Mathematical Toolkit **Autumn 2025**

Lecture 9: October 28, 2025

Lecturer: Madhur Tulsiani

1 Bounding the eigenvalues: Gershgorin Disc Theorem

We will now see a simple but extremely useful bound on the eigenvalues of a matrix, given by the Gershgorin disc theorem. Many useful variants of this bound can also be derived from the observation that for any invertible matrix S, the matrices $S^{-1}MS$ and M have the same eigenvalues (prove it!).

Theorem 1.1 (Gershgorin Disc Theorem) *Let $M \in \mathbb{C}^{n\times n}$. Let $R_i = \sum_{j\neq i} |M_{ij}|$. Define the set*

$$\mathrm{Disc}(M_{ii}, R_i) \;:=\; \{z \mid z \in C, |x - M_{ii}| \leq R_i\}\;.$$

If λ is an eigenvalue of M, then

$$\lambda \;\in\; \bigcup_{i=1}^{n} \mathrm{Disc}(M_{ii}, R_i)\,.$$

Proof: Let $x \in \mathbb{C}^n$ be an eigenvector corresponding to the eigenvalue λ. Let $i_0 = \mathrm{argmax}_{i\in[n]}\{|x_i|\}$. Since x is an eigenvector, we have

$$Mx = \lambda x \quad \Rightarrow \quad \forall i \in [n] \;\; \sum_{j=1}^{n} M_{ij} z_j \;=\; \lambda x_i\,.$$

In particular, we have that for $i = i_0$,

$$\sum_{j=1}^{n} M_{i_0 j} x_j \;=\; \lambda x_{i_0} \;\;\Rightarrow\;\; \sum_{j=1}^{n} M_{i_0 j}\frac{x_j}{x_{i_0}} \;=\; \lambda \;\;\Rightarrow\;\; \sum_{j\neq i_0} M_{i_0 j}\frac{x_j}{x_{i_0}} \;=\; \lambda - M_{i_0 i_0}\,.$$

Thus, we have

$$|\lambda - M_{i_0 i_0}| \;\leq\; \sum_{j\neq i_0} |M_{i_0 j}| \cdot \left|\frac{x_j}{x_{i_0}}\right| \;\leq\; \sum_{j\neq i_0} |M_{i_0 j}| \;=\; R_{i_0}\,.$$

■

1.1 An application to compressed sensing

The Gershgorin disc theorem is quite useful in compressed sensing, to ensure what is known as the "Restricted Isometry Property" for the measurement matrices.

Definition 1.2 *A matrix $A \in \mathbb{R}^{k \times n}$ is said to have the* restricted isometry property *with parameters (t, δ_t) if*

$$(1 - \delta_t) \cdot \|x\|^2 \;\leq\; \|Ax\|^2 \;\leq\; (1 + \delta_t) \cdot \|x\|^2$$

for all $x \in \mathbb{R}^n$ which satisfy $|\{i \mid x_i \neq 0\}| \leq t$.

Thus, we want the transformation A to be approximately norm preserving for all *sparse* vectors x. This can of course be ensured for all x by taking $A = \mathsf{id}$, but we require $k \ll n$ for the applications in compressed sensing. In general, the restricted isometry property is NP-hard to verify and can thus also be hard to reason about for a given matrix. The Gershgorin Disc Theorem lets us derive a much easier condition which is sufficient to ensure the restricted isometry property.

Definition 1.3 *Let $A \in \mathbb{R}_{k \times n}$ be such that $\left\|A^{(i)}\right\| = 1$ for each column $A^{(i)}$ of A. Define the* coherence *of A as*

$$\mu(A) \;=\; \max_{i \neq j} \left| \left\langle A^{(i)}, A^{(j)} \right\rangle \right| .$$

We will prove the following

Proposition 1.4 *Let $A \in \mathbb{R}^{k \times n}$ be such that $\left\|A^{(i)}\right\| = 1$ for each column $A^{(i)}$ of A. Then, for any t, the matrix A has the restricted isometry property with parameters $(t, (t-1)\mu(A))$.*

Note that the bound becomes meaningless if $s \geq 1 + \frac{1}{\mu(A)}$. However, the above proposition shows that it may be sufficient to bound $\mu(A)$ (which is also easier to check in practice).

Proof: Consider any x such that $|\{i \mid x_i \neq 0\}| \leq t$. Let S denote the support of x i.e., $S = \{i \mid x_i \neq 0\}$. Let A_S denote the $k \times |S|$ submatrix where we only keep the columns corresponding to indices in S. Let x_S denote x restricted to the non-zero entries. Then

$$\|Ax\|^2 \;=\; \|A_S x_S\|^2 \;=\; \left\langle A_S^T A_S x_S, x_S \right\rangle .$$

Thus, it suffices to bound the eigenvalues of the matrix $A_S^T A_S$. Note that $(A_S)_{ij} = \left\langle A^{(i)}, A^{(j)} \right\rangle$. Thus the diagonal entries are 1 and the off-diagonal entries are bounded by $\mu(A)$ in absolute value. By the Gershgorin Disc Theorem, for any eigenvalue λ of A, we have

$$|\lambda - 1| \;\leq\; (t-1) \cdot \mu(A) .$$

Thus, we have

$$(1-(t-1)\cdot\mu(A))\cdot\|x\|^2 \;\leq\; \|Ax\|^2 \;\leq\; (1+(t-1)\cdot\mu(A))\cdot\|x\|^2\,,$$

as desired. ∎

The theorem is also very useful for bounding how much the eigenvalues of matrix change due to a perturbation. We will see an example of this in the homework.

2 Solving systems of linear equations: Gaussian elimination

Given a system of linear equations $Ax = b$ for $A \in \mathbb{F}^{m\times n}, b \in \mathbb{F}^m$, recall that we can solve the system or determine that there is no solution by converting the matrix $[A \mid b]$ to a row-reduced form using elementary row operations.

Definition 2.1 *A matrix $M \in \mathbb{F}^{m\times n}$ is said to be in* row-reduced *form if*

- *The first non-zero entry in each row (known as the* leading entry*) is 1.*
- *If the leading entry in row i_0 is in column j_0, then $M_{ij} = 0$ for all $i > i_0$ and $j \leq j_0$.*
- *All non-zero rows occur above the zero rows.*

Notice that a matrix in the row-reduced form is always upper triangular. The system has no solution if and only if there is a non-zero row with a leading entry in the last column (corresponding to the entries of b). Also, if the system has a solution, then it can easily be found using back-substitution, starting from the last non-zero row.

Also, recall that an elementary row operations consist of the following (using M_i to denote the i^{th} row of M):

- Swapping the rows M_i and M_j, for some $i, j, \in [m]$.
- $M_i \leftarrow c \cdot M_i$ for some $i \in [m], c \in \mathbb{F} \setminus \{0\}$.
- $M_i \leftarrow M_i + c \cdot M_j$ for some $i, j \in [m], c \in \mathbb{F}$.

A matrix M can always be converted to a row-reduced form using elementary row operations, which gives a general algorithm for solving a system of linear equations over any field. However, the time taken by this algorithm can be as large as $\Omega(n^3)$, which is prohibitive for large matrices. In the next lecture, we will discuss methods which can take advantage of sparsity to significantly speed up the solution of linear systems.

Exercise 2.2 *Prove that performing elementary row operations on a given matrix M changes neither the row rank, nor the column rank of M. Use this to prove that for any matrix M, the row-rank and column-rank are equal.*

3 Solving sparse systems of linear equations

Given $A \in \mathbb{R}^{m \times n}$, if we have a representation of the non-zero entries of A in "list form" i.e., a list of the non-zero entries in each row, then the for any vector v, if the matrix has a total of N non-zero entries, then for any vector v, the product Av can be computed using $O(N)$ arithmetic operations. We will keep this as our base cost and try to compute a solution to $Ax = b$ using as few matrix-vector multiplications as possible.

For the purposes of the discussion below, we will assume that $A \in \mathbb{R}^{n \times n}$ is a symmetric, positive-definite matrix (written as $A \succ 0$). This assumption is not as restrictive as it sounds, and in particular is no more restrictive than assuming that A is invertible. Given a system $A_0 x = b_0$, we can always multiply both sides by A_0^T and obtain $A_0^T A_0 x = A_0^T b_0$, where the matrix $A_0^T A_0$ is now positive-definite (if A_0 is invertible). Note that $A_0^T A_0$ may not be sparse, but we can still compute $A_0^T A_0 v$ in $O(N)$ operations for any vector v using only $O(N)$ operations (we will also need the list of non-zero entries in every column for this). Taking $A = A_0^T A_0$ and $b = A_0^T b_0$ satisfies the required assumptions.

Remark 3.1 *The methods we discuss here will require analyzing distances and inner products, and thus we will work with matrices with real entries (though everything we say will extend easily to complex matrices).*

3.1 Steepest descent

Given a system $Ax = b$ with $A \succ 0$, we apply a method for minimizing the function

$$f(x) \;=\; \frac{1}{2} \cdot \langle Ax, x \rangle - \langle b, x \rangle + c$$

for some arbitrary constant $c \in \mathbb{R}$. This can be motivated by recalling that we originally had the system $A_0 x = b_0$ and $Ax = b$ was obtained by multiplying both sides by A_0^T. If we consider minimizing the least square distance, we get

$$\|A_0 x - b_0\|^2 \;=\; \langle A_0 x, A_0 x \rangle - 2 \langle b_0, A_0 x \rangle + \|b_0\|^2 \;=\; \langle Ax, x \rangle - 2 \langle b, x \rangle + \|b_0\|^2 \,.$$

Of course, scaling by a factor of 2 and changing the constant term does not change the minimizer. If u is the solution to the linear system, we can also re-write the above as

$$\|A_0(x-u)\|^2 \;=\; \langle A(x-u), (x-u) \rangle \;=\; \langle x-u, x-u \rangle_A \,,$$

where $\langle x, y \rangle_A$ denotes the function $\langle Ax, y \rangle$.

Exercise 3.2 *Let $A \in \mathbb{R}^{n \times n}$ be a positive definite matrix. Let the function $\mu : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ be defined as $\mu(x, y) = \langle Ax, y \rangle$. Check that μ is an inner product. We will use $\langle \cdot, \cdot \rangle_A$ to this inner product.*

The inner product $\langle \cdot, \cdot \rangle_A$ and the associated norm are sometimes more convenient for measuring the distance to the solution u since this distance actually measures the least square error in the "output" $A_0 x$ rather than the "input" x. We will need this inner product when working with the conjugate gradient method.

We will use the following algorithm for solving the linear system:

- Start with an arbitrary vector x_0.
- At time t, update

$$x_{t+1} \;=\; x_t - \eta \cdot \nabla f(x_t) \;=\; \eta \cdot (Ax_t - b)\;.$$

The method can also be analyzed by choosing an optimal step size η_t at each time t but we will work with the simpler variant here. Let u be the solution to the system $Ax = b$. We note that

$$x_{t+1} - u \;=\; x_t - u - \eta \cdot A(x_t - u) \;=\; (I - \eta A)(x_t - u)\,.$$

By induction,

$$x_t - u \;=\; (I - \eta A)^t (x_0 - u) \quad \Rightarrow \quad \|x_t - u\|_2 \;\leq\; \|I - \eta A\|_2^t \, \|x_0 - u\|_2\;,$$

where we used the fact that if λ is an eigenvalue of M, then λ^t is an eigenvalue of M^t, which gives that $\left\|(I - \eta A)^t\right\|_2 = \|I - \eta A\|_2^t$. Thus, if $\|I - \eta A\|_2$ is small, we can reach a point close to the solution u in a small number of steps. We now choose η to minimize $\|I - \eta A\|_2$. Let $0 < \lambda_1 \leq \cdots \leq \lambda_n$ denote the eigenvalues of A. Then, the eigenvalues of $I - \eta A$ are $1 - \eta\lambda_1 \geq \cdots \geq 1 - \eta\lambda_n$. Thus, we have

$$\|I - \eta A\|_2 \;=\; \max\left\{\left|1 - \eta\lambda_1\right|, \left|1 - \eta\lambda_n\right|\right\}\;.$$

Check that this is minimized for $\lambda = \frac{2}{\lambda_1 + \lambda_n}$. Plugging this, we get that

$$\|I - \eta A\|_2 \;=\; 1 - \frac{2}{\frac{\lambda_n}{\lambda_1} + 1} \;=\; 1 - \frac{2}{\kappa + 1}\,.$$

Here $\kappa = \lambda_n / \lambda_1$ is known as the condition number of the matrix A. Using this, we get that $\|x_t - u\| \leq \varepsilon \, \|x_0 - u\|$ after $O(\kappa \log(1/\varepsilon))$ iterations. Notice that the cost of each iteration is just $O(1)$ matrix-vector multiplications.

Exercise 3.3 *Obtain a similar bound for the distance $\|x_t - u\|_A$ defined as $\sqrt{\langle (x-u), (x-u) \rangle_A}$.*

In the next lecture, we will discuss the conjugate gradient method, which can obtain a similar guarantee in $O(\sqrt{\kappa} \log(1/\varepsilon))$ iterations.