Linear Algebra Intro

Linear Algebra Intro

- Linear algebra is one of the main building blocks of modern RL and dynamical systems
  - We will cover important properties as we go but we won't have time to go in much depth
- A scalar $x \in \mathbb{R}$ is just a real number
- A p-dimensional vector $\boldsymbol{x} \in \mathbb{R}^p$ is a list of p scalars, i.e.,

$$\boldsymbol{x} = \begin{bmatrix} x_1 \\ x_2 \\ \dots \\ x_p \end{bmatrix}$$

  - where x_i denotes the ith element of $\boldsymbol{x}$

Linear Algebra Intro, cont'd

Rensselaer

- A $p \times n$ matrix $\boldsymbol{A} \in \mathbb{R}^{p \times n}$ consists of n p-dimensional vectors, i.e.,

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \dots & \dots & \dots & \dots \\ a_{p1} & a_{p2} & \dots & a_{pn} \end{bmatrix}$$

  - where a_{ij} denotes the element in row i and column j
  - where a_i denotes the ith column vector of $\boldsymbol{A}$
  - where a_i^r denotes the ith row vector of $\boldsymbol{A}$
- Why do we need matrices?
  - Store data
  - Represent multi-dimensional data (e.g., images)
  - Perform operations in multiple dimensions (e.g., rotation)

Linear Algebra Intro, cont'd

Rensselaer

- Vectors are by default represented as columns
- The transpose of a vector $\boldsymbol{x} \in \mathbb{R}^p$, written $\boldsymbol{x}^T$, is a row vector:

$$\boldsymbol{x}^T = [x_1 \quad x_2 \quad \dots \quad x_p]$$

- Similarly, the transpose of a matrix $\boldsymbol{A}$ is

$$\boldsymbol{A}^T = \begin{bmatrix} a_{11} & a_{21} & \dots & a_{p1} \\ a_{12} & a_{22} & \dots & a_{p2} \\ \dots & \dots & \dots & \dots \\ a_{1n} & a_{2n} & \dots & a_{pn} \end{bmatrix}$$

  – or, equivalently

$$\boldsymbol{A}^T = \begin{bmatrix} \boldsymbol{a}_1^T \\ \boldsymbol{a}_2^T \\ \dots \\ \boldsymbol{a}_n^T \end{bmatrix}$$

Multiplication

Rensselaer

- The inner product of two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^p$ is

$$\boldsymbol{x}^T\boldsymbol{y} = x_1y_1 + x_2y_2 + \cdots + x_py_p$$

  - Note they must have the same dimension
  - The inner product is a scalar

- The product of two matrices $\boldsymbol{A} \in \mathbb{R}^{p\times n}$ and $\boldsymbol{B} \in \mathbb{R}^{n\times m}$ is the inner product of all of $\boldsymbol{A}$'s rows with all of $\boldsymbol{B}$'s columns:

$$\boldsymbol{AB} = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ a_{21} & \dots & a_{2n} \\ a_{p1} & \dots & a_{pn} \end{bmatrix} \begin{bmatrix} b_{11} & \dots & b_{1m} \\ \dots & \dots & \dots \\ b_{n1} & \dots & b_{nm} \end{bmatrix} = \begin{bmatrix} c_{11} & \dots & c_{1m} \\ \dots & \dots & \dots \\ c_{p1} & \dots & c_{pm} \end{bmatrix}$$

- Note that dimensions must match!
- What are the dimensions of the output matrix?

$$p \times m$$

Multiplication Example

Rensselaer

$$\boldsymbol{AB} = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \begin{bmatrix} 5 & 6 \\ 7 & 8 \end{bmatrix}$$

$$= \begin{bmatrix} 5 + 14 & 6 + 16 \\ 15 + 28 & 18 + 32 \end{bmatrix}$$

$$= \begin{bmatrix} 19 & 22 \\ 43 & 50 \end{bmatrix}$$

- Note that in general $\boldsymbol{AB} \neq \boldsymbol{BA}$

Multiplication, cont'd

- Note that for any matrix $\boldsymbol{A}$ and vector $\boldsymbol{x}$, the following is true

$$\boldsymbol{A}\boldsymbol{x} = x_1\boldsymbol{a}_1 + \cdots + x_n\boldsymbol{a}_n$$

- Let $\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}$

$$\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x} = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ a_{21} & \dots & a_{2n} \\ a_{p1} & \dots & a_{pn} \end{bmatrix} \begin{bmatrix} x_1 \\ \dots \\ x_n \end{bmatrix}$$

  – What is b_1?

$$\begin{bmatrix} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n \\ 0 \\ \dots \\ 0 \end{bmatrix}$$

Multiplication, cont'd

- Note that for any matrix $\boldsymbol{A}$ and vector $\boldsymbol{x}$, the following is true

$$\boldsymbol{A}\boldsymbol{x} = x_1\boldsymbol{a}_1 + \cdots + x_n\boldsymbol{a}_n$$

- Let $\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}$

$$\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x} = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ a_{21} & \dots & a_{2n} \\ a_{p1} & \dots & a_{pn} \end{bmatrix} \begin{bmatrix} x_1 \\ \dots \\ x_n \end{bmatrix}$$

  - What about the rest?

$$\begin{bmatrix} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n \\ \dots \\ a_{p1}x_1 + a_{p2}x_2 + \cdots + a_{pn}x_n \end{bmatrix} = x_1\boldsymbol{a}_1 + \cdots + x_n\boldsymbol{a}_n$$

Transpose Property

- Note that

$$(\boldsymbol{AB})^T = \boldsymbol{B}^T\boldsymbol{A}^T$$

  - Why?

$$\boldsymbol{AB} = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ a_{21} & \dots & a_{2n} \\ a_{p1} & \dots & a_{pn} \end{bmatrix} \begin{bmatrix} b_{11} & \dots & b_{1m} \\ \dots & \dots & \dots \\ b_{n1} & \dots & b_{nm} \end{bmatrix}$$

- First column of $(\boldsymbol{AB})^T$ is $[\boldsymbol{a}_1^r\boldsymbol{b}_1 \;\; \boldsymbol{a}_1^r\boldsymbol{b}_2 \cdots \boldsymbol{a}_1^r\boldsymbol{b}_n]^T$
  - In general, column i is $\left[\boldsymbol{a}_i^r\boldsymbol{b}_1 \;\; \boldsymbol{a}_i^r\boldsymbol{b}_2 \cdots \boldsymbol{a}_i^r\boldsymbol{b}_n\right]^T$
- First column of $\boldsymbol{B}^T\boldsymbol{A}^T$ $\left[\boldsymbol{b}_1^T\boldsymbol{a}_1^r \; \boldsymbol{b}_2^T\boldsymbol{a}_1^r \cdots \boldsymbol{b}_n^T\boldsymbol{a}_1^r\right]^T$
  - In general, column i is $\left[\boldsymbol{b}_1^T\boldsymbol{a}_i^r \; \boldsymbol{b}_2^T\boldsymbol{a}_i^r \cdots \boldsymbol{b}_n^T\boldsymbol{a}_i^r\right]^T$
- Matrices are the same (note that for vectors $\boldsymbol{x}^T\boldsymbol{y} = \boldsymbol{y}^T\boldsymbol{x}$)

The identity matrix

- There is a special square matrix $\boldsymbol{I} \in \mathbb{R}^{n \times n}$ with 1's on the diagonal and 0's everywhere else:

$$\boldsymbol{I} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & 1 \end{bmatrix}$$

- We call $\boldsymbol{I}$ the identity matrix
- Among other things, multiplication by $\boldsymbol{I}$ does not modify a matrix
  – i.e., for any $\boldsymbol{A} \in \mathbb{R}^{n \times n}$:

$$\boldsymbol{AI} = \boldsymbol{IA} = \boldsymbol{A}$$

Symmetric Matrices

Rensselaer

- A square matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ is symmetric if its values are symmetric about the diagonal, i.e.,

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & a_{21} & \dots & a_{n1} \\ a_{21} & a_{22} & \dots & a_{n2} \\ \dots & \dots & \dots & \dots \\ a_{n1} & a_{n2} & \dots & a_{nn} \end{bmatrix}$$

- Note that if $\boldsymbol{A}$ is symmetric, then $\boldsymbol{A} = \boldsymbol{A}^T$

Notation

- I will use capital letters for random variables, e.g., X
- Vectors are in bold lowercase, e.g., $\boldsymbol{x}$
  - But random vectors will be uppercase bold, i.e., $\boldsymbol{X}$
  - When clear from context, capital bold letters will also indicate matrices, e.g., $\boldsymbol{W}$
- I will use lowercase letters for sampled data points, e.g., $\boldsymbol{x}$
- Subscripts typically indicate the example index in a dataset, e.g., $\boldsymbol{x}_i$ is the ith example in the dataset
- When clear from context, a subscript will also denote the specific element in a vector
  - E.g., x_i is the ith element of vector $\boldsymbol{x}$

Linearly Independent Vectors

Rensselaer

- A sequence of vectors $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k$ is *linearly dependent* if there exist coefficients $a_1, \dots, a_k$, not all zero, such that
$$a_1\boldsymbol{v}_1 + \cdots + a_k\boldsymbol{v}_k = \boldsymbol{0}$$
- For example, if $a_1 \neq 0$, then
$$\boldsymbol{v}_1 = -\frac{a_2}{a_1}\boldsymbol{v}_2 - \cdots - \frac{a_k}{a_1}\boldsymbol{v}_k$$
  - i.e., $\boldsymbol{v}_1$ can be written as a linear combination of other $\boldsymbol{v}_i$'s
- The $\boldsymbol{v}_i$'s are *linearly independent* if there exist no such a_i's
- Linear independence is a central concept in linear algebra
- For example, if each $\boldsymbol{v}_i \in \mathbb{R}^k$, then the $\boldsymbol{v}_i$'s form a basis for $\mathbb{R}^k$
  - Every other vector in $\mathbb{R}^k$ can be written as a linear combination of $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k$

Linear independence, examples

- Are the following vectors linear independent:

$$\boldsymbol{x} = [1,2], \boldsymbol{y} = [2,4]$$

  - No, because $\boldsymbol{y} = 2\boldsymbol{x}$

$$\boldsymbol{x} = [1,0], \boldsymbol{y} = [0,1]$$

  - Yes, there is no way to express $\boldsymbol{y}$ as a multiple of $\boldsymbol{x}$

$$\boldsymbol{x} = [1,2,3], \boldsymbol{y} = [4,5,6], \boldsymbol{z} = [5,7,9]$$

  - No, because $\boldsymbol{z} = \boldsymbol{x} + \boldsymbol{y}$

Matrix Rank

- Suppose $\boldsymbol{A} \in \mathbb{R}^{m \times n}$
  - i.e., $\boldsymbol{A}$ consists of n m-dimensional columns
- The *rank* of $\boldsymbol{A}$ is the maximal number of linearly independent columns in $\boldsymbol{A}$
- A matrix $\boldsymbol{A}$ is said to be full rank if its rank is equal to the number of columns
- Is $\boldsymbol{A} = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$ full rank?
  - Yes, its columns are independent
- Is $\boldsymbol{A} = \begin{bmatrix} 1 & 4 & 5 \\ 2 & 5 & 7 \\ 3 & 6 & 9 \end{bmatrix}$ full rank?
  - No. $\boldsymbol{A}$ has a rank of 2

Matrix Inverse

- Let $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ be a square matrix
- If $\boldsymbol{A}$ is full rank, then there exists a unique matrix $\boldsymbol{B} \in \mathbb{R}^{n \times n}$ such that

$$\boldsymbol{B}\boldsymbol{A} = \boldsymbol{I}$$

  - where $\boldsymbol{I}$ is the identity matrix
  - We say $\boldsymbol{B}$ is the inverse of $\boldsymbol{A}$, written $\boldsymbol{A}^{-1}$
- If $\boldsymbol{A}$ is not full rank, the inverse does not exist

Matrix Pseudo-Inverse

- Suppose $\boldsymbol{A}$ is not square, i.e., $\boldsymbol{A} \in \mathbb{R}^{m \times n}$
  - Assume first $m > n$, i.e., $\boldsymbol{A}$ is a tall matrix
- If $\boldsymbol{A}$ is full rank, then $\boldsymbol{A}^T\boldsymbol{A}$ is full rank (and square)
  - However, $\boldsymbol{A}\boldsymbol{A}^T$ is not!
- Consider the matrix $\left(\boldsymbol{A}^T\boldsymbol{A}\right)^{-1}\boldsymbol{A}^T\boldsymbol{A}$
  - What is it equal to?
  - We say $\left(\boldsymbol{A}^T\boldsymbol{A}\right)^{-1}\boldsymbol{A}^T$ is the pseudo-inverse of $\boldsymbol{A}$
  - Called "pseudo-inverse" because it is not unique
- What about the case $m < n$?
  - The pseudo-inverse is $\boldsymbol{A}^T\left(\boldsymbol{A}\boldsymbol{A}^T\right)^{-1}$, on the right:

$$\boldsymbol{A}\boldsymbol{A}^T\left(\boldsymbol{A}\boldsymbol{A}^T\right)^{-1} = \boldsymbol{I}$$

Eigenvectors and Eigenvalues

Rensselaer

- Suppose we are given a square matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$
- A vector $\boldsymbol{v}$ is said to be an eigenvector of $\boldsymbol{A}$ if

$$\boldsymbol{A}\boldsymbol{v} = \lambda\boldsymbol{v}$$

  – where $\lambda \in \mathbb{R}$ is a corresponding eigenvalue
- If the matrix $\boldsymbol{A}$ is full rank, it has n eigenvectors, $\boldsymbol{v}_i$
  – And n corresponding eigenvalues, λ_i
  – If eigenvalues are not repeated, the eigenvectors form a basis in $\mathbb{R}^n$
  – i.e., any $\boldsymbol{x} \in \mathbb{R}^n$ can be written as a linear combination

$$\boldsymbol{x} = c_1\boldsymbol{v}_1 + \cdots + c_n\boldsymbol{v}_n$$

- There may be repeated eigenvalues
- $\boldsymbol{A}$ is full rank iff $\lambda_i \neq 0$ for all i

Eigenvectors and Eigenvalues, cont'd

- Suppose a square matrix $\boldsymbol{A}$ has eigenvalues $\lambda_1, \dots, \lambda_n$
- What are the eigenvalues of $\boldsymbol{A}^2$?

$$\lambda_1^2, \dots, \lambda_n^2$$

- Take any eigenvalue λ_i and corresponding eigenvector $\boldsymbol{v}_i$

$$\begin{aligned} \boldsymbol{A}\boldsymbol{A}\boldsymbol{v}_i &= \boldsymbol{A}\lambda_i\boldsymbol{v}_i \\ &= \lambda_i^2\boldsymbol{v}_i \end{aligned}$$

- In general, the eigenvalues of $\boldsymbol{A}^k$ are

$$\lambda_1^k, \dots, \lambda_n^k$$

  - The eigenvectors are the same as those of $\boldsymbol{A}$

Eigenvectors and Eigenvalues Examples

Rensselaer

- Consider the identity matrix

$$\boldsymbol{I} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

- What are the eigenvectors of $\boldsymbol{I}$?
  - Trick question. Every vector is an eigenvector of $\boldsymbol{I}$
  - E.g., unit vectors $\boldsymbol{e}_1 = [1\ 0\ 0]^T, \boldsymbol{e}_2 = [0\ 1\ 0]^T, \boldsymbol{e}_3 = [0\ 0\ 1]^T$
- How about the eigenvalues?

$$\lambda_1 = \lambda_2 = \lambda_3 = 1$$

- Given a vector $\boldsymbol{v} = [v_1\ v_2\ v_3]^T$, how do we express $\boldsymbol{v}$ as a linear combination of the unit vectors?

$$\boldsymbol{v} = v_1\boldsymbol{e}_1 + v_2\boldsymbol{e}_2 + v_3\boldsymbol{e}_3$$

Linear Systems

- Consider a general discrete-time linear system

$$\boldsymbol{x}_k = \boldsymbol{A}\boldsymbol{x}_{k-1}$$

- i.e.,

$$\boldsymbol{x}_k = \boldsymbol{A}^k\boldsymbol{x}_0$$

  – for some initial $\boldsymbol{x}_0$

- Suppose $\boldsymbol{A}$ has non-repeated eigenvalues
  – Recall that the eigenvectors of $\boldsymbol{A}$ form a basis in $\mathbb{R}^n$, so

$$\boldsymbol{x}_0 = c_1\boldsymbol{v}_1 + \cdots + c_n\boldsymbol{v}_n$$

- Then

$$\boldsymbol{A}^k\boldsymbol{x}_0 = c_1\lambda_1^k\boldsymbol{v}_1 + \cdots + c_n\lambda_n^k\boldsymbol{v}_n$$

- Under what conditions does $\boldsymbol{x}_k$ converge to $\boldsymbol{0}$?
  – need $|\lambda_i| < 1$, for all i