

6.1 Introduction: Matrix Inversion

In the last note, we considered a system of pumps and reservoirs where the water in each reservoir is represented as a vector and the pumps, represented as a matrix, act on the reservoirs to move water into a new state. If we know the current state of the reservoirs, $\vec{v}[t]$, and we know the state transition matrix describing the pumps, A , we can find the water at the next time step through matrix-vector multiplication:

$$\vec{v}[t + 1] = A\vec{v}[t]$$

However, suppose we'd like to find the water in the reservoirs at a *previous* timestep, $\vec{v}[t - 1]$. Is there a state transition matrix B , that can take us backwards in time?

$$\vec{v}[t - 1] = B\vec{v}[t]$$

It turns out that the matrix that “undoes” the effects of A is its inverse! In this note, we'll define matrix inverses, introduce some properties, and investigate when matrix inverses exist (and when they don't exist).

6.1.1 Definition and properties of matrix inverses

Definition 6.1 (Inverse): A square matrix A is said to be invertible if there exists an matrix B such that

$$AB = BA = I. \tag{1}$$

where I is the identity matrix. In this case, we call the matrix B the inverse of the matrix A , which we denote as A^{-1} .

Example 6.1 (Matrix inverse): Consider the 2×2 matrix $A = \begin{bmatrix} 1 & 1 \\ 2 & 1 \end{bmatrix}$. Then $A^{-1} = \begin{bmatrix} -1 & 1 \\ 2 & -1 \end{bmatrix}$. We can verify that the following holds

$$AA^{-1} = \begin{bmatrix} 1 & 1 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} -1 & 1 \\ 2 & -1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{2}$$

$$A^{-1}A = \begin{bmatrix} -1 & 1 \\ 2 & -1 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 2 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}. \tag{3}$$

Let's show an important property of matrix inverses: **If A is an invertible matrix, then its inverse must be unique.**

Proof. Suppose B_1 and B_2 are both inverses of the matrix A . Then we have

$$AB_1 = B_1A = I \quad (4)$$

$$AB_2 = B_2A = I \quad (5)$$

Now take the equation

$$AB_1 = I. \quad (6)$$

Multiplying both sides of the equation by B_2 from the left, we have

$$B_2(AB_1) = B_2I = B_2. \quad (7)$$

Notice that by associativity of matrix multiplication, the left hand side of the equation above becomes

$$B_2(AB_1) = (B_2A)B_1 = IB_1 = B_1. \quad (8)$$

Hence we have

$$B_1 = B_2. \quad (9)$$

We see that B_1 and B_2 must be equal, so the inverse of any invertible matrix is unique. \square

In discussion, you will see a few more useful properties of matrix inverses.

Some of the next natural questions to ask are:

- How do we know whether or not a matrix is invertible?
- If a matrix is invertible, how do we go about finding its inverse?

It turns out Gaussian elimination could help us answer these questions!

6.1.2 Finding Inverses With Gaussian Elimination

A square matrix M and its inverse M^{-1} will always satisfy the following conditions $MM^{-1} = I$ and $M^{-1}M = I$, where I is the identity matrix.

$$\text{Let } M = \begin{bmatrix} 1 & 1 \\ 2 & 1 \end{bmatrix} \text{ and } M^{-1} = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix}$$

We want to find the values of b_{ij} such that the equation $MM^{-1} = I$ would be satisfied.

$$\begin{bmatrix} 1 & 1 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

Just as we did for solving equations of the form $A\vec{x} = \vec{b}$, we can write the above as an **augmented matrix**, which joins the left and right numerical matrices together and hides the variable matrix, as shown below.

$$\left[\begin{array}{cc|cc} 1 & 1 & 1 & 0 \\ 2 & 1 & 0 & 1 \end{array} \right]$$

Now, to find the inverse matrix M^{-1} using Gaussian elimination, we have to transform the left numerical matrix (left half of the augmented matrix) to the identity matrix, then the right numerical matrix (right half of the augmented matrix) becomes our solution. In equation form $MM^{-1} = I$, we are transforming M and I simultaneously using row operations so that the equation becomes $IM^{-1} = A$, where A is the resulting numerical matrix from the Gaussian elimination. Since M^{-1} is multiplied by the identity matrix I , the resulting numerical matrix A must equal to M^{-1} , and we have the values for the elements in our inverse matrix. We will now do the actual computation below:

$$\begin{aligned} \left[\begin{array}{cc|cc} 1 & 1 & 1 & 0 \\ 2 & 1 & 0 & 1 \end{array} \right] &\Rightarrow R_2 - 2R_1 \Rightarrow \left[\begin{array}{cc|cc} 1 & 1 & 1 & 0 \\ 0 & -1 & -2 & 1 \end{array} \right] \Rightarrow -1(R_2) \Rightarrow \left[\begin{array}{cc|cc} 1 & 1 & 1 & 0 \\ 0 & 1 & 2 & -1 \end{array} \right] \\ &\Rightarrow R_1 - R_2 \Rightarrow \left[\begin{array}{cc|cc} 1 & 0 & -1 & 1 \\ 0 & 1 & 2 & -1 \end{array} \right] \end{aligned}$$

M^{-1} is the right half of the augmented matrix. Therefore $M^{-1} = \begin{bmatrix} -1 & 1 \\ 2 & -1 \end{bmatrix}$. More generally, for any $n \times n$ matrix M , we can perform Gaussian elimination on the augmented matrix

$$\left[\begin{array}{c|c} M & I_n \end{array} \right]$$

If at termination of Gaussian elimination, we end up with an identity matrix on the left, then the matrix on the right is the inverse of the matrix M .

$$\left[\begin{array}{c|c} I_n & M^{-1} \end{array} \right]$$

If we don't end up with an identity matrix on the left after running Gaussian elimination, we know that the matrix is not invertible.

Knowing if a matrix is invertible can tell us about the rows/columns of a matrix, and knowing about the rows/columns can tell us if a matrix is invertible - let's look at how.

Additional Resources For more on matrix inverses, read *Strang* pages 83-85 and try Problem Set 2.5; or read *Schum*'s pages 33-34 and try Problems 2.17 to 2.20, 2.54, 2.55, 2.57, and 2.58.

6.2 Connecting invertibility with matrix rows and columns

First let's consider how the rows of the matrix relate to invertibility.

Example 6.2 (Invertibility Intuition – Rows): Suppose we have a black and white image with two pixels. We cannot directly see the shade of each pixel, but we can measure *linear combinations* of the light the two pixels absorb. Define the light absorbed by the two pixels as x_1 and x_2 . We take the following two measurements, y_1, y_2 :

$$\begin{aligned} y_1 &= x_1 + x_2 \\ y_2 &= 2x_1 + 2x_2 \end{aligned}$$

Written in matrix form:

$$\vec{y} = \begin{bmatrix} 1 & 1 \\ 2 & 2 \end{bmatrix} \vec{x} \quad (10)$$

Can we invert our measurement matrix to solve for the shade of each pixel? No, our second measurement y_2 does not provide any new information, since it is linearly dependent with the first measurement. In our matrix representation, each measurement corresponds to a row, so we can guess that we need linearly independent rows to have an invertible matrix.

We won't prove this rigorously, but we can extend this intuition by examining the Gaussian elimination method for finding matrix inverses: If we run Gaussian elimination on a matrix M and do not end up with the identity matrix, this means that the matrix is not invertible. If we don't get the identity matrix, we will have a row of zeros, which indicates that the rows of M are linearly dependent.

Now let's look from the column perspective.

Consider A as an operator on any vector $\vec{x} \in \mathbb{R}^n$. What does it mean for A to have an inverse? It suggests that we can find a matrix that "undoes" the effect of matrix A operating on any vector $\vec{x} \in \mathbb{R}^n$. What property should A have in order for this to be possible? A should map any two distinct vectors to distinct vectors in \mathbb{R}^n , i.e., $A\vec{x}_1 \neq A\vec{x}_2$ for vectors \vec{x}_1, \vec{x}_2 such that $\vec{x}_1 \neq \vec{x}_2$.

Consider this example:

Example 6.3 (Invertibility Intuition – Columns):

Is the matrix $A = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$ invertible? Intuitively, it is not because A can map two distinct vectors into the same vector.

$$\begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 3 \end{bmatrix} = 1 \times \begin{bmatrix} 1 \\ 0 \end{bmatrix} + 3 \times \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 4 \\ 0 \end{bmatrix} \quad (11)$$

$$\begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 2 \\ 2 \end{bmatrix} = 2 \times \begin{bmatrix} 1 \\ 0 \end{bmatrix} + 2 \times \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 4 \\ 0 \end{bmatrix}. \quad (12)$$

We cannot recover the vector uniquely after it is operated by A . This is connected with the fact that the columns are linearly dependent – different weighted combinations of columns could generate the same vector.

Theorem 6.1: A matrix A is invertible if and only if its columns are linearly independent.

Let's prove this theorem. The statement "if and only if" means that we need to prove two things:

1. If A is invertible, then its columns are linearly independent.
2. If A 's columns are linearly independent, then it is invertible.

For the first statement, we're trying to prove the following:

$$A^{-1} \text{ exists} \implies \text{the columns of } A \text{ are linearly independent}$$

We know that if the columns of \mathbf{A} are linearly independent, then $\mathbf{A}\vec{x} = \vec{0}$ only when $\vec{x} = \vec{0}$, so we can rephrase what we're trying to prove as

$$\mathbf{A}^{-1} \text{ exists} \implies (\mathbf{A}\vec{x} = \vec{0} \text{ only when } \vec{x} = \vec{0})$$

To prove this, assume that \mathbf{A} is invertible. Let \vec{v} be some vector such that $\mathbf{A}\vec{v} = \vec{0}$:

$$\begin{aligned} \mathbf{A}\vec{v} = \vec{0} &\longleftarrow \text{left-multiply by } \mathbf{A}^{-1} \\ \mathbf{A}^{-1}\mathbf{A}\vec{v} = \mathbf{I}\vec{v} = \vec{0} \\ \vec{v} &= \vec{0} \end{aligned}$$

Hooray! We've successfully proven the first statement.

The second statement is a little trickier to prove:

$$(\text{columns of } \mathbf{A} \text{ are linearly independent}) \implies \mathbf{A}^{-1} \text{ exists}$$

or, rewritten

$$(\mathbf{A}\vec{x} = \vec{b} \text{ has a unique solution } \vec{x}) \implies \mathbf{A}^{-1} \text{ exists}$$

Now suppose that $\{\vec{e}_1, \vec{e}_2, \dots, \vec{e}_n\}$ are the columns of the identity matrix. We want to see if there's a matrix \mathbf{M} which is \mathbf{A} 's inverse, i.e. it obeys the following properties:

$$\begin{aligned} \mathbf{A}\mathbf{M} = \mathbf{M}\mathbf{A} &= \mathbf{I} \\ &= \begin{bmatrix} | & | & \cdots & | \\ \vec{e}_1 & \vec{e}_2 & \cdots & \vec{e}_n \\ | & | & \cdots & | \end{bmatrix} \end{aligned}$$

Using the definition of matrix-matrix multiplication as a series of stacked matrix-vector multiplications:

$$\mathbf{A} \begin{bmatrix} | & | & \cdots & | \\ \vec{m}_1 & \vec{m}_2 & \cdots & \vec{m}_n \\ | & | & \cdots & | \end{bmatrix} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{A}\vec{m}_1 & \mathbf{A}\vec{m}_2 & \cdots & \mathbf{A}\vec{m}_n \\ | & | & \cdots & | \end{bmatrix} = \begin{bmatrix} | & | & \cdots & | \\ \vec{e}_1 & \vec{e}_2 & \cdots & \vec{e}_n \\ | & | & \cdots & | \end{bmatrix}$$

What we need to prove now is that there are vectors $\{\vec{m}_1, \vec{m}_2, \dots, \vec{m}_n\}$ such that

$$\mathbf{A}\vec{m}_i = \vec{e}_i \forall i \in \{1, 2, \dots, n\}$$

but we already know this is true because this is a square matrix and because of the definition of linear independence above! We also already know from lecture that the left and right inverses are identical. Thus, we've proven the second statement because the m vectors exist, i.e. \mathbf{M} exists.

6.3 Practice Problems

These practice problems are also available in an interactive form on the course website.

1. Find the inverse of $\begin{bmatrix} 1 & 0 & 2 \\ -1 & 1 & -1 \\ 0 & -1 & 0 \end{bmatrix}$.
2. Find the inverse of $\begin{bmatrix} 1 & 0 & 2 \\ -1 & 1 & -1 \\ 0 & -1 & -1 \end{bmatrix}$.
3. Suppose $\mathbf{A} = \mathbf{BC}$, where \mathbf{B} is a 4×2 matrix and \mathbf{C} is a 2×4 matrix. Is \mathbf{A} invertible?
 - (a) Yes, \mathbf{A} is invertible.
 - (b) No, \mathbf{A} is not invertible.
 - (c) Depends on \mathbf{C} only.
 - (d) Depends on \mathbf{B} and \mathbf{C} .
4. Let the matrix \mathbf{A} be the state transition matrix for some system. Given some state after n steps $\vec{x}[n]$, can we always find $\vec{x}[n+1]$?
 - (a) Yes, we simply apply the matrix \mathbf{A} on $\vec{x}[n]$.
 - (b) No, we need to know the initial state $\vec{x}[0]$.
 - (c) No, we don't have enough information about the system.
5. Let the matrix \mathbf{A} be the state transition matrix for some system. Given some state after n steps $\vec{x}[n]$, can we always find $\vec{x}[n-1]$?
 - (a) Yes, we can use Gaussian elimination to find the initial state.
 - (b) Yes, we simply apply the matrix \mathbf{A}^{-1} on $\vec{x}[n]$.
 - (c) No, we don't know whether the matrix \mathbf{A} is invertible.
6. Suppose that the state transition matrix for a system is given by $\begin{bmatrix} 1 & 0.5 \\ 0 & 0 \end{bmatrix}$. Given some state after n steps $\vec{x}[n]$, can we find $\vec{x}[n-1]$?
7. True or False: The inverse of a diagonal matrix, where all of the diagonal entries are non-zero, is another diagonal matrix.
8. True or False: If $\mathbf{A}^n = \mathbf{0}$, where $\mathbf{0}$ is the zero matrix, for some $n \in \mathbb{R}$, then \mathbf{A} is not invertible.