,} \quad \boldsymbol{R}^{-1}\boldsymbol{B}^\top \tilde{\boldsymbol{x}} = \tilde{\boldsymbol{f}}. \tag{1.6}$$

If the net flow constraint is satisfied, this yields,

$$\underbrace{\boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{B}^\top}_{\text{Laplacian}}\tilde{\boldsymbol{x}} = \boldsymbol{B}\tilde{\boldsymbol{f}} = \boldsymbol{d}. \tag{1.7}$$

Definition 1.4 (Laplacian matrix)**.** The *Laplacian matrix* of an oriented graph G is defined as,

$$\boldsymbol{L} \doteq \boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{B}^\top = \boldsymbol{B}\boldsymbol{W}\boldsymbol{B}^\top \in \mathbb{R}^{|V|\times|V|}, \tag{1.8}$$

where $\boldsymbol{W} \doteq \boldsymbol{R}^{-1}$ is a diagonal matrix of weights $\boldsymbol{w}(e) \doteq \frac{1}{\boldsymbol{r}(e)}$.

Intuitively, the weight of an edge can be understood as how "connected" the two vertices at its endpoints are. In contrast, the resistance of an edge is smaller when endpoints are well-connected.

We will now learn a little more about Laplacian matrices.

Definition 1.5 (Weighted adjacency matrix)**.** The *weighted adjacency matrix* of a graph G, $\boldsymbol{A} \in \mathbb{R}^{|V|\times|V|}$, is defined as,

$$\boldsymbol{A}(u,v) \doteq \begin{cases} w(\{u,v\}) & \text{if } u \sim v \\ 0 & \text{otherwise.} \end{cases} \tag{1.9}$$

Lemma 1.6. *$\boldsymbol{A}$ is symmetric, that is, $\boldsymbol{A} = \boldsymbol{A}^\top$.*

Proof. This follows immediately from the fact that G is undirected. □

Definition 1.7 (Weighted degree)**.** The *weighted degree* of a vertex $v \in V$ is given as,

$$\boldsymbol{d}(v) \doteq \sum_{\{u,v\}\in E} \boldsymbol{w}(\{u,v\}). \tag{1.10}$$

We write $\boldsymbol{D} \doteq \operatorname{diag}\{\boldsymbol{d}(v)\}_{v\in V}$.

Lemma 1.8. *$\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{A}$.*

Proof. The proof is identical to the proof of lemma 1.3, only that every entry is now weighted, due to the additional factor $\boldsymbol{W}$. □

Corollary 1.9. *$\boldsymbol{L}$ is symmetric.*

Proof. This directly follows from the fact that $\boldsymbol{D}$ and $\boldsymbol{A}$ are symmetric.[3] □

[3] Diagonal matrices like $\boldsymbol{D}$ are trivially symmetric.

Lemma 1.10. *For any $\boldsymbol{x} \in \mathbb{R}^{|V|}$, we have,*

$$\boldsymbol{x}^\top \boldsymbol{L}\boldsymbol{x} = \sum_{\{u,v\}\in E} \boldsymbol{w}(\{u,v\})[\boldsymbol{x}(u) - \boldsymbol{x}(v)]^2 \geq 0. \tag{1.11}$$

Proof. We have,

$$\begin{aligned}
&\boldsymbol{x}^\top \boldsymbol{L}\boldsymbol{x} = \boldsymbol{x}^\top \boldsymbol{D}\boldsymbol{x} - \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}. \\
&\boldsymbol{x}^\top \boldsymbol{D}\boldsymbol{x} = \sum_{v\in V} \boldsymbol{d}(v)\boldsymbol{x}(v)^2 = \sum_{\{u,v\}\in E} \boldsymbol{w}(\{u,v\})[\boldsymbol{x}(u)^2 + \boldsymbol{x}(v)^2]. \\
&\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} = \sum_{v\in V} \boldsymbol{x}(v)(\boldsymbol{A}\boldsymbol{x})(v) \\
&\qquad = \sum_{v\in V} \boldsymbol{x}(v) \sum_{u\in V} \boldsymbol{A}(v,u)\boldsymbol{x}(u) \\
&\qquad = \sum_{v,u\in V} \boldsymbol{w}(\{u,v\})\boldsymbol{x}(u)\boldsymbol{x}(v) \\
&\qquad = 2 \sum_{\{u,v\}\in E} \boldsymbol{w}(\{u,v\})\boldsymbol{x}(u)\boldsymbol{x}(v).
\end{aligned}$$

Combining the above equalities, we obtain,

$$\begin{aligned} \boldsymbol{x}^\top \boldsymbol{L}\boldsymbol{x} &= \sum_{\{u,v\}\in E} \boldsymbol{w}(\{u,v\})[\boldsymbol{x}(u)^2 + \boldsymbol{x}(v)^2] - 2\sum_{\{u,v\}\in E} \boldsymbol{w}(\{u,v\})\boldsymbol{x}(u)\boldsymbol{x}(v) \\ &= \sum_{\{u,v\}\in E} \boldsymbol{w}(\{u,v\})[\boldsymbol{x}(u) - \boldsymbol{x}(v)]^2. \end{aligned}$$

□

Corollary 1.11. *$\boldsymbol{L}$ is positive semi-definite.*[4]

[4] By the previous lemma, $\boldsymbol{L}$ satisfies the definition of positive semi-definiteness, which we will introduce in the following section.

Exercise 1.12. *A matrix $\boldsymbol{M}$ is a Laplacian matrix iff it satisfies the following conditions:*

1. $\boldsymbol{M}^\top = \boldsymbol{M}$;
2. *the diagonal entries of $\boldsymbol{M}$ are non-negative, and the off-diagonal entries of $\boldsymbol{M}$ are non-positive; and*
3. $\boldsymbol{M}\mathbf{1} = \mathbf{0}$.

It is often to useful to look at a normalized Laplacian matrix, where weighted vertex degrees are normalized to one and edges are weighted based on the degrees of their endpoints.

Definition 1.13 (Normalized Laplacian matrix)**.** The *normalized Laplacian matrix* of an oriented graph G is defined as,

$$\boldsymbol{N}(i,j) \doteq \begin{cases} 1 & \text{if } i = j \\ -\frac{1}{\sqrt{\boldsymbol{d}(i)\boldsymbol{d}(j)}} & \text{if } i \sim j \\ 0 & \text{otherwise.} \end{cases} \tag{1.12}$$

This characterization is equivalent to,[5]

[5] $\boldsymbol{A}^{-1/2}$, where $\boldsymbol{A}$ is a diagonal matrix, is the diagonal matrix $\mathrm{diag}_i\{\boldsymbol{A}(i,i)^{-1/2}\}$.

$$\boldsymbol{N} = \boldsymbol{D}^{-1/2}\boldsymbol{L}\boldsymbol{D}^{-1/2} = \boldsymbol{I} - \boldsymbol{D}^{-1/2}\boldsymbol{A}\boldsymbol{D}^{-1/2}. \tag{1.13}$$

As we will see later, the normalized Laplacian matrix is intimately related to the probability transition matrix of a random walk on G, where transition probabilities are proportional to edge weights.

Lemma 1.14. *$\boldsymbol{N}$ is positive semi-definite.*

Proof. We have for any $\boldsymbol{x} \in \mathbb{R}^n$,

$$\boldsymbol{x}^\top \boldsymbol{N}\boldsymbol{x} = \boldsymbol{x}^\top \boldsymbol{D}^{-1/2}\boldsymbol{L}\boldsymbol{D}^{-1/2}\boldsymbol{x} = (\boldsymbol{D}^{-1/2}\boldsymbol{x})^\top \boldsymbol{L}(\boldsymbol{D}^{-1/2}\boldsymbol{x}) \geq 0,$$

using positive semi-definiteness of $\boldsymbol{L}$. □

1.2 An Optimization Problem

We have seen that finding electrical voltages $\tilde{\boldsymbol{x}}$ or the electrical flow $\tilde{\boldsymbol{f}}$ is equivalent, we can go from one to the other and back. So let us first focus on how we can find electrical voltages $\tilde{\boldsymbol{x}}$.

In eq. (1.7), we saw that electrical voltages satisfy $\boldsymbol{L}\tilde{\boldsymbol{x}} = \boldsymbol{d}$, obeying by Ohm's law and satisfying the net flow constraint. A standard approach to reframe the solution to such a system of linear equations as the result of an optimization problem, is to consider the cost function,

$$c(\boldsymbol{x}) \doteq \frac{1}{2}\boldsymbol{x}^\top \boldsymbol{L}\boldsymbol{x} - \boldsymbol{x}^\top \boldsymbol{d}. \tag{1.14}$$

Observe that $\nabla c(\boldsymbol{x}) = \boldsymbol{L}\boldsymbol{x} - \boldsymbol{d} \stackrel{!}{=} 0$ iff $\boldsymbol{L}\boldsymbol{x} = \boldsymbol{d}$.

Claim 1.15. *c is convex, hence, its critical point coincides with its minimizer.*[6]

[6] We will develop tools to show this in the next part.

We have therefore recast the problem of finding electrical voltages to the convex optimization,

$$\tilde{\boldsymbol{x}} = \operatorname*{arg\,min}_{\boldsymbol{x}\in\mathbb{R}^{|V|}} c(\boldsymbol{x}). \tag{1.15}$$

1.3 *Energy & Duality*

Let us now look at the same problem through a different lens. Transporting current through a network of resistors requires energy, which is dissipated as heat by the resistor. By *Joule's law*, sending a current $\tilde{f}$ across a resistor with potential drop $\tilde{x}$, spends $\tilde{f}\cdot\tilde{x}$ units of energy per unit time.[7] Using Ohm's law, we have,

[7] We will think about everything as if happening in one unit of time.

$$\tilde{f}\cdot\tilde{x} = \frac{\tilde{x}^2}{r} = r\cdot\tilde{f}^2. \tag{1.16}$$

We can therefore write the *electrical energy* dissipated by routing a current $\tilde{\boldsymbol{f}}$ (or equivalently with electrical voltages $\tilde{\boldsymbol{x}}$) as,

$$\mathcal{E}(\tilde{\boldsymbol{f}}) \doteq \sum_{e\in E} \boldsymbol{r}(e)\tilde{\boldsymbol{f}}(e)^2 = \tilde{\boldsymbol{f}}^\top \boldsymbol{R}\tilde{\boldsymbol{f}} = \tilde{\boldsymbol{x}}^\top \boldsymbol{L}\tilde{\boldsymbol{x}} \doteq \mathcal{E}(\tilde{\boldsymbol{x}}). \tag{1.17}$$

using Ohm's law, $\tilde{\boldsymbol{f}} = \boldsymbol{R}^{-1}\boldsymbol{B}^\top\tilde{\boldsymbol{x}}$

Remark 1.16. Given electrical voltages $\tilde{\boldsymbol{x}} \perp \mathbf{1}$, we can also write the electrical energy as

$$\mathcal{E}(\tilde{\boldsymbol{x}}) = \tilde{\boldsymbol{x}}^\top \boldsymbol{L}\tilde{\boldsymbol{x}} = \boldsymbol{d}^\top\tilde{\boldsymbol{x}} = \boldsymbol{d}^\top \boldsymbol{L}^+\boldsymbol{d} \doteq \mathcal{E}(\boldsymbol{d}), \tag{1.18}$$

using $\boldsymbol{L}\boldsymbol{x} = \boldsymbol{d}$ and $\boldsymbol{L}^+\boldsymbol{L}\boldsymbol{x} = \boldsymbol{x}$ as $\boldsymbol{x} \perp \ker \boldsymbol{L}$

using the pseudoinverse $\boldsymbol{L}^+$ of $\boldsymbol{L}$.

Let us consider the *electrical energy-minimizing flow*:[8]

[8] Note that we could instead (and equivalently) characterize the optimization problem using electrical voltages.

$$\boldsymbol{f}^* \doteq \operatorname*{arg\,min}_{\substack{\boldsymbol{f}\in\mathbb{R}^{|E|}\\ \boldsymbol{B}\boldsymbol{f}=\boldsymbol{d}}} \mathcal{E}(\boldsymbol{f}). \tag{1.19}$$

Exercise 1.17. *$\boldsymbol{f}^*$ is precisely the electrical flow $\tilde{\boldsymbol{f}}$, that is, $\boldsymbol{f}^*$ satisfies Ohm's law.*

This indicates that the two above optimization problems are intimately related: both yield the electrical flow (or equivalently, electrical voltages). In fact, it can be shown that,

Exercise 1.18. $\mathcal{E}(\tilde{\boldsymbol{f}}) = -c(\tilde{\boldsymbol{x}})$,

where we think about maximizing $-c(\tilde{\boldsymbol{x}})$ instead of minimizing $c(\tilde{\boldsymbol{x}})$. More generally,

Exercise 1.19. $\mathcal{E}(\boldsymbol{f}) \geq -c(\boldsymbol{x})$ *for any flow* $\boldsymbol{f}$ *routing* $\boldsymbol{d}$ *and any voltages* $\boldsymbol{x}$.

So, for any voltages $\boldsymbol{x}$, the value of $-c(\boldsymbol{x})$ is a lower bound on the minimum electrical energy $\mathcal{E}(\tilde{\boldsymbol{f}})$.

This is an example of a much broader phenomenon known as Lagrangian duality, where we have a minimization problem and a related maximization problem that gives lower bounds on the optimal value of the minimization problem.

2

Linear Algebra

Claim 2.1. *If a square matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ is symmetric*[1]*, then $\boldsymbol{A}$ has n real* eigenvalues $\lambda_1, \ldots, \lambda_n$ *and* eigenvectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in \mathbb{R}^n$ *such that $\boldsymbol{A}\boldsymbol{v}_i = \lambda_i \boldsymbol{v}_i$ and the $\boldsymbol{v}_i$ are orthogonal*[2]*.*

[1] $\boldsymbol{A}$ is symmetric iff $\boldsymbol{A} = \boldsymbol{A}^\top$.

[2] that is, $\boldsymbol{v}_i^\top \boldsymbol{v}_j = 0$ for $i \neq j$

Definition 2.2 (Positive (semi-)definiteness)**.** Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be a symmetric matrix. We say $\boldsymbol{A}$ is

1. *positive definite* iff $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} > 0$ for any $\boldsymbol{x} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}$;
2. *positive semi-definite* iff $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} \geq 0$ for any $\boldsymbol{x} \in \mathbb{R}^n$;
3. if neither $\boldsymbol{A}$ nor $-\boldsymbol{A}$ is positive semi-definite, $\boldsymbol{A}$ is *indefinite*.

We denote by S^n the set of symmetric $n \times n$ matrices, by S^n_+ the set of such matrices that are positive semi-definite, and by S^n_{++} the set of such matrices that are positive definite.

Theorem 2.3. *Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be a symmetric matrix. Then,*

1. *$\boldsymbol{A}$ is positive definite iff all eigenvalues are positive; and*
2. *$\boldsymbol{A}$ is positive semi-definite iff all eigenvalues are non-negative.*

This theorem is a corollary of the Courant-Fischer theorem, which we will work towards now.

Fact 2.4 (Spectral theorem for symmetric matrices)**.** *For all symmetric matrices $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ there exist*

$$\boldsymbol{V} = \begin{bmatrix} \boldsymbol{v}_1 & \cdots & \boldsymbol{v}_n \end{bmatrix} \in \mathbb{R}^{n\times n}, \quad \textit{and} \quad \boldsymbol{\Lambda} = \operatorname{diag}\{\lambda_i\}_{i\in[n]} \in \mathbb{R}^{n\times n}, \tag{2.1}$$

where λ_i and $\boldsymbol{v}_i$ are the eigenvalues and corresponding (normalized) eigenvectors of $\boldsymbol{A}$, such that

1. *$\boldsymbol{A} = \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^\top = \sum_{i=1}^n \lambda_i \boldsymbol{v}_i \boldsymbol{v}_i^\top$; and*
2. *$\boldsymbol{V}^\top \boldsymbol{V} = \boldsymbol{I}$, i.e., the columns of $\boldsymbol{V}$ form an orthonormal basis of $\mathbb{R}^n$.*

Theorem 2.5 (Courant-Fischer min-max theorem)**.** *For symmetric*

matrices $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ *with eigenvalues* $\lambda_1 \leq \cdots \leq \lambda_n$,

$$\lambda_i = \min_{\substack{\text{subspace } W \subseteq \mathbb{R}^n \\ \dim(W)=i}} \max_{\substack{\boldsymbol{x}\in W \\ \boldsymbol{x}\neq \boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} \tag{2.2}$$

$$= \max_{\substack{\text{subspace } W \subseteq \mathbb{R}^n \\ \dim(W)=n-i+1}} \min_{\substack{\boldsymbol{x}\in W \\ \boldsymbol{x}\neq \boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}}. \tag{2.3}$$

Proof. We show eq. (2.2). The proof of the other equation proceeds analogously.

- "$\geq$": We choose $W = \operatorname{span}\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_i\}$ We can write $\boldsymbol{x}$ in the basis of eigenvectors,

  $$\boldsymbol{x} = \sum_{j=1}^{i} \boldsymbol{c}(j)\boldsymbol{v}_j$$

  for some $\boldsymbol{c} \in \mathbb{R}^i$. We have,

  $$\boldsymbol{x}^\top \boldsymbol{x} = \|\boldsymbol{x}\|_2^2 = \sum_{j=1}^{i}\sum_{k=1}^{i} \boldsymbol{c}(j)\boldsymbol{c}(k)\boldsymbol{v}_j^\top \boldsymbol{v}_k = \sum_{j=1}^{i} \boldsymbol{c}(j)^2,$$

  using that $\boldsymbol{v}_j^\top \boldsymbol{v}_k = 0$ if $j \neq k$ and $\boldsymbol{v}_j^\top \boldsymbol{v}_k = 1$ otherwise

  and,

  $$\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} = \boldsymbol{x}^\top \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^\top \boldsymbol{x} = (\boldsymbol{V}^\top \boldsymbol{x})^\top \boldsymbol{\Lambda}(\underbrace{\boldsymbol{V}^\top \boldsymbol{x}}_{\boldsymbol{c}})$$

  $$= \boldsymbol{c}^\top \boldsymbol{\Lambda}\boldsymbol{c} = \sum_{j=1}^{i} \lambda_j \boldsymbol{c}(j)^2 \leq \lambda_i \sum_{j=1}^{i} \boldsymbol{c}(j)^2.$$

  Altogether,

  $$\frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} \leq \lambda_i.$$

- "$\leq$": Consider any subspace $W \subseteq \mathbb{R}^n$ with $\dim(W) = i$ and fix the subspace $T \doteq \operatorname{span}\{\boldsymbol{v}_i, \ldots, \boldsymbol{v}_n\}$ with $\dim(T) = n - i + 1$. We have that $\dim(W \cap T) = \dim(W) + \dim(T) - \dim(W \cup T)$ and $\dim(W \cup T) \leq \dim(\mathbb{R}^n) = n$, so, $\dim(W \cap T) \geq 1$. Therefore,

  $$\max_{\substack{\boldsymbol{x}\in W \\ \boldsymbol{x}\neq \boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} \geq \max_{\substack{\boldsymbol{x}\in W\cap T \\ \boldsymbol{x}\neq \boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}}$$

  $$\geq \min_{\substack{\text{subspace } V \subseteq T \\ \dim(V)=1}} \max_{\substack{\boldsymbol{x}\in V \\ \boldsymbol{x}\neq \boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}}.$$

  For the last inequality note that V can be chosen as $W \cap T$.
  We choose $V \doteq \operatorname{span}\{\boldsymbol{v}_i\}$. For some $c \in \mathbb{R}$, we can write $\boldsymbol{x} = c\boldsymbol{v}_i$. Similarly to the previous part, we obtain,

  $$\frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} = \frac{\lambda_i c^2}{c^2} = \lambda_i.$$

□

Proof of theorem 2.3. Using Courant-Fischer, we have for the smallest eigenvalue λ_1 of $\boldsymbol{A}$,

$$\lambda_1 = \min_{\substack{\boldsymbol{x}\in\mathbb{R}^n \\ \boldsymbol{x}\neq\boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}}.$$

Thus, if λ_1 is positive, then $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} > 0$ for all $\boldsymbol{x} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}$. In contrast, if for any such $\boldsymbol{x}$, $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} > 0$, then λ_1 must be positive. The proof of positive semi-definiteness is analogous. □

Corollary 2.6. *For a symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ and any $\boldsymbol{x} \in \mathbb{R}^n$,*

$$\lambda_{\min}(\boldsymbol{A}) \|\boldsymbol{x}\|_2^2 \leq \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}, \tag{2.4}$$

where $\lambda_{\min}(\boldsymbol{A})$ is the smallest eigenvalue of $\boldsymbol{A}$.

Proof. By Courant-Fischer, we have for any $\boldsymbol{x} \in \mathbb{R}^n$ such that $\boldsymbol{x} \neq \boldsymbol{0}$,

$$\lambda_{\min}(\boldsymbol{A}) = \min_{\substack{\boldsymbol{y}\in\mathbb{R}^n \\ \boldsymbol{y}\neq\boldsymbol{0}}} \frac{\boldsymbol{y}^\top \boldsymbol{A}\boldsymbol{y}}{\|\boldsymbol{y}\|_2^2} \leq \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\|\boldsymbol{x}\|_2^2}.$$

If $\boldsymbol{x} = \boldsymbol{0}$, the inequality trivially holds. □

Claim 2.7. *For any matrix $\boldsymbol{M}$ and invertible matrix $\boldsymbol{T}$, $\boldsymbol{M}$ and $\boldsymbol{T}\boldsymbol{M}\boldsymbol{T}^{-1}$ have the same eigenvalues.*

2.1 *Square Roots*

We now state a useful fact for positive semi-definite matrices.

Lemma 2.8. *Any symmetric and positive semi-definite matrix $\boldsymbol{A}$ has a positive semi-definite square root $\boldsymbol{A}^{1/2}$ such that $\boldsymbol{A}^{1/2}\boldsymbol{A}^{1/2} = \boldsymbol{A}$.*

Proof. By the spectral theorem, $\boldsymbol{A} = \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^\top$, where $\boldsymbol{V}$ is an orthonormal matrix of eigenvectors and $\boldsymbol{\Lambda}$ is a diagonal matrix of eigenvalues. Let $\boldsymbol{A}^{1/2} \doteq \boldsymbol{V}\boldsymbol{\Lambda}^{1/2}\boldsymbol{V}^\top$, where $\boldsymbol{\Lambda}^{1/2} = \mathrm{diag}_i\{\boldsymbol{\Lambda}(i,i)^{1/2}\}$. Then,

$$\begin{aligned}
\boldsymbol{A}^{1/2}\boldsymbol{A}^{1/2} &= \boldsymbol{V}\boldsymbol{\Lambda}^{1/2}\boldsymbol{V}^\top\boldsymbol{V}\boldsymbol{\Lambda}^{1/2}\boldsymbol{V}^\top \\
&= \boldsymbol{V}\boldsymbol{\Lambda}^{1/2}\boldsymbol{\Lambda}^{1/2}\boldsymbol{V}^\top \\
&= \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^\top = \boldsymbol{A}.
\end{aligned}$$

It is immediately clear from the definition that $\boldsymbol{A}^{1/2}$ is positive semi-definite. □

2.2 *Matrix Norms*

Definition 2.9 (Matrix norm)**.** Given a matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ and norms $\|\cdot\|_\alpha$ and $\|\cdot\|_\beta$ on $\mathbb{R}^n$, the *(induced) norm* of $\boldsymbol{A}$ is defined as,[3]

$$\|\boldsymbol{A}\|_{\alpha\to\beta} \doteq \sup_{\substack{\boldsymbol{x}\in\mathbb{R}^n \\ \boldsymbol{x}\neq\boldsymbol{0}}} \frac{\|\boldsymbol{A}\boldsymbol{x}\|_\beta}{\|\boldsymbol{x}\|_\alpha}. \tag{2.5}$$

[3] If you think of $\boldsymbol{A}$ as a linear map, you can think of $\|\cdot\|_\alpha$ as a norm of the input space and $\|\cdot\|_\beta$ as a norm of the output space.

We write $\|\boldsymbol{A}\|_\alpha \doteq \|\boldsymbol{A}\|_{\alpha\to\alpha}$.

Lemma 2.10. *For any matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$, any $\boldsymbol{x} \in \mathbb{R}^n$, and any norm $\|\cdot\|$ on $\mathbb{R}^n$,*

$$|\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}| \leq \|\boldsymbol{A}\| \, \|\boldsymbol{x}\|^2. \tag{2.6}$$

Proof. We have,

$$\begin{aligned} |\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}| &\leq \|\boldsymbol{x}\| \, \|\boldsymbol{A}\boldsymbol{x}\| && \text{using Cauchy-Schwarz} \\ &\leq \|\boldsymbol{A}\| \, \|\boldsymbol{x}\|^2. && \text{using the definition of the induced matrix norm} \end{aligned}$$

□

Lemma 2.11. *For a symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$,*

$$\|\boldsymbol{A}\|_2 = \max\{|\lambda_{\min}(\boldsymbol{A})|, |\lambda_{\max}(\boldsymbol{A})|\}. \tag{2.7}$$

$\|\boldsymbol{A}\|_2$ is called the spectral norm *of $\boldsymbol{A}$.*

Proof. We have,

$$\begin{aligned} \|\boldsymbol{A}\|_2^2 &= \sup_{\substack{\boldsymbol{x}\in\mathbb{R}^n \\ \boldsymbol{x}\neq\boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} \\ &= \sup_{\substack{\boldsymbol{x}\in\mathbb{R}^n \\ \boldsymbol{x}\neq\boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}^2 \boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} && \text{using that } \boldsymbol{A} \text{ is symmetric, } \boldsymbol{A}^\top = \boldsymbol{A} \\ &= \sup_{\substack{\boldsymbol{x}\in\mathbb{R}^n \\ \boldsymbol{x}\neq\boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{V}\boldsymbol{\Lambda}^2\boldsymbol{V}^\top \boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} && \text{using that } \boldsymbol{V} \text{ is orthogonal, } \boldsymbol{V}^\top = \boldsymbol{V}^{-1} \\ &= \sup_{\substack{\boldsymbol{x}\in\mathbb{R}^n \\ \boldsymbol{x}\neq\boldsymbol{0}}} \frac{\boldsymbol{V}^\top \boldsymbol{x}^\top \boldsymbol{\Lambda}^2 (\boldsymbol{V}^\top \boldsymbol{x})}{\boldsymbol{V}^\top \boldsymbol{x}^\top (\boldsymbol{V}^\top \boldsymbol{x})} \\ &= \sup_{\substack{\boldsymbol{y}\in\mathbb{R}^n \\ \boldsymbol{y}\neq\boldsymbol{0}}} \frac{\boldsymbol{y}^\top \boldsymbol{\Lambda}^2 \boldsymbol{y}}{\boldsymbol{y}^\top \boldsymbol{y}} = \|\boldsymbol{\Lambda}\|_2^2. && \text{using that the columns of } \boldsymbol{V} \text{ form a basis of } \mathbb{R}^n, \text{ set } \boldsymbol{y} \doteq \boldsymbol{V}^\top \boldsymbol{x} \end{aligned}$$

Finally,

$$\|\boldsymbol{\Lambda}\|_2^2 = \sup_{\substack{\boldsymbol{x}\in\mathbb{R}^n \\ \boldsymbol{x}\neq\boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{\Lambda}^2 \boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} = \sup_{\substack{\boldsymbol{x}\in\mathbb{R}^n \\ \boldsymbol{x}\neq\boldsymbol{0}}} \frac{\sum_{i=1}^n \lambda_i^2 \boldsymbol{x}(i)^2}{\sum_{i=1}^n \boldsymbol{x}(i)^2} = \max_{i\in[n]} \lambda_i^2. \qquad \square$$

2.3 *Loewner Order*

The Loewner order (or positive semi-definite order) is a partial ordering on symmetric matrices.

Definition 2.12 (Loewner order)**.** Given symmetric matrices $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{n\times n}$, $\boldsymbol{A} \preceq \boldsymbol{B}$ iff $\forall \boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{x}^\top \boldsymbol{B}\boldsymbol{x}$.

Remark 2.13. $\boldsymbol{A} \succeq \boldsymbol{0}$ iff $\boldsymbol{A}$ is positive semi-definite.

Lemma 2.14 (Properties of the Loewner order)**.** *We have that for any symmetric $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C} \in \mathbb{R}^{n\times n}$,*

1. $\boldsymbol{A} \preceq \boldsymbol{A}$ *(reflexivity);*
2. $\boldsymbol{A} \preceq \boldsymbol{B}, \boldsymbol{B} \preceq \boldsymbol{A} \implies \boldsymbol{A} = \boldsymbol{B}$ *(antisymmetry);*
3. $\boldsymbol{A} \preceq \boldsymbol{B}, \boldsymbol{B} \preceq \boldsymbol{C} \implies \boldsymbol{A} \preceq \boldsymbol{C}$ *(transitivity);*
4. $\boldsymbol{A} \preceq \boldsymbol{B} \implies \boldsymbol{A} + \boldsymbol{C} \preceq \boldsymbol{B} + \boldsymbol{C}$*;*
5. *if $\boldsymbol{A} \succeq \boldsymbol{0}$, then $\frac{1}{\alpha}\boldsymbol{A} \preceq \boldsymbol{A} \preceq \alpha\boldsymbol{A}$ for any $\alpha \geq 1$; and*
6. *if $\boldsymbol{A} \preceq \boldsymbol{B}$, then $\forall i \in [n] : \lambda_i(\boldsymbol{A}) \leq \lambda_i(\boldsymbol{B})$, where $\lambda_i(\boldsymbol{M})$ is the i-th largest eigenvalue of $\boldsymbol{M}$.*[4]

[4] The converse is false:

$$\boldsymbol{A} \doteq \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}, \quad \boldsymbol{B} \doteq \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$$

have equal eigenvalues, but $\boldsymbol{A} \not\preceq \boldsymbol{B}$ and $\boldsymbol{B} \not\preceq \boldsymbol{A}$.

Remark 2.15. Properties (1), (2), and (3) together imply that the Loewner order is a partial ordering of symmetric matrices.

Proof. Properties (1) through (5) follow directly from the definition using only elementary operations. For property (6), we have by Courant-Fischer,

$$\begin{aligned}\lambda_i(\boldsymbol{A}) &= \min_{\substack{\text{subspace } W \subseteq \mathbb{R}^n \\ \dim(W)=i}} \max_{\substack{\boldsymbol{x}\in W \\ \boldsymbol{x}\neq \boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} \\ &\leq \min_{\substack{\text{subspace } W \subseteq \mathbb{R}^n \\ \dim(W)=i}} \max_{\substack{\boldsymbol{x}\in W \\ \boldsymbol{x}\neq \boldsymbol{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{B}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{x}} = \lambda_i(\boldsymbol{B}). \qquad \square\end{aligned}$$

Loewner Order on Graphs

Definition 2.16. We write $G \preceq H$ iff we have for the Laplacian matrices $\boldsymbol{L}_G$ and $\boldsymbol{L}_H$ that $\boldsymbol{L}_G \preceq \boldsymbol{L}_H$. For any $c > 0$, we write cG in place of $c\boldsymbol{L}_G$, which corresponds to scaling the weight of every edge of G by c.

Lemma 2.17. *For subgraphs $H \subseteq G$, we have $H \preceq G$.*

Proof. Dropping edges can only decrease the quadratic form of the Laplacian of G (see eq. (1.11)). $\square$

2.4 Kernel and Image

Definition 2.18 (Kernel)**.** The *kernel* (or *null space*) of a linear map $\boldsymbol{A} \in \mathbb{R}^{n\times m}$ is the linear subspace,

$$\ker \boldsymbol{A} \doteq \{\boldsymbol{x} \in \mathbb{R}^m \mid \boldsymbol{A}\boldsymbol{x} = \boldsymbol{0}\} \subseteq \mathbb{R}^m. \tag{2.8}$$

Definition 2.19 (Image)**.** The *image* (or *range*) of a linear map $\boldsymbol{A} \in \mathbb{R}^{n\times m}$ is the linear subspace,

$$\operatorname{im} \boldsymbol{A} \doteq \{\boldsymbol{A}\boldsymbol{x} \mid \boldsymbol{x} \in \mathbb{R}^m\} = \operatorname{span}\{\boldsymbol{A}(:,1), \ldots, \boldsymbol{A}(:,m)\} \subseteq \mathbb{R}^n, \tag{2.9}$$

where $\boldsymbol{A}(:,i)$ denotes the i-th column vector of $\boldsymbol{A}$.

We have the following useful property relating image and kernel.

Lemma 2.20. *For a matrix $\boldsymbol{A} \in \mathbb{R}^{n\times m}$, we have $\operatorname{im}(\boldsymbol{A}^\top)^\perp = \ker \boldsymbol{A}$ and $(\operatorname{im} \boldsymbol{A})^\perp = \ker(\boldsymbol{A}^\top)$.*[5]

[5] Note that $(W^\perp)^\perp = W$ for any subspace W.

Proof. If $\boldsymbol{x} \in \operatorname{im}(\boldsymbol{A}^\top)^\perp$, we have $\boldsymbol{A}(i,:)^\top \boldsymbol{x} = 0$ for all $i \in [n]$. Hence, $\boldsymbol{A}\boldsymbol{x} = 0$ and $\boldsymbol{x} \in \ker \boldsymbol{A}$.

Conversely, if $\boldsymbol{x} \in \ker \boldsymbol{A}$, then $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{0}$. Therefore, $\boldsymbol{A}(i,:)^\top \boldsymbol{x} = 0$ for all $i \in [n]$. As $\operatorname{im}(\boldsymbol{A}^\top) = \operatorname{span}\{\boldsymbol{A}(1,:), \ldots, \boldsymbol{A}(n,:)\}$, we have $\boldsymbol{z}^\top \boldsymbol{x} = 0$ for any $\boldsymbol{z} \in \operatorname{im}(\boldsymbol{A}^\top)$, yielding, $\boldsymbol{x} \in \operatorname{im}(\boldsymbol{A}^\top)^\perp$.

The same argument goes through for the transpose of $\boldsymbol{A}$ on both sides. □

2.5 Matrix Functions

First, let us remind ourselves of the notion of the trace of a matrix.

Definition 2.21 (Trace)**.** The *trace* of a symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ is defined as,

$$\operatorname{tr} \boldsymbol{A} \doteq \sum_{i=1}^{n} \boldsymbol{A}(i,i). \tag{2.10}$$

Lemma 2.22. *The trace is invariant under cyclic permutations. That is, for any $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{n\times n}$,*

$$\operatorname{tr} \boldsymbol{A}\boldsymbol{B} = \operatorname{tr} \boldsymbol{B}\boldsymbol{A}. \tag{2.11}$$

Proof. We have,

$$\operatorname{tr} \boldsymbol{A}\boldsymbol{B} = \sum_{i=1}^{n} (\boldsymbol{A}\boldsymbol{B})(i,i)$$

$$
\begin{aligned}
&= \sum_{i=1}^{n} \sum_{j=1}^{n} \boldsymbol{A}(i,j)\boldsymbol{B}(j,i) \\
&= \sum_{j=1}^{n} \sum_{i=1}^{n} \boldsymbol{B}(j,i)\boldsymbol{A}(i,j) \\
&= \sum_{j=1}^{n} (\boldsymbol{BA})(j,j) \\
&= \operatorname{tr} \boldsymbol{BA}. \qquad \square
\end{aligned}
$$

Lemma 2.23. *For $\boldsymbol{A} \in \mathbb{R}^{n \times n}$,*

$$\operatorname{tr} \boldsymbol{A} = \sum_{i=1}^{n} \lambda_i, \tag{2.12}$$

where λ_i are the eigenvalues of $\boldsymbol{A}$.

Proof. By the spectral theorem for symmetric matrices and using the cycle property of the trace,

$$\operatorname{tr} \boldsymbol{A} = \operatorname{tr} (\boldsymbol{V\Lambda V}^\top) = \operatorname{tr} (\boldsymbol{\Lambda} \underbrace{\boldsymbol{V}^\top \boldsymbol{V}}_{\boldsymbol{I}}) = \operatorname{tr} \boldsymbol{\Lambda} = \sum_{i=1}^{n} \lambda_i,$$

where $\boldsymbol{\Lambda}$ is a diagonal matrix of eigenvalues of $\boldsymbol{A}$ and $\boldsymbol{V}$ is an orthonormal matrix of the corresponding eigenvectors. $\square$

Definition 2.24 (Matrix function)**.** A *matrix function* $f : S^n \to S^n$ given the scalar function $f : \mathbb{R} \to \mathbb{R}$ is defined as,

$$f(\boldsymbol{A}) \doteq \boldsymbol{V} \operatorname{diag}_i \{f(\lambda_i)\} \boldsymbol{V}^\top, \tag{2.13}$$

where $\boldsymbol{A} = \boldsymbol{V} \operatorname{diag}_i \{\lambda_i\} \boldsymbol{V}^\top$ is the spectral decomposition of $\boldsymbol{A}$. We say,

1. a function $f : \mathcal{S} \to \mathbb{R}$ for $\mathcal{S} \subseteq S^n$ is *monotonically increasing* iff $\boldsymbol{A} \preceq \boldsymbol{B}$ implies $f(\boldsymbol{A}) \leq f(\boldsymbol{B})$; and similarly
2. a matrix function $f : \mathcal{S} \to \mathcal{T}$ for $\mathcal{S}, \mathcal{T} \subseteq S^n$ is *monotonically increasing* iff $\boldsymbol{A} \preceq \boldsymbol{B}$ implies $f(\boldsymbol{A}) \preceq f(\boldsymbol{B})$.

Lemma 2.25. *If $f : T \to \mathbb{R}$ for $T \subseteq \mathbb{R}$ is monotone, then $\boldsymbol{X} \mapsto \operatorname{tr} f(\boldsymbol{X})$ is monotone.*

Proof. Let $\boldsymbol{A} \preceq \boldsymbol{B}$. Then, $\lambda_i(\boldsymbol{A}) \leq \lambda_i(\boldsymbol{B})$ for all $i \in [n]$. Thus, $f(\lambda_i(\boldsymbol{A})) \leq f(\lambda_i(\boldsymbol{B}))$, and hence,

$$\operatorname{tr} f(\boldsymbol{A}) = \sum_{i=1}^{n} f(\lambda_i(\boldsymbol{A})) \leq \sum_{i=1}^{n} f(\lambda_i(\boldsymbol{B})) = \operatorname{tr} f(\boldsymbol{B}).$$

The analogous argument works when f is monotonically decreasing. $\square$

Claim 2.26 (Facts about matrix functions)**.**

1. $\boldsymbol{X} \mapsto \boldsymbol{X}^{-1}$ *is monotonically decreasing on* S^n_{++}*;*
2. $\log$ *is monotonically increasing on* S^n_{++}*;*
3. $\exp(\boldsymbol{A}) \preceq \boldsymbol{I} + \boldsymbol{A} + \boldsymbol{A}^2$ *for* $\|\boldsymbol{A}\|_2 \leq 1$*;*
4. $\log(\boldsymbol{I} + \boldsymbol{A}) \preceq \boldsymbol{A}$ *for* $\boldsymbol{A} \succeq -\boldsymbol{I}$*; and*
5. *(*Lieb's theorem*) if* $f(\boldsymbol{A}) = \operatorname{tr}(\exp(\boldsymbol{H} + \log(\boldsymbol{A})))$ *for* $\boldsymbol{A} \in S^n_{++}$ *and some* $\boldsymbol{H} \in S^n$*, then* $-f$ *is convex.*

2.6 *Pseudoinverses*

We often want to solve systems of linear equations such as

$$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} \tag{2.14}$$

where $\boldsymbol{A}$ and $\boldsymbol{b}$ are given and we seek to identify $\boldsymbol{x}$. If $\boldsymbol{A}$ is invertible, a solution to eq. (2.14) is given by $\boldsymbol{x} \doteq \boldsymbol{A}^{-1}\boldsymbol{b}$. Recall that a square matrix $\boldsymbol{A}$ is invertible iff $\det \boldsymbol{A} \neq 0$. Can we still explicitly write $\boldsymbol{x}$ when $\boldsymbol{A}$ is not invertible?

When $\boldsymbol{A}$ is not invertible, the corresponding system of linear equations does not have a unique solution. However, we may still be able to find *some* solution. The Moore-Penrose inverse is such a generalization of the inverse.

Definition 2.27 (Moore-Penrose inverse)**.** Given a symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$,[6] its *Moore-Penrose inverse* (or simply *pseudoinverse*) is a matrix $\boldsymbol{A}^+ \in \mathbb{R}^{n\times n}$ such that

1. $\boldsymbol{A}^+$ is symmetric;
2. $\ker \boldsymbol{A}^+ = \ker \boldsymbol{A}$;[7] and
3. for $\boldsymbol{v} \perp \ker \boldsymbol{A}$, $\boldsymbol{A}^+\boldsymbol{A}\boldsymbol{v} = \boldsymbol{v}$ and $\boldsymbol{A}\boldsymbol{A}^+\boldsymbol{v} = \boldsymbol{v}$.

Thus, provided $\boldsymbol{b} \perp \ker \boldsymbol{A}$, $\boldsymbol{x} \doteq \boldsymbol{A}^+\boldsymbol{b}$ is a solution to eq. (2.14).

[6] The Moore-Penrose inverse can also be defined for non-symmetric and even non-square matrices, but this will not be important for us.

[7] that is, for any $\boldsymbol{v} \in \mathbb{R}^n$,

$$\boldsymbol{A}^+\boldsymbol{v} = \boldsymbol{0} \iff \boldsymbol{A}\boldsymbol{v} = \boldsymbol{0}$$

Lemma 2.28. *The pseudoinverse of the symmetric matrix* $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ *is (uniquely*[8]*) characterized as,*

$$\boldsymbol{A}^+ = \boldsymbol{V}\boldsymbol{\Lambda}^+\boldsymbol{V}^\top, \tag{2.15}$$

where $\boldsymbol{V}$ *is the matrix of orthogonal eigenvectors and*

$$\boldsymbol{\Lambda}^+ = \operatorname{diag}_{i\in[n]} \begin{cases} \lambda_i^{-1} & \lambda_i \neq 0 \\ 0 & \textit{otherwise} \end{cases} \tag{2.16}$$

for eigenvalues λ_i*.*

[8] this we will not prove

Proof. The properties of a pseudoinverse follow from the definition immediately. □

Corollary 2.29. $\boldsymbol{A} \succeq \boldsymbol{B}$ *implies* $\boldsymbol{A}^+ \preceq \boldsymbol{B}^+$ *when* $\ker \boldsymbol{A} = \ker \boldsymbol{B}$*.*

2.7 Solving Linear Systems

Linear systems such as eq. (2.14) can often be solved much faster than by computing the pseudoinverse of $\boldsymbol{A}$.

We first need to remind ourselves of the notion of a projection.

Definition 2.30 (Projection matrix)**.** A *projection matrix* is a matrix $\boldsymbol{\Pi} \in \mathbb{R}^{n \times n}$ such that $\boldsymbol{\Pi}^2 = \boldsymbol{\Pi}$.[9] We say that $\boldsymbol{\Pi}$ is *orthogonal* iff $\operatorname{im} \boldsymbol{\Pi} \perp \ker \boldsymbol{\Pi}$.

[9] This property is called *idempotency*.

Remark 2.31. Note that $\boldsymbol{\Pi}$ is the identity operator on $\operatorname{im} \boldsymbol{\Pi}$, i.e., for all $\boldsymbol{x} \in \operatorname{im} \boldsymbol{\Pi}$ we have that $\boldsymbol{\Pi}\boldsymbol{x} = \boldsymbol{x}$.

Definition 2.32 (Orthogonal projection to the complement of the kernel)**.** Given a matrix $\boldsymbol{A}$, $\boldsymbol{\Pi}_{\boldsymbol{A}}$ is the *orthogonal projection to* $(\ker \boldsymbol{A})^{\perp}$, that is,

1. $\boldsymbol{\Pi}_{\boldsymbol{A}} \boldsymbol{v} = \boldsymbol{0}$ if $\boldsymbol{v} \in \ker \boldsymbol{A}$; and
2. $\boldsymbol{\Pi}_{\boldsymbol{A}} \boldsymbol{v} = \boldsymbol{v}$ if $\boldsymbol{v} \in (\ker \boldsymbol{A})^{\perp}$.[10]

[10] Recall that $(\ker \boldsymbol{A})^{\perp} = \operatorname{im} \boldsymbol{A}^{\top}$.

Claim 2.33. *For $\boldsymbol{A} \in S^n$ with spectral decomposition $\boldsymbol{A} = \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^{\top} = \sum_i \lambda_i \boldsymbol{v}_i \boldsymbol{v}_i^{\top}$ where $\boldsymbol{V}$ is orthonormal, we have,*

$$\boldsymbol{\Pi}_{\boldsymbol{A}} = \sum_{\substack{i \\ \lambda_i \neq 0}} \boldsymbol{v}_i \boldsymbol{v}_i^{\top} = \boldsymbol{A}^{+/2} \boldsymbol{A} \boldsymbol{A}^{+/2} = \boldsymbol{A}\boldsymbol{A}^{+} = \boldsymbol{A}^{+}\boldsymbol{A}. \tag{2.17}$$

Corollary 2.34. *In particular, $\boldsymbol{\Pi}_{\boldsymbol{A}}^{+} = \boldsymbol{\Pi}_{\boldsymbol{A}}$.*

Claim 2.35. *Consider a real symmetric matrix $\boldsymbol{A} = \boldsymbol{X}\boldsymbol{Y}\boldsymbol{X}^{\top}$, where $\boldsymbol{X}$ is real and invertible and $\boldsymbol{Y}$ is real and symmetric. Then,*

$$\boldsymbol{A}^{+} = \boldsymbol{\Pi}_{\boldsymbol{A}} (\boldsymbol{X}^{\top})^{-1} \boldsymbol{Y}^{+} \boldsymbol{X}^{-1} \boldsymbol{\Pi}_{\boldsymbol{A}}. \tag{2.18}$$

Now, let us return to solving linear systems.

Lemma 2.36. *Given an invertible square lower triangular matrix $\mathcal{L}$ or an invertible square upper triangular matrix $\mathcal{U}$, we can solve the linear systems $\mathcal{L}\boldsymbol{y} = \boldsymbol{b}$ and $\mathcal{U}\boldsymbol{z} = \boldsymbol{b}$ in time $\mathcal{O}(\operatorname{nnz} \mathcal{L})$ and $\mathcal{O}(\operatorname{nnz} \mathcal{U})$, respectively, where $\operatorname{nnz} \boldsymbol{A}$ denotes the number of non-zero entries of $\boldsymbol{A}$.*[11]

[11] We need the matrices to be stored as an adjacency list to have fast access to their non-zero entries.

Proof sketch. We can iteratively solve the linear equations of $\mathcal{L}\boldsymbol{y} = \boldsymbol{b}$. As $\mathcal{L}$ has full rank and is lower triangular, there is always one linear equation in a single variable.

These algorithms are known as forward and back substitution for lower and upper triangular matrices, respectively. □

Thus, if we can decompose $\boldsymbol{A} = \mathcal{L}\mathcal{L}^\top$ such that $\mathcal{L}$ is invertible and lower triangular, we can solve the linear system of eq. (2.14) in time $\mathcal{O}(\text{nnz}\,\mathcal{L})$ by solving the two linear systems,

$$\mathcal{L}\boldsymbol{y} = \boldsymbol{b} \quad \text{and} \tag{2.19}$$

$$\mathcal{L}^\top \boldsymbol{x} = \boldsymbol{y}. \tag{2.20}$$

Fact 2.37 (Cholesky decomposition)**.** *For any positive semi-definite matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$, there is a decomposition of the form $\boldsymbol{A} = \mathcal{L}\mathcal{L}^\top$ where $\mathcal{L} \in \mathbb{R}^{n\times n}$ is lower triangular and positive semi-definite.*

We will see that the Cholesky decomposition can be computed efficiently when $\boldsymbol{A}$ is a graph Laplacian.

3

Probability

3.1 Random Walks

In this section, we study random walks on undirected weighted graphs $G = (V, E, \boldsymbol{w})$ with self-loops. A *random walk* visits a random sequence of vertices $X_1, X_2, \ldots$, where

$$\mathbb{P}[X_{t+1} = v \mid X_t = u] = \frac{\boldsymbol{w}(\{v,u\})}{\boldsymbol{d}(u)} \tag{3.1}$$

and $\boldsymbol{d}(u)$ is the weighted degree of vertex u, as we have defined previously.

Remark 3.1. The random walks considered here satisfy the *Markov property*, that is,

$$X_{t+1} \perp X_1, \ldots, X_{t-1} \mid X_t. \tag{3.2}$$

[1] So you can think of these random walks as Markov chains.

Moreover, we restrict our attention to time-homogeneous random walks, that is, the transition probabilities remain constant over time.

We can therefore model the update of a single round using the linear map $\boldsymbol{W} \in \mathbb{R}^{|V| \times |V|}$ (called *transition matrix*),

$$\boldsymbol{W} \doteq \begin{bmatrix} \frac{\boldsymbol{w}(\{1,1\})}{\boldsymbol{d}(1)} & \cdots & \frac{\boldsymbol{w}(\{n,1\})}{\boldsymbol{d}(n)} \\ \vdots & \ddots & \vdots \\ \frac{\boldsymbol{w}(\{1,n\})}{\boldsymbol{d}(1)} & \cdots & \frac{\boldsymbol{w}(\{n,n\})}{\boldsymbol{d}(n)} \end{bmatrix} = \boldsymbol{A}\boldsymbol{D}^{-1}. \tag{3.3}$$

A probability distribution over vertices is a vector $\boldsymbol{p} \in \mathbb{R}^{|V|}$ such that $\mathbf{1}^\top \boldsymbol{p} = 1$ and $\boldsymbol{p} \geq \mathbf{0}$. If our initial distribution is $\boldsymbol{p}_0$, we have,

$$\boldsymbol{p}_t = \boldsymbol{W}^t \boldsymbol{p}_0. \tag{3.4}$$

Observe that $\boldsymbol{W}^t(u, v)$ denotes the probability to reach v from u in exactly t steps.

Definition 3.2 (Mixing)**.** We say that a random walk $\boldsymbol{W}$ is *mixing* at step t iff for each $u, v \in V$, $\boldsymbol{W}^t(u,v) \geq \frac{1}{2n}$.[2]

[2] That is, the random walk is "half way" to being completely mixed.

Definition 3.3 (Stationary distribution)**.** A distribution $\boldsymbol{\pi} \in \mathbb{R}^{|V|}$ is *stationary* iff $\boldsymbol{\pi} = \boldsymbol{W}\boldsymbol{\pi}$.

Lemma 3.4. *Every graph has the stationary distribution* $\boldsymbol{\pi} \doteq \frac{\boldsymbol{d}}{\mathbf{1}^\top \boldsymbol{d}}$.

Proof. First, π is a distribution as,

$$\mathbf{1}^\top \boldsymbol{\pi} = \frac{\mathbf{1}^\top \boldsymbol{d}}{\mathbf{1}^\top \boldsymbol{d}}$$

and clearly $\boldsymbol{\pi} \geq \mathbf{0}$. We have,

$$\boldsymbol{W}\boldsymbol{\pi} = \frac{1}{\mathbf{1}^\top \boldsymbol{d}} \boldsymbol{A}\boldsymbol{D}^{-1}\boldsymbol{d} = \frac{1}{\mathbf{1}^\top \boldsymbol{d}} \boldsymbol{A}\mathbf{1} = \frac{\boldsymbol{d}}{\mathbf{1}^\top \boldsymbol{d}} = \boldsymbol{\pi}. \qquad \square$$

Remark 3.5. When the graph is connected, this is the unique stationary distribution.[3]

[3] In the context of Markov chains, irreducibility is sufficient for a unique stationary condition and equivalent to the transition graph being connected.

Lazy Random Walk

We would also like to have that we converge to this stationary distribution regardless of the initial distribution $\boldsymbol{p}_0$, but this is not true for general graphs, as is shown in fig. 3.1. A sufficient condition for convergence to the stationary distribution is, however, that all vertices have self-loops.[4]

Figure 3.1: Consider the initial distribution $\boldsymbol{p}_0(1) = 1, \boldsymbol{p}_0(2) = 0$. Clearly, the random walk will forever oscillate between the two states.

[4] It is easy to check that this ensures that the Markov chain is aperiodic, which together with irreducibility implies convergence to the unique stationary distribution.

Given the random walk $\boldsymbol{W}$, the associated *lazy random walk* is given by,

$$\tilde{\boldsymbol{W}} \doteq \frac{1}{2}\boldsymbol{I} + \frac{1}{2}\boldsymbol{W},$$

that is, we add self-loops to each vertex with weight $1/2$ and halve all other weights. Observe that this does not change the stationary distribution of the random walk. This ensures that the following holds.

Theorem 3.6 (Convergence of lazy random walk)**.** *For a connected graph, the lazy random walk converges to its unique stationary distribution irrespectively of the initial distribution* $\boldsymbol{p}_0$*,*

$$\lim_{t\to\infty} \tilde{\boldsymbol{W}}^t \boldsymbol{p}_0 = \tilde{\boldsymbol{\pi}} = \boldsymbol{\pi}. \tag{3.5}$$

To prove this theorem, let us first understand the transition matrix in terms of the graph Laplacian.

Lemma 3.7. *When $\nu_1, \dots, \nu_n$ are the eigenvalues and $\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_n$ the corresponding eigenvectors of the normalized Laplacian matrix $\boldsymbol{N}$, then $\tilde{\boldsymbol{W}}$ has eigenvalues $1 - \nu_i/2$ and (not necessarily orthogonal) eigenvectors $\boldsymbol{D}^{1/2}\boldsymbol{\psi}_i$.*

Proof. Let us first express the transition matrix of the original random walk in terms of the normalized graph Laplacian,

$$\begin{aligned}
\boldsymbol{W} = \boldsymbol{A}\boldsymbol{D}^{-1} &= \boldsymbol{D}^{1/2}(\boldsymbol{D}^{-1/2}\boldsymbol{A}\boldsymbol{D}^{-1/2})\boldsymbol{D}^{-1/2} \\
&= \boldsymbol{I} + \boldsymbol{D}^{1/2}\underbrace{(\boldsymbol{D}^{-1/2}\boldsymbol{A}\boldsymbol{D}^{-1/2} - \boldsymbol{I})}_{-\boldsymbol{N}}\boldsymbol{D}^{-1/2} \\
&= \boldsymbol{I} - \boldsymbol{D}^{1/2}\boldsymbol{N}\boldsymbol{D}^{-1/2} \qquad (3.6) \\
&= \boldsymbol{D}^{1/2}(\boldsymbol{I} - \boldsymbol{N})\boldsymbol{D}^{-1/2}.
\end{aligned}$$

By claim 2.7, $\boldsymbol{D}^{1/2}(\boldsymbol{I} - \boldsymbol{N})\boldsymbol{D}^{-1/2}$ and $\boldsymbol{I} - \boldsymbol{N}$ have the same eigenvalues, namely $1 - \nu_i$. We also have,

$$\tilde{\boldsymbol{W}} = \frac{1}{2}\boldsymbol{I} + \frac{1}{2}(\boldsymbol{I} - \boldsymbol{D}^{1/2}\boldsymbol{N}\boldsymbol{D}^{-1/2}) = \boldsymbol{I} - \frac{1}{2}\boldsymbol{D}^{1/2}\boldsymbol{N}\boldsymbol{D}^{-1/2}, \qquad (3.7)$$

implying that the eigenvalues of $\tilde{\boldsymbol{W}}$ are $1 - \nu_i/2$. Finally, we have,

$$\begin{aligned}
\tilde{\boldsymbol{W}}\boldsymbol{D}^{1/2}\boldsymbol{\psi}_i &= (\boldsymbol{I} - \frac{1}{2}\boldsymbol{D}^{1/2}\boldsymbol{N}\boldsymbol{D}^{-1/2})\boldsymbol{D}^{1/2}\boldsymbol{\psi}_i \\
&= \boldsymbol{D}^{1/2}\boldsymbol{\psi}_i - \frac{1}{2}\boldsymbol{D}^{1/2}\boldsymbol{N}\boldsymbol{\psi}_i \\
&= \boldsymbol{D}^{1/2}\boldsymbol{\psi}_i - \frac{\nu_i}{2}\boldsymbol{D}^{1/2}\boldsymbol{\psi}_i \qquad \text{using that } \boldsymbol{\psi}_i \text{ is an eigenvector of } \boldsymbol{N} \text{ with corresponding eigenvalue } \nu_i \\
&= \left(1 - \frac{\nu_i}{2}\right)\boldsymbol{D}^{1/2}\boldsymbol{\psi}_i. \qquad \square
\end{aligned}$$

Proof of theorem 3.6. TBD $\square$

Theorem 3.8 (Convergence rate of lazy random walk)**.** *For any unweighted connected graph G, we have that at time step t,*[5]

$$\|\boldsymbol{p}_t - \boldsymbol{\pi}\|_\infty \leq e^{-\nu_2 t/2}\sqrt{n}. \qquad (3.8)$$

[5] We will later see that ν_2 is an indicator of the "connectedness" of G.

Proof. TBD $\square$

Hitting Time

Definition 3.9 (Hitting time)**.** The *hitting time*,

$$H_{a,s} \doteq \min\{t \geq 1 \mid X_t = s, X_0 = a\}, \qquad (3.9)$$

is the number of steps to reach s starting from a. We have,

$$\boldsymbol{h}_s(a) \doteq \mathbb{E}[H_{a,s}] = 1 + \sum_{b \sim a} \frac{\boldsymbol{w}(\{a,b\})}{\boldsymbol{d}(a)}\boldsymbol{h}_s(b). \qquad (3.10)$$

Lemma 3.10. *If $\tilde{\boldsymbol{x}}$ is a solution to $\boldsymbol{L}\tilde{\boldsymbol{x}} = \boldsymbol{d} - \|\boldsymbol{d}\|_1 \mathbf{1}_s$,*[6] *then*

$$\boldsymbol{h}_s = \tilde{\boldsymbol{x}} - \tilde{\boldsymbol{x}}(s)\mathbf{1}. \tag{3.11}$$

[6] We use $\mathbf{1}_s$ as a shorthand notation for $\mathbf{1}_{\{s\}}$.

Proof. For any $a \neq s$, we can equivalently write eq. (3.10) as,

$$\mathbf{1}_a^\top \boldsymbol{h}_s = 1 + (\boldsymbol{W}\mathbf{1}_a)^\top \boldsymbol{h}_s \iff \mathbf{1}_a^\top (\boldsymbol{I} - \boldsymbol{W}^\top)\boldsymbol{h}_s = 1.$$

This yields a linear system of $n-1$ equations,

$$\mathbf{1} - \alpha \mathbf{1}_s = (\boldsymbol{I} - \underbrace{\boldsymbol{W}^\top}_{\boldsymbol{D}^{-1}\boldsymbol{A}})\boldsymbol{h}_s, \tag{3.12}$$

where α is due to the remaining degree of freedom, as the entry corresponding to s is not fixed. Multiplying from the left with $\boldsymbol{D}$, we obtain,

$$\boldsymbol{d} - \alpha \boldsymbol{d}(s)\mathbf{1}_s = (\boldsymbol{D} - \boldsymbol{A})\boldsymbol{h}_s = \boldsymbol{L}\boldsymbol{h}_s.$$

Recall that $\ker \boldsymbol{L} = \operatorname{span}\{\mathbf{1}\}$, and hence, for $\boldsymbol{h}_s$ to exist, we must choose α such that $\boldsymbol{d} - \alpha \boldsymbol{d}(s)\mathbf{1}_s \perp \mathbf{1}$. We have,

$$\mathbf{1}^\top(\boldsymbol{d} - \alpha \boldsymbol{d}(s)\mathbf{1}_s) = \|\boldsymbol{d}\|_1 - \alpha \boldsymbol{d}(s) \overset{!}{=} 0 \iff \alpha = \frac{\|\boldsymbol{d}\|_1}{\boldsymbol{d}(s)}.$$

Finally, note that the solution $\tilde{\boldsymbol{x}}$ to $\boldsymbol{L}\tilde{\boldsymbol{x}} = \boldsymbol{d} - \|\boldsymbol{d}\|_1 \mathbf{1}_s$ is not unique. Given that $\boldsymbol{h}_s$ is one solution, we have that any $\tilde{\boldsymbol{x}} = \boldsymbol{h}_s + c\mathbf{1}$ for $c \in \mathbb{R}$ is also a solution.[7] Yet, we know that $\boldsymbol{h}_s(s) = 0$, implying that $\boldsymbol{h}_s = \tilde{\boldsymbol{x}} - \tilde{\boldsymbol{x}}(s)\mathbf{1}$. □

[7] This follows directly from the fact that $\ker \boldsymbol{L} = \operatorname{span}\{\mathbf{1}\}$.

Commute Time

An issue with hitting times is that they do not need to be symmetric. This motivates the consideration of commute times, which correspond to the number of steps it takes to reach b from a and return to a.

TBD
Figure 3.2: Example where hitting times are not symmetric.

Definition 3.11 (Commute time)**.** The *commute time* between a and b is defined as,

$$C_{a,b} \doteq H_{a,b} + H_{b,a}. \tag{3.13}$$

Remark 3.12. By definition, commute times are symmetric.

Lemma 3.13. *If $\tilde{\boldsymbol{x}}$ is a solution to $\boldsymbol{L}\tilde{\boldsymbol{x}} = \|\boldsymbol{d}\|_1 (\mathbf{1}_a - \mathbf{1}_b)$, then*

$$\mathbb{E}[C_{a,b}] = (\mathbf{1}_a - \mathbf{1}_b)^\top \tilde{\boldsymbol{x}} = \tilde{\boldsymbol{x}}(a) - \tilde{\boldsymbol{x}}(b). \tag{3.14}$$

The $\tilde{\boldsymbol{x}}$ can be interpreted as electrical voltages inducing flow that routes $\|\boldsymbol{d}\|_1$ units from a to b.

Proof. We write $\boldsymbol{b}_v \doteq \boldsymbol{d} - \|\boldsymbol{d}\|_1 \mathbf{1}_v$ and let $\boldsymbol{L}\tilde{\boldsymbol{y}} = \boldsymbol{b}_b$ and $\boldsymbol{L}\tilde{\boldsymbol{z}} = \boldsymbol{b}_a$. Then,

$$\begin{aligned}\mathbb{E}[C_{a,b}] &= \boldsymbol{h}_b(a) + \boldsymbol{h}_a(b) \\ &= \tilde{\boldsymbol{y}}(a) - \tilde{\boldsymbol{y}}(b) + \tilde{\boldsymbol{z}}(b) - \tilde{\boldsymbol{z}}(a) \\ &= (\mathbf{1}_a - \mathbf{1}_b)^\top (\tilde{\boldsymbol{y}} - \tilde{\boldsymbol{z}}).\end{aligned}$$

Observe that $\tilde{\boldsymbol{x}} \doteq \tilde{\boldsymbol{y}} - \tilde{\boldsymbol{z}}$ solves $\boldsymbol{L}\tilde{\boldsymbol{x}} = \boldsymbol{b}_b - \boldsymbol{b}_a = \|\boldsymbol{d}\|_1 (\mathbf{1}_a - \mathbf{1}_b)$. □

We will see in chapter 12 that the expected commute time is intimately related to the electrical energy required to route flow from a to b, also called the effective resistance between a and b.

3.2 Concentration

Theorem 3.14 (Markov's inequality)**.** *For any random variable $X \geq 0$ and $t > 0$,*

$$\mathbb{P}[X \geq t] \leq \frac{\mathbb{E}[X]}{t}. \tag{3.15}$$

Proof. We have,

$$\mathbb{E}[X] = \int_0^\infty x f(x)\, dx \geq \int_t^\infty x f(x)\, dx \geq t \int_t^\infty f(x)\, dx = t\mathbb{P}[X \geq t],$$

where f is the probability density function of X. □

Fact 3.15 (Jensen's inequality)**.** *For a random variable X, if f is convex, then $\mathbb{E}[f(X)] \geq f(\mathbb{E}[X])$.*[8]

[8] We prove the finite form in theorem A.1.

TBD
Figure 3.3: Jensen's inequality.

Theorem 3.16 (Bernstein concentration bound)**.** *Given independent real-valued random variables $X_1, \ldots, X_k \in \mathbb{R}$ such that $\mathbb{E}[X_i] = 0$ and $|X_i| \leq R$. Let $X \doteq \sum_i X_i$ and $\sigma^2 \doteq \mathrm{Var}[X] = \sum_i \mathbb{E}\left[X_i^2\right]$. Then, for $t > 0$,*

$$\mathbb{P}[|X| \geq t] \leq 2\exp\left(\frac{-t^2}{2Rt + 4\sigma^2}\right). \tag{3.16}$$

Proof. TBD □

Corollary 3.17 (Chernoff bound)**.** *TBD*

Proof. TBD □

Theorem 3.18 (Bernstein matrix concentration bound)**.** *Suppose $\mathbf{X}_1, \ldots, \mathbf{X}_k \in \mathbb{R}^{n \times n}$ are independent symmetric matrix-valued random variables satisfying $\mathbb{E}[\mathbf{X}_i] = \mathbf{0}$ and $\|\mathbf{X}_i\|_2 \leq R$. Let $\mathbf{X} \doteq \sum_i \mathbf{X}_i$ and $\sigma^2 \doteq \mathrm{Var}[\mathbf{X}] = \sum_i \mathbb{E}\left[\mathbf{X}_i^2\right]$. Then, for $t > 0$,*

$$\mathbb{P}[\|\mathbf{X}\|_2 \geq t] \leq 2n\exp\left(\frac{-t^2}{2Rt + 4\sigma^2}\right). \tag{3.17}$$

Proof. TBD □

3.3 *Martingales*

Definition 3.19 (Martingale)**.** A *martingale* is a sequence of random variables $Z_0, \ldots, Z_k$ such that

$$\mathbb{E}[Z_i \mid Z_0, \ldots, Z_{i-1}] = Z_{i-1}. \tag{3.18}$$

That is, conditional on the outcome of all the previous random variables, the expectation of Z_i equals Z_{i-1}.

Typically, we use martingales to show a statement such as "Z_k is concentrated around $\mathbb{E}[Z_k]$".

We can alternatively think of a martingale as the sequence of changes in $\{Z_i\}_i$. Let $X_i \doteq Z_i - Z_{i-1}$. The sequence of $\{X_i\}_i$ is called *martingale difference sequence*. The martingale condition is equivalent to,

$$\mathbb{E}[X_i \mid Z_0, \ldots, Z_{i-1}] = \mathbb{E}[X_i \mid Z_0, X_1, \ldots, X_{i-1}] = 0. \tag{3.19}$$

We can write,

$$Z_k = Z_0 + \sum_{i=1}^{k} Z_i - Z_{i-1} = Z_0 + \sum_{i=1}^{k} X_i. \tag{3.20}$$

Theorem 3.20. *Suppose that $\{X_i\}_i$ form a scalar martingale difference sequence and $|X_i| \leq R$. Define the "pseudo-variance",*

$$W_i \doteq \sum_{j=1}^{i} \mathbb{E}\left[X_j^2 \mid X_1, \ldots, X_{j-1}\right], \tag{3.21}$$

Then,

$$\mathbb{P}\left[\left|\sum_{i=1}^{k} X_i\right| \geq k \text{ and } W_k \leq \sigma^2\right] \leq C_2 \exp\left(-C_1 \frac{t^2}{Rt + \sigma^2}\right), \tag{3.22}$$

for constants C_1, C_2.

By a simple union bound,

$$\mathbb{P}\left[\left|\sum_{i=1}^{k} X_i\right| \geq k\right] \leq \mathbb{P}\left[\left|\sum_{i=1}^{k} X_i\right| \geq k \text{ and } W_k \leq \sigma^2\right] + \mathbb{P}\left[W_k > \sigma^2\right]. \tag{3.23}$$

Proof. TBD □

4
Analysis

In this chapter, we study functions $f : S \to \mathbb{R}$ where $S \subseteq \mathbb{R}^n$.

4.1 First-order Taylor Approximations

Definition 4.1 (Gradient)**.** The *gradient* of a function $f : S \to \mathbb{R}$ at point $\boldsymbol{x} \in S$ is,

$$\boldsymbol{\nabla} f(\boldsymbol{x}) \doteq \begin{bmatrix} \frac{\partial f(\boldsymbol{x})}{\partial \boldsymbol{x}(1)} & \cdots & \frac{\partial f(\boldsymbol{x})}{\partial \boldsymbol{x}(n)} \end{bmatrix}^\top. \tag{4.1}$$

For a single-variable function $f : \mathbb{R} \to \mathbb{R}$ that is differentiable, we have for any $x, \delta \in \mathbb{R}$,

$$f(x+\delta) = f(x) + \frac{df(x)}{dx} + o(|\delta|), \quad \text{where } \lim_{\delta \to 0} \frac{o(|\delta|)}{|\delta|} = 0,$$

using a first-order Taylor approximation around x. We can use a similar approximation when f is a multi-variable function.

Definition 4.2 (Fréchet differentiable)**.** A function $f : S \to \mathbb{R}$ is *(Fréchet) differentiable* at $\boldsymbol{x} \in S$ if there exists $\boldsymbol{g} \in \mathbb{R}^n$ such that,[1]

$$\lim_{\substack{\boldsymbol{\delta} \in \mathbb{R}^n \\ \boldsymbol{\delta} \to \mathbf{0}}} \frac{|f(\boldsymbol{x}+\boldsymbol{\delta}) - [f(\boldsymbol{x}) + \boldsymbol{g}^\top \boldsymbol{\delta}]|}{\|\boldsymbol{\delta}\|_2} = 0. \tag{4.2}$$

[1] Here, $\boldsymbol{g}$ can be understood as a candidate for $\boldsymbol{\nabla} f(\boldsymbol{x})$ and $f(\boldsymbol{x}) + \boldsymbol{g}^\top \boldsymbol{\delta}$ is a linear approximation of f around $\boldsymbol{x}$.

This is equivalent to,

$$f(\boldsymbol{x}+\boldsymbol{\delta}) = f(\boldsymbol{x}) + \boldsymbol{g}^\top \boldsymbol{\delta} + o(\|\boldsymbol{\delta}\|_2), \tag{4.3}$$

for any $\boldsymbol{x} \in S$ and $\boldsymbol{\delta} \in \mathbb{R}^n$ where $\lim_{\boldsymbol{\delta} \to \mathbf{0}} \frac{o(\|\boldsymbol{\delta}\|_2)}{\|\boldsymbol{\delta}\|_2} = 0$. This is also called a *first-order expansion* of f around $\boldsymbol{x}$.

Remark 4.3. If this holds, $\boldsymbol{g} = \boldsymbol{\nabla} f(\boldsymbol{x})$.

Definition 4.4 (Continuously differentiable)**.** We say that $f : S \to \mathbb{R}$ is *continuously differentiable* on S if it is differentiable and its gradient is continuous on S.

Fact 4.5 (Taylor's theorem, first-order form)**.** *If $f : S \to \mathbb{R}$ is continuously differentiable, then for all $\boldsymbol{x}, \boldsymbol{y} \in S$,*

$$f(\boldsymbol{y}) = f(\boldsymbol{x}) + \nabla f(\boldsymbol{z})^\top (\boldsymbol{y} - \boldsymbol{x}), \tag{4.4}$$

for some $\boldsymbol{z} \in [\boldsymbol{x}, \boldsymbol{y}] \doteq \{\theta \boldsymbol{x} + (1-\theta)\boldsymbol{y} \mid \theta \in [0,1]\}$.

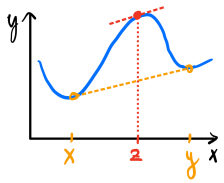

Figure 4.1: Illustration of Taylor's theorem. The affine approximation is shown in orange.

Taylor's theorem implies that f can be approximated by the affine function,

$$\boldsymbol{y} \to f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}),$$

when $\boldsymbol{y}$ is "close to" $\boldsymbol{x}$.

First-order Optimality Conditions

Definition 4.6 (Stationary point)**.** Given a function $f : S \to \mathbb{R}$, a point $\boldsymbol{x} \in S$ where $\nabla f(\boldsymbol{x}) = 0$ is called a *stationary point* of f.[2]

Theorem 4.7 (First-order optimality condition)**.** *If $\boldsymbol{x} \in S$ is a local extremum of a differentiable function $f : S \to \mathbb{R}$, then $\nabla f(\boldsymbol{x}) = 0$.*[3]

Proof. Assume $\boldsymbol{x}$ is a local minimum of f. Then, for all $\boldsymbol{d} \in \mathbb{R}^n$ and for all small enough $\lambda \in R$, we have $f(\boldsymbol{x}) \leq f(\boldsymbol{x} + \lambda \boldsymbol{d})$, so,

$$\begin{aligned} 0 &\leq f(\boldsymbol{x} + \lambda \boldsymbol{d}) - f(\boldsymbol{x}) \\ &= \lambda \nabla f(\boldsymbol{x})^\top \boldsymbol{d} + o(\lambda \left\|\boldsymbol{d}\right\|_2). \end{aligned}$$

using a first-order expansion of f around x

Dividing by λ and taking the limit $\lambda \to 0$, we obtain,

$$0 \leq \nabla f(\boldsymbol{x})^\top \boldsymbol{d} + \lim_{\lambda \to 0} \frac{o(\lambda \left\|\boldsymbol{d}\right\|_2)}{\lambda} = \nabla f(\boldsymbol{x})^\top \boldsymbol{d}.$$

Take $\boldsymbol{d} \doteq -\nabla f(\boldsymbol{x})$.[4] Then, $0 \leq -\left\|\nabla f(\boldsymbol{x})\right\|_2^2$, so $\nabla f(\boldsymbol{x}) = 0$. □

[2] Being a stationary point is not sufficient for optimality. Take for example the point $x \doteq 0$ of $f(x) \doteq x^3$.

[3] Here it is important that we have chosen $S \subseteq \mathbb{R}^n$ to be open. When S is not open, an extremum could be on the boundary of the domain, where the gradient is non-zero.

[4] We can only take this step because we assumed that S is open.

4.2 *Directional Derivatives*

Definition 4.8 (Jacobian)**.** Given a vector-valued function,

$$\boldsymbol{g} : \mathbb{R}^n \to \mathbb{R}^m, \quad \boldsymbol{x} \mapsto \begin{bmatrix} g_1(\boldsymbol{x}) \\ \vdots \\ g_m(\boldsymbol{x}) \end{bmatrix},$$

where $g_i : \mathbb{R}^n \to \mathbb{R}$, the *Jacobian* of $\boldsymbol{g}$ at $\boldsymbol{x} \in \mathbb{R}^n$ is,

$$D\boldsymbol{g}(\boldsymbol{x}) \doteq \begin{bmatrix} Dg_1(\boldsymbol{x}) \\ \vdots \\ Dg_m(\boldsymbol{x}) \end{bmatrix} \doteq \begin{bmatrix} \frac{\partial \boldsymbol{g}(\boldsymbol{x})}{\partial x(1)} & \cdots & \frac{\partial \boldsymbol{g}(\boldsymbol{x})}{\partial x(n)} \end{bmatrix}. \tag{4.5}$$

Remark 4.9. For $f : S \to \mathbb{R}$ and any $\boldsymbol{x} \in S$, $Df(\boldsymbol{x}) = \boldsymbol{\nabla} f(\boldsymbol{x})^\top$.

Definition 4.10. For a vector-valued function $\boldsymbol{g} : \mathbb{R}^n \to \mathbb{R}^m$, we define,

$$\boldsymbol{\nabla} \boldsymbol{g}(\boldsymbol{x}) \doteq D\boldsymbol{g}(\boldsymbol{x})^\top = \begin{bmatrix} \boldsymbol{\nabla} g_1(\boldsymbol{x}) & \cdots & \boldsymbol{\nabla} g_m(\boldsymbol{x}) \end{bmatrix}. \tag{4.6}$$

Definition 4.11 (Directional derivative)**.** Let $f : S \to \mathbb{R}$ be differentiable at $\boldsymbol{x} \in \mathbb{R}^n$. Given $\boldsymbol{d} \in \mathbb{R}^n$, the *directional derivative* of f at $\boldsymbol{x}$ in the direction $\boldsymbol{d}$ is,

$$Df(\boldsymbol{x})[\boldsymbol{d}] \doteq \lim_{\lambda \to 0} \frac{f(\boldsymbol{x} + \lambda \boldsymbol{d}) - f(\boldsymbol{x})}{\lambda}. \tag{4.7}$$

Lemma 4.12. $Df(\boldsymbol{x})[\boldsymbol{d}] = \boldsymbol{\nabla} f(\boldsymbol{x})^\top \boldsymbol{d} = Df(\boldsymbol{x})\boldsymbol{d}$.

Proof. Using a first-order expansion, we have,

$$f(\boldsymbol{x} + \lambda \boldsymbol{d}) = f(\boldsymbol{x}) + \lambda \boldsymbol{\nabla} f(\boldsymbol{x})^\top \boldsymbol{d} + o(\lambda \left\| \boldsymbol{d} \right\|_2).$$

Dividing by λ yields,

$$\frac{f(\boldsymbol{x} + \lambda \boldsymbol{d}) - f(\boldsymbol{x})}{\lambda} = \boldsymbol{\nabla} f(\boldsymbol{x})^\top \boldsymbol{d} + \underbrace{\frac{o(\lambda \left\| \boldsymbol{d} \right\|_2)}{\lambda}}_{\to 0}.$$

Taking the limit $\lambda \to 0$ gives the desired result. □

4.3 *Second-order Taylor Approximations*

Example 4.13. Given $f : S \to \mathbb{R}$ and any $\boldsymbol{x} \in S$, consider the vector-valued function $\boldsymbol{\nabla} f$. We have,

$$D\boldsymbol{\nabla} f(\boldsymbol{x}) = \begin{bmatrix} \frac{\partial \boldsymbol{\nabla} f(\boldsymbol{x})}{\partial \boldsymbol{x}(1)} & \cdots & \frac{\partial \boldsymbol{\nabla} f(\boldsymbol{x})}{\partial \boldsymbol{x}(n)} \end{bmatrix} = \begin{bmatrix} \frac{\partial^2 f(\boldsymbol{x})}{\partial \boldsymbol{x}(1)\, \partial \boldsymbol{x}(1)} & \cdots & \frac{\partial^2 f(\boldsymbol{x})}{\partial \boldsymbol{x}(n)\, \partial \boldsymbol{x}(1)} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 f(\boldsymbol{x})}{\partial \boldsymbol{x}(1)\, \partial \boldsymbol{x}(n)} & \cdots & \frac{\partial^2 f(\boldsymbol{x})}{\partial \boldsymbol{x}(n)\, \partial \boldsymbol{x}(n)} \end{bmatrix}. \tag{4.8}$$

Definition 4.14 (Hessian)**.** Given $f : S \to \mathbb{R}$ that is twice differentiable coordinatewise, we define the *Hessian* $\boldsymbol{H}_f(\boldsymbol{x}) \in \mathbb{R}^{n \times n}$ of f at a point $\boldsymbol{x} \in S$ as,

$$\boldsymbol{H}_f(\boldsymbol{x})(i, j) \doteq \frac{\partial^2 f(\boldsymbol{x})}{\partial \boldsymbol{x}(i)\, \partial \boldsymbol{x}(j)} \tag{4.9}$$

Remark 4.15. By eq. (4.8), $\boldsymbol{H}_f(\boldsymbol{x}) = (D\boldsymbol{\nabla} f(\boldsymbol{x}))^\top$.[5]

[5] Sometimes, $\boldsymbol{H}_f(\boldsymbol{x}) = \boldsymbol{\nabla}^2 f(\boldsymbol{x})$ is written (informally!).

Definition 4.16 (Twice Fréchet Differentiable)**.** We say a function $f : S \to \mathbb{R}$ is *twice (Fréchet) differentiable* at $\boldsymbol{x} \in S$ if there exists $\boldsymbol{g} \in \mathbb{R}^n$ and a matrix $\boldsymbol{M} \in \mathbb{R}^{n \times n}$ such that,[6]

$$\lim_{\substack{\boldsymbol{\delta} \in \mathbb{R}^n \\ \boldsymbol{\delta} \to \mathbf{0}}} \frac{|f(\boldsymbol{x}+\boldsymbol{\delta}) - [f(\boldsymbol{x}) + \boldsymbol{g}^\top \boldsymbol{\delta} + \frac{1}{2}\boldsymbol{\delta}^\top \boldsymbol{M}\boldsymbol{\delta}]|}{\|\boldsymbol{\delta}\|_2^2} = 0. \tag{4.10}$$

[6] You can think of $\boldsymbol{g}$ as a candidate for $\nabla f(\boldsymbol{x})$, $\boldsymbol{M}$ as a candidate for $\boldsymbol{H}_f(\boldsymbol{x})$, and $f(\boldsymbol{x}) + \boldsymbol{g}^\top \boldsymbol{\delta} + \frac{1}{2}\boldsymbol{\delta}^\top \boldsymbol{M}\boldsymbol{\delta}$ is a quadratic approximation of f around $\boldsymbol{x}$.

This is equivalent to,

$$f(\boldsymbol{x}+\boldsymbol{\delta}) = f(\boldsymbol{x}) + \boldsymbol{g}^\top \boldsymbol{\delta} + \frac{1}{2}\boldsymbol{\delta}^\top \boldsymbol{M}\boldsymbol{\delta} + o(\|\boldsymbol{\delta}\|_2^2), \tag{4.11}$$

for any $\boldsymbol{x} \in S$ and $\boldsymbol{\delta} \in \mathbb{R}^n$ where $\lim_{\boldsymbol{\delta} \to \mathbf{0}} \frac{o(\|\boldsymbol{\delta}\|_2^2)}{\|\boldsymbol{\delta}\|_2^2} = 0$. This is also called a *second-order expansion* of f around $\boldsymbol{x}$.

Remark 4.17. If this holds,

1. $\boldsymbol{g} = \nabla f(\boldsymbol{x})$,
2. $\boldsymbol{M} = \boldsymbol{H}_f(\boldsymbol{x})$, and
3. $\boldsymbol{H}_f(\boldsymbol{x}) = \boldsymbol{H}_f(\boldsymbol{x})^\top$.

Definition 4.18 (Twice continuously differentiable)**.** We say that $f : S \to \mathbb{R}$ is *twice continuously differentiable* on S if it is twice differentiable and the gradient and Hessian are continuous on S.

Fact 4.19 (Taylor's theorem, second-order form)**.** *If $f : S \to \mathbb{R}$ is twice continuously differentiable, then for all $\boldsymbol{x}, \boldsymbol{y} \in S$,*

$$f(\boldsymbol{y}) = f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top (\boldsymbol{y}-\boldsymbol{x}) + \frac{1}{2}(\boldsymbol{y}-\boldsymbol{x})^\top \boldsymbol{H}_f(\boldsymbol{z})(\boldsymbol{y}-\boldsymbol{x}), \tag{4.12}$$

for some $\boldsymbol{z} \in [\boldsymbol{x}, \boldsymbol{y}]$.

Second-order Optimality Conditions

Theorem 4.20 (Necessary second-order optimality condition)**.** *Let $f : S \to \mathbb{R}$ be twice continuously differentiable at $\boldsymbol{x} \in S$ and S be open. Then, if $\boldsymbol{x}$ is a local minimum, $\boldsymbol{H}_f(\boldsymbol{x})$ is positive semi-definite.*[7]

[7] In other words, if $\boldsymbol{H}_f(\boldsymbol{x})$ has a negative eigenvalue, $\boldsymbol{x}$ cannot be a minimum. Intuitively, you can think of $\boldsymbol{H}_f(\boldsymbol{x})$ as the curvature of f around $\boldsymbol{x}$, and therefore, a negative eigenvalue indicates that the function value can be decreased by moving in some direction.

Proof. Suppose $\boldsymbol{x} \in S$ is a local minimum. By the first-order optimality condition, $\nabla f(\boldsymbol{x}) = 0$. For any direction $\boldsymbol{d} \in \mathbb{R}^n$ and small enough $\lambda \in [-\epsilon, \epsilon] \setminus \{0\}$,

$$\begin{aligned} 0 &\leq f(\boldsymbol{x} + \lambda \boldsymbol{d}) - f(\boldsymbol{x}) && \text{using that } f \text{ is locally minimized at } \boldsymbol{x} \\ &= \frac{1}{2}\lambda^2 \boldsymbol{d}^\top \boldsymbol{H}_f(\boldsymbol{x})\boldsymbol{d} + o(\lambda^2 \|\boldsymbol{d}\|_2^2). && \text{using a second-order expansion} \end{aligned}$$

Multiplying both sides by $2/\lambda^2$ and taking the limit $\lambda \to 0$, we obtain,

$$0 \leq \boldsymbol{d}^\top \boldsymbol{H}_f(\boldsymbol{x})\boldsymbol{d} + \lim_{\lambda \to 0} \frac{o(\lambda^2 \|\boldsymbol{d}\|_2^2)}{\lambda^2} = \boldsymbol{d}^\top \boldsymbol{H}_f(\boldsymbol{x})\boldsymbol{d}. \qquad \square$$

Theorem 4.21 (Sufficient second-order optimality condition)**.** *Let $f : S \to \mathbb{R}$ be twice continuously differentiable at $\boldsymbol{x} \in S$ and S be open. Then, if $\boldsymbol{x}$ is a stationary point and $\boldsymbol{H}_f(\boldsymbol{x})$ is positive definite, $\boldsymbol{x}$ is a local minimum.*

Proof. Suppose that $\boldsymbol{x} \in S$ is stationary, i.e., $\nabla f(\boldsymbol{x}) = \boldsymbol{0}$, and $\boldsymbol{H}_f(\boldsymbol{x})$ is positive definite. For any direction $\boldsymbol{d} \in \mathbb{R}^n$ and small enough $\lambda \in [-\epsilon, \epsilon] \setminus \{0\}$,

$$
\begin{aligned}
f(\boldsymbol{x} + \lambda \boldsymbol{d}) &= f(\boldsymbol{x}) + \underbrace{\nabla f(\boldsymbol{x})^\top \boldsymbol{d}}_{0} + \frac{1}{2}\lambda^2 \boldsymbol{d}^\top \boldsymbol{H}_f(\boldsymbol{x}) \boldsymbol{d} + o(\lambda^2 \left\|\boldsymbol{d}\right\|_2^2) && \text{using a second-order expansion} \\
&\geq f(\boldsymbol{x}) + \frac{1}{2}\lambda^2 \lambda_{\min}(\boldsymbol{H}_f(\boldsymbol{x})) \left\|\boldsymbol{d}\right\|_2^2 + o(\lambda^2 \left\|\boldsymbol{d}\right\|_2^2) && \text{by corollary 2.6} \\
&\geq f(\boldsymbol{x}) + \frac{1}{4}\lambda_{\min}(\boldsymbol{H}_f(\boldsymbol{x})) \left\|\boldsymbol{d}\right\|_2^2 && \text{for small enough } \lambda \\
&> f(\boldsymbol{x}) && \text{using positive definiteness of } \boldsymbol{H}_f(\boldsymbol{x})
\end{aligned}
$$

□

PART II

Convex Optimization

5

Convex Geometry

We want to develop a better understanding of optimization problems. The general form of an *optimization problem* is,

$$\min_{\substack{y \in \mathbb{R}^n \\ g(y) \leq b}} f(y), \tag{5.1}$$

where $f : \mathbb{R}^n \to \mathbb{R}$ is the function to be minimized, $g : \mathbb{R}^n \to \mathbb{R}^m$ is a vector-valued function of m constraints with thresholds $b \in \mathbb{R}^m$.[1]

[1] It suffices to consider minimization problems. If we want to maximize a function f, this is equivalent to minimizing the function $-f$.

Definition 5.1 (Feasible set)**.** We call

$$\mathcal{F} \doteq \{x \in \mathbb{R}^n \mid g(x) \leq b\} \tag{5.2}$$

the *feasible set*. We call

- $x \in \mathcal{F}$ a *feasible point*; and
- $x \notin \mathcal{F}$ an *infeasible point*.

Definition 5.2 (Optimal solution)**.** We say that $x^* \in \mathbb{R}^n$ is *optimal* if $x^* \in \mathcal{F}$ and $\forall x \in \mathcal{F} : f(x^*) \leq f(x)$.

Let us look at a sufficient condition for optimal solutions.

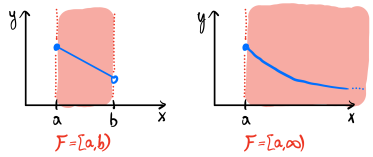

Figure 5.1: Examples of optimization problems without an optimal solution.

Theorem 5.3 (Extreme value theorem)**.** *Let $f : \mathbb{R}^n \to \mathbb{R}$ be continuous, and let $\mathcal{F} \subseteq \mathbb{R}^n$ be non-empty, bounded, and closed. Then, f is bounded on $\mathcal{F}$ and has an optimal solution.*

5.1 Convex Sets & Functions

Definition 5.4 (Convex set)**.** A set $S \subseteq \mathbb{R}^n$ is *convex* iff

$$\forall x, y \in S : \forall \theta \in [0,1] : \theta x + (1-\theta) y \in S. \tag{5.3}$$

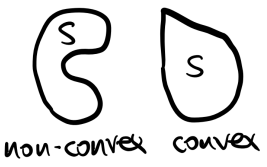

Figure 5.2: Example of a non-convex and a convex set.

Definition 5.5 (Convex function)**.** For a convex set $S \subseteq \mathbb{R}^n$, a function $f : S \to \mathbb{R}$ is *convex* on S iff

$$\forall \boldsymbol{x}, \boldsymbol{y} \in S : \forall \theta \in [0,1] : f(\theta \boldsymbol{x} + (1-\theta)\boldsymbol{y}) \leq \theta f(\boldsymbol{x}) + (1-\theta) f(\boldsymbol{y}). \quad (5.4)$$

Similarly, we call f *strictly convex* on S iff

$$\forall \boldsymbol{x}, \boldsymbol{y} \in S : \forall \theta \in [0,1] : f(\theta \boldsymbol{x} + (1-\theta)\boldsymbol{y}) < \theta f(\boldsymbol{x}) + (1-\theta) f(\boldsymbol{y}). \quad (5.5)$$

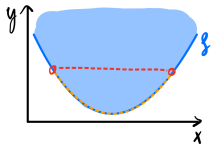

Figure 5.3: Example of a convex function. Any line between two points on f, lies "above" f. The epigraph of f is shown in blue.

Remark 5.6. If the function f is convex on S, we say that the function $-f$ is *concave* on S.

Remark 5.7. We call the optimization problem from eq. (5.1) *convex* iff f and g_i are convex functions. Sometimes it is useful to equivalently write

$$\min_{\boldsymbol{y} \in S} f(\boldsymbol{y}) \quad \text{for } S \doteq \{\boldsymbol{y} \in \mathbb{R}^n \mid \boldsymbol{g}(\boldsymbol{y}) \leq \boldsymbol{b}\}. \quad (5.6)$$

Observe that these characterizations are equivalent, as any convex set S can be characterized in terms of convex constraints.

Fact 5.8. *A differentiable and convex function f, whose domain $S \subseteq \mathbb{R}^n$ is open and convex, is always continuously differentiable.*

In the following, we will assume that $S \subseteq \mathbb{R}^n$ is open.

Definition 5.9 (Epigraph)**.** The *epigraph* of a function $f : S \to \mathbb{R}$ is

$$\operatorname{epi}(f) \doteq \{(\boldsymbol{x}, y) \mid f(\boldsymbol{x}) \leq y\} \subseteq \mathbb{R}^{n+1}. \quad (5.7)$$

Exercise 5.10. *The function $f : S \to \mathbb{R}$ is convex iff* $\operatorname{epi}(f)$ *is convex.*

Definition 5.11 ((Sub-)level set)**.** Given a function $f : S \to \mathbb{R}$, we call,

$$S_\alpha(f) \doteq \{\boldsymbol{x} \in S \mid f(\boldsymbol{x}) \leq \alpha\}, \quad (5.8)$$
$$L_\alpha(f) \doteq \{\boldsymbol{x} \in S \mid f(\boldsymbol{x}) = \alpha\}, \quad (5.9)$$

its *α-sub-level set* and *α-level set*, respectively.

Exercise 5.12. *Any α-sub-level set of a convex function is convex.*[2]

[2] Note that the other direction does not hold! Take $f(x) \doteq x^3$ as an example. Functions whose sub-level sets are convex are called *quasiconvex*.

5.2 First-order Characterization of Convexity

Theorem 5.13 (First-order characterization of convexity)**.** *Consider a differentiable function $f : S \to \mathbb{R}$. Then, f is convex iff*

$$f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}) \quad (5.10)$$

for all $\boldsymbol{x}, \boldsymbol{y} \in S$. Moreover, f is strictly convex iff

$$f(\boldsymbol{y}) > f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}) \quad (5.11)$$

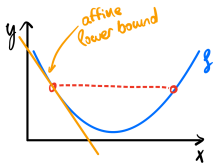

Figure 5.4: The first-order characterization characterizes convexity in terms of affine lower bounds.

Proof. We first prove the statement about convexity.

- "⇒": Fix any $x, y \in S$. As f is convex,

  $$f((1-\theta)x + \theta y) \leq (1-\theta)f(x) + \theta f(y),$$

  for all $\theta \in [0,1]$. We can rearrange to,

  $$f(\underbrace{(1-\theta)x + \theta y}_{x+\theta(y-x)}) - f(y) \leq \theta(f(x) - f(y)).$$

  Dividing by θ yields,

  $$\frac{f(x + \theta(y-x)) - f(x)}{\theta} \leq f(y) - f(x).$$

  Taking the limit $\theta \to 0$ on both sides gives the directional derivative at x in direction $y - x$,

  $$\nabla f(x)^\top (y - x) = Df(x)[y - x] \leq f(y) - f(x).$$

- "⇐": Fix any $x, y \in S$ and let $z \doteq \theta y + (1-\theta)x$. We have,

  $$\begin{aligned} f(y) &\geq f(z) + \nabla f(z)^\top (y - z), \quad \text{and} \\ f(x) &\geq f(z) + \nabla f(z)^\top (x - z). \end{aligned}$$

  We also have $(1-\theta)(y-x) = y - z$ and $\theta(y-x) = x - z$. Hence,

  $$\begin{aligned} \theta f(y) + (1-\theta) f(x) &\geq f(z) + \nabla f(z)^\top (\underbrace{\theta(y-z) + (1-\theta)(x-z)}_{0}) \\ &= f(\theta y + (1-\theta)x). \end{aligned}$$

Finally, observe that the statement about strict convexity can be proven analogously by making the inequalities strict. □

Theorem 5.14. *Let $f : S \to \mathbb{R}$ be a convex and differentiable function. Then, if $x \in S$ is a stationary point of f, then x is a global minimum of f.*

Proof. By the first-order characterization of convexity, we have for any $y \in S$,

$$f(y) \geq f(x) + \underbrace{\nabla f(x)^\top}_{0}(y - x) = f(x).$$ □

5.3 Second-order Characterization of Convexity

Theorem 5.15 (Second-order characterization of convexity)**.** *Consider a twice continuously differentiable function $f : S \to \mathbb{R}$.*[3]

1. *f is convex iff $H_f(x)$ is positive semi-definite for all $x \in S$.*
2. *f is strictly convex iff $H_f(x)$ is positive definite for all $x \in S$.*

[3] Here we need our assumption that S is open.

Proof. We first prove the statement about convexity.

- "⇐": Fix any $\boldsymbol{x}, \boldsymbol{y} \in S$. By the second-order form of Taylor's theorem,

$$\begin{aligned} f(\boldsymbol{y}) &= f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}) + \frac{1}{2} \underbrace{(\boldsymbol{y} - \boldsymbol{x})^\top \boldsymbol{H}_f(\boldsymbol{z})(\boldsymbol{y} - \boldsymbol{x})}_{\geq 0} \\ &\geq f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}), \end{aligned}$$

  for some $\boldsymbol{z} \in [\boldsymbol{x}, \boldsymbol{y}]$. This coincides with the first-order characterization of convexity.
- "⇒": Fix any $\boldsymbol{x} \in S$ and $\boldsymbol{d} \in \mathbb{R}^n$. Note that, as S is open, for small enough $\lambda \in [-\epsilon, \epsilon] \setminus \{0\}$, $\boldsymbol{x} + \lambda \boldsymbol{d} \in S$. We have,

$$\begin{aligned} 0 &\leq f(\boldsymbol{x} + \lambda \boldsymbol{d}) - [f(\boldsymbol{x}) + \lambda \nabla f(\boldsymbol{x})^\top \boldsymbol{d}] && \text{using the first-order characterization of convexity} \\ &= \frac{1}{2} \lambda^2 \boldsymbol{d}^\top \boldsymbol{H}_f(\boldsymbol{x}) \boldsymbol{d} + o(\lambda^2 \|\boldsymbol{d}\|_2^2). && \text{using a second-order expansion} \end{aligned}$$

  Multiplying both sides by $2/\lambda^2$ and taking the limit $\lambda \to 0$, we obtain,

$$0 \leq \boldsymbol{d}^\top \boldsymbol{H}_f(\boldsymbol{x}) \boldsymbol{d} + \lim_{\lambda \to 0} \frac{o(\lambda^2 \|\boldsymbol{d}\|_2^2)}{\lambda^2} = \boldsymbol{d}^\top \boldsymbol{H}_f(\boldsymbol{x}) \boldsymbol{d}.$$

The statement about strict convexity follows by using the first-order characterization of strict convexity instead and replacing inequalities with strict inequalities. □

6

Gradient Descent

Gradient descent is a method for solving minimization problems such as eq. (5.1).

Definition 6.1 (Approximate solution)**.** We say that a solution $\boldsymbol{x}_k$ to the optimization problem $\min_{\boldsymbol{x}\in S} f(\boldsymbol{x})$ is *ϵ-approximate* iff

$$f(\boldsymbol{x}_k) - f(\boldsymbol{x}^*) \leq \epsilon \tag{6.1}$$

for some $\epsilon > 0$, where $\boldsymbol{x}^* \in \arg\min_{\boldsymbol{x}\in S} f(\boldsymbol{x})$.

For this chapter, we assume that the optimization problem is unconstrained, i.e., $S = \mathbb{R}^n$. In chapter 8, we explore how we can solve constrained optimization problems using Lagrange multipliers.

The idea of gradient descent is to iteratively take a step in the opposite direction of the gradient starting from some initial point $\boldsymbol{x}_0 \in \mathbb{R}^n$,

$$\boldsymbol{x}_{i+1} \leftarrow \boldsymbol{x}_i - \alpha \nabla f(\boldsymbol{x}_i), \tag{6.2}$$

where $\alpha > 0$ is some learning rate.

TBD
Figure 6.1: Non-convex function.

When does gradient descent work? Clearly, we need that f is convex, otherwise gradient descent might not converge to the global minimum at all. But this is not enough! We also need to ensure that the gradient of f does not change arbitrarily when making very small steps, else the gradient direction would not be useful for us. This property is often called *smoothness*.

TBD
Figure 6.2: Function whose gradient is close to zero at a non-optimal point.

Finally, it is intuitively clear that we can do much better when we can ensure that the gradient of f is only close to zero around its minimum. If not, that is, we (almost) have "saddle points", the step size of gradient descent will slow down and depending on the stopping criterion we might even return a point that is not the minimizer. To

exclude such functions from our analysis, we often assume that f satisfies the *PL condition*, or else is *strongly convex*.

You can think of smoothness as providing a quadratic upper bound to our function, whereas strong convexity provides a quadratic lower bound.

6.1 Smoothness

Definition 6.2 (Smoothness)**.** Let $f : S \to \mathbb{R}$ be continuously differentiable. We say, f is *β-smooth* for some $\beta > 0$ iff for any $\boldsymbol{x}, \boldsymbol{y} \in S$

$$\|\nabla f(\boldsymbol{x}) - \nabla f(\boldsymbol{y})\|_2 \leq \beta \|\boldsymbol{x} - \boldsymbol{y}\|_2 . \tag{6.3}$$

In other words, the gradient of f is β-Lipschitz.

Lemma 6.3. *A twice continuously differentiable function $f : S \to \mathbb{R}$ is β-smooth iff for any $\boldsymbol{x} \in S$, $\lambda_{\max}(\boldsymbol{H}_f(\boldsymbol{x})) \leq \beta$.*

Proof. TBD □

Lemma 6.4. *A continuously differentiable function $f : S \to \mathbb{R}$ is β-smooth iff for any $\boldsymbol{x}, \boldsymbol{y} \in S$,*

$$f(\boldsymbol{y}) \leq f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}) + \frac{\beta}{2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 . \tag{6.4}$$

In words, $f(\boldsymbol{y})$ is upper bounded by a quadratic approximation based at $f(\boldsymbol{x})$.

Proof. TBD □

Analysis of Gradient Descent

A natural approach is to choose the gradient step of each iteration such that we minimize the upper bound (due to smoothness) based at the current solution,

$$\nabla_\delta \left(f(\boldsymbol{x}_i) + \nabla f(\boldsymbol{x}_i)^\top \delta + \frac{\beta}{2} \|\delta\|_2^2 \right) = \nabla f(\boldsymbol{x}_i) + \beta\delta \stackrel{!}{=} 0, \tag{6.5}$$

which is achieved for $\delta = -\frac{1}{\beta} \nabla f(\boldsymbol{x}_i)$. Thus,

$$f(\boldsymbol{x}_{i+1}) - f(\boldsymbol{x}_i) \leq \underbrace{\nabla f(\boldsymbol{x}_i)^\top \delta}_{-\frac{1}{\beta}\|\nabla f(\boldsymbol{x}_i)\|_2^2} + \underbrace{\frac{\beta}{2} \|\delta\|_2^2}_{\frac{1}{2\beta}\|\nabla f(\boldsymbol{x}_i)\|_2^2} = -\frac{1}{2\beta} \|\nabla f(\boldsymbol{x}_i)\|_2^2 . \tag{6.6}$$

Moreover, due to the first-order characterization of convexity,

$$f(\boldsymbol{x}_i) - f(\boldsymbol{x}^*) \leq \nabla f(\boldsymbol{x}_i)^\top (\boldsymbol{x}_i - \boldsymbol{x}^*) \leq \|\nabla f(\boldsymbol{x}_i)\|_2 \|\boldsymbol{x}_i - \boldsymbol{x}^*\|_2 , \tag{6.7}$$

where the second inequality follows from Cauchy-Schwarz. Combining the previous two inequalities,

$$f(\boldsymbol{x}_{i+1}) - f(\boldsymbol{x}_i) \leq -\frac{1}{2\beta}\left(\frac{f(\boldsymbol{x}_i) - f(\boldsymbol{x}^*)}{\|\boldsymbol{x}_i - \boldsymbol{x}^*\|_2}\right)^2 \leq -\frac{1}{2\beta}\left(\frac{f(\boldsymbol{x}_i) - f(\boldsymbol{x}^*)}{\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_2}\right)^2. \tag{6.8}$$

using that $\|\boldsymbol{x}_i - \boldsymbol{x}^*\|_2 \leq \|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_2$, which follows from f decreasing in every iteration and convexity

Theorem 6.5 (Convergence of gradient descent)**.** *Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is convex and β-smooth. The gradient descent scheme,*

$$\boldsymbol{x}_{i+1} \doteq \boldsymbol{x}_i - \frac{1}{\beta}\nabla f(\boldsymbol{x}_i), \tag{6.9}$$

yields an ϵ-approximate solution $\boldsymbol{x}_k$ for any

$$k \geq \frac{2\beta\,\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_2^2}{\epsilon}.$$

Proof. We prove $f(\boldsymbol{x}_k) - f(\boldsymbol{x}^*) \leq \frac{2\beta\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_2^2}{k+1}$ by induction on the length of the computation k. Suppose $k = 0$, then by the smoothness of f,

$$f(\boldsymbol{x}_0) \leq f(\boldsymbol{x}^*) - \underbrace{\nabla f(\boldsymbol{x}^*)^\top}_{\mathbf{0}}(\boldsymbol{x}_0 - \boldsymbol{x}^*) + \frac{\beta}{2}\,\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_2^2.$$

For the induction step, suppose that the statement holds for the k-th iterate. We write $\text{gap}_i \doteq f(\boldsymbol{x}_i) - f(\boldsymbol{x}^*)$. Using eq. (6.8),

$$\text{gap}_{k+1} - \text{gap}_k \leq -\frac{\text{gap}_k^2}{2\beta\,\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_2^2}.$$

Dividing by $\text{gap}_k \cdot \text{gap}_{k+1}$ and using that $\text{gap}_{k+1} > 0$ and $\text{gap}_k \geq \text{gap}_{k+1}$, we have,

$$\frac{1}{\text{gap}_k} - \frac{1}{\text{gap}_{k+1}} \leq -\frac{\text{gap}_k^2}{2\beta\,\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_2^2\,\text{gap}_k\text{gap}_{k+1}} \leq -\frac{1}{2\beta\,\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_2^2}.$$

Thus,

$$\frac{1}{\text{gap}_{k+1}} \geq \frac{1}{2\beta\,\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_2^2} + \frac{1}{\text{gap}_k} \geq \frac{(k+1)+1}{2\beta\,\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_2^2}$$

using the induction hypothesis

□

6.2 Strong Convexity

We can improve our analysis, when we assume that f is strongly convex, that is lower bounded by a quadratic. Intuitively, this ensures that our steps are large when we are far away from the optimum.

Definition 6.6 (Strong convexity)**.** Let $f : S \to \mathbb{R}$ be continuously differentiable. We say, f is *μ-strongly convex* for some $\mu > 0$ iff for any $\boldsymbol{x}, \boldsymbol{y} \in S$,

$$f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + \boldsymbol{\nabla} f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}) + \frac{\mu}{2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 . \tag{6.10}$$

Lemma 6.7. *A twice continuously differentiable function $f : S \to \mathbb{R}$ is μ-strongly convex iff for any $\boldsymbol{x} \in S$, $\lambda_{\min}(\boldsymbol{H}_f(\boldsymbol{x})) \geq \mu$.*

Proof. TBD □

Corollary 6.8. *If f is β-smooth and μ-strongly convex, then $\mu \leq \beta$.*

Definition 6.9 (Condition number)**.** We call $\kappa \doteq \frac{\beta}{\mu}$ the *condition number* of a function f that is β-smooth and μ-strongly convex.

Often, a weaker condition known as *PL condition* is sufficient to design fast algorithms.

Definition 6.10 (Polyak-Łojasiewicz inequality)**.** A continuously differentiable function $f : S \to \mathbb{R}$ satisfies the PL inequality with parameter $\mu > 0$ iff

$$\frac{1}{2} \|\boldsymbol{\nabla} f(\boldsymbol{x})\|_2^2 \geq \mu (f(\boldsymbol{x}) - f(\boldsymbol{x}^*)), \tag{6.11}$$

where $\boldsymbol{x}^* \in \arg\min_{\boldsymbol{x} \in S} f(\boldsymbol{x})$.

Intuitively, the norm of the gradient is tied to the suboptimality of the current solution.

Lemma 6.11. *Let $f : S \to \mathbb{R}$ be continuously differentiable and μ-strongly convex. Then, f satisfies the PL condition.*

Proof. As f is μ-strongly convex,

$$f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + \boldsymbol{\nabla} f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}) + \frac{\mu}{2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 ,$$

for any $\boldsymbol{x}, \boldsymbol{y} \in S$. Taking the minimum with respect to $\boldsymbol{y}$ on both sides yields,

$$f(\boldsymbol{x}^*) \geq f(\boldsymbol{x}) - \frac{1}{2\mu} \|\boldsymbol{\nabla} f(\boldsymbol{x})\|_2^2 ,$$

as the right-hand side is minimized for $\boldsymbol{y} = \boldsymbol{x} - \frac{1}{\mu} \boldsymbol{\nabla} f(\boldsymbol{x})$. The PL inequality follows from rearranging the terms. □

Analysis of Gradient Descent

Theorem 6.12 (Convergence of gradient descent with a strongly convex objective)**.** *Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is μ-strongly convex and β-smooth. The gradient descent scheme,*

$$\boldsymbol{x}_{i+1} \doteq \boldsymbol{x}_i - \frac{1}{\beta}\nabla f(\boldsymbol{x}_i), \tag{6.12}$$

yields an ϵ-approximate solution $\boldsymbol{x}_k$ for any

$$k \geq \kappa \log\left(\frac{\beta \left\|\boldsymbol{x}_0 - \boldsymbol{x}^*\right\|_2^2}{2\epsilon}\right).$$

Proof. See the first graded homework. □

Remark 6.13. It turns out that the PL condition is sufficient to establish this convergence rate.

6.3 Acceleration

We can get an algorithm that converges substantially faster than vanilla gradient descent, using a method known as *accelerated gradient descent*. The key idea is to — instead of only tracking the upper bound that is due to smoothness — also use track lower bounds. In one iteration we might not make much progress in terms of reducing the upper bound (that is, improving our current solution), but instead increase the upper bound, which still reduces the error.

Theorem 6.14 (Convergence of accelerated gradient descent)**.** *Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is convex and β-smooth. The accelerated gradient descent scheme,*

$$\begin{aligned}
a_i &\doteq \frac{i+1}{2}, \quad A_i \doteq \frac{(i+1)(i+2)}{4} \\
\boldsymbol{v}_0 &\doteq \boldsymbol{x}_0 - \frac{1}{2\beta}\nabla f(\boldsymbol{x}_0) \\
\boldsymbol{y}_i &\doteq \boldsymbol{x}_i - \frac{1}{\beta}\nabla f(\boldsymbol{x}_i) \\
\boldsymbol{x}_{i+1} &\doteq \frac{A_i \boldsymbol{y}_i + a_{i+1}\boldsymbol{v}_i}{A_{i+1}} \\
\boldsymbol{v}_{i+1} &\doteq \boldsymbol{v}_i - \frac{a_{i+1}}{\beta}\nabla f(\boldsymbol{x}_{i+1}),
\end{aligned} \tag{6.13}$$

yields an ϵ-approximate solution $\boldsymbol{x}_k$ for any

$$k \geq \sqrt{\frac{2\beta \left\|\boldsymbol{x}_0 - \boldsymbol{x}^*\right\|_2^2}{\epsilon}}.$$

Here, $\boldsymbol{y}_i$ is the current solution (i.e., an upper bound), $\boldsymbol{v}_i$ is a lower bound, and $\boldsymbol{x}_i$ a point that trades improving the lower/upper bounds.

Proof. TBD □

Acceleration with Strongly Convex Objectives

Theorem 6.15 (Convergence of accelerated gradient descent with a strongly convex objective)**.** *Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is μ-strongly convex and β-smooth. The accelerated gradient descent scheme,*

$$\begin{aligned} \boldsymbol{y}_0 &\doteq \boldsymbol{x}_0 \\ \boldsymbol{y}_{i+1} &\doteq \boldsymbol{x}_i - \frac{1}{\beta}\boldsymbol{\nabla} f(\boldsymbol{x}_i) \\ \boldsymbol{x}_{i+1} &\doteq (1+\theta)\boldsymbol{y}_{i+1} + \theta \boldsymbol{y}_i \end{aligned} \tag{6.14}$$

for $\theta \doteq \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}$ yields an ϵ-approximate solution $\boldsymbol{x}_k$ for any

$$k \geq \sqrt{\kappa} \log\left(\frac{\beta \left\|\boldsymbol{x}_0 - \boldsymbol{x}^*\right\|_2^2}{\epsilon}\right).$$

Proof. See the first graded homework. □

7
Non-Euclidean Geometries

8

Lagrange Multipliers and Duality

8.1 Separating Hyperplanes

Definition 8.1 (Hyperplane)**.** A *hyperplane* of dimension n is the subset,

$$H(\boldsymbol{n}, \mu) \doteq \{\boldsymbol{x} \in \mathbb{R}^n \mid \boldsymbol{n}^\top \boldsymbol{x} = \mu\}, \tag{8.1}$$

for some *normal* $\boldsymbol{n} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}$ and *threshold* $\mu \in \mathbb{R}$.

Every hyperplane divides $\mathbb{R}^n$ into two half-spaces $\{\boldsymbol{x} \mid \boldsymbol{n}^\top \boldsymbol{x} \leq \mu\}$ and $\{\boldsymbol{x} \mid \boldsymbol{n}^\top \boldsymbol{x} \geq \mu\}$. It separates two sets, if they lie in different half-spaces.

Definition 8.2 (Separating hyperplane)**.** We say a hyperplane H *separates* two sets A, B iff

$$\begin{aligned} &\forall a \in A : \boldsymbol{n}^\top \boldsymbol{a} \leq \mu \quad \text{and} \\ &\forall b \in B : \boldsymbol{n}^\top \boldsymbol{b} \geq \mu. \end{aligned} \tag{8.2}$$

If the inequalities are strict, we say that H *strictly* separates A and B.

If A, B are non-convex, we are not guaranteed that a separating hyperplane exists (e.g., a point cannot be separated from a ring around it). However, if we assume that $\boldsymbol{A}$ and $\boldsymbol{B}$ are convex, a separating hyperplane always exists.

Fact 8.3 (Separating hyperplane theorem)**.** *Given two disjoint and non-empty convex subsets $A, B \subseteq \mathbb{R}^n$, there exists a separating hyperplane.*

However, it is not true that there always exists a strictly separating hyperplane. Consider $A \doteq \{(x, y) \mid x \leq 0\}$ and $B \doteq \{(x, y) \mid x > 0 \text{ and } y \geq \frac{1}{x}\}$. Clearly they are disjoint and convex; however, the only separating hyperplane is $H = \{(x, y) \mid x = 0\}$, which intersects A.

TBD
Figure 8.1: Example where no strictly separating hyperplane exists.

When we also assume that A and B are closed and bounded, a strictly separating hyperplane always exists.

Theorem 8.4 (Separating hyperplane theorem; closed, bounded sets)**.** *Given two disjoint, closed, bounded, and non-empty convex subsets $A, B \subseteq \mathbb{R}^n$, there exists a strictly separating hyperplane.*

If $\boldsymbol{c} \in A, \boldsymbol{d} \in B$ are minimizers of $\min_{\boldsymbol{a}\in A,\ \boldsymbol{b}\in B} \|\boldsymbol{a}-\boldsymbol{b}\|_2$, then one such hyperplane is given by,

$$\boldsymbol{n} \doteq \boldsymbol{d} - \boldsymbol{c} \quad \textit{and} \quad \mu \doteq \frac{1}{2}\Big(\|\boldsymbol{d}\|_2^2 - \|\boldsymbol{c}\|_2^2\Big). \tag{8.3}$$

TBD
Figure 8.2: Illustration of strictly separating hyperplane.

Proof. We want to show that $\boldsymbol{n}^\top \boldsymbol{b} > \mu$ for all $\boldsymbol{b} \in B$. Then, $\boldsymbol{n}^\top \boldsymbol{a} < \mu$ for all $\boldsymbol{a} \in A$ follows by symmetry. We have,

$$\begin{aligned}
\boldsymbol{n}^\top \boldsymbol{d} - \mu &= (\boldsymbol{d}-\boldsymbol{c})^\top \boldsymbol{d} - \frac{1}{2}\Big(\|\boldsymbol{d}\|_2^2 - \|\boldsymbol{c}\|_2^2\Big) \\
&= \|\boldsymbol{d}\|_2^2 - \boldsymbol{d}^\top \boldsymbol{c} - \frac{1}{2}\|\boldsymbol{d}\|_2^2 + \frac{1}{2}\|\boldsymbol{c}\|_2^2 \\
&= \frac{1}{2}\|\boldsymbol{d}-\boldsymbol{c}\|_2^2 > 0.
\end{aligned}$$

using the asusmption that A, B are disjoint, close, and bounded, their distance is positive

Suppose for a contradiction that there exists a $\boldsymbol{u} \in B$ such that $\boldsymbol{n}^\top \boldsymbol{u} - \mu \leq 0$.

Consider the line defined by the distance minimizer $\boldsymbol{d}$ and the point on the "wrong side" $\boldsymbol{u}$, $\boldsymbol{b}(\lambda) \doteq \boldsymbol{d} + \lambda(\boldsymbol{u}-\boldsymbol{d})$. Taking the derivative of the distance between $\boldsymbol{b}(\lambda)$ and $\boldsymbol{c}$ and evaluating it at $\lambda = 0$ (which is when $\boldsymbol{b}(\lambda) = \boldsymbol{d}$), we obtain,

$$\begin{aligned}
\frac{d}{d\lambda}\|\boldsymbol{b}(\lambda)-\boldsymbol{c}\|_2^2\Big|_{\lambda=0} &= 2(\boldsymbol{d} - \lambda\boldsymbol{d} + \lambda\boldsymbol{u} - \boldsymbol{c})^\top(\boldsymbol{u}-\boldsymbol{d})\Big|_{\lambda=0} \\
&= 2(\boldsymbol{d}-\boldsymbol{c})^\top(\boldsymbol{u}-\boldsymbol{d}).
\end{aligned}$$

However,

$$\boldsymbol{n}^\top \boldsymbol{u} - \mu = (\boldsymbol{d}-\boldsymbol{c})^\top(\boldsymbol{u}-\boldsymbol{d}) + \underbrace{\boldsymbol{n}^\top \boldsymbol{d} - \mu}_{>0} \leq 0,$$

implies that $(\boldsymbol{d}-\boldsymbol{c})^\top(\boldsymbol{u}-\boldsymbol{d})$, and hence, the gradient are negative, which contradicts the minimality of $\boldsymbol{d}$. □

8.2 Lagrange Multipliers and KKT Conditions

We will now discuss how we can treat constraints in a convex optimization problem,

$$\alpha^* \doteq \min_{\substack{\boldsymbol{y}\in\mathbb{R}^n \\ \boldsymbol{A}\boldsymbol{y}=\boldsymbol{b} \\ \boldsymbol{g}(\boldsymbol{y})\leq \boldsymbol{0}}} f(\boldsymbol{y}), \tag{8.4}$$

where $f : \mathbb{R}^n \to \mathbb{R}$ is convex, $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, $\boldsymbol{b} \in \mathbb{R}^m$, and we have k convex constraints $g_i : \mathbb{R}^n \to \mathbb{R}$.

Remark 8.5. The linear constraints $\boldsymbol{A}\boldsymbol{y} = \boldsymbol{b}$ are not necessary, as they can also be modeled using the (convex) constraints $\boldsymbol{A}\boldsymbol{y} - \boldsymbol{b} \leq \boldsymbol{0}$ and $\boldsymbol{b} - \boldsymbol{A}\boldsymbol{y} \leq 0$. We include them here to see later that programs with only linear constraints can be handled in a slightly different way.

We call this optimization problem the *primal program* and we will later see that it has an associated dual problem. We say that $\boldsymbol{y} \in \mathbb{R}^n$ is *primal feasible* iff $\boldsymbol{A}\boldsymbol{y} = \boldsymbol{b}$ and $\boldsymbol{g}(\boldsymbol{y}) \leq \boldsymbol{0}$.

An Intuition

TBD
Figure 8.3: Illustration of simple constrained optimization.

In the following, we want to answer the question: "When is a feasible point optimal?" To simplify things a bit, let us consider the optimization problem of minimizing f under the single constraint g. Suppose we know the feasible optimum $\boldsymbol{y}^*$. Then for infinitesimal $\boldsymbol{\delta}$, if $\boldsymbol{\delta} \perp \boldsymbol{\nabla} g(\boldsymbol{y}^*)$, then $\boldsymbol{y}^* + \boldsymbol{\delta}$ and $\boldsymbol{y}^* - \boldsymbol{\delta}$ are feasible. But then we must have that $\boldsymbol{\delta} \perp \boldsymbol{\nabla} f(\boldsymbol{y}^*)$, or else one direction would improve the objective. This tells us that,

$$\boldsymbol{\delta}^\top \boldsymbol{\nabla} f(\boldsymbol{y}^*) = 0 = \boldsymbol{\delta}^\top \boldsymbol{\nabla} g(\boldsymbol{y}^*),$$

and hence, there exists some $\lambda \in \mathbb{R}$ such that $\boldsymbol{\nabla} f(\boldsymbol{y}^*) = \lambda \boldsymbol{\nabla} g(\boldsymbol{y}^*)$. This is the fundamental intuition behind a *Lagrange multiplier*: the gradient of the objective at an optimal point is a linear combination of the gradients of the (tight) constraints.

We say that the constraint g_i is *tight* at $\boldsymbol{y}$ iff $g_i(\boldsymbol{y}) = 0$. We know that if $\boldsymbol{y}$ is feasible, then $\boldsymbol{y} + \boldsymbol{\delta}$ is feasible for some infinitesimal $\boldsymbol{\delta}$ if for all constraints $i \in [k]$ we have that either the constraint is not tight, $g_i(\boldsymbol{y}) < 0$, or $\boldsymbol{y} + \boldsymbol{\delta}$ is "more" feasible, $\boldsymbol{\delta}^\top \boldsymbol{\nabla} g_i(\boldsymbol{y}) \leq 0$. The last condition tells us that for each tight constraint g_i, we get a half-space of feasible directions.

TBD
Figure 8.4: Each tight constraint yields a half-space of feasible directions.

On the other hand, if $\boldsymbol{y}^* + \boldsymbol{\delta}$ is feasible, then the objective f must "worsen", $\boldsymbol{\delta}^\top \boldsymbol{\nabla} f(\boldsymbol{y}^*) \geq 0$.

We can therefore write the gradient of f at $\boldsymbol{y}^*$ as a linear combination of the gradients of tight constraints.[1] Therefore, for some $\boldsymbol{x} \in \mathbb{R}^m$ and $\boldsymbol{s} \in \mathbb{R}^k$ with $\boldsymbol{s}(i) \geq 0$ if $g_i(\boldsymbol{y}^*) = 0$ and $\boldsymbol{s}(i) = 0$ otherwise,

[1] Note that linear constraints are always tight.

$$-\boldsymbol{\nabla} f(\boldsymbol{y}^*) = \sum_{i=1}^{k} \boldsymbol{s}(i) \boldsymbol{\nabla} g_i(\boldsymbol{y}^*) + \sum_{j=1}^{m} \boldsymbol{x}(j) \boldsymbol{A}(j,:). \tag{8.5}$$

The variables $\boldsymbol{s}$ and $\boldsymbol{x}$ are called the *dual variables* and are said to be *dual feasible* iff $\boldsymbol{s} \geq \boldsymbol{0}$. If $\boldsymbol{y}$ is also primal feasible, we say that $(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s})$ are *primal-dual feasible*. Equation (8.5) is equivalent to

$$\boldsymbol{\nabla} f(\boldsymbol{y}^*) + \boldsymbol{\nabla} g(\boldsymbol{y}^*)^\top \boldsymbol{s} + \boldsymbol{A}^\top \boldsymbol{x} = \boldsymbol{0} \tag{8.6}$$

along with the condition $\boldsymbol{g}(\boldsymbol{y}^*)^\top \boldsymbol{s} = \boldsymbol{0}$, ensuring that $\boldsymbol{s}(i) = 0$ when constraint g_i is not tight.

Definition 8.6 (Karush-Kuhn-Tucker (KKT) conditions and Lagrangian)**.** Points $\boldsymbol{y} \in \mathbb{R}^n, \boldsymbol{x} \in \mathbb{R}^m, \boldsymbol{s} \in \mathbb{R}^k$ satisfy the *Karush-Kuhn-Tucker conditions* iff,

$$\nabla_{\boldsymbol{y}} L(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s}) = \boldsymbol{0} \qquad \textit{gradient condition} \tag{8.7}$$

$$\boldsymbol{g}(\boldsymbol{y})^\top \boldsymbol{s} = \boldsymbol{0} \qquad \textit{complementary slackness} \tag{8.8}$$

ensures that $\boldsymbol{s}(i)$ is forced to 0 when constraint g_i is not tight

$$\boldsymbol{g}(\boldsymbol{y}) \leq \boldsymbol{0} \quad \text{and} \quad \boldsymbol{A}\boldsymbol{y} = \boldsymbol{b} \qquad \text{primal feasibility} \tag{8.9}$$

$$\boldsymbol{s} \geq \boldsymbol{0} \qquad \text{dual feasibility,} \tag{8.10}$$

where

$$L(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s}) \doteq f(\boldsymbol{y}) + \boldsymbol{s}^\top \boldsymbol{g}(\boldsymbol{y}) + \boldsymbol{x}^\top (\boldsymbol{b} - \boldsymbol{A}\boldsymbol{y}) \tag{8.11}$$

is the *Lagrangian* of the optimization problem.

Remark 8.7. Note that,

$$\nabla_{\boldsymbol{y}} L(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s}) = \nabla f(\boldsymbol{y}) + \nabla \boldsymbol{g}(\boldsymbol{y})^\top \boldsymbol{s} + \boldsymbol{A}^\top \boldsymbol{x},$$

which coincides with our intuition from eq. (8.6).

Intuitively, if $(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s})$ satisfy the KKT conditions, then $\boldsymbol{y}$ is optimal. It turns out that this intuition is correct, and we will prove this in the following section.

8.3 Lagrangian Duality

We have that if $(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s})$ are primal-dual feasible,

$$L(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s}) = f(\boldsymbol{y}) + \underbrace{\boldsymbol{s}^\top}_{\geq \boldsymbol{0}} \underbrace{\boldsymbol{g}(\boldsymbol{y})}_{\leq \boldsymbol{0}} + \boldsymbol{x}^\top (\underbrace{\boldsymbol{b} - \boldsymbol{A}\boldsymbol{y}}_{\boldsymbol{0}}) \leq f(\boldsymbol{y}). \tag{8.12}$$

We can also write the primal problem as a two-player game in terms of the Lagrangian,

$$\alpha^* = \min_{\substack{\boldsymbol{y} \in \mathbb{R}^n \\ \boldsymbol{A}\boldsymbol{y} = \boldsymbol{b} \\ \boldsymbol{g}(\boldsymbol{y}) \leq \boldsymbol{0}}} f(\boldsymbol{y}) = \min_{\boldsymbol{y} \in \mathbb{R}^n} \max_{\substack{\boldsymbol{x} \in \mathbb{R}^m \\ \boldsymbol{s} \in \mathbb{R}^k \\ \boldsymbol{s} \geq \boldsymbol{0}}} L(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s}). \tag{8.13}$$

This is because for a minimizing $\boldsymbol{y}$ all constraints have to be satisfied and the Lagrangian simplifies to $L(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s}) = f(\boldsymbol{y})$. If $\boldsymbol{b} - \boldsymbol{A}\boldsymbol{y} = \boldsymbol{0}$ was violated, making $\boldsymbol{x}$ large sends $L(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s}) \to \infty$. If $\boldsymbol{g}(\boldsymbol{y}) \leq \boldsymbol{0}$ was violated, making $\boldsymbol{s}$ large sends $L(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s}) \to \infty$.

We therefore have for any dual feasible $(\boldsymbol{x}, \boldsymbol{s})$,

$$\alpha^* = f(\boldsymbol{y}^*) \geq L(\boldsymbol{y}^*, \boldsymbol{x}, \boldsymbol{s}) \geq \min_{\boldsymbol{y} \in \mathbb{R}^n} L(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s}) \doteq L(\boldsymbol{x}, \boldsymbol{s}). \tag{8.14}$$

Definition 8.8 (Dual program)**.** The *dual program* is given as,

$$\beta^* \doteq \max_{\substack{\boldsymbol{x}\in\mathbb{R}^m \\ \boldsymbol{s}\in\mathbb{R}^k \\ \boldsymbol{s}\geq\mathbf{0}}} L(\boldsymbol{x},\boldsymbol{s}) = \max_{\substack{\boldsymbol{x}\in\mathbb{R}^m \\ \boldsymbol{s}\in\mathbb{R}^k \\ \boldsymbol{s}\geq\mathbf{0}}} \min_{\boldsymbol{y}\in\mathbb{R}^n} L(\boldsymbol{y},\boldsymbol{x},\boldsymbol{s}). \tag{8.15}$$

Theorem 8.9 (Weak duality)**.** $\beta^* \leq \alpha^*$.

Proof. This follows immediately from eq. (8.14). □

Remark 8.10. Observe that the dual program is a convex optimization problem in disguise. We can equivalently consider the optimization problem,

$$-\beta^* = \min_{\substack{\boldsymbol{x}\in\mathbb{R}^m \\ \boldsymbol{s}\in\mathbb{R}^k \\ \boldsymbol{s}\geq\mathbf{0}}} -L(\boldsymbol{x},\boldsymbol{s}) = \min_{\substack{\boldsymbol{x}\in\mathbb{R}^m \\ \boldsymbol{s}\in\mathbb{R}^k \\ \boldsymbol{s}\geq\mathbf{0}}} \max_{\boldsymbol{y}\in\mathbb{R}^n} -L(\boldsymbol{y},\boldsymbol{x},\boldsymbol{s}). \tag{8.16}$$

Note that $-L(\boldsymbol{y},\boldsymbol{x},\boldsymbol{s})$ is a linear function in the dual variables, hence convex, and $-L(\boldsymbol{x},\boldsymbol{s})$ is a maximum of these functions, so also convex.

Definition 8.11 (Strong duality)**.** If $\alpha^* = \beta^*$, we say that *strong duality* holds.

Remark 8.12. It is immediately clear that if we consider linear programs, that is, we have only linear constraints, strong duality always holds. This is because at any primal-dual feasible point $(\boldsymbol{y},\boldsymbol{x})$, we have that by definition $L(\boldsymbol{y},\boldsymbol{x}) = f(\boldsymbol{y})$.

Before we analyze when strong duality holds, let us return to the KKT conditions and confirm our intuition from the previous section.

Theorem 8.13 (KKT conditions are necessary)**.** *If strong duality holds, then the KKT conditions hold for primal-dual optimal* $(\boldsymbol{y}^*,\boldsymbol{x}^*,\boldsymbol{s}^*)$.

Proof. By strong duality,

$$L(\boldsymbol{y}^*,\boldsymbol{x}^*,\boldsymbol{s}^*) = \alpha^* = \beta^*.$$

As $L(\boldsymbol{y},\boldsymbol{x}^*,\boldsymbol{s}^*)$ is a convex function in $\boldsymbol{y}$, we have that

$$\nabla_{\boldsymbol{y}} L(\boldsymbol{y},\boldsymbol{x}^*,\boldsymbol{s}^*)\big|_{\boldsymbol{y}=\boldsymbol{y}^*} = \mathbf{0},$$

so the gradient condition is satisfied. Moreover, we have,

$$f(\boldsymbol{y}^*) = \alpha^* = L(\boldsymbol{y}^*,\boldsymbol{x}^*,\boldsymbol{s}^*) = f(\boldsymbol{y}^*) + \boldsymbol{s}^\top \boldsymbol{g}(\boldsymbol{y}^*) + \boldsymbol{x}^\top (\underbrace{\boldsymbol{b} - \boldsymbol{A}\boldsymbol{y}}_{\mathbf{0}}),$$

so $\boldsymbol{s}^\top \boldsymbol{g}(\boldsymbol{y}^*) = \mathbf{0}$ (complementary slack) holds. By assumption, $(\boldsymbol{y}^*,\boldsymbol{x}^*,\boldsymbol{s}^*)$ are primal-dual feasible. □

Theorem 8.14 (KKT conditions are sufficient)**.** *If the KKT conditions hold at $(\hat{\boldsymbol{y}}, \hat{\boldsymbol{x}}, \hat{\boldsymbol{s}})$, then they are primal-dual optimal and strong duality holds.*

Proof. Because $L(\boldsymbol{y}, \hat{\boldsymbol{x}}, \hat{\boldsymbol{s}})$ is a convex function, by the gradient condition, $\hat{\boldsymbol{y}}$ is its global minimizer. Therefore,

$$L(\hat{\boldsymbol{y}}, \hat{\boldsymbol{x}}, \hat{\boldsymbol{s}}) = \min_{\boldsymbol{y} \in \mathbb{R}^n} L(\boldsymbol{y}, \hat{\boldsymbol{x}}, \hat{\boldsymbol{s}}) = L(\hat{\boldsymbol{x}}, \hat{\boldsymbol{s}}) \leq \beta^*.$$

At the same time, using primal-dual feasibility and complementary slack,

$$L(\hat{\boldsymbol{y}}, \hat{\boldsymbol{x}}, \hat{\boldsymbol{s}}) = f(\hat{\boldsymbol{y}}) + \underbrace{\hat{\boldsymbol{s}}^\top \boldsymbol{g}(\hat{\boldsymbol{y}})}_{\mathbf{0}} + \hat{\boldsymbol{x}}^\top (\underbrace{\boldsymbol{b} - \boldsymbol{A}\hat{\boldsymbol{y}}}_{\mathbf{0}}) = f(\hat{\boldsymbol{y}}) \geq \alpha^*.$$

Therefore, $\alpha^* \leq \beta^*$. By weak duality, we get the opposite inequality, and hence strong duality holds. □

8.4 Slater's Condition

We will now discuss under which circumstances strong duality holds. In general, we can have that strong duality does not hold.

Example 8.15. TBD

Definition 8.16 (Slater's condition)**.** *Slater's condition* is satisfied iff there exists some $\boldsymbol{y} \in \mathbb{R}^n$ such that $\boldsymbol{A}\boldsymbol{y} = \boldsymbol{b}$ and $\boldsymbol{g}(\boldsymbol{y}) < \mathbf{0}$.[2] We say that $\boldsymbol{y}$ is *strictly feasible*.

[2] This is satisfied by most useful optimization problems.

Theorem 8.17. *If Slater's condition holds, then strong duality holds.*

Proof. TBD □

8.5 Example: Duality of Maximum Flow and Minimum Cut

We can write the (directed) maximum flow problem in a graph G as,

$$\max_{\substack{F \in \mathbb{R} \\ \boldsymbol{f} \in \mathbb{R}^{|E|} \\ \boldsymbol{B}\boldsymbol{f} = F(\mathbf{1}_t - \mathbf{1}_s) \\ \mathbf{0} \leq \boldsymbol{f} \leq \boldsymbol{c}}} F = - \min_{\substack{F \in \mathbb{R} \\ \boldsymbol{f} \in \mathbb{R}^{|E|} \\ \boldsymbol{B}\boldsymbol{f} = F(\mathbf{1}_t - \mathbf{1}_s) \\ \mathbf{0} \leq \boldsymbol{f} \leq \boldsymbol{c}}} -F. \tag{8.17}$$

Here, $\mathbf{0} \leq \boldsymbol{f}$ ensures that directions are respected, and $\boldsymbol{f} \leq \boldsymbol{c}$ ensures that edge capacities are respected. $F(\mathbf{1}_t - \mathbf{1}_s)$ is the demand of a flow routing F units from t to s. Observe that when there exists an s-t path, then there are flows $\boldsymbol{f}$ strictly satisfying the constraint $\mathbf{0} \leq \boldsymbol{f} \leq \boldsymbol{c}$, and hence, Slater's criterion is satisfied. By strong duality, the above program is equivalent to its dual program,

$$- \max_{\substack{\boldsymbol{x} \in \mathbb{R}^{|V|} \\ \boldsymbol{s} \in \mathbb{R}^{|E|} \\ \boldsymbol{s} \geq \mathbf{0}}} \min_{\substack{F \in \mathbb{R} \\ \boldsymbol{f} \in \mathbb{R}^{|E|} \\ \boldsymbol{f} \geq \mathbf{0}}} -F + \boldsymbol{s}^\top (\boldsymbol{f} - \boldsymbol{c}) + \boldsymbol{x}^\top (F(\mathbf{1}_t - \mathbf{1}_s) - \boldsymbol{B}\boldsymbol{f}) \tag{8.18}$$

$$= \min_{\substack{\boldsymbol{x}\in\mathbb{R}^{|V|}\\ \boldsymbol{s}\in\mathbb{R}^{|E|}\\ \boldsymbol{s}\geq\boldsymbol{0}\\ (\mathbf{1}_t-\mathbf{1}_s)^\top\boldsymbol{x}=1\\ \boldsymbol{s}\geq\boldsymbol{B}^\top\boldsymbol{x}}} \boldsymbol{s}^\top\boldsymbol{c} \tag{8.19}$$

$$= \min_{\substack{\boldsymbol{x}\in\mathbb{R}^{|V|}\\ (\mathbf{1}_t-\mathbf{1}_s)^\top\boldsymbol{x}=1}} \sum_{e\in E} \max\{(\boldsymbol{B}^\top\boldsymbol{x})(e), 0\}\cdot\boldsymbol{c}(e). \tag{8.20}$$

Observe that the dual program is a linear program computing the minimum cut.

8.6 Fenchel Conjugates

Definition 8.18 (Fenchel conjugate)**.** Given a function $f : S \to \mathbb{R}$, its *Fenchel conjugate* is the function $f^* : \mathbb{R}^n \to \mathbb{R}$,[3]

$$f^*(\boldsymbol{z}) \doteq \sup_{\boldsymbol{y}\in S} \boldsymbol{z}^\top\boldsymbol{y} - f(\boldsymbol{y}). \tag{8.21}$$

[3] In principle, f^* is a function defined over the dual space of S, but this will not be very important to us.

Remark 8.19. f^* is convex, as it is the maximum of linear functions.

Example 8.20. Let us consider the function $f(\boldsymbol{y}) \doteq \|\boldsymbol{y}\|$. Then we have,

$$\begin{aligned} f^*(\boldsymbol{z}) &= \sup_{\boldsymbol{y}\in\mathbb{R}^n} \boldsymbol{z}^\top\boldsymbol{y} - \underbrace{\|\boldsymbol{y}\|}_{\doteq\theta} \\ &= \sup_{\theta\geq 0} \theta(\|\boldsymbol{z}\|_* - 1) \\ &= \begin{cases} \infty & \|\boldsymbol{z}\|_* > 1 \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \tag{8.22}$$

Recall that

$$\theta\,\|\boldsymbol{z}\|_* = \max_{\substack{\boldsymbol{y}\in\mathbb{R}^n\\ \|\boldsymbol{y}\|=\theta}} \boldsymbol{z}^\top\boldsymbol{y}.$$

Example 8.21. For the function $f(\boldsymbol{y}) \doteq \frac{1}{p}\|\boldsymbol{y}\|_p^p$, we have that $f^*(\boldsymbol{z}) = \frac{1}{q}\|\boldsymbol{z}\|_q^q$, where $\frac{1}{p}+\frac{1}{q}=1$.

Example 8.22. When we have a primal program with only linear constraints,

$$\min_{\substack{\boldsymbol{y}\in\mathbb{R}^n\\ \boldsymbol{A}\boldsymbol{y}=\boldsymbol{b}}} f(\boldsymbol{y}),$$

we can write the dual program as,

$$\begin{aligned} \max_{\boldsymbol{x}\in\mathbb{R}^m}\min_{\boldsymbol{y}\in\mathbb{R}^n} f(\boldsymbol{y}) + \boldsymbol{x}^\top(\boldsymbol{b}-\boldsymbol{A}\boldsymbol{y}) &= \max_{\boldsymbol{x}\in\mathbb{R}^m} \boldsymbol{b}^\top\boldsymbol{x} - \max_{\boldsymbol{y}\in\mathbb{R}^n}(\boldsymbol{x}^\top\boldsymbol{A}\boldsymbol{y} - f(\boldsymbol{y})) \\ &= \max_{\boldsymbol{x}\in\mathbb{R}^m} \boldsymbol{b}^\top\boldsymbol{x} - f^*(\boldsymbol{A}^\top\boldsymbol{x}). \end{aligned} \tag{8.23}$$

Lemma 8.23 (Properties of the Fenchel conjugate). *If f is strictly convex (i.e., $\boldsymbol{H}_f \succ \boldsymbol{0}$) and $\boldsymbol{\nabla} f$ is a surjective mapping onto $\mathbb{R}^n$,*

1. $\nabla f(\nabla f^*(\boldsymbol{z})) = \boldsymbol{z}$ *and* $\nabla f^*(\nabla f(\boldsymbol{y})) = \boldsymbol{y}$*;*
2. $(f^*)^* = f$*; and*
3. $\boldsymbol{H}_{f^*}(\nabla f(\boldsymbol{y})) = \boldsymbol{H}_f^{-1}(\boldsymbol{y})$.[4]

Proof. TBD □

[4] Thus, if f is β-smooth, f^* is β-strongly convex. In other words, f^* has the "opposite" curvature of f.

9
Newton's Method

PART III

Spectral Graph Theory

10

Introduction to Spectral Graph Theory

Spectral graph theory studies graphs through linear algebra. The fundamental object that we will work with is the Laplacian matrix that we introduced in definition 1.4.

10.1 Eigenvalues of the Laplacian Matrix

Lemma 10.1. *Denote by $X_1 \cup \cdots \cup X_k = V$ the connected components of a graph G. Then we have for the Laplacian matrix $\boldsymbol{L}$ of G,*

$$\ker \boldsymbol{L} = \operatorname{span}\{\mathbf{1}_{X_1}, \ldots, \mathbf{1}_{X_k}\}, \tag{10.1}$$

where $\mathbf{1}_X(v) = \mathbb{1}\{v \in X\}$. In particular, $\ker \boldsymbol{L} = \operatorname{span}\{\mathbf{1}\}$ if G is connected.

Proof. TBD □

Corollary 10.2. *When G has k connected components, then $\lambda_i(\boldsymbol{L}) = 0$ for all $i \in [k]$.*

In the following, we will study connected graphs. If a graph consists of multiple connected components, we may study each individually.

Lemma 10.3. $\boldsymbol{L} \preceq 2\boldsymbol{D}$.

Proof. It can be shown that $\boldsymbol{D} + \boldsymbol{A} \succeq \mathbf{0}$,[1] so, $\boldsymbol{D} \succeq -\boldsymbol{A}$, and hence, $\boldsymbol{L} \preceq 2\boldsymbol{D}$. □

[1] The proof is similar to the proof that $\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{A} \succeq \mathbf{0}$. $\boldsymbol{D} + \boldsymbol{A}$ is also called the *signless Laplacian matrix* and is an interesting object in its own right: its eigenvalues can be used to count edges and identify bipartite components.

Electrical Energy

Lemma 10.4. *For any voltages $\boldsymbol{x} \in \mathbb{R}^{|V|}$,*

$$\lambda_2(\boldsymbol{L}) \|\boldsymbol{x}\|_2^2 \leq \mathcal{E}(\boldsymbol{x}) \leq \lambda_n(\boldsymbol{L}) \|\boldsymbol{x}\|_2^2 . \tag{10.2}$$

Proof. TBD □

Lemma 10.5. *Voltages $\boldsymbol{x} \in \mathbb{R}^{|V|}$ routing demands $\boldsymbol{d} \perp \mathbf{1}$ satisfy,*

$$\frac{\|\boldsymbol{d}\|_2^2}{\lambda_n(\boldsymbol{L})} \leq \mathcal{E}(\boldsymbol{x}) \leq \frac{\|\boldsymbol{d}\|_2^2}{\lambda_2(\boldsymbol{L})}. \tag{10.3}$$

Proof. Recall from eq. (1.18) that $\mathcal{E}(\boldsymbol{x}) = \boldsymbol{d}^\top \boldsymbol{L}^+ \boldsymbol{d}$. Using Courant-Fischer,

$$\begin{aligned}
\lambda_n(\boldsymbol{L}^+) &= \max_{\boldsymbol{x} \neq \mathbf{0}} \frac{\boldsymbol{x}^\top \boldsymbol{L}^+ \boldsymbol{x}}{\|\boldsymbol{x}\|_2^2} \geq \frac{\boldsymbol{d}^\top \boldsymbol{L}^+ \boldsymbol{d}}{\|\boldsymbol{d}\|_2^2} \\
\lambda_2(\boldsymbol{L}^+) &= \min_{\substack{\boldsymbol{x} \neq \mathbf{0} \\ \boldsymbol{x} \perp \mathbf{1}}} \frac{\boldsymbol{x}^\top \boldsymbol{L}^+ \boldsymbol{x}}{\|\boldsymbol{x}\|_2^2} \leq \frac{\boldsymbol{d}^\top \boldsymbol{L}^+ \boldsymbol{d}}{\|\boldsymbol{d}\|_2^2}, && \text{using that } \ker \boldsymbol{L}^+ = \ker \boldsymbol{L} = \operatorname{span}\{\mathbf{1}\}
\end{aligned}$$

using that $\boldsymbol{d} \neq \mathbf{0}$ and $\boldsymbol{d} \perp \mathbf{1}$. Rearranging the inequalities, we obtain,

$$\lambda_2(\boldsymbol{L}^+) \|\boldsymbol{d}\|_2^2 \leq \boldsymbol{d}^\top \boldsymbol{L}^+ \boldsymbol{d} \leq \lambda_n(\boldsymbol{L}^+) \|\boldsymbol{d}\|_2^2 .$$

Finally, recall from eq. (2.16) that $\lambda_2(\boldsymbol{L}^+) = \lambda_n(\boldsymbol{L})^{-1}$ and $\lambda_n(\boldsymbol{L}^+) = \lambda_2(\boldsymbol{L})^{-1}$. □

Useful Inequalities

Lemma 10.6 (Path inequality)**.** *We have that*

$$(n-1) P_n \succeq G_{1,n}, \tag{10.4}$$

where P_n is the unit weight path graph on n vertices and $G_{i,j}$ is the unit weight graph on n vertices with the single edge $\{i, j\}$.

Proof. Fix any $\boldsymbol{x} \in \mathbb{R}^n$ and let $\boldsymbol{\Delta}(i) \doteq \boldsymbol{x}(i+1) - \boldsymbol{x}(i)$. We have,

$$\begin{aligned}
\boldsymbol{x}^\top G_{1,n} \boldsymbol{x} &= (\boldsymbol{x}(n) - \boldsymbol{x}(1))^2 \\
&= \left(\sum_{i=1}^{n-1} \boldsymbol{\Delta}(i) \right)^2 \\
&= (\mathbf{1}_{n-1}^\top \boldsymbol{\Delta})^2 \\
&\leq \|\mathbf{1}_{n-1}\|_2^2 \, \|\boldsymbol{\Delta}\|_2^2 && \text{using Cauchy-Schwarz} \\
&= (n-1) \sum_{i=1}^{n-1} \boldsymbol{\Delta}(i)^2 \\
&= (n-1) \sum_{i=1}^{n-1} (\boldsymbol{x}(i+1) - \boldsymbol{x}(i))^2 \\
&= (n-1) \boldsymbol{x}^\top P_n \boldsymbol{x}.
\end{aligned}$$

□

Lemma 10.7. *For any unit weight graph G on n vertices with diameter*[2] *D,*

$$\lambda_2(G) \geq \frac{1}{nD}. \tag{10.5}$$

[2] The *diameter* of a graph is the maximum shortest distance between any two vertices.

Proof. We denote by $G^{i,j}$ the subgraph of G consisting of the shortest $i-j$ path. We have,

$$\begin{aligned} K_n = \sum_{i<j} G_{i,j} &\preceq \sum_{i<j} \underbrace{(j-i)}_{\leq D} \underbrace{G^{i,j}}_{\subseteq G} \\ &\preceq n^2 DG. \end{aligned}$$

analogously to the path inequality

Thus, $n^2 D\lambda_2(G) \geq \lambda_2(K_n) = n$, and hence, $\lambda_2(G) \geq \frac{1}{nD}$. □

10.2 Examples

Lemma 10.8 (Spectrum of the complete graph)**.**

$$\lambda_2(K_n) = \cdots = \lambda_n(K_n) = n. \tag{10.6}$$

Proof. We have, $\boldsymbol{A} = \mathbf{1}\mathbf{1}^\top - \boldsymbol{I}$ and $\boldsymbol{D} = (n-1)\boldsymbol{I}$, so $\boldsymbol{L} = n\boldsymbol{I} - \mathbf{1}\mathbf{1}^\top$. For any $\boldsymbol{x} \perp \mathbf{1}$, $\boldsymbol{L}\boldsymbol{x} = n\boldsymbol{x}$. □

We now want to better understand λ_2 and λ_n for some common graphs. The tools we will use are the following:

- to lower bound $\lambda_2(G)$: Relate the eigenvalues of K_n and G, yielding $K_n \preceq f(n)G$. Knowing that $\lambda_2(K_n) = n$, we have, $f(n)\lambda_2(G) \geq n$, and hence, $\lambda_2(G) \geq \frac{n}{f(n)}$.
- to upper bound $\lambda_2(G)$: Due to Courant-Fischer,

  $$\lambda_2(G) = \min_{\substack{\boldsymbol{x}\perp\mathbf{1} \\ \boldsymbol{x}\neq\mathbf{0}}} \frac{\boldsymbol{x}^\top \boldsymbol{L}\boldsymbol{x}}{\boldsymbol{x}^\top\boldsymbol{x}} \leq \frac{\boldsymbol{y}^\top \boldsymbol{L}\boldsymbol{y}}{\boldsymbol{y}^\top\boldsymbol{y}},$$

  for any $\boldsymbol{y} \perp \mathbf{1}, \boldsymbol{y} \neq \mathbf{0}$. We can therefore find a so-called *test vector* $\boldsymbol{y}$ with these properties.
- to lower bound $\lambda_n(G)$: Similarly, due to Courant-Fischer,

  $$\lambda_n(G) = \max_{\boldsymbol{x}\neq\mathbf{0}} \frac{\boldsymbol{x}^\top \boldsymbol{L}\boldsymbol{x}}{\boldsymbol{x}^\top\boldsymbol{x}} \geq \frac{\boldsymbol{y}^\top \boldsymbol{L}\boldsymbol{y}}{\boldsymbol{y}^\top\boldsymbol{y}},$$

  for any $\boldsymbol{y} \neq \mathbf{0}$.
- to upper bound $\lambda_n(G)$: Using that $\boldsymbol{L} \preceq 2\boldsymbol{D}$, we have, $\lambda_n(G) \leq 2\max_{v\in V} \boldsymbol{d}(v)$, where $\boldsymbol{d}(v)$ is the weighted degree of v.

Path Graph

Exercise 10.9. $\lambda_2(P_n) = \Theta\left(\frac{1}{n^2}\right)$.

Exercise 10.10. $\lambda_n(P_n) \in [1, 4]$.

Complete Binary Tree

Exercise 10.11. $\lambda_2(T_d) = \Theta\left(\frac{1}{n}\right)$.

Exercise 10.12. $\lambda_n(T_d) \in [1, 6]$.

11

Conductance and Expanders

11.1 Conductance

Definition 11.1 (Volume)**.** The *volume* of a set of vertices $S \subseteq V$ is the sum of weighted degrees,

$$\mathrm{vol}(S) \doteq \sum_{v \in S} \boldsymbol{d}(v) = \mathbf{1}_S^\top \boldsymbol{d} = \mathbf{1}_S^\top \boldsymbol{D} \mathbf{1}_S. \tag{11.1}$$

A *cut* $(S, V \setminus S)$ is a proper subset of vertices, $\emptyset \subset S \subset V$ partitioning vertices into two sets S and $V \setminus S$.

Definition 11.2 (Value of a cut)**.** The *value* of a cut $(S, V \setminus S)$ is the sum of weights of crossing edges,

$$\begin{aligned} c(S) &\doteq \sum_{\substack{\{a,b\} \in E \\ a \in S,\ b \in V \setminus S}} \boldsymbol{w}(\{a,b\}) \\ &= \sum_{\{a,b\} \in E} \boldsymbol{w}(\{a,b\}) [\mathbf{1}_S(a) - \mathbf{1}_S(b)]^2 = \mathbf{1}_S^\top \boldsymbol{L} \mathbf{1}_S. \end{aligned} \tag{11.2}$$

Definition 11.3 (Conductance of a cut)**.** The *conductance* of a cut $(S, V \setminus S)$ is,

$$\phi(S) \doteq \frac{c(S)}{\min\{\mathrm{vol}(S), \mathrm{vol}(V \setminus S)\}} = \phi(V \setminus S) \in [0, 1], \tag{11.3}$$

where, if the graph has unit weights, $c(S) = |E(S, V \setminus S)|$ counts the number of crossing edges.

Remark 11.4. If $\mathrm{vol}(S) \leq \mathrm{vol}(V \setminus S)$, then we can write

$$\phi(S) = \frac{\mathbf{1}_S^\top \boldsymbol{L} \mathbf{1}_S}{\mathbf{1}_S^\top \boldsymbol{D} \mathbf{1}_S}. \tag{11.4}$$

Definition 11.5 (Conductance of a graph)**.** The *conductance* of a graph G is the smallest conductance of all cuts,

$$\phi(G) \doteq \min_{\emptyset \subset S \subset V} \phi(S). \tag{11.5}$$

Thus, $\phi(G)$ is small if there is a "good" cut (with few crossing edges relative to the volume of the parts). In contrast, if $\phi(G)$ is large, then G is well-connected, i.e., there is no "good" cut.

Definition 11.6 (Expander and expander decomposition)**.** For any $\phi \in (0, 1]$, if $\phi(G) \geq \phi$, then G is called a *ϕ-expander*.

A *ϕ-expander decomposition* of *quality q* is a partition of the vertex set $V = X_1 \cup \cdots \cup X_k$ such that

1. $G[X_i]$ is a ϕ-expander; and
2. the number of edges not contained in any $G[X_i]$ is at most $q\phi m$, i.e. only "few" edges cross the parts.

In chapter 17, we discuss how we can efficiently find an expander decomposition. Let us consider a few examples.

Exercise 11.7. *$\phi(K_n) = \frac{n}{2(n-1)}$. So K_n is a $\frac{1}{2}$-expander.*

Exercise 11.8. *$\phi(P_n) = \frac{1}{n-1}$. So P_n is a $\frac{1}{n}$-expander.*

Lemma 11.9. *If G is a connected ϕ-expander with unit weights, then we have for the diameter D of G,*

$$D = \mathcal{O}\left(\frac{\log m}{\phi}\right). \tag{11.6}$$

Proof. Fix any pair of vertices $s, t \in V$. Let $B(s, d)$ be the closed ball around s of radius d. Let $E(B(s, d))$ be the internal edge set of $B(s, d)$.

Observe that since G is connected, we have $|E(B(s, 0))| \geq 1$. Moreover, for each $d \geq 0$ where $|E(B(s, d))| \leq \frac{m}{2}$, we have by the definition of a ϕ-expander that

$$|E(B(s, d), V \setminus B(s, d))| \geq \phi \cdot |E(B(s, d))|.$$

Thus, $|E(B(s, d+1))| \geq (1+\phi)|E(B(s, d))|$. Let $d_{\max}$ be the largest integer such that $|E(B(s, d_{\max}))| \leq m/2$. We have $d_{\max} \leq 2\log m/\phi$ as otherwise,

$$|E(B(s, d_{\max}))| > (1+\phi)^{2\log m/\phi} \geq (1 + \phi/2 + (\phi/2)^2)^{2\log m/\phi} \geq e^{\log m} = m, \qquad \text{using } e^x < 1 + x + x^2 \text{ for } x < 1.79$$

which gives a contradiction to the fact that the number of edges in G is m. Thus, for a radius of $2\log m/\phi + 1$, the ball centered at s has more than $m/2$ edges.

Finally, follow the same argument from t. As both balls contain more than $m/2$ edges, they must intersect in at least one edge. But this implies that there is a s-t path of length $\mathcal{O}(\log m/\phi)$. □

11.2 Cheeger's Inequality

Theorem 11.10 (Cheeger's inequality)**.** *We have for a graph G and its normalized Laplacian matrix $\boldsymbol{N}$,*

$$\frac{\lambda_2(\boldsymbol{N})}{2} \leq \phi(G) \leq \sqrt{2\lambda_2(\boldsymbol{N})}. \tag{11.7}$$

In words, $\lambda_2(\boldsymbol{N})$ approximates the conductance $\phi(G)$ up to a square root. This is why we say that $\lambda_2(\boldsymbol{N})$ is a measure of connectivity of a graph.

Proof. TBD □

11.3 Sparsity

A concept related to conductance is the notion of *sparsity*.

Definition 11.11 (Sparsity)**.** The *sparsity* of a cut $(S, V \setminus S)$ is,

$$\psi(S) \doteq \frac{c(S)}{\min\{|S|, |V \setminus S|\}} = \psi(V \setminus S) \in [0, \max_{v \in V} \boldsymbol{d}(v)]. \tag{11.8}$$

The sparsity of a graph is again defined as the smallest sparsity of all cuts.

Remark 11.12. If $|S| \leq |V \setminus S|$, then we can write

$$\psi(S) = \frac{\mathbf{1}_S^\top \boldsymbol{L} \mathbf{1}_S}{\mathbf{1}_S^\top \mathbf{1}_S}. \tag{11.9}$$

Intuitively, sparsity corresponds to a non-normalized variant of conductance. Sometimes it is easier to reason about sparsity than it is to reason about conductance.

Lemma 11.13. *We have for any cut $(S, V \setminus S)$ in a connected unit weight graph that $\psi(S) \geq \phi(S)$.*

Proof. As the graph is connected and has unit weights, $\mathrm{vol}(S) = \sum_{v \in S} \boldsymbol{d}(v) \geq |S|$. □

An alternative version of Cheeger's inequality relates the second eigenvalue of $\boldsymbol{L}$ (not $\boldsymbol{N}$!) to the sparsity of the graph.

Fact 11.14 (Cheeger's inequality for sparsity)**.** *We have for a graph G and its Laplacian matrix $\boldsymbol{L}$,*

$$\frac{\lambda_2(\boldsymbol{L})}{2} \leq \psi(G) \leq \sqrt{2\lambda_2(\boldsymbol{L}) \max_{v \in V} \boldsymbol{d}(v)}. \tag{11.10}$$

12

Effective Resistance

Definition 12.1 (Effective resistance)**.** The *effective resistance*,

$$R_{\text{eff}}(a,b) \doteq \min_{\substack{\boldsymbol{f}\in\mathbb{R}^{|E|}\\ \boldsymbol{B}\boldsymbol{f}=\mathbf{1}_b-\mathbf{1}_a}} \mathcal{E}(\boldsymbol{f}) = \min_{\substack{\boldsymbol{f}\in\mathbb{R}^{|E|}\\ \boldsymbol{B}\boldsymbol{f}=\mathbf{1}_b-\mathbf{1}_a}} \boldsymbol{f}^\top \boldsymbol{R}\boldsymbol{f}, \tag{12.1}$$

is the minimum electrical energy required to route one unit of flow from a to b.

Remark 12.2. Per definition of electrical energy, routing F units of flow from a to b costs $F^2 R_{\text{eff}}(a,b)$.

Let us first consider a few examples.

TBD
Figure 12.1: Sequential resistors.

Example 12.3. For the graph of fig. 12.1, $R_{\text{eff}}(1,k+1) = \sum_{i=1}^{k} \boldsymbol{r}(i)$.

Proof sketch. For the flow to be 1, by Ohm's law, the voltage difference across edge i must be $\boldsymbol{r}(i)$. □

TBD
Figure 12.2: Parallel resistors.

Example 12.4. For the graph of fig. 12.2, $R_{\text{eff}}(1,2) = \frac{1}{\sum_{i=1}^{k} 1/\boldsymbol{r}(i)}$.

Proof sketch. For the flow to be 1, by Ohm's law, we must have,

$$1 = \sum_{i=1}^{k} \frac{\tilde{\boldsymbol{x}}(\{1,2\})}{\boldsymbol{r}(i)},$$

where $\tilde{\boldsymbol{x}}(\{1,2\})$ is the voltage difference between vertices 1 and 2. Note that $R_{\text{eff}}(1,2) = \tilde{\boldsymbol{x}}(\{1,2\})$. □

Lemma 12.5. $R_{\text{eff}}(a,b) = \left\| \boldsymbol{L}^{+/2}(\mathbf{1}_b - \mathbf{1}_a) \right\|_2^2$.

Proof. As the electrical flow $\tilde{\boldsymbol{f}}$ is energy-minimizing, we have that $R_{\text{eff}}(a,b) = \tilde{\boldsymbol{f}}^\top \boldsymbol{R}\tilde{\boldsymbol{f}}$. Recall that by Ohm's law this flow corresponds to

voltages $\tilde{\boldsymbol{x}}$ solving $\boldsymbol{L}\tilde{\boldsymbol{x}} = \mathbf{1}_b - \mathbf{1}_a$, that is, $\tilde{\boldsymbol{x}} = \boldsymbol{L}^+(\mathbf{1}_b - \mathbf{1}_a)$. We obtain,

$$\begin{aligned} R_{\text{eff}}(a,b) = \tilde{\boldsymbol{f}}^\top \boldsymbol{R}\tilde{\boldsymbol{f}} = \tilde{\boldsymbol{x}}^\top \boldsymbol{L}\tilde{\boldsymbol{x}} &= (\mathbf{1}_b - \mathbf{1}_a)^\top \boldsymbol{L}^+ \boldsymbol{L}\boldsymbol{L}^+(\mathbf{1}_b - \mathbf{1}_a) \\ &= (\mathbf{1}_b - \mathbf{1}_a)^\top \boldsymbol{L}^+(\mathbf{1}_b - \mathbf{1}_a) && \text{using that } \mathbf{1}_b - \mathbf{1}_a \perp \mathbf{1} \\ &= \left\| \boldsymbol{L}^{+/2}(\mathbf{1}_b - \mathbf{1}_a) \right\|_2^2. \end{aligned}$$ □

Lemma 12.6. *If G is a ϕ-expander, then*

$$R_{\text{eff}}(a,b) \leq 2\phi^{-2}\left(\frac{1}{\boldsymbol{d}(b)} + \frac{1}{\boldsymbol{d}(a)}\right). \tag{12.2}$$

Proof. By Cheeger's inequality,

$$\phi \leq \phi(G) \leq \sqrt{2\lambda_2(\boldsymbol{N})} \implies \frac{\phi^2}{2} \leq \lambda_2(\boldsymbol{N}).$$

By Courant-Fischer, we have that for any $\boldsymbol{y} \perp \ker \boldsymbol{N}$,

$$\frac{\phi^2}{2} \leq \lambda_2(\boldsymbol{N}) \leq \frac{\boldsymbol{y}^\top \boldsymbol{N}\boldsymbol{y}}{\boldsymbol{y}^\top \boldsymbol{y}} \implies \frac{\phi^2}{2}\boldsymbol{y}^\top \boldsymbol{y} \leq \boldsymbol{y}^\top \boldsymbol{N}\boldsymbol{y}.$$

Equivalently,

$$\frac{\phi^2}{2}\boldsymbol{\Pi}_{\boldsymbol{N}} \preceq \boldsymbol{N}, \qquad \text{using that } \boldsymbol{\Pi}_{\boldsymbol{N}}\boldsymbol{v} = \boldsymbol{v} \text{ for } \boldsymbol{v} \perp \ker \boldsymbol{N}, \text{ and } \boldsymbol{\Pi}_{\boldsymbol{N}}\boldsymbol{v} = 0 \text{ if } \boldsymbol{v} \in \ker \boldsymbol{N}$$

where $\boldsymbol{\Pi}_{\boldsymbol{N}}$ is the projection orthogonal to the kernel of $\boldsymbol{N}$. From this we conclude that

$$2\phi^{-2}\boldsymbol{\Pi}_{\boldsymbol{N}} = 2\phi^{-2}\boldsymbol{\Pi}_{\boldsymbol{N}}^+ \succeq \boldsymbol{N}^+, \qquad \text{using } \boldsymbol{\Pi}_{\boldsymbol{N}}^+ = \boldsymbol{\Pi}_{\boldsymbol{N}}$$

as $\boldsymbol{A} \succeq \boldsymbol{B}$ implies $\boldsymbol{A}^+ \preceq \boldsymbol{B}^+$ when $\ker \boldsymbol{A} = \ker \boldsymbol{B}$. By eq. (2.18),

$$\boldsymbol{N}^+ = (\boldsymbol{D}^{-1/2}\boldsymbol{L}\boldsymbol{D}^{-1/2})^+ = \boldsymbol{\Pi}_{\boldsymbol{N}}\boldsymbol{D}^{1/2}\boldsymbol{L}^+\boldsymbol{D}^{1/2}\boldsymbol{\Pi}_{\boldsymbol{N}} \tag{12.3}$$

Therefore, for any $\boldsymbol{y} \perp \ker \boldsymbol{N}$,

$$2\phi^{-2}\boldsymbol{y}^\top \boldsymbol{y} \geq \boldsymbol{y}^\top \boldsymbol{N}^+ \boldsymbol{y} = \boldsymbol{y}^\top \boldsymbol{D}^{1/2}\boldsymbol{L}^+\boldsymbol{D}^{1/2}\boldsymbol{y}.$$

Substituting $\boldsymbol{z} \doteq \boldsymbol{D}^{-1/2}\boldsymbol{y}$, we obtain,

$$2\phi^{-2}\boldsymbol{z}^\top \boldsymbol{D}^{-1}\boldsymbol{z} \geq \boldsymbol{z}^\top \boldsymbol{L}^+ \boldsymbol{z}.$$

Finally, observe that for $\boldsymbol{z} \doteq \mathbf{1}_b - \mathbf{1}_a$, we have that $\boldsymbol{y} = \boldsymbol{D}^{1/2}(\mathbf{1}_b - \mathbf{1}_a) \perp \ker \boldsymbol{N}$ as $\mathbf{1}_b - \mathbf{1}_a \perp \ker \boldsymbol{L}$ and therefore,[1]

$$R_{\text{eff}}(a,b) = \boldsymbol{z}^\top \boldsymbol{L}^+ \boldsymbol{z} \leq 2\phi^{-2}\boldsymbol{z}^\top \boldsymbol{D}^{-1}\boldsymbol{z} = 2\phi^{-2}\left(\frac{1}{\boldsymbol{d}(b)} + \frac{1}{\boldsymbol{d}(a)}\right). \qquad \square$$

[1] We have for the kernel of the normalized Laplacian matrix, $\boldsymbol{N} = \boldsymbol{D}^{-1/2}\boldsymbol{L}\boldsymbol{D}^{-1/2}$, that $\ker \boldsymbol{N} = \boldsymbol{D}^{1/2} \ker \boldsymbol{L} = \operatorname{span}\{\boldsymbol{D}^{1/2}\mathbf{1}\}$.

Lemma 12.7. $\mathbb{E}[C_{a,b}] = \|\boldsymbol{d}\|_1 R_{\text{eff}}(a,b)$.

Proof. Recall that $\mathbb{E}[C_{a,b}] = (\mathbf{1}_a - \mathbf{1}_b)^\top \tilde{\boldsymbol{x}}$ for a solution $\tilde{\boldsymbol{x}}$ to $\boldsymbol{L}\tilde{\boldsymbol{x}} = \|\boldsymbol{d}\|_1 (\mathbf{1}_a - \mathbf{1}_b)$, that is, $\tilde{\boldsymbol{x}} = \|\boldsymbol{d}\|_1 \boldsymbol{L}^+(\mathbf{1}_a - \mathbf{1}_b)$. Now, observe that,

$$R_{\text{eff}}(b,a) = (\mathbf{1}_a - \mathbf{1}_b)^\top \boldsymbol{L}^+(\mathbf{1}_a - \mathbf{1}_b) = \frac{1}{\|\boldsymbol{d}\|_1}(\mathbf{1}_a - \mathbf{1}_b)^\top \tilde{\boldsymbol{x}}.$$

Thus, $\mathbb{E}[C_{a,b}] = \|\boldsymbol{d}\|_1 R_{\text{eff}}(b,a)$. Using symmetry of the commute time, $\mathbb{E}[C_{a,b}] = \mathbb{E}[C_{b,a}] = \|\boldsymbol{d}\|_1 R_{\text{eff}}(a,b)$. □

Corollary 12.8. *Effective resistance is symmetric.*

Remark 12.9. For an edge $e = \{a,b\} \in E$, we write,

$$R_{\text{eff}}(e) \doteq R_{\text{eff}}(a,b) = R_{\text{eff}}(b,a). \tag{12.4}$$

12.1 Effective Resistance as a Metric

Before showing that effective resistance is a metric on the set of vertices, we consider the following lemma. We will write,

$$\tilde{\boldsymbol{x}}_{a,b} \doteq \boldsymbol{L}^+(\mathbf{1}_b - \mathbf{1}_a), \tag{12.5}$$

for the electrical voltages required to route one unit of current from a to b.

Lemma 12.10. *If $\tilde{\boldsymbol{x}}_{a,b}$ is a solution to $\boldsymbol{L}\tilde{\boldsymbol{x}}_{a,b} = \mathbf{1}_b - \mathbf{1}_a$, then we have for all $c \in V$ that $\tilde{\boldsymbol{x}}_{a,b}(b) \geq \tilde{\boldsymbol{x}}_{a,b}(c) \geq \tilde{\boldsymbol{x}}_{a,b}(a)$.*

Proof sketch. Consider any $c \in V \setminus \{a,b\}$. Then, $(\boldsymbol{L}\tilde{\boldsymbol{x}}_{a,b})(c) = 0$. Thus,

$$\left(\sum_{v \sim c} \boldsymbol{w}(\{v,c\})\right)\tilde{\boldsymbol{x}}_{a,b}(c) - \left(\sum_{v \sim c} \boldsymbol{w}(\{v,c\})\tilde{\boldsymbol{x}}_{a,b}(v)\right) = 0.$$

So, we have,

$$\tilde{\boldsymbol{x}}_{a,b}(c) = \frac{\sum_{v \sim c} \boldsymbol{w}(\{v,c\})\tilde{\boldsymbol{x}}_{a,b}(v)}{\sum_{v \sim c} \boldsymbol{w}(\{v,c\})}.$$

In words, the electrical voltage of c is a weighted average of the voltages of its neighbors. It follows that the voltages of a and b take the largest absolute values. □

Definition 12.11 (Metric)**.** A *metric* on a set S is a function $d : S \times S \to \mathbb{R}$ such that for any $a, b, c \in S$,

1. $d(a,b) = 0 \iff a = b$;
2. $d(a,b) \geq 0$;
3. $d(a,b) = d(b,a)$; and
4. $d(a,b) \leq d(a,c) + d(c,b)$.

Lemma 12.12. *Effective resistance is a metric on V.*

Proof. It is easy to check that properties (1) and (2) are satisfied. We have that property (3) is satisfied by corollary 12.8.

Let us therefore consider property (4), the triangle inequality. We have,

$$\tilde{\boldsymbol{x}}_{a,b} = \boldsymbol{L}^{+}(\mathbf{1}_b - \mathbf{1}_a) = \boldsymbol{L}^{+}(\mathbf{1}_c - \mathbf{1}_a + \mathbf{1}_b - \mathbf{1}_c) = \tilde{\boldsymbol{x}}_{a,c} + \tilde{\boldsymbol{x}}_{c,b},$$

This we can use to rephrase the effective resistance,

$$\begin{aligned} R_{\text{eff}}(a,b) = (\mathbf{1}_b - \mathbf{1}_a)^\top \tilde{\boldsymbol{x}}_{a,b} &= (\mathbf{1}_b - \mathbf{1}_a)^\top (\tilde{\boldsymbol{x}}_{a,c} + \tilde{\boldsymbol{x}}_{c,b}) \\ &= \tilde{\boldsymbol{x}}_{a,c}(b) - \tilde{\boldsymbol{x}}_{a,c}(a) + \tilde{\boldsymbol{x}}_{c,b}(b) - \tilde{\boldsymbol{x}}_{c,b}(a) \\ &\leq \tilde{\boldsymbol{x}}_{a,c}(c) - \tilde{\boldsymbol{x}}_{a,c}(a) + \tilde{\boldsymbol{x}}_{c,b}(b) - \tilde{\boldsymbol{x}}_{c,b}(c) \qquad \text{using lemma 12.10} \\ &= (\mathbf{1}_c - \mathbf{1}_a)^\top \tilde{\boldsymbol{x}}_{a,c} + (\mathbf{1}_b - \mathbf{1}_c)^\top \tilde{\boldsymbol{x}}_{c,b} \\ &= R_{\text{eff}}(a,c) + R_{\text{eff}}(c,b). \end{aligned}$$

□

13

Spectral Graph Sparsification

Many combinatorial graph algorithms perform better on sparse graphs. In this chapter, we will see that for any dense graph, we can find a sparse graph with approximately the same Laplacian matrix as measured by quadratic forms.

Definition 13.1 (Matrix approximation)**.** Given $\boldsymbol{A}, \boldsymbol{B} \in S_+^n$ and $\epsilon > 0$, we say,

$$\boldsymbol{A} \approx_\epsilon \boldsymbol{B} \quad \text{iff} \quad \frac{1}{1+\epsilon}\boldsymbol{A} \preceq \boldsymbol{B}(1+\epsilon)\boldsymbol{A}. \tag{13.1}$$

Given some graph $G = (V, E, \boldsymbol{w})$, our goal is to find a graph $\tilde{G} = (V, \tilde{E}, \tilde{\boldsymbol{w}})$ such that $|\tilde{E}| \ll |E|$ and $\boldsymbol{L}_G \approx_\epsilon \boldsymbol{L}_{\tilde{G}}$. We will write $\boldsymbol{L} \doteq \boldsymbol{L}_G$ and $\tilde{\boldsymbol{L}} \doteq \boldsymbol{L}_{\tilde{G}}$.

Lemma 13.2. *If $\boldsymbol{L} \approx_\epsilon \tilde{\boldsymbol{L}}$, then for any cut $(S, V \setminus S)$,*

$$\frac{1}{1+\epsilon} c_G(S) \leq c_{\tilde{G}}(S) \leq (1+\epsilon) c_G(S). \tag{13.2}$$

Proof. Recall that $c_G(S) = \mathbf{1}_S^\top \boldsymbol{L} \mathbf{1}_S$. The statement follows immediately by comparing the quadratic forms. □

Theorem 13.3 (Spectral graph approximation by sampling)**.** *For any $\epsilon, \delta \in (0,1)$ and sampling probabilities,*

$$p_e \doteq \min\{\alpha \boldsymbol{w}(e) R_{\text{eff}}(e), 1\}, \tag{13.3}$$

of each edge $e \in E$ for some scaling parameter $\alpha > 0$ such that if $e \in \tilde{E}$ with probability p_e and $\tilde{\boldsymbol{w}}(e) \doteq {}^{w(e)}/_{p_e}$, then with probability at least $1 - \delta$ the graph $\tilde{G} = (V, \tilde{E}, \tilde{\boldsymbol{w}})$ satisfies,[1]

$$\boldsymbol{L} \approx_\epsilon \tilde{\boldsymbol{L}} \quad \textit{and} \quad |\tilde{E}| = \mathcal{O}\Big(\frac{n}{\epsilon} \log\Big(\frac{n}{\delta}\Big)\Big).$$

[1] Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011

Proof. TBD □

Remark 13.4. It can be shown that to compute the sampling probabilities p_e, it is sufficient to solve Laplacian linear systems to find sufficiently good approximations to the effective resistances.

14

Solving Laplacian Linear Systems

We have seen by now that many problems can be reduced to solving a Laplacian linear system $\boldsymbol{L}\boldsymbol{x} = \boldsymbol{d}$ where $\boldsymbol{d} \in \operatorname{im} L = (\ker \boldsymbol{L})^{\perp}$. Recall that we can obtain a Cholesky decomposition $\boldsymbol{L} = \mathcal{L}\mathcal{L}^{\top}$ where $\mathcal{L}$ is lower-triangular and positive semi-definite. If $\boldsymbol{L}$ (and therefore $\mathcal{L}$) were invertible, then we have seen that the linear system can be solved in time $\mathcal{O}(\operatorname{nnz} \mathcal{L})$. However, we know that $\boldsymbol{L}$ is not invertible as $\ker \boldsymbol{L} \neq \{\mathbf{0}\}$. It turns out that we can still use of a Cholesky decomposition in solving Laplacian linear systems because we have a simple characterization of the kernel of $\boldsymbol{L}$. This we will discuss first, then we discuss how to efficiently compute the Cholesky decomposition.

Theorem 14.1. *We can solve the linear system $\boldsymbol{L}\boldsymbol{x} = \boldsymbol{d}$ where $\boldsymbol{d} \perp \mathbf{1}$ in time $\mathcal{O}(n^3)$ by first computing a Cholesky decomposition of $\mathbf{L}$ using Gaussian elimination and then applying the pseudoinverse $\boldsymbol{L}^{+}$ to $\boldsymbol{d}$.*

Finally, we will see that we can find approximate solutions in almost linear time.

14.1 Dealing with pseudoinverses

A natural approach is to characterize the pseudoinverse $\boldsymbol{L}^{+}$ in terms of the lower triangular matrix $\mathcal{L}$.

Lemma 14.2. *Given a factorization $\mathbf{L} = \mathcal{L}\mathcal{L}^{\top}$ where $\mathcal{L}$ is lower triangular and all diagonal entries are strictly non-zero except that $\mathcal{L}(n, n) = 0$, we can apply $\boldsymbol{L}^{+}$ in time $\mathcal{O}(n)$.*

Proof. Consider the matrix $\hat{\mathcal{L}}$, which is identical to $\mathcal{L}$ except that $\hat{\mathcal{L}}(n, n) = 1$. Let $\mathcal{D}$ be the diagonal matrix with $\mathcal{D}(i, i) = 1$ for $i < n$ and $\mathcal{D}(n, n) = 0$. Then, $\mathcal{L}\mathcal{L}^{\top} = \hat{\mathcal{L}}\mathcal{D}\hat{\mathcal{L}}^{\top}$ and $\hat{\mathcal{L}}$ is invertible. By

claim 2.35, we have

$$\boldsymbol{L}^{+} = \boldsymbol{\Pi}_{\boldsymbol{L}}(\hat{\mathcal{L}}^{\top})^{-1}\mathcal{D}^{+}\hat{\mathcal{L}}^{-1}\boldsymbol{\Pi}_{\boldsymbol{L}},$$

where $\boldsymbol{\Pi}_{\boldsymbol{L}}$ is the orthogonal projection to the kernel of $\boldsymbol{L}$. Note that $\boldsymbol{\Pi}_{\boldsymbol{L}} = \boldsymbol{I} - \frac{1}{n}\mathbf{1}\mathbf{1}^{\top}$, as this satisfies $\boldsymbol{\Pi}_{\boldsymbol{L}}\boldsymbol{v} = \boldsymbol{v}$ for $\boldsymbol{v} \perp \mathbf{1}$ and $\boldsymbol{\Pi}_{\boldsymbol{L}}\boldsymbol{v} = \mathbf{0}$ for $\boldsymbol{v} \in \operatorname{span}\{\mathbf{1}\}$. We also have $\mathcal{D}^{+} = \mathcal{D}$.

Finally, observe that $\boldsymbol{\Pi}_{\boldsymbol{L}}$ and $\mathcal{D}$ can be applied in time $\mathcal{O}(n)$; and by lemma 2.36, we can apply $\hat{\mathcal{L}}^{-1}$ and $(\hat{\mathcal{L}}^{\top})^{-1}$ in time $\mathcal{O}(\operatorname{nnz}\mathcal{L})$. □

14.2 *Computing the Cholesky Decomposition*

Theorem 14.3 (Cholesky decomposition on graph Laplacians)**.** *Using Gaussian elimination, we can compute in $\mathcal{O}(n^3)$ time a factorization $\boldsymbol{L} = \mathcal{L}\mathcal{L}^{\top}$, where $\mathcal{L}$ is lower triangular and has positive diagonal entries except that $\mathcal{L}(n,n) = 0$.*

Proof. Let $\boldsymbol{L}^{(0)} \doteq \boldsymbol{L}$. For $i \in [n-1]$, we define,

$$\boldsymbol{l}_i \doteq \frac{1}{\sqrt{\boldsymbol{L}^{(i-1)}(i,i)}}\boldsymbol{L}^{(i-1)}(:,i) \quad \text{and}$$
$$\boldsymbol{L}^{(i)} \doteq \boldsymbol{L}^{(i-1)} - \boldsymbol{l}_i\boldsymbol{l}_i^{\top}.$$

Claim 14.4. *Fix some $i < n$. Let $U \doteq \{i+1,\ldots,n\}$. Then, $\boldsymbol{L}^{(i)}(i,j) = 0$ if $i \notin U$ or $j \notin U$; and $\boldsymbol{L}^{(i)}$ is a graph Laplacian matrix of a connected graph on the vertex set U.*

It follows that $\boldsymbol{L}^{(n-1)} = \mathbf{0}_{n\times n}$ because the only graph on a single vertex is the empty graph. From this, we have that $\boldsymbol{L} = \sum_{i=1}^{n-1}\boldsymbol{l}_i\boldsymbol{l}_i^{\top}$, so $\mathcal{L} = [\boldsymbol{l}_1 \cdots \boldsymbol{l}_{n-1}\ \mathbf{0}]$, provided that $\boldsymbol{l}_i$ is well-defined, i.e., $\boldsymbol{L}^{(i-1)}(i,i) \neq 0$ for all $i < n$. But this follows immediately from the claim, as the diagonal of the Laplacian matrix of a connected graph on more than one vertex must be strictly positive (as the degrees must be non-zero).

In each iteration, we compute $\boldsymbol{L}^{(i)}$ in time $\mathcal{O}(n^2)$ and we proceed for $\mathcal{O}(n)$ iterations. □

Proof sketch of claim 14.4. We focus on the first elimination, the remaining are similar. We write,

$$\boldsymbol{L}^{(0)} = \boldsymbol{L} \doteq \begin{bmatrix} w & -\boldsymbol{a}^{\top} \\ -\boldsymbol{a} & \operatorname{diag}(\boldsymbol{a}) + \boldsymbol{L}_{-1} \end{bmatrix},$$

where $\boldsymbol{L}_{-1}$ is defined to make the equality work. We have that,

$$\boldsymbol{l}_1 = \frac{1}{\sqrt{w}}\begin{bmatrix} w \\ -\boldsymbol{a} \end{bmatrix} \quad \text{and} \quad \boldsymbol{l}_1\boldsymbol{l}_1^{\top} = \begin{bmatrix} w & -\boldsymbol{a}^{\top} \\ -\boldsymbol{a} & \frac{1}{w}\boldsymbol{a}\boldsymbol{a}^{\top} \end{bmatrix}$$

Therefore,

$$\boldsymbol{L}^{(1)} = \boldsymbol{L}^{(0)} - \boldsymbol{l}_1 \boldsymbol{l}_1^\top = \begin{bmatrix} 0 & \boldsymbol{0}^\top \\ \boldsymbol{0} & \boldsymbol{S}^{(1)} \end{bmatrix},$$

where $\boldsymbol{S}^{(1)} \doteq \boldsymbol{L}_{+1} + \boldsymbol{L}_{-1}$ and $\boldsymbol{L}_{+1} \doteq \operatorname{diag}(\boldsymbol{a}) - \frac{1}{w}\boldsymbol{a}\boldsymbol{a}^\top$. $\boldsymbol{S}^{(1)}$ is also called the *Schur complement*. We can also phrase it differently as,

$$\begin{aligned} \boldsymbol{l}_1 \boldsymbol{l}_1^\top &= \begin{bmatrix} w & -\boldsymbol{a}^\top \\ -\boldsymbol{a} & \operatorname{diag}(\boldsymbol{a}) \end{bmatrix} - \begin{bmatrix} 0 & \boldsymbol{0}^\top \\ \boldsymbol{0} & \operatorname{diag}(\boldsymbol{a}) - \frac{1}{w}\boldsymbol{a}\boldsymbol{a}^\top \end{bmatrix} && (14.1) \\ &\doteq \textsc{Star}(1, \boldsymbol{L}) - \textsc{Clique}(1, \boldsymbol{L}). && (14.2) \end{aligned}$$

TBD
Figure 14.1: $\boldsymbol{S}^{(1)}$ is the Laplacian matrix of the graph, where the first vertex was removed and all of its neighbors are made to be in a clique.

It remains to show that $\boldsymbol{S}^{(1)}$ is the Laplacian matrix of a connected graph on the vertex set $\{2, \ldots, n\}$.

Observe that the sum of two Laplacian matrices is a again a graph Laplacian. Therefore, $\boldsymbol{L}_{-1}$ is a graph Laplacian by definition. It is easy to check that by the characterization of exercise 1.12, $\boldsymbol{L}_{+1}$ also is a graph Laplacian, which represents a clique formed by the neighbors of vertex 1. Hence, $\boldsymbol{S}^{(1)} = \boldsymbol{L}_{+1} + \boldsymbol{L}_{-1}$ is a graph Laplacian.

It remains to show that the underlying graph is connected. Consider any $i, j \in V \setminus \{1\}$. There exists an i-j path P in the graph of $\boldsymbol{L}$. If P does not use vertex 1, then P is a path in the graph of $\boldsymbol{L}_{-1}$. If P does use vertex 1, it does so using some edges $(u, 1)$ and $(1, v)$. Replacing the two edges with the edge (u, v), which appears in $\boldsymbol{L}_{+1}$ as $\boldsymbol{L}_{+1}(u, v) < 0$, yields a path in the graph of $\boldsymbol{S}^{(1)}$. □

14.3 *Approximate Almost Linear-Time Solvers*

Definition 14.5 (Approximate solution to linear system)**.** We say that $\tilde{\boldsymbol{x}}$ is an *ϵ-approximate solution* to the linear system $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ iff

$$\|\tilde{\boldsymbol{x}} - \boldsymbol{x}^*\|_{\boldsymbol{A}}^2 \leq \epsilon \|\boldsymbol{x}^*\|_{\boldsymbol{A}}^2, \tag{14.3}$$

where $\boldsymbol{x}^*$ is a solution and $\|\boldsymbol{y}\|_{\boldsymbol{A}} \doteq \sqrt{\boldsymbol{y}^\top \boldsymbol{A}\boldsymbol{y}}$ denotes the *Mahalanobis norm* with respect to $\boldsymbol{A}$.

The fast algorithm has two main steps.

Theorem 14.6 (Approximate Cholesky decomposition on graph Laplacians)**.** *We can find $\mathcal{L}\mathcal{L}^\top \approx_{1/2} \boldsymbol{L}$ such that $\mathcal{L}$ is lower triangular and* $\operatorname{nnz} \mathcal{L} = \mathcal{O}\left(m \log^3 n\right)$*, with probability at least $1 - 3/n^5$ in time $\mathcal{O}\left(m \log^3 n\right)$.*[1]

Proof. TBD □

[1] Rasmus Kyng and Sushant Sachdeva. Approximate gaussian elimination for laplacians-fast, sparse, and simple. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 573–582. IEEE, 2016

Here, the main idea is to replace the clique of the neighbors of the eliminated vertex in each round of Gaussian elimination by a sparse approximation.

Theorem 14.7. *We can find an ϵ-approximate solution $\tilde{\boldsymbol{x}}$ to $\boldsymbol{L}\boldsymbol{x} = \boldsymbol{d}$, using an algorithm that takes $\mathcal{O}\left(m \log^3 n \log(1/\epsilon)\right)$ time and succeeds with probability at least $1 - 1/n^{10}$.*

Proof. TBD □

PART IV

Combinatorial Graph Algorithms

15

Algorithms for Maximum Flow

In this chapter, we study classical combinatorial algorithms for the maximum flow problem. Many of the algorithms can be adapted to solve the more general minimum-cost flow problem. In recent years, significant progress was made by using convex optimization and interior point methods. To begin with, we will derive a combinatorial proof for the strong duality of maximum flow and minimum cut.[1]

[1] We have already seen a proof of this in section 8.5.

We consider directed graphs $G = (V, E, \boldsymbol{c})$ with edge capacities $\boldsymbol{c} \in \mathbb{R}^{|E|}$. Recall that a *flow* is a vector $\boldsymbol{f} \in \mathbb{R}^{|E|}$ and $\boldsymbol{f}$ routes demands $\boldsymbol{d}$ iff $\boldsymbol{B}\boldsymbol{f} = \boldsymbol{d}$. We say that $\boldsymbol{f}$ is *feasible* iff $\mathbf{0} \leq \boldsymbol{f} \leq \boldsymbol{c}$, where $\mathbf{0} \leq \boldsymbol{f}$ ensures that the flow respects edge directions and $\boldsymbol{f} \leq \boldsymbol{c}$ ensures that the flow respects edge capacities.

We call a flow $\boldsymbol{f}$ that routes demands $F(\mathbf{1}_t - \mathbf{1}_s)$ for some $F \in \mathbb{R}$ and vertices $s, t \in V$, so F units of flow from s to t, an *s-t flow* with *value* $\mathrm{val}(\boldsymbol{f}) = F$. We say that $\boldsymbol{f}$ is *optimal* iff there is no feasible *s-t* flow $\boldsymbol{f}'$ such that $\mathrm{val}(\boldsymbol{f}') > \mathrm{val}(\boldsymbol{f})$.

We can decompose any *s-t* flow into two kinds of flows: path flows and cycle flows.

Definition 15.1 (Path flow)**.** An *s-t path flow* $\boldsymbol{f}$ is an *s-t* flow with $\mathrm{val}(\boldsymbol{f}) = \alpha$ for some $\alpha > 0$ that can be expressed as,

$$\boldsymbol{f} = \alpha \sum_{e \in P} \mathbf{1}_e, \tag{15.1}$$

for some *s-t* path P.

TBD
Figure 15.1: Example of an *s-t* path flow.

Definition 15.2 (Cycle flow)**.** A *cycle flow* $\boldsymbol{f}$ is a flow routing demands $\mathbf{0}$ that can be expressed as,

$$\boldsymbol{f} = \alpha \sum_{e \in C} \mathbf{1}_e, \tag{15.2}$$

for some $\alpha > 0$, and cycle C.

TBD
Figure 15.2: Example of a cycle flow.

Lemma 15.3 (Path-cycle decomposition)**.** *Any s-t flow $\boldsymbol{f} \geq \mathbf{0}$ can be decomposed into $k \leq m$ s-t path flows and l cycle flows.*

Proof. TBD □

Lemma 15.4. *There exists an optimal flow with a path-cycle decomposition that has only paths and no cycles.*

Proof. TBD □

Lemma 15.5. *There exists an s-t flow $\boldsymbol{f} \geq \mathbf{0}$ iff there exists a directed s-t path.*

Proof. TBD □

Recall that a *cut* is a proper subset of vertices $\emptyset \subset S \subset V$. An *s-t cut* is a cut $(S, V \setminus S)$ separating s and t, i.e., $s \in S, t \in V \setminus S$. We say that the capacity of a cut is the sum of capacities of crossing edges,

$$\operatorname{cap}(S) \doteq \sum_{\substack{\{a,b\} \in E \\ a \in S,\ b \in V \setminus S}} \boldsymbol{c}(\{a,b\}). \tag{15.3}$$

Theorem 15.6 (Weak duality of maximum flow/minimum cut)**.** *For any feasible s-t flow $\boldsymbol{f}$ and any s-t cut $(S, V \setminus S)$,*

$$\operatorname{val}(\boldsymbol{f}) \leq \operatorname{cap}(S). \tag{15.4}$$

Proof. Let $(S, V \setminus S)$ be any *s-t* cut. Suppose for a contradiction that $\operatorname{val}(\boldsymbol{f}) > \operatorname{cap}(S)$ for some flow $\boldsymbol{f}$. But this contradicts feasibility of $\boldsymbol{f}$ because the crossing edges of the cut form a bottleneck. □

An important concept in the analysis of flow algorithms is the so-called residual graph.

Definition 15.7 (Residual graph)**.** The *residual graph* $G_{\boldsymbol{f}}$ of some *s-t* flow $\boldsymbol{f} \geq \mathbf{0}$ is the graph G with edge capacities $[-\boldsymbol{f}, \boldsymbol{c} - \boldsymbol{f}]$. That is, we say that a flow $\hat{\boldsymbol{f}}$ is *feasible* in the residual graph iff $-\boldsymbol{f} \leq \hat{\boldsymbol{f}} \leq \boldsymbol{c} - \boldsymbol{f}$.

Intuitively, sending positive flow $\boldsymbol{c} - \boldsymbol{f}$ along an edge in $G_{\boldsymbol{f}}$ corresponds to sending the maximum additional flow without violating the capacity constraint within G, whereas sending negative flow $-\boldsymbol{f}$ along an edge in $G_{\boldsymbol{f}}$ corresponds to "undoing" the flow that was send along this edge by $\boldsymbol{f}$. This simple argument shows that if $\hat{\boldsymbol{f}}$ is feasible in $G_{\boldsymbol{f}}$, then $\boldsymbol{f} + \hat{\boldsymbol{f}}$ is feasible in G.

We call an *s-t* flow $\hat{\boldsymbol{f}}$ in $G_{\boldsymbol{f}}$ an *augmenting flow* of $\boldsymbol{f}$.

Lemma 15.8 (Flow optimality condition)**.** *A feasible s-t flow f in G is optimal iff there is no s-t path in G_f, or equivalently, iff there is no f-augmenting flow.*

Proof. TBD □

Theorem 15.9 (Strong duality of maximum flow/minimum cut)**.** *We have that,*

$$\max_{\substack{F\in\mathbb{R},\ f\in\mathbb{R}^{|E|}\\ Bf=F(\mathbf{1}_t-\mathbf{1}_s)}} F = \min_{s\text{-}t \text{ cut } (S,V\setminus S)} \operatorname{cap}(S). \tag{15.5}$$

Proof. By weak duality of maximum flow and minimum cut, we have the direction $\leq$. For the direction $\geq$, let f^* be an optimal flow and consider the cut,

$$S \doteq \{v \in V \mid \text{there exists an } s\text{-}v \text{ path in } G_{f^*}\}.$$

We make two observations.

1. By definition, there are no edges from S to $V \setminus S$ in G_{f^*}, that is, f^* *saturates*[2] all crossing edges of the cut.
2. By definition, f^* routes no flow from $V \setminus S$ to S.

This implies that $\operatorname{val}(f^*) \geq \operatorname{cap}(S)$ for this cut $(S, V \setminus S)$. □

[2] That is, f^* sends flow equal to the capacity of the edge.

15.1 The Ford-Fulkerson Algorithm

Our prior discussion gives rise to a very natural algorithm.

Algorithm 15.10: FORDFULKERSON(G)

```
f ← 0
while there exists any s-t path flow f̂ in G_f do
    f ← f + f̂
return f
```

Given a feasible flow f, we can find an f-augmenting flow, or determine that none exists, in time $\mathcal{O}(m)$ using breadth-first search or depth-first search.

Theorem 15.11 (Ford-Fulkerson)**.** *If capacities are integral,* FORDFULKERSON *converges to an optimal flow f^* in $\operatorname{val}(f^*)$ iterations.*

Proof. Observe that the initial flow $\mathbf{0}$ is trivially feasible. In each iteration, we add the augmenting flow $\hat{f}$ with $\operatorname{val}(\hat{f}) > 0$, and due to

the integral capacities, $\mathrm{val}(\hat{f}) \geq 1$. Therefore, the flow value $\mathrm{val}(f)$ increases in each iteration by at least one. □

Improving Ford-Fulkerson

It turns out that if we always choose the shortest augmenting path, we converge in time $\mathcal{O}(nm^2)$. This is known as the *Edmonds-Karp algorithm*.

We can do still better, by choosing that path with the maximum bottleneck capacity. That is, we choose,

$$P^* = \operatorname*{arg\,max}_{\text{augmenting paths } P} \min_{e \in P} c(e). \tag{15.6}$$

Within the framework of Ford-Fulkerson this corresponds to the augmenting path that allows us to route the most additional flow.

Theorem 15.12. FordFulkerson, *where in each iteration we choose the augmenting path with maximum bottleneck capacity, converges in time* $\mathcal{O}(m^2 \log mU)$.

Proof. We can find P^* using a binary search on $[1, U]$, where $U \doteq \max_e c(e)$, by removing all edges with absolute capacity in G_f below the current threshold and testing if an s-t path in G_f exists: if it does, we increase the threshold; if it does not, we decrease the threshold. This procedure takes $\mathcal{O}(m \log U)$ time. If we only consider the occurring capacities, the runtime improves to $\mathcal{O}(m \log m)$.

Suppose $\hat{F}$ is the flow left in G_f. By the path decomposition lemma, this flow can be decomposed into at most m path flows (the "best" of which is P^*) and P^* must route at least the average amount of flow. Hence, P^* routes at least $\hat{F}/m$ units of flow. Thus, the algorithm converged if,

$$\left(1 - \frac{1}{m}\right)^T F^* < 1,$$

where T is the number of augmentations and F^* is the value of an optimal flow. So, some $T = \mathcal{O}(m \log F^*)$ is sufficient.

Overall, we get,

$$\mathcal{O}(m \log m \cdot T) = \mathcal{O}\Big(m^2 \log(m + F^*)\Big) = \mathcal{O}\Big(m^2 \log mU\Big) \qquad \square \qquad \text{using } F^* \leq mU$$

15.2 Dinitz's Algorithm

We will now look at Dinitz's algorithm, which also belongs to the family of Ford-Folkerson algorithms. That is, we still start with the

initial feasible flow $f \leftarrow \mathbf{0}$ and iteratively improve this flow by finding an augmenting flow in G_f.

In Dinitz's algorithm, our strategy to finding an augmenting path is to find so-called *blocking flows* in G_f, which in each iteration "block" one shortest s-t path in G_f. Dinitz's algorithm runs in $\mathcal{O}\left(n^2 + nm \log^2 n\right)$ time in general graphs and in time $\mathcal{O}\left(\min\{m^{3/2}, mn^{2/3}\}\right)$ in unit capacity graphs.

Remark 15.13. It can also be shown that Dinitz's algorithm converges in time $\mathcal{O}(m\sqrt{n})$ in bipartite matching graphs, but we will not show this here.

Definition 15.14 (Level)**.** The *level* $\ell(v)$ of a vertex $v \in V$ in G_f is the length (i.e., number of edges) of the shortest s-v path in G_f.

We call an edge $e = (u, v)$ in G_f *admissible* iff $\ell(v) = \ell(u) + 1$. Intuitively, e is admissible iff it is part of a shortest s-u path.

The *level graph* L_f of G_f is the subgraph of only admissible edges.

TBD
Figure 15.3: Illustration of admissible edges.

Definition 15.15 (Blocking flow)**.** A *blocking flow* in G_f is a flow $\hat{f}$ such that

1. $\hat{f}$ is feasible;
2. $\hat{f}$ uses only admissible edges; and
3. for every s-t path in the level graph, $\hat{f}$ saturates at least one edge.

Remark 15.16. By definition, a blocking flow in G_f is a f-augmenting flow.

TBD
Figure 15.4: Examples of blocking flows.

Algorithm 15.17: DINITZ(G)

```
f ← 0
repeat
    f̂ ← blocking flow in G_f
    f ← f + f̂
until G_f is disconnected
return f
```

Lemma 15.18. *Let f be a feasible flow, $\hat{f}$ a blocking flow in G_f, and let $f' \doteq f + \hat{f}$. Then $\ell_{G_{f'}}(t) > \ell_{G_f}(t)$.*

Proof. TBD □

Theorem 15.19 (Dinitz)**.** *Dinitz's algorithm converges in $\mathcal{O}(n)$ iterations.*

Proof. By the previous lemma, $\ell(t)$ increases in every iteration by at least one. As path can contain each vertex at most once, the level of any vertex can never be larger than n. □

Finding Blocking Flows

A naïve approach is to repeatedly use depth-first search to find an unsaturated s-t path in the level graph L_f. Whenever we find such a path, we route the maximum possible flow along this path, saturating at least one of its edges.

Algorithm 15.20: FINDBLOCKINGFLOW(L_f)

```
1 f̂ ← 0
2 while there exists an s-t path P in L_f do
3     Let f̂' be a flow saturating P
4     f̂ ← f̂ + f̂'
5     Remove all edges saturated by f̂' from L_f
6 return f̂
```

Lemma 15.21. FINDBLOCKINGFLOW *returns a blocking flow in* $\mathcal{O}(mn)$ *time.*

Remark 15.22. This can be improved to $\mathcal{O}\left(m \log^2 n + n\right)$ by using link-cut trees, which we will discuss in the next chapter.

Proof. TBD □

Unit Capacity Graphs

Lemma 15.23. *In unit capacity graphs,* FINDBLOCKINGFLOW *returns a blocking flow in* $\mathcal{O}(m)$ *time.*

Proof. TBD □

Theorem 15.24. *In unit capacity graphs, Dinic's algorithm terminates in* $\mathcal{O}\left(\min\{m^{1/2}, n^{2/3}\}\right)$ *time.*

Proof. TBD □

15.3 The Push-Relabel Algorithm

15.4 Outlook

We have seen two approaches to solving maximum flow problems: Ford-Fulkerson maintains a feasible flow and augments this flow until it is optimal. In contrast, Push-Relabel maintains that there is no augmenting path and terminates when the flow is feasible.

Currently, the best known algorithm for real-valued capacities is due to Orlin and takes $\mathcal{O}(mn)$ time.[3]

Geometrically speaking, Ford-Fulkerson is analogous to a simplex method and Push-Relabel is analogous to an exterior-point method. In recent years, interior-point methods were used to find efficient algorithms when capacities are integral, the best known algorithm taking $\mathcal{O}\left(m^{1+o(1)} \log U\right)$ time.[4]

[3] James B Orlin. Max flows in o (nm) time, or better. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 765–774, 2013

[4] Li Chen, Rasmus Kyng, Yang P Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. *arXiv preprint arXiv:2203.00671*, 2022

16

Link-Cut Trees

17
Finding Expanders using Maximum Flow

18

Distance Oracles

PART V

Further Topics

19

Interior Point Methods for Maximum Flow

A

Solutions

A.1 Part I

Electrical Flows

Solution to exercise 1.12. TBD □

Proof of claim 1.15. We have that $\boldsymbol{H}_c(\boldsymbol{x}) = \boldsymbol{L}(\boldsymbol{x}) \succeq \boldsymbol{0}$. □

Solution to exercise 1.17. Consider any $\boldsymbol{f} \in \mathbb{R}^{|E|}$ satisfying $\boldsymbol{B}\boldsymbol{f} = \boldsymbol{d}$. We have for any $\boldsymbol{x} \in \mathbb{R}^{|V|}$,

$$\begin{aligned}\frac{1}{2}\boldsymbol{f}^\top \boldsymbol{R}\boldsymbol{f} &= \frac{1}{2}\boldsymbol{f}^\top \boldsymbol{R}\boldsymbol{f} - \boldsymbol{x}^\top \underbrace{(\boldsymbol{B}\boldsymbol{f} - \boldsymbol{d})}_{0} \\ &\geq \min_{\boldsymbol{f}' \in \mathbb{R}^{|E|}} \underbrace{\frac{1}{2}\boldsymbol{f}'^\top \boldsymbol{R}\boldsymbol{f}' - \boldsymbol{x}^\top(\boldsymbol{B}\boldsymbol{f}' - \boldsymbol{d})}_{\doteq g(\boldsymbol{f}')}.\end{aligned}$$

Note that g is convex. Taking the gradient gives,

$$\boldsymbol{\nabla}_{\boldsymbol{f}'} g(\boldsymbol{f}') = \boldsymbol{R}\boldsymbol{f}' - \boldsymbol{B}^\top \boldsymbol{x}.$$

So $\boldsymbol{f}' = \boldsymbol{R}^{-1}\boldsymbol{B}^\top \boldsymbol{x}$ is the minimizer of g. We obtain,

$$\frac{1}{2}\boldsymbol{f}^\top \boldsymbol{R}\boldsymbol{f} \geq -\frac{1}{2}\boldsymbol{x}^\top \boldsymbol{L}\boldsymbol{x} + \boldsymbol{d}^\top \boldsymbol{x},$$

but $\tilde{\boldsymbol{f}}^\top \boldsymbol{R}\tilde{\boldsymbol{f}} = \tilde{\boldsymbol{x}}^\top \boldsymbol{L}\tilde{\boldsymbol{x}} = \boldsymbol{d}^\top \tilde{\boldsymbol{x}}$ for electrical voltages $\tilde{\boldsymbol{x}}$ and electrical flow $\tilde{\boldsymbol{f}}$ (see eq. (1.18)). Thus,

$$\frac{1}{2}\boldsymbol{f}^\top \boldsymbol{R}\boldsymbol{f} \geq \frac{1}{2}\tilde{\boldsymbol{x}}^\top \boldsymbol{L}\tilde{\boldsymbol{x}} = \frac{1}{2}\tilde{\boldsymbol{f}}^\top \boldsymbol{R}\tilde{\boldsymbol{f}}.$$

Hence, $\tilde{\boldsymbol{f}}$ is the minimum electrical energy flow among all flows routing $\boldsymbol{d}$. □

Solution to exercise 1.18. TBD □

Solution to exercise 1.19. TBD □

Linear Algebra

Proof of claim 2.1. Because the characteristic polynomial of $\boldsymbol{A}$ is of degree n, it has n complex roots, which are the eigenvalues $\lambda_1, \ldots, \lambda_n$ of $\boldsymbol{A}$. We will first prove that the λ_i are real. Then, we will prove that the corresponding eigenvectors $\boldsymbol{v}_i$ are orthogonal.

1. Let λ be any eigenvalue of $\boldsymbol{A}$. We denote by $\bar{\lambda}$ the complex conjugate of λ. Clearly, if $\lambda = \bar{\lambda}$, then $\lambda \in \mathbb{R}$. By the definition of the eigenvalue λ with associated eigenvector $\boldsymbol{v}$, we have,

$$\lambda \bar{\boldsymbol{v}}^\top \boldsymbol{v} = \bar{\boldsymbol{v}}^\top \boldsymbol{A} \boldsymbol{v}.$$

   Taking the complex conjugate and transpose of both sides gives,

$$\bar{\lambda} \bar{\boldsymbol{v}}^\top \boldsymbol{v} = \bar{\boldsymbol{v}}^\top \bar{\boldsymbol{A}}^\top \boldsymbol{v} = \bar{\boldsymbol{v}}^\top \boldsymbol{A} \boldsymbol{v} = \lambda \bar{\boldsymbol{v}}^\top \boldsymbol{v}.$$

   using that $\boldsymbol{A}$ is real and symmetric, $\bar{\boldsymbol{A}}^\top = \boldsymbol{A}$

   We have $\lambda = \bar{\lambda}$ as desired.
2. It remains to show that for eigenvalues λ_i, λ_j with associated eigenvectors $\boldsymbol{v}_i, \boldsymbol{v}_j$ and $i \neq j$, we have $\boldsymbol{v}_i^\top \boldsymbol{v}_j = 0$. By the definition of an eigenvalue, we have,

$$\begin{aligned} \lambda_i \boldsymbol{v}_j^\top \boldsymbol{v}_i = \boldsymbol{v}_j^\top \boldsymbol{A} \boldsymbol{v}_i &= (\boldsymbol{v}_i^\top \boldsymbol{A}^\top \boldsymbol{v}_j)^\top \\ &= (\boldsymbol{v}_i^\top \boldsymbol{A} \boldsymbol{v}_j)^\top = \lambda_j (\boldsymbol{v}_i^\top \boldsymbol{v}_j)^\top = \lambda_j \boldsymbol{v}_j^\top \boldsymbol{v}_i. \end{aligned}$$

   using that $\boldsymbol{A}$ is symmetric

   We get that $\boldsymbol{v}_j^\top \boldsymbol{v}_i = 0$ if $\lambda_i \neq \lambda_j$. □

Proof of claim 2.7. Suppose λ is an eigenvalue of $\boldsymbol{M}$ with associated eigenvector $\boldsymbol{v}$. Also suppose that $\boldsymbol{T}\boldsymbol{v} = \boldsymbol{w}$. Then,

$$\boldsymbol{T}\boldsymbol{M}\boldsymbol{T}^{-1}\boldsymbol{w} = \boldsymbol{T}\boldsymbol{M}\boldsymbol{v} = \lambda \boldsymbol{T}\boldsymbol{v} = \boldsymbol{w}.$$

It is easy to check that the other direction holds too. □

Proof of claim 2.26. TBD □

Proof of claim 2.33. TBD □

Proof of claim 2.35. TBD □

Probability

Theorem A.1 (Jensen's inequality, finite form)**.** *Let $f : S \to \mathbb{R}$ be a convex function on the convex set $S \subseteq \mathbb{R}^n$. Suppose that $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k S$ and $\theta_1, \ldots, \theta_k \geq 0$ with $\theta_1 + \cdots + \theta_k = 1$. Then,*

$$f(\theta_1 \boldsymbol{x}_1 + \cdots + \theta_k \boldsymbol{x}_k) \leq \theta_1 f(\boldsymbol{x}_1) + \cdots + \theta_k f(\boldsymbol{x}_k). \tag{A.1}$$

Proof. We prove the statement by induction on k. The base case, $k = 2$, follows trivially from the convexity of f. For the induction step, suppose that the statement holds for some $k \geq 2$. Assume w.l.o.g. that $\theta_{k+1} \in (0, 1)$. We have,

$$\begin{aligned}
\sum_{i=1}^{k+1} \theta_i f(\boldsymbol{x}_i) &= (1 - \theta_{k+1}) \left(\sum_{i=1}^{k} \frac{\theta_i}{1 - \theta_{k+1}} f(\boldsymbol{x}_i) \right) + \theta_{k+1} f(\boldsymbol{x}_{k+1}) \\
&\geq (1 - \theta_{k+1}) f\left(\sum_{i=1}^{k} \frac{\theta_i}{1 - \theta_{k+1}} \boldsymbol{x}_i \right) + \theta_{k+1} f(\boldsymbol{x}_{k+1}) && \text{using the induction hypothesis} \\
&\geq f\left(\sum_{i=1}^{k+1} \theta_i \boldsymbol{x}_i \right). && \text{using convexity of } f
\end{aligned}$$

□

A.2 Part II

Convex Geometry

Proof of theorem 5.3. In our proof, we will use the following two theorems.

Fact A.2 (Bolzano-Weierstrass theorem)**.** *Every bounded sequence in $\mathbb{R}^n$ has a convergent subsequence.*

Fact A.3 (Boundedness theorem)**.** *Let $f : \mathbb{R}^n \to \mathbb{R}$ be a continuous function and $\mathcal{F} \subseteq \mathbb{R}^n$ be non-empty, bounded, and closed. Then f is bounded on $\mathcal{F}$.*

Let α be the infimum of f over $\mathcal{F}$, i.e., the largest value for which any $\boldsymbol{x} \in \mathcal{F}$ satisfies $f(\boldsymbol{x}) \geq \alpha$. By the boundedness theorem, the infumum exists, as f is lower bounded and the set of lower bounds has a greatest lower bound, α.

Let $\mathcal{F}_k \doteq \{\boldsymbol{x} \in \mathcal{F} \mid \alpha \leq f(\boldsymbol{x}) \leq \alpha + 2^{-k}\}$. $\mathcal{F}_k$ cannot be empty, since if it were, then $\alpha + 2^{-k}$ would be a strictly greater lower bound on f than α. For each k, let $\boldsymbol{x}_k$ be some $\boldsymbol{x} \in \mathcal{F}_k$. $\{\boldsymbol{x}_k\}_{k=1}^{\infty}$ is a bounded sequence as $\mathcal{F}_k \subseteq \mathcal{F}$, so by the Bolzano-Weierstrass theorem, there exists a convergent subsequence, $\{\boldsymbol{y}_k\}_{k=1}^{\infty}$, with limit $\bar{\boldsymbol{y}}$. Because the

set is closed, $\bar{\boldsymbol{y}} \in \mathcal{F}$. By continuity, $f(\bar{\boldsymbol{y}}) = \lim_{k\to\infty} f(\boldsymbol{y}_k)$, and by construction, $\lim_{k\to\infty} f(\boldsymbol{y}_k) = \alpha$.

Thus, the optimal solution is $\bar{\boldsymbol{y}}$. □

Solution to exercise 5.10. TBD □

Solution to exercise 5.12. TBD □

A.3 Part III

Introduction to Spectral Graph Theory

Solution to exercise 10.9. TBD □

Solution to exercise 10.10. TBD □

Solution to exercise 10.11. TBD □

Solution to exercise 10.12. TBD □

Conductance and Expanders

Solution to exercise 11.7. TBD □

Solution to exercise 11.8. TBD □

Summary of Notation

We follow these general rules:

- uppercase italic for constants N
- lowercase italic for indices i and scalar variables x
- lowercase italic bold for vectors $\boldsymbol{x}$
- uppercase italic bold for matrices $\boldsymbol{M}$
- uppercase italic for random variables X
- uppercase bold for random vectors $\mathbf{X}$
- uppercase italic for sets A

$\doteq$	equality by definition
iff	if and only if
$\mathbb{N}$	set of natural numbers $\{1, 2, \dots\}$
$\mathbb{N}_0$	set of natural numbers, including 0, $\mathbb{N} \cup \{0\}$
$\mathbb{R}$	set of real numbers
$\mathbb{C}$	set of complex numbers
$[m]$	set of natural numbers from 1 to m, $\{1, 2, \dots, m-1, m\}$
$(a, b]$	real interval between a and b including b but not including a
$f : A \to: B$	function f from elements of set A to elements of set B
$\mathbb{1}\{predicate\}$	indicator function ($\mathbb{1}\{predicate\} \doteq 1$ if the *predicate* is true, else 0)
$\leftarrow$	assignment

LINEAR ALGEBRA

S^n	set of symmetric $n \times n$ matrices
S^n_+	set of symmetric and positive semi-definite $n \times n$ matrices
S^n_{++}	set of symmetric and positive definite $n \times n$ matrices
$\boldsymbol{A} \preceq \boldsymbol{B}$	Loewner order on symmetric matrices, $\forall \boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{x} \leq \boldsymbol{x}^\top \boldsymbol{B} \boldsymbol{x}$
$\boldsymbol{x} \perp \boldsymbol{y}$	$\boldsymbol{x}$ and $\boldsymbol{y}$ are orthogonal, i.e., $\boldsymbol{x}^\top \boldsymbol{y} = 0$
$\boldsymbol{x} \perp W$	$\boldsymbol{x}$ is orthogonal to every vector $\boldsymbol{y}$ in subspace W
$W^\perp$	orthogonal complement of subspace W, $\{\boldsymbol{x} \in \mathbb{R}^n \mid \boldsymbol{x} \perp W\}$
$\text{span}\{\boldsymbol{x}_1, \dots, \boldsymbol{x}_n\}$	smallest subspace containing $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n$
$\dim(W)$	number of vectors in a basis of a subspace W
$H(\boldsymbol{n}, \mu)$	hyperplane with normal $\boldsymbol{n}$ and threshold μ

$\mathbf{1}_S$	vector such that $\mathbf{1}_S(i) = \mathbb{1}\{i \in S\}$
$\boldsymbol{A}^\top$	transpose of matrix $\boldsymbol{A}$
$\boldsymbol{A}^{-1}$	inverse of matrix $\boldsymbol{A}$
$\boldsymbol{A}^+$	pseudoinverse of matrix $\boldsymbol{A}$
$\boldsymbol{A}^{1/2}$	square root of symmetric positive semi-definite matrix $\boldsymbol{A}$
$\boldsymbol{\Pi}_{\boldsymbol{A}}$	orthogonal projection to $(\ker \boldsymbol{A})^\perp$
$\operatorname{nnz} \boldsymbol{A}$	number of non-zero entries of $\boldsymbol{A}$
$\operatorname{tr} \boldsymbol{A}$	trace of $\boldsymbol{A}$, $\sum_i \boldsymbol{A}(i,i)$
$\operatorname{diag}_{i \in I}\{a_i\}$	diagonal matrix with elements a_i, indexed according to the set I
$\ker \boldsymbol{A}$	kernel (or null space) of $\boldsymbol{A}$, $\{\boldsymbol{x} \in \mathbb{R}^n \mid \boldsymbol{A}\boldsymbol{x} = \mathbf{0}\}$
$\operatorname{im} \boldsymbol{A}$	image of $\boldsymbol{A}$, $\operatorname{span}\{\boldsymbol{A}(:,i)\}_{i \in [n]}$
$\lambda_i(\boldsymbol{A})$	i-th smallest eigenvalue of $\boldsymbol{A}$
$\|\boldsymbol{A}\|_{\alpha \to \beta}$	matrix norm of $\boldsymbol{A}$ induced by norms $\|\cdot\|_\alpha$ and $\|\cdot\|_\beta$

PROBABILITY

$\mathbb{P}[X = x]$	probability of a random variable X taking on the value x
$X \sim F$	random variable X follows the distribution F
$x \sim F$	value x is sampled according to distribution F
$X \perp Y$	random variable X is independent of random variable Y
$X \perp Y \mid Z$	random variable X is conditionally independent of random variable Y given random variable Z
$\mathbb{E}[X]$	expected value of random variable X
$\operatorname{Var}[X]$	variance of random variable X
$\boldsymbol{W} \in \mathbb{R}^{\lvert V\rvert \times \lvert V\rvert}$	transition matrix of random walk, $\boldsymbol{A}\boldsymbol{D}^{-1} = \boldsymbol{I} - \boldsymbol{D}^{1/2}\boldsymbol{N}\boldsymbol{D}^{-1/2}$
$\tilde{\boldsymbol{W}} \in \mathbb{R}^{\lvert V\rvert \times \lvert V\rvert}$	transition matrix of lazy random walk, $\frac{\boldsymbol{I}}{2} + \frac{\boldsymbol{W}}{2} = \boldsymbol{I} - \frac{1}{2}\boldsymbol{D}^{1/2}\boldsymbol{N}\boldsymbol{D}^{-1/2}$
$\boldsymbol{p}_t \in \mathbb{R}^{\lvert V\rvert}$	probability distribution of a random walk at time t, $\boldsymbol{W}^t \boldsymbol{p}_0$
$H_{a,s}$	hitting time of s starting from a
$\boldsymbol{h}_s \in \mathbb{R}^{\lvert V\rvert}$	expected hitting times of s, $\boldsymbol{h}_s(a) = \mathbb{E}[H_{a,s}]$
$C_{a,b}$	commute time between a and b, $H_{a,b} + H_{b,a}$

ANALYSIS

$\boldsymbol{\nabla} f(\boldsymbol{x}) \in \mathbb{R}^{n \times 1}$	gradient of a function f at a point $\boldsymbol{x}$, $\begin{bmatrix} \frac{\partial f(\boldsymbol{x})}{\partial \boldsymbol{x}(1)} & \cdots & \frac{\partial f(\boldsymbol{x})}{\partial \boldsymbol{x}(n)} \end{bmatrix}^\top$
$[\boldsymbol{x}, \boldsymbol{y}]$	set of convex combinations of $\boldsymbol{x}$ and $\boldsymbol{y}$, $\{\theta \boldsymbol{x} + (1-\theta)\boldsymbol{y} \mid \theta \in [0,1]\}$
$Df(\boldsymbol{x})[\boldsymbol{d}]$	directional derivative of f at $\boldsymbol{x}$ in direction $\boldsymbol{d}$, $\lim_{\lambda \to 0} \frac{f(\boldsymbol{x}+\lambda \boldsymbol{d}) - f(\boldsymbol{x})}{\lambda}$
$D\boldsymbol{g}(\boldsymbol{x}) \in \mathbb{R}^{m \times n}$	Jacobian of vector-valued function $\boldsymbol{g} : \mathbb{R}^n \to \mathbb{R}^m$, $\begin{bmatrix} \frac{\partial \boldsymbol{g}(\boldsymbol{x})}{\partial \boldsymbol{x}(1)} & \cdots & \frac{\partial \boldsymbol{g}(\boldsymbol{x})}{\partial \boldsymbol{x}(n)} \end{bmatrix}$
$\boldsymbol{H}_f(\boldsymbol{x}) \in \mathbb{R}^{n \times n}$	Hessian of f, $(D\boldsymbol{\nabla} f(\boldsymbol{x}))^\top$
$\operatorname{epi}(f)$	epigraph of a function f, $\{(\boldsymbol{x}, y) \mid f(\boldsymbol{x}) \leq y\} \subseteq \mathbb{R}^{n+1}$

$S_\alpha(f)$	α-sub-level set of a function f, $\{\boldsymbol{x} \in S \mid f(\boldsymbol{x}) \leq \alpha\}$
$L_\alpha(f)$	α-level set of a function f, $\{\boldsymbol{x} \in S \mid f(\boldsymbol{x}) = \alpha\}$

GRAPHS

V	set of vertices
E	set of edges
n	number of vertices, $\lvert V\rvert$
m	number of edges, $\lvert E\rvert$
$G[X]$	subgraph of G induced by $X \subseteq V$
$u \sim v$	vertices u and v are adjacent
$\deg(v)$	degree of vertex v
$\boldsymbol{r} \in \mathbb{R}^{\lvert E\rvert}$	resistances
$\boldsymbol{w} \in \mathbb{R}^{\lvert E\rvert}$	weights, $1/\boldsymbol{r}(e)$
$\boldsymbol{d} \in \mathbb{R}^{\lvert V\rvert}$	weighted degrees, $\sum_{\{u,v\}\in E} \boldsymbol{w}(\{u,v\})$
$\tilde{\boldsymbol{A}} \in \mathbb{R}^{\lvert V\rvert\times\lvert V\rvert}$	adjacency matrix
$\boldsymbol{A} \in \mathbb{R}^{\lvert V\rvert\times\lvert V\rvert}$	weighted adjacency matrix
$\boldsymbol{B} \in \mathbb{R}^{\lvert V\rvert\times\lvert E\rvert}$	incidence matrix
$\boldsymbol{R} \in \mathbb{R}^{\lvert E\rvert\times\lvert E\rvert}$	diagonal matrix of resistances, $\operatorname{diag}\{\boldsymbol{r}(e)\}_{e\in E}$
$\boldsymbol{W} \in \mathbb{R}^{\lvert E\rvert\times\lvert E\rvert}$	diagonal matrix of weights, $\operatorname{diag}\{\boldsymbol{w}(e)\}_{e\in E}$
$\boldsymbol{D} \in \mathbb{R}^{\lvert V\rvert\times\lvert V\rvert}$	diagonal matrix of weighted degrees, $\operatorname{diag}\{\boldsymbol{w}(v)\}_{v\in V}$
$\boldsymbol{L} \in \mathbb{R}^{\lvert V\rvert\times\lvert V\rvert}$	Laplacian matrix, $\boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{B}^\top = \boldsymbol{B}\boldsymbol{W}\boldsymbol{B}^\top = \boldsymbol{D} - \boldsymbol{A}$
$\operatorname{vol}(S)$	volume of a set of vertices S, $\sum_{v\in S} \boldsymbol{d}(v) = \mathbf{1}_S^\top \boldsymbol{D}\mathbf{1}_S$
$c(S)$	value of a cut S, $\sum_{\{a,b\}\in E:\ a\in S,\ b\in V\setminus S} \boldsymbol{w}(\{a,b\}) = \mathbf{1}_S^\top \boldsymbol{L}\mathbf{1}_S$
$\phi(S)$	conductance of a cut S, $\frac{c(S)}{\min\{\operatorname{vol}(S),\operatorname{vol}(V\setminus S)\}}$
$\phi(G)$	conductance of a graph G, $\min_{\emptyset\subset S\subset V} \phi(S)$
$\psi(S)$	sparsity of a cut S, $\frac{c(S)}{\min\{\lvert S\rvert,\lvert V\setminus S\rvert\}}$
$\psi(G)$	sparsity of a graph G, $\min_{\emptyset\subset S\subset V} \psi(S)$
K_n	unit weight complete graph on n vertices
P_n	unit weight path graph on n vertices
T_d	unit weight complete binary tree with d levels
$G_{i,j}$	unit weight graph with single edge $\{i,j\}$
$G^{i,j}$	subgraph of G consisting of the shortest i,j path

FLOWS

$\boldsymbol{x} \in \mathbb{R}^{\lvert V\rvert}$	voltages
$\boldsymbol{x}(e)$	voltage difference of edge $e = \{u,v\}$, $\boldsymbol{x}(u) - \boldsymbol{x}(v)$
$\boldsymbol{f} \in \mathbb{R}^{\lvert E\rvert}$	flow

$\boldsymbol{d} \in \mathbb{R}^{\lvert V \rvert}$	demands, modeling the net flow
$\tilde{\boldsymbol{x}} \in \mathbb{R}^{\lvert V \rvert}$	electrical voltages
$\tilde{\boldsymbol{x}}_{a,b} \in \mathbb{R}^{\lvert V \rvert}$	electrical voltages routing demands $\mathbf{1}_b - \mathbf{1}_a$
$\tilde{\boldsymbol{f}} \in \mathbb{R}^{\lvert E \rvert}$	electrical flow
$\mathcal{E}(\tilde{\boldsymbol{f}}), \mathcal{E}(\tilde{\boldsymbol{x}}), \mathcal{E}(\boldsymbol{d})$	electrical energy, $\tilde{\boldsymbol{f}}^\top \boldsymbol{B}^\top \tilde{\boldsymbol{x}} = \tilde{\boldsymbol{f}}^\top \boldsymbol{R} \tilde{\boldsymbol{f}} = \tilde{\boldsymbol{x}}^\top \boldsymbol{L} \tilde{\boldsymbol{x}} = \boldsymbol{d}^\top \boldsymbol{x} = \boldsymbol{d}^\top \boldsymbol{L}^+ \boldsymbol{d}$
$\boldsymbol{c} \in \mathbb{R}^{\lvert E \rvert}$	edge capacities
$\mathrm{val}(\boldsymbol{f})$	value (routed units of flow) of an s-t flow $\boldsymbol{f}$
$G_{\boldsymbol{f}}$	residual graph with respect to the flow $\boldsymbol{f}$
$\hat{\boldsymbol{f}}$	flow within the residual graph $G_{\boldsymbol{f}}$
U	maximum edge capacity, $\max_e \boldsymbol{c}(e)$
$L_{\boldsymbol{f}}$	level graph of residual graph $G_{\boldsymbol{f}}$

OPTIMIZATION

$L(\boldsymbol{y}, \boldsymbol{x}, \boldsymbol{s})$	Lagrangian of an optimization problem with primal variables $\boldsymbol{y}$ and dual variables $\boldsymbol{x}, \boldsymbol{s}$

Bibliography

Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

Henri Cartan. *Differential calculus*. Hermann, 1983.

Li Chen, Rasmus Kyng, Yang P Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. *arXiv preprint arXiv:2203.00671*, 2022.

Rasmus Kyng and Sushant Sachdeva. Approximate gaussian elimination for laplacians-fast, sparse, and simple. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 573–582. IEEE, 2016.

James B Orlin. Max flows in o (nm) time, or better. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 765–774, 2013.

Ralph Tyrell Rockafellar. Convex analysis. In *Convex analysis*. Princeton university press, 2015.

Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.

Index